Chiara Paci

Sans Security Essential. GSEC Practical Assignment Version 1.3

.

Abstract *My first firewall was a little one, with only two networks and three-four rules. Then, I built more complex firewalls, with many networks and rules. This paper illustrates my firewall planning guideline: it is a questions list, with explication and guidance for their answers. The questions are grouped into three sections:*

1. ***Who are you?*** *This part help you to collect information about the actual state of network.*
2. ***What do you want?*** *This part help you to write down policies.*
3. ***What do you need?*** *This part help you to design firewall architecture.*

I am writing about firewalls, because currently my work is building firewalls. Nevertheless, I think this is a good example of applying general security rules to a particular case.

Perimeter defense planning: a guideline

> You can think of it two ways. Either there are no bad firewalls, only good firewalls used in silly ways, or there are no good firewalls, only bad firewalls used in places where their weaknesses are acceptable. Either way, the trick is to match the firewall to the need [Zwicky, p. 162].

1 Who are you?

Your security strategy must be proportionate to your business. Different companies need different security levels. This can be true also for different departments or systems on the same company.

Different companies need different security strategies too. You must know how people work and how the company (or better, the network) interoperate with other networks.

> You are not looking for the perfect firewall; you are looking for the firewall that best solves your particular problem. [...] On a large network, the best solution will almost always involve a combination of technologies. On a small network, the best solution may well involve something that's said to be "insecure" or "low performance" or "unmaintainable" [Zwicky, p.162].

1.1 What are your front doors?

It's important that you know your front doors, because you are trying to build a perimetrical protection. This may be not so simple as you can think.

A front door is any point of connection between your network and external world: in a network, you can find several kinds of front door: Internet accesses, direct connections between private network, modems and so on.

a. **Internet accesses** are easy to detect. Generally, you have only one, well known, Internet access point. Some companies have more Internet access points, but all are well known: Internet accesses cost time and money, and need some bureaucratic and official stuff, like a registering at a Registration

Authority.

b. A **direct connection between two networks** is expensive, but it does not require an official registration and is more simple to obtain and configure than an Internet connection. Network connections can be made through a VPN: in this case, it's also cheaper than an Internet connection. Generally, you know if you have some direct connected networks, but only appropriate policies can assure that you have only known connections.
c. **Modem connections** are more hardly detected and more dangerous. At this moment, it's simpler to buy a computer equipped with a 56kbps modem than without it. In any case, a modem is cheap and easy to obtain. Only appropriate policies can avoid or regulate this type of connection.
d. **Other types** of connection are hidden or indirect. Good examples are VPN (client from your network and server in a remote network) and file sharing systems (like Gnutella). Both systems can create holes and access points in your network. Discovering them is hard, but you can reduce them with appropriate policies and firewall configuration.

1.2 What is your resources value?

Starting from the same principles - i.e., you have to guarantee confidentiality, integrity and availability - the consequences may be quite different. A golf club web site needs less security structures than a bank network and the golf club has less money to spend than a bank. At the same time, both the golf club web site may have restricted information (associates personal data) and the bank may have unrestricted information (bank description web site).

It's important that you estimate the actual value of your resources and information. The worst thing you could do is overestimating or underestimating information value. Obviously, overestimating it's better than underestimating, but overestimating can be dangerous too. Security systems cost time and money, slow down connections and need management. The golf club web site does not need a complex system with multiple firewall levels, external log collection and printing, dedicated intrusion detection cluster, honeypot network and centralized management console (all on high availability systems). A more more simple firewall with intrusion detection software can be sufficient, with or without a DMZ, depending on the size and structure of web site. But the complex system is the minimum requirement for the bank network and often you must use several such a system, one for each bank department.

a. **Are some computers or systems more confidential than others?** For example, authentication or accounting databases or classified projects servers.
b. **Are some computers or systems less confidential than others?** For example, external web servers or public ftp sites.
c. **How many classifications levels are there?** Some companies have only two levels: public access systems (web, mail servers) and internal network. Others have three levels: public access systems, internal network and systems protected from internal network. Others have even more levels: for example, mail and web server are less confidential than internal network, but mail server information are more confidential than web server information.

1.3 How do people work?

Security is the medium, not the goal (this is very hard to remember, I know). You must taylor your rules to the environment. The best computer for pure security is an off computer. Even better no computer at all. But this does not work, because people have some difficulties using an off or non-existent computer.

If your company network is connected with many other networks and people use these connections in their work, it could be a bad idea closing them because it's more secure to have a single access point. It could be a better idea to secure all connections, but leave them in place. In other words, you need a trade-off between security and operability.

