JULY | 2011

INFORMATION SECURITY

Preventing threats, protecting your data

800.808.4239 | CDWG.com/securityguide

CDW·G REFERENCE GUIDE

A guide to the latest technology for people who get IT

CDW·G PEOPLE WHO GET IT™

INFORMATION SECURITY REFERENCE GUIDE

IN THIS ISSUE:

WHAT IS A CDW·G REFERENCE GUIDE?

At CDW·G, we're committed to getting you everything you need to make the right purchasing decisions — from products and services to information about the latest technology.

Our Reference Guides are designed to provide you with an in-depth look at topics that relate directly to the IT challenges you face. Consider them an extension of your account manager's knowledge and expertise. We hope you find this guide to be a useful resource.

SECURITY: ENABLING SUCCESS

Well-organized security addressing today's threats supports an organization's overall productivity.

One of the few constants in the IT world is that it is always changing. Nowhere is this more true than in security. And three of the larger developments in IT, social networking, remote access and cloud computing, are having a direct impact on the security equation for organizations.

Many users have grown up with the Internet, social networking and an always-on IT environment. As their computing demands evolve, so does the threat landscape. Providing safe access to the services and data an organization's users need requires a thoughtful rethinking of security.

Social networking sites focus primarily on building interpersonal community. While we tend to think of these sites in terms of our personal lives, they are steadily becoming a part of the educational and government IT landscape. Organizations of all types are turning to such tools to better connect with the public.

Remote access allows network access to offsite staff, vendor partners and other guests. It also invites uncontrolled and poorly understood systems into the enterprise network. While IT managers can acknowledge the benefits of these technologies, they do have security trade-offs, and appropriate controls need to be implemented to maintain the integrity of data and systems.

Cloud computing, essentially providing scalable end-user access to applications, infrastructure and platforms via a front-end interface, such as a web browser, is changing the way that organizations arrange their computing operations. This dynamic approach to computing comes with unique security concerns, but also presents some unique opportunities for improving an organization's overall security.

Security = Productivity

Fortunately, when done right, security enables productivity. Social networking sites and remote access shouldn't engender fear; they simply require security based on sound policy, along with organizational buy-in. Savvy IT and security leaders can give their staffs the tools they need to be productive and still maintain security.

In organizations of every size, IT now forms the basis for almost every activity: public-facing operations, distribution of services and information, the organizational back office, financial and personnel systems. When the IT foundation is secured properly, the organization can move ahead smoothly and pursue its mission and goals.

Staff need to have confidence that their normal daily tasks are not going to put their organization at risk. When security is built into an organization's IT infrastructure, work — not worry — is the result.

While the idea of security as an operations enabler seems simple, achieving a properly secured organization can be challenging because it requires the entire operation, not just the security team, to step back and think globally about information security.

Threats in 2011

Spam, viruses, malware and phishing

SECURITY: KNOW THYSELF

The secret to proper security is very simple: Understand the organization's operations. Without that understanding, information security can't be effective.

When security investments and security strategies are closely tied to the organization's processes, everything else follows. Investment is proportionate, risk mitigation is appropriate, and costs and benefits are aligned.

Going from idea to implementation requires three key strategies: full integration, appropriate investment, and continuous feedback.

are never going away — in fact, they're only getting worse. While traditional security techniques to keep these threats at bay are evolving, there's not a lot of new ground being broken in these areas. IT managers with well-thought-out endpoint security programs and perimeter defenses shouldn't spend time re-evaluating those defenses against old-school attacks.

The onset of focused attacks, rather than these more random threats, is changing the game of security. Organizations are under attack all the time. But some organizations are under direct, targeted attack, and this type of threat requires a different way of thinking.

Targeted Intrusions

Targeted attacks can come from both outside and inside an organization. As the disclosure of classified federal government information to WikiLeaks reminds us, even trusted staff continue to present a threat when it comes to data leaks, whether it's accidental or deliberate.

When the threat comes from outside the organization, the security industry has coined a new term: advanced persistent threat, or APT. APTs are not broad-scale attacks aimed at anyone and everyone. They're specific, directed threats aimed at a single target, possibly incorporating zero-day vulnerability exploits and stacking multiple attacks into a single package.

For example, a hacker using a two-year-old attack to break into a web server to store a personal collection of MP3s is a threat. A hacker writing an exploit specifically to break into an organization's web server to steal its CAD files is an APT.

A New Breed of Hacker

The second big change in the threat landscape comes from a newly motivated attacker. In the 1990s, hackers were the computing equivalent of a graffiti tagger, largely motivated by fame and doing little damage.

In the last ten years, financial incentive became a larger motivation for attacks. Hackers aimed to steal data, such as Social Security numbers, that they could sell; or they attempted to take over an army of computers to build a botnet that could be rented out to the highest bidder.

Today, with political upheaval around the world, threats are also coming from attackers who want to expose confidential data that they feel will be incriminating toward the organization, or who want to punish organizations they feel are on the wrong side of their particular cause. Since these attackers have an entirely different motivation, warding them off requires different defenses.

IT managers should consider the targeted intrusion by a motivated attacker as their primary new threat. Existing security products and policies should remain in place, but IT departments must evaluate whether current security plans adequately protect the organization from such intrusions.

As the sidebar *Strategies for Security Success* shows, IT managers can employ three key strategies to fully integrate security, tie it to the organization's operations and continuously evaluate the organizational security posture. ■

STRATEGIES FOR SECURITY SUCCESS

Strategy	Implementation
Full integration	Security must be baked into the network, into applications, and into systems and processes. Layering security on top is rarely effective or efficient. Building security in from the beginning makes it part of every team's responsibility.
Appropriate investment	To achieve the appropriate level of investment, remember that security is about reducing risk. Security specialists with a global view must work with other teams to help map out risks to the organization. This exposes areas where security is needed, thereby simplifying planning and justifying security investments.
Continuous feedback	Deploying security technologies must be part of a continuous cycle. Protection is important, but accountability and visibility are just as crucial. By measuring value and pruning what doesn't work, IT managers can avoid ever-escalating maintenance costs and complexity.

ASSESSING RISK

Figuring out the costs associated with risk yields strong security.

Good security starts with a simple idea: reduction of risk. At its core, information security is about reducing risk. An economist would go even further and say that security is about reducing costs. When bad things happen, there are costs. The costs can be direct, such as compensation to people for losses, or indirect, such as damage to reputation. Reduce the risk, reduce the cost.

On the other hand, security has a cost, too — staff and budget. If security delays a project, there's an opportunity cost. And if security gets in the way of staff getting their work done, that's yet another cost.

So how do IT managers decide the best way to balance these costs — the cost of bad things happening against the cost of reducing risks through security? The concept is easy: Use subtraction.

Good security results when the cost of the security is lower than the cost of the risk. When the cost of security exceeds the cost of risk, the result is poor security.

Calculating Risk and Security Costs

Balancing the costs of risk and security can be tricky. As a result, many IT managers have given up trying to quantify security and risk, instead relying on best practices, peer benchmarking and regulatory requirements. That's dangerous.

Security investments must be connected to the organization's operations. And even if the process is difficult, the connection must be quantifiable and defensible. Without a clear connection, organizations risk overinvesting in some aspects of security and underinvesting in others.

