Title: Principles of Secure Network Design
Document Number: CASE-260-002
Version: 1.1, August 27th, 2001
OS Ver. this Paper Applies to: All
HW Platforms this Paper Applies to: All
Audience (Internal or External): External

Principles of Secure Network Design

Introduction

NetScreen network security devices are among the fastest and most effective in the world. However, even the most efficient and powerful tool can be rendered useless by poor implementation. Network security, by its very definition, is a difficult subject for which to obtain information. This whitepaper attempts to address that issue by informing NetScreen customers about effective and secure network design, thereby maximizing the usefulness of their investment in NetScreen products. Be sure to check out the Reference section at the bottom for links to additional information not covered in this whitepaper.

Concepts

Throughout this paper, terms and concepts common to network security professionals will be used. The following is a brief glossary of terms used, as well as other security terms and concepts you should be aware of.

Hardened System – A server that has had all appropriate security patches installed, bug fixes applied, and which has been configured securely. These systems have been designed to resist penetration. Each port a server listens on is a potential vulnerability – if you are not using that service, or do not have it patched, patch it or remove it.

Bastion Host – A Hardened System taken one step further - configured with the minimal software to support a single network service. These hosts and the data they possess must be considered expendable. A loss of any one Bastion Host should not compromise the security of any other host.

Demilitarized Zone (DMZ) – A subnet or group of subnets separated (typically physically) from the more sensitive areas of a network infrastructure, and populated with Hardened Systems and Bastion Hosts. Any system put on a DMZ must also be considered expendable.

Stateful Inspection – A firewall process that checks the TCP header for information on the session's state – whether it is initializing (SYN), ongoing (SYN/ACK), or terminating (FIN). A stateful inspector firewall will typically track each session flowing through it. Packets from unknown sessions that appear to be part of an ongoing session (illegal) are dropped. All NetScreen network security devices are stateful inspectors.

Packet Filtering – A router/firewall process that contains access control lists ("**ACL's**") that restrict flow of information through it based upon protocol characteristics such as source/destination IP address, protocol or port used. Generally, packet-filtering routers do not track sessions through them unless the router is also doing a NAT process, and the NAT process would track the session for NAT purposes.

Network Intrusion Detection System (NIDS) – A system that passively monitors all data flowing past its network interface, looking for hostile data patterns. A flexible NIDS will have user-definable data patterns or "**Signatures**" which allow it to detect new types of attacks as they evolve. A NIDS will send an alert to an administrator or to an active firewall for evaluation and action. Most any network security system should have at least one NIDS.

Honeypot – A special-purpose system designed to be hacked. This may be a copy of a real server, but with special software to track, contain, or otherwise occupy the attacker while an evaluation of what the attacker is doing can be made or session tracing performed. These are generally found in very complex network security configurations. Honeypots can be placed in high-profile areas, or areas where attackers usually search for targets. Some very advanced systems will even deflect attacks from a real server to a Honeypot.

Whitehat Hacker – An individual who is very interested in network security. Also known as an "**Ethical Hacker**". These people will work hard to break down popular servers and protocols, looking for vulnerabilities in them. Once a vulnerability is discovered, they work with the vendor of the affected system to fix the problem. They may write a non-malicious exploit - a program that takes advantage of the problem - demonstrating the vulnerability. These people see themselves as the "Good Guys" in cyberspace – the good guys in old Westerns wore white hats. They also consider themselves 'ethical' since they know of bad things to do to others, but do not act on them.

Blackhat Hacker – An individual who is very interested in network security. Also known as a "**Cracker**" for cracking open a hardened system. These people will work hard to break down popular servers and protocols, looking for vulnerabilities in them. Once a vulnerability is discovered, they use that vulnerability to create a malicious exploit. They do this illegal activity for fun, profit, or political gain ("**Hacktivism**"). Since these guys are obviously the "Bad Guys" of cyberspace, the Blackhat moniker applies.

Greyhat Hacker – An individual who sometimes acts like a Whitehat, and sometimes a Blackhat Hacker. Many people who tout themselves as Whitehat Hackers are actually Greyhats – they will do a malicious (Blackhat) act against someone they dislike, or if they think they won't be caught.

