

# Chapter 7

## Infrastructure Security

The purpose of this chapter is to:

- Help you decide where to put those firewalls
- Help you decide how many of them to deploy
- Describe the other key components of infrastructure security that you need to make it look like you know what you're doing

## Anecdote

*I may be a violent, raging psycho—but I'm sensitive. And I always buy my wife a nice birthday present and something nice for Valentine's Day. My wife used to be an executive officer at the local council but gave it all up to teach kiddies arts and crafts (design and technology, they say—but whatever floats your boat!). When she was finishing her teaching course, we were watching a DIY program on TV where some bozo makes a £40 table out of £700 of wood, and she announces, "Buy me a router for Valentine's Day." I know she has no idea about money, so I didn't remark that the router would cost quite a bit more than the average card and chocolates.*

*At work, I recouped the losses by gaining full comic value from all the lads around. "Look at my wife's Valentine present!" I shouted as all the blokes pawed the JCB (nothing but the best) yellow monster while all the women flicked hair and tsch'ed to no effect. All the blokes wanted to be as Cool 'n Hard as me.*

*At home, I gave this beautiful router to She Who Must Be Obeyed—and got no response. Finally: "Oh," she said. In my best Tarzan voice, I said, "You want router, I get router." She screamed, "No! I read your latest research article—I want 802.11g, not 802.11b! I want 54Mbit connections!" Obviously, the light of my life wanted a new wireless broadband router. (I kept the one I got her—why not!)*

*The lesson here: When buying routers, proxies, and firewalls, you need to understand what they do and where you are going to put them, plus what they can plug into.*

## Introduction

When you are designing a hugely expensive and critical infrastructure, no matter how experienced you are, you will need to validate and cross-check your work. Well, you will if you are a thinking person. This chapter covers this less-than-common subject.

## Network Perimeter Security

The approach and techniques detailed here step through a series of decisions, both financial and risk based, to help you find a solid, cost-effective architecture. I have also provided a series of example designs, rules of thumb, and design paradigms that are good practice, if not best practice. These will help you get a robust design.

Here is the **first rule of thumb**: When designing secure firewalls, try to separate systems used for e-commerce from those used for browsing or other corporate communications. Indeed, as a broad paradigm, where possible you should separate the security arrangements for different types of traffic.

This statement might sound lame, but operations and network management will try to squeeze everything into one domestic home firewall hanging off a small DSL line. Then they will leave you holding the baby.

The e-commerce and corporate browsing infrastructures should be segregated because they have different:

- **Availability requirements** Loss of your browsing capability for half a morning will get you a nasty phone call from the marketing director, who spends all day “connected.” Loss of your online shop for half a morning could cost you dozens of customers and serious revenue.
- **Bandwidth requirements** Do you really want your e-channel being ground to a halt by staff watching video streams of the latest sporting event?
- **Regulation requirements** Data protection, money laundering?
- **Security requirements** Different containment and separation requirements, different confidentiality requirements that should be reflected in access controls and authorization procedures and the like.

But I’m not suggesting every corporate brochureware Web site has its own firewall and management platform. I usually use the following acid tests to help me decide:

- **Does it take transactions?** If it does, it will be a revenue center on someone’s balance sheet. These people will almost certainly demand a level of service greater than can be expected for a simple corporate Web server. They should pay for that greater service.
- **Has it got customers?** Not all customers create transactions, but they might have paid for a service. So the same argument applies.

