Global Information Assurance Certification Paper

Todd Ginther
June 12, 2000
GIAC Firewall and Perimeter Protection Curriculum
SANS 2000 – San Jose

NOTE: The IP addresses of my company have been sanitized throughout this assignment to 10.x.y.z addresses. For the purposes of this assignment, assume 10.x.y.z addresses are Internet routable addresses.

1 – Egress filter

A simple way for companies to be good net citizens exists and is called egress filtering. Simply put, it is protecting the rest of the Internet from spoofed traffic originating from within your company network. This is usually done by egress filters on the border routers of your company.

The benefits of egress filtering can be extended further when the filters are set up correctly. For example, the following Cisco access control list, when applied to traffic leaving the company network destined for the Internet, simply protects against spoofed traffic by allowing only the firewall (10.20.30.40) to appear as the source address in packets:

```
access-list 110 permit ip host 10.20.30.40 any
access-list 110 deny   ip any any
```

A very useful extension of this access control list is shown below:

```
access-list 110 permit ip host 10.20.30.40 any
access-list 110 deny   ip any any log
```

The 'log' statement can serve to alert you to packets the somehow got though the firewall unchecked but where blocked at the router. It can also alert you to trojans installed on inside systems that attempt to send out spoofed traffic.

The first line of the filter permits all ip traffic out to the internet only if the source address is that of the firewall, 10.20.30.40.

The second line blocks and logs all other traffic.

The steps to applying this filter on a Cisco router are (from enable mode):

```
#config t
#access-list 110 permit ip host 10.20.30.40 any
#access-list 110 deny   ip any any log
#int eth0
#ip access-group 110 out
#exit
#exit
```

Please note that this assumes eth0 is the interface connecting your router to the Internet. When this access list has been verified it should be saved:

```
#wr mem
```

This particular access list could be tested by allowing one inside machine access to the Internet through your firewall and then view the router log files to ensure that this attempt was blocked and logged.

2 – Firewall policy violations

Router log entries as well as firewall log entries will be used for this portion of the assignment because our router is part of our security solution and, because traffic into our company is so locked down, our firewall rarely sees strange outside traffic.

Violation 1:

```
*Jun  7 11:02:24: %SEC-6-IPACCESSLOGP: list 120 denied tcp
24.27.232.200(53) -> 10.20.38.19(53), 1 packet
```

Here we see an external user most likely attempting a DNS zone transfer in an effort to get a look at our company network layout. With a company's DNS tables, a potential attacker may glean information on which systems are financial system, domain controllers, etc, due to bad (or, obvious) machine naming practices.

'list 120' is the router access control list that blocked and recorded this traffic. In this case, it is the list applied against Internet traffic entering our company network.

The next two fields, 'denied' and 'tcp' show the action taken against the traffic and what type of traffic it was.

The first address we see is the source IP address in the packet with the port number in brackets (53 here, DNS). The second IP address is the destination IP address and again, the port number is in brackets (53 here, DNS).

This violation was caught by the following rule on our border router:

```
deny ip any any log
```

This is a very simple rule and it sits as the last entry in the access list that controls the Internet traffic destined for our company network. If the default implicit 'deny' was used instead of this explicit one, logging would not have taken place and the volume and types of attacks against us would never be known unless our router failed and the traffic all hit the firewall.

Violation 2:

```
Jun  9 09:13:06 firewall.company.com unix: securityalert: udp if=hme2
from 172.16.1.13:39176 to 10.20.30.40 on unserved port 33450
```

Judging by the protocol and port number used, a traceroute was most likely being attempted to or through the firewall.

First, besides date/time, we see the firewall's full name. Next we see where the alert came from (on the box) and what type of alert it is. From unix, type is securityalert.

Next up is the protocol, udp. Then the interface that this packet arrived on, interface hme2. Finally we see the source IP address:port, then the destination IP address with the destination port after the 'unserved port' message.

This violation was caught by no explicit rule on the firewall. Rather, it was caught by the fact that a packet arrived on a port the was not serviced by the firewall and was thus denied and logged (default firewall action).

Although no large damage to the company or firewall would likely result by allowing traceroutes themselves to the firewall, the policy states that if something is not needed it is not put in place (or, enabled). There is no good reason for giving out any more information than is absolutely needed by the firewall to outside users. Allowing this traffic through the firewall would be more troublesome as it would give outsiders information on where routers are placed in our network.

