

SECURITY REFERENCE GUIDE



Effective responses to help keep business networks safe from threats.

CDW.com/securityguide | 888.509.4239



The Right Technology. Right Away.®

SECURITY REFERENCE GUIDE

TABLE OF CONTENTS

CHAPTER

WHAT IS A CDW REFERENCE GUIDE?

IT professionals are acutely aware of the need for network security. And at CDW, we're committed to getting you everything needed to meet your security requirements — from products and services to information about the latest technology. For example, our reference guides are designed to provide an in-depth look at the topics that relate directly to IT challenges. Consider them an extension of your account manager's knowledge and expertise.

01	Security as a Business Enabler	3
	• The Current Threat Environment	
	• Roadmap for Success	
02	Risk Assessment and Compliance	5
	• Determining Risk	
	• Mitigating Risk	
	• Compliance Considerations	
	• Security as a Process	
03	Improving Network Threat Defenses	9
	• Improve Visibility and Internal Controls	
	• Implement Identity Management Policies	
	• Don't Overlook the Basics	
	• Guard Against Application Attacks	
	• Consider Outside Help	
04	Data Loss Prevention	13
	• Data Loss Opportunities	
	• Loss Vectors	
	• Preventing Loss Through Technology	
05	Secure Remote Access	25
	• Varieties of Remote Access	
	• Threats to Remote Access	
06	Endpoint Security	29
	• Varieties of Endpoints	
	• Threats to Endpoints	
	GLOSSARY	33
	INDEX	35



SECURITY AS A BUSINESS ENABLER

CHAPTER 1:

The Current Threat Environment

Roadmap for Success

No doubt, CIOs and CEOs read about the latest IT security threats, data breaches and their associated losses. And they likely come away asking themselves a number of questions.

"Are we devoting the right amount of resources to information security? Are we focusing on the right areas with our security investments?" And, most important — and most difficult to answer — "Are we getting value for the money we are spending?"

Answering these questions is never easy. However, there is a strategy that will allow IT professionals to sit in the toughest meeting and provide defensible and understandable answers. It's not some arcane formula, and there's no particular magic to it. In fact, it's simple: Tie security investments and security strategies to the business, and everything else will follow.

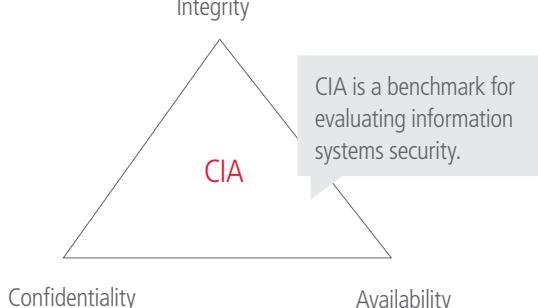
The idea is to make security a business enabler. When every security product, service and policy is mapped to the business, defending existing investments is straightforward. And it makes it easier to lay out a roadmap detailing future security direction.

Although the strategy is simple, getting there can be challenging. This is because it requires not only the security team, but also the entire enterprise to change the way it thinks about information security.

CURRENT THREAT ENVIRONMENT

Before enterprise networks, client/server computing and the Internet, IT security was a simple and straightforward discipline. Simply apply the traditional confidentiality, integrity and availability

(CIA) security triad to mainframe applications, add in physical access controls and impediments, and the issue is handled.



Unfortunately, scaling this methodology up to the world of the Internet and extranets, 100 gigabyte USB drives that fit in a wallet, and ever-multiplying online records and databases simply doesn't work. Clearly, the media is filled with reports of threats and data breaches.

Every time a significant security problem is revealed in the international press, security managers rush to identify the root cause of the loss and then collectively heave a sigh of relief that they have not made the exact same mistakes. Meanwhile, an ever-increasing array of regulatory regimes adds even more layers of security to burdened networks.

For businesses that haven't experienced a serious security problem, the question is often not "if," but "when." When will a privileged person decide to betray fiduciary trust? When will a determined attacker put a specific company in his or her sights?

What's more, when will the right data fall into the wrong hands because a car got broken into, or a phone was lost, or a notebook got left in a meeting room over lunch? When will someone click the wrong box in a graphical user interface (GUI) and misconfigure a firewall?

Many information security managers lose sleep knowing that they're only a few seconds away from a data disaster. Yet they don't have a clear strategy to reduce this risk. The reason they don't have a security game plan is they've isolated themselves from the business.

SECURITY FIRST

Security is rarely considered a core part of every project and every product. Instead, it's something separate, stuck in at the end after all the real work has been done.

Security isn't integrated; it's treated as a shell, added — often in haste and with tight budgets — to already-designed applications, networks and processes. Because security isn't built in, its costs become overhead, something with no clear return on investment (ROI) and an obvious target anytime there is a need for cost-cutting.

The result is usually less than optimal: Tools are installed without concrete measures of value. Security is seen as a blockade to business. And even with all that work, some security problems will remain because there is always some corner or edge left unprotected.

ROADMAP FOR SUCCESS

Business IT chiefs can set the right course through the world of information security by adopting a new model. Three key guideposts can serve as a roadmap to an effective security strategy:

Network Security Success		
1. Bake in Security	2. Tie Security to the Business	3. View Security as a Cycle

Bake in Security — Often core assumptions about security are wrong. Therefore, the most basic strategies need to be rethought and reengineered. Security must be baked into not only systems, but data and applications as well.

Whether talking identity access management, security information management or integrated threat management, security is rarely effective or efficient when it is added as an afterthought. Security must be part of every team's responsibility.

Properly applied, information security is not a business blocker. It should be a business enabler, not an impediment to getting things done. When security is baked into applications and networks at the right time and in the right places, it can be nearly transparent, yet dramatically reduce the potential for loss.

Security can't be layered on at the end; it has to be built in, at the beginning, to every piece of the business. The security team can advise on strategies and techniques, but the final responsibility for securing networks, applications and systems has to rest with the teams building and managing that infrastructure.

Tie Security to the Business — Security is all about reduction in risk. Proactively mapping out risks to the business makes it easier to see where security is needed — and where it isn't — making it easier to plan and justify security investments.

Good information security practice measures that risk and then finds a cost-effective way to mitigate it. If the risk is high, then a high investment is justified. If the risk is low, then the protection should match the risk.

Don't just look at the bottom line on your purchase order for products or services. Examine the total cost to the company. Sometimes, organizations will take on a larger share of risk than is normal to jump ahead of their competition or to get a product to market sooner.

Taking risks is fine, as long as the risks are calculated. So be certain to present these costs and options in a balanced and unbiased way.

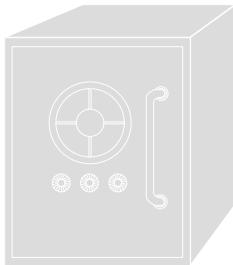
View Security as a Cycle — Think of security as a continuous process. Just like any investment, security investments need to be periodically reevaluated. The threat landscape changes continuously, which means that old products may not be providing the right protection.

Business risks are also constantly changing. The risks that caused concern and stimulated investment in security a few years ago may have dissipated — or they may have gotten much worse. Therefore, reevaluate assumptions that justified earlier security investments to see if they still hold true.

And finally, have clear visibility into the performance of security tools. This is a way to continually measure their effectiveness.

Following these principles of building information security will help in avoiding the missteps of many IT managers. They can assist in making the most effective investments to protect the enterprise and in successfully meeting goals. And they can assist in communicating those successes throughout the organization. ♦

RISK ASSESSMENT AND COMPLIANCE



CHAPTER 2:

Determining Risk

Mitigating Risk

Compliance Considerations

Security as a Process

All security is about reduction in risk. Security has a cost: a resource cost, such as people and capital; opportunity cost, such as products that are delayed because of the time it takes to incorporate security; and intrusion cost, which is the extent to which business processes are impeded from speedy execution by the addition of security.

At the same time, risk has a cost. If mitigating security measures are not in place and an event occurs, the enterprise could face direct expenditures and long-term indirect harm including damage to its reputation, fines, devalued stock price and/or loss of customers.

Simply put, good security is that for which the cost is lower than that of the risk which it mitigates. Bad security is that for which the cost outweighs the cost of the risk.

Comparing costs is easy — just subtract one from the other. But measuring costs is difficult. To see how hard it is to match security investment and the cost of risk, look no further than a local airport.

Anyone who has flown can recite a litany of security measures related to aviation. There are security guards to pay as well as equipment that must be bought and maintained. Valuable airport space is taken up by all of this — space that could be used for more productive purposes.

And there are more subtle costs. Business travelers must come to the airport hours before the flight, thereby often increasing unproductive time. Increase the aggravation and wasted time related to flying and some people simply won't do it, or will do it less frequently. This will change aircraft load ratios and schedules. It's a complex set of variables.

Calculating the cost of security, though, is a cakewalk compared with calculating the cost of risk. Suppose airport security was like security at a typical city bus station: Walk on the bus whenever you want, no metal detectors, no guards.

How many more security incidents would occur? How often would they occur? And what would the cost of the incidents be? How many people would avoid travel because of the risks, and how would that affect costs?

The answers to these questions range from the actuarial to the emotional. Obviously, measuring the risk cost is a complex endeavor.

The good news: No business IT security decision is as complicated as aviation security. With quality information from the business side of the firm, good decisions about where — and how much — to invest in security can be made.

DETERMINING RISK

While assigning metrics to risk management is difficult, a popular framework can help IT managers make risk mitigation decisions.

Start by identifying assets that are at risk. These could be physical or virtual IT assets, such as information or notebook PCs. They can also be more intangible assets, such as reputation and goodwill. Then identify threats or risks to these assets.

Consider the risk of an attack that takes down an external website that hosts e-commerce. In this example, let's use an externally generated attack, such as a successful SQL injection attack, that causes the back-end database to be corrupted.

Bases and Consequences of Risk



First, determine the value of each asset. In this example, it's easy because an e-commerce site generates revenue around the clock. Figuring out how much the site is worth can be accomplished by looking at the profit or margin figures from that site.

