

CCNP SWITCH Workbook - CCNP - Basic Switching Review Lab

1.1 Preliminary Switch Configuration

Load the *CCNP-Switch-Task1-1* initial configurations before starting.

Tasks

- On each of your switches, perform the following:
 - As you access each switch, issue a command to view the quantity, and naming convention, of all interfaces in that switch. You may want to write down the starting and ending interface names/numbers of each switch.
 - Assign a descriptive Hostname of your choosing for each device shown in the topology diagram.
 - Configure an enable password of **ine** (all letters lowercase).
 - Configure a command so that if you mistype a command in the future, the switch will not attempt to perform a DNS resolution/lookup.
 - Shut down all interfaces, then enable only those interfaces that connect to devices shown in the topology diagram.
 - Configure the switch with a command so that if, while you are typing something, a syslog message is displayed, the switch will automatically repeat what you were typing.
 - Configure each switch to allow incoming Telnet sessions for remote management. Configure a local password of **ine** to authenticate any incoming Telnet sessions.
- On each of your routers, configure the following:
 - As you access each router, issue a command to view the quantity, and naming convention, of all interfaces in that router. You may want to write down the starting and ending interface names/numbers of each router.
 - Configure each router to allow incoming Telnet sessions for remote management. Configure a local password of **ine** to authenticate any incoming Telnet sessions.

Rack1SW1#show interface status

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/1		connected	1	a-full	a-100	10/100BaseTX
Fa0/2		connected	1	a-full	a-100	10/100BaseTX
Fa0/3		notconnect	1	auto	auto	10/100BaseTX
Fa0/4		notconnect	1	auto	auto	10/100BaseTX
Fa0/5		notconnect	1	auto	auto	10/100BaseTX
Fa0/6		notconnect	1	auto	auto	10/100BaseTX
Fa0/7		notconnect	1	auto	auto	10/100BaseTX
Fa0/8		notconnect	1	auto	auto	10/100BaseTX
Fa0/9		notconnect	1	auto	auto	10/100BaseTX
Fa0/10		connected	trunk	a-full	a-100	10/100BaseTX
Fa0/11		connected	trunk	a-full	a-100	10/100BaseTX
Fa0/12		connected	trunk	a-full	a-100	10/100BaseTX
Fa0/13		connected	trunk	a-full	a-100	10/100BaseTX
Fa0/14		connected	trunk	a-full	a-100	10/100BaseTX
Fa0/15		connected	trunk	a-full	a-100	10/100BaseTX
Fa0/16		notconnect	1	auto	auto	10/100BaseTX
Fa0/17		notconnect	1	auto	auto	10/100BaseTX
Fa0/18		notconnect	1	auto	auto	10/100BaseTX
Fa0/19		notconnect	1	auto	auto	10/100BaseTX
Fa0/20		notconnect	1	auto	auto	10/100BaseTX
Fa0/21		connected	1	a-full	a-100	10/100BaseTX

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/22		connected	1	a-full	a-100	10/100BaseTX
Fa0/23		notconnect	1	auto	auto	10/100BaseTX
Fa0/24		connected	1	a-full	a-100	10/100BaseTX
Gi0/1		notconnect	1	auto	auto	Not Present
Gi0/2		notconnect	1	auto	auto	Not Present

Switch-1 Configuration

```
Switch#config t
Enter configuration commands, one per line. End with CNTL/Z. Switch(config)#hostname Switch-1
Switch-1(config)#enable password INE
Switch-1(config)#no ip domain-lookup
Switch-1(config)#interface range fast0/1 - 24 , gig 0/1 - 2
Switch-1(config-if-range)#shutdown
Switch-1(config-if-range)#

Switch-1(config-if-range)#
Switch-1(config-if-range)#exit
```

```
Switch-1(config)#interface range fast 0/1 - 2 , fast 0/10 - 15
Switch-1(config-if-range)#no shutdown
Switch-1(config-if-range)#
Switch-1(config-if-range)#exit Switch-1(config)#line console 0
Switch-1(config-line)#logging synchronous
Switch-1(config-line)#Switch-1(config-line)#line vty 0 5
Switch-1(config-line)#password INE

Switch-1(config-line)#end
Switch-1#
```

Perform the same commands above on Switch-2 and Switch-3, substituting an appropriate Hostname and interface numbers.

```
Router-1#show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	10.10.10.9	YES	manual	up	up
FastEthernet0/1	10.10.10.34	YES	manual	up	up
FastEthernet0/0/0	unassigned	YES	unset	administratively down	down
FastEthernet0/0/1	unassigned	YES	unset	administratively down	down
FastEthernet0/0/2	unassigned	YES	unset	administratively down	down
FastEthernet0/0/3	unassigned	YES	unset	administratively down	down
Serial1/0	unassigned	YES	unset	administratively down	down
Serial1/1	unassigned	YES	unset	administratively down	down
Serial1/2	1.2.1.253	YES	SLARP	up	up
Serial1/3	unassigned	YES	unset	administratively down	down
Vlan1	unassigned	YES	unset	up	down

Router-1 Configuration

```
Rtr-1#conf t
Enter configuration commands, one per line. End with CNTL/Z. Rtr-1(config)#line vty 0 5
Rtr-1(config-line)#password ine

Rtr-1(config-line)#end
Rtr-1#
```

Perform the same commands above on Router-2 through Router-4, substituting an appropriate Hostname and interface numbers.

Verification

At this time, the only VLANs that should exist in your switches are the default VLANs, so all routers should be in VLAN-1 (although they are configured for different IP subnets).

If you have configured the preceding tasks correctly, you should be able to ping and telnet between any pair of routers that are in the same subnet. For example, from Router-1 you should be able to ping and telnet to the IP address of 10.10.10.10 (Router-2's IP address) as well as 10.10.10.33 (also pre-configured on Router-2).

Router-1 Verification

```
Rtr-1#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 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
Rtr-1#Rtr-1#telnet 10.10.10.10
Trying 10.10.10.10 ... Open
|
|
User Access Verification
Password:


Rtr-2>enable
Password:
Rtr-2#
Rtr-2#exit

[Connection to 10.10.10.10 closed by foreign host]
Rtr-1#
```

CCNP SWITCH Workbook - 1. CCNP - Basic Switching Review Lab

1.2 VLANs and VTP

Load the *CCNP-Switch-Task1-2* initial configurations before starting.



Tasks

- Configure VTP on all of your switches using the following guidelines:
 - All switches should utilize a version of VTP that provides support for Token Ring VLANs should you ever implement those in the future.
 - Switch-1 is the only switch that should allow you to manually add (or remove) VLANs. This switch will dynamically propagate any VLAN changes you make to other switches within the same VTP domain.
 - All remaining switches can dynamically add or remove VLANs via receipt of VTP messages from a VTP server, but you should not be able to manually add (or remove) VLANs from any other switches other than Switch-1.
 - All switches should be placed in the VTP domain **INE**.
 - Your switches should not accept any un-authenticated VTP information from any rogue switches. VTP messages should be sent with an MD5 hash of **INE**.
- On Switch-1, configure VLANs 2, 3, 4, and 5 and verify that these VLANs have been propagated to Switch-2 and Switch-3 via VTP.

All preconfigured Enable passwords are **ine**

Normally, VTP information would not propagate between switches unless you had first, pre-configured VLAN Trunking. However you'll notice that (when using the INE CCNP Lab Racks) VLAN Trunks will dynamically form without any user intervention between all three switches. This is due to the hardware models of these switches and their default DTP modes. If replicating this task on non-INE equipment you may have to create VLAN Trunks between switches first, before continuing with this VTP Task.

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#vtp version 2
Switch-1(config)#vtp mode server
Device mode already VTP Server for VLANS. Switch-1(config)#vtp domain INE
Changing VTP domain name from NULL to INE
Switch-1(config)#
*Mar  1 00:22:58.223: %SW_VLAN-6-VTP_DOMAIN_NAME_CHG: VTP domain name changed to INE. Switch-1(config)#
vtp password INE
Setting device VTP password to INE Switch-1(config)#vlan 2-5

Switch-1(config-vlan)#exit
Switch-1(config)#end
Switch-1#
```

Switch-2 Configuration

```
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#vtp password INE
Setting device VLAN database password to INE Switch-2(config)#vtp mode client

Setting device to VTP CLIENT mode.
Switch-2(config)#end
Switch-2#
```

Switch-3 Configuration

```
Switch-3#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-3(config)#vtp password INE
Setting device VLAN database password to INESwitch-3(config)#vtp mode client

Setting device to VTP CLIENT mode.
Switch-3(config)#end
Switch-3#
```

Verification

The first criteria required you to run VTP version-2. VTP version-3 would also have met the stated criteria; however, the switches in this lab do not support VTP version-3. In the above configurations, you'll notice that the VTP version and VTP domain name were only manually configured on Switch-1. This information was then dynamically learned by Switch-2 and Switch-3 (although it would not have hurt anything had you configured this information manually on those two switches as well).

The VTP password of **INE** had to be manually configured on all three switches. Finally, VTP Client mode had to be specified manually on Switch-2 and Switch-3. Had you selected VTP Transparent mode for these two switches, they would not have had the capability to dynamically learn of any new VLANs from the VTP Server (Switch-1).

To verify that VTP successfully propagated the new VLANs created on Switch-1 (the VTP Server), you would use the command **show vtp status** to confirm that all switches were synchronized to the same VTP Configuration Revision number, and the command **show vlan** to confirm that Switch-2 and Switch-3 had learned of VLANs 2-5.

Switch-1 Verification

```
Switch-1#show vtp status
VTP Version capable      : 1 to 3 VTP version running : 2
VTP Domain Name         : INE
VTP Pruning Mode        : Disabled
VTP Traps Generation     : Disabled
Device ID               : 0019.2f45.ec00
Configuration last modified by 0.0.0.0 at 3-1-93 00:23:28
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----VTP Operating Mode : Server
```

```
Maximum VLANs supported locally : 1005
Number of existing VLANs : 9 Configuration Revision : 2

MD5 digest : 0xBE 0x7E 0x34 0x0A 0xA4 0x67 0x5C 0x2C
            0x22 0x5C 0xD3 0x91 0x4D 0x06 0x94 0x6B

Switch-1#
```

Switch-1#**show vlan**

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, Fa0/16, Fa0/17, Fa0/18 Fa0/19, Fa0/20, Fa0/21, Fa0/22 Fa0/23, Fa0/24, Gi0/1, Gi0/2
2 VLAN0002	active	
3 VLAN0003	active	
4 VLAN0004	active	
5 VLAN0005	active	
1002 fddi-default	act/unsup	
1003 trcrf-default	act/unsup	
1004 fddinet-default	act/unsup	
1005 trbrf-default	act/unsup	

<output omitted for brevity>

Switch-2 Verification

Switch-2#**show vtp status**

```
VTP Version : running VTP2
Configuration Revision : 2
Maximum VLANs supported locally : 1005
Number of existing VLANs : 9 VTP Operating Mode : Client
VTP Domain Name : INE

VTP Pruning Mode : Disabled
VTP V2 Mode : Enabled
VTP Traps Generation : Disabled
MD5 digest : 0xBE 0x7E 0x34 0x0A 0xA4 0x67 0x5C 0x2C
Configuration last modified by 0.0.0.0 at 3-1-93 00:23:28
```

Switch-2#

Switch-2#

Switch-2#show vlan

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, Fa0/13, Fa0/14, Fa0/15 Fa0/19, Fa0/20, Fa0/21, Fa0/22 Fa0/23, Fa0/24, Gi0/1, Gi0/2
2	VLAN0002	active	
3	VLAN0003	active	
4	VLAN0004	active	
5	VLAN0005	active	
1002	fddi-default	act/unsup	
1003	trcrf-default	act/unsup	
1004	fddinet-default	act/unsup	
1005	trbrf-default	act/unsup	

<output omitted for brevity>

Switch-3 Verification

Switch-3#show vtp status

VTP Version : running VTP2

Configuration Revision : 2

Maximum VLANs supported locally : 1005

Number of existing VLANs : 9 VTP Operating Mode : Client

VTP Domain Name : INE

VTP Pruning Mode : Disabled

VTP V2 Mode : Enabled

VTP Traps Generation : Disabled

MD5 digest : 0xBE 0x7E 0x34 0x0A 0xA4 0x67 0x5C 0x2C

Configuration last modified by 0.0.0.0 at 3-1-93 00:23:28

Switch-3#

Switch-3#

Switch-3#show vlan


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, Fa0/13, Fa0/14, Fa0/15 Fa0/19, Fa0/20, Fa0/21, Fa0/22 Fa0/23, Fa0/24, Gi0/1, Gi0/2
2 VLAN0002	active	
3 VLAN0003	active	
4 VLAN0004	active	
5 VLAN0005	active	
1002 fddi-default	act/unsup	
1003 trcrf-default	act/unsup	
1004 fddinet-default	act/unsup	
1005 trbrf-default	act/unsup	

<output omitted for brevity>

CCNP SWITCH Workbook - 1. CCNP - Basic Switching Review Lab

1.3 VLAN Trunks

Load the *CCNP-Switch-Task1-3* initial configurations before starting.



Tasks

- Configure all inter-switch links shown in the topology diagram above as 802.1q VLAN Trunks using the following criteria:
 - 802.1q Trunks should be negotiated between switches using some combination of DTP modes.
 - If an 802.1q Trunk stops receiving DTP keepalives from its peer switch, that Trunk should fallback to an Access Port in VLAN-2.
 - The Native VLAN of all 802.1q Trunks should be set to VLAN-3.

All preconfigured Enable passwords are **ine**

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch-1(config)#interface range fastethernet 0/10 - 15
Switch-1(config-if-range)#
switchport trunk encapsulation dot1q
```

```
Switch-1(config-if-range)#switchport mode dynamic desirable
Switch-1(config-if-range)#switchport access vlan 2
Switch-1(config-if-range)#switchport trunk native vlan 3

Switch-1(config-if-range)#no shutdown
Switch-1(config-if-range)#end
Switch-1#
```

Repeat the same steps above on Switch-2 and Switch-3, substituting the appropriate interface numbers.

Verification

Many switches that support both ISL and 802.1q Trunking methods will first require you to select one or the other before configuring the `switchport mode` command. That's why this example started with the `switchport trunk encapsulation dot1q` command.

The first two criteria required you to use some form of the `switchport mode dynamic` command. In the example, all of the switchports have been configured as `dynamic desirable`, but you could also have selected `dynamic auto` on one side of any given link (as long as the other side of that same link was `desirable`).

The command `switchport access vlan 2` ensured that if this interface ever fell back to being an Access Switchport, it would be in Access VLAN-2. Finally, you needed to use the `switchport trunk native vlan` command to ensure that the Native VLAN was set to VLAN-3.

To verify that your interfaces were operational as 802.1q VLAN Trunks (and utilizing VLAN-3 as the Native VLAN), there are several commands you could have issued. Shown below is the output of the command `show interface trunk`.

Switch-1 Verification

```
Switch-1#show interface trunk
```

Port	Mode	Encapsulation	Status	Native vlan
Fa0/10	desirable	802.1q	trunking	3
Fa0/11	desirable	802.1q	trunking	3
Fa0/12	desirable	802.1q	trunking	3
Fa0/13	desirable	802.1q	trunking	3
Fa0/14	desirable	802.1q	trunking	3
Fa0/15	desirable	802.1q	trunking	3

<output omitted for brevity>

The command `show interface switchport | include (Name|Access)` could be used to verify that if any trunk fell back to being an Access Switchport, it would be in Access VLAN-2.

```
Switch-1#show interface switchport | include (Name|Access)
```

```
Name: Fa0/1
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/2
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/3
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/4
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/5
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/6
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/7
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/8
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/9
```

```
Access Mode VLAN: 1 (default)Name: Fa0/10
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
Name: Fa0/11
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
Name: Fa0/12
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
Name: Fa0/13
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
Name: Fa0/14
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
Name: Fa0/15
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
Name: Fa0/16
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/17
```

```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/18
```


```
Access Mode VLAN: 1 (default)
```

```
Name: Fa0/19
Access Mode VLAN: 1 (default)
Name: Fa0/20
Access Mode VLAN: 1 (default)
Name: Fa0/21
Access Mode VLAN: 1 (default)
Name: Fa0/22
Access Mode VLAN: 1 (default)
Name: Fa0/23
Access Mode VLAN: 1 (default)
Name: Fa0/24
Access Mode VLAN: 1 (default)
Name: Gi0/1
Access Mode VLAN: 1 (default)
Name: Gi0/2
Access Mode VLAN: 1 (default)
Switch-1#
```

CCNP SWITCH Workbook - CCNP - Basic Switching Review Lab

1.4 EtherChannel

Load the *CCNP-Switch-Task1-4* initial configurations before starting.



Tasks

- Configure the three links connecting Switch-1 to Switch-2 into an EtherChannel bundle using PAgP to dynamically negotiate and maintain this EtherChannel.

All preconfigured Enable passwords are **ine**

Switch-1 Configuration

```
Switch-1#config t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#
interface range fast 0/10 - 12
Switch-1(config-if-range)#shutdown
Switch-1(config-if-range)#
*Mar  1 19:19:19.702: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to down
Switch-1(config-if-range)#
*Mar  1 19:19:21.590: %LINK-5-CHANGED: Interface FastEthernet0/10, changed state to administratively down
*Mar  1 19:19:21.632: %LINK-5-CHANGED: Interface FastEthernet0/11, changed state to administratively down
*Mar  1 19:19:21.674: %LINK-5-CHANGED: Interface FastEthernet0/12, changed state to administratively down
```

```
*Mar  1 19:19:22.596: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/10, changed state to down
Switch-1(config-if-range)#
*Mar  1 19:19:22.638: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to down
*Mar  1 19:19:22.680: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/12, changed state to down
Switch-1(config-if-range)#Switch-1(config-if-range)#channel-protocol pagp
Switch-1(config-if-range)#channel-group 1 mode desirable
Creating a port-channel interface Port-channel 1

Switch-1(config-if-range)#
*Mar  1 19:19:49.700: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to up
Switch-1(config-if-range)#no shutdown

Switch-1(config-if-range)#end
Switch-1#
```

Switch-2 Configuration

```

Switch-2#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-2(config)#int range fast 0/10 - 12
Switch-2(config-if-range)#shutdown
Switch-2(config-if-range)#channel-protocol pagp
Switch-2(config-if-range)#channel-group 1 mode auto

Creating a port-channel interface Port-channel 1
Switch-2(config-if-range)#no shutdown

Switch-2(config-if-range)#end
Switch-2#
Switch-2#
Switch-2#
Switch-2#
Switch-2#
*Mar  1 19:22:52.971: %SYS-5-CONFIG_I: Configured from console by console
Switch-2#
Switch-2#
*Mar  1 19:22:54.203: %LINK-3-UPDOWN: Interface FastEthernet0/10, changed state to up
*Mar  1 19:22:54.215: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to up
*Mar  1 19:22:54.227: %LINK-3-UPDOWN: Interface FastEthernet0/12, changed state to up
Switch-2#
*Mar  1 19:22:59.667: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/12, changed state to up
*Mar  1 19:22:59.739: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to up
*Mar  1 19:23:00.563: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/10, changed state to up
Switch-2#
*Mar  1 19:23:00.655: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
*Mar  1 19:23:01.655: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel1, changed state to up

Switch-2#

```

Verification

Similar to the use of DTP for VLAN Trunks, both PAgP and LACP can be used to dynamically negotiate and maintain EtherChannels. In the example above, the command `channel-protocol pagp` was not absolutely necessary because PAgP is inferred in the next command, `channel-group 1 mode auto` (the `auto` and `desirable` keywords can only be used with PAgP). However, the `channel-protocol` command was inserted so that anyone reading this configuration file in the future would have no doubt about our intent to use PAgP and not LACP. Also, it is always wise to first administratively disable (`shutdown`) any interfaces before making a Layer-1 or Layer-2 change to those interfaces.

The command `show etherchannel summary` is used below to verify that the EtherChannel between Switch-1 and Switch-2 is functional. Notice that this has been configured as a Layer-2 EtherChannel.

Switch-1 Verification

```
Switch-1#show etherchannel summary
Flags: D - down P - bundled 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

       M - not in use, minimum links not met
       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)        PAgP        Fa0/10(P) Fa0/11(P) Fa0/12(P)
```

Switch-1#

CCNP SWITCH Workbook - CCNP - Basic Switching Review Lab

1.5 Switchport and Switch IP Configuration

Load the *CCNP-Switch-Task1-5* initial configurations before starting.



Tasks

- Configure the interfaces on Switch-1, Switch-2, and Switch-3 that are connected to routers into the appropriate VLANs as shown in the topology diagram above.
- Switchports connected to routers should never transmit, or respond to, DTP frames.
- Configure the following IP addresses on Interface VLAN-1 of each switch:
 - Switch-1: 1.1.1.11 /24
 - Switch-2: 1.1.1.22 /24
 - Switch-3: 1.1.1.33 /24
- Verify that any two routers that share a common subnet can ping each other.
- Verify that from any switch you can ping and telnet into any other switch.

All preconfigured Enable passwords are **ine**

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-1(config)#
interface range fast 0/1 - 2
Switch-1(config-if-range)#switchport mode access
Switch-1(config-if-range)#switchport access vlan 2
Switch-1(config-if-range)#no shutdown
Switch-1(config-if-range)#exit Switch-1(config)#interface vlan 1
Switch-1(config-if)#ip address 1.1.1.11 255.255.255.0
Switch-1(config-if)#no shutdown

Switch-1(config-if)#end
Switch-1#
*Mar  1 21:09:07.503: %SYS-5-CONFIG_I: Configured from console by console
Switch-1#
```

Switch-2 Configuration

```
Switch-2#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-2(config)#
interface range fast 0/1 - 2
Switch-2(config-if-range)#switchport mode access
Switch-2(config-if-range)#switchport access vlan 3
Switch-2(config-if-range)#no shutdown
Switch-2(config-if-range)#exit Switch-2(config)#interface range fast0/3 - 4
Switch-2(config-if-range)#switchport mode access
Switch-2(config-if-range)#switchport access vlan 5
Switch-2(config-if-range)#no shutdown
Switch-2(config-if-range)#exit Switch-2(config)#interface vlan 1
Switch-2(config-if)#ip address 1.1.1.22 255.255.255.0
Switch-2(config-if)#no shutdown

Switch-2(config-if)#end
Switch-2#
```

Switch-3 Configuration

```
Switch-3#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-3(config)#int range fast 0/3 - 4
Switch-3(config-if-range)#switchport mode access
Switch-3(config-if-range)#switchport access vlan 4
Switch-3(config-if-range)#no shutdown
Switch-3(config-if-range)#exit Switch-3(config)#interface vlan 1
Switch-3(config-if)#ip address 1.1.1.33 255.255.255.0

Switch-3(config-if)#no shutdown
Switch-3(config-if)#end
Switch-3#
```

Verification

The stated criteria, "**Switchports connected to routers should never transmit, or respond to, DTP frames,**" hopefully made you aware that you needed to use the command `switchport mode access` on these interfaces. Had you neglected to type that command, your interfaces would have remained in the default state of **switchport mode dynamic <auto|desirable>** and would have been in a state where they could respond to incoming DTP frames.

The command `show interface switchport | include (Name|Access)` could be used to verify your configuration:

Switch-1 Verification

```
Switch-1#show interface switchport | i (Name|Access)
```

```
Name: Fa0/1
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
Name: Fa0/2
```

```
Access Mode VLAN: 2 (VLAN0002)
```

```
<Output omitted for brevity>
```

```
Switch-1#ping 1.1.1.22
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.22, timeout is 2 seconds:
```

```
.!!!!
```

```
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

```
Switch-1#
```

```
Switch-1#ping 1.1.1.33
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.33, timeout is 2 seconds:
```

```
.!!!!
```

```
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
```


```
Switch-1#
```

Repeat the commands above on Switch-2 and Switch-3 for verification, pinging the appropriate IP addresses.

CCNP SWITCH Workbook - 2. Multilayer Switching, DHCP, and SDM

2.1 Switched Virtual Interface Configuration

Load the *CCNP-Switch-Task2-1* initial configurations before starting.



Tasks

- In this task, Switch-1 will be the central device able to route packets between the subnets shown in the topology diagram. As such, configure a Switched Virtual Interface(s) for VLAN-2, VLAN-3, VLAN-4, and VLAN-5 using the IP addresses and subnet masks shown below.

SVI	IP Address
2	10.10.10.9 /29
3	10.10.10.33 /29
4	30.30.30.1 /26
5	20.20.20.97 /27

- Verify that the SVIs you just created in Switch-1 are in the state UP/UP.
- Enable IP routing on Switch-1.

All preconfigured Enable passwords are **ine**.

Switch-1 Configuration

```
<output omitted for brevity>
!ip routing
!
interface Vlan1
 ip address 1.1.1.11 255.255.255.0
!
interface Vlan2
 ip address 10.10.10.9 255.255.255.248
!
interface Vlan3
 ip address 10.10.10.33 255.255.255.248
!
interface Vlan4
 ip address 30.30.30.1 255.255.255.192
!
interface Vlan5 ip address 20.20.20.97 255.255.255.224
!
<output omitted for brevity>
```

Switch-1 Verification

```
Switch-1#sho ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
Vlan1	1.1.1.11	YES	manual	up	up
Vlan2	10.10.10.9	YES	manual	up	up
Vlan3	10.10.10.33	YES	manual	up	up
Vlan4	30.30.30.1	YES	manual	up	up
Vlan5	20.20.20.97	YES	manual	up	up

```
<Output omitted for brevity>
```

The command output above would verify that you correctly configured the SVIs specified in this task, but not that the switch has the ability to route packets to/from these SVIs. To verify that Switch-1 is now capable of routing packets, you would have to inspect the IP routing table, looking for **Connected** routes to the subnets you assigned to the SVIs.

```
Switch-1#show ip route
```

Codes: 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

Gateway of last resort is not set

1.0.0.0/24 is subnetted, 1 subnets C 1.1.1.0 is directly connected, Vlan1

20.0.0.0/27 is subnetted, 1 subnets C 20.20.20.96 is directly connected, Vlan5

10.0.0.0/29 is subnetted, 2 subnets C 10.10.10.8 is directly connected, Vlan2

C 10.10.10.32 is directly connected, Vlan3


30.0.0.0/26 is subnetted, 1 subnets C 30.30.30.0 is directly connected, Vlan4

```
Switch-1#
```

CCNP SWITCH Workbook - 2. Multilayer Switching, DHCP, and SDM

2.2 Configuring the Switch as a DHCP Server

Load the *CCNP-Switch-Task2-2* initial configurations before starting.



Tasks

- In this task, Switch-1 will be configured as a DHCP Server, and all four routers will be DHCP Clients.
- Configure DHCP Server functionality on Switch-1 using the following criteria:

Pool Name	Network Address	Subnet Mask	Default-Router	Lease Time
vlan2	10.10.10.8	/29	10.10.10.9	1-day
vlan3	10.10.10.32	/29	10.10.10.33	1-day
vlan4	30.30.30.0	/26	30.30.30.1	1-day
vlan5	20.20.20.96	/27	20.20.20.97	1-day

- Configure all FastEthernet interfaces on the routers (connected to the switches shown in the topology diagram) to obtain their IP address dynamically via DHCP.
- Verify that your routers have obtained an appropriate DHCP IP address.
- Verify that any one of your routers can ping the IP address of any other router.

All preconfigured Enable passwords are **ine**.

Switch-1 Configuration

```
<output omitted for brevity>
!
!
ip dhcp pool vlan2
    network 10.10.10.8 255.255.255.248
    default-router 10.10.10.9
!
ip dhcp pool vlan3
    network 10.10.10.32 255.255.255.248
    default-router 10.10.10.33
!
ip dhcp pool vlan4
    network 30.30.30.0 255.255.255.192
    default-router 30.30.30.1
!
ip dhcp pool vlan5
    network 20.20.20.96 255.255.255.224    default-router 20.20.20.97
!
<output omitted for brevity>
```

Router-2 Configuration

```
Rtr-2#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Rtr-2(config)#interface range fast0/0 - 1
Rtr-2(config-if-range)#shutdown
Rtr-2(config-if-range)#ip address dhcp
Rtr-2(config-if-range)#no shutdown
Rtr-2(config-if-range)#end
Rtr-2#
Rtr-2#
Oct 29 13:25:05.413: %SYS-5-CONFIG_I: Configured from console by console
```

Rtr-2#

Oct 29 13:25:06.649: %LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up

Oct 29 13:25:07.649: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up

Rtr-2#

Oct 29 13:25:07.673: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up

Rtr-2#

Rtr-2#

Oct 29 13:25:45.165: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/0 assigned DHCP address 10.10.10.11, mask 255.255.255.0

Rtr-2#

Oct 29 13:25:49.265: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 10.10.10.35, mask 255.255.255.0

Rtr-2#

Repeat the command **ip address dhcp** on all other Router FastEthernet interfaces shown in the topology diagram.

Switch-1 Verification

Switch-1#**sho ip dhcp server statistics**

Memory usage 19669 **Address pools 4**

Database agents 0 **Automatic bindings 8**

Manual bindings 0

Expired bindings 0

Malformed messages 0

Secure arp entries 0

Renew messages 0

Message Received

BOOTREQUEST 0 **DHCPDISCOVER 12**

DHCPREQUEST 9

DHCPDECLINE 0

DHCPRELEASE 3

DHCPINFORM 0

Message Sent

BOOTREPLY 0 **DHCPOFFER 12**

DHCPACK 9

DHCPNAK 0

Switch-1#

Switch-1#**sho ip dhcp binding**

Bindings from all pools not associated with VRF:

IP address	Client-ID/ Hardware address/ User name	Lease expiration	Type
------------	--	------------------	------

10.10.10.10	0063.6973.636f.2d30.	Mar 03 1993 12:08 AM	Automatic
--------------------	-----------------------------	-----------------------------	------------------

3031.662e.6361.3035.

2e65.6162.302d.4661.

302f.30

10.10.10.11	0063.6973.636f.2d30.	Mar 03 1993 12:49 AM	Automatic
--------------------	-----------------------------	-----------------------------	------------------

3031.612e.3663.3330.

2e38.6664.652d.4661.

302f.30

10.10.10.34	0063.6973.636f.2d30.	Mar 03 1993 12:46 AM	Automatic
--------------------	-----------------------------	-----------------------------	------------------

3031.662e.6361.3035.

2e65.6162.312d.4661.

302f.31

10.10.10.35 0063.6973.636f.2d30. Mar 03 1993 12:49 AM Automatic

3031.612e.3663.3330.