Script Kiddie – [Derogatory term] – A person who dabbles in computers and network security who typically has limited knowledge about how the whole system works. They take pre-written code (or "Scripts") and blindly throw them at systems on the network in an attempt to affect them in some way. It has been estimated that over half of the world's well-known hacker attacks have been perpetrated by these sorts of individuals.

Denial of Service (DoS) Attack – An attack designed to disrupt a network service. Typically in a DoS attack, a flood of information from the attacker will overwhelm a serving system's resources, causing it to be unable to field valid network requests. Other DoS attacks can cause the serving process to crash, also denying the service.

Distributed Denial of Service (DDoS) Attack – A DoS attack (typically a flood) from multiple source points. This is more effective, as it is no longer one attacker against one server in an effort to overwhelm the server. Now, many low-bandwidth connections can be added together to attack a high-bandwidth site. Additionally, having random floods from multiple paths make backtracking extremely difficult, if not impossible. Sometimes, a personal home computer (on a cable modem or DSL line) is an unwilling participant in a DDoS attack through the use of a "**Zombie**" agent resident on their computer – typically put there by a Blackhat by hacking into that person's computer.

IP Address Spoofing – A common technique used by attackers to do dirty work. Depending on the circumstances, a spoofed IP can be used to perform difficult-to-trace DDoS attacks, hide their true address in a clutter of bogus addresses, or in rare occasions to take advantage of IP address related trusted relationships between two hosts. The attacker sends "**Crafted Packets**" (Packets made from the ground up, and not created and processed normally through the IP Stack) that have source IP addresses other than what has been assigned to the interface. These IP addresses can be anything – they are typically completely random. IP address spoofing works best when the spoofed addresses used are not currently in use, as otherwise, the real host may reply and reset the session. The target system receives these packets, and responds as appropriate, with the return traffic being sent to the actual address' owner, and not the attacker. In order for an attacker to receive return traffic from a spoofed address, the attacker must be in the path of transmission to the actual owner's address.

Port Scanning – A technique for determining which ports a server is listening to. A port-scanning program (such as Nmap – see the Reference –Tools section) will barrage a server with connection requests on a range or list of ports, and report back which ports the server responded on. This is useful when an attacker knows of a vulnerability of a particular protocol or daemon, and wishes to find targets to exploit. They are also useful for system administrators to determine the security level of their servers, and to test the effectiveness of firewall access policies.

Trojan Horse – A program with functionality (typically malicious) not made known to an end-user. A common example of this would be perhaps a cute "Elf Bowling" game you received as an email attachment. This 'Trojaned' program might also secretly install a remote administration (AKA "Back Door") program that allowed an attacker access to your computer. Programs Trojaned are usually games or other amusements that compel you to run them. Quite often, you will get a copy from a friend who thought it was "neat" – neither of you knowing you were giving an attacker complete access to your computer.

Worm – A self-replicating attack program. Worms differ from typical viruses in that they are completely automatic – no interaction with a user is required. A good example is the recent "Code Red Worm". The Code Red Worm starts on one infected host, scanning 100 systems simultaneously, looking for other computers with the Internet Information Server/Index Server Buffer Overflow vulnerability. When a vulnerable target is found, it immediately and automatically infects the new host with the code. The newly infected host starts this process all over again. Each infected host will attempt to infect 100 more hosts. In less than five days in July 2001, this worm had spread to over 196,000 hosts.

Premises

Security is a process, not something you can buy in a shrink-wrapped box. A corollary: **Security is never an absolute quantity**. It is a moving target – as software and hardware development continue, and as new products emerge (with new bugs), hackers will seek those vulnerabilities, and discover new and innovative ways of exploiting them. As new products are developed and sent to market to counter hacker attacks, hackers discover new ways to get around those products and keep hacking. It is an arms race, and one you need to be prepared to win.