For each threat, come up with two numbers. One is the exposure factor (EF) and the other is the annual rate of occurrence (ARO). For example, if the company relies on a single e-commerce server in a single location, then a successful attack has an EF of 100 percent.

On the other hand, with two servers in two different data centers in two locations, the EF could be lower because the attacker may not actually hit both data centers. If operating a recovery data center that wasn't active unless the main servers were down, the EF would be 50 percent (at least as long as the attacker didn't go after the recovery site). If the data centers were load sharing, it might be somewhere in the middle, maybe 75 percent.

With EF and asset value, single loss expectancy (SLE) can be calculated: Given the value of an asset per minute, hour, day or week, multiply the EF by the value and factor in downtime. That's how much this threat will cost the company.

Note: Other qualitative and quantitative factors may enter into the equation. For example, damage to customer goodwill, brand equity and company reputation could affect asset value. Timing of an incident can also be critical. An e-commerce site outage on Black Friday would have a greater affect on asset value than one happening on a different day.

The second metric for each threat, ARO, is how often a particular

threat will happen in a given year. If it happens once a year, that's 100 percent. If it happens every other year, it's 50 percent.

Multiply the SLE by the ARO to derive the annual loss expectancy (ALE). That's the amount of money the firm will lose every year if the threat is not mitigated.

MITIGATING RISK

Knowing that ALE costs can be effectively calculated, begin to look at risk mitigation. Websites are susceptible to SQL injection attacks because they are sometimes poorly written.

Therefore, one mitigation might be better training for application programmers. Or a third-party software audit firm could be engaged to find these types of bugs first, before a malicious attacker does. External risk mitigation tools could also be installed, such as application-specific firewalls or intrusion prevention systems (IPSs).

Each risk mitigation approach will alter the EF and ARO and, hence, the ALE. Keep in mind, some risk mitigation solutions are designed to minimize the consequences of incidents.

For example, let's say the firm invests in better programmer training designed to reduce security problems and have a third-party software audit. That might change the ARO from once a year to once every five years. In other words, ARO goes from 100 percent to 20 percent.

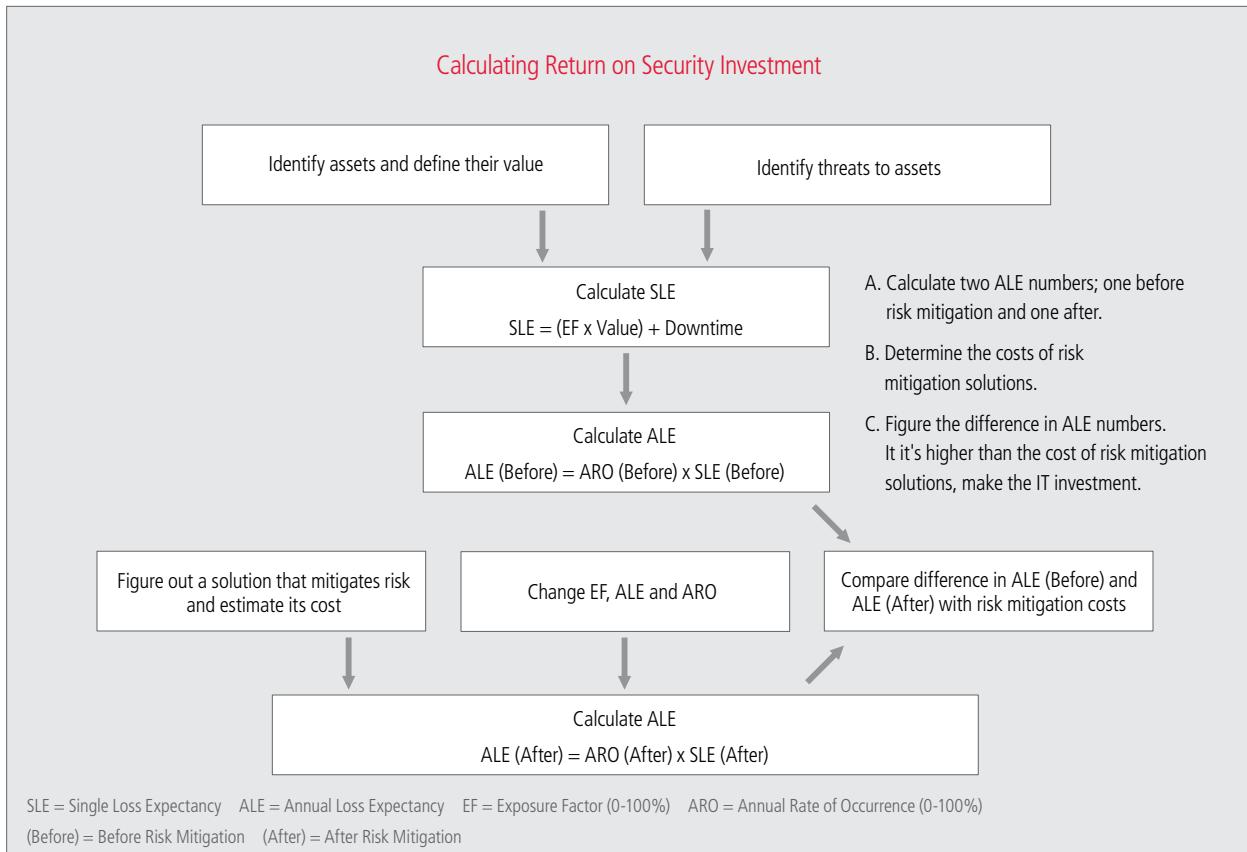
On the other hand, adding web-application firewalls, intrusion prevention technology and better programmer training will reduce the ARO further, maybe to 10 percent or even 5 percent. This is because multiple risk mitigation technologies provide distinct areas of coverage, even if there is some overlap.

The calculation will result in two ALE numbers: one before risk mitigation and one after risk mitigation. Hopefully, risk mitigation costs, including initial capital, maintenance and eventual replacement expense, will be known.

Take the difference in the ALE numbers, and see if it's higher or lower than the cost of the risk mitigation solutions. If it's higher, go ahead and make the IT investment — by spending money, you will actually save money. But if the costs of the risk mitigation solutions are higher than the difference in ALE, then the investment won't pay off.

It's pretty obvious that calculating these numbers can be challenging. While there is a significant body of academic research on the effectiveness of risk mitigation technologies, applying that research to a specific environment often involves estimation.

For example, does an IPS along with an application-specific firewall take ARO to 5 percent or 10 percent — or even 15 percent? But as hard as it is to determine these numbers,



operating without them can make it difficult to justify security investments.

QUADRANTS OF RISK

It can also be useful to classify security measures in terms of the ways that they reduce risk. Recall that risk can be defined as annualized loss expectancy, which is the product of the frequency and the severity of security incidents.

These two factors function as “amplifiers” of risk: If either of them increases, risk increases. Likewise, decreasing the probability of security incidents or limiting their impact, risk decreases.

Along the same line of reasoning, we can think of threats and vulnerabilities as “sources” of risk: Security measures are applied to reduce risk by negating one or both of these factors.

Both of these pairings — threats/vulnerabilities and frequency/severity — can be viewed as axes on a coordinate plane. See page 8. This allows use of the quadrants (labeled 1, 2, 3 and 4) to better understand the nature of the security tools.

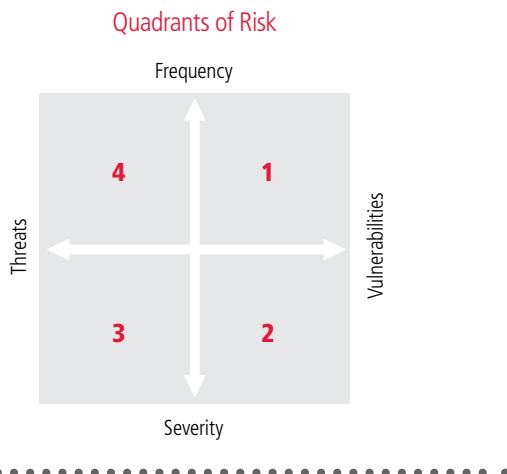
Every security measure addresses two of these factors. Here are some examples fitting into each quadrant.

QUADRANT 1. Patch management: A critical element of a strong security program, patch management ensures that systems are not running software with known vulnerabilities. It reduces the likelihood of system compromise by lowering the frequency of security incidents.

QUADRANT 2. Breaking password trust relationships: Once an attacker gains access to a system, the next step is to see where else that access might lead. On the assumption that users tend to use the same password in multiple places, an attacker will typically crack the passwords of system users and try them out on other likely targets.

Breaking password trust relationships between systems — whether by enforcing specific password policies, implementing two-factor authentication or eliminating local password databases — falls into quadrant two. These measures don’t necessarily prevent system compromise, but they drastically reduce the impact of such an event.

QUADRANT 3. Back-end database restructuring: Many web applications start out as customer-facing front ends to databases designed for an organization’s internal use. Examples include



employment websites where job seekers register profiles to receive notices of new jobs, social-networking sites allowing members to post personal information and the like.

Such web applications become an attractive target for attack by identity thieves because the database may contain personal information. Restructuring the website's back-end database to include only information needed for web functionality doesn't address whatever technical vulnerabilities the web application may have. It nevertheless reduces the impact of a security incident by making the website a far less interesting target for attackers.

QUADRANT 4. User education and awareness: Ensuring that users remain well-informed doesn't address any technical vulnerability. However, making them wary about existing threats reduces the probability that they will fall prey to fraudulent e-mails.

Each of these quadrants represents a different approach to the same goal: reduction of risk. The field of threat prevention deals with quadrants three and four. It realistically acknowledges that systems will always have vulnerabilities and focuses on keeping risk at a tolerable level by finding ways to address potential exploits.

COMPLIANCE CONSIDERATIONS

Recently, compliance has become the focus of a great deal of security effort. Compliance deals with the extent to which an organization conforms to a given standard.

Actual regulations form the basis for some compliance initiatives — the Health Insurance Portability and Accountability Act (HIPAA) for healthcare-related entities, the Gramm-Leach-Bliley Act (GLBA) for financial institutions, the Sarbanes-Oxley Act (SOX) for publicly held entities and various state privacy laws.

