

Firewall

Introduction to a Firewall

A firewall is basically the first line of defense for your network. It is a hardware device or software program that restricts and controls the flow of traffic between networks, typically between an internal corporate network and the Internet. Firewalls can provide secure gateway services between internal networks and protect networked computers from intentional intrusion that can compromise confidentiality or result in data corruption or denial of service. The primary purpose of a firewall is to keep uninvited guests from browsing your network.

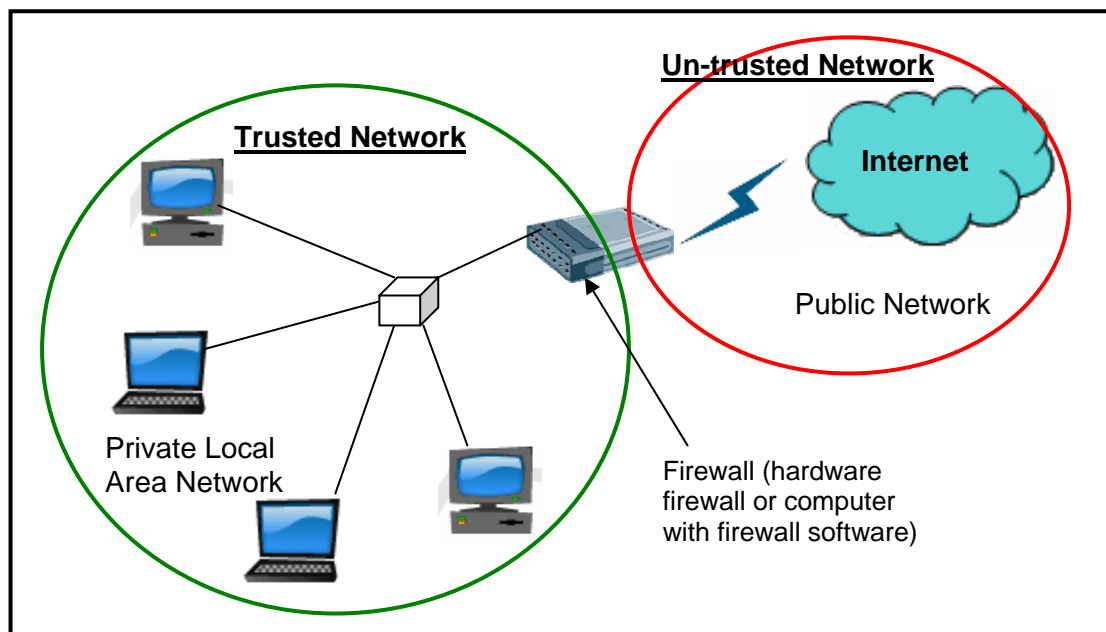


Figure 1. Firewall providing protection to a Local Network

A firewall can be thought of as a pair of mechanisms: one which exists to block traffic, and the other which exists to permit traffic. It allows you to establish certain rules to determine what traffic should be allowed in or out of your network. Depending on the type of firewall implemented, you can restrict access to specific IP addresses or domain names, or block certain types of traffic by disallowing access to the TCP/IP ports they use.

A firewall has the following characteristics:

- **Single Point of Contact**

A firewall is a single point of contact between two or more networks. All traffic must pass through this single point of contact. This is to ensure that unauthorized users from the Internet have no access to confidential information.

