SomNOG9

Track 1 · Network Infrastructure

STUDENT LAB NOTEBOOK

19 – 21 September 2026

SNU KM4 Campus, Mogadishu, Somalia

Somali Network Operators' Group · SomNOG

Name: _____

Group number (G): _____ Pod: _____

Contents

How to use this notebook

This notebook is written so you can **read the comment first and type the command second**.

Every configuration block looks like this:

```
! Enter global configuration mode
configure terminal
!
! Set the hostname so you always know which device you are on
hostname SW1-GRP1
```

Lines that begin with an exclamation mark are **comments**. Cisco IOS ignores them completely. That means you can copy an entire block — comments and all — and paste it into the terminal, and the router or switch will execute only the real commands. The comments become part of your configuration file as you type, which is exactly how professional network engineers document their work.

Before every block you will find a line like this:

> **Where to type this:** SW1 · console cable · start from the `SW1-GRP1#` prompt

That line tells you three things: **which device**, **how you are connected to it**, and **which prompt you must be sitting at** before you paste anything. If your prompt does not match, you are on the wrong device or in the wrong mode. Stop and fix that first.

Rules for this workshop

- Type the commands yourself the first time. Paste only when you are repeating something you already understand.
- Never skip a **Verify** section. A configuration you cannot verify is a configuration you cannot trust.
- Some labs are designed to **fail on purpose**. When something does not work, that is the lesson. Discuss it in your group and write down your theory before you turn the page.
- Answer the **Questions** out loud as a group before starting the next lab.

Your group number

Your instructor will give your group a number from 1 to 6. Everywhere in this notebook you see the letter **G**, type your own number instead.

> Group 3 types `10.3.100.0`. Group 6 types `10.6.100.0`. Group 3's switch is called `SW1-GRP3`.

My group number is G = _____

1. Your pod

Device	Model	Its hostname
Switch 1	Cisco Catalyst 3560v2 — 24 × FastEthernet	`SW1-GRPG`
Switch 2	Cisco Catalyst 3560v2 — 24 × FastEthernet	`SW2-GRPG`
Router	Cisco 2800 — only Fa0/0 and Fa0/1	`R-GRPG`
Laptops	5 per group	—

At the front of the room the instructor runs:

Device	Role
`CORE`	Cisco 2800 — the upstream router that performs NAT to the Internet
`CORE-SW`	Cisco 3560v2 — the classroom backbone switch every group plugs into

2. Your peer group

Your router has exactly two interfaces. Fa0/0 goes down to your own switch. Fa0/1 goes **directly to another group's router** with a single cable. The pairs are fixed:

Your group	Your peer group
Group 1	Group 2
Group 2	Group 1
Group 3	Group 4
Group 4	Group 3
Group 5	Group 6
Group 6	Group 5

My peer group is _____. Their router is called _____.

> **Use a crossover cable for the router-to-router link.** Two FastEthernet ports facing each other with no switch between them need the transmit and receive pairs swapped. Many 2800 FastEthernet ports do not auto-detect this. If the link will not come up, this is the first thing to check.

3. Physical cabling map

Cable your pod exactly like this before Lab 1. Every group is identical.

Switch 1 (SW1)

Port	Connected to	Purpose
Fa0/1 – Fa0/6	Student laptops	Access ports, VLAN 100 (USERS)
Fa0/7 – Fa0/12	Student laptops	Access ports, VLAN 200 (SERVERS)
Fa0/13 – Fa0/18	Student laptops	Access ports, VLAN 99 (MGT)
Fa0/19 – Fa0/20	nothing	Spare — shut down
Fa0/21	`CORE-SW`	Routed uplink to the classroom backbone (from Lab 12)
Fa0/22	SW2 Fa0/22	Inter-switch link 1 — STP and EtherChannel
Fa0/23	SW2 Fa0/23	Inter-switch link 2 — stays shut until Lab 8
Fa0/24	Router Fa0/0	Routed port down to your router

Switch 2 (SW2)

Port	Connected to	Purpose
Fa0/1 – Fa0/6	Student laptops	Access ports, VLAN 100 (USERS)
Fa0/7 – Fa0/12	Student laptops	Access ports, VLAN 200 (SERVERS)
Fa0/13 – Fa0/18	Student laptops	Access ports, VLAN 99 (MGT)
Fa0/19 – Fa0/21	nothing	Spare — shut down
Fa0/22	SW1 Fa0/22	Inter-switch link 1
Fa0/23	SW1 Fa0/23	Inter-switch link 2 — stays shut until Lab 8
Fa0/24	nothing	Spare — shut down

Router (R)

Port	Connected to	Purpose
Fa0/0	SW1 Fa0/24	Routed link down into your own LAN
Fa0/1	**Your peer group's router Fa0/1**	Point-to-point WAN link — crossover cable

4. Addressing plan

Your group owns the block **`10.G.0.0/16`**. The second octet is always your group number. That rule never breaks.

Purpose	Subnet	Addresses
Router Loopback0	`10.G.0.1/32`	R = `10.G.0.1`
Router Fa0/0 ↔ SW1 Fa0/24	`10.G.1.0/30`	R = `.1`, SW1 = `.2`
VLAN 99 — MGT	`10.G.99.0/24`	Gateway (SW1) = `.1`, SW2 = `.12`
VLAN 100 — USERS	`10.G.100.0/24`	Gateway (SW1) = `.1`, DHCP from `.21`
VLAN 200 — SERVERS	`10.G.200.0/24`	Gateway (SW1) = `.1`, SW2 = `.2`, DHCP from `.21`
SW1 Fa0/21 ↔ backbone	`10.0.0.0/24`	SW1 = `10.0.0.G`, CORE = `10.0.0.254`

The router-to-router WAN links

A link between exactly two routers uses a **/30**: one network address, two usable host addresses, one broadcast address. Nothing is wasted. **The lower-numbered group owns the /30 and hands the second address to its peer.**

Link	Subnet	Lower group takes	Higher group takes
Group 1 ↔ Group 2	`10.1.255.0/30`	G1 = `10.1.255.1`	G2 = `10.1.255.2`
Group 3 ↔ Group 4	`10.3.255.0/30`	G3 = `10.3.255.1`	G4 = `10.3.255.2`
Group 5 ↔ Group 6	`10.5.255.0/30`	G5 = `10.5.255.1`	G6 = `10.5.255.2`

My WAN subnet is _____. My address is _____. My peer's is _____.

5. Lab index

Lab	Title	Devices
1	Basic housekeeping configuration	SW1, SW2
2	VLAN creation and port assignment	SW1
3	Verifying your VLANs	SW1
4	DHCP for VLAN 100	SW1
5	First connectivity test — the same-VLAN rule	SW1 + laptops
6	Extending VLANs to SW2 — discovering the trunk	SW1, SW2
7	Inter-VLAN routing — discovering the SVI	SW1
8	Spanning Tree Protocol	SW1, SW2
9	EtherChannel	SW1, SW2
10	Router housekeeping and the routed port	R, SW1
11	Static routing to your peer group	R
12	Removing statics and deploying OSPF	R, SW1
13	Reaching the Internet — NAT on the core	CORE (instructor)