Violation 3:

```
*Jun 11 16:00:33: %SEC-6-IPACCESSLOGP: list 110 denied udp
10.23.132.21(138) -> 24.65.148.255(138), 1 packet
```

This violation would have allowed netbios information from an inside machine out to a user on the Internet. This is an example of strict egress filtering – only certain internal IP addresses are allowed onto the Internet and this was not one of them.

The format of this style of log entry is defined in '*Violation 1*'. It should be noted that this access control list is an egress list whereas the access control list in '*Violation 1*' was ingress.

This traffic was blocked at the border router with the following simple rule:

```
deny ip any any log
```

As in '*Violation 1*', this is the last entry in the access control list and is used instead of the implicit deny so that logging can be performed.

Besides being a good net citizen by not allowing spoofed IP addresses out onto the Internet from your company, strictly allowing only certain devices on the Internet (if you have that luxury) will allow you to protect your company. If an internal server gets misconfigured or mixed up and starts to transmit data to an Internet IP address (as in this example) it will be blocked by the egress filter and thus will be prevented from actually transmitting that potentially sensitive data over the Internet. That server can be fixed and no harm was done.

Violation 4:

```
Jun 11 11:41:25 firewall.company.com unix: securityalert: tcp if=hme0
from 10.54.102.23:1902 to 38.9.24.255 on unserved port 119
```

An internal user is attempting to use a news server that they are not allowed to access.

The format of this log entry is defined in '*Violation 2*'. This packet arrived on interface hme0, the internal interface. This, with the source address, confirm that an internal user has generated this error. Also to be noted is that the destination port is 119, nntp

This packet was 'caught' by a rule on the firewall that explicitly specifies the news server that internal users are allowed to connect to. As this internal user tried a different news server, a security alert was generated.

No damage would have resulted from this traffic had the firewall not stopped it. This policy was put in place not so much from a security perspective as from a corporate policy perspective.

Violation 5:

```
*Jun  7 03:22:14: %SEC-6-IPACCESSLOGDP: list 120 denied icmp
24.27.234.229 -> 10.44.177.255 (8/0), 1 packet
*Jun  7 03:22:17: %SEC-6-IPACCESSLOGDP: list 120 denied icmp
24.27.234.229 -> 10.44.178.255 (8/0), 1 packet
*Jun  7 03:22:20: %SEC-6-IPACCESSLOGDP: list 120 denied icmp
24.27.234.229 -> 10.44.179.255 (8/0), 1 packet
```

This example is a small portion of a what is most likely a smurf attack attempt.

This log format is defined in '*Violation 1*'. Please note that there are no port numbers as this is icmp. The (8/0) in each listing indicates the ICMP type and code. In this case it indicates Echo Request. Each packet is destined for a different class-C broadcast address.

As in '*Violation 1*', this violation was caught by the following rule on our border router:

```
deny ip any any log
```

This type of traffic is almost certainly malicious and would have caused grief had it been allowed through the firewall and had the machines on the destination subnets all responded with Echo Replies. Had this traffic been accidentally allowed in, the replies would have all been stopped at the firewall and would not have been allowed out to the Internet address (which was most likely spoofed). This amount of traffic could have potentially overwhelmed the firewall with thousands of responses and could have temporarily frozen the firewall or worse, crashed the firewall. In either case, a denial of service would be felt. Had the responses been allowed through the firewall onto the Internet, the poor individual at 24.27.234.229 would have certainly felt the denial of

service attack succeed against him as his machine is almost certainly less powerful than a typical corporate firewall (with less bandwidth too).

3 – Defense in depth architecture

First Question:

The network design shown below is set out to be DDOS resistant by the following features:

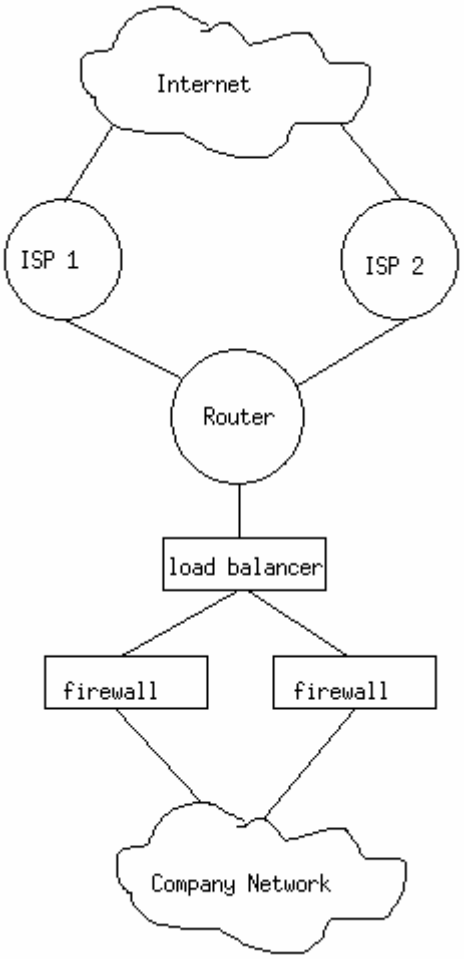

- Two different ISP's are used to connect the site to the Internet. Should one ISP have a DDOS attack launched against one of its own clients, this site will not be indirectly affected to a large degree.
- If possible, a load monitoring agreement should be set up with each ISP so that early warning may be possible of a DDOS.
- If possible, choose ISP's that do ingress filtering.
- The router should be large enough to handle all the traffic that may arrive at it over both ISP links so that while attacks may use up the bandwidth, they will not kill the router (who's logs will be needed to deal with the DDOS attack in a quicker manner).
- The router should have unneeded services (echo, chargen, etc) turned off and should deny as much broadcast and multicast traffic as is possible.
- Filtering on the router should be done by access control lists placed inbound on the interfaces whenever possible.
- The link between the router and the load balancer should have the bandwidth to contend with the possibility of such a volume of valid traffic as fills the pipes from ISP 1 and ISP 2. Should this not be realistic, then a committed access rate for various protocols should be set up on the router to reduce the chances of this happening.
- As firewall are generally slower than routers at processing packets, more than one firewall would most likely be a wise choice (with a load balancer to ensure even use of the firewalls).

- From the router down to the internal machines, all should be up to date with the latest patches.

Reference: www.sans.org/ddos_roadmap.htm

Second Question:

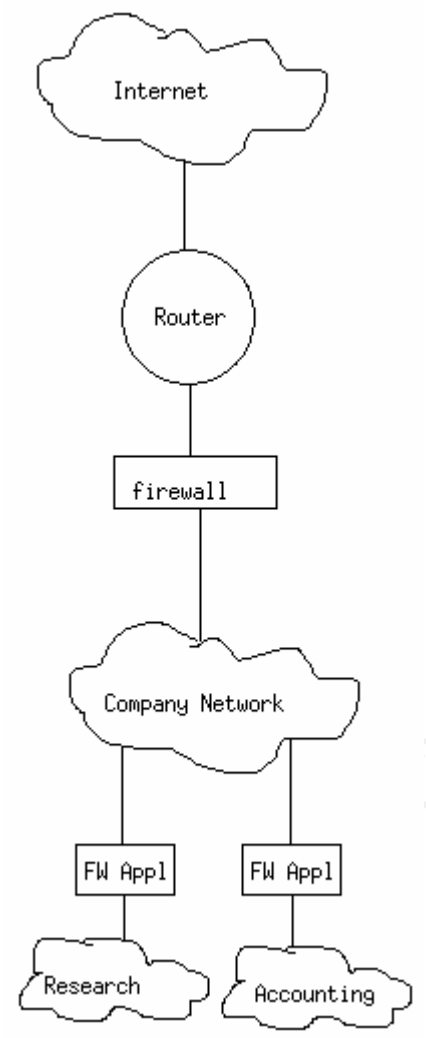

In this design diagram the four pieces of purchased hardware – router, firewall, 2 x firewall appliances (FW Appl) – have been laid out in such a fashion as to provide the internal research and accounting subnets a high degree of protection.

A router is used as the first line of defense with the firewall behind it. The router first because it can filter large amounts of traffic before that traffic hits the firewall. The firewall, being proxy-based, can then handle the remaining traffic with more attention to the detail of each protocol to ensure that malicious traffic is not accidentally passed through. More complex rule sets can be in place on the firewall than can be placed on the router.

Inside of the company's network, the two firewall appliances are used to isolate the research and accounting subnets from general company traffic. This is done to protect against unethical employees and general leaking of information.

Thus, using this design, the two critical subnetworks are protected from both the Internet and from the general company network.

