

A Methodology for Secure Network Design

Cecilia Zanni¹ (cecilia.zanni@lsis.org) – Martin Silva² (msilva@prmarg.org)

Introduction

The Internet has undoubtedly become the largest public data network; enabling and facilitating both personal and business communications worldwide. The volume moving over the Internet, as well as corporate networks; is expanding exponentially every day. More and more communication is taking place via e-mail; mobile workers, telecommuters, and branch offices are using the Internet to remotely connect to their corporate networks; and commercial transactions completed over the Internet, via de World Wide Web, now account for large portions of corporate revenue [Cis01].

E-business applications such as e-commerce, supply-chain management, and remote access allow companies to streamline processes, lower operating costs, and increase customer satisfaction. Such applications require mission-critical networks that accommodate voice, video and data traffic, and these networks must be scalable to support increasing numbers of users and the need for greater capacity and performance. However, as networks enable more and more applications and are available to more and more users, they become ever more vulnerable to a wider range of security threats. To combat those threats and ensure that e-business transactions are not compromised, security technology must play a major role in today's networks [Cis03a].

The main premise of this work is that even the most efficient and powerful tool can be rendered useless by poor implementation [Net01]. Because of that, it will be intended to describe an in-depth approach to secure network design, focused on the expected threats and methods mitigation, as we think security is in the basis of network design, even prior to decision-making on the buying of new or most costly equipment.

Premises

[Net01] states the following:

1. Security is a process, not something that can be bought in a shrink-wrapped box:
2. Effective security is Security-in-Depth
3. If you don't know what you are protecting and why, you can't protect it

Practice

¹ LSIS – UMR CNRS 6168 – Campus Universitaire de Saint Jérôme – 13397 Marseille Cedex 20 – France

² Programa Regional de Meteorología, CONICET (Consejo Nacional de Investigaciones Científicas y Técnicas) – Casilla de Correo 330 – (5500) Mendoza – Argentina

The objective of network security is to protect networks and their applications against attacks, ensuring information availability, confidentiality and integrity. When organizations design their network security architectures to meet this objective, they must consider a number of factors. Not all networks and their associated applications have the same risks of attacks or possible costs of repairing attack damages. Therefore, companies must perform cost-benefit analyses to evaluate the potential returns on investment for various network security technologies and components versus the opportunity costs of not implementing those items. In the process, enterprises should make sure to consider their network security implementations as competitive advantages that can attract customers, employees, and partners [Cis03a].

Security Policy

Usually, the primary prerequisite for implementing network security, and the driver for the security design process, is the security policy. A security policy is a formal statement, supported by a company's highest levels of management, regarding the rules by which employees who have access to any corporate resource abide. The security policy should address two main issues: the security requirements as driven by the business needs of the organization, and the implementation guidelines regarding the available technology. In addressing these issues, the security policy typically includes several elements. For example, the security policy usually includes an authentication policy that defines the levels of passwords and rights required for each type of user (corporate, remote, dial-in, VPN, administrators, and so forth). Because business requirements and security technologies are always evolving, the security policy should be a living document that is updated regularly (at least once per year) [Cis03a].

Security Architecture

The security architecture should be developed by both the network design and the IT security teams. It is typically integrated into the existing enterprise network and is dependent on the IT services that are offered through the network infrastructure. The access and security requirements of each IT service should be defined before the network is divided into modules with clearly identified trust levels. Each module can be treated separately and assigned a different security model. The goal is to have layers of security so that a "successful" intruder's access is constrained to a limited part of the network. Just as the bulkhead design in a ship can contain a leak so that the entire ship does not sink, the layered security design limits the damage a security breach has on the health of the entire network. In addition, the architecture should define common security services to be implemented across the network. Typical services include [Cis03a]:

- Password authentication, authorization, and accounting (AAA)
- Confidentiality provided by virtual private networks (VPNs)
- Access (trust model)
- Security monitoring by intrusion detection systems (IDSs)

After the key decisions have been made, the security architecture should be deployed in a phased format, addressing the most critical areas first [Cis03a].

Security Technologies

As noted earlier, network security design requires that corporations determine the level of implementation investment and the total cost of intrusion they can withstand. Then corporations must decide how to allocate their available network security budgets to adequately secure their networks. To ensure the most comprehensive level of protection possible, every network should include security components that address the following five aspects of network security [Cis03a].

