


ZCNA

ZyeemZia™ Certified Network Associate

Build networks that just work.



BEFORE YOU BEGIN

How to use this book

A modern, hands-on associate certification for the people who keep networks running: how packets travel, IPv4 and IPv6 addressing, switching and routing, everyday IP services, Wi-Fi, security fundamentals and the automation every new engineer now needs.

01

Learn

Work through 40 lessons. Each ends with a knowledge check.

02

Practise

Fill in the templates, answer the scenarios and case studies, review the glossary.

03

Certify

Take the 40-question online exam at zcna.zeemzia.com. Pass mark 80%.

04

Get certified

Download your certificate online. 90%+ and the capstone earn Distinction.

Eight-week study plan

WHEN	MODULE	DO THIS	DONE
Week 1	m1 Wire	Build a lab, capture packets	☐
Week 2	m2 Address	Subnet 50 addresses, plan IPv6	☐
Week 3	m3 Switch	VLANs, trunks, STP lab	☐
Week 4	m4 Route	Static and OSPF lab	☐
Week 5	m5 Serve	DHCP, NAT, DNS lab	☐
Week 6	m6 Air	Design a small office WLAN	☐
Week 7	m7 Guard	ACL and port-security lab	☐
Week 8	m8 Automate	Backup script, capstone, exam	☐

Learner name

Start date Target exam date

ZCNA is an independent certification programme created and issued by ZyeemZia. It is not affiliated with, endorsed by, sponsored by or equivalent to any certification from Cisco Systems, Inc. (including CCNA®, CCNP®, CCIE® and CCDE®) or from any other vendor or certifying body. All product names and trademarks belong to their respective owners. Configuration examples use widely taught, vendor-neutral style syntax.

CONTENTS

What you will learn

01 Wire — Network Fundamentals

- | | |
|--|---|
| 1. What Networks Do — and the ZCNA Packet Journey · 22 min | 4. Network Topologies and Architectures · 25 min |
| 2. Layered Models, Encapsulation and PDUs · 24 min | 5. Transport Protocols and First Troubleshooting Tools · 30 min |
| 3. Ethernet: Frames, MAC Addresses, Media and PoE · 28 min | |

02 Address — IPv4 & IPv6 Addressing

- | | |
|---|--|
| 1. IPv4 Addresses: Structure, Classes and Private Ranges · 26 min | 4. IPv6 Fundamentals: Format, Compression and Address Types · 26 min |
| 2. Subnetting Step by Step · 32 min | 5. IPv6 Configuration: EUI-64, SLAAC, DHCPv6 and Dual Stack · 30 min |
| 3. VLSM Design and Verifying Host Configuration · 30 min | |

03 Switch — Switching & VLANs

- | | |
|--|--|
| 1. How Switches Learn and Forward Frames · 24 min | 4. Spanning Tree: Loop Prevention with STP and RSTP · 32 min |
| 2. VLANs, Access Ports and 802.1Q Trunks · 30 min | 5. EtherChannel with LACP and Neighbour Discovery with LLDP · 26 min |
| 3. Inter-VLAN Routing: Router-on-a-Stick and SVIs · 28 min | |

04 Route — Routing Essentials

- | | |
|---|---|
| 1. The Routing Table: Longest Match, Administrative Distance and Metrics · 28 min | 3. OSPFv2 Concepts: Neighbours, Router ID, DR/BDR and Cost · 30 min |
| 2. Static and Default Routes for IPv4 and IPv6 · 28 min | 4. Configuring and Verifying Single-Area OSPF · 32 min |
| | 5. First-Hop Redundancy with VRRP-Style Gateways · 24 min |

05 Serve — IP Services

- | | |
|--|--|
| 1. DHCP and DNS — Addresses and Names on Demand · 30 min | 3. NTP, SNMP and Syslog — Time, Health and Evidence · 30 min |
| 2. NAT and PAT — Sharing Public Addresses · 30 min | 4. Secure Management and File Transfer · 25 min |
| | 5. QoS Basics — Protecting Voice and Video · 30 min |

06 Air — Wireless Networking

- | | |
|--|---|
| 1. RF Fundamentals — The Physics Behind Wi-Fi · 30 min | 4. Wireless Security — WPA2, WPA3 and 802.1X · 30 min |
| 2. 802.11 Standards and WLAN Building Blocks · 25 min | 5. Designing a Small Office WLAN · 30 min |
| 3. AP Modes, WLAN Architectures and Roaming · 30 min | |

07 Guard — Security Fundamentals

- | | |
|---|---|
| 1. Threats, Vulnerabilities and the Human Factor · 25 min | 4. Layer 2 Security — Port Security, DHCP Snooping and DAI · 30 min |
| 2. AAA and Device Hardening · 30 min | 5. VPNs, Firewalls, IPS and Zero Trust · 30 min |
| 3. Access Control Lists — Filtering Traffic · 35 min | |

08 Automate — Automation, Programmability and Readiness


- | | |
|---|---|
| 1. Why Automate — Traditional vs Controller-Based Networking · 25 min | 4. Intent-Based Networking and AI-Assisted Operations · 25 min |
| 2. REST APIs, JSON and YAML · 30 min | 5. The Whole Packet Journey, the Capstone and Exam Readiness · 35 min |
| 3. Configuration Management, Ansible and a First Python Script · 35 min | |

[+ Glossary](#) · [Capstone guide](#) · [Practice exam & answers](#) · [Certificate](#) · [Notes](#)

THE FRAMEWORK

The ZCNA Packet Journey

Follow one packet across eight hops: the wire it rides on, its address, the switch and router that move it, the services it uses, the air it crosses, the guards that inspect it and the code that automates it all.



01 Wire

Models, Ethernet, media, topologies and transport — the first stage of every packet's journey

03 Switch

MAC learning, VLANs, trunks, inter-VLAN routing, spanning tree, EtherChannel and LLDP

05 Serve

DHCP, DNS, NAT, time, monitoring, secure management and QoS — the services every packet depends on.

07 Guard

Threats, AAA, hardening, ACLs, layer 2 protections, VPNs, firewalls and zero trust.

02 Address

Binary, subnetting, VLSM and IPv6 — giving every interface a correct, planned address

04 Route

Routing tables, static and default routes, single-area OSPFv2 and first-hop redundancy

06 Air

RF fundamentals, 802.11 generations, WLAN architectures, roaming and security — then a real design.

08 Automate

Controllers, APIs, data formats, configuration as code, AI-assisted operations — then the whole journey and the capstone.

01

MODULE 1 · WIRE

Network Fundamentals

Models, Ethernet, media, topologies and transport — the first stage of every packet's journey

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Explain what networks do and use the eight-stage ZCNA Packet Journey to reason about problems
- Map OSI and TCP/IP layers, describe encapsulation and name each PDU
- Describe Ethernet frames, MAC addresses, duplex, cabling choices and PoE budgets
- Compare campus, data-centre, WAN, SOHO and cloud topologies
- Contrast TCP and UDP and use ping, traceroute, ipconfig/ip and arp to isolate faults

LESSONS

1. What Networks Do — and the ZCNA Packet Journey · 22 min
2. Layered Models, Encapsulation and PDUs · 24 min
3. Ethernet: Frames, MAC Addresses, Media and PoE · 28 min
4. Network Topologies and Architectures · 25 min
5. Transport Protocols and First Troubleshooting Tools · 30 min

MODULE 1 · WIRE · LESSON 1 · 22 MIN

What Networks Do — and the ZCNA Packet Journey

Devices, roles and the eight-stage story every packet follows

- Explain what a network does in terms of connectivity, sharing and services
- Identify the job of hosts, switches, routers, access points and firewalls
- Describe the eight stages of the ZCNA Packet Journey
- Distinguish a LAN, a WAN and a cloud network

Open a banking app in Dhaka, join a video class from Sylhet or scan a parcel in a Chattogram warehouse and the same thing happens underneath: your device chops information into small units, labels them with addresses, and hands them to a chain of devices that carry them — often across several countries — to the right destination and back again, in milliseconds. That chain of devices, cables, radio links and rules is a **network**. This course teaches you to design, build, verify and fix the most common parts of it.

DEFINITION · NETWORK

A **computer network** is a set of devices connected by wired or wireless links that follow shared rules (**protocols**) so they can exchange data and share resources such as internet access, files, printers and applications.

Who does what in a typical network

DEVICE	MAIN JOB	WORKS MOSTLY WITH	EVERYDAY EXAMPLE
Host / endpoint	Creates and consumes data	Applications, IP addresses	Laptop, phone, server, IP camera
Switch	Connects devices inside one LAN and forwards frames	MAC addresses	48-port switch in a floor wiring cabinet
Router	Connects different networks and chooses paths	IP addresses, routing table	Branch router linking an office to the internet
Wireless access point	Bridges Wi-Fi clients onto the wired LAN	Radio + MAC addresses	Ceiling AP in a classroom
Firewall	Allows or blocks traffic according to policy	Addresses, ports, sessions	Edge firewall in front of a data centre
Controller / cloud dashboard	Configures and monitors many devices centrally	APIs, policies	Wireless or SD-WAN controller

The ZCNA Packet Journey — eight stages

S1

Wire

Bits leave the device over copper, fibre or radio inside an Ethernet frame. Models, cables, speeds and transport protocols live here.

S2

Address

Every interface needs a correct IPv4 and/or IPv6 address, mask or prefix, and gateway — or the journey never starts.

S3

Switch

Inside the LAN, switches learn MAC addresses, separate traffic into VLANs and keep the topology loop-free.

S4

Route

Routers compare the destination with their routing table and forward the packet hop by hop between networks.

S5

Serve

Supporting services — DHCP, DNS, NAT, NTP, logging, QoS — make the journey automatic, named, translated and measurable.

S6

Air

Wireless LANs carry the first or last hop for most users today, with their own rules for radio, roaming and security.

S7

Guard

Security controls decide which packets are allowed, from which ports and users, and protect the devices themselves.

S8

Automate

Controllers, APIs and scripts configure and check hundreds of devices consistently instead of one at a time.

INSIGHT · USE THE JOURNEY AS A TROUBLESHOOTING MAP

When something breaks, ask *which stage failed?* A bad cable is a Wire problem; a wrong mask is an Address problem; a port in the wrong VLAN is a Switch problem; a missing route is a Route problem. Naming the stage narrows the search from "the network is down" to one layer of checks.

A student opens the university portal

01

Wire

The laptop's Wi-Fi radio sends frames to the access point, which puts them on a copper uplink.

→

02

Address

The laptop already has 10.20.30.45/24 and gateway 10.20.30.1 from DHCP.

→

03

Switch

Access and distribution switches forward the frame in the student VLAN toward the gateway MAC.

→

04

Route

The gateway routes the packet toward the data-centre subnet that hosts the portal.

→

05

Serve

DNS has already turned the portal name into an IP address.

→

06

Guard

A firewall checks that students may reach TCP port 443 on the portal server.

LAN, WAN and cloud network**LAN (local area network)**

- One building or campus
- Owned and run by the organisation
- Ethernet and Wi-Fi at 1–100 Gbps
- Switches and access points dominate

WAN (wide area network)

- Connects sites across cities or countries
- Usually rides on provider circuits or the internet
- Lower speed per cost, higher latency
- Routers, SD-WAN edges and VPNs dominate

Cloud network

- Virtual networks inside a provider's data centres
- Built from software objects: virtual networks, subnets, route tables, security groups
- Reached over VPN, private links or the internet
- Same IP ideas, different tools

What a well-run network delivers

- **Connectivity** — any authorised device can reach the services it needs.
- **Performance** — enough bandwidth and low enough delay for voice, video and apps.
- **Availability** — no single cable, device or link failure should stop the business.
- **Security** — only the right users and devices reach the right resources.
- **Manageability** — engineers can see, change and verify the network quickly.

A first look from a Linux host

Linux host

```
$ ip -br addr
lo                UNKNOWN          127.0.0.1/8 ::1/128
eth0              UP                10.20.30.45/24 fe80::5054:ff:fe1a:2b3c/64
$ ip route
default via 10.20.30.1 dev eth0 proto dhcp metric 100
10.20.30.0/24 dev eth0 proto kernel scope link src 10.20.30.45 metric 100
```

Two commands already tell you three Journey facts: the host's address (Address), the link state UP (Wire) and the default gateway it will send off-subnet traffic to (Route).

SCENARIO · DECIDE · WHERE WOULD YOU START?

A teacher says: "The internet is down in Room 204." Other rooms work. In Room 204 one laptop shows a valid address but cannot open any website; the projector PC in the same room shows a link light that is off.

- ☐ A Restart the core router so every room gets a fresh start
- ☐ B Treat it as two separate problems: check the projector PC's cable and port (Wire) and the laptop's gateway and DNS (Address/Serve)
- ☐ C Call the internet provider because websites will not load
- ☐ D Replace the room's access point

Circle your choice. Answer and feedback: module answer key.

PRO TIP · DRAW BEFORE YOU TYPE

Even a rough sketch — devices, links, subnets, which port goes where — saves hours. Professionals keep an up-to-date diagram and an addressing plan for every site, and they update both after every change.

KNOWLEDGE CHECK

1. Which device's main job is to forward packets **between different IP networks**?

- ☐ A Access switch
- ☐ B Wireless access point
- ☐ C Patch panel
- ☐ D Router

2. In the ZCNA Packet Journey, a laptop with the wrong subnet mask is mainly a problem at which stage?

- ☐ A Address
- ☐ B Wire
- ☐ C Switch
- ☐ D Automate

3. What best describes a cloud network in an IaaS provider?

- ☐ A A WAN circuit leased from a mobile operator
- ☐ B A Wi-Fi network shared by several buildings
- ☐ C Virtual networks, subnets and route tables defined in software
- ☐ D A physical switch rack the customer installs on site

4. Which statement about LANs and WANs is most accurate?

- ☐ A A LAN always uses fibre and a WAN always uses copper
- ☐ B A WAN links sites that are geographically far apart
- ☐ C A WAN is always faster than a LAN
- ☐ D A LAN can only connect up to 24 devices

5. Most rooms in a school work, but one room has no connectivity. What is the most sensible first move?

- ☐ A Reload the core router during class hours
- ☐ B Open an urgent fault ticket with the internet service provider
- ☐ C Reset every switch in the building to defaults
- ☐ D Check that room's local links and host settings

□ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 1 · WIRE · LESSON 2 · 24 MIN

Layered Models, Encapsulation and PDUs

OSI and TCP/IP as maps for building and troubleshooting

- Name the seven OSI layers and the four TCP/IP layers and map them to each other
- Explain encapsulation and de-encapsulation and name the PDU at each layer
- Use a layered approach to troubleshoot a connectivity problem

Networking is complicated because many jobs must happen at once: turning bits into light, finding the next device, finding the right network, delivering to the right application and making sure nothing is lost. **Layered models** split those jobs into stacked layers. Each layer offers a service to the one above and relies on the one below. That separation is why a web browser does not care whether you are on fibre or Wi-Fi.

The OSI reference model (top layer first)

7 • Application	Interface to network services	HTTP, DNS, SMTP, SSH
6 • Presentation	Data format, encoding, encryption	TLS is often placed here
5 • Session	Opens, manages and closes dialogues	Mostly handled inside apps today
4 • Transport	End-to-end delivery between applications	TCP, UDP, port numbers
3 • Network	Logical addressing and path selection	IPv4, IPv6, routers
2 • Data link	Frames on one link or LAN	Ethernet, 802.11, MAC addresses, switches
1 • Physical	Bits as voltage, light or radio	Cables, connectors, transceivers


OSI vs TCP/IP and the PDU at each layer

TCP/IP LAYER	OSI LAYERS	PDU NAME	ADDRESS OR IDENTIFIER USED
Application	7, 6, 5	Data / message	Names, URLs
Transport	4	Segment (TCP) / datagram (UDP)	Port numbers
Internet	3	Packet	IP addresses
Link (network access)	2 and 1	Frame (L2) / bits (L1)	MAC addresses

DEFINITION · PDU

A **protocol data unit** is the chunk of data a layer works with, including that layer's header (and trailer, if any). Using the right name — segment, packet, frame — tells colleagues instantly which layer you mean.

Encapsulation on the sending host



Two directions of the same idea

Encapsulation (sending)

- Moves down the stack
- Each layer adds its own header
- Ethernet also adds a trailer (FCS)
- Result: a frame on the wire

De-encapsulation (receiving)

- Moves up the stack
- Each layer reads and removes its header
- A layer checks the field that says which protocol is next (EtherType, protocol number, port)
- Result: data delivered to the right application

EXAMPLE · WHAT A ROUTER ACTUALLY REWRITES

As a packet crosses a router, the **frame** is rebuilt for every hop — new source and destination MAC addresses for each link — while the **packet's** source and destination IP addresses stay the same end to end (unless NAT is in the path, covered in the Serve module). The router also lowers the TTL by one.

Why layers help engineers

- **Troubleshooting** — test one layer at a time instead of guessing.
- **Interoperability** — any vendor's switch can carry any vendor's IP traffic.
- **Change** — upgrade copper to fibre without rewriting applications.
- **Shared vocabulary** — "it's a Layer 2 issue" means the same thing everywhere.

Bottom-up troubleshooting in five checks

1. **Layer 1** Is the link up? Check cable, port lights, interface status and error counters.
2. **Layer 2** Is the port in the right VLAN? Does the switch learn the host's MAC address?
3. **Layer 3** Does the host have the right IP, mask and gateway? Can it ping the gateway?
4. **Layer 4** Is the service listening on the expected port? Is a firewall blocking it?
5. **Layer 7** Does the name resolve? Does the application return an error or a certificate warning?

Seeing the layers in one capture summary

Linux host

```
# summary of one captured frame (fields shown per layer)
Ethernet II, Src: 52:54:00:1a:2b:3c, Dst: 00:11:22:33:44:55, Type: IPv4 (0x0800)
Internet Protocol Version 4, Src: 10.20.30.45, Dst: 10.50.0.10, TTL: 64
Transmission Control Protocol, Src Port: 51544, Dst Port: 443, Flags: [SYN]
# Layer 2 -> Layer 3 -> Layer 4: the EtherType says IPv4 follows, the protocol field says TCP follows
```

Packet analysers display each header on its own line. Reading them top to bottom is reading the stack from Layer 2 upward.

SCENARIO · DECIDE · CHOOSING A STARTING LAYER

A user can ping the file server by IP address but cannot open `\\files.campus.local`. Other users can.

- ☐ A Start at Layer 1 and replace the user's patch cable
- ☐ B Reboot the access switch
- ☐ C Change the server's IP address
- ☐ D Check name resolution (DNS) on the user's machine

Circle your choice. Answer and feedback: module answer key.

WATCH OUT · MODELS ARE MAPS, NOT LAWS

Real protocols do not always fit neatly. TLS, QUIC and ARP are argued about in every classroom. Use the layers to organise your thinking — do not lose points arguing about edge cases.

KNOWLEDGE CHECK

6. What is the PDU called at the transport layer when TCP is used?

☐ A Frame	☐ B Segment
☐ C Packet	☐ D Bit stream
7. Which TCP/IP layer corresponds to OSI Layer 3?

☐ A Transport	☐ B Application
☐ C Internet	☐ D Link
8. As a packet crosses three routers to reach a server, which fields normally change at every hop?

☐ A Source and destination MAC addresses	☐ B Source and destination IPv4 or IPv6 addresses
☐ C TCP source and destination ports	☐ D The application data payload
9. A host pings its gateway successfully but a web app times out on TCP 443. Which layer should you test next?

☐ A Layer 1 — reseal and re-test the patch cable	☐ B Layer 2 — clear the switch MAC address table
☐ C Layer 4 — test whether TCP 443 is reachable	☐ D Layer 1 — reinstall the network card driver
10. During de-encapsulation, how does the receiving host know that the Ethernet payload is IPv4?

☐ A The EtherType field value	☐ B The destination MAC address
☐ C The FCS trailer	☐ D The TTL value

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 1 · WIRE · LESSON 3 · 28 MIN

Ethernet: Frames, MAC Addresses, Media and PoE

The Wire stage in detail — from frame fields to fibre and powered ports

- Describe the fields of an Ethernet frame and the structure of a MAC address
- Explain speed and duplex negotiation and recognise a duplex mismatch
- Choose between copper, multimode and single-mode fibre
- Size Power over Ethernet for phones, APs and cameras

Ethernet (IEEE 802.3) is the language of wired LANs, from a home router to a 400 Gbps data-centre fabric. Every wired frame you will troubleshoot in this course is an Ethernet frame, so it pays to know its parts by heart.

Ethernet II frame, field by field

FIELD	SIZE	PURPOSE
Preamble + SFD	7 + 1 bytes	Synchronises the receiver; marks the start of the frame (not counted in frame size)
Destination MAC	6 bytes	Who should receive the frame (unicast, multicast or broadcast)
Source MAC	6 bytes	Who sent it — switches learn from this field
802.1Q tag (optional)	4 bytes	VLAN ID and priority on trunk links
EtherType	2 bytes	Next protocol: 0x0800 IPv4, 0x86DD IPv6, 0x0806 ARP
Payload	46–1500 bytes	The packet; padded if shorter than 46 bytes
FCS	4 bytes	CRC checksum; receivers discard frames that fail it

DEFINITION · MAC ADDRESS

A **MAC address** is a 48-bit Layer 2 identifier, written as 12 hexadecimal digits (for example `00:1A:2B:3C:4D:5E` or `001a.2b3c.4d5e`). The first 24 bits are normally the manufacturer's **OUI**; the last 24 bits are assigned by that manufacturer.

MAC address facts you will use

- **Unicast** — one interface. The lowest bit of the first byte is 0.
- **Multicast** — a group; lowest bit of the first byte is 1 (e.g. `01:00:5E:...` for IPv4 multicast, `33:33:...` for IPv6).
- **Broadcast** — `FF:FF:FF:FF:FF:FF`, every device in the VLAN.
- **Locally administered** addresses (second-lowest bit of the first byte set) are common on virtual machines and phones that randomise their Wi-Fi MAC for privacy.
- Standard frame size is **64–1518 bytes** untagged, not counting the preamble (1522 with an 802.1Q tag); larger "jumbo" frames must be enabled end to end.

Reading speed, duplex and errors

Switch SW1 · IOS-style CLI

```
SW1# show interfaces status
Port      Name           Status      Vlan Duplex Speed Type
Gi1/0/1   PC-Accounts    connected   10   a-full a-1000 10/100/1000BaseTX
Gi1/0/2   Printer-2F     connected   10   half   100    10/100/1000BaseTX
Gi1/0/3                   notconnect  1     auto   auto   10/100/1000BaseTX

SW1# show interfaces Gi1/0/2 | include duplex|CRC|collision
Half-duplex, 100Mb/s, media type is 10/100/1000BaseTX
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
  0 output errors, 214 collisions, 0 interface resets
  37 late collision, 0 deferred
```

The `a-` prefix means the value was auto-negotiated. Gi1/0/2 has been hard-set to 100/half; late collisions on a half-duplex port are the classic sign that the other end is running full duplex.

WATCH OUT · DUPLEX MISMATCH: THE SLOW, SILENT FAULT

If one end is fixed at full duplex and the other auto-negotiates, the auto side often falls back to **half duplex**. The link comes up and pings work, but throughput collapses: the half-duplex side logs late collisions and the full-duplex side logs CRC/run errors. Best practice: leave both ends on auto, or hard-set **both** ends identically.

Choosing the medium

Copper twisted pair

- RJ-45, up to 100 m per channel
- Cat5e: 1 Gbps; Cat6A: 10 Gbps at 100 m
- Carries PoE
- Cheapest for desks, phones, APs, cameras

Multimode fibre (MMF)

- Larger core, cheaper short-reach optics (e.g. 850 nm)
- Typical 10G short-reach: ~300 m on OM3, ~400 m on OM4
- Inside buildings and data centres
- Immune to electrical interference

Single-mode fibre (SMF)

- Tiny core, laser optics (1310/1550 nm)
- Kilometres to tens of kilometres
- Between buildings, campuses and to providers
- Higher optic cost, lowest loss

Common Ethernet standards

STANDARD	SPEED	MEDIUM	TYPICAL REACH
1000BASE-T	1 Gbps	Cat5e or better copper, 4 pairs	100 m
2.5G/5GBASE-T	2.5 / 5 Gbps	Cat5e/Cat6 copper	100 m — popular for Wi-Fi 6/7 AP uplinks
10GBASE-T	10 Gbps	Cat6A copper	100 m (Cat6 shorter)
10GBASE-SR	10 Gbps	Multimode fibre	Up to ~300–400 m (OM3/OM4)
10GBASE-LR	10 Gbps	Single-mode fibre	10 km
25G/100G short-reach	25 / 100 Gbps	Multimode fibre or direct-attach copper	Tens of metres — server and leaf links

Power over Ethernet at a glance

IEEE STANDARD	COMMON NAME	POWER AT THE SWITCH PORT (PSE)	TYPICAL DEVICES
802.3af (Type 1)	PoE	15.4 W	Desk IP phones, basic cameras
802.3at (Type 2)	PoE+	30 W	Most Wi-Fi access points, PTZ-lite cameras
802.3bt (Type 3)	PoE++ / 4PPoE	60 W	High-end APs, video phones, small displays
802.3bt (Type 4)	PoE++ / 4PPoE	90 W	Large PTZ cameras, LED lighting, thin clients

PRO TIP · BUDGET THE WHOLE SWITCH, NOT JUST THE PORT

A 48-port PoE switch may support 30 W per port but only, say, 740 W in total. Add up the worst-case draw of every powered device plus headroom before you buy — APs and PTZ cameras draw most during boot and heater cycles.

SCENARIO · DECIDE · THE SLOW PRINTER PORT

Users say printing to the second-floor printer takes minutes. The switch port (shown above as Gi1/0/2) is hard-coded to 100/half; the printer's NIC is set to auto. Pings succeed.

- ☐ A Replace the printer because its NIC is clearly faulty
- ☐ B Move the printer to a different VLAN
- ☐ C Set the switch port back to auto speed and duplex (or match both ends manually) and clear the counters
- ☐ D Enable PoE on the port

Circle your choice. Answer and feedback: module answer key.

Copper and fibre cabling checklist

- ☐ Copper channel length (patch + horizontal + patch) is 100 m or less
- ☐ Cable category matches the speed (Cat6A for 10GBASE-T at full length)
- ☐ Every cable is labelled at both ends with the same ID as the patch-panel port
- ☐ Fibre type (OM3/OM4/OS2) matches the optics at both ends
- ☐ Optic wavelength and standard are identical at both ends (e.g. SR to SR)
- ☐ Fibre ends are inspected and cleaned before connecting
- ☐ PoE budget calculated with at least 20% headroom
- ☐ Speed/duplex left on auto at both ends, or hard-set identically
- ☐ Interface error counters checked after installation (CRC, input errors, collisions)
- ☐ Test results (certifier or light-meter readings) saved with the site documentation

KNOWLEDGE CHECK

11. Which field in an Ethernet frame does a switch use to **learn** where a device is connected?
 - ☐ A Destination MAC address
 - ☐ B Source MAC address
 - ☐ C EtherType
 - ☐ D Frame check sequence
12. A port shows many late collisions and is running half duplex. What is the most likely cause?
 - ☐ A The cable is longer than 10 metres
 - ☐ B The access VLAN ID is higher than 1000
 - ☐ C The port is supplying PoE+
 - ☐ D The far end is running full duplex
13. Two buildings are 4 km apart and need a 10 Gbps link. Which medium fits best?
 - ☐ A Single-mode fibre with long-reach optics
 - ☐ B Cat6A copper with 10GBASE-T
 - ☐ C Multimode OM4 fibre with 850 nm short-reach optics
 - ☐ D Cat5e copper with a PoE injector
14. An access point needs up to 25 W. Which is the lowest PoE standard that can power it?
 - ☐ A 802.3af (PoE)
 - ☐ B 802.3at (PoE+)
 - ☐ C 802.3bt Type 4
 - ☐ D 802.3bt Type 3
15. Which destination MAC address reaches every device in the same VLAN?
 - ☐ A 00:00:00:00:00:00
 - ☐ B 01:00:5E:00:00:01
 - ☐ C 33:33:00:00:00:01
 - ☐ D FF:FF:FF:FF:FF:FF

16. Why is a duplex mismatch hard to spot?

- ☐ A The interface goes down every few seconds
- ☐ B The switch disables the port immediately
- ☐ C Pings still succeed while throughput collapses
- ☐ D The host loses its IP address

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 1 · WIRE · LESSON 4 · 25 MIN

Network Topologies and Architectures

Campus tiers, spine-leaf, WAN, SOHO and cloud — choosing the right shape

- Compare two-tier (collapsed core) and three-tier campus designs
- Explain why data centres use spine-leaf fabrics
- Describe common WAN, SOHO and cloud connectivity options

A **topology** is the shape of a network: which devices connect to which, and how traffic flows between them. Good topologies are predictable — every access switch is the same number of hops from the core, every path has a backup, and growth means adding a block rather than redesigning everything.

Three-tier campus (top layer first)

Core	Very fast, very reliable backbone	Connects distribution blocks, data centre and WAN edge	No user ports, minimal policy
Distribution	Aggregates access switches per building or floor	Default gateways for VLANs, routing, policy	Usually deployed in redundant pairs
Access	Ports for users, phones, APs, printers, cameras	VLAN assignment, PoE, port security	Dual uplinks to distribution
Endpoints	Laptops, phones, IoT, access points		

Two-tier vs three-tier campus

Two-tier (collapsed core) • Core and distribution combined in one redundant pair • Fewer boxes, lower cost • Ideal for a single building or small campus • Typical for a 100–1,500-user office	Three-tier • Separate core layer joins many distribution blocks • Scales to large multi-building campuses • Failure in one building does not touch the core design • More devices and cabling to manage
--	---

Data centres moved from tree-shaped designs to **spine-leaf** fabrics because modern applications send most traffic *between* servers (east-west) rather than in and out (north-south). In a spine-leaf fabric every leaf switch connects to every spine switch, and leaves never connect to each other directly.

Spine-leaf rules of thumb

- Any server is exactly **leaf** → **spine** → **leaf** away from any other: predictable latency.
- All uplinks are active at once using equal-cost routing — no links sit blocked.
- Need more servers? Add a leaf. Need more bandwidth between leaves? Add a spine.
- Servers, storage, firewalls and border routers all attach to leaves.

WAN connectivity options

OPTION	WHAT IT IS	STRENGTHS	WATCH OUT FOR
Leased line / point-to-point	Dedicated circuit between two sites	Predictable, private	Cost per Mbps, single path

OPTION	WHAT IT IS	STRENGTHS	WATCH OUT FOR
Provider MPLS-style VPN	Provider-run private routed network	SLAs, QoS across the WAN	Long lead times, price
Internet + IPsec VPN	Encrypted tunnels across the public internet	Cheap, fast to deploy	No end-to-end performance guarantee
SD-WAN	Controller-managed edges using several links at once	Uses broadband, fibre and 4G/5G together; app-aware paths	Needs good design of policies and security
4G/5G cellular	Mobile data link	Quick backup or temporary site link	Variable latency, data caps

EXAMPLE · A SOHO NETWORK IN ONE BOX

A small shop in Mirpur typically has one **all-in-one router**: a router with NAT to the ISP, a small built-in switch, a Wi-Fi access point, a DHCP server and a basic firewall — all in one device. The same functions exist in an enterprise; they are just split across dedicated, redundant devices.

On-premises vs cloud networking

On-premises

- You buy, rack and cable switches and routers
- VLANs and subnets configured device by device
- Capacity changes take weeks
- You control every layer

Public cloud (IaaS)

- Virtual networks and subnets created through a console or API
- Route tables, gateways and security groups are software objects
- Capacity changes take minutes
- Provider controls the physical layers; you control the logical ones

INSIGHT · HYBRID IS THE NORMAL CASE

Most organisations in 2026 run some workloads on-premises, some in one or more clouds and many as SaaS. The WAN is therefore no longer "branch to head office" but "every site and user to every cloud". That is why IP addressing must be planned so on-prem and cloud ranges never overlap.

CASE STUDY · A DHAKA HEAD OFFICE REFRESH

A garments buying house in Gulshan, Dhaka has 220 staff over three floors of one building, a small server room, two internet providers and an online ERP hosted in a cloud provider. Today every floor has one unmanaged switch daisy-chained to the next, and when the second-floor switch fails, the third floor loses everything.

1. Which campus design (two-tier or three-tier) fits, and why?

2. How should the floor switches connect to remove the daisy-chain weakness?

3. How should the office reach the cloud ERP, and what protects that path?

Model answer: module answer key.

SCENARIO · DECIDE · CHOOSING A DATA-CENTRE SHAPE

A fintech start-up in Banani is building a small data hall with 12 racks of virtualised servers that talk to each other constantly. The vendor proposes a classic tree: two core switches, each rack switch connected to only one of them.

- A Ask for a spine-leaf fabric: each rack (leaf) switch uplinks to every spine, using equal-cost paths
- B Accept the tree — it is cheaper and simpler
- C Connect all rack switches to each other in a ring
- D Put all servers on one large switch

Circle your choice. Answer and feedback: module answer key.

TEMPLATE Site topology and inventory worksheet

Record every network device, its role and its uplinks so diagrams and troubleshooting start from facts.

DEVICE NAME	ROLE / TIER	LOCATION	MODEL CATEGORY & PORTS	UPLINK 1 → PEER:PORT	UPLINK 2 → PEER:PORT	MANAGEMENT IP
DHK-F2-AS1	Access	Floor 2 IDF	48-port PoE+ L2 switch	DHK-CORE1 Gi1/0/2	DHK-CORE2 Gi1/0/2	10.10.99.12

KNOWLEDGE CHECK

17. In a spine-leaf fabric, how are leaf switches connected?

- ☐ A Each leaf connects to every other leaf
- ☐ B Each leaf connects to every spine
- ☐ C Leaves connect in a ring to one spine
- ☐ D Each leaf connects to exactly one spine

18. A single-building office of 300 users needs redundancy but a small budget. Which design fits best?

- ☐ A Three-tier with a separate core
- ☐ B Spine-leaf data-centre fabric
- ☐ C Two-tier collapsed core
- ☐ D Daisy-chained unmanaged switches

19. Which layer of a three-tier campus normally hosts the default gateways for user VLANs?

- ☐ A Access
- ☐ B Core
- ☐ C Data-centre server block
- ☐ D Distribution

20. A branch needs a WAN link this week at low cost. Which option is quickest?

- ☐ A Internet broadband with an IPsec VPN
- ☐ B A new provider MPLS-style circuit
- ☐ C A dedicated leased line
- ☐ D A dark-fibre build all the way to head office

21. Why must on-premises and cloud IP ranges be planned together?

- ☐ A Overlapping ranges make routing between them fail
- ☐ B Cloud providers forbid private address ranges
- ☐ C Cloud subnets must always use public addresses
- ☐ D Overlaps only affect wireless users

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 1 · WIRE · LESSON 5 · 30 MIN

Transport Protocols and First Troubleshooting Tools

TCP vs UDP, ports, the three-way handshake — and the commands you will run every day

- Compare TCP and UDP and match common applications to each
- Describe the TCP three-way handshake and the role of port numbers
- Use ping, traceroute, ipconfig/ip and arp to locate a fault

IP gets a packet to the right **host**. The transport layer gets the data to the right **application** on that host and decides how careful delivery must be. Two protocols do almost all of this work: **TCP** for reliable, ordered streams and **UDP** for fast, lightweight messages.

TCP vs UDP

TCP

- Connection-oriented: handshake before data
- Reliable: sequence numbers, acknowledgements, retransmission
- Flow control (window) and congestion control
- 20-byte minimum header
- Web (HTTP/1.1, HTTP/2), SSH, email, file transfer

UDP

- Connectionless: just send
- No built-in retransmission or ordering
- 8-byte header, very low overhead
- Applications add reliability if they need it
- DNS queries, DHCP, VoIP/video media, NTP, SNMP, QUIC (HTTP/3)

TCP three-way handshake (client to web server)

01 SYN

Client → server: "I want to connect; my starting sequence number is x."



02 SYN-ACK

Server → client: "Agreed; my sequence number is y; I acknowledge x+1."



03 ACK

Client → server: "I acknowledge y+1." The connection is established.



04 Data

Both sides send data; every byte is acknowledged.



05 FIN / ACK

Each side closes its direction with FIN and ACK (or a RST aborts at once).

Port numbers you should recognise

PORT	TRANSPORT	SERVICE
20/21	TCP	FTP data/control
22	TCP	SSH (secure remote CLI)
23	TCP	Telnet (insecure — avoid)
25	TCP	SMTP (mail transfer)
53	UDP and TCP	DNS
67/68	UDP	DHCP server/client
69	UDP	TFTP
80 / 443	TCP (443 also UDP for QUIC)	HTTP / HTTPS
123	UDP	NTP
161/162	UDP	SNMP / SNMP traps
514	UDP	Syslog

Numbers behind the transport layer

MEASURE	FORMULA	WHAT IT TELLS YOU
Port ranges	0–1023 well-known · 1024–49151 registered · 49152–65535 dynamic	Servers listen on well-known or registered ports; clients pick a temporary (ephemeral) source port.
Socket	socket = IP address + protocol + port	A TCP connection is identified by both sockets: client IP:port and server IP:port.
Bandwidth-delay product	BDP = bandwidth × round-trip time	How much data must be "in flight" to fill a link; e.g. 100 Mbps × 0.04 s = 4 Mb ≈ 500 KB.

ping and traceroute from Linux

Linux host

```
$ ping -c 3 10.20.30.1
64 bytes from 10.20.30.1: icmp_seq=1 ttl=64 time=0.61 ms
64 bytes from 10.20.30.1: icmp_seq=2 ttl=64 time=0.55 ms
64 bytes from 10.20.30.1: icmp_seq=3 ttl=64 time=0.58 ms
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
$ traceroute -n 203.0.113.80
 1  10.20.30.1      0.70 ms  0.66 ms  0.63 ms
 2  100.64.12.1    4.91 ms  5.02 ms  4.89 ms
 3  * * *
 4  203.0.113.80  18.44 ms 18.39 ms 18.51 ms
```

Each traceroute line is one router hop, found by sending probes with TTL 1, 2, 3... and listening for ICMP "time exceeded" replies. A * * * line only means that hop did not answer — the trace still reached the target.

ipconfig and arp on Windows

Windows host

```
C:\> ipconfig /all
IPv4 Address. . . . . : 10.20.30.45(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 10.20.30.1
DNS Servers . . . . . : 10.50.0.53
C:\> arp -a
Interface: 10.20.30.45 --- 0x7
Internet Address      Physical Address      Type
10.20.30.1            00-11-22-33-44-55    dynamic
```

ARP maps an IPv4 address to a MAC address. Seeing the gateway in the ARP table proves the host reached it at Layer 2.

Same job, three platforms

TASK	WINDOWS	LINUX	IOS-STYLE DEVICE
Show IP settings	ipconfig /all	ip addr / ip -br addr	show ip interface brief
Show routes	route print	ip route	show ip route
Show ARP / neighbour cache	arp -a	ip neigh	show ip arp
Test reachability	ping	ping	ping
Trace the path	tracert (ICMP probes)	traceroute (UDP probes by default)	traceroute

TASK	WINDOWS	LINUX	IOS-STYLE DEVICE
Test a TCP port	Test-NetConnection host - Port 443	nc -vz host 443	telnet host 443 (as a quick test)

A five-ping method for "I can't reach the server"

1. **Ping 127.0.0.1** Confirms the host's own TCP/IP stack works.
2. **Ping your own IP** Confirms the interface is configured and up.
3. **Ping the default gateway** Confirms the local LAN, VLAN and gateway are fine.
4. **Ping a remote IP** Confirms routing beyond the gateway (a firewall may block ICMP — note it).
5. **Ping the server by name** Confirms DNS. If the IP works but the name fails, the problem is name resolution.

SCENARIO · DECIDE · READING THE CLUES

A host's `ipconfig` shows IP 10.20.30.45/24 and gateway 10.20.31.1. It can ping other PCs in 10.20.30.0/24 but nothing outside.

- ☐ A Replace the switch port because the link is unstable
- ☐ B Change the mask to /16 so everything becomes local
- ☐ C Correct the default gateway to the router's address inside 10.20.30.0/24
- ☐ D Flush the DNS cache

Circle your choice. Answer and feedback: module answer key.

REFLECT

Think of the last time "the internet" failed for you at home, work or university. Using the Packet Journey stages, which stage do you now think failed, and which single command would have proved it fastest?

Key terms

Protocol

A set of agreed rules for formatting and exchanging data between devices.

PDU

Protocol data unit — the data unit at a given layer: data, segment, packet, frame, bits.

EtherType

Ethernet field naming the payload protocol, e.g. 0x0800 IPv4, 0x86DD IPv6.

Duplex

Whether a link can send and receive at the same time (full) or only one at a time (half).

Spine-leaf

Data-centre topology where every leaf switch connects to every spine switch.

Encapsulation

Adding a layer's header (and trailer) around data as it moves down the stack.

MAC address

48-bit Layer 2 interface identifier; first 24 bits usually the manufacturer OUI.

FCS

Frame check sequence — a CRC at the end of a frame used to detect corruption.

PoE

Power over Ethernet — IEEE 802.3af/at/bt delivery of DC power over twisted-pair cable.

Collapsed core

Two-tier campus design where core and distribution roles share one switch pair.

Three-way handshake

TCP connection set-up: SYN, SYN-ACK, ACK.

TTL

IP time-to-live field, decreased by each router; traceroute uses it to find hops.

Ephemeral port

Temporary client-side port chosen from the dynamic range for one connection.

ARP

Address Resolution Protocol — finds the MAC address for a known IPv4 address on the local network.

KNOWLEDGE CHECK

22. Which application traffic is **usually** carried by UDP?

- ☐ A An SSH session to a router
- ☐ B Sending email between mail servers
- ☐ C An HTTP/1.1 file download
- ☐ D Voice media in a VoIP call

23. What is the correct order of the TCP three-way handshake?

- ☐ A SYN, ACK, FIN
- ☐ B ACK, SYN, SYN-ACK
- ☐ C SYN, SYN-ACK, ACK
- ☐ D SYN, ACK, SYN-ACK

24. A trace shows * * * at hop 3 and then reaches the destination at hop 4. What does hop 3 tell you?

- ☐ A Every packet is being dropped at hop 3 by an ACL
- ☐ B That router did not answer the probes
- ☐ C The destination is unreachable
- ☐ D The TTL was set too high

25. A host pings its gateway and remote IPs but fails when using server names. Which service should you check?

- ☐ A ARP
- ☐ B DNS
- ☐ C DHCP
- ☐ D STP

26. Which port and transport pair is correct?

- ☐ A SSH — TCP 22
- ☐ B DNS — TCP 80
- ☐ C HTTPS — UDP 161
- ☐ D NTP — TCP 25

27. Which command on a Linux host shows the ARP/neighbour cache?

- ☐ A `ip route`
- ☐ B `ss -t`
- ☐ C `traceroute -n`
- ☐ D `ip neigh`

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 1 · WIRE

Answer key

Knowledge checks

- 1. D** A router reads the destination IP address and consults its routing table to move packets between networks. A switch forwards frames inside a LAN, an access point bridges radio clients onto the LAN, and a patch panel is passive cabling.
- 2. A** The mask is part of the host's IP configuration, so it belongs to the Address stage. Wire covers physical links and frames, Switch covers MAC learning and VLANs, and Automate covers programmability.
- 3. C** Cloud networks are built from software objects — virtual networks, subnets, route tables and security rules — on the provider's hardware. The other options describe physical WAN, wireless or on-premises gear.
- 4. B** WANs connect distant sites, usually over provider services. Media choice is not fixed to either, LANs are usually faster per cost than WANs, and LAN size is not limited to 24 devices.
- 5. D** When the fault is isolated to one area, start where the symptoms are. Reloading the core or resetting switches causes wider outages, and the provider link clearly works for the other rooms.
- 6. B** TCP's PDU is a segment (UDP's is usually called a datagram). Frames belong to Layer 2, packets to Layer 3, and bits to Layer 1.
- 7. C** OSI Layer 3 (network) maps to the TCP/IP internet layer, where IP lives. Transport is OSI 4, application covers OSI 5–7, and link covers OSI 1–2.
- 8. A** Each link has its own frame, so MAC addresses are rewritten per hop. IP addresses and ports stay the same end to end unless NAT is used, and the payload is untouched.
- 9. C** A working ping to the gateway shows Layers 1–3 are fine locally, so the next check is Layer 4 reachability to port 443. The other options re-test layers already proven to work.
- 10. A** EtherType (0x0800 for IPv4, 0x86DD for IPv6) identifies the next protocol. The MAC address identifies the receiver, the FCS detects errors, and the TTL limits hop count.
- 11. B** Switches record the source MAC and the arriving port in their MAC table. The destination MAC is used for forwarding, EtherType identifies the payload protocol, and the FCS detects errors.
- 12. D** Late collisions on a half-duplex port almost always indicate a duplex mismatch with a full-duplex neighbour. A 10 m cable is well within limits, VLAN numbers do not cause collisions, and PoE does not affect duplex.
- 13. A** Single-mode fibre and LR optics reach 10 km. Copper is limited to about 100 m and multimode short-reach optics to a few hundred metres.
- 14. B** 802.3at delivers 30 W at the switch port, enough for a 25 W device after cable loss. 802.3af offers only 15.4 W; the 802.3bt types would work but are more than the minimum required.
- 15. D** All ones is the Ethernet broadcast address. All zeros is not a valid destination, 01:00:5E... is IPv4 multicast, and 33:33... is IPv6 multicast.
- 16. C** Small pings rarely trigger collisions, so basic tests pass while real traffic suffers. The link stays up, the switch does not shut the port, and IP settings are unaffected.
- 17. B** Full leaf-to-spine connectivity gives equal-cost paths and a fixed hop count. Leaf-to-leaf links, rings and single uplinks break that predictability.
- 18. C** Collapsing core and distribution into a redundant pair suits one building. A separate core adds cost without benefit here, spine-leaf targets server fabrics, and daisy-chains create single points of failure.
- 19. D** Distribution switches aggregate access blocks and usually provide layer 3 gateways and policy. The core focuses on fast transport, access provides user ports, and the data-centre block hosts servers rather than user gateways.
- 20. A** Internet VPNs can be deployed as soon as broadband is available. Provider private circuits and fibre builds usually take weeks or months.
- 21. A** If the same prefix exists on both sides, routers cannot decide where to send traffic. Cloud networks commonly use private ranges, do not require public addressing for subnets, and overlaps affect every user.
- 22. D** Real-time voice prefers UDP because late retransmissions are useless. SSH, SMTP and HTTP/1.1 all rely on TCP's reliability.
- 23. C** The client sends SYN, the server answers SYN-ACK and the client confirms with ACK. The other sequences are not valid TCP set-up exchanges.
- 24. B** Many routers rate-limit or block ICMP time-exceeded messages; traffic still passes, as hop 4 proves. Nothing is dropped, the destination answered, and TTL is set correctly by the tool.
- 25. B** Working IP connectivity with failing names points to name resolution. ARP and DHCP already worked (the host reaches its gateway), and STP is a switch loop-prevention protocol.
- 26. A** SSH uses TCP 22. DNS uses port 53, HTTPS uses 443, SNMP uses 161 and NTP uses UDP 123; TCP 25 is SMTP.
- 27. D** `ip neigh` lists IP-to-MAC mappings. `ip route` shows routes, `ss -t` lists TCP sockets and `traceroute -n` traces the path.

Scenarios & case studies

Scenario · Where would you start?

B is best. Correct. The symptoms point to different Journey stages. Splitting them avoids chasing one imaginary fault.

Scenario · Choosing a starting layer

D is best. Correct. IP reachability is proven; the failing piece is turning the name into an address — an application-layer service.

Scenario · The slow printer port

C is best. Correct. Matching duplex removes the late collisions and restores throughput; clearing counters lets you confirm the fix.

Case study · A Dhaka head office refresh

A **two-tier (collapsed-core)** design fits one building with 220 users: a redundant pair of layer 3 core/distribution switches in the server room and one managed access switch (or stack) per floor. Each access switch gets **two uplinks — one to each core switch** — so no floor depends on another. The core pair hosts the VLAN gateways. For the ERP, use both ISPs: an **SD-WAN or dual-ISP edge** with encrypted tunnels (or a private cloud interconnect if budget allows) to the cloud network, with automatic failover between providers and addressing that does not overlap the cloud ranges.

Scenario · Choosing a data-centre shape

A is best. Correct. Spine-leaf gives every server pair the same hop count, keeps all links active and grows by adding leaves or spines.

Scenario · Reading the clues

C is best. Correct. 10.20.31.1 is not in the host's own subnet, so the host cannot ARP for it and off-subnet traffic has nowhere to go.

