



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Application Layer Firewalls vs. Network Layer Firewalls: Which Is the Better Choice?

Keith D. Maxon
August 13, 2000

The purpose of this paper is to explain the classical definitions of both a network firewall and an application firewall, and compare/contrast the two. In the process of doing so, some assumptions have to be made. Many of the benefits and drawbacks that are stated do not really come into play, as an administrator should not set up their network in this manner. The pros, cons and some of the vulnerabilities will be discussed about each firewall type. To conclude the paper, an explanation of modern firewall technology will be examined, and how the various technologies differ from the classic definitions.

Network layer firewalls run at layer 3 (Network) and sometimes 4 (Transport) of the OSI Model and are only able to make “decisions” that fall under these two layers. “One thing that is an important distinction about many network level firewalls is that they route traffic directly through them.”¹ Meaning they scan for source and destination information and allow or disallow packets based on this information. Network layer firewalls typically fall under one of the following two categories: packet filters and circuit layer gateways.

“A packet filter examines IP packets and makes a decision to accept or deny traffic based upon criteria such as source and destination IP addresses and source and destination TCP/UDP port numbers.”² Circuit layer gateways take this a step further and operate in layer 4. “As such, they can make basic authorization decisions based on source and destination IP address as well as protocol type and port.”³ This provides a higher level of flexibility in that they can make decisions on whether inbound requests to ports are valid. VLSI (very large scale integration) devices, such as routers and switches have the ability to function as network firewalls.

Network firewalls are typically used when speed is essential. Since packets are not passed to the application layer and the contents of the packet are not being analyzed, packets can be processed quicker. This can be advantageous for firewalls that scan for connections to web and email servers, especially ones that have high amounts of traffic. This is due to the fact that latency is your enemy when it comes to people accessing your site. This offers a layer of protection to your network and does not impede connectivity. Generally speaking, network firewalls are a cheaper alternative. Most logical network devices offer at least some level of packet filtering. This would allow use of pre-existing equipment to perform firewall duties. Some network operating systems also come with the ability to do packet filtering. This may prove to be an inexpensive solution, but can often produce problems. The most evident is that the firewall would be susceptible to any attacks or vulnerabilities that the operating system possesses.

Network level firewalls run on an access control list and do not provide the same high level of protection that application firewalls do, since they cannot monitor the contents of packets. The list simply verifies if the source and destination data are valid. This can

present a problem if you are actively trying to scan for vulnerabilities in the data itself. Typically network level firewalls do not provide a high level of auditing or logging. Based on how closely the traffic needs to be scanned, this may present a problem.

Network firewalls are susceptible to different exploits. Three common ones are buffer overruns, IP spoofing and ICMP tunneling. Buffer overruns typically occur when data sizes inside a buffer exceed what was allotted. “A buffer overflow condition would normally cause a segmentation violation to occur.”⁴ If we were to assume that a buffer was created with a fixed length of 500 bytes, we could send the process data exceeding that size. If carefully crafted, executable code could be inserted and ran. For example, if one were running sendmail behind the firewall, “an attacker could send specific code that will overflow the buffer of a command like VRFY and execute /bin/sh. If sendmail is running at root, /bin/sh will have root access.”⁵ Since these exploits take advantage of the application layer, a network firewall could not scan them and disallow them. IP spoofing is simply sending your data to a source, in this case a firewall and faking a source address that the firewall will trust. In this particular scenario, the hacker would be able to access internal machines since he compromised the firewall. ICMP tunneling allows a hacker to insert his data into a legitimate ICMP packet. Since the network firewall cannot probe the packet past the IP headers, it cannot deny the connection. In order for an exploit like this to work properly, a process must be in place on the other side of the firewall to strip the data out of this packet. The system has already been compromised if it has reached this point. In real life, an intelligent administrator would drop all ICMP traffic at the firewall. However, for purposes of this discussion, we see how the firewall would not be able to stop this exploit in the long term.

Application level firewalls, as the name implies, operate in the Application Layer of the OSI model. They view information as a data stream and not as a series of packets. In this way, they are able to scan information being passed over them and to ensure that the information is acceptable, based on its own set of rules. “They generally are hosts running proxy servers, which permit no traffic directly between networks, and which perform elaborate logging and auditing of traffic passing through them.”⁶

As stated earlier, these firewalls work at the application level, so they tend to be equipped with a certain level of logic. This allows the firewall to make some intelligent decisions about what to do with packets that are passing through it. An example of this ability follows: “In an early incarnation of sendmail, the original implementation of an SMTP mail server, a backdoor command was inserted to assist in debugging the application. SMTP is based on a simple, human-readable, text-based dialog between the client and server, using commands such as ‘HELO’, ‘QUIT’, and ‘DATA’. The backdoor command was ‘WIZ’, which allowed the client machine to gain root shell access on the remote sendmail server. Since neither Packet Filters nor Circuit Layer Gateways examine application data, they were vulnerable to this backdoor exploit.”⁷ In this example, an application firewall can be configured to check for a “known” vulnerability. This may prove to be cumbersome, as an administrator would have to stay on top of all possible vulnerabilities, but the option is available. Another benefit of application level firewalls is that they typically do a large amount of logging, which makes it easier to track when a

potential vulnerability happens. Another major benefit of application firewalls is that they typically support the ability to report to intrusion detection software. This allows third party software to take control of an intrusive situation and perform tasks above the capabilities of the firewall itself. This is useful if you want to monitor a hacker once they get inside instead of just blocking them or have the system send a page when an intrusion is detected.

