CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Comprehensive Coverage of the CCIE Service Exam

Table of Contents

Module 1 – OSPF	
Lab 1	Configure OSPF on Ethernet - Area 10
Lab 2	Configuring OSPF on Serial Links - Area 10
Lab 3	Configuring OSPF in Area 0
Lab 4	Configuring Unicast-based OSPF
Lab 5	Configuring an OSPF ASBR
Lab 6	Configuring a Multi-Area / Multi-Domain Topology
Lab 7	Configuring Inter-Area Route Summarization
Lab 8	Configuring External Route Summarization
Lab 9	LSA Type 3 Filtering
Lab 10	Configuring OSPF Authentication
Lab 11	Configuring OSPF Area Types
Lab 12	Configuring Virtual Link
Lab 13	Configuring BFD for OSPF
Lab 14	Configuring IP FRR - OSPF
Module 2 – IS-IS	
Lab 1	Basic IS-IS Configuration with Areas
Lab 2	Optimizing IS-IS
Lab 3	IS-IS Authentication
Lab 4	Configure Inter-area Interfaces
Lab 5	IS-IS Multi-Area / Multi-Domain Configuration
Lab 6	Configure Route Leaking
Lab 7	Route Summarization
Lab 8	Configuring BFD for IS-IS
Module 3 – BGP	
Lab 1	Configuring eBGP
Lab 2	Configuring eBGP Multi-Hop
Lab 3	Redistributing Networks into BGP
Lab 4	Configuring BGP Authentication
Lab 5	Configuring iBGP with Route Reflectors
Lab 6	Route Filtering using ACLs
Lab 7	Route Filtering using Prefix-Lists
Lab 8	Route Filtering using AS Path-Filter
Lab 9	Configuring Route Aggregation – Summary Only

Lab 10	Configuring Route Aggregation – Manual Filtering
Lab 11	Configuring Route Aggregation – Suppress Maps
Lab 12	Configuring Base BGP Topology – eBGP & iBGP
Lab 13	Configuring BGP Attributes – Local Preference
Lab 14	Configuring BGP Attributes – MED
Lab 15	Configuring BGP Attributes – Weight
Lab 16	Configuring BGP Attributes – AS-Path
Lab 17	Configuring BGP Attributes – No-Export Community
Lab 18	Configuring BGP Attributes – No-Advertise Community
Lab 19	Configuring BGP Conditional Advertisement
Lab 20	Configuring BGP Multi-Path – eBGP – iBGP
Lab 21	Configuring to Redistribute iBGP Routes into IGP
Lab 22	Configuring BGP Route Reflector with Next-Hop Changed
Lab 23	Configuring BGP Route Reflection based on Dynamic Neighbors
Lab 24	Working with Private AS Numbers
Lab 25	Configuring the Local-AS Command
Lab 26	Configuring BFD for BGP
Lab 27	Configuring BGP Confederations
Module 4 – IPv6	
Lab 1	Configuring IPv6 Addressing
Lab 2	Configuring OSPFv3
Lab 3	Configuring EIGRP for IPv6
Lab 4	Configuring IS-IS for IPv6
Lab 5	Configuring BGP for IPv6
Lab 6	Configuring IPv6IP Tunneling
Lab 7	Configuring NAT64
Module 5 – Configuring MPLS Unicast Routing	
Lab 1	Configuring MPLS Unicast Routing
Lab 2	Authenticating LDP Peers
Module 6 – Configuring Intra-AS MPLS VPNs	
Lab 1	Configuring MPLS VPN – PE-CE Using Static Routing
Lab 2	Configuring MPLS VPN – PE-CE Using EIGRP Routing
Lab 3	Configuring MPLS VPN – PE-CE Using IS-IS Routing
Lab 4	Configuring MPLS VPN – PE-CE Using BGP Routing – 1
Lab 5	Configuring MPLS VPN – PE-CE Using BGP Routing – 2
Lab 6	Configuring MPLS VPN – PE-CE Using OSPF

Lab 7	Configuring MPLS VPN – PE-CE Using OSPF – Domain-ID
Lab 8	Configuring MPLS VPN – PE-CE Using OSPF – Sham-link
Lab 9	Configuring MPLS VPN Extranets
Module 7 - Configuring MPLS VPNs for IPv6 Networks	
Lab 1	Configuring MPLS Unicast Routing – IPv4
Lab 2	Configuring MPLS VPNv6
Module 8 – Configuring Inter-AS MPLS VPNs	
Lab 1	Configuring MPLS Unicast Routing
Lab 2	Configuring Intra-AS MPLS VPNv4 – AS 100 & AS 200
Lab 3	Configuring Inter-AS MPLS VPN – Option A – VRF-based
Lab 4	Configuring Inter-AS MPLS VPN – Option B – MP-eBGP on ASBR
Lab 5	Configuring Inter-AS MPLS VPN – Option C MP-eBGP on RR
Lab 6	Configuring Inter-AS MPLS VPN – Using a Non-VPN Transit Provider
Module 9 - Carrier Supporting Carrier (CSC)	
Lab 1	Configuring Carrier Supporting Carriers (CSC) – IP Only
Lab 2	CSC – MPLS Unicast Routing
Lab 3	CSC – MPLS VPN (H-MPLS VPN)
Module 10 - Configuring MPLS on IOS-XR	
Lab 1	Configuring SP Core Networks with MPLS & MP-iBGP
Lab 2	Configuring Intra-AS MPLS VPN within AS 100 & AS 200
Lab 3	Configuring Inter-AS MPLS VPN – Option C
Module 11 – Segment Routing & Large Scale MPLS	
Lab 1	Basic Intra-AS MPLS VPN using LDP
Lab 2	Configuring Segment Routing using OSPF
Lab 3	Configuring Segment Routing using IS-IS
Lab 4	Configuring SR-LDP Mapping
Lab 5	Configuring Segment Routing using BGP
Lab 6	Configuring Large Scale MPLS / Unified MPLS
Module 12 - Configuring Multicast VPN	
Lab 1	Configuring the SP Core Network with Intra-AS MPLS VPN
Lab 2	Configuring Multicast VPN (M-VPN) using Static RP
Lab 3	Configuring Multicast VPN (M-VPN) using Auto RP
Lab 4	Configuring Mutticast VPN using mLDP

Module 13 - Configuring MPLS TE	
Lab 1	Configuring the SP Core Network with MPLS
Lab 2	Configuring MPLS TE - Static Tunnels
Lab 3	Configuring MPLS TE - Dynamic Tunnels
Lab 4	Configuring MPLS TE - Dynamic Tunnels with Pre-emption
Lab 5	Configuring MPLS TE with IS-IS
Lab 6	Configuring MPLS TE - Link Protection with Fast Re-Route (FRR)
Module 14 – Quality of Service (QoS)	
Lab 1	Configuring the SP Core Network with MPLS
Lab 2	Configuring MPLS QoS – Uniform Mode
Lab 3	Configuring MPLS QoS – Long Pipe Mode
Lab 4	Configuring MPLS QoS – Short Pipe Mode
Module 15 - Configuring Layer 2 MPLS VPNs	
Lab 1	Implementing AToM – Ethernet VLAN over MPLS
Lab 2	Implementing AToM – Ethernet over MPLS
Lab 3	Implementing AToM – PPP over MPLS
Lab 4	Implementing AToM – Interworking PPP-Ethernet
Lab 5	Configuring VPLS

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring OSPF

Lab 1 – Configure OSPF on Ethernet – Area 10

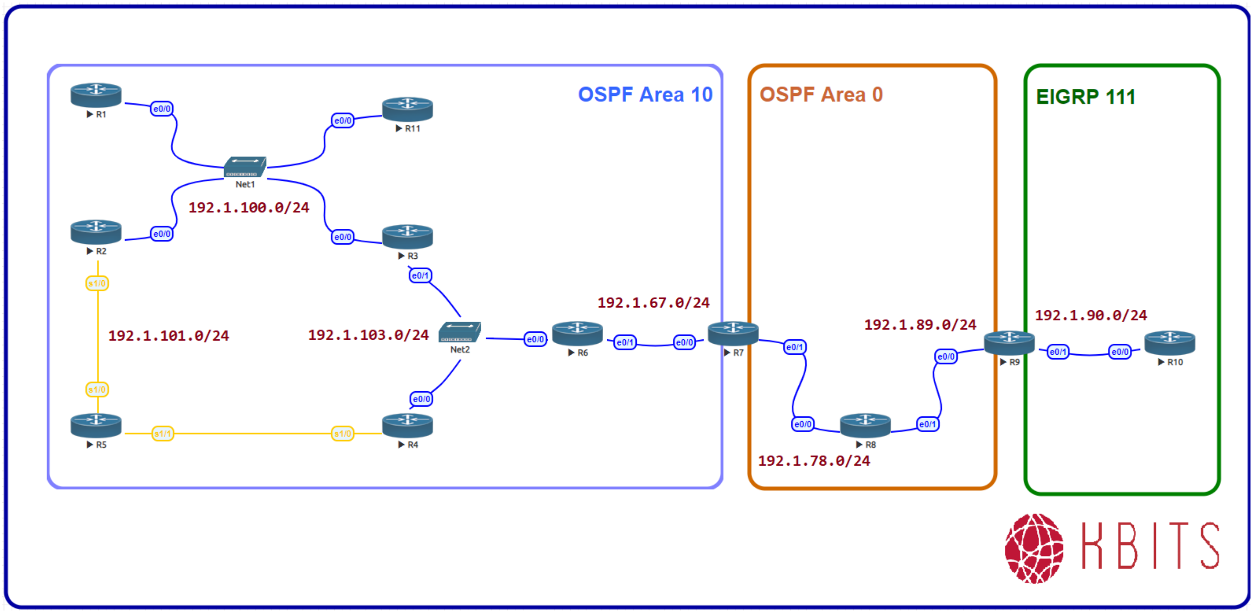

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
E 0/0	192.1.100.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
E 0/0	192.1.100.2	255.255.255.0
S 1/0	192.1.101.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
E 0/0	192.1.100.3	255.255.255.0
E 0/1	192.1.103.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
E 0/0	192.1.103.4	255.255.255.0
S 1/0	192.1.102.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
S 1/0	192.1.101.5	255.255.255.0
S 1/1	192.1.102.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
E 0/0	192.1.103.6	255.255.255.0
E 0/1	192.1.67.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
E 0/0	192.1.67.7	255.255.255.0
E 0/1	192.1.78.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
E 0/0	192.1.78.8	255.255.255.0
E 0/1	192.1.89.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	9.9.9.9	255.0.0.0
E 0/0	192.1.89.9	255.255.255.0

E 0/1	192.1.90.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	10.10.10.10	255.0.0.0
E 0/0	192.1.90.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	11.11.11.11	255.0.0.0
E 0/0	192.1.100.11	255.255.255.0

Task 1

Configure OSPF all the Broadcast Multi-Access (BMA) Ethernet network in Area 10. Enable OSPF on all loopbacks on all routers. Hard Code the Router-id based on the following:

R1 – 0.0.0.1
R2 – 0.0.0.2
R3 – 0.0.0.3
R4 – 0.0.0.4
R6 – 0.0.0.6
R7 – 0.0.0.7
R11 – 0.0.0.11

R1 Router OSPF 1 Router-id 0.0.0.1 Network 1.0.0.0 0.255.255.255 area 10 Network 192.1.100.0 0.0.0.255 area 10	**R2** Router OSPF 1 Router-id 0.0.0.2 Network 2.0.0.0 0.255.255.255 area 10 Network 192.1.100.0 0.0.0.255 area 10
R3 Router OSPF 1 Router-id 0.0.0.3 Network 3.0.0.0 0.255.255.255 area 10 Network 192.1.100.0 0.0.0.255 area 10 Network 192.1.103.0 0.0.0.255 area 10	**R4** Router OSPF 1 Router-id 0.0.0.4 Network 4.0.0.0 0.255.255.255 area 10 Network 192.1.103.0 0.0.0.255 area 10
R6 Router OSPF 1 Router-id 0.0.0.6	**R7** Router OSPF 1 Router-id 0.0.0.7

Network 6.0.0.0 0.255.255.255 area 10 Network 192.1.103.0 0.0.0.255 area 10 Network 192.1.67.0 0.0.0.255 area 10	Network 7.0.0.0 0.255.255.255 area 10 Network 192.1.67.0 0.0.0.255 area 10

Task 2

Configure the routers such that R1 becomes the DR and R2 as the BDR on the 192.1.100.0/24 Network. R3 should be the DR & R4 should be the BDR for the 192.1.103.0/24 network.

R1 Interface E 0/0 Ip ospf priority 100	**R2** Interface E 0/0 Ip ospf priority 50
R3 Interface E 0/1 Ip ospf priority 100	**R4** Interface E 0/0 Ip ospf priority 50

Note: Issue the **Clear ip ospf process** command to reset the OSPF process for the change to take effect.

Lab 2 – Configuring OSPF on Serial Links – Area 10

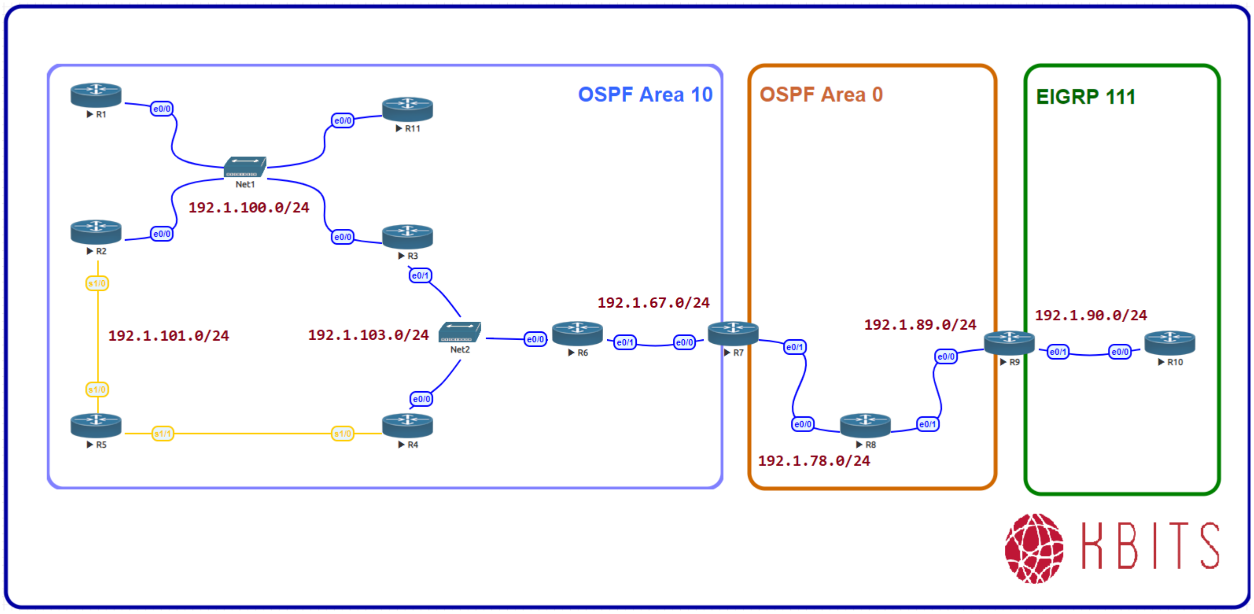

Task 1

Run OSPF as your Routing Protocol on the Serial Networks between R2, R4 & R5 in Area 10. Enable OSPF on the Loopback interface on R5. Configure the Router ID of R5 as 0.0.0.5.

R2 router ospf 1 network 192.1.101.0 0.0.0.255 area 10
R4 router ospf 1 network 192.1.102.0 0.0.0.255 area 10
R5 router ospf 1 router-id 0.0.0.5 network 5.0.0.0 0.255.255.255 area 10 network 192.1.101.0 0.0.0.255 area 10 network 192.1.102.0 0.0.0.255 area 10

Lab 3 – Configuring OSPF in Area 0

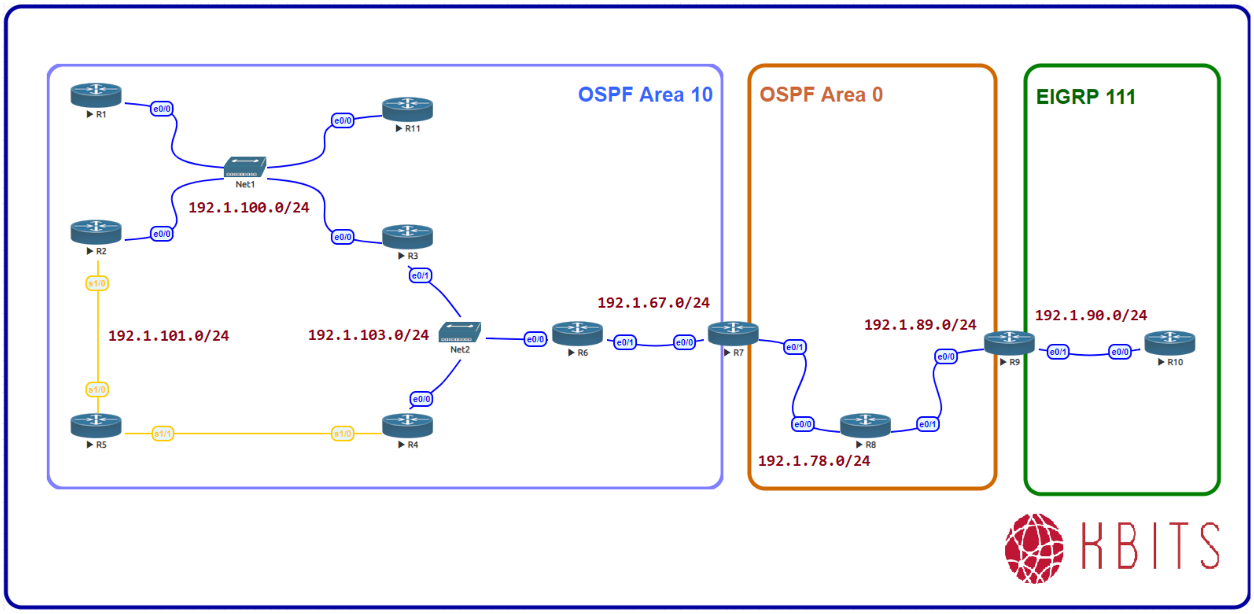

Task 1

Configure R7, R8 & R9 in Area 0. Don't enable the Loopback Interface of R9 in OSPF. The Router ID's for R8 & R9 should be 0.0.0.8 & 0.0.0.9 respectively. Make sure that the neighbor relationships in Area 0 are established bypassing the DR & BDR election wait time.

R7

```
router ospf 1
 network 192.1.78.0 0.0.0.255 area 0
!
Interface E 0/1
 Ip ospf network point-to-point
```

R8

```
router ospf 1
 router-id 0.0.0.8
 network 8.0.0.0 0.255.255.255 area 0
 network 192.1.78.0 0.0.0.255 area 0
 network 192.1.89.0 0.0.0.255 area 0
!
Interface E 0/0
 Ip ospf network point-to-point
```

! Interface E 0/1 Ip ospf network point-to-point
R9 router ospf 1 router-id 0.0.0.9 network 192.1.89.0 0.0.0.255 area 0

Task 2

Make sure that all OSPF Loopbacks networks appear with the Interface mask. They should not appear as a Host Route.

R1 Interface Loopback0 Ip ospf network point-to-point	**R2** Interface Loopback0 Ip ospf network point-to-point
R3 Interface Loopback0 Ip ospf network point-to-point	**R4** Interface Loopback0 Ip ospf network point-to-point
R5 Interface Loopback0 Ip ospf network point-to-point	**R6** Interface Loopback0 Ip ospf network point-to-point
R7 Interface Loopback0 Ip ospf network point-to-point	**R8** Interface Loopback0 Ip ospf network point-to-point
R9 Interface Loopback0 Ip ospf network point-to-point	**R11** Interface Loopback0 Ip ospf network point-to-point

Lab 4 – Configuring Unicast-based OSPF

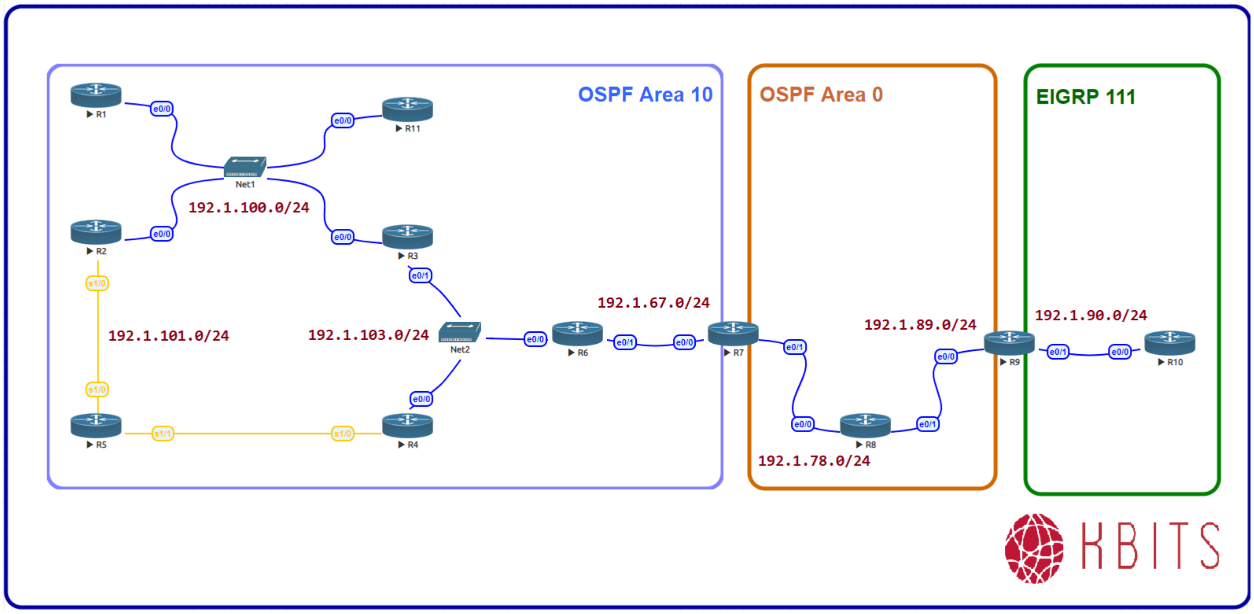

Task 1

Configure Unicast-based OSPF between R6 & R7.

R6 Interface E 0/1 Ip ospf network non-broadcast ! Router ospf 1 Neighbor 192.1.67.7
R6 Interface E 0/1 Ip ospf network non-broadcast ! Router ospf 1 Neighbor 192.1.67.6

Lab 5 – Configuring an OSPF ASBR

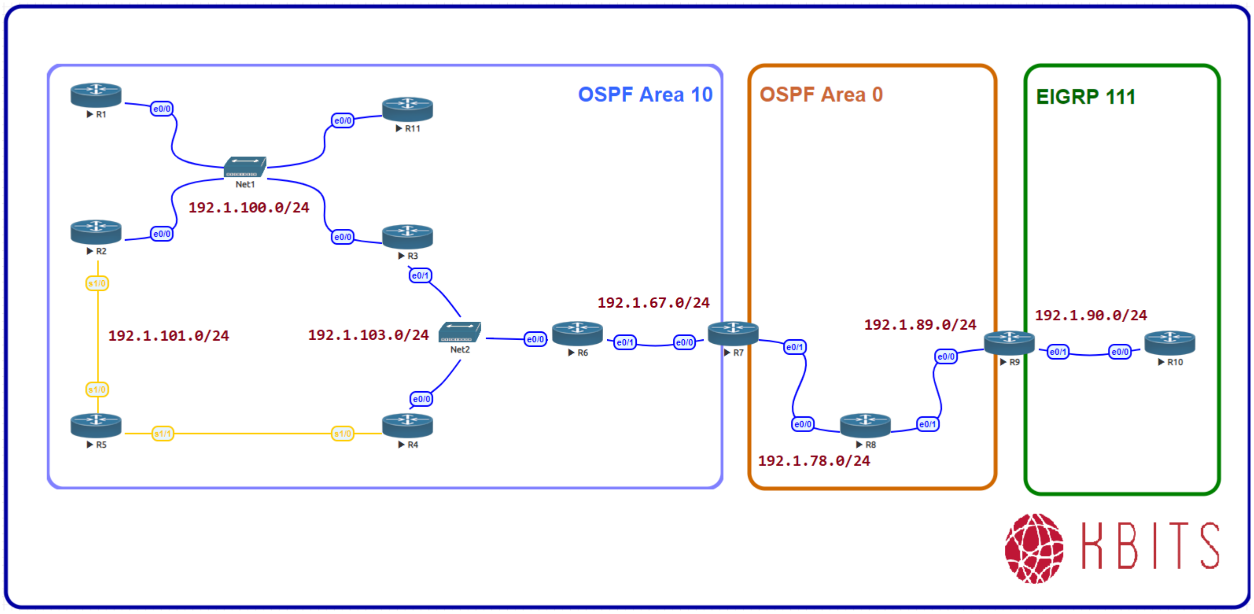

Task 1

Configure EIGRP in AS 111 between R9 & R10. Enable all loopbacks on the 2 routers in EIGRP.

R9

```
Router eigrp 111
 Network 192.1.90.0
 Network 9.0.0.0
```

R10

```
Router eigrp 111
 Network 192.1.90.0
 Network 10.0.0.0
```

Task 2

Configure Mutual Route Redistribution between OSPF & EIGRP on R9. Use Seed Metrics of your choice.

R9

```
Router eigrp 111
 Redistribute ospf 1 metric 10 10 10 10 10
```

```
!
Router ospf 1
 Redistribute eigrp 111 subnets
```

Verification:

Verify the OSPF Database for appropriate LSAs on the appropriate routers.

Lab 6 – Configuring a Multi-Area / Multi-Domain Topology

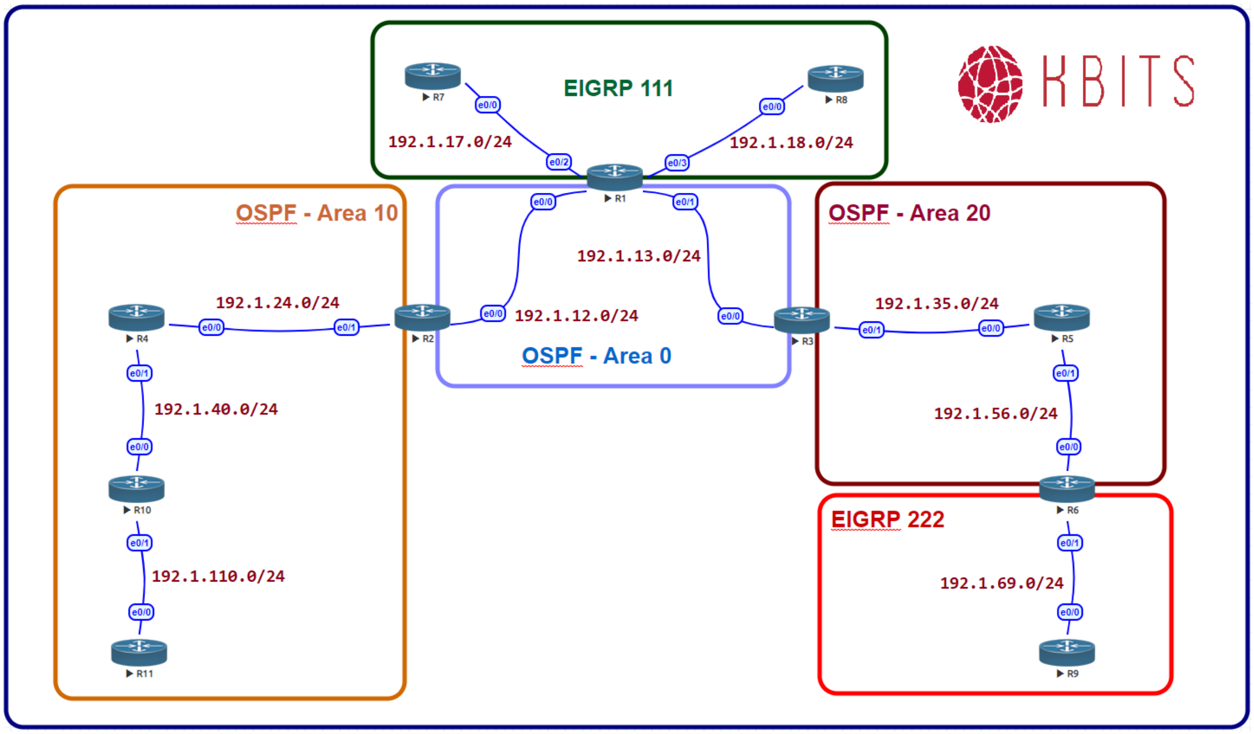

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 1	11.11.11.11	255.0.0.0
E 0/0	192.1.12.1	255.255.255.0
E 0/1	192.1.13.1	255.255.255.0
E 0/2	192.1.17.1	255.255.255.0
E 0/3	192.1.18.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
Loopback 1	22.22.22.22	255.0.0.0

E 0/0	192.1.12.2	255.255.255.0
E 0/1	192.1.24.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
Loopback 1	33.33.33.33	255.0.0.0
E 0/0	192.1.13.3	255.255.255.0
E 0/1	192.1.35.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
E 0/0	192.1.24.4	255.255.255.0
E 0/1	192.1.40.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
E 0/0	192.1.35.5	255.255.255.0
E 0/1	192.1.56.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
Loopback 1	66.66.66.66	255.0.0.0
E 0/0	192.1.56.6	255.255.255.0
E 0/1	192.1.69.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
Loopback 1	107.7.72.1	255.255.255.0
Loopback 2	107.7.73.1	255.255.255.0
Loopback 3	107.7.74.1	255.255.255.0

Loopback 4	107.7.75.1	255.255.255.0
E 0/0	192.1.17.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
Loopback 1	88.88.88.88	255.0.0.0
E 0/0	192.1.18.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	9.9.9.9	255.0.0.0
Loopback 1	99.99.99.99	255.0.0.0
E 0/0	192.1.69.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	10.10.10.10	255.0.0.0
E 0/0	192.1.40.10	255.255.255.0
E 0/1	192.1.110.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	111.111.100.11	255.255.255.0
Loopback 1	111.111.101.11	255.255.255.0
Loopback 2	111.111.102.11	255.255.255.0
Loopback 3	111.111.103.11	255.255.255.0
E 0/0	192.1.110.11	255.255.255.0

Task 1

Configure OSPF in Area 0 between R1, R2 & R3. Besides the physical links, enable the Loopback 0 interfaces of all 3 routers in Area 0. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R1 – 0.0.0.1
R2 – 0.0.0.2
R3 – 0.0.0.3

R1	R2
Router OSPF 1 Router-id 0.0.0.1 Network 1.0.0.0 0.255.255.255 area 0 Network 192.1.12.0 0.0.0.255 area 0 Network 192.1.13.0 0.0.0.255 area 0 ! Interface Loopback0 Ip ospf network point-to-point	Router OSPF 1 Router-id 0.0.0.2 Network 2.0.0.0 0.255.255.255 area 0 Network 192.1.12.0 0.0.0.255 area 0 ! Interface Loopback0 Ip ospf network point-to-point
R3 Router OSPF 1 Router-id 0.0.0.3 Network 3.0.0.0 0.255.255.255 area 0 Network 192.1.13.0 0.0.0.255 area 0 ! Interface Loopback0 Ip ospf network point-to-point	

Task 2

Configure OSPF in Area 10 between R2, R4, R10 & R11. Besides the physical links, enable the Loopback 1 interface on R2 and all the loopbacks of the other 3 routers in Area 10. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R4 – 0.0.0.4
R10 – 0.0.0.10
R11 – 0.0.0.11

R2	R4
Router OSPF 1 Network 192.1.24.0 0.0.0.255 area 10	Router OSPF 1 Router-id 0.0.0.4 Network 4.0.0.0 0.255.255.255 area 10 Network 192.1.24.0 0.0.0.255 area 10 Network 192.1.40.0 0.0.0.255 area 10 ! Interface Loopback0 Ip ospf network point-to-point
R10 Router OSPF 1 Router-id 0.0.0.10 Network 10.0.0.0 0.255.255.255 area 10	**R11** Router OSPF 1 Router-id 0.0.0.11 Network 111.111.0.0 0.0.255.255 area 10

Network 192.1.40.0 0.0.0.255 area 10 Network 192.1.110.0 0.0.0.255 area 10 ! Interface Loopback0 Ip ospf network point-to-point	Network 192.1.110.0 0.0.0.255 area 10 ! Interface Loopback0 Ip ospf network point-to-point ! Interface Loopback1 Ip ospf network point-to-point ! Interface Loopback2 Ip ospf network point-to-point ! Interface Loopback3 Ip ospf network point-to-point

Task 3

Configure OSPF in Area 20 between R3, R5 & R6. Besides the physical links, enable the Loopback 0 interface on R3 & R6 and all the loopbacks on R5 in Area 20. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R5 – 0.0.0.5
R6 – 0.0.0.6

R5 Router OSPF 1 Network 192.1.35.0 0.0.0.255 area 20	**R5** Router OSPF 1 Router-id 0.0.0.5 Network 5.0.0.0 0.255.255.255 area 20 Network 192.1.35.0 0.0.0.255 area 20 Network 192.1.56.0 0.0.0.255 area 20 ! Interface Loopback0 Ip ospf network point-to-point
R6 Router OSPF 1 Router-id 0.0.0.6 Network 6.0.0.0 0.255.255.255 area 20 Network 192.1.56.0 0.0.0.255 area 20 ! Interface Loopback0 Ip ospf network point-to-point	

Task 4

Configure EIGRP is AS 111 between R1, R7 & R8. Enable all loopbacks on R7 & R8 in EIGRP 111. Enable Loopback 1 on R1 in EIGRP 111.

R1 Router EIGRP 111 Network 192.1.17.0 Network 192.1.18.0 Network 11.0.0.0	**R7** Router EIGRP 111 Network 192.1.17.0 Network 7.0.0.0 Network 107.0.0.0
R8 Router EIGRP 111 Network 192.1.18.0 Network 8.0.0.0 Network 88.0.0.0	

Task 5

Configure EIGRP is AS 222 between R6 & R9. Enable all loopbacks on R9 in EIGRP 222. Enable Loopback 1 on R6 in EIGRP 222.

R6 Router EIGRP 222 Network 192.1.69.0 Network 66.0.0.0	**R9** Router EIGRP 111 Network 192.1.69.0 Network 9.0.0.0 Network 99.0.0.0

Task 6

Configure Mutual Redistribution between the appropriate routers to allow end-to-end connectivity between all routing domains. Use Seed metric of your choice.

R1 Router ospf 1 Redistribute eigrp 111 subnets ! Router eigrp 111 Redistribute ospf 1 metric 10 10 10 10 10	**R6** Router ospf 1 Redistribute eigrp 222 subnets ! Router eigrp 222 Redistribute ospf 1 metric 10 10 10 10 10

Lab 7 – Configuring Inter-Area Route Summarization

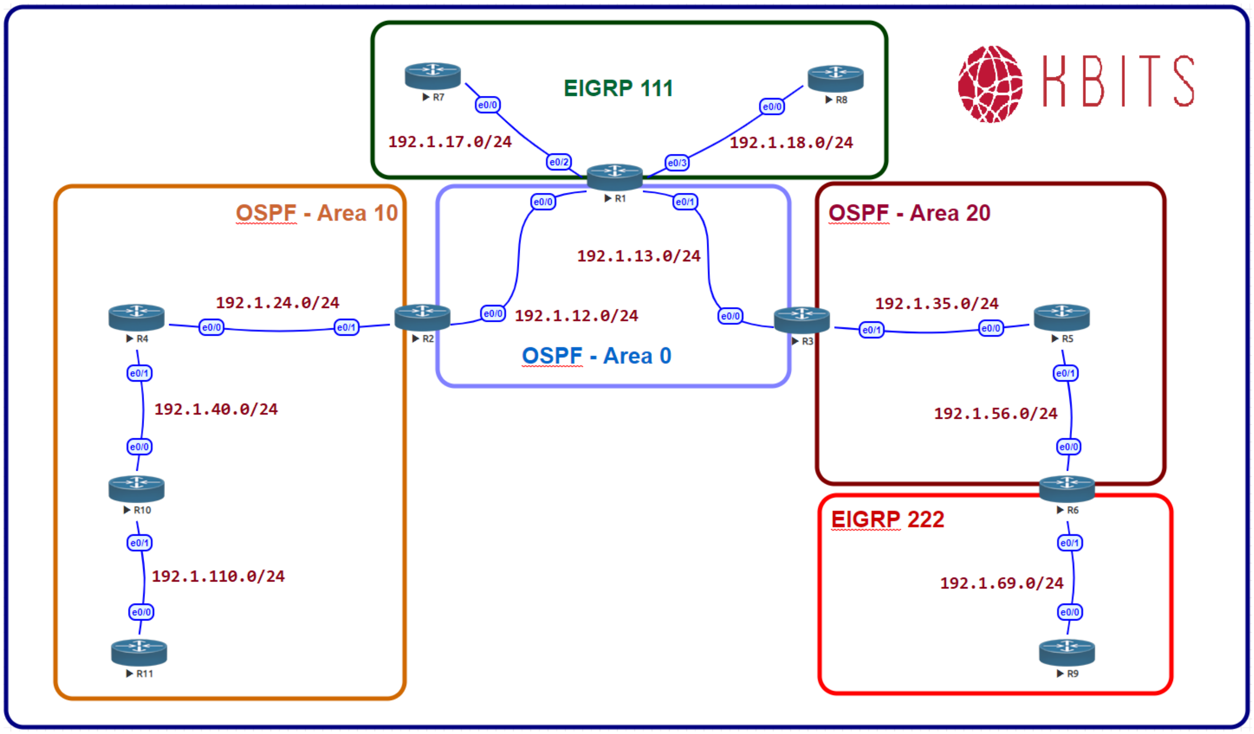

Task 1

Configure Route Summarization on the appropriate ABR to summarize all the R11 Loopbacks.

R2 Router ospf 1 Area 10 range 111.111.100.0 255.255.252.0

Lab 8 – Configuring External Route Summarization

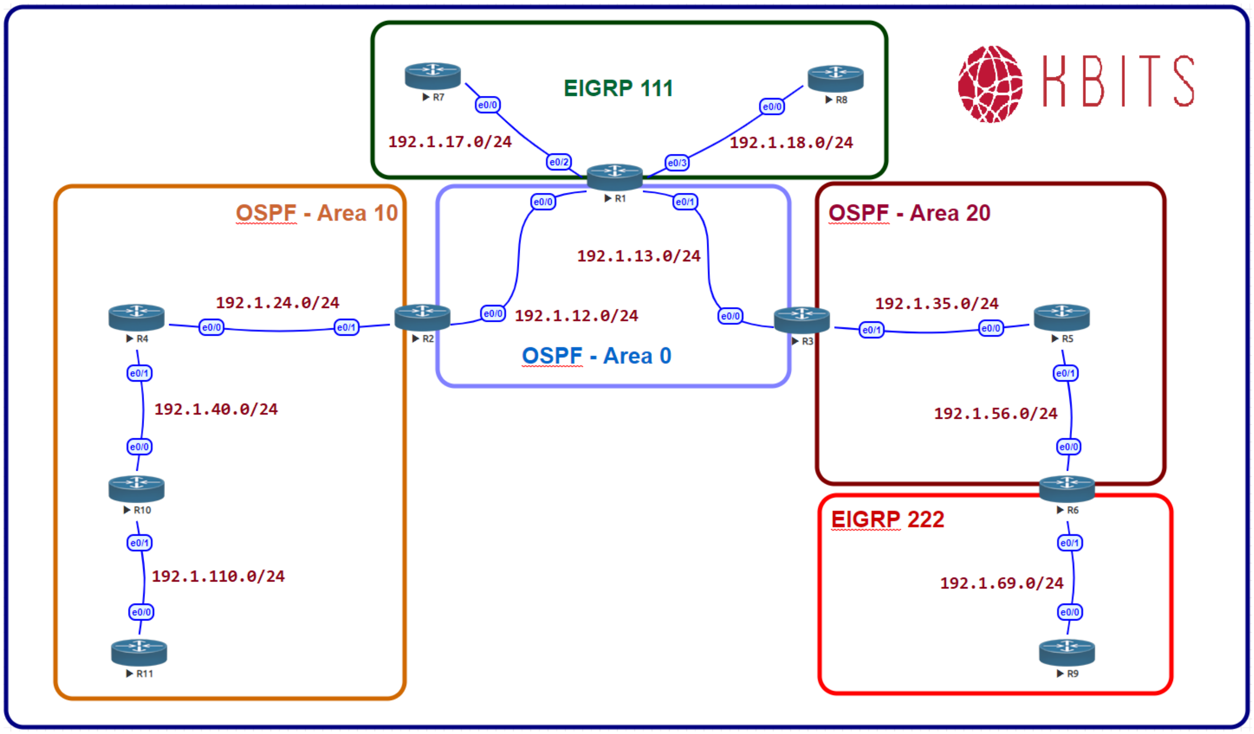

Task 1

Configure Route Summarization on the appropriate ASBR to summarize all the routes from the 107.0.0.0/8 major network towards OSPF. Use the longest mask for Route Summarization.

```
R1

Router ospf 1
 Summary-address 107.7.72.0 255.255.252.0
```

Lab 9 – Route Summarization and LSA Filtering

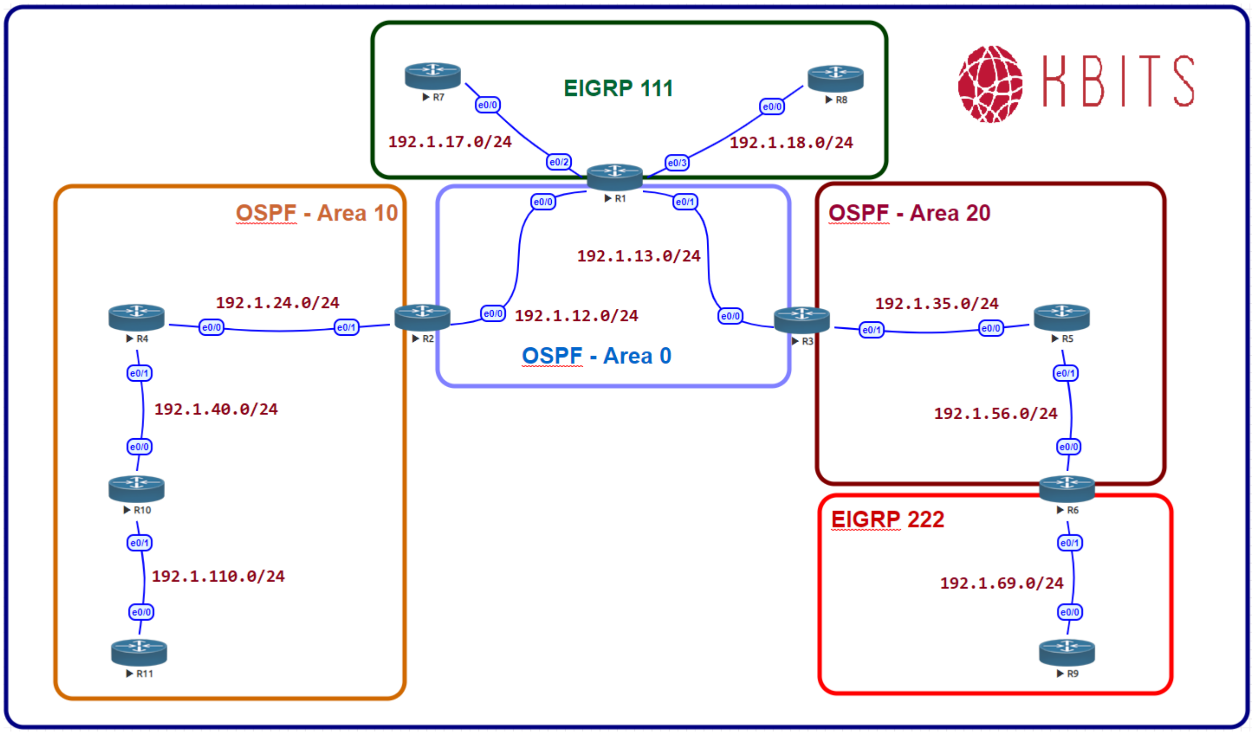

Task 1

Configure LSA Filtering such that network 4.0.0.0/8 is not allowed to leave Area 10.

R2

```
Ip prefix-list FILTER1 deny 4.0.0.0/8
Ip prefix-list FILTER1 permit 0.0.0.0/0 le 32
!
Router ospf 1
 Area 10 filter-list prefix FILTER1 out
```

Lab 10 – Configuring OSPF Authentication

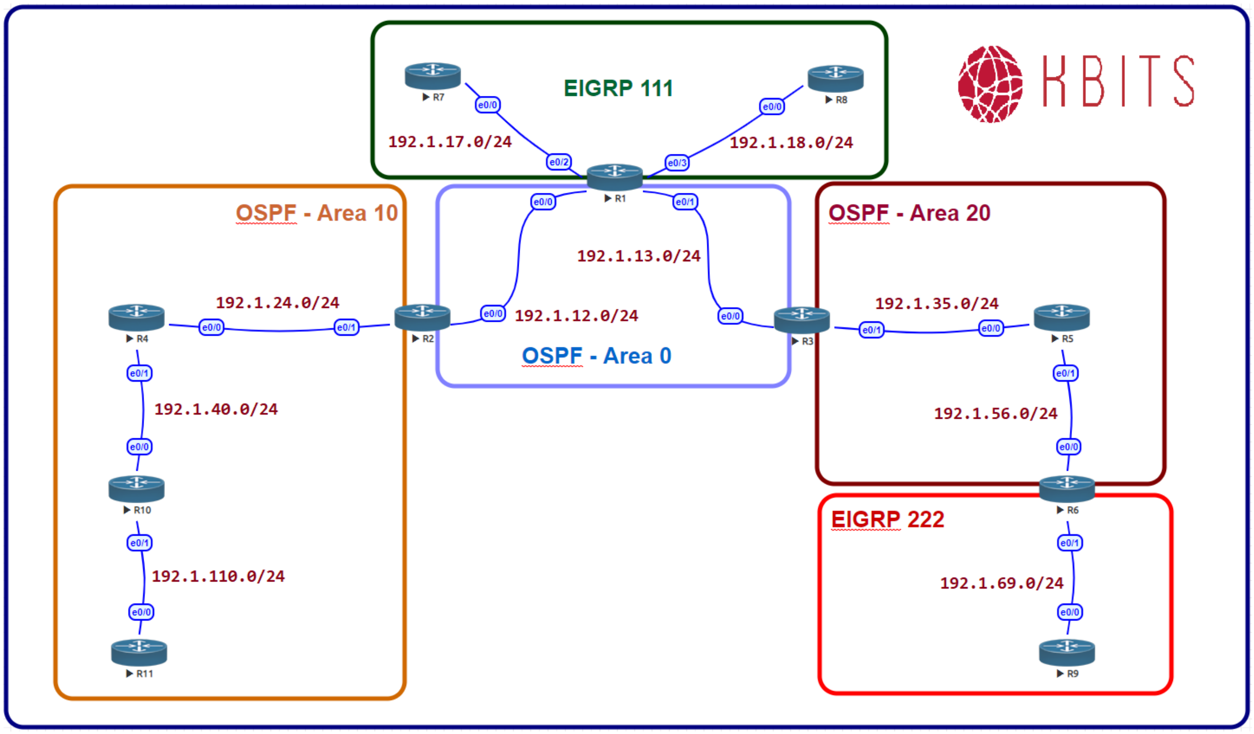

Task 1

Configure the most secure authentication on all routers in Area's 0. Use a key of 1 and a key-string **ccie123**.

R1

```
interface E 0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 ccie123
!
interface E 0/1
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 ccie123
```

R2

```
interface E 0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 ccie123
```

R3

```
interface E 0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 ccie123
```

Task 2

Configure text authentication on all routers in 10. Use a key-string **cisco**.

R2

```
interface E 0/1
 ip ospf authentication
 ip ospf authentication-key cisco
```

R4

```
interface E 0/0
 ip ospf authentication
 ip ospf authentication-key cisco
!
interface E 0/1
 ip ospf authentication
 ip ospf authentication-key cisco
```

R10

```
interface E 0/0
 ip ospf authentication
 ip ospf authentication-key cisco
!
interface E 0/1
 ip ospf authentication
 ip ospf authentication-key cisco
```

R11

```
interface E 0/0
 ip ospf authentication
 ip ospf authentication-key cisco
```

Lab 11 – Configuring OSPF Area Types

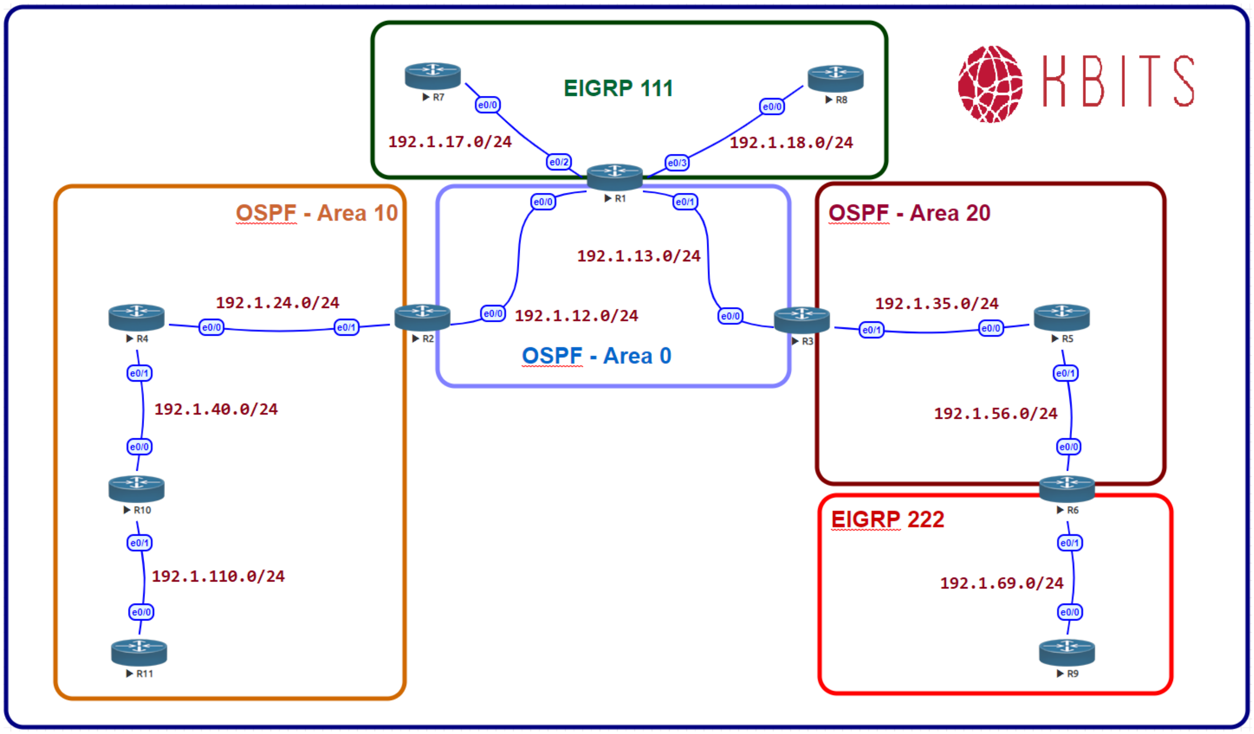

Task 1

Configure Area 10 such that it does not receive any External Routes. It should maintain connectivity to the External Routes. **(Stub Area)**

R2 Router ospf 1 Area 10 stub	**R4** Router ospf 1 Area 10 stub
R10 Router ospf 1 Area 10 stub	**R11** Router ospf 1 Area 10 stub
Note: The ABR will block the External Routes from EIGRP 111 & EIGRP 222 from reaching Area 10 Internal Routers. R2 will inject a default route instead. This is a Stub Area. Verify it on R4, R10 & R11 by checking the Routing table.	