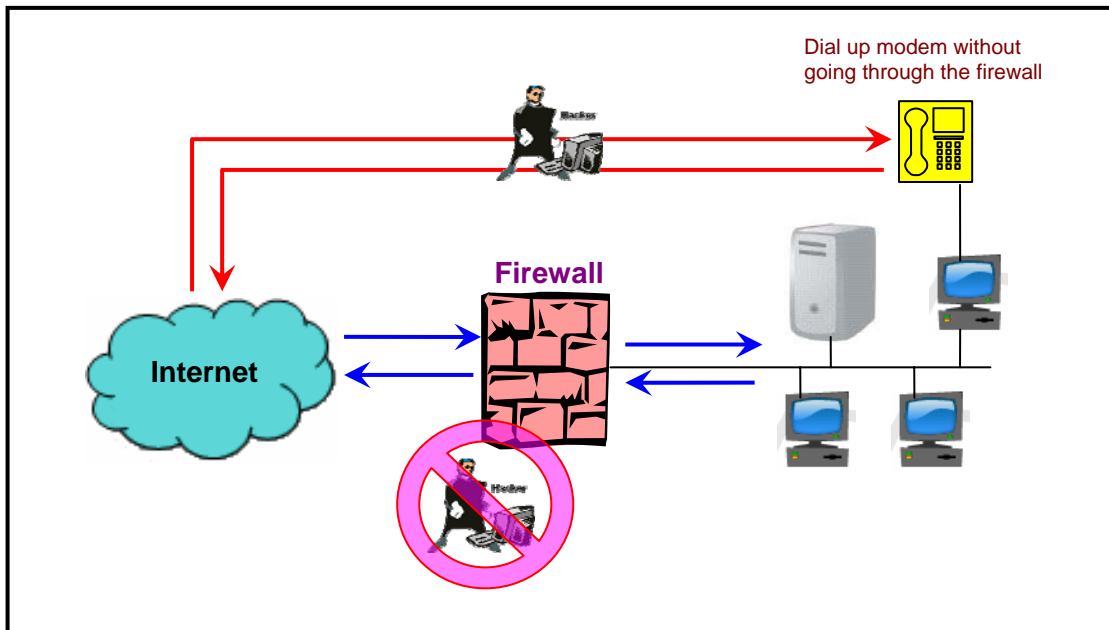


Figure 2. Firewall with Single Point of Contact

- **Controlled Traffic**

A firewall can control and authenticate the traffic that passes through it. The administrator must enable the user authentication at the firewall so that only those who are authenticated by with a valid user name and password are allowed to access to the Internet.

- **Logged Traffic**

All traffic that passes the firewall is logged. You can examine the logs to determine the details of requests that the firewall has processed. You can also analyze the log files to determine whether any security breach was attempted on the network.

Type of Firewalls

There are basically four mechanisms used by firewalls to restrict traffic. One device or application may use more than one of them to provide more in-depth protection. The four mechanisms are **packet filtering**, **circuit-level gateway**, **proxy server** and **application gateway**.

Packet Filter

The most basic and fundamental type of firewall is the packet filter. A packet filter intercepts all traffic to and from the network and evaluates it against the rules you provide. Packets are analyzed against a set of filters. Typically, the packet filter can assess the source IP address, destination IP address, TCP or UDP source port, TCP or UDP destination port and the protocol number that resides in a field in the IP header. It is these criteria that you can filter to allow or disallow traffic from certain IP addresses or certain ports. It is implemented inside the firewalls and prevents suspicious traffic from reaching the destination network. Filtered routers protect all the machines on the destination network from suspicious traffic.

