Two-Site Enterprise Network Design

Cisco Packet Tracer Lab — Configuration Report

VLAN Segmentation | OSPF Routing | DHCP | Extended ACL

Topology: Sydney <-> Melbourne | Cisco 2911 Routers + 2960 Switches

Tools: Cisco IOS CLI • OSPF Area 0 • 802.1Q Trunking • Extended ACLs

01
Project Overview

This project demonstrates a complete multi-site enterprise network build using Cisco Packet Tracer. Starting from a single-site layout, the network evolved into a fully routed, dual-site design featuring VLAN segmentation, dynamic routing via OSPF, automated IP addressing through DHCP, and traffic policy enforcement using extended ACLs.

Sites	Routers	Switches	VLANs
2	**2**	**2**	**4**
Sydney · Melbourne	Cisco 2911	Cisco 2960	VLAN 10 · VLAN 20

Routing Protocol	DHCP Pools	ACL Type	Transit Link
OSPF	**4**	**Extended**	**10.0.0.0/30**
Area 0	Per VLAN per site	ACL 100 inbound	GigE point-to-point

Network Topology

The final topology includes two sites connected over a routed GigabitEthernet point-to-point link. Each site has a Cisco 2911 router performing inter-VLAN routing and DHCP services, and a Cisco 2960 switch providing VLAN access ports and 802.1Q trunk connectivity to the router.

