

CCIE Enterprise Infrastructure Lab Test Section 1.10-16



Hello, last time we have talked about CCIE enterprise infrastructure lab section 1.1-1.9 so many students asked us to share the follow-up part. Therefore, BestCiscoDumps continues to bring follow-up sections to all CCIE EI lab candidates. In order to ensure that the examination candidates will not reduce their learning efficiency due to too much content, we only talk about CCIE EI lab section 1.10-1.16 today.

According to the Convention, since there are many pictures to verify the answers, students who need to verify the pictures can contact us to obtain a more complete CCIE EI lab workbook. If you want to use real remote rack to practice lab, you can also contact us through the contact information on the website!

Now let's get to the point.

Click the link to view the complete [CCIE Enterprise Infrastructure Lab topology](#)

Click the link to view the [CCIE Enterprise Infrastructure Lab Section 1.1-1.9](#)

The current page is CCIE Enterprise Infrastructure Lab Section 1.10-1.16
Click the link to view the [CCIE Enterprise Infrastructure Lab Section 2.1-2.3](#)
Click the link to view the [CCIE Enterprise Infrastructure Lab Section 2.4-2.6](#)
Click the link to view the [CCIE Enterprise Infrastructure Lab Section 3.1-3.4](#)

SECTION 1.10: Bringing up VPNv4/VPNv6 in SP#1

Configure routers r3, r4, r5 and r6 in SP#1 according to these requirements:

1. Configure r3 through r6 for mutual VPNv4 and VPNv6 route exchange without the use of a route-reflector. Use Lo0 IPv4 addresses for peerings.
2. Configure r3 through r6 to assign (allocate/bind) as few unique MPLS labels to all existing and future VPNv4 and VPNv6 routes as possible.
3. On routers r3 through r6, prevent any existing and future customer from discovering details about the inner topology of SP#1, It is not allowed to use ACLs to accomplish this requirement.

Solution

On r3

```
r3(config)#router bgp 10000
```

```
r3(config-router)#neighbor 100.255.254.4 remote-as 10000
```

```
r3(config-router)#neighbor 100.255.254.4 update-source Loopback0
```

```
r3(config-router)#neighbor 100.255.254.5 remote-as 10000
```

```
r3(config-router)#neighbor 100.255.254.5 update-source Loopback0
```

```
r3(config-router)#neighbor 100.255.254.6 remote-as 10000
```

```
r3(config-router)#neighbor 100.255.254.6 update-source Loopback0
```

```
r3(config-router)#address-family vpnv4
```

```
r3(config-router-af)#neighbor 100.255.254.4 activate
```

```
r3(config-router-af)#neighbor 100.255.254.5 activate
```