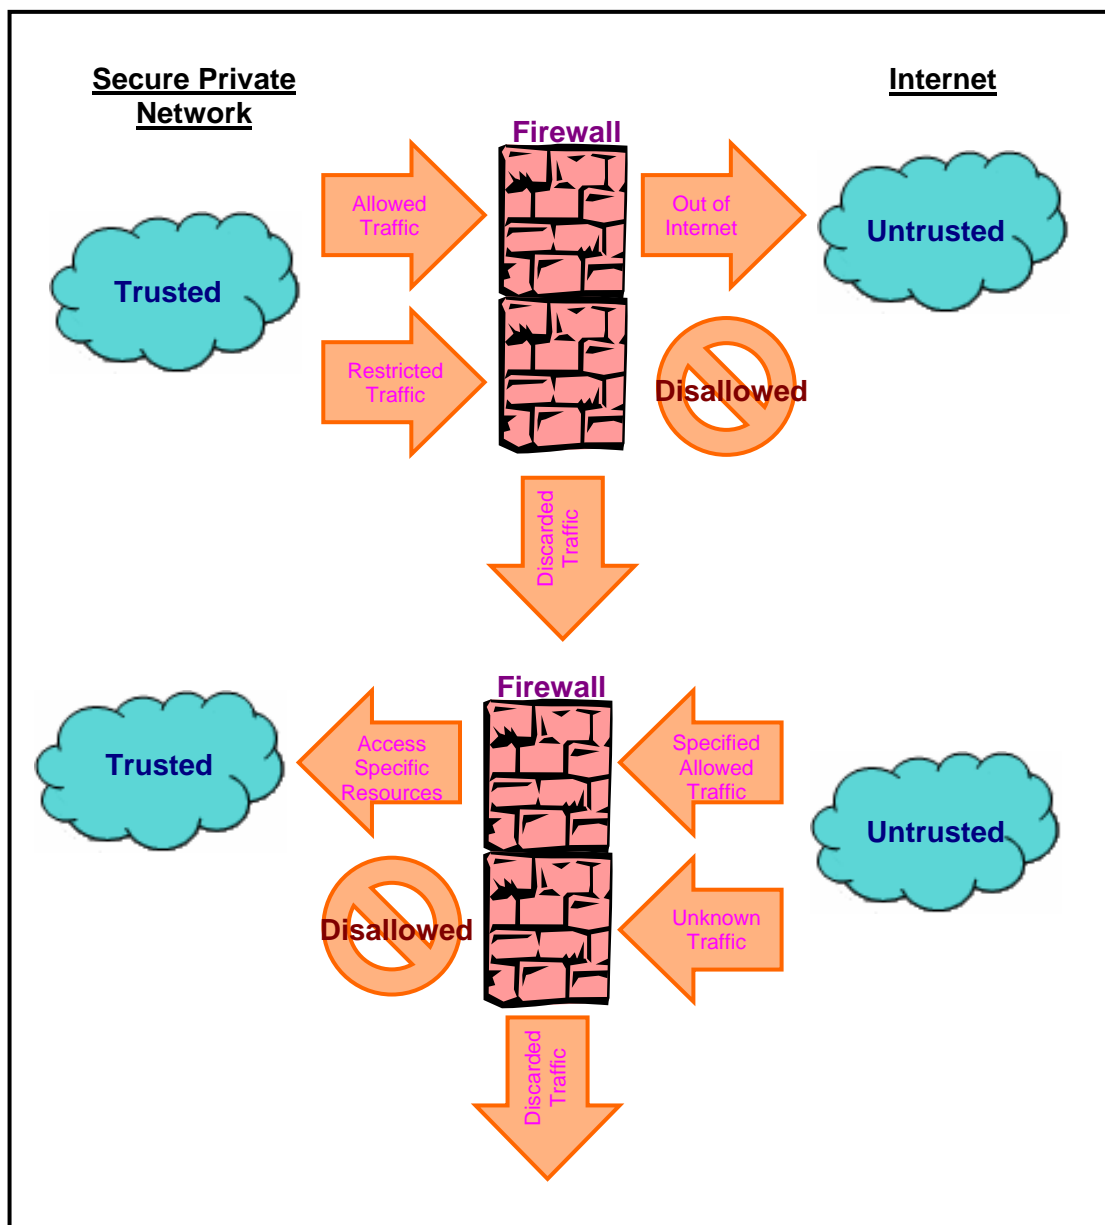


Figure 3. Packet Filtering

Packet filtering firewalls work at the Internet Protocol (IP) layer of the TCP/IP model or Network layer of the OSI model. They are usually part of a router. It receives packets from one network and forwards them to another network. This layer is concerned with routing packets to their destination and it is this layer where the firewall determines whether a packet is from a trusted source, even though it is unable to know what it contains or what other packets it is associated with. Depending on the packet and the criteria, the firewall can drop the packet, forward it or send a message to the originator. The rules will include source and destination IP address, source and destination port number and protocol used. Packet filter looks at one packet at a time and considers each of them in order to make a forwarding decision.

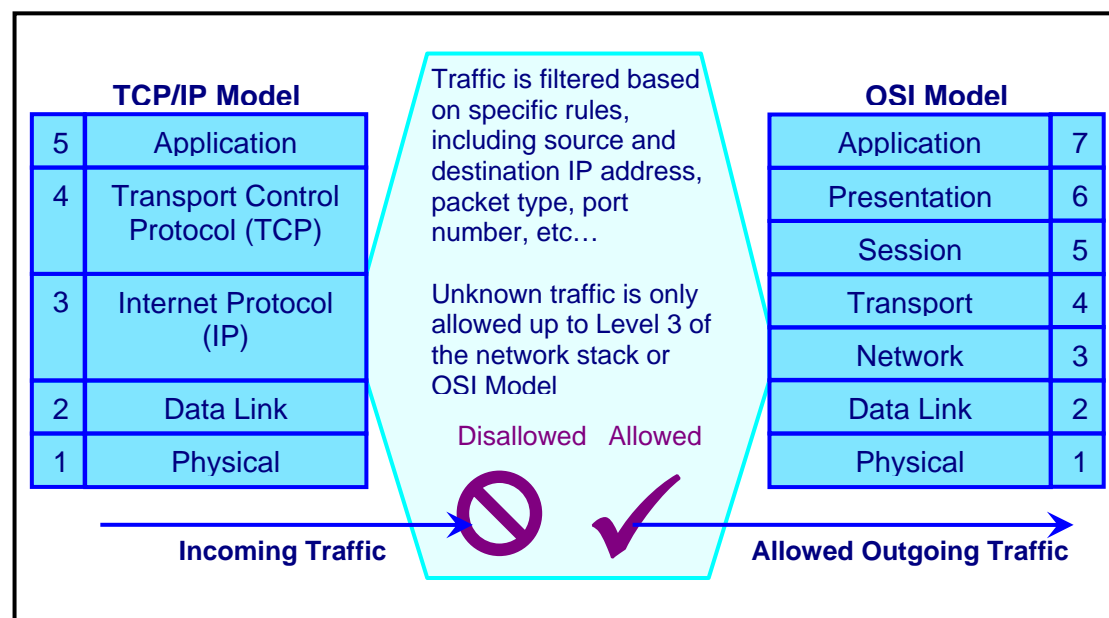


Figure 4. Packet Filtering in TCP/IP and OSI Models

The advantages of packet filtering firewalls are low cost and low impact on network performance for a very large number of user connections. The rules set are less complex and less likely to contain unintentional access routes. Most routers support packet filtering. Packet filtering is implemented for an initial degree of security at a low network layer. Since packet filters usually do not check data above Layer 3 of the OSI model, they can operate very quickly. Its fast, flexibility and simplicity allow packet filters to be deployed into nearly any enterprise network infrastructure. It is able to block denial-of-service and related attacks which makes it suitable to be placed at the outermost boundary of an untrusted network. Packet filter is also referred to as boundary router that can block certain attacks and filter unwanted

protocols, perform simple access control and then pass the traffic into other firewalls to check higher layers of the OSI model.