At other times, industry drives the standards. The Payment Card Industry Data Security Standard (PCI DSS) is the most notorious example of a nonregulatory compliance standard.

It's important to bear in mind that while all these standards encourage good general security practices, there is a fundamental difference between compliance and security. Security always deserves attention beyond just the minimum effort needed to comply with a standard.

SECURITY AS A PROCESS

Security is not a static consideration that can be quickly addressed and then forgotten. It's an evolving, never-ending process, on both an organization-wide scale and at the level of individual initiatives.

Initially, this notion may be a bit overwhelming. Like any other organizational activity, though, security can be seen as a process that consists of a series of steps or phases. This process, or lifecycle, consists of four phases: design, implementation, testing and monitoring.

- 1. Design:** A system's security begins with a definition of its functional requirements and the development of a design that's able to meet them. Given that security is a fundamental property of a system, taking security into account at the time of design is essential. Adding security features on as an afterthought is ineffective and costly.
- 2. Implementation:** Even the best design may not turn out as planned when executed. Therefore, it's critical to pay attention to security concerns as a design is translated into a working system. It's difficult to think of every detail of a project before starting, so be prepared during implementation to take course corrections as needed.
- 3. 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. A security assessment can serve as an important sanity check or quality assurance gateway before a new application or network site goes live. Testing is an opportunity to answer the question: "Do our security measures accomplish what they're intended to do?"
- 4. Monitoring:** Once a system is deployed, it should be monitored in order to detect and respond to security incidents and to measure whether the security is mitigating risk. Nearly everything in the IT environment has the capacity to keep some sort of log or send alert messages.

Monitoring also includes reevaluating the cost of security — whether this solution is costing more or less than expected. Monitoring security elements will indicate when it becomes necessary to tweak system design or implement changes. Then test to make sure the changes are successful. This cycle must be ongoing if security is to be effective. ♦



IMPROVING NETWORK THREAT DEFENSES

CHAPTER 3:

Improve Visibility and Internal Controls

Implement Identity Management Policies

Don't Overlook the Basics

Guard Against Application Attacks

Consider Outside Help

Everyone involved with IT has witnessed the increasing variety and virulence of threats. As attackers have moved from joyriding to profit-motivated harvesting of systems into "bot" armies, the sophistication of attacks has increased.

At the same time, growing corporate reliance on the Internet also increases vulnerability. In fact, according to the Symantec Global Internet Security Threat Report, Trends for 2008, published in April 2009, "Malicious activity has increasingly become web-based with attackers targeting end users instead of computers."

According to the experts, even economic distress, tight budgets and reduced staffs are not reasons for businesses to ignore security concerns. In times of escalating attacks, IT managers must bring their "A" game to security and risk mitigation.

IMPROVE VISIBILITY AND INTERNAL CONTROLS

Just when IT chiefs thought that companies at high risk for data theft had their online gateways locked down, it seems another mind-boggling breach occurs. In fact, the consensus among industry experts is that 2009 was worse than 2008 in terms of data loss. And that's saying a lot.

According to the "2009 Data Breach Investigations Report," conducted by the Verison Business RISK Team, 285 million records were breached in 2008. And this number exceeded the combined total of exposed records from 2004 through 2007. This obviously shows that current network security initiatives are far from foolproof.

Most IT security is based on blocking controls at the perimeter. For example, firewalls and antimalware programs filter traffic before it enters the network. Perimeter controls focus on the most toxic part of the IT world: the Internet.

Using perimeter controls is a great strategy. However, they don't tell much about network traffic and whether it's secure on the inside. Getting to the next level of IT security requires improved visibility and internal controls.

Visibility means knowing what is happening on the network from a security point of view — which can also mean knowing what's happening from a management point of view. Control means enabling control points throughout the network, not just at the perimeter, to direct and manage traffic.

Control at the edge is a start, but not sufficient for today's networks with their semipermeable borders, branch-office virtual private networks (VPNs), mobile devices and near-universal connectivity. Adding controls into the network changes it into a secure organizational asset.

IMPLEMENT IDENTITY MANAGEMENT POLICIES

The insider threat to security is another important reason to increase both visibility and control beyond what perimeter protection provides. According to John Kindervag, senior analyst for security and risk management with Forrester Research, "Insider threats are always big because people have unfettered access to all kinds of resources they shouldn't."

Social Networking and Web 2.0 Security

Web applications used to facilitate interactive information sharing present new security challenges for businesses. Along with postings to popular sites like Twitter, Facebook and LinkedIn, these can include Internet forums, blogs, wikis, podcasts, pictures and more.

Companies often wish to control employee use of the Internet. This is either because some uses are deemed inappropriate or because of the greater risk of malware in some types of websites.

A primary tool for controlling web usage by end users is URL filtering. When this security tool is in place, either as a dedicated web security gateway or as part of a unified threat management (UTM) device, each URL is compared against a database and categorized. The security manager can allow or block traffic based on broad categories.

Another new development in Internet security is the application-aware firewall. It is able to identify a wide variety of web and non-web applications.

IT managers sometimes find that Web 2.0 sites are handled poorly by URL filtering. This is because many different applications may be layered on the same website, and only some of them might be considered inappropriate.

Application-aware firewalls, on the other hand, can see further into the data stream of these multiapplication websites and apply additional controls, such as bandwidth management. Application-aware firewalls are especially useful when close management of user Internet access is needed.

Employees typically have greater access to data than is required to perform their jobs. Therefore, many will find valuable information simply through curiosity and "looking around" at company computer resources.

Obviously, staff should have access to the resources they need to accomplish their job. However, there is a fine line between allowing excessive access and being restrictive to the point of rendering employees unable to do what they need to do. In many cases, this trade off comes down to procedural issues, and eventually becomes a human resources issue.

A second, related issue is managing the risk posed by a terminated insider. This is someone who may want to extract retribution from the former employer either by taking information and redistributing it or simply causing damage.

While many security strategies can apply to insiders as well as outsiders, identity management (IDM) policies and procedures can help close the gap. One way to start is by answering these three questions when developing security policies and procedures:

1. How does IDM interact with the hiring and firing process?

For greatest security, IDM should be tightly integrated so that employees are given access to only the systems and physical spaces they need to do their jobs, with a corresponding removal of access when their position no longer needs it. In addition, many IDM products can track users within their authorized areas of responsibility. Therefore, should a malicious event occur, it can be linked to a specific individual.

2. How does IDM interact with nonemployee access?

Many nonemployees, including contractors and consultants, may wander through a facility as easily as staff. Guest users, another category of nonemployee, should have no IT access but may be given Internet or wireless access. IDM should account for these types of nonemployees by restricting access somewhat, but not so much that staff are tempted to share credentials or passwords.

3. Quis custodiet ipsos custodes?

Or, "Who watches the watchers?" Any IDM solution must account for the dichotomy of IT staff members who must be able to administer servers, yet should not be able to see the information on those servers.

DON'T OVERLOOK THE BASICS

When working to improve the quality of an organization's security posture, the basics of firewalls, antimalware and good patch discipline cannot be ignored. While these technologies haven't changed much, there are a few differences to be aware of:

- **Firewalls are now unified threat management (UTM) devices.**

Vendors have added antimalware, intrusion prevention, URL filtering and other security services to their products. These are all available, typically for a small subscription fee. However, be careful of turning capabilities on willy-nilly, as the impact on network performance can be dramatic. Furthermore, it can be cost effective to consider the security offered by a UTM first, especially when hardware has not been upgraded.

- **Robust endpoint security suites.**

Now it's client endpoint security, which adds in personal firewalls, host intrusion prevention and even network access control (NAC). Manufacturers have done an excellent job of integrating many different security products into a single console and a single client.

- **Patch discipline is more than just worrying about Microsoft's "Patch Tuesday."**

With many applications from

numerous vendors spread across a number of systems, and patches coming out on a daily basis, don't take a chance with a fragmented patch management system.

GUARD AGAINST APPLICATION ATTACKS

IT veterans have likely heard about all sorts of bad things people can send your way, such as Smurf attacks, pings of death, IP source routing and fragmented IP packets. But none of those matter anymore because companies have done such a good job securing their networks with perimeter firewalls.

The significant attacks are now up at Layer 7, the application layer of the network. This has happened for two reasons. One is that IT drove the attackers there by solving simpler network- and transport-layer security issues. The other is that users now inhabit a vast forest of applications, many of which are built on dozens or even hundreds of other, often open-source components.

At the same time, there's an equally vast, desolate desert of application security that most developers won't bother — or don't know how — to cross. The result is that there's not much left to attack at the network layer. However, there are ample opportunities at the application layer.

Take, for example, the use of SQL injection and persistent cross-site scripting flaws in websites to compromise end-user workstations. See the illustration below. A typical dynamic website divides its functionality into tiers, with an inner database tier, a

middle application tier and an outer web front end.

SQL injection and cross-site scripting attacks corrupt the web front-end and tamper with the database. This can allow hackers and thieves to vandalize and replace web pages, steal credit card numbers and other private data, and manipulate databases.

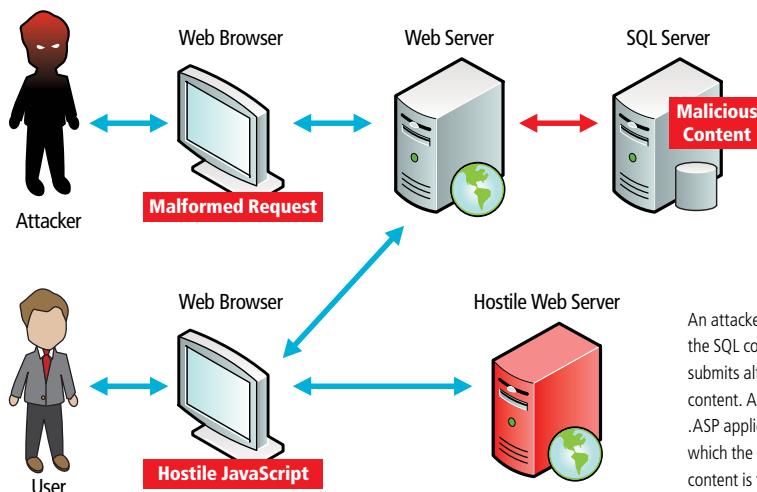
For example, an attacker constructs a request that appends a piece of HTML code to the end of every last name in a database, exploiting a bug in the web front-end. Each web front-end retrieves the user's name, polluted with the additional HTML, and sends it to the user's browser to welcome them to the site.

