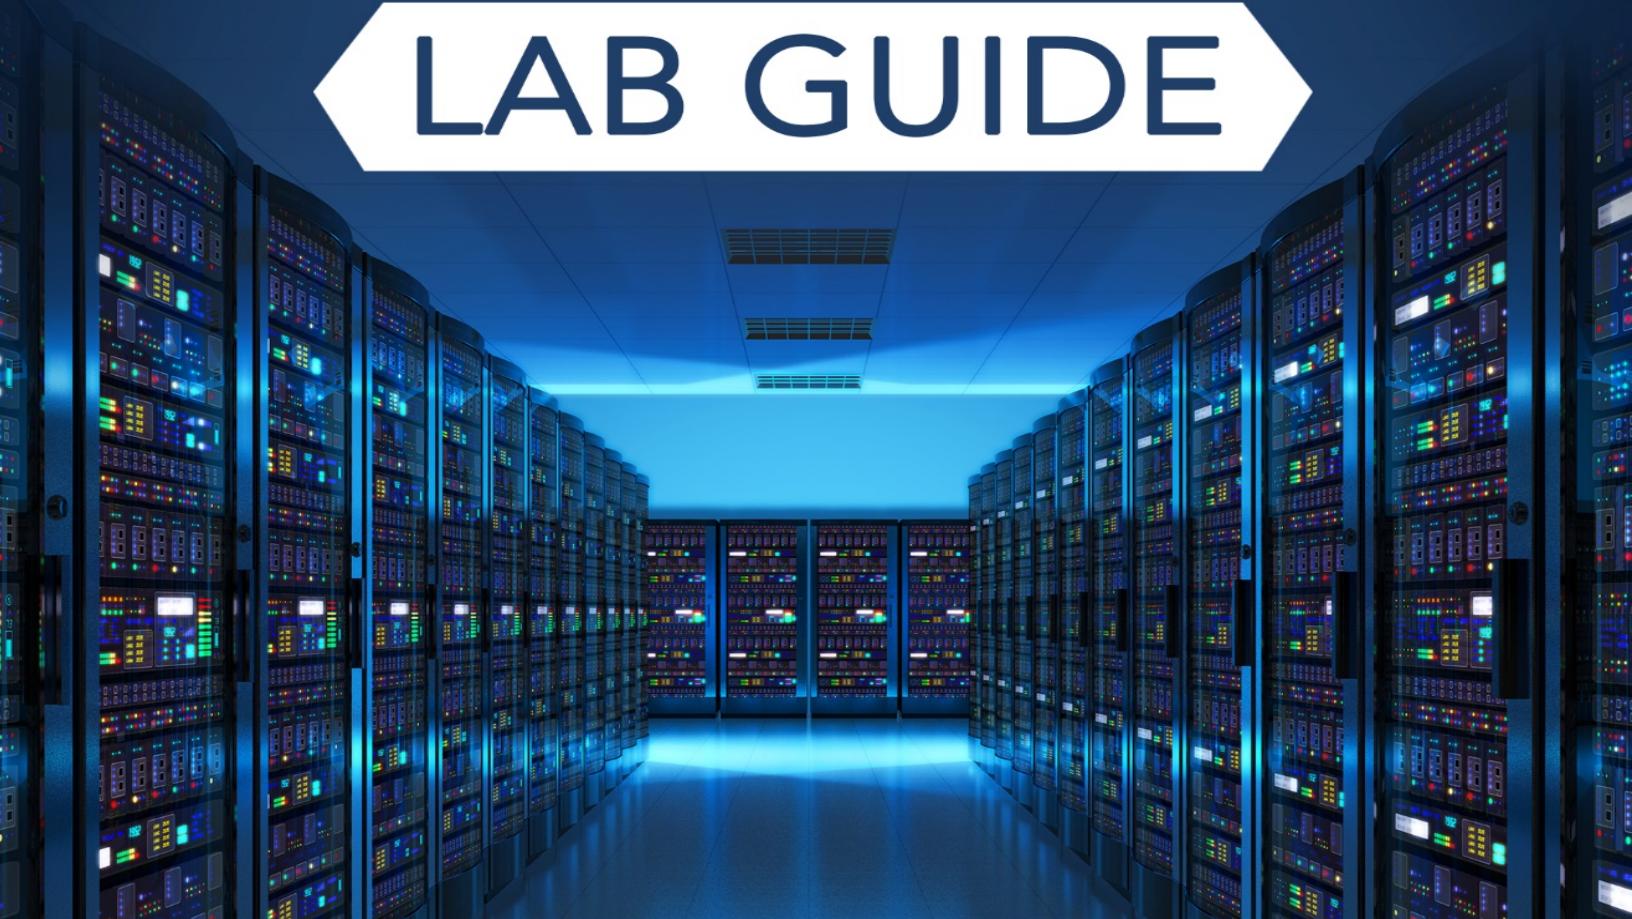




FLACKBOX
www.flackbox.com

CISCO CCNA

LAB GUIDE



NEIL ANDERSON

Clickable Table of Contents

Clickable Table of Contents	2
Introduction.....	4
How the Lab Works	4
Get the Complete Course	5
About the Author	5
04 The IOS Operating System - Lab Exercise	6
11 Cisco Device Functions – Lab Exercise.....	19
12 The Life of a Packet - Lab Exercise.....	26
12 The Life of a Packet – Answer Key.....	28
13 The Cisco Troubleshooting Methodology - Lab Exercise.....	31
13 The Cisco Troubleshooting Methodology – Answer Key.....	33
14 Cisco Router and Switch Basics - Lab Exercise.....	38
14 Cisco Router and Switch Basics - Answer Key	41
15 Cisco Device Management - Lab Exercise	47
15 Cisco Device Management - Answer Key	50
16 Routing Fundamentals - Lab Exercise.....	60
16 Routing Fundamentals - Answer Key.....	64
17 Dynamic Routing Protocols – Lab Exercise.....	79
17 Dynamic Routing Protocols - Answer Key	84
18 Connectivity Troubleshooting – Lab Exercise.....	103
18 Connectivity Troubleshooting - Answer Key	105
19-1 IGP Interior Gateway Protocol Fundamentals Configuration – Lab Exercise.....	107
19-1 IGP Interior Gateway Protocol Fundamentals Configuration - Answer Key	109
20-1 OSPF Configuration – Lab Exercise.....	114
20-1 OSPF Configuration - Answer Key	119
22-1 VLAN and Inter-VLAN Routing Configuration – Lab Exercise	137
22-1 VLAN and Inter-VLAN Routing Configuration - Answer Key.....	141
23-1 DHCP Configuration – Lab Exercise	152
23-1 DHCP Configuration - Answer Key.....	155
24-1 HSRP Configuration – Lab Exercise.....	160
24-1 HSRP Configuration - Answer Key	163
25-1 Spanning Tree Troubleshooting – Lab Exercise.....	170
25-1 Spanning Tree Troubleshooting - Answer Key	173

25-2 Spanning Tree Configuration – Lab Exercise	184
25-2 Spanning Tree Configuration - Answer Key.....	186
26-1 EtherChannel Configuration – Lab Exercise	190
26-1 EtherChannel Configuration - Answer Key.....	194
27-1 Port Security Configuration Lab Exercise	208
27-1 Port Security Configuration Answer Key.....	210
28-1 ACL Configuration – Lab Exercise	215
28-1 ACL Configuration - Answer Key.....	218
29-1 NAT Configuration – Lab Exercise	229
29-1 NAT Configuration - Answer Key.....	235
30-1 IPv6 Configuration - Lab Exercise	246
30-1 IPv6 Configuration - Answer Key.....	249
33-1 Cisco Device Security Configuration - Lab Exercise.....	265
33-1 Cisco Device Security Configuration - Answer Key.....	269
34 Network Device Management – Lab Exercise	276
34 Network Device Management – Answer Key.....	278
37 Wireless Fundamentals Configuration – Lab Exercise.....	280
37 Wireless Fundamentals Configuration - Answer Key	285
SPECIAL OFFER – Cisco CCNA Gold Bootcamp	312
SPECIAL OFFER – AlphaPrep 10 Day Trial	315

Disclaimer and Copyright

The information contained in this guide is for informational purposes only. Any advice that I give is my opinion based on my own experience. You should always seek the advice of a professional before acting on something that I have published or recommended.

The material in this guide may include information, products or services by third parties. Third Party Materials comprise of the products and opinions expressed by their owners. As such, I do not assume responsibility or liability for any Third Party material or opinions. You are responsible for complying with any legal requirements such as licensing of Third Party software.

No part of this publication shall be reproduced, transmitted, or sold in whole or in part in any form, without the prior written consent of the author. All trademarks and registered trademarks appearing in this guide are the property of their respective owners.

Introduction

Thanks very much for taking the time to download this free eBook. It contains complete configuration lab exercises and solutions to help you pass the Cisco CCNA 200-301 exam. You can also use it as a configuration reference for Cisco devices. I hope you can make use of it to expand your networking knowledge and further your career.

How the Lab Works

I wanted to make this a completely free resource and as simple to use as possible so the free software Packet Tracer is used for the labs.

To download the software, enroll in the 'Introduction to Packet Tracer' course at <https://www.netacad.com/campaign/ptdt-4>

Then go to <https://www.netacad.com>, click on 'Resources' near the top-right corner then 'Download Packet Tracer'

The downloadable projects were created in Packet Tracer version 7.3.0.838. Please ensure you are using this or a newer version of the software.

I've provided a lab startup file for each exercise so you can get up and running with the labs immediately. A download link is provided at the start of each exercise.

Please watch my short free video showing how to install and use Packet Tracer first if you haven't used the software before:

[How to install and use Packet Tracer](#)

If you have issues logging in to Packet Tracer, try clearing your web browser cache or using a different browser.

If you have any other problems with installation or logging in, Cisco have their own dedicated support team who can provide the best help:

https://learningnetwork.cisco.com/community/networking_academy

or <https://www.facebook.com/cisconetworkingacademy>

If you find any errors in the book, please let me know so I can correct them. You can email me at neil@flackbox.com

Get the Complete Course

The lab exercises here can be used on their own or as a complement to my [Cisco CCNA Gold Bootcamp course](#). It has the highest review ratings of any CCNA course online and includes over 30 hours of video tutorials, quizzes, study notes and advanced hands-on lab exercises. You can shortcut your path to CCNA certification by getting the course here:
<https://www.flackbox.com/cisco-ccna-training-course>

For practice tests I recommend [AlphaPrep](#). They partner with Cisco and the CCNA test provider Pearson to bring you the most accurate preparation tests, and their advanced test engine lets you know when you're ready for the exam.
[Click here for a 10 day free trial.](#)

About the Author



I'm Neil Anderson, you can visit my blog at
<https://www.flackbox.com>.

The main focus of my current role is delivery of technical training and development of course content for large enterprise and service provider customers such as Cisco, NetApp, Verizon and IBM.

I dropped out of school with no qualifications or future plans at the age of 15. When I got a little bit older and wiser I realised I should make a career for myself so I learned about IT technologies through books and online resources. It's my passion now to help you do the same.

Connect with me on social media:



04 The IOS Operating System - Lab Exercise

This lab explores basic navigation of the Cisco IOS operating system CLI (Command Line Interface). Only a single device is required.

This lab is a guided walkthrough of the IOS command line interface. Exercises for later sections will be split into two parts - first the tasks for you to complete on your own (without step by step instructions), and then an answer key showing you the solution.

Load the Startup Configuration

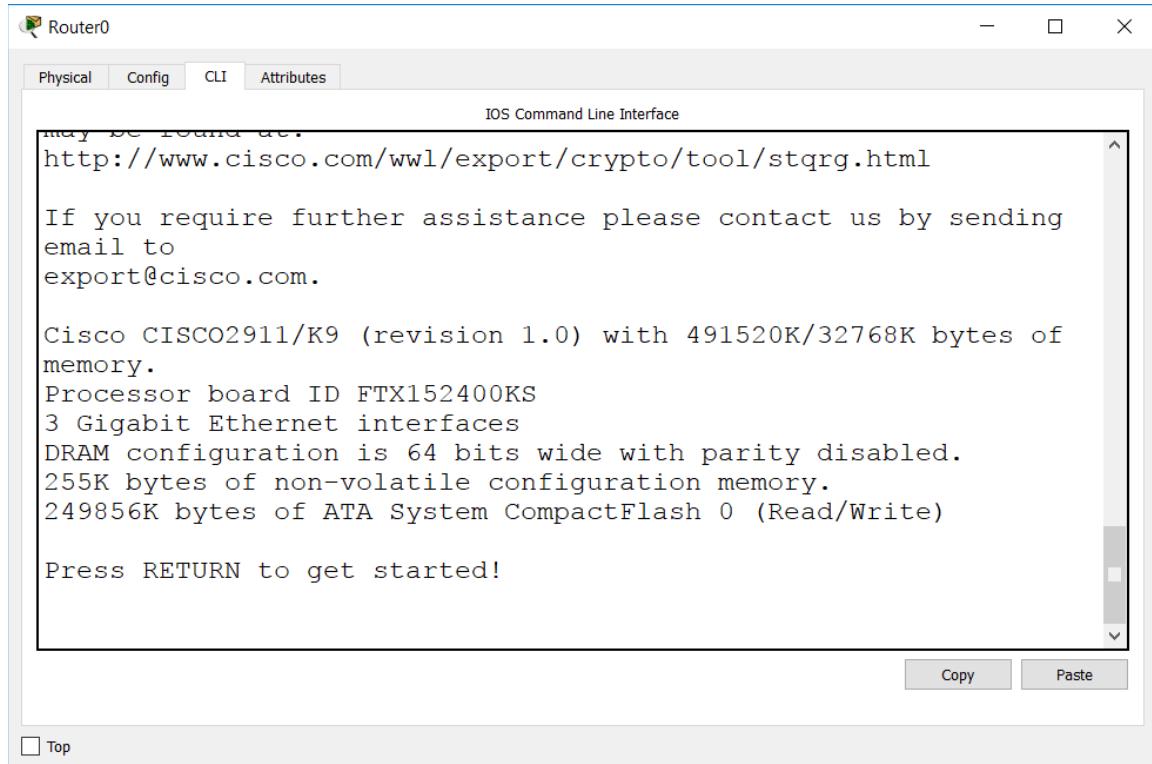
[Download the '04 The IOS Operating System.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Please watch my short free course showing how to install and use Packet Tracer first if you haven't used the software before:

[How to install and use Packet Tracer](#)

Connect To Your Device

Click on Router0 and then the CLI tab to access the console.



Router0

Physical Config CLI Attributes

IOS Command Line Interface

```
may be found at.
http://www.cisco.com/wwl/export/crypto/tool/stqrg.html

If you require further assistance please contact us by sending
email to
export@cisco.com.

Cisco CISCO2911/K9 (revision 1.0) with 491520K/32768K bytes of
memory.
Processor board ID FTX152400KS
3 Gigabit Ethernet interfaces
DRAM configuration is 64 bits wide with parity disabled.
255K bytes of non-volatile configuration memory.
249856K bytes of ATA System CompactFlash 0 (Read/Write)

Press RETURN to get started!
```

Copy Paste

Top

Press Return to get started, then enter Privileged Exec mode.

```
Router>enable
Router#
```

Reboot the device.

```
Router#reload
Proceed with reload? [confirm]
```

Observe the device going through the bootup process in the command line output. This is possible because we are using a console connection (we could not see this if we connected to an IP address on the device.)

If prompted to enter the initial configuration dialog after the device has booted up, enter 'no'.

```
Would you like to enter the initial configuration dialog?
[yes/no]: no
```

Explore User Exec Mode and CLI command help

Notice that you are in User Exec mode as indicated by the 'Router>' prompt. ('Router' will be replaced with the device hostname after you configure one.)

Router>

Enter a question mark to explore the commands that are available in User Exec mode.

Router>?

Exec commands:

```
<1-99> Session number to resume
connect Open a terminal connection
disable Turn off privileged commands
disconnect Disconnect an existing network connection
enable Turn on privileged commands
exit Exit from the EXEC
logout Exit from the EXEC
ping Send echo messages
resume Resume an active network connection
show Show running system information
ssh Open a secure shell client connection
telnet Open a telnet connection
terminal Set terminal line parameters
traceroute Trace route to destination
```

Only a very limited set of informational commands are available in User Exec mode and we won't typically be working here.

Enter the 'show run' command.

```
RouterX>show run
^
% Invalid input detected at '^' marker.
```

'show run' is a valid command but it's run at Privileged Exec mode, not User Exec, so the command fails.

This is the most common issue to trip up beginners at the IOS command line. If you see the 'invalid input' error then check you are at the correct level for the command you are trying to run.

Exploring Privileged Exec (Enable) Mode and Context Sensitive Help

Enter Privileged Exec mode. This mode is often commonly known as Enable mode. Notice that the prompt changes to 'Router#'

```
Router>enable  
Router#
```

Drop back to User Exec mode.

```
Router#disable  
Router>
```

Go back to Privileged Exec mode by using shortened command abbreviation.

```
Router>en  
Router#
```

Command abbreviation only works when you enter letters which could only match one unique command. Attempt to return to User Exec mode by entering the command 'di'

```
Router#di  
% Ambiguous command: "di"
```

Check to see all the possible commands which begin with the letters 'di'

```
Router#di?  
dir disable disconnect
```

We can see that the shortest combination we could use for Disable would be 'disa'

We can access detailed informational and debug output in Privileged Exec mode.

Check to see all commands that begin with 'sh'

```
Router#sh?  
Show
```

'show' is the only command that begins with 'sh' so we can use that as the abbreviation.

Enter 'sh ?' to see all available show commands. Notice that we have now included a space before the question mark. This enters context sensitive help for the 'show' command.

```
Router#sh ?
aaa Show AAA values
access-lists List access lists
arp Arp table
cdp CDP information
class-map Show QoS Class Map
clock Display the system clock
controllers Interface controllers status
crypto Encryption module
debugging State of each debugging option
dhcp Dynamic Host Configuration Protocol status
dot11 IEEE 802.11 show information
file Show filesystem information
flash: display information about flash: file system
flow Flow information
frame-relay Frame-Relay information
history Display the session command history
hosts IP domain-name, lookup style, nameservers, and host
table
interfaces Interface status and configuration
ip IP information
ipv6 IPv6 information
license Show license information
line TTY line information
--More--
```

Press the Enter key when you see '—More—' to cycle through the additional output one line at a time.

```
- Output truncated -
history Display the session command history
hosts IP domain-name, lookup style, nameservers, and host
table
interfaces Interface status and configuration
ip IP information
ipv6 IPv6 information
license Show license information
line TTY line information
lldp LLDP information
--More--
```

One line at a time is a very slow way to view additional output so press the Space Bar to cycle through it one page at a time instead.

```
- Output truncated -  
sessions Information about Telnet connections  
snmp snmp statistics  
spanning-tree Spanning tree topology  
ssh Status of SSH server connections  
standby standby configuration  
startup-config Contents of startup configuration  
storm-control Show storm control configuration  
tcp Status of TCP connections  
tech-support Show system information for Tech-Support  
--More--
```

Keep hitting the Space Bar until you reach the end of the output.

```
- Output truncated -  
standby standby configuration  
startup-config Contents of startup configuration  
storm-control Show storm control configuration  
tcp Status of TCP connections  
tech-support Show system information for Tech-Support  
terminal Display terminal configuration parameters  
users Display information about terminal lines  
version System hardware and software status  
vlan-switch VTP VLAN status  
vtp Configure VLAN database  
Router#sh
```

Check the possible options for the 'show aaa' command. (We're using aaa for illustrative purposes here. Don't worry about the meaning of the individual aaa commands, they're not important for this exercise.)

```
Router#sh aaa ?  
local Show AAA local method options  
sessions Show AAA sessions as seen by AAA Session MIB  
user Show users active in AAA subsystem
```

Context sensitive help can be very useful if you're not sure about the exact command you need to use. Unfortunately its use may be disabled in the simulator questions on the CCNA exam so you'll need to actually know the commands.

Enter 'sh aaa us' and then hit the Tab key to see Tab completion in action.

```
Router#sh aaa user
```

The Tab key will complete a partially entered command for you. Again this will only work if you've entered enough letters to be a unique match.

Enter the command 'sh aaa user all'.

```
Router#sh aaa user all
Router#
```

Notice that you do not get any output when you enter the command. This is not an error - AAA has not been configured. The CLI simply returns to the Enable prompt because there is nothing to show.

Enter the command 'sh aaa usor all'

```
Router#sh aaa usor all
^
% Invalid input detected at '^' marker.
```

If you enter an illegal command you will get an error message.

Here we made a typo. The CLI warns us that invalid input was detected and shows us the location of the typo is at the 'o' of usor. We typed usor instead of user.

Enter the command 'sh aaa' and hit Enter.

```
Router#sh aaa
% Incomplete command.
```

The router warns us that we've entered an incomplete command, we need to enter additional input. We could enter 'sh aaa ?' again to see the available options.

Explore Global Configuration Mode

Enter Global Configuration mode. (The command can be abbreviated to 'conf t'.)

```
Router#configure terminal  
Enter configuration commands, one per line. End with  
CNTL/Z.  
Router(config)#
```

Notice that the prompt changes to 'Router(config)#'

Global Configuration mode is where we can enter configuration which affects the device as a whole (as opposed to configuring a particular interface for example).

Add a couple of host entries. (Don't worry what this command does for now, we're going to use it to illustrate command history in a second.)

```
Router(config)#ip host Server1 1.1.1.1  
Router(config)#ip host Server2 2.2.2.2
```

Attempt to change the hostname of the device to R1 by entering the command 'R1'.

```
Router(config)#R1  
^  
% Invalid input detected at '^' marker.
```

Oops we forgot to enter the 'hostname' keyword at the start of the command.

Hit the Up Arrow on your keyboard to cycle back to the previous command.

```
Router(config)#R1
```

Enter Ctrl-A to bring the cursor to the beginning of the line and change the entry to 'hostname R1'. This is quicker than typing the command again.
(We can also use 'Ctrl-E' to bring the cursor to the end of the line, and the left and right arrows to move the cursor one character at a time.)

```
Router(config)#hostname R1  
R1(config)#
```

Notice that the command prompt changes to show the router's hostname.

Hit the Up Arrow repeatedly to cycle back through your previous command history, and then the Down Arrow to cycle back again. Notice that command

history is specific to your current level in the command hierarchy - only the commands you previously entered in Global Configuration mode are shown.

Enter the command ‘show ip interface brief’ to check which interfaces are available in the router.

```
R1(config)#show ip interface brief
^
% Invalid input detected at '^' marker.
```

You receive the ‘invalid input detected’ error message but we haven’t made a typo. We’re getting the error because you have to be at the correct level whenever you enter a command. We’re in Global Configuration mode but ‘show’ commands are run in Privileged Exec mode.

We can override this for ‘show’ commands by entering ‘do’ at the start of the command. This works from any level in the command hierarchy.

Enter the correct command to check what interfaces are available from Global Configuration mode.

```
R1(config)#do show ip interface brief
Interface IP-Address OK? Method Status Protocol
GigabitEthernet0/0 unassigned YES NVRAM administratively down down
GigabitEthernet0/1 unassigned YES NVRAM administratively down down
GigabitEthernet0/2 unassigned YES NVRAM administratively down down
Vlan1 unassigned YES NVRAM administratively down down
```

Enter Interface Configuration mode for one of your interfaces.

```
R1(config)#interface gigabitEthernet 0/0
R1(config-if)#
```

Notice that the prompt changes to indicate you are in Interface Configuration mode.

Drop back down to Global Configuration mode.

```
R1(config-if)#exit
R1(config)#
```

The ‘exit’ command drops back down one level.

Hit the Up Arrow and go back to Interface Configuration mode.

```
R1(config)#interface gigabitEthernet 0/0
R1(config-if)#
```

Drop all the way back down to Privilege Exec mode with a single command.

```
R1(config-if)#end  
R1#
```

The 'end' command drops back down to Privilege Exec mode from any level. You can also achieve this by entering 'Ctrl-C'

View the entire device configuration.

```
R1#show running-config  
Building configuration...  
  
Current configuration : 737 bytes  
!  
version 15.1  
no service timestamps log datetime msec  
no service timestamps debug datetime msec  
no service password-encryption  
!  
hostname R1  
!  
Output truncated -
```

View the entire configuration, starting from the hostname.

```
R1#sh run | begin hostname  
hostname R1  
!  
!  
!  
!  
!  
!  
!  
!  
!  
no ip cef  
no ipv6 cef  
!  
Output truncated -
```

Note that the IOS command line is not case sensitive, except when we pipe commands.

```
R1#sh run | begin Hostname  
R1#
```

Here we entered 'Hostname' with a capital letter at the start, but this is not how it is shown in the configuration. The router could find no instance of 'Hostname' so it returns no output.

View configuration lines which include the word 'interface'.

```
R1#show run | include interface
interface GigabitEthernet0/0
interface GigabitEthernet0/1
interface GigabitEthernet0/2
interface Vlan1
```

View all configuration lines which do not include the word 'interface'.

```
R1#show run | exclude interface
Building configuration...

Current configuration : 737 bytes
!
version 15.1
no service timestamps log datetime msec
no service timestamps debug datetime msec
no service password-encryption
!
hostname R1
!
Output truncated -
```

IOS Configuration Management

Copy the running configuration to the startup configuration.

```
R1#copy run start
Destination filename [startup-config]?
Building configuration...
[OK]
```

Change the hostname of the router to RouterX

```
R1#config t
Enter configuration commands, one per line. End with
CNTL/Z.
R1(config)#hostname RouterX
RouterX(config)#
```

Notice that when you enter a command in IOS it takes effect immediately, we can see the command prompt changes to show the new hostname.

Check what hostname will be used when the system reboots.

```
RouterX(config)#do show startup-config
Using 737 bytes
!
version 15.1
no service timestamps log datetime msec
no service timestamps debug datetime msec
no service password-encryption
!
hostname R1
!
```

Commands take effect immediately but are not persistent across a reboot until we save them.

Save the current running configuration so it will be applied next time the router is reloaded.

```
RouterX#copy run start
Destination filename [startup-config]?
Building configuration...
[OK]
```

Verify the new hostname will be applied following a reboot.

```
RouterX#show start
Using 742 bytes
!
version 15.1
no service timestamps log datetime msec
no service timestamps debug datetime msec
no service password-encryption
!
hostname RouterX
!
```

Backup the current running configuration to flash memory in the router.

```
RouterX#copy run flash:
Destination filename [running-config]? config-backup
Building configuration...
[OK]
```

It's not a good idea to back up a device to the device itself, so enter the command to back the running configuration up to an external TFTP server.

```
RouterX#copy run tftp
Address or name of remote host []? 10.10.10.10
Destination filename [RouterX-config]?

Writing running-config.....
%Error opening tftp://10.10.10.10/RouterX-config (Timed out)
```

(The command will try to run for a while and then time out and fail because we didn't set up connectivity to a TFTP server in the lab.)

Reload the device and check it comes back up with the expected configuration with hostname RouterX.

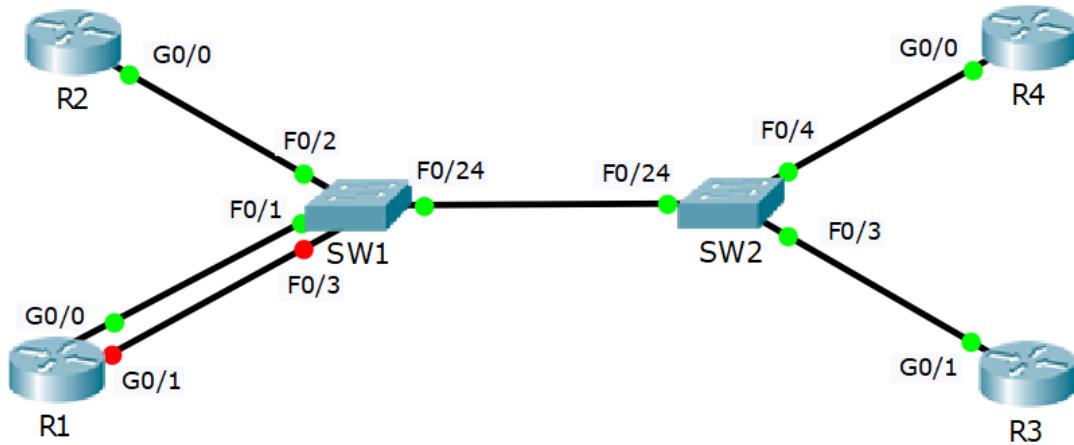
```
RouterX#reload
Proceed with reload? [confirm]
```

11 Cisco Device Functions – Lab Exercise

This lab explores the MAC address table on Cisco IOS switches and routing table on Cisco IOS routers.

This lab is a guided walkthrough of Cisco device functions. You will explore the commands used here in much more detail as you go through the rest of the course.

Lab Topology



Load the Startup Configurations

[Download the '11 Cisco Device Functions.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

This preconfigures each router with an IP address in the 10.10.10.0/24 network.

Verify the Switch MAC Address Tables

- 1) Log into routers R1 to R4 and verify which interface is configured on the 10.10.10.0/24 network.

```
R1#show ip interface brief
Interface IP-Address OK? Method Status Protocol
GigabitEthernet0/0 10.10.10.1 YES manual up up
GigabitEthernet0/1 unassigned YES unset administratively down down
GigabitEthernet0/2 unassigned YES unset administratively down down
Vlan1 unassigned YES unset administratively down down
```

```
R2#show ip interface brief
Interface IP-Address OK? Method Status Protocol
GigabitEthernet0/0 10.10.10.2 YES manual up up
GigabitEthernet0/1 unassigned YES unset administratively down down
GigabitEthernet0/2 unassigned YES unset administratively down down
Vlan1 unassigned YES unset administratively down down
```

```
R3#show ip interface brief
Interface IP-Address OK? Method Status Protocol
GigabitEthernet0/0 unassigned YES unset administratively down down
GigabitEthernet0/1 10.10.10.3 YES manual up up
GigabitEthernet0/2 unassigned YES unset administratively down down
Vlan1 unassigned YES unset administratively down down
```

```
R4#show ip interface brief
Interface IP-Address OK? Method Status Protocol
GigabitEthernet0/0 10.10.10.4 YES manual up up
GigabitEthernet0/1 unassigned YES unset administratively down down
GigabitEthernet0/2 unassigned YES unset administratively down down
Vlan1 unassigned YES unset administratively down down
```

R1, R2 and R4 are using GigabitEthernet0/0, R3 is using GigabitEthernet0/1.

- 2) Note down the MAC addresses of these interfaces.

```
R1#show interface gig0/0
GigabitEthernet0/0 is up, line protocol is up (connected)
Hardware is CN Gigabit Ethernet, address is 0090.2b82.ab01
(bia 0090.2b82.ab01)
```

```
R2#show interface gig0/0
GigabitEthernet0/0 is up, line protocol is up (connected)
Hardware is CN Gigabit Ethernet, address is 0060.2fb3.9152
(bia 0060.2fb3.9152)
```

```
R3#show interface gig0/1
GigabitEthernet0/1 is up, line protocol is up (connected)
Hardware is CN Gigabit Ethernet, address is 0001.9626.8970
(bia 0001.9626.8970)
```

```
R4#show interface gig0/0
GigabitEthernet0/0 is up, line protocol is up (connected)
Hardware is CN Gigabit Ethernet, address is 00d0.9701.02a9
(bia 00d0.9701.02a9)
```

Note: the MAC addresses in your lab may be different.

- 3) Verify connectivity between the routers by pinging R2, R3 and R4 from R1.

```
R1#ping 10.10.10.2
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 0/0/3 ms
```

```
R1#ping 10.10.10.3
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 0/0/1 ms
```

```
R1#ping 10.10.10.4
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 0/0/1 ms
```

- 4) Ping R3 and R4 from R2.

```
R2#ping 10.10.10.3
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 0/0/1 ms
```

```
R2#ping 10.10.10.4
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 0/0/1 ms
```

5) View the dynamically learned MAC addresses on SW1 and verify that the router's MAC addresses are reachable via the expected ports. Ignore any other MAC addresses in the table.

```
SW1#show mac address-table dynamic
Mac Address Table
-----
Vlan Mac Address Type Ports
-----
1 0001.9626.8970 DYNAMIC Fa0/24
1 000c.cf84.8418 DYNAMIC Fa0/24
1 0060.2fb3.9152 DYNAMIC Fa0/2
1 0090.2b82.ab01 DYNAMIC Fa0/1
1 00d0.9701.02a9 DYNAMIC Fa0/24
```

6) Repeat on SW2.

```
SW2#show mac address-table dynamic
Mac Address Table
-----
Vlan Mac Address Type Ports
-----
1 0001.9626.8970 DYNAMIC Fa0/3
1 000b.be53.6418 DYNAMIC Fa0/24
1 0060.2fb3.9152 DYNAMIC Fa0/24
1 0090.2b82.ab01 DYNAMIC Fa0/24
1 00d0.9701.02a9 DYNAMIC Fa0/4
```

7) Clear the dynamic MAC Address Table on SW1.

```
SW1#clear mac address-table dynamic
```

8) Show the dynamic MAC Address Table on SW1. Do you see any MAC addresses? Why or why not?

```
SW1#show mac address-table dynamic
Mac Address Table
-----
Vlan Mac Address Type Ports
-----
1 0001.9626.8970 DYNAMIC Fa0/24
1 000c.cf84.8418 DYNAMIC Fa0/24
1 0060.2fb3.9152 DYNAMIC Fa0/2
1 0090.2b82.ab01 DYNAMIC Fa0/1
1 00d0.9701.02a9 DYNAMIC Fa0/24
```

Devices in a real world network tend to be chatty and send traffic frequently, this causes the MAC address table to update (you may see less entries in Packet Tracer).

The switch will periodically flush old entries.

Examine a Routing Table

- 1) View the routing table on R1. What routes are present and why?

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.10.10.0/24 is directly connected, GigabitEthernet0/0
L 10.10.10.1/32 is directly connected, GigabitEthernet0/0
```

The router has a connected route for the 10.10.10.0/24 network and a local route for 10.10.10.1/32. These routes were automatically created when the IP address 10.10.10.1/24 was configured on interface GigabitEthernet0/0

- 2) Configure IP address 10.10.20.1/24 on interface GigabitEthernet0/1

```
R1(config)#interface GigabitEthernet 0/1
R1(config-if)#ip address 10.10.20.1 255.255.255.0
R1(config-if)#no shutdown
```

- 3) What routes are in the routing table now?

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C 10.10.10.0/24 is directly connected, GigabitEthernet0/0
L 10.10.10.1/32 is directly connected, GigabitEthernet0/0
C 10.10.20.0/24 is directly connected, GigabitEthernet0/1
L 10.10.20.1/32 is directly connected, GigabitEthernet0/1
```

The router has routes for both interfaces and can route traffic between hosts on the 10.10.10.0/24 and 10.10.20.0/24 networks.

4) Configure a static route to 10.10.30.0/24 with a next hop address of 10.10.10.2

```
R1(config)#ip route 10.10.30.0 255.255.255.0 10.10.10.2
```

5) What routes are in the routing table now?

```
R1(config)#do show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route
```

Gateway of last resort is not set

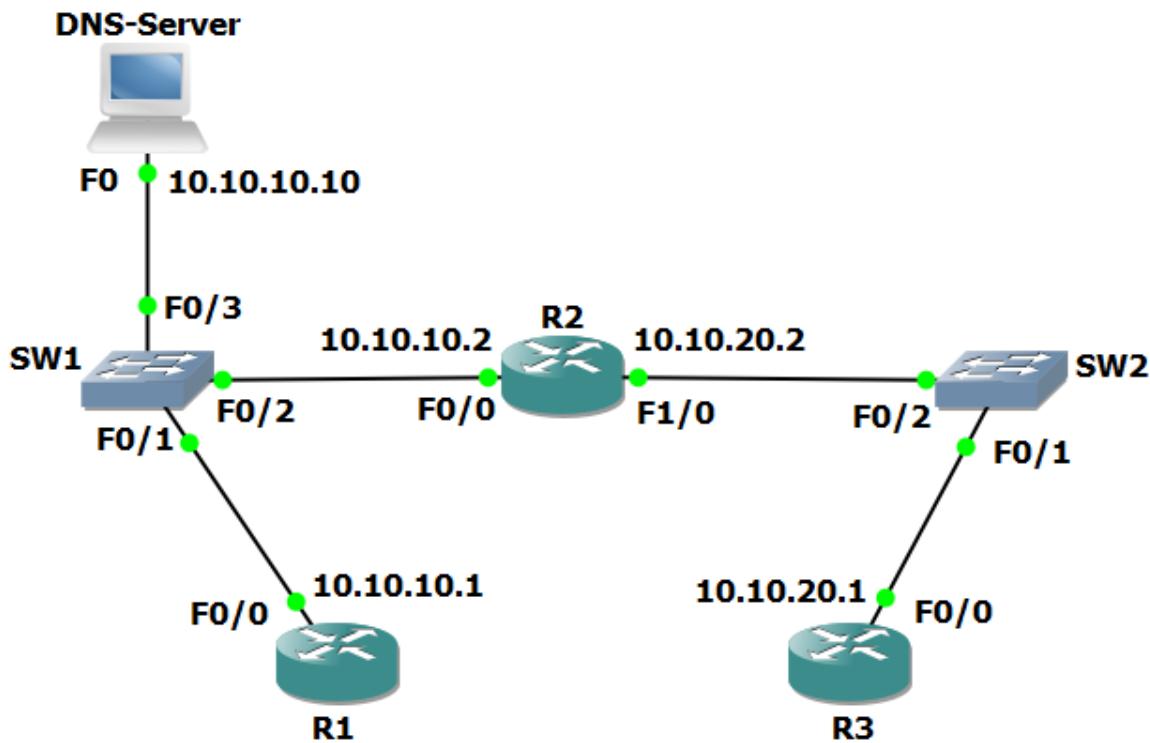
```
10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C 10.10.10.0/24 is directly connected, GigabitEthernet0/0
L 10.10.10.1/32 is directly connected, GigabitEthernet0/0
C 10.10.20.0/24 is directly connected, GigabitEthernet0/1
L 10.10.20.1/32 is directly connected, GigabitEthernet0/1
S 10.10.30.0/24 [1/0] via 10.10.10.2
```

The router has routes to its locally connected networks, and also to 10.10.30.0/24 which is available via 10.10.10.2

12 The Life of a Packet - Lab Exercise

This lab explores DNS configuration on Cisco routers and the ARP cache.

Lab Topology



Load the Startup Configurations

[Download the '12 The Life of a Packet.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

This configures the lab topology as shown above and adds static routes between R1 and R3.

You can learn the theory for this section and shortcut your path to CCNA certification by getting my CCNA Gold Bootcamp course:
<https://www.flackbox.com/cisco-ccna-training-course>

Configure the Routers as DNS Clients

Note that routers cannot be DNS servers in Packet Tracer (it does not support the 'ip dns server' command) so we are using a Packet Tracer server device as the DNS server.

The host with IP address 10.10.10.10 has been configured as a DNS server and is able to resolve DNS requests for 'R1', 'R2' and 'R3'.

A domain name is not in use.

- 1) Configure R1, R2 and R3 to use 10.10.10.10 as their DNS server. You do not need to configure a domain-name or domain-list.
- 2) Verify that you can ping R2 and R3 from R1 using their hostnames 'R1' and 'R3' (it may take some time for the DNS server to resolve the DNS request).
- 3) Verify that you can ping R1 and R2 from R3 using their hostnames 'R1' and 'R2'.

Examine the ARP Cache on the Routers

- 4) Do you expect to see an entry for R3 in the ARP cache of R1? Why or why not?
- 5) Verify the ARP cache on R1, R2 and R3. What do you see?

12 The Life of a Packet – Answer Key

This lab explores DNS client configuration on Cisco routers and the ARP cache.

Configure the Routers as DNS Clients

Note that routers cannot be DNS servers in Packet Tracer (it does not support the 'ip dns server' command) so we are using a Packet Tracer server device as the DNS server.

The host with IP address 10.10.10.10 has been configured as a DNS server and is able to resolve DNS requests for 'R1', 'R2' and 'R3'.

A domain name is not in use.

- 1) Configure R1, R2 and R3 to use 10.10.10.10 as their DNS server. You do not need to configure a domain-name or domain-list.

```
R1(config)#ip domain-lookup
R1(config)#ip name-server 10.10.10.10
```

```
R2(config)#ip domain-lookup
R2(config)#ip name-server 10.10.10.10
```

```
R3(config)#ip domain-lookup
R3(config)#ip name-server 10.10.10.10
```

- 2) Verify that you can ping R2 and R3 from R1 using their hostnames 'R1' and 'R3' (it may take some time for the DNS server to resolve the DNS request).

```
R1#ping R2
Translating "R2"...domain server (10.10.10.10)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.2, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
0/0/1 ms
```

```
R1#ping R3
```

```
Translating "R3"...domain server (10.10.10.10)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.20.1, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
0/0/1 ms
```

- 3) Verify that you can ping R1 and R2 from R3 using their hostnames 'R1' and 'R2'.

```
R3#ping R1
Translating "R1"...domain server (10.10.10.10)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.1, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
0/0/2 ms
```

```
R3#ping R2
Translating "R2"...domain server (10.10.10.10)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.2, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
0/0/2 ms
```

Examine the ARP Cache on the Routers

4) Do you expect to see an entry for R3 in the ARP cache of R1? Why or why not?

ARP requests use broadcast traffic so they are not forwarded by a router. R1 will have entries in its ARP cache for all hosts it has seen ***on its directly connected networks*** (10.10.10.0/24).

R1 is not directly connected to the 10.10.20.0/24 network so it will not have an entry in the ARP cache for R3 at 10.10.20.1.

R1 can reach R3 via R2's IP address 10.10.10.2 – this IP address is included in the ARP cache.

The DNS server at 10.10.10.10 is also in the same IP subnet as R1 so will also appear in the ARP cache.

5) Verify the ARP cache on R1, R2 and R3.

```
R1#show arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 10.10.10.1 - 0090.0CD7.0D01 ARPA FastEthernet0/0
Internet 10.10.10.2 4 0004.9A96.A9A5 ARPA FastEthernet0/0
Internet 10.10.10.10 2 0090.21C6.D284 ARPA FastEthernet0/0
```

```
R2#show arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 10.10.10.1 4 0090.0CD7.0D01 ARPA FastEthernet0/0
Internet 10.10.10.2 - 0004.9A96.A9A5 ARPA FastEthernet0/0
Internet 10.10.10.10 1 0090.21C6.D284 ARPA FastEthernet0/0
Internet 10.10.20.1 4 0030.F2BA.30E7 ARPA FastEthernet1/0
Internet 10.10.20.2 - 0060.2FCA.ACA0 ARPA FastEthernet1/0
```

```
R3#show arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 10.10.20.1 - 0030.F2BA.30E7 ARPA FastEthernet0/0
Internet 10.10.20.2 4 0060.2FCA.ACA0 ARPA FastEthernet0/0
```

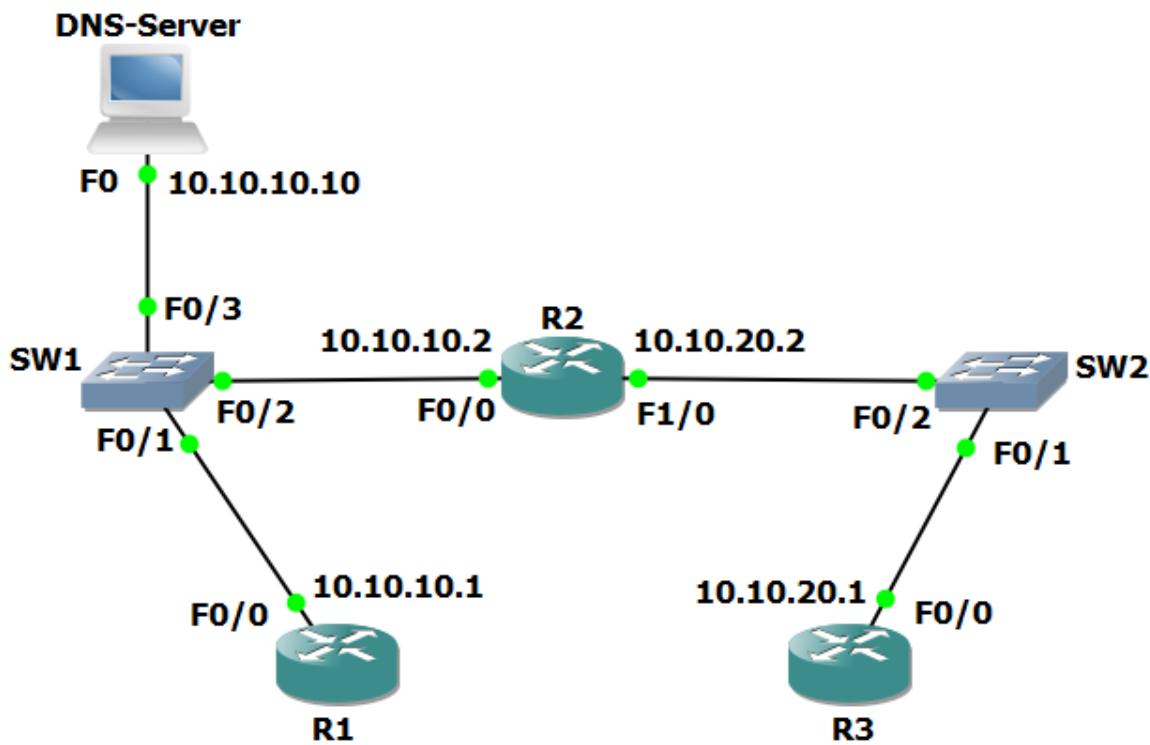
R2 is directly connected to 10.10.10.0/24 and 10.10.20.0/24 so it has entries in its ARP cache for both networks.

R3 is directly connected to the 10.10.20.0/24 network so it has entries in its ARP cache for that network only. It does not have any entries for the 10.10.10.0/24 network.

13 The Cisco Troubleshooting Methodology - Lab Exercise

This lab tests your network connectivity troubleshooting skills.

Lab Topology



Load the Startup Configurations

[Download the '13 The Cisco Troubleshooting Methodology.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

This configures the lab topology as shown above with **R3** as a DNS server and adds static routes between **R1** and **R3**.

Troubleshoot Connectivity to DNS Server

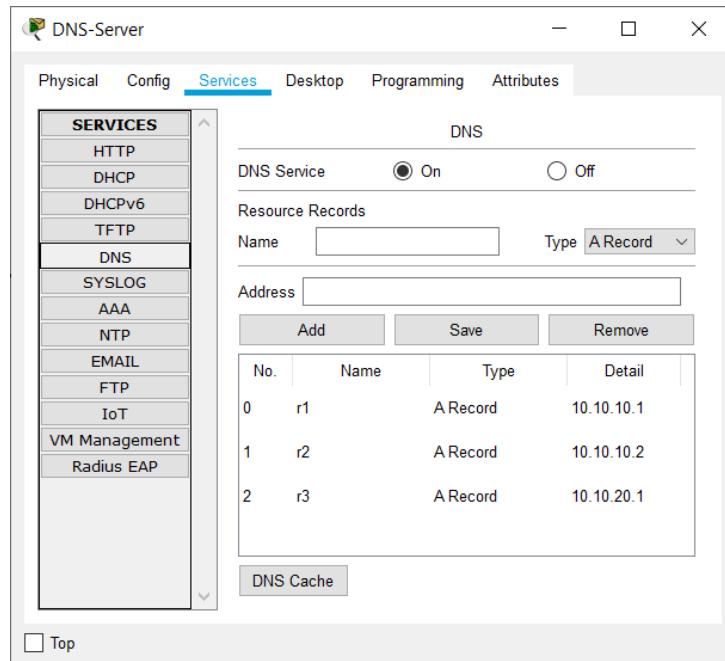
Note that routers cannot be DNS servers in Packet Tracer (it does not support the 'ip dns server' command) so we are using a Packet Tracer server device as the DNS server.

- 1) The host with IP address 10.10.10.10 has been configured as a DNS server and should be able to resolve requests for 'R1', 'R2' and 'R3'. Members of staff have complained that DNS is not working.
- 2) From R3, use Telnet to check if the DNS service appears operational on the DNS server at 10.10.10.10.

```
R3#telnet 10.10.10.10
Trying 10.10.10.10 ...
% Connection timed out; remote host not responding
```

- 3) When you have verified that DNS is not working, troubleshoot and fix the problem. You have fixed the problem when R3 can ping R1 by hostname. Note that there may be more than one issue causing the problem.

(You can click on the DNS server and then the 'Services' tab to check the server's DNS configuration.)



Hint: you can use the `show ip interface brief` command to verify interfaces are operational on routers and switches. This command will be covered in more detail in the next section.

13 The Cisco Troubleshooting Methodology – Answer Key

This lab tests your network troubleshooting skills.

Troubleshoot Connectivity to DNS Server

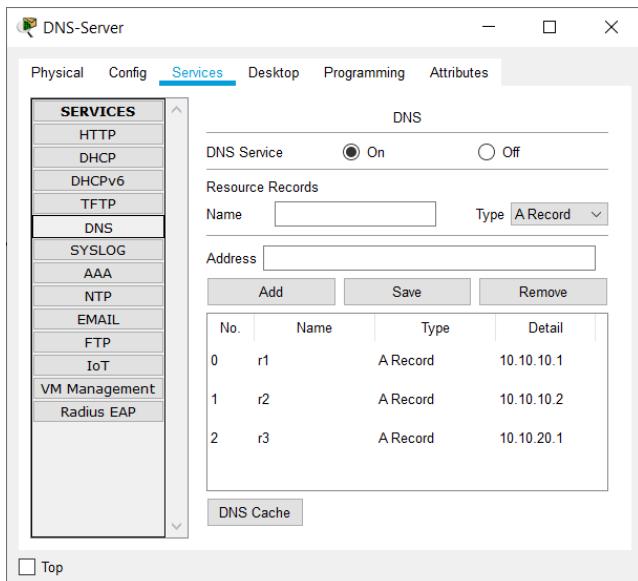
Note that routers cannot be DNS servers in Packet Tracer (it does not support the 'ip dns server' command) so we are using a Packet Tracer server device as the DNS server.

- 9) The host with IP address 10.10.10.10 has been configured as a DNS server and should be able to resolve requests for 'R1', 'R2' and 'R3'. Members of staff have complained that DNS is not working.
- 10) From R3, use Telnet to check if the DNS service appears operational on the DNS server at 10.10.10.10.

```
R3#telnet 10.10.10.10
Trying 10.10.10.10 ...
% Connection timed out; remote host not responding
```

- 11) When you have verified that DNS is not working, troubleshoot and fix the problem. You have fixed the problem when R3 can ping R1 by hostname. Note that there may be more than one issue causing the problem.

(You can click on the DNS server and then the 'Services' tab to check the server's DNS configuration.)



There is more than one way to troubleshoot the issue. A suggested workflow is shown below.

The first two questions to ask when troubleshooting a problem are:

1. Was it working before? If so, has something changed which could cause the problem? This will usually direct you to the cause.

This question is not particularly useful for our example as the DNS server has just been brought online for the first time.

2. Is the problem affecting everybody or just one particular user? If it's affecting just one user, the likelihood is that the problem is at their end.

In this case the problem is affecting all users, so the problem is likely on the server end or with the network.

The error message when we tried to Telnet was 'remote host not responding', so it looks like a connectivity issue.

Ping from R3 to the DNS server.

```
R3#ping 10.10.10.10
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 10.10.10.10, timeout is 2  
seconds:  
U.U.U  
Success rate is 0 percent (0/5)
```

The ping fails at the network layer so there is little point in checking the DNS service at higher layers until we fix this problem.

Rather than checking connectivity hop by hop, we can possibly save a little time by using traceroute.

```
R3#traceroute 10.10.10.10  
Type escape sequence to abort.  
Tracing the route to 10.10.10.10  
  
1 10.10.20.2 0 msec 0 msec 0 msec  
2 10.10.20.2 !H * !H  
3 * *
```

The traceroute got as far as R2, which lets us know that R3 has the correct route to get to the DNS server, and the problem is probably between R2 and the DNS server.

R2 has an interface connected to the 10.10.10.0/24 network, so we don't need to check it has a route to the DNS server. We do need to check that the interface is up though.

```
R2#sh ip int brief
Interface          IP-Address      OK? Method Status      Protocol
FastEthernet0/0    10.10.10.2    YES  NVRAM  administratively down down
FastEthernet0/1    unassigned      YES  NVRAM  administratively down down
FastEthernet1/0    10.10.20.2    YES  NVRAM  up        up
FastEthernet1/1    unassigned      YES  NVRAM  administratively down down
Vlan1             unassigned      YES  NVRAM  administratively down down
```

There's the problem – FastEthernet0/0 facing the DNS server is administratively shutdown. Let's fix it.

```
R2(config)#interface f0/0
R2(config-if)#no shutdown
```

Next we'll try pinging from R3 to the DNS server again to verify we fixed connectivity.

```
R3#ping 10.10.10.10
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.10, timeout is 2
seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max =
0/0/0 ms
```

That looks better – don't worry if the first one or two pings are dropped, there may be a delay while the ARP cache is updated. Next we'll verify DNS is working.

```
R3#ping R1
Translating "R1"...domain server (10.10.10.1)
% Unrecognized host or address or protocol not running.
```

The error message tells us the problem if we take the time to really read it – R3 is using 10.10.10.1 as its DNS server, but the correct address is 10.10.10.10.

We fix that next. Don't forget to remove the incorrect entry first.

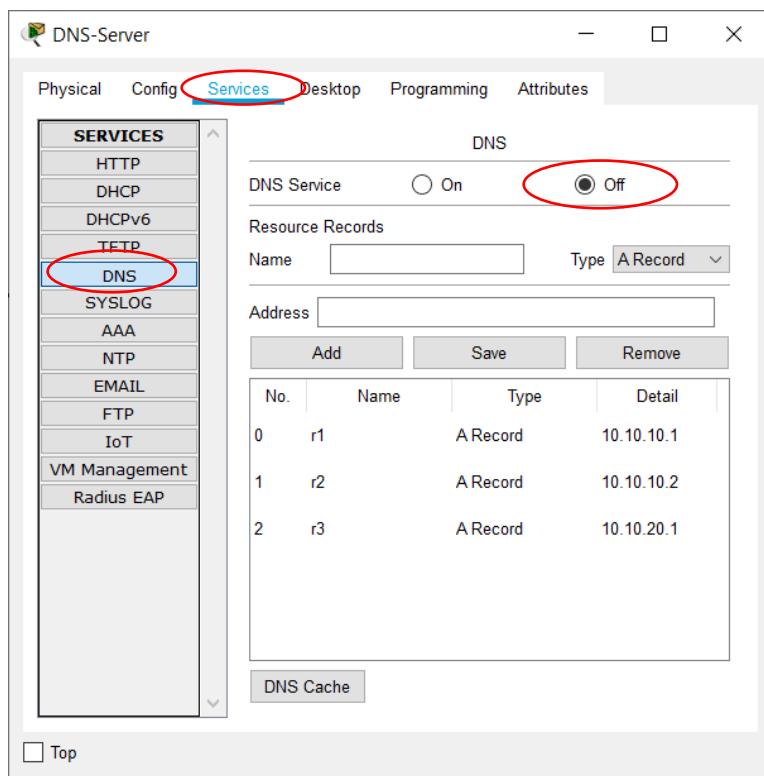
```
R3(config)#no ip name-server 10.10.10.1
R3(config)#ip name-server 10.10.10.10
```

Then test again.

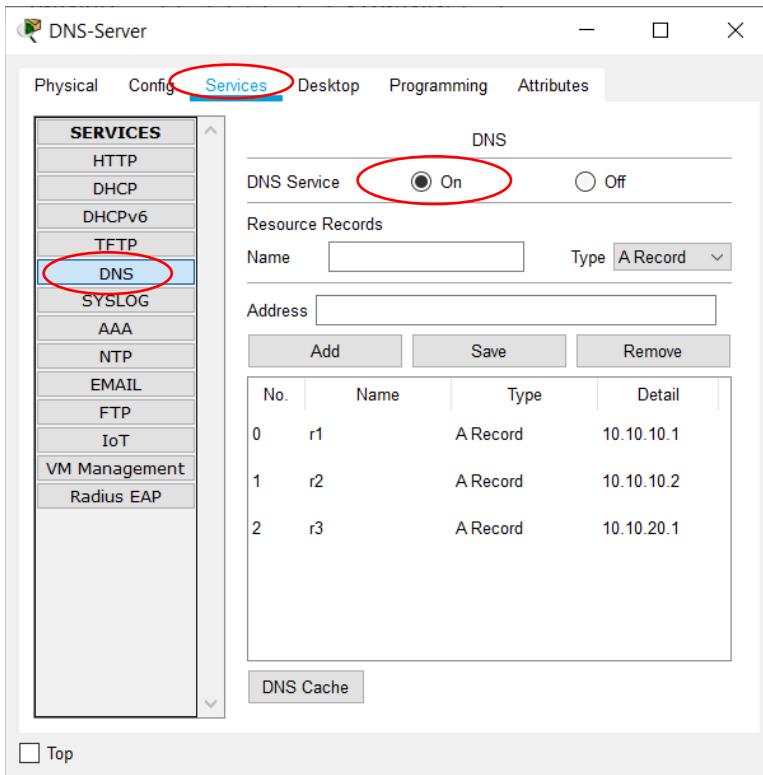
```
R3#ping R1
Translating "R1"...
domain server (10.10.10.1)
% Unrecognized host or address or protocol not running.
```

The error message is still there. We know we have connectivity and the DNS server configured correctly on R3, so the problem looks like it's on the DNS server.

Check the DNS service is running on the 10.10.10.10 host and that address records are configured for 'R1', 'R2' and 'R3'.



The address records are there but the DNS service is turned off. Turn it back on.



Time to test it from R3 again.

```
R3#ping R1
Translating "R1"...domain server (10.10.10.10)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.1, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
0/0/4 ms
```

That's the problem solved.

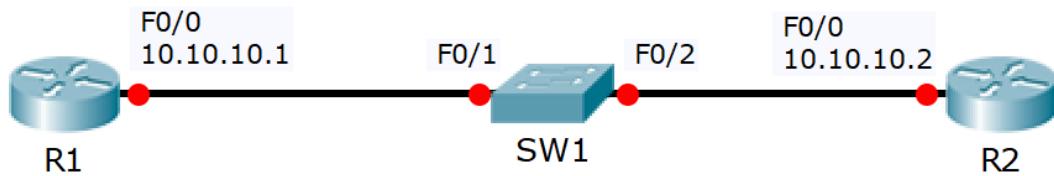
To summarise the issues: port FastEthernet0/0 was shut down on R2, R3 was using the wrong IP address for the DNS server, and the DNS service was not running on the server.

Problems in the real world are usually caused by just one error rather than three as in this case. This can still occur though, particularly when a new service is being deployed.

14 Cisco Router and Switch Basics - Lab Exercise

In this lab you will complete a basic configuration on a switch, verify Cisco Discovery Protocol CDP and analyse the effects of interface speed and duplex configuration.

Lab Topology



Load the Startup Configurations

[Download the '14 Cisco Router and Switch Basics.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Cisco Router and Switch Initial Configuration

- 1) Configure Router 1 with the hostname 'R1'
- 2) Configure Router 2 with the hostname 'R2'
- 3) Configure Switch 1 with the hostname 'SW1'
- 4) Configure the IP address on R1 according to the topology diagram
- 5) Configure the IP address on R2 according to the topology diagram
- 6) Give SW1 the management IP address 10.10.10.10/24
- 7) The switch should have connectivity to other IP subnets via R2
- 8) Verify the switch can ping its default gateway
- 9) Enter suitable descriptions on the interfaces connecting the devices
- 10) On SW1, verify that speed and duplex are automatically negotiated to 100 Mbps full duplex on the link to R1
- 11) Manually configure full duplex and FastEthernet speed on the link to R2
- 12) What version of IOS is the switch running?

