

Producing Your Network Security Policy

Executive Summary

Network security experts agree that well-run corporations need a written security policy. The policy sets appropriate expectations regarding the use and administration of corporate IT assets. However, the conventional wisdom holds that composing and maintaining these documents bog down in a morass of bureaucratic inefficiency and pointless wrangling, which never ends and produces nothing useful.

This paper lays out a common-sense approach to writing corporate security policies that makes them easier to draft, maintain, and enforce. Our "question and answer" approach requires no outside consultants. Instead, you can use your in-house knowledge and resources to yield a brief, usable, and – most importantly – understandable policy document, in a reasonable amount of time. To help you generate such a policy, this paper clears away some misconceptions about the purpose of network security; details the process of writing the policy; then explains how to keep refining the drafted policy.

Introduction

It is the rare organization that is happy with its security policy. Many will admit to not even having one. But, security policies are like noses: everyone has one. Every organization follows either a formal or an informal security policy, even if it is what we jokingly refer to as the Primordial Network Security Policy: "Allow anyone in here to get out, for anything, but keep people out there from getting in."

Realistically, many security policies are ineffective. Sometimes an organization gets lucky and has a security policy that is pretty good – but not usually. To be effective, a security policy (and, let's reset that right now to "security policies," because we are talking about a set of policies) should be consistent, relevant, and useable. The goal of this white paper is to help you create such documents.

Armed with this paper, your small- or medium-sized enterprise (SME) can either create your first computer network security policy, or beef up what you already have. This paper covers policy but not procedures. Computer and network security policies define proper and improper behavior; they spell out what is permitted and what is denied. Procedures detail the methods to support and enforce the policies, and usually describe specific steps to take in regular system administration. For example, your policy might state, "Server administrators must adhere to the company's operating system configuration standards." A separate procedures document would specify what all those settings are.

This paper will help you set policy. First, we correct some misconceptions to help you understand what your real goals are. Then we describe the process for writing your policy, and end with some thoughts on what to do after completing your initial draft.

Four Common Misconceptions

1. "The goal of network security is to secure the network" (or "the computers"). Securing the network is easy, but it's not your goal. Your real goal — and a more difficult job — is securing the business.

The goal of network security is to support the network and computer business requirements, using methods that reduce risk. Security policies describe what you must secure, and the ways you secure them, to support your business or mission. Firewalls, intrusion detection systems

(IDS), anti-virus (AV), backup and restore strategies, locked doors, and system administration checklists are all some of the things you might use. Security policies provide the blueprint for using them: the what, how, why, when, and by whom.

2. "Security policies must be long and complex." In fact, just the opposite is true. We believe the well-known security axiom, "Complexity and security are inversely proportional." Complex systems are usually less secure than simple systems. Complex policies are usually ignored; simple policies might live.

A good security policy is really a set of documents, each addressing a specific need. By breaking your overall policy into smaller pieces, each managed separately, you greatly simplify the process of creating effective, consistent, relevant, and useable documents. This is not to say that the entire set of security policies will or should be just a few pages. But each individual element — each policy — should be usable by the target audience. "Usable" does not mean merely "understandable," or even "readable" and "memorable." It also has to take into account your corporate culture. So keep it real. Don't write academic tomes (unless that is your corporate culture). Write something your target audience can read and understand, in the amount of time their duties permit them.

3. "Security policies have to be nearly perfect, or 100% complete." No. Good enough security now is better than perfect security never.

For some reason organizations treat security as something sacred, when this is exactly the area where practicality should reign. There is not one right way to write a security policy. You are also allowed to modify it later. General George S. Patton said, "A good plan, violently executed right now, is far better than a perfect plan executed next week." That is not to say that your goal should be to produce something shoddy or incomplete and call it "whole." But it is perfectly fine to build security policies in parts, refining each part separately in the ongoing process of security policy development. Some parts will seem fully baked before other parts do. That's OK. That's how the process works.

4. "Security policies only have to be written once." Until there are no more bad guys in the world and everyone agrees to mind his or her own business, the process of managing a security policy never ends.

The threats your organization faces will change over time. As the threats to your business change, so too will your company's business requirements. The vulnerabilities will change as well, and so will the risks you are willing to take to do business, and so will the tools you use to reduce or counter those risks. Because of all this, the security policy process is never really done. It only lies dormant for a time.