02

MODULE 2 · ADDRESS

IPv4 & IPv6 Addressing

Binary, subnetting, VLSM and IPv6 — giving every interface a correct, planned address

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Convert IPv4 addresses and masks between decimal and binary and recognise special ranges
- Subnet any IPv4 block and calculate network, broadcast and usable range quickly
- Design and document a VLSM addressing plan and verify hosts against it
- Read, compress and classify IPv6 addresses and build EUI-64 interface IDs
- Configure IPv6 with SLAAC or DHCPv6 and plan a dual-stack rollout

LESSONS

1. IPv4 Addresses: Structure, Classes and Private Ranges · 26 min
2. Subnetting Step by Step · 32 min
3. VLSM Design and Verifying Host Configuration · 30 min
4. IPv6 Fundamentals: Format, Compression and Address Types · 26 min
5. IPv6 Configuration: EUI-64, SLAAC, DHCPv6 and Dual Stack · 30 min

MODULE 2 · ADDRESS · LESSON 1 · 26 MIN

IPv4 Addresses: Structure, Classes and Private Ranges

Binary, masks, the classful past, CIDR and the ranges you must recognise

- Convert between dotted-decimal and binary for any octet
- Explain network and host portions using a mask or prefix length
- Contrast the historical class system with CIDR
- Recognise private, loopback, link-local, CGN and documentation ranges

An IPv4 address is **32 bits**, written as four 8-bit **octets** in dotted decimal — for example `192.168.10.25`. Part of those bits identifies the **network** (which subnet) and the rest identifies the **host** (which interface on that subnet). The **subnet mask** or **prefix length** tells everyone where the dividing line falls.

Binary place values of one octet: 128 · 64 · 32 · 16 · 8 · 4 · 2 · 1

DECIMAL	BINARY	PLACE VALUES THAT ARE 1
192	11000000	128 + 64
172	10101100	128 + 32 + 8 + 4
10	00001010	8 + 2
248	11111000	128 + 64 + 32 + 16 + 8
255	11111111	all eight bits

Decimal to binary by subtraction

1. **Start at 128** If the number is 128 or more, write 1 and subtract 128; otherwise write 0.
2. **Move right** Repeat with 64, 32, 16, 8, 4, 2 and 1 using what is left.
3. **Check** Add the place values that hold a 1 — the total must equal the original number.
4. **Example: 200** $200 - 128 = 72 \rightarrow 1$; $72 - 64 = 8 \rightarrow 1$; $32 \rightarrow 0$; $16 \rightarrow 0$; $8 - 8 = 0 \rightarrow 1$; then $0,0,0 \rightarrow 11001000$.

DEFINITION · SUBNET MASK AND PREFIX LENGTH

A **subnet mask** is 32 bits with all network bits set to 1 and all host bits set to 0. `255.255.255.0` has 24 ones, so it is written in **CIDR prefix notation** as `/24`. `10.1.4.20/22` means "the first 22 bits identify the network".

The classful system (historical)

CLASS	FIRST OCTET	DEFAULT MASK	ORIGINAL PURPOSE
A	1–126	/8 (255.0.0.0)	Very large networks (127 reserved for loopback)
B	128–191	/16 (255.255.0.0)	Medium networks
C	192–223	/24 (255.255.255.0)	Small networks
D	224–239	—	Multicast groups
E	240–255	—	Reserved / experimental

Classful thinking vs CIDR

Classful (1980s)

- Mask implied by the first octet
- Only /8, /16 or /24 networks
- Huge waste: a 300-host site needed a whole /16 (65,534 hosts)
- Still useful to recognise multicast (D) and reserved (E) space

CIDR (1993 onward, today)

- Prefix length written explicitly: /8 to /32
- Any boundary: /22, /27, /30...
- Allows route summarisation (aggregation)
- The only model routers and clouds use today

Special IPv4 ranges to recognise on sight

RANGE	NAME	USE
10.0.0.0/8	Private (RFC 1918)	Large enterprise and cloud networks
172.16.0.0/12	Private (RFC 1918)	172.16.0.0 – 172.31.255.255
192.168.0.0/16	Private (RFC 1918)	Homes, SOHO, labs
100.64.0.0/10	Shared address space	Carrier-grade NAT between ISP and customer
127.0.0.0/8	Loopback	The host itself (127.0.0.1)
169.254.0.0/16	Link-local	Self-assigned when DHCP fails (often called APIPA)
192.0.2.0/24, 198.51.100.0/24, 203.0.113.0/24	Documentation	Examples in books and training — never routed
224.0.0.0/4	Multicast	One-to-many groups (e.g. routing protocols)

INSIGHT · PRIVATE ADDRESSES NEED TRANSLATION TO REACH THE INTERNET

Private ranges are not routed on the public internet, so thousands of organisations reuse `10.0.0.0/8` safely. To browse the web, traffic from private hosts is translated to a public address by NAT — you will configure that in the Serve module.

Counting addresses

MEASURE	FORMULA	WHAT IT TELLS YOU
Total addresses	$2^{(32 - \text{prefix})}$	A /24 has $2^8 = 256$ addresses; a /22 has $2^{10} = 1,024$.
Usable hosts	$2^h - 2$ (h = host bits)	Subtract the network address (all host bits 0) and broadcast address (all host bits 1). A /24 gives 254 hosts.
Exceptions	/31 → 2 usable · /32 → 1 address	A /31 is used on point-to-point links with no broadcast; a /32 identifies a single host or loopback.

Seeing addresses on a router

R1 · IOS-style CLI

```
R1# show ip interface brief
Interface                IP-Address      OK? Method Status      Protocol
GigabitEthernet0/0       192.168.10.1    YES manual up          up
GigabitEthernet0/1       10.0.12.1       YES manual up          up
GigabitEthernet0/2       unassigned      YES unset  administratively down down
Loopback0                1.1.1.1        YES manual up          up
R1# show running-config interface GigabitEthernet0/0
interface GigabitEthernet0/0
 ip address 192.168.10.1 255.255.255.0
```

`show ip interface brief` lists addresses but not masks; use `show running-config interface` or `show ip interface` to see the mask.

SCENARIO · DECIDE · IS IT PRIVATE?

A new engineer plans to number a branch LAN as `172.32.10.0/24`, believing all 172.x addresses are private.

- ☐ A Approve it — every 172.x.x.x address is private
- ☐ B Use 169.254.10.0/24 instead
- ☐ C Use 203.0.113.0/24 instead
- ☐ D Change it to a subnet inside 172.16.0.0/12, such as 172.20.10.0/24

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

- What is 11110000 in decimal?
 - ☐ A 224
 - ☐ B 248
 - ☐ C 240
 - ☐ D 192
- Which address is in RFC 1918 private space?
 - ☐ A 172.31.200.5
 - ☐ B 172.15.200.5
 - ☐ C 172.32.0.1
 - ☐ D 192.169.1.1
- A host shows 169.254.33.7/16. What most likely happened?
 - ☐ A It received a public address from the ISP
 - ☐ B It is using carrier-grade NAT space
 - ☐ C It was set to a documentation address
 - ☐ D It failed to get an address from DHCP
- How many usable host addresses does a /23 provide?
 - ☐ A 512
 - ☐ B 254
 - ☐ C 510
 - ☐ D 1022
- Why did CIDR replace classful addressing?
 - ☐ A Classful masks could not be written in binary or dotted decimal
 - ☐ B Fixed class sizes wasted addresses and table space
 - ☐ C Class C networks could not carry TCP traffic
 - ☐ D Routers could not forward multicast traffic
- Which range should appear in a training diagram so it can never clash with a real network?
 - ☐ A 192.168.100.0/24
 - ☐ B 198.51.100.0/24
 - ☐ C 100.64.0.0/10
 - ☐ D 11.0.0.0/8

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 2 · ADDRESS · LESSON 2 · 32 MIN

Subnetting Step by Step

The block-size method with fully worked examples

- Find the network, broadcast and usable range for any IPv4 address and prefix
- Choose a prefix that meets a host or subnet requirement
- Split a network into equal subnets and list their ranges

Subnetting means borrowing host bits to create more, smaller networks. It is the single most tested and most used skill in addressing. You do not need long binary conversions every time — the **block-size method** below lets you work most problems in your head in under a minute.

The three formulas you need

MEASURE	FORMULA	WHAT IT TELLS YOU
Subnets created	$\text{subnets} = 2^s$	s = bits borrowed. Borrowing 2 bits from a /24 gives $2^2 = 4$ subnets of /26.
Usable hosts	$\text{hosts} = 2^h - 2$	h = host bits left. A /26 leaves 6 host bits: $64 - 2 = 62$ hosts.
Block size	$\text{block} = 256 - \text{mask value in the interesting octet}$	For 255.255.255.224 the block is $256 - 224 = 32$: subnets start at 0, 32, 64, 96...

Prefix cheat sheet (last octet)

PREFIX	MASK	BLOCK SIZE	USABLE HOSTS
/24	255.255.255.0	256	254
/25	255.255.255.128	128	126
/26	255.255.255.192	64	62
/27	255.255.255.224	32	30
/28	255.255.255.240	16	14
/29	255.255.255.248	8	6
/30	255.255.255.252	4	2
/31	255.255.255.254	2	2 (point-to-point only)

The block-size method

1. **Find the interesting octet** The octet where the mask is neither 255 nor 0. For /27 that is the fourth octet; for /20 it is the third.
2. **Find the block size** 256 minus the mask value in that octet.
3. **Find the network address** The highest multiple of the block size that is less than or equal to the address's value in that octet. Octets to the right become 0.
4. **Find the broadcast address** Next network address minus 1 (octets to the right become 255).
5. **List the usable range** Network + 1 up to broadcast – 1.

EXAMPLE · WORKED EXAMPLE 1 – 192.168.50.77/27

Interesting octet: 4th. Mask 255.255.255.224 → block **32**. Multiples of 32: 0, 32, **64**, 96. 77 falls in the 64 block.

Network 192.168.50.64 · **Broadcast** 192.168.50.95 (96 – 1) · **Usable** 192.168.50.65 – 192.168.50.94 · **30 hosts**.

EXAMPLE · WORKED EXAMPLE 2 – 172.16.45.200/20

/20 = 255.255.240.0, so the interesting octet is the 3rd and the block is 256 – 240 = **16**. Multiples of 16: 0, 16, **32**, 48.

45 falls in the 32 block. **Network** 172.16.32.0 · **Broadcast** 172.16.47.255 · **Usable** 172.16.32.1 – 172.16.47.254 ·

$2^{12} - 2 = 4,094$ hosts.

Splitting 192.168.10.0/24 into four /26 subnets

SUBNET	NETWORK	FIRST HOST	LAST HOST	BROADCAST
1	192.168.10.0/26	192.168.10.1	192.168.10.62	192.168.10.63
2	192.168.10.64/26	192.168.10.65	192.168.10.126	192.168.10.127
3	192.168.10.128/26	192.168.10.129	192.168.10.190	192.168.10.191
4	192.168.10.192/26	192.168.10.193	192.168.10.254	192.168.10.255

PRACTICE Practise: check the worked examples

Enter 192.168.50.77/27 and then 172.16.45.200/20. Confirm the network, broadcast and host count match the worked examples, then try 10.9.130.14/18 on your own before checking.

ADDRESS / PREFIX	NETWORK	MASK	FIRST HOST	LAST HOST	BROADCAST	HOSTS
192.168.10.77/26						
10.45.130.9/20						
172.20.5.200/27						
10.0.0.129/25						
192.168.1.5/30						
172.16.99.1/22						

Answers: use the subnet calculator on the course website, or check with the answer key at the end of the book.

From a requirement to a subnet**01 Hosts needed**

e.g. 50 devices + growth

→

02 Host bits

smallest h where $2^h - 2 \geq \text{need} \rightarrow h = 6$

→

03 Prefix

$32 - 6 = /26$

→

04 Block size

64

→

05 Ranges

write network, gateway, usable range and broadcast

WATCH OUT · THE FOUR CLASSIC SUBNETTING MISTAKES

1) Forgetting the –2 for network and broadcast. 2) Choosing a network address that is not a multiple of the block size (192.168.10.40/27 is **not** a network – .32 is). 3) Assigning the network or broadcast address to a host. 4) Using the right block size in the wrong octet on prefixes shorter than /24.

Applying subnets to router interfaces

R1 · IOS-style CLI

```
R1(config)# interface GigabitEthernet0/0
R1(config-if)# ip address 192.168.10.1 255.255.255.192
R1(config-if)# no shutdown
R1(config-if)# interface GigabitEthernet0/1
R1(config-if)# ip address 192.168.10.65 255.255.255.192
R1(config-if)# no shutdown
R1(config-if)# interface GigabitEthernet0/2
R1(config-if)# ip address 192.168.10.129 255.255.255.192
R1(config-if)# no shutdown
R1(config-if)# end
R1# show ip route connected
C      192.168.10.0/26 is directly connected, GigabitEthernet0/0
C      192.168.10.64/26 is directly connected, GigabitEthernet0/1
C      192.168.10.128/26 is directly connected, GigabitEthernet0/2
```

Each interface takes the first usable address of its /26 as the gateway for that subnet. The router refuses overlapping subnets on two interfaces — a useful safety net.

SCENARIO · DECIDE · CHOOSING A PREFIX

A new lab needs 50 PCs today and the lab manager expects 8 more next year. You must use the smallest subnet that fits.

- ☐ A /27
- ☐ B /26
- ☐ C /25
- ☐ D /28

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

7. What is the network address of 10.1.1.100/26?
 - ☐ A 10.1.1.96
 - ☐ B 10.1.1.0
 - ☐ C 10.1.1.100
 - ☐ D 10.1.1.64
8. What is the broadcast address of 192.168.7.130/28?
 - ☐ A 192.168.7.143
 - ☐ B 192.168.7.255
 - ☐ C 192.168.7.135
 - ☐ D 192.168.7.159
9. Which prefix is the smallest that supports 25 hosts?
 - ☐ A /28
 - ☐ B /26
 - ☐ C /27
 - ☐ D /29
10. How many /27 subnets fit inside one /24?
 - ☐ A 4
 - ☐ B 16
 - ☐ C 8
 - ☐ D 32
11. A PC is configured as 172.16.40.0/22 with gateway 172.16.40.1. Why does it fail?
 - ☐ A 172.16.40.0 is the broadcast address of its /22
 - ☐ B 172.16.40.0 is the network address of its /22
 - ☐ C The gateway must be the last usable address
 - ☐ D A /22 mask is not allowed on a host
12. For 172.16.45.200/20, what is the last usable host?
 - ☐ A 172.16.47.254
 - ☐ B 172.16.45.254
 - ☐ C 172.16.63.254
 - ☐ D 172.16.47.255

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 2 · ADDRESS · LESSON 3 · 30 MIN

VLSM Design and Verifying Host Configuration

Right-sized subnets for every segment — then proving each host is set up correctly

- Design a VLSM plan that allocates the largest subnets first without overlap
- Document an addressing plan that others can follow
- Verify a host's IP, mask, gateway and DNS settings and spot common errors

Real sites never need equal-sized subnets. A staff VLAN may need 100 addresses while a router-to-router link needs two. **Variable-length subnet masking (VLSM)** lets you cut one address block into differently sized subnets so nothing is wasted and the plan can still be summarised as one route.

The VLSM procedure

1. **List requirements** Every segment with its host count, including growth, gateways and infrastructure devices.
2. **Sort largest first** Big subnets have big block sizes; placing them first keeps every subnet on a valid boundary.
3. **Allocate from the start of the block** Give the largest segment the first network address; the next subnet starts right after its broadcast.
4. **Check boundaries** Every network address must be a multiple of its own block size.
5. **Record spare space** Note what is left so future growth comes from the same block.

Worked VLSM: a university lab block 192.168.40.0/24

SEGMENT	HOSTS NEEDED	PREFIX (USABLE)	SUBNET	USABLE RANGE	BROADCAST
Lab A	100	/25 (126)	192.168.40.0/25	.1 – .126	.127
Lab B	50	/26 (62)	192.168.40.128/26	.129 – .190	.191
Staff	20	/27 (30)	192.168.40.192/27	.193 – .222	.223
Printers	10	/28 (14)	192.168.40.224/28	.225 – .238	.239
Router link	2	/30 (2)	192.168.40.240/30	.241 – .242	.243
Spare	—	/30 + /29	192.168.40.244/30, 192.168.40.248/29	—	—

PRACTICE Check every VLSM subnet

Enter each subnet from the lab table (for example 192.168.40.192/27) and confirm the usable range and broadcast. Then enter 192.168.40.200/28 — the calculator reports the network as 192.168.40.192, proving .200 is not a valid /28 boundary.

ADDRESS / PREFIX	NETWORK	MASK	FIRST HOST	LAST HOST	BROADCAST	HOSTS
192.168.10.77/26						
10.45.130.9/20						
172.20.5.200/27						
10.0.0.129/25						
192.168.1.5/30						
172.16.99.1/22						

Answers: use the subnet calculator on the course website, or check with the answer key at the end of the book.

WATCH OUT · OVERLAP IS THE SILENT KILLER

If 192.168.40.128/25 and 192.168.40.192/27 both exist, the /27 sits inside the /25. Routers will prefer the longer prefix for some addresses and the shorter for others, producing "half the hosts work" faults. Always check new subnets against the plan before configuring them.

CASE STUDY · A CHATTOGRAM WAREHOUSE ADDRESSING PLAN

A logistics company is opening a bonded warehouse near Chattogram port. The IT team has been given **172.20.8.0/22** for the site. Requirements: handheld scanners and floor Wi-Fi 300 hosts; office staff 100; CCTV cameras 50; IP phones 25; switch and AP management 12; two router links to two WAN providers.

1. Which prefix does each segment need?

2. Allocate the subnets in VLSM order and give each network address.

3. How much of the /22 is left for growth, and in what blocks?

Model answer: module answer key.

TEMPLATE IPv4 addressing plan

One row per subnet. Keep it with the site documentation and update it before every change.

VLAN / SEGMENT	PURPOSE	SUBNET / PREFIX	GATEWAY	DHCP RANGE	RESERVED / STATIC	NOTES
10	Office staff	172.20.10.0/25	172.20.10.1	.20 – .120	.2 – .19 printers, servers	Growth to 110 users expected

Once a plan exists, every host must match it. Most "network" tickets on day one turn out to be one wrong field on a host: a mistyped mask, a gateway from another subnet, or a stale DNS server. Verification means checking each field against the plan, not just checking that *something* is filled in.

Verifying a Linux host against the plan

Linux host

```
$ ip -4 addr show eth0
    inet 172.20.10.57/25 brd 172.20.10.127 scope global dynamic eth0
$ ip route
default via 172.20.10.1 dev eth0 proto dhcp metric 100
172.20.10.0/25 dev eth0 proto kernel scope link src 172.20.10.57
$ resolvectl status eth0 | grep 'DNS Servers'
    DNS Servers: 172.20.10.10
$ ping -c 2 172.20.10.1
2 packets transmitted, 2 received, 0% packet loss
```

Address, prefix and broadcast match the plan for VLAN 10; the gateway is inside the subnet; DNS is set; the gateway answers.

What "correct" looks like on a host

- The IP address is a **usable** address in the planned subnet (not network, not broadcast).
- The mask/prefix **matches the plan** exactly — the same on every host in the subnet.
- The default gateway is **inside the same subnet** and is the router's real interface address.
- DNS servers are reachable and are the ones in the plan.
- No `169.254.x.x` address — that means DHCP failed.
- No duplicate address warnings in the OS logs.

SCENARIO · DECIDE · HALF THE SUBNET WORKS

PCs in VLAN 20 should be 192.168.40.128/26 with gateway 192.168.40.129. PC-7 is set to 192.168.40.170 with mask 255.255.255.240 (/28). It can reach PCs at .161–.174 but not the gateway or anything else.

- ☐ A Correct PC-7's mask to 255.255.255.192 (/26) to match the plan
- ☐ B Change PC-7's gateway to 192.168.40.161 so it is inside its /28
- ☐ C Replace PC-7's network card
- ☐ D Move PC-7 to VLAN 10

Circle your choice. Answer and feedback: module answer key.

Host configuration verification checklist

- ☐ IP address is inside the planned subnet and is not the network or broadcast address
- ☐ Prefix/mask matches the addressing plan
- ☐ Default gateway is in the same subnet and answers ping
- ☐ ARP/neighbour table shows the gateway's MAC address
- ☐ DNS servers match the plan and resolve an internal and an external name
- ☐ Address source is as planned (DHCP lease or documented static)
- ☐ No 169.254.x.x or duplicate-address warnings
- ☐ Trace to a remote subnet leaves through the expected gateway

KNOWLEDGE CHECK

13. When designing VLSM, why allocate the largest subnet first?

- ☐ A Routers require subnets in descending order
- ☐ B Small subnets cannot use private addresses
- ☐ C It reduces the number of broadcast addresses
- ☐ D It keeps every subnet aligned to its block boundary

14. Which pair of subnets overlaps?

- ☐ A 10.5.0.0/24 and 10.5.1.0/24
- ☐ B 10.5.2.0/25 and 10.5.2.128/25
- ☐ C 10.5.0.0/23 and 10.5.1.0/24
- ☐ D 10.5.4.0/23 and 10.5.6.0/23

15. A host has 192.168.40.200/27 and gateway 192.168.40.129. What is wrong?

- ☐ A The gateway is outside the host's subnet
- ☐ B The host address is the broadcast address of its /27
- ☐ C A /27 cannot be used on hosts
- ☐ D The gateway must end in .254

16. A Windows PC shows 169.254.12.9. What should you check first?

- ☐ A Whether the DNS server name is spelled correctly
- ☐ B Whether DHCP is reachable from the PC's VLAN
- ☐ C Whether the default route exists on the core
- ☐ D Whether the PC's browser proxy is configured

17. A site was given 10.30.0.0/22. Which subnet can be part of it?

- ☐ A 10.30.4.0/24
- ☐ B 10.30.0.0/21
- ☐ C 10.31.0.0/24
- ☐ D 10.30.3.0/24

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 2 · ADDRESS · LESSON 4 · 26 MIN

IPv6 Fundamentals: Format, Compression and Address Types

Reading, writing and recognising 128-bit addresses

- Write IPv6 addresses in full and compressed form correctly
- Identify global unicast, unique local, link-local, multicast, loopback and unspecified addresses
- Explain the /48 – /64 prefix structure used by most organisations

IPv4's 4.3 billion addresses ran out at the regional registries years ago; NAT stretched them further but made networks more complex. **IPv6** uses **128-bit** addresses — enough for every device to have a globally unique address — and drops broadcast entirely. Large mobile networks, content providers and clouds already carry much of their traffic over IPv6, so every junior engineer must read it fluently.

DEFINITION · IPV6 ADDRESS FORMAT

128 bits written as **eight groups of four hexadecimal digits** (hextets) separated by colons, e.g.

2001:0db8:0acd:0010:0000:0000:0000:0001 . A **prefix length** such as /64 marks the network part, just like CIDR in IPv4.

Compressing an IPv6 address

1. **Drop leading zeros** In each hextet, remove zeros at the front: 0db8 → db8 , 0010 → 10 , 0000 → 0 .
2. **Replace one run of zero hextets with ::** The longest run of consecutive all-zero hextets becomes :: .
3. **Only once** :: may appear only once — otherwise the reader cannot tell how many zeros each one hides.
4. **Tie-break and style** If two runs are equal, compress the first (leftmost). Do not use :: for a single zero hextet; write letters in lowercase.

EXAMPLE · COMPRESSION IN PRACTICE

2001:0db8:0acd:0010:0000:0000:0000:0001 → 2001:db8:acd:10::1 .

2001:0db8:0000:0000:00a0:0000:0000:0001 has two equal runs of two zero hextets, so the first is compressed → 2001:db8::a0:0:0:1 .

fe80:0000:0000:0000:0000:0000:0000:0001 → fe80::1 .

IPv6 address types

TYPE	PREFIX	SCOPE AND USE
Global unicast (GUA)	2000::/3	Publicly routable; most begin with 2 or 3
Unique local (ULA)	fc00::/7 (fd00::/8 in practice)	Private, inside an organisation; similar role to RFC 1918
Link-local	fe80::/10	Only on one link; every IPv6 interface has one; used by routing protocols and as next hop
Multicast	ff00::/8	One-to-many; ff02::1 all nodes, ff02::2 all routers
Loopback	::1/128	The host itself
Unspecified	::/128	"No address yet" — source during address set-up
Documentation	2001:db8::/32	Examples and training only

A typical global unicast address (left to right, top first)

Global routing prefix · 48 bits	Assigned by a provider or registry	Example: 2001:db8:acad::/48
Subnet ID · 16 bits	Chosen by the organisation	65,536 possible /64 subnets per site
Interface ID · 64 bits	Identifies the interface on the subnet	Random, stable-private, EUI-64 or manual

IPv6 sizing

MEASURE	FORMULA	WHAT IT TELLS YOU
Subnets in a /48	$2^{(64 - 48)} = 2^{16} = 65,536$	A single site allocation holds 65,536 standard /64 subnets.
Addresses in a /64	$2^{64} \approx 1.8 \times 10^{19}$	Never "run out" on a LAN; /64 is the standard LAN size and is required for SLAAC.
Subnets in a /56	$2^{(64 - 56)} = 256$	A common home or small-office delegation from an ISP.

IPv4 vs IPv6 at a glance

IPv4

- 32-bit, dotted decimal
- Broadcast exists
- ARP maps IP to MAC
- Header 20–60 bytes with checksum
- DHCP for automatic addressing
- NAT is everywhere

IPv6

- 128-bit, hexadecimal with colons
- No broadcast — multicast instead
- Neighbor Discovery (ICMPv6) replaces ARP
- Fixed 40-byte header, no header checksum
- SLAAC and/or DHCPv6
- Global addresses end to end; firewalls still required

INSIGHT · EVERY INTERFACE HAS SEVERAL IPV6 ADDRESSES

Expect a link-local address **plus** one or more global (or unique local) addresses, and often temporary privacy addresses too. Seeing three IPv6 addresses on one laptop is normal, not a misconfiguration.

Reading IPv6 addresses on a router and a host

R1 · IOS-style CLI

```
R1# show ipv6 interface brief
GigabitEthernet0/0      [up/up]
    FE80::1
    2001:DB8:ACAD:10::1
GigabitEthernet0/1      [up/up]
    FE80::21A:2BFF:FE3C:4D5E
    2001:DB8:ACAD:20:21A:2BFF:FE3C:4D5E
GigabitEthernet0/2      [administratively down/down]
    unassigned
```

IOS-style devices print IPv6 in uppercase; hosts usually print lowercase. Both are the same address.

SCENARIO · DECIDE · SPOT THE INVALID ADDRESS

A colleague's addressing sheet lists four server addresses. Which one must be fixed before it is configured?

- ☐ A 2001:db8:acad:10::25
- ☐ B fd12:3456:789a:1::25
- ☐ C 2001:db8::10::25
- ☐ D 2001:db8:0:10:0:0:0:25

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

18. What is the correct compressed form of 2001:0db8:0000:0042:0000:0000:0000:0001?

- ☐ A 2001:db8::42::1
- ☐ B 2001:db8:0:42:0::1
- ☐ C 2001:db8:0:42::1
- ☐ D 2001:db8::42:0:0:0:1

19. Which prefix identifies link-local IPv6 addresses?

- ☐ A ff00::/8
- ☐ B fc00::/7
- ☐ C 2001:db8::/32
- ☐ D fe80::/10

20. What replaces IPv4 broadcast in IPv6?

- ☐ A Anycast to the nearest router on the link
- ☐ B Multicast groups such as ff02::1
- ☐ C Unicast to every host in turn
- ☐ D Broadcast to ffff::/16

21. An ISP delegates a /48 to a campus. How many /64 subnets can the campus create?

- ☐ A 65,536
- ☐ B 256
- ☐ C 4,096
- ☐ D 16,384

22. Which address is a unique local address?

- ☐ A fe80::abcd:12
- ☐ B 2001:db8:acad:12::10
- ☐ C ff02::2
- ☐ D fd00:abcd:12:5::10

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 2 · ADDRESS · LESSON 5 · 30 MIN

IPv6 Configuration: EUI-64, SLAAC, DHCPv6 and Dual Stack

How IPv6 hosts get addresses — and how to run IPv4 and IPv6 side by side

- Build an EUI-64 interface ID from a MAC address
- Compare SLAAC, stateless DHCPv6 and stateful DHCPv6 and the RA flags that select them
- Configure and verify IPv6 addressing on a router and explain dual stack

IPv6 was designed so hosts can configure themselves. A router announces the subnet prefix; hosts build their own addresses and check nobody else uses them. The machinery behind this is **Neighbor Discovery Protocol (NDP)**, which runs over ICMPv6 and also replaces ARP.

Neighbor Discovery messages

ICMPV6 TYPE	MESSAGE	PURPOSE
133	Router Solicitation (RS)	Host asks routers to announce themselves (sent to ff02::2)
134	Router Advertisement (RA)	Router announces prefix, flags and itself as default gateway
135	Neighbor Solicitation (NS)	"Who has this IPv6 address?" — address resolution and duplicate address detection
136	Neighbor Advertisement (NA)	Reply with the MAC address
137	Redirect	Tells a host a better first hop exists

Building an EUI-64 interface ID from MAC 00:1A:2B:3C:4D:5E

1. **Split the MAC in half** 001A2B | 3C4D5E
2. **Insert FFFE in the middle** 001A2B FFFE 3C4D5E → 64 bits.
3. **Flip the 7th bit of the first byte** 00 = 00000000 → 00000010 = 02 . The interface ID is 021A:2BFF:FE3C:4D5E .
4. **Add the /64 prefix** With prefix 2001:db8:acad:20::/64 the address is 2001:db8:acad:20:21a:2bff:fe3c:4d5e .

INSIGHT · EUI-64 IS NOW MOSTLY A ROUTER THING

Because EUI-64 embeds the hardware MAC, a laptop using it could be tracked across networks. Modern operating systems therefore create **random or stable-private interface IDs** by default. You will still see EUI-64 on routers and some IoT devices — and in exams.

Three ways a host gets IPv6 settings

SLAAC only

- RA carries prefix with A flag set
- Host builds its own address
- DNS server can be sent in the RA (RDNSS option)
- Default gateway = router's link-local address


Stateless DHCPv6 (SLAAC + DHCPv6)

- RA: A flag set, O flag (other config) set
- Address from SLAAC
- DNS and domain from DHCPv6
- DHCPv6 server keeps no address leases

Stateful DHCPv6

- RA: M flag (managed) set
- Address assigned and recorded by DHCPv6
- Gateway still learned from the RA
- Some mobile operating systems do not support it — plan carefully

SLAAC from power-on to first packet



Configuring IPv6 on a router

R1 · IOS-style CLI

```

R1(config)# ipv6 unicast-routing
R1(config)# interface GigabitEthernet0/0
R1(config-if)# ipv6 address 2001:db8:acad:10::1/64
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# no shutdown
R1(config-if)# interface GigabitEthernet0/1
R1(config-if)# ipv6 address 2001:db8:acad:20::/64 eui-64
R1(config-if)# ipv6 nd other-config-flag
R1(config-if)# no shutdown
  
```

`ipv6 unicast-routing` makes the device forward IPv6 and send RAs. A manual link-local such as fe80::1 makes the gateway easy to recognise. `other-config-flag` sets the O flag so hosts on Gi0/1 ask DHCPv6 for DNS (stateless DHCPv6).

Verifying from a Linux host

Linux host

```

$ ip -6 addr show eth0
    inet6 2001:db8:acad:10:9c4e:51ff:fe2a:7b11/64 scope global dynamic mngtmpaddr
    inet6 fe80::9c4e:51ff:fe2a:7b11/64 scope link
$ ip -6 route | grep default
default via fe80::1 dev eth0 proto ra metric 100
$ ping -c 2 2001:db8:acad:10::1
2 packets transmitted, 2 received, 0% packet loss
  
```

`proto ra` proves the default route came from a Router Advertisement; the next hop is the router's link-local fe80::1.

Dual stack means every device runs IPv4 and IPv6 at the same time, each with its own addresses, routes and policies. It is the most common way organisations adopt IPv6: nothing breaks for IPv4-only services, and IPv6 traffic grows as applications and providers support it. Modern hosts try IPv6 first and fall back to IPv4 within a fraction of a second if IPv6 is slow (the "happy eyeballs" approach).

Dual-stack rules for the real world

- Plan IPv6 per VLAN the same way you plan IPv4: one /64 per VLAN, recorded in the addressing plan.
- Apply equivalent **security policy** to both stacks — an open IPv6 path around an IPv4 firewall is a classic hole.
- Monitor both: a broken IPv6 path can slow every website even when IPv4 works.
- Where IPv4 is scarce, IPv6-only networks with NAT64/DNS64 translation are increasingly common (especially on mobile).

SCENARIO · DECIDE · NO IPV6 ADDRESSES ON THE LAN

After configuring `ipv6 address 2001:db8:acad:10::1/64` on R1's LAN interface, hosts only have fe80:: addresses. The interface is up/up and hosts can ping fe80::1.

- ☐ A Change the prefix to /48 so it is larger
- ☐ B Configure a DHCPv4 pool for the LAN
- ☐ C Remove the link-local address
- ☐ D Enable `ipv6 unicast-routing` globally so R1 forwards IPv6 and sends Router Advertisements

Circle your choice. Answer and feedback: module answer key.

REFLECT

Your organisation still runs IPv4 only. Write three sentences for your manager explaining one business risk of staying IPv4-only and one practical first step toward dual stack.

Key terms**CIDR**

Classless Inter-Domain Routing — explicit prefix lengths of any size instead of fixed classes.

RFC 1918

The private IPv4 ranges 10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16.

VLSM

Variable-length subnet masking — different prefix lengths inside one address block.

Link-local (IPv6)

fe80::/10 address valid on one link only; every IPv6 interface has one.

Unique local

Private IPv6 range fc00::/7, used as fd00::/8.

SLAAC

Stateless address autoconfiguration — hosts build IPv6 addresses from RA prefixes.

Dual stack

Running IPv4 and IPv6 at the same time on the same devices and links.

Prefix length

The number of leading network bits, written /n (e.g. /26, /64).

Block size

256 minus the mask value in the interesting octet; the gap between subnet network addresses.

Broadcast address

The address with all host bits set to 1; reaches every host on the IPv4 subnet.

Global unicast

Publicly routable IPv6 address, currently allocated from 2000::/3.

EUI-64

Method of building a 64-bit interface ID from a MAC by inserting FFFE and flipping the 7th bit.

Router Advertisement

ICMPv6 type 134 message announcing prefixes, flags and a default gateway.

KNOWLEDGE CHECK

23. What is the EUI-64 interface ID for MAC 00:50:56:AB:CD:EF?

- ☐ A 0050:56FF:FEAB:CDEF
- ☐ B 0250:56FF:FEAB:CDEF
- ☐ C 0250:56FE:FFAB:CDEF
- ☐ D 0050:56AB:FFFE:CDEF

24. Which RA flag tells hosts to obtain their **address** from a DHCPv6 server?
- ☐ A O (other configuration)
 - ☐ B A (autonomous) in the prefix option
 - ☐ C M (managed address configuration)
 - ☐ D L (on-link) in the prefix option
25. Where does an IPv6 host normally learn its default gateway?
- ☐ A From a Router Advertisement
 - ☐ B From a stateful DHCPv6 lease
 - ☐ C From an ARP reply
 - ☐ D From the DNS server
26. Which NDP message does a host send to check that its new address is not already in use?
- ☐ A Router Solicitation
 - ☐ B Neighbor Solicitation
 - ☐ C Router Advertisement
 - ☐ D Neighbor Advertisement
27. A dual-stack firewall has strict IPv4 rules but allows all IPv6. What is the risk?
- ☐ A IPv4 traffic will automatically be dropped
 - ☐ B Hosts will stop receiving IPv6 addresses
 - ☐ C DNS will stop resolving IPv4 names
 - ☐ D Attackers can reach hosts over the open IPv6 path
28. Why should LAN subnets in IPv6 normally be /64?
- ☐ A Routers cannot route any prefix longer than /64
 - ☐ B IPv6 hosts reject addresses with a /48 prefix
 - ☐ C SLAAC and most host features expect a 64-bit interface ID
 - ☐ D A /64 allows exactly 254 hosts per VLAN

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 2 · ADDRESS

Answer key

Knowledge checks

- 1. C** $128 + 64 + 32 + 16 = 240$. 224 is 11100000, 248 is 11111000 and 192 is 11000000.
- 2. A** The 172 private block is 172.16.0.0 – 172.31.255.255, so 172.31.200.5 is private. 172.15.x, 172.32.x and 192.169.x are all public space.
- 3. D** Hosts self-assign 169.254.0.0/16 when DHCP does not answer. ISP public addresses, CGN space (100.64.0.0/10) and documentation ranges look different.
- 4. C** A /23 has 9 host bits: $2^9 = 512$ addresses, minus network and broadcast = 510. 512 forgets the two reserved addresses, 254 is a /24 and 1022 is a /22.
- 5. B** CIDR allows any prefix length, saving addresses and enabling summarisation. Masks were always binary, Class C carried TCP normally, and multicast is unrelated.
- 6. B** 198.51.100.0/24 is reserved for documentation. 192.168.100.0/24 is private space that real networks use, 11.0.0.0/8 is real public space, and 100.64.0.0/10 is used by ISPs for carrier-grade NAT.
- 7. D** Block size 64: 0, 64, 128. 100 falls in the 64 block. 10.1.1.96 would be the answer for /27, 10.1.1.0 for /24 or /25, and 10.1.1.100 is a host address.
- 8. A** Block 16: the 128 block covers .128–.143, so the broadcast is .143. .255 would be /24, .135 is /29 and .159 is /27.
- 9. C** A /27 gives 30 usable hosts. /28 gives 14 and /29 gives 6 (both too small); /26 works but is larger than needed.
- 10. C** Borrowing 3 bits ($27 - 24$) gives $2^3 = 8$ subnets. 4 would be /26, 16 would be /28, and 32 is the block size, not the count.
- 11. B** For /22 the third-octet block is 4, so the subnet is 172.16.40.0 – 172.16.43.255; .40.0 is the network address and cannot be a host. The broadcast is 172.16.43.255, gateways may be any usable address, and /22 is a valid host mask.
- 12. A** Block 16 in the third octet: the subnet runs 172.16.32.0 – 172.16.47.255, so the last host is .47.254. .45.254 ignores the /20, .63.254 uses the wrong block, and .47.255 is the broadcast.
- 13. D** Placing large blocks first avoids gaps and misaligned networks. Routers do not care about configuration order, small subnets can be private, and every subnet has one broadcast regardless.
- 14. C** 10.5.0.0/23 covers 10.5.0.0 – 10.5.1.255, which contains 10.5.1.0/24. The other pairs are adjacent but separate.
- 15. A** 192.168.40.200/27 lives in .192 – .223, so .129 is not on the local subnet and cannot be reached by ARP. .200 is a normal host address, /27 is valid and gateways need not end in .254.
- 16. B** A 169.254 address means no DHCP offer arrived, so check the port's VLAN, the DHCP server or relay. DNS, core routes and proxies matter only after the PC has an address.
- 17. D** 10.30.0.0/22 spans 10.30.0.0 – 10.30.3.255, so 10.30.3.0/24 fits. 10.30.4.0 is outside, a /21 is larger than the /22, and 10.31.x is a different block.
- 18. C** Leading zeros drop and the longest zero run (three hextets) becomes ::. Two :: is invalid, the third option does not compress the longest run fully, and the fourth compresses the shorter run instead of the longest.
- 19. D** Link-local addresses use fe80::/10. ff00::/8 is multicast, fc00::/7 is unique local and 2001:db8::/32 is documentation space.
- 20. B** IPv6 has no broadcast; all-nodes multicast ff02::1 and other groups serve that role. Anycast targets the nearest of several interfaces, sequential unicast is not used, and ffff::/16 is not a broadcast range.
- 21. A** $64 - 48 = 16$ bits of subnet ID: $2^{16} = 65,536$. 256 would be a /56, 4,096 a /52 and 16,384 a /50.
- 22. D** fd00::/8 sits inside fc00::/7, the unique local range. fe80:: is link-local, 2001:db8:: is documentation global space and ff02::2 is the all-routers multicast group.
- 23. B** Insert FFFE between 005056 and ABCDEF, then flip the 7th bit (00 → 02). The second option forgets the flip, the third reverses FFFE, and the fourth inserts FFFE in the wrong place.
- 24. C** The M flag signals stateful DHCPv6 addressing. O signals DHCPv6 for other settings like DNS, A permits SLAAC, and L says the prefix is on-link.
- 25. A** IPv6 default gateways come from RAs, even when DHCPv6 is used. DHCPv6 does not carry a gateway option, ARP is IPv4 only, and DNS resolves names.
- 26. B** Duplicate address detection uses a Neighbor Solicitation for the tentative address. RS asks for routers, RA is sent by routers, and a Neighbor Advertisement is the reply to a solicitation.
- 27. D** Hosts prefer IPv6 and are reachable on it, so an unprotected IPv6 path bypasses the IPv4 policy. The other effects do not follow from an open IPv6 rule.
- 28. C** SLAAC builds 64-bit interface IDs, so the LAN prefix must be /64. Routers can route longer prefixes (e.g. /127 links), hosts accept other prefixes manually, and a /64 holds far more than 254 hosts.

Scenarios & case studies

Scenario · Is it private?

D is best. Correct. That keeps the branch in RFC 1918 space and avoids clashing with real public addresses.

Scenario · Choosing a prefix

B is best. Correct. 62 usable hosts covers 58 devices plus the gateway with a little room left.

Case study · A Chattogram warehouse addressing plan

Largest first: **Scanners/Wi-Fi** /23 (510) → 172.20.8.0/23 (8.1 – 9.254). **Office** /25 (126) → 172.20.10.0/25. **CCTV** /26 (62) → 172.20.10.128/26. **Voice** /27 (30) → 172.20.10.192/27. **Management** /28 (14) → 172.20.10.224/28. **WAN 1** /30 → 172.20.10.240/30 and **WAN 2** /30 → 172.20.10.244/30 (or /31s on routers that support them). Spare: 172.20.10.248/29 and the whole of **172.20.11.0/24** — enough to add another /24 VLAN later while the site is still summarised as one 172.20.8.0/22 route.

Scenario · Half the subnet works

A is best. Correct. With /28 the PC believes its subnet is .160–.175, so the gateway at .129 looks remote and unreachable. The planned /26 puts .129 back on-link.

Scenario · Spot the invalid address

C is best. Correct — this one is invalid. Two `::` make the address ambiguous; the number of hidden zeros cannot be worked out.

Scenario · No IPv6 addresses on the LAN

D is best. Correct. Without it the device acts like a host: no RAs, so clients have no prefix for SLAAC.

03

MODULE 3 · SWITCH

Switching & VLANs

MAC learning, VLANs, trunks, inter-VLAN routing, spanning tree, EtherChannel and LLDP

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Explain how switches learn, forward, filter and flood frames and trace a host to its port
- Design VLANs and configure access ports and 802.1Q trunks safely
- Implement inter-VLAN routing with router-on-a-stick or layer 3 switch SVIs
- Control spanning tree root placement and protect edge ports with PortFast and BPDU guard
- Build and verify LACP EtherChannels and document neighbours with LLDP

LESSONS

1. How Switches Learn and Forward Frames · 24 min
2. VLANs, Access Ports and 802.1Q Trunks · 30 min
3. Inter-VLAN Routing: Router-on-a-Stick and SVIs · 28 min
4. Spanning Tree: Loop Prevention with STP and RSTP · 32 min
5. EtherChannel with LACP and Neighbour Discovery with LLDP · 26 min

MODULE 3 · SWITCH · LESSON 1 · 24 MIN


How Switches Learn and Forward Frames

MAC learning, forwarding, filtering and flooding — and how to find any device

- Explain how a switch builds and uses its MAC address table
- Predict whether a frame will be forwarded, filtered or flooded
- Distinguish collision domains from broadcast domains
- Locate a host's switch port from its IP address

A switch is a frame-forwarding machine. It knows nothing about IP subnets on its access ports; it only reads MAC addresses. Yet with one simple rule — **remember where each source address was seen** — it turns a pile of cables into an efficient LAN where traffic goes only where it needs to.

What a switch does with every frame



Forwarding decisions

DESTINATION MAC IS...	SWITCH ACTION	WHY
Known, on a different port	Forward out that one port	Normal unicast switching
Known, on the same port it arrived on	Filter (drop)	The receiver is already on that segment
Unknown unicast	Flood to all ports in the VLAN except the source	The switch must find the receiver
Broadcast (FF:FF:FF:FF:FF:FF)	Flood within the VLAN	Broadcasts are meant for everyone in the VLAN
Multicast	Flood within the VLAN unless multicast snooping limits it	Snooping learns which ports joined the group

DEFINITION · MAC ADDRESS TABLE

The **MAC address table** (often called the CAM table because of the fast memory it uses) maps **VLAN + MAC address** → **port**. Dynamic entries are learned from source addresses; static entries are configured by an engineer.

Reading the MAC address table

Switch SW1 · IOS-style CLI

```
SW1# show mac address-table dynamic
      Mac Address Table
```

```
-----
Vlan    Mac Address      Type      Ports
----    -
10      0050.56a1.1b2c   DYNAMIC   Gi1/0/1
10      0050.56a1.77e0   DYNAMIC   Gi1/0/2
20      6c4b.90de.0012   DYNAMIC   Gi1/0/14
10      0011.2233.4401   DYNAMIC   Gi1/0/48
20      0011.2233.4401   DYNAMIC   Gi1/0/48
Total Mac Addresses for this criterion: 5
```

The same router MAC appears in VLAN 10 and VLAN 20 on Gi1/0/48 — the uplink trunk. Many MACs on one port usually means that port leads to another switch, AP or router.

Hub vs switch vs router

Hub (legacy)

- Layer 1 repeater
- Every bit goes out every port
- One collision domain for all ports
- Half duplex only — not used today

Switch

- Layer 2 (some also route at Layer 3)
- Forwards by MAC address
- Each port is its own collision domain
- A VLAN is one broadcast domain

Router

- Layer 3
- Forwards by IP address and routing table
- Each interface is its own broadcast domain
- Does not forward broadcasts between networks

Collision vs broadcast domains

- A **collision domain** is where two frames can collide; full-duplex switch ports eliminate collisions in practice.
- A **broadcast domain** is everywhere a broadcast reaches — one VLAN.
- Switches separate collision domains; **routers and VLANs** separate broadcast domains.
- Huge broadcast domains (thousands of hosts in one VLAN) waste bandwidth and CPU on every host.

Walk-through: PC-A (VLAN 10) talks to PC-B for the first time

1. **ARP broadcast** PC-A sends an ARP request to FF:FF:FF:FF:FF:FF. The switch learns PC-A's MAC on Gi1/0/1 and floods the frame to every VLAN 10 port.
2. **ARP reply** PC-B replies by unicast to PC-A's MAC. The switch learns PC-B on Gi1/0/2 and forwards the reply only to Gi1/0/1.
3. **Data** From now on frames between A and B go port to port; no flooding is needed while the entries stay fresh.

INSIGHT · STORE-AND-FORWARD VS CUT-THROUGH

Most enterprise switches **store and forward**: they receive the whole frame and check the FCS before sending it on, so corrupted frames die at the first switch. Some data-centre switches use **cut-through** to start forwarding as soon as the destination MAC is read, saving microseconds but passing on bad frames.

Finding which port a host is on

R1 and SW1 · IOS-style CLI

```
R1# show ip arp 10.10.20.57
Protocol Address      Age (min)  Hardware Addr  Type   Interface
Internet 10.10.20.57          3    6c4b.90de.0012  ARPA   GigabitEthernet0/0.20
SW1# show mac address-table address 6c4b.90de.0012
Vlan    Mac Address      Type        Ports
----    -
20      6c4b.90de.0012   DYNAMIC     Gi1/0/14
```

IP → MAC from the gateway's ARP table, then MAC → port from the switch. If the port is an uplink, repeat on the next switch until you reach an access port.

SCENARIO · DECIDE · A SUSPICIOUS DEVICE

Security asks you to find the wall port used by IP 10.10.20.57. You have CLI access to the gateway router and every switch.

- ☐ A Read the gateway's ARP entry for the IP, then search for that MAC in the switch MAC tables until you reach an access port
- ☐ B Ping the IP and wait for someone to complain
- ☐ C Shut down every port in VLAN 20 one by one
- ☐ D Check the DNS server for the device name

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

- A switch receives a frame whose destination MAC is not in its table. What does it do?
 - ☐ A Floods it out all ports in that VLAN except the source
 - ☐ B Drops it and sends an error to the source host
 - ☐ C Sends it only to the uplink port toward the core
 - ☐ D Floods it out every port in every VLAN on the switch
- Which part of an incoming frame does a switch use to **populate** its MAC table?
 - ☐ A Destination MAC address
 - ☐ B Source IP address
 - ☐ C Source MAC address
 - ☐ D VLAN priority bits
- How many broadcast domains exist on one switch with VLANs 10, 20 and 30 in use?
 - ☐ A One
 - ☐ B Three
 - ☐ C As many as there are ports
 - ☐ D Two
- Twenty MAC addresses are learned on port Gi1/0/48. What is the most likely explanation?
 - ☐ A The port connects to another switch or AP
 - ☐ B The port has a duplex mismatch
 - ☐ C The MAC table has overflowed
 - ☐ D Twenty users are sharing one patch cable through a splitter
- Which device separates broadcast domains?
 - ☐ A A hub
 - ☐ B A layer 2 switch without VLANs
 - ☐ C A repeater
 - ☐ D A router

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 3 · SWITCH · LESSON 2 · 30 MIN

VLANs, Access Ports and 802.1Q Trunks

Segmenting the LAN and carrying many VLANs over one link

- Explain why VLANs are used and what a VLAN changes in the network
- Configure access ports and 802.1Q trunks with a safe native VLAN
- Verify VLAN membership and trunk status and fix allowed-VLAN errors

A **VLAN (virtual LAN)** splits one physical switch network into several logical networks. Devices in different VLANs cannot talk at Layer 2 even when they share a switch; to communicate they must go through a router or layer 3 switch, where policy can be applied. In practice, **one VLAN = one IP subnet = one broadcast domain**.