Another premise is this: **Effective security is Security-in-Depth**. How many locks do you have on your front door? Just one? Or one for the doorknob, a deadbolt, and a chain? Do you have an alarm system as well? How about a bat by the bed? If you have all of this, then you already know what 'Defense-in-Depth' means. Network security is no different. Having a NetScreen firewall protecting your front door is a good start to an overall effective network security system. However, it is the whole system working together that makes you more secure.

Our final premise is this: **If you don't know what you are protecting and why, you can't protect anything**. An audit needs to be done against your company's network to determine what's important, what's not, and why. A small company could complete such an audit in five minutes for free, while an audit of a multi-national conglomerate could take months and millions of dollars. What it really boils down to is this: How important is this system? How likely will it be compromised with its current configuration? And finally, can the company afford a compromise of this system?

Practice

Audit It!

Our first step begins with the audit. Decisions need to be made about every resource connected to the network:

How important is it? Consider all political, technological, financial, and privacy issues. This relates to the content of the system. A credit card or medical records database would be very important, while a desktop computer used mostly for playing games might not be.

How vulnerable is it? This relates to how easy it is to compromise a server in a default (out of the box) configuration, or how difficult it is to properly harden. This also relates to how often new vulnerabilities are announced. More popular operating systems and services (i.e. Microsoft™ Windows™ and Internet Information Server™) generally get more attention by hackers, increasing the likelihood of a vulnerability being discovered.

How expensive is it to replace? This relates to the time, hassle, and expense of rebuilding or replacing a compromised system. This can also relate to initial installation costs.

Based upon the results of these questions, a matrix can be made of every network device. The examples provided vary from site to site – For example, a Macintosh computer would typically be considered more critical/valuable at a graphic design company than it would be at a finance company.

Important	Vulnerable	Expensive	Resulting Suggested Security Profile	Examples
High	High	High	Trusted/Management Area, double-authentication, no outward or inward access	Windows NT/2000 Servers
Low	High	High	Trusted Area, single-authentication, no inward access	Printers/Webcams
High	Low	High	DMZ Area, double-authentication	Solaris
Low	Low	High	Any Area, double-authentication	Routers
High	High	Low	Trusted Area, no inward access	Windows NT/2000
Low	High	Low	Trusted Area, no inward access	Windows 95/98/Me
High	Low	Low	DMZ Area, single-authentication	Linux/FreeBSD
Low	Low	Low	Untrusted Area, unrestricted access	Macintosh

Once we have our matrix completed, we can see groups of systems come together that can be treated the same way. More decisions need to be made, as you may not want or be able to put all of these similar machines in the same places. However, they can still be treated the same. For the purposes of these recommendations, double-authentication could mean anywhere from two usernames and passwords for entry, or an encryption protocol and a password, or even a one-time-use (i.e. SecurID token) password and a point-to-point permit for access.

Make some notes on the suggestions presented in the matrix:

1. If a system is vulnerable, that system needs access restrictions placed on that system, especially inward access. Once an attacker gets a foothold inside your network, other systems not normally accessible to the Internet could be attacked.
2. If a system is important, multiple authentication systems protecting access to that system are warranted.
3. If a system is expensive to repair or replace, maximize its security to protect your investment.
4. If a system is not considered important, make sure more important systems get the proper security they need. These unimportant systems should be considered expendable.

Partition It!

Next we need to lay out our security zones. These zones will use the groups developed during the audit, with modifications for geographical, departmental, financial, or political constraints. These zones, combined with monitoring systems you install in them, give you “Defense-in-Depth” – no one failure will compromise your entire network.

Subnets work best to separate different security zones, with either stateful inspecting or packet filtering firewalls dividing the zones. NetScreen Virtual Systems (found on our NS-500 and NS-1000 models) are ideal for this task. Another good idea is to put similar networks on similar switches divided by VLANs – Untrusted subnets on the Untrusted switch, DMZ subnets on the DMZ switch, and Trusted subnets on the Trusted switch (this will more than likely be your biggest switch, or group of switches). Cables from these switches should be going to firewalls, either in route mode or NAT mode.

Furthermore, your DMZ area can be grouped into subnets by function – placing Bastion Host FTP servers on one subnet and Bastion Host web servers on a different subnet is a good way to assure that any new FTP exploit you may have missed will not take down your web servers. Putting a firewall between these subnets is a good way to keep them protected from each other.