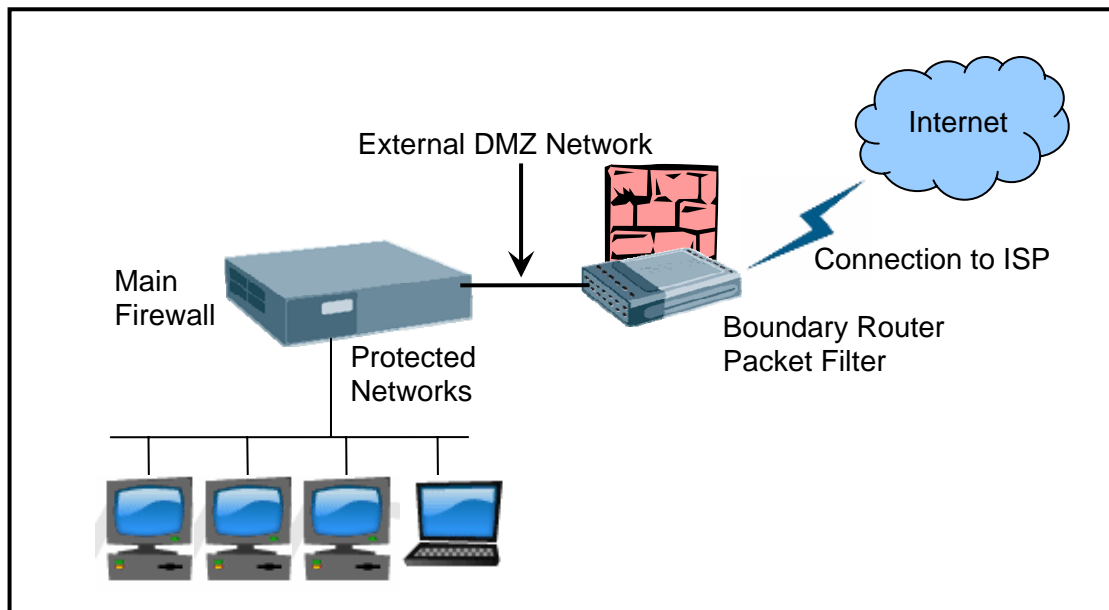


Figure 5. Packet Filter as a Boundary Router

This boundary router accepts packets from the untrusted network connection which is from another router of the ISP. The router will then perform access control according to the policy in place, like block SNMP, permit HTTP, etc. It then passes the packets to other more powerful firewalls for more access control and filtering operations at the higher layers of the OSI model. The internal less trusted network between the boundary router and an inner firewall is referred to as the external Demilitarized Zone (DMZ) network.

There are limitations to packet filtering. Packet filter firewalls do not examine upper-layer data. They cannot prevent attacks that employ application-specific vulnerabilities or functions. This firewall cannot block specific application commands. If a packet filter allows an application, all functions available within application will be permitted. The logging function in the packet filter firewall is limited. Packet filter logs normally contain the same information (source address, destination address, type of traffic) used to make access control decisions. Most packet filter firewalls do not support advanced user authentication schemes. This function is normally available in the upper layer of the OSI model.

Packet filter firewalls are generally vulnerable to attacks and exploits that take advantage of problems within the TCP/IP specification and protocol stack, such as

network address spoofing. Most packet filter firewalls cannot detect a network packet in which the address information has been altered. Intruders employ spoofing attacks to bypass the security controls implemented in a firewall.

Packet filter firewalls are suitable for environments where high speed is given priority over logging and user authentication.

For packet filter firewalls, whenever a rule that permits or denies the packet is found, one of the following actions is taken:

- **Accept** – the firewall passes the packet through the firewall as requested.
- **Deny** – the firewall drops the packet without passing it through the firewall. Once the packet is dropped, an error message is returned to the source system.
- **Discard** – the firewall not only drops the packet, it will not return an error message to the source system. This action is known as “black hole” so that a firewall is not revealed to the outsider.

Filtering can occur on outbound as well as inbound traffic. An organization can choose to restrict the types of traffic originating from within the organization. In practice, outbound filtering is often employed on IP addresses and application traffic. For example, blocking all users, internal and external, from connecting to certain systems such as the packet filter itself, backup servers and other sensitive systems.

Circuit-Level Gateway Firewalls

A circuit-level gateway is sometimes described as a second generation firewall. It relays TCP connections, but does not perform any other processing or filtering of protocol. It works at the session layer of the OSI model or the TCP layer of TCP/IP. It is a fast unrestricted passage through the firewall that allows and denies packets based on administrator’s predefined rules maintained in the TCP/IP kernel. It monitors TCP handshaking between packets to determine whether a requested session is legitimate.

