Redes de Computadores (RCOMP)

Theoretical-Practical (TP) Lesson 08

2017/2018

- Static firewall.
- Cisco IOS standard and extended Access Control Lists.
- IP Named Access Control Lists.

Firewalls

When connected to the internet, we are exposed to all types of traffic from all over the world. This represents a serious security threat. Firewalls have the mission of analysing packets and block any traffic that might represent a threat.

Nevertheless, the golden rule on setting up a firewall is: **let desired (required)** traffic pass through and <u>block all the rest</u>.

Each end-node (e.g. workstation or server) should run a local firewall, taking special attention to incoming traffic, however, by running the firewall on intermediate nodes like routers, many nodes can be protect at once.

Users' workstations are often out of control from network administrators, nevertheless, they can be protected if traffic is blocked before reaching them. We must, however, remember this will not block traffic between workstations.

Servers can also be protected by firewalls running on routers, in this case we should place all servers in an isolated network apart from users' workstations. Because we assume all servers are trustable, we can even get rid of local firewalls on each server, due to this, that isolated network is usually know as Demilitarized zone (DMZ).

Internet, intranet and DMZ

Under the security point of view at least three network zones can be pointed out:

- Internet from where most threats come from
- Intranet local networks for internal users
- Demilitarized zone (DMZ) trusted local network for servers only

Because they present different levels of security risk, these zones are separated by firewalls.



Static firewalls

There are several different types of firewall, they all have the same objective: let pass only desired traffic. Static firewalls, also called stateless firewalls or packet filters work by analysing each individual packet's layer three and layer four headers. This means the analysis will be on IP source and destination addresses (layer three), and also, ICMP message types, UDP port numbers, and TCP port numbers (layer four). They are called static and stateless because they don't change behaviour along time.

State full or dynamic firewalls, on the other hand, track application layer network transactions, while an authorised transaction is ongoing, they allow the required traffic to pass through. They are called dynamic because rules change depending on what is happening.

Some types of attacks can be prevented only by dynamic firewalls, this is the case of DoS (denial-of-service) attacks. A DoS attack is performed by flooding with requests a service that is supposed to be accessible to everyone.

A static firewall can't handle DoS attacks because it must allow access to that service for normal use. A dynamic firewall, however, can detect the attack and change the rules to block it.

Nevertheless, static firewalls can block many other attack types.

Access Control List (ACL)

In Cisco IOS, a static firewall can be deployed in a router by using access control lists. An ACL is a sequence of rules, each rule contains a matching criterion and an action to perform to every packet matching that criterion. Typical actions are **permit** or **deny**.

Each packet's characteristics are confronted sequentially with each rule in the ACL until a match is found, if so, the corresponding action is performed on the packet (**permit** or **deny**).

With Cisco IOS access control lists, if a packet doesn't match any rule in the ACL, then **the packet is ultimately denied**.

There are several types of ACL in Cisco IOS, to start with, there are numbered access lists (identified by numbers) and named access lists (identified by names) but they are very similar. We will start by numbered ACLs.

Each numbered ACL has a unique number that identifies it (within the router), the number also settles the ACL type.

1 up to 99 – for **standard** access lists

100 up to 199 – for **extended** access lists

Standard access lists

For numbered access lists, one ACL is known to be of the standard type if its identifier is between one and 99. These access lists support a single criterion for packet matching: **the packet's source address**.

The **access-list** command is used to create a numbered access list:

access-list number action source-address

number is the ACL identifier (1 to 99)

action is either permit or deny

source-address is an address matching criterion for packet's source address

The access-list command adds (appends) one rule to the ACL, a numbered ACL can't be edited, commands must be typed in the exact order we want rules to be in the ACL.

If a wrong rule is added (entered), there is no other option but to remove the entire ACL (by using the **no access-list number** command), and then type all the commands again.

ACL addresses matching criterions

Matching criterions for IPv4 addresses can be specified in one of three ways:

host IPv4-ADDRESS- matches the exact provided IPv4-ADDRESSany- matches any IPv4 address

And the general form:

IPv4-ADDRESS IPv4-wildcard - matches the provided IPv4-ADDRESS, however, **it ignores bits that have value one on the IPv4-wildcard**.

