Introduction So far in this book we have explored various threats to networks. And in Chapter 9, “Computer Security Technology,” we examined a variety of technical defenses against such attacks. However, the fact is that technology by itself cannot solve all network security problems. There are some issues that technology cannot stop. Examples of this include the following: ■ Virus software won’t prevent a user from manually opening an attachment and releasing a virus. ■ A technologically secured network is still very vulnerable if former employees (perhaps some unhappy with the company) still have working passwords or if passwords are simply put on Post-it notes on computer monitors. ■ A server is not secure if it is in a room that nearly everyone in the company has access to. ■ Your network is not secure if end users are vulnerable to social engineering. 250 Defining User Policies 251 Another reason that technology alone is not the answer is that technology must be appropriately applied. Policies are used to guide you in how to implement and manage security, including security technology. In this chapter, we will examine computer security policies, including the elements that go into creating good security policies as well as examples of how to establish a network security policy. What Is a Policy? A security policy is a document that defines how an organization will deal with some aspect of security. There can be policies regarding end-user behavior, IT response to incidents, or policies for specific issues and incidents. Security policies can also be created to deal with regulatory requirements. These types of policies direct members of the organization as to how to comply with certain regulations. A good example would be a policy informing healthcare workers how to comply with HIPAA when using electronic medical records software. Or policies can simply be advisory, suggesting to employees how they should handle certain items, but not requiring compliance. For example, a policy might advise users that emailing from a smart phone using a Wi-Fi hotspot can be unsecure, but not forbid it. Defining User Policies When discussing user policies, there is one rule you must keep in mind: You should have a policy for every foreseeable situation. Failure to have policies that address a given problem will usually result in that problem being exacerbated. Something may seem like common sense to you but may not be to someone with no training or experience in computer networks or network security. The misuse of systems is a major problem for many organizations. A large part of the problem comes from the difficulty in defining exactly what is misuse. Some things might be obvious misuse, such as using company time and computers to search for another job or to view illicit websites. However, other areas are not so clear, such as an employee using her lunchtime to look up information about a car she is thinking of buying. Generally, good user policies will outline specifically how people are to use the system and how they should not. For a policy to be effective, it needs to be very clear and quite specific. Vague statements such as “computers and Internet access are only for business use” are simply inadequate. I would recommend something more clear and perhaps more enforceable, perhaps something like “computers and Internet access are only for business purposes during business hours. However, employees may use the computer/Internet access for personal use during nonwork time such as breaks, lunch, and before work. However, such use must be in compliance with Internet usage policies.” That is clear, direct, and enforceable. Other areas for potential misuse are also covered by user policies, including password sharing, copying data, leaving accounts logged on while employees go to lunch, and so on. All of these issues ultimately 252 CHAPTER 10 Security Policies have a significant impact on your

network's security and must be clearly spelled out in your user policies. We will now examine several areas that effective user policies must cover: ■ Passwords ■ Internet use ■ Email usage ■ Installing/uninstalling software ■ Instant messaging ■ Desktop configuration ■ Bring Your Own Device Passwords Keeping passwords secure is critical. In Chapter 8, "Encryption," appropriate passwords were discussed as part of operating system hardening. A good password is at least eight characters long, uses numbers and special characters, and has no obvious relevance to the end user. For example, a Dallas Cowboys fan would be ill advised to use a password like cowboys or godallas but might be well advised to use a password such as %trEe987 or 123DoG$$ since those don't reflect the person's personal interests and therefore would not be easily guessed. Issues such as minimum password length, password history, and password complexity come under administrative policies, not user policies. User policies dictate how the end user should behave. For reliable security, I recommend a passphrase that has been altered to include numbers and special characters. This can be something easy to remember but altered so that it will not be vulnerable to guessing or brute-force attacks. An example would be taking the phrase "I like double cheese burgers" and altering it to be IliK3double3ch33$eburg3r$. Notice the Es were changed to 3s, the Ss to $s, and two random letters were capitalized. You now have a 25-character password that is also complex. It is easy to remember and very difficult to break. However, no password is secure, no matter how long or how complex, if it is listed on a sticky note stuck to the user's computer monitor. This may seem obvious, but it is not at all uncommon to go into an office and find a password either on the monitor or in the top drawer of the desk. Every janitor or anyone who simply passes by the office can get that password. It is also not uncommon to find employees sharing passwords. For example, Bob is going to be out of town next week, so he gives Juan his password so that Juan can get into his system, check email, and more. The problem is that now two people have that password. And what happens if during the week Bob is gone, Juan gets ill and decides he will share the password with Shelly so that she can keep checking that system while Juan is out sick? It does not take long for a password to get to so many people that it is no longer useful at all from a security perspective. Issues like minimum length of passwords, password age, and password history are ones of administrative policies. System administrators can force these requirements. However, none