2e38.6664.662d.4661.

302f.31

20.20.20.99 0063.6973.636f.2d30. Mar 03 1993 12:51 AM Automatic

3031.382e.6239.6261.

2e36.6464.382d.4661.

302f.30

20.20.20.100 0063.6973.636f.2d30. Mar 03 1993 12:51 AM Automatic

3031.632e.3538.3965.

2e37.6165.312d.4661.

302f.31

30.30.30.2 0063.6973.636f.2d30. Mar 03 1993 12:51 AM Automatic

3031.382e.6239.6261.

2e36.6464.392d.4661.

302f.31

30.30.30.3 0063.6973.636f.2d30. Mar 03 1993 12:51 AM Automatic

3031.632e.3538.3965.

2e37.6165.302d.4661.

302f.30

Switch-1#

CCNP SWITCH Workbook - Multilayer Switching, DHCP, and SDM

2.3 CEF and TCAMs

Load the *CCNP-Switch-Task2-3* initial configurations before starting.

Tasks

In this task, you'll gain exposure to some of the commands used to monitor the CEF and TCAM tables in switches. There's little you typically need to do to configure or modify these tables, but knowing how to view their contents can prove to be useful in a production environment.

Cisco Express Forwarding (CEF) uses two tables, the Forwarding Information Base (FIB) and the Adjacency Table. Let's first look at the FIB. In Switch-1, issue the command `show ip cef` and answer the following four questions.

Question 1

Do you have any "**Drop**" entries? If so, do you understand why they are there?

All preconfigured Enable passwords are **ine**.

Switch-1 Verification

```
Switch-1#show ip cef
```

Prefix	Next Hop	Interface
0.0.0.0/0	no route	0.0.0.0/8 drop
0.0.0.0/32	receive	
1.1.1.0/24	attached	Vlan1
1.1.1.0/32	receive	Vlan1
1.1.1.11/32	receive	Vlan1
1.1.1.22/32	attached	Vlan1
1.1.1.33/32	attached	Vlan1
1.1.1.255/32	receive	Vlan1
10.10.10.8/29	attached	Vlan2

10.10.10.8/32	receive	Vlan2
10.10.10.9/32	receive	Vlan2
10.10.10.10/32	attached	Vlan2
10.10.10.11/32	attached	Vlan2
10.10.10.15/32	receive	Vlan2
10.10.10.32/29	attached	Vlan3
10.10.10.32/32	receive	Vlan3
10.10.10.33/32	receive	Vlan3
10.10.10.34/32	attached	Vlan3
10.10.10.35/32	attached	Vlan3
10.10.10.39/32	receive	Vlan3
20.20.20.96/27	attached	Vlan5
Prefix	Next Hop	Interface
20.20.20.96/32	receive	Vlan5
20.20.20.97/32	receive	Vlan5
20.20.20.99/32	attached	Vlan5
20.20.20.100/32	attached	Vlan5
20.20.20.127/32	receive	Vlan5
30.30.30.0/26	attached	Vlan4
30.30.30.0/32	receive	Vlan4
30.30.30.1/32	receive	Vlan4
30.30.30.2/32	attached	Vlan4
30.30.30.3/32	attached	Vlan4
30.30.30.63/32	receive	Vlan4 127.0.0.0/8 drop
224.0.0.0/4	drop	
224.0.0.0/24	receive	240.0.0.0/4 drop
255.255.255.255/32	receive	

Switch-1#

In a router or switch that is NOT running CEF, if a packet were received that ultimately needed to be dropped/discarded (such as a packet with a destination IP address of 127.anything), the CPU of the device would first have to be interrupted from its current task, invoke the **IP Input** process, look up the destination of the packet, discover that there was no matching route in the IP Routing table, and then discard the packet. That's a lot of work to do on a packet that will ultimately be thrown away!

When using CEF (and especially in switches that download CEF tables into hardware TCAM tables for faster processing), 99% of packets never need to be seen by the CPU at all. Instead, in this same example, a packet that was received with a destination IP address of 127.anything would be matched by the "**Drop**" entry in the CEF table and be dropped...without ever bothering the CPU to look up that

packet.

Question 2

In that same output of `show ip cef`, you should see that most of your entries fall into the categories of **Attached** or **Receive**. Do you understand the differences between the two?

Switch-1 Verification

```
Switch-1#show ip cef
```

Prefix	Next Hop	Interface
0.0.0.0/0	no route	
0.0.0.0/8	drop	
0.0.0.0/32	receive	
1.1.1.0/24	attached	Vlan1
1.1.1.0/32	receive	Vlan1
1.1.1.11/32	receive	Vlan1
1.1.1.22/32	attached	Vlan1
1.1.1.33/32	attached	Vlan1
1.1.1.255/32	receive	Vlan1
10.10.10.8/29	attached	Vlan2
10.10.10.8/32	receive	Vlan2
10.10.10.9/32	receive	Vlan2
10.10.10.10/32	attached	Vlan2
10.10.10.11/32	attached	Vlan2
10.10.10.15/32	receive	Vlan2

<output omitted for brevity>

In a CEF FIB table, any ingress IP packet matching an entry marked **Receive** will be sent to the CPU for processing. As an example, from a previous task you configured the IP address **1.1.1.11 on Interface VLAN-1** of Switch-1. If a packet were received by the switch with an exact 32-bit match of the destination IP address of 1.1.1.11, this packet would naturally have to be sent to the CPU for processing because the packet was, indeed, meant for the switch itself. So the CPU would need to "receive" this packet to inspect its contents (the packet might contain a Telnet request for the switch, or maybe someone was trying to ping the switch).

Notice the entry for **1.1.1.22/32**. This is the IP address of interface VLAN-1 on **Switch-2**. For this entry to be visible in the FIB table of Switch-1, we know that, at some point in time, Switch-1 generated an ARP Request for 1.1.1.22 and received

an ARP Reply (most likely because an administrator, logged in to the CLI of Switch-1, issued a Ping or Telnet to Switch-2). Switch-1 recognized the following about the IP address of 1.1.1.22: * That IP address was not locally configured on Switch-1 itself (thus it was not classified as a "Receive" entry). * That IP address was a host belonging to the subnet 1.1.1.0/24, which is "attached" (directly-connected to) Switch-1.

So an entry marked "**Attached**" is either: * An entry for an IP subnet (route) that is directly connected (or "attached") to the switch itself or... * An entry for a host address, in an IP subnet that is directly connected to the switch itself.

Question 3

Imagine a scenario in which one of your switches started having a series of problems (IGP Routing peering relationships periodically dropping, spanning tree loops coming and going, etc.), and these problems were traced back to an extremely high CPU utilization on your switch. With some investigation (and looking at the output of the `show process cpu` command), you notice that the **IP Input** process has been consuming most of your CPU utilization. This specific process is used when the CPU is forced to look up the destination (route) of a packet. However, in a switch, 99% (or more) of all packets should be looked up, and forwarded in hardware, so the IP Input process should rarely be invoked.

An extremely high CPU utilization, caused by massive consumption of the IP Input process on a switch, often indicates a routing table and/or ARP table that is so massive in size that the hardware TCAM memory could not contain all of the entries. If you suspect that this is the case, you may want to obtain a count of the total number of entries in your CEF/FIB table on the switch having problems...and compare that against other switches in your network (at the same layer, Distribution or Core) that are not having problems. This could confirm your suspicions that a switch upgrade might be needed to obtain more TCAM memory, or possibly some IP routing adjustments might be necessary (filtering or summarization).

- Count the total number of entries you see in the output of `show ip cef` . If this switch were a Distribution or Core switch in a live, production environment, manually counting these entries (which could easily number into the tens of thousands) would be impossible. Wouldn't it be great if there were a command that gave you the total count of entries? There is!
- Compare the count you just made with the output of the command `show ip cef epoch` .

Switch-1 Verification

Switch-1#show ip cef

Prefix	Next Hop	Interface
0.0.0.0/0	no route	
0.0.0.0/8	drop	
0.0.0.0/32	receive	
1.1.1.0/24	attached	Vlan1
1.1.1.0/32	receive	Vlan1
1.1.1.11/32	receive	Vlan1
1.1.1.22/32	attached	Vlan1
1.1.1.33/32	attached	Vlan1
1.1.1.255/32	receive	Vlan1
10.10.10.8/29	attached	Vlan2
10.10.10.8/32	receive	Vlan2
10.10.10.9/32	receive	Vlan2
10.10.10.10/32	attached	Vlan2
10.10.10.11/32	attached	Vlan2
10.10.10.15/32	receive	Vlan2
10.10.10.32/29	attached	Vlan3
10.10.10.32/32	receive	Vlan3
10.10.10.33/32	receive	Vlan3
10.10.10.34/32	attached	Vlan3
10.10.10.35/32	attached	Vlan3
10.10.10.39/32	receive	Vlan3
20.20.20.96/27	attached	Vlan5
Prefix	Next Hop	Interface
20.20.20.96/32	receive	Vlan5
20.20.20.97/32	receive	Vlan5
20.20.20.99/32	attached	Vlan5
20.20.20.100/32	attached	Vlan5
20.20.20.127/32	receive	Vlan5
30.30.30.0/26	attached	Vlan4
30.30.30.0/32	receive	Vlan4
30.30.30.1/32	receive	Vlan4
30.30.30.2/32	attached	Vlan4
30.30.30.3/32	attached	Vlan4
30.30.30.63/32	receive	Vlan4
127.0.0.0/8	drop	
224.0.0.0/4	drop	
224.0.0.0/24	receive	
240.0.0.0/4	drop	
255.255.255.255/32	receive	

Switch-1#

There are 38 entries in the sample output above, which is also reflected in the `show ip cef epoch` output below.

```
Switch-1#show ip cef epoch
Table: Default Database epoch:      2(38 entries at this epoch)

Switch-1#
```

Question 4

The CEF FIB table contains all of the IP destinations that the switch knows about, as well as the egress interface for that destination (physical interface or VLAN interface). But it does not contain the Layer 2 rewrite information to create a new Ethernet header for those packets. That information is contained in the CEF **Adjacency Table**.

Imagine that a packet is received by your switch with a destination IP address of **10.10.10.14** (which does not currently exist in your topology). If this were the first time the switch had ever needed to forward a packet to that particular host (10.10.10.14), there would NOT be a /32 FIB entry matching that destination. Instead, the closest match would be the following:

```
Switch-1#show ip cef
Prefix                Next Hop                Interface
<output omitted for brevity> 10.10.10.8/29          attached              Vlan2
```

What does the adjacency look like that is associated with this FIB entry? Issue the command `show ip cef adjacency Vlan2 10.10.10.8`. Your output should resemble the following:

```
Switch-1#show ip cef adj Vlan2 10.10.10.8
% No adjacency for 10.10.10.8 on Vlan2

Switch-1#
```

Does this output indicate that packets going to unknown hosts in the 10.10.10.8/29 subnet will be discarded/dropped? No. When using this command with an IP address, you will only see adjacencies for host addresses. 10.10.10.8 is not a host address in this case; it is a subnet address. So in our example, if a packet were received with a destination of 10.10.10.14 and the only matching FIB entry was an "attached" entry for 10.10.10.8/29, this would point to a **Glean adjacency** (not a normal adjacency used for Host addresses).

Packets with destinations that match a **Glean adjacency** are packets that:

- Match a destination subnet that is directly connected to the switch.
- Match an unknown host within that directly connected subnet.

So when a packet matches a **Glean adjacency** the following will occur:

- Temporarily, a **Punt Adjacency** will be created for that specific destination address.
- The packet that was received will be "punted" to the CPU of the switch, which will trigger the switch to transmit an ARP Request for that host (the original packet will then be dropped).
- When the ARP Response is returned, the switch will have all of the Layer 2 rewrite information it needs to change that Punt Adjacency into a normal adjacency for all future packets received for that host (until the ARP entry expires or is deleted).

To view this special type of **Glean Adjacency**, issue the command `show ip cef adjacency glean`.

```
Switch-1#show ip cef adjacency glean

Prefix          Next Hop      Interface
1.1.1.0/24      attached     Vlan1
10.10.10.8/29   attached     Vlan2
10.10.10.32/29  attached     Vlan3
20.20.20.96/27  attached     Vlan5
30.30.30.0/26   attached     Vlan4
Switch-1#
```

Finally, let's look at a valid adjacency for a known host IP address. Issue the following commands and compare their output:

- `show ip cef adjacency vlan1 1.1.1.22 detail`
- `show ip cef adjacency vlan1 1.1.1.22 internal`
- `show adjacency 1.1.1.22 detail`

Switch-1 Verification

```
Switch-1#Switch-1#show ip cef adjacency vlan1 1.1.1.22 detail
```

```
IPv4 CEF is enabled for distributed and running
```

```
VRF Default:
```

```
38 prefixes (38/0 fwd/non-fwd)
```

```
Table id 0
```

```
Database epoch:          2 (38 entries at this epoch)
```

```
1.1.1.22/32, epoch 2, flags attached
```

```
Adj source: IP adj out of Vlan1, addr 1.1.1.22 056D6920
```

```
Dependent covered prefix type adjfib cover 1.1.1.0/24
```

```
attached to Vlan1
```

```
Switch-1#
```

```
Switch-1#Switch-1#show ip cef adjacency vlan1 1.1.1.22 internal
```

```
IPv4 CEF is enabled for distributed and running
```

```
VRF Default:
```

```
38 prefixes (38/0 fwd/non-fwd)
```

```
Table id 0
```

```
Database epoch:          2 (38 entries at this epoch)
```

```
1.1.1.22/32, epoch 2, flags attached, refcount 4, per-destination sharing
```

```
sources: Adj
```

```
subblocks:
```

```
Adj source: IP adj out of Vlan1, addr 1.1.1.22 056D6920
```

```
Dependent covered prefix type adjfib cover 1.1.1.0/24
```

```
ifnums:
```

```
Vlan1(965): 1.1.1.22
```

```
path 0462E0F0, path list 04623D18, share 1/1, type adjacency prefix, for IPv4
```

```
attached to Vlan1, adjacency IP adj out of Vlan1, addr 1.1.1.22 056D6920
```

```
output chain: IP adj out of Vlan1, addr 1.1.1.22 056D6920
```

```
Switch-1#
```

```
Switch-1#Switch-1#show adjacency 1.1.1.22 detail
```

Protocol	Interface	Address
IP	Vlan1	1.1.1.22(8)
		0 packets, 0 bytes
		epoch 0
		sourced in sev-epoch 0
		Encap length 14
		000C8581A50000192F45EC400800

```
L2 destination address byte offset 0
L2 destination address byte length 6
Link-type after encap: ip
ARP
```

Switch-1#

CCNP SWITCH Workbook - Multilayer Switching, DHCP, and SDM

2.4 Switching Database Manager (SDM)

Load the *CCNP-Switch-Task2-4* initial configurations before starting.

Tasks

In this task, you will gain exposure to some of the commands used to monitor and change the Switch Database Manager (SDM) within Cisco switches.

- Issue a command to view the current SDM template in use on Switch-1. Notice the memory allocation for each section of the TCAM.

All preconfigured Enable passwords are **ine**.

Switch-1 Verification

```
Switch-1#show sdm prefer
```

```
The current template is "desktop default" template.
```

The selected template optimizes the resources in the switch to support this level of features for 8 routed interfaces and 1024 VLANs.

number of unicast mac addresses:	6K
number of IPv4 IGMP groups + multicast routes:	1K
number of IPv4 unicast routes:	8K
number of directly-connected IPv4 hosts:	6K
number of indirect IPv4 routes:	2K
number of IPv4 policy based routing aces:	0
number of IPv4/MAC qos aces:	0.5K
number of IPv4/MAC security aces:	1K

```
Switch-1#
```

Next, modify the TCAM memory utilization in Switch-1 so that it is optimized to provide more memory for IPv4 Unicast routes.

- Issue a command to view the SDM template called Routing. Notice the different TCAM allocation in this template as compared to your current template.

Switch-1 Verification

```
Switch-1#show sdm prefer routing
```

```
"desktop routing" template:
```

The selected template optimizes the resources in the switch to support this level of features for 8 routed interfaces and 1024 VLANs.

number of unicast mac addresses:	3K
number of IPv4 IGMP groups + multicast routes:	1K
number of IPv4 unicast routes:	11K
number of directly-connected IPv4 hosts:	3K
number of indirect IPv4 routes:	8K
number of IPv4 policy based routing aces:	0.5K
number of IPv4/MAC qos aces:	0.5K
number of IPv4/MAC security aces:	1K

```
Switch-1#
```

Finally, change over to the Routing template.

- Ensure that your Running-Config has been saved to the Startup-Config.
- Issue the commands needed to change over, and utilize, the Routing SDM template.
- Verify that the change has taken place and that Routing is now the current SDM template.

Switch-1 Verification

```
Switch-1#copy running-config startup-config
```

```
Destination filename [startup-config]?
```

```
Building configuration...
```

```
[OK]
```

```
0 bytes copied in 1.393 secs (0 bytes/sec)
```

```
Switch-1#
```

```
Switch-1#config t
```

```
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#sdm prefer routing
```

```
Changes to the running SDM preferences have been stored, but cannot take effect until the next reload.
```

```
Use 'show sdm prefer' to see what SDM preference is currently active.
```

```
Switch-1(config)#end
```

```
Switch-1#
```

```
*Mar  2 20:52:58.569: %SYS-5-CONFIG_I: Configured from console by console
```

```
Switch-1#
```

```
Switch-1#reload
```

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

```
Switch-1#sho sdm prefer
```

```
The current template is "desktop routing" template.
```

```
The selected template optimizes the resources in the switch to support this level of features for 8 routed interfaces and 1024 VLANs.
```


number of unicast mac addresses:	3K
number of IPv4 IGMP groups + multicast routes:	1K
number of IPv4 unicast routes:	11K
number of directly-connected IPv4 hosts:	3K
number of indirect IPv4 routes:	8K
number of IPv4 policy based routing aces:	0.5K
number of IPv4/MAC qos aces:	0.5K
number of IPv4/MAC security aces:	1K

Switch-1#

CCNP SWITCH Workbook - PVST 802.1d Spanning-Tree

3.1 802.1d Root Bridge Selection

Load the *CCNP-Switch-Task3-1* initial configurations before starting.



Tasks

In this task, you will manipulate 802.1d PVST Spanning-Tree parameters to ensure that pre-determined switches take on the role of Spanning-Tree Root Bridge for certain VLANs.

- Issue a command to ensure that Switch-2 is the Spanning-Tree Root Bridge for VLAN-3 with a Bridge Priority of 8192.
- Issue a command to ensure that Switch-1 is the Spanning-Tree Root Bridge for VLAN-4 with a Bridge Priority of 8192.

All preconfigured Enable passwords are **ine**.

Switch-2 Configuration and Verification

```
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#
spanning-tree vlan 3 priority 8192
```

```
Switch-2(config)#end
Switch-2#
*Mar  2 21:48:26.404: %SYS-5-CONFIG_I: Configured from console by console
Switch-2#
```

```
Switch-2#show spanning-tree vlan 3
VLAN0003
Spanning tree enabled protocol ieee
Root ID    Priority    8195
           Address      000c.8581.a500 This bridge is the root
           Hello Time  2 sec    Max Age 20 sec    Forward Delay 15 sec

<output omitted for brevity>
```

Switch-1 Configuration and Verification

```
Switch-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-1(config)#
spanning-tree vlan 4 priority 8192

Switch-1(config)#end
Switch-1#
```


```
Switch-1#show spanning-tree vlan 4
VLAN0004
Spanning tree enabled protocol ieee
Root ID    Priority    8196
           Address      0019.2f45.ec00 This bridge is the root
           Hello Time  2 sec    Max Age 20 sec    Forward Delay 15 sec

<output omitted for brevity>
```

CCNP SWITCH Workbook - PVST 802.1d Spanning-Tree

3.2 PVST Path Manipulation (Part 1)

Load the *CCNP-Switch-Task3-2* initial configurations before starting.



Tasks

In this task, you will manipulate 802.1d PVST Spanning-Tree commands to force a specific forwarding path for VLAN-3 traffic. Notice the topology diagram above. When this task is complete, traffic (pings) from Router-2 to Router-3 (both in VLAN-3) will take the path indicated by the dashed arrows.

- Configure a single command, on only a single switch, to achieve the desired Spanning-Tree forwarding path for VLAN-3 traffic from Router-3 to Router-2.
 - Notice that traffic sent between Switch-3 and Switch-1 should be transmitted on **FastEthernet0/15**.
 - The specific link selected in the EtherChannel between Switch-1 and Switch-2 for transmitting frames is irrelevant.
 - You are NOT allowed to modify Root Bridge assignment or disable any interfaces to accomplish this task.

All preconfigured Enable passwords are **ine**.


Switch-3 Configuration

```
Switch-3#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch-3(config)#int fast 0/15Switch-3(config-if)#spanning-tree vlan 3 cost 9

Switch-3(config-if)#end
Switch-3#
```

Verification

Before any configuration was completed, frames sent to/from Router-3 to Router-2 (on VLAN-3) would have taken the following path, because this was the only path that wasn't blocked.



To force the forwarding path that we wanted (with only a single command and limited to our original criteria), the only method was to make Switch-3 believe that the cumulative cost of getting to the Root Bridge (Switch-2) via Switch-1 was lower (less) than its directly connected cost to the Root.

Because Switch-3 had a directly connected FastEthernet link to the Root Bridge (with an associated cost of 19), you had to make Switch-3 believe that the cumulative upstream cost to reach the Root Bridge by way of Switch-1 was **LESS than 19**.


Moving our attention to Switch-1 for a moment, hopefully you noticed that Switch-1 has a 3-port FastEthernet EtherChannel as its Root Port. Locally viewed by Switch-1, this EtherChannel has a cost of nine (9).

```
Switch-1#  
show spanning-tree vlan 3  
  
VLAN0003  
  Spanning tree enabled protocol ieee  
    Root ID    Priority    8195  
              Address      000c.8581.a500 Cost      9  
    Port      56 (Port-channel1)
```

Factoring in the additional FastEthernet link between Switch-3 and Switch-1, the total cumulative cost of this path (for Switch-3 to reach the VLAN-3 Root Bridge) was $19 + 9 = 27$. Clearly, the cost of 27 was not as good as the direct cost to the Root Bridge of 19 that Switch-3 selected.

The solution was (on Switch-3) to locally reduce the cost of FastEthernet 0/15 to a value of "x" so that $9 + x < 19$. So "x" could take any value of nine (9) or less. In the configuration example, a value of nine (9) was selected so that the total cost to reach the VLAN-3 Root Bridge (by way of Switch-1) would be 18...which was less than Switch-3's directly connected cost of 19.

Verification of Spanning-Tree Forwarding Path for VLAN-3



```
Switch-1#sho spanning-tree vlan 3
```

```
VLAN0003
```

```
Spanning tree enabled protocol ieee
```

```

Root ID    Priority    8195
           Address    000c.8581.a500
           Cost       9
           Port       56 (Port-channel1)
           Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

```

```

Bridge ID  Priority    32771 (priority 32768 sys-id-ext 3)
           Address    0019.2f45.ec00
           Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
           Aging Time  300 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/13	Desg	FWD	19	128.15	P2p
Fa0/14	Desg	FWD	19	128.16	P2p
Fa0/15	Desg	FWD	19	128.17	P2p
Po1	Root	FWD	9	128.56	P2p

```
Switch-2#sho spanning-tree vlan 3
```

VLAN0003

Spanning tree enabled protocol ieee

Root ID Priority 8195

Address 000c.8581.a500 This bridge is the root

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 8195 (priority 8192 sys-id-ext 3)

Address 000c.8581.a500

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Aging Time 300 sec

Interface		Role	Sts	Cost	Prio.Nbr	Type
----- Fa0/1 Desg FWD						
19	128.1	P2p	Fa0/16	Desg FWD		
19	128.16	P2p	Fa0/17	Desg FWD		
19	128.17	P2p	Fa0/18	Desg FWD		
19	128.18	P2p	Po1	Desg FWD		
9	128.65	P2p				

Switch-3

```
#sho spanning-tree vlan 3
```

```
VLAN0003
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    8195
           Address    000c.8581.a500
           Cost       18
           Port       15 (FastEthernet0/15)
           Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Bridge ID  Priority    32771 (priority 32768 sys-id-ext 3)
           Address    000e.830d.f680
           Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
           Aging Time  300 sec
```


Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/3	Desg	FWD	19	128.3	P2p Fa0/13Altn BLK 19
128.13	P2p	Fa0/14Altn	BLK 19		
128.14	P2p	Fa0/15	Root FWD 9		
128.15	P2p	Fa0/16Altn	BLK 19		
128.16	P2p	Fa0/17Altn	BLK 19		
128.17	P2p	Fa0/18Altn	BLK 19		
128.18	P2p				

CCNP SWITCH Workbook - PVST 802.1d Spanning-Tree

3.3 PVST Path Manipulation (Part 2)

Load the *CCNP-Switch-Task3-3* initial configurations before starting.



Tasks

In this task, you will manipulate 802.1d PVST Spanning-Tree commands to force a specific forwarding path for VLAN-4 traffic. Notice the topology diagram above. When this task is complete, traffic (pings) from Router-1 to Router-4 (both in VLAN-4) will take the path indicated by the dashed arrows.

- With only a single spanning-tree command (on any single switch), ensure that Switch-3 selects interface FastEthernet0/14 as its Root Port for VLAN-4.
 - You are not allowed to modify port cost on any switch to accomplish this.
 - You are not allowed to disable any interfaces to accomplish this.
 - You are not allowed to change Root Bridge selection to accomplish this.
- Using no more than two spanning-tree commands (on any two switches), ensure that Switch-2 selects interface FastEthernet0/18 as its Root Port for VLAN-4.
 - You are not allowed to disable any interfaces to accomplish this.
 - You are not allowed to change Root Bridge selection to accomplish this.

When you have completed the tasks above, ensure that your Spanning-Tree forwarding path for VLAN-4 resembles the topology image at the top of this page.

All preconfigured Enable passwords are **ine**.

Switch-1 Configuration (this accomplished the first objective)

```
Switch-1#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#int fast 0/14
Switch-1(config-if)#spanning-tree vlan 4 port-priority 16

Switch-1(config-if)#end
Switch-1#
```

Switch-3 Verification

```
Switch-3#show spanning-tree vlan 4

VLAN0004
  Spanning tree enabled protocol ieee
  Root ID    Priority    8196
             Address     0019.2f45.ec00
             Cost        19 Port      14 (FastEthernet0/14)

             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
```

Because you were not allowed to manipulate any port cost values to accomplish the first objective, the only other method available was to manipulate the Sending Port-ID of interface FastEthernet0/14 on Switch-1.

Recall that when a switch has two or more equal-cost paths to the root, the next logical tie breaker is to view the Sending Bridge-IDs of BPDUs received on those paths. If one Sending Bridge-ID of an upstream switch is numerically lower than another upstream switch, the tie is broken.

However, in this case, BDPUs received on ports 0/13 through 0/15 (on Switch-3) all contained the same Sending Bridge-ID: the Bridge-ID of Switch-1.

So the next (and last) tie breaker in this instance was for Switch-3 to look at the

Sending Port-IDs in the BPDUs it was receiving from Switch-1. Recall that the Sending Port-ID field within a BPDU is a two-part field containing a Port-ID and a Port-Priority. The Port-ID section cannot be manipulated or changed, but the Port-Priority field can be changed via Cisco IOS CLI commands.

The default Port Priority of Cisco switches (using 12.x or 15.x IOS) is 128.

Lowering this value to something less than 128 (sixteen (16) was selected in this example) on port 0/14 of Switch-1 made port 0/14 more preferable (from a Spanning-Tree perspective) than ports 0/13 or 0/15 to any downstream switch connected to these three ports (Switch-3).

Switch-2 Configuration (this accomplished part 1 of the second objective)

```
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#interface port-channel 1
Switch-2(config-if)#spanning-tree vlan 4 cost 39

Switch-2(config-if)#end
Switch-2#
```

By artificially increasing the cost of the EtherChannel upstream to the Root Bridge to 39, this path was now less preferable than two hops of FastEthernet (total, cumulative cost of 38) by way of Switch-3 to reach the Root Bridge.

However, left at this stage, interface FastEthernet0/16 was selected as the Root Port on Switch-2 and the remaining interfaces connected to Switch-3 were blocked. The objectives stated that we needed traffic to flow on FastEthernet0/18; therefore, a second step is needed.

```
Switch-2#sho spanning-tree vlan 4

VLAN0004
  Spanning tree enabled protocol ieee
  Root ID    Priority    8196
             Address     0019.2f45.ec00
             Cost        38 Port 16 (FastEthernet0/16)
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32772 (priority 32768 sys-id-ext 4)
             Address     000c.8581.a500
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec
```

Interface	Role	Sts	Cost	Prio.	Nbr	Type

Fa0/4	Desg	FWD	19	128	.4	P2p
Fa0/16	Root	FWD	19	128	.16	P2p
Fa0/17	Altn	BLK	19	128	.17	P2p
Po1	Altn	BLK	39	128	.65	P2p

Switch-2 Configuration (this completed the second objective)


```
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#int fast 0/18
Switch-2(config-if)#spanning-tree vlan 4 cost 18

Switch-2(config-if)#end
Switch-2#
```

The same objective could have been achieved by reducing the Port-Priority (to something less than 128) on interface FastEthernet0/18 of Switch-3.

Verification

Your forwarding path for VLAN-4 should now resemble this.



Switch-2

```
#sho spanning-tree vlan 4
```

```
VLAN0004
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    8196
           Address    0019.2f45.ec00
           Cost      37
           Port      18 (FastEthernet0/18)
           Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Bridge ID  Priority    32772 (priority 32768 sys-id-ext 4)
           Address    000c.8581.a500
           Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
           Aging Time 300 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/4	Desg	FWD	19	128.4	P2p
Fa0/16	Altn	BLK	19	128.16	P2p
Fa0/17	Altn	BLK	19	128.17	P2p
Fa0/18	Root	FWD	18	128.18	P2p

Po1 Altn BLK 39 128.65 P2p

Switch-3

#sho spanning-tree vlan 4

VLAN0004

Spanning tree enabled protocol ieee

Root ID Priority 8196
 Address 0019.2f45.ec00
 Cost 19
 Port 14 (FastEthernet0/14)
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32772 (priority 32768 sys-id-ext 4)
 Address 000e.830d.f680
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
 Aging Time 300 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/13	Altn BLK	19	128.13	P2p	Fa0/14 Root FWD 19 128.14 P2p
Fa0/15	Altn BLK	19	128.15	P2p	
Fa0/16	Desg FWD	19	128.16	P2p	
Fa0/17	Desg FWD	19	128.17	P2p	
Fa0/18	Desg FWD	19	128.18	P2p	

Switch-1#

sho spanning-tree vlan 4

VLAN0004

Spanning tree enabled protocol ieee

Root ID Priority 8196
 Address 0019.2f45.ec00 This bridge is the root
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 8196 (priority 8192 sys-id-ext 4)
 Address 0019.2f45.ec00
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
 Aging Time 300 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----------	------	-----	------	----------	------

Fa0/1	Desg FWD 19	128.3	P2p
Fa0/13	Desg FWD 19	128.15	P2p
Fa0/14	Desg FWD 19	16.16	P2p
Fa0/15	Desg FWD 19	128.17	P2p
Po1	Desg FWD 9	128.56	P2p

CCNP SWITCH Workbook - PVST 802.1d Spanning-Tree

3.4 Portfast, UplinkFast, & BackboneFast

Load the *CCNP-Switch-Task3-4* initial configurations before starting.

