Chapter 11

IT Infrastructure Security Plan

Solutions in this chapter:

- Infrastructure Security Assessment
- Project Parameters
- Project Team
- Project Organization
- Project Work Breakdown Structure
- Project Risks and Mitigation Strategies
- Project Constraints and Assumptions
- Project Schedule and Budget
- Infrastructure Security Project Outline

☑ Summary

☑ Solutions Fast Track

Introduction

Infrastructure security is at the root of your entire corporate security plan. Other individual security area plans (ISAPs) may overlap with your infrastructure security plan to some extent. For example, a wireless network is part of your infrastructure, but it's also a large enough area to be addressed in a separate project plan. You'll need to ensure that your corporate IT security project and your ISAPs cover all the bases, but be aware that there are overlapping areas that should be clearly delineated if you're working on several projects in parallel. You don't want project teams wrestling over ownership of one part of your network or another. In this chapter, we'll look at the basic infrastructure components and how to secure them; then we'll create a project plan utilizing this information.

Infrastructure Security Assessment

There are two distinct processes: audit and assessment. An *assessment* is intended to look for issues and vulnerabilities that can be mitigated, remediated, or eliminated prior to a security breach. An *audit* is normally conducted after an assessment with the goal of measuring compliance with policies and procedures. Typically, someone is held accountable for audit results. Some people don't like the term *auditing;* perhaps it's too reminiscent of ol' Uncle Sam scouring through your tax return from three years ago when you claimed that one vacation as a business trip because you talked to your boss on your cell phone while waiting at the shuttle to your beachfront hotel. Though the terms *assessment* and *audit* are often used interchangeably, in this chapter we focus on assessments.

As we've discussed throughout this book, there are three primary components of IT security: *people*, *process,* and *technology*. A balanced approach addresses all three areas, because focusing on one area to the exclusion of others creates security holes. People, including senior management, must buy into the importance of security, and they must understand and participate in their role in maintaining security. Process includes all the practices and procedures that occur and reoccur to keep the network secure. Technology obviously includes all hardware and software that comprises the network infrastructure. Part of the technology assessment required to assess and harden infrastructure security includes deploying the right technological solutions for your firm and not the "one size fits all" or the "it was all we could afford" solution. In IT, we often focus a disproportionate amount of time and energy on securing the technology and overlook the importance of both people and process to the overall security environment.

To secure your infrastructure, you need to understand its building blocks. These include:

- Network perimeter protection
- Internal network protection
- Intrusion monitoring and prevention
- Host and server configuration
- Protection against malicious code
- Incident response capabilities
- Security policies and procedures
- Employee awareness and training
- Physical security and monitoring

We'll discuss policies, procedures, and training in the chapter on operational security later in this book, so we won't discuss that material here.

We can look at the infrastructure security assessment in three segments, as shown in Figure 11.1.

Figure 11.1 Infrastructure Assessment Overview

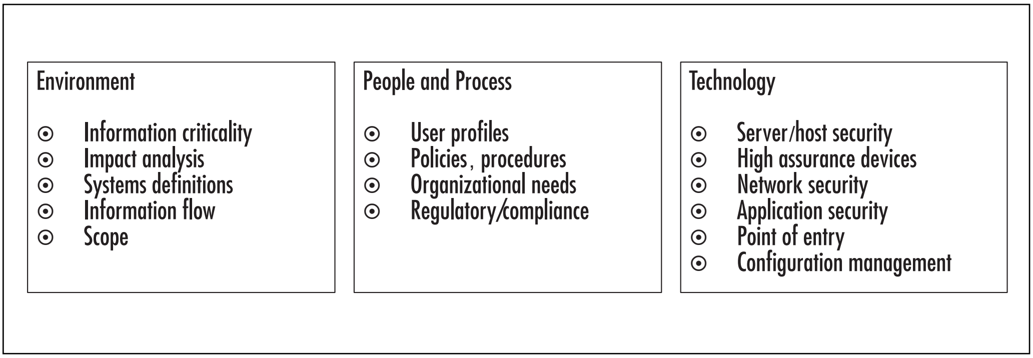

Environment

- Information criticality
- Impact analysis
- Systems definitions
- Information flow
- Scope

People and Process

- User profiles
- Policies, procedures
- Organizational needs
- Regulatory/compliance

Technology

- Server/host security
- High assurance devices
- Network security
- Application security
- Point of entry
- Configuration management

Internal Environment

Security assessments should begin by looking at the overall environment in which security must be implemented, since security does not exist in a vacuum. Looking at the relative importance of your company's information is a good starting point, because you need to find the right balance between security and information criticality. As part of that analysis, you also need to look at the impact of a network infrastructure intrusion and what that would cost to defend and repair. You need to define the various systems you have in place and look at how information flows through your organization to understand the infrastructure you're trying to protect. Finally, you need to create an initial assessment of scope to define what *is* and *is not* included in your project. We'll look at scope later in the chapter, when we begin developing our project plan.

Information Criticality

It's important to begin by looking at information criticality. We've discussed this topic throughout this book, and it will continue to be a common theme because there's really no point in securing something that no one wants. It's why a new Lexus RX-330 comes with a lo-jack system, but a 1993 Dodge Dart with serious body damage is not likely to need any protection (in fact, there might be an economic benefit to having such a vehicle stolen—no offense intended to any 1993 Dodge Dart owners). Information criticality is an assessment of what your network holds and how important that is in the overall scheme of things. Not all data is created equal, and if your company manufactures steel troughs for horse feed, there's a good chance your network data is not nearly as interesting to a potential attacker as the data in an online stock brokerage firm or a bank or credit card processing house network. Therefore, you need to look at the criticality of your information and decide how much you're willing to spend to secure that information. No one ever wants a security breach, but it would not make good business sense to spend $15 million to secure a network for a company that pulls in $5 million annually and doesn't store sensitive personal data such as credit card numbers or medical records. That said, just because your company makes $5 million annually doesn't mean that you *shouldn't* look seriously at the criticality of your data, to be sure you don't have excessive exposure. If you are storing credit card numbers or medical records, you'd better be sure your security solutions are up to standards, because your legal liability could significantly outstrip that $5 million annually in a big hurry.

Impact Analysis

You'll notice as you read the chapters for the individual security area plans that some of this information overlaps. It's hard to perform an impact analysis on an infrastructure breach without also seeing how it would impact your wireless network components, your Web site, or your policies and procedures. However, in looking at the impact to your infrastructure, you'll need to understand how a breach could impact the very foundation of your organization. The impact analysis should include:

- **Cost of network infrastructure—failure (downtime)** Server down, database server down, routers down, etc.

- **Cost of network infrastructure—unavailable (slow or unresponsive)** Denial-of-service attacks, packet flooding, etc.
- **Cost of network infrastructure breach—data confidentiality, integrity, availability** Man-in-the-middle, spoofing, phishing, etc.
- **Cost to company reputation** Lost sales, lost customers, loss of long-term business relationships.
- **Cost to company** Cost of remediation, cost of litigation.

You should combine information criticality with the findings of your impact analysis to form a clear picture of what you're trying to protect and why. When you understand the impact, you can see where the important areas are in your organization, and you can use this information, in part, to prioritize your approach to securing the network.

System Definitions

Infrastructure systems clearly include the "backbone" services, including DHCP servers, DNS servers, Directory Services servers, e-mail servers, database servers, firewalls, DMZs, routers/switches, operating systems, Web servers, and security applications (antivirus, antispyware, IDS/IPS, etc.). If it's helpful, you can also look at your systems from the OSI model perspective—from the physical layer all the way up through the application layer, whatever makes the most sense to you and your team.

Creating (or updating) network diagrams can also be included in the system definitions overview, since the way everything fits together is part of understanding the whole.

Information Flow

One area that is sometimes overlooked in the assessment phase is the flow of information through the infrastructure. This area can be used in conjunction with your systems definitions to help map your network and to discover the key areas that need to be protected and how an attacker would get to those assets.

It sometimes helps to look at information flow from different perspectives. For example, how does information from a user computer flow? How does DNS or DHCP traffic flow through the network? How is external traffic coming into the network managed, and where and how does it enter? How is traffic leaving the network for the public network (Internet) managed? Creating a map of your network infrastructure and information flow will help you visualize your network and identify potential weak spots.

Scope

You might want to limit the scope of your infrastructure security project for a variety of reasons. While you're looking at your internal environment, you might choose to limit the scope. "Scoping" is often done at this point when you're engaging an external security consultant. However, if you're doing this work internally, you may limit your scope here, or you may choose to do a full assessment and then limit the scope after you see what's what.

People and Process

Clearly, people and processes impact network security in a big way. Most security breaches occur from the inside, not the outside, despite the media's sensationalized focus on external security breaches. The people in your organization can be your defenders or your downfall, depending on how they approach security. Savvy, well-informed users can augment the technical security measures by avoiding becoming victims of social engineering, by reporting suspicious activity, by avoiding responding to phishing e-mail, or by not leaving their computer logged in and unattended. All the security in the world can't prevent problems if users are not pulling their weight. There are many ways to inform and involve users, and unfortunately, many IT departments don't leverage these opportunities very successfully, because they often fall victim to a "user as pain in the hind quarters" mentality. Let's look at how users and organizational processes should be reviewed during an infrastructure assessment.

User Profiles

What kinds of users do you have? Where and how do they work? If you begin by looking at your user population, you will see segments that have higher and lower risk profiles. The clerk in the mailroom might only have access to e-mail and the mailroom application, but does he or she also have Internet access and the ability to download and install programs? What about the marketing staff who travel worldwide? What kinds of information do they keep on their laptops (usernames, passwords, domain names, sensitive documents, contacts, and the like), and how does this impact your network security?

Users can be categorized in whatever ways work for you in your organization, but here's a list of potential risks by employee type, to get you thinking:

- **Executive** High-profile targets, often not extremely "tech savvy," potentially easy to get information about (from press releases, public filings, legal filings, and so on).
- **Director** High-profile targets, may travel extensively with sensitive information, may need to connect to the network in a variety of insecure locations.
- **Finance**, **marketing**, **HR, legal** Access to extremely sensitive data, may be high-profile targets due to their access to sensitive data, may travel extensively and be desirable targets of social engineering.
- **IT staff** Access to network resources, ability to grant/deny access, potentially desirable targets of social engineering (especially via help desk), highly desirable targets (IT usernames and passwords with administrative privileges are the Holy Grail for hackers).
- **Users** Access to sensitive company information, often targets of social engineering.

In addition to these categories, you may have user groups defined in your network security management system (which manages access control) that you want to use. Microsoft defines users as administrators, power users, and the like, and that might also work for you. Again, the point is to use a categorization method that's meaningful to the way your company and your existing network infrastructure are organized, so you can understand the risks users bring into the organization and the strategies for keeping the network secure in light of the way various users work.

Policies and Procedures

We won't spend a lot of time discussing policies and procedures in this chapter; we'll focus on them in an upcoming chapter on operational security. As we've discussed, no single security topic exists in a vacuum or silo, and as you move through your project planning, you'll notice areas of intersection and overlap. There are few hard-and-fast rules about where these overlapping elements should be placed; the important factor is to be sure they *are* included *someplace*.

Infrastructure policies and procedures touch on the day-to-day operations of the IT staff, including the way security is monitored (auditing functions, log files, alerts) and how it is maintained (backups, updates, upgrades). Policies regarding user behavior are also crucial to ensuring that the network infrastructure remains safe. Finally, corporate policies regarding the use of data, computer and electronic equipment, and building access, to name just three, are areas that should be reviewed and revised to support and enhance security across the enterprise. For more specific information on polices and procedures, see Chapter 13, "Operational Security."

Organizational Needs

The internal environment is shaped by the organization's business profile, including the type of business, the nature of sales and marketing functions, the types of customers, the kinds of employees, and the flow of work through the company. What does your company require from the network services you provide, and how can these needs be secured? If you believe your organization's network, data, and computer needs are being met, delineate what they are and check with a few users to see if you're on the mark or if you're really off-center by a wide margin. Make sure that you understand how the network fits into the organization, not the other way around, and then design your security solution around it.

Regulatory/Compliance

Any infrastructure assessment and security plan must incorporate regulatory and compliance requirements. These vary greatly from state to state and country to country, and as you're probably well aware, keeping up with them can be more than a full-time job. Many companies are hiring compliance officers whose primary job is to manage corporate compliance. If your company has a compliance officer, you should certainly make sure he or she is a member of your IT project team, at least during the definition phase, when you're developing your functional and technical requirements, since these are often the method by which compliance occurs. We've included a short list here with a few Web site links, but it's not exhaustive; you should seek legal advice regarding regulatory and compliance requirements for your firm if you don't have a knowledgeable and experienced compliance officer in place.

Business Intelligence…

Common Compliance Standards

There are numerous compliance issues facing organizations today. Below are just a few of the compliance standards you should be aware of and should evaluate whether your firm is subject to these regulations or not.

British Standard 7799 (BS7799), eventually evolved into ISO17799.

Child Online Protection Act (COPA), www.copacommission.org.

Continued

Health Insurance Portability and Accountability Act (HIPAA), www.cms.hhs.gov/hipaa/hipaa1/content/more.asp.

Family Educational Rights and Privacy Act (FERPA), www.ed.gov/policy/gen/guid/fpco/ferpa/index.html.

Federal Information Security Mgmt Act (FISMA), csrc.nist.gov/seccert/.

Gramm-Leach Bliley Act (GLBA), www.ftc.gov/privacy/glbact/.

Homeland Security Presidential Directive 7 (HSPD-7), www.whitehouse.gov/news/release/2003/12/20031217-5.html.

ISO 17799, www.iso.org (International Organization for Standardization's INFOSEC recommendations).

National Strategy to Secure Cyberspace, www.whitehouse.gov/pcipb/.

Sarbanes-Oxley Act (SOX), www.aicpa.org/sarbanes/index.asp.

Technology

The technology assessment involves the three elements: people, process, and technology. However, the technology portion of the assessment will probably take up 80% of your time due to the vast number of technological components involved in securing the infrastructure. Servers and hosts must be updated, patched, and secured. Applications must be updated, patched, and secured. The perimeter of your network must be secured, tested, and monitored. Remote access and wireless access must be secured, tested, and monitored. Data traveling across the network needs to be secured against a variety of attack types, which is done through various protocols at different network layers, depending on where the data originates, where it's headed, and what it contains. We'll spend the remainder of the chapter looking at the technology components of infrastructure security, holding off discussing the policies and procedures (which impact user behavior and the *people* aspect) until a later chapter.

Establishing Baselines

The point of performing these assessments is not to prove that your network is secure or insecure but to find out exactly what level of security you actually have and to establish baselines. When you know the starting point, you can improve security incrementally and document it as you go. Baselines are created by establishing a known starting point, in this case your current settings.

It might be tempting to correct problems as you perform this assessment, but it's not the best way to proceed. As you know, making a configuration change at Point A can cause a ripple effect through your network and show up at Point C in a strange and unexpected way. As you develop your project plan, be clear with your project team that they need to document existing configurations, settings, versions, and so on, without making changes. If a team member finds a serious security hole, it should be brought to your attention immediately for action. The point is that if a serious problem is found, it should be quickly addressed but not in an ad hoc manner. It should be assessed and addressed in a calm, rational, thoughtful manner, and possibly incorporated into your project plan. Does that mean that you wait until your project planning is complete to address a serious security hole? Absolutely not. You should, however, use a well thought out strategy for addressing it outside the project planning cycle, then document the changes and incorporate them into your project plan. What you want to avoid is having every person looking at the network making small tweaks here and there to "tighten up security" as they go, because you'll end up with a mess at the end of your evaluation period. Serious problems should be brought to your immediate, and minor issues should be well documented.