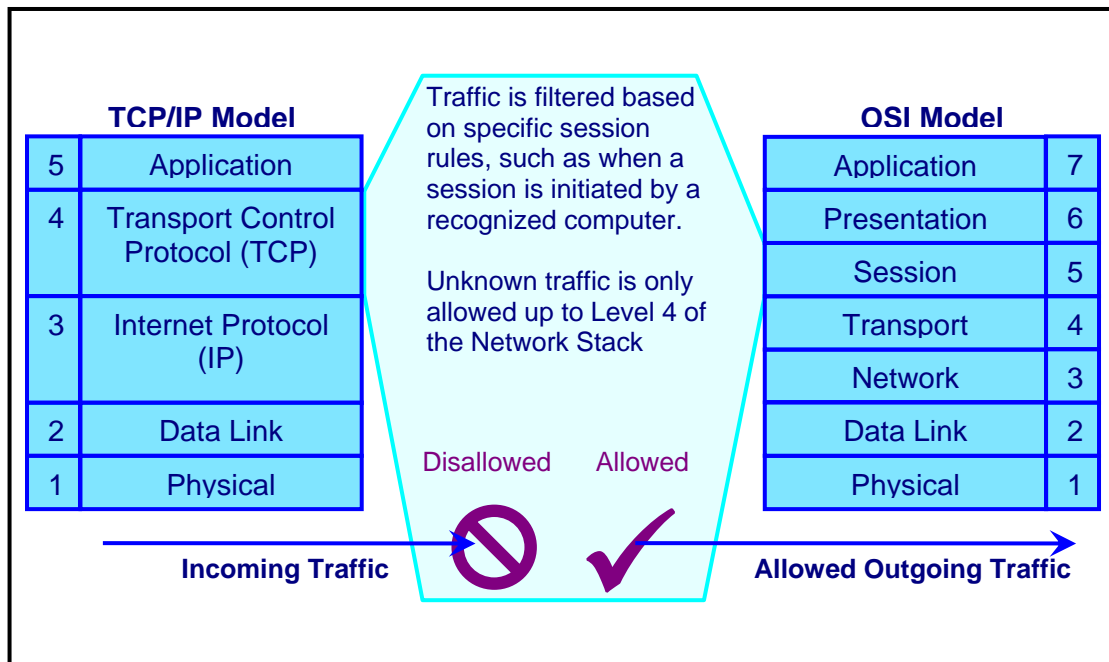


Figure 6. Circuit Level Firewall in TCP/IP and OSI Models

Information passed to the remote computer through a circuit level gateway appears to have originated from the gateway. It is basically used for TCP connections. It examines each connection setup to ensure that it follows a legitimate handshake for the transport layer protocol being used. It typically stores the following information:

- A unique session identifier (for tracking purposes)
- The state of the connection (handshaking, established or closing)
- Sequencing information
- Source IP address
- Destination IP address
- Physical network interface through which the packet arrives
- Physical network interface through which the packet leaves

The firewall checks whether the sending computer has permission to send to the destination and whether the receiving computer has permission to receive from the sender. If the connection is allowed, all associated packets are routed through the firewall with no further security checks.

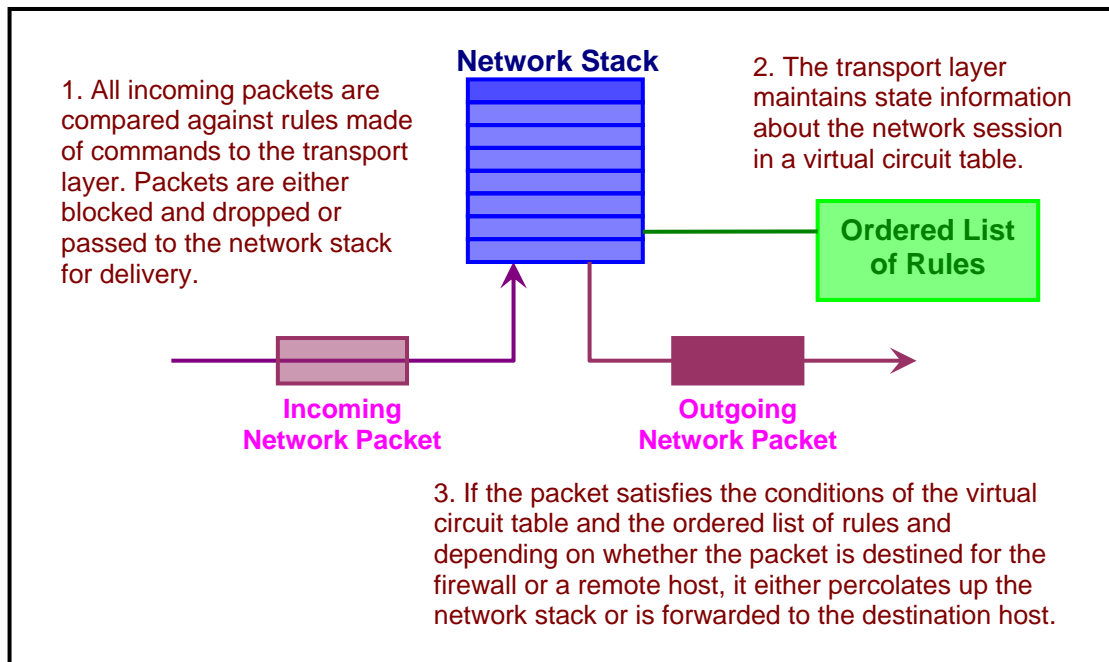


Figure 7. Circuit-Level Firewall Operation

This is useful for hiding information about protected networks by prohibiting connections between specific Internet sources and internal computers. Circuit level gateways are relatively inexpensive and have the advantage of hiding information about the private network they protect. But, they do not filter individual packets. It performs **Network Address Translation (NAT)** to shield internal IP addresses from external users.

