Stateful Inspection

Rainer Kampp

Darmstadt, March 2016

Abstract. Stateful Inspection, invented by CheckPoint Software Technologies, is the de facto technology for enterprise-class network security solution firewalls. In order to provide robust security, a firewall must track and control the flow of communication passing through it. To reach control decision for TCP/IP based services (whether to accept, reject, authenticate, encrypt and/or log communication) it is not sufficient to examine packets in isolation. State information from past communications or applications, is an essential factor in making the network secure. Over the past several years, enterprise firewalls have become the staple of network security architectures. Designed primarily to provide access control to network resources, firewalls have been successfully deployed in large majority of networks.

Table of Contents

1 OSI and TCP/IP

1.1 The OSI-Model

The OSI-Model (Open System Interconnection) describes the communication between computer systems with 7 independent protocol layers.

- Layer 7 Application
- Layer 6 Presentation
- Layer 5 Session
- Layer 4 Transport
- Layer 3 Network
- Layer 2 Data Link
- Layer 1 Physical

It is very useful to know these layers because this is essential for understanding all security measures handled by firewalls. The Physical Layer describes the electrical specifications of the connection. The Data Link Layer connects the physical part of the network with the NIC (Network Interface Card) of the computer system. Each network card has its own unique physical address, called MAC (media access control) address.

The Network Layer is responsible for the routing of the data stream. It also handles the relationship between the MAC address (physical address) and the IP address (logical address). The Transport Layer ensures reliable connectivity from end-to-end, The Session Layer makes sure that information is in synchronization on both sides and the Presentation Layer guarantees that the received format of the data is useful for the system. The Application Layer for example determines if the running application needs network connection and then manages the requests from the running program to the other layers.

1.2 The TCP/IP-Model

The TCP/IP-Model describes the communication between computers in abstraction to the OSI-Model with 4 independent protocol layers.

- Layer 4 Application
- Layer 3 Transport
- Layer 2 Internet
- Layer 1 Physical

TCP (the transmisson control protocol) is responsible for breaking up the messages into datagrams, putting a header at the front of each datagram, reassemble them at the designation computer, resending anything that get lost and putting the datagrams in the right order. The TCP-Layer sends each of these datagrams to the Internet (IP) - Layer. The IP-Layer does not care about what is in the datagram or even in the TCP header. The flags SYN (Synchronize) and ACK

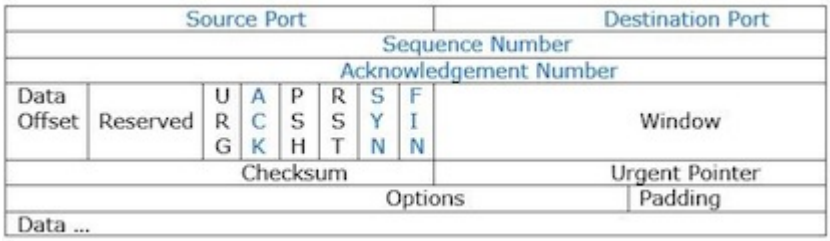

Fig. 1. TCP-Header

(Acknowledge) are used to initiate a normal TCP connection. The flag FIN (Finish) is used to finish a connection. For example a web-client initiates the connection to the addressed web-server by sending a SYN packet to the well-know port 80 (http). The server then responds with a SYN/ACK packet the client finally responds with an ACK packet and the connection is established. This procedure is called the "three-way handshake". The TCP-Layer sends each of these datagrams to the Internet(IP)-Layer. The IP-Layer does not care about what is in the datagram or even in the TCP header. IP is simply responsible for the routing of the datagrams. The IP-Header contains additional fields. The flags

Fig. 2. IP-Header

and fragment offset are used to keep track of the pieces when a datagram has to slip up. This can happen when datagrams are forwarded through a network for which they are too big. The time to live is a number that is decremented whenever a datagram passes through a system. When it goes to zero, the datagram is discarded and this prevents routing-loopbacks. The type of service describes what protocol is used where the service TCP is indicated with 6, UDP with 17 and ICMP with 1. UDP (the user datagram protocol) is designed for applications where no sequences of datagrams need to put together. It is a connectionless transport protocol. UDP is used by applications that do not require the level of service of TCP or that wish to use communications services (multicast or broadcast delivery) not available from TCP.

