



KaozhengPro

IT認證考試題庫 專業平臺


考證寶提供最新考古題與模擬試題
協助您高效通過認證考試

www.kaozhengpro.com

Exam : 350-001-LAB

Title : CCIE-ROUTING AND SWITCHING QUALIFICATION (Lab exam)

Version : Demo



Part-1 : Bridge and Switching

1.1 vtp

Sw1/Sw2/Sw3/Sw4

VTP domain:VTP+YY

VTP mode:Transparent

RackYYSw1/SW2/SW3/SW4:

vtp domain VTPYY
vtp mode transparent
that's not need to use s/c mode in VT

1.2
VLAN

Sw1:
40 VLAN_BB1
55 VLAN_55
60 VLAN_BB3
100 VLAN_100
200 VLAN_200

Sw2:
50 VLAN_BB2
100 VLAN_100
200 VLAN_200

Sw3
60 VLAN_BB3
200 VLAN_200
Sw4
200 VLAN_200

RackYYSw1:
VLAN Name Status Ports
1 default active F0/7,F0/8,F0/9,F0/11,F0/12,F0/13,F0/14Fa0/15,Fa0/17
Fa0/18,Gi0/1,Gi0/2
40 VLAN_BB1 active Fa0/4, Fa0/10
55 VLAN_55 active Fa0/5
60 VLAN_BB3 active Fa0/6
100 VLAN_100 active Fa0/1, Fa0/2
200 VLAN_200 active

RackYYSw2:
VLAN Name Status Ports
1 default active F0/4,F0/5,F0/7,F0/8,F0/9,F0/11,F0/12,F0/13,F0/14,F0/16,F0/17
F0/18,Gi0/1,Gi0/2
50 VLAN_BB2 active F0/6, F0/10
100 VLAN_100 active F0/1, F0/2
200 VLAN_200 active Po10

RackYYSw3:

VLAN Name Status Ports

```
1 default active      F0/1,F0/2,F0/3,F0/4 F0/5, F0/6, F0/7,F0/8,F0/9,F0/11,F0/12,F0/13,
F0/14, F0/15, F0/16,F0/17,F0/18,Gi0/1,Gi0/2
60 VLAN_BB3 active   Fa0/10
200 VLAN_200 active
```

RackYYSw4:

VLAN Name Status Ports

```
1 default active      F0/1,F0/2,F0/3,F0/4,F0/5,F0/6,F0/7,F0/8,F0/9,F0/10F0/11,F0/12
F0/13,F0/14,F0/15,F0/16,F0/17,F0/18,Gi0/1,Gi0/2
200 VLAN_200 active   Po10
```

Verify: show vlan brief

When u done, u must verify careful.

1.3 Port-channel between Sw1 a Sw2

Sw1 and Sw2 layer 3 Ether channel use port-channel 21 only

Assign YY.YY.100.1/24 to Sw1 Port-channel 21

Assign YY.YY.100.2/24 to Sw2 Port-channel 21

ensure interface F0/23 and F0/24 are channel member for both Switch

Do not rely on PAgP or LACP to facilitate the connection

Verify layer 2 and layer 3 connectivity via the channel link

RackYYSw1:

```
interface range FastEthernet0/23-24
shutdown
no Switchport
channel-group 21 mode on
no shutdown
interface Port-channel21
ip address YY.YY.100.1 255.255.255.0
```

RackYYSw2:

```
interface range FastEthernet0/23-24
shutdown
no Switchport
channel-group 21 mode on
no shutdown
interface Port-channel21
ip address YY.YY.100.2 255.255.255.0
```

Verify: show vtp status; show etherchannel summary; show etherchannel port-channel;

1.4 Port-channel between Sw1-Sw3 and Sw2-Sw4

Sw1-Sw3 Sw2-Sw4 layer 2 Ether-channels use port 10 only F0/19 and F0/20 are members of the switches

On Sw2 and Sw4 assign all interface in the EC as static-access port on VLAN_200 unconditional enable PAGP to facility the connection

verify layer 2 connectivity via the channel link

RackYYSw1:

```
interface range FastEthernet0/19-20
shutdown
Switchport trunk encapsulation isl
Switchport mode trunk
channel-group 10 mode desirable
no shutdown
```

RackYYSw3:

```
interface range FastEthernet0/19-20
shutdown
Switchport trunk encapsulation isl
Switchport mode trunk
channel-group 10 mode desirable
no shutdown
```

RackYYSw2:

```
interface range FastEthernet0/19-20
shutdown
Switchport mode access
Switchport access VLAN 200
channel-group 10 mode desirable
no shutdown
```

RackYYSw4:

```
interface range FastEthernet0/19-20
shutdown
Switchport mode access
Switchport access VLAN 200
channel-group 10 mode desirable
no shutdown
```

Verify: show etherchannel summary; show etherchannel portchannel;

Combining the VLAN below, you can see the Sw2 and Sw4 po10 also in VLAN200, before you do on any group - 10 before Switchport access desirable mode or the VLAN 200 VLAN po10 see table.