Why every enterprise uses VLANs

- **Security** — guests, cameras and finance PCs are separated and filtered at the gateway.
- **Smaller broadcast domains** — less noise per host.
- **Flexibility** — a user moves desks; only the port's VLAN changes, not the cabling.
- **Quality of service** — voice gets its own VLAN that is easy to prioritise.
- **Clean addressing** — the VLAN plan and the IP plan line up one to one.

The 802.1Q tag (4 bytes inserted after the source MAC)

FIELD	BITS	MEANING
TPID	16	0x8100 — marks the frame as tagged
PCP (priority)	3	Class of service 0–7 for QoS
DEI	1	Drop eligible indicator
VID	12	VLAN ID: 1–4094 usable (0 and 4095 reserved)

Access port vs trunk port

Access port

- Belongs to one data VLAN
- Frames sent and received untagged
- Connects end devices: PCs, printers, cameras
- Can add a separate voice VLAN for an IP phone

Trunk port

- Carries many VLANs over one link
- Frames tagged with 802.1Q (except the native VLAN)
- Connects switch-switch, switch-router, switch-AP, switch-server host
- Should list allowed VLANs explicitly

Creating VLANs and access ports

Switch SW1 · IOS-style CLI

```
SW1(config)# vlan 10
SW1(config-vlan)# name STAFF
SW1(config-vlan)# vlan 20
SW1(config-vlan)# name STUDENTS
SW1(config-vlan)# vlan 50
SW1(config-vlan)# name VOICE
SW1(config-vlan)# vlan 99
SW1(config-vlan)# name NATIVE-UNUSED
SW1(config-vlan)# exit
SW1(config)# interface range GigabitEthernet1/0/1 - 12
SW1(config-if-range)# switchport mode access
SW1(config-if-range)# switchport access vlan 10
SW1(config-if-range)# switchport voice vlan 50
SW1(config-if-range)# spanning-tree portfast
```

Hard-coding `switchport mode access` stops the port from ever becoming a trunk. The voice VLAN lets an IP phone tag its own traffic while the PC behind it stays untagged in VLAN 10.

Configuring an 802.1Q trunk

Switch SW1 · IOS-style CLI

```
SW1(config)# interface GigabitEthernet1/0/24
SW1(config-if)# switchport mode trunk
SW1(config-if)# switchport trunk native vlan 99
SW1(config-if)# switchport trunk allowed vlan 10,20,50,99
SW1(config-if)# switchport nonegotiate
```

Some multi-protocol switches also need `switchport trunk encapsulation dot1q` before `switchport mode trunk`. `switchport nonegotiate` turns off automatic trunk negotiation messages.

Verifying VLANs and trunks

Switch SW1 · IOS-style CLI

```
SW1# show vlan brief
VLAN Name                Status    Ports
-----
1    default                active    Gi1/0/13, Gi1/0/14, Gi1/0/15
10   STAFF                   active    Gi1/0/1, Gi1/0/2, Gi1/0/3
20   STUDENTS                active    Gi1/0/16, Gi1/0/17
99   NATIVE-UNUSED           active

SW1# show interfaces trunk
Port      Mode      Encapsulation  Status      Native vlan
Gi1/0/24  on        802.1q         trunking    99
Port      Vlans allowed on trunk
Gi1/0/24  10,20,50,99
```

Trunk ports do not appear in `show vlan brief` — check them with `show interfaces trunk`. Ports still in VLAN 1 are a sign of unconfigured ports.

WATCH OUT · NATIVE VLAN PITFALLS

Frames in the **native VLAN** cross a trunk untagged. If the two ends disagree on the native VLAN, untagged traffic silently lands in the wrong VLAN on the other side. Use the **same, unused** native VLAN at both ends (e.g. 99), never put users in it, and never leave VLAN 1 as the native VLAN carrying user traffic.

PRO TIP · HARDEN EVERY SWITCH ON DAY ONE

Put unused ports in an unused "parking" VLAN and shut them down; set every user port to `switchport mode access`; disable trunk negotiation; restrict allowed VLANs on trunks. These four habits close the most common layer 2 attacks (covered fully in the Guard module).

SCENARIO · DECIDE · A NEW VLAN DOES NOT WORK ON THE SECOND FLOOR

VLAN 30 (Accounts) works on SW1 but PCs in VLAN 30 on SW2 cannot reach their gateway on SW1. `show interfaces trunk` on SW1 shows allowed VLANs `10,20,50,99` on the link to SW2. VLAN 30 exists on both switches.

- ☐ A Run `switchport trunk allowed vlan 30` on the trunk
- ☐ B Run `switchport trunk allowed vlan add 30` on both ends of the trunk
- ☐ C Change the native VLAN to 30
- ☐ D Set the SW2 ports to `switchport mode trunk`

Circle your choice. Answer and feedback: module answer key.

TEMPLATE VLAN plan

One row per VLAN, agreed before any switch is configured. Pair it with the IP addressing plan.

VLAN ID	NAME	PURPOSE / USERS	IPV4 SUBNET	IPV6 /64	GATEWAY (SVI)	ALLOWED ON TRUNKS TO
20	STUDENTS	Lab PCs and student Wi-Fi	10.10.20.0/23	2001:db8:acad:20::/64	DS1 Vlan20 10.10.20.1	All access switches, WLC

KNOWLEDGE CHECK

6. How many bits does the 802.1Q tag use for the VLAN ID?

- ☐ A 8
- ☐ B 16
- ☐ C 3
- ☐ D 12

7. Which frames cross an 802.1Q trunk **untagged**?

- ☐ A Frames in VLAN 4095
- ☐ B Frames in the native VLAN
- ☐ C All broadcast frames
- ☐ D Frames sent by IP phones in the voice VLAN

8. A trunk allows VLANs 10,20. An engineer types `switchport trunk allowed vlan 40`. What is the result?

- ☐ A Only VLAN 40 is now allowed
- ☐ B VLANs 10, 20 and 40 are allowed
- ☐ C The command is rejected as a duplicate
- ☐ D All VLANs become allowed

9. Two switches have native VLAN 1 on one end and 99 on the other. What is the main risk?

- ☐ A The trunk speed drops to 100 Mbps
- ☐ B All tagged VLANs stop forwarding
- ☐ C Untagged traffic leaks into the wrong VLAN
- ☐ D The switches reboot repeatedly

10. Which command shows the VLANs actually allowed on a trunk?

- ☐ A `show vlan brief`
- ☐ B `show interfaces trunk`
- ☐ C `show mac address-table`
- ☐ D `show ip interface brief`

11. Where should a desk IP phone with a PC plugged into it connect?

- ☐ A A trunk port allowing every VLAN configured on the switch
- ☐ B An access port in the native VLAN
- ☐ C A routed port with an IP address
- ☐ D An access port with a data VLAN and a voice VLAN

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 3 · SWITCH · LESSON 3 · 28 MIN

Inter-VLAN Routing: Router-on-a-Stick and SVIs

Letting VLANs talk — through a gateway where policy can be applied

- Compare router-on-a-stick, SVIs on a layer 3 switch and separate routed interfaces
- Configure router-on-a-stick subinterfaces and layer 3 switch SVIs
- Troubleshoot an SVI that is up/down

VLANs keep traffic apart; **inter-VLAN routing** brings it back together in a controlled way. Every VLAN needs a default gateway — an IP interface that belongs to that VLAN's subnet. Where that gateway lives decides how fast inter-VLAN traffic flows and where you enforce policy.

Three ways to route between VLANs

Separate router interfaces

- One physical router port per VLAN
- Simple to understand
- Runs out of ports quickly
- Rare except in very small labs


Router-on-a-stick

- One trunk to the router, one subinterface per VLAN
- Cheap; fine for a small branch
- All inter-VLAN traffic crosses one link twice
- Throughput limited by that link and the router

Layer 3 switch with SVIs

- One virtual interface (SVI) per VLAN inside the switch
- Routing in hardware at line rate
- Standard for campus distribution/core
- Uplinks to router or firewall use routed ports

A packet from VLAN 10 to VLAN 20 via router-on-a-stick



Router-on-a-stick

R1 · IOS-style CLI

```

R1(config)# interface GigabitEthernet0/0
R1(config-if)# no shutdown
R1(config-if)# interface GigabitEthernet0/0.10
R1(config-subif)# description STAFF
R1(config-subif)# encapsulation dot1Q 10
R1(config-subif)# ip address 10.10.10.1 255.255.255.0
R1(config-subif)# interface GigabitEthernet0/0.20
R1(config-subif)# description STUDENTS
R1(config-subif)# encapsulation dot1Q 20
R1(config-subif)# ip address 10.10.20.1 255.255.255.0
R1(config-subif)# interface GigabitEthernet0/0.99
R1(config-subif)# encapsulation dot1Q 99 native
  
```

The subinterface number is only a label — the `encapsulation dot1Q` VLAN ID is what matters. Matching them (Gi0/0.10 ↔ VLAN 10) keeps life simple. The switch port facing the router must be a trunk carrying 10, 20 and 99.

SVIs on a layer 3 switch

Switch DS1 · IOS-style CLI

```
DS1(config)# ip routing
DS1(config)# vlan 10,20,30
DS1(config-vlan)# exit
DS1(config)# interface vlan 10
DS1(config-if)# ip address 10.10.10.1 255.255.255.0
DS1(config-if)# no shutdown
DS1(config-if)# interface vlan 20
DS1(config-if)# ip address 10.10.20.1 255.255.255.0
DS1(config-if)# no shutdown
DS1(config-if)# interface GigabitEthernet1/0/48
DS1(config-if)# no switchport
DS1(config-if)# ip address 10.0.0.1 255.255.255.252
```

`ip routing` turns on IPv4 routing between SVIs. `no switchport` converts Gi1/0/48 into a routed port for the uplink to the edge router or firewall.

An SVI comes up/up only when...

- The VLAN **exists** in the switch's VLAN database.
- At least one port in that VLAN is **up and forwarding** (an access port, or a trunk that carries the VLAN).
- The SVI is **not administratively shut down** (`no shutdown`).
- On many platforms, `ip routing` must be enabled for SVIs to route between each other.

Spotting a down SVI

Switch DS1 · IOS-style CLI

```
DS1# show ip interface brief | include Vlan
Vlan10          10.10.10.1      YES manual up
Vlan20          10.10.20.1      YES manual up
Vlan30          10.10.30.1      YES manual up
DS1# show vlan id 30
VLAN Name
-----
30    VLAN0030
      Status
      active
      Ports
```

VLAN 30 exists but has no ports, and no trunk carries it yet — so the SVI is up/down.

SCENARIO · DECIDE · ACCOUNTS CANNOT REACH THEIR GATEWAY

On DS1, `Vlan30` shows **up/down**. VLAN 30 exists. The accounts PCs are on access switch AS3, whose ports are in VLAN 30, and AS3 connects to DS1 through trunk Gi1/0/3 that allows VLANs 10,20,99.

- ☐ A Run `shutdown` then `no shutdown` on interface Vlan30
- ☐ B Change the SVI's IP address to 10.10.30.254
- ☐ C Enable `ip routing` again
- ☐ D Add VLAN 30 to the allowed list on the DS1–AS3 trunk (both ends)

Circle your choice. Answer and feedback: module answer key.

WATCH OUT · SAME SUBNET, TWO GATEWAYS

When migrating from router-on-a-stick to SVIs, never leave both the router subinterface and the new SVI configured with the same gateway IP — you create a duplicate address that breaks the VLAN intermittently. Remove one, then enable the other, in a planned window.

CASE STUDY · A UNIVERSITY COMPUTER-LAB NETWORK

A private university in Dhaka's Bashundhara area has two student computer labs (90 PCs each), a staff room (30 PCs), lab printers and a file server. Everything is in one flat VLAN 1 on one large layer 3-capable switch plus three access switches. Students have been browsing staff shares, and last week a student plugged a small home switch into two wall ports, taking the whole lab down for an hour.

1. Propose a VLAN design for this site.

2. Where should inter-VLAN routing happen and why?

3. Which switch features stop the "home switch in two wall ports" outage from repeating?

Model answer: module answer key.

KNOWLEDGE CHECK

12. On a router-on-a-stick, what decides which VLAN a subinterface serves?

- | | |
|--|--|
| ☐ A The subinterface number | ☐ B The interface description |
| ☐ C The <code>encapsulation dot1Q</code> VLAN ID | ☐ D The IP address's third octet |

13. An SVI shows up/down. What is the most likely cause?

- | | |
|---|--|
| ☐ A No active port is forwarding in that VLAN | ☐ B The SVI has been administratively shut down by an engineer |
| ☐ C <code>ip routing</code> has been enabled | ☐ D The SVI address is a private address |

14. Why are layer 3 switch SVIs preferred over router-on-a-stick in a campus?

- | | |
|--|--|
| ☐ A SVIs do not need IP addresses | ☐ B SVIs remove the need to create any VLANs on the switch |
| ☐ C Hardware routing with no single-trunk bottleneck | ☐ D SVIs work only with IPv6 traffic |

15. Which command turns a layer 3 switch port into a routed port?

- | | |
|--|---|
| ☐ A <code>no switchport</code> | ☐ B <code>switchport mode access</code> |
| ☐ C <code>ip routing</code> | ☐ D <code>switchport trunk native vlan 1</code> |

16. PCs in VLAN 10 can reach each other but not VLAN 20. The router-on-a-stick config looks correct. What should you check next?

- | | |
|--|---|
| ☐ A The PCs' network card drivers are the latest available version | ☐ B The router's hostname matches the switch's |
| ☐ C The MAC table aging timer is 300 seconds | ☐ D The router-facing switch port trunks both VLANs |

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 3 · SWITCH · LESSON 4 · 32 MIN

Spanning Tree: Loop Prevention with STP and RSTP

Root bridge election, port roles and states, PortFast and BPDU guard

- Explain why layer 2 loops are destructive
- Predict the root bridge, root ports, designated ports and blocked ports
- Compare 802.1D STP with Rapid STP (802.1w)
- Configure root priority, PortFast and BPDU guard and verify the result

Redundant links between switches are essential — but Ethernet frames have **no TTL**. A broadcast that enters a loop circles forever, multiplying at every switch until links and CPUs are saturated. **Spanning Tree Protocol (STP)** prevents this by logically blocking just enough ports to leave one loop-free path, and unblocking them automatically if the active path fails.

Symptoms of a layer 2 loop

- **Broadcast storm** — link utilisation jumps to 100%, every host slows down.
- **MAC flapping** — the same MAC address is learned on different ports again and again.
- **Duplicate frames** — applications receive the same data twice.
- Switches become slow to answer on the CLI; port LEDs blink in unison.

How STP builds the tree

1. **Elect the root bridge** Switches exchange BPDUs. The lowest **bridge ID** (priority, then MAC address) wins. Default priority is 32768, so the oldest/lowest MAC often wins by accident.
2. **Choose a root port on every other switch** The port with the lowest total **path cost** to the root. Ties: lowest sender bridge ID, then lowest sender port ID.
3. **Choose a designated port on every segment** The port on each link with the lowest cost back to the root forwards for that link. All root bridge ports are designated.
4. **Block the rest** Every remaining port is non-designated (alternate in RSTP) and discards data frames, breaking the loop.

STP port costs by link speed

LINK SPEED	SHORT (802.1D-1998) COST	LONG (802.1T) COST
10 Mbps	100	2,000,000
100 Mbps	19	200,000
1 Gbps	4	20,000
10 Gbps	2	2,000

DEFINITION · BRIDGE ID AND THE EXTENDED SYSTEM ID

The **bridge ID** is 8 bytes: a 2-byte priority and the switch's 6-byte MAC. With per-VLAN spanning tree, 12 bits of the priority carry the VLAN number, so priority can only be set in **steps of 4096** and the displayed value is priority + VLAN ID (32768 + 10 = 32778 for VLAN 10).

802.1D STP vs Rapid STP (802.1w)

Classic STP

- Port states: blocking, listening, learning, forwarding, disabled
- Timer-driven: 30–50 seconds to converge
- Roles: root, designated, non-designated
- Topology changes relayed through the root

Rapid STP

- Port states: discarding, learning, forwarding
- Handshake-driven (proposal/agreement): typically under a few seconds
- Adds alternate and backup roles — pre-computed backups
- Edge ports go straight to forwarding

RSTP port roles

ROLE	STATE	MEANING
Root	Forwarding	Best path toward the root bridge (one per non-root switch)
Designated	Forwarding	Best port on a segment toward the root; forwards for that segment
Alternate	Discarding	Backup path to the root via another switch — takes over quickly
Backup	Discarding	Redundant port to a segment this switch already serves (rare)
Edge (PortFast)	Forwarding	Connects to an end device; no BPDUs expected

Placing the root and protecting edge ports

Switch DS1 / AS1 · IOS-style CLI

```
DS1(config)# spanning-tree mode rapid-pvst
DS1(config)# spanning-tree vlan 10,20,30 root primary
DS2(config)# spanning-tree vlan 10,20,30 root secondary
AS1(config)# spanning-tree mode rapid-pvst
AS1(config)# interface range GigabitEthernet1/0/1 - 44
AS1(config-if-range)# spanning-tree portfast
AS1(config-if-range)# spanning-tree bpduguard enable
```

`rapid-pvst` runs a separate RSTP instance per VLAN in IOS-style CLI; the standards-based multi-instance alternative is MSTP (802.1s). `root primary` sets a priority low enough to win (24576 when no lower priority exists), and `root secondary` sets 28672 for the backup.

Verifying spanning tree for one VLAN

Switch AS1 · IOS-style CLI

```
AS1# show spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol rstp
    Root ID    Priority    24586
              Address    0011.2233.4401
              Cost        4
              Port        23 (GigabitEthernet1/0/23)
    Bridge ID  Priority    32778 (priority 32768 sys-id-ext 10)
  Interface   Role Sts Cost    Prio.Nbr Type
  Gi1/0/1     Desg FWD 4      128.1   P2p Edge
  Gi1/0/23    Root FWD 4      128.23  P2p
  Gi1/0/24    Altn BLK 4      128.24  P2p
```

Root priority 24586 = 24576 + VLAN 10, so DS1 is root. Gi1/0/23 leads to the root; Gi1/0/24 is the alternate path, blocked until Gi1/0/23 fails. Gi1/0/1 is an edge port. (Short cost method shown; the long method would show 20000 for 1 Gbps.)

WATCH OUT · PORTFAST ONLY ON EDGE PORTS

PortFast skips the listening/learning wait. On a port connected to another switch it can create a temporary loop. Always pair it with **BPDU guard**, which err-disables the port the moment a BPDU arrives — turning a potential network-wide outage into one dead wall port.

SCENARIO · DECIDE · THE ROOT BRIDGE IS IN A CUPBOARD

`show spanning-tree` reveals that an old 100 Mbps access switch in a storeroom is the root bridge for every VLAN, because all switches use default priority and its MAC is lowest. Traffic between floors hairpins through it.

- ☐ A Unplug the old switch so another becomes root
- ☐ B Disable spanning tree to stop the hairpin
- ☐ C Set the core/distribution switches as root primary and secondary for all VLANs
- ☐ D Enable PortFast on the uplinks

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

17. Switch A has priority 32768 and MAC 0011.1111.1111; Switch B has priority 28672 and MAC 0011.9999.9999. Which becomes root?
 - ☐ A Switch A, because its MAC address is lower
 - ☐ B Switch B, because its priority is lower
 - ☐ C Both switches share the root role
 - ☐ D Whichever switch booted first
18. Which RSTP port role is a pre-computed backup path toward the root bridge?
 - ☐ A Designated
 - ☐ B Edge
 - ☐ C Root
 - ☐ D Alternate
19. With per-VLAN spanning tree, which priority value can you configure?
 - ☐ A 10000
 - ☐ B 8192
 - ☐ C 32770
 - ☐ D 100
20. A user plugs a small switch into a port with PortFast and BPDU guard. What happens?
 - ☐ A The new switch becomes the root bridge
 - ☐ B The port moves to the alternate role
 - ☐ C The port is err-disabled when it receives a BPDU
 - ☐ D Spanning tree is disabled on the VLAN
21. How long does classic 802.1D STP typically take to move a port to forwarding after a topology change?
 - ☐ A 30–50 seconds
 - ☐ B Less than one second
 - ☐ C Two to five minutes
 - ☐ D Exactly 2 seconds
22. Which is a classic sign of a layer 2 loop?
 - ☐ A CRC and input errors climbing on a single access port
 - ☐ B One VLAN missing from a trunk
 - ☐ C An SVI showing up/down
 - ☐ D The same MAC flaps between two ports

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 3 · SWITCH · LESSON 5 · 26 MIN

EtherChannel with LACP and Neighbour Discovery with LLDP

Bundling links for bandwidth and resilience — and mapping who is connected to whom

- Explain why EtherChannel is used and how LACP negotiates a bundle
- Configure and verify a layer 2 LACP EtherChannel
- Use LLDP to discover and document neighbouring devices

Connect two switches with two parallel links and spanning tree will block one of them. **EtherChannel** (link aggregation, IEEE 802.1AX — originally 802.3ad) bundles several physical links into one logical **port-channel**. STP sees one link, so all members forward, bandwidth adds up and losing one cable only reduces capacity.

What link aggregation gives you

- **More bandwidth** — up to 8 active members in one bundle on most platforms.
- **Resilience** — a failed member is removed in a moment; traffic shifts to the others.
- **Simpler STP** — one logical port, no blocked parallel links.
- **One place to configure** — trunk settings live on the port-channel interface.

LACP mode combinations

SIDE A	SIDE B	BUNDLE FORMS?
active	active	Yes
active	passive	Yes
passive	passive	No — nobody starts negotiation
on (static)	on (static)	Yes, but with no negotiation or error checking
on (static)	active	No — modes must both be static or both LACP

Building a layer 2 LACP bundle

Switch AS1 · IOS-style CLI

```
AS1(config)# interface range GigabitEthernet1/0/23 - 24
AS1(config-if-range)# channel-group 1 mode active
AS1(config-if-range)# exit
AS1(config)# interface Port-channel1
AS1(config-if)# switchport mode trunk
AS1(config-if)# switchport trunk native vlan 99
AS1(config-if)# switchport trunk allowed vlan 10,20,30,50,99
```

Configure the same on the neighbour (active or passive). Settings applied to Port-channel1 are inherited by the members.

Verifying the bundle

Switch AS1 · IOS-style CLI

```
AS1# show etherchannel summary
Flags:  D - down          P - bundled in port-channel
        I - stand-alone  s - suspended
        S - Layer2       U - in use
Number of channel-groups in use: 1
Group  Port-channel  Protocol  Ports
-----+-----+-----+-----
1      Po1(SU)      LACP      Gi1/0/23(P)  Gi1/0/24(P)
```

Po1(SU) = layer 2, in use. (P) on each member = bundled. A member showing (s) or (I) needs attention.

INSIGHT · LOAD BALANCING IS PER FLOW, NOT PER PACKET

Each frame is assigned to a member by hashing addresses (MAC, IP and/or ports). One large file transfer between two hosts therefore uses **one** member link; many flows spread across all members. A 2 × 1 Gbps bundle does not make a single download run at 2 Gbps.

WATCH OUT · MEMBERS MUST MATCH

Every member must have the same speed and duplex, the same mode (access or trunk), the same access VLAN or the same native and allowed VLANs. A mismatched port is suspended or left out of the bundle — and if bundling fails on only one side, you risk a loop.

LLDP (Link Layer Discovery Protocol, IEEE 802.1AB) lets each device advertise its name, port, capabilities and management address to its directly connected neighbours. It is vendor-neutral, so switches, routers, APs, phones and servers from different makers can all see each other. **LLDP-MED** extensions tell IP phones which voice VLAN and power level to use.

Discovering neighbours with LLDP

Switch AS1 · IOS-style CLI

```
AS1(config)# lldp run
AS1# show lldp neighbors
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (W) WLAN Access Point, (S) Station
Device ID      Local Intf      Hold-time  Capability  Port ID
DS1             Gi1/0/23        120        B,R         Gi1/0/5
DS1             Gi1/0/24        120        B,R         Gi1/0/6
AP-2F-01        Gi1/0/7         120        W           eth0
Total entries displayed: 3
```

LLDP advertisements are sent every 30 seconds by default and held for 120 seconds. `show lldp neighbors detail` adds IP addresses and software versions.

PRO TIP · LLDP IS YOUR FREE DOCUMENTATION TOOL

Run LLDP on all infrastructure ports and collect neighbour tables regularly — they rebuild a physical topology diagram automatically. Consider disabling it on ports facing untrusted users or the internet, since it reveals device details.

SCENARIO · DECIDE · ONE MEMBER WILL NOT BUNDLE

`show etherchannel summary` shows Po1(SU) LACP Gi1/0/23(P) Gi1/0/24(s) . On Gi1/0/24 someone earlier configured `switchport trunk native vlan 1` directly on the member, while Port-channel1 uses native VLAN 99.

- ☐ A Remove Gi1/0/24 from the bundle and leave it as a separate trunk
- ☐ B Make Gi1/0/24's settings match the port-channel (native VLAN 99), then confirm it shows (P)
- ☐ C Change the whole bundle to mode `on` to force it
- ☐ D Replace the cable on Gi1/0/24

Circle your choice. Answer and feedback: module answer key.

Switch deployment checklist (VLANs, trunks, STP, bundles)

- ☐ VLANs created and named exactly as in the VLAN plan
- ☐ All user ports set to `switchport mode access` with the correct data/voice VLAN
- ☐ Unused ports shut down and placed in a parking VLAN
- ☐ Trunks: same unused native VLAN both ends, explicit allowed list, negotiation off
- ☐ RSTP (or MSTP) enabled; root primary/secondary set on distribution/core
- ☐ PortFast + BPDU guard on every edge port, never on switch-to-switch links
- ☐ EtherChannel members identical; bundle shows (SU) with all members (P)
- ☐ LLDP enabled on infrastructure ports and neighbour table saved to documentation
- ☐ `show interfaces status` reviewed for err-disabled ports and duplex issues
- ☐ Configuration saved and backed up off the device

REFLECT

Picture a switch closet you have seen (or your university lab). If one uplink cable were pulled out right now, what would happen to users — and which feature from this module (STP, RSTP, EtherChannel) would decide how long they noticed?

Key terms**MAC address table**

Switch table mapping VLAN + MAC address to a port; also called the CAM table.

Broadcast domain

The set of devices reached by a layer 2 broadcast — one VLAN.

802.1Q

IEEE standard that inserts a 4-byte tag with a 12-bit VLAN ID into Ethernet frames on trunks.

SVI

Switched virtual interface — a layer 3 interface for a VLAN on a layer 3 switch.

Flooding

Sending a frame out every port in the VLAN except the arrival port (unknown unicast, broadcast).

VLAN

A logical layer 2 network on shared switches; normally one IP subnet per VLAN.

Native VLAN

The VLAN whose frames cross an 802.1Q trunk untagged; must match at both ends.

Router-on-a-stick

Inter-VLAN routing using one router trunk link with one subinterface per VLAN.

Root bridge

The STP reference switch with the lowest bridge ID; all paths are measured toward it.

BPDU guard

Feature that err-disables an edge port if it receives a spanning-tree BPDU.

LLDP

IEEE 802.1AB vendor-neutral protocol for advertising device identity to direct neighbours.

Alternate port

RSTP discarding port that provides a ready backup path to the root.

LACP

Link Aggregation Control Protocol — negotiates EtherChannel bundles (active/passive).

KNOWLEDGE CHECK

23. Why do two parallel trunks between switches not double usable bandwidth without EtherChannel?

- ☐ A STP blocks one link to prevent a loop
- ☐ B Trunks can only carry one VLAN each
- ☐ C LLDP disables the second link
- ☐ D 802.1Q tags halve the bandwidth

24. Both sides of a link are set to LACP **passive**. What happens?

- ☐ A The bundle forms normally
- ☐ B No bundle forms
- ☐ C The ports become static **on** members
- ☐ D One side automatically becomes active

25. A 2 × 10 Gbps LACP bundle carries one large backup flow between two servers. What speed can that single flow reach?

- ☐ A Up to 20 Gbps
- ☐ B Up to 5 Gbps
- ☐ C Up to 10 Gbps
- ☐ D Up to 1 Gbps

26. Which flag in **show etherchannel summary** means a member is actively bundled?

- ☐ A s
- ☐ B P
- ☐ C I
- ☐ D D

27. What information does LLDP provide about a directly connected neighbour?

- ☐ A Its name, port ID and capabilities
- ☐ B Its full routing table
- ☐ C Every MAC address it has learned in each VLAN
- ☐ D Its saved configuration file

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 3 · SWITCH

Answer key

Knowledge checks

- 1. A** Unknown unicast frames are flooded within the VLAN. Switches do not return errors, have no default uplink for unknown MACs, and never flood across VLANs.
- 2. C** Learning uses the source MAC and arrival port. Forwarding uses the destination MAC; IP addresses and priority bits are not used for learning.
- 3. B** Each VLAN is its own broadcast domain. One would mean no VLANs, ports correspond to collision domains, and two does not match the three VLANs.
- 4. A** Many MACs on one port is normal on uplinks and AP ports. Duplex problems do not create MAC entries, overflow would affect all ports, and one cable serves one device.
- 5. D** Routers do not forward broadcasts between networks. Hubs and repeaters forward everything, and a switch without VLANs is one broadcast domain.
- 6. D** The VID is 12 bits, giving IDs 1–4094. 16 bits is the TPID, 3 bits is the priority field, and 8 bits would allow only 256 values.
- 7. B** Only native-VLAN frames are sent untagged. VLAN 4095 is reserved, broadcasts are tagged with their VLAN, and phone traffic is tagged in the voice VLAN.
- 8. A** Without `add`, the command replaces the list. `switchport trunk allowed vlan add 40` would keep 10 and 20.
- 9. C** A native mismatch merges two VLANs for untagged frames. It does not change speed, block tagged VLANs or reboot devices (switches usually log a warning).
- 10. B** `show interfaces trunk` lists mode, native VLAN and allowed VLANs. `show vlan brief` omits trunk ports, the MAC table shows learned addresses, and `show ip interface brief` lists layer 3 interfaces.
- 11. D** Access ports with a voice VLAN carry the phone's tagged voice and the PC's untagged data. A full trunk to a desk is a security risk, the native VLAN should hold no users, and a routed port does not switch.
- 12. C** The dot1Q encapsulation sets the VLAN. Subinterface numbers and descriptions are labels, and the IP address does not select a VLAN.
- 13. A** Protocol down means the VLAN has no active, forwarding port. An admin-shut SVI would show administratively down, `ip routing` does not bring SVIs down, and private addresses are normal.
- 14. C** SVIs route at line rate inside the switch. They do need IP addresses, they depend on VLANs, and they support IPv4 and IPv6.
- 15. A** `no switchport` makes the port a layer 3 interface. Access mode keeps it a layer 2 port, `ip routing` enables routing globally, and the native VLAN command applies to trunks.
- 16. D** If the switch-to-router link is an access port or lacks VLAN 20, frames never reach the subinterface. Drivers, hostnames and aging timers do not cause this symptom.
- 17. B** Priority is compared before MAC, and $28672 < 32768$. MAC only breaks ties, there is one root per instance, and boot order is irrelevant.
- 18. D** Alternate ports discard but can take over quickly if the root port fails. Designated and root ports forward, and edge ports face end devices.
- 19. B** Priority must be a multiple of 4096 (0, 4096, 8192 ...). The other values are not multiples of 4096.
- 20. C** BPDU guard shuts the port as soon as a BPDU arrives. Root election, alternate roles and disabling STP do not occur.
- 21. A** Listening and learning each take 15 seconds of forward delay, plus up to 20 seconds of max age. Sub-second behaviour is RSTP-like; minutes or 2 seconds are wrong.
- 22. D** MAC flapping happens when looping frames arrive from different directions. CRC errors point to cabling or duplex, a missing VLAN to trunk config, and an up/down SVI to missing active ports.
- 23. A** Spanning tree blocks the redundant path. Trunks carry many VLANs, LLDP only advertises information, and tags add just 4 bytes.
- 24. B** Passive only responds; with nobody initiating, LACP never negotiates. Modes do not change automatically and static `on` is a separate setting.
- 25. C** Hashing pins a flow to one member link, so a single flow is limited to one member's speed. Many flows together can use the full 20 Gbps.
- 26. B** P = bundled, s = suspended, I = stand-alone and D = down.
- 27. A** LLDP advertises identity, port, capabilities and management address. It does not share routing tables, MAC tables or configurations.

Scenarios & case studies

Scenario · A suspicious device

A is best. Correct. ARP maps IP to MAC; MAC tables map MAC to port. Following uplinks leads you to the exact port.

Scenario · A new VLAN does not work on the second floor

B is best. Correct. `add` appends VLAN 30 while keeping 10, 20, 50 and 99.

Scenario · Accounts cannot reach their gateway

D is best. Correct. Once the trunk carries VLAN 30 and forwards, DS1 has an active port in the VLAN and the SVI comes up/up.

Case study · A university computer-lab network

Create VLANs per role: **Lab A** and **Lab B** (or one student VLAN of /24 with room to grow), **Staff, Printers, Servers** and a **Management** VLAN, plus an unused native VLAN for trunks. Route between them with **SVIs on the layer 3 switch** — wire-speed routing in hardware — and apply ACLs so student VLANs can reach printers and the internet but not staff shares (ACL details come in the Guard module). Make the access-switch uplinks 802.1Q trunks with explicit allowed lists. To stop the loop, run **RSTP**, set the core switch as **root bridge**, and enable **PortFast with BPDU guard** on every student-facing port so a port that hears a BPDU from a rogue switch shuts down instead of looping the network.

Scenario · The root bridge is in a cupboard

C is best. Correct. Lowering their priority makes the root deterministic and puts it at the centre of the topology.

Scenario · One member will not bundle

B is best. Correct. Members must match; once consistent, LACP bundles the port.

04

MODULE 4 · ROUTE

Routing Essentials

Routing tables, static and default routes, single-area OSPFv2 and first-hop redundancy

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Read a routing table and apply longest-prefix match, administrative distance and metric
- Configure IPv4 and IPv6 static, default and floating static routes
- Explain OSPF neighbours, router ID, DR/BDR election and cost
- Configure and verify single-area OSPFv2 with passive interfaces and default-route origination
- Design and verify a VRRP-style redundant default gateway

LESSONS

1. The Routing Table: Longest Match, Administrative Distance and Metrics · 28 min
2. Static and Default Routes for IPv4 and IPv6 · 28 min
3. OSPFv2 Concepts: Neighbours, Router ID, DR/BDR and Cost · 30 min
4. Configuring and Verifying Single-Area OSPF · 32 min
5. First-Hop Redundancy with VRRP-Style Gateways · 24 min

MODULE 4 · ROUTE · LESSON 1 · 28 MIN


The Routing Table: Longest Match, Administrative Distance and Metrics

How a router decides where every packet goes

- Describe the steps a router takes to forward a packet
- Read an IOS-style routing table, including codes, AD and metric
- Apply longest-prefix match to choose a route
- Explain how administrative distance and metric select routes for the table

A router has one question for every packet: **which next hop and exit interface take this destination closest to its goal?** The answer comes from the **routing table**, a list of known destination prefixes and how to reach each one. Routes enter the table from three sources: directly connected interfaces, static routes typed by an engineer, and dynamic routing protocols such as OSPF.

What a router does with each packet



Reading a routing table

R1 · IOS-style CLI

```

R1# show ip route
Codes: L - local, C - connected, S - static, O - OSPF, * - candidate default
Gateway of last resort is 203.0.113.1 to network 0.0.0.0
S* 0.0.0.0/0 [1/0] via 203.0.113.1
    10.0.0.0/8 is variably subnetted, 7 subnets, 5 masks
C    10.0.12.0/30 is directly connected, GigabitEthernet0/1
L    10.0.12.1/32 is directly connected, GigabitEthernet0/1
C    10.0.13.0/30 is directly connected, GigabitEthernet0/2
L    10.0.13.1/32 is directly connected, GigabitEthernet0/2
S    10.2.0.0/16 [1/0] via 10.0.13.2
O    10.2.20.0/24 [110/200] via 10.0.12.2, 00:12:44, GigabitEthernet0/1
S    10.2.20.128/25 [1/0] via 10.0.13.2
C    192.168.10.0/24 is directly connected, GigabitEthernet0/0
L    192.168.10.1/32 is directly connected, GigabitEthernet0/0
  
```

In `[110/200]`, 110 is the administrative distance and 200 is the metric. `L` routes are the router's own interface addresses as /32 host routes. (The internet-facing connected route is trimmed from this listing.)

DEFINITION · LONGEST-PREFIX MATCH

When several routes contain the destination, the router uses the one with the **longest prefix length** (the most specific route), regardless of where it came from. A /25 beats a /24, which beats a /16, which beats the /0 default.

Longest-prefix match on the table above

DESTINATION	MATCHING ROUTES	CHOSEN ROUTE	NEXT HOP
10.2.20.130	/0, /16, /24, /25	10.2.20.128/25 (static)	10.0.13.2

DESTINATION	MATCHING ROUTES	CHOSEN ROUTE	NEXT HOP
10.2.20.50	/0, /16, /24	10.2.20.0/24 (OSPF)	10.0.12.2
10.2.99.7	/0, /16	10.2.0.0/16 (static)	10.0.13.2
8.8.4.4	/0 only	0.0.0.0/0 (default)	203.0.113.1
192.168.10.77	/0, /24	192.168.10.0/24 (connected)	Direct on Gi0/0

Default administrative distances (IOS-style platforms)

ROUTE SOURCE	AD	NOTES
Connected interface	0	Most trusted
Static route	1	Can be raised manually (floating static)
External BGP	20	Routes from another organisation's network
OSPF	110	Link-state interior protocol
IS-IS	115	Link-state interior protocol
RIP	120	Legacy distance-vector protocol
Internal BGP	200	Routes learned from peers in the same organisation
Unknown / unusable	255	Never installed

WATCH OUT · AD VALUES ARE NOT UNIVERSAL

These are common defaults on IOS-style devices. Other platforms use different numbers (some give OSPF internal routes 10, for example). The **idea** is universal: a lower value means a more trusted source. Always check the platform you are working on.

Administrative distance vs metric

Administrative distance

- Compares **different sources** offering the same prefix
- Example: static (1) vs OSPF (110) for 10.5.0.0/16
- Local to one router; not advertised
- Lower wins

Metric

- Compares **paths within one protocol**
- Example: OSPF cost 20 vs cost 30
- Calculated by the protocol (OSPF: sum of interface costs)
- Lower wins; equal metrics can load-share (ECMP)

How routes get into the table and get used

1. **Each source offers its best path** OSPF runs SPF and offers its lowest-cost path; static routes offer themselves.
2. **Same prefix from several sources?** The source with the lowest administrative distance is installed.
3. **Same prefix, same source, equal metrics?** Several next hops are installed (equal-cost multipath).
4. **Forwarding** For each packet, longest-prefix match picks among everything installed.

INSIGHT · DIFFERENT PREFIXES NEVER COMPETE ON AD

10.2.20.0/24 from OSPF and 10.2.0.0/16 from a static route are **both** installed — they are different prefixes. AD only matters when the prefix *and* the length are identical. Mixing these two ideas up is the most common routing exam mistake.

SCENARIO · DECIDE · THE UNEXPECTED PATH

Traffic to server 10.2.20.140 goes through R3 (10.0.13.2), but the team expected it to follow OSPF through R2. The table is the one shown above.

- ☐ A The static 10.2.20.128/25 is the longest match for 10.2.20.140, so it is used
- ☐ B OSPF has a worse administrative distance than static, so static always wins
- ☐ C The default route is preferred for unknown servers
- ☐ D OSPF routes are ignored when a static default exists

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

- A router has routes 172.16.0.0/16, 172.16.4.0/22 and 172.16.4.0/24. Which route does a packet to 172.16.5.9 use?
 - ☐ A 172.16.4.0/24
 - ☐ B 172.16.0.0/16
 - ☐ C The default route
 - ☐ D 172.16.4.0/22
- The same prefix 10.9.0.0/16 is learned from OSPF and from a static route with default settings. Which is installed?
 - ☐ A The OSPF route
 - ☐ B Both, load-sharing traffic
 - ☐ C The static route
 - ☐ D Neither, because they conflict
- In `0 10.2.20.0/24 [110/200] via 10.0.12.2`, what does 200 represent?
 - ☐ A The administrative distance
 - ☐ B The route age in seconds
 - ☐ C The interface bandwidth
 - ☐ D The OSPF cost metric
- What does a router do when a packet's TTL reaches 0?
 - ☐ A Forwards it using the default route
 - ☐ B Resets the TTL to 255 and forwards it
 - ☐ C Drops it and sends ICMP time exceeded to the source
 - ☐ D Floods it to every interface
- What do `L` entries in an IOS-style routing table represent?
 - ☐ A Routes learned from LLDP neighbours
 - ☐ B The router's own interface addresses as /32s
 - ☐ C Load-balanced routes
 - ☐ D Legacy routes from a deleted protocol

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 4 · ROUTE · LESSON 2 · 28 MIN

Static and Default Routes for IPv4 and IPv6

Next-hop, exit-interface, default and floating static routes

- Configure IPv4 and IPv6 static routes, including default routes
- Choose between next-hop and exit-interface forms
- Build a floating static route as a backup path and explain its limits
- Troubleshoot missing return routes

A **static route** is a route you type by hand. It never changes unless you change it, uses no bandwidth or CPU for routing updates, and is perfectly predictable. That makes it ideal for small or stub networks, for the default route to an ISP and for backup paths — and a poor fit for large networks that change often.

Static vs dynamic routing

Static

- Configured manually per router
- No overhead, very predictable
- Does not react to failures beyond the local link
- Best for stub sites, defaults and backups

Dynamic (e.g. OSPF)

- Routers exchange reachability automatically
- Adapts to failures and new subnets
- Uses some CPU, memory and bandwidth
- Best for networks with several paths or many prefixes

IPv4 static and default routes

R1 · IOS-style CLI

```
! network route via a next-hop address
R1(config)# ip route 10.2.0.0 255.255.0.0 10.0.12.2
! host route (/32) to one server
R1(config)# ip route 10.9.9.10 255.255.255.255 10.0.13.2
! fully specified: exit interface plus next hop
R1(config)# ip route 10.3.0.0 255.255.0.0 GigabitEthernet0/2 10.0.13.2
! default route to the ISP
R1(config)# ip route 0.0.0.0 0.0.0.0 203.0.113.1
```

A next-hop route needs a second lookup to find the exit interface (a recursive lookup). The fully specified form avoids that and is clearest on multi-access Ethernet links.

WATCH OUT · EXIT-INTERFACE-ONLY ROUTES ON ETHERNET

`ip route 10.3.0.0 255.255.0.0 GigabitEthernet0/2` (no next hop) tells the router every destination in 10.3.0.0/16 is directly on that segment. On Ethernet it must then ARP for each destination and relies on the neighbour answering by proxy ARP. Use it only on true point-to-point links; on Ethernet give a next hop.


IPv6 static and default routes

R1 · IOS-style CLI

```
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 route 2001:db8:acad:200::/56 2001:db8:0:12::2
R1(config)# ipv6 route 2001:db8:acad:300::/56 GigabitEthernet0/2 fe80::3
R1(config)# ipv6 route ::/0 2001:db8:ffff:1::1
R1(config)# end
R1# show ipv6 route static
S    2001:DB8:ACAD:200::/56 [1/0]
    via 2001:DB8:0:12::2
S    2001:DB8:ACAD:300::/56 [1/0]
    via FE80::3, GigabitEthernet0/2
S    ::/0 [1/0]
    via 2001:DB8:FFFF:1::1
```

When the next hop is a link-local address you *must* name the exit interface too, because the same fe80:: address can exist on every link.

Floating static route: primary and backup



Configuring a floating static default

R1 · IOS-style CLI

```
R1(config)# ip route 0.0.0.0 0.0.0.0 203.0.113.1
R1(config)# ip route 0.0.0.0 0.0.0.0 198.51.100.1 10
R1# show ip route static | include 0.0.0.0
S*    0.0.0.0/0 [1/0] via 203.0.113.1
```

Only the AD 1 route appears while it is valid. The trailing **10** raises the backup's AD; to float behind a route learned by OSPF (AD 110), the backup's AD must be higher than 110.

INSIGHT · FLOATING STATICS ONLY SEE LOCAL FAILURES

A static route is removed when its exit interface goes down or its next hop becomes unreachable in the table. If the provider fails *beyond* your directly connected link, the interface stays up and the primary route stays — traffic is black-holed. Production designs add **reachability tracking** (probes that remove the route when a far-end test fails) or use a dynamic protocol.

Kinds of static route

TYPE	EXAMPLE PREFIX	TYPICAL USE
Network route	10.2.0.0/16	Reach a remote site or block
Host route	10.9.9.10/32	Steer traffic to one server or loopback
Default route	0.0.0.0/0 ::/0	Send everything unknown to the ISP or core
Summary route	10.0.0.0/8	One route covering many subnets behind one next hop
Floating static	Any, with a raised AD	Backup path that appears only when the primary disappears

SCENARIO • DECIDE • PING TIMES OUT ONE WAY

The branch router BR1 has a default route to HQ. A branch PC (10.40.1.20) pings an HQ server (10.1.1.10): the request reaches the server (a capture there proves it) but replies never arrive. HQ router HQ1 has routes for its own subnets and a default to the internet.

- ☐ A Add a route on BR1 for 10.1.1.0/24
- ☐ B Change the server's default gateway to BR1
- ☐ C Add a route on HQ1 for the branch prefix 10.40.0.0/16 via the WAN link to BR1
- ☐ D Increase the AD of HQ1's default route

Circle your choice. Answer and feedback: module answer key.

TEMPLATE Routing table worksheet

Plan or audit every route on a router: what it reaches, how, and why it exists.

[illegible]

KNOWLEDGE CHECK

6. Which command creates an IPv4 default route to 198.51.100.1?

- ☐ A ip route 0.0.0.0 0.0.0.0 198.51.100.1
 ☐ B ip route 198.51.100.1 0.0.0.0 0.0.0.0
- ☐ C ip default 198.51.100.1 255.255.255.255
 ☐ D ip route 0.0.0.0 255.255.255.255 198.51.100.1

7. A primary default route is static (AD 1). What AD should a backup static default use to float?

- ☐ A Exactly 1, the same as the primary
- ☐ B Any value above 1, such as 10
- ☐ C 0, so it is always preferred
- ☐ D 255, so it is always installed

8. Why must an IPv6 static route using a link-local next hop also name the exit interface?

- ☐ A The same link-local address can exist on many links
 ☐ B Link-local addresses cannot be pinged
- ☐ C IPv6 static routes always need two next hops
 ☐ D Link-local next hops are only valid for multicast traffic

9. A floating static default should take over when the ISP fails beyond your directly connected link. What is missing?

- ☐ A A lower AD on the floating route
 ☐ B A /32 host route to the ISP router
- ☐ C IPv6 enabled on the WAN link
 ☐ D Reachability tracking for the primary route

10. Branch pings reach HQ but replies never return. What is the most likely cause?

- ☐ A The branch PC's DNS server is wrong
- ☐ B The HQ switch has PortFast enabled
- ☐ C HQ has no route back to the branch subnet
- ☐ D The branch router uses a host route instead of a network route

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 4 · ROUTE · LESSON 3 · 30 MIN

OSPFv2 Concepts: Neighbours, Router ID, DR/BDR and Cost

How single-area OSPF discovers neighbours and computes best paths

- Explain how OSPF forms adjacencies and builds a link-state database
- List the parameters that must match for neighbours to form
- Predict the router ID, the DR and the BDR on a segment
- Calculate OSPF cost from interface bandwidth and reference bandwidth

OSPF (Open Shortest Path First, OSPFv2 for IPv4) is the most widely deployed open interior routing protocol. Each router describes its links in **link-state advertisements (LSAs)**; all routers in an area hold the same **link-state database (LSDB)** and each runs the **SPF (Dijkstra) algorithm** to compute its own shortest path to every prefix. In this course you work with a single area: **area 0, the backbone**.

OSPF in one minute

- Open standard (OSPFv2 is defined in RFC 2328); runs directly over IP as protocol 89.
- Metric is **cost**, based on interface bandwidth; lower total cost wins.
- Hellos to multicast **224.0.0.5** discover neighbours; DR/BDR listen on **224.0.0.6**.
- Converges fast and sends updates only when something changes (plus periodic refresh).
- Areas let large networks scale; small and medium networks often use area 0 only.

