

16. IPv4 DHCP

DHCP basics

DHCP = Dynamic Host Configuration Protocol: gives a host an IP address **plus** the settings it needs: subnet mask, default gateway, DNS servers, lease time, other options. **Lease** = permission to use an address for a limited time (IOS default with no Lease line = **1 day**). Server **UDP 67**, client **UDP 68**.


Setting	Role
IP address	identifies the host
Subnet mask	which part is the network
Default gateway	way out of the subnet
DNS server	name → IP lookup

Gateway and DNS solve different problems: IP to remote network works but names fail → check DNS. Lease OK but remote network fails → check gateway/routing.

Client's first broadcast stays **inside its VLAN**. Local DHCP server answers directly; a server across a Layer 3 boundary needs a **relay** on the client's gateway.

Four steps (DORA)

Msg	By	Src	Dst	Port
Discover	client	0.0.0.0	→ 255.255.255.255	68 → 67
Offer	server	server IP	→ bcst/client	67 → 68
Request	client	0.0.0.0	→ 255.255.255.255	68 → 67
ACK	server	server IP	→ bcst/client	67 → 68



Offer = proposal only; several servers may offer, client picks one. Request is broadcast so the **other servers know they were not chosen**. Only after ACK does the client configure the address.

Slides draw all four as broadcasts; real servers may unicast Offer/ACK, and renewals are normally unicast (RFC 2131). Initial Request ≠ lease renewal.

DHCP server on a switch/router

Gateway first: VLAN needs an active SVI (interface vlan 10) or router subinterface.

```
ip dhcp excluded-address 10.1.10.1
10.1.10.20 ! range, both ends incl.
ip dhcp excluded-address 10.1.10.254 ! single address
ip dhcp pool VLAN10 ! any name
network 10.1.10.0 255.255.255.0 ! subnet ID + mask
default-router 10.1.10.1 ! VLAN gateway (SVI)
dns-server 8.8.8.8 8.8.4.4
lease 7 ! days; lease 0 12 = 12 h; lease infinite
exit
show ip dhcp binding ! who got which IP
show ip dhcp pool / show ip dhcp conflict
```

Command	Common mistake
excluded-address	with d (slide typo "exclude"); excluded = never handed out, not a reservation
network	host IP instead of subnet ID
default-router	DHCP server's IP instead of VLAN gateway
dns-server	getting an IP ≠ names work
lease	leased address is not permanent

Order: exclusions (global, before pool) → pool → network → default-router → dns → lease → exit → verify.


One pool per VLAN/subnet. Server picks the pool matching the subnet of the interface the request arrived on; for a relayed request it uses the relay's interface IP.

Ex: R1 0/0.10 = 10.0.10.1/24, G0/0.20 = 10.0.20.1/24 → exclude 10.0.10.1 and 10.0.20.1, pool V10 (network 10.0.10.0, default-router 10.0.10.1), pool V20 (10.0.20.0, 10.0.20.1).

DHCP relay (ip helper-address)

Routers do not forward broadcasts. Helper on the **client-facing** L3 interface turns the broadcast into a **unicast** to the server and adds the relay IP (giaddr) so the server knows which pool.

```
interface vlan 10 ! client side, NOT uplink
ip address 10.1.10.1 255.255.255.0
ip helper-address 20.5.5.5
```



Server still needs: pool for **10.1.10.0/24** with default-router **10.1.10.1**, and routes both ways. Helper creates no pool or route.

Only VLAN 20 gets no lease while VLAN 10 works → VLAN 20 SVI missing helper, or no pool for that subnet. Client shows 169.254.x.x (APIPA) = no DHCP reply.

Symptom	Check
No lease	VLAN membership, gateway up, pool subnet + free addresses, helper placement, server reachability
Lease OK, remote fails	delivered gateway, routing
IP OK, names fail	delivered DNS, DNS reachability

IPv6 SLAAC

SLAAC = Stateless Address Autoconfiguration: host builds its own IPv6 address from the /64 prefix in a **Router Advertisement (RA)**; no DHCPv6 lease. **M flag** = get address from DHCPv6; **O flag** = get other info (DNS) from DHCPv6. M=0, O=0 → SLAAC only (DNS via RA option).


```
ipv6 unicast-routing ! global; without it no RA
interface vlan 10
ipv6 address 2001:DB8:10::1/64
ipv6 enable
no ipv6 nd managed-config-flag ! M=0
no ipv6 nd other-config-flag ! O=0
```

LAB 13 uses IPv4 DHCP and IPv6 auto together. Keep the lab's own IPv6 addresses in the lab.