In this case, the HTML includes a reference to a second, hostile web server, retrieving a fragment of JavaScript in the user's browser. These types of attacks have been used to plant malware or spyware on the user's PC allowing access to sensitive data like credit card or Social Security numbers or even medical records.

A key technique for dealing with application security flaws is to use intrusion prevention systems and application-layer firewalls. An IPS is a network device that monitors network activity for malicious or undesirable activity and can react to block or prevent it in real time.

Application-layer firewalls act to protect web and other application servers from network-based attacks. These are usually distinguished from IPSs (and their brethren, intrusion detection systems, or IDSs) because they are focused on a narrow range of applications, such as web or SQL applications.

SQL Injection/Cross-Site Scripting Example



An attacker constructs a request designed to tamper with the SQL composed by the .ASP application. The application submits altered SQL to the database, inserting malicious content. A user makes a standard website request and the .ASP application issues normal SQL to the database — of which the contents have been altered. The malicious content is then returned.

If running business-critical web services in an Internet-facing data center, without absolute confidence in the ability of developers to write great code, then an application-layer firewall should be part of a network security strategy.

IPSs and IDSs can both prevent and detect intrusions. However, they have additional benefits, including the ability to discover security policy violations, infected systems, misconfigured applications and firewalls, information leakage, and unauthorized servers and clients.

CONSIDER OUTSIDE HELP

For most companies, security isn't a part of their business. Security is a necessary support function, not unlike printing paychecks and purchase orders or keeping the lights on.

When the investment in IT begins to become disproportionate to the other needs of the organization — or when there is an inability to manage the risk with the resources available — it's time to think about alternatives.

Many organizations have begun to look for ways to shift burdensome security tasks, such as monitoring security events or managing patches, from their own IT staff onto managed security service providers (MSSPs). The cost efficiencies of this move have begun to make economic sense.

Hosting and managed service offerings have become mature, stable commodities. This fact allows organizations to weigh the cost of a service provider against the benefits of freeing their own staff for other work.

Risk transference is a second motivator for this strategy. If, for example, an enterprise can move all handling of payment card transactions to a third-party service provider, it can greatly reduce the scope of the process — and therefore also the cost — required for PCI DDS compliance.

Testing can be another reason to consider outside help. Security testing is often quite different from ordinary functional testing. As a result, outside expertise is frequently needed to make sure that the testing covers all the areas it should address and that it is sufficiently rigorous.

SECURING VIRTUAL ENVIRONMENTS

The strong trend toward virtualization in the data center brings the obvious question of network security: What do we need to do differently in a virtual environment? While network security with virtual systems is no different from security around physical systems, the short answer may be nothing. However, in some ways, the verdict is still out.

For the most part, security is the same. However, it's not clear, for example, whether it's safe to use one virtual machine (VM) host to house both a VM guest that resides on the DMZ and a VM guest that resides on the internal network. In theory, it should be impossible for a compromised VM to attack the hypervisor and get to other VMs. Still, the history of network security is littered with disproven theories.

While virtualization offers numerous cost-saving and efficiency-building benefits, experts recommend staying attuned to special security issues that could put data at risk.

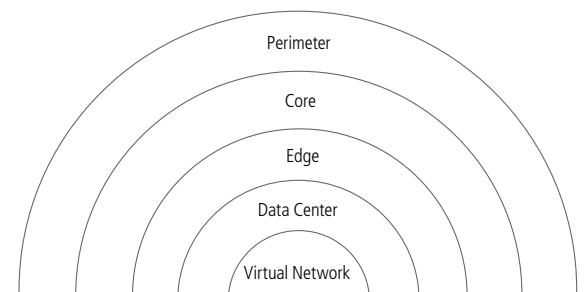
It's also worthwhile to consider the twin issues of performance and availability. These are even more important in a virtual environment than they are in a physical environment.

This is true because more applications are packed into a smaller space. Therefore, it's important to design security for the highest availability possible. This will typically include high-availability pairs of firewalls, built to maximize uptime.

More systems in smaller spaces can also mean greater network stress on individual links. If 10Gbps links are inconceivably fast for a single physical server, they certainly aren't overkill when 20, 200 or even 400 virtual servers are packed into a single cabinet. Network infrastructure upgrade, including firewalls, may be needed to handle multiple 1Gbps links or 10Gbps links into virtualization hosts. ♦

Network Layers Requiring Security

While virtualization offers numerous cost-saving and efficiency-building benefits, experts recommend staying attuned to special security issues that could put data at risk.





DATA LOSS PREVENTION

CHAPTER 4:

- Data Loss Opportunities
- Loss Vectors
- Preventing Loss Through Technology

When businesses take stock of their digital assets, they quickly conclude that data is their most valuable resource. For a few companies, the primary asset may be the availability of some service or the capacity for certain throughput. But as a general rule, servers are purchased, applications are designed and networks are built with the goal of collecting and manipulating data.

And if information is the chief asset, we need to keep at least some of it confidential. As various members of a firm handle private data, it tends to accumulate in unexpected places. And it may be handled unsafely or disclosed inadvertently. Data loss prevention deals with these risks.

DATA LOSS OPPORTUNITIES

Let's walk through a brief scenario. See the illustration on page 14. Staffer A works in the home office of corporation XYZ. Staffer B works in the firm's accounts receivable department in another state. One afternoon, Staffer B calls Staffer A with a question about the latest purchase order from Acme Corp., a customer that has been having trouble with payments.

Because of the time difference, Staffer A isn't at his desk, so Staffer B leaves a voicemail. Unified messaging brings the voicemail to Staffer A's smartphone, from which he listens to it. Staffer A then uses his phone's web browser to check the status of the order via the company's extranet, and forwards an e-mail with the relevant information from his inbox to Staffer B.

This scenario represents a typical transaction in today's fast-paced

work environment and serves as a perfect example of how well-integrated technologies enable projects to move forward, even when people are out for an espresso. A closer look, however, reveals potential problems.

The aforementioned diagram falls far short of identifying every possible juncture where data can leak. However, it does demonstrate that each step in the conversation involves some sensitive information that is subject to potential compromise.

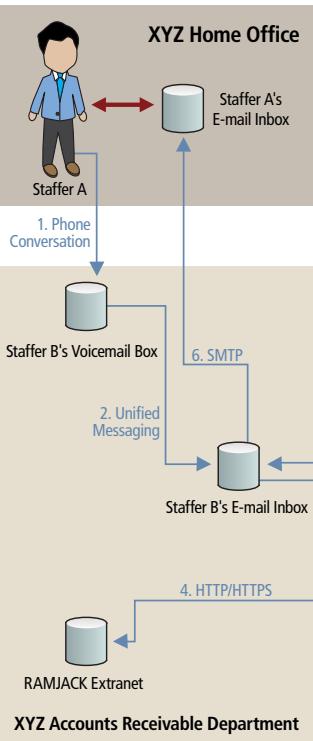
The diagram delineates at least three zones in the transaction described above: the home office, the remote office and the coffee shop from which Staffer A handles part of his job. It also clearly shows that portions of critical information migrated from zone to zone in the process of answering Staffer B's question.

Every time a file or document is stored or moved, the potential exists for it to be intercepted or simply left open for later discovery. Data loss prevention usually focuses on restricting the flow of private information across organizational boundaries. We want to set policies that govern how information is handled.

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

LOSS VECTORS

Sadly, most data loss is self-inflicted. In many cases, data falling into the wrong hands is a result of what a company's own users



Opportunities for Data Leakage

accidentally discard as side effects of their day-to-day digital lives.

Information simply proliferates too quickly and in too many ways for users to keep track of it. Most data leakage pathways have two fundamental things in common:

- Technology that makes it very easy to exchange information
- Security limitations on the technology that make it unsafe for certain types of information transmission

Risk increases when users handle sensitive data in unsafe ways. This is because they don't know better, because they think it will be expedient or because operations processes force them to do so.

Exhaustively listing all the ways that information can leak from the enterprise is impossible. However, most experts offer a short list of vectors that should be given additional attention, including:

- Removable and portable media, such as USB drives, CDs and DVDs, and digital media players. These media are easily lost or reused without full erasure. (Note: Today, highly effective data security methods are available for protecting portable media so as not to infringe on the incremental productivity they offer.)
- Portable devices, such as smartphones. These are treated differently from corporate assets such as notebooks, but they

may carry the same information unsecured.

- Electronic mail. The ease with which e-mail chains can leave a company containing sensitive information deep inside is obvious. But the propensity for people to attach rich documents such as spreadsheets and presentations is also a concern.
- Peer-to-peer communications, such as instant messaging (IM) and social networking; and personal writing, such as blogging and microblogging. These are all avenues where information can be disclosed without corporate intent.

PREVENTING LOSS THROUGH TECHNOLOGY

The wide variety of potential leakage vectors has created a huge market for products that claim some sort of data loss prevention (DLP) benefit. While these products cover a wide range of features, a few basic strategies for DLP deployment will help to ensure greater success in reducing inappropriate information leakage.

There are two critical points to keep in mind when beginning any DLP project:

- DLP products are designed to help honest people stay honest. Someone intent on sneaking information out of an organization, in defiance of policy, will very likely be successful. When information flows like water, it is difficult to plug every leak.
- DLP products can be used to identify leakage as well as stop leakage. DLP can help to identify end users who are exposing information carelessly or against policy. At the same time, it can be used to set different lifetime security levels on data, specifying required actions, i.e., allow but report, quarantine for review before acting or block all together.

Begin with Policy: Most DLP deployments start by looking at products, such as USB control and protection, file or drive encryption for notebooks or desktops, e-mail content scanning or network-based DLP products. However, that's the wrong starting point.

Instead, begin by identifying the main sources of business risk caused by data loss. Then, identify policies to help contain that risk. Only after the sources and policies are identified should IT start looking for technological solutions to assist in supporting these policies.