Using a framework to compare costs and risk helps balance the two. Not every asset will fit into such a framework, and not every organization will use the same one. But having a simple framework to compare security investments with costs to the organization can help leadership identify which investments make sense and which don't. It can also identify areas where important assets have

CALCULATING RETURN ON SECURITY INVESTMENT

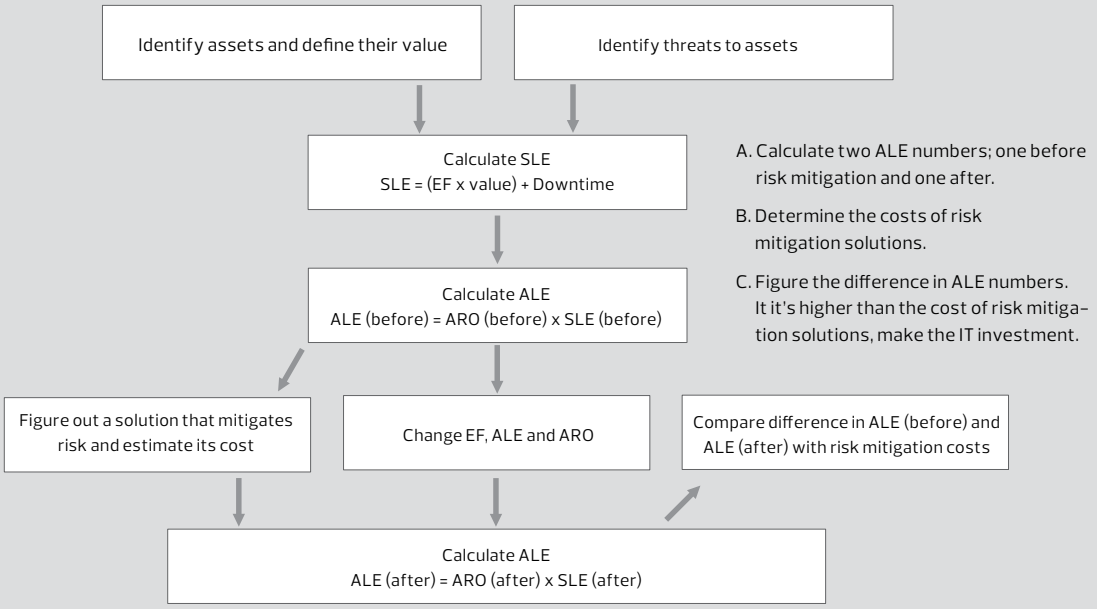

SLE = Single Loss Expectancy ALE = Annual Loss Expectancy EF = Exposure Factor (0–100%) ARO = Annual Rate of Occurrence (0–100%)
(Before) = Before Risk Mitigation (After) = After Risk Mitigation

significant risks that are not being mitigated.

Compliance and Security

Compliance and security go hand in hand at many organizations, but IT managers should understand that they don't have the same goals. Compliance is a process of conforming to a particular standard or regulatory regime.

Being in compliance doesn't necessarily mean that risk is reduced — other than the risk of failing a compliance audit. This regulatory risk presents yet another cost that must be factored into any security (or compliance) decision-making.

Compliance rules come from different sources. In some cases, legislation forms the basis for compliance initiatives. For example, the Health Insurance Portability and Accountability Act (HIPAA) applies to any organization remotely related to providing healthcare, while the Gramm–Leach–Bliley Act (GLBA) applies to banks, and Sarbanes–Oxley (SOX) dictates rules for publicly held entities.

States also have an array of privacy laws related to security. Organizations serving end users across many or all 50 states have discovered that they need to comply with the most restrictive laws among all of them. Often, California is the role model for privacy and security, so it is worth paying attention to the legislative news from the West Coast.

Other compliance requirements come from industry regulations. The Payment Card Industry Data Security Standard (PCI DSS) is the most familiar example of an industry standard that many organizations that handle financial transactions must comply with.

Compliance efforts can play a valuable role in security planning. When IT managers link compliance to correct intentions, regulatory risk and other risks are all reduced. In other words, compliance gives better security when the organization follows the spirit and intention of the compliance program, rather than blindly following the letter of the law.

An organization concerned about security asks itself: What do we have to do in order to be safe? An organization focused on compliance asks a different question: What do we have to do in order to meet this set of requirements? Or, more cynically: How bad can our security be and still comply?

The key to making compliance efforts an overall benefit instead of an exercise in frustration is to combine the two: How can we meet this set of requirements, while increasing our overall security?

FRAMING THE ANSWER

If the framework in the diagram *Calculating Return on Security Investment* on page 6 doesn't work for an organization, there's no shortage of formulas available for calculating security risk.

- **Factor Analysis of Information Risk:** FAIR is a proprietary framework that focuses on statistical methods of analyzing risk. It is available free of charge for noncommercial use for those who want to analyze their organization. It is available at **fairwiki.riskmanagementinsight.com**
- **Operationally Critical Threat, Asset and Vulnerability Evaluation:** OCTAVE was created by Carnegie Mellon's Software Engineering Institute. This suite of tools and procedures for security strategy planning comes in three different formats, including one for smaller organizations. It is available at **cert.org/octave**
- **Risk Management Framework:** RMF, developed by the U.S. National Institute of Standards and Technology, will be especially of interest to federal agencies because it links directly to Federal Information Processing Standards and can help ensure FIPS compliance in this area. It is available at **csrc.nist.gov**
- **Threat Agent Risk Assessment:** TARA documents the process that Intel uses in its own risk assessment, focusing on distilling the complicated set of potential threats into the most likely and important ones. A Department of Homeland Security version of Intel's work is available as the IT Sector Baseline Risk Assessment (ITSRA). It is also available at **intel.com**

Security as a Process

As emphasized in Chapter 1, security really needs to be fully integrated into IT and built into projects from the very beginning. The investment should be appropriate to operational needs and be part of a cycle of continuous evaluation and feedback. These ideas can be formalized into a security process that begins with policy and then continues for the life of the organization.

The security process has four phases: design, implementation, testing and monitoring. Because every project has its own timetable, a security team may be participating in all four phases at once. Generally, though, most of the process revolves around monitoring security. Projects are executed once, but maintained for a very long time. Unless the monitoring phase shows a need to revisit security, the first three phases may be brief.

1. Design: During system or application design, security has to be a core consideration. A member of the security team should participate during system design meetings to help point out areas where security may be a concern.

Application developers often don't have the background to identify the attack surface of their systems and applications. So security expertise is needed to both understand what is being proposed and to make sure necessary security elements are built into the design.

2. Implementation: As system development begins, the security team must be prepared to make midcourse corrections to compensate for last-minute changes and as fuzzy concepts in the original design are fleshed out. Interfaces between systems and applications, as well as any user-facing components, need to be carefully examined.

3. Testing: Separate teams should help test security to determine whether the security measures work as they're intended to. Automated tools, such as vulnerability assessment packages, are useful at this stage because they allow for repeatable and exhaustive analysis. This is a good opportunity to have some outside parties look at the project to be sure that nothing is missed and that assumptions still hold.

4. Monitoring: The most significant and continuous part of the security process is monitoring. Once everything is in place, security output needs to be integrated into the organization's normal procedures. This makes detecting incidents and responding to them a normal part of operations, rather than something exceptional for this project or process.

Periodically, the monitoring results should be analyzed to measure whether security components are actually mitigating any risks and are within the expected costs. If something is out of line, it's time to jump back to the first phase and realign the security investment with the operational need. ■

CHAPTER 3

KEEPING UP WITH NETWORK SECURITY

Changes to the network require security adjustments.

In some ways, networks have become the circulatory systems of many organizations. Information critical to day-to-day operations passes over the network at some point, making networks a critical asset. If the network stops working, so does the organization. Protecting networks and the data that flow over them is one of the highest priorities for IT managers.

This chapter covers many familiar security tools, such as firewalls, intrusion detection systems (IDSs) and intrusion prevention systems (IPSs). And it covers some newer ones, such as network behavior anomaly detection (NBAD) systems, network access control (NAC) and identity management.

The basic tools for network protection have changed only slightly in the past few years. Firewalls and unified threat management (UTM) devices have become faster and more sophisticated. IDS and IPS solutions have become more accurate and more manageable. The improvements in these devices make standard security tasks easier, cheaper and more effective.

Unfortunately, tackling only the standard security tasks is very much the wrong thing to do, because the threat environment that the organization operates in has completely changed. New approaches are required to keep up with these changes, as documented in the diagram *Threat Environment Changes* on page 9.

Meeting these environmental changes is not as simple as ordering a device and dropping it into the network. These changes require rethinking and redeployment of existing tools, along with adjustments in security policy and in where security teams put their efforts.