This does not mean that you have to adjust all the rules to the environment. You must change something into environment if you are making a security system. For example, you must close Internet accesses to the bank network with ftp and telnet, also if employees use these services to connect at the bank network from home.

It depends on your business. On the golf club web site, you can leave ftp access from Internet, if golf club editors use them for updates. You have to secure it, with patches, access control lists, strong passwords and so on, but you don't really need to stop it. On the bank network, an ftp access from Internet can be a serious hazard. If you really need to access the bank network from Internet, it's better to use a VPN.

1.4 Who do you trust and who you don't?

You must split people in many groups. The first group is the more trusted. They can access to data at the higher security level. The last group is the less trusted. They can't access to anything. You can have one or more groups lists.

You have to split people which pass through your perimeter into groups. You can start dividing the world into two groups: *internal network* and *others*. Group *internal network* is the most trusted, group *others* is the least trusted. But even the golf club web site can need another level: editors want to connect to the web site from their homes to update with ftp. Therefore you create another group, *editors*, which is more trusted than *others*, but less trusted than *internal network*. This is a simple, but common scenario.

The bank network is more complex. The bank has several customers, several suppliers and also several employees. They may need to connect to the bank network, from other networks or from home. Furthermore, the bank does not trust all the people at the same level and for the same data. Therefore you need more groups and groups lists.

1.5 What are acceptable risks?

At this moment, you know your resources and their value and also how your company use them. In other words, you know what is important for your company. As usual, you can have one or more importance levels.

You have three possible strategies about risks: you can accept, reduce or transfer.

a. **Accepting risk** means that the lost or damage of your resources costs less than placing a security system to protect them. This is more true if you have an effective disaster recovery procedure (strictly speaking, a disaster recovery procedure is a security measure). To the golf club web editors, closing ftp access costs more than restore the web site from backup, then they decide to accept the risk of running ftp service.
b. **Reducing risk** means that your resources are critical and you decide to protect them. Resources in the bank network have a great value and closing ftp costs less (in time, money and reputation) than restoring systems. Then the bank decide to close ftp accesses from Internet.
c. You **transfer risks** when you estimate that your data have a sufficient value to need protection, but you can not have a significant risk reduction. The bank may offer an on-line trading site to its customers. This is a risk, because through the web site someone can violate the system, but the bank can not avoid or reduce it. Instead, the bank assure all resources: if someone violate the system, the insurance refund the damage suffered by the bank.

2 What do you want?

In the previous section, you have collected information about your actual situation. Now you must decide how to modify this situation to obtain the wanted security degree.

At this moment, you have four principal lists about your network:

a. Resources, classified by value and restriction.
b. People groups, classified by trusting and needs.
c. Front doors.
d. Risk acceptance decisions.

Each front door involves a part of people and resources. Depending on risk acceptance, you need two access matrices for each front door, which illustrate the relationship between people and resources. Each front door has two possible traffic directions: from external network to internal and from internal network to external. In the first matrix, people are in external and resources in internal. In the second, people are in internal and resources in external. External resources can be, for example, the Internet connection, if you are billed on a traffic basis. You may also include that you don't want that your employees connect to file sharing systems, or you may state what external services are allowed (for example, only http and e-mail).

The golf club has a single front door (its Internet connection), four people groups (internal network, editors, associates and others) and four restriction levels (system, server data, personal associates information, other pages). Depending on decisions illustrated above, the golf club has the following access matrix of inbound connections. In this access matrix, the group *internal network* does not appear, because it does not access golf club web site through the front door.

<table>
<tr><th rowspan="3"></th><th colspan="7">Golf club web site - Accesses from Internet</th></tr>
<tr><th colspan="2">editors</th><th colspan="3">associates</th><th colspan="2">others</th></tr>
<tr><th>prot.</th><th>perm.</th><th colspan="2">prot.</th><th>perm.</th><th>prot.</th><th>perm.</th></tr>
<tr><th>system</th><td colspan="2">no access</td><td colspan="3">no access</td><td colspan="2">no access</td></tr>
<tr><th>web server data</th><td>ftp
http</td><td>read
write</td><td colspan="3">no access</td><td colspan="2">no access</td></tr>
<tr><th rowspan="2">personal associates data</th><td rowspan="2">ftp
http</td><td rowspan="2">read
write</td><td rowspan="2">http</td><td>owner</td><td>write</td><td rowspan="2" colspan="2">no access</td></tr>
<tr><td>all</td><td>read</td></tr>
<tr><th>other data</th><td>ftp
http</td><td>read
write</td><td>http</td><td colspan="2">read</td><td>http</td><td>read</td></tr>
</table>

Access matrices is the first step to planning policies. You need to write policies before selecting technology and architecture. This is the correct practice, because this way you can select the best product for your needs. At the end, you must update the policies by adjusting them to the selected architecture and product capabilities. In this step, you are writing the alpha version of your policies.