Identity

Identity is the accurate and positive identification of network users, hosts, applications, services and resources. Identity mechanisms are important because they ensure that authorized users gain access to the enterprise computing resources they need, while unauthorized users are denied access [Cis03a].

Perimeter Security

Perimeter security solutions control access to critical network applications, data, and services so that only legitimate users and information can pass through the network. This access control is handled by routers and switches with access control lists (ACLs) and by dedicated firewall appliances. A firewall provides a barrier to traffic crossing a network's "perimeter" and permits only authorized traffic to pass, according to a predefined security policy. Complementary tools, including virus scanners and content filters, also help control network perimeters. Firewalls are generally the first security products that organizations deploy to improve their security postures [Cis03a].

Secure Connectivity

Companies must protect confidential information from eavesdropping or tampering during transmission. By implementing Virtual Private Networks (VPNs) enterprises can establish private, secure communications across a public network—usually the Internet—and extend their corporate networks to remote offices, mobile users, telecommuters, and extranet partners. Encryption technology ensures that messages traveling across a VPN cannot be intercepted or read by anyone other than the authorized recipient by using advanced mathematical algorithms to "scramble" messages and their attachments [Cis03a].

Security Monitoring

To ensure that their networks remain secure, companies should continuously monitor for attacks and regularly test the state of their security infrastructures. Network vulnerability scanners can proactively identify areas of weakness, and intrusion detection systems can monitor and reactively respond to security events as they occur.

Intrusion detection systems and vulnerability scanners provide an additional layer of network security. While firewalls permit or deny traffic based on source, destination, port, or other criteria, they do not actually analyze traffic for attacks or search the network for existing vulnerabilities. In addition, firewalls typically do not address the internal threat presented by "insiders." [Cis03a].

Security Policy Management

As networks grow in size and complexity, the requirement for centralized security policy management tools that can administer security elements is paramount. Sophisticated tools that can specify, manage, and audit the state of security policy through browser-based user interfaces enhance the usability and effectiveness of network security solutions [Cis03a].

Top Ten Security Tips

1. Encourage or require employees to choose passwords that are not obvious
2. Require employees to change passwords every 90 days
3. Make sure your virus protection subscription is current
4. Educate employees about the security risks of e-mail attachments
5. Implement a complete and comprehensive network security solution
6. Assess your security posture regularly
7. When an employee leaves a company; remove that employee's network access immediately
8. If you allow people to work from home, provide a secure, centrally managed server for remote traffic
9. Update your Web server software regularly
10. Do not run any unnecessary network services

These tips have been extracted from [Cis03a].

Implementation

1 - Audits

Decisions need to be made about every resource connected to the network [Net01].

- How important is it? All political, technological; financial, and privacy issues have to be considered. This relates to the content of the system.
- How vulnerable is it? This relates to how easy it is to compromise a server in a default configuration, or how difficult it is to properly harden it. This also relates to how often new vulnerabilities are announced. More popular operating systems and services 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 [Net01].

| 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 ³ 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 the matrix completed, groups of systems that come together can be treated the same way. More decisions are needed to be made, as these similar machines may be put or not 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 .

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.
1. If a system is expensive to repair or replace, maximize its security to protect your investment.
2. 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.

³ DMZ (Demilitarized Zone): 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 be considered expendable.

2 – Network segmentation

Next, security zones need to be laid out. These zones will use the groups developed during the audit, with modifications for geographical, departmental, financial, or political constraints. These zones, combined with installed monitoring systems, give “Defense-in-Depth” – no one failure will compromise the entire network [Net01].

Subnets work best to separate different security zones, with either stateful inspecting⁴ or packet filtering⁵ firewalls dividing the zones. 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 the biggest switch, or group of switches) [Net01].

Cables from these switches should be going to firewalls, either in route mode or NAT mode.

Furthermore, the DMZ area can be grouped into subnets by function – placing Bastion Host⁶ 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 that may have been missed will not take down the web servers. Putting a firewall between these subnets is a good way to keep them protected from each other [Net01].

One subnet not mentioned until now is the Administration/Management LAN. 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 the support infrastructure will reside – the management side of the IDS⁷, the authentication servers, configuration servers, and logging servers. Outside (Internet) access to this LAN, whether inbound or outbound, should be denied [Net01].

Figure 1 shows a typical network system architecture.