r3(config-router-af)#neighbor 100.255.254.6 activate

r3(config-router-af)#exit-address-family

r3(config-router)#address-family vpnv6

r3(config-router-af)#neighbor 100.255.254.4 activate

r3(config-router-af)#neighbor 100.255.254.5 activate

r3(config-router-af)#neighbor 100.255.254.6 activate

r3(config-router-af)#exit-address-family

r3(config-router)#exit

r3(config)#

r3(config)#no mpls ip propagate-ttl

r3(config)#mpls label mode all-vrfs protocol bgp-vpnv4 per-vrf

r3(config)#mpls label mode all-vrfs protocol bgp-vpnv6 per-vrf

On r4

r4(config)#router bgp 10000

r4(config-router)#neighbor 100.255.254.3 remote-as 10000

r4(config-router)#neighbor 100.255.254.3 update-source Loopback0

r4(config-router)#neighbor 100.255.254.5 remote-as 10000

r4(config-router)#neighbor 100.255.254.5 update-source Loopback0

r4(config-router)#neighbor 100.255.254.6 remote-as 10000

r4(config-router)#neighbor 100.255.254.6 update-source Loopback0

r4(config-router)#address-family vpnv4

r4(config-router-af)#neighbor 100.255.254.3 activate

r4(config-router-af)#neighbor 100.255.254.5 activate

r4(config-router-af)#neighbor 100.255.254.6 activate

r4(config-router-af)#exit-address-family

r4(config-router)#address-family vpnv6

r4(config-router-af)#neighbor 100.255.254.3 activate

r4(config-router-af)#neighbor 100.255.254.5 activate

r4(config-router-af)#neighbor 100.255.254.6 activate

r4(config-router-af)#exit-address-family

r4(config-router)#exit

r4(config)#

r4(config)#no mpls ip propagate-ttl

r4(config)#mpls label mode all-vrfs protocol bgp-vpnv4 per-vrf

r4(config)#mpls label mode all-vrfs protocol bgp-vpnv6 per-vrf

On r5

r5(config)#router bgp 10000

r5(config-router)#neighbor 100.255.254.3 remote-as 10000

r5(config-router)#neighbor 100.255.254.3 update-source Loopback0

r5(config-router)#neighbor 100.255.254.4 remote-as 10000

r5(config-router)#neighbor 100.255.254.4 update-source Loopback0

r5(config-router)#neighbor 100.255.254.6 remote-as 10000

r5(config-router)#neighbor 100.255.254.6 update-source Loopback0

r5(config-router)#address-family vpnv4

r5(config-router-af)#neighbor 100.255.254.3 activate

r5(config-router-af)#neighbor 100.255.254.4 activate

r5(config-router-af)#neighbor 100.255.254.6 activate

r5(config-router-af)#exit-address-family

r5(config-router)#address-family vpnv6

r5(config-router-af)#neighbor 100.255.254.3 activate

r5(config-router-af)#neighbor 100.255.254.4 activate

r5(config-router-af)#neighbor 100.255.254.6 activate

r5(config-router-af)#exit-address-family

r5(config-router)#exit

r5(config)#

r5(config)#no mpls ip propagate-ttl

r5(config)#mpls label mode all-vrfs protocol bgp-vpnv4 per-vrf

r5(config)#mpls label mode all-vrfs protocol bgp-vpnv6 per-vrf

On r6

r6(config)#router bgp 10000

```
r6(config-router)#bgp router-id 10.255.254.6

r6(config-router)#bgp log-neighbor-changes

r6(config-router)#neighbor 100.255.254.3 remote-as 10000

r6(config-router)#neighbor 100.255.254.3 update-source Loopback0

r6(config-router)#neighbor 100.255.254.4 remote-as 10000

r6(config-router)#neighbor 100.255.254.4 update-source Loopback0

r6(config-router)#neighbor 100.255.254.5 remote-as 10000

r6(config-router)#neighbor 100.255.254.5 update-source Loopback0

r6(config-router)#address-family vpnv4

r6(config-router-af)#neighbor 100.255.254.3 activate

r6(config-router-af)#neighbor 100.255.254.4 activate

r6(config-router-af)#neighbor 100.255.254.5 activate

r6(config-router-af)#exit-address-family

r6(config-router)#address-family vpnv6

r6(config-router-af)#neighbor 100.255.254.3 activate

r6(config-router-af)#neighbor 100.255.254.4 activate

r6(config-router-af)#neighbor 100.255.254.5 activate

r6(config-router-af)#exit-address-family

r6(config-router)#exit

r6(config)#
```

```
r6(config)#no mpls ip propagate-ttl
```

```
r6(config)#mpls label mode all-vrfs protocol bgp-vpnv4 per-vrf
```

```
r6(config)#mpls label mode all-vrfs protocol bgp-vpnv6 per-vrf
```

SECTION 1.11: Fixing Broken DMVPN between DC and Branches #3 & #4

Correct the configuration issues resulting in broken DMVPN tunnel connectivity between DC, Branch3 and Branch4 according to these requirements:

1. The DMVPN must operate in IPsec-protected Phase 3 mode.
2. Using the FVRF approach, safeguard the DMVPN operation against any potential recursive routing issues involving the tunnel.
3. Do not create any new VRFs.
4. Do not change the tunnel source commands on tunnel interfaces.
5. On spokes do not add new BGP neighbors; reuse those that are currently up while changing their VRF membership as needed.
6. It is not allowed to modify configuration on DC r24 to complete this entire task.