DLP is an especially interesting piece 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.

Keep in mind, it's essential to start with clear identification of the sensitive data needing protection and policies for protecting it.

Without this, success cannot be assured.

Deployment policies must also include incident management, response and procedures. Incidents should be routed to appropriate handlers based on the policy violated, breach severity, the data user plus other criteria. As is evident, DLP is designed on protecting information and, therefore, requires a dedicated workflow.

Involve non-IT security staff in policy development and product selection. Since HR and legal departments will be responsible for the final resolution of some issues brought to light by DLP products, they should be involved early on. In addition, the potential intrusion of DLP products (especially USB protection tools) into day-to-day operations can torpedo a project without having broad support from the entire organization.

Use a Full Strength DLP Solution: Many existing security products have DLP features. But without an organization-wide view of data at rest (sometimes called content-at-rest, content scanning or content discovery), data in motion leaving the network, and a corresponding data protection policy, DLP efforts will be fragmented at best.

A full DLP solution integrates content discovery on the network (such as identification of organization-controlled credit card numbers, personal identification numbers or files with sensitive data in the wrong parts of the network) with scanning of outbound traffic (typically combined with outbound web proxy).

To be effective, the DLP solution must look at all types of traffic leaving the network. These include e-mail, web traffic, file transfer and instant messaging.

Using point solutions that may solve particular problems, such as scanning e-mail or scanning instant messaging, may be attractive from a budget point of view. However, this will cost more money and implementation time in the long run. A single broad-based DLP solution covers multiple bases at once and unifies it under a single policy management console.

Get Identity Management House In Order: For most enterprises, knowing that a file was blocked from outbound transmission is great loss-prevention news. But knowing who sent that file is even better information, because that helps in the big picture of loss prevention through education and, if necessary, a change in job responsibilities.

This means identity management must be under control. Knowing who is on the network at any moment, and how to track an IP address to a person, is a critical part of DLP deployment.

When a DLP solution includes content discovery, identity management is even more important. This is because a file with inappropriate information found on a network share somewhere

has to be tied back to a person.

Take Small Steps, Confidently: DLP deployments can be difficult, not so much because of the technology involved, but because of policy and workflow development. It's better at first to start protecting small bits of information. This enables learning how well the tools work before trying to scale up to a final DLP solution.

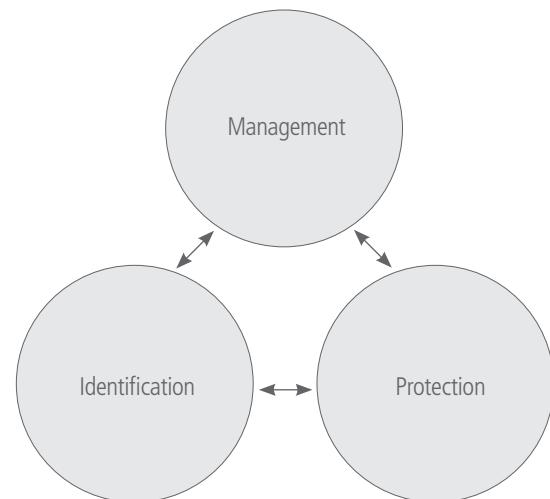
DLP technology deployments will generate false positives and true positives. Only after some experience with the tools will IT managers be able to tune them to keep false positives — and true positives — to a manageable level. The lessons to be learned are not only technological, but also procedural and organizational.

Start small, with alert-only or monitor-only policies. Then escalate deployment to active blocking and policy enforcement as toolsets are learned and trusted, and business requirements are met.

With proper implementation, DLP technologies can offer a plethora of information regarding data being stored, transferred and used across the network. Analysis of this information allows a firm to determine where to apply extra layers of protection. ◊

DLP Functions

Data loss prevention systems offer workings to implement three specific functions: management, identification and protection.





SECURE REMOTE ACCESS

CHAPTER 5:

Varieties of Remote Access

Threats to Remote Access

A few IT professionals still remember the days when it was essential to be in the office in order to perform most work functions. Today, staff is expected to keep up with some tasks while on the road or at home, and notebooks have become standard-issue in many companies.

This means that people are accessing organizations' networks remotely. Therefore, remote access must be secure. And there are a number of issues key to ensuring this security.

Perhaps the most important is that many companies don't have a good handle on all the types of remote access they need to be concerned about. (Another important issue, traffic interception, is covered in Chapter 6, because the problem of eavesdropping isn't unique to remote access.)

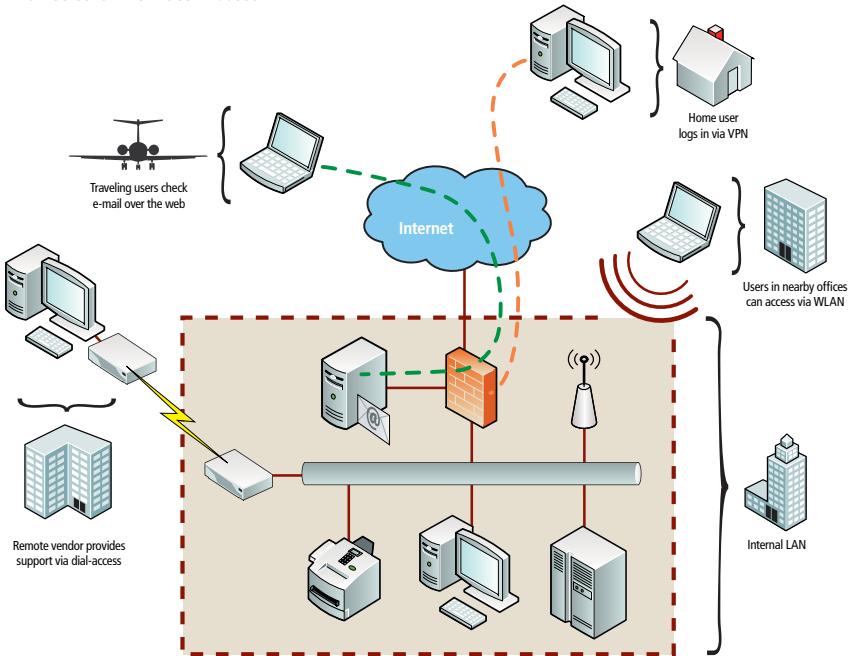
VARIETIES OF REMOTE ACCESS

Strictly speaking, remote access encompasses any access to private network resources from beyond the firm's physical perimeter. This broad characterization includes some modes that most organizations don't take into consideration in their security planning.

As noted in the diagram at the right, remote access can include a number of scenarios:

- Users may log via VPN to do work on the firm's internal systems with notebooks and desktop PCs. These devices may be managed by the firm's IT staff, or they may be personal devices, such as a home PC. This access normally comes through a VPN concentrator, which typically runs either Internet Protocol Security (IPSec) or Secure Sockets Layer (SSL) protocols to protect the data stream.

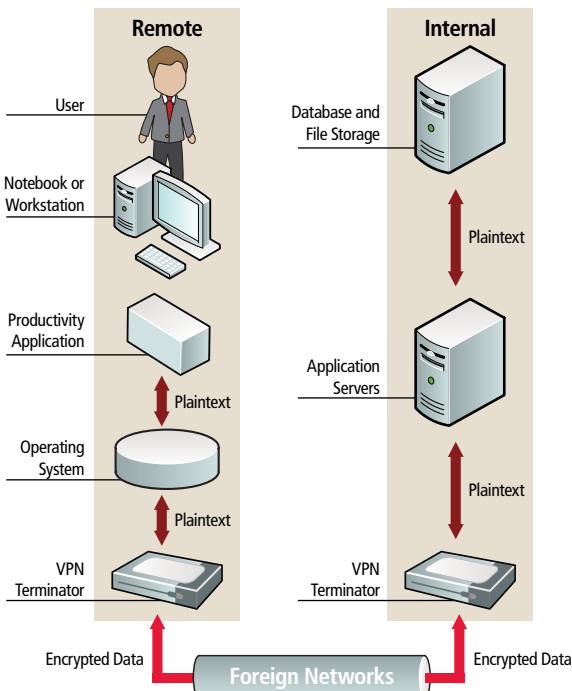
Varieties of Remote Access



- Vendors and partners may have access to support systems on the firm's network. This is often accomplished with VPN connections over the web, and may include dedicated circuits. This remote access is distinguished by the level of access granted, which is usually very restricted. Otherwise, the technology may be identical to that described earlier.
- Mobile staff may use devices such as PDAs and smartphones to gain access to particular applications. This access often requires a special-purpose gateway specific to the mobile device, such as a BlackBerry Enterprise Server.
- Traveling users might make use of lightweight remote-access solutions such as the web, e-mail or various support tools. These systems may not be thought of as being true remote access because they are intentionally made accessible through the firewalls. At the same time, these systems often pull information from within the network, granting accessibility to data that is considered internal.

Each of these modes of remote access might have many variations. For example, SSL VPN can be used to grant network layer access (often called network extension), single IP/port access (called port forwarding), or web-only application extension (called reverse proxy), all in the same VPN concentrator.

VPN Remote Access



While the specific technologies in play at a given business each have different strengths and weaknesses, their purposes remain similar. The goal is to facilitate access to internal resources without exposing them to risk.

THREATS TO REMOTE ACCESS

Having identified the basic types of remote access, it's time to start thinking about what might go wrong. Envision a typical remote-access scenario — in this case, a VPN — as shown in the figure at bottom left.

The diagram depicts the various participants in a remote-access event: the remote side, the internal side and whatever networks carry information between the two. Both sides have some device or software that manages the remote connection.

Typically, this task is handled by software at the remote user's end and an appliance on the organization's end. However, variations abound. The illustration highlights the fact that 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.

Remote-access solutions are subject to a variety of threats, most of which are independent of the technology used. For example:

Denial of Service (DoS) Attacks — These can overwhelm the VPN concentrator so that users cannot connect. They can also cause account lockout by intentionally sending an incorrect password over and over for a remote-access user.

Traffic Analysis — An eavesdropper can learn a great deal from encrypted traffic, despite the fact that the actual contents of packets might not be accessible.