Addressing Risks to the Corporate Network

Once you have created a prioritized list of risks to your network as well as their associated costs, your next step will be to determine a course of action in handling each risk. When deciding how to address risks to your network, you typically have one of four options:

- **Avoidance** You can avoid a risk by changing the scope of the project so that the risk in question no longer applies, or change the features of the software to do the same. In most cases, this is not a viable option, since eliminating a network service such as e-mail to avoid risks from viruses would usually not be seen as an appropriate measure. (Network services exist for a reason; your job as a security professional is to make those services as secure as possible.) One example of how avoidance would be a useful risk management tactic is a case where a company has a single server that acts as both a Web server and a database server housing confidential personnel records, when there is no interaction whatsoever between the Web site and personnel information. In this scenario, purchasing a second server to house the employee database, removing the personnel database from the Web server entirely, and placing the employee database server on a private network segment with no contact to the Internet would be a way of avoiding Web-based attacks on personnel records, since this plan of action "removes" a feature of the Web server (the personnel files) entirely.
- **Transference** You can transfer a risk by moving the responsibility to a third party. The most well-known example of this solution is purchasing some type of insurance—let's say flood insurance—for the contents of your server room. Although the purchase of this insurance does not diminish the likelihood that a flood will occur in your server room, it does ensure that the monetary cost of the damage will be borne by the insurance company in return for your policy premiums. It's important to note that transference is not a 100-percent solution—in the flood example, your company will likely still incur some financial loss or decreased productivity in the time it takes you to restore your server room to working order. As with most risk management tactics, bringing the risk exposure down to zero is usually an unattainable goal.

- **Mitigation** This is what most IT professionals think of when implementing a risk management solution. Mitigation involves taking some positive action to reduce the likelihood that an attack will occur or to reduce the potential damage that would be caused by an attack, without removing the resource entirely, as is the case with avoidance. Patching servers, disabling unneeded services, and installing a firewall are some solutions that fall under the heading of risk mitigation.
- **Acceptance** After you have delineated all the risks to your infrastructure that can be avoided, transferred, or mitigated, you are still left with a certain amount of risk that you won't be able to reduce any further without seriously impacting your business (taking an e-mail server offline as a means to combat viruses, for example). Your final option is one of acceptance, where you decide that the residual risks to your network have reached an acceptable level, and you choose to monitor the network for any signs of new or increased risks that might require more action later.

There is no one right way to address all risks to your infrastructure; you'll most likely take a blended approach to security. There are some risks you absolutely need to avoid, other risks you can reasonably transfer or mitigate, and still others that you simply accept because the cost of avoiding them is just not worth it.

Business Intelligence…

Depth in Defense

Depth in defense is a key concept to understand before heading into an infrastructure security project. The concept is a fairly straightforward one: Security comes not from one source but from many layers of protection. Almost any attacker can find a way in through a single-defense

Continued

system, but it's much more difficult (but not impossible) to find a way in through a maze of security measures. When security measures are used in combination, it's like having a deadbolt, a padlock, a keypad, a card reader, and a biometric scanner attached to the network. An attacker can get through one or two, maybe even three, but it's the fourth and fifth layers that finally stop the would-be intruder and cause him or her to look for another, easier target. In the world of IT security, nothing is 100-percent secure unless it's powered off and locked in an isolated box, at which point it becomes completely useless. Understanding the depth-in-defense approach will help you as you try to evaluate the measures you should take to secure your network infrastructure. You may choose to implement something less iron-clad (at a drastically lower cost) in one area, knowing that the "layering" effect will likely give you a strong enough level of defense against most known threats.

External Environment

The external environment includes the changes in technology that might impact your business, the changes in the regulatory and legal environments that could impact your business, and the changing landscape of threats to your network. It's not a static picture; you'll need to implement policies and procedures that allow you and your IT staff to remain up to date with these changes so that you can continually monitor, assess, and address these changes in a proactive and positive manner.

We've talked about the legal implications of compliance and the importance of understanding those compliance issues when you're planning your IT security project. Because these issues are numerous, industry specific, and ever changing, we're not going to get into specific compliance data in this book. We recap some of the more common ones in this section, just in case you missed them earlier. We've also provided some Web links for you to learn more about these standards. There may be serious legal issues involved with compliance and noncompliance, so be sure to check with your firm's legal counsel to determine the regulations that apply to your firm. You might want to complete the internal assessment prior to contacting your attorney, so that you have a clear understanding of the kinds of information your network stores and the criticality of that information. For example, if your company recently started storing segments of people's medical records as part of a new business partnership with another firm, you will most likely have to comply with HIPAA standards (and possibly others). Recent changes to your company's business may have pulled you into areas in which regulation and compliance are mandatory, so be sure to do a full assessment here.

Threats

Predicting network threats and analyzing the risks they present to your infrastructure are among the cornerstones of the network security design process. Understanding the types of threats that your network will face helps you in designing appropriate countermeasures and in obtaining the necessary money and resources to create a secure network framework. Members of an organization's management structure will likely be resistant to spending money on a threat that they don't understand; this process will also help them understand the very real consequences of network threats so they can make informed decisions about the types of measures to implement. In this section, we discuss some common network attacks that you will likely face when you're designing a secure network and how each of these attacks can adversely affect your network.

When classifying network threats, many developers and security analysts have taken to using a model called STRIDE, which is an acronym for the following terms:

- **Spoofing identity** These include attacks that involve illegally accessing and using account information that isn't yours, such as shoulder-surfing someone's password while he types it on his keyboard. This type of attack affects the confidentiality of data.
- **Tampering with data** These attacks involve a malicious modification of data, interfering with the integrity of an organization's data. The most common of these is a man-in-the-middle (MITM) attack, where a third party intercepts communications between two legitimate hosts and tampers with the information

as it is sent back and forth. This is akin to sending an e-mail to Mary that says, "The meeting is at 3:00 P.M.", but a malicious attacker intercepts and changes the message to, "The meeting has been cancelled."

- **Repudiation** These threats occur when a user can perform a malicious action against a network resource and then deny that she did so, and the owners or administrators of the data have no way of proving otherwise. A repudiation threat can attack any portion of the confidentiality, integrity, and availability (CIA) triad.
- **Information disclosure** This occurs when information is made available to individuals who should not have access to it. Information disclosure can occur through improperly applied network permissions that allow a user to read a confidential file or give an intruder the ability to read data being transmitted between two networked computers. Information disclosure affects the confidentiality of your company's data and resources.
- **Denial of service** So-called DoS attacks do not attempt to alter a company's data; rather, they attack a network by denying access to valid users, by flooding a Web server with phony requests so that legitimate users cannot access it, for example. DoS attacks affect the availability of your organization's data and resources. A new variation is a distributed DoS (DDoS), also called a *zombie net* or *zombie attack*.
- **Elevation of privilege** This type of attack takes place when an unprivileged, nonadministrative user gains administrative or "root level" access to an entire system, usually through a flaw in the system software. When this occurs, an attacker has the ability to alter or even destroy any data that he finds, since he is acting with administrative privileges.

This type of threat affects all portions of the CIA triad, since the attacker can access, change, and remove any data that he or she sees fit. When you are analyzing a potential network threat, try to remember the

STRIDE acronym as a means of classifying and reacting to the threat. You can use the STRIDE model throughout the life of your corporate network when you're designing and maintaining security policies and procedures.

Recognizing External Threats

Now that we've discussed a model for classifying network threats, we can look at some of the common attacks in more detail. Entire books can be (and have been) written that solely discuss the kinds of threats that we look at in this section, so we'll be giving you a "birds-eye" view of the kinds of attacks that your network security design will need to guard against.

Denial-of-Service Attacks

As we've already mentioned, the DoS attack (and its first cousin, the DDoS attack) works to disrupt services on a network so that legitimate users cannot access resources they need. Some examples include attempts to disrupt the connection between two specific machines, or more commonly, attempts to flood an entire network with traffic, thereby overloading the network and preventing legitimate traffic from being transmitted. There can also be instances in which an illegitimate use of resources can result in denial of service. For example, if an intruder uses a vulnerability in your FTP server to upload and store illegal software, this can consume all available disk space on the FTP server and prevent legitimate users from storing their files. A DoS attack can effectively disable a single computer or an entire network.

A common venue of attack for DoS is against an organization's network bandwidth and connectivity; the attacker's goal is to prevent other machines from communicating due to the traffic flood. An example of this type of attack is the *SYN flood attack*. In a SYN flood, the attacker begins to establish a connection to the victim machine but in such a way that the connection is never completed. Since even the most powerful server has only a certain amount of memory and number processor cycles to devote to its workload, legitimate connection attempts can be denied while the victim machine is trying to complete these fake "half-open" connections.

Another common DoS is the so-called *Ping of Death,* where an attacker sends so many *PING* requests to a target machine that it is overloaded and unable to process legitimate network requests. An intruder might also attempt to consume network resources in other ways, including generating a massive number of e-mail messages, intentionally generating system errors that need to be included in Event Viewer logs, or misusing FTP directories or network shares to overload available disk space. Basically, anything that allows data, whether on a network cable or hard drive, to be written at will (without any type of control mechanism) can create a DoS when the attack has exhausted a system's finite resources.

Distributed Denial-of-Service Attacks

Distributed denial-of-service (DDoS) attacks are a relatively new development, made possible (and attractive to attackers) by the ever-expanding number of machines that are attached to the Internet. The first major wave of DDoS attacks on the Internet appeared in early 2000 and targeted such major e-commerce and news sites as Yahoo!, eBay, Amazon, Datek, and CNN. In each case, the Web sites belonging to these companies were unreachable for several hours at a time, causing a severe disruption to their online presence and effectiveness. Many more DDoS attacks have occurred since then, affecting networks and Web sites large and small.

WARNING

Most publicity surrounding DDoS attacks has focused on Web servers as a target, but remember that any computer attached to the Internet can fall victim to the effects of a DDoS attack. This can include everything from file servers or e-mail servers to your users' desktop workstations.

The DDoS attack begins with a human attacker using a small number of computers, called *masters*. The master computers use network scanners to find as many weakly secured computers as it can, and they use system vulnerabilities (usually well-known ones) to install a small script or a service (referred to in the UNIX world as a *daemon*) onto the insecure computer. This machine becomes a *zombie* and can now be triggered by the master computer to attack any computer or network attached to the Internet. Once the organizer of the DDoS attack has a sufficient number of zombie machines under control, he or she will use the "zombi-fied" machines to send a stream of packets to a designated target computer or network, called the *victim*. For most of these attacks, these packets are directed at the victim machine. The distributed nature of the DDoS attack makes it extremely difficult to track down the person or persons who began it; the actual attacks are coming from zombie machines, and the owners of these machines are often not even aware that their machines have been compromised. Making matters even more difficult, most network packets used in DDoS attacks use forged source addresses, which means that they are essentially lying about where the attack is coming from.

Viruses, Worms, and Trojan Horses

Viruses, Trojans, and worms are quite possibly the most disruptive of all security threats that we discuss in this section. These three types of threats, working alone or in combination, can alter or delete data files and executable programs on your network shares, flood e-mail servers and network connections with malicious traffic, and even create a "back door" into your systems that can allow a remote attacker to entirely take over control of a computer. You'll often hear these three terms used interchangeably, but each type of threat is slightly different. A *virus* is a piece of code that will alter an existing file and then use that alteration to recreate itself many times over. A *worm* simply makes copies of itself over and over again for the purpose of exhausting available system resources. A worm can target both hard drive space and processor cycles.

Business Intelligence...

Even Symantec Is Vulnerable

On May 24, 2006, a research company, eEye Digital Security, announced it had discovered a *high severity* security vulnerability in the Symantec antivirus program used by 200 million computers worldwide. The vulnerability was characterized as severe because it didn't require any user interaction to be exploited, making it highly susceptible to worm attacks. The irony is, of course, that this vulnerability was discovered not a week after the CEO of Symantec slammed Microsoft's "security monoculture" as a source of vulnerability. Since no one single product or defense will provide adequate security in today's threat environment, this finding underscores the need for depth in defense. And, as Symantec's CEO discovered, it also underscores the danger of tossing rocks at security "glass houses."

Software Vulnerabilities

Some network attacks target vulnerabilities in the way that a software application or entire operating system has been programmed. For example, a buffer overflow attack occurs when a malicious user sends more data to a program than it knows how to handle. For example, we've all seen Web forms that ask you to fill in personal information: first name, last name, telephone number, and so forth. A careless developer might program the "First Name" field to only be able to handle 10 characters; that is, a name that is 10 letters long. If the Web application does not check for buffer overflows, an attacker can input a long string of gibberish into the First Name field in an attempt to cause a buffer overflow error. At this point, the attacker could even embed the name of an executable file into that long string of text and actually pass commands to the system as if he or she were sitting at the server console itself. A similar software vulnerability is a format string vulnerability that would allow an attacker to insert random data into a file or database, including malicious code that can be executed against the server as though the attacker were sitting right in front of the keyboard.

Another attack that is specifically common to Web and FTP servers is a *directory traversal vulnerability*. This type of vulnerability allows a user to gain access to a directory on a server that he hasn't been specifically given permissions to, by virtue of having permissions to a parent or child directory. Say that someone goes to the URL www.airplanes.com/biplanes/cessna/model1.html. He decides to manually change this URL (in other words, not following an <HREF> link on the site itself) to www.airplanes.com/biplanes/piper, to see if the directory structure holds any information there. If the Web site hasn't been properly patched and configured with the correct security settings, the user might find that he now has access to every single file in the piper/ directory. Even worse, he can once again execute a command from the Web browser by changing the URL to something like www.airplanes.com/biplanes/piper/del%20*.*. (*%20* is used in HTML to represent a space, so that command would read *del* *.* on a regular command line.)

Another common attack also occurred in NetMeeting and Windows Media Player some time ago, where an attacker could insert special characters during a file transfer that would allow him to browse an unsuspecting user's hard drive directory structure.

Unfortunately, the breadth and depth of software vulnerabilities grows almost daily due to the wonderfully wide variety of applications available on the market. This variety provides new and useful functionality to users, but it obviously can create headaches for IT staff just trying to keep up.

Nontechnical Attacks

A final category of attack that we'll discuss here are those that use less technical means to circumvent network security. *Social engineering attacks* rely on an unsuspecting user's lack of security consciousness. In some cases, the attacker will rely on someone's goodwill, using a tactic like, "I've really got to get this done and I don't have access to these files. Can you help me?" (This works because most of us, at heart, really want to be helpful to those around us.) Other social engineering attacks use a more threat-based approach, insisting that the attacker is the secretary for Mr. Big-Shot VP who needs his password reset right away and heaven help you if you keep him waiting. This method relies on the assumption that a show of authority will cause someone without adequate training to bypass security procedures, to keep the "big-shot important user/client" happy. Since social engineering attacks are nontechnical in nature, the measures required to defend against them are more administrative than anything else. It's critical to have well-understood security policies in place that apply to everyone, regardless of their position in your company. This will assist in preventing an attacker from circumventing security procedures because a help desk or other staff member is unaware of them. We discuss user education and awareness campaigns later in this book.

Top 20 Threats

The SANS organization publishes and maintains a top-20 list of network threats. You might want to refer to this list as you're developing your infrastructure security plan; it will give you excellent insight into the latest threats and how to address them. For the most up-to-date list, visit www.sans.org/top20/#threatindex. The current Top Vulnerabilities in Windows Systems list contains the following categories:

- W1. Windows Services
- W2. Internet Explorer
- W3. Windows Libraries
- W4. Microsoft Office and Outlook Express
- W5. Windows Configuration Weaknesses

The Top Vulnerabilities in Cross-Platform Applications list:

- C1. Backup Software
- C2. Antivirus Software
- C3. PHP-based Applications

- C4. Database Software
- C5. File Sharing Applications
- C6. DNS Software
- C7. Media Players
- C8. Instant Messaging Applications
- C9. Mozilla and Firefox Browsers
- C10. Other Cross-platform Applications

The Top Vulnerabilities in UNIX Systems list:

- U1. UNIX Configuration Weaknesses
- U2. Mac OS X

The Top Vulnerabilities in Networking Products list:

- N1. Cisco IOS and non-IOS Products
- N2. Juniper, CheckPoint and Symantec Products
- N3. Cisco Devices Configuration Weaknesses

Later in this chapter, we'll look at the major threats and vulnerabilities so you can build these into your project plan, as appropriate, based on your own unique network configuration.