Task 2

This step is a continuation of Task 1. Area 10 should also block Inter-Area routes maintaining reachability to them. **(Totally Stubby Area)**

R2

Router ospf 1 Area 10 stub no-summary
Note: The ABR will block the Inter-Area Routes from getting propagating into Area 10. Instead R2 will inject a default route instead. This is a Totally Stubby Area. Verify it on R4, R10 & R11 by checking the Routing table.

Task 3

Configure Area 20 such that it does not receive any external routes from the backbone. The External routes from EIGRP 222 should continue to be received in Area 20 and propagated into the Backbone. **(NSSA Area)**

R3
Router ospf 1 Area 20 nssa
R5 Router ospf 1 Area 20 nssa
R6 Router ospf 1 Area 20 nssa
Note: The ABR will block the External routes from the Backbone (EIGRP). Area 20 will continue to receive the external routes from EIGRP 222 as N routes. These routes will continue to be propagated towards the backbone. The ABR will convert the N routes into E routes as it propagates it into the Backbone. You will receive reachability to the External Routes from the Backbone as the ABR does not inject a default route in this configuration.

Task 4

This step is a continuation of Task 3. Configure Area 20 such that the previous requirement is maintained but Area 20 should also have reachability to the external routes from the backbone (EIGRP Routes). **(NSSA-Stub Area)**

R3
Router ospf 1 Area 20 stub default-information-originate
Note: This builds on the NSSA area by regaining reachability to the Backbone external routes. This is done by having the ABR injecting

the default route into Area 20.

Task 5

Configure Area 20 such that the Inter-Area routes are also blocked in addition to the external routes from the backbone. **(NSSA-Totally Stubby Area)**

R3

```
Router ospf 1
 Area 20 nssa no-summary
```

Lab 12 – Configuring Virtual Link

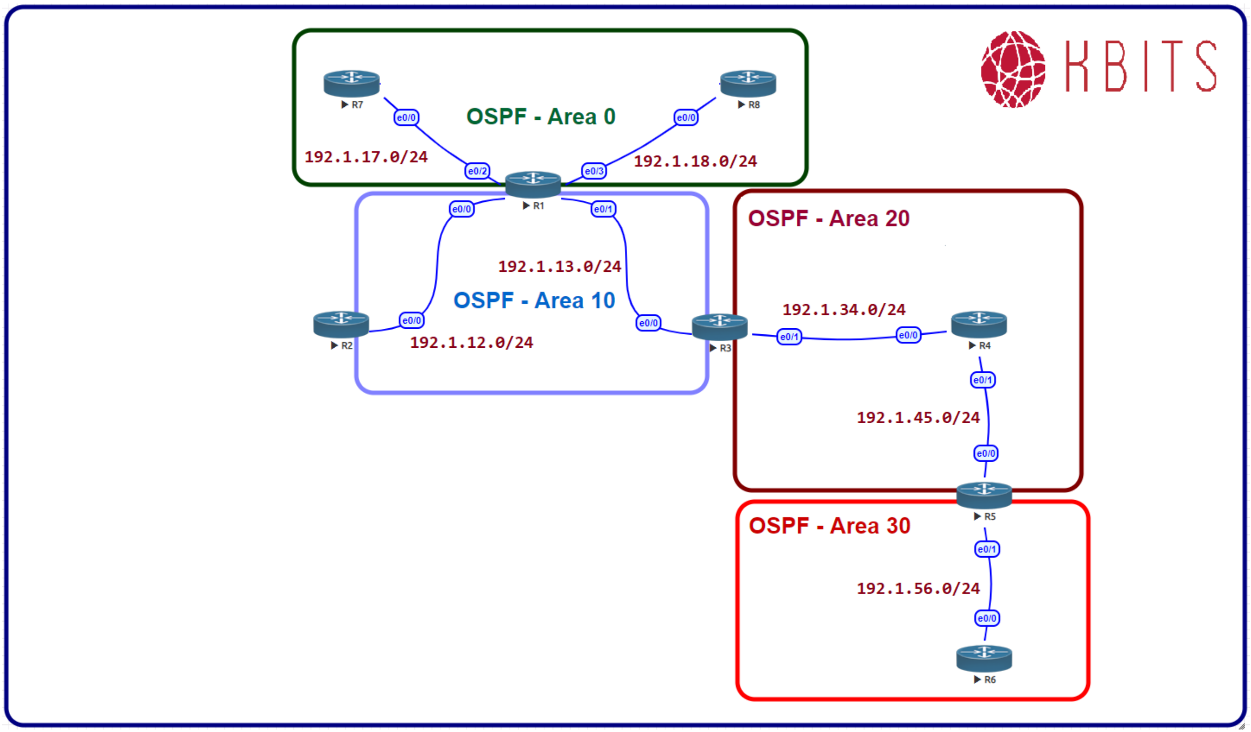

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 1	11.11.11.11	255.0.0.0
E 0/0	192.1.12.1	255.255.255.0
E 0/1	192.1.13.1	255.255.255.0
E 0/2	192.1.17.1	255.255.255.0
E 0/3	192.1.18.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
E 0/0	192.1.12.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
Loopback 1	33.33.33.33	255.0.0.0
E 0/0	192.1.13.3	255.255.255.0
E 0/1	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
E 0/0	192.1.34.4	255.255.255.0
E 0/1	192.1.45.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
Loopback 1	55.55.55.55	255.0.0.0
E 0/0	192.1.45.5	255.255.255.0
E 0/1	192.1.56.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
E 0/0	192.1.56.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
E 0/0	192.1.17.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
E 0/0	192.1.18.8	255.255.255.0

Task 1

Configure OSPF in Area 0 between R1, R7 & R8. Besides the physical links, enable the Loopback 0 interfaces of all 3 routers in Area 0. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R1 – 0.0.0.1
R7 – 0.0.0.7
R8 – 0.0.0.8

<table>
<tr>
<td>

R1

```
Router OSPF 1
 Router-id 0.0.0.1
 Network 1.0.0.0 0.255.255.255 area 0
 Network 192.1.12.0 0.0.0.255 area 0
 Network 192.1.13.0 0.0.0.255 area 0
!
Interface Loopback0
 Ip ospf network point-to-point
```

</td>
<td>

R7

```
Router OSPF 1
 Router-id 0.0.0.7
 Network 7.0.0.0 0.255.255.255 area 0
 Network 192.1.17.0 0.0.0.255 area 0
!
Interface Loopback0
 Ip ospf network point-to-point
```

</td>
</tr>
<tr>
<td>

R8

```
Router OSPF 1
 Router-id 0.0.0.8
 Network 8.0.0.0 0.255.255.255 area 0
 Network 192.1.18.0 0.0.0.255 area 0
!
Interface Loopback0
 Ip ospf network point-to-point
```

</td>
<td></td>
</tr>
</table>

Task 2

Configure OSPF in Area 10 between R1, R2 & R3. Besides the physical links, enable the Loopback 0 interfaces of R2 & R3 in Area 10. R1 Loopback1 should also be enabled in area 10. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R2 – 0.0.0.2
R3 – 0.0.0.3

<table>
<tr>
<td>

R1

```
Router OSPF 1
 Network 192.1.12.0 0.0.0.255 area 10
 Network 192.1.13.0 0.0.0.255 area 10
 Network 11.0.0.0 0.255.255.255 area 10
```

</td>
<td>

R2

```
Router OSPF 1
 Router-id 0.0.0.2
 Network 2.0.0.0 0.255.255.255 area 10
 Network 192.1.12.0 0.0.0.255 area 10
```

</td>
</tr>
</table>

<table>
<tr><td>!
Interface Loopback1
 Ip ospf network point-to-point</td><td>!
Interface Loopback0
 Ip ospf network point-to-point</td></tr>
<tr><td>R3

Router OSPF 1
 Router-id 0.0.0.3
 Network 3.0.0.0 0.255.255.255 area 10
 Network 192.1.13.0 0.0.0.255 area 10
!
Interface Loopback0
 Ip ospf network point-to-point</td><td></td></tr>
</table>

Task 3

Configure OSPF in Area 20 between R3, R4 & R5. Besides the physical links, enable the Loopback 0 interfaces of R4 & R5 in Area 20. R3 Loopback1 should also be enabled in area 20. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R4 – 0.0.0.4
R5 – 0.0.0.5

<table>
<tr><td>R3

Router OSPF 1
 Network 192.1.34.0 0.0.0.255 area 20
 Network 33.0.0.0 0.255.255.255 area 20
!
Interface Loopback1
 Ip ospf network point-to-point</td><td>R4

Router OSPF 1
 Router-id 0.0.0.4
 Network 4.0.0.0 0.255.255.255 area 20
 Network 192.1.34.0 0.0.0.255 area 20
 Network 192.1.45.0 0.0.0.255 area 20
!
Interface Loopback0
 Ip ospf network point-to-point</td></tr>
<tr><td>R5

Router OSPF 1
 Router-id 0.0.0.5
 Network 5.0.0.0 0.255.255.255 area 20
 Network 192.1.45.0 0.0.0.255 area 20
!
Interface Loopback0
 Ip ospf network point-to-point</td><td></td></tr>
</table>

Task 4

Configure a Virtual Link between the appropriate devices to allow Area 20 to communicate to the rest of the network.

```
R1

router ospf 1
  area 10 virtual-link 0.0.0.3
```

```
R3

router ospf 1
 area 10 virtual-link 0.0.0.1
```

Task 5

Configure OSPF in Area 30 between R5 & R6. Besides the physical links, enable the Loopback 0 interfaces of R6 in Area 30. R5 Loopback1 should also be enabled in area 30. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R6 – 0.0.0.6

R5	R6
Router OSPF 1 Network 192.1.56.0 0.0.0.255 area 30 Network 55.0.0.0 0.255.255.255 area 30 ! Interface Loopback1 Ip ospf network point-to-point	Router OSPF 1 Router-id 0.0.0.6 Network 6.0.0.0 0.255.255.255 area 30 Network 192.1.56.0 0.0.0.255 area 30 ! Interface Loopback0 Ip ospf network point-to-point

Task 6

Configure a Virtual Link between the appropriate devices to allow Area 30 to communicate to the rest of the network.

```
R3

router ospf 1
  area 20 virtual-link 0.0.0.5
```

```
R5

router ospf 1
 area 20 virtual-link 0.0.0.3
```

Lab 13 – Configuring BFD for OSPF

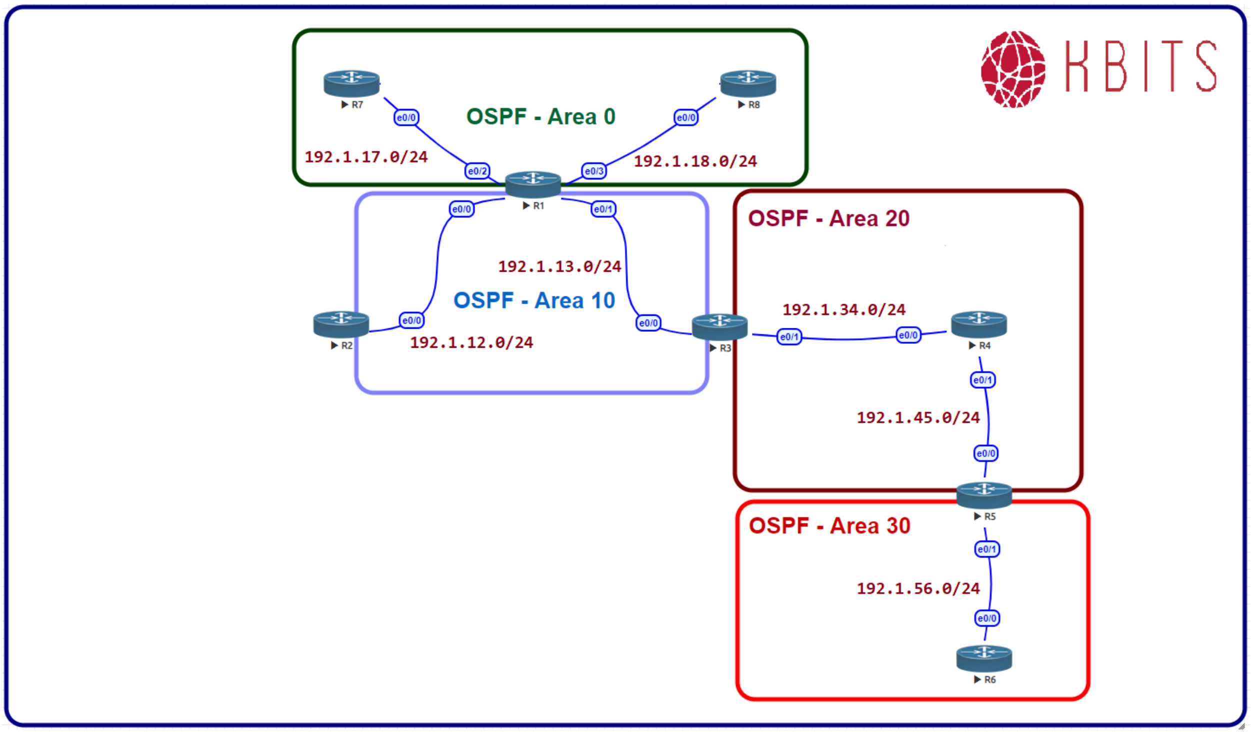

Task 1

Configure BFD between all routers in area 0. Configure the BFD Interface interval to be 300 for sending and receiving. A neighbor should be deemed dead is the router misses 3 hellos.

R1 Interface E 0/2 bfd interval 300 min_rx 300 multiplier 3 ! Router ospf 1 bfd all-interfaces	**R7** Interface E 0/0 bfd interval 300 min_rx 300 multiplier 3 ! Router ospf 1 bfd all-interfaces
R8 Interface E 0/0 bfd interval 300 min_rx 300 multiplier 3 ! Router ospf 1 bfd all-interfaces	

Lab 14 – Configuring IP FRR - OSPF

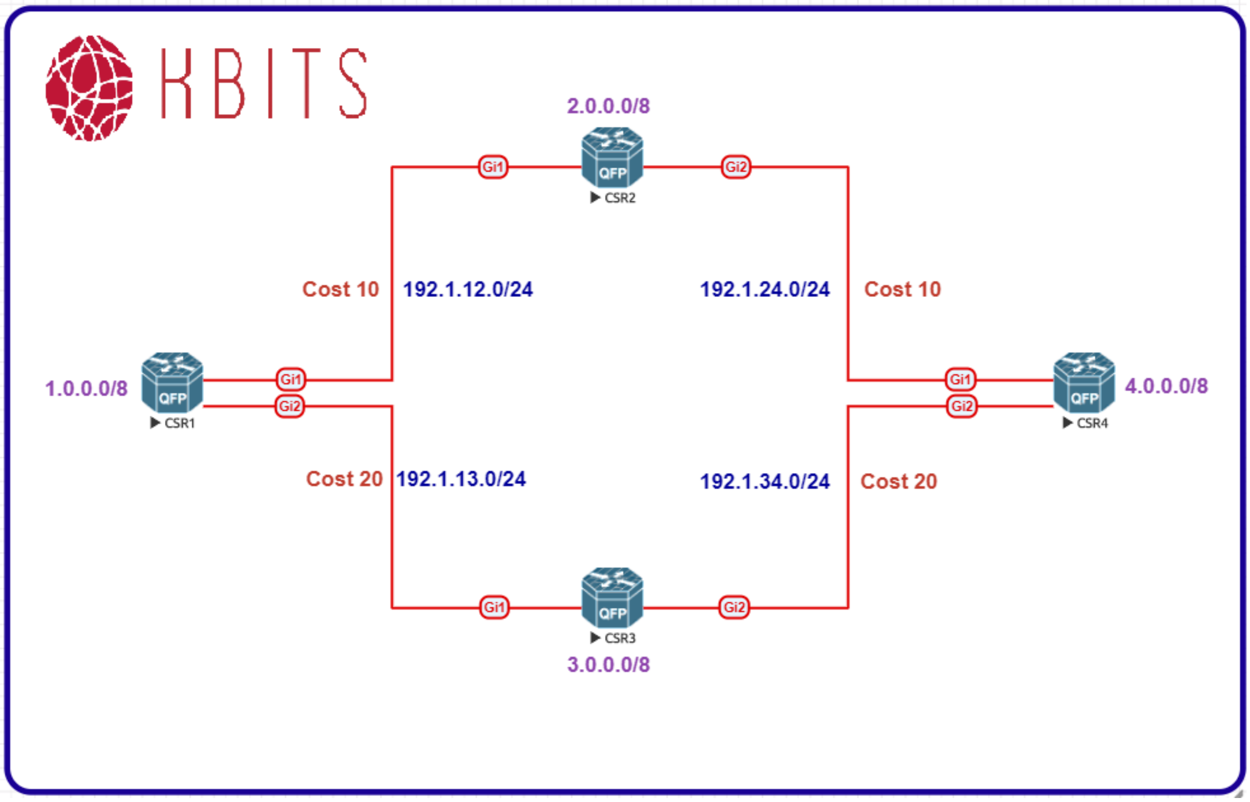

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
G 1	192.1.12.1	255.255.255.0
G 2	192.1.13.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
G 1	192.1.12.2	255.255.255.0
G 2	192.1.24.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
G 1	192.1.13.3	255.255.255.0
G 2	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
G 1	192.1.24.4	255.255.255.0
G 2	192.1.34.4	255.255.255.0

Task 1

Configure OSPF in Area 0 between R1, R2, R3 & R4. Besides the physical links, enable the Loopback 0 interfaces of all 4 routers in Area 0. Loopbacks should be advertised with the Interface Mask. Hard Code the Router-id based on the following:

R1 – 0.0.0.1
R2 – 0.0.0.2
R3 – 0.0.0.3
R4 – 0.0.0.4

R1	R2
Router OSPF 1 Router-id 0.0.0.1 Network 1.0.0.0 0.255.255.255 area 0 Network 192.1.12.0 0.0.0.255 area 0 Network 192.1.13.0 0.0.0.255 area 0 ! Interface Loopback0 Ip ospf network point-to-point	Router OSPF 1 Router-id 0.0.0.2 Network 2.0.0.0 0.255.255.255 area 0 Network 192.1.12.0 0.0.0.255 area 0 Network 192.1.24.0 0.0.0.255 area 0 ! Interface Loopback0 Ip ospf network point-to-point
R3 Router OSPF 1 Router-id 0.0.0.3 Network 3.0.0.0 0.255.255.255 area 0 Network 192.1.13.0 0.0.0.255 area 0 Network 192.1.34.0 0.0.0.255 area 0 ! Interface Loopback0	**R4** Router OSPF 1 Router-id 0.0.0.4 Network 4.0.0.0 0.255.255.255 area 0 Network 192.1.24.0 0.0.0.255 area 0 Network 192.1.34.0 0.0.0.255 area 0 ! Interface Loopback0

Ip ospf network point-to-point	Ip ospf network point-to-point

Task 2

Configure the link cost based on the Diagram.

R1 Interface Gig1 Ip ospf cost 10 ! Interface Gig2 Ip ospf cost 20	**R2** Interface Gig1 Ip ospf cost 10 ! Interface Gig2 Ip ospf cost 10
R3 Interface Gig1 Ip ospf cost 20 ! Interface Gig2 Ip ospf cost 20	**R4** Interface Gig1 Ip ospf cost 10 ! Interface Gig2 Ip ospf cost 20

Task 3

Verify the routing table and CEF on R1 for Network 4.0.0.0/8. It should have a single path via R2 (Lower cost)

R1 Show IP route 4.0.0.0 **Note:** It should have a single path via 192.1.12.2 Show ip cef 4.0.0.0 **Note:** It should have a single path via 192.1.12.2

Task 4

Enable Fast-reroute on all routers in area 0. Configure the Priority as low that creates the backup route for all networks in the OSPF Database.

R1 Router ospf 1 fast-reroute per-prefix enable area 0 prefix-priority low
R2 Router ospf 1 fast-reroute per-prefix enable area 0 prefix-priority low
R3 Router ospf 1 fast-reroute per-prefix enable area 0 prefix-priority low
R4 Router ospf 1 fast-reroute per-prefix enable area 0 prefix-priority low

Task 5

Verify the routing table and CEF on R1 for Network 4.0.0.0/8. It should have a repair path via R3 (higher cost) installed and ready in case the lower cost route goes down.

R1 Show IP route 4.0.0.0 **Note:** It should have a repair path via 192.1.13.3 Show ip cef 4.0.0.0 **Note:** It should have a repair path via 192.1.13.3

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring IS-IS

Lab 1 – Basic IS-IS Configuration with Areas

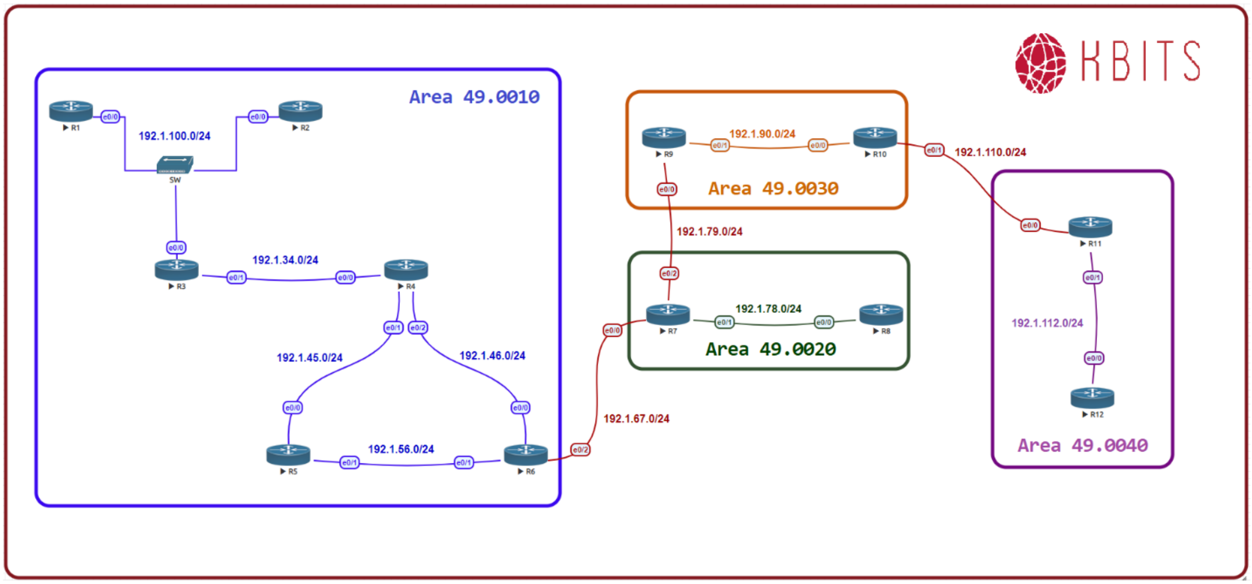

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 1	201.1.4.1	255.255.255.0
Loopback 2	201.1.5.1	255.255.255.0
Loopback 3	201.1.6.1	255.255.255.0
Loopback 4	201.1.7.1	255.255.255.0
E 0/0	192.1.100.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
E 0/0	192.1.100.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
E 0/0	192.1.100.3	255.255.255.0

E 0/1	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
E 0/0	192.1.34.4	255.255.255.0
E 0/1	192.1.45.4	255.255.255.0
E 0/2	192.1.46.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
E 0/0	192.1.45.5	255.255.255.0
E 0/1	192.1.56.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
E 0/0	192.1.46.6	255.255.255.0
E 0/1	192.1.56.6	255.255.255.0
E 0/2	192.1.67.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
E 0/0	192.1.67.7	255.255.255.0
E 0/1	192.1.78.7	255.255.255.0
E 0/2	192.1.79.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
E 0/0	192.1.78.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	9.9.9.9	255.0.0.0

E 0/0	192.1.79.9	255.255.255.0
E 0/1	192.1.90.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	10.10.10.10	255.0.0.0
E 0/0	192.1.90.10	255.255.255.0
E 0/1	192.1.110.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	11.11.11.11	255.0.0.0
E 0/0	192.1.110.11	255.255.255.0
E 0/1	192.1.112.11	255.255.255.0

R12

Interface	IP Address	Subnet Mask
Loopback 0	12.12.12.12	255.0.0.0
E 0/0	192.1.112.12	255.255.255.0

Task 1

Configure IS-IS on all 6 routers in Area **49.0010**. Use XXXX.XXX.XXXX as the System ID. Advertise all the Loopbacks in IS-IS. Make sure that the Routers only establish L1 Adjacencies with each other. Also, make sure that R6 is capable of Intra-area as well as Inter-area adjencies.

R1	R2
Router isis Net 49.0010.1111.1111.1111.00 Is-type level-1 ! Int lo0 Ip router isis Int lo1 Ip router isis Int lo2 Ip router isis Int lo3 Ip router isis	Router isis Net 49.0010.2222.2222.2222.00 Is-type level-1 ! Int lo0 Ip router isis Int E 0/0 Ip router isis

Int lo4 Ip router isis Int E 0/0 Ip router isis	
R3 Router isis Net 49.0010.3333.3333.3333.00 Is-type level-1 ! Int lo0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis	**R4** Router isis Net 49.0010.4444.4444.4444.00 Is-type level-1 ! Int lo0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis Int E 0/2 Ip router isis
R5 Router isis Net 49.0010.5555.5555.5555.00 Is-type level-1 ! Int lo0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis	**R6** Router isis Net 49.0010.6666.6666.6666.00 ! Int lo0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis

Task 2

Configure IS-IS on the 2 routers in Area **49.0020**. Use XXXX.XXX.XXXX as the System ID. Advertise all the Loopbacks in IS-IS. Make sure that the Routers only establish L1 Adjacencies with each other. Also, make sure that R7 is capable of Intra-area as well as Inter-area adjencies.

R7 Router isis Net 49.0020.7777.7777.7777.00 ! Int lo0 Ip router isis	**R8** Router isis Net 49.0020.8888.8888.8888.00 Is-type level-1 ! Int lo0

Int E 0/1 Ip router isis	Ip router isis Int E 0/0 Ip router isis

Task 3

Configure IS-IS on the 2 routers in Area **49.0030**. Advertise all the Loopbacks in IS-IS. Make sure that the Routers only establish L1 Adjacencies with each other. Also, make sure that R11 is capable of Intra-area as well as Inter-area adjencies. Configure the System ID's based on the following:

R11 – 0011.0011.0011
R12 – 0012.0012.0012

R11	**R12**
Router isis Net 49.0040.0011.0011.0011.00 ! Int lo0 Ip router isis Int E 0/1 Ip router isis	Router isis Net 49.0040.0012.0012.0012.00 Is-type level-1 ! Int lo0 Ip router isis Int E 0/0 Ip router isis

Task 4

Configure IS-IS on the 2 routers in Area **49.0030**. Advertise all the Loopbacks in IS-IS. Make sure that the Routers only establish L2 Adjacencies with each other. Configure the System ID's based on the following:

R9 – 9999.9999.9999
R10 – 1010.1010.1010

R9	**R10**
Router isis Net 49.0030.9999.9999.9999.00 Is-type level-2 ! Int lo0 Ip router isis Int E 0/1 Ip router isis	Router isis Net 49.0030.1010.1010.1010.00 Is-type level-2 ! Int lo0 Ip router isis Int E 0/0 Ip router isis

Lab 2 – Optimizing IS-IS

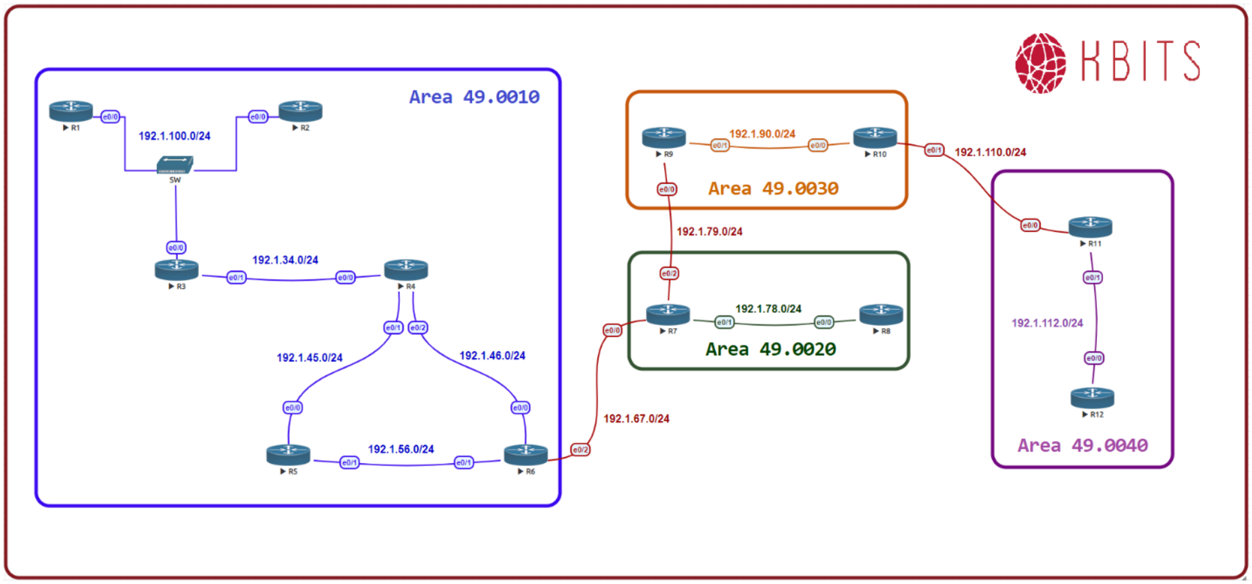

Task 1

Make sure the R1 and R3 are the DIS for their respective Multi-Access Segments.

R1	**R3**
Interface E 0/0 Isis priority 100	Interface F 0/0 Isis priority 100

Task 2

Configure the Hello between R3 and R4 to be 5 seconds with a dead timer of 15 seconds.

R3	**R4**
Interface E 0/1 Isis hello-interval 5 isis hello-multiplier 3	Interface E 0/0 Isis hello-interval 5 isis hello-multiplier 3

Task 3

Configure all the Routers such that MPLS-TE is supported on them.

R1	**R2**
Router isis	Router isis

Metric-style wide	Metric-style wide
R3 Router isis Metric-style wide	**R4** Router isis Metric-style wide
R5 Router isis Metric-style wide	**R6** Router isis Metric-style wide
R7 Router isis Metric-style wide	**R8** Router isis Metric-style wide
R9 Router isis Metric-style wide	**R10** Router isis Metric-style wide
R11 Router isis Metric-style wide	**R12** Router isis Metric-style wide

Task 4

Configure the link between R5 & R6 to be a low priority link. It should only be used in case R5 & R6 have lost their respective links towards R4.

R5 Interface E 0/1 Isis metric 50	**R6** Interface E 0/1 Isis metric 50

Lab 3 – IS-IS Authentication

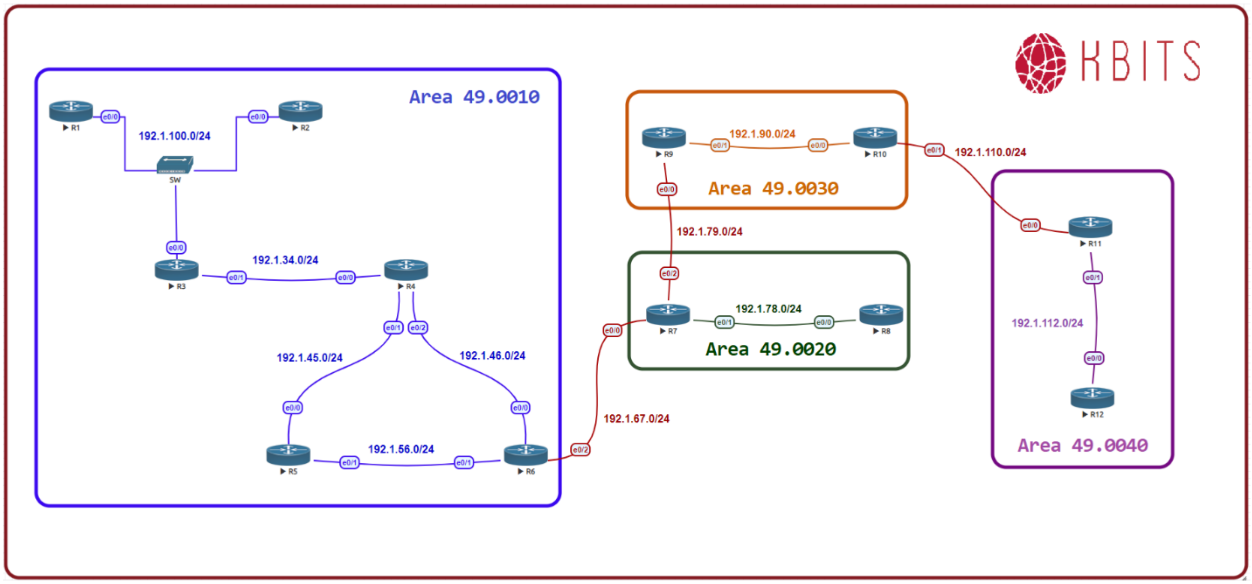

Task 1

Configure MD5 authentication for the Link between R3 & R4. Use ccie as the key-string with a key-id of 1.

R3

```
Key chain AUTH
 Key 1
  Key-string ccie
!
Interface E 0/1
 Isis authentication key-chain AUTH
 Isis authentication mode MD5
```

R4

```
Key chain AUTH
 Key 1
  Key-string ccie
!
Interface E 0/0
 Isis authentication key-chain AUTH
 Isis authentication mode MD5
```

Task 2

Configure Text authentication for the Link between R5 & R6. Use ccie as the key-string with a key-id of 1.

R5 Key chain AUTH Key 1   Key-string ccie ! Interface E 0/1 Isis authentication key-chain AUTH Isis authentication mode text
R6 Key chain AUTH Key 1   Key-string ccie ! Interface E 0/1 Isis authentication key-chain AUTH Isis authentication mode text

Lab 4 – Configure Inter-area Interfaces

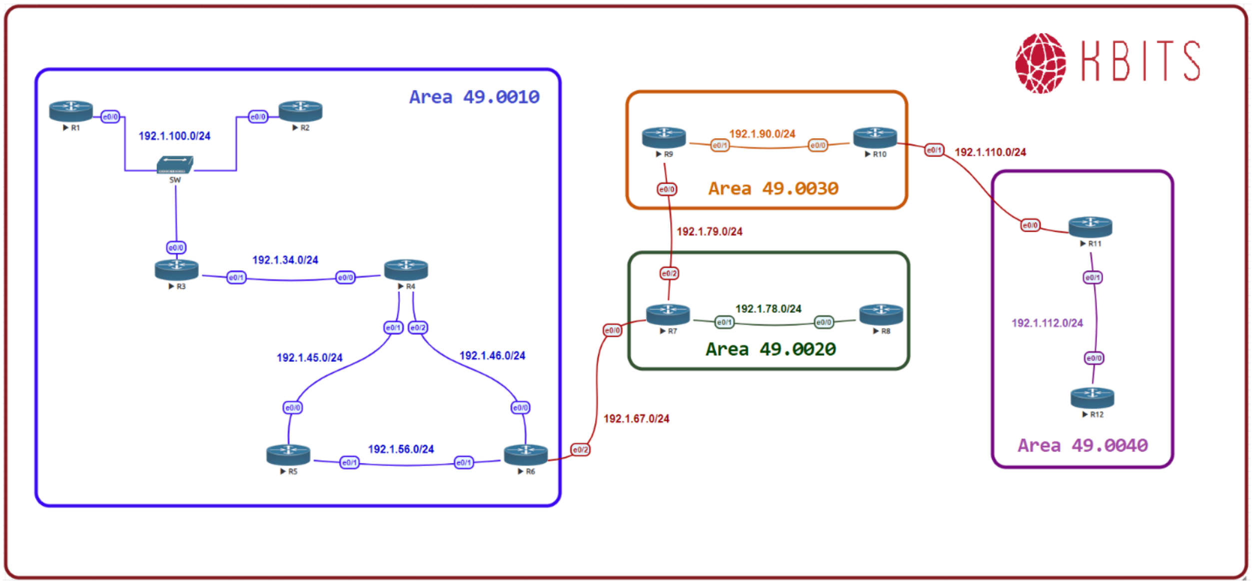

Task 1

Configure IS-IS between R6 & R7 to connect Area 49.0010 to Area 49.0020.

R6 Interface E 0/2 Ip router isis
R7 Interface E 0/0 Ip router isis

Task 2

Configure IS-IS between R7 & R9 to connect Area 49.0020 to Area 49.0030.

R7 Interface E 0/2 Ip router isis
R9 Interface E 0/0 Ip router isis

Task 3

Configure IS-IS between R10 & R11 to connect Area 49.0030 to Area 49.0040.

R10 Interface E 0/1 Ip router isis
R11 Interface E 0/0 Ip router isis

Task 4

Verify connectivity from R1 to R12 using Ping. What type of routes do you have in the routing tables? Do you have full connectivity?

Lab 5 –IS-IS Multi-Area / Multi-Domain Configuration

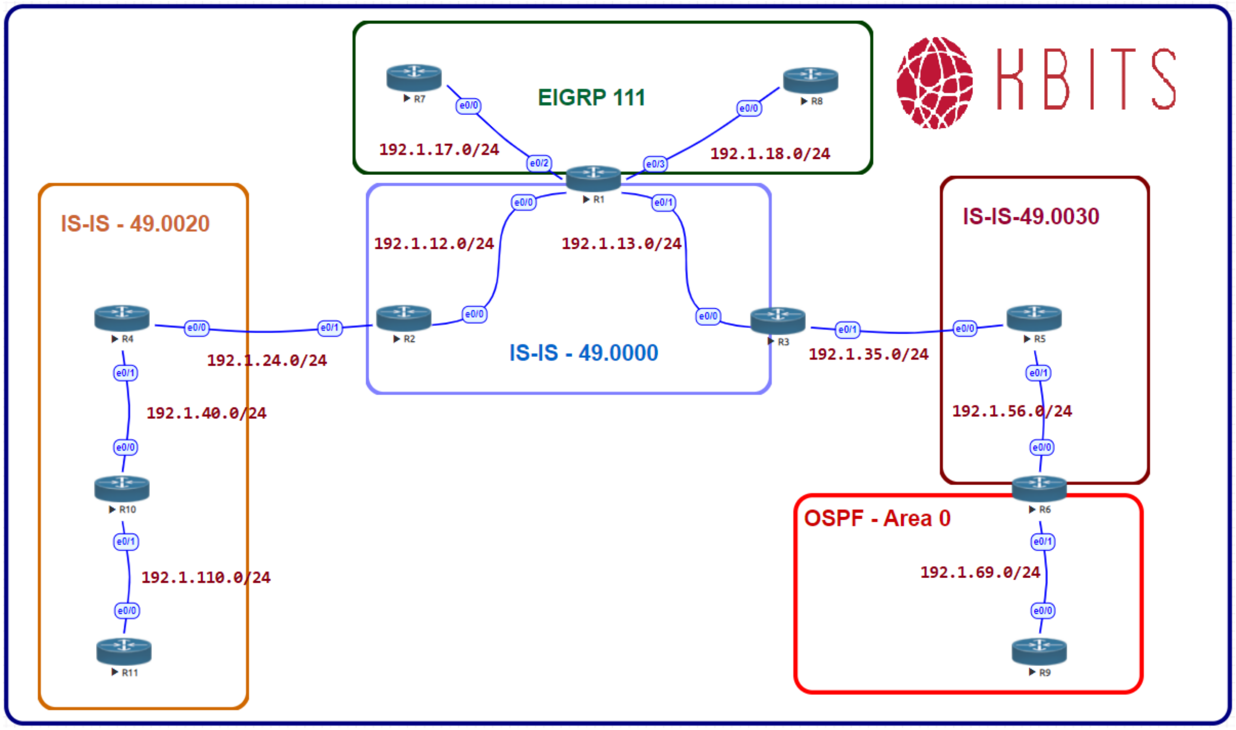

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 1	11.11.11.11	255.0.0.0
E 0/0	192.1.12.1	255.255.255.0
E 0/1	192.1.13.1	255.255.255.0
E 0/2	192.1.17.1	255.255.255.0
E 0/3	192.1.18.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
E 0/0	192.1.12.2	255.255.255.0

E 0/1	192.1.24.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
E 0/0	192.1.13.3	255.255.255.0
E 0/1	192.1.35.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
E 0/0	192.1.24.4	255.255.255.0
E 0/1	192.1.40.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
E 0/0	192.1.35.5	255.255.255.0
E 0/1	192.1.56.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
Loopback 1	66.66.66.66	255.0.0.0
E 0/0	192.1.56.6	255.255.255.0
E 0/1	192.1.69.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
Loopback 1	107.7.72.1	255.255.255.0
Loopback 2	107.7.73.1	255.255.255.0
Loopback 3	107.7.74.1	255.255.255.0
Loopback 4	107.7.75.1	255.255.255.0
E 0/0	192.1.17.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
E 0/0	192.1.18.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	9.9.9.9	255.0.0.0
E 0/0	192.1.69.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	10.10.10.10	255.0.0.0
E 0/0	192.1.40.10	255.255.255.0
E 0/1	192.1.110.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	111.111.100.11	255.255.255.0
Loopback 1	111.111.101.11	255.255.255.0
Loopback 2	111.111.102.11	255.255.255.0
Loopback 3	111.111.103.11	255.255.255.0
E 0/0	192.1.110.11	255.255.255.0

Task 1

Configure IS-IS in Area 49.0000 on R1, R2 & R3. Besides the physical links, enable IS-IS on the Loopback 0 interfaces of all 3 routers. Configure the routers as Level-2 routers. Configure the System-IDs based on the following:

R1 – 1111.1111.1111
R2 – 2222.2222.2222
R3 – 3333.3333.3333

R1	R2
Router isis Net 49.0000.1111.1111.1111.00 Is-type level-2 ! Interface loopback0	Router isis Net 49.0000.2222.2222.2222.00 Is-type level-2 ! Interface loopback0

Ip router isis ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis	Ip router isis ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis
R3 Router isis Net 49.0000.3333.3333.3333.00 Is-type level-2 ! Interface loopback0 Ip router isis ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis	

Task 2

Configure IS-IS in Area 49.0010 on R4, R10 & R11. Besides the physical links, enable IS-IS on Loopback 0 interfaces of R4 & R10. Configure all the interfaces on R11 for IS-IS. Configure R10 & R11 routers as Level-1 routers. Configure R4 to such that it can establish either a Level-1 or Level-2 neighbor relationships. Configure the System-IDs based on the following:

R4 – 4444.4444.4444
R10 – 1010.1010.1010
R11 – 0011.0011.0011

R4 Router isis Net 49.0010.4444.4444.4444.00 ! Interface loopback0 Ip router isis ! Interface E 0/0 Ip router isis ! Interface E 0/1	**R10** Router isis Net 49.0010.1010.1010.1010.00 Is-type level-1 ! Interface loopback0 Ip router isis ! Interface E 0/0 Ip router isis !

Ip router isis	Interface E 0/1 Ip router isis
R11 Router isis Net 49.0010. 0011.0011.0011.00 Is-type level-1 ! Interface loopback0 Ip router isis ! Interface loopback1 Ip router isis ! Interface loopback2 Ip router isis ! Interface loopback3 Ip router isis ! Interface E 0/0 Ip router isis	

Task 3

Configure IS-IS in Area 49.0020 on R5 & R6. Besides the physical links, enable IS-IS on Loopback 0 interfaces of R5 & R6. Configure R5 & R6 routers as Level-2 routers. Configure the System-IDs based on the following:

R5 – 5555.5555.5555
R5 – 6666.6666.6666

R5	**R6**
Router isis Net 49.0020.5555.5555.5555.00 Is-type level-2 ! Interface loopback0 Ip router isis ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis	Router isis Net 49.0020.6666.6666.6666.00 Is-type level-2 ! Interface loopback0 Ip router isis ! Interface E 0/0 Ip router isis

Task 4

Configure EIGRP is AS 111 between R1, R7 & R8. Enable all loopbacks on R7 & R8 in EIGRP 111. Enable Loopback 1 on R1 in EIGRP 111.

R1 Router EIGRP 111 Network 192.1.17.0 Network 192.1.18.0 Network 11.0.0.0	**R7** Router EIGRP 111 Network 192.1.17.0 Network 7.0.0.0 Network 107.0.0.0
R8 Router EIGRP 111 Network 192.1.18.0 Network 8.0.0.0	

Task 5

Configure OSPF in Area 0 between R6 & R9. Enable all loopbacks on R9 in OSPF. Enable Loopback 1 on R6 in OSPF.

R6 Router ospf 1 Network 192.1.69.0 0.0.0.255 area 0 Network 66.0.0.0 0.255.255.255 area 0	**R9** Router ospf 1 Network 192.1.69.0 0.0.0.255 area 0 Network 9.0.0.0 0.255.255.255 area 0

Task 6

Configure Mutual Redistribution between the appropriate routers to allow end-to-end connectivity between all routing domains. Use Seed metric of your choice.

R1 Router isis Redistribute eigrp 111 ! Router eigrp 111 Redistribute isis metric 10 10 10 10 10	**R6** Router isis Redistribute ospf 1 ! Router ospf 1 Redistribute isis subnets

Lab 6 – Configuring Route Leaking

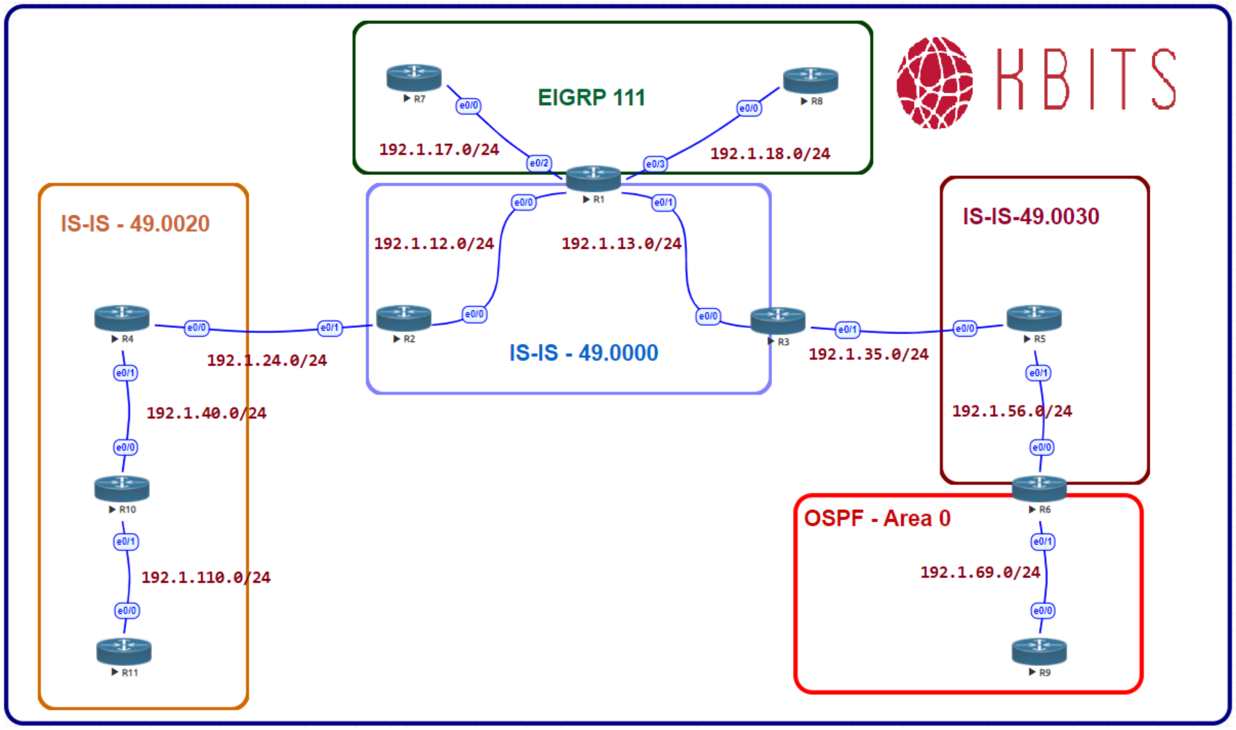

Task 1

Configure Area's 49.0010 & 49.0020 to receive all routes from all areas.