of that will be particularly helpful if the users don't manage their passwords in a secure fashion. Defining User Policies 253 All of this means you need explicit policies regarding how users secure their passwords. Those policies should specify the following: ■ Passwords are never to be kept written down in an accessible place. The preference is that they not be written down at all, but if they are, they should be in a secure area such as a lock box at your home (not in the office right next to your computer). ■ Passwords must never be shared with another person for any reason. ■ If an employee believes his password has been compromised, he should immediately contact the IT department so his password can be changed and so that logon attempts with the old password can be monitored and traced. Internet Use Most organizations provide their users with some sort of Internet access. There are several reasons for this. The most obvious reason is email. However, that is hardly the only reason to have Internet access in a business or academic setting. There is also the Web, and even chat rooms. (Believe it or not, they can, and in some cases are being used for business communications.) The Internet can be

used for legitimate purposes within any organization, but it can also bring about serious security problems. Appropriate polices must be in place to govern the use of Internet technologies. The World Wide Web is a wonderful resource for a tremendous wealth of data. Throughout this book, we have frequently referenced websites where one can find valuable security data and useful utilities. The Internet is also replete with useful tutorials on various technologies. However, even nontechnology–related business interests can be served via the Web. Here are a few examples of legitimate business uses of the Web: ■ Sales staff checking competitor's websites to see what products or services they offer, in what areas, perhaps even getting prices ■ Creditors checking the business's AM Best or Standard and Poor's rating to see how their business financial rating is doing ■ Business travelers checking weather conditions and getting prices for travel ■ Online training with webinars ■ Web meetings ■ Online bill payment or in some cases even filing regulatory and government documents Of course, there are other web activities that are clearly not appropriate on a company's network: ■ Using the Web to search for a new job ■ Any pornographic use ■ Any use that violates local, state, or federal laws 254 CHAPTER 10 Security Policies ■ Use of the Web to conduct your own business (you have another enterprise you are involved in other than the company's business) In addition, there are gray areas. Some activities might be acceptable to some organizations, but not to others. Such activities might include ■ Online shopping during the employee's lunch or break time ■ Reading news articles online during lunch or break time ■ Viewing humorous websites What one person might view as absurdly obvious might not be to another. It is critical that any organization have very clear policies detailing specifically what is and what is not acceptable use of the Web at work. Giving clear examples of what is acceptable use and what is not is also important. You should also remember that most proxy servers and many firewalls can block certain websites. This will help prevent employees from misusing the company's web connection. Email Usage Most business and even academic activity now occurs via email. As we have discussed in several previous chapters, email also happens to be the primary vehicle for virus distribution. This means that email security is a significant issue for any network administrator. Clearly, you cannot simply ban all email attachments. However, you can establish some guidelines for how to handle email attachments. Users should only open an attachment if it meets the following criteria: ■ It was expected. (Someone requested documents from a colleague or client.) ■ If it was not expected, did it come from a known source? If so, first send that person an email (or phone them) and ask if she sent the attachment. If so, open it. ■ It appears to be a legitimate business document (a spreadsheet, a document, a presentation, and so on). It should be noted that some people might find such criteria unrealistic. There is no question they are inconvenient. However, with the prevalence of viruses, often attached to email, these measures are prudent. Many people choose not to go to this level to try to avoid viruses, and that may be your choice as well. Just bear in mind that millions of computers are infected with some sort of virus every single year. No one should ever open an attachment that meets any of the following criteria: ■ It comes from an unknown source. ■ It is some active code or executable. ■ It is an animation/movie. Defining User Policies 255 ■ The email itself does not appear legitimate. (It seems to entice you to open the attachment rather than simply being a legitimate business communication that happens to have an attachment.) If the end user has any doubt whatsoever, then she should not open

the email. Rather, she should contact someone in the IT department who has been designated to handle security. That person can then either compare the email subject line to known viruses or can simply come check out the email personally. Then if it appears legitimate the user can open the attachment. FYI: About Attachments I frequently follow the “better safe than sorry” axiom on this matter. This means that when forwarded some joke, image, flash animation, and so on circulating the Internet, I simply delete it. That may mean that I will miss many humorous images and stories, but it also means I will miss many viruses. You would do well to consider emulating this practice. Installing/Uninstalling Software This is one matter that does have an absolute answer. End users should not be allowed to install anything on their machine. This includes wallpapers, screensavers, utilities—anything. The best approach is to limit their login privileges so that they cannot install anything. However, this should be coupled with a strong policy statement prohibiting the installation of anything on their PC. If they wish to install something, it should first be scanned by the IT department and approved. This process might be cumbersome, but it is necessary. Some organizations go so far as to remove or at least disable media drives (CD, USB, and so on) from end-user PCs, so installations can only occur from files that the IT department has put on some network drive. This