CCIE R&S v5 Lab Workbook PDF Section 2

Hello, everyone. Today, BestCiscoDumps continues to share CCIE R&S v5 lab workbook PDF section 2. Most of the basic configurations in Section 1 are for switching, and some routing related configurations will be configured from Section 2.1. These configurations and ideas of problem solutions are often seen in the daily work of network engineers. Therefore, I hope you can read it several times and understand the principle.

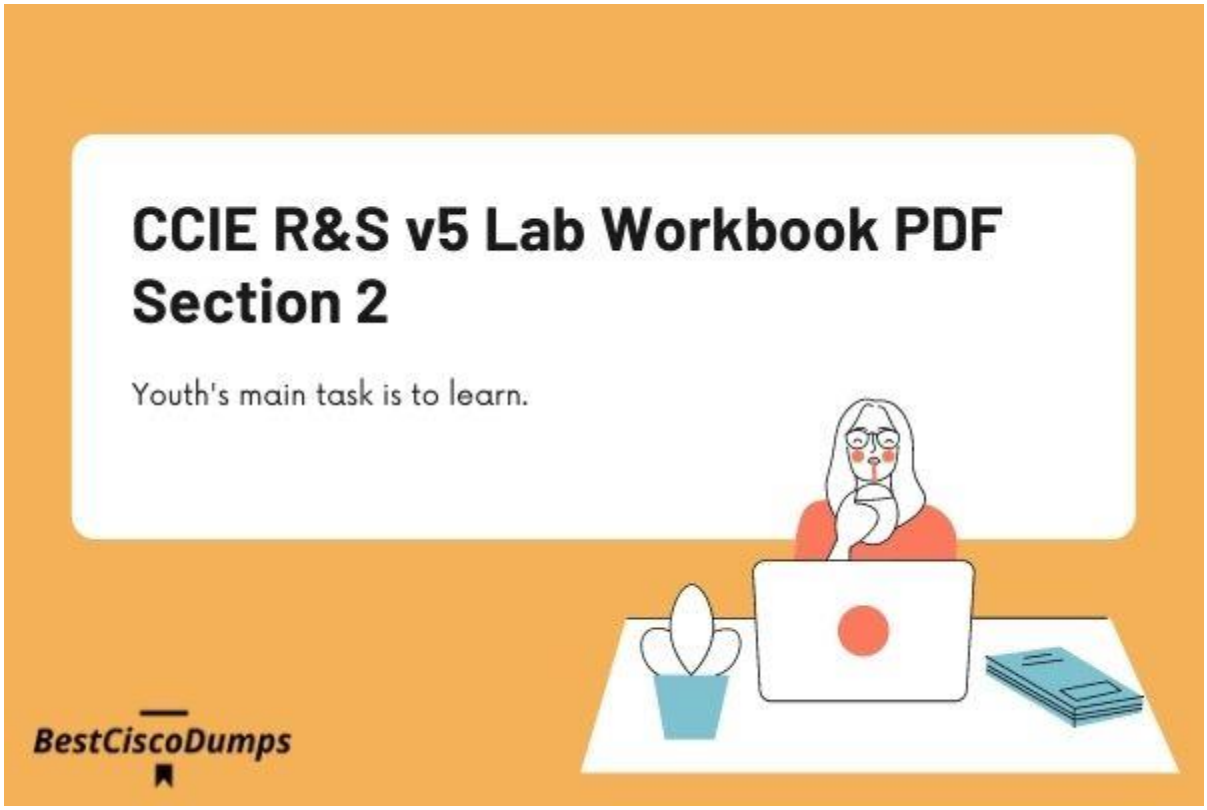

Click the link to view the CCIE R&S Lab Topology and CCIE R&S v5 Lab Workbook PDF Section 1
The current page is CCIE R&S v5 Lab Workbook PDF Section 2
Click the link to view the CCIE R&S v5 Lab Workbook PDF Section 3-5

SECTION 2 – Layer 3 Technologies

SECTION 2.1: OSPF in AS 12345

Configure OSPFv2 Area 0 in the ACME Headquarters (AS 12345) according to the following requirements

- Configure the OSPF process ID to 12345 and set the router-id interface lo0 on all seven routers
- The interface lo0 of each router must be seen as an internal OSPF prefix by all other routers
- Ensure that OSPF is not running on any interface that is facing another AS. use any method to accomplish this requirement
- SW1 and SW2 must not participate in routing at all
- Do not change the default OSPF cost of any interface in AS 12345
- R1 must see the following OSPF routes in its routing table

Solution:

R1:

router ospf 12345

router-id 123.1.1.1

network 123.1.1.1 0.0.0.0 area 0

network 123.10.1.2 0.0.0.0 area 0

network 123.10.1.5 0.0.0.0 area 0

R2:

router ospf 12345

router-id 123.2.2.2

network 123.2.2.2 0.0.0.0 area 0

network 123.10.1.9 0.0.0.0 area 0

network 123.10.1.17 0.0.0.0 area 0

R3:

router ospf 12345

router-id 123.3.3.3

network 123.3.3.3 0.0.0.0 area 0

network 123.10.1.10 0.0.0.0 area 0

network 123.10.1.13 0.0.0.0 area 0

R4:

router ospf 12345

router-id 123.4.4.4

network 123.4.4.4 0.0.0.0 area 0

network 123.10.1.18 0.0.0.0 area 0

network 123.10.1.21 0.0.0.0 area 0

network 123.10.1.1 0.0.0.0 area 0

R5:

router ospf 12345

router-id 123.5.5.5

network 123.5.5.5 0.0.0.0 area 0

network 123.10.1.14 0.0.0.0 area 0

network 123.10.1.6 0.0.0.0 area 0

network 123.10.1.29 0.0.0.0 area 0

R6:

router ospf 12345

router-id 123.6.6.6

network 123.6.6.6 0.0.0.0 area 0

network 123.10.1.22 0.0.0.0 area 0

network 123.10.1.25 0.0.0.0 area 0

R7:

router ospf 12345

router-id 123.7.7.7

network 123.7.7.7 0.0.0.0 area 0

network 123.10.1.26 0.0.0.0 area 0

network 123.10.1.30 0.0.0.0 area 0

2.2 EIGRP in AS 34567

Configure EIGRP for ipv4 in the New York office (AS 34567) according to the following requirements

- The EIGRP autonomous system number is 34567
- Do not configure any EIGRP virtual instance name in EIGRP AS 34567
- The interface lo0 of each router must be seen as an internal EIGRP prefix by all other routers
- Ensure that EIGRP is not running on any interface that is facing another AS , use any method to accomplish this requirement
- Using a single command in one switch only , ensure that R8 installs two equal-cost paths for the following three prefixes:

vlan 411

interface Lo0 of SW4 interface Lo0 of R11

- Using a single command in one switch only , ensure that R9 installs two equal-cost paths for the following three prefixes:

vlan 310

interface Lo0 of SW3 interface Lo0 of R10

Solution:

R8:

router eigrp 34567 //Automatic aggregation in lasted IOS is disabled by Default network 123.8.8.8 0.0.0.0

network 123.10.2.1 0.0.0.0

network 123.10.2.5 0.0.0.0

R9:

router eigrp 34567

network 123.9.9.9 0.0.0.0

network 123.10.2.2 0.0.0.0

network 123.10.2.9 0.0.0.0

R10:

router eigrp 34567

network 123.10.10.10 0.0.0.0

network 123.10.2.18 0.0.0.0

network 123.10.2.25 0.0.0.0

R11:

router eigrp 34567

network 123.11.11.11 0.0.0.0

network 123.10.2.22 0.0.0.0

network 123.10.2.26 0.0.0.0

SW3:

router eigrp 34567

network 123.33.33.33 0.0.0.0

network 123.10.2.6 0.0.0.0

network 123.10.2.13 0.0.0.0

network 123.10.2.17 0.0.0.0

SW4:

router eigrp 34567

network 123.44.44.44 0.0.0.0

network 123.10.2.10 0.0.0.0

network 123.10.2.14 0.0.0.0

network 123.10.2.21 0.0.0.0

SW3/SW4:

interface vlan 34

delay 100

Note:

The bandwidth of SW3's SVI34 interface is 1G, which is different from that of R9's e0/1 bandwidth is 10M.

R8 to R11 with only one S and one FS

There is only one S (the next hop is SW4) to lo0 of SW3 destined for R11, and there is no FS. The next hop to R10 does not match the FC

When metric is calculated, the interface bandwidth is the smallest one involved in the calculation. All of this is not related to the bandwidth, only the parameters that change the delay.

The delay unit in the interface is 10 times us, so it is changed to 100, that is, 1000us.

SECTION 2.3: EIGRP in AS 45678

Configure EIGRP for ipv4 in the Sydney office (BGP AS 45678) according to the following requirements

- The EIGRP autonomous system number is 45678
- The interface lo0 of each device must be seen as an internal EIGRP prefix by all other routers
- Ensure the EIGRP is not running on any interface that is facing another AS, use any method to accomplish this requirement
- Sw5 and Sw6 are Layer 3 switches and must also run EIGRP on all three routers (R15, R16and R17), ensure that EIGRP uses 64 bits for the metric calculation of any prefix
- Do not change the default bandwidth or delay on any physical links in AS 45678

Note: R15, R16, and R17 use the named 64-bit eigrp configuration mode. SW5 and SW6 use the normal eigrp configuration mode.

