


Prepared by Dr. Mohammed Anbar

Network Security (CYB 630)
Network Infrastructure Security (part two)

Lecture 2





Some of the contents in this presentation is adopted from the resources provided by the authors of the following book:

Wong, Angus, and Alan Yeung. *Network infrastructure security*. Springer Science & Business Media, 2009.

Chapter


01

Learning objectives

- Explain Network Infrastructure and its security,
- Differentiate between Information Security and Network Infrastructure security,
- Explain Network Infrastructure components ,
- Explain switch security,
- Explain Dynamic Host Configuration Protocol (DHCP) security ,
- Explain Domain Name System (DNS) security .

- Layer 2 of the seven-layer OSI model is the data link layer which is on top of the physical layer (layer 1).
- Since the physical layer is only concerned with the transmission of a raw bit stream over a physical medium, the data link layer attempts to provide reliable data transfer across the physical link.
- The services provided by the link layer include flow control, acknowledgment, error recovery, and maintenance (activation and deactivation) of the link.
- Switch is the most used layer 2 device in today's local area networks.

A network switch connects devices within a network (often a local area network, or LAN*) and forwards data packets to and from those devices. Unlike a router, a switch only sends data to the single device it is intended for (which may be another switch, a router, or a user's computer), not to networks of multiple devices.



- Network switches can operate at either OSI layer 2 (the data link layer) or layer 3 (the network layer).
- Layer 2 switches forward data based on the destination MAC address while layer 3 switches forward data based on the destination IP address. Some switches can do both.
- Layer 2 switches most often connect to the devices in their networks using Ethernet cables. Ethernet cables are physical cables that plug into devices via Ethernet ports.

When a switch receives a frame:

index switch table using MAC dest address

if entry found for the destination

then{

if dest on a segment from which frame arrived

then drop the frame

else forward the frame on the interface indicated

}

else flood

forward on all but the interface
on which the frame arrived



Port Table

Destination Addresses	Port
EE	1
AA	2
BA	3
EA	4
AC	5
AB	6

Two terms are associated with frames entering or leaving an interface:

- **Ingress** – entering the interface
- **Egress** – exiting the interface



Port Table

Destination Addresses	Port
EE	1
AA	2
BA	3
EA	4
AC	5
AB	6

A switch forwards based on the ingress interface and the destination **MAC address**. A switch uses its MAC address table to make forwarding decisions.

Note: A switch will never allow traffic to be forwarded out the interface it received the traffic.

A MAC (Media Access Control) address, sometimes referred to as a hardware or physical address, is **a unique, 12-character alphanumeric attribute that is used to identify individual electronic devices on a network**. An example of a MAC address is: 00-B0-D0-63-C2-26

- A switch will use the destination MAC address to determine the egress interface.
- Before a switch can make this decision, it must learn what interface the destination is located.
- A switch builds a MAC address table, also known as a Content Addressable Memory (CAM) table, by recording the source MAC address into the table along with the port it was received.

Frame Forwarding

The Switch Learn and Forward Method (cont.)

The switch uses a two-step process:

Step 1. Learn – Examines Source Address

- Adds the source MAC if not in table.
- Resets the time out setting back to 5 minutes if source is in the table.

Step 2. Forward – Examines Destination Address

- If the destination MAC is in the MAC address table, it is forwarded out the specified port.
- If a destination MAC is not in the table, it is flooded out all interfaces except the one it was received.

A switch will use one of two methods to make forwarding decisions after it receives a frame:

Store-and-forward switching

Cut-through switching

Ethernet II Frame Structure

7 Bytes	1 Byte	6 Bytes	6 Bytes	2 Bytes	46 – 1500 Bytes	4 Bytes
Preamble	SFD	Destination Address	Source Address	Type	Data Payload	Frame Check Sequence (FCS)

- **Preamble** – informs the receiving system that a frame is starting and enables synchronization.
- **SFD (Start Frame Delimiter)** – signifies that the Destination MAC Address field begins with the next byte.
- **Destination MAC** – identifies the receiving system.
- **Source MAC** – identifies the sending system.
- **Type** – defines the type of protocol inside the frame, for example, IPv4 or IPv6.
- **Data and Pad** – contains the payload data. Padding data is added to meet the minimum length requirement for this field (46 bytes).
- **FCS (Frame Check Sequence)** – contains a 32-bit Cyclic Redundancy Check (CRC) which allows the detection of corrupted data.