is usually a more extreme measure than most organizations will require, but it is an option you should be aware of. In fact, Windows allows the administrator to disable allowing new USB devices. So the admin can install some USB devices that are approved for corporate use and then disallow any additional devices being added. Instant Messaging Instant messaging is also widely used and abused by employees in companies and organizations. In some cases, instant messaging can be used for legitimate business purposes. However, it does pose a significant security risk. There have been viruses that specifically propagated via instant messaging. In one incident, the virus would copy everyone on the user’s buddy list with the contents of all conversations. Thus, a conversation you thought was private was being broadcast to everyone you knew and had messaged with. Instant messaging is also a threat from a purely informational security perspective. Nothing stops an end user from instant messaging trade secrets or confidential information without the traceability of email going through the corporate email server. It is recommended that instant messaging simply be 256 CHAPTER 10 Security Policies banned from all computers within an organization. If you find you absolutely must have it, then you must establish very strict guidelines for its use, including the following: ■ Instant messaging can only be used for business communications, no personal conversations. Now, this might be a bit difficult to enforce. Rules like this often are. More common rules, such as prohibiting personal web browsing, are also quite difficult to enforce. However, it is still a good idea to have those rules in place. Then if you find a person violating them, you do have a company policy that you can refer to that prohibits such actions. However, you should be aware that in all likelihood you won’t catch most violations of this rule. ■ No confidential or private business information should be sent via instant messaging. Desktop Configuration Many users like to reconfigure their desktop. This means changing the background, screensaver, font size, and resolution. Theoretically speaking, this is not a security hazard. Simply changing your computer’s background image cannot compromise your computer’s security. However, there are other issues involved. The first issue is where the background image comes from. Frequently, end users will download images from the Internet. This

means there is a chance of getting a virus or Trojan horse, particularly one using a hidden extension (that is, it appears to be a mypic.jpg but is really mypic.jpg.exe). There are also human resources/harassment issues if an employee uses a backdrop or screensaver that is offensive to other employees. Some organizations simply decide to prohibit any changes to the system configuration for this reason. The second problem is technical. In order to give users access to change screensavers, background images, and resolution, you must give them rights that will also allow them to change other system settings you might not want changed. The graphical display options are not separated from all other configuration options. This means that allowing users to change their screensaver might open the door for them to alter other settings (such as the network card configuration or the Windows Internet Connection Firewall) that would compromise security. Bring Your Own Device Bring your own device (BYOD) has become a significant issue for most organizations. Most, if not all, of your employees will have their own smart phones, tablets, smart watches, and Fitbits that they will carry with them into the workplace. When they connect to your wireless network, this introduces a host of new security concerns. You have no idea what networks that device previously connected to, what software was installed on them, or what data might be exfiltrated by these personal devices. In highly secure environments, the answer may be to forbid personally owned devices. However, in many organizations, such a policy is impractical. A workaround for that is to have a Wi-Fi network that is dedicated to BYOD and is not connected to the company's main network. Another approach, albeit more technologically complex, is to detect the device on connection, and if it is not a companyissued device, significantly limit its access. Defining User Policies 257 Whatever approach you take, you must have some policy regarding personal devices. They are already ubiquitous. Just a few years ago smart phones were around, but smart watches were not. It is difficult to predict what new smart devices might loom on the horizon. Final Thoughts on User Policies This section has provided an overview of appropriate and effective user policies. It is critical that any organization implement solid user policies. However, these policies will not be effective unless you have clearly defined consequences for violating them. Many organizations find it helpful to spell out specific consequences that escalate with each incident, such as the following: ■ The first incident of violating any of these policies will result in a verbal warning. ■ A second incident will result in a written warning. ■ The third incident will result in suspension or termination. (In the case of academic settings, this would be suspension or expulsion.) You must clearly list the consequences, and all users should sign a copy of the user policies upon being hired. This prevents employees claiming they were not aware of the policies. CAUTION Termination or Expulsion Any policy that can lead to expulsion from a school or termination from a job (or even a demotion) should first be cleared by your legal advisor. There can be significant legal ramifications for wrongful termination or expulsion. I am not an attorney or an expert in legal matters and cannot provide you with legal advice. It is imperative that you do consult an attorney about these matters. It is also important to realize that there is another cost to misuse of corporate Internet access. That cost is lost productivity. How much time does the average employee spend reading personal email, doing nonbusiness web activities, or instant messaging? It is hard to say. However, for an informal view, go to www.yahoo.com on any given business day, during business hours, and click on one of the news stories. At the bottom of the story you will see a message

