Chapter 1

Defining Perimeter and Internal Segments

Solutions in this Chapter:

- Internal versus External Segments
- Footprinting: Finding the IP Addresses Assigned to Your Company

Related Chapters:

- Chapter 2 Assessing Your Current Network
- Chapter 10 Perimeter Network Design
- Chapter 11 Internal Network Design

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

With the proliferation of wireless access points (WAPs), virtual private networks (VPNs), and extranets, it's becoming increasingly difficult to determine where your network begins and ends. Add this complexity to common economic factors, such as company mergers and acquisitions, and now you have a tangled web of interconnected segments and networks that you will need to understand. While this book aims at providing you the necessary tools to protect your network infrastructure assets, it is imperative that before we dive into the details you have a good understanding of how your network is designed.

Having a commanding knowledge of your network topology today is no simple feat. We are often reminded of a financial services company at which we performed some consulting work. This company has grown over the past few years by acquiring related financial companies. At the end of the day, this team of network engineers had to manage over 300 Frame Relay lines, over 100 Microsoft Windows NT 4.0 domains, and numerous Internet access points (IAPs). To add insult to injury, these networks are not static environments; in fact, there are numerous routing changes and firewall modifications made on a daily basis. The only saving grace this team of dedicated foot soldiers has are solid topology diagrams detailing each Frame Relay network and IAP, and a comprehensive list of all of their outwardly facing IP addresses.

While these tools sound like networking basics, we are constantly surprised at the number of IT departments that are without this information. Without knowing how your network is laid out, or understanding which segments touch the Internet directly, it will be nearly impossible for you to begin locking down your network devices. If you are not armed with these tools already, this chapter will help you find your external IP address presence and help you get a handle on understanding the differences between your core network segments and those that lie on your perimeter. Chapter 2, "Assessing Your Current Network," will help provide you with those all-important topology maps if you aren't fortunate enough to have them in your toolbox already. Furthermore, the end goal of this chapter is to arrive at common language that can be easily understood, and used throughout the entirety of the book.

Internal versus External Segments

Most of the time it might be quite simple to define your network segments as internal or external, core or perimeter; in larger, more heterogeneous

organizations, this is not an easy task. Corporate acquisitions, multiple Internet service providers (ISPs), and remote offices offer areas of complexity that might result in some uncertainty as to which network is connected and where it leads. The following section will help you define and piece together those segments that will lead to a better understanding of your network topology.

Explaining the External Segment or Perimeter Segment

Simply defined, an external, or perimeter segment, is any network that exists in a low security zone of your environment. In other words, any network that connects your physical environment to another untrusted network, such as the Internet. A good example could be a network that is attached to the external interface of your firewall and connects to the external interface of you ISP's router. In this scenario, the network is untrusted from the standpoint of your organization because it is ultimately controlled by the ISP.

This definition could extend to other network segments as well, such as a demilitarized zone (DMZ) that houses and provides Web or application services to other untrusted networks. In many cases, this type of network would be considered external, or on the perimeter, since many of those services map directly to external or public IP addresses. This class of service would still fit in our description because the firewall is passing certain types of untrusted traffic to that DMZ network; thus, you cannot always guarantee the safety of those devices from Internet traffic.

If you begin to think about your network from the perspective of a potential attacker on the Internet, the definition of the external segment will become clearer. An untrusted Internet attacker will only have access to devices or services that are directly connected to the Internet. With this in mind, you now have a clear picture of what we would consider a perimeter network or device. Does it serve content to the Internet? Can anyone PING or connect to the device?

Wireless Access Points: Extending the Perimeter

As wireless technology has matured over the years, so has its acceptance in corporate America. More and more, companies are turning to wireless technology to extend usability to employees and management. While this increase in usability can drive efficiency in the workplace, it also adds risk to the IT department that is working to protect the corporate assets.

Without diving into too much detail on how WAPs work, each device emits a radio frequency (RF) that is used to pass network communication and protocols. Many of these devices have a substantial range, meaning that people who are physically located far from the access point will still be able to communicate with it. Additionally, in many companies these WAPs are located on internal segments, providing connectivity to corporate mail servers, payrolls servers, intranet sites, and potentially users' desktops.