Must be done in the Port. See display any up.

1.5

Catalyst layer 3 configuration

Configure Sw1 and Sw2 IP address as outlined n diagram

Connectivity to R3 uses route ports

R1 and R2 are members of vlan 100 on Sw1 and Sw2

1.6

Catalyst layer 3 configuration

Configure Sw3 and Sw4 IP addressing

Configure VLAN_200 in Sw1 with IP address YY.YY.34.1/24

Configure VLAN_200 in Sw2 with IP address YY.YY.43.1/24

Verify the connectivity between Sw1 and Sw2

RackYYSw1:

VLAN 100 YY.YY.12.254/24

VLAN 200 YY.YY.34.1/24

RackYYSw2:

VLAN 100 YY.YY.21.254/24

VLAN 200 YY.YY.43.1/24

RackYYSw3:

VLAN 200 YY.YY.34.254/24

RackYYSw4:

VLAN 200 YY.YY.43.254/24

RackYYSw1:

ip routing

interface VLAN100

ip address YY.YY.12.254 255.255.255.0

interface VLAN200

ip address sYY.YY.34.1 255.255.255.0

RackYYSw2:

ip routing

interface VLAN100

ip address YY.YY.21.254 255.255.255.0

interface VLAN200

ip address YY.YY.43.1 255.255.255.0

RackYYSw3:

ip routing

```
interface VLAN200
ip address YY.YY.34.254 255.255.255.0
```

```
RackYYSw4:
ip routing
interface VLAN200
ip address YY.YY.43.254 255.255.255.0
Verify: show ip interface brief; show ip route
```

```
RackYYSw1:
interface FastEthernet0/3
no Switchport
ip address YY.YY.13.2 255.255.255.0
```

```
RackYYSw2:
interface FastEthernet0/3
no Switchport
ip address YY.YY.31.2 255.255.255.0
Verify: show interface status; show ip interface brief; show ip route
```

1.7Catalyst feature

Configure Sw1-F0/1 so that the interface will stop forwarding unicast traffic if the input rate exceeds 65 Mbps

```
RackYYSw1:
interface Fa0/1
Storm-control unicast level 55.00
Verify: show storm-control unicast
```

1.8 Catalyst tuning

Configure the amount of time a neighbour should hold CDP information sent by Sw2 before discarding it to 2 minutes

```
RackYYSw1:
cdp holdtime 120
Verify: show cdp
```

1.9 Catalyst Feature

Configure Sw1 to control and block the flood of unknown Multicast traffic on the interface F0/5

```
RackYYSw1:
interface Fa0/5
```

Switchport block multicast

Ip Igmp snooping

Or 'ip cg mp enab'•

Verify: show interface interface-id switchport

Part-2 : IGP and BGP

IGP

2.1 OSPF Bbackbones

The link between Sw1 and Sw2

All interface in VLAN_100 on Sw1 Sw2 R1 and R2

R3 G0/0 and G0/1 and the fa0/3 on Sw1 and Sw2

Loop back 0 interface on Sw1 Sw2 R2 and R3

Verifying that all OSPF neighbor have built their adjacencies

RackYYR1:

Router ospf YY

network YY.YY.12.1 0.0.0.0 area 0

network YY.YY.21.1 0.0.0.0 area 0

RackYYR2:

Router ospf YY

network YY.YY.2.2 0.0.0.0 area 0

network YY.YY.12.2 0.0.0.0 area 0

network YY.YY.21.2 0.0.0.0 area 0

RackYYSw1:

Router ospf YY

network YY.YY.7.7 0.0.0.0 area 0

network YY.YY.12.254 0.0.0.0 area 0

network YY.YY.13.2 0.0.0.0 area 0

network YY.YY.100.1 0.0.0.0 area 0

RackYYSw2:

Router ospf YY

network YY.YY.8.8 0.0.0.0 area 0

network YY.YY.21.254 0.0.0.0 area 0

network YY.YY.31.2 0.0.0.0 area 0

network YY.YY.100.2 0.0.0.0 area 0

RackYYR3:

Router ospf YY

network YY.YY3.3 0.0.0.0 area 0

```
network YY.YY.13.1 0.0.0.0 area 0
network YY.YY.31.1 0.0.0.0 area 0
Verify: show ip ospf interface brief; show ip ospf neighbor
```