Solution

On r24

```
r24(config)#crypto isakmp policy 10
```

```
r24(config-isakmp)#hash md5
```

```
r24(config-isakmp)#exit
```

```
r24(config)#
```

```
r24(config)#interface tunnel0
```

```
r24(config-if)#ip nhrp map multicast dynamic
```

```
r24(config-if)#exit
```

```
r24(config)#
```

On r61

```
r61(config)#crypto keyring KR vrf WAN
```

```
r61(conf-keyring)# pre-shared-key address 0.0.0.0 0.0.0.0 key cisco
```

```
r61(conf-keyring)# exit
```

```
r61(config)#
```

```
r61(config)#interface loopback 0
```

```
r61(config-if)#vrf forwarding WAN
```

```
r61(config-if)#ip address 10.6.255.61 255.255.255.255
```

```
r61(config-if)#exit
```

```
r61(config)#
```

```
r61(config)#interface GigabitEthernet0/0
```

```
r61(config-if)#vrf forwarding WAN
```

```
r61(config-if)#ip address 100.5.61.2 255.255.255.252
```

```
r61(config-if)#exit
```

```
r61(config)#
```

```
r61(config)#interface tunnel 0
```

```
r61(config-if)#ip mtu 1440
```

```
r61(config-if)#ip nhrp shortcut
```

```
r61(config-if)#no ip nhrp map 10.2.255.24 10.200.0.1
```

```
r61(config-if)#ip nhrp map 10.200.0.1 10.2.255.24
```


r61(config-if)#tunnel vrf WAN

r61(config-if)#exit

r61(config)#

r61(config)#router bgp 65006

r61(config-router)#no network 10.6.255.61 mask 255.255.255.255

r61(config-router)#no neighbor 100.5.61.1 remote-as 10000

r61(config-router)#address-family ipv4 vrf WAN

r61(config-router-af)#network 10.6.255.61 mask 255.255.255.255

r61(config-router-af)#neighbor 100.5.61.1 remote-as 10000

r61(config-router-af)#neighbor 100.5.61.1 activate

r61(config-router-af)#exit-address-family

r61(config-router)#exit

r61(config)#

On r62

r62(config)#crypto keyring KR vrf WAN

r62(conf-keyring)# pre-shared-key address 0.0.0.0 0.0.0.0 key cisco

r62(conf-keyring)#exit

r62(config)#

r62(config)#interface Loopback0

r62(config-if)# vrf forwarding WAN

r62(config-if)# ip address 10.6.255.62 255.255.255.255

r62(config-if)#exit

r62(config)#

r62(config)#interface GigabitEthernet0/0

r62(config-if)# vrf forwarding WAN

r62(config-if)# ip address 100.6.62.2 255.255.255.252

r62(config-if)#exit

r62(config)#

r62(config)#interface tunnel 0

r62(config-if)#ip mtu 1440

r62(config-if)#ip nhrp shortcut

r62(config-if)#ip nhrp network-id 1010

r62(config-if)#tunnel vrf WAN

r62(config-if)#exit

r62(config)#

r62(config)#router bgp 65006

r62(config-router)#no network 10.6.255.62 mask 255.255.255.255

r62(config-router)#no neighbor 100.6.62.1 remote-as 10000

r62(config-router)# address-family ipv4 vrf WAN

r62(config-router-af)# network 10.6.255.62 mask 255.255.255.255

```
r62(config-router-af)# neighbor 100.6.62.1 remote-as 10000
```

```
r62(config-router-af)# neighbor 100.6.62.1 activate
```

```
r62(config-router-af)# exit-address-family
```

```
r62(config-router)#exit
```

```
r62(config)#
```

```
On r70
```

```
r70(config)#crypto isakmp key cisco address 0.0.0.0
```

```
r70(config)#interface Tunnel0
```

```
r70(config-if)# ip mtu 1440
```

```
r70(config-if)#ip nhrp shortcut
```

```
r70(config-if)#ip nhrp redirect
```

```
r70(config-if)#tunnel vrf WAN
```

```
r70(config-if)#exit
```

```
r70(config)#
```

SECTION 1.12: Turning EIGRP on DMVPN and DMVPN-enabled Sites

BestCiscoDumps

Optimize the DMVPN operation according to these requirements:

1. Ensure that Branch 3 & Branch 4 can receive only a default route over EIGRP in DMVPN.
2. The default route origination must be done on r24 without the use of any static routes, redistribution, or route filtering.
3. It is not allowed to modify the configuration of r61 and r62 in Branch#3 to accomplish this task.

