

Prepared by Dr. Mohammed Anbar


CYB 603
Firewall

Lecture 4



- Explain firewall concepts,
- Explain firewall actions and types,
- Explain the attacks on packet filtering firewalls and the appropriate countermeasures,
- Explain firewall deployment architecture,
- Understand firewall policies.

- A firewall is a part of a computer system or network designed to stop unauthorized traffic from flowing from one network to another.
- In firewall, all inbound and outbound network traffic will be checked by the firewall.
- If the firewall identifies packets as legitimate packets, then it allows them to go to the other side of the firewall otherwise the firewall will not allow the packet to the intended destination.



IPv4 Firewall

IPv6 Firewall

Hide empty chains

Reset Counters

Restart Firewall

Table: Filter

Chain INPUT (Policy ACCEPT) (Traffic)

Pkts.	Traffic	Accepted	Denied	Rejected
36.84 K	3.26 MB			
4.75 M	1.56 GB			
3.80 M	892.10 MB	ACCEPT		
145.63 K	9.72 MB	syn_flood		
372.08 K	332.92 MB	zone_lan_input		
14.84 K	874.72 KB	zone_wan_input		
506.23 K	329.45 MB	zone_DMZ_input		
58.72 K	3.60 MB	zone_Kids_input		

- Packets flowing through a firewall can have one of three outcomes:
 - **Accepted:** permitted through the firewall
 - **Denied:** not allowed through with no indication of failure
 - **Rejected:** not allowed through, accompanied by an attempt to inform the source that the packet was rejected
- Policies used by the firewall to handle packets are based on several properties of the packets being inspected, including the
 - Protocol used, such as TCP or UDP
 - The source and destination IP addresses
 - The source and destination ports
 - The application-level payload of the packet (e.g., whether it contains a virus).

Firewall actions rely on two types of filtering which are

Ingress filtering


Egress filtering

Firewall actions rely on two types of filtering. The type is ingress filtering which inspects the incoming traffic to safeguard an internal network and prevent attacks from outside while the second type is egress filtering which inspects the outgoing network traffic and prevents the user in the internal network to reach out to the outside network.

- **Ingress filtering:** *Inspects the incoming traffic to safeguard an internal network and prevent attacks from outside.*
- **Egress filtering:** *Inspects the outgoing network traffic and prevents the users in the internal network to reach out to the outside network. For example, blocking social networking sites in school*

- Firewall checks every single packet traversing the network. For example, when any user downloads a big-size file from the internet, the data of the file is divided into small chunks called packets. Each packet consists of two parts which are the packet header and packet payload.
- Packet header consists of info such as source and destination IP while the payload consists of the packet content

What does data packet mean?



Two fundamental approaches to creating firewall policies (or rulesets)

- **Blacklist** approach (default-allow)

All packets are allowed through except those that fit the rules defined specifically in a blacklist.

- Pros: flexible in ensuring that service to the internal network is not disrupted by the firewall
- Cons: unexpected forms of malicious traffic could go through

- **Whitelist** approach (default-deny)

Packets are dropped or rejected unless they are specifically allowed by the firewall

- Pros: A safer approach to defining a firewall ruleset
- Cons: must consider all possible legitimate traffic in rulesets

1. The firewall is required to be used for all internal and external traffic.
2. Only traffic that has been authorized, as specified by the local security policy, will be permitted to pass.
Different firewalls kinds are utilized, and they apply different security policies.
3. Trusted computer systems are suitable for hosting a firewall and are often required in government applications.

General Techniques That Firewalls Use To Control Access And Enforce The Site's Security Policy


- **Service control:** determines the types of Internet services that can be accessed, inbound or outbound. The firewall may filter traffic on the basis of IP address, protocol, or port number; may provide proxy software that receives and interprets each service request before passing it on; or may host the server software itself, such as a Web or mail service.
- **Direction control:** establishes the direction in which specific service requests may be initiated and permitted to pass through the firewall.


General Techniques That Firewalls Use To Control Access And Enforce The Site's Security Policy (cont.)

- **User control:** regulates access to a service according to the user making the request. Users inside the firewall perimeter normally have access to this capability (local users).
- **Behavior control:** regulates how specific services are used. For instance, the firewall might filter email to filter out spam, or it might allow outside access to just some of the data on a local Web server.



- This firewall controls traffic based on the information in packet headers, without looking into the payload that contains application data.
- Doesn't maintain the states of packets. Also called Stateless Firewall.
- Also, in this type of firewall, predetermined entries that contain the IP address and port are saved into the access control entries (ACEs) table.





