Realtime publishers

The Shortcut Guide™ To

Prioritizing Security Spending

sponsored by

IBM®

Dan Sullivan

Copyright Statement

Chapter 4: Best Practices for Prioritizing Security Spending

Businesses and other organizations face difficult challenges when protecting their information assets and infrastructure. As we have discussed throughout this guide, the combination of security threats, constantly evolving IT technologies, and limited resources require us to prioritize security spending. There are so many factors that can influence our choices, it helps to first clearly articulate our goal and then develop a methodology for meeting the goal.

We start with assuming our goal is to implement security practices that align with business strategy. Security is not an end in itself and it is not the driving goal. That assertion may sound a little strange especially in a book about security, but it actually addresses the heart of the problem in prioritizing security spending.

We do not implement security controls for the sake of having controls. If we did, we would disconnect our computers from the Internet, severely limit physical contact with servers, and implement access controls so stringent that even if potential users somehow gained access to a system they could not do any substantive work with applications or data. Once again, security is a means to an end, not an end in itself. Our real goal is to support the business strategy of the organization. Formulating that strategy is outside the scope of security practices, although, as we have seen throughout this guide, security concerns can influence that strategy and constrain options for implementing it.

There are so many options to choose from when considering how to best allocate security spending that a methodology for organizing and categorizing those decisions can be a significant aid to the process. The methodology advocated in this book is depicted in Figure 4.1. The methodology is similar to other IT life cycle-based methodologies that start with understanding business requirements and follow through multiple stages to implementation and maintenance.

The methodology also includes an education component, which is not typical of other life cycle-based methodologies. This education is unlike the training provided to administrators or end users of a new application, which tends to focus on particular functionality of a system. The purpose of security education is to develop an awareness of security threats and how to detect and avoid them. It is important that all stakeholders understand fundamentals of security threats because the specifics are constantly changing. Email viruses were more of a problem 10 years ago than today but malware has not gone away; it instead uses additional methods to propagate.

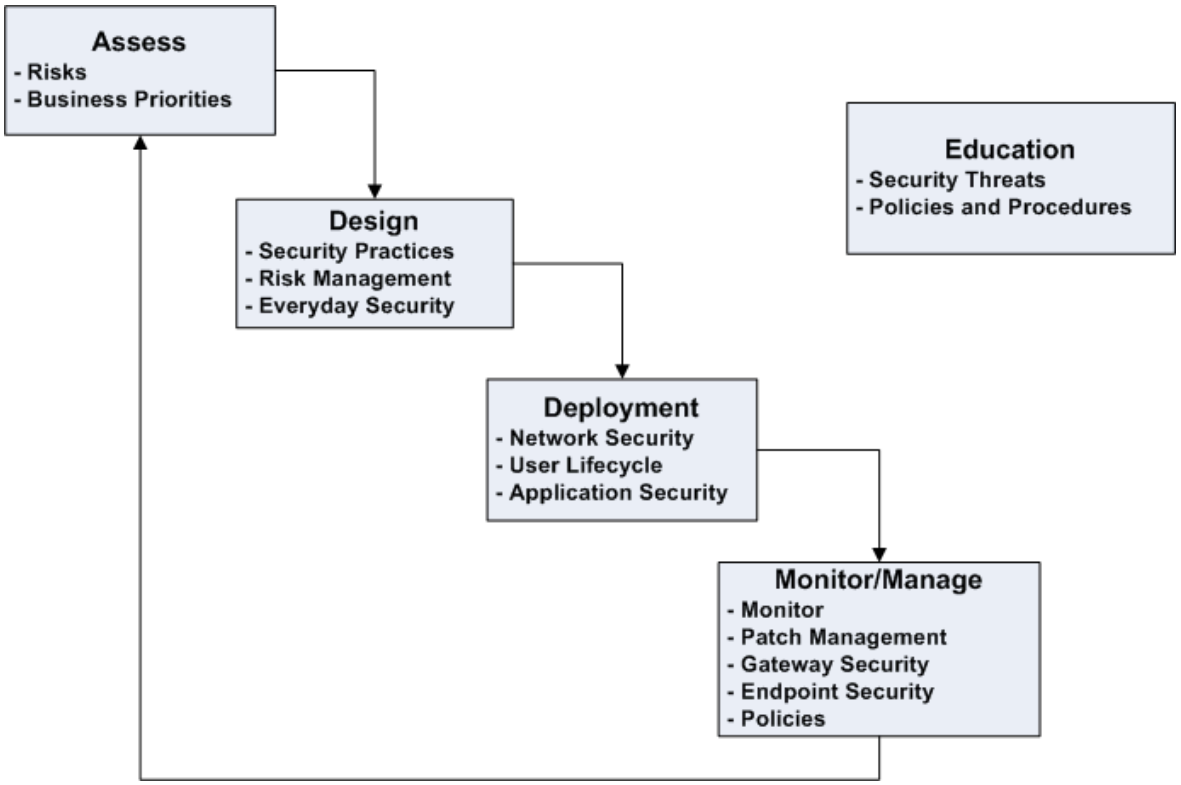

Figure 4.1: A methodology such as the model shown here helps with prioritizing security spending. The process begins with assessing risks and business priorities followed by steps to design, deploy, and manage security controls. Note this process is cyclical and knowledge gained in the process can help guide later assessments. Education is a key element of the methodology and spans the other stages.