⁴ Stateful Inspection: A firewall process that checks the TCP header for information on the session’s state. 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.

⁵ 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.

⁶ Bastion Host: A hardened server configured with the minimal software to support a single network service. A loss of any one Bastion Host should not compromise the security of any other host.

⁷ IDS (Intrusion Detection System): A system that passively monitors all data flowing past its network interface, looking for hostile data patterns. A flexible IDS will have user-definable data patterns or “signatures”, which allow it to detect new types of attacks as they evolve. An IDS 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 IDS.

Network System Architecture

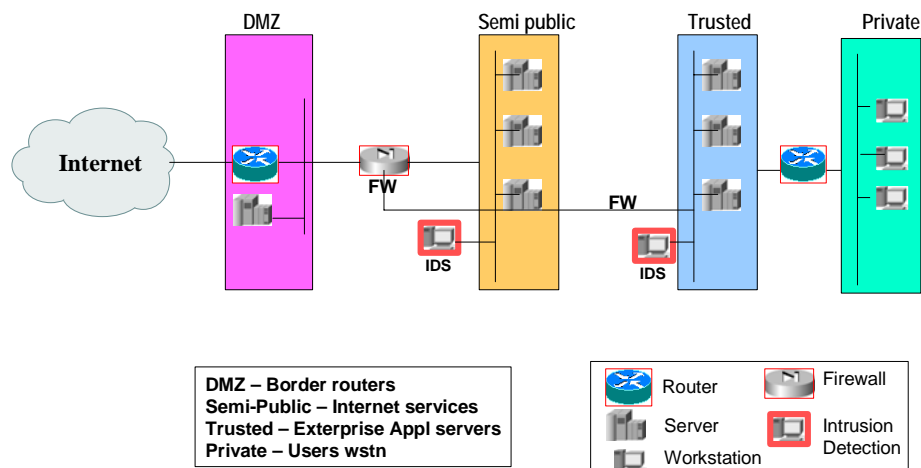


Figure 1

3 - Hardening

Many systems come out of the box in a very insecure state. Service packs, Hotfixes, and Patches must be applied and kept current if systems are expected 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 [Net01].

By turning off unneeded services, updating needed services with patches against vulnerabilities, and placing these systems on purpose-specific subnets, these systems can be made very secure [Net01].

4 - Monitor

A firewall alone will not effectively answer all of the security needs. The 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), Intrusion Detection Systems (IDS), authentication servers (RADIUS, TACACS+, SecurID/ACE, LDAP, etc.) and, in some extreme cases, Honeypots⁸. Primaries of all these systems should reside in the

⁸ 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.

Management LAN, where the company's Network Security Manager and the Network Security Team keep track of the health of the network. Secondary systems should reside on other trusted segments.

IDSs 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. IDSs 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 the Management LAN with an IP address assigned. All out-of-band management should also reside on this ultra-trusted LAN.

Intrusion Detection Systems

IDS⁹ are a topic that has recently garnered much interest in the computer security community. The need for some element that could alert and inform administrators about something strange in near real time resulted in intrusion detection systems. When presented with different types of IDS one might be tempted to assume that one approach or another was inherently superior. The mixture of approaches used for IDS offers the security analyst a unique opportunity in terms of the synergies inherent in combined techniques, by mixing strengths of existing approaches cancel out many of their weaknesses to produce an extremely reliable IDS capability. Basically there are three major approaches: the host-based intrusion detection system (HIDS), the network-based intrusion system (NIDS) and the new concept of hybrid-IDS [Ran01].

Host-based intrusion detection systems

Historically, this was the first IDS, based on academic research in the late 1980's [Den87]. Typically, they reside within the host operating environment.

Early HIDS only dealt with data that had already been collected for them by existing audit facilities. The current generation HIDS are more aggressive about collecting information – they install agents that monitor system processes, check and watch registry entries, observe who is logged in and where they logged in from, and collect data about what programs are running, including their CPU usage and file accesses [Ran01].

HIDS are part of the target and are therefore able to give very good information about the state of the system during an attack, but since this fact, any information it provides becomes suspect immediately when an attack succeeds against the target. More than that, logs the HIDS relies on may be altered or deleted.

On the other side, HIDS will have difficulty detecting attacks that completely wipe out the target system. When the operating system is crashed, the HIDS has crashed along with it and no alert is generated [Ran01].