Neighbour states from first hello to full adjacency

STATE	WHAT IS HAPPENING
Down	No hellos received from this neighbour yet
Init	A hello arrived, but it does not list our router ID yet
2-Way	Each router sees itself in the other's hello — bidirectional; DR/BDR election happens here
ExStart	Routers agree who leads the database exchange (MTU problems often stall here)
Exchange	Database description packets summarise each LSDB
Loading	Missing LSAs are requested and received
Full	Databases are synchronised — the adjacency is complete

Must match for an adjacency to form

PARAMETER	DEFAULT ON ETHERNET	IF MISMATCHED
Area ID	—	Hellos ignored; no neighbour
Subnet and mask on the link	—	No neighbour on broadcast links
Hello / dead intervals	10 s / 40 s	No neighbour
Authentication type and key	None	No neighbour
Area type (e.g. stub flag)	Normal	No neighbour
Interface MTU	1500	Neighbour stuck in ExStart/Exchange

PARAMETER	DEFAULT ON ETHERNET	IF MISMATCHED
Router IDs	Must be unique	Unstable adjacencies and route errors

How a router chooses its OSPF router ID

1. **Manual** The `router-id` command — always the best practice.
2. **Highest loopback** If none is set, the highest IPv4 address on any up loopback interface.
3. **Highest active interface** If no loopback exists, the highest IPv4 address on an up physical interface.
4. **Changing it** A new router ID takes effect only after the OSPF process restarts (e.g. `clear ip ospf process`).

DEFINITION · DR AND BDR

On multi-access segments such as Ethernet, OSPF elects a **designated router (DR)** and a **backup DR (BDR)**. Every other router (a **DROTHER**) forms a full adjacency only with the DR and BDR, not with each other. This cuts the number of adjacencies and LSAs on busy segments.

DR/BDR election rules

- Highest interface **OSPF priority** wins (default 1; range 0–255).
- Priority **0** means the router never becomes DR or BDR.
- Tie: highest **router ID** wins.
- Election is **not pre-emptive** — a better router joining later does not take over until the DR fails or OSPF restarts.
- DROTHERs stay in **2-Way** with each other; that is normal, not a fault.
- Point-to-point links (or interfaces set to point-to-point) elect no DR/BDR.

OSPF cost

MEASURE	FORMULA	WHAT IT TELLS YOU
Interface cost	<code>cost = reference bandwidth ÷ interface bandwidth (minimum 1)</code>	Default reference is 100 Mbps, so 100 Mbps, 1 Gbps and 10 Gbps all get cost 1.
Better reference	<code>auto-cost reference-bandwidth 1000000 (Mbps)</code>	With a 100 Gbps reference: 100 Gbps = 1, 10 Gbps = 10, 1 Gbps = 100, 100 Mbps = 1000.
Path cost	<code>sum of outgoing interface costs along the path</code>	The router adds the cost of each interface the packet leaves through toward the destination.

Cost with each reference bandwidth

INTERFACE	REFERENCE 100 MBPS (DEFAULT)	REFERENCE 100 GBPS
100 Mbps	1	1000
1 Gbps	1	100
10 Gbps	1	10
100 Gbps	1	1

WATCH OUT · SET THE SAME REFERENCE BANDWIDTH EVERYWHERE

With the default reference, OSPF cannot tell a 1 Gbps link from a 100 Gbps link and may choose a slow path. Raise the reference on **every** router in the area, otherwise routers compute costs on different scales and paths become inconsistent.

SCENARIO · DECIDE · NEIGHBOURS THAT NEVER MEET

R1 and R2 share 10.0.12.0/30. `show ip ospf neighbor` on both is empty. Debug output shows R1 receiving hellos from R2 with hello 5 s, dead 20 s; R1 uses 10 s / 40 s. Area, subnet and authentication all match.

- ☐ A Make the hello and dead intervals identical on both interfaces
- ☐ B Change R1's router ID to be higher than R2's
- ☐ C Set the OSPF priority to 0 on both
- ☐ D Increase the reference bandwidth on R1

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

11. Which OSPF neighbour state shows that databases are fully synchronised?

- ☐ A 2-Way
- ☐ B Full
- ☐ C ExStart
- ☐ D Loading

12. A router has no `router-id` command, loopbacks 10.255.0.3 and 10.255.0.9, and interfaces up with 192.168.1.1 and 172.16.0.1. What is its router ID?

- ☐ A 192.168.1.1
- ☐ B 10.255.0.3
- ☐ C 10.255.0.9
- ☐ D 172.16.0.1

13. Two DROTHERs on the same Ethernet segment show each other in state 2-WAY. What does this mean?

- ☐ A Their MTUs do not match
- ☐ B Their hello and dead timers do not match
- ☐ C Authentication has failed
- ☐ D This is normal behaviour for DROTHERs

14. Using the default reference bandwidth, what is the OSPF cost of a 10 Gbps interface?

- ☐ A 1
- ☐ B 10
- ☐ C 100
- ☐ D 0.01

15. On a shared Ethernet segment, which router becomes DR when all priorities are default?

- ☐ A The router with the highest router ID
- ☐ B The router with the lowest router ID
- ☐ C The router that has the most neighbours
- ☐ D The router with the fastest interface

16. Which mismatch typically leaves two OSPF neighbours stuck in ExStart or Exchange?

- ☐ A Interface description
- ☐ B Interface MTU
- ☐ C OSPF process ID
- ☐ D Router hostname

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 4 · ROUTE · LESSON 4 · 32 MIN

Configuring and Verifying Single-Area OSPF

From `router ospf` to a full routing table — and how to prove it works

- Enable OSPFv2 with network statements or per-interface commands
- Use wildcard masks, passive interfaces and default-route origination correctly
- Verify neighbours, interfaces, costs and OSPF routes with show commands

Configuring OSPF takes a handful of lines. The skill lies in enabling it on exactly the right interfaces, keeping hellos off user LANs, setting sensible costs and then **proving** with show commands that every neighbour and route is what you expect.

DEFINITION · WILDCARD MASK

OSPF `network` statements use a **wildcard mask** — the inverse of a subnet mask. A 0 bit means "must match" and a 1 bit means "don't care". For a /24 the wildcard is `0.0.0.255` ; for a /30 it is `0.0.0.3` . Quick method: 255.255.255.255 minus the subnet mask.

OSPF with network statements

R1 · IOS-style CLI

```
R1(config)# router ospf 1
R1(config-router)# router-id 1.1.1.1
R1(config-router)# auto-cost reference-bandwidth 100000
R1(config-router)# network 10.0.12.0 0.0.0.3 area 0
R1(config-router)# network 10.0.13.0 0.0.0.3 area 0
R1(config-router)# network 192.168.10.0 0.0.0.255 area 0
R1(config-router)# passive-interface GigabitEthernet0/0
R1(config-router)# default-information originate
R1(config-router)# interface GigabitEthernet0/1
R1(config-if)# ip ospf network point-to-point
```

Any interface whose address matches a `network` statement joins area 0. The process ID (1) is local only. `passive-interface` advertises the LAN subnet but sends no hellos to users. `default-information originate` advertises R1's static default route into OSPF. Gi0/1 is set to point-to-point to match R2's end of the link.

OSPF enabled per interface

R2 · IOS-style CLI

```
R2(config)# router ospf 1
R2(config-router)# router-id 2.2.2.2
R2(config-router)# auto-cost reference-bandwidth 100000
R2(config-router)# passive-interface GigabitEthernet0/0
R2(config-router)# interface GigabitEthernet0/0
R2(config-if)# ip ospf 1 area 0
R2(config-if)# interface GigabitEthernet0/1
R2(config-if)# ip ospf 1 area 0
R2(config-if)# ip ospf network point-to-point
```

Per-interface enablement is explicit and easy to read. Setting a router-to-router Ethernet link to point-to-point skips DR/BDR election and speeds up adjacency.

Options you will use often

- `passive-interface default` then `no passive-interface <uplink>` — safest on routers with many LAN interfaces.
- `ip ospf cost <n>` on an interface overrides the bandwidth-based cost.
- `ip ospf priority 0` keeps a router from ever becoming DR on that segment.
- Loopback interfaces are advertised as /32 host routes by default — ideal stable router IDs and management addresses.

Verifying neighbours and interfaces

R1 · IOS-style CLI

```
R1# show ip ospf neighbor
Neighbor ID    Pri  State           Dead Time   Address        Interface
2.2.2.2        0    FULL/-          00:00:36    10.0.12.2      GigabitEthernet0/1
3.3.3.3        1    FULL/DR         00:00:33    10.0.13.2      GigabitEthernet0/2
R1# show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
Gi0/0      1    0     192.168.10.1/24  100   DR    0/0
Gi0/1      1    0     10.0.12.1/30     100   P2P   1/1
Gi0/2      1    0     10.0.13.1/30     100   BDR   1/1
```

FULL/- is a point-to-point neighbour (no DR). On Gi0/2, R3 is DR and R1 is BDR. Cost 100 for 1 Gbps confirms the 100 Gbps reference. The passive LAN shows 0 neighbours, as intended.

Verifying OSPF routes

R1 · IOS-style CLI

```
R1# show ip route ospf
 10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
0       10.2.20.0/24 [110/200] via 10.0.12.2, 00:04:10, GigabitEthernet0/1
0       10.3.30.0/24 [110/200] via 10.0.13.2, 00:03:55, GigabitEthernet0/2
R1# show ip protocols | include Router ID|Routing for|Passive
Router ID 1.1.1.1
Routing for Networks:
Passive Interface(s):
```

Cost 200 = R1's outgoing 1 Gbps link (100) + R2's 1 Gbps LAN interface (100). `show ip protocols` confirms router ID, network statements and passive interfaces.

A verification routine after any OSPF change

1. **Interfaces** `show ip ospf interface brief` — every intended interface listed, in area 0, with the expected cost.
2. **Neighbours** `show ip ospf neighbor` — every expected neighbour FULL (or 2-WAY between DROTHERs).
3. **Routes** `show ip route ospf` — every remote subnet present with sensible costs.
4. **Defaults** Branches show `0xE2 0.0.0.0/0` if a default is being originated.
5. **Data plane** Ping and trace from a user subnet to a remote subnet to prove the forwarding path.

SCENARIO · DECIDE · A SUBNET THAT WILL NOT ADVERTISE

R3's LAN is 10.3.30.0/25 on Gi0/0. R3 has `network 10.3.30.128 0.0.0.127 area 0`. Neighbours are FULL, but no other router learns 10.3.30.0/25.

- ☐ A Make Gi0/0 passive
- ☐ B Restart the OSPF process on every router
- ☐ C Raise the reference bandwidth on R3
- ☐ D Fix the statement to `network 10.3.30.0 0.0.0.127 area 0` (or use `ip ospf 1 area 0` on Gi0/0)

Circle your choice. Answer and feedback: module answer key.

CASE STUDY · REPLACING STATIC ROUTES ACROSS THREE SITES

A pharmaceutical distributor has its head office in Tejgaon, Dhaka and warehouses in Khulna and Rajshahi. Each pair of sites is linked by a provider point-to-point circuit, forming a triangle. Every router uses static routes. When the Dhaka–Khulna circuit failed last month, Khulna was cut off for six hours even though the Khulna–Rajshahi–Dhaka path was healthy. Only Dhaka has internet access.

1. Why did the static design fail to use the healthy path?

2. Outline a single-area OSPF design for the three routers.

3. How do the warehouses learn the way to the internet?

Model answer: module answer key.

OSPF deployment and verification checklist

- ☐ Router ID set manually and unique on every router
- ☐ Same reference bandwidth configured on every router in the area
- ☐ Only intended interfaces enabled; wildcard masks double-checked
- ☐ User-facing LAN interfaces passive
- ☐ Router-to-router Ethernet links set to point-to-point where appropriate
- ☐ Hello/dead timers, area, subnet, MTU and authentication consistent on each link
- ☐ Every expected neighbour FULL (DROTHER pairs may be 2-WAY)
- ☐ Every remote subnet present in `show ip route ospf` with the expected cost
- ☐ Default route originated from the internet edge only
- ☐ Configuration saved; routing table snapshot stored for comparison

KNOWLEDGE CHECK

17. What is the wildcard mask for the subnet 10.0.12.0/30?

- ☐ A 255.255.255.252
- ☐ B 0.0.0.4
- ☐ C 0.0.0.252
- ☐ D 0.0.0.3

18. What does `passive-interface GigabitEthernet0/0` do in OSPF?
- ☐ A Removes Gi0/0's subnet from OSPF entirely
 - ☐ B Disables Gi0/0 so it shuts down
 - ☐ C Advertises Gi0/0's subnet but sends no hellos on it
 - ☐ D Makes Gi0/0 the DR on its segment
19. Branch routers need a default route learned via OSPF from the HQ router that has a static default. Which command on HQ achieves this?
- ☐ A `redistribute connected`
 - ☐ B `default-information originate`
 - ☐ C `ip route 0.0.0.0 0.0.0.0 null0`
 - ☐ D `passive-interface default`
20. `show ip ospf neighbor` shows `FULL/ -` for a neighbour. What does the dash mean?
- ☐ A No DR/BDR role on a point-to-point link
 - ☐ B The neighbour is down
 - ☐ C The neighbour has not been assigned a router ID yet
 - ☐ D Authentication is disabled
21. Two routers both have `router ospf 1` and `router ospf 20`. Must process IDs match for adjacency?
- ☐ A Yes, they must be identical
 - ☐ B Yes, but only on point-to-point links
 - ☐ C No, but both process IDs must be even numbers
 - ☐ D No, the process ID is locally significant

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 4 · ROUTE · LESSON 5 · 24 MIN

First-Hop Redundancy with VRRP-Style Gateways

Keeping the default gateway alive when a router or switch fails


- Explain why hosts need a redundant default gateway
- Describe how a virtual IP, virtual MAC, master and backup work together
- Configure and verify a VRRP group with priority, preemption and tracking

Hosts have exactly one default gateway address, usually set by DHCP. If the router or layer 3 switch holding that address fails, every host on the subnet is cut off — even if a second router sits right next to it. A **first-hop redundancy protocol (FHRP)** lets two or more routers share one **virtual gateway address**, so the hosts' setting never has to change.

DEFINITION · VRRP

Virtual Router Redundancy Protocol (VRRPv3 in RFC 5798, supporting IPv4 and IPv6) is the open-standard FHRP. Routers in a VRRP group share a **virtual IP** and a **virtual MAC**. One router is **master** and answers for the virtual address; the others are **backups** that take over if the master's advertisements stop.

How a VRRP failover unfolds



VRRP facts worth remembering

ITEM	VALUE
Virtual MAC (IPv4)	00-00-5E-00-01-XX, where XX is the group (VRID) in hex
Default priority	100 (range 1–254 configurable; 255 = the router that owns the real address)
Election	Highest priority wins; tie goes to the highest primary IP address
Preemption	Enabled by default in VRRP
Advertisements	Sent by the master to 224.0.0.18 (IP protocol 112), every 1 s by default
Roles	One master, one or more backups per group

Configuring VRRP on two gateways

R1 and R2 · IOS-style CLI

```
R1(config)# track 1 interface GigabitEthernet0/1 line-protocol
R1(config)# interface GigabitEthernet0/0
R1(config-if)# ip address 192.168.10.2 255.255.255.0
R1(config-if)# vrrp 10 ip 192.168.10.1
R1(config-if)# vrrp 10 priority 110
R1(config-if)# vrrp 10 track 1 decrement 20
R2(config)# interface GigabitEthernet0/0
R2(config-if)# ip address 192.168.10.3 255.255.255.0
R2(config-if)# vrrp 10 ip 192.168.10.1
```

Hosts use 192.168.10.1 as their gateway. If R1's uplink Gi0/1 goes down, tracking lowers R1 to 90, R2 (100) preempts and becomes master. Syntax varies between platforms and software versions; the concepts do not.

Verifying the group

R1 · IOS-style CLI

```
R1# show vrrp brief
Interface      Grp Pri Time   Own Pre State  Master addr  Group addr
Gi0/0          10 110 3570      Y Master 192.168.10.2 192.168.10.1
R2# show vrrp brief
Interface      Grp Pri Time   Own Pre State  Master addr  Group addr
Gi0/0          10 100 3609      Y Backup 192.168.10.2 192.168.10.1
```

Time is the master-down interval in milliseconds: three advertisement intervals plus a small skew that is shorter for higher priorities.

Without and with first-hop redundancy

Single gateway

- Gateway failure isolates the whole subnet
- Engineers must re-address or reconfigure DHCP during an outage
- Second router sits idle
- Simple but fragile

VRRP group

- Virtual gateway survives a router failure in seconds
- Hosts and DHCP settings never change
- Load can be shared by running two groups with different masters
- Small extra configuration per VLAN

Design tips

- Make the VRRP master for a VLAN the **same switch that is STP root** for that VLAN, so layer 2 and layer 3 paths line up.
- Use **tracking** of uplinks so a router that has lost its way out hands over mastership.
- To share load, make R1 master for odd VLANs and R2 master for even VLANs.
- Keep the virtual IP consistent across sites (e.g. always .1) so documentation is predictable.
- Vendor-specific FHRPs work on the same idea; VRRP is the choice when devices come from different makers.

SCENARIO · DECIDE · MASTER WITH NO WAY OUT

R1 is VRRP master for VLAN 10. Its WAN uplink fails, but its LAN interface stays up, so R1 remains master. Users' traffic now goes to R1, which has to send everything across to R2 — or drops it if it has no route.

- ☐ A Lower R2's priority so it never becomes master
- ☐ B Configure interface or reachability tracking on R1 to decrement its priority below R2's when the uplink fails
- ☐ C Disable preemption on both routers
- ☐ D Change hosts' gateway to R2's real address

Circle your choice. Answer and feedback: module answer key.

REFLECT

Look at a network you know — your office, campus or home. List every single point of failure between a user's laptop and the internet (cables, switches, gateway, ISP). Which one would you fix first with what you learned in this module, and how?

Key terms**Routing table**

A router's list of destination prefixes with next hops, exit interfaces, AD and metric.

Administrative distance

Trust value used to choose between different route sources for the same prefix; lower wins.

Static route

Manually configured route that does not change unless an engineer changes it.

Floating static route

Static route with a raised AD that is installed only when the preferred route disappears.

Router ID

32-bit dotted-decimal value that uniquely identifies an OSPF router.

Wildcard mask

Inverse mask used in OSPF network statements and ACLs; 0 = match, 1 = ignore.

VRRP

Open-standard first-hop redundancy protocol that shares a virtual gateway IP between routers.

Longest-prefix match

Rule that the most specific matching route (longest prefix) is used for forwarding.

Metric

Value a routing protocol uses to compare paths; OSPF's metric is cost.

Default route

0.0.0.0/0 or ::/0 — matches any destination not covered by a more specific route.

OSPF

Open Shortest Path First — link-state interior routing protocol using cost and the SPF algorithm.

DR / BDR

Designated and backup designated routers that reduce adjacencies on multi-access OSPF segments.

Passive interface

OSPF interface whose subnet is advertised but on which no hellos are sent.

KNOWLEDGE CHECK

22. What problem does a first-hop redundancy protocol solve?

- ☐ A Routing loops between OSPF areas
- ☐ B Duplex mismatches on access ports
- ☐ C Hosts losing their only default gateway
- ☐ D Broadcast storms caused by layer 2 loops

23. Two VRRP routers have priorities 120 and 100. Which becomes master?

- ☐ A The router with priority 120
- ☐ B The router with priority 100
- ☐ C The router with the lower IP address
- ☐ D Whichever router booted first

24. What is the virtual MAC address for VRRP group 10 (IPv4)?

- ☐ A 00-00-5E-00-01-10
- ☐ B 01-00-5E-00-00-12
- ☐ C 00-00-5E-00-01-0A
- ☐ D FF-FF-FF-FF-FF-0A

25. Which address do hosts configure as their default gateway in a VRRP design?

- ☐ A The master router's real interface IP
- ☐ B The backup router's real interface IP
- ☐ C The VRRP multicast address 224.0.0.18
- ☐ D The shared virtual IP

26. Why track the uplink interface of a VRRP master?

- ☐ A So it sends advertisements more often
- ☐ B So it hands over mastership when it loses its way out
- ☐ C So it can use a second virtual IP address
- ☐ D So hosts receive a new DHCP lease automatically

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 4 · ROUTE

Answer key

Knowledge checks

- 1. D** 172.16.4.0/22 covers 172.16.4.0 – 172.16.7.255, so it contains .5.9 and is longer than /16. The /24 covers only 172.16.4.x, and the default is used only when nothing else matches.
- 2. C** Same prefix, different sources: the lower AD (static 1 vs OSPF 110) wins. Load sharing happens only within one source with equal metrics, and conflicts are resolved by AD.
- 3. D** The second number is the metric. 110 is the AD, the age appears later in the line, and bandwidth is not shown here.
- 4. C** TTL prevents endless routing loops; at 0 the packet is discarded and the sender is informed. Routers never reset TTL or flood packets.
- 5. B** Local routes mark the exact addresses configured on the router's interfaces. They are not from LLDP, not a load-balancing marker, and not legacy entries.
- 6. A** A default route is prefix 0.0.0.0 with mask 0.0.0.0, followed by the next hop. The other options reverse the fields, use invented syntax or give a /32 mask.
- 7. B** The backup must have a worse (higher) AD than the primary. Equal AD would load-share, 0 is reserved for connected routes, and 255 means never installed.
- 8. A** Link-local addresses are unique only per link, so the router needs the interface to know which link. They can be pinged, one next hop is enough, and they are unicast.
- 9. D** Without tracking, the primary stays installed while its interface is up. Lowering the backup's AD would make it primary, a host route does not detect remote failure, and IPv6 is unrelated.
- 10. C** One-way success points to a missing return route. DNS is not used for ping by IP, PortFast does not drop replies, and a host route at the branch would not stop return traffic.
- 11. B** Full means the LSDBs match. 2-Way is bidirectional hello contact, ExStart negotiates the exchange, and Loading is still requesting LSAs.
- 12. C** With no manual ID, the highest up loopback address is used: 10.255.0.9. Physical addresses are considered only when no loopback exists.
- 13. D** DROTHERs fully adjoint only the DR and BDR, so 2-Way between them is expected. MTU problems stall in ExStart/Exchange, and timer or authentication mismatches prevent any neighbour state.
- 14. A** 100 Mbps ÷ 10,000 Mbps is below 1, and cost cannot go below 1. 10 would require a 100 Gbps reference, 100 a 1 Tbps reference, and fractional costs are not allowed.
- 15. A** With equal priority 1, the highest router ID wins. Neighbour count and interface speed are not election criteria.
- 16. B** MTU mismatches break database description exchange. Descriptions and hostnames are irrelevant, and process IDs are locally significant and need not match.
- 17. D** $255.255.255.255 - 255.255.255.252 = 0.0.0.3$. The second is the subnet mask, 0.0.0.4 is the block size, and 0.0.0.252 inverts the wrong bits.
- 18. C** Passive interfaces stay in OSPF (their prefix is advertised) but form no neighbours. They are not removed, not shut down, and the DR role is unrelated.
- 19. B** `default-information originate` injects the existing default into OSPF. Redistributing connected routes adds subnets not a default, a null0 default would discard traffic, and passive-interface stops hellos.
- 20. A** On point-to-point network types there is no DR/BDR, so the role shows as a dash. FULL proves the neighbour is up, every router has an ID, and authentication is not shown here.
- 21. D** Process IDs only identify the OSPF instance inside one router. Area IDs must match, not process IDs, and there is no even-number rule.
- 22. C** FHRPs keep the gateway address alive if one router fails. Routing loops, duplex problems and broadcast storms are handled by routing design, autonegotiation and STP.
- 23. A** Highest priority wins in VRRP. IP address only breaks ties (highest wins), and boot order does not override priority while preemption is on.
- 24. C** The group number is written in hex in the last byte: 10 = 0x0A. 0x10 would be group 16, 01-00-5E... is an IPv4 multicast MAC, and FF-... is broadcast.
- 25. D** Hosts point to the virtual IP so failover is invisible. Real addresses tie hosts to one router, and 224.0.0.18 is where VRRP advertisements go.
- 26. B** Tracking lowers priority on uplink failure, letting a better-connected router preempt. It does not change timers, add virtual IPs or trigger DHCP renewals.

Scenarios & case studies

Scenario · The unexpected path

A is best. Correct. .140 falls inside .128/25, which is more specific than the OSPF /24. Remove or correct the /25 if it is no longer wanted.

Scenario · Ping times out one way

C is best. Correct. Without a return route, HQ1 sends replies to its internet default. Routing must work in both directions.

Scenario · Neighbours that never meet

A is best. Correct. Timer mismatches cause hellos to be rejected; matching them lets the routers reach Full.

Scenario · A subnet that will not advertise

D is best. Correct. The statement matches .128–.255, not the interface address in .0–.127, so Gi0/0 never joined OSPF.

Case study · Replacing static routes across three sites

Static routes pointed only at the Dhaka–Khulna circuit, and nothing recalculated a path when it failed. With **OSPF in area 0** on all three routers, the failure removes the link from every LSDB and SPF moves traffic over Khulna–Rajshahi–Dhaka within seconds.

Design: manual **router IDs** on loopbacks (e.g. 1.1.1.1, 2.2.2.2, 3.3.3.3), WAN links set to **point-to-point** network type, **passive** LAN interfaces, a consistent **reference bandwidth** on all routers and, if one circuit is faster, costs that prefer it. Dhaka keeps its static default to the ISP and uses `default-information originate`, so both warehouses learn an OSPF default route that also fails over automatically.

Scenario · Master with no way out

B is best. Correct. With priority $110 - 20 = 90$, R2 (100) preempts and hosts' gateway moves to the router that can still reach the network.

05

MODULE 5 · SERVE

IP Services

DHCP, DNS, NAT, time, monitoring, secure management and QoS — the services every packet depends on.

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Configure and troubleshoot DHCP servers and relays, and explain DNS resolution and record types.
- Configure static NAT, dynamic NAT and PAT and read a translation table confidently.
- Set up NTP, SNMPv3 and syslog so devices are observable and their logs are trustworthy.
- Manage devices securely with SSH and choose safe protocols for file transfer.
- Explain QoS marking, queuing, policing and shaping and read a basic QoS policy.

LESSONS

1. DHCP and DNS — Addresses and Names on Demand · 30 min
2. NAT and PAT — Sharing Public Addresses · 30 min
3. NTP, SNMP and Syslog — Time, Health and Evidence · 30 min
4. Secure Management and File Transfer · 25 min
5. QoS Basics — Protecting Voice and Video · 30 min

MODULE 5 · SERVE · LESSON 1 · 30 MIN

DHCP and DNS — Addresses and Names on Demand

How a host gets its IP settings automatically and turns names into addresses.

- Walk through the four DHCP messages (DORA) and explain why the first two are broadcasts.
- Configure a router as a DHCP server and as a DHCP relay with `ip helper-address`.
- Describe recursive DNS resolution and choose the right DNS record type for a job.

In the **ZCNA Packet Journey** the packet has a wire, an address, a switch path and a route. Now it needs **services**. The first two it meets are almost invisible when they work and painfully obvious when they fail: **DHCP** gives a host its IP settings, and **DNS** turns the name a user types into an address the network can route. A large share of "the network is down" tickets are really DHCP or DNS problems, so an associate engineer must be able to prove quickly whether they are working.

DORA — a client joins VLAN 10 and asks for an address

01

Discover

Client → broadcast. Source 0.0.0.0:68, destination 255.255.255.255:67. "Is there a DHCP server out there?"

→

02

Offer

Server → client. Proposes 192.168.10.21/24, gateway, DNS servers and a lease time.

→

03

Request

Client → broadcast. "I accept the offer from this server" — broadcast so any other servers that made offers withdraw them.

→

04

Ack

Server → client. Confirms the lease; the client configures the address and (on most systems) probes it to detect a duplicate before using it.

What DHCP hands out (common options)

- IP address and subnet mask (the lease).
- Default gateway (option 3) and DNS server list (option 6).
- Domain name (option 15) and lease duration (option 51).
- Special-purpose options, for example option 43 or 60 for vendor-specific data, or option 150/66 to point IP phones at a file server.

INSIGHT · LEASES ARE BORROWED, NOT OWNED

A client tries to **renew** directly with its server at 50% of the lease (T1) and, if that fails, **rebinds** by broadcast at 87.5% (T2). If the lease expires with no answer the host must stop using the address. Short leases suit busy guest Wi-Fi; longer leases suit stable office desktops.

Router as a DHCP server

R1 · IOS-style CLI

```
! Reserve addresses for the gateway, servers and printers
R1(config)# ip dhcp excluded-address 192.168.10.1 192.168.10.20
R1(config)# ip dhcp pool VLAN10-STAFF
R1(dhcp-config)# network 192.168.10.0 255.255.255.0
R1(dhcp-config)# default-router 192.168.10.1
R1(dhcp-config)# dns-server 192.168.10.5 192.168.10.6
R1(dhcp-config)# domain-name branch.example
R1(dhcp-config)# lease 7
R1(dhcp-config)# end
R1# show ip dhcp binding
IP address      Client-ID/Hardware address    Lease expiration           Type
192.168.10.21   0100.5056.a1b2.c3             Oct 26 2026 09:14 AM      Automatic
```

Exclusions are configured globally, outside the pool. Forgetting them is the classic cause of duplicate-address conflicts with statically addressed devices.

DHCP Discover is a broadcast, and routers do not forward broadcasts. In most real networks the DHCP server sits in a data centre or server VLAN, not on every user subnet. The fix is a **DHCP relay agent**: the router interface that faces the clients listens for DHCP broadcasts, rewrites them as **unicast** to the server, and fills in the **giaddr** (gateway IP address) field with its own interface address. The server uses giaddr to pick the correct scope and sends the reply back to the relay.

DHCP relay on the client-facing interface

R2 · IOS-style CLI

```
! Clients live on VLAN 30; the DHCP server is 10.0.50.10
R2(config)# interface GigabitEthernet0/1.30
R2(config-subif)# encapsulation dot1Q 30
R2(config-subif)# ip address 192.168.30.1 255.255.255.0
R2(config-subif)# ip helper-address 10.0.50.10
R2(config-subif)# end
```

The helper goes on the interface where the client broadcasts arrive — never on the interface facing the server.

Watching DORA from the client

Linux host

```
$ sudo dhclient -v eth0
DHCPDISCOVER on eth0 to 255.255.255.255 port 67 interval 3
DHCPOFFER of 192.168.30.44 from 192.168.30.1
DHCPREQUEST for 192.168.30.44 on eth0 to 255.255.255.255 port 67
DHCPACK of 192.168.30.44 from 192.168.30.1
bound to 192.168.30.44 -- renewal in 43200 seconds.
```

The offer appears to come from 192.168.30.1 — the relay — even though the real server is in another subnet.

WATCH OUT · APIPA IS A SYMPTOM, NOT AN ADDRESS PLAN

A Windows or macOS host showing a **169.254.x.x** address got no DHCP reply and assigned itself a link-local address. Check the access port VLAN, the trunk, the helper address, the server scope and whether the pool is exhausted — in that order, following the packet.

DNS is a distributed, hierarchical database. A host asks its configured **resolver** (often a local DNS server) for `portal.branch.example`. If the resolver has no cached answer it performs **recursive resolution**: it asks a **root** server, which refers it to the **TLD** servers for `.example`, which refer it to the **authoritative** servers for

`branch.example`, which return the answer. The resolver caches the result for the record's **TTL** and replies to the host.

DNS record types an associate must recognise

RECORD	HOLDS	TYPICAL USE
A	Name → IPv4 address	<code>www</code> → 203.0.113.10
AAAA	Name → IPv6 address	<code>www</code> → 2001:db8:10::10
CNAME	Alias → another name	<code>intranet</code> → <code>portal.branch.example</code>
MX	Mail exchanger for a domain, with preference	Where email for the domain is delivered
PTR	Address → name (reverse lookup)	Log readability, some mail checks
NS	Authoritative name servers for a zone	Delegation from the parent zone
TXT	Free text	Domain verification, email policies (SPF, DMARC)
SOA	Start of authority: zone serial and timers	Zone transfers and caching behaviour

Testing resolution

Linux host and R1

```
$ dig portal.branch.example A +short
192.168.10.80
$ dig -x 192.168.10.80 +short
portal.branch.example.
! On the router: let it resolve names, or stop it trying
R1(config)# ip name-server 192.168.10.5
R1(config)# no ip domain-lookup
```

`no ip domain-lookup` stops the router treating a mistyped command as a hostname and waiting for DNS. Use it on lab devices; on production devices that must resolve names, configure `ip name-server` instead.

PRO TIP · PORTS TO REMEMBER

DHCP: UDP 67 (server) and 68 (client). DNS: UDP 53 for most queries, **TCP 53** for zone transfers and large responses. Firewalls that allow only UDP 53 can break DNS in subtle ways.

SCENARIO · DECIDE · HELPDESK CALL: NO ADDRESSES ON THE NEW FLOOR

A new third-floor VLAN 40 (192.168.40.0/24) was added yesterday. Every PC on the floor shows a 169.254.x.x address. The DHCP server at 10.0.50.10 has a scope for 192.168.40.0/24 with plenty of free addresses. Ports are in VLAN 40 and the trunk to the router allows VLAN 40. What should you check next?

- ☐ A Whether the router subinterface for VLAN 40 has `ip helper-address 10.0.50.10`.
- ☐ B Whether the DHCP server's own interface has a helper address.
- ☐ C Whether the lease time on the server is too long.
- ☐ D Whether DNS has an A record for each PC.

Circle your choice. Answer and feedback: module answer key.

CASE STUDY · MIRPUR CLINIC: MONDAY MORNING, NO NETWORK

A small clinic in Mirpur, Dhaka, has 35 PCs on VLAN 10 and 12 IP phones on VLAN 20. Over the weekend the ISP replaced the edge router and restored its configuration from an old backup. On Monday, staff PCs obtain addresses but

cannot open any website; the phones have 169.254.x.x addresses. The internal DHCP server (10.0.50.10) and DNS server (10.0.50.11) are healthy, and pinging 10.0.50.11 from a PC succeeds.

1. Why do the PCs get addresses while the phones do not?

2. What could stop PCs that can reach the DNS server from opening websites?

3. Which three show commands would you run first on the new router, and what are you looking for?

Model answer: module answer key.

TEMPLATE DHCP scope plan

Plan every scope once, centrally, so exclusions and options are consistent and reviewable.

VLAN / SUBNET	GATEWAY	EXCLUDED RANGE	DNS SERVERS	LEASE	RELAY (HELPER) ON
10 Staff – 192.168.10.0/24	192.168.10.1	.1–.20	10.0.50.11, 10.0.50.12	7 days	R1 Gi0/1.10
20 Voice – 192.168.20.0/24	192.168.20.1	.1–.10	10.0.50.11	1 day	R1 Gi0/1.20

KNOWLEDGE CHECK

1. Why is the DHCP Request message sent as a broadcast even though the client already knows the server's address?

- ☐ A Because the client has no MAC address until the Ack arrives

☐ C Because routers forward broadcasts more reliably than unicasts

☐ B So every server that made an offer learns which one was accepted

☐ D Because DHCP uses TCP and needs a broadcast handshake first

2. Where must `ip helper-address` be configured?

- ☐ A On the DHCP server's network interface

☐ C On the router interface that faces the Internet service provider

☐ B On every access switch port in the VLAN

☐ D On the router interface that receives the client broadcasts

3. Which DNS record lets a log server show a hostname instead of 10.0.5.20?

- ☐ A An MX record for the organisation's domain

☐ C A PTR record in the reverse zone

☐ B A CNAME pointing to the address

☐ D An NS record for the host

4. A PC shows 169.254.33.7 /16. What does this most directly tell you?

- ☐ A The PC sent DHCP requests and received no usable reply
- ☐ B The DHCP server assigned it a link-local reservation
- ☐ C The DNS server could not resolve the PC's name
- ☐ D The default gateway rejected the PC's ARP request

5. A resolver has no cached answer for `shop.example`. Which server does it contact first in a full recursive lookup?

- ☐ A The authoritative server for `shop.example` directly
- ☐ B The DHCP server that supplied its settings
- ☐ C A root server, which refers it to the TLD servers
- ☐ D The client's default gateway for the answer

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 5 · SERVE · LESSON 2 · 30 MIN

NAT and PAT — Sharing Public Addresses

Translate private addresses at the edge with static NAT, dynamic NAT and overload.

- Use the four NAT address terms correctly when reading a translation table.
- Configure static NAT, a dynamic NAT pool and PAT on an edge router.
- Troubleshoot common NAT faults with show and clear commands.

Branch offices use RFC 1918 private addresses (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16) inside, but Internet routers will not route them. **Network Address Translation (NAT)** rewrites the source address of outbound packets at the edge and reverses the change for returning traffic. **Port Address Translation (PAT)**, also called **NAT overload**, also rewrites ports so hundreds of inside hosts can share a single public IPv4 address.


The four NAT address names

TERM	MEANING	IN OUR EXAMPLE
Inside local	Address of the inside host as seen inside	192.168.10.21
Inside global	Address of the inside host as seen outside (after translation)	203.0.113.1
Outside local	Address of the outside host as seen from inside	198.51.100.80
Outside global	Real address of the outside host on the Internet	198.51.100.80

DEFINITION · LOCAL VS GLOBAL, INSIDE VS OUTSIDE

Inside/outside says *whose* address it is (our host or the remote one). **Local/global** says *where you are looking from* (inside the edge router or outside it). With ordinary source NAT, outside local and outside global are the same.

PAT in action — one public address, many hosts



Three ways to translate

Static NAT

- One inside address ↔ one public address, permanently
- Allows inbound connections to servers
- Uses one public IP per server

Dynamic NAT

- Inside hosts borrow addresses from a public pool
- One public IP per active host
- Pool exhaustion blocks new hosts

PAT (overload)

- Many hosts share one address using unique ports
- Default choice for branch Internet access
- Inbound connections need static or port-forward entries

Static NAT plus PAT on a branch edge

EDGE1 · IOS-style CLI

```
! ISP gave us 203.0.113.0/29; .1 is our outside interface, .6 the ISP gateway
EDGE1(config)# interface GigabitEthernet0/0
EDGE1(config-if)# ip address 203.0.113.1 255.255.255.248
EDGE1(config-if)# ip nat outside
EDGE1(config-if)# interface GigabitEthernet0/1
EDGE1(config-if)# ip address 192.168.10.1 255.255.255.0
EDGE1(config-if)# ip nat inside
EDGE1(config-if)# exit
! Static NAT: publish the web server
EDGE1(config)# ip nat inside source static 192.168.10.10 203.0.113.2
! PAT: everyone else shares the interface address
EDGE1(config)# access-list 1 permit 192.168.10.0 0.0.0.255
EDGE1(config)# ip nat inside source list 1 interface GigabitEthernet0/0 overload
EDGE1(config)# ip route 0.0.0.0 0.0.0.0 203.0.113.6
```

Three ingredients every time: mark inside and outside interfaces, define who is translated (the ACL), and define what they become (static address, pool or interface).

Dynamic NAT with a pool

EDGE1 · IOS-style CLI

```
! Alternative to PAT: a dynamic pool of three public addresses
EDGE1(config)# ip nat pool PUBLIC-POOL 203.0.113.3 203.0.113.5 netmask 255.255.255.248
EDGE1(config)# ip nat inside source list 1 pool PUBLIC-POOL
! Add the keyword overload to the line above to PAT across the pool
```

Without `overload`, the fourth simultaneous inside host gets no translation and its traffic is dropped.

Verifying translations

EDGE1 · IOS-style CLI

```
EDGE1# show ip nat translations
Pro Inside global      Inside local      Outside local     Outside global
tcp  203.0.113.1:51515   192.168.10.21:51515 198.51.100.80:443 198.51.100.80:443
tcp  203.0.113.1:1024    192.168.10.37:51515 198.51.100.80:443 198.51.100.80:443
---  203.0.113.2         192.168.10.10      ---              ---
EDGE1# show ip nat statistics
Total active translations: 3 (1 static, 2 dynamic; 2 extended)
Outside interfaces: GigabitEthernet0/0
Inside interfaces: GigabitEthernet0/1
EDGE1# clear ip nat translation *
```

Two PCs used the same source port 51515, so the router kept it for the first and chose another port (1024) for the second. The static entry shows -- because it exists even when idle. `clear` removes dynamic entries only.

WATCH OUT · THE USUAL NAT FAULTS

Interfaces marked inside/outside the wrong way round; an ACL that does not match the real inside subnet; a missing default route; or a pool on a subnet the ISP does not route to you. Check the interface roles first — `show ip nat statistics` lists them.

INSIGHT · NAT IS NOT A FIREWALL — AND IPV6 RARELY NEEDS IT

PAT blocks unsolicited inbound connections only as a side effect. Use a real stateful firewall policy for protection. With IPv6 there are enough global addresses that translation is normally unnecessary; security comes from firewall policy instead.

SCENARIO · DECIDE · A SERVER THAT NOBODY OUTSIDE CAN REACH

A Dhaka office publishes its booking server 192.168.10.10 as 203.0.113.2 with static NAT. Staff inside browse the Internet fine through PAT. Customers report the server is unreachable. `show ip nat translations` shows the static entry. What is the most likely remaining problem?

- ☐ A Inbound traffic to 203.0.113.2 is not permitted by the edge firewall or ACL policy.
- ☐ B PAT must be removed because static NAT and PAT cannot coexist.
- ☐ C The inside interface needs `ip nat outside` as well.
- ☐ D The server needs a public address configured on its own NIC.

Circle your choice. Answer and feedback: module answer key.

TEMPLATE NAT translation plan

Record every published service and shared pool before touching the edge router, so changes are reviewable.

SERVICE / GROUP	INSIDE LOCAL	INSIDE GLOBAL	TYPE	PORTS	OWNER / TICKET
Booking web server	192.168.10.10	203.0.113.2	Static	TCP 443	IT-OPS / CHG-1042
Staff Internet access	192.168.10.0/24	203.0.113.1 (interface)	PAT	All	IT-OPS / CHG-1042

KNOWLEDGE CHECK

6. In `show ip nat translations`, which column holds the private address of your own PC?

- ☐ A Inside global
- ☐ B Inside local
- ☐ C Outside local
- ☐ D Outside global

7. A branch has 120 users and one public IPv4 address. Which configuration lets all of them browse at the same time?

- ☐ A One static NAT entry for each of the 120 users
- ☐ B A dynamic pool containing that single address, no overload
- ☐ C Marking the outside interface as `ip nat inside`
- ☐ D An ACL-matched inside source list with `overload`

8. Which command shows which interfaces the router treats as inside and outside?

- ☐ A `show ip nat statistics`
- ☐ B `show ip nat translations`
- ☐ C `show ip interface brief`
- ☐ D `show running-config | include pool`

9. Why does `show ip nat translations` list the static entry even when nobody is using the server?

- ☐ A The server sends keepalives that refresh the entry
- ☐ B Static entries are created by the first inbound packet
- ☐ C Static entries are configured, not created by traffic
- ☐ D The router copies it from the DHCP binding table

10. PCs on 192.168.20.0/24 cannot reach the Internet, but 192.168.10.0/24 can. NAT uses `access-list 1 permit 192.168.10.0 0.0.0.255`. What is the fix?

- ☐ A Change the outside interface to `ip nat inside`
- ☐ B Configure a second default route for VLAN 20
- ☐ C Enable `ip helper-address` on the VLAN 20 interface
- ☐ D Add a permit for 192.168.20.0 0.0.0.255 to ACL 1

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 5 · SERVE · LESSON 3 · 30 MIN

NTP, SNMP and Syslog — Time, Health and Evidence

Synchronise clocks, poll and receive device health, and collect logs you can trust.

- Explain NTP stratum and configure devices as NTP clients.
- Compare SNMPv2c with SNMPv3 and configure a secure SNMPv3 user.
- Interpret syslog severity levels and send the right ones to a server.

When something breaks at 02:14, you will rebuild the story from logs on ten devices. If their clocks disagree by four minutes, the story is fiction. **NTP** makes clocks agree, **SNMP** lets a monitoring system ask devices how they are (and devices shout when something changes), and **syslog** sends a timestamped record of events to a central server. Together they turn a network from a black box into something you can observe.

The observability trio

NTP

Time

UDP 123. Clients sync to servers arranged by **stratum**: 0 is a reference clock (GPS, atomic), 1 is a server attached to it, each hop adds one, and 16 means unsynchronised.

SNMP

Health

UDP 161 for manager polls (get, set) and UDP 162 for agent notifications (traps and informs). Data is organised in a **MIB** and addressed by **OIDs**.

LOG

Events

Syslog, traditionally UDP 514. Each message carries a facility, a **severity** **0–7** and text. Central storage survives a device reboot or theft.

Router as an NTP client (Bangladesh Standard Time, UTC+6)

R1 · IOS-style CLI

```
R1(config)# ntp server 10.0.99.123 prefer
R1(config)# ntp server 10.0.99.124
R1(config)# clock timezone BST 6 0
R1(config)# service timestamps log datetime msec localtime show-timezone
R1(config)# end
R1# show ntp status
Clock is synchronized, stratum 3, reference is 10.0.99.123
R1# show ntp associations
  address      ref clock      st  when  poll reach  delay  offset  disp
*~10.0.99.123  192.0.2.10     2   37   64  377   1.20   0.315   0.92
~10.0.99.124  192.0.2.11     2   21   64  377   1.45   0.402   1.03
```

The `*` marks the peer currently selected. Reach 377 (octal) means the last eight polls all succeeded. Many teams instead keep every device on UTC to avoid time-zone confusion across regions — pick one standard and apply it everywhere.

PRO TIP · TWO SOURCES, AUTHENTICATED

Point devices at **at least two** internal NTP servers, and have those servers use several upstream sources. Enable NTP authentication where supported so an attacker cannot shift your clocks and invalidate certificates or logs.

SNMPv2c vs SNMPv3

SNMPv2c

- Authenticates with a **community string**
- Community travels in clear text
- No encryption of data
- Still common for read-only polling on isolated management networks

SNMPv3

- Named **users** grouped with access rights
- Message integrity and authentication (for example SHA)
- Optional encryption (for example AES) — `priv` level
- The recommended choice for anything new

SNMPv3 security levels

LEVEL	AUTHENTICATION	ENCRYPTION
noAuthNoPriv	Username only	No
authNoPriv	Yes (hash-based)	No
authPriv	Yes (hash-based)	Yes

SNMPv3 with authentication and encryption

R1 · IOS-style CLI

```
! Only the monitoring server may query
R1(config)# ip access-list standard NMS-ONLY
R1(config-std-nacl)# permit host 10.0.99.50
R1(config-std-nacl)# exit
R1(config)# snmp-server group NMS-GRP v3 priv access NMS-ONLY
R1(config)# snmp-server user nms-poller NMS-GRP v3 auth sha <auth-pass> priv aes 128 <priv-pass>
R1(config)# snmp-server host 10.0.99.50 version 3 priv nms-poller
R1(config)# snmp-server enable traps
```

SNMPv3 users are deliberately not shown in the running configuration. Store the passphrases in your password vault when you create them. `snmp-server enable traps` with no keywords turns on many trap types — in production, enable only the ones your monitoring actually uses.

Syslog severity levels

LEVEL	NAME	EXAMPLE
0	Emergency	System unusable
1	Alert	Immediate action needed (e.g. temperature critical)
2	Critical	Critical hardware or software condition
3	Error	Interface or process error
4	Warning	Configuration or resource warning
5	Notification	Normal but significant (e.g. interface up/down, config change)
6	Informational	Routine events, ACL log hits
7	Debugging	Output of debug commands

EXAMPLE · A MEMORY HOOK

Every **Alert** Can **Escalate** When **Nobody** Investigates **Daily** — Emergency, Alert, Critical, Error, Warning, Notification, Informational, Debugging. Lower number = more severe. Configuring a level sends **that level and everything more severe**.

Sending logs to a server

R1 · IOS-style CLI

```
R1(config)# logging host 10.0.99.60
R1(config)# logging trap notifications
R1(config)# logging buffered 32768 informational
R1(config)# logging source-interface Loopback0
R1(config)# end
R1# show logging | include UPDOWN
Oct 19 10:15:32.441 BST: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to down
Oct 19 10:15:33.442 BST: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/2, changed state to down
```

`logging trap notifications` sends levels 0–5 to the server. The message format is `FACILITY-SEVERITY-MNEMONIC`: `%LINK-3-UPDOWN` is facility `LINK` at severity 3 (Error).

SCENARIO · DECIDE · LOGS THAT DON'T LINE UP

During an outage review, the firewall log shows a WAN link failure at 14:02:10, but the router log shows its interface going down at 14:06:45. Engineers are arguing about which device failed first. What should be fixed so this never happens again?