Business Intelligence…

Hackers Turn to Security Software

An article in the *Washington Post* in late 2005 highlighted a new and growing trend among hackers: the new focus on security software used by millions of end users. In the "old days," hackers focused on attacking operating systems and exploiting known vulnerabilities. Although that still occurs, the new threat front is in the very software you rely on to secure your computer from the bad guys. As hackers look for and exploit

Continued

these vulnerabilities, they expose users to a whole new realm of risk. Operating systems such as Windows and Linux are now regularly updated and patched, but security software programs typically were only updating virus signature files, not the program itself. Now security software program makers are finding their products under attack and are having to respond as operating system companies once did.

For more information and to read the whole article, head to this URL: www.washingtonpost.com/wp-dyn/content/article/2005/11/21/AR2005112101424.html.

Network Security Checklist

This section is a lengthy one and is intended to provide you with a thorough review of the types of things you should review, assess, and think about when you prepare your infrastructure security project plan. Even though we've created a detailed list, there's always a chance there are additional elements your plan will need. Certainly, there's also a strong likelihood that there are things in these checklists that you don't have and don't need. That's okay. The point is to try to help you think through all the details you possibly can about your network infrastructure, to ensure that you are thorough and don't leave any stone unturned. At the end of this process, you may decide not to address some aspects of infrastructure security, or you might choose to work on some of these items in a Phase 2 or Phase 3 project plan. This should give you a great start in thinking all this through.

We've divided the infrastructure project into four main areas, though you may choose to parse it out differently. We'll look at devices and media and ways to secure network devices (excluding servers and user computers) and the network media. Media could mean secure network area storage devices (NAS), backup media, or other storage devices. The "Topologies" section includes how you segment the network for security, including creating DMZs and implementing firewalls, and how you secure network traffic. Intrusion detection and prevention systems are pretty popular these days (for good reason), so we'll look at best practices for implementing IDS/IPS that you can utilize in your project plan.

Finally, we'll look at system hardening, including hardening infrastructure servers (DNS, DHCP, and so on), application and database servers, and other computers on the network. Keep in mind that this is not a "how to" as much as it is a list of things to consider and include in your project plan. There are volumes filled with information on these topics; it would be far outside the scope of this book to talk about how you do these things. Our intent is to provide a framework and a solid starting point for your infrastructure security project-planning process. If you're not sure what some of these things are or if you're uncertain as to how to address these issues, you'll need to do further research on these topics.

Devices and Media

Network devices typically include routers, switches, firewalls, and other communication devices. We cover these items extensively at the end of this section (we placed it there because it's a long, wide-ranging list). The short story is that routers, switches, and other communication devices should be:

1. **Physically secured** Place devices in a locked cabinet, locked room, and locked building, where possible. Where that's not possible, devices should be closely monitored or access should be controlled or limited.
2. **Physically inspected** Remove extra cables, disable external ports, and disconnect unused connections.
3. **Hardened** Remove unused software, disable unused ports, stop or uninstall unused protocols and services, disable unused functionality, remove unused user accounts, change default settings, use strong passwords, and remove or limit all but one administrative account.
4. **Monitored** Audit, log, and monitor all access to devices, both physical and logical; monitor all successful logons; monitor all failed logons; review log files frequently; and store configuration data in a safe, secure location.

5. **Encrypted** Encrypt sensitive data files; encrypt and secure all removable media; create a secure system for handling removable media, including backup files; create a log file to track media handling; secure removable media in locked, access-controlled location; and store archives in a secure, off-site location.

Topologies

Network infrastructure security:

1. Create secure boundaries using firewalls, DMZs, and proxy servers.
2. Create secure remote access.
3. Create secure wireless access.
4. Implement a segmented network.
5. Implement network traffic security protocols for sensitive network traffic.
6. Deploy network security technologies.

- Use Encrypting File System (EFS) or similar file encryption.
- Require and use strong user authentication, passwords and account policies.
- Employ the concept of "least privileges" when assigning user rights.

Security infrastructure components include routers, proxy servers, firewalls, and DMZs. Firewalls are pretty straightforward and can be implemented as hardware or software solutions. Let's take a side street and take a quick look at DMZs.

Demilitarized zones, or *DMZs*, are isolated network segments that typically sit between the Internet and your network, whether in front of or behind your firewall (or between two firewalls). There are many different ways to set up a DMZ; again, it's outside the scope of this book to discuss the design, implementation, and configuration of a DMZ. However, it might be helpful to discuss a few highlights of DMZ design that might help as you look at implementing or tightening a DMZ for your network.

Designing DMZs

DMZ design, like security design, is always a work in progress. As in security planning and analysis, we find DMZ design carries great flexibility and change potential to keep the protection levels we put in place in an effective state. The ongoing work is required so that the system's security is always as high as we can make it within the constraints of time and budget, while still allowing appropriate users and visitors to access the information and services we provide. You will find that the time and funds spent in the design process and preparation for the implementation are very good investments if the process is focused and effective; this will lead to a high level of success and a good level of protection for your network.

In this section of the chapter, we explore the fundamentals of the design process. We incorporate the information we discussed in relation to security and traffic flow to make decisions about how our initial design should look. Additionally, we'll build on that information and review some other areas of concern that could affect the way you design your DMZ structure.

Design of the DMZ is critically important to the overall protection of your internal network—and the success of your firewall and DMZ deployment. The DMZ design can incorporate sections that isolate incoming VPN traffic, Web traffic, partner connections, employee connections, and public access to information provided by your organization. Design of the DMZ structure throughout the organization can protect internal resources from internal attack. As we discussed in the security section, it has been well documented that much of the risk of data loss, corruption, and breach actually exists inside the network perimeter. Our tendency is to protect assets from external harm but to disregard the dangers that come from our own internal equipment, policies, and employees.

These attacks or disruptions do not arise solely from disgruntled employees. Many of the most damaging conditions occur because of inadvertent mistakes made by well-intentioned employees. Each of these entry points is a potential source of loss for your organization and ultimately can provide an attack point to defeat your other defenses. Additionally, the design of your DMZ will allow you to implement a multilayered approach to securing your resources that does not leave a single point of failure in your plan. This minimizes the problems and loss of protection that can occur because of misconfiguration of rule sets or ACL lists, as well as reducing the problems that can occur due to hardware configuration errors.

Remote Access

Remote access is granted in a number of different ways, so the way it should be secured varies widely. The basics are that the remote access servers should be physically secured (as should all infrastructure servers) in an access-controlled location. The number of accounts that are authorized to log onto the server for administrative purposes should be limited and audited. The communication link between the RAS and the remote users should be secured, as should the data on that link, if needed. The network traffic security methods include signing, encryption, and tunneling. The level of these methods is determined by the system with the least capabilities. Older operating systems cannot utilize the latest encryption technologies, for example, so you might include policies that require that remotely connecting users use the latest version of Windows XP Professional, to enable the entire end-to-end communication link to use the strongest available encryption. You can also require strong authentication across remote links. Different operating systems implement this differently; in Windows Server 2003, for example, it's implemented through policies set in Administrative Tools | Routing and Remote Access.

Wireless Access

We've devoted a whole chapter to wireless security, so we will only discuss the top-level items here:

- Change access point default settings.
- Disable SSID broadcasting; create a closed system (does not respond to clients with "Any" SSID assigned).
- Transmission power control (limiting the amount of power used for transmission to control the signal range).
- Enable MAC address filtering.
- Enable WEP or WPA.
- Filter protocols.
- Define IP allocations for the WLAN.
- Use VPNs.
- Secure users'computers.

All these choices have pros and cons, distinct advantages and disadvantages; you'll need to decide the right approach for your organization. As with all things in IT security, it's important that you understand the result of the solutions you're using, understand the configuration and maintenance of these elements, and be sure you test them well in a lab or isolated setting before implementing them across the enterprise.

Intrusion Detection Systems/ Intrusion Prevention Systems (IDS/IPS)

First, let's define IDS and IPS, because they're not one and the same. *Intrusion detection systems* (IDS) are passive in nature; they let you know an intrusion is taking place or has occurred. They do nothing to stop an intrusion. On the other hand, an *intrusion prevention system* (IPS) is an active system that works to stop an intrusion or to prevent one when "it thinks" one is occurring. How does "it" think? It does so based on how you configure it, so we end up back at that persistent *people* problem we've mentioned once or twice. An IPS has one major drawback, and that is the high likelihood of false positives. Depending on how you configure the IPS, the results of a response to a false positive might be far more devastating than an actual intrusion, so you're walking a fine line with IPS. That said, some excellent hardware and software solutions are available on the market today, many of which are a great improvement over IDS/IPS systems of the past. It is far outside the scope of this book to discuss the pros and cons, the highlights and lowlights of these systems, so we're not going there. However, we will mention a few different ways you can implement and secure your IDS/IPS systems and leave it up to you to develop a specific plan for implementing these systems, since they are so varied.

A word of caution: IDS/IPS is not a standalone defense. You should implement it with the understanding that it contributes to your depth of defense, but alone it will not keep your network safe. It's a great tool to have in your security toolkit, but it's not the magic bullet everyone wishes they had.

IPSs introduce fundamental performance and stability issues within the network or system they are designed to protect. The act of implementing automatic controls in response to detecting attacks does not come without a price. For example, an inline network IPS will not forward packets before inspecting Application-layer data. This inspection takes time and can result in a slowdown in the responsiveness and throughput of the local network. A host IPS that has been charged with the inspection and validation of an application's system calls can impact a kernel's ability to quickly service system calls, which may only be 1 to 15 percent but is probably noticeable.

Network Active Response System

A *network active response system* has the ability to interact with network traffic indirectly through the modification of firewall policies and router Access Control Lists (ACLs). They also have the ability to take down switch ports (for locally generated attacks) and to spoof error code packets such as Transmission Control Protocol (TCP), RST (Reset), or Internet Control Message Protocol (ICMP) unreachable packets. Such an active response system is commonly implemented directly within a network IDS, where it can easily take advantage of its detection capabilities.

This is useful for tearing down individual sessions or for trying to convince an attacking host that the target is unreachable due to ICMP errors. However, there is not usually much time between these measures and the goal of the attack. It's unclear whether the countermeasure will be successful.

There are four classes of countermeasure that a network IPS can utilize to thwart a network-based attack. Each class applies to one layer of the protocol stack, beginning at the Data Link layer:

- **Data Link layer countermeasures** Administratively shut down a switch port interface associated with a system from which attacks are being launched. This approach is feasible only for attacks that are generated from a local system. Having the ability to timeout the downed switch port is important, since the port probably should not be shut down indefinitely.
- **Network layer countermeasures** Interact with the external firewall or router to add a general rule to block all communication from individual IP addresses or entire networks. An inline IPS can accomplish the same thing without having to appeal to an external device, since packets from specific IP addresses can simply be blocked after an attack has been detected. Similarly to Data Link layer responses, timeouts are important at the Network layer, since the firewall rule set or router ACL modifications should be removed after a configurable amount of time.
- **Transport layer countermeasures** Generate TCP RST packets to tear down malicious TCP sessions, or issues any of several available ICMP error-code packets in response to malicious UDP traffic. (Note that ICMP is strictly a Network layer protocol and is the standard method of communicating various errors to clients that utilize UDP). Timeouts are not applicable here, because countermeasures are leveraged against an attacker on a per-session or per-packet basis.

- **Application layer countermeasures** Alter malicious Application layer data so as to render it harmless before it reaches the target system. This countermeasure requires that the IPS be in line in the communication path. Any previously calculated Transport layer checksum must be recalculated. Similarly to the Network layer, timeouts are not applicable here, since the effects of replacing Application layer data are transitory and do not linger once an altered packet is forwarded through the IPS.

Later in this chapter, we'll walk through a number of "generic" countermeasures and hardening tasks related to these layers when we look at various ways routers, switches, and other network devices can be hardened in conjunction with whatever IDS/IPS system you implement.

Host Active Response System

A *host active response system* is usually implemented in software and is deployed directly on a host system. Once a suspicious event has been detected on a host (through any number of means, such as log file analysis, detection of specific files or registry keys associated with known exploits, or a suspicious server running on a high port), a host active response system is charged with taking an action. As with network active response, the expectation for a host active response system is that countermeasures will not necessarily prevent an attack from initially being successful. The emphasis is on trying to mitigate the effects and damage caused by an attack after detection. After an attack is detected, automated responses can include alteration of file system permissions, changes in access that a system grants to users, automated removal of worms or viruses (anti-virus), and additions of new rules to a local firewall subsystem.

Before we move into system hardening, let's take a look at how IDS/IPS systems are implemented in the network infrastructure. Figure 11.2 shows the IDS system as part of the infrastructure. The IDS server, in this case, would be connected to a span port so that it would monitor all traffic on the local network. The IDS system is capable of spoofing a TCP RST or ICMP error code packet to thwart the attack but would not be effective against single-packet attacks.

Figure 11.2 IDS System Placement in Infrastructure

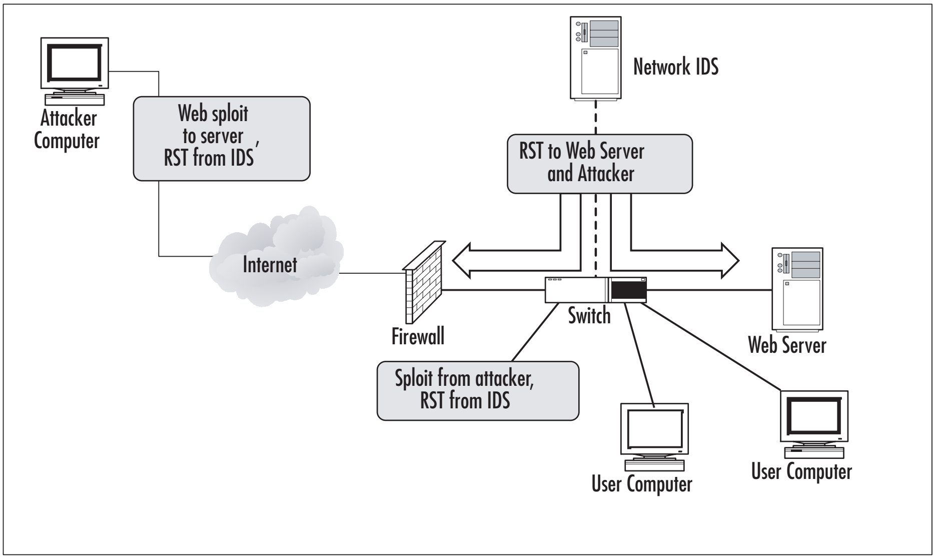
Network IDS
Attacker Computer
Web sploit to server , RST from IDS
RST to Web Server and Attacker
Internet
Firewall
Switch
Web Server
Sploit from attacker, RST from IDS
User Computer
User Computer

An inline system performs a bit differently, as shown in Figure 11.3. In this case, the inline system captures the sploit (short for *exploit*) and modifies it to protect the local network. A typical deployment of the IPS occurs just inside the firewall. In this position, it captures all incoming traffic before it goes to the local network, providing ubiquitous protection, even for single-packet attacks. Because all traffic flows through an inline IPS, downsides such as false positives and slower response times must be factored in.