Assessing Business Environment

Prioritizing security spending begins with understanding the environment in which business operates. This entails understanding both external and internal threats as well as how those threats could impact your business priorities. Lessons from the 6th century B.C. Chinese military strategist Sun Tzu serve as apt advice for today's security professionals:

> So it is said that if you know your enemies and know yourself, you will fight without danger in battles. If you only know yourself, but not your opponent, you may win or may lose. If you know neither yourself nor your enemy, you will always endanger yourself. (Source: Sun Tzu, *The Art of War*, http://en.wikipedia.org/wiki/The_Art_of_War).

Know Your Enemy: Identifying Security Risks

Let's categorize security risks according to the three broad objectives of information security: protecting the confidentiality, integrity, and availability of information and systems and networks.

Threats to Confidentiality

Threats to confidentiality come from a wide array of persons and technical means:

- Data breaches by outside attackers
- Unauthorized access or use of information by employees or other insiders
- Loss of storage media, such as backup tapes
- Social engineering attacks, such as phishing

The risks to a business if the confidentiality of information were compromised include loss of customer confidence, brand damage, and possible fines or other consequences of not complying with government and industry regulations.

Threats to Integrity

Threats to integrity are similar to threats to confidentiality. In this case, however, we are concerned about protecting the accuracy and reliability of data. For example, a disgruntled employee might tamper with database records in retaliation for a poor performance review or other perceived slight, or an authorized user might make an unintentional mistake that damages the data. If adequate monitoring and auditing procedures are not in place, the unauthorized changes might not be detected. Similarly, an outside attacker could sabotage a business by writing bogus data to database records. Encryption techniques can also be used to determine whether data has been modified from its original values.

The risks associated with loss of integrity can be costly. For example, if a database were tampered with, how could you (1) identify which parts of the database had been tampered with and which data elements were still intact and (2) recover the accurate data? Backups might help, but when was the last known good backup made? How will legitimate changes to the database since the last known good backup be handled? Is there a sufficient paper trail to reconstruct the lost information? In addition to disrupting day-to-day operations and taking employees from their regular jobs to triage and clean up the compromised data, there may be additional consequences when auditors become involved.

Threats to Availability

Availability is the third category of risks. Availability is threatened by natural disaster as well as human-made incidents. Backup and recovery operations and disaster recovery planning can mitigate the risks of natural disasters, accidents, and human error (for example, accidently deleting a file). Intentional attacks on availability, such as botnets, spam, and malware infections, require additional responses; fortunately, many of the same security controls one would put in place to protect the confidentiality and integrity of information help to protect availability as well.

Availability risks to business can hit directly at day-to-day operations. How quickly a business can restore services, either in the main office environment or in a disaster recovery facility, will directly impact the extent of the damage to the business.

Cross Reference

Chapters 2 and 3 of this shortcut guide provide more detail on specific risks and how emerging and recently adapted technologies, such as cloud computing and virtualization, introduce variations on well-known risks.

Understanding the types of security risks to a business is only part of the assessment process; one must be able to map those risks to business priorities in order to appropriately prioritize your own security spending.

Mapping Risks to Business Priorities

Business priorities often entail some amount of IT involvement. This could range from the relatively simple, such as moving employees from a remote office to headquarters, to strategic initiatives, such as reducing support costs by launching a customer self-service Web site. One way to map risks to business priorities is to start with a business goal and outline IT's responsibility for meeting that goal. Next, for each responsibility, determine what risks affect those responsibilities, especially if the tasks are not performed correctly. For example, if a goal of the business is to optimize the allocation of staff in the business, some employees may be moved to different departments. Unless all of those employees' access controls are updated, some may retain access to information that their new positions do not warrant. (See Table 4.1 for additional examples).