One subnet not mentioned until now is your Administration/Management LAN. This is your “Holy of Holies” – your inner sanctum of networking. Access to and from this LAN should be extremely protected – point-to-point policies on specific protocols only -- preferably encrypted. This is where a majority of your support infrastructure will reside – the management side of your NIDS, your authentication servers, configuration servers, and logging servers. Outside (Internet) access to this LAN, whether inbound or outbound, should be denied.

Fix It!

Many systems come out of the box in a very insecure state. Service packs, Hotfixes, and Patches must be applied and kept current if you expect your systems to stay safe. This is known as “**Hardening**” a system – making it difficult to penetrate. A system running a service with a well-known vulnerability (i.e. Code Red/IIS Index Server vulnerability) is disaster waiting to happen, even if the server is behind a firewall. In all probability, if an attack comes in on an appropriate port the server is supposed to be listening on, the firewall will not notice the attack come in.

By turning off unneeded services, updating needed services with patches against vulnerabilities, and placing these systems on purpose-specific subnets, you can make these systems very secure. Consider this: Should a security analyst be concerned seeing FTP traffic coming from a subnet that is only supposed to contain Web servers? Adding a firewall rule to enforce that concept would be even more effective.

Monitor It!

Once your structure is in place, it is time to add the good stuff. A firewall alone will not effectively answer all of your security needs. Your firewall, as powerful as it may be, needs support to run at its peak. This support comes in the form of secure (Bastion Host) system loggers (syslog servers), Network Intrusion Detection Systems (NIDS), authentication servers (RADIUS, TACACS+, SecurID/ACE, LDAP, etc.) and, in some extreme cases, Honeypots. Primaries of all these systems should reside in your Management LAN, where the company’s Network Security Manager and the Network Security Team keep track of the health of your network. Secondary systems should reside on other Trusted segments.

NIDS ideally would be placed on every subnet, but in practice should be placed at a minimum at ingress points to major networks such as in front of the Untrusted firewall, in front of the Trusted area, and on the major access point to the DMZ subnets. NIDS should be configured with at least two network interface cards (NICs), with one NIC placed on the subnet it is monitoring (but NOT given an IP Address), and the other NIC placed on your Management LAN with an IP address assigned. All out-of-band management should also reside on this ultra-trusted LAN.

Protect It!

Before something bad happens to your network, you need to formulate an Action Plan. Lists of phone numbers and addresses of important people should be in handy and well known locations at your Network Operations Center (NOC), along with instructions for when and who to call when something happens. Your 24x7 NOC personnel (or equivalent on-call group) should be trained to handle immediate response to attacks and “Hold the Fort” until more experienced help arrives.

Check It!

Now that everything is set up, you need to put on your “Black Hat” and look at your network from the perspective of an attacker. If you are not the official Network Security Administrator, you might want to save this task for him or her. In addition, that person needs to check with management before proceeding – a misunderstanding of what is to be done and by who can have legal consequences. Ask before attacking your own network! Once the appropriate management personnel are informed and approving, and all other legal issues are cleared out, it’s time to get evil!

Download some common, non-destructive hacking tools (see the Reference - Tools section for good ideas), and perform spot checks with port scanners and vulnerability scanners to see how well your design holds up. Things to check:

1. Are you seeing the services you expected to see? (This is good)
2. Are you seeing the protected services you should not be seeing? (This is bad)
3. How are your logging and alert systems working?
4. Did your NIDS pick up the appropriate attacks at places you would expect it to, or did the scan penetrate deeper than expected?
5. Did you scare the person in the NOC? You should have.

Update It!

Back to our first premise – Security is a process – a continual process of periodic review and updating as changing conditions dictate. New vulnerabilities, attacks and attack types come out, systems need patching, new security products come on the market – keeping abreast of all of these can easily be a full time job. Be sure to budget for the time and cost to maintain your security system. If you ignore your company’s security long enough, it **will** go away.