2.2 OSPF over NBMA

OSPF area 11 consist of the follow interface and attributes
The Frame Relay network between R3 R4 R5
Loop back 0 on R4 and R5
VLAN_55
Ensure there is no DR/BDR

RackYYR3:

```
interface s0/0/0.3
ip ospf network point-to-multipoint non-broadcast
Router ospf YY
network YY.YY.11.3 0.0.0.0 area 11
nei YY.YY.11.4
nei YY.YY.11.5
```

RackYYR4:

```
interface s0/0/0.4
ip ospf network point-to-multipoint non-broadcast
Router ospf YY
network YY.YY.4.4 0.0.0.0 area 11
network YY.YY.11.4 0.0.0.0 area 11
```

RackYYR5:

```
interface s0/0/0.5
ip os net point-to-multipoint non-broadcast
Router ospf YY
network YY.YY.5.5 0.0.0.0 area 11
network YY.YY.11.5 0.0.0.0 area 11
network YY.YY.55.254 0.0.0.0 area 11
Verify: show ip ospf interface brief; show ip ospf neighbor
```

2.3 OSPF ASBR and RIP version 2

Configure R4 to receive RIP v2 routes from Backbone 1
When properly configured you will receives RIP v2 routes in the class B address range 199.172.Z.Z
Configure R4 so that the external RIP routes are injected into area 11 and appear throughout that OSPF domain
Ensure external routes originates from Autonomous Systems Boundary Routers (ASBR) outside area 11 cannot be flooded within the area

Permit OSPF type-3 routes into area 11(Look in R5)

RackYYR4:

```
ip prefix-list fbb1 per 199.172.0.0/16 le 32
```

```
Router rip
```

```
version 2
```

```
no auto-summary
```

```
network 150.1.0.0
```

```
distribute-list prefix fbb1 in Fa0/0
```

```
Router ospf YY
```

```
redistribute rip metric-type 1 subnets
```

```
area 11 nssa
```

RackYYR3:

```
Router ospf YY
```

```
area 11 nssa
```

RackYYR5:

```
Router ospf YY
```

```
area 11 nssa
```

Verify: show ip protocol; show ip route rip; show ip ospf; show ip route ospf;

2.4 Area 34 and Area 43

OSPF area 34 consists of the VLAN_200 interfaces on Sw1 and Sw3 and loopback 0 in Sw3

OSPF area 43 consists of the VLAN_200 interfaces on Sw2 and Sw4 and loopback 0 in Sw4

RackYYSw1:

```
Router ospf YY
```

```
network YY.YY.34.1 0.0.0.0 area 34
```

RackYYSw2:

```
Router ospf YY
```

```
network YY.YY.43.1 0.0.0.0 area 43
```

RackYYSw3:

```
Router ospf YY
```

```
network YY.YY.9.9 0.0.0.0 area 34
```

```
network YY.YY.34.254 0.0.0.0 area 34
```

RackYYSw4:

```
Router ospf YY
```

```
network YY.YY.10.10 0.0.0.0 area 43
network YY.YY.43.254 0.0.0.0 area 43
Verify: show ip ospf interface brief; show ip ospf neighbor
```

2.5 OSPF ABR

Static routes are not permitted for this question
inject a default route into area 0 area 11 area 34 area 43
Use fewest number of steps or commands to completes this

```
RackYYR3:
Router ospf YY
area 11 nssa default-information-originate
default-information originate always
Verify: show ip route ospf; show ip ospf database
```

2.6 OSPF Summary

Add the following interface on R2 to Area 0
Loopback 22 180.88.22.254/24
Loopback 32 180.88.32.254/24
Loopback 47 180.88.47.254/24
Summarize the above address into a single route
Your summary route must be compact and not waste address space
Verify the Summary is in the OSPF routing table on R5 and you can ping all the host
address

R3, Sw 1, Sw2 to make regional summary.

```
RackYYR2:
int lo22
ip address 180.88.22.254 255.255.255.0
int lo32
ip address 180.88.32.254 255.255.255.0
int lo47
ip address 180.88.47.254 255.255.255.0
Router ospf YY
network 180.88.0.0 0.0.63.255 area 0
Here I choose loopback interface directly in three IOS12.4 versions used after the interface
declared ospf support. That fast, and not an error.
```

```
RackYYR3/Sw1/Sw2:
Router ospf YY
area 0 range 180.88.0.0 255.255.192.0
Verify; show ip ospf; show ip route ospf; show ip ospf database
```

(If the announcement of OSPF loop, IP add is 24 bits, I use all is point-to-point type)

2.7 RIP version 2

Advertise all the individual YY.YY.0.0 network prefixes generated within your lab topology to backbone 1

Instruct the backbone 1 router that your networks are 5 hops away

Filter all other prefixes to backbone 1

1: RackYYR4:

Access-list 4 per YY.YY.0.0 0.0.255.255

Router rip

Redistribute ospf yy metric 1

Offset-list 4 out 4 g0/0

Distribute-list 4 out g0/0

Verify: debug ip rip

2: router rip

Redistribute os 8 metric 5 route-map fromOSPF

Route-map fromOSPF per 10

Match ip add prefix-list fromOSPF

Ip prefix-list fromOSPF per 8.8.0.0/16 le 32

Using a heavy distributed orders will solve the three requirements. Our slogan is the strategy to solve the problem with a minimum of job, to minimize CPU consumption. And into the routing database RIP 8.8.0.0 only 16.