Next Generation Security Devices

As you look at your current implementation of IDS or IPS (or if you're considering an implementation), you should also keep an eye on recent developments in the world of security devices. *Network processors* can be deployed in various architectures including parallel, where each processor handles 1/*N* of the total load or pipeline, where, as a packet moves through the pipeline, each processor typically handles a single specific repetitive task. The network processor was originally targeted to the routing market, but it is easy to see how it can be applied to the increased demands of packet inspection in network security. For example, one processor could handle the pattern matching for known worm signatures, another could analyze for protocol standards compliance, and yet another could look for protocol or usage anomalies. The network processor would have direct access to fast memory that stores policies and signatures, whereas slower, larger memory would store state information and heuristics information. New attacks could be mitigated by adding new code to the network processor. A separate processor can handle management functions such as logging and policy management. Network processors also offer the ability to scale, much like CPUs on computer systems.

Figure 11.3 IPS System Inline Placement in Infrastructure

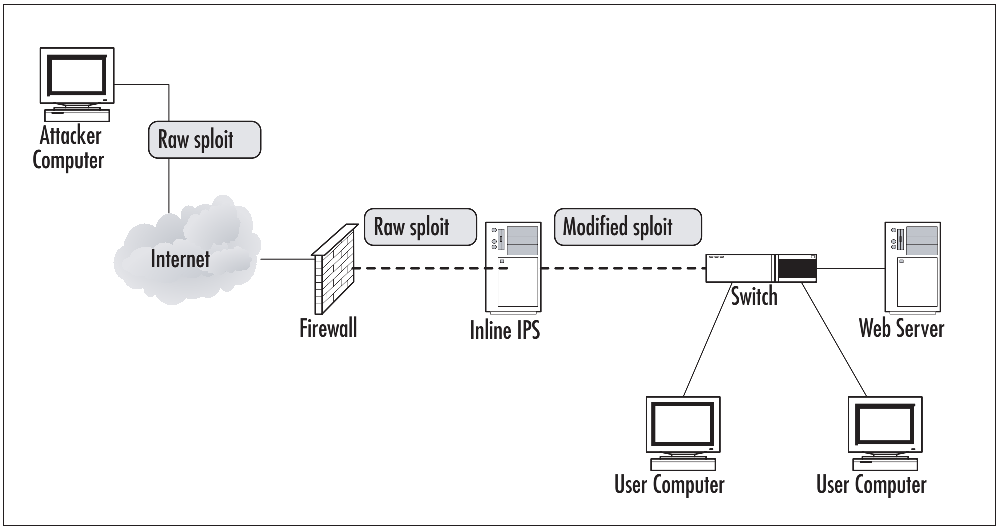

Business Intelligence…

Intrusion Prevention and Detection Resource

At the risk of sounding a bit self-serving, if you have any desire to understand more about IDS/IPS, you really should check out another Syngress book. There may be other excellent IDS/IPS resources out there, but *Intrusion Detection and Active Response: Deploying Network and Host IPS*, by Michael Rash, Angela Orebaugh, Graham Clark, Becky Pinkard, and Jake Babbin, with a foreword by Stephen Northcutt (Syngress Publishing, Inc., 2005), is a great resource. If you're like most IT professionals, you're inundated with technical information on a daily (okay, hourly) basis and it's hard to stay up to date on every topic in the computer world. This book provides excellent background information and helps you understand the wild world of IDS/IPS so you can make informed decisions about how, when, and where to implement it in your organization. If you're looking for an excellent resource on this topic, do yourself a favor and check out this one-stop-shopping trip for an excellent IDS/IPS education.

System Hardening

Server security:

1. Always control physical and network access to critical servers, especially domain controllers, DNS servers, DHCP servers, and other infrastructure servers. Keep infrastructure servers in an access-controlled location.
2. Always perform tasks on the servers with the least possible privileges. Do not perform tasks with Administrator privileges, if possible. Use the *Run As* command (or equivalent) when needed.
3. Restrict user and machine access to groups that have loose security settings. Provide users and computers with the least possible

permissions while still meeting their needs to access and use network resources.

4. Secure the data on the computers using strong ACLs and, if needed, the *syskey* utility. The *syskey* utility provides protection against password-cracking software that targets the Security Access Management (SAM) database or directory services. It uses strong encryption that is much more difficult (if not close to impossible) and time consuming to crack.
5. Require the use of strong passwords via password policy settings.
6. Restrict the downloading and installation of programs that do not come from known, trusted sources.
7. Maintain up-to-date virus protection on all systems.
8. Keep all software patches up to date. Patches often address newly discovered security holes. Applying patches in a timely manner on all affected machines can prevent problems that are easily avoided.
9. Deploy server, application and client-side security technologies:
   - Secure server traffic traveling on the network.
   - Secure application and user data traveling on the network.
   - Secure network access points and network access.
   - Secure client devices including desktops, laptops, and PDAs.
   - Implement automatically updating virus and spyware protection systems.

Other Infrastructure Issues

1. Deploy network monitoring and auditing.
2. Develop a disaster recovery plan that includes creating backups, documenting recovery options and using repair and recovery tools. (See Chapter 13 for more on disaster planning and recovery.)

3. Develop standard operating procedures that include strong monitoring, auditing, and documentation.

Business Intelligence…

Rootkits

There's been a lot of news in the recent past about the problems presented by rootkit attacks. As you're well aware, those little pieces of malware reside so deep in the system that you can't possibly remove them without completely starting from scratch. After a system is compromised, all the affected software must be reinstalled from known "clean" sources. Since it can be difficult to determine precisely which pieces of software have been affected, the best way to guarantee security is to reinstall the entire operating system (OS) and all applications. OS kernels can also be compromised (see www.rootkit.com), and when they are, nothing on the system (even the most basic file system, memory, and network status information) can be trusted. An after-the-fact forensic analysis of the file system may turn up useful information if the disk is mounted underneath an uncompromised OS, but this is a time-consuming operation.

Other Network Components: Routers, Switches, RAS, NMS, IDS

There are numerous components that should be checked during an infrastructure security project. The list in this section was compiled, in part, from a network checklist developed by the Defense Information Systems Agency (DISA) for the Department of Defense (DoD). Although not all items listed will apply to your network and it's possible that not all items that apply to your network appear on this list, this is an extensive list that you can use as the starting point for your own checklist. Some of the items in this list contain brief explanations included to help you understand their importance. Our assumption is that you're familiar with the ins and outs of network security, but there are a few places where a quick clarification will help, and we've included them as well. These are written in language that reflects problems you would find that should be remedied (for instance, highlighting the problem you're looking for, not necessarily the solution you should implement). The list is organized by device type, beginning with routers and other network devices and moving on to firewalls, VLANs, RAS servers, and so on.

Network

- **Network infrastructure is not properly documented** You should begin with a clear understanding of how your network infrastructure is currently configured. This should be well documented and kept up to date.
- **Network connections exist without approval** All network connections should exist only with explicit approval or knowledge of the IT department. This is typically a problem with modems, wireless access points, and USB-type network devices.
- **Unmanaged backdoor connections, backdoor network connections bypass perimeter** Every network in the world has a variety of backdoor connections that network administrators use (or that software developers build in). When unmanaged, these connections create security problems for your network infrastructure. These are especially problematic when these backdoors bypass perimeter security systems. If you can use them, so can the bad guys.
- **Circuit location is not secure** The location of network circuitry, including the backbone and other highly critical components, should be secured physically.
- **Network devices are not stored in secure communications room** This is part of physical security; to the extent possible, network devices should be stored in a secure communications room. This should certainly be true for mission-

critical devices. Physical security of the company's premises, coupled with physical security of key network devices, is part of a depth-in-defense strategy.

- **Minimum operating system release level** All network devices—from desktop computers to servers to firewalls to routers—should have the latest updates and patches for the operating system they are running. As seen from the top-20 threat list, many are threats to portions of the operating system, so all device operating systems should be kept up to date. Where possible, you may also choose to upgrade the operating system itself to a newer, more secure version, where appropriate. This OS release-level maintenance should also apply to routers and other devices that have operating systems, firmware, or other embedded software functionality.
- **DNS servers must be defined for client resolver** If a router or similar network device is specified as a client resolver (resolves DNS to IP address), the router should have a DNS server defined. If the DNS server is specified, it makes it more difficult for an attacker to substitute his or her IP address for that of the destination host. If this type of man-in-the-middle attack is successful, the unsuspecting host user could transmit sensitive information, including logon, authentication, and password data, to the attacker.

External Communications (also see "Remote Access")

- **Modems are not disconnected** In Chapter 12 ("Wireless Security"), we discuss the problem with unsecured modems. Briefly, these can be attacked by wardialers who simply look for modems connected to corporate networks. These can create significant security holes and are often overlooked in our quest to lock down the wired network.

- **An ISP connection exists without written approval** In most companies, this might be a difficult trick to achieve, but it certainly warrants examination to ensure that the ISP connection(s) is managed by the IT department and not some errant user who managed to get the local ISP provider to run a cable into the office on a Saturday morning.
- **Communications devices are not password protected** This seems like a giant "Duh!" but you'd probably be surprised how often communication devices such as modems, routers, switches, and other "smart" devices are left unprotected by even a simple password or that use the default password that came with the device out of the box.
- **No warning banner** Failure to display the required login banner prior to logon attempts will limit the site's ability to prosecute unauthorized access. It also presents the potential for criminal and civil liability for systems administrators and information systems managers. Not displaying the proper banner will also hamper the site's ability to monitor device usage. Displaying a banner warning users of the consequences of unauthorized access helps warn off the bad guys and draws a line in the legal sand that you might need later.

TCP/IP (Some TCP/IP Information Also Found in the "Routers" Section)

- **LAN addresses are not protected from the public** In later versions of the Windows operating system, even home users were able to easily implement Network Address Translation (NAT) to protect internal IP addresses from Internet users. Most businesses these days have implemented some method of protecting internal IP addresses so that hackers can't use this information to decipher the network structure and plan an attack.

- **The DHCP server is not configured to log hostnames** To identify and combat IP address spoofing, it is highly recommended that the DHCP server log MAC addresses or hostnames on the DHCP server.
- **TCP and UDP small server services are not disabled** TCP and UDP services are often available on network devices, including routers and servers. Disabling these services if they're not used helps reduce the attack footprint. TCP and UDP protocols include services that routers can support; however, they are not required for operation. Attackers have used these services to cause network DoS attacks.
- **TCP keepalives for Telnet session must be enabled** Enabling TCP keepalives on incoming connections can help guard against both malicious attacks and orphaned sessions caused by remote system crashes. Enabling the TCP keepalives causes the router to generate periodic keepalive messages, letting it detect and drop broken Telnet connections.
- **Identification support is enabled** Identification support allows you to query a TCP port for identification. This feature enables an unsecured protocol to report the identity of a client initiating a TCP connection and a host responding to the connection. With identification support, you can connect a TCP port on a host, issue a simple text string to request information, and receive a simple text-string reply. This is another mechanism to learn the router vendor, model number, and software version being run. Identification support should be disabled on routers and other network devices that provide this functionality.

Business Intelligence…

Whitelisting

Whitelisting is the ability to easily specify IP addresses or networks that should never be the subject of an automated response in an IDS/IPS system. For example, IP addresses associated with systems that are critical to a network (for example, the Domain Name Server, or DNS, or upstream router) should not be automatically blocked by an active response system, nor should sessions be altered by an inline IPS. Some active response systems include the ability to whitelist IP addresses and networks and to specify which protocols should be ignored. For example, if a DNS server sends an attack across the network to a Web server, it may be permissible for an active response system to capture the individual TCP session on port 80 but ignore everything else.

- **IP-directed broadcasts are not disabled** An *IP-directed broadcast* is a datagram sent to the broadcast address of a subnet that is not directly attached to the sending machine.The directed broadcast is routed through the network as a Unicast packet until it arrives at the target subnet, where it is converted into a link layer broadcast. Due to the nature of the IP addressing architecture, only the last router in the chain, which is connected directly to the target subnet, can conclusively identify a directed broadcast. IP-directed broadcasts are used in the extremely common and popular *smurf*, or DoS, attacks. In a smurf attack, the attacker sends ICMP echo requests from a falsified source address to a directed broadcast address, causing all the hosts on the target subnet to send replies to the falsified source. By sending a continuous stream of such requests, the attacker can create a much larger stream of replies, which can completely inundate the host whose address is being falsified.This service should be disabled on all interfaces when it's not needed to prevent smurf and DoS attacks.

- **Ingress filtering inbound spoofing addresses** Inbound spoofing occurs when someone outside the network uses an internal IP address to gain access to systems or devices on the internal network. If the intruder is successful, they can intercept data, passwords, and the like and use that information to perform destructive acts on network devices or network data.
- **Egress outbound spoofing filter** You should restrict the router from accepting any outbound IP packet that contains an illegitimate address in the source address field via egress ACLs or by enabling Unicast Reverse Path Forwarding. ACLs are the first line of defense in a layered security approach. They permit authorized packets and deny unauthorized packets based on port or service type. They enhance the network's posture by not allowing packets to even reach a potential target within the security domain. Auditing packets attempting to penetrate the network but that are stopped by an ACL will allow network administrators to broaden their protective ring and more tightly define the scope of operation.

Administration

- **Devices exist that have standard default passwords** This is another major "Duh!" item; again, it's surprising how easy it is to get into a large number of devices just by using the default password that the device shipped with. Want to know the default password? Go up on the manufacturer's Web site, look for the user guide for the specific device, and the default password is almost guaranteed to be listed in the first five pages of the manual.
- **Group accounts or user accounts without passwords** Without passwords on user accounts for network devices, one

level of complexity is removed from gaining access to the routers. If a default user ID has not been changed or is guessed by an attacker, the network could be easily compromised, since the only remaining step would be to crack the password. Sharing group accounts on any network device should also be prohibited. If these group accounts are not changed when someone leaves the group, that person could possibly gain control of the device. Having group accounts does not allow for proper auditing of who is accessing or changing the network. Only allow individual user account access and require each user to have a unique user ID and a strong password.

- **Assign lowest privilege level to user accounts** Across the enterprise, you should always assign the least privilege possible for all users. This prevents users from getting into places they shouldn't, and it also prevents hackers from upgrading their privileges if they manage to get in on a user account that has too many privileges. Even IT staff should have user accounts with least privileges for most day-to-day network tasks, and they should only log on with administrative privileges when needed. Network outages and security holes can be created by users with too many permissions or even by a well-meaning but inexperienced net admin.
- **Strong password policies are not enforced** Strong passwords is an inadequate defense on its own, but it slows down a would-be intruder and can also alert a net admin to a potential problem if failed password attempts are monitored and accounts are locked down after too many failed attempts. Requiring users to use strong passwords, to change them periodically, and to prevent them from repeating old passwords too frequently are all parts of strong password policy. In addition, you can audit failed attempts, notify a net admin of too many failed attempts, and lock out an account with too many failed accounts as part of your strong password policy implementation.

- **Passwords are not recorded and stored properly** User passwords should not be recorded and stored, but certain administrative ones absolutely should be. You can probably think of several scenarios where someone who doesn't normally require administrative access requires it. For example, suppose as part of your disaster recovery plan, you have an executive VP who is responsible for coordinating recovery efforts. He or she should have access to these passwords only for these emergency situations, because on a day-to-day basis, you operate on the principle of "least access" and the EVP really has nothing more than the equivalent rights of a power user. Having these passwords on a network server in plain sight or in a paper file someplace obvious is not a good idea. Making sure these emergency passwords are recorded and stored properly ensures security for the network on a day-to-day basis but provides an important fail-safe option in emergencies as well.
- **Passwords are viewable when displaying the router or other device** Many attacks on computer systems are launched from within the network by unsatisfied or disgruntled employees. It's vital that all router passwords be encrypted so they cannot be intercepted by viewing the console. If the router network is compromised, large parts of the network could be incapacitated with just a few simple commands.
- **Passwords are transmitted in clear text** There are many types of situations in which passwords are transmitted in clear text. This creates an opportunity for an attacker to seize passwords. Review how and where passwords are transmitted and secure the communication lines if the passwords themselves are transmitted in clear text.
- **Emergency accounts should be limited to one** Emergency accounts on devices such as routers or switches should be limited to one. Authentication for administrative access to the router should obviously be required at all times. A single account can be created on the router's local database for use in an emergency,

such as when the authentication server is down or connectivity between the router and the authentication server is not operable. Verify that there is one and only one emergency account to prevent unnecessary opportunities for attack.