R3

```
Access-list 101 permit ip any any
!
router isis
 redistribute isis ip level-2 into level-1 distribute-list 101
```

R4

```
Access-list 101 permit ip any any
!
router isis
 redistribute isis ip level-2 into level-1 distribute-list 101
```

Lab 7 – Route Summarization

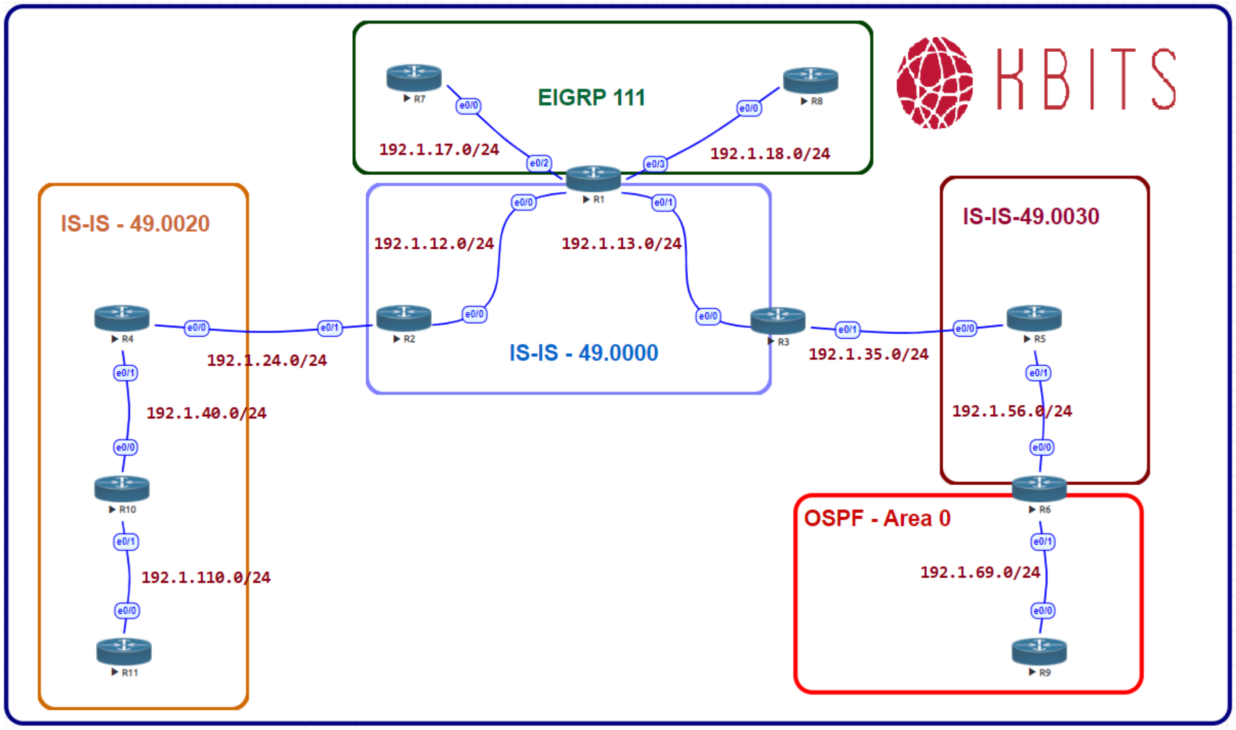

Task 1

Configure Area 49.0010 such that all the 111.0.0.0/8 routes are summarized out of the area.

```
R4

router isis
 summary-address 111.111.100.0 255.255.252.0
```

Task 2

Configure R1 such that all the 107.0.0.0/8 routes are summarized in IS-IS.

```
R1

router isis
 summary-address 107.7.72 255.255.252.0
```

Lab 8 – Configuring BFD for IS-IS

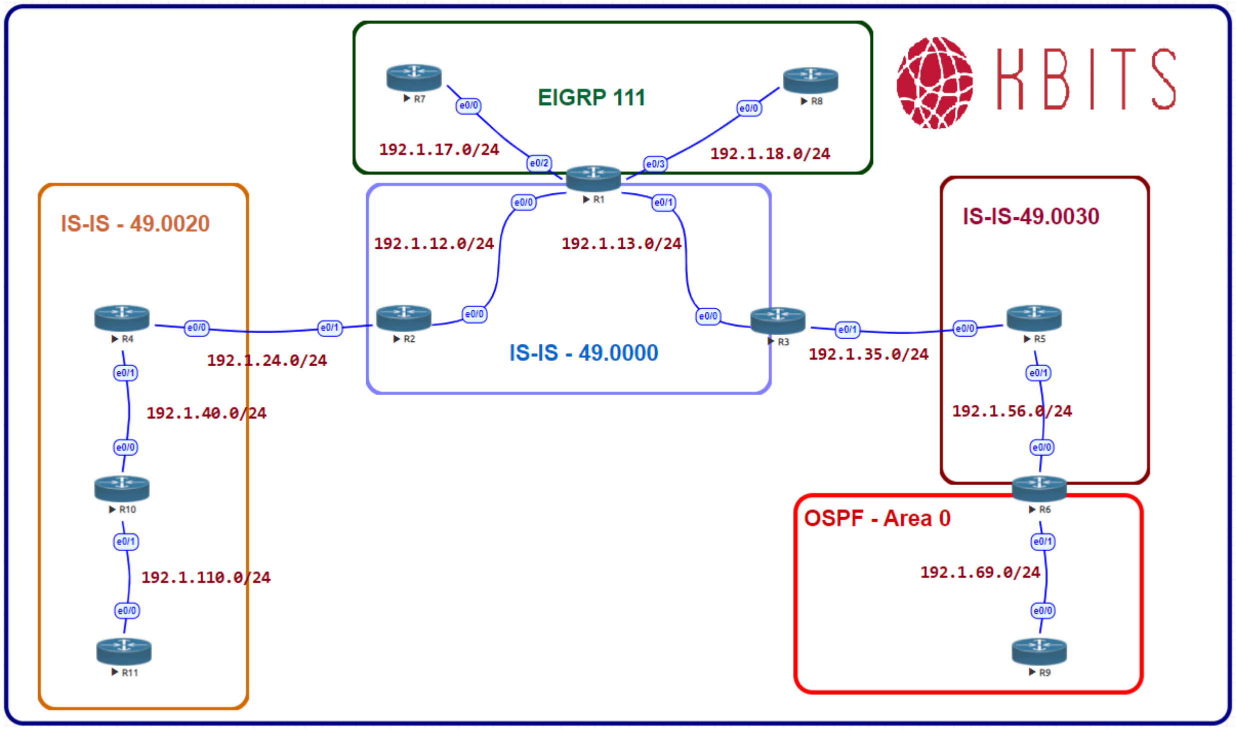

Task 1

Configure BFD between all routers in area 49.0000. Configure the BFD Interface interval to be 300 for sending and receiving. A neighbor should be deemed dead is the router misses 3 hellos.

R1	**R2**
Interface E 0/0 bfd interval 300 min_rx 300 multiplier 3 ! Interface E 0/1 bfd interval 300 min_rx 300 multiplier 3 ! Router isis bfd all-interfaces	Interface E 0/0 bfd interval 300 min_rx 300 multiplier 3 ! Router isis bfd all-interfaces
R3 Interface E 0/0 bfd interval 300 min_rx 300 multiplier 3 ! Router isis bfd all-interfaces bfd all-interfaces	

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring BGP

Lab 1 – Configuring eBGP

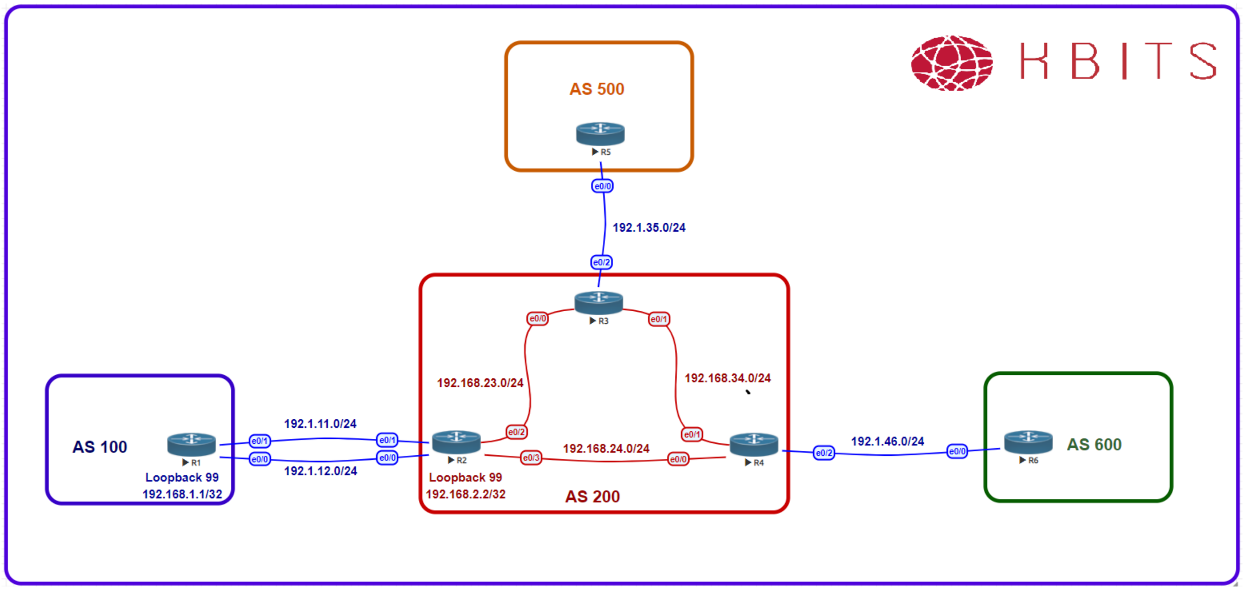

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 1	11.1.1.1	255.255.255.0
Loopback 99	192.168.1.1	255.255.255.255
E 0/0	192.1.12.1	255.255.255.0
E 0/1	192.1.11.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
Loopback 1	22.2.2.2	255.255.255.0
Loopback 10	10.2.2.2	255.255.255.255
Loopback 99	192.168.2.2	255.255.255.255
E 0/0	192.1.12.2	255.255.255.0
E 0/1	192.1.11.2	255.255.255.0
E 0/2	192.168.23.2	255.255.255.0
E 0/3	192.168.24.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
Loopback 1	33.3.3.3	255.255.255.0
Loopback 10	10.3.3.3	255.255.255.255
E 0/0	192.168.23.3	255.255.255.0
E 0/1	192.168.34.3	255.255.255.0
E 0/2	192.1.35.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
Loopback 1	44.4.4.4	255.255.255.0
Loopback 10	10.4.4.4	255.255.255.255
E 0/0	192.168.24.4	255.255.255.0
E 0/1	192.168.34.4	255.255.255.0
E 0/2	192.1.46.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
Loopback 1	55.5.5.5	255.255.255.0
E 0/0	192.1.35.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
Loopback 1	66.6.6.6	255.255.255.0
E 0/0	192.1.46.6	255.255.255.0

Task 1

Configure a BGP neighbor relationship between R3 and R5. R3 should be in AS 200 and R5 should be in AS 500. Advertise the loopback networks in BGP. Hard-code the Router ID for the BGP routers as 33.33.33.33 for R3 and 55.55.55.55 for R5.

R3	**R5**
Router BGP 200 bgp router-id 33.33.33.33 Network 3.0.0.0 Network 33.3.3.0 mask 255.255.255.0 Neighbor 192.1.35.5 remote-as 500	Router BGP 500 bgp router-id 55.55.55.55 Network 5.0.0.0 Network 55.5.5.0 mask 255.255.255.0 Neighbor 192.1.35.3 remote-as 200

Task 2

Configure a BGP neighbor relationship between R4 and R6. R4 should be in AS 200 and R6 should be in AS 600. Advertise the loopback networks in BGP. Hard-code the Router ID for the BGP routers as 44.44.44.44 for R4 and 66.66.66.66 for R6.

R4	**R6**
Router BGP 200 bgp router-id 44.44.44.44 Network 4.0.0.0 Network 44.4.4.0 mask 255.255.255.0 Neighbor 192.1.46.6 remote-as 600	Router BGP 600 bgp router-id 66.66.66.66 Network 6.0.0.0 Network 66.6.6.0 mask 255.255.255.0 Neighbor 192.1.46.4 remote-as 200

Lab 2 – Configuring eBGP Multi-Hop

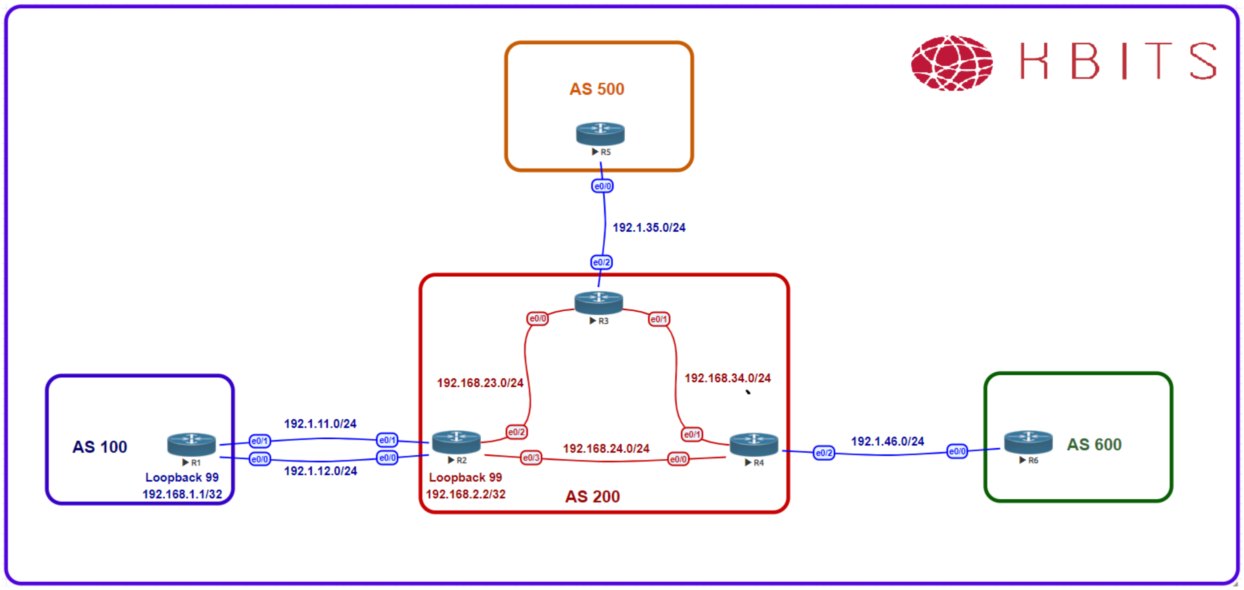

Task 1

Configure a Static route on R1 & R2 to reach each others Loopback 99 via the 2 directed connected links.

R1 Ip route 192.168.2.2 255.255.255.255 192.1.11.2 Ip route 192.168.2.2 255.255.255.255 192.1.12.2
R2 Ip route 192.168.1.1 255.255.255.255 192.1.11.1 Ip route 192.168.1.1 255.255.255.255 192.1.12.1

Task 2

Configure a BGP neighbor relationship between R1 & R2 in AS 100 & AS 200 respectively. Use Loopback99 address for the peering.

R1 Router BGP 100 Neighbor 192.168.2.2 remote-as 200 Neighbor 192.168.2.2 ebgp-multihop	**R2** Router BGP 200 Neighbor 192.168.1.1 remote-as 100 Neighbor 192.168.1.1 ebgp-multihop

Lab 3 – Redistributing Networks into BGP

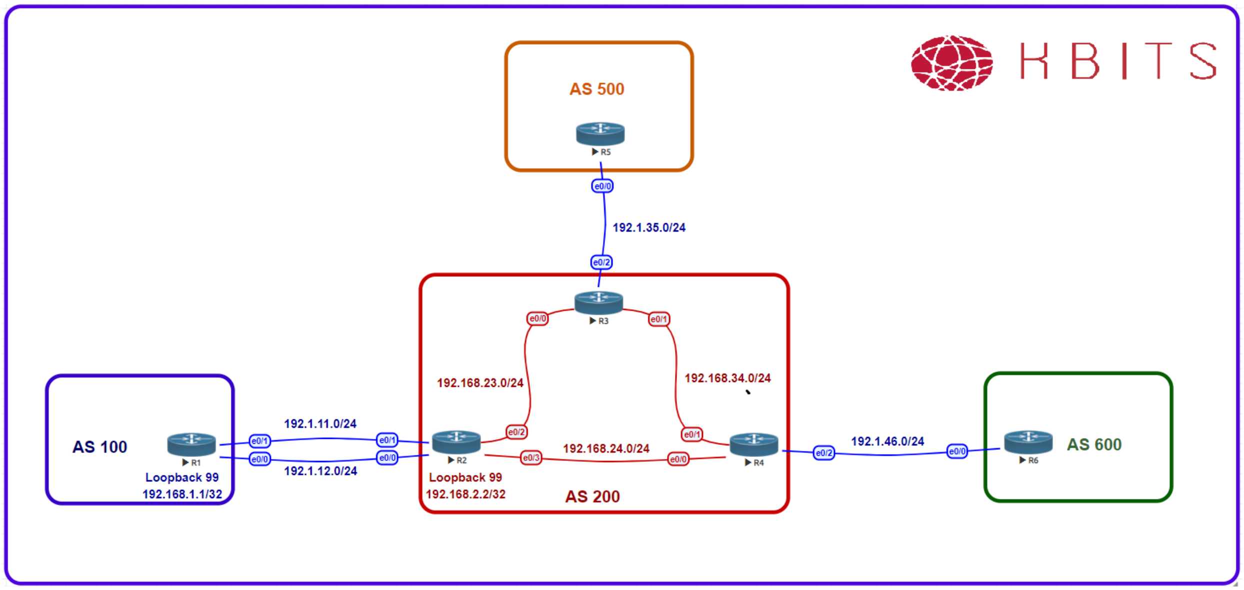

Task 1

Inject Loopback0 & Loopback1 networks on R1 into BGP. Make sure that the routes appear with an origin code of **"i"** in the BGP table.

R1

```
Ip prefix-list RC permit 1.0.0.0/8
Ip prefix-list RC permit 11.1.1.0/24
!
Route-map RC
 Match ip address prefix RC
 Set origin igp
!
Router bgp 100
 Redistribute connected route-map RC
```

Lab 4 – Configuring BGP Authentication

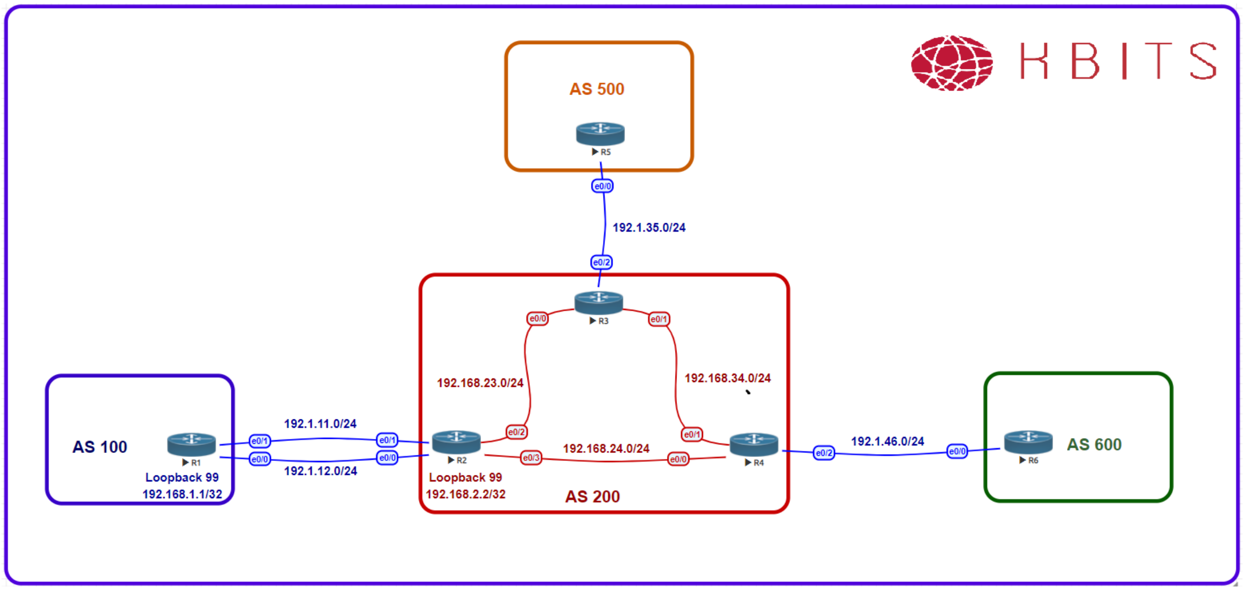

Task 1

Configure MD5 Authentication between all eBGP peers using a password of **ccie123.**

R1 Router BGP 100 Neighbor 192.168.2.2 password ccie123	**R2** Router BGP 200 Neighbor 192.168.2.2 password ccie123
R3 Router BGP 200 Neighbor 192.1.35.5 password ccie123	**R5** Router BGP 500 Neighbor 192.1.35.3 password ccie123
R4 Router BGP 200 Neighbor 192.1.46.6 password ccie123	**R6** Router BGP 600 Neighbor 192.1.46.4 password ccie123

Lab 5 – Configuring iBGP with Route Reflectors

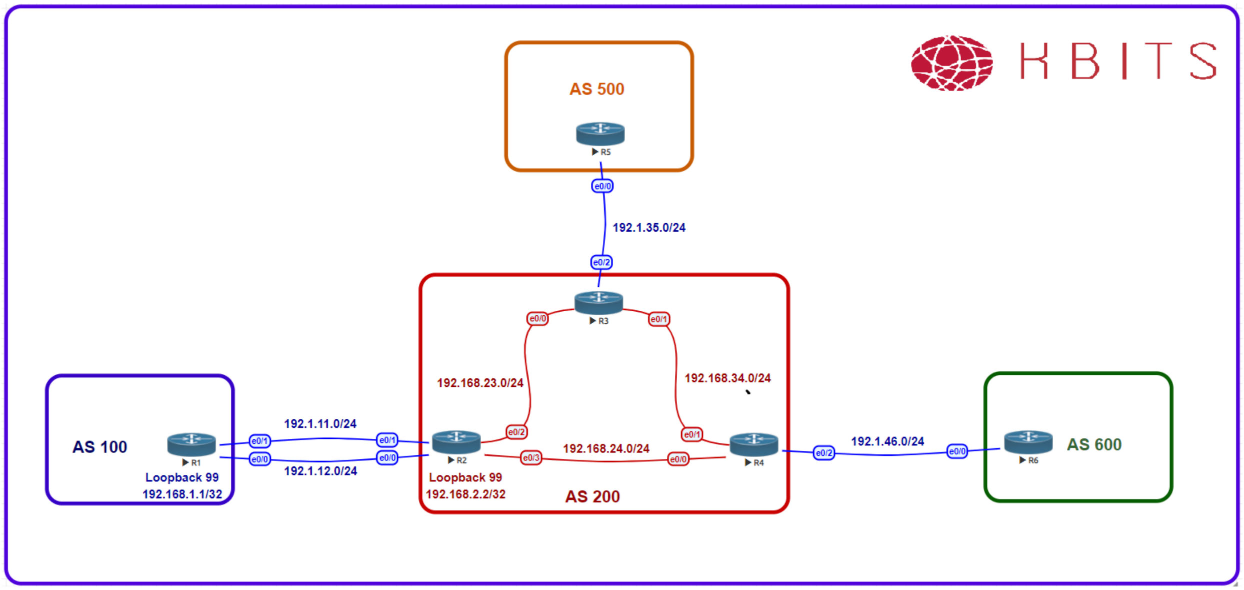

Task 1

Configure IS-IS as the IGP to route the Loopback10 networks within AS 200. Configure IS-IS with a 24-bit metric. The IS-IS neighbors should maintain a Level-2 database only. Use the following for the NET address:

R2 – 49.0000.2222.2222.2222.00
R3 – 49.0000.3333.3333.3333.00
R4 – 49.0000.4444.4444.4444.00

R2	**R3**
Router isis	Router isis
Net 49.0000.2222.2222.2222.00	Net 49.0000.3333.3333.3333.00
Is-type level-2	Is-type level-2
Metric-style wide	Metric-style wide
!	!
Interface loopback10	Interface loopback10
Ip router isis	Ip router isis
!	!
Interface E 0/2	Interface E 0/0
Ip router isis	Ip router isis
!	!
Interface E 0/3	Interface E 0/1
Ip router isis	Ip router isis

R4 Router isis Net 49.0000.4444.4444.4444.00 Is-type level-2 Metric-style wide ! Interface loopback10 Ip router isis ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis	

Task 2

Configure an iBGP neighbor relationship between R2 & R3. The neighbor relationship should be configured with redundancy in mind. Make sure that the eBGP routes are propagated and injected into the BGP table.

R2 Router BGP 200 Neighbor 10.3.3.3 remote-as 200 Neighbor 10.3.3.3 update-source loop10 Neighbor 10.3.3.3 next-hop-self	**R2** Router BGP 200 Neighbor 10.2.2.2 remote-as 200 Neighbor 10.2.2.2 update-source loop10 Neighbor 10.2.2.2 next-hop-self

Task 3

Configure an iBGP neighbor relationship between R3 & R4. The neighbor relationship should be configured with redundancy in mind. Make sure that the eBGP routes are propagated and injected into the BGP table.

R3 Router BGP 200 Neighbor 10.4.4.4 remote-as 200 Neighbor 10.4.4.4 update-source loop10 Neighbor 10.4.4.4 next-hop-self	**R4** Router BGP 200 Neighbor 10.3.3.3 remote-as 200 Neighbor 10.3.3.3 update-source loop10 Neighbor 10.3.3.3 next-hop-self

Verification:

- Make sure that AS 500 Loopbacks can reach the Loopback interfaces in AS 100, AS 200 & AS 600.

- Try the reachability between the AS 100 & AS 600 loopbacks? Are they reachable?

Task 4

Re-configure R3 such that it propagates the routes from R2 towards R4 and vice versa. Use Peer-group to accomplish this task.

R3

```
Router BGP 200
 No neighbor 10.2.2.2
 No neighbor 10.4.4.4
 Neighbor IBGP peer-group
 Neighbor IBGP remote-as 200
 Neighbor IBGP update-source Loopback10
 Neighbor IBGP next-hop-self
 Neighbor IBGP route-reflector-client
 Neighbor 10.2.2.2 peer-group IBGP
 Neighbor 10.4.4.4 peer-group IBGP
```

Verification:

- Try the reachability between the AS 100 & AS 600 loopbacks? Are they reachable?

- Trace a packet from R1 to R6 (1.1.1.1 to 6.6.6.6). What path does it take?

Task 5

You would like the RR to be an “inline RR”. This is for the purpose of future Data Filtering. Configure R3 to accomplish this.

R3

```
Router BGP 200
 Neighbor IBGP next-hop-self all
```

Verification:

- Trace a packet from R1 to R6 (1.1.1.1 to 6.6.6.6). What path does it take?

Lab 6 – Route Filtering using ACLs

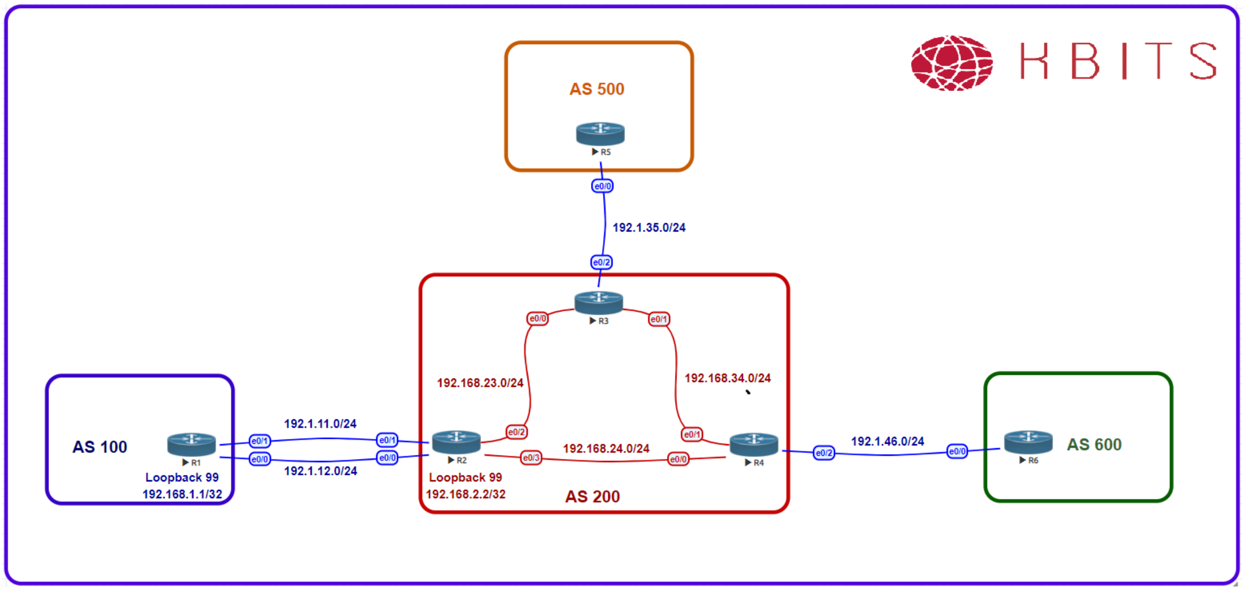

Task 1

Create the following Loopbacks on R2

Loopback 1 – 192.2.1.1/24
Loopback 2 – 192.2.2.1/24
Loopback 3 – 192.2.3.1/24
Loopback 4 – 192.2.4.1/24
Loopback 5 – 192.2.5.1/24
Loopback 6 – 192.2.6.1/24
Loopback 7 – 192.2.7.1/24
Loopback 8 – 192.2.8.1/24

R2

```
interface Loopback1
 ip address 192.2.1.1 255.255.255.0
!
interface Loopback2
 ip address 192.2.2.1 255.255.255.0
!
interface Loopback3
 ip address 192.2.3.1 255.255.255.0
!
interface Loopback4
 ip address 192.2.4.1 255.255.255.0
!
```

```
interface Loopback5
 ip address 192.2.5.1 255.255.255.0
!
interface Loopback6
 ip address 192.2.6.1 255.255.255.0
!
interface Loopback7
 ip address 192.2.7.1 255.255.255.0
!
interface Loopback8
 ip address 192.2.8.1 255.255.255.0
```

Task 2

Advertise the newly created routes in BGP. Do not use the network command to accomplish this. These routes should have an origin code of "igp".

R2

```
Access-list 1 permit 192.2.1.1 0.0.0.255
Access-list 1 permit 192.2.2.1 0.0.0.255
Access-list 1 permit 192.2.3.1 0.0.0.255
Access-list 1 permit 192.2.4.1 0.0.0.255
Access-list 1 permit 192.2.5.1 0.0.0.255
Access-list 1 permit 192.2.6.1 0.0.0.255
Access-list 1 permit 192.2.7.1 0.0.0.255
Access-list 1 permit 192.2.8.1 0.0.0.255
!
Route-map RC permit 10
 Match address 1
 Set origin igp
!
Router bgp 200
 Redistribute connected route-map RC
```

Task 3

Configure R2 such that it blocks all the 192.2.X.0 routes that have an odd number in the third octet from propagating outside the local AS. Use the distribute-list command to accomplish this task.

```
R2

Access-list 2 deny 192.2.1.0 0.0.254.255
Access-list 2 permit any
!
Router bgp 200
 Neighbor 192.168.1.1 distribute-list 2 out
```

Task 4

Configure R3 such that it blocks all the 192.2.X.0 routes that have an even number in the third octet from propagating from R2. Use the distribute-list command to accomplish this task.

```
R3

Access-list 1 deny 192.2.0.0 0.0.254.255
Access-list 1 permit any
!
Router bgp 200
 Neighbor 10.2.2.2 distribute-list 1 in
```

Lab 7 – Route Filtering using Prefix-Lists

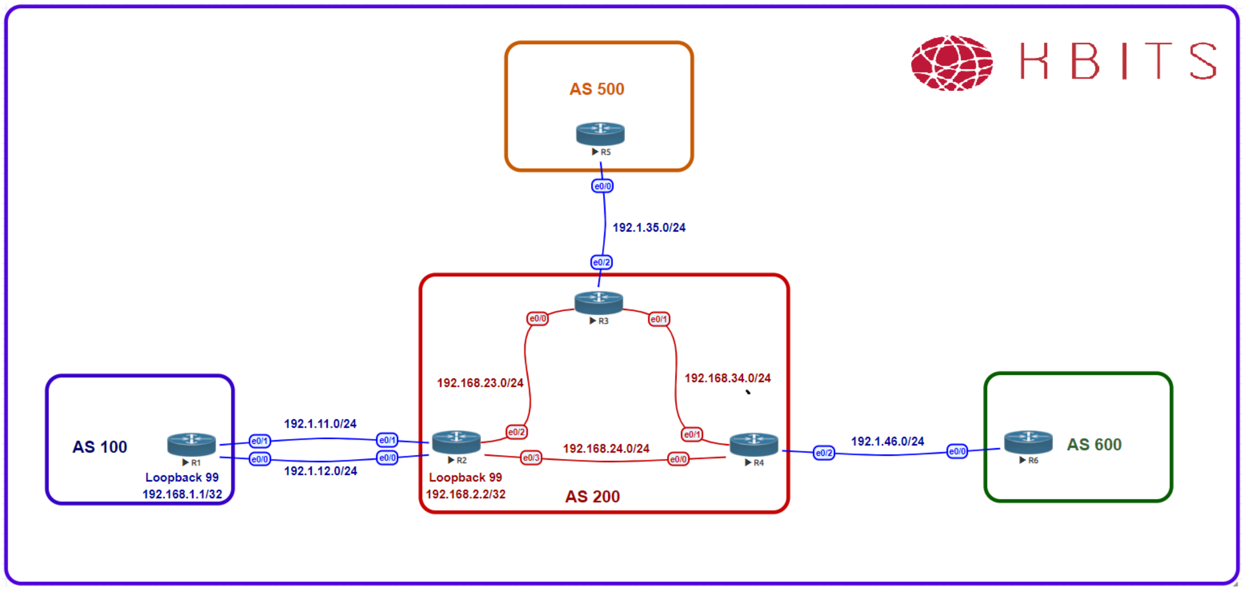

Task 1

Create the following Loopbacks on R3

Loopback 1 – 150.3.16.1/20
Loopback 2 – 150.3.36.1/22
Loopback 3 – 150.3.40.1/22
Loopback 4 – 150.3.50.1/23
Loopback 5 – 150.3.65.1/24
Loopback 6 – 150.13.0.1/16
Loopback 7 – 150.14.64.1/18

R3

```
interface Loopback1
 ip address 150.3.16.1 255.255.240.0
!
interface Loopback2
 ip address 150.3.36.1 255.255.252.0
!
interface Loopback3
 ip address 150.3.40.1 255.255.252.0
!
interface Loopback4
 ip address 150.3.50.1 255.255.254.0
!
interface Loopback5
```

```
 ip address 150.3.65.1 255.255.255.0
!
interface Loopback6
 ip address 150.13.0.1 255.255.0.0
!
interface Loopback7
 ip address 150.14.64.1 255.255.192.0
```

Task 2

Advertise the newly created routes in BGP using the Network command.

```
R3

Router bgp 200
 Network 150.3.16.0 mask 255.255.240.0
 Network 150.3.36.0 mask  255.255.252.0
 Network 150.3.40.0 mask 255.255.252.0
 Network 150.3.50.0 mask  255.255.254.0
 Network 150.3.65.0 mask 255.255.255.0
 Network 150.13.0.0
 Network 150.14.64.0 mask 255.255.192.0
```

Task 3

Configure R2 such that it blocks all the 150.X.X.0 routes that have a subnet mask between 17 and 23 bits coming in from R3.

```
R2

IP Prefix-list PLIST1 deny 150.0.0.0/8 ge 17 le 23
IP Prefix-list PLIST1 permit 0.0.0.0/0 le 32
!
Router bgp 200
 Neighbor 10.3.3.3 prefix-list PLIST1 in
```

Lab 8 – Route Filtering using AS Path-Filter

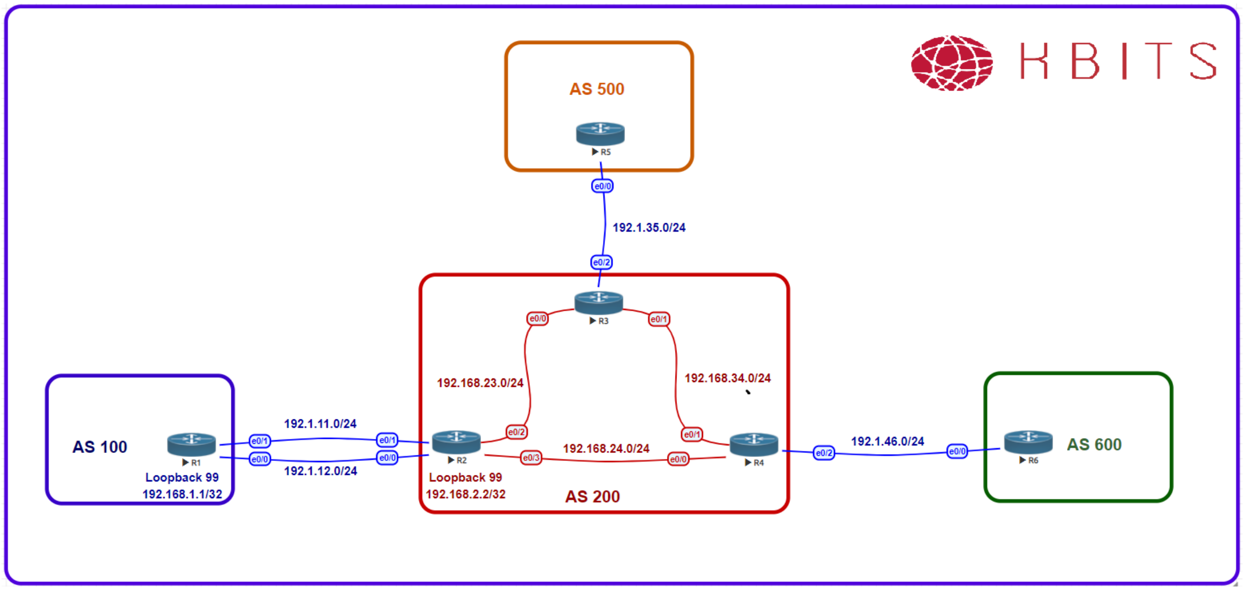

Task 1

Configure BGP such that AS 500 does not use AS 200 as a transit AS. Configuration should be done in AS 500.

```
R5

Ip as-path access-list 1 permit ^200$
!
Router BGP 500
 Neighbor 192.1.35.3 filter-list 1 in
```

Task 2

Configure BGP such that AS 100 does not use AS 200 to get AS 500 routes. Configuration should be done in AS 200. You are only allowed a single line in the AS-path filter.

```
R2

Ip as-path access-list 1 permit ^$
!
Router BGP 200
 Neighbor 192.168.1.1 filter-list 1 out
```

Lab 9 – Configuring Route Aggregation – Summary Only

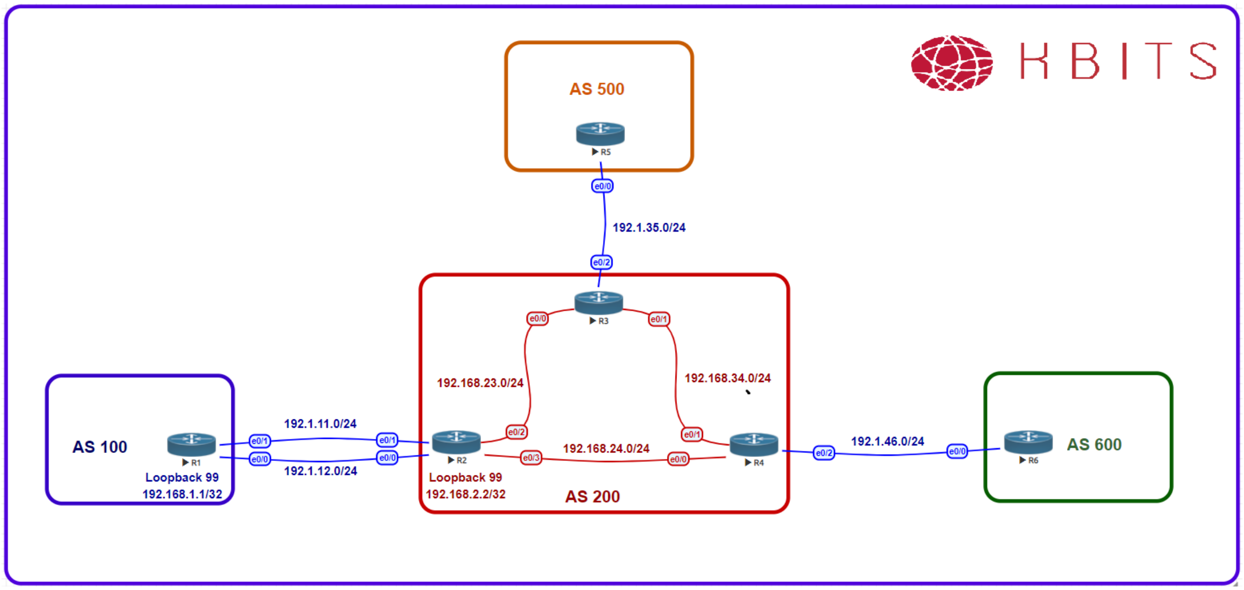

Task 1

Create the following Loopbacks on R3 and advertise them under BGP:

Loopback 1 – 203.1.4.1/24
Loopback 2 – 203.1.5.1/24
Loopback 3 – 203.1.6.1/24
Loopback 4 – 203.1.7.1/24

R3

```
interface Loopback1
 ip address 203.1.4.1 255.255.255.0
!
interface Loopback2
 ip address 203.1.5.1 255.255.255.0
!
interface Loopback3
 ip address 203.1.6.1 255.255.255.0
!
interface Loopback4
 ip address 203.1.7.1 255.255.255.0
!
Router BGP 200
 Network 203.1.4.0
 Network 203.1.5.0
```

```
 Network 203.1.6.0
 Network 203.1.7.0
```

Task 2

Configure Route Aggregation on R3 such that these routes are summarized as a single route. Only the Summary route should be send to R3's neighbors.

R3

```
Router bgp 200
 Aggregate-address 203.1.4.0 255.255.252.0 summary-only
```

Lab 10 – Configuring Route Aggregation – Manual Filtering

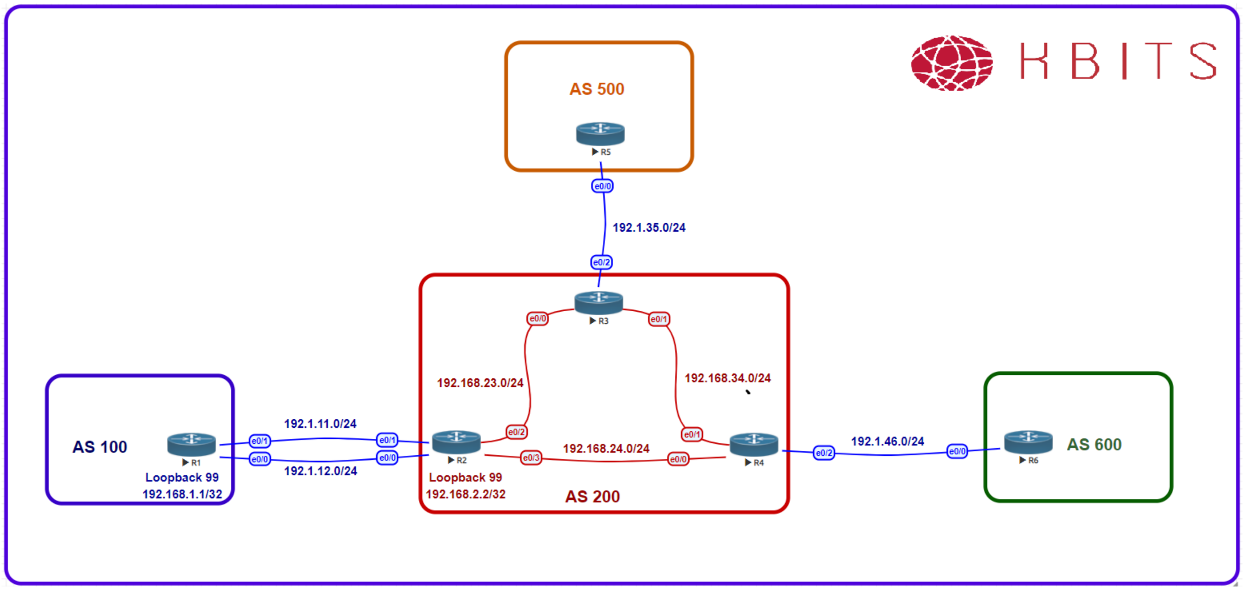

Task 1

Create the following Loopbacks on R4 and advertise them under BGP:

Loopback 1 – 204.1.4.1/24
Loopback 2 – 204.1.5.1/24
Loopback 3 – 204.1.6.1/24
Loopback 4 – 204.1.7.1/24

R4

```
interface Loopback1
 ip address 204.1.4.1 255.255.255.0
!
interface Loopback2
 ip address 204.1.5.1 255.255.255.0
!
interface Loopback3
 ip address 204.1.6.1 255.255.255.0
!
interface Loopback4
 ip address 204.1.7.1 255.255.255.0
!
Router BGP 200
```

```
 Network 204.1.4.0
 Network 204.1.5.0
 Network 204.1.6.0
 Network 204.1.7.0
```

Task 2

Configure Route Aggregation on R4 such that these routes are summarized as a single route. Only the Summary Route should be sent towards the eBGP neighbor (R6). Only the Specific Routes should be sent towards the iBGP neighbor (R3). The routes should not be seen as suppressed on R4.

R4

```
IP Prefix-list PLIST-R6 deny 204.1.4.0/22 ge 24
IP Prefix-list PLIST-R6 permit 0.0.0.0/0 le 32
!
IP Prefix-list PLIST-R3 deny 204.1.4.0/22
IP Prefix-list PLIST-R3 permit 0.0.0.0/0 le 32
!
Router bgp 200
 Aggregate-address 204.1.4.0 255.255.252.0
 Neighbor 192.1.46.6 prefix-list PLIST-R6 out
 Neighbor 10.3.3.3 prefix-list PLIST-R3 out
```

Lab 11 – Configuring Route Aggregation – Suppress Maps

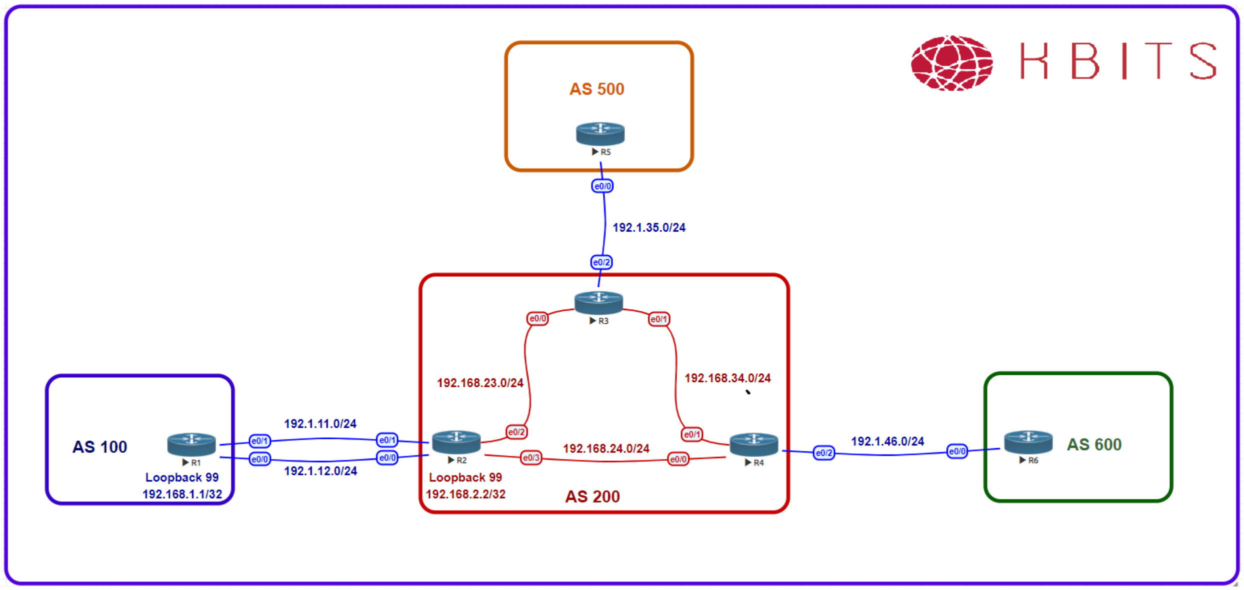

Task 1

Create the following Loopbacks on R2 and advertise them under BGP:

Loopback 1 – 202.1.4.1/24
Loopback 2 – 202.1.5.1/24
Loopback 3 – 202.1.6.1/24
Loopback 4 – 202.1.7.1/24

R2

```
interface Loopback1
 ip address 202.1.4.1 255.255.255.0
!
interface Loopback2
 ip address 202.1.5.1 255.255.255.0
!
interface Loopback3
 ip address 202.1.6.1 255.255.255.0
!
interface Loopback4
 ip address 202.1.7.1 255.255.255.0
!
Router BGP 234
```

```
 Network 202.1.4.0
 Network 202.1.5.0
 Network 202.1.6.0
 Network 202.1.7.0
```

Task 2

Configure Route Aggregation on R2 such that these routes are summarized as a single route. Only the Summary route and the 202.1.5.0 route should be send to R2's neighbors. The other specific routes should be seen as suppressed on R2.

R2

```
Access-list 5 permit 202.1.5.0 0.0.0.255
!
Route-map SUPMAP deny 10
 Match address 5
Route-map SUPMAP permit 20
!
Router bgp 200
 Aggregate-address 202.1.4.0 255.255.252.0 supress-map SUPMAP
```

Lab 12 – Configuring Base BGP Topology – eBGP & iBGP

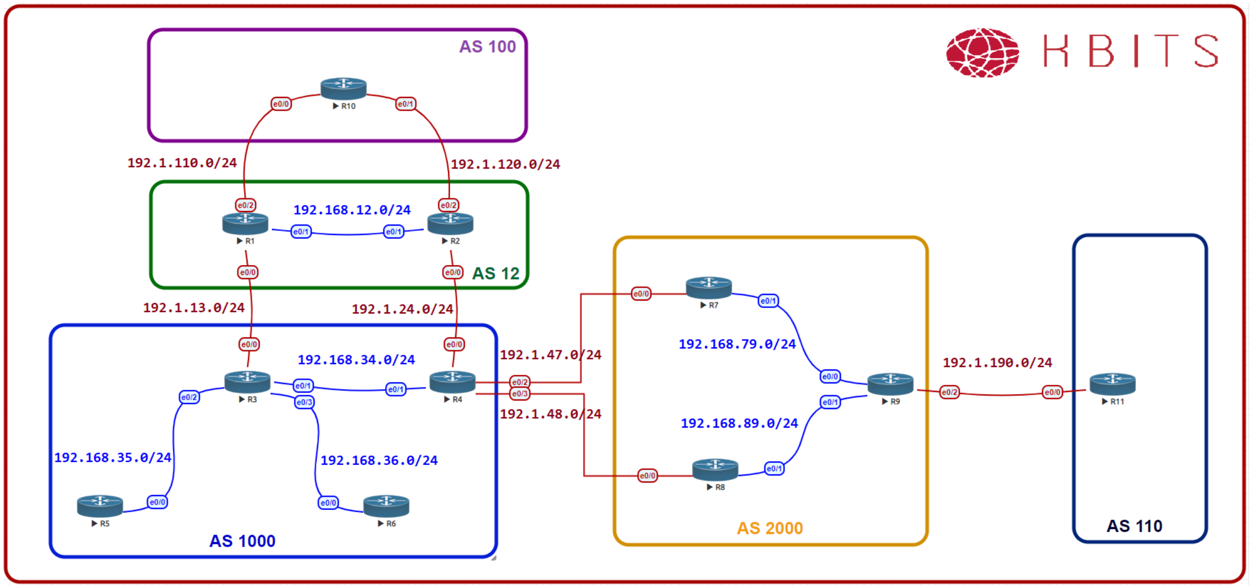

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 1	11.1.1.1	255.255.255.0
Loopback 10	10.1.1.1	255.255.255.255
E 0/0	192.1.13.1	255.255.255.0
E 0/1	192.168.12.1	255.255.255.0
E 0/2	192.1.110.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
Loopback 1	22.2.2.2	255.255.255.0
Loopback 10	10.2.2.2	255.255.255.255
E 0/0	192.1.24.2	255.255.255.0
E 0/1	192.168.12.2	255.255.255.0
E 0/2	192.1.120.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
Loopback 1	33.3.3.3	255.255.255.0
Loopback 10	10.3.3.3	255.255.255.255
E 0/0	192.1.13.3	255.255.255.0
E 0/1	192.168.34.3	255.255.255.0
E 0/2	192.168.35.3	255.255.255.0
E 0/3	192.168.36.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
Loopback 1	44.4.4.4	255.255.255.0
Loopback 10	10.4.4.4	255.255.255.255
E 0/0	192.1.24.4	255.255.255.0
E 0/1	192.168.34.4	255.255.255.0
E 0/2	192.1.47.4	255.255.255.0
E 0/3	192.1.48.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
Loopback 1	55.5.5.5	255.255.255.0
E 0/0	192.168.35.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
Loopback 1	66.6.6.6	255.255.255.0
E 0/0	192.168.36.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
Loopback 1	77.7.7.7	255.255.255.0
Loopback 10	10.7.7.7	255.255.255.255
E 0/0	192.1.47.7	255.255.255.0
E 0/1	192.168.79.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
Loopback 1	88.8.8.8	255.255.255.0
Loopback 10	10.8.8.8	255.255.255.255
E 0/0	192.1.48.8	255.255.255.0
E 0/1	192.168.89.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	9.9.9.9	255.0.0.0
Loopback 1	99.9.9.9	255.255.255.0
Loopback 10	10.9.9.9	255.255.255.255
E 0/0	192.168.79.9	255.255.255.0
E 0/1	192.168.89.9	255.255.255.0
E 0/2	192.1.190.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	100.100.100.10	255.0.0.0
Loopback 1	101.101.101.10	255.255.255.0
E 0/0	192.1.110.10	255.255.255.0
E 0/1	192.1.120.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	111.111.111.11	255.0.0.0
Loopback 1	112.112.112.11	255.255.255.0
E 0/0	192.1.190.11	255.255.255.0

Task 1

Configure eBGP neighbor relationships between R10 in AS 100 with R1 & R2 in AS 12. Advertise all public loopback networks in BGP.