3. router rip

Redis os 8 route-map fromOSPF

Default-metric 5

Route-map fromOSFP per 10

Match ip add fromOSPF

Ip access-list standard fromOSPF

Per 8.8.0.0 0.0.255.255

4. router rip

Redis os 8 route-map fromOSPF

Route-map fromOSPF per 10

Set metric 5

Distribute-list prefix fromOSPF out os 8

Ip prefix-list fromOSPF per 8.8.0.0/16 le 32

The source method the worst, To execute the job finished three strategy needs.

2.8 EIGRP

EIGRP 100 AS 100 consists of the following interface

The Frame Relay network between R1 and R6

Loopback0 on R1 and R6

The BB2 interface on R6 should appear as an external EIGRP route on R1

R6 must have a single 16 bit prefix via R1 to the YY.YY.0.0 network. Do not use route filters or automatic summary

Redistribute EIGRP routes into ospf area

RackYYR1:

Ip prefix-list eto per YY.YY.0.0/16

Route-map eto deny 10

Match ip add pre eto

Route-map eto per 20

ip prefix-list ote seq 5 permit 0.0.0.0/0

route-map ote deny 10

match ip address prefix-list ote

route-map ote permit 20

Router eigrp 100

No au

Net YY.YY.16.1 0.0.0.0

Net YY.YY.1.1 0.0.0.0

Redistribute ospf YY metric 10000 100 255 1 1500 route-map ote

Router os yy

Redistribute eigrp 100 subnets metric-type 1 route-map eto

Int s0/0/0

Ip summary ei 100 YY.YY.0.0 255.255.0.0

RackYYR6:

Route-map con per 10

Match interface E0/1

Router eigrp 100

No au

Net YY.YY.16.6 0.0.0.0

Net YY.YY.6.6 0.0.0.0

Redistribute connected route-map CON metric 10000 100 255 1 1500

Verify: show ip protocol; show ip route eigrp; show ip route ospf;

2.9 EIGRP over BB3

The backbone 3 router will be sending some class A,B and C IP prefixes

Create a prefix-list and apply it so that the EIGRP process will only accept prefixes in the class C address might on the routing table

Deny all routes to BB3

(Here should notice first octet=192-200)

Prefix list access list

A:0.0.0.0/1 le 32 0.0.0.0 127.255.255.255

B:128.0.0.0/2 le 32 128.0.0.0 63.255.255.255

C:192.0.0.0/3 le 32 192.0.0.0 31.255.255.255

RackYYR6:

Ip prefix-list fbb3 per 192.0.0.0/5 le 32

Ip prefix-list fbb3 per 200.0.0.0/8 le 32

Ip prefix-list tbb3 deny 0.0.0.0/0 le 32

Router eigrp 100

Net 150.3.YY.1 0.0.0.0

Distribute-list prefix fbb3 in F0/0

Distribute-list prefix tbb3 out F0/0

Verify: show ip protocol; show ip route eigrp

2.10 IPV6

R1 G0/1 2033:YY:YY:21::1

S0/0/0 2033:YY:YY:16::1(FE80::217:94FF:FE15:8C90)

R6 f0/1 2033:YY:YY:62::6

S0/3/0 2033:YY:YY:16::6(FE80::215:C6FF:FE4A:6210)

All the interface run OSPF v3

RackYYR1#show ipv6 interface brief

Gi0/0 [up/up]

FE80::ZZZZ:ZZZZ:ZZZZ //link-local address

2038:YY:YY:11::1

Serial0/0/0 [up/up]

FE80::ZZZZ:ZZZZ:ZZZZ

2038:YY:YY:61::1

RackYYR6#show ipv6 interface brief

Gi0/0 [up/up]

FE80::ZZZZ:ZZZZ:ZZZZ

2038:YY:YY:66::6

Serial0/0/0 [up/up]