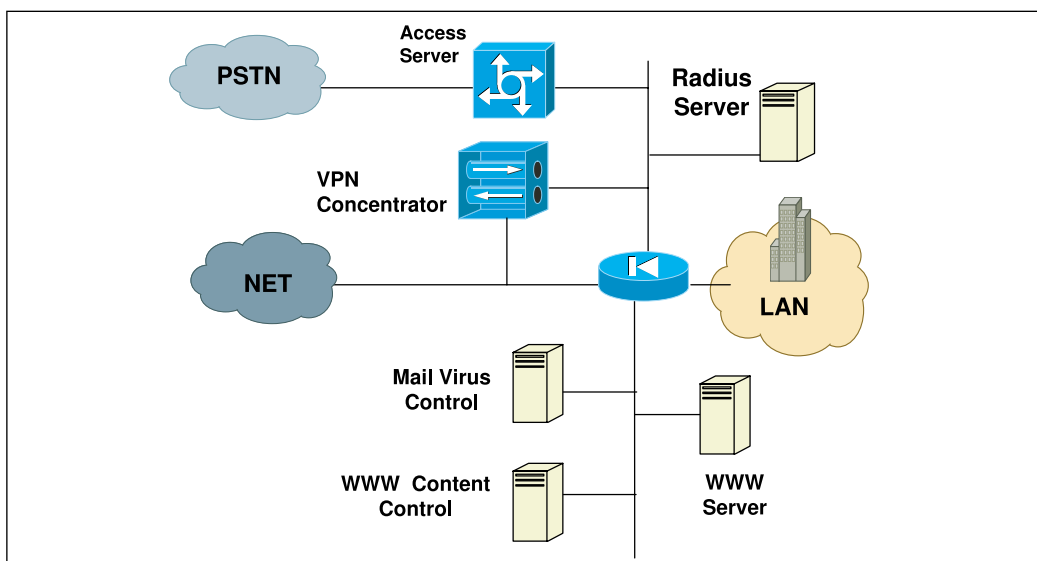
- **Has it got a nontrivial database with personal details on it?**  
This is a little avant garde, but in my opinion, systems that have databases need DBAs, so they are not trivial.

If any of these answers is yes, I tend to look for separation between the e-commerce system and the corporate firewall.

## The Corporate Firewall

The corporate firewall typically handles a number of services that the organization uses to communicate with the outside world. Here we are usually talking about internal Web browsing and the ubiquitous e-mail. However, few organizations these days are without remote access for roaming or home workers. This requires that we include IPSec-based or SSL-based VPN and/or PSTN dial-in (see Figure 7.1).

**Figure 7.1** The Corporate Firewall



If you went to a professional security boutique and asked for a design, you might receive something like this. It is aimed at a small organization so it hasn't used a defense-in-depth firewall configuration, but this layout is more than fit-for-purpose. And your designer has done it so many times before, he

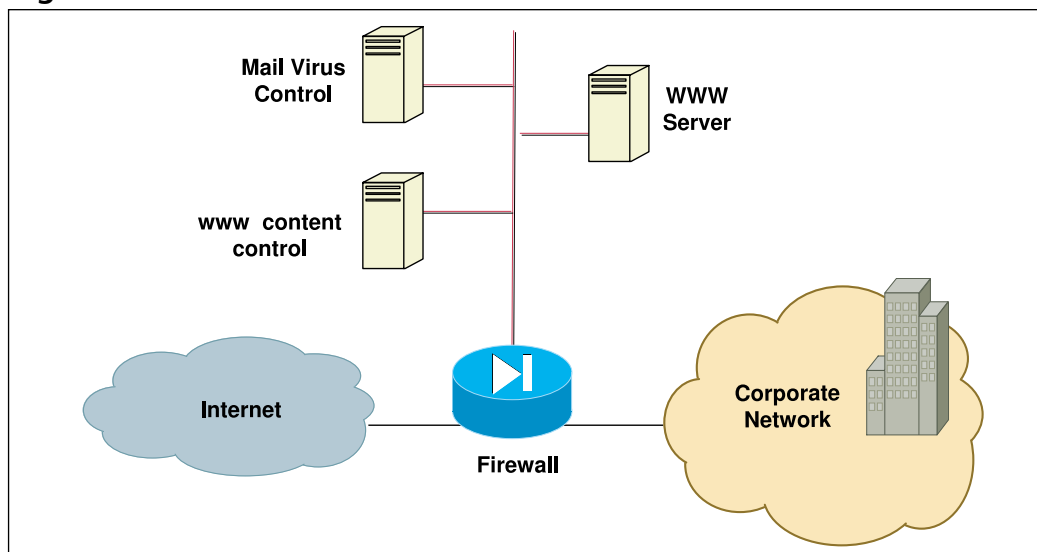
probably did it automatically, without thinking. But let's take a moment to analyze the services and the ways the inherent risks are managed.