Both IPv4-ADDRESS and IPv4-wildcard are dot-decimal representations of 32 bits. The IPv4-wildcard acts as a bits mask, for each bit with value one on the IPv4-wildcard, the same bit at IPv4-ADDRESS is ignored for matching criterion.

Thus:

host IPv4-ADDRESS	is equivalent to	IPv4-ADDRESS 0.0.0.0
any	is equivalent to	0.0.0.0 255.255.255.255

IPv4 wildcards and network masks

To create an IPv4 addresses matching criterion encompassing an entire IPv4 network we can use the IPv4 network address as IPv4-ADDRESS and the inverted network mask (reversed bits) as IPv4-wildcard.

Examples:

To match all addresses of network 192.168.50.0/24 we can use 192.168.50.0 0.0.0.255

(24 bits prefix-length means the network mask is 255.255.255.0)

To match all addresses of network **10.0.0/8** we can use **10.0.0.0 0.255.255.255**

(8 bits prefix-length means the network mask is 255.0.0.0)

To match all addresses of network **170.2.64.0/23** we can use **170.2.64.0 0.0.1.255** (23 bits prefix-length means the network mask is 255.255.254.0)

As it comes clear in the last example, to fully understand which addresses are matched, a binary analysis may be required.

IPv4 wildcards

As seen, wildcards can be used as network masks equivalents, yet, wildcards are much more powerful because unlike network masks bits one and zero can be alternated.

Examples:

The criterion 192.4.50.0 0.1.0.255 matches networks 192.4.50.0/24 and 192.5.50.0/24

The criterion 170.5.7.10 0.0.0.64 matches addresses 170.5.7.10 and 170.5.7.74

The criterion **10.1.0.5 0.32.0.64** matches **10.1.0.5**, **10.33.0.5**, **10.1.0.69**, and **10.33.0.69**

Looking in more detail into this last example:

- because bit with value 32 on the second octet is one, it's ignored for matching purpose, therefore, addresses with value 1 on that octet match the address, but if the value is 33 they also match.
- because bit with value 64 on the fourth octet is one, it is ignored for matching purpose, likewise, addresses with value 5 on that octet match the address, but if the value is 69 they also match.

A rule may by used to check our solution: **the number of matched addresses is always the power of two of the number of one bits on the wildcard**. We can also conclude no wildcard can match an odd number of addresses.

Enforcing access lists

Whatever type of access list we are talking about, just creating it has no immediate effect, it must be applied to a specific network interface to filter the incoming or outgoing traffic on that interface.

Until we apply an ACL to an interface, all traffic is allowed on that interface, once we apply, only traffic that matches a **permit** rule on the ACL will be allowed.

An existing ACL is associated to an interface by using the **ip access-group** command, within the interface's configuration:

ip access-group ACCESS-LIST-ID in out

Where ACCESS-LIST-ID is the ACL identifier (either a number or a name).

The ACL is either applied to incoming traffic (in) or outgoing traffic (out). Incoming traffic stands for traffic being received from the network on this interface. Outgoing traffic stands for traffic being emitted to the network from this interface. Each traffic direction is independent of the other, enforcing an ACL on incoming traffic does not affect the outgoing traffic of the interface. Two access lists cannot be applied to the same interface in the same direction, the second application removes the first.

Anti spoofing with standard access lists

Standard access lists only analyse packets' source address, nevertheless, that's enough to prevent IP spoofing attacks.

IP spoofing stands for forging the source address of packets, they may serve several illicit purposes, including:

- Misleading static firewalls that allow access to some services only if traffic comes from allowed source addresses.
- Misleading dynamic firewalls on DoS attacks by making them think the requests flooding the network are coming from different sources.

The network administrator is obliged to block two types of spoofing:

External spoofing: this applies to incoming traffic to a local network, for the sake of local network's security, any incoming packet with a source address belonging to the local network is spoofing and should be blocked.