5.1 Lab diagram

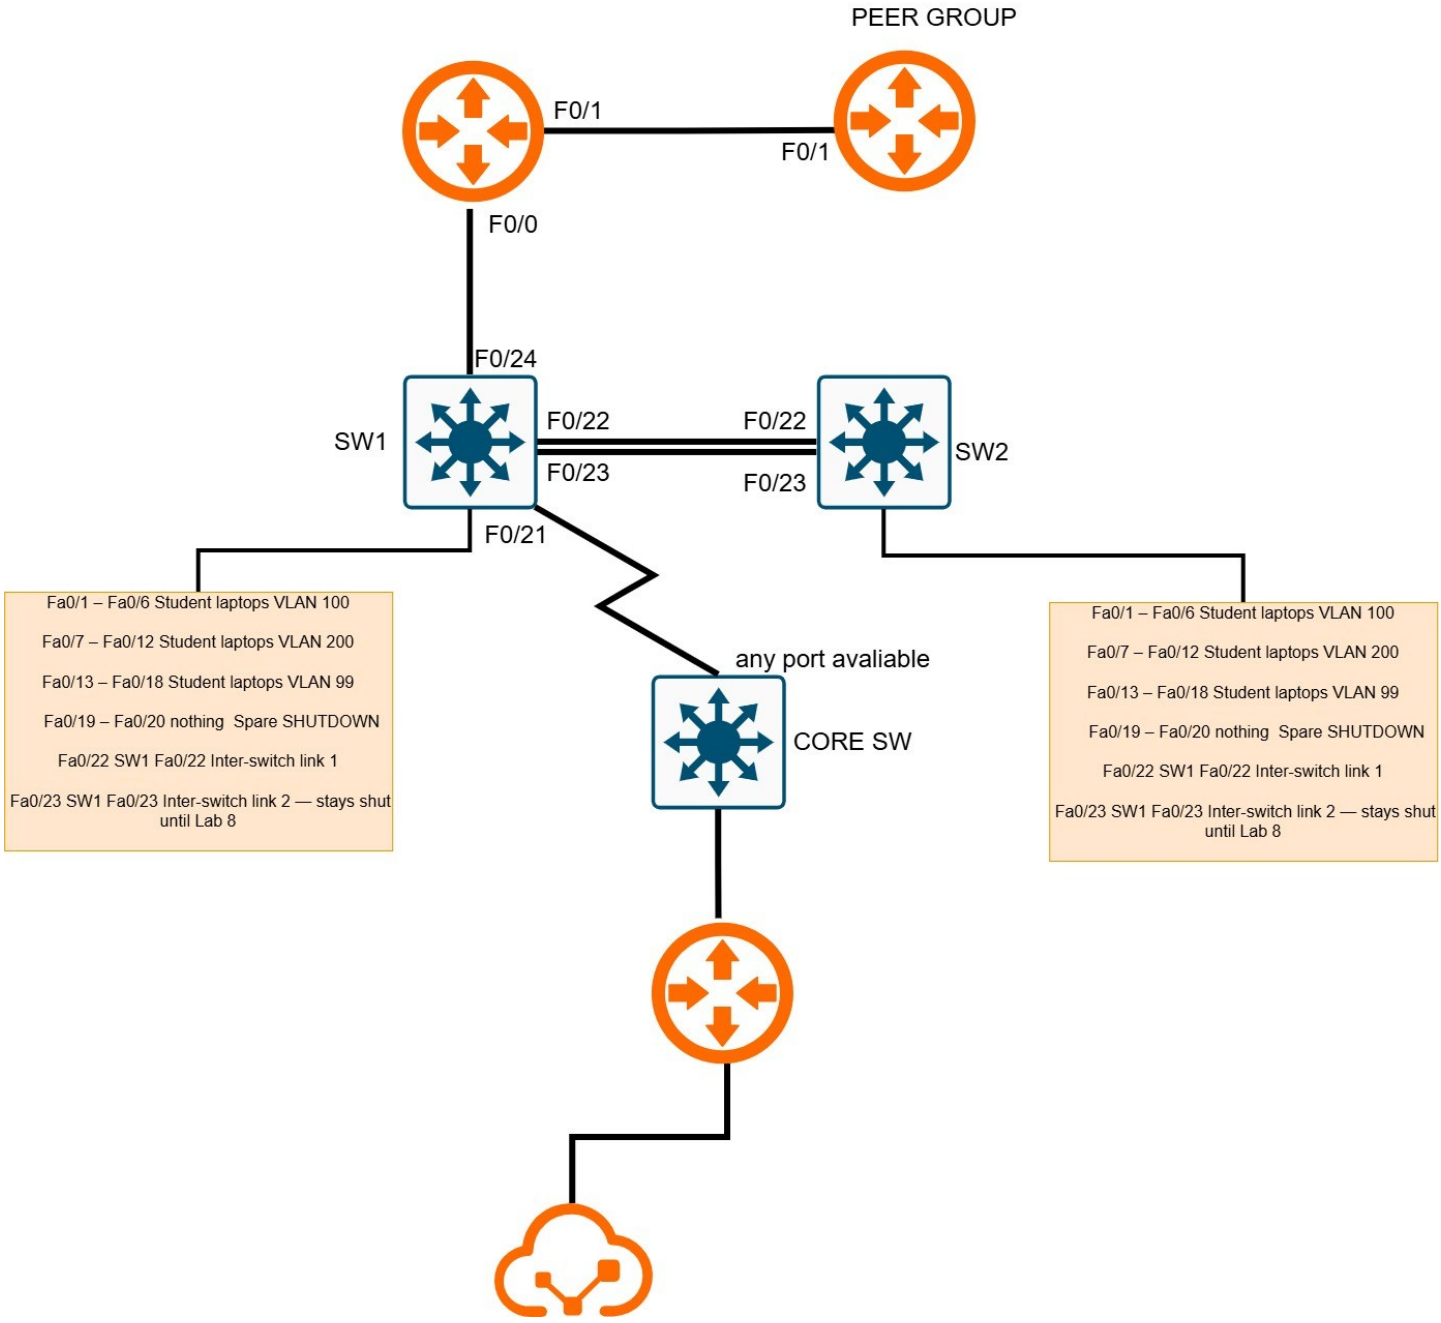

Lab 1 — Basic housekeeping configuration

Devices: SW1 and SW2 **Time:** 30 minutes

Objective

Bring both switches to a known, clean, named state. Every lab after this one assumes your switches carry the correct hostname and can be reached for management.

Task 1.1 — Connect your console

Take the light blue rollover cable. Plug the RJ45 end into the switch port marked **CONSOLE** and the USB end into your laptop.

Open PuTTY (Windows), or use `screen` / `minicom` (Linux and macOS). Serial settings:

Setting	Value
Speed	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

Press Enter a few times. You should see a prompt. If the screen stays blank, you have the wrong COM port or the wrong speed.

Task 1.2 — Erase the old configuration

> **Where to type this:** SW1 · console cable · start from the `Switch>` prompt

```
! Enter privileged EXEC mode
enable
!
! Erase the saved startup configuration
write erase
!
! Delete the VLAN database — VLANs are NOT stored in the startup config
delete flash:vlan.dat
!
! Restart the switch with an empty configuration
reload
```

Answer `no` when it asks whether to save, and press Enter to confirm the reload. When the switch comes back, answer `no` to the initial configuration dialog.

Do the same on SW2 before continuing.

Task 1.3 — Housekeeping on Switch 1

> **Where to type this:** SW1 · console cable · start from the `Switch>` prompt
> Remember to replace **G** with your group number in the hostname and the banner.

```
! Enter privileged EXEC mode on Switch 1
enable
!
! Enter global configuration mode
configure terminal
!
! Configure the hostname so you always know which device you are on
hostname SW1-GRPG
!
! Set the encrypted enable password
enable secret cisco
!
! Stop the switch trying to DNS-resolve your typing mistakes
no ip domain-lookup
!
! Encrypt plain-text passwords stored in the configuration file
service password-encryption
!
! Display a legal banner before login
banner motd #Unauthorised access prohibited - SomNOG9 Group G#
!
! Configure the console line
line console 0
 password cisco
 login
 logging synchronous
 exec-timeout 30 0
 exit
!
! Configure the virtual terminal lines used by Telnet and SSH
line vty 0 15
 password cisco
 login
 exit
!
! Set the timezone to East Africa Time
clock timezone EAT 3
!
! Leave configuration mode
end
!
! Set the software clock so that log messages carry a useful timestamp
clock set 09:00:00 19 September 2026
!
! Save the running configuration to NVRAM
write memory
```

Task 1.4 — Housekeeping on Switch 2

> **Where to type this:** SW2 · move your console cable to the second switch · start from the `Switch>` prompt
> The only difference from Task 1.3 is the hostname.

```
! Enter privileged EXEC mode on Switch 2
enable
!
! Enter global configuration mode
configure terminal
!
! Configure the hostname SW2
hostname SW2-GRPG
!
! Set the encrypted enable password
enable secret cisco
!
! Stop DNS lookups on mistyped commands
no ip domain-lookup
!
! Encrypt plain-text passwords
service password-encryption
```

```
!
! Display a legal banner before login
banner motd #Unauthorised access prohibited - SomNOG9 Group G#
!
! Configure the console line
line console 0
 password cisco
 login
 logging synchronous
 exec-timeout 30 0
 exit
!
! Configure the virtual terminal lines
line vty 0 15
 password cisco
 login
 exit
!
! Set the timezone
clock timezone EAT 3
!
! Leave configuration mode and save
end
write memory
```

Task 1.5 — Shut down the unused ports

An unused port that is administratively down cannot be used by somebody who walks into the room with a laptop.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Select the spare access ports
interface range fa0/19 - 20
 description UNUSED
 shutdown
 exit
!
! Keep the second link to SW2 down until Lab 8 - we want one link only for now
interface fa0/23
 description SECOND LINK TO SW2 - DOWN UNTIL LAB 8
 shutdown
 exit
!
! Leave configuration mode and save
end
write memory
```

> **Where to type this:** SW2 · console cable · start from the `SW2-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Select the spare access ports - SW2 has one more spare than SW1
interface range fa0/19 - 21
 description UNUSED
 shutdown
 exit
!
! Fa0/24 on SW2 is not connected to anything
interface fa0/24
 description UNUSED
 shutdown
 exit
```

```
!
! Keep the second link to SW1 down until Lab 8
interface fa0/23
 description SECOND LINK TO SW1 - DOWN UNTIL LAB 8
 shutdown
 exit
!
! Leave configuration mode and save
end
write memory
```

Verify

Where to type this: SW1 and SW2 · console cable · from the # prompt

```
! Confirm the hostname, banner, passwords and line settings are all present
show running-config
!
! Confirm the IOS version and that the configuration register is 0x2102
show version
!
! Confirm Fa0/19-20 and Fa0/23 are administratively down
show ip interface brief
!
! Confirm the date and time are correct
show clock
```

Questions

1. Why did you have to delete `vlan.dat` separately? Where does `write erase` not reach?
2. What is the difference between `enable password` and `enable secret`? Which one did you configure and why?
3. Your switch has no IP address yet, but you configured `line vty`. Can anyone Telnet to it right now? Explain.
4. What does `logging synchronous` fix? Wait until a port changes state while you are typing and you will see it for yourself.

Lab 2 — VLAN creation and port assignment

Devices: SW1 **Time:** 25 minutes

Objective

Create three VLANs on SW1 and place the access ports into them.

At the end of this lab SW1 will have **no IP address in any of these VLANs.** That is deliberate. Do not create a switched virtual interface yet, even if you already know how.