Circuit-level firewalls are usually faster than the static packet-filter firewalls because they don't keep re-inspecting the communication streams. It is also generally faster than application layer firewalls because they perform fewer evaluations. The rules tend to be straightforward. Most applications will not be aware of such firewall that is monitoring the network communication. It can log TCP connections and analyze the logs easily.

However, circuit-level firewall does not inspect the complete conversation and does not keep watch on the conversation once it has been "approved" as valid, so communication session hi-jacking can still be an issue. This firewall does not validate communication occurring across the port. It cannot restrict the protocol subsets other than TCP and it cannot perform strict security checks on a higher-level protocol.

Network Address Translation

The Internet is expanding at an exponential rate. As the amount of information and resources increases, it has become a necessity for small businesses and homes to connect to the Internet. One known problem with connecting the IP networks to the Internet is that local hosts must have globally unique addresses in order to be identified by the backbone routers of the Internet. These Internet backbone routers will not route IP packets if the IP addresses do not process registered and unique incoming packet. End users connected to hosts with duplicate addresses will not be reached and cannot establish any application sessions.

Network Address Translation (NAT) solves this problem by reassigning IP addresses and port numbers. It is a method of connecting multiple computers to the Internet using one IP address simultaneously. This allows home users and small businesses to connect their network to the Internet cheaply and efficiently. It contains a pool of available ports so that one public IP address may be constantly reused. This means that only a single unique IP address is required to represent an entire group of computers.

NAT allows you to take advantage of the reserved address blocks. Typically, your internal network will be setup to use one or more of these network blocks. They are

10.0.0.0 / 8	(10.0.0.0 – 10.255.255.255)
172.16.0.0 /12	(172.16.0.0 – 172.31.255.255)
192.168.0.0 / 16	(192.168.0.0 – 192.168.255.255)

NAT Modes of Operation

NAT is like the receptionist in a large office. The customer calls the main number, the only number they know, of the office. When the customer tells the receptionist who he is looking for, the receptionist will check a lookup table with the name and extension number. The receptionist will then know when and how to forward the customer to the correct extension.

There are different forms of NAT operations are:

- Port Address Translation, PAT (many-to-one)

When a host on the private network sends an IP packet to the public network via the NAT, the NAT will keep track of the actual private address of that host, but substitutes the assigned public address into the IP address before it is sent into the public network.

When a reply comes back from the public network, the NAT restores the actual private address before sending the reply to the host.

Nothing originating from the public network can get through to the private network.

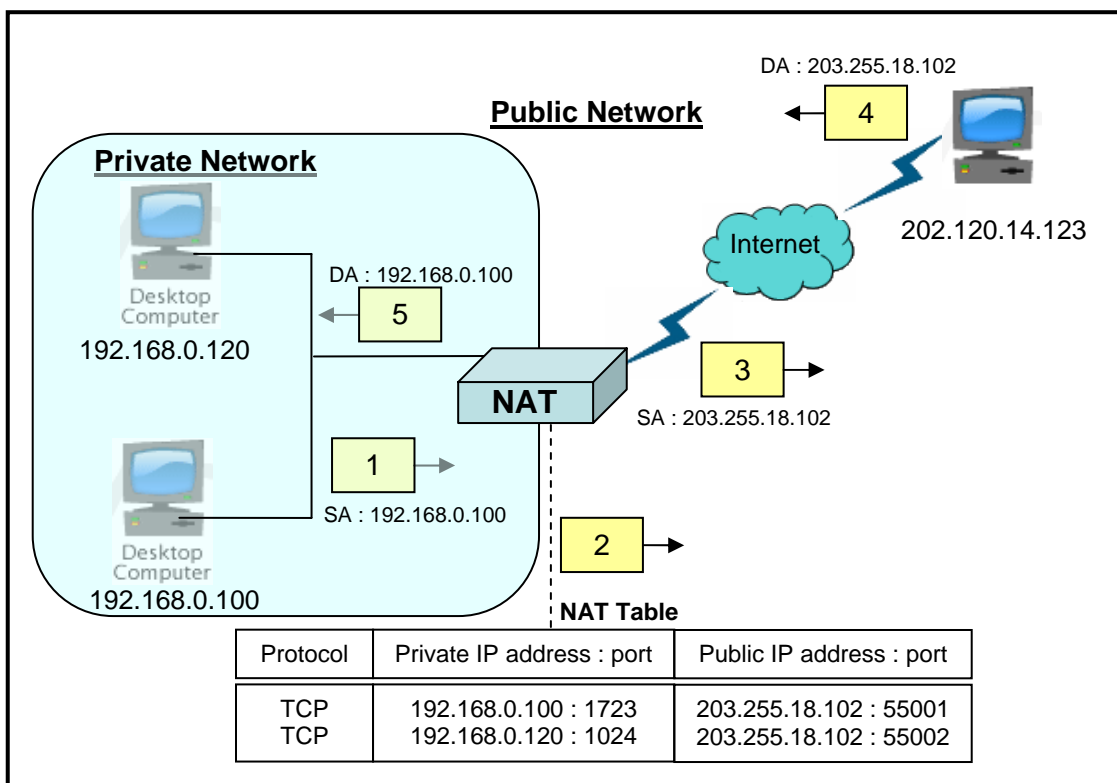


Figure 8. Port Address Translation

Operational Steps:

1. The Private Host 192.168.0.100 opens a connection to a server on the public network (202.120.14.123). This public server may be any UDP/TCP based server.
2. The first packet that the NAT receives from the private host (192.168.0.100) for forwarding to the NAT interface causes a check at the NAT table. All packets from the private network destined for the public network must then go through this NAT.

The NAT looks at the source IP address of this incoming packet and check if an entry exists for that specific IP address in the NAT table. If no entry is found in the table, a new entry will be added. Once done, NAT routines will do the port mapping and modify the source IP address.

In the dynamic mode, NAT will translate the private IP address to a legal public address.

3. The router replaces the private source IP address of 192.168.0.100 with the NAT-enabled public IP address (203.255.18.102) and forwards the packet to its destination, 202.120.14.123. More importantly, the source TCP port is also replaced from 1723 to 55001. This port translation will be used later to route the responding packet back to the correct originator (192.168.0.100).
 4. 202.120.14.123 receives the packet and replies to the private host 192.168.0.100 by using the NAT-enabled public address as its destination address 203.255.18.102.
 5. When the NAT router receives the reply destined to its publicly configured IP address (203.255.18.102), it performs a NAT table lookup. NAT takes the destination IP address of the incoming packet and checks its table to see if an entry exists. Once an entry is found, it uses the destination port of the incoming packet to remap the port and destination IP address. The NAT translates the destination IP address back to the corresponding private IP address of 192.168.0.100 based on the port number and forwards the packet to the private host 192.168.0.100.
 6. The private host 192.168.0.100 receives the packet and continues the conversation. The NAT performs Steps 2 through 5 for each packet.
- Static NAT mode (one-to-one)
- The NAT can provide support for servers in the private network that need to be “seen” from the public network.

Using this feature, a Web or FTP server located on the private network can use an illegal (private) IP address and still allow public clients access to these services.