17. Enterprise Campus Network Design

Why hierarchy

Old shared Ethernet (hub): all hosts share bandwidth, one collision domain + one broadcast domain. Switch → each port own collision domain; router/L3 switch → splits broadcast domains. Network grown without a plan → nobody can tell which path traffic takes. Good design: known traffic flows, enough capacity, room to grow, known failure plan.



Layer	Role	Choose for
Access	connect end devices; VLAN membership, filtering, traffic priority	low cost/port, high port density, scalable uplinks, redundant paths up
Distribution	aggregate access; inter-VLAN routing; security + connectivity policy (ACL, QoS); L2/L3 boundary	high L3 throughput, fast ports, ACLs, QoS, resilient links
Core	connect distribution blocks; carry campus traffic fast	very high L3 throughput, redundancy, fast recovery, QoS, simple forwarding

Keep ACLs and heavy packet processing out of the core: core just forwards fast. Policy and routing belong at distribution.

Service	Resource	Path
Local	same segment/VLAN as user	access only
Remote	other VLAN in same block	access → distribution
Enterprise	central, shared by campus	access → distr. → core

Server in another switch block is always reached through the core.

Switch block

Switch block = access switches + the distribution switches they connect to. **Core** joins blocks. Routing normally starts at distribution: VLANs + broadcasts stop there, so **STP runs only inside each block** (failure stays in one block).

Two cables to one upstream switch protect a link, not the switch. Use a pair of distribution switches; each access switch uplinks to **both**. Core also needs redundancy (one core switch = single point of failure).

Three switch-block designs

Design	Result
Shared VLAN (VLAN A on both access switches)	trunks between all 4 switches → loops → STP blocks some links; larger L2 domain
Local VLANs (A only left, B only right)	access uplinks still L2 trunks; link between distribution switches routed (L3) → no L2 loop, nothing to block
Routed access (L3 moves down)	access switches route: all uplinks L3; routing protocol (OSPF) fails over; both uplinks carry traffic (equal-cost)

Sizing a block

- Access: ports for users/devices + uplink capacity.
 - Distribution: all access traffic + inter-VLAN routing + ACL/QoS work.
 - VLAN size: reach, broadcast/multicast load, STP domain size.
 - Growth: new groups/apps and where their servers are.
- ~2,000 users per block = slides' historical guideline, not a protocol limit.

Block too large: distribution overloaded by routing/ACL work, or broadcast/multicast slows the block → add capacity or split the block; adding access ports makes it worse.

Core designs

Design	Structure / reason
Dual core	2 multilayer core switches; distribution has redundant L3 paths into core; no single core failure point; equal-cost paths active
Multinode core	several core pairs joined; each block to one pair; very large campus kept modular
Collapsed core	same device pair does distribution + core; access dual-homed to the pair; smaller campus, less cost. Merges layers, keeps redundancy

Design review

- Mark access / distribution / core.
- Draw VLAN and L3 boundaries; trunk vs routed links.
- Each access switch → two different distribution switches.

- Link the distribution pair together; do not link access switches (unless stacked).
- Test link failure and switch failure; know what recovers (STP, HSRP, routing).
- Remaining link must carry all traffic after a failure.

18. Integrating Wireless LAN

Half-duplex, CSMA/CA

One channel = shared airtime, only one sender at a time (half-duplex). **CSMA/CA** (Collision Avoidance):

- Wait for idle channel + **DIFS**.
- Count down random backoff while idle.
- Busy → pause countdown; resume when free.
- Transmit at 0. Collision still possible (two counters hit 0 together).

Ex A/B/C: A sends first, B and C defer; C has shorter backoff → C next, B waits again. Avoidance reduces risk, does not remove it.