4. It is allowed to add commands to the configuration of r70 in Branch#4 to accomplish this task; none of the existing configuration on r70 may be removed to accomplish this task.

Configure sw601 and sw602 at Branch#3 according to these requirements:

1. Routers r61 and r62 must not send EIGRP queries to sw601 and sw602.
2. Switches sw601 and sw602 must allow advertising any current or future directly connected network to r61 and r62 after the network is added to EIGRP.
3. Switches sw601 and sw602 must continue to propagate the default route received from r61 and r62 to each other. To select the default route, use a prefix list with a "permit" type entry only.
4. Switches sw601 and sw602 must not propagate the default route back to r61 and r62.
5. If the prefix list that allows the propagation of selected EIGRP learned networks between sw601 and sw602 is modified in the future, the same set of networks must be disallowed from being advertised back to r61 and r62 automatically, without any additional configuration.

Solution

On r24

```
r24(config)#router eigrp ccie
```

```
r24(config-router)#address-family ipv4 unicast autonomous-system 65006
```

```
r24(config-router-af)#af-interface tunnel 0
```

```
r24(config-router-af-interface)#no passive-interface
```

```
r24(config-router-af-interface)#summary-address 0.0.0.0 0.0.0.0
```

```
r24(config-router-af-interface)#exit-af-interface
```

```
r24(config-router-af)#exit-address-family
```

```
r24(config-router)#exit
```

```
r24(config)#
```

On r70

```
r70(config)#router eigrp ccie
```

```
r70(config-router)#address-family ipv4 unicast autonomous-system 65006
```

```
r70(config-router-af)#af-interface tunnel 0
```

```
r70(config-router-af-interface)#no passive-interface
```

```
r70(config-router-af-interface)#exit-af-interface
```

```
r70(config-router-af)#exit-address-family
```

```
r70(config-router)#exit
```

```
r70(config)#
```

On sw601

```
sw601(config)#ip prefix-list PERMIT seq 1 permit 0.0.0.0/0
```

```
sw601(config)#route-map PERMIT permit 1
```

```
sw601(config-route-map)#match ip address prefix-list PERMIT
```

```
sw601(config-route-map)#exit
```

```
sw601(config)#
```

```
sw601(config)#access-list 1 permit 10.6.0.0 0.0.255.255
```

```
sw601(config)#access-list 1 deny any
```

```
sw601(config)#router eigrp ccie
```

```
sw601(config-router)# address-family ipv4 unicast autonomous-system 65006
```

```
sw601(config-router-af)#eigrp stub connected leak-map PERMIT
```

```
sw601(config-router-af)#topology base
sw601(config-router-af-topology)#distribute-list 1 out Ethernet0/1
sw601(config-router-af-topology)#distribute-list 1 out Ethernet0/2
sw601(config-router-af-topology)#exit-af-topology
sw601(config-router-af)#exit-address-family
sw601(config-router-af)#exit
sw601(config)#
```

On sw602

```
sw602(config)#ip prefix-list PERMIT seq 1 permit 0.0.0.0/0
sw602(config)#route-map PERMIT permit 1
sw602(config-route-map)#match ip address prefix-list PERMIT
sw602(config-route-map)#exit
sw602(config)#
sw602(config)#access-list 1 permit 10.6.0.0 0.0.255.255
sw602(config)#access-list 1 deny any
sw602(config)#router eigrp ccie
sw602(config-router)#address-family ipv4 unicast autonomous-system 65006
sw602(config-router-af)#eigrp stub connected leak-map PERMIT
sw602(config-router-af)#topology base
sw602(config-router-af-topology)#distribute-list 1 out Ethernet0/1
```

```
sw602(config-router-af-topology)#distribute-list 1 out Ethernet0/2
```

```
sw602(config-router-af-topology)#exit-af-topology
```

```
sw602(config-router-af)#exit-address-family
```

```
sw602(config-router)#exit
```

```
sw602(config)#
```

SECTION 1.13 IPv4 Networks on Legacy Branches

On sw211 in DC, complete the DHCP server configuration according to these requirements:

1. Create IPv4 DHCP pools named br3_v2000 and br3_v2001 for Branch#3 VLANs 2000 (10.6.100.0/24) and 2001 (10.6.101.0/24), respectively.
2. Create IPv4 DHCP pool named br_v1 for the subnet 10.7.1.0/24 on Branch#4
3. In each subnet assign addresses from .101 up to .254 inclusively and the appropriate gateway to clients.