CDP Configuration

- 13) Verify the directly attached Cisco neighbors using Cisco Discovery Protocol
- 14) Prevent R1 from discovering information about Switch 1 via CDP
- 15) Flush the CDP cache on R1 by entering the 'no cdp run' then 'cdp run' commands in global configuration mode
- 16) Verify that R1 cannot see SW1 via CDP

Switch Troubleshooting

- 17) Verify the status of the switch port connected to R2 with the `show ip interface brief` command. It should show status and protocol up/up.
- 18) Shut down the interface connected to R2 and issue a `show ip interface brief` command again. The status and protocol should show administratively down/down.
- 19) Bring the interface up again. Verify the speed and duplex setting.
- 20) Set the duplex to half on Switch 1. Leave the settings as they are on R2.
- 21) Verify the state of the interface.
- 22) Set the duplex back to full duplex.
- 23) Set the speed to 10 Mbps.
- 24) Check if the interface is still operational.
- 25) Check if the interface is operational on R2. What is the status of the interface?

14 Cisco Router and Switch Basics - Answer Key

In this lab you will complete a basic configuration on a switch, verify Cisco Discovery Protocol CDP and analyse the effects of interface speed and duplex configuration.

Cisco Router and Switch Initial Configuration

- 1) Configure Router 1 with the hostname 'R1'

```
Router(config)#hostname R1
R1(config)#
```

- 2) Configure Router 2 with the hostname 'R2'

```
Router(config)#hostname R2
R2(config)#
```

- 3) Configure Switch 1 with the hostname 'SW1'

```
Switch(config)#hostname SW1
SW1(config)#
```

- 4) Configure the IP address on R1 according to the topology diagram

```
R1(config)#interface FastEthernet0/0
R1(config-if)#ip address 10.10.10.1 255.255.255.0
R1(config-if)#no shutdown
```

- 5) Configure the IP address on R2 according to the topology diagram

```
R2(config)#interface FastEthernet0/0
R2(config-if)#ip address 10.10.10.2 255.255.255.0
R2(config-if)#no shutdown
```

- 6) Give SW1 the management IP address 10.10.10.10/24

```
SW1(config)#interface vlan1
SW1(config-if)#ip address 10.10.10.10 255.255.255.0
SW1(config-if)#no shutdown
```

7) The switch should have connectivity to other IP subnets via R2

```
SW1(config)#ip default-gateway 10.10.10.2
```

8) Verify the switch can ping its default gateway

```
SW1#ping 10.10.10.2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.10.10.2, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms

9) Enter suitable descriptions on the interfaces connecting the devices

```
R1(config)#interface FastEthernet 0/0
R1(config-if)#description Link to SW1
```

```
R2(config-if)#interface FastEthernet 0/0
R2(config-if)#description Link to SW1
```

```
SW1(config)#interface FastEthernet 0/1
SW1(config-if)#description Link to R1
SW1(config-if)#interface FastEthernet 0/2
SW1(config-if)#description Link to R2
```

10) On SW1, verify that speed and duplex are automatically negotiated to 100 Mbps full duplex on the link to R1

```
SW1#show interface f0/1
FastEthernet0/1 is up, line protocol is up (connected)
Hardware is Lance, address is 00e0.8fd6.8901 (bia
00e0.8fd6.8901)
Description: Link to R1
BW 100000 Kbit, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 100Mb/s
```

11) Manually configure full duplex and FastEthernet speed on the link to R2

```
SW1(config)#interface FastEthernet 0/2
SW1(config-if)#speed 100
SW1(config-if)#duplex full
```

Don't forget to configure matching settings on R2!

```
R2(config)#interface FastEthernet 0/0
R2(config-if)#speed 100
R2(config-if)#duplex full
```

12) What version of IOS is the switch running?

```
SW1#show version
Cisco IOS Software, C2960 Software (C2960-LANBASE-M),
Version 12.2(25)FX, RELEASE SOFTWARE (fc1)
```

CDP Configuration

13) Verify the directly attached Cisco neighbors using Cisco Discovery Protocol

```
SW1#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
Device ID Local Intrfce Holdtme Capability Platform Port ID
R1        Fas 0/1        170        R        C2800  Fas 0/0
R2        Fas 0/2        134        R        C2800  Fas 0/0
```

14) Prevent R1 from discovering information about Switch 1 via CDP

```
SW1(config)#interface FastEthernet 0/1
SW1(config-if)#no cdp enable
```

15) Flush the CDP cache on R1 by entering the 'no cdp run' then 'cdp run' commands in global configuration mode

```
R1(config)#no cdp run
R1(config)#cdp run
```

16) Verify that R1 cannot see SW1 via CDP

```
R1#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID      Local Intrfce      Holdtme      Capability      Platform      Port ID
R1#
```

Switch Troubleshooting

17) Verify the status of the switch port connected to R2 with the `show ip interface brief` command. It should show status and protocol up/up.

```
SW1#show ip interface brief
Interface          IP-Address      OK? Method Status
Protocol
Vlan1              10.10.10.10    YES manual up
FastEthernet0/1    unassigned      YES unset  up
FastEthernet0/2    unassigned      YES unset  up
```

18) Shut down the interface connected to R2 and issue a `show ip interface brief` command again. The status and protocol should show administratively down/down.

```
SW1(config)#interface FastEthernet 0/2
SW1(config-if)#shutdown
*Mar  1 00:44:34.212: %LINK-5-CHANGED: Interface
FastEthernet0/2, changed state to administratively down
*Mar  1 00:44:35.219: %LINEPROTO-5-UPDOWN: Line protocol on
Interface FastEthernet0/2, changed state to down
```

```
SW1(config-if)#do show ip interface brief
Interface          IP-Address      OK? Method Status
Protocol
Vlan1              10.10.10.10    YES manual up
FastEthernet0/1    unassigned      YES unset  up
FastEthernet0/2    unassigned      YES unset  administratively down down
```

19) Bring the interface up again. Verify the speed and duplex setting.

```
SW1(config)#interface FastEthernet 0/2
SW1(config-if)#no shutdown
SW1(config-if)#
*Mar  1 00:45:52.637: %LINK-3-UPDOWN: Interface
FastEthernet0/2, changed state to up
*Mar  1 00:45:53.644: %LINEPROTO-5-UPDOWN: Line protocol on
Interface FastEthernet0/2, changed state to up
```

```
SW1#sh interface f0/2
FastEthernet0/2 is up, line protocol is up (connected)
Hardware is Lance, address is 00e0.8fd6.8902 (bia
00e0.8fd6.8902)
BW 100000 Kbit, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 100Mb/s
```

20) Set the duplex to half on Switch 1. Leave the settings as they are on R2.

```
SW1(config-if)#duplex half
SW1(config-if)#
%LINK-5-CHANGED: Interface FastEthernet0/2, changed state
to down

%LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/2, changed state to down
```

21) Verify the state of the interface.

The interface is down/down. It will not forward traffic.

```
SW1#show ip interface brief
Interface      IP-Address OK? Method Status Protocol
FastEthernet0/1 unassigned YES manual up up
FastEthernet0/2 unassigned YES manual down down
```

22) Set the duplex back to full duplex.

```
SW1(config)#int f0/2
SW1(config-if)#duplex full
SW1(config-if)#
%LINK-5-CHANGED: Interface FastEthernet0/2, changed state
to up

%LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/2, changed state to up
```

23) Set the speed to 10 Mbps.

```
SW1(config)#int f0/2
SW1(config-if)#speed 10
SW1(config-if)#
%LINK-5-CHANGED: Interface FastEthernet0/2, changed state
to down

%LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/2, changed state to down
```

24) Check if the interface is still operational.

```
SW1#show ip interface brief
Interface          IP-Address      OK? Method Status
Protocol
Vlan1              10.10.10.10    YES manual up
FastEthernet0/1    unassigned      YES unset up
FastEthernet0/2  unassigned      YES unset down
```

The interface status is down/down.

25) Check if the interface is operational on R2. What is the status of the interface?

```
R2#show ip interface brief
Interface          IP-Address      OK? Method Status
Protocol
FastEthernet0/0    10.10.10.2     YES manual up
down
```

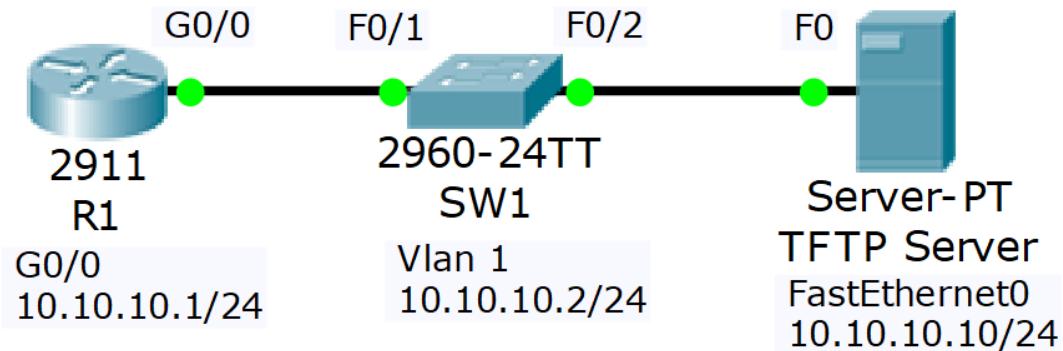
The interface status is up/down on R2.

15 Cisco Device Management - Lab Exercise

In this lab you will perform a factory reset, password recovery, configuration backup, and system image backup and recovery on a Cisco router. You will also perform an IOS upgrade on a Cisco switch.

Use Cisco Packet Tracer for this exercise. The generic server in Packet Tracer (as shown in the topology diagram below) has built-in TFTP server software.

Lab Topology



Load the Startup Configurations

[Download the '15 Cisco Device Management.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Factory Reset

- 12) View the running configuration on R1. Note that the hostname and interface have been configured
- 13) Factory reset R1 and reboot
- 14) Watch the boot up process as the router boots
- 15) The router should boot into the Setup Wizard. Exit out of the wizard and then confirm the startup and running configurations are empty.
- 16) Paste the configuration for R1 from the '15 Cisco Device Management Configs.zip' file back into the configuration and save

Password Recovery

- 17) Set the enable secret 'Flackbox1' on R1
- 18) Configure the router to boot into the rommon prompt on next reload, and reboot the router. In a real world scenario you would enter the Break sequence on the keyboard when first powering up the router to access the rommon prompt
- 19) Configure the router to ignore the startup-config when booting up, and reload the router
- 20) The router should boot into the Setup Wizard. Exit out of the wizard
- 21) What do you expect to see if you view the running and startup configurations? Confirm this.
- 22) Copy the startup config to the running config. Do not miss this step or you will factory reset the router!
- 23) Remove the enable secret
- 24) Ensure the router will reboot normally on the next reload and you will be able to access the router
- 25) Reboot the router to confirm

Configuration Backup

- 26) Backup the running configuration to Flash on R1. Use a suitable name for the backup file. Verify the configuration has been backed up.
- 27) Backup the R1 startup configuration to the TFTP server. Use a suitable name for the backup file. Verify the configuration has been backed up.

IOS System Image Backup and Recovery

- 28) Backup the IOS system image on R1 to the TFTP server. Verify the configuration has been backed up.
- 29) Delete the system image from Flash and reload.
- 30) Use Internet search to find system recovery instructions for your model of router. Recover the system image using the TFTP server.
- 31) If you are using the latest version of Packet Tracer the download will time out because the GigabitEthernet interface stays down in rommon mode. You have completed this part of the lab once you have entered the `tftpdownld` command.

IOS Image Upgrade

- 1) Verify SW1 is running C2960 Software (C2960-LANBASE-M), Version 12.2(25)FX
- 2) Use the TFTP server to upgrade to C2960-lanbasek9-mz.150-2.SE4.bin
- 3) Reboot and verify the switch is running the new software version

15 Cisco Device Management - Answer Key

In this lab you will perform a factory reset, password recovery, configuration backup, and system image backup and recovery on a Cisco router. You will also perform an IOS upgrade on a Cisco switch.

Factory Reset

32) View the running configuration on R1. Note that the hostname and interface have been configured

```
R1#sh run
Building configuration...

Current configuration : 696 bytes
!
hostname R1
!
interface GigabitEthernet0/0
ip address 10.10.10.1 255.255.255.0
duplex auto
speed auto
```

33) Factory reset R1 and reboot

```
R1#write erase
Erasing the nvram filesystem will remove all configuration
files! Continue? [confirm]
[OK]
Erase of nvram: complete
%SYS-7-NV_BLOCK_INIT: Initialized the geometry of nvram
R1#reload
Proceed with reload? [confirm]
```

34) Watch the boot up process as the router boots

```
System Bootstrap, Version 15.1(4)M4, RELEASE SOFTWARE (fc1)

Readonly ROMMON initialized

IOS Image Load Test

Digitally Signed Release Software

Self decompressing the image :
#####
##### [OK]
```

35) The router should boot into the Setup Wizard. Exit out of the wizard and then confirm the startup and running configurations are empty.

```
--- System Configuration Dialog ---

Continue with configuration dialog? [yes/no]: no

Router>enable
Router#show run
Building configuration...
hostname Router
!
interface GigabitEthernet0/0
no ip address
duplex auto
speed auto
shutdown

Router#show start
startup-config is not present
```

36) Paste the configuration for R1 from the '15 Cisco Device Management Configs.zip' file back into the configuration and save

```
Router#configure terminal
Router(config)#hostname R1
R1(config)#
R1(config)#interface GigabitEthernet0/0
R1(config-if)# ip address 10.10.10.1 255.255.255.0
R1(config-if)# duplex auto
R1(config-if)# speed auto
R1(config-if)# no shutdown
R1(config-if)#!
```

```
R1(config-if)#line con 0
R1(config-line)# exec-timeout 30 0
R1(config-line)#end
R1#copy run start
Destination filename [startup-config]?
Building configuration...
[OK]
```

Password Recovery

37) Set the enable secret 'Flackbox1' on R1

```
R1(config)#enable secret Flackbox1
```

38) Configure the router to boot into the rommon prompt on next reload, and reboot the router. (In a real world scenario you would enter the Break sequence on the keyboard when first powering up the router to access the rommon prompt)

```
R1(config)#config-register 0x2120
R1(config)#end
R1#copy run start
R1#reload
Proceed with reload? [confirm]
```

39) Configure the router to ignore the startup-config when booting up, and reload the router

```
rommon 1 > confreg 0x2142
rommon 2 > reset
```

40) The router should boot into the Setup Wizard. Exit out of the wizard

```
--- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: no
```

41) What do you expect to see if you view the running and startup configurations? Confirm this.

The running configuration should be empty because the router bypassed loading the startup config on boot up. The startup config should remain unchanged and all previous configuration should still be there.

```
Router#sh run
Building configuration...

hostname Router
!
interface GigabitEthernet0/0
no ip address
duplex auto
speed auto

Router#sh start
!
hostname R1
!
enable secret 5 $1$mERr$J2XZHMOgpVVXdLjC9lYtE1
!
interface GigabitEthernet0/0
ip address 10.10.10.1 255.255.255.0
duplex auto
speed auto
```

42) Copy the startup config to the running config. Do not miss this step or you will factory reset the router!

```
Router#copy start run
Destination filename [running-config]?
```

43) Remove the enable secret

```
Router(config)#no enable secret
```

44) Ensure the router will reboot normally on the next reload and you will be able to access the router

```
Router(config)#config-register 0x2102
Router(config)#end
Router#copy run start
Destination filename [startup-config]?
Building configuration...
[OK]
```

45) Reboot the router to confirm

```
Router#reload
Proceed with reload? [confirm]

R1>en
R1#sh run
Building configuration...
hostname R1
!
interface GigabitEthernet0/0
ip address 10.10.10.1 255.255.255.0
duplex auto
speed auto
```

Configuration Backup

Important: Filenames are case sensitive – you must enter them *exactly* as shown.
(**c2900** is different to **C2900**)

46) Backup the running configuration to Flash on R1. Use a suitable name for the backup file. Verify the configuration has been backed up.

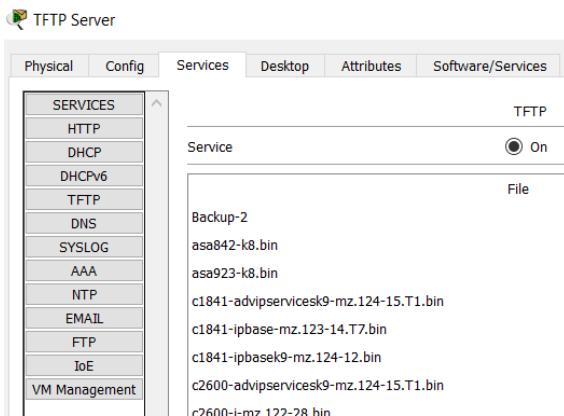
```
R1#copy run flash
Destination filename [running-config]? Backup-1
Building configuration...
[OK]

R1#show flash

System flash directory:
File Length Name/status
5 728 Backup-1
3 33591768 c2900-universalk9-mz.SPA.151-4.M4.bin
2 28282 sigdef-category.xml
1 227537 sigdef-default.xml
[33848315 bytes used, 221895685 available, 255744000 total]
249856K bytes of processor board System flash (Read/Write)
```

47) Backup the R1 startup configuration to the TFTP server. Use a suitable name for the backup file. Verify the configuration has been backed up.

```
R1#copy start tftp
Address or name of remote host [ ]? 10.10.10.10
Destination filename [R1-config]? Backup-2
Writing startup-config....!!
[OK - 728 bytes]
728 bytes copied in 3.007 secs (242 bytes/sec)
```



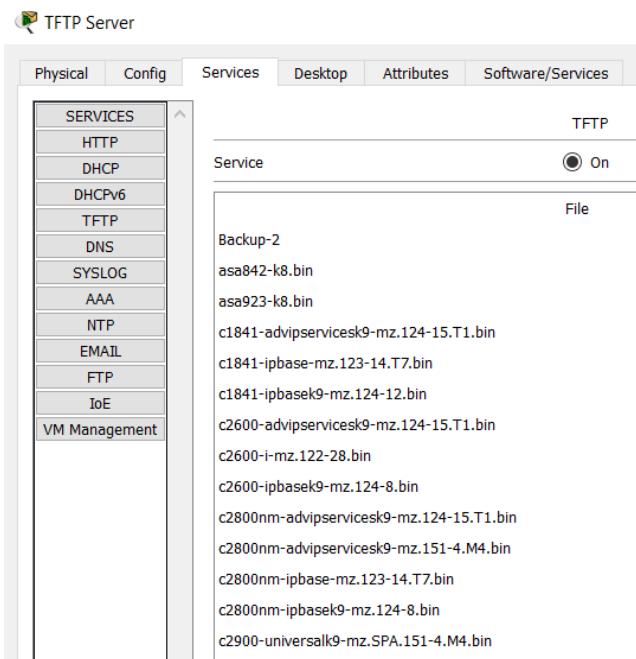
IOS System Image Backup and Recovery

48) Backup the IOS system image on R1 to the TFTP server. Verify the configuration has been backed up.

```
R1#show flash
```

```
System flash directory:  
File Length Name/status  
5 728 Backup-1  
3 33591768 c2900-universalk9-mz.SPA.151-4.M4.bin  
2 28282 sigdef-category.xml  
1 227537 sigdef-default.xml  
[33848315 bytes used, 221895685 available, 255744000 total]  
249856K bytes of processor board System flash (Read/Write)
```

```
R1#copy flash tftp  
Source filename []? c2900-universalk9-mz.SPA.151-4.M4.bin  
Address or name of remote host []? 10.10.10.10  
Destination filename [c2900-universalk9-mz.SPA.151-4.M4.bin]?  
  
Writing c2900-universalk9-mz.SPA.151-4.M4.bin.....!!!!!!!!!!!!!!  
!!!!!!!!!!!!!![OK - 33591768 bytes]
```



49) Delete the system image from Flash and reload.

```
R1#delete flash:c2900-universalk9-mz.SPA.151-4.M4.bin
Delete filename [c2900-universalk9-mz.SPA.151-4.M4.bin]?
Delete flash:/c2900-universalk9-mz.SPA.151-4.M4.bin?
[confirm]

R1#reload
Proceed with reload? [confirm]
Boot process failed...
The system is unable to boot automatically. The BOOT
environment variable needs to be set to a bootable
image.
rommon 1 >
```

50) Use Internet search to find system recovery instructions for your model of router. Recover the system image using the TFTP server.

http://www.cisco.com/c/en/us/td/docs/routers/access/1900/software/configuration/guide/Software_Configuration/appendixCrommon.html is the first hit when searching for 'Cisco 2900 rommon recovery'.

Go to the "["Recovering the System Image \(tftpdnld\)" section.](#)

Use **ALL CAPITAL LETTERS** for this configuration:

```
rommon 1 > IP_ADDRESS=10.10.10.1
rommon 2 > IP_SUBNET_MASK=255.255.255.0
rommon 3 > DEFAULT_GATEWAY=10.10.10.1
rommon 4 > TFTP_SERVER=10.10.10.10
rommon 5 > TFTP_FILE=c2900-universalk9-mz.SPA.151-4.M4.bin
rommon 6 > TFTP_DESTINATION=flash:
rommon 7 > TFTP_TIMEOUT=120
rommon 8 > tftpdnld
```

```
IP_ADDRESS: 10.10.10.1
IP_SUBNET_MASK: 255.255.255.0
DEFAULT_GATEWAY: 10.10.10.1
TFTP_SERVER: 10.10.10.10
TFTP_FILE: c2900-universalk9-mz.SPA.151-4.M4.bin
Invoke this command for disaster recovery only.
WARNING: all existing data in all partitions on flash will
be lost!
```

Do you wish to continue? y/n: [n]: y

.....[TIMED OUT] TFTP: Operation terminated.

51) If you are using the latest version of Packet Tracer the download will time out because the GigabitEthernet interface stays down in rommon mode. You have completed this part of the lab once you have entered the tftpdownld command.

This is as far as we can go with this part of the lab. You've seen how to recover the system image, the only thing left to do in a real world scenario after the download has completed is:

```
rommon 9 > reset
```

IOS Image Upgrade

4) Verify SW1 is running C2960 Software (C2960-LANBASE-M), Version 12.2(25)FX

```
SW1#sh version
Cisco IOS Software, C2960 Software (C2960-LANBASE-M),
Version 12.2(25)FX
```

5) Use the TFTP server to upgrade to c2960-lanbasek9-mz.150-2.SE4.bin

```
SW1#copy tftp flash
Address or name of remote host []? 10.10.10.10
Source filename []? c2960-lanbasek9-mz.150-2.SE4.bin
Destination filename [c2960-lanbasek9-mz.150-2.SE4.bin]?

Accessing tftp://10.10.10.10/c2960-lanbasek9-mz.150-
2.SE4.bin....
Loading c2960-lanbasek9-mz.150-2.SE4.bin from 10.10.10.10:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
[OK - 4670455 bytes]

4670455 bytes copied in 3.086 secs (121674 bytes/sec)
```

```
SW1#show flash
Directory of flash:/

1 -rw- 4414921 <no date> c2960-lanbase-mz.122-25.FX.bin
3 -rw- 4670455 <no date> c2960-lanbasek9-mz.150-2.SE4.bin
2 -rw- 1054 <no date> config.text

64016384 bytes total (54929954 bytes free)
```

```
SW1#config t
SW1(config)#boot system c2960-lanbasek9-mz.150-2.SE4.bin
```

6) Reboot and verify the switch is running the new software version

```
SW1#reload
Proceed with reload? [confirm]
```

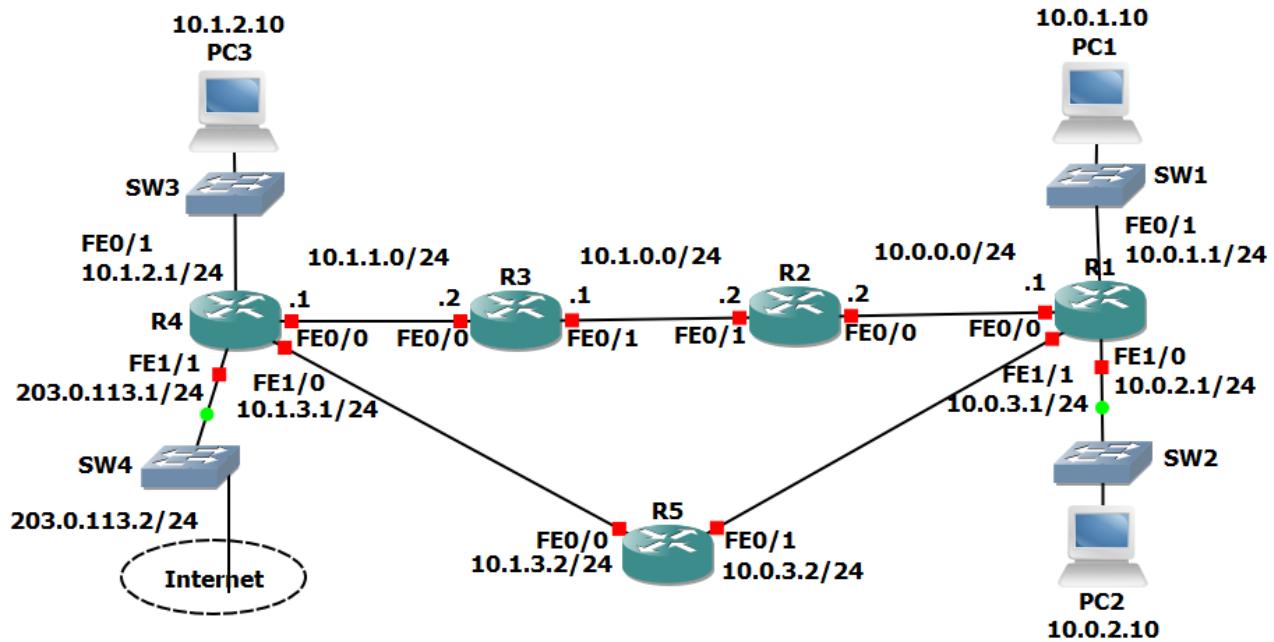
```
SW1#show version
Cisco IOS Software, C2960 Software (C2960-LANBASEK9-M) ,
Version 15.0(2)SE4, RELEASE SOFTWARE (fc1)
```

16 Routing Fundamentals - Lab Exercise

In this lab you will configure and verify connected, local, static, summary and default routes. You will also examine the effect of longest prefix match routing.

All routers and switches are unconfigured at the start of the lab. The PCs have been configured with their network settings.

Lab Topology

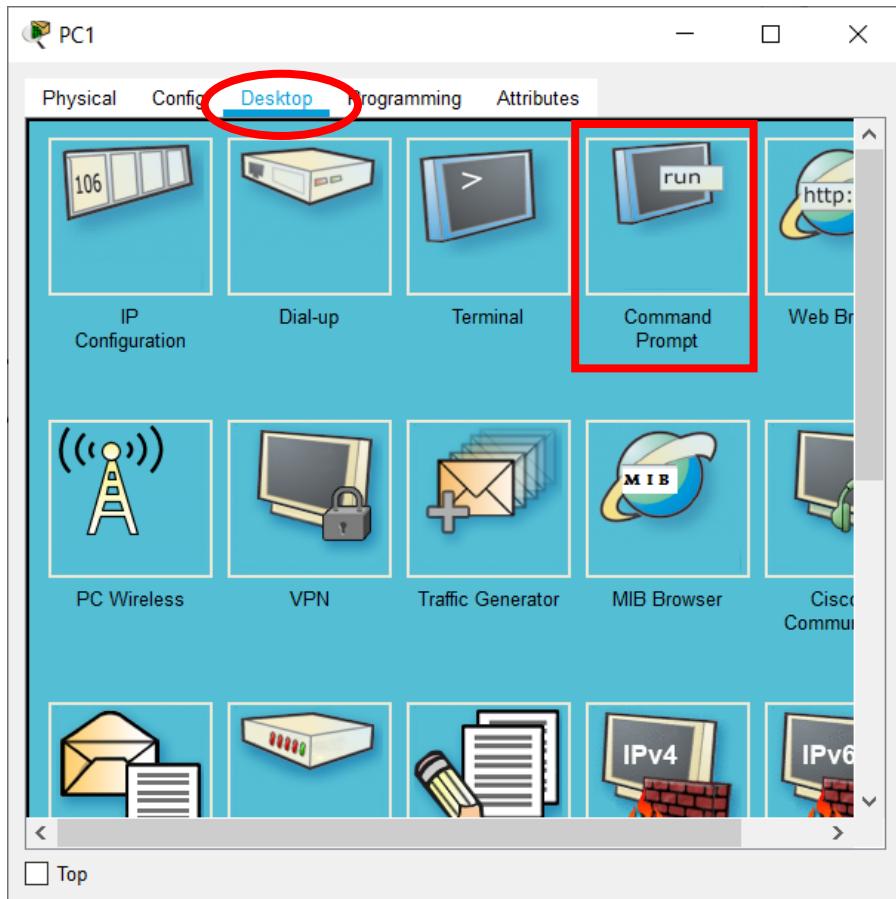


Load the Startup Configurations

[Download the '16 Routing Fundamentals.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Connected and Local Routes

- 1) Say no when asked if you would like to enter the initial configuration dialog on each router.
- 2) Configure hostnames on the routers according to the Lab Topology diagram.
- 3) Configure IP addresses on R1 according to the Lab Topology diagram
- 4) Verify routes have been automatically added for the connected and local networks (note that local routes only appear from IOS 15)
- 5) Do you see routes for all networks that R1 is directly connected to? Why or why not?
- 6) Should you be able to ping from PC1 to PC2? Verify this.
(Click on PC1 then 'Desktop' and 'Command Prompt' to access its command line interface.)



- 7) Verify the traffic path from PC1 to PC2. Use the 'tracert' command.
- 8) Should you be able to ping from PC1 to PC3? Verify this.

Static Routes

- 9) Configure IP addresses on R2, R3 and R4 according to the Lab Topology diagram. Do not configure the Internet FastEthernet 1/1 interface on R4. Do not configure R5.
- 10) Verify PC3 can ping its default gateway at 10.1.2.1
- 11) Configure static routes on R1, R2, R3 and R4 to allow connectivity between all their subnets. Use /24 prefixes for each network.
- 12) Verify connectivity between PC1, PC2 and PC3.
- 13) Verify the path traffic takes from PC1 to PC3.

Summary Routes

- 14) Remove all the static routes on R1
- 15) Verify that PC1 loses connectivity to PC3
- 16) Restore connectivity to all subnets with a single command on R1.
- 17) Verify the routing table on R1 does not contain /24 routes to remote subnets.
- 18) Ensure that connectivity is restored between PC1 and PC3.

Longest Prefix Match

- 19) Configure IP addresses on R5 according to the Lab Topology diagram
- 20) Do not add any additional routes. Does PC1 have reachability to the FastEthernet 0/0 interface on R5? If so, which path will the traffic take?
- 21) Ensure reachability over the shortest possible path from R5 to all directly connected networks on R1. Achieve this with a single command.
- 22) Verify the path traffic takes from PC1 to the FastEthernet 0/0 interface on R5.
- 23) Verify the path the return traffic takes from R5 to PC1.
- 24) Ensure that traffic between PC1 and the FastEthernet 0/0 interface on R5 takes the most direct path in both directions.
- 25) Verify that traffic between PC1 and the FastEthernet 0/0 interface on R5 takes the most direct path in both directions.

Default Route and Load Balancing

- 26) Configure an IP address on the Internet FastEthernet 1/1 interface on R4 according to the lab topology diagram.
- 27) Ensure that all PCs have a route out to the internet through the Internet Service Provider connection on R4. (Note that the lab does not actually have Internet connectivity.)
- 28) Traffic from PC1 and PC2 going to the internet should be load balanced over R2 and R5.

16 Routing Fundamentals - Answer Key

In this lab you will configure and verify connected, local, static, summary and default routes. You will also examine the effect of longest prefix match routing.

All routers and switches are unconfigured at the start of the lab. The PCs have been configured with their network settings.

Connected and Local Routes

- 1) Say no when asked if you would like to enter the initial configuration dialog on each router.

Would you like to enter the initial configuration dialog?
[yes/no]: no

- 2) Configure hostnames on the routers according to the Lab Topology diagram.

On R1:

```
Router(config)# hostname R1
```

Repeat to configure the correct hostname on the other routers.

- 3) Configure IP addresses on R1 according to the Lab Topology diagram

```
R1(config)#int f0/0
R1(config-if)#ip address 10.0.0.1 255.255.255.0
R1(config-if)#no shut
R1(config-if)#int f0/1
R1(config-if)#ip address 10.0.1.1 255.255.255.0
R1(config-if)#no shut
R1(config-if)#int f1/0
R1(config-if)#ip address 10.0.2.1 255.255.255.0
R1(config-if)#no shut
R1(config-if)#int f1/1
R1(config-if)#ip address 10.0.3.1 255.255.255.0
R1(config-if)#no shut
```

4) Verify routes have been automatically added for the connected and local networks (note that local routes only appear from IOS 15)

```
Router#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile,
B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS
inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route
```

Gateway of last resort is not set

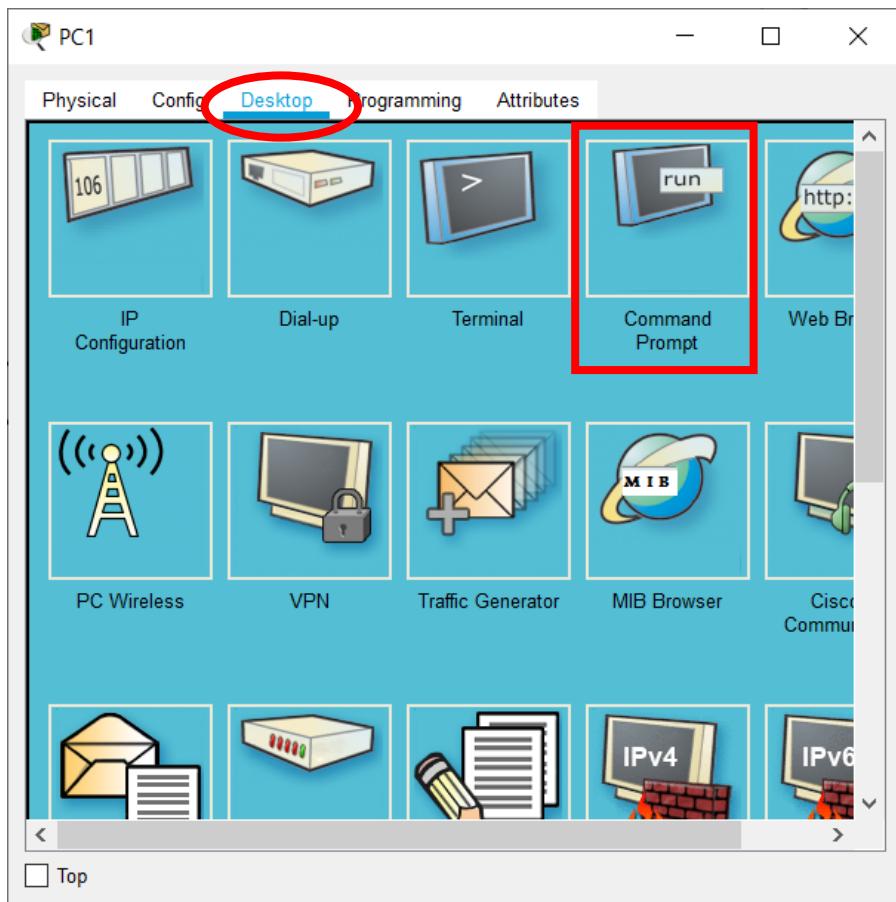
```
10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
```

5) Do you see routes for all networks that R1 is directly connected to? Why or why not?

You cannot see routes for the links connected to R2 and R5 (10.0.0.0/24 and 10.0.3.0/24). The interfaces on R2 and R5 are shutdown by default so the links are down. Both sides of the link must be up for it to be live. A router will not insert routes in its routing table which use links which are down.

You can see routes for the links which are connected to the switches SW1 and SW2 (10.0.1.0/24 and 10.0.2.0/24). Switch ports are not shutdown by default so those links are up.

6) Should you be able to ping from PC1 to PC2? Verify this.
(Click on PC1 then 'Desktop' and 'Command Prompt' to access its command line interface.)



Ping from PC1 to PC2 should be successful as both PCs are in networks which R1 is directly connected to.

```
C:\>ping 10.0.2.10
```

```
Pinging 10.0.2.10 with 32 bytes of data:
```

```
Reply from 10.0.2.10: bytes=32 time<1ms TTL=127
Reply from 10.0.2.10: bytes=32 time<1ms TTL=127
Reply from 10.0.2.10: bytes=32 time<1ms TTL=127
Reply from 10.0.2.10: bytes=32 time=1ms TTL=127
```

```
Ping statistics for 10.0.2.10:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

7) Verify the traffic path from PC1 to PC2. Use the 'tracert' command.

The ping went via R1 at 10.0.1.1

```
C:\>tracert 10.0.2.10
```

```
Tracing route to 10.0.2.10 over a maximum of 30 hops:
```

```
1 0 ms 0 ms 0 ms 10.0.1.1  
2 0 ms 0 ms 0 ms 10.0.2.10
```

```
Trace complete.
```

8) Should you be able to ping from PC1 to PC3? Verify this.

Ping from PC1 to PC3 should fail as R1 does not have a route to the 10.1.2.0 network.

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Reply from 10.0.1.1: Destination host unreachable.  
Request timed out.  
Reply from 10.0.1.1: Destination host unreachable.  
Reply from 10.0.1.1: Destination host unreachable.
```

```
Ping statistics for 10.1.2.10:
```

```
Bytes: Sent = 4, Received = 0, Lost = 4 (100% loss)
```

Static Routes

9) Configure IP addresses on R2, R3 and R4 according to the Lab Topology diagram. Do not configure the Internet FastEthernet 1/1 interface on R4. Do not configure R5.

```
R2(config)#int f0/0
R2(config-if)#ip add 10.0.0.2 255.255.255.0
R2(config-if)#no shut
R2(config-if)#int f0/1
R2(config-if)#ip add 10.1.0.2 255.255.255.0
R2(config-if)#no shut

R3(config)#int f0/1
R3(config-if)#ip add 10.1.0.1 255.255.255.0
R3(config-if)#no shut
R3(config-if)#int f0/0
R3(config-if)#ip add 10.1.1.2 255.255.255.0
R3(config-if)#no shut

R4(config)#int f0/0
R4(config-if)#ip add 10.1.1.1 255.255.255.0
R4(config-if)#no shut
R4(config-if)#int f0/1
R4(config-if)#ip add 10.1.2.1 255.255.255.0
R4(config-if)#no shut
R4(config-if)#int f1/0
R4(config-if)#ip add 10.1.3.1 255.255.255.0
R4(config-if)#no shut
```

10) Verify PC3 can ping its default gateway at 10.1.2.1

```
C:\>ping 10.1.2.1

Pinging 10.1.2.1 with 32 bytes of data:

Reply from 10.1.2.1: bytes=32 time<1ms TTL=255

Ping statistics for 10.1.2.1:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

11) Configure static routes on R1, R2, R3 and R4 to allow connectivity between all their subnets. Use /24 prefixes for each network.

```
R1(config)#ip route 10.1.0.0 255.255.255.0 10.0.0.2
R1(config)#ip route 10.1.1.0 255.255.255.0 10.0.0.2
R1(config)#ip route 10.1.2.0 255.255.255.0 10.0.0.2
R1(config)#ip route 10.1.3.0 255.255.255.0 10.0.0.2

R2(config)#ip route 10.0.1.0 255.255.255.0 10.0.0.1
R2(config)#ip route 10.0.2.0 255.255.255.0 10.0.0.1
R2(config)#ip route 10.0.3.0 255.255.255.0 10.0.0.1
R2(config)#ip route 10.1.1.0 255.255.255.0 10.1.0.1
R2(config)#ip route 10.1.2.0 255.255.255.0 10.1.0.1
R2(config)#ip route 10.1.3.0 255.255.255.0 10.1.0.1

R3(config)#ip route 10.0.0.0 255.255.255.0 10.1.0.2
R3(config)#ip route 10.0.1.0 255.255.255.0 10.1.0.2
R3(config)#ip route 10.0.2.0 255.255.255.0 10.1.0.2
R3(config)#ip route 10.0.3.0 255.255.255.0 10.1.0.2
R3(config)#ip route 10.1.2.0 255.255.255.0 10.1.1.1
R3(config)#ip route 10.1.3.0 255.255.255.0 10.1.1.1

R4(config)#ip route 10.1.0.0 255.255.255.0 10.1.1.2
R4(config)#ip route 10.0.0.0 255.255.255.0 10.1.1.2
R4(config)#ip route 10.0.1.0 255.255.255.0 10.1.1.2
R4(config)#ip route 10.0.2.0 255.255.255.0 10.1.1.2
R4(config)#ip route 10.0.3.0 255.255.255.0 10.1.1.2
```

12) Verify connectivity between PC1, PC2 and PC3.

Ping PC2 and PC3 from PC1.

```
C:\>ping 10.0.2.10
```

```
Pinging 10.0.2.10 with 32 bytes of data:
```

```
Reply from 10.0.2.10: bytes=32 time=1ms TTL=127
Reply from 10.0.2.10: bytes=32 time=1ms TTL=127
Reply from 10.0.2.10: bytes=32 time<1ms TTL=127
Reply from 10.0.2.10: bytes=32 time=1ms TTL=127
```

```
Ping statistics for 10.0.2.10:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
Reply from 10.1.2.10: bytes=32 time=1ms TTL=124
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
```

```
Ping statistics for 10.1.2.10:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

Ping PC3 from PC2.

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
Reply from 10.1.2.10: bytes=32 time=3ms TTL=124
```

```
Ping statistics for 10.1.2.10:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 3ms, Average = 0ms
```

Ping verifies two way reachability so we have now verified full connectivity.

13) Verify the path traffic takes from PC1 to PC3.

```
C:\>tracert 10.1.2.10
```

```
Tracing route to 10.1.2.10 over a maximum of 30 hops:
```

```
1 0 ms 0 ms 1 ms 10.0.1.1
2 0 ms 1 ms 0 ms 10.0.0.2
3 1 ms 0 ms 0 ms 10.1.0.1
4 1 ms 0 ms 0 ms 10.1.1.1
5 0 ms 0 ms 0 ms 10.1.2.10
```

Trace complete.

The traffic goes via the path R1 > R2 > R3 > R4

Note that IP return traffic (PC3 to PC1 in this case) does not necessarily have to come back along the same path, although it typically will.

Summary Routes

14) Remove all the static routes on R1

```
R1(config)#no ip route 10.1.0.0 255.255.255.0 10.0.0.2
R1(config)#no ip route 10.1.1.0 255.255.255.0 10.0.0.2
R1(config)#no ip route 10.1.2.0 255.255.255.0 10.0.0.2
R1(config)#no ip route 10.1.3.0 255.255.255.0 10.0.0.2
```

15) Verify that PC1 loses connectivity to PC3

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Reply from 10.0.1.1: Destination host unreachable.
```

```
Ping statistics for 10.1.2.10:
```

```
packets: Sent = 4, Received = 0, Lost = 4 (100% loss)
```

16) Restore connectivity to all subnets with a single command on R1.

A summary route to 10.1.0.0/16 will add all remote subnets with one command.

```
R1(config)#ip route 10.1.0.0 255.255.0.0 10.0.0.2
```

17) Verify the routing table on R1 does not contain /24 routes to remote subnets.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, 1 - LISP
      + - replicated route, % - next hop override
```

```
Gateway of last resort is not set
```

```
10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
S 10.1.0.0/16 [1/0] via 10.0.0.2
```

18) Ensure that connectivity is restored between PC1 and PC3.

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
```

```
Ping statistics for 10.1.2.10:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

Longest Prefix Match

19) Configure IP addresses on R5 according to the Lab Topology diagram

```
R5(config)#int f0/0
R5(config-if)#ip add 10.1.3.2 255.255.255.0
R5(config-if)#no shut
R5(config-if)#int f0/1
R5(config-if)#ip add 10.0.3.2 255.255.255.0
R5(config-if)#no shut
```

20) Do not add any additional routes. Does PC1 have reachability to the FastEthernet 0/0 interface on R5? If so, which path will the traffic take?

The summary route on R1 to 10.1.0.0/16 will provide a route to R5 over the path R1 > R2 > R3 > R4 > R5, but R5 does not have a route back to PC1.

A ping from PC1 to 10.1.3.2 on R5 will fail.

```
C:\>ping 10.1.3.2
```

```
Pinging 10.1.3.2 with 32 bytes of data:
```

```
Request timed out.  
Request timed out.  
Request timed out.  
Request timed out.
```

```
Ping statistics for 10.1.3.2:
```

```
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)
```

Traceroute will show replies from R1 > R2 > R3 > R4 before failing. (You can hit Ctrl-C on the keyboard to break out of the command.)

```
C:\>tracert 10.1.3.2
```

```
Tracing route to 10.1.3.2 over a maximum of 30 hops:
```

```
1 0 ms 0 ms 1 ms 10.0.1.1  
2 0 ms 0 ms 0 ms 10.0.0.2  
3 1 ms 0 ms 0 ms 10.1.0.1  
4 0 ms 2 ms 0 ms 10.1.1.1  
5 * * * Request timed out.  
6 *  
Control-C  
^C
```

- 21) Ensure reachability over the shortest possible path from R5 to all directly connected networks on R1. Achieve this with a single command.

Add a summary route on R5 for all the directly connected networks on R1.

```
R5(config)#ip route 10.0.0.0 255.255.0.0 10.0.3.1
```

22) Verify the path traffic takes from PC1 to the FastEthernet 0/0 interface on R5.

```
C:\>tracert 10.1.3.2
```

```
Tracing route to 10.1.3.2 over a maximum of 30 hops:
```

```
1 0 ms 0 ms 0 ms 10.0.1.1
2 0 ms 0 ms 0 ms 10.0.0.2
3 0 ms 0 ms 0 ms 10.1.0.1
4 0 ms 0 ms 1 ms 10.1.1.1
5 * 0 ms 1 ms 10.1.3.2
```

Trace complete.

Traffic takes the path R1 > R2 > R3 > R4 > R5

23) Verify the path the return traffic takes from R5 to PC1.

```
R5#traceroute 10.0.1.10
Type escape sequence to abort.
Tracing the route to 10.0.1.10
```

```
1 10.0.3.1 2 msec 0 msec 0 msec
2 10.0.1.10 1 msec 0 msec 1 msec
```

Traffic takes the path R5 > R1. This shows that routers make independent decisions based on their routing table and it is possible (though not common) for return traffic to take an asynchronous path.

24) Ensure that traffic between PC1 and the FastEthernet 0/0 interface on R5 takes the most direct path in both directions.

A route from R1 to the 10.1.3.0/24 network on R5 will achieve this.

```
R1(config)#ip route 10.1.3.0 255.255.255.0 10.0.3.2
```

Traffic to 10.1.3.0/24 will now match two routes in the routing table.

```
S          10.1.0.0/16 [1/0] via 10.0.0.2
S          10.1.3.0/24 [1/0] via 10.0.3.2
```

The new route has a longer prefix match so will be preferred.

25) Verify that traffic between PC1 and the FastEthernet 0/0 interface on R5 takes the most direct path in both directions.

On PC1:

```
C:\>ping 10.1.3.2
```

```
Pinging 10.1.3.2 with 32 bytes of data:
```

```
Reply from 10.1.3.2: bytes=32 time=10ms TTL=254
Reply from 10.1.3.2: bytes=32 time=1ms TTL=254
Reply from 10.1.3.2: bytes=32 time<1ms TTL=254
Reply from 10.1.3.2: bytes=32 time<1ms TTL=254
```

```
Ping statistics for 10.1.3.2:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 10ms, Average = 2ms
```

```
C:\>tracert 10.1.3.2
```

```
Tracing route to 10.1.3.2 over a maximum of 30 hops:
```

```
1 2 ms 0 ms 0 ms 10.0.1.1
2 0 ms 0 ms 0 ms 10.1.3.2
```

```
Trace complete.
```

On R5:

```
R5#traceroute 10.0.1.10
Type escape sequence to abort.
Tracing the route to 10.0.1.10

1 10.0.3.1 0 msec 1 msec 0 msec
2 10.0.1.10 0 msec 0 msec 0 msec
```

Default Route and Load Balancing

26) Configure an IP address on the Internet FastEthernet 1/1 interface on R4 according to the lab topology diagram.

```
R4(config)#int f1/1
R4(config-if)#ip add 203.0.113.1 255.255.255.0
R4(config-if)#no shut
```

27) Ensure that all PCs have a route out to the internet through the Internet Service Provider connection on R4. (Note that the lab does not actually have Internet connectivity.)

```
R1(config)#ip route 0.0.0.0 0.0.0.0 10.0.0.2
R2(config)#ip route 0.0.0.0 0.0.0.0 10.1.0.1
R3(config)#ip route 0.0.0.0 0.0.0.0 10.1.1.1
R4(config)#ip route 0.0.0.0 0.0.0.0 203.0.113.2
R5(config)#ip route 0.0.0.0 0.0.0.0 10.1.3.1
```

All routers should have a default route to the next hop IP on the path to the Internet.

```
R4#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is 203.0.113.2 to network 0.0.0.0

```
10.0.0.0/24 is subnetted, 8 subnets
S 10.0.0.0 [1/0] via 10.1.1.2
S 10.0.1.0 [1/0] via 10.1.1.2
S 10.0.2.0 [1/0] via 10.1.1.2
S 10.0.3.0 [1/0] via 10.1.1.2
S 10.1.0.0 [1/0] via 10.1.1.2
C 10.1.1.0 is directly connected, FastEthernet0/0
C 10.1.2.0 is directly connected, FastEthernet0/1
C 10.1.3.0 is directly connected, FastEthernet1/0
```

```
C 203.0.113.0/24 is directly connected, FastEthernet1/1
S* 0.0.0.0/0 [1/0] via 203.0.113.2
```

28) Traffic from PC1 and PC2 going to the internet should be load balanced over R2 and R5.

Add an additional default route on R1 to send Internet traffic via R5.

```
R1(config)#ip route 0.0.0.0 0.0.0.0 10.0.3.2
```

The routing table will show that R1 will load balance traffic over both paths.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

```
Gateway of last resort is 10.0.0.2 to network 0.0.0.0
```

```
10.0.0.0/8 is variably subnetted, 10 subnets, 3 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
S 10.1.0.0/16 [1/0] via 10.0.0.2
S 10.1.3.0/24 [1/0] via 10.0.3.2
S* 0.0.0.0/0 [1/0] via 10.0.0.2
                           [1/0] via 10.0.3.2
```

Add additional routes on R4 to load balance the return traffic to PC1 and PC2 from the Internet.

```
R4(config)#ip route 10.0.1.0 255.255.255.0 10.1.3.2
R4(config)#ip route 10.0.2.0 255.255.255.0 10.1.3.2
```

R4's routing table should show both paths back to 10.0.1.0/24 and 10.0.2.0/24

R4#sh ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
+ - replicated route, % - next hop override

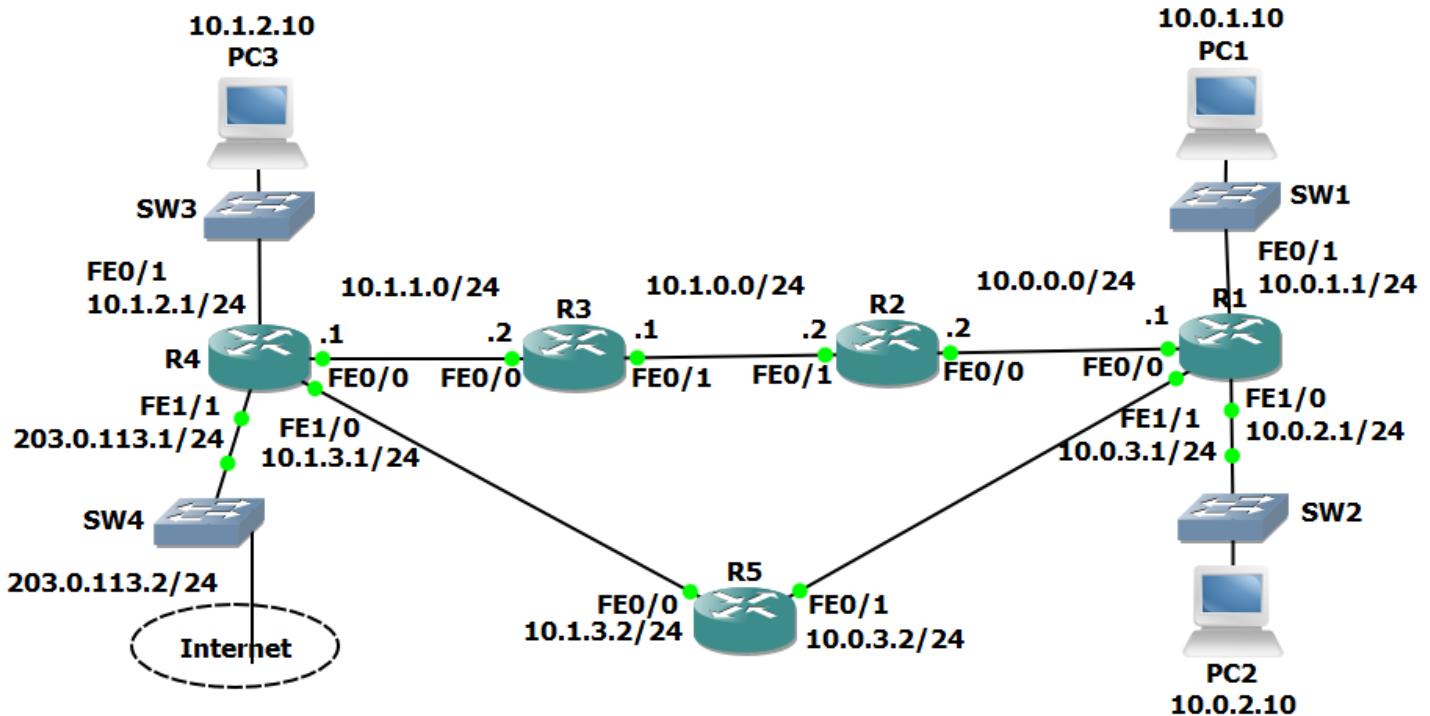
Gateway of last resort is 203.0.113.2 to network 0.0.0.0

```
10.0.0.0/8 is variably subnetted, 11 subnets, 2 masks
S      10.0.0.0/24 [1/0] via 10.1.1.2
S      10.0.1.0/24 [1/0] via 10.1.3.2
          [1/0] via 10.1.1.2
S      10.0.2.0/24 [1/0] via 10.1.3.2
          [1/0] via 10.1.1.2
S      10.0.3.0/24 [1/0] via 10.1.1.2
S      10.1.0.0/24 [1/0] via 10.1.1.2
C      10.1.1.0/24 is directly connected, FastEthernet0/0
L      10.1.1.1/32 is directly connected, FastEthernet0/0
C      10.1.2.0/24 is directly connected, FastEthernet0/1
L      10.1.2.1/32 is directly connected, FastEthernet0/1
C      10.1.3.0/24 is directly connected, FastEthernet1/0
L      10.1.3.1/32 is directly connected, FastEthernet1/0
203.0.113.0/24 is variably subnetted, 2 subnets, 2 masks
C      203.0.113.0/24 is directly connected, FastEthernet1/1
L      203.0.113.1/32 is directly connected, FastEthernet1/1
S*     0.0.0.0/0 [1/0] via 203.0.113.2
```

17 Dynamic Routing Protocols – Lab Exercise

In this lab you will examine features which are common to all Interior Gateway Protocols. Basic configuration for various IGPs is included but is not the focus of this lab exercise. Configuration for each of the IGPs will be covered in detail later.

Lab Topology



Load the Startup Configurations

[Download the '17 Dynamic Routing Protocols.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Routing Protocol Updates

- 1) Enter the commands below on each router to provision a basic RIPv1 configuration and enable RIP on every interface.

```
router rip
network 10.0.0.0
no auto-summary
```

- 2) Debug the routing protocol updates on R1 with the 'debug ip rip' command. Observe the updates being sent and received. What kind of traffic is used (unicast, broadcast or multicast)?
- 3) Enter the commands below to enable RIPv2 on every router.

```
router rip
version 2
```

- 4) What kind of traffic is used for the updates now?
- 5) Turn off all debugging on R1.
- 6) Check that RIP routes have been added to R1 and it has a route to every subnet in the lab.
- 7) Why are there two routes to the 10.1.1.0/24 network in the routing table?

Comparing Routing Protocols

- 8) View the RIP database on R1.
- 9) Enter the commands below on each router to provision a basic OSPF configuration and enable OSPF on every interface.

```
router ospf 1
network 10.0.0.0 0.255.255.255 area 0
```

- 10) Give OSPF time to converge. Are RIP routes included in the routing table on R1 now? Why or why not?
- 11) Why is there now only one route to the 10.1.1.0/24 network?

- 12) Disable interface FastEthernet 0/0 on R2. What do you expect to happen to R1's routing table?
- 13) Verify your expected changes to R1's routing table.
- 14) Aside from the next hop address, what else has changed on the routing table?
- 15) View the OSPF database on R1 with the 'show ip ospf database' command. What is different between it and the RIP database? Why?

Routing Protocol Metrics and Administrative Distance

- 16) Enter the command below to remove OSPF on every router

```
no router ospf 1
```

- 17) Will R1 still have connectivity to R4?
- 18) What is the metric to the 10.1.1.0/24 network on R1?
- 19) Why is there only one route on R1 to the 10.1.1.0/24 network now?
- 20) Make the required change so that there are two routes to the 10.1.1.0/24 network in the routing table on R1.
- 21) Enter the commands below on each router to provision a basic EIGRP configuration and enable EIGRP on every interface.

```
router eigrp 100
no auto-summary
network 10.0.0.0 0.255.255.255
```

- 22) What changes do you expect to see in the routing tables? Why?
- 23) Verify the changes to the routing table on R1.
- 24) What is the metric to the 10.1.1.0/24 network on R1?
- 25) Why is there only one route to the 10.1.1.0/24 network on R1?

26) Disable RIP and EIGRP on R5 with the commands below.

```
R5(config)#no router rip  
R5(config)#no router eigrp 100
```

27) Configure the network so that there is still connectivity between all subnets if the link between R1 and R2 goes down. Accomplish this with six commands. Do not enable EIGRP on R5 but note that the routing protocol is expected to be enabled there in the future.

28) What changes do you expect to see to the routing table on R1?

29) Verify the changes to the routing table on R1.

30) Verify that traffic from PC1 to PC3 still goes via R2.

31) Shut down interface FastEthernet 0/0 on R2.

32) What changes do you expect to see on R1's routing table?

33) Verify the changes to the routing table on R1.

34) Verify connectivity between PC1 and PC3.

35) Verify the traffic goes via R5.

36) Bring interface FastEthernet 0/0 on R2 back up.

37) Enter the commands below on R5 to provision a basic EIGRP configuration and enable EIGRP on every interface.

```
R5(config)#router eigrp 100  
R5(config-router)#no auto-summary  
R5(config-router)#network 10.0.0.0 0.255.255.255
```

Loopback Interfaces

38) Configure loopback interface 0 on each router. Assign the IP address 192.168.0.x/32, where 'x' is the router number (for example 192.168.0.3/32 on R3.)

39) Is there connectivity to the loopback interfaces from the PCs? Why or why not?

40) Enter the commands below on each router to include the loopback interfaces in EIGRP.

```
R1(config)#router eigrp 100
R1(config-router)#network 192.168.0.0 0.0.0.255
```

41) Verify the loopback interfaces are in the routing table on R1.

42) Verify connectivity from PC1 to the loopback interface on R5.

Adjacencies and Passive Interfaces

43) Enter the command below to verify that R1 has established EIGRP adjacencies with R2 and R5.

```
R1#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
          Address           Interface      Hold Uptime      SRTT      RTO      Q      Seq
          (sec)           (sec)        (ms)      Cnt  Num
 1  10.0.3.2           Fa1/1          14 00:17:21    33  198    0  16
 0  10.0.0.2           Fa0/0          11 00:19:21    36  216    0  32
```

44) Verify that traffic from R5 to the directly connected interfaces on R1 goes via the FastEthernet 0/1 interface.

45) Enter the commands below to configure the loopback interface and the link to R5 as passive interfaces on R1.