board for this story. It lists date and time of posts. See how many posts are done during business hours. It is unlikely that all of those people are out of work, retired, or at home sick. Let me be completely clear. The Internet is the single greatest communication tool in human history. And it can have a tremendous positive effect on any business. I conduct almost all of my business activities through the Web. However, many employees do abuse the Internet, and it does decrease productivity for those who cannot be self-disciplined. Here are just a few studies supporting this assertion: ■ A 2008 study showed employees wasting an average of 5.3 hours a week on idle Internet surfing. While this study is a bit old, the problem persists to this day. It is often the case that widespread Internet access actually inhibits productivity rather than enhances it. 258 CHAPTER 10 Security Policies ■ Ohio State University researchers found regular Facebook users had a lower GPA than nonusers (http://researchnews.osu.edu/archive/facebookusers.htm). ■ There have been multiple studies showing that Facebook at work has a negative impact on productivity. Most of these studies are a few years old, because it has become well established that social media at work hurts productivity. http://www.cnbc.com/2016/02/04/facebook-turns-12--trillions-in-time-wasted.html http://www.riskmanagementmonitor.com/the-risks-of-social-media-decreased-workerproductivity/ http://www.shellypalmer.com/2011/05/social-media-use-drastically-reduces-work-productivity/ Defining System Administration Policies In addition to determining policies for users, you must have some clearly defined policies for system administrators. There must be a procedure for adding users, removing users, dealing with security issues, changing any system, and so on. There must also be standards for handling any deviation. New Employees When a new employee is hired, the system administration policy must define specific steps to safeguard company security. New employees must be given access to the resources and applications their job function requires. The granting of that access must be documented (possibly in a log). It is also critical that the new employee receive a copy of the company's computer security/acceptable use polices and sign a document acknowledging receipt of such. Before a new employee starts to work, the IT department (specifically network administration) should receive a written request from the business unit that person will be working for. That request should specify exactly what resources this user will need, when she will start, and have the signature of someone in the business unit with authority to approve such a request. Then the person who is managing network administration or network security should approve and sign the request. After you have implemented the new user on the system with the appropriate rights, you can file a copy of the request. Departing Employees When an employee leaves, it is critical to make sure all of his logins are terminated and all access to all systems is discontinued immediately. Unfortunately, this is an area of security that all too many organizations do not give enough attention to. You cannot be certain which employees will bear the company ill will and which won't upon leaving the company. It is imperative to have all of the former employees' access shut down on their last day of work. This includes physical access to the building. If a former employee has keys and is disgruntled, nothing can stop him from returning to steal or vandalize computer equipment. When an employee leaves the company, you need to ensure that on his last day the following actions take place: Defining System Administration Policies 259 ■ All logon accounts to any server, VPN, network, or other resource are disabled. ■ All keys to the facility are returned. ■ All accounts for email, Internet

access, wireless Internet, cell phones, and so on are shut off. ■ Any accounts for mainframe resources are canceled. ■ The employee's workstation hard drive is searched. The last item might seem odd. But if an employee was gathering data to take with him (proprietary company data) or conducting any other improper activities, you need to find out right away. If you do see any evidence of such activity, you need to secure that workstation and keep it for evidence in any civil or criminal proceedings. All of this might seem a bit extreme to some readers. It is true that with the vast majority of exiting employees you will have no issues to be concerned about. However, if you do not make it a habit of securing employees' access when they depart, you will eventually have an unfortunate situation that could have been easily avoided. Change Requests The nature of information technology is change. Not only do end users come and go, but requirements change frequently. Business units request access to different resources, server administrators upgrade software and hardware, application developers install new software, and web developers change the website. Change is occurring all the time. Therefore, it is important to have a change control process. This process not only makes the change run smoothly, but it also allows the IT security personnel to examine the change for any potential security problems before it is implemented. A change control request should go through the following steps: ■ An appropriate manager within the business unit signs the request, signifying approval of the request. In other words, there is no point in pursuing the change request process if the immediate supervisor of the requestor has not approved the request. ■ The appropriate IT unit (database administration, network admin, email admin, cloud administration) verifies that the request is one it can fulfill technologically, fits within budget constraints, and does not violate IT policies. ■ The IT security unit verifies that this change will not cause security problems. This is becoming more and more critical in modern times. ■ The appropriate IT unit formulates a plan to implement the change and a plan to roll back the change in the event of some failure. That latter part is very critical and is often overlooked. There must be some mechanism to roll back the change should it cause any problems. ■ The date and time for the change is scheduled, and all relevant parties are notified. 