R1	**R2**
Router BGP 12 Network 1.0.0.0 Network 11.1.1.0 mask 255.255.255.0 Neighbor 192.1.110.10 remote-as 100	Router BGP 12 Network 2.0.0.0 Network 22.2.2.0 mask 255.255.255.0 Neighbor 192.1.120.10 remote-as 100
R10	
Router BGP 100 Network 100.0.0.0 Network 101.101.101.0 mask 255.255.255.0 Neighbor 192.1.110.1 remote-as 12 Neighbor 192.1.120.1 remote-as 12	

Task 2

Configure iBGP neighbor relationships between R1 & R2 in AS 12. Configure the neighbor relationship based on a private loopback address. Use EIGRP 12 as the underlay IGP.

R1	**R2**
Router eigrp 12 Network 192.168.12.0 Network 10.0.0.0 ! Router BGP 12 Neighbor 10.2.2.2 remote-as 12 Neighbor 10.2.2.2 update-source Lo10 Neighbor 10.2.2.2 next-hop-self	Router eigrp 12 Network 192.168.12.0 Network 10.0.0.0 ! Router BGP 12 Neighbor 10.1.1.1 remote-as 12 Neighbor 10.1.1.1 update-source Lo10 Neighbor 10.1.1.1 next-hop-self

Task 3

Configure eBGP neighbor relationships between R1 in AS 12 and R3 in AS 1000. Advertise all public loopback networks on R3 in BGP.

R1	R3
Router BGP 12 Neighbor 192.1.13.3 remote-as 1000	Router BGP 1000 Network 3.0.0.0 Network 33.3.3.0 mask 255.255.255.0 Neighbor 192.1.13.1 remote-as 12

Task 4

Configure eBGP neighbor relationships between R2 in AS 12 and R4 in AS 1000. Advertise all public loopback networks on R4 in BGP.

R2	R4
Router BGP 12 Neighbor 192.1.24.4 remote-as 1000	Router BGP 1000 Network 4.0.0.0 Network 44.4.4.0 mask 255.255.255.0 Neighbor 192.1.24.2 remote-as 12

Task 5

Configure iBGP neighbor relationships between R3 & R4 in AS 1000. Configure the neighbor relationship based on the physical link.

R3	R4
Router BGP 1000 Neighbor 192.168.34.4 remote-as 1000 Neighbor 192.168.34.4 next-hop-self	Router BGP 1000 Neighbor 192.168.34.3 remote-as 1000 Neighbor 192.168.34.3 next-hop-self

Task 6

Configure OSPF as the IGP to connect R3 to R6 in Area 0. Only enable OSPF on the R3-R6 physical link on R3. Enable OSPF on all interfaces on R6 in area 0. Configure Mutual Redistribution on R3 between OSPF and BGP

R3	**R6**
Router ospf 1 Network 192.168.36.0 0.0.0.255 area 0 Redistribute bgp 1000 ! Router bgp 1000 Redistribute ospf 1	Router ospf 1 Network 6.0.0.0 0.255.255.255 area 0 Network 66.0.0.0 0.255.255.255 area 0 Network 192.168.36.0 0.0.0.255 area 0

Task 7

Configure iBGP neighbor relationships between R3 & R5 in AS 1000. Configure the neighbor relationship based on the physical link.

R3	**R5**
Router BGP 1000 Neighbor 192.168.35.5 remote-as 1000 Neighbor 192.168.35.5 next-hop-self	Router BGP 1000 Neighbor 192.168.35.3 remote-as 1000 Neighbor 192.168.35.3 next-hop-self

Task 8

Configure eBGP neighbor relationships between R4 in AS 1000 with R7 & R8 in AS 2000. Advertise all public loopback networks in BGP on R7 & R8.

R7	**R8**
Router BGP 2000 Network 7.0.0.0 Network 77.7.7.0 mask 255.255.255.0 Neighbor 192.1.47.4 remote-as 1000	Router BGP 2000 Network 8.0.0.0 Network 88.8.8.0 mask 255.255.255.0 Neighbor 192.1.48.4 remote-as 1000
R4 Router BGP 1000 Neighbor 192.1.47.4 remote-as 2000 Neighbor 192.1.48.4 remote-as 2000	

Task 9

Configure iBGP neighbor relationships between R7, R8 & R9 in AS 2000. Advertise the public loopback addresses of R9 in BGP. Configure the neighbor relationship based on a private loopback address. Use IS-IS in area 49.0000 as the underlay IGP. Use System-ID on your choice. Configure R9 as the Route Reflector for R7 & R8. Do not configure a direct BGP peering between R7 & R8.

R7	**R8**
Router isis Net 49.0000.7777.7777.7777.00 Is-type level-2 ! Interface E 0/1 Ip router isis ! Interface Loopback10 Ip router isis ! Router BGP 2000 Neighbor 10.9.9.9 remote-as 2000 Neighbor 10.9.9.9 update-source Lo10 Neighbor 10.9.9.9next-hop-self	Router isis Net 49.0000.8888.8888.8888.00 Is-type level-2 ! Interface E 0/1 Ip router isis ! Interface Loopback10 Ip router isis ! Router BGP 2000 Neighbor 10.9.9.9 remote-as 2000 Neighbor 10.9.9.9 update-source Lo10 Neighbor 10.9.9.9next-hop-self
R9 Router isis Net 49.0000.9999.9999.9999.00 Is-type level-2 ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis ! Interface Loopback10 Ip router isis ! Router BGP 2000 Network 9.0.0.0 Network 99.9.9.0 mask 255.255.255.0 Neighbor IBGP peer-group Neighbor IBGP remote-as 2000 Neighbor IBGP update-source Lo10	

```
 Neighbor IBGP next-hop-self
 Neighbor IBGP route-reflector-client
 Neighbor 10.7.7.7 peer-group IBGP
 Neighbor 10.8.8.8 peer-group IBGP
```

Task 10

Configure eBGP neighbor relationships between R9 in AS 2000 and R11 in AS 110. Advertise all public loopback networks on R11 in BGP.

R9

```
Router BGP 2000
 Neighbor 192.1.190.11 remote-as 110
```

R11

```
Router BGP 110
 Network 111.0.0.0
 Network 112.112.112.0 mask 255.255.255.0
 Neighbor 192.1.190.9 remote-as 2000
```

Lab 13 – Configuring BGP Attributes – Local Preference

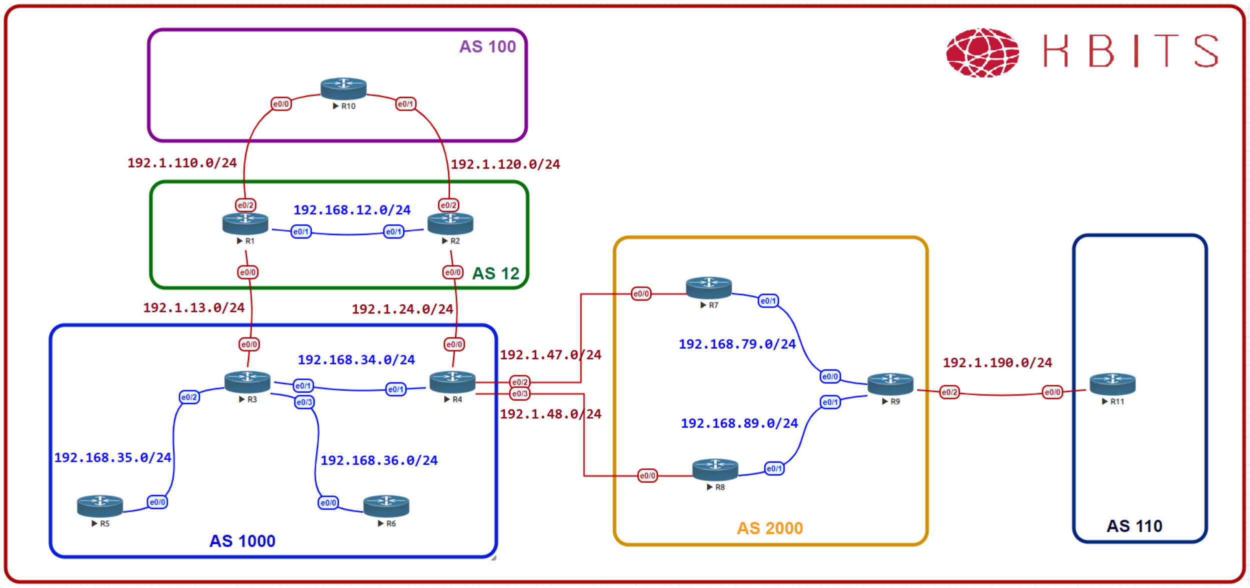

Task 1

Configure AS 2000 such that it prefers the Link between R4-R7 for traffic leaving AS 2000 towards AS 1000.

R7

```
route-map SETATT permit 10
 set local-preference 111
!
router bgp 2000
 neighbor 192.1.47.4 route-map SETATT in
```

Task 2

Configure AS 2000 such that it prefers the Link between R4-R8 for traffic destined towards 1.4.1.0/24 & 1.4.2.0/24 leaving AS 2000 towards AS 1000. R4-R8 should be preferred link only for 1.4.1.0/24 & 1.4.2.0/24. The rest should continue to use R4-R7.

R8

```
ip prefix-list PL1 permit 1.4.1.0/24
ip prefix-list PL1 permit 1.4.2.0/24
!
route-map SETATT permit 10
 match ip address prefix PL1
 set local-preference 115
route-map SETATT permit 20
!
router bgp 2000
 neighbor 192.1.48.4 route-map SETATT in
```

Lab 14 – Configuring BGP Attributes – MED

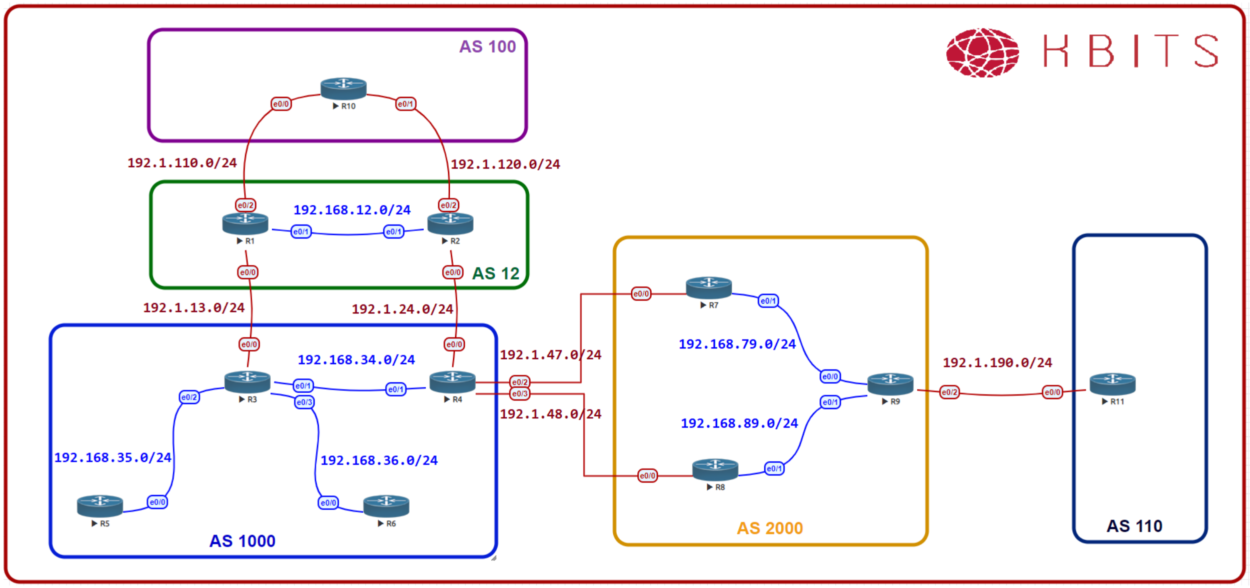

Task 1

Configure AS 2000 such that it prefers the Link between R4-R8 for traffic entering AS 2000 from AS 1000.

R7

```
route-map SETMED permit 10
 set metric 77
!
router bgp 2000
 neighbor 192.1.47.4 route-map SETMED out
```

Lab 15 – Configuring BGP Attributes – Weight

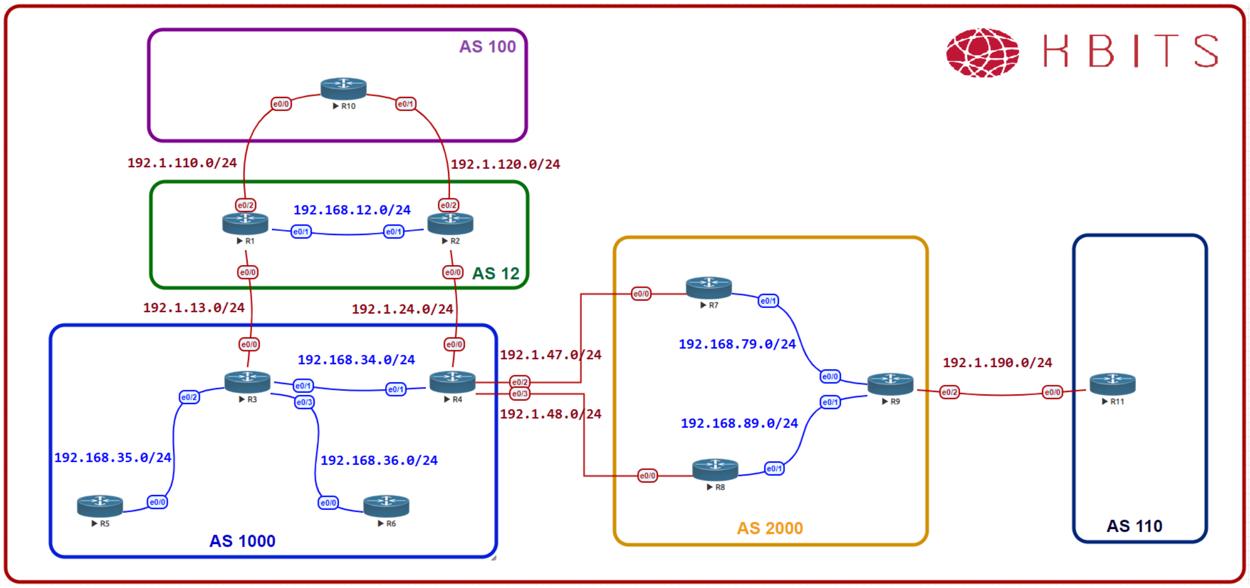

Task 1

Configure R8 such that all traffic towards AS1000 should use the Link between R4 & R8 as the preferred link. This should only affect the local router and not the rest of the AS.

R8

```
route-map SETWT
 set weight 88
!
router bgp 2000
 neighbor 192.1.48.4 route-map SETWT in
```

Lab 16 – Configuring BGP Attributes – AS-Path

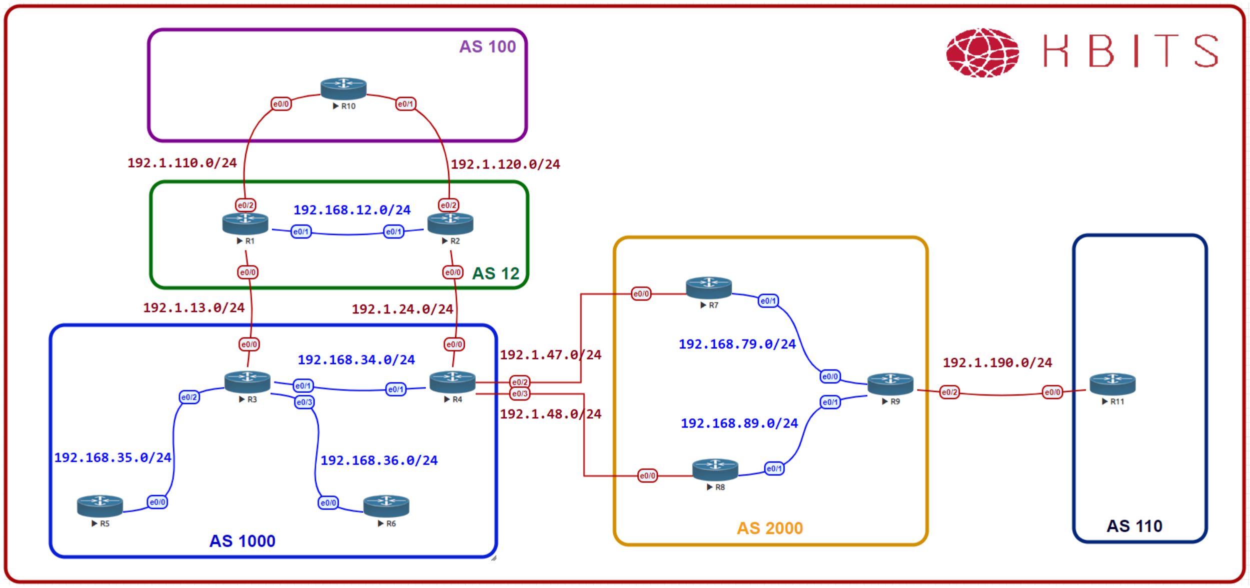

Task 1

De-configure the Route-map from the previous 3 labs. This is done so that we can accomplish the same tasks using the AS-Path attribute

R7

```
No route-map SETATT
No route-map SETMED
!
router bgp 2000
 no neighbor 192.1.47.4 route-map SETATT in
 no neighbor 192.1.47.4 route-map SETMED out
```

R8

```
No route-map SETATT
No route-map SETWT
!
router bgp 2000
 no neighbor 192.1.48.4 route-map SETATT in
 no neighbor 192.1.48.4 route-map SETWT in
```

Task 2

Configure AS 2000 such that it prefers the Link between R4-R7 for traffic leaving AS2000 towards AS1000. Use the AS-Path attribute to accomplish this task.

R8

```
route-map SETAS permit 10
 set as-path prepend 1000
!
router bgp 2000
 neighbor 192.1.48.4 route-map SETAS in
```

Task 3

Configure AS 2000 such that it prefers the Link between R4-R8 for traffic entering AS2000 from AS1000. Use the AS-Path attribute to accomplish this task.

R7

```
route-map SETAS permit 10
 set as-path prepend 2000
!
router bgp 2000
 neighbor 192.1.47.4 route-map SETAS out
```

Lab 17 – Configuring BGP Attributes – No-Export Community Attribute

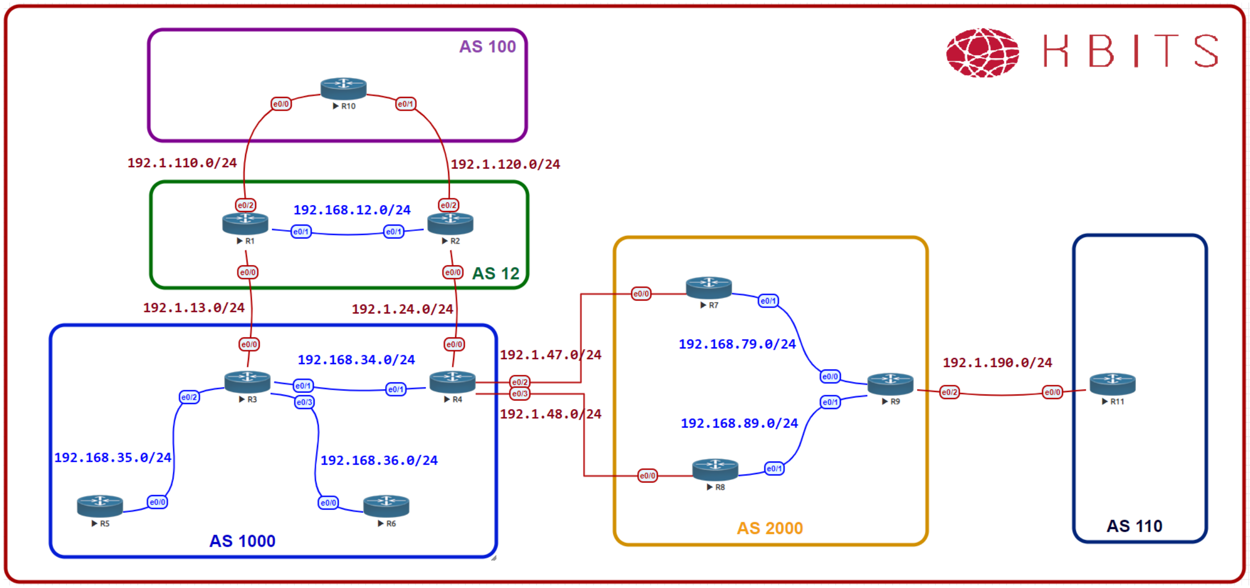

Task 1

AS110 wants to limit the propagation of 111.0.0.0/8 network to AS2000 only. AS2000 should not export this route outside AS2000. Use the appropriate Community attribute to accomplish this.

R11

```
ip prefix-list PL1 permit 111.0.0.0/8
!
route-map SETCOMM permit 10
 match ip address prefix PL1
 set community no-export
route-map SETCOMM permit 20
!
router bgp 110
 neighbor 192.1.190.9 route-map SETCOMM out
 neighbor 192.1.190.9 send-community standard
```

R9

```
router bgp 110
 neighbor 10.7.7.7 send-community standard
 neighbor 10.8.8.8 send-community standard
```

Lab 18 – Configuring BGP Attributes – No-Advertise Community Attribute

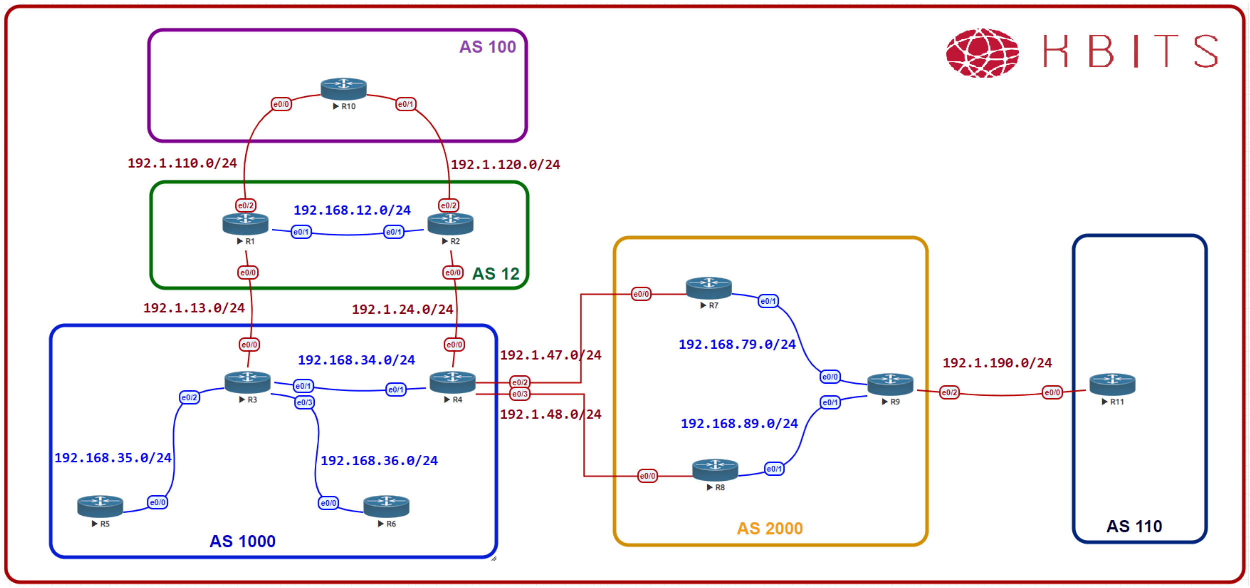

Task 1

AS110 wants to limit the propagation of 112.112.1112.0/24 network to R9 only. R9 should not forward this network to anyone including the iBGP Neighbors. Use the appropriate Community attribute to accomplish this.

R11

```
ip prefix-list PL2 permit 112.112.112.0/24
!
route-map SETCOMM permit 5
 match ip address prefix PL2
 set community no-advertise
```

Lab 19 – Configuring BGP Conditional Advertisement

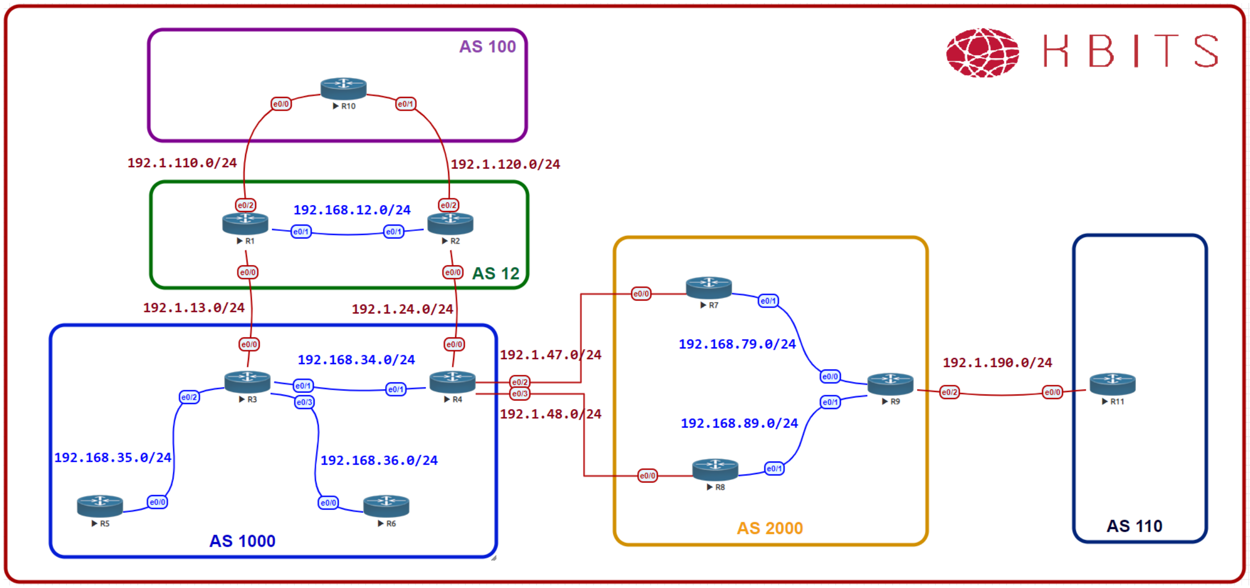

Task 1

De-configure the Route-map from the previous 3 labs. This is done so that we have all the routes present for the next set of labs.

R8 no route-map SETAS ! router bgp 2000 no neighbor 192.1.48.4 route-map SETAS in
R7 no route-map SETAS permit ! router bgp 2000 no neighbor 192.1.47.4 route-map SETAS out
R11 no route-map SETCOMM ! router bgp 110 no neighbor 192.1.190.9 route-map SETCOMM out

Task 2
Configure a loopback on R7 and advertise it thru BGP. This will be used to check the status of R7 in a later step.

R7

```
Interface Loopback99
 ip address 10.77.77.77 255.255.255.255
!
router bgp 2000
 network 10.77.77.77 mask 255.255.255.255
```

Task 3
Configure a route-map on R8 to classify the route that will be conditionally advertised.

R8

```
ip prefix-list PL2 permit 111.0.0.0/8
ip prefix-list PL2 permit 112.112.112.0/24
!
route-map AMAP
 match ip address prefix PL2
```

Task 4
Configure a route-map on R8 to call an ACL that will indicate the absence of the 10.77.77.77/32 route.

R8

```
ip prefix-list PL3 permit 10.77.77.77/32
!
route-map NEM
 match ip address prefix PL3
```

Task 5
Configure the Conditional Advertisement of the 111.0.0.0/8 & 112.112.112.0/24 routes from R8 to R4 only if R7 is down.

R8

```
router bgp 2000
 neighbor 192.1.48.4 advertise-map AMAP non-exist-map NEM
```

Lab 20 – Configuring BGP Multi-Path – eBGP – iBGP

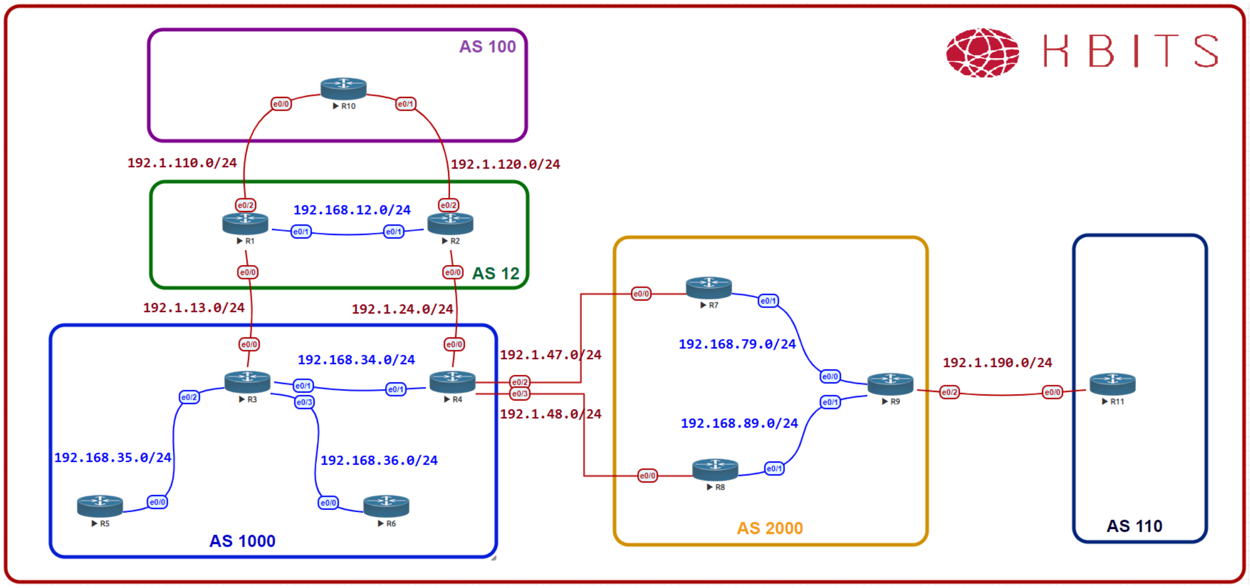

Task 1

Configure R10 to allow it to inject multiple routes on the Links between R10-R1 & R10-R2 (eBGP Neighbors).

```
R10

Router bgp 100
 maximum-paths 2
```

Task 2

Configure R9 to allow it to inject multiple routes on the Links between R9-R7 & R9-R8 (iBGP Neighbors).

```
R9

Router bgp 2000
 maximum-paths ibgp 2
```

Task 2

Configure R3 to allow it to inject multiple routes on the Links between R1-R3(eBGP) & R3-R4 (iBGP Neighbors)

R3
Router bgp 1000 maximum-paths eibgp 2

Lab 21 – Configuring to Redistribute iBGP Routes into IGP

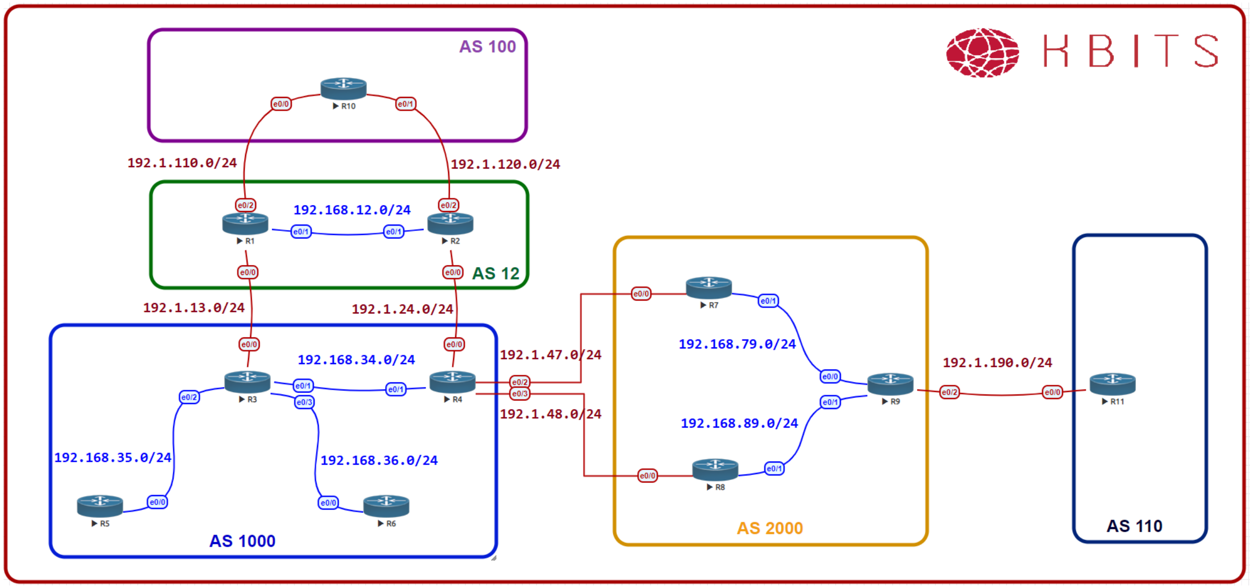

Task 1

Check the routing table of R6. Does it have all the routes from AS2000 & AS110?

Task 2

Configure R3 to redistribute iBGP routes into IGP.

R3

```
router bgp 1000
 bgp redistribute-internal
```

Lab 22 – Configuring BGP Route Reflector with Next-Hop Changed

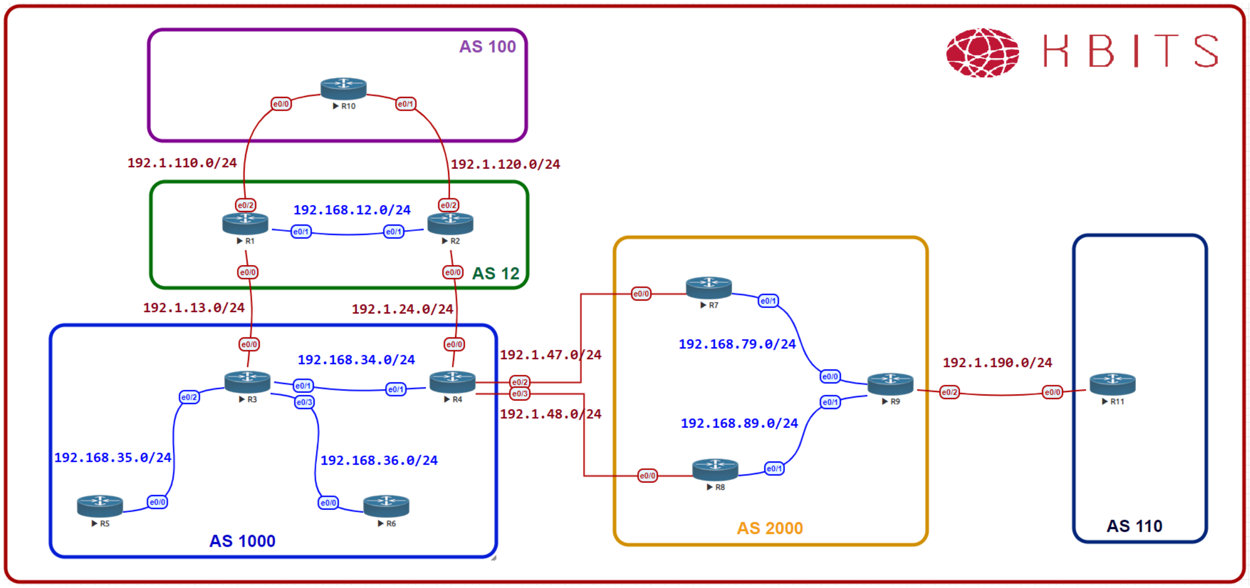

Task 1

Configure R3 as the Route Reflector between R4 & R5. Make sure to change the next-hop to R3.

```
R3

Router bgp 1000
 Neighbor 192.168.35.5 route-reflector-client
 Neighbor 192.168.35.5 next-hop-self all
 Neighbor 192.168.34.4 route-reflector-client
 Neighbor 192.168.34.4 next-hop-self all
```

Lab 23 – Configuring BGP Route Reflection based on Dynamic Neighbors

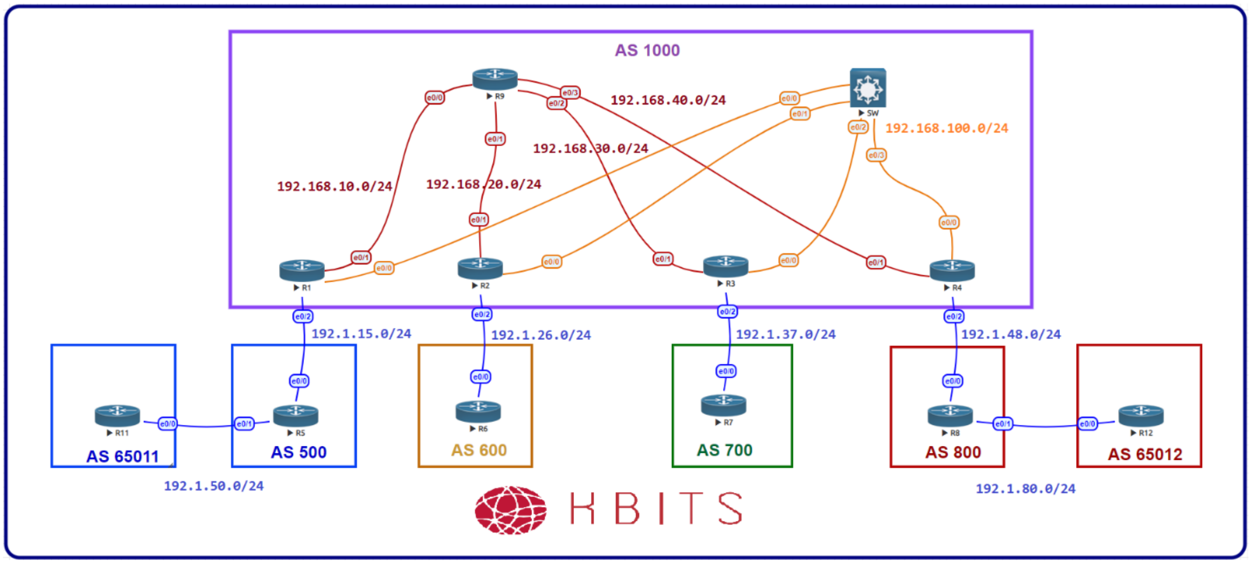

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 10	172.16.1.1	255.255.255.255
E 0/0	192.168.100.1	255.255.255.0
E 0/1	192.168.10.1	255.255.255.0
E 0/2	192.1.15.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
Loopback 10	172.16.1.2	255.255.255.255
E 0/0	192.168.100.2	255.255.255.0
E 0/1	192.168.20.2	255.255.255.0
E 0/2	192.1.26.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
Loopback 10	172.16.1.3	255.255.255.255
E 0/0	192.168.100.3	255.255.255.0
E 0/1	192.168.30.3	255.255.255.0
E 0/2	192.1.37.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
Loopback 10	172.16.1.4	255.255.255.255
E 0/0	192.168.100.4	255.255.255.0
E 0/1	192.168.40.4	255.255.255.0
E 0/2	192.1.48.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
E 0/0	192.1.15.5	255.255.255.0
E 0/1	192.1.50.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
E 0/0	192.1.26.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
E 0/0	192.1.37.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
E 0/0	192.1.48.8	255.255.255.0
E 0/2	192.1.80.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 10	172.16.1.9	255.255.255.255
E 0/0	192.168.10.9	255.255.255.0
E 0/1	192.168.20.9	255.255.255.0
E 0/2	192.168.30.9	255.255.255.0
E 0/3	192.168.40.9	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	11.11.11.11	255.0.0.0
E 0/0	192.1.50.11	255.255.255.0

R12

Interface	IP Address	Subnet Mask
Loopback 0	12.12.12.12	255.0.0.0
E 0/0	192.1.80.12	255.255.255.0

Task 1

Configure EIGRP 111 as the underlay IGP to route the Loopback 10 networks on the underlay networks

R1 Router eigrp 1000 network 192.168.100.0 network 192.168.10.0 network 172.16.1.0 0.0.0.255	**R2** Router eigrp 1000 network 192.168.100.0 network 192.168.20.0 network 172.16.1.0 0.0.0.0
R3 Router eigrp 1000	**R4** Router eigrp 1000

network 192.168.100.0 network 192.168.30.0 network 172.16.1.0 0.0.0.0	network 192.168.100.0 network 192.168.40.0 network 172.16.1.0 0.0.0.0
R9 Router eigrp 1000 network 192.168.10.0 network 192.168.20.0 network 192.168.30.0 network 192.168.40.0 network 172.16.1.0 0.0.0.0	

Task 2

Configuring iBGP between the ASBR (R1,R2,R3 & R4) and the RR (R9) based on Loopbacks. Configure R9 such that it accepts neighbor requests from any router from the 172.16.1.0/24 subnet. Authenticate the neighbor relationship with a password of **ccie12353**. Advertise the Loopback 0 networks on ASBRs in BGP.

R1 Router BGP 1000 Network 1.0.0.0 Neighbor 172.16.1.9 remote-as 1000 Neighbor 172.16.1.9 update-source Lo10 Neighbor 172.16.1.9 next-hop-self Neighbor 172.16.1.9 password ccie12353	**R2** Router BGP 1000 Network 2.0.0.0 Neighbor 172.16.1.9 remote-as 1000 Neighbor 172.16.1.9 update-source Lo10 Neighbor 172.16.1.9 next-hop-self Neighbor 172.16.1.9 password ccie12353
R3 Router BGP 1000 Network 3.0.0.0 Neighbor 172.16.1.9 remote-as 1000 Neighbor 172.16.1.9 update-source Lo10 Neighbor 172.16.1.9 next-hop-self Neighbor 172.16.1.9 password ccie12353	**R4** Router BGP 1000 Network 4.0.0.0 Neighbor 172.16.1.9 remote-as 1000 Neighbor 172.16.1.9 update-source Lo10 Neighbor 172.16.1.9 next-hop-self Neighbor 172.16.1.9 password ccie12353
R9 router bgp 1000 neighbor IBGP peer-group neighbor IBGP remote-as 1000 neighbor IBGP update-source loopback10 neighbor IBGP route-reflector-client neighbor IBGP password ccie12353 bgp listen range 172.16.1.0/24 peer-group IBGP	

Task 3

Configuring eBGP neighbor relationship between AS 1000 and the connected ASs on the appropriate ASBRs. Advertise Loopack0 networks on R5, R6, R7 & R8.

R1 router bgp 1000 neighbor 192.1.15.5 remote-as 500
R2 router bgp 1000 neighbor 192.1.26.6 remote-as 600
R3 router bgp 1000 neighbor 192.1.37.7 remote-as 700
R4 router bgp 1000 neighbor 192.1.48.8 remote-as 800
R5 router bgp 500 network 5.0.0.0 neighbor 192.1.15.1 remote-as 1000
R6 router bgp 600 network 6.0.0.0 neighbor 192.1.26.2 remote-as 1000
R7 router bgp 700 network 7.0.0.0 neighbor 192.1.37.3 remote-as 1000
R8 router bgp 800 network 8.0.0.0 neighbor 192.1.48.4 remote-as 1000

Lab 24 – Working with Private AS Numbers

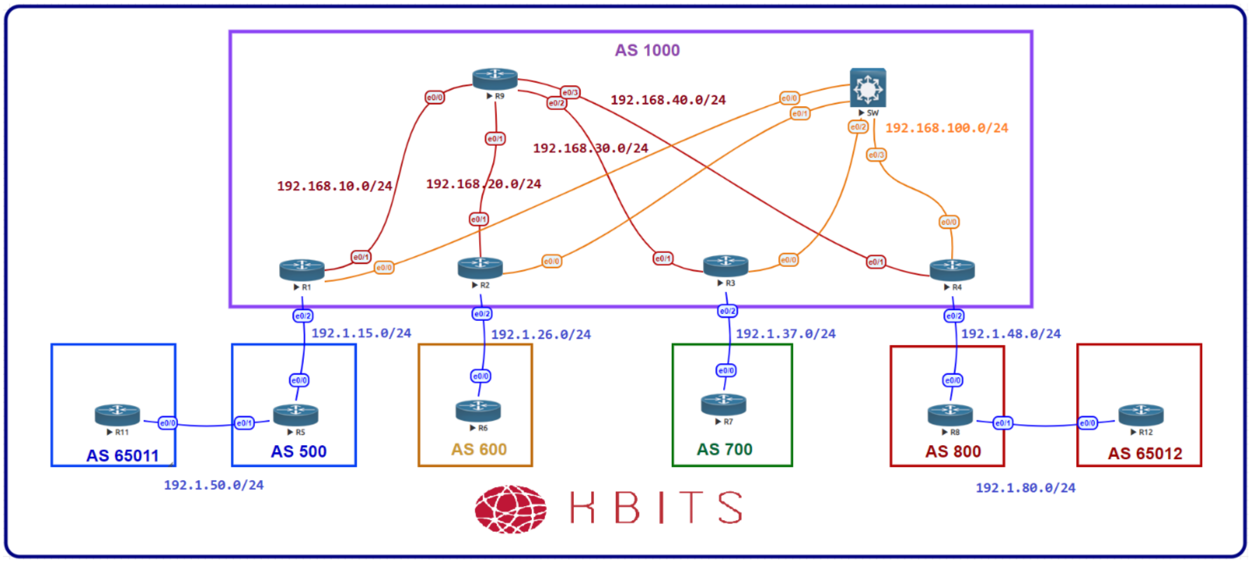

Task 1

Configure a relationship between the Customer (R12) and AS 800. The Customer should use AS 65012 as the AS #. Advertise the Loopback 0 network on R12.

R8 router bgp 800 neighbor 192.1.80.12 remote-as 65012
R12 router bgp 65012 neighbor 192.1.80.8 remote-as 800 network 12.0.0.0

Task 2

Configure R8 such that it removes the Private AS # from the AS Path before propagating the route towards AS 1000

R8 router bgp 800 neighbor 192.1.48.4 remove-private-as

Lab 25 – Configuring the Local-AS Command

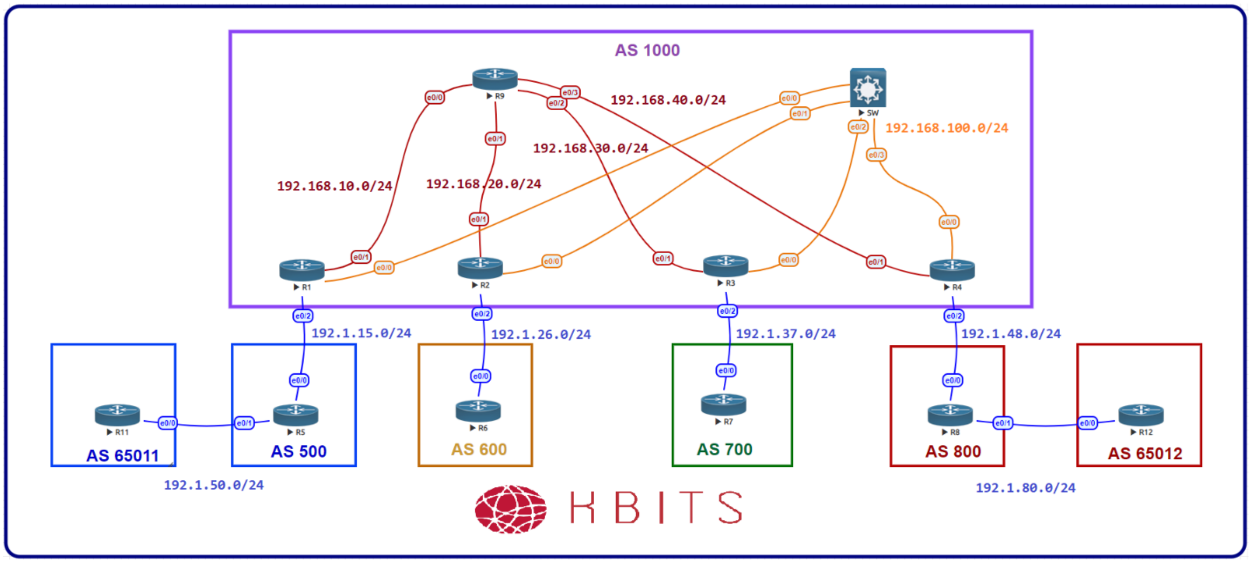

Task 1

Configure a relationship between the Customer (R11) and AS 500. The Customer should use AS 65011 as the AS #. Advertise the Loopback 0 network on R11.

R5 router bgp 500 neighbor 192.1.50.11 remote-as 65011
R11 router bgp 65011 neighbor 192.1.50.5 remote-as 500 network 11.0.0.0

Task 2

Configure R5 such that it removes the Private AS # from the AS Path before propagating the route towards AS 1000

R5 router bgp 500 neighbor 192.1.15.1 remove-private-as

Task 3

R11 acquires and configures a new AS #. It is a Public AS# 110. R5 will change the neighbor relationship in a maintenance window after 5 days. In the meanwhile R11 needs to change the AS # to the new to establish a new neighbor relationship with a new SP. Allow R11 to establish both neighbor relationships.