VLAN ID	Name	Purpose
99	MGT	Management — switch and router administration
100	USERS	Staff laptops
200	SERVERS	Servers and printers

Task 2.1 — Create the VLANs

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Create VLAN 99 and give it a human-readable name
vlan 99
 name MGT
 exit
!
! Create VLAN 100 for the staff laptops
vlan 100
 name USERS
 exit
!
! Create VLAN 200 for the servers
vlan 200
 name SERVERS
 exit
!
! Stay in configuration mode - Task 2.2 continues from here
```

Task 2.2 — Assign the access ports

> **Where to type this:** SW1 · console cable · you should still be at the `SW1-GRPG(config)#` prompt

```
! Six ports for VLAN 100
interface range fa0/1 - 6
 description USERS ACCESS
 switchport mode access
 switchport access vlan 100
 spanning-tree portfast
 no shutdown
 exit
!
! Six ports for VLAN 200
interface range fa0/7 - 12
 description SERVERS ACCESS
```

```
 switchport mode access
 switchport access vlan 200
 spanning-tree portfast
 no shutdown
 exit
!
! Six ports for the management VLAN
interface range fa0/13 - 18
 description MGT ACCESS
 switchport mode access
 switchport access vlan 99
 spanning-tree portfast
 no shutdown
 exit
!
! Leave configuration mode and save
end
write memory
```

What each of those commands did:

Command	Why it is there
`switchport mode access`	Forces the port to be an access port. It will never negotiate itself into a trunk.
`switchport access vlan 100`	Places the port in VLAN 100. Without this it stays in VLAN 1.
`spanning-tree portfast`	Moves the port straight to forwarding when a laptop is plugged in, instead of waiting ~30 seconds.

Verify

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! List every VLAN with its name, status and assigned ports
show vlan brief
!
! Confirm the layer 2 settings of a single port
show interfaces fa0/1 switchport
```

VLANs 99, 100 and 200 must appear as **active** with the correct ports listed underneath each one. If a port is missing from the list, it is still in VLAN 1 — go back and find the typo.

Questions

5. VLAN 1 still exists and you did not delete it. Which ports are still in VLAN 1 on SW1?
6. What does `switchport mode access` protect you from that leaving the default would not?
7. You created VLAN 99 and called it MGT, but the switch still has no management IP address. Is VLAN 99 doing anything useful right now?

Lab 3 — Verifying your VLANs

Devices: SW1 **Time:** 15 minutes

Objective

Learn the six `show` commands that answer almost every VLAN question you will ever be asked.

Task 3.1 — Run the verification commands

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt
> These are all `show` commands. None of them change anything, so run them freely.

```
! Every VLAN, its name, its status, and the ports assigned to it
show vlan brief
!
! The detail of one VLAN only
show vlan id 100
!
! Every port with its status, VLAN, duplex and speed in one screen
show interfaces status
!
! The layer 2 settings of a single port: mode, access VLAN, negotiation state
show interfaces fa0/1 switchport
!
! Which MAC addresses the switch has learned, on which port, in which VLAN
show mac address-table
!
! Which ports are trunks - you should see nothing at this stage
show interfaces trunk
```

Task 3.2 — Record your baseline

Write these down for your group's records:

- Number of ports in VLAN 100: _____
- Number of ports in VLAN 200: _____
- Number of ports still in VLAN 1: _____
- Output of `show interfaces trunk`: _____

Questions

8. `show vlan brief` lists Fa0/22 under VLAN 1. What does that tell you about the link to SW2 right now?
9. The MAC address table is empty or nearly empty. Why?
10. Which single command would you run first if a user complained "my laptop has no network"?

Lab 4 — DHCP for VLAN 100

Devices: SW1 **Time:** 25 minutes

Objective

Turn SW1 into the DHCP server for VLAN 100, so a laptop plugged into Fa0/1 receives an address automatically.

Read this before you configure

A Cisco device can only serve DHCP for a subnet in which it **has an interface**. So SW1 needs an IP address inside VLAN 100. That interface is called a **switched virtual interface**, or SVI.

Creating an SVI does **not** turn on routing. The 3560 has routing disabled by default. In this lab you are giving SW1 one foot in VLAN 100 so it can answer DHCP requests — nothing more. Remember this when Lab 5 fails.

Task 4.1 — Create the SVIs and the DHCP pool

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt
> Replace **G** with your group number in all four addresses below.

```
! Enter global configuration mode
configure terminal
!
! Create the switched virtual interface for VLAN 100
! This gives SW1 an address inside VLAN 100 so it can serve DHCP there
interface vlan 100
 description USERS GATEWAY
 ip address 10.G.100.1 255.255.255.0
 no shutdown
 exit
!
! Create the management SVI - SW1 is host .11 in the management subnet
interface vlan 99
 description MANAGEMENT
 ip address 10.G.99.11 255.255.255.0
 no shutdown
 exit
!
! Reserve the first 20 addresses for infrastructure - never hand these out
ip dhcp excluded-address 10.G.100.1 10.G.100.20
!
! Create the DHCP pool for VLAN 100
ip dhcp pool VLAN100-USERS
 network 10.G.100.0 255.255.255.0
 default-router 10.G.100.1
 dns-server 8.8.8.8
 domain-name somnog.lab
 lease 0 8
 exit
!
! Leave configuration mode and save
end
write memory
```

What each pool command did:

Command	Why it is there
`network`	The range of addresses this pool may hand out.
`default-router`	The gateway address given to every client.

Command	Why it is there
`dns-server`	The DNS resolver given to every client.
`lease 0 8`	Lease time of 0 days and 8 hours.

Verify

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Vlan99 and Vlan100 should be listed as up/up with the right addresses
show ip interface brief
!
! The pool exists and shows 0 leased addresses so far
show ip dhcp pool
!
! Empty for now - it fills up in Lab 5
show ip dhcp binding
!
! The switch can ping its own SVI
ping 10.G.100.1
```

If `Vlan100` shows as **down/down**, no access port in VLAN 100 is up yet. Plug a laptop into Fa0/1 and check again. An SVI stays down until at least one port in that VLAN is up.

Questions

11. Why did you exclude `10.G.100.1` to `10.G.100.20` from the pool?
12. The pool hands out a default gateway of `10.G.100.1`. What device is that, and can it currently forward traffic anywhere?
13. What would happen if a second DHCP server appeared on VLAN 100?

Lab 5 — First connectivity test: the same-VLAN rule

Devices: SW1 and three laptops **Time:** 25 minutes

Objective

Prove with your own laptops what a VLAN actually does. **Half of this lab is designed to fail.** Do not fix anything until the instructor says so.

Task 5.1 — Plug in

Student	Switch port	VLAN	Expected address
Student A	SW1 Fa0/1	100 USERS	`10.G.100.21` or higher
Student B	SW1 Fa0/2	100 USERS	`10.G.100.22` or higher
Student C	SW1 Fa0/7	200 SERVERS	none yet

On every laptop: set the wired adapter to **obtain an IP address automatically**, turn Wi-Fi **off**, and either turn off the firewall or allow inbound ICMP. Windows Defender blocks ping by default and will waste twenty minutes of your life.

Task 5.2 — Check what each laptop received

> **Where to type this:** on the student laptops, in a command prompt or terminal

```
rem Windows - show the address, mask, gateway and DHCP server received
ipconfig /all
```

```
# Linux - show the address and default route received
ip addr show
ip route
```

```
# macOS - show the address and default route received
ifconfig
netstat -rn
```

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! List every address the switch has leased, and to which MAC address
show ip dhcp binding
!
! Confirm the switch has learned the laptop MAC addresses on the right ports
show mac address-table vlan 100
```

Student C received no address. Write down why you think that is before continuing.

Task 5.3 — Test 1: ping inside the same VLAN

> **Where to type this:** on Student A's laptop

```
rem Ping Student B - both laptops are in VLAN 100 on the same switch
ping 10.G.100.22
```

This works. Write down the round-trip time: _____

Task 5.4 — Test 2: ping across VLANs

First give Student C a **static** address so there is something to ping: address `10.G.200.50`, mask `255.255.255.0`, **no gateway**.

> **Where to type this:** on Student A's laptop

```
rem Ping Student C - Student A is in VLAN 100, Student C is in VLAN 200
ping 10.G.200.50
```

This fails. You will see `Destination host unreachable` or `Request timed out`.

Task 5.5 — Investigate before you fix anything

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Look at the routing table - notice how short it is
show ip route
!
! SW1 has an address in VLAN 100 and VLAN 99, but nothing in VLAN 200
show ip interface brief
!
! Search the configuration for the ip routing command - it is not there
show running-config | include ip routing
```

> **Where to type this:** on Student A's laptop

```
rem Student A has no ARP entry for the VLAN 200 address
arp -a
```

Questions — discuss as a group and write your answers down

14. Both laptops are plugged into the **same physical switch**. Why can A reach B but not C?
15. Student A's laptop was given a default gateway of `10.G.100.1`. Did the ping to `10.G.200.50` even reach that gateway? How could you prove it?
16. Student C never received a DHCP address. Which two things are missing on SW1 for VLAN 200?
17. In one sentence: what is a VLAN?

> **Do not configure anything yet.** The fix is Lab 7. First, in Lab 6, you are going to make the problem bigger.

Lab 6 — Extending VLANs to SW2: discovering the trunk

Devices: SW1 and SW2 **Time:** 35 minutes

Objective

Build the same VLANs on SW2, put the second switch's DHCP service in place, and discover for yourselves why a single access link between two switches is not enough.

Task 6.1 — Create the same VLANs on SW2

> **Where to type this:** SW2 · console cable · start from the `SW2-GRPG#` prompt
> The VLAN **numbers and names must match SW1 exactly.** A VLAN is only the same VLAN if both switches call it by the same number.

