9. Access to select Cisco online documentation is available from your desktop.
10. When you finish the lab exam make sure that all devices are accessible for the grading proctor,byhaving them in EXEC mode and dose the console windows.A device that is not accessible forreading cannot be marked,and this may cause you to lose substantial points.For Cisco IOS XRdevices,make sure that you commit all of your changes and have the prompt in the EXEC mode.Failure to do this,can delay your lab grading.
11. You have 5 hours to complese tis modue Upon finishing the exam,ensure that all devices are accessible.Any device is not accessible for grading purpose may cause you lose substantial

www.passccielab.com DEMO

Section 1:Core routing(25%)

Section 1.1.IS-IS(4 points)

Question：

The Emerald Tier 1 service provider operates in different geographical locations. Enable IS-IS in this network across three Instances (AGG1, CORE and AGG2), given these requirements:

- Make sure all three IS-IS instances are in Layer 2.
- Derive the NET-ID where the last three octets of the device loopback0 address should match the system ID of the node. For example loopback0 address of 10.255.0.1 should have net ID of 49.0001.0010.0255.0000.0001.00
- Only use flexible CLI configuration groups called CCIE-ISIS to simplify the configuration. Use AGG1, CORE and AGG2 IS-IS instances for both IPv4 and IPv6 respectively.
- Use the same IS-IS database for IPv4 and IPv6.
- Make sure to allow IS-IS to be able to use metric values higher than 63. Use a Layer 2 metric of 200 for IPv4 and 400 for IPv6.
- Make sure IS-IS adjacency formation/maintenance is optimized from the bandwidth utilization point of view using flexible CLI configuration groups.
- Make sure not to leak IGP routes between the three IS-IS Instances.
- Interfaces connected between nodes need to be P2P.
- Make sure both route reflectors (RR) are reachable from the IS-IS CORE instance.
- Make sure LDP is enabled in each IS-IS instance.
- Make sure IS-IS only uses a link for forwarding once LDP has converged on the link.(Enable ISIS overload bit (3min=180s)

www.passccielab.com DEMO

Solution:

With requirement:Make sure EDP is enabled in each iS-S instance we should remove the globa

onpreconfigure and configure mpls auto-configurelupder the IS-IS process.

X-PE5, X-PE6, X-P1, X-P2, X-P3, X-P4, X-ASBR1, X-ASBR2:

Remove the globally pre-configured LDP configuration

```
conf t
no mpls ldp
!
mpls ldp"
```

Configure group CCIE-ISIS for those devices

```
group CCIE-ISIS
router isis *.*
set-overload-bit on-startup 180 level-2
is-type level-2-only
log adjacency changes
address-family ipv4 unicast
metric-style wide
mpls ldp auto-config
!
address-family ipv4 unicast
metric-style wide
mpls ldp auto-config
!
address-family ipv6 unicast
metric-style wide
single-topology
!
interface 'Loopback.*'
passive
```

www.passccielab.com DEMO

```
address-family ipv4 unicast
metric 200 level 2
!
address-family ipv6 unicast
metric 400 level 2
!
!
interface 'GigabitEthernet.*'
point-to-point
hello-padding disable
address-family ipv4 unicast
metric 200 level 2
mpls ldp sync level 2
!
address-family ipv6 unicast
metric 400 level 2
mpls ldp sync level 2
!
!
!
end-group
```

XAE5:

Configure ISIS instance.

```
router isis AGG1
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0007.00
interface Loopback0
!
```

www.passccielab.com DEMO

```
interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/1
!
```

X-PE6:

Configure ISIS instances.

```
router isis AGG1
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0003.00
interface Loopback0
!
interface GigabitEthernet0/0/0/2
!
!
router isis CORE
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0003.00
interface Loopback0
!
interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/1
!
interface GigabitEthernet0/0/0/3
!
interface GigabitEthernet0/0/0/4
!
```

www.passccielab.com DEMO

```
!
```

Note: X-P1 needs to configure two ISIS instances (CORE and AGG1) to make sure not to leak IGP routes between the IS-IS Instances.