- ☐ (A) Configure every device to synchronise with the same internal NTP servers and use one time-zone standard.
- ☐ (B) Raise the router's logging level to debugging.
- ☐ (C) Switch the firewall to SNMPv3.
- ☐ (D) Increase the logging buffer on the router.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

11. An NTP client reports stratum 16. What does that mean?
 - ☐ (A) It is directly attached to a GPS reference clock
 - ☐ (B) The device is not synchronised to any valid time source
 - ☐ (C) It is sixteen hops from the best server and still accurate
 - ☐ (D) It is acting as the master clock for the network
12. A router has `logging trap warnings`. Which message is sent to the syslog server?
 - ☐ (A) A level 3 error about an interface
 - ☐ (B) A level 5 notification of an interface state change
 - ☐ (C) A level 6 informational ACL hit
 - ☐ (D) A level 7 debug message
13. Why is SNMPv3 at the authPriv level preferred over SNMPv2c?
 - ☐ (A) It authenticates users and encrypts the SNMP payload
 - ☐ (B) It switches SNMP to TCP so that polls and replies are never lost
 - ☐ (C) It removes the need for a monitoring server
 - ☐ (D) It sends traps on UDP 161 instead of 162
14. In `%LINEPROTO-5-UPDOWN`, what does the 5 represent?
 - ☐ (A) The interface slot number 5
 - ☐ (B) Severity level 5, Notification
 - ☐ (C) Facility number 5 in the MIB
 - ☐ (D) Five repeated occurrences of the event
15. Which port pair is correct for SNMP?
 - ☐ (A) TCP 161 for polling, TCP 162 for traps and informs
 - ☐ (B) UDP 123 for polling, UDP 514 for traps and informs
 - ☐ (C) UDP 161 for polling, UDP 162 for traps and informs
 - ☐ (D) UDP 514 for polling, UDP 161 for traps

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 5 · SERVE · LESSON 4 · 25 MIN

Secure Management and File Transfer

Replace Telnet with SSH and move images and configurations safely.

- Configure SSH version 2 for remote management and disable Telnet.
- Choose between TFTP, FTP, SCP and SFTP for configuration and image transfers.
- Follow a safe workflow to back up configurations and stage a software image.

Every device you manage is reached over the network — so the management path is itself an attack surface. **Telnet** (TCP 23) sends the username, password and every command in clear text; anyone who can capture traffic on the path can read them. **SSH** (TCP 22) encrypts the session and authenticates the server with a host key. For a ZCNA engineer the rule is simple: **SSHv2 only, local or AAA accounts, and a management ACL**.

Enabling SSH on an IOS-style device

1. **Name the device** A hostname and domain name are needed to label the RSA key pair.
2. **Generate keys** Create an RSA key of at least 2048 bits; this enables the SSH server.
3. **Force version 2** SSHv1 has known weaknesses and must not be used.
4. **Create an account** Use `secret` (hashed) rather than `password`.
5. **Lock down the VTY lines** Authenticate locally (or via AAA), allow SSH only, set an idle timeout and apply a source ACL.

SSH for remote management

SW1 · IOS-style CLI

```
SW1(config)# hostname SW1
SW1(config)# ip domain-name branch.example
SW1(config)# crypto key generate rsa modulus 2048
SW1(config)# ip ssh version 2
SW1(config)# username netadmin privilege 15 algorithm-type scrypt secret <strong-passphrase>
SW1(config)# line vty 0 15
SW1(config-line)# login local
SW1(config-line)# transport input ssh
SW1(config-line)# exec-timeout 10 0
SW1(config-line)# access-class MGMT-ONLY in
SW1(config-line)# end
SW1# show ip ssh
SSH Enabled - version 2.0
```

`transport input ssh` disables Telnet on those lines. The `MGMT-ONLY` standard ACL (built in the Guard module) limits which source addresses may even attempt to log in.

First connection from an admin workstation

Linux host

```
$ ssh netadmin@10.0.99.11
The authenticity of host '10.0.99.11' can't be established.
RSA key fingerprint is SHA256:Qe3f...truncated...
Are you sure you want to continue connecting (yes/no)? yes
netadmin@10.0.99.11's password:
SW1#
```

The first-connection prompt is your chance to verify the host key fingerprint against your inventory. If it changes later without a planned rebuild, stop and investigate.

Moving files to and from devices

PROTOCOL	TRANSPORT	AUTHENTICATION	ENCRYPTION	USE IT FOR
TFTP	UDP 69	None	None	Isolated lab or a locked-down management VLAN only
FTP	TCP 21 (control) and 20 or a passive data port	Username/password in clear text	None	Legacy servers; avoid on untrusted paths
SCP	TCP 22 (over SSH)	SSH accounts or keys	Yes	Copying images and configs to and from devices
SFTP	TCP 22 (over SSH)	SSH accounts or keys	Yes	File servers and automation pipelines
HTTPS	TCP 443	Accounts or tokens	Yes	Downloads from controllers and repositories

Backups and image staging

R1 · IOS-style CLI

```
! Back up the running configuration to a TFTP server in the management VLAN
R1# copy running-config tftp:
Address or name of remote host []? 10.0.99.70
Destination filename [r1-config]? r1-2026-10-19.cfg
!!
2847 bytes copied in 0.412 secs
! Stage a new image securely with SCP
R1# copy scp://netadmin@10.0.99.70/images/new-image.bin flash:
R1# show flash: | include new-image
```

Each **!** during a TFTP copy is a successful block. Before booting a new image, compare its hash with the value published by the software provider and confirm there is enough flash for both old and new images.

WATCH OUT · NEVER UPGRADE WITHOUT A WAY BACK

Keep the previous image on flash until the new one is proven, save the configuration before the change, schedule a maintenance window, and make sure you have **out-of-band console access** — a remote upgrade that fails without console access means a long drive to the site.

In-band vs out-of-band management

In-band

- Management traffic shares the production network
- Simple and cheap
- If the network breaks, you may lose access exactly when you need it

Out-of-band

- Separate management network or console server
- Cellular or separate ISP backup is common
- Access survives production outages and misconfigurations

SCENARIO · DECIDE · CHOOSING A TRANSFER METHOD

A junior engineer must copy a new image to 40 branch routers across the WAN. The only server ready is a TFTP server in head office. The WAN crosses an ISP network. What should the team do?

- ☐ A Use SCP or SFTP from an internal server so transfers are authenticated and encrypted, then verify each image hash.
- ☐ B Use TFTP because it is fastest and needs no accounts.
- ☐ C Use FTP with a strong password.
- ☐ D Email the image to each branch manager to load from USB.

Circle your choice. Answer and feedback: module answer key.

Secure management baseline (per device)

- ☐ SSH version 2 enabled with a 2048-bit (or stronger) key; Telnet disabled on all VTY lines.
- ☐ Local fallback account uses a hashed `secret`; shared accounts avoided where AAA is available.
- ☐ VTY lines protected by a management source ACL and an idle timeout.
- ☐ NTP synchronised to two internal servers; logs timestamped.
- ☐ Syslog sent to the central server at the agreed level.
- ☐ SNMPv3 authPriv only; no default or guessable community strings.
- ☐ Configuration backed up after every change, stored off-device.
- ☐ Console or out-of-band access tested this quarter.

KNOWLEDGE CHECK

16. Which command on the VTY lines stops Telnet while still allowing SSH?

- | | |
|--|--|
| ☐ A <code>ip ssh version 2</code> | ☐ B <code>login local</code> |
| ☐ C <code>ip ssh authentication-retries 2</code> | ☐ D <code>transport input ssh</code> |

17. Why must a hostname and domain name be set before generating the RSA key on an IOS-style device?

- | | |
|--|---|
| ☐ A They are used to name the key pair the SSH server uses | ☐ B They become the default SSH username and password |
| ☐ C They tell the device which DNS server to trust | ☐ D They are required for TFTP backups to work |

18. Which protocol transfers files with no authentication at all?

- | | |
|------------------------------|-------------------------------|
| ☐ A SFTP | ☐ B SCP |
| ☐ C TFTP | ☐ D HTTPS |

19. A remote router upgrade fails and the device will not boot the new image. What would have prevented a site visit?

- | | |
|---|---|
| ☐ A A longer SSH idle timeout | ☐ B Tested out-of-band console access to the router |
| ☐ C A larger logging buffer | ☐ D An SNMPv2c community string |

☐ Lesson complete · Date _____ · Knowledge check score ____ / 4

MODULE 5 · SERVE · LESSON 5 · 30 MIN

QoS Basics — Protecting Voice and Video

Why congestion hurts real-time traffic and how classification, marking and queuing help.

- Describe bandwidth, delay, jitter and loss and their effect on applications.
- Classify and mark traffic using DSCP and 802.1p CoS values.
- Explain queuing, policing and shaping, and read a basic QoS policy.

Most of the time a network has spare capacity and every packet is forwarded immediately. **Quality of Service (QoS)** matters in the moments when it does not: a 1 Gbps LAN feeds a 50 Mbps WAN link, everyone joins a video call at 10:00, or a backup starts during business hours. Queues build, and the network must decide **which packets wait and which are dropped**. Without a policy that decision is first-come, first-served — and real-time voice loses.

Four characteristics QoS manages

BW

Bandwidth

How much capacity a class of traffic can use on a link.

DLY

Delay

Time from sender to receiver. For interactive voice, aim for well under 150 ms one-way.

JIT

Jitter

Variation in delay between packets. Voice and video need it low and steady (a common target is under 30 ms).

LOSS

Loss

Packets dropped. TCP recovers by retransmitting; real-time UDP media simply sounds or looks broken.

The QoS toolkit, in the order a packet meets it

01

Classify

Identify the traffic: by interface, ACL, application or existing marking

→

02

Mark

Write a value into the header so later devices need not reclassify

→

03

Police / shape

Enforce or smooth a rate at the edge

→

04

Queue

Place packets in per-class queues when the link is congested

→

05

Schedule / drop

Decide which queue sends next; drop early to avoid collapse

DEFINITION · TRUST BOUNDARY

The point in the network where you start believing markings. Usually the access switch port: you may trust an IP phone's markings, but you re-mark whatever a user PC claims. Anything beyond the boundary uses the marks already set.

Where the marks live

MARKING	HEADER	BITS	VALUES	SCOPE
CoS / PCP	802.1Q VLAN tag (layer 2)	3	0–7	Lost when the frame leaves the trunk or is routed
DSCP	IP header DS field (layer 3)	6	0–63	Travels end to end with the packet (unless re-marked)

Common DSCP values

TRAFFIC	PER-HOP BEHAVIOUR	DSCP NAME	DECIMAL
Voice media	Expedited Forwarding	EF	46
Interactive video	Assured Forwarding class 4, low drop	AF41	34
Business-critical data	Assured Forwarding class 2, low drop	AF21	18
Routing and network control	Class Selector 6	CS6	48
Everything else	Default (best effort)	DF / CS0	0
Bulk / scavenger	Class Selector 1 (lower than best effort)	CS1	8

Working out DSCP numbers

MEASURE	FORMULA	WHAT IT TELLS YOU
Assured Forwarding	$DSCP = 8 \times \text{class} + 2 \times \text{drop precedence}$	AF31 = $8 \times 3 + 2 \times 1 = 26$; AF43 = $8 \times 4 + 2 \times 3 = 38$. A higher drop precedence is dropped first inside the same class.
Class Selector	$DSCP = 8 \times \text{class}$	CS3 = 24, CS6 = 48. These map to the old 3-bit IP precedence values.

Policing vs shaping

Policing

- Measures traffic against a rate
- Excess is dropped or re-marked immediately
- No added delay, but can cause TCP retransmissions
- Typical at a provider's ingress or on untrusted traffic

Shaping

- Buffers excess and sends it later
- Smooths bursts to match a contracted rate
- Adds delay but avoids provider drops
- Typical on the customer side of a WAN link

Queuing concepts

- **FIFO** — one queue, first in first out; no differentiation.
- **Priority queue (low-latency queuing)** — a strict-priority queue for voice, always served first but capped by a policer so it cannot starve others.
- **Class-based weighted fair queuing** — each class gets a guaranteed share of bandwidth during congestion.
- **Congestion avoidance (e.g. WRED)** — drops some lower-priority TCP packets *before* the queue is full so flows slow down gracefully instead of all at once.

A modular QoS policy on the WAN edge

WAN1 · IOS-style CLI

```
WAN1(config)# class-map match-any VOICE
WAN1(config-cmap)# match dscp ef
WAN1(config-cmap)# class-map match-any VIDEO
WAN1(config-cmap)# match dscp af41
WAN1(config-cmap)# policy-map WAN-EDGE-OUT
WAN1(config-pmap)# class VOICE
WAN1(config-pmap-c)# priority percent 20
WAN1(config-pmap-c)# class VIDEO
WAN1(config-pmap-c)# bandwidth percent 30
WAN1(config-pmap-c)# class class-default
WAN1(config-pmap-c)# fair-queue
WAN1(config-pmap-c)# interface GigabitEthernet0/0
WAN1(config-if)# service-policy output WAN-EDGE-OUT
WAN1# show policy-map interface GigabitEthernet0/0 output
```

Three building blocks: **class-map** (what), **policy-map** (what to do), **service-policy** (where and in which direction). Queuing policies act on **output** — the point where congestion actually happens.

PRO TIP · QOS CANNOT CREATE BANDWIDTH

QoS chooses who suffers during congestion; it does not make a link faster. If a link is congested all day, the answer is more capacity or less traffic — QoS just keeps voice and critical apps usable until then.

SCENARIO · DECIDE · CHOPPY CALLS AT THE BRANCH

A Chattogram branch has a 50 Mbps WAN link. Calls break up every afternoon when staff upload scanned documents to head office. Phones mark voice EF, and the access switches trust the phones. The WAN router has no QoS policy. What is the best first step?

- ☐ A Apply an outbound policy on the WAN interface with a priority queue for EF traffic.
- ☐ B Re-mark all traffic to EF so everything gets priority.
- ☐ C Apply the policy inbound on the LAN interface.
- ☐ D Change phone marking from DSCP to CoS only.

Circle your choice. Answer and feedback: module answer key.

Key terms

DORA

Discover, Offer, Request, Acknowledge — the four-message DHCP exchange that leases an address to a client.

TTL (DNS)

How long a resolver may cache a DNS answer before asking again.

Inside global

The public address that represents an inside host to the outside world after NAT.

MIB / OID

The structured database of values an SNMP agent exposes, and the numeric identifier of each value.

Syslog severity

A level from 0 (Emergency) to 7 (Debugging) attached to each log message.

DHCP relay

A router function (`ip helper-address`) that forwards client DHCP broadcasts as unicast to a server in another subnet, adding the giaddr field.

PAT

Port Address Translation (NAT overload): many inside hosts share one public address, distinguished by port numbers.

Stratum

An NTP server's distance from a reference clock; 16 means unsynchronised.

SNMP trap

An unsolicited notification an agent sends to a manager on UDP 162; an inform is the acknowledged version.

SCP / SFTP

File transfer protocols that run over SSH, giving authentication and encryption.

DSCP

The 6-bit field in the IP header used to mark a packet's QoS class end to end.

Trust boundary

The network point where QoS markings start being accepted rather than rewritten.

Shaping

Buffering excess traffic to send it later at a target rate, smoothing bursts.

REFLECT

Think of a network you use daily — a campus, office or home. Which IP service (DHCP, DNS, NAT, NTP, logging or QoS) do you think fails or misbehaves most often there, and how would you prove it with one command?

KNOWLEDGE CHECK

20. Which DSCP value is normally used for voice media?

- | | |
|--|---|
| ☐ A AF41, decimal 34 | ☐ B CS6, decimal 48 |
| ☐ C CS1, decimal 8 | ☐ D EF, decimal 46 |

21. What is the decimal DSCP value of AF32?

- | | |
|----------------------------|----------------------------|
| ☐ A 26 | ☐ B 30 |
| ☐ C 34 | ☐ D 28 |

22. Why is CoS alone not enough to protect voice across a routed WAN?

- | | |
|---|---|
| ☐ A CoS is in the 802.1Q tag, which is lost when routed | ☐ B CoS has more values than DSCP, which confuses routers on the path |
| ☐ C CoS is only supported on wireless networks | ☐ D CoS is applied after queuing, so it is too late |

23. A provider drops all traffic above the contracted 50 Mbps. Which customer-side tool avoids those drops?

- | | |
|---|--|
| ☐ A Policing inbound traffic to 50 Mbps | ☐ B Marking all traffic CS1 |
| ☐ C Shaping outbound traffic to 50 Mbps | ☐ D Increasing the DHCP lease time |

24. Where should a queuing policy for a congested WAN link be applied?

- | | |
|--|--|
| ☐ A Inbound on the WAN interface | ☐ B Outbound on the WAN interface |
| ☐ C Inbound on every access port | ☐ D On the DHCP server's interface |

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 5 · SERVE

Answer key

Knowledge checks

- 1. B** Broadcasting the Request tells all servers which offer was chosen so the others can release theirs. The client always has a MAC address, routers do not forward broadcasts at all, and DHCP runs over UDP, not TCP.
- 2. D** The relay listens on the client-facing interface and converts the broadcast to unicast. Server, access-port and Internet-facing interfaces never see the client broadcast.
- 3. C** Reverse lookups map addresses to names using PTR records. MX is for mail, a CNAME aliases one name to another name, and NS names authoritative servers.
- 4. A** A 169.254.0.0/16 address is self-assigned when DHCP fails. Servers do not hand out link-local addresses, and DNS or ARP failures do not change how the host is addressed.
- 5. C** Recursive resolution starts at the root, follows referrals to the TLD and then the authoritative server. The resolver does not yet know the authoritative server; DHCP and the gateway play no part in name resolution.
- 6. B** Inside local is the inside host's address as seen inside — its private address. Inside global is the translated public address; the outside columns describe the remote host.
- 7. D** PAT (overload) multiplexes many hosts onto one address using ports. Static entries need one address each, a one-address pool without overload serves one host at a time, and reversed interface roles break NAT.
- 8. A** The statistics output lists inside and outside interfaces plus hit counters. The translation table shows entries, `show ip interface brief` shows status only, and the pool filter shows pool definitions.
- 9. C** Static NAT entries exist as soon as they are configured. They are not refreshed by keepalives, not created by traffic and have no relationship with DHCP bindings.
- 10. D** Only traffic matched by the NAT ACL is translated, so the new subnet must be added. The interface roles and default route already work for VLAN 10, and a helper address relates to DHCP, not NAT.
- 11. B** Stratum 16 is the NTP value for unsynchronised. Stratum 0/1 sit closest to a reference clock, valid strata stop at 15, and a master clock advertises a low stratum.
- 12. A** The configured level sends itself and everything more severe (lower numbers): 0–4. Levels 5, 6 and 7 are less severe and are filtered out.
- 13. A** authPriv adds per-user authentication and encryption. SNMP still uses UDP, still needs a manager, and traps still use UDP 162.
- 14. B** The format is FACILITY-SEVERITY-MNEMONIC, so 5 is the severity. It is not a slot, a MIB facility number or a repeat counter.
- 15. C** Managers poll agents on UDP 161 and agents send notifications to UDP 162. UDP 123 is NTP and UDP 514 is syslog; SNMP does not normally use TCP.
- 16. D** `transport input ssh` limits the protocols accepted on the lines. `ip ssh version 2` selects the SSH version, `login local` chooses where credentials are checked, and `ip ssh authentication-retries` limits failed login attempts.
- 17. A** The key pair is labelled from the host and domain names. They are not credentials, do not configure DNS servers and have nothing to do with TFTP.
- 18. C** TFTP has neither authentication nor encryption. SFTP and SCP run over SSH and HTTPS uses TLS, all of which authenticate.
- 19. B** Out-of-band console access works even when the device cannot boot normally. Timeouts, log buffers and SNMP all depend on the device already running.
- 20. D** Voice uses Expedited Forwarding (46). AF41 is typical for interactive video, CS6 for network control, and CS1 for scavenger traffic below best effort.
- 21. D** $AF = 8 \times \text{class} + 2 \times \text{drop precedence} = 8 \times 3 + 2 \times 2 = 28$. 26 is AF31, 30 is AF33 and 34 is AF41.
- 22. A** The layer 2 tag, and its CoS bits, disappear at the router. CoS has fewer values (0–7) than DSCP, works on wired trunks, and is a marking that queuing reads rather than something applied afterwards.
- 23. C** Shaping buffers bursts so the provider never sees excess. Policing drops excess itself, marking alone does not limit rate, and DHCP is unrelated.
- 24. B** Queues form when packets wait to leave a slower interface, so queuing is applied outbound there. Inbound and access-port policies classify or police but cannot schedule the congested link.

Scenarios & case studies

Scenario · Helpdesk call: no addresses on the new floor

A is best. Correct. The scope exists and layer 2 is fine, so the broadcast is reaching the router but going no further. A new client subnet needs its own helper address.

Case study · Mirpur clinic: Monday morning, no network

The phones' VLAN 20 subinterface was probably missing from the old backup or lacks `ip helper-address 10.0.50.10`, so their DHCP broadcasts never reach the server. PCs can reach internal DNS, so local routing works; the Internet failure points to the edge: NAT may be missing or have reversed inside/outside roles, the NAT ACL may not match 192.168.10.0/24 (old addressing), or the default route to the ISP may be wrong. First commands: `show running-config interface` for each subinterface (helper addresses), `show ip nat statistics` (inside/outside roles and hits) with `show ip nat translations`, and `show ip route` (a default route via the new ISP next hop). Lesson: a restored backup must be compared against the current design before go-live.

Scenario · A server that nobody outside can reach

A is best. Correct. The translation exists, so the next thing on the packet's path is the security policy on the outside interface or upstream firewall.

Scenario · Logs that don't line up

A is best. Correct. Without common time, ordering events across devices is guesswork; NTP plus a consistent time zone (or UTC) gives a single timeline.

Scenario · Choosing a transfer method

A is best. Correct. Secure transport plus hash verification protects both confidentiality and integrity of the image.

Scenario · Choppy calls at the branch

A is best. Correct. Markings already exist; what is missing is a queuing decision where congestion occurs — the WAN egress.

06

MODULE 6 · AIR

Wireless Networking

RF fundamentals, 802.11 generations, WLAN architectures, roaming and security — then a real design.

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Explain RF behaviour, calculate dBm and SNR, and choose bands and channels sensibly.
- Describe 802.11 generations through Wi-Fi 7 and the SSID, BSSID, BSS and ESS building blocks.
- Compare autonomous, controller-based and cloud-managed WLANs and explain client roaming.
- Select WPA2/WPA3 personal or enterprise security and describe 802.1X roles.
- Plan, document and validate a small office WLAN.

LESSONS

1. RF Fundamentals — The Physics Behind Wi-Fi · 30 min
2. 802.11 Standards and WLAN Building Blocks · 25 min
3. AP Modes, WLAN Architectures and Roaming · 30 min
4. Wireless Security — WPA2, WPA3 and 802.1X · 30 min
5. Designing a Small Office WLAN · 30 min

MODULE 6 · AIR · LESSON 1 · 30 MIN

RF Fundamentals — The Physics Behind Wi-Fi

Frequency, power, channels, interference and signal-to-noise ratio in plain language.

- Convert between milliwatts and dBm and apply the 3 dB and 10 dB rules.
- Calculate signal-to-noise ratio and judge whether a reading is good enough for the task.
- Compare the 2.4, 5 and 6 GHz bands and pick non-overlapping channels.

Up to now the packet has travelled on copper and fibre, where a cable guarantees a private, predictable path. The **Air** stage is different: Wi-Fi shares radio spectrum with every nearby device, signals weaken with distance and walls, and only one transmitter on a channel can talk at a time. A good wireless engineer thinks first about **radio frequency (RF)** behaviour and only then about configuration.

RF vocabulary

- **Frequency** — how many times per second the wave cycles, in hertz. Wi-Fi uses the 2.4 GHz, 5 GHz and 6 GHz bands.
- **Wavelength** — the physical length of one cycle. Higher frequency means shorter wavelength, which generally means less range and poorer wall penetration.
- **Amplitude / power** — the strength of the signal, measured in milliwatts (mW) or decibel-milliwatts (dBm).
- **Channel** — a defined slice of a band (for example 20, 40, 80, 160 or 320 MHz wide) that a radio uses.
- **Half duplex** — on one channel, one radio transmits at a time; Wi-Fi uses **CSMA/CA** (listen, back off, then send) to share the air.

Power and quality numbers

MEASURE	FORMULA	WHAT IT TELLS YOU
dBm	$\text{dBm} = 10 \times \log_{10}(\text{power in mW})$	1 mW = 0 dBm, 10 mW = 10 dBm, 100 mW = 20 dBm. Received Wi-Fi signals are tiny, so they are negative: -65 dBm is stronger than -80 dBm.
Rule of 3s and 10s	+3 dB $\approx$ $\times 2$ power; +10 dB = $\times 10$ power	Going from -70 to -67 dBm roughly doubles received power; from -70 to -60 dBm is ten times more.
SNR	$\text{SNR (dB)} = \text{signal (dBm)} - \text{noise floor (dBm)}$	Signal -65 dBm with noise -92 dBm gives 27 dB. Higher SNR lets radios use faster modulation.
Channel centre	2.4 GHz: $2407 + 5 \times \text{channel}$ MHz; 5 GHz: $5000 + 5 \times \text{channel}$ MHz	Channel 6 = 2437 MHz; channel 36 = 5180 MHz. Useful when a tool shows frequency instead of channel number.

Rule-of-thumb targets

USE	SIGNAL (RSSI)	SNR
Basic browsing and email	-70 dBm or better	20 dB or more
Voice and video calls	-67 dBm or better	25 dB or more
Dense, high-throughput areas	-65 dBm or better	25–30 dB or more

WATCH OUT · RSSI IS NOT THE WHOLE STORY

A strong signal on a busy, noisy channel can perform worse than a weaker signal on a clean one. Always read signal, noise, channel utilisation and retry rates together.

The three Wi-Fi bands

2.4 GHz

- Best range and wall penetration
- Only three non-overlapping 20 MHz channels: **1, 6, 11**
- Crowded: neighbours, Bluetooth, microwave ovens
- Use 20 MHz channels only

5 GHz

- Many more 20 MHz channels (varies by country)
- Some channels need **DFS** (radar detection) and may move
- Shorter range than 2.4 GHz
- 40 or 80 MHz channels practical in many offices

6 GHz

- Large clean spectrum where allowed (up to 1200 MHz)
- Only Wi-Fi 6E and Wi-Fi 7 clients can use it
- WPA3 or Enhanced Open is mandatory
- Rules and power levels vary by country

INSIGHT · WHY 1, 6 AND 11?

A 2.4 GHz channel is about 20 MHz wide, but channel centres are only 5 MHz apart. Channels 1, 6 and 11 are the set that do not overlap. Using channel 3 next to a neighbour on channel 1 creates **adjacent-channel interference**, which is worse than simply sharing channel 1.

What weakens or distorts a signal

EFFECT	WHAT HAPPENS	TYPICAL CAUSE
Free-space path loss	Signal spreads out and weakens with distance	Every link — it is unavoidable
Absorption	Material soaks up energy	Concrete, brick, water, people, dense foliage
Reflection	Signal bounces off a surface	Metal shelves, lift shafts, tinted glass
Multipath	Reflected copies arrive at different times	Warehouses, factories; modern MIMO can exploit it
Co-channel interference	Radios on the same channel must take turns	Too many APs on one channel
Non-Wi-Fi interference	Energy that is not Wi-Fi raises the noise floor	Microwave ovens, cordless phones, video senders

Reading RF from a client

Linux laptop

```
$ iw dev wlan0 link
Connected to 3c:84:6a:11:22:33 (on wlan0)
    SSID: Office-Staff
    freq: 5180
    signal: -61 dBm
    tx bitrate: 433.3 MBit/s VHT-MCS 9 80MHz short GI VHT-NSS 1
$ nmcli -f BSSID,SSID,CHAN,SIGNAL,SECURITY dev wifi list
BSSID      SSID          CHAN  SIGNAL  SECURITY
3C:84:6A:11:22:33 Office-Staff  36    82      WPA3
3C:84:6A:11:22:34 Office-Guest  36    80      OWE
58:D9:C3:44:55:66 Neighbour-Net 6     47      WPA2
```

`freq: 5180` is channel 36. `nmcli` shows `SIGNAL` as a 0–100 quality percentage, not dBm; use `iw` when you need the real value. Notice one AP radio advertises two SSIDs with two BSSIDs.

SCENARIO · DECIDE · SLOW WI-FI NEAR THE CANTEEN

Users near an office canteen complain that Wi-Fi becomes unusable around lunchtime. The AP there uses 2.4 GHz channel 6 and 5 GHz channel 44. Most complaining laptops are connected on 2.4 GHz. Signal is -58 dBm. What is the most likely cause and fix?

- ☐ A Microwave ovens raising the 2.4 GHz noise floor; steer capable clients to 5 GHz and reduce reliance on 2.4 GHz.
- ☐ B The AP is too far away; add a second AP on channel 6.
- ☐ C Switch the 2.4 GHz radio to channel 3 to sit between the others.
- ☐ D Increase AP transmit power to maximum.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

1. A client reads -72 dBm signal with a -95 dBm noise floor. What is the SNR?

- ☐ A 23 dB
- ☐ B -23 dB
- ☐ C 167 dB
- ☐ D 72 dB

2. Which set of 2.4 GHz channels does not overlap?

- ☐ A 1, 5 and 9
- ☐ B 1, 6 and 11
- ☐ C 2, 7 and 12
- ☐ D 1, 3 and 5

3. Transmit power rises from 10 dBm to 20 dBm. How has power in milliwatts changed?

- ☐ A It has doubled, from 10 mW to 20 mW
- ☐ B It has risen by 10 mW, from 10 mW to 20 mW
- ☐ C It is ten times higher, from 10 mW to 100 mW
- ☐ D It is a hundred times higher, from 1 mW to 100 mW

4. Why can a laptop that only supports Wi-Fi 6 not see a 6 GHz-only SSID?

- ☐ A The SSID is hidden on 6 GHz by default
- ☐ B Wi-Fi 6 clients support only WPA2, which 6 GHz SSIDs do not offer
- ☐ C 6 GHz uses 2.4 GHz channels for discovery
- ☐ D 6 GHz operation requires Wi-Fi 6E or newer hardware

5. What does DFS mean for a 5 GHz channel?

- ☐ A The channel is reserved for 802.11b clients only
- ☐ B The AP must transmit at double power to overcome radar
- ☐ C The AP must leave the channel if it detects radar
- ☐ D The channel is used only for mesh backhaul links

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 6 · AIR · LESSON 2 · 25 MIN

802.11 Standards and WLAN Building Blocks

From 802.11a/b/g to Wi-Fi 7, and what SSIDs, BSSIDs and service sets really are.

- Place each 802.11 amendment in its band and explain the key technology it added.
- Distinguish SSID, BSSID, BSS, ESS and the distribution system.
- Describe how a client discovers, authenticates and associates to an AP.

IEEE 802.11 is the family of standards for wireless LANs; the **Wi-Fi Alliance** certifies products and gives the generations friendly names such as Wi-Fi 6. Every new generation keeps backward compatibility on the bands it shares, but older clients on a channel slow everyone down because they use more airtime to send the same data.

802.11 generations at concept level

AMENDMENT	NAME	BANDS	KEY ADDITIONS	THEORETICAL MAX
802.11b	Legacy	2.4 GHz	DSSS, up to 11 Mbps	11 Mbps
802.11a / g	Legacy	5 GHz / 2.4 GHz	OFDM, up to 54 Mbps	54 Mbps
802.11n	Wi-Fi 4	2.4 and 5 GHz	MIMO, 40 MHz channels, channel bonding	600 Mbps
802.11ac	Wi-Fi 5	5 GHz	80/160 MHz channels, 256-QAM, downlink MU-MIMO	≈6.9 Gbps
802.11ax	Wi-Fi 6 / 6E	2.4, 5 (+6 GHz for 6E)	OFDMA, uplink and downlink MU-MIMO, 1024-QAM, BSS colouring, target wake time	≈9.6 Gbps
802.11be	Wi-Fi 7	2.4, 5 and 6 GHz	320 MHz channels (6 GHz), 4096-QAM, multi-link operation (MLO)	≈46 Gbps

WATCH OUT · THEORETICAL MAXIMUMS ARE MARKETING NUMBERS

They assume the widest channel, every spatial stream and perfect conditions. A typical phone with two spatial streams in a busy office sees a small fraction. Design for **airtime and capacity**, not headline speed.

Technologies worth knowing

MIMO

Multiple-input multiple-output

Several antennas send parallel **spatial streams**, multiplying throughput for capable clients.

MU

Multi-user MIMO

The AP sends to (or receives from) several clients at the same moment using separate spatial streams.

OFDMA

Orthogonal frequency-division multiple access

Wi-Fi 6 splits a channel into resource units so several clients share one transmission — great for many small packets.

QAM

Modulation density

Higher QAM packs more bits per symbol but needs a cleaner signal (higher SNR).

MLO

Multi-link operation

Wi-Fi 7 lets a device use links on more than one band together for throughput or reliability.

TWT

Target wake time

Clients and AP agree when to wake, saving battery for phones and IoT sensors.

Service set building blocks (largest first)

ESS — Extended Service Set	Several BSSs sharing one SSID	Joined by a wired distribution system	Lets clients roam
BSS — Basic Service Set	One AP radio and its associated clients	Identified by a BSSID	
BSSID	A MAC address for each SSID on each radio	Different on every AP and band	
SSID	The network name users see	Up to 32 bytes	Same on every AP in an ESS

DEFINITION · DISTRIBUTION SYSTEM (DS)

The wired network that connects APs to each other and to the rest of the LAN. In an ESS the DS is usually your switched Ethernet network, with each SSID mapped to a VLAN.

Infrastructure vs ad hoc


Infrastructure mode

- Clients talk through an AP
- Central security and VLAN mapping
- Normal for every business WLAN

Independent BSS (ad hoc)

- Clients talk directly to each other
- No AP, no central control
- Rare; mostly temporary or device-to-device links

How a client joins a WLAN



Three kinds of 802.11 frame

TYPE	PURPOSE	EXAMPLES
Management	Join, leave and advertise networks	Beacon, probe, authentication, association, deauthentication
Control	Help deliver frames on a shared medium	ACK, RTS, CTS, block ACK
Data	Carry user traffic	Your encrypted IP packets

PRO TIP · FEWER SSIDS, MORE AIRTIME

Each SSID sends its own beacons on every radio, typically about ten per second, at the lowest basic rate. Five SSIDs on one channel can waste a noticeable share of airtime before any user sends data. Aim for three or four SSIDs at most and separate users by policy (VLAN, role) instead.

SCENARIO · DECIDE · MIXED CLIENTS ON ONE AP

A training centre's 5 GHz radio serves 30 modern laptops and four very old barcode scanners that only support 802.11a. Throughput for everyone drops sharply when the scanners are busy. What is the best long-term fix?

- ☐ A Replace or isolate the legacy scanners, for example on a separate SSID or band with a planned retirement date.
- ☐ B Increase the channel width to 160 MHz.
- ☐ C Hide the SSID so the scanners cannot find it.
- ☐ D Add a second SSID for the laptops on the same radio.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

6. What uniquely identifies a single SSID on a single AP radio?
 - ☐ A Its BSSID
 - ☐ B Its ESS name
 - ☐ C Its channel number
 - ☐ D Its VLAN ID
7. Which Wi-Fi 6 feature lets one transmission carry data for several clients using different resource units?
 - ☐ A Channel bonding
 - ☐ B OFDMA
 - ☐ C Target wake time
 - ☐ D DSSS
8. Which frame type does an AP send about ten times per second to advertise each SSID?
 - ☐ A ACK, a control frame
 - ☐ B RTS, a control frame
 - ☐ C Probe response, a data frame
 - ☐ D Beacon, a management frame
9. A user's phone joins the office SSID and then receives an IP address. Which step must happen just before DHCP on a WPA2-Personal network?
 - ☐ A An 802.1X exchange with a RADIUS server
 - ☐ B A DNS lookup for the AP's name
 - ☐ C A spanning-tree convergence on the AP port
 - ☐ D The 4-way handshake that derives encryption keys
10. Which statement about Wi-Fi 7 is correct?
 - ☐ A It works only in the 6 GHz band and ignores 2.4 and 5 GHz
 - ☐ B It can use links on several bands at once (MLO)
 - ☐ C It removes the need for WPA3 security
 - ☐ D It replaces Ethernet as the distribution system

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 6 · AIR · LESSON 3 · 30 MIN

AP Modes, WLAN Architectures and Roaming

Autonomous, controller-based and cloud-managed APs — and how clients move between them.

- Compare autonomous, controller-based and cloud-managed WLAN architectures.
- Match common AP operating modes to their purpose.
- Explain client roaming and the role of 802.11k, 802.11v and 802.11r.

A single home router with built-in Wi-Fi is easy to manage. Fifty APs across three buildings are not — channels, power, security and VLANs must stay consistent, and users expect to walk from one end of the site to the other on a video call. The answer is an **architecture** that centralises management and a design that supports **roaming**.

Three WLAN architectures

Autonomous

- Each AP configured and managed individually
- Traffic bridged locally to its own VLANs
- Fine for one to three APs
- Hard to keep consistent at scale

Controller-based

- Lightweight APs managed by a WLAN controller (appliance, virtual or embedded)
- APs tunnel to the controller using **CAPWAP** (control UDP 5246, data UDP 5247)
- Central RF management, roaming and policy
- Traffic can be tunnelled centrally or switched locally

Cloud-managed

- APs managed from a cloud dashboard
- User data stays on the local network; only management goes to the cloud
- Fast multi-site rollout and monitoring
- Depends on Internet reachability for changes and visibility

DEFINITION · SPLIT-MAC

In controller-based designs, time-critical 802.11 functions (beacons, ACKs, encryption on many platforms) stay on the AP, while management functions (authentication policy, RF management, roaming decisions) live on the controller. The AP and controller split the work.

A lightweight AP comes online

01

Power and address

AP boots via PoE and gets an IP address by DHCP



02

Find a controller

Uses DHCP options, DNS or a configured address to discover controllers



03

Join

Builds an encrypted CAPWAP control tunnel and downloads image and configuration



04

Serve clients

Starts beaconing SSIDs; client data tunnelled or switched locally per policy

Common AP operating modes (generic names)

MODE	SERVES CLIENTS?	PURPOSE
Local / client-serving	Yes	Normal operation; scans other channels briefly between client traffic
Local switching (remote site)	Yes	Controller manages the AP, but client traffic is bridged locally at the branch
Monitor	No	Full-time scanning for rogues, interference and attacks (wireless IDS)
Sniffer	No	Captures frames on a channel and forwards them to an analyser

MODE	SERVES CLIENTS?	PURPOSE
Bridge / mesh	Yes or backhaul only	Wireless links between APs or buildings where cabling is impossible
Sensor	No	Behaves like a client to test connectivity, DHCP, DNS and application response

How the AP's switch port is configured

Autonomous or local switching

- AP bridges each SSID to its own VLAN
- Switch port is an **802.1Q trunk**
- Native/management VLAN for the AP itself

Central tunnelling

- Client traffic rides inside CAPWAP to the controller
- Switch port is usually an **access port** in the AP management VLAN
- SSID-to-VLAN mapping happens at the controller

Access switch port for a locally switching AP

SW2 · IOS-style CLI

```
! Port for an AP that bridges Staff (VLAN 10) and Guest (VLAN 30) locally
SW2(config)# interface GigabitEthernet1/0/12
SW2(config-if)# description AP-FL00R2-EAST
SW2(config-if)# switchport mode trunk
SW2(config-if)# switchport trunk native vlan 99
SW2(config-if)# switchport trunk allowed vlan 10,30,99
SW2(config-if)# spanning-tree portfast trunk
SW2(config-if)# power inline auto
SW2# show power inline GigabitEthernet1/0/12
```

VLAN 99 carries the AP's own management address. Prune the trunk to exactly the VLANs the AP needs; `portfast trunk` avoids waiting for spanning tree when the AP reboots.

PoE budget for APs (power at the switch port)

STANDARD	NAME	MAX AT SWITCH (PSE)	TYPICAL AP FIT
802.3af	PoE (Type 1)	15.4 W	Older or small single-band APs
802.3at	PoE+ (Type 2)	30 W	Most Wi-Fi 5 and many Wi-Fi 6 APs
802.3bt	PoE++ Type 3	60 W	High-end Wi-Fi 6E and Wi-Fi 7 APs
802.3bt	PoE++ Type 4	90 W	Specialist APs, cameras with heaters, small switches

Roaming is decided by the **client**, not the AP. When the current signal drops below the client's threshold, it scans, chooses a better BSSID with the same SSID and reassociates. Good roaming needs about **15–20% cell overlap** at the target signal level (often –67 dBm for voice), consistent security settings, and the same VLAN or a controller that hides the subnet change.

Standards that help roaming

11k

802.11k — neighbour reports

The AP tells the client which nearby APs and channels to scan, so it scans less.

11v

802.11v — BSS transition

The network can suggest a better AP to a client (for example to balance load or move off a weak AP).

11r

802.11r — fast transition

Keys are pre-derived so the client skips a full 802.1X exchange when it moves — vital for voice on enterprise SSIDs.

Layer 2 vs layer 3 roaming

Layer 2 roam

- New AP maps the SSID to the same VLAN/subnet
- Client keeps its IP address
- Simple and common in one building

Layer 3 roam

- New AP's VLAN is a different subnet
- Controller anchors or tunnels traffic so the client keeps its IP
- Without that help, sessions drop and DHCP restarts

SCENARIO · DECIDE · CALLS DROP IN THE CORRIDOR

Nurses using Wi-Fi handsets on an enterprise (802.1X) SSID report that calls drop for a few seconds whenever they walk from Ward A to Ward B. Survey data shows -66 dBm at the overlap and both APs map the SSID to the same VLAN. What should you enable first?

- ☐ A 802.11r fast transition, with 802.11k neighbour reports, on the voice SSID.
- ☐ B Increase both APs to maximum power.
- ☐ C Put Ward B's AP on a different VLAN.
- ☐ D Switch the SSID to WPA2-Personal to avoid RADIUS.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

- Which protocol and ports does a lightweight AP use to talk to its controller?
 - ☐ A CAPWAP on UDP 5246 for control and UDP 5247 for data
 - ☐ B SNMP on UDP 161 for control and UDP 162 for data
 - ☐ C SSH on TCP 22 for control and TCP 23 for data
 - ☐ D RADIUS on UDP 1812 for control and UDP 1813 for data
- An AP must scan all channels full time for rogue devices and serve no clients. Which mode fits?
 - ☐ A Local mode
 - ☐ B Bridge mode
 - ☐ C Monitor mode
 - ☐ D Local switching mode
- Who makes the final decision to roam to another AP?
 - ☐ A The WLAN controller
 - ☐ B The client device
 - ☐ C The RADIUS server
 - ☐ D The access switch
- A new Wi-Fi 7 AP powers up but runs with radios disabled on an 802.3af switch. What is the likely issue?
 - ☐ A The trunk port to the AP is missing the guest VLAN
 - ☐ B The AP has no DNS server configured
 - ☐ C The port's PoE budget is too low for the AP
 - ☐ D The SSID uses WPA3 instead of WPA2
- What does a cloud-managed WLAN normally send to the cloud?
 - ☐ A All client data traffic for inspection
 - ☐ B Only DHCP requests from clients
 - ☐ C Every 802.11 beacon frame
 - ☐ D Management, configuration and monitoring data

□ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 6 · AIR · LESSON 4 · 30 MIN

Wireless Security — WPA2, WPA3 and 802.1X

Protect the air: personal vs enterprise modes, SAE, EAP and guest access.

- Explain why WEP and WPA are obsolete and what WPA2 and WPA3 changed.
- Contrast personal (passphrase) and enterprise (802.1X) authentication.
- Name the three 802.1X roles and choose a suitable EAP method.

On a cable an attacker needs a port; in the air they need only to be within range — in the car park, in the next building, on the floor above. Wireless security therefore has two jobs: **authenticate** who may join, and **encrypt** every frame so eavesdroppers learn nothing. Modern networks use **WPA2** at minimum and **WPA3** wherever clients support it.

A short history of Wi-Fi security

GENERATION	ENCRYPTION	STATUS IN 2026
Open	None	Only acceptable with Enhanced Open (OWE) or for captive-portal guest access with care
WEP	RC4 with static keys	Broken — must never be used
WPA	TKIP (RC4-based)	Deprecated — disable it
WPA2	AES-CCMP	Still widely used; acceptable with a strong passphrase or 802.1X
WPA3	AES (CCMP or GCMP), SAE, mandatory PMF	Recommended; required on 6 GHz

Personal vs enterprise

Personal (PSK / SAE)

- Everyone shares one passphrase
- No per-user identity or easy revocation
- WPA3-Personal uses **SAE**, which resists offline guessing
- Suits homes, small shops, some IoT

Enterprise (802.1X)

- Each user or device has its own credential or certificate
- Authentication by a **RADIUS** server
- Revoke one user without touching others; assign VLAN or role per user
- Standard for organisations

The three 802.1X roles

SUP

Supplicant

The client device and its software that presents credentials.

AUTH

Authenticator

The AP or WLAN controller (or a switch port on wired networks) that relays EAP and opens the port when told.

AS

Authentication server

Usually a RADIUS server that checks credentials against a directory and returns accept or reject, plus policy such as VLAN.

Enterprise join, step by step



EAP methods you will meet

METHOD	CREDENTIALS	NOTES
EAP-TLS	Certificates on both server and client	Strongest common option; needs a certificate authority and device enrolment
PEAP (with MSCHAPv2 inside)	Server certificate + username/password	Easy to deploy; clients must validate the server certificate
EAP-TTLS	Server certificate + various inner methods	Flexible; common on non-Windows clients

WATCH OUT · VALIDATE THE SERVER CERTIFICATE

If clients accept any RADIUS certificate, an attacker can run a fake AP with the same SSID and harvest password hashes. Push a profile that pins the expected certificate authority and server name.

What WPA3 adds

- **SAE** replaces the PSK exchange in personal mode, so capturing a handshake no longer allows offline dictionary attacks, and past sessions stay safe if the passphrase leaks later (forward secrecy).
- **Protected Management Frames (802.11w)** are mandatory, blocking forged deauthentication attacks.
- **WPA3-Enterprise 192-bit mode** for high-security environments.
- **Enhanced Open (OWE)**: encryption for open guest networks without a password.
- **Transition mode** allows WPA2 and WPA3 clients on one SSID during migration — at WPA2's security level for the WPA2 clients.

WPA3-Personal on an open-source AP daemon

Linux AP · hostapd.conf

```
# Small-office AP: WPA3-Personal only, 5 GHz channel 36
interface=wlan0
ssid=Office-Staff
country_code=BD
ieee80211d=1
hw_mode=a
channel=36
ieee80211n=1
ieee80211ac=1
wpa=2
wpa_key_mgmt=SAE
rsn_pairwise=CCMP
ieee80211w=2
sae_password=<long-random-passphrase>
```

`wpa=2` selects RSN (the WPA2/WPA3 frame format), `wpa_key_mgmt=SAE` makes it WPA3-Personal and `ieee80211w=2` makes management-frame protection mandatory. The country code makes the AP obey local channel and power rules.

802.1X on the same AP

Linux AP · hostapd.conf

```
# Same AP, WPA2-Enterprise using a RADIUS server
wpa=2
wpa_key_mgmt=WPA-EAP
rsn_pairwise=CCMP
ieee8021x=1
auth_server_addr=10.0.99.10
auth_server_port=1812
auth_server_shared_secret=<radius-shared-secret>
ieee80211w=1
```

The AP is the authenticator: it relays EAP to the RADIUS server and never sees the user's password when a TLS-based EAP method is used.

`ieee80211w=1` makes PMF optional so older clients can still join.

Guest and IoT access patterns

- **Guest:** separate SSID mapped to a guest VLAN with Internet-only access, client isolation, rate limits, and OWE or a captive portal with acceptable-use terms.
- **IoT:** devices that cannot do 802.1X get their own SSID and VLAN, ideally with a unique passphrase per device or group where the platform supports it, and firewall rules limited to what they need.
- **Rogue detection:** monitor-mode APs or scanning radios flag unknown APs plugged into your wired network or impersonating your SSID.

SCENARIO · DECIDE · A DHAKA START-UP OUTGROWS ITS WI-FI PASSWORD

A 60-person software company uses one WPA2-Personal passphrase that has not changed in three years. Two developers left on bad terms last month. Management wants a fix that will scale. What should you recommend?

- ☐ A Move staff to WPA3-Enterprise (802.1X) with per-user credentials or certificates from the company directory, keep a separate guest SSID.
- ☐ B Change the passphrase every month and email it to staff.
- ☐ C Hide the SSID so former staff cannot find it.
- ☐ D Filter by MAC address and keep the old passphrase.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

16. In 802.1X on a WLAN, which device is the authenticator?

- ☐ A The AP or WLAN controller
- ☐ B The laptop's Wi-Fi software
- ☐ C The RADIUS server
- ☐ D The DHCP server

17. Which WPA3-Personal feature prevents attackers cracking a captured handshake offline?

- ☐ A Simultaneous Authentication of Equals (SAE)
- ☐ B TKIP per-packet key mixing with a longer passphrase
- ☐ C MAC address filtering
- ☐ D Hiding the SSID in beacons

18. Which EAP method requires a certificate on every client device?

- ☐ A PEAP with MSCHAPv2
- ☐ B EAP-TLS
- ☐ C EAP-TTLS with PAP
- ☐ D WPA3-Personal SAE

19. A 6 GHz SSID will not come up on a new AP configured for WPA2-Personal. Why?

- ☐ A 6 GHz only supports 802.1X with EAP-TLS certificates
- ☐ B 6 GHz SSIDs cannot use AES encryption
- ☐ C 6 GHz requires WPA3 or Enhanced Open security
- ☐ D 6 GHz requires the SSID to be hidden

20. Why is mandatory PMF (802.11w) valuable?

- ☐ A It doubles throughput on crowded channels
- ☐ B It replaces the RADIUS server in enterprise mode
- ☐ C It lets WEP clients join WPA3 networks
- ☐ D It blocks forged deauthentication frames

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 6 · AIR · LESSON 5 · 30 MIN

Designing a Small Office WLAN

A practical walk-through from requirements to validation survey.