There are really two structures here—the hardware (plus software) deployed to provide:

- Remote access into the organization
- Internet surfing and mail

There is nothing wrong in combining them if the throughput is negligible, but let's see how they look when deployed as two separate structures, the way a medium-sized organization should. (Why? Well, the functions require different services and therefore have a different risk profile. Remember the first rule of thumb.) Figure 7.2 diagrams an Internet access firewall.

**Figure 7.2** An Internet Access Firewall



## Threat Analysis

From Table 7.1 we get another design **rule of thumb**: Terminate all inward traffic in a firewall DMZ. This can be achieved by either placing the server in the DMZ or using a protocol-aware proxy.

**Table 7.1** Threat Analysis of a Corporate Firewall

| Activity                                                              | Threat                               | Countermeasure                            |
|-----------------------------------------------------------------------|--------------------------------------|-------------------------------------------|
| Public network connectivity                                           | Hacking/unauthorized access          | Firewall                                  |
| External access caused by security flaws in operating systems or apps | Hacking                              | DMZ; all inward services terminate in DMZ |
| Web Browsing by staff                                                 | Looking at porno, time wasting       | Content control and reporting             |
| Mail inward                                                           | Virus and worms<br>Spam              | Mail virus scanning<br>Spam filters       |
| Mail outward                                                          | Offensive comment or company secrets | Banned word list on<br>Virus software     |

## E-mail Protection

Protection against mail viruses is *essential*. About 80 percent of all security incidents arise from the lack of such protection. Although as a function this protection is available on leading firewalls, it historically has not performed as well there. However, this hasn't deterred suppliers, and now the successors to PIX, NetScreen, and Proventia all have it as a feature.

However, I still believe e-mail protection is much more scalable when included as a mail relay in the DMZ. This has the following advantages:

- Your mail that isn't time sensitive can be offloaded to another processor.
- You can stack multiple engines. I recommend CLAM (freeware) and at least one commercial library.
- Extra features like disclaimer and spam control are available from the separate products.

Recent tests (from sources that must remain unnamed; there are many in security) support my theory, showing that certain combined appliances missed over 20 percent of the polymorphic viruses passing through them. Other security appliances only check a hot-list of most frequently encountered

viruses, meaning that old copies of Code-Red that are still circulating can still infect you. All-in-one appliances are good in many situations, but they are often an excuse to combine three poor products into one that is marketable. See the sidebar for more details.

## Tools & Traps...

### Notes on Security Appliances

Many industry analysts and vendors recommend that you get key security features bundled on one single security appliance firewall. These products are available, and sometimes they're appropriate. Here's an analogy.

I have a Swiss Army knife and it's great. It cuts nearly as well as any other knife on the market. It also has a little can and bottle opener. These aren't very good, but in an emergency, they will do the job at the risk of hacking a lump out of my fingers. Lastly, my trusty Swiss Army knife has a little saw that is truly hopeless, but it looks good. It's thrown in for free; it doesn't cut a tree down like my chain saw, nor will it cut mitres like my cross-pull power saw. But it looks good and it's free. Where's the harm as long as you don't need it every day?

So it is with multifunction security appliances:

- Some features are good (for example, the firewall and encryption functions). These features are equally as good as sole-purpose products.
- Some functions, such as URL blocking for Web sites, do an acceptable job *and can be compared to standalone products*.
- Some functions, such as IDS or virus protection, generally *don't do the job to the required standard*. If you buy the MessageLabs service, it uses three separate virus signature libraries. I use MIMESweeper, with a minimum of two. The risk of infection by mail is high, so you really don't want to make false economies here. These functions also detract from the main function of the firewall. Often you find that a 100mb/s firewall turns into a 50mb/s one because of the extra CPU requirements of these other functions.

Continued

You are the security officer and you are responsible for the maintenance of the security of your organization's very expensive IT, so don't let these important functions be trivialized for the saving of a few hundred dollars and one shelf in a rack.