```
! Enter global configuration mode
configure terminal
!
! The same three VLANs as SW1
vlan 99
 name MGT
 exit
vlan 100
 name USERS
 exit
vlan 200
 name SERVERS
 exit
!
! Stay in configuration mode - Task 6.2 continues from here
```

Task 6.2 — Assign the access ports on SW2

> **Where to type this:** SW2 · console cable · you should still be at the `SW2-GRPG(config)#` prompt

```
! Six ports for VLAN 100
interface range fa0/1 - 6
 description USERS ACCESS
 switchport mode access
 switchport access vlan 100
 spanning-tree portfast
 no shutdown
 exit
!
! Six ports for VLAN 200
interface range fa0/7 - 12
 description SERVERS ACCESS
 switchport mode access
 switchport access vlan 200
 spanning-tree portfast
 no shutdown
 exit
!
! Six ports for the management VLAN
interface range fa0/13 - 18
 description MGT ACCESS
 switchport mode access
 switchport access vlan 99
 spanning-tree portfast
 no shutdown
 exit
```

Task 6.3 — SW2 becomes the DHCP server for VLAN 200

VLAN 100 is served by SW1. VLAN 200 will be served by SW2. Each switch carries one pool, so both switches matter.

> **Where to type this:** SW2 · console cable · you should still be at the `SW2-GRPG(config)#` prompt
> Note that SW2 takes `.2` in VLAN 200. The address `.1` is reserved for a gateway that does not exist yet.

```
! SW2 needs an address inside VLAN 200 so it can serve DHCP there
! This is a server interface, not a gateway - so it takes .2
interface vlan 200
 description DHCP SERVER INTERFACE
 ip address 10.G.200.2 255.255.255.0
 no shutdown
 exit
!
! The management SVI - SW2 is host .12
interface vlan 99
 description MANAGEMENT
 ip address 10.G.99.12 255.255.255.0
 no shutdown
 exit
!
! SW2 is not routing, so it needs a default gateway for its own traffic
ip default-gateway 10.G.99.1
!
! Reserve the first 20 addresses of VLAN 200
ip dhcp excluded-address 10.G.200.1 10.G.200.20
!
! The DHCP pool for VLAN 200
! Note the default-router address - it does not answer yet. That is deliberate.
ip dhcp pool VLAN200-SERVERS
 network 10.G.200.0 255.255.255.0
 default-router 10.G.200.1
 dns-server 8.8.8.8
 lease 0 8
 exit
!
! Leave configuration mode and save
end
write memory
```

Task 6.4 — Connect the two switches

Cable **SW1 Fa0/22 to SW2 Fa0/22**. Then configure both ends as an access port in VLAN 100.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Configure the link to SW2 as a plain access port in VLAN 100
interface fa0/22
 description LINK TO SW2 Fa0/22
 switchport mode access
 switchport access vlan 100
 no shutdown
 exit
!
! Leave configuration mode
end
```

> **Where to type this:** SW2 · console cable · start from the `SW2-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
```

```
!
! The same configuration on the other end of the cable
interface fa0/22
 description LINK TO SW1 Fa0/22
 switchport mode access
 switchport access vlan 100
 no shutdown
 exit
!
! Leave configuration mode
end
```

Task 6.5 — Two more students plug in, this time on SW2

Student	Switch port	VLAN	Test to run
Student D	SW2 Fa0/1	100 USERS	Ping Student A on SW1
Student E	SW2 Fa0/7	200 SERVERS	Ping Student C on SW1

Where to type this: on Student D's laptop

```
rem Did Student D receive an address from SW1's pool on the other switch?
ipconfig /all
!
rem Ping Student A - VLAN 100 on SW2 to VLAN 100 on SW1
ping 10.G.100.21
```

Where to type this: on Student E's laptop

```
rem Did Student E receive an address from SW2's own pool?
ipconfig /all
!
rem Ping Student C - VLAN 200 on SW2 to VLAN 200 on SW1
ping 10.G.200.50
```

Test 1 works. Test 2 fails. Both students are in a VLAN that exists on both switches, and both are using the same physical cable between the switches. Only one of them can pass.

Task 6.6 — Investigate

Where to type this: SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Note the Operational Mode: static access, and Access Mode VLAN: 100
show interfaces fa0/22 switchport
!
! Every MAC learned over this link is in VLAN 100 and nothing else
show mac address-table interface fa0/22
!
! Still empty
show interfaces trunk
```

Questions — answer these before the fix

18. The cable between SW1 and SW2 is a single physical wire. How does it know that a frame belongs to VLAN 100 and not VLAN 200?
19. Student D received a DHCP address from SW1, which is on the other switch. Describe the path that the DHCP DISCOVER took.
20. Student E is in VLAN 200 on SW2 and got an address from SW2's own pool, so DHCP worked. Why can E not reach VLAN 200 on SW1?
21. You have three VLANs and one cable between the switches. How many cables would you need if an access port were the only tool available?

Task 6.7 — The fix: make it a trunk

An **access** port carries exactly one VLAN and sends frames untagged. A **trunk** port carries many VLANs and adds an 802.1Q tag to every frame so the far end knows which VLAN it belongs to.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Create an unused VLAN to act as the native VLAN
! Untagged traffic will land here instead of in VLAN 1
vlan 999
 name NATIVE-UNUSED
 exit
!
! Convert the link to SW2 into a trunk
interface fa0/22
 description TRUNK TO SW2
!
! Use the 802.1Q standard for tagging - required on the 3560
 switchport trunk encapsulation dot1q
!
! Force the port to be a trunk instead of negotiating
 switchport mode trunk
!
! Put untagged traffic in the unused VLAN, not VLAN 1
 switchport trunk native vlan 999
!
! Allow only the VLANs we actually use across this link
 switchport trunk allowed vlan 99,100,200
!
! Do not send DTP frames - we have already decided this is a trunk
 switchport nonegotiate
 exit
!
! Leave configuration mode and save
end
write memory
```

> **Where to type this:** SW2 · console cable · start from the `SW2-GRPG#` prompt
> The configuration must be identical on both ends of the cable.

```
! Enter global configuration mode
configure terminal
!
! The same native VLAN must exist on this switch too
vlan 999
 name NATIVE-UNUSED
 exit
!
! The other end of the trunk
interface fa0/22
```

```
 description TRUNK TO SW1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 999
 switchport trunk allowed vlan 99,100,200
 switchport nonegotiate
 exit
!
! Leave configuration mode and save
end
write memory
```

Verify

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Fa0/22 should be listed - mode on, encapsulation 802.1q,
! native VLAN 999, allowed VLANs 99,100,200
show interfaces trunk
!
! Operational Mode should now read: trunk
show interfaces fa0/22 switchport
```

Now repeat the two ping tests from Task 6.5. **Both succeed.**

Questions — after the fix

22. What exactly does 802.1Q add to an Ethernet frame, and how many bytes is it?
23. Why did we move the native VLAN away from VLAN 1?
24. What would break if SW1 used native VLAN 999 and SW2 used native VLAN 1? What message would CDP print?
25. Test 1 already worked before the trunk. Does making it a trunk change anything for VLAN 100? Explain in terms of tagging.

Lab 7 — Inter-VLAN routing: discovering the SVI

Devices: SW1 **Time:** 30 minutes

Objective

Every VLAN now works perfectly — inside itself. VLAN 100 still cannot speak to VLAN 200. In this lab you turn SW1 into the Layer 3 device for the whole pod.

Task 7.1 — Confirm the problem is still there

> **Where to type this:** on Student A's laptop

```
rem Still fails, even though the trunk is now up
ping 10.G.200.50
```

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Only directly connected networks appear, and VLAN 200 is not one of them
show ip route
!
! SW1 has no interface in VLAN 200
show ip interface brief
```

The trunk solved a **Layer 2** problem — frames now travel between switches in the correct VLAN. Routing between VLANs is a **Layer 3** problem, and nothing you have configured so far does Layer 3 forwarding.

Task 7.2 — Give SW1 a foot in VLAN 200

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Create the SVI for VLAN 200 and take the .1 address
! This is the address SW2's DHCP pool has been advertising all along
interface vlan 200
 description SERVERS GATEWAY
 ip address 10.G.200.1 255.255.255.0
 no shutdown
 exit
!
! Change the management SVI to .1 so SW1 is the gateway there too
interface vlan 99
 ip address 10.G.99.1 255.255.255.0
 exit
!
! Leave configuration mode
end
```

> SW2 keeps `10.G.99.12` and uses `ip default-gateway 10.G.99.1`, which now exists.

Task 7.3 — Test again. It still fails.

> **Where to type this:** on Student A's laptop

```
rem Still fails
ping 10.G.200.50
```

Where to type this: SW1 · console cable · from the `SW1-GRPG#` prompt

```
! The switch ITSELF can reach VLAN 200
ping 10.G.200.50
!
! The switch ITSELF can reach VLAN 100
ping 10.G.100.21
```

SW1 can reach both VLANs, but it will not carry a packet from one to the other. Something is still switched off.