The policies document must be simple and short (one or two pages). If your document is longer, try to split policies in smaller parts.

2.1 What is the background?

This is an introduction and illustrate motivations and circumstances of your policies.

a. **Version and date**

   Just to keep them ordered and track changes.

b. **Why do you need these policies?**

   This section help reader to know what is the policies context.

c. **Do exist other versions of these policies?**

   If you are starting building your perimeter protection, you have not other perimeter policies. But if you are modifying the perimeter protection, you have old policies, possibly not written. This is an alpha version of the new project policies, not of whole perimeter policies.

d. **What are related documents?**

   You can put into this section any document that is related to policies and perimeter defense. You can have internal communications, budget documents, other related policies and procedures and any other paper that helps your reader to understand how to insert these policies into the company structure.

2.2 What is covered by these policies?

This section is very important: you define the areas covered by the policies. At the end of this part, your reader must know if these are the policies that apply to his situation or if he must read other documents, because this one does not cover his case.

a. **People**

   In this section you put the people groups lists which you have collected before, with trusting and needs.

b. **Resources**

   In this section you put the resources lists, with value and restriction level.

c. **Who are responsible of these policies?**

   This section includes two groups. The policies writers are in the first group: they can help the reader to understand the policies (i.e., system administrators). People in the second group are these that can authorize policies updates (i.e., managers).

d. **What are contacts and responsibilities outside company?**

   You need this part if you have connections with other networks. For example, if you have a direct connection between your network and a supplier network, in this section you insert the connection responsible in the supplier company.

2.3 What do policies state?

This section is the heart of document. You must remember that you are writing policies for a perimeter defense. Don't insert in this section all possible accesses to resources, but only the accesses through your front doors.

a. **What are the default policies?**

   You use these policies when any other policy does not apply. You can choose two default policy types:

   deny all;
   permit all.

   You have several default policies, at least two for each front door: a default policy about inbound connections and one about outbound connections.

b. **What are general principles?**

   Into this section you put any general rules that you apply. For example, you may state that another network must have appropriate security measures in order to authorize a direct connection between your network and the other network. This applies to all connections and is independent from the particular permission you assign to the remote network owners.

c. **What are rules?**

   Her you put the access matrices you made before. For each front door, you write two rules list, specifying who can access what, when he can do it, how and where. You can also add some notes about risks that each rule involves. This can help you and your reader to understand how to update policies to increase security.

d. **What are policies updating rules?**

   This section states who can change the policies and when you need to release another version of the policies. For example, if you have a stable connections list, you may state that you need a new policies version whenever you add another connection. Otherwise, if you have a dynamic connections list, but with a stable list of connections types, you may state that you need a new policies version only when you add another connection type and not when you only add a connection of a known type.

e. **How do you handle violations?**

   This illustrate how your company prosecutes policies violations. It depends on business, resources value and violation gravity.

3 What do you need?

In the previous sections you have analyzed your company and your network and you have decided what your security system must do. The last step is designing your perimeter defense.

3.1 What is your firewall?

A perimeter defense is based on a firewall.

> Perhaps it is best to describe first what a firewall is not: a firewall is not simply a router, host system, or collection of systems that provides security to a network. Rather, a firewall is an approach to security; it helps implement a larger security policy that defines the services and access to be permitted, and it is an implementation of that policy in terms of a network configuration, one or more host systems and routers, and other security measures such as advanced authentication in place of static passwords. The main purpose of a firewall system is to control access to or from a protected network (i.e., a site). It implements a network access policy by forcing connections to pass through the firewall, where they can be examined and evaluated [Wack].

a. **Packet filter**

   A packet filter is the most simple firewall. It consider only single packets: if a packet matches some characteristics, it forwards packet, else it drops packet. A packet filter is fast, but not very effective [Guttman].

b. **State-full inspection packet filter**

   A state-full inspection packet filter is an evolution of the packet filter. Instead of single packets, it considers entire sessions, from the first SYN packet to the last FIN packet. It is more effective than a simple packet filter and, provided with sufficient memory and cpu speed, it is as fast as a simple packet filter [Cronje].

c. **Circuit-level gateway**

   A circuit-level gateway is similar to the packet filter. Instead of simple filtering, it runs an application that monitors the connections. This application knows the communication protocol (for example, http) and validates the commands instead of packets [Cronje].

d. **Application proxy**

   An application proxy is the most effective tool, but it is slower and less flexible than a packet filter. It does not forward traffic, it forwards requests. A program, that runs on the firewall, examines external requests. If a request is legitimate and correct, then the program forwards the request to the appropriate internal server. This program acts like a server with actual clients and like a client with actual servers. For this reason, you need a program for each service that you want to proxy [Guttman].