4 – Create a test

Problem:

A company has expanded and now consists of two sites, each in a separate city. Management has decided that a VPN solution will be the cheapest way to get the users at the secondary site connected to the network at the main site. The main site has never had VPN and is connected to the Internet through a firewall then a router and allows outbound connections only. As a VPN server will be installed, management wants you to

include information on using this server for employees who wish to work from home. Keeping security in mind, design a network solution that would (a) connect the secondary site to the Internet, (b) add a VPN server to the main site, (c) diagram how at-home employees would connect to the VPN server.

Solution:

Given that the company's main site is already protected by a router-firewall system, it makes sense to use this scheme for the remote location for more reasons than just that it is a good solution. This company likely has administrators already trained on using the router and firewall effectively at the main site thus fewer mistakes are likely to be made when setting up the secondary site. As has already been mentioned, the router-firewall type of setup is a good solution and is in fact used by a great many companies.

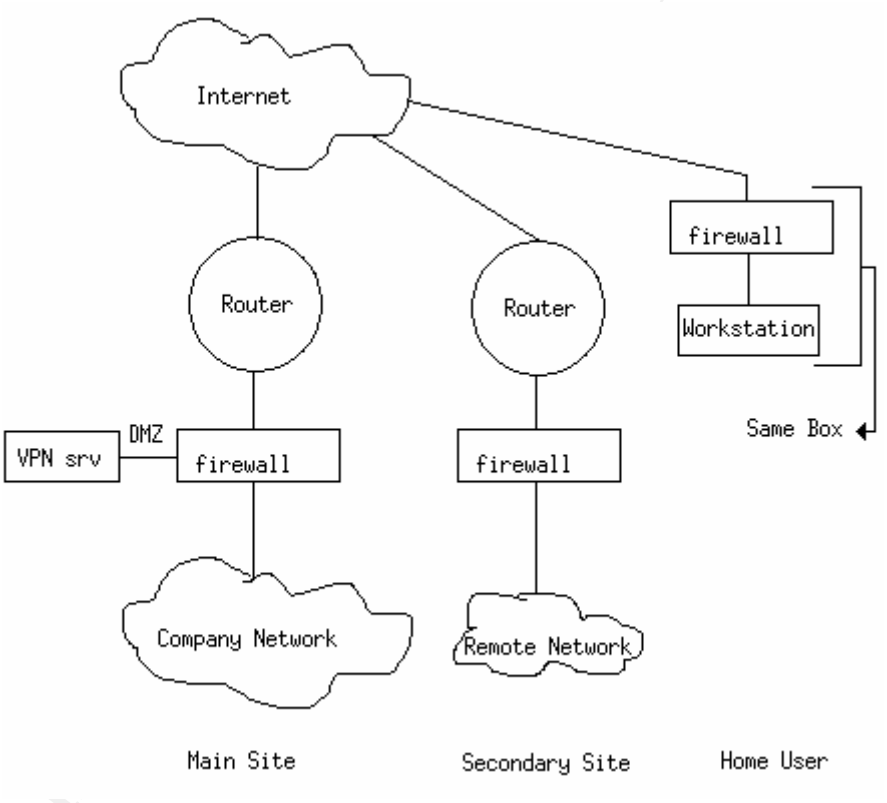

The choice to install the VPN server (VPN srv) on a new, protected network at the main site (DMZ) was made because, while the VPN server is still protected by the router and firewall combination, the company network is also protected from the VPN server. The VPN server, by its nature, accepts inbound connections from the Internet and is therefore an easier potential target than internal company network machines which never see connections straight from the Internet. By placing the firewall between the VPN server and the company network, it (firewall) can be used to limit access by VPN users to only certain internal machines. Router and firewall changes will have to be made to allow for inbound Internet traffic destined for the VPN server – either destined directly for the VPN server, or destined to a port on the firewall that then forwards the traffic to the VPN server.

The home users (as well a secondary site users) would use VPN software installed on their workstations. As more and more home users are connecting to the Internet with

high bandwidth connections, security for home users is of increasing importance. Home users are usually, when VPN'd in, easy launching points into the company network. By installing and correctly configuring a software firewall on the home user's machine, a large degree of protection is added for both the user and the company. In addition to protecting the users from attack from others on the Internet, the firewall could also be used to prevent traffic destined for internal company addresses other than the VPN server from leaving the workstation. This could happen if the VPN software decides not to tunnel some traffic that ought to be tunneled or if the user mistakenly thinks that they are VPN connected when in fact they are not and tries to send company traffic over the Internet.