R11

```
no router bgp 65011
router bgp 110
 network 1.11.1.0 mask 255.255.255.0
 neighbor 192.1.50.5 remote-as 500
 neigbhor 192.1.50.5 local-as 65011
```

Lab 26 – Configuring BFD for BGP

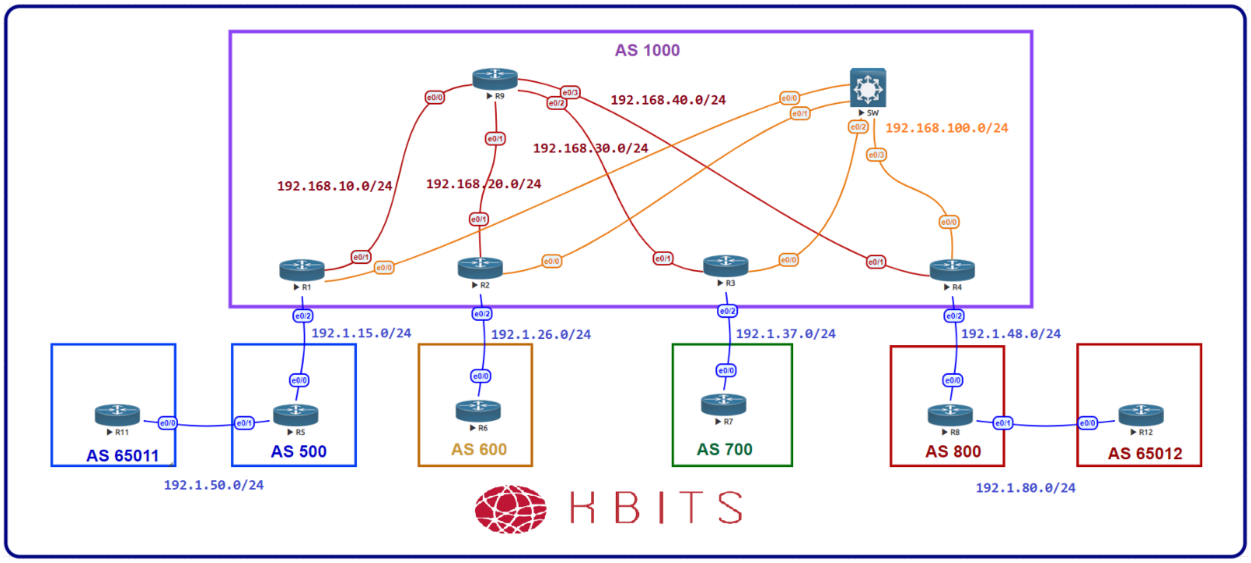

Task 1

Configure BFD using a send and receive interval of 350 ms. A neighbor is deemed dead if 3 hellos are missed. Configure it for the following eBGP neighbor relationships:

R1 - R5
R2 - R6
R3 - R7
R4 - R8

R1

```
Interface E 0/2
 bfd interval 350 min_rx 350 multiplier 3
!
router bgp 1000
 neighbor 192.1.15.5 fall-over bfd
```

R5

```
Interface E 0/0
 bfd interval 350 min_rx 350 multiplier 3
!
router bgp 500
 neighbor 192.1.15.1 fall-over bfd
```

R2

Interface E 0/2 bfd interval 350 min_rx 350 multiplier 3 ! router bgp 1000 neighbor 192.1.26.6 fall-over bfd
R6 Interface E 0/0 bfd interval 350 min_rx 350 multiplier 3 ! router bgp 600 neighbor 192.1.26.2 fall-over bfd
R3 Interface E 0/2 bfd interval 350 min_rx 350 multiplier 3 ! router bgp 1000 neighbor 192.1.37.7 fall-over bfd
R7 Interface E 0/0 bfd interval 350 min_rx 350 multiplier 3 ! router bgp 700 neighbor 192.1.37.3 fall-over bfd
R4 Interface E 0/2 bfd interval 350 min_rx 350 multiplier 3 ! router bgp 1000 neighbor 192.1.48.8 fall-over bfd
R8 Interface E 0/0 bfd interval 350 min_rx 350 multiplier 3 ! router bgp 800 neighbor 192.1.48.4 fall-over bfd

Lab 27 – Configuring BGP Confederations

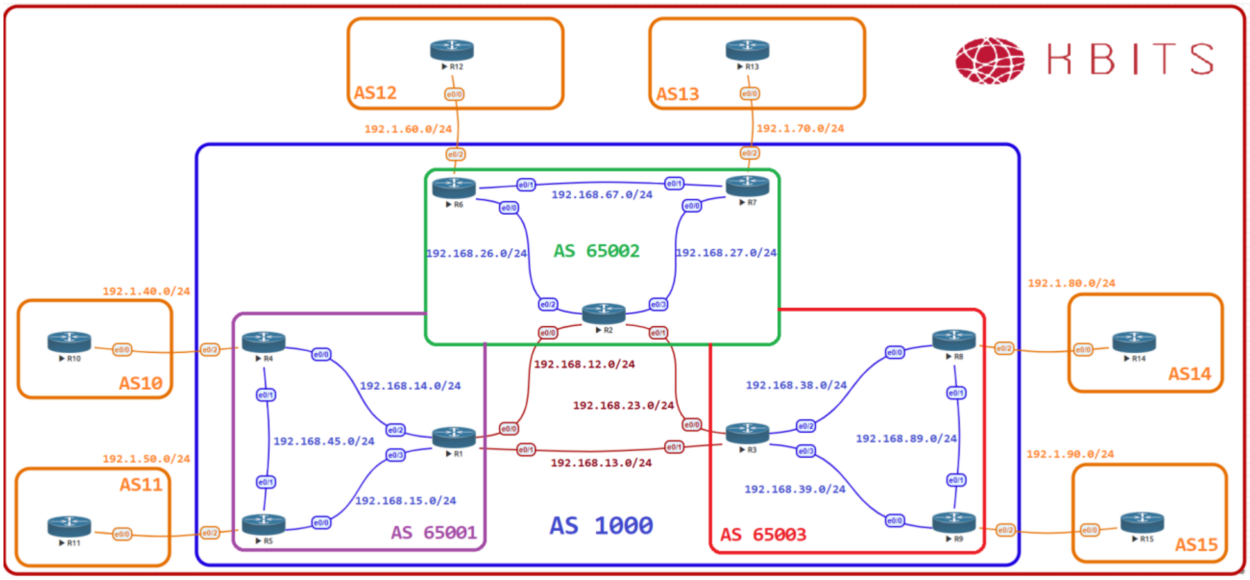

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.0.0.0
Loopback 10	172.16.1.1	255.255.255.255
E 0/0	192.168.12.1	255.255.255.0
E 0/1	192.168.13.1	255.255.255.0
E 0/2	192.168.14.1	255.255.255.0
E 0/3	192.168.15.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.0.0.0
Loopback 10	172.16.1.2	255.255.255.255
E 0/0	192.168.12.2	255.255.255.0
E 0/1	192.168.23.2	255.255.255.0
E 0/2	192.168.26.2	255.255.255.0
E 0/3	192.168.27.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.0.0.0
Loopback 10	172.16.1.3	255.255.255.255
E 0/0	192.168.23.3	255.255.255.0
E 0/1	192.168.13.3	255.255.255.0
E 0/2	192.168.38.3	255.255.255.0
E 0/3	192.168.39.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.0.0.0
Loopback 10	172.16.1.4	255.255.255.255
E 0/0	192.168.14.4	255.255.255.0
E 0/1	192.168.45.4	255.255.255.0
E 0/2	192.1.40.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.0.0.0
Loopback 10	172.16.1.5	255.255.255.255
E 0/0	192.168.15.5	255.255.255.0
E 0/1	192.168.45.5	255.255.255.0
E 0/2	192.1.50.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.0.0.0
Loopback 10	172.16.1.6	255.255.255.255
E 0/0	192.168.26.6	255.255.255.0
E 0/1	192.168.67.6	255.255.255.0
E 0/2	192.1.60.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.0.0.0
Loopback 10	172.16.1.7	255.255.255.255
E 0/0	192.168.27.7	255.255.255.0
E 0/1	192.168.67.7	255.255.255.0
E 0/2	192.1.70.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.0.0.0
Loopback 10	172.16.1.8	255.255.255.255
E 0/0	192.168.38.8	255.255.255.0
E 0/1	192.168.89.8	255.255.255.0
E 0/2	192.1.80.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	9.9.9.9	255.0.0.0
Loopback 10	172.16.1.9	255.255.255.255
E 0/0	192.168.39.9	255.255.255.0
E 0/1	192.168.89.9	255.255.255.0
E 0/2	192.1.90.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	100.100.100.100	255.0.0.0
E 0/0	192.1.40.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	111.111.111.111	255.0.0.0
E 0/0	192.1.50.11	255.255.255.0

R12

Interface	IP Address	Subnet Mask
Loopback 0	112.112.112.112	255.0.0.0
E 0/0	192.1.60.12	255.255.255.0

R13

Interface	IP Address	Subnet Mask
Loopback 0	113.113.113.113	255.0.0.0
E 0/0	192.1.70.13	255.255.255.0

R14

Interface	IP Address	Subnet Mask
Loopback 0	114.114.114.114	255.0.0.0
E 0/0	192.1.80.14	255.255.255.0

R15

Interface	IP Address	Subnet Mask
Loopback 0	115.115.115.115	255.0.0.0
E 0/0	192.1.90.15	255.255.255.0

Task 1

Configure the underlay IGP as EIGRP in AS 1000 between R1, R2 & R3. These routers represent their respective Sub-AS's.

R1 router eigrp 1000 network 192.168.12.0 network 192.168.13.0 network 172.16.1.0 0.0.0.255	**R2** router eigrp 1000 network 192.168.12.0 network 192.168.23.0 network 172.16.1.0 0.0.0.255
R3 router eigrp 1000 network 192.168.13.0 network 192.168.23.0 network 172.16.1.0 0.0.0.255	

Task 2

Configure the underlay IGP as EIGRP in Sub-AS 65001 between R1, R4 & R5. Advertise the links with the Sub-AS and the Loopback 10 networks in EIGRP.

R1 router eigrp 65001 network 192.168.14.0 network 192.168.15.0 network 172.16.1.0 0.0.0.255	**R4** router eigrp 65001 network 192.168.14.0 network 192.168.45.0 network 172.16.1.0 0.0.0.255
R5 router eigrp 65001 network 192.168.15.0 network 192.168.45.0 network 172.16.1.0 0.0.0.255	

Task 3

Configure the underlay IGP as EIGRP in Sub-AS 65002 between R2, R6 & R7. Advertise the links with the Sub-AS and the Loopback 10 networks in EIGRP.

R1 router eigrp 65002 network 192.168.26.0 network 192.168.27.0 network 172.16.1.0 0.0.0.255	**R6** router eigrp 65002 network 192.168.26.0 network 192.168.67.0 network 172.16.1.0 0.0.0.255
R7 router eigrp 65002 network 192.168.27.0 network 192.168.67.0 network 172.16.1.0 0.0.0.255	

Task 4

Configure the underlay IGP as EIGRP in Sub-AS 65003 between R3, R8 & R9. Advertise the links with the Sub-AS and the Loopback 10 networks in EIGRP.

R3 router eigrp 65003 network 192.168.38.0 network 192.168.39.0 network 172.16.1.0 0.0.0.255	**R8** router eigrp 65003 network 192.168.38.0 network 192.168.89.0 network 172.16.1.0 0.0.0.255
R9 router eigrp 65003 network 192.168.39.0 network 192.168.89.0 network 172.16.1.0 0.0.0.255	

Task 5

Configure AS 65001 with iBGP. Configure R1 as the RR. Set the relationship based on Loopback10. The Confederation Identifier is 1000. R1 is peering up only with 65002 in its confederation.

R1 router bgp 65001 bgp confederation identifier 1000 bgp confederation peer 65002 network 1.0.0.0 neighbor IBGP peer-group neighbor IBGP remote-as 65001 neighbor IBGP update-source Loopback10 neighbor IBGP next-hop-self neighbor IBGP route-reflector-client neighbor 172.16.1.4 peer-group IBGP neighbor 172.16.1.5 peer-group IBGP	
R4 router bgp 65001 bgp confederation identifier 1000 network 4.0.0.0 neighbor 172.16.1.1 remote-as 65001 neighbor 172.16.1.1 update-source Lo10 neighbor 172.16.1.1 next-hop-self	**R5** router bgp 65001 bgp confederation identifier 1000 network 5.0.0.0 neighbor 172.16.1.1 remote-as 65001 neighbor 172.16.1.1 update-source Lo10 neighbor 172.16.1.1 next-hop-self

Task 6

Configure AS 65002 with iBGP. Configure R2 as the RR. Set the relationship based on Loopback10. The Confederation Identifier is 1000. R2 is peering up with 65001 & 65003 in its confederation.

<table>
<tr><td colspan="2">

R2

```
router bgp 65002
 bgp confederation identifier 1000
 bgp confederation peer 65001 65003
 network 2.0.0.0
 neighbor IBGP peer-group
 neighbor IBGP remote-as 65002
 neighbor IBGP update-source Loopback10
 neighbor IBGP next-hop-self
 neighbor IBGP route-reflector-client
 neighbor 172.16.1.6 peer-group IBGP
 neighbor 172.16.1.7 peer-group IBGP
```

</td></tr>
<tr><td>

R6

```
router bgp 65002
 bgp confederation identifier 1000
 network 6.0.0.0
 neighbor 172.16.1.2 remote-as 65002
 neighbor 172.16.1.2 update-source Lo10
 neighbor 172.16.1.2 next-hop-self
```

</td><td>

R7

```
router bgp 65002
 bgp confederation identifier 1000
 network 7.0.0.0
 neighbor 172.16.1.2 remote-as 65002
 neighbor 172.16.1.2 update-source Lo10
 neighbor 172.16.1.2 next-hop-self
```

</td></tr>
</table>

Task 7

Configure AS 65003 with iBGP. Configure R3 as the RR. Set the relationship based on Loopback10. The Confederation Identifier is 1000. R3 is peering up only with 65002 in its confederation.

R3

```
router bgp 65003
 bgp confederation identifier 1000
 bgp confederation peer 65002
 network 3.0.0.0
 neighbor IBGP peer-group
 neighbor IBGP remote-as 65003
 neighbor IBGP update-source Loopback10
 neighbor IBGP next-hop-self
 neighbor IBGP route-reflector-client
 neighbor 172.16.1.8 peer-group IBGP
 neighbor 172.16.1.9 peer-group IBGP
```

R8

```
router bgp 65003
 bgp confederation identifier 1000
 network 8.0.0.0
 neighbor 172.16.1.3 remote-as 65003
 neighbor 172.16.1.3 update-source Lo10
 neighbor 172.16.1.3 next-hop-self
```

R9

```
router bgp 65003
 bgp confederation identifier 1000
 network 9.0.0.0
 neighbor 172.16.1.3 remote-as 65003
 neighbor 172.16.1.3 update-source Lo10
 neighbor 172.16.1.3 next-hop-self
```

Task 8

Configure eBGP neighbor relationships with Remote-AS's (10,11,12,13 & 14). Use the appropriate ASBR to configure the relationship. Have the Remote AS's advertise the Loopback0 interface network.

R4 router bgp 65001 neighbor 192.1.40.10 remote-as 10	**R10** router bgp 10 network 100.0.0.0 neighbor 192.1.40.4 remote-as 1000
R5 router bgp 65001 neighbor 192.1.50.11 remote-as 11	**R11** router bgp 11 network 111.0.0.0 neighbor 192.1.50.5 remote-as 1000
R6 router bgp 65002 neighbor 192.1.60.12 remote-as 12	**R12** router bgp 12 network 112.0.0.0 neighbor 192.1.60.6 remote-as 1000
R7 router bgp 65002 neighbor 192.1.70.13 remote-as 13	**R13** router bgp 13 network 113.0.0.0 neighbor 192.1.70.7 remote-as 1000
R8 router bgp 65003 neighbor 192.1.80.14 remote-as 14	**R14** router bgp 14 network 114.0.0.0 neighbor 192.1.80.8 remote-as 1000
R9 router bgp 65003 neighbor 192.1.90.15 remote-as 15	**R15** router bgp 15 network 115.0.0.0 neighbor 192.1.90.9 remote-as 1000

Task 9

Configure eBGP neighbor relationships between the Confederation Peers. (R1-R2) & (R2-R3). These are eBGP neighbor relationships that are on Loopbacks. Make sure to allow the ebgp-multihop.

R1

```
router bgp 65001
 neighbor 172.16.1.2 remote-as 65002
 neighbor 172.16.1.2 update-source Loopback10
 neighbor 172.16.1.2 next-hop-self
 neighbor 172.16.1.2 ebgp-multihop
```

R2

```
router bgp 65002
 neighbor 172.16.1.1 remote-as 65001
 neighbor 172.16.1.1 update-source Loopback10
 neighbor 172.16.1.1 next-hop-self
 neighbor 172.16.1.1 ebgp-multihop
 neighbor 172.16.1.3 remote-as 65003
 neighbor 172.16.1.3 update-source Loopback10
 neighbor 172.16.1.3 next-hop-self
 neighbor 172.16.1.3 ebgp-multihop
```

R3

```
router bgp 65003
 neighbor 172.16.1.2 remote-as 65002
 neighbor 172.16.1.2 update-source Loopback10
 neighbor 172.16.1.2 next-hop-self
 neighbor 172.16.1.2 ebgp-multihop
```

Verification:

Use Ping to verify end-to-end reachability between AS's 10,11,12,13 & 14 via AS 1000.

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring IPv6

Lab 1 – Configuring IPv6 Addressing

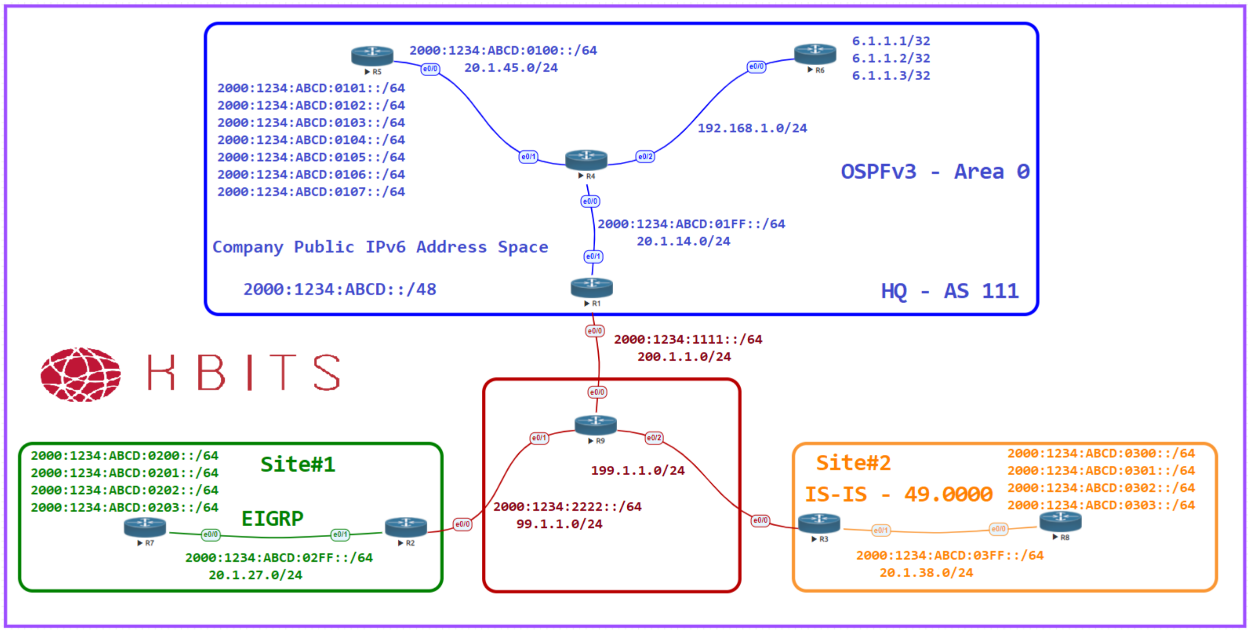

Task 1

Configure Headquarters with IPv6 addressing based on the Network Diagram. The Network between R4 & R6 will remain as IPv4 only. Configure the rest of the routers with IPv6 addressing based on the Network Diagram. Configure a default route on the Edge Router (R1) towards the ISP.

R1

```
ipv6 unicast-routing
!
Interface E 0/0
 ipv6 address 2000:1234:1111::1/64
 no shut
!
Interface E 0/1
 ipv6 address 2000:1234:ABCD:01FF::1/64
 no shut
!
ipv6 route ::/0 2000:1234:1111::9
```

R4

```
ipv6 unicast-routing
!
Interface E 0/0
 ipv6 address 2000:1234:ABCD:01FF::4/64
 no shut
!
```

```
Interface E 0/1
 ipv6 address 2000:1234:ABCD:0100::4/64
 no shut
```

R5

```
ipv6 unicast-routing
!
Interface E 0/0
 ipv6 address 2000:1234:ABCD:0100::5/64
 no shut
!
Interface Loopback1
 ipv6 address 2000:1234:ABCD:0101::5/64
!
Interface Loopback2
 ipv6 address 2000:1234:ABCD:0102::5/64
!
Interface Loopback3
 ipv6 address 2000:1234:ABCD:0103::5/64
!
Interface Loopback4
 ipv6 address 2000:1234:ABCD:0104::5/64
!
Interface Loopback5
 ipv6 address 2000:1234:ABCD:0105::5/64
!
Interface Loopback6
 ipv6 address 2000:1234:ABCD:0106::5/64
!
Interface Loopback7
 ipv6 address 2000:1234:ABCD:0107::5/64
```

Task 2

Configure **Site#1** with IPv6 addressing based on the Network Diagram. Configure a default route on the Edge Router (R2) towards the ISP.

R2

```
Interface E 0/0
 ipv6 address 2000:1234:2222::2/64
 no shut
!
Interface E 0/1
 ipv6 address 2000:1234:ABCD:02FF::2/64
 no shut
!
```

ipv6 route ::/0 2000:1234:2222::9
R7 ipv6 unicast-routing ! Interface E 0/0 ipv6 address 2000:1234:ABCD:02FF::7/64 no shut ! Interface Loopback1 ipv6 address 2000:1234:ABCD:0200::7/64 ! Interface Loopback2 ipv6 address 2000:1234:ABCD:0201::7/64 ! Interface Loopback3 ipv6 address 2000:1234:ABCD:0202::7/64 ! Interface Loopback4 ipv6 address 2000:1234:ABCD:0203::7/64

Task 3
Configure **Site#2** with IPv6 addressing based on the Network Diagram.

R3 ipv6 unicast-routing ! Interface E 0/1 ipv6 address 2000:1234:ABCD:03FF::3/64 no shut
R8 ipv6 unicast-routing ! Interface E 0/0 ipv6 address 2000:1234:ABCD:03FF::8/64 no shut ! Interface Loopback1 ipv6 address 2000:1234:ABCD:0300::8/64 ! Interface Loopback2 ipv6 address 2000:1234:ABCD:0301::8/64 ! Interface Loopback3

```
 ipv6 address 2000:1234:ABCD:0302::8/64
!
Interface Loopback4
 ipv6 address 2000:1234:ABCD:0303::8/64
```

Task 4

Configure IPv4 IP Addresses based on the network diagram. Configure Static Routing to provide full reachability for IPv4 networks. You are allowed to use static routes.

R1

```
Interface E 0/0
 Ip address 200.1.1.1 255.255.255.0
 No shut
!
Interface E 0/1
 Ip address 20.1.14.1 255.255.255.0
 No shut
!
Ip route 0.0.0.0 0.0.0.0 200.1.1.9
Ip route 20.1.45.0 255.255.255.0 20.1.14.4
Ip route 6.1.1.0  255.255.255.0 20.1.14.4
```

R2

```
Interface E 0/0
 Ip address 99.1.1.2 255.255.255.0
 No shut
!
Interface E 0/1
 Ip address 20.1.27.2 255.255.255.0
 No shut
!
Ip route 0.0.0.0 0.0.0.0 99.1.1.9
```

R3

```
Interface E 0/0
 Ip address 199.1.1.3 255.255.255.0
 No shut
!
Interface E 0/1
 Ip address 20.1.38.3 255.255.255.0
 No shut
!
```

Ip route 0.0.0.0 0.0.0.0 199.1.1.9
R4 Interface E 0/0 Ip address 20.1.14.4 255.255.255.0 No shut ! Interface E 0/1 Ip address 20.1.45.4 255.255.255.0 No shut ! Interface E 0/2 Ip address 192.168.1.4 255.255.255.0 No shut ! Ip route 0.0.0.0 0.0.0.0 20.1.14.1 Ip route 6.1.1.0 255.255.255.0 192.168.1.6
R5 Interface E 0/0 Ip address 20.1.45.5 255.255.255.0 No shut ! Ip route 0.0.0.0 0.0.0.0 20.1.45.4
R6 Interface E 0/0 Ip address 192.168.1.6 255.255.255.0 No shut ! Interface Loo1 Ip address 6.1.1.1 255.255.255.255 ! Interface Loo2 Ip address 6.1.1.2 255.255.255.255 ! Interface Loo3 Ip address 6.1.1.3 255.255.255.255 ! Ip route 0.0.0.0 0.0.0.0 192.168.1.4 ! Line vty 0 4 Password cisco Login

Transport input all
R7 Interface E 0/0 Ip address 20.1.27.7 255.255.255.0 No shut ! Ip route 0.0.0.0 0.0.0.0 20.1.27.2
R8 Interface E 0/0 Ip address 20.1.38.8 255.255.255.0 No shut ! Ip route 0.0.0.0 0.0.0.0 20.1.38.3
R9 Interface E 0/0 Ip address 200.1.1.9 255.255.255.0 No shut ! Interface E 0/1 Ip address 99.1.1.9 255.255.255.0 No shut ! Interface E 0/2 Ip address 199.1.1.9 255.255.255.0 No shut ! Ip route 20.1.14.0 255.255.255.0 200.1.1.1 Ip route 20.1.45.0 255.255.255.0 200.1.1.1 Ip route 6.1.1.0 255.255.255.0 200.1.1.1 Ip route 20.1.27.0 255.255.255.0 99.1.1.2 Ip route 20.1.38.0 255.255.255.0 199.1.1.3

Lab 2 – Configuring OSPFv3

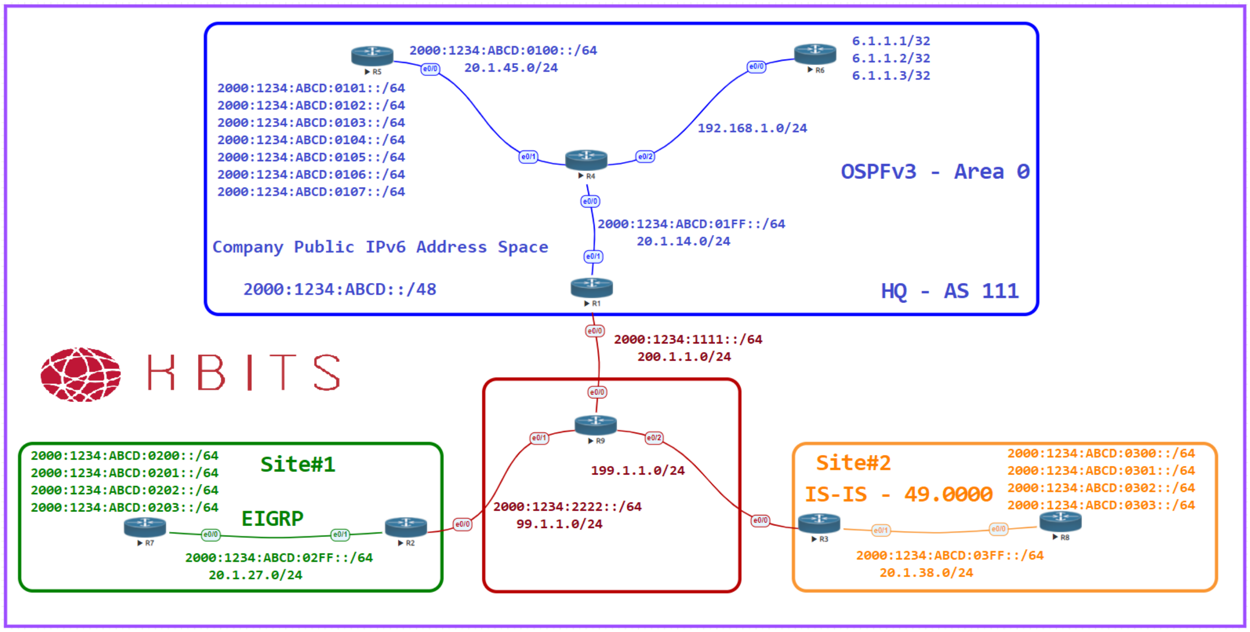

Task 1

Configure Headquarters with OSPFv3 within the HQ Site. Use X.X.X.X. as the router-id. (X stands for the Router #). Enable all the IPv6 addresses within the HQ site in OSPF. Have R1 inject a default route towards R4. The loopback interfaces should appear in the routing table using the interface mask.

R1

```
ipv6 router ospf 1
 router-id 1.1.1.1
 default-information originate always
!
Interface E 0/1
 ipv6 ospf 1 area 0
```

R4

```
ipv6 router ospf 1
 router-id 4.4.4.4
!
Interface E 0/0
 ipv6 ospf 1 area 0
!
Interface E 0/1
 ipv6 ospf 1 area 0
```

R5

```
ipv6 router ospf 1
```

```
 router-id 5.5.5.5
!
Interface E 0/0
 ipv6 ospf 1 area 0
!
Interface Loopback 1
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
!
Interface Loopback 2
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
!
Interface Loopback 3
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
!
Interface Loopback 4
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
!
Interface Loopback 5
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
!
Interface Loopback 6
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
!
Interface Loopback 7
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
```

Lab 3 – Configuring EIGRP for IPv6

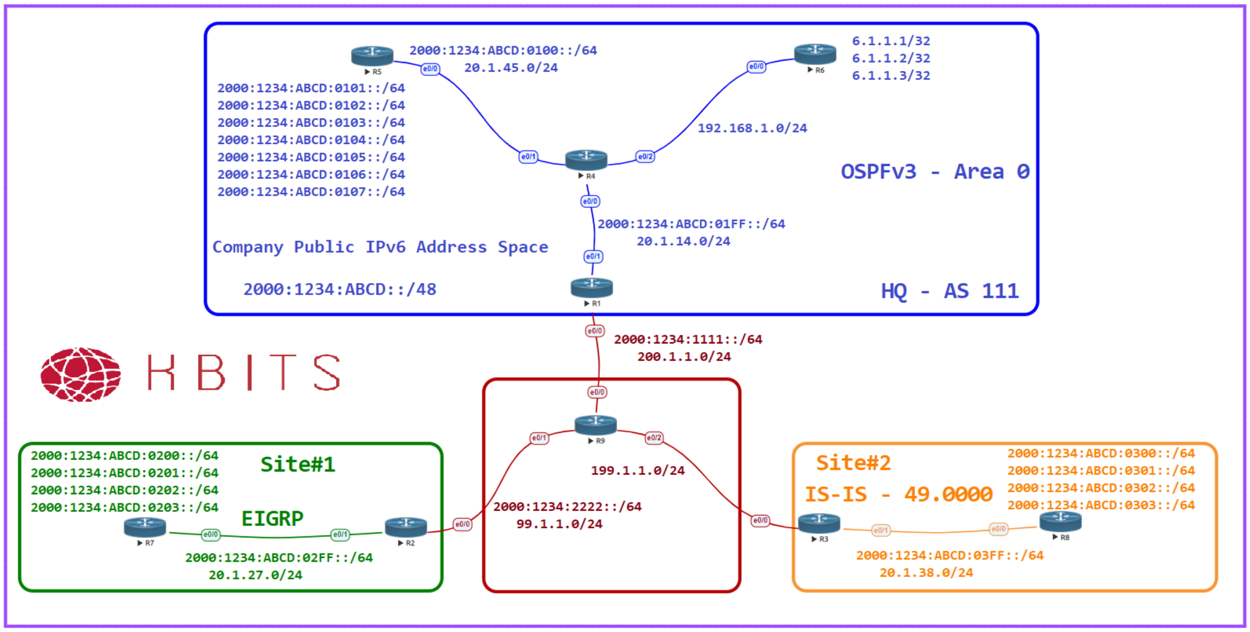

Task 1

Configure EIGRP 222 within Site#1. Use X.X.X.X. as the router-id. (X stands for the Router #). Enable all the IPv6 addresses within Site#1 in EIGRP. Configure a default route on R7 towards R2.

R2

```
ipv6 router eigrp 222
 router-id 2.2.2.2
!
Interface E 0/1
 ipv6 eigrp 222
```

R7

```
ipv6 router eigrp 222
 router-id 7.7.7.7
!
Interface E 0/0
 ipv6 eigrp 222
!
Interface Loopback 1
 ipv6 eigrp 222
!
Interface Loopback 2
```

```
 ipv6 eigrp 222
!
Interface Loopback 3
 ipv6 eigrp 222
!
Interface Loopback 4
 ipv6 eigrp 222
!
Ipv6 route ::/0 2000:1234:ABCD:02FF::2
```

Lab 4 – Configuring IS-IS for IPv6

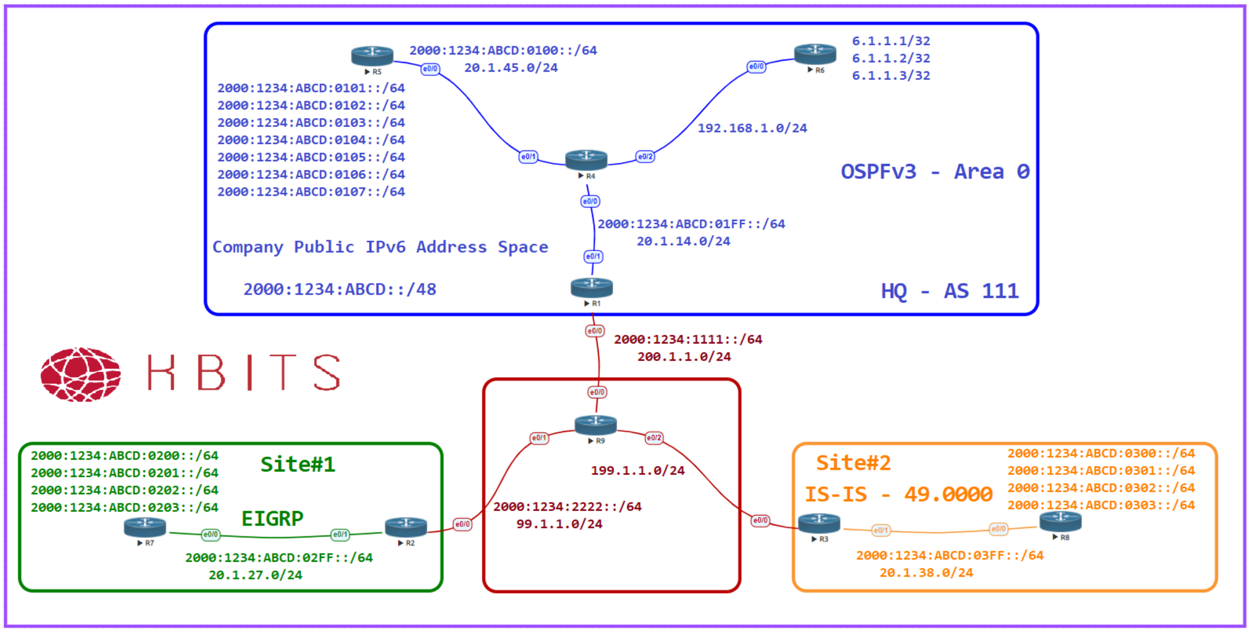

Task 1

Configure IS-IS within Site#1 based on the diagram. Use XXXX.XXXX.XXXX. as the System-id. (X stands for the Router #). Enable all the IPv6 addresses within Site#1 in IS-IS. Configure the Routers as Level-2 Routers with a metric-style of wide. Configure a default route on R7 towards R2.

R3

```
router isis
 net 49.0000.3333.3333.3333.00
 is-type level-2-only
 metric-style wide
!
address-family ipv6
  multi-topology
!
Interface E 0/1
 Ipv6 router isis
```

R8

```
router isis
 net 49.0000.8888.8888.8888.00
 is-type level-2-only
```

```
 metric-style wide
!
 address-family ipv6
  multi-topology
!
Interface E 0/0
 Ipv6 router isis
!
Interface Loopback 1
Ipv6 router isis
!
Interface Loopback 2
 Ipv6 router isis
!
Interface Loopback 3
 Ipv6 router isis
!
Interface Loopback 4
 Ipv6 router isis
!
Ipv6 route ::/0 2000:1234:ABCD:03FF::3
```

Lab 5 – Configuring BGP for IPv6

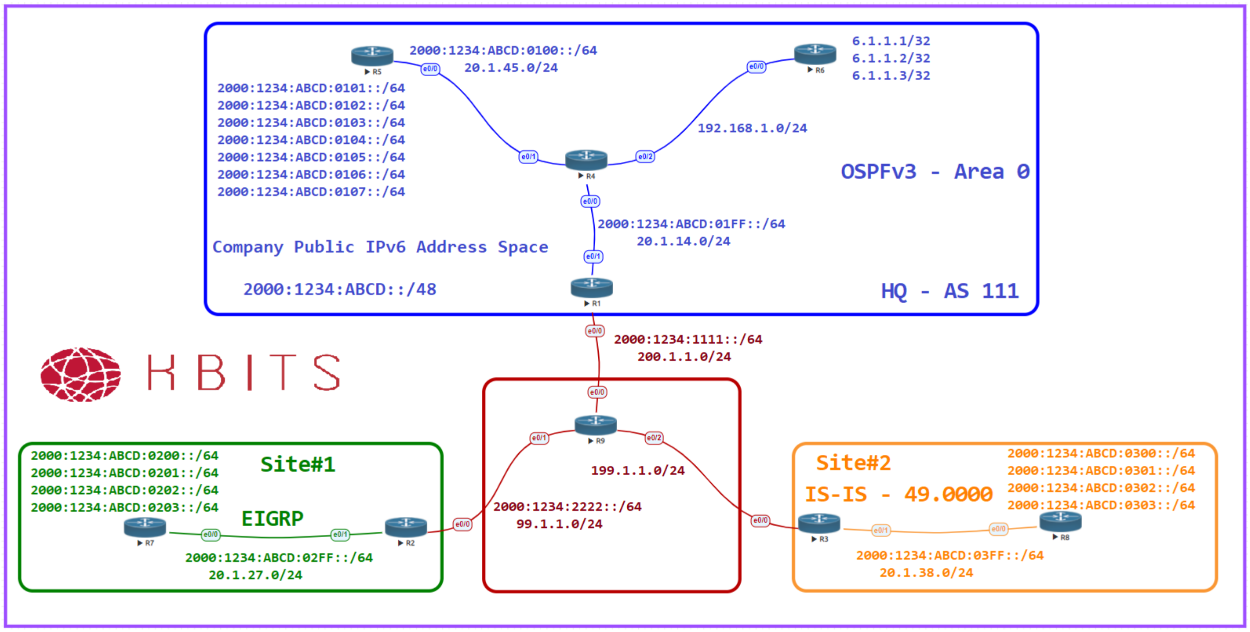

Task 1

Configure BGP between R1 & R9. Configure R1 in AS 111. Redistribute the internal networks to BGP and vice versa.

R1

```
router bgp 111
 neighbor 2000:1234:1111::9 remote-as 1000
 address-family ipv6
  neighbor 2000:1234:1111::9 activate
  redistribute ospf 1
!
ipv6 router ospf 1
 redistribute bgp 111
```

R9

```
router bgp 1000
 neighbor 2000:1234:1111::1 remote-as 111
 address-family ipv6
  neighbor 2000:1234:1111::1 activate
```

Task 2

Configure BGP between R2 & R9. Configure R2 in AS 222. Redistribute the internal networks to BGP and vice versa.

R2

```
router bgp 111
 neighbor 2000:1234:2222::9 remote-as 1000
 address-family ipv6
  neighbor 2000:1234:2222::9 activate
  redistribute eigrp 222
!
ipv6 router eigrp 222
 redistribute bgp 222 metric 10 10 10 10 10
```

R9

```
router bgp 1000
 neighbor 2000:1234:2222::2 remote-as 222
 address-family ipv6
  neighbor 2000:1234:2222::2 activate
```

Lab 6 – Configuring IPv6IP Tunneling

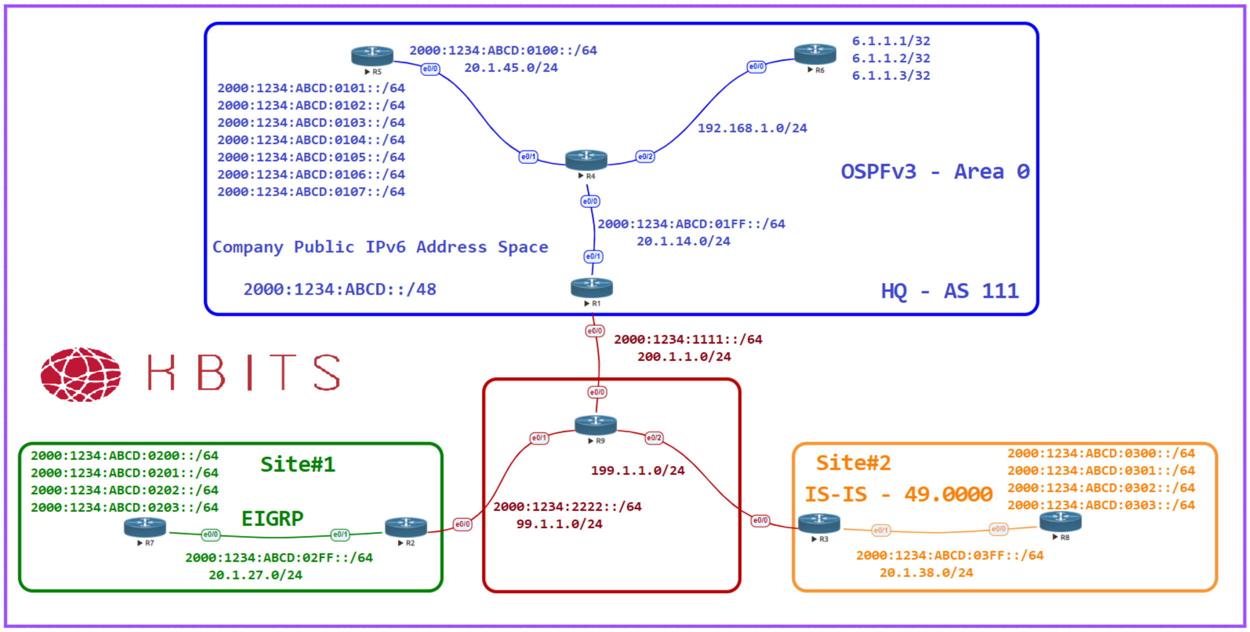

Task 1

Configure a IPv6IP tunnel to connect R1 to R3. Use the 2000:1234:ABCD:01FE::/64 as the Tunnel Network. Enable the Tunnel Interface in OSPF.

R1

```
Interface tunnel 1
 tunnel source 200.1.1.1
 tunnel destination 199.1.1.3
 tunnel mode ipv6ip
 ipv6 address 2000:1234:ABCD:01FE::1/64
 ipv6 ospf 1 area 0
```

R3

```
Interface tunnel 1
 tunnel source 199.1.1.3
 tunnel destination 200.1.1.1
 tunnel mode ipv6ip
 ipv6 address 2000:1234:ABCD:01FE::3/64
 ipv6 ospf 1 area 0
```

Task 2
Configure route redistribution on R3 between OSPF and IS-IS.

```
R3

Ipv6 router ospf 1
 Redistribute isis
!
Router isis
 Address-family ipv6 unicast
  Redistribute ospf 1
```

Task 3
Configure route redistribution on R1 between OSPF and BGP for external OSPF routes as well.

```
R1

Router bgp 111
 Address-family ipv6
  Redistribute ospf 1 match internal external
```

Lab 7 – Configuring NAT64 – R4

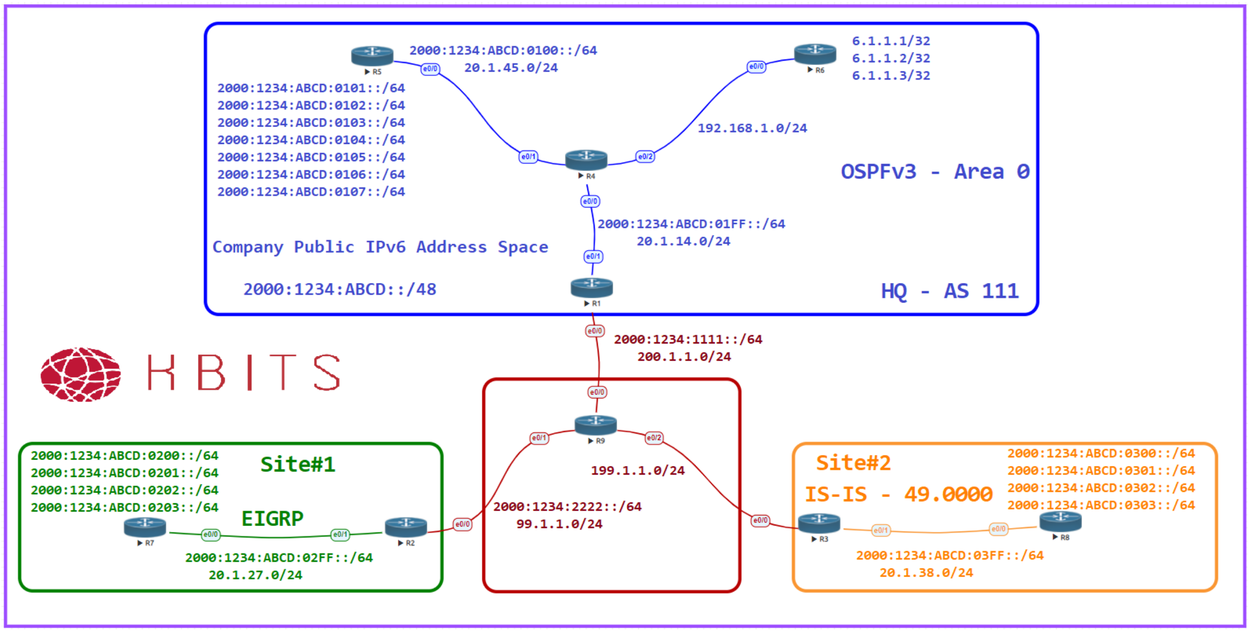

Task 1

Enable NAT64 on all Interfaces on R4.

R4

```
Interface E0/0
 nat64 enable
!
Interface E0/1
 nat64 enable
!
Interface E0/2
 nat64 enable
```

Task 2

Dedicate an IPv6 network prefix for NAT64

R4

```
nat64 prefix stateful 2000:1234:ABCD:0400::/64
```

Task 3

Inject the NAT64 into the IPv6 network by creating a Null 0 route for it and redistributing it into BGP. Allow this route to redistributed into BGP on R1.

```
R4

ipv6 route 2000:1234:ABCD:400::/64 Null0
ipv6 router ospf 1
 redistribute static
```

```
R1

Router bgp 111
 Address-family ipv6
  Redistribute ospf 1 match internal external
```

Task 4

Configure Static NAT for IPv4 Servers. Translate to the following:

- 6.1.1.1 – 2000:1234:ABCD:0400::1
- 6.1.1.2 – 2000:1234:ABCD:0400::2
- 6.1.1.2 – 2000:1234:ABCD:0400::3

```
R4

nat64 v4v6 static 6.1.1.1 2000:1234:ABCD:0400::1
nat64 v4v6 static 6.1.1.2 2000:1234:ABCD:0400::2
nat64 v4v6 static 6.1.1.3 2000:1234:ABCD:0400::3
```

Task 5

Configure Dynamic PAT for your networks (2000:1234:ABCD::/64 to a pool of 10.10.10.1 & 10.10.10.2.

```
R4

ipv6 access-list IPV6LIST
 permit ip 2000:1234:ABCD::/48 any
!
nat64 v4 pool V4POOL 10.10.10.1 10.10.10.2
!
nat64 v6v4 list IPV6LIST pool V4POOL overload
```

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring MPLS Unicast Routing

Lab 1 – Configuring MPLS Unicast Routing

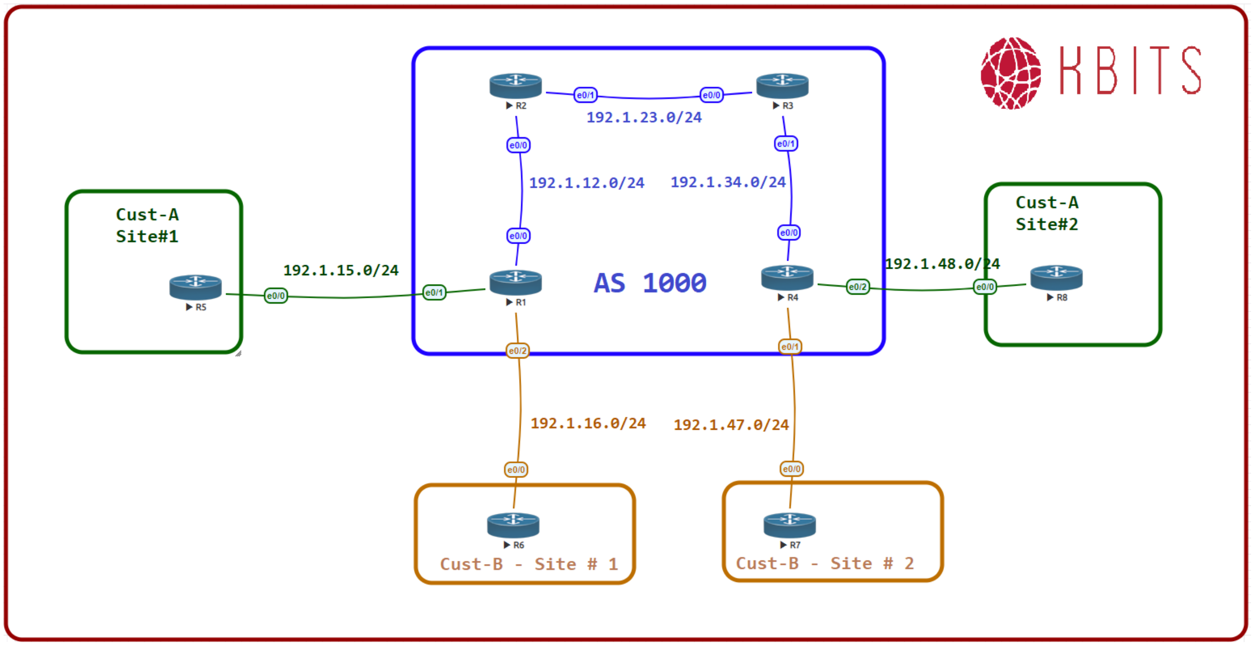

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
E 0/0	192.1.12.1	255.255.255.0
E 0/1	192.1.15.1	255.255.255.0
E 0/2	192.1.16.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
E 0/0	192.1.12.2	255.255.255.0
E 0/1	192.1.23.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
E 0/0	192.1.23.3	255.255.255.0
E 0/1	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
E 0/0	192.1.34.4	255.255.255.0
E 0/1	192.1.47.4	255.255.255.0
E 0/2	192.1.48.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	10.5.5.5	255.255.255.0
E 0/0	192.1.15.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	10.6.6.6	255.255.255.0
E 0/0	192.1.16.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	10.7.7.7	255.255.255.0
E 0/0	192.1.47.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	10.8.8.8	255.255.255.0
S 0/0	192.1.48.8	255.255.255.0

Task 1

Configure OSPF between all the SP routers (R1, R2, R3, R4). Use x.x.x.x as the router-id, where x is the Router number. Advertise all Internal links in OSPF in area 0.

R1 Router ospf 1 Router-id 1.1.1.1 Network 1.1.1.1 0.0.0.0 area 0 Network 192.1.12.0 0.0.0.255 area 0	**R2** Router ospf 1 Router-id 2.2.2.2 Network 2.2.2.2 0.0.0.0 area 0 Network 192.1.12.0 0.0.0.255 area 0 Network 192.1.23.0 0.0.0.255 area 0
R3 Router ospf 1 Router-id 3.3.3.3 Network 3.3.3.3 0.0.0.0 area 0 Network 192.1.23.0 0.0.0.255 area 0 Network 192.1.34.0 0.0.0.255 area 0	**R4** Router ospf 1 Router-id 4.4.4.4 Network 4.4.4.4 0.0.0.0 area 0 Network 192.1.34.0 0.0.0.255 area 0

Task 2

Configure MPLS on all the physical links in the SP Network. Use LDP to distribute labels. The LDP neighbour relationships should be formed based on the most reliable interface. The Labels should be assigned from the range X00 – X99, where X is the router number.