Task 7.4 — The fix: enable routing

Where to type this: SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Enable IPv4 routing on the Layer 3 switch
! This is DISABLED by default on a 3560 - the switch has addresses but
! will not forward a packet from one VLAN to another until you type this
ip routing
!
! Leave configuration mode and save
end
write memory
```

Verify

Where to type this: SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Three connected routes now appear: 10.G.99.0, 10.G.100.0, 10.G.200.0
show ip route
```

Where to type this: on Student A's laptop

```
rem Now succeeds
ping 10.G.200.50
!
rem One hop through the VLAN 100 gateway
tracert 10.G.200.50
```

Now release and renew DHCP on Student C's laptop so it stops using a static address:

```
rem Windows
ipconfig /release
ipconfig /renew
```

```
# Linux
sudo dhclient -r eth0
```

```
sudo dhclient eth0
```

Questions

26. You created the SVI in Task 7.2 and the ping still failed. What did `ip routing` actually change inside the switch?
27. SW2 also has SVIs. Why did we **not** enable `ip routing` on SW2?
28. Before the DHCP renew, Student C's laptop had no default gateway. Would the inter-VLAN ping have worked from C to A? Why?
29. Draw your pod. Mark every device that now has an IP address, and write every address next to it. This drawing is your reference for the rest of the workshop.

Lab 8 — Spanning Tree Protocol

Devices: SW1 and SW2 **Time:** 40 minutes

Objective

Add a second cable between your switches, watch Spanning Tree save you from a broadcast storm, then take deliberate control of which switch is the root bridge.

Task 8.1 — Look at STP before you change anything

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Which STP mode is running, and how many VLANs are protected
show spanning-tree summary
!
! The root bridge, this switch's bridge ID, and the role of every port
show spanning-tree vlan 100
```

Run the same two commands on SW2 and compare. Record:

- Root bridge for VLAN 100: _____
- Its priority: _____
- Its MAC address: _____

By default every switch uses priority 32768, so the switch with the **lowest MAC address** wins. That is an accident of manufacturing, not a design decision.

Task 8.2 — Bring up the second link

Cable **SW1 Fa0/23 to SW2 Fa0/23**.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Configure the second link exactly like the first, then enable it
interface fa0/23
 description SECOND TRUNK TO SW2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 999
 switchport trunk allowed vlan 99,100,200
 switchport nonegotiate
 no shutdown
 exit
!
! Leave configuration mode
end
```

> **Where to type this:** SW2 · console cable · start from the `SW2-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! The other end of the second link
interface fa0/23
 description SECOND TRUNK TO SW1
```

```
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 999
 switchport trunk allowed vlan 99,100,200
 switchport nonegotiate
 no shutdown
 exit
!
! Leave configuration mode
end
```

Task 8.3 — Find the blocked port

Where to type this: SW1 and then SW2 · console cable · from the # prompt

```
! Look at the role and state of Fa0/22 and Fa0/23
show spanning-tree vlan 100
!
! List every port that STP has put into blocking state
show spanning-tree blockedports
```

One port on the **non-root** switch will be in `BLK` — blocking, or alternate. STP has found the loop and broken it. The link is not wasted; it is standby.

Task 8.4 — Take control of the root bridge

Do not let a MAC address decide your topology. SW1 is your Layer 3 switch, so SW1 should be the root.

Where to type this: SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Rapid PVST+ converges in seconds instead of about 50 seconds
spanning-tree mode rapid-pvst
!
! Lower the priority so SW1 is elected root for all three VLANs
spanning-tree vlan 99,100,200 priority 4096
!
! Apply PortFast automatically to every access port
spanning-tree portfast default
!
! Shut down any PortFast port that receives a BPDU - that would be a rogue switch
spanning-tree portfast bpduguard default
!
! Leave configuration mode and save
end
write memory
```

Where to type this: SW2 · console cable · start from the `SW2-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Both switches must run the same STP mode
spanning-tree mode rapid-pvst
!
! A higher number than SW1 makes SW2 the deliberate backup root
spanning-tree vlan 99,100,200 priority 8192
!
! The same edge-port protection
```

```
spanning-tree portfast default
spanning-tree portfast bpduguard default
!
! Leave configuration mode and save
end
write memory
```

Task 8.5 — Test the failover

Start a continuous ping from Student A to Student D and leave it running.

```
rem Windows - continuous ping, stop it later with Ctrl+C
ping -t 10.G.100.30
```

Where to type this: SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Simulate a cable failure on the forwarding link
interface fa0/22
 shutdown
!
! Wait, watch the pings, then restore the link
 no shutdown
 exit
end
```

While it is down, run `show spanning-tree vlan 100` on SW2 and watch Fa0/23 move from blocking to forwarding.
How many pings were lost? _____

Verify

Where to type this: SW1 · console cable · from the `SW1-GRPG#` prompt

```
! SW1 should say: This bridge is the root
show spanning-tree vlan 100
!
! Mode should read rapid-pvst, with PortFast and BPDU Guard enabled by default
show spanning-tree summary
!
! Fa0/1 should be shown as an edge port
show spanning-tree interface fa0/1 detail
```

Questions

30. What are the three parts of a bridge ID, and which one did you change?
31. Why is `4096` a valid priority but `5000` is not?
32. What is the danger of enabling PortFast on a port that connects to another switch, and which command protects you from it?
33. You now have two 100 Mbps links between your switches but only one is forwarding. How much bandwidth are you actually getting? Lab 9 fixes this.

Lab 9 — EtherChannel

Devices: SW1 and SW2 **Time:** 30 minutes

Objective

Bundle Fa0/22 and Fa0/23 into a single logical link, so Spanning Tree sees one port and both cables carry traffic.

Task 9.1 — Rebuild the two ports as a channel group

Both interfaces must be identical and clean before they can be bundled. Shutting them first prevents a temporary loop while you work.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Select both inter-switch links together
interface range fa0/22 - 23
!
! Shut them before rebuilding - this prevents a temporary loop
 shutdown
!
! Strip the trunk settings off the physical ports
! From now on the trunk lives on the logical interface, not on these
 no switchport trunk allowed vlan
 no switchport trunk native vlan
 no switchport mode trunk
!
! Use LACP, the IEEE 802.3ad standard, rather than Cisco's PAgP
 channel-protocol lacp
!
! Put both ports into port-channel 1 and actively ask the far end to bundle
 channel-group 1 mode active
!
! Bring both physical ports back up
 no shutdown
 exit
!
! Stay in configuration mode - Task 9.2 continues from here
```

> **Where to type this:** SW2 · console cable · start from the `SW2-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Exactly the same on the far end - both sides active is the safest combination
interface range fa0/22 - 23
 shutdown
 no switchport trunk allowed vlan
 no switchport trunk native vlan
 no switchport mode trunk
 channel-protocol lacp
 channel-group 1 mode active
 no shutdown
 exit
```

Creating the channel group automatically creates the interface `Port-channel1`.

Task 9.2 — Configure the trunk on the logical interface

From now on you configure `Port-channel1`, never the physical ports.

> **Where to type this:** SW1 · console cable · you should still be at the `SW1-GRPG(config)#` prompt

```
! The bundle itself is the trunk now
interface port-channel 1
 description ETHERCHANNEL TO SW2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 999
 switchport trunk allowed vlan 99,100,200
 exit
!
! Leave configuration mode and save
end
write memory
```

> **Where to type this:** SW2 · console cable · you should still be at the `SW2-GRPG(config)#` prompt

```
! The same configuration on the far end of the bundle
interface port-channel 1
 description ETHERCHANNEL TO SW1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 999
 switchport trunk allowed vlan 99,100,200
 exit
!
! Leave configuration mode and save
end
write memory
```

Verify

> **Where to type this:** SW1 · console cable · from the `SW1-GRPG#` prompt

```
! Look for Po1(SU) with both Fa0/22(P) and Fa0/23(P)
show etherchannel summary
!
! Confirm the protocol is LACP and both ports are members
show etherchannel 1 port-channel
!
! The far switch should be visible as an LACP partner
show lacp neighbor
!
! Only Po1 is listed now, not Fa0/22 and Fa0/23
show interfaces trunk
!
! Only one port - Po1 - and it is FORWARDING. No blocked ports
show spanning-tree vlan 100
```

Reading the flags in `show etherchannel summary`:

Flag	Meaning
`(P)`	Bundled in the port-channel — this is what you want
`(I)`	Stand-alone — the far end is not configured, or the settings do not match

Flag	Meaning
`(s)`	Suspended — a mismatch in speed, duplex, VLAN or trunk settings
`(D)`	Down

Task 9.3 — Failure test

Start a continuous ping between Student A and Student D again.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Pull one cable out of the bundle
interface fa0/23
 shutdown
!
! Now check the summary in another window, then restore it
 no shutdown
 exit
end
```

How many pings were lost this time, compared with the STP test in Lab 8? _____

Questions

34. Spanning Tree no longer blocks a port. Is the loop gone, or is STP simply not seeing it? Explain.
35. What is the difference between `mode active`, `mode passive`, `mode on` and `mode desirable`? Which combinations actually form a channel?
36. Your bundle is 200 Mbps. Will a single file transfer between two laptops run at 200 Mbps? Explain how EtherChannel load balancing really works.
37. Why did you shut the ports down before configuring the channel group?

Lab 10 — Router housekeeping and the routed port

Devices: Router and SW1 **Time:** 35 minutes

Objective

Bring your router into service and connect it to SW1 with a **routed port** — a switch port that behaves like a router interface. This link is the road that OSPF will later travel into your LAN.

Task 10.1 — Router housekeeping

Move your console cable to the router. The serial settings are the same as the switches.

> **Where to type this:** R · console cable · start from the `Router>` prompt
> Replace **G** with your group number in the hostname and the banner.