Strategic Business Goal	IT Responsibilities	Security Risks
Consolidate Remote Offices	Move or decommission servers and workstations	Data leaks due to data left on decommissioned hardware
		Lost or stolen hardware
	Migrate applications and files to existing servers	Improperly configured applications
		Improperly patched OSs for new applications
		Incorrect access controls on migrated files
	Update authentication and authorization information	Old authorizations left in place
		Reassigned employees with incorrect authorizations
		Weak passwords or default password used during migration
Reassign employees to minimize labor costs	Move employees to different departments	Failure to remove authorizations for previously-used applications
		Employees leave sensitive material on flash drives or other ad hoc backup devices
		New employees not sufficiently educated on security policies and procedures

Strategic Business Goal	IT Responsibilities	Security Risks
	De-provision terminated employees	User-specific accounts left accessible on applications
		Passwords not changed on shared access accounts
		Employees retaining copies of sensitive or confidential information
		Terminated employee tampering with data
Launch new customer self-service Web application	Protect back-end systems	SQL, HTML, and other forms of injection attacks through Web application interface
		Unencrypted data in databases
		OSs on application servers and database servers are not hardened
		Improperly configured gateway
	Ensure confidentiality during data transmission	Insufficient use of encryption and SSL
	Ensure availability of service	DoS attack
		Insufficient available network resources due to spam, botnet, or other unwanted network activity

Table 4.1: The risks to business objectives can be discovered, in part, by determining IT responsibilities for realizing those objectives and identifying the risks if those responsibilities are not properly met.

As we can see from Table 4.1, there is no shortage of security risks associated even with routine business operations. How should one prioritize among all the different ways something could go wrong? We want to reduce the risk of anything that could adversely impact the business, so we focus on:

1. Identifying the business goals that have the potential for greatest gains or cost reductions. For example, launching a customer self service application may save more over time than an office consolidation; thus, the former should have higher priority than the latter.

2. Within the prioritized set of business goals, what is the potential adverse impact of each risk and what is the likelihood that the risk will be realized? For example, a terminated employee with administrator access to a customer database is a high impact and possibly high-likelihood risk. An employee moved from one department to another who retains read/write access to a shared drive for her former department probably falls toward the low impact and low likelihood end of the spectrum.

By ordering on a combination of business priorities, risks to those priorities, and the likely impact if a particular set of risks were realized, we can distinguish issues most demanding of resources from those that present less of a risk to the organization. As we delve into more details, we move from the assessment stage to the design phase of the methodology.

Designing for Security

Designing for security extends the alignment of security practices with business operations to develop a more detailed analysis of risks, costs, and likelihoods of adverse events. We'll divide our discussion of the design stage into two parts: conventional risk management practices and everyday security practices.

Risk Management Practices

Risk management practices provide the means to determine the potential cost of various risks to an organization. In its most basic form, risk analysis is a methodical analysis process made up of several steps, including:

- Assigning values to information assets
- Identifying threats to those assets
- Estimating the value of a loss to each type of asset from each type of threat
- Analyzing threats and estimating their likelihood
- Calculating the potential loss from each threat
- Making a decision on how to respond to the risk, which can include reducing the risk, transferring the risk through insurance, or accepting the risk

At the end of the process, we should have a rational, cost-justified response to the set of threats a business faces. Although this is true in theory, in practice there are a number of difficulties with this model. Nonetheless, it is a solid methodology for at least starting a risk management assessment.

Caution: Do Not Be Lulled into a False Sense of Certainty by the Formality of Risk Analysis

Risk analysis is a valuable and important tool, but there is an easily overlooked aspect of risk analysis that should be called out. Risk analysis entails a logical sequence of analytic steps that can sometimes mask its limitations. Some of those steps are shrouded in technical jargon such as "Single Loss Expectancy" and "Annualized Rate of Occurrence," which are further elevated by their use in formulas. Do not let the formalisms fool you. The outputs are only as good as the inputs, and good inputs to rare events (for example, what is the probability a database administrator will steal your customer credit card data?) are hard to come by. Building spreadsheets with a lot of calculations can be a useful exercise, in part due to the calculations it renders but also for highlighting just how soft some of the our working assumptions are.

It is important to remember that prioritizing security spending and analyzing security risks is not as easily quantified and calculated as it might first appear when presented with risk analysis equations.

To assign a value, we have to consider a number of aspects of an asset. For example, what is the replacement value of the asset should the business have to replace it? Is that value different from the value a competitor may place on the asset? The two valuations are not always the same. For example, if a company lost the design for a product that is in its third generation, the engineers within the company experienced with previous designs could redesign the product much faster (and therefore in a less costly manner) than a competitor starting more or less from scratch. Other valuations could be based on how much the asset cost to develop or how much profit it generates for the company. There is no single answer on how to value an asset and for complex cases it may be safer to consider several valuation methods before setting a value.

We have considered some of the types of threats to assets, but to reiterate, here are some of the most significant threats: data theft; data destruction; data tampering; loss of infrastructure due to natural disaster; disruption of operations due to malware, DoS attack, sabotage by disgruntled insider, and so on; and privacy breach. Within the risk management framework, we methodically consider how each relevant threat can affect each asset. Obviously, some threats, such as data theft, do not directly impact all assets, such as network infrastructure, even though those assets may be used in an attack. The substantive cost of data theft is related to the stolen data, not the network capacity consumed by the thieves as they copied files across the network.