FE80::ZZZZ:ZZZZ:ZZZZ

2038:YY:YY:61::6

RackYYR1#show ipv6 route

IPv6 Routing Table -7 entries

Codes: C -Connected, L -Local, S -Static, R -RIP, B -BGP U -Per-user Static route I1 -ISIS L1, I2 -ISIS L2, IA - ISIS inte area, IS -ISIS summary O - OSPF intr OI - OSPF inter, OE1 - OSPF ext 1, OE2 -OSPF ext 2 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2

C 2038:YY:YY:11::/64 [0/0]

via ::, Gi0/0

L 2038:YY:YY:11::1/128 [0/0]

via ::, Gi0/0

C 2038:YY:YY:61::/64 [0/0]

via ::, Serial0/0/0

L 2038:YY:YY:61::1/128 [0/0]

via ::, Serial0/0/0

O 2038:YY:YY:66::/64 [110/65]

via FE80::ZZZZ:ZZZZ:ZZZZ, Serial0/0/0

L FE80::/10 [0/0]

via ::, Null0

L FF00::/8 [0/0]

via ::, Null0

RackYYR6#show ipv6 route

IPv6 Routing Table -7 entries Codes: C -Connected, L -Local, S -Static, R -RIP, B -BGP U -Per-user Static route I1 -ISIS L1, I2 -ISIS L2, IA -ISIS inter area, IS -ISIS summary O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 -OSPF ext 2 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2

O 2038:YY:YY:11::/64 [110/65]

via FE80::ZZZZ:ZZZZ:ZZZZ, Serial0/0

C 2038:YY:YY:61::/64 [0/0]

via ::, Serial0/0/0

L 2038:YY:YY:61::6/128 [0/0]

via ::, Serial0/0/0

C 2038:YY:YY:66::/64 [0/0]

via ::, Gi0/0

L 2038:YY:YY:66::6/128 [0/0]

via ::, Gi0/0

L FE80::/10 [0/0]

via ::, Null0

L FF00::/8 [0/0]

via ::, Null0

RackYYR1:

ipv6 unicast-routing

ipv6 Router ospf YY

```
Router-id YY.YY.1.1
interface Serial0/0/0
ipv6 address 2033:YY:YY:16::1/64
ipv6 ospf network point-to-point
ipv6 ospf 8 area 0
Frame-Relay map ipv6 2038:YY:YY:16::6 106 broadcast
Frame-Relay map ipv6 "9i nk_l ocđj± 106 br oadcæ
interface g0/1
ipv6 address 2033:YY:YY:21::1/64
ipv6 ospf 1 area 0
```

```
RackYYR6:
ipv6 unicast-routing
ipv6 Router ospf YY
Router-id YY.YY.6.6
interface Serial0/0/0
ipv6 address 2033:YY:YY:16::6/64
ipv6 ospf network point-to-point
ipv6 ospf 6 area 0
Frame-Relay map ipv6 2033:YY:YY:16::1 601 broadcast
Frame-Relay map ipv6 "9i nk_l ocđj± 601 br oadcæ
interface g0/1
ipv6 address 2033:YY:YY:62::6/64
Ipv6 ospf 6 area 0
Verify: show ipv6 interface brief; show ipv6 route
```

BGP

```
R1: Loopback 200: 200.1YY.101.1/32
R2: Loopback 200: 200.1YY.102.1/32
R3: Loopback 200: 200.YY.3.1/32
R4: Loopback 200: 200.YY.4.1/32
R5: Loopback 200: 200.YY.5.1/32
R6: Loopback 200: 200.1YY.106.1/32
```

7.1 IBGP

Configure IBGP as follows

AS YY: Configure only R3 R4, and R5 to be part of the AS YY ,R3 is the Route-Reflector for this AS

AS 1YY: Configure only R1, R2 and R6 to be part of the AS 1YY. Don't configure RR or confederation in the AS

You can use any IP address to form the IBGP peers

Advertise the loopback 200 on all BGP routers through BGP and make sure you are able to ping these loopbacks from inside each AS

Loopback 200:

AS YY: 200.YY.X.1/32

AS 1YY:200:1YY.10X.1/32

RackYYR3:

Router bgp YY

no auto-summary

no synchronization

bgp Router-id YY.YY.3.3

network 200.YY.3.1 mask 255.255.255.255

neighbor YY.YY.4.4 remote-as YY

neighbor YY.YY.4.4 update-source loop0

neighbor YY.YY.4.4 route-reflector-client

neighbor YY.YY.5.5 remote-as yy

neighbor YY.YY.5.5 update-source loop0

neighbor YY.YY.5.5 route-reflector-client

RackYYR4:

Router bgp YY

no auto-summary

no synchronization

bgp Router-id YY.YY.4.4

network 200.YY.4.1 mask 255.255.255.255

neighbor YY.YY.3.3 remote-as YY

neighbor YY.YY.3.3 update-source Loopback0

RackYYR5:

Router bgp YY

no auto-summary

no synchronization

bgp Router-id YY.YY.5.5

network 200.YY.5.1 mask 255.255.255.255

neighbor YY.YY.3.3 remote-as YY

neighbor YY.YY.3.3 update-source Loopback0

If use peer - group, configuration is as follows:

RACK08R3#router bgp 8

no synchronization

bgp router-id 8.8.3.3

bgp log-neighbor-changes

network 200.8.3.1 mask 255.255.255.255

neighbor zhenglei peer-group

neighbor zhenglei remote-as 8

```
neighbor zhenglei update-source Loopback0
neighbor zhenglei route-reflector-client
neighbor 8.8.4.4 peer-group zhenglei
neighbor 8.8.5.5 peer-group zhenglei
no auto-summary
```

RackYYR1:

```
Router bgp 10YY
no auto-summary
no synchronization
bgp Router-id YY.YY.1.1
network 200.1YY.101. 1 mask 255.255.255.255
neighbor YY.YY.2.2 remote-as 10YY
neighbor YY.YY.2.2 update-source Loopback0
neighbor YY.YY.6.6 remote-as 10YY
neighbor YY.YY.6.6 update-source Loopback0
```

RackYYR2:

```
Router bgp 10YY
no auto-summary
no synchronization
bgp Router-id YY.YY.2.2
network 200.1YY.102.1 mask 255.255.255.255
neighbor YY.YY.1.1 remote-as 10YY
neighbor YY.YY.1.1 update-source Loopback0
neighbor YY.YY.6.6 remote as 10YY
neighbor YY.YY.6.6 update-source Loopback0
```

RackYYR6:

```
Router bgp 10YY
no auto-summary
no synchronization
bgp Router-id YY.YY.6..6
network 200.1YY.106.1 mask 255.255.255.255
neighbor YY.YY.1.1 remote-as 10YY
neighbor YY.YY.1.1 update-source Loopback0
neighbor YY.YY.2.2 remote-as 10YY
neighbor YY.YY.2.2 update-source Loopback0
Verify: show ip bgp summary; show ip bgp
```

7.2 EGP

Configure EBGP as follows

R6 EBGP peers with BB2 IP address 150.2.YY.254 AS 254

R1 EBGPeers with R3

R2 EBGPeers with R3

You can use any IP address to form the EBGPeers

Make sure all routers in AS YY have the EBGPeers routes from AS 254 via 1YY on their BGP and IP routing tables. You do not need to ping these routes

Make sure you are able to ping the loop back 200 from all BGP routers on both AS. You are permitted to use 4 static routes within minimum mask to fulfill this Requirement

RackYYR6:

Router bgp 10YY

neighbor 150.2.YY.254 remote-as 254

neighbor 150.2.YY.254 local-as YY no-prepend

RackYYR1:

Router bgp 10YY

neighbor YY.YY.3.3 remote-as YY

neighbor YY.YY.3.3 ebgp-multihop 255

neighbor YY.YY.3.3 update-source Loopback0

RackYYR2:

Router bgp 10YY

neighbor YY.YY.3.3 remote-as YY

neighbor YY.YY.3.3 ebgp-multihop 255

neighbor YY.YY.3.3 update-source Loopback0

RackYYR3:

Router bgp YY

Neighbor YY.YY.1.1 remote-as 10YY

Neighbor YY.YY.2.2 remote-as 10YY

Neighbor YY.YY.1.1 update-source loop0

Neighbor YY.YY.2.2 update-source loop0

Neighbor YY.YY.1.1 ebgp-multihop 255

Neighbor YY.YY.2.2 ebgp-multihop 255

RackYYSw1:

Ip route 200.1YY.100.0 255.255.252.0 valn 100

Ip route 200.1yy.106.1 255.255.255.255 Y.Y.1.1

RackYYSw2:

Ip route 200.1YY.100.0 255.255.252.0 valn 100

Ip route 200.1yy.106.1 255.255.255.255 Y.Y.1.1

Verify: show ip bgp; ping all lo200 in ASYY and AS1YY

7.3 Path Selection

Configure R1 so it informs AS YY that the routes 200.1YY.101.1 and 200.1YY.106.1 are to preferable be reached via R1

Configure R2 so it informs AS YY ,that the routes 200.1YY.102.1 are to preferable be reached via R2

Route filtering is not permitted, DO NOT change any attributes coming from BGP AS 254

RackYYR1: ;

```
ip prefix-list r2loop seq 5 permit 200.1YY.102.1/32
route-map MED permit 10
match ip address prefix r2loop
set metric 100
route-map MED permit 20
```