R1 Mpls ldp router-id Loopback0 ! Mpls label range 100 199 ! Interface E 0/0 Mpls ip	**R2** Mpls ldp router-id Loopback0 ! Mpls label range 200 299 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip
R3 Mpls ldp router-id Loopback0 ! Mpls label range 300 399 ! Interface E 0/0	**R4** Mpls ldp router-id Loopback0 ! Mpls label range 400 499 ! Interface E 0/0

Mpls ip ! Interface E 0/1 Mpls ip	Mpls ip

Lab 2 – Authenticating LDP Peers

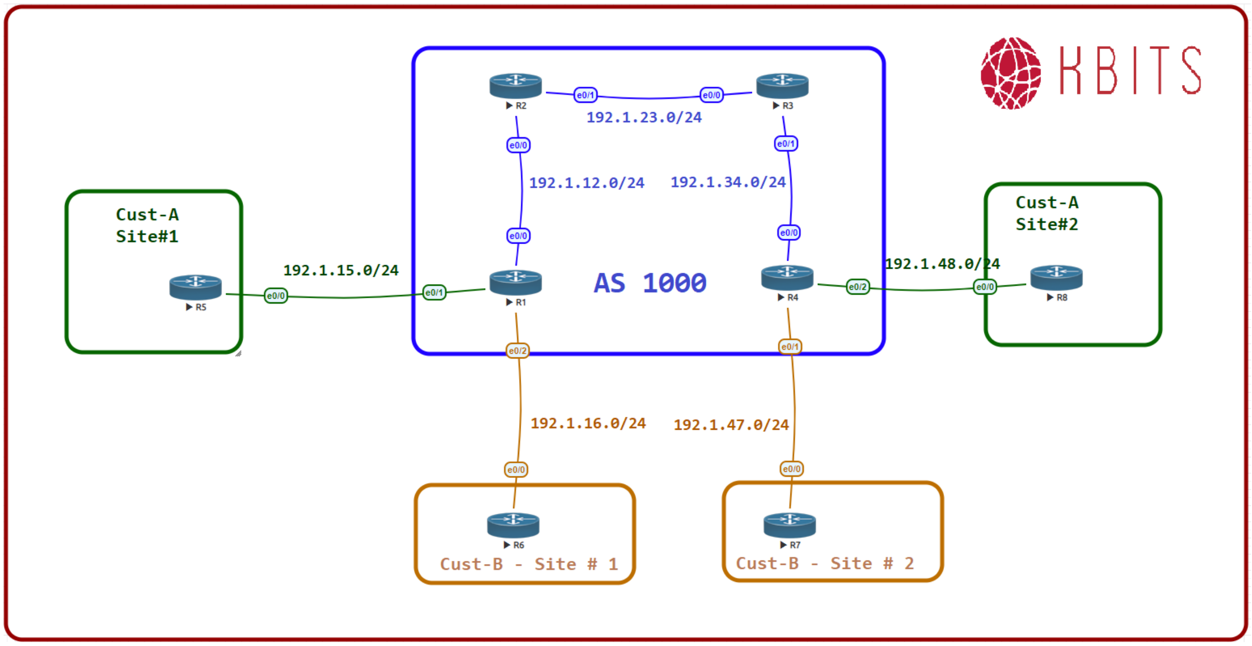

Task 1

All LDP neighbor relationships should be authenticated using a password of **ccie12353**.

R1 Mpls ldp neighbor 2.2.2.2 password ccie12353
R2 Mpls ldp neighbor 1.1.1.1 password ccie12353 Mpls ldp neighbor 3.3.3.3 password ccie12353
R3 Mpls ldp neighbor 2.2.2.2 password ccie12353 Mpls ldp neighbor 4.4.4.4 password ccie12353
R4 Mpls ldp neighbor 3.3.3.3 password ccie12353

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring Intra-AS MPLS VPNs

Lab 1 – Configuring MPLS VPN – PE-CE using Static Routing

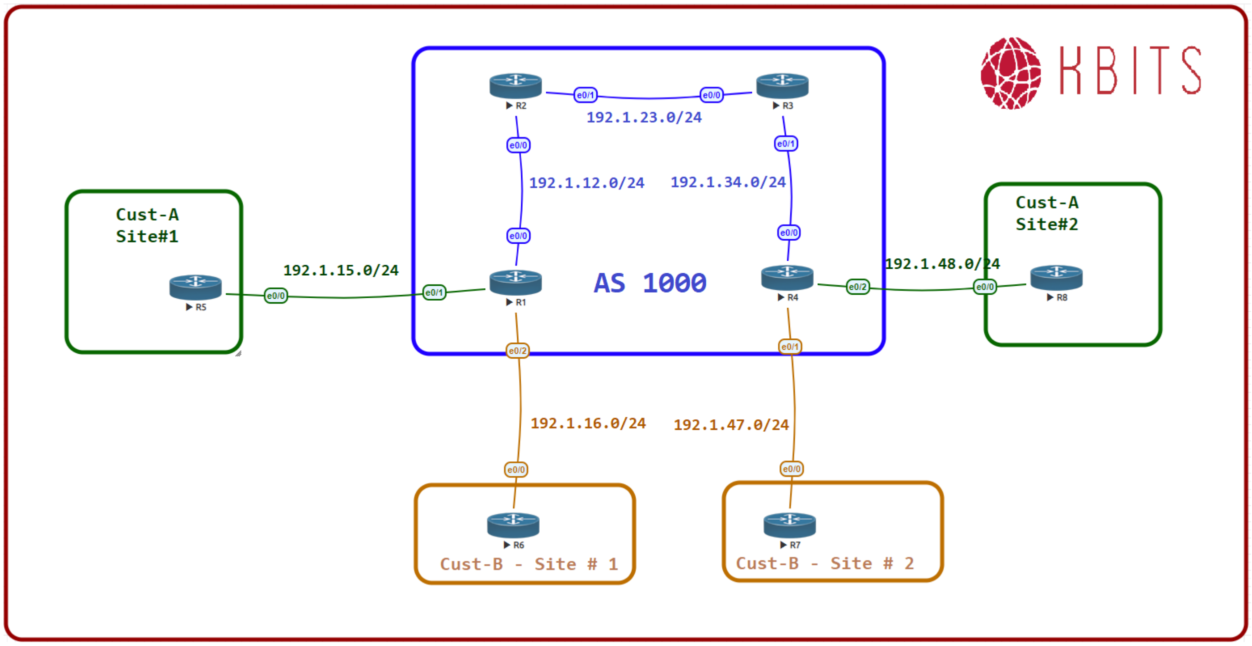

Note:
Save the Configs on all the routers. **Do not save the configs during the labs.** At the completion of this lab, **reload the routers without saving**. This will allow you to do the next lab based on the same topology.

Task 1

Configure a VPNv4 (MP-iBGP) neighbor relationship between R1 and R4.

R1 Router BGP 1000 Neighbor 4.4.4.4 remote-as 1000 Neighbor 4.4.4.4 update-source loopback0 ! Address-family vpnv4 Neighbor 4.4.4.4 activate
R4 Router BGP 1000 Neighbor 1.1.1.1 remote-as 1000 Neighbor 1.1.1.1 update-source loopback0 ! Address-family vpnv4 Neighbor 1.1.1.1 activate

Task 2

Configure a VRF **Cust-A** with a RD value of 1000:1 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R4.

R1
 	
Vrf definition Cust-A
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E 0/1
 vrf forwarding Cust-A
 Ip address 192.1.15.1 255.255.255.0
 No shut
	R4
	
	Vrf definition Cust-A
	 rd 1000:1
	 address-family ipv4
	  route-target both 1000:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.48.4 255.255.255.0
	 No shut

Task 3

Configure a static route on R1 in the Cust-A vrf to reach the 10.5.5.0 on R5. Inject this route into BGP such that it should be reachable from Cust-A VRF on R4. Configure a default Route on R5 towards R1.

R1

```
ip route vrf Cust-A 10.5.5.0 255.255.255.0 192.1.15.5
!
Router BGP 1000
!
Address-family ipv4 vrf Cust-A
 Redistribute static
```

R5

```
ip route 0.0.0.0 0.0.0.0 192.1.15.1
```

Task 4

Configure a static route on R4 in the Cust-A vrf to reach the 10.8.8.0 on R8. Inject this route into BGP such that it should be reachable from Cust-A VRF on R1. Configure a default Route on R8 towards R4.

R4

```
ip route vrf Cust-A 10.8.8.0 255.255.255.0 192.1.48.8
!
Router BGP 1000
!
Address-family ipv4 vrf Cust-A
 Redistribute static
```

R8

```
ip route 0.0.0.0 0.0.0.0 192.1.48.4
```

Task 5

Configure a VRF **Cust-B** with a RD value of 1000:2 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-B sites on R1 and R4.

R1

```
Vrf definition Cust-B
 rd 1000:2
 address-family ipv4
  route-target both 1000:2
!
Interface E 0/2
 Ip vrf forwarding Cust-B
 Ip address 192.1.16.1 255.255.255.0
 No shut
```

R4

```
Vrf definition Cust-B
 rd 1000:2
 address-family ipv4
  route-target both 1000:2
!
Interface E 0/1
 Ip vrf forwarding Cust-B
 Ip address 192.1.47.4 255.255.255.0
 No shut
```

Task 6

Configure a static route on R1 in the Cust-B vrf to reach the 10.6.6.0 on R6. Inject this route into BGP such that it should be reachable from Cust-B VRF on R4. Configure a default Route on R6 towards R1.

R1 ip route vrf Cust-B 10.6.6.0 255.255.255.0 192.1.16.6 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-B Redistribute static
R6 ip route 0.0.0.0 0.0.0.0 192.1.16.1

Task 7

Configure a static route on R4 in the CUST-B vrf to reach the 10.7.7.0 on R7. Inject this route into BGP such that it should be reachable from CUST-B VRF on R1. Configure a default Route on R7 towards R4.

R4 ip route vrf Cust-B 10.7.7.0 255.255.255.0 192.1.47.7 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-B Redistribute static
R7 ip route 0.0.0.0 0.0.0.0 192.1.47.4

NOTE:

Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 2 – Configuring MPLS VPN – PE-CE using EIGRP

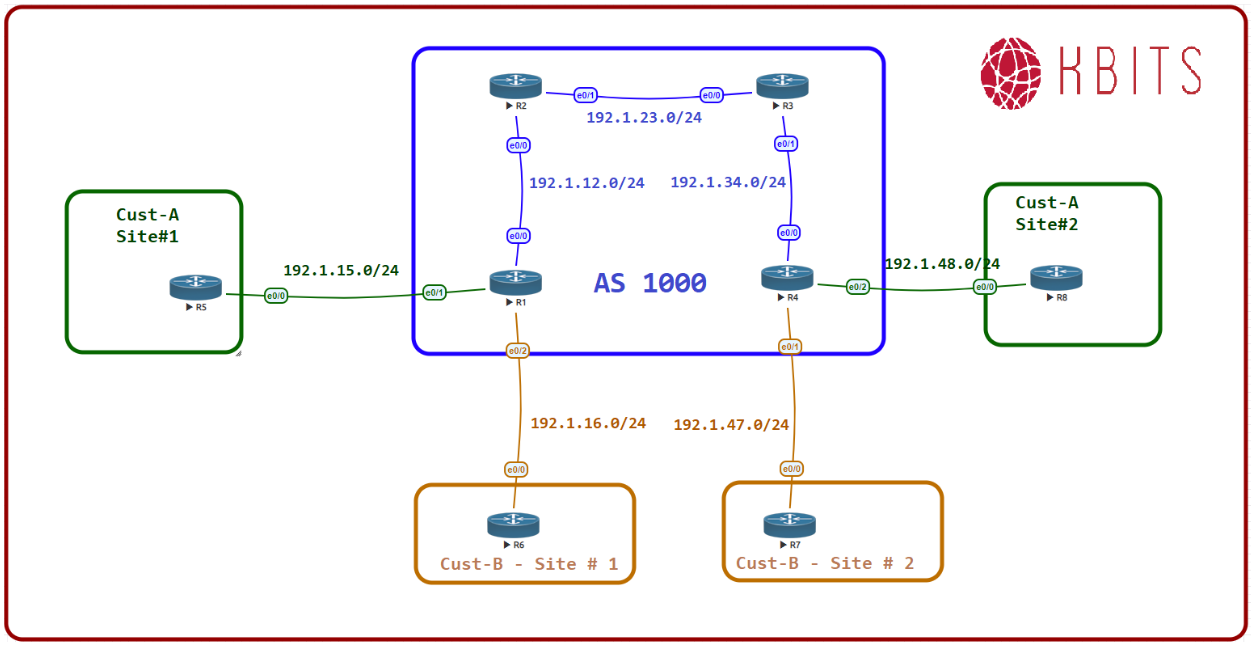

Note:

Save the Configs on all the routers. **Do not save the configs during the labs**. At the completion of this lab, **reload the routers without saving**. This will allow you to do the next lab based on the same topology.

Task 1

Configure a VPNv4 neighbor relationship between R1 and R4.

R1
Router BGP 1000 Neighbor 4.4.4.4 remote-as 1000 Neighbor 4.4.4.4 update-source loopback0 ! Address-family vpnv4 Neighbor 4.4.4.4 activate
R4
Router BGP 1000 Neighbor 1.1.1.1 remote-as 1000 Neighbor 1.1.1.1 update-source loopback0 ! Address-family vpnv4 Neighbor 1.1.1.1 activate

Task 2

Configure a VRF **Cust-A** with a RD value of 1000:1 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R4.

R1	R4

Vrf definition Cust-A
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E 0/1
 vrf forwarding Cust-A
 Ip address 192.1.15.1 255.255.255.0
 No shut
	Vrf definition Cust-A
	 rd 1000:1
	 address-family ipv4
	  route-target both 1000:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.48.4 255.255.255.0
	 No shut

Task 3

Configure EIGRP 100 as the Routing Protocol between R5 and R1-vrf Cust-A. Advertise all the routes on R5 in EIGRP. Advertise the VRF link in EIGRP on R1 under the appropriate address family. Make sure the VRF Cust-A on R4 has reachability to routes learned from R5.

R1

Router EIGRP 1 ! Address-family ipv4 vrf Cust-A Autonomous-system 100 Network 192.1.15.0 Redistribute BGP 1000 metric 10 10 10 10 10 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-A Redistribute eigrp 100
R5 Router EIGRP 100 Network 192.1.15.0 Network 10.0.0.0

Task 4

Configure EIGRP 100 as the Routing Protocol between R4 and R8-vrf CUST-A. Advertise all the routes on R8 in EIGRP. Advertise the VRF link in RIP on R4 under the appropriate address family. Make sure the VRF CUST-A on R1 has reachability to routes learned from R8.

R4
Router EIGRP 1 ! Address-family ipv4 vrf Cust-A Autonomous-system 100 Network 192.1.48.0 Redistribute BGP 1000 metric 10 10 10 10 10 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-A Redistribute eigrp 100
R8 Router EIGRP 100 Network 192.1.48.0 Network 10.0.0.0

Task 5

Configure a VRF **Cust-B** with a RD value of 1000:2 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-B sites on R1 and R4.

R1	**R4**

Vrf definition Cust-B
 rd 1000:2
 address-family ipv4
  route-target both 1000:2
!
Interface E 0/2
 Ip vrf forwarding Cust-B
 Ip address 192.1.16.1 255.255.255.0
 No shut
	Vrf definition Cust-B
	 rd 1000:2
	 address-family ipv4
	  route-target both 1000:2
	!
	Interface E 0/1
	 Ip vrf forwarding Cust-B
	 Ip address 192.1.47.4 255.255.255.0
	 No shut

Task 6

Configure EIGRP 200 as the Routing Protocol between R6 and R1-vrf Cust-B. Advertise all the routes on R6 in EIGRP 200. Advertise the VRF link in EIGRP on R1 under the appropriate address family. Make sure the VRF Cust-B on R4 has reachability to routes learned from R6.

R1
Router EIGRP 1 ! Address-family ipv4 vrf Cust-B Autonomous-system 200 Network 192.1.16.0 Redistribute BGP 100 metric 10 10 10 10 10 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-B Redistribute eigrp 200
R6
Router EIGRP 200 Network 192.1.16.0 Network 10.0.0.0

Task 7

Configure EIGRP 222 as the Routing Protocol between R7 and R4-vrf Cust-B. Advertise all the routes on R7 in EIGRP 222. Advertise the VRF link in EIGRP on R4 under the appropriate address family. Make sure the VRF Cust-B on R1 has reachability to routes learned from R7.

R4 Router EIGRP 1 ! Address-family ipv4 vrf Cust-B Autonomous-system 222 Network 192.1.47.0 Redistribute BGP 1000 metric 10 10 10 10 10 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-B Redistribute eigrp 2222
R7 Router EIGRP 222 Network 192.1.47.0 Network 10.0.0.0

NOTE:
Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 3 – Configuring MPLS VPN – PE-CE using IS-IS

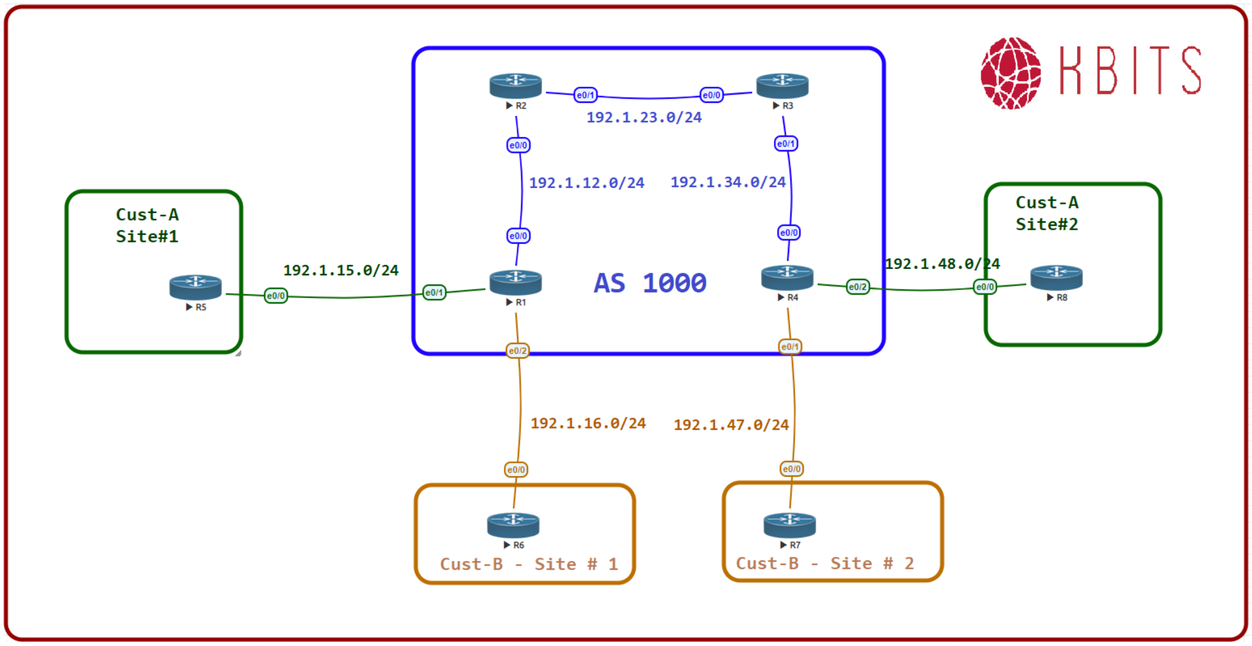

Note:
Save the Configs on all the routers. **Do not save the configs during the labs**. At the completion of this lab, **reload the routers without saving**. This will allow you to do the next lab based on the same topology.

Task 1

Configure a VPNv4 neighbor relationship between R1 and R4.

R1
Router BGP 1000 Neighbor 4.4.4.4 remote-as 1000 Neighbor 4.4.4.4 update-source loopback0 ! Address-family vpnv4 Neighbor 4.4.4.4 activate
R4 Router BGP 1000 Neighbor 1.1.1.1 remote-as 1000 Neighbor 1.1.1.1 update-source loopback0 ! Address-family vpnv4 Neighbor 1.1.1.1 activate

Task 2

Configure a VRF **Cust-A** with a RD value of 1000:1 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R4.

R1	**R4**

Vrf definition Cust-A
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E 0/1
 vrf forwarding Cust-A
 Ip address 192.1.15.1 255.255.255.0
 No shut
	Vrf definition Cust-A
	 rd 1000:1
	 address-family ipv4
	  route-target both 1000:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.48.4 255.255.255.0
	 No shut

Task 3

Configure IS-IS as the Routing Protocol between R5 and R1-vrf Cust-A. Advertise all the routes on R5 in IS-IS. Use 49.0000 as the area ID and XXXX.XXXX.XXXX as the system id, where x is the router #. Configure the IS-IS routers as Level-2. Use the Wide Metric-Style. Advertise the VRF link in IS-IS on R1 under the appropriate address family. Make sure the VRF Cust-A on R4 has reachability to routes learned from R5.

R1

```
Router isis 1
 Vrf Cust-A
 Net 49.0000.1111.1111.1111.00
 Is-type level-2
 Metric-style wide
 Redistribute bgp 1000
!
Interface E 0/1
 Ip router isis 1
!
Router BGP 1000
!
Address-family ipv4 vrf Cust-A
 Redistribute isis 1
```

R5

```
Router isis
 Net 49.0000.5555.5555.5555.00
 Is-type level-2
 Metric-style wide
!
Interface loopback0
 Ip router isis
!
Interface E 0/0
 Ip router isis
```

Task 4

Configure IS-IS as the Routing Protocol between R8 and R4-vrf Cust-A. Advertise all the routes on R8 in IS-IS. Use 49.0000 as the area ID and XXXX.XXXX.XXXX as the system id, where x is the router #. Configure the IS-IS routers as Level-2. Use the Wide Metric-Style. Advertise the VRF link in IS-IS on R4 under the appropriate address family. Make sure the VRF Cust-A on R1 has reachability to routes learned from R8.

R4
R4 Router isis 1 Vrf Cust-A Net 49.0000.4444.4444.4444.00 Is-type level-2 Metric-style wide Redistribute bgp 1000 ! Interface E 0/2 Ip router isis 1 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-A Redistribute isis 1
R8 Router isis Net 49.0000.5555.5555.5555.00 Is-type level-2 Metric-style wide ! Interface loopback0 Ip router isis ! Interface E 0/0 Ip router isis

Task 5

Configure a VRF **Cust-B** with a RD value of 1000:2 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-B sites on R1 and R4.

R1	**R4**

Vrf definition Cust-B
 rd 1000:2
 address-family ipv4
  route-target both 1000:2
!
Interface E 0/2
 Ip vrf forwarding Cust-B
	Vrf definition Cust-B
	 rd 1000:2
	 address-family ipv4
	  route-target both 1000:2
	!
	Interface E 0/1
	 Ip vrf forwarding Cust-B

Ip address 192.1.16.1 255.255.255.0 No shut	Ip address 192.1.47.4 255.255.255.0 No shut

Task 6

Configure IS-IS as the Routing Protocol between R6 and R1-vrf Cust-B. Advertise all the routes on R6 in IS-IS. Use 49.0000 as the area ID and XXXX.XXXX.XXXX as the system id, where x is the router #. Configure the IS-IS routers as Level-2. Use the Wide Metric-Style. Advertise the VRF link in IS-IS on R1 under the appropriate address family. Make sure the VRF Cust-B on R4 has reachability to routes learned from R6.

R1

```
Router isis 2
 Vrf Cust-B
 Net 49.0000.1111.1111.1111.00
 Is-type level-2
 Metric-style wide
 Redistribute bgp 1000
!
Interface E 0/2
 Ip router isis 2
!
Router BGP 1000
!
Address-family ipv4 vrf Cust-B
 Redistribute isis 2
```

R6

```
Router isis
 Net 49.0000.6666.6666.6666.00
 Is-type level-2
 Metric-style wide
!
Interface loopback0
 Ip router isis
!
Interface E 0/0
 Ip router isis
```

Task 7

Configure IS-IS as the Routing Protocol between R7 and R4-vrf Cust-B. Advertise all the routes on R7 in IS-IS. Use 49.0000 as the area ID and XXXX.XXXX.XXXX as the system id, where x is the router #. Configure the IS-IS routers as Level-2. Use the Wide Metric-Style. Advertise the VRF link in IS-IS on R4 under the appropriate address family. Make sure the VRF Cust-B on R1 has reachability to routes learned from R7.

R4

```
Router isis 2
 Vrf Cust-B
 Net 49.0000.4444.4444.4444.00
 Is-type level-2
 Metric-style wide
 Redistribute bgp 1000
!
Interface E 0/1
 Ip router isis 2
!
Router BGP 1000
!
Address-family ipv4 vrf Cust-B
 Redistribute isis 2
```

R7

```
Router isis
 Net 49.0000.7777.7777.7777.00
 Is-type level-2
 Metric-style wide
!
Interface loopback0
 Ip router isis
!
Interface E 0/0
 Ip router isis
```

NOTE:

Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 4 – Configuring MPLS VPN – PE-CE using BGP – 1

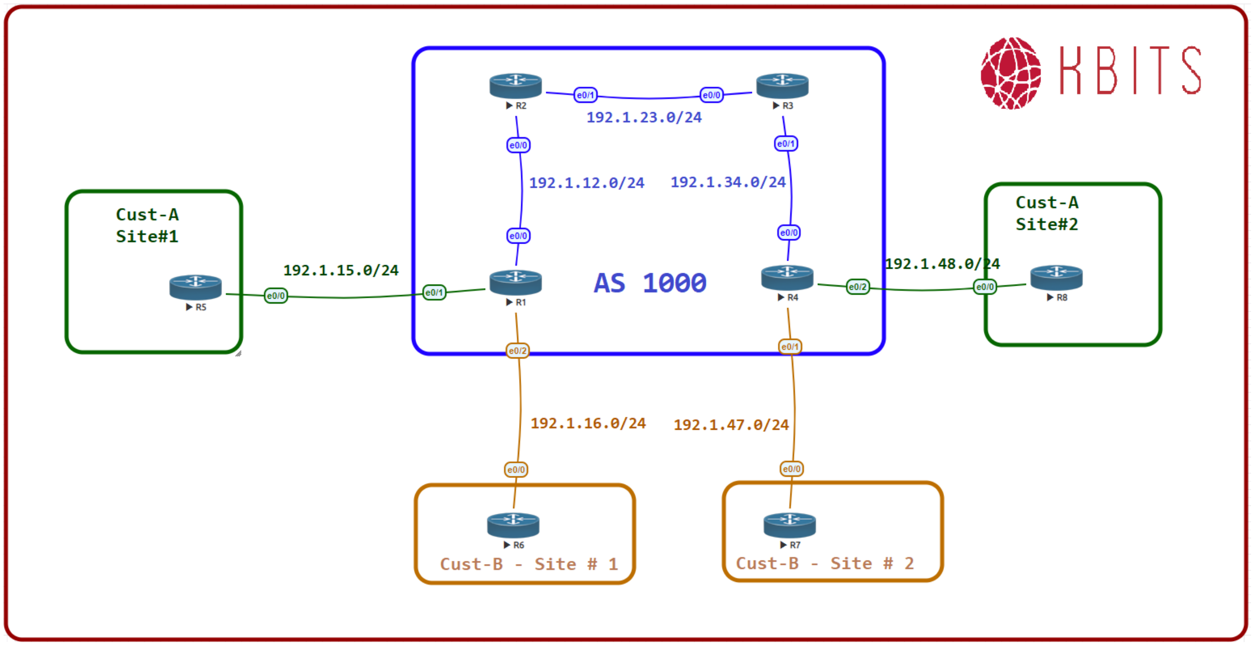

Note:

Save the Configs on all the routers. **Do not save the configs during the labs**. At the completion of this lab, **reload the routers without saving**. This will allow you to do the next lab based on the same topology.

Task 1

Configure a VPNv4 neighbor relationship between R1 and R4.

```
R1

Router BGP 1000
 Neighbor 4.4.4.4 remote-as 1000
 Neighbor 4.4.4.4 update-source loopback0
!
Address-family vpnv4
 Neighbor 4.4.4.4 activate
```

```
R4

Router BGP 1000
 Neighbor 1.1.1.1 remote-as 1000
 Neighbor 1.1.1.1 update-source loopback0
!
Address-family vpnv4
 Neighbor 1.1.1.1 activate
```

Task 2

Configure a VRF **Cust-A** with a RD value of 1000:1 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R4.

R1	R4

Vrf definition Cust-A
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E 0/1
 vrf forwarding Cust-A
 Ip address 192.1.15.1 255.255.255.0
 No shut
	Vrf definition Cust-A
	 rd 1000:1
	 address-family ipv4
	  route-target both 1000:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.48.4 255.255.255.0
	 No shut

Task 3

Configure BGP as the Routing Protocol between R5 and R1-vrf Cust-A. Advertise all the routes on R5 in BGP. Configure R5 with an AS # of 65005. Configure the BGP neighbor relationship on R1 for the Cust-A VRF. Make sure the VRF Cust-A on R4 has reachability to routes learned from R5.

R1 Router BGP 1000 ! Address-family ipv4 vrf Cust-A Neighbor 192.1.15.5 remote-as 65005
R5 Router bgp 65005 Network 10.5.5.0 mask 255.255.255.0 Neighbor 192.1.15.1 remote-as 1000

Task 4

Configure BGP as the Routing Protocol between R8 and R4-vrf Cust-A. Advertise all the routes on R8 in BGP. Configure R8 with an AS # of 65008. Configure the BGP neighbor relationship on R4 for the Cust-A VRF. Make sure the VRF Cust-A on R1 has reachability to routes learned from R8.

R4 Router BGP 1000 ! Address-family ipv4 vrf Cust-A Neighbor 192.1.48.8 remote-as 65008
R8 Router 65008 Network 10.8.8.0 mask 255.255.255.0 Neighbor 192.1.48.4 remote-as 1000

Task 5

Configure a VRF **Cust-B** with a RD value of 1000:2 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-B sites on R1 and R4.

R1	**R4**

Vrf definition Cust-B
 rd 1000:2
 address-family ipv4
  route-target both 1000:2
!
Interface E 0/2
 Ip vrf forwarding Cust-B
 Ip address 192.1.16.1 255.255.255.0
 No shut
	Vrf definition Cust-B
	 rd 1000:2
	 address-family ipv4
	  route-target both 1000:2
	!
	Interface E 0/1
	 Ip vrf forwarding Cust-B
	 Ip address 192.1.47.4 255.255.255.0
	 No shut

Task 6

Configure BGP as the Routing Protocol between R6 and R1-vrf Cust-B. Advertise all the routes on R6 in BGP. Configure R6 with an AS # of 65006. Configure the BGP neighbor relationship on R1 for the Cust-B VRF. Make sure the VRF Cust-B on R4 has reachability to routes learned from R6.

R1
Router BGP 1000 ! Address-family ipv4 vrf Cust-B Neighbor 192.1.16.6 remote-as 65006
R6
Router bgp 65006 Network 10.6.6.0 mask 255.255.255.0 Neighbor 192.1.16.1 remote-as 1000

Task 7

Configure BGP as the Routing Protocol between R7 and R4-vrf Cust-B. Advertise all the routes on R7 in BGP. Configure R7 with an AS # of 65007. Configure the BGP neighbor relationship on R4 for the Cust-B VRF. Make sure the VRF Cust-B on R1 has reachability to routes learned from R7.

R4
Router BGP 1000

! Address-family ipv4 vrf Cust-B Neighbor 192.1.47.7 remote-as 65007
R7 Router bgp 65007 Network 10.7.7.0 mask 255.255.255.0 Neighbor 192.1.47.4 remote-as 1000

NOTE:
Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 5 – Configuring MPLS VPN – PE-CE using BGP – 2

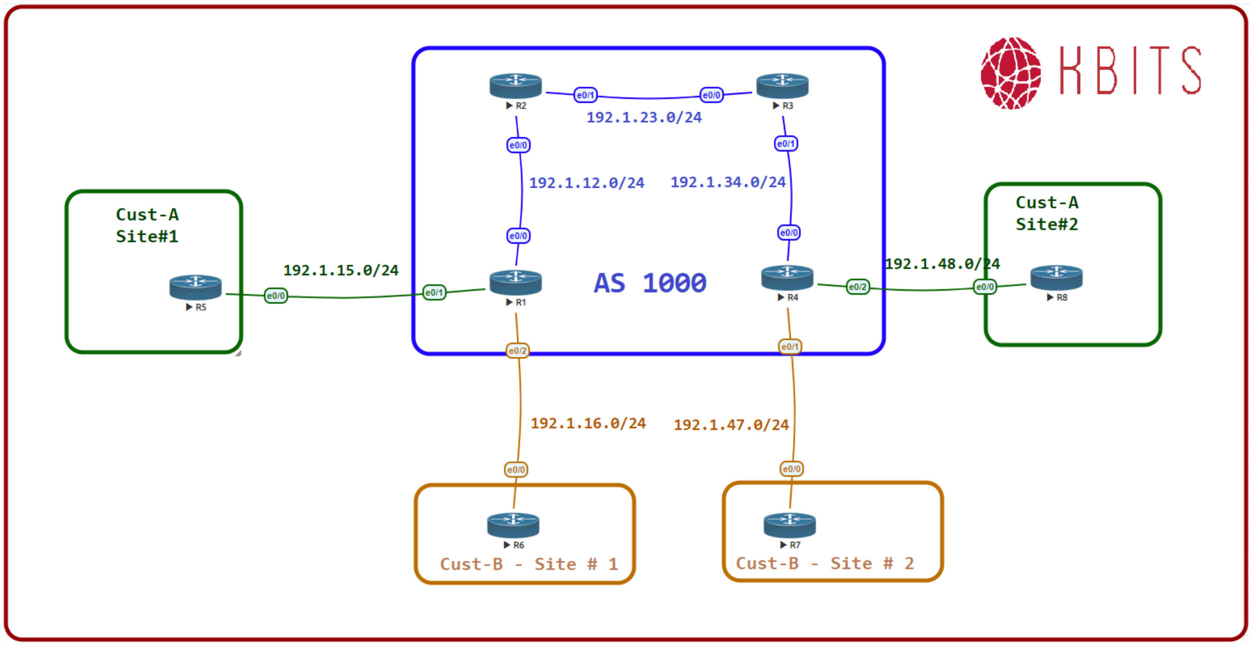

Note:

Save the Configs on all the routers. **Do not save the configs during the labs**. At the completion of this lab, **reload the routers without saving**. This will allow you to do the next lab based on the same topology.

Task 1

Configure a VPNv4 neighbor relationship between R1 and R4.

R1 Router BGP 1000 Neighbor 4.4.4.4 remote-as 1000 Neighbor 4.4.4.4 update-source loopback0 ! Address-family vpnv4 Neighbor 4.4.4.4 activate
R4 Router BGP 1000 Neighbor 1.1.1.1 remote-as 1000 Neighbor 1.1.1.1 update-source loopback0 ! Address-family vpnv4 Neighbor 1.1.1.1 activate

Task 2

Configure a VRF **Cust-A** with a RD value of 1000:1 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R4.

R1	**R4**

Vrf definition Cust-A
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E 0/1
 vrf forwarding Cust-A
 Ip address 192.1.15.1 255.255.255.0
 No shut
	Vrf definition Cust-A
	 rd 1000:1
	 address-family ipv4
	  route-target both 1000:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.48.4 255.255.255.0
	 No shut

Task 3

Configure BGP as the Routing Protocol between R5 and R1-vrf Cust-A. Advertise all the routes on R5 in BGP. Configure R5 with an AS # of 65001. Configure the BGP neighbor relationship on R1 for the Cust-A VRF. Make sure the VRF Cust-A on R4 has reachability to routes learned from R5.

R1 Router BGP 1000 ! Address-family ipv4 vrf Cust-A Neighbor 192.1.15.5 remote-as 65001
R5 Router bgp 65001 Network 10.5.5.0 mask 255.255.255.0 Neighbor 192.1.15.1 remote-as 1000

Task 4

Configure BGP as the Routing Protocol between R8 and R4-vrf Cust-A. Advertise all the routes on R8 in BGP. Configure R8 with an AS # of 65001. Configure the BGP neighbor relationship on R4 for the Cust-A VRF. Make sure the VRF Cust-A on R1 has reachability to routes learned from R8.

R4 Router BGP 1000 ! Address-family ipv4 vrf Cust-A Neighbor 192.1.48.8 remote-as 65001
R8 Router 65001 Network 10.8.8.0 mask 255.255.255.0 Neighbor 192.1.48.4 remote-as 1000

Task 5

Configure the PE's (R1 & R4) such that R5 routes are injected into R8's BGP table and vice versa.

R1 Router BGP 1000 ! Address-family ipv4 vrf Cust-A Neighbor 192.1.15.5 as-override
R4 Router BGP 1000 ! Address-family ipv4 vrf Cust-A Neighbor 192.1.48.8 as-override

Task 6

Configure a VRF **Cust-B** with a RD value of 1000:2 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-B sites on R1 and R4.

R1	**R4**
Vrf definition Cust-B rd 1000:2 address-family ipv4 route-target both 1000:2 ! Interface E 0/2 Ip vrf forwarding Cust-B Ip address 192.1.16.1 255.255.255.0 No shut	Vrf definition Cust-B rd 1000:2 address-family ipv4 route-target both 1000:2 ! Interface E 0/1 Ip vrf forwarding Cust-B Ip address 192.1.47.4 255.255.255.0 No shut

Task 7

Configure BGP as the Routing Protocol between R6 and R1-vrf Cust-B. Advertise all the routes on R6 in BGP. Configure R6 with an AS # of 65002. Configure the BGP neighbor relationship on R1 for the Cust-B VRF. Make sure the VRF Cust-B on R4 has reachability to routes learned from R6.

R1 Router BGP 1000 ! Address-family ipv4 vrf Cust-B Neighbor 192.1.16.6 remote-as 65002
R6 Router bgp 65002 Network 10.6.6.0 mask 255.255.255.0 Neighbor 192.1.16.1 remote-as 1000

Task 8

Configure BGP as the Routing Protocol between R7 and R4-vrf Cust-B. Advertise all the routes on R7 in BGP. Configure R7 with an AS # of 65002. Configure the BGP neighbor relationship on R4 for the Cust-B VRF. Make sure the VRF Cust-B on R1 has reachability to routes learned from R7.

R4 Router BGP 1000 ! Address-family ipv4 vrf Cust-B Neighbor 192.1.47.7 remote-as 65002
R7 Router bgp 65002 Network 10.7.7.0 mask 255.255.255.0 Neighbor 192.1.47.4 remote-as 1000

Task 9

Configure the CE's (R6 & R7) to allow routes from the remote site to be injected into BGP.

R6 Router BGP 65002 Neighbor 192.1.16.1 allowas-in
R7 Router BGP 65002 Neighbor 192.1.47.4 allowas-in

NOTE:

Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 6 – Configuring MPLS VPN – PE-CE using OSPF

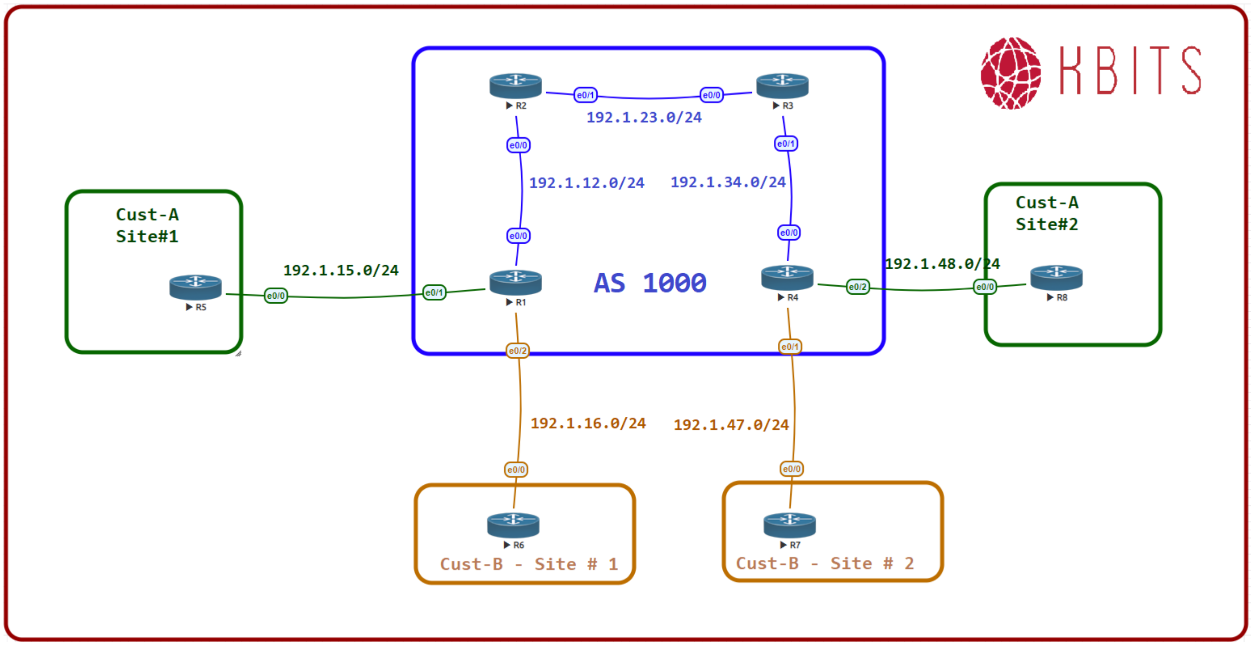

Note:
Save the Configs on all the routers. **Do not save the configs during the labs**. At the completion of this lab, **reload the routers without saving**. This will allow you to do the next lab based on the same topology.

Task 1

Configure a VPNv4 neighbor relationship between R1 and R4.

R1 Router BGP 1000 Neighbor 4.4.4.4 remote-as 1000 Neighbor 4.4.4.4 update-source loopback0 ! Address-family vpnv4 Neighbor 4.4.4.4 activate	
R4 Router BGP 1000 Neighbor 1.1.1.1 remote-as 1000 Neighbor 1.1.1.1 update-source loopback0 ! Address-family vpnv4 Neighbor 1.1.1.1 activate	

Task 2

Configure a VRF **Cust-A** with a RD value of 1000:1 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R4.

R1	**R4**

Vrf definition Cust-A
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E 0/1
 vrf forwarding Cust-A
 Ip address 192.1.15.1 255.255.255.0
 No shut
	Vrf definition Cust-A
	 rd 1000:1
	 address-family ipv4
	  route-target both 1000:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.48.4 255.255.255.0
	 No shut

Task 3

Configure OSPF as the PE-CE Routing protocol in Area 0 between R1 & R5. Advertise all networks on R5 in OSPF. Enable the R1-R5 link on R1 under the Cust-A VRF. Use OSPF process ID 58 on R1. Make sure the VRF Cust-A on R4 has reachability to routes learned from R5.

R1 Router ospf 58 vrf Cust-A Network 192.1.15.0 0.0.0.255 area 0 Redistribute bgp 1000 ! Router bgp 1000 Address-family ipv4 vrf Cust-A Redistribute ospf 58
R5 Router ospf 1 Network 10.5.5.0 0.0.0.255 area 0 Network 192.1.15.0 0.0.0.255 area 0

Task 4

Configure OSPF as the PE-CE Routing protocol in Area 0 between R4 & R8. Advertise all networks on R8 in OSPF. Enable the R4-R8 link on R4 under the Cust-A VRF. Use OSPF process ID 58 on R4. Make sure the VRF Cust-A on R1 has reachability to routes learned from R8.

R4 Router ospf 58 vrf Cust-A Network 192.1.48.0 0.0.0.255 area 0 Redistribute bgp 1000 ! Router bgp 1000 Address-family ipv4 vrf Cust-A Redistribute ospf 58
R8 Router ospf 1 Network 10.8.8.0 0.0.0.255 area 0 Network 192.1.48.0 0.0.0.255 area 0

Task 5

Configure a VRF **Cust-B** with a RD value of 1000:2 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-B sites on R1 and R4.

R1	**R4**
Vrf definition Cust-B rd 1000:2 route-target both 1000:2 ! Interface E 0/2 Ip vrf forwarding Cust-B Ip address 192.1.16.1 255.255.255.0 No shut	Vrf definition Cust-B rd 1000:2 route-target both 1000:2 ! Interface E 0/1 Ip vrf forwarding Cust-B Ip address 192.1.47.4 255.255.255.0 No shut

Task 6

Configure OSPF as the PE-CE Routing protocol in Area 0 between R1 & R6. Advertise all networks on R6 in OSPF. Enable the R1-R6 link on R1 under the Cust-B VRF. Use OSPF process ID 6 on R1. Make sure the VRF Cust-B on R4 has reachability to routes learned from R6.

R1
Router ospf 6 vrf Cust-B Network 192.1.16.0 0.0.0.255 area 0 Redistribute bgp 1000 ! Router bgp 1000 Address-family ipv4 vrf Cust-B Redistribute ospf 6
R6
Router ospf 1 Network 10.6.6.0 0.0.0.255 area 0 Network 192.1.16.0 0.0.0.255 area 0

Task 7

Configure OSPF as the PE-CE Routing protocol in Area 0 between R4 & R7. Advertise all networks on R7 in OSPF. Enable the R4-R7 link on R4 under the Cust-B VRF. Use OSPF process ID 7 on R4. Make sure the VRF Cust-B on R1 has reachability to routes learned from R7.

R4 Router ospf 7 vrf Cust-B Network 192.1.47.0 0.0.0.255 area 0 Redistribute bgp 1000 ! Router bgp 1000 Address-family ipv4 vrf Cust-B Redistribute ospf 7
R7 Router ospf 1 Network 10.7.7.0 0.0.0.255 area 0 Network 192.1.47.0 0.0.0.255 area 0

NOTE:

For the Cust-A VRF, the OSPF routes form the other site appears as O IA (Inter-Area) routes. This is since PE Routers are using the same process ID (58). The MPLS network is treated as the OSPF Super-Backbone.

For the Cust-B VRF, the OSPF routes form the other site appears as O E2 (External) routes. This is since PE Routers are using different Process ID for the Address Family OSPF process.

Lab 7 – Configuring MPLS VPN – PE-CE using OSPF – Domain-ID

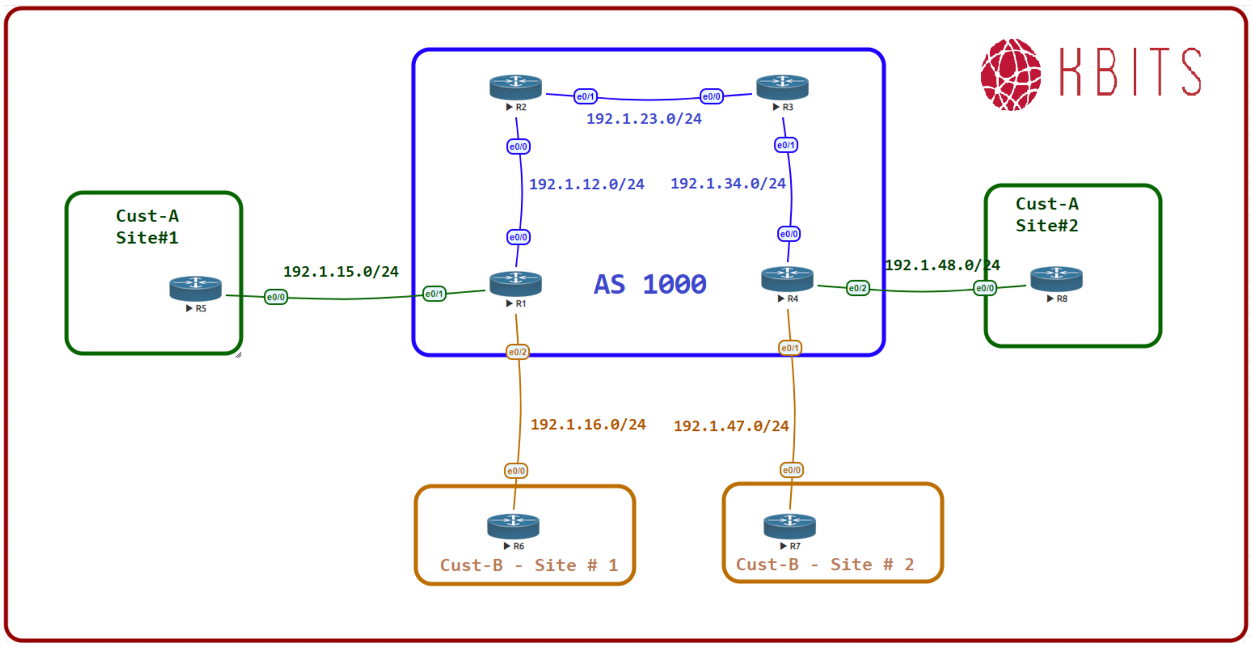

Task 1

Configure a Domain-id under OSPF for Cust-B VRF on R1 and R4 as 0.0.0.67 to ensure that OSPF routes are injected as O IA routes on the Customer Routers.

R1 Router ospf 6 vrf Cust-B Domain-id 0.0.0.67
R4 Router ospf 7 vrf Cust-B Domain-id 0.0.0.67

NOTE:

For the Cust-B VRF, the OSPF routes from the other site now appear as O IA routes.

Lab 8 – Configuring MPLS VPN – PE-CE using OSPF – Sham-Link

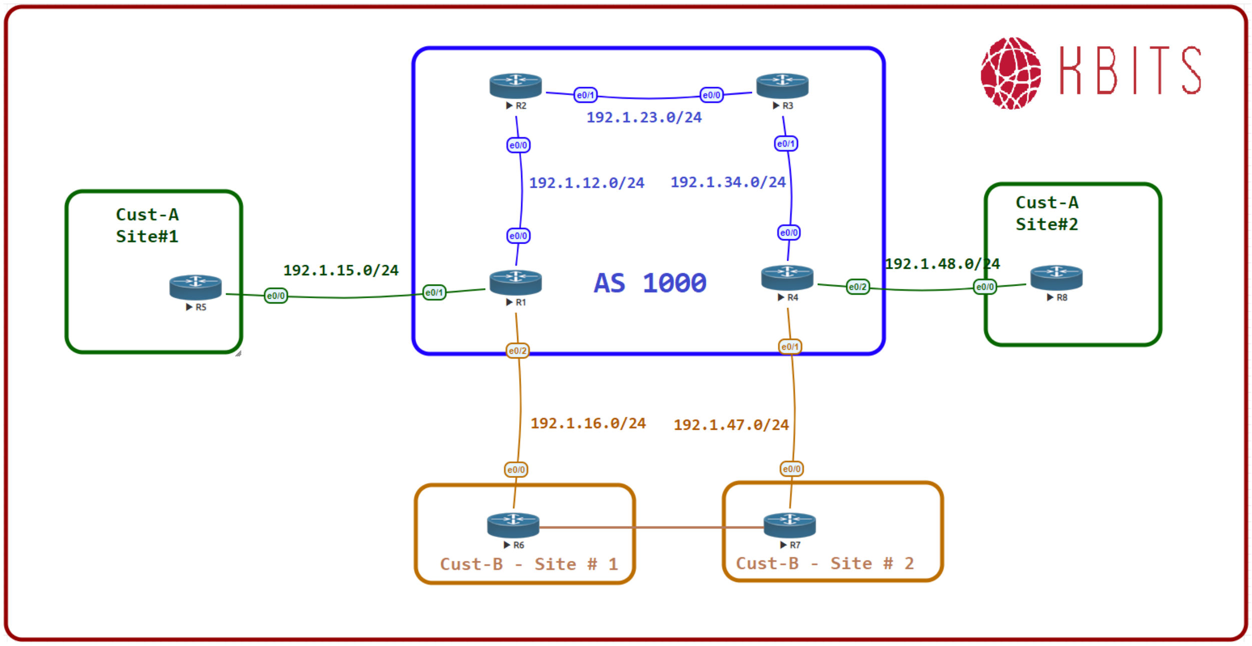

Task 1

Configure a Link between R6 and R7 as 10.67.67.0/24. Advertise this link in OSPF. E 0/1 on both routers to connect. As this is a backup (backdoor) link, set the cost on both sides to be 1000.