Summary

This paper has discussed several important considerations for designing a secure network. Specifically:

1. Security is a process.
2. Effective security is Security-in-Depth.
3. An audit is needed to determine the importance of assets in your network.

The paper also covered the 7 steps to a more secure network design:

1. Audit – Determine what is important and why.
2. Partition – Separate the important from the unimportant.
3. Fix – Make your default-configured systems more secure.
4. Monitor – Add monitoring and logging systems to round out your security.

5. Protect – Make a plan, and a contact sheet for when attacks happen.
6. Check – Do a dry run – and attack your own network.
7. Update – Keep current, and re-evaluate your design as things change.

All of this should be taken in the context that a company's business is its lifeblood, and if there is a conflict between business opportunities and security, security will lose. Application of even a few of these principles will take you far in securing your network from those who would rob you of your company's most valuable data. The most important step, the magic of security, is to design your network securely while empowering your business, not hindering it. Finding that balance is beyond the scope of this Whitepaper, and varies from customer to customer. A truly effective security plan is transparent to the network user doing the right thing, and very apparent to the user doing the wrong thing. Using NetScreen network security devices is the cornerstone to that effective security plan.

Reference

The Reference section has been divided into three major groups: Online, Offline (hardcopy), and Tools. The Online section has been subdivided into three categories: Whitehat, Greyhat, and Blackhat. Most of the online sites have cross-links to other similar sites. Additionally, Google has a great index of network security sites at: http://directory.google.com/Top/Computers/Security/News/

Online - Whitehat:

SANS Institute
http://www.sans.org
The SANS (System Administration, Networking, and Security) Institute is a cooperative research and education organization comprised of system administrators, security professionals, and network administrators. They have great network security courses, as well as certification in network security.

CERT Coordination Center
http://www.cert.org
Federally funded network security research organization out of Carnegie Mellon University.

Security Focus
http://www.securityfocus.com
For-profit, private organization with a lot of security information available for free. Great email newsletters.

NT Bugtraq
http://www.ntbugtraq.com
Non-profit, non-Microsoft™ organization dedicated to detecting and tracking bugs (and vulnerabilities) in Microsoft products. This site also has great email newsletters.

Razor Bindview™
http://razor.bindview.com
Razor is another great site for tools and advisories. It is a subdivision of Bindview Corporation – a security management services and products company.

Online - Greyhat:

Packetstorm
http://packetstormsecurity.org

Packetstorm is one of the many 'greyhat' sites included because it has such great potential for ultimate good or ultimate evil. Source code as well as compiled binaries for very malicious programs can be obtained from this site – enabling both the network administrator and the hacker alike.

Packet Nexus
http://www.packetnexus.com

A good 'hub' website with links and content to a whole world of network security – both light side and dark side.

DefCon
http://www.defcon.org

World's largest hacking party and conference held every year in July in Las Vegas, Nevada. Organized by "The Dark Tangent" (AKA Jeff Moss), it has a legitimate half-brother – "The Blackhat Briefs" held at Caesars' Palace the two days before DefCon. The year 2001 saw the 9th annual DefCon convention at the Alexis Park Hotel. See http://www.defconpics.org for pictures from that event. Talks given over the three-day conference range from the "Newbie" group through "General" all the way up to "Uber Haxor" levels. Cost is $50 for full access to the three-day seminar. Wireless 802.11b and 10/100BaseT wired network access to the Internet is also provided through the "Capture the Flag" network (CTFnet) – a free-for-all competitive hacking zone, where points are given for style, ingenuity, and difficulty of a particular hack.

Online - Blackhat:

Rain Forest Puppy
http://www.wiretrip.net/rfp/7/index.asp

RFP is one of those awkward Greyhat/Blackhat sort of individuals that is hard to categorize. He has been placed here due to the fact that all his exploits are malicious – they crash systems. He is very adept at crashing Windows systems.

Cult of the Dead Cow (CDC)
http://www.cultdeadcow.com

CDC is one of the more prolific and socially outrageous hacker groups. Responsible for the Back Orifice and BO2K "Remote Administration" tools (AKA "Back Doors"), their presentations at hacking conventions are typically loud, raunchy, and crude. They also write very good code. Back Orifice and BO2K are among the most popular back door programs in the world.