Wi-Fi standards

IEEE	Gen	GHz	Max PHY	MIMO
802.11	orig.	2.4	2 Mbps	no
802.11b		2.4	11 Mbps	no
802.11a		5	54 Mbps	no
802.11g		2.4	54 Mbps	no
802.11n	Wi-Fi 4	2.4/5	600 Mbps	SU-MIMO
802.11ac	Wi-Fi 5	5	6.9 Gbps	MU (Wave 2)
802.11ax	Wi-Fi 6	2.4/5	9.6 Gbps	MU-MIMO
802.11ax	Wi-Fi 6E	2.4/5/6	9.6 Gbps	MU-MIMO
802.11be	Wi-Fi 7	2.4/5/6	~46 Gbps	MU-MIMO

PHY rate = theoretical; real throughput lower (width, streams, clients, signal, shared airtime).

Service sets

Term	Meaning
IBSS	ad hoc, no AP, stations talk directly
BSS	one AP + its associated stations (AP = cell centre)
ESS	several BSSs joined by a distribution system (wired)
SSID	network name; AP maps SSID → VLAN (Staff → VLAN 10)
Association	client membership with AP; needs SSID, supported rate, credentials

Infrastructure mode: clients never talk directly; even two clients on the same AP go through the AP (exception: Direct Link Setup). Several SSIDs on autonomous AP → switch port = trunk.

Cells, channels, coloring

Cell = one AP's coverage. Neighbouring cells overlap a little (roaming), no coverage holes; grid vs honeycomb layout; plan channels between floors too. **Full power ≠ capacity**: bigger cell = more clients share one AP's airtime. Lower power = microcell, very small = picocell (denser). Neighbours on **non-overlapping channels 1, 6, 11** (2.4 GHz).

BSS coloring (Wi-Fi 6): 6-bit colour in PHY header → tell own vs neighbour BSS without reading 48-bit BSSID. Different colour + weak signal → may transmit (spatial reuse); still CSMA/CA. Not security.

Autonomous vs controller-based

	Autonomous	Unified (WLC)
Mgmt	each AP separately	WLC manages all LAPs
Forwarding	AP bridges client traffic onto wired VLAN	LAP tunnels all traffic to WLC; WLC puts it on wired net
VLANs	trunk with every SSID's VLAN to each AP	LAP port = access port; only WLC port is trunk
Scale	hard to change channel/power one AP at a time	central channel/power, clients, roaming

Autonomous roaming in one SSID needs the same VLAN on every AP port (end-to-end VLAN). With WLC the client VLAN

only exists at the WLC.


Split-MAC (CUWN)

AP's work split between two devices. **LAP** (lightweight AP, cannot work alone) = real-time radio: tx/rx frames, beacons, probe responses, 802.11 encryption. **WLC** (separate controller for many LAPs) = authentication, security policy, association + roaming, RF channel + power, QoS. WLC also: dynamic channel/power, raise power over failed AP's area, client load balancing, RF monitoring, wireless IPS.

CAPWAP

CAPWAP = Control And Provisioning of Wireless Access Points; LAP → WLC over IP (can be different subnets); replaced LWAPP.

Tunnel	Carries	Protection
Control (UDP 5246)	config, AP management	authenticated + encrypted
Data (UDP 5247)	client traffic	optional, DTLS



Mutual auth with X.509 certificates. WPA2/WPA3 protects client → AP; CAPWAP DTLS protects AP ↔ WLC (separate).

LAP joins a WLC

- Get IP from DHCP.
- Discover WLCs, learn addresses.
- Join one (fail → try next).
- Match image; download + reboot if needed.
- Secure CAPWAP; download SSIDs, VLANs, channels.

LAP loses WLC → clients dropped, LAP rediscovered (FlexConnect / backup WLC can keep clients up).

Central switching path: client → LAP → CAPWAP → WLC → campus; two clients on same AP: A → LAP → WLC → LAP → B. Troubleshoot: client VLAN/IP → tunnel/WLC → routing beyond WLC.

Roaming

Autonomous: each AP keeps own client list; old AP forwards leftover frames to new AP. WLC: controller keeps list and moves client.