260 CHAPTER 10 Security Policies Your change control process might not be identical to this one; in fact, you might be much more specific. However, the key to remember is that in order for your network to be secure, you simply cannot have changes happening without some process for examining the impact of those changes prior to implementing them. Change management activities are frequently managed through a Change Control Board (CCB) process, sometimes also called a Change Approval Board (CAB). The change process was detailed previously in this section, but the basic process can be summarized as follows: ■ Initiated with RFC document (Request for Comments or Request for Change) ■ RFC sent for approval ■ Priority is set ■ Assigned to whomever makes the change ■ Document decisions ■ Evaluate by CAB ■ RFC scheduled ■ Complete when change owner and requester verify successful implementation ■ Review of RFC This can be a process with a CAB meeting formally and documentation being extensive, or it can be informal and conducted via emails to the appropriate parties. In Practice Extremes of Change Control Anyone with even a few years of experience in the IT profession can tell you that when it comes to change control there are all sorts of different approaches. The real problem is those IT groups that implement unreasonable extremes. I have seen both. Without using the real names of the companies involved, let's examine a real case of each

extreme. Software consultant Company X was a small company that did custom financial applications for various companies. It had a staff of less than 20 developers who frequently traveled to client locations around the country. It literally had ■ No documentation for any of its applications—not even a few notes. ■ No change control process. When someone did not like a setting on a server or some part of the network configuration, he simply changed it. ■ No process for handling former employee access. In one case, a person had been gone for six months and still had a valid logon account. Defining System Administration Policies 261 Now, clearly this is alarming from several perspectives, not just from a security viewpoint. However, that is one extreme, one that makes for a chaotic environment that is very unsecure. Security-minded network administrators tend to move toward the opposite extreme, one that can have a negative impact on productivity. Company B had over 2,000 employees, with an IT staff of about 100 people. In this company, however, the bureaucracy had overwhelmed the IT department to the point that their productivity was severely impacted. In one case, a person was a web server administrator, and the decision had been made that he also needed database administration rights on a single database server. The process, however, took three months with one face-to-face meeting between his manager and the CIO, as well as two phone conferences and a dozen emails between his manager and the manager of the database group. The company's convoluted change control process had a severely negative impact on productivity. Some employees informally estimated that even the low-level IT supervisors spent 40% of their time in meetings/conferences, reporting on meetings/conferences, or preparing for meetings/ conferences. And the further one went up the IT ladder, the more one's time became consumed in bureaucratic activities. Both of these examples are meant to illustrate two extremes in change control management that you should try to avoid. Your goal in implementing change control management is simply to have an orderly and safe way of managing change, not to be an impediment to productivity. Security Breaches Unfortunately, the reality is that your network will probably, at some point, have a security breach of some kind. This could mean that you are the target of a denial of service (DoS) attack, your system is infected with a virus, or perhaps a hacker gains entrance and destroys or copies sensitive data. You must have some sort of plan for how to respond should any such event occur. This book cannot tell you specifically how to deal with each and every event that might occur; however, we can discuss some general guidelines for what to do in certain, general, situations. We will look at each of the main types of security breaches and what actions you should take for each. Virus Infection When a virus strikes your system, immediately quarantine the infected machine or machines. This means literally unplugging the machines from the network. If it is a subnet, then unplug its switch or disconnect wireless access. Isolate the infected machines (unless your entire network is infected, in which case simply shut down your router/ISP connection to close you off from the outside world and prevent spread beyond your network). After implementing the quarantine, you can safely take the following steps. ■ Scan and clean each and every infected machine. Since the machines are now off the network, this will be a manual scan. 262 CHAPTER 10 Security Policies ■ Log the incident, the hours/resources taken to clean the systems, and the systems that were affected. ■ When you are certain the systems are clean, bring them online in stages (a few at a time). With each stage, check all machines to see that they are patched, updated, and have properly configured/running

antivirus. ■ Notify the appropriate organization leaders of the event and the actions you have taken. ■ After you have dealt with the virus and notified the appropriate people, you should then have a meeting with appropriate IT staff to discuss what can be learned from this breach and how you might prevent it from occurring in the future. DoS Attacks ■ If you have taken the steps outlined earlier in this book (such as properly configuring your router and your firewall to reduce the impact of any attempted DoS), then you will already be alleviating some of the damage from this type of attack. Use your firewall logs or IDS to find out which IP address (or addresses) originated the attacks. Note the IP addresses, and then (if your firewall supports this feature, and most do) deny that IP address access to your network. ■ Use online resources (interNIC.net and so on) to find out who the address belongs to. ■ Contact that organization and inform it of what is occurring. ■ Log all of these activities and inform the appropriate organizational leaders. ■ After you have dealt with the DoS and notified the appropriate people, you should have a meeting with appropriate IT staff to discuss what can be learned from this attack and how you might prevent it from occurring in the future. Intrusion by a Hacker ■ Immediately copy the logs of all affected systems (firewall, targeted servers, and so on) for use as evidence. ■ Immediately scan all systems for Trojan horses, changes to firewall settings, changes to port filtering, new services running, and so on. In essence, you are performing an emergency audit to see what damage has been done. ■ Document everything. Of all of your documentation, this