The inherent risk from these devices comes from that fact that they might not be properly secured. Unsecured WAPs provide a gateway into the internal network for untrusted users. Potential attackers could take advantage of misconfigurations or lax security policies on these devices and begin to communicate on your internal network. Because of the increased range capabilities of these devices, the untrusted user might be walking by your building, sitting in your parking lot, or on a different floor in your office building. Regardless of the user's location, this unsecured device just opened the door to your internal network.

So, how do WAPs extend the perimeter? If you recall our basic definition of an external segment (providing services or connectivity to an untrusted network or user), this technology falls into that scenario. This device could potentially allow an untrusted user with no privilege access to your company's internal assets and resources, thereby extending the perimeter onto your internal segments. What's worse is that any type of elaborate firewall setup (that might be air-tight) has been completely circumvented and done so from the comfort of the untrusted user's '83 Toyota across the street.

The Internal Segment Explained

Using the information already presented in this chapter, it is quite simple to deduce what the definition of an internal segment is. For the purposes of this book, we define an internal segment as any network that resides in the secured portion of your environment and provides resources or services that are only for internal use (that is, should not be accessible by untrusted Internet users).

Similar to how we thought about our external properties, if you think about the internal segments as providing resources only to internal assets, you will get a clearer picture of how the network should be defined. Most of the networks within your corporate environment will be internal, as many companies have only a few IAPs.

Assigning Criticality to Internal Segments

Since most of your networks are going to be internal segments, they cannot all have the same importance for your organization. Prioritizing these segments is an important step in aligning your network for security and business continuity plans. For example, many of your network segments will only house employee desktops or laptops, while some might contain mission-critical servers, such as mail, payroll, software development source code, customer databases, or HR applications. While you will want to provide the most comprehensive security policy and defense for your entire environment, it is not practical when the latest security tsunami hits.

Assigning network and device criticality is an essential step in planning for how you are going to handle security patches, network recovery, and continuity. For example, a few months ago a serious design flaw was discovered in the Cisco Internet Operating System (IOS) that runs on all Cisco routers and some other Cisco network devices. Many organizations have hundreds, if not thousands, of Cisco routers in use on their network. Instantly, those companies had a massive project on their hands. The use of network and device criticality helped those administrators put together a plan of action on which Cisco devices needed to be updated first and which were less important.

For the perfect example, we refer back to our favorite financial services company that we previously mentioned. When the Cisco IPv4 vulnerability hit the wire in July 2003, this company was not prepared for the chaos and damage that could potentially ensue from such a threat. With nearly 700 Cisco devices deployed across their worldwide enterprise, this bank only had a few spreadsheets with asset information, mainly comprised of IP addresses and physical asset location. What's worse, the security team had zero information as to which department or person was in charge of the maintenance of each device. Any inkling of network device criticality at this point was nothing but a distant dream.

Within a few hours, reports started to surface as to the dire circumstances surrounding this vulnerability. The security team was feverishly trying to make heads or tails of the asset inventory information they did possess. Questions similar to, "Is that our router or does the Telco maintain it?" were shouted from offices. Spreadsheets were being circulated through e-mail like a bad Outlook virus! Alas, IT personnel had very few answers and a tremendous amount of questions. Almost four hours into the crisis, they had made zero progress on their remediation efforts.

All told, it took nearly six business days for the bank to fully remediate their Cisco devices. The main reason for this delay was not policy or change control, but rather, the network engineers did not have accurate inventories of the network device assets and their respective owners/maintainers. Essentially, it took them six days just to find all of their routers and the corresponding individual who administered the device. It was not an impressive showing, but thankfully the vulnerability turned out to be nothing more of a scare, so little damage was actually realized.

Nevertheless, had they moved on from this incident without learning anything this anecdote would not have made the pages of this book. The security staff spent many weeks after the Cisco scare working on assembling all of the asset information into a consolidated spreadsheet. They documented their network architectures and spent time going through all of their telecommunications contracts to understand where their responsibilities ended and the ISP demarcation began. Their data collection did not stop with networking devices, but stretched to the desktop where they inventoried systems down to the OS revision. With this information in hand, they began to decide which devices and networks were most important to the business. While this information didn't prove useful immediately, it wasn't long until the next Microsoft worm exploded onto the scene.

When the Microsoft Messenger Service Buffer Overflow began to make headlines in October 2003, this security team was well poised to respond. Even with thousands more Windows devices to patch (compared to only 700 Cisco devices), the total time for complete remediation was only three days–a significant improvement in their processes. Part of the reason why they were able to act so swiftly this time was the asset inventory spreadsheets and the asset criticality information. Rather than spinning their wheels on less critical Microsoft systems, they focused on the business-critical servers and workstations first, and then broadened their approach outward as resources became available. This allowed them to ensure the continuity of the business through the security threat, and lessen the potential impact across the enterprise.