Solution:

R15:

router eigrp CCIE

address-family ipv4 unicast autonomous-system 45678

network 123.15.15.15 0.0.0.0

network 123.20.1.1 0.0.0.0

network 123.20.1.9 0.0.0.0

R16:

router eigrp CCIE

address-family ipv4 unicast autonomous-system 45678

network 123.16.16.16 0.0.0.0

network 123.20.1.2 0.0.0.0

network 123.20.1.17 0.0.0.0

R17:

router eigrp CCIE

address-family ipv4 unicast autonomous-system 45678

network 123.17.17.17 0.0.0.0

network 123.20.1.10 0.0.0.0

network 123.20.1.18 0.0.0.0

SW5:

router eigrp 45678

network 123.55.55.55 0.0.0.0

network 123.20.1.3 0.0.0.0

SW6:

router eigrp 45678

network 123.66.66.66 0.0.0.0

network 123.20.1.11 0.0.0.0

SECTION 2.4: EIGRP in AS 65222

Configure EIGRP for ipv4 in the Tokyo and Singapore offices (BGP AS 65222) according to the following requirements

- The EIGRP autonomous system number is 45678
- The interface lo0 of each router belonging to EIGRP AS 45678 must be seen as an internal EIGRP prefix by all routers in AS 65222
- Ensure that EIGRP is not running on any physical interface that is facing another AS , use any method to accomplish this requirement
- The LAN segment connected to interface E0/0 on both R18 and R19 must be advertised into EIGRP as an internal prefix
- R17 must establish an EIGRP peering with both R18 and R19 via the same interfaceTunnel0
- R17 will not send query to R18 and R19

Note: R18, R19 do not need to use 64bit eigrp to configure

Solution:

R18:

router eigrp 45678

network 123.18.18.18 0.0.0.0

network 10.2.18.0 0.0.0.255

eigrp stub

R19:

router eigrp 45678

network 123.19.19.19 0.0.0.0

network 10.2.19.0 0.0.0.255

eigrp stub

SECTION 2.5: BGP in AS 12345

BGP is partially preconfigured in ACME Headquarters. Complete the configuration as required.

Configure IBGP in ACME's headquarters (AS 12345) according to the following requirements

- R4 and R5 must not establish any BGP session at any time
- All BGP routers must use their interface lo0 as the BGP router ID
- Disable the default ipv4 unicast address family for peering session establishment in all BGP routers
- R1 must be the ipv4 route reflector for BGP AS 12345
- R1 must use the peer-group named 'iBGP' for all internal peerings

Solution:

IBGP neighbor configuration in AS 12345

R1:

router bgp 12345

bgp router-id 123.1.1.1

no bgp default ipv4-unicast

neighbor iBGP peer-group // Note that i is lowercase

neighbor iBGP remote-as 12345

neighbor iBGP update-source Loopback0

neighbor 123.2.2.2 peer-group iBGP

neighbor 123.3.3.3 peer-group iBGP

neighbor 123.6.6.6 peer-group iBGP

neighbor 123.7.7.7 peer-group iBGP

address-family ipv4

neighbor 123.2.2.2 activate

neighbor 123.3.3.3 activate

neighbor 123.6.6.6 activate

neighbor 123.7.7.7 activate

neighbor IBGP route-reflector-client

R2:

router bgp 12345

bgp router-id 123.2.2.2

no bgp default ipv4-unicast

neighbor 123.1.1.1 remote-as 12345

neighbor 123.1.1.1 update-source Loopback0

address-family ipv4

neighbor 123.1.1.1 activate

R3:

router bgp 12345

bgp router-id 123.3.3.3

no bgp default ipv4-unicast

neighbor 123.1.1.1 remote-as 12345

neighbor 123.1.1.1 update-source Loopback0

address-family ipv4

neighbor 123.1.1.1 activate

R6:

router bgp 12345

bgp router-id 123.6.6.6

no bgp default ipv4-unicast

neighbor 123.1.1.1 remote-as 12345

neighbor 123.1.1.1 update-source Loopback0

address-family ipv4

neighbor 123.1.1.1 activate

R7:

router bgp 12345

bgp router-id 123.7.7.7

no bgp default ipv4-unicast

neighbor 123.1.1.1 remote-as 12345

neighbor 123.1.1.1 update-source Loopback0

address-family ipv4

neighbor 123.1.1.1 activate

Configure EBGP between ACME's San Francisco and San Jose sites according to the following requirements

- R20 is a CE router and uses EBGP to connect to the managed services that are provided by the PE routers R2 and R3
- R20 must establish a separate EBGP peering with both R2 and R3 for every VRF
- R20 must advertise the following prefix to all of its BGP peers
  - 123.0.0.0/8, summary-only
  - 10.0.0.0/8, summary-only
- R20 must advertise a default route to all of its BGP peers except to 10.120.99.1 and10.120.99.5

Solution:

EBGP neighbor configuration in AS 12345

R2: TO-R20

router bgp 12345

address-family ipv4 vrf GREEN

neighbor 10.120.12.2 remote-as 65112

address-family ipv4 vrf BLUE

neighbor 10.120.13.2 remote-as 65112

address-family ipv4 vrf RED

neighbor 10.120.14.2 remote-as 65112

address-family ipv4 vrf YELLOW

neighbor 10.120.15.2 remote-as 65112

address-family ipv4 vrf INET

neighbor 10.120.99.2 remote-as 65112

R3: TO-R20

router bgp 12345

address-family ipv4 vrf GREEN

neighbor 10.120.12.6 remote-as 65112

address-family ipv4 vrf BLUE

neighbor 10.120.13.6 remote-as 65112

address-family ipv4 vrf RED

neighbor 10.120.14.6 remote-as 65112

address-family ipv4 vrf YELLOW

neighbor 10.120.15.6 remote-as 65112

address-family ipv4 vrf INET

neighbor 10.120.99.6 remote-as 65112

R20: TO-R2/R3

router bgp 65112

bgp router-id 123.20.20.20

neighbor 10.120.12.1 remote-as 12345

neighbor 10.120.12.1 default-originate

neighbor 10.120.13.1 remote-as 12345

neighbor 10.120.13.1 default-originate

neighbor 10.120.14.1 remote-as 12345

neighbor 10.120.14.1 default-originate

neighbor 10.120.15.1 remote-as 12345

neighbor 10.120.15.1 default-originate

neighbor 10.120.99.1 remote-as 12345

neighbor 10.120.12.5 remote-as 12345

neighbor 10.120.12.5 default-originate

neighbor 10.120.13.5 remote-as 12345

neighbor 10.120.13.5 default-originate

neighbor 10.120.14.5 remote-as 12345

neighbor 10.120.14.5 default-originate

neighbor 10.120.15.5 remote-as 12345

neighbor 10.120.15.5 default-originate

neighbor 10.120.99.5 remote-as 12345

Config summary on R20:

R20:

router bgp 65112

network 10.0.0.0 mask 255.0.0.0

network 123.0.0.0 mask 255.0.0.0

auto-summary

Other EBGP neighbor relationships to be established:

R2: TO-ISP1

router bgp 12345

address-family ipv4 vrf BLUE

neighbor 101.1.123.1 remote-as 10001

address-family ipv4 vrf GREEN

neighbor 101.1.123.1 remote-as 10001

address-family ipv4 vrf INET

neighbor 101.1.123.1 remote-as 10001

address-family ipv4 vrf RED

neighbor 101.1.123.1 remote-as 10001

address-family ipv4 vrf YELLOW

neighbor 101.1.123.1 remote-as 10001

R3: TO-ISP2

router bgp 12345

address-family ipv4 vrf BLUE

neighbor 102.2.123.1 remote-as 10002

address-family ipv4 vrf GREEN

neighbor 102.2.123.1 remote-as 10002

address-family ipv4 vrf INET

neighbor 102.2.123.1 remote-as 10002

address-family ipv4 vrf RED

neighbor 102.2.123.1 remote-as 10002

address-family ipv4 vrf YELLOW

neighbor 102.2.123.1 remote-as 10002

R6: TO-ISP4

router bgp 12345

address-family ipv4 vrf BLUE

neighbor 201.1.123.1 remote-as 20001

address-family ipv4 vrf GREEN

neighbor 201.1.123.1 remote-as 20001

address-family ipv4 vrf INET

neighbor 201.1.123.1 remote-as 20001

R7: TO-ISP5

router bgp 12345

address-family ipv4 vrf BLUE

neighbor 202.2.123.1 remote-as 20002

address-family ipv4 vrf INET

neighbor 202.2.123.1 remote-as 20002

address-family ipv4 vrf RED

neighbor 202.2.123.1 remote-as 20002

Note the following pre-configuration in the simulator:

ISP1: TO-ISP2

The following sub-interface configuration is missing

interface Ethernet0/3.99

encapsulation dot1Q 99

ip vrf forwarding INET

ip address 100.1.2.1 255.255.255.252

ISP2: This interface belongs to the VRF INET

int e0/0

ip vrf forwarding INET

ip address 33.10.2.1 255.255.255.252

SECTION 2.6: BGP in AS 34567

BGP is partially preconfigured in ACME New York office (AS 34567). Complete the configuration as required

Configure IBGP in AS 34567 according to the following requirements

- SW3 and SW4 must not establish any BGP session at any time
- All four BGP routers must always use their interface Lo0 as the BGP router ID
- Disable the default ipv4 unicast address family for peering session establishment in all BGP routers
- Configure full-mesh IBGP peerings between all four routers, use any configuration method
- R9 must be selected as the preferred exit point for traffic destined to remote AS's
- R11 must selected as the next preferred exit point in case R9 fails
- No BGP speaker in AS 34567 may use the ‘network’ statement under the BGP router configuration
- Ensure that the BGP next-hop is never marked as unreachable as long as the interface Lo0of the remote peer is known via IGP

Solution:

IBGP neighbor configuration in AS 34567

R8:

router bgp 34567

bgp router-id 123.8.8.8

neighbor 123.9.9.9 remote-as 34567

neighbor 123.9.9.9 update-source Loopback0

neighbor 123.10.10.10 remote-as 34567

neighbor 123.10.10.10 update-source Loopback0

neighbor 123.11.11.11 remote-as 34567

neighbor 123.11.11.11 update-source Loopback0

neighbor 123.9.9.9 next-hop-self

neighbor 123.10.10.10 next-hop-self

neighbor 123.11.11.11 next-hop-self

no bgp default ipv4-unicast

R9:

router bgp 34567

bgp router-id 123.9.9.9

neighbor 123.8.8.8 remote-as 34567

neighbor 123.8.8.8 update-source Loopback0

neighbor 123.10.10.10 remote-as 34567

neighbor 123.10.10.10 update-source Loopback0

neighbor 123.11.11.11 remote-as 34567

neighbor 123.11.11.11 update-source Loopback0

neighbor 123.8.8.8 next-hop-self

neighbor 123.10.10.10 next-hop-self

neighbor 123.11.11.11 next-hop-self

no bgp default ipv4-unicast

router bgp 34567

bgp router-id 123.10.10.10

neighbor 123.8.8.8 remote-as 34567

neighbor 123.8.8.8 update-source Loopback0

neighbor 123.9.9.9 remote-as 34567

neighbor 123.9.9.9 update-source Loopback0

neighbor 123.11.11.11 remote-as 34567

neighbor 123.11.11.11 update-source Loopback0

neighbor 123.9.9.9 next-hop-self

neighbor 123.8.8.8 next-hop-self

neighbor 123.11.11.11 next-hop-self

no bgp default ipv4-unicast

R11:

router bgp 34567

bgp router-id 123.11.11.11

neighbor 123.8.8.8 remote-as 34567

neighbor 123.8.8.8 update-source Loopback0

neighbor 123.9.9.9 remote-as 34567

neighbor 123.9.9.9 update-source Loopback0

neighbor 123.10.10.10 remote-as 34567

neighbor 123.10.10.10 update-source Loopback0

neighbor 123.8.8.8 next-hop-self

neighbor 123.9.9.9 next-hop-self

neighbor 123.10.10.10 next-hop-self

no bgp default ipv4-unicast

EBGP neighbor configuration in AS 34567

R8:

router bgp 34567

neighbor 101.1.34.1 remote-as 10001

R9:

router bgp 34567

neighbor 102.1.34.1 remote-as 10002

neighbor 33.34.4.1 remote-as 30000

R10:

router bgp 34567

neighbor 201.1.34.1 remote-as 20001

neighbor 201.1.34.1 pre 61 in

ip pre 61 deny 61.61.61.61/32

ip pre 61 per 0.0.0.0/0 le 32

New requirements in the examination room: 61.61.61.61/32 routes need to be filtered

R11:

router bgp 34567

neighbor 33.34.3.1 remote-as 30000

neighbor 202.2.34.1 remote-as 20002

R9:

router bgp 34567

bgp default local-preference 200

R11:

router bgp 34567

bgp default local-preference 110 //Backup

clear ip bgp * soft

- All four BGP routers redistribute EIGRP into BGP
- R9 and R11 must redistribute only the BGP default route into EIGRP

- Ensure that R9 is the only router that sees the default route as a BGP route and that all other routers (R8, R10, R11) see it as an EIGRP external route

Solution:

R8/R9/R10/R11:

router bgp 34567

redistribute eigrp 34567

R9 /R11:

router eigrp 34567

redistribute bgp 34567 metric 1000 100 255 1 1500 route-map DEFAULT

route-map DEFAULT permit 10

match ip address prefix-list DEFAULT

ip prefix-list DEFAULT permit 0.0.0.0/0

For the default route, the routing table entry of each router has met the last requirement by default:

Because the EBGP route (AD=20) is loaded in the routing table of R9, BGP entries are presented.

In the BGP table of R8, R10, and R11, the IBGP route is preferred (the next hop is 123.9.9.9, AD=200), and the AD of eigrp is 170. Therefore, the external D EX table entry of eigrp is loaded.

SECTION 2.7: BGP in AS 45678 and AS 65222

Configure EBGP in the ACME APAC region (AS 45678 and AS 65222) according to the following requirements

- SW5 and SW6 must not establish any BGP session at any time
- All BGP routers must always use their interface Lo0 as the BGP router ID

- No IBGP peering sessions are allowed in AS 45678
- R15 must establish an EBGP peering with AS 10003 and must receive a default route as well as other prefixes.
- R15 must redistribute BGP into EIGRP and vice-versa
- R15 must also advertise an aggregate prefix 123.20.1.0/24 to AS 10003 and must suppress all component prefixes
- R16, R17, R18 and R19 must establish an EBGP peering with AS 20003 and must receive a default route as well as other prefixes
- R16, R17, R18 and R19 must not advertise any prefix to AS 20003
- As long as R15 is operational, R16, R17, R18 and R19 must prefer the EIGRP default route over the EBGP default route
- Do not create any VRF anywhere in order to accomplish the above requirements
- Be aware that the completion of this task is contingent on the completion of some other tasks

Solution:

R15:

router bgp 45678

bgp router-id 123.15.15.15

neighbor 103.2.45.1 remote-as 10003

redistribute eigrp 45678

aggregate-address 123.20.1.0 255.255.255.0 summary-only

router eigrp CCIE

address-family ipv4 unicast autonomous-system 45678

topology base

redistribute bgp 45678 metric 1000 100 255 1 1500

R16:

router bgp 45678

bgp router-id 123.16.16.16

neighbor 203.3.16.1 remote-as 20003

R17:

router bgp 45678

bgp router-id 123.17.17.17

neighbor 203.3.17.1 remote-as 20003

R18:

router bgp 65222

bgp router-id 123.18.18.18

neighbor 203.3.18.1 remote-as 20003

R19:

router bgp 65222

bgp router-id 123.19.19.19

neighbor 203.3.19.1 remote-as 20003

R16/R17:

router bgp 45678

network 0.0.0.0 backdoor

Note: The 0.0.0.0 route of the DEX in the R16 and R17 routing tables is normally added to the table before it can be passed to the R18 and R19.