On Branch#3 complete and correct the configuration on switches sw601, sw602 and sw610 to allow HSRP and DHCP Relay operation in VLANs 2000 and 2001 according to these requirements:

1. HSRP must implicitly use the vMAC address range of 0000.0c9f.f000 through 0000.0c9f.ffff
2. The group number must be 100 for VLAN 2000 and 101 for VLAN 2001
3. Sw601 must be Active gateway for VLAN 2000 with a priority of 110; the Active role ownership must be deterministic.
4. Sw602 must be Active gateway for VLAN 2001 with a priority of 110; the Active role ownership must be deterministic.
5. Each Active switch must track its uplink interface g0/1 and g0/2. If either of these interface goes down, the Active switch must allow the other switch to become Active. However, it is not allowed for the tracking to modify the HSRP priority to accomplish this requirements.
6. Both sw601 and sw602 must be configured as DHCP relay agents in both VLANs 2000 and 2001, pointing toward the DHCP server 10.2.255.211 at sw211. However, at anytime, only the Active router in the particular VLAN should relay the DHCP messages.
7. Place host61 and host62 into VLANs 2000 and 2001 respectively and make sure they are assigned their correct IPv4 configuration.

8. It is not permitted to use any kind of scripting to complete this task.

On Branch#3 complete the configuration of the router r70 according to these requirements:

1. Assign IP address 10.7.1.1/24 to g0/2
2. Enable DHCP relay on this interface and point it to the DHCP server 10.2.255.211 at sw211.
3. It is allowed to add one additional missing command to the r70 configuration to allow clients connected to g0/2 obtain their IPv4 configuration.
4. Make sure that host r71 and host r72 are assigned their correct IPv4 configuration

Solution

On sw211

```
sw211(config)#ip dhcp pool br3_v2000
```

```
sw211(dhcp-config)#network 10.6.100.0 255.255.255.0
```

```
sw211(dhcp-config)#default-router 10.6.100.1
```

```
sw211(dhcp-config)#exit
```

```
sw211(config)#
```

```
sw211(config)#ip dhcp excluded-address 10.6.100.1 10.6.100.100
```

```
sw211(config)#ip dhcp excluded-address 10.6.100.255
```

```
sw211(config)#ip dhcp pool br3_v2001
```

```
sw211(dhcp-config)#network 10.6.101.0 255.255.255.0
```

```
sw211(dhcp-config)#default-router 10.6.101.1
```

```
sw211(dhcp-config)#exit
```

```
sw211(config)#
```