Typical products are:

- Spam Assassin
- MIMESweeper
- Mail Marshall

Alternatively, you can outsource all activity to a third party that will do the spam and virus removal for you. Leaders in the field are MessageLabs, BlackSpider, and Postini.

## Browser Content Control and Logging

Content control products prevent employees browsing unsuitable content, such as porno or hacking sites. Although there are no laws that require these control facilities, corporations increasingly buy them to limit legal exposure. A number of corporations have been sued for constructive dismissal on the grounds that they did not provide a suitable workplace because staff viewing pornography and other offensive material was making the office an "unfit workplace."

Additionally, it is clear that many employees who work the hardest and have no time to help others always seem to have a browser open on their machine. Targeting them or even viewing packets with a packet sniffer could infringe their human rights, but these automatic devices have generally been recognized by courts in sensible countries as reasonable controls.

Typical products are:

- CensorNet
- Surf Control
- WebWasher
- Web Sense
- Blue Coat

## Web and FTP Server

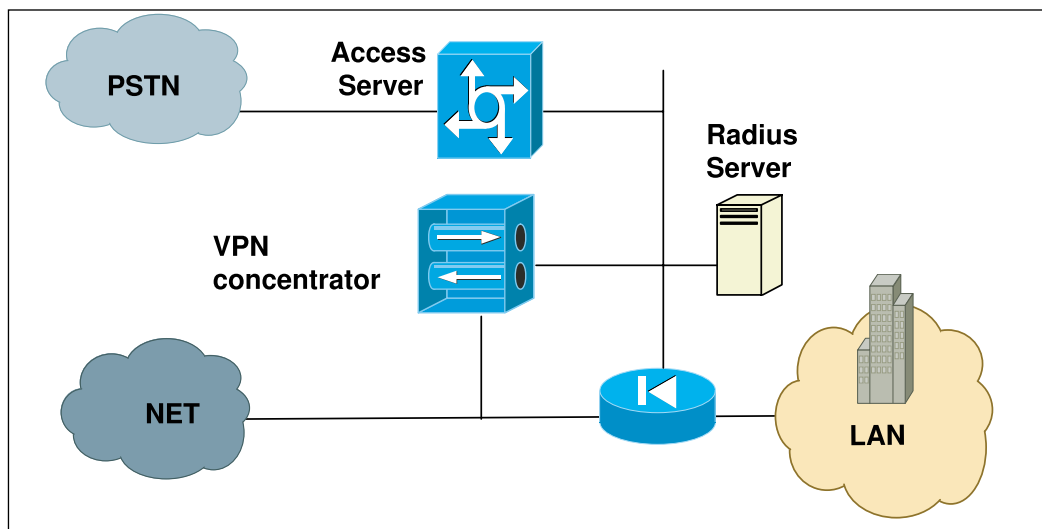
Typically, the DMZ also holds a simple corporate Web server or FTP server, even if the company has an e-commerce site elsewhere. Make sure that these servers are suitably hardened and administered separately.

## Remote Access DMZ

Many organizations have a legacy dial-in remote access server. This server will need access to a radius-based authentication server (see Figure 7.3).

More modern remote access will be enabled by either an SSL or IPSec VPN server, using the ubiquitous Internet. It is good practice to “DMZ” both of these to allow better access control. I’ll come right out and say it: Many VPN servers provide lame firewall service and aren’t certified to EAL4. Rant over.

**Figure 7.3** Remote Access DMZ



## Threat Analysis

The threats and countermeasures look something like the list in Table 7.2.

**Table 7.2** Threat Analysis of a Remote DMZ

| Activity                    | Threat                                                                                                        | Countermeasure                        |
|-----------------------------|---------------------------------------------------------------------------------------------------------------|---------------------------------------|
| Public network connectivity | Hacking/unauthorized access through insecure, nonmandated protocols                                           | Firewall                              |
|                             | Hacking/brute-force attack providing entry via authorized protocols by repeated trial of user ID and password | Strong authentication                 |
|                             | Getting access to confidential information                                                                    | Encryption                            |
|                             | Man-in-the-middle attack                                                                                      | Encryption plus strong authentication |
| Mail inward                 | Viruses and worms                                                                                             | Optional mail virus scanning          |

## Remote Access Design Options

Personally, these days I usually run encryption over the PSTN. The local operators in some parts of the world aren't very secure and run it through IP over the Internet, so it makes sense. To achieve this goal, a direct connection into the VPN concentrator from the access server is required. Typically, a port on the same virtual LAN (VLAN) or a new VLAN and an extra network interface card (NIC) in the concentrator will do it.