R6	R7
Interface E 0/1 Ip address 10.67.67.6 255.255.255.0 Ip ospf cost 1000 No shut ! Router OSPF 1 Network 10.67.67.0 0.0.0.255 area 0	Interface E 0/1 Ip address 10.67.67.6 255.255.255.0 Ip ospf cost 1000 No shut ! Router OSPF 1 Network 10.67.67.0 0.0.0.255 area 0

Task 2

Configure a new loopback each on R1 and R4. This newly created loopback should be part of vrf Cust-B. Advertise this loopback under BGP for the Cust-B vrf. The Loopback information is as follows:

- R1 – Loopback 67 – 172.16.67.1/32
- R4 – Loopback 67 – 172.16.67.4/32

R1 Interface Loopback 67 Ip vrf forwarding Cust-B Ip address 172.16.67.1 255.255.255.255 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-B Network 172.16.67.1 mask 255.255.255.255
R4 Interface Loopback 67 Ip vrf forwarding Cust-B Ip address 172.16.67.4 255.255.255.255 ! Router BGP 1000 ! Address-family ipv4 vrf Cust-B Network 172.16.67.4 mask 255.255.255.255

Task 3

Traffic between Cust-B Sites should be using the new link (Back door) although the cost is much higher than the MPLS cloud. You would like the traffic to go thru the MPLS link instead. Configure a Sham-Link between R1 and R4 based on the new Loopbacks created in the previous step.

R1 Router ospf 6 vrf Cust-B area 0 sham-link 172.16.67.1 172.16.67.4
R4 Router ospf 8 vrf Cust-B area 0 sham-link 172.16.67.4 172.16.67.1

Lab 9 – Configuring MPLS VPN Extranets

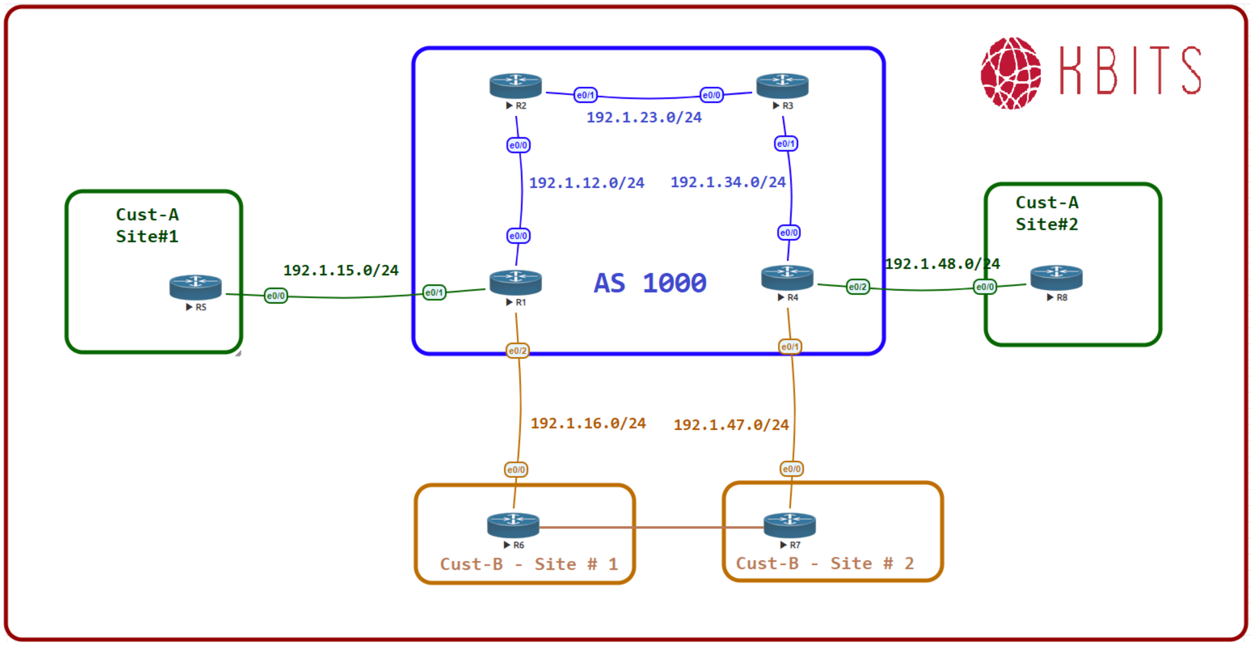

Task 1

Configure R1 such that it sets the RT for the 10.5.5.0/24 route in the Cust-A vrf using a Route-Target of 1000:99. These routes will be later imported into Cust-B.

R1

```
access-list 55 permit 10.5.5.0 0.0.0.255
!
route-map EM-CustA permit 10
 match ip address 55
 set extcommunity rt  1000:99
```

Task 2

Configure R1 such that it sets the RT for the 10.6.6.0/24 route in the Cust-B vrf using a Route-Target of 1000:99. These routes will be later imported into Cust-A.

```
R1

access-list 66 permit 10.6.6.0 0.0.0.255
!
route-map EM-CustB permit 10
 match ip address 66
 set extcommunity rt  1000:99
```

Task 3

Configure R1 Cust-A & Cust-B vrf's to export routes using the Route-map create in the previous steps. Also import the common RT to allow routes to be inter-exchanged between them.

```
R1

Vrf definition Cust-A
 Address-family ipv4
  Export map  EM-CustA
  Route-target import 1000:99
!
Vrf definition Cust-B
 Address-family ipv4
  Export map  EM-CustB
  Route-target import 1000:99
```

Task 4

Configure R4 such that it sets the RT for the 10.8.8.0/24 route in the Cust-A vrf using a Route-Target of 1000:99. These routes will be later imported into Cust-B.

```
R4

access-list 88 permit 10.8.8.0 0.0.0.255
!
route-map EM-CustA permit 10
 match ip address 88
 set extcommunity rt  1000:99
```

Task 5

Configure R4 such that it sets the RT for the 10.7.7.0/24 route in the Cust-B vrf using a Route-Target of 1000:99. These routes will be later imported into Cust-A.

R4

```
access-list 77 permit 10.7.7.0 0.0.0.255
!
route-map EM-CustB permit 10
 match ip address 77
 set extcommunity rt  1000:99
```

Task 6

Configure R4 Cust-A & Cust-B vrf's to export routes using the Route-map create in the previous steps. Also import the common RT to allow routes to be inter-exchanged between them.

R4

```
Vrf definition Cust-A
 Address-family ipv4
  Export map  EM-CustA
  Route-target import 1000:99
!
Vrf definition Cust-B
 Address-family ipv4
  Export map  EM-CustB
  Route-target import 1000:99
```

NOTE:

Reload the Routers without saving the configs. This will setup the topology for the next lab.

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring MPLS VPN for IPv6 Networks

Lab 1 – Configuring MPLS Unicast Routing

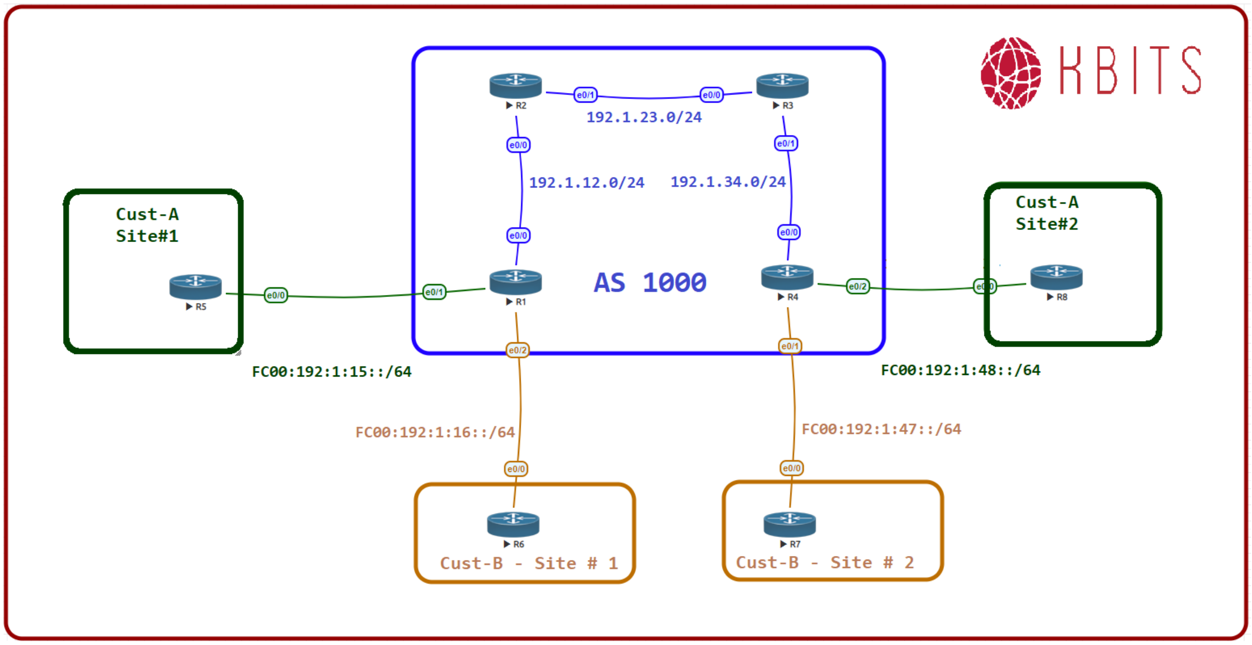

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
E 0/0	192.1.12.1	255.255.255.0
E 0/1	FC00:192:1:15::1	/64
E 0/2	FC00:192:1:16::1	/64

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
E 0/0	192.1.12.2	255.255.255.0
E 0/1	192.1.23.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
E 0/0	192.1.23.3	255.255.255.0
E 0/1	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
E 0/0	192.1.34.4	255.255.255.0
E 0/1	FC00:192:1:47::4	/64
E 0/2	FC00:192:1:48::4	/64

R5

Interface	IP Address	Subnet Mask
Loopback 0	FC00:10:5:5::5	/64
E 0/0	FC00:192:1:15::5	/64

R6

Interface	IP Address	Subnet Mask
Loopback 0	FC00:10:6:6::6	/64
E 0/0	FC00:192:1:16::6	/64

R7

Interface	IP Address	Subnet Mask
Loopback 0	FC00:10:7:7::7	/64
E 0/0	FC00:192:1:47::7	/64

R8

Interface	IP Address	Subnet Mask
Loopback 0	FC00:10:8:8::8	/64
E 0/0	FC00:192:1:48::8	/64

Task 1

Configure OSPF between all the SP routers (R1, R2, R3, R4). Use x.x.x.x as the router-id, where x is the Router number. Advertise all Internal links in OSPF in area 0.

R1 Router ospf 1 Router-id 1.1.1.1 Network 1.1.1.1 0.0.0.0 area 0 Network 192.1.12.0 0.0.0.255 area 0	**R2** Router ospf 1 Router-id 2.2.2.2 Network 2.2.2.2 0.0.0.0 area 0 Network 192.1.12.0 0.0.0.255 area 0 Network 192.1.23.0 0.0.0.255 area 0
R3 Router ospf 1 Router-id 3.3.3.3 Network 3.3.3.3 0.0.0.0 area 0 Network 192.1.23.0 0.0.0.255 area 0 Network 192.1.34.0 0.0.0.255 area 0	**R4** Router ospf 1 Router-id 4.4.4.4 Network 4.4.4.4 0.0.0.0 area 0 Network 192.1.34.0 0.0.0.255 area 0

Task 2

Configure MPLS on all the physical links in the SP Network. Use LDP to distribute labels. The LDP neighbour relationships should be formed based on the most reliable interface. The Labels should be assigned from the range X00 – X99, where X is the router number.

R1 Mpls ldp router-id Loopback0 ! Mpls label range 100 199 ! Interface E 0/0 Mpls ip	**R2** Mpls ldp router-id Loopback0 ! Mpls label range 200 299 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip
R3 Mpls ldp router-id Loopback0 ! Mpls label range 300 399 ! Interface E 0/0	**R4** Mpls ldp router-id Loopback0 ! Mpls label range 400 499 ! Interface E 0/0

Mpls ip ! Interface E 0/1 Mpls ip	Mpls ip

Lab 2 – Configuring MPLS VPNv6

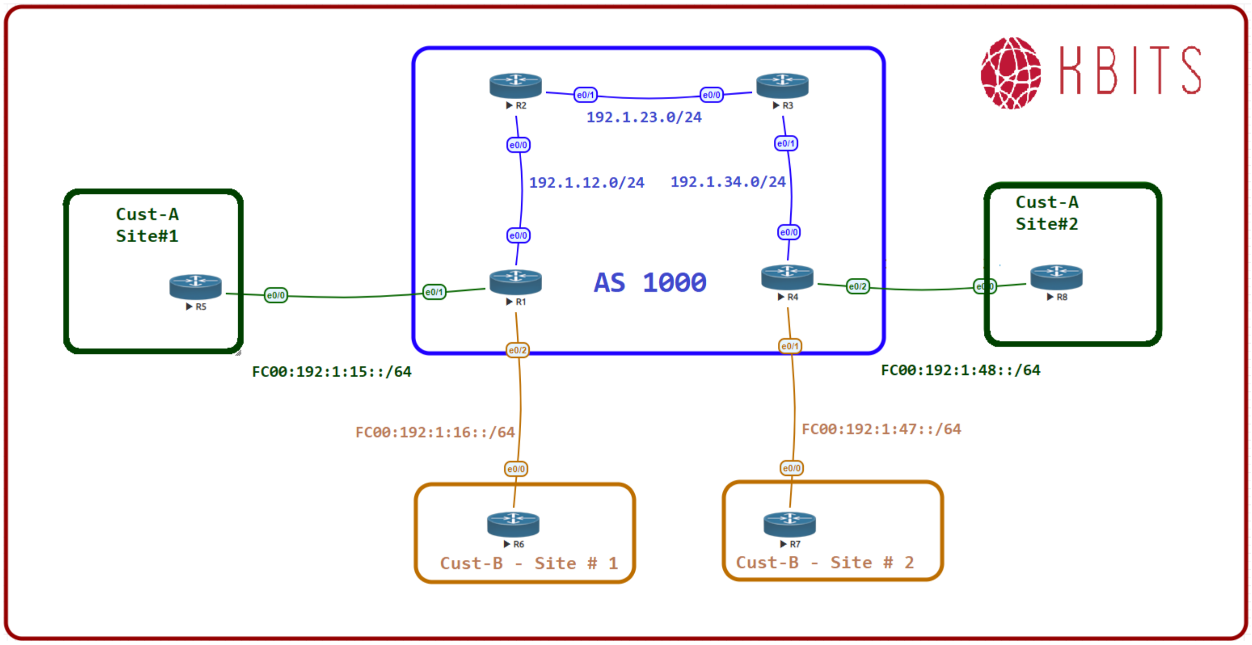

Task 1

Configure a VPNv4 neighbor relationship between R1 and R4.

R1 Ipv6 unicast-routing ! Router BGP 1000 Neighbor 4.4.4.4 remote-as 1000 Neighbor 4.4.4.4 update-source loopback0 ! Address-family vpnv6 Neighbor 4.4.4.4 activate
R4 Ipv6 unicast-routing ! Router BGP 1000 Neighbor 1.1.1.1 remote-as 1000 Neighbor 1.1.1.1 update-source loopback0 ! Address-family vpnv6 Neighbor 1.1.1.1 activate

Task 2

Configure a VRF **Cust-A** with a RD value of 1000:1 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R4.

R1 Vrf definition Cust-A rd 1000:1 address-family ipv6   route-target both 1000:1 ! Interface E 0/1 vrf forwarding Cust-A Ipv6 address fc00:192:1:15::1/64 No shut
R4 Vrf definition Cust-A rd 1000:1

```
 address-family ipv6
  route-target both 1000:1
!
Interface E 0/2
 vrf forwarding Cust-A
 Ipv6 address fc00:192:1:48::4/64
 No shut
```

Task 3

Configure BGP as the Routing Protocol between R5 and R1-vrf Cust-A. Advertise all the routes on R5 in BGP. Configure R5 with an AS # of 65005. Configure the BGP neighbor relationship on R1 for the Cust-A VRF. Make sure the VRF Cust-A on R4 has reachability to routes learned from R5.

R1

```
Router BGP 1000
!
 address-family ipv6 vrf Cust-A
  neighbor fc00:192:1:15::5 remote-as 65005
```

R5

```
ipv6 unicast-routing
!
Interface Loo0
 ipv6 address fc00:10:5:5::5/64
!
Interface E 0/0
 ipv6 address fc00:192:1:15::5/64
 no shut
!
router bgp 65005
 address-family ipv6
  neighbor fc00:192:1:15::1 remote-as 1000
  network fc00:10:5:5::/64
```

Task 4

Configure BGP as the Routing Protocol between R8 and R4-vrf Cust-A. Advertise all the routes on R8 in BGP. Configure R8 with an AS # of 65008. Configure the BGP neighbor relationship on R4 for the Cust-A VRF. Make sure the VRF Cust-A on R1 has reachability to routes learned from R8.

R4

```
Router BGP 1000
!
 address-family ipv6 vrf Cust-A
  neighbor fc00:192:1:48::8 remote-as 65008
```

R8

```
ipv6 unicast-routing
!
Interface Loo0
 ipv6 address fc00:10:8:8::8/64
!
Interface E 0/0
 ipv6 address fc00:192:1:48::8/64
 no shut
!
router bgp 65008
 address-family ipv6
  neighbor fc00:192:1:48::4 remote-as 1000
  network fc00:10:8:8::/64
```

Task 5

Configure a VRF **Cust-B** with a RD value of 1000:2 on R1 and R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-B sites on R1 and R4.

R1

```
Vrf definition Cust-B
 rd 1000:2
 address-family ipv6
  route-target both 1000:2
!
Interface E 0/2
 vrf forwarding Cust-B
 Ipv6 address fc00:192:1:16::1/64
 No shut
```

```
R4

Vrf definition Cust-B
 rd 1000:2
 address-family ipv6
  route-target both 1000:2
!
Interface E 0/1
 vrf forwarding Cust-B
 Ipv6 address fc00:192:1:47::4/64
 No shut
```

Task 6

Configure BGP as the Routing Protocol between R6 and R1-vrf Cust-B. Advertise all the routes on R6 in BGP. Configure R6 with an AS # of 65006. Configure the BGP neighbor relationship on R1 for the Cust-B VRF. Make sure the VRF Cust-B on R4 has reachability to routes learned from R6.

```
R1

Router BGP 1000
!
 address-family ipv6 vrf Cust-B
  neighbor fc00:192:1:16::6 remote-as 65006
```

```
R6

ipv6 unicast-routing
!
Interface Loo0
 ipv6 address fc00:10:6:6::6/64
!
Interface E 0/0
 ipv6 address fc00:192:1:16::6/64
 no shut
!
router bgp 65006
 address-family ipv6
  neighbor fc00:192:1:16::1 remote-as 1000
  network fc00:10:6:6::/64
```

Task 7

Configure BGP as the Routing Protocol between R7 and R4-vrf Cust-B. Advertise all the routes on R7 in BGP. Configure R7 with an AS # of 65007. Configure the BGP neighbor relationship on R4 for the Cust-B VRF. Make sure the VRF Cust-B on R1 has reachability to routes learned from R7.

R4 Router BGP 1000 ! address-family ipv6 vrf Cust-B   neighbor fc00:192:1:47::7 remote-as 65007

R7
 	
ipv6 unicast-routing
!
Interface Loo0
 ipv6 address fc00:10:7:7::7/64
!
Interface E 0/0
 ipv6 address fc00:192:1:47::7/64
 no shut
!
router bgp 65007
 address-family ipv6
  neighbor fc00:192:1:44::4 remote-as 1000
  network fc00:10:7:7::/64

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring Inter-AS MPLS VPN

Lab 1 – Configuring MPLS Unicast Routing

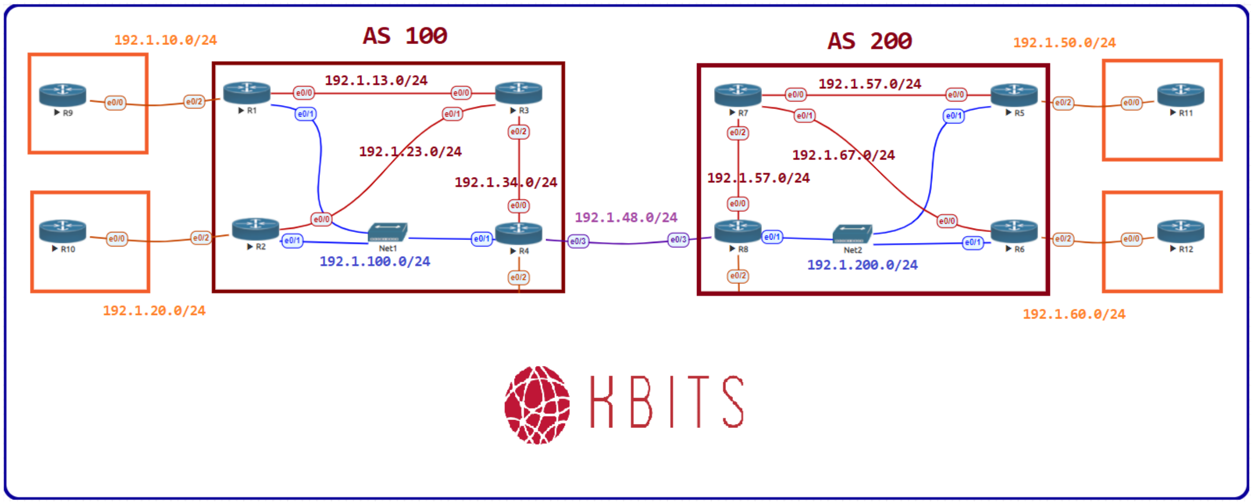

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
E 0/0	192.1.13.1	255.255.255.0
E 0/1	192.1.100.1	255.255.255.0
E 0/2	192.1.10.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
E 0/0	192.1.23.2	255.255.255.0
E 0/1	192.1.100.2	255.255.255.0
E 0/2	192.1.20.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
E 0/0	192.1.13.3	255.255.255.0
E 0/1	192.1.23.3	255.255.255.0
E 0/2	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
E 0/0	192.1.34.4	255.255.255.0
E 0/1	192.1.100.4	255.255.255.0
E 0/3	192.1.48.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.255.255.255
E 0/0	192.1.57.5	255.255.255.0
E 0/1	192.1.200.5	255.255.255.0
E 0/2	192.1.50.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.255.255.255
E 0/0	192.1.67.6	255.255.255.0
E 0/1	192.1.200.6	255.255.255.0
E 0/2	192.1.60.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.255.255.255
E 0/0	192.1.57.7	255.255.255.0
E 0/1	192.1.67.7	255.255.255.0
E 0/2	192.1.78.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.255.255.255
E 0/0	192.1.78.8	255.255.255.0
E 0/1	192.1.200.8	255.255.255.0
E 0/3	192.1.48.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	10.9.9.9	255.255.255.0
E 0/0	192.1.10.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	10.10.10.10	255.255.255.0
E 0/0	192.1.20.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	10.11.11.11	255.255.255.0
E 0/0	192.1.50.11	255.255.255.0

R12

Interface	IP Address	Subnet Mask
Loopback 0	10.12.12.12	255.255.255.0
S 0/0	192.1.60.12	255.255.255.0

Task 1

Configure OSPF as the underlay routing protocol between all AS 100 routers. (R1, R2, R3, R4). Use x.x.x.x as the router-id, where x is the Router number. Enable all Internal links in OSPF in area 0.

R1	**R2**
Router ospf 1 Router-id 1.1.1.1 Network 1.1.1.1 0.0.0.0 area 0 Network 192.1.13.0 0.0.0.255 area 0 Network 192.1.100.0 0.0.0.255 area 0	Router ospf 1 Router-id 2.2.2.2 Network 2.2.2.2 0.0.0.0 area 0 Network 192.1.23.0 0.0.0.255 area 0 Network 192.1.100.0 0.0.0.255 area 0
R3 Router ospf 1 Router-id 3.3.3.3 Network 3.3.3.3 0.0.0.0 area 0 Network 192.1.13.0 0.0.0.255 area 0 Network 192.1.24.0 0.0.0.255 area 0 Network 192.1.34.0 0.0.0.255 area 0	**R4** Router ospf 1 Router-id 4.4.4.4 Network 4.4.4.4 0.0.0.0 area 0 Network 192.1.34.0 0.0.0.255 area 0 Network 192.1.100.0 0.0.0.255 area 0

Task 2

Configure MPLS (LDP) on all the physical links in AS 100.

R1	**R2**
Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip	Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip
R3 Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip !	**R4** Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip

Interface E 0/2 Mpls ip	

Task 3

Configure IS-IS as the underlay routing protocol between all AS 200 routers. (R5, R6, R7, R8). Use xxxx.xxxx.xxxx as the system-id, where x is the Router number. Use 49.0000 as the Area ID. Configure the routers as Level-2 Routers. Enable the wide metric style.

R5 Router isis Net 49.0000.5555.5555.5555.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis ! Interface Loopback0 Ip router isis	**R6** Router isis Net 49.0000.6666.6666.6666.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis ! Interface Loopback0 Ip router isis
R7 Router isis Net 49.0000.7777.7777.7777.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis ! Interface E 0/2 Ip router isis ! Interface Loopback0 Ip router isis	**R8** Router isis Net 49.0000.8888.8888.8888.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis ! Interface Loopback0 Ip router isis

Task 4
Configure MPLS (LDP) on all the physical links in AS 200.

R5	R6
R5 Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip	**R6** Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip
R7 Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip ! Interface E 0/2 Mpls ip	**R8** Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip

Lab 2 – Configuring Intra-AS MPLS VPNv4 – AS 100 & AS 200

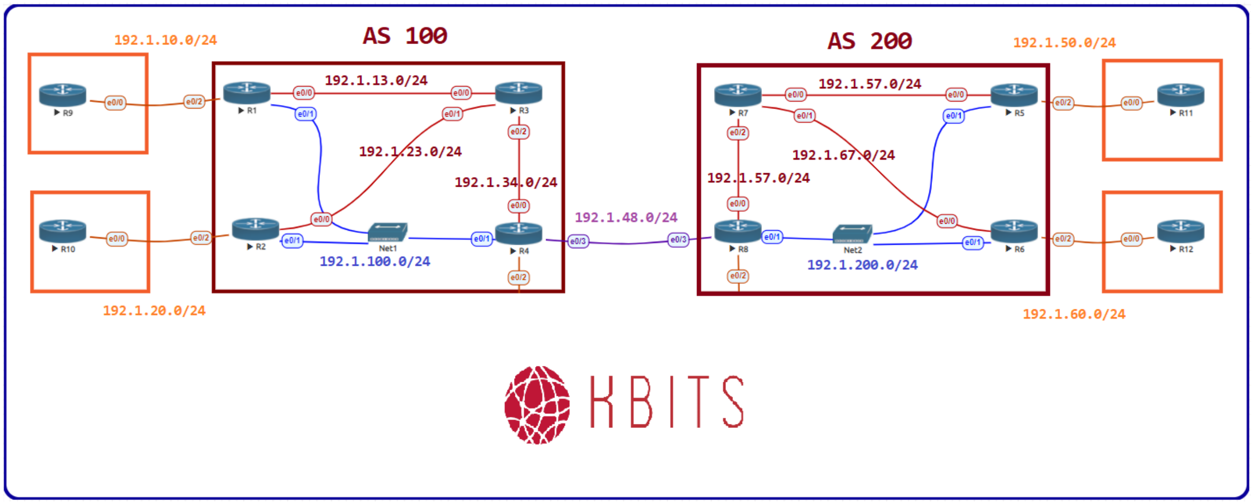

AS 100

Task 1

Configure R3 as the VPNv4 Route Reflector for R1, R2 & R3.

R1
 	
Router bgp 100
 Neighbor 3.3.3.3 remote-as 100
 Neighbor 3.3.3.3 update-source loo 0
!
 address-family vpnv4
  Neighbor 3.3.3.3 activate
	R2
	
	Router bgp 100
	 Neighbor 3.3.3.3 remote-as 100
	 Neighbor 3.3.3.3 update-source loo 0
	!
	 address-family vpnv4
	  Neighbor 3.3.3.3 activate

R4
 	
Router bgp 100
 Neighbor 3.3.3.3 remote-as 100
 Neighbor 3.3.3.3 update-source loo 0
!
 address-family vpnv4
  Neighbor 3.3.3.3 activate
	R3
	
	Router bgp 100
	 Neighbor IBGP peer-group
	 Neighbor IBGP remote-as 100
	 Neighbor IBGP update-source loo 0
	!
	 address-family vpnv4
	  Neighbor IBGP route-reflector-client
	  Neighbor 1.1.1.1 activate
	  Neighbor 2.2.2.2 activate
	  Neighbor 4.4.4.4 activate

Task 2

Configure a VRF **Cust-A** with a RD value of 100:1 on R1, R2 & R4. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1, R2 and R4.

R1	**R2**

Vrf definition Cust-A
 rd 100:1
 address-family ipv4
  route-target both 100:1
!
Interface E 0/2
 vrf forwarding Cust-A
 Ip address 192.1.10.1 255.255.255.0
 No shut
	Vrf definition Cust-A
	 rd 100:1
	 address-family ipv4
	  route-target both 100:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.20.2 255.255.255.0
	 No shut

Task 3

Configure BGP as the Routing Protocol between the PE Routes (R1, R2 & R4) and respective CE Routres (R9, R10 & R13) in vrf Cust-A. Advertise all the routes on CE routers in BGP. Configure CEs with an AS # of 650XX, where XX is the 2 digit Router # (65009, 65010 & 65013).

R1	**R2**

Router BGP 100
!
 address-family ipv4 vrf Cust-A
  neighbor 192.1.10.9 remote-as 65009
	Router BGP 100
	!
	 address-family ipv4 vrf Cust-A
	  neighbor 192.1.20.10 remote-as 65010

R9
 	
Router BGP 65009
!
  Network 10.9.9.0 mask 255.255.255.0
  neighbor 192.1.10.1 remote-as 100
	R10
	
	Router BGP 65010
	!
	  Network 10.10.10.0 mask 255.255.255.0
	  neighbor 192.1.20.2 remote-as 100

AS 200

Task 1

Configure R7 as the VPNv4 Route Reflector for R5, R6 & R7.

R5
 	
Router bgp 200
 Neighbor 7.7.7.7 remote-as 200
 Neighbor 7.7.7.7 update-source loo 0
!
 address-family vpnv4
  Neighbor 7.7.7.7 activate
	R6
	
	Router bgp 200
	 Neighbor 7.7.7.7 remote-as 200
	 Neighbor 7.7.7.7 update-source loo 0
	!
	 address-family vpnv4
	  Neighbor 7.7.7.7 activate

R8
 	
Router bgp 200
 Neighbor 7.7.7.7 remote-as 200
 Neighbor 7.7.7.7 update-source loo 0
!
 address-family vpnv4
  Neighbor 7.7.7.7 activate
	R7
	
	Router bgp 200
	 Neighbor IBGP peer-group
	 Neighbor IBGP remote-as 200
	 Neighbor IBGP update-source loo 0
	!
	 address-family vpnv4
	  Neighbor IBGP route-reflector-client
	  Neighbor 5.5.5.5 activate
	  Neighbor 6.6.6.6 activate
	  Neighbor 8.8.8.8 activate

Task 2

Configure a VRF **Cust-A** with a RD value of 200:1 on R5, R6 & R8. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R5, R6 and R8.

R5
 	
Vrf definition Cust-A
 rd 200:1
 address-family ipv4
  route-target both 200:1
!
Interface E 0/2
 vrf forwarding Cust-A
 Ip address 192.1.50.5 255.255.255.0
 No shut
	R6
	
	Vrf definition Cust-A
	 rd 200:1
	 address-family ipv4
	  route-target both 200:1
	!
	Interface E 0/2
	 vrf forwarding Cust-A
	 Ip address 192.1.60.6 255.255.255.0
	 No shut

Task 3

Configure BGP as the Routing Protocol between the PE Routes (R1, R2 & R4) and respective CE Routres (R9, R10 & R13) in vrf Cust-A. Advertise all the routes on CE routers in BGP. Configure CEs with an AS # of 650XX, where XX is the 2 digit Router # (65009, 65010 & 65013).

R5
 	
Router BGP 200
!
 address-family ipv4 vrf Cust-A
  neighbor 192.1.50.11 remote-as 65011
	R6
	
	Router BGP 200
	!
	 address-family ipv4 vrf Cust-A
	  neighbor 192.1.60.12 remote-as 65012

R11
 	
Router BGP 65011
!
  Network 10.11.11.0 mask 255.255.255.0
  neighbor 192.1.50.5 remote-as 200
	R12
	
	Router BGP 65012
	!
	  Network 10.12.12.0 mask 255.255.255.0
	  neighbor 192.1.60.6 remote-as 200

Note:

Save the Configs on all the routers. **Do not save the configs during the labs**. At the completion of this lab, **reload the routers without saving**. This will allow you to do the next lab based on the same topology.

Lab 3 – Configuring Inter-AS MPLS VPN – Option A Back-to-Back VRF

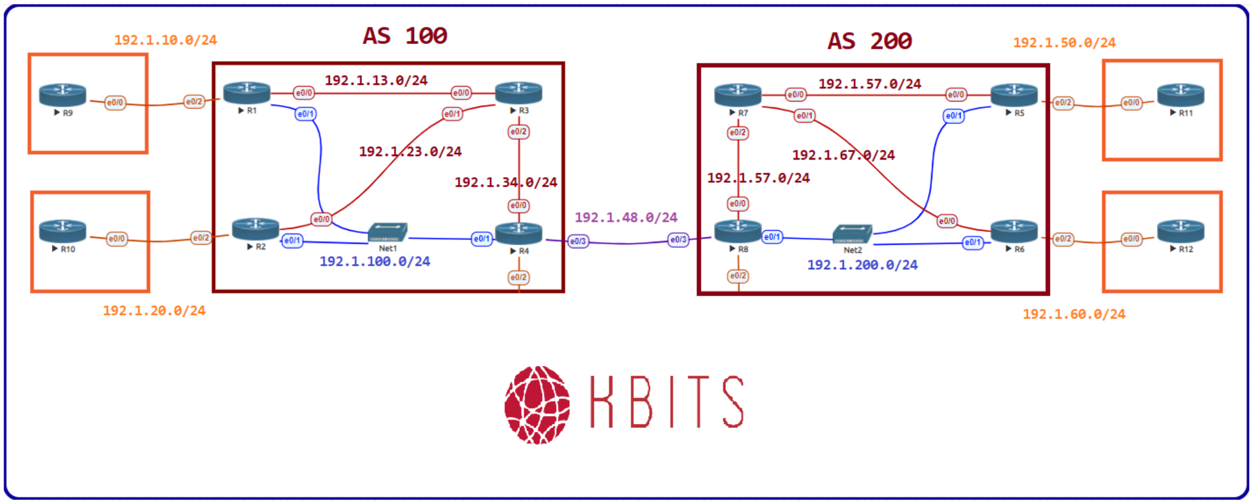

Task 1
Connect the ASBR's (R4 & R8) to each other on the E 0/3 Interface. Configure the peer ASBR as a CE router within a VRF. Run the PE-CE protocol just like you would on a PE-CE Link.

R4	**R8**
Interface E0/3	Interface E0/3
vrf forwarding Cust-A	vrf forwarding Cust-A
ip address 192.1.48.4 255.255.255.0	ip address 192.1.48.8 255.255.255.0
no shut	no shut
!	!
Router bgp 100	Router bgp 200
address-family ipv4 vrf Cust-A	address-family ipv4 vrf Cust-A
neighbor 192.1.48.8 remote-as 200	neighbor 192.1.48.4 remote-as 100

Verification:
Make sure you have end-to-end reachability between the 4 Cust-A Sites (R9 – R12).

NOTE:
Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 4 – Configuring Inter-AS MPLS VPN – Option B MP-eBGP on ASBR

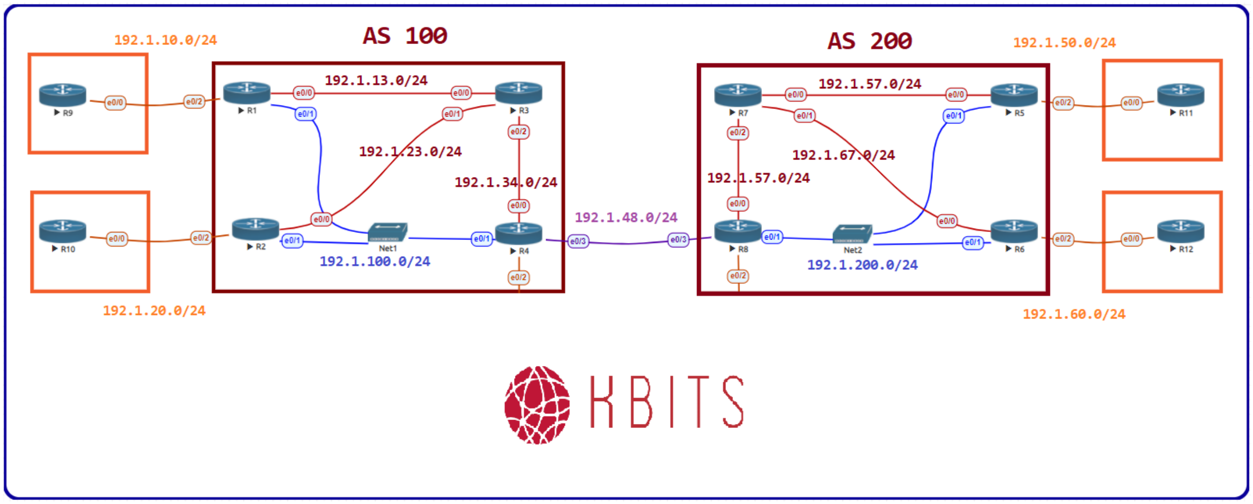

Task 1

By default, the ASBR will reject all routes that do not match an Import RT on a local VRF. In this option, the ASBR generally does not have a VRF configured. Disable the Route-target filter on the ASBRs to receive all VPNv4 routes to all all routes to come in.

R4	R8
router bgp 100 no bgp default route-target filter	router bgp 200 no bgp default route-target filter

Task 2

Configure an MP-eBGP relationship between the ASBRs. Make sure to set the Next-hop-self when sending VPNv4 routes towards the RR.

R4	R8

Interface E0/3
 ip address 192.1.48.4 255.255.255.0
 no shut
!
router bgp 100
 neighbor 192.1.48.8 remote-as 200
 address-family vpnv4
  neighbor 192.1.48.8 activate
  neighbor 10.3.3.3 next-hop-self
	Interface E0/3
	 ip address 192.1.48.8 255.255.255.0
	 no shut
	!
	router bgp 200
	 neighbor 192.1.48.4 remote-as 100
	 address-family vpnv4
	  neighbor 192.1.48.4 activate
	  neighbor 10.7.7.7 next-hop-self

Task 3

Configure the PE Routers to import the RT from the Remote-AS.

R1
 	
vrf definition Cust-A
 address-family ipv4
  route-target import 200:1
	R2
	
	vrf definition Cust-A
	 address-family ipv4
	  route-target import 200:1

R5
 	
vrf definition Cust-A
 address-family ipv4
  route-target import 100:1
	R6
	
	vrf definition Cust-A
	 address-family ipv4
	  route-target import 100:1

Verification:

Make sure you have end-to-end reachability between the 4 Cust-A Sites (R9 – R12).

NOTE:

Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 5 – Configuring Inter-AS MPLS VPN – Option C MP-eBGP on RR

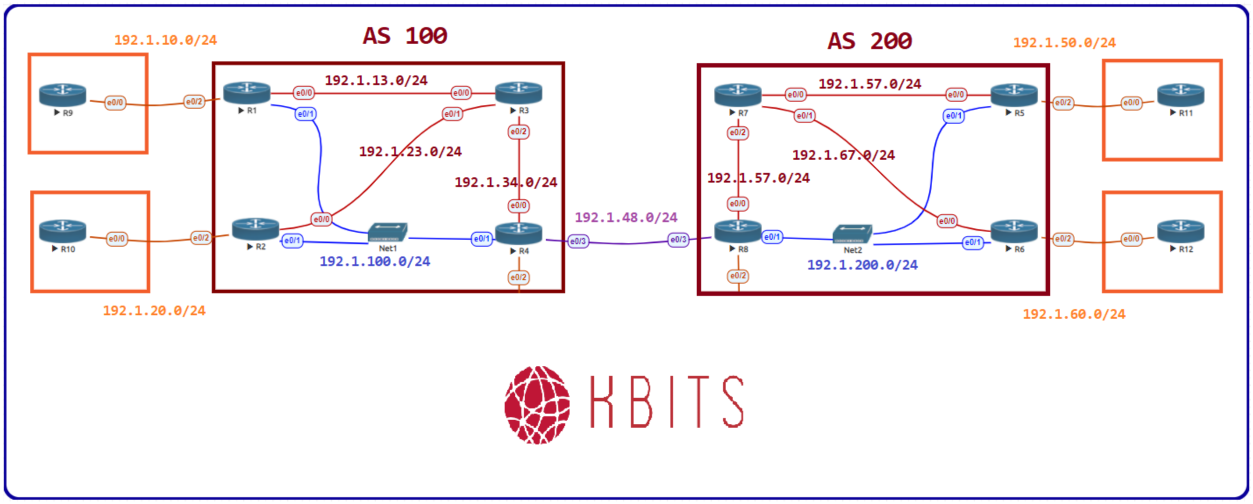

Task 1

Configure BGP between the 2 AS's on the ASBRs (R4 & R8). Make sure to send the label for any routes that are exchanged. Disbable the BGP neighbor relationship between the RR & the ASBR as it is not required for this Option.

R4	**R8**
Interface E0/3	Interface E0/3
ip address 192.1.48.4 255.255.255.0	ip address 192.1.48.8 255.255.255.0
no shut	no shut
!	!
router bgp 100	router bgp 200
neighbor 192.1.48.8 remote-as 200	neighbor 192.1.48.4 remote-as 100
neighbor 192.1.48.8 send-label	neighbor 192.1.48.4 send-label
no neighbor 3.3.3.3	no neighbor 7.7.7.7

Task 2

Configure Filtered Route Redistribution on the ASBR to redistribute the RR loopback address to the remote AS. This is required to configure a MP-eBGP relationship between the RR's. Make sure to redistribute BGP routes into the local IGP.

R4	**R8**
Access-list 1 permit 3.3.3.3 0.0.0.0 Access-list 2 permit 7.7.7.7 0.0.0.0 ! route-map O2B match ip address 1 ! Route-map B2O Match ip address 2 ! router bgp 100 redistribute ospf 1 route-map O2B ! Router ospf 1 Redistribute bgp 100 route-map B2O	Access-list 1 permit 7.7.7.7 0.0.0.0 Access-list 2 permit 3.3.3.3 0.0.0.0 ! route-map I2B match ip address 1 ! Route-map B2I Match ip address 2 ! router bgp 200 redistribute isis route-map I2B ! Router isis Redistribute bgp 200 route-map B2I

Task 3

Configure MP-eBGP between RRs. Make sure to enable the BGP Multi-hop capability.

R3	**R2**

router bgp 100
 neighbor 7.7.7.7 remote-as 200
 neighbor 7.7.7.7 update-source lo 0
 neighbor 7.7.7.7 ebgp-multihop
!
 address-family vpnv4
  neighbor 7.7.7.7 activate
	router bgp 200
	 neighbor 3.3.3.3 remote-as 100
	 neighbor 3.3.3.3 update-source lo 0
	 neighbor 3.3.3.3 ebgp-multihop
	!
	 address-family vpnv4
	  neighbor 3.3.3.3 activate

Task 4
Configure the PE Routers to import the RT from the Remote-AS.

R1
 	
vrf definition Cust-A
 address-family ipv4
  route-target import 200:1
	R2
	
	vrf definition Cust-A
	 address-family ipv4
	  route-target import 200:1

R5
 	
vrf definition Cust-A
 address-family ipv4
  route-target import 100:1
	R6
	
	vrf definition Cust-A
	 address-family ipv4
	  route-target import 100:1

Verification:
Make sure you have end-to-end reachability between the 4 Cust-A Sites (R9 – R12).

NOTE:
Reload the Routers without saving the configs. This will setup the topology for the next lab.

Lab 6 – Configuring Inter-AS MPLS VPN – Use a Non-VPN Transit Provider

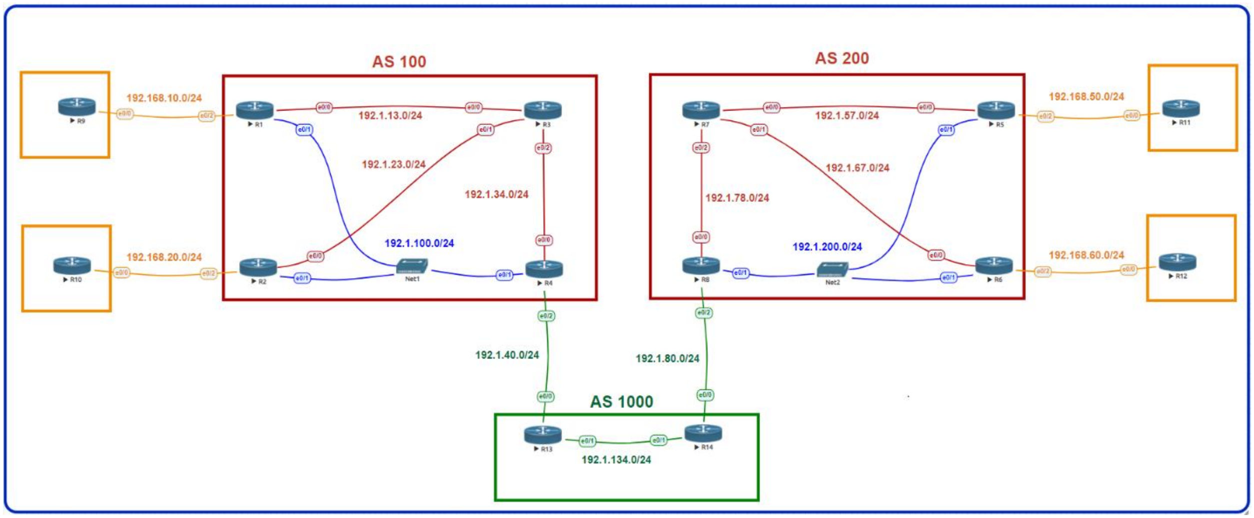

Task 1

Configure BGP between AS 100 & AS 200 via AS 1000. AS 1000 is a Non-VPN SP providing labeled reachability between the 2 AS's. Make sure to send the label for any routes that are exchanged. Disbable the BGP neighbor relationship between the RR & the ASBR within AS 100 & AS 200 as it is not required for this Option.

R4 Interface E0/2 ip address 192.1.40.4 255.255.255.0 no shut ! router bgp 100 neighbor 192.1.40.13 remote-as 1000 neighbor 192.1.40.13send-label no neighbor 3.3.3.3	**R8** Interface E0/2 ip address 192.1.80.8 255.255.255.0 no shut ! router bgp 200 neighbor 192.1.80.14 remote-as 1000 neighbor 192.1.80.14 send-label no neighbor 3.3.3.3
R13 Interface E 0/0 Ip address 192.1.40.13 255.255.255.0 No shut ! Interface E 0/1 Ip address 192.1.134.13 255.255.255.0	**R14** Interface E 0/0 Ip address 192.1.80.14 255.255.255.0 No shut ! Interface E 0/1 Ip address 192.1.134.14 255.255.255.0

Mpls ip No shut ! Interface Loopback0 Ip add 13.13.13.13 255.255.255.255 ! Mpls ldp router-id Loopback0 ! router ospf 1 router-id 0.0.0.13 network 13.13.13.13 0.0.0.0 area 0 network 192.1.134.0 0.0.0.255 area 0 ! router bgp 1000 neighbor 14.14.14.14 remote-as 1000 neighbor 14.14.14.14 update-source lo10 neighbor 14.14.14.14 send-label neighbor 14.14.14.14 next-hop-self neighbor 192.1.40.4 remote-as 100 neighbor 192.1.40.4 send-label	Mpls ip No shut ! Interface Loopback0 Ip add 14.14.14.14 255.255.255.255 ! Mpls ldp router-id Loopback0 ! router ospf 1 router-id 0.0.0.13 network 14.14.14.14 0.0.0.0 area 0 network 192.1.134.0 0.0.0.255 area 0 ! router bgp 1000 neighbor 13.13.13.13 remote-as 1000 neighbor 13.13.13.13 update-source lo10 neighbor 13.13.13.13 send-label neighbor 13.13.13.13 next-hop-self neighbor 192.1.80.8 remote-as 200 neighbor 192.1.80.8 send-label

Task 2

Configure Filtered Route Redistribution on the ASBR to redistribute the RR loopback address to the remote AS. This is required to configure a MP-eBGP relationship between the RR's. Make sure to redistribute BGP routes into the local IGP. We will configure MP-eBGP to keep the next-hop-unchanged to avoid the RR in the Data path. Redistribute the PE Loopbacks to provide reachability to the PE Loopbacks in the Remote AS.

R4	R8
Access-list 1 permit 1.1.1.1 0.0.0.0 Access-list 1 permit 2.2.2.2 0.0.0.0 Access-list 1 permit 3.3.3.3 0.0.0.0 Access-list 2 permit 5.5.5.5 0.0.0.0 Access-list 2 permit 6.6.6.6 0.0.0.0 Access-list 2 permit 7.7.7.7 0.0.0.0 ! route-map O2B match ip address 1 ! Route-map B2O Match ip address 2 ! router bgp 100 redistribute ospf 1 route-map O2B ! Router ospf 1 Redistribute bgp 100 route-map B2O	Access-list 1 permit 5.5.5.5 0.0.0.0 Access-list 1 permit 6.6.6.6 0.0.0.0 Access-list 1 permit 7.7.7.7 0.0.0.0 Access-list 2 permit 1.1.1.1 0.0.0.0 Access-list 2 permit 2.2.2.2 0.0.0.0 Access-list 2 permit 3.3.3.3 0.0.0.0 ! route-map I2B match ip address 1 ! Route-map B2I Match ip address 2 ! router bgp 200 redistribute isis route-map I2B ! Router isis Redistribute bgp 200 route-map B2I

Task 3

Configure MP-eBGP between RRs. Make sure to enable the BGP Multi-hop capability.