Router bgp 10YY

```
neighbor YY.YY.3.3 route-map MED out
```

RackYYR2:

```
ip prefix-list r1r6loop per 200.1yy.101.1/32
ip prefix-list r1r6loop per 200.1yy.106.1/32
route-map MED permit 10
match ip address prefix r1r6loop
set metric 100
route-map MED permit 20
```

Router bgp 10YY

```
neighbor YY.YY.3.3 route-map MED out
```

RackYYR6:

Router bgp 10YY

```
Neighbor YY.YY.1.1 send-community
```

```
Neighbor YY.YY.2.2 send-community
```

RackYYR1:

Router bgp 10YY

```
Neighbor YY.YY.2.2 send-community
```

```
Neighbor YY.YY.3.3 send-community
```

RackYYR2:

Router bgp 10YY

```
Neighbor YY.YY.1.1 send-community
```

```
Neighbor YY.YY.3.3 send-community
```

RackYYR3:

Router bgp YY

Neighbor YY.YY.4.4 send-community

Neighbor YY.YY.5.5 send-community

Verify : show ip bgp; show ip bgp community

Part-3 : IP Feature

(Multicast eight points, eight points, safety QOS8 points, IP properties eight points, total 32 points)

IP IOS feature

3.1 Exception handling

Configure R4 to enable exception handling

Filename:R4-DUMP Username:ccie Password:cisco

Ftp address: 150.1.YY.254

RackYYR4:

ip ftp username ccie

ip ftp password cisco

exception protocol ftp

exception dump150.1.YY.254

exception corefile R4-DUMP

3.2 System logging

- Buffer alert critical emergencies and error
- Set the buffer size to 8192
- Indicate the date and time for each logged entry

RackYYR5:

logging on

logging buffered 8192 errors

clock timezone GMT 8

clock set hh:mm:ss month year

service timestamps log datetime local-time year show-timezone

Verify: show logging;

(All the time is like all equipment in good, I finally see configuration of all the equipment has been a time, and Windows of time)

3.3 DHCP

Configure R5 to provide the following parameters for DHCP client on VLAN_55

IP address
DNS server YY.YY.55.60 and YY.YY.55.67
Domain:cisco.com
Default gateway
Hosts must retain DHCP assigned address 10 days
Permit only secure ARP entries to be installed in R5's ARP table

RackYYR5:

```
Service dhcp
ip dhcp excluded address YY.YY.55.254
ip dhcp excluded address YY.YY.55.60
ip dhcp excluded address YY.YY.55.67
ip dhcp pool cisco
network YY.YY.55.0 255.255.255.0
default-router YY.YY.55.254
dns-server YY.YY.55.60 YY.YY.55.67
domain-name ccie.com
lease 10
update arp
```

Security

6.1 Tracing Traffic Source to Device under Attack

It is suspected that Dos attack is being launched at host 150.3.YY.254 select an appropriate device to configure so that you can start tracing the source of this attack.

Your solution must meet the following criteria

The result of the trace must be sent to syslog once a day

This device is limited to trace to one IP address only

DO NOT configure ACL to achieve this

RackYYR5:

```
ip source-track 150.1.YY.254
ip source-track address-limit 1
ip source-track syslog-interval 1440
Verify: show ip source-track; show ip source-track
```

6.2 IP Fragment Attacking

R4 picked up an BB1 from source is to address, with the attack and destination is a web server: 10.1.y.5 requirements

R4 stop these attacks, and allow the other traffic flow through

Rack11R4:

```
ip access-list extended FRAGMENT
deny ip any host 10.1.yy.5 fragment
permit ip any any
```

```
int g0/0
ip access-group FRAGMENT in
```

6.3 Catalyst Security

On Sw1-Fa0/7 configure 802.1.x authentication meeting the following When clients that do not

```
RackYYSw1:
aaa new-model
aaa authentication dot1x default group radius
aaa authorization network default group radius
dot1x system-auth-control
dot1x guest-vlan supplicant
```

```
int Fa0/7
Switchport mode access
dot1x port-control auto
dot1x guest-vlan 55
dot1x host-mode multi-host
Verify: show dot1x all; show dot1x interface interface-id details
```

QOS

4.1 Congestion Avoidance Notification

Configure R1-S0/0/0 such that is out bound traffic has utilized 75% of total bandwidth.

R1 should sign that the network is congested and the recipients need to slow down sending packets.

DO NOT configure Frame Relay BECN or FECN for this question

```
RackYYR1:
Ip tcp ecn
Policy-map QOS
class class-default
bandwidth percent 75
random-detect
random-detect ecn
interface s0/0/0
no random-detect
service-policy output QOS
Verify: show policy-map interface interface-id
```

4.2 Traffic policing

Client on VLAN_BB1 and VLAN_55 access a URL located on VLAN_BB2 frequently. This

URL is http://www.this website.com/directory.