For strong authentication, I am particularly fond of the one-time password generators, so you should consider deploying:

- Cryptocard
- RSA SecurID
- Vasco token—Digipass

These can be used with the PSTN, SSL VPN, or IPsec VPN. They should speak “radius” as a native tongue so they require no further integration.

As a one-time evangelist of PKI, I should recommend digital certificates, but I'm not, simply because they are too difficult to manage for a typical

medium-sized organization. (Please note that I'm not suggesting that risk/threat is a function of size only; capital expenditure and head count to maintain large infrastructure *often* are. Many investment banks are “medium-sized” but need and use Rolls Royce countermeasures.)

If you are dying to use an integrated firewall appliance, this is your chance. As mentioned before, many of them have virus capability, and *extra* virus scanning of your remote access connection is a useful addition. Many of them can also speak SSL-VPN and IPSec-VPN, so you can end up with all the functionality in one box. That has to be good for the flexible enterprise. But please make sure that virus scanning occurs on the decrypted VPN stream. Use the dummy virus EICAR to check.

Lastly, and as mentioned earlier in this chapter, it would be good to have IDS included. As an old fossil, I would not jump at the chance to integrate IDS into an appliance. An IDS is a detective control, and therefore I like it to have complete separation (in other words, segregation of duties).

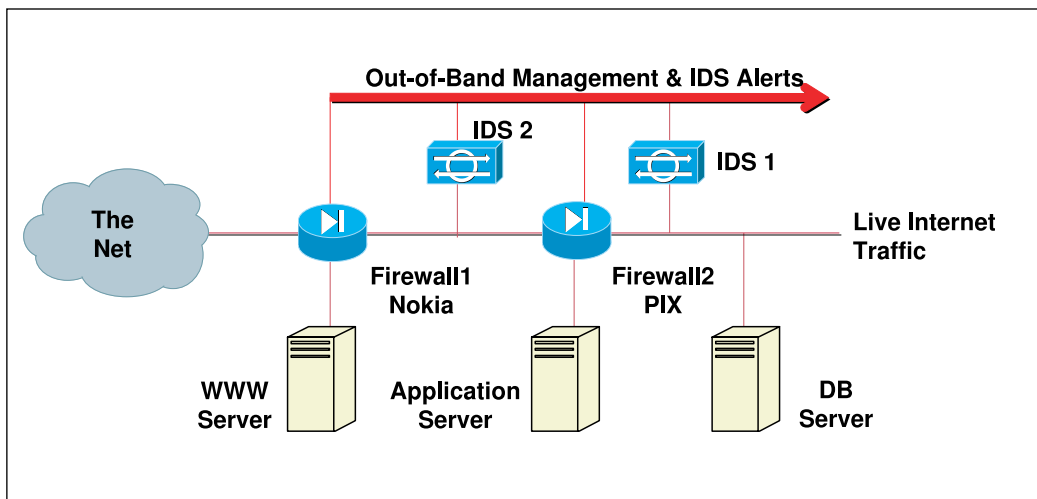
## E-commerce

Designing an e-commerce configuration is not very much more difficult, but it does require an understanding of the application that is going to be hosted and the products that are to be used. I once worked at a hosting organization that applied one standard £50,000 infrastructure onto the typical £2 million application, with no regard to the protocols and the risks. Surprise, surprise—the net result was not just poor firewalling. It extended to a virtual war between support and sales, with perplexed and unhappy customers to boot.

The fact is that application of defense in depth is a lot more complex than it might seem. Issues that are most likely to be encountered are:

- Cost—as always
- The need for resilience
- Better detective controls

A good nonresilient configuration might look like the diagram in Figure 7.4.