- Gather WLAN requirements for coverage, capacity, roaming and security.
- Produce a simple channel, power and SSID plan for a small site.
- Plan and record a validation survey.

Design is where everything in the Air stage comes together. A small office does not need a consultant's 80-page report, but it does need a clear answer to five questions: **who** will use the WLAN, **where**, **with what devices**, **for which applications**, and **how securely**. Answer those first and the AP count, channels and SSIDs almost choose themselves.

The ZCNA small-office WLAN walk-through

1. **Requirements** Count users and devices (most people carry two or three), list critical applications (voice, video, cloud apps) and note special areas such as meeting rooms or warehouses.
2. **Floor plan and walls** Get a scaled plan and mark wall materials. Concrete and metal change everything; plasterboard barely matters.
3. **Predictive design** Estimate AP count for both coverage (signal ≥ -67 dBm where voice is needed) and capacity (devices per radio).
4. **Channel and power plan** 2.4 GHz: 1/6/11 at 20 MHz, turn some 2.4 GHz radios off if cells overlap too much. 5 GHz: non-adjacent channels, 20–40 MHz in dense areas. Power similar to the weakest client, not maximum.
5. **SSID and security plan** Usually Staff (802.1X or WPA3-Personal), Guest (isolated) and possibly IoT; each mapped to its own VLAN.
6. **Wired and power** Cable each AP to a switch with enough PoE budget and the correct trunk or access configuration.
7. **Install and validate** Mount APs as designed (ceiling, antennas the right way) and run a post-install validation survey.

Quick capacity estimate

MEASURE	FORMULA	WHAT IT TELLS YOU
Devices	$\text{devices} = \text{users} \times \text{devices per user}$	40 users $\times$ 2.5 = 100 devices.
Active devices	$\text{active} = \text{devices} \times \text{concurrency}$	Not everyone is busy at once; at 60% concurrency, 60 active devices.
Radios needed	$\text{radios} = \text{active} \div \text{target per radio}$	With a target of 25 active devices per 5 GHz radio, $60 \div 25 = 2.4$, so plan 3 APs — then check coverage.

INSIGHT · COVERAGE VS CAPACITY

A single powerful AP might *cover* a whole floor, but 100 devices sharing its airtime will crawl. Small offices are usually **capacity-bound**: more APs at lower power beat one AP at full power.

Three survey types

Predictive

- Software model from floor plan and wall types
- Before buying hardware
- Only as good as the wall data

Passive

- Walk the site listening to beacons
- Measures real signal, noise and channel overlap
- Used for validation and troubleshooting

Active

- Client associated and passing traffic
- Measures throughput, loss, roaming
- Proves the network meets application needs

Example plan: three-AP office (single floor)

AP	LOCATION	2.4 GHZ	5 GHZ	WIDTH (5 GHZ)	POWER
AP1	Reception and meeting room	Ch 1	Ch 36	40 MHz	Medium
AP2	Open-plan centre	Off (AP1 and AP3 cover 2.4 GHz)	Ch 52	40 MHz	Medium
AP3	Back offices and store	Ch 11	Ch 149	40 MHz	Medium

Checking the PoE budget before go-live

SW2 · IOS-style CLI

```
SW2# show power inline
```

Module	Available (Watts)	Used (Watts)	Remaining (Watts)			
-----	-----	-----	-----			
1	370.0	76.5	293.5			
Interface	Admin	Oper	Power(W)	Device	Class	Max(W)
Gi1/0/10	auto	on	25.5	AP1-RECEPTION	4	30.0
Gi1/0/11	auto	on	25.5	AP2-OPENPLAN	4	30.0
Gi1/0/12	auto	on	25.5	AP3-BACKOFFICE	4	30.0

Each AP draws its PoE+ class 4 allocation. The switch still has about 293 W spare, so a fourth AP or IP phones will fit. If an AP shows **off** or a low class, check the cable and the AP's power requirement before blaming the radio configuration.

PRO TIP · TURNING 2.4 GHZ RADIOS OFF IS A FEATURE

In a compact office, three 2.4 GHz radios on three APs often hear each other everywhere, creating co-channel contention with no gain. Leaving one radio off is a common, sensible design choice — as long as 2.4-only devices still have coverage.

CASE STUDY · SYLHET TEA-GARDEN OFFICE WI-FI

A tea estate near Sylhet is replacing a single consumer router in its two-storey brick office (18 staff, 6 laptops shared in the weighing room, CCTV on a separate wired network). The office also needs Wi-Fi in a nearby tin-roofed weighing shed 40 m away, where staff use tablets to record harvest weights. Electricity is unstable, and the Internet comes from one fibre link plus a mobile-data backup. Staff complain of dead spots upstairs and that visitors use the staff password.

1. Which architecture would you choose, and why?

2. How would you cover the weighing shed?

3. What SSIDs and security would you use?

4. What power and resilience considerations matter here?

Model answer: module answer key.

TEMPLATE **WLAN survey sheet**

Record each measurement point during a validation survey so problems can be traced to a location and fixed.

POINT / ROOM	BAND / CHANNEL	STRONGEST BSSID	SIGNAL (DBM)	SNR (DB)	THROUGHPUT TEST	NOTES
Upstairs office, desk 4	5 GHz / 36	AP3 5 GHz	-63	29	180 Mbps down	Good; AP1 also heard at -72
Weighing shed, scale 2	5 GHz / 149	AP-SHED 5 GHz	-70	22	45 Mbps down	Tablet OK; check behind steel racks

WLAN go-live checklist

- ☐ Country code set correctly on every AP or the controller.
- ☐ Channel and power plan applied; no adjacent-channel use in 2.4 GHz.
- ☐ No more than three or four SSIDs; each mapped to the intended VLAN.
- ☐ Staff SSID on WPA3 (or WPA2/WPA3 transition with a migration date); WEP/TKIP disabled.
- ☐ Guest SSID isolated from internal networks and rate-limited.
- ☐ 802.11k/v/r enabled where clients support them (test old devices first).
- ☐ PoE budget confirmed; APs running with all radios enabled.
- ☐ Validation survey completed and saved with the as-built AP map.
- ☐ Monitoring and rogue detection enabled; admin access secured.

SCENARIO · DECIDE · ONE AP OR THREE?

A 35-person NGO office (one open floor, 600 m², plasterboard walls) asks you to install a single top-end Wi-Fi 7 AP at maximum power in the centre to save money. Staff make many video calls. What do you advise?

- ☐ A Install two or three APs at moderate power, sized for the number of active devices and video calls.
- ☐ B Accept the single AP at maximum power because Wi-Fi 7 has a 46 Gbps maximum.
- ☐ C Use one AP but set every client to 2.4 GHz for better range.
- ☐ D Install three APs all on the same 5 GHz channel at maximum power.

Circle your choice. Answer and feedback: module answer key.

Key terms**dBm**

Power relative to 1 milliwatt on a logarithmic scale; 0 dBm = 1 mW, -70 dBm is a typical usable Wi-Fi signal.

Co-channel interference

Contention caused by several radios sharing one channel and taking turns to transmit.

BSSID

The MAC address identifying one SSID on one AP radio.

OFDMA

Wi-Fi 6+ technique that splits a channel into resource units shared by several clients in one transmission.

CAPWAP

IETF protocol used between lightweight APs and a WLAN controller (UDP 5246 control, 5247 data).

SAE

Simultaneous Authentication of Equals: WPA3-Personal key exchange that resists offline password guessing.

PMF

Protected Management Frames (802.11w): authentication of management frames, mandatory in WPA3.

SNR

Signal-to-noise ratio: received signal minus noise floor, in dB. Higher is better.

DFS

Dynamic Frequency Selection: rules that make 5 GHz APs detect radar and vacate some channels.

ESS

Extended Service Set: several BSSs with the same SSID joined by a distribution system, enabling roaming.

MLO

Multi-link operation: Wi-Fi 7 devices using links on several bands at once.

802.11r

Fast BSS transition: pre-derived keys that make roaming between APs quick on secured SSIDs.

802.1X

Port-based network access control using EAP between supplicant, authenticator and authentication server.

REFLECT

Walk through a building you know well — your campus, office or home. Where would you put access points, which walls would cause the most trouble, and which band would you rely on for each area?

KNOWLEDGE CHECK

21. An office has 50 users with two devices each and expects 50% concurrency. Targeting 25 active devices per radio, how many 5 GHz radios are needed for capacity?

- ☐ A One
- ☐ B Four
- ☐ C Two
- ☐ D Five

22. Which survey type proves that real applications achieve the expected throughput after installation?

- ☐ A An active survey
- ☐ B A predictive survey
- ☐ C A passive survey
- ☐ D A spectrum-only survey

23. Why might a designer switch off some 2.4 GHz radios in a compact office?

- ☐ A Because 2.4 GHz cannot carry encrypted traffic
- ☐ B To cut co-channel contention from overlapping cells
- ☐ C Because 2.4 GHz requires 802.3bt power
- ☐ D Because 2.4 GHz radios disable 802.11r roaming for all clients

24. What transmit power strategy suits a small, dense office?

- ☐ A Maximum power on every AP to create the largest possible cells
- ☐ B Minimum power with one AP per building
- ☐ C Power that changes randomly each hour
- ☐ D Moderate power matched to clients, with more APs

☐ Lesson complete · Date _____ · Knowledge check score ____ / 4

MODULE 6 · AIR

Answer key

Knowledge checks

- 1. A** SNR = signal – noise = $-72 - (-95) = 23$ dB. A negative or summed value comes from subtracting in the wrong order or adding, and 72 ignores the noise floor.
- 2. B** Only 1, 6 and 11 are spaced far enough apart for 20 MHz channels in a three-channel design. The other sets overlap their neighbours.
- 3. C** +10 dB is a factor of ten, so 10 dBm (10 mW) becomes 20 dBm (100 mW). Doubling would be +3 dB, and 10 dBm is 10 mW, not 1 mW.
- 4. D** Wi-Fi 6 (802.11ax) originally covers 2.4 and 5 GHz; the 6 GHz extension is Wi-Fi 6E. SSIDs are not hidden by default, Wi-Fi 6 clients support WPA3, and 6 GHz discovery does not use 2.4 GHz channels.
- 5. C** Dynamic Frequency Selection shares some 5 GHz channels with radar, so APs must detect radar and move. It has nothing to do with 802.11b, higher power or mesh-only use.
- 6. A** Each SSID on each radio has its own BSSID, a MAC address. The ESS shares one SSID across many APs, a channel is shared by many BSSs, and a VLAN is a wired-side mapping.
- 7. B** OFDMA divides a channel into resource units for multiple clients. Channel bonding widens a channel, TWT schedules sleep, and DSSS is the old 802.11b modulation.
- 8. D** Beacons are management frames advertising the BSS. ACK and RTS are control frames for delivery, and probe responses are management (not data) frames sent only when asked.
- 9. D** With a pre-shared key the 4-way handshake completes the security setup, then encrypted data such as DHCP can flow. RADIUS is used only in enterprise mode, DNS comes after addressing, and STP is unrelated to client joins.
- 10. B** MLO is a headline Wi-Fi 7 feature. Wi-Fi 7 also uses 2.4 and 5 GHz, still relies on WPA3, and still needs a wired distribution system.
- 11. A** CAPWAP (an IETF standard) uses UDP 5246 and 5247. SNMP is for monitoring, SSH and Telnet are management sessions, and RADIUS ports are for authentication and accounting.
- 12. C** Monitor mode dedicates the radios to scanning. Local and local-switching modes serve clients, and bridge mode builds wireless links between sites or APs.
- 13. B** Roaming is client-driven; 802.11k and 802.11v let the network give hints, but the client chooses. The controller, RADIUS server and switch cannot force a normal roam.
- 14. C** 802.3af supplies up to 15.4 W, well below what modern multi-radio APs need, so they boot in a reduced-power state. VLANs, DNS and the security mode do not decide which radios get power.
- 15. D** Cloud-managed designs keep user traffic local and send management and telemetry to the cloud. They do not haul all user data, DHCP or beacons to the provider.
- 16. A** The authenticator relays EAP between client and server and enforces the result. The laptop is the supplicant, RADIUS is the authentication server, and DHCP takes no part.
- 17. A** SAE makes each guess require a live exchange with the AP. TKIP belongs to obsolete WPA, and MAC filtering and hidden SSIDs are not cryptographic protections.
- 18. B** EAP-TLS uses mutual certificate authentication. PEAP and TTLS need only a server certificate plus inner credentials, and SAE is a personal-mode handshake, not an EAP method.
- 19. C** Rules for 6 GHz forbid WPA2 and legacy security, requiring WPA3 or OWE with PMF. WPA3-Personal is allowed (not only EAP-TLS), AES is used, and SSIDs need not be hidden.
- 20. D** Protected Management Frames authenticate frames such as deauthentication. PMF does not change throughput, does not replace RADIUS and has nothing to do with WEP compatibility.
- 21. C** $50 \times 2 = 100$ devices; 50% active = 50; $50 \div 25 = 2$ radios. One radio is undersized, while four and five ignore concurrency or the per-radio target.
- 22. A** Active surveys pass real traffic while associated. Predictive surveys model before installation, passive surveys only listen to beacons, and spectrum analysis shows RF energy but not application performance.
- 23. B** With only three channels, overlapping 2.4 GHz cells mostly add contention. Encryption, PoE class and roaming support are not reasons to disable the band.
- 24. D** Clients transmit at lower power than APs; balanced, moderate cells improve uplink and capacity. Maximum power creates huge overlapping cells, one minimum-power AP leaves gaps, and random changes cause instability.

Scenarios & case studies

Scenario · Slow Wi-Fi near the canteen

A is best. Correct. Microwave ovens leak energy around 2.45 GHz. The signal is strong, so the problem is noise and interference, which 5 GHz avoids.

Scenario · Mixed clients on one AP

A is best. Correct. Slow legacy clients consume far more airtime per byte, dragging the whole BSS down. Isolating then retiring them restores capacity.

Scenario · Calls drop in the corridor

A is best. Correct. Coverage and VLAN are fine; the gap is the full 802.1X re-authentication on every roam. Fast transition removes most of that delay.

Scenario · A Dhaka start-up outgrows its Wi-Fi password

A is best. Correct. Per-user identity means leavers are disabled in one place, and the network logs who connected. Guests stay separate.

Case study · Sylhet tea-garden office Wi-Fi

A cloud-managed or small controller-based design with three or four ceiling-mounted APs (two per floor, brick walls absorb a lot) gives central management for a site with no on-site engineer; autonomous APs would be harder to keep consistent. For the shed, run outdoor-rated fibre or Ethernet if within 100 m of a switch (40 m is fine for copper, ideally in conduit with surge protection), or use a point-to-point outdoor bridge link if trenching is impossible; the metal roof means the AP must be inside or under the roof, not relying on signal from the office. SSIDs: Staff on WPA3-Personal at minimum (WPA3-Enterprise if they have a directory), a separate Guest SSID on an isolated VLAN with OWE or a rotating passphrase, and change the leaked staff passphrase. Power: a UPS for the PoE switch and router so the WLAN survives short cuts, PoE budget checked for all APs, and the mobile backup tested. Validate with a passive survey upstairs and an active test in the shed with the tablets actually used.

Scenario · One AP or three?

A is best. Correct. The space is capacity-bound; several moderate cells give more total airtime and better uplink from phones and laptops.

07

MODULE 7 · GUARD

Security Fundamentals

Threats, AAA, hardening, ACLs, layer 2 protections, VPNs, firewalls and zero trust.

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Describe threats, vulnerabilities and social engineering using the CIA triad, and recommend strong authentication.
- Configure AAA with local fallback and harden routers and switches against common attacks.
- Write, place and verify standard and extended ACLs with correct wildcard masks.
- Configure port security, DHCP snooping and dynamic ARP inspection on access switches.
- Explain IPsec VPNs, firewall and IPS roles, and the principles of zero trust.

LESSONS

1. Threats, Vulnerabilities and the Human Factor · 25 min
2. AAA and Device Hardening · 30 min
3. Access Control Lists — Filtering Traffic · 35 min
4. Layer 2 Security — Port Security, DHCP Snooping and DAI · 30 min
5. VPNs, Firewalls, IPS and Zero Trust · 30 min

MODULE 7 · GUARD · LESSON 1 · 25 MIN

Threats, Vulnerabilities and the Human Factor

The vocabulary of security, the CIA triad, social engineering and strong authentication.

- Distinguish threat, vulnerability, exploit and risk using network examples.
- Map common attacks to the part of the CIA triad they damage.
- Recommend a modern password and multi-factor authentication policy.

The **Guard** stage is where the packet meets security controls. Before configuring any of them, an engineer needs a shared vocabulary — and an honest view that many incidents start not with clever hacking but with a reused password, an unpatched device or a convincing phone call. Network security is technical **and** human.

Security vocabulary

TERM	MEANING	NETWORK EXAMPLE
Asset	Something of value to protect	Core banking server, customer data, the WAN link
Vulnerability	A weakness that could be used	Router running software with a published flaw; Telnet enabled
Threat	Someone or something that could cause harm	Criminal group, disgruntled insider, power surge
Exploit	The method or code that uses a vulnerability	A script that abuses the router flaw to gain access
Risk	Likelihood × impact of a threat using a vulnerability	High: internet-facing, unpatched, critical service
Mitigation / control	A measure that reduces risk	Patch, disable Telnet, add an ACL, monitor logs

The CIA triad

C

Confidentiality

Only authorised people can read data.
Broken by eavesdropping, stolen credentials, misconfigured access.
Protected by encryption and access control.

I

Integrity

Data is not altered without detection.
Broken by man-in-the-middle tampering or unauthorised config changes. Protected by hashes, signatures, change control.

A

Availability

Systems are usable when needed.
Broken by DoS attacks, failures, ransomware. Protected by redundancy, capacity, backups, DoS protection.

Common attacks and what they damage

ATTACK	HOW IT WORKS	MAINLY HITS
DoS / DDoS	Floods a target with traffic or requests from one or many sources	Availability
Spoofing	Fakes an IP, MAC or identity to bypass trust	Integrity, confidentiality
Man-in-the-middle	Attacker sits between two parties (e.g. ARP spoofing, rogue AP)	Confidentiality, integrity
Reconnaissance	Scans and gathers information before an attack	Prepares later attacks
Malware / ransomware	Malicious code; ransomware encrypts data for payment	Availability, confidentiality
Password attacks	Guessing, dictionary, credential stuffing with leaked passwords	Confidentiality
Layer 2 attacks	MAC flooding, rogue DHCP, ARP poisoning, VLAN hopping	Confidentiality, availability

Social engineering — attacks on people

- **Phishing** — mass emails that imitate a bank, courier or IT team to steal credentials.
- **Spear phishing / whaling** — tailored messages to a specific person or senior executive.
- **Vishing and smishing** — the same trick by voice call or SMS (very common with mobile financial services).
- **Pretexting** — a believable story: "I'm from the ISP, I need your router password to fix the outage."
- **Tailgating** — following an authorised person through a secure door.
- **Baiting** — a USB stick labelled "Salary 2026" left in the car park.

PRO TIP · A SECURITY PROGRAMME, NOT A SINGLE PRODUCT

Strong programmes combine **user awareness** (everyone recognises phishing), **training** (staff know procedures), **physical access control** (locked comms rooms, badges) and technical controls. The network team owns part of each.

Password policy: then and now

Outdated habits

- Short passwords with forced symbols and digits
- Mandatory change every 30–90 days
- Security questions for recovery
- Users write variants like Dhaka@2025, Dhaka@2026

Current good practice

- Long passphrases are the priority
- Check new passwords against lists of breached passwords
- Change only when compromise is suspected
- Allow password managers and paste
- Rate-limit and lock out repeated failures

Multi-factor authentication

KNOW

Something you know

Password, passphrase or PIN.

HAVE

Something you have

Authenticator app, hardware security key, smart card, phone.

ARE

Something you are

Fingerprint, face or other biometric — usually unlocks a device-held key.

INSIGHT · NOT ALL MFA IS EQUAL

SMS codes are better than nothing but can be intercepted or socially engineered. Push approvals can be abused by spamming prompts until a tired user taps "approve". **Phishing-resistant** methods — hardware keys and passkeys bound to the real website — are the strongest choice, especially for network administrators.

SCENARIO · DECIDE · THE FRIENDLY ISP ENGINEER

A caller says he is from your ISP, mentions your circuit ID correctly and asks the junior engineer on duty to read out the edge router's admin password so he can "fix packet loss before the CEO notices". What should the engineer do?

- ☐ A Decline, end the call and contact the ISP through the number in your contract to verify the request.
- ☐ B Give the password because he knows the circuit ID.
- ☐ C Create a temporary admin account for him instead.
- ☐ D Ask him to email the request first, then send the password by email.

Circle your choice. Answer and feedback: module answer key.

REFLECT

Think about the accounts you use for network or lab work. Which ones still rely on a password alone, and what would it take to add phishing-resistant MFA to the most important one?

KNOWLEDGE CHECK

1. Telnet is enabled on an Internet-facing router. In security terms, what is the enabled Telnet service?

☐ A A vulnerability	☐ B A threat
☐ C An exploit	☐ D A compensating control
2. A DDoS attack makes the online banking portal unreachable. Which part of the CIA triad is primarily affected?

☐ A Confidentiality	☐ B Integrity
☐ C Availability	☐ D Non-repudiation
3. An employee receives an SMS that appears to come from her mobile wallet asking her to confirm a PIN via a link. What is this?

☐ A Tailgating	☐ B Smishing
☐ C Vishing	☐ D Whaling
4. Which password practice reflects current guidance?

☐ A Force a password change every 30 days	☐ B Require exactly one symbol and one digit
☐ C Block password managers from pasting	☐ D Screen new passwords against known-breached lists
5. Which MFA method is most resistant to phishing?

☐ A A six-digit code sent by SMS	☐ B A push notification that the user taps to approve on a phone
☐ C A security question about a first pet	☐ D A hardware security key or passkey bound to the site

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 7 · GUARD · LESSON 2 · 30 MIN

AAA and Device Hardening

Centralise who can log in and what they can do, then lock down every device.

- Explain authentication, authorisation and accounting and why AAA is centralised.
- Compare RADIUS and TACACS+ and configure a device to use a AAA server with local fallback.
- Apply a device-hardening baseline to routers and switches.

With five devices, local usernames are manageable. With five hundred, they are a liability: when an engineer leaves, someone must remove their account everywhere, and nobody can prove who changed what. **AAA** moves identity and policy to central servers so every device asks the same question in the same way — and records the answer.

AAA — three questions

AuthN

Authentication

Who are you? Prove it with a password, certificate or MFA.

AuthZ

Authorisation

What are you allowed to do? Privilege level, permitted commands, VLAN or role.

Acct

Accounting

What did you do and when? Login times, commands entered, session data — evidence for audits and incidents.

RADIUS vs TACACS+ (concept level)

FEATURE	RADIUS	TACACS+
Standard	Open IETF standard	Openly documented protocol (RFC 8907)
Transport	UDP 1812 (auth) and 1813 (accounting)	TCP 49
Encryption	Only the password field is hidden	Entire payload encrypted
AAA separation	Authentication and authorisation combined	Each A handled separately
Per-command authorisation	Limited	Yes — ideal for device administration
Typical use	Network access: 802.1X, Wi-Fi, VPN users	Administrator logins to routers, switches, firewalls

INSIGHT · USE BOTH, FOR DIFFERENT JOBS

Many organisations run RADIUS for **users joining the network** (Wi-Fi, wired 802.1X, VPN) and TACACS+ for **engineers managing devices**, often on the same policy server platform.

Device administration through TACACS+ with local fallback

R1 · IOS-style CLI

```
R1(config)# aaa new-model
R1(config)# tacacs server TAC-1
R1(config-server-tacacs)# address ipv4 10.0.99.11
R1(config-server-tacacs)# key <shared-secret>
R1(config-server-tacacs)# exit
R1(config)# username breakglass privilege 15 algorithm-type scrypt secret <vaulted-secret>
R1(config)# aaa authentication login default group tacacs+ local
R1(config)# aaa authorization exec default group tacacs+ local
R1(config)# aaa accounting commands 15 default start-stop group tacacs+
R1(config)# line vty 0 15
R1(config-line)# transport input ssh
```

`local` after the group means the device uses its local account **only if the server is unreachable** — not if the server rejects the user. Keep the break-glass password in a vault and test it regularly.

WATCH OUT · LOCK-OUT RISK

Typing `aaa new-model` immediately changes how logins work. On a remote device, keep an existing session open, configure and test a new SSH login in a second window, and only then close the first session.

Device hardening baseline

- Use `enable algorithm-type scrypt secret` and hashed user secrets; `service password-encryption` only hides plain passwords weakly (type 7 is reversible).
- Enforce `security passwords min-length` and slow brute force with `login block-for`.
- SSHv2 only, management ACL on VTY lines, idle timeouts, login banner.
- Disable services you do not use, such as the plain HTTP server; use HTTPS only if web management is needed.
- Shut down unused switch ports and place them in an unused, isolated VLAN.
- Disable discovery protocols (such as LLDP) on ports facing untrusted networks.
- Keep software patched against published vulnerabilities; back up configurations off-device.
- Send logs to a central server, synchronise time, and monitor configuration changes.

Hardening commands in practice

SW1 · IOS-style CLI

```
SW1(config)# enable algorithm-type scrypt secret <strong-secret>
SW1(config)# security passwords min-length 12
SW1(config)# login block-for 120 attempts 5 within 60
SW1(config)# no ip http server
SW1(config)# banner motd ^Authorised access only. Activity is logged.^
SW1(config)# vlan 999
SW1(config-vlan)# name UNUSED-PARKING
SW1(config-vlan)# interface range GigabitEthernet1/0/20 - 24
SW1(config-if-range)# switchport mode access
SW1(config-if-range)# switchport access vlan 999
SW1(config-if-range)# shutdown
```

`login block-for 120 attempts 5 within 60` blocks new logins for two minutes after five failures within one minute. Banners should warn, not welcome — avoid words like "Welcome".

Management plane, control plane, data plane — protect each

Management plane

- SSH, SNMP, NTP, logging, APIs
- Protect with AAA, ACLs, encryption
- Compromise = attacker controls the device

Control plane

- Routing protocols, STP, ARP
- Protect with authentication and rate limits
- Compromise = traffic misdirected

Data plane

- User traffic being forwarded
- Protect with ACLs, firewalls, port security
- Compromise = data exposed or blocked

Router and switch hardening checklist

- ☐ Hostname, domain, SSHv2 with 2048-bit (or stronger) keys; Telnet off.
- ☐ AAA to TACACS+ or RADIUS with tested local break-glass account.
- ☐ Enable secret using a strong hash; minimum password length set.
- ☐ VTY lines: management ACL, idle timeout, login rate limiting.
- ☐ Unused services disabled (HTTP, unneeded discovery on edge ports).
- ☐ Unused ports shut down and in a parking VLAN.
- ☐ NTP, syslog and SNMPv3 configured to management servers.
- ☐ Software version checked against current security advisories.
- ☐ Configuration backed up and compared with the approved baseline.
- ☐ Login banner with a legal warning in place.

SCENARIO · DECIDE · LEAVER ON FRIDAY

A network engineer leaves the company on Friday. The company has 120 switches and routers. Some use local accounts only; others use TACACS+ with local fallback. Which approach removes the leaver's access most reliably and makes the next leaver easy?

- ☐ A Disable the user in the central directory/AAA server now, remove local accounts by automation, and move every device to AAA with a vaulted break-glass account.
- ☐ B Change the enable secret on every device.
- ☐ C Wait for the next quarterly audit to clean up accounts.
- ☐ D Block the leaver's laptop MAC address on the management VLAN.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

6. Which protocol encrypts the whole payload and supports per-command authorisation for device administrators?

- ☐ A RADIUS over UDP 1812
- ☐ B SNMPv2c over UDP 161
- ☐ C TACACS+ over TCP 49
- ☐ D Syslog over UDP 514

7. With `aaa authentication login default group tacacs+ local`, when is the local account used?

- ☐ A Only when no TACACS+ server responds
- ☐ B Whenever the TACACS+ server rejects the password
- ☐ C Always before the TACACS+ server is tried
- ☐ D Only for console logins, never for SSH

8. Which accounting record helps most when investigating who changed a router's routing configuration?

- ☐ A Exec accounting of session start times only
- ☐ B Command accounting for privilege level 15
- ☐ C DHCP lease accounting
- ☐ D Interface bandwidth counters

9. Why is `service password-encryption` not enough to protect passwords in a configuration?

- ☐ A It only works on switches, not routers
- ☐ B It deletes passwords after 30 days
- ☐ C Its type 7 encoding is easily reversed
- ☐ D It sends passwords to the syslog server

10. What is the safest way to enable AAA on a remote router?

- ☐ A Keep one session open and test a new login before closing it
- ☐ B Enable AAA and immediately reload the router
- ☐ C Delete the local account first so nobody can bypass the AAA servers
- ☐ D Configure AAA only through a Telnet session

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 7 · GUARD · LESSON 3 · 35 MIN

Access Control Lists — Filtering Traffic

Standard and extended ACLs, wildcard masks, placement and verification.

- Calculate wildcard masks for hosts, subnets and ranges.
- Write correct standard and extended ACLs, numbered and named.
- Place and apply ACLs in the right direction, and verify them with hit counters.

An **access control list (ACL)** is an ordered list of permit and deny statements. Routers and switches use ACLs to filter traffic, but also to **identify** traffic for NAT, QoS, VTY access and routing policy. ACL mistakes are among the most common causes of self-inflicted outages, so precision matters.

How every ACL is processed

- Entries are checked **top-down**; the **first match wins** and processing stops.
- Every ACL ends with an invisible **implicit deny any** — if nothing matched, the packet is dropped.
- An ACL with no permit statements therefore blocks everything.
- One ACL per interface, **per direction (in/out)**, **per protocol (IPv4/IPv6)**.
- Inbound ACLs check packets as they arrive, before routing; outbound ACLs check packets about to leave.

Wildcard masks

MEASURE	FORMULA	WHAT IT TELLS YOU
Rule	0 bit = must match; 1 bit = ignore	A wildcard mask tells the device which address bits to compare.
From a subnet mask	wildcard = 255.255.255.255 - subnet mask	/24 → 0.0.0.255; /27 (255.255.255.224) → 0.0.0.31; /20 (255.255.240.0) → 0.0.15.255.
Shortcuts	host A.B.C.D = A.B.C.D 0.0.0.0; any = 0.0.0.0 255.255.255.255	Use the keywords for readability.

Wildcard practice

MATCH THIS	ADDRESS	WILDCARD
Single host 10.1.1.5	10.1.1.5	0.0.0.0 (or host 10.1.1.5)
Subnet 192.168.10.0/24	192.168.10.0	0.0.0.255
Subnet 172.16.32.0/20	172.16.32.0	0.0.15.255
Subnet 10.10.4.64/26	10.10.4.64	0.0.0.63
Point-to-point 203.0.113.4/30	203.0.113.4	0.0.0.3

Standard vs extended

Standard ACL

- Matches **source address only**
- Numbers 1–99 and 1300–1999, or named
- Place **close to the destination** so it does not block the source from other places
- Typical use: VTY access, NAT and SNMP source lists

Extended ACL

- Matches protocol, source, destination and ports
- Numbers 100–199 and 2000–2699, or named
- Place **close to the source** to drop unwanted traffic early
- Typical use: traffic policy between subnets

Standard named ACL protecting management access

R1 · IOS-style CLI

```
! Only the IT subnet may manage this router
R1(config)# ip access-list standard MGMT-ONLY
R1(config-std-nacl)# 10 permit 10.0.99.0 0.0.0.255
R1(config-std-nacl)# 20 deny any log
R1(config-std-nacl)# line vty 0 15
R1(config-line)# access-class MGMT-ONLY in
! Numbered equivalent of line 10: access-list 10 permit 10.0.99.0 0.0.0.255
```

On VTY lines the command is `access-class`, not `access-group`. The explicit `deny any log` adds logging for refused attempts; without it the implicit deny still applies silently.

Extended named ACL on the guest VLAN, close to the source

R2 · IOS-style CLI

```
! Guests (192.168.30.0/24) may use DNS and the web, but never reach internal 10.0.0.0/8
R2(config)# ip access-list extended GUEST-IN
R2(config-ext-nacl)# 5 permit udp any any eq bootps
R2(config-ext-nacl)# 10 permit udp 192.168.30.0 0.0.0.255 host 10.0.50.11 eq domain
R2(config-ext-nacl)# 20 deny ip 192.168.30.0 0.0.0.255 10.0.0.0 0.255.255.255 log
R2(config-ext-nacl)# 30 permit tcp 192.168.30.0 0.0.0.255 any eq 80
R2(config-ext-nacl)# 40 permit tcp 192.168.30.0 0.0.0.255 any eq 443
R2(config-ext-nacl)# 50 permit udp 192.168.30.0 0.0.0.255 any eq 443
R2(config-ext-nacl)# 60 deny ip any any log
R2(config-ext-nacl)# interface GigabitEthernet0/1.30
R2(config-subif)# ip access-group GUEST-IN in
```

Order matters: the DNS permit (line 10) must come **before** the deny to 10.0.0.0/8 (line 20), or guests could never resolve names. UDP 443 allows QUIC-based web traffic. Line 5 matters too: DHCP Discovers arrive from source 0.0.0.0 and renewals go to the DHCP server in 10.0.0.0/8, so without an early bootps permit an inbound ACL silently breaks DHCP.

Verifying and editing with sequence numbers

R2 · IOS-style CLI

```
R2# show access-lists GUEST-IN
Extended IP access list GUEST-IN
  5 permit udp any any eq bootps (962 matches)
 10 permit udp 192.168.30.0 0.0.0.255 host 10.0.50.11 eq domain (4127 matches)
 20 deny ip 192.168.30.0 0.0.0.255 10.0.0.0 0.255.255.255 log (12 matches)
 30 permit tcp 192.168.30.0 0.0.0.255 any eq www (8830 matches)
 40 permit tcp 192.168.30.0 0.0.0.255 any eq 443 (51204 matches)
 50 permit udp 192.168.30.0 0.0.0.255 any eq 443 (9921 matches)
 60 deny ip any any log (37 matches)
R2# configure terminal
R2(config)# ip access-list extended GUEST-IN
R2(config-ext-nacl)# 15 permit icmp 192.168.30.0 0.0.0.255 host 192.168.30.1 echo
R2(config-ext-nacl)# no 60
```

Hit counters prove which line is matching. Named ACLs let you insert a line (15) between existing ones or delete a single line (`no 60`) without rewriting the list. Removing line 60 leaves the implicit deny in place, just without logging.

WATCH OUT · DIRECTION IS THE MOST COMMON MISTAKE

Stand on the router and look at the interface. **In** = traffic arriving from that subnet; **out** = traffic leaving towards it. An ACL written with the subnet as the *source* belongs **inbound** on the subnet's interface.

ZCNA ACL method

1. **State the policy in words** "Tellers may reach core banking on TCP 443 and DNS; nothing else."
2. **Identify the source and destination** Subnets, hosts, protocols and ports.
3. **Choose type and placement** Extended close to source; standard close to destination.
4. **Write specific lines first** Permits and denies for exceptions before broader rules.
5. **Decide on the final line** Explicit `deny ip any any log` for visibility, or rely on the implicit deny.
6. **Apply in a change window and verify** Test allowed and blocked flows; watch hit counters; keep a rollback ready.

CASE STUDY · MOTIJHEEL BANK BRANCH: THE 09:00 LOCKOUT

A bank branch in Motijheel, Dhaka, has a teller VLAN 192.168.50.0/24 on router subinterface Gi0/1.50. Policy says tellers may only reach the core banking servers in 10.20.0.0/24 on TCP 443 and the DNS server 10.20.1.53. On Thursday night an engineer added:

```
ip access-list extended TELLER-POLICY
10 permit tcp 192.168.50.0 0.0.0.255 10.20.0.0 0.0.0.255 eq 443
interface Gi0/1.50
ip access-group TELLER-POLICY out
```

At 09:00 on Sunday no teller can open the banking application, and queues form in the banking hall. The ACL shows zero matches.

1. Why does the ACL show zero matches and block everything?

2. What is missing from the policy itself?

3. Write the corrected configuration.

4. What process failures allowed this to reach a working day?

Model answer: module answer key.

TEMPLATE ACL planner

Write every ACL as policy first, then as lines. Reviewers can check intent, order and placement before it touches a live interface.

SEQ	ACTION	PROTOCOL	SOURCE	DESTINATION	PORT	REASON / TICKET
10	permit	udp	192.168.50.0 0.0.0.255	host 10.20.1.53	eq 53	Teller DNS — CHG-2210
30	permit	tcp	192.168.50.0 0.0.0.255	10.20.0.0 0.0.0.255	eq 443	Core banking HTTPS — CHG-2210

SCENARIO · DECIDE · WHERE SHOULD IT GO?

Users in 10.1.10.0/24 (behind R1) must not reach the HR server 10.3.3.10 (behind R3) on any protocol, but must still reach everything else, including other servers behind R3. You will use one extended ACL. Where is the best place to apply it?

- ☐ A Inbound on R1's interface facing 10.1.10.0/24.
- ☐ B Outbound on R3's interface facing the HR server.
- ☐ C Inbound on R3's WAN interface as a standard ACL denying 10.1.10.0/24.
- ☐ D On the VTY lines of R1.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

11. What wildcard mask matches every host in 172.20.16.0/21?

- ☐ A 0.0.15.255
- ☐ B 0.0.7.255
- ☐ C 255.255.248.0
- ☐ D 0.0.3.255

12. An ACL contains only `deny tcp any host 10.0.5.5 eq 23` and is applied inbound. What happens to all other traffic on that interface?

- ☐ A It is permitted because nothing denies it
- ☐ B It is permitted only if it is UDP
- ☐ C It is logged and then forwarded
- ☐ D It is dropped by the implicit deny

13. Which command applies a named ACL to restrict SSH access to the router itself?

- ☐ A `ip access-group MGMT-ONLY in` under the VTY lines
☐ B `ip access-group MGMT-ONLY out` on the LAN interface
☐ C `access-list MGMT-ONLY permit ssh` in global mode
☐ D `access-class MGMT-ONLY in` under the VTY lines

14. Where should a standard ACL normally be placed?

- ☐ A Close to the destination
☐ B Close to the source
☐ C On every interface in both directions
☐ D Only on the Internet-facing interface

15. Line 30 of a named ACL shows 0 matches while users complain the traffic it permits is blocked. Lines 10 and 20 have many matches. What should you suspect first?

- ☐ A The router has run out of memory for ACLs
☐ B Line 30 needs to be moved to the end
☐ C An earlier line matches that traffic first
☐ D Hit counters only update after a reload

16. Which extended ACL line lets hosts in 10.5.0.0/16 reach web server 10.9.9.9 over HTTPS?

- ☐ A `permit udp 10.5.0.0 0.0.255.255 host 10.9.9.9 eq 443 tcp`
☐ B `permit tcp 10.5.0.0 0.0.255.255 host 10.9.9.9 eq 443`
☐ C `permit tcp host 10.9.9.9 10.5.0.0 0.0.255.255 eq 443`
☐ D `permit tcp 10.5.0.0 255.255.0.0 host 10.9.9.9 eq 443`

☐ Lesson complete · Date _____ · Knowledge check score ____ / 6

MODULE 7 · GUARD · LESSON 4 · 30 MIN

Layer 2 Security — Port Security, DHCP Snooping and DAI

Stop the attacks that start on an access port.

- Configure port security with sticky MACs and choose a violation mode.
- Explain how DHCP snooping blocks rogue servers and builds a binding table.
- Configure dynamic ARP inspection and explain its dependency on DHCP snooping.

Most routers and firewalls sit at layer 3, but an attacker who plugs into an office wall port lands on layer 2 — inside the VLAN, beside the victims. Several classic attacks live here: flooding the MAC table, running a rogue DHCP server, and poisoning ARP caches to intercept traffic. Access switches have built-in features to stop each of them.

Layer 2 attacks and their counters

ATTACK	WHAT IT DOES	SWITCH FEATURE
MAC flooding	Fills the MAC table so the switch floods frames to every port, exposing traffic	Port security (MAC limits)
Rogue DHCP server	Hands out a fake gateway or DNS to redirect traffic	DHCP snooping
DHCP starvation	Requests every address so real clients get none	DHCP snooping rate limit, port security
ARP spoofing / poisoning	Claims to be the gateway's IP so traffic flows through the attacker	Dynamic ARP inspection (DAI)
IP spoofing on the LAN	Uses another host's IP address	IP source guard (uses the snooping table)
VLAN hopping	Negotiates a trunk or double-tags frames to reach other VLANs	Static access ports, DTP disabled, unused native VLAN

Port security for a desk with a phone and a PC

SW1 · IOS-style CLI

```
SW1(config)# interface GigabitEthernet0/5
SW1(config-if)# switchport mode access
SW1(config-if)# switchport access vlan 20
SW1(config-if)# switchport port-security
SW1(config-if)# switchport port-security maximum 2
SW1(config-if)# switchport port-security mac-address sticky
SW1(config-if)# switchport port-security violation restrict
SW1(config-if)# end
SW1# show port-security interface GigabitEthernet0/5
Port Security          : Enabled
Port Status            : Secure-up
Violation Mode         : Restrict
Maximum MAC Addresses  : 2
Total MAC Addresses    : 2
Sticky MAC Addresses   : 2
Security Violation Count : 0
```

Port security works only on a port statically set to access (or trunk) mode — not on a dynamically negotiating port. A maximum of 2 allows an IP phone plus the PC behind it. Sticky addresses are learned dynamically and written into the running configuration; save it to keep them after a reload.

Violation modes

MODE	DROPS OFFENDING FRAMES	LOGS / COUNTS	PORT STATE
Protect	Yes	No	Stays up
Restrict	Yes	Yes (syslog, SNMP, counter)	Stays up
Shutdown (default)	Yes	Yes	Error-disabled until recovered


Recovering an error-disabled port

SW1 · IOS-style CLI

```
SW1# show interfaces status err-disabled
Port      Name      Status      Reason
Gi0/7     Desk-17   err-disabled psecure-violation
! After removing the unauthorised device:
SW1(config)# interface GigabitEthernet0/7
SW1(config-if)# shutdown
SW1(config-if)# no shutdown
! Or recover automatically after 5 minutes:
SW1(config)# errdisable recovery cause psecure-violation
SW1(config)# errdisable recovery interval 300
```

Find out **why** the violation happened before re-enabling. Automatic recovery is convenient but brings a port back even if the rogue device is still attached.

DHCP snooping in action



DHCP snooping on an access switch

SW1 · IOS-style CLI

```
SW1(config)# ip dhcp snooping
SW1(config)# ip dhcp snooping vlan 20,30
SW1(config)# interface GigabitEthernet0/24
SW1(config-if)# description UPLINK-TO-DISTRIBUTION
SW1(config-if)# ip dhcp snooping trust
SW1(config-if)# interface range GigabitEthernet0/1 - 22
SW1(config-if-range)# ip dhcp snooping limit rate 15
SW1(config-if-range)# end
SW1# show ip dhcp snooping binding
MacAddress      IpAddress      Lease(sec)  Type           VLAN  Interface
00:50:56:A1:B2:C3  192.168.20.31  84211      dhcp-snooping  20    GigabitEthernet0/5
Total number of bindings: 1
```

Trust only the ports that lead towards legitimate DHCP servers. The rate limit (packets per second) error-disables a port that floods DHCP messages.

PRO TIP · WHEN SNOOPING BREAKS DHCP

Many switches insert **DHCP option 82** (relay agent information) into client requests when snooping is enabled. If the upstream relay or server rejects packets carrying option 82, clients stop getting addresses the moment snooping goes live. Check this first, and adjust the option 82 setting on the switch or the trust settings on the relay to match your design.

Dynamic ARP inspection**SW1 · IOS-style CLI**

```
SW1(config)# ip arp inspection vlan 20,30
SW1(config)# interface GigabitEthernet0/24
SW1(config-if)# ip arp inspection trust
SW1(config-if)# end
SW1# show ip arp inspection statistics vlan 20
```

Vlan	Forwarded	Dropped	DHCP Drops	ACL Drops
----	-----	-----	-----	-----
20	5321	14	14	0

DAI checks every ARP on untrusted ports against the DHCP snooping binding table. An ARP claiming the gateway's IP from a PC port does not match any binding and is dropped (the 14 DHCP drops above). Hosts with static IPs need ARP ACL entries or they will be blocked.

INSIGHT · A CHAIN OF TRUST

DHCP snooping builds the table; DAI and IP source guard **use** it. Enable and verify snooping first, then DAI. Turning DAI on in a VLAN with static-addressed servers and no ARP ACLs is a reliable way to cause an outage.

SCENARIO · DECIDE · ODD GATEWAY ON THE SECOND FLOOR

Several PCs on VLAN 30 suddenly have the default gateway 192.168.30.200 instead of 192.168.30.1 and slow, strange web pages. A student laptop running a home router image was found connected to a desk port. DHCP snooping is enabled on VLAN 20 only. What is the most effective prevention?

- ☐ A Enable DHCP snooping on VLAN 30 too, with only the uplink trusted.
- ☐ B Reduce the DHCP lease time on the real server.
- ☐ C Enable port security with maximum 1 MAC on the uplink.
- ☐ D Add a static ARP entry for the gateway on the router.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

17. Which port-security violation mode drops offending frames, logs the event and keeps the port up?

- ☐ A Protect
- ☐ B Shutdown
- ☐ C Restrict
- ☐ D Sticky

18. With DHCP snooping enabled, which message is dropped when it arrives on an untrusted port?

- ☐ A DHCP Discover
- ☐ B DHCP Offer
- ☐ C DHCP Request
- ☐ D ARP request for the gateway

19. What does dynamic ARP inspection use to decide whether an ARP reply is valid?

- ☐ A The DHCP snooping binding table and any ARP ACLs
- ☐ B The default gateway's routing table and its ARP cache
- ☐ C The spanning-tree topology table
- ☐ D The DNS server's PTR records

20. After enabling DAI on VLAN 40, a server with a static IP loses connectivity, but DHCP clients are fine. Why?

- ☐ A DAI blocks all servers by design
- ☐ B DAI requires the server to use IPv6
- ☐ C The server's port must be in shutdown mode
- ☐ D The server has no snooping binding or ARP ACL entry

21. Which configuration helps prevent VLAN hopping through trunk negotiation on a user port?

- ☐ A Enable sticky MAC learning on the user's access port
- ☐ B Configure the port as a trusted DHCP port
- ☐ C Static access mode with negotiation disabled
- ☐ D Increase the port-security maximum to 10

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 7 · GUARD · LESSON 5 · 30 MIN

VPNs, Firewalls, IPS and Zero Trust

Protect traffic in transit and control what crosses network boundaries.

- Compare site-to-site and remote-access VPNs and outline how IPsec protects traffic.
- Distinguish stateless ACLs, stateful firewalls, next-generation firewalls, IDS and IPS.
- Explain the core principles of zero trust and how they change network design.

So far the Guard stage has protected devices and access ports. The final layer protects **traffic crossing untrusted networks** — the Internet between a branch and head office — and **traffic crossing boundaries** between zones. Then it questions the boundary itself: in a zero-trust design, being "inside" the network earns no automatic trust.

Two kinds of VPN

Site-to-site

- Connects whole networks (branch ↔ head office, office ↔ cloud)
- Gateways build the tunnel; users are unaware
- Usually IPsec, always-on
- Replaces or backs up private WAN links

Remote access

- Connects one user device to the organisation
- VPN client software (IPsec or TLS-based) or a browser portal
- Authenticates the user, ideally with MFA
- Full tunnel (all traffic) or split tunnel (only corporate traffic)

What IPsec provides

CONF

Confidentiality

Encryption (for example AES) so captured packets are unreadable.

INT

Integrity

A keyed hash detects any change in transit.

AUTH

Authentication

Peers prove identity with pre-shared keys or certificates.

ANTI

Anti-replay

Sequence numbers stop an attacker re-sending captured packets.

Building an IPsec tunnel (concept)

01

Interesting traffic

Branch gateway sees traffic for head-office subnets

→

02

IKE negotiation

Peers authenticate and agree algorithms on UDP 500 (UDP 4500 if NAT is detected)

→

03

IPsec SAs

Security associations created for each direction

→

04

ESP protects data

Packets encrypted and authenticated inside ESP (IP protocol 50)

→

05

Rekey / expire

Keys refreshed on schedule; tunnel torn down when idle or on failure

IPsec terms

TERM	MEANING
IKEv2	Current key-exchange protocol that negotiates and authenticates IPsec tunnels
ESP	Encapsulating Security Payload — encryption plus integrity; the normal choice
AH	Authentication Header (IP protocol 51) — integrity only, no encryption, does not work through NAT

TERM	MEANING
Tunnel mode	Whole original packet encrypted and given a new outer IP header (site-to-site)
Transport mode	Only the payload protected; original IP header kept (host-to-host)
NAT-T	Wraps ESP in UDP 4500 so it can pass through NAT devices

Verifying a site-to-site tunnel

BR1 · IOS-style CLI

```
BR1# ping 10.20.0.10 source 192.168.50.1
!!!!
BR1# show crypto ipsec sa | include pkts
    #pkts encaps: 1520, #pkts encrypt: 1520, #pkts digest: 1520
    #pkts decaps: 1498, #pkts decrypt: 1498, #pkts verify: 1498
```

Rising *encrypt* **and** *decrypt* counters prove traffic is flowing both ways through the tunnel. Encrypts rising with zero decrypts usually means the far end is not sending back — check its policy, routing or a firewall blocking ESP/UDP 4500.