- **Unnecessary or unauthorized router or device accounts exist** This point is related to the previous item. You should eliminate any unused, unnecessary, or unauthorized device accounts except for one authorized emergency account.
- **Disable unused ports and services** On every server, every firewall, and every device, disable unused ports and services. Microsoft took a giant leap forward in the more recent versions of the Windows operating system when the company changed the default configuration from "open" to "closed." This meant that the net admin had to consciously enable and open services and ports after installation. Earlier versions came open and unlocked out of the box, and the net admin had to sift through the system to lock it down. For all devices, disable unused ports and services, uninstall unused applications, and remove unused hardware.
- **Auditing and logging files are not set to record *denied* events, not set to record system activity** Auditing and logging are key components of any security architecture. It is essential that security personnel know what is being done, being attempted, and by whom in order to compile an accurate risk assessment. Auditing the actions, particularly *denied* events, on routers provides a means to identify potential attacks or threats. Maintaining an audit trail of system activity logs (*syslog*) can help you identify configuration errors, understand past intrusions, troubleshoot service disruptions, and react to probes and scans of the network.

- **Configurations are stored in unsecured locations** To ensure network and data availability, the configuration data of key network infrastructure components should be maintained in a secure, offsite location. This is part of good disaster recovery planning practices and adds to security if these configurations are stored in secured locations offsite rather than in an unlocked file cabinet in the mailroom. Access to these configuration files should be restricted and logged to prevent unauthorized access.

Network Management

- **Out-of-band network management not implemented or required** It's outside the scope of this chapter (and book) to get into a deep discussion of in-band and out-of-band network management, but we will toss out a couple of quick explanations before discussing the infrastructure security implications of both. In-band network management uses the same network infrastructure as the devices and data being managed. Most networking equipment basically sends out IP traffic for network management on the same medium as the traffic it's managing (routers, switches, and so forth). Out-of-band network management uses a separate connection, often a serial RS-232 port, instead of the network port used for in-band management. There are security pros and cons to both, so the key is to secure whichever method(s) you implement.

  Without secure out-of-band management implemented with authenticated access controls, strong two-factor authentication, encryption of the management session, and audit logs, unauthorized users may gain access to network managed devices such as routers or communications servers (CS). If the router network is compromised, large parts of the network could be incapacitated with only a few commands. If a CS is compromised, unauthorized users could gain access to the network and its attached systems. The CS could be disabled, therefore disallowing authorized subscribers from supporting mission critical functions.

From an architectural point of view, providing out-of-band management of network systems is the best first step in any management strategy. No network production traffic resides on an out-of-band network.

- **Use of in-band management is not limited, restricted, or encrypted** It is imperative that communications used for administrative access to network components are limited to emergency situations or where out-of-band management would hinder daily operational requirements. In-band management introduces the risk of an attacker gaining access to the network internally or even externally. In-band management should be restricted to a limited number of authorized IP addresses to improve security. The in-band access should also be encrypted for added security. Without encrypted in-band management connections, unauthorized users may gain access to network managed devices such as routers, firewalls, or remote access servers. If any of these devices are compromised, the entire network could also be compromised. Administrative access requires the use of encryption on all communication channels between the remote user and the system being accessed. It is imperative to protect communications used for administrative access because an attacker who manages to hijack the link would gain immediate access to the network.
- **Log all in-band management access attempts** Since in-band traffic travels on the same pathways as normal network traffic, be sure that all in-bound management access attempts are logged. This will give you an indication as to whether an intruder is attempting to gain control of key network devices. These attempts should not go unnoticed and should be verified against legitimate management activity of that device. For example, if the

access attempts happen after business hours, it's possible (or likely) that the attempts are unauthorized.

- **Two-factor authentication is not used for in-band or out-of-band network management** Without strong two-factor authorization, unauthorized users may gain access to network managed devices such as routers, firewalls, and remote access servers. If any of these devices are compromised, the entire network could also be compromised.
- **Filter ICMP on external interface** The Internet Control Message Protocol (ICMP) supports IP traffic by relaying information about paths, routes, and network conditions. ICMP unreachable notifications, mask replies, and redirects should be disabled on all externally-interfaced routers to prevent hackers using these messages to perform network mapping and infrastructure discovery.
- **SNMP access is not restricted by IP address** Detailed information about the network is sent across the network via SNMP. If this information is discovered by attackers, it could be used to trace the network, show the network topology, and possibly gain access to network devices. Access to SNMP should be for specific IP addresses only.
- **SNMP is blocked at all external interfaces** Clearly, using SNMP to map a network and discover the network infrastructure is a great hacker tool that should be secured to the greatest extent possible. This includes blocking SNMP on all external interfaces.
- **SNMP write access to the router is enabled** This allows an intruder to set various configuration settings to allow him or her greater access to the router and hence to the network. SNMP write access should be disabled.
- **Block identified inbound ICMP messages** Using inbound ICMP Echo, Information, Net Mask, and Timestamp requests, an attacker can create a map of the subnets and hosts behind the

router. An attacker can perform a DoS attack by flooding the router or internal hosts with Echo packets. With inbound ICMP Redirect packets, the attacker can change a host's routing tables.

- **Block identified outbound ICMP traffic** An attacker from the internal network (behind the router) may be able to launch DoS attacks with outbound ICMP packets. It is important to block all unnecessary ICMP traffic message types.
- **Block all inbound *traceroutes*** If you're ever had to troubleshoot a network or Internet connection, you're familiar with the *traceroute* command. This is a helpful tool in troubleshooting, but it also provides great information to a would-be attacker to create a map of the subnets and hosts behind the router. These should not be allowed into the network through the router or other externally facing devices.
- **Secure NMS traffic using IPSec** To securely protect the network, Network Management Systems (NMS) and access to them must be controlled to guard against outside or unauthorized intrusion, which could result in system or network compromise. Allowing any device to send traps or information may create a false positive and having site personnel perform unneeded or potentially hazardous actions on the network in response to these false traps. These sessions must be controlled and secured by IPSec.
- **An insecure version of SNMP is being used** SNMP Versions 1 and 2 are not considered secure and are not recommended. Instead, use SNMP Version 3, which provides the User-based Security Model (USM), which gives strong authentication and privacy. Without Version 3, it's possible an attacker could gain unauthorized access to detailed network management information that can be used to map and subsequently attack the network.

- **SNMP standard operating procedures are not documented** Standard operating procedures will ensure consistency and will help prevent errors or omissions that could create a security hole.
- **NMS security alarms not defined by violation type or severity** Ensure that security alarms are set up within the managed network's framework. At a minimum, these will include the following:

  - **Integrity violation** Indicates that network contents or objects have been illegally modified, deleted, or added.
  - **Operational violation** Indicates that a desired object or service could not be used.
  - **Physical violation** Indicates that a physical part of the network (such as a cable) has been damaged or modified without authorization.
  - **Security mechanism violation** Indicates that the network's security system has been compromised or breached.
  - **Time domain violation** Indicates that an event has happened outside its allowed or typical time slot.

  Also ensure that alarms are categorized by severity using the following guidelines:

  - Critical and major alarms are given when a condition that affects service has arisen. For a critical alarm, steps must be taken immediately to restore the service that has been lost completely.
  - A major alarm indicates that steps must be taken as soon as possible because the affected service has degraded drastically and is in danger of being lost completely.
  - A minor alarm indicates a problem that does not yet affect service but may do so if the problem is not corrected.

- A warning alarm is used to signal a potential problem that may affect service.
- An indeterminate alarm is one that requires human intervention to decide its severity.

Without the proper categories of security alarm being defined on the NMS, responding to critical outages or attacks on the network may not be coordinated correctly with the right personnel, hardware, software, or vendor maintenance. Delays will inevitably occur that will cause network outages to last longer than necessary or expose the network to larger, more extensive attacks or outages.

- **The NMS is not located in a secure environment** Any network management server (or any other highly critical network component) should be kept in a physically secure location with restricted access. Since many attacks come from inside an organization, by people who are authorized to be on the premises, it's important to physically secure all critical network components to the greatest degree possible. Using keypad or card-swipe access control can also help identify specific administrative access, to allow you to further control and monitor access.

  Access to NMS and other network critical components should be restricted via access controls as well, and all activity, including all successful and failed attempts to log on, should be logged. The log file, as with all log files, should be reviewed regularly, stored for 30 days, and archived for a year, unless regulatory or compliance requirements differ.

- **NMS accounts are not properly maintained** Only those accounts necessary for the operation of the system and for access logging should be maintained. This is true for all servers and network devices. Good "housekeeping" is an essential element to network security, and removing or disabling unused accounts as well as removing and investigating unauthorized accounts is critical.

Routers and Routing

- **No documented procedures and maintenance for MD5 keys** Routing protocols should use MD5 to authenticate neighbors prior to exchanging route table updates, to ensure that route tables are not corrupted or compromised.
- **MD5 Key Lifetime expiration is set to never expire** MD5 is a public key encryption algorithm that uses the exchange of encryption keys across a network link. If these keys are not managed properly, they could be intercepted by unauthorized users and used to break the encryption algorithm. This check is in place to ensure that keys do not expire, creating a DoS due to adjacencies being dropped and routes being aged out. The recommendation is to use two rotating six-month keys, with a third key set as infinite lifetime. The lifetime key should be changed seven days after the rotating keys have expired.
- **Console port is not configured to time out** Console ports on routers or other network devices should be set to time out after some specified period of inactivity. In most cases, a 5- or 10-minute timeout is appropriate. A router is a highly desirable asset to an intruder, so setting a low threshold on timeout will help increase security.
- **Modems are connected to the console or aux port** There may be valid reasons to have a modem connected to the console or auxiliary port of a router or other network device, but you should first ensure that this connection is absolutely necessary. If not, remove it. If it is needed, be sure to secure it by requiring a username and password (and other security measures) and avoid default configurations.
- **The router or network device's auxiliary port is not disabled** If the router or other network device has an auxiliary port, be sure it is disabled it if it's not in use. These are the kinds of welcome backdoors hackers look for.

- **Login is not limited to three attempts** Login attempts for any network device that exceed three tries are likely the work of a hacker. Limiting login attempts to three is a reasonable limit, and most net admins will stop after three attempts if they cannot recall the appropriate login. This won't stop a hacker who is willing to try three times, wait some specified interval, and try again, but it will prevent automated attacks from going through quickly (or at all).
- **Secure Shell timeout is not 60 seconds or less** Many routes and network management devices use the Secure Shell (SSH) protocol to secure communications to the device. Reducing the broken Telnet session expiration time to 60 seconds or less strengthens the router or network device from being attacked using an expired session.
- **Key services are not disabled on all routers** The DHCP, finger service, HTTP, FTP, and BSD *r*-commands and *bootp* services should be disabled on routers and network devices for added security. All unused protocols and services should be disabled to prevent unauthorized use of these services.
- **Configuration autoloading must be disabled** The routers can find their startup configuration in their own NVRAM or load it over the network via TFTP or Remote Copy (*rcp*). Obviously, loading in across the network is a security risk. If an attacker intercepted the startup configuration, it could be used to gain access to the router and take control of network traffic.
- **IP source routing is not disabled on all routers** IP source routing is a process whereby individual packets can specify routing. This is a method that attackers can exploit, so this ability should be disabled on routers and network devices with this capability.

- **Proxy ARP is not disabled** When proxy ARP is enabled on some routers, it allows that router to extend the network (at Layer 2) across multiple interfaces (LAN segments). Because proxy ARP allows hosts from different LAN segments to look like they are on the same segment, proxy ARP is safe only when it's used between trusted LAN segments. Attackers can leverage the trusting nature of proxy ARP by spoofing a trusted host and then intercepting packets. You should always disable proxy ARP on router interfaces that do not require it, unless the router is being used as a LAN bridge.
- **Gratuitous ARP is not disabled** A gratuitous ARP is an ARP broadcast in which the source and destination MAC addresses are the same. It is used to inform the network about a host's IP address. A spoofed gratuitous ARP message can cause network mapping information to be stored incorrectly, causing network malfunction and resulting in various types of service denials, leading to an *availability* issue.
- **Routers are not set to intercept TCP *SYN* attacks** The TCP *SYN* attack involves transmitting a volume of connections that cannot be completed at the destination. This attack causes the connection queues to fill up, thereby denying service to legitimate TCP users. Routers and similar network devices should be configured to intercept TCP *SYN* attacks to prevent DoS attacks from an outside network.
- **Router is not configured to block known DDoS ports** Several high-profile DDoS attacks have been launched across the Internet. Although routers cannot prevent DDoS attacks in general, it is usually sound security practice to discourage the activities of specific DDoS agents (a.k.a. *zombies*) by adding access list rules that block their particular ports.

- **TFTP used without specific need or approval, access is not restricted** Trivial File Transfer Protocol (TFTP) is a simple form of FTP that uses the User Datagram Protocol (UDP) and provides no security features at all (not even a password). It is often used by routers, X-terminals, and servers to boot diskless workstations, but by its very nature it is an insecure protocol. It should not be implemented without a very specific need to do so, and access to the TFTP server should be restricted and monitored.
- **The FTP username and password are not configured** The FTP server should require the use of usernames and passwords to prevent anonymous use of the FTP functionality on the network.

Firewall

- **Firewall not implemented and configured properly** You should ensure that one or more firewalls are installed and properly configured. The default configuration should be the most restrictive configuration, *deny-by-default,* so that only specifically allowed traffic is allowed into the network.
- **A screened subnet (DMZ) is not implemented** Without the dual-homed screened subnet (a DMZ), architecture traffic that would be normally destined for the DMZ would have to be redirected to the site's internal network. Computers on the inside of the firewall should send outbound requests through the firewall and into the DMZ. The DMZ, in turn, routes or redirects these outbound requests. Typically, a firewall will not accept inbound requests from the DMZ computers, which adds another layer of protection to the network clients.
- **Using an application-level firewall** All networks should use an application-level gateway or firewall to proxy all traffic to external networks. Devices such as SSL gateways, e-mail gateways that will proxy services to protect the network, are also accept-

able. A Layer 4 or stateful inspection firewall, in collaboration with application-level proxy devices, can be used to secure all connections.

- **Firewall does not require authentication, does not lock out after three attempts** Firewalls are the enforcement mechanisms of the security on the network, and they are ideal targets for attackers. Firewall placement in the network and the level of access granted to the users accessing the device also increase the risk profile associated with remote management. Therefore, all personnel who access the firewall both locally and remotely should be granted the minimum privilege level needed to perform their duties. The standard three-attempt lockout should be enforced, with the exception that when a firewall administrator is locked out, the senior net admin (or network security officer, if one exists) should be responsible for unlocking the account.
- **Firewall remote access is not restricted** Only the firewall administrator should be able to access the firewall remotely. Remove unused accounts and remove access for all staff other than the administrator.
- **Firewall is not configured to protect the network** Ensure that the firewall is actually configured to protect the network. Configuration of the firewall will vary from site to site, but in general, it should at least be configured to prevent TCP *SYN* flooding and the Ping of Death attacks.
- **Firewall has unnecessary services enabled** As with all network devices, disable, uninstall, and deconfigure any unused or unnecessary services. The fewer services that are enabled, the smaller the attack footprint.
- **Firewall version is not a supported or current** As with all network devices, it's critical to keep the firewall software (and hardware, if appropriate) up to date with current versions, patches, and updates. It's extremely common for attackers to

exploit known security issues days, weeks, or even months after a patch is available. This type of hacking is pretty lazy stuff and is a bit of an embarrassment if it occurs, because it's 100-percent preventable. Keep your firewall up to date.