R18/R19:

router bgp 65222

network 0.0.0.0 backdoor

SECTION 2.8: BGP routing policies

Configure the ACME network as per the following requirements

- All ACME border routers in AS 12345 must filter the BGP prefixes that are advertised to their SP in VRF INET and must allow only all prefixes that belong to the class A 123.0.0.0/8 and all other VRF's must propagate all prefix
- All ACME border routers in AS 34567 must filter the BGP prefixes that are advertised to their SP and must allow only all prefixes that belong to the class A 123.0.0.0/8
- Do not use any route-map or access-list to accomplish the above requirements
- R13 must route traffic preferably via AS 20002, use any method to accomplish this requirement
- All three remote sites in AS 65111 must be able to ping 1.2.3.4 and traceroute must reveal the exact same path as shown in the following output

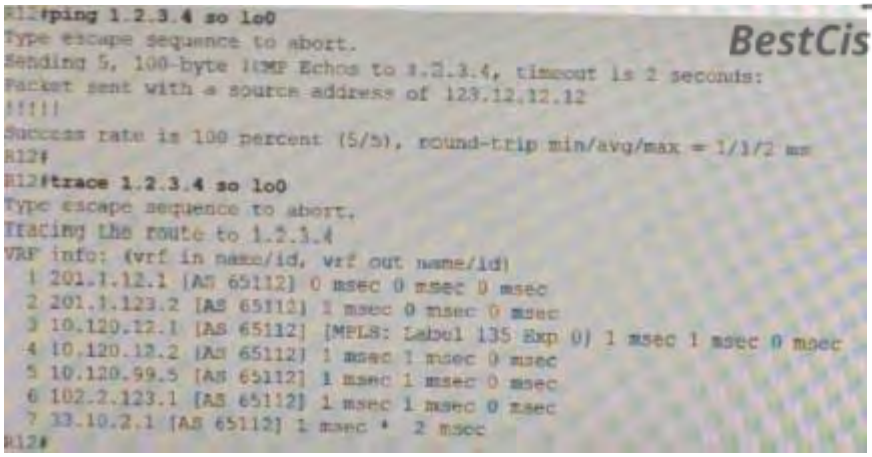

Solution:

INET VRF Based BGP Route Filtering in AS 12345

R2:

router bgp 12345

address-family ipv4 vrf INET

neighbor 101.1.123.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

R3:

router bgp 12345

address-family ipv4 vrf INET

neighbor 102.2.123.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

R6:

router bgp 12345

address-family ipv4 vrf INET

neighbor 201.1.123.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

R7:

router bgp 12345

address-family ipv4 vrf INET

neighbor 202.2.123.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

BGP Route Filtering Based on Global Routing Table in AS 34567

R8:

router bgp 34567 address-family ipv4

neighbor 101.1.34.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

R9:

router bgp 34567 address-family ipv4

neighbor 102.1.34.1 prefix-list FILTER out

neighbor 33.34.4.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

R10:

router bgp 34567

address-family ipv4

neighbor 201.1.34.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

R11:

router bgp 34567 address-family ipv4

neighbor 202.1.34.1 prefix-list FILTER out

neighbor 33.34.3.1 prefix-list FILTER out

ip prefix-list FILTER permit 123.0.0.0/8 le 32

EBGP neighbor configuration in AS 65111

R12: TO- ISP4

router bgp 65111

bgp router-id 123.12.12.12

neighbor 201.1.12.1 remote-as 20001

redistribute connect

R13: TO-ISP4/ISP5

router bgp 65111

bgp router-id 123.13.13.13

neighbor 201.1.13.1 remote-as 20001

neighbor 202.2.13.1 remote-as 20002

neighbor 202.2.13.1 weight 1000 // R13 preferred AS 20002 for export

redistribute connected

R14: TO- ISP5

router bgp 65111

bgp router-id 123.14.14.14

neighbor 202.2.14.1 remote-as 20002

redistribute connected

R20:

router bgp 65112

neighbor 10.120.99.5 weight 1000 //Make the next hop of R12, R13, R14
trace 1.2.3.4 source lo must be e0/3 of R3

SECTION 2.9: OSPFV3

Refer to “Diagram 5: IPv6 Topology”

ACME Corp is considering deploying IPv6 and wants to start implementing OSPFv3 in three sites only: New York, San Diego and Miami Both regional service providers offer IPv6 services via xx Diego and Miami) by offering two sub interfaces, one for IPv4’s managed services and one for IPv6