Tasks

In this task, you will configure some optional Spanning-Tree features that provide faster reconvergence when changes are detected in the STP topology.

- **Objective 1:** With only a single spanning-tree command (performed on Switch-1, Switch-2, and Switch-3), ensure that when any Access Switchport comes online, it will bypass the Listening and Learning states of Spanning-Tree and proceed directly to the Forwarding state.
 - You are not allowed to implement any interface-level commands to accomplish this objective.
- **Objective 2:** Using a single spanning-tree command (performed on Switch-2 only), ensure that if this switch physically loses its Root Port (port becomes disabled) it will rapidly recover a new Root Port in only 1-2 seconds.
- **Objective 3:** Using a single spanning-tree command (performed on all switches in your topology), ensure that if any switch detects an indirect link failure on the Designated Bridge it is connected to, the Spanning-Tree topology will be able to reconverge in roughly 30 seconds.

Verify that all of the objectives above are functioning as expected.

All preconfigured Enable passwords are **ine**.

Switch Configuration (this accomplished the first objective)

```
Switch-1#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#
spanning-tree portfast default
```

```
%Warning: this command enables portfast by default on all interfaces. You
should now disable portfast explicitly on switched ports leading to hubs,
switches and bridges as they may create temporary bridging loops.
```

```
Switch-1(config)#end
Switch-1#
```

The same command should also be done on Switch-2 and Switch-3.

Verification of Portfast

The command `show spanning-tree detail active` shows you which ports have portfast enabled by default, but, as you can see from the output below, the output of that command is pretty verbose.

```
Switch-1#show spanning-tree detail active
...
<output omitted for brevity>
...
Port 3 (FastEthernet0/1) of VLAN0004 is designated forwarding

Port path cost 19, Port priority 128, Port Identifier 128.3.
Designated root has priority 8196, address 0019.2f45.ec00
Designated bridge has priority 8196, address 0019.2f45.ec00
Designated port id is 128.3, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1 The port is in the portfast mode by default

Link type is point-to-point by default
BPDU: sent 6604, received 0
```

The output of this command can be reduced by using the pipe symbol, the "include" keyword, and creative use of regular expressions (not something you're expected to know as part of the CCNP exam, but still fun to play around with).

```
Switch-1#show spanning-tree detail active | i (designated forwarding|portfast)
Port 15 (FastEthernet0/13) of VLAN0001 is designated forwarding
Port 16 (FastEthernet0/14) of VLAN0001 is designated forwarding
Port 17 (FastEthernet0/15) of VLAN0001 is designated forwarding
Port 15 (FastEthernet0/13) of VLAN0002 is designated forwarding
Port 16 (FastEthernet0/14) of VLAN0002 is designated forwarding
Port 17 (FastEthernet0/15) of VLAN0002 is designated forwarding
Port 15 (FastEthernet0/13) of VLAN0003 is designated forwarding
```

```
Port 16 (FastEthernet0/14) of VLAN0003 is designated forwarding
Port 17 (FastEthernet0/15) of VLAN0003 is designated forwarding
Port 3 (FastEthernet0/1) of VLAN0004 is designated forwarding
The port is in the portfast mode by default
```

```
Port 15 (FastEthernet0/13) of VLAN0004 is designated forwarding
Port 16 (FastEthernet0/14) of VLAN0004 is designated forwarding
Port 17 (FastEthernet0/15) of VLAN0004 is designated forwarding
Port 56 (Port-channel1) of VLAN0004 is designated forwarding
Port 15 (FastEthernet0/13) of VLAN0005 is designated forwarding
Port 16 (FastEthernet0/14) of VLAN0005 is designated forwarding
Port 17 (FastEthernet0/15) of VLAN0005 is designated forwarding
Switch-1#
```

Switch Configuration (this accomplished the second objective)

```
Switch-2#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#spanning-tree uplinkfast

Switch-2(config)#end
Switch-2#
```

Verification of UplinkFast

```
Switch-2#show spanning-tree summary
Switch is in pvst mode
Root bridge for: VLAN0003
Extended system ID          is enabled
PortFast Default            is enabled
PortFast BPDU Guard Default is disabled
PortFast BPDU Filter Default is disabled
Loopguard Default           is disabled
EtherChannel misconfig guard is enabled UplinkFast is enabled

...
<output omitted for brevity>
```

```
Switch-2#show spanning-tree uplinkfast
```

```
UplinkFast is enabled
```

```
Station update rate set to 150 packets/sec.
```

```
UplinkFast statistics
```

```
-----
```

```
Number of transitions via uplinkFast (all VLANs)          : 0
```

```
Number of proxy multicast addresses transmitted (all VLANs) : 0
```

```
Name                Interface List
```

```
-----
```

```
VLAN0001            Fa0/16(fwd), Fa0/17, Fa0/18, Po1
```

```
VLAN0002            Fa0/16(fwd), Fa0/17, Fa0/18, Po1
```

```
VLAN0003
```

```
VLAN0004            Fa0/18(fwd), Fa0/16, Fa0/17, Po1
```

```
VLAN0005            Fa0/16(fwd), Fa0/17, Fa0/18, Po1
```

```
Switch-2#
```

```
Switch-2#debug spanning-tree uplinkfast
```

```
Spanning Tree uplinkfast debugging is on
```

```
Switch-2#
```

To test UplinkFast, will shut down/disable the Root Port for VLAN-1 (FastEthernet0/16).

```
Switch-2#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch-2(config)#interface fast0/16
```

```
Switch-2(config-if)#shutdown
```

```
Switch-2(config-if)#^Z
```

```
Switch-2#
```

```
*Mar 3 01:18:58.079: STP FAST: UPLINKFAST: make_forwarding on VLAN0001 FastEthernet0/17 root port id new: 128.17 pr
```

```
*Mar 3 01:18:58.079: %SPANTREE_FAST-7-PORT_FWD_UPLINK: VLAN0001 FastEthernet0/17 moved to Forwarding (UplinkFast)
```

```
*Mar 3 01:18:58.079: STP: UFAST: removing prev root port Fa0/16 VLAN0001 port-id 8010
```

```
*Mar 3 01:18:58.079: STP FAST: UPLINKFAST: make_forwarding on VLAN0002 FastEthernet0/17 root port id new: 128.17 pr
```

```
...
```

```
<output omitted for brevity>
```

```
...
```

```
Switch-2#un all
All possible debugging has been turned off
Switch-2#
```

Switch Configuration (this accomplished the final objective)

```
Switch-1#
Switch-1#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#
spanning-tree backbonefast

Switch-1(config)#end
Switch-1#
```

To complete the objective, this command would need to be repeated on the remaining switches.

Verification of BackboneFast

```
Switch-1#show spanning-tree summary
Switch is in pvst mode
Root bridge for: VLAN0004
Extended system ID          is enabled
Portfast Default            is enabled
PortFast BPDU Guard Default is disabled
Portfast BPDU Filter Default is disabled
Loopguard Default           is disabled
EtherChannel misconfig guard is enabled
UplinkFast                  is disabled BackboneFast is enabled

...
<output omitted for brevity>
```

CCNP SWITCH Workbook - PVST 802.1d Spanning-Tree

3.5 STP Topology Changes

Load the *CCNP-Switch-Task3-5* initial configurations before starting.

In this section, you will gain exposure to the 802.1d Spanning-Tree topology change process and how it is affected by STP timers.

Task 1

Disable the Portfast feature on interface Fastethernet0/3 of Switch-3.

All preconfigured Enable passwords are **ine**.

Switch-3 Configuration

```
Switch-3#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-3(config)#int fast 0/3
Switch-3(config-if)#no spanning-tree portfast

Switch-3(config-if)#end
```

Switch-3 Verification

```
Switch-3#sho spanning-tree interface fast0/3 portfast
VLAN0003 disabled

Switch-3#
```

Task 2

When a port that is NOT configured for Portfast is disabled, moves from learning to forwarding, or moves from learning to blocking, the switch is triggered to send a **Spanning-Tree Topology Change Notification** toward the Root Bridge. Let's view this in action:

- On Switch-3, issue the command `debug spanning-tree event` .
- While this debug is running, disable (`shutdown`) interface FastEthernet0/3 and look for debug output related to a topology change.
- Enable (`no shutdown`) interface FastEthernet0/3 and watch the debug output again. Notice how only after the port moves from **Learning** to **Forwarding** will you see transmission of an STP Topology Change Notice.

Switch-3 Configuration and Verification

```
Switch-3(config)#int fast 0/3
Switch-3(config-if)#shut
Switch-3(config-if)#*Mar 3 19:48:21.467: STP: VLAN0003 sent Topology Change Notice on Fa0/15
Switch-3(config-if)#
*Mar 3 19:48:23.383: %LINK-5-CHANGED: Interface FastEthernet0/3, changed state to administratively down
*Mar 3 19:48:24.383: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/3, changed state to down
Switch-3(config-if)#
Switch-3(config-if)#Switch-3(config-if)#no shut
Switch-3(config-if)#
Switch-3(config-if)#
*Mar 3 19:48:50.519: set portid: VLAN0003 Fa0/3: new port id 8003*Mar 3 19:48:50.519:
STP: VLAN0003 Fa0/3 -> listening
Switch-3(config-if)#
*Mar 3 19:48:50.903: %LINK-3-UPDOWN: Interface FastEthernet0/3, changed state to up
*Mar 3 19:48:51.903: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/3, changed state to up
Switch-3(config-if)#
Switch-3(config-if)#*Mar 3 19:49:05.519: STP: VLAN0003 Fa0/3 -> learning
Switch-3(config-if)#
Switch-3(config-if)#*Mar 3 19:49:20.519: STP: VLAN0003 sent Topology Change Notice on Fa0/15
*Mar 3 19:49:20.519: STP: VLAN0003 Fa0/3 -> forwarding
Switch-3(config-if)#
```

Task 3

When the STP Root Bridge receives a topology change notification within a particular VLAN, it responds by setting the "**Topology Change**" flag within its next

few BPDUs that it transmits into that VLAN. As switches receive BPDUs with the **TC-Flag** set, they will reduce their MAC Address-Table Aging timer (sometimes called the CAM Aging Time) down to the value of the STP **Forwarding Delay timer** for all MAC addresses learned within that VLAN. The Root Bridge will do this as well for its own MAC addresses learned within that VLAN.

The Root Bridge will continue to transmit BPDUs with the TC-Flag set until the **Forwarding-Delay + Max-Age timers expire (35 seconds total)**, at which point the remaining BPDUs sent will reset the TC-Flag back to zero and all MAC Address-Table timers for that VLAN will increase back to their default value (300 seconds).

You should know that the Spanning-Tree Forwarding Delay timer is 15 seconds by default, and the Max-Age timer is set to 20 seconds.

The next task will enable you to see this process in action.

- If you haven't already done so, ensure that you have enabled (`no shutdown`) interface FastEthernet0/3 of Switch-3.
- While still within interface configuration mode (`Switch(config-if)#`), issue the command `do show spanning-tree vlan 3` and notice the "**Aging Time**" (which reflects the MAC Address-Table Aging Time for this particular VLAN).

```
Switch-3(config)#int fast 0/3
Switch-3(config-if)#do show spanning-tree vlan 3

VLAN0003
  Spanning tree enabled protocol ieee
  Root ID    Priority    8195
             Address     000c.8581.a500
             Cost        18
             Port        15 (FastEthernet0/15)
             Hello Time   2 sec    Max Age 20 sec    Forward Delay 15 sec

  Bridge ID  Priority    32771 (priority 32768 sys-id-ext 3)
             Address     000e.830d.f680
             Hello Time   2 sec    Max Age 20 sec    Forward Delay 15 sec Aging Time 300 sec
```

- If you disabled the previous debug, turn it back on again with the command `do debug spanning-tree events` .

```
Switch-3(config-if)#do debug spanning-tree events
```

```
Spanning Tree event debugging is on
```

```
Switch-3(config-if)#
```

With a text editor like Notepad or Wordpad, type the commands "**do show clock**" and "**do show spanning-tree vlan 3**". Then copy both commands so that you can rapidly paste them into your Telnet client over and over again without typing anything.

- Disable (`shutdown`) interface FastEthernet0/3.
- The moment the debug indicates transmission of a Spanning-Tree Topology Change from your switch, issue the command `do show clock` (notice the timestamp) followed immediately by `do show spanning-tree vlan 3` and notice the value of the "**Aging Time**".
- Repeat the commands `do show clock` and `do show spanning-tree vlan 3` (use the Up arrow on your keyboard to repeat the commands from the command history or paste them in repeatedly from your text editor) until you see the Aging Time increase back to the default value of 300 seconds.
 - How much total time elapsed from the moment the Aging Time first reduced to 15 seconds and finally increased back to its default value? It should have been roughly 35 seconds (Forwarding-Delay + Max-Age).

Switch-3 Configuration and Verification

```
Switch-3(config-if)#shutdown
```

```
Switch-3(config-if)#*Mar 3 20:34:16.550: STP: VLAN0003 sent Topology Change Notice on Fa0/15
```

```
*Mar 3 20:34:18.466: %LINK-5-CHANGED: Interface FastEthernet0/3, changed state to administratively down
```

```
*Mar 3 20:34:19.466: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/3, changed state to down
```

```
Switch-3(config-if)#do show spanning-tree vlan 3
```

```
VLAN0003
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    8195
```

```
Address    000c.8581.a500
```

```
Cost        18
```

```
Port        15 (FastEthernet0/15)
```

```
Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Bridge ID  Priority    32771 (priority 32768 sys-id-ext 3)
```

```
Address      000e.830d.f680
Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec Aging Time 15 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----	-----	-----	-----	-----	-----
Fa0/13	Altn	BLK	19	128.13	P2p
Fa0/14	Altn	BLK	19	128.14	P2p
Fa0/15	Root	FWD	9	128.15	P2p
Fa0/16	Altn	BLK	19	128.16	P2p
Fa0/17	Altn	BLK	19	128.17	P2p
Fa0/18	Altn	BLK	19	128.18	P2p

<output omitted for brevity>

```
Switch-3(config-if)#do show spanning-tree vlan 3
```

VLAN0003

Spanning tree enabled protocol ieee

```
Root ID    Priority    8195
Address     000c.8581.a500
Cost        18
Port        15 (FastEthernet0/15)
Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Bridge ID  Priority    32771 (priority 32768 sys-id-ext 3)
Address     000e.830d.f680
Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec Aging Time 300 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----	-----	-----	-----	-----	-----
Fa0/13	Altn	BLK	19	128.13	P2p
Fa0/14	Altn	BLK	19	128.14	P2p
Fa0/15	Root	FWD	9	128.15	P2p
Fa0/16	Altn	BLK	19	128.16	P2p
Fa0/17	Altn	BLK	19	128.17	P2p
Fa0/18	Altn	BLK	19	128.18	P2p

```
Switch-3(config-if)#do show clock
```

```
*20:34:55.510 UTC Wed Mar 3 1993
```

When you're done with this page, ensure that you have:

- **Re-configured the STP Portfast feature** for interface FastEthernet0/3 of Switch-3
- Enabled (`no shutdown`) interface FastEthernet0/3 on Switch-3
- Disabled all debugs with the command `undebug all` (or simply "`u all`")

Switch-3#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Switch-3(config)#**int fast 0/3**

Switch-3(config-if)#**spanning-tree portfast**

%Warning: portfast should only be enabled on ports connected to a single host. Connecting hubs, concentrators, switches, bridges, etc... to this interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION

%Portfast has been configured on FastEthernet0/3 but will only have effect when the interface is in a non-trunking mode. Switch-3(config-if)#**no shutdown**

Switch-3(config-if)#end

Switch-3#unde

*Mar 3 20:51:48.478: %SYS-5-CONFIG_I: Configured from console by console

*Mar 3 20:51:49.026: set portid: VLAN0003 Fa0/3: new port id 8003

*Mar 3 20:51:49.026: STP: VLAN0003 Fa0/3 ->jump to forwarding from blockingSwitch-3#**undebug all**

*Mar 3 20:51:49.398: %LINK-3-UPDOWN: Interface FastEthernet0/3, changed state to up

*Mar 3 20:51:50.398: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/3, changed state to up

All possible debugging has been turned off

Switch-3#

Switch-3#

CCNP SWITCH Workbook - PVST 802.1d Spanning-Tree

3.6 BPDUGuard & Err-Disable Recovery

Load the *CCNP-Switch-Task3-6* initial configurations before starting.

In this section, you will gain exposure to an optional security feature that can protect your switch from receiving, and processing, BPDUs from rogue devices.

Task 1

- Verify that interface FastEthernet0/1 on Switch-2 still has Portfast operational.
- Configure an additional feature on Fast0/1 so that if this port receives an incoming BPDU, it will go into the err-disable state.
- Configure another feature that will automatically bring any port out of err-disable after 30 seconds, no matter what caused the port to become err-disabled.

All preconfigured Enable passwords are **ine**.

Switch-2 Configuration

```
Switch-2#sho spanning-tree int fast 0/1 portfast
```

```
VLAN0003          enabled
```

```
Switch-2#
```

```
Switch-2(config)#int fast 0/1
Switch-2(config-if)#spanning-tree bpduguard enable

Switch-2(config-if)#exit Switch-2(config)#errdisable recovery cause all
Switch-2(config)#errdisable recovery interval 30

Switch-2(config)#end
```

Switch-3 Verification

```
Switch-2#show errdisable recovery
```

ErrDisable Reason	Timer Status
-----	-----
arp-inspection	Enabled
bpduguard	Enabled
channel-misconfig	Enabled
dhcp-rate-limit	Enabled
dtp-flap	Enabled
gbic-invalid	Enabled
l2ptguard	Enabled
link-flap	Enabled
mac-limit	Enabled
link-monitor-failure	Enabled
loopback	Enabled
oam-remote-failure	Enabled
pagp-flap	Enabled
port-mode-failure	Enabled
psecure-violation	Enabled
security-violation	Enabled
sfp-config-mismatch	Enabled
storm-control	Enabled
udld	Enabled
unicast-flood	Enabled
vmps	Enabled

```
Timer interval: 30 seconds
```

```
Interfaces that will be enabled at the next timeout:
```

```
Switch-2#
```

Task 2

To see the BPDUGuard feature in action, we need the device connected to FastEthernet0/1 on Switch-2 (which is Router-2 in this case) to send a BPDU. Normally, routers do not participate in Spanning-Tree, but with Cisco IOS, there is almost always a way to work around any problem.

The process of configuring a router to participate in 802.1d Spanning-Tree is not something you need to know for your CCNP SWITCH

exam. This process shown only to demonstrate the Spanning-Tree BPDUGuard feature.

First, ensure that you have two simultaneous Telnet windows open. One should be to Router-2, and the other should be to Switch-2.

- Configure Router-2 to participate in Spanning-Tree and generate BPDUs.
- Watch the BPDUGuard feature on Switch-2 react to these incoming BPDUs, and place the port into ERR-DISABLE state.
- Watch the errdisable recovery feature (that you previously configured) dynamically recover this port after 30 seconds.

Router-2 Configuration

```
Rtr-2#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Rtr-2(config)#bridge 1 protocol ieee
Rtr-2(config)#bridge 1 priority 4096
Rtr-2(config)#interface fast0/1
Rtr-2(config-if)#shutdown
Rtr-2(config-if)#no ip address
Rtr-2(config-if)#
Oct 31 12:31:52.702: %LINK-5-CHANGED: Interface FastEthernet0/1, changed state to administratively down
Oct 31 12:31:53.702: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down
Rtr-2(config-if)#bridge-group 1

Rtr-2(config-if)#
```

```
Rtr-2(config-if)#no shut

Rtr-2(config-if)#
Oct 31 12:35:37.298: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
Oct 31 12:35:38.298: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
Rtr-2(config-if)#
Oct 31 12:35:45.218: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down
Rtr-2(config-if)#
```

BDPUGuard Verification on Switch-2

```
Switch-2#
```

```
*Mar  3 23:58:53.038: %SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port Fa0/1 with BPDU Guard enabled. Disabling po
```

```
Switch-2#
```

```
*Mar  3 23:58:53.038: %PM-4-ERR_DISABLE: bpduguard error detected on Fa0/1, putting Fa0/1 in err-disable state
```

```
Switch-2#
```

30-seconds later, the **errdisable recovery feature** attempts to bring the port back up, but moments later ANOTHER BPDU is received, causing **BPDUGuard** to errdisable the port.

```
Switch-2# *Mar  3 23:59:23.038:
```

```
%PM-4-ERR_RECOVER: Attempting to recover from bpduguard err-disable state on Fa0/1
```

```
Switch-2# *Mar  3 23:59:26.042:
```

```
%SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port Fa0/1 with BPDU Guard enabled. Disabling port.
```

```
Switch-2# *Mar  3 23:59:26.042:
```

```
%PM-4-ERR_DISABLE: bpduguard error detected on Fa0/1, putting Fa0/1 in err-disable state
```

```
Switch-2#
```

Task 3

Go back to Router-2 and remove the commands that caused it to participate in Spanning-Tree.

```
Rtr-2(config)#int fast0/1
```

```
Rtr-2(config-if)#shut
```

```
Oct 31 12:42:59.650: %LINK-5-CHANGED: Interface FastEthernet0/1, changed state to administratively down
```

```
Rtr-2(config-if)#no bridge-group 1
```

```
Rtr-2(config-if)#no shut
```


```
Rtr-2(config-if)#end
```

```
Rtr-2#
```

CCNP SWITCH Workbook - PVST 802.1d Spanning-Tree

3.7 BPDUFilter, LoopGuard, & RootGuard

Load the *CCNP-Switch-Task3-7* initial configurations before starting.



In this section, you will gain exposure to the BPDUFilter, LoopGuard, and RootGuard Spanning-Tree features.

Task 1: Preliminary Configuration

- Disable the following interfaces on Switch-2: FastEthernet0/17 - 18.
- Disable the following interfaces on Switch-3: FastEthernet0/14 - 15.
- Ensure that your Spanning-Tree topology for VLAN-4 resembles the diagram above.

All preconfigured Enable passwords are **ine**.

Switch-3 Configuration

```
Switch-3#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-3(config)#
int range fast 0/14 - 15 , fast 0/17 - 18
Switch-3(config-if-range)#shutdown

Switch-3(config-if-range)#end
Switch-3#
```

Spanning-Tree Topology Verification (Switch-1)

```
Switch-1#sho spanning-tree vlan 4

VLAN0004
  Spanning tree enabled protocol ieee
  Root ID    Priority    8196
             Address      0019.2f45.ec00 This bridge is the root
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8196  (priority 8192 sys-id-ext 4)
             Address      0019.2f45.ec00
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface                Role Sts Cost      Prio.Nbr Type
-----
Fa0/1                  Desg FWD 19      128.3    P2p Edge
Fa0/13                 Desg FWD 19      128.15   P2p
Po1                    Desg FWD 9       128.56   P2p
```

Spanning-Tree Topology Verification (Switch-2)

```
Switch-2#sho spanning-tree vlan 4
```

```
VLAN0004
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    8196
           Address    0019.2f45.ec00
           Cost        3038
           Port        16 (FastEthernet0/16)
           Hello Time   2 sec   Max Age 20 sec   Forward Delay 15 sec
```

```
Bridge ID  Priority    49156 (priority 49152 sys-id-ext 4)
           Address    000c.8581.a500
           Hello Time   2 sec   Max Age 20 sec   Forward Delay 15 sec
           Aging Time   300 sec
```

```
Uplinkfast enabled
```

```
Interface          Role Sts Cost      Prio.Nbr Type
```

```
-----
Fa0/4              Desg FWD 3019      128.4   P2p Edge
Fa0/16             Root FWD 3019      128.16  P2p
Po1                Altn BLK 3039      128.65  P2p
```

Spanning-Tree Topology Verification (Switch-3)

```
Switch-3#sho spanning-tree vlan 4
```

```
VLAN0004
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    8196
           Address    0019.2f45.ec00
           Cost        19
           Port        13 (FastEthernet0/13)
           Hello Time   2 sec   Max Age 20 sec   Forward Delay 15 sec
```

```
Bridge ID  Priority    32772 (priority 32768 sys-id-ext 4)
           Address    000e.830d.f680
           Hello Time   2 sec   Max Age 20 sec   Forward Delay 15 sec
           Aging Time   300 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/13	Root	FWD	19	128.13	P2p
Fa0/16	Desg	FWD	19	128.16	P2p

Task 2: Creating a Spanning-Tree Loop

- Enable the BPDUFilter feature on Switch-1's Port-Channel interface leading to Switch-2. This will prevent Switch-1 from transmitting BPDUs to Switch-2.
- Notice that after Max-Age expires, Switch-2 will take its EtherChannel out of the Blocking state and into the Forwarding state. You will now have a Spanning-Tree Loop.

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#int port-channel 1
Switch-1(config-if)#spanning-tree bpdufilter enable

Switch-1(config-if)#end
```

Spanning-Tree Loop Verification

```
Switch-1#sho spanning-tree vlan 4

<output omitted for brevity>

Interface                Role Sts Cost      Prio.Nbr Type
-----
19          128.3      P2p Edge Fa0/13 Desg FWD
19          128.15    P2p
Po1
```

```
Switch-2#sho spanning-tree vlan 4

<output omitted for brevity>
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
----- Fa0/4 Desg FWD					
3019	128.4	P2p	Edge Fa0/16	Root	FWD
3019	128.16	P2p	Po1	Desg	FWD
3039	128.65	P2p			

Switch-3#**sho spanning-tree vlan 4**

<output omitted for brevity>

Interface	Role	Sts	Cost	Prio.Nbr	Type
----- Fa0/13 Root FWD					
19	128.13	P2p	Fa0/16	Desg	FWD
19	128.16	P2p			

If your loop is up for any length of time, you will also start seeing messages that resemble this: %SW_MATM-4-MACFLAP_NOTIF: Host xxxx.xxxx.xxxx in vlan x is flapping...

Task 3: LoopGuard

- Implement the LoopGuard feature to prevent this type of STP Loop from forming:
 - Remove the BPDUFilter feature from the Port-Channel interface of Switch-1.
 - Wait for Spanning-Tree to reconverge by Blocking the EtherChannel on Switch-2.
 - Configure the LoopGuard feature on interface Port-Channel 1 of Switch-2.
 - Re-configure BPDUFilter on interface Port-Channel 1 of Switch-1.

This time, after Switch-2 stops receiving BPDUs on its Port-Channel (after Max-Age expires), its EtherChannel should NOT go into the Forwarding state (which caused the loop prior to implementation of LoopGuard) but instead should remain in the Blocking state.

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-1(config)#int port-channel 1
Switch-1(config-if)#spanning-tree bpdufilter disable

Switch-1(config-if)#end
```

Switch-2 Configuration

```
Switch-2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Switch-2(config)#int port-channel 1Switch-2(config-if)#spanning-tree guard loop

Switch-2(config-if)#
```

Switch-2 LoopGuard Verification

```
Switch-2#sho spanning-tree int port-channel 1 detail | i (VLAN|Loop)
Port 65 (Port-channel1) of VLAN0001 is alternate blocking Loop guard is enabled on the port
Port 65 (Port-channel1) of VLAN0002 is alternate blocking Loop guard is enabled on the port
Port 65 (Port-channel1) of VLAN0003 is designated forwarding Loop guard is enabled on the port
Port 65 (Port-channel1) of VLAN0004 is alternate blocking Loop guard is enabled on the port
Port 65 (Port-channel1) of VLAN0005 is alternate blocking Loop guard is enabled on the port

Switch-2#
```

Watching LoopGuard in Action

```
Switch-1#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#int port-channel 1
Switch-1(config-if)#spanning-tree bpdupfilter enable

Switch-1(config-if)#end
Switch-1#
Switch-2#
*Mar  4 01:12:27.633: %SPANTREE-2-LOOPGUARD_BLOCK: Loop guard blocking port Port-channel1 on VLAN0001.

Switch-2#
Switch-2#
```

- Disable the BPDUFILTER feature on the Port-Channel interface of Switch-1.

```
Switch-1#
conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch-1(config)#int port-channel 1Switch-1(config-if)#spanning-tree bpdupfilter disable

Switch-1(config-if)#end
Switch-1#
```

Task 4: RootGuard

- Implement the RootGuard feature to provide our STP topology with a secured perimeter, beyond which no unauthorized switch will be able to take over the position of Root Bridge.
 - Configure the Spanning-Tree Root Guard feature on interface FastEthernet0/4 of Switch-2.
 - Move to **Router-4** and disable interface FastEthernet0/1.
 - While still within the configuration of Router-4, enable Spanning-Tree on this router and give it a Bridge-Priority of 4096.
 - Enable interface FastEthernet0/1 on Router-4. As soon as it transmits a Superior BPDU to Switch-2, the Root Guard feature on Switch-2 should place port 0/4 into the Root Inconsistent state.

Switch-2 Root Guard Configuration

```
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch-2(config)#int fast 0/4Switch-2(config-if)#spanning-tree guard root
Switch-2(config-if)#end
Switch-2# *Mar  4 01:28:35.881:
%SPANTREE-2-ROOTGUARD_CONFIG_CHANGE: Root guard enabled on port FastEthernet0/4.

Switch-2#
```

Root Guard Verification

```
Rtr-4#
conf t
Enter configuration commands, one per line. End with CNTL/Z.
Rtr-4(config)#int fast 0/1
Rtr-4(config-if)#shut
Rtr-4(config-if)#exitRtr-4(config)#bridge 1 protocol ieee
Rtr-4(config)#bridge 1 priority 4096
Rtr-4(config)#int fast 0/1
Rtr-4(config-if)#bridge-group 1

Rtr-4(config-if)#
Rtr-4(config-if)#
Rtr-4(config-if)#no shutdown
```

```
Switch-2#
*Mar  4 01:33:35.201: %SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port FastEthernet0/4 on VLAN0004.


Switch-2#
*Mar  4 01:33:35.789: %LINK-3-UPDOWN: Interface FastEthernet0/4, changed state to up
*Mar  4 01:33:36.789: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/4, changed state to up
Switch-2#
```

- Remove all bridging configuration from Router-4.