Frame Forwarding Switch Forwarding Methods

- **Store-and-forward switching** - Receives the entire frame and ensures the frame is valid. Store-and-forward switching is Cisco's preferred switching method.
- **Cut-through switching** – Forwards the frame immediately after determining the destination MAC address of an incoming frame and the egress port.



- The security of the Ethernet switch is essential since it is the most commonly used layer 2 device in current LANs. Unfortunately, the security of switches is often disregarded.
- Network administrators rarely monitor or investigate the operation of layer 2 infrastructures until there is a connectivity issue since they normally consider layer 2 networks can be trusted.
- Security measures typically concentrate on layer 3. For instance, to detect any suspicious activity, intrusion detection systems often inspect layers 3 and 4 communications, particularly TCP/IP packets.


- Switch attacks are difficult to detect. One of the reasons is that detecting such attacks requires a thorough understanding of how data link protocols work.
- For example, to exploit the spanning tree protocol (STP), one must first understand the roles of a switch port (can be a *root*, *designated*, or *blocked* port) and their possible states (such as *forwarding*, *listening*, and *learning*) in an STP network.

How Switches can be Attacked?

MAC Flooding

As can be seen, a new entry for the new address will be added to the CAM table whenever a frame arrives with a source MAC address that the switch has never encountered. The table's size is limited. The switch, unlike a hub, will broadcast all the frames to all outbound ports when the table is overloaded rather than relying on it.





Limiting the amount of MAC addresses that can be connected to a switch port is a frequent strategy for preventing the MAC flooding attack. High-end switches have a feature called port security that does this. The administrator can determine how many MAC addresses can be linked to a switch port or which MAC addresses are allowed to access a particular port.

How Switches can be Attacked?

ARP spoofing

When a host wants to send an IP packet to another host on an Ethernet/IP network, simply knowing the destination IP is insufficient. To build and send Ethernet packets into the network, it also needs to know the destination host's MAC address. Therefore, a procedure is needed to determine a host's MAC address from its IP. For this, the Address Resolution Protocol (ARP) is employed.

1. When Host-A wants to interact with Host-B, in a typical ARP operation, Host-A broadcasts an ARP Request packet asking Host-B, "If you have IP-B, please give your MAC to me."
2. The only device that would reply to Host-A directly with an ARP Reply packet in unicast declaring, "I am Host-B, and my MAC is MAC-B," is the device with the given IP address, IP-B.

- Operating systems typically cache resolved addresses for a brief period to minimize the number of ARP requests.
- A host will examine the ARP cache to see if the required IP-MAC mapping is present before sending an ARP request.
- If the mapping is present, it won't need to send another ARP request to obtain the corresponding MAC address from the cache.
- For viewing, adding, or removing items from the ARP cache, Windows has the ARP command which is `arp -a`

```
Interface: 192.168.56.1 --- 0xc
  Internet Address      Physical Address      Type
  224.0.0.22           01-00-5e-00-00-16    static


Interface: 192.168.17.1 --- 0x15
  Internet Address      Physical Address      Type
  224.0.0.22           01-00-5e-00-00-16    static

Interface: 10.207.161.246 --- 0x16
  Internet Address      Physical Address      Type
  10.207.160.1          c0-62-6b-e2-26-40    dynamic
  10.207.160.108        3c-2a-f4-ad-cc-58    dynamic
  10.207.160.116        00-0c-29-d3-22-e4    dynamic
  10.207.161.129        3c-8d-20-69-b0-fa    dynamic
  172.31.108.101        00-ae-6b-64-6d-a8    dynamic
  224.0.0.22           01-00-5e-00-00-16    static
```


The result of execution of `arp -a` command

Since ARP does not require authentication, whenever a computer receives an ARP reply, it will update its cache regardless of whether it has sent out an ARP request. Therefore, one can easily poison other hosts by sending out an ARP Reply packet with wrong information.

1. Before the attack, the ARP caches of Host-A and Host-B contain the correct mappings



2. To poison the ARP cache of Host-A, Host-X sends Host-A an ARP Reply packet with the wrong mapping: IP-B-to-MAC-X.
3. Similarly, Host-X sends to Host-B an ARP Reply with the wrong mapping: IP-A-to-MAC-X.
4. Since ARP does not require authentication, Host-A and Host-B believe that the mappings are valid and update their ARP tables as usual. Now, their ARP caches are poisoned.




Source: [1]

3. As a result, the traffic between Host-A and Host B will be redirected to flow through Host-X first, instead of directly to each other. In this case, Host-x can intercept the traffic between them, even though they are connected using a switch.


The attacker can perform the following attack with its cache poisoning.