sw211(config)#ip dhcp excluded-address 10.6.101.1 10.6.101.100

sw211(config)#ip dhcp excluded-address 10.6.101.255

sw211(config)#ip dhcp pool br_v1

sw211(dhcp-config)#network 10.7.1.0 255.255.255.0

sw211(dhcp-config)#default-router 10.7.1.1

sw211(dhcp-config)#exit

sw211(config)#

sw211(config)#ip dhcp excluded-address 10.7.1.1 10.7.1.100

sw211(config)#ip dhcp excluded-address 10.7.1.255

sw211(config)#

On sw601

sw601(config)#interface vlan 2000

sw601(config-if)#standby 100 preempt

sw601(config-if)#ip helper-address 10.2.255.211

sw601(config-if)#exit

sw601(config)#

sw601(config)#interface vlan 2001

sw601(config-if)#standby version 2

sw601(config-if)#exit

sw601(config)#

sw601(config)#track 1 interface gigabitEthernet 0/1 line-protocol

sw601(config-track)#exit

sw601(config)#

sw601(config)#track 2 interface gigabitEthernet 0/2 line-protocol

sw601(config-track)#exit

sw601(config)#

sw601(config)#interface vlan 2000

sw601(config-if)#standby 100 track 1 shutdown

sw601(config-if)#standby 100 track 2 shutdown

sw601(config-if)#exit

sw601(config)#

On sw602

sw602(config)#interface Vlan2000

sw602(config-if)#standby version 2

sw602(config-if)#exit

sw602(config)#

sw602(config)#interface Vlan2001

sw602(config-if)#standby version 2

sw602(config-if)#no standby 0 ip 10.6.101.1

sw602(config-if)#standby 101 ip 10.6.101.1

sw602(config-if)#standby 101 priority 110

sw602(config-if)#standby 101 preempt

sw602(config-if)#ip helper-address 10.2.255.211

sw602(config-if)#exit

sw602(config)#

sw602(config)#track 1 interface gigabitEthernet 0/1 line-protocol

sw602(config-track)#exit

sw602(config)#

sw602(config)#track 2 interface gigabitEthernet 0/2 line-protocol

sw602(config-track)#exit

sw602(config)#

sw602(config)#interface vlan 2001

sw602(config-if)#standby 101 track 1 shutdown

sw602(config-if)#standby 101 track 2 shutdown

sw602(config-if)#exit

sw602(config)#

On sw610

sw610(config)#vlan 2000

sw610(config-vlan)#exit

sw610(config)#

sw610(config)#interface gigabitEthernet 0/0

sw610(config-if)#switchport mode access

sw610(config-if)#switchport access vlan 2000

sw610(config-if)#exit

sw610(config)#

sw610(config)#interface gigabitEthernet 0/1

sw610(config-if)#switchport mode access

sw610(config-if)#switchport access vlan 2001

sw610(config-if)#exit

sw610(config)#

sw610(config)#interface GigabitEthernet2/0

sw610(config-if)#no switchport trunk allowed vlan 1,2000

sw610(config-if)#switchport trunk allowed vlan 1,2000-2001

sw610(config-if)#exit

sw610(config)#

sw610(config)#interface GigabitEthernet2/1

sw610(config-if)#no switchport trunk allowed vlan 1,2000

sw610(config-if)#switchport trunk allowed vlan 1,2000-2001

sw610(config-if)#exit

sw610(config)#

On r70

```
r70(config)#interface gigabitEthernet 0/2
```

```
r70(config-if)#ip address 10.7.1.1 255.255.255.0
```

```
r70(config-if)#ip helper-address 10.2.255.211
```

```
r70(config-if)#exit
```

```
r70(config)#
```

```
r70(config)#router eigrp ccie
```

```
r70(config-router)# address-family ipv4 unicast autonomous-system 65006
```

```
r70(config-router-af)#network 10.7.1.0 0.0.0.255
```

```
r70(config-router-af)#network 172.31.70.0 0.0.0.255
```

```
r70(config-router-af)#network 192.168.0.0 0.0.255.255
```

```
r70(config-router-af)#exit-address-family
```

```
r70(config-router)#exit
```

```
r70(config)#
```

SECTION 1.14: Multicast in FADB2

FABD2 is preparing to enable PIM Sparse Mode multicast routing in its network. As a part of validating the runbooks, FABD2 requires a sanity check to prevent inappropriate use of multicast related configuration commands on different router types.

- First Hop Routers – routers where multicast sources are connected.
- Last Hop Routers – routers where multicast receivers (subscribers) are connected.
- Intermediary Hop Routers – routers on the path between First Hop and Last Hop routers.

In the table below, for each configuration command, select all router types where the use of the command is appropriate. (Select all that apply)

Command	Router Type		
	Last Hop Router	First Hop Router	Intermediary Hop Router
ip igmp spt-threshold			
ip igmp version			
ip pim register-source			
ip pim sparse-mode			
ip pim rp-address			

Solution

Command	Router Type		
	Last Hop Router	First Hop Router	Intermediary Hop Router
ip igmp spt-threshold	✓		
ip igmp version	✓		
ip pim register-source		✓	
ip pim sparse-mode	✓	✓	✓
ip pim rp-address	✓	✓	✓

SECTION 1.15: Extending connectivity to IaaS

Extend the IPv6 connectivity from HQ through the SP into the giosk VRF on the IaaS site according to these requirements.