CCNP SWITCH Workbook - 802.1w RSTP and 802.1s MST

4.1 Comparing RSTP & 802.1d Convergence

Load the *CCNP-Switch-Task4-1* initial configurations before starting.



Tasks

In this task, you will enable 802.1w Rapid Spanning-Tree on two of your switches and compare the behavior of this protocol against another switch running 802.1d PVST.

Objective-1: Preliminary Configuration

- Configure Switch-1 and Switch-3 for Rapid PVST mode.
- Ensure that Switch-1 is the STP Root Bridge for VLANs 1-5.

Switch-1 Configuration

```
Switch-1#config t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#
spanning-tree mode rapid-pvst
Switch-1(config)#spanning-tree vlan 1 root primary
Switch-1(config)#spanning-tree vlan 2 root primary
Switch-1(config)#spanning-tree vlan 3 root primary
Switch-1(config)#spanning-tree vlan 4 root primary
```

```
Switch-1(config)#spanning-tree vlan 5 root primary
```

```
Switch-1(config)#end
```

```
Switch-1#
```

Switch-3 Configuration

```
Switch-3#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z. Switch-3(config)#
```

```
spanning-tree mode rapid-pvst
```

```
Switch-3(config)#end
```

```
Switch-3#
```

Spanning-Tree Verification (Switch-1)

```
Switch-1#sho spanning-tree bridge protocol
```

```
VLAN0001 rstp
```

```
VLAN0002 rstp
```

```
VLAN0003 rstp
```

```
VLAN0004 rstp
```

```
VLAN0005 rstp
```

```
Switch-1#
```

```
Switch-1#sho spanning-tree root
```

Vlan	Root ID	Root Cost	Hello Time	Max Age	Fwd Dly	Root Port
VLAN0001	24577 0019.2f45.ec000					
2 20 15	VLAN0002	24578	0019.2f45.ec000			
2 20 15	VLAN0003	24579	0019.2f45.ec000			
2 20 15	VLAN0004	24580	0019.2f45.ec000			
2 20 15	VLAN0005	24581	0019.2f45.ec000			
2 20 15						

```
Switch-1#
```

A value of zero (0) under the Root Cost column indicates that this local switch is the Root Bridge.

Spanning-Tree Verification (Switch-2)

```
Switch-2#show spanning-tree bridge protocol
VLAN0001 ieee
VLAN0002 ieee
VLAN0003 ieee
VLAN0004 ieee
VLAN0005 ieee

Switch-2#
```

"ieee" in this context refers to 802.1d.

Spanning-Tree Verification (Switch-3)

```
Switch-3#show spanning-tree bridge protocolVLAN0001 rstp
VLAN0002 rstp
VLAN0003 rstp
VLAN0004 rstp
VLAN0005 rstp

Switch-3#
```

Objective-2: Comparing 802.1d and 802.1w Convergence

- Issue the command `show spanning-tree vlan 1` and notice which neighboring switches are running 802.1w (Rapid PVST) with Switch-1, and which neighboring switches are running 802.1d STP with Switch-1.
- Disable ports 0/13 and 0/10 on Switch-1.
- Turn on the IOS debug, `debug spanning-tree switch state`.
- Enable port 0/10 (leading to your 802.1d neighbor) and take note of the total time taken to move to the Forwarding state on this link.
- Enable port 0/13 (leading to your 802.1w RSTP neighbor) and take note of the total time taken to move to the Forwarding state on this link.

Switch-1 Verification

```
Switch-1#sho spanning-tree vlan 1
```

```
VLAN0001 Spanning tree enabled protocol rstp
```

```
Root ID    Priority    24577
           Address    0019.2f45.ec00
           This bridge is the root
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    24577 (priority 24576 sys-id-ext 1)
           Address    0019.2f45.ec00
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
           Aging Time 300 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/10	Desg	FWD	19	128.12	P2p Peer(STP)
Fa0/13	Desg	FWD	19	128.15	P2p
Fa0/14	Desg	FWD	19	128.16	P2p

In the output above, notice that Switch-1 is running 802.1w RSTP and has noticed that a neighboring switch on port 0/10 is in 802.1d Spanning-Tree mode.

```
Switch-1#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#
```

```
interface range fast 0/10 , fast 0/13
```

```
Switch-1(config-if-range)#shutdown
```

```
*Mar 1 00:37:30.034: %LINK-5-CHANGED: Interface FastEthernet0/10, changed state to administratively down
```

```
*Mar 1 00:37:30.076: %LINK-5-CHANGED: Interface FastEthernet0/13, changed state to administratively down
```

```
*Mar 1 00:37:31.040: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/10, changed state to down
```

```
Switch-1(config-if-range)#end
```

```
Switch-1#
```

```
Switch-1#debug spanning-tree switch state
```

```
Spanning Tree Port state changes debugging is on
```

```
Switch-1#
```

```
Switch-1#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#int fast 0/10
```

```
Switch-1(config-if)#no shut
```

```
Switch-1(config-if)#
```

```
*Mar 1 00:42:54.262: %LINK-3-UPDOWN: Interface FastEthernet0/10, changed state to up
```

```
Switch-1(config-if)#*Mar 1 00:42:56.804: STP SW: Fa0/10 newblocking
```

```
req for 1 vlans
```

```
*Mar 1 00:42:57.810: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/10, changed state to up
```

```
Switch-1(config-if)#*Mar 1 00:43:11.811: STP SW: Fa0/10 newlearning
```

```
req for 1 vlans
```

```
Switch-1(config-if)#*Mar 1 00:43:26.818: STP SW: Fa0/10 newforwarding
```

```
req for 1 vlans
```

```
Switch-1(config-if)#
```

The above output demonstrates that it took roughly 30 seconds on a switch running 802.1w RSTP to transition a port to the Forwarding State when that port was connected to an 802.1d switch.

```
Switch-1(config-if)#
```

```
Switch-1(config-if)#interface fast0/13
```

```
Switch-1(config-if)#no shut
```

```
Switch-1(config-if)#^Z
```

```
Switch-1#
```

```
Switch-1#
```

```
*Mar 1 00:46:37.013: %LINK-3-UPDOWN: Interface FastEthernet0/13, changed state to up
```

```
Switch-1#*Mar 1 00:46:39.563: STP SW: Fa0/13 newblocking
```

```
req for 1 vlans*Mar 1 00:46:39.580: STP SW: Fa0/13 newforwarding
```

```
req for 1 vlans
```

```
*Mar 1 00:46:40.570: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/13, changed state to up
```

```
Switch-1#
```

The above output demonstrates that it took roughly 1 second on a switch running 802.1w RSTP to transition a port to the Forwarding State when that port was connected to another 802.1w RSTP switch.

4.2 802.1s MST Initial Configuration

IST = No VLANs
Instance-1 = VLANs 1-3
Instance-2 = VLANs 4-4094

MST Region "INE"

PVST

Sw1

Sw3

Sw2

0/13

0/14

0/10

0/16

VLAN-4

Fa0/3

Fa0/4

Fa0/1

4.4.4.3/24

4.4.4.4/24

3

4

In this task, you will implement some preliminary configurations to convert part of your switched topology over to 802.1s MST.

- ## Switch-1 Configuration

<https://t.me/learningnets>

```
Switch-1(config)#end
```

Switch-3 Configuration

```
Switch-3(config)#Switch-3(config)#spanning-tree mst configuration
Switch-3(config-mst)#name INE
Switch-3(config-mst)#instance 1 vlan 1-3
Switch-3(config-mst)#instance 2 vlan 4-4094
Switch-3(config-mst)#exit
Switch-3(config)#spanning-tree mode mst

Switch-3(config)#end
Switch-3#
```

Switch-1 MST Verification

```
Switch-1#show spanning-tree mst configuration
Name [INE]
Revision 0      Instances configured 3

Instance  Vlans mapped
-----
0 none
1 1-3
2 4-4094
-----

Switch-1#
```

```
Switch-1#show spanning-tree bridge protocol
MST0 mstp
MST1 mstp
MST2 mstp

Switch-1#
```

Switch-3 MST Verification

Switch-3#

sho spanning-tree mst configuration#b3Name [INE]

Revision 0 Instances configured 3

Instance Vlans mapped

-----0 none

1 1-3

2 4-4094

Switch-3#

Switch-3#sho spanning-tree bridge protocol

MST0 mstp

MST1 mstp


MST2 mstp

Switch-3#

CCNP SWITCH Workbook - 802.1w RSTP and 802.1s MST

4.3 802.1s MST Load Balancing

Load the *CCNP-Switch-Task4-3* initial configurations before starting.



Tasks

In this task, you will modify your existing 802.1s MST configuration to accomplish per-instance load balancing of traffic.

OBJECTIVE 1: Minor changes to existing configuration

- Look at the topology diagram above and notice the change of VLAN interface assignment on Switch-3 and the change of IP address on Router-3, as well as the SVIs configured on Switch-1 and Switch-2.
- The above-mentioned changes have not been configured yet; configure them yourself now.

Switch-1 Configuration and Verification

```
Switch-1#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#int vlan 4
*Mar 1 19:00:09.918: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan4, changed state to up
Switch-1(config-if)#ip address 4.4.4.1 255.255.255.0
```

```

Switch-1(config-if)#end
Switch-1#
*Mar  1 19:00:29.036: %SYS-5-CONFIG_I: Configured from console by console
Switch-1#Switch-1#show ip interface brief | i (Vlan4)
Vlan4                4.4.4.1                YES manual up        up

Switch-1#

```

Switch-3 Configuration and Verification

```

Switch-3#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-3(config)#interface fast0/3
Switch-3(config-if)#switchport access vlan 3
Switch-3(config-if)#end
Switch-3#
*Mar  1 19:02:37.471: %SYS-5-CONFIG_I: Configured from console by console Switch-3#
sho interface fast0/3 switchport | i Access
Access Mode VLAN: 3 (VLAN0003)

Switch-3#

```

Switch-2 Configuration and Verification

```

Switch-2#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-2(config)#interface vlan 3
Switch-2(config-if)#ip address 3.3.3.2 255.255.255.0
Switch-2(config-if)#no shutdown
Switch-2(config-if)#exit
Switch-2(config)#end
Switch-2#Switch-2#show ip int brief | i (Vlan3)
Vlan3                3.3.3.2                YES manual up        up

Switch-2#

```

Router-3 Configuration and Verification

```

Router-3#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Router-3(config)#int fast 0/1
Router-3(config-if)#ip add 3.3.3.3 255.255.255.0

```

```
Router-3(config-if)#end
```

```
Router-3#
```

OBJECTIVE 2: Configuration of 802.1s MST Load Balancing

- View the topology diagram again and notice the dashed arrows, which indicate the STP Forwarding paths for VLAN-3 and VLAN-4 traffic.
- Configure your switches to accomplish the load balancing paths shown in the diagram, using the following criteria:
 - Ensure that Switch-1 is the STP Root Bridge for the MST Instance containing VLAN-3.
 - Ensure that Switch-3 is the STP Root Bridge for the MST Instance containing VLAN-4.
 - All commands configured to manipulate VLAN-3 traffic should ONLY be configured on Switch-1.

MST Load Balancing for VLAN-3 (Configuration and Verification)

```
Switch-1#conf tSwitch-1(config)#spanning-tree mst 1 root primary
Switch-1(config)#interface FastEthernet0/14
    Switch-1(config-if)#spanning-tree mst 1 port-priority 16
Switch-1(config-if)#end
Switch-1#Switch-1#show spanning-tree mst 1

##### MST1 vlans mapped: 1-3
Bridge          address 0019.2f45.ec00  priority          24577 (24576 sysid 1)
Root            this switch for MST1
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/10	Desg	FWD	200000	128.12	P2p Bound(PVST)
Fa0/13	Desg	FWD	200000	128.15	P2p
Fa0/14	Desg	FWD	200000	16.16	P2p

Switch-1#

When manipulating 802.1s MST values, hopefully you remembered that you must manipulate values related to MST Instances, not individual VLANs. When VLAN-3 is

part of Instance-1, there is no way to manipulate the STP forwarding path ONLY for VLAN-3. You must manipulate the STP forwarding path for the instance that VLAN-3 is a part of.

In the configuration above, the command `spanning-tree mst 1 root primary` was used to ensure that Switch-1 would become the STP Root Bridge for Instance-1. As an alternative, you could have used the command `spanning-tree mst 1 priority <value>` and then selected some priority value less than 32768.

Because the criteria stated that any load-balancing configuration you implemented to affect Instance-1 could ONLY be done on Switch-1, you did not have the option of manipulating STP port-cost values. Instead, your only option was to decrease the STP port-priority of interface FastEthernet0/14 on Switch-1 so that this interface would be viewed as the preferred path by downstream Switch-3.

MST Load-Balancing Verification for Instance-1 (VLAN-3)

```
Switch-3#sho spanning-tree mst 1
##### MST1 vlans mapped: 1-3
Bridge      address 000e.830d.f680  priority      32769 (32768 sysid 1)
Root        address 0019.2f45.ec00  priority      24577 (24576 sysid 1)
            port Fa0/14      cost          200000    rem hops 19

Interface    Role Sts Cost      Prio.Nbr Type
-----
Fa0/3        Desg FWD 200000    128.3    P2p Edge Fa0/13 Altn BLK
200000      128.13 P2p Fa0/14 Root FWD
200000      128.14 P2p
Fa0/16       Desg FWD 200000    128.16    P2p Bound(PVST)

Switch-3#
```

```
Switch-2#sho spanning-tree vlan 3

VLAN0003
Spanning tree enabled protocol ieee
Root ID    Priority      32768
           Address      000e.830d.f680
           Cost        19
           Port        16 (FastEthernet0/16)
           Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```

Bridge ID Priority    32771 (priority 32768 sys-id-ext 3)
Address      000c.8581.a500
Hello Time   2 sec   Max Age 20 sec   Forward Delay 15 sec
Aging Time   300 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/10	Altn	BLK			
19	128.10	P2p	Fa0/16	Root	FWD
19	128.16	P2p			

MST Load-Balancing for VLAN-4 (Configuration and Verification)

```

Switch-3#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-3(config)#
spanning-tree mst 2 priority 8192
Switch-3(config)#end
Switch-3#
*Mar  1 19:28:11.943: %SYS-5-CONFIG_I: Configured from console by console
Switch-3#Switch-3#sho spanning-tree mst 2
##### MST2      vlans mapped:  4-4094
Bridge      address 000e.830d.f680  priority      8194  (8192 sysid 2)
Root        this switch for MST2

```

Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/13	Desg	FWD	200000	128.13	P2p
Fa0/14	Desg	FWD	200000	128.14	P2p
Fa0/16	Desg	FWD	200000	128.16	P2p Bound(PVST)

```
Switch-3#
```

In order to accomplish the STP Forwarding Path shown for VLAN-4, you only needed to ensure that Switch-3 was configured as the STP Root Bridge for Instance-2. In this example, we manually lowered the Bridge Priority to accomplish that objective. Once that was completed, Instance-2 (and VLAN-4) traffic naturally took the paths shown in the diagram without further configuration.

MST Load-Balancing Verification for Instance-2 (VLAN-4)

```
Switch-1#sho spanning-tree mst 2
```

```
##### MST2 vlans mapped: 4-4094
```

```
Bridge      address 0019.2f45.ec00 priority 32770 (32768 sysid 2)
Root        address 000e.830d.f680 priority 8194 (8192 sysid 2)
            port Fa0/13 cost 200000 rem hops 19
```

Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/10	Desg	FWD	200000	128.12	P2p Bound(PVST) Fa0/13 Root FWD
200000	128.15	P2p	Fa0/14	Altn BLK	
200000	128.16	P2p			

```
Switch-1#
```

```
Switch-2#sho spanning-tree vlan 4
```

```
VLAN0004
```

```
Spanning tree enabled protocol ieee
```

```
Root ID      Priority    32768
              Address     000e.830d.f680
              Cost        19
              Port        16 (FastEthernet0/16)
              Hello Time   2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID    Priority    32772 (priority 32768 sys-id-ext 4)
              Address     000c.8581.a500
              Hello Time   2 sec Max Age 20 sec Forward Delay 15 sec
              Aging Time   300 sec
```


Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/4	Desg	FWD	19	128.4	P2p Edge Fa0/10 Altn BLK
19	128.10	P2p	Fa0/16	Root FWD	
19	128.16	P2p			

CCNP SWITCH Workbook - VLAN Access-Lists

5.1 Initial VACL Configuration

Load the **CCNP-Switch-Task5-1** initial configurations before starting.



Tasks

In this task, you will gain exposure to the VLAN Access-Maps feature (often called VACLs) to restrict intra-VLAN traffic between hosts. Note that although the topology diagram shows four "hosts," you will actually be using routers as your hosts.

- Verify that all four hosts can ping and telnet to each other. All pre-configured Telnet passwords are **ine**.
- After performing the verification above, create, and apply, a VACL in **Switch-2** that meets the following criteria:
 - Allows Host-1 to initiate ICMP pings all devices except Host-3.

- Neither Host-2 or Host-3 should be able to initiate Telnet sessions to each other.
- All other traffic, not explicitly denied above, should be forwarded.

Switch-2 Configuration

```
<output omitted for brevity>
!
vlan access-map INE 10
  action drop
  match ip address 101
vlan access-map INE 20
  action forward
!
vlan filter INE vlan-list 1
!
<output omitted for brevity>
!
access-list 101 permit icmp host 1.1.1.1 host 1.1.1.3 echo
access-list 101 permit tcp host 1.1.1.2 host 1.1.1.3 eq telnet
access-list 101 permit tcp host 1.1.1.3 host 1.1.1.2 eq telnet
!
<output omitted for brevity>
```

When configuring VLAN Access-Maps (VACLs), there is almost always more than a single approach that can be taken. Configuration of access-lists can sometimes be more of an art than a science. In this example, a couple of things that were mentioned (**or inferred**) in the initial criteria needed to be noticed:

- The first criterion that referenced ICMP pings said that Host-1 should not be allowed to **initiate pings to Host-3**, but it NEVER mentioned that Host-3 could not initiate ICMP Pings to Host-1. So a correct VACL configuration should have allowed Host-3 to ping Host-1, but NOT vice-versa.
- You could have created a second access-list (ACL 102 perhaps) that permitted all traffic, and then associated that ACL with Sequence-20 of the VLAN Access-Map. However, recall that when a VLAN Access-Map has no "match" statement it will, by default, match everything. So an extra ACL permitting any/any would have been redundant.

VACL Verification

```
Switch-2#sho vlan access-map

Vlan access-map "INE" 10
  Match clauses: ip address: 101
  Action: drop
Vlan access-map "INE" 20 Match clauses:
  Action: forward
Switch-2#sho vlan filter
VLAN Map INE is filtering VLANs:
1

Switch-2#
```

```
Host-1#ping 1.1.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 msHost-1#ping 4.4.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 msHost-1#ping 1.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.3, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)

Host-1#
```

```
Host-3#ping 1.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Host-3#
```

```
Host-2#telnet 1.1.1.1
Trying 1.1.1.1 ... Open
User Access Verification
```

Password:

Host-1>exit

[Connection to 1.1.1.1 closed by foreign host] Host-2# **telnet 1.1.1.11**

Trying 1.1.1.11 ... Open

User Access Verification

Password:

Switch-1>exit

[Connection to 1.1.1.11 closed by foreign host] Host-2# **telnet 1.1.1.3**


Trying 1.1.1.3 ... **% Connection timed out; remote host not responding**

Host-2#

CCNP SWITCH Workbook - VLAN Access-Lists

5.2 Using ACL "Deny" statements within a VACL

Load the **CCNP-Switch-Task5-2** initial configurations before starting.



Tasks

In this task, you will gain exposure to how a "deny" keyword in an access-list affects VACL processing.

- Verify that all four hosts can ping and telnet to each other. All pre-configured Telnet passwords are **ine**.
- After performing the verification above, create, and apply, a VACL in **Switch-2** that meets the following criteria:
 - Host-1 and Host-2 should be allowed to telnet to each other.
 - All other IP traffic within VLAN-1, not explicitly allowed above, should be

dropped by the VACL.

- You may only configure Access-List 101 for use in this lab (the quantity of lines, or Access Control Entries (ACEs), within ACL 101 are up to you), but **only a SINGLE line (ACE) of that ACL is allowed to use the permit keyword**. All other lines must use the **deny** keyword.

Switch-2 Configuration

```
<output omitted for brevity>
!
vlan access-map INE 10
  action drop
  match ip address 101
vlan access-map INE 20
  action forward
!
vlan filter INE vlan-list 1
!
<output omitted for brevity>
! access-list 101 deny
  tcp host 1.1.1.1 host 1.1.1.2 eq telnet access-list 101 deny
  tcp host 1.1.1.1 eq telnet host 1.1.1.2 access-list 101 deny
  tcp host 1.1.1.2 host 1.1.1.1 eq telnet access-list 101 deny
  tcp host 1.1.1.2 eq telnet host 1.1.1.1 access-list 101 permit
ip any any
!
<output omitted for brevity>
```

Although the criteria specified in the objectives forced you to configure your access-list and VACL in a way that was rather inefficient, there was a reason for this. The objective of this lab was to reinforce the concept that if a packet, inspected by a VACL Sequence, matches a "deny" statement within a access-list, processing of the VACL is done for that sequence, and the packet must now be inspected against the next sequence of that VACL. In other words, a packet that matches a "deny" in the ACL is "**denied**" (prevented) **from taking the Action** specified in that VACL Sequence, and must now be re-processed by the next Sequence of that same VACL.

VACL Verification

```
Switch-2#show vlan access-map

Vlan access-map "INE" 10
  Match clauses:
    ip address: 101
  Action:
    drop
Vlan access-map "INE" 20
  Match clauses:
  Action:
    forward
Switch-2#Switch-2#show vlan filter

VLAN Map INE is filtering VLANs:
  1
Switch-2#
```

```
Host-1#telnet 1.1.1.2
Trying 1.1.1.2 ...Open

User Access Verification

Password:
Host-2>exit
```

```
Host-2#telnet 1.1.1.1
Trying 1.1.1.1 ...Open

User Access Verification

Password:
Host-1>exit
```

All other types of traffic between hosts would fail at this point.

CCNP SWITCH Workbook - VLAN Access-Lists

5.3 Troubleshooting an existing VACL

Load the **CCNP-Switch-Task5-3** initial configurations before starting.

Tasks


In this task, you will troubleshoot a misconfigured VACL.

A customer who is relatively new to configuration of VLAN access-maps has opened a trouble ticket with you. He has configured a VACL that was supposed to accomplish the following objectives, but it is not working correctly.

- Host-1 and Host-4 should not be allowed to telnet to each other.
- No hosts should be able to telnet to any IP address configured on Switch-1.
- Hosts in VLAN-1 should be able to ping Host-4, but not ping the SVI configured for VLAN-4 on Switch-1.
- Host-4 should be allowed to ping its own default gateway.
- All other traffic, not specifically denied above, should be allowed to be forwarded.

After exhaustive troubleshooting on his own, the customer is too frustrated to clearly identify the specific problems to you. Instead, he is insisting that you log in to his equipment and figure out for yourself what is not working correctly and report back with the problems that you have identified (**don't fix anything yet**).

So **do that now**: Log in to the various devices in the topology and try to determine which of the customer's stated objectives are NOT being met by the pre-configured VACL.



At first glance, only a single problem appears to be presenting itself. Currently, Host-4 is allowed to telnet to either of the IP addresses configured on Switch-1, but the second bullet states that NO hosts should be allowed to telnet to ANY IP address on Switch-1.

After sharing this problem with the customer, he is reassured that you actually know what you're talking about and have identified the same problem that he did. He now trusts you to attempt to fix this problem.

The customer has one additional demand: After each change you apply to his switch, he insists that you test all stated objectives again. If you find that any of your changes have created ANOTHER problem, he wants to know about it immediately and will re-assess his troubleshooting contract with you.

Read no further until you have solved the first problem that was identified.

Fixing the First Problem

When initially configured, the VLAN access-map had only been applied against VLAN-1. By not having the VACL applied against VLAN-4, there was nothing to prevent Host-4 from telnetting to either of the IP addresses on Switch-1, and that was supposed to be denied. The fix to this problem was to **add VLAN-4** to the **vlan filter** command.

```
Switch-2#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-2(config)#
vlan filter INE vlan-list 1,4

Switch-2(config)#end
Switch-2#
```

```
Host-4#telnet 1.1.1.11
Trying 1.1.1.11 ... % Connection timed out; remote host not responding
Host-4#telnet 4.4.4.11
Trying 4.4.4.11 ... % Connection timed out; remote host not responding
```

So the fix above did solve the initial problem that was identified. However, it also created a new problem, and the customer is not happy.

You have noticed that Host-4 cannot ping its own default gateway of 4.4.4.11, although the criteria specifically stated that this traffic should be allowed. The customer slams down the telephone and walks away angry; however, after a few minutes he calls back and states that if you can fix the new problem you caused

without causing any additional problems, he will maintain his contract with you.

New Problem Created

```
Host-4#ping 4.4.4.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.11, timeout is 2 seconds:
.....Success rate is 0 percent (0/5)

Host-4#
```

Identifying the Cause of the New Problem

```
Switch-2#sho ip access-list
Extended IP access list 101
 10 permit tcp host 1.1.1.1 host 4.4.4.4 eq telnet
 20 permit tcp host 4.4.4.4 host 1.1.1.1 eq telnet
 30 permit tcp any host 1.1.1.11 eq telnet
 40 permit tcp any host 4.4.4.11 eq telnet 50 permit icmp any host 4.4.4.11 echo
```

As you can see, the last line of the access-list means that any packet, from any source, is **permitted to be dropped** by the first sequence of the VACL named, "INE". Because you just applied this VACL against VLAN-4, now packets from Host-4 are matching this line of the ACL and being dropped by the VACL.

Somehow you must modify the last line of this ACL without causing any other problems in the process.

Fixing the Problem without Creating Another Problem

```
Switch-2#conf tSwitch-2(config)#ip access-list extended 101
Switch-2(config-ext-nacl)#55 permit icmp 1.1.1.0 0.0.0.255 host 4.4.4.11 echo
Switch-2(config-ext-nacl)#do show ip access-list

Extended IP access list 101
 10 permit tcp host 1.1.1.1 host 4.4.4.4 eq telnet
 20 permit tcp host 4.4.4.4 host 1.1.1.1 eq telnet
 30 permit tcp any host 1.1.1.11 eq telnet
 40 permit tcp any host 4.4.4.11 eq telnet
```

```
50 permit icmp any host 4.4.4.11 echo 55 permit icmp 1.1.1.0 0.0.0.255 host 4.4.4.11 echo
```

First, treat this numbered access-list as if it were a named access-list, thereby giving you the ability to edit it. Adding the new sequence number of 55, in this case, will cause the VACL only to match on packets with a source address of 1.1.1.anything and drop those packets if they are going to the SVI for VLAN-4. This line will NOT match packets sourced from Host-4. However, Sequence-50 (which is dropping packets from Host-4) is still being processed BEFORE Sequence-55.

The last step is to remove Sequence-50 from this ACL, and it is for this reason that you decided to modify this ACL as if it were a named access-list. There is no way to remove a single line of an ACL if you treat it as a numbered ACL without first deleting the entire ACL and then re-copying it back to the device...which would create additional problems.

```
Switch-2(config-ext-nacl)#no 50
Switch-2(config-ext-nacl)#do show ip access-list

Extended IP access list 101
 10 permit tcp host 1.1.1.1 host 4.4.4.4 eq telnet
 20 permit tcp host 4.4.4.4 host 1.1.1.1 eq telnet
 30 permit tcp any host 1.1.1.11 eq telnet
 40 permit tcp any host 4.4.4.11 eq telnet 55 permit icmp 1.1.1.0 0.0.0.255 host 4.4.4.11 echo

Switch-2(config-ext-nacl)#end
Switch-2#
```

Verification that Host-4 Can Ping Its Default Gateway


```
Host-4#
ping 4.4.4.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.11, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Host-4#
```

CCNP SWITCH Workbook - Private-VLANs

6.1 Community Private VLANs

Load the **CCNP-Switch-Task6-1** initial configurations before starting.



Tasks

In this task, you will configure and test Private VLANs using Community VLANs.

You are an Internet Service Provider who also offers the additional service of hosting a customer's physical server(s) at your location. You have decided to implement a Private VLAN solution so that servers belonging to multiple customers can all share the same physical switch and share the same IP subnet, yet maintain privacy at the Datalink Layer.

In this lab, even though the term "**servers**" are mentioned, in reality **Router-1, Router-2, Switch-2, and Switch-3** will be your devices emulating servers. To initiate pings from one "server" to another, you will need to access the CLI of these devices.

- Before you configure anything on Switch-1, verify that all servers can ping each other. Your Private VLAN solution will change this behavior.
- Configure a Private VLAN solution on Switch-1, using the topology diagram as your guide:
 - Servers belonging to the same company should be able to ping each other.
 - Servers belonging to different companies should not be able to ping each other.
 - You may select any VLAN numbers you wish, keeping in mind the rules of Private VLANs.
 - At this point in time, no servers will have a default gateway.

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-1(config)#vtp mode transparent
Setting device to VTP Transparent mode for VLANs. Switch-1(config)#vlan 12
Switch-1(config-vlan)#private-vlan community
Switch-1(config-vlan)#exit Switch-1(config)#vlan 2233
Switch-1(config-vlan)#private-vlan community
Switch-1(config-vlan)#exit Switch-1(config)#vlan 111
Switch-1(config-vlan)#private-vlan primary
Switch-1(config-vlan)#private-vlan association 12,2233
Switch-1(config-vlan)#exit Switch-1(config)#int range fast 0/1 - 2
Switch-1(config-if-range)#switchport mode private-vlan host
Switch-1(config-if-range)#switchport private-vlan host-association 111 12
Switch-1(config-if-range)#exit
Switch-1(config)#Switch-1(config)#int range fast 0/10 , fast 0/15
Switch-1(config-if-range)#switchport mode private-vlan host
Switch-1(config-if-range)#switchport private-vlan host-association 111 2233

Switch-1(config-if-range)#end
Switch-1#
```

Verification

To begin, you needed to remember that to configure Private VLANs, a switch must first be in VTP Transparent mode. Also, no matter which VLAN numbers you selected, you needed to adhere to these guidelines to meet the objectives of this lab:

- Create two Community VLANs, because you are hosting servers from two different companies.
- Create a single Primary VLAN, because all of those servers had to be in the same IP Subnet.
- Associate both of your Community VLANs to your Primary VLAN.
- Access the physical interfaces on Switch-1 connected to your "servers" and configure them as follows:
 - Configure as Private-VLAN Host Ports (**switchport mode private-vlan host**)
 - Inform each interface about the Private VLAN pair it was associated with (**switchport private-vlan host-association**)

The first step in verification is to ensure that:

- Switch-1 recognizes that the VLANs you created are not "normal" VLANs but either Primary or Community Private-VLANs.
- Switch-1 recognizes that the physical interfaces connected to your servers are not "Access" switchports, but rather private-vlan host ports, and in the correct Private-VLAN assignments.