A modern open-source remote-access VPN

Linux laptop · WireGuard

```
# /etc/wireguard/corp.conf — split-tunnel remote access
[Interface]
PrivateKey = <laptop-private-key>
Address = 10.200.0.15/32
DNS = 10.0.50.11
#
[Peer]
PublicKey = <gateway-public-key>
Endpoint = vpn.branch.example:51820
AllowedIPs = 10.0.0.0/8, 192.168.0.0/16
```

`AllowedIPs` lists only corporate ranges, so this is a **split tunnel**: Internet traffic goes directly. Setting it to `0.0.0.0/0` would make it a **full tunnel**. Private keys never leave the device.

Inspection devices compared

DEVICE	WHAT IT INSPECTS	ACTION	NOTES
ACL (stateless)	Each packet's addresses, protocol and ports	Permit or deny	No memory of connections; return traffic needs its own rule
Stateful firewall	Tracks connections in a state table	Allows replies automatically	Default policy: inside may start sessions out; outside may not start sessions in
Next-generation firewall	Applications, users, URLs, files; optional TLS inspection	Granular allow/deny, often integrated IPS	Policy by app and identity, not just ports
IDS	Copy of traffic (SPAN or tap)	Alerts only	Cannot stop the packet it detects
IPS	Traffic inline	Can drop and block in real time	A false positive can block legitimate traffic — tune carefully

EXAMPLE · ZONES IN A BRANCH FIREWALL

Inside (staff), **Guest**, **DMZ** (public booking server) and **Outside** (Internet). Typical policy: Inside → Outside allowed with inspection; Outside → DMZ only HTTPS to the booking server; Guest → Outside only; nothing from Guest or Outside to Inside. Every rule has an owner and a reason.

Zero trust in five principles

ID

Verify explicitly

Every request is authenticated and authorised using identity, device health and context — every time.

LP

Least privilege

Users and devices reach only the applications they need, not whole networks.

SEG

Micro-segmentation

Small zones limit how far an intruder can move laterally.

BRCH

Assume breach

Design as if an attacker is already inside: monitor, log and limit blast radius.

CONT

Continuous evaluation

Trust is re-checked during the session as risk changes.

Traditional perimeter vs zero trust

Perimeter model

- Inside the firewall = trusted
- VPN puts a user on the whole network
- Flat internal networks
- One breach exposes much of the estate

Zero-trust model

- Location alone grants nothing
- Access per application after identity and device checks
- Segmented networks and policy enforcement points
- Breaches are contained and visible

SCENARIO · DECIDE · REMOTE STAFF AND THE OLD VPN

A Khulna NGO gives field staff a full-tunnel VPN that places them on the head-office LAN with access to every server. A stolen laptop with a saved VPN password was used to browse the finance share. Which change best applies zero-trust thinking?

- ☐ A Require MFA and device-health checks, and grant per-application access instead of whole-network access.
- ☐ B Switch from full tunnel to split tunnel.
- ☐ C Move the VPN to a different port number.
- ☐ D Add an IDS to alert on finance share access.

Circle your choice. Answer and feedback: module answer key.

Key terms

Vulnerability

A weakness in a system, configuration or process that a threat could use.

Phishing

Social engineering through deceptive messages that steal credentials or install malware.

AAA

Authentication, authorisation and accounting — usually provided centrally by RADIUS or TACACS+.

Extended ACL

An ACL that matches protocol, source, destination and ports; normally placed close to the source.

Port security

Switch feature that limits and learns the MAC addresses allowed on a port.

CIA triad

Confidentiality, integrity and availability — the three core goals of information security.

MFA

Multi-factor authentication: proof from two or more categories — know, have, are.

TACACS+

AAA protocol on TCP 49 that encrypts the full payload and supports per-command authorisation.

Implicit deny

The invisible final rule of every ACL that drops anything not matched earlier.

DHCP snooping

Switch feature that drops DHCP server messages on untrusted ports and builds an IP-MAC binding table.

DAI

Dynamic ARP inspection: drops ARP messages that do not match trusted IP-to-MAC bindings.

IPS

Intrusion prevention system: inline device that detects and blocks malicious traffic.

IPsec

Suite of protocols (IKE, ESP, AH) that authenticate and encrypt IP traffic, typically for VPNs.

Zero trust

Security model that grants no implicit trust based on network location and verifies every request.

KNOWLEDGE CHECK

22. Which IPsec protocol provides encryption and is normally used for VPN data?

- ☐ A AH, IP protocol 51
- ☐ B IKE, UDP 500
- ☐ C GRE, IP protocol 47
- ☐ D ESP, IP protocol 50

23. A site-to-site tunnel shows encrypt counters rising but decrypt counters stuck at zero. What is the most likely situation?

- ☐ A The tunnel is healthy and idle in both directions
- ☐ B Packets go out through the tunnel but none come back
- ☐ C The local router cannot encrypt any traffic
- ☐ D The peers have not yet started the IKE negotiation at all

24. What is the key difference between an IDS and an IPS?

- ☐ A An IPS sits inline and can block traffic; an IDS only alerts
- ☐ B An IDS encrypts traffic; an IPS only decrypts it
- ☐ C An IPS works only on wireless networks, an IDS only on wired ones
- ☐ D An IDS replaces the need for a firewall

25. A remote-access VPN sends only traffic for 10.0.0.0/8 through the tunnel, while web browsing goes direct. What is this called?

- ☐ A Full tunnelling
- ☐ B IPsec transport mode
- ☐ C Split tunnelling
- ☐ D NAT traversal

26. Which statement best reflects zero trust?

- ☐ A Devices on the internal LAN are trusted automatically
- ☐ B A strong perimeter firewall removes the need for segmentation
- ☐ C VPN users should reach the entire internal network
- ☐ D Each request is verified using identity, device and context

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 7 · GUARD

Answer key

Knowledge checks

- 1. A** An enabled clear-text management service is a weakness — a vulnerability. The threat is whoever could use it, the exploit is the method, and a compensating control would be something that reduces the risk, such as disabling it.
- 2. C** DDoS stops legitimate users reaching the service. It does not by itself expose or alter data, and non-repudiation is not one of the three CIA elements.
- 3. B** Phishing by SMS is smishing. Tailgating is physical, vishing uses voice calls, and whaling targets senior executives.
- 4. D** Breached-password screening stops users choosing passwords attackers already have. Forced periodic change and rigid composition rules produce predictable variants, and blocking paste discourages long unique passwords.
- 5. D** Origin-bound keys will not authenticate to a fake site. SMS codes and push approvals can be relayed or spammed, and security questions are just weak extra passwords.
- 6. C** TACACS+ encrypts the full body and separates authorisation, enabling per-command control. RADIUS hides only the password and combines authN/authZ, while SNMP and syslog are not AAA protocols.
- 7. A** Methods are tried in order and the next method is used only on an error such as no response. A reject is a final answer; local is not tried first; and the default list applies to all lines without their own list.
- 8. B** Command accounting records each configuration command and who entered it. Session times alone do not show the commands, and DHCP leases or interface counters do not identify administrators.
- 9. C** Type 7 is a weak reversible encoding meant to stop shoulder-surfing. It works on routers and switches alike, does not expire passwords and does not log them.
- 10. A** A spare open session lets you undo a mistake. Reloading, deleting the fallback account or using Telnet each increase the chance of lock-out or exposure.
- 11. B** A /21 mask is 255.255.248.0, and 255.255.255.255 – 255.255.248.0 = 0.0.7.255. 0.0.15.255 is a /20, 255.255.248.0 is the subnet mask itself, and 0.0.3.255 is a /22.
- 12. D** Every ACL ends with an implicit deny any, so a deny-only ACL blocks everything. Nothing is permitted by default, protocol type does not matter, and there is no automatic logging.
- 13. D** VTY lines use `access-class . ip access-group` applies ACLs to interfaces, an outbound LAN ACL would filter transit traffic, and the last command is not valid syntax.
- 14. A** Because a standard ACL matches only the source, placing it near the source would block that source from every destination. Extended ACLs go near the source; applying everywhere or only at the edge ignores the policy.
- 15. C** First match wins, so a broader deny above line 30 is the usual cause. Memory exhaustion is unlikely, moving a permit lower makes things worse, and counters update in real time.
- 16. B** The correct line uses TCP, source subnet with wildcard 0.0.255.255, then the destination host and port. The others mix protocols, reverse source and destination, or use a subnet mask instead of a wildcard.
- 17. C** Restrict drops and logs while keeping the port up. Protect drops silently, shutdown error-disables the port, and sticky is a learning method, not a violation mode.
- 18. B** Offers and Acks are server messages and are dropped on untrusted ports. Discover and Request are client messages and pass; ARP is checked by DAI, not DHCP snooping.
- 19. A** DAI validates IP-to-MAC pairs against snooping bindings (or ARP ACLs for static hosts). Routing, STP and DNS hold no IP-to-MAC bindings for this purpose.
- 20. D** Static hosts never appear in the DHCP binding table, so their ARPs fail validation unless an ARP ACL permits them. DAI does not block servers by design, needs no IPv6, and shutting the port would cut it off completely.
- 21. C** A static access port that does not negotiate cannot become a trunk. Sticky MACs, DHCP trust and MAC limits address other attacks.
- 22. D** ESP encrypts and authenticates the payload. AH provides integrity only, IKE negotiates keys rather than carrying data, and GRE is an unencrypted tunnelling protocol.
- 23. B** Encrypts prove local packets enter the tunnel; zero decrypts mean nothing returns, pointing to the far end's policy, routing or a blocked path. An idle tunnel would not show rising encrypts, encryption is working locally, and without IKE there would be no SAs or counters.
- 24. A** IPS is inline and can drop packets, while IDS watches a copy and alerts. Neither is an encryption device, IPS is not wireless-only, and neither replaces a firewall.
- 25. C** Sending only corporate prefixes through the tunnel is split tunnelling. Full tunnel sends everything, transport mode is an IPsec encapsulation choice, and NAT-T lets ESP cross NAT.
- 26. D** Zero trust removes implicit trust based on location and verifies each request. The other options describe the traditional perimeter model zero trust replaces.

Scenarios & case studies

Scenario · The friendly ISP engineer

A is best. Correct. A known detail like a circuit ID is not proof of identity, and legitimate providers never need your admin password. Verify through a trusted channel.

Scenario · Leaver on Friday

A is best. Correct. Central identity means one change removes access everywhere; a single shared break-glass account in a vault handles AAA outages.

Case study · Motijheel bank branch: the 09:00 lockout

Applied **out** on Gi0/1.50, the ACL checks traffic leaving the router towards the tellers — replies whose source is 10.20.x.x. No line matches source 192.168.50.0/24 in that direction, so the implicit deny drops every reply (zero matches on line 10). It belongs **inbound** on Gi0/1.50, close to the source. The policy also forgot DNS: the application is opened by name, so even correctly applied the ACL would break it. Corrected:

```
ip access-list extended TELLER-POLICY
5 permit udp any any eq bootps
10 permit udp 192.168.50.0 0.0.0.255 host 10.20.1.53 eq domain
20 permit tcp 192.168.50.0 0.0.0.255 host 10.20.1.53 eq domain
30 permit tcp 192.168.50.0 0.0.0.255 10.20.0.0 0.0.0.255 eq 443
40 deny ip any any log
interface Gi0/1.50
no ip access-group TELLER-POLICY out
ip access-group TELLER-POLICY in
```

Line 5 keeps teller DHCP working through the relay once the ACL is inbound. Process failures: no test of allowed flows after the change, no peer review of direction and DNS, no monitoring of hit counters, and no rollback plan or post-change check before the branch opened.

Scenario · Where should it go?

A is best. Correct. An extended ACL can match both source and destination exactly, so placing it close to the source stops the traffic before it crosses the WAN.

Scenario · Odd gateway on the second floor

A is best. Correct. A rogue DHCP server on an untrusted port would then have its offers dropped before any client accepts them.

Scenario · Remote staff and the old VPN

A is best. Correct. Identity plus device verification would have stopped the stolen laptop, and per-app access limits what any single account can reach.

08

MODULE 8 · AUTOMATE

Automation, Programmability and Readiness

Controllers, APIs, data formats, configuration as code, AI-assisted operations — then the whole journey and the capstone.

BY THE END OF THIS MODULE YOU WILL BE ABLE TO

- Compare traditional and controller-based networking and explain planes and APIs.
- Use REST concepts, HTTP verbs and status codes, and read JSON, YAML and YANG-modelled data.
- Explain idempotent configuration management with Ansible and read a short Python audit script.
- Describe intent-based networking and apply verification and safety to AI-assisted operations.
- Trace and troubleshoot a packet through all eight stages and plan the ZCNA capstone and exam.

LESSONS

1. Why Automate — Traditional vs Controller-Based Networking · 25 min
2. REST APIs, JSON and YAML · 30 min
3. Configuration Management, Ansible and a First Python Script · 35 min
4. Intent-Based Networking and AI-Assisted Operations · 25 min
5. The Whole Packet Journey, the Capstone and Exam Readiness · 35 min

MODULE 8 · AUTOMATE · LESSON 1 · 25 MIN

Why Automate — Traditional vs Controller-Based Networking

Planes, controllers, SDN and the APIs that connect them.

- Explain the business and technical reasons to automate network operations.
- Separate the data, control and management planes and describe what a controller centralises.
- Distinguish northbound from southbound interfaces and name common southbound protocols.

Everything you have configured so far — VLANs, OSPF, NAT, ACLs, SSIDs — you typed into one device at a time. That works for a lab. For a bank with 300 branches it means thousands of nearly identical changes, each a chance for a typo at 23:00. The **Automate** stage is about making the network **consistent, fast to change, and provable**: the computer does the repetitive work, and engineers design, review and verify.

Why organisations automate

CONS

Consistency

Every device gets the same tested configuration; drift from the standard is detected automatically.

SPD

Speed

A change to 300 branches takes minutes, not weeks, and new sites are built from a template.

ERR

Fewer errors

Most outages are caused by human change errors; scripts do not mistype IP addresses.

AUD

Evidence

Every change is recorded in version control: who, what, when and why.

SCAL

Scale

A small team can run a large network and spend time on design instead of copy-paste.

The three planes of a network device

Management plane

How engineers and tools configure and monitor the device

SSH, SNMP, syslog, NETCONF, RESTCONF, APIs

Control plane

How the device decides where traffic goes

OSPF, STP, ARP, routing and MAC tables built here

Data plane

How the device actually forwards traffic

Frame and packet forwarding, ACLs and QoS at line rate

Traditional vs controller-based networking

Traditional (distributed)

- Each device has its own control and management plane
- Configured box by box through the CLI
- Network-wide view lives in engineers' heads and spreadsheets
- Changes are slow and inconsistent at scale

Controller-based / SDN

- A **controller** holds the network-wide view and policy
- Devices are programmed by the controller through APIs
- Engineers express policy once; the controller pushes it everywhere
- Telemetry flows back for assurance and troubleshooting

DEFINITION · SOFTWARE-DEFINED NETWORKING (SDN)

An approach where the control and management of the network are centralised in software (the controller) and exposed through APIs, while devices focus on forwarding. Pure SDN moves the control plane entirely; most enterprise products today keep distributed routing protocols on devices and centralise management, policy and visibility.

Where the APIs sit

Applications and orchestration

Scripts, IT service management, dashboards, automation pipelines

Northbound API

Usually a REST API with JSON

Lets applications ask the controller for data or request changes

Controller

Network-wide view, policy, inventory, assurance

Southbound interfaces

NETCONF, RESTCONF, gNMI, OpenFlow, or SSH/CLI and SNMP for older devices

Network devices

Switches, routers, APs, firewalls

Underlay and overlay in a controller-managed fabric

Underlay

- The physical routed network
- Simple, stable IP connectivity between devices (often OSPF or another IGP)
- Rarely changes

Overlay

- Virtual networks built on top of the underlay with tunnels (for example VXLAN)
- Segments and policies follow users and workloads
- Changes frequently, driven by the controller

INSIGHT · AUTOMATION DOES NOT REPLACE FUNDAMENTALS

A controller can push a broken design to 300 sites in two minutes. You still need to understand subnets, VLANs, routing and ACLs — to design the intent, to review what the automation will do and to troubleshoot when the overlay hides a broken underlay.

SCENARIO · DECIDE · WHICH APPROACH FOR THE SCHOOL NETWORK?

A district education office manages 45 schools, each with one router, two switches and four APs. Two engineers make every change by SSH. Last term, an ACL typo at one school blocked exam results uploads for a day. They have a small budget. What is the most sensible first step?

- ☐ A Standardise each school's configuration into a template and push changes with an automation tool from version control, after a peer review.
- ☐ B Buy a full SDN fabric for all 45 schools immediately.
- ☐ C Keep using SSH but ask engineers to type more carefully.
- ☐ D Stop making changes during term time.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

1. OSPF calculating best paths belongs to which plane?

- ☐ A Data plane
- ☐ B Management plane
- ☐ C Control plane
- ☐ D Application plane

2. An IT ticketing system asks the controller for a list of devices with high CPU. Which interface is it using?

- ☐ A A southbound interface to the switches
- ☐ B The controller's northbound API
- ☐ C The console port of each device
- ☐ D The data plane of the access switches

3. Which of these is a southbound protocol a controller might use to configure devices?

- ☐ A HTTP status code 200
- ☐ B DHCP option 43
- ☐ C Syslog severity 5
- ☐ D NETCONF

4. In a fabric, what is the underlay?

- ☐ A The physical IP network that carries the overlay tunnels
- ☐ B The virtual networks that users and devices are placed into
- ☐ C The northbound REST API of the controller
- ☐ D The wireless SSIDs that clients join

☐ Lesson complete · Date _____ · Knowledge check score ____ / 4

MODULE 8 · AUTOMATE · LESSON 2 · 30 MIN

REST APIs, JSON and YAML

Talk to controllers and devices programmatically and read the data they return.

- Map HTTP verbs to create, read, update and delete operations and interpret status codes.
- Read and write valid JSON and YAML and spot common syntax errors.
- Describe how model-driven interfaces (YANG with NETCONF or RESTCONF) structure device data.

A **REST API** lets software talk to a controller or device using the same HTTP that web browsers use. Instead of screen-scraping CLI output, a script sends a request to a URL and receives **structured data** — usually JSON — that it can process reliably. Reading API documentation and responses is now a core associate skill.

REST in practice

- **Client-server:** your script is the client; the controller or device is the server.
- **Stateless:** every request carries everything needed, including authentication.
- **Resources and URIs:** each object has a URL, for example `/api/v1/devices/42` .
- **Uniform verbs:** HTTP methods say what you want to do with the resource.
- **Representations:** data travels as JSON (or XML), described by headers such as `Content-Type` and `Accept` .

HTTP verbs and CRUD

VERB	CRUD ACTION	EXAMPLE	SAFE TO REPEAT?
GET	Read	List interfaces or one device's details	Yes — read only
POST	Create	Add a new VLAN or site	No — may create duplicates
PUT	Update (replace)	Replace a whole object with the body sent	Yes — same result each time
PATCH	Update (partial)	Change only the description field	Usually
DELETE	Delete	Remove a VLAN	Yes — it is gone after the first call

HTTP status codes you will see

CODE	MEANING	WHAT TO DO
200 OK	Request succeeded	Read the body
201 Created	New resource created (POST)	Note the new ID or location
204 No Content	Succeeded, nothing to return	Common after PUT, PATCH or DELETE
400 Bad Request	Malformed request or invalid data	Check the JSON body and field names
401 Unauthorized	Missing or invalid authentication	Check the token or credentials
403 Forbidden	Authenticated but not permitted	Check the account's role
404 Not Found	Resource or URL does not exist	Check the URI and object ID
429 Too Many Requests	Rate limit hit	Slow down; respect any retry header
500 Internal Server Error	Server-side failure	Retry later, check server logs

A GET request with a bearer token

```
Linux host

$ curl -s -H "Accept: application/json" \
  -H "Authorization: Bearer $TOKEN" \
  https://ctrl.branch.example/api/v1/devices?site=sylhet
{"devices": [{"name": "r1", "ip": "10.0.99.1", "role": "router", "reachable": true}, {"name": "sw1", "ip": "10.0.99.11", "role": "switch", "reachable": false}]}
```

The token is passed in a header from an environment variable — never hard-coded in scripts or committed to version control. The response is compact JSON; tools like `jq` or Python make it readable.

The same data in JSON and YAML

JSON

- `{`
- `"hostname": "r1",`
- `"vlangs": [10, 20],`
- `"ntp": {"server": "10.0.99.123", "prefer": true}`
- `}`

YAML

- `hostname: r1`
- `vlangs:`
- `- 10`
- `- 20`
- `ntp:`
- `server: 10.0.99.123`
- `prefer: true`

Reading structured data

ELEMENT	JSON	YAML
Object / dictionary	<code>{ "key": value }</code>	<code>key: value</code> lines at the same indent
Array / list	<code>[1, 2, 3]</code>	Lines starting with <code>-</code>
String	Always in double quotes	Quotes optional (needed for special characters)
Boolean / null	<code>true</code> , <code>false</code> , <code>null</code>	<code>true</code> , <code>false</code> , <code>null</code> (or <code>~</code>)
Comments	Not allowed	<code># comment</code>
Structure	Braces and brackets; whitespace ignored	Indentation defines structure; use spaces, never tabs

WATCH OUT · CLASSIC SYNTAX MISTAKES

JSON: single quotes, a trailing comma after the last item, or a missing comma between items. YAML: mixing tabs and spaces, inconsistent indentation, or forgetting the space after `-` or `:` . A validator or linter catches these in seconds — run one before every commit.

DEFINITION · MODEL-DRIVEN PROGRAMMABILITY

YANG (RFC 7950) is a modelling language that describes device data — interfaces, routes, VLANs — as a precise tree. **NETCONF** (RFC 6241, SSH on TCP 830, XML) and **RESTCONF** (RFC 8040, HTTPS, JSON or XML) read and change data that follows those models. Standard models such as `ietf-interfaces` work across many platforms.

RESTCONF: reading a standard YANG model

Linux host

```
$ curl -s -u netauto --cacert ca.pem \
-H "Accept: application/yang-data+json" \
https://10.0.99.1/restconf/data/ietf-interfaces:interfaces/interface=Loopback0
{
  "ietf-interfaces:interface": [
    {
      "name": "Loopback0",
      "type": "iana-if-type:softwareLoopback",
      "enabled": true
    }
  ]
}
```

The URL follows the YANG tree: module `ietf-interfaces`, container `interfaces`, list `interface` keyed by name. `-u netauto` without a password makes curl prompt for it, and `--cacert` verifies the device's certificate instead of disabling TLS checks.

SCENARIO · DECIDE · A 401 AT MIDNIGHT

An automation job that adds guest VLANs through the controller's API worked yesterday. Tonight every POST returns `401 Unauthorized`; GET requests from an engineer's browser session still work. What is the most likely cause?

- ☐ A The job's API token has expired or been revoked, so its requests are no longer authenticated.
- ☐ B The JSON body has a trailing comma.
- ☐ C The VLAN already exists.
- ☐ D The controller is overloaded.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

5. Which HTTP verb should a script use to read the list of VLANs without changing anything?

- ☐ A POST
- ☐ B GET
- ☐ C PUT
- ☐ D PATCH

6. An API returns `403 Forbidden`. What does this tell you?

- ☐ A Your credentials were not sent or are invalid
- ☐ B The URL you requested does not exist
- ☐ C The server crashed while processing the request
- ☐ D Authenticated, but not permitted to do this

7. Which JSON snippet is valid?

- ☐ A `{"vlan": 20, "name": "GUEST"}`
- ☐ B `{'vlan': 20, 'name': 'GUEST'}`
- ☐ C `{"vlan": 20, "name": "GUEST",}`
- ☐ D `{vlan: 20, name: GUEST}`

8. In YAML, what defines which keys belong inside `ntp:` ?

- ☐ A Curly braces that must wrap around the nested keys
- ☐ B A semicolon at the end of each line
- ☐ C Consistent indentation with spaces beneath it
- ☐ D Tab characters before each nested key

9. Which statement about RESTCONF is correct?

- ☐ A It uses SNMP traps to push configuration
- ☐ B It replaces YANG with free-form CLI text
- ☐ C It uses HTTPS to work with YANG-modelled data
- ☐ D It runs only over Telnet on TCP 23

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 8 · AUTOMATE · LESSON 3 · 35 MIN

Configuration Management, Ansible and a First Python Script

Treat configuration as code: templates, playbooks, idempotency and version control.

- Explain idempotency and why it makes automation safe to re-run.
- Read an Ansible inventory and playbook and predict the result of a check run.
- Read and adapt a short Python script that audits device data.

Configuration management means describing how devices *should* be configured, storing that description in version control, and letting a tool make reality match it. The description becomes the **source of truth**; the tool reports and corrects **drift**. This is the heart of "network as code".

A network-as-code workflow



DEFINITION · IDEMPOTENCY

An operation is **idempotent** if running it once or ten times leaves the system in the same state. An idempotent task that adds `ntp server 10.0.99.123` changes nothing on a device that already has it and reports `ok`, not `changed`. This makes automation safe to re-run and turns every run into a compliance check.

Ansible at a glance

How it works

- **Agentless:** connects over SSH (or HTTPS APIs) — nothing installed on network devices
- Runs from a control node such as a Linux VM
- Playbooks written in YAML

Key parts

- **Inventory:** which devices, in which groups, with which variables
- **Playbook:** ordered plays and tasks
- **Modules:** units of work, such as push config or gather facts
- **Templates:** Jinja2 files that generate config

Inventory: who and how

Ansible control node · inventory.yml

```

all:
  children:
    school_routers:
      hosts:
        sch-rtr-01: { ansible_host: 10.30.1.1 }
        sch-rtr-02: { ansible_host: 10.30.2.1 }
      vars:
        ansible_connection: ansible.netcommon.network_cli
        ansible_network_os: <platform-collection>.<os>  # set for your platform
        ansible_user: netauto
        ntp_server: 10.0.99.123
  
```

Group variables apply to every host in the group. The network OS value comes from the collection for your platform; credentials come from a vault or prompt, never from plain text in the file.

Playbook: what to do

Ansible control node · baseline.yml

```
- name: Apply NTP and logging baseline to school routers
  hosts: school_routers
  gather_facts: false
  tasks:
    - name: Ensure standard NTP and syslog lines are present
      ansible.netcommon.cli_config:
        config: |
          ntp server {{ ntp_server }}
          logging host 10.0.99.60
```

`{{ ntp_server }}` is filled from the inventory. Platform-specific resource modules give the strongest idempotency; this generic module relies on the platform comparing the running configuration before pushing.

A dry run before the real change

Ansible control node

```
$ ansible-playbook -i inventory.yml baseline.yml --check --diff
PLAY [Apply NTP and logging baseline to school routers] ****
TASK [Ensure standard NTP and syslog lines are present] ****
ok: [sch-rtr-01]
changed: [sch-rtr-02]
PLAY RECAP ****
sch-rtr-01 : ok=1  changed=0  unreachable=0  failed=0  skipped=0  rescued=0  ignored=0
sch-rtr-02 : ok=1  changed=1  unreachable=0  failed=0  skipped=0  rescued=0  ignored=0
```

`--check` makes no changes; `--diff` shows what would change. Router 01 already complies (`ok`), router 02 has drifted (`changed`). Run again without `--check` in the change window, then once more — the second real run should report zero changes.

One template, many switches

Jinja2 template · access-switch.j2

```
hostname {{ hostname }}
{% for vlan in vlans %}
vlan {{ vlan.id }}
  name {{ vlan.name }}
{% endfor %}
ntp server {{ ntp_server }}
```

Feed it a different hostname and VLAN list per site and every switch gets a consistent, correctly formatted configuration.

A tiny, vendor-neutral audit script

Python 3 · ntp_audit.py

```
import json
STANDARD_NTP = "10.0.99.123"
with open("devices.json", encoding="utf-8") as f:
    devices = json.load(f)          # a list of dictionaries
for dev in devices:
    status = "OK" if dev["ntp"] == STANDARD_NTP else "FIX"
    print(f"{dev['name']:<10} {dev['ntp']:<15} {status}")
```

The script reads exported device data, compares one field with the standard and prints a report. Indentation defines the loop body in Python, just as indentation defines nesting in YAML.

Running the audit

Linux host

```
$ cat devices.json
[{"name": "sch-rtr-01", "ntp": "10.0.99.123"},
 {"name": "sch-rtr-02", "ntp": "192.168.1.1"}]
$ python3 ntp_audit.py
sch-rtr-01 10.0.99.123      OK
sch-rtr-02 192.168.1.1    FIX
```

Next steps would be to fetch the data live through an API instead of a file, and to feed the FIX list into the Ansible playbook.

CASE STUDY · RANGPUR SCHOOL-DISTRICT AUTOMATION PROJECT

A district education office near Rangpur manages 45 schools. Each has a router, two access switches and four APs. Two engineers make all changes by SSH using notes in a shared document. Recently: an ACL typo blocked exam-result uploads at one school; three schools still use an old NTP server; and nobody can say which switches have Telnet enabled. The office has one Linux server, a small budget and engineers who know basic CLI but not programming.

1. What would you automate first, and why?

2. What is the minimum source of truth they need?

3. How would you make changes safe?

4. How would you measure success after one term?

Model answer: module answer key.

TEMPLATE Automation inventory

Capture what each device needs for automation before writing any playbook.

DEVICE	SITE	ROLE	MGMT IP	ACCESS METHOD	PLATFORM / OS	BASELINE COMPLIANT?
<i>sch-rtr-01</i>	<i>Pirganj High School</i>	<i>Router</i>	<i>10.30.1.1</i>	<i>SSH (network_cli)</i>	<i>IOS-style router</i>	<i>Yes</i>
<i>sch-sw-01a</i>	<i>Pirganj High School</i>	<i>Access switch</i>	<i>10.30.1.11</i>	<i>SSH (network_cli)</i>	<i>IOS-style switch</i>	<i>No — Telnet on</i>

DEVICE	SITE	ROLE	MGMT IP	ACCESS METHOD	PLATFORM / OS	BASELINE COMPLIANT?

SCENARIO · DECIDE · THE SECOND RUN

You run a playbook that should enforce SSH-only VTY lines on 45 routers. The first real run reports `changed` on 12 routers. You immediately run it again with no other changes, and it again reports `changed` on the same 12. What does this most likely mean?

- ☐ A The task is not idempotent — something in the desired lines never matches how the device stores them, so it re-pushes every time.
- ☐ B The 12 routers are reverting the change automatically.
- ☐ C This is normal; Ansible always reports changed.
- ☐ D The other 33 routers failed silently.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

10. A task runs twice on the same device and reports `changed` the first time and `ok` the second. What property does this show?

- ☐ A Statelessness
- ☐ B Idempotency
- ☐ C Encapsulation
- ☐ D Redundancy

11. Why is Ansible described as agentless for network devices?

- ☐ A It needs an agent only on routers, not on switches
- ☐ B It runs entirely inside the network controller's own software
- ☐ C It uses SNMP traps instead of connections
- ☐ D It uses SSH or APIs and installs nothing on devices

12. What does `--check` add to an `ansible-playbook` run?

- ☐ A It predicts changes without applying them
- ☐ B It applies changes and then rolls them back
- ☐ C It checks only YAML syntax and exits
- ☐ D It forces every task to report changed

13. In the Python audit script, what does `json.load(f)` return for the sample file?

- ☐ A A single long string of text
- ☐ B A list of dictionaries, one per device
- ☐ C A YAML document
- ☐ D A dictionary of lists, one per field

14. Which file in an Ansible project lists the devices and their groups?

- ☐ A The playbook
- ☐ B The Jinja2 template
- ☐ C The inventory
- ☐ D The module documentation

☐ Lesson complete · Date _____ · Knowledge check score ____ / 5

MODULE 8 · AUTOMATE · LESSON 4 · 25 MIN


Intent-Based Networking and AI-Assisted Operations

From business intent to verified configuration — with machines helping and humans in charge.

- Describe the intent-based networking loop: translate, activate, assure.
- Identify realistic uses of AI-assisted tools in network operations.
- Apply verification and safety controls to any AI-suggested change.

Automation executes what you tell it. **Intent-based networking** goes one step further: you state **what** you want — "guest devices may reach only the Internet", "voice gets priority at every branch" — and the system works out **how**, applies it, and continuously checks that reality still matches the intent. AI-assisted tools add analysis and suggestions on top. Both are powerful; both need engineers who can verify the result.

The intent-based loop




EXAMPLE · INTENT IN ACTION

Intent: "Point-of-sale terminals at every branch may reach only the payment servers on TCP 443." The system generates ACLs or segment policies for 120 branches, deploys them, and later raises an alert when a new branch's terminal is seen talking to a file server — before an auditor finds it.

Realistic uses of AI-assisted network operations

USE	WHAT IT HELPS WITH	MAIN RISK
Anomaly detection	Spotting unusual traffic, latency or error patterns against a learned baseline	False alarms or missed slow-burn issues
Log and alert summarisation	Grouping thousands of events into a few likely incidents	Important detail summarised away
Natural-language queries	"Which branches had WAN loss above 2% this week?"	Wrong query translated into wrong data
Configuration drafting	First draft of an ACL, template or script	Plausible but wrong or non-existent commands
Root-cause hints	Correlating a routing change with a jump in tickets	Correlation mistaken for cause
Capacity forecasting	Predicting when links or APs will saturate	Forecasts based on unusual history

The ZCNA safe-assist loop



TEST**Test**

Try it in a lab, simulator or digital twin; run a check-mode dry run.

APPL**Apply**

Deploy through the normal change process, with a rollback ready.

VER**Verify**

Confirm with show commands, tests and monitoring that the intent is met.

WATCH OUT · DATA AND ACCESS SAFETY

Never paste configurations containing passwords, keys, SNMP strings or customer data into tools your organisation has not approved. Give automated agents **read-only** credentials by default, require human approval for any change, and log every action they take. An assistant that can change 300 routers can also break 300 routers.

Healthy and unhealthy use**Healthy**

- Summarise logs, then confirm the key events yourself
- Draft a template, then review and lab-test it
- Ask for likely causes, then prove one with evidence
- Keep humans accountable for every change

Unhealthy

- Paste a suggested command straight into production
- Trust a summary without reading the source
- Share full configurations with unapproved tools
- Let a tool make changes with admin rights and no approval

Reviewing a suggested change

R1 · IOS-style CLI

```
! Suggested by an assistant: 'permit tcp any host 10.20.0.10 eq 443'
! Review: does this match the policy 'tellers only'? No - 'any' is too broad.
! Corrected after review:
R1(config)# ip access-list extended TELLER-POLICY
R1(config-ext-nacl)# 25 permit tcp 192.168.50.0 0.0.0.255 host 10.20.0.10 eq 443
R1(config-ext-nacl)# end
R1# show access-lists TELLER-POLICY | include 25
    25 permit tcp 192.168.50.0 0.0.0.255 host 10.20.0.10 eq 443 (0 matches)
```

The suggestion was syntactically valid but violated the intent. Review compares each line with the policy, not just with the syntax — then hit counters confirm the line is used.

Before applying any AI-suggested or auto-generated change

- ☐ Secrets and customer data removed from anything shared with the tool.
- ☐ Every command exists on the target platform and version (checked in documentation or lab).
- ☐ The change matches the written policy or intent, not just the symptom.
- ☐ Blast radius understood: which devices, users and services are affected.
- ☐ Dry run or lab test completed, with output saved.
- ☐ Peer review and change approval recorded.
- ☐ Rollback steps written and tested.
- ☐ Post-change verification commands and success criteria defined.

SCENARIO · DECIDE · THE CONFIDENT SUGGESTION

During an outage at a Gazipur factory, an AI-assisted tool summarises the logs and suggests that an STP root change caused the problem. It proposes a command to set a new root priority on the core switch. Production is down and the plant manager is pressing for action. What should the engineer do?

- ☐ A Quickly verify the claim with show commands (STP root, topology change counts, logs), then apply a reviewed fix through the emergency change process.
- ☐ B Apply the command immediately because the tool analysed all the logs.
- ☐ C Ignore the tool and reload the core switch.
- ☐ D Wait until the next scheduled change window.

Circle your choice. Answer and feedback: module answer key.

KNOWLEDGE CHECK

15. In intent-based networking, what does the assurance stage do?

- ☐ A Continuously checks that the network behaves as intended
- ☐ B Converts business language into device configuration
- ☐ C Pushes the generated configuration out to every relevant device
- ☐ D Writes the business intent for the organisation

16. An assistant drafts an ACL line that is syntactically valid. What must the engineer still check?

- ☐ A Only that the line has a sequence number
- ☐ B Only that the device CLI accepts the command without errors
- ☐ C Nothing further if the tool is approved
- ☐ D That it matches the written policy and test results

17. Which credential model is safest for a new AI-assisted monitoring agent?

- ☐ A Read-only access, with changes needing human approval
- ☐ B Full admin rights so it can fix issues quickly
- ☐ C The shared break-glass account for convenience
- ☐ D No authentication on a trusted management VLAN

18. What is the main risk of relying on an automated log summary during an incident?

- ☐ A Summaries always take longer than reading logs
- ☐ B Key details may be summarised away or misinterpreted
- ☐ C Logs cannot be summarised by software
- ☐ D Summaries delete the original log entries from the server

☐ Lesson complete · Date _____ · Knowledge check score ____ / 4

MODULE 8 · AUTOMATE · LESSON 5 · 35 MIN


The Whole Packet Journey, the Capstone and Exam Readiness

Follow one packet through all eight stages, then prepare to prove your skills.


- Trace a single user request through all eight stages of the ZCNA Packet Journey.
- Use the journey as a structured troubleshooting method.
- Plan the ZCNA capstone and a personal exam-readiness strategy.

Eight modules ago the journey started with a signal on a wire. Now you can follow a real request end to end — and, more importantly, you can use that path as a **troubleshooting map**. When something fails, ask "how far did the packet get?" and test each stage in order instead of guessing.

The ZCNA Packet Journey



Nadia opens the bank portal from a Sylhet branch laptop



Journey-based troubleshooting

SYMPTOM	STAGE TO SUSPECT	FIRST CHECKS
Link light off, interface down/down	Wire	Cable, transceiver, <code>show interfaces status</code>
169.254.x.x address	Serve (DHCP) or Switch	Access VLAN, trunk, <code>ip helper-address</code> , scope

SYMPTOM	STAGE TO SUSPECT	FIRST CHECKS
Can ping gateway but not remote subnet	Route	<code>show ip route</code> , return route, OSPF neighbours
Can reach IPs but not names	Serve (DNS)	Client DNS server, <code>nslookup</code> / <code>dig</code> , ACL on UDP/TCP 53
Inside works, Internet fails	Serve (NAT) or Route	<code>show ip nat statistics</code> , default route
Wi-Fi connects then drops when walking	Air	Signal at cell edge, roaming settings, 802.11r
One application blocked, others fine	Guard	<code>show access-lists</code> hit counters, firewall logs
Same fault on many sites after a change	Automate	Recent commits, template diff, rollback

PRO TIP · BOTTOM-UP, TOP-DOWN OR FOLLOW THE PATH

For a single user, follow the packet path from their device outward. For a whole-site outage, start at the shared dependencies (WAN, core, DHCP, DNS). For a fault right after a change, start with what changed. The journey gives you the checklist either way.

SCENARIO · DECIDE · INTEGRATED TROUBLESHOOTING

A new Wi-Fi 6E AP was installed in a Barishal branch meeting room on Monday. It bridges the Branch-Staff SSID to VLAN 10 locally. Laptops associate and pass WPA3-Enterprise authentication, but then show 169.254.x.x addresses. Wired PCs in VLAN 10 on the same switch work normally. The AP's switch port shows `switchport mode trunk` , `switchport trunk native vlan 99` and `switchport trunk allowed vlan 30,99` . What is the most likely cause?

- ☐ A VLAN 10 is missing from the trunk's allowed list, so wireless clients' DHCP broadcasts never reach the wired VLAN.
- ☐ B The 6 GHz SSID must use WPA2 instead of WPA3.
- ☐ C The router needs an `ip helper-address` for VLAN 10.
- ☐ D The AP's transmit power is too low for the meeting room.

Circle your choice. Answer and feedback: module answer key.

The ZCNA capstone: build the Barishal branch

- Design** Address plan (IPv4 and IPv6), VLANs for staff, voice, guest and management, and a diagram.
- Switch and route** Access and distribution switching with trunks, STP and EtherChannel; OSPF to head office with a default route and gateway redundancy.
- Serve** DHCP with relay, DNS, PAT to the Internet, NTP, syslog, SNMPv3, SSH-only management and a WAN QoS policy.
- Air** A two-AP WLAN with staff (WPA3) and guest SSIDs, a channel plan and a short survey record.
- Guard** Hardened devices, AAA with fallback, ACLs for guest and teller policies, port security, DHCP snooping and DAI.
- Automate** A YAML inventory, one template and one playbook or script that audits the baseline.
- Prove it** Verification evidence for every stage: show outputs, test results and a one-page journey walk-through.

TEMPLATE Capstone evidence log

Collect proof for each stage as you build, so your final submission is a record, not a reconstruction.

STAGE	REQUIREMENT	VERIFICATION COMMAND / TEST	RESULT	SCREENSHOT OR FILE
Serve	Guest VLAN gets DHCP via relay	<code>show ip dhcp binding</code> on server; <code>client ipconfig / ip addr</code>	Pass — 192.168.30.21	serve-dhcp-01.txt
Guard	Guest cannot reach 10.0.0.0/8	Ping 10.0.50.11 from guest; <code>show access-lists</code> GUEST-IN	Pass — line 20 matches	guard-acl-02.txt

STAGE	REQUIREMENT	VERIFICATION COMMAND / TEST	RESULT	SCREENSHOT OR FILE

Exam strategy that works

- **Read the last sentence first** — know whether it asks for the **first**, **best**, **most likely** or **incorrect** option before reading the story.
- **Find the stage**: decide which part of the Packet Journey the scenario is really about.
- **Read outputs carefully**: interface state, hit counters, "stratum 16", 169.254 — the evidence usually points at one answer.
- **Eliminate** options that are technically false, then those that are true but do not solve this problem.
- **Manage time**: mark long questions, move on, return later; never leave a question blank.
- **Practise in a lab**: configure and break things yourself; commands you have typed are commands you remember.
- **Check the current ZCNA exam guide** for format, duration and any lab components before you book.

Exam-readiness checklist

- ☐ I can subnet an IPv4 /24 into any prefix in under two minutes and read IPv6 prefixes confidently.
- ☐ I can configure VLANs, trunks and a basic EtherChannel from memory and verify STP root.
- ☐ I can configure OSPFv2 single-area and static routes, and read a routing table.
- ☐ I can configure DHCP relay, PAT and an SSH-only management baseline.
- ☐ I can write and place standard and extended ACLs with correct wildcard masks.
- ☐ I can explain WPA2 vs WPA3 and personal vs enterprise, and 802.1X roles.
- ☐ I can read JSON and YAML and explain REST verbs and status codes.
- ☐ I have completed the capstone with verification evidence for every stage.
- ☐ I have taken at least two timed practice exams and reviewed every wrong answer.

Key terms

Data plane

The part of a device that forwards packets and frames.

Management plane

How engineers and tools configure and monitor a device: SSH, SNMP, APIs.

Control plane

The part of a device that builds forwarding decisions, using protocols such as OSPF, STP and ARP.

SDN controller

Software that holds the network-wide view and programs devices through APIs.

Northbound API

The interface applications use to talk to a controller, usually REST with JSON.

REST API

An HTTP-based interface where resources have URLs and verbs such as GET and POST act on them.

YAML

A human-friendly data format where indentation defines structure; used by Ansible.

Idempotency

The property that repeating an operation leaves the system in the same final state.

Intent-based networking

Defining desired outcomes and letting the system translate, apply and continuously assure them.

Southbound interface

How a controller talks to devices: NETCONF, RESTCONF, gNMI, OpenFlow or CLI.

JSON

A text data format of objects, arrays, strings, numbers, booleans and null, with double-quoted strings.

YANG

A data-modelling language that describes device configuration and state for NETCONF and RESTCONF.

Source of truth

The authoritative, version-controlled record of how the network should be.

REFLECT

Look back at the eight stages of the ZCNA Packet Journey. Which stage do you feel least confident troubleshooting today, and what two lab exercises will you do this week to fix that?

KNOWLEDGE CHECK

19. A user's laptop has a valid address and can ping its gateway but not a server in another subnet. Which stage do you investigate first?

- | | |
|--|--|
| ☐ A Wire — cable and transceiver | ☐ B Air — channel and signal |
| ☐ C Serve — DHCP scope options for the user's VLAN | ☐ D Route — routes to the destination and back |

20. Many branches fail in the same way within minutes of a change window. Where should troubleshooting start?

- | | |
|--|---|
| ☐ A With each branch's physical cabling, optics and patch panels | ☐ B With the RF channel plan at each site |
| ☐ C With the change that was just automated to those sites | ☐ D With the clients' operating systems |

21. Clients can reach websites by IP address but not by name. Which stage is most likely failing?

- | | |
|---|--|
| ☐ A Switch, at the trunk | ☐ B Route, at OSPF |
| ☐ C Serve, at the DNS service | ☐ D Wire, at the physical link |

22. In an exam question asking for the FIRST action, which option is usually best?

- | | |
|--|--|
| ☐ A The safe step that gathers evidence about the likely stage | ☐ B The most expensive hardware replacement |
| ☐ C A reload of every device in the path | ☐ D The longest and most detailed configuration change |

☐ Lesson complete · Date _____ · Knowledge check score ____ / 4

MODULE 8 · AUTOMATE

Answer key

Knowledge checks

- 1. C** Routing protocols build the forwarding decisions, which is the control plane. The data plane forwards packets, the management plane configures and monitors, and the application plane sits above the controller.
- 2. B** Applications talk to the controller through its northbound API. Southbound interfaces connect the controller to devices, and console ports or the data plane are not APIs for applications.
- 3. D** NETCONF is a device configuration protocol used southbound. A status code is part of HTTP responses, DHCP option 43 passes vendor data to clients, and syslog severities classify messages.
- 4. A** The underlay is the routed physical network; overlays (such as VXLAN segments) ride on top of it. APIs and SSIDs are other layers entirely.
- 5. B** GET is the read-only method. POST creates, PUT replaces and PATCH partially updates resources.
- 6. D** 403 is an authorisation failure for a known identity. Missing or invalid credentials give 401, a wrong URL gives 404, and server faults give 5xx codes.
- 7. A** JSON requires double-quoted keys and strings and forbids trailing commas. Single quotes, a trailing comma and unquoted keys or strings are all invalid.
- 8. C** YAML nesting is expressed by indentation using spaces. Braces are optional flow style, semicolons are not YAML syntax, and tabs are not allowed for indentation.
- 9. C** RESTCONF is an HTTPS-based interface to YANG data. It is not SNMP, it depends on YANG rather than replacing it, and it never uses Telnet.
- 10. B** Same end state on repeated runs is idempotency. Statelessness describes REST requests, encapsulation is about headers, and redundancy is about backup paths.
- 11. D** Ansible works through existing management interfaces. It installs nothing on any device type, runs from a control node rather than inside a controller, and does not rely on SNMP traps.
- 12. A** Check mode is a dry run that reports what would change. It does not apply and roll back, it goes further than a syntax check, and it does not force changed results.
- 13. B** The file is a JSON array of objects, which Python loads as a list of dicts. It is parsed rather than left as a string, it is not YAML, and its top level is an array, not an object of lists.
- 14. C** The inventory defines hosts, groups and variables. Playbooks define tasks, templates generate config text, and module docs describe options.
- 15. A** Assurance uses telemetry to verify ongoing compliance with intent. Translation converts intent to policy, activation pushes it, and the business — not the system — defines the intent.
- 16. D** Valid syntax can still violate intent, as with an overly broad `any`. Sequence numbers and CLI acceptance are trivial checks, and tool approval does not remove engineer accountability.
- 17. A** Least privilege limits damage from wrong suggestions or compromise. Full admin rights, sharing break-glass credentials or skipping authentication all enlarge the blast radius.
- 18. B** A summary can omit or distort the crucial event, so engineers confirm key evidence. Summaries are usually faster, are technically feasible, and do not delete source logs.
- 19. D** Reaching the gateway proves the wire, address and wireless access work, so the next stage is routing. DHCP has already succeeded.
- 20. C** A simultaneous fault across sites after a change points to the change itself — check the commit and diff, and roll back if needed. Cabling, RF and client OS problems would not appear everywhere at once.
- 21. C** IP reachability proves layers below DNS work; name resolution is the failing service. A trunk, routing or link fault would also break access by IP.
- 22. A** "First" questions reward evidence-based, low-risk steps. Replacing hardware, reloading devices or making large changes come later, if at all.