Rule Set A

action	ourhost	port	theirhost	port	comment
block	*	*	SPIGOT	*	we don't trust these people
allow	OUR-GW	25	*	*	connection to our SMTP port

Rule Set B

action	ourhost	port	theirhost	port	comment
block	*	*	*	*	default

Rule Set C

action	ourhost	port	theirhost	port	comment
allow	*	*	*	25	connection to their SMTP port

Rule Set A

action	ourhost	port	theirhost	port	comment
block	*	*	SPIGOT	*	we don't trust these people
allow	OUR-GW	25	*	*	connection to our SMTP port

Inbound mail is allowed (port 25 is for SMTP incoming), but only to a gateway host. However, packets from a particular external host, SPIGOT, are blocked, because that host has a history of sending massive files in e-mail messages.

Rule Set B

action	ourhost	port	theirhost	port	comment
block	*	*	*	*	default


This is an explicit statement of the default policy. All rulesets include this rule implicitly as the last rule.

Rule Set C

action	ourhost	port	theirhost	port	comment
allow	*	*	*	25	connection to their SMTP port

This ruleset is intended to specify that any inside host can send mail to the outside. A TCP packet with a destination port of 25 is routed to the SMTP server on the destination machine. The problem with this rule is that the use of port 25 for SMTP receipt is only a default; an outside machine could be configured to have some other application linked to port 25. As this rule is written, an attacker could gain access to internal machines by sending packets with a TCP source port number of 25.

- Stateful firewalls can tell when packets are part of legitimate sessions originating within a trusted network.
- Stateful firewalls maintain tables containing information on each active connection, including the IP addresses, ports, and sequence numbers of packets.
- Using these tables, stateful firewalls can allow only inbound TCP packets that are in response to a connection initiated from within the internal network.



Source Address	Source Port	Destination Address	Destination Port	Connection State
192.168.1.100	1030	210.22.88.29	80	Established
192.168.1.102	1031	216.32.42.123	80	Established
192.168.1.101	1033	173.66.32.122	25	Established
192.168.1.106	1035	177.231.32.12	79	Established
223.43.21.231	1990	192.168.1.6	80	Established
2122.22.123.32	2112	192.168.1.6	80	Established
210.922.212.18	3321	192.168.1.6	80	Established
24.102.32.23	1025	192.168.1.6	80	Established
223.21.22.12	1046	192.168.1.6	80	Established

- Controls input, output, and access from/to an application or service.
- Acts as an intermediary by impersonating the intended recipient.
- The client's connection terminates at the proxy and a separate connection is initiated from the proxy to the destination host.
- Data on the connection is analyzed up to the application layer to determine if the packet should be allowed or rejected.



- Tiny fragment attacks: The hacker creates incredibly small fragments using the IP fragmentation option and forces the TCP header data into a separate packet fragment. The goal of this attack is to get around filtering policies that rely on TCP header data.
- A packet filter will typically decide whether to filter a packet based on its initial fragment. All following packet fragments are filtered out merely because they are part of the packet whose first fragment was rejected.
- The attacker hopes that the filtering firewall examines only the first fragment and that the remaining fragments are passed through.
- Imposing a requirement that the first fragment of a packet must have a predetermined minimum amount of the transport header will prevent a tiny fragment attack.

- **IP address spoofing:** The intruder transmits packets from outside with an inside host's address in the source IP address field. The attacker thinks that using a spoofed address will enable infiltration of systems that use simple source address security, which accepts packets from only a small number of trusted internal hosts.
- The countermeasure is to discard packets with an inside source address if the packet arrives on an external interface. In fact, this countermeasure is often implemented at the router external to the firewall.


- **Option 1: Bastion Host**

The firewall is positioned between the internal and external networks in this basic configuration. Simple networks are ideally suited to this topology. Since there is only one border, if someone manages to get through the firewall, they have unrestricted access to the secured network.


- **Option 2: Staging Area or Demilitarized Zone (DMZ)**

This topology allows organizations to host servers that face the internet directly and separates the trusted network and the Internet, thus allowing the users to access the internet securely. If a malicious user manages to compromise the firewall, he or she will not have access to the intranet services (provided the firewall is properly configured).

- This is the most commonly deployed architecture. The DMZ hosts all the servers offering public services, which face the Internet.
- The private zone contains all internal network resources such as the file server, the application server, the database servers, user workstations, and printers, which do not have any business connecting to the Internet.
- The DMZ zone hosts your public Web server, mail server, DNS servers, and other similar systems.



To protect private networks and individual machines from the dangers of the greater Internet, a firewall can be employed to filter incoming or outgoing traffic based on a predefined set of rules called **firewall policies**.



- A firewall policy dictates how firewalls should handle network traffic for specific IP addresses and address ranges, protocols, applications, and content types (e.g., active content) based on the organization's information security policies.
- Firewall policy should be documented in the system security plan and maintained and updated frequently as classes of new attacks or vulnerabilities arise, or as the organization's needs regarding network applications change. The policy should also include specific guidance on how to address changes to the ruleset. There are four types of policies which are :
 - Policies Based on IP Addresses and Protocols
 - Policies Based on Applications
 - Policies Based on User Identity
 - Policies Based on Network Activity

- Firewall policies should only permit appropriate source and destination IP addresses to be used. Specific recommendations for IP addresses include: Traffic with invalid source or destination addresses should always be blocked, regardless of the firewall location.
- Examples of relatively common invalid IPv4 addresses are 127.0.0.0 to 255.255.255.255 (also known as the localhost addresses) and 0.0.0.0 (interpreted by some operating systems as localhost or a broadcast address).
- These have no legitimate use on a network. Also, traffic using link-local addresses (169.254.0.0 to 169.254.255.255) should be blocked.