- *Man-in-the-middle attack:* If Host-X re-routes packets to the correct destination in both directions, Host-A and Host-B will not be aware that the traffic between them is being monitored (or modified) by Host-X.
- *DoS attack:* If Host-X chooses not to re-route packets, it will launch a Denial of Service (DoS) attack because Host-A and Host-B could not communicate with each other anymore. However, since no traffic happens to Host-A and Host-B, the corresponding entries in the ARP caches will be timeout. To keep denying of the service, Host-X must continue poisoning their ARP caches (i.e., repeat the above process) regularly.

A man-in-the-middle (MITM) attack is a general term for when an attacker positions himself in a conversation between a user and an application—either to eavesdrop or to impersonate one of the parties, making it appear as if a normal exchange of information is underway. The goal of an attack is to steal personal information, such as login credentials, account details, and credit card numbers



A denial-of-service (DoS) attack floods a server with traffic, making a website or resource unavailable. A distributed denial-of-service (DDoS) attack is a DoS attack that uses multiple computers or machines to flood a targeted resource. Both types of attacks overload a server or web application with the goal of interrupting services.



There are other differences, however, involving either their nature or detection, including:

1. Ease of detection/mitigation:

Since a DoS comes from a single location, it is easier to detect its origin and sever the connection. On the other hand, a DDoS attack comes from multiple remote locations, disguising its origin.

2. Speed of attack:

Because a DDoS attack comes from multiple locations, it can be deployed much faster than a DoS attack that originates from a single location. The increased speed of attack makes detecting it more difficult, meaning increased damage or even a catastrophic outcome.

3. Traffic volume:

A DDoS attack uses multiple remote machines (zombies or bots), which enables it to deliver significantly more traffic from various places at once and quickly overload a server while evading detection.

4. Manner of execution:

Multiple hosts that have been infected with malware (bots) are coordinated by a DDoS attack to form a botnet that is controlled by a command-and-control (C&C) server. A DoS attack, on the other hand, usually employs a script or tool to execute the attack from a single machine.

5. Tracing of source(s):

The use of a botnet in a DDoS attack means that tracing the actual origin is much more complicated than tracing the origin of a DoS attack.

➤ *Static ARP Entries:*

- To prevent spoofing, the ARP cache can store static IP-MAC mappings. As the entries are non-changing, any spoofed ARP replies will be ignored. This concept works; however, requiring the ARP cache to always keep up-to-date IP-MAC mappings of every computer on the network makes this solution impractical for most networks.

➤ *Detection:*

- To detect possible ARP spoofing-based attack, the IP-MAC mappings for the machines on the network can be monitored. There are programs, such as ARPWatch (<http://www.securityfocus.com/tools/142>), to do that. When any suspicious changes of the mappings occur, the program will notify the administrator to investigate.


➤ *No Cache Update:*

- A simple and practical solution is to only accept ARP replies and update the entries in the cache when they are expired (after the timeout period). In this case, in order to poison the ARP cache, the attacker must send an ARP reply packet faster than the legitimate host do, making the spoofing process difficult.

- Dynamic Host Configuration Protocol (DHCP) is used to automatically configure client machines with a dynamically assigned IP address and other network configuration parameters, such as the default gateway and DNS server addresses, during their boot time. It eliminates the need for network administrators to keep track of individual client IP addresses.
- Network administrators used to statically assign an IP address to each client machine in the past (before DHCP), and users had to manually set up their machines with the provided IP address. This technique makes incorrect configuration common and necessitates educated users.


The proper functioning of DHCP is essential. Users may be configured with false network information, such as the incorrect name server address if it is exploited, posing a security risk.

1. The network administrator only needs to set up the DHCP server with the required network configuration data and a pool of IP addresses.
2. If DHCP is enabled, a client computer will automatically request the DHCP server for an IP address when it boots up.
3. The server will then select an available address from the pool and response the client with the selected address and other network configuration information.



Source: [1]

1. The attacker can carry out this attack by repeatedly sending DHCP requests to the DHCP server with different (forged) MAC addresses.
2. Since the server treats each request with a new MAC as a new network client and assigns it an IP address, the address pool will eventually be fully allocated after the attacker sends a significant number of requests.
3. Any new valid client won't be able to get an IP address as a result because there are no more addresses accessible. It is a kind of denial of service attack -- denial of clients from having DHCP service.





There is no DHCP server authentication and DHCP clients do not know the DHCP server address when they are booted up. Therefore, an attacker can set up its own DHCP server pretending to be the legitimate one and return the client with a fake default gateway address. It can allow the attacker to intercept all network traffic.