- **The firewall logs are not being reviewed daily** There's really no point in creating log files if you're not going to review them. Reviewing and analyzing log files is part art, part science, but the only way you'll ever know what's going on is to actually review those files on a regular basis. If you don't know that a hacker was chopping away at your network security last night, you'll probably be surprised when he or she manages to hack in tomorrow night.
- **Firewall log retention does not meet policy** The firewall logs can be used for forensic analysis in support of incidents (after the fact) as well as to aid in normal traffic analysis. It can take numerous days to recover from a firewall outage when a proper backup scheme is not used. Firewall logs should be stored in secure locations; they should be stored for 30 days and archived for one year.
- **The firewall configuration is not backed up weekly** It's quite a chore to properly configure a corporate firewall, as you probably well know. Therefore, it's wise to back up the configuration data for the firewall on a weekly basis or whenever the firewall configuration changes. This provides excellent forensic support and helps in disaster recovery efforts.
- **The firewall is not configured to alarm the admin** If someone is knocking at the door but no one's home, an intruder may well decide to just barge right in. That's the net result of having a firewall that is not configured to alarm the administrator to unusual traffic.

- **The firewall is not configured properly** The firewall should be configured to protect the network. The following are suggested settings:
    - Log unsuccessful authentication attempts.
    - Stamp audit trail data with the date and time it was recorded.
    - Record the source IP, destination IP, protocol used, and the action taken.
    - Log administrator logons, changes to the administrator group, and account lockouts.
    - Protect audit logs from deletion and modification.

Intrusion Detection/Intrusion Prevention

- **The company does not have an incident response policy** An IDS is pretty worthless if you don't also have an incident response policy in place. Develop an incident response policy so there are clear lines of responsibility and reporting. Also clearly delineate how, where, and to whom to report suspicious activity.
- **Unauthorized traffic is not logged** Audit logs are necessary to provide a trail of evidence in case the network is compromised. With this information, the network administrator can devise ways to block the attack and possibly identify and prosecute the attacker. Information supplied by an IDS can be used for forensic analysis in support of an incident as well as to aid in normal traffic analysis.
- **No established weekly backup procedures** IDS data needs to be backed up to ensure that it is preserved in the event of a hardware failure of the IDS or in the event the IDS is breached.
- **IDS antivirus updates procedures not in the standard operating procedure** IDS systems require antivirus updates. Be sure that these updates are in the standard operating procedures

for IT staff. Sometimes it's the little things we overlook that bite us the hardest; this one's a no-brainer but easy to overlook.

- **Switches and cross-connects are not secure** Since the intrusion detection and prevention system includes all hardware required to connect horizontal wiring to the backbone wiring, it's important that all switches and associated cross-connect hardware are kept in a secured location, a locked room or an enclosed cabinet that is locked. This will also prevent an attacker from gaining privilege mode access to the switch. Several switch products require only a reboot of the switch to reset or recover the password.

Remote Access

- **The management VLAN is not secured** In a VLAN-based network, switches use VLAN1 as the default VLAN for in-band management and to communicate with other networking devices using Spanning-Tree Protocol (STP), Cisco Discovery Protocol (CDP), Dynamic Trunking Protocol (DTP), VLAN Trunking Protocol (VTP), and Port Aggregation Protocol (PAgP)—all untagged traffic. As a consequence, VLAN1 may unwisely span the entire network if it's not appropriately pruned. If its scope is large enough, the risk of compromise can increase significantly.
- **Remote Access Servers do not require encryption for end-user access** You should ensure that only users who require remote access are granted it and that all remote access traffic is encrypted to the fullest extent possible.
- **RAS does not use two-factor authentication** Without strong two-factor authorization, unauthorized users may gain access to network services, devices, and data. Clearly, if an intruder gains control of network infrastructure devices, he or she could inflict damage to either the data or the network, causing loss of confidentiality, integrity, or availability.

- **Remote Access Server connectivity isn't logged** Logging is your friend; keeping a log file of RAS connectivity is critical to keep track of who is attempting to log in, who did log in and when, and how long they were logged in. Reviewing log files daily will help you notice patterns and problems earlier in the cycle than reviewing log files infrequently (or never).
- **RAS session exceeds 30-minute inactivity** An RAS session that is inactive should be terminated to prevent session hijacking. Terminate idle connections after no more than 30 minutes of inactivity.
- **RAS log retentions do not meet requirements** Depending on organizational, legal, or regulatory requirements, you should keep log files for 30 days and archive them for one year.
- **The logs are not viewed on a weekly basis** Reviewing log files daily will help you notice patterns and problems earlier in the cycle than reviewing log files infrequently (or never).
- **Modems are not physically protected** Limiting the access to infrastructure modems and keeping accurate records of the deployed modems will limit the chance that unauthorized modems will be placed into the infrastructure. If an unauthorized person has physical access to a site's modems, the switch or software settings can be changed to affect the security of a system.
- **An accurate list of all modems isn't maintained** Keeping accurate records of the deployed modems will limit the chance that unauthorized modems will be placed into the infrastructure. It will also help you keep track of modems that are no longer used so they can be physically removed or disabled.
- **Modems are not restricted to single-line operation** Modems should be connected to phone lines that have very basic capabilities. If a phone line has advanced features such as call forwarding, it's possible an intruder could take control of a modem, computer, or network. Keep it simple for better security.

- **Proper call logs are not being maintained** Logs of all in-bound and out-bound calls for modems and phone lines should be logged and reviewed on a regular basis. Hijacked modems could conceivably allow an attacker to steal phone time and incur long-distance charges on your company's dime. Make sure you know what's going on with modems and phone lines to avoid big phone bills or network intrusion.
- **Callback procedures are not configured correctly** One way to increase security is to implement a callback feature on the modem so that a caller's call disconnects and the modem calls back a preprogrammed number. Ensure that if callback procedures are used, on establishment of the callback connection the communications device requires the user to authenticate to the system.
- **RAS/NAS server is not located in a screened subnet** Allowing a remote connection to the private network unchecked by the firewall enables a mobile user to violate the security policy and put the network infrastructure in a vulnerable position. The risk would be magnified if a remote access session were hijacked.
- **The RAS/NAS is not configured to use PPP** To securely protect the network, Network Access Servers (NAS) and access to them must be controlled to guard against outside or unauthorized intrusion, which could result in system or network compromise. If the NAS is accessed remotely, the risk of compromising a password or user ID increases. The authentication of the remote nodes must be controlled by encryption such as CHAP with MD5 or MS-CHAP with MD4.
- **VPN gateway is located behind the firewall** Allowing a remote connection to the private network unchecked by the firewall enables a mobile user to violate the security policy and put the network infrastructure in a vulnerable position. The risk would be magnified if the VPN connection were hijacked.

- **The VPN connection is not using IPSec's ESP tunnel** Ensure that remote access via VPN uses IPSec ESP in tunnel mode. For legacy support, L2TP may be used if IPSec provides encryption or another technology that provides security such as AES, 3DES, SSH, or SSL.
- **VPN is not configured as a tunnel type VPN** Be sure that VPNs are established as tunnel type VPNs, which terminate outside the firewall (in other words, between the router and the firewall, or connected to an outside interface of the router). If VPNs terminate inside the firewall, you basically have taken the firewall out of the security mix and reduced your line of defense by one. Improperly deployed VPNs take away a firewall's ability to audit useful information.

We've walked through a lot of very specific security information in this section, some of which might be relevant to your organization, some of which might not be. What is highly likely, though, is that if you even scanned this section, you thought of a few things you might otherwise have overlooked, or it sparked you to make a note to check one thing or another. The key is to be thorough, and to that end, this list should have helped you make sure you covered some of the nitty-gritty details of network infrastructure security.

Project Parameters

It's time now to plan your infrastructure security project. We've covered a lot of detail, and now we'll try to focus it down into a project plan that you can use to secure your infrastructure. Let's start with our problem and mission (outcome) statements. Remember, this is a good time to gather your core IT project team together to help you begin defining the basic project parameters. You probably could do some (or all) of this preliminary work on your own, but there's a lot to be said for getting the core team fired up and engaged with the project from the very start. You're less likely to have gaps in the project plan if you start relying on the "two heads are better" theory right from the start. Here are two sample problem statements you can use to begin developing your own:

> Our network infrastructure is vulnerable to attack because our security technologies have not kept pace with changes in the external environment. We currently do not have a meaningful approach to security, and all measures in place have been ad hoc or reactive. We are not confident of our level of security across the enterprise.
>
> We recently experienced a security breach that caused a network outage for three days. We were fortunate that no sensitive data appears to have been stolen or compromised. We took remedial measures, but we are not confident that our data or our network is secure.

Next, let's look at the mission or outcome statement for these problem statements. What's the desired outcome in both of these cases? The short answer is a secure network. We can probably use a single outcome statement for both problems, and it might look something like this:

> We want to create and implement a comprehensive infrastructure security plan so that we are confident we have developed and can maintain as secure a network environment as is reasonably possible.

Your possible solutions run the gamut, but we're going to assume that you've made the decision to secure your network infrastructure by developing a security plan focused on network infrastructure. That being the case, let's look at the requirements for this project.

Requirements

If you haven't done so, gather your core project team to work with you on the requirements. Requirements are those areas the project must address; in some respects, this is the real foundation of your project. Whatever is defined in the requirements should be implemented in the project, and whatever's in the project should be defined by the requirements. Success factors, those things required to make the project successful, can be defined within your requirements here or as part of your assumptions (discussed later in this chapter).

Functional Requirements

The functional requirements for your infrastructure security project will vary from the list we're providing, but your list should have the same overall elements. Where we would expect to see more divergence is in the technical requirements, which we'll discuss in a moment. Functional requirements might include:

- Physically secure premises
- Secure network infrastructure servers
- Secure network components (firewalls, routers, switches)
- Secure local communication (authentication, access control, encryption)
- Secure remote communication (authentication, access control, encryption)
- Secure user devices (operating system, antivirus, antispyware, application, file system)
- Create secure operating procedures
- Create documentation

Technical Requirements

Clearly, the technical requirements for your infrastructure security project will vary greatly from whatever list we provide, because the technical requirements are based on the specific network topology, server types, server operating systems, communication methods, authentication methods, and more. The lengthy list of items presented earlier in this chapter should provide you with plenty of ideas and material for creating your technical requirements for the project. Instead of going into detail here, we present the categories that you should include; the details under those categories are up to you, based on your unique requirements.

Remember, technical requirements should describe the "how" of your functional requirements. So, as you work through this section, keep in mind that you should be describing, very specifically, how you will implement the functional requirements via technology.Your technical requirements should be detailed in describing how these things will be accomplished, but be careful to stick to describing the *requirement*. Let's take "Physically secure the premises" as an example.You don't need to describe how the premises will actually be secured ("Phil will get a screwdriver and install a deadbolt …"); you need to describe the *technical details* of how it will be secured. So, let's say you're the only tenant in a building.You might describe your technical requirement in this manner:

- Upgrade all external entry doors to card-swipe system. Card-swipe system should be compatible with the existing employee card system, XYZ. (You might include the technical specs of this system here as well.)
- Install security monitoring system (with cameras) focused on parking lot and all external doorways (3). System should be able to record continuously for 24 hours, cameras should be able to record in slow motion and high resolution, the system should be able to "respond" to potential incidents, and the system should record events and have at least three methods of administrator alert.

These are just some of the ways you can capture technical requirements. Clearly, if you're talking about a server, you would include processor speed, memory specifications, disk drive specifications, operating system, and so on. However, other kinds of less technical elements, such as how to secure the premises, might look like the example provided. If your card-swipe system, for example, must conform to certain standards, those standards should be included as well.

Legal/Compliance Requirements

Create a list of the functional, technical, and administrative requirements for your infrastructure security project based on the legal, regulatory, and compliance requirements. Taking time to translate these requirements into project requirements at this juncture will help ensure that you build compliance requirements into your project. In standard project management, it's always easier to build something in at the front end than to add it at the back end (it reduces errors, omissions, time, and cost), so now's the time to add these requirements to the greatest extent possible. Also, be sure to add milestones and documentation requirements to your project plan based on compliance needs.

Policy Requirements

Policy requirements may fall under functional requirements, but there's no rule that you can't include policy requirements as a distinct category of requirements if doing so helps you cover all the bases. We'll look at policies in more detail in a later chapter, but for now, let's walk through a few ideas for policies related to securing the infrastructure:

- User policies
- Network access policies
- Remote access policies
- Wireless policies
- Network administration/network management policies
- Server policies
- Firewall, IDS/IPS, DMZ policies
- Regulatory/compliance policies
- Corporate policies
- Legal policies

Scope

At this point, you should have an idea of the scope of your project. You could choose to address your complete infrastructure security needs during this project, or you might choose to parse it out into smaller subprojects and time them in stages or phases to meet organizational needs. Making changes to the infrastructure comes with risk, and you'll need to be careful to take this fact into consideration as you plan your project. This starts with determining the proper scope for your project. For example, you might have recently implemented an IDS that you're satisfied with, so you could choose to include IDS in your project only to the extent that it ties in with other infrastructure security measures. However, you might feel that your biggest exposure is on network servers such as DHCP, DNS, and directory servers, so your primary focus will be to harden these servers and related network traffic. Your assessment should tell you where you need to focus and what must be included in the plan and perhaps what can safely be omitted from your plan. Then clearly define what is and what is not part of your project so that you leave nothing open to interpretation.

Schedule

Since we haven't created a detailed work breakdown structure (WBS) yet, we can't develop a detailed schedule, but we can begin to develop a higher-level schedule. First, you should take a look at your organization and see if there are any events or timelines that might come into play. You certainly don't want to be in the middle of a network outage (due to an upgrade) when an important client is visiting, when your marketing department has a big presentation coming up, or when your manufacturing group is working overtime to get a large order out. Taking organizational needs into consideration is critical to project success and helps grease the political gears as well.

In addition to organizational needs, you might know about other timelines or constraints that should be considered in the schedule. Are other security initiatives being planned or under way? If so, is there a logical order to the plans themselves? It might make sense to complete early phases of an infrastructure security project before implementing a new wireless network security plan, for example. Also, look at other IT projects to determine how they might impact the infrastructure security plan or how the infrastructure security activities might disrupt or alter other projects that are in the planning or implementation stages.

Finally, look at your talented IT staff and determine if there are any scheduling issues that would impact your project, such as your best wireless or IP person heading out for vacation or your encryption specialist planning to be out for a month on paternity leave. Whatever the case, if you already know about these scheduling issues, you might as well begin addressing them here.

You might have a rough idea of how long this project will take, given what you've looked at thus far, and you may be able to see where it will fit in your overall IT schedule. You'll have to balance the demands for your IT resources with the need to secure the infrastructure, so this is a good point to try to get a handle on some of those schedule constraints.

Budget

Your budget will be large or small depending on how well secured your infrastructure currently is and how large your company is. For example, if you already know that you're going to have to upgrade some of those old servers still running Windows NT (gasp), then your budget is going to have to include a whole host of things like the server box itself, the operating system, license upgrades, and updated applications. As with scheduling, you won't have an exact amount in mind yet, but you might have some large segments defined or at least identified. Begin making a list of the components you believe you've identified in terms of purchases so you can verify (or modify) this list after you've created your WBS. Also keep in mind that with an infrastructure security project, a large percentage of your budget might be expended on labor costs (if you track internal labor costs in your project) because much of the work entails checking configurations and modifying settings.

