Protection at the Perimeter – One Link in the Defense-in-Depth Chain

By Paul Kincaid, wnipegjets@yahoo.com

In this age of Internet worms, more needs to be done at an organization level. Nearly all worms today will infect a server on a particular network listening port and then will attempt to propagate itself out to other vulnerable servers. However, while most of these servers may have a legitimate reason for listening on the port (such as TCP port 80 for a web server), this same web server should not need to send out requests to other web servers. The concept is the same for a server that has been taken over by a malicious attacker – an organization has the ability to limit what an attacker can launch from an exploited box to other companies and organizations. This article will address securing the edge firewall for not only incoming traffic, but outgoing traffic as well.

Firewalls have become an essential tool in the fight against hackers and malicious activities. However, an improperly located, configured, or monitored firewall can give a false sense of security for the organization. Simply placing a firewall and blocking access into the network is only one step in battle. This paper will also address some additional steps necessary to better protect yourself and prevent an exploited box from becoming a platform for attack or the latest warez server.

Assumptions

A couple of assumptions need to be made for this concept to be fully implemented and understood. First is the robustness of the firewall installed in the organization. Most modern firewalls understand the concept of established TCP connections. An established session through a firewall means the firewall has seen a complete TCP 3-way handshake from the source to the destination (i.e. client to server) and will allow return traffic from the server to the client, typically on a TCP port greater than 1024. If a firewall did not understand the concept of an established session, all ports greater than 1024 would need to be allowed out of the firewall. This is commonly referred to as a "stateful" firewall.

One layer of defense that will not be addressed is the edge router. Most modern routers or multi-layer switches have firewall features built-in (commonly referred to as Access Control Lists or ACLs). This paper will specifically address securing the firewall device, but the concepts within this document can also be applied to an edge device.

One of the basic concepts to get out of this article is you not only need to protect what is coming into the organization, but also what is going out. Let's look at the example network:

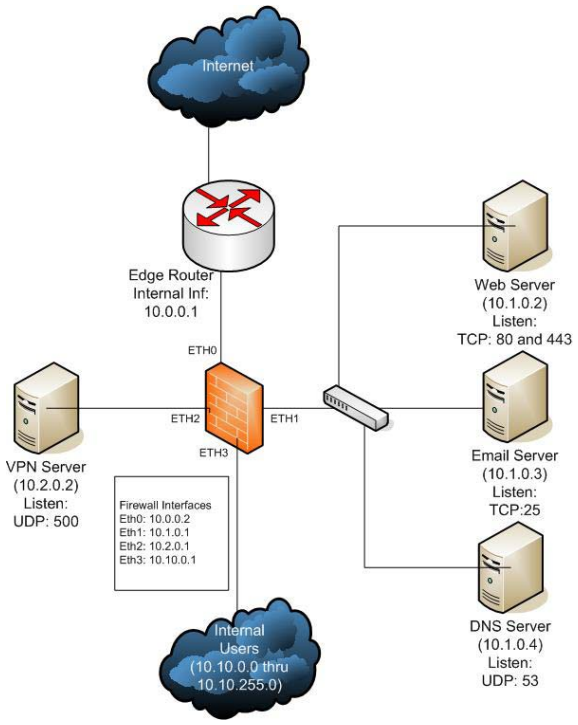

With the example network we have 3 separate segments to protect, each segment will have 2 (one inbound and one outbound) policy. Most organizations and people realize that you should only allow in what you want to allow in and deny everything else. This same concept needs to apply to the outbound traffic. Let's look at each segment individually.

DMZ

The only traffic allowed in should be TCP ports 80 and 443 to the web server, TCP port 25 to the mail server, and UDP port 53 to the DNS server. That is it, nothing else. Now, on the outbound, you should only allow the return traffic from those servers to the clients. For example, the web server should only be sending traffic out from source port 80 to clients, it should not be sending anything else out. Although this may seem simple enough, having intimate knowledge of the traffic patterns to and from the individual servers is essential to the process. Prior to implementing any new firewall policy, the systems administrators should spend some time monitoring the network (i.e. traffic sniffing) so they know the traffic patterns of the segment.

In our example network, the firewall policy on the DMZ interface should look something like this:

Allow in:

- All sources to 10.1.0.2 (web server) on TCP destination ports 80 and 443
- All sources to 10.1.0.3 (mail server) on TCP destination port 25
- All sources to 10.1.0.4 (dns server) on UDP destination port 53
- Deny all other inbound traffic

Allow out:

- 10.1.0.2 (web server) to all sources on TCP source ports 80 and 443
- 10.1.0.3 (mail server) to all sources on TCP source port 25
- 10.1.0.4 (dns server) to all sources on UDP source port 53
- Deny all other outbound traffic

VPN Subnet

The same concept as the DMZ applies to the VPN subnet – only allow in traffic to the VPN device on the appropriate VPN ports and protocols. Typically this means IP Protocols 50 and 51 and UDP destination port 500.

In our example network, the firewall policy on the VPN interface should look something like this:

Allow In:

- All source to 10.2.0.2 (VPN device) on UDP destination port 500
- All sources to 10.2.0.2 (VPN device) on IP Protocols 50 and 51
- Deny all other inbound traffic

Allow Out:

- 10.2.0.2 (VPN device) to all sources on UDP source port 500
- 10.2.0.2 (VPN device) to all sources on IP Protocols 50 and 51
- Deny all other outbound traffic

Internal Subnet

This one can get a little tricky. Although an organization may not limit where an employee can go, there are some definite services (such as Peer-to-Peer file sharing and default Microsoft services) that should absolutely be blocked at the internal side of the firewall. However, it is a better idea to only allow those ports outbound from your user's workstations that you specifically allow. The other problem with this is an educated user can simply bypass the rules by tunneling traffic across approved ports. This is a topic of an entirely different and long article.