Select one suitable router to configure, so as to conserve bandwidth meeting the following criteria.

Traffic from this URL back to these clients should not exceed 640000 bits per second.

If the files download from this URL are image file then drop the traffic

You may assume image the names end with the suffix:*.gif*.jpg or *.jpeg

RackYYR6:

```
ip cef
ip access-list extended TRAFFIC
permit ip 150.2.YY.0 0.0.0.255 150.1.YY.0 0.0.0.255
permit ip 150.2.YY.0 0.0.0.255 YY.YY.55.0 0.0.0.255
class-map match-all url
match access-group name TRAFFIC
match protocol http host www.thiSwebsite.com
match protocol http url /directory /*
class-map match-all pic
match class-map url
match protocol http url *.jpg|.jpeg|.gif
policy-map NBAR
class pic
drop
class url
police cir 64000
interface Gi0/1
service-policy input NBAR
ip nbar protocol-discovery
Verify: show policy-map interface interface-id
```

4.3 Discard Eligible and Traffic Shaping

The Frame Relay link on R5 is experiencing heavy congesting. Configure R5 so that the Frame Relay provider does not drop any routing protocol packets during

congesting and if the number of packets in R5's Frame Relay interface queue exceeds 10, then the traffic rate will reduced to 32000 bps.

RackYYR5:

```
access-list 105 deny ospf any any
access-list 105 deny tcp any eq 179 any
access-list 105 deny tcp any any eq 179
access-list 105 deny pim any any
access-list 105 permit ip any any
frame-relay de-list 1 protocol ip list 105
interface s0/0.5
```

```
frame-relay de-group 1 503
map-class frame-relay FRTS
Frame-Relay adaptive-shaping interface-congestion 10
Frame-Relay mincir 32000
interface Serial0/0
Frame-Relay traffic-shaping
interface Serial0/0.5
Frame-Relay interface-dlci 503
class FRTS
Verify: show frame-relay pvc dlci
```

Multicast

5.1 Sparse Mode Multicasting

There is a multicast source for group 224.2.2.2 located at VLAN_BB2 and another source for group 224.3.3.3 located at VLAN_BB3. There are clients on VLAN_55 that would like to access these two groups.

Configure R5,R3, Sw1,R1 and R6 to meet the following requirements

Configure all devices using sparse mode

R1 will be the RP for both multicast groups and R3 will be backup RP. Use the most reliable way to achieve this objective and do not configure RP information statically

R5 needs to be able to ping both 224.2.2.2 and 224.3.3.3

RackYYR6:

```
ip multicast-routing
int g0/1
ip pim sparse-mode
ip igmp join-group 224.2.2.2 (This configuration is used)
int g0/0
ip pim sparse-mode
ip igmp join-group 224.3.3.3 (This configuration is used)
int s0/0/0
ip pim sparse-mode
ip pim nbma-mode
```

RackYYR1:

```
ip multicast-routing
int s0/0/0
ip pim sparse-mode
ip pim nbma-mode
int g0/0
ip pim sparse-mode
int lo200
ip pim sparse-mode
```

```
ip pim send-rp-ann LO200 sco 10 group-list 11
ip pim send-rp-dis LO200 sco 10
access-list 11 per 224.2.2.2
access-list 11 per 224.3.3.3
```

RackYYSw1:

```
ip multicast-routing
int VLAN 100
ip pim sparse-mode
int Fa0/3
ip pim sparse-mode
```

RackYYR3:

```
ip multicast routing
int g0/0
ip pim sparse-mode
int s0/0/0.3
ip pim sparse-mode
ip pim nbma-mode
ip pim dr-priority 200
int lo0
ip pim sparse-mode
ip pim send-rp-ann Loopback0 sco 10 group-list 33
ip pim send-rp-dis Loopback0 sco 10
access-list 33 permit 224.2.2.2
access-list 33 permit 224.3.3.3
```

RackYYR5:

```
ip multicast-routing
int s0/0/0.5
ip pim sparse-mode
ip pim nbma-mode
```

```
int g0/0
```

```
ip pim sparse-mode
```

Verify: show ip pim neighbor; show ip pim rp mapping; R5 ping group addresses 224.2.2.2 and 224.3.3.3;

5.2 Defense against Multicast Dos Attack

There is a concern that hacker launch Dos attack against R5 with multicast group membership traffic. Configure R5 so that accept only 100 IGMP reports at any time but this limit does not apply to the group 224.3.3.3.

RackYYR5:

```
ip access-list extended 105
```

```
permit igmp any host 224.3.3.3
```

```
int g0/0
```

```
ip igmp limit 110 except 105
```

```
Verify: show ip igmp interface interface-id
```