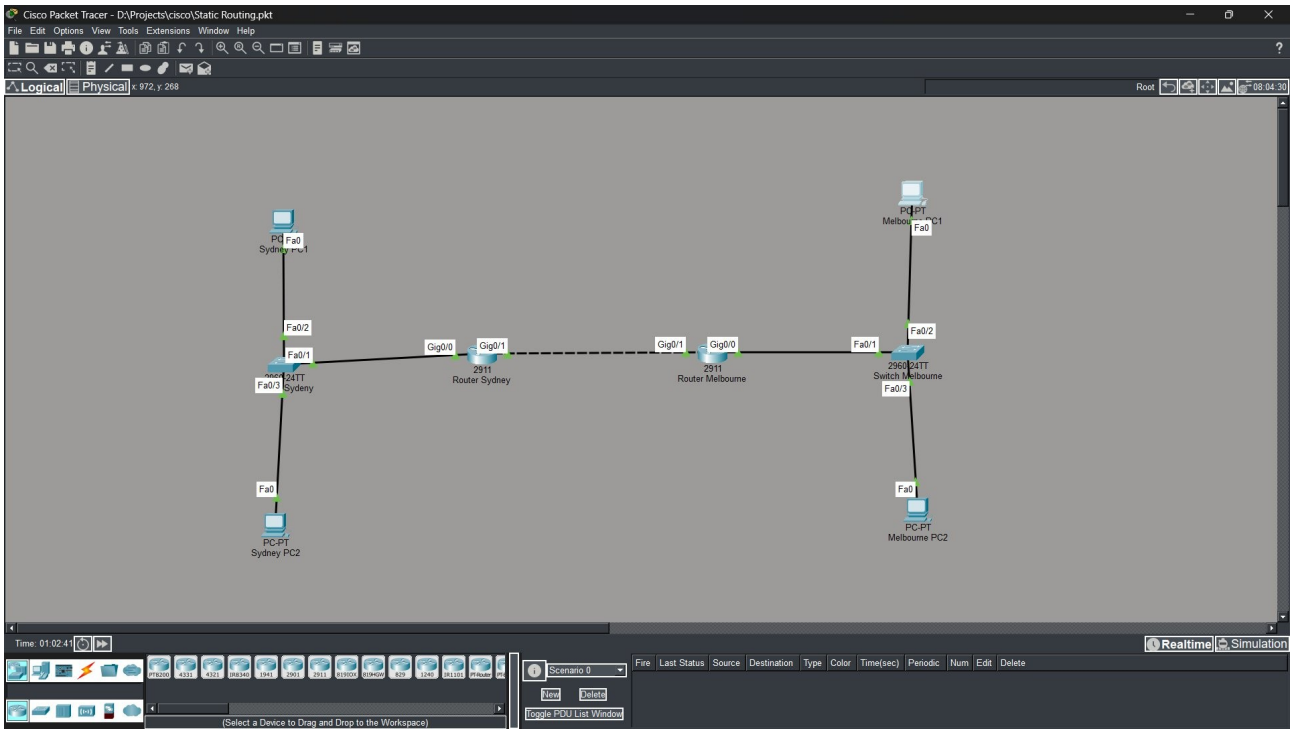

Figure 1 — Full two-site logical topology: Sydney (left) and Melbourne (right) connected via GigabitEthernet0/1

02
VLAN Segmentation & Trunking

Each site uses two VLANs to separate traffic by department. VLAN 10 (Sales) carries traffic from PC1, while VLAN 20 (IT) carries traffic from PC2. FastEthernet access ports are assigned per VLAN, and FastEthernet0/1 is configured as an 802.1Q trunk toward the router on both switches.

VLAN	Name	Sydney Interface	Melbourne Interface	Subnet
10	Sales	Fa0/2	Fa0/2	192.168.10.0/24 · 192.168.30.0/24
20	IT	Fa0/3	Fa0/3	192.168.20.0/24 · 192.168.40.0/24
Trunk	802.1Q	Fa0/1 → Router	Fa0/1 → Router	VLANs 1,10,20 active

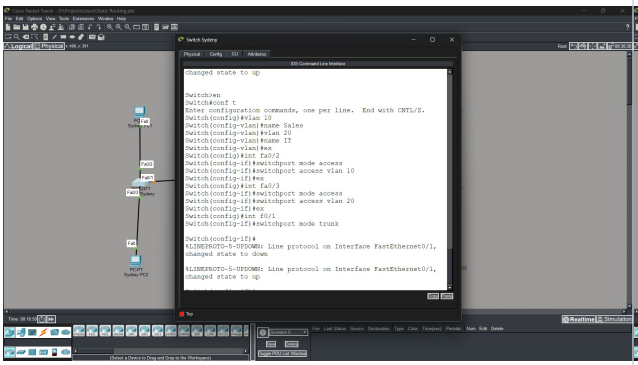

Figure 2 — Sydney switch: VLAN 10 (Sales) & VLAN 20 (IT) creation, access port assignment, trunk on Fa0/1

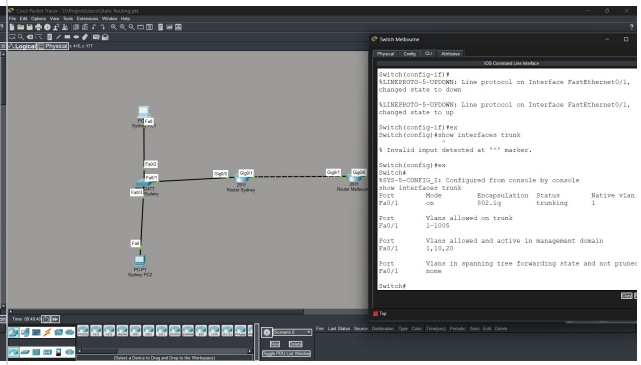

Figure 3 — Melbourne switch trunk verification: Fa0/1 trunking 802.1Q, VLANs 1,10,20 active in management domain

03 Router-on-a-Stick Inter-VLAN Routing

Both routers use the router-on-a-stick technique, creating logical subinterfaces on GigabitEthernet0/0 with 802.1Q encapsulation. Each subinterface is assigned the default gateway IP for its respective VLAN subnet and acts as the Layer 3 gateway for end devices.