As you begin to map out your network, it would be wise to begin thinking about how important that segment is to your business. Documenting this information will help when crisis strikes and you and your team need to act swiftly.

Footprinting: Finding the IP Addresses Assigned to Your Company

Now that you have a clear understanding of where your perimeter networks are, and more importantly what they are connected to, the next important step is to ensure that you haven't missed any of them. Since your perimeter networks should be the only gateway for untrusted Internet attackers to enter your network, you will want to make certain that there aren't any other IAPs out there that were acquired through a business merger or a new remote office. The following sections will help you begin to collect information about the public IP addresses assigned to your organization.

Using *whois* to Understand Who You Are

The International Corporation for Assigned Names and Numbers, better known as ICANN, defines the Address Supporting Organization (ASO), which maintains databases of assigned public IP addresses. These databases are broken down into Regional Internet Registries (RIR). Each geographic region has an organization that is responsible for tying the publicly assigned IP addresses with the corresponding company. In other words, when you or your ISP purchases a new network block, the company and contact information is stored in these databases. These providers correlate the IP address block information with your public company information. The following is some sample output of a RIR IP block record:

```
OrgName:     BrianCorp Inc.
OrgID:       BrianCI
Address:     One Brian Way
City:        Newport Beach
StateProv:   CA
PostalCode: 92660
Country:     US

NetRange:    192.0.2.0 - 192.0.2.255
CIDR:        192.0.2.0/24
NetName:     BCorp
NetHandle:   NET-192-0-2-0-1
Parent:      NET-192-0-0-0-0
```

```
NetType:     Direct Assignment
NameServer: NS1.US.bkhome.COM
NameServer: NS2.US.bkhome.COM
Comment:
RegDate:     2002-09-26
Updated:     2004-03-01

TechHandle: BK763
TechName:   Kenyon, Brian
TechPhone:
TechEmail:  dns@bkhome.com

# ARIN WHOIS database, last updated 2004-03-03 19:15
# Enter ? for additional hints on searching ARIN's WHOIS database.
```

There are currently four active RIRs and one pending approval. The RIRs are as follows:

- **ARIN** North and South America Registry also serving parts of Sub-Sahara Africa
- **APNIC** Registry serving the Asia Pacific region
- **LACNIC** Latin America and parts of the Caribbean
- **RIPE** Registry for Europe, Middle East, Central Asia, and parts of Africa
- **AfriNIC** Pending approval, will serve African regions

Unless your organization is located in several different countries, you will most likely be using ARIN for the majority of *whois* queries.

RIRs can be queried by using IP address or domain name to provide specific company information. Only UNIX-based operating systems come with an embedded *whois* client; however, there are several freeware utilities available for the Windows platform. For the most part, you could use various Web sites to handle the *whois* query for you, such as www.network-tools.com or www.dnsstuff.com. The Network-Tools site will allow you to search through the ARIN, RIPE, and APNIC databases only, while the DnsStuff site will attempt to ascertain the appropriate RIR to query before giving you an error. For further searching capabilities you can go directly to the particular RIR's Web site, such as www.arin.net or www.apnic.net.

Using DNS Interrogation for More Information

What happens if you do not know all of the domains or IP addresses that might be assigned by your company? If your organization, or parent company, is a publicly traded company, you can use the U.S. Securities and Exchange Commission's (SEC) Web site to gather information about potential subsidiaries. The SEC has a search utility named EDGAR used for searching through public SEC filings. Using this utility, you can query your company name for a detailed list of all the SEC filings. For simplicity, we typically look at the 10-Q filings for any given organization. These filings take place each quarter and will have the most up-to-date information. Once you open the filing, search for the term *subsidiary*, or any variation of it, to find other related entities to your organization.

For example, a search on a fictional company, BrianCorp Inc, might yield the subsidiary, Brian-Ventures. With this information, we are going to do a little more digging.