Further, with the still all too common practice of users clicking on unknown attachments, a malicious user could deliver an email to an unsuspecting user that clicks on the attachment, which is not a virus, but a specific program that will establish a reverse SSH tunnel with the attackers remote server. This tunnel will allow the attacker to connect directly to the user's workstation without having to attack and bypass the firewall – the tunnel is established and cannot be blocked by a firewall if the tunnel is across a legitimate service – such as HTTP. Only through the use of monitoring software (IDS, IPS, etc), proxy servers, and user education could this attack be detected and stopped.

Ports that should be allowed outbound:

- Web and HTTPS (TCP/80 and TCP/443) – other ports may be necessary based on user requests. There are some web servers on the Internet that listen on port

81, 8080, or other non-standard ports. Additionally, use of a proxy server internally would provide not only better network performance, but the firewall could further be locked down to only allow the proxy server outbound on port 80, 443, etc.

- DNS (UDP/53) – only from the authorized internal DNS server or for all users to the authorized DNS server
- Mail (TCP/25) – only from the authorized internal Mail server. There should be an internal email server that users normally connect to for their mail, the mail server within the DMZ should pass mail from itself into the internal mail server over a secure connection.
- FTP (TCP/20 and 21) – this one gets a little tricky since there is passive and active FTP. Active FTP only uses TCP ports 20 and 21; however, passive FTP utilizes high ports above 1024, thus it is more difficult to limit the outbound connections.

In reality, those are the only ports that an employee should be allowed to use. However, there are some other ports that could be considered, such as AOL IM (TCP/5190), SSH (TCP/22), etc.

Overall Controls

There are a number of overall controls that also need to be addressed; however, will not be addressed in-depth in this paper.

- All server, routers, switches, and any other device should be configured to only allow authorized services to be available on the network. This means that if a web server is running on a Windows server, there is no reason the NetBIOS ports, server service, NetBIOS over TCP/IP, etc need to be running or listening. This assures that if a system is compromised in a subnet, the attacker can not easily attack other servers in that subnet – Defense-in-Depth.
- Routinely monitor and review log files of all servers and devices within the network. Consider the creation of a SYSLOG server to collect all log files out-of-band (a secondary network or serial connection) to ensure that if an attacker compromises a system, they cannot erase all evidence of their presence. Automated tools should be implemented to assist the administrators in the monitoring and review of log files.
- Utilize Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) in all subnets – especially in the DMZ, VPN, and Internal subnets. Monitor and response to alert messages. Again, the use of automated tools and signature tuning can assist the administrator in monitoring and responding to attacks against the network.
- Absolutely, positively do not use the same administrator password on all servers and devices within the network. If an attacker compromises a system and obtains the password file, the attacker can attempt to crack the passwords, at which time all servers and devices would be exposed to attack. If you must write down the passwords to remember all of the different passwords, maintain the list in a secure location (and not on a PC connected to the network – locate them within a secure safe within the organization).

- Ensure all users and administrators with remote access to the network maintain the same vigilance and policies with their remote workstations as with a system located within the Internal network. Train remote users in the monitoring, detection, and other methods to detect malicious activity – such as maintaining burglarer alarms on their homes, be observant for signs of a physical key logger, and maintain up-to-date anti-virus, host-based firewall, and host-based Intrusion Detection/Prevention. Implement systems/solutions that will force remote users to be up to date on their AV, HIDS and remote access software.

I fully understand that by utilizing this concept automatic updates from the manufacturers of the software will be broken. However, I can argue that automatic updating is inherently both insecure and could cause major problems with operational systems. The solution that I always point to is for the administrator to download all of the latest patches from, say Microsoft, to a system, verify the patches are correct through MD5 sums or PGP keys, burn the patches to a CD, install the patches on an internal testing system and perform rigorous security and operational testing, then and only then should the patches be installed from the trusted CD-ROM to production systems. This will still allow for the tight control on what goes in and comes out of your network, but will also ensure an intermediate level of testing to help avoid costly and damaging failures of patches to install properly on production systems.

No network can be 100% secure if it is connected to the Internet and providing various types of content and services to the world. However, actions can be taken by the firewall administrator to minimize the impact of a malicious attack against a server or service. In addition to a good, solid firewall policy, Intrusion Detection (IDS) and Intrusion Prevention Systems (IPS) should be employed and monitored as well as in-depth review and monitoring of the log files of all devices. This can be a daunting task for a large network; however, with some upfront work in creating scripts and other automated tools, the monitoring can become easier for the administrators and provide a greater level of vigilance. Defense-in-depth can and should be a vital concept in any network – if an attacker gets through one layer of defense, there are multiple other layers to protect the network and thwart any other malicious activities.

This paper is by no means completely comprehensive. It was designed to get the administrators to start thinking more about not just what comes into the network, but also on what goes out. Also, do not get so focused on the latest worm that is clogging the network that you completely ignore the more dangerous targeted attack against a specific server, service, or company. These attacks are typically directed only a single or very few servers within the organization utilizing 0-day (unpublished) or specific attacks – not the latest worm that is making the rounds. Only through the use of Defense-in-Depth can you provide a higher level of protection of the network against both known and unknown attacks.

Some reference material that can help with firewall design and implementation:

Building Internet Firewalls by D. Zwicky, Simon Cooper and D. Brent Chapman from O'Reilly Books

Linux Firewalls by Robert L. Ziegler from O'Reilly books

TCP/IP Illustrated Vol 1 by W. Richard Stevens from Addison-Wesley Pearson Education