- Outbound traffic with invalid source addresses should be blocked (this is often called egress filtering).
Systems that have been compromised by attackers can be used to attack other systems on the Internet; using invalid source addresses makes these kinds of attacks more difficult to stop.
- Blocking this type of traffic at an organization's firewall helps reduce the effectiveness of these attacks.
- Incoming traffic with a destination address of the firewall itself should be blocked unless the firewall is offering services for incoming traffic that require direct connections. For example, if the firewall is acting as an application proxy.

- Every organization needs a firewall that can filter IPv6 traffic, regardless of whether it enables IPv6 traffic to enter its internal network.
- The administrative interface should allow administrators to clone IPv4 rules to IPv6 addresses to make administration easier.
- The firewall should be able to block IPv6-related protocols such as 6-to-4 and 4-to-6 tunnelling, Teredo, and Intra-site Automatic Tunnel Addressing Protocol (ISATAP) if they are not required.

- An application server typically listens on one or more fixed TCP or UDP ports. Some applications use a single port, but many applications use multiple ports. For example, although SMTP uses TCP port 25 for sending mail, it uses TCP port 587 for mail submission.
- Similarly, FTP uses at least two ports, one of which can be unpredictable, and while most web servers use only TCP port 80, it is common to have websites that also use additional ports such as TCP port 8080.
- Some applications use both TCP and UDP; for example, DNS lookups can occur on UDP port 53 or TCP port 53.

- For incoming TCP and UDP traffic, deny-by-default policies should be utilized, just like for other components of firewall rulesets. Outgoing TCP and UDP traffic typically has less restrictive policies since most organizations allow users to access a large variety of external apps hosted on millions of different external hosts.
- In addition to allowing and blocking UDP and TCP traffic, many firewalls are also able to report or block malformed UDP and TCP traffic directed towards the firewall or to hosts protected by the firewall.
- This traffic is frequently used to scan for hosts and may also be used in certain types of attacks. The firewall can help block such activity—or at least report when such activity is happening.

- Attackers can use various ICMP types and codes to perform reconnaissance or manipulate the flow of network traffic. However, ICMP is needed for many useful things, such as getting reasonable performance across the Internet.
- Some firewall policies block all ICMP traffic, but this often leads to problems with diagnostics and performance.
- Other common policies allow all outgoing ICMP traffic but limit incoming ICMP to those types and codes needed for Path Maximum Transmission Unit (PMTU) discovery (ICMP code 3) and destination reachability.

- To prevent malicious activity, firewalls at the network perimeter should deny all incoming and outgoing ICMP traffic except for those types and codes specifically permitted by the organization. For ICMP in IPv4, ICMP type 3 messages should not be filtered because they are used for important network diagnostics.
- The ping command (ICMP code 8) is an important network diagnostic, but incoming pings are often blocked by firewall policies to prevent attackers from learning more about the internal topology of the organization's network.
- For ICMP in IPv6, many types of messages must be allowed in specific circumstances to enable various IPv6 features.

- The application firewall policy is defined by a collection of rule sets. A rule set defines the rules that match the application ID detected, based on the application signature.
- After an application firewall policy is created by adding rules, you can select that policy to be the active policy on your device.
- The application firewall policy identifies the application ID as an unknown application ID under the following circumstances:
 - No application ID matches the traffic.
 - The system encounters an error when identifying the application.
 - Application ID is not identified during failover sessions.

- Traditional packet filtering does not see the identities of the users who are communicating in the traffic traversing the firewall.
- So firewall technologies without more advanced capabilities cannot have policies that allow or deny access based on those identities.
- However, many other firewall technologies can see these identities and therefore enact policies based on user authentication.
- One of the most common ways to enforce user identity policy at a firewall is by using a VPN.

- Many firewalls allow the administrator to block established connections after a certain period of inactivity. For example, if a user on the outside of a firewall has logged into a file server but has not made any requests during the past 15 minutes, the policy might be to block any further traffic on that connection.
- Time-based policies are useful in thwarting attacks caused by a logged-in user walking away from a computer and someone else sitting down and using the established.
- However, these policies can also be bothersome for users who make connections but do not use them frequently. For instance, a user might connect to a file server to read a file and then spend a long time editing the file.

Wack, J., Cutler, K., & Pole, J. (2002). *Guidelines on firewalls and firewall policy*. BOOZ-ALLEN AND HAMILTON INC MCLEAN VA.

Stallings, W. (2006). *Cryptography and network security principles and practices* 4th edition



Thank You