Switch-1 Configuration Verification

```
Switch-1#show vlan private-vlan
```

Primary	Secondary	Type	Ports
111	12	community	Fa0/1, Fa0/2
111	2233	community	Fa0/10, Fa0/15

```
Switch-1#
```

Verification that Customer-A's Devices Are Segregated from Customer-B's Devices

```
Cust-A-Server-1>
```

```
Cust-A-Server-1>ping 1.1.1.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
```

```
.!!!! Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

```
Cust-A-Server-1>Cust-A-Server-1>ping 1.1.1.22
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.22, timeout is 2 seconds:
```

```
..... Success rate is 0 percent (0/5)
```

```
Cust-A-Server-1>ping 1.1.1.33
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.33, timeout is 2 seconds:
```

```
..... Success rate is 0 percent (0/5)
```

```
Cust-A-Server-1>
```

```
Cust-B-Server-1>ping 1.1.1.33
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.33, timeout is 2 seconds:
```

```
.!!!! Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
```

```
Cust-B-Server-1>Cust-B-Server-1>ping 1.1.1.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
```

```
..... Success rate is 0 percent (0/5)
```

```
Cust-B-Server-1>ping 1.1.1.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
```


```
..... Success rate is 0 percent (0/5)
```

```
Cust-B-Server-1>
```

CCNP SWITCH Workbook - Private-VLANs

6.2 Isolated Private LANs

Load the **CCNP-Switch-Task6-2** initial configurations before starting.



Tasks

In this task, you will configure and test Private VLANs using a combination of Community and Isolated Secondary VLANs.

You are an Internet Service Provider who also offers the additional service of hosting a customer's physical server(s) at your location. You have decided to implement a Private VLAN solution so that servers belonging to multiple customers can all share the same physical switch and the same IP subnet, yet maintain privacy at the Datalink Layer.

In this lab, even though the term "**servers**" is mentioned, in reality **Router-1, Router-2, Switch-2, and Switch-3** will be your devices emulating servers. To initiate pings from one "server" to another, you will need to access the CLI of these devices.

If you are working on this lab immediately following completion of the previous task, note in the topology diagram that Switch-3 is now a server owned by **Customer-C**.

- Before you configure anything on Switch-1, verify that all servers can ping each other. Your Private VLAN solution will change this behavior.
- Configure a Private VLAN solution on Switch-1, using the topology diagram as your guide:
 - Servers belonging to Customer-A should be able to ping each other.
 - Customer-B and Customer-C each own a single server. Place both of these servers into a **single** Secondary VLAN so that they may neither ping each other nor the servers owned by Customer-A.
 - You may select any VLAN numbers you wish, keeping in mind the rules of Private VLANs.
 - At this point in time, no servers will have a default gateway.

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-1(config)#vtp mode transparent
Setting device to VTP Transparent mode for VLANs. Switch-1(config)#vlan 12
Switch-1(config-vlan)#private-vlan community
Switch-1(config-vlan)#exit Switch-1(config)#vlan 2233
Switch-1(config-vlan)#private-vlan isolated
Switch-1(config-vlan)#exit Switch-1(config)#vlan 111
Switch-1(config-vlan)#private-vlan primary
Switch-1(config-vlan)#private-vlan association 12,2233
Switch-1(config-vlan)#exit Switch-1(config)#int range fast 0/1 - 2
Switch-1(config-if-range)#switchport mode private-vlan host
Switch-1(config-if-range)#switchport private-vlan host-association 111 12
Switch-1(config-if-range)#exit Switch-1(config)#int range fast 0/10 , fast 0/15
Switch-1(config-if-range)#switchport mode private-vlan host
Switch-1(config-if-range)#switchport private-vlan host-association 111 2233

Switch-1(config-if-range)#end
Switch-1#
```

Verification

To begin, you needed to remember that to configure Private VLANs, a switch must first be in VTP Transparent mode. Also, no matter which VLAN numbers you selected, you needed to adhere to these guidelines to meet the objectives of this lab:

- Create two Secondary VLANs:
 - A Community PVLAN for the servers owned by Customer-A
 - An Isolated PVLAN for the servers owned by Customer-B and Customer-C
- Create a single Primary VLAN, because all of those servers needed to be in the same IP Subnet.
- Associate both of your Secondary VLANs to your Primary VLAN.
- Access the physical interfaces on Switch-1 connected to your "servers" and configure them as follows:
 - Configure as Private-VLAN Host Ports (**switchport mode private-vlan host**)
 - Inform each interface about the Private VLAN pair it was associated with (**switchport private-vlan host-association**)

The first step in verification is to ensure that:

- Switch-1 recognizes that the VLANs you created are not "normal" VLANs but either Primary, Community, or Isolated Private-VLANs.
- Switch-1 recognizes that the physical interfaces connected to your servers are not "Access" switchports, but rather private-vlan host ports, and in the correct Private-VLAN assignments.

Switch-1 Configuration Verification

```
Switch-1#show vlan private-vlan
```

Primary	Secondary	Type	Ports
111	12	community	Fa0/1, Fa0/2
111	2233	isolated	Fa0/10, Fa0/15

```
Switch-1#
```

Verification that Customer-A's Devices Are Segregated

from Customer-B's Devices

Cust-A-Server-1>

Cust-A-Server-1>ping 1.1.1.2

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:

.!!!! Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

Cust-A-Server-1>Cust-A-Server-1>ping 1.1.1.22

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.22, timeout is 2 seconds:

..... Success rate is 0 percent (0/5)

Cust-A-Server-1>ping 1.1.1.33

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.33, timeout is 2 seconds:

..... Success rate is 0 percent (0/5)


Cust-A-Server-1>

At this point, because they are in an Isolated Private VLAN and have no default gateway, the servers owned by Customer-B and Customer-C should not be able to ping anything.

CCNP SWITCH Workbook - Private-VLANs

6.3 PVLAN Promiscuous Ports (SVIs)

Load the **CCNP-Switch-Task6-3** initial configurations before starting.



Tasks

In this task, you will configure and test Private VLAN Promiscuous Ports.

The topology shown in the diagram, including all Private-VLANs, has been pre-configured.

In this lab, even though the term "**servers**" is mentioned, in reality **Router-1, Router-2, Switch-2, and Switch-3** will be your devices emulating servers. To initiate pings from one "server" to another, you'll need to access the CLI of these devices.

- Before you configure anything on Switch-1, verify that your Private VLANs exist and are assigned to the correct ports.
- Currently, all servers have a default-gateway configured of **1.1.1.111**, but they cannot reach this IP address nor can they reach any other IP addresses outside of their own subnet.
- Configure interface VLAN 111 on Switch-1 as the Private-VLAN Promiscuous Port such that:

- All servers should be able to ping the IP address of their default-gateway (1.1.1.111).
- The Network Administrator's PC (Router-3) should be able to ping any of the servers.

Switch-1 Configuration

```
Switch-1#conf tSwitch-1(config)#interface vlan 111
Switch-1(config-if)#private-vlan mapping 12,2233
Switch-1(config-if)#end
Switch-1#*Mar  1 02:38:31.355: %PV-6-PV_MSG: Created a private vlan mapping, Primary 111, Secondary 12
*Mar  1 02:38:31.355: %PV-6-PV_MSG: Created a private vlan mapping, Primary 111, Secondary 2233

Switch-1#
```

Switch-1 Configuration Verification

```
Switch-1#sho int vlan 111 private-vlan mapping
Interface Secondary VLANs
-----vlan111  12, 2233

Switch-1#
```

PVLAN Outside Reachability Verification

```
Cust-A-Server-1>ping 1.1.1.111
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Cust-A-Server-1>
```

```
Cust-A-Server-2>ping 1.1.1.111
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:
```

```
!!!!!!Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
Cust-A-Server-2>
```

```
Cust-B-Server-1>ping 1.1.1.111
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:
```

```
!!!!!!Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
Cust-B-Server-1>
```

```
Cust-C-Server-1>ping 1.1.1.111
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:
```

```
!!!!!!Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
Cust-C-Server-1>
```

Admin-PC>ping 1.1.1.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:

!!!!!!**Success rate is 100 percent**

(5/5), round-trip min/avg/max = 1/2/4 ms**Admin-PC>ping 1.1.1.2**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:

!!!!!!**Success rate is 100 percent**

(5/5), round-trip min/avg/max = 1/3/8 ms**Admin-PC>ping 1.1.1.22**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.22, timeout is 2 seconds:

!!!!!!**Success rate is 100 percent**

(5/5), round-trip min/avg/max = 1/2/4 ms**Admin-PC>ping 1.1.1.33**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.33, timeout is 2 seconds:

!!!!!!**Success rate is 100 percent**


(5/5), round-trip min/avg/max = 1/2/4 ms

Admin-PC>

CCNP SWITCH Workbook - Private-VLANs

6.4 PVLAN Promiscuous Ports (L2 Interfaces)

Load the **CCNP-Switch-Task6-4** initial configurations before starting.



Tasks

In this task, you will configure a physical switchport as a Private VLAN Promiscuous Port.

The topology shown in the diagram, including all Private-VLANs, has been pre-configured.

In this lab, even though the term "**servers**" is mentioned, in reality **Router-1, Router-2, Switch-2, and Switch-3** will be your devices emulating servers. To initiate pings from one "server" to another, you will need to access the CLI of these devices.

- Before you configure anything on Switch-1, verify that your Private VLANs exist and are assigned to the correct ports.
- Currently, all servers have a default-gateway configured of **1.1.1.111**, but they are unable to reach this IP address, nor can they reach any other IP addresses outside of their own subnet.
- Configure interface FastEthernet0/14 on Switch-1 as the Private-VLAN Promiscuous Port so that:

- All servers should be able to ping the IP address of their default-gateway (1.1.1.111).
- All servers should be able to ping the IP address of 33.33.33.3.

Switch-1 Configuration

```
Switch-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Switch-1(config)#int fast0/14
Switch-1(config-if)#switchport mode private-vlan promiscuous
Switch-1(config-if)#switchport private-vlan mapping 111 12
Switch-1(config-if)#switchport private-vlan mapping 111 2233

Switch-1(config-if)#end
Switch-1#
```

Switch-1 Configuration Verification

```
Switch-1#show interface fast0/14 private-vlan mapping
Private vlan mapping information is not available for FastEthernet0/14
```

Notice that the command above, which is useful for verifying an SVI configured as a PVLAN Promiscuous Port, does not help when verifying a physical interface as the Promiscuous Port.

```
Switch-1#show interface fast0/14 switchport
Name: Fa0/14
Switchport: EnabledAdministrative Mode: private-vlan promiscuous
Operational Mode: private-vlan promiscuous
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: 111 (VLAN0111) 12 (VLAN0012) 2233 (VLAN2233)

Administrative private-vlan trunk native VLAN: none
```

```
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk mappings: none
Operational private-vlan:
111 (VLAN0111) 12 (VLAN0012) 2233 (VLAN2233)
```

PVLAN Outside Reachability Verification

```
Cust-A-Server-1>ping 1.1.1.111
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/2/4 ms
Cust-A-Server-1>ping 33.33.33.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.33.33.3, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/1/1 ms
Cust-A-Server-1>
```

```
Cust-A-Server-2>ping 1.1.1.111
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/2/4 ms
Cust-A-Server-2>ping 33.33.33.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.33.33.3, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/1/1 ms
Cust-A-Server-2>
```

```
Cust-B-Server-1>ping 1.1.1.111
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/2/4 ms
Cust-B-Server-1>ping 33.33.33.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.33.33.3, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/1/1 ms
```

Cust-B-Server-1>

Cust-C-Server-1>ping 1.1.1.111

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.1.1.111, timeout is 2 seconds:

!!!!!!**Success rate is 100 percent (5/5)**

, round-trip min/avg/max = 1/2/4 ms Cust-C-Server-1>**ping 33.33.33.3**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 33.33.33.3, timeout is 2 seconds:

!!!!!!**Success rate is 100 percent (5/5)**


, round-trip min/avg/max = 1/1/1 ms

Cust-C-Server-1>

CCNP SWITCH Workbook - Switching Security Features

7.1 The Need for DHCP Snooping

Load the **CCNP-Switch-Task7-1** initial configurations before starting.



Task

Most of the topology shown in the diagram, including the Corporate DHCP Server, VLANs, Routing, and IP addresses, have been pre-configured. You will configure the Rogue DHCP Server in this task.

In this lab, even though the terms "**servers**" and "**clients**" are mentioned, in reality, your routers are acting as these devices.

- Before you configure anything, ensure that the Corporate DHCP Server is reachable and functional by:

- Disabling interface FastEthernet0/1 on DHCP Client-R2 and DHCP Client-R2.
- Enabling these interfaces.
- Waiting and watching for IP interface assignment via DHCP.

Corporate DHCP Verification

Client-R2>

en

Client-R2#conf t

Enter configuration commands, one per line. End with CNTL/Z. Client-R2(config)#**int fast 0/1**

Client-R2(config-if)#**shutdown**

Client-R2(config-if)#

*Nov 6 06:57:45.839: %LINK-5-CHANGED: Interface FastEthernet0/1, changed state to administratively down

*Nov 6 06:57:46.839: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down

Client-R2(config-if)#**no shutdown**

Client-R2(config-if)#end

Client-R2#

Client-R2#

*Nov 6 06:57:52.511: %SYS-5-CONFIG_I: Configured from console by console

Client-R2#

*Nov 6 06:57:56.991: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up

Client-R2#

*Nov 6 06:58:02.155: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 24.24.24.3, mask 255.255.255.0

Client-R2#

Client-R4>

en

Client-R4#conf t

Enter configuration commands, one per line. End with CNTL/Z. Client-R4(config)#**int fast0/1**

Client-R4(config-if)#**shutdown**

Client-R4(config-if)#

Nov 6 07:05:11.002: %LINK-5-CHANGED: Interface FastEthernet0/1, changed state to administratively down

Nov 6 07:05:12.002: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down

Client-R4(config-if)#**no shutdown**

Client-R4(config-if)#end

Client-R4#

Nov 6 07:05:19.966: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up

Client-R4#

Client-R4#

```
Nov  6 07:05:20.070: %SYS-5-CONFIG_I: Configured from console by console
Nov  6 07:05:20.966: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
Client-R4#
Nov  6 07:05:29.122: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 24.24.24.4, mask 255.255.255.0
Client-R4#
```

Task

- Currently, any broadcasts sent by either Client-R2 or Client-R4 are flooded by Switch-2 and visible to the other clients as well as Router-3 (Rogue DHCP Server). This means that when (as an example) Client-R2 transmits a broadcast DHCP Discover packet, and broadcast DHCP Request packet, these are visible to other devices within VLAN-24.
- Verify this for yourself by doing the following:
 - Ensure that you have two Telnet windows open at the same time: a window to R2, and another window to R4.
 - On Client-R4, enter the command `debug ip udp port 67`. This command will display any IP packets that Client-R4 has received and forwarded to its own CPU, with a source, or destination, UDP port 67 (the DHCP Server port).
 - On Client-R2, **shutdown** interface **FastEthernet0/1** and then enable this interface. The debug running on Client-R4 will prove that it has received the DHCP broadcasts sent to and from Client-R2.
 - Turn off all debugging on Client-R4 with the command `undebug all` (or simply `u all`).

```
Client-R4>
enClient-R4#debug ip udp port 67
UDP packet debugging is on
Client-R4#
```

```
Client-R2(config)#int fast 0/1
Client-R2(config-if)#shutdown
Client-R2(config-if)#no shutdown
```

```
Client-R4
# Nov  6 07:23:41.881: UDP: rcvd src=0.0.0.0(68), dst=255.255.255.255(67)
, length=320
Client-R4# Nov  6 07:23:43.885: UDP: rcvd src=24.24.24.22(67), dst=255.255.255.255(68)
, length=308 Nov  6 07:23:43.885: UDP: rcvd src=0.0.0.0(68), dst=255.255.255.255(67)
, length=332 Nov  6 07:23:43.889: UDP: rcvd src=24.24.24.22(67), dst=255.255.255.255(68)
, length=308
Client-R4#
...
Client-R4#undebg all
All possible debugging has been turned off
```

Task

- Complete the configuration of R3 as a Rogue DHCP Server by doing the following:
 - Enable the command `service dhcp`.
 - Create a **DHCP Pool** (using the name **INE**).
 - Your pool should provide IP addresses within the correct subnet of **24.24.24.0/24**.
 - Your pool should be configured to **intentionally offer an INCORRECT IP address** of **24.24.24.33** (the Rogue DHCP Server) as the Default-Router for all DHCP clients.
 - Your pool should be configured with a DHCP **Lease** time of **7 days**.

Rogue DHCP Server Configuration

```
Rogue-Server-R3>en
Rogue-Server-R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z. Rogue-Server-R3(config)#service dhcp
Rogue-Server-R3(config)#ip dhcp pool INE
Rogue-Server-R3(dhcp-config)#network 24.24.24.0 /24
Rogue-Server-R3(dhcp-config)#default-router 24.24.24.33
Rogue-Server-R3(dhcp-config)#lease 7
```

```
Rogue-Server-R3(dhcp-config)#end
Rogue-Server-R3#
```

Ensure that the Rogue DHCP is reachable by pinging it from Switch-2.

```
Switch-2#ping 24.24.24.33

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24.24.24.33, timeout is 2 seconds:
.!!!!Success rate is 80 percent (4/5)
, round-trip min/avg/max = 1/2/4 ms
Switch-2#
```

Typically, when DHCP clients receive more than a single offer in response to their DHCP Discover packet, they will accept **the very first offer they received**.

In this case, because the Rogue DHCP Server is located in the same VLAN as the DHCP clients, any DHCP offer sent from this device should be accepted before any legitimate offer sent from the Corporate DHCP Server. This will result in DHCP clients being given an IP address in the correct subnet, but their default-gateway assignment will be wrong. Any packets they attempt to send off-subnet will be sent to R3 (the Rogue DHCP Server) because the clients believe that device is their Default-Gateway.

To verify this:

- Disable interface FastEthernet0/1 on Client-R2.
- Re-enable this interface.
- Watch for this interface to be assigned a DHCP Address.
- To prove that this address came from the Rogue DHCP Server, issue the command `show ip route` and notice that you now have a default route to 24.24.24.33. This route SHOULD have pointed to 24.24.24.22 (your DHCP Relay Agent/Default-Gateway).

If the above steps do not work, and you still have a default-route pointing at your legitimate Default-Gateway of 24.24.24.22, disable FastEthernet0/1 on Client-R2, wait at least 30 seconds, and then enable this interface again.

Verification

Client-R2

```
(config)#int fast 0/1Client-R2(config-if)#shutdown
Client-R2(config-if)#Client-R2(config-if)#no shutdown
Client-R2(config-if)#
Client-R2(config-if)#end
Client-R2#
```

```
*Nov  6 07:38:43.343: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 24.24.24.2, mask 255.255.255.0
```

Client-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
       a - application route
       + - replicated route, % - next hop override
```

```
Gateway of last resort is 24.24.24.33 to network 0.0.0.0
```

```
S* 0.0.0.0/0 [254/0] via 24.24.24.33
```

```
24.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
```

```
C 24.24.24.0/24 is directly connected, FastEthernet0/1
```

```
L 24.24.24.2/32 is directly connected, FastEthernet0/1
```

```
Client-R2#
```

Ping an off-subnet address like 1.1.1.11. Does your ping succeed? Why or why not?

Client-R2#ping 1.1.1.11

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.11, timeout is 2 seconds:
```

```
!!!!!!Success rate is 100 percent (5/5)
```

```
, round-trip min/avg/max = 1/1/4 ms
```

```
Client-R2#
```

You can see in the above example that pinging an off-subnet address DID work; but how was that possible if Client-R2 **now believes the Rogue DHCP Server (R3) is the default-gateway?**

Here is the explanation:

- Any IP packets that Client-R2 sends to an off-subnet address are (at Layer 2) being addressed to the destination MAC address of its default-gateway (the Rogue DHCP Server in this case).
- When those Ethernet frames arrive on R3 (Rogue DHCP Server), R3 believes those frames are meant for itself so the Ethernet header is removed, revealing the IP packet inside.
- Our Rogue DHCP Server (R3 in this example) has been pre-configured for IP routing. So when it inspects the packet it received (originally sent by Client-R2) and realizes that the packet is NOT meant for itself, **it routes that packet to its OWN default-gateway of 24.24.24.22.**

In this manner, packets sent off-subnet by the DHCP clients are actually getting routed, but they are getting routed by the wrong device (Rogue DHCP Server). This can be viewed by doing the following steps:

- In the Rogue DHCP Server (R3), enable the command `debug ip packet detail`.
- Issue another ICMP **Ping from Client-R2** to any off-subnet address (1.1.1.11 is used in the example below).
- Look at the debug output captured by R3. This proves that these ICMP pings sent by Client-R2 were actually sent to, and routed by, the Rogue DHCP Server.

Verification

```
Rogue-Server-R3>
enRogue-Server-R3#debug ip packet detail
IP packet debugging is on (detailed)

Rogue-Server-R3#
```

```
Client-R2#ping 1.1.1.11
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.11, timeout is 2 seconds:
```

```
!!!!!!Success rate is 100 percent (5/5)
```

```
, round-trip min/avg/max = 1/2/4 ms
```

```
Client-R2#
```

```
Rogue-Server-R3#Nov  6 08:00:28.017: IP: s=24.24.24.2 (FastEthernet0/0), d=1.1.1.11
, len 100, input featureNov  6 08:00:28.017:      ICMP type=8, code=0, MCI Check(94), rtype 0,
forus FALSE
, sendself FALSE, mtu 0, fwdchk FALSENov  6 08:00:28.017: FIBipv4-packet-proc:
route packet from FastEthernet0/0 src 24.24.24.2 dst 1.1.1.11
Nov  6 08:00:28.017: FIBfwd-proc: packet routed by adj to FastEthernet0/0 24.24.24.22
Nov  6 08:00:28.021: FIBipv4-packet-proc: packet routing succeededNov  6 08:00:28.021: IP:
s=24.24.24.2 (FastEthernet0/0), d=1.1.1.11 (FastEthernet0/0), len 100, redirected

Nov  6 08:00:28.021:      ICMP type=8, code=0
...
<output omitted for brevity>
...
Rogue-Server-R3#
```


In summary, this task has demonstrated the need for DHCP snooping by showing how:

- All DHCP broadcast-based packets are seen by all devices within a VLAN.
- The above observation allows placement of a Rogue DHCP Server within a VLAN that can offer DHCP clients incorrect information.

CCNP SWITCH Workbook - Switching Security Features

7.2 DHCP Snooping Configuration

Load the **CCNP-Switch-Task7-2** initial configurations before starting.



Task

All of the topology shown in the diagram, including the Corporate DHCP Server, Rogue DHCP Server, VLANs, routing, and IP addresses, have been pre-configured. You will configure the DHCP Snooping feature in this task.

In this lab, even though the terms "**servers**" and "**clients**" are mentioned, in reality, your routers are acting as these devices.

- Configure DHCP Snooping on Switch-2.
- Verify that DHCP Snooping is functional by:
 - **1:** Verify that the DHCP Snooping feature has been enabled correctly by viewing the output of the command `show ip dhcp snooping`.
 - **2:** View the DHCP Snooping Binding Database and look for entries created by DHCP Clients R2 and R4.
 - **3:** Verify that all DHCP clients received their DHCP information from the Corporate DHCP Server.
 - **4:** Verify that the Rogue DHCP Server never had any visibility to DHCP broadcasts generated by the clients.
 - **5:** Verify that DHCP clients did not have visibility to each other's DHCP broadcasts.

Switch-2 Configuration

```
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#ip dhcp snooping
Switch-2(config)#ip dhcp snooping vlan 24
```

Verification 1

```
Switch-2#show ip dhcp snooping
```

```
Switch DHCP snooping is enabled
```

```
DHCP snooping is configured on following VLANs: 24
```

```
DHCP snooping is operational on following VLANs: 24
```

DHCP snooping is configured on the following L3 Interfaces:

Insertion of option 82 is enabled

circuit-id default format: vlan-mod-port

remote-id: 000c.8581.a500 (MAC)

Option 82 on untrusted port is not allowed

Verification of hwaddr field is enabled

Verification of giaddr field is enabled

DHCP snooping trust/rate is configured on the following Interfaces:

Interface	Trusted	Allow option	Rate limit (pps)
-----	-----	-----	-----
Switch-2#			

Verification 2

```
Client-R2#
```

```
conf tClient-R2(config)#int fast 0/1
```

```
Client-R2(config-if)#shutdown
```

```
Client-R2(config-if)#no shutdown
```

```
Client-R2(config-if)#end
```

```
Client-R2#
```

```
Client-R2#
```

```
*Nov 6 09:08:09.903: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 24.24.24.9, mask 255.255.255.0
```

Client-R4

```
#conf t Client-R4(config)#int fast 0/1
```

```
Client-R4(config-if)#shutdown
```

```
Client-R4(config-if)#no shutdown
```

```
Client-R4(config-if)#end
```

```
Client-R4#
```

```
Client-R4#
```

```
Nov 6 09:14:51.205: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 24.24.24.4, mask 255.255.255.0
```

```
Switch-2#show ip dhcp snooping binding
```

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
00:1A:6C:30:8F:DF	24.24.24.9	604611	dhcp-snooping	24	FastEthernet0/1
00:1C:58:9E:7A:E1	24.24.24.4	604659	dhcp-snooping	24	FastEthernet0/4

```
Total number of bindings: 2
```

```
Switch-2#
```

Verification 3

```
DHCP-Server-R1#show ip dhcp pool
```

```
Pool INE :
```

```
Utilization mark (high/low) : 100 / 0
```

```
Subnet size (first/next) : 0 / 0
```

```
Total addresses : 254 Leased addresses : 2
```

```
Pending event : none
```

```
1 subnet is currently in the pool : Current index IP address range Leased addresses
```

```
24.24.24.10 24.24.24.1 - 24.24.24.254 2
```

```
DHCP-Server-R1#
```

```
DHCP-Server-R1#show ip dhcp binding
```

```
Bindings from all pools not associated with VRF:
```

IP address	Client-ID/ Hardware address/ User name	Lease expiration	Type
------------	--	------------------	------

24.24.24.4	0063.6973.636f.2d30.	Nov 13 2014 05:33 PM	Automatic
------------	----------------------	----------------------	-----------

3031.632e.3538.3965.

2e37.6165.312d.4661.

302f.31

24.24.24.9 0063.6973.636f.2d30. Nov 13 2014 05:33 PM Automatic

3031.612e.3663.3330.

2e38.6664.662d.4661.

302f.31

DHCP-Server-R1#

Client-R2# show dhcp lease

Temp IP addr: 24.24.24.9 for peer on Interface: FastEthernet0/1

Temp sub net mask: 255.255.255.0 DHCP Lease server: 1.1.1.1

, state: 5 Bound

DHCP transaction id: 1FB

Lease: 604800 secs, Renewal: 302400 secs, Rebind: 529200 secs

Temp default-gateway addr: 24.24.24.22

Next timer fires after: 3d11h

Retry count: 0 Client-ID: cisco-001a.6c30.8fdf-Fa0/1

Client-ID hex dump: 636973636F2D303031612E366333302E

386664662D4661302F31

Hostname: Client-R2

Client-R2#

Client-R4# show dhcp lease

Temp IP addr: 24.24.24.4 for peer on Interface: FastEthernet0/1

Temp sub net mask: 255.255.255.0 DHCP Lease server: 1.1.1.1

, state: 5 Bound

DHCP transaction id: 265B

Lease: 604800 secs, Renewal: 302400 secs, Rebind: 529200 secs

Temp default-gateway addr: 24.24.24.22

Next timer fires after: 3d11h

Retry count: 0 Client-ID: cisco-001c.589e.7ae1-Fa0/1

Client-ID hex dump: 636973636F2D303031632E353839652E

376165312D4661302F31

Hostname: Client-R4

Client-R4#

Verification 4

```
Rogue-Server-R3>
```

```
en
```

```
Rogue-Server-R3#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z. Rogue-Server-R3(config)#
```

```
logging buffer debug
```

```
Rogue-Server-R3(config)#end
```

```
Rogue-Server-R3#
```

```
Nov 6 09:25:29.333: %SYS-5-CONFIG_I: Configured from console by console Rogue-Server-R3#
```

```
debug ip udp port 67
```

```
UDP packet debugging is on
```

```
Rogue-Server-R3#clear log
```

```
Clear logging buffer [confirm]
```

```
Rogue-Server-R3#
```

Above, we are enabling debug output to be logged to a memory buffer so that we can come back to this device and view it at a later time. Then we are enabling the command `debug ip udp port 67` so that we can view any DHCP packets (UDP Port 67 is the DHCP Server port) that this device had visibility to.