Quality

You could define quality as the level of protection you're willing to accept, though it might be difficult to quantify. As we've stated, quality is a mindset, and you should instill this mindset in your IT project team. As you define your project plan, you'll have the opportunity to create specific quality metrics related to your infrastructure and incorporate them into your task details. Remember that security comes from depth of defense, so you want each layer you build to be as strong as it can be, within the defined constraints (time, cost, criticality, and so on) and understanding that no system is 100-percent secure.

Once you've defined the project parameters, it's essential that you develop your priorities. If you haven't checked in with your project sponsor, this would be a good time to discuss the priorities. What's the least flexible element here? Are you expected to meet a particular deadline, or were you handed a set budget? Understanding which of these parameters must be met will tell you how to make decisions during the project work phase. If you have a deadline, you'll focus your efforts on making sure that the schedule stays on track, which might mean spending a bit more on overtime than was in the original budget. As you know, in projects something always changes and something *has* to give. Understanding where that flexibility should come from will help you meet organizational requirements. Understand which parameter is least flexible and which is most flexible. Discuss this with your project sponsor, and make sure you're clear and in agreement. That way, when project work is under way, you won't have to keep going back to your sponsor to make basic decisions, and you'll know you're making decisions that support these priorities.

Key Skills Needed

For your network infrastructure security project, you're going to need a very wide variety of skills. Here we list some of the obvious ones; you can add to (or modify) this list as you define your project:

- **Network services** Securing the infrastructure requires a solid look at security settings on infrastructure servers such as DHCP, DNS, and directory services servers. These key servers require a deep understanding of the services they provide as well as an understanding of best practices in each of these areas.
- **Network perimeter services** Securing the perimeter involves installing, configuring, and managing components such as firewalls, routers, proxy servers, and DMZs. These typically require a strong background and ability to work with various protocols, including SNMP, ICMP, TCP, IP, FTP, HTTPS, SSH, SSL and more.
- **Intrusion detection/intrusion prevention** Installing and configuring IDS/IPS systems require a strong skill set in networking, understanding how information and IP traffic flows through the network infrastructure, and understanding the kinds of threats that are commonly launched (TCP, ICMP, etc.).
- **Remote access** Securing remote access requires an understanding of communication devices and protocols as well as of various authentication and encryption standards and methods. RAS, VLAN, VPN, and tunneling are just a few of the concepts needed in this area.
- **Wireless access** Securing a wireless LAN is discussed in a later chapter, but the skills here are the ability to understand and use various wireless network tools (the same ones the hackers use) and an understanding of how wireless networks are vulnerable and can be protected using a variety of tools.
- **Servers and hosts** These entail understanding operating systems, patches, upgrades, and vulnerabilities as well as how to secure files, folders, data, and user accounts.
- **Network administration** A strong understanding of network administration tools and techniques, including the ability to audit,

review, and manage user and group accounts, access control lists, services and protocols, and other administrative tasks, is critical.

- **Documentation** You need people who are excellent at documenting the systems, the proposed changes, the implemented changes, and the final infrastructure configuration data. Documentation may be required for legal or compliance requirements as well.
- **Communication** You should have one or more people on your team who are good at communicating and creating connections within the organization. Since infrastructure changes can have a big impact on business operations, you will need to effectively and proactively communicate with various stakeholders and users during the course of the project.
- **Training** You might need to train users or IT staff on new methods, technologies, or other changes that occur as a result of implementing the infrastructure project plan.

Key Personnel Needed

Now that you've developed your list of needed skills, you can develop your wish list for project team members. You might want to create your "A" list and a backup list, but for political purposes, you might want to call it your Primary and Secondary teams, to avoid ruffling any feathers. Be sure to highlight the skills needed for which there are no internal matches. This will indicate places you need to seek staff training or external expertise. Also look at your personnel needed and get a sense for how much you're relying on a two or three people. It's often the case that we want the three people who are best to work on everything, but that will slow your project down tremendously. Develop your personnel list, then determine where your gaps are and how you'll address those gaps.

Project Processes and Procedures

As with any project, you should identify the processes and procedures you'll use during the duration of the project. We are assuming you have a whole stash of those at your disposal, so we won't run through the basics here. However, there are five areas to keep in mind when you're working on an infrastructure security project; you might want to check that these are in your processes and procedures. If not, add them as needed:

- **Testing procedures** Define how, when, and where you'll test security solutions before implementing them on the live network. Clearly, some things can be done on live systems; others should be tested before going live. Define what should be tested offline and how those testing scenarios should proceed.
- **Rollback procedures** For any major changes you're making, be sure you've identified and tested reliable rollback plans so you can roll back to a known good state in case things go wrong.
- **Escalation procedures** These are standard in any project, but when you're dealing with the infrastructure, you might need to beef this area up a bit to make sure you have the right people on standby or on call when critical portions of the project are being implemented.
- **Critical issue reporting** As with other security projects, you might want to review your critical issue-reporting procedures to see whether you need to create a "team red response" that will enable you and your project team to quickly address any vulnerabilities or issues that are deemed critical, severe, or extreme. Develop the process for defining these kinds of issues, and develop an agreed-on scale or measurement system so your team can quickly deal with imminent or urgent issues that may arise.
- **Documentation** Depending on the nature of your project as well as your regulatory or compliance requirements, you might need to revisit your typical project documentation processes and

procedures so that you can create the kinds of documentation your project requires, without sorting back through the project to develop the documentation after the fact.

Project Team

We always recommend gathering at least your core team together to help define the project, but at this point, you have defined the skills and personnel you need, so you should be ready to create your project team. As with almost all IT projects, you should involve subject matter experts from outside the IT department so that you get a well-rounded view of life within the organization, not just within the IT department. Also keep in mind that with infrastructure issues, many areas are technically beyond the grasp of many employees within the organization (well, most everyone outside the IT department). Be patient, be prepared to explain things in nontechnical ways, and don't discuss technical details when it isn't absolutely necessary. You want to involve users and stakeholders in appropriate ways, but there will be a more limited role for them in this type of project than in most of the other kinds of IT security projects you might undertake.

An infrastructure security project spans the entire enterprise and is as deep as it is wide. To be successfully completed, the project requires an extensive set of first-rate skills. Be sure that your project team has all the skills delineated in your skills assessment. If it does not, your project is at risk because gaps in skills or skills that are not up to standards will create problems in quality or scheduling down the line. Be sure you have all the requisite skills; if you don't, be sure to create a plan to address those gaps. Whether you need additional external staff or just some training for internal staff, be sure to add this to your budget and possibly your schedule.

Develop your team, create a team roster, and get the team ready to create the detailed project plan, especially the work breakdown structure.

Project Organization

You might want to organize your project by the topic areas we've defined so that you have subteams dedicated to: devices and media, topologies, IDS/IPS, and system hardening. You could choose to parse it out differently, or you might have the whole team, if it is small, work through each stage of the project plan together. It's up to you, and it's certainly somewhat dependent on the size of your company, the size of the project, and the size of your IT staff. Be sure that everyone is clear about what their roles are within the team and within the project. Organizing your team will provide the necessary structure for the team to be productive. Since we're assuming you've managed a lot of projects (or have read up on your project management skills recently), we won't delve into the details of organizing this project except to make one note: The work in the various segments of this project overlaps a lot, and if your project and team are not well organized, you're going to have people working at cross-purposes and creating a big mess. Keep the project and your team organized to avoid this scenario.

Project Work Breakdown Structure

Your approach to creating your work breakdown structure (WBS) might be different from the method we provide; that's fine as long as you cover the basics. Our recommended approach is to start with your mission statement and your selected solution and create three to five high-level objectives. From there, you can parse each of those objectives down into smaller components until you have tasks that actually make sense and are understandable. Tasks should be broken down until they represent an understandable and manageable unit of work. The 80/8 rule is a good one to keep in mind; it states that no task should exceed 80 hours or be less than 8 hours. If a task is longer than 80 hours, it needs to be broken down into smaller components. If you define tasks of less than 8 hours, you'll end up with a scheduling nightmare on your hands.

We'll start with the four major areas we discussed at the opening of this chapter:

1. Devices and media
2. Topologies
3. Intrusion detection/intrusion prevention
4. System hardening

If you recall from prior chapters, these are not properly written as tasks or even as objectives; they're topic labels. So, let's fix that and create the top-level objectives based on these four areas of concern:

1. Audit and secure devices and media
2. Audit and secure network topology
3. Implement or harden intrusion prevention/detection systems
4. Harden systems

Now we have a better starting point for our WBS. From here, we can break these down into smaller tasks. We're not going to dig down as deep as you'll need to, because once you get beyond a certain level of detail, the plan is very much dependent on the nature and structure of your organization and how you and your team decide to approach the project. So, don't fight with the structure presented here; use it as a guide to create one that works for you. Also note that where servers or other devices may be called out, the numbers or types of devices may not track with standard networking practices. They are presented as examples of a WBS tree, not necessarily examples of best practices in networking. In reality, you will have more or fewer DNS servers, but we only mention one. You will have a long list of tasks under Task 3.4, "Assess and harden routers, switches, and other network communication devices." We didn't dig down at all levels of the WBS but provided samples of how or where you might develop additional tasks and subtasks. And, while this list is long, it's not as long as your infrastructure security project plan's WBS will end up being. However, this should give you a running start:

1. Audit and secure devices and media.

2. Audit and secure network topology.

2.1 Create secure boundaries using firewalls, DMZs, and proxy servers.

2.2 Create secure remote access.

2.2.1 Secure all Remote Access Servers.

2.2.1.1 Physically secure Remote Access Servers.

2.2.1.2 Secure Remote Access Servers.

2.2.1.2.1 Remove excess administrative accounts.

2.2.1.2.2 Disable all unused services, ports, and protocols.

2.2.1.2.3 Remove all unused applications.

2.2.1.2.4 Disable all unused modems.

2.2.2 Secure remote communications.

2.2.2.1 Evaluate the feasibility and desirability of implementing VLAN.

2.2.2.1 Evaluate the feasibility and desirability of implementing VPN.

2.3 Create secure wireless access.

2.3.1 Change all wireless access points' default settings.

2.3.2 Disable SSID broadcasting, create a closed system.

2.3.3 Enable MAC address filtering.

2.3.4 Evaluate and implement encryption (WEP or WPA).

2.3.5 Filter wireless protocols.

2.3.5 Define IP allocations for the WLAN.

2.3.6 Evaluate VPNs for possible implementation.

2.3.7 Secure users' wireless devices.

2.3.8 Develop wireless policies for users.

2.3.9 Develop wireless policies for IT operations.

2.4 Implement a segmented network.

2.5 Implement network traffic security protocols for sensitive network traffic.

2.6 Deploy network security technologies.

2.6.1 Use Encrypting File System (EFS) or similar file encryption.

2.6.2 Require and use strong user authentication, passwords, and account policies.

2.6.3 Employ the concept of "least privileges" when assigning user rights.

3. Implement or harden intrusion prevention/detection systems.

3.1 Assess security of current IDS/IPS system or evaluate need for implementing IDS/IPS system.

3.1.1 Evaluate intrusion detection system feasibility and desirability.

3.1.2 Inline intrusion prevention system feasibility and desirability.

3.1.3 Network active response system feasibility and desirability.

3.1.4 Host active response system feasibility and desirability.

3.1.5 Network processors feasibility and desirability.

3.2 Assess and harden DMZ or evaluate need for implementing DMZ.

3.3 Assess and harden firewall or evaluate need for implementing additional firewalls.

3.4 Assess and harden routers, switches, and other network communication devices.

4. Harden systems.

4.1 Evaluate physical security and access control to critical servers.

4.1.1 Evaluate and secure access to domain controllers.

4.1.1.1 Evaluate and secure domain controller 1.

4.1.1.2 Evaluate and secure domain controller 2.

4.1.1.3 Evaluate and secure domain controller 3.

4.1.2 Evaluate and secure access to DHCP server.

4.1.3 Evaluate and secure access to DNS server.

4.2 Review and revise administrative accounts on infrastructure servers.

4.2.1 Remove unused or superfluous administrative accounts.

4.2.2 Remove unused or unnecessary non-administrative accounts.

4.2.3 Remove unused rights and privileges.

4.3 Implement strong authentication and password policies on all infrastructure devices.

4.4 Review, record and update (as needed) operating system and application version levels.

4.4.1 Review and record operating system versions on all infrastructure servers.

4.4.1.1 Review and record operating system version on domain controller 1.

4.4.1.2 Review and record operating system version on domain controller 2.

4.4.1.3 Review and record operating system version on domain controller 3.

4.4.1.4 Review and record operating system version on DHCP server.

4.4.1.5 Review and record operating system version on DNS server.

4.4.2 Update operating systems on all infrastructure servers.

4.4.2.1 Update operating system on domain controller 1.

4.4.2.2 Update operating system on domain controller 2.

4.4.2.3 Update operating system on domain controller 3.

4.4.2.4 Update operating system on DHCP server.

4.4.2.5 Update operating system on DNS server.

4.5 Review current status of virus protection software installed on servers.

4.6 Assess and implement server, application, and client-side security technologies.

4.6.1 Secure server traffic traveling on the network.

4.6.2 Secure application and user data traveling on the network.

4.6.3 Secure network access points and network access.

4.6.4 Secure client devices including desktops, laptops, and PDAs.

4.6.4.1 Upgrade all insecure "legacy" operating systems.

4.6.4.2 Update all operating systems with latest revisions, patches, and updates.

4.6.4.3 Update all applications with latest revisions, patches, and updates.

4.6.4.4 Update all virus protection programs.

4.6.4.4.1 Ensure latest virus definition file is loaded.

4.6.4.4.2 Ensure virus program is configured to automatically download the latest definition file from secure server or Internet site (WSUS in Windows or vendor Web site).

4.6.4.5 Enable file encryption for mobile devices.

4.6.4.6 Implement strong passwords.

4.6.4.7 Update user policies to prevent downloading or installing of unsigned programs.

5. Document all infrastructure changes.

5.1 Document changes to all infrastructure configuration settings.

5.2 Document changes to network topology, layout, or structure.

5.3 Document changes to standard operating procedures.

5.4 Document changes to user policies and procedures.

6. Perform compliance audit.

Once you've completed the WBS, you need to go through with your subject matter experts and develop the task details. Details can include task owners, resources, known constraints, or requirements for the task, task duration, task cost or budget, tools or equipment needed for the task, completion criteria, deadline or due date, and any other data relevant to the task and its successful completion. Remember that the functional, technical, and legal requirements should be fully incorporated into the project task detail or they will get lost. This is a great opportunity to review your requirements and go through your task details to ensure that everything is included, before project work starts.

This is also a point at which you should do a scope check and make sure that the WBS describes your intended scope. It's fairly common for the scope described by the WBS to be larger than the stated scope. In fact, this is often the first source of "scope creep." Look at your scope statement and at your WBS and reconcile any discrepancies. For example, you might have stated in your scope statement that something was not part of the project scope but that element shows up in the WBS. Decide if that element should be in or out, then adjust either your scope statement or your WBS accordingly. If there are substantive changes to your scope, check in with your project sponsor to gain agreement as to the modified or updated scope and WBS.

Project Risks and Mitigation Strategies

This section of your project plan defines the risks to your project and the strategies you'll use to avoid or mitigate your risk. There are always risks with every project, and it's important to take time to identify those risks while you're calm, cool, and collected. There are some projects for which the risks outweigh the benefits and you decide, as a team or an organization, to not go down that path. Securing the infrastructure is not likely to fall into that category, but it's always important to keep this in mind—that sometimes doing nothing is a better choice.

However, you've decided to strengthen security on your network infrastructure and there are attendant risks. Let's look at one risk you might have, and you can then use this structure to develop additional risk and mitigation strategies. We'll use the following ranking system: 1 = Extremely high, 5 = Extremely low.

Risk: Improper configuration could completely disable network.