Scenarios & case studies

Scenario · Which approach for the school network?

A is best. Correct. It tackles the real causes — inconsistency and unreviewed manual changes — with low cost and builds a foundation for a controller later.

Scenario · A 401 at midnight

A is best. Correct. 401 means authentication failed for that client. The engineer's separate session still being valid points to the job's own credential.

Case study · Rangpur school-district automation project

Start with **read-only** automation: a scheduled playbook or script that collects configurations and reports compliance (NTP server, Telnet disabled, SSH v2, standard ACLs). It is low-risk, builds skills and immediately answers the Telnet question. The minimum source of truth is a version-controlled inventory (YAML) listing each school's devices, management IPs, subnets and VLANs, plus standard values such as NTP and syslog servers. Make changes safe with templates for the standard ACL and baseline, peer review of every change in Git, `--check --diff` dry runs, a pilot on two schools before all 45, and saved pre-change configurations for rollback. Measure: number of devices compliant with the baseline (target 100%), change-related incidents per term, time to apply a district-wide change, and time to build a new school from template.

Scenario · The second run

A is best. Correct. A properly idempotent task would report `ok` on the second run. Compare the pushed lines with the running configuration format to find the mismatch.

Scenario · The confident suggestion

A is best. Correct. Emergency pressure shortens the process but does not remove verification. A wrong change during an outage can make it worse.

Scenario · Integrated troubleshooting

A is best. Correct. Air and authentication succeed, and DHCP works for wired VLAN 10, so the break is at the Switch stage on the AP's trunk: add VLAN 10 to the allowed list.

REFERENCE

Glossary

8

802.11r Fast BSS transition: pre-derived keys that make roaming between APs quick on secured SSIDs. (Air)

802.1Q IEEE standard that inserts a 4-byte tag with a 12-bit VLAN ID into Ethernet frames on trunks. (Switch)

802.1X Port-based network access control using EAP between supplicant, authenticator and authentication server. (Air)

A

AAA Authentication, authorisation and accounting — usually provided centrally by RADIUS or TACACS+. (Guard)

Administrative distance Trust value used to choose between different route sources for the same prefix; lower wins. (Route)

Alternate port RSTP discarding port that provides a ready backup path to the root. (Switch)

ARP Address Resolution Protocol — finds the MAC address for a known IPv4 address on the local network. (Wire)

B

Block size 256 minus the mask value in the interesting octet; the gap between subnet network addresses. (Address)

BPDU guard Feature that err-disables an edge port if it receives a spanning-tree BPDU. (Switch)

Broadcast address The address with all host bits set to 1; reaches every host on the IPv4 subnet. (Address)

Broadcast domain The set of devices reached by a layer 2 broadcast — one VLAN. (Switch)

BSSID The MAC address identifying one SSID on one AP radio. (Air)

C

CAPWAP IETF protocol used between lightweight APs and a WLAN controller (UDP 5246 control, 5247 data). (Air)

CIA triad Confidentiality, integrity and availability — the three core goals of information security. (Guard)

CIDR Classless Inter-Domain Routing — explicit prefix lengths of any size instead of fixed classes. (Address)

Co-channel interference Contention caused by several radios sharing one channel and taking turns to transmit. (Air)

Collapsed core Two-tier campus design where core and distribution roles share one switch pair. (Wire)

Control plane The part of a device that builds forwarding decisions, using protocols such as OSPF, STP and ARP. (Automate)

D

DAI Dynamic ARP inspection: drops ARP messages that do not match trusted IP-to-MAC bindings. (Guard)

Data plane The part of a device that forwards packets and frames. (Automate)

dBm Power relative to 1 milliwatt on a logarithmic scale; 0 dBm = 1 mW, -70 dBm is a typical usable Wi-Fi signal. (Air)

Default route 0.0.0.0/0 or ::/0 — matches any destination not covered by a more specific route. (Route)

DFS Dynamic Frequency Selection: rules that make 5 GHz APs detect radar and vacate some channels. (Air)

DHCP relay A router function (`ip helper-address`) that forwards client DHCP broadcasts as unicast to a server in another subnet, adding the giaddr field. (Serve)

DHCP snooping Switch feature that drops DHCP server messages on untrusted ports and builds an IP-MAC binding table. (Guard)

DORA Discover, Offer, Request, Acknowledge — the four-message DHCP exchange that leases an address to a client. (Serve)

DR / BDR Designated and backup designated routers that reduce adjacencies on multi-access OSPF segments. (Route)

DSCP The 6-bit field in the IP header used to mark a packet's QoS class end to end. (Serve)

Dual stack Running IPv4 and IPv6 at the same time on the same devices and links. (Address)

Duplex Whether a link can send and receive at the same time (full) or only one at a time (half). (Wire)

E

Encapsulation Adding a layer's header (and trailer) around data as it moves down the stack. (Wire)

Ephemeral port Temporary client-side port chosen from the dynamic range for one connection. (Wire)

ESS Extended Service Set: several BSSs with the same SSID joined by a distribution system, enabling roaming. (Air)

EtherType Ethernet field naming the payload protocol, e.g. 0x0800 IPv4, 0x86DD IPv6. (Wire)

EUI-64 Method of building a 64-bit interface ID from a MAC by inserting FFFE and flipping the 7th bit. (Address)

Extended ACL An ACL that matches protocol, source, destination and ports; normally placed close to the source. (Guard)

F

FCS Frame check sequence — a CRC at the end of a frame used to detect corruption. (Wire)

Floating static route Static route with a raised AD that is installed only when the preferred route disappears. (Route)

Flooding Sending a frame out every port in the VLAN except the arrival port (unknown unicast, broadcast). (Switch)

G

Global unicast Publicly routable IPv6 address, currently allocated from 2000::/3. (Address)

I

Idempotency The property that repeating an operation leaves the system in the same final state. (Automate)

Implicit deny The invisible final rule of every ACL that drops anything not matched earlier. (Guard)

Inside global The public address that represents an inside host to the outside world after NAT. (Serve)

Intent-based networking Defining desired outcomes and letting the system translate, apply and continuously assure them. (Automate)

IPS Intrusion prevention system: inline device that detects and blocks malicious traffic. (Guard)

IPsec Suite of protocols (IKE, ESP, AH) that authenticate and encrypt IP traffic, typically for VPNs. (Guard)

J

JSON A text data format of objects, arrays, strings, numbers, booleans and null, with double-quoted strings. (Automate)

L

LACP Link Aggregation Control Protocol — negotiates EtherChannel bundles (active/passive). (Switch)

Link-local (IPv6) fe80::/10 address valid on one link only; every IPv6 interface has one. (Address)

LLDP IEEE 802.1AB vendor-neutral protocol for advertising device identity to direct neighbours. (Switch)

Longest-prefix match Rule that the most specific matching route (longest prefix) is used for forwarding. (Route)

M

MAC address 48-bit Layer 2 interface identifier; first 24 bits usually the manufacturer OUI. (Wire)

MAC address table Switch table mapping VLAN + MAC address to a port; also called the CAM table. (Switch)

Management plane How engineers and tools configure and monitor a device: SSH, SNMP, APIs. (Automate)

Metric Value a routing protocol uses to compare paths; OSPF's metric is cost. (Route)

MFA Multi-factor authentication: proof from two or more categories — know, have, are. (Guard)

MIB / OID The structured database of values an SNMP agent exposes, and the numeric identifier of each value. (Serve)

MLO Multi-link operation: Wi-Fi 7 devices using links on several bands at once. (Air)

N

Native VLAN The VLAN whose frames cross an 802.1Q trunk untagged; must match at both ends. (Switch)

Northbound API The interface applications use to talk to a controller, usually REST with JSON. (Automate)

O

OFDMA Wi-Fi 6+ technique that splits a channel into resource units shared by several clients in one transmission. (Air)

OSPF Open Shortest Path First — link-state interior routing protocol using cost and the SPF algorithm. (Route)

P

Passive interface OSPF interface whose subnet is advertised but on which no hellos are sent. (Route)

PAT Port Address Translation (NAT overload): many inside hosts share one public address, distinguished by port numbers. (Serve)

PDU Protocol data unit — the data unit at a given layer: data, segment, packet, frame, bits. (Wire)

Phishing Social engineering through deceptive messages that steal credentials or install malware. (Guard)

PMF Protected Management Frames (802.11w): authentication of management frames, mandatory in WPA3. (Air)

PoE Power over Ethernet — IEEE 802.3af/at/bt delivery of DC power over twisted-pair cable. (Wire)

Port security Switch feature that limits and learns the MAC addresses allowed on a port. (Guard)

Prefix length The number of leading network bits, written /n (e.g. /26, /64). (Address)

Protocol A set of agreed rules for formatting and exchanging data between devices. (Wire)

R

REST API An HTTP-based interface where resources have URLs and verbs such as GET and POST act on them. (Automate)

RFC 1918 The private IPv4 ranges 10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16. (Address)

Root bridge The STP reference switch with the lowest bridge ID; all paths are measured toward it. (Switch)

Router Advertisement ICMPv6 type 134 message announcing prefixes, flags and a default gateway. (Address)

Router ID 32-bit dotted-decimal value that uniquely identifies an OSPF router. (Route)

Router-on-a-stick Inter-VLAN routing using one router trunk link with one subinterface per VLAN. (Switch)

Routing table A router's list of destination prefixes with next hops, exit interfaces, AD and metric. (Route)

S

SAE Simultaneous Authentication of Equals: WPA3-Personal key exchange that resists offline password guessing. (Air)

SCP / SFTP File transfer protocols that run over SSH, giving authentication and encryption. (Serve)

SDN controller Software that holds the network-wide view and programs devices through APIs. (Automate)

Shaping Buffering excess traffic to send it later at a target rate, smoothing bursts. (Serve)

SLAAC Stateless address autoconfiguration — hosts build IPv6 addresses from RA prefixes. (Address)

SNMP trap An unsolicited notification an agent sends to a manager on UDP 162; an inform is the acknowledged version. (Serve)

SNR Signal-to-noise ratio: received signal minus noise floor, in dB. Higher is better. (Air)

Source of truth The authoritative, version-controlled record of how the network should be. (Automate)

Southbound interface How a controller talks to devices: NETCONF, RESTCONF, gNMI, OpenFlow or CLI. (Automate)

Spine-leaf Data-centre topology where every leaf switch connects to every spine switch. (Wire)

Static route Manually configured route that does not change unless an engineer changes it. (Route)

Stratum An NTP server's distance from a reference clock; 16 means unsynchronised. (Serve)

SVI Switched virtual interface — a layer 3 interface for a VLAN on a layer 3 switch. (Switch)

Syslog severity A level from 0 (Emergency) to 7 (Debugging) attached to each log message. (Serve)

T

TACACS+ AAA protocol on TCP 49 that encrypts the full payload and supports per-command authorisation. (Guard)

Three-way handshake TCP connection set-up: SYN, SYN-ACK, ACK. (Wire)

Trust boundary The network point where QoS markings start being accepted rather than rewritten. (Serve)

TTL IP time-to-live field, decreased by each router; traceroute uses it to find hops. (Wire)

TTL (DNS) How long a resolver may cache a DNS answer before asking again. (Serve)

U

Unique local Private IPv6 range fc00::/7, used as fd00::/8. (Address)

V

VLAN A logical layer 2 network on shared switches; normally one IP subnet per VLAN. (Switch)

VLSM Variable-length subnet masking — different prefix lengths inside one address block. (Address)

VRP Open-standard first-hop redundancy protocol that shares a virtual gateway IP between routers. (Route)

Vulnerability A weakness in a system, configuration or process that a threat could use. (Guard)

W

Wildcard mask Inverse mask used in OSPF network statements and ACLs; 0 = match, 1 = ignore. (Route)

Y

YAML A human-friendly data format where indentation defines structure; used by Ansible. (Automate)

YANG A data-modelling language that describes device configuration and state for NETCONF and RESTCONF. (Automate)

Z

Zero trust Security model that grants no implicit trust based on network location and verifies every request. (Guard)

CAPSTONE

Capstone: Build a branch office network end to end

Design, configure and document the network for a realistic 60-person branch office (for example a bank branch in Chattogram or a software house in Dhaka), using a virtual lab or packet simulator.

Deliverables

- ☐ Requirements summary and simple topology diagram
- ☐ IPv4 addressing plan with VLSM and an IPv6 plan
- ☐ VLAN, trunk and spanning-tree design with configurations
- ☐ Inter-VLAN routing and static/OSPF routing with verification output
- ☐ DHCP, DNS, NAT/PAT and NTP services configured
- ☐ Secure small-office Wi-Fi design (SSIDs, WPA3, channels)
- ☐ ACLs, port security and device hardening checklist
- ☐ A small automation script or playbook that backs up configurations
- ☐ Troubleshooting log for at least three faults you introduced and fixed
- ☐ One-page handover document for the support team


Self-assessment rubric

CRITERION	DEVELOPING	PROFICIENT	EXCELLENT
Addressing	○ Errors or overlaps	○ Correct plan	○ Efficient VLSM plus IPv6 with growth room
Switching	○ VLANs only	○ VLANs, trunks and STP correct	○ Plus protection features and verified root placement
Routing	○ Partial reachability	○ Full reachability	○ Verified OSPF with sensible design choices
Services	○ Some services missing	○ All services work	○ Documented and resilient (relay, redundancy)
Wireless	○ Not addressed	○ Working design	○ Secure, interference-aware design
Security	○ Defaults left	○ ACLs and hardening applied	○ Least-privilege policy with evidence
Automation	○ None	○ Script runs	○ Reusable, documented and safe
Troubleshooting & docs	○ No evidence	○ Log and handover present	○ Clear, reproducible and professional

CAPSTONE

Planning pages

Sketch your design, timeline or architecture here.



EXAM PREPARATION

Practice exam - 60 questions

The online certification exam draws 40 of these questions at random, with a 60-minute limit and a 80% pass mark. Practise here first: time yourself, mark your answers, then check the key and read every explanation, including for the questions you got right.

Wire — Network Fundamentals

1. A 45 W PTZ camera is connected to a switch port that supplies 802.3at (PoE+). It boots but restarts whenever it pans and zooms. What should the engineer do?
 - ☐ A Hard-set the port to 100 Mbps full duplex
 - ☐ B Put the camera in its own voice VLAN
 - ☐ C Enable PortFast and BPDU guard on the port
 - ☐ D Move it to an 802.3bt port that can supply 60 W
2. A backup server's switch port shows `Gi1/0/9 connected 20 a-half a-100` while the server NIC is hard-set to 100 Mbps full duplex. Backups are extremely slow. What is the best fix?
 - ☐ A Replace the server NIC with a faster 10 Gbps network card
 - ☐ B Move the server into a separate VLAN
 - ☐ C Set both ends to auto, or hard-set both identically
 - ☐ D Increase the switch MAC aging timer
3. A new 10 Gbps fibre link between two floors will not come up. One end has a 10GBASE-SR optic and the other a 10GBASE-LR optic on OM4 multimode cable. What is the problem?
 - ☐ A The optics do not match each other or the cable
 - ☐ B OM4 cannot carry 10 Gbps over any distance
 - ☐ C The link needs PoE to power the optics
 - ☐ D Both fibre ends must be set to half duplex at 10 Gbps
4. A capture on a web server shows a packet from client 10.40.1.20, but the frame's source MAC belongs to the server's own gateway router. Why?
 - ☐ A The client is spoofing the router's MAC address
 - ☐ B Routers build a new frame with their own MAC on each hop
 - ☐ C NAT has replaced the client's MAC address
 - ☐ D The access switch rewrote the source MAC to save MAC table space
5. A developer is building an app that sends a small sensor reading every second. Occasional loss is acceptable, but overhead must be minimal. Which transport fits best?
 - ☐ A TCP with a large receive window
 - ☐ B TCP with keepalives enabled
 - ☐ C UDP
 - ☐ D ICMP echo requests
6. Security policy says switches must be managed only over an encrypted CLI. Which rule should the management firewall allow?
 - ☐ A TCP 23 to the switch management addresses
 - ☐ B UDP 161 to the switch management addresses
 - ☐ C TCP 80 to the switch management addresses
 - ☐ D TCP 22 to the switch management addresses
7. A university with eight buildings and 4,000 users wants a campus design where adding a building does not change the existing backbone. Which design fits?
 - ☐ A Two-tier collapsed core in the main building
 - ☐ B Three-tier: access, distribution per building, and a core
 - ☐ C A single large switch stack for the campus
 - ☐ D Daisy-chained access switches between buildings
8. A Linux PC cannot reach anything. `ip -br addr` shows `eth0 DOWN` and the switch shows the port as `notconnect`. What should be checked first?
 - ☐ A The cable, wall jack and patch-panel connection
 - ☐ B The DNS server configured on the PC
 - ☐ C The default route on the core router
 - ☐ D The OSPF neighbour states on the head-office WAN routers

Address — IPv4 & IPv6 Addressing

9. A host is configured as `10.10.5.130/25` with gateway `10.10.5.1`. It reaches local hosts but nothing remote. What is wrong?
 - ☐ A The host is using the subnet's broadcast address
 - ☐ B A /25 mask cannot be used on a host
 - ☐ C The host needs a public IP address to route
 - ☐ D The gateway is outside the host's subnet

10. A new VLAN must hold 120 devices and the address space is tight. Which prefix should be used?
- (A) /25 (B) /24
(C) /26 (D) /27
11. 192.168.100.0/25 is already allocated to Lab A. The next segment needs 50 hosts. Which subnet should it get?
- (A) 192.168.100.64/26 (B) 192.168.100.128/26
(C) 192.168.100.100/26 (D) 192.168.100.128/25
12. On-premises networks use 192.168.0.0/16. A team proposes 192.168.0.0/16 for a new cloud network that will be connected by VPN. What should you recommend?
- (A) Keep the same block so firewall rules can be reused (B) Use 169.254.0.0/16 so the cloud is isolated
(C) Use a non-overlapping block such as 10.200.0.0/16 (D) Use 203.0.113.0/24 because it is never routed
13. Which is a correct compressed form of `2001:0db8:0000:0000:0000:ff00:0042:8329` ?
- (A) 2001:db8:0:0::ff00:0042:8329:0 (B) 2001:db8::ff:42:8329
(C) 2001:db8::ff00:42:8329 (D) 2001:db8:::ff00:42:8329
14. Hosts on an IPv6 VLAN form global addresses with SLAAC but have no DNS server. The router sends RAs with only the A flag set. What is the best fix?
- (A) Set the O flag and provide stateless DHCPv6 for DNS (B) Set the M flag and remove the prefix information from the RA
(C) Change the VLAN prefix from /64 to /56 (D) Disable SLAAC and configure static IPv4 DNS
15. A router interface has MAC `0c1d.34aa.bbcc` and the command `ipv6 address 2001:db8:5:1::/64 eui-64`. What address does it get?
- (A) 2001:db8:5:1:c1d:34ff:feaa:bbcc (B) 2001:db8:5:1:e1d:34fe:ffaa:bbcc
(C) 2001:db8:5:1:c1d:34aa:ffe:bbcc (D) 2001:db8:5:1:e1d:34ff:feaa:bbcc
16. What is the broadcast address of the subnet that contains 10.44.77.200/21?
- (A) 10.44.77.255 (B) 10.44.79.255
(C) 10.44.71.255 (D) 10.44.80.255

Switch — Switching & VLANs

17. SW1 logs a native VLAN mismatch on its trunk to SW2. `show interfaces trunk` shows native VLAN 1 on SW1 and 99 on SW2. Some untagged management traffic appears in the wrong VLAN. What should you do?
- (A) Remove VLAN 99 from the allowed list on SW2 (B) Convert the trunk to an access port in VLAN 1
(C) Enable PortFast on both trunk ports (D) Set the same unused native VLAN on both ends
18. PCs in VLAN 20 cannot reach their gateway 10.10.20.1. The router shows `interface Gi0/0.20`, `encapsulation dot1Q 30`, `ip address 10.10.20.1 255.255.255.0`. What is the fault?
- (A) Subinterface numbers must start at 100 (B) The gateway must use the last usable address
(C) The subinterface is tagged for VLAN 30, not VLAN 20 (D) Router-on-a-stick cannot route VLAN 20
19. SW1 is root for VLAN 10. SW2 and SW3 each have a direct 1 Gbps link to SW1 and a 1 Gbps link to each other. All priorities are 32768; SW2's MAC is lower than SW3's. Which port is blocked?
- (A) SW3's port toward SW2 (B) SW2's port toward SW3
(C) SW2's port toward SW1 (D) SW3's port toward SW1
20. A student connects a patch cable between two wall ports in a lab. The VLAN floods with traffic and MAC addresses flap between ports. Which feature on edge ports would have stopped this?
- (A) LLDP on every access port (B) BPDU guard with PortFast
(C) A higher MAC aging timer (D) LACP active mode on access ports

21. An EtherChannel will not form. SW1 members use `channel-group 1 mode on` and SW2 members use `channel-group 1 mode active`. What is the fix?
- (A) Change SW2 to `mode passive`
 - (B) Add more member links to SW2
 - (C) Set SW1's native VLAN to 1
 - (D) Run LACP on both ends
22. You search an access switch for a laptop's MAC and find it on Gi1/0/48, which is the uplink to the distribution switch. What should you do next?
- (A) Shut down Gi1/0/48 to isolate the laptop
 - (B) Clear the whole MAC table and wait
 - (C) Follow the uplink to the next switch and repeat the search
 - (D) Assume the laptop is plugged directly into port Gi1/0/48 itself
23. A PC on Gi1/0/7 receives an address from the wrong subnet. `show vlan brief` lists Gi1/0/7 under VLAN 1, while the plan says VLAN 10. What fixes it?
- (A) `switchport access vlan 10` on Gi1/0/7
 - (B) `switchport trunk native vlan 10` on Gi1/0/7
 - (C) `spanning-tree portfast` on Gi1/0/7
 - (D) `no switchport` on Gi1/0/7

Route — Routing Essentials

24. R1 has `0.0.0.0/0` via A, `10.0.0.0/8` via B, `10.20.0.0/16` via C and `10.20.30.0/24` via D. Where does a packet to 10.20.31.5 go?
- (A) Via D
 - (B) Via C
 - (C) Via B
 - (D) Via A
25. An engineer adds `ip route 10.50.0.0 255.255.0.0 10.0.12.9 120` while OSPF already provides 10.50.0.0/16. Which route is installed?
- (A) The static route, because static routes always win over OSPF
 - (B) Both, with equal-cost load sharing
 - (C) Neither, until one of them is removed
 - (D) The OSPF route, because 110 is lower than 120
26. A branch has a primary static default to ISP-1 and a floating static (AD 10) to ISP-2. ISP-1's network fails beyond its modem, but the router's link to the modem stays up. Traffic is lost. Why?
- (A) The floating route needs an AD lower than 1
 - (B) Floating static routes only work for IPv6
 - (C) Its interface stays up, so the primary route stays
 - (D) The branch needs a separate /32 host route for every website
27. Two OSPF routers on a link stay in `EXSTART/EXCHANGE`. Hellos, area and subnet match. One interface has `mtu 1400`. What is the fix?
- (A) Make the interface MTU match on both routers
 - (B) Change one router's OSPF process ID
 - (C) Set both interface priorities to 0
 - (D) Raise the reference bandwidth on both routers
28. On one Ethernet segment, R1 (priority 1, RID 1.1.1.1), R2 (priority 10, RID 2.2.2.2) and R3 (priority 0, RID 9.9.9.9) elect a DR. Later R4 joins with priority 50. Which router is DR after R4 joins?
- (A) R4
 - (B) R2
 - (C) R3
 - (D) R1
29. After an engineer adds `passive-interface default` under `router ospf 1` on R2, all of R2's neighbours disappear. What else was needed?
- (A) `default-information originate` on R2
 - (B) A new router ID on R2
 - (C) `ip ospf priority 0` on R2's uplinks
 - (D) `no passive-interface` on R2's uplink interfaces
30. Both VRRP routers for VLAN 10 show `State Master` in `show vrrp brief`, and hosts report intermittent gateway problems. What is the most likely cause?
- (A) Both routers are using the default priority of 100
 - (B) They cannot hear each other's VRRP advertisements
 - (C) Preemption is enabled on both routers
 - (D) The virtual IP is the first address in the subnet

Serve — IP Services

31. PCs in a new VLAN 60 (192.168.60.0/24) receive 169.254.x.x addresses. The router subinterface Gi0/1.60 has `ip address 192.168.60.1 255.255.255.0` and `ip helper-address 10.0.50.10`. The DHCP server at 10.0.50.10 shows no Discover packets arriving from 192.168.60.0/24 scope. Other VLANs work. What should you check next?
- (A) Whether the DHCP server's lease time is longer than seven days
 - (B) Whether the router has `no ip domain-lookup` configured globally
 - (C) Whether VLAN 60 is allowed on the switch-to-router trunk
 - (D) Whether the PCs have the correct DNS suffix configured locally
32. Users on 192.168.10.0/24 browse the Internet, but users on the newly added 192.168.11.0/24 cannot. Both subnets reach internal servers.
- (A) The outside interface is configured with `ip nat inside` instead of outside
 - (B) The NAT ACL does not match the new 192.168.11.0/24 subnet
 - (C) The default route points to the wrong next hop at the ISP side
 - (D) The DNS server is not resolving names for the new subnet's users
33. Output from an edge router:
- ```
Pro Inside global Inside local Outside local Outside global
tcp 203.0.113.1:4096 192.168.1.50:51000 198.51.100.9:443 198.51.100.9:443
```
- Which address will the web server 198.51.100.9 see as the source of the connection?
- (A) 192.168.1.50, port 51000
  - (B) 198.51.100.9, port 443
  - (C) 203.0.113.1, port 51000
  - (D) 203.0.113.1, port 4096
34. An engineer configures `logging trap informational` on a router. Which messages will be sent to the syslog server?
- (A) Severity levels 0 through 6
  - (B) Severity level 6 only
  - (C) Severity levels 6 and 7
  - (D) Severity levels 0 through 7
35. `show ntp status` on a switch reports `Clock is unsynchronized, stratum 16, no reference clock`. The configured NTP server is reachable by ping. What is the most likely cause?
- (A) UDP 123 is blocked or the server itself is unsynchronised
  - (B) The switch needs SNMPv3 enabled before NTP will start
  - (C) The switch's syslog level is set too high to accept time
  - (D) The DHCP scope does not include an NTP option for switches
36. A security audit finds `snmp-server community public RO` on 80 switches. Monitoring must continue. What is the best remediation?
- (A) Rename the community string to `private` on every switch
  - (B) Use SNMPv3 authPriv users and remove the community
  - (C) Change the community from RO to RW for better control
  - (D) Move SNMP polling from UDP 161 to UDP 162 on all switches
37. Voice calls across a branch's 100 Mbps WAN break up during large file uploads. Phones mark EF, and the WAN router has a policy with `priority percent 20` for EF applied **inbound** on the WAN interface. What should be changed?
- (A) Remark voice traffic from EF to AF41 on the LAN
  - (B) Apply the queuing policy inbound on every LAN port
  - (C) Apply the queuing policy outbound on the WAN interface
  - (D) Raise the voice priority queue from 20 to 100 percent
38. A remote router must be upgraded across the Internet. Which approach best protects the image and the session?
- (A) Copy the image with TFTP and check its file size afterwards
  - (B) Copy the image with FTP using a long, complex password
  - (C) Copy the image over Telnet after enabling privilege level 15
  - (D) Copy it by SCP and verify its hash before booting

## Air — Wireless Networking

39. A survey reading shows signal -74 dBm and noise -90 dBm at a reception desk where staff make Wi-Fi calls. What does this mean for voice?
- (A) SNR is 164 dB, far above the voice target
  - (B) SNR is 16 dB, comfortably above the voice target
  - (C) SNR is -16 dB, so the reading must be invalid
  - (D) SNR is 16 dB, below the usual voice target

40. A small office has three APs. Their 2.4 GHz radios are on channels 1, 4 and 8. Users report slow 2.4 GHz performance. What should be done?
- (A) Move all three radios to channel 6 together
  - (B) Re-plan the radios to channels 1, 6 and 11
  - (C) Widen the 2.4 GHz channels to 40 MHz each
  - (D) Increase each 2.4 GHz radio to maximum power
41. A Wi-Fi 6E AP broadcasts an SSID on 6 GHz. An engineer tries to set the SSID to WPA2-Personal for compatibility, but the 6 GHz radio refuses to enable it. Why?
- (A) 6 GHz operation requires WPA3 or Enhanced Open with PMF
  - (B) 6 GHz radios can broadcast only hidden SSIDs by default
  - (C) WPA2 requires 802.3bt power to run at 6 GHz
  - (D) WPA2 works at 6 GHz only with 20 MHz channels
42. Lightweight APs at a branch cannot join the controller in head office. A firewall between them permits SSH, HTTPS and SNMP. What should be permitted?
- (A) RADIUS on UDP 1812 and 1813
  - (B) IPsec IKE on UDP 500 and 4500
  - (C) CAPWAP on UDP 5246 and 5247
  - (D) Syslog on UDP 514 and 162
43. Clients using an 802.1X SSID experience two-second gaps in calls each time they move between APs on the same VLAN. Coverage overlap meets the design. Which feature most directly addresses this?
- (A) 802.11d country information
  - (B) 802.11h dynamic frequency selection
  - (C) 802.11r fast BSS transition
  - (D) 802.11w protected management frames
44. A 60-person company still uses one WPA2-Personal passphrase that several former employees know. Management wants per-user control and the ability to remove individuals instantly. What should be deployed?
- (A) WPA3-Enterprise with 802.1X and a RADIUS server
  - (B) WPA3-Personal with a new shared passphrase
  - (C) WPA2-Personal with MAC address filtering added
  - (D) An open SSID with the network name hidden
45. A new AP installed on an 802.3af switch port boots, but its management page shows one radio disabled and USB power off. The design specifies 802.3at. What is the likely cause?
- (A) The trunk to the AP is missing the guest VLAN on this port
  - (B) The port cannot supply enough power for the AP
  - (C) The AP's country code is wrong and is blocking one radio
  - (D) The SSID uses WPA3 and the AP needs firmware for WPA2

## Guard — Security Fundamentals

46. An engineer applies `ip access-group 110 in` on the LAN interface. ACL 110 contains only `deny tcp 192.168.5.0 0.0.0.255 any eq 23`. Shortly afterwards, every user on that LAN loses all connectivity. Why?
- (A) Telnet blocking automatically blocks all TCP traffic on the interface
  - (B) Extended ACLs must always be applied outbound to work correctly
  - (C) ACL numbers between 100 and 199 are reserved for standard ACLs
  - (D) The implicit deny drops everything not explicitly permitted
47. Which wildcard mask matches only hosts in 10.10.8.0/22?
- (A) 0.0.3.255
  - (B) 0.0.7.255
  - (C) 255.255.252.0
  - (D) 0.0.1.255
48. A router has:
- ```
ip access-list standard MGMT
permit 10.0.99.0 0.0.0.255
interface Gi0/0
ip access-group MGMT in
```
- Engineers from 10.0.99.0/24 can still SSH to the router, but so can users in 192.168.20.0/24 who reach it through Gi0/2. What should be done?
- (A) Convert the ACL to an extended ACL applied outbound on Gi0/0
 - (B) Add `deny any` before the permit statement in the ACL
 - (C) Apply the ACL to the VTY lines with `access-class MGMT in`
 - (D) Apply the ACL to Gi0/0 outbound instead of inbound

49. Users on VLAN 20 suddenly receive the default gateway 192.168.20.254 from an unknown DHCP server. Which switch feature prevents this?
- (A) Dynamic ARP inspection with the uplink trusted (B) Port security with sticky MAC addresses
(C) Spanning tree root guard on the uplink (D) DHCP snooping with only the uplink trusted
50. After enabling dynamic ARP inspection on VLAN 30, DHCP clients work, but a printer with a static IP address cannot be reached. What is the best fix?
- (A) Disable DHCP snooping on VLAN 30 permanently (B) Permit the printer's IP and MAC with an ARP ACL for DAI
(C) Trust every access port in VLAN 30 for DAI (D) Enable port security on the printer's port
51. `show interfaces status err-disabled` lists Gi0/9 with reason `psecure-violation`. The port has `switchport port-security maximum 1` and a sticky MAC. A user connected a small unmanaged switch with two PCs. What is the correct sequence?
- (A) Raise the maximum to 50 and leave the switch connected (B) Re-enable the port first, then investigate the violation
(C) Disable port security on every access port in the whole building (D) Remove the unmanaged switch, then shut and re-enable the port
52. A company wants engineers' router logins authorised per command and every configuration command recorded centrally. Which choice fits best?
- (A) RADIUS with a shared password for all engineers (B) Local accounts on each router with type 7 passwords
(C) TACACS+ with command authorisation and accounting (D) SNMPv3 users with authPriv on every router
53. A site-to-site IPsec tunnel to a branch behind an ISP's NAT device fails to pass data, although IKE negotiation completes. Which traffic must the path permit?
- (A) IP protocol 51 for AH encryption (B) UDP 4500 for NAT traversal of ESP
(C) TCP 49 for the IKE key exchange messages (D) UDP 1812 for IPsec key exchange

Automate — Automation, Programmability and Readiness

54. A script sends a request to the controller to add a new site and receives `201 Created`. Which HTTP verb did it most likely use?
- (A) POST (B) GET
(C) DELETE (D) HEAD
55. A playbook run reports:
- ```
r1 : ok=3 changed=0
r2 : ok=3 changed=2
```
- The same playbook is immediately run again with no other changes. What should a well-written, idempotent playbook report for r2?
- (A) changed=2, because the same tasks run again (B) failed=2, because the lines already exist  
(C) changed=0, because r2 now matches the desired state (D) unreachable=1, because the device is locked
56. Which YAML snippet correctly defines a list of two VLAN IDs under the key `vlan`?
- (A) `vlan:` followed by indented lines `- 10` and `- 20` (B) `vlan:` followed by tab-indented lines `- 10` and `- 20`  
(C) `vlan = (10; 20)` on one line with parentheses (D) `vlan:` followed by unindented lines `10,` and `20,`
57. In a controller-based architecture, an IT service-management platform asks the controller to apply a QoS policy to 40 sites. Which interface carries this request?
- (A) The southbound NETCONF session to each switch (B) The northbound REST API of the controller  
(C) The data plane of the WAN routers (D) The console port of the controller
58. A RESTCONF GET to a router returns `401 Unauthorized`. The URL and model name were copied from working documentation. What should be checked first?
- (A) Whether the JSON body has a trailing comma in it (B) Whether the YANG model exists on the router at all  
(C) Whether the router's NTP server is reachable now (D) The credentials or token sent with the request

59. An AI-assisted tool suggests a command to fix a routing loop during an outage. The engineer has never seen the command before. What should happen first?
- ☐ A Apply it immediately because outages need fast action
  - ☐ B Verify the command and diagnosis against evidence
  - ☐ C Apply it to all routers so the network stays consistent
  - ☐ D Reload the router so the command takes effect cleanly
60. Laptops on the staff SSID get valid addresses and can resolve names, but cannot open the Internet; wired PCs on the same VLAN can. Using the Packet Journey, which stage should you examine first?
- ☐ A Guard — an ACL or policy applied to wireless traffic
  - ☐ B Wire — the physical cable to the laptops
  - ☐ C Address — the DHCP scope for the VLAN
  - ☐ D Route — the default route on the edge

## EXAM PREPARATION

# Practice exam answers & explanations

**1. D** PoE+ provides 30 W at the port, well below the camera's 45 W peak, so it browns out under load; an 802.3bt Type 3 port (60 W) fixes it. Speed/duplex, VLAN choice and spanning-tree features do not affect the power available.

**2. C** The switch fell back to half duplex because the server does not negotiate, causing a duplex mismatch. Matching both ends removes the collisions. A faster NIC, a new VLAN or MAC aging changes leave the mismatch in place.

**3. A** SR (850 nm, multimode) and LR (1310 nm, single-mode) are different standards; both ends must use the same type, matched to the fibre. OM4 supports 10 Gbps short-reach, optics are powered by the switch, and 10G Ethernet is full duplex only.

**4. B** Layer 2 headers are rebuilt at every router, so the last router's MAC appears as the source while IP addresses stay end to end. This is normal, not spoofing; NAT changes IP addresses, not MACs; and switches do not rewrite source MACs.

**5. C** UDP's 8-byte header and lack of connection set-up suit frequent small messages where retransmission is not needed. TCP options add reliability and overhead the app does not need, and ICMP is not an application transport.

**6. D** SSH (TCP 22) provides an encrypted CLI. Telnet (23) is clear text, SNMP (UDP 161) is a monitoring protocol, and HTTP (80) is unencrypted web access.

**7. B** A separate core joins per-building distribution blocks, so new buildings plug in as new blocks. A collapsed core suits one building, one stack is a single point of failure, and daisy-chains make buildings depend on each other.

**8. A** A down link on both ends is a Layer 1 problem, so start with the physical path. DNS, core routing and OSPF only matter once the link is up.

**9. D** 10.10.5.130/25 is in 10.10.5.128 – 10.10.5.255, so 10.10.5.1 is not on-link; the gateway must be inside that range (e.g. 10.10.5.129). .130 is a normal host, /25 is valid, and private addresses route internally.

**10. A** A /25 provides 126 usable hosts, the smallest that fits 120. A /24 wastes 128 addresses, while /26 (62) and /27 (30) are too small.

**11. B** The /25 uses .0 – .127, so the next free /26 starts at .128. .64/26 overlaps Lab A, .100 is not a /26 boundary, and a /25 is larger than needed.

**12. C** Overlapping ranges make routing across the VPN ambiguous. Reusing rules is not worth broken routing, link-local space is not routable, and documentation space must not be used in production.

**13. C** Drop leading zeros and replace the run of three zero hexets with ::. The first adds a hexet, the second wrongly shortens ff00 (only leading zeros may be removed), and ::: is invalid.

**14. A** The O flag tells hosts to get other settings such as DNS from DHCPv6 while keeping SLAAC addresses (RDNSS in the RA is another option). The M-flag change rebuilds addressing unnecessarily, /56 breaks SLAAC, and IPv4 DNS settings do not configure IPv6 clients.

**15. D** Insert FFFE in the middle (0c1d:34ff:feaa:bbcc) and flip the 7th bit: 0c becomes 0e. The second option skips the flip, the third reverses FFFE, and the fourth inserts it in the wrong place.

**16. B** /21 gives a block of 8 in the third octet: 77 sits in 72 – 79, so the subnet is 10.44.72.0 and the broadcast 10.44.79.255. .77.255 treats it as /24, .71.255 is the previous subnet's broadcast, and .80.255 is in the next subnet.

**17. D** The native VLAN must match so untagged frames stay in one VLAN. Removing VLAN 99 or converting to access breaks the multi-VLAN link, and PortFast does not belong on trunks between switches.

**18. C** The dot1Q VLAN ID decides which VLAN a subinterface serves; it must be 20. Subinterface numbers are labels, gateways can be any usable address, and router-on-a-stick supports any VLAN ID.

**19. A** Both have equal cost to the root on the SW2–SW3 link, so the tie goes to the lower bridge ID: SW2's port becomes designated and SW3's becomes alternate (blocked). Ports toward the root are root ports and forward.

**20. B** With BPDU guard, an edge port that receives the looped BPDU is err-disabled, breaking the loop. LLDP only reports neighbours, MAC aging does not stop loops, and LACP builds bundles between switches.

**21. D** Static **on** sends no LACP, so it cannot pair with an LACP side; both must run LACP (active/active or active/passive) or both be static. Passive still requires LACP from the other end, extra links do not help, and native VLAN is unrelated.

**22. C** An uplink port means the laptop is further along; use LLDP to identify the neighbour and search its table. Shutting the uplink cuts off many users, clearing tables only delays you, and a trunk is not the laptop's access port.

**23. A** The port is in the default VLAN; assigning it to VLAN 10 puts the PC in the right subnet. Native VLAN settings apply to trunks, PortFast only speeds up forwarding, and **no switchport** turns it into a routed port.

**24. B** 10.20.31.5 is outside 10.20.30.0/24 but inside 10.20.0.0/16, the longest matching prefix. D's /24 does not contain it, and B's /8 and the default are shorter matches.

**25. D** For the same prefix, the lower AD wins; the static's AD was raised to 120, so it floats behind OSPF (110). Statics win only with their default AD of 1, routes from different sources do not load-share, and conflicts are resolved by AD.

**26. C** Static routes are only removed when the exit interface or next hop is lost locally; reachability tracking is needed to detect remote failure. An AD below 1 is impossible for statics, floating statics work for IPv4, and per-site host routes are not a design.

**27. A** MTU mismatches stall the database exchange. Process IDs are local, priority 0 only affects DR election, and reference bandwidth changes costs, not adjacency.

**28. B** R2 wins the original election on priority. DR election is not pre-emptive, so R4 does not take over when it joins. R3 can never be DR with priority 0, and R1 has the lowest non-zero priority.

**29. D** Passive-interface default stops hellos everywhere; uplinks must be re-enabled with `no passive-interface`. Originating a default, a new router ID or priority 0 do not restore hellos.

**30. B** Two masters means advertisements to 224.0.0.18 are not getting through (for example a VLAN or filtering problem between them). Equal priorities are resolved by IP address, preemption is normal, and .1 is a common virtual IP.

**31. C** The relay is configured but the server sees nothing, so the broadcasts are probably not reaching the router — a missing VLAN on the trunk is the classic cause. Lease time and DNS suffix do not stop a Discover reaching the server, and `no ip domain-lookup` only affects how the router resolves names.

**32. B** Because the original subnet still reaches the Internet, the NAT roles and default route must be working. The only per-subnet NAT element is the ACL in `ip nat inside source list`. A DNS failure would also affect access by IP, and the question says browsing fails only for the new subnet.

**33. D** The inside global address and port are what outside hosts see after PAT. 192.168.1.50:51000 is the private inside local address, 198.51.100.9:443 is the server itself, and the translated port in the table is 4096, not 51000.

**34. A** A configured level sends that level and every more severe (lower-numbered) level: 0–6. It is not limited to one level, and level 7 (debugging) is excluded because it is less severe than informational.

**35. A** Ping proves IP reachability but not that UDP 123 is permitted or that the server has valid time to give. NTP does not depend on SNMP or syslog settings, and a switch with a configured `ntp server` does not need a DHCP option.

**36. B** SNMPv3 with authentication and encryption fixes the clear-text, guessable credential. `private` is equally guessable, RW adds write risk, and port 162 is for traps rather than polling.

**37. C** Congestion happens when packets leave the slower WAN interface, so queuing must be applied outbound there. Remarking voice to AF41 removes priority, LAN inbound policies cannot schedule the WAN, and a 100% priority queue could starve all other traffic.

**38. D** SCP authenticates and encrypts, and a hash comparison proves integrity. TFTP has no authentication or encryption and file size does not prove integrity, FTP exposes credentials in clear text, and Telnet is unencrypted.

**39. D**  $SNR = -74 - (-90) = 16$  dB. Voice designs typically aim for about 25 dB with -67 dBm or better, so this is too low. Adding magnitudes gives the meaningless 164, and subtracting in reverse gives a negative value.

**40. B** Channels 1, 4 and 8 overlap, causing adjacent-channel interference; 1, 6 and 11 do not. One shared channel adds contention, 40 MHz in 2.4 GHz consumes most of the band, and more power increases interference.

**41. A** Rules for 6 GHz exclude WPA2 and legacy security, requiring WPA3 or OWE with protected management frames. Hidden SSIDs, PoE class and channel width are unrelated to the security mode allowed.

**42. C** Lightweight APs use CAPWAP control (5246) and data (5247) to reach the controller. RADIUS is for client authentication, IKE is for VPN negotiation, and syslog and SNMP traps are monitoring traffic.

**43. C** 802.11r pre-derives keys so the client avoids a full 802.1X exchange on each roam. 802.11d advertises regulatory domain, 802.11h handles DFS and power control, and 802.11w protects management frames.

**44. A** Enterprise mode gives each user their own credential, so access can be revoked individually. A new shared passphrase still leaks with the next leaver, MAC filtering is easily bypassed, and a hidden open SSID is not access control.

**45. B** 802.3af supplies up to 15.4 W, so a PoE+ AP falls back to a reduced-power mode. A missing VLAN affects traffic, not radio power state; a wrong country code changes allowed channels, and security mode does not disable hardware.

**46. D** With no permit statement, the implicit deny drops everything. Denying TCP 23 does not affect other ports, extended ACLs work in either direction, and 100–199 is the extended range.

**47. A** A /22 mask is 255.255.252.0; subtracting it from 255.255.255.255 gives 0.0.3.255. 0.0.7.255 is a /21, 255.255.252.0 is the subnet mask itself, and 0.0.1.255 is a /23.

**48. C** Restricting management sessions to the router is done on the VTY lines with `access-class`, which covers every interface. Interface ACLs only filter traffic on that interface, and placing `deny any` first would block the engineers too.

**49. D** DHCP snooping drops server messages (Offer, Ack) arriving on untrusted ports, silencing a rogue server. DAI validates ARP, port security limits MAC addresses, and root guard protects STP — none filter DHCP offers.

**50. B** Static hosts have no DHCP snooping binding, so DAI needs an ARP ACL entry for them. Disabling snooping breaks DAI, trusting every port removes protection, and port security does not affect ARP validation.

**51. D** Remove the cause first, then recover the port with shutdown/no shutdown. Raising the limit legitimises an unauthorised device, re-enabling first simply triggers another violation, and disabling port security removes the control everywhere.

**52. C** TACACS+ separates authorisation and supports per-command authorisation and command accounting. RADIUS lacks per-command control and a shared password removes accountability, local type 7 passwords are weak and not central, and SNMP is not a login protocol.

**53. B** When NAT is detected, ESP is encapsulated in UDP 4500. AH (protocol 51) provides no encryption and does not survive NAT, TCP 49 is TACACS+, and UDP 1812 is RADIUS.

**54. A** POST creates a new resource, and 201 confirms creation. GET reads, DELETE removes, and HEAD retrieves headers only.

**55. C** Idempotent tasks change nothing when the device already matches. Reporting changed again would indicate a non-idempotent task, existing lines are not failures, and nothing suggests the device is unreachable.

**56. A** YAML lists use a dash and space under an indented key. Tabs are not allowed for indentation, the parenthesis form is not YAML, and unindented values with commas are not list items.

**57. B** Applications talk to the controller through its northbound API. Southbound interfaces connect controller to devices, the data plane forwards user traffic, and the console is a local management port.

**58. D** 401 means authentication failed. A GET normally has no body, a missing model would give 404 or a similar error, and NTP problems do not cause authentication rejections in this way.

**59. B** Suggestions are proposals; unfamiliar commands must be verified, even in an emergency, because a wrong change can deepen the outage. Applying widely, applying blindly or reloading all skip verification and add risk.

**60. A** Addressing and DNS work and wired users on the same VLAN reach the Internet, so routing and NAT are fine; the difference is a wireless-specific policy. Laptops have no cable, and the scope already works.

## ANSWER KEY


# Subnetting worksheet answers

| ADDRESS / PREFIX | NETWORK       | MASK            | FIRST HOST    | LAST HOST      | BROADCAST      | HOSTS |
|------------------|---------------|-----------------|---------------|----------------|----------------|-------|
| 192.168.10.77/26 | 192.168.10.64 | 255.255.255.192 | 192.168.10.65 | 192.168.10.126 | 192.168.10.127 | 62    |
| 10.45.130.9/20   | 10.45.128.0   | 255.255.240.0   | 10.45.128.1   | 10.45.143.254  | 10.45.143.255  | 4094  |
| 172.20.5.200/27  | 172.20.5.192  | 255.255.255.224 | 172.20.5.193  | 172.20.5.222   | 172.20.5.223   | 30    |
| 10.0.0.129/25    | 10.0.0.128    | 255.255.255.128 | 10.0.0.129    | 10.0.0.254     | 10.0.0.255     | 126   |
| 192.168.1.5/30   | 192.168.1.4   | 255.255.255.252 | 192.168.1.5   | 192.168.1.6    | 192.168.1.7    | 2     |
| 172.16.99.1/22   | 172.16.96.0   | 255.255.252.0   | 172.16.96.1   | 172.16.99.254  | 172.16.99.255  | 1022  |

# Your certification path

Each ZyeemZia™ track stands on its own, and together they form a path from first network to expert design, plus leadership and AI architecture. Take the next step at the sites below.

## Networking ladder



## Leadership & AI



ZyeemZia™

**ZCNA**

Certified Network Associate

**CERTIFICATE OF ACHIEVEMENT**

This certifies that

---

has successfully completed all modules of the ZyeemZia™  
Certified Network Associate track and passed the ZCNA  
certification exam

IPv4/IPv6 addressing &amp; subnetting

Switching &amp; VLANs

Routing &amp; OSPF

IP services &amp; Wi-Fi

Security &amp; automation basics

**ISSUED****CREDENTIAL ID**

---

*Earn the official certificate online at [zcna.zyeemzia.com](https://zcna.zyeemzia.com); complete every lesson and  
pass the exam. This page is for your printed record.*

## NOTES

## My notes