```
! Enter privileged EXEC mode on the router
enable
!
! Enter global configuration mode
configure terminal
!
! Configure the hostname
hostname R-GRPG
!
! Set the encrypted enable password
enable secret cisco
!
! Stop DNS lookups on mistyped commands
no ip domain-lookup
!
! Encrypt plain-text passwords
service password-encryption
!
! Legal banner
banner motd #SomNOG9 Group G Router - authorised access only#
!
! Configure the console line
line console 0
 password cisco
 login
 logging synchronous
 exec-timeout 30 0
 exit
!
! Configure the virtual terminal lines - a router has 5, not 16
line vty 0 4
 password cisco
 login
 exit
!
! Set the timezone
clock timezone EAT 3
!
! Leave configuration mode and save
end
write memory
```

Task 10.2 — Create the loopback and configure Fa0/0

A loopback interface never goes down. It is the router's permanent identity for OSPF, for management and for logging.

> **Where to type this:** R · console cable · start from the `R-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! A /32 host address - one address, nothing wasted
interface loopback 0
 description ROUTER ID AND MANAGEMENT
 ip address 10.G.0.1 255.255.255.255
 exit
!
! The link down to your own switch
! 255.255.255.252 is a /30 - the classic mask for a link between two devices
interface fa0/0
 description LINK TO SW1 Fa0/24
 ip address 10.G.1.1 255.255.255.252
!
! Router interfaces are DOWN by default - unlike switch ports
 no shutdown
 exit
!
! Leave configuration mode and save
end
write memory
```

Task 10.3 — Turn SW1 Fa0/24 into a routed port

This is the one place in the workshop where you take a switch port **out** of switching.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Convert the port facing the router from Layer 2 to Layer 3
interface fa0/24
 description ROUTED LINK TO R-GRPG Fa0/0
!
! This single command removes all switching behaviour from the port
! It becomes a router interface that happens to live inside a switch
 no switchport
!
! The second usable address of the /30
 ip address 10.G.1.2 255.255.255.252
 no shutdown
 exit
!
! Leave configuration mode and save
end
write memory
```

> `no switchport` only works because you enabled `ip routing` in Lab 7. On a switch with routing disabled, a routed port has nowhere to send packets.

Task 10.4 — Temporary reachability

Until OSPF runs in Lab 12, two static routes let the router and the switch see each other's networks.

> **Where to type this:** R · console cable · start from the `R-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Send everything inside your own /16 down to SW1
ip route 10.G.0.0 255.255.0.0 10.G.1.2
!
```

```
! Leave configuration mode and save
end
write memory
```

Where to type this: SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! SW1 sends everything it does not know about up to the router
ip route 0.0.0.0 0.0.0.0 10.G.1.1
!
! Leave configuration mode and save
end
write memory
```

Verify

Where to type this: R · console cable · from the `R-GRPG#` prompt

```
! Fa0/0 should be up/up with 10.G.1.1, and Loopback0 up/up
show ip interface brief
!
! The router can reach the routed port on SW1
ping 10.G.1.2
!
! The router can reach your VLAN 100 gateway, sourced from its loopback
ping 10.G.100.1 source loopback 0
!
! SW1 should appear as a neighbour on Fa0/0
show cdp neighbors
```

Where to type this: on Student A's laptop

```
rem A laptop deep in VLAN 100 can reach the router loopback
ping 10.G.0.1
!
rem Two hops: the VLAN 100 gateway, then the router
tracert 10.G.0.1
```

Questions

38. Why is a /30 the right mask for a link between exactly two routers? How many addresses does it waste?
39. What is the difference between an SVI and a routed port? When would you choose each?
40. Your laptop's ping to `10.G.0.1` succeeded. List every device the packet touched, in order, and say what each one did to it.
41. Why do router interfaces need `no shutdown` when switch ports were already up?

Lab 11 — Static routing to your peer group

Devices: Router **Time:** 40 minutes

Objective

Connect your network directly to another group's network with a single cable, and reach it using static routes. This is the first time your traffic leaves your own pod.

Task 11.1 — Agree the addresses with your peer group

Walk over to your peer group and agree the numbers face to face. **The lower-numbered group owns the /30 and hands the second address to the other group.** This conversation is exactly what two ISPs do before they interconnect.

My group (G)	_____
My peer group (H)	_____
Our shared /30	_____
My address on the link	_____
My peer's address on the link	_____

Task 11.2 — Cable the link

Run **one crossover cable** from your router's **Fa0/1** to your peer group's router **Fa0/1**. There is no switch in between. Two FastEthernet ports facing each other need the transmit and receive pairs swapped, and many 2800 ports will not do that for you.

Task 11.3 — Configure the WAN interface

The example below is for **Group 1**, which owns `10.1.255.0/30` and takes `.1`. Substitute your own values.

> **Where to type this:** R · console cable · start from the `R-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! The point-to-point link to your peer group's router
interface fa0/1
 description P2P LINK TO GROUP 2 ROUTER
!
! The first usable address of the /30 - the lower group takes .1
 ip address 10.1.255.1 255.255.255.252
!
! Disable CDP on a WAN link - you do not advertise yourself to another network
 no cdp enable
 no shutdown
 exit
!
! Leave configuration mode and save
end
write memory
```

Your peer group configures the mirror image on their own router:

Where to type this: the peer group's router · their console cable · from their `#` prompt

```
! Enter global configuration mode
configure terminal
!
! The other end of the same cable
interface fa0/1
 description P2P LINK TO GROUP 1 ROUTER
!
! The second usable address, given to you by Group 1
 ip address 10.1.255.2 255.255.255.252
 no cdp enable
 no shutdown
 exit
!
! Leave configuration mode and save
end
write memory
```

Task 11.4 — Test Layer 3 before you route

Where to type this: R · console cable · from the `R-GRPG#` prompt

```
! Ping your peer's side of the /30
! This MUST work before anything else - no routing protocol saves a broken link
ping 10.1.255.2
!
! The subinterface should be up/up
show ip interface brief
!
! The /30 should appear as a connected network
show ip route connected
```

If this ping fails, stop here. In order of likelihood: you are using a straight-through cable instead of a crossover; one side has the wrong address or wrong mask; one side is still shut down.

Task 11.5 — Add the static routes

Your peer owns the whole of `10.H.0.0/16`. One static route covers all of it.

Where to type this: R · console cable · start from the `R-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Reach Group 2's entire network via Group 2's address on the /30
! Destination network .. mask .. next hop
ip route 10.2.0.0 255.255.0.0 10.1.255.2
!
! Leave configuration mode and save
end
write memory
```

The peer group adds the mirror route:

```
! Enter global configuration mode
configure terminal
!
! Reach Group 1's entire network via Group 1's address on the /30
ip route 10.1.0.0 255.255.0.0 10.1.255.1
!
```

```
! Leave configuration mode and save
end
write memory
```

Verify

Where to type this: R · console cable · from the `R-GRPG#` prompt

```
! Your static route should be listed with an S
show ip route static
!
! The router explains which route it would use for a given destination, and why
show ip route 10.2.100.0
!
! You can reach your peer's VLAN 100 gateway
ping 10.2.100.1 source loopback 0
!
! Shows the path your packets take
traceroute 10.2.100.1
```

Where to type this: on Student A's laptop

```
rem A laptop in your VLAN 100 reaches a laptop in your peer's VLAN 100
ping 10.2.100.21
```

Questions

42. Your peer's laptop can now ping yours. How many routing decisions did that packet pass through in each direction?
43. What is the administrative distance of a static route? What is it for OSPF? Which one wins, and why?
44. You wrote one static route for your peer's entire `/16`. How many routes would you need if you had listed each VLAN separately? Now multiply that by six groups.
45. Your peer reboots their router. How long does it take your router to notice the route is dead? What does that tell you about static routing?

Lab 12 — Removing statics and deploying OSPF

Devices: Router and SW1 **Time:** 50 minutes

Objective

Delete every static route you wrote, connect your pod to the classroom backbone, and let OSPF discover the whole room automatically.

Understand the topology first

Your router's two interfaces are both in use: Fa0/0 to your switch, Fa0/1 to your peer group. There is no port left for the classroom backbone.

So the backbone uplink goes on **SW1 Fa0/21**, as a second routed port. SW1 is a Layer 3 switch and will run OSPF just like the router does. Every group's SW1 plugs into `CORE-SW`, and together with the core router they all share one broadcast segment: `10.0.0.0/24`. Your SW1 takes `10.0.0.G`. The core router takes `10.0.0.254`.

This is a real OSPF broadcast network, so you will also see a **designated router** election happen.

Task 12.1 — Remove the static routes

> **Where to type this:** R · console cable · start from the `R-GRPG#` prompt
> Substitute your own peer group's network and next hop.

```
! List every static route so you know exactly what you are removing
show ip route static
!
! Enter global configuration mode
configure terminal
!
! Remove the route to your peer group
no ip route 10.2.0.0 255.255.0.0 10.1.255.2
!
! Remove the route down to your own LAN
no ip route 10.G.0.0 255.255.0.0 10.G.1.2
!
! Leave configuration mode
end
!
! Confirm you really did break it - this should now fail
ping 10.2.100.1
```

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Remove the default route as well
no ip route 0.0.0.0 0.0.0.0 10.G.1.1
!
! Leave configuration mode
end
```

Everything is broken again. Good. Now build it properly.

Task 12.2 — Connect SW1 to the classroom backbone

Run one cable from **SW1 Fa0/21** to the port on `CORE-SW` that the instructor gives you.