Internal spoofing: this applies to outgoing traffic from a local network, for the sake of other networks' security, any outgoing packet with a source address that does not belong to the local network is spoofing and should be blocked.

Anti spoofing - example

Imagine the following scenario.

The local network is 194.20.30.0/24

We are to block external spoofing coming from the internet and internal spoofing coming from the local network.



The first issue is: in which interfaces the access lists should be enforced. We could block external spoofing on Fa0/1 incoming traffic or Fa0/0 outgoing traffic. We could block internal spoofing in Fa0/0 incoming traffic or Fa0/1 outgoing traffic.

The general rule is: if there's no specific advantage in doing otherwise, we should always enforce access lists as close to the traffic source as possible, so usually on incoming traffic.

Following this principle, in this case, we will block external spoofing on FaO/1 incoming traffic and block internal spoofing in FaO/0 incoming traffic.

Anti spoofing - example

So we will define one access list for each spoofing type:

```
no access-list 10
access-list 10 deny 194.20.30.0 0.0.0.255
access-list 10 permit any
no access-list 20
access-list 20 permit 194.20.30.0 0.0.0.255
interface Fa0/0
ip access-group 20 in
interface Fa0/1
ip access-group 10 in
```



Access list 10 says any traffic with source address belonging to the local network should be blocked and all others should pass.

Access list 20 says only traffic with source address belonging to the local network is allowed.

Access list 20 is applied to Fa0/0 incoming traffic (block internal spoofing) and access list 10 is applied to Fa0/1 incoming traffic (block external spoofing).

Remember, at the end of every access list there's an invisible, although implicit, rule **deny any**.

Before starting the definition of an ACL with **access-list** commands, is wise to make sure it's empty by issuing the command **no access-list**.

Extended access lists

Numbered extended access lists work the same way as numbered standard access lists, however, they are identified by a number between 100 and 199.

While the packet's source IP address is the only matching criterion on standard access lists, extended access lists also <u>require</u> a protocol criterion and a destination address criterion:

access-list number action protocol source-address destination-address

- number is the ACL identifier (100 to 199)
- action is either permit or deny
- **protocol** is layer three or layer four protocol name identifier
- source-address is a matching criterion for packets' source address
- destination-address is a matching criterion for packets' destination address
- The rule only matches a packet if all rule criterions are meet at the same time.
- The protocol name can be, among others: **ip**, **icmp**, **udp** or **tcp**.
- Because all these protocols are transported by IP, specifying **ip** as a matching protocol will match any packet's protocol.

ICMP protocol extended access lists

Depending on the matching criterion for the protocol, additional matching criterions are available. For the ICMP protocol:

access-list number action icmp source-address destination-address [icmp-type [icmp-code] | icmp-message]

We can set the criterions through either the decimal values of ICMP message type and code or through a match criterion for the name of the ICMP message. Some frequently used ICMP message names (**icmp-message**) are: **echo** to match ICMP echo requests and **echo-reply** to match ICMP echo replies.

Specifying **echo** is equivalent to specify **80**, because the ICMP echo request is a type 8 message with code 0.

Specifying **echo-reply** is equivalent to specify **0**, because the ICMP echo reply is a type 0 message with code 0.

If no ICMP specific criterion is used, then the rule matches any ICMP message.

UDP protocol extended access lists

If the matching protocol is UDP we can also define matching criterions based on source and destination UDP port numbers:

access-list number action udp source-address [operator port] destination-address [operator port]

The criterion after source-address is for matching the source port number and the one after destination-address is for matching the destination port number. You may specify only one criterion, both, or none. If no port number is specified it will match any port number.

The **port** argument can either be a port decimal number or a name representing the standard service for that port number.

The **operator** can be one of:

- It less than
- **gt** greater than
- eq equal to
- neq not equal to
- range within the inclusive range

(in this last case two space separated port numbers or service names are to be used)

TCP protocol extended access lists

If the matching protocol is TCP we can, as well, define matching criterions based on source and destination port numbers (now for TCP port numbers):

access-list number action tcp source-address [operator port] destination-address [operator port] [established]

Criteria specification for ports is the same as before with UDP. Additionally, one new criterion can be used specifically, and only, for TCP.