Meeting the Challenge of Borderless Networks

Firewalls were originally designed as barriers between the Internet and existing enterprise networks. The thinking behind firewall deployment was based on a critical assumption: the existing network

THREAT ENVIRONMENT CHANGES

Borderless Networks	
The change: There is no longer a "crunchy shell around a soft chewy center" network environment. Instead, the network has become a web of intersecting connections, with branch-office virtual private networks (VPNs) and end-user remote access, connections to partners and vendors, multiple sites and Internet links, and differing levels of security within the enterprise itself.	**Effect on security:** The single-perimeter firewall is a thing of the past. Network managers must now consider multiple types of access control devices at multiple points in the network, turning the network itself into a "crunchy shell," a secured highway for transmission of critical data.
Pervasive Computing	
The change: The organization-issued desktop is now only a tiny subset of the devices connecting to the network. In addition to an organization's assets, such as embedded devices (printers, energy management systems, storage area networks, network appliances), staff bring with them their personal PCs, smartphones, tablets and netbooks — all connected to a ubiquitous wireless network, whether it's an 802.11 WLAN or some type of high-speed mobile telephony data service.	**Effect on security:** Organization-owned embedded devices have opaque security characteristics, and even fully controlled servers may be untouchable from a security point of view because of the complex wave of interlocking version dependencies in the applications they host. Network managers must adjust their security models to accommodate a huge variety of essentially uncontrolled devices. This may be assisted by increasing the reach of intrusion prevention technologies.
Cloud Computing	
The change: This technological innovation moves the locus of some or all of an organization's information assets outside the well-controlled data center and into the realm of someone else's security policy and deployment. As organizations look to cut costs and focus on their core operations, more and more applications are becoming candidates for outsourcing to the cloud.	**Effect on security:** Moving applications to the cloud also means refocusing security and changes in emphasis. Encryption enforcement, an afterthought in most networks today, takes on a critical role when valuable data is moving across the Internet. At the same time, access controls, authentication and authorization have to be firmed up, as old assumptions about what's "inside" versus "outside" are irrelevant. Audit and compliance requirements add another layer of complexity to any cloud outsourcing exercise.

was firm and unchanging. It didn't matter how bad or good internal security was because the only purpose of the firewall was to control traffic to or from the Internet.

Most enterprise networks are largely undifferentiated from a security point of view, and traffic flows across them unhindered. Because of the enormity of the task of addressing internal access controls, many IT managers haven't yet tackled this job.

The concept of borderless networks simply acknowledges what most IT and security teams have known for some time: Internal networks need internal access controls and greater visibility.

The most extreme example of internal access control is NAC. Different manufacturers have different strategies, but in general, NAC shrinks the border of the network down to the port level. Every point of connection to the network becomes a miniature firewall, with rules for access set up on the fly based on who and what is connected to that port.

NAC offers user-focused access control: What the user is allowed to do on the network is a function of who the user is (and what groups the user belongs to) and the state of the user's endpoint device.

Organizations have been slow to embrace NAC for a variety of reasons. It stretches the limits of technology at both the network switch and the connecting device. It requires that the access control policies actually get defined, a challenging

task in an organization where no access controls have existed in the past. And when NAC fails, it denies users access to the network, potentially blocking staff from getting their work done.

Even if NAC adoption has been slow, it remains the single most important advance in access control for enterprise networks. By moving access controls as close to the user as possible, the borderless network becomes secure. Network managers who aren't ready to take on the challenges of access control themselves can use some of the ideas and principles behind NAC to push access controls deep into their networks.

IT managers should be careful to match the level of mitigation to the threat to avoid the twin mistakes of underkill and overkill. For example, when connecting sensitive networks to the Internet, a continuous state of attack from hostile and aggressive outsiders must be assumed.

Therefore, a very strong type of access control in the form of a stateful firewall with extensive logging, high-end management, denial-of-service (DoS) mitigation and sophisticated proxies is needed (avoiding underkill). On the other hand, when trying to control access by trusted staff between WAN sites, a less sophisticated access control, such as a simple router or switch access control list, would be more appropriate (avoiding overkill).

When borderless networks call for increased installation of access controls, this doesn't mean one should install as many firewalls as possible, everywhere in the network. It could mean install some firewalls in some places, and use the security features in Layer 3 switches in other places.

Decisions on access control type and location also have to be made in the context of the organization's IT staff responsibilities. For example, if the network and security teams are integrated, access controls implemented as access control lists (ACLs) on switches are easy to accomplish and cost-effective.

If the security team is completely separate from the network team, additional hardware may be needed to apply access controls because of the former team's lack of direct access to the switch fabric.

The best approach to the challenge of borderless networks is a calm and deliberative one. Recognition must be given to the fact that an organization's networks may have been running successfully for decades without added security, so there may be no need for a sudden rush to disrupt a well-behaved network.

On the other hand, the organizational inertia of "why fix something that isn't broken" has to be overcome with the realization that the threat environment has changed, and the network has to change to remain secure.

Pervasive and Autonomous Computing

As the cost of building network-enabled computing devices drops, manufacturers are responding by flooding the market with products. From smartphones to Wi-Fi-enabled sneakers, from wireless thermostats to security turnstiles with network-enabled main controllers, and from printers to projectors, Ethernet and wireless connections are becoming ubiquitous in our day-to-day lives. IT managers now need to secure networks on which the majority of devices are not running Windows, are not joined to an Active Directory domain, and may not have installed antimalware.

At the same time that connectivity is rising, devices and systems are also becoming more autonomous, constantly performing actions on the user's behalf without being directed to do so. IT managers must anticipate this increase in autonomous behavior and provide proper security against the inevitable DoS consequences.

These challenges come both inside and outside the network security perimeter, from both staff and end-user computing devices. IT managers can respond to these challenges by increasing visibility and control of network traffic, using techniques such as IPS deployment, NBAD, and application and network whitelisting to allow only specified traffic to flow across the backbone.

Using Intrusion Prevention Effectively

IPS units are essentially firewalls turned inside out. A firewall follows a positive security model: Nothing is permitted by default, unless a rule has been added to specifically allow data to pass.

An IPS is the opposite, following a negative security model: All traffic is permitted by default, unless a signature is found identifying the traffic as malicious or dangerous and thus blocking it.

A firewall may have hundreds of rules; any IPS will start with thousands. Fortunately, those rules are created and updated by the IPS manufacturer.

IPS devices have been criticized for their

inability to block zero-day attacks and for how difficult they are to manage effectively. But the challenge that pervasive computing offers is not about zero-day attacks or detailed forensics, so much as it is about old attack vectors that aren't going away.

Microsoft's research in this area is very instructive. Each month, Microsoft releases a patch that is actually a malicious software scanner. The Microsoft Windows Malicious Software Removal Tool looks for known problems, alerts the user if malicious software is found, and also reports the results to Microsoft. This gives Microsoft unprecedented knowledge about malicious threats actually found on Windows systems.

Last year, seven of the top 10 threats identified by Microsoft's tool were more than six months old. Overall, 82 percent of the identified malware infections were for malware that had been on the list for more than six months.

It's impossible to translate malware threat statistics directly into IPS signature effectiveness, but the simple conclusion is that the largest number of infections aren't from new attacks but from old ones. This may change from day to day, month to month and incident to incident. And the cost of a zero-day attack can be high. But those are rare events.

What IT managers need to guard against every single day are the old attacks that are still floating around. Old-attack persistence is especially strong in embedded devices that don't have the same software updating capability or antimalware that more recent operating systems have. IT managers should use IPS units to help block and contain threats by deploying the technology at strategic junction points detailed in the *Key IPS Deployment Locations* diagram.

KEY IPS DEPLOYMENT LOCATIONS

IPS Location	Deployment Purpose
Inside the organization's firewall	In user-protection mode, this deployment blocks threats to end-user devices and embedded systems from Internet sources such as defaced web pages, phishing sites or malware-infected web servers.
The edge of the data center	In server-protection mode, this deployment blocks attacks that may be directed against organizational servers, from both Internet attackers and malicious or infected internal attackers.
Next to wireless LAN controllers	In high-sensitivity mode, accepting possible false positives, this deployment blocks any malicious attack or behavior that may originate from wireless networks.
Next to WAN VPN concentrators	This deployment watches for infections that may have originated in branch offices, with special attention given to branch-to-branch traffic.