Router	Subinterface	Encapsulation	IP Address	VLAN
Sydney	G0/0.10	802.1Q 10	192.168.10.1/24	10 (Sales)
Sydney	G0/0.20	802.1Q 20	192.168.20.1/24	20 (IT)
Melbourne	G0/0.10	802.1Q 10	192.168.30.1/24	10 (Sales)
Melbourne	G0/0.20	802.1Q 20	192.168.40.1/24	20 (IT)

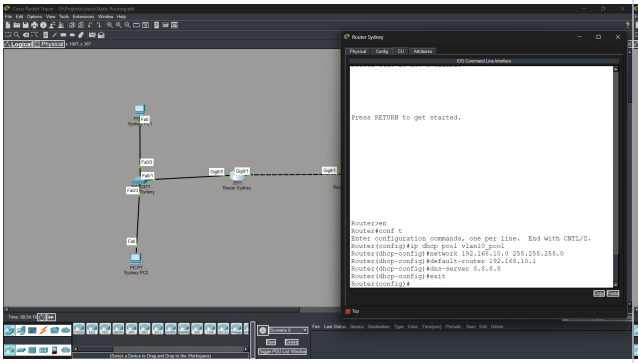

Figure 4 — Sydney router subinterface configuration: G0/0.10 (192.168.10.1) and G0/0.20 (192.168.20.1) with OSPF update

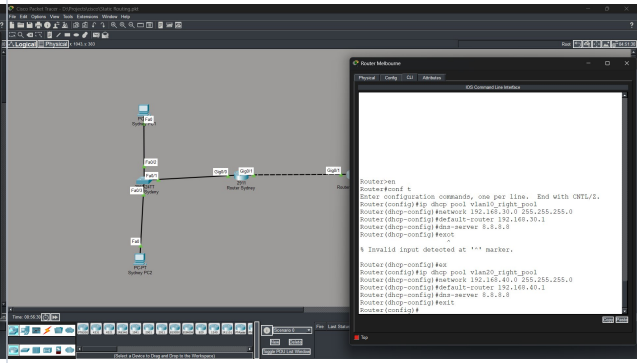

Figure 5 — Melbourne router subinterface configuration: G0/0.10 (192.168.30.1) and G0/0.20 (192.168.40.1) with OSPF routes learned

04 DHCP Configuration

Each router serves as a DHCP server for its local VLANs. Four pools were created — two per site — covering all four VLAN subnets. All pools use Google DNS (8.8.8.8). Successful DHCP leases were confirmed by checking the binding table and the IP configuration of end devices.

Pool Name	Router	Network	Default Gateway	DNS
vlan10_pool	Sydney	192.168.10.0/24	192.168.10.1	8.8.8.8
vlan20_pool	Sydney	192.168.20.0/24	192.168.20.1	8.8.8.8
vlan10_right_pool	Melbourne	192.168.30.0/24	192.168.30.1	8.8.8.8
vlan20_right_pool	Melbourne	192.168.40.0/24	192.168.40.1	8.8.8.8

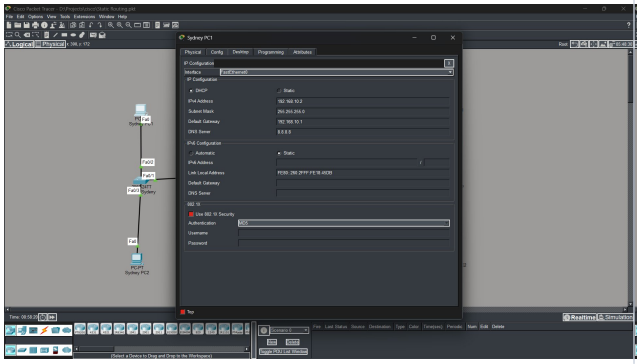

Figure 6 — Sydney router: vlan10_pool DHCP configuration (192.168.10.0/24)

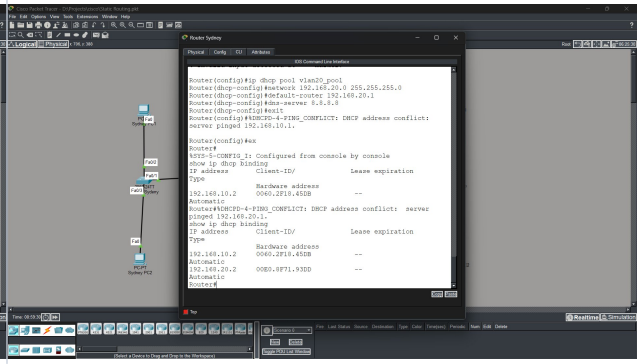

Figure 7 — Melbourne router: vlan10_right_pool and vlan20_right_pool DHCP configuration