X-P2:

Configure ISIS instances.

```
router isis AGG2
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0004.00
interface Loopback0
!
interface GigabitEthernet0/0/0/2
!
!
router isis CORE
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0004.00
interface Loopback0
!
interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/1
!
interface GigabitEthernet0/0/0/4
!
interface GigabitEthernet0/0/0/5
!
!
```

www.passccielab.com DEMO

Note: X-P2 needs to configure two ISIS instances (CORE and AGG2) to make sure not to leak IGP routes between the IS-IS Instances.

X-P3:

Configure ISIS instances.

```
router isis AGG1
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0005.00
interface Loopback0
!
interface GigabitEthernet0/0/0/2
!
!
router isis CORE
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0005.00
interface Loopback0
!
interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/1
!
interface GigabitEthernet0/0/0/4
!
!
```

www.passccielab.com DEMO

Note: X-P3 needs to configure two ISIS instances (CORE and AGG1) to make sure not to leak IGP routes between the IS-IS Instances.

X-P4:

Configure ISIS instances.

```
router isis AGG2
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0006.00
interface Loopback0
!
interface GigabitEthernet0/0/0/1
!
!
router isis CORE
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0006.00
interface Loopback0
!
interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/2
!
interface GigabitEthernet0/0/0/4
!
!
```

www.passccielab.com DEMO

Note: X-P4 needs to configure two ISIS instances (CORE and AGG2) to make sure not to leak IGP routes between the IS-IS Instances.

X-ASBR1:

Configure ISIS instances.

```
router isis AGG2
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0001.00
interface Loopback0
!
interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/1
!
!
```

X-ASBR2:

Configure ISIS instances.

```
router isis AGG2
apply-group CCIE-ISIS
net 49.0001.0001.0000.0100.0002.00
interface Loopback0
!
interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/2
!
!
```

X-RR1:

Configure ISIS instance and enable ISIS under relevant interfaces.

```
router isis
```

www.passccielab.com DEMO

```
net 49.0001.0001.0000.0100.0009.00
passive-interface Loopback0
is-type level-2-only
metric-style wide
metric 200 level-2
!
address-family ipv6
metric 400 level-2
exit-address-family
!
interface GigabitEthernet1
ip router isis
ipv6 router isis
isis network point-to-point
no isis hello padding always
!
```

X-RR2:

Configure ISIS instance and enable ISIS under relevant interfaces.

```
router isis
net 49.0001.0001.0000.0100.0010.00
passive-interface Loopback0
is-type level-2-only
metric-style wide
metric 200 level-2
!
address-family ipv6
```

www.passccielab.com DEMO

```
metric 400 level-2
exit-address-family
!
interface GigabitEthernet1
ip router isis
ipv6 router isis
isis network point-to-point
no isis hello padding always
!
```