Using the utility NSLOOKUP, which is on all versions of Windows and UNIX operating systems, do a quick lookup for Brian-Ventures.com, Brian-Ventures.org, and so forth.

```
C:\>nslookup brian-ventures.com
Server:  dns.bkhome.com
Address:  192.0.2.111

Non-authoritative answer:
Name:    brian-ventures.com
Address:  192.0.2.21
```

Our results show that the domain brian-ventures.com does exist and it resides at the IP address 192.0.2.21 (not a public IP address and used for example only). Using this information we go to the ARIN Web site and do a quick lookup on the IP address to see what the entire network block is and to determine if it actually belongs to the company. The following is some sample output:

```
Search results for: 192.0.2.21

OrgName:     BrianCorp Inc.
OrgID:       BrianCI
Address:     One Brian Way
City:        Newport Beach
```

```
StateProv:  CA
PostalCode: 92660
Country:    US

NetRange:   192.0.2.0 - 192.0.2.255
CIDR:       192.0.2.0/24
NetName:    BCorp
NetHandle:  NET-192-0-2-0-1
Parent:     NET-192-0-0-0-0
NetType:    Direct Assignment
NameServer: NS1.US.bkhome.COM
NameServer: NS2.US.bkhome.COM
Comment:
RegDate:    2002-09-26
Updated:    2004-03-01

TechHandle: BK763
TechName:   Kenyon, Brian
TechPhone:
TechEmail:  dns@bkhome.com

# ARIN WHOIS database, last updated 2004-03-03 19:15
```

From this information provided by the ARIN database, we are able to ascertain that the Web site is owned by BrianCorp, and we own the entire 192.0.2.0 Class B network. Keep in mind, however, that BrianCorp might not own the entire Class B range, as they might just lease a small subset of the Class B from their upstream ISP or Web hosting provider. However, with this information we can cross-reference our network topologies and make sure that we accounted for this public (external facing) network.

Tools & Traps…

The DNS Zone Transfer

DNS has always provided a volume of information regarding which domains belong to a company and on which network it resides. While this information is generally used so that the general public can access your public Web sites by mapping an IP address to the domain name, it can also provide a lot of useful information in tracking down which domains are owned by the company.

If you do not have access to your DNS zone information, you can try to obtain it through a common DNS feature called the *zone transfer*. Zone transfers were previously used to share updated information with other DNS servers, primarily for redundancy in case the primary DNS server were to fail. While open Internet zone transfers aren't common practice anymore, some DNS servers and networks are still misconfigured to allow this. The most common attribute of a DNS server that allows zone transfers is the presence of TCP port 53 being open. Since common DNS queries are performed on UDP port 53, TCP does not need to be open and can be blocked, thereby disabling zone transfers on the network layer.

Using a utility like NSLOOKUP will provide the mechanism for the zone transfer.

```
C:\>nslookup
Default Server:  dns.corp.com
Address:  10.22.164.12


> server dns.bkhome.com
Default Server:  dns.bkhome.com
Address:  192.0.2.111
> set type=any
> ls -d bkhome.com
[dns.bkhome.com]
bkhome.com           SOA      dns.bkhome.com
bkhome.com           MX       30    mail.bkhome.com
bkhome.com           NS       dns.bkhome.com
bkhome.com           A        192.0.2.2
```

Continued

```
mail                    A        192.0.2.3
www                     A        192.0.2.2
brian-ventures          A        10.162.183.21
brian-invest            A        10.162.183.22
```

The preceding output shows all the subdomains and the mail record for the bkhome.com domain. From this information, we can see that there are two different networks that are providing Web services: 192.0.2.0 and 10.162.183.0.

While we used this for internal IP address allocation reasons, attackers can use this information to learn about your networks and your topology. As a general rule, you want to disable zone transfers from both the Internet and internal segments.

DNS zone transfers can be disabled both from a networking and an application perspective. To block zone transfers on the network you can filter TCP port 53 to the DNS server. While this will block the zone transfer from occurring over the network, the DNS application would still allow it if you could connect to the server on that port. Each DNS application, such as the Windows DNS implementation and BIND (Berkley Internet Name Domain) for UNIX, have different remediation steps to disable zone transfers. The zone transfer can be disabled entirely, or it can be enabled to only allow transfers to particular hosts, which is a more common implementation method.

Checklist

- Take the time to make an accurate diagram of your network infrastructure, including IAPs and leased lines from your telecommunications provider.
- Use vulnerability assessment (VA) tools, or port scanners to discover and record devices on your network.
- Using VA tools look for WAPs and examine their security policies.
- Check Regional Internet Registries (RIRs) for detailed information on your company's network blocks and assigned IP addresses.
- Query and examine your DNS servers regularly to determine if there is any unneeded information leakage or the possibility of zone transfers.