If you're willing to believe the truths of the previous four paragraphs, then press on. We will next discuss:

- f* The general structure of the security policy
- f* The process of putting the policy documents together
- f* What documents you should create
- f* What they should say
- f* How and when to review and revise them

The Process

The first step in writing your policies is to gather a team. Writing a set of security policies is usually a top-down process, but it does not have to be, and may combine bottom-up and top-down approaches. Your policy development team should be made up of people who work with your network and the Internet, but come from different functional areas of the company. Each manager in your company has a unique view of the company's needs and risks. You need people who know something about the technology, but also some who know about business. Include some people from the trenches, too. There is nothing less useful than a painstakingly documented security policy that, when implemented, makes the shipping department unable to track packages, or blocks the sales reps from network resources they need from the road.

However, don't let the process of forming the committee halt all progress. Remember, well begun is half done; you can start developing the drafts with just a few knowledgeable IT staffers.

Before writing any policies, scope out your business requirements. What regulations apply to your industry (GLBA, HIPAA, Sarbanes-Oxley, ISO17799, new state or local laws, etc.)? Get familiar with penalties for any non-compliance, as this will help you prioritize your policies and gauge the proper level of discipline for employees who do not adhere to policy. To consider other business issues, ask yourself:

- f* What services are required for your business, and how might you provide them securely?
- f* How much do employees depend on Internet access, use of email and availability of intranet services?
- f* Do your users need remote access to the internal network?
- f* Is there a business requirement for everyone to have access to the Web?
- f* Do customers access your data (technical support, order status, etc.) via the Internet?

It takes discipline to ask repeatedly, "Is there a business requirement?" for every service. But the business requirements are the most important drivers of your security policies. Business drivers help you distinguish between what the organization really needs, as opposed to what a few employees want. If you have trouble getting started, look at what you are already doing and ask, "Why are we doing that?" The answer will kick-start your response to the questions above.

Policy Creation

In this section we will interactively start to draft your policies, beginning with the Root Security Policy and then working through a few of the others. Each policy sets out the definitive answer to a set of key questions.

The Root Security Policy

The first document you'll draft is the "Root Security Policy." This is also the easiest to write, as it is the framework which points to the other policy documents.

As you draft a Root Security Policy, you will also enumerate the initial list of subordinate policies that you should produce next.

Your list will be specific to your organization, but will probably include the following subordinate policies:

- f* **Computer Acceptable Use.** A general document covering all computer use by employees and contractors, including desktop, mobile, home PCs, and servers.
- f* **Password.** A description of the requirements for password protecting computer systems, the rules for choosing passwords, and how the password policy is enforced.

- f* **Email.** This policy covers the use of email sent from any company email address and received at any company computer system.
- f* **Web.** A specification of what browsers may be used, how they should be configured, and any restrictions on which sites employees can visit.
- f* **Mobile Computing and Portable Storage.** A description of who owns the mobile computing and portable storage on your network, how they are supported, and what specific devices (if any) are authorized for use on the company network.
- f* **Remote Access.** A policy stating who can access what information from which locations under what circumstances.
- f* **Internet.** A description of your Internet-facing gateway configuration, stating what is allowed in and out, and why.
- f* **Wireless.** A specification stating how wireless access will be managed on your network; how access points will be plugged in, secured, and maintained; who is allowed to use them; and under what circumstances.
- f* **Servers.** A statement of the company standards for servers, what services are enabled or disabled by default, and important distinctions between production, test, and development environments.
- f* **Incident Response Plan.** No policy is complete until it also specifies what to do when defenses fail: what is considered a security incident; who gets called; who is authorized to shut things down if needed; who is responsible for enforcing applicable local laws; who speaks for the company.

The Root Security Policy will also specify numerous easy-to-describe items, such as:

- f* **The company name.** Does this policy apply to your whole company, or a distinct division, office, or locale?
- f* **The purpose of the policy.** What is it for? What do you hope to gain by creating a policy?
- f* **The individuals or organizations responsible for the policy.** Who is responsible for overall network security? The head of the IT department? The information systems security officer? Some other executive? (Tip: In the drafting phase, you are allowed to write “To Be Determined.”) Eventually, you will also want a “Site Security Committee,” a small group of managers and technical people representing various user groups in the organization, including the HR department. This committee is responsible for maintaining the policy so that it is current and relevant.
- f* **The scope of the policies.** Make sure that you state what geographies, organizations, and assets you are covering. State explicitly which geographies, organizations, and assets the policy does not cover, as well.