Behavior Anomaly Detection

Network behavior anomaly detection is an evolving area of technology development. Security manufacturers are building the technology into stand-alone products and adding it to IPS solutions. It has a natural affinity with intrusion prevention technology because it uses the same negative security model. NBAD even appears in some firewalls, typically in the form of DoS protections.

The idea is simple: By looking at the behavior of systems, such as what ports they are connecting to, how much data they are moving and how often they connect, threats will become apparent without even looking at the contents of the traffic. Academics have shown how NBAD can be used to detect and block malicious traffic, but security vendors have been slow to turn this research into effective products.

The best success with NBAD has occurred in the area of DoS and distributed DoS (DDoS) detection and mitigation. IT managers operating mission-critical web services, in particular, should investigate DDoS detection and mitigation tools and services for their Internet-facing servers.

However, even small-scale organizations can benefit from DoS detection and mitigation, which is being added to most UTM firewalls. IT managers should methodically revisit their existing firewalls to uncover rate-limiting and DoS-evasion features that are built in, or update to recent firmware to get the latest features in this area.

Most NBAD tools depend on a constant stream of network flow data from existing network devices. Flow data comprises summary records that provide a snapshot of traffic across an interface, such as IP addresses, port numbers, and traffic byte and packet counts. Flow data doesn't include actual packet contents.

Cisco's NetFlow is the most popular data format, and led to the creation of the IP Flow Information Export (IPFIX) standard. IT managers will want to

ensure that all key devices, especially routers and core switches, can send NetFlow or IPFIX data to internal collector devices. This information is valuable not just for security purposes, but also for problem solving, debugging and network performance engineering.

With or without NBAD, IT managers should begin collecting network flow data and analyzing it with commercial or open-source tools to gain greater visibility into network operations.

Secure Cloud Computing

Cloud computing presents a significant change in organizational IT. Organizations may come to the cloud in many ways. It can start with software as a service (SaaS) applications such as hosted antispam services, customer relationship management (CRM) or e-mail. It can come with application development projects hosted on platform as a service (PaaS), or with computing and storage through infrastructure as a service (IaaS).

However an organization chooses to utilize cloud computing, security is a major concern. When applications and data move outside the firewall, across the Internet and into an outside data center, suddenly everything security-related gets much more complicated. IT managers should focus on encryption and identity management, and work with auditing and compliance teams to be sure that nothing falls through the cracks.

Obviously, every connection to cloud-based services and applications should be protected with strong encryption. Unfortunately, this point can be difficult to make with application developers who are accustomed to allowing unencrypted data to move across the network because it's simply "inside the firewall."

It's easy to get lost in arguments about which data and connections need to be protected and which are not sensitive, but these arguments can (and should) be short-circuited by blanket policies that require strong encryption everywhere, for everything.

It's simply more efficient to encrypt everything as a matter of policy rather than attempt to make individual decisions about different types of connections or data flows. And there's little reason not to — in this era of high-speed computing desktops, servers and security appliances, there's no performance hit from protecting everything.

Any encryption policy should be enforced by firewalls and web proxies. Organizations want to make it impossible to make an unencrypted connection to any type of cloud service. Where possible, this should be enforced on the cloud side as well, blocking any unencrypted connections.

IT managers may even elect for double encryption: building a site-to-site virtual private network (VPN) tunnel to the cloud computing provider and running encrypted session connections on top of that. The reason is simple: As users migrate in and out of the organizational network, they should not have to change encryption procedures, whether they are on the road, at home or in the office.

Encryption should also be mandatory for Internet-based services provided to the public or visitors — whether those services are in the organizational data center or from cloud-based providers. The safest way to preserve user privacy is to simply require secure connections for all websites and applications. The information being provided may well be completely public and intended for wide dissemination, but the fact that a trusted partner or constituent is downloading that information is not.

Most modern web applications and even static pages use session cookies to help provide tracking, customize pages and give a better user experience. Encrypting those pages avoids the type of headlines that some social networking sites have had to deal with when end users discovered what security experts already knew:

that session cookies could easily be captured over public wireless networks, making impersonation and data compromise a matter of a few clicks.

Adding encryption does have one downside: Network-based security tools such as IDS, IPS and data loss prevention (DLP) won't be able to handle encrypted traffic without special help, such as Secure Sockets Layer (SSL) decryption appliances that let these security monitoring and enforcement points continue to do their jobs.

Improving Identity Management

A 2011 study, conducted by the McAfee computer security firm, found that of the organizations that have rolled out enterprisewide directories, more than 80 percent base their identity management on Windows Active Directory. That's good news for IT managers, because having a single widely agreed upon standard makes it easier to link cloud applications with organizational authentication and authorization systems.

Identity management is a general concept that usually includes:

- Automating and tracking the provisioning (and deprovisioning) of accounts for users;
- Handling the problem of password resets and password synchronization across multiple directories or authentication systems; and
- Defining access controls or group membership across the organization.

Identity management can be as simple as a set of scripts, policies and procedures for handling account creation and deletion, or it can be a full-fledged solution from one of several vendors to help ease the task of adding and deleting users from directories and authentication systems.

Sometimes, identity management means solving problems by changing the way things work. For example, merging redundant (or partially overlapping) directories into a single directory is a common part of an identity management project, and can be done virtually with some identity management products.

A basic requirement of any organizational cloud application is seamless linkage to identity management systems. Organizations that have made the investment in public key infrastructure (PKI) or two-factor tokens for user authentication should screen cloud service providers carefully to be sure that their advanced authentication systems can integrate with cloud applications.

Many common SaaS cloud applications include authentication (validation of user credentials) only when they link to identity management. A better approach is to include both authentication and authorization information, although more sophisticated cloud-based applications with delegated management and complex role definitions may have trouble fitting all that information into existing organizational directories.

Opening up authentication systems to cloud-based applications does raise other security concerns. Once a cloud-based application accepts a username and password validated against an organizational directory, the potential for brute-force attacks increases.

An outsider can now start guessing usernames and passwords on the cloud application. This makes brute-force detection (often called break-in detection) a mandatory feature for identity management.

Security teams are already responsible for ensuring that solid password management is properly communicated to the user community. Cloud computing raises the stakes by stretching the web of trust beyond organizational boundaries. Sometimes even security professionals reuse passwords inappropriately across systems. Identity management should apply technological enforcement to these essentially human problems.

Compliance in the Cloud

Cloud service providers understand that they need to open up their logging systems to share information with organizations for debugging and compliance purposes. IT managers should make the linkage between external service providers to their in-house logging systems as early as possible to identify any holes that need to be filled.

Because cloud logging is hard to vary, the natural tendency is to collect more information than necessary, just in case. This highlights the need for good log management.

A security information management (SIM) tool, sometimes called security event management (SEM) or security information and event management (SIEM), can be useful for log management to meet compliance efforts and to understand the security posture of networks. These systems accept log messages from different devices, including firewalls, IPSs, vulnerability analyzers, hosts and cloud service providers, and correlates them to identify security events of interest.

In a network where hundreds of thousands of security events can be logged in a single day, SIMs are helpful tools that boil down the unmanageably large pile of events into actionable security information.

SIMs are especially useful in the process of securing borderless networks, because they provide a single source of information about access control violations. While many SIMs are installed under the rubric of a compliance project, the usefulness of the technology is far greater than simple compliance. And as the complexity of security deployments increases, both security and network managers will be able to pull useful information, reports and alerts from SIMs. ■

PROTECTING THE APPLICATION LAYER

New attack vectors require new security approaches.

Security threats are quickly moving up the network stack toward the application layer. Attackers are finding that web browsers, web applications and all of the associated helper applications are fertile ground for malicious actions. And sometimes, the application itself is the problem.

To address this growing area of concern, security vendors have focused on protecting the application layer. They've responded with two broad classes of products: end-user protective firewalls, which are now being billed as "next-generation firewalls," and application firewalls, used to protect servers and the applications that run on them.