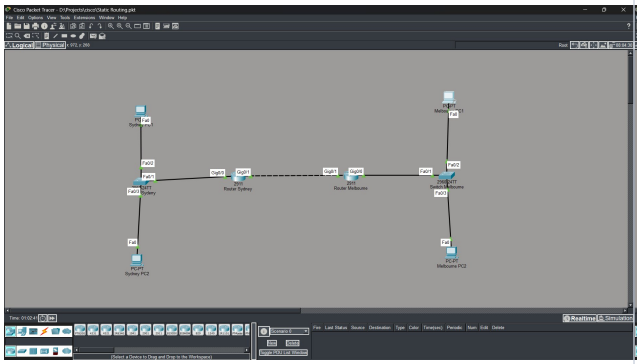

Figure 8 — Sydney DHCP binding table showing active leases: 192.168.10.2 and 192.168.20.2

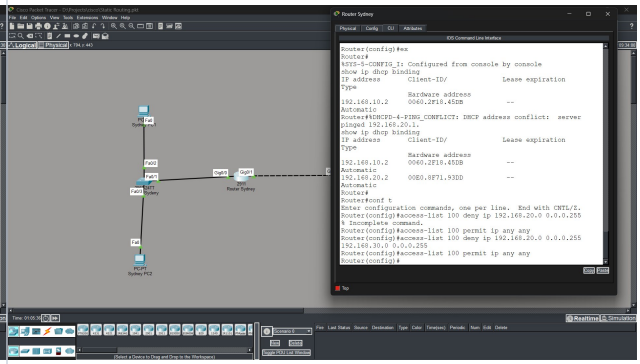

Figure 9 — Sydney PC1 IP config panel confirming DHCP lease: 192.168.10.2/24, gateway 192.168.10.1, DNS 8.8.8.8

05 OSPF Dynamic Routing

Static routes were replaced with OSPF process 1 in area 0 on both routers. Sydney uses router ID 1.1.1.1 and Melbourne uses 2.2.2.2. After OSPF was enabled and the transit link (10.0.0.0/30) was advertised, a full adjacency formed between the two routers. LAN-facing subinterfaces are set as passive to prevent unnecessary OSPF hellos on end-device segments.

Router	Router ID	Networks Advertised	Passive Interfaces
Sydney	1.1.1.1	192.168.10.0/24, 192.168.20.0/24, 10.0.0.0/30	G0/0.10, G0/0.20
Melbourne	2.2.2.2	192.168.30.0/24, 192.168.40.0/24, 10.0.0.0/30	G0/0.10, G0/0.20

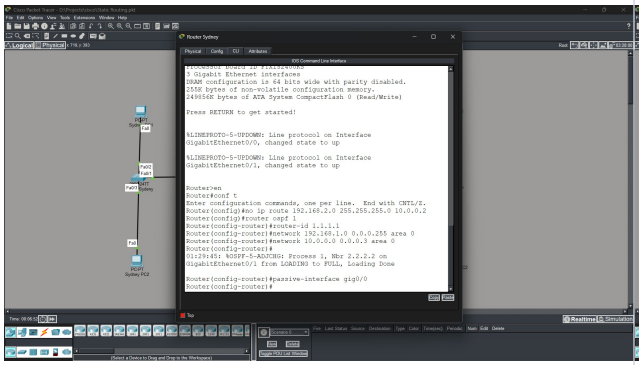

Figure 10 — Sydney router: OSPF process 1 configured, router ID 1.1.1.1, static route removed