- Set up global IPv6 addressing on the link between r11 and r3.
 - On r11 assign 2001:2710:311::2/64 to g0/0
 - On r3 assign 2001:2710:311::1/64 to g1
- Enable the existing IPv4 BGP session between r11 and r3 to also advertise IPv6 prefixes. Do not configure a standalone IPv6 BGP session between these two routers.
- Perform bidirectional route redistribution between the IPv6 EIGRP and BGP processes on r11
- Ensure that all current and future IPv6 prefixes advertised between r11 and r3 will be installed into the RIB of these routers with the next hop address set to the proper global unicast address on their interconnection. Any policy that accomplishes this requirement must be applied in the inbound direction
- The giosk VRF on r4 that extends the IPv6 connectivity from r4 to r30 on the IaaS site is a separate VRF independent of fabd2 VRF. Any route leaking from fabd2 VRF into giosk VRF must be done on a per-site basis and only for this FABD2 sites that need connectivity with the IaaS site.
- By configuring r3 and r4 only, ensure that the HQ FABD2 site will have mutual visibility with the IaaS site while preventing

- any other FABD2 site from possibly learning about the routes on the laaS site
 - the laaS site from possibly learning about the routes on any other FABD2 site
7. Use the minimum number of commands necessary to accomplish this requirements. Do not remove any existing configuration. If necessary, you are allowed to use an additional route-target with the value of 10000:3681
 8. Verify that host11 and host12 can ping 2001:db8:14::1 located at the laaS site. It is permitted to modify one existing configuration command on one of the SP routers to meet this requirement

Solution

On r11

```
r11(config)#interface Ethernet0/0
```

```
r11(config-if)# ipv6 address 2001:2710:311::2/64
```

```
r11(config-if)#exit
```

```
r11(config)#
```

```
r11(config)#router bgp 65001
```

```
r11(config-router)#neighbor 2001:2710:311::1 remote-as 10000
```

```
r11(config-router)#address-family ipv4
```

```
r11(config-router-af)#neighbor 2001:2710:311::1 activate
```

```
r11(config-router-af)#redistribute eigrp 65001
```

```
r11(config-router-af)#exit-address-family
```

```
r11(config-router)#exit
```

```
r11(config)#
```

```
r11(config)#router eigrp ccie
```

```
r11(config-router)# address-family ipv6 unicast autonomous-system 65001
```

```
r11(config-router-af)# topology base
```

```
r11(config-router-af-topology)#redistribute bgp 65001
```

```
r11(config-router-af-topology)#exit-af-topology
```

```
r11(config-router-af)#exit-address-family
```

```
r11(config-router)#exit
```

```
r11(config)#
```

On r3

```
r3(config)#interface Ethernet0/0
```

```
r3(config-if)# ipv6 address 2001:2710:311::1/64
```

```
r3(config-if)#exit
```

```
r3(config)#
```

```
r3(config)#router bgp 10000
```

```
r3(config-router)#address-family ipv4 vrf fabd2
```

```
r3(config-router-af)#redistribute connected
```

```
r3(config-router-af)#neighbor 2001:2710:311::2 remote-as 65001
```

```
r3(config-router-af)#neighbor 2001:2710:311::2 activate
```

```
r3(config-router-af)#exit-address-family
```

```
r3(config-router)#exit
```

```
r3(config)#
```


SECTION 1.16: Enabling Internet Access for FABD2

Enable highly available internet access for the FABD2 company network according to these requirements:

1. On routers r12, r23 and r24, bring up IPv4 BGP peerings with the ISP. Make sure that a default route is received over these peerings.
2. On routers r12 and r23 inject default route into OSPF if it is present in the routing table from a different routing source than the OSPFv2 process 1. On each router, this requirement must be completed using minimum possible number of commands.
3. On router r24 inject a default route into OSPF if and only if it is learned from ISP over BGP. To accomplish this requirement, it is allowed to use a route-map that references both a prefix-list and a tag. This requirement must be completed using minimum possible number of commands.
4. Router r12 may be used as an internet exit for the FABD2 company network only if neither r23 nor r24 are advertising a default route in OSPF. This requirement must be accomplished exclusively in "router ospf" mode on router r12 without changing the default parameters on routers r23 and r24.
5. On routers r12, r23 and r24 configure PAT and translate the entire FABD2 internal network 10.0.0.0/8 to the router address on the link towards the ISP. Create a standard ACL named NAT for this purpose. Do not use NAT pools.
6. Ensure that the internet connectivity of the FABD2 company network makes use of high availability provided by r12, r23 and r24