Next-generation Firewalls

Traditional firewall technology is largely based on the transport layer of the network stack, using a combination of network-layer IP addresses, transport-layer protocols such as transmission control protocol (TCP) and user datagram protocol (UDP), and transport-layer port numbers to make access control decisions.

For traffic coming into the enterprise network, traditional technology has typically been good enough: Allow traffic to a specific IP address, on a specific port number. That approach safely controls incoming services such as e-mail, domain name system (DNS) and organizational web servers.

Outbound traffic, on the other hand, requires greater granularity in today's threat environment. Unfortunately, traditional firewalls don't offer IT managers much control. Organizations with loose policies often operate with a default outbound "allow" policy and have no idea what users are really doing.

Organizations with strict policies block almost everything, forcing users to send traffic through a proxy that can provide malware scanning and URL controls. While some proxies have a degree of application-layer intelligence, most simply add malware scanning, basic IP-based and URL-based access controls, and authentication,

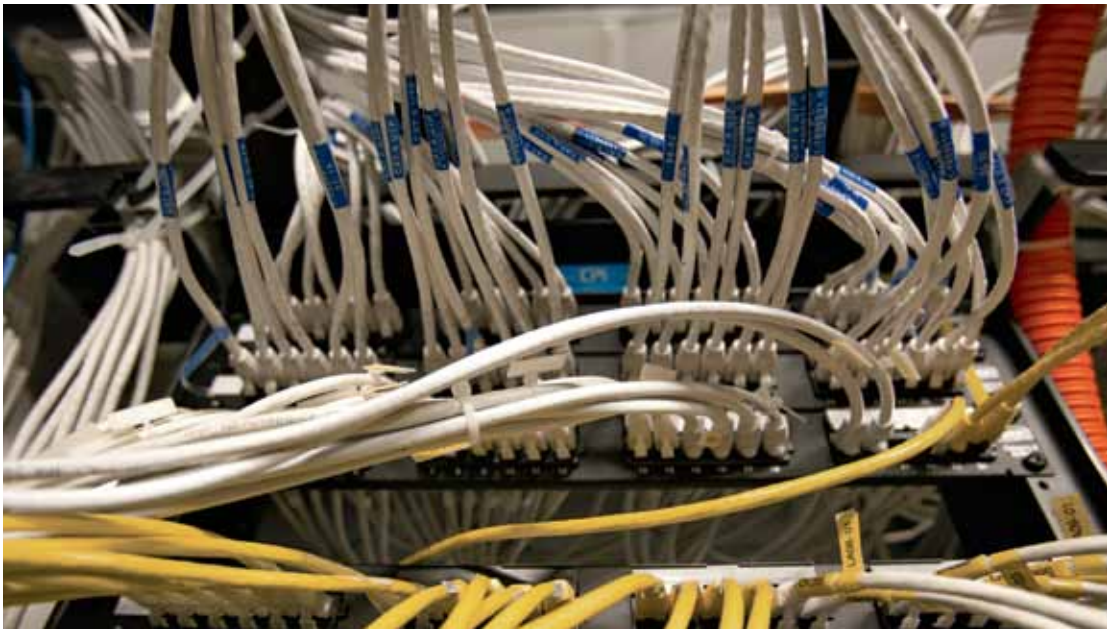

focusing almost entirely on web-browsing controls.

These additional security controls add value, but IT managers are asking for more control and more visibility. This is where next-generation firewalls (NGFWs) come in.

NGFWs take traditional firewalls and bundle additional security services, such as intrusion prevention and reputation filtering. These devices can push the function of the web proxy into the firewall. But the significant difference that makes these devices valuable is their application-layer controls and application-layer visibility.

Consider social networking sites. An IT manager may choose to simply block these websites entirely. This is easy to do using URL filters. However, security policy may recognize that there is value in some access to social networking websites.

An NGFW should have visibility into the social networking application itself. This lets the IT manager, for example, allow network users to read postings, but not to post information. Or the organization can allow users to read and post information, but not run associated games. Another option is to allow all social networking website activity, but to gain visibility into which users are utilizing this access and what parts of the sites they are using.

Another example where NGFWs prove helpful is webmail. Organizations with strict security policies often forbid the use of webmail. URL filtering will catch the big providers, such as Google Gmail and Yahoo! Mail, but cannot catch every webmail application. An NGFW that understands the common webmail applications can identify the application, no matter where it is running, and then apply organizational access controls.

NGFWs also go beyond proxies by identifying and controlling nonweb applications, such as peer-to-peer file sharing (BitTorrent), Voice over IP (Skype) and instant messaging (MSN Messenger).

NGFW capabilities are not for every network, and a network that has proxy servers for outbound traffic may already be well protected. But when granular application-layer access controls and visibility are needed to protect and manage organizational users, the features being built into NGFWs are the wave of the future.

Application-layer Firewalls

When placed in front of a well-managed server, a traditional firewall does not offer a tremendous amount of protection. The goal of the firewall is to permit or deny traffic to specific TCP/IP port numbers. But once the traffic is permitted, the firewall does not inspect the traffic for further threats.

The idea of application-layer firewalls is to add strong application-layer knowledgeable protections to a traditional firewall. For example, in a UTM firewall, the IPS can provide considerable application-layer protections.

In a full application-layer firewall, though, the firewall has a close knowledge of the types of applications it is protecting, which offers the opportunity for even stronger threat mitigation than an IPS can offer.

In the application-layer firewall market, web-application firewalls are the most common, with more than a dozen products on the market. These range from software products to security appliances to cloud-based services.

Architecturally, most web-application firewalls look more like reverse proxy servers than firewalls, and can provide protections that a normal IPS or firewall cannot. For example, an IPS always uses a negative security model: Every web transaction is allowed, unless a rule or signature blocks that transaction.

This is good for detecting automated attacks and attacks using known vulnerabilities, but it does not help against custom attacks and zero-day attacks. Some web-application firewalls employ a positive security model: Every web transaction is blocked, unless the firewall has a rule or signature to permit it.

Obviously, building positive security models for web applications can be difficult, because the firewall has to understand and allow all of the legal transactions of the web applications behind it. When an organization depends on the security of their web applications, though, this extra effort is worth the trouble. Better web-application firewalls have very advanced tools for building and updating their positive security models.

Because web applications vary so much from organization to organization, web-application firewalls do not all have the same feature sets and are difficult to compare. The Web Application Security Consortium (www.webappsec.org) is an unbiased starting point for anyone considering a web-application firewall.

One valuable resource that the Web Application Security Consortium offers is a Web Application Firewall Evaluation Criteria resource, which both helps explain what a web-application firewall is and assists IT managers in making educated decisions when selecting web-application firewalls. ■

ACRONYM SPEAK: IS A UTM AN NGFW?

If unified threat management (UTM) is a firewall plus security services, and a next-generation firewall (NGFW) is a firewall plus security services, why are there two acronyms for essentially the same product? The answer has everything to do with marketing and industry analysts, and nothing to do with technology.

Originally, there were just firewalls. However, about a decade ago, a tradition of adding security services to enterprise-class firewalls was firmly cemented when all of the firewall manufacturers added VPN tools to their products — effectively killing the nascent stand-alone, site-to-site VPN appliance business.

Over time, firewall manufacturers continued to add services until an industry analyst created the UTM moniker to help differentiate between products that were moving forward with additional services from those that were standing still. UTM meant "firewall plus antimalware and intrusion prevention, and maybe some other services."

However, IT managers of large networks weren't interested in running antimalware and IPSs in their firewalls for two reasons. First, they already had those security services in other dedicated products. Second, they didn't want to deal with the performance uncertainties that UTMs introduced.

Big networks didn't use UTMs; small networks did. Thus, over time, UTMs quickly became stigmatized as a branch-office or small network firewall technology.

But the need for additional security services in large networks continued. It morphed into a slightly different set of services, with emphasis on application visibility rather than IPS and antimalware. To resolve this conflict, a different industry analyst coined a different term: NGFW, focusing more on application-layer than network-layer security services.