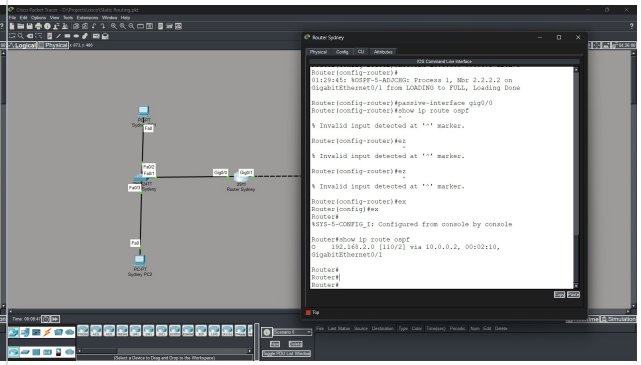

Figure 11 — OSPF adjacency confirmation: Nbr 2.2.2.2 reaches FULL state on GigabitEthernet0/1

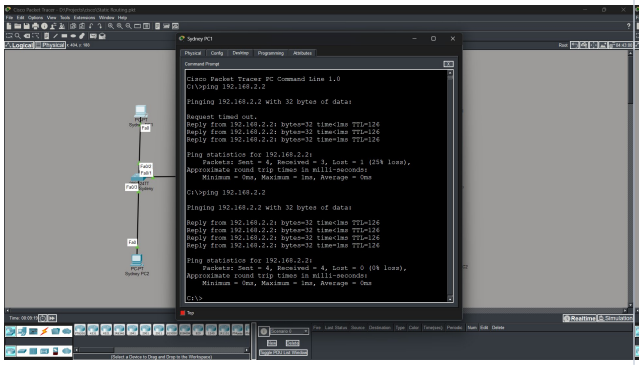

Figure 12 — Sydney OSPF routing table: 192.168.2.0 learned via 10.0.0.2 on GigabitEthernet0/1

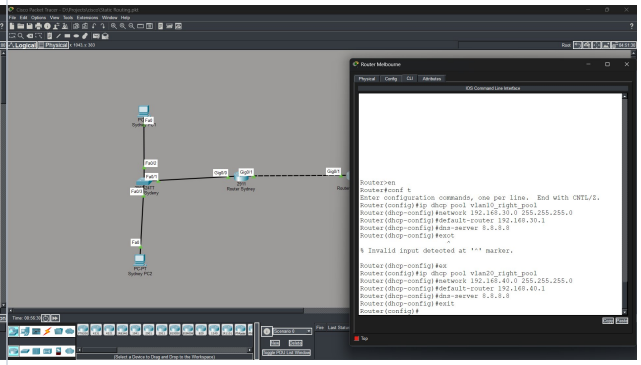

Figure 13 — Melbourne OSPF routing table: 192.168.10.0 and 192.168.20.0 learned via 10.0.0.1

06
Extended ACL — Traffic Policy

An extended ACL was applied to enforce a traffic policy: all IP traffic from Sydney VLAN 20 (192.168.20.0/24) destined for Melbourne VLAN 10 (192.168.30.0/24) is denied. All other traffic is permitted. The ACL is applied inbound on subinterface GigabitEthernet0/0.20 on the Sydney router.

ACL Configuration

```
access-list 100 deny ip 192.168.20.0 0.0.0.255 192.168.30.0 0.0.0.255
access-list 100 permit ip any any
interface GigabitEthernet0/0.20
ip access-group 100 in
```

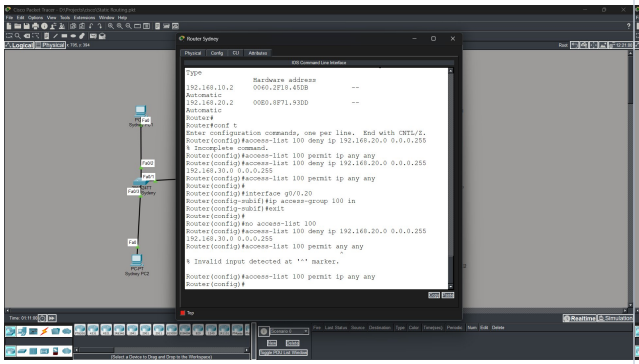

Figure 14 — ACL 100 configuration on Sydney router and application to G0/0.20 inbound