must be the most thorough. You must specify which IT personnel took what actions at what times. Some of this data may be part of later court proceedings, so absolute accuracy is necessary. It is probably a good idea to log all activities taken during this time and to have at least two people verify and sign the log. ■ Change all affected passwords. Repair any damage done. ■ Inform the appropriate business leaders of what has happened. Defining Access Control 263 ■ After you have dealt with the breach and notified the appropriate people, you should have a meeting with appropriate IT staff to discuss what can be learned from this breach and how you might prevent it from occurring in the future. These are just general guidelines, and some organizations may have much more specific actions they want taken in the event of some security breach. You should also bear in mind that throughout this book when we have discussed various sorts of threats to network security, we have mentioned particular steps and policies that should be taken. The policies in this chapter are meant to be in addition to any already outlined in this book. It is an unfortunate fact that some organizations have no plan for what to do in case of an emergency. It is important that you do have at least some generalized procedures you can implement. Defining Access Control An important area of security policies that usually generates some controversy in any organization is access control. There is always a conflict between users' desire for unfettered access to any data or resources on the network and the security administrator's desire to protect that data and resources. This means that extremes in policies are not practical. You cannot simply lock down every resource as completely as possible since that would impede the user's access to those resources. Conversely, you cannot simply allow anyone and everyone complete access to everything. The core of access control is the concept introduced in Chapter 1, "Introduction to Computer Security": least privileges. Each person is given the minimum privileges necessary to do her job. No more and no less. This is where the least privileges concept comes into play. The idea is simple. Each user, including IT

personnel, gets the least access he can have and still effectively do his job. Rather than ask the question, “why not give this person access to X?” you should ask, “why give this person access to X?” And if you don’t have a very good reason, then don’t. This is one of the fundamentals of computer security. The more people that have access to any resource, the more likely some breach of security is to occur. Along with, and related to, least privileges is the concept of implicit deny. Implicit deny means that all users are implicitly denied access to network resources until an administrator explicitly grants them. Separation of duty, job rotation, and mandatory vacations are also important and related concepts. Separation of duty means that no one person can perform critical tasks; at least two individuals are needed. This prevents one person from accidently, or intentionally, causing some security breach via inappropriate use of critical functions. Both job rotation and mandatory vacations are used to make sure that, periodically, the person performing a given job changes. This makes it more difficult for one person to exploit her position to breach security. Obviously, trade-offs must be made between access and security. Examples abound. One common example involves sales contact information. Clearly, a company’s marketing department needs access to this data. However, what happens if your competitors get all of your company’s contact information? That could allow them to begin targeting your current client list. This requires a trade-off 264 CHAPTER 10 Security Policies between security and access. In this case, you would probably give salespeople access only to the contacts that are within their territory. No one other than the sales manager should have complete access to all the marketing data. Developmental Policies Many IT departments include programmers and web developers. Unfortunately, many security policies do not address secure programming. No matter how good your firewalls, proxy server, virus scanning, and policies are, if your developers create code that is flawed, you will have security breaches. Clearly, the topic of secure programming requires a separate volume to explore thoroughly. Nonetheless, we can consider a brief checklist for defining secure development policies. If your company currently has no secure programming initiatives, this checklist is certainly better than developing in a vacuum. It can also serve as a starting point to get you thinking and talking about secure programming: ■ All code, especially code done by outside parties (contractors, consultants, and so on) must be checked for backdoors/Trojan horses. ■ All buffers must have error handling that prevents buffer overruns. ■ All communication (such as using TCP sockets to send messages) must adhere to your organization’s secure communications guidelines. ■ Any code that opens any port or performs any sort of communication is thoroughly documented, and the IT security unit is apprised of the code, what it will do, and how it will be used. ■ All input is filtered for items that might facilitate an attack, such as an SQL injection attack. ■ All vendors should supply you with a signed document verifying that there are no security flaws in their code. Following these guidelines will not guarantee that flawed code is immune from being introduced into your system, but it will certainly lower the odds significantly. And the unfortunate fact is that these simple steps alone are more than most organizations are taking. A very good place to look at security policies is the SANS Institute (www.sans.org/security-resources/policies/). Standards, Guidelines, and Procedures Related to policies are standards, guidelines, and procedures. All of these documents are related to security policies and in fact support those policies. A standard is a general statement of the desired level of operation. For example, requiring 99.5% network uptime would