Summary

This chapter helped provide some of the basic information that can later be used in diagramming and understanding the network topology. While much of this information is not ground breaking, we have established a common language that will be used throughout the book. The use of *external* or *perimeter segments* will be used to refer to untrusted networks, or those that can be easily accessed from the Internet, while the *internal segment* will be used to describe the protected internal company networks.

We also provided some valuable information on tracking down domains and rogue networks that your IT department might not be aware of. The Regional Internet Registries will provide detailed information on the network blocks owned by your company. This information is extremely valuable, as it will help you understand what is publicly available and to where your perimeter extends. Additionally, we touched on the notion of assigning a criticality value to each of your internal and external network segments. This data will help you decide how to react when a serious security vulnerability emerges and you are forced to react to protect you company's networks.

Ultimately, having these data points will help you apply the techniques and procedures in this book. Having a solid knowledge of where all your devices are and how they interconnect will be essential in providing a solid defense-in-depth strategy to protecting your environment.

Solutions Fast Track

Internal versus External Segments

- ☑ External or perimeter segments are networks that are directly connected to an untrusted network, such as the Internet.
- ☑ Internal segments are networks that are highly protected and secured and provide interior resources that should not be available to untrusted networks.
- ☑ Wireless access points (WAPs) extend the perimeter into the internal segments, as they can allow untrusted and unprivileged users access to internal resources.

- ☑ Network and asset criticality is an important data point allowing you to prioritize your work in remediating security vulnerabilities across the enterprise.

Footprinting: Finding the IP Addresses Assigned to Your Company

- ☑ Regional Internet Registries (RIRs) provide detailed information regarding IP blocks assigned to your company.
- ☑ These RIRs can be queried using a *whois* client or through various Web sites.
- ☑ DNS information can be a valuable source in finding rogue domains and networks in use by your company.

Links to Sites

- **www.arin.net** The RIR site for North America.
- **www.apnic.net** The RIR site for the Asia Pacific region.
- **www.ripe.net** The RIR site for Europe, the Middle East, and Africa regions.
- **www.lacnic.net** The RIR site for Latin America.
- **www.network-tools.com** A basic network management site featuring multiple network lookup features.
- **www.dnsstuff.com** A site with various CGI network management and DNS-related tools.
- **www.sec.gov** The U.S. Government Securities and Exchange Commission used for publicly traded companies and their filings.
- **www.freeedgar.com** A site dedicated to searching the SEC filings.

Mailing Lists

- **www.apnic.net/community/lists/index.html** Provides general discussions on the APNIC registry.

- **www.arin.net/mailing_lists/index.html** Provides numerous mailing lists regarding the North America Internet registry.
- **www.cisco.com/offer/newsletter/123668_1** This Cisco mailing list provides quick information on Cisco products and vulnerabilities.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

Q: How should I begin to discover and map devices on my network?

A: Port scanners and vulnerability assessment tools offer a great way to discover live devices on your network. Most tools allow you to export the results into a CSV or XML for further manipulation. Refer to Chapter 2 for more details.

Q: I have multiple Frame Relay lines in my network, but very little information on them; what should I do?

A: As boring as this sounds, digging up and reading your telecommunication contracts can be extremely beneficial in uncovering details about your leased lines.

Q: I do not have a DMZ and do not provide any services out to the Internet, so do I have a perimeter?

A: Yes, you do. Even if you have a drop-all policy on your firewall, and no DMZ connected to it, you still have devices that are connected to the Internet and could potentially be compromised. For example, at the very least your firewall has an untrusted interface connected to the Internet. This interface can fall victim to some firewall exploits and provide a door into your internal network. If you have a router connected to your Internet lines, that would be a perimeter device and poses some risk to your infrastructure.

Q: All of my network segments are critical; how can I differentiate them and assign different values?

A: This is actually simpler than you might think. Take some time and set up a meeting with your CFO or Risk Management person and ask him or her what the most critical aspects of the business are, and what could potentially cause your business to come to a crashing halt if it were to stop working or become unavailable. Then, examine your networks with these factors in mind. When you isolate those segments or devices that provide value to those vital business factors, than you have decided on which are most critical networks and devices. Everything else cascades down from there.