Roam	Change	How
Intra-controller	AP changes, WLC same	WLC updates entry to new LAP; <10 ms
L2 inter-controller	AP + WLC change, same subnet	entry handed to new WLC; IP kept (local-to-local); <20 ms
L3 inter-controller	new WLC on different subnet	old = anchor , new = foreign ; EoIP tunnel; IP kept (local-to-foreign)

Ex: Client-1 192.168.100.199 roams AP-1/WLC-1 (VLAN 100) → AP-2/WLC-2 (subnet 192.168.200.0/24): WLC-1 anchor, WLC-2 foreign, keeps 100.199. Roles belong to WLCs, not APs. Anchor stays anchor on later roams. Guest anchor: guest traffic tunneled on purpose to one WLC (DMZ).

Fast roaming + mobility groups

Joining new AP is fast; full re-login (802.1X + RADIUS + new keys) is slow. Both client + network must support:

Method	Idea
CCKM	Cisco; WLC stores client keys; needs CCX client
PKC / OKC	cached key shared with other APs
802.11r (FT)	IEEE standard; keys set up during the move

"Remembers previous APs, up to 8" on the slide = Sticky Key Caching (SKC).

Mobility group = WLCs sharing client info for roaming. **Mobility list** (MAC, IP, group name) configured manually; no automatic trust; best within the same group.

19. High Availability

HA = network keeps working almost all the time even when a part fails. Campus HA: two of every device, two of every link, switch pairs joined into one logical switch (STP blocks nothing), backup supervisor that takes over fast.

Pairs still waste capacity

Each access switch uplinks to each distribution switch: survives link/distribution failure, but (1) access switches do not back each other up; (2) redundant uplinks form loops → **STP blocks** paid-for bandwidth.

Fix: two physical switches act as **one logical switch** → links to it bundle into one **EtherChannel**, STP sees one link.

Stage	Result
4 separate switches	STP blocks uplinks; failover waits for STP
Logical access (A1+A2), separate D1, D2	one channel to each D → loop through D1-D2 → STP may block a whole channel
Logical access + logical distribution	one MEC ; no loop; all links forward; losing a member only cuts capacity

MEC = multichassis EtherChannel: member links end on different physical switches of a logical switch.

StackWise vs VSS

	StackWise / Plus	VSS
Joins	up to 9 fixed switches, stack cables in a ring	exactly 2 identical chassis over a VSL
Control	stack master; one IP	one active supervisor, other standby
Layer	access (3750-E/X, 3850)	distribution / core

Supervisor redundancy

Supervisor = brain of a chassis switch: management, routing protocols, forwarding tables. First to boot = active, other = standby (power supplies same idea).

Mode	Standby state	Failover
RPR	partially booted; must reload all modules , init route engine, L2, L3, FIB flushed	> 2 min (ports drop)
RPR+	booted, route engine ready; modules not reloaded (ports keep state); L2/L3 protocols start after failure	> 30 s
SSO	fully initialized; startup + running config + L2 state synced; FIB synced	≈ 1 s, links stay up

Slides typo "RPR" = RPR. Log says "booted standby, ports kept, L2/L3 started after failure, ~40 s" → RPR+.

RIB, FIB, NSF

Table	Built by	Used for
RIB	routing protocols + static/connected	choosing best routes (control plane)
FIB	CEF, from RIB (+ adjacency table = next-hop MAC)	hardware forwarding (data plane)

NSF (Nonstop Forwarding) with SSO: standby has FIB copy → packets keep flowing during switchover; new active rebuilds RIB; **NSF-aware** neighbours keep the adjacency and resend routes. Protocols: BGP, EIGRP, OSPF, IS-IS. Slide typos "NFS", "EIRP".

Want links up during supervisor failure → SSO (+ NSF for routed traffic).

20. First Hop Redundancy: FHRP and HSRP

Gateway = single point of failure

Host sends every off-subnet packet to one default gateway.
Ex: VLAN 10 (10.1.1.0/24) hosts use R1 10.1.1.9; R1 or its LAN interface fails → nothing leaves the subnet. Adding R2 10.1.1.129 = second exit, but hosts still point at .9. Changing every host / waiting for DHCP for R2 is not favorable.

FHRP: routers keep own interface IPs and share a **virtual IP + virtual MAC**; hosts use the VIP as gateway; one router answers, another takes over on failure.