be a standard. A guideline is a general suggestion on how to achieve some standard. Guidelines are broad and are sometimes optional (not mandatory). Procedures are specific instructions on how to handle a specific issue. Data Classification 265 Data Classification It is critical to classify information within your organization. This process is common in defense department–related agencies and organizations. It is less common in the civilian sector. Classifying information provides employees with guidance on how to handle data. Classification can be as simple as two categories: ■ Public information is information that can be disseminated publically to anyone. There are no restrictions on who can view the data. ■ Private information is intended only for use internally in the organization. This type of information can potentially embarrass the company, disclose trade secrets, reveal corporate strategy, expose private personal data of employees or customers, or otherwise reveal information that your organization does not want revealed. This two-tier approach to data classification is rather elementary. Most organizations will have multiple tiers. Each tier is defined by the damage that information disclosure could cause. We will take a look at Department of Defense clearance levels. These provide some insight into varying security classifications. Even if you work in an entirely civilian environment, reviewing the DoD approach can give you some suggestions on how to classify your data and how to properly evaluate which personnel should have access. DoD Clearances The terms secret and top secret have specific meanings. The United States has a specific hierarchy of classification. The lowest is confidential. This is information that might damage national security if disclosed. Secret information is data that might cause serious damage to national security if disclosed. Top secret information is data that could be expected to cause exceptionally grave damage to national security if disclosed. There is another designation: Top Secret SCI or Sensitive Compartmented Information. Each of these clearances requires a different level of investigation. For a secret clearance, a complete background check including criminal, work history, credit check, and check with various national agencies (Department of Homeland Security, Immigration, State Department, and so on) is required. This is referred to as an NACLC or National Agency Check with Law and Credit. The check for employment will cover the last 7 years. The secret clearance may or may not include a polygraph. The top secret clearance is more rigorous, as you may imagine. It uses a Single Scope Background Investigation (SSBI). This means a complete NACL for the subject and spouse that goes back at least 10 years. It will also involve a subject interview conducted by a trained investigator. Direct verification of employment, education, birth, and citizenship are also required. At least four references are necessary, and at least two of those will be interviewed by investigators. A polygraph is also used. The SSBI is repeated every 5 years. Sensitive Compartmented Information is assigned only after a complete SSBI has been completed. An SCI may have its own process for evaluating access; therefore, a standard description of what is involved is not available. 266 CHAPTER 10 Security Policies Disaster Recovery Before we can discuss disaster recovery, we have to define what a disaster is. A disaster is any event that significantly disrupts your organization's operations. A hard drive crash on a critical server is a disaster. Other examples would include fire, earthquake, your telecom provider being down, a labor strike that affects shipping to and from your business, and a hacker deleting critical files. Just keep in mind that any event that can significantly disrupt your organization's operations is a disaster. Disaster Recovery Plan A disaster recovery plan

(DRP) is the plan you have in place to return business to normal operations. This must include a number of items. You must address personnel issues, including being able to find temporary personnel if needed and being able to contact the personnel you have employed. It also includes having specific people assigned to specific tasks. If there is a disaster, who in your organization is tasked with the following: ■ Locating alternative facilities ■ Getting equipment to those facilities ■ Installing and configuring software ■ Setting up the network at the new facility ■ Contacting staff, vendors, and customers These are just a few questions that a disaster recovery plan must address. Business Continuity Plan A business continuity plan (BCP) is similar to a DRP but with a different focus. The DRP is designed to get the organization back to full functionality as quickly as possible. A BCP is designed to get minimal business functions back up and running at least at some level so you can conduct some type of business. An example would be a retail store whose credit card processing system is down. Disaster recovery is concerned with getting the system back up and running, and business continuity is concerned with simply getting a temporary solution, such as processing credit cards manually. To successfully formulate a BCP, you must consider those systems that are most critical for your business and have an alternative plan in case those systems are down. The alternative plan need not be perfect, just functional. Impact Analysis? Before you can create a realistic DRP or BCP, you have to do an impact analysis of what damage to your organization a given disaster might be. This is called a business impact analysis or business impact assessment (BIA). Consider a web server crash. If your organization is an e-commerce business, then Disaster Recovery 267 a web server crash is a very serious disaster. However, if your business is an accounting firm and the website is just a way for new customers to find you, then a web server crash is less critical. You can still do business and earn revenue while the web server is down. You should make a spreadsheet of various likely or plausible disasters and do a basic BIA for each. A few things go into your BIA. One item to consider is the maximum tolerable downtime (MTD). How long can a given system be down before the effect is catastrophic and the business is unlikely to recover? Another item to consider is the mean time to repair (MTTR). How long is it likely to take to repair a given system if it is down? These factors help you determine the business impact of a given disaster. Fault Tolerance The fact is that equipment fails. At some point all equipment fails. So fault tolerance is important. At the most basic level, fault tolerance for a server means a backup. If the server fails, did you back up the data so you can restore it? While database administrators may use a number of different types of data backups, from a security point of view there are three primary backup types we are concerned with: ■ Full: All changes ■ Differential: All changes since last full backup ■ Incremental: All changes since last backup of any type Consider a scenario where you do a full backup at 2 a.m. each morning. But you are concerned about the possibility of a server crash before the next full backup, so you want to do a backup every two hours. Well, the type of backup you choose will determine the efficiency of doing those frequent backups and the time needed to restore. So let's consider each scenario and what would happen if the system crashes at 10:05 a.m. Full: In this scenario, you do a full backup at 4 a.m., 6 a.m., …10 a.m., and then the system crashes. Well, to restore, you just have to restore the last full backup, done at 10 a.m. This makes restoration much simpler. However, running a full backup every 2 hours is very time consuming and resource intensive and will have a significant