Looking at the product feature list for UTM and NGFW devices, it's hard to tell the difference. Generally, when a vendor labels a device UTM, it is aiming at smaller networks. And when the vendor labels a device NGFW, it has a larger, enterprise focus — even if it's the same device.

DATA LOSS PREVENTION

Data must be protected from both inside and outside threats.

Most security initiatives focus on keeping visitors with malicious intent out of enterprise networks. Firewalls, IPS units and malware detection tools all focus on identifying and blocking threats. However, threats don't come only from the outside.

When someone inside the organization acts maliciously by breaking into systems or stealing data, the threat is just as significant. In fact,interior threats are the most dangerous to organizations, because they may involve trusted individuals with access to sensitive and valuable information, confidential systems and network pathways that are not carefully monitored.

Insider threats aren't all malicious. A staffer may mistype an e-mail address or not pay attention to a long e-mail with several embedded replies and accidentally send sensitive information out across the Internet. Or someone may launch a peer-to-peer application to download music without realizing that valuable documents found on the desktop are now subject to sharing as well.

Protecting against insider threats requires a combination of techniques, including education, policy and technology. One valuable technique that offers visibility into potential insider threats is to ensure that IPS appliances can see internal flows between servers and users. Simply placing an IPS close to the firewall doesn't track these types of threats, which is why most experts recommend placing IPS sensors in front of enterprise assets.

Data loss protection (DLP) tools zero in on one aspect of interior threats: the loss of sensitive data. These security tools don't guard against every type of insider threat, but they present an outstanding complement to other threat-focused tools.

Guarding Against Loss

DLP products fit into three primary categories: channel-specific solutions, endpoint and data-at-rest solutions, and network-based solutions.

Channel-specific DLP solutions: These products are integrated with other security tools, such as antispam gateways.

They are called channel-specific because they only look at one channel for data loss, rather than across the entire network. For example, a channel-specific DLP application integrated with an antispam gateway will be able to identify data leakage via e-mail, but it won't help at all with File Transfer Protocol (FTP) or web browsing.

Endpoint and data-at-rest DLP products: These tools work in conjunction with endpoint security products, such as desktop antimalware packages, to help enforce policy. The most common example of this kind of interaction would start with a prohibition against using flash memory devices (thumb drives) via USB.

If policy prohibits writing data to a thumb drive, then DLP capabilities in the endpoint can enforce that policy. The data-at-rest part of the DLP tool can also act by scanning network drives (and local hard drives on personal workstations) in search of sensitive data in the wrong places.

Network-based DLP products: These tools examine traffic passing over the network, and can report on (or block) transactions that violate organizational policy. For example, if policy does not allow personal ID numbers to be sent off-network, then DLP apps in the network can identify prohibited traffic and either block it or send a notification when it occurs.

The DLP market also makes a distinction between content-aware and content-neutral products. In the example above, blocking flash drive access is content neutral, while blocking personal ID numbers is content aware.

In practice, both content-neutral and content-aware loss protection occur in endpoint protection and network-based products. However, endpoint protection appliances tend to be less content aware because of the difficulty of making each endpoint aware of all the potential types of content that need to be blocked.

Some solutions have both endpoint protection and network-based components, while other tools work only in one of the two areas. The benefit of having an integrated solution is that a single policy can work across both endpoint protection and network-based tools. This simplifies the task of deploying DLP capabilities across the organization.

However, having the vastly different types of protections available in the endpoint and in the network can cause compromises. Organizations especially concerned about data loss should separate their network-based DLP solutions from their endpoint protection DLP solutions to gain the best protection in both areas.

DLP Tool Deployment

With channel-specific endpoint protection and network-based DLP products all widely available, there is no shortage of solutions to evaluate. A few basic strategies for deploying DLP applications will help to ensure greater success in reducing inappropriate information leakage.

When launching any DLP project, there are two critical points to keep in mind. First, DLP products are designed to help honest people stay honest. Someone intent on sneaking information out of an organization in defiance of policy will probably be successful. When information flows like water, it is difficult to stop every leak.

Second, DLP products are better at identifying leakage than stopping leakage. DLP tools can help IT departments identify users who are exposing information carelessly or against policy, and this can be valuable in educating end users and solving user behavior problems. Organizations should not expect DLP solutions to actually stop leakage at the moment it occurs.

Organizations without experience in DLP tool deployment should take advantage of consulting support because it will dramatically simplify implementation and help ensure even coverage across all product features. As with IPS deployments, there is no substitute for experience to speed deployment time and offer the most effective solution.

Begin with Policy

Most rollouts of DLP apps begin with the IT team evaluating solutions, such as USB control and protection, file or drive encryption for notebooks and desktops, e-mail content scanning, or network-based DLP products. That's the wrong starting point.

Instead, IT managers should start by identifying the main sources of operational risk caused by data loss and leakage. Then they should work to identify policies to help contain that risk. Only after the sources and policies are

identified should technology come into play.

DLP is a unique aspect of security policy because the most effective DLP programs are based on user education and training, not on technological enforcement. Adding technology to help people comply with policy and to identify when they break policy is an additional benefit. But if an organization doesn't start with clear identification of the sensitive data it needs to protect and the policies for protecting it, it will not achieve great success with DLP.

DLP initiatives must also include workflows for when an incident occurs. Every DLP tool will identify potential leakage; it's up to the organization to decide what to do with this information. A best practice is to involve non-IT security staff in DLP policy development and product selection.

Because human resource and legal departments will be responsible for the final resolution of some issues brought to light by DLP products, they should be involved in the process early on. In addition, the potential intrusion of DLP products (especially USB protection tools) into day-to-day staff operations can torpedo a project if it does not have broad support from the entire organization.

Protect Channels with Full-strength Solutions

Many security products contain channel-specific DLP capabilities, which can be a valuable part of an enterprise DLP strategy. However, channel-specific DLP is generally only sufficient in compliance efforts, where the goal is to comply with audits rather than solve a real security threat.

Using point solutions that perform particular security tasks, such as scanning e-mail or instant messages, may be attractive from a budget point of view, but they will cost more money and implementation time in the long run. Organizations need to look across their desktops, servers and networks to get full DLP coverage and give the protection the policy requires.

Generally, it is useful to treat the two main detection paths for DLP solutions (endpoint/data-at-rest and network-based) as having separate policies, even if they are part of a single DLP product offering. A single broad-based DLP solution for both endpoint and network threats covers two bases at once, unified under a single policy console. However, the policies for endpoints are generally different from network-based DLP, so there is little harm in separating the two functions.

A full DLP solution goes beyond endpoint content-neutral protection and adds content-aware detection. Content-aware detection integrates content discovery (such as identification of organization-controlled credit card numbers or personal identification numbers, or files with sensitive data in the wrong parts of the network) with the scanning of outbound traffic. To be effective, the DLP solution must scrutinize all types of traffic leaving the network, including e-mail, web traffic, file transfer and instant messaging.

SMALL STEPS, TAKEN CONFIDENTLY

The policy and workflow development aspects of DLP tool deployments can be challenging. At first, it's better to start protecting small bits of information to learn how well the tools work before trying to scale up to a final DLP solution.

DLP technology deployments will generate false positives and true positives. Experience with the tools is the only way to learn how to tune them to keep false positives — as well as true positives — to a manageable level. The lessons to learn are not only technological, but also procedural and organizational.

This suggests that a winning organizational strategy begins with a small start, followed by escalating deployment with more and more features and greater and greater policy coverage. Rather than going for the "big bang" deployment, success with DLP solutions comes from gradual successive refinement of policy and constant assessment of results against operational requirements.

Get Identity Management in Order

For most organizations, having a file blocked from outbound transmission is a loss-prevention win. However, few content-aware DLP tool implementations provide true prevention. This means that figuring out who sent a file or attempted to send it is extremely important in resolving incidents. This knowledge aids in the big picture of loss prevention through education and supports, if necessary, a change in access to sensitive data.

Successful DLP initiatives depend on a strong foundation of identity management. Knowing who is on the network at any moment and how to track an IP address to a person is a critical component of DLP strategies. ■

ENDPOINT SECURITY AND REMOTE ACCESS

End-user device security is constantly evolving.