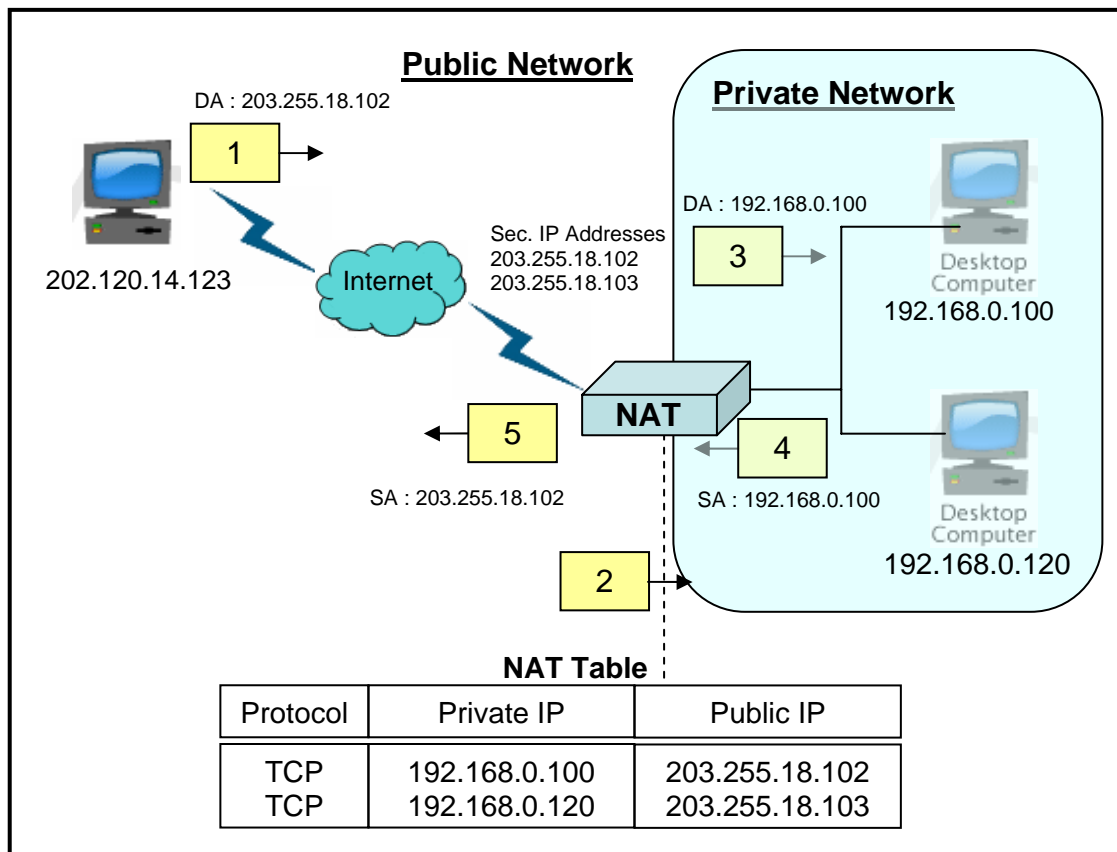


Figure 9. Static NAT

The diagram above shows how NAT translates virtual public IP addresses to private IP addresses. Note that no port translation takes place in static mode. The NAT should have two interfaces with IP addresses assigned to them. The private interface has IP address 192.168.0.1 and the public interface has another IP address 203.255.18.101 assigned. Two secondary IP addresses are also added to the public interface (203.255.18.102, 203.255.18.103) to use with NAT.

Static mode is required when a host on the private networks needs to be accessed by hosts on the public networks.

Operational Steps:

1. The user on host at the public network (202.120.14.123) opens a connection to virtual host at 203.255.18.102. This IP address is a secondary IP address bound to the public NAT-enabled interface and is statically mapped to the private host (192.168.0.100).

2. The NAT receives the connection request and checks the contents of its translation table to see if an entry for that destination address exists. If it does, the NAT routine determines that destination address (202.120.14.123) must be translated. It then sets up a translation of virtual public IP address (202.120.14.123) to a private IP address of the internal host needing to be accessed (192.168.0.100).
 3. The NAT replaces the virtual destination address (202.120.14.123) with the translated private host IP address (192.168.0.100) and forwards the packet to that node on the private network.
 4. The private host (192.168.0.100) receives the packet and responds. The destination IP address in the response being the public IP address (202.120.14.123) that initiated the original connection request.
 5. The NAT receives the packet, performs a NAT table lookup using the private IP address. The NAT then translates the source address (192.168.0.100) to the address of the virtual host (202.120.14.123) and forwards the packets.
- Dynamic NAT mode (one-to-one)
- It maps an unregistered IP address to a registered IP address from a group of registered IP address.

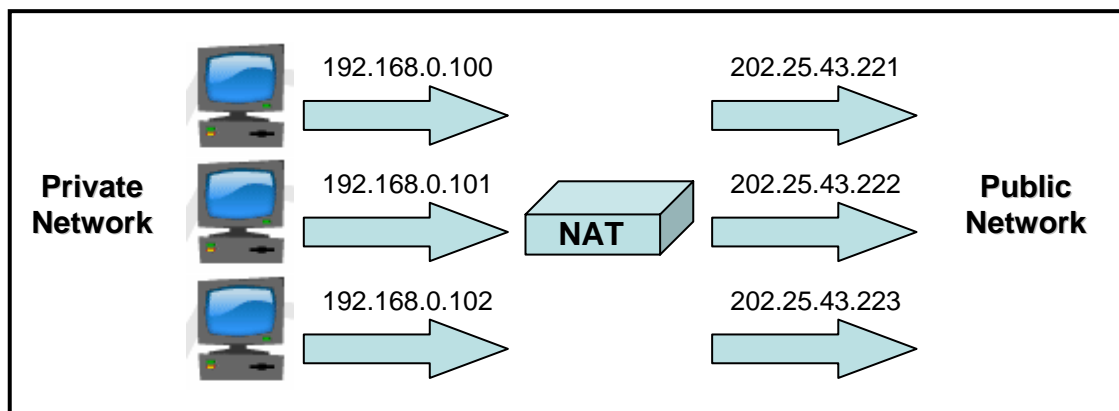


Figure 10. Dynamic NAT

Application Gateway

Known as application proxy or application-level proxy, an application gateway is an application program that runs on a firewall system between two networks. When a client program establishes a connection to a destination service, it connects to an application gateway or a proxy. The client then negotiates with the proxy server in order to communicate with the destination service. In effect, the proxy establishes the connection with the destination and acts on behalf of the client, hiding and protecting individual computers on the network behind the firewall. This creates two connections: one between the client and the proxy server and one between the proxy server and the destination. Once connected, the proxy makes all packet-forwarding decisions. Since all communication is conducted through the proxy server, computers behind the firewall are protected.

Summary for Firewall

The primary function of a firewall is to block “bad” traffic, while allowing “good” traffic to pass through. Its access control features are used to distinguish between the good and bad traffic. Firewall can be software that runs on computers or standalone hardware.

A firewall is usually installed between the computers and the Internet. It can allow web page request, file download, chat, etc. However, it can also make sure that the other people on the web cannot get access to the internal servers and computers for file sharing or printing.

Everyone who is connecting to the Internet should be protected by a firewall. There are many programs on the Internet that can scan huge ranges of IP address for vulnerabilities. These programs can be downloaded and run easily without much network knowledge. They are harmful to the computer or network. A firewall can keep you safe from such attacks.