Protocol	Origin	Model	Balances
HSRP	Cisco	1 active, others standby/listen	per subnet (group per VLAN)
VRRP	RFC 5798	1 master, others backup	per subnet
GLBP	Cisco	active/active	per host (different vMACs)


Election (who is active)

- No other HSRP router → local router active.
- No active yet → highest **priority** (default **100**).
- Tie → highest IP on the HSRP interface.
- Active already exists → keeps role unless a higher-priority router has **preempt**.

"Highest priority is always active" = false. Preempt is OFF by default; configure it on the router that should reclaim. HSRP = highest wins, STP = lowest wins.

Ex: R1 pri 110 (no P) reboots → R2 (100) active; R1 returns → R2 stays active, R1 standby. Fix: standby 1 preempt on R1.

Failover



Hello **3 s**, hold **10 s**. Standby stops hearing hellos → hold expires → becomes active, uses VIP + vMAC, sends **gratuitous ARP** (unrequested ARP reply) from the vMAC → switches update MAC tables → frames for gateway go to R2's port. **Hosts change nothing** (same gateway IP, same ARP entry).

HSRP v1 vs v2

	v1	v2
IPv6	no	yes
Hello timers	seconds	milliseconds
Groups	0–255	0–4095
vMAC	0000.0C07.ACxx	0000.0C9F.Fxxx
Multicast	224.0.0.2	224.0.0.102
Unique router ID	no	yes

vMAC = fixed prefix + **group number in hex** (*16: quotient, remainder; 10–15 = A–F).

Ex: group 10 = 0x0A → v1 0000.0C07.AC0A, v2 0000.0C9F.F00A. 0000.0C07.AC10 = group 16, not 10. Group 30 = 0x1E. Both routers same version.

Configuring HSRP

```
! R1 (wants active)
interface g0/0
 ip address 192.168.1.2 255.255.255.0
 standby version 2
 standby 10 ip 192.168.1.1 ! group + VIP
 standby 10 priority 110 ! > default 100
 standby 10 preempt ! take role back
 standby 10 name GW10 ! optional
! R2
interface g0/0
 ip address 192.168.1.3 255.255.255.0
 standby version 2
 standby 10 ip 192.168.1.1
 show standby brief / show standby
```

Must match on both: group number, VIP, version. Each keeps own interface IP. Group number does **not** have to equal VLAN number (convention only; v1 max 255).

Reading show standby brief

Interface	Grp	Pri	P	State	Active
Standby	Virtual IP				
G10/0	1	100	Active	local	10.1.1.9
10.1.1.9	10.1.1.1				
Field	Meaning				
State	this router's own role				
Active / Standby	real IP of router in that role; local = this router				
P	blank = no preempt; P = preempt				
Virtual IP	hosts' gateway, same on both				

Long show standby adds vMAC, hello/hold timers, preemption, state history, active router priority.

Both routers 'Active', standby 'unknown' → no hellos: group/version mismatch, link down, multicast filtered. Preempt is never the cause.

Load sharing: DSW1 active VLAN 10 (pri 110 + preempt in vlan 10 only), DSW2 active VLAN 20; one group per VLAN.

21. Introduction to Quality of Service

QoS = methods to protect and prioritize time-critical/important traffic: classify, mark, limit rate, decide what to queue, send first or drop when congested. Fights **delay**, **jitter** (delay variation), **packet loss**. Cannot add bandwidth.

	Data	Voice	Video
Bandwidth	high	low	high
Delay	not critical	<150 ms	<150 ms (real-time)
Jitter	n/a	low (<30 ms)	low
Loss	<5%, TCP resends	<1%	low
Packets	~300 B avg, steady/bursty	≤480 B, steady, UDP	800–1500 B, bursty VBR, UDP

Jitter OK: VoIP <30 ms, video conf 30–50, gaming <30, streaming <100, data >100+. VDI: mixed sizes, TCP (RDP, ICA) or UDP (PCoIP), sensitive.

Service models

Model	Guarantee	How
Best effort	none	all alike, one FIFO
IntServ	hard QoS	per-flow reservation on whole path (RSVP)
DiffServ	soft QoS	classes; per-hop preferential treatment

Campus uses DiffServ: scales (each hop only reads marking, no per-flow state).

Switch QoS pipeline

classify police mark queue sel. schedule

Classify → internal DSCP; police → over-rate dropped/marked down; mark/re-mark; choose egress queue; congestion management (which queue next) + avoidance (WRED drops early). No avoidance → full queue just drops new packets.