For example, someone sniffing a wireless network can see what the endpoints of the conversation are. They might also be able to deduce where and when users work, bank and shop.

Dictionary Attacks — These are used to guess the passwords of remote-access users. Where account lockout is in place, as it should be, these can be run as "low and slow" attacks that may not trigger IDS and log analysis alerts.

Some of the threats to remote access can be mitigated by smart system design. For example, the use of time-based password tokens dramatically reduces the possibility that someone will be able to guess or steal a password.

Other threats are nearly impossible to protect against. For example, a truly determined attacker with a "bot" army at their disposal can lock up any VPN concentrator by throwing sufficient resources at it.

When designing a remote-access solution, the security analysis

should focus on the greatest risks. Because remote-access users are geographically dispersed and because they may be connecting over wireless networks, the most significant risk to secure remote access is the man-in-the-middle attack.

Man in the Middle — The threat of impersonating the remote access system itself must be considered. In a man-in-the-middle (MITM) attack on a remote-access system, the adversary convinces the remote user that an imposter system is, in fact, a 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, as shown in the figure at right.

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. Fortunately, a well-designed remote-access solution can entirely mitigate the potential for a man-in-the-middle attack.

SECURING REMOTE ACCESS

The simplest way to secure a remote-access solution is to assume that every aspect of the remote path is compromised and untrusted at all times. If a design accounts for this, then the company won't be vulnerable to predictable (and unpredictable) threats.

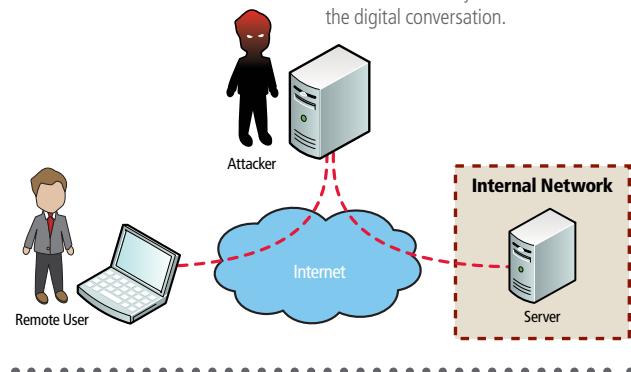
For example, much has been made in the press about "evil twin" access points. These wireless access points mimic either commercial Wi-Fi services or even a firm's own enterprise APs. With a proper design, your answer to this is, "So what?"

If end users are connecting via commercial Wi-Fi services while on the road, the assumption should be that the entire path is untrusted. In that case, it doesn't matter whether the access point is legitimate or an "evil twin," when assuming that all traffic is compromised and should be watched.

For MITM threats, use these five guidelines:

1. All remote-access solutions must include bidirectional authentication. In other words, the remote-access VPN concentrator should authenticate the user, but at the same time, the user should authenticate the remote-access VPN concentrator he or she is connecting to.
2. All digital certificates in a firm's deployment must be issued either by a trusted third party or, better, your own corporate certification authority (CA). Users should never see a certificate error at any time because these errors are the first sign that a MITM attack is ongoing. Indeed, telling a user to click "OK"

Man-in-the-Middle Attack



in response to a certificate error should be considered a firing offense among security staff.

3. Remote-access clients and VPN concentrators should have nonessential CA information removed.
4. Strong authentication — either a two-factor authentication system or digital certificates — is important for end users. Windows Active Directory has made issuing digital certificates to end users fairly simple. Still, most enterprises prefer to use two-factor systems, such as time-based password or challenge-response tokens, or one-time-password scratch cards or tokens. Shoulder surfing is a persistent problem in remote-access environments, and password reuse must also be guarded against.
5. When given two or more configuration choices, always select the most secure possible. Security parameters in the VPN concentrator should be selected to increase security, rather than performance.

For example, when selecting encryption algorithms, high key length Advanced Encryption Standard (AES) is always appropriate; Data Encryption Standard (DES) is never the correct answer.

With the high-speed processors available in today's notebooks and PCs, there is no reason to select a less secure algorithm. Old concerns about export controls on long key length for authentication also do not apply anymore, which means that certificates should be generated based on key lengths of at least 2,048 bits.

When in doubt about security, it's always less expensive to ask for help from a consultant during rollout than it is to clean up after a break-in. Locking the doors before there is a disaster is always preferred to the alternative. ♦



ENDPOINT SECURITY

CHAPTER 6:

Varieties of Endpoints

Threats to Endpoints

The true value of technology lies in the degree to which it enables coworkers to handle information or engage in transactions that would otherwise remain inaccessible. Despite the fact that the servers that store and process information may be locked in the data center, the vast majority of our meaningful interactions with computers takes place on network endpoints: workstations, notebooks, smartphones, PDAs and the like.

As users, we type in passwords, connect via public networks, view sensitive reports, engage in private communications and generally conduct business on network endpoints. As IT professionals, we recognize that the involvement of network endpoints in the storage and processing of sensitive information means they need protection.

VARIETIES OF ENDPOINTS

Before diving into what it takes to secure a network endpoint, it's important to define what an endpoint is. The figure on page 30 depicts a fairly typical network ecosystem and calls attention to the variety of systems that might be considered endpoints.

The list of endpoint varieties is longer than most people expect:

- Workstations, terminals and notebooks are the most obvious endpoints and represent the main means by which users interact with network resources. As a general rule, end-user systems on an organization's network belong to the firm and have standard operating system builds, security policies, antimalware protection installed and so on.

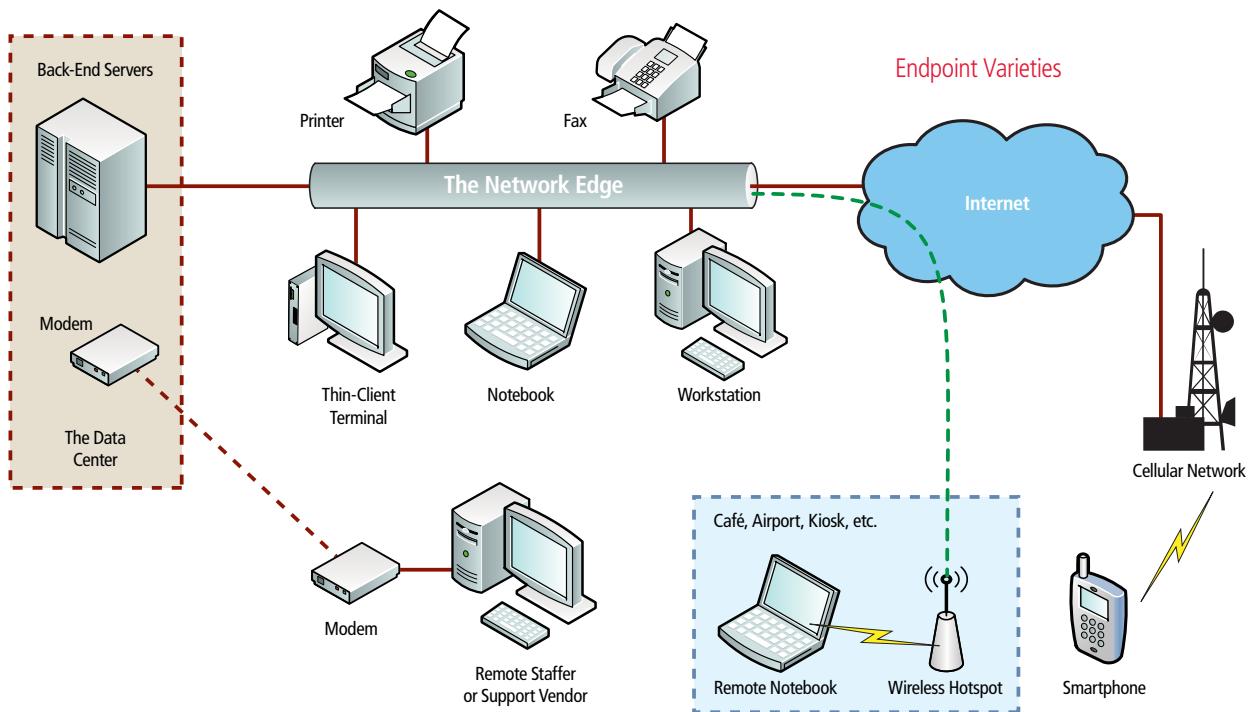
- Workstations and notebooks can also be used remotely and might not belong to the company or be under the company's complete control. For instance, a staffer might check e-mail from a notebook while traveling; an IT worker might log in from his home PC; or a vendor might connect remotely to support a particular product or system.
- Though often overlooked, network-connected multifunction printers serve as endpoints as well. When you print or fax through such a device, sensitive information can leave the safe boundaries of the organizational network.
- Smartphones and PDAs also represent endpoints. Although less fully featured than workstations, they have large storage capacities and can be used to handle confidential files or e-mail.

Each of these families of devices has different security limitations. Consequently, varying constraints exist around what network and system administrators can do with them. With this in mind, we can dive into the threats facing network endpoints.

MITIGATING THREATS TO ENDPOINTS

In today's business world, securing corporate endpoints is increasingly essential. Not only are threats omnipresent, complicating matters is the ever-expanding mobility of endpoint devices.

Because IT tends 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 just as surely as attacks on the central resources themselves.

Equipment Loss

The loss of physical equipment (euphemistically referred to as "unauthorized acquisition"), while unsettling, is insubstantial compared with the loss of the information that it houses. Nearly all endpoint devices store some valuable information, whether proprietary data, internal communications or cached passwords. When equipment is lost, repaired, recycled or stolen, this information is at risk.

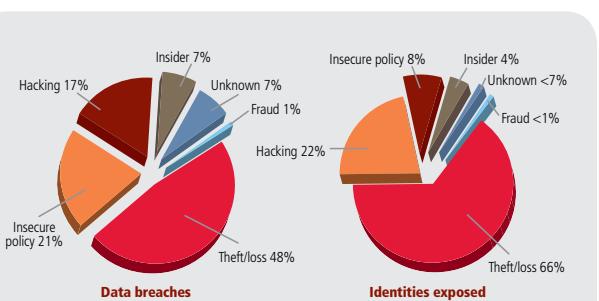
Most multifunction printers contain some sort of persistent storage where images of recently processed documents are cached, perhaps along with network credentials. The potential exists for someone in physical possession of one of these devices to extract sensitive company data from it, or have the images retransmitted via fax or e-mail at the touch of a few buttons.