3.2 What is your architecture?

You have some basic perimeter architectures and you may combine them to build your own perimeter. Your limits are budget and fantasy.

a. **Multi-homed host**

   This is a firewall with more than one network interface. Each interface is connected to a network segment. In the most common situation, you have a dual-homed host with an interface connected to the internal network and the other connected to Internet. The firewall then acts like an intermediary [Guttman, Zwicky p. 122-126].

b. **Screened host**

In the screened host architecture all inbound connections are redirected to a special host, called a bastion host. The bastion host is the interface between the internal and the external network. The external network can access some data in the internal network, but only through the bastion host [Guttman, Zwicky p. 126-128].

c. **Screened subnet**

A screened subnet is a logical evolution of the screened host. This architecture adds another security layer by building a special network, called DMZ, where all bastion hosts reside. A DMZ (DeMilitarized Zone), or perimeter network, is a network added between external and internal network. By this way, the external network can not directly access resources into the internal network. A DMZ is an application of Defense-in-Depth principle:

> A DMZ is a glowing example of the Defense-in-Depth principle. The Defense-in-Depth principle states that no one thing, no two things, will ever provide total security. It states that the only way for a system to be reasonably secured is to consider every aspect of the systems existence and secure them all. A DMZ is a step towards defense in depth because it adds an extra layer of security beyond that of a single perimeter[Young].

In the implementation cost of a DMZ you must calculate also the performance degradation, but usually you have a little effect on performance and a significant increase in security. In a typical scenario, the external network is Internet and the DMZ contains web server or mail server [Young]. On a more complex situation you may need multiple DMZ. If you have a direct connection with another network, you can build a DMZ between two networks. You can place in this DMZ some common servers, like ftp or file servers [Zwicky, p.155].

3.3 What else should you need?

a. **Intrusion detection**

You must insert an intrusion detection system in your perimeter defense. This can be a simple program installed in your firewall or a complex system with multiple hosts. Intrusion detection systems can be based on:

> **anomalies detection**: if something not appears like a normal activity, then the IDS sends an alert to the administrator; this may happen, for example, when a user connects to system at the night time [Buonocore];
> **patterns**: if something matches a pattern, then the IDS sends an alert to the administrator; this may happen, for example, when someone is scanning your network [Buonocore].

Both anomaly detection and signature IDS based produce false positive alerts, i.e. they could send alert even for legitimate use. “However it is better to have an IDS produce a manageable number of such alerts then ignore possible intrusions” [Buonocore].

We have two types of IDS.

> **Network based IDS**: the IDS scans the network and analyzes all packets on the network. In a perimeter defense, the best place for network based IDS is on the firewall or on a separate host in front of or behind the firewall.

> In front of the firewall the IDS will record attacks that may be blocked by the configuration of the firewall but this data can be used to tune the IDS. Behind the firewall only attacks that got though the firewall would be reported and require attention [Buonocore].

You may add a network based IDS in each DMZ, but it could become useless if you have some protection against sniffers (like switches or encrypted traffic), because network IDS rely on sniffing.

A network based IDS is cheap, easy to manage and independent form hosts operating systems, but it slows down the network [Buonocore].

Host based IDS: it works on a single host and examines the host status (file system accesses, permissions changes, user logon, administration activity and so on). In a perimeter defense, the best place for network based IDS is on the interface bastion hosts (web, mail servers). Because an host based IDS works on single hosts, it can detect attacks also if you have switches or encryption and does not affect the network performance, but it is expensive and not easy managed [Buonocore].

Using both network based and host based IDS is an example of correct diversity of defense strategy.

b. **Honeypots**

> [...] we will define a honeypot as "*a resource who's value is in being attacked or compromised*". This means that whatever we designate as a honeypot, it is our expectation and goal to have the system probed, attacked, and potentially exploited. [...]
> A honeypot may be a system that merely emulates other systems or applications, creates a jailed environment, or may be a standard built system. Regardless of how you build and use the honeypot, it's value lies in the fact that it is attacked. [Spitzner]

You can group honeypot into two types:

production honeypots are built to take an attacker away from the real system and to increase security [Spitzner];
research honeypots are built to obtain information about threats and on how to defend against these threats [Spitzner].