The quality of a risk analysis is highly dependent on accurate estimates of the likelihood of events. How often will a business experience a particular threat? It is fairly safe to say that exposure to viruses, worms, Trojans, and other forms of malware would be a frequent problem without basis security measures. Similarly, spam would litter the inboxes of employees unless email filtering is in place. Guessing how often a systems administrator will steal confidential or sensitive data is a more difficult task; however, given the potential cost of such a breach, even infrequent losses such as that would warrant low-cost, high-impact security measures such as appropriate access controls, separation of duties, and rotation of duties.

Risk Analysis Calculations

As noted earlier, risk analysis is a formal process and as with most formal systems, there are equations involved. Here are some basic terms and formulas for getting started with risk analysis:

- Exposure factor—The percentage of an asset's value that is lost when a threat is realized. For example, if a botnet is consuming 10% of network bandwidth, the exposure factor is 10% from a botnet. A natural disaster that completely knocks out network services has an exposure factor of 100%
- Single Loss Expectancy (SLE)—The estimated value of a loss due to a single incident, i.e., the asset value multiplied by the exposure factor
- Annualized Rate of Occurrence (ARO)—The estimated number of times an event will occur within a single year. (Fractions are used when the rate of occurrence is less than once per year). Some AROs will be high (for example, the number of times a workstation will be exposed to malware), and some will be low (for example, the number of times a data center will burn to the ground).
- Annualized Loss Expectancy (ALE)—How much each threat will cost you over the course of a year; formally, the SLE multiplied by the ARO.

At the end of a risk management analysis, we have a set of risks and costs associated with those risks. Now what? Ideally, we can mitigate many of these risks for reasonable costs. Potentially expensive risks do not necessarily require expensive countermeasures. Here are some examples:

- Risk: Total loss of data because a disgruntled systems or database administrator could wipe out the database. Mitigation: Separate duties so that different employees are responsible for creating backups and performing other day-to-day operations on the database.
- Risk: Keyboard loggers could be used to capture usernames and passwords on critical servers. Mitigations: Install an anti-malware solution, change passwords frequently, and use two-factor authentication.

- Risk: Fire destroys the data center. Mitigations: Use a cloud computing provider for disaster recovery, storing backups in the cloud and using cloud computers as a disaster recovery site; purchase insurance for the cost of destroyed physical assets.

In general, one has the option of mitigating a risk by implementing security controls, transferring the risk by purchasing insurance, or accepting the risk and living with the consequences. The best choice for your situation can be determined by taking into account the particulars of your business and your risk tolerance. It is worth noting, though, that security measures can address multiple risks; it is not simply a matter of every risk requiring a separate security measure. For example, a good spam filter may eliminate virus-carrying email messages even if it did not scan for viruses. Applying OS patches may eliminate a security vulnerability as well as improve the reliability of an application, thus lowering support costs.

Fortunately, not all security measures are expensive or time consuming. Sometimes, day-to-day practices can contribute to an improved security posture as well.

Everyday Security

The risk management practices just described tackle security challenges from the top down; everyday security takes a more bottom-up approach to the problem. Rather than ask "What are all the risks to an organization?" we start with the question "What steps can be readily taken to improve overall security?" The goal in this case is to leverage security practices and applications for the greatest benefit. Risk management and everyday security practices are complementary. We do not have to choose one or the other; in fact, we want to practice both.

Everyday security is an emerging set of practices that focus on quickly performed tasks that help mitigate risks. These include:

- Running anti-malware scans and making sure anti-malware software is up to date
- Reviewing email and content filtering logs to watch for spikes in spam or malicious content trying to enter the network
- Checking changes to access controls to determine whether any accounts have had privileges elevated or unauthorized accounts have been created
- Reviewing the status of patch logs to determine whether there were any failures in automatically applied patches to OSs or applications
- Scanning network logs for unusual events, such as a large number of failed attempts to access the network over a virtual private network (VPN)

As we can see from this list, everyday security depends on having systems and practices in place to monitor both scheduled events, such as patching processes, and for anomalous events, such as unauthorized changes to access controls.

Risk analysis helps us understand the big picture of information security and prioritize our security spending. Everyday security leverages the tools and procedures that risk management puts in place to go after additional "low hanging fruit" at minimal additional cost.

Before moving on from design considerations, we should mention that both risk analysis and everyday security practices should incorporate the goal of defense in depth. No security measure is perfect. Complex systems, including security applications, often harbor vulnerabilities, errors or misconfigurations. One way to deal with this is to have multiple measures protecting against a risk. In this way, if one measure fails, the others may still provide protection. For example, a network appliance may be in place to scan for malware but running antivirus locally on individual workstations provides an additional layer of protection should the network filter fail.

By considering the broad range of risks to the business, we can improve the chances of understanding the comprehensive range of threats that can adversely affect a business. By considering the everyday security practices available to us, we may find relatively easy-to-implement, high-impact security procedures. To make sure we have the tools and procedures in place for a sustainable security strategy, we next turn our attention to deployment issues.