```
R1(config-if)#router eigrp 100
R1(config-router)#passive-interface loopback0
R1(config-router)#passive-interface fastethernet1/1
```

46) What changes do you expect to see in the routing table on R5 and why?

47) Verify the expected changes to the routing table on R5.

17 Dynamic Routing Protocols - Answer Key

In this lab you will examine features which are common to all Interior Gateway Protocols. Basic configuration for various IGPs is included but is not the focus of this lab exercise. Configuration for each of the IGPs will be covered in detail later.

Routing Protocol Updates

- 1) Enter the commands below on each router to provision a basic RIPv1 configuration and enable RIP on every interface.

```
router rip
network 10.0.0.0
no auto-summary
```

- 2) Debug the routing protocol updates on R1 with the 'debug ip rip' command. Observe the updates being sent and received. What kind of traffic is used (unicast, broadcast or multicast)?

The updates are being sent on the broadcast address 255.255.255.255. All hosts on the subnet must process the packets.

```
R1#debug ip rip
RIP protocol debugging is on
R1#RIP: sending v1 update to 255.255.255.255 via
FastEthernet0/0 (10.0.0.1)
RIP: build update entries
    network 10.0.1.0 metric 1
    network 10.0.2.0 metric 1
```

- 3) Enter the commands below to enable RIPv2 on every router.

```
router rip
version 2
```

- 4) What kind of traffic is used for the updates now?

The updates are being sent on the RIPv2 multicast address 224.0.0.9. Only RIPv2 routers will process the packets beyond layer 3.

```
RIP: sending v2 update to 224.0.0.9 via FastEthernet1/0
(10.0.2.1)
```

- 5) Turn off all debugging on R1.

```
R1#undebug all
All possible debugging has been turned off
```

- 6) Check that RIP routes have been added to R1 and it has a route to every subnet in the lab.

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
R 10.1.0.0/24 [120/1] via 10.0.0.2, 00:00:08, FastEthernet0/0
R 10.1.1.0/24 [120/2] via 10.0.0.2, 00:00:08, FastEthernet0/0
              [120/2] via 10.0.3.2, 00:00:24, FastEthernet1/1
R 10.1.2.0/24 [120/2] via 10.0.3.2, 00:00:24, FastEthernet1/1
R 10.1.3.0/24 [120/1] via 10.0.3.2, 00:00:24, FastEthernet1/1
```

- 7) Why are there two routes to the 10.1.1.0/24 network in the routing table?

Two paths to 10.1.1.0/24 have an equal metric – a hop count of 2. Both routes are installed in the routing table and the router will perform Equal Cost Load Balancing between the next hops of 10.0.3.2 and 10.0.0.2.

Comparing Routing Protocols

- 8) View the RIP database on R1.

```
R1#sh ip rip database
10.0.0.0/24 auto-summary
10.0.0.0/24 directly connected, FastEthernet0/0
10.0.1.0/24 auto-summary
10.0.1.0/24 directly connected, FastEthernet0/1
10.0.2.0/24 auto-summary
10.0.2.0/24 directly connected, FastEthernet1/0
10.0.3.0/24 auto-summary
10.0.3.0/24 directly connected, FastEthernet1/1
10.1.0.0/24 auto-summary
10.1.0.0/24
[1] via 10.0.0.2, 00:00:01, FastEthernet0/0
10.1.1.0/24 auto-summary
10.1.1.0/24
[2] via 10.0.0.2, 00:00:01, FastEthernet0/0 [2] via
10.0.3.2, 00:00:12, FastEthernet1/1
10.1.2.0/24 auto-summary
10.1.2.0/24
[2] via 10.0.3.2, 00:00:12, FastEthernet1/1
10.1.3.0/24 auto-summary
10.1.3.0/24
[1] via 10.0.3.2, 00:00:12, FastEthernet1/1
```

- 9) Enter the commands below on each router to provision a basic OSPF configuration and enable OSPF on every interface.

```
router ospf 1
network 10.0.0.0 0.255.255.255 area 0
```

10) Give OSPF time to converge. Are RIP routes included in the routing table on R1 now? Why or why not?

The RIP routes are replaced by OSPF because its Administrative Distance of 110 is preferred to RIP's AD of 120.

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2] via 10.0.0.2, 00:01:01, FastEthernet0/0
O 10.1.1.0/24 [110/3] via 10.0.0.2, 00:01:01, FastEthernet0/0
O 10.1.2.0/24 [110/4] via 10.0.0.2, 00:01:01, FastEthernet0/0
O 10.1.3.0/24 [110/13] via 10.0.0.2, 00:01:01, FastEthernet0/0
```

11) Why is there now only one route to the 10.1.1.0/24 network?

OSPF uses cost as its metric which takes into account interface bandwidth. The interfaces on R5 have a configured bandwidth of 10Mbps. The interfaces along the top path of the network topology all have the default FastEthernet bandwidth of 100Mbps so this route is preferred. All traffic will go via the next hop 10.0.0.2.

```
R5#show run | section interface
interface FastEthernet0/0
bandwidth 10000
ip address 10.1.3.2 255.255.255.0
duplex auto
speed auto
interface FastEthernet0/1
bandwidth 10000
ip address 10.0.3.2 255.255.255.0
duplex auto
speed auto
! output truncated
```

12) Disable interface FastEthernet 0/0 on R2. What do you expect to happen to R1's routing table?

OSPF will reconverge. The routes to the 10.1.0.0 networks via R2 will be removed from the routing table and replaced with routes via R5.

```
R2(config)#interface fastEthernet 0/0
R2(config-if)#shutdown

%LINK-5-CHANGED: Interface FastEthernet0/0, changed state
to administratively down

%LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/0, changed state to down

00:27:22: %OSPF-5-ADJCHG: Process 1, Nbr 10.0.3.1 on
FastEthernet0/0 from FULL to DOWN, Neighbor Down: Interface
down or detached
```

13) Verify your expected changes to R1's routing table.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/22] via 10.0.3.2, 00:01:05, FastEthernet1/1
O 10.1.1.0/24 [110/21] via 10.0.3.2, 00:01:05, FastEthernet1/1
O 10.1.2.0/24 [110/21] via 10.0.3.2, 00:01:05, FastEthernet1/1
O 10.1.3.0/24 [110/20] via 10.0.3.2, 00:01:05, FastEthernet1/1
```

14) Aside from the next hop address, what else has changed on the routing table?

The new routes have a higher metric. This is why they were not in the routing table when the path via R2 was up.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
 10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/22] via 10.0.3.2, 00:01:05, FastEthernet1/1
O 10.1.1.0/24 [110/21] via 10.0.3.2, 00:01:05, FastEthernet1/1
O 10.1.2.0/24 [110/21] via 10.0.3.2, 00:01:05, FastEthernet1/1
O 10.1.3.0/24 [110/20] via 10.0.3.2, 00:01:05, FastEthernet1/1
```

15) View the OSPF database on R1 with the 'show ip ospf database' command. What is different between it and the RIP database? Why?

RIP is a Distance Vector routing protocol so it only knows its directly connected neighbors and the lists of networks those neighbors have advertised.

OSPF is a Link State routing protocol so it knows the state of every link on every router in its area.

```
R1#show ip ospf database
      OSPF Router with ID (10.0.3.1) (Process ID 1)

      Router Link States (Area 0)

Link ID ADV Router Age Seq# Checksum Link count
10.1.1.2 10.1.1.2 563 0x80000004 0x003cb6 2
203.0.113.1 203.0.113.1 558 0x80000005 0x00df85 3
10.1.3.2 10.1.3.2 558 0x80000004 0x00844e 2
10.0.3.1 10.0.3.1 218 0x80000007 0x00775c 3
10.1.0.2 10.1.0.2 218 0x80000005 0x00c15c 1

      Net Link States (Area 0)
Link ID ADV Router Age Seq# Checksum
10.1.0.1 10.1.1.2 572 0x80000001 0x00bc09
10.1.1.1 203.0.113.1 563 0x80000001 0x005efe
10.0.3.2 10.1.3.2 559 0x80000001 0x0049f2
10.1.3.1 203.0.113.1 558 0x80000002 0x00025c
```

Routing Protocol Metrics and Administrative Distance

16) Enter the command below to remove OSPF on every router

```
no router ospf 1
```

17) Will R1 still have connectivity to R4?

Yes. RIP is still running so RIP routes will replace the removed OSPF routes in the routing table.

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
R 10.1.0.0/24 [120/3] via 10.0.3.2, 00:00:12, FastEthernet1/1
R 10.1.1.0/24 [120/2] via 10.0.3.2, 00:00:12, FastEthernet1/1
R 10.1.2.0/24 [120/2] via 10.0.3.2, 00:00:12, FastEthernet1/1
R 10.1.3.0/24 [120/1] via 10.0.3.2, 00:00:12, FastEthernet1/1
```

18) What is the metric to the 10.1.1.0/24 network on R1?

A hop count of 2.

19) Why is there only one route on R1 to the 10.1.1.0/24 network now?

Interface FastEthernet 0/0 on R2 is still shut down so no routes go through it.

20) Make the required change so that there are two routes to the 10.1.1.0/24 network in the routing table on R1.

```
R2(config)#interface f0/0
R2(config-if)#no shut

R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile,
B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS
inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
R 10.1.0.0/24 [120/1] via 10.0.0.2, 00:00:03, FastEthernet0/0
R 10.1.1.0/24 [120/2] via 10.0.3.2, 00:00:15, FastEthernet1/1
              [120/2] via 10.0.0.2, 00:00:03, FastEthernet0/0
R 10.1.2.0/24 [120/2] via 10.0.3.2, 00:00:15, FastEthernet1/1
R 10.1.3.0/24 [120/1] via 10.0.3.2, 00:00:15, FastEthernet1/1
```

21) Enter the commands below on each router to provision a basic EIGRP configuration and enable EIGRP on every interface.

```
router eigrp 100
no auto-summary
network 10.0.0.0 0.255.255.255
```

22) What changes do you expect to see in the routing tables? Why?

The RIP routes will be replaced by EIGRP because its Administrative Distance of 90 is preferred to RIP's AD of 120.

23) Verify the changes to the routing table on R1.

```
R1#show ip route
Codes: L - local, C - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
D 10.1.0.0/24 [90/30720] via 10.0.0.2, 00:00:32, FastEthernet0/0
D 10.1.1.0/24 [90/33280] via 10.0.0.2, 00:00:29, FastEthernet0/0
D 10.1.2.0/24 [90/35840] via 10.0.0.2, 00:00:25, FastEthernet0/0
D 10.1.3.0/24 [90/261120] via 10.0.3.2, 00:00:19, FastEthernet1/1
```

24) What is the metric to the 10.1.1.0/24 network on R1?

A composite metric of 33280.

25) Why is there only one route to the 10.1.1.0/24 network on R1?

EIGRP uses a composite metric which takes into account interface bandwidth and delay. The interfaces on R5 have a configured bandwidth of 10Mbps. The interfaces along the top path of the network topology all have the default FastEthernet bandwidth of 100Mbps so this route is preferred. All traffic will go via the next hop 10.0.0.2.

26) Disable RIP and EIGRP on R5 with the commands below.

```
R5(config)#no router rip
R5(config)#no router eigrp 100
```

27) Configure the network so that there is still connectivity between all subnets if the link between R1 and R2 goes down. Accomplish this with six commands. Do not enable EIGRP on R5 but note that the routing protocol is expected to be enabled there in the future.

Floating static routes need to be added as a backup to the EIGRP routes. We want to ensure EIGRP routes are preferred when available so set the AD to be higher than EIGRP's AD of 90.

```
R1(config)#ip route 10.1.0.0 255.255.0.0 10.0.3.2 95
```

```
R2(config)#ip route 10.0.0.0 255.255.0.0 10.1.0.1 95
```

```
R3(config)#ip route 10.0.0.0 255.255.0.0 10.1.1.1 95
```

```
R4(config)#ip route 10.0.0.0 255.255.0.0 10.1.3.2 95
```

```
R5(config)#ip route 10.0.0.0 255.255.0.0 10.0.3.1 95
```

```
R5(config)#ip route 10.1.0.0 255.255.0.0 10.1.3.1 95
```

R5 is not running EIGRP so it is not currently necessary to set the Administrative Distance for its routes to 95. It is required to prevent the floating static routes from being preferred when EIGRP is enabled in the future however.

Summary routes need to be used to accomplish the task in six commands.

28) What changes do you expect to see to the routing table on R1?

The summary route will be added to the routing table but not used because it has a prefix length of /16, compared to the EIGRP routes which have a longer prefix length of /24.

If individual floating static routes had been added for each of the /24 destination networks then these would not have appeared in the routing table (unless a link went down) because EIGRP has a better Administrative Distance.

29) Verify the changes to the routing table on R1.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
S 10.1.0.0/16 [95/0] via 10.0.3.2
D 10.1.0.0/24 [90/30720] via 10.0.0.2, 00:04:48, FastEthernet0/0
D 10.1.1.0/24 [90/33280] via 10.0.0.2, 00:04:45, FastEthernet0/0
D 10.1.2.0/24 [90/35840] via 10.0.0.2, 00:04:41, FastEthernet0/0
D 10.1.3.0/24 [90/266240] via 10.0.0.2, 00:03:02, FastEthernet0/0
```

30) Verify that traffic from PC1 to PC3 still goes via R2.

```
C:\>tracert 10.1.2.10
```

```
Tracing route to 10.1.2.10 over a maximum of 30 hops:
```

```
1 1 ms 0 ms 1 ms 10.0.1.1
2 0 ms 3 ms 0 ms 10.0.0.2
3 1 ms 0 ms 0 ms 10.1.0.1
4 0 ms 1 ms 0 ms 10.1.1.1
5 * 0 ms 0 ms 10.1.2.10
```

Trace complete.

31) Shut down interface FastEthernet 0/0 on R2.

```
R2(config)#interface f0/0
R2(config-if)#shutdown
```

32) What changes do you expect to see on R1's routing table?

The EIGRP routes will be removed.

33) Verify the changes to the routing table on R1.

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
S 10.1.0.0/16 [95/0] via 10.0.3.2
```

34) Verify connectivity between PC1 and PC3.

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Reply from 10.1.2.10: bytes=32 time=1ms TTL=125
Reply from 10.1.2.10: bytes=32 time=1ms TTL=125
Reply from 10.1.2.10: bytes=32 time=1ms TTL=125
Reply from 10.1.2.10: bytes=32 time<1ms TTL=125
```

```
Ping statistics for 10.1.2.10:
```

```
Bytes: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

35) Verify the traffic goes via R5.

```
C:\>tracert 10.1.2.10
```

```
Tracing route to 10.1.2.10 over a maximum of 30 hops:
```

```
1 0 ms 0 ms 1 ms 10.0.1.1  
2 0 ms 0 ms 0 ms 10.0.3.2  
3 0 ms 0 ms 0 ms 10.1.3.1  
4 0 ms 0 ms 1 ms 10.1.2.10  
Trace complete.
```

36) Bring interface FastEthernet 0/0 on R2 back up.

```
R2(config)#interface f0/0  
R2(config-if)#no shutdown
```

37) Enter the commands below on R5 to provision a basic EIGRP configuration and enable EIGRP on every interface.

```
R5(config)#router eigrp 100  
R5(config-router)#no auto-summary  
R5(config-router)#network 10.0.0.0 0.255.255.255
```

Loopback Interfaces

38) Configure loopback interface 0 on each router. Assign the IP address 192.168.0.x/32, where 'x' is the router number (for example 192.168.0.3/32 on R3.)

```
R1(config)#interface loopback0  
R1(config-if)#ip address 192.168.0.1 255.255.255.255
```

```
R2(config)#interface loopback0  
R2(config-if)#ip address 192.168.0.2 255.255.255.255
```

```
R3(config)#interface loopback0  
R3(config-if)#ip address 192.168.0.3 255.255.255.255
```

```
R4(config)#interface loopback0  
R4(config-if)#ip address 192.168.0.4 255.255.255.255
```

```
R5(config)#interface loopback0  
R5(config-if)#ip address 192.168.0.5 255.255.255.255
```

39) Is there connectivity to the loopback interfaces from the PCs? Why or why not?

There is no connectivity from the PCs to the loopback interfaces because they are not in the routing tables of the routers (apart from the local loopback interface on each router). The loopback interfaces are not in the routing tables because they are in the 192.168.0.0/24 range which has not been included in the routing protocol.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
S 10.1.0.0/16 [95/0] via 10.0.3.2
D 10.1.0.0/24 [90/30720] via 10.0.0.2, 00:03:01, FastEthernet0/0
D 10.1.1.0/24 [90/33280] via 10.0.0.2, 00:03:01, FastEthernet0/0
D 10.1.2.0/24 [90/35840] via 10.0.0.2, 00:03:01, FastEthernet0/0
D 10.1.3.0/24 [90/261120] via 10.0.3.2, 00:02:32, FastEthernet1/1
  192.168.0.0/32 is subnetted, 1 subnets
C 192.168.0.1/32 is directly connected, Loopback0
```

40) Enter the commands below on each router to include the loopback interfaces in EIGRP.

```
R1(config)#router eigrp 100
R1(config-router)#network 192.168.0.0 0.0.0.255
```

41) Verify the loopback interfaces are in the routing table on R1.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
S 10.1.0.0/16 [95/0] via 10.0.3.2
D 10.1.0.0/24 [90/30720] via 10.0.0.2, 00:04:53, FastEthernet0/0
D 10.1.1.0/24 [90/33280] via 10.0.0.2, 00:04:53, FastEthernet0/0
D 10.1.2.0/24 [90/35840] via 10.0.0.2, 00:04:53, FastEthernet0/0
D 10.1.3.0/24 [90/261120] via 10.0.3.2, 00:04:24, FastEthernet1/1
192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
D 192.168.0.2/32 [90/156160] via 10.0.0.2, 00:00:26, FastEthernet0/0
D 192.168.0.3/32 [90/158720] via 10.0.0.2, 00:00:23, FastEthernet0/0
D 192.168.0.4/32 [90/161280] via 10.0.0.2, 00:00:20, FastEthernet0/0
D 192.168.0.5/32 [90/386560] via 10.0.3.2, 00:00:17, FastEthernet1/1
```

42) Verify connectivity from PC1 to the loopback interface on R5.

```
C:\>ping 192.168.0.5
```

```
Pinging 192.168.0.5 with 32 bytes of data:
```

```
Reply from 192.168.0.5: bytes=32 time<1ms TTL=254
Reply from 192.168.0.5: bytes=32 time=1ms TTL=254
Reply from 192.168.0.5: bytes=32 time<1ms TTL=254
Reply from 192.168.0.5: bytes=32 time=4ms TTL=254
```

```
Ping statistics for 192.168.0.5:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 4ms, Average = 1ms
```

Adjacencies and Passive Interfaces

43) Enter the command below to verify that R1 has established EIGRP adjacencies with R2 and R5.

```
R1#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for process 100
      H   Address           Interface      Hold Uptime      SRTT      RTO  Q  Seq
          (sec)          (ms)          Cnt Num
1   10.0.3.2           Fa1/1           14 00:17:21    33  198  0  16
0   10.0.0.2           Fa0/0           11 00:19:21    36  216  0  32
```

44) Verify that traffic from R5 to the directly connected interfaces on R1 goes via the FastEthernet 0/1 interface.

```
R5#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 3 masks
S 10.0.0.0/16 [95/0] via 10.0.3.1
D 10.0.0.0/24 [90/261120] via 10.0.3.1, 00:07:30, FastEthernet0/1
D 10.0.1.0/24 [90/261120] via 10.0.3.1, 00:07:30, FastEthernet0/1
D 10.0.2.0/24 [90/261120] via 10.0.3.1, 00:07:30, FastEthernet0/1
C 10.0.3.0/24 is directly connected, FastEthernet0/1
L 10.0.3.2/32 is directly connected, FastEthernet0/1
S 10.1.0.0/16 [95/0] via 10.1.3.1
D 10.1.0.0/24 [90/263680] via 10.0.3.1, 00:07:30, FastEthernet0/1
[90/263680] via 10.1.3.1, 00:07:30, FastEthernet0/0
D 10.1.1.0/24 [90/261120] via 10.1.3.1, 00:07:30, FastEthernet0/0
D 10.1.2.0/24 [90/261120] via 10.1.3.1, 00:07:30, FastEthernet0/0
C 10.1.3.0/24 is directly connected, FastEthernet0/0
L 10.1.3.2/32 is directly connected, FastEthernet0/0
192.168.0.0/32 is subnetted, 5 subnets
D 192.168.0.1/32 [90/386560] via 10.0.3.1, 00:03:37, FastEthernet0/1
D 192.168.0.2/32 [90/389120] via 10.0.3.1, 00:03:32, FastEthernet0/1
D 192.168.0.3/32 [90/389120] via 10.1.3.1, 00:03:29, FastEthernet0/0
D 192.168.0.4/32 [90/386560] via 10.1.3.1, 00:03:26, FastEthernet0/0
C 192.168.0.5/32 is directly connected, Loopback0
```

45) Enter the commands below to configure the loopback interface and the link to R5 as passive interfaces on R1.

```
R1(config-if)#router eigrp 100
R1(config-router)#passive-interface loopback0
R1(config-router)#passive-interface fastethernet1/1
```

46) What changes do you expect to see in the routing table on R5 and why?

The EIGRP adjacency between R1 and R5 will go down. All EIGRP routes via R1 will be removed from the routing table and replaced with routes via R4. Configuring the loopback interface as a passive interface on R1 does not affect any routing tables but is a best practice.

```
R1(config-router)#
%DUAL-5-NBRCHANGE: IP-EIGRP 100: Neighbor 10.0.3.2
(FastEthernet1/1) is down: holding time expired
```

47) Verify the expected changes to the routing table on R5.

```
R5#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override

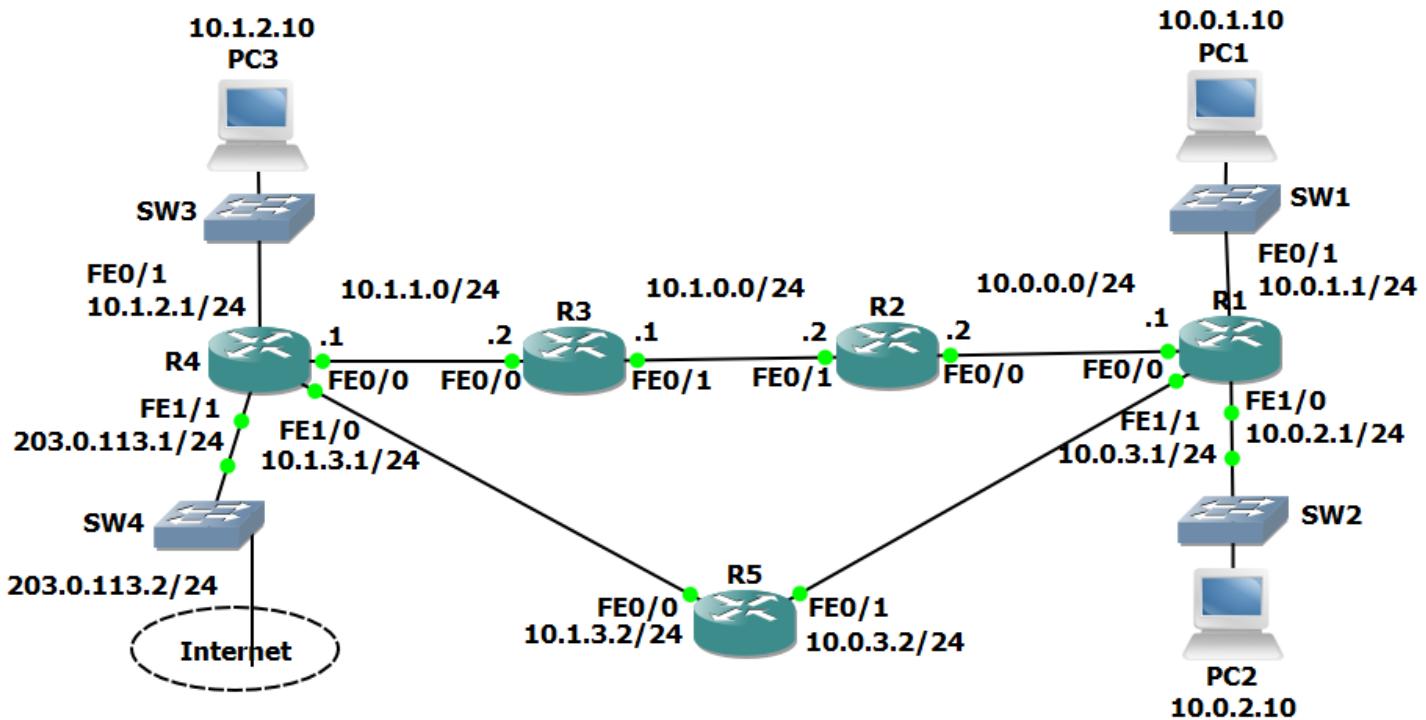
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 12 subnets, 3 masks
S 10.0.0.0/16 [95/0] via 10.0.3.1
D 10.0.0.0/24 [90/266240] via 10.1.3.1, 00:13:54, FastEthernet0/0
D 10.0.1.0/24 [90/268800] via 10.1.3.1, 00:13:54, FastEthernet0/0
D 10.0.2.0/24 [90/268800] via 10.1.3.1, 00:13:54, FastEthernet0/0
C 10.0.3.0/24 is directly connected, FastEthernet0/1
L 10.0.3.2/32 is directly connected, FastEthernet0/1
S 10.1.0.0/16 [95/0] via 10.1.3.1
D 10.1.0.0/24 [90/263680] via 10.1.3.1, 00:13:54, FastEthernet0/0
D 10.1.1.0/24 [90/261120] via 10.1.3.1, 00:13:54, FastEthernet0/0
D 10.1.2.0/24 [90/261120] via 10.1.3.1, 00:13:54, FastEthernet0/0
C 10.1.3.0/24 is directly connected, FastEthernet0/0
L 10.1.3.2/32 is directly connected, FastEthernet0/0
192.168.0.0/32 is subnetted, 5 subnets
D 192.168.0.1/32 [90/394240] via 10.1.3.1, 00:10:01, FastEthernet0/0
D 192.168.0.2/32 [90/391680] via 10.1.3.1, 00:09:56, FastEthernet0/0
D 192.168.0.3/32 [90/389120] via 10.1.3.1, 00:09:53, FastEthernet0/0
D 192.168.0.4/32 [90/386560] via 10.1.3.1, 00:09:50, FastEthernet0/0
C 192.168.0.5/32 is directly connected, Loopback0
```

18 Connectivity Troubleshooting – Lab Exercise

In this lab you will troubleshoot a connectivity issue.

Lab Topology



Load the Startup Configurations

[Download the '18 Connectivity Troubleshooting.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Troubleshoot Connectivity

- 1) Use ping to test connectivity from PC1 to PC3.
- 2) Use traceroute to determine where the problem is likely to be.
- 3) Determine the issue and fix it to restore connectivity between PC1 and PC3.

18 Connectivity Troubleshooting - Answer Key

In this lab you will troubleshoot a connectivity issue.

Troubleshoot Connectivity

- 1) Use ping to test connectivity from PC1 to PC3.

Connectivity is down.

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Request timed out.  
Request timed out.  
Reply from 10.1.0.1: Destination host unreachable.  
Reply from 10.1.0.1: Destination host unreachable.
```

```
Ping statistics for 10.1.2.10:  
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)
```

- 2) Use traceroute to determine where the problem is likely to be.

Traceroute is making it as far as 10.1.0.1 on R3 so that is a good place to continue troubleshooting.

```
C:\>tracert 10.1.2.10
```

```
Tracing route to 10.1.2.10 over a maximum of 30 hops:
```

```
1 0 ms 0 ms 0 ms 10.0.1.1  
2 0 ms 0 ms 0 ms 10.0.0.2  
3 0 ms 0 ms 0 ms 10.1.0.1  
4 1 ms * 1 ms 10.1.0.1  
5 * 0 ms  
Control-C  
^C
```



- 3) Determine the issue and fix it to restore connectivity between PC1 and PC3.

The first thing to do is check that R3 has a route to PC3's network 10.1.2.0/24

```
R3#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
S 10.0.0.0/24 [1/0] via 10.1.0.2
S 10.0.1.0/24 [1/0] via 10.1.0.2
S 10.0.2.0/24 [1/0] via 10.1.0.2
S 10.0.3.0/24 [1/0] via 10.1.0.2
C 10.1.0.0/24 is directly connected, FastEthernet0/1
L 10.1.0.1/32 is directly connected, FastEthernet0/1
C 10.1.1.0/24 is directly connected, FastEthernet0/0
L 10.1.1.2/32 is directly connected, FastEthernet0/0
S 10.1.3.0/24 [1/0] via 10.1.1.1
```

The network is not in the routing table. From the 'show ip route' command we can also see that static routes are being used. We need to add a static route for the 10.1.2.0/24 network with R4 as the next hop.

```
R3(config)#ip route 10.1.2.0 255.255.255.0 10.1.1.1
```

Lastly we need to verify connectivity is restored from PC1.

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Request timed out.
```

```
Request timed out.
```

```
Reply from 10.1.2.10: bytes=32 time<1ms TTL=124
```

```
Reply from 10.1.2.10: bytes=32 time=1ms TTL=124
```

```
Ping statistics for 10.1.2.10:
```

```
Packets: Sent = 4, Received = 2, Lost = 2 (50% loss),
```

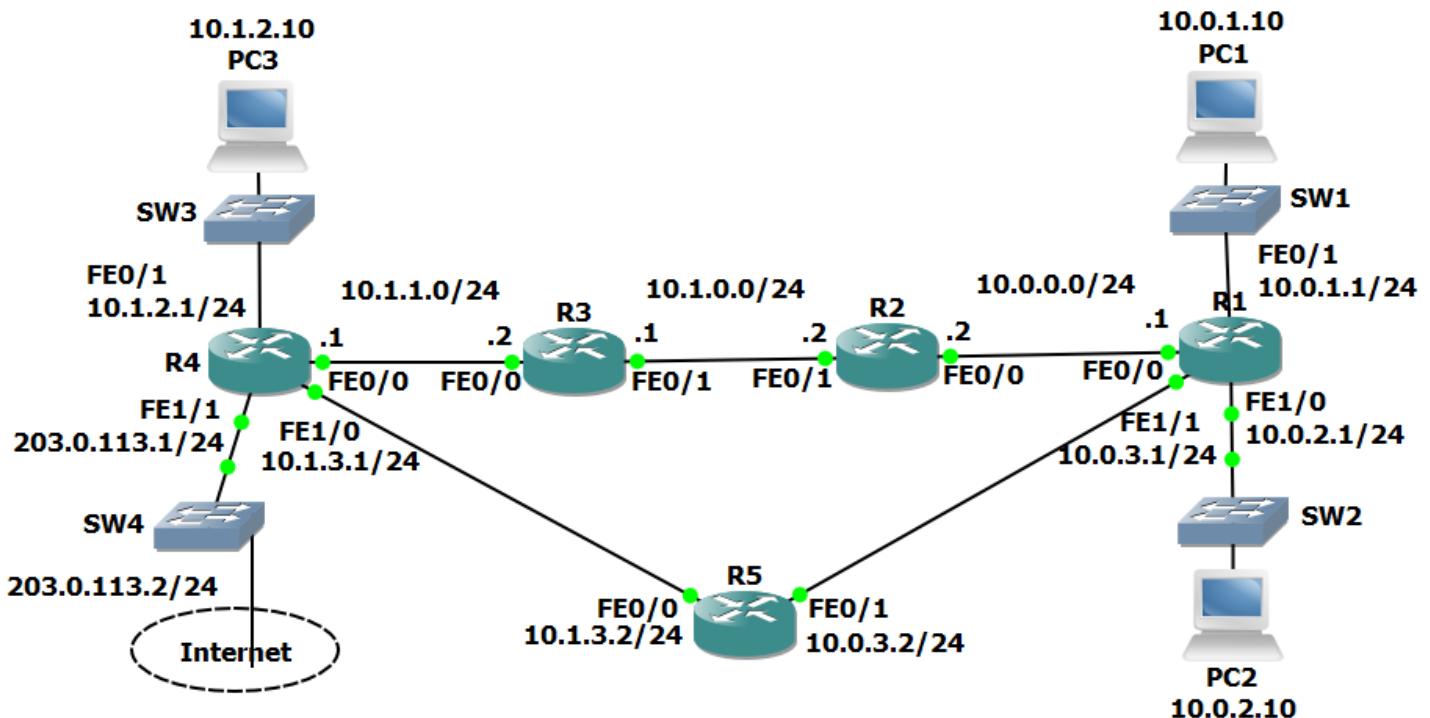
```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

19-1 IGP Interior Gateway Protocol Fundamentals Configuration – Lab Exercise

In this lab you will configure the RIPv2 and EIGRP routing protocols. IP addresses are already configured on the routers.

Lab Topology



Load the Startup Configurations

[Download the '19-1 IGP Fundamentals.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

RIP Configuration

- 1) Enable RIPv2 on every router. Ensure all networks except 203.0.113.0/24 are advertised. Do not perform any summarisation.
- 2) Verify all networks are in the router's routing tables.
- 3) Verify that routing is working by checking that PC1 has connectivity to PC3.
- 4) Ensure that all routers have a route to the 203.0.113.0/24 network. Internal routes must not be advertised to the Service Provider at 203.0.113.2.
- 5) Verify that all routers have a path to the 203.0.113.0/24 network.
- 6) Configure a default static route on R4 to the Internet via the service provider at 203.0.113.2
- 7) Ensure that all other routers learn via RIP how to reach the Internet.
- 8) Verify all routers have a route to the Internet.

EIGRP Configuration

- 9) Enable EIGRP AS 100 on every router. Ensure all networks except 203.0.113.0/24 are advertised in EIGRP.
- 10) Verify the routers have formed adjacencies with each other.
- 11) Which routing protocol (RIP or EIGRP) do you expect routes to the 10.x.x.x networks to be learned from in the routing tables?
- 12) Do you expect to see any routes from the other routing protocol in the routing tables?
- 13) View the routing tables to verify your answers.

19-1 IGP Interior Gateway Protocol Fundamentals Configuration - Answer Key

In this lab you will configure the RIPv2 and EIGRP routing protocols. IP addresses are already configured on the routers.

RIP Configuration

- 1) Enable RIPv2 on every router. Ensure all networks except 203.0.113.0/24 are advertised. Do not perform any summarisation.

On every router:

```
R1(config)#router rip
R1(config-router)#version 2
R1(config-router)#no auto-summary
R1(config-router)#network 10.0.0.0
```

- 2) Verify all networks are in the router's routing tables.

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
R 10.1.0.0/24 [120/1] via 10.0.0.2, 00:00:00, FastEthernet0/0
R 10.1.1.0/24 [120/2] via 10.0.0.2, 00:00:00, FastEthernet0/0
              [120/2] via 10.0.3.2, 00:00:10, FastEthernet1/1
R 10.1.2.0/24 [120/2] via 10.0.3.2, 00:00:10, FastEthernet1/1
R 10.1.3.0/24 [120/1] via 10.0.3.2, 00:00:24, FastEthernet1/1
```

3) Verify that routing is working by checking that PC1 has connectivity to PC3.

```
C:\>ping 10.1.2.10
```

```
Pinging 10.1.2.10 with 32 bytes of data:
```

```
Request timed out.
```

```
Reply from 10.1.2.10: bytes=32 time=1ms TTL=125
```

```
Reply from 10.1.2.10: bytes=32 time<1ms TTL=125
```

```
Reply from 10.1.2.10: bytes=32 time<1ms TTL=125
```

```
Ping statistics for 10.1.2.10:
```

```
_packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

4) Ensure that all routers have a route to the 203.0.113.0/24 network.
Internal routes must not be advertised to the Service Provider at 203.0.113.2.

The 203.0.113.0/24 network must be added to the RIP process on R4, and interface FastEthernet 1/1 facing the service provider configured as a passive interface to avoid sending out internal network information.

```
R4(config)#router rip
R4(config-router)#passive-interface f1/1
R4(config-router)#network 203.0.113.0
```

5) Verify that all routers have a path to the 203.0.113.0/24 network.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

```
Gateway of last resort is not set
```

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
```

```

C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
R 10.1.0.0/24 [120/1] via 10.0.0.2, 00:00:24, FastEthernet0/0
R 10.1.1.0/24 [120/2] via 10.0.0.2, 00:00:24, FastEthernet0/0
                [120/2] via 10.0.3.2, 00:00:14, FastEthernet1/1
R 10.1.2.0/24 [120/2] via 10.0.3.2, 00:00:14, FastEthernet1/1
R 10.1.3.0/24 [120/1] via 10.0.3.2, 00:00:14, FastEthernet1/1
R 203.0.113.0/24 [120/2] via 10.0.3.2, 00:00:12, FastEthernet1/1

```

- 6) Configure a default static route on R4 to the Internet via the service provider at 203.0.113.2

```
R4(config)#ip route 0.0.0.0 0.0.0.0 203.0.113.2
```

- 7) Ensure that all other routers learn via RIP how to reach the Internet.

Inject the default static route into RIP on R4.

```
R4(config)#router rip
R4(config-router)#default-information originate
```

- 8) Verify all routers have a route to the Internet.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is 10.0.3.2 to network 0.0.0.0

```

10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
R 10.1.0.0/24 [120/1] via 10.0.0.2, 00:00:20, FastEthernet0/0
R 10.1.1.0/24 [120/2] via 10.0.0.2, 00:00:20, FastEthernet0/0
                [120/2] via 10.0.3.2, 00:00:13, FastEthernet1/1
R 10.1.2.0/24 [120/2] via 10.0.3.2, 00:00:13, FastEthernet1/1
R 10.1.3.0/24 [120/1] via 10.0.3.2, 00:00:13, FastEthernet1/1
R 203.0.113.0/24 [120/2] via 10.0.3.2, 00:00:13, FastEthernet1/1
R* 0.0.0.0/0 [120/2] via 10.0.3.2, 00:00:13, FastEthernet1/1

```

EIGRP Configuration

9) Enable EIGRP AS 100 on every router. Ensure all networks except 203.0.113.0/24 are advertised.

On every router:

```
R1(config)#router eigrp 100
R1(config-router)#network 10.0.0.0
```

10) Verify the routers have formed adjacencies with each other.

```
R1#sh ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
  H   Address           Interface      Hold Uptime    SRTT    RTO    Q    Seq
      (sec)          (ms)          Cnt Num
  0   10.0.0.2          Fa0/0          11 00:00:20  21    126    0   10
  1   10.0.3.2          Fa1/1          11 00:00:10  44    264    0    6
```

11) Which routing protocol (RIP or EIGRP) do you expect routes to the 10.x.x.x networks to be learned from in the routing tables?

Both RIP and EIGRP are advertising the 10.x.x.x networks. EIGRP has a better (lower) administrative distance of 90 compared to RIP's AD of 120, so the EIGRP routes will be installed in the router's routing tables.

12) Do you expect to see any routes from the other routing protocol in the routing tables?

Only RIP (not EIGRP) is advertising the 203.0.113.0/24 network and injecting the default static route. Those routes will remain unchanged in the routing tables.

13) View the routing tables to verify your answers.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override

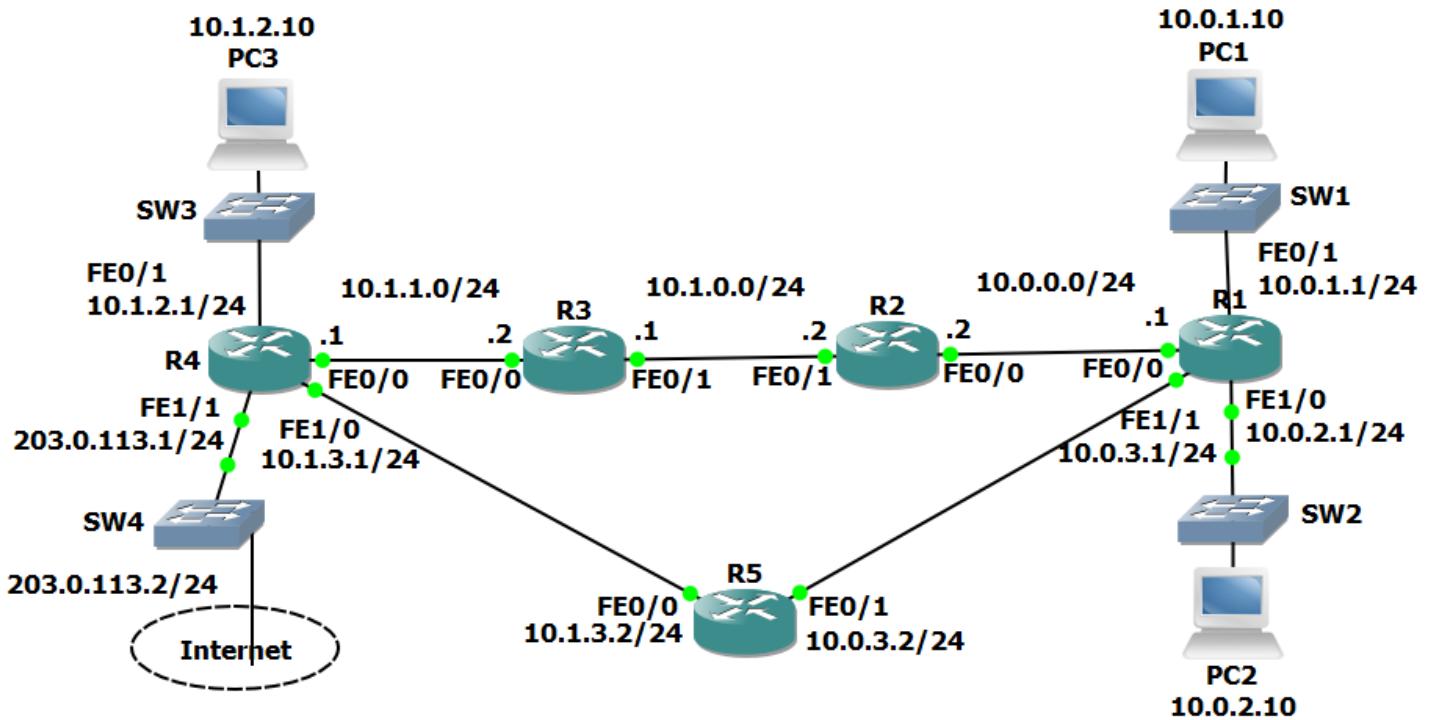
Gateway of last resort is 10.0.3.2 to network 0.0.0.0

  10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
D 10.1.0.0/24 [90/30720] via 10.0.0.2, 00:06:39, FastEthernet0/0
D 10.1.1.0/24 [90/33280] via 10.0.0.2, 00:06:21, FastEthernet0/0
D 10.1.2.0/24 [90/35840] via 10.0.0.2, 00:06:15, FastEthernet0/0
D 10.1.3.0/24 [90/261120] via 10.0.3.2, 00:06:09, FastEthernet1/1
R 203.0.113.0/24 [120/2] via 10.0.3.2, 00:00:22, FastEthernet1/1
R* 0.0.0.0/0 [120/2] via 10.0.3.2, 00:00:22, FastEthernet1/1
```

20-1 OSPF Configuration – Lab Exercise

In this lab you will configure the OSPF routing protocol. IP addresses have already been configured on the router interfaces.

Lab Topology



Load the Startup Configurations

[Download the '20-1 OSPF Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

You can learn the theory for this section and shortcut your path to CCNA certification by getting my CCNA Gold Bootcamp course:
<https://www.flackbox.com/cisco-ccna-training-course>

OSPF Basic Configuration

- 1) Enable a loopback interface on each router. Use the IP address 192.168.0.x/32, where 'x' is the router number. For example 192.168.0.3/32 on R3.
- 2) Enable single area OSPF on every router. Ensure all networks except 203.0.113.0/24 are advertised.
- 3) What do you expect the OSPF Router ID to be on R1? Verify this.
- 4) Verify the routers have formed adjacencies with each other.
- 5) Verify all 10.x.x.x networks and loopbacks are in the router's routing tables.
- 6) Set the reference bandwidth so that a 100 Gbps interface will have a cost of 1.
- 7) What will the OSPF cost be on the FastEthernet links? Verify this.
- 8) What effect does this have on the cost to the 10.1.2.0/24 network from R1?

OSPF Cost

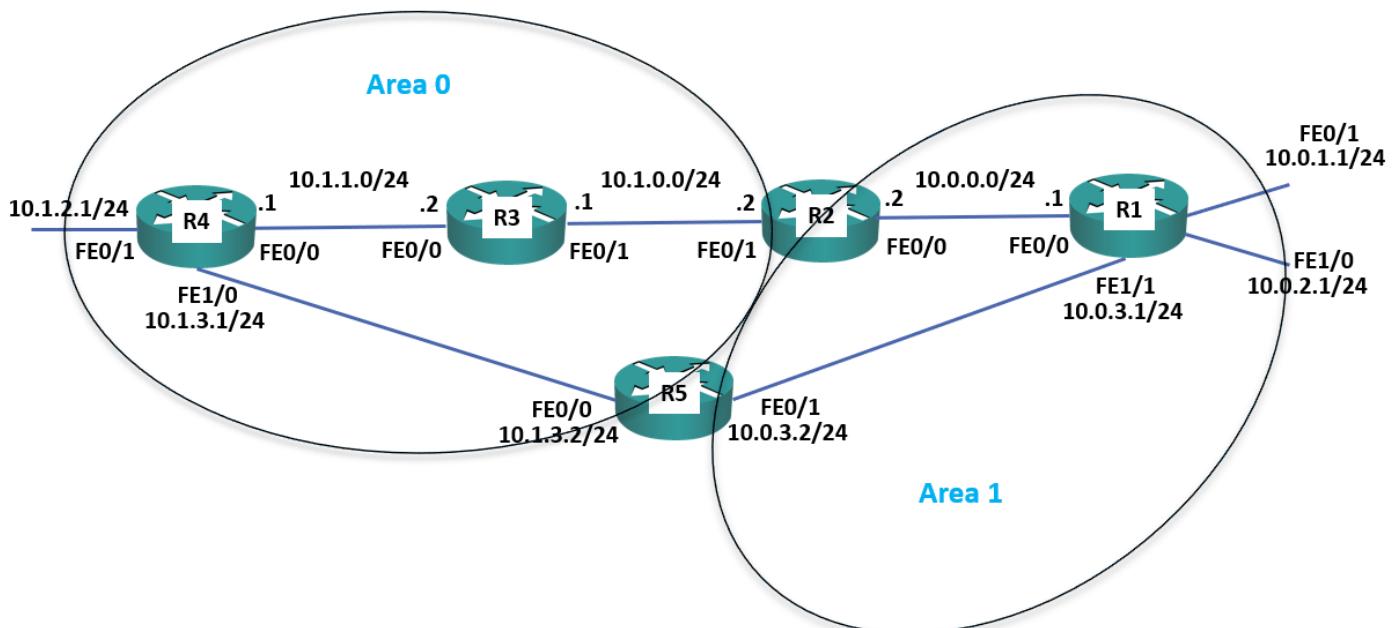
- 9) There are two possible paths which R1 could use to reach the 10.1.2.0/24 network – either through R2 or R5. Which route is in the routing table?
- 10) Change this so that traffic from R1 to 10.1.2.0/24 will be load balanced via both R2 and R5.
- 11) Verify that traffic to the 10.1.2.0/24 network from R1 is load balanced via both R2 and R5.

Default Route Injection

- 12) Ensure that all routers have a route to the 203.0.113.0/24 network. Internal routes must not be advertised to the Service Provider at 203.0.113.2.
- 13) Verify that all routers have a path to the 203.0.113.0/24 network.
- 14) Configure a default static route on R4 to the Internet via the service provider at 203.0.113.2
- 15) Ensure that all other routers learn via OSPF how to reach the Internet.
- 16) Verify all routers have a route to the Internet.

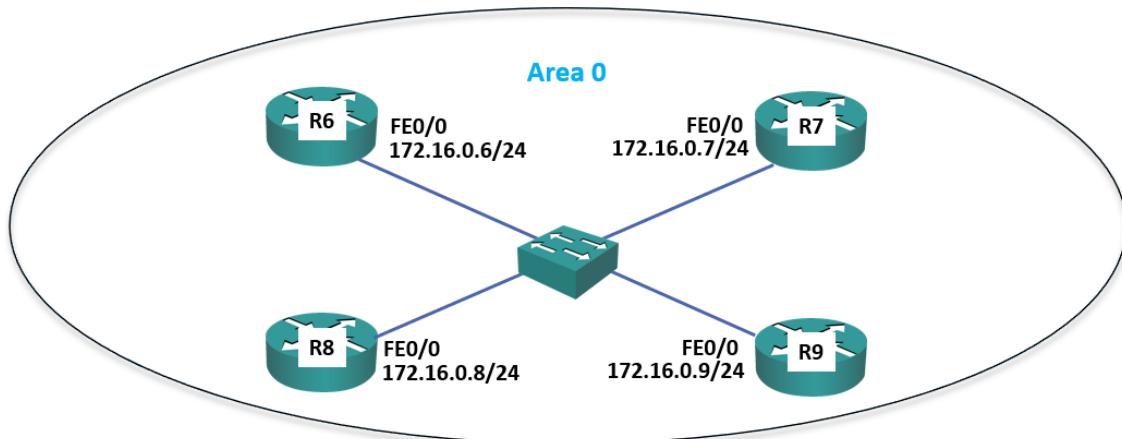
Multi-Area OSPF

- 17) Convert the network to use multi-area OSPF. R3 and R4 should be backbone routers, R1 a normal router in Area 1, and R2 and R5 ABRs as shown in the diagram below.
Save your changes to the startup config and reboot the routers to ensure the changes take effect.



- 18) Verify the router's interfaces are in the correct areas.
- 19) Verify the routers have formed adjacencies with each other.
- 20) What change do you expect to see on R1's routing table? Verify this (give the routing table a few seconds to converge).
- 21) Do you see less routes in R1's routing table? Why or why not?
- 22) Configure summary routes on the Area Border Routers for the 10.0.0.0/16 and 10.1.0.0/16 networks.
- 23) Verify R1 now sees a single summary route for 10.1.0.0/16 rather than individual routes for the 10.1.x.x networks.
- 24) Verify R1 is receiving a summary route for the 10.1.0.0/16 network from both R2 and R5.
- 25) R1 is routing traffic to 10.1.0.0/16 via R2 only. Why is it not load balancing the traffic through both R2 and R5?

DR and BDR Designated Routers



- 26) Enable a loopback interface on routers R6 to R9. Use the IP address 192.168.0.x/32, where 'x' is the router number. For example 192.168.0.6/32 on R6.
- 27) Enable OSPF for Area 0 on the Loopback 0 and FastEthernet 0/0 interfaces on routers R6 to R9.

- 28) Set the reference bandwidth on routers R6 to R9 so that a 100 Gbps interface will have a cost of 1.
- 29) Which routers do you expect to be the DR and BDR on the Ethernet segment? Verify this.
- 30) Set R6 as the Designated Router without changing any IP addresses.
- 31) Verify R6 is the Designated Router.

20-1 OSPF Configuration - Answer Key

In this lab you will configure the OSPF routing protocol. IP addresses have already been configured on the router interfaces.

OSPF Basic Configuration

- 1) Enable a loopback interface on routers R1 to R5. Use the IP address 192.168.0.x/32, where 'x' is the router number. For example 192.168.0.3/32 on R3.

On routers R1 to R5:

```
R1(config)#interface loopback0
R1(config-if)#ip address 192.168.0.1 255.255.255.255
```

- 2) Enable single area OSPF on routers R1 to R5. Ensure all networks except 172.16.0.0/24 and 203.0.113.0/24 are advertised.

On routers R1 to R5:

```
R1(config)#router ospf 1
R1(config-router)#network 10.0.0.0 0.255.255.255 area 0
R1(config-router)#network 192.168.0.0 0.0.0.255 area 0
```

You can use different network statements, as long as they cover the range of IP addresses configured on the router interfaces.

3) What do you expect the OSPF Router ID to be on R1? Verify this.

The loopback address is used for the Router ID, 192.168.0.1

```
R1#sh ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
Router ID 192.168.0.1
  Number of areas in this router is 1. 1 normal 0 stub 0
nssa
  Maximum path: 4
  Routing for Networks:
    10.0.0.0 0.255.255.255 area 0
    192.168.0.0 0.0.0.255 area 0
  Routing Information Sources:
    Gateway          Distance      Last Update
    192.168.0.1      110          00:00:25
    192.168.0.2      110          00:00:25
    192.168.0.3      110          00:00:25
    192.168.0.4      110          00:00:25
    192.168.0.5      110          00:00:25
  Distance: (default is 110)
```

4) Verify routers R1 to R5 have formed adjacencies with each other.

```
R1#show ip ospf neighbor

Neighbor ID      Pri  State            Dead Time      Address          Interface
192.168.0.5      1    FULL/BDR        00:00:31      10.0.3.2        FastEthernet1/1
192.168.0.2      1    FULL/DR         00:00:39      10.0.0.2        FastEthernet0/0
```

5) Verify all 10.x.x.x networks and loopbacks are in the routing tables on R1 to R5.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2] via 10.0.0.2, 00:03:13, FastEthernet0/0
O 10.1.1.0/24 [110/3] via 10.0.0.2, 00:02:51, FastEthernet0/0
      [110/3] via 10.0.3.2, 00:02:51, FastEthernet1/1
O 10.1.2.0/24 [110/3] via 10.0.3.2, 00:02:51, FastEthernet1/1
O 10.1.3.0/24 [110/2] via 10.0.3.2, 00:02:51, FastEthernet1/1
      192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O 192.168.0.2/32 [110/2] via 10.0.0.2, 00:03:25, FastEthernet0/0
O 192.168.0.3/32 [110/3] via 10.0.0.2, 00:03:13, FastEthernet0/0
O 192.168.0.4/32 [110/3] via 10.0.3.2, 00:02:51, FastEthernet1/1
O 192.168.0.5/32 [110/2] via 10.0.3.2, 00:03:25, FastEthernet1/1
```

6) Set the reference bandwidth so that a 100 Gbps interface will have a cost of 1.

Remember to do this on all routers R1 to R5.

```
R1(config)#router ospf 1
R1(config-router)#auto-cost reference-bandwidth 100000
```

7) What will the OSPF cost be on the FastEthernet links? Verify this.

OSPF Cost = Reference bandwidth / Interface bandwidth.

$$100000 / 100 = 1000$$

```
R1#show ip ospf interface FastEthernet 0/0

FastEthernet0/0 is up, line protocol is up
Internet address is 10.0.0.1/24, Area 0
Process ID 1, Router ID 192.168.0.1, Network Type BROADCAST, Cost: 1000
```

8) What effect does this have on the cost to the 10.1.2.0/24 network from R1?

The cost changes from 3 to 3000.

Before reference bandwidth change:

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2] via 10.0.0.2, 00:03:13, FastEthernet0/0
O 10.1.1.0/24 [110/3] via 10.0.0.2, 00:02:51, FastEthernet0/0
      [110/3] via 10.0.3.2, 00:02:51, FastEthernet1/1
O 10.1.2.0/24 [110/3] via 10.0.3.2, 00:02:51, FastEthernet1/1
O 10.1.3.0/24 [110/2] via 10.0.3.2, 00:02:51, FastEthernet1/1
      192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O 192.168.0.2/32 [110/2] via 10.0.0.2, 00:03:25, FastEthernet0/0
O 192.168.0.3/32 [110/3] via 10.0.0.2, 00:03:13, FastEthernet0/0
O 192.168.0.4/32 [110/3] via 10.0.3.2, 00:02:51, FastEthernet1/1
O 192.168.0.5/32 [110/2] via 10.0.3.2, 00:03:25, FastEthernet1/1
```

After reference bandwidth change:

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2000] via 10.0.0.2, 00:01:04, FastEthernet0/0
O 10.1.1.0/24 [110/3000] via 10.0.0.2, 00:01:04, FastEthernet0/0
              [110/3000] via 10.0.3.2, 00:01:04, FastEthernet1/1
O 10.1.2.0/24 [110/3000] via 10.0.3.2, 00:01:04, FastEthernet1/1
O 10.1.3.0/24 [110/2000] via 10.0.3.2, 00:01:04, FastEthernet1/1
192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O 192.168.0.2/32 [110/1012] via 10.0.0.2, 00:01:04, FastEthernet0/0
O 192.168.0.3/32 [110/2012] via 10.0.0.2, 00:01:04, FastEthernet0/0
O 192.168.0.4/32 [110/2012] via 10.0.3.2, 00:01:04, FastEthernet1/1
O 192.168.0.5/32 [110/1012] via 10.0.3.2, 00:01:04, FastEthernet1/1
```

OSPF Cost

9) There are two possible paths which R1 could use to reach the 10.1.2.0/24 network – either through R2 or R5. Which route is in the routing table?

The path via R5 at 10.0.3.2.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2000] via 10.0.0.2, 00:01:04, FastEthernet0/0
O 10.1.1.0/24 [110/3000] via 10.0.0.2, 00:01:04, FastEthernet0/0
              [110/3000] via 10.0.3.2, 00:01:04, FastEthernet1/1
O 10.1.2.0/24 [110/3000] via 10.0.3.2, 00:01:04, FastEthernet1/1
O 10.1.3.0/24 [110/2000] via 10.0.3.2, 00:01:04, FastEthernet1/1
      192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O 192.168.0.2/32 [110/1012] via 10.0.0.2, 00:01:04, FastEthernet0/0
O 192.168.0.3/32 [110/2012] via 10.0.0.2, 00:01:04, FastEthernet0/0
O 192.168.0.4/32 [110/2012] via 10.0.3.2, 00:01:04, FastEthernet1/1
O 192.168.0.5/32 [110/1012] via 10.0.3.2, 00:01:04, FastEthernet1/1
```

10) Change this so that traffic from R1 to 10.1.2.0/24 will be load balanced via both R2 and R5.

Since we changed the reference bandwidth, all interfaces have a cost of 1000. The current path from R1 > R5 > R4 has a cost of 3000 (the cost of the destination interface itself is also counted in the total cost).

The path from R1 > R2 > R3 > R4 has a cost of 4000.

The easiest way to configure both paths to have the same cost is to configure the links from R1 > R5 and R5 > R4 to have a cost of 1500 each. (R1 > R5 = 1500, plus R5 > R4 = 1500, plus cost of 10.1.2.0/24 interface on R4 = 1000. Total = 4000).

```
R1(config)#int f1/1
R1(config-if)#ip ospf cost 1500

R5(config)#int f0/0
R5(config-if)# ip ospf cost 1500
R5(config)#int f0/1
R5(config-if)# ip ospf cost 1500

R4(config)#int f1/0
R4(config-if)# ip ospf cost 1500
```

11) Verify that traffic to the 10.1.2.0/24 network from R1 is load balanced via both R2 and R5.

```
R1#sh ip route
Codes: L - local, C - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2000] via 10.0.0.2, 00:05:57, FastEthernet0/0
O 10.1.1.0/24 [110/3000] via 10.0.0.2, 00:00:38, FastEthernet0/0
O 10.1.2.0/24 [110/4000] via 10.0.3.2, 00:00:25, FastEthernet1/1
               [110/4000] via 10.0.0.2, 00:00:25, FastEthernet0/0
O 10.1.3.0/24 [110/3000] via 10.0.3.2, 00:00:25, FastEthernet1/1
               192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O 192.168.0.2/32 [110/1012] via 10.0.0.2, 00:05:57, FastEthernet0/0
O 192.168.0.3/32 [110/2012] via 10.0.0.2, 00:05:57, FastEthernet0/0
O 192.168.0.4/32 [110/3012] via 10.0.3.2, 00:00:25, FastEthernet1/1
               [110/3012] via 10.0.0.2, 00:00:25, FastEthernet0/0
O 192.168.0.5/32 [110/1512] via 10.0.3.2, 00:00:38, FastEthernet1/1
```

Default Route Injection

12) Ensure that routers R1 to R5 have a route to the 203.0.113.0/24 network.
Internal routes must not be advertised to the Service Provider at 203.0.113.2.

The 203.0.113.0/24 network must be added to the OSPF process on R4, and interface FastEthernet 1/1 facing the service provider configured as a passive interface to avoid sending out internal network information.