1. The attacker first grabs an IP address from the legitimate DHCP server for future use. It runs its rogue DHCP server.
2. When a DHCP client broadcast a DHCP DISCOVERY packet, both the legitimate and the rogue servers send a DHCPOFFER to the client.

3. The client accepts the response from whichever DHCP responds first. If the client uses the one from the legitimate server, this attack will be in failure. To ensure the client accepts the response from the rogue server, the attacker can first deny the service from the legitimate server using the address starvation attack mentioned before.
4. In the rogue server's response, the default gateway address points to the attacker's own machine.
5. Then, whenever the client has packets to destinations outside the local network, the packets will be sent to the rogue default gateway that would capture the content of the packets.

- The rogue default gateway needs to rewrite the packets and forward them to the correct default gateway so that the destination will not know there is a "man" in the middle intercepting the packets.
- However, the packets from the destination will go through the correct default gateway, which will forward the packets directly to the attacked client.
- Therefore, the rogue gateway cannot see any incoming packets, which is the limitation of this attack.



Man-in-the-middle attack using a rogue DHCP server

- Some Cisco switches have a DHCP security feature to prevent this exploit. The switch can be informed of which ports are attached to a DHCP server using a feature known as DHCP snooping.
- These are known as trusted ports. The switch disables DHCP answers on untrusted ports and only permits those ports to respond to DHCP queries. It stops client computers from pretending as DHCP servers when they are not connected to trusted ports.

- The Domain Name System (DNS) is the phonebook of the Internet. Humans access information online through domain names, like nytimes.com or espn.com. Web browsers interact through Internet Protocol (IP) addresses.
- DNS translates domain names to IP addresses so browsers can load Internet resources. Each device connected to the Internet has a unique IP address that other machines use to find the device.
- DNS servers eliminate the need for humans to memorize IP addresses such as 192.168.1.1 (in IPv4), or more complex newer alphanumeric IP addresses such as 2400:cb00:2048:1::c629:d7a2 (in IPv6).

Source : Cloudflare

What is DNS security?

- DNS security is the practice of protecting DNS infrastructure from cyberattacks in order to keep it performing quickly and reliably.
- An effective DNS security strategy incorporates a number of overlapping defenses, including establishing redundant DNS servers, applying security protocols like DNSSEC, and requiring rigorous DNS logging.

- Like many Internet protocols, the DNS system was not designed with security in mind and contains several design limitations.
- These limitations, combined with advances in technology, make DNS servers vulnerable to a broad spectrum of attacks, including spoofing, amplification, DoS (Denial of Service), or the interception of private personal information. And since DNS is an integral part of most Internet requests, it can be a prime target for attacks.


Source : Cloudflare

DNS spoofing/cache poisoning: This is an attack where forged DNS data is introduced into a DNS resolver's cache, resulting in the resolver returning an incorrect IP address for a domain. Instead of going to the correct website, traffic can be diverted to a malicious machine or anywhere else the attacker desires; often this will be a replica of the original site used for malicious purposes such as distributing malware or collecting login information.

DNS tunneling: This attack uses other protocols to tunnel through DNS queries and responses. Attackers can use SSH, TCP, or HTTP to pass malware or stolen information into DNS queries, undetected by most firewalls.

Source : Cloudflare

DNS hijacking: In DNS hijacking the attacker redirects queries to a different domain name server. This can be done either with malware or with the unauthorized modification of a DNS server. Although the result is similar to that of DNS spoofing, this is a fundamentally different attack because it targets the DNS record of the website on the nameserver, rather than a resolver's cache.



Source : Cloudflare

What is DNSSEC?

- DNS Security Extensions (DNSSEC) is a security protocol created to mitigate DNS based attacks. DNSSEC protects against attacks by digitally signing data to help ensure its validity.
- In order to ensure a secure lookup, the signing must happen at every level in the DNS lookup process.
- DNSSEC implements a hierarchical digital signing policy across all layers of DNS.

- For example, in the case of a 'google.com' lookup, a root DNS server would sign a key for the.COM nameserver, and the.COM nameserver would then sign a key for google.com's authoritative nameserver.
- DNSSEC is meant to work with other security measures like SSL/TLS as part of a holistic Internet security strategy.

- [1] Wong, Angus, and Alan Yeung. *Network infrastructure security*. Springer Science & Business Media, 2009.
- [2] Vincent Nestler, Gregory White, wm. Arthur Conklin, Principles of Computer Security: CompTIA Security+ and Beyond – Lab Manual, ISBN: 978-0-07-174856-8, MHID: 0-07-174856-3, McGraw Hill, 2011



Thank You