The price you pay for the ability to scan packets for rogue data comes in performance. Since the firewall operates at the application layer, the datagram has to be passed through all the subordinate layers. The difference may not appear substantial, but when the system is scanning thousands of packets, it becomes more evident. Many people insist that the “bit stripping” or the removing of headers and passing the data up to the next level, that occurs while passing packets up and down the layers, is not at all significant. However, with the speed of machines today, the task of moving through the OSI model is typically negligible. The application firewalls will suffer a higher rate of diminishing utility. As more connections are being made to the firewall, its rate of degradation will decrease faster than the available bandwidth. By today’s standards, if an application layer firewall were to suffer a solid performance hit, it is more likely that it is related to I/O cycles required for logging and auditing than “bit stripping.”

Due to the amount of work the firewalls must do, application firewalls are less susceptible to attacks that hide data in legitimate traffic and more susceptible to distributed denial of service (DDOS) attacks. If enough data is forced on the firewall it can cease to operate. The high number of service level vulnerabilities that currently exist can also compromise application firewalls. For example, sendmail and DNS have numerous well-known exploits. If the firewall is allowing SMTP traffic or DNS traffic to pass through and a hacker has access to one of the many exploits, typically the firewall will allow the data to pass, unless elaborate rules are established. Setting such elaborate rules usually proves to be burdensome to most administrators, so this type of exploit is usually left unchecked.

In print, it would appear that what one firewall has as a benefit, the other has as a drawback. In reality, the delineation between network layer firewalls and application layer firewalls is quickly diminishing. Modern firewalls perform some tasks in both the network and application layer. Many network IOS’s have the ability to scan traffic for vulnerabilities beyond layer 3, even though it may be a layer 3 device. “When viewed as a whole, Circuit Layer Gateways do not operate purely at layer 4. They have become hybrid software implementations to address the need for stringent Internet security. It is generally marketed as ‘Stateful Multi-Layer Inspection’, which means the software operates at many layers. Conversely, Application Layer firewalls do not solely function at the application layer. For example, in the Axent Raptor Firewall, it is possible to pass traffic through local-tunnels, a stateless layer 3 mechanism, or layer 4 Generic Service Proxies with no application data scanning.”⁸ Firewalls that fully function in the network and application layer are not developed fully as of yet, but the advances in the technology should be considered. It is also important to note that many application level firewalls

offer some level of clustering that allow the firewall to overcome its speed issue. This allows one to add more machines as needed.

No one firewall will meet one hundred percent of everyone's needs. Before purchasing a secure firewall solution, make sure to fully analyze the pros and cons. As a general rule, if speed is the most important feature, look into the network layer firewall. If security is a top concern, then look into an application layer firewall. "In a perfect world, you would have an application proxy securing your corporate network, but a network layer firewall to protect your web presence, without impeding performance."⁹ Ultimately a firewall serves more for piece of mind than a security device. In the end, a hacker is more likely to look for another way in, such as social engineering passwords from the staff of a company, using a war dialer to locate modems on a network to dial in and bypass the firewall entirely or look for exploits on a mail or web server that would allow them to pass through the firewall legitimately. This is due in part, to the high level of security that firewalls provide. Hackers will always look for the easiest route into the system first. It's very similar to locking the car doors even though a thief can still get into your car by breaking the windows. The locked doors have forced the thief to go in a different route. This does not mean that a firewall should not be put into place. Make sure that policies are set up to cover all security related aspects of the LAN. Also remember that no matter how powerful the firewall is it is only as strong as the policy enforcement. Ensure that the firewall is up to date on security vulnerabilities and all access lists are accurate. If this is not done, it will quickly become another doorstop in the organization.

Sources:

¹ Curtin, Matt. "Internet Firewalls: Frequently Asked Questions." November 25, 1999.

URL: <http://www.interhack.net/pubs/fwfaq/>

² Unknown. "High Availability Security Solutions." June 8, 2000. URL:

<http://www.radware.com/support/papers/fwlb.html>

³ Unknown. "Application Proxy vs. Stateful Inspection Firewall Technology." June 1,

2000. URL:<http://www.firetower.com/applicationproxy.html>

⁴ McClure, Stuart "Buffer Overflow Attacks." Hacking Exposed. Page 214-215.

September 10, 1999. <http://www.hackingexposed.com/>

⁵ McClure, Stuart "Buffer Overflow Attacks." Hacking Exposed. Page 214-215.

September 10, 1999. <http://www.hackingexposed.com/>

⁶ Curtin, Matt. "Internet Firewalls: Frequently Asked Questions." November 25, 1999.

URL: <http://www.interhack.net/pubs/fwfaq/>

⁷ Unknown. "Application Proxy vs. Stateful Inspection Firewall Technology." June 1,

2000. URL:<http://www.firetower.com/applicationproxy.html>

⁸ Unknown. "Application Proxy vs. Stateful Inspection Firewall Technology." June 1,

2000. URL:<http://www.firetower.com/applicationproxy.html>

⁹ Unknown. "Application Proxy vs. Stateful Inspection Firewall Technology." June 1,

2000. URL:<http://www.firetower.com/applicationproxy.html>