**Figure 7.4** A Network Configured for E-commerce

This is a classic nonresilient design. Let's observe the key design characteristics to help us form some rules of thumb for good practice:

- **Dual “defense in-depth” firewall configuration** Two firewalls from different manufacturers. In this design, we place a PIX with its robust stack and limited OS as the external firewall. We would place a more flexible, software-based firewall, such as Check Point FireWall-1 running on a Nokia platform, as the internal firewall. This firewall is feature rich and can provide some excellent enhancements.
- **Use of tiered design** Each application layer has its own security zone, or tier. This extends the concept of defense in depth into the realm of the principle of least privilege. Each class of server can only communicate to servers in the same tier (and therefore with the same risk profile) or to specific servers outside that tier on a specific port allowed by a firewall rule. This means that if an exposure develops, it is contained in the appropriate tier.
- **Use of detective and preventive controls** IDs and firewalls.

- **Out of band management** Access to the firewall, servers, and IDS is via a special separate network. This provides a degree of separation discussed in the previous section.

### Notes from the Underground...

#### Out-of-Band Networks

Out-of-band (OOB) networks have to be made especially secure. Many times during my years as a penetration tester, I found that an OOB had not been prepared properly, so it represented a high-speed back door for hackers into the heart of the corporate network. OOB networks should always be terminated on a firewall. This really is the application of our first rule of thumb that we mentioned earlier.

In the earlier example, we have achieved a degree of separation in the most common and economic manner. We run the OOB network to the IDS and firewalls but from then on use the existing infrastructure to reach the application servers. This means that management traffic could be entering into the same server interface as data, and this might be unacceptable.

If more separation is required, each server must have a separate NIC installed. Define each DMZ as a separate Security Zone and give it a separate physical switch to prevent VLAN hopping. In our example:

- Zone 1 (outside) = Firewall1
- Zone 2 (middle) = WWW server, IDS2, and Firewall2
- Zone 3 (inner) = Application DB servers and IDS1

If you are completely paranoid, you can do what I have seen recently and give each device a separate VLAN—doable, but nuts.

## Threat Analysis

The threats and countermeasures look something like the list in Table 7.3.

**Table 7.3** Threat Analysis of a Network Configured for E-commerce

| Activity                       | Threat                                                                                                        | Countermeasure                                                                   |
|--------------------------------|---------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Public network<br>Connectivity | Hacking/unauthorized access causing reputation loss because the Web site is defaced                           | External firewall                                                                |
|                                | Hacking/unauthorized access causing regulatory or financial loss because personal data is accessed or damaged | Internal firewall and IDS/IPS                                                    |
|                                | All admin/privileged access allowed only via out-of-band network                                              | Out-of-band admin network, strong authentication                                 |
|                                | DDOS attack                                                                                                   | No countermeasure                                                                |
| Identity theft                 | Unavailability due to configuration error                                                                     | No countermeasure                                                                |
|                                | Getting access to confidential information                                                                    | Encryption (SSL)                                                                 |
| Unauthorized access from LAN   | Man-in-the-middle attack                                                                                      | Encryption plus authentication                                                   |
|                                | Insider attack                                                                                                | Internal firewall, IDS/IPS, out-of-band admin network with strong authentication |