When organizations have high-level hackers setting their sights on networks via undisclosed zero-day vulnerability exploits and advanced persistent threats, the mundane task of keeping endpoints safe becomes critical.

Savvy IT managers know that a system with well-protected endpoints is more secure than a system with out-of-date signature packs — or no protection at all.

Is there anything new to say about endpoint security and remote access? Absolutely. A combination of new information, new strategies and new endpoints makes the task of securing clients worth revisiting each year.

Antimalware Efficacy

With nearly identical core functionality, antimalware products compete for IT dollars based on obscure feature differentiation and intangibles such as customer relationships, support quality and volume pricing.

The best reason to choose a particular antimalware product is more difficult to gauge: How well does the package identify malware on desktops and notebooks? Such comparisons can be challenging because they vary over time from year to year and even from week to week. Efficacy also varies based on the deployment environment. Different organizations in different parts of the world are exposed to different types of malware.

With existing products installed on the network, it's easy to measure efficacy by asking a simple question: Are there malware infections on the network? If the answer is yes, then clearly there's a problem with the installed antimalware, and an evaluation of alternatives is appropriate.

To determine the right replacement, feature testing and evaluation is fairly simple. Discovering how well the core engine will detect malware in a particular, however, is a bit trickier.

A small number of independent test labs based in Europe provide unbiased information on the core efficacy rate of each product and are worth researching. IT managers should regularly evaluate their antimalware products to be sure that they have installed the right product, that it's doing the job and that it has a high efficacy rate.

Always-on Remote Access

One current mobility trend is to tightly integrate endpoint security with remote access. Endpoint security manufacturers have tried to build integrated clients, focusing largely on desktop security needs such as antimalware, host intrusion prevention and personal firewalls. Most organizations select a single integrated client rather than installing three (or more) separate clients on each endpoint.

The next step is to further integrate endpoint security with remote-access VPNs and wired/wireless supplicants. These features don't all go into a single client (yet), but are integrated from a management and compliance point of view.

When SSL VPNs first became popular, one of the significant differences between them and legacy IP security (IPsec) VPNs was posture checking, which validates

that the endpoint connecting to the VPN complies with the security policy of the organization. Checking the posture of the endpoint is also considered a common requirement for NAC products.

Mobility product manufacturers are taking the idea of controlling endpoint posture one step further. The traditional barrier between a VPN client and an endpoint security client is dropping, as these products become more tightly integrated and compatible.

This higher level of integration presents an "always-on" security strategy for IT managers. For example, location awareness (Is the user in the organization's building? At home? At a hotel or café hotspot?) can be integrated into the endpoint client security policy to provide more continuous protection, regardless of the environment.

This always-on strategy can even be used to require VPN connections to help protect the endpoint. Most IT managers have encouraged staff to use their VPN connections only when they are actively working. However, some are using a new strategy: bringing up the VPN connection all the time and routing all Internet-bound traffic over the VPN.

It's not a particularly efficient strategy, but with the high speed and low cost of most Internet connections, a little loss of efficiency may be worth it for the additional security it brings. By routing traffic back through the organization's network, all of the security policies and protections available to devices in the building can be extended to devices on the road or at home.

Combining endpoint security, remote access and wired/wireless supplicants into a single client does require a higher level of cooperation among different IT management groups, such as desktop support, security and telecommunications.

However, as organizations realize the strong overlap in the goals of these different groups, many of the previous problems associated with a tightly integrated remote access and endpoint security strategy are disappearing.

Not every organization needs to re-evaluate its endpoint security and VPN client footprints. But as staff become more mobile, and as organizations adopt a "work anyplace" strategy, finding ways to leverage tighter integration between VPN, wired/wireless supplicant and endpoint security offers a secure anytime, anyplace option.

Protecting Mobile Devices

It's obvious that the IT group must secure notebooks at home and on the road, but smaller devices such as smartphones and tablets also need protection. These devices can store plenty of sensitive data, making device loss as big a potential problem as losing a notebook.

Securing mobile devices requires a slightly different strategy from that of traditional desktops and notebooks. Begin with these five key steps:

1. **Start with policy.** Everything from device selection to deployment, use and retirement should be laid out so that expectations on all sides are understood.
2. **Require encryption for all data communications.** Trying to decide what needs security and what doesn't is a waste of time. Encrypt everything moving over the network, and get rid of the uncertainty. Unlike notebooks, traditional VPN connections may not work for all applications, so paying close attention here is important.
3. **Encrypt data on the device, and require passwords to unlock the device.** Most lost and stolen devices are quickly wiped and resold. Don't leave anything interesting lying around. This is one area where mobile devices are ahead of notebooks, although the new generation of hardware-encrypted hard drives is a move forward.
4. **Invest in malware protection and central management for mobile devices.** Threat vectors such as e-mail may be protected at e-mail gateways. However, Wi-Fi and Bluetooth both open up devices to attack from anyone who can get within a few feet. A little control goes a long way. Central management must include remote device wipe as well. The constrained operating system environment on a mobile device makes it easier to provide all-around protection than on a typical notebook.
5. **Require authentication.** Every device needs to have its auto-lock feature (with break-in detection and an auto-wipe after a certain number of missed passwords) turned on to get one more level of protection. ■

CLOUD SECURITY ADDS CONFIDENCE

When on the road, it's easy for busy staff to fall behind in their malware protection on mobile devices. One solution to this problem is to use cloud-based security software as a service to protect endpoints. The idea is simple: Tunnel the most problematic traffic (web browsing) to a cloud-based security system and let the provider take on the burden of staying up to date, filtering malware and even enforcing acceptable-use policies.

Using the cloud isn't a complete solution. Personal firewalls are still critical, and the provider can't protect what it doesn't see. But as an alternative, the cloud is becoming a popular option.

FINESSING PHYSICAL SECURITY

Collaboration between information security and physical security has benefits.

While physical security has many areas of concentration, one area that overlaps closely with information security is access control: authenticating users and granting access according to policy. As part of an overall convergence between physical and information security, organizations have been exploring ways to use the same authentication and access control systems.

Integrating these IT systems with physical authentication devices such as proximity badges, magnetic stripe cards and hardware tokens provides an organization with a more consistent view of security across all levels.

A second common area of overlap between physical security and IT stems from the increasing use of digital tools to implement physical security, such as surveillance cameras, networked physical devices (gates and door locks) and physical plant management (lights and HVAC).

Achieving better coordination between the teams implementing information security and physical security in each organization brings obvious advantages, along with some unanticipated benefits. In the area of compliance, for example, the ability to pair physical security data and information security data adds layers of assurance when demonstrating compliance with a particular regulatory regime.

Facilitating Security Coordination

Information security and physical security tend to be separate disciplines within organizations, even reporting to different members of the executive team. These groups may have little interaction. Successful coordination of these security divisions and their staff requires some careful planning.

The best approach is to establish formal lines of communication between the groups, including specific shared responsibilities. A monthly planning meeting is a good start, but for true success, the teams must have more frequent interaction and

clear operations-support tasks to accomplish.

Because the teams themselves may be separate (and likely have very different views on security within the enterprise), shared knowledge and shared tasks will help. Joining these teams under a chief security officer can help to better align the operational side of security with the strategic side of the organization.

One place to start combining physical security and information security is in information and education programs. All IT managers know that the biggest risks they have to worry about are human ones: staff members making careless mistakes or errors in judgment, trusted staff acting against the best interests of the organization, or staff serving as the leverage point for an attacker to gain unauthorized access.

Staff education is an important tool to help reduce risks, so every security group tends to create its own program. Bringing these together into a single, coordinated message can help increase awareness and improve security overall.

A second opportunity for coordination is to combine authentication systems, such as tokens or badges, into a single view and a single product line. By integrating physical and software authentication, organizations can avoid the embarrassing consequences of locking someone out of the building — but letting them log on and wreak havoc from home.

These are just starting points for discussion. Most organizations with a commitment to good security will find many opportunities for physical and information security teams to work together and add value.