Deployment of Security Measures

The goal of the deployment phase is to roll out security practices and solutions in such a way as to meet the security objectives defined in earlier stages and to do it in a cost-effective manner. Indeed, these security solutions and practices can help control costs in the long run. To see this, we will consider three different sets of deployment issues:

- Network security life cycle
- User life cycle management
- Web application security

These three areas are chosen as representative of a broader array of deployment areas and, although the list is not comprehensive, they do provide numerous examples of the types of considerations we face during deployment.

Network Security Life Cycle

In the realm of network security, we have two general areas that we tend to focus on: protecting the organization from the "threat of the day" and protecting the organization against threats that create a continuous risk to the organization. We will consider each in turn.

Episodic Threats

Episodic threats, or "threats of the day," are rapidly emerging threats that may require a specific, targeted response. For example, a vulnerability may be found in a popular email client that is exploited by a rapidly spreading worm. The worm contains multiple payloads, including keyloggers for capturing usernames and passwords. Even with the security procedures and applications in place, the malware is able to avoid detection and threatens to spread throughout the corporate network. An immediate response is required. By definition, we cannot plan for an unanticipated event, such as the particular worm just described, but we can plan for the need to respond rapidly to a new threat.

Preparing for episodic threats is more about the ability to gather information and change operational states rapidly than it is about planning for all possible worst-case scenarios. Some of the top-priority considerations for dealing with episodic events include having

- Ready access to device configurations and patch levels to help identify specific devices that may be vulnerable
- The ability to quickly isolate an infected device or subnet to reduce the chance of spreading malicious software
- Consolidated reporting systems that provide access to multiple logs and monitoring detail, which help verify anomalous events using evidence from multiple systems
- An incident response plan in place so that well-defined procedures are followed to address a threat rather than the formulation of an ad hoc response to each new incident

Fortunately, there is significant overlap between the security procedures and applications one would put in place for dealing with episodic threats and the ones we want in place for continuous, long-term threats.

Continuous Long-Term Threats

There is no shortage of continuous long-term threats to information security. The risk analysis process outlined earlier can be used to identify and prioritize risks, while the everyday security practices can contribute to mitigating those threats and limiting the need for additional time-consuming procedures. When considering how to prioritize the need to address these long-term threats, we need to consider business strategy, risk tolerance, and the types of risks we face. Some that should be considered include:

- Data loss due to a variety of factors, such as vulnerabilities in applications and networks as well as weaknesses in business operations and protections against employees and other insiders with malicious intent
- Unavailability of services due to natural disaster, technical failure, or malicious activity against a business
- Loss of computing, storage, and network resources due to botnets, spyware, and other unwanted programs

- Employees using business systems for inappropriate and non-business uses, especially activities that could create a hostile workplace environment
- Inability to demonstrate compliance with government and industry regulations, which could result in fines or other penalties

Priority in security spending should be given to areas such as these that are most relevant to a particular business' needs. We should note, however, long-term threats are not limited to technical threats, such as malware, and cannot be addressed only by spending money on security technology. Insufficient governance and management is also a potential threat. One area where poor oversight can have significant consequences is the realm of user life cycle management. Another non-technical area that deserves attention when setting priorities is education, which will be addressed later.

User Life Cycle Management

User life cycle management is one area of technology where a combination of technical controls and sound practices can significantly reduce vulnerabilities related to data loss, data tampering, violation of privacy regulations, and unauthorized use of business computing resources. When prioritizing security operations, consider (at least) three aspects of user life cycle management:

- User provisioning
- Roles and privileges
- De-provisioning

User Provisioning

User provisioning is the process by which employees, consultants, business partners, and others are provided access to business systems. Optimizing the user provision process can improve security, but there are plenty of other business drivers that make this a worthy area of investment. For example, a new employee may need access to multiple systems including email, shared network drives, operational systems, human resources and other administrative applications, and network access control systems. Streamlining the provisioning process can reduce the cost of provisioning by reducing demands on support staff as well as providing employees with the tools they need in a timely manner. From a security perspective, an automated user provisioning process can reduce the chance of error over manual process where support staff manually creates individual accounts across multiple systems. Automated user provisioning can also help with managing roles and privileges.