The Root Security Policy should also briefly describe:

- f* **Penalties for breaking policy.** Determining this will probably require talking to upper management and the Human Resources department, but will almost always include words such as, “disciplinary action up to and including termination of employment.”
- f* **Who enforces the policy.** All managers and supervisors should be responsible for the administration and enforcement of the policy.

- f* **Who must abide by the policy.** All employees should be responsible for adhering to the policy. Will there be exceptions? Under what circumstances? Take the time to define an exceptions process, making sure that the process calls for periodic review of the exceptions.
- f* **The other documents listed in the policy.** The root policy forms the framework for the rest of the document. This entry can be as simple as a bulleted list or as detailed as a bibliography.
- f* **How to request policy changes.** There will and should be changes as the policy matures. Specify when and how changes can be made, and who can make them.
- f* **How often your policies must be reviewed.** For most SMEs, a year is too long to go without a policy review. Monthly reviews are too frequent. If you are not sure what interval to specify, start with “quarterly, or as needed.” (“As needed,” for example, might be after a request for a change, or when the requirements or the apparent threats change.)

How many pages will all this fill up? As many as it takes, but usually no more than two to five pages (initially) for the Root Security Policy, and one or two pages for each individual sub-policy. In school, you might have gotten used to padding your documents to meet some minimum page requirement. In security policies, follow the opposite route: brevity reads like wisdom.

Acceptable Use Policies

After you have the Root Security Policy, producing the subordinate or AUP policy largely means making lists and asking questions. List the assets you must protect. For example, your list might include:

- f* Desktop computers
- f* Mobile computers
- f* Servers
- f* Routers
- f* Email systems
- f* Application data

For each asset, ask:

- f* Who administers the asset?
- f* Who uses it?
- f* How critical is it to the mission of your enterprise?
- f* How do you manage it?
- f* How do you protect it?

This exercise gives you a checklist of assets to cover with AUPs. All of your AUPs will answer similar, overlapping core questions, and will have the same format. Each AUP will address:

- f* **Objective.** What is the purpose of this particular AUP?
- f* **Target.** Describe the systems to which this AUP pertains.
- f* **Responsible parties.** To whom does the policy pertain?
- f* **Policy.** A detailed statement of what the policy is. What does the policy permit? What does it deny? What are the user responsibilities? What must the employee report?

As you go through the process of creating your AUPs, you must answer these questions at least. If other good questions occur to you, answer them as well (For example, according to policy, where and how is data stored? Is an employee allowed to store the sales forecast on a home computer? When and how should data be destroyed?). Some information in the AUPs will be redundant, but one of the goals is for each AUP to be somewhat self-contained.

Answer each question with a simple declarative sentence or paragraph. You want to reduce the possibility of your text being misinterpreted. Therefore, as you answer the questions, don't merely write to be understood. Try to write so that you cannot be misunderstood. Prefer specific words over vague words. Prefer active sentences over passive sentences. For example, in a passive sentence like "This policy is to be adhered to at all times," who must adhere to the policy? We can't tell. Clearer: "All IT team members must follow this policy, always."

For the balance of this guide, as we describe each AUP we'll assume that you have answered the fundamental questions above. The rest of our description will concentrate on issues unique to each AUP.

Personal Computer AUP

Questions to answer include:

- f* Who owns the PC? (Most of the time it's the company, but you'll need the policy to also address, for example, the home PC of a telecommuter.)
- f* Are there any restrictions on non-business use? (For example, may a company computer be used for games, or personal email?)
- f* Who is authorized to use the PC? Only the employee to whom it is issued? Any employee of the company? An employee's immediate family?
- f* How should the user protect the computer data? (For example, must the user encrypt files?)
- f* How should the user protect the computer from unauthorized access? (Passwords? Password-protected screensaver? How many minutes before the screen saver times out?)
- f* What software must be running on the PC? (Examples: Antivirus? Personal firewall? A spyware detector? What versions? etc.)
- f* Are there restrictions on software installation? (For example, is the user permitted to install anything downloaded from the Internet?)
- f* What special protection is required for mobile computers?
- f* What activities or classes of activities are prohibited?
- f* Will you monitor keystrokes or communications? If so, how will you notify employees of this?
- f* Who is responsible to back up computer data?
- f* How must the user protect or mark personal information on a company-owned computer?