Physical Security on the Network

The traditional tools of physical security — surveillance video, door lock controls, alarms, and access tokens and cards — often use equally traditional infrastructure: their own dedicated wiring, control systems and central data capture tools.

In an era of virtualization and widespread IP deployment, traditional physical security infrastructure doesn't make sense. Cabling for a specific application, such as video surveillance, is expensive and inflexible, while the data capture tools for applications, such as digital video recorders, are usually pricey and have low capacity.

Bringing physical security onto the enterprise network will reduce costs and increase overall flexibility. Switching to IP-based cameras gives the physical security team the ability to place a camera anywhere it can get an Ethernet port — in some cases, anywhere it can get a wireless signal. That means that reconfiguring a security monitoring system to meet changes in traffic flow or building design is as easy as plugging a device into the network.

At the same time, leveraging storage area networks (SANs) for digital video storage can give the physical security team the ability to store more data, more reliably, for less cost than it was paying for dedicated DVR devices. Upgrading to IP video also usually includes motion technology, which saves video only when something is happening in the frame — a huge savings in storage costs, and a very popular feature with the physical security team.

IT managers may find that linking up professionally with the physical security team offers the opportunity to add physical security within important areas such as the data center. For example, many IT managers would love to have video surveillance within the data center to help with compliance, documentation and troubleshooting efforts.

For the physical security team, adding more cameras is just adding more cameras — it's not a demanding project for them the way it might be for the information security team, which would have to start from scratch.

Cost savings from automating and networking physical security can also include staff changes. When all video is IP, central monitoring of multiple facilities is easier, and onsite physical security staff during off hours can be redirected to more valuable projects. ■

STREAMLINING AUTHENTICATION

In many organizations, there are nearly as many authentication databases as there are applications requiring authentication. On the software side, identity management techniques and application interfaces such as Remote Authentication Dial In User Service (RADIUS), Lightweight Directory Access Protocol (LDAP) and Security Assertion Markup Language (SAML) are being used to reduce the number of databases and systems.

This glossary serves as a quick reference to some of the essential terms touched on in this guide. Please note that acronyms are commonly used in the IT field and that variations exist.

GLOSSARY

Access control

Access control can be either a technique or technology for authenticating users and granting access according to policy.

Advanced persistent threat (APT)

This new term signifies targeted attacks that combine multiple threat vectors, often including zero-day vulnerabilities, to achieve a specific intrusion, control a specific set of systems, or acquire a particular level of access. Compare this with typical Internet-based attacks, which simply seek to exploit a particular vulnerability without caring about a specific target.

Autonomous computing

This term refers to computing systems that act and react on their own, adapting to environmental conditions without direct immediate input from the user or systems manager.

Borderless network

Sometimes called deperimeterization, borderless networks have many components and access methods that resist any attempt to create a single chokepoint for all ingress and egress traffic. Borderless networks also have varying levels of access control internally, reflecting the different levels of sensitivity and trust within an organization.

Cloud computing

Broadly defined, cloud computing offers organizations the ability to gain access to computing resources, on demand, using networks such as the Internet.

Compliance regime

This term refers to a lumping together of whatever rules, regulations and laws apply to a particular organization operating in a particular geographic region. Each organization has one or more compliance regimes, which may include a wide variety of official or industry sources with many different goals — all of which must be applied as a whole.

Data loss prevention (DLP)

DLP encompasses a family of security products aimed at mitigating the threat of sensitive or critical data being taken outside of organizational control. DLP products help to protect against both malicious and unintentional loss or leakage of sensitive information.

Denial of service (DoS)

A DoS attack prevents legitimate users from accessing system resources. DoS is a common attack method in which a target (server) is saturated with requests so that the natural processing flow is slowed or stopped entirely.

Endpoint

Network endpoints are devices,

such as workstations, notebooks, smartphones and PDAs (as well as printers and fax machines), that connect users to the network. They all require unique security consideration because of their access to the network.

Health Insurance Portability and Accountability Act (HIPAA)

HIPAA is federal legislation passed in 1996 that includes a privacy rule creating national standards to protect personal health information.

Identity management

Within IT, identity management handles the creation, updating and revocation of user credentials; the assigning of access rights and group identification to users; and the authentication and authorization of users to applications, operating systems and other IT systems.

Intrusion prevention system (IPS)

An IPS is an in-line device that inspects network traffic, compares against known signatures and blocks traffic that matches the signatures according to policy. When not in-line (or when blocking is not enabled), the same product should be referred to as an intrusion detection system (IDS).

IP security (IPsec)

IPsec refers to a suite of protocols that are used to secure IP communications via authentication and encryption of each IP packet within a data stream. IPsec supports both transport and tunnel encryption modes.

NetFlow

Originally developed by Cisco for the purpose of generating accounting information for ISPs, the NetFlow standard has since grown to become a heavily used management tool for providing sampled statistics about network traffic directly from routers and switches without the need for additional taps or sensors.

Network access control (NAC)

NAC refers to network-based, user-focused access controls embedded into network equipment. These access controls are dynamically established based on the identification of the user connecting to the network.

Network behavior anomaly detection (NBAD)

This technique protects networks based on statistical problem detection. NBAD technology is built into some IPS/IDS products, as well as stand-alone tools.

Next-generation firewalls (NGFW)

These new devices take traditional firewalls and bundle additional security services, such as intrusion prevention and reputation filtering. The key capability NGFWs provide is application-layer control and application-layer visibility.

Payment Card Industry Data Security Standard (PCI DSS)

PCI DSS is a set of security standards created to guide credit card processing companies in defending against fraud, hacking and other security threats. This set of standards has since been adopted outside of the credit card processing industry.

Pervasive computing

The organization-issued desktop is now only a tiny subset of the devices connecting to the network. Users connect to the enterprise network via home PCs, smartphones, tablets and netbooks. This is pervasive computing.

Platform as a service (PaaS)

PaaS is a type of cloud computing in which the cloud service provider offers a common platform for application development and deployment as part of its service.

Regulatory risk

Regulatory risk is a type of risk resulting from the failure to comply with a required regulatory regime. Rather than a risk generated by a negative event or a failure of threat mitigation, regulatory risk is suffered when a compliance failure (through an audit, for example) results in a penalty to the organization.

Secure Sockets Layer virtual private network (SSL VPN)

SSL is a remote access virtual private network technology that encrypts user traffic using the SSL and Transport Layer Security (TLS) encryption protocols. SSL VPNs are generally more flexible than IPsec VPNs, although leading vendors are building both technologies into the most recent wave of equipment.

Security information management (SIM)

Sometimes called security event management (SEM) or security information and event management (SIEM), the purpose of SIM is to gather security alerts and log messages from different devices, including firewalls, IPSs, vulnerability analyzers, hosts and cloud service providers, and correlate them to identify security events of interest.

Software as a service (SaaS)

The most common type of cloud computing, SaaS refers to a specific application that is completely owned, controlled and managed by the service provider but used by the organization renting the service. The most popular SaaS options are e-mail and CRM applications.

Unified threat management (UTM)

UTMs add antimalware, intrusion prevention and occasionally other services to basic firewall functionality. These appliances are typically targeted for branch offices or small networks.

INDEX

ABOUT THE CONTRIBUTORS

AARON COLWELL is an Inside Solution Architect specializing in network security for CDW. With over five years' experience in IT, he works with organizations to help them refine their security policies and practices and assists them with industry and government regulatory compliance. He often speaks at conferences and has been quoted in many industry publications.

PEYTON ENGEL leads a team of security engineers at CDW and has been with the company for over thirteen years. Since 2001, he has been responsible for team growth and management, including sales and marketing. Over the years, Peyton has presented his security research at national conferences throughout the country.

JOEL SNYDER, Ph.D., is a senior IT consultant with 30 years of practice. An internationally recognized expert in the areas of security, messaging and networks, Dr. Snyder is a popular speaker and author and is known for his unbiased and comprehensive tests of security and networking products. His clients include major organizations around the world.

LOOK INSIDE FOR MORE INFORMATION ON:

- Addressing borderless network threats
- Securing a cloud environment
- Defending against application-layer attacks
- Protecting endpoint devices