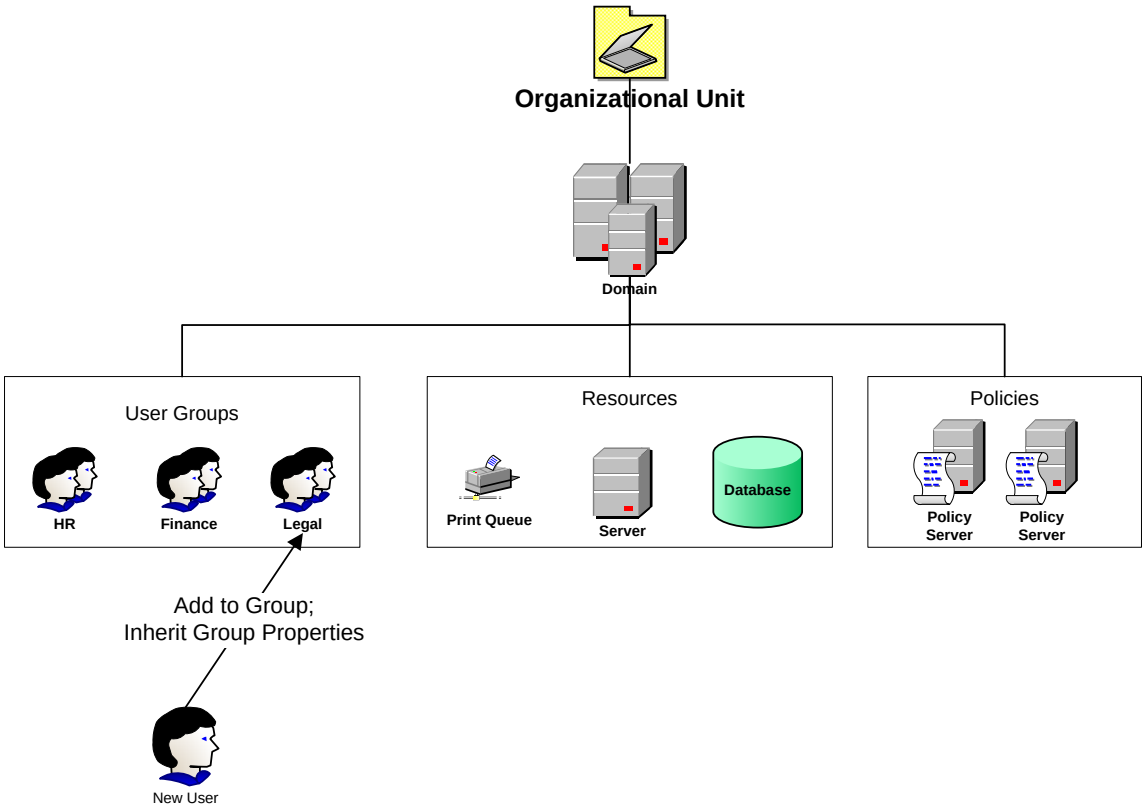

Figure 4.2: User provisioning with automated systems can help ensure proper access controls and policies governing a new user are correct and complete by adding users to existing management structures governed by pre-defined policies.

Roles and Privileges

A top priority in user life cycle management is controlling roles and privileges assigned to users. In a worst case scenario, systems administrators assign privileges to individual users granting them the ability to work with particular resources ranging from printers and shared drives to enterprise applications and network access control systems. An only slightly better version of this scenario is that roles are established but are so narrowly tailored that new groups are constantly added to accommodate new sets of users. Using too few roles risks having to assign privileges for some members of a group that really are not warranted for other members. The Goldilocks' Principle applies here: not too much, not too little when it comes to defining roles.

Once established, there should be policies that require periodic recertification to ensure that users are still entitled to their existing roles and privileges—a useful audit measure that is often required by compliance regulations. Over time, the number of roles and privileges that a user could accumulate might create a separation of duties exposure that should be considered when assigning new roles or when recertifying user capabilities.

When prioritizing for security spending, do not overlook the value of a good system supporting user provisioning and role management. The ability to clearly define access control policies and assign those to logically organized groups can help insure consistent application of the business' access control policy. It can simplify user management, reduce the time for creating new user accounts, and significantly reduce the cost of user administration during the user life cycle. It also helps during the deprovisioning process.

Deprovisioning

Businesses and the workforce are dynamic. Most sizeable businesses can reasonably expect a constant stream of new hires, terminations, and movement within an organization. Just as it is important to get new users access to the systems they need, it is important to ensure that such access is terminated when it is no longer required. Deprovisioning, the process of removing access to systems, is also supported by automated user management systems. This supports security objectives by helping ensure that only those with a well-defined business requirement for a system will have that access. In combination with strong password policies, strong authentication, and account monitoring, sound user life cycle management practices can mitigate risks stemming from insufficient or poorly managed access controls.

Federated Identity Management

User access to applications can extend beyond corporate boundaries and into other businesses. Once again, the need to align business strategy with security comes into focus. For example, a product design firm may contract the manufacture of a device to a manufacturing partner. Employees at the manufacturing firm need access to design documents and other resources from the designer. Should the designer start creating accounts on its network for all the manufacturer's employees? Of course not, and they should not have to create accounts for the few engineers at the manufacturing company with legitimate need to access the designer's systems. An alternate approach is to establish a federated identity management system.

With a federated identity management system, participating organizations trust their partners to verify identity of users and assign them to roles appropriate for their business needs. Partners can then use those role-based attributes to determine what kinds of access extra-organization users can have to resources. Standards, such as the Security Assertion Markup Language (SAML), enable non-vendor specific solutions for federated identity management in support of Single Sign-On (SSO) and role management.