Verification

1.Verify that ISIS neighbors are"Up"and has routes for both ipv4/ipv6.Loopback

X-RR1,X-RR2 adver into ISIS Core

X-PE5

```
RP/0/0/CPU0:X-PE5#show isis neighbors
Sat Feb 22 06:24:09.057 UTC

IS-IS AGG1 neighbors:
System Id      Interface        SNPA           State Holdtime Type IETF-NSF
X-P1           Gi0/0/0/0        *PtoP*         Up    28       L2   Capable
X-PE6          Gi0/0/0/1        *PtoP*         Up    28       L2   Capable

Total neighbor count: 2

RP/0/0/CPU0:X-PE5#show ip route isis
Sat Feb 22 06:24:31.896 UTC

i L2 1.0.100.3/32 [115/400] via 1.100.5.1, 00:03:22, GigabitEthernet0/0/0/0
i L2 1.0.100.5/32 [115/600] via 1.100.9.6, 00:03:00, GigabitEthernet0/0/0/1
i L2 1.0.100.8/32 [115/400] via 1.100.9.6, 00:03:00, GigabitEthernet0/0/0/1
i L2 1.100.6.0/24 [115/400] via 1.100.9.6, 00:03:00, GigabitEthernet0/0/0/1

RP/0/0/CPU0:X-PE5#show route ipv6 isis
Sat Feb 22 06:24:57.574 UTC

i L2 2620:fc7:1000::3/128
      [115/600] via fe80::52b3:e5ff:fe00:c03, 00:03:49, GigabitEthernet0/0/0/0
i L2 2620:fc7:1000::5/128
      [115/800] via fe80::5274:a0ff:fe00:1002, 00:03:27, GigabitEthernet0/0/0/1
i L2 2620:fc7:1000::8/128
      [115/600] via fe80::5274:a0ff:fe00:1002, 00:03:27, GigabitEthernet0/0/0/1
i L2 2620:fc7:1006::/64
```

```
      [115/600] via fe80::5274:a0ff:fe00:1002, 00:03:27, GigabitEthernet0/0/0/1
```

www.passccielab.com DEMO

X-RR1

```
X-RR1#show isis neighbors
Load for five secs: 0%/0%; one minute: 0%; five minutes: 0%
Time source is hardware calendar, *06:25:52.527 UTC Sat Feb 22 2025


System Id      Type Interface   IP Address      State Holdtime Circuit Id
X-P1           L2   Et0/1       1.100.11.1      UP    23       00

X-RR1#show ip route isis
Load for five secs: 0%/0%; one minute: 0%; five minutes: 0%
Time source is hardware calendar, *06:26:06.524 UTC Sat Feb 22 2025

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 not set

      1.0.0.0/8 is variably subnetted, 15 subnets, 2 masks
i L2     1.0.100.3/32 [115/400] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.0.100.4/32 [115/600] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.0.100.5/32 [115/600] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.0.100.6/32 [115/600] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.0.100.10/32 [115/16777614] via 1.100.11.1, 00:02:27, Ethernet0/1
i L2     1.100.3.0/24 [115/400] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.100.4.0/24 [115/600] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.100.7.0/24 [115/400] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.100.8.0/24 [115/600] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.100.12.0/24 [115/600] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.100.32.0/24 [115/600] via 1.100.11.1, 00:02:46, Ethernet0/1
i L2     1.100.42.0/24 [115/400] via 1.100.11.1, 00:02:46, Ethernet0/1
```

```
IPv6 Routing Table - default - 16 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

www.passccielab.com All rights reserved

```
       la - LISP alt, lr - LISP site-registrations, ld - LISP dyn-eid
       a - Application
I2  2620:FC7:1000::3/128 [115/600]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1000::4/128 [115/800]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1000::5/128 [115/800]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1000::6/128 [115/800]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1000::11/128 [115/16777614]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1003::/64 [115/600]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1004::/64 [115/800]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1007::/64 [115/600]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1008::/64 [115/800]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1012::/64 [115/800]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1032::/64 [115/800]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
I2  2620:FC7:1042::/64 [115/600]
     via FE80::52B3:E5FF:FE00:C04, Ethernet0/1
```