Markings

Marking	Layer	Where	Bits	Values
CoS	2	802.1Q tag (tagged frames only)	3	0–7
IP Prec.	3	top 3 bits ToS byte	3	0–7
DSCP	3	top 6 bits same byte	6	0–63

b5 b4 b3 class	b2 b1 drop	0	ECN (2)
DSCP (6 bits) · IPP = b5–b3			

DSCP = IPP × 8 (IPP 5 = 101 → 101000 = 40). Last 2 bits of the byte = ECN (slide says "last 2 bits of DSCP"). IP phone

voice = **DSCP 46 (101110)**, **IPP 5, CoS 5**.

Per-hop behaviors

PHB	DSCP	Treatment
Default	0	best effort
EF	46	low delay/jitter/loss, priority queue (voice)
AF (RFC 2597)	AFxy = 8x + 2y	x class 1–4 (higher = better queue), y drop 1–3 (higher = dropped sooner)
CS	CSn = 8n	IPP-compatible, only top 3 bits set

Class	Drop 1	Drop 2	Drop 3
4 best	AF41 34	AF42 36	AF43 38
3	AF31 26	AF32 28	AF33 30
2	AF21 18	AF22 20	AF23 22
1 worst	AF11 10	AF12 12	AF13 14

Decode: 46 → EF. Else value ÷ 8 = class, remainder ÷ 2 = drop; remainder 0 → Csn. Binary: 3 | 2 | 1 bits; drop bits 01 = 1, 10 = 2, 11 = 3.

Ex: 28 = 011 10 0 → AF32.30 = 011 11 0 → AF33.34 → AF41.22 → AF23.24 → CS3.48 → CS6 (routing protocols).56 → CS7. Class 5–7 never AF.

Trust and classification

Every packet gets an **internal DSCP**. Default port **untrusted** → DSCP 0. **trust dscp** keeps packet's DSCP. Trust CoS/IPP → default map ×8: 0,8,16,24,32,40,48,56.

Ex: host sends CoS 5 → trusted CoS port = 40, untrusted = 0. CoS 5 → 40, not 46 (EF); for voice trust DSCP or change map. CoS 3 → 24.

Classify by: trusted marking, port default value, ACL, **NBAR** (app recognition). Classify as close to the user as possible (access switch). Marking = write new DSCP/CoS/IPP (whole port or policy map; ex video server → DSCP 40).

Shaping vs policing

	Shaping	Policing
Excess	delayed (buffered)	dropped / re-marked now
Bursts	smoothed, fewer drops later	cut at rate (burst size allows some)
Delay?	yes	no
Good for	bursty TCP out to slower ISP link	hard limit, user traffic at access
Bad for	voice/video	traffic that can't lose packets

Leaky token bucket: packet pours tokens = its size; bucket drains at configured rate; size = burst. Overflow → **out of profile** → drop/re-mark. Policser as **dropper** (≤CIR sent, rest dropped) or **marker** (≤CIR AF11, CIR–PIR AF12, >PIR AF13).

Ex: 1 Gbps edge, provider contract 200 Mbps polices above → **shape** to 200 Mbps on your egress.

Queueing / scheduling

Mechanism	Key point
FIFO	default, one queue, tail drop, no QoS
WRR	queues in turn; share = weight ÷ total
Priority queue	served first; WRR ignored for it (voice)
Custom queueing	% per traffic type, unused borrowed; WAN; static, not for fast WAN
LLQ	CBWFQ + one strict-priority queue

Ex: weights 50/10/25/15 → 50/10/25/15 %. Weights 4/2/1/1 → Q1 = 4/8 = 50%, Q2 25%, Q3 = Q4 = 12.5%. Divide by the total, not 100.

Buffers split too: priority queue small (voice leaves at once), low-priority queues bigger.

Congestion avoidance

Tail drop (default): full queue drops every arriving packet → many TCP flows lose packets and slow-start together (**TCP synchronization**), aggressive flows hog the queue, dead queues add delay/jitter, premium dropped like best effort.

RED: random drops **before** full, more as queue fills; senders back off at different times; still tail drop when full; helps

only TCP. **WRED:** thresholds/drop probability per DSCP/IPP.