```
R4(config)#router ospf 1
R4(config-router)#passive-interface f1/1
R4(config-router)#network 203.0.113.0 0.0.0.255 area 0
```

13) Verify that routers R1 to R5 have a path to the 203.0.113.0/24 network.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2000] via 10.0.0.2, 00:08:40, FastEthernet0/0
O 10.1.1.0/24 [110/3000] via 10.0.0.2, 00:03:21, FastEthernet0/0
O 10.1.2.0/24 [110/4000] via 10.0.3.2, 00:03:08, FastEthernet1/1
              [110/4000] via 10.0.0.2, 00:03:08, FastEthernet0/0
O 10.1.3.0/24 [110/3000] via 10.0.3.2, 00:03:08, FastEthernet1/1
              192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O 192.168.0.2/32 [110/1012] via 10.0.0.2, 00:08:40, FastEthernet0/0
O 192.168.0.3/32 [110/2012] via 10.0.0.2, 00:08:40, FastEthernet0/0
O 192.168.0.4/32 [110/3012] via 10.0.3.2, 00:03:08, FastEthernet1/1
              [110/3012] via 10.0.0.2, 00:03:08, FastEthernet0/0
O 192.168.0.5/32 [110/1512] via 10.0.3.2, 00:03:21, FastEthernet1/1
O 203.0.113.0/24 [110/3001] via 10.0.3.2, 00:00:03, FastEthernet1/1
              [110/3001] via 10.0.0.2, 00:00:03, FastEthernet0/0
```

14) Configure a default static route on R4 to the Internet via the service provider at 203.0.113.2

```
R4(config)#ip route 0.0.0.0 0.0.0.0 203.0.113.2
```

15) Ensure that routers R1 to R5 learn via OSPF how to reach the Internet.

```
R4(config)#router ospf 1
R4(config-router)#default-information originate
```

16) Verify routers R1 to R5 have a route to the Internet.

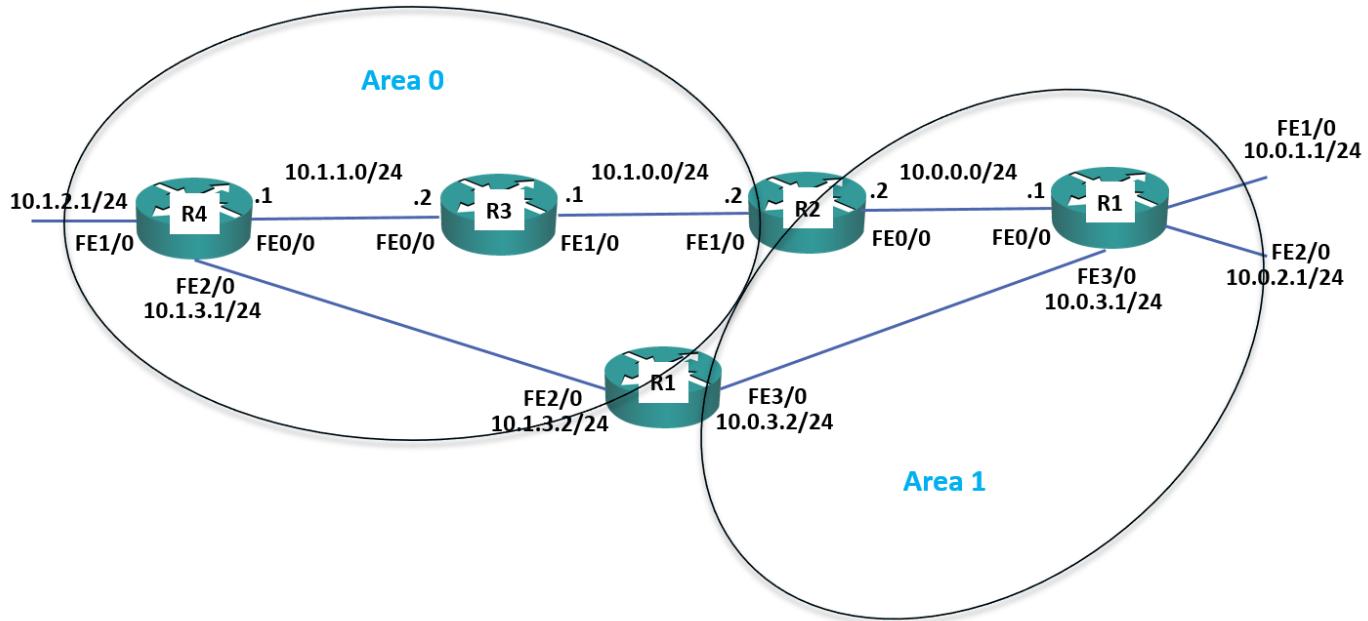
```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is 10.0.3.2 to network 0.0.0.0

```
 10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O 10.1.0.0/24 [110/2000] via 10.0.0.2, 00:10:54, FastEthernet0/0
O 10.1.1.0/24 [110/3000] via 10.0.0.2, 00:05:35, FastEthernet0/0
O 10.1.2.0/24 [110/4000] via 10.0.3.2, 00:05:22, FastEthernet1/1
              [110/4000] via 10.0.0.2, 00:05:22, FastEthernet0/0
O 10.1.3.0/24 [110/3000] via 10.0.3.2, 00:05:22, FastEthernet1/1
              192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O 192.168.0.2/32 [110/1012] via 10.0.0.2, 00:10:54, FastEthernet0/0
O 192.168.0.3/32 [110/2012] via 10.0.0.2, 00:10:54, FastEthernet0/0
O 192.168.0.4/32 [110/3012] via 10.0.3.2, 00:05:22, FastEthernet1/1
              [110/3012] via 10.0.0.2, 00:05:22, FastEthernet0/0
O 192.168.0.5/32 [110/1512] via 10.0.3.2, 00:05:35, FastEthernet1/1
O 203.0.113.0/24 [110/3001] via 10.0.3.2, 00:02:17, FastEthernet1/1
              [110/3001] via 10.0.0.2, 00:02:17, FastEthernet0/0
O*E2 0.0.0.0/0 [110/1] via 10.0.3.2, 00:00:06, FastEthernet1/1
              [110/1] via 10.0.0.2, 00:00:06, FastEthernet0/0
```

Multi-Area OSPF

17) Convert the network to use multi-area OSPF. R3 and R4 should be backbone routers, R1 a normal router in Area 1, and R2 and R5 ABRs as shown in the diagram below.
Save your changes to the startup config and reboot the routers to ensure the changes take effect.



R3 and R4 require no change as all their interfaces are already in Area 0.

R1's interfaces need to be reconfigured to be in Area 1 rather than Area 0.

```
R1#show run | section ospf
  ip ospf cost 1500
  router ospf 1
    log-adjacency-changes
    auto-cost reference-bandwidth 100000
    network 10.0.0.0 0.255.255.255 area 0
    network 192.168.0.0 0.0.0.255 area 0
```

```
R1(config)#router ospf 1
R1(config-router)#network 10.0.0.0 0.255.255.255 area 1
R1(config-router)#network 192.168.0.0 0.0.0.255 area 1
R1#copy run start
R1#reload
```

R2 interface FastEthernet 0/1 should remain in Area 0. FastEthernet 0/0 needs to be reconfigured to be in Area 1. I used a 10.0.0.0/8 network statement originally so I need to remove that and add more granular statements.

```
R2#sh run | section ospf
router ospf 1
  log-adjacency-changes
  auto-cost reference-bandwidth 100000
  network 10.0.0.0 0.255.255.255 area 0
  network 192.168.0.0 0.0.0.255 area 0

R2(config)#router ospf 1
R2(config-router)#no network 10.0.0.0 0.255.255.255 area 0
R2(config-router)#network 10.1.0.0 0.0.0.255 area 0
R2(config-router)#network 10.0.0.0 0.0.0.255 area 1
R2#copy run start
R2#reload
```

R5 interface FastEthernet 0/0 should remain in Area 0. FastEthernet 0/1 needs to be reconfigured to be in Area 1.

```
R5#sh run | section ospf
  ip ospf cost 1500
  ip ospf cost 1500
router ospf 1
  log-adjacency-changes
  auto-cost reference-bandwidth 100000
  network 10.0.0.0 0.255.255.255 area 0
  network 192.168.0.0 0.0.0.255 area 0

R5(config)#router ospf 1
R5(config-router)#no network 10.0.0.0 0.255.255.255 area 0
R5(config-router)#network 10.1.3.0 0.0.0.255 area 0
R5(config-router)#network 10.0.3.0 0.0.0.255 area 1
R5#copy run start
R5#reload
```

18) Verify the router's interfaces are in the correct areas.

```
R2#show ip ospf interface
```

```
Loopback0 is up, line protocol is up
  Internet address is 192.168.0.2/32, Area 0
FastEthernet0/1 is up, line protocol is up
  Internet address is 10.1.0.2/24, Area 0
FastEthernet0/0 is up, line protocol is up
  Internet address is 10.0.0.2/24, Area 1
! Output truncated
```

19) Verify routers R1 to R5 have formed adjacencies with each other.

```
R1#sh ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.0.5	1	FULL/DR	00:00:33	10.0.3.2	FastEthernet1/1
192.168.0.2	1	FULL/DR	00:00:31	10.0.0.2	FastEthernet0/0

20) What change do you expect to see on R1's routing table? Verify this (give the routing table a few seconds to converge).

The networks beyond R2 and R5 will appear as Inter Area routes (apart from the default route which will appear as an external route as it was redistributed into OSPF).

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override

Gateway of last resort is 10.0.0.2 to network 0.0.0.0

      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O IA 10.1.0.0/24 [110/2000] via 10.0.0.2, 00:03:10, FastEthernet0/0
O IA 10.1.1.0/24 [110/3000] via 10.0.0.2, 00:03:10, FastEthernet0/0
O IA 10.1.2.0/24 [110/4000] via 10.0.0.2, 00:03:10, FastEthernet0/0
O IA 10.1.3.0/24 [110/4500] via 10.0.0.2, 00:02:59, FastEthernet0/0
      192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O IA 192.168.0.2/32 [110/1012] via 10.0.0.2, 00:03:10, FastEthernet0/0
O IA 192.168.0.3/32 [110/2012] via 10.0.0.2, 00:03:10, FastEthernet0/0
O IA 192.168.0.4/32 [110/3012] via 10.0.0.2, 00:03:10, FastEthernet0/0
O IA 192.168.0.5/32 [110/4512] via 10.0.0.2, 00:02:59, FastEthernet0/0
O IA 203.0.113.0/24 [110/3001] via 10.0.0.2, 00:03:10, FastEthernet0/0
O*E2 0.0.0.0/0 [110/1] via 10.0.0.2, 00:02:59, FastEthernet0/0
```

21) Do you see less routes in R1's routing table? Why or why not?

R1 has the same amount of routes in its routing table because OSPF does not perform automatic summarisation. You must configure manual summarisation to reduce the size of the routing table.

22) Configure summary routes on the Area Border Routers for the 10.0.0.0/16 and 10.1.0.0/16 networks.

```
R2(config)#router ospf 1
R2(config-router)#area 0 range 10.1.0.0 255.255.0.0
R2(config-router)#area 1 range 10.0.0.0 255.255.0.0

R5(config-if)#router ospf 1
R5(config-router)#area 0 range 10.1.0.0 255.255.0.0
R5(config-router)#area 1 range 10.0.0.0 255.255.0.0
```

23) Verify R1 now sees a single summary route for 10.1.0.0/16 rather than individual routes for the 10.1.x.x networks.

```
R1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override

Gateway of last resort is 10.0.0.2 to network 0.0.0.0

      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
C 10.0.0.0/24 is directly connected, FastEthernet0/0
L 10.0.0.1/32 is directly connected, FastEthernet0/0
C 10.0.1.0/24 is directly connected, FastEthernet0/1
L 10.0.1.1/32 is directly connected, FastEthernet0/1
C 10.0.2.0/24 is directly connected, FastEthernet1/0
L 10.0.2.1/32 is directly connected, FastEthernet1/0
C 10.0.3.0/24 is directly connected, FastEthernet1/1
L 10.0.3.1/32 is directly connected, FastEthernet1/1
O IA 10.1.0.0/24 [110/2000] via 10.0.0.2, 00:00:04, FastEthernet0/0
      192.168.0.0/32 is subnetted, 5 subnets
C 192.168.0.1/32 is directly connected, Loopback0
O IA 192.168.0.2/32 [110/1012] via 10.0.0.2, 00:05:27, FastEthernet0/0
O IA 192.168.0.3/32 [110/2012] via 10.0.0.2, 00:05:27, FastEthernet0/0
O IA 192.168.0.4/32 [110/3012] via 10.0.0.2, 00:05:27, FastEthernet0/0
O IA 192.168.0.5/32 [110/4512] via 10.0.0.2, 00:05:16, FastEthernet0/0
O IA 203.0.113.0/24 [110/3001] via 10.0.0.2, 00:05:27, FastEthernet0/0
O*E2 0.0.0.0/0 [110/1] via 10.0.0.2, 00:05:16, FastEthernet0/0
```

24) Verify R1 is receiving a summary route for the 10.1.0.0/16 network from both R2 and R5.

R1#sh ip ospf database

OSPF Router with ID (192.168.0.1) (Process ID 1)

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link
count					
192.168.0.1	192.168.0.1	18	0x80000005	0x00536E	5
192.168.0.2	192.168.0.2	27	0x80000003	0x0069ED	1
192.168.0.5	192.168.0.5	1890	0x80000003	0x00C490	1

Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.0.1	192.168.0.1	18	0x80000002	0x00DF0E
10.0.3.1	192.168.0.1	18	0x80000002	0x00E8FE

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.0.5	192.168.0.5	408	0x80000006	0x00e987
10.0.3.0	192.168.0.5	408	0x80000007	0x007e7d
192.168.0.4	192.168.0.5	408	0x80000009	0x00bbd1
203.0.113.0	192.168.0.5	408	0x8000000a	0x00ebdb
192.168.0.3	192.168.0.5	408	0x8000000b	0x00f4ab
192.168.0.2	192.168.0.5	408	0x8000000c	0x003084
10.0.0.0	192.168.0.5	408	0x8000000d	0x002c09
192.168.0.2	192.168.0.2	1079	0x80000005	0x001c5c
192.168.0.3	192.168.0.2	1079	0x80000006	0x004446
192.168.0.4	192.168.0.2	1079	0x80000008	0x006932
203.0.113.0	192.168.0.2	1079	0x80000009	0x00993c
192.168.0.5	192.168.0.2	398	0x80000006	0x00308a
10.0.3.0	192.168.0.2	398	0x80000007	0x00c282
192.168.0.1	192.168.0.5	393	0x80000010	0x003c74
10.1.0.0	192.168.0.5	82	0x80000015	0x007679
10.1.0.0	192.168.0.2	67	0x8000001f	0x000ee4

Summary ASB Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.0.4	192.168.0.2	27	0x80000002	0x00EE9B
192.168.0.4	192.168.0.5	1889	0x80000002	0x00433A

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	192.168.0.4	207	0x80000002	0x00152F	1

25) R1 is routing traffic to 10.1.0.0/16 via R2 only. Why is it not load balancing the traffic through both R2 and R5?

We configured the link from R1 to R5 to have a higher cost than the link from R1 to R2 earlier.

```
R1#sh run | begin interface FastEthernet1/1
Building configuration...
Current configuration : 100 bytes
!
interface FastEthernet1/1
  ip address 10.0.3.1 255.255.255.0
  ip ospf cost 1500
```

DR and BDR Designated Routers

26) Enable a loopback interface on routers R6 to R9. Use the IP address 192.168.0.x/32, where 'x' is the router number. For example 192.168.0.6/32 on R6.

On routers R6 to R9:

```
R6(config)#interface loopback0
R6(config-if)#ip address 192.168.0.6 255.255.255.255
```

27) Enable OSPF for Area 0 on the Loopback 0 and FastEthernet 0/0 interfaces on routers R6 to R9.

On routers R6 to R9:

```
R1(config)#router ospf 1
R1(config-router)#network 172.16.0.0 0.0.0.255 area 0
R1(config-router)# network 192.168.0.0 0.0.0.255 area 0
```

You can use different network statements, as long as they cover the range of IP addresses configured on the router interfaces.

28) Set the reference bandwidth on routers R6 to R9 so that a 100 Gbps interface will have a cost of 1.

Remember to do this on all routers R6 to R9.

```
R6(config)#router ospf 1
R6(config-router)#auto-cost reference-bandwidth 100000
```

29) Which routers do you expect to be the DR and BDR on the Ethernet segment? Verify this.

OSPF priority has not been set so all routers will have the default of 1. R9 and R8 will be elected as the DR and BDR respectively because they have the highest Router IDs (because they have the highest IP addresses on their loopback interfaces).

```
R6#show ip ospf interface FastEthernet 0/0
```

```
FastEthernet0/0 is up, line protocol is up
  Internet address is 172.16.0.6/24, Area 0
  Process ID 1, Router ID 192.168.0.6, Network Type BROADCAST, Cost: 1000
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 192.168.0.9, Interface address 172.16.0.9
  Backup Designated Router (ID) 192.168.0.8, Interface address 172.16.0.8
```

```
R6#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.0.8	1	FULL/BDR	00:00:31	172.16.0.8	FastEthernet0/0
192.168.0.7	1	2WAY/DROTHER	00:00:39	172.16.0.7	FastEthernet0/0
192.168.0.9	1	FULL/DR	00:00:39	172.16.0.9	FastEthernet0/0

```
R9#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.0.8	1	FULL/BDR	00:00:31	172.16.0.8	FastEthernet0/0
192.168.0.7	1	2WAY/DROTHER	00:00:39	172.16.0.7	FastEthernet0/0
192.168.0.6	1	FULL/DROTHER	00:00:39	172.16.0.6	FastEthernet0/0

30) Set R6 as the Designated Router without changing any IP addresses.

Configure a higher OSPF priority on R6.

```
R6(config)#interface FastEthernet0/0
R6(config-if)#ip ospf priority 100
R6(config-if)#end
R6#clear ip ospf process
```

31) Verify R6 is the Designated Router.

```
R6#show ip ospf interface FastEthernet 0/0
```

```
FastEthernet0/0 is up, line protocol is up
Internet address is 172.16.0.6/24, Area 0
Process ID 1, Router ID 192.168.0.6, Network Type BROADCAST, Cost: 1000
Transmit Delay is 1 sec, State DR, Priority 100
Designated Router (ID) 192.168.0.6, Interface address 172.16.0.6
Backup Designated Router (ID) 192.168.0.8, Interface address 172.16.0.8
```

```
R6#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.0.8	1	FULL/BDR	00:00:31	172.16.0.8	FastEthernet0/0
192.168.0.7	1	2WAY/DROTHER	00:00:39	172.16.0.7	FastEthernet0/0
192.168.0.9	1	FULL/DROTHER	00:00:39	172.16.0.9	FastEthernet0/0

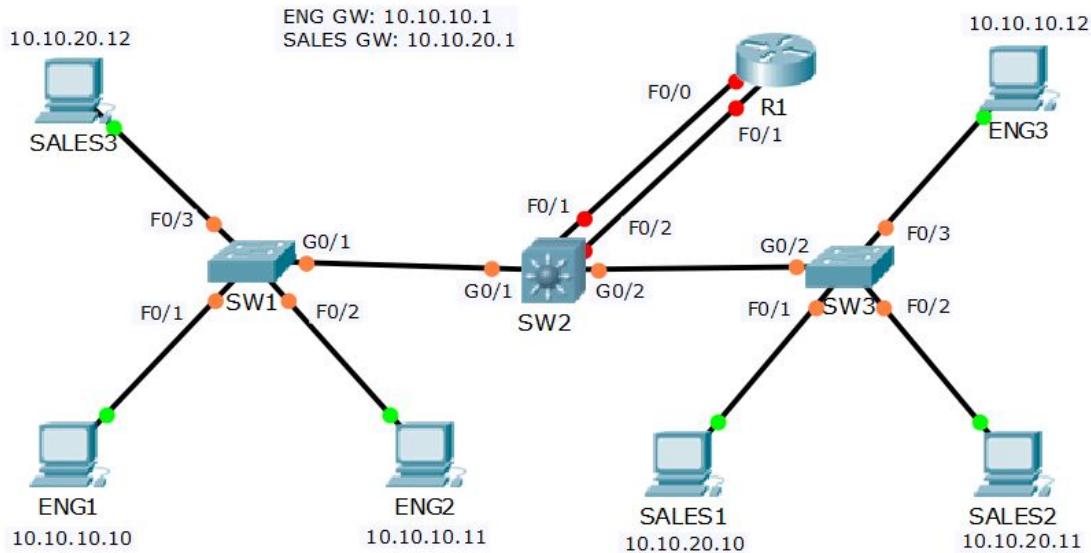
```
R9#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.0.8	1	FULL/BDR	00:00:31	172.16.0.8	FastEthernet0/0
192.168.0.7	1	2WAY/DROTHER	00:00:39	172.16.0.7	FastEthernet0/0
192.168.0.6	1	FULL/DR	00:00:39	172.16.0.6	FastEthernet0/0

22-1 VLAN and Inter-VLAN Routing Configuration – Lab Exercise

In this lab you will perform a VLAN configuration for a campus network, including Virtual Trunking Protocol, Access and Trunk ports, and inter-VLAN routing.

Lab Topology



Load the Startup Configurations

[Download the '22-1 VLAN and Inter-VLAN Routing Configuration.zip' file here.](#)
Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

VTP, Access and Trunk Ports

- 1) All routers and switches are in a factory default state. View the VLAN database on SW1 to verify no VLANs have been added.
- 2) View the default switchport status on the link from SW1 to SW2.
- 3) Configure the links between switches as trunks.
- 4) Configure SW1 as a VTP Server in the VTP domain Flackbox.
- 5) SW2 must not synchronise its VLAN database with SW1.
- 6) SW3 must learn VLAN information from SW1. VLANs should not be edited on SW3.
- 7) Add the Eng, Sales and Native VLANs on all switches.
- 8) Verify the VLANs are in the database on each switch.
- 9) Configure the trunk links to use VLAN 199 as the native VLAN for better security.
- 10) Configure the switchports connected to the PCs with the correct VLAN configuration.
- 11) Verify the Eng1 PC has connectivity to Eng3.
- 12) Verify Sales1 has connectivity to Sales3.

Inter-VLAN Routing – Option 1

Separate Interfaces on Router

- 13) Configure interface FastEthernet0/0 on R1 as the default gateway for the Eng PCs.
- 14) Configure interface FastEthernet0/1 on R1 as the default gateway for the Sales PCs.
- 15) Configure SW2 to support inter-VLAN routing using R1 as the default gateway.
- 16) Verify the Eng1 PC has connectivity to the VLAN 20 interface on R1.
- 17) Verify the Eng1 PC has connectivity to Sales1.
- 18) Clean-up: Shut down interface FastEthernet0/1 on R1.

Inter-VLAN Routing – Option 2

Router on a Stick

- 19) Configure sub-interfaces on FastEthernet0/0 on R1 as the default gateway for the Eng and Sales PCs.
- 20) Configure SW2 to support inter-VLAN routing using R1 as the default gateway.
- 21) Verify the Eng1 PC has connectivity to the VLAN 20 interface on R1.
- 22) Verify the Eng1 PC has connectivity to Sales1.
- 23) Clean-up: Shut down interface FastEthernet0/0 on R1.

Inter-VLAN Routing – Option 3

Layer 3 Switch

- 24) Enable layer 3 routing on SW2.
- 25) Configure SVIs on SW2 to support inter-VLAN routing between the Eng and Sales VLANs.
- 26) Verify the Eng1 PC has connectivity to the VLAN 20 interface on SW2.
- 27) Verify the Eng1 PC has connectivity to Sales1.

22-1 VLAN and Inter-VLAN Routing Configuration - Answer Key

In this lab you will perform a VLAN configuration for a campus network, including Virtual Trunking Protocol, Access and Trunk ports, and inter-VLAN routing.

VTP, Access and Trunk Ports

- 1) All routers and switches are in a factory default state. View the VLAN database on SW1 to verify no VLANs have been added.

```
SW1#show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Fa0/1, Fa0/2, Fa0/3, Fa0/4 Fa0/5, Fa0/6, Fa0/7, Fa0/8 Fa0/9, Fa0/10, Fa0/11, Fa0/12 Fa0/13, Fa0/14, Fa0/15, Fa0/16 Fa0/17, Fa0/18, Fa0/19, Fa0/20 Fa0/21, Fa0/22, Fa0/23, Fa0/24 Gig0/1, Gig0/2
1002	fdci-default	active	
1003	token-ring-default	active	
1004	fdinnet-default	active	
1005	trnet-default	active	

- 2) View the default switchport status on the link from SW1 to SW2.

```
SW1#show interface gig0/1 switchport
Name: Gig0/1
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
Protected: false
Appliance trust: none
```

The trunking mode is set to dynamic auto and the interface is currently in the access port operational mode using the default VLAN 1.

- 3) Configure the links between switches as trunks.

```
SW1(config)#int g0/1
SW1(config-if)#switch mode trunk

SW2(config)#int g0/1
SW2(config-if)#switch trunk encapsulation dot1q
SW2(config-if)#switch mode trunk
SW2(config-if)#int g0/2
SW2(config-if)#switch trunk encapsulation dot1q
SW2(config-if)#switch mode trunk

SW3(config)#int g0/2
SW3(config-if)#switch mode trunk
```

- 4) Configure SW1 as a VTP Server in the VTP domain Flackbox.

```
SW1(config)#vtp domain Flackbox
Changing VTP domain name from NULL to Flackbox
SW1(config)#vtp mode server
Device mode already VTP SERVER.
```

- 5) SW2 must not synchronise its VLAN database with SW1.

```
SW2(config)#vtp mode transparent
Setting device to VTP TRANSPARENT mode.
```

- 6) SW3 must learn VLAN information from SW1. VLANs should not be edited on SW3.

```
SW3(config)#vtp mode client
Setting device to VTP CLIENT mode.
SW3(config)#vtp domain Flackbox
Changing VTP domain name from NULL to Flackbox
```

7) Add the Eng, Sales and Native VLANs on all switches.

VLANs must be configured on the VTP Server SW1 and on VTP Transparent SW2. VTP Client SW3 will learn the VLANs from SW1.

```
SW1(config)#vlan 10
SW1(config-vlan)#name Eng
SW1(config-vlan)#vlan 20
SW1(config-vlan)#name Sales
SW1(config-vlan)#vlan 199
SW1(config-vlan)#name Native
```

```
SW2(config)#vlan 10
SW2(config-vlan)#name Eng
SW2(config-vlan)#vlan 20
SW2(config-vlan)#name Sales
SW2(config-vlan)#vlan 199
SW2(config-vlan)#name Native
```

8) Verify the VLANs are in the database on each switch.

```
SW3#sh vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Fa0/1, Fa0/2, Fa0/3, Fa0/4 Fa0/5, Fa0/6, Fa0/7, Fa0/8 Fa0/9, Fa0/10, Fa0/11, Fa0/12 Fa0/13, Fa0/14, Fa0/15, Fa0/16 Fa0/17, Fa0/18, Fa0/19, Fa0/20 Fa0/21, Fa0/22, Fa0/23, Fa0/24 Gig0/1
10	Eng	active	
20	Sales	active	
199	Native	active	
1002	fdci-default	active	
1003	token-ring-default	active	
1004	fdinnet-default	active	
1005	trnet-default	active	

9) Configure the trunk links to use VLAN 199 as the native VLAN for better security.

```
SW1(config)#interface gig0/1
SW1(config-if)#switch trunk native vlan 199

SW2(config)#int gig0/1
SW2(config-if)#switch trunk native vlan 199
SW2(config-if)#int gig0/2
SW2(config-if)#switch trunk native vlan 199

SW3(config)#int gig0/2
SW3(config-if)#switch trunk native vlan 199
```

10) Configure the switchports connected to the PCs with the correct VLAN configuration.

Eng PCs should be in VLAN 10, Sales PCs in VLAN 20.

```
SW1(config)#int range f0/1 - 2
SW1(config-if-range)#switch mode access
SW1(config-if-range)#switch access vlan 10
SW1(config-if-range)#int f0/3
SW1(config-if)#switch mode access
SW1(config-if)#switch access vlan 20

SW3(config)#int range f0/1 - 2
SW3(config-if-range)#switch mode access
SW3(config-if-range)#switch access vlan 20
SW3(config-if-range)#int f0/3
SW3(config-if)#switch mode access
SW3(config-if)#switch access vlan 10
```

11) Verify the Eng1 PC has connectivity to Eng3.

```
C:\>ping 10.10.10.12

Pinging 10.10.10.12 with 32 bytes of data:

Reply from 10.10.10.12: bytes=32 time<1ms TTL=128
Reply from 10.10.10.12: bytes=32 time<1ms TTL=128
Reply from 10.10.10.12: bytes=32 time<1ms TTL=128
Reply from 10.10.10.12: bytes=32 time=1ms TTL=128

Ping statistics for 10.10.10.12:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

12) Verify Sales1 has connectivity to Sales3.

```
C:\>ping 10.10.20.12

Pinging 10.10.20.12 with 32 bytes of data:

Reply from 10.10.20.12: bytes=32 time=1ms TTL=128
Reply from 10.10.20.12: bytes=32 time<1ms TTL=128
Reply from 10.10.20.12: bytes=32 time<1ms TTL=128
Reply from 10.10.20.12: bytes=32 time<1ms TTL=128

Ping statistics for 10.10.20.12:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

Inter-VLAN Routing – Option 1

Separate Interfaces on Router

13) Configure interface FastEthernet0/0 on R1 as the default gateway for the Eng PCs.

```
R1(config)#interface FastEthernet 0/0
R1(config-if)#ip address 10.10.10.1 255.255.255.0
R1(config-if)#no shutdown
```

14) Configure interface FastEthernet0/1 on R1 as the default gateway for the Sales PCs.

```
R1(config)#interface FastEthernet 0/1
R1(config-if)#ip address 10.10.20.1 255.255.255.0
R1(config-if)#no shutdown
```

15) Configure SW2 to support inter-VLAN routing using R1 as the default gateway.

```
SW2(config)#interface FastEthernet 0/1
SW2(config-if)#switchport mode access
SW2(config-if)#switchport access vlan 10
SW2(config-if)#interface FastEthernet 0/2
SW2(config-if)#switchport mode access
SW2(config-if)#switchport access vlan 20
```

16) Verify the Eng1 PC has connectivity to the VLAN 20 interface on R1.

```
C:\>ping 10.10.20.1

Pinging 10.10.20.1 with 32 bytes of data:

Reply from 10.10.20.1: bytes=32 time<1ms TTL=255
Reply from 10.10.20.1: bytes=32 time=1ms TTL=255
Reply from 10.10.20.1: bytes=32 time<1ms TTL=255
Reply from 10.10.20.1: bytes=32 time=1ms TTL=255

Ping statistics for 10.10.20.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

17) Verify the Eng1 PC has connectivity to Sales1.

```
C:\>ping 10.10.20.10

Pinging 10.10.20.10 with 32 bytes of data:

Request timed out.
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127

Ping statistics for 10.10.20.10:
    Packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

18) Clean-up: Shut down interface FastEthernet0/1 on R1.

```
R1(config)#int f0/1
R1(config-if)#shutdown
```

Inter-VLAN Routing – Option 2

Router on a Stick

19) Configure sub-interfaces on FastEthernet0/0 on R1 as the default gateway for the Eng and Sales PCs.

```
R1(config)#interface FastEthernet 0/0
R1(config-if)#no ip address
R1(config-if)#no shutdown
R1(config-if)#interface FastEthernet 0/0.10
R1(config-subif)#encapsulation dot1q 10
R1(config-subif)#ip address 10.10.10.1 255.255.255.0
R1(config-subif)#interface FastEthernet 0/0.20
R1(config-subif)#encapsulation dot1q 20
R1(config-subif)#ip address 10.10.20.1 255.255.255.0
```

20) Configure SW2 to support inter-VLAN routing using R1 as the default gateway.

```
SW2(config)#interface FastEthernet 0/1
SW2(config-if)#switch trunk encap dot1q
SW2(config-if)#switchport mode trunk
```

21) Verify the Eng1 PC has connectivity to the VLAN 20 interface on R1.

```
C:\>ping 10.10.20.1

Pinging 10.10.20.1 with 32 bytes of data:

Reply from 10.10.20.1: bytes=32 time<1ms TTL=255
Reply from 10.10.20.1: bytes=32 time=1ms TTL=255
Reply from 10.10.20.1: bytes=32 time<1ms TTL=255
Reply from 10.10.20.1: bytes=32 time=1ms TTL=255

Ping statistics for 10.10.20.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

22) Verify the Eng1 PC has connectivity to Sales1.

```
C:\>ping 10.10.20.10

Pinging 10.10.20.10 with 32 bytes of data:

Request timed out.
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127

Ping statistics for 10.10.20.10:
    Packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

23) Clean-up: Shut down interface FastEthernet0/0 on R1.

```
R1(config)#int f0/0
R1(config-if)#shutdown
```

Inter-VLAN Routing – Option 3

Layer 3 Switch

24) Enable layer 3 routing on SW2.

```
SW2(config)#ip routing
```

25) Configure SVIs on SW2 to support inter-VLAN routing between the Eng and Sales VLANs.

```
SW2(config)#interface vlan 10
SW2(config-if)#ip address 10.10.10.1 255.255.255.0
SW2(config-if)#interface vlan 20
SW2(config-if)#ip address 10.10.20.1 255.255.255.0
```

26) Verify the Eng1 PC has connectivity to the VLAN 20 interface on SW2.

```
C:\>ping 10.10.20.1

Pinging 10.10.20.1 with 32 bytes of data:

Reply from 10.10.20.1: bytes=32 time<1ms TTL=255
Reply from 10.10.20.1: bytes=32 time=1ms TTL=255
Reply from 10.10.20.1: bytes=32 time<1ms TTL=255
Reply from 10.10.20.1: bytes=32 time=1ms TTL=255

Ping statistics for 10.10.20.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

27) Verify the Eng1 PC has connectivity to Sales1.

```
C:\>ping 10.10.20.10

Pinging 10.10.20.10 with 32 bytes of data:

Request timed out.
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127

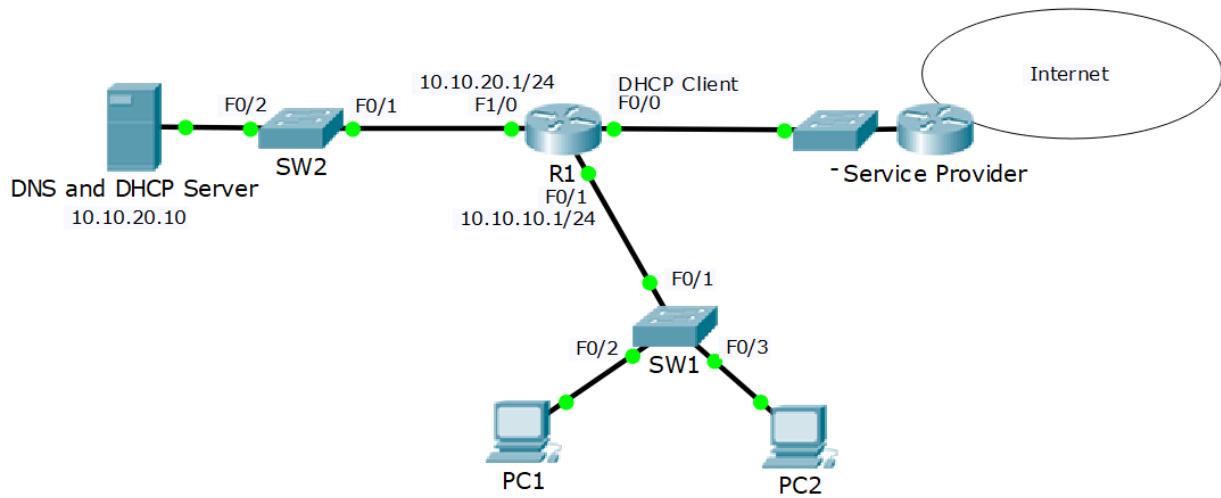
Ping statistics for 10.10.20.10:
    Packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

23-1 DHCP Configuration – Lab Exercise

In this lab you will perform a DHCP configuration for a small campus network. You will configure a router's outside interface as a DHCP client. You will then set up DHCP services, using a Cisco router first and then an external DHCP server. The external DHCP server is inside the campus LAN but outside the router.

Note that the external DHCP server at 10.10.20.10 will not be used until the last part of the lab.

Lab Topology



Load the Startup Configurations

[Download the '23-1 DHCP Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Cisco DHCP Client

- 1) You have not acquired a static public IP address from the Internet service provider. Configure the outside interface FastEthernet 0/0 on R1 to receive its IP address using DHCP. The Service Provider is already configured and you have no access to it.
- 2) Verify that R1 received its public IP address via DHCP (you may need to wait a few minutes for the address to be assigned).
- 3) What is the IP address of R1's DHCP server?

Cisco DHCP Server

- 4) Enable the DHCP service on R1 so it gives out IP addresses to the PCs in the 10.10.10.0/24 subnet. Leave IP addresses 10.10.10.1 – 10 free to be assigned to servers and printers. 10.10.20.10 is the DNS server.
- 5) Verify the clients received their IP information via DHCP.
- 6) Verify the clients can ping the DNS server by its hostname 'DNSserver' (it might take some time for DNS to resolve the hostname).
- 7) On R1, verify both clients received an IP address via DHCP.
- 8) Cleanup – remove the DHCP server configuration on R1. You will use an external DHCP server instead in the next section.
- 9) Enter the command 'ipconfig /release' on the PCs to release their IP addresses.
- 10) Enter the command 'ipconfig /renew' on the PCs and verify they can no longer obtain an IP address via DHCP

External DHCP Server

- 11) The server at 10.10.20.10 has been configured as a DHCP server with a scope of IP addresses for the 10.10.10.0/24 subnet, but the PCs there are not receiving IP addresses. Why is this?

12) Configure the network to allow the PCs to receive their IP addresses from the DHCP server.

13) Verify the clients received their IP information via DHCP.

23-1 DHCP Configuration - Answer Key

In this lab you will perform a DHCP configuration for a small campus network. You will configure a router's outside interface as a DHCP client. You will then set up DHCP services, using a Cisco router first and then an external DHCP server. The external DHCP server is inside the campus LAN but outside the router.

Note that the external DHCP server at 10.10.20.10 will not be used until the last part of the lab.

Cisco DHCP Client

- 1) You have not acquired a static public IP address from the Internet service provider. Configure the outside interface FastEthernet 0/0 on R1 to receive its IP address using DHCP. The Service Provider is already configured and you have no access to it.

```
R1(config)#interface f0/0
R1(config-if)#ip address dhcp
R1(config-if)#no shutdown
```

- 2) Verify that R1 received its public IP address via DHCP (you may need to wait a few minutes for the address to be assigned).

```
%DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/0 assigned
DHCP address 203.0.113.2, mask 255.255.255.0, hostname R1
```

```
R1#sh ip int brief
Interface IP-Address OK? Method Status Protocol
FastEthernet0/0 203.0.113.2 YES DHCP up up
FastEthernet0/1 unassigned YES unset administratively down down
FastEthernet1/0 unassigned YES unset administratively down down
FastEthernet1/1 unassigned YES unset administratively down down
Vlan1 unassigned YES unset administratively down down
```

3) What is the IP address of R1's DHCP server?

The DHCP server is at 203.0.113.1. We can get this information by viewing the DHCP lease information.

```
R1#show dhcp lease
Temp IP addr: 203.0.113.2 for peer on Interface:
FastEthernet0/0
Temp sub net mask: 255.255.255.0
DHCP Lease server: 203.0.113.1 , state: Bound
DHCP Transaction id: 64B8EE07
Lease: 86400 secs, Renewal: 43200 secs, Rebind: 75600 secs
Temp default-gateway addr: 203.0.113.1
Next timer fires after: 11:53:13
Retry count: 0 Client-ID:cisco-0001.63C2.9701-Fa0/0
Client-ID hex dump: 636973636F2D303030312E363343322E
93730312D4661302F30
Hostname: R1
```

Cisco DHCP Server

4) Enable the DHCP service on R1 so it gives out IP addresses to the PCs in the 10.10.10.0/24 subnet. Leave IP addresses 10.10.10.1 – 10 free to be assigned to servers and printers. 10.10.20.10 is the DNS server.

```
R1(config)#ip dhcp excluded-address 10.10.10.1 10.10.10.10
R1(config)#ip dhcp pool Flackbox
R1(dhcp-config)#network 10.10.10.0 255.255.255.0
R1(dhcp-config)#default-router 10.10.10.1
R1(dhcp-config)#dns-server 10.10.20.10
```

5) Verify the clients received their IP information via DHCP.

```
C:\>ipconfig /all

FastEthernet0 Connection: (default port)

  Connection-specific DNS Suffix...:
  Physical Address.....: 0000.0CA0.A359
  Link-local IPv6 Address.....:
  FE80::200:CFF:FEA0:A359
    IP Address.....: 10.10.10.12
    Subnet Mask.....: 255.255.255.0
    Default Gateway.....: 10.10.10.1
    DNS Servers.....: 10.10.20.10
    DHCP Servers.....: 10.10.10.1
    DHCPv6 IAID.....: 28262
    DHCPv6 Client DUID.....:
  00-01-00-01-61-91-76-98-00-0C-A0-A3-59
```

6) Verify the clients can ping the DNS server by its hostname 'DNSserver' (it might take some time for DNS to resolve the hostname).

```
C:\>ping dnsserver

Pinging 10.10.20.10 with 32 bytes of data:

Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time=1ms TTL=127
Reply from 10.10.20.10: bytes=32 time<1ms TTL=127
Reply from 10.10.20.10: bytes=32 time=1ms TTL=127

Ping statistics for 10.10.20.10:
  Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
  Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

7) On R1, verify both clients received an IP address via DHCP.

R1#show ip dhcp binding			
IP address	Client-ID/ Hardware address	Lease expiration	Type
10.10.10.12	0000.0CA0.A359	--	Automatic
10.10.10.11	0002.164E.D116	--	Automatic

8) Cleanup – remove the DHCP server configuration on R1. You will use an external DHCP server instead in the next section.

```
R1(config)#no ip dhcp excluded-address 10.10.10.1 10.10.10.10
R1(config)#no ip dhcp pool Flackbox
```

9) Enter the command ‘ipconfig /release’ on the PCs to release their IP addresses.

```
C:\>ipconfig /release

IP Address.....: 0.0.0.0
Subnet Mask.....: 0.0.0.0
Default Gateway.: 0.0.0.0
DNS Server.....: 0.0.0.0
```

10) Enter the command ‘ipconfig /renew’ on the PCs and verify they can no longer obtain an IP address via DHCP

```
C:\>ipconfig /renew
DHCP request failed.
```

External DHCP Server

11) The server at 10.10.20.10 has been configured as a DHCP server with a scope of IP addresses for the 10.10.10.0/24 subnet, but the PCs there are not receiving IP addresses. Why is this?

DHCP requests use broadcast traffic. R1 is not forwarding the requests on to the DHCP server as routers do not forward broadcast traffic by default.

12) Configure the network to allow the PCs to receive their IP addresses from the DHCP server.

On the interface where they are received, configure the router to forward DHCP requests to the server.

```
R1(config)#interface f0/1
R1(config-if)#ip helper-address 10.10.20.10
```

13) Verify the clients received their IP information via DHCP.

```
C:\>ipconfig /all

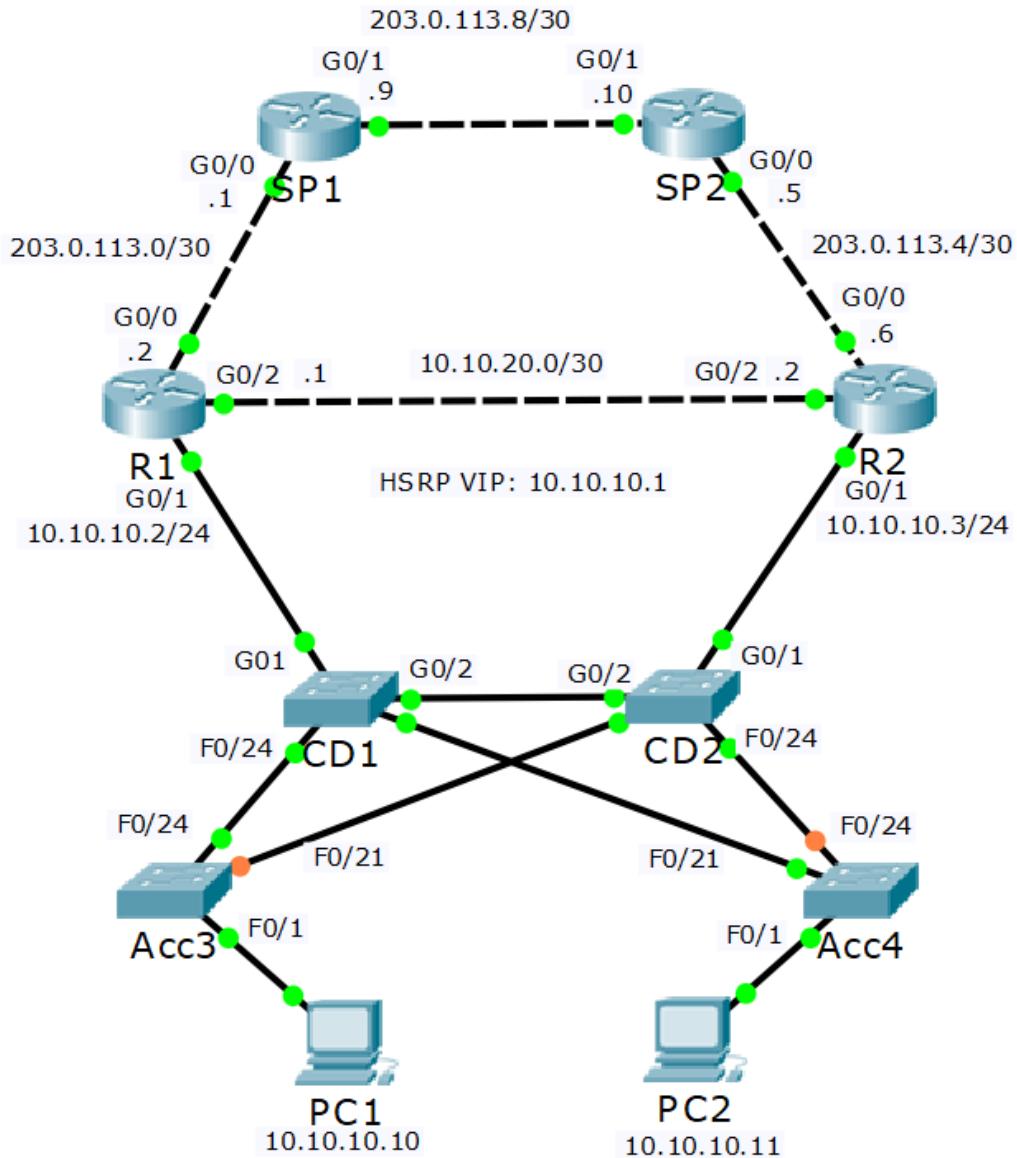
FastEthernet0 Connection: (default port)

  Connection-specific DNS Suffix...:
  Physical Address.....: 0000.0CA0.A359
  Link-local IPv6 Address.....:
  FE80::200:CFF:FEA0:A359
    IP Address.....: 10.10.10.101
    Subnet Mask.....: 255.255.255.0
    Default Gateway.....: 10.10.10.1
    DNS Servers.....: 10.10.20.10
    DHCP Servers.....: 10.10.20.10
    DHCPv6 IAID.....: 28262
    DHCPv6 Client DUID.....:
  00-01-00-01-61-91-76-98-00-00-0C-A0-A3-59
```

24-1 HSRP Configuration – Lab Exercise

In this lab you will configure and test HSRP for a small campus network.

Lab Topology



Load the Startup Configurations

[Download the '24-1 HSRP Configuration.zip' file here.](#)

Basic HSRP

- 1) Configure basic HSRP for the 10.10.10.0/24 network using the IP addresses shown in the topology diagram.
- 2) Wait for HSRP to come up on both routers and then check which is the active router.
- 3) Verify that the PCs can ping their default gateway using the HSRP address 10.10.10.1.
- 4) Verify that the PCs have upstream connectivity via their HSRP default gateway. Ping SP1 at 203.0.113.1
- 5) What is the MAC address on the physical interface of the active router?
- 6) What is the MAC address of the HSRP virtual interface?
- 7) Verify the PCs are using the virtual MAC address for their default gateway.

Priority and Pre-emption

- 8) Configure HSRP so that R1 will be the preferred router. Use a single command.
- 9) Which router do you expect will be active now? Verify this.
- 10) Ensure that R1 is the active router. Do not reboot.

Test HSRP

- 11) Run a continuous ping to the HSRP IP address from PC1 with the 'ping 10.10.10.1 -n 1000' command.
- 12) Save the configuration on R1 then reboot.
- 13) View the ping output on PC1. You should see a few dropped pings as R2 transitions to active following the outage of R1.

- 14) Verify R2 has transitioned to HSRP active.
- 15) Wait for R1 to complete booting and HSRP to come up. Verify R1 transitions to HSRP active because pre-emption is enabled.
- 16) Hit 'Ctrl-C' to cancel the ping on PC1. If you scroll back you should see a dropped ping or two as R1 transitioned back to HSRP active.

24-1 HSRP Configuration - Answer Key

In this lab you will configure and test HSRP for a small campus network.

Basic HSRP

- 1) Configure basic HSRP for the 10.10.10.0/24 network using the IP addresses shown in the topology diagram.

Check to see if the physical IP addresses for the HSRP interface GigabitEthernet0/1 have been configured on R1 and R2.

```
R1#sh ip int brief
Interface          IP-Address      OK? Method Status          Protocol
GigabitEthernet0/0 203.0.113.2    YES manual up           up
GigabitEthernet0/1 unassigned      YES unset  administratively down  down
GigabitEthernet0/2 10.10.20.1     YES manual up           up
Vlan1              unassigned      YES unset  administratively down  down
```

They haven't been configured yet so correct that. We'll do R1 first.

```
R1(config)#interface g0/1
R1(config-if)#ip address 10.10.10.2 255.255.255.0
R1(config-if)#no shutdown
```

Then add the virtual IP address.

```
R1(config-if)#standby 1 ip 10.10.10.1
```

Repeat on R2. Use a different physical address and the same virtual IP address.

```
R2(config)#interface g0/1
R2(config-if)#ip address 10.10.10.3 255.255.255.0
R2(config-if)#no shutdown
R2(config-if)#standby 1 ip 10.10.10.1
```

- 2) Wait for HSRP to come up on both routers and then check which is the active router.

Both routers have the same default priority so the router with the highest IP address will be preferred. If you complete the configuration on both routers before HSRP comes up on both then R2 will be active.
(If you configure HSRP on R1 before R2, and HSRP has already come up on R1 before HSRP comes up on R2, then R1 will be active.)

```
R1#show standby
GigabitEthernet0/1 - Group 1
State is Standby
7 state changes, last state change 00:16:34
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
Local virtual MAC address is 0000.0C07.AC01 (v1 default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 0.327 secs
Preemption disabled
Active router is 10.10.10.3
Standby router is local
Priority 100 (default 100)
Group name is hsrp-Gig0/1-1 (default)
```

- 3) Verify that the PCs can ping their default gateway using the HSRP address 10.10.10.1.

```
C:\>ping 10.10.10.1

Pinging 10.10.10.1 with 32 bytes of data:

Reply from 10.10.10.1: bytes=32 time=1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=43ms TTL=255
Reply from 10.10.10.1: bytes=32 time=1ms TTL=255
Reply from 10.10.10.1: bytes=32 time<1ms TTL=255

Ping statistics for 10.10.10.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 43ms, Average = 11ms
```

4) Verify that the PCs have upstream connectivity via their HSRP default gateway. Ping SP1 at 203.0.113.1

```
C:\>ping 203.0.113.1

Pinging 203.0.113.1 with 32 bytes of data:

Request timed out.
Reply from 203.0.113.1: bytes=32 time<1ms TTL=254
Reply from 203.0.113.1: bytes=32 time<1ms TTL=254
Reply from 203.0.113.1: bytes=32 time=1ms TTL=254

Ping statistics for 203.0.113.1:
    Packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

5) What is the MAC address on the physical interface of the active router?

```
R2#show interface g0/1
GigabitEthernet0/1 is up, line protocol is up (connected)
Hardware is CN Gigabit Ethernet, address is 0001.6470.2502
truncated
```

MAC addresses may be different in your lab.

6) What is the MAC address of the HSRP virtual interface?

```
R2#show standby
GigabitEthernet0/1 - Group 1
State is Active
6 state changes, last state change 00:16:22
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
Local virtual MAC address is 0000.0C07.AC01 (vl default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 2.475 secs
Preemption disabled
Active router is local
Standby router is 10.10.10.2
Priority 100 (default 100)
Group name is hsrp-Gig0/1-1 (default)
```

7) Verify the PCs are using the virtual MAC address for their default gateway.

```
C:\>arp -a
  Internet Address        Physical Address        Type
  10.10.10.1              0000.0c07.ac01      dynamic
  10.10.10.2              00d0.ffe2.2d02      dynamic
```

Priority and Pre-emption

8) Configure HSRP so that R1 will be the preferred router. Use a single command.

```
R1(config)#interface g0/1
R1(config-if)#standby 1 priority 110
```

9) Which router do you expect will be active now? Verify this.

R2 will remain active because we didn't configure pre-emption.

```
R1#show standby
GigabitEthernet0/1 - Group 1
State is Standby
7 state changes, last state change 00:06:06
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
Local virtual MAC address is 0000.0C07.AC01 (vl default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 0.782 secs
Preemption disabled
Active router is 10.10.10.3
Standby router is local
Priority 110 (configured 110)
Group name is hsrp-Gig0/1-1 (default)
```

10) Ensure that R1 is the active router. Do not reboot.

```
R1(config)#int g0/1
R1(config-if)#standby 1 preempt

R1#show standby
GigabitEthernet0/1 - Group 1
State is Active
9 state changes, last state change 00:27:09
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
Local virtual MAC address is 0000.0C07.AC01 (v1 default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 1.275 secs
Preemption enabled
Active router is local
Standby router is 10.10.10.3, priority 100 (expires in 8 sec)
Priority 110 (configured 110)
Group name is hsrp-Gig0/1-1 (default)
```

Test HSRP

11) Run a continuous ping to the HSRP IP address from PC1 with the 'ping 10.10.10.1 -n 1000' command.

```
C:\>ping 10.10.10.1 -n 1000

Pinging 10.10.10.1 with 32 bytes of data:

Reply from 10.10.10.1: bytes=32 time=13ms TTL=255
Reply from 10.10.10.1: bytes=32 time=32ms TTL=255
Reply from 10.10.10.1: bytes=32 time=2ms TTL=255
```

12) Save the configuration on R1 then reboot.

```
R1#copy run start
Destination filename [startup-config]?
Building configuration...
[OK]
R1#reload
Proceed with reload? [confirm]
```

13) View the ping output on PC1. You should see a few dropped pings as R2 transitions to active following the outage of R1.

```
Reply from 10.10.10.1: bytes=32 time<1ms TTL=255
Reply from 10.10.10.1: bytes=32 time<1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=15ms TTL=255
Request timed out.
Request timed out.
Reply from 10.10.10.1: bytes=32 time<1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=11ms TTL=255
Reply from 10.10.10.1: bytes=32 time=14ms TTL=255
Reply from 10.10.10.1: bytes=32 time=12ms TTL=255
```

14) Verify R2 has transitioned to HSRP active.

```
R2#show standby
GigabitEthernet0/1 - Group 1
State is Active
9 state changes, last state change 00:33:44
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
Local virtual MAC address is 0000.0C07.AC01 (vl default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 0.074 secs
Preemption disabled
Active router is local
Standby router is unknown, priority 110
Priority 100 (default 100)
Group name is hsrp-Gig0/1-1 (default)
```

15) Wait for R1 to complete booting and HSRP to come up. Verify R1 transitions to HSRP active because pre-emption is enabled.

```
R1#show standby
GigabitEthernet0/1 - Group 1
State is Active
5 state changes, last state change 00:00:38
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
Local virtual MAC address is 0000.0C07.AC01 (vl default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 2.533 secs
Preemption enabled
Active router is local
Standby router is 10.10.10.3
Priority 110 (configured 110)
Group name is hsrp-Gig0/1-1 (default)
```

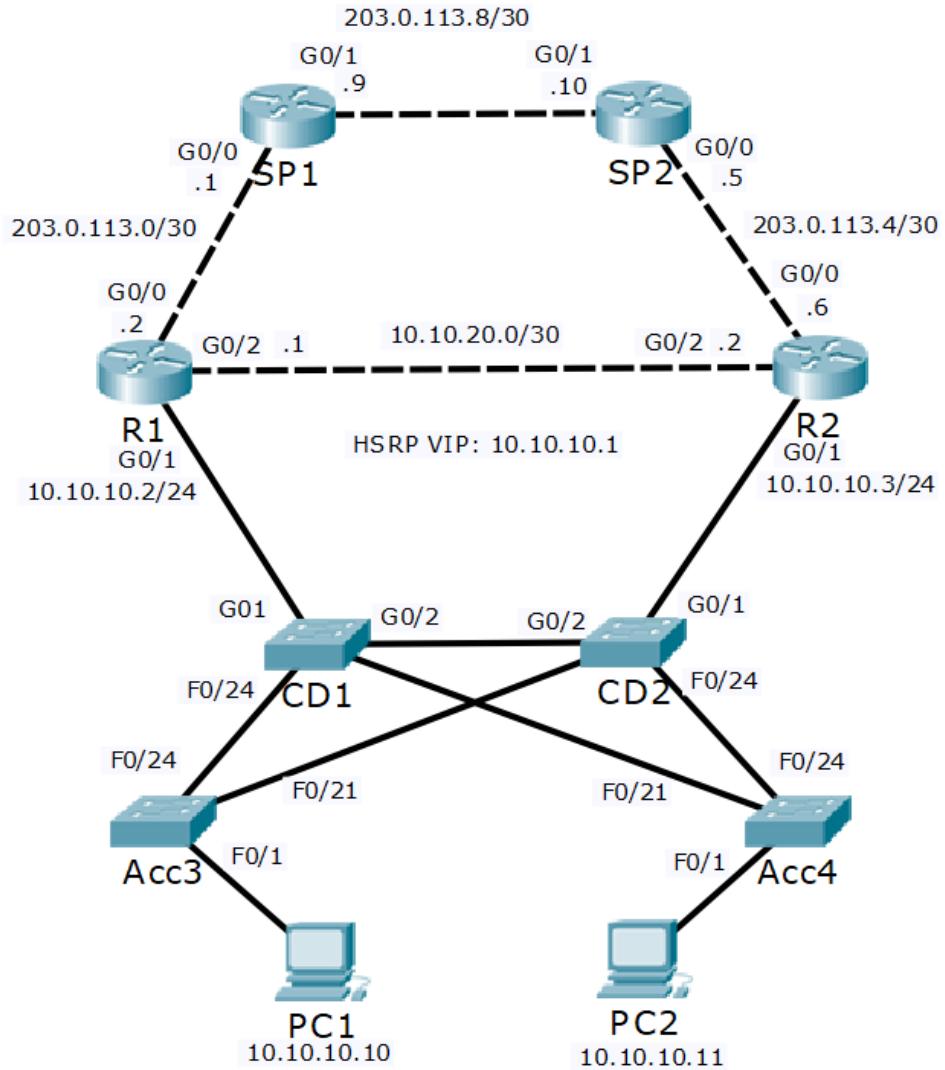
16) Hit 'Ctrl-C' to cancel the ping on PC1. If you scroll back you should see a dropped ping or two as R1 transitioned back to HSRP active.