In the previous task, we saw how flooded DHCP packets were captured by this debug. Now with DHCP Snooping enabled on the switch, those packets will **NOT** be flooded and this debug on the Rogue DHCP Server (R3) should pick up...nothing.

```
Client-R2#
```

```
conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z. Client-R2(config)#
```

```
Client-R2(config-if)#shutdown
```

```
Client-R2(config-if)#Client-R2(config-if)#no shutdown
```

```
Client-R2(config-if)#end
```

```
Client-R2#
```

```
Client-R2#
```

```
*Nov 6 09:23:18.031: %DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 24.24.24.9, mask 255.255.255.0
```

```
Rogue-Server-R3#show log
```

```
Syslog logging: enabled (0 messages dropped, 3 messages rate-limited, 0 flushes, 0 overruns, xml disabled, filtering disabled)
```

```
No Active Message Discriminator.
```

```
No Inactive Message Discriminator.
```

```
Console logging: level debugging, 105 messages logged, xml disabled,  
                  filtering disabled
```

```
Monitor logging: level debugging, 0 messages logged, xml disabled,  
                  filtering disabled
```

```
Buffer logging:  level debugging, 1 messages logged, xml disabled,  
                  filtering disabled
```

```
Exception Logging: size (4096 bytes)
```

```
Count and timestamp logging messages: disabled
```

```
Persistent logging: disabled
```

```
No active filter modules.
```

```
Trap logging: level informational, 48 message lines logged
```

```
Logging Source-Interface:      VRF Name:
```

```
Log Buffer (4096 bytes):
```

```
Rogue-Server-R3#
```

The debug we enabled has not captured any UDP Port-67 packets, as expected.

```
Rogue-Server-R3#undebug all
```

```
All possible debugging has been turned off
```

```
Rogue-Server-R3#
```


Verification 5

Just follow the exact same steps as in Verification 4, but this time enable `debug ip udp port 67` on DHCP Client-R4. Just like the previous verification step, R4 should receive no output from this debug.

Task

Notice that in the previous steps, you enabled DHCP Snooping on a switch (Switch-2) that was configured as a DHCP Relay Agent. In other words, when Switch-2 receives DHCP broadcasts on its VLAN-24 Switched Virtual Interface, it encapsulates those and routes them toward the Corporate DHCP Server as unicasts.

Now, change the configuration of Switch-1 and Switch-2 so they match the topology diagram below (notice that Switch-1 is now configured as the DHCP Relay Agent, and the link between Switch-1 and Switch-2 is a Layer-2 Access Switchport. Do NOT change any of your existing DHCP Snooping configuration on Switch-2.



Switch-2 Configuration

```
Switch-2#  
conf t  
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#int fast 0/10  
Switch-2(config-if)#switchport  
Switch-2(config-if)#switchport mode access  
Switch-2(config-if)#switchport access vlan 24  
Switch-2(config-if)#exit Switch-2(config)#no interface vlan 24  
  
Switch-2(config)#end  
Switch-2#
```

Switch-1 Configuration

```
Switch-1#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#interface vlan 24
Switch-1(config-if)#ip address 24.24.24.22 255.255.255.0
Switch-1(config-if)#ip helper-address 1.1.1.1
Switch-1(config-if)#no shut
Switch-1(config-if)#exit Switch-1(config)#vlan 24
Switch-1(config-vlan)#exit Switch-1(config)#interface fast 0/10
Switch-1(config-if)#switchport
Switch-1(config-if)#switchport access vlan 24
Switch-1(config-if)#switchport mode access

Switch-1(config-if)#end
Switch-1#
```

Task

Test to see if DHCP Snooping is still functional on Switch-2. If not, do you know WHY?

Verification

```
Client-R2#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Client-R2(config)#int fast 0/1
Client-R2(config-if)#shutdown
Client-R2(config-if)#Client-R2(config-if)#no shutdown
Client-R2(config-if)#end
Client-R2#
*Nov  6 10:03:17.267: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up

Client-R2#
```

You should have noticed on both Client-R2 and Client-R4 that they are no longer able to obtain DHCP information from any DHCP Server.

Task

- Fix DHCP Snooping on Switch-2 so that both DHCP clients can, once again, obtain DHCP information from the Corporate DHCP Server.

Switch-2 Configuration

```
Switch-2#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#
no ip dhcp snooping information option
Switch-2(config)#interface FastEthernet0/10
Switch-2(config-if)#ip dhcp snooping trust

Switch-2(config-if)#end
Switch-2#
```

DHCP Client Verification

```
Client-R2#show ip int br
```

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	unassigned	YES	NVRAM	administratively down	down
FastEthernet0/1	24.24.24.9	YES	DHCP		
up	up				

What did you learn about DHCP Snooping after changing the topology?


1. When a switch is operating as a Layer-3 switch and a DHCP Relay Agent, the Layer-3 SVI that receives inbound DHCP broadcasts from clients is, by default, trusted by DHCP Snooping. So there is no need to configure any interfaces (physical or virtual) as DHCP Snooping Trusted interfaces.
2. When a switch is operating as a Layer-2 switch, and all DHCP clients as well as any ports that lead to trusted DHCP servers are Layer-2 Switchports, then the physical interfaces leading upstream to trusted DHCP servers must be configured with the `ip dhcp snooping trust` command.
3. When a switch is operating as a Layer-2 switch, it will, by default, insert the **DHCP Option-82** into any DHCP client packets it receives. However, most DHCP servers (including Cisco routers and switches configured as DHCP servers) cannot recognize

Option-82 and will drop any DHCP client packets that contain this option. To fix this, we needed to configure the DHCP snooping switch to NOT insert Option-82 into DHCP Client packets by using the command `no ip dhcp snooping information option` .

CCNP SWITCH Workbook - Switching Security Features

7.3 Demonstrating the Need for Dynamic ARP Inspection

Load the **CCNP-Switch-Task7-3** initial configurations before starting.



All of the topology shown in the diagram, including the Corporate DHCP Server, Malicious User, VLANs, routing, and IP addresses, have been pre-configured.

In this lab, even though the terms "**servers**" and "**clients**" are

mentioned, in reality, your routers are acting as these devices.

Task

In this first task, you'll learn why the Dynamic ARP Inspection feature can be useful within a network.

- Within either R2 or R4, ping the default-gateway of 24.24.24.22.
- After the successful ping, within the router where you just issued the ping, issue the command `show ip arp 24.24.24.22` and write down somewhere the MAC address you see associated with that IP address.
- On the Malicious User (R3):
 - Issue the command `debug arp` .
 - **Change the IP address** of interface FastEthernet0/0 to be the same as the default-gateway for VLAN-24 (**24.24.24.22**).
 - Watch the gratuitous ARP response that the Malicious User sends.
 - Disable all debugs on the Malicious User (R3) with the command `undebug all` .
- On Client-R2, issue the command `show ip arp 24.24.24.22` several times repeatedly over the course of a minute or so; what do you notice?
- On Client-R2, **perform an extended ping of 1.1.1.1 using a count of 2,000** and note the results.

Obtain ARP Information for default-gateway

```
Client-R2#ping 24.24.24.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24.24.24.22, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/2/4 ms
Client-R2#sho ip arp 24.24.24.22
Protocol Address          Age (min)  Hardware Addr   Type   InterfaceInternet
24.24.24.22          1         0019.2f45.ec41
ARPA    FastEthernet0/1
Client-R2#
```

Malicious User Configuration

```
Malicious-User#debug arp
```

```
ARP packet debugging is on
```

```
Malicious-User#
```

```
conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z. Malicious-User(config)#int fast 0/0
```

```
Malicious-User(config-if)#ip add 24.24.24.22 255.255.255.0
```

```
Malicious-User(config-if)#end
```

```
Malicious-User#
```

Malicious User Debug Verification

Malicious-User#

```
Nov  6 13:21:00.971: IP ARP: sent rep src 24.24.24.22 0018.b9ba.6dd8,  
dst 24.24.24.22 ffff.ffff.ffff  
FastEthernet0/0
```

```
Nov  6 13:21:00.971: IP ARP: sent rep src 24.24.24.22 0018.b9ba.6dd8,  
dst 24.24.24.22 ffff.ffff.ffff FastEthernet0/0
```

Malicious-User#

Malicious-User#

```
Nov  6 13:21:01.559: %SYS-5-CONFIG_I: Configured from console by console
```

Malicious-User#

```
Nov  6 13:21:05.487: IP ARP: sent rep src 24.24.24.22 0018.b9ba.6dd8,  
dst 24.24.24.22 ffff.ffff.ffff FastEthernet0/0 Nov  6 13:21:05.495:
```

```
IP ARP: rcvd rep src 24.24.24.22 0019.2f45.ec41
```

```
, dst 24.24.24.22 FastEthernet0/0 Nov  6 13:21:05.495:
```

```
%IP-4-DUPADDR: Duplicate address 24.24.24.22 on FastEthernet0/0, sourced by 0019.2f45.ec41
```

Malicious-User#

Malicious-User#

```
Nov  6 13:21:05.495: IP ARP: Gratuitous ARP throttled.
```

```
Nov  6 13:21:05.495: IP ARP: 24.24.24.22 added to arp_defense_Q
```

Malicious-User#

```
Nov  6 13:21:06.487: IP ARP: 24.24.24.22 removed from arp_defense_Q
```

```
Nov  6 13:21:06.487: IP ARP: sent rep src 24.24.24.22 0018.b9ba.6dd8,  
dst 24.24.24.22 0018.b9ba.6dd8 FastEthernet0/0
```

Cisco routers and switches automatically send a gratuitous ARP response either the moment you change an IP address on an interface or after that interface comes up from previously being disabled.

In the output above, we can see that when you duplicated Switch-1's IP address (by configuring it on FastEthernet0/0 of R3, the Malicious User), both R3 and Switch-1 got into an ARP war, whereby they continually sent gratuitous ARP responses for their IP addresses. Both devices realized (through the reception of the other's ARP response) that there was a duplicate IP addressing conflict; however, this realization did nothing to stop them from perpetually sending gratuitous ARP responses.

So what do we see on the clients?

Client ARP Table

```
Client-R2#sho ip arp 24.24.24.22
```

```

Protocol Address      Age (min)  Hardware Addr  Type   Interface Internet
24.24.24.22      0  0019.2f45.ec41
  ARPA   FastEthernet0/1
Client-R2#sho ip arp 24.24.24.22
Protocol Address      Age (min)  Hardware Addr  Type   Interface Internet
24.24.24.22      0  0018.b9ba.6dd8
  ARPA   FastEthernet0/1
Client-R2#sho ip arp 24.24.24.22
Protocol Address      Age (min)  Hardware Addr  Type   Interface Internet
24.24.24.22      0  0019.2f45.ec41
  ARPA   FastEthernet0/1

```

On the clients, we see that the ARP table is now being "poisoned" by the Malicious User. In this case, because both the Malicious User and the device that legitimately owns the address of 24.24.24.22 are constantly sending gratuitous ARPs, the client's ARP cache will keep bouncing back and forth between both MAC addresses.

What effect does it have on the client when, temporarily at least, it has the wrong MAC information for its default-gateway?

```

Client-R2#sho ip route
<output omitted for brevity>

Gateway of last resort is 24.24.24.22 to network 0.0.0.0
S* 0.0.0.0/0 [254/0] via 24.24.24.22

1.0.0.0/32 is subnetted, 1 subnets
S    1.1.1.1 [254/0] via 24.24.24.22, FastEthernet0/1
24.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C    24.24.24.0/24 is directly connected, FastEthernet0/1
L    24.24.24.1/32 is directly connected, FastEthernet0/1
Client-R2#

```

The static route above was dynamically created when Client-R2 obtained a DHCP IP address, and along with it learned of the default-gateway of 24.24.24.22. But because the MAC address learned for that same gateway is fluctuating between Switch-1 and the Malicious User, this is what happens when we actually try to USE that default-gateway for packet routing:

```


Client-R2#ping 1.1.1.1 repeat 2000
Type escape sequence to abort.
Sending 2000, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:

```


CCNP SWITCH Workbook - Switching Security Features

7.4 Dynamic ARP Inspection Configuration

Load the **CCNP-Switch-Task7-4** initial configurations before starting.



All of the topology shown in the diagram, including the Corporate DHCP Server, Malicious User, VLANs, routing, and IP addresses, have been pre-configured.

In this lab, even though the terms "**servers**" and "**clients**" are mentioned, in reality, your routers are acting as these devices.

Task

In the previous task, you learned why the Dynamic ARP Inspection feature can be useful within a network. Now you'll configure it and watch it solve that problem.

- Configure the Dynamic ARP Inspection feature on Switch-2.
- Ensure that Client R2 can now ping the DHCP Server at 1.1.1.1 with a count of 1000 packets and a 100% success rate.
- Verify that the feature is working by:
 - 1: View the output of the command `show ip arp inspection`.
 - 2: View the output of `show ip arp 24.24.24.22` in Client R2, and ensure that the MAC address is the correct MAC for SVI VLAN-24 on Switch-2 (the correct default-gateway).
 - 3: **Issue a ping from Client-R2 to the address 1.1.1.1**, with a repeat count of 1000, and ensure that you have 100% success.

Switch-2 Configuration

```
Switch-2#  
conf t  
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#ip arp inspection vlan 24  
Switch-2(config)#int fast Switch-2(config)#int fast 0/10  
Switch-2(config-if)#ip arp inspect trust  
  
Switch-2(config-if)#end
```

Verification 1: Verification of Dynamic ARP Inspection on Switch-2

```
Switch-2#show ip arp inspection
```

```
Source Mac Validation      : Disabled  
Destination Mac Validation : Disabled  
IP Address Validation     : Disabled
```

```
Vlan      Configuration      Operation
ACL Match      Static ACL
-----
----- 24 Enabled Active
...
<output omitted for brevity>
```

Verification 2

```
Switch-1#show int vlan 24
Vlan24 is up, line protocol is up   Hardware is EtherSVI, address is 0019.2f45.ec41
(bia 0019.2f45.ec41) Internet address is 24.24.24.22/24
```

```
Client-R2#sho ip arp 24.24.24.22
```

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
	24.24.24.22	5	0019.2f45.ec41		FastEthernet0/1

```
Client-R2#
```

Verification 3

[illegible]

Why did this work? Remember that Dynamic ARP Inspection works alongside DHCP Snooping. DHCP Snooping (which was pre-configured for you in this lab) builds a DHCP Snooping Binding Table from the DHCP Client-Server transactions it observes.

When an ARP request or an ARP reply arrives on a switch configured for Dynamic ARP Inspection, DAI looks into that DHCP Snooping Binding Table to see if it can verify that the device transmitting the ARP is actually valid.

In this case, the Malicious User had a static IP address, so whenever the Malicious User transmitted a spoofed Gratuitous ARP reply, Switch-2 (running DAI) would intercept it and attempt to verify the validity of that user. However, DAI within Switch-2 was unable to find a corresponding DHCP Snooping Binding entry for that IP address/MAC Address learned on port 0/3. Because DAI couldn't verify the validity of the sender of that Gratuitous ARP Reply (from the Malicious User), the switch dropped the packet and would log the following message:


```
*Mar  1 01:13:28.719: %SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Res) on Fa0/3, vlan 24
. ([0018.b9ba.6dd8/24.24.24.22/0018.b9ba.6dd8/24.24.24.22/01:13:28 UTC Mon Mar  1 1993])
```

The legitimate owner of 24.24.24.22 (Switch-1) also was not participating in DHCP. So we had to use the command `ip arp inspection trust` on Switch-2's interface leading toward Switch-1 (FastEthernet0/10).

CCNP SWITCH Workbook - Switching Security Features

7.5 IP Source Guard

Load the **CCNP-Switch-Task7-5** initial configurations before starting.



All of the topology shown in the diagram, including the DHCP Snooping feature on Switch-2, have been pre-configured.

In this lab, even though the terms "**servers**" and "**clients**" are mentioned, in reality, your routers are acting as these devices.

Task

In this task, you'll see how the IP Source Guard feature can prevent IP spoofing attacks.

First, ensure that your current topology is up and functional (**you may have to enable some interfaces**).

- Verify that both DHCP clients (R2 and R3) have obtained IP addresses and a default-gateway via DHCP.
- Verify that DHCP Snooping is configured and active on Switch-2, VLAN-24, and that this switch has some entries in the DHCP Snooping Binding Table for the two DHCP clients.
- Verify that both DHCP clients can ping the DHCP server (at 1.1.1.1) as well as the Corporate Email Server at 24.24.24.100.

Initial Topology Verification (DHCP Assignment on Clients)

```
Client-R2#sho ip int br
```

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	unassigned	YES	unset	administratively down	down
FastEthernet0/1	24.24.24.1	YES	DHCP	up	up

```
Client-R2#sho ip route
```

```
<output omitted for brevity>
```

```
Gateway of last resort is 24.24.24.22 to network 0.0.0.0
```

```
S* 0.0.0.0/0 [254/0] via 24.24.24.22
```

```
1.0.0.0/32 is subnetted, 1 subnets
```

```
S      1.1.1.1 [254/0] via 24.24.24.22, FastEthernet0/1
```

Malicious-User#sho ip int br

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	24.24.24.2	YES	DHCP		
up	up				

Malicious-User#sho ip route

<output omitted for brevity>

Gateway of last resort is 24.24.24.22 to network 0.0.0.0

S* 0.0.0.0/0 [254/0] via 24.24.24.22

1.0.0.0/32 is subnetted, 1 subnets

S 1.1.1.1 [254/0] via 24.24.24.22, FastEthernet0/0

Initial Topology Verification (DHCP Snooping on Switch-2)

Switch-2#show ip dhcp snooping

Switch DHCP snooping is enabled

DHCP snooping is configured on following VLANs:

24

DHCP snooping is operational on following VLANs:

24

<output omitted for brevity>

DHCP snooping trust/rate is configured on the following Interfaces:

Interface	Trusted	Allow option	Rate limit (pps)
-----	-----	-----	-----
FastEthernet0/4	yes	unlimited	yes
Custom circuit-ids: FastEthernet0/10	yes	unlimited	

Switch-2#show ip dhcp snooping binding

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
-----	-----	-----	-----	-----	-----
00:1A:6C:30:8F:DF	24.24.24.1	604476	dhcp-snooping	24	FastEthernet0/1
00:18:B9:BA:6D:D8	24.24.24.2	604506	dhcp-snooping	24	FastEthernet0/3

Initial Topology Verification (ICMP Connectivity from DHCP Clients)

```
Client-R2#
Client-R2#ping 1.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/2/4 msClient-R2#ping 24.24.24.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24.24.24.100, timeout is 2 seconds:
.!!!!!!Success rate is 80 percent (4/5)
, round-trip min/avg/max = 1/1/1 ms
Client-R2#
```

```
Malicious-User#ping 1.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/2/4 msMalicious-User#ping 24.24.24.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24.24.24.100, timeout is 2 seconds:
.!!!!!!Success rate is 80 percent (4/5)
, round-trip min/avg/max = 1/1/4 ms
Malicious-User#
```

Task: Demonstrating the Need for IP Source Guard

In this topology, the Malicious User has decided that he wants to initiate a Denial of Service (DoS) attack on the Corporate DHCP Server. Knowing that servers must process received ICMP pings via their CPUs, his intent is to bring down this server by flooding it with thousands of pings (called a Ping Attack).

To prevent the attack from being traced back to his machine, the Malicious User decides to spoof the source IP address of his pings. He has no need to receive responses to his pings, he just wants his pings to reach the server.

- Enable the following commands on the DHCP server (Router-1):

- logging buffer debug
- debug ip icmp
- Initiate pings from the Malicious User to the DHCP server.
- View the debug output on the DHCP server. This verifies that the pings were successfully received by this server.

Initial Ping Tests

```
DHCP-Server-R1(config)#logging buffer debug
DHCP-Server-R1(config)#end DHCP-Server-R1#debug ip icmp
ICMP packet debugging is on

DHCP-Server-R1#
```

```
Malicious-User#ping 1.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!!Success rate is 100 percent (5/5)
, round-trip min/avg/max = 1/2/4 ms
Malicious-User#
```

```
DHCP-Server-R1#show log
<output omitted for brevity>

Log Buffer (4096 bytes):

Nov  7 12:38:56.423: %SYS-5-CONFIG_I: Configured from console by consoleNov  7 12:39:19.039: ICMP:
echo reply sent, src 1.1.1.1, dst 24.24.24.2
, topology BASE, dscp 0 topoid 0Nov  7 12:39:19.039: ICMP:echo reply sent, src 1.1.1.1, dst 24.24.24.2
, topology BASE, dscp 0 topoid 0Nov  7 12:39:19.043: ICMP:echo reply sent, src 1.1.1.1, dst 24.24.24.2
, topology BASE, dscp 0 topoid 0Nov  7 12:39:19.043: ICMP:echo reply sent, src 1.1.1.1, dst 24.24.24.2
, topology BASE, dscp 0 topoid 0Nov  7 12:39:19.047: ICMP:echo reply sent, src 1.1.1.1, dst 24.24.24.2
, topology BASE, dscp 0 topoid 0
DHCP-Server-R1#
```

Initiation of Ping Denial of Service Attack

- On the Malicious User, **configure interface loopback0 and assign it the IP address of 3.3.3.3 /32.**

- The icmp debug should still be running on the DHCP server. If not, re-enable it.
- Initiate an extended ping from the Malicious User to the DHCP server using the following criteria:
 - **Use a repeat count of 50 pings.**
 - **Use a source address of 3.3.3.3.**
- View the debug output on the DHCP server. This verifies that, after using a "spoofed" IP address, the pings from the Malicious User still made their way to the DHCP server.

DoS Ping Attack Configuration and Verification

```
Malicious-User#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Malicious-User(config)#
interface loopback 0
Malicious-User(config-if)#ip add 3.3.3.3 255.255.255.255

Malicious-User(config-if)#end
Malicious-User#
```

```
Malicious-User#ping 1.1.1.1 repeat 50 source 3.3.3.3
Type escape sequence to abort.
Sending 50, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 3.3.3.3
```

```
DHCP-Server-R1#
Nov  7 12:45:45.235: ICMP: echo reply sent, src 1.1.1.1, dst 3.3.3.3
, topology BASE, dscp 0 topoid 0 Nov  7 12:45:45.239:
ICMP: dst (1.1.1.1) host unreachable rcv from 1.1.1.11

DHCP-Server-R1#
DHCP-Server-R1#
```

Task: Configuring and Verifying IP Source Guard

As a network administrator, you don't have the luxury of knowing in advance where hosts reside that might initiate a Spoofing attack on your network by changing their

source IP address. However, the IP Source Guard feature is one that is enabled on individual interfaces, so you would most likely enable it on all interfaces that connect to training rooms, offices, and office workspaces.

- On Switch-2, **enable the IP Source Guard feature** on the two interfaces connected to your DHCP clients (FastEthernet0/1 and FastEthernet0/3).
- Issue the command `show ip verify source` to confirm that this feature has been activated on the correct interfaces.

```
Switch-2#conf tSwitch-2(config)#interface range fast 0/1 , fast 0/3
Switch-2(config-if-range)#ip verify source

Switch-2(config-if-range)#end
Switch-2#
```

```
Switch-2#Switch-2#sho ip verify source
```

Interface	Filter-type	Filter-mode	IP-address	Mac-address	Vlan
Fa0/1	ip	active	24.24.24.1		
		24	Fa0/3	ip	active
		24			24.24.24.2

```
Switch-2#
```

- Once again, initiate the extended ping from the Malicious User, using the source IP address of 3.3.3.3.

The IP Source Guard feature works in hardware, so you will not see any syslogs or debug output to indicate that it is working.

To verify that the IP Source Guard feature successfully dropped the "spoofed" packets from the Malicious User:

- View the output of `debug ip icmp` on the DHCP server. You should see...nothing. This indicates that no ICMP packets were received.

```
Malicious-User#ping 1.1.1.1 repeat 50 source 3.3.3.3
```

```
Type escape sequence to abort.
```

```
Sending 50, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
```

```
Packet sent with a source address of 3.3.3.3
```

```
DHCP-Server-R1#show debug
```

```
Generic IP:
```

```
_ICMP packet debugging is on_
```

```
<no output is seen>
```

```
DHCP-Server-R1#
```

IP Source Guard on Non-DHCP Ports

Currently, interface FastEthernet0/4 on Switch-2 leads to the Corporate Email Server. This server has a static IP address, so this interface was configured as a DHCP snooping trusted interface.

```
Switch-2#sho ip dhcp snooping
```

```
<output omitted for brevity>
```

```
Interface Trusted
```

```
Allow option    Rate limit (pps)
```

```
-----  
FastEthernet0/4    yes    yes    unlimited
```


```
Custom circuit-ids: FastEthernet0/10    yes
```

```
yes    unlimited
```

```
Custom circuit-ids:
```

```
Switch-2#
```

Imagine that interface FastEthernet0/4 was connected to a hub. This hub is not only connected to the Corporate Email Server, but also to a training room as shown below. You probably would NOT want that interface to be a DHCP snooping trusted interface, and you probably WOULD want to enable IP Source Guard on this interface.



Do that now.

- On Switch-2, convert interface FastEthernet0/4 to a DHCP Snooping untrusted interface.
- Configure the IP Source Guard feature on this interface.

Switch-2 Configuration

```
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#int fast 0/4
Switch-2(config-if)#no ip dhcp snooping trust
Switch-2(config-if)#ip verify source

Switch-2(config-if)#end
Switch-2#
```

- Initiate a ping from the Corporate Email Server (R4) to the Corporate DHCP Server. The ping should fail. Do you understand why?

```
Corporate-Email-Server#ping 1.1.1.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
```

```
.....Success rate is 0 percent (0/5)
```

```
Corporate-Email-Server#
```

- Issue the command `show ip verify source` on Switch-2. Pay special attention to the status of port 0/4. Do you understand the output you are viewing?

```
Switch-2#sho ip verify source
```

Interface	Filter-type	Filter-mode	IP-address	Mac-address	Vlan
-----	-----	-----	-----	-----	----
Fa0/1	ip	active	24.24.24.1		24
Fa0/3	ip	active	24.24.24.2		24
Fa0/4	ip	active	deny-all		
		24			

```
Switch-2#
```

As mentioned earlier, IP Source Guard relies on being able to match the source IP address and incoming interface of IP packets against what it finds in the DHCP Snooping Binding Table. However, the Corporate Email Server is not participating in DHCP; it has a static address.

Because IP Source Guard cannot find any entry in the DHCP Snooping Binding Table for interface FastEthernet0/4, yet it knows that this interface is a DHCP snooping **untrusted** interface, it automatically places a dynamic ACL blocking all inbound IP packets received on this interface.

To fix this, we need to create a static IP entry that IP Source Guard can use to verify the validity of packets received from the Corporate Email Server.

- In Switch-2, configure a static IP entry so that IP Source Guard will no longer block packets received from the Corporate Email Server.

Switch-2, Static IP Entry Configuration

```
Corporate-Email-Server#show interface Fast0/1 | i bia
```

```
Hardware is MV96340 Ethernet, address is 001c.589e.7ae1 (bia 001c.589e.7ae1
```

```
)
```

```
Corporate-Email-Server#
```

```
Switch-2(config)#ip source binding 001C.589E.7AE1 vlan 24 24.24.24.100 interface Fa0/4
```

```
Switch-2#show ip verify source
```

Interface	Filter-type	Filter-mode	IP-address	Mac-address	Vlan
-----	-----	-----	-----	-----	----
Fa0/1	ip	active	24.24.24.1		24
Fa0/3	ip	active	24.24.24.2		24
Fa0/4	ip	active	24.24.24.100		24

```
Switch-2#
```

```
Corporate-Email-Server#ping 1.1.1.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
```

```
!!!!!!Success rate is 100 percent (5/5)
```


```
, round-trip min/avg/max = 1/2/4 ms
```

```
Corporate-Email-Server#
```

CCNP SWITCH Workbook - Switching Security Features

7.6 Storm-Control

Load the **CCNP-Switch-Task7-6** initial configurations before starting.



Task

In this task, you'll see how the Storm-Control feature can prevent the devastating effects of bridging loops.

- Configure the Storm-Control feature on Switch-2 following these criteria:
 - The feature should only be configured on interfaces FastEthernet0/10 - 12.
 - Storm-control should only be monitoring broadcast traffic.
 - Storm-control should be triggered when broadcasts exceed 0.50% of the interface bandwidth.
 - When Storm-control is triggered, it should continue to block traffic until broadcast traffic has diminished to no greater than 0.20% of the interface bandwidth.

Switch-2 Configuration

```
Switch-2#config t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#int range fast0/10 - 12
Switch-2(config-if-range)#storm-control broadcast level ? <0.00 - 100.00> Enter rising threshold
pps          Enter suppression level in packets per second

Switch-2(config-if-range)#storm-control broadcast level 0.5 ? <0.00 - 100.00> Enter falling threshold
<cr>

Switch-2(config-if-range)#storm-control broadcast level 0.5 0.2 ?
<cr>
Switch-2(config-if-range)#storm-control broadcast level 0.5 0.2

Switch-2(config-if-range)#end
Switch-2#
```

Storm-Control Configuration Verification

```
Switch-2#show storm-control
```

Interface	Filter	State	Upper	Lower	Current
Fa0/10		Forwarding	0.50%	0.20%	
	0.00%	Fa0/11	Forwarding	0.50%	0.20%
	0.00%	Fa0/12	Forwarding	0.50%	0.20%
	0.00%				

```
Switch-2#
```

Task

Right now, there are not many broadcasts going on in this topology. To make things interesting (and trigger Storm-Control), let's induce a broadcast storm.

To begin, we'll intentionally create a bridging loop in this topology (do not try this in your production network).

- **Ensure that Switch-1 is the Spanning-Tree Root Bridge for VLAN-1.**
- **Configure the Spanning-Tree BPDUFilter feature on Switch-2, interfaces FastEthernet0/10 - 12.**

The above configurations should cause all three interfaces (FastEthernet0/10 - 12) connecting Switch-1 to Switch-2 to be in Forwarding state on both sides of the links.

Inducing a Bridging Loop

```
Switch-1#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-1(config)#
spanning-tree vlan 1 priority 0
Switch-1(config)#end
Switch-1#
Switch-1#show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    1
             Address     0019.2f45.ec00
             This bridge is the root
             Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```

Bridge ID  Priority    1      (priority 0 sys-id-ext 1)
Address    0019.2f45.ec00
Hello Time  2 sec    Max Age 20 sec    Forward Delay 15 sec
Aging Time  15 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/10	Desg	FWD			
19	128.12	P2p	Fa0/11	Desg	FWD
19	128.13	P2p	Fa0/12	Desg	FWD
19	128.14	P2p			

Switch-2#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#**int range fast0/10 - 12**

Switch-2(config-if-range)#**spanning-tree bpdudfilter enable**

Switch-2(config-if-range)#end

Switch-2#

Switch-2#**show spanning-tree vlan 1**

VLAN0001

Spanning tree enabled protocol ieee

Root ID Priority 32769

Address 000c.8581.a500

This bridge is the root

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)

Address 000c.8581.a500

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Aging Time 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/4	Desg	FWD	19	128.4	P2p Edge Fa0/10 Desg FWD
19	128.10	P2p	Fa0/11	Desg	FWD
19	128.11	P2p	Fa0/12	Desg	FWD
19	128.12	P2p			

Task

Now that we have a broadcast loop between Switch-1 and Switch-2, any broadcasts

that are received by either of these switches (on VLAN-1) will be constantly replicated and forwarded around this loop. Any other ports, also in VLAN-1, will suffer from broadcasts as well.

Let's test this.

- View the running-configuration of Router-4. You will notice that the pre-config for this lab created a static ARP entry on this router for 1.1.1.2 with a broadcast MAC address.