negative impact on your server's performance. Differential: In this scenario, you do a differential backup at 4 a.m., 6 a.m., …10 a.m., and then the system crashes. To restore, you will need to restore the last full backup done at 2 a.m. and the most recent differential backup done at 10 a.m. This is just a little more complicated than the full backup strategy. However, those differential backups are going to get larger each time you do them, and thus more time consuming and resource intensive. While they won't have the impact of the full backups, they will still slow down your network. Incremental: In this scenario, you do an incremental backup at 4 a.m., 6 a.m., …10 a.m., and then the system crashes. To restore, you need to restore the last full backup done at 2 a.m. and then each incremental backup done since then, and they must be restored in order. This is much more complex to restore, but each incremental backup is small and does not take much time or consume many resources. 268 CHAPTER 10 Security Policies There is no "best" backup strategy. Which one you select will depend on your organization's needs. Whatever backup strategy you choose, you must periodically test it. The only effective way to test your backup strategy is to actually restore the backup data to a test machine. The other fundamental aspect of fault tolerance is RAID, or redundant array of independent disks. RAID allows your servers to have more than one hard drive, so that if the main hard drive fails, the system keeps functioning. The primary RAID levels are described here: ■ RAID 0 (striped disks) distributes data across multiple disks in a way that gives improved speed at any given instant. There is no fault tolerance. ■ RAID 1 mirrors the contents of the disks, making a form of 1:1 ratio real–time backup. This is also called mirroring. ■ RAID 3 or 4 (striped disks with dedicated parity) combines three or more disks in a way that protects data against loss of any one disk. Fault tolerance is achieved by adding an extra disk to the array and dedicating it to storing parity information. The storage capacity of the array is reduced by one disk. ■ RAID 5 (striped disks with distributed parity) combines three or more disks in a way that protects data against the loss of any one disk. It is similar to RAID 3, but the parity is not stored on one dedicated drive; instead, parity information is interspersed across the drive array. The storage capacity of the array is a function of the number of drives minus the space needed to store parity. ■ RAID 6 (striped disks with dual parity) combines four or more disks in a way that protects data against loss of any two disks. ■ RAID 1+0 (or 10) is a mirrored data set (RAID 1) that is then striped (RAID 0), hence the "1+0" name. A RAID 1+0 array requires a minimum of four drives: two mirrored drives to hold half of the striped data, plus another two mirrored for the other half of the data. A server without at least RAID level 1 is gross negligence on the part of the network administrator. RAID 5 is actually very popular with servers. While RAID and backup strategies are the fundamental issues of fault tolerance, any backup system provides additional fault tolerance. This can include uninterruptable power supplies, backup generators, and redundant Internet connections. Important Laws There are a number of computer laws in various countries, states, and provinces. It is important to be familiar with the laws that are relevant to your jurisdiction. However, there are a few laws that are most critical in the United States. We will discuss each of those here. Important Laws 269 HIPAA The Health Insurance Portability and Accountability Act (HIPAA) is a regulation that mandates national standards and procedures for the storage, use, and transmission of personal medical information. Passed into law in 1996, HIPAA has caused a great deal of change in healthcare record keeping. HIPAA covers three areas—confidentiality, privacy, and security

of patient records—and was implemented in phases to make the transition easier. Confidentiality and privacy of patient records had to be implemented by a set date, followed by security of patient records. Standards for transaction codes in medical record transmissions had to be completed by a given date as well. The penalties for HIPAA violations are very stiff: They can be as high as $250,000 based on the circumstances. Medical practices are required to appoint a security officer. All related parties, such as billing agencies and medical records storage facilities, are required to comply with these regulations. Sarbanes–Oxley The legislation came into force in 2002 and introduced major changes to the regulation of financial practice and corporate governance. Named after Senator Paul Sarbanes and Representative Michael Oxley, this law is designed to make publically traded corporations more accountable. The legislation focuses primarily on financial issues, but it also affects the IT departments whose job it is to store a corporation's electronic records. The Sarbanes–Oxley Act states that all business records, including electronic records and electronic messages, must be saved for "not less than five years." The consequences for noncompliance are fines, imprisonment, or both. Payment Card Industry Data Security Standards While not a law, the Payment Card Industry Data Security Standards (PCI DSS) are certainly something any IT security professional who works for a company that handles credit cards and debit cards should be familiar with. PCI DSS is a proprietary information security standard for organizations that handle branded credit cards from the major card schemes including Visa, MasterCard, American Express, and Discover.