```
Reply from 10.10.10.1: bytes=32 time=1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=10ms TTL=255
Reply from 10.10.10.1: bytes=32 time=19ms TTL=255
Request timed out.
Reply from 10.10.10.1: bytes=32 time<1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=13ms TTL=255
Reply from 10.10.10.1: bytes=32 time=13ms TTL=255
Reply from 10.10.10.1: bytes=32 time<1ms TTL=255
```

25-1 Spanning Tree Troubleshooting – Lab Exercise

Spanning Tree is enabled by default so the troubleshooting lab comes before the configuration lab for this section. In this lab you will troubleshoot Spanning Tree for a small campus network.

Lab Topology

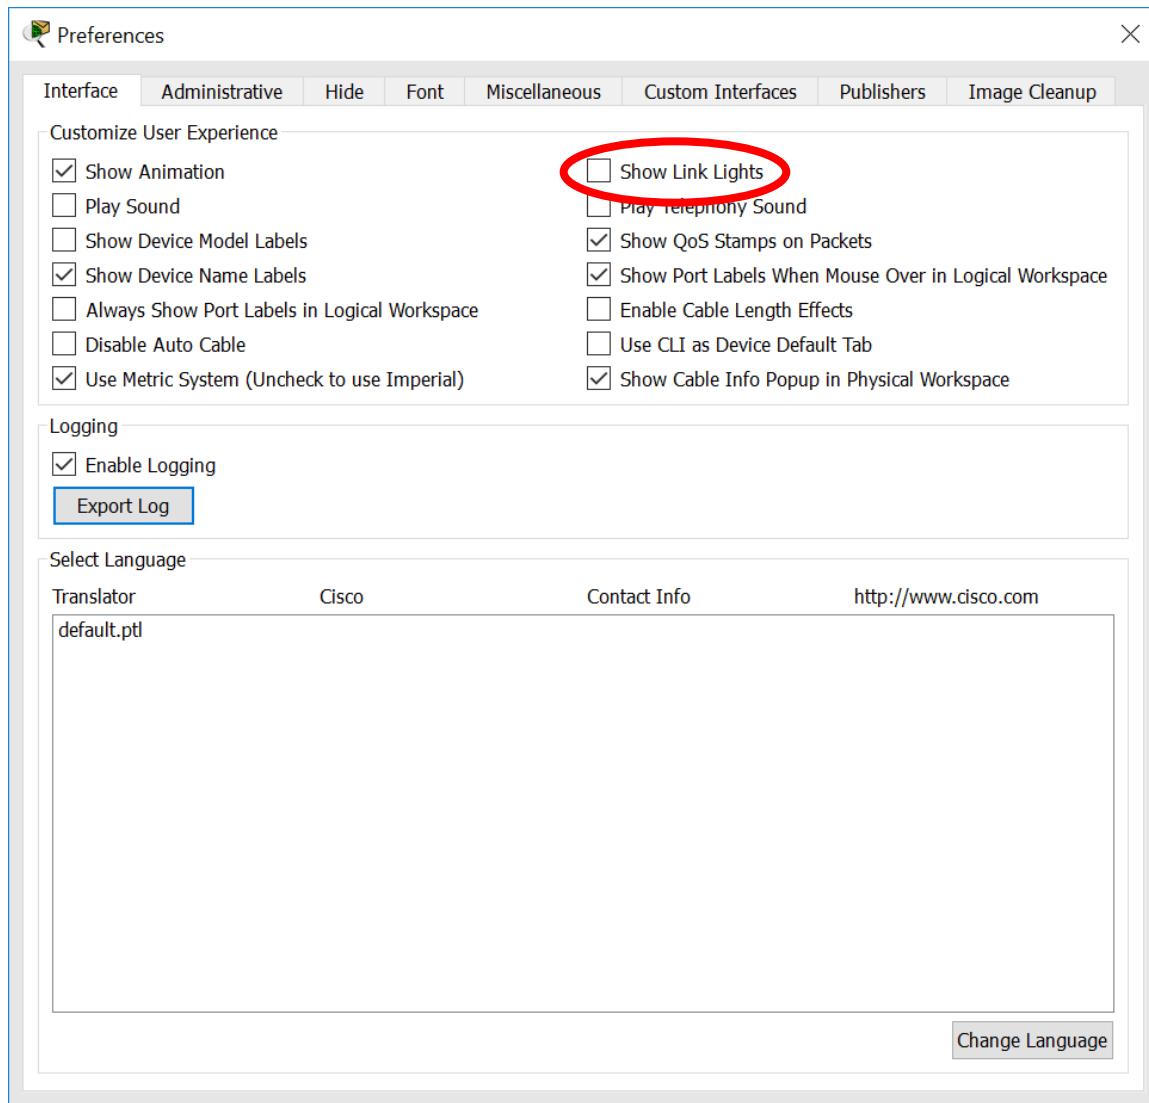


Load the Startup Configurations

IMPORTANT: Before loading the lab exercise, disable 'Show Link Lights' in Packet Tracer. (You can change it back at the end of the lab.)

In Packet Tracer, click 'Options > Preferences'

On the 'Interface' tab, uncheck 'Show Link Lights'



[Download the '25-1 STP Troubleshooting.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Troubleshooting the Spanning Tree Protocol

- 1) The Network Operations Centre has reported that traffic is not following the most direct path from the branch office PCs to the Internet. Your task at this stage is to verify this.

Your task is to report which path traffic is currently taking from both PC1 and PC2 to 203.0.113.9. Do not change any configuration.

25-1 Spanning Tree Troubleshooting - Answer Key

In this lab you will troubleshoot Spanning Tree for a small campus network.

Troubleshooting the Spanning Tree Protocol

- 1) The Network Operations Centre has reported that traffic is not following the most direct path from the branch office PCs to the Internet. Your task at this stage is to verify this.

Your task is to report which path traffic is currently taking from both PC1 and PC2 to 203.0.113.9. Do not change any configuration.

There is not a set order of actions to troubleshoot this scenario. Troubleshooting in a logical fashion will however make it easier and quicker. This is how I would do it.

Check which router is the HSRP active gateway for the 10.10.10.0/24 network.

```
R1#show standby
GigabitEthernet0/1 - Group 1
State is Active
7 state changes, last state change 00:28:52
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
Local virtual MAC address is 0000.0C07.AC01 (v1 default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 2.276 secs
Preemption enabled
Active router is local
Standby router is 10.10.10.3, priority 100 (expires in 7 sec)
Priority 110 (configured 110)
Group name is hsrp-Gig0/1-1 (default)
```

R1 has been preconfigured with a higher HSRP priority and pre-emption enabled.

Check the PCs have connectivity to 203.0.113.9. Ping from both PCs.

```
C:\>ping 203.0.113.9

Pinging 203.0.113.9 with 32 bytes of data:

Request timed out.
Reply from 203.0.113.9: bytes=32 time=15ms TTL=254
Reply from 203.0.113.9: bytes=32 time=10ms TTL=254
Reply from 203.0.113.9: bytes=32 time<1ms TTL=254

Ping statistics for 203.0.113.9:
    Packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 15ms, Average = 8ms
```

Check which Layer 3 path the PCs are using to get to 203.0.113.9. Run traceroute on both PCs.

```
C:\>tracert 203.0.113.9

Tracing route to 203.0.113.9 over a maximum of 30 hops:
  1  0 ms        0 ms        0 ms        10.10.10.2
  2  11 ms       0 ms        0 ms        203.0.113.9

Trace complete.
```

The PCs are taking the most direct path via their HSRP default gateway to get to 203.0.113.9. The Layer 3 configuration and operations all look good.

Check which VLAN the PCs are in on the Acc3 and Acc4 switches.

```
Acc3#show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Fa0/2, Fa0/3, Fa0/4, Fa0/5 Fa0/6, Fa0/7, Fa0/8, Fa0/9 Fa0/10, Fa0/11, Fa0/12, Fa0/13 Fa0/14, Fa0/15, Fa0/16, Fa0/17 Fa0/18, Fa0/19, Fa0/20, Fa0/22 Fa0/23, Gig0/1, Gig0/2
10	Eng	active	Fa0/1
20	Sales	active	
199	Native	active	
1002	fdci-default	active	
1003	token-ring-default	active	
1004	fdinnet-default	active	
1005	trnet-default	active	

The PCs are in the Eng VLAN 10.

On CD1 and CD2, check the switch ports connecting to the routers have also been configured as access ports in VLAN 10.

```
CD1#sh run
! truncated
interface GigabitEthernet0/1
switchport access vlan 10
switchport mode access
```

On all switches, check the interfaces connecting switches have been configured as trunks with matching native VLANs.

```
CD1#sh run
! truncated
!
interface FastEthernet0/21
switchport trunk native vlan 199
switchport mode trunk
!
interface FastEthernet0/24
switchport trunk native vlan 199
switchport mode trunk
!
interface GigabitEthernet0/2
switchport trunk native vlan 199
switchport mode trunk
```

The existing HSRP and VLAN configuration looks good.

The next thing to verify is Spanning Tree. One of the central core/distribution switches should be used as the Root Bridge to ensure Layer 2 traffic uses the most direct path over the network.

Check the Spanning Tree topology for VLAN 10 on switch Acc3.

```
Acc3#show spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    32778
              Address     0001.C962.D43D
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778 (priority 32768 sys-id-ext 10)
              Address     0001.C962.D43D
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time   20

  Interface      Role Sts Cost      Prio.Nbr Type
  -----  -----
  Fa0/1          Desg FWD 19        128.1    P2p
  Fa0/24         Desg FWD 19        128.24   P2p
  Fa0/21         Desg FWD 19        128.21   P2p
```

Acc3 is the Root Bridge. That is not good. Check to see if this is because it has a configured Bridge Priority.

```
Acc3#sh run | include priority
Acc3#
```

Bridge Priority has not been configured on Acc3. Check the other switches.

```
CD1#sh run | include priority
CD1#
```

```
CD2#sh run | include priority
CD2#
```

```
Acc4#sh run | include priority
Acc4#
```

Bridge Priority has not been configured anywhere so the switch with the lowest MAC address should be selected as the Root Bridge.

We learned that Acc3's MAC address is 0001.C962.D43D from the output of the 'show spanning tree vlan 10' command we ran on it.

Verify that all switches have matching Spanning Tree Root Bridge information. They should all have the same default Bridge Priority value (as one was not manually set), and agree that Acc3 has the lowest MAC address and is the Root Bridge.

```
CD1#sh spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    32778
              Address     0001.C962.D43D
              Cost         19
              Port        24 (FastEthernet0/24)
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
              Address     0090.0CA0.3902
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  20

  Interface   Role Sts Cost      Prio.Nbr Type
  -----  -----
  Fa0/21      Desg FWD 19        128.21   P2p
  Fa0/24      Root FWD 19        128.24   P2p
  Gi0/2      Altn BLK 4         128.26   P2p
  Gi0/1      Desg FWD 4         128.25   P2p
```

CD1's MAC address is 0090.0CA0.3902. It agrees that Acc3 is the Root Bridge.

```
CD2#sh spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    32778
              Address     0001.C962.D43D
              Cost         19
              Port        21 (FastEthernet0/21)
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
              Address     0090.0C16.7A9B
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  20

  Interface   Role Sts Cost      Prio.Nbr Type
  -----  -----
  Fa0/21      Root FWD 19        128.21   P2p
  Fa0/24      Desg FWD 19        128.24   P2p
  Gi0/2      Desg FWD 4         128.26   P2p
  Gi0/1      Desg FWD 4         128.25   P2p
```

CD2's MAC address is 0090.0C16.7A9B. It agrees that Acc3 is the Root Bridge.

```

Acc4#show spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    32778
  Address    0001.C962.D43D
  Cost       38
  Port       24 (FastEthernet0/24)
  Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
  Address    0060.708A.D564
  Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
  Aging Time 20

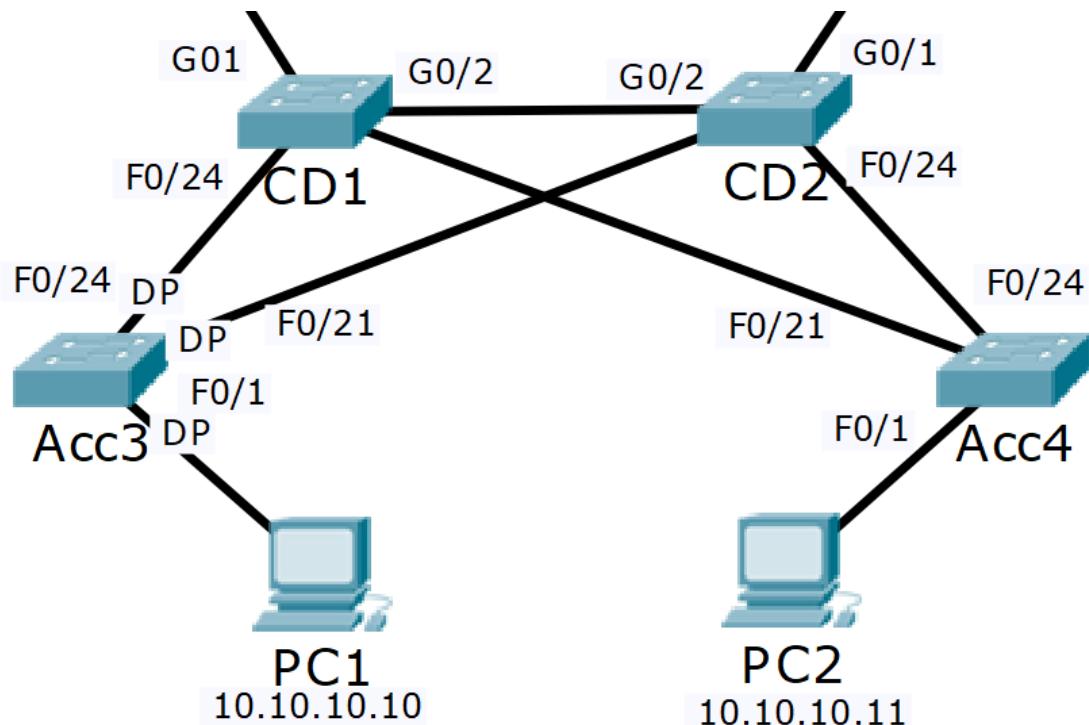
  Interface      Role Sts Cost      Prio.Nbr Type
  -----  -----
  Fa0/1          Desg FWD 19      128.1    P2p
  Fa0/21         Altn BLK 19      128.21   P2p
  Fa0/24         Root FWD 19      128.24   P2p

```

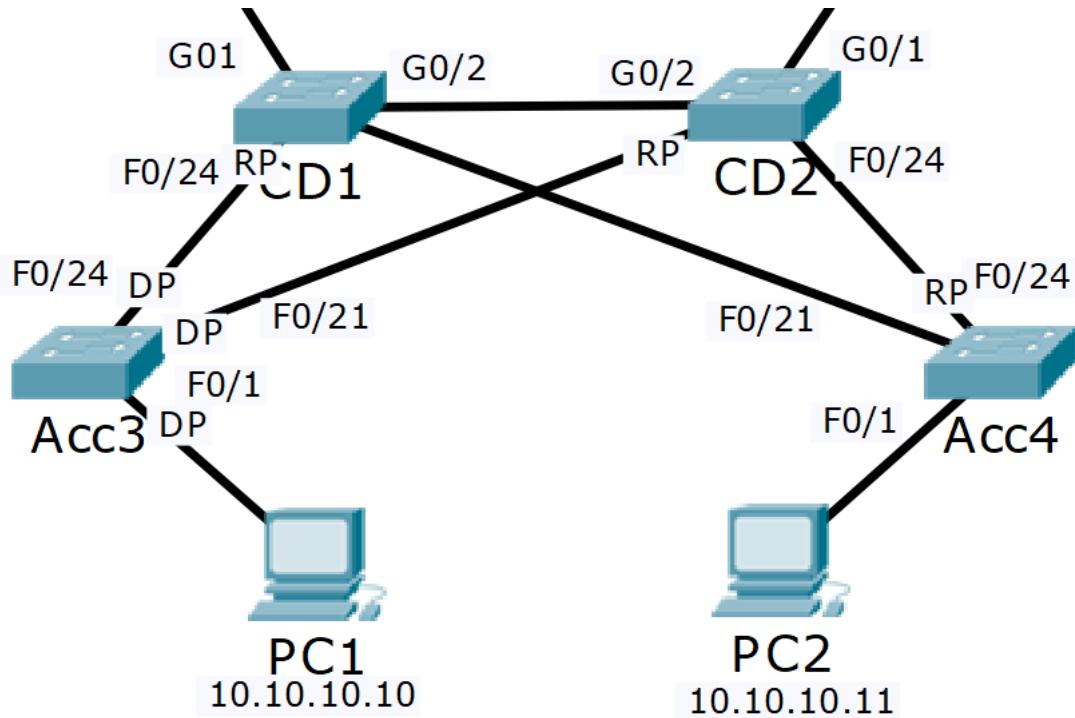
Acc4's MAC address is 0060.708A.D564. It agrees that Acc3 is the Root Bridge.

We were asked to report on the forwarding paths currently being used. Use the output of the 'show spanning-tree vlan 10' commands on each switch to diagram the Spanning Tree.

Acc3 is the Root Bridge so all its ports are Designated Ports.



Add each switch's Root Port to the diagram. The 'Port' value in the 'Root ID' section of the 'show spanning-tree vlan 10' output shows this information.



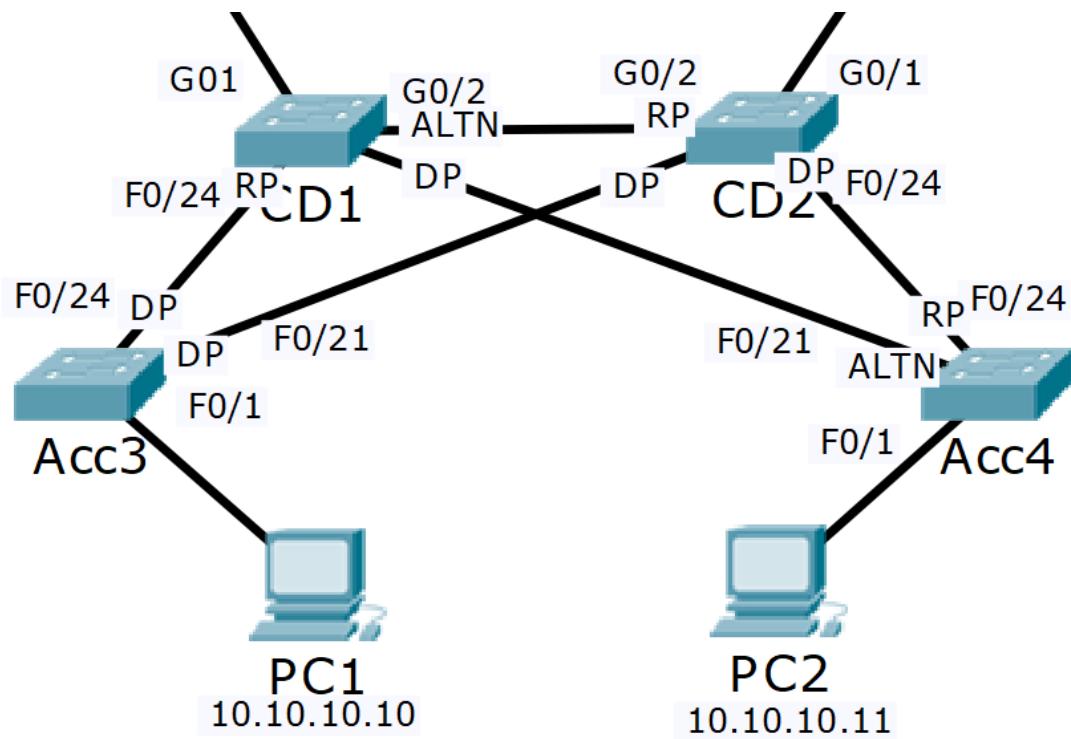
The port on the other side of Root Ports is always a Designated Port. Label F0/24 on CD2 as a Designated Port in the diagram.

The links which are left must be blocking. These are CD1 G0/2 to CD2 G0/2, and CD1 F0/21 to Acc4 F0/21. Check which side is blocking from the output of the 'show spanning-tree vlan 10' commands.

CD1 G0/2 is the blocking Alternate port, CD2 G0/2 is the forwarding Designated Port.

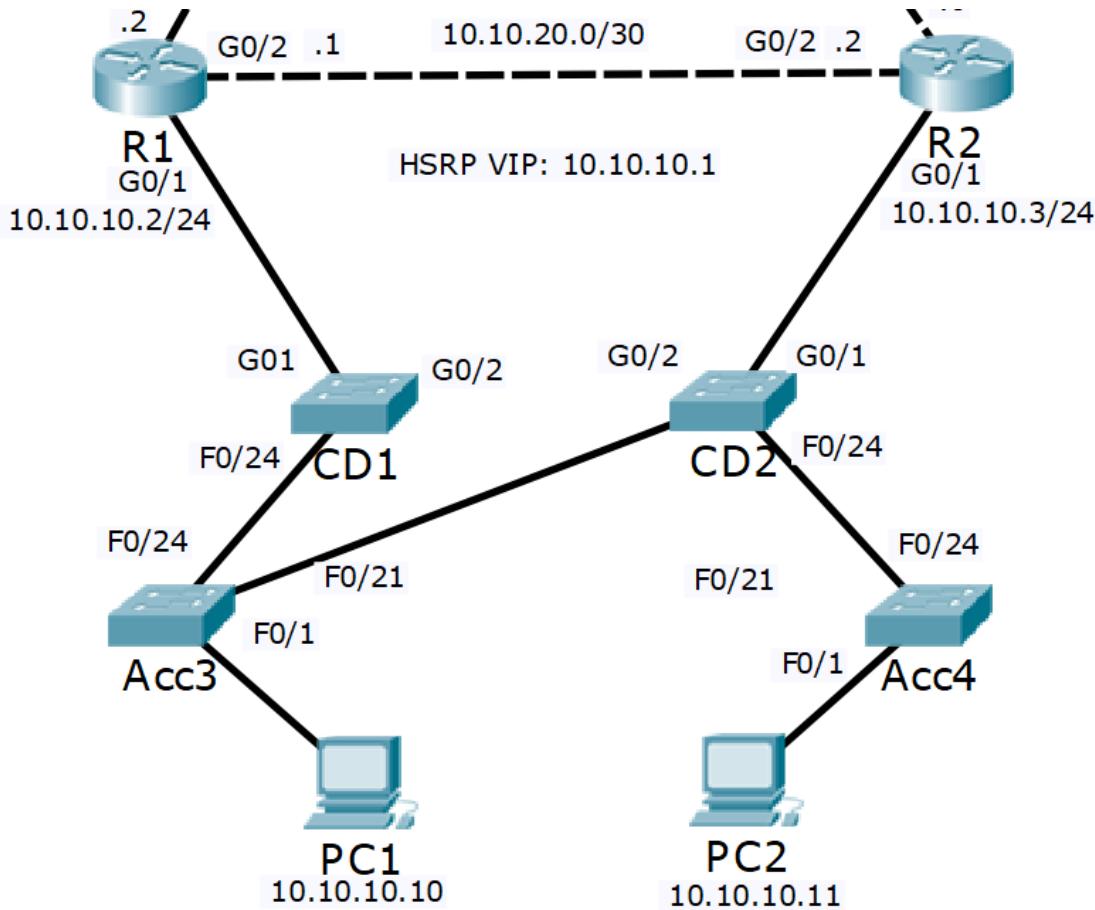
CD1 F0/21 is the forwarding Designated Port, Acc4 F0/21 is the blocking Alternate port.

Add this information to the diagram.



Continued on next page...

By removing the blocking links from the diagram we can see the Spanning Tree.



We can see that PC1 will take the path PC1 > Acc3 > CD1 > R1 to reach its HSRP default gateway. This is the most direct path and is good.

The end to end path from PC1 to 203.0.113.9 is PC1 > Acc3 > CD1 > R1 > SP1

PC2 however will use the path PC2 > Acc4 > CD2 > Acc3 > CD1 > R1. This is not the most direct path as traffic is transiting CD2 and Acc3 rather than going directly over the link from Acc4 to CD1.

The end to end path from PC2 to 203.0.113.9 is PC2 > Acc4 > CD2 > Acc3 > CD1 > R1 > SP1

We can verify the path being used by checking the MAC address tables on the switches. First verify the HSRP virtual MAC address.

```
R1#sh standby
GigabitEthernet0/1 - Group 1
State is Active
5 state changes, last state change 00:00:30
Virtual IP address is 10.10.10.1
Active virtual MAC address is 0000.0C07.AC01
```

Then clear the ARP cache on PC2, and ping the virtual IP address to generate traffic.

```
C:\>arp -d
C:\>ping 10.10.10.1

Pinging 10.10.10.1 with 32 bytes of data:

Reply from 10.10.10.1: bytes=32 time=22ms TTL=255
Reply from 10.10.10.1: bytes=32 time<1ms TTL=255
Reply from 10.10.10.1: bytes=32 time=11ms TTL=255
Reply from 10.10.10.1: bytes=32 time=11ms TTL=255

Ping statistics for 10.10.10.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 22ms, Average = 11ms
```

Then check the MAC address table on Acc4.

```
Acc4#show mac address-table
      Mac Address Table
-----
Vlan      Mac Address          Type      Ports
---      -----
 1        0002.1690.e018    DYNAMIC   Fa0/24
10        0000.0c07.ac01    DYNAMIC   Fa0/24
10        0001.6310.5d05    DYNAMIC   Fa0/1
10        0001.6470.2502    DYNAMIC   Fa0/24
10        0002.1690.e018    DYNAMIC   Fa0/24
10        00d0.ffe2.2d02    DYNAMIC   Fa0/24
20        0002.1690.e018    DYNAMIC   Fa0/24
199       0002.1690.e018    DYNAMIC   Fa0/24
```

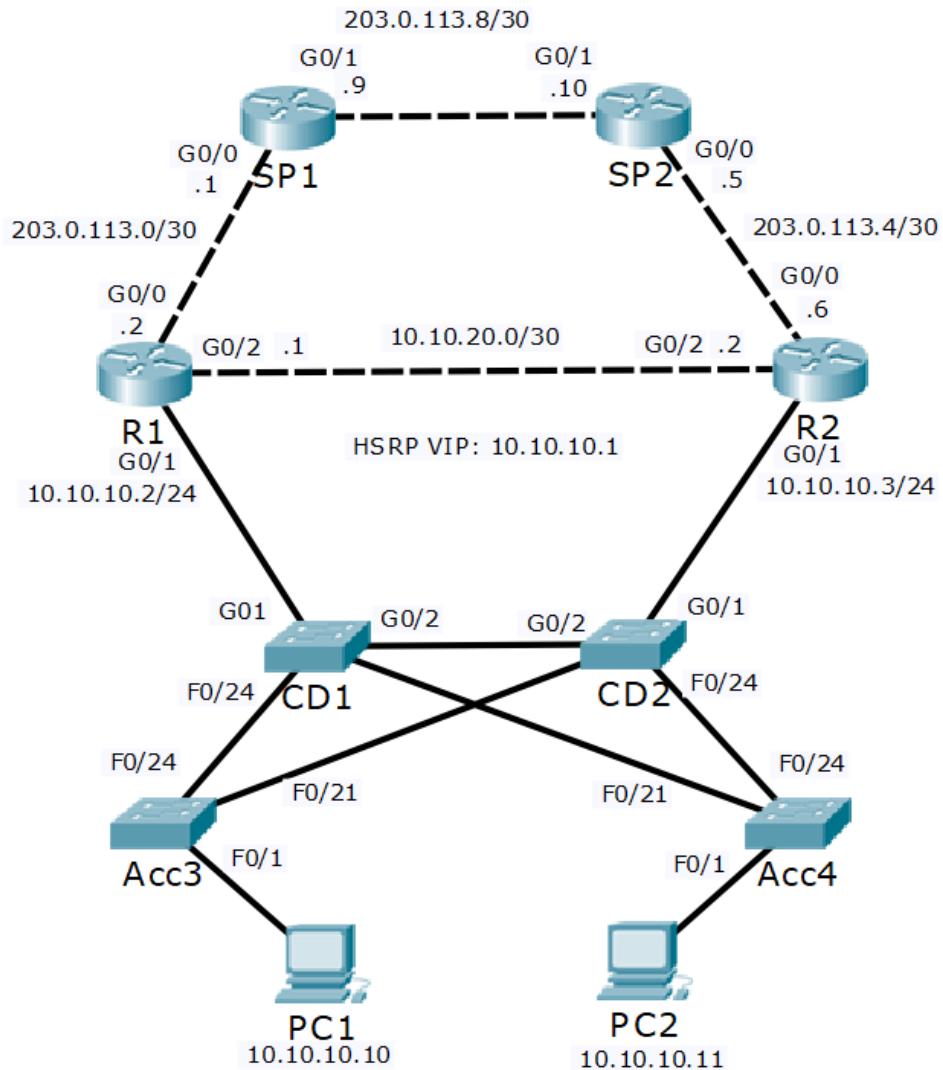
We can see that the HSRP virtual MAC address 0000.0c07.ac01 is reached through interface F0/24 to CD2, rather than on the direct link to CD1 over interface F0/21.

We can go hop by hop using the 'show mac address-table' command to verify the traffic path end to end across the switched network.

25-2 Spanning Tree Configuration – Lab Exercise

You discovered that switch Acc3 is the Root Bridge in the previous Spanning Tree Troubleshooting lab and traffic is not taking the most direct path across the network. You will correct the configuration in this lab.

Lab Topology



Load the Startup Configurations

[Download the '25-2 EtherChannel Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Spanning Tree Root Bridge Configuration

- 1) Configure the network so that traffic between the PCs and the Internet travels along the shortest available path. If a core/distribution switch fails traffic should failover to the next shortest available path. Do not change any Layer 3 configuration such as HSRP settings.

Portfast and BPDU Guard Configuration

- 2) A Layer 2 loop cannot be formed on a port where a single end host is connected. Disable spanning tree on these ports.

You are concerned that a user may introduce a loop into the network by adding additional switches or changing the cabling. Also ensure that the ports where Spanning Tree is disabled will be automatically shut down if a switch is detected on the other side of the link.

25-2 Spanning Tree Configuration - Answer Key

You discovered that switch Acc3 is the Root Bridge in the previous Spanning Tree Troubleshooting lab and traffic is not taking the most direct path across the network. You will correct the configuration in this lab.

Spanning Tree Root Bridge Configuration

- 1) Configure the network so that traffic between the PCs and the Internet travels along the shortest available path. If a core/distribution switch fails traffic should failover to the next shortest available path. Do not change any Layer 3 configuration such as HSRP settings.

We need to configure the Spanning Tree so it aligns with the HSRP configuration. R1 is the HSRP active gateway. R1 is directly connected to the core/distribution switch CD1 (but not CD2) so we should make this the Spanning Tree Root Bridge.

```
CD1(config)#spanning-tree vlan 10 root primary
```

If CD1 fails we need to ensure that the Spanning Tree Root Bridge will failover to CD2 rather than an access layer switch.

```
CD2(config)#spanning-tree vlan 10 root secondary
```

Verify CD1 has the best Bridge Priority and becomes the Root Bridge.

```
CD1#show spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
              Address     0090.0CA0.3902
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24586 (priority 24576 sys-id-ext 10)
              Address     0090.0CA0.3902
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time   20

  Interface   Role Sts Cost      Prio.Nbr Type
  -----  -----
  Fa0/21      Desg FWD 19        128.21   P2p
  Fa0/24      Desg FWD 19        128.24   P2p
  Gi0/1       Desg FWD 4         128.25   P2p
  Gi0/2       Desg FWD 4         128.26   P2p
```

Check the other switches to verify CD2 has the next best Bridge Priority.

```
CD2#show spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
              Address     0090.0CA0.3902
              Cost        4
              Port        26(GigabitEthernet0/2)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    28682 (priority 28672 sys-id-ext 10)
              Address     0090.0C16.7A9B
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time   20

  Interface   Role Sts Cost      Prio.Nbr Type
  -----  -----
  Gi0/2       Root FWD 4         128.26   P2p
  Fa0/21      Desg FWD 19        128.21   P2p
  Gi0/1       Desg FWD 4         128.25   P2p
  Fa0/24      Desg FWD 19        128.24   P2p
```

```

Acc3#show spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
              Address     0090.0CA0.3902
              Cost         19
              Port        24 (FastEthernet0/24)
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
              Address     0001.C962.D43D
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  20

  Interface   Role Sts Cost      Prio.Nbr Type
  -----  -----
  Fa0/1        Desg FWD 19      128.1    P2p
  Fa0/21       Altn BLK 19      128.21   P2p
  Fa0/24       Root FWD 19      128.24   P2p

```

```

Acc4#show spanning-tree vlan 10
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
              Address     0090.0CA0.3902
              Cost         19
              Port        21 (FastEthernet0/21)
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
              Address     0060.708A.D564
              Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  20

  Interface   Role Sts Cost      Prio.Nbr Type
  -----  -----
  Fa0/1        Desg FWD 19      128.1    P2p
  Fa0/21       Root FWD 19      128.21   P2p
  Fa0/24       Altn BLK 19      128.24   P2p

```

Verify the end to end traffic path between the PCs and the Internet by using the 'show spanning-tree vlan 10' and 'show mac address-table' commands as shown in the last lab exercise.

Portfast and BPDU Guard Configuration

- 2) A Layer 2 loop cannot be formed on a port where a single end host is connected. Disable spanning tree on these ports.

You are concerned that a user may introduce a loop into the network by adding additional switches or changing the cabling. Also ensure that the ports where Spanning Tree is disabled will be automatically shut down if a switch is detected on the other side of the link.

```
Acc3(config)#int f0/1
Acc3(config-if)#spanning-tree portfast
Acc3(config-if)#spanning-tree bpduguard enable
```

```
Acc4(config)#int f0/1
Acc4(config-if)#spanning-tree portfast
Acc4(config-if)#spanning-tree bpduguard enable
```

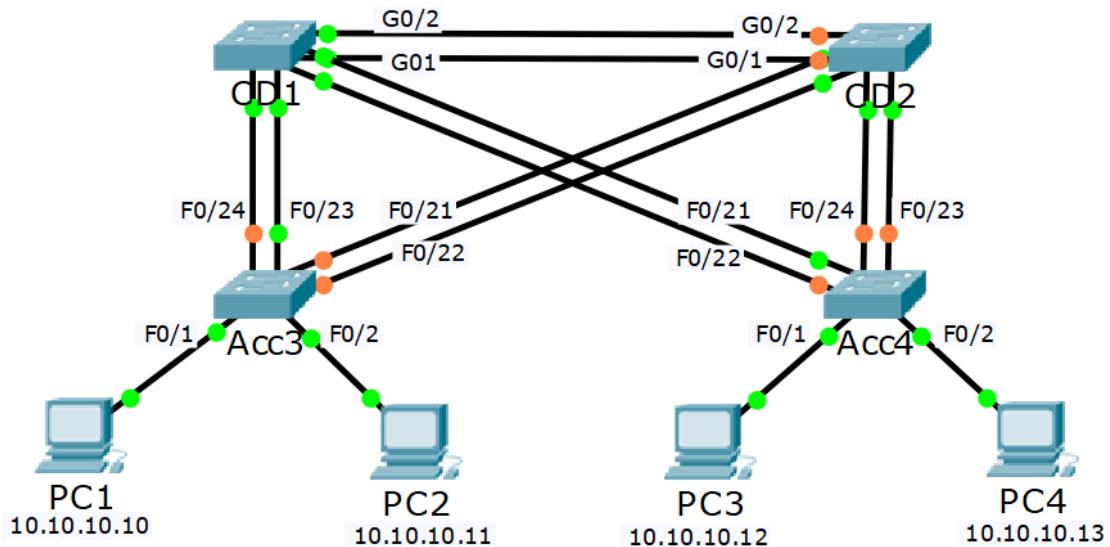
```
CD1(config)#int g0/1
CD1(config-if)#spanning-tree portfast
CD1(config-if)#spanning-tree bpduguard enable
```

```
CD2(config)#int g0/1
CD2(config-if)#spanning-tree portfast
CD2(config-if)#spanning-tree bpduguard enable
```

26-1 EtherChannel Configuration – Lab Exercise

In this lab you will configure Etherchannel links in a campus LAN.

Lab Topology



Load the Startup Configurations

[Download the '26-1 EtherChannel Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

LACP EtherChannel Configuration

- 1) The access layer switches Acc3 and Acc4 both have two FastEthernet uplinks. How much total bandwidth is available between the PCs attached to Acc3 and the PCs attached to Acc4?
- 2) Convert the existing uplinks from Acc3 to CD1 and CD2 to LACP EtherChannel. Configure descriptions on the port channel interfaces to help avoid confusion later.
- 3) Verify the EtherChannels come up.

PAgP EtherChannel Configuration

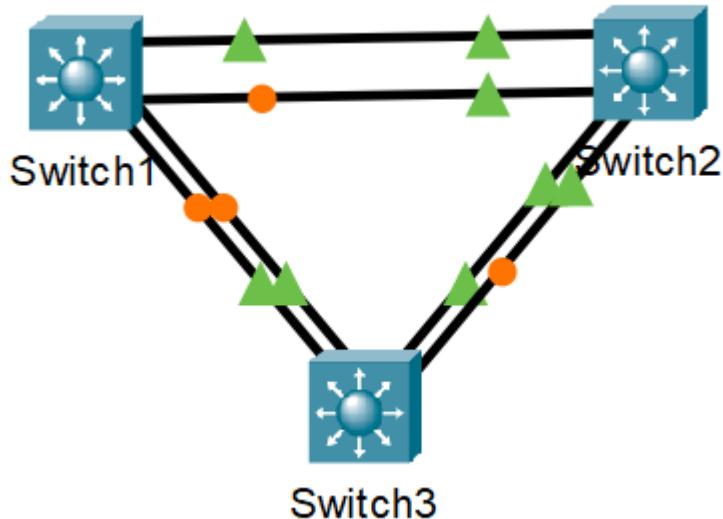
- 4) Convert the existing uplinks from Acc4 to CD1 and CD2 to PAgP EtherChannel. (Note that in a real world environment you should always use LACP if possible.)
- 5) Verify the EtherChannels come up.

Static EtherChannel Configuration

- 6) Convert the existing uplinks between CD1 and CD2 to static EtherChannel.
- 7) Verify the EtherChannel comes up.
- 8) How much total bandwidth is available between the PCs attached to Acc3 and the PCs attached to Acc4 now?

Layer 3 EtherChannel Configuration

The Layer 3 switches Switch1, Switch2 and Switch3 are physically separate from the switches you configured earlier in this lab exercise.



- 9) Switch1 and Switch2 are connected together with their GigabitEthernet1/0/1 and 1/0/2 interfaces.
Configure these interfaces as a Layer 3 Etherchannel with LACP.
Configure IP address 192.168.0.1/30 on Switch1 and 192.168.0.2/30 on Switch2.
- 10) Switch1 and Switch3 are connected together with their GigabitEthernet1/0/3 and 1/0/4 interfaces.
Configure these interfaces as a Layer 3 Etherchannel with LACP.
Configure IP address 192.168.0.5/30 on Switch1 and 192.168.0.6/30 on Switch3.
- 11) Switch2 and Switch3 are connected together with their GigabitEthernet1/0/5 and 1/0/6 interfaces.
Configure these interfaces as a Layer 3 Etherchannel with LACP.
Configure IP address 192.168.0.9/30 on Switch2 and 192.168.0.10/30 on Switch3.
- 12) Verify the EtherChannels come up.
- 13) Configure Switch1, Switch2 and Switch3 to advertise the IP subnets configured on their Etherchannel interfaces in OSPF Area 0.
- 14) Verify the OSPF adjacencies are formed successfully.

- 15) Verify Switch1, Switch2 and Switch3 have routes to all configured networks in their routing tables.
- 16) Which physical ports on which switches do you expect the Spanning Tree protocol to disable? Verify this.

26-1 EtherChannel Configuration - Answer Key

In this lab you will configure EtherChannel links in a campus LAN.

LACP EtherChannel Configuration

- 1) The access layer switches Acc3 and Acc4 both have two FastEthernet uplinks. How much total bandwidth is available between the PCs attached to Acc3 and the PCs attached to Acc4?

Spanning tree shuts down all but one uplink on both switches so the total bandwidth available between them is a single FastEthernet link – 100 Mbps.

- 2) Convert the existing uplinks from Acc3 to CD1 and CD2 to LACP EtherChannel. Configure descriptions on the port channel interfaces to help avoid confusion later.

The uplinks go to two separate redundant switches at the core/distribution layer so we need to configure two EtherChannels, one to CD1 and one to CD2.

We'll configure the Acc3 side of the EtherChannel to CD1 first. Don't forget to set the native VLAN on the new port channel interface.

```
Acc3(config)#interface range f0/23 - 24
Acc3(config-if-range)#channel-group 1 mode active
Acc3(config-if-range)#exit
Acc3(config)#interface port-channel 1
Acc3(config-if)#description Link to CD1
Acc3(config-if)#switchport mode trunk
Acc3(config-if)#switchport trunk native vlan 199
```

Then configure switch CD1 with matching settings.

```
CD1(config)#interface range f0/23 - 24
CD1(config-if-range)#channel-group 1 mode active
CD1(config-if-range)#exit
CD1(config)#interface port-channel 1
CD1(config-if)#description Link to Acc3
CD1(config-if)#switchport mode trunk
CD1(config-if)#switchport trunk native vlan 199
```

Next configure the Acc3 side of the EtherChannel to CD2. Remember to use a different port channel number.

```
Acc3(config)#interface range f0/21 - 22
Acc3(config-if-range)#channel-group 2 mode active
Acc3(config-if-range)#exit
Acc3(config)#interface port-channel 2
Acc3(config-if)#description Link to CD2
Acc3(config-if)#switchport mode trunk
Acc3(config-if)#switchport trunk native vlan 199
```

Then configure switch CD2 with matching settings.

```
CD2(config)#interface range f0/21 - 22
CD2(config-if-range)#channel-group 2 mode active
CD2(config-if-range)#exit
CD2(config)#interface port-channel 2
CD2(config-if)#description Link to Acc3
CD2(config-if)#switchport mode trunk
CD2(config-if)#switchport trunk native vlan 199
```

3) Verify the EtherChannels come up.

The port channels should show flags (SU) (Layer 2, in use) with member ports (P) (in port-channel). Verify on both sides of the port channel.

Acc3:

```
Acc3#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
|
Number of channel-groups in use: 2
Number of aggregators:          2
Group  Port-channel  Protocol      Ports
-----+-----+-----+
1      Po1 (SU)       LACP      Fa0/23 (P)  Fa0/24 (P)
2      Po2 (SU)       LACP      Fa0/21 (P)  Fa0/22 (P)
```

CD1:

```
CD1#show etherchannel summary
Flags:  D - down          P - in port-channel
       I - stand-alone  s - suspended
       H - Hot-standby (LACP only)
       R - Layer3         S - Layer2
       U - in use          f - failed to allocate aggregator
       u - unsuitable for bundling
       w - waiting to be aggregated
       d - default port
```

Number of channel-groups in use: 1
Number of aggregators: 1

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	LACP	Fa0/23 (P) Fa0/24 (P)

CD2:

```
CD2#show etherchannel summary
Flags:  D - down          P - in port-channel
       I - stand-alone  s - suspended
       H - Hot-standby (LACP only)
       R - Layer3         S - Layer2
       U - in use          f - failed to allocate aggregator
       u - unsuitable for bundling
       w - waiting to be aggregated
       d - default port
```

Number of channel-groups in use: 1
Number of aggregators: 1

Group	Port-channel	Protocol	Ports
2	Po2 (SU)	LACP	Fa0/21 (P) Fa0/22 (P)

PAgP EtherChannel Configuration

- 4) Convert the existing uplinks from Acc4 to CD1 and CD2 to PAgP EtherChannel. (Note that in a real world environment you should always use LACP if possible.)

It's good practice to use the same port channel number on both sides of the link. CD1 is already using port channel 1 to Acc3, and CD2 is using port channel 2 to Acc3. From Acc4 to CD1 we'll use port channel 2, and from Acc4 to CD2 we'll use port channel 1.

We'll configure the Acc4 side of the EtherChannel to CD2 first.

```
Acc4(config)#interface range f0/23 - 24
Acc4(config-if-range)#channel-group 1 mode desirable
Acc4(config-if-range)#exit
Acc4(config)#interface port-channel 1
Acc4(config-if)#description Link to CD2
Acc4(config-if)#switchport mode trunk
Acc4(config-if)#switchport trunk native vlan 199
```

Then configure switch CD2 with matching settings.

```
CD2(config)#interface range f0/23 - 24
CD2(config-if-range)#channel-group 1 mode desirable
CD2(config-if-range)#exit
CD2(config)#interface port-channel 1
CD2(config-if)#description Link to Acc4
CD2(config-if)#switchport mode trunk
CD2(config-if)#switchport trunk native vlan 199
```

Next configure the Acc4 side of the EtherChannel to CD1. Remember to use a different port channel number.

```
Acc4(config)#interface range f0/21 - 22
Acc4(config-if-range)#channel-group 2 mode desirable
Acc4(config-if-range)#exit
Acc4(config)#interface port-channel 2
Acc4(config-if)#description Link to CD1
Acc4(config-if)#switchport mode trunk
Acc4(config-if)#switchport trunk native vlan 199
```

Then configure switch CD1 with matching settings.

```
CD1(config)#interface range f0/21 - 22
CD1(config-if-range)#channel-group 2 mode desirable
CD1(config-if-range)#exit
CD1(config)#interface port-channel 2
CD1(config-if)#description Link to Acc4
CD1(config-if)#switchport mode trunk
CD1(config-if)#switchport trunk native vlan 199
```

5) Verify the EtherChannels come up.

The port channels should show flags (SU) (Layer 2, in use) with member ports (P) (in port-channel). Verify on both sides of the port channel.

On the core/distribution layer switches you should see both the LACP and PAgP port channels up.

Acc4:

```
Acc4#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

Number of channel-groups in use: 2

Number of aggregators: 2

Group	Port-channel	Protocol	Ports	
-----	-----	-----	-----	-----
1	Po1 (SU)	PAgP	Fa0/23 (P)	Fa0/24 (P)
2	Po2 (SU)	PAgP	Fa0/21 (P)	Fa0/22 (P)

CD1:

```
CD1#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  S - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

Number of channel-groups in use: 2

Number of aggregators: 2

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	LACP	Fa0/23 (P) Fa0/24 (P)
2	Po2 (SU)	PAgP	Fa0/21 (P) Fa0/22 (P)

CD2:

```
CD2#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  S - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

Number of channel-groups in use: 2

Number of aggregators: 2

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	PAgP	Fa0/23 (P) Fa0/24 (P)
2	Po2 (SU)	LACP	Fa0/21 (P) Fa0/22 (P)

Static EtherChannel Configuration

- 6) Convert the existing uplinks between CD1 and CD2 to static EtherChannel.

Port channels 1 and 2 are already in use so we'll use port channel 3.

Configure the CD1 side first.

```
CD1(config)#interface range g0/1 - 2
CD1(config-if-range)#channel-group 3 mode on
CD1(config-if-range)#exit
CD1(config)#interface port-channel 3
CD1(config-if)#description Link to CD2
CD1(config-if)#switchport mode trunk
CD1(config-if)#switchport trunk native vlan 199
```

Then configure switch CD2 with matching settings.

```
CD2(config)#interface range g0/1 - 2
CD2(config-if-range)#channel-group 3 mode on
CD2(config-if-range)#exit
CD2(config)#interface port-channel 3
CD2(config-if)#description Link to CD1
CD2(config-if)#switchport mode trunk
CD2(config-if)#switchport trunk native vlan 199
```

7) Verify the EtherChannel comes up.

CD1:

```
CD1#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

Number of channel-groups in use: 3

Number of aggregators: 3

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	LACP	Fa0/23 (P) Fa0/24 (P)
2	Po2 (SU)	PAgP	Fa0/21 (P) Fa0/22 (P)
3	Po3 (SU)	-	Gig0/1 (P) Gig0/2 (P)

CD2:

```
CD2#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

Number of channel-groups in use: 3

Number of aggregators: 3

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	PAgP	Fa0/23 (P) Fa0/24 (P)
2	Po2 (SU)	LACP	Fa0/21 (P) Fa0/22 (P)
3	Po3 (SU)	-	Gig0/1 (P) Gig0/2 (P)

8) How much total bandwidth is available between the PCs attached to Acc3 and the PCs attached to Acc4 now?

The port channels from the Acc3 and Acc4 switches towards the root bridge CD1 are up and forwarding. Spanning tree shuts down the port channels toward CD2 to prevent a loop.

The port channels from Acc3 and Acc4 facing the root bridge comprise two FastEthernet interfaces, so the total bandwidth available between the PCs attached to the different access layer switches is 200 Mbps.

Layer 3 EtherChannel Configuration

The switches Switch1, Switch2 and Switch3 are physically separate from the switches you configured earlier in this lab exercise.

9) Switch1 and Switch2 are connected together with their GigabitEthernet1/0/1 and 1/0/2 interfaces.
Configure these interfaces as a Layer 3 Etherchannel with LACP.
Configure IP address 192.168.0.1/30 on Switch1 and 192.168.0.2/30 on Switch2.

You can use your choice of number for the channel-group. Starting at '1' and using the same number on both sides keeps the configuration logical and easier to troubleshoot.

```
Switch1(config)#interface range GigabitEthernet 1/0/1 - 2
Switch1(config-if-range)#no switchport
Switch1(config-if-range)#channel-group 1 mode active
Switch1(config-if-range)#exit
```

```
Switch1(config)#interface port-channel 1
Switch1(config-if)#ip address 192.168.0.1 255.255.255.252
Switch1(config-if)#no shutdown
```

```
Switch2(config)#interface range GigabitEthernet 1/0/1 - 2
Switch2(config-if-range)#no switchport
Switch2(config-if-range)#channel-group 1 mode active
Switch2(config-if-range)#exit
```

```
Switch2(config)#interface port-channel 1
Switch2(config-if)#ip address 192.168.0.2 255.255.255.252
Switch2(config-if)#no shutdown
```

10) Switch1 and Switch3 are connected together with their GigabitEthernet1/0/3 and 1/0/4 interfaces.
Configure these interfaces as a Layer 3 Etherchannel with LACP.
Configure IP address 192.168.0.5/30 on Switch1 and 192.168.0.6/30 on Switch3.

```
Switch1(config)#interface range GigabitEthernet 1/0/3 - 4
Switch1(config-if-range)#no switchport
Switch1(config-if-range)#channel-group 2 mode active
Switch1(config-if-range)#exit

Switch1(config)#interface port-channel 2
Switch1(config-if)#ip address 192.168.0.5 255.255.255.252
Switch1(config-if)#no shutdown

Switch3(config)#interface range GigabitEthernet 1/0/3 - 4
Switch3(config-if-range)#no switchport
Switch3(config-if-range)#channel-group 2 mode active
Switch3(config-if-range)#exit

Switch3(config)#interface port-channel 2
Switch3(config-if)#ip address 192.168.0.6 255.255.255.252
Switch3(config-if)#no shutdown
```

11) Switch2 and Switch3 are connected together with their GigabitEthernet1/0/5 and 1/0/6 interfaces.
Configure these interfaces as a Layer 3 Etherchannel with LACP.
Configure IP address 192.168.0.9/30 on Switch2 and 192.168.0.10/30 on Switch3.

```
Switch2(config)#interface range GigabitEthernet 1/0/5 - 6
Switch2(config-if-range)#no switchport
Switch2(config-if-range)#channel-group 3 mode active
Switch2(config-if-range)#exit

Switch2(config)#interface port-channel 3
Switch2(config-if)#ip address 192.168.0.9 255.255.255.252
Switch2(config-if)#no shutdown

Switch3(config)#interface range GigabitEthernet 1/0/5 - 6
Switch3(config-if-range)#no switchport
Switch3(config-if-range)#channel-group 3 mode active
Switch3(config-if-range)#exit

Switch3(config)#interface port-channel 3
Switch3(config-if)#ip address 192.168.0.10 255.255.255.252
Switch3(config-if)#no shutdown
```

12) Verify the EtherChannels come up.

```
Switch1#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  S - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

Number of channel-groups in use: 2

Number of aggregators: 2

Group	Port-channel	Protocol	Ports
1	Po1 (RU)	LACP	Gig1/0/1(P) Gig1/0/2(P)
2	Po2 (RU)	LACP	Gig1/0/3(P) Gig1/0/4(P)

```
Switch2#show etherchannel summary
```

```
Flags:  D - down          P - in port-channel
        I - stand-alone  S - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

Number of channel-groups in use: 2

Number of aggregators: 2

Group	Port-channel	Protocol	Ports
1	Po1 (RU)	LACP	Gig1/0/1(P) Gig1/0/2(P)
3	Po3 (RU)	LACP	Gig1/0/5(P) Gig1/0/6(P)

```

Switch3#sh etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  S - suspended
        H - Hot-standby (LACP only)
        R - Layer3         S - Layer2
        U - in use          f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 2
Number of aggregators:          2

Group  Port-channel  Protocol      Ports
-----+-----+-----+
2      Po2 (RU)      LACP         Gig1/0/3 (P)  Gig1/0/4 (P)
3      Po3 (RU)      LACP         Gig1/0/5 (P)  Gig1/0/6 (P)

```

13) Configure Switch1, Switch2 and Switch3 to advertise the IP subnets configured on their Etherchannel interfaces in OSPF Area 0.

On Switch1, Switch2 and Switch3:

```

Switch1(config)#ip routing
Switch1(config)#router ospf 1
Switch1(config-router)#network 192.168.0.0 0.0.0.255 area 0

```

14) Verify the OSPF adjacencies are formed successfully.

On Switch1, Switch2 and Switch3:

```

Switch1#show ip ospf neighbor

Neighbor ID      Pri   State          Dead Time      Address          Interface
192.168.0.10     1     FULL/BDR      00:00:38      192.168.0.6    Port-channel12
192.168.0.9      1     FULL/BDR      00:00:30      192.168.0.2    Port-channell1

```

15) Verify Switch1, Switch2 and Switch3 have routes to all configured networks in their routing tables.

```
Switch1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
      i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
      * - candidate default, U - per-user static route, o - ODR
      P - periodic downloaded static route
```

Gateway of last resort is not set

```
      192.168.0.0/30 is subnetted, 3 subnets
C        192.168.0.0 is directly connected, Port-channel1
C        192.168.0.4 is directly connected, Port-channel2
O        192.168.0.8 [110/2] via 192.168.0.2, 00:01:57, Port-channel1
                  [110/2] via 192.168.0.6, 00:01:57, Port-channel2
```

```
Switch2#show ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
      i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
      * - candidate default, U - per-user static route, o - ODR
      P - periodic downloaded static route
```

Gateway of last resort is not set

```
      192.168.0.0/30 is subnetted, 3 subnets
C        192.168.0.0 is directly connected, Port-channel1
O        192.168.0.4 [110/2] via 192.168.0.1, 00:02:45, Port-channel1
                  [110/2] via 192.168.0.10, 00:02:45, Port-channel3
C        192.168.0.8 is directly connected, Port-channel3
```

```
Switch3#show ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
      i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
      * - candidate default, U - per-user static route, o - ODR
      P - periodic downloaded static route
```

Gateway of last resort is not set

```
      192.168.0.0/30 is subnetted, 3 subnets
O        192.168.0.0 [110/2] via 192.168.0.5, 00:03:13, Port-channel2
                  [110/2] via 192.168.0.9, 00:03:13, Port-channel3
C        192.168.0.4 is directly connected, Port-channel2
C        192.168.0.8 is directly connected, Port-channel3
```

16) Which physical ports on which switches do you expect the Spanning Tree protocol to disable? Verify this.

Spanning Tree only runs on Layer 2 interfaces. It will not run on or shut any of the ports down as they are all Layer 3 ports. The Layer 3 switches' routing tables will handle path selection, redundancy and load balancing.

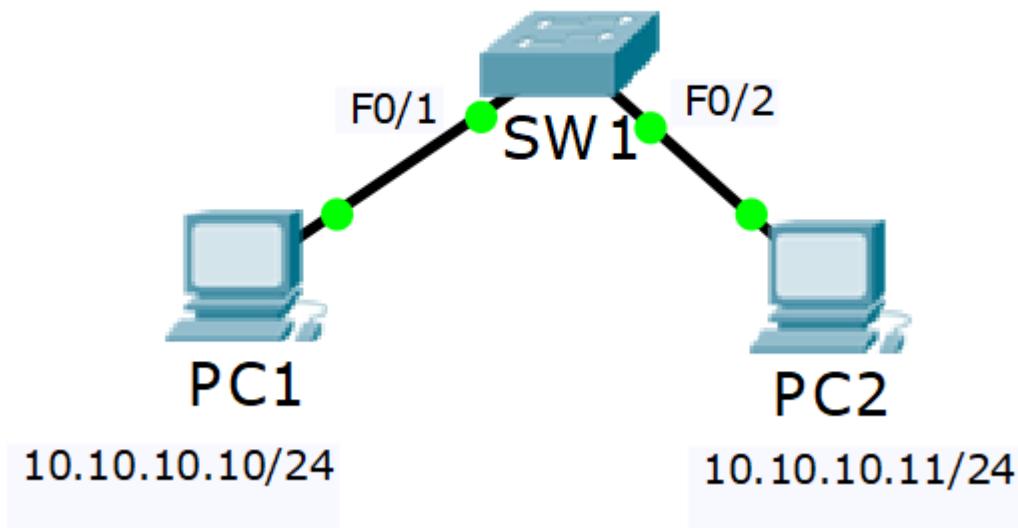
```
Switch1#show spanning-tree
```

```
No spanning tree instance exists.
```

27-1 Port Security Configuration Lab Exercise

In this lab you will configure Port Security on a small campus network.

Lab Topology



Load the Startup Configurations

[Download the '27-1 Port Security Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Disable Unused Ports

- 1) Disable all unused ports on SW1. This prevents unauthorised hosts plugging in to them to gain access to the network.

Port Security Configuration

- 2) Configure port security on interface FastEthernet 0/1. Allow a maximum of two MAC addresses and manually add PC1's MAC address to the configuration.
- 3) Enable Port Security on interface FastEthernet 0/2 with the default settings.
- 4) Use a 'show port-security' command to verify the MAC address on PC2.
- 5) Verify the full Port Security configuration on both interfaces.

27-1 Port Security Configuration Answer Key

In this lab you will configure Port Security on a small campus network.

Disable Unused Ports

- 1) Disable all unused ports on SW1. This prevents unauthorised hosts plugging in to them to gain access to the network.

'show ip interface brief' shows ports FastEthernet 0/1 – 24 and GigabitEthernet0/1 – 2. Interfaces FastEthernet 0/1 and 0/2 are in use.

```
SW1#sh ip int brief
Interface IP-Address OK? Method Status Protocol
FastEthernet0/1 unassigned YES manual up up
! truncated
FastEthernet0/24 unassigned YES manual down down
GigabitEthernet0/1 unassigned YES manual down down
GigabitEthernet0/2 unassigned YES manual down down
Vlan1 unassigned YES manual administratively down down
```

Interfaces FastEthernet 0/1 and 0/2 are in use.

```
SW1(config)#interface range f0/3 - 24
SW1(config-if-range)#shutdown
```

```
SW1(config-if-range)#interface range g0/1 - 2
SW1(config-if-range)#shutdown
```

Port Security Configuration

- 2) Configure port security on interface FastEthernet 0/1. Allow a maximum of two MAC addresses and manually add PC1's MAC address to the configuration.

We need to discover PC1's MAC address. We can get this information from the PC itself or from the switch. Use 'ipconfig /all' to find it on the PC.

```
C:\>ipconfig /all

FastEthernet0 Connection: (default port)

  Connection-specific DNS Suffix...:
  Physical Address.....: 0000.1111.1111
  Link-local IPv6 Address....: FE80::200:CF:FEA0:A359
  IP Address.....: 10.10.10.10
  Subnet Mask.....: 255.255.255.0
  Default Gateway.....: 0.0.0.0
  DNS Servers.....: 0.0.0.0
  DHCP Servers.....: 0.0.0.0
  DHCPv6 IAID.....: 28262
  DHCPv6 Client DUID.....:
  00-01-00-01-61-91-76-98-00-00-11-11-11-11
```

Use 'show mac address-table' to find it on the switch. Use ping to generate some traffic from the PC if it does not show up in the MAC address table.