```
Host#sh run
Building configuration...

Current configuration : 1653 bytes
!
<output omitted for brevity>
...
! interface FastEthernet0/1
ip address 1.1.1.4 255.255.255.0
duplex auto
speed auto
...
...
!
!
! arp 1.1.1.2 ffff.ffff.ffff ARPA
!
```

In theory, when you try to ping 1.1.1.2 from the host (Router-4), all pings to this address should be sent at Layer 2 to the broadcast destination MAC address (because of the static ARP entry). These broadcasts will be replicated infinitely by Switch-1 and Switch-2 as they are propagated across the bridging loop.

To view this, do the following:

- On Switch-2, **shutdown interfaces FastEthernet0/11 - 12.**
- While still in Switch-2, change the **load-interval on interface FastEthernet0/10** to the minimum value (30 seconds).
- **Clear the interface counters** of Switch-2.
- On Switch-2, issue `show interface fastethernet0/10 | i packets output .`

- Repeat the command above on Switch-2 three or four times and notice that right now, there are very few (if any) packets being transmitted from FastEthernet0/10.

```
Switch-2#conf t
Switch-2(config)#int range fast 0/11 - 12
Switch-2(config-if-range)#shutdown
Switch-2(config-if-range)#exit
Switch-2(config)#Switch-2(config)#int fast 0/10
Switch-2(config-if)#load-interval 30
Switch-2(config-if)#exit
Switch-2(config)#end
Switch-2#clear counters

Clear "show interface" counters on all interfaces [confirm]
Switch-2#
```

- **Initiate a ping from the Host to 1.1.1.2 with a repeat count of 1000.** You will not see any replies to these pings; this is simply to generate broadcast traffic on VLAN-1.
- Repeat the command on Switch-2, `show interface fastethernet0/10 | i packets output`, every couple of seconds.
- Notice that the continuous ping you are transmitting from the host has a default timeout of 2 seconds. So the output of the command above should be incrementing every couple of seconds.

```
Host#ping 1.1.1.2 repeat 1000
Type escape sequence to abort. Sending 1000, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds
:
.....
```

```
Switch-2#show int fast 0/10 | i packets output
35 packets output
, 4679 bytes, 0 underruns
Switch-2#show int fast 0/10 | i packets output 36 packets output
, 4827 bytes, 0 underruns
Switch-2#show int fast 0/10 | i packets output 37 packets output
, 4975 bytes, 0 underruns
Switch-2#show int fast 0/10 | i packets output 39 packets output
, 5217 bytes, 0 underruns
```

Now we will demonstrate how even a few broadcasts can quickly become a broadcast storm when there is a bridging loop in the topology.

- On Switch-2, enter the command `interface range fast0/11 - 12` and press Enter.
- While within interface range configuration mode, enter the command `do show interface fastethernet0/10 | i packets output`.
- Enable interfaces **FastEthernet0/11 - 12** with the `no shutdown` command. This will reinstate your bridging loop.
- While still within interface configuration mode, every second or two, repeat the command in your history buffer, `do show interface fastethernet0/10 | i packets output`.

The moment that Spanning-Tree places FastEthernet0/11 and 0/12 into the Forwarding state and your bridging loop is induced, you should notice that the quantity of packets output on Fast0/10 dramatically increases.

- Shut down port 0/11 and 0/12 as soon as you start seeing the "packets output" counter increasing out of control.

Bridging Loop Ramifications - Verification

```
Switch-2(config)#int range fast 0/11 - 12
Switch-2(config-if-range)#do show int fast 0/10 | i packets output
    369 packets output, 52282 bytes, 0 underrunsSwitch-2(config-if-range)#no shutdown
Switch-2(config-if-range)#do show int fast 0/10 | i packets output
    375 packets output, 53116 bytes, 0 underruns
Switch-2(config-if-range)#
*Mar  1 02:53:45.071: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to up
*Mar  1 02:53:45.083: %LINK-3-UPDOWN: Interface FastEthernet0/12, changed state to up
Switch-2(config-if-range)#
*Mar  1 02:53:48.619: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to up
*Mar  1 02:53:48.631: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/12, changed state to up
Switch-2(config-if-range)#do show int fast 0/10 | i packets output
    378 packets output
, 53506 bytes, 0 underrunsSwitch-2(config-if-range)#do show int fast 0/10 | i packets output
    381 packets output
, 53896 bytes, 0 underrunsSwitch-2(config-if-range)#do show int fast 0/10 | i packets output
    383 packets output
, 54427 bytes, 0 underrunsSwitch-2(config-if-range)#do show int fast 0/10 | i packets output
    386 packets output
, 54817 bytes, 0 underrunsSwitch-2(config-if-range)#do show int fast 0/10 | i packets output
    388 packets output
, 55113 bytes, 0 underrunsSwitch-2(config-if-range)#do show int fast 0/10 | i packets output
    391 packets output
, 55503 bytes, 0 underruns
```

```

Switch-2(config-if-range)#do show int fast 0/10 | i packets output
392 packets output
, 55651 bytes, 0 underrunsSwitch-2(config-if-range)#do show int fast 0/10 | i packets output
394 packets output
, 55947 bytes, 0 underrunsSwitch-2(config-if-range)#do show int fast 0/10 | i packets output
396 packets output
, 56243 bytes, 0 underruns
Switch-2(config-if-range)#*Mar 1 02:54:19.631:
%STORM_CONTROL-3-FILTERED: A Broadcast storm detected on Fa0/10. A packet filter action has been applied on the interface
Switch-2(config-if-range)#
*Mar 1 02:54:19.651: %SW_MATM-4-MACFLAP_NOTIF: Host 001c.589e.7ae1 in vlan 1 is flapping between port Fa0/12 and po
Switch-2(config-if-range)#do show int fast 0/10 | i packets output
5378 packets output
, 793525 bytes, 10108 underruns
Switch-2(config-if-range)#
*Mar 1 02:54:23.703: %STORM_CONTROL-3-FILTERED: A Broadcast storm detected on Fa0/10. A packet filter action has been applied on the interface
Switch-2(config-if-range)#do show int fast 0/10 | i packets output 15730 packets output
, 2325567 bytes, 31841 underruns
Switch-2(config-if-range)#do show int fast 0/10 | i packets output 29262 packets output
, 4328303 bytes, 60766 underruns
Switch-2(config-if-range)#
*Mar 1 02:54:27.775: %STORM_CONTROL-3-FILTERED: A Broadcast storm detected on Fa0/10. A packet filter action has been applied on the interface
Switch-2(config-if-range)#do show int fast 0/10 | i packets output 51821 packets output
, 7666981 bytes, 108650 underruns
Switch-2(config-if-range)#
*Mar 1 02:54:31.867: %STORM_CONTROL-3-FILTERED: A Broadcast storm detected on Fa0/10. A packet filter action has been applied on the interface
Switch-2(config-if-range)#shutdown

```

Task

- If they are not already down, **shut down all three interfaces on Switch-2 connecting it to Switch-1** (FastEthernet0/10 - 12).
- Change the **action of Storm-Control** to automatically **shutdown** these three interfaces when broadcast storms trigger this feature.
- Configure the **errdisable recovery feature** as follows:
 - The feature should be allowed to **automatically recover a port** that was placed into the err-disable state **due to Storm-Control**.
 - The feature should automatically recover ports that meet the criteria above **after 30 seconds**.

Storm-Control Modification - Configuration

```
Switch-2(config)#interface range fast 0/10 - 12  
Switch-2(config-if-range)#shutdown  
Switch-2(config-if-range)#storm-control action shutdown  
Switch-2(config-if-range)#exit  
Switch-2(config)#errdisable recovery cause storm-control  
Switch-2(config)#errdisable recovery interval 30  
  
Switch-2(config)#
```

Task

In this final task, you will observe the Storm-Control feature as it shuts down some interfaces. You will then see the Error-Disable Recovery feature enable these ports automatically after 30 seconds...only to see them shut down again moments later by Storm-Control.

- Ensure that your continuous ping from the host to 1.1.1.2 is still active. If not, restart it.
- In Switch-2, using the `interface range` command, enable ports 0/10 - 12.

```

Switch-2(config)#interface range fast 0/10 - 12
Switch-2(config-if-range)#no shutdown

Switch-2(config-if-range)#end
Switch-2#

Switch-2#
*Mar  1 03:09:39.063: %SYS-5-CONFIG_I: Configured from console by console
Switch-2#
*Mar  1 03:09:40.235: %LINK-3-UPDOWN: Interface FastEthernet0/10, changed state to up
*Mar  1 03:09:40.247: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to up
*Mar  1 03:09:40.259: %LINK-3-UPDOWN: Interface FastEthernet0/12, changed state to up
Switch-2# *Mar  1 03:09:43.799: %LINEPROTO-5-UPDOWN:
Line protocol on Interface FastEthernet0/11, changed state to up
*Mar  1 03:09:43.919: %LINEPROTO-5-UPDOWN:
Line protocol on Interface FastEthernet0/12, changed state to up
*Mar  1 03:09:43.959: %LINEPROTO-5-UPDOWN:
Line protocol on Interface FastEthernet0/10, changed state to up

Switch-2#

```

Notice that about 30 seconds after the Line protocol of all three interfaces changes state to up, and Spanning-Tree places the ports into the Forwarding state, a Broadcast storm is detected by Storm-Control.

```

Switch-2#
Switch-2#
Switch-2#
*Mar  1 03:10:13.703: %SW_MATM-4-MACFLAP_NOTIF: Host 001c.589e.7ae1 in vlan 1 is flapping between port Fa0/12 and po
*Mar  1 03:10:14.135: %PM-4-ERR_DISABLE:
storm-control error detected on Fa0/10, putting Fa0/10 in err-disable state
*Mar  1 03:10:14.247: %STORM_CONTROL-3-SHUTDOWN: A packet storm was detected on Fa0/10. The interface has been disab
Switch-2# *Mar  1 03:10:14.247: %PM-4-ERR_DISABLE:
storm-control error detected on Fa0/11, putting Fa0/11 in err-disable state
*Mar  1 03:10:14.355: %PM-4-ERR_DISABLE:
storm-control error detected on Fa0/12, putting Fa0/12 in err-disable state

Switch-2#
*Mar  1 03:10:15.183: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/10, changed state to down
*Mar  1 03:10:15.291: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to down
*Mar  1 03:10:15.367: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/12, changed state to down
Switch-2#
*Mar  1 03:10:16.243: %LINK-3-UPDOWN: Interface FastEthernet0/10, changed state to down

```

```
*Mar  1 03:10:16.351: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to down *Mar  1 03:10:16.379: %LINK-3-UPDOWN: Interface FastEthernet0/12, changed state to down
Switch-2#
Switch-2#
```

About 30 seconds after Storm-Control has disabled the interfaces and they go down, the Error-Disable Recovery feature enables the ports again.

```
Switch-2#
Switch-2# *Mar  1 03:10:44.247: %PM-4-ERR_RECOVER:
Attempting to recover from storm-control err-disable state on Fa0/10 *Mar  1 03:10:44.355: %PM-4-ERR_RECOVER:
Attempting to recover from storm-control err-disable state on Fa0/11 *Mar  1 03:10:44.371: %PM-4-ERR_RECOVER:
Attempting to recover from storm-control err-disable state on Fa0/12
Switch-2#
Switch-2#
Switch-2#
*Mar  1 03:10:47.899: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to up
*Mar  1 03:10:47.979: %LINK-3-UPDOWN: Interface FastEthernet0/12, changed state to up
*Mar  1 03:10:48.063: %LINK-3-UPDOWN: Interface FastEthernet0/10, changed state to up
Switch-2#
*Mar  1 03:10:49.915: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to up
*Mar  1 03:10:49.995: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/12, changed state to up
*Mar  1 03:10:50.075: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/10, changed state to up
Switch-2#
Switch-2#
```

And again, as soon as the interfaces go up and pass into the Spanning-Tree Forwarding state, Storm-Control is invoked and disables all three interfaces.

```
Switch-2#
*Mar  1 03:11:19.691: %SW_MATM-4-MACFLAP_NOTIF: Host 001c.589e.7ae1 in vlan 1 is flapping between port Fa0/12 and po
*Mar  1 03:11:19.939: %PM-4-ERR_DISABLE:
storm-control error detected on Fa0/10, putting Fa0/10 in err-disable state
*Mar  1 03:11:20.055: %STORM_CONTROL-3-SHUTDOWN: A packet storm was detected on Fa0/10. The interface has been disabled.
*Mar  1 03:11:20.055: %PM-4-ERR_DISABLE:
storm-control error detected on Fa0/11, putting Fa0/11 in err-disable state
```

```
*Mar  1 03:11:20.163: %PM-4-ERR_DISABLE:
```

```
storm-control error detected on Fa0/12, putting Fa0/12 in err-disable state
```

```
*Mar  1 03:11:20.987: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/10, changed state to down
```

```
*Mar  1 03:11:21.099: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to down
```

```
*Mar  1 03:11:21.175: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/12, changed state to down
```

```
Switch-2#
```

```
*Mar  1 03:11:22.047: %LINK-3-UPDOWN: Interface FastEthernet0/10, changed state to down
```

```
*Mar  1 03:11:22.159: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to down
```

```
*Mar  1 03:11:22.187: %LINK-3-UPDOWN: Interface FastEthernet0/12, changed state to down
```

```
Switch-2#
```

```
Switch-2#
```

```
Switch-2#
```


- On Switch-2, **disable interfaces FastEthernet0/10 - 12.**
- **Stop the continuous ping** from the Hhst.

To stop a continuous ping, on your keyboard type Control-Shift-6-6. If you are using INE's web-based GUI to access the CLI of the devices, this command sequence might not work. You will have to be using your own Telnet client (PuTTY, Hyperterminal, etc.) to use this command sequence.

CCNP SWITCH Workbook - First Hop Redundancy Protocols

8.1 HSRP Configuration

Load the **CCNP-Switch-Task8-1** initial configurations before starting.



All of the physical connections, VLANs, IP addresses, DHCP configuration, and Routing Protocols (EIGRP) shown in the topology diagram have been pre-configured for you.

Task

In this first task, you will complete a basic HSRP configuration.

First, ensure that your current topology is up and functional (**you may have to enable some interfaces**).

- Verify that your DHCP client (R4) has obtained an IP address via DHCP.

The IP address of the default-gateway (1.2.3.123) will not be reachable yet because you have not configured HSRP on any routers at this point.

DHCP Client Verification

```
Host#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Host(config)#int fast 0/1
Host(config-if)#shutdown
Host(config-if)#no shutdown
Host(config-if)#
Nov 17 09:38:13.531: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
Nov 17 09:38:14.531: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
Nov 17 09:38:20.639:
%DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 1.2.3.4, mask 255.255.255.0, hostname Host
```

Task

Configure HSRP on Router-1, Router-2, and Router-3 following these guidelines:

- HSRP will provide gateway redundancy for the DHCP client (R4).
- Use an HSRP **Group number of one (1)**.
- Use an HSRP **Virtual-IP address of 1.2.3.123**.
- Without using any "show" commands, **can you predict** which router will be the HSRP Active router? The HSRP Standby router?
- **Verify** whether your prediction about the routers playing the roles of HSRP Active and Standby were correct.

HSRP Initial Configuration

Router-1#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-1(config)#**int fast 0/1**

Router-1(config-if)#**standby 1 ip 1.2.3.123**

Router-1(config-if)#end

Router-1#

Router-2#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-2(config)#**int fast 0/1**

Router-2(config-if)#**standby 1 ip 1.2.3.123**

Router-2(config-if)#end

Router-2#

Router-3#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-3(config)#**int fast 0/0**

Router-3(config-if)#**standby 1 ip 1.2.3.123**

Router-3(config-if)#end

Router-3#

HSRP Verification

By default, the very first router to be configured with HSRP would assume the role of the HSRP Active router. Subsequent to that, even though HSRP Preemption is disabled by default, all remaining routers connected to the same HSRP Group will elect amongst each other the HSRP Standby router based on highest IP address (because all routers have the same, default HSRP priority).

If a new router comes online later with a higher IP address than the current HSRP **Standby** router, the new router will assume the role of HSRP Standby.

In the example below, Router-1 was the first router to be configured with HSRP so it became the HSRP Active router. Router-3 was the last router to be configured with HSRP, but it has a higher interface IP address than Router-2 so Router-3 became the HSRP Standby router.

Router-1#show standby

```
FastEthernet0/1 - Group 1 State is Active
    2 state changes, last state change 00:09:39
Virtual IP address is 1.2.3.123
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 2.080 secs
Preemption disabled Active router is local
Standby router is 1.2.3.3
, priority 100 (expires in 9.440 sec)
    Priority 100 (default 100)
    Group name is "hsrp-Fa0/1-1" (default)
Router-1#
```

Router-2#show standby

```
FastEthernet0/1 - Group 1 State is Listen
    4 state changes, last state change 00:03:23
Virtual IP address is 1.2.3.123
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
Preemption disabled Active router is 1.2.3.1
, priority 100 (expires in 9.792 sec) Standby router is 1.2.3.3
, priority 100 (expires in 11.040 sec)
    Priority 100 (default 100)
    Group name is "hsrp-Fa0/1-1" (default)
Router-2#
```

Router-3#show standby

```
FastEthernet0/0 - Group 1 State is Standby
    3 state changes, last state change 00:03:28
Virtual IP address is 1.2.3.123
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 1.616 secs
```

```
Preemption disabled Active router is 1.2.3.1
, priority 100 (expires in 10.096 sec) Standby router is local

Priority 100 (default 100)
Group name is "hsrp-Fa0/0-1" (default)
Router-3#
```

HSRP Virtual MAC Address

As you know, HSRP utilizes a virtual MAC address at the OSI Datalink Layer. This MAC address is monitored at any given time by the HSRP Active Router.

- View the HSRP Virtual MAC address in the MAC Address-Table of Switch-2.

```
Switch-2#sho mac address-table
```

Mac Address Table

```
-----
Vlan    Mac Address      Type    Ports
----    -
All     000c.8581.a500   STATIC  CPU
...
<output omitted for brevity>
... 24 0000.0c07.ac01 DYNAMIC Fa0/2
```

```
Switch-2#show mac address-table address 0000.0c07.ac01
```

Mac Address Table

```
-----
Vlan    Mac Address      Type    Ports
----    -
24      0000.0c07.ac01   DYNAMIC Fa0/2

Total Mac Addresses for this criterion: 1
Switch-2#
```

Test Basic HSRP Functionality

For this next task, you'll need to have two Telnet windows open simultaneously. One window should allow you to view the output of the DHCP client (R4) while the other

allows you to make changes to Switch-2.

- **Begin a continuous ping** from the DHCP client (R4) to the IP address of **111.111.111.111 (repeat count of 10,000)**.
- **Shut down** the FastEthernet interface of Switch-2 that connects to your HSRP Active router.
- **Monitor** how many (if any) pings are lost (and the length of time until pings resume) until one of the other routers assumes the role of HSRP Active Router.
- How long did it take for the HSRP Standby Router to switch over to HSRP Active?

```
Host# Host# ping 111.111.111.111 repeat 10000  
Type escape sequence to abort.  
Sending 10000, 100-byte ICMP Echos to 111.111.111.111, timeout is 2 seconds:  
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!  
  
...  
<output omitted for brevity>
```

```
Switch-2(config)#int fast 0/2
Switch-2(config-if)#shut

Switch-2(config-if)#
```

Output seen on Host:


```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
<output omitted for brevity>Success rate is 99 percent
(9994/10000), round-trip min/avg/max = 1/2/28 ms
Host#
```

In the test above, pings were stopped for about 10 seconds and then resumed.

CCNP SWITCH Workbook - First Hop Redundancy Protocols

8.2 Configuring Sub-Second HSRP Timers

Load the **CCNP-Switch-Task8-2** initial configurations before starting.



All of the physical connections, VLANs, IP addresses, DHCP configuration, and Routing Protocols (EIGRP) shown in the topology diagram have been pre-configured for you.

Task

In this task, you will configure HSRP Preemption so that Router-2 always becomes the HSRP Active router.

First, ensure that your current topology is up and functional (**you may have to enable some interfaces**).

- **Disable** interface FastEthernet0/1 on Router-2.
- **Increase the HSRP Priority of Router-2** (for HSRP Group-1 on interface

FastEthernet0/1) to **101**.

- **Enable HSRP Preemption** on Router-2.
- On Router-2, while still within interface configuration mode, enable the command `do debug standby` so you can view the HSRP State Machine as Router-2 becomes part of the HSRP Group.
- Enable interface FastEthernet0/1 and view the debug output.
- After viewing the debug output, disable all debugs with the privileged EXEC command `undebug all`.

Router-2 Configuration

```
Router-2#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Router-2(config)#int fast 0/1
Router-2(config-if)#shutdown
Router-2(config-if)#
Nov 18 08:03:28.943: %HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Standby -> Init
Router-2(config-if)#
Nov 18 08:03:30.939: %LINK-5-CHANGED: Interface FastEthernet0/1, changed state to administratively down
Nov 18 08:03:31.939: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down
Router-2(config-if)#standby 1 priority 101
Router-2(config-if)#standby 1 preempt
Router-2(config-if)#do debug standby
HSRP debugging is on

Router-2(config-if)#
```

Monitoring HSRP Debug Output

```
Router-2(config-if)#no shutdown
Router-2(config-if)#^Z
Router-2#
Router-2#
Nov 18 08:06:22.231: HSRP: Fa0/1 Interface UP
Nov 18 08:06:22.231: HSRP: Fa0/1 Starting minimum intf delay (1 secs)
Nov 18 08:06:22.231: HSRP: Fa0/1 ARP reload
Nov 18 08:06:22.231: HSRP: Fa0/1 ARP reload
Nov 18 08:06:22.295: HSRP: Fa0/1 Nbr 1.2.3.1 Passive timer expired
Nov 18 08:06:22.295: HSRP: Fa0/1 Nbr 1.2.3.1 is no longer passive
Nov 18 08:06:22.295: HSRP: Fa0/1 Nbr 1.2.3.1 destroyed
Nov 18 08:06:23.223: HSRP: Fa0/1 Intf min delay expired
```

```

Nov 18 08:06:23.223: HSRP: Fa0/1 Grp 1 Init: a/HSRP enabled
Nov 18 08:06:23.223: HSRP: Fa0/1 Grp 1 Init -> Listen
Nov 18 08:06:23.223: HSRP: Fa0/1 Interface adv out, Passive, active 0 passive 1Nov 18 08:06:23.223:
HSRP: Fa0/1 Grp 1 Redundancy "hsrp-Fa0/1-1" state Init -> Backup
Nov 18 08:06:23.223: HSRP: Fa0/1 IP Redundancy "hsrp-Fa0/1-1" update, Init -> Backup
Router-2#
Router-2#
Nov 18 08:06:23.315: %SYS-5-CONFIG_I: Configured from console by console
Nov 18 08:06:24.219: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
Router-2#
Nov 18 08:06:24.475: HSRP: Fa0/1 Grp 1 MAC addr update Delete from SMF 0000.0c07.ac01
Nov 18 08:06:25.155: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.1 Standby pri 100 vIP 1.2.3.123
Nov 18 08:06:25.155: HSRP: Fa0/1 Grp 1 Listen: l/Hello rcvd from lower pri Standby router (100/1.2.3.1)
Nov 18 08:06:25.155: HSRP: Fa0/1 Grp 1 Standby router is 1.2.3.1
Nov 18 08:06:25.155: HSRP: Fa0/1 Nbr 1.2.3.1 created
Nov 18 08:06:25.155: HSRP: Fa0/1 Nbr 1.2.3.1 standby for group 1Nov 18 08:06:25.155:
HSRP: Fa0/1 Grp 1 Listen -> Speak
Nov 18 08:06:25.155: HSRP: Fa0/1 Grp 1 Redundancy "hsrp-Fa0/1-1" state Backup -> Speak
Nov 18 08:06:25.159: HSRP: Fa0/1 Grp 1 Hello out 1.2.3.2 Speak pri 101 vIP 1.2.3.123
Nov 18 08:06:25.159: HSRP: Fa0/1 IP Redundancy "hsrp-Fa0/1-1" standby, unknown -> 1.2.3.1
Nov 18 08:06:25.159: HSRP: Fa0/1 IP Redundancy "hsrp-Fa0/1-1" update, Backup -> Speak
Router-2#
Nov 18 08:06:25.219: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
Nov 18 08:06:26.147: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.3 Active pri 100 vIP 1.2.3.123
Nov 18 08:06:26.151: HSRP: Fa0/1 Grp 1 Active router is 1.2.3.3
Nov 18 08:06:26.151: HSRP: Fa0/1 Nbr 1.2.3.3 is no longer passiveNov 18 08:06:26.151:
HSRP: Fa0/1 Nbr 1.2.3.3 active for group 1
Nov 18 08:06:26.151: HSRP: Fa0/1 Interface adv out, Passive, active 0 passive 1Nov 18 08:06:26.151:
HSRP: Fa0/1 Grp 1 Speak: h/Hello rcvd from lower pri Active router (100/1.2.3.3)
Nov 18 08:06:26.151: HSRP: Fa0/1 Grp 1 Active router is local, was 1.2.3.3
Nov 18 08:06:26.151: HSRP: Fa0/1 Nbr 1.2.3.3 no longer active for group 1 (Speak)
Nov 18 08:06:26.151: HSRP: Fa0/1 Nbr 1.2.3.3 Was active or standby - start passive holddown
Nov 18 08:06:26.151: HSRP: Fa0/1 Interface adv out, Active, active 1 passive 1
Nov 18 08:06:26.151: HSRP: Fa0/1 Grp 1 Coup out 1.2.3.2 Speak pri 101 vIP 1.2.3.123
Nov 18 08:06:26.151: HSRP: Fa0/1 Grp 1 Speak -> ActiveNov 18 08:06:26.151:
%HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Speak -> Active
Nov 18 08:06:26.151: HSRP: Fa0/1 Interface adv out, Active, active 1 passive 0
Nov 18 08:06:26.151: HSRP: Fa0/1 Grp 1 Redundancy "hsrp-Fa0/1-1" state Speak -> Active
Nov 18 08:06:26.151: HSRP: Fa0/1 Grp 1 Hello out 1.2.3.2 Active pri 101 vIP 1.2.3.123
Nov 18 08:06:26.155: HSRP: Fa0/1 Grp 1 Added 1.2.3.123 to ARP (0000.0c07.ac01)
Nov 18 08:06:26.155: HSRP: Fa0/1 Grp 1 Activating MAC 0000.0c07.ac01
Nov 18 08:06:26.155: HSRP: Fa0/1 Grp 1 Adding 0000.0c07.ac01 to MAC address filter
Nov 18 08:06:26.155: HSRP: Fa0/1 IP Redundancy "hsrp-Fa0/1-1" update, Speak -> Active
Nov 18 08:06:26.155: HSRP: Fa0/1 Interface adv in, Passive, active 0, passive 1, from 1.2.3.1
Router-2#
Nov 18 08:06:26.155: HSRP: Fa0/1 Interface adv in, Passive, active 0, passive 1, from 1.2.3.3

```

```

Nov 18 08:06:26.163: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.3 Speak pri 100 vIP 1.2.3.123
Router-2#
Nov 18 08:06:28.903: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.3 Speak pri 100 vIP 1.2.3.123
Nov 18 08:06:29.115: HSRP: Fa0/1 Grp 1 Hello out 1.2.3.2 Active pri 101 vIP 1.2.3.123
Nov 18 08:06:29.147: HSRP: Fa0/1 IP Redundancy "hsrp-Fa0/1-1" update, Active -> Active
Router-2#
Nov 18 08:06:31.559: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.3 Speak pri 100 vIP 1.2.3.123
Nov 18 08:06:31.915: HSRP: Fa0/1 Grp 1 Hello out 1.2.3.2 Active pri 101 vIP 1.2.3.123
Router-2#
Nov 18 08:06:34.427: HSRP: Fa0/1 Grp 1 Hello out 1.2.3.2 Active pri 101 vIP 1.2.3.123
Nov 18 08:06:34.551: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.3 Speak pri 100 vIP 1.2.3.123
Router-2#Nov 18 08:06:36.699: HSRP: Fa0/1 Grp 1 Standby router is unknown, was 1.2.3.1
Nov 18 08:06:36.699: HSRP: Fa0/1 Nbr 1.2.3.1 no longer standby for group 1 (Active)
Nov 18 08:06:36.699: HSRP: Fa0/1 Nbr 1.2.3.1 Was active or standby - start passive holddown
Nov 18 08:06:37.275: HSRP: Fa0/1 Grp 1 Hello out 1.2.3.2 Active pri 101 vIP 1.2.3.123
Nov 18 08:06:37.431: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.3 Speak
pri 100 vIP 1.2.3.123
Router-2#Nov 18 08:06:37.975: HSRP: Fa0/1 Grp 1 Hello in 1.2.3.3 Standby
pri 100 vIP 1.2.3.123Nov 18 08:06:37.975: HSRP: Fa0/1 Grp 1 Standby router is 1.2.3.3

Nov 18 08:06:37.975: HSRP: Fa0/1 Nbr 1.2.3.3 is no longer passive
Nov 18 08:06:37.975: HSRP: Fa0/1 Nbr 1.2.3.3 standby for group 1

```

Task

- Modify the HSRP timers on Router-2 so that:
 - This router transmits HSRP **Hello packets every 100msecs**.
 - This router uses an HSRP **Holdtime of 300msecs**.

Router-2 Configuration

```

Router-2#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Router-2(config)#int fast 0/1
Router-2(config-if)#standby 1 timers msec 100 msec 300

Router-2(config-if)#end
Router-2#

```

Task

Verify whether your changes have reduced the HSRP convergence time by doing the following:

- On Router-3:
 - In global configuration mode, enter `logging buffer debug` . This will copy all debug output to a memory buffer (log).
 - While still in global configuration mode, enter `service timestamps debug datetime msec` . This command will ensure that timestamps, down to millisecond values, will be applied to any debug output.
 - Enter the command `no logging console debug` to ensure that any debug you enable will not have its output logged to the console.
 - Exit back to privileged EXEC mode on Router-3 and clear any existing contents of the logging buffer with the command `clear log` .
 - On Router-3, enable the command `debug standby` .

Router-3 Debug Preparation