Least Privilege and Rotation of Duties

In addition to user life cycle management practices that apply to all users, there are two security principles that should be used when managing users with elevated access, such as system managers, application managers, and database administrators. The *principle of least privilege* dictates that users (and programs for that matter) should only have the privileges they need to do their job. This applies in general to all users but is especially important when elevate privileges are assigned to some.

The *principle of rotation of duties* calls for employees to change responsibilities to reduce the chance of an employee successfully committing a malicious act against a business. For example, the systems administrators in two departments may exchange responsibilities every 6 months. Knowing that someone else will have access to the same systems, files, and processes you have and could discover unauthorized activity may deter would-be attackers. Of course, there is still a risk of collusion between employees that rotate among the same position, but following this principle does introduce another hurdle to overcome for an attacker

Web Application Security

Another area of concern when prioritizing security spending is how to meet the demands of Web application security. Applications vary in complexity, the environments in which they operate, and their level of exposure to threats, but there are several practices that can serve the security needs of many Web applications. First, let's look at some of the business drivers that demand attention to Web application security.

Highly damaging attacks continue to focus on application vulnerabilities and the number of vulnerabilities affecting Web applications is one of the fastest growing security problems. This is not surprising given that:

- 54.9% of vulnerabilities are Web application vulnerabilities
- SQL Injection attacks increased by 30x over a 6-month period from late 2008 to early 2009
- 74% of Web application vulnerabilities had no patch by the end of 2008

Vulnerabilities provide the means while valuable information provides the motive for many attacks. Stealing sensitive information is the second largest motivation for Web application attacks (Source: Web Incidents Hacking Database 2008 Annual Report). These attacks are costly, too. An average security breach costs $6.6 million, including the cost of customer notifications, fines, lost customer trust, brand loss, lawsuits, and disruption of business operations.

In the spirit of starting at the beginning, businesses should build security testing into the software development life cycle. Programming best practices, such as validating user input, checking boundary conditions, and performing code reviews can help mitigate threats from injection attacks, buffer overflows, and erroneous or malicious code slipping into production. It is also significantly cheaper to build it in from the beginning. Testing is also essential and should include white box and black box testing as feasible.

Organizations can improve application security by focusing efforts on several practices, including:

- Understanding the most common Web vulnerabilities
- Designing and building security into applications from their inception in order to help mitigate the risk of internal and external threats once applications are deployed
- Testing all applications before they go into production
- Ensuring built-in security defenses are effective before and after deployment
- Protecting vital customer data and information assets from external and internal threats
- Building internal security awareness and providing secure coding training for developers

White box testing, also known as structural testing, uses knowledge about the structure of an application to develop test cases. Black box testing starts with an understanding of how the system is supposed to function and what input it accepts and then tests for appropriate output. Black box testing is useful for ensuring a program meets specifications but can also help identify weaknesses, such as non-validated user input. More sophisticated application security scanners can probe applications in search of known vulnerabilities and can provide more comprehensive reporting.

Resource

For more information about application security best practices and the most common hack attacks, see the Open Web Application Security Project (OWASP) at http://www.owasp.org.

Security managers and developers should concentrate on identifying and resolving security vulnerabilities while Web applications are still in the design and development process. In this way, they can help protect the organization from threats that could compromise data and systems once these applications are deployed. A comprehensive approach to Web application security—in which security concerns play an important role in every stage of the application's life cycle—can keep an organization a step ahead.

Do not overlook developer training when it comes to prioritizing security spending. Investments in developers who can learn to avoid common security pitfalls can have payoffs well into the future.

Moving from design consideration to implementation issues brings what will likely be a familiar list of security best practices.

Monitoring and Management

Day-to-day operations require their own particular set of security practices that should be familiar to readers of this guide. The practices outlined here fit into risk management practices as a means of mitigating risks to a business; they are also elements of everyday security management practices. The following sections highlight those that should be in place in most well-established security practices.

Monitoring and Reporting

Monitoring and reporting systems should be in place to allow security professionals, network administrators, and systems managers to measure performance and events in the infrastructure. Measurements should align with control structures so that they provide feedback on the effectiveness of controls. For example, one metric is the average time required to provision a new user or the average time required to resolve a first-tier security support issue. The objective is to maximize benefit to the organization while reducing costs by measuring and tuning key security-related processes.

In an ideal world, we would have a consolidated reporting system that could populate a security management dashboard, highlight critical events, and provide indicators of potentially problematic situations. We are not there yet, but basic security management reporting that can collect logs and alerts from multiple systems and provide a single point of access is a start and outsourced solutions that provide security reporting and management are readily available. Truly consolidated, comprehensive security data reporting is not here yet but we can make use of existing tools along the way.