As UDP is a connectionless protocol an UDP-Header is shorter than a TCP-header. ICMP (the Internet control message protocol) is used for error messages for the TCP/IP software itself and fit into one datagram. This protocol is used for network troubleshooting, commands like "ping" or "tracert" use parts of the available message types such as "echo request", "echo reply", "destination unreachable", "time exceeded" and are very useful for network administrators to monitor the route the packets take to reach the destination IP address.

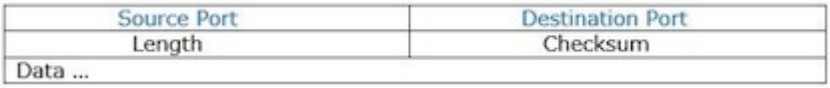

Fig. 3. UDP-Header

2 Packet Filter Firewall

2.1 Static Packet Filter Firewall

The static packet filter, historically implemented in routers, examine each and every packet at the network layer - OSI layer 3 - and compare them to the configured access lists. The administrator can define rules that determine which packets are accepted and which packets are denied. These rules are called the security policy.

2.2 Dynamic Packet Filter Firewall

The dynamic (stateful) packet filter is an advanced packet filter that operates up into the transport layer - OSI layer 4 - to collect additional state information. In simplest terms, the typical dynamic packet filter is aware of the difference between a new and an established connection. Once a connection is established, it is entered into a table and all following packets are compared to this table. When the packet is found to be an existing connection, it is allowed to pass without any further inspection. This so called state awareness provides measurable performance benefit.

3 Application Level Gateway Firewall

An Application Level Gateway intercepts incoming and outgoing packets, runs proxy servers that copy and forward information across the gateway preventing any direct connection between a trusted server or client and an untrusted host. The proxies are application specific (such as http, ftp, smtp etc.) and examine the entire packet and can filter at the application layer of the OSI model.

4 Stateful Inspection Firewall

[1] Stateful inspection combines the many aspects of dynamic packet filtering and application level gateways (Proxies). While stateful inspection has inherent ability to inspect all seven layers of the OSI model, most installations only operate as a dynamic packet filter at the network layer because of the dramatic impact of performance.

4.1 True Stateful Inspection

To track context, a firewall must inspect the content of packet payloads to ensure that each packet entering the network meets the expected parameters and attributes of the communication session. This guarantees that malicious packets that do not fit the context of the communication cannot circumvent firewall security. The following are examples of state and context-related information that a firewall should track and analyze.

- Packet header information (source address, destination address, service, source port, destination port, packet length)
- connection state information (which ports are being opened for which connection)
- sequence and acknowledgement numbers, fragment offset
- packet reassemble

True Stateful Inspection means tracking the state and context of all communications.

5 Network and Application Security

If a network relies on an application level security method, incoming and outgoing packets cannot access services for which there is no proxy. An application level gateway that runs FTP and HTTP proxies, only those packets generated by these services could pass through the firewall. All other services would be blocked. Strong application proxy that inspects header length can eliminate an entire class of buffer overrun attacks. But these proxies must be written securely. Historically some vendors have introduced buffer overruns within the application gateway itself. With proxy firewalls you must establish a TCP session with the firewall itself if you want to access a service on the other side of the firewall. If the proxy detects no problems, the firewall establishes another connection with the destination device.

This is the primary advantage of application level gateway firewalls because no direct connections are allowed through them. The disadvantage is that these firewalls are not transparent for the internal hosts, which want to connect to an external server. Each internal host must be configured to be aware of the firewall

and must have a client software that is designed to be capable of communicating with the proxy software on the firewall. Nowadays in modern secure environments the client do not either have to be aware of the firewall or run special software to communicate with external network.