Ex: 50% full → drop DSCP 0–20 randomly; 80% → also 21–51; full → tail drop. At 60%: only DSCP 18 candidates, not 26.

Best practice: voice in strict priority queue, WRR for the rest, WRED on data queues (esp. congested WAN links). **AutoQoS:** one command builds standard classes/policies.

22. Monitoring Switches: Syslog, NTP, SNMP, Port Mirroring

Syslog messages

*Oct 20 15:31:07.412 ICT: %SYS-5-CONFIG_I: Configured from console by console

Field	Example	Meaning
Timestamp	*Oct 20 15:31:07.412 ICT	date/time or uptime
Facility	%SYS	module that generated it
Severity	5	0–7, lower = more severe
Mnemonic	CONFIG_I	exact event
Text	Configured from...	description

Lvl	Keyword	Typical events
0	emergencies	crash, stopped processes
1	alerts	platform errors
2	critical	hardware, port security, STP
3	errors	ACL/TCAM, PAgP, interface up/down (LINK-3)
4	warnings	DHCP snooping
5	notifications	802.1X, DTP EtherChannel, PoE, line protocol (LINEPROTO-5)
6	informational	stack, DAI, VTP, UDLD, diagnostics
7	debugging	debug output

Every Awesome Cisco Engineer Will Need Ice-cream Daily. Threshold shows **its level and all lower numbers**; keyword = number (logging console 4 = logging console warnings → 0–4).

```
logging console 4 ! default console = 7
logging buffered 6 ! buffer off by default
logging buffered 16384 ! size 4096–2147483647
logging buffered 16384 warnings ! size + level
service timestamps log datetime msec
localtime show-timezone year
service timestamps log uptime ! 00:12:44 style
show logging
```

logging console warnings: LINK-3 shown, LINEPROTO-5 hidden, logging console 3: LINEPROTO-5 not shown.

Clock and NTP

Boot clock = **March 1, 1993**. clock set 15:30:00 20 October 2026 = **privileged EXEC**; clock timezone ICT 7 = **global config** (offset –23...23, min 0–59). show clock.

```
ntp server 200.1.1.10 prefer
ntp server 200.1.1.240 version 4 ! v4 adds IPv6
show ntp status ! synchronized? stratum
show ntp associations
```

Stratum = hops from reference clock; attached = 1; client = peer + 1; **16 = unsynchronized**. Symbols: * sys.peer (syncd to), * candidate, * selected, ~ outlier, x falseticker, ~ configured. **reach 377** = last 8 polls OK.

Ex: ~10.0.0.1 st 2, ~10.0.0.2 st 4 → **SW reports 3. Use the * line, not the lowest stratum.**

SNMP

Manager (NMS: Cacti, Munin) polls and receives alerts; **agent** = process on device; **MIB** = tree of device variables in memory; **OID** = dotted ID, ASN.1. Ex 1.3.6.1.2.1.2.1.10 = interface inbound bytes (IF-MIB). Vendors under 1.3.6.1.4.1: Cisco, 9, Microsoft, 311, Juniper, 2636.

Op	Dir	Port	Purpose
Get	M → A	161	read one variable
Get-next	M → A	161	next variable (walk table)
Get-bulk	M → A	161	whole table (v2c+)
Set	M → A	161	write a variable
Trap	A → M	162	alert, not acknowledged
Inform	A → M	162	alert, manager acks (v2c+)

	v1	v2c	v3
RFC	1157	1901	3410–3415
Auth	community	community	username (in groups)
Counters	32-bit	64-bit	64-bit
Extra	get/set, traps	+ get-bulk, inform	HMAC SHA/MD5, DES/3DES/AES, views

v3 level	Keyword	Auth	Encrypt
noAuthNoPriv	noauth	no	no
authNoPriv	auth	yes	no
authPriv	priv	yes	yes

```
access-list 10 permit 192.168.3.99
snmp-server community PUBLIC ro 10 ! v1/v2c
snmp-server host 192.168.3.99 PUBLIC ! v1
trap
snmp-server host 192.168.3.99 informs
version 2c PUBLIC
snmp-server view v1 iso included ! optional
snmp-server group NetOps v3 priv read v1
snmp-server user mymonitor NetOps v3 auth sha
AUTHPM priv aes 128 PRIVPM access 10
snmp-server host 192.168.3.99 informs
version 3 priv mymonitor
```

Order v3: group (level) → user (passwords) → host. No view = all MIB visible. Use v3; if v1 is a must → read-only + ACL (community is clear text).