```
Router-3#conf t
Enter configuration commands, one per line. End with CNTL/Z. Router-3(config)#logging buffer debug
Router-3(config)#no logging console debug
Router-3(config)#service timestamps debug datetime msec
Router-3(config)#end Router-3#clear log
Clear logging buffer [confirm]
Router-3#debug standby
HSRP debugging is on

Router-3#
```

Go to Switch-2, and disable the FastEthernet0/1 interface.

Removing the Current HSRP Active Router from the Topology

```
Switch-2>
en
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#int fast 0/1
Switch-2(config-if)#shutdown

Switch-2(config-if)#end
```

- Go back to **Router-3** and disable all debugs with the command `undebug all`
- **View the debug output** on Router-3 by entering the command `show log`.
 - How did Router-3 first learn that the HSRP Active Router (Router-2) was no longer alive?
 - After determining that the HSRP Active Router was no longer alive, how long did it take Router-3 to assume the role of HSRP Active?

Viewing the Debug Output on Router-3

```
Router-3#undebug all
```

```
All possible debugging has been turned offRouter-3#show log
```

```
...
```

```
<output omitted for brevity>
```

```
Log Buffer (4096 bytes):
```

```
...
```

```
<output omitted for brevity>
```

```
...
```

```
Nov 18 08:51:54.647: HSRP: Fa0/0 Grp 1 Hello in 1.2.3.2 Active pri 101 vIP 1.2.3.123Nov 18
```

```
08:51:54.743: HSRP: Fa0/0 Grp 1Hello in 1.2.3.2 Active pri 101
```

```
vIP 1.2.3.123
```

```
Nov 18 08:51:56.295: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123
```

```
Nov 18 08:51:57.935: HSRP: Fa0/0 Interface adv in, Passive, active 0, passive 1, from 1.2.3.1
```

```
Nov 18 08:51:58.951: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123
```

```
Nov 18 08:52:01.703: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123
```

```
Nov 18 08:52:04.567: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123Nov 18
```

```
08:52:05.527: HSRP: Fa0/0 Grp 1 Standby: c/Active timer expired (1.2.3.2)
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Grp 1Active router is local, was 1.2.3.2
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Nbr 1.2.3.2 no longer active for group 1 (Standby)
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Nbr 1.2.3.2 Was active or standby - start passive holddown
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Grp 1 Standby router is unknown, was localNov 18
```

```
08:52:05.527: HSRP: Fa0/0 Grp 1 Standby -> Active
```

```
Nov 18 08:52:05.527: %HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Standby -> Active
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Interface adv out, Active, active 1 passive 0
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Grp 1 Redundancy "hsrp-Fa0/0-1" state Standby -> Active
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Active pri 100 vIP 1.2.3.123
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Grp 1 Added 1.2.3.123 to ARP (0000.0c07.ac01)
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Grp 1 Activating MAC 0000.0c07.ac01
```

```
Nov 18 08:52:05.527: HSRP: Fa0/0 Grp 1 Adding 0000.0c07.ac01 to MAC address filter
```

```
Nov 18 08:52:05.531: HSRP: Fa0/0 IP Redundancy "hsrp-Fa0/0-1" standby, local -> unknown
```

```
Nov 18 08:52:05.531: HSRP: Fa0/0 IP Redundancy "hsrp-Fa0/0-1" update, Standby -> Active
```

In the debug output on Router-3, we notice the following:

1. The last HSRP Hello that Router-3 receives from the HSRP Active Router (Router-2) is at timestamp **08:51:54.743**.
2. The "Active Timer" (Holdtime) on Router-3 **does not expire until 08:52:05.527**, which is roughly **10 cseconds** after it received the last HSRP Hello from Router-2.
3. At that same timestamp of 08:52:05.527, Router-3 assumes the HSRP Active Router role.

Notice the HSRP Packet format, shown below:

Version	Op Code	State	Hellotime
Holdtime	Priority	Group	Reserved
Authentication Data			
Authentication Data			
Virtual IP Address			

You can see from the image above that the HSRP Active Router DOES include its **Hello Advertisement Interval**, and also tells its HSRP peers how long to wait until declaring the Active router to be down (the **Holdtime**) as fields within the HSRP Hello packet. So if Router-2 was advertising sub-second timers, why did Router-3 wait 10 seconds before transitioning its state to Active?

The problem here is that **both of these fields within the packet are only 1-byte in length**. These fields are too small to contain any value greater than 255 (or 11111111 in binary). How could they be used to advertise something like 300 milliseconds? Or 800 milliseconds? When you configure sub-second values on the HSRP Active router (as you did in Router-2), the value of both the Hello and Holdtime fields will be set to one (1).

When an HSRP Router receives a HSRP Hello packet with timers of one (1), there is no way for that receiving router to determine whether the advertised Holdtime is really 1-second...or something LESS than 1-second. The receiving router will basically ignore those advertised timers and stick with the default timers (Hello = 3-seconds and Holdtime = 10-seconds).

In summary, even though you configured Router-2 to transmit its Hello packets every 100 milliseconds and configured a Holdtime of 300 milliseconds, it did not

decrease the HSRP convergence time because Router-2 had no way to place those specific values within the HSRP Hello packet.

Task

- Configure both Router-1 and Router-3 with the exact same HSRP timers you configured on Router-2.
- Perform the same set of steps you did in the proceeding task and view the new convergence time of HSRP.

Configuration of Router-1 and Router-3

```
Router-1#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Router-1(config)#int fast 0/1
Router-1(config-if)#standby 1 timers msec 100 msec 300
Router-1(config-if)#end
Router-1#
Nov 18 09:28:53.723: %SYS-5-CONFIG_I: Configured from console by console
Nov 18 09:28:54.375: %HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Speak -> Standby
Router-1#
RS-CCNP-Cage1-AS#3
[Resuming connection 3 to rsprack1-r3 ... ]

Router-3>enRouter-3#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Router-3(config)#int fast 0/0
Router-3(config-if)#standby 1 timers msec 100 msec 300
Router-3(config-if)#end
Router-3#
Router-3#Router-3#show standby
FastEthernet0/0 - Group 1
  State is Standby
    13 state changes, last state change 00:32:19
  Virtual IP address is 1.2.3.123
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default) Hello time 100 msec, hold time 300 msec
  Next hello sent in 0.032 secs

  Preemption disabled
  Active router is 1.2.3.2, priority 101 (expires in 0.352 sec)
  Standby router is local
  Priority 100 (default 100)
```

```
Group name is "hsrp-Fa0/0-1" (default)
```

```
Router-3#
```

Router-3 Debug Preparation

Because we now have the HSRP Active router transmitting Hellos every 100 milliseconds, we will probably need to increase the size of the logging buffer on Router-3 so that it can hold more information.

```
Router-3(config)#logging buffer 100000
Router-3(config)#end Router-3#clear log
Clear logging buffer [confirm]
Router-3#debug standby
HSRP debugging is on

Router-3#
```

Removing the Current HSRP Active Router from the Topology

```
Switch-2>
en
Switch-2#conf t
Enter configuration commands, one per line. End with CNTL/Z. Switch-2(config)#int fast 0/1
Switch-2(config-if)#shutdown

Switch-2(config-if)#end
Switch-2#
```

Viewing the Debug Output on Router-3

```
Router-3#show log
...
<output omitted for brevity>
...
Nov 18 09:35:49.647: HSRP: Fa0/0 Grp 1 Hello in 1.2.3.2 Active pri 101 vIP 1.2.3.123
Nov 18 09:35:49.695: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123 Nov 18
09:35:49.743: HSRP: Fa0/0 Grp 1 Hello in 1.2.3.2 Active
pri 101 vIP 1.2.3.123
Nov 18 09:35:49.807: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123
Nov 18 09:35:49.919: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123
Nov 18 09:35:50.015: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Standby pri 100 vIP 1.2.3.123 Nov 18
09:35:50.079: HSRP: Fa0/0 Grp 1 Standby: c/Active timer expired (1.2.3.2)
```

```

Nov 18 09:35:50.079: HSRP: Fa0/0 Grp 1 Active router is local, was 1.2.3.2

Nov 18 09:35:50.079: HSRP: Fa0/0 Nbr 1.2.3.2 no longer active for group 1 (Standby)
Nov 18 09:35:50.079: HSRP: Fa0/0 Nbr 1.2.3.2 Was active or standby - start passive holddown
Nov 18 09:35:50.079: HSRP: Fa0/0 Grp 1 Standby router is unknown, was local
Nov 18 09:35:50.079: HSRP: Fa0/0 Grp 1 Standby -> Active
Nov 18 09:35:50.079: %HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Standby -> Active
Nov 18 09:35:50.079: HSRP: Fa0/0 Interface adv out, Active, active 1 passive 0
Nov 18 09:35:50.079: HSRP: Fa0/0 Grp 1 Redundancy "hsrp-Fa0/0-1" state Standby -> Active
Nov 18 09:35:50.079: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Active pri 100 vIP 1.2.3.123
Nov 18 09:35:50.079: HSRP: Fa0/0 Grp 1 Added 1.2.3.123 to ARP (0000.0c07.ac01)
Nov 18 09:35:50.079: HSRP: Fa0/0 Grp 1 Activating MAC 0000.0c07.ac01Nov 18
09:35:50.079: HSRP: Fa0/0 Grp 1 Adding 0000.0c07.ac01 to MAC address filter

Nov 18 09:35:50.079: HSRP: Fa0/0 IP Redundancy "hsrp-Fa0/0-1" standby, local -> unknown
Nov 18 09:35:50.083: HSRP: Fa0/0 IP Redundancy "hsrp-Fa0/0-1" update, Standby -> Active
Nov 18 09:35:50.083: HSRP: Fa0/0 Grp 1 Hello in 1.2.3.1 Speak pri 100 vIP 1.2.3.123
Nov 18 09:35:50.175: HSRP: Fa0/0 Grp 1 Hello out 1.2.3.3 Active pri 100 vIP 1.2.3.123

```


Now that we have changed the Hello and Hold timers locally on both Router-1 and Router-3, they are aware that HSRP is expected to receive HSRP Hello packets much more quickly than every 3 seconds.

Above we can see that the **last HSRP Hello received** from the Active Router (Router-2) was at timestamp 09:35:49.743. Then, about **300 milliseconds later**, that router is declared dead (the **Active timer expired at 50.079**) and Router-3 assumes the role of the HSRP Active router.

CCNP SWITCH Workbook - First Hop Redundancy Protocols

8.3 HSRP Object Tracking

Load the **CCNP-Switch-Task8-3** initial configurations before starting.



All of the physical connections, VLANs, IP addresses, DHCP configuration, and Routing Protocols (EIGRP) shown in the topology diagram have been pre-configured for you.

Task

In this task, you will configure HSRP to monitor a Track Object.

First, ensure that your current topology is up and functional (**you may have to enable some interfaces**).

- Configure a **Track Object** that tracks the **reachability of the ip route 111.111.111.0/24**.
- Configure HSRP on Routers-1, 2, and 3 so that:

- HSRP watches the Track Object.
- If the tracked route becomes unreachable, a router will **decrement its HSRP priority by fifty (50)**.

Router Configuration

```
Router-1#conf t
Router-1(config)#track 1 ip route 111.111.111.0/24 reachability

Router-1(config-track)#exit
Router-1(config)#int fast 0/1
Router-1(config-if)#standby 1 track 1 decrement 50

Router-1(config-if)#end
Router-1#
```

```
Router-2>enable
Router-2#config t
Enter configuration commands, one per line. End with CNTL/Z.
Router-2(config)#track 1 ip route 111.111.111.0/24 reachability

Router-2(config-track)#exit
Router-2(config)#int fast 0/1
Router-2(config-if)#standby 1 track 1 decrement 50

Router-2(config-if)#end
Router-2#
Router-2#
Nov 18 11:25:07.194: %SYS-5-CONFIG_I: Configured from console by console
Router-2#
```

```
Router-3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Router-3(config)#track 1 ip route 111.111.111.0/24 reachability

Router-3(config-track)#exit
Router-3(config)#int fast 0/0
Router-3(config-if)#standby 1 track 1 decrement 50

Router-3(config-if)#end
Router-3#
```

Configuration Verification

```
Router-3#show standby
```

```
FastEthernet0/0 - Group 1
```

```
State is Active
```

```
17 state changes, last state change 01:51:17
```

```
Virtual IP address is 1.2.3.123
```

```
Active virtual MAC address is 0000.0c07.ac01
```

```
Local virtual MAC address is 0000.0c07.ac01 (v1 default)
```

```
Hello time 100 msec, hold time 300 msec
```

```
Next hello sent in 0.016 secs
```

```
Preemption disabled
```

```
Active router is local
```

```
Standby router is 1.2.3.1, priority 100 (expires in 0.256 sec)
```

```
Priority 100 (default 100) Track object 1 state Up decrement 50
```

```
Group name is "hsrp-Fa0/0-1" (default)
```

```
Router-3#
```

```
Router-3#show track
```

```
Track 1
```

```
IP route 111.111.111.0 255.255.255.0 reachability
```

```
Reachability is Up (EIGRP)
```

```
1 change, last change 00:01:46
```

```
First-hop interface is FastEthernet0/1 Tracked by:
```

```
HSRP FastEthernet0/0 1
```

```
Router-3#
```

Task

- Configure HSRP Preemption on Routers-1, 2, and 3.
- Ensure that Router-1, Router-2, and Router-3 have the following HSRP Priority values:
 - Router-1: 110
 - Router-2: 120
 - Router-3: 130
- Ensure that Router-3 is the current HSRP Active Router.

Configuration

```
Router-1(config)#int fast 0/1
Router-1(config-if)#standby 1 priority 110
Router-1(config-if)#standby 1 preempt

Router-1(config-if)#end
Router-1#
```

```
Router-2#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Router-2(config)#int fast 0/1
Router-2(config-if)#standby 1 priority 120
Router-2(config-if)#standby 1 preempt

Router-2(config-if)#end
```

```
Router-3
(config)#int fast 0/0 Router-3(config-if)#standby 1 priority 130
Router-3(config-if)#standby 1 preempt

Router-3(config-if)#end
Router-3#
```

Verification

```
Router-3#show standby
```

```
FastEthernet0/0 - Group 1
```

```
State is Active
```

```
22 state changes, last state change 00:00:28
```

```
Virtual IP address is 1.2.3.123
```

```
Active virtual MAC address is 0000.0c07.ac01
```

```
Local virtual MAC address is 0000.0c07.ac01 (v1 default)
```

```
Hello time 100 msec, hold time 300 msec
```

```
Next hello sent in 0.032 secs
```

```
Preemption enabled Active router is local
```

```
Standby router is 1.2.3.2, priority 120 (expires in 0.256 sec)
```

```
Priority 130 (configured 130)
```

```
Track object 1 state Up decrement 50
```

```
Group name is "hsrp-Fa0/0-1" (default)
```

```
Router-3#
```

Task

Verify that HSRP Object Tracking is functional by performing the following tasks:

- Log in to Switch-3 and disable interface FastEthernet0/13
- The above action should cause Router-3 to lose its route to 111.111.111.0/24 (the route being tracked).
 - Observe any syslogs on Router-3. Is there any indication that HSRP has decremented its priority due to the loss of the tracked object?
- Identify the current HSRP Active Router after Router-3 has lost its tracked object.

Verification

```
Switch-3(config)#interface FastEthernet0/13
```

```
Switch-3(config-if)#shut
```

```
Nov 18 11:40:16.630: %TRACK-6-STATE: 1 ip route 111.111.111.0/24 reachability Up -> Down
```

```
Router-3#
```

```
Nov 18 11:40:16.646: %HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Active -> Speak
```

```
Router-3#
Router-3#show standby
FastEthernet0/0 - Group 1
  State is Listen
    25 state changes, last state change 00:00:57
  Virtual IP address is 1.2.3.123
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 100 msec, hold time 300 msec
  Preemption enabled
  Active router is 1.2.3.2, priority 120 (expires in 0.272 sec)
  Standby router is 1.2.3.1, priority 110 (expires in 0.320 sec) Priority 80 (configured 130)

  Track object 1 state Down decrement 50
  Group name is "hsrp-Fa0/0-1" (default)
Router-3#
```


```
Router-2#show standby
FastEthernet0/1 - Group 1 State is Active
  21 state changes, last state change 00:02:32
  Virtual IP address is 1.2.3.123
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 100 msec, hold time 300 msec
    Next hello sent in 0.016 secs
  Preemption enabled Active router is local

  Standby router is 1.2.3.1, priority 110 (expires in 0.336 sec)
  Priority 120 (configured 120)
    Track object 1 state Up decrement 50
  Group name is "hsrp-Fa0/1-1" (default)
Router-2#
```

CCNP SWITCH Workbook - First Hop Redundancy Protocols

8.4 VRRP

Load the **CCNP-Switch-Task8-4** initial configurations before starting.



All of the physical connections, VLANs, IP addresses, DHCP configuration, and Routing Protocols (EIGRP) shown in the topology diagram have been pre-configured for you.

Task

In this first task, you will complete a basic VRRP configuration.

First, ensure that your current topology is up and functional (**you may have to enable some interfaces**).

- Verify that your DHCP client (R4) has obtained an IP address via DHCP.

The IP address of the default-gateway (1.2.3.123) will not be reachable yet because you have not configured VRRP on any routers.

DHCP Client Verification

```
Host#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Host(config)#int fast 0/1
Host(config-if)#shutdown
Host(config-if)#no shutdown
Host(config-if)#
Nov 17 09:38:13.531: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
Nov 17 09:38:14.531: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
Nov 17 09:38:20.639:
%DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 1.2.3.4, mask 255.255.255.0, hostname Host
```

Task

Configure VRRP on Router-1, Router-2, and Router-3 following these guidelines:

- VRRP will provide gateway redundancy for the DHCP client (R4).
- Use a VRRP Group number of one (1).
- Use a VRRP Virtual-IP address of 1.2.3.123.
- Without using any "show" commands, **can you predict** which router will be the VRRP Master router?
- Verify whether your prediction about the router playing the role of VRRP Master was correct.

VRRP Initial Configuration

Router-1#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-1(config)#**int fast 0/1**

Router-1(config-if)#**vrrp 1 ip 1.2.3.123**

Router-1(config-if)#end

Router-1#

Router-2#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-2(config)#**int fast 0/1**

Router-2(config-if)#**vrrp 1 ip 1.2.3.123**

Router-2(config-if)#end

Router-2#

Router-3#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-3(config)#**int fast 0/0**

Router-3(config-if)#**vrrp 1 ip 1.2.3.123**

Router-3(config-if)#end

Router-3#

VRRP Verification

By default, the very first router to be configured with VRRP would assume the role of the VRRP Master router.

Unlike HSRP, VRRP does have Preemption enabled by default. If a new router comes online with a higher IP address or priority than the current VRRP **Master** router, the new router will assume the role of VRRP Master.

As you configured VRRP you would have noticed that, regardless of the order in which you configured VRRP onto routers-1, 2, and 3, Router-3 would have become the VRRP Master simply by virtue of it having the highest IP address in the VRRP Group.

```

Router-3#
Nov 18 12:02:43.202: %VRRP-6-STATECHANGE: Fa0/0 Grp 1 state Init -> Backup
Router-3#
Nov 18 12:02:44.182: %SYS-5-CONFIG_I: Configured from console by console
Router-3#Nov 18 12:02:46.818: %VRRP-6-STATECHANGE: Fa0/0 Grp 1 state Backup -> Master
Router-3#
Router-3#show vrrp
FastEthernet0/0 - Group 1   State is Master

Virtual IP address is 1.2.3.123
Virtual MAC address is 0000.5e00.0101
Advertisement interval is 1.000 sec Preemption enabled

Priority is 100
Master Router is 1.2.3.3 (local), priority is 100
Master Advertisement interval is 1.000 sec
Master Down interval is 3.609 sec

Router-3#

```

Test Basic VRRP Functionality

For this task, you will need to have two Telnet windows opened simultaneously. One window should allow you to view the output of the DHCP client (R4) while the other allows you to make changes to Switch-2.

- **Begin a continuous ping** from the DHCP client (R4) to the IP address of **111.111.111.111 (repeat count of 10,000)**.
- **Shut down** the FastEthernet interface of Switch-2 that connects to your VRRP Master router.
- **Monitor** how many (if any) pings are lost (and the length of time until pings resume) until one of the other routers assumes the role of VRRP Master Router.
- How long did it take for a new VRRP Master Router to be elected and become active?


```
Host#ping 111.111.111.111 repeat 10000
Type escape sequence to abort.
Sending 10000, 100-byte ICMP Echos to 111.111.111.111, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
<output omitted for brevity> !!!!!!!!!!!!!!!...
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

In the test above, pings were stopped for about 3 seconds and then resumed.

CCNP SWITCH Workbook - First Hop Redundancy Protocols

8.5 GLBP Basic Configuration

Load the **CCNP-Switch-Task8-5** initial configurations before starting.



All of the physical connections, VLANs, IP addresses, DHCP configuration, and Routing Protocols (EIGRP) shown in the topology diagram have been pre-configured for you.

Task

In this lab, you will complete a basic GLBP configuration.

First, ensure that your current topology is up and functional (**you may have to enable some interfaces**).

- Verify that your DHCP clients (Sw1 and R2) have obtained an IP address via DHCP.

The IP address of the default-gateway (1.2.3.123) will not be reachable yet because you have not configured GLBP on any routers.

DHCP Client Verification

```
Host#
conf t
Enter configuration commands, one per line. End with CNTL/Z. Host(config)#int fast 0/1
Host(config-if)#shutdown
Host(config-if)#no shutdown
Host(config-if)#
Nov 17 09:38:13.531: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
Nov 17 09:38:14.531: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
Nov 17 09:38:20.639:
%DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address 1.2.3.5, mask 255.255.255.0, hostname Host
```

Task

Configure GLBP on Router-1, Router-3, and Router-4 following these guidelines:

- GLBP will provide gateway redundancy for the DHCP client (R4).
- Use a GLBP Group number of one (1).
- Use a GLBP Virtual-IP address of 1.2.3.123.
- Without using any "show" commands, **can you predict** which router will be the GLBP Active Virtual Gateway (AVG) router?
- Verify whether your prediction about the router playing the role of GLBP AVG was correct.

GLBP Initial Configuration

Router-1#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-1(config)#**int fast 0/1**

Router-1(config-if)#**glbp 1 ip 1.2.3.123**

Router-1(config-if)#end

Router-1#

Router-4#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-4(config)#**int fast 0/1**

Router-4(config-if)#**glbp 1 ip 1.2.3.123**

Router-4(config-if)#end

Router-4#

Router-3#

conf t

Enter configuration commands, one per line. End with CNTL/Z. Router-3(config)#**int fast 0/0**

Router-3(config-if)#**glbp 1 ip 1.2.3.123**

Router-3(config-if)#end

Router-3#

GLBP Verification

By default, the very first router to be configured with GLBP would assume the role of the GLBP AVG router.

Similar to HSRP, GLBP does not have Preemption enabled by default. After the GLBP AVG is initially elected, that router will remain as the GLBP AVG, even if a new router comes online with a higher IP address or priority than the current GLBP **AVG** router.

As you configured GLBP you would have noticed that, regardless of the order in which you configured GLBP onto routers-1, 3, and 4, the first router you configured for GLBP would have become the GLBP AVG.

In the examples below, GLBP was first configured on Router-3, followed by Router-1 and then Router-4. Notice that the first router configured (Router-3) remained the

GLBP AVG, even though Router-4 has a higher IP address.

In the output below, any time you see the term **Active**, this refers to the **Active Virtual Gateway (AVG)**. It does **NOT** mean Active Virtual Forwarder.

Router-3#show glbp

FastEthernet0/0 - Group 1 State is Active

1 state change, last state change 00:04:20

Virtual IP address is 1.2.3.123

Hello time 3 sec, hold time 10 sec

Next hello sent in 1.952 secs

Redirect time 600 sec, forwarder timeout 14400 sec

Preemption disabled Active is local

Standby is 1.2.3.4, priority 100 (expires in 8.256 sec)

Priority 100 (default)

Weighting 100 (default 100), thresholds: lower 1, upper 100

Load balancing: round-robin Group members:

0018.b9ba.6dd8 (1.2.3.3) local

001c.589e.7ae1 (1.2.3.4)

001f.ca05.eab1 (1.2.3.1)

There are 3 forwarders (1 active) Forwarder 1

State is Active

1 state change, last state change 00:04:08

MAC address is 0007.b400.0101 (default)

Owner ID is 0018.b9ba.6dd8

Redirection enabled

Preemption enabled, min delay 30 sec Active is local

, weighting 100

Forwarder 2

State is Listen

MAC address is 0007.b400.0102 (learnt)

Owner ID is 001f.ca05.eab1

Redirection enabled, 598.656 sec remaining (maximum 600 sec)

Time to live: 14398.656 sec (maximum 14400 sec)

Preemption enabled, min delay 30 sec

Active is 1.2.3.1 (primary), weighting 100 (expires in 9.376 sec)

Forwarder 3

State is Listen

MAC address is 0007.b400.0103 (learnt)

Owner ID is 001c.589e.7ae1

Redirection enabled, 598.272 sec remaining (maximum 600 sec)

Time to live: 14398.272 sec (maximum 14400 sec)

Preemption enabled, min delay 30 sec

Active is 1.2.3.4 (primary), weighting 100 (expires in 10.240 sec)

Test Basic GLBP Functionality

By default, GLBP operates on a round-robin basis. The router serving as the AVG will listen for ARP requests from hosts for the GLBP Virtual-IP address and respond with the MAC addresses of each AVF. To test this, we should see that even though Host-1 (Sw1) and Host-2 (R2) have learned of the same Default-Gateway (1.2.3.123), each host should have a different MAC address (that was learned via ARP) for its Default-Gateway.

- In Host-1 (Sw1), initiate a ping to the default-gateway address of 1.2.3.123. This ping is simply a mechanism to force Host-1 to ARP for its default-gateway.
- Have Host-2 (R2) also ping its default-gateway address of 1.2.3.123.
- Look in the ARP tables of both Host-1 and Host-2 and notice the MAC address that was learned for 1.2.3.123.
 - Did both hosts learn of the same MAC address?
 - Can you identify which AVF has been assigned to each host?

Host-1#ping 1.2.3.123

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.2.3.123, timeout is 2 seconds:

.!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

Host-1#

Host-1#sho ip arp 1.2.3.123

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
ARPA	1.2.3.123	0	0007.b400.0101		Internet

Host-1#

Host-2>

enHost-2#ping 1.2.3.123

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1.2.3.123, timeout is 2 seconds:

..!!!

Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms

Host-2#

Host-2#show ip arp 1.2.3.123

Protocol	Address	Age (min)	Hardware Addr	Type	Interface	Internet
1.2.3.123	0	0007.b400.0102				
ARPA	FastEthernet0/1					
Host-2#						

Verification of GLBP AVF Allocation

Router-3#show glbp

FastEthernet0/0 - Group 1

State is Active

1 state change, last state change 00:17:45

Virtual IP address is 1.2.3.123

Hello time 3 sec, hold time 10 sec

Next hello sent in 1.824 secs

Redirect time 600 sec, forwarder timeout 14400 sec

Preemption disabled

Active is local

Standby is 1.2.3.4, priority 100 (expires in 8.096 sec)

Priority 100 (default)

Weighting 100 (default 100), thresholds: lower 1, upper 100

Load balancing: round-robin

Group members: 0018.b9ba.6dd8 (1.2.3.3) local

001c.589e.7ae1 (1.2.3.4) 001f.ca05.eab1 (1.2.3.1)

There are 3 forwarders (1 active)

Forwarder 1

State is Active

1 state change, last state change 00:17:33

MAC address is 0007.b400.0101 (default) Owner ID is 0018.b9ba.6dd8

Redirection enabled

Preemption enabled, min delay 30 sec

Active is local, weighting 100 Client selection count: 1

Forwarder 2

State is Listen

MAC address is 0007.b400.0102 (learnt) Owner ID is 001f.ca05.eab1

Redirection enabled, 598.784 sec remaining (maximum 600 sec)

Time to live: 14398.784 sec (maximum 14400 sec)

Preemption enabled, min delay 30 sec

Active is 1.2.3.1 (primary), weighting 100 (expires in 10.624 sec) Client selection count: 1

Forwarder 3

State is Listen

MAC address is 0007.b400.0103 (learnt)

Owner ID is 001c.589e.7ae1

Redirection enabled, 598.112 sec remaining (maximum 600 sec)

Time to live: 14398.112 sec (maximum 14400 sec)

Preemption enabled, min delay 30 sec

Active is 1.2.3.4 (primary), weighting 100 (expires in 8.672 sec)

Router-3#

CCNP SWITCH Workbook - CCNP Switch Workbook Introduction

CCNP Switch Introduction

Welcome!

Thank you for using this workbook as part of your preparations for pursuing your CCNP SWITCH certification. The sections and tasks within this workbook are designed to give you hands-on experience with the majority of topics defined as "**Configure and verify**" within the CCNP SWITCH version 2.0 blueprint.


Although it is advisable to begin with the first task in each section and then work progressively through that section, the individual tasks were designed in such a way that you can start with any task you wish without following any specific order. After you have downloaded the initial configurations for a task, you may begin working on that task, even if you have not completed the tasks that preceded it.

Diagrams

There is one main diagram supplied with this workbook that should be used to give you a complete understanding of the network topology. Often, you will find that there are individual diagrams for each section, but these are all permutations of the main diagram shown below.

Aside from the links shown below, there are other links (not displayed in the topology diagram) that lead to other devices not used in this workbook. To prevent unexpected behavior, **it is always recommended that you shut down all links on all three switches as your first step, and then enable only those links displayed in the topology for any given task.**

Also, please remember that interfaces on routers are administratively disabled by default. So for most tasks that utilize routers, you will need to administratively enable any of their interfaces that are used for any given task.



Feedback

Please let us know how we're doing! In the upper-right corner of your screen, you will see a **Feedback** link. If you found any errors in this workbook or have any suggestions for improvement, we'd like to know. Also, if you enjoyed this workbook, we'd like to know that as well.

CCNP SWITCH Workbook - CCNP Switch Workbook Introduction

CCNA/CCNP Rack Rental Guide

[Click here to access the CCNA/CCNP Rack Rental Guide.](#)