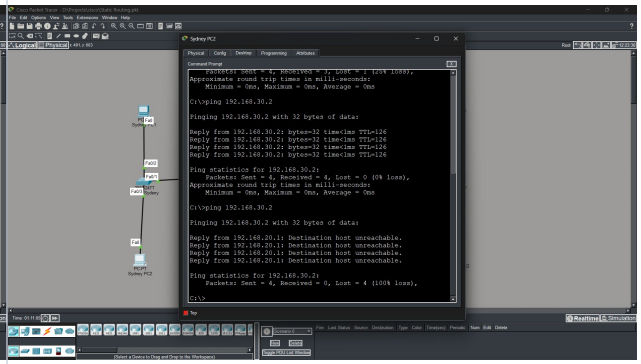

Figure 15 — Ping test from Sydney PC2: first test succeeds (before ACL), second test blocked with "Destination host unreachable" after ACL applied

07 Connectivity Testing & Verification

End-to-end connectivity was validated using ping tests between PCs across both sites. The first test from each PC typically shows one timeout due to ARP resolution, followed by 0% loss on repeat tests. After the ACL was applied, Sydney PC2 correctly received "Destination host unreachable" responses from 192.168.20.1 when attempting to reach Melbourne VLAN 10.

Source	Destination	Result	Notes
Sydney PC1	192.168.2.2	Pass (0% loss)	Initial 1 timeout (ARP), then stable
Melbourne PC1	192.168.1.2	Pass (0% loss)	Early routing phase
Melbourne PC1	192.168.10.2	Pass (0% loss)	Post-VLAN phase, stable after first atternpt
Sydney PC1	10.0.0.2 / 192.168.30.1	Pass (0% loss)	Transit and remote gateway reachable
Sydney PC2	192.168.30.2 (pre-ACL)	Pass (0% loss)	Before ACL applied
Sydney PC2	192.168.30.2 (post-ACL)	**BLOCKED**	ACL 100 denying VLAN 20 → VLAN 30

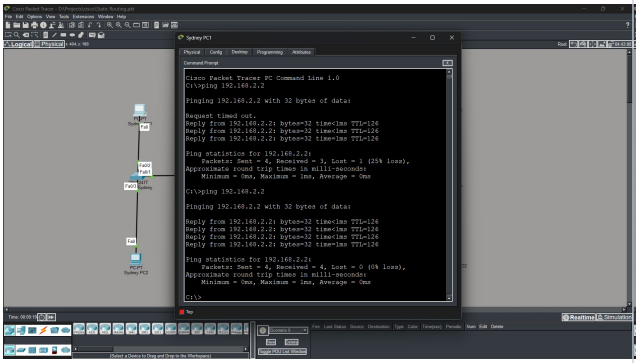

Figure 16 — Sydney PC1 ping to 192.168.2.2: first attempt shows 1 timeout (ARP), second run achieves 0% loss

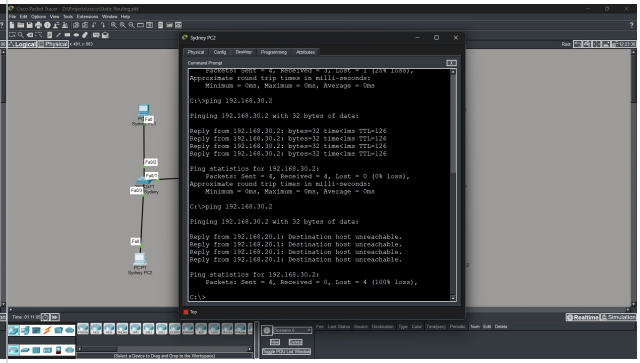

Figure 17 — Sydney PC2 ping to 192.168.30.2: successful before ACL, then blocked after ACL 100 applied

08

Troubleshooting & Issues Encountered

The build followed an iterative process. Several configuration mistakes were encountered and corrected, which is typical in Packet Tracer labs and demonstrates practical troubleshooting skills.

Incomplete ACL command

Initial ACL entry was entered without the destination network, resulting in "% Incomplete command." Corrected by re-entering the full extended ACL syntax with source and destination.