```
RP/0/0/CPU0:X-P1#show mpls ldp igp sync brief
Sat Mar 18 09:39:35.616 UTC

Interface                VRF Name                         Status
-----------------------  -------------------------------  ---------
Gi0/0/0/0                default                          Ready
Gi0/0/0/1                default                          Ready
Gi0/0/0/2                default                          Ready
Gi0/0/0/3                default                          Not ready
Gi0/0/0/4                default                          Ready
```

Note: Not ready which the state of interface connected to RR.

Note: **LDP-IGP Sync shows "Not Ready" in the interface connected to X-RRs on X-P1 and X-P2 and it's fine, there is no LDP neighbor between X-P1,X-P2 and X-RRs.**

www.passc

1.2.Section 1.2:BGP-Unified MPLS(3points)

Question

Implement Unified MPLS across the core and aggregation metro fabricat the Emerald Tier 1SP asper given requirements:

- **Al the X-PE and X-ASBR devices loopbackO should be learnedin different IS-IS domain and notvia IGP**
- **All the X-PE's andX-ASBR devices loopbacko should be reachable via MPLS path.**
- **Create iBGP-LUin core IS-IS process usingX-RR1 and provide redundancy on X-RR2.**
- **Use suitable inline Route reflectors for AGG1and AGG2 processes**

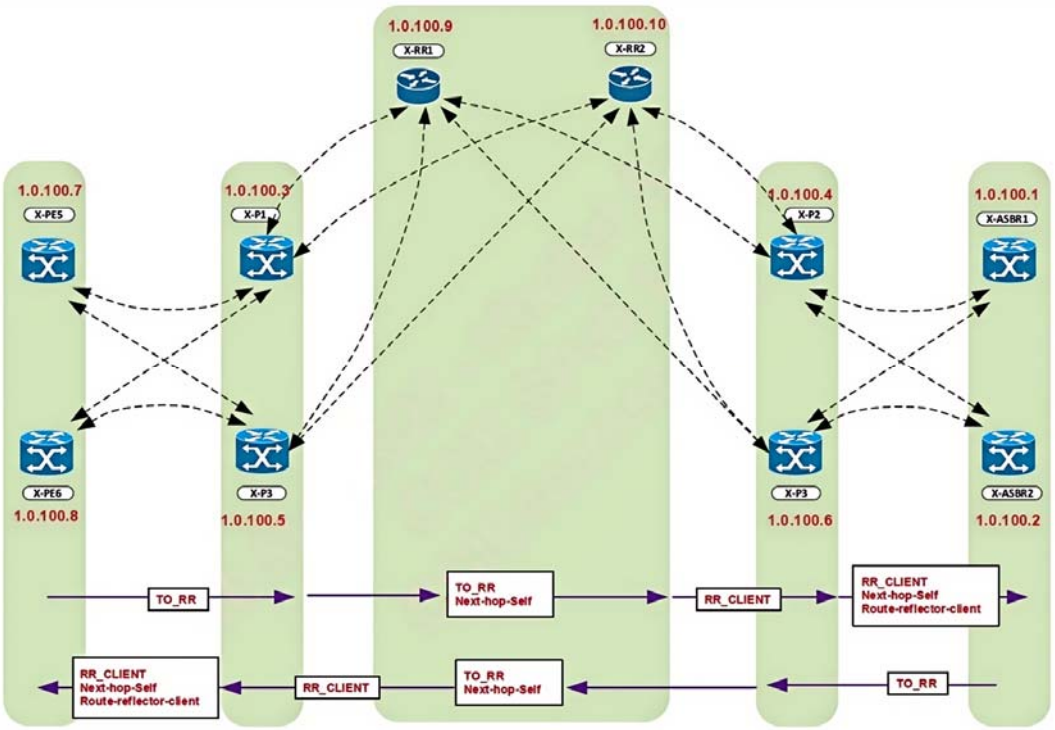

DEMO

Solution

X-PE5:

Configure iBGP-LU.

```
router bgp 100
bgp router-id 1.0.100.7
address-family ipv4 unicast
network 1.0.100.7/32
allocate-label all
!
neighbor-group To-RR
remote-as 100
update-source Loopback0
address-family ipv4 labeled-unicast
!
!
neighbor 1.0.100.3
use neighbor-group To-RR
!
neighbor 1.0.100.5
use neighbor-group To-RR
!
```

www.passccielab.com DEMO

X-PE6:

Configure iBGP-LU.

```
router bgp 100
bgp router-id 1.0.100.8
```