Encrypting the data on a physical device or portable media can help limit the exposure to the business should it be lost. It's another case of risk mitigation minimizing the consequences of incidents.

(Note: If a physical device is powered on when lost or stolen, any

data in use is likely to be unencrypted. And the same may hold true for recent keystrokes, call logs and the like.)

Some strategies, such as whole-disk encryption, involve encrypting all data on a device, while others selectively encrypt only sensitive



Computer loss leading cause of data and identity theft

In 2008, the primary cause of data breaches that could facilitate identity theft was the theft or loss of a computer or other medium on which data is stored or transmitted, such as a USB key or a backup medium.

Source: "Symantec Global Internet Security Threat Report — Trends for 2008," published April 2009.

Diagram: Based on data provided by OSF DataLoss DB.

materials. The difficulty with the latter approach is that it's not always easy to ensure that all sensitive information receives protection. In addition, security experts have questioned the ability of partial-disk encryption or folder-by-folder encryption to protect data in complex operating system environments.

When investigating system- and disk-encryption products, look for features such as key escrow or shared keys. They would allow the IT department to decrypt data in case the original key — or original key holder — is unavailable.

Be on the lookout for fantasy encryption claims as well. These include claims that an encryption algorithm is better than standard and analyzed algorithms like AES.

Subversion and Infection

As noted, encryption might serve as a suitable protection for information stored on powered-off systems. But information is generally unencrypted when in use.

Spyware, for example, generally has access to whatever materials the logged-in user can see. And bots running as background services might be able to access anything that the operating system has access to, including files on other devices on the network.

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.

As in any security situation, one can never presume that any effort will prove 100 percent successful. Instead, one should approach the compromise of network endpoints as a realistic possibility and plan for it with strategies such as security information management, network segmentation and network admission control.

On the software side, protection against subversion and infection begins with a broad category of products now called endpoint security products. Most of these were marketed for years as antivirus tools, but with the addition of personal firewall and host-based intrusion prevention — and the broadening of antivirus to include multiple types of malware — the familiar antivirus moniker has changed.

Recently, the task of selecting endpoint security tools has turned into a beauty contest. Product evaluators have found that dominant products are largely identical in feature and form, yet have slight differences in management style.

Of course, each product will have moments when it leapfrogs the competition with a new feature or coverage of a new type of

threat. But over the past three years, endpoint security products have converged to a high and fairly uniform standard.

That's good news for IT security managers, because the choice of an endpoint security product can focus on the features that really matter in a particular environment, or issues such as cost, vendor relationship and product extensions, rather than concerns about whether the product will work.

When choosing and deploying endpoint security tools, pay particular attention to some of the newer features being added, such as NAC, application whitelisting and mobile-device support. In these areas, compatibility with existing infrastructure is critical, and there are still differences that may turn into showstoppers for a particular business.

Eavesdropping and Impersonation

Network endpoints rely on other network resources in order to be productive. This characteristic means, as a rule, that sensitive information (the files and data that the endpoint handles, as well as authentication credentials) necessarily traverse the network.

Consequently, unencrypted network traffic can be vulnerable to some degree of eavesdropping. Passively listening to network traffic reveals a great deal of private data, system administration practices and even passwords for services, as anyone who has ever operated an IDS has learned.

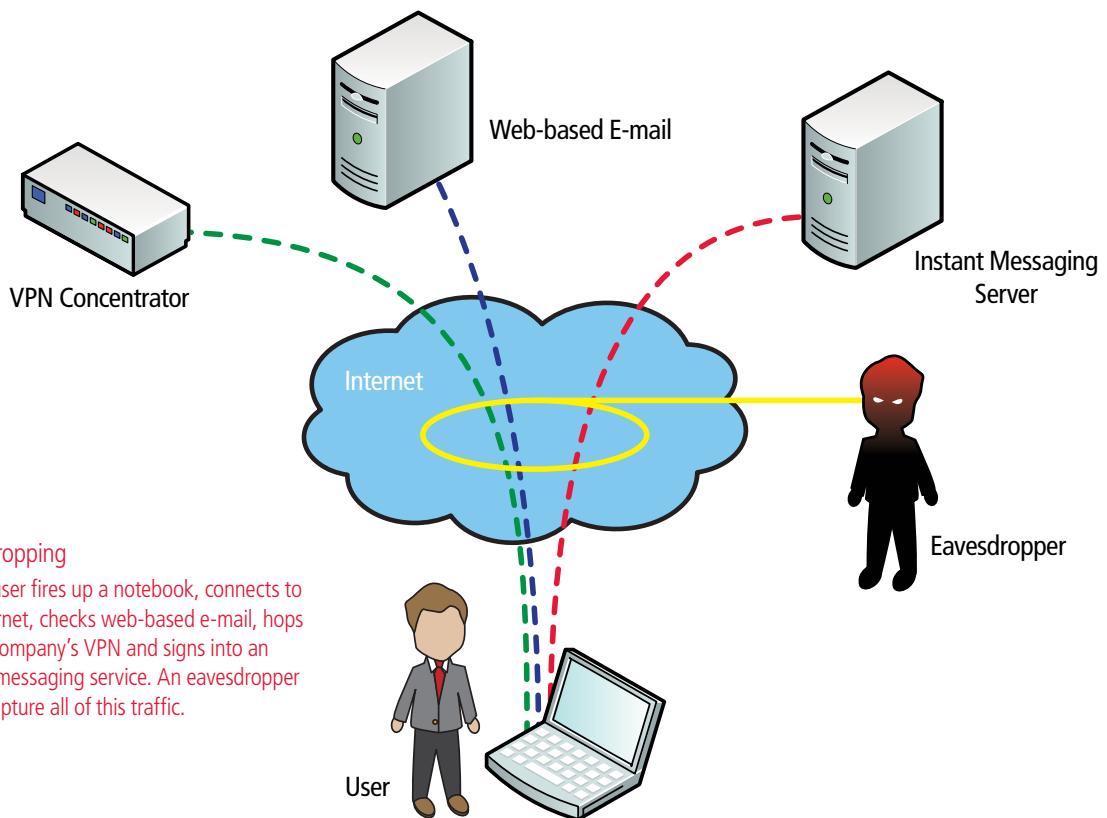
Remote network endpoints have always dealt with the problem of eavesdropping: VPNs and cryptographically protected application protocols such as SSH or Hypertext Transfer Protocol over Secure Sockets Layer (HTTPS) make it possible to conduct private transactions on public networks without major information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

However, these measures also create a false sense of security because users often aren't aware of all the avenues for disclosure of their passwords. In the figure on page 32, a user fires up a notebook, connects to the Internet, checks web-based e-mail, hops on the company's VPN and signs into an instant-messaging service. An eavesdropper can capture all of this traffic.

The VPN logon is secure, but the other two may not be. Moreover, the instant-messaging logon might happen automatically, without the user even knowing about it.

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

Eavesdropping isn't required for password disclosure. If the user



uses the same password for enterprise e-mail and personal e-mail, online shopping accounts, financial services accounts, iTunes and airline frequent-flyer accounts, then many organizations have possession of the password. While Amazon.com may be trusted not to lose all of those passwords, the same trust shouldn't be given to CheapDealsOnSocks.com.

The solution to the problem of password theft and password reuse is to ensure that stolen passwords won't be useful to an eavesdropper in any way. Use of digital certificates and smartcards is one approach, but if the overhead of a public key infrastructure (PKI) deployment seems out of reach, then token-based systems, such as one-time passwords or time-based passwords, are a reasonable alternative.

An even simpler approach is to educate users to protect their passwords from disclosure. For example, they should never reuse their enterprise password for anything else.

Protecting passwords isn't the only eavesdropping issue. Data in transit across enterprise networks may need to be encrypted as well. While one typically considers the internal network in an

organization to be safe, the possible presence of both hostile software (such as an infected PC) and hostile people on a network means that eavesdropping should be a concern.

Unfortunately, there is little agreement on how to resolve this issue. While Microsoft Windows has supported automatic system-to-system communications encryption for many years, few enterprises (other than Microsoft) have rolled out this valuable feature.

Third parties have also come up with on-the-LAN encryption solutions that go beyond the Microsoft boundary. And the Institute of Electrical and Electronics Engineers (IEEE) is working on a standard to push encryption down to the Ethernet layer.

One of the concerns with such widespread encryption is that it can act counter to good security policy, as some network eavesdroppers might be IPSs, malware-detecting firewalls or network auditing tools. All of these tools become ineffective in the face of widespread on-network encryption. ♦

GLOSSARY



This glossary serves as a quick reference to some of the most essential terms noted in this reference guide.

ACCESS CONTROL LIST (ACL)

An ACL is a data set that dictates user permissions and access to particular objects within a network. The list consists of users and what actions each user is permitted to perform on a network object.

ADVANCED ENCRYPTION STANDARD (AES)

AES, also known as the Rijndael algorithm, was first published in 1998 by the Belgian cryptographers Joan Daemen and Vincent Rijmen. The U.S. government adopted AES for all encryption purposes in 2002. AES can be used with key sizes of 128, 192 and 256 bits. AES users should select 256-bit key size if it is supported in their equipment.

ANNUAL LOSS EXPECTANCY (ALE)

Used in risk assessment, it defines the annual loss expectancy for a particular asset. This is calculated by multiplying the single loss expectancy (SLE) by the annual rate of occurrence (ARO). For example, an asset with an SLE of \$6,000 that is affected once every four years has an ALE of $\$6,000 \times 1/4 = \$1,500$.

ANNUAL RATE OF OCCURRENCE (ARO)

Used in risk assessment, it defines how often a negative event occurs per year.