2600
http://www.2600.com

An old school (phone phreaking days) monolith of the hacking underground. 2600 gets its name from the 2600 Hz tone used by old analog-controlled pay phones to go into 'maintenance mode', to include free phone calls anywhere in the world. A founding member of 2600 discovered that the free whistle given away in a box of Capt'n Crunch™ cereal made precisely the proper 2600 Hz frequency to enable the free call mode. His 'handle' (nickname) became Capt'n Crunch.

Offline/Hardcopy:

Hacking Exposed, 2nd Edition, by Joel Scambray, Stuart McClure, George Kurtz
Paperback - 703 pages 2nd edition (October 11, 2000)
McGraw-Hill Professional Publishing; ISBN: 0072127481

A great overview of hacking techniques, and concepts.

Network Intrusion Detection: An Analyst's Handbook, 2nd Edition, by Stephen Northcutt, Donald McLachlan, and Judy Novak.
Paperback - 450 pages 2nd edition (September 22, 2000)
New Riders Publishing; ISBN: 0735710082

Great book on the forensics of hacking: How to handle a break-in, how to keep one from happening.

Maximum Security, 3rd Edition, by Anonymous
Paperback - 864 pages 2nd edition (September 15, 1998)
Sams; ISBN: 0672313413
Paperback - 896 pages 3rd edition (May 17, 2001)
Sams; ISBN: 0672318717

In-depth details of exactly how to hack, different editions cover different techniques, get all of them to have complete coverage. These are big books!

The Process of Network Security, by Thomas Wadlow
Paperback - 304 pages 1st edition (April 15, 2000)
Addison-Wesley Pub Co; ISBN: 0201433176

Good book on security processes, policies, design.

Tools - Freeware

Ethereal
http://www.ethereal.com

Ethereal is a commercial-grade packet sniffer and analyzer freely available under the Gnu Public License (GPL). There are both Windows and Unix (Linux/Solaris/FreeBSD/etc) versions available. The source code is also available. This is a great tool for troubleshooting network problems, as well as analyzing hacker attacks.

Snort
http://www.snort.org

Snort is a easily configurable Network Intrusion Detection System (NIDS) freely available under the Gnu Public License (GPL). There are many freeware, shareware, and commercial add-ons for this very fast and flexible IDS. Snort runs on Unix variants (Linux, FreeBSD, Solaris, etc).

Nmap
http://www.insecure.org/nmap/index.html

Nmap is a powerful port-scanning program used by Whitehats and Blackhats alike. A very informative program, it tells you which ports a system (or subnet(s) of systems) is listening on. Certain Blackhat features include: IP Spoofing, Stealth Scanning, Christmas-Tree Scanning, and Decoys. It is freely available under the Gnu Public License (GPL).

Nessus
http://www.nessus.org

Nessus is an incredible commercial-grade vulnerability scanner also freely available under the Gnu Public License (GPL). Nessus can use Nmap to further probe networks for holes. Nessus can selectively scan for over 675 (and growing) known security problems. The resulting reports are organized by host, categorized by severity, and can be exported in a variety of formats, to include a very slick cross-linked HTML including pie charts. Links to fixes for known security problems are included.

Legal Notice

This paper has made references to many websites ("Linked Sites"). The Linked Sites are not under the control of NetScreen and NetScreen is not responsible for the contents of any Linked Site, including without limitation any link contained in a Linked Site, or any changes or updates to a Linked Site. NetScreen is not responsible for web casting or any other form of transmission received from any Linked Site nor is NetScreen responsible if the Linked Site is not working appropriately. NetScreen is providing these links to you only as a convenience as a reference to this paper, and the inclusion of any link does not imply endorsement by NetScreen of the site or any association with its operators. You are responsible for viewing and abiding by the privacy statements and terms of use posted at the Linked Sites.

Furthermore, certain tools and programs may be illegal to use in certain countries, or on networks you are not responsible for. NetScreen will not be held liable for any activities you learn about from this document or its suggested sources. This paper is for educational purposes only.