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt
> Replace **G** with your group number in the address.

```
! Enter global configuration mode
configure terminal
!
! A second routed port, this time facing the classroom backbone
interface fa0/21
 description ROUTED UPLINK TO CORE-SW
 no switchport
!
! Every group shares 10.0.0.0/24. Your host address is your group number.
 ip address 10.0.0.G 255.255.255.0
 no shutdown
 exit
!
! Leave configuration mode
end
!
! Confirm you can reach the core router before configuring OSPF
ping 10.0.0.254
```

Task 12.3 — Configure OSPF on the router

The whole classroom is a single OSPF **area 0**.

> **Where to type this:** R · console cable · start from the `R-GRPG#` prompt
> The wildcard mask is the inverse of the subnet mask. `255.255.255.252` becomes `0.0.0.3`. `255.255.255.255` becomes `0.0.0.0`. `255.255.255.0` becomes `0.0.0.255`.

```
! Enter global configuration mode
configure terminal
!
! Start OSPF process 1 - the number is local to this router only
router ospf 1
!
! Pin the router ID to the loopback so that it never changes
 router-id 10.G.0.1
!
! Send no hellos anywhere by default, then open only the links you mean to
! This is the safe way round: deny everything, then permit
 passive-interface default
!
! Speak OSPF down to SW1
 no passive-interface fa0/0
!
! Speak OSPF across to your peer group
 no passive-interface fa0/1
!
! Advertise the loopback - a /32 needs wildcard 0.0.0.0
 network 10.G.0.1 0.0.0.0 area 0
!
! Advertise the link to SW1
 network 10.G.1.0 0.0.0.3 area 0
!
! Advertise the peer /30 - use YOUR link's subnet here
 network 10.1.255.0 0.0.0.3 area 0
!
! Make cost meaningful above 100 Mbps - every device in the room uses 1000
 auto-cost reference-bandwidth 1000
 exit
!
! Leave configuration mode and save
end
```

```
write memory
```

Task 12.4 — Configure OSPF on SW1

> **Where to type this:** SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! Start OSPF on the Layer 3 switch
router ospf 1
!
! Pin the router ID to the management SVI address
 router-id 10.G.99.1
!
! Do not send hellos into the user VLANs - nobody there speaks OSPF
 passive-interface default
!
! Speak OSPF on the routed port to your router
 no passive-interface fa0/24
!
! Speak OSPF on the routed port to the classroom backbone
 no passive-interface fa0/21
!
! Advertise the backbone segment
 network 10.0.0.0 0.0.0.255 area 0
!
! Advertise the routed link to the router
 network 10.G.1.0 0.0.0.3 area 0
!
! Advertise all three of your VLANs
 network 10.G.99.0 0.0.0.255 area 0
 network 10.G.100.0 0.0.0.255 area 0
 network 10.G.200.0 0.0.0.255 area 0
!
! The same reference bandwidth as everybody else
 auto-cost reference-bandwidth 1000
 exit
!
! Leave configuration mode and save
end
write memory
```

Task 12.5 — Secure the adjacencies

Apply this to every interface that has an OSPF neighbour. Both ends of every link must use the **same key number and the same password**, or the adjacency will drop.

> **Where to type this:** R · console cable · start from the `R-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! The link down to SW1
interface fa0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 somnog9
 exit
!
! The link across to your peer group - agree this with them first
interface fa0/1
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 somnog9
 exit
!
```

```
! Leave configuration mode and save
end
write memory
```

Where to type this: SW1 · console cable · start from the `SW1-GRPG#` prompt

```
! Enter global configuration mode
configure terminal
!
! The routed port to your router
interface fa0/24
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 somnog9
 exit
!
! The routed port to the classroom backbone
interface fa0/21
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 somnog9
 exit
!
! Leave configuration mode and save
end
write memory
```

Verify

Where to type this: R · console cable · from the `R-GRPG#` prompt

```
! You should have neighbours in state FULL: SW1, and your peer group's router
show ip ospf neighbor
!
! Every intended interface, in area 0, with its cost
show ip ospf interface brief
!
! Routes marked O for networks you never typed a route for
show ip route ospf
```

Where to type this: SW1 · console cable · from the `SW1-GRPG#` prompt

```
! On the backbone you will see several neighbours, and a DR and BDR election
show ip ospf neighbor
!
! Which switch won the designated router election on 10.0.0.0/24
show ip ospf interface fa0/21
!
! The link-state database - every device in the classroom appears here
show ip ospf database
!
! You should now be able to reach ANY group, not just your peer
ping 10.4.100.1
```

Task 12.6 — Watch a failure and a recovery

Start a continuous ping to a distant group's gateway, then break your peer link.

Where to type this: R · console cable · start from the `R-GRPG#` prompt

```
! Watch adjacency events happen in real time
```

```
debug ip ospf adj
!
! Enter global configuration mode
configure terminal
!
! Break the link to your peer group
interface fa0/1
 shutdown
 exit
end
!
! The path has changed - traffic now goes via the backbone instead
show ip route 10.2.0.0
!
! Restore the link
configure terminal
interface fa0/1
 no shutdown
 exit
end
!
! ALWAYS turn debugging off when you are finished
undebug all
```

Questions

46. You deleted your static routes and now you have **more** routes than before. Where did they come from?
47. What does `passive-interface default` protect you from? Give a real example of what could go wrong without it.
48. Why did we set the router ID manually instead of letting OSPF choose one?
49. On the `10.0.0.0/24` backbone there is a DR and a BDR, but on your peer /30 there is neither. Why does OSPF treat the two links differently?
50. Compare your answer to Lab 11 question 4 with what you just saw. Write two sentences on why service providers do not run networks on static routes.

Lab 13 — Reaching the Internet: NAT on the core router

Devices: CORE — configured by the instructor, observed and verified by every group **Time:** 40 minutes

Objective

Understand how a whole classroom of private addresses reaches the public Internet through one router, and confirm from your own laptop that it works.

Before the instructor starts

The core router must already see a route to all six groups.

> **Where to type this:** R · console cable · from the `R-GRPG#` prompt

```
! Your adjacency with SW1 and with your peer must be FULL
show ip ospf neighbor
!
! You can reach the core router
ping 10.0.0.254
```

If your group's network is missing from the core's routing table, your group is the one holding up the lab. Go back to Lab 12.

Task 13.1 — What the instructor configures on CORE

Follow along on the projector. Copy this block into your notebook.

> **Where to type this:** CORE · instructor's console · from the `CORE#` prompt

```
! Enter global configuration mode
configure terminal
!
! The interface facing the Internet uplink
interface fa0/1
 description INTERNET UPLINK
 ip address dhcp
!
! Mark this as the OUTSIDE of the NAT boundary
 ip nat outside
 no shutdown
 exit
!
! The interface facing the classroom backbone
interface fa0/0
 description CLASSROOM BACKBONE
 ip address 10.0.0.254 255.255.255.0
!
! Mark this as the INSIDE of the NAT boundary
 ip nat inside
 no shutdown
 exit
!
! Define who is allowed to be translated
ip access-list standard NAT-CLASSROOM
 permit 10.0.0.0 0.255.255.255
 exit
!
! Translate every permitted source to the single outside address,
! using port numbers to tell the sessions apart. This is PAT.
ip nat inside source list NAT-CLASSROOM interface fa0/1 overload
!
! Tell all six groups where the exit is
```

```
router ospf 1
 router-id 10.0.0.254
 network 10.0.0.0 0.0.0.255 area 0
 default-information originate always
 exit
!
! Leave configuration mode and save
end
write memory
```

Task 13.2 — A static NAT for one group's server

One group volunteers a laptop in VLAN 200 as a web server. The instructor publishes it to the Internet.

Where to type this: CORE · instructor's console · from the `CORE#` prompt

```
! Enter global configuration mode
configure terminal
!
! Map port 8080 on the public address to port 80 on the inside server
ip nat inside source static tcp 10.G.200.50 80 interface fa0/1 8080
!
! Leave configuration mode
end
!
! The static mapping appears immediately, before any traffic flows
show ip nat translations
```

Task 13.3 — What every group verifies

Where to type this: R · console cable · from the `R-GRPG#` prompt

```
! A default route learned from OSPF, marked O*E2
show ip route 0.0.0.0
!
! The router reaches the Internet
ping 8.8.8.8
!
! Sourcing from your private loopback still works - NAT translated it
ping 8.8.8.8 source loopback 0
```

Where to type this: on Student A's laptop

```
rem A laptop in VLAN 100 reaches the Internet
ping 8.8.8.8
!
rem Count the hops inside the classroom before the packet leaves
tracert 8.8.8.8
```

Task 13.4 — Watch the translation table

Where to type this: CORE · instructor's console · on the projector, while the whole room generates traffic

```
! Every active session: inside address and port, mapped to outside address and port
show ip nat translations
!
! Totals, hits, misses, and which interfaces are inside and outside
show ip nat statistics
```

```
!
! Clear all dynamic translations and watch them rebuild
clear ip nat translation *
```

Questions

51. Thirty laptops share one public IP address. What field in the packet lets the core router tell one laptop's session from another's?
52. Your laptop's address is `10.G.100.21`. What source address does a server on the Internet actually see?
53. Why can a host on the Internet not start a connection to your laptop, and what did the instructor have to do in Task 13.2 to make one host reachable?
54. The core advertises a default route with `default-information originate always`. What does `always` change, and why is that both convenient and dangerous?
55. NAT breaks the end-to-end principle of the Internet. Name two protocols or applications that struggle because of it, and say how IPv6 changes the picture.