6 Network and Application Threads

[4]Any traffic not adhering to strict protocol or application standard must be closely analyzed before it is permitted into the network, otherwise business-critical applications may be put at risk. For example binary data in HTTP headers are prohibited by the official HTTP standard. Most of the firewall in the past do not check this and as a result many hackers launch attacks by including executable code in HTTP headers. Also the HTTP standard does not limit header length, excessive length should be blocked or flagged to reduce the chance of buffer overflows and to limit the size of code that can be inserted using the overflow threat. Malicious data can also enter the internal network by embedding itself in URLs. An application such as an email-client could automatically execute an HTML-embedded URL. If the URL was malicious, damage to the network or the users system may occur. Therefore access to potentially malicious URLs should be blocked and limited.

Not only application-layer communications introduce malicious data to a network, the application itself might perform unauthorized operations. A network security solution must have the ability to identify and control such operations. A firewall should place connection restrictions on particular file names and controls potentially hazardous FTP commands like PUT, GET, SITE and REST. For example a security policy may require operational restrictions on all files containing the word "payroll".

Preventing malicious manipulation of network-protocols (e.g. ICMP) is a crucial requirement for multi-level security firewalls. ICMP allows one network node to ping or send an echo request to other network nodes to determine their operational status. This capability can be used to start a "smurf" DoS (Denial of Service) attack. The smurf attack is possible because standard ICMP does not match requests to replies. Therefore an attacker can send a ping with spoofed source IP address to an IP broadcast address. The IP broadcast address reaches all IP addresses in a given network. All machines within the pinged network send echo replies to the spoofed and innocent source IP. Too many pings and responses can flood the spoofed network and deny access to legitimate traffic. Dropping replies that do not match requests can block this type of attack. Stateful inspection handles this attack by creating virtual session information for connectionless protocols like UDP and ICMP.

Another network-layer event is the PortScan. A port scan does what the name implies: a hacker for example scans a range of ports on a target host in hopes of identifying and exploiting weakness on running applications. The reconnaissance that a port scan performs is a hazard then can lead to an attack. A security gateway must be able to raise alerts and block or shutdown communications from source of the scan.

well-known ports	service	ports	service
20	ftp data-tranfer	110	pop3
21	ftp commands	143	imap
23	telnet	443	https
25	smtp	3306	mysql
43	whois	5060	sip
53	dns	8080	http
80	http	10000	webadmin

Fig. 4. well-known ports

At least a firewall must defend against variations of well-known attacks (Code Red or Nimda). There are firewall devices where the attack patterns (malicious URL) must match identically to those in the firewall database. Any variation of an attack, no matter how trivial, will traverse the firewall undetected. This is a direct result of the fact that some firewalls does not support regular expression matching, which gives the administrators the ability to look for attack variants using wild card definitions.

7 Performance and Security Considerations

Performance (packet throughput and simultaneous connections) and security on the other side are both aspects that must be considered when you use a firewall as a security measure. [3] The security expert Bill Stout wrote on the firewall mailing list: "The purpose of a security device is to protect a network, not to be fast. Fast is what airline travelers want when passing through airport security, secure is what they want when they tumble through the air after their plane blows up."

As we learned the highest level of protection is achieved by Application Level Gateway (Proxy) firewall because all the 7 OSI Layers were inspected and analyzed.The Stateful Inspection firewall also looks up to the 7 layers but not so deep in content. It provides "light proxies", that do not intercept the client/server communication. Therefore these firewalls should be faster than Proxy firewalls but cannot support the same protection level. The packet firewall even the dynamic packet filter have the highest packet throughput but less security protection, because only the OSI layer 3 respective layer 4 is examined.

8 Conclusion

Stateful packet filters are faster than application level gateways. They have better performance but less security. To achieve the highest level of protection in combination with the highest network performance you can use both technologies. The application level gateway (proxy) as external firewall to the Internet, to achieve the highest security by an adequate throughput and a True Stateful Inspection packet filter as internal firewall to the Intranet, where high TCP/IP traffic throughput is more necessary and important.

References

[1] Paul Henry, "An Examination of Firewall Architectures", April 2001

[2] Rainer Kampp, "Stateful Inspection", GIAC Certified Professionals, June 2003 `http://www.giac.org/certified-professional/rainer-kampp/105165`

[3] Trusted Information System Inc., "Application Gateways and Stateful Inspection", Januar 1998

[4] Check Point Software Technologies Ltd., "Check Point Application Intelligence", 2003