Email

Questions to answer include:

- f* May employees use email accounts for non-business-related email?
- f* Must employees include a disclaimer when they send non-business-related email? When they post to public email forums?
- f* Must they (or may they) encrypt and sign messages? If so, how?

- f* What restrictions apply to sending email? (For example, you should generally prohibit spam, illegal transmissions, chain letters, etc.)
- f* What, if any, attachment types are prohibited (sending or receiving)?
- f* Is email subject to monitoring? If so, how will you notify employees of this fact?
- f* What email client software is permitted?
- f* May users access outside email accounts (other ISPs, Hotmail, Hushmail, etc.)? If so, under what conditions?
- f* May employees access web-based email accounts from company PCs? What rules govern the use of POP? IMAP?
- f* Who may use corporate email systems or email clients?

Web Access

Questions to answer include:

- f* Are there any restrictions on accessing external web sites? (To answer this, you might want to consult the web site of one of the many web filtering services. They'll list categories of objectionable material you might not have thought of, including sports sites, hate sites, gambling sites, etc.)
- f* May users access non-business related email accounts via the Web, using company machines?
- f* Do you restrict certain types of web content, such as streaming media, or content that might violate copyright laws?
- f* What browsers may or may not be used?
- f* What, if any, rules govern the use of browser add-ins or Java applets?
- f* Are there any required browser configurations?

Mobile Computing & Portable Storage

Questions to answer include:

- f* May users bring their own PDA to work? If so, are they permitted to access the company network via their PDA?
- f* What PDAs are supported?
- f* May users install software on the PDA?
- f* Are there restrictions on devices with wireless capabilities (see Wireless AUP)?
- f* May corporate data be stored on user-owned portable storage (USB jump drives, smart phones, iPods, etc.)? If so, how must it be protected?
- f* Should the storage device be password-locked? Should it be encrypted?
- f* Do you restrict the type of information you permit to be stored on portable devices?
- f* Is any particular software prohibited?
- f* Many devices use portable SD or CF memory cards. Are employees permitted to use these? What rules govern their usage? Can company data be stored on these cards? (Consider also unorthodox "storage devices" such as digital cameras, novelty devices, toys with RAM, DVDs, etc.)

Remote Access

Questions to answer include:

- f* Do you restrict remote access to the enterprise to just authorized users (probably—i.e., not employee's family members)?
- f* Is wireless access permitted? (See Wireless AUP.) Is access from Internet cafés permitted? If so, under what circumstances and with what safeguards?
- f* What software and hardware combinations and configurations are required for remote access?
- f* May users access the enterprise network from devices your IT department has not issued, or has not authorized?
- f* How will you authenticate (confirm the identity of) the person accessing the network?
- f* How is access controlled? By passwords, security tokens, VPNs, or what?
- f* Is remote access granted to all employees, or must they apply for it? How do they apply?
- f* What activities are prohibited? What activities are permitted?
- f* Is account activity monitored? If so, how will you notify employees of this fact?
- f* In what ways must a user protect the remote access account?

Internet-facing Gateway Configuration

This could apply to a router, switch, firewall, or UTM. If you have more than one, write an AUP for each one, but use the same format for each. Identify which gateway you are writing about in the Target section of each AUP.

Questions to answer include:

- f* How is the device accessed? What authentication is used? Are there other required safeguards for access? What protocols are permitted or denied for administrative access? (For example, “SSH required, TELNET prohibited.”)
- f* What outgoing protocols are permitted? Which are denied? (If Trojan horse code infests your network, proper egress filtering can prevent the code from contacting its author to establish a back door into your network.)
- f* What incoming protocols are permitted? Which are denied? (Ingress filtering.)
- f* What application-level controls are in place for filtering? For example, do you filter HTTP (Web) traffic? SMTP (mail) traffic? How about Domain Name Services?
- f* Will you control Internet access based on protocol and target-system? For example, will the gateway permit email- or web-related protocols from the Internet to any internal machine, or just to email and web servers?
- f* Will you control outgoing connections to ensure that only appropriate connections are made? For example, will email-related protocols be allowed outbound only from the enterprise email server(s) and not from desktop computers? What about other services, such as DNS and NTP?