```
C:\>ping 10.10.10.11

Pinging 10.10.10.11 with 32 bytes of data:

Reply from 10.10.10.11: bytes=32 time=14ms TTL=128
Reply from 10.10.10.11: bytes=32 time<1ms TTL=128
Reply from 10.10.10.11: bytes=32 time=19ms TTL=128
Reply from 10.10.10.11: bytes=32 time<1ms TTL=128

Ping statistics for 10.10.10.11:
  Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
  Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 19ms, Average = 8ms
```

```
SW1#show mac address-table
      Mac Address Table
-----
Vlan      Mac Address          Type      Ports
----      -----
1        0000.1111.1111      DYNAMIC   Fa0/1
1        0000.2222.2222      DYNAMIC   Fa0/2
```

You need to make the interface an access port before the switch will accept Port Security configuration. No VLANs are configured on the switch or specified in the lab task so leave it in the default VLAN 1.

```
SW1(config)#int f0/1
SW1(config-if)#switch mode access
```

Add the Port Security configuration.

```
SW1(config-if)#switchport port-security
SW1(config-if)#switchport port-security maximum 2
SW1(config-if)#switchport port-security mac-address
0000.1111.1111
```

- 3) Enable Port Security on interface FastEthernet 0/2 with the default settings.

```
SW1(config)#int f0/2
SW1(config-if)#switch mode access
SW1(config-if)#switchport port-security
```

4) Use a 'show port-security' command to verify the MAC address on PC2.

You may have to ping from PC2 first to generate some traffic.

```
C:\>ping 10.10.10.10

Pinging 10.10.10.10 with 32 bytes of data:

Reply from 10.10.10.10: bytes=32 time=1ms TTL=128
Reply from 10.10.10.10: bytes=32 time<1ms TTL=128
Reply from 10.10.10.10: bytes=32 time=13ms TTL=128
Reply from 10.10.10.10: bytes=32 time=17ms TTL=128

Ping statistics for 10.10.10.10:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 17ms, Average = 7ms
```

```
SW1#show port-security address
          Secure Mac Address Table
-----
Vlan  Mac Address  Type          Ports          Remaining Age
          (mins)
-----
1    0000.1111.1111  SecureConfigured  FastEthernet0/1
1    0000.2222.2222  DynamicConfigured  FastEthernet0/2
-----
Total Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 1024
```

PC2's MAC address is 0000.2222.2222

5) Verify the full Port Security configuration on both interfaces.

```
SW1#show port-security int f0/1
Port Security : Enabled
Port Status : Secure-up
Violation Mode : Shutdown
Aging Time : 0 mins
Aging Type : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses : 2
Total MAC Addresses : 1
Configured MAC Addresses : 1
Sticky MAC Addresses : 0
Last Source Address:Vlan : 0000.1111.1111:1
Security Violation Count : 0
```

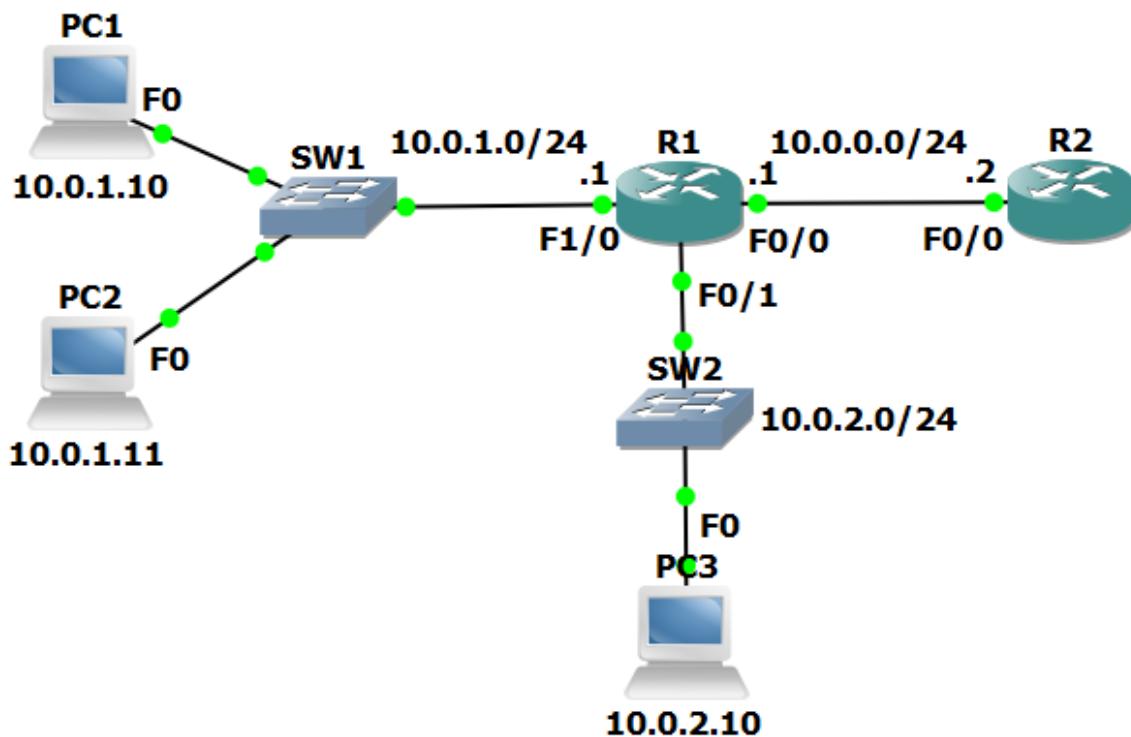
```
SW1#show port-security int f0/2
Port Security : Enabled
Port Status : Secure-up
Violation Mode : Shutdown
Aging Time : 0 mins
Aging Type : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses : 1
Total MAC Addresses : 1
Configured MAC Addresses : 0
Sticky MAC Addresses : 0
Last Source Address:Vlan : 0000.2222.2222:1
Security Violation Count : 0
```

28-1 ACL Configuration – Lab Exercise

In this lab you will configure standard and extended Access Control Lists.

The routers and PCs have been configured with their network addressing settings, and R2 has a static route for the internal 10.0.1.0/24 and 10.0.2.0/24 networks.

Lab Topology



Load the Startup Configurations

[Download the '28-1 ACL Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Numbered Standard ACL

- 1) Verify that all PCs have connectivity to each other, to R1 and to R2.
- 2) Configure and apply a numbered standard ACL on R1 which denies traffic from all hosts in the 10.0.2.0/24 subnet to R2.

The PCs in the 10.0.1.0/24 and 10.0.2.0/24 subnets must maintain connectivity to each other.

The PCs in the 10.0.1.0/24 subnet must maintain connectivity to R2.

- 3) Test that traffic is secured exactly as required.

Verify PC1 and PC2 can ping R2.

PC3 cannot ping R2.

PC3 can ping PC1 and PC2.

Numbered Extended ACL

- 4) Configure and apply a numbered extended ACL on R1 which permits Telnet access from PC1 to R2. Telnet to R2 must be denied for all other PCs in the network.

All other connectivity must be maintained.

Do not change the existing ACL.

Telnet access has already been enabled on R2. The password is 'Flackbox'

- 5) Test that traffic is secured exactly as required. Use the command 'telnet 10.0.0.2' from the PCs to test and the password 'Flackbox'. Type 'exit' to leave the Telnet session.

Verify that PC1 can ping and Telnet to R2.

PC2 can ping R2 but not Telnet to it.

PC3 cannot ping or Telnet to R2.

The PCs can all ping each other.

- 6) How many Telnet packets were permitted by the ACL?

Named Extended ACL

7) Remove the numbered extended ACL you just configured from the interface. Do not delete the ACL.

8) Configure and apply a named extended ACL on R1 as follows:

Permit Telnet from PC1 to R2. Telnet to R2 must be denied for all other PCs in the network.

Permit ping from PC2 to R2. Ping to R2 must be denied for all other PCs in the network.

All other connectivity must be maintained.

9) Test that traffic is secured exactly as required.

Verify that PC1 cannot ping R2 but can Telnet to it.

PC2 can ping R2 but cannot Telnet to it.

PC3 cannot ping or Telnet to R2.

The PCs can all ping each other.

28-1 ACL Configuration - Answer Key

In this lab you will configure standard and extended Access Control Lists.

The routers and PCs have been configured with their network addressing settings, and R2 has a static route for the internal 10.0.1.0/24 and 10.0.2.0/24 networks.

Numbered Standard ACL

- 1) Verify that all PCs have connectivity to each other, to R1 and to R2.

From PC1, ping PC2, PC3, R1 and R2.

```
C:\>ping 10.0.1.11
```

```
Pinging 10.0.1.11 with 32 bytes of data:
```

```
Reply from 10.0.1.11: bytes=32 time=1ms TTL=128
Reply from 10.0.1.11: bytes=32 time<1ms TTL=128
Reply from 10.0.1.11: bytes=32 time<1ms TTL=128
Reply from 10.0.1.11: bytes=32 time<1ms TTL=128
```

```
Ping statistics for 10.0.1.11:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

```
C:\>ping 10.0.2.10
```

```
Pinging 10.0.2.10 with 32 bytes of data:
```

```
Request timed out.
Reply from 10.0.2.10: bytes=32 time<1ms TTL=127
Reply from 10.0.2.10: bytes=32 time<1ms TTL=127
Reply from 10.0.2.10: bytes=32 time<1ms TTL=127
```

```
Ping statistics for 10.0.2.10:
Packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

```
C:\>ping 10.0.1.1
```

```
Pinging 10.0.1.1 with 32 bytes of data:
```

```
Reply from 10.0.1.1: bytes=32 time<1ms TTL=255
```

```
Ping statistics for 10.0.1.1:
```

```
Bytes: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

```
C:\>ping 10.0.0.2
```

```
Pinging 10.0.0.2 with 32 bytes of data:
```

```
Request timed out.
```

```
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254
```

```
Ping statistics for 10.0.0.2:
```

```
Bytes: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

We have already verified connectivity between the PCs in both subnets. Ping R1 and R2 from PC3.

```
C:\>ping 10.0.2.1
```

```
Pinging 10.0.2.1 with 32 bytes of data:
```

```
Reply from 10.0.2.1: bytes=32 time=1ms TTL=255
Reply from 10.0.2.1: bytes=32 time<1ms TTL=255
Reply from 10.0.2.1: bytes=32 time=1ms TTL=255
Reply from 10.0.2.1: bytes=32 time<1ms TTL=255
```

```
Ping statistics for 10.0.2.1:
```

```
Bytes: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

```
C:\>ping 10.0.0.2

Pinging 10.0.0.2 with 32 bytes of data:

Request timed out.
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254
Reply from 10.0.0.2: bytes=32 time=1ms TTL=254
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254

Ping statistics for 10.0.0.2:
Packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

- 2) Configure and apply a numbered standard ACL on R1 which denies traffic from all hosts in the 10.0.2.0/24 subnet to R2.

The PCs in the 10.0.1.0/24 and 10.0.2.0/24 subnets must maintain connectivity to each other.

The PCs in the 10.0.1.0/24 subnet must maintain connectivity to R2.

The task specifies that a numbered standard ACL be used on R1. This checks the source address only. This prevents us from configuring an ACL inbound on the F0/1 interface which denies traffic from the 10.0.2.0/24 subnet to R2 but permits it to the 10.0.1.0/24 network – that would require an extended ACL.

The only way the task can be completed is by applying the ACL outbound on the F0/0 interface facing R2.

Configure a numbered standard ACL that denies traffic from 10.0.2.0/24. Do not forget to permit from 10.0.1.0/24 as the implicit deny any at the end of the ACL would block the traffic otherwise.

```
R1(config)#access-list 1 deny 10.0.2.0 0.0.0.255
R1(config)#access-list 1 permit 10.0.1.0 0.0.0.255
```

Do not forget to apply the ACL to the interface.

```
R1(config)#interface f0/0
R1(config-if)#ip access-group 1 out
```

3) Test that traffic is secured exactly as required.

Verify PC1 and PC2 can ping R2.
PC3 cannot ping R2.
PC3 can ping PC1 and PC2.

PC1 and PC2 should be able to ping R2.

```
C:\>ping 10.0.0.2
```

Pinging 10.0.0.2 with 32 bytes of data:

```
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254
```

Ping statistics for 10.0.0.2:

```
packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

PC3 should not be able to ping R2.

```
C:\>ping 10.0.0.2
```

Pinging 10.0.0.2 with 32 bytes of data:

```
Reply from 10.0.2.1: Destination host unreachable.
```

Ping statistics for 10.0.0.2:

```
packets: Sent = 4, Received = 0, Lost = 4 (100% loss)
```

PC3 should be able to ping PC1 and PC2.

```
C:\>ping 10.0.1.10
```

Pinging 10.0.1.10 with 32 bytes of data:

```
Reply from 10.0.1.10: bytes=32 time<1ms TTL=127
Reply from 10.0.1.10: bytes=32 time<1ms TTL=127
Reply from 10.0.1.10: bytes=32 time<1ms TTL=127
Reply from 10.0.1.10: bytes=32 time=1ms TTL=127
```

Ping statistics for 10.0.1.10:

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

```
C:\>ping 10.0.1.11
```

Pinging 10.0.1.11 with 32 bytes of data:

```
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127
Reply from 10.0.1.11: bytes=32 time=5ms TTL=127
Reply from 10.0.1.11: bytes=32 time=1ms TTL=127
```

Ping statistics for 10.0.1.11:

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 5ms, Average = 1ms
```

Numbered Extended ACL

- 4) Configure and apply a numbered extended ACL on R1 which permits Telnet access from PC1 to R2. Telnet to R2 must be denied for all other PCs in the network.

All other connectivity must be maintained.

Do not change the existing ACL.

Telnet access has already been enabled on R2. The password is 'Flackbox'

All traffic from the 10.0.2.0/24 subnet to R2 is already denied by the numbered standard ACL we configured.

We need to create an ACL which will allow Telnet traffic to R2 at 10.0.0.2 from PC1 at 10.0.1.10 but deny it from all other hosts in the 10.0.1.0/24 subnet.

The implicit deny all at the end of the ACL would block Telnet traffic from the other hosts to R2, but it would also block all other traffic from the 10.0.1.0/24 subnet including other applications to R2 and traffic to the 10.0.2.0/24 subnet. We need to explicitly block the Telnet traffic but allow other traffic.

```
R1(config)#access-list 100 permit tcp host 10.0.1.10 host  
10.0.0.2 eq telnet  
R1(config)#access-list 100 deny tcp 10.0.1.0 0.0.0.255 host  
10.0.0.2 eq telnet  
R1(config)#access-list 100 permit ip any any
```

The access list should be applied as close to the source as possible, so apply it inbound on interface F1/0. We already have an outbound ACL on F0/0 so we could not apply it there anyway.

```
R1(config)#interface f1/0  
R1(config-if)#ip access-group 100 in
```

- 5) Test that traffic is secured exactly as required. Use the command ‘telnet 10.0.0.2’ from the PCs to test and the password ‘Flackbox’. Type ‘exit’ to leave the Telnet session.

Verify that PC1 can ping and Telnet to R2.
PC2 can ping R2 but not Telnet to it.
PC3 cannot ping or Telnet to R2.
The PCs can all ping each other.

PC1 should be able to Telnet to R2.

```
PC1#telnet 10.0.0.2  
Trying 10.0.0.2 ... Open
```

```
User Access Verification  
Password: Flackbox  
R2>
```

PC2 should be able to ping but not Telnet to R2.

```
C:\>ping 10.0.0.2
```

```
Pinging 10.0.0.2 with 32 bytes of data:
```

```
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254
```

```
Ping statistics for 10.0.0.2:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

```
C:\>telnet 10.0.0.2
```

```
Trying 10.0.0.2 ...
% Connection timed out; remote host not responding
```

PC3 should not be able to ping or Telnet to R2.

```
C:\>ping 10.0.0.2
```

```
Pinging 10.0.0.2 with 32 bytes of data:
```

```
Reply from 10.0.2.1: Destination host unreachable.
```

```
Ping statistics for 10.0.0.2:
```

```
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
```

```
C:\>telnet 10.0.0.2
```

```
Trying 10.0.0.2 ...
% Connection timed out; remote host not responding
```

PC3 should be able to ping PC1 and PC2.

```
C:\>ping 10.0.1.10
```

Pinging 10.0.1.10 with 32 bytes of data:

```
Reply from 10.0.1.10: bytes=32 time<1ms TTL=127
```

Ping statistics for 10.0.1.10:

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

```
C:\>ping 10.0.1.11
```

Pinging 10.0.1.11 with 32 bytes of data:

```
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127
Reply from 10.0.1.11: bytes=32 time=1ms TTL=127
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127
```

Ping statistics for 10.0.1.11:

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

6) How many Telnet packets were permitted by the ACL?

Enter 'show access-list' to check the hit counts. Note that your values may be different.

```
R1#show access-lists 100
Extended IP access list 100
 permit tcp host 10.0.1.10 host 10.0.0.2 eq telnet (23 match(es))
 deny tcp 10.0.1.0 0.0.0.255 host 10.0.0.2 eq telnet (12 match(es))
 permit ip any any (12 match(es))
```

Named Extended ACL

- 7) Remove the numbered extended ACL you just configured from the interface. Do not delete the ACL.

```
R1(config)#int f1/0
R1(config-if)#no ip access-group 100 in
```

- 8) Configure and apply a named extended ACL on R1 as follows:

Permit Telnet from PC1 to R2. Telnet to R2 must be denied for all other PCs in the network.

Permit ping from PC2 to R2. Ping to R2 must be denied for all other PCs in the network.

All other connectivity must be maintained.

All traffic from the 10.0.2.0/24 subnet to R2 is already denied by the numbered standard ACL we configured.

We do need to configure an ACL to secure traffic from the 10.0.1.0/24 subnet.

```
R1(config)#ip access-list extended F1/0_in
R1(config-ext-nacl)#permit tcp host 10.0.1.10 host 10.0.0.2
eq telnet
R1(config-ext-nacl)#deny tcp 10.0.1.0 0.0.0.255 host
10.0.0.2 eq telnet
R1(config-ext-nacl)#permit icmp host 10.0.1.11 host
10.0.0.2 echo
R1(config-ext-nacl)#deny icmp 10.0.1.0 0.0.0.255 host
10.0.0.2 echo
R1(config-ext-nacl)#permit ip any any
```

Don't forget to apply the ACL to the interface.

```
R1(config)#int f1/0
R1(config-if)#ip access-group F1/0_in in
```

9) Test that traffic is secured exactly as required.

Verify that PC1 cannot ping R2 but can Telnet to it.

PC2 can ping R2 but cannot Telnet to it.

PC3 cannot ping or Telnet to R2.

The PCs can all ping each other.

PC1 cannot ping R2 but can Telnet to it.

```
C:\>ping 10.0.0.2
```

```
Pinging 10.0.0.2 with 32 bytes of data:
```

```
Reply from 10.0.1.1: Destination host unreachable.
```

```
Ping statistics for 10.0.0.2:
```

```
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
```

```
C:\>telnet 10.0.0.2
```

```
Trying 10.0.0.2 ...Open
```

User Access Verification

Password:

```
R2>
```

PC2 can ping R2 but cannot Telnet to it.

```
C:\>ping 10.0.0.2
```

```
Pinging 10.0.0.2 with 32 bytes of data:
```

```
Reply from 10.0.0.2: bytes=32 time<1ms TTL=254
```

```
Ping statistics for 10.0.0.2:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

```
C:\>telnet 10.0.0.2
```

```
Trying 10.0.0.2 ...
```

```
% Connection timed out; remote host not responding
```

PC3 cannot ping or Telnet to R2.

```
C:\>ping 10.0.0.2
```

```
Pinging 10.0.0.2 with 32 bytes of data:
```

```
Reply from 10.0.2.1: Destination host unreachable.  
Reply from 10.0.2.1: Destination host unreachable.  
Reply from 10.0.2.1: Destination host unreachable.  
Reply from 10.0.2.1: Destination host unreachable.
```

```
Ping statistics for 10.0.0.2:
```

```
packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
```

```
C:\>telnet 10.0.0.2
```

```
Trying 10.0.0.2 ...  
% Connection timed out; remote host not responding
```

PC3 should be able to ping PC1 and PC2.

```
C:\>ping 10.0.1.10
```

```
Pinging 10.0.1.10 with 32 bytes of data:
```

```
Reply from 10.0.1.10: bytes=32 time<1ms TTL=127  
Reply from 10.0.1.10: bytes=32 time=1ms TTL=127  
Reply from 10.0.1.10: bytes=32 time<1ms TTL=127  
Reply from 10.0.1.10: bytes=32 time<1ms TTL=127
```

```
Ping statistics for 10.0.1.10:
```

```
packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

```
C:\>ping 10.0.1.11
```

```
Pinging 10.0.1.11 with 32 bytes of data:
```

```
Reply from 10.0.1.11: bytes=32 time=1ms TTL=127  
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127  
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127  
Reply from 10.0.1.11: bytes=32 time<1ms TTL=127
```

```
Ping statistics for 10.0.1.11:
```

```
packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

29-1 NAT Configuration – Lab Exercise

In this lab you will configure static and dynamic Network Address Translation, and Port Address Translation.

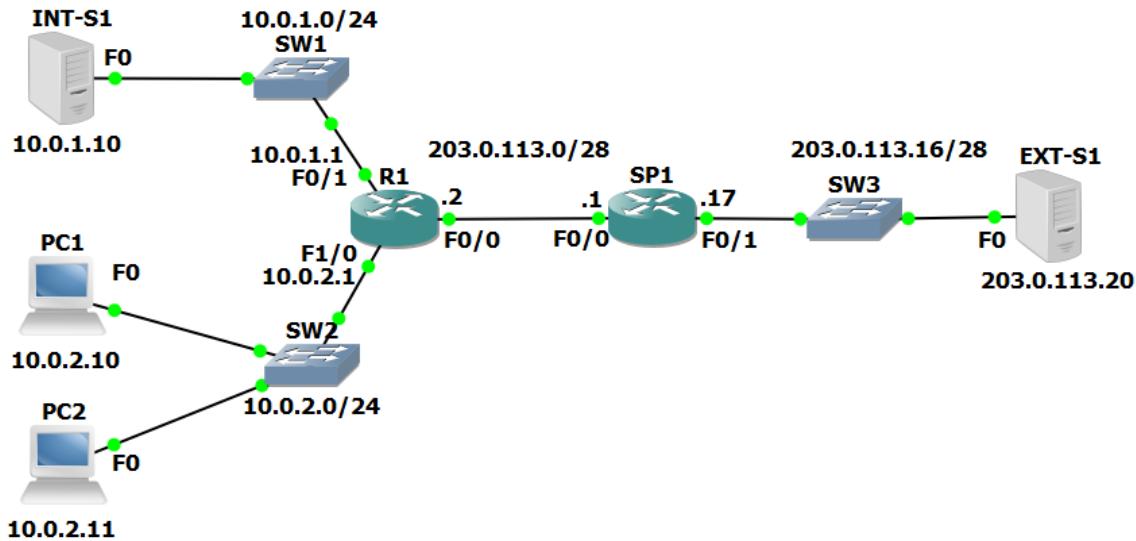
The routers, servers and PCs have been configured with their network addressing settings.

R1 is the WAN edge router at your company, it has a default route pointing to the Service Provider router SP1.

You have bought the range of public IP addresses 203.0.113.0/28.
203.0.113.1 is assigned to the Service Provider default gateway SP1,
203.0.113.2 is assigned to the Internet facing F0/0 interface on your router R1.
203.0.113.3 – 203.0.113.14 are your remaining available public IP addresses.

Note that entries in the NAT translation table age out quickly. Send the traffic again if you do not see the expected results in the table.

Lab Topology

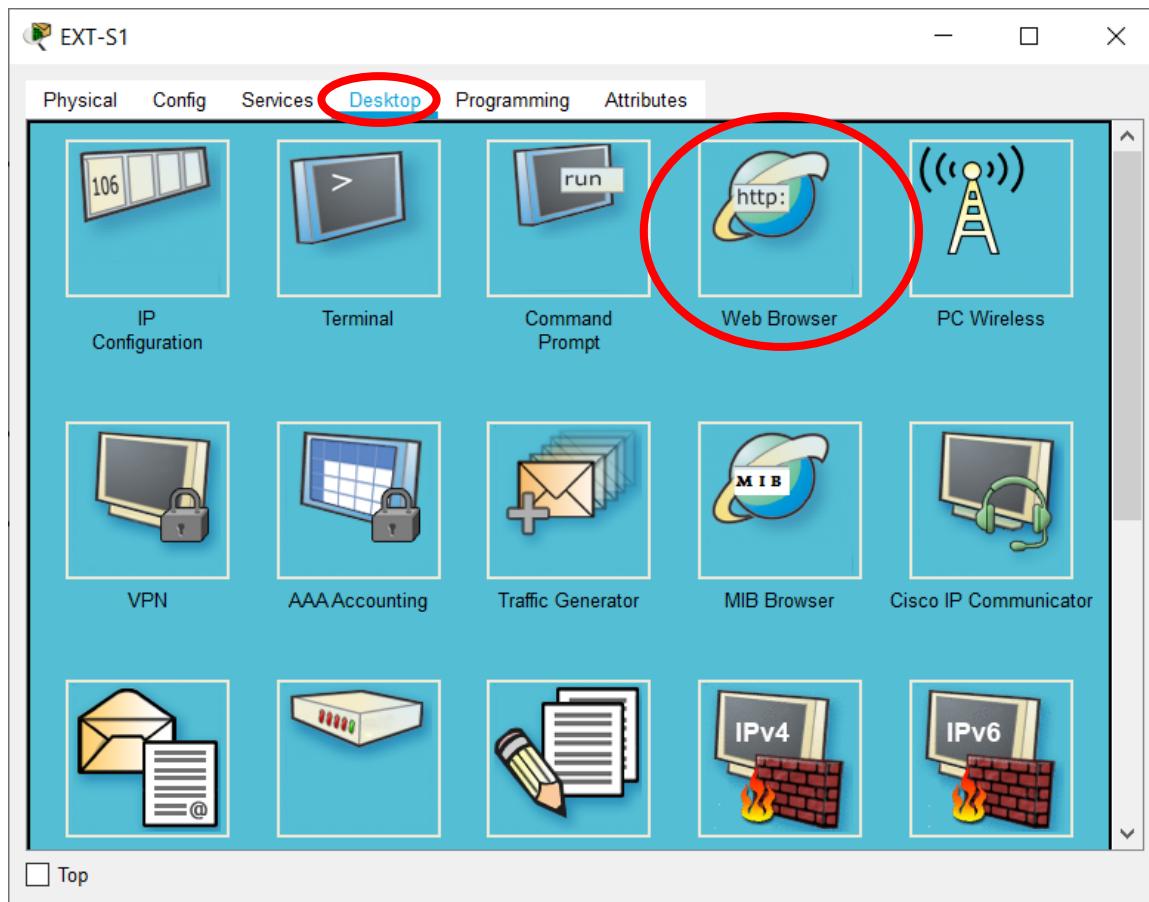


Load the Startup Configurations

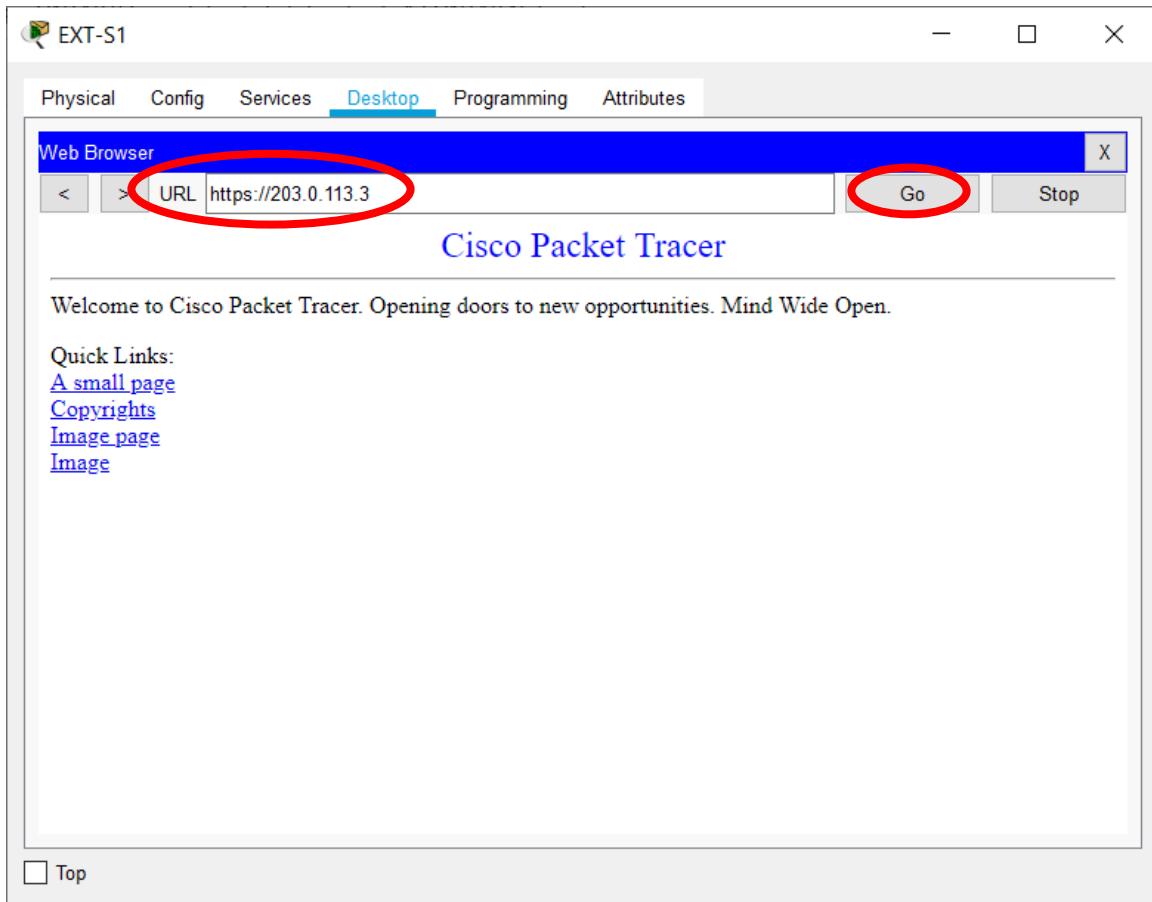
[Download the '29-1 NAT Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Static NAT

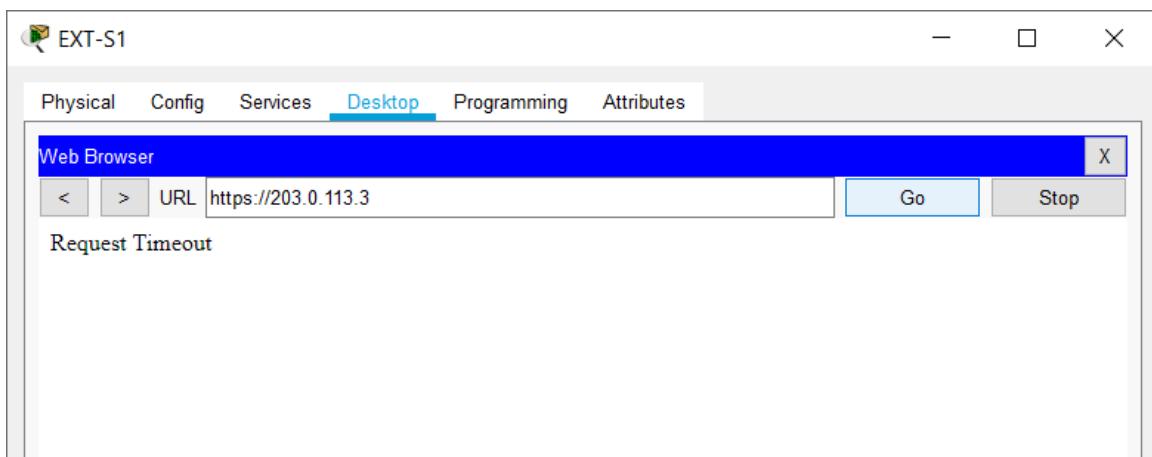
- 1) Int-S1 is your company's web server. It must be reachable from external customers browsing the Internet. Configure NAT on R1 so that external customers can reach the server using the public IP address 203.0.113.3. Do not change any IP addressing or routing information.
- 2) Ping Ext-S1 from Int-S1 to check the NAT rule is working and you have connectivity.
- 3) On Ext-S1, click on 'Desktop' then open 'Web Browser'. Open the NAT'd public IP address of Int-S1 at <https://203.0.113.3> in the browser to verify that external customers on the Internet can reach your web server.



You will see the output shown below if your configuration is working:



You will see a 'Request Timeout' error message if your configuration is **not** working (you need to fix it):



- 4) Verify the connection appears in the NAT translation table. Note that entries age out quickly so generate the traffic again if you did not check the table quickly enough.

Dynamic NAT

- 5) Configure NAT on R1 so that PCs in the 10.0.2.0/24 subnet have connectivity to external networks on a first come first served basis. Assign global addresses from the range 203.0.113.4 to 203.0.113.12. Do not enable Port Address Translation.
- 6) Turn on NAT debugging on R1. Ping Ext-S1 from PC1. View the debug output on R1. You should see NAT entries for the 5 pings. Which global address was PC1 translated to?
- 7) Verify the ping connection appears in the NAT translation table.
- 8) When all the addresses in the pool 203.0.113.4 to 203.0.113.12 have been allocated, what will happen when the next PC tries to send traffic to an external host?
- 9) Enable Port Address Translation so that the last IP address in the range can be reused when all IP addresses have been allocated to clients.
- 10) Cleanup: Completely remove the access list and all NAT configuration from R1. Use the commands 'show run | section nat' and 'show access-list' to verify all configuration is removed.

Port Address Translation PAT

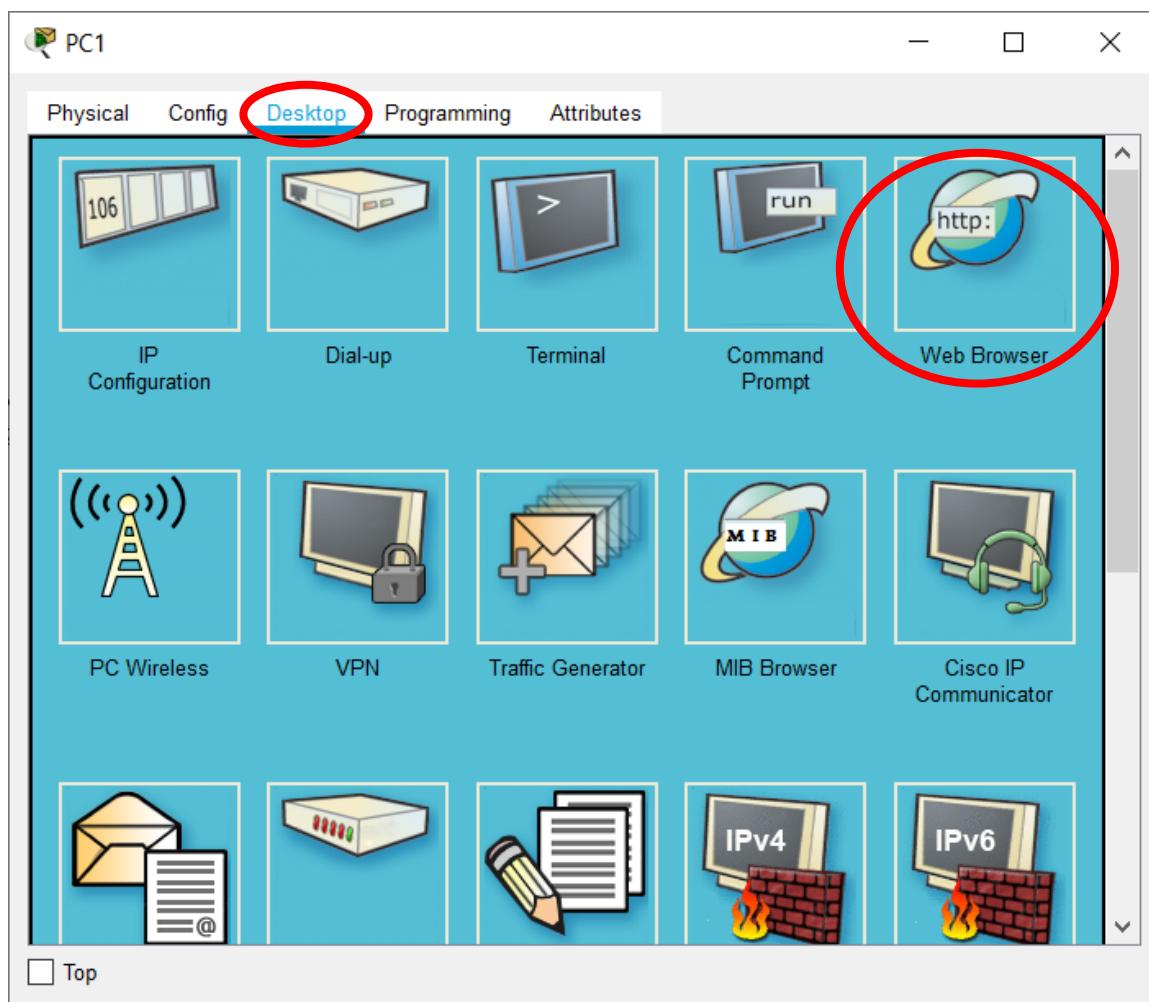
- 11) Your company no longer has a range of public IP addresses. Instead you will receive a single public IP address via DHCP from your service provider.

Shutdown interface F0/0 on R1 and remove its IP address. Reconfigure it to receive its IP address via DHCP from the service provider router SP1.

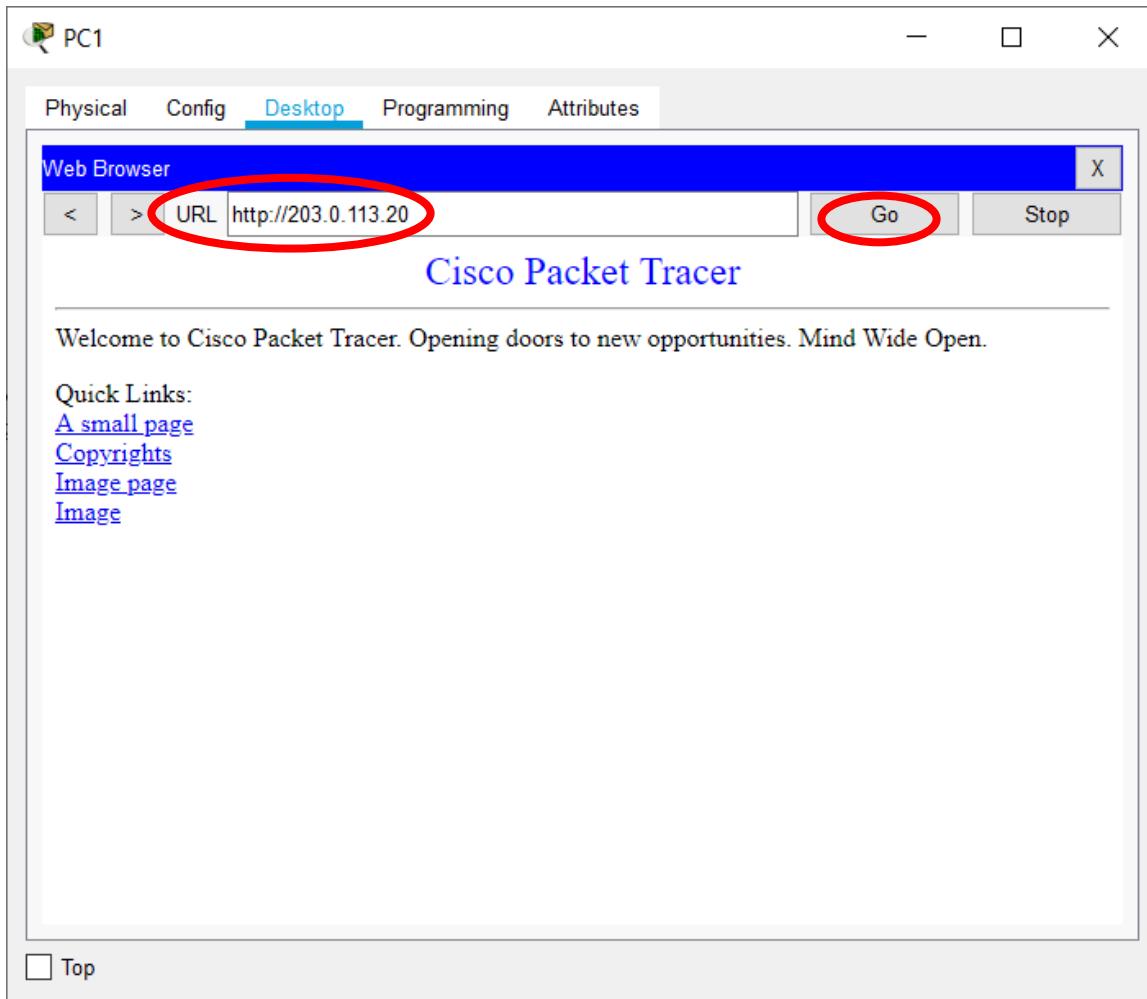
12) Bring the interface back up and wait for DHCP. What IP address is it assigned?

13) Configure NAT on R1 so that PCs in the 10.0.2.0/24 subnet have connectivity to external networks on a first come first served basis. Allow multiple PCs to reuse the IP address on interface F0/0 on R1. Ensure that NAT continues to work if the DHCP address assigned by the provider changes.

14) Ensure NAT debugging on R1 is still enabled. On PC1, click on 'Desktop' then open 'Web Browser'. Open the public IP address of Ext-S1 at <https://203.0.113.20> in the browser to verify that PAT is working and you can reach external servers.



You will see the output shown below if your configuration is working:



You will see a 'Request Timeout' error message if your configuration is **not** working (you need to fix it).

- 15) View the debug output on R1. Which global IP address was PC1 translated to?
- 16) On PC2, click on 'Desktop' then open 'Web Browser'. Open the public IP address of Ext-S1 at <http://203.0.113.20> in the browser. Which global IP address is PC2 translated to?
- 17) Verify the connections in the NAT translation table.
- 18) Show the NAT statistics on R1.

29-1 NAT Configuration - Answer Key

In this lab you will configure static and dynamic Network Address Translation, and Port Address Translation.

The routers, servers and PCs have been configured with their network addressing settings.

R1 is the WAN edge router at your company, it has a default route pointing to the Service Provider router SP1.

You have bought the range of public IP addresses 203.0.113.0/28. 203.0.113.1 is assigned to the Service Provider default gateway SP1, 203.0.113.2 is assigned to the Internet facing F0/0 interface on your router R1. 203.0.113.3 – 203.0.113.14 are your remaining available public IP addresses.

Note that entries in the NAT translation table age out quickly. Send the traffic again if you do not see the expected results in the table.

Static NAT

- 1) Int-S1 is your company's web server. It must be reachable from external customers browsing the Internet. Configure NAT on R1 so that external customers can reach the server using the public IP address 203.0.113.3. Do not change any IP addressing or routing information.

Int-S1 needs a fixed public IP address so we must configure static NAT. Configure the F0/0 interface facing the Internet as the NAT outside interface.

```
R1(config)#int f0/0
R1(config-if)#ip nat outside
```

Configure the F0/1 interface facing Int-S1 as a NAT inside interface.

```
R1(config)#int f0/1
R1(config-if)#ip nat inside
```

Configure a static NAT rule mapping the inside local address 10.0.1.10 to the global IP address 203.0.113.3.

```
R1(config)#ip nat inside source static 10.0.1.10
203.0.113.3
```

2) Ping Ext-S1 from Int-S1 to check the NAT rule is working and you have connectivity.

```
C:\>ping 203.0.113.20
```

```
Pinging 203.0.113.20 with 32 bytes of data:
```

```
Request timed out.
```

```
Request timed out.
```

```
Reply from 203.0.113.20: bytes=32 time<1ms TTL=126
```

```
Reply from 203.0.113.20: bytes=32 time<1ms TTL=126
```

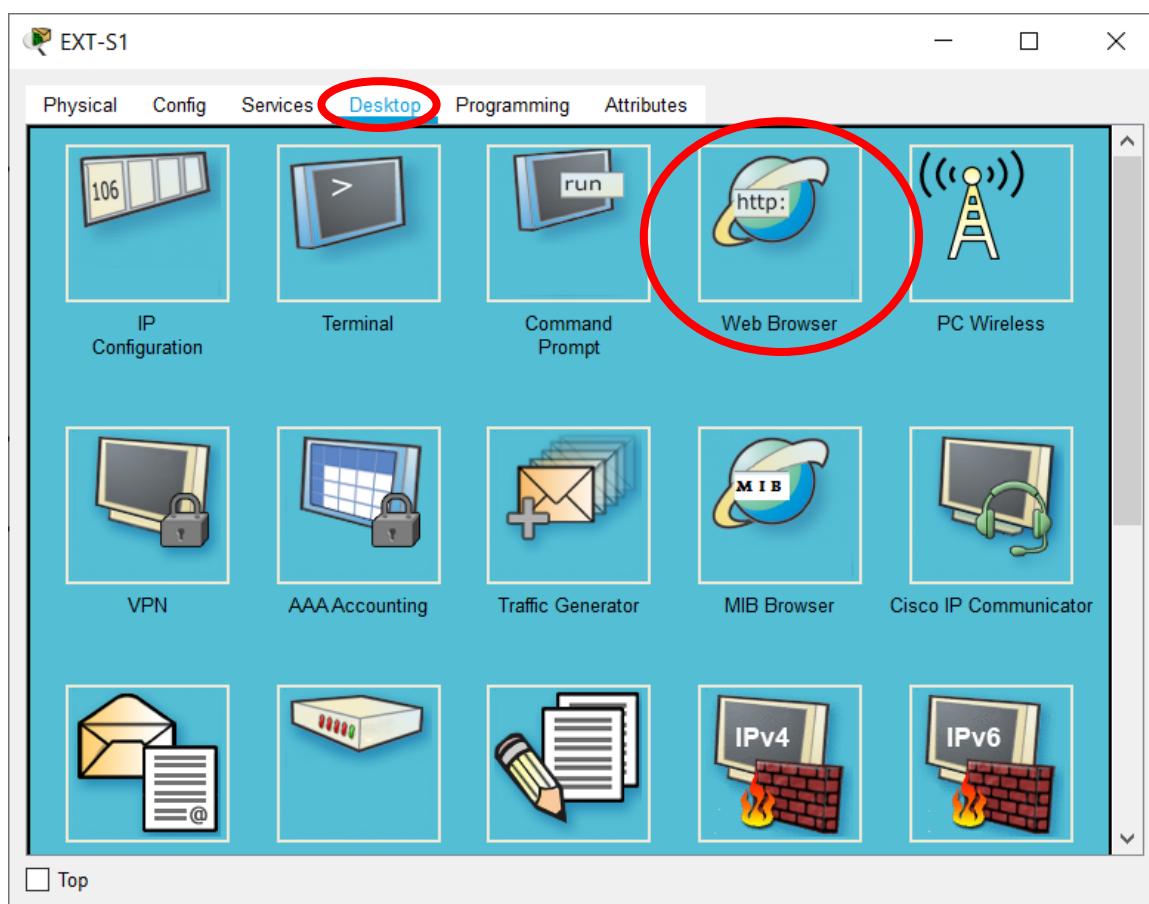
```
Ping statistics for 203.0.113.20:
```

```
Packets: Sent = 4, Received = 2, Lost = 2 (50% loss),
```

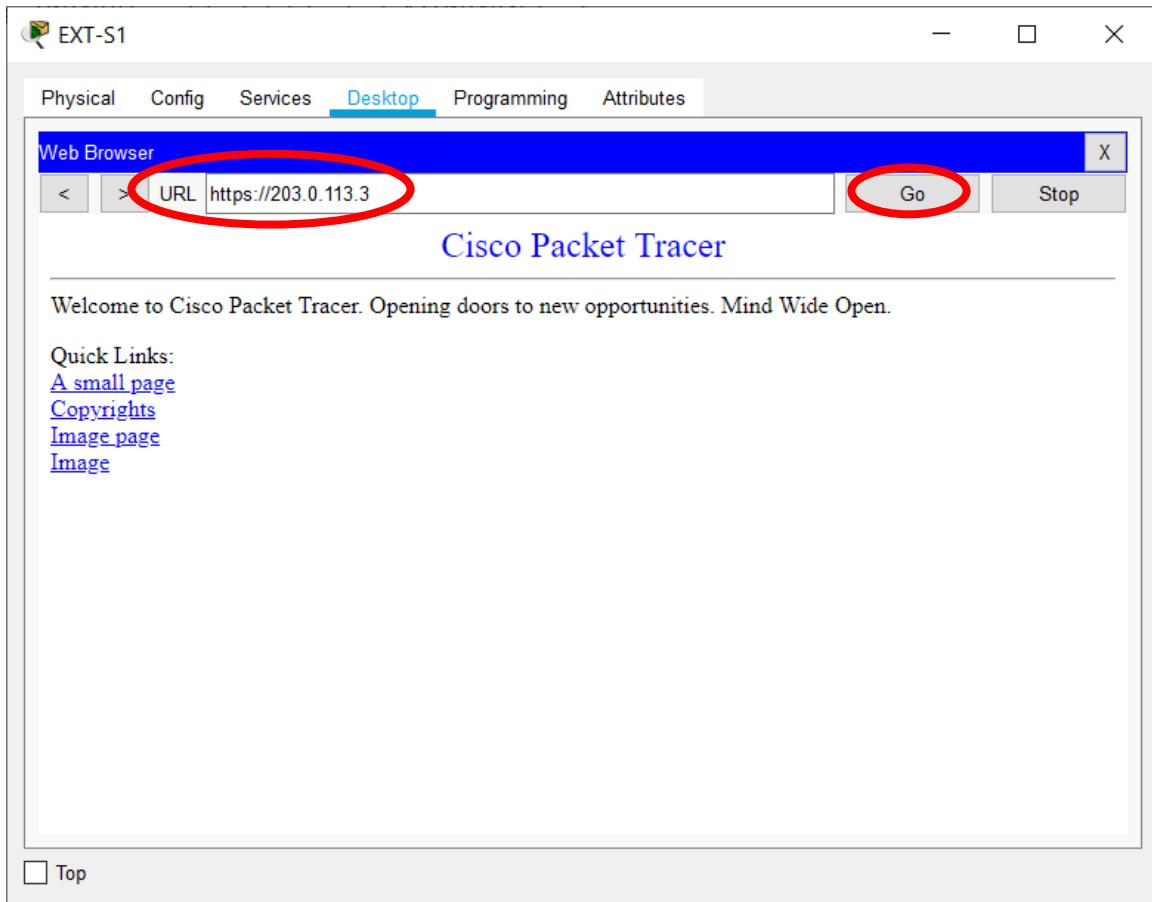
```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

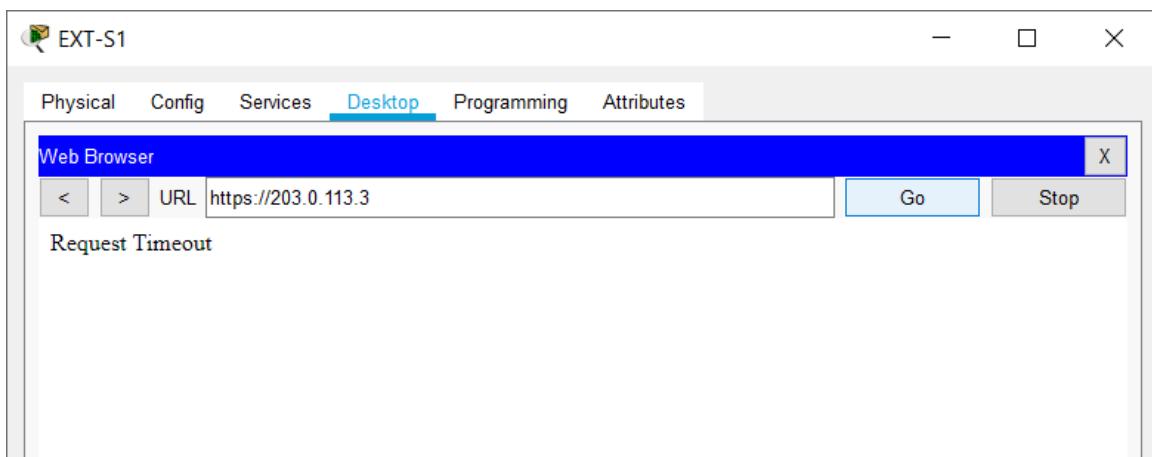
3) On Ext-S1, click on 'Desktop' then open 'Web Browser'. Open the NAT'd public IP address of Int-S1 at <https://203.0.113.3> in the browser to verify that external customers on the Internet can reach your web server.



You will see the output shown below if your configuration is working:



You will see a 'Request Timeout' error message if your configuration is **not** working (you need to fix it):



- 4) Verify the connection appears in the NAT translation table. Note that entries age out quickly so generate the traffic again if you did not check the table quickly enough.

Your output should look similar to below. The source port numbers may be different.

```
R1#sh ip nat translation
Pro Inside global      Inside local      Outside local      Outside global
--- 203.0.113.3        10.0.1.10        ---              ---
tcp 203.0.113.3:443   10.0.1.10:443    203.0.113.20:1027 203.0.113.20:1027
```

Dynamic NAT

- 5) Configure NAT on R1 so that PCs in the 10.0.2.0/24 subnet have connectivity to external networks on a first come first served basis. Assign global addresses from the range 203.0.113.4 to 203.0.113.12. Do not enable Port Address Translation.

We already set the F0/0 interface facing the Internet as the NAT outside interface when we configured static NAT for Int-S1.

```
R1(config)#int f0/0
R1(config-if)#ip nat outside
```

Configure the F1/0 interface facing the PCs as a NAT inside interface.

```
R1(config)#int f1/0
R1(config-if)#ip nat inside
```

Configure the pool of global addresses.

```
R1(config)#ip nat pool Flackbox 203.0.113.4 203.0.113.12
netmask 255.255.255.240
```

Create an access list which references the internal IP addresses we want to translate.

```
R1(config)#access-list 1 permit 10.0.2.0 0.0.0.255
```

Associate the access list with the NAT pool to complete the configuration.

```
R1(config)#ip nat inside source list 1 pool Flackbox
```

6) Turn on NAT debugging on R1. Ping Ext-S1 from PC1. View the debug output on R1. You should see NAT entries for the 4 pings. Which global address was PC1 translated to?

PC1 will be translated to the first available address in the pool, 203.0.113.4.

```
R1#debug ip nat
IP NAT debugging is on
```

On PC1:

```
C:\>ping 203.0.113.20
```

```
Pinging 203.0.113.20 with 32 bytes of data:
```

```
Request timed out.
```

```
Reply from 203.0.113.20: bytes=32 time<1ms TTL=126
```

```
Reply from 203.0.113.20: bytes=32 time<1ms TTL=126
```

```
Reply from 203.0.113.20: bytes=32 time=1ms TTL=126
```

```
Ping statistics for 203.0.113.20:
```

```
packets: Sent = 4, Received = 3, Lost = 1 (25% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

```
R1#
*Mar 01, 03:19:48.1919: NAT: s=10.0.2.10->203.0.113.4, d=203.0.113.20 [1]
R1#
*Mar 01, 03:19:54.1919: NAT: s=10.0.2.10->203.0.113.4, d=203.0.113.20 [2]
*Mar 01, 03:19:54.1919: NAT*: s=203.0.113.20, d=203.0.113.4->10.0.2.10 [30]
R1#
*Mar 01, 03:19:55.1919: NAT: s=10.0.2.10->203.0.113.4, d=203.0.113.20 [3]
*Mar 01, 03:19:55.1919: NAT*: s=203.0.113.20, d=203.0.113.4->10.0.2.10 [31]
R1#
*Mar 01, 03:19:56.1919: NAT: s=10.0.2.10->203.0.113.4, d=203.0.113.20 [4]
*Mar 01, 03:19:56.1919: NAT*: s=203.0.113.20, d=203.0.113.4->10.0.2.10 [32]
```

7) Verify the ping connection appears in the NAT translation table.

```
R1#sh ip nat translation
Pro Inside global      Inside local      Outside local      Outside global
--- 203.0.113.3        10.0.1.10        ---             ---
icmp 203.0.113.4:1  10.0.2.10:1    203.0.113.20:1  203.0.113.20:1
--- 203.0.113.4        10.0.2.10        ---             ---
```

8) When all the addresses in the pool 203.0.113.4 to 203.0.113.12 have been allocated, what will happen when the next PC tries to send traffic to an external host?

It will not get a global IP address because they have all been given out. The traffic will fail. The user could try again after waiting for a translation to time out and its global address to be released back into the pool.

9) Enable Port Address Translation so that the last IP address in the range can be reused when all IP addresses have been allocated to clients.

```
R1#clear ip nat translation *
R1#config t
R1(config)#ip nat inside source list 1 pool Flackbox overload
```

10) Cleanup: Completely remove the access list and all NAT configuration from R1. Use the commands 'show run | section nat' and 'show access-list' to verify all configuration is removed.

```
R1(config)#int f0/0
R1(config-if)#no ip nat outside
R1(config-if)#int f0/1
R1(config-if)#no ip nat inside
R1(config-if)#int f1/0
R1(config-if)#no ip nat inside
R1(config-if)#no ip nat inside source static 10.0.1.10
203.0.113.3
R1(config)#end
R1#clear ip nat translation *
R1#config t
R1(config)#no ip nat inside source list 1 pool Flackbox overload
R1(config)#no ip nat pool Flackbox 203.0.113.4 203.0.113.12
netmask 255.255.255.240
R1(config)#no access-list 1

R1#show run | section nat
R1#

R1#show access-list
R1#
```

Port Address Translation PAT

11) Your company no longer has a range of public IP addresses. Instead you will receive a single public IP address via DHCP from your service provider.

Shutdown interface F0/0 on R1 and remove its IP address. Reconfigure it to receive its IP address via DHCP from the service provider router SP1.

```
R1(config)#int f0/0
R1(config-if)#shutdown
R1(config-if)#no ip address
R1(config-if)#ip address dhcp
```

12) Bring the interface back up and wait for DHCP. What IP address is it assigned?

```
R1(config)#int f0/0
R1(config-if)#no shutdown

R1#sh ip int brief
Interface          IP-Address      OK? Method Status
Protocol
FastEthernet0/0  203.0.113.13  YES  DHCP    up
FastEthernet0/1    10.0.0.1.1    YES  NVRAM   up
FastEthernet1/0    10.0.0.2.1    YES  NVRAM   up
FastEthernet1/1    unassigned    YES  NVRAM   administratively down down
```

13) Configure NAT on R1 so that PCs in the 10.0.2.0/24 subnet have connectivity to external networks on a first come first served basis. Allow multiple PCs to reuse the IP address on interface F0/0 on R1. Ensure that NAT continues to work if the DHCP address assigned by the provider changes.

Configure the F0/0 interface facing the Internet as the NAT outside interface.

```
R1(config)#int f0/0
R1(config-if)#ip nat outside
```

Configure the F1/0 interface facing Int-S1 as a NAT inside interface.