Configure OSPFv3 in the ACME New York office as per the following requirements

- Configure the ospf process ID to 1 and set the router-id to interface Lo0 on all seven routers
- Do not enable OSPFv3 on any interfaces other than the interfaces that are indicated on the “Diagram 5: IPv6 Topology”
- Place the interfaces in the OSPFv3 area as indicated on the diagram. Do not create any new area. No other interface may be included in OSPFv3. Do not forget the Lo0 interfaces
- SW4 must be selected as the designated router on VLAN 34 and must have the best chance of retaining this role even if a new OSPFv3 device added to the VLAN in the future
- SW3 must be selected as the backup designated router on VLAN 34 and must take over the designated router role if SW4 is down

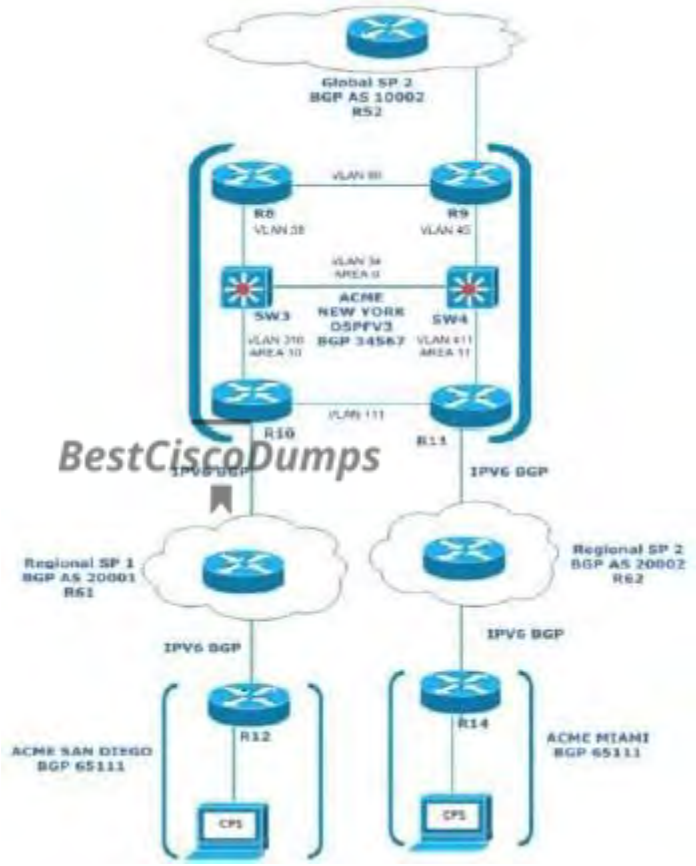

Solution:

SW3:

ipv6 unicast-routing

ipv6 cef //default config

ipv6 router ospf 1

router-id 123.33.33.33

int lo0

ipv6 enable //In the simulator, there is no ipv6 address under this interface. Add this command

ipv6 ospf 1 area 0

int vlan 34

ipv6 ospf priority 254

ipv6 ospf 1 area 0

int vlan 310

ipv6 ospf 1 area 10

SW4:

ipv6 unicast-routing

ipv6 cef

ipv6 router ospf 1

router-id 123.44.44.44

int lo0

ipv6 enable

ipv6 ospf 1 area 0

int vlan 34

ipv6 ospf priority 255

ipv6 ospf 1 area 0

int vlan 411

ipv6 ospf 1 area 11

R10:

ipv6 unicast-routing

ipv6 cef

ipv6 router ospf 1

router-id 123.10.10.10

int lo0

ipv6 enable

ipv6 ospf 1 area 10

int e0/2

ipv6 ospf 1 area 10

R11:

ipv6 unicast-routing

ipv6 cef

ipv6 router ospf 1

router-id 123.11.11.11

int lo0

ipv6 enable

ipv6 ospf 1 area 11

int e0/1

ipv6 ospf 1 area 11

SECTION 2.10: BGP for IPv6

Refer to “Diagram 5: IPv6 Topology”

Configure the ACME network as per the following requirements

- Establish the four EBGP peering as indicated on ‘Diagram 5: IPV6 routing’
- Do not use the network command under the BGP address-family ipv6 on either R10 or R11
- Both regional service providers will advertise the necessary prefixes
- Advertise the IPv6 prefix of interface E0/0 into BGP on both R12 and R14
- R12 and R14 can ping each other’s E0/0

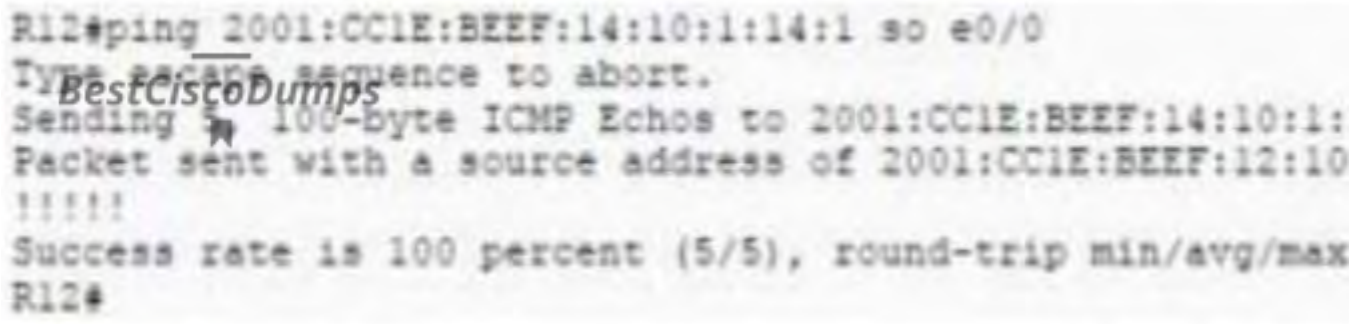

Solution:

R10:

router bgp 34567

neighbor 2001:CC1E:BEF:10:201:1:34:1 remote-as 20001

address-family ipv6

neighbor 2001:CC1E:BEF:10:201:1:34:1 activate // Must be manually activated under the ipv6 address family

redistribute ospf 1 match internal external include-connected

ipv6 router ospf 1

redistribute bgp 34567

R11:

router bgp 34567

neighbor 2001:CC1E:BEF:11:202:2:34:1 remote-as 20002

address-family ipv6

neighbor 2001:CC1E:BEF:11:202:2:34:1 activate

redistribute ospf 1 match internal external include-connected

ipv6 router ospf 1 redistribute bgp 34567

R12:

ipv6 unicast-routing

router bgp 65111

neighbor 2001:CC1E:BEF:12:201:1:12:1 remote-as 20001

address-family ipv6

neighbor 2001:CC1E:BEF:12:201:1:12:1 activate

redistribute connected

R14:

ipv6 unicast-routing

router bgp 65111

neighbor 2001:CC1E:BEF:14:202:2:14:1 remote-as 20002

address-family ipv6

neighbor 2001:CC1E:BEF:14:202:2:14:1 activate

redistribute connected

show bgp ipv6 unicast summary //Check ipv6-BGP neighbors

show bgp ipv6 unicast //Check ipv6-BGP neighbors

SECTION 2.11: Layer 3 Multicast

- A Streaming server is connected to vlan 5 on sw5 and that receivers at the DMVPN spokes, bebind R18 and R19
- The ACME APAC region network as per the following requirements
- Only network segments with Active receiver that explicitly requested the data mustreceive multicast traffic
- The loopback 0 of R15 must be configured as the Rendezvous Point
- A standard method of dynamically distributing electing RP
- R16 and R17 must participate in the multicast routing
- Testing purpose, configure interface E0/0 of both R18 and R19 join to the group 232.1.1.1
- Put unused port on sw5 into vlan 5 and confirm that multicast is working as required byusing the following test SW5 must receive from both R18 and R19


```
SW5#Ping 232.1.1.1 so vl 5
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 232.1.1
Reply to request 1 from 10.2.18.1, 84 ms
Reply to request 2 from 10.2.19.1, 72 ms
```


New demand:

Put an unused interface into VLAN 5 on the SW5.

In the pre-configuration, the SVI5 port of SW5 is down because no interface is assigned to vlan5.

Multicast source: SW5' SVI5--123.55.55.0/24

Receiver: R18 and R19's e0/0

RP: R15's lo0--123.15.15.15/32

Solution:

Prerequisite: All multicast routers have unicast routes to source and RP address segments.

R15:

ip multicast-routing

int lo0

ip pim sparse-mode

int e0/1

ip pim sparse-mode

int e0/2

ip pim sparse-mode

ip pim rp-candidate loopback 0

ip pim bsr-candidate loopback 0

SW5:

ip multicast-routing

int vlan 5

ip pim sparse-mode

int vlan 55

ip pim sparse-mode

interface range [unusedinterface]

switchport mode access

switchport access vlan 5

SW6:

ip multicast-routing

int vlan 6

ip pim sparse-mode

int vlan 66

ip pim sparse-mode

R16:

ip multicast-routing

int lo0

ip pim sparse-mode //could not be configured

int e0/1

ip pim sparse-mode

int e0/2

ip pim sparse-mode

R17:

ip multicast-routing

int lo0

ip pim sparse-mode //could not be configured

int e0/1

ip pim sparse-mode

int e0/2

ip pim sparse-mode

int tun 0

ip pim sparse-mode

R18:

ip multicast-routing

int lo0

ip pim sparse-mode //could not be configured

int tun 0

ip pim sparse-mode

int e0/0

ip pim sparse-mode

ip igmp join-group 232.1.1.1

R19:

ip multicast-routing

int lo0

ip pim sparse-mode //could not be configured

int tun 0

ip pim sparse-mode

int e0/0

ip pim sparse-mode

ip igmp join-group 232.1.1.1

show ip pim rp //Check active RP information

Note: The route table of R18 and R19 goes to 123.15.15.15/32 and 123.55.55.0/24. It cannot be a BGP entry. The interface IP (10.18.19.1) whose route next hop should be directly to tun0 of R17.

Well, today, CCIE R&S v5 lab workbook PDF section 2 will be shared here. All network engineers must read it several times. Only if you learn the principles of these solutions will the content we share be meaningful. BestCiscoDumps sincerely hopes that every Cisco test candidate can successfully pass the test with our CCIE Lab Dumps help!