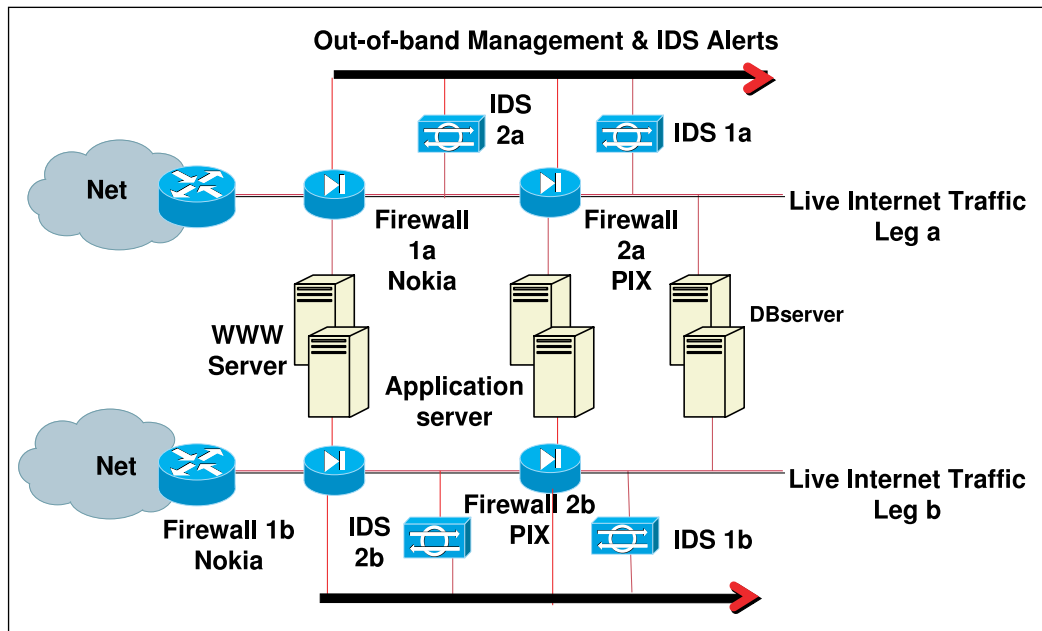
The previous configuration deals well with confidentiality and integrity but not availability. Everything is a single point of failure. A small error in a firewall rule or a disk failure will see the whole e-business offline for a number of hours while lengthy restores are conducted or replacement firewalls found.

In the configuration shown in Figure 7.5, the single points of failure have been removed by the duplication of all key components. In effect, it is two of the simple configuration “bolted” together to form two legs. This gives rise to another rule of thumb.

**NOTE**

Where possible, eradicate single points of failure.

**Figure 7.5** Duplicating Key Components in a Network Configured for E-commerce



Obviously, this should be done based on needs, mean time to failure (MTTF), and service level agreements (SLAs). Redundant configurations can be created very simply, usually with the minimum of effort in active/passive mode. This is where you buy two pieces of kit (firewall, router and switch) but only use the second of each (nominated the secondary) when the first (nominated as the primary) fails.

Management, however, tends to find this solution unacceptable because they pay twice as much but 50 percent of the purchase sits in a corner gathering dust (figuratively speaking, unless it is a very cold spare).

Management, therefore, tends to prefer active/active configurations, because they gain a perceived performance benefit from the having two

## 138 Chapter 7 • Infrastructure Security

pieces of kit in use at once. The presumption is that *two is better than one*. This tends to be a myth, because typically, the performance bottleneck isn't the routers, firewalls, and switches but the application—but *c'est la guerre*.

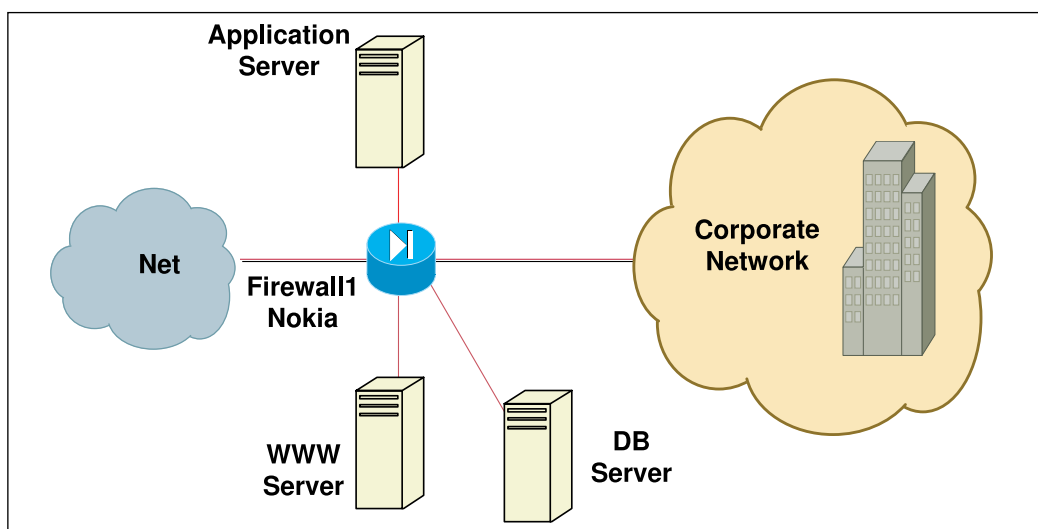
When you're trying to create active/passive structures, beware of:

- **Asymmetric routing** Your firewall and your IDS will probably not like an inward path being different from your outward path.
- **Nonsticky sessions** Your application will almost certainly need to be served continuously from one server; it needs to be sticky. Otherwise cookies and session data go walkies.
- **Memory creep** Data retrieval time increases through loss of locality of reference.
- **Span tree or route convergence problems**
- **Problems with client-side certificates**

And despite it all, the application will probably run slower.

However, if your organization is a little strapped for cash, the single-firewall configuration shown in Figure 7.6 has a degree of the protection that a bigger configuration gives.

**Figure 7.6** A Single-Firewall Configuration



Here we have a single firewall. However, each tier is given a separate interface to form a DMZ—a reasonably low-cost option that provides a high degree of containment between the presentation tier, the application tier, and the database tier. This means that each can have a separate access list that prevents direct access to standing data from the outside.

## Threat Analysis

The threats and countermeasures look something like the list in Table 7.4.

**Table 7.4** Threat Analysis for a Single Firewall Configuration

| Activity                     | Threat                                                                                                        | Countermeasure                                                                                                                                |
|------------------------------|---------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Public network connectivity  | Hacking/unauthorized access causing reputation loss because the Web site is defaced                           | The only firewall access list on interface 1                                                                                                  |
|                              | Hacking/unauthorized access causing regulatory or financial loss because personal data is accessed or damaged | The only firewall access list on interface 2; outside addresses are never allowed to touch the databases; access must be from the application |
|                              | All admin/privileged access allowed only via out-of-band network                                              | No countermeasure                                                                                                                             |
|                              | DDoS attack                                                                                                   | No countermeasure                                                                                                                             |
|                              | Unavailability due to configuration error                                                                     | No countermeasure                                                                                                                             |
| Identity theft               | Getting access to confidential information                                                                    | Encryption (SSL)                                                                                                                              |
|                              | Man-in-the-middle attack                                                                                      | Encryption plus authentication                                                                                                                |
| Unauthorized access from LAN | Insider attack                                                                                                | No countermeasure                                                                                                                             |

## Just Checking

When you think you've finished your design, just ask yourself if you've remembered:

- **NTP** You will need a time server or at least a time source; otherwise your logs will be useless.
- **Syslog server** Oh yeah, *logs*. You'll need a syslog server for the UNIX box's switches and routers.
- **Console** If the data center is remote, you need to get to the console, so you'll need a console server.
- **Authentication** Are you going to centralize user accounts or use strong authentication? This can be very problematic with Windows and Active Directory.
- **Backup** Have you got backup software and a tape drive installed (on your management LAN) to back up the app?

## Summary

The purpose of this chapter was not to show you how to design complex DMZs and firewall solutions for your organization. That's not your job (although if you have absorbed the information as I have hoped, you would do a better job than most). The purpose of this chapter was to show you the rationale behind common designs and the considerations behind the designs. This knowledge will enable you to ask searching questions about the design and the protection it affords to your company—and that *is* your job. How can you do it without this knowledge?

In the chapter, we covered my basic rules for design. These are only guidelines but are listed here again. When you begin asking the questions about your particular setup, these points provide a very useful start. Use:

- Separate firewalls for e-commerce and general corporate use
- Proxies or DMZ-servers to terminate all inward traffic to the corporation

- Dual defense-in-depth firewall configuration, where required
- A tiered firewall design
- Detective and preventative controls
- Out-of-band management
- Duplicate components to eradicate single points of failure

Now on to firewalls.