⁹ Typical widely used tools for IDS include Tripwire (<http://www.tripwire.com/>, <http://www.tripwire.org/>), PortSentry, Swatch, Snort (<http://www.snort.org/>) and Prelude (<http://www.prelude-ids.org/>).

Network-based intrusion detection systems

A NIDS operates by accessing traffic off a network broadcast medium. With an ethernet-type network, this entails placing the interface card into « promiscuous » mode - a mode in which the card collects all the traffic crossing the network segment whether or not it is destined to the system which is listening. The listening system, the NIDS, then tries to detect attack patterns within the collected traffic, treating the network traffic as data to be examined passively. The NIDS approach has a number of attractive properties, which have made them the predominant IDS technique in the past years: NIDS are installed per network segment rather than per host. NIDS deal with traffic as abstract data; a denial of service or «death packet» which might crash a target host will not affect the NIDS. NIDS have a few areas in which they are weak when compared to a HIDS: NIDS may miss packets due to congestion on the network link that they are monitoring; NIDS do not have a good notion of user identity; NIDS' weaknesses primarily have to do with their ability to understand what is going on within the host: who the user is, how the host is interpreting the attacks as seen, and whether the attack worked on the host. By itself a NIDS is still a valuable tool, but a sophisticated attacker might be able to exploit its shortcomings to mask their actions [Ran01].

Hybrid intrusion detection systems

These are an interesting blend of the strengths and weaknesses of HIDS and NIDS. They operate much like a NIDS – they collect traffic at a packet level, process it, and detect or deflect attacks. But, like a HIDS, they do it on a per-host basis. They have most of the advantages of a NIDS except for in the area of deployment, since they must be deployed on every host in order to function. Additionally, since they act as a shim, they may interfere with other applications that shim the TCP/IP stack in the operating system – firewalls and VPNs may not function correctly with shim-type IDS. Another popular form of hybrid is NIDS/HIDS that have been designed to work together. Such cooperative IDS allow a great deal of flexibility in selecting where and how to cover analysis within a network and across hosts [Ran01].

5 - Protect

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

6 - Check

Now that everything is set up, it is necessary to look at the network from the perspective of an attacker. In addition, the person in charge needs to check with management before proceeding – a misunderstanding of what is to be done and by who can have legal consequences. Once the appropriate management personnel are informed and approving, and all other legal issues are cleared out, it's time to download some

common, non-destructive hacking tools, and perform spot checks with port scanners¹⁰ and vulnerability scanners to see how well the design holds up.

Things to check:

1. Are the expected services visible?
2. Are the protected services visible?
3. How are the logging and alert systems working?
4. Did the IDS pick up the appropriate attacks at the expected places, or did the scan penetrate deeper than expected?
5. Was the person in the NOC concerned?

7 - Update

Back to the 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. A budget for the time and cost to maintain the security system is needed.

Conclusions

Although these steps can successfully be used to create a secure environment, they are not the only factor for an optimum network security.

To ensure that their networks remain secure, organizations should continuously monitor for attacks and regularly test the state of their security infrastructures. A firewall alone will not effectively answer all of the security needs, as now crackers intend to attack the target networks using the vulnerabilities of available services. That's why IDS strategically located are of crucial need.

An awareness of the importance of security and accountability within the organization should be created. Establishing good security policy, staying up to date on the latest development in the hacker and security communities, maintaining and monitoring all system with sound system administration practices are among the heart of best practices in network security.

¹⁰ Port scanning: A technique for determining which ports a server is listening to. A port-scanning program 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.

References

- [Cis01] Cisco Systems
A Beginner's Guide to Network Security
2001
- [Cis03a] Cisco Systems
Network Security: An Executive Overview
http://www.cisco.com/warp/public/cc/so/neso/netsp_pl.htm
- [Cis03b] Cisco Systems
Action Steps for Improving Information Security
http://www.cisco.com/warp/public/cc/so/neso/saso/roi5_wp.htm
- [Den87] Denning, Dorothy
An Intrusion-Detection Model
IEEE Transactions on Software Engineering
February 1987
- [Net01] NetScreen Technologies
Principles of Secure Network Design
2001
- [Ran01] Ranum, Marcus J.
Coverage in Intrusion Detection Systems
NFR Security Inc Technical Publications
June 2001