1. Criticality: 1
2. Likelihood of occurrence: 3
3. Relative risk ranking: 2
4. Mitigation strategy 1: Test all configurations in lab prior to rollout.
5. Risk of mitigation 1: Not all lab tests will completely mirror actual conditions.

6. Mitigation strategy 2: Develop fail-safe rollback plans for all critical configuration changes.
7. Risk of mitigation 2: Rollback will take time and set back project completion timelines.
8. Trigger 1: One week prior to scheduled configuration change.
9. Trigger 2: Forty-five minutes after network outage occurs.
10. Notes: All configuration changes will be tested in the lab first, but the there is still a chance that the configuration change could cause the network to crash. If this occurs, rollback plans will be implemented after 45 minutes of network downtime have elapsed.

As you can see from this single example, you can develop sound contingency plans for risks you decide are worth planning for. Some risks are too small to bother planning for; other risks are significant but unlikely to occur, and planning for them would also not be a good use of time.

Once you've listed every risk you can think of, you can develop a ranked list based on both criticality and likelihood of occurrence. From there, you can develop mitigation strategies for just the most critical and most likely-to-occur risks. You may choose to develop more than one mitigation strategy. In our example, our first choice was to test in the lab, but we also conceded that testing in the lab may not mirror real-world results and the risk of our mitigation strategy had to be addressed as well. In this case, we developed a secondary or backup mitigation strategy as a fail-safe option. Both mitigation strategies require a defined trigger—how will you know when to implement your risk mitigation plan? In this case, you might choose to build lab testing into your project plan for all configuration change tasks and avoid this first mitigation strategy. However, you would still need the second mitigation strategy and trigger in the event that you ran into a configuration problem that you couldn't immediately find. In this case, after a 45-minute outage you'd go to Plan B, your predetermined rollback plan. It's nice to have Plan B ready to go when you're running around like you hair is on fire because you have 47 different people asking you when the network will be back up, what the problem is, and why they can't log onto the network.

Project Constraints and Assumptions

Constraints for an infrastructure security project might come in all shapes and sizes. You may well face budgetary constraints that limit the scope of your project. You might face scope problems because the infrastructure needs a lot of upgrading but your company isn't willing to implement all the changes needed, for a variety of financial, political, and organizational reasons. You might face resistance within the organization because some changes impact users' computing behaviors, and this can cause problems. In addition, you may face specific constraints or limits within portions of your project. For a variety of reasons, you might not be able to make changes to application servers or database servers due to other projects under way or other organizational issues. The infrastructure security is core to your network's security, so constraints to the project should be clearly identified and discussed. If the constraints are too great or impede your project too much, you should have a talk with your project sponsor. Although every project has to deal with a variety of constraints, you must decide whether the constraints are reasonable or if they place an undue burden on your project. You are responsible for project success and, ultimately, for the security of the infrastructure, so it's up to you to clear away these obstacles or get your project sponsor involved with removing them so the project can be successful.

Equally important is delineating the assumptions you're working under as you move into the project. For example, if you assume certain resources will be available or if you assume that other projects will be completed first, you should state that clearly. Your assumptions should be clearly delineated so that you and your team can *challenge*, *clarify*, or *confirm* those assumptions before proceeding with the project. The most dangerous assumptions are the ones we don't know we're making.

For example, if you've been in the midst of deploying a particular encryption method, your project would work on the assumption that the encryption scheme was already in place or was being deployed. This fact is critical to note in your project because, if something outside your control changes the encryption scheme on which your infrastructure project is based, you'll have to rework your project plan. This could (and probably would) impact both your schedule and your budget. It's hard enough to bring a project in on time and on budget without having the project environment shift around on you. Listing the assumptions you're making going into the project will help you as you develop your plan as well because others on the team or in the organization can challenge your assumptions, if they know what they are. To go back to the encryption example, if you list this as an assumption and someone on your project team lets you know that the encryption project was put on hold for one reason or another, it's good to know that ahead of time rather than planning based on that incorrect assumption.

Some project managers like to list project success factors in their list of assumptions because they are assuming these factors will be in place or will occur. You can also discuss and define success factors at the front of your project-planning process, if that's a more logical flow for you. Some teams don't know what it will take to be successful until they've neared the end of their definition and planning work; others like to define these elements right up front. Whatever works for you is fine. Just be sure to define these so you'll know what it will take for your project to succeed.

You and your project team will have to look at the project environment and list the constraints and the assumptions you're making in order for your project to get off to a good start and to have a better-than-average chance for success. These elements are unique to each project and each company, so we can't give you a list of things to place in this section, but now that we've discussed them in general terms, you and your team should be able to dig in and find the constraints and assumptions for your own, unique infrastructure security project.

Project Schedule and Budget

You can see from the lengthy project plan we've created that your schedule and budget are going to be challenging to develop. Once you've created your WBS, you can look through your task details and begin developing your schedule. The schedule is best developed in a project management software program since you will have a lot of moving parts to handle. If you have subteams working in parallel on different aspects of the project plan, be sure you address this in your schedule. First, you'll have to be sure you're not double-booking someone and throwing your schedule off. Second, you want to keep an eye on how different segments of the project will impact other segments so you don't end up working at cross-purposes, or worse, damaging something another team just implemented. If one team is upgrading the firewalls and another team is working on IPS, it's entirely possible one team's work will greatly impact the other team's work and cause confusion, problems, errors, or omissions.

Be sure to check your critical path tasks after you've loaded your schedule into the software program, since these tasks will determine the longest, least-flexible path through your project. Although we haven't discussed the more technical aspects of scheduling (we assume you know them), recall that you can indicate lead and lag times as well as float to create a more realistic schedule. If everything in your project plan ends up on the critical path, or if none of your tasks end up on the critical path, there's a good chance you have a fundamental problem with how your schedule is set up.

As for budget, you should have a pretty clear idea of what this project will cost at this point, with one notable exception. If you are using your infrastructure security project plan to evaluate the need for an IDS, IPS, DMZ, or other network equipment, you might not yet have sufficient data with which to get bids for these systems. In that case, you need some sort of placeholder to indicate that a system will be purchased but the system has yet to be clearly defined and therefore cannot be spec'd out or priced. If you know the order of magnitude, it might be good to add a dollar-amount placeholder. For example, suppose you know that one type of system you're looking for costs about $18,000, plus or minus $2,000. You might want to put $20,000 into your budget as a placeholder so that when your project budget is approved, you have that cost built in. It's usually difficult to get your budget increased after it's been approved, unless you specifically get your budget approved with the understanding that it does *not* include the cost of new hardware or software solutions that may be recommended as a result of the project assessment.

If you've made it this far, you've made it to the end of the chapter and the end of your planning cycle for your infrastructure security project plan. It's a lot to cover because the infrastructure is wide and deep, but if you take time to step through your planning in a measured, thoughtful manner, you'll end up with better results than if you just rush headlong into the project work. That's a guarantee. Your project might not be perfect, it could come in late or over budget, but whatever result you turn in will be far better than if you used no consistent approach or framework at all.

IT Infrastructure Security Project Outline

- Audit and secure devices and media.
- Audit and secure network topology.
  - Create secure boundaries using firewalls, DMZs and proxy servers.
  - Create secure remote access.
  - Create secure wireless access.
  - Implement a segmented network.
  - Implement network traffic security protocols for sensitive network traffic.
  - Deploy network security technologies.

- Implement or harden intrusion prevention/detection systems.
  - Assess security of current IDS/IPS system *or* evaluate need for implementing IDS/IPS system.
  - Assess and harden DMZ *or* evaluate need for implementing DMZ.
  - Assess and harden firewall *or* evaluate need for implementing additional firewalls.
  - Assess and harden routers, switches, and other network communication devices.
- Harden systems.
  - Evaluate physical security and access control to critical servers.
  - Review and revise administrative accounts on infrastructure servers.
  - Implement strong authentication and password policies on all infrastructure devices.
  - Review, record, and update (as needed) operating system and application version levels.
  - Review current status of virus protection software installed on servers.
  - Assess and implement server, application, and client-side security technologies.
- Document all infrastructure changes.
  - Document changes to all infrastructure configuration settings.
  - Document changes to network topology, layout or structure.
  - Document changes to standard operating procedures.
  - Document changes to user policies and procedures.
- Perform compliance audit.

Summary

We've covered a lot of ground in this chapter because your network infrastructure is literally and figuratively the backbone of your network. Infrastructure security touches every aspect of your network, and a thorough assessment will take time and careful effort to complete so that your network is as secure as it can reasonably be, given the organizational constraints and considerations you'll have to deal with. It's often helpful to break the network infrastructure down into it systems or areas to help ensure that you cover all the areas, including devices and media, topology, intrusion detection and prevention, system hardening, and all the network components such as routers, switches, and modems. Once you've identified all the areas, you need to take a top-to-bottom look at how security is currently implemented and what threats exist. By looking at issues such as information criticality and performing an impact analysis, you can decide what should be included in your project and what can reasonably be left out or delayed for a later phase if needed. Understanding the threat environment and your network's vulnerabilities is also important during your planning phase.

Requirements need to be thoroughly developed because they form the foundation of your project's scope. Functional requirements should be developed first, followed by technical, legal, and policy requirements. Be sure to build these into your task details when you create your WBS so that all required elements will be present and accounted for in your project plan.

In an infrastructure security project, you'll need a wide variety of skills that span the depth and breadth of networking knowledge. Be sure you define the skills you'll need so that you can assess your team and your organization to identify skills gaps. These will have to be addressed before your project can proceed, and this often requires hiring outside contractors or providing training for internal staff members. Either way, this can impact both your budget and your schedule, so be sure you do a gap analysis between needed and available skills prior to proceeding with your project.

The WBS defines the scope of your project, so once you've identified all the work through delineating the tasks, be sure to do a scope check. If the defined scope is smaller than the scope outlined in your WBS, you need to reconcile the differences. Also be sure to discuss any scope changes with your project sponsor so that you start off with the same expectations about project results.

Scheduling an infrastructure security project can be challenging due to all the moving parts involved. You'll run into scheduling conflicts, resource usage conflicts, timing issues, and more. These should be resolved to the greatest degree possible before starting the project, because things will only get more complicated and difficult to resolve once project work is under way. One important scheduling note is that with all areas of your network being poked and prodded, you'll need to make sure subproject teams are not working at cross-purposes and undoing work just done or inadvertently injecting false indicators into the process through their own task work.

When it's all said and done, you should be able to define, implement, manage, and close a very successful infrastructure security project, if you follow a consistent methodology and make teamwork and quality topmost priorities. This is the foundation of all other security projects; it touches on everything in your organization, so success here will create the framework for a very secure network that will help you sleep at night, knowing you've done everything possible to keep your organization's assets secure.

Solutions Fast Track

Infrastructure Security Auditing

☑ Auditing or assessing the infrastructure security is a large task that encompasses every aspect of your network.

☑ Infrastructure projects cross several boundaries, and you should be sure that any overlap is addressed so you are not working at cross-purposes.

☑ The infrastructure project can be parsed out in numerous ways. One way is to look at it in terms of these systems: network perimeter, internal network, intrusion monitoring and prevention, host and server configuration, malicious code protection, incident response capabilities, security policies and procedures, employee awareness and training, and physical security and monitoring.

☑ The internal and external environments should be assessed thoroughly prior to planning your project.

☑ Internal factors include understanding information criticality, the potential impact of a breach, the information flow, policies and procedures, user needs, and regulatory/compliance issues.

☑ Externally, you need to consider the types of threats your network is vulnerable to, including spoofing, repudiation, data tampering, denial of service, and elevation of privileges.

☑ The SAN Institute publishes a list of the top 20 vulnerabilities that can serve as a great starting point for assessing your network's vulnerabilities.

☑ The assessment should look at devices and media, topologies, intrusion detection/intrusion prevention systems, and system hardening.

☑ Devices and media include all the network infrastructure devices that must be secured, including routers, switches, and modems.

Project Parameters

☑ Defining the functional and technical specifications for your infrastructure project will define the scope of work you need to accomplish.

☑ There could be specific legal or regulatory compliance issues to be addressed within the scope of your infrastructure security plan; you should include these issues in the early stages of project planning.

☑ Defining scope, initial budget, initial schedule, and quality guidelines based on the functional, technical, legal, and policy requirements gives you a solid starting point.

☑ Developing the relative priority of your parameters and gaining project sponsor agreement is important in helping you know how to make decisions for your project, moving forward. The least flexible parameter will be your constraint; the most flexible parameter will be what "gives" when things change in the course of project work.

☑ Technical skills needed for the project include network services, network perimeter, intrusion detection/prevention, remote access, wireless access, server and host administration, familiarity with protocols, ports, and services as well as skills in documentation, communication, and training.

Project Team

☑ Your core project team should help in defining the project.

☑ Your infrastructure project team should include people with the needed skills, which, in an infrastructure security project, are extensive.

☑ Your project is at risk if you don't have the skills you need on your team. Be sure to address skills gaps before you start project work and add any associated costs into your project budget.

Project Organization

☑ The project should be organized around areas of the network and areas of expertise.

☑ Typical organizational methods should form the foundation of your project organization.

☑ Infrastructure projects require extra coordination to ensure that subteams are not working at cross-purposes.

Project Work Breakdown Structure

☑ The work breakdown structure, or WBS, for an infrastructure project should begin with the high-level objectives for your project, which might include securing devices and media, securing the perimeter, securing infrastructure components, or whatever way you choose to segment the work in this project.

☑ Task details should reflect the functional, technical, and regulatory requirements for your project. Check task details against requirements to be sure everything is included at the outset.

☑ The scope statement and the scope described by the WBS might not be in sync. Compare them and make any modifications needed to either your scope statement or your WBS before proceeding.

☑ If there are significant changes to your scope statement, check in with your project sponsor before proceeding to ensure that you're both on the same page with regard to scope.

Project Risks and Mitigation Strategies

☑ The risks inherent in an infrastructure project are many because this type of project touches every aspect of your network.

☑ Identify all potential risks and rank them according to criticality and likelihood of occurrence, then look over the list and make any reasonable adjustments.

☑ Determine how far down your risk list you will plan, then develop mitigation strategies and triggers for each defined risk.

☑ If you determine that there are one or more significant risks, you should sit down and talk them over with your project sponsor. In some cases, the risks outweigh any potential benefit and the project should be canceled, redefined, or postponed until those risks can be more clearly evaluated and addressed.

Project Constraints and Assumptions

☑ Constraints are present in every project, but in an infrastructure security project, you could have constraints on several fronts.

☑ If the constraints are too great, they can hinder or prevent project success. Discuss major constraints with your project sponsor to determine the best course of action. Otherwise, develop ways to address these constraints or plan around them.

☑ Clearly delineating assumptions is an important part of your infrastructure security plan because you have to work from a known good point.

☑ If your assumptions are stated, they can be challenged, clarified, or confirmed.

☑ Assumptions may also include success factors, since you might be assuming certain factors must be in place for the project to succeed. Success factors are also sometimes developed in the requirements phase of the project-planning process, depending on how you go about identifying them.

Project Schedule and Budget

- ☑ After you've developed your project's WBS, you should have sufficient data to create a fairly realistic and feasible project schedule and budget.
- ☑ Keep in mind that the project schedule has a lot of moving parts, and you're likely to run into issues around conflicting resource demands or subteams working at cross-purposes.
- ☑ Be sure your budget provides for adequate training or hiring of needed resources.
- ☑ Large purchases such as IDS/IPS or other major components might not be decided at the outset of the project. Create ballpark estimates or insert placeholders in your budget so it is clear whether or not large-ticket items are included or specifically excluded from your project budget.

Infrastructure Security Project Outline

- ☑ Audit and secure network topology.
- ☑ Audit and secure network topology.
- ☑ Implement or harden intrusion prevention/detection systems.
- ☑ Harden systems.
- ☑ Document all infrastructure changes.
- ☑ Perform compliance audit.