Servers

This AUP should apply to all dedicated servers used by multiple users within the enterprise that may also communicate with other servers (Internet or intranet). This is a general policy. You might have more specific ones for HTTP servers, FTP servers, email servers, etc.

Questions to answer include:

- f* What is the main purpose of the server?
- f* What information must be kept about the server (IP address, MAC address, physical location, operating system (OS) version, patch level, services offered, person responsible, etc.)?
- f* May individuals deploy internal servers, or is that right restricted to IT alone?
- f* What are the OS configuration policies that must be followed? (Usually, the answer points to another document for OS-specific procedures.)
- f* How is administrative access controlled?
- f* How is the server physically protected?
- f* How is the server monitored, and by whom?
- f* What is the back-up policy?
- f* Who may access the server (non-administrative access)? Who may not?
- f* What is the patch-control policy for the server? (This may point to another “patch control” policy or procedure document.)
- f* What change control procedures must be followed? (This may point to another “change control” policy or procedure document.)

Wireless Devices

The Target section of your Wireless AUP might look different than the Target section of other AUPs. You must specify if your wireless policy covers:

- f* All wireless devices, i.e., mobile phones, PDAs, computers (This is the most comprehensive option).
- f* Only those that directly connect to the enterprise network (This omits basic mobile phones.)
- f* Those that indirectly or occasionally connect to the enterprise network, such as PDAs, PDA/phone combinations, and devices like the Blackberry. (This omits wireless enabled laptops.)

Depending on the degree to which you use wireless technologies, it may be necessary to draft a Wireless AUP specific to each type of device (e.g., Blackberry® wireless devices, wireless enabled laptops, 802.11 capable phones).

Once you've defined the Target, questions to answer include:

- f* Is wireless access to your network allowed?
- f* Are any kinds of data or communication prohibited over wireless?
- f* What protection must be in place before wireless communication is authorized?
- f* How must the user protect the wireless device (physically and logically)?
- f* For wireless data, what hardware is approved, permitted, and/or required?

- f* For wireless data, what software is approved, permitted, and/or required?
- f* What rules govern wireless access points? Must they be deployed, configured, and installed by IT, or are other employees permitted such activities?
- f* What are the configuration requirements for wireless access points? (For example, "The password and username must be changed from the manufacturer's default.")
- f* Is VPN software required? Which or what kind? Who must or may install it?
- f* Where must wireless connections terminate on the network? In the DMZ? On a segregated VLAN?
- f* How must wireless connections authenticate and encrypt?
- f* Is wireless permitted for remote users connecting to the Internet with enterprise equipment?
- f* Is wireless access permitted for mobile users connecting into the enterprise network from outside the enterprise?
- f* What security devices and controls must be in place on authorized wireless devices?
- f* What configuration rules must users follow when connecting home wireless networks to enterprise assets?

Incident Response Plan (IRP)

“The best laid plans of mice and men often go awry.” The sections above deal with your Acceptable Use Policy. But even with the best security devices, practices, and policies, you still might find yourself dealing with a network security intrusion or incident. Since mid-incident is the worst time to develop a plan, your next step is to formulate an Incident Response Policy.

Initial questions to answer in your IRP include:

- f* What do you consider a security incident? (You probably will consider web site defacement or a virus outbreak a security incident. But, is a port scan of all your Internet-facing systems a security incident? How about if they are port scanned once a day for a week? What if you discover the LAN room was left unlocked overnight?)
- f* If an incident occurs, who are you going to call? Everyone in the organization should know who to call. Everyone who is on the call list should know what to do with a suspected security incident.
- f* When must you call the police, FBI, Secret Service, or other local or federal civil authorities? Talking with a lawyer or your local FBI field office will help here.
- f* Which are your most important systems? Which are most difficult to recover? Which are least important or easiest to recover? If a security incident brings systems down, balancing the importance of each system against how long it takes to recover it will help you prioritize your triage efforts.