```
R1(config)#int f1/0
R1(config-if)#ip nat inside
```

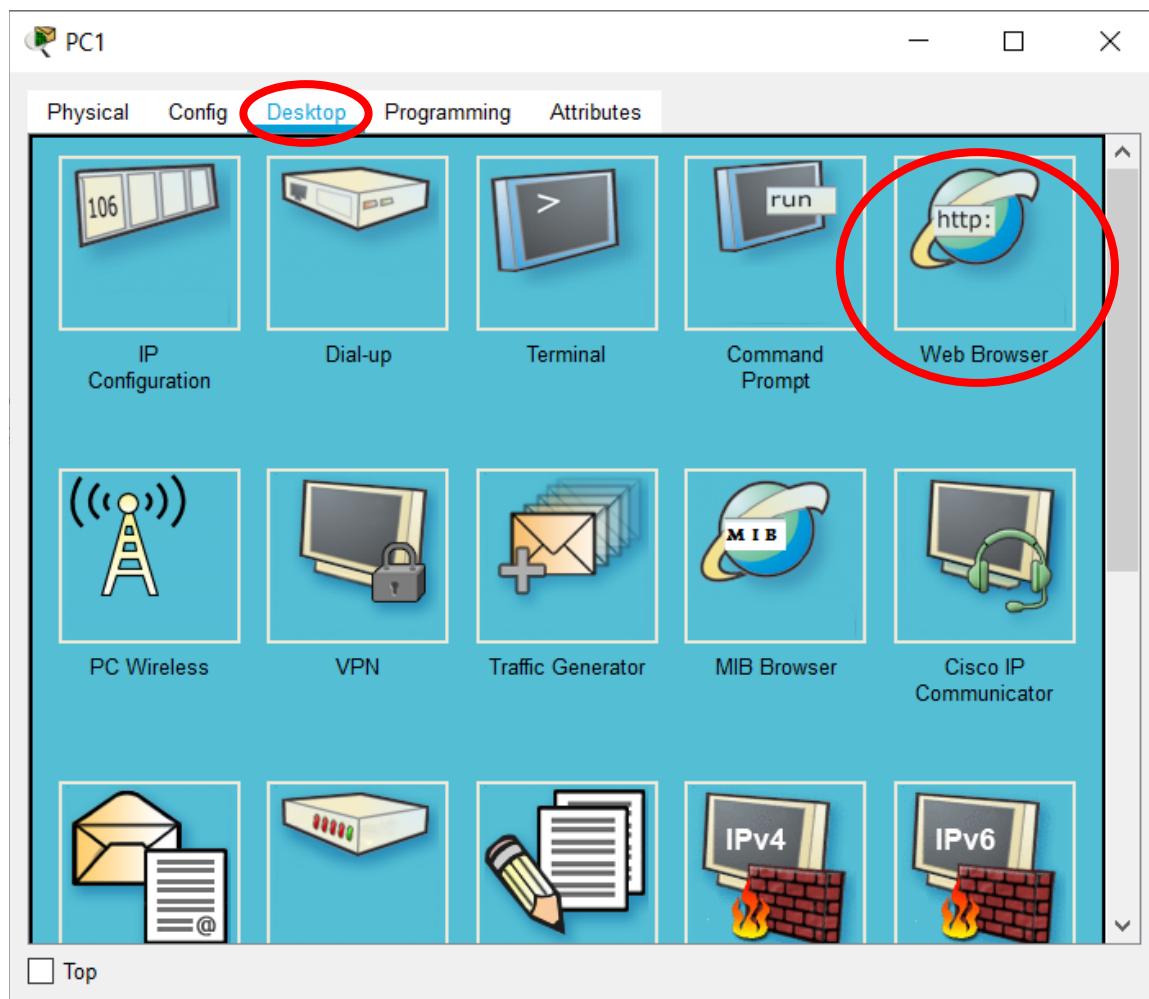
Create an access list which references the internal IP addresses we want to translate.

```
R1(config)#access-list 1 permit 10.0.2.0 0.0.0.255
```

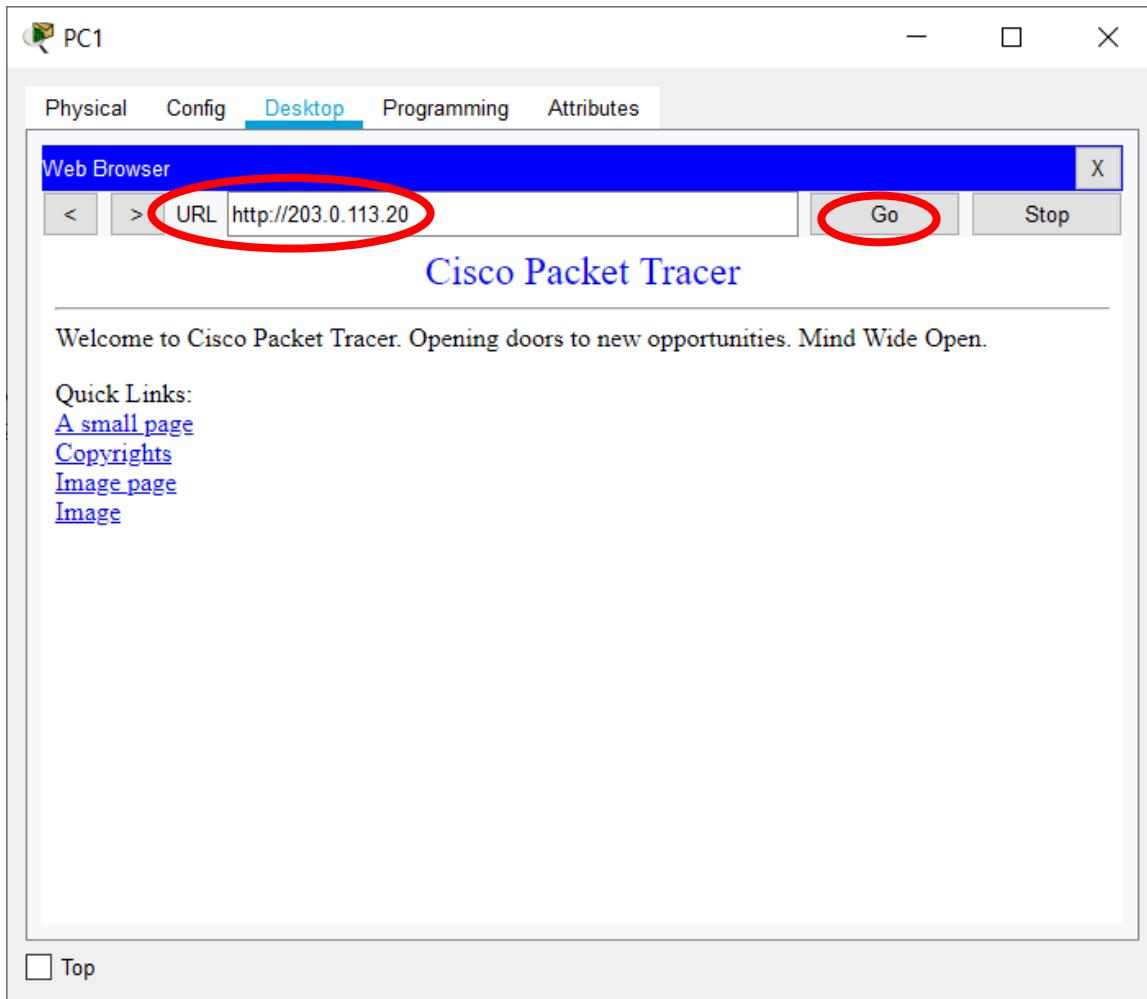
Associate the access list with the outside interface and enable PAT to complete the configuration.

```
R1(config)#ip nat inside source list 1 interface f0/0  
overload
```

14) Ensure NAT debugging on R1 is still enabled. On PC1, click on 'Desktop' then open 'Web Browser'. Open the public IP address of Ext-S1 at <http://203.0.113.20> in the browser to verify that PAT is working and you can reach external servers.



You will see the output shown below if your configuration is working:



You will see a 'Request Timeout' error message if your configuration is **not** working (you need to fix it).

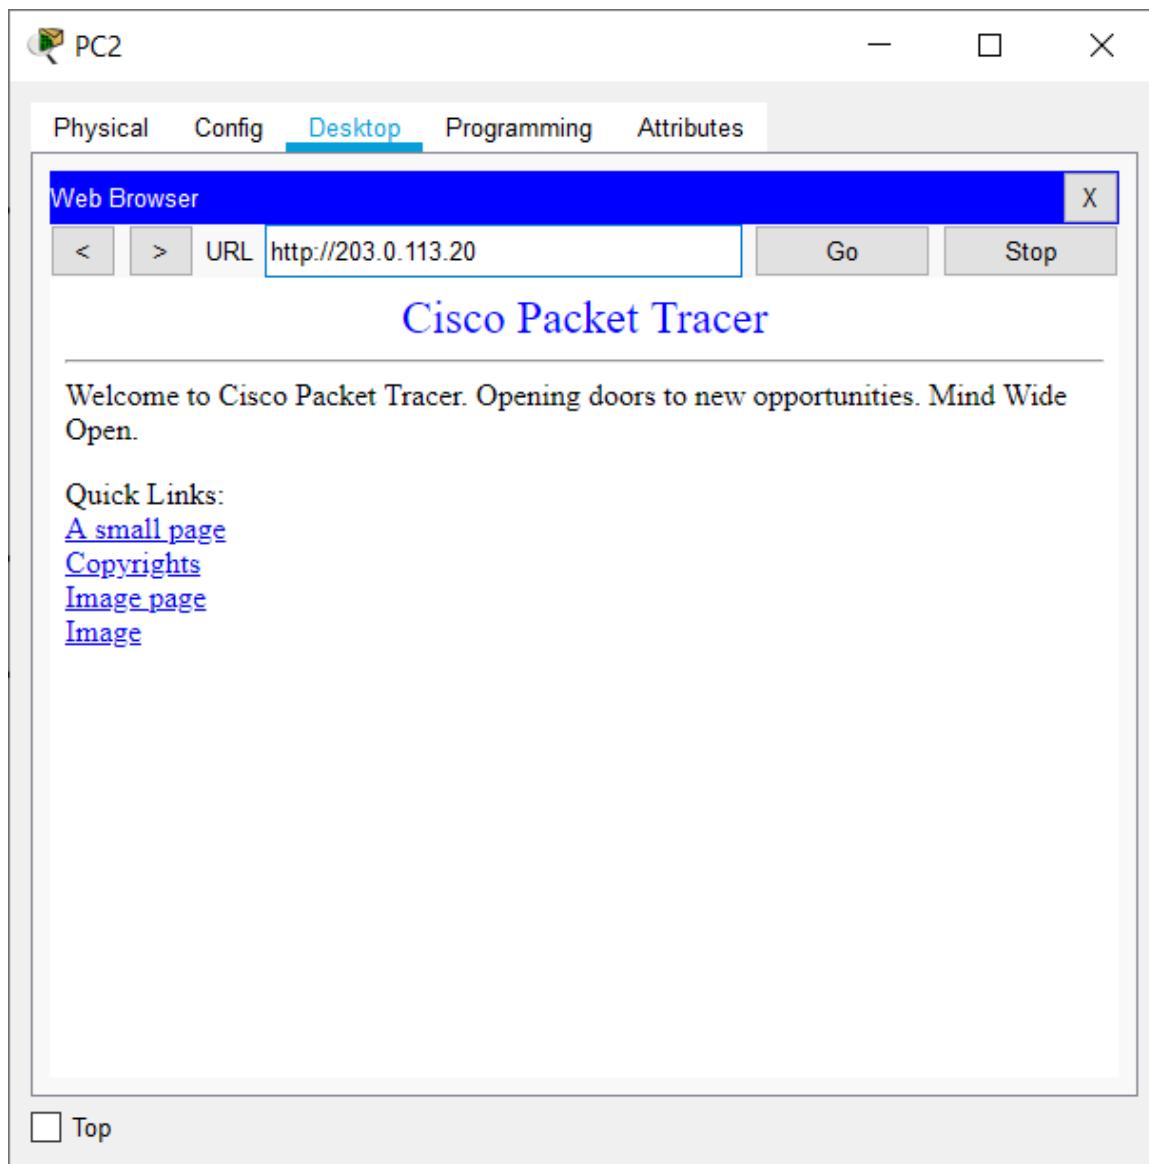
- 15) View the debug output on R1. Which global IP address was PC1 translated to?

PC1 is translated to the address on R1's F0/0 interface, 203.0.113.13.

```
*Mar 01, 03:40:43.4040: NAT*: s=10.0.2.10->203.0.113.13, d=203.0.113.20 [10]
*Mar 01, 03:40:43.4040: NAT*: s=203.0.113.20, d=203.0.113.13->10.0.2.10 [37]
*Mar 01, 03:40:43.4040: NAT*: s=10.0.2.10->203.0.113.13, d=203.0.113.20 [11]
*Mar 01, 03:40:43.4040: NAT*: s=10.0.2.10->203.0.113.13, d=203.0.113.20 [12]
```

16) On PC2, click on 'Desktop' then open 'Web Browser'. Open the public IP address of Ext-S1 at <http://203.0.113.20> in the browser. Which global IP address is PC2 translated to?

PC2 is also translated to 203.0.113.13.



```
*Mar 01, 03:47:51.4747: NAT: s=10.0.2.11->203.0.113.13, d=203.0.113.20 [1]
*Mar 01, 03:47:51.4747: NAT*: s=203.0.113.20, d=203.0.113.13->10.0.2.11 [40]
*Mar 01, 03:47:51.4747: NAT*: s=10.0.2.11->203.0.113.13, d=203.0.113.20 [2]
*Mar 01, 03:47:51.4747: NAT*: s=10.0.2.11->203.0.113.13, d=203.0.113.20 [3]
```

17) Verify the connections in the NAT translation table.

PC1 and PC2 are translated to different source ports.

```
R1#sh ip nat translation
Pro Inside global      Inside local        Outside local        Outside global
tcp 203.0.113.13:1025  10.0.2.10:1025    203.0.113.20:80    203.0.113.20:80
tcp 203.0.113.13:1024  10.0.2.11:1025    203.0.113.20:80    203.0.113.20:80
```

18) Show the NAT statistics on R1.

The numbers in your output may be different.

```
R1#sh ip nat statistics
Total translations: 2 (0 static, 2 dynamic, 2 extended)
Outside Interfaces: FastEthernet0/0
Inside Interfaces: FastEthernet1/0
Hits: 41 Misses: 32
Expired translations: 16
Dynamic mappings:
```

30-1 IPv6 Configuration - Lab Exercise

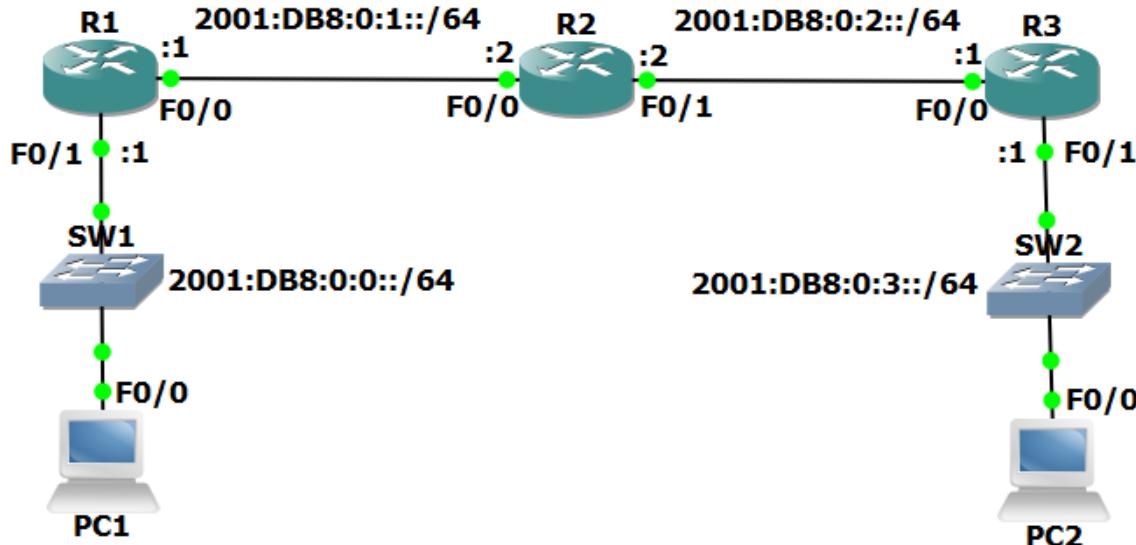
In this lab you will configure IPv6 addressing and static routing for a small campus network. The company already has IPv4 configured on their existing network but they will be introducing a new IPv6 application soon. Your task is to configure the network to be dual stack by adding support for IPv6 addresses. IPv6 routing will be configured in the next lab exercise.

Routers R1, R2 and R3 have been configured with IPv4 addresses and connectivity has been established with between all networks.

PC1 and PC2 are Cisco IOS routers mimicking end hosts. They have been configured with IPv4 addresses and an IPv4 default static route to their default gateways.

There is no IPv6 configuration on any of the devices.

Lab Topology



Load the Startup Configurations

[Download the '30-1 IPv6 Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Verify IPv4 Connectivity

- 1) Verify R1, R2, R3, PC1 and PC2 have been configured with IPv4 addresses as shown in the topology diagram.
- 2) View the routing tables on R1, R2 and R3 to verify connectivity has been established between all networks. What routing protocol is being used?
- 3) Verify PC1 and PC2 have been configured with the correct default gateway.
- 4) Ping PC2 from PC1 to verify end to end reachability.

IPv6 Addressing

- 5) Configure global unicast IPv6 addresses on R1, R2, and R3, according to the network topology diagram. Do not enable ipv6 unicast-routing.
- 6) Configure global unicast EUI-64 IPv6 addresses on the FastEthernet 0/0 interfaces on PC1 and PC2.
- 7) Will the routers have IPv6 link local addresses on the interfaces where you just configured global unicast addresses? What about the other interfaces? Verify this.
- 8) Note down the EUI-64 global unicast addresses on PC1 and PC2.
- 9) Configure link local addresses on R1, R2 and R3. For each router, use the address shown below on each of its interfaces.
R1: FE80::1/64
R2: FE80::2/64
R3: FE80::3/64
- 10) Verify the global unicast and link local addresses have been configured correctly on R1, R2 and R3.
- 11) Ping R1 and R3 on their link local addresses from R2.
- 12) View the IPv6 neighbors known by R2.

Static Routing

- 13) Verify which IPv6 dynamic routing protocols are running on R1, R2 and R3. Do not use the 'show run' command.
- 14) Use the 'show run | include ipv6 route' command on R1, R2 and R3 to verify if they have been configured with IPv6 static routes.
- 15) Do you expect to see any routes in the IPv6 routing tables? Why or why not? Verify this.
- 16) Do you expect PC1 to be able to ping PC2 on its IPv6 address? Why or why not? Verify this.
- 17) Configure PC1 to use R1 as its IPv6 default gateway.
- 18) Configure PC2 to use R3 as its IPv6 default gateway.
- 19) Verify PC1 and PC2 can ping their default gateways.
- 20) Add a static route on R2 for the 2001:db8::/64 network.
- 21) PC1 has reachability to its default gateway R1, and R2 has a route to the 2001:db8::/64 network. Do you expect PC1 be able to ping R2 on 2001:db8:0:1::2? Why or why not? Verify this.
- 22) Fix the problem to allow PC1 to ping R2 on 2001:db8:0:1::2. Enter the command which will fix the problem on R1, R2 and R3.
- 23) Verify PC1 can ping R2 on 2001:db8:0:1::2 now.
- 24) Will PC1 be able to ping PC2 by IPv6 address? Why or why not? Verify this.
- 25) Configure static routes to allow reachability between all IPv6 networks.
- 26) Verify the IPv6 routing tables on R1, R2 and R3.
- 27) Verify PC1 can now ping PC2 by its IPv6 address.

30-1 IPv6 Configuration - Answer Key

In this lab you will configure IPv6 addressing and static routing for a small campus network. The company already has IPv4 configured on their existing network but they will be introducing a new IPv6 application soon. Your task is to configure the network to be dual stack by adding support for IPv6 addresses. IPv6 routing will be configured in the next lab exercise.

Routers R1, R2 and R3 have been configured with IPv4 addresses and connectivity has been established with between all networks.

PC1 and PC2 are Cisco IOS routers mimicking end hosts. They have been configured with IPv4 addresses and an IPv4 default static route to their default gateways.

There is no IPv6 configuration on any of the devices.

Verify IPv4 Connectivity

- 1) Verify R1, R2, R3, PC1 and PC2 have been configured with IPv4 addresses as shown in the topology diagram.

```
R1#sh ip interface brief
Interface          IP-Address      OK? Method Status
Protocol
FastEthernet0/0   10.10.1.1      YES  NVRAM  up
FastEthernet0/1   10.10.0.1      YES  NVRAM  up
FastEthernet1/0   unassigned     YES  NVRAM  administratively down down
FastEthernet1/1   unassigned     YES  NVRAM  administratively down down
Vlan1            unassigned     YES  NVRAM  administratively down down
```

```
R2#sh ip interface brief
Interface          IP-Address      OK? Method Status
Protocol
FastEthernet0/0   10.10.1.2      YES  NVRAM  up
FastEthernet0/1   10.10.2.2      YES  NVRAM  up
FastEthernet1/0   unassigned     YES  NVRAM  administratively down down
FastEthernet1/1   unassigned     YES  NVRAM  administratively down down
Vlan1            unassigned     YES  NVRAM  administratively down down
```

```
R3#sh ip interface brief
Interface          IP-Address      OK? Method Status
Protocol
FastEthernet0/1   10.10.2.1      YES  NVRAM  up
FastEthernet1/0   10.10.3.1      YES  NVRAM  up
FastEthernet1/1   unassigned     YES  NVRAM  administratively down down
FastEthernet0/0   unassigned     YES  NVRAM  administratively down down
Vlan1            unassigned     YES  NVRAM  administratively down down
```

```
PC1#sh ip interface brief
```

Interface Protocol	IP-Address	OK?	Method	Status
FastEthernet0/0	10.10.0.10	YES	NVRAM	up up
FastEthernet1/0	unassigned	YES	NVRAM	administratively down down
FastEthernet2/0	unassigned	YES	NVRAM	administratively down down
FastEthernet3/0	unassigned	YES	NVRAM	administratively down down
Vlan1	unassigned	YES	NVRAM	administratively down down

PC2#sh ip interface brief	Interface Protocol	IP-Address	OK?	Method	Status
FastEthernet0/0	10.10.3.10	YES	NVRAM	up up	
FastEthernet1/0	unassigned	YES	NVRAM	administratively down down	
FastEthernet2/0	unassigned	YES	NVRAM	administratively down down	
FastEthernet3/0	unassigned	YES	NVRAM	administratively down down	
Vlan1	unassigned	YES	NVRAM	administratively down down	

2) View the routing tables on R1, R2 and R3 to verify connectivity has been established between all networks. What routing protocol is being used?

EIGRP is the routing protocol.

```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
C 10.10.0.0/24 is directly connected, FastEthernet0/1
L 10.10.0.1/32 is directly connected, FastEthernet0/1
C 10.10.1.0/24 is directly connected, FastEthernet0/0
L 10.10.1.1/32 is directly connected, FastEthernet0/0
D 10.10.2.0/24 [90/30720] via 10.10.1.2, 00:00:32, FastEthernet0/0
D 10.10.3.0/24 [90/33280] via 10.10.1.2, 00:00:21, FastEthernet0/0
```

```
R2#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
D 10.10.0.0/24 [90/30720] via 10.10.1.1, 00:07:18, FastEthernet0/0
C 10.10.1.0/24 is directly connected, FastEthernet0/0
L 10.10.1.2/32 is directly connected, FastEthernet0/0
C 10.10.2.0/24 is directly connected, FastEthernet0/1
L 10.10.2.2/32 is directly connected, FastEthernet0/1
D 10.10.3.0/24 [90/30720] via 10.10.2.1, 00:01:23, FastEthernet0/1
```

```
R3#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
D 10.10.0.0/24 [90/33280] via 10.10.2.2, 00:02:31, FastEthernet0/0
D 10.10.1.0/24 [90/30720] via 10.10.2.2, 00:02:31, FastEthernet0/0
C 10.10.2.0/24 is directly connected, FastEthernet0/0
L 10.10.2.1/32 is directly connected, FastEthernet0/0
C 10.10.3.0/24 is directly connected, FastEthernet0/1
L 10.10.3.1/32 is directly connected, FastEthernet0/1
```

- 3) Verify PC1 and PC2 have been configured with the correct default gateway.

PC1 is using R1 as its default gateway.

```
PC1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is 10.10.0.1 to network 0.0.0.0

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.10.0.0/24 is directly connected, FastEthernet0/0
L 10.10.0.10/32 is directly connected, FastEthernet0/0
S* 0.0.0.0/0 [1/0] via 10.10.0.1
```

PC2 is using R3 as its default gateway.

```
PC2#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override
```

Gateway of last resort is 10.10.3.1 to network 0.0.0.0

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.10.3.0/24 is directly connected, FastEthernet0/0
L 10.10.3.10/32 is directly connected, FastEthernet0/0
S* 0.0.0.0/0 [1/0] via 10.10.3.1
```

- 4) Ping PC2 from PC1 to verify end to end reachability.

```
PC1#ping 10.10.3.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.3.10, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
80/80/84 ms
```

IPv6 Addressing

5) Configure global unicast IPv6 addresses on R1, R2, and R3, according to the network topology diagram. Do not enable ipv6 unicast-routing.

```
R1(config)#int f0/1
R1(config-if)#ipv6 address 2001:db8::1/64
R1(config-if)#no shutdown
R1(config-if)#exit
R1(config)#int f0/0
R1(config-if)#ipv6 address 2001:db8:0:1::1/64
R1(config-if)#no shutdown

R2(config)#int f0/0
R2(config-if)#ipv6 address 2001:db8:0:1::2/64
R2(config-if)#no shut
R2(config-if)#exit
R2(config)#int f0/1
R2(config-if)#ipv6 address 2001:db8:0:2::2/64
R2(config-if)#no shut

R3(config)#int f0/0
R3(config-if)#ipv6 address 2001:db8:0:2::1/64
R3(config-if)#no shut
R3(config-if)#exit
R3(config)#int f0/1
R3(config-if)#ipv6 address 2001:db8:0:3::1/64
R3(config-if)#no shut
```

6) Configure global unicast EUI-64 IPv6 addresses on the FastEthernet 0/0 interfaces on PC1 and PC2.

```
PC1(config)#int f0/0
PC1(config-if)#ipv6 address 2001:db8::/64 eui-64
PC1(config-if)#no shut

PC2(config)#int f0/0
PC2(config-if)#ipv6 address 2001:db8:0:3::/64 eui-64
PC2(config-if)#no shut
```

7) Will the routers have IPv6 link local addresses on the interfaces where you just configured global unicast addresses? What about the other interfaces? Verify this.

EUI-64 link local addresses will be automatically configured when IPv6 is enabled on the interfaces (by applying the global unicast addresses). The other interfaces will not have link local addresses.

```
R1#show ipv6 interface brief
FastEthernet0/0          [up/up]
FE80::20D:BDFF:FE2D:27D4
  2001:DB8:0:1::1
FastEthernet0/1          [up/up]
FE80::2D0:97FF:FE64:3118
  2001:DB8::1
FastEthernet1/0          [administratively down/down]
  unassigned
FastEthernet1/1          [administratively down/down]
  unassigned
Vlan1                   [administratively down/down]
  unassigned
```

8) Note down the EUI-64 global unicast addresses on PC1 and PC2.

Note that these addresses are generated based on the interface MAC address and may be different in your lab.

```
PC1#show ipv6 interface brief
FastEthernet0/0          [up/up]
FE80::200:CFF:FE47:14C0
  2001:DB8::200:CFF:FE47:14C0
```

```
PC2#show ipv6 interface brief
FastEthernet0/0          [up/up]
FE80::201:C7FF:FE50:8E8A
  2001:DB8:0:3:201:C7FF:FE50:8E8A
```

9) Configure link local addresses on R1, R2 and R3. For each router, use the address shown below on each of its interfaces.

R1: FE80::1/64

R2: FE80::2/64

R3: FE80::3/64

```
R1(config)#int f0/0
R1(config-if)#ipv6 address fe80::1 link-local
R1(config-if)#exit
R1(config)#int f0/1
R1(config-if)#ipv6 address fe80::1 link-local

R2(config)#int f0/0
R2(config-if)#ipv6 address fe80::2 link-local
R2(config-if)#exit
R2(config)#int f0/1
R2(config-if)#ipv6 address fe80::2 link-local

R3(config)#int f0/0
R3(config-if)#ipv6 address fe80::3 link-local
R3(config-if)#exit
R3(config)#int f0/1
R3(config-if)#ipv6 address fe80::3 link-local
```

10) Verify the global unicast and link local addresses have been configured correctly on R1, R2 and R3.

```
R1#sh ipv6 int brief
FastEthernet0/0          [up/up]
  FE80::1
  2001:DB8:0:1::1
FastEthernet0/1          [up/up]
  FE80::1
  2001:DB8::1
FastEthernet1/0          [administratively down/down]
  unassigned
FastEthernet1/1          [administratively down/down]
  Unassigned
Vlan1                   [administratively down/down]
  unassigned
```

```
R2#sh ipv6 int brief
FastEthernet0/0          [up/up]
  FE80::2
  2001:DB8:0:1::2
FastEthernet0/1          [up/up]
  FE80::2
  2001:DB8:0:2::2
FastEthernet1/0          [administratively down/down]
  unassigned
FastEthernet1/1          [administratively down/down]
  Unassigned
Vlan1                   [administratively down/down]
  unassigned
```

```
R3#sh ipv6 int brief
FastEthernet0/0          [up/up]
  FE80::3
  2001:DB8:0:2::1
FastEthernet0/1          [up/up]
  FE80::3
  2001:DB8:0:3::1
FastEthernet1/0          [administratively down/down]
  unassigned
FastEthernet1/1          [administratively down/down]
  Unassigned
Vlan1                   [administratively down/down]
  unassigned
```

11) Ping R1 and R3 on their link local addresses from R2.

```
R2#ping fe80::1
Output Interface: FastEthernet0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::1, timeout is 2
seconds:
Packet sent with a source address of
FE80::2%FastEthernet0/0
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
16/22/40 ms
```

```

R2#ping fe80::3
Output Interface: FastEthernet0/1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::3, timeout is 2
seconds:
Packet sent with a source address of
FE80::2%FastEthernet1/0
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
20/26/52 ms

```

12) View the IPv6 neighbors known by R2.

Note that your output may be different depending on what IPv6 traffic has been seen recently by R2.

```

R2#show ipv6 neighbors
IPv6 Address                               Age Link-layer Addr State Interface
2001:DB8:0:0:1::1                           0 000D.BD2D.27D4  REACH Fa0/0
2001:DB8:0:0:2::1                           0 0030.F2BA.30E7  REACH Fa0/1
FE80::1                                     0 000D.BD2D.27D4  REACH Fa0/0
FE80::3                                     0 0030.F2BA.30E7  REACH Fa0/1

```

Static Routing

13) Verify which IPv6 dynamic routing protocols are running on R1, R2 and R3. Do not use the 'show run' command.

The routers are not running any IPv6 dynamic routing protocol.

```

R1#sh ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"

```

```

R2#sh ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"

```

```

R3#sh ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"

```

14) Use the 'show run | include ipv6 route' command on R1, R2 and R3 to verify if they have been configured with IPv6 static routes.

The routers have not been configured with any IPv6 static routes.

```
R1#show run | include ipv6 route
R1#
```

```
R2#show run | include ipv6 route
R2#
```

```
R3#show run | include ipv6 route
R3#
```

15) Do you expect to see any routes in the IPv6 routing tables? Why or why not? Verify this.

The routers have been configured with IPv6 addresses so connected and local routes will appear in the routing table.

```
R1#show ipv6 route
IPv6 Routing Table - default - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
      B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
      I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
      EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
      NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
      OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
C  2001:DB8::/64 [0/0]
    via FastEthernet0/1, directly connected
L  2001:DB8::1/128 [0/0]
    via FastEthernet0/1, receive
C  2001:DB8:0:1::/64 [0/0]
    via FastEthernet0/0, directly connected
L  2001:DB8:0:1::1/128 [0/0]
    via FastEthernet0/0, receive
L  FF00::/8 [0/0]
    via Null0, receive
```

```
R2#sh ipv6 route
IPv6 Routing Table - default - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
C  2001:DB8:0:1::/64 [0/0]
    via FastEthernet0/0, directly connected
L  2001:DB8:0:1::2/128 [0/0]
    via FastEthernet0/0, receive
C  2001:DB8:0:2::/64 [0/0]
    via FastEthernet0/1, directly connected
L  2001:DB8:0:2::2/128 [0/0]
    via FastEthernet0/1, receive
L  FF00::/8 [0/0]
    via Null0, receive
```

```
R3#sh ipv6 route
IPv6 Routing Table - default - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
C  2001:DB8:0:2::/64 [0/0]
    via FastEthernet0/0, directly connected
L  2001:DB8:0:2::1/128 [0/0]
    via FastEthernet0/0, receive
C  2001:DB8:0:3::/64 [0/0]
    via FastEthernet0/1, directly connected
L  2001:DB8:0:3::1/128 [0/0]
    via FastEthernet0/1, receive
L  FF00::/8 [0/0]
    via Null0, receive
```

16) Do you expect PC1 to be able to ping PC2 on its IPv6 address? Why or why not? Verify this.

No IPv6 routing has been configured so there is no connectivity between different IPv6 subnets. (Note that PC2's EUI-64 IPv6 address may be different in your lab. Use 'show ipv6 interface brief' to check its address.)

```
PC1#ping 2001:DB8:0:3:201:C7FF:FE50:8E8A
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to  
2001:DB8:0:3:201:C7FF:FE50:8E8A, timeout is 2 seconds:  
.....  
Success rate is 0 percent (0/5)
```

17) Configure PC1 to use R1 as its IPv6 default gateway.

```
PC1(config)#ipv6 route ::/0 2001:db8::1
```

18) Configure PC2 to use R3 as its IPv6 default gateway.

```
PC2(config)#ipv6 route ::/0 2001:db8:0:3::1
```

19) Verify PC1 and PC2 can ping their default gateways.

```
PC1#ping 2001:db8::1  
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 2001:DB8::1, timeout is 2  
seconds:  
!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max =  
20/31/44 ms
```

```
PC2#ping 2001:DB8:0:3::1  
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::1, timeout  
is 2 seconds:  
!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max =  
16/26/56 ms
```

20) Add a static route on R2 for the 2001:db8::/64 network.

```
R2(config)#ipv6 route 2001:db8::/64 2001:db8:0:1::1
```

21) PC1 has reachability to its default gateway R1, and R2 has a route to the 2001:db8::/64 network. Do you expect PC1 be able to ping R2 on 2001:db8:0:1::2? Why or why not? Verify this.

PC1 cannot ping R2 on 2001:db8:0:1::2 because ipv6 unicast-routing has not been enabled on the routers.

```
PC1#ping 2001:db8:0:1::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:1::2, timeout
is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

22) Fix the problem to allow PC1 to ping R2 on 2001:db8:0:1::2. Enter the command which will fix the problem on R1, R2 and R3.

```
R1(config)#ipv6 unicast-routing
R2(config)#ipv6 unicast-routing
R3(config)#ipv6 unicast-routing
```

23) Verify PC1 can ping R2 on 2001:db8:0:1::2 now.

```
PC1#ping 2001:db8:0:1::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:1::2, timeout
is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
60/60/64 ms
```

24) Will PC1 be able to ping PC2 by IPv6 address? Why or why not? Verify this.

PC1 cannot ping PC2 because routing has not been configured between the 2001:db8::/64 and 2001:db8:0:3::/64 networks.

```
PC1#ping 2001:DB8:0:3:201:C7FF:FE50:8E8A
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to
2001:DB8:0:3:201:C7FF:FE50:8E8A, timeout is 2 seconds:
UUUUU
Success rate is 0 percent (0/5)
```

25) Configure static routes to allow reachability between all IPv6 networks.

```
R1(config)#ipv6 route 2001:db8:0:2::/64 2001:db8:0:1::2
R1(config)#ipv6 route 2001:db8:0:3::/64 2001:db8:0:1::2
```

```
R2(config)#ipv6 route 2001:db8::/64 2001:db8:0:1::1
R2(config)#ipv6 route 2001:db8:0:3::/64 2001:db8:0:2::1
```

```
R3(config)#ipv6 route 2001:db8::/64 2001:db8:0:2::2
R3(config)#ipv6 route 2001:db8:0:1::/64 2001:db8:0:2::2
```

26) Verify the IPv6 routing tables on R1, R2 and R3.

```
R1#sh ipv6 route
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
        I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
        EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
        NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
        OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
C  2001:DB8::/64 [0/0]
    via FastEthernet0/1, directly connected
L  2001:DB8::1/128 [0/0]
    via FastEthernet0/1, receive
C  2001:DB8:0:1::/64 [0/0]
    via FastEthernet0/0, directly connected
L  2001:DB8:0:1::1/128 [0/0]
    via FastEthernet0/0, receive
S  2001:DB8:0:2::/64 [1/0]
    via 2001:DB8:0:1::2
S  2001:DB8:0:3::/64 [1/0]
    via 2001:DB8:0:1::2
L  FF00::/8 [0/0]
    via Null0, receive
```

```
R2#sh ipv6 route
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
S  2001:DB8::/64 [1/0]
    via 2001:DB8:0:1::1
C  2001:DB8:0:1::/64 [0/0]
    via FastEthernet0/0, directly connected
L  2001:DB8:0:1::2/128 [0/0]
    via FastEthernet0/0, receive
C  2001:DB8:0:2::/64 [0/0]
    via FastEthernet0/1, directly connected
L  2001:DB8:0:2::2/128 [0/0]
    via FastEthernet0/1, receive
S  2001:DB8:0:3::/64 [1/0]
    via 2001:DB8:0:2::1
L  FF00::/8 [0/0]
    via Null0, receive
```

```
R3#sh ipv6 route
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
S  2001:DB8::/64 [1/0]
    via 2001:DB8:0:2::2
S  2001:DB8:0:1::/64 [1/0]
    via 2001:DB8:0:2::2
C  2001:DB8:0:2::/64 [0/0]
    via FastEthernet0/0, directly connected
L  2001:DB8:0:2::1/128 [0/0]
    via FastEthernet10/0, receive
C  2001:DB8:0:3::/64 [0/0]
    via FastEthernet0/1, directly connected
L  2001:DB8:0:3::1/128 [0/0]
    via FastEthernet0/1, receive
L  FF00::/8 [0/0]
    via Null0, receive
```

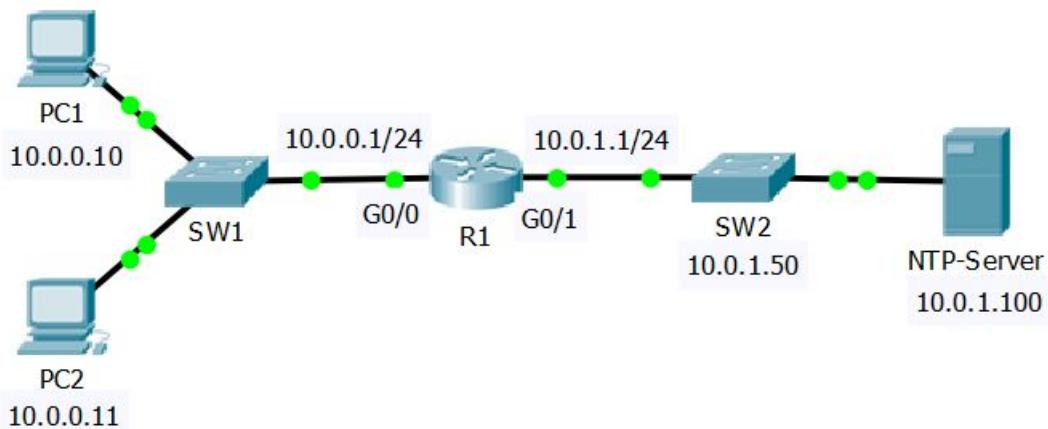
27) Verify PC1 can now ping PC2 by its IPv6 address.

```
PC1#2001:DB8:0:3:201:C7FF:FE50:8E8A
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to
2001:DB8:0:3:201:C7FF:FE50:8E8A, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
80/80/84 ms
```

33-1 Cisco Device Security Configuration - Lab Exercise

In this lab you will secure administrative access to the Cisco router in a small campus network.

Lab Topology



Load the Startup Configurations

[Download the '33-1 Cisco Device Security Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

Secure Privileged Exec Mode

- 1) Set the enable password **Flackbox2** on R1 to secure access to Privileged Exec (Enable) mode.
- 2) Exit to User Exec mode.
- 3) Enter Privileged Exec mode.
- 4) Set the enable secret **Flackbox1**.
- 5) Exit to User Exec mode.
- 6) Do you expect to be able to enter Privileged Exec mode using the password **Flackbox2**? Why or why not? Verify this.
- 7) Show the running configuration on R1. Can you read the enable password and secret in plain text?
- 8) Ensure that passwords will not show in plain text in the output of 'show' commands.
- 9) Verify the enable password is now encrypted when you show the running configuration.

Secure Remote Telnet and SSH Access

- 10) Enable synchronous logging on R1 and ensure administrators are logged out after 15 minutes of activity on the console and virtual terminal lines 0-15.
- 11) Allow the administrator workstation at 10.0.0.10 to Telnet into R1 using the password **Flackbox3**. Ensure no other host has Telnet access to the router.
- 12) Ensure that users attempting to Telnet into the router see the message "Authorised users only"
- 13) Verify you can Telnet into R1 from PC1 and enter Privileged Exec mode. Close the Telnet session when done.
- 14) Verify Telnet access fails from PC2.

- 15) Configure R1 so that administrators will be prompted to enter a username and password when they attempt to Telnet into the router. Use username **admin** and password **Flackbox4**.
- 16) Verify you are prompted for a username and password when you attempt to Telnet to the router.
- 17) Allow the administrator workstation at 10.0.0.10 to SSH into R1. Use the domain name **flackbox.com** and a 768 bit key.
- 18) Verify you can SSH into R1 from PC1. Close the session when done.
- 19) Do you expect to be able to SSH to R1 from PC2? Why or why not? Verify this.
- 20) You can currently access R1 using either Telnet or SSH. Telnet is an insecure protocol as all communication is sent in plain text. Configure R1 so that only SSHv2 remote access is allowed.
- 21) Verify you cannot Telnet into R1 from PC1 but can SSH. Exit when done.
- 22) What username and password do you need to use to login when you connect directly to R1 with a console cable?
- 23) Configure R1 to require no username but a password of **Flackbox5** to login over the console connection.
- 24) Verify you can access R1 over the console connection and enter Privileged Exec mode.

NTP Network Time Protocol

- 25) Configure R1 to synchronise its time with the NTP server at 10.0.1.100. Set the timezone as Pacific Standard Time which is 8 hours before UTC.
- 26) Check the current time on the router and verify it is synchronised with the NTP server.

Switch Management

27) Configure SW2 with IP address 10.0.1.50 for management on VLAN 1.
Ensure the switch has connectivity to other IP subnets.

(Note that it is best practice to NOT use VLAN 1 for any production traffic in a real world network and we would normally have a separate dedicated IP subnet for management traffic. We are using VLAN 1 in our lab environment to simplify the topology)

33-1 Cisco Device Security Configuration - Answer Key

In this lab you will secure administrative access to the Cisco router in a small campus network.

Secure Privileged Exec Mode

- 1) Set the enable password **Flackbox2** on R1 to secure access to Privileged Exec (Enable) mode.

```
R1(config)#enable password Flackbox2
```

- 2) Exit to User Exec mode.

```
R1#exit
```

- 3) Enter Privileged Exec mode.

```
R1>enable  
Password:Flackbox2  
R1#
```

- 4) Set the enable secret **Flackbox1**.

```
R1(config)#enable secret Flackbox1
```

- 5) Exit to User Exec mode.

```
R1#exit
```

- 6) Do you expect to be able to enter Privileged Exec mode using the password **Flackbox2**? Why or why not? Verify this.

You cannot enter Privileged Exec mode using the enable password because it has been superseded by the enable secret.

```
R1>enable  
Password:Flackbox2  
Password:Flackbox1  
R1#
```

7) Show the running configuration on R1. Can you read the enable password and secret in plain text?

The enable password is shown in plain text but the enable secret is encrypted.

```
R1#show run
Building configuration...

Current configuration : 762 bytes
!
version 15.1
no service timestamps log datetime msec
no service timestamps debug datetime msec
no service password-encryption
!
hostname R1
!
enable secret 5 $1$mERr$J2XZHMOgpVVXdLjC9lYtE1
enable password Flackbox2
```

8) Ensure that passwords will not show in plain text in the output of 'show' commands.

```
R1(config)#service password-encryption
```

9) Verify the enable password is now encrypted when you show the running configuration.

```
R1#show running-config
Building configuration...

Current configuration : 772 bytes
!
version 15.1
no service timestamps log datetime msec
no service timestamps debug datetime msec
service password-encryption
!
hostname R1
!
enable secret 5 $1$mERr$J2XZHMOgpVVXdLjC9lYtE1
enable password 7 0807404F0A1207180A59
```

Secure Remote Telnet and SSH Access

- 10) Enable synchronous logging on R1 and ensure administrators are logged out after 15 minutes of activity on the console and virtual terminal lines 0-15.

```
R1(config)#line console 0
R1(config-line)#logging synchronous
R1(config-line)#exec-timeout 15

R1(config)#line vty 0 15
R1(config-line)#logging synchronous
R1(config-line)#exec-timeout 15
```

- 11) Allow the administrator workstation at 10.0.0.10 to Telnet into R1 using the password **Flackbox3**. Ensure no other host has Telnet access to the router.

```
R1(config)#access-list 1 permit host 10.0.0.10

R1(config)#line vty 0 15
R1(config-line)#login
R1(config-line)#password Flackbox3
R1(config-line)#access-class 1 in
```

- 12) Ensure that users attempting to Telnet into the router see the message "Authorised users only"

```
R1(config)#banner login "
Enter TEXT message. End with the character ' '.
Authorised users only"
```

13) Verify you can Telnet into R1 from PC1 and enter Privileged Exec mode.
Close the Telnet session when done.

```
C:\> telnet 10.0.0.1
Trying 10.0.0.1 ...Open
```

Authorised users only

User Access Verification

```
Password: Flackbox3
R1>enable
Password: Flackbox1
R1#exit
```

[Connection to 10.0.0.1 closed by foreign host]

14) Verify Telnet access fails from PC2.

```
C:\> telnet 10.0.0.1
Trying 10.0.0.1 ...
% Connection refused by remote host
```

15) Configure R1 so that administrators will be prompted to enter a username and password when they attempt to Telnet into the router. Use username **admin** and password **Flackbox4**.

```
R1(config)#username admin secret Flackbox4
R1(config)#line vty 0 15
R1(config-line)#login local
```

16) Verify you are prompted for a username and password when you attempt to Telnet to the router.

```
C:\> telnet 10.0.0.1
Trying 10.0.0.1 ...Open
```

Authorised users only

User Access Verification

Username: admin

```
Password: Flackbox4
R1>exit
```

[Connection to 10.0.0.1 closed by foreign host]

17) Allow the administrator workstation at 10.0.0.10 to SSH into R1. Use the domain name **flackbox.com** and a 768 bit key.

```
R1(config)#ip domain-name flackbox.com
R1(config)#crypto key generate rsa
The name for the keys will be: R1.flackbox.com
Choose the size of the key modulus in the range of 360 to
2048 for your
General Purpose Keys. Choosing a key modulus greater than
512 may take
a few minutes.
```

```
How many bits in the modulus [512]: 768
% Generating 768 bit RSA keys, keys will be non-
exportable...[OK]
```

18) Verify you can SSH into R1 from PC1. Close the session when done.

```
C:\> ssh -l admin 10.0.0.1
Open
Password: Flackbox4
R1>exit
```

```
[Connection to 10.0.0.1 closed by foreign host]
```

19) Do you expect to be able to SSH to R1 from PC2? Why or why not?
Verify this.

You will not be able to SSH to R1 from PC2. Telnet and SSH access are both controlled by the 'line vty' configuration which has an access list applied only allowing access from PC1.

```
C:\> ssh -l admin 10.0.0.1
Trying 10.0.0.1 ...
% Connection refused by remote host
```

20) You can currently access R1 using either Telnet or SSH. Telnet is an insecure protocol as all communication is sent in plain text. Configure R1 so that only SSHv2 remote access is allowed.

```
R1(config)#line vty 0 15
R1(config-line)#transport input ssh
R1(config-line)#exit
R1(config)#ip ssh version 2
```

21) Verify you cannot Telnet into R1 from PC1 but can SSH. Exit when done.

```
C:\> telnet 10.0.0.1
Trying 10.0.0.1 ...Open

[Connection to 10.0.0.1 closed by remote host]

C:\> ssh -l admin 10.0.0.1
Open
Password: Flackbox4
R1>exit

[Connection to 10.0.0.1 closed by foreign host]
```

22) What username and password do you need to use to login when you connect directly to R1 with a console cable?

No username and password are currently required to login to the console. The virtual terminal lines which control Telnet and SSH access have been secured but console access has not.

23) Configure R1 to require no username but a password of **Flackbox5** to login over the console connection.

```
R1(config)#line console 0
R1(config-line)#login
% Login disabled on line 0, until 'password' is set
R1(config-line)#password Flackbox5
```

24) Verify you can access R1 over the console connection and enter Privileged Exec mode.

```
R1(config-line)#end
R1#logout

R1 con0 is now available
Press RETURN to get started.
Authorised users only
User Access Verification
Password: Flackbox5
R1>enable
Password: Flackbox1
R1#
```

NTP Network Time Protocol

25) Configure R1 to synchronise its time with the NTP server at 10.0.1.100.
Set the timezone as Pacific Standard Time which is 8 hours before UTC.

```
R1(config)#clock timezone PST -8
R1(config)#ntp server 10.0.1.100
```

26) Check the current time on the router and verify it is synchronised with the NTP server.

```
R1#show clock
16:19:36.51 PST Mon Oct 2 2017

R1#show ntp status
Clock is synchronized, stratum 2, reference is 10.0.1.100
nominal freq is 250.0000 Hz, actual freq is 249.9990 Hz,
precision is 2**19
reference time is DD53255C.0000039C (00:16:28.924 UTC Tue
Oct 3 2017)
clock offset is 0.00 msec, root delay is 0.00 msec
root dispersion is 0.02 msec, peer dispersion is 0.02 msec.
```

Switch Management

27) Configure SW2 with IP address 10.0.1.50 for management on VLAN 1.
Ensure the switch has connectivity to other IP subnets.

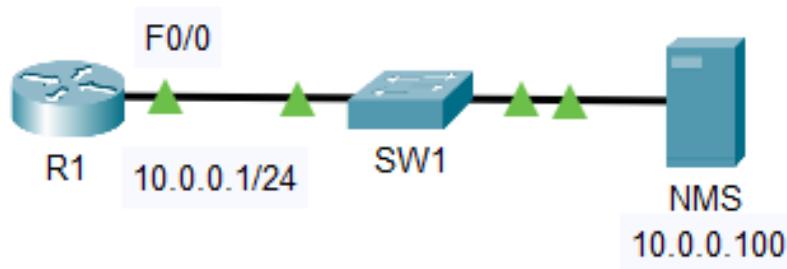
(Note that it is best practice to NOT use VLAN 1 for any production traffic in a real world network and we would normally have a separate dedicated IP subnet for management traffic. We are using VLAN 1 in our lab environment to simplify the topology).

```
SW2(config)#int vlan 1
SW2(config-if)#ip address 10.0.1.50 255.255.255.0
SW2(config-if)#no shutdown
SW2(config-if)#exit
SW2(config)#ip default-gateway 10.0.1.1
```

34 Network Device Management – Lab Exercise

You will configure Syslog and SNMP (Simple Network Management Protocol) logging in this lab exercise. The NMS server is acting as an external Syslog destination.

Lab Topology



Load the Startup Configurations

[Download the '34 Network Device Management.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

SNMP and Syslog

- 1) Configure SNMP communities on R1. Use **Flackbox1** as the Read Only community string, and **Flackbox2** as the Read Write community string.
- 2) Configure R1 so it will show events from all severity levels to the external Syslog server at 10.0.0.100.
- 3) Verify you have set the correct severity level.
- 4) Enable then disable the FastEthernet 0/1 interface on R1.
- 5) On the Syslog server at 10.0.0.100, click 'Services' then 'SYSLOG' and check you can see events for the interface coming up then back down.

34 Network Device Management – Answer Key

You will configure Syslog and SNMP (Simple Network Management Protocol) logging in this lab exercise. The NMS server is acting as an external Syslog destination.

SNMP and Syslog

- 1) Configure SNMP communities on R1. Use **Flackbox1** as the Read Only community string, and **Flackbox2** as the Read Write community string.

```
R1(config)#snmp-server community Flackbox1 ro
R1(config)#snmp-server community Flackbox2 rw
```

- 2) Configure R1 so it will show events from all severity levels to the external Syslog server at 10.0.0.100.

```
R1(config)#logging 10.0.0.100
R1(config)#logging trap debugging
```

- 3) Verify you have set the correct severity level.

```
R1#show logging
Syslog logging: enabled (0 messages dropped, 0 messages rate-limited,
0 flushes, 0 overruns, xml disabled, filtering disabled)
```

No Active Message Discriminator.

No Inactive Message Discriminator.

```
Console logging: level debugging, 3 messages logged, xml disabled,
filtering disabled
```

```
Monitor logging: level debugging, 3 messages logged, xml disabled,
filtering disabled
```

```
Buffer logging: disabled, xml disabled,
filtering disabled
```

```
Logging Exception size (4096 bytes)
```

```
Count and timestamp logging messages: disabled
```

```
Persistent logging: disabled
```

No active filter modules.

```
ESM: 0 messages dropped
```

```
Trap logging: level debugging, 3 message lines logged
```

```
Logging to 10.0.0.100 (udp port 514, audit disabled,
authentication disabled, encryption disabled, link up),
2 message lines logged,
0 message lines rate-limited,
```

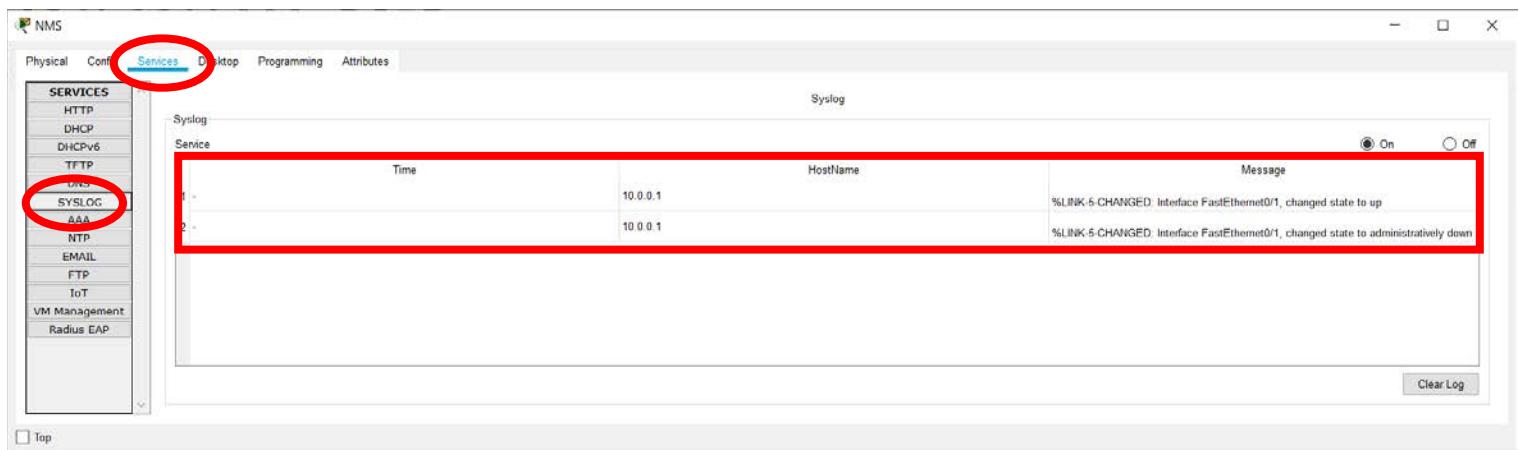
```
0 message lines dropped-by-MD,  
xml disabled, sequence number disabled  
filtering disabled
```

4) Enable then disable the FastEthernet 0/1 interface on R1.