Appendix A — Instructor notes

Why the labs are ordered this way

The notebook order is not the agenda order, and that is deliberate. The agenda teaches STP before VLANs; the labs teach VLANs first, because STP is only interesting once there is a trunk worth protecting.

Session	Labs to run
Day 1, 14:00 — Lab setup overview	Lab 1
Day 1, 16:00 — Cisco config essentials	Finish Lab 1, plus a `show` command drill
Day 2, 08:00 — STP theory	Theory only; tell them the lab comes after the trunk exists
Day 2, 09:00 — VLANs	Labs 2, 3, 4, 5
Day 2, 11:00 — Advanced layer 2	Labs 6, 7
Day 2, 14:00 — Routing basics	Labs 8, 9, then Lab 10
Day 2, 16:00 — Static routing	Lab 11
Day 3, 08:00 — OSPF	Lab 12
Day 3, 09:30 — BGP basics	Theory, plus the optional extension below
Day 3, 11:00 — NAT/PAT	Lab 13

The three deliberate failures

These are the spine of the workshop. Protect them. If a fast group fixes them early, the lesson evaporates.

56. **Lab 5** — VLAN 100 pings VLAN 100; VLAN 100 cannot ping VLAN 200. Both laptops are on the same switch. This is what a VLAN *is*.
57. **Lab 6** — the SW1–SW2 link is an access port in VLAN 100, so VLAN 100 crosses it and VLAN 200 does not. One wire, one VLAN. This is what a trunk is *for*.
58. **Lab 7** — the SVI exists, and the switch itself can reach both VLANs, but hosts still cannot cross. `ip routing` is disabled by default on the 3560. This separates "has an IP address" from "forwards packets", and it is the one students remember.

Do not let anyone type `ip routing` during Lab 4.

Why the backbone uplink sits on SW1 and not on the router

The router has exactly two FastEthernet ports. Fa0/0 is the pod uplink and Fa0/1 is the peer-group link, which leaves nothing for the classroom backbone. SW1 is a Layer 3 switch running OSPF, so putting the backbone uplink on SW1 Fa0/21 is both practical and realistic — in a real campus the distribution switch often holds the uplink.

It also buys a teaching point for free. The `10.0.0.0/24` backbone is a broadcast segment shared by seven OSPF speakers, so students see a DR and BDR election, while their peer /30 is a point-to-point link with neither. That contrast is hard to demonstrate otherwise.

Why SW1 takes `.1` and SW2 takes `.2` in VLAN 200

SW2 must have an IP address inside VLAN 200 to serve DHCP for it — a Cisco DHCP server selects its pool from the receiving interface's subnet. But SW1 is the pod's Layer 3 device and must own the gateway address `10.G.200.1`. Giving SW2 `.2` lets both facts be true without an address conflict.

SW2's pool hands out `10.G.200.1` as the default router from the very beginning, and that address simply does not answer until Lab 7. That silence is the lesson.

Pre-workshop checklist

- Erase and reload all 12 pod switches and 6 pod routers. Delete `vlan.dat` on every switch.
- Confirm every 3560v2 runs an IP Base or IP Services image, and that `ip routing` and `router ospf` are available. Test-configure one unit.
- **Make or buy six crossover cables** for the router-to-router links, and label them clearly. This is the single most likely cause of a stuck lab on Day 2.
- Cable `CORE-SW` to all six groups' SW1 Fa0/21, plus the core router's Fa0/0.
- Pre-configure `CORE` with `10.0.0.254/24`, OSPF area 0 and MD5 key `somnog9`, but leave NAT until Lab 13.
- Count USB-to-serial adapters: one per group minimum, two is much better.
- Check laptops for Ethernet ports; have USB NIC adapters ready for the ones without.
- Print one notebook per student, and one copy of Appendix B per pod.

Common errors and what to say

Symptom	Cause	Prompt to give the student
Router-to-router link stays down	Straight-through cable instead of crossover	"What is different about a cable between two routers?"
SVI stays `down/down`	No access port in that VLAN is up	"Which physical ports are in that VLAN, and is anything plugged in?"
Laptop gets no DHCP address	Firewall, Wi-Fi still on, or wrong VLAN on the port	"Show me `show ip dhcp binding` and `show interfaces status`"
Trunk will not come up	One side never got `switchport mode trunk`	"Run `show interfaces trunk` on both sides and compare"
EtherChannel shows `(s)`	Mismatched trunk, VLAN, speed or duplex between members	"Compare the running config of Fa0/22 and Fa0/23"
OSPF neighbour stuck in `EXSTART`	MTU mismatch	"What MTU does each side report in `show ip ospf interface`?"
OSPF neighbour stuck in `INIT`	One side has authentication, the other does not	"Read your own log messages out loud"
Ping works from the router but not from a laptop	No default gateway, or the return route is missing	"Where does the reply packet need to go, and does that device know the way back?"

Optional extension — eBGP for the Day 3, 09:30 session

For groups that finish Lab 12 early. Group G becomes AS `6500G`; the core is AS `65000`. The BGP session runs from SW1, which is the device with the backbone adjacency.

```
! Enter global configuration mode
configure terminal
!
! Start BGP in your own autonomous system
router bgp 6500G
!
! Use a stable address as the BGP router ID
 bgp router-id 10.G.99.1
!
! Peer with the core router across the backbone segment
 neighbor 10.0.0.254 remote-as 65000
 neighbor 10.0.0.254 description UPSTREAM CORE
!
! Advertise your aggregate - nothing more specific
 network 10.G.0.0 mask 255.255.0.0
 exit
end
!
! The session state should reach a number, not Active or Idle
show ip bgp summary
!
! Every other group's /16, with its AS path
show ip bgp
```

Ask a student to read an AS path out loud and explain what each number means. Then have one group apply a prefix-list filter and watch a neighbour's prefix disappear.

Appendix B — Per-group quick reference

Cut this page out and tape it to your pod.

	G1	G2	G3	G4	G5	G6
Block	10.1.0.0/16	10.2.0.0/16	10.3.0.0/16	10.4.0.0/16	10.5.0.0/16	10.6.0.0/16
Loopback0	10.1.0.1	10.2.0.1	10.3.0.1	10.4.0.1	10.5.0.1	10.6.0.1
R Fa0/0	10.1.1.1/30	10.2.1.1/30	10.3.1.1/30	10.4.1.1/30	10.5.1.1/30	10.6.1.1/30
SW1 Fa0/24	10.1.1.2/30	10.2.1.2/30	10.3.1.2/30	10.4.1.2/30	10.5.1.2/30	10.6.1.2/30
VLAN 99 gw	10.1.99.1	10.2.99.1	10.3.99.1	10.4.99.1	10.5.99.1	10.6.99.1
VLAN 100 gw	10.1.100.1	10.2.100.1	10.3.100.1	10.4.100.1	10.5.100.1	10.6.100.1
VLAN 200 gw	10.1.200.1	10.2.200.1	10.3.200.1	10.4.200.1	10.5.200.1	10.6.200.1
SW2 VLAN 200	10.1.200.2	10.2.200.2	10.3.200.2	10.4.200.2	10.5.200.2	10.6.200.2
SW1 Fa0/21	10.0.0.1/24	10.0.0.2/24	10.0.0.3/24	10.0.0.4/24	10.0.0.5/24	10.0.0.6/24
Peer group	G2	G1	G4	G3	G6	G5
R Fa0/1	10.1.255.1/30	10.1.255.2/30	10.3.255.1/30	10.3.255.2/30	10.5.255.1/30	10.5.255.2/30

The core router is `10.0.0.254` on the backbone segment.

Appendix C — Command reference

Verification commands worth memorising

Command	What it tells you
`show ip interface brief`	Every interface, its address, and whether it is up at Layer 1 and Layer 2
`show vlan brief`	Every VLAN and which ports are in it
`show interfaces status`	Port status, VLAN, duplex and speed in one screen
`show interfaces trunk`	Which ports are trunks, their native VLAN and allowed VLANs
`show mac address-table`	Which MAC is on which port, in which VLAN
`show spanning-tree vlan <id>`	Root bridge, port roles and port states
`show etherchannel summary`	Bundle members and their flags
`show ip route`	The full routing table
`show ip route <prefix>`	Why the device chose a particular path
`show ip ospf neighbor`	Adjacency states
`show ip ospf interface brief`	Which interfaces run OSPF, in which area, at what cost
`show ip nat translations`	Active NAT sessions
`show cdp neighbors detail`	What is on the other end of every cable, and its IP address
`show logging`	The log buffer — read this before asking for help

Recovery and housekeeping

Command	What it does
`terminal monitor`	Show log messages on a Telnet or SSH session
`undebug all`	Turn off all debugging — always do this
`clear ip ospf process`	Restart OSPF and force new adjacencies
`clear mac address-table dynamic`	Flush learned MAC addresses
`clear ip dhcp binding *`	Release all DHCP leases
`copy running-config startup-config`	The long form of `write memory`
`show running-config interface fa0/22`	The configuration of one interface only

Wildcard masks for OSPF

Subnet mask	Prefix	Wildcard mask
`255.255.255.255`	/32	`0.0.0.0`

Subnet mask	Prefix	Wildcard mask
255.255.255.252	/30	0.0.0.3
255.255.255.0	/24	0.0.0.255
255.255.0.0	/16	0.0.255.255

SomNOG9 Track 1 — Network Infrastructure. Prepared for the Somali Network Operators' Group.