Solution

On r12

```
r12(config)#router bgp 65001
```

```
r12(config-router)#neighbor 200.99.12.1 remote-as 19999
```

```
r12(config-router)#exit
```

```
r12(config)#
```

```
r12(config)#router ospf 1
```

```
r12(config-router)#default-information originate metric 50000
```

```
r12(config-router)#exit
```

```
r12(config)#  
  
r12(config)#interface gigabitEthernet 0/0  
  
r12(config-if)#ip nat outside  
  
r12(config-if)#exit  
  
r12(config)#  
r12(config)#interface gigabitEthernet 0/1  
r12(config-if)#ip nat inside  
  
r12(config-if)#exit  
  
r12(config)#  
r12(config)#interface gigabitEthernet 0/2  
r12(config-if)#ip nat inside  
  
r12(config-if)#exit  
  
r12(config)#  
r12(config)#interface gigabitEthernet 0/3  
r12(config-if)#ip nat inside  
r12(config-if)#exit  
  
r12(config)#  
r12(config)#ip access-list standard NAT  
r12(config-std-nacl)#permit 10.0.0.0 0.255.255.255  
r12(config-std-nacl)#exit
```

r12(config)#

r12(config)# ip nat inside source list NAT interface gigabitEthernet 0/0
overload

On r23

r23(config)#router bgp 65002

r23(config-router)#neighbor 200.99.23.1 remote-as 19999

BestCiscoDumps
r23(config-router)#exit

r23(config)#

r23(config)#router ospf 1

r23(config-router)#default-information originate

r23(config-router)#exit

r23(config)#

r23(config)#interface gigabitEthernet 1

r23(config-if)#ip nat outside

r23(config-if)#exit

r23(config)# *BestCiscoDumps*

r23(config)#interface gigabitEthernet 2

r23(config-if)#ip nat inside

r23(config-if)#exit

r23(config)#

r23(config)#interface gigabitEthernet 3

r23(config-if)#ip nat inside

r23(config-if)#exit

r23(config)#

r23(config)#ip access-list standard NAT

r23(config-std-nacl)#permit 10.0.0.0 0.255.255.255

r23(config-std-nacl)#exit

r23(config)#

r23(config)#ip nat inside source list NAT interface gigabitEthernet 0/0
overload

On r24

r24(config)#router bgp 65002

r24(config-router)#neighbor 200.99.24.1 remote-as 19999

r24(config-router)#distance 4 200.99.24.1 0.0.0.0

r24(config-router)#exit

r24(config)#

r24(config)#ip prefix-list INTERNET seq 1 permit 0.0.0.0/0

r24(config)#route-map INTRENET permit 1

r24(config-route-map)#match ip address prefix-list INTERNET

r24(config-route-map)#match tag 19999

r24(config-route-map)#exit

r24(config)#

```
r24(config)#route-map INTRENET permit 2

r24(config-route-map)#exit

r24(config)#

r24(config)#router ospf 1

r24(config-router)#default-information originate route-map INTERNET

r24(config-router)#exit

r24(config)#

r24(config)#interface gigabitEthernet 1

r24(config-if)#ip nat outside

r24(config-if)#exit

r24(config)#

r24(config)#interface gigabitEthernet 2

r24(config-if)#ip nat inside

r24(config-if)#exit

r24(config)#

r24(config)#interface gigabitEthernet 3

r24(config-if)#ip nat inside

r24(config-if)#exit

r24(config)#

r24(config)#interface tunnel 0
```

```
r24(config-if)#ip nat inside
```

```
r24(config-if)#exit
```

```
r24(config)#
```

```
r24(config)#ip access-list standard NAT
```

```
r24(config-std-nacl)#permit 10.0.0.0 0.255.255.255
```

```
r24(config-std-nacl)#exit
```

```
r24(config)#
```

```
r24(config)#ip nat inside source list NAT interface gigabitEthernet 0/0  
overload
```

The above is today's [CCIE enterprise infrastructure lab](#) practical questions sharing. If you want to download more CCIE EI lab workbooks for free, please always pay attention to our website, because we will bring benefits to candidates from time to time!

If you want to get more verification about the questions and a more complete [CCIE EI lab workbook](#), please contact us!

BestCiscoDumps