Honeypots rarely can increase prevention, but can help on detection and reaction. In detection, because an IDS can miss new or unknown attacks, but an honeypot “happily capture any attacks thrown their way” [Spitzner]. In reaction, because an honeypot is a system useless and not used: then, if you have a violation, you can take off-line the honeypot and full analyze what happened (because it is useless), without users interference (because it is not used). Afterwards, you can apply the acquired knowledge to all compromised systems [Spitzner].

c. **Centralized management host**

If you have a complex system, with many firewalls, IDS hosts, honeypots and other things, you need a centralized management host. This host can collect logs, backup hosts, maintain configuration databases, monitor the system and so on.

d. **Support services**

> You should carefully examine all cases where the firewall is getting information from external machines, get rid of as many dependencies as possible, and move other services into the firewall wherever possible [Zwicky, p. 163].

Common support services are DNS and time synchronizations. A good place for these services is the centralized management host, if you have one. On a simpler situation, you may try to eliminate these services, for example using ip addresses instead than host names [Zwicky, p.163].

3.4 What is the best firewall?

a. **What firewall features do you need?**

At this point you know what is appropriate to your network, i.e. if you need a little simple system or a complex system. Any way, you can consider some features in order to select the more appropriate product.

> **Reliability and redundancy**: your firewall is often a critical place, then you may need an high availability system [Zwicky, p. 160].
> **Scalability**: if your company grows, you may need both that your firewall grows and add other perimeter defenses. If you think that your network will become larger, then you may prefer a modular system or a firewall that can be updated or expanded instead of changed [Zwicky, p. 160].
> **Adaptability**: if you have a non-standard network and you need some customization to your perimeter defense, then you may prefer a more adaptable firewall [Zwicky, p. 160].

b. **How do you manage the firewall?**

When you manage a firewall, you may prefer a tool because it is more secure than others (for example, ssh vs telnet), but often it's a matter of taste (for example, command line vs a secure graphical gui).

> **Configuration**: you need to configure your firewall depending on your network in order to use it. You may need to see the details of configuration, in order to debug or audit the configuration [Zwicky, p. 161]. You must verify if the firewall configuration capabilities meet your project decisions (for example, does the firewall need name resolution?).
> **Auditability**: it's very important that you can know what your firewall is doing. You may need only the predefined log levels, or you may need to configure the log levels. On some situations, you may need to collect logs in another host [Zwicky, p. 160-164].
> **Back up**: depending on firewall and on situations, you may need to backup either the entire firewall or its configuration only. You only need to save the configuration if you can quickly rebuild the firewall from the installation set, otherwise you need to backup the entire firewall. You must know how the backup is performed [Zwicky, p. 162].

c. **How much does the firewall cost?**

Like for any other system, a firewall cost can be split in four parts:

> hardware;
> software;
> support and upgrades;
> administration and installation.

Unlike other systems, you can buy either a software only solution or a complete solution, depending on supplier. If you have a software only solution, then you need to buy also hardware and operating system [Zwicky, p. 160-161].

4 References

1. [Walt] van der Walt, Charl, “Introduction to Security Policies, Part Two: Creating a Supportive Environment ”, September 24, 2001, URL: http://online.securityfocus.com/infocus/1473, (17 Mar 2002)
2. [Wack] Wack, John P. and Carnahan, Lisa J., “Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls”, February 9, 1995 URL: http://csrc.nist.gov/publications/nistpubs/800-10/main.html (16 Mar 2002).
3. [Zwicky] Zwicky, Elizabeth D. and Cooper, Simon and Chapman, D. Brent, Building Internet Firewalls, Sebastopol:O'Reilly & Associates, 2000
4. [Guttman] Guttman, Barbara and Bagwill, Robert “Internet Security Policy: A Technical Guide”, July 31, 1997, URL:http://csrc.nist.gov/isptg/html/ (16 Mar 2002)
5. [Cronje] Cronje, Gerhard, “Choosing The Best Firewall” , April 10, 2001, URL: http://rr.sans.org/firewall/best.php (16 Mar 2002)
6. [Young] Young, Scott, "Designing a DMZ", March 26, 2001, URL: http://rr.sans.org/firewall/DMZ.php (26 Feb 2002).
7. [Buonocore] Buonocore, Kathleen, “Selecting an Intrusion Detection System”, August 19, 2001, http://rr.sans.org/intrusion/selecting.php (16 Mar 2002).
8. [Spitzner] Spitzner, Lance, “Honeypots. Definitions and Value of Honeypots”, March 8, 2002, URL: http://www.enteract.com/~lspitz/honeypot.html (17 Mar 2002)