Mistyped permit statement

"permit any any" was entered without the required protocol keyword. Corrected to "permit ip any any".

Command mode error

"show ip route ospf" was entered while still in router config mode, producing "% Invalid input detected". Exited to privileged EXEC mode before running show commands.

DHCP address conflict

A DHCPD-4-PING_CONFLICT error appeared when the server pinged 192.168.20.1 (the gateway address) during pool initialisation. This is expected behaviour — the router detects the gateway is already reachable and logs a conflict before skipping it from the pool.

Typo exiting DHCP config

"#exot" was entered instead of "#exit" in DHCP configuration mode, producing an invalid input error. Corrected to "#ex".

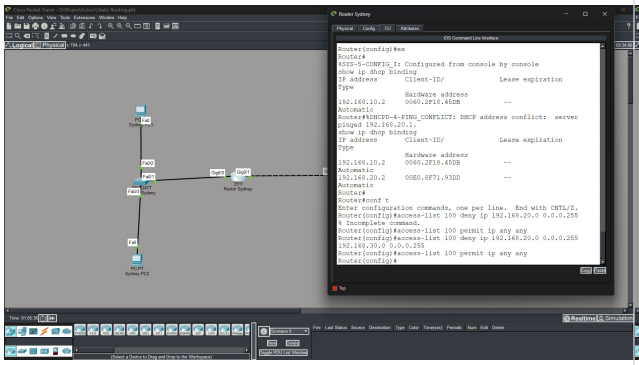

Figure 18 — ACL syntax errors and corrections visible in Sydney router CLI history

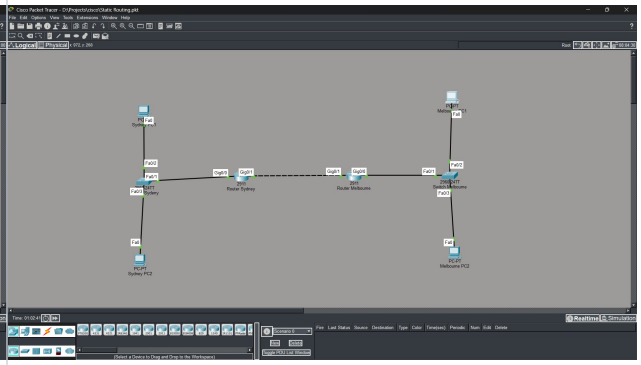

Figure 19 — DHCP conflict message and resolved binding table showing both leases active

09
Assessment & Summary

The project successfully demonstrates a full enterprise network build progression across all key areas of network engineering:

Phase 1 — Physical topology

Two-site layout with Cisco 2911 routers and 2960 switches, connected over a GigabitEthernet point-to-point link.

Phase 2 — VLAN segmentation

VLANs 10 (Sales) and 20 (IT) configured on both switches with 802.1Q trunk ports toward each router.

Phase 3 — Inter-VLAN routing

Router-on-a-stick subinterfaces on both routers enabling Layer 3 forwarding between VLANs without a Layer 3 switch.

Phase 4 — DHCP automation

Four DHCP pools providing automatic addressing to all four VLAN subnets, with verified bindings on all clients.

Phase 5 — OSPF dynamic routing

OSPF area 0 replacing static routes, full adjacency confirmed, and remote VLAN networks learned dynamically.

Phase 6 — Access control

Extended ACL blocking VLAN 20 traffic from Sydney reaching Melbourne VLAN 10, verified with ping tests before and after.

Technologies Demonstrated

✓ Cisco IOS CLI configuration	✓ 802.1Q VLAN trunking	✓ Router-on-a-stick inter-VLAN routing
✓ OSPF dynamic routing (area 0)	✓ DHCP server configuration	✓ Extended access control lists
✓ IPv4 subnetting and addressing	✓ Network connectivity verification	✓ Iterative troubleshooting

The final network state shows Sydney VLAN 20 traffic correctly blocked from reaching Melbourne VLAN 10 by the extended ACL, while all other inter-site traffic continues to flow via OSPF-learned routes through properly addressed VLAN gateways — confirming the policy is working as designed.