The **established** criterion matches TCP segments with either the ACK bit or the RST bit on. This excludes the initial SYN request segment sent with the purpose of establishing a new TCP connection, therefore, **it will match only TCP traffic regarding already established TCP connections**.

Although we can specify port numbers matching criterions for both source ports and destination ports, they are most useful on destination ports. Servers provide network services in standard fixed number ports, if we want to allow clients access to specific servers we can use the destination port criterion. Clients on the other hand use dynamically assigned port numbers.

Port numbers can be replaced by the service name, for instance the name **www** represent number 80 and the name **domain** represents number 53.

Extended access lists application example

We now have this scenario:

We still want to block spoofing, but now there's an HTTP server on a DMZ. The DMZ is trusted, so no spoofing will be coming from there.

The DMZ must be protected, access to the HTTP server is the only



allowed.

```
no access-list 100
access-list 100 permit tcp 194.10.8.0 0.0.0.255 host 194.10.9.7 eq 80
access-list 100 deny ip any 194.10.9.0 0.0.0.255
access-list 100 permit ip 194.10.8.0 0.0.0.255 any
no access-list 120
access-list 120 deny ip 194.10.8.0 0.0.1.255 any
access-list 120 permit tcp any host 194.10.9.7 eq 80
access-list 120 deny ip any 194.10.9.0 0.0.0.255
access-list 120 permit ip any any
interface Fa0/0
ip access-group 100 in
interface Fa1/0
ip access-group 120 in
```

Named access lists

Named access lists have one major advantage over numbered access lists, they can be edited. A named access list can be created and later edited with the **ip access-list** command:

ip access-list standard | extended ACCESS-LIST-NAME

This will make CLI enter in **named access list configuration mode**, either (configstd-nacl) for a standard access list or (config-ext-nacl) for an extended access list.

In **named access list configuration mode**, we add rules by using the same syntax as with numbered access list except the access-list command is omitted:

[SEQUENCE-NUMBER] permit | deny MATCH-CRITERIONS-SPECIFICATION

The MATCH-CRITERIONS-SPECIFICATION is exactly the same as for numbered access-list, depending on being standard or extended.

Every rule in the named access list has a sequence number, if none is specified when adding the rule it will be added (appended) to the end. The first rule to be added will have sequence number 10 and next rules will be added with a 10 increment. Sequence numbers can be from 1 up to 2147483647.

If the SEQUENCE-NUMBER is specified the rule will be placed in that position, if the position is occupied the command fails.

Named access lists

We can remove a rule by using the command:

no SEQUENCE-NUMBER

On a real router (not in Packet Trace) you can also redefine the sequence numbers of rules in a named ACL:

ip access-list resequence ACCESS-LIST-NAME starting-sequence-number increment-value

We can see the present sequence numbers of a named ACL by exiting the configuration mode (end) and running the command:

show ip access-lists ACCESS-LIST-NAME

Named access lists example



Additional notes regarding anti spoofing

We must be aware there are some cases where network administrators are not able to avoid IP spoofing.

Internet incoming traffic – except for the local networks addresses, we can't tell if one source address is spoofing or not, so all we can do is blocking packets with local networks' source addresses. This is why it's so important for each network to play it's role and block internal spoofing.

Network traffic within local networks – traffic between nodes in the same IPv4 network is not handled by any router and escapes the administrator's control. To be able to do something about this we would have to use layer three switches. A layer three switch operates as a layer two switch but understands IPv4 and access control lists can be enforced to block any local IP spoofing.

Local transit networks – transit networks are used for traffic passing, this means some traffic travelling through them is neither originated by them nor destined to them. This presents an issue for anti-spoofing rules because we can no longer enforce that every incoming packet from that network must have a source address belonging to the network.



We can see there are two spoofing issues about transit Network B:

- It can send to the internet packets with forged source addresses belonging to Network A
- It can send to Network A packets with any forged source address, except for addresses belonging to Network A.