Log Management

Priority can be given to improving event reporting in order to control IT operations costs. Effective log management, for example, can help reduce false alarms by providing multiple pieces of data around an event. What may be a triggered event in one system might not be supported by data provided by other systems. The key is to have data from multiple systems available in one place and that is the reason to emphasize consolidated reporting. In addition to the immediate feedback on the validity of reported events, this type of data may also be useful for tuning intrusion prevention systems (IPS).

Change and Patch Management

Analogous to the user life cycle is the life cycle of infrastructure, devices, and applications. Devices are added to networks, virtual hosts are provisioned, network segments are redesigned, and applications are updated. The practices of change management and patch management contribute to improved security by maintaining up-to-date information about devices and the software that runs on them as well as keeping them up to date with the appropriate software patches. With the growth of zero-day security exploits, it becomes more and more critical that security patches are applied immediately across your enterprise to prevent hackers from exploiting known vulnerabilities when they are externalized.

Network Security

Of the many aspects of network security, we should ensure adequate attention is given to gateway and endpoint security along with network access control policies. This is one area of security where different businesses may have different requirements, but it is an area where there is likely to be many common practices. Also, at the level of network access control policies, there is limited need for close alignment with fine-grained business strategies; network access control policies are a good candidate for outsourcing. The advantages of outsourcing are, first, it may reduce costs, and second, a security service provider may be better able to keep abreast of emerging threats, more efficiently monitor the network, and ensure policies are implemented as expected.

Creating and Maintaining Policies

A policy is a set of rules and practices a business chooses to enforce as part of the effort to maintain the confidentiality, integrity, and availability of its information resources. Policy-based management of network devices can be a cost-effective way to automate policy enforcement and improve compliance. Of course, policies are not just for devices; policies should be in place to cover an array of topics, ranging from the type of data that can be transferred to employee-owned smartphones and PDAs to OS and antivirus software requirements for any device connecting to the corporate network.

Countering the "threat of the day" and addressing long-term continuous threats requires a multifaceted approach combining technology, policies, practices, and people. It is now time to turn to the final topic: education.

Education and Security

When considering security spending priorities, it is understandable that many will focus on technical solutions; some others will add business practices, such as monitoring and auditing to their topics of interest. It is not at all clear how education fairs in comparison and this is unfortunate.

Security education is important across the organization. C-level executives need to understand governance responsibilities and be able to hold executives responsible for their stewardship of information assets. Software developers must be aware of sound design and coding practices. End users should be familiar with common threats and learn enough about the functional aspects of IT to detect obvious malicious activity. We should not expect the average user to be able to diagnose a malware infection on his or her computer, but they should know the basics of secure communications, appropriate use of company systems, and the reasons behind security policies.

In general, we do not just want employees who can remember rote rules ("Don't click on any link in an email") but understand basic principles, like the association of URLs to businesses and how SSL and extended validation digital certificates provide assurances related to identity. Phishers, malware pushers, and social engineers are constantly crafting new ways to attack and a well-educated workforce can be a significant contribution to the security effort. Ultimately, many companies consider employees personally responsible such that security compliance is a condition of employment—if employees carelessly cause or enable a significant security exposure, they can be candidates for dismissal.

Summary

The key to properly prioritizing security spending is remembering that business strategy drives security. We, as IT professionals, do not just protect systems and data as an end in themselves and we do not secure them so much that they are unusable. We manage risks as much as we can within the rational boundaries set by business strategy.

Businesses have an opportunity to align business strategy with security practices. Concerns about maintaining security with a dynamic workforce, the increase in cybercrime, the evolution of technology adoption, and regulatory compliance requirements are some of the factors driving the need for efficient and effective security strategies. Properly implemented, security strategies can provide a framework of controls that sustain a business and mitigate risks sufficiently that innovative business plans are supported rather than curtailed.

The challenges to maintaining a secure environment are increasing. IT infrastructure is complex and new technologies, such as virtualization and cloud computing, are introducing new operational considerations, although the fundamentals of information security have not changed. Instead, security practices and security spending adapt to the dynamic business environment. This chapter has explored best practices and topics to consider in security spending by framing them around four stages of IT management: assessing, designing, deploying, and monitoring/managing. We have also considered the need for attention to security education. The key points of these stages can be summarized as:

- Assess the business environment to identify risks and map those risks to business priorities
- Align security practices with business operations and manage risks across people, data, infrastructure, applications, and business processes
- Use security solutions to control costs by working within established frameworks, such as the network security life cycle, user management life cycle, and Web application security practices
- Align security measures to control structures to maximize benefit while controlling costs

Again, there is no single rule for prioritizing security spending that applies to all businesses. The principles outlined here along with several specific areas of security technologies and practices can provide a starting point for establishing your own security spending priorities.

In spite of the demands to address the evolving threat landscape, the dictates of regulatory compliance, and dynamic market factors, businesses can use security practices to their advantage while mitigating risks. The process begins with prioritizing security spending as outlined in this guide.