The IRP should contain contact information for everyone who must be contacted when an incident is discovered or suspected. People to consider adding to the "call list" include:

- f* System and network administrators
- f* Senior management
- f* Managed service providers
- f* Help desk
- f* Lawyers
- f* Public relations

f Law enforcement

And then from that list, answer:

- f* Who must be contacted immediately?
- f* Who can be contacted later? What are the outer boundaries of "later"?
- f* Who is responsible to contact whom?

Note that any mention of a security incident made to people outside your organization can have unforeseen repercussions. Your policy should state who is authorized to discuss your company's security with outsiders. All other insiders should be prohibited from divulging information, especially to the press.

Modern legislation means that every security incident can have legal implications. Even if you do not intend to prosecute anyone, the law might obligate you to disclose the incident. Consider having your IRP specify the title of the person in your company who is responsible for having functional familiarity with relevant laws in your country, state, and municipality; in other words, they know what actions the laws support and which actions they prohibit, and can advise others on what to do.

Your IRP should remind you of the steps to take during a security incident:

- f* The most important thing to remember: take good notes. During an incident, you will be tempted to wait until things are contained to document the event. Take the time to take good notes as events unfold, for a whole lot of reasons ranging from helping you understand the incident later, to defending your actions in court. You don't have to write down every little command you type. Write what symptoms made you take action. Document each major milestone in your response, noting each new tool used (including version numbers) and why you did what you did.
- f* Try to answer as many of the W's for any incident as possible: Who, What, Where, When, Why, and How.
- f* Make sure you're addressing the entire problem by assessing its scope. How bad is it? How many systems are affected? How certain are you that you're aware of the entire problem? (Axiom: Evil usually multi-tasks.) How bad might it become?
- f* Secure the systems affected. You need printed procedures for securing each system. It may be as simple as unplugging the network connection or pulling the power plug. It may require posting a guard. Seriously.
- f* If you plan on going to court over the incident, now is the time to call in experts. Do not touch anything else. Do not do anything else. "Computer forensics" is a specialty. Forensics requires evidence—certifiably unchanged data. Take no action that might change data.
- f* If you are not going to prosecute the incident, secure the vulnerabilities, if any, that led to the incident. Restore the system from a known good baseline and return it to use.

You will want to add and improve upon this rudimentary IRP. Incident response planning deserves a course of its own; this paper presents just a start. Like all security-related policies, your IRP should be tested, periodically evaluated (especially after use), and revised.

Next Steps

What a lot of work! Yet the task is feasible, isn't it?

If you go through the steps laid out in this paper, you will end up with a draft security policy. At this stage, there is danger on two fronts.

First, you might foolishly think that you are finished. This is just the starting point. You now have an imperfect blueprint. You must refine it, and then — when you are satisfied with it — you must periodically pick it up and view it again. Review with a critical eye, asking, “What is right? What is wrong? What needs to be changed? What should be completely scrapped?”

The other mistake is to think that the whole process is rubbish. If you have gone through the exercise, no matter what you ended up with, you are almost certainly further along than when you started. Even if most of what you have needs major revisions, by answering the questions you have learned a lot about what you are doing and what your enterprise needs in security policies.

In either case, take a break from the process and come back at it with a fresh perspective in two to four weeks. As you review what you’ve done in the previous round, look for areas where your answers could be misinterpreted, where your needs have changed, or where your original ideas aren’t working out. As with owning a house, part of owning a security policy is tinkering with it. Don’t be afraid to make necessary changes to meet your evolving business needs, and don’t forget to schedule your next periodic review.

Conclusion

Improving your existing security policy (even if it’s the primordial variant) need not overwhelm you. You can do it. In fact, there’s probably no one better suited. Ask questions, challenge your assumptions, write it all down, and give yourself permission to be less than perfect.

Following these simple steps will give you a great shot at producing a brief, usable, and most importantly understandable policy document in a reasonable amount of time. If no better benefits emerge, “I helped create my company’s security policy” will look great on your resume.

For more information about WatchGuard security solutions, visit us at www.watchguard.com or contact your reseller.