BOT

Short for "robot," a bot in the context of network security is a compromised computer that is under the control of a third party (sometimes called a "bot herder"). Typically, bots are used to launch distributed denial-of-service (DDoS) attacks, send spam or host websites used as part of phishing scams. Collections of bots, called "bot armies," can number in the hundreds or hundreds of thousands.

DENIAL OF SERVICE (DoS)

A DoS attack prevents legitimate users from accessing system resources. A common method consists of saturating a target (server) with requests so 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.

EXPOSURE FACTOR (EF)

Used in risk assessment, it defines how much of an asset will be impaired when a negative event occurs. For example, if a virus is expected to infect 25 percent of your systems, your EF would be 25 percent.

GRAMM-LEACH-BLILEY ACT (GLBA)

GLBA is a federal law enacted in 1999 that sets out rules and provisions for financial institutions that protect a citizen's financial records and information. These provisions include doing security assessments; developing and implementing security solutions that detect, prevent and allow timely incident response; and performing auditing and monitoring of the institution's security environment.

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.

INTERNET PROTOCOL SECURITY (IPSec)

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

LOW AND SLOW

A term used in IDS and IPS parlance, a "low and slow" attack is one that operates below the alerting level of most break-in detection systems. For example, if an attacker were to try to brute-force guess a password using low and slow techniques, they might try only a single password each day. Low and slow attacks are particularly difficult

to detect and mitigate because they operate below the threshold of most attack evasion systems.

MAN-IN-THE-MIDDLE ATTACK

This term refers to an attack where the adversary impersonates a remote-access system itself. Remote users are led to believe that an imposter system is, in fact, the legitimate source of remote-access connectivity. The attacker can inspect and alter all traffic between the remote user and the organization's network.

NETWORK ACCESS CONTROL (NAC)

Also known as Network Admission Control, the term is used to describe a user-focused, network-based access control. With NAC, any user attempting to connect to the network is authenticated and has the security posture of their workstation checked for compliance with corporate standards. Based on the authentication and the endpoint posture, the network itself will enforce access controls as defined by the NAC manager.

ONE-TIME PASSWORD (OTP)

A password that can be used only once. Typically, systems that are based on OTPs generate lists of passwords for each user. As each password is used, it is crossed off the list and the next password is the one that is expected. Variations using hardware tokens are also common.

PAYMENT CARD INDUSTRY DATA SECURITY STANDARD (PCI DSS)

A set of security standards created to guide credit card processing companies in defending against fraud, hacking and other security threats. Processing companies must be PCI DSS compliant if they are processing, storing or transmitting credit card payments.

PHISHING

A form of online scam that attempts to trick people into disclosing private information, such as credit card numbers. Well-known brands are often used to lure subjects to spoofed websites — or even hijacked domains — that look legitimate.

REGULATORY RISK

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

SARBANES-OXLEY ACT (SOX)

SOX is a federal law passed in 2002 that defines legislative audit requirements for corporations' financial reporting to improve the accuracy and reliability of corporate disclosures. From an IT perspective, SOX pushes corporations to archive and store any finance-related document or e-mail.

SECURE SOCKETS LAYER (SSL)

SSL is a cryptographic protocol for communications over TCP/IP networks. This protocol encrypts segments of the transport layer protocols for an end-to-end connection across the network. SSL has been superceded by Transport Layer Security (TLS).

SHOULDER SURFING

Password theft from looking over someone's shoulder (or simply at their keyboard) while they type in their password. Shoulder surfing is a popular way to collect passwords, and is easily mitigated through the use of nonreusable passwords.

SINGLE LOSS EXPECTANCY (SLE)

Used in risk assessment, it represents the expected loss to the organization for a single negative event, such as a virus infection or a break-in. The SLE is roughly calculated by multiplying the value of the asset affected by the exposure factor (EF), and factoring in the total downtime. For example, if an asset generates \$28,000 in value each day, has an EF of 50 percent, and is down for two days, then the SLE is $\$28,000 \times 50\% \times 2 = \$28,000$.

TIME-BASED PASSWORD TOKEN

The most popular of the token-based password schemes, time-based password tokens generate passwords that are a direct function of the time. Time-based password tokens typically show the password in a display, changing it at specified time intervals (usually every minute). The actual password entered is normally a concatenation of the displayed password and a second fixed password (or PIN) known to the token owner.

TRAFFIC ANALYSIS

Traffic analysis is a security threat when carried out by outside parties. Even if the contents of the traffic remain encrypted, eavesdroppers can make inferences simply because the traffic is going on, and potentially put a firm at risk.

UTM FIREWALL

A term originally coined by IDC's Charles Kolodgy, UTM has become the dominant firewall type for the small- and medium-size business (SMB) network. At its core, UTM brings together three main ideas: multiple security features, integrated on the basis of a mature firewall, deployed in an appliance form factor. Most UTM firewalls include firewall, VPN, IPS and antivirus features at a minimum, with URL and content filtering, antispam and application-aware firewalls as common additions.

WHOLE-DISK ENCRYPTION

Whole-disk encryption can happen through either a hardware or software solution. This approach encrypts every bit of data on a disk. The term "whole" refers to the fact that this is an all-or-nothing approach to encryption. Users cannot pick and choose the files that they wish to encrypt; everything on the disk is secured.

INDEX



Advanced Encryption Standard (AES).....	27, 31
Annual Loss Expectancy (ALE).....	6-7
Annual Rate of Occurrence (ARO).....	6-7
Bot	9, 26, 31
Business enabler	3-4
CIA security triad.....	3
Compliance	5, 8, 12
Data Loss Prevention (DLP)	13-15
Denial of Service (DoS).....	26
Dictionary attack.....	26
Eavesdropping	31-32
Encryption.....	30-32
Endpoint	10, 26-27, 29-31
Exposure Factor (EF)	6-7
Gramm-Leach-Bliley Act (GLBA)	8
Health Insurance Portability and Accountability Act (HIPAA).....	8
Identity management.....	9-10, 15
Internet Protocol Security (IPSec)	25
Intrusion Detection System (IDS).....	11-12, 26
Intrusion Prevention System (IPS).....	6, 11-12, 32
Managed Security Service Providers (MSSPs).....	12
Man-in-the-Middle (MITM) attack.....	26-27
Network Access Control (NAC)	10, 31
One-time Password (OTP).....	27
Payment Card Industry Data Security Standard (PCI DSS).....	8
Perimeter controls	9
Risk assessment	5-8
Sarbanes-Oxley Act (SOX).....	8
Shoulder surfing.....	27
Single Loss Expectancy (SLE).....	6-7
Social networking (Web 2.0).....	10
SQL injection	11
Time-based password token	26
Traffic analysis.....	26
Virtual environments.....	12
Virtual Private Network (VPN)	9, 31-32

Disclaimer

The terms and conditions of product sales are limited to those contained on CDW's website at CDW.com. Notice of objection to and rejection of any additional or different terms in any form delivered by customer is hereby given. For all products, services and offers, CDW® reserves the right to make adjustments due to changing market conditions, product/service discontinuation, manufacturer price changes, errors in advertisements and other extenuating circumstances. CDW®, CDW-G® and The Right Technology. Right Away® are registered trademarks of CDW LLC. All other trademarks and registered trademarks are the sole property of their respective owners. CDW and the Circle of Service logo are registered trademarks of CDW LLC. LLC. Intel Trademark Acknowledgement: Celeron, Celeron Inside, Centrino, Centrino Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel Viiv, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, Viiv Inside, vPro Inside, Xeon and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries. Intel's processor ratings are not a measure of system performance. For more information please see www.intel.com/go/rating. AMD Trademark Acknowledgement: AMD, the AMD Arrow, AMD Opteron, AMD Phenom, AMD Athlon, AMD Turion, AMD Sempron, AMD Geode, Cool 'n' Quiet and PowerNow! and combinations thereof are trademarks of Advanced Micro Devices, Inc. This document may not be reproduced or distributed for any reason. Federal law provides for severe and criminal penalties for the unauthorized reproduction and distribution of copyrighted materials. Criminal copyright infringement is investigated by the Federal Bureau of Investigation (FBI) and may constitute a felony with a maximum penalty of up to five (5) years in prison and/or a \$250,000 fine. Title 17 U.S.C. Sections 501 and 506. This reference guide is designed to provide readers with information regarding network security. CDW makes no warranty as to the accuracy or completeness of the information contained in this reference guide nor specific application by readers in making decisions regarding network security implementation. Furthermore, CDW assumes no liability for compensatory, consequential or other damages arising out of or related to the use of this publication. The content contained in this publication represents the views of the authors and not necessarily those of the publisher. ©2010 CDW LLC. All rights reserved.

CDW

One CDW Way
200 N. Milwaukee Avenue
Vernon Hills, IL 60061



Your account number is

Key Code

PRESORTED STANDARD
U.S. POSTAGE
PAID
PERMIT NO. 1
CAROL STREAM, IL 60188

100421
Flyer 75758



CDW.com/securityguide
888.509.4239

ABOUT THE CONTRIBUTORS

“**BUSINESSES
HAVE CHANGED
COURSE.**

Instead of devoting their efforts to protecting computers, the focus has shifted to protecting the critical data on them.



PEYTON ENGEL
CDW Security Expert



PEYTON ENGEL leads a team of security engineers at CDW. With the CDW (formerly Berbee) team since 1998, he has been responsible for its growth and management, including sales and marketing, since 2001. Peyton has presented his security research at national conferences including DEFCON (2004, 2006), ToorCon (2002, 2005) and USENIX/LISA (invited speaker: 2003, 2005). Peyton's chief technical interests are software security and the security relationships between systems in large networked environments. He works in Madison, Wis.



JOEL SNYDER is a senior partner with an IT consultancy. With 30 years of practice, Dr. Snyder is an internationally recognized expert in the areas of security, messaging and networks. A popular speaker and author, he is known for his unbiased and comprehensive tests of security and networking products. His clients include major organizations on six continents, from a home base in Tucson, Ariz.

SECURITY REFERENCE GUIDE

LOOK INSIDE for more information on:

- Using security as a business enabler
- Ways to assess risk and compliance
- Improving network visibility and internal controls
- Data loss prevention
- Remote access security
- Security endpoints and portable media