Acknowledged + encrypted → informs version 3 priv. **Acknowledged only** → informs v2c. v1 cannot send informs.

Local SPAN


SPAN copies traffic from source ports/VLANs to a destination port with an analyzer; normal forwarding unchanged.

```
monitor session 1 source interface g1/0/1
both ! rx/tx/both
monitor session 2 source vlan 10 rx ! VLANs
own session
monitor session 1 destination interface
g1/0/48
monitor session 2 destination interface
g1/0/47 encapsulation replicate
show monitor session 1 detail
no monitor session 1 ! or all
```

Default copies untagged, no BPDU/CDP/VTP/DTP/PagP; **encapsulation replicate** keeps tags + those frames.

Destination port: STP disabled → never connect it back into the network. **Sources faster than destination** → mirrored frames dropped.

Remote SPAN (RSPAN)



```
vlan 99 ! on A, B and C
remote-span
! A (source)
monitor session 1 source interface g1/0/1
both
monitor session 1 destination remote vlan 99
! C (destination)
monitor session 1 source remote vlan 99
monitor session 1 destination interface
g1/0/48
```

No MAC learning in RSPAN VLAN (copies keep real hosts' MACs) → flooded to analyzer. Middle switch: no session, only VLAN 99, allowed on trunks; never on access ports. VTP can carry it; pruning keeps it off extra trunks.

VLAN 99 missing on B → copies dropped at B. show monitor 'Remote Source Session' = analyzer on another switch.


23. Basic Security

CIA	Meaning	Threat
Availability	data/services reachable reliably, on time	DoS flood
Integrity	only authorized can change data	altered config/records
Confidentiality	only authorized can access data	eavesdropping

HA/HSRP → availability; SNMPv3 auth → integrity; SNMPv3 priv, SSH → confidentiality.

Firewall

Inline between internal and external networks; every packet crosses it. Rules match: source/destination address; transport port; text in HTTP URI; TCP state (**stateful** = remembers sessions, allows replies of inside-started sessions).



Judged by who **initiates**: inside → outside allowed; outside → DMZ allowed; outside → inside denied; replies to inside sessions allowed (stateful). DMZ → inside: denied unless explicitly permitted.

DMZ = less-secured zone for servers outside users may reach (public web for www.example.com). DMZ server compromised → attacker still has to cross the firewall to inside.

IPS / IDS

IPS (Intrusion Prevention System): compares traffic with **attack signatures** (DoS, DDoS, virus, worm); inline; drops matching packets or redirects to another security tool; signature DB must stay **updated**. **IDS**: usually watches a **copy** (SPAN port) and alerts; cannot block inline.

	Firewall	IPS
Decides by	admin rules; address, port, URI, state	attack signatures (+ reputation/context in NGIPS)
Placement	inline at edge	inline at edge
Updating	rules change with policy	signatures updated continuously

New worm on an allowed port → port-based firewall lets it pass; IPS/NGIPS with updated signature stops it.

NGFW / NGIPS

NGFW = all traditional firewall functions + **AVC** (Application Visibility and Control) + advanced malware protection + URL filtering + NGIPS.

NGIPS = IPS + AVC + **contextual awareness** (how info to judge vulnerabilities) + **reputation-based filtering** (known-bad URLs/IPS) + **event impact levels** (rank alerts).

Products: ASA 5500-X with FirePOWER Services; Firepower 1000/2100/4100/9000; NGIPSv (VMware); Firepower Threat Defense for ISR.

Ports and addresses (ch 16–22)

Protocol	Transport / address
DHCP server / client	UDP 67 / UDP 68
CAPWAP control / data	UDP 5246 / UDP 5247
HSRP v1 / v2	UDP 1985, 224.0.0.2 / 224.0.0.102
VRRP	IP protocol 112, 224.0.0.18
SNMP get/set / trap+inform	UDP 161 / UDP 162
Syslog to server	UDP 514
NTP	UDP 123
SSH / Telnet	TCP 22 / TCP 23