```
R1(config)#int f0/1  
R1(config-if)#no shutdown  
%LINK-5-CHANGED: Interface FastEthernet0/1, changed state  
to up
```

```
R1(config-if)#shutdown  
%LINK-5-CHANGED: Interface FastEthernet0/1, changed state  
to administratively down
```

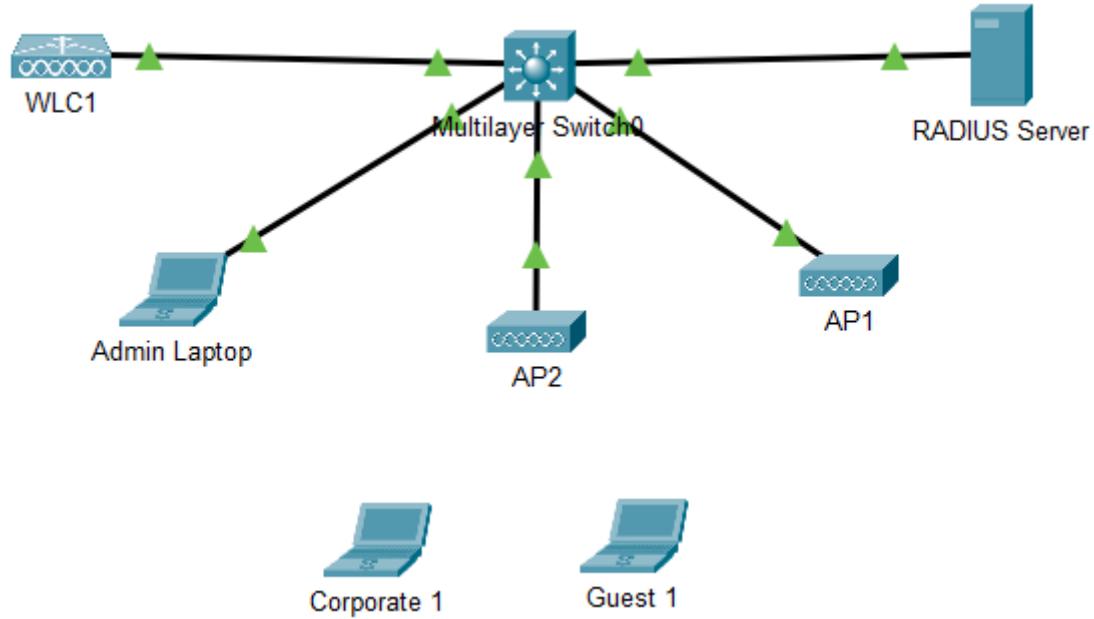
5) On the Syslog server at 10.0.0.100, click 'Services' then 'SYSLOG' and check you can see events for the interface coming up then back down.



37 Wireless Fundamentals Configuration – Lab Exercise

In this lab you will configure Corporate and Guest WLANs in a company campus.

Lab Topology



Load the Startup Configurations

[Download the '37 Wireless Fundamentals Configuration.zip' file here.](#) Extract the project .pkt file then open it in Packet Tracer. Do not try to open the project from directly inside the zip file.

VLANs and IP subnets have already been set up for the company servers and IT administrators to connect via wired connections:

VLAN Name	VLAN Number	IP Subnet	Gateway (on switch)
Server	11	192.168.11.0/24	192.168.11.1
Admin	21	192.168.21.0/24	192.168.21.1

The IT administrators are restricted to wired connections for security reasons, an 'Admin' WLAN will not be created.

A new Wireless LAN Controller has been added to the network. Your colleague has already performed the initial setup at the command line to give the device IP address 192.168.10.11/24

Two Lightweight Wireless Access Points have just been unboxed and cabled to the Multilayer Switch.

Your job is to configure the new Corporate and Guest WLANs.

Note: Packet Tracer does not support a trunk port to the WLC so you will configure the VLAN information on 'dummy' ports on the switch. The devices are really connected to interfaces GigabitEthernet1/0/11 – 15. Do not change this.

Switch Configuration

- 1) On the multilayer switch, create a new VLAN for management of the wireless infrastructure devices. Use VLAN number 10 and name the VLAN 'Management'.
- 2) Create a VLAN interface on the multilayer switch to be used as the default gateway for the Management VLAN. Use IP address 192.168.10.1/24
- 3) Create a DHCP scope on the multilayer switch to allocate IP addresses to Wireless Access Points on the Management VLAN.
Use an address range of 192.168.10.101 to 192.168.10.254.
The default gateway is 192.168.10.1 and the Wireless APs should learn the address of the Wireless LAN Controller.
(A DNS server is not required in this lab environment.)
- 4) You will create a WLAN for Corporate users (staff members) later in this lab exercise. Create a new VLAN for the staff users. Use VLAN number 22 and name the VLAN 'Corporate'.
- 5) Create a VLAN interface on the multilayer switch to be used as the default gateway for the Corporate VLAN. Use IP address 192.168.22.1/24

- 6) You will also create a WLAN for guest users (non-staff members) later in this lab exercise. Create a new VLAN for the guest users. Use VLAN number 23 and name the VLAN 'Guest'.
- 7) Create a VLAN interface on the multilayer switch to be used as the default gateway for the Guest VLAN. Use IP address 192.168.23.1/24
- 8) Verify you now have these VLANs and VLAN interfaces configured (do not worry about the VLAN interface status being up/down, that is expected in this lab environment):

VLAN Name	VLAN Number	IP Subnet	Gateway (on switch)
Management	10	192.168.10.0/24	192.168.10.1
Server	11	192.168.11.0/24	192.168.11.1
Admin	21	192.168.21.0/24	192.168.21.1
Corporate	22	192.168.22.0/24	192.168.22.1
Guest	23	192.168.23.0/24	192.168.23.1

- 9) Port GigabitEthernet1/0/5 on the multilayer switch is connected to the Wireless LAN Controller.
Configure the switchport to support the Corporate and Guest WLANs and management of the Wireless Access Points.
The spanning tree protocol should not check for possible layer 2 loops on the port.
- 10) Port GigabitEthernet1/0/3 and GigabitEthernet1/0/4 on the multilayer switch are connected to the Lightweight Access Points.
Configure the switchports to support the Corporate and Guest WLANs and management of the Wireless Access Points.
The spanning tree protocol should not check for possible layer 2 loops on the port.

Wireless LAN Controller and RADIUS Server Integration

- 11) Check you can ping the Wireless LAN Controller at 192.168.10.11 from the Admin laptop.
- 12) Open <https://192.168.10.11> (use https, not http) in a web browser window on the Admin laptop to open the Wireless LAN Controller administration GUI.
Login with username **admin** and password **Flackbox1**
If you get a 'Host Name Unresolved' error message then close the web browser window, then reopen it and try again.
- 13) On the dashboard Summary page, verify the two Access Points have registered with the WLC.
- 14) Add the RADIUS AAA server at 192.168.10.10 to the Wireless LAN Controller.
Your colleague has already added the Wireless LAN Controller as a client on the RADIUS server with shared secret **Flackbox1**.

DHCP on Wireless LAN Controller

- 15) Wireless DHCP clients can receive their IP address from an external DHCP server or from the Wireless LAN Controller.
Configure a DHCP scope on the WLC for Corporate wireless clients with the address range 192.168.22.101 to 192.168.22.254.
Enter all other relevant details (a DNS server is not required in this lab environment.)
- 16) Configure a DHCP scope on the WLC for Guest wireless clients with the address range 192.168.23.101 to 192.168.23.254.
Enter all other relevant details (a DNS server is not required in this lab environment.)

Logical Interfaces on the Wireless LAN Controller

- 17) Create a logical interface on the Wireless LAN Controller in the Corporate VLAN, with IP address 192.168.22.2.
Wireless clients on the Corporate VLAN should get an IP address from the Wireless LAN Controller.
- 18) Create a logical interface in the Guest VLAN with IP address 192.168.23.2.
Wireless clients on the Guest VLAN should get an IP address from the Wireless LAN Controller.

Wireless LANs

- 19) Create the wireless LAN named 'Corporate'. Clients should be authenticated by the 192.168.10.10 RADIUS server you added earlier, and WPA2 AES encryption should be used.
- 20) Create the wireless LAN named 'Guest'. WPA2 AES encryption should be used, and clients should authenticate with the pre-shared key **Flackbox3**.
- 21) Save the configuration on the Wireless LAN Controller.

Join Clients to the Wireless LANs

- 22) A username **Flackbox** with password **Flackbox2** has been configured on the RADIUS server.
Connect to the 'Corporate' WLAN from the Corporate1 laptop using this username.
- 23) Connect to the 'Guest' WLAN from the Guest1 laptop.

Note that the wireless clients will be assigned IP addresses from the 192.168.10.0/24 subnet in this Packet Tracer lab, rather than the Corporate and Guest DHCP scopes as would happen in a real world environment.

37 Wireless Fundamentals Configuration - Answer Key

In this lab you will configure Corporate and Guest WLANs in a company campus. VLANs and IP subnets have already been set up for the company servers and IT administrators to connect via wired connections:

VLAN Name	VLAN Number	IP Subnet	Gateway (on switch)
Server	11	192.168.11.0/24	192.168.11.1
Admin	21	192.168.21.0/24	192.168.21.1

The IT administrators are restricted to wired connections for security reasons, an 'Admin' WLAN will not be created.

A new Wireless LAN Controller has been added to the network. Your colleague has already performed the initial setup at the command line to give the device IP address 192.168.10.11/24

Two Lightweight Wireless Access Points have just been unboxed and cabled to the Multilayer Switch.

Your job is to configure the new Corporate and Guest WLANs.

Note: Packet Tracer does not support a trunk port to the WLC so you will configure the VLAN information on 'dummy' ports on the switch. The devices are really connected to interfaces GigabitEthernet1/0/11 – 15. Do not change this.

Switch Configuration

- 1) On the multilayer switch, create a new VLAN for management of the wireless infrastructure devices. Use VLAN number 10 and name the VLAN 'Management'.

```
Switch(config)#vlan 10
Switch(config-vlan)#name Management
```

- 2) Create a VLAN interface on the multilayer switch to be used as the default gateway for the Management VLAN. Use IP address 192.168.10.1/24

```
Switch(config)#interface vlan 10
Switch(config-if)#ip address 192.168.10.1 255.255.255.0
```

- 3) Create a DHCP scope on the multilayer switch to allocate IP addresses to Wireless Access Points on the Management VLAN.
Use an address range of 192.168.10.101 to 192.168.10.254.
The default gateway is 192.168.10.1 and the Wireless APs should learn the address of the Wireless LAN Controller.
(A DNS server is not required in this lab environment.)

```
Switch(config)#ip dhcp excluded-address 192.168.10.1
192.168.10.100
Switch(config)#ip dhcp pool Management
Switch(dhcp-config)# network 192.168.10.0 255.255.255.0
Switch(dhcp-config)# default-router 192.168.10.1
Switch(dhcp-config)# option 43 ip 192.168.10.11
```

- 4) You will create a WLAN for Corporate users (staff members) later in this lab exercise. Create a new VLAN for the staff users. Use VLAN number 22 and name the VLAN 'Corporate'.

```
Switch(config)#vlan 22
Switch(config-vlan)#name Corporate
```

- 5) Create a VLAN interface on the multilayer switch to be used as the default gateway for the Corporate VLAN. Use IP address 192.168.22.1/24

```
Switch(config)#interface vlan 22
Switch(config-if)#ip address 192.168.22.1 255.255.255.0
```

- 6) You will also create a WLAN for guest users (non-staff members) later in this lab exercise. Create a new VLAN for the guest users. Use VLAN number 23 and name the VLAN 'Guest'.

```
Switch(config)#vlan 23
Switch(config-vlan)#name Guest
```

- 7) Create a VLAN interface on the multilayer switch to be used as the default gateway for the Guest VLAN. Use IP address 192.168.23.1/24

```
Switch(config)#interface vlan 23
Switch(config-if)#ip address 192.168.23.1 255.255.255.0
```

8) Verify you now have these VLANs and VLAN interfaces configured (do not worry about the VLAN interface status being up/down, that is expected in this lab environment):

VLAN Name	VLAN Number	IP Subnet	Gateway (on switch)
Management	10	192.168.10.0/24	192.168.10.1
Server	11	192.168.11.0/24	192.168.11.1
Admin	21	192.168.21.0/24	192.168.21.1
Corporate	22	192.168.22.0/24	192.168.22.1
Guest	23	192.168.23.0/24	192.168.23.1

```
Switch#show vlan

VLAN Name          Status    Ports
---- --
1    default        active   Gig1/0/3, Gig1/0/4, Gig1/0/5, Gig1/0/6
                           Gig1/0/7, Gig1/0/8, Gig1/0/9, Gig1/0/10
                           Gig1/0/16, Gig1/0/17, Gig1/0/18, Gig1/0/19
                           Gig1/0/20, Gig1/0/21, Gig1/0/22, Gig1/0/23
                           Gig1/0/24, Gig1/1/1, Gig1/1/2, Gig1/1/3
                           Gig1/1/4
10   Management     active   Gig1/0/11, Gig1/0/12, Gig1/0/13, Gig1/0/14
                           Gig1/0/15
11   Server         active   Gig1/0/2
21   Admin          active   Gig1/0/1
22   Corporate      active
23   Guest          active
1002 fddi-default  active
1003 token-ring-default  active
1004 fddinet-default  active
1005 trnet-default   active
```

```

Switch#show ip interface brief
Interface          IP-Address      OK? Method Status      Protocol
GigabitEthernet1/0/1  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/2  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/3  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/4  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/5  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/6  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/7  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/8  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/9  unassigned     YES NVRAM  down       down
GigabitEthernet1/0/10 unassigned     YES NVRAM  down       down
GigabitEthernet1/0/11 unassigned     YES NVRAM  up        up
GigabitEthernet1/0/12 unassigned     YES NVRAM  up        up
GigabitEthernet1/0/13 unassigned     YES NVRAM  up        up
GigabitEthernet1/0/14 unassigned     YES NVRAM  up        up
GigabitEthernet1/0/15 unassigned     YES NVRAM  up        up
GigabitEthernet1/0/16 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/17 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/18 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/19 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/20 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/21 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/22 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/23 unassigned     YES NVRAM  down      down
GigabitEthernet1/0/24 unassigned     YES NVRAM  down      down
GigabitEthernet1/1/1  unassigned     YES NVRAM  down      down
GigabitEthernet1/1/2  unassigned     YES NVRAM  down      down
GigabitEthernet1/1/3  unassigned     YES NVRAM  down      down
GigabitEthernet1/1/4  unassigned     YES NVRAM  down      down
Vlan1              unassigned     YES unset  administratively down down
Vlan10             192.168.10.1  YES manual  up        up
Vlan11             192.168.11.1  YES manual  up        down
Vlan22             192.168.22.1  YES manual  up        down
Vlan23             192.168.23.1  YES manual  up        down

```

9) Port GigabitEthernet1/0/5 on the multilayer switch is connected to the Wireless LAN Controller.
 Configure the switchport to support the Corporate and Guest WLANs and management of the Wireless Access Points.
 The spanning tree protocol should not check for possible layer 2 loops on the port.

The switchport connected to the WLC should be configured as a trunk which carries the AP management and WLAN traffic.

```

Switch(config)#interface GigabitEthernet1/0/5
Switch(config-if)#switchport trunk encapsulation dot1q
Switch(config-if)#switchport mode trunk
Switch(config-if)#switchport trunk allowed vlan 10,22,23
Switch(config-if)#spanning-tree portfast trunk

```

10) Port GigabitEthernet1/0/3 and GigabitEthernet1/0/4 on the multilayer switch are connected to the Lightweight Access Points. Configure the switchports to support the Corporate and Guest WLANs and management of the Wireless Access Points. The spanning tree protocol should not check for possible layer 2 loops on the port.

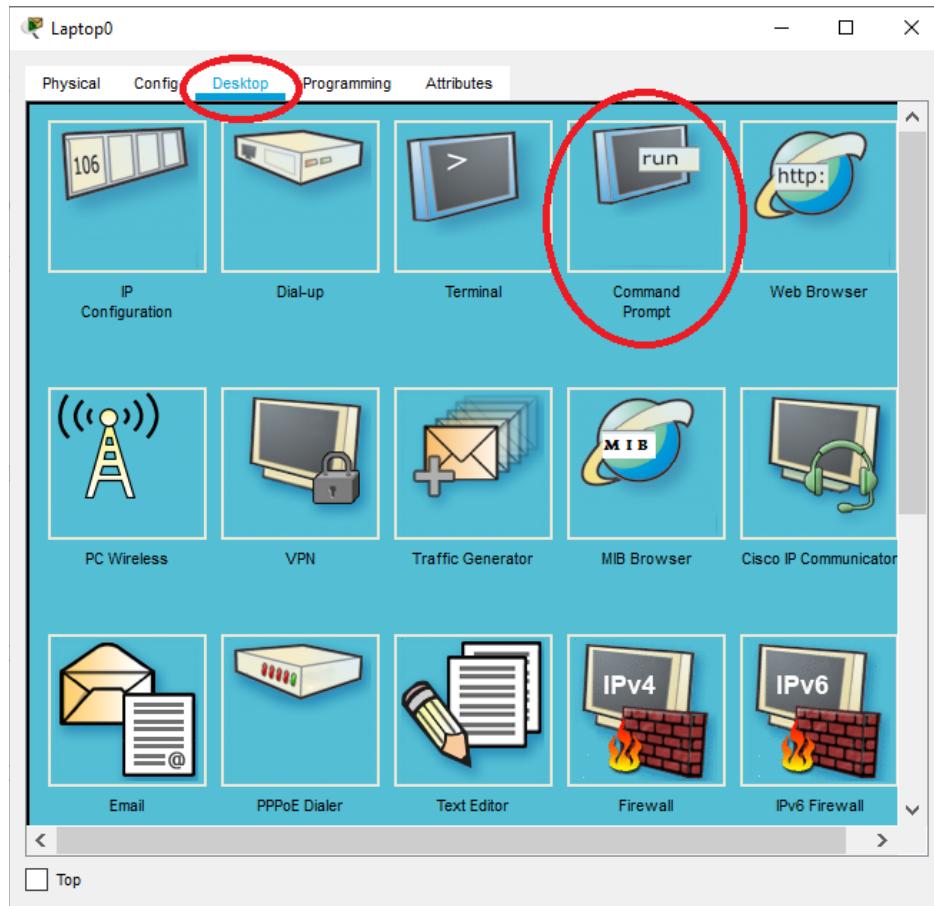
The switchports connected to the Access Points should be configured as access ports for the AP management VLAN. Traffic will be carried inside a CAPWAP tunnel to the WLC.

```
Switch(config)#interface range GigabitEthernet1/0/3 - 4
Switch(config-if)#switchport mode access
Switch(config-if)#switchport access vlan 10
Switch(config-if)#spanning-tree portfast
```

Wireless LAN Controller and RADIUS Server Integration

- 11) Check you can ping the Wireless LAN Controller at 192.168.10.11 from the Admin laptop.

Open a command prompt on the Admin laptop.



```
C:\>ping 192.168.10.11
```

```
Pinging 192.168.10.11 with 32 bytes of data:
```

```
Request timed out.
```

```
Request timed out.
```

```
Reply from 192.168.10.11: bytes=32 time<1ms TTL=254
```

```
Reply from 192.168.10.11: bytes=32 time<1ms TTL=254
```

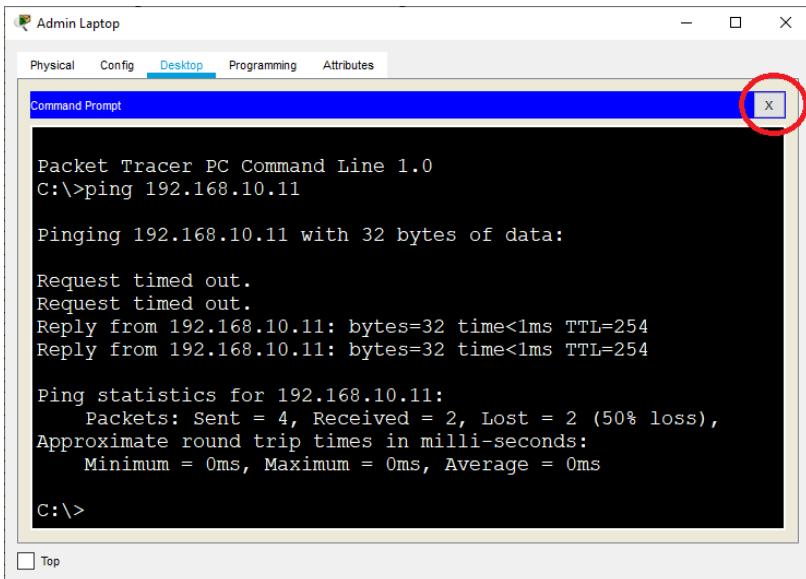
```
Ping statistics for 192.168.10.11:
```

```
Packets: Sent = 4, Received = 2, Lost = 2 (50% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

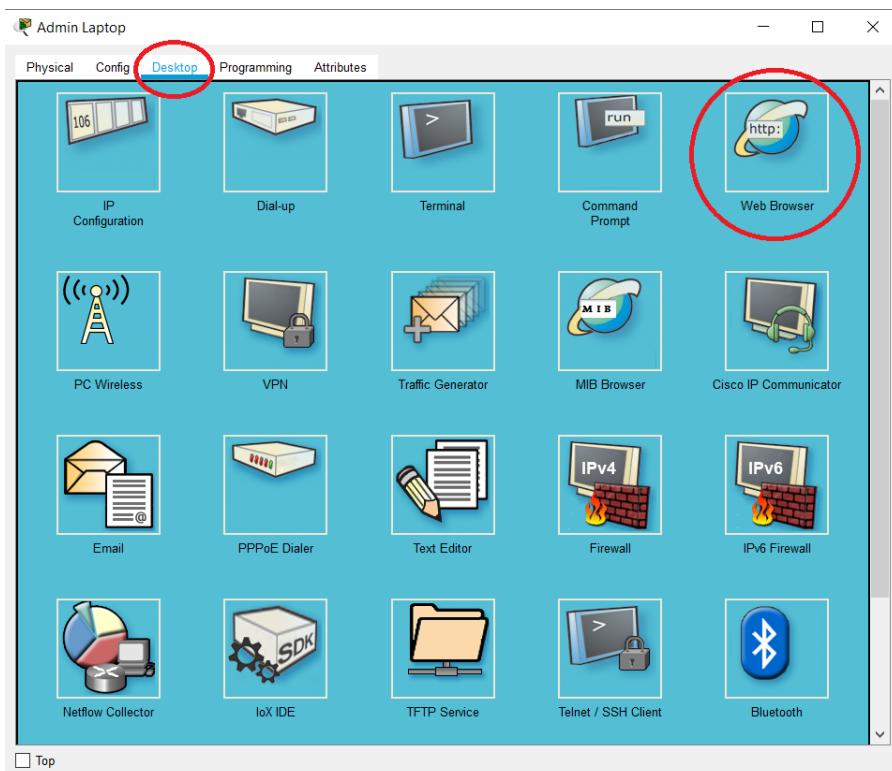
Close the command prompt window.

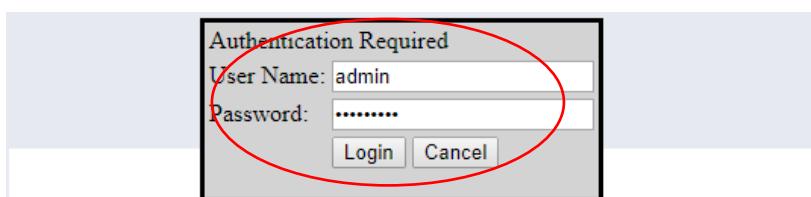
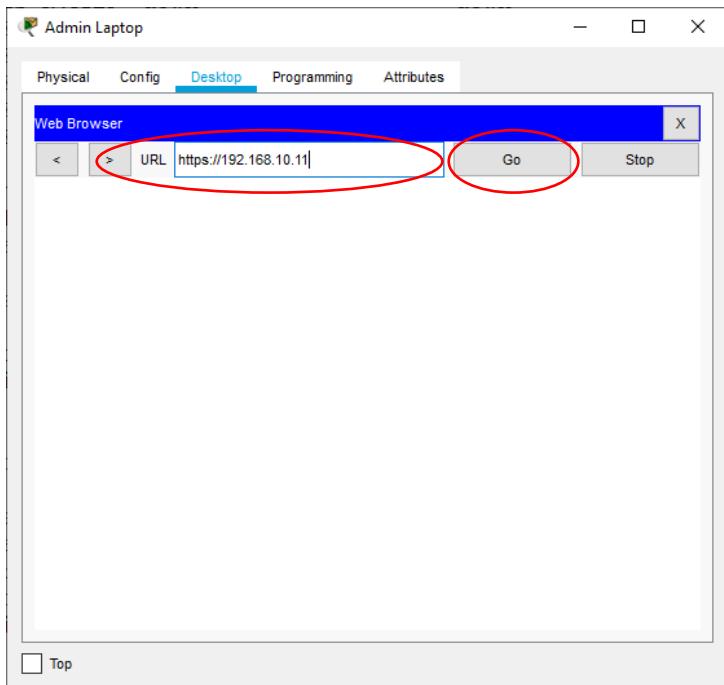


12) Open <https://192.168.10.11> (use https, not http) in a web browser window on the Admin laptop to open the Wireless LAN Controller administration GUI.

Login with username **admin** and password **Flackbox1**

If you get a 'Host Name Unresolved' error message then close the web browser window, then reopen it and try again.





Wireless LAN Controller

Welcome! Please click the login button to enter your user name and password

Login

© 2005 - 2017 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All third party trademarks are the property of their respective owners.

13) On the dashboard Summary page, verify the two Access Points have registered with the WLC.

	Total	Up	Down	
802.11a/n/ac Radios	2	2	0	Detail
802.11b/g Radios	2	2	0	Detail
Dual-band Radios	0	0	0	Detail
All APs	2	2	0	Detail

14) Add the RADIUS AAA server at 192.168.10.10 to the Wireless LAN Controller.
 Your colleague has already added the Wireless LAN Controller as a client on the RADIUS server with shared secret **Flackbox1**.

Click 'Security' > 'AAA' > 'RADIUS' > 'Authentication' then 'New'

Enter the details for the RADIUS server then click 'Apply'.

RADIUS Authentication Servers > New

Server Index (Priority)	<input type="text" value="1"/>
Server IP Address(Ipv4/Ipv6)	<input type="text" value="192.168.10.10"/>
Shared Secret Format	<input type="text" value="ASCII"/>
Shared Secret	<input type="text" value="*****"/>
Confirm Shared Secret	<input type="text" value="*****"/>
Key Wrap	<input type="checkbox"/> (Designed for FIPS customers and requires a key wrap compliant RADIUS server)
Port Number	<input type="text" value="1812"/>
Server Status	<input type="text" value="Enabled"/>
Support for CoA	<input type="text" value="Disabled"/>
Server Timeout	<input type="text" value="2"/> seconds
Network User	<input checked="" type="checkbox"/> Enable
Management	<input checked="" type="checkbox"/> Enable
Management Retransmit Timeout	<input type="text" value="2"/> seconds
IPSec	<input type="checkbox"/> Enable

Verify the RADIUS server is added.

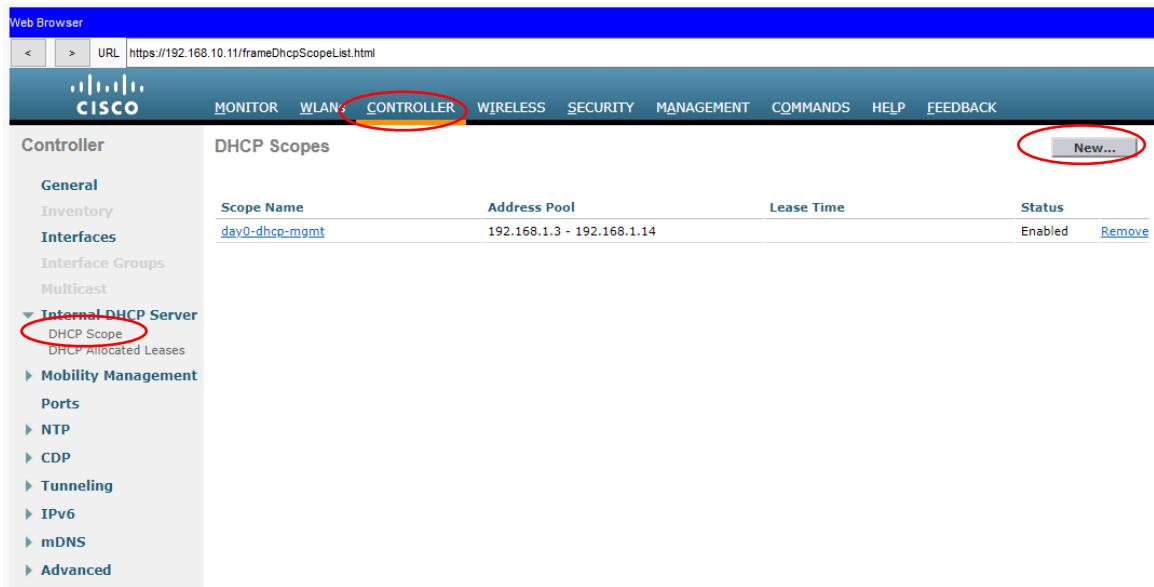
RADIUS Authentication Servers

Auth Called Station ID Type	<input type="text" value="IP Address"/>						
Use AES Key Wrap	<input type="checkbox"/> (Designed for FIPS customers and requires a key wrap compliant RADIUS server)						
MAC Delimiter	<input type="text" value="Hyphen"/>						
Framed MTU	<input type="text" value="1300"/>						
<hr/>							
Network User	Management	Server Index	Server Address(Ipv4/Ipv6)	Port	IPSec	Admin Status	Remove
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="192.168.10.10"/>	<input type="text" value="1812"/>	<input type="text" value="Disabled"/>	<input type="text" value="Enabled"/>	Remove

DHCP on Wireless LAN Controller

15) Wireless DHCP clients can receive their IP address from an external DHCP server or from the Wireless LAN Controller.
Configure a DHCP scope on the WLC for Corporate wireless clients with the address range 192.168.22.101 to 192.168.22.254.
Enter all other relevant details (a DNS server is not required in this lab environment.)

Click 'Controller' > 'Internal DHCP Server' > 'DHCP Scope' then 'New'



Web Browser
< > URL https://192.168.10.11/frameDhcpScopeList.html

CISCO MONITOR WLAN **CONTROLLER** WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Controller DHCP Scopes **New...**

Scope Name	Address Pool	Lease Time	Status
day0-dhcp-mgmt	192.168.1.3 - 192.168.1.14		Enabled Remove

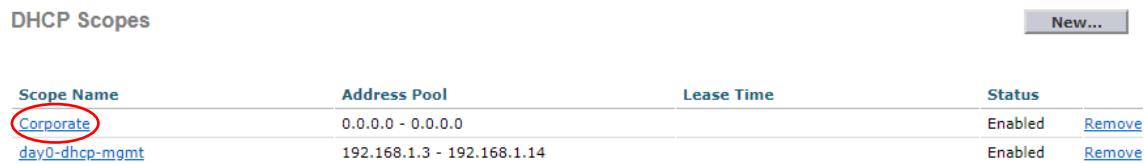
General
Inventory
Interfaces
Interface Groups
Multicast
Internal DHCP Server **DHCP Scope** [DHCP Allocated Leases](#)
Mobility Management
Ports
NTP
CDP
Tunneling
IPv6
mDNS
Advanced

Name the scope 'Corporate' then click 'Apply'.

DHCP Scope > New

Scope Name **Corporate**

Click on the Corporate DHCP scope to configure it.



DHCP Scopes **New...**

Scope Name	Address Pool	Lease Time	Status
Corporate	0.0.0 - 0.0.0		Enabled Remove
day0-dhcp-mgmt	192.168.1.3 - 192.168.1.14		Enabled Remove

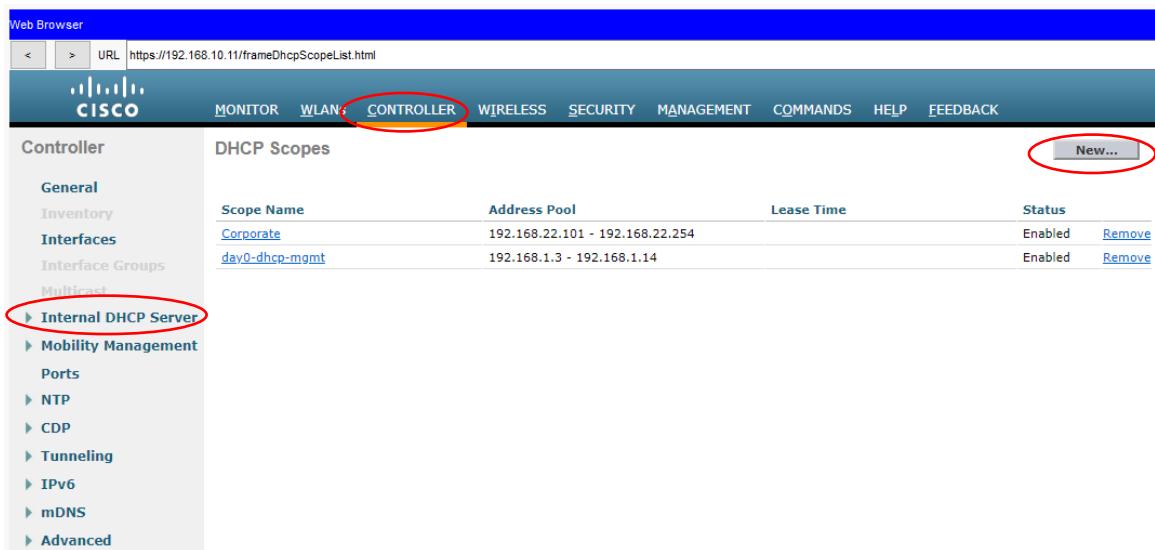
Enter the details then click 'Apply'

DHCP Scope > Edit

Scope Name	Corporate
Pool Start Address	192.168.22.101
Pool End Address	192.168.22.254
Network	192.168.22.0
Netmask	255.255.255.0
Lease Time (seconds)	86400
Default Routers	192.168.22.1
DNS Domain Name	Not Supported
DNS Servers	0.0.0.0
Netbios Name Servers	0.0.0.0
Status	Enabled ▾

16) Configure a DHCP scope on the WLC for Guest wireless clients with the address range 192.168.23.101 to 192.168.23.254.
Enter all other relevant details (a DNS server is not required in this lab environment.)

Click 'Controller' > 'Internal DHCP Server' > 'DHCP Scope' then 'New'



Web Browser

URL: https://192.168.10.11/frameDhcpScopeList.html

CISCO

MONITOR WLAN CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Controller DHCP Scopes New...

Scope Name	Address Pool	Lease Time	Status
Corporate	192.168.22.101 - 192.168.22.254		Enabled Remove
day0-dhcp-mgmt	192.168.1.3 - 192.168.1.14		Enabled Remove

Internal DHCP Server

- General
- Inventory
- Interfaces
- Interface Groups
- Port
- NTP
- CDP
- Tunneling
- IPv6
- mDNS
- Advanced

Name the scope 'Guest' then click 'Apply'.

DHCP Scope > New

Scope Name

Click on the Corporate DHCP scope to configure it.

DHCP Scopes

[New...](#)

Scope Name	Address Pool	Lease Time	Status
Guest	0.0.0.0 - 0.0.0.0		Enabled Remove
Corporate	192.168.22.101 - 192.168.22.254		Enabled Remove
day0-dhcp-mgmt	192.168.1.3 - 192.168.1.14		Enabled Remove

Enter the details then click 'Apply'

DHCP Scope > Edit

Scope Name	Guest		
Pool Start Address	192.168.23.101		
Pool End Address	192.168.23.254		
Network	192.168.23.0		
Netmask	255.255.255.0		
Lease Time (seconds)	86400		
Default Routers	192.168.23.1	0.0.0.0	0.0.0.0
DNS Domain Name	Not Supported		
DNS Servers	0.0.0.0	0.0.0.0	0.0.0.0
Netbios Name Servers	0.0.0.0	0.0.0.0	0.0.0.0
Status	Enabled ▼		

Verify both scopes are enabled.

DHCP Scopes

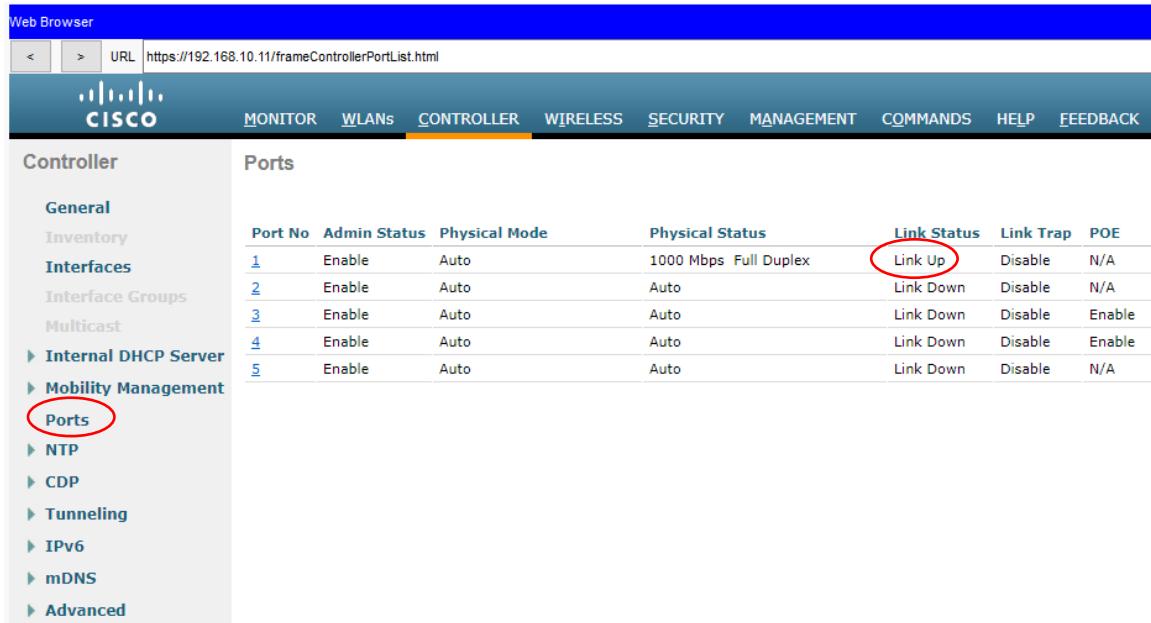
[New...](#)

Scope Name	Address Pool	Lease Time	Status
Guest	192.168.23.101 - 192.168.23.254		Enabled Remove
Corporate	192.168.22.101 - 192.168.22.254		Enabled Remove
day0-dhcp-mgmt	192.168.1.3 - 192.168.1.14		Enabled Remove

Logical Interfaces on the Wireless LAN Controller

17) Create a logical interface on the Wireless LAN Controller in the Corporate VLAN, with IP address 192.168.22.2.
Wireless clients on the Corporate VLAN should get an IP address from the Wireless LAN Controller.

Click 'Ports' to check which physical interface is connected to the switch.



Port No	Admin Status	Physical Mode	Physical Status	Link Status	Link Trap	POE
1	Enable	Auto	1000 Mbps Full Duplex	Link Up	Disable	N/A
2	Enable	Auto	Auto	Link Down	Disable	N/A
3	Enable	Auto	Auto	Link Down	Disable	Enable
4	Enable	Auto	Auto	Link Down	Disable	Enable
5	Enable	Auto	Auto	Link Down	Disable	N/A

Port 1 is connected.

Click 'Controller' > 'Interfaces' then 'New'



Interface Name	VLAN Identifier	IP Address	Interface Type	Dynamic AP Management	IPv6 Address
testvif	untagged	192.168.10.11	Static	Enabled	192.168.10.129

Enter Interface Name 'Corporate' and VLAN ID '22' then click 'Apply'

Interfaces > New

Interface Name	Corporate
VLAN Id	22

Enter the details for the VLAN interface. It should be associated with Port Number 1, and the 192.168.10.11 management address of the WLC should be configured as the DHCP server.

Interfaces > Edit

General Information

Interface Name	Corporate
MAC Address	00:D0:BC:6E:BD:49

Configuration

Guest Lan	<input type="checkbox"/>
Quarantine	<input type="checkbox"/>
Quarantine Vlan Id	0
NAS-ID	

Physical Information

Port Number	1
Backup Port	0
Active Port	0
Enable Dynamic AP Management	<input type="checkbox"/>

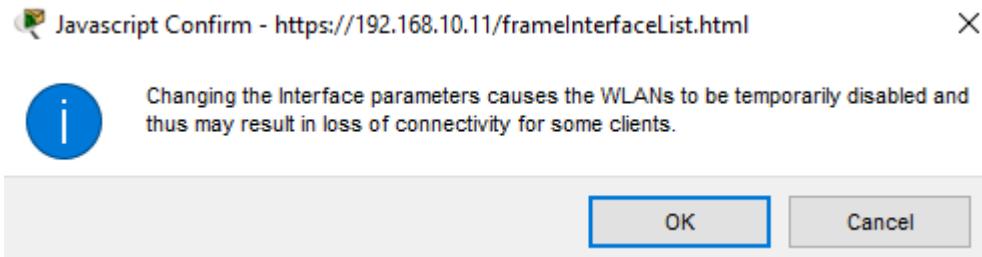
Interface Address

VLAN Identifier	22
IP Address	192.168.22.2
Netmask	255.255.255.0
Gateway	192.168.22.1

DHCP Information

Primary DHCP Server	192.168.10.11
Secondary DHCP Server	
DHCP Proxy Mode	Global ▾
Enable DHCP Option 82	<input type="checkbox"/>

Click 'OK' on the warning message. No wireless clients are connected yet so there will be no disruption.



18) Create a logical interface in the Guest VLAN with IP address 192.168.23.2.

Wireless clients on the Guest VLAN should get an IP address from the Wireless LAN Controller.

Click 'Controller' > 'Interfaces' then 'New'



Enter Interface Name 'Guest' and VLAN ID '23' then click 'Apply'

Interfaces > New

Interface Name	<input type="text" value="Guest"/>
VLAN Id	<input type="text" value="23"/>

Enter the details for the VLAN interface. It should be associated with Port Number 1, and the 192.168.10.11 management address of the WLC should be configured as the DHCP server.

Interfaces > Edit

General Information

Interface Name	Guest
MAC Address	00:04:9A:CE:DD:26

Configuration

Guest Lan	<input type="checkbox"/>
Quarantine	<input type="checkbox"/>
Quarantine Vlan Id	0
NAS-ID	<input type="text"/>

Physical Information

Port Number	1
Backup Port	0
Active Port	0
Enable Dynamic AP Management	<input type="checkbox"/>

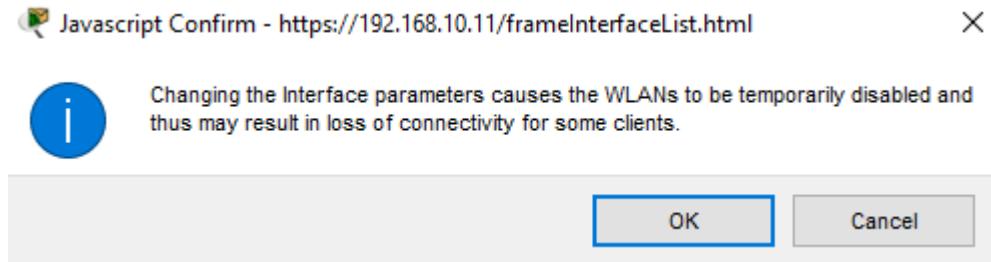
Interface Address

VLAN Identifier	23
IP Address	192.168.23.2
Netmask	255.255.255.0
Gateway	192.168.23.1

DHCP Information

Primary DHCP Server	192.168.10.11
Secondary DHCP Server	<input type="text"/>
DHCP Proxy Mode	Global
Enable DHCP Option 82	<input type="checkbox"/>

Click 'OK' on the warning message. No wireless clients are connected yet so there will be no disruption.



Verify both interfaces have been created.

Interfaces

Interface Name	VLAN Identifier	IP Address	Interface Type	Dynamic AP Management	IPv6 Address	
Corporate	22	192.168.22.2	Dynamic	Disabled		Remove
Guest	23	192.168.23.2	Dynamic	Disabled		Remove
management	untagged	192.168.10.11	Static	Enabled	::/128	
virtual	N/A	192.0.2.1	Static	Not Supported		

Wireless LANs

19) Create the wireless LAN named 'Corporate'. Clients should be authenticated by the 192.168.10.10 RADIUS server you added earlier, and WPA2 AES encryption should be used.

Click on 'WLANs', select 'Create New' in the drop-down then click 'Go'

Web Browser

URL: https://192.168.10.11/frameWlan.html

CISCO

MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

WLANs

WLANs

Current Filter: [Change Filter] [Clear Filter]

Create New ▾ Go

WLAN ID Type Profile Name WLAN SSID Admin Status Security Policies

Enter the details then click 'Apply'

WLANS > New

Type	<input type="button" value="WLAN ▾"/>
Profile Name	<input type="text" value="Corporate"/>
SSID	<input type="text" value="Corporate"/>
ID	<input type="button" value="1 ▾"/>

Associate the WLAN with the 'Corporate' interface. Don't enable the status as you haven't configured the security settings yet. Click 'Apply'.

WLANS > Edit 'Corporate'

<input type="button" value="General"/>	<input type="button" value="Security"/>	<input type="button" value="QoS"/>	<input type="button" value="Policy-Mapping"/>	<input type="button" value="Advanced"/>
Profile Name	<input type="text" value="Corporate"/>			
Type	<input type="text" value="WLAN"/>			
SSID	<input type="text" value="Corporate"/>			
Status	<input type="checkbox"/> Enabled			
Security Policies	None (Modifications done under security tab will appear after applying the changes.)			
Radio Policy	<input type="button" value="All ▾"/>			
Interface/Interface Group(G)	<input type="button" value="Corporate ▾"/>			
Multicast Vlan Feature	<input type="checkbox"/> Enabled			
Broadcast SSID	<input checked="" type="checkbox"/> Enabled			
NAS-ID	<input type="text"/>			

Click on the 'Security' tab and ensure Layer 2 Security is 'WPA + WPA2', the WPA2 Policy is applied with AES encryption, and Authentication Key Management is 802.1X then click 'Apply'.

WLANS > Edit 'Corporate'

General Security QoS Policy-Mapping Advanced

Layer 2 Layer 3 AAA Servers

Layer 2 Security **WPA+WPA2**

MAC Filtering

Fast Transition

Fast Transition

Protected Management Frame

PMF **Disabled**

WPA+WPA2 Parameters

WPA Policy

WPA2 Policy

WPA2 Encryption AES TKIP

Authentication Key Management

802.1X Enable

CCKM Enable

PSK Enable

Click on the 'Security' then 'AAA Servers' tabs, select the RADIUS server you added earlier 'IP:192.168.10.10, Port:1812' as Server 1, and click 'Apply'.

WLANS > Edit 'Corporate'

General Security QoS Policy-Mapping Advanced

Layer 2 Layer 3 AAA Servers

Select AAA servers below to override use of default servers on this WLAN

Radius Servers

Radius Server Overwrite interface Enabled

Authentication Servers	Accounting Servers	EAP Parameters
<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	Enable <input type="checkbox"/>
Server 1 IP:192.168.10.10, Port:1812	None	
Server 2 None	None	
Server 3 None	None	
Server 4 None	None	
Server 5 None	None	
Server 6 None	None	

Radius Server Accounting

Interim Update

On the 'General' tab, tick the 'Enabled' checkbox to enable the WLAN and click 'Apply'.

WLANS > Edit 'Corporate'

General		Security	QoS	Policy-Mapping	Advanced
Profile Name	Corporate				
Type	WLAN				
SSID	Corporate				
Status	<input checked="" type="checkbox"/> Enabled				
Security Policies	[WPA2][Auth(802.1X)] (Modifications done under security tab will appear after applying the changes.)				
Radio Policy	All				
Interface/Interface Group(G)	Corporate				
Multicast Vlan Feature	<input type="checkbox"/> Enabled				
Broadcast SSID	<input checked="" type="checkbox"/> Enabled				
NAS-ID					

20) Create the wireless LAN named 'Guest'. WPA2 AES encryption should be used, and clients should authenticate with the pre-shared key **Flackbox3**.

Click on 'WLANS', select 'Create New' in the drop-down then click 'Go'



Web Browser
 URL: https://192.168.10.11/frameWlan.html

CISCO MONITOR WLANS CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

WLANS WLANS

Current Filter: [Change Filter](#) [Clear Filter](#)

WLAN ID	Type	Profile Name	WLAN SSID	Admin Status	Security Policies
1	WLAN	Corporate	Corporate	Disabled	[WPA2][Auth(802.1X)]

[Create New](#) [Go](#) [Remove](#)

Enter the details then click 'Apply'

WLANs > New

Type	WLAN ▼
Profile Name	Guest
SSID	Guest
ID	2 ▼

Associate the WLAN with the 'Guest' interface and click 'Apply'. Do not enable the status as you haven't configured the security settings yet.

WLANs > Edit 'Guest'

General		Security	QoS	Policy-Mapping	Advanced
Profile Name	Guest				
Type	WLAN				
SSID	Guest				
Status	<input type="checkbox"/> Enabled				
Security Policies	None (Modifications done under security tab will appear after applying the changes.)				
Radio Policy	All ▼				
Interface/Interface Group(G)	Guest ▼				
Multicast Vlan Feature	<input type="checkbox"/> Enabled				
Broadcast SSID	<input checked="" type="checkbox"/> Enabled				
NAS-ID					

Click on the 'Security' tab and ensure Layer 2 Security is 'WPA + WPA2', the WPA2 Policy is applied with AES encryption, Authentication Key Management is PSK and enter the pre-shared key **Flackbox3**, then click 'Apply'. You may need to scroll down to see the field to enter the pre-shared key in.

WLANS > Edit 'Guest'

General Security QoS Policy-Mapping Advanced

Layer 2 Layer 3 AAA Servers

Fast Transition

Fast Transition

Protected Management Frame

PMF

WPA+WPA2 Parameters

WPA Policy

WPA2 Policy

WPA2 Encryption AES TKIP

Authentication Key Management

802.1X Enable

CCKM Enable

PSK Enable

FT 802.1X Enable

FT PSK Enable

PSK Format
*****|

On the 'General' tab, tick the 'Enabled' checkbox to enable the WLAN and click 'Apply'.

WLANS > Edit 'Guest'

General Security QoS Policy-Mapping Advanced

Profile Name Guest

Type WLAN

SSID Guest

Status Enabled

Security Policies [WPA2][Auth(PSK)]
(Modifications done under security tab will appear after applying the changes.)

Radio Policy

Interface/Interface Group(G) Guest

Multicast Vlan Feature Enabled

Broadcast SSID Enabled

NAS-ID

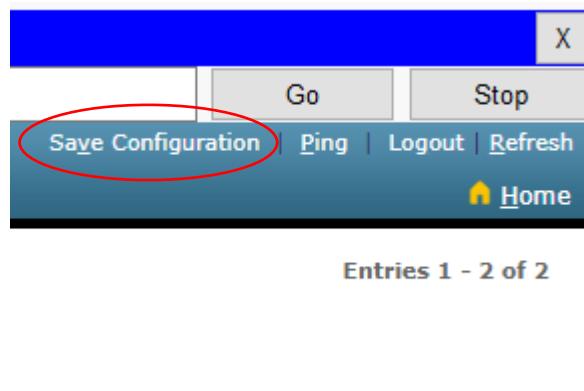
Click 'WLANS' to verify both WLANS are enabled.



WLAN ID	Type	Profile Name	WLAN SSID	Admin Status	Security Policies
1	WLAN	Corporate	Corporate	Enabled	[WPA2][Auth(802.1X)]
2	WLAN	Guest	Guest	Enabled	[WPA2][Auth(PSK)]

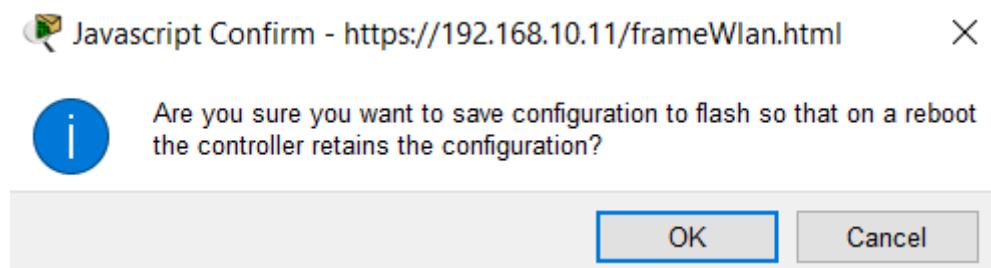
21) Save the configuration on the Wireless LAN Controller.

Click 'Save Configuration' near the top-right corner.



Entries 1 - 2 of 2

Click 'OK' when you see the warning message (this does NOT reboot the controller).



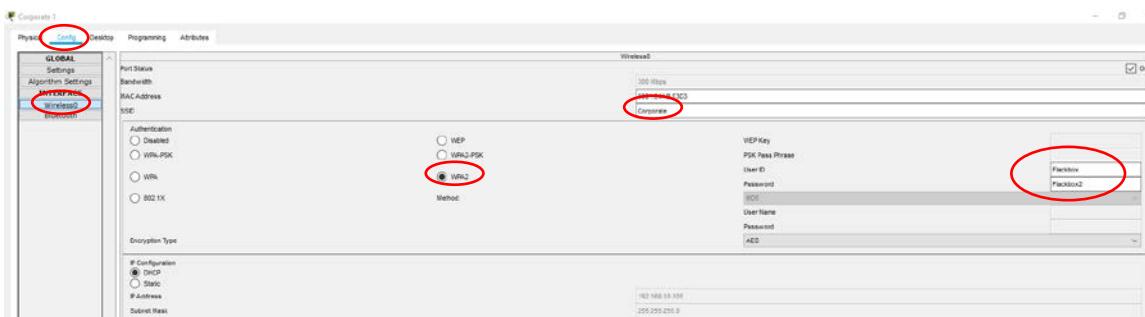
Join Clients to the Wireless LANs

Note that the wireless clients will be assigned IP addresses from the 192.168.10.0/24 subnet in this Packet Tracer lab, rather than the Corporate and Guest DHCP scopes as would happen in a real world environment.

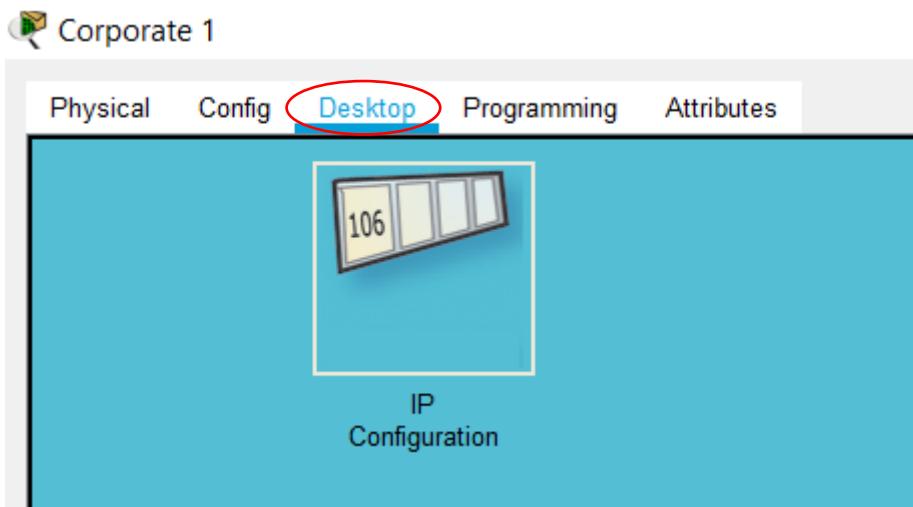
22) A username **Flackbox** with password **Flackbox2** has been configured on the RADIUS server.

Connect to the 'Corporate' WLAN from the Corporate1 laptop using this username.

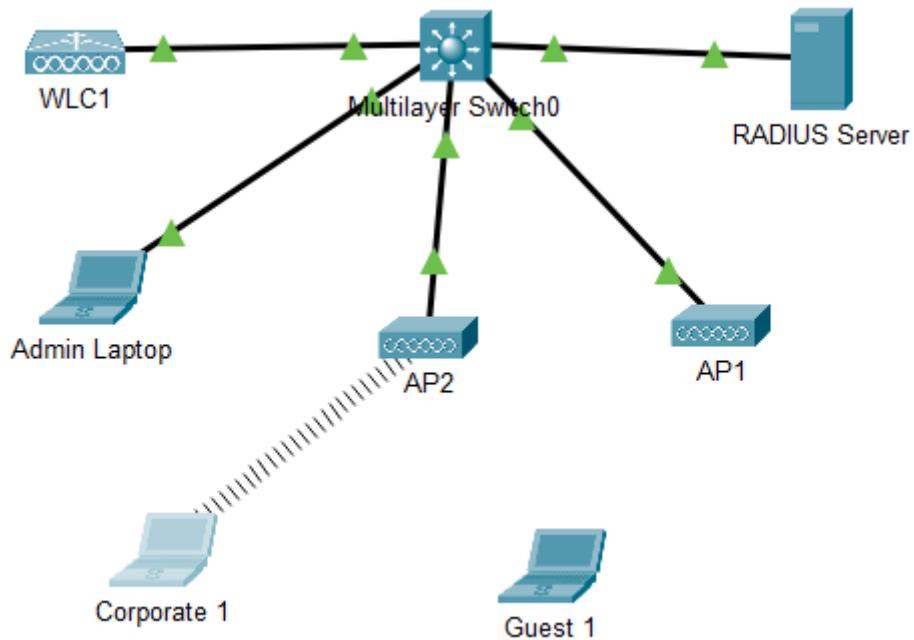
Click on the Corporate1 laptop in the Packet Tracer main window, then 'Config' and 'Wireless0'. Enter the SSID 'Corporate', select WPA2 authentication then enter the user ID Flackbox and password Flackbox2.



Click out of the 'Config' tab to ensure the changes take effect.

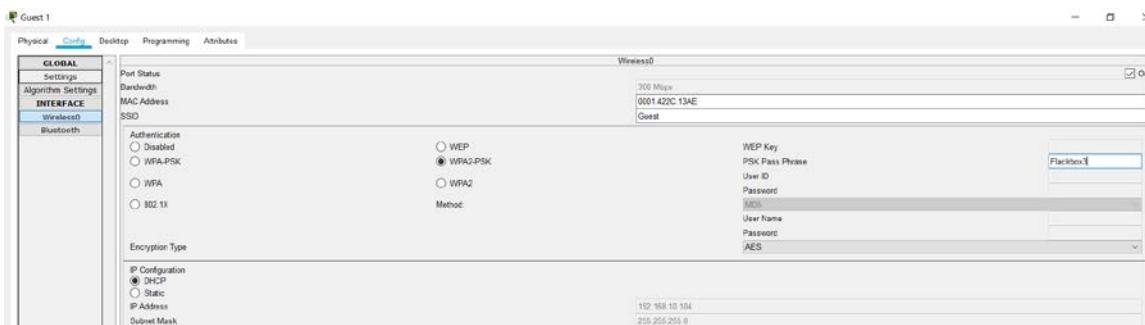


Verify the laptop connects in the Packet Tracer main window.



23) Connect to the 'Guest' WLAN from the Guest1 laptop.

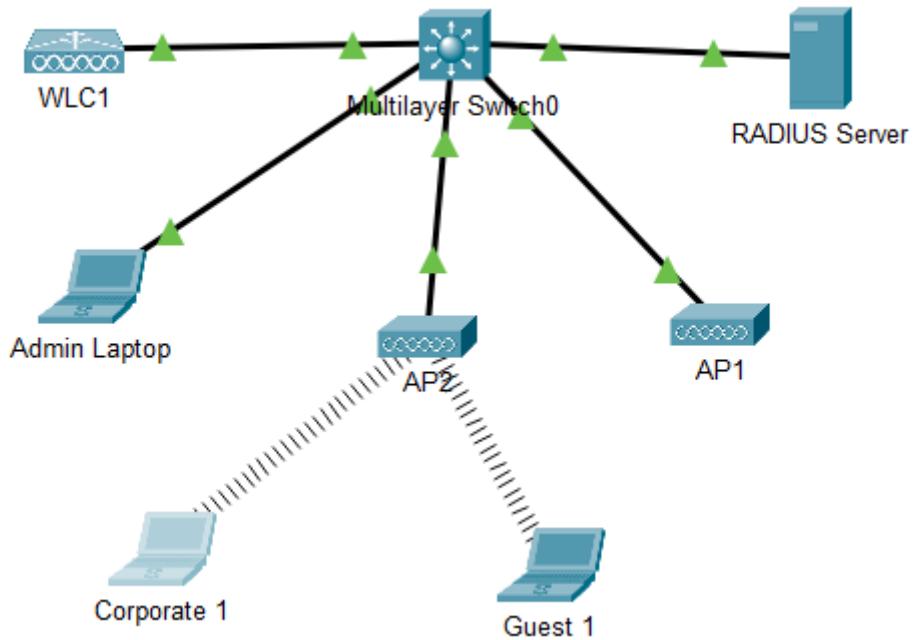
Click on the Guest1 laptop in the Packet Tracer main window, then 'Config' and 'Wireless0'. Enter the SSID 'Guest', select WPA2-PSK authentication then enter the pre-shared key **Flackbox3**



Click out of the 'Config' tab to ensure the changes take effect.



Verify the laptop connects in the Packet Tracer main window.



SPECIAL OFFER – Cisco CCNA Gold Bootcamp

[Click to see the video promo:](#)



My [Cisco CCNA Gold Bootcamp](#) course features comprehensive video tutorials, live demos, quizzes, study notes and advanced lab exercises and covers **everything** you need to pass the Cisco CCNA.

I explain each topic with a detailed tutorial first, show you how to implement it with a live demo, then provide configuration and advanced troubleshooting lab exercises to give you hands on practice. You get the full theoretical and practical knowledge of the technology.

The course offers an **easy to understand, structured approach to shortcut your path** to mastering Cisco networking to the CCNA level and beyond.

[Special Offer: Click Here to enroll in the ‘Cisco CCNA 200-301 Gold Bootcamp’](#)

HD video tutorials covering all topics on the CCNA 200-301 exam

Live lab demo video for every feature



Quizzes and Study Notes in every section to prepare you for the CCNA exam

Complete configuration and troubleshooting hands-on lab exercises, with Solution Guides which double as an on-the-job reference library

Q. Is the course right for me?

I could go on and on about why I think my course is the best Cisco CCNA training resource available, but instead I'll let you hear from the students who've already enrolled.

The course has achieved an absolutely stellar 4.7 out of 5 rating from over 10,000 public student reviews.



Cisco CCNA 200-301 – The Complete Guide

BESTSELLER 306 lectures • 38 total hours

The top rated CCNA course online and only one where all questions get a response. Full lab exercises included. | By Neil Anderson

★★★★★ 4.7
(14,511 Ratings)
42,391 students

"I kept hearing how getting your CCNA would change your life. Well it's true.



I 'Crushed' the CCNA exam thanks to Neil's course, and went from being a mail courier to my first tech job as a network lab engineer working for Cisco."

- John Salmon

Q. How long do I have access to the course for?

You have access to the course forever, your enrolment never expires. And you have access to me for as long as I'm still breathing, I genuinely love to help.

Q. What if I buy the course and realise it is not for me?

The course comes with an unconditional risk-free 30 day money-back guarantee – that's my personal promise of your success.

Q. Is the course worth it?

Let's break it down...

The course contains over 30 hours of High Definition video tutorials, live demos, quizzes, study notes and advanced lab exercises and covers **everything** you need to know to pass the Cisco CCNA. Yours to access forever.

It offers an **easy to understand, structured approach to shortcut your path** to mastering Cisco networking to the CCNA level and beyond.

If you took a comparable course in the classroom it would cost you thousands of dollars.

The course will quickly earn back your investment in yourself many times over. Here's the average base salary when you get your CCNA:

Network Engineer CCNA Salaries

14,333 Salaries Updated Mar 21, 2020



Average Base Pay

\$72,362 /yr



The course delivers INSANE value!

Special Offer: Click Here to enroll in the 'Cisco CCNA 200-301 Gold Bootcamp'

SPECIAL OFFER – AlphaPrep 10 Day Trial

I recommend AlphaPrep for the best CCNA practice tests. They partner with Cisco and the CCNA test provider Pearson to bring you the most accurate questions which are incredibly similar in style and difficulty to the questions you will face in the real exam.

The advanced test engine adapts to your answers to determine your current readiness level and generate the most relevant questions in future. Detailed feedback tells you exactly what topics you know well enough to pass and where you need further study.

Additional book, video tutorial and quiz resources are also included to help you best prepare for your CCNA exam.

[Click here for a 10 day trial.](#)