R3	R2

router bgp 100
 neighbor 7.7.7.7 remote-as 200
 neighbor 7.7.7.7 update-source lo 0
 neighbor 7.7.7.7 ebgp-multihop
!
 address-family vpnv4
  neighbor 7.7.7.7 activate
  neighbor 7.7.7.7 next-hop-unchanged
	router bgp 200
	 neighbor 3.3.3.3 remote-as 100
	 neighbor 3.3.3.3 update-source lo 0
	 neighbor 3.3.3.3 ebgp-multihop
	!
	 address-family vpnv4
	  neighbor 3.3.3.3 activate
	  neighbor 3.3.3.3 next-hop-unchanged

Task 4

Configure the PE Routers to import the RT from the Remote-AS.

R1	**R2**

vrf definition Cust-A
 address-family ipv4
  route-target import 200:1
	vrf definition Cust-A
	 address-family ipv4
	  route-target import 200:1

R5	**R6**

vrf definition Cust-A
 address-family ipv4
  route-target import 100:1
	vrf definition Cust-A
	 address-family ipv4
	  route-target import 100:1

Verification:

Make sure you have end-to-end reachability between the 4 Cust-A Sites (R9 – R12).

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Carrier Supporting Carrier (CSC)

Lab 1 – Configuring Carrier Supporting Carriers (CSC) – IP Only

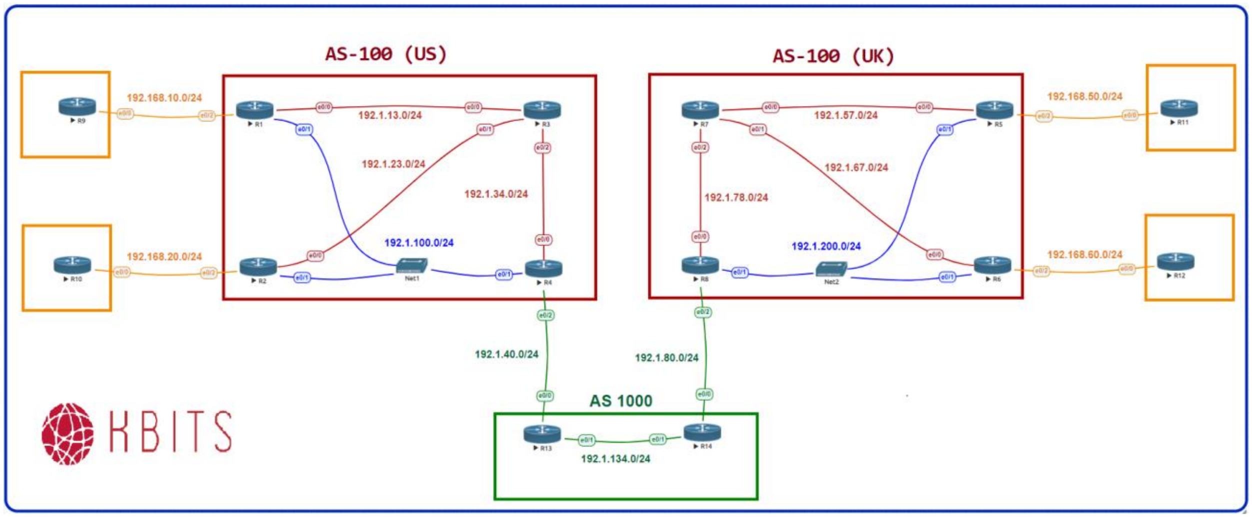

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
E 0/0	192.1.13.1	255.255.255.0
E 0/1	192.1.100.1	255.255.255.0
E 0/2	192.168.10.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
E 0/0	192.1.23.2	255.255.255.0
E 0/1	192.1.100.2	255.255.255.0
E 0/2	192.168.20.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
E 0/0	192.1.13.3	255.255.255.0
E 0/1	192.1.23.3	255.255.255.0
E 0/2	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
E 0/0	192.1.34.4	255.255.255.0
E 0/1	192.1.100.4	255.255.255.0
E 0/2	192.1.40.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.255.255.255
E 0/0	192.1.57.5	255.255.255.0
E 0/1	192.1.200.5	255.255.255.0
E 0/2	192.168.50.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.255.255.255
E 0/0	192.1.67.6	255.255.255.0
E 0/1	192.1.200.6	255.255.255.0
E 0/2	192.168.60.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.255.255.255
E 0/0	192.1.57.7	255.255.255.0
E 0/1	192.1.67.7	255.255.255.0
E 0/2	192.1.78.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.255.255.255
E 0/0	192.1.78.8	255.255.255.0
E 0/1	192.1.200.8	255.255.255.0
E 0/2	192.1.80.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	10.9.9.9	255.255.255.0
E 0/0	192.168.10.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	10.10.10.10	255.255.255.0
E 0/0	192.168.20.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	10.11.11.11	255.255.255.0
E 0/0	192.168.50.11	255.255.255.0

R12

Interface	IP Address	Subnet Mask
Loopback 0	10.12.12.12	255.255.255.0
S 0/0	192.168.60.12	255.255.255.0

R13

Interface	IP Address	Subnet Mask
Loopback 0	13.13.13.13	255.255.255.0
E 0/0	192.1.40.13	255.255.255.0
E 0/1	192.1.134.13	255.255.255.0

R14

Interface	IP Address	Subnet Mask
Loopback 0	14.14.14.14	255.255.255.0
E 0/0	192.1.80.14	255.255.255.0
E 0/1	192.1.134.14	255.255.255.0

Task 1

Configure OSPF as the underlay routing protocol between all AS 100 (US) routers. (R1, R2, R3, R4). Use x.x.x.x as the router-id, where x is the Router number. Enable all Internal links in OSPF in area 0.

R1	R2
Router ospf 1 Router-id 1.1.1.1 Network 1.1.1.1 0.0.0.0 area 0 Network 192.1.13.0 0.0.0.255 area 0 Network 192.1.100.0 0.0.0.255 area 0	Router ospf 1 Router-id 2.2.2.2 Network 2.2.2.2 0.0.0.0 area 0 Network 192.1.23.0 0.0.0.255 area 0 Network 192.1.100.0 0.0.0.255 area 0
R3	**R4**
Router ospf 1 Router-id 3.3.3.3 Network 3.3.3.3 0.0.0.0 area 0 Network 192.1.13.0 0.0.0.255 area 0 Network 192.1.24.0 0.0.0.255 area 0 Network 192.1.34.0 0.0.0.255 area 0	Router ospf 1 Router-id 4.4.4.4 Network 4.4.4.4 0.0.0.0 area 0 Network 192.1.34.0 0.0.0.255 area 0 Network 192.1.100.0 0.0.0.255 area 0

Task 2

Configure IS-IS as the underlay routing protocol between all AS 100 (UK) routers. (R5,R6,R7 & R8). Use xxxx.xxxx.xxxx as the system-id, where x is the Router number. Use 49.0000 as the Area ID. Configure the routers as Level-2 Routers. Enable the wide metric style.

R5	R6
Router isis Net 49.0000.5555.5555.5555.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis !	Router isis Net 49.0000.6666.6666.6666.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis !

Interface E 0/1 Ip router isis ! Interface Loopback0 Ip router isis	Interface E 0/1 Ip router isis ! Interface Loopback0 Ip router isis
R7 Router isis Net 49.0000.7777.7777.7777.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis ! Interface E 0/2 Ip router isis ! Interface Loopback0 Ip router isis	**R8** Router isis Net 49.0000.8888.8888.8888.00 Is-type level-2 Metric-style wide ! Interface E 0/0 Ip router isis ! Interface E 0/1 Ip router isis ! Interface Loopback0 Ip router isis

Task 3

Configure the Backbone-SP (Provider SP) with MPLS VPN to connect the AS 100 (US) & AS 100 (UK) sites to each other. Run BGP as the PE-CE protocol between AS 1000 and AS 100 (US & UK).

R13

```
router eigrp 100
 network 13.0.0.0
 network 192.1.134.0
!
mpls ldp router-id loopback0
!
Interface E0/1
 mpls ip
!
router bgp 1000
 neighbor 14.14.14.14 remote-as 1000
 neighbor 14.14.14.14 update-source Loopback0
!
```

```
 address-family vpnv4
  neighbor 14.14.14.14 activate
!
vrf definition SP-100
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E0/0
 vrf forwarding SP-100
 ip address 192.1.40.13 255.255.255.0
 no shut
!
router bgp 1000
 address-family ipv4 vrf SP-100
  neighbor 192.1.40.4 remote-as 100
```

R14

```
router eigrp 100
 network 13.0.0.0
 network 192.1.134.0
!
mpls ldp router-id loopback0
!
Interface E0/1
 mpls ip
!
router bgp 1000
 neighbor 13.13.13.13 remote-as 1000
 neighbor 13.13.13.13 update-source Loopback0
!
 address-family vpnv4
  neighbor 13.13.13.13 activate
!
vrf definition SP-100
 rd 1000:1
 address-family ipv4
  route-target both 1000:1
!
Interface E0/0
 vrf forwarding SP-100
 ip address 192.1.80.14 255.255.255.0
 no shut
!
```

router bgp 1000 address-family ipv4 vrf SP-100   neighbor 192.1.80.8 remote-as 100
R4 Interface E 0/2 ip address 192.1.40.4 255.255.255.0 no shut ! router bgp 100 neighbor 192.1.40.13 remote-as 1000 neighbor 192.1.40.13 allowas-in redistribute ospf 1 ! router ospf 1 redistribute bgp 100
R8 Interface E 0/2 ip address 192.1.80.8 255.255.255.0 no shut ! router bgp 100 neighbor 192.1.80.14 remote-as 1000 neighbor 192.1.80.14 allowas-in redistribute isis ! router isis redistribute bgp 100

Verification:

- Make sure you have end-to-end reachability between the 2 AS 100 Sites (US & UK) by using a Ping between R1 – R8 Loopback interfaces.
- Use the Traceroute command to verify that the packets are IP Packets within the AS 100 Sites.

Lab 2 – CSC – MPLS Unicast Routing

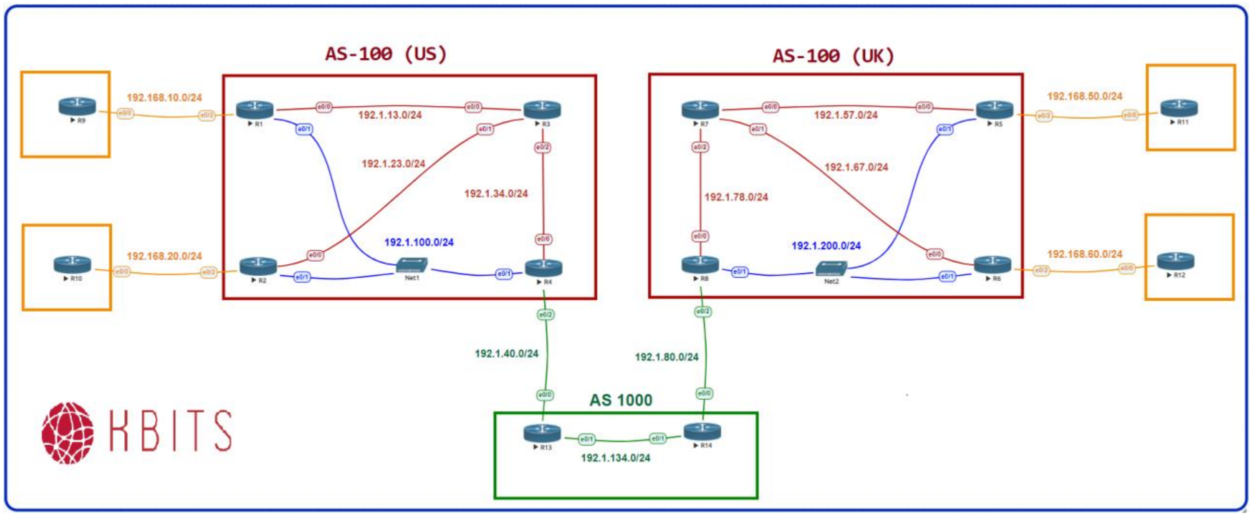

Task 1

Configure LDP within the AS 100 US Site.

R1	R2
mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip	mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip
R3 mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip ! Interface E 0/2	**R4** mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip

Task 2

Configure LDP within the AS 100 UK Site.

R5 mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip	**R6** mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip
R7 mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip ! Interface E 0/2	**R8** mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip

Task 3

Configure the PE-CE Routing such that all routes exchanged via BGP get assigned a label by BGP.

R4
 	
router bgp 100
  neighbor 192.1.40.13 send-label
	R8
	
	router bgp 100
	  neighbor 192.1.80.14 send-label

R13
 	
router bgp 1000
 address-family ipv4 vrf SP-100
  neighbor 192.1.40.4 send-label
	R14
	
	router bgp 1000
	 address-family ipv4 vrf SP-100
	  neighbor 192.1.80.8 send-label

Verification:

- Make sure you have end-to-end reachability between the 2 AS 100 Sites (US & UK) by using a Ping between R1 – R8 Loopback interfaces.

- Use the Traceroute command to verify that the packets are labeled Packets end-to-end.

Lab 3 – CSC – MPLS VPN (H-MPLS VPN)

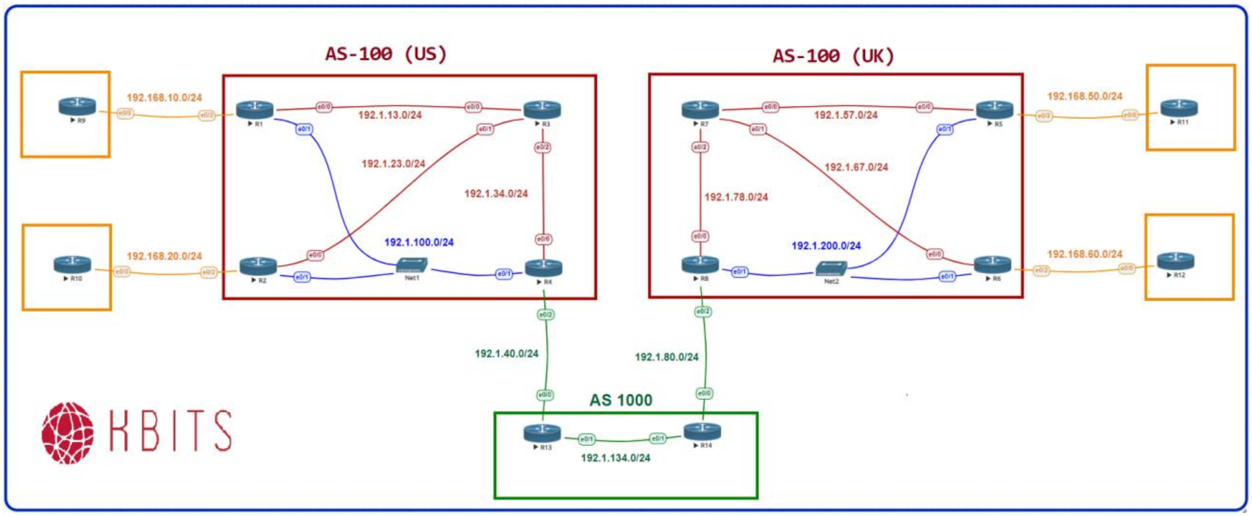

Task 1

Configure an Intra-AS MPLS VPN within AS 100 US site to connect Cust-A Sites (R9 & R10).

R1	R2

router bgp 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 3.3.3.3 update-source lo 0
!
 address-family vpnv4
  neighbor 3.3.3.3 activate
!
vrf definition Cust-A
 rd 100:1
 address-family ipv4
  route-target both 100:1
!
interface E0/2
 vrf forwarding Cust-A
 ip address 192.168.10.1 255.255.255.0
 no shut
!
router bgp 100
 address-family ipv4 vrf Cust-A
  neighbor 192.168.10.9 remote-as 65009
	router bgp 100
	 neighbor 3.3.3.3 remote-as 100
	 neighbor 3.3.3.3 update-source lo 0
	!
	 address-family vpnv4
	  neighbor 3.3.3.3 activate
	!
	vrf definition Cust-A
	 rd 100:1
	 address-family ipv4
	  route-target both 100:1
	!
	interface E0/2
	 vrf forwarding Cust-A
	 ip address 192.168.20.2 255.255.255.0
	 no shut
	!
	router bgp 100
	 address-family ipv4 vrf Cust-A
	  neighbor 192.168.20.10 remote-as 65010

R3	

router bgp 100
 neighbor IBGP peer-group
 neighbor IBGP remote-as 100
 neighbor IBGP update-source lo 0
 neighbor 1.1.1.1 peer-group IBGP
 neighbor 2.2.2.2 peer-group IBGP
!
 address-family vpnv4
  neighbor IBGP route-reflector-client
  neighbor 1.1.1.1 activate
  neighbor 2.2.2.2 activate	

R9 router bgp 65009 neighbor 192.168.10.1 remote-as 100 network 10.9.9.0 mask 255.255.255.0	**R10** router bgp 65010 neighbor 192.168.20.2 remote-as 100 network 10.10.10.0 mask 255.255.255.0

Task 2

Configure an Intra-AS MPLS VPN within AS 100 UK site to connect Cust-A Sites (R1 & R12).

R5	**R6**

router bgp 100
 neighbor 7.7.7.7 remote-as 100
 neighbor 7.7.7.7 update-source lo0
!
 address-family vpnv4
  neighbor 7.7.7.7 activate
!
vrf definition Cust-A
 rd 100:1
 address-family ipv4
  route-target both 100:1
!
interface E0/2
 vrf forwarding Cust-A
 ip address 192.168.50.5 255.255.255.0
 no shut
!
router bgp 100
 address-family ipv4 vrf Cust-A
  neighbor 192.168.50.11 remote-as 65011
	router bgp 100
	 neighbor 7.7.7.7 remote-as 100
	 neighbor 7.7.7.7 update-source lo0
	!
	 address-family vpnv4
	  neighbor 7.7.7.7 activate
	!
	vrf definition Cust-A
	 rd 100:1
	 address-family ipv4
	  route-target both 100:1
	!
	interface E0/2
	 vrf forwarding Cust-A
	 ip address 192.168.60.6 255.255.255.0
	 no shut
	!
	router bgp 100
	 address-family ipv4 vrf Cust-A
	  neighbor 192.168.60.12 remote-as 65012

R7	

R7
 	
router bgp 100
 neighbor IBGP peer-group
 neighbor IBGP remote-as 100
 neighbor IBGP update-source lo 0
 neighbor 5.5.5.5 peer-group IBGP
 neighbor 6.6.6.6 peer-group IBGP
!
 address-family vpnv4
  neighbor IBGP route-reflector-client
  neighbor 5.5.5.5 activate
  neighbor 6.6.6.6 activate	

R11 router bgp 65011 neighbor 192.168.50.5 remote-as 100 network 10.11.11.0 mask 255.255.255.0	**R12** router bgp 65012 neighbor 192.168.60.6 remote-as 100 network 10.12.12.0 mask 255.255.255.0

Task 3

Configure a MP-iBGP relationship between the R3 & R7, the Local Site RRs.

R3	R6

router bgp 100
 neighbor 7.7.7.7 remote-as 100
 neighbor 7.7.7.7 update-source loo 0
!
 address-family vpnv4
  neighbor 7.7.7.7 activate
  neighbor 7.7.7.7 route-reflector-client
	router bgp 100
	 neighbor 3.3.3.3 remote-as 100
	 neighbor 3.3.3.3 update-source lo0
	!
	 address-family vpnv4
	  neighbor 3.3.3.3 activate
	  neighbor 3.3.3.3 route-reflector-client

Verification:

Make sure you have end-to-end reachability between the 4 Cust-A Sites (R9 – R12).

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring MPLS on IOS-XR

Lab 1 – Configuring SP Core Networks with MPLS Unicast Routing & MP-iBGP

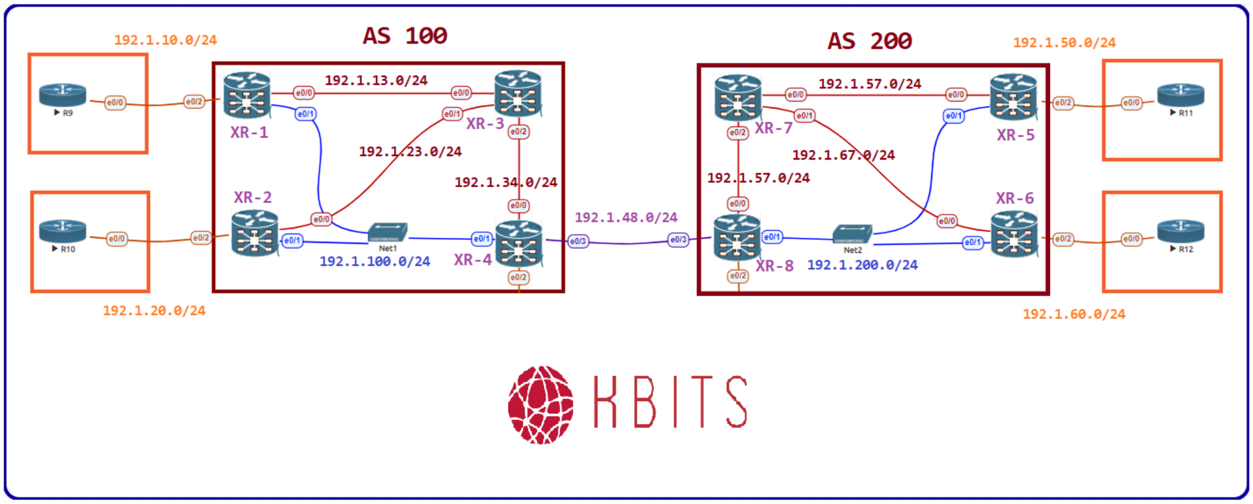

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
Gig0/0/0/0	192.1.13.1	255.255.255.0
Gig0/0/0/1	192.1.100.1	255.255.255.0
Gig0/0/0/2	192.1.10.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
Gig0/0/0/0	192.1.23.2	255.255.255.0
Gig0/0/0/1	192.1.100.2	255.255.255.0
Gig0/0/0/2	192.1.20.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
G 0/0/0/0	192.1.13.3	255.255.255.0
G 0/0/0/1	192.1.23.3	255.255.255.0
G 0/0/0/2	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
G 0/0/0/0	192.1.34.4	255.255.255.0
G 0/0/0/1	192.1.100.4	255.255.255.0
G 0/0/0/2	192.1.48.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.255.255.255
G 0/0/0/0	192.1.57.5	255.255.255.0
G 0/0/0/1	192.1.200.5	255.255.255.0
G 0/0/0/2	192.1.50.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	6.6.6.6	255.255.255.0
G 0/0/0/0	192.1.67.6	255.255.255.0
G 0/0/0/1	192.1.200.6	255.255.255.0
G 0/0/0/2	192.1.60.6	255.255.255.0

R7

Interface	IP Address	Subnet Mask
Loopback 0	7.7.7.7	255.255.255.255
G 0/0/0/0	192.1.57.7	255.255.255.0
G 0/0/0/1	192.1.67.7	255.255.255.0
G 0/0/0/2	192.1.78.7	255.255.255.0

R8

Interface	IP Address	Subnet Mask
Loopback 0	8.8.8.8	255.255.255.255
G 0/0/0/0	192.1.34.4	255.255.255.0
G 0/0/0/1	192.1.200.8	255.255.255.0
G 0/0/0/2	192.1.48.8	255.255.255.0

R9

Interface	IP Address	Subnet Mask
Loopback 0	10.9.9.9	255.255.255.0
E 0/0	192.1.10.9	255.255.255.0

R10

Interface	IP Address	Subnet Mask
Loopback 0	10.10.10.10	255.255.255.0
E 0/0	192.1.20.10	255.255.255.0

R11

Interface	IP Address	Subnet Mask
Loopback 0	10.11.11.11	255.255.255.0
E 0/0	192.1.50.11	255.255.255.0

R12

Interface	IP Address	Subnet Mask
Loopback 0	10.12.12.12	255.255.255.0
E 0/0	192.1.60.12	255.255.255.0

AS 100

Task 1

Configure IP Addresses on R1, R2, R3 & R4 based on the above table. Run OSPF as the IGP for AS 100.

R1
 	
hostname XR1
!
Interface Gig0/0/0/0
 ip address 192.1.13.1 255.255.255.0
 no shut
!
Interface Gig0/0/0/1
 ip address 192.1.100.1 255.255.255.0
 no shut
!
Interface loopback0
 ip address 1.1.1.1 255.255.255.255
!
router ospf 1
 router-id 0.0.0.1
 area 0
  Interface Gig0/0/0/0
    exit
  Interface Gig0/0/0/1
    exit
  Interface Loopback0
    exit
!
commit
	R2
	
	hostname XR2
	!
	Interface Gig0/0/0/0
	 ip address 192.1.23.2 255.255.255.0
	 no shut
	!
	Interface Gig0/0/0/1
	 ip address 192.1.100.2 255.255.255.0
	 no shut
	!
	Interface loopback0
	 ip address 2.2.2.2 255.255.255.255
	!
	router ospf 1
	 router-id 0.0.0.2
	 area 0
	  Interface Gig0/0/0/0
	    exit
	  Interface Gig0/0/0/1
	    exit
	  Interface Loopback0
	    exit
	!
	commit

R3 hostname XR3 ! Interface Gig0/0/0/0 ip address 192.1.13.3 255.255.255.0 no shut ! Interface Gig0/0/0/1 ip address 192.1.23.3 255.255.255.0 no shut !	**R4** hostname XR4 ! Interface Gig0/0/0/0 ip address 192.1.34.4 255.255.255.0 no shut ! Interface Gig0/0/0/1 ip address 192.1.100.4 255.255.255.0 no shut !

Interface Gig0/0/0/2
 ip address 192.1.34.3 255.255.255.0
 no shut
!
Interface loopback0
 ip address 3.3.3.3 255.255.255.255
!
router ospf 1
 router-id 0.0.0.3
 area 0
  Interface Gig0/0/0/0
    exit
  Interface Gig0/0/0/1
    exit
  Interface Gig0/0/0/2
    exit
  Interface Loopback0
    exit
!
commit
	Interface loopback0
	 ip address 4.4.4.4 255.255.255.255
	!
	router ospf 1
	 router-id 0.0.0.4
	 area 0
	  Interface Gig0/0/0/0
	    exit
	  Interface Gig0/0/0/1
	    exit
	  Interface Loopback0
	    exit
	!
	commit

Task 2

Configure LDP for MPLS Unicast Routing using Loopback 0 as the Router-id.

R1
 	
mpls ldp
 router-id 1.1.1.1
 interface gig0/0/0/0
  exit
 interface gig0/0/0/1
  exit
!
commit
	R2
	
	mpls ldp
	 router-id 2.2.2.2
	 interface gig0/0/0/0
	  exit
	 interface gig0/0/0/1
	  exit
	!
	commit

R3
 	
mpls ldp
 router-id 3.3.3.3
 interface gig0/0/0/0
  exit
 interface gig0/0/0/1
  exit
interface gig0/0/0/2
  exit
!
commit
	R4
	
	mpls ldp
	 router-id 4.4.4.4
	 interface gig0/0/0/0
	  exit
	 interface gig0/0/0/1
	  exit
	!
	commit

Task 3

Configure MP-iBGP between R1, R2 & R3. Configure R3 to be the RR for R1 & R2 for VPNv4 routes.

<table>
<tr><td>

R1

```
router bgp 100
 address-family vpnv4 unicast
  exit
 !
 neighbor 3.3.3.3
  remote-as 100
  update-source Loopback0
  address-family vpnv4 unicast
!
commit
```

</td><td>

R2

```
router bgp 100
 address-family vpnv4 unicast
  exit
 !
 neighbor 3.3.3.3
  remote-as 100
  update-source Loopback0
  address-family vpnv4 unicast
!
commit
```

</td></tr>
<tr><td>

R3

```
router bgp 100
 address-family vpnv4 unicast
  exit
 !
 neighbor-group MP-iBGP
  remote-as 100
  update-source Loopback0
  address-family vpnv4 unicast
   route-reflector-client
     exit
      exit
  !
 neighbor 1.1.1.1
  use neighbor-group MP-iBGP
   exit
 !
 neighbor 2.2.2.2
  use neighbor-group MP-iBGP
   exit
!
 commit
```

</td><td></td></tr>
</table>

Verification:

Verify the neighbor relationship on XR3 by using the "sh bgp vpnv4 unicast neighbors" command.

AS 200

Task 1

Configure IP Addresses on R5, R6, R7 & R8 based on the above table. Run IS-IS as the IGP for AS 200.

R5	R6

hostname XR5
!
Interface Gig0/0/0/0
 ip address 192.1.57.5 255.255.255.0
 no shut
!
Interface Gig0/0/0/1
 ip address 192.1.200.5 255.255.255.0
 no shut
!
Interface loopback0
 ip address 5.5.5.5 255.255.255.255
!
router isis 1
 net 49.0000.5555.5555.5555.00
 is-type level-2-only
 address-family ipv4 unicast
  metric-style wide
   exit
!
 Interface Gig0/0/0/0
  address-family ipv4 unicast
   exit
  exit
!
 Interface Gig0/0/0/1
  address-family ipv4 unicast
   exit
  exit
!
 Interface Loopback0
  address-family ipv4 unicast
   exit
  exit
!
Commit
	hostname XR6
	!
	Interface Gig0/0/0/0
	 ip address 192.1.67.6 255.255.255.0
	 no shut
	!
	Interface Gig0/0/0/1
	 ip address 192.1.200.6 255.255.255.0
	 no shut
	!
	Interface loopback0
	 ip address 6.6.6.6 255.255.255.255
	!
	router isis 1
	 net 49.0000.6666.6666.6666.00
	 is-type level-2-only
	 address-family ipv4 unicast
	  metric-style wide
	   exit
	!
	 Interface Gig0/0/0/0
	  address-family ipv4 unicast
	   exit
	  exit
	!
	 Interface Gig0/0/0/1
	  address-family ipv4 unicast
	   exit
	  exit
	!
	 Interface Loopback0
	  address-family ipv4 unicast
	   exit
	  exit
	!
	Commit

R7	R8

hostname XR7
!
Interface Gig0/0/0/0
 ip address 192.1.57.7 255.255.255.0
 no shut
!
Interface Gig0/0/0/1
 ip address 192.1.67.7 255.255.255.0
 no shut
!
Interface Gig0/0/0/2
 ip address 192.1.78.7 255.255.255.0
 no shut
!
Interface loopback0
 ip address 7.7.7.7 255.255.255.255
!
router isis 1
 net 49.0000.7777.7777.7777.00
 is-type level-2-only
 address-family ipv4 unicast
  metric-style wide
   exit
!
 Interface Gig0/0/0/0
  address-family ipv4 unicast
   exit
  exit
!
 Interface Gig0/0/0/1
  address-family ipv4 unicast
   exit
  exit
!
 Interface Gig0/0/0/2
  address-family ipv4 unicast
   exit
  exit
!
 Interface Loopback0
  address-family ipv4 unicast
   exit
  exit
	hostname XR8
	!
	Interface Gig0/0/0/0
	 ip address 192.1.78.8 255.255.255.0
	 no shut
	!
	Interface Gig0/0/0/1
	 ip address 192.1.200.8 255.255.255.0
	 no shut
	!
	Interface loopback0
	 ip address 8.8.8.8 255.255.255.255
	!
	router isis 1
	 net 49.0000.8888.8888.8888.00
	 is-type level-2-only
	 address-family ipv4 unicast
	  metric-style wide
	   exit
	!
	 Interface Gig0/0/0/0
	  address-family ipv4 unicast
	   exit
	  exit
	!
	 Interface Gig0/0/0/1
	  address-family ipv4 unicast
	   exit
	  exit
	!
	 Interface Loopback0
	  address-family ipv4 unicast
	   exit
	  exit
	!
	Commit

! commit	

Task 2

Configure LDP for MPLS Unicast Routing using Loopback 0 as the Router-id.

R5	**R6**

mpls ldp
 router-id 5.5.5.5
 interface gig0/0/0/0
  exit
 interface gig0/0/0/1
  exit
!
commit
	mpls ldp
	 router-id 6.6.6.6
	 interface gig0/0/0/0
	  exit
	 interface gig0/0/0/1
	  exit
	!
	commit

R7
 	
mpls ldp
 router-id 7.7.7.7
 interface gig0/0/0/0
  exit
 interface gig0/0/0/1
  exit
interface gig0/0/0/2
  exit
!
Commit
	R8
	
	mpls ldp
	 router-id 8.8.8.8
	 interface gig0/0/0/0
	  exit
	 interface gig0/0/0/1
	  exit
	!
	commit

Task 3

Configure MP-iBGP between R5, R6 & R7. Configure R7 to be the RR for R5 & R6 for VPNv4 routes.

<table>
<tr>
<td>

R5

```
router bgp 200
 address-family vpnv4 unicast
  exit
 !
 neighbor 7.7.7.7
  remote-as 200
  update-source Loopback0
  address-family vpnv4 unicast
!
commit
```

</td>
<td>

R6

```
router bgp 200
 address-family vpnv4 unicast
  exit
 !
 neighbor 7.7.7.7
  remote-as 200
  update-source Loopback0
  address-family vpnv4 unicast
!
commit
```

</td>
</tr>
<tr>
<td>

R7

```
router bgp 200
 address-family vpnv4 unicast
  exit
 !
 neighbor-group MP-iBGP
  remote-as 200
  update-source Loopback0
  address-family vpnv4 unicast
   route-reflector-client
     exit
      exit
  !
 neighbor 5.5.5.5
  use neighbor-group MP-iBGP
   exit
 !
 neighbor 6.6.6.6
  use neighbor-group MP-iBGP
   exit
!
 Commit
```

</td>
<td></td>
</tr>
</table>

Verification:

Verify the neighbor relationship on XR7 by using the “sh bgp vpnv4 unicast neighbors” command.

Lab 2 – Configuring Intra-AS MPLS VPN within AS 100 & AS 200

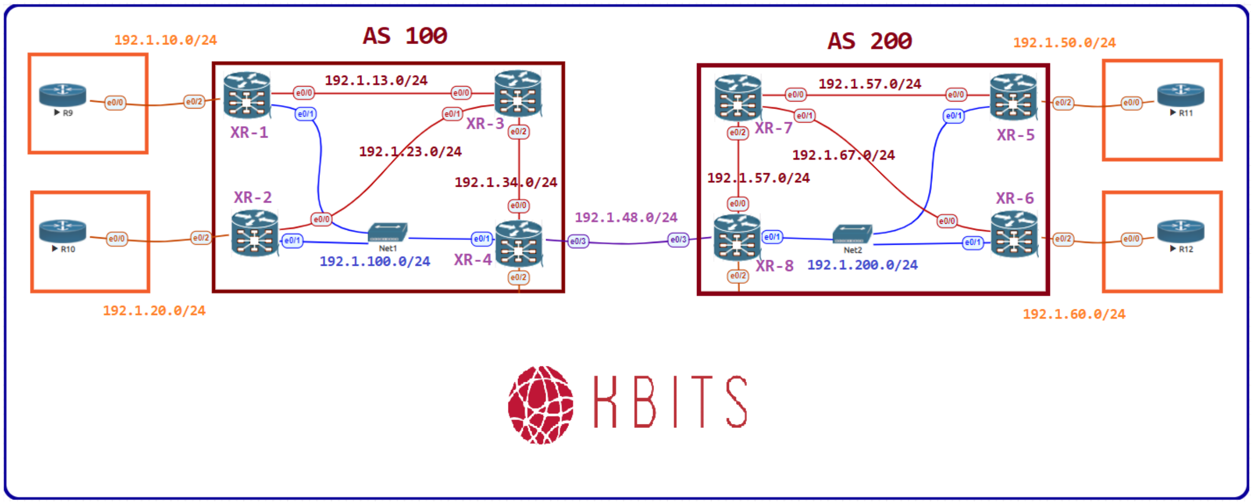

AS 100 – Configuring Intra-AS MPLS VPN to connect R9 & R10

Task 1

Configure a VRF Cust-A with a RD value of 100:1 on R1 and R2. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R1 and R2.

R1	**R2**
vrf Cust-A address-family ipv4 unicast import route-target 100:1 exit export route-target 100:1 exit commit ! Interface gig0/0/0/2 vrf Cust-A ip address 192.1.10.1 255.255.255.0 no shut ! Commit	vrf Cust-A address-family ipv4 unicast import route-target 100:1 exit export route-target 100:1 exit commit ! Interface gig0/0/0/2 vrf Cust-A ip address 192.1.20.2 255.255.255.0 no shut ! Commit

Task 2
Configure BGP as the PE-CE Routing Protocol on the PE Routers. Use 65009 as the AS # for Site 1 and 65010 as the AS # for Site 2.

R1	**R2**

route-policy PASSALL
 pass
 exit
!
router bgp 100
 address-family ipv4 unicast
  exit
 vrf CUST-A
  rd 100:1
  address-family ipv4 unicast
   exit
  neighbor 192.1.10.9
   remote-as 65009
   address-family ipv4 unicast
    route-policy PASSALL in
    route-policy PASSALL out
!
commit
	route-policy PASSALL
	 pass
	 exit
	!
	router bgp 100
	 address-family ipv4 unicast
	  exit
	 vrf CUST-A
	  rd 100:1
	  address-family ipv4 unicast
	   exit
	  neighbor 192.1.20.10
	   remote-as 65010
	   address-family ipv4 unicast
	    route-policy PASSALL in
	    route-policy PASSALL out
	!
	commit

Task 3
Configure BGP as the PE-CE Routing Protocol on the CE Routers. Use 65009 as the AS # for Site 1 and 65010 as the AS # for Site 2. Advertise the Loopback in BGP.

R9	**R10**
Interface Loopback0 ip address 10.9.9.9 255.255.255.0 ! Interface E 0/0 ip address 192.1.10.9 255.255.255.0 no shut ! router bgp 65009 network 10.9.9.0 mask 255.255.255.0 neighbor 192.1.10.1 remote-as 100	Interface Loopback0 ip address 10.10.10.10 255.255.255.0 ! Interface E 0/0 ip address 192.1.20.10 255.255.255.0 no shut ! router bgp 65010 network 10.10.10.0 mask 255.255.255.0 neighbor 192.1.20.2 remote-as 100

AS 200 – Configuring Intra-AS MPLS VPN to connect R11 & R12

Task 1

Configure a VRF Cust-A with a RD value of 200:1 on R5 and R6. Use the same extended community for your Route-target import and export. Assign this VRF to the links that connect to Cust-A sites on R5 and R6.

R1	**R2**

vrf CUST-A
 address-family ipv4 unicast
   import route-target
    200:1
    exit
   export route-target
    200:1
     exit
  commit
!
Interface gig0/0/0/2
 vrf CUST-A
 ip address 192.1.50.5 255.255.255.0
 no shut
!
Commit
	vrf CUST-A
	 address-family ipv4 unicast
	   import route-target
	    200:1
	    exit
	   export route-target
	    200:1
	     exit
	  commit
	!
	Interface gig0/0/0/2
	 vrf CUST-A
	 ip address 192.1.60.6 255.255.255.0
	 no shut
	!
	Commit

Task 2

Configure BGP as the PE-CE Routing Protocol on the PE Routers. Use 65011 as the AS # for Site 3 and 65012 as the AS # for Site 4.

R5	**R6**

route-policy PASSALL
 pass
 exit
!
router bgp 200
 address-family ipv4 unicast
   exit
 vrf CUST-A
  rd 200:1
  address-family ipv4 unicast
   exit
  neighbor 192.1.50.11
	route-policy PASSALL
	 pass
	 exit
	!
	router bgp 200
	 address-family ipv4 unicast
	   exit
	 vrf CUST-A
	  rd 200:1
	  address-family ipv4 unicast
	   exit
	  neighbor 192.1.60.12

remote-as 65011 address-family ipv4 unicast route-policy PASSALL in route-policy PASSALL out ! commit	remote-as 65012 address-family ipv4 unicast route-policy PASSALL in route-policy PASSALL out ! commit

Task 3

Configure BGP as the PE-CE Routing Protocol on the CE Routers. Use 65009 as the AS # for Site 1 and 65010 as the AS # for Site 2. Advertise the Loopback in BGP.

R11	**R12**
Interface Loopback0 ip address 10.11.11.11 255.255.255.0 ! Interface E 0/0 ip address 192.1.50.11 255.255.255.0 no shut ! router bgp 65011 network 10.11.11.0 mask 255.255.255.0 neighbor 192.1.50.5 remote-as 200	Interface Loopback0 ip address 10.12.12.12 255.255.255.0 ! Interface E 0/0 ip address 192.1.60.12 255.255.255.0 no shut ! router bgp 65012 network 10.12.12.0 mask 255.255.255.0 neighbor 192.1.60.6 remote-as 200

Verification:

- Verify the connectivity between R9 & R10 in AS 100.
- Verify the connectivity between R11 & r12 in AS 200.

Lab 3 – Configuring Inter-AS MPLS VPN – Option C (MP-eBGP between RRs)

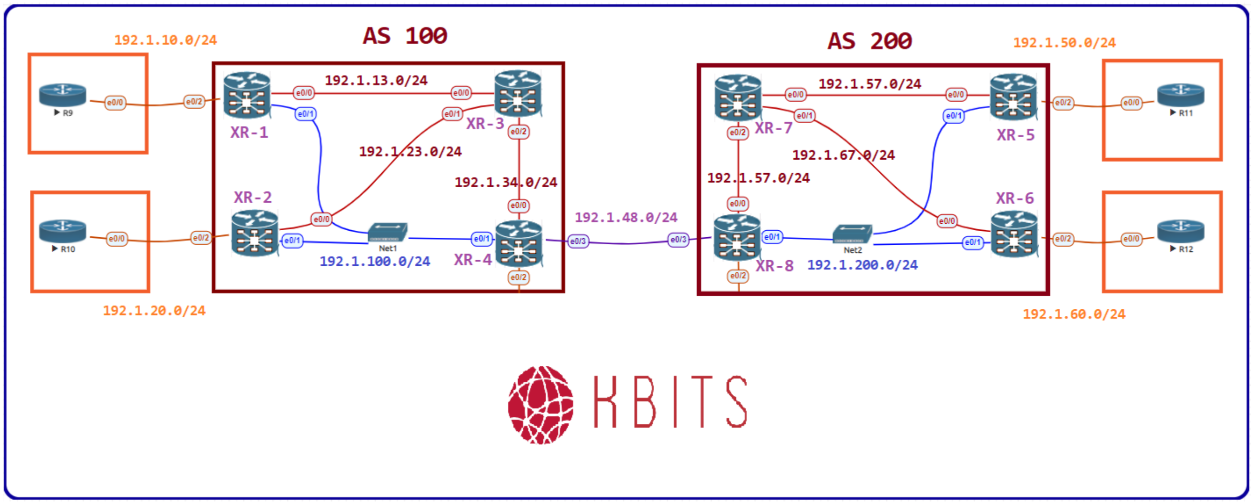

Task 1
Configure the Interface between the ASBRs.

R4

```
Interface Gig 0/0/0/2
 ip address 192.1.48.4 255.255.255.0
 no shut
!
commit
```

R8

```
Interface Gig 0/0/0/2
 ip address 192.1.48.8 255.255.255.0
 no shut
!
commit
```

Task 2
Configure the Route Policy for Route Leaking.

R4

```
route-policy O2B
 if destination in (1.1.1.1/32,2.2.2.2/32,3.3.3.3/32) then
```

pass
endif
end-policy
!
route-policy B2O
if destination in (5.5.5.5/32,6.6.6.6/32,7.7.7.7/32) then
pass
endif
end-policy
!
route-policy PASSALL
pass
exit
!
Commit

R8
 	
route-policy B2I
if destination in (1.1.1.1/32,2.2.2.2/32,3.3.3.3/32) then
pass
endif
end-policy
!
route-policy I2B
if destination in (5.5.5.5/32,6.6.6.6/32,7.7.7.7/32) then
pass
endif
end-policy
!
route-policy PASSALL
pass
exit
!
commit

Task 3

Configure BGP between the ASBRs to provide underlay routing for MP-eBGP.

R4	**R8**

router bgp 100
 address-family ipv4 unicast
  redistribute ospf 1 route-policy O2B
  allocate-label all
    exit
!
 neighbor 192.1.48.8
  remote-as 200
  address-family ipv4 labeled-unicast
   route-policy PASSALL in
   route-policy PASSALL out
   exit
  exit
exit
!
commit
	router bgp 200
	 address-family ipv4 unicast
	  redistribute isis 1 route-policy I2B
	  allocate-label all
	    exit
	!
	 neighbor 192.1.48.4
	  remote-as 100
	  address-family ipv4 labeled-unicast
	   route-policy PASSALL in
	   route-policy PASSALL out
	   exit
	  exit
	exit
	!
	commit

Task 4

Configure Redistribution of BGP routes into IGP.

R4 router ospf 1 redistribute bgp 100 route-policy B2O ! router static address-family ipv4 unicast   192.1.48.8/32 GigabitEthernet0/0/0/2 ! commit! router ospf 1 redistribute bgp 100 exit commit

R8
 	
router isis 1
 address-family ipv4 unicast
  redistribute bgp 200 route-policy B2I
  exit
!
router static
 address-family ipv4 unicast
  192.1.48.4/32 GigabitEthernet0/0/0/2
!
commit

Task 5

Configure MP-eBGP between the RRs to exchange the VPNv4 routes.

R3
 	
route-policy PASSALL
 pass
 exit
!
router bgp 100
 neighbor 7.7.7.7
  remote-as 200
  update-source Loopback00
  ebgp-multihop
	R8
	
	route-policy PASSALL
	 pass
	 exit
	!
	router bgp 200
	 neighbor 3.3.3.3
	  remote-as 100
	  update-source Loopback00
	  ebgp-multihop

address-family vpnv4 unicast route-policy PASSALL in route-policy PASSALL out next-hop-unchanged exit exit exit ! commit	address-family vpnv4 unicast route-policy PASSALL in route-policy PASSALL out next-hop-unchanged exit exit exit ! commit

Task 6

Add the 200:1 route-target as a import route-target on all PE Routers (R1 and R2) for VRF Cust-A in AS 100. Add the 100:1 route-target as a import route-target on all PE Routers (R5 and R6) for VRF Cust-A in AS 200.

R1

```
vrf Cust-A
 address-family ipv4 unicast
   import route-target
    200:1
!
commit
```

R2

```
vrf Cust-A
 address-family ipv4 unicast
   import route-target
    200:1
!
commit
```

R5

```
vrf Cust-A
 address-family ipv4 unicast
   import route-target
    100:1
!
commit
```

R6

```
vrf Cust-A
 address-family ipv4 unicast
   import route-target
    100:1
!
commit
```

Verification:

- Verify the connectivity between the 4 sites. R9 & R10 should be able to reach R11 & R12 Loopbacks and vice versa.

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Segment Routing and Large Scale MPLS

Lab 1 – Basic Intra-AS MPLS VPN using LDP

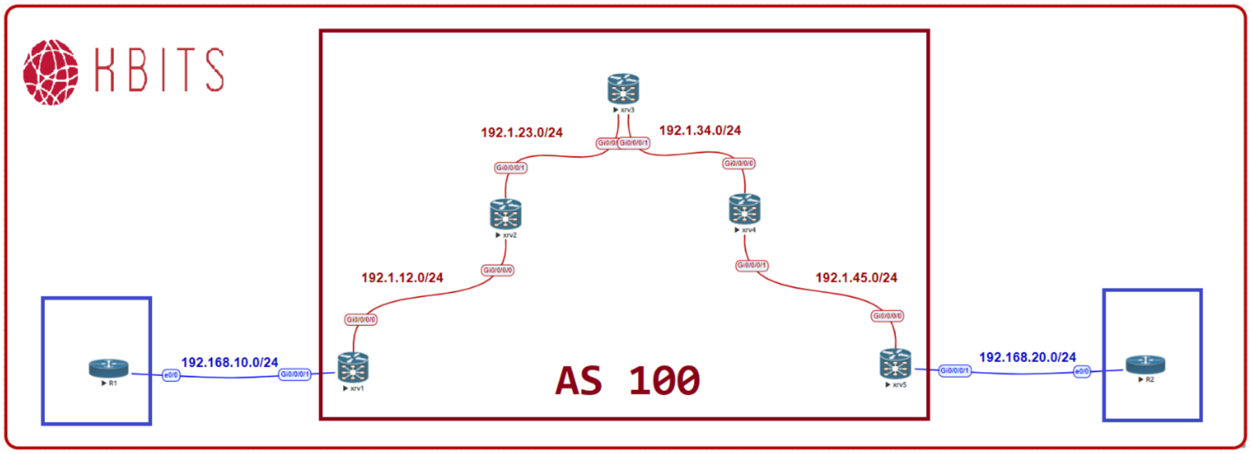

Interface IP Address Configuration

XR1

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.1	255.255.255.255
Gig0/0/0/0	192.1.12.1	255.255.255.0
Gig0/0/0/1	192.168.10.1	255.255.255.0

XR2

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.2	255.255.255.255
Gig0/0/0/0	192.1.12.2	255.255.255.0
Gig0/0/0/1	192.1.23.2	255.255.255.0

XR3

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.3	255.255.255.255
G 0/0/0/0	192.1.23.3	255.255.255.0
G 0/0/0/1	192.1.34.3	255.255.255.0

XR4

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.4	255.255.255.255
G 0/0/0/0	192.1.34.4	255.255.255.0
G 0/0/0/1	192.1.45.4	255.255.255.0

XR5

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.5	255.255.255.255
G 0/0/0/0	192.1.45.5	255.255.255.0
G 0/0/0/1	192.168.20.5	255.255.255.0

R1

Interface	IP Address	Subnet Mask
Loopback 0	10.11.11.11	255.255.255.0
G 0/0/0/0	192.168.10.11	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	10.22.22.22	255.255.255.255
G 0/0/0/0	192.168.20.22	255.255.255.0

AS 100

Task 1

Configure IP Addresses on R1, R2, R3, R4 & R5 based on the above table. Run OSPF as the IGP for AS 100.

XR1
 	
hostname XR1
!
interface gig0/0/0/0
 ip address 192.1.12.1 255.255.255.0
 no shut
!
interface loopback0
 ip address 10.1.1.1 255.255.255.255
!
router ospf 1
 router-id 0.0.0.1
 area 0
  interface gig0/0/0/0
   exit
  interface Loopback0
   exit
commit
	XR2
	
	hostname XR2
	!
	interface gig0/0/0/0
	 ip address 192.1.12.2 255.255.255.0
	 no shut
	!
	interface gig0/0/0/1
	 ip address 192.1.23.2 255.255.255.0
	 no shut
	!
	interface loopback0
	 ip address 10.1.1.2 255.255.255.255
	!
	router ospf 1
	 router-id 0.0.0.2
	 area 0
	  interface gig0/0/0/0
	   exit
	  interface gig0/0/0/1
	   exit
	  interface Loopback0
	   exit
	commit

XR3 hostname XR3 ! interface gig0/0/0/0 ip address 192.1.23.3 255.255.255.0 no shut ! interface gig0/0/0/1 ip address 192.1.34.3 255.255.255.0 no shut ! interface loopback0	**XR4** hostname XR4 ! interface gig0/0/0/0 ip address 192.1.34.4 255.255.255.0 no shut ! interface gig0/0/0/1 ip address 192.1.45.4 255.255.255.0 no shut ! interface loopback0

ip address 10.1.1.3 255.255.255.255 ! router ospf 1 router-id 0.0.0.3 area 0 interface gig0/0/0/0 exit interface gig0/0/0/1 exit interface Loopback0 exit commit	ip address 10.1.1.4 255.255.255.255 ! router ospf 1 router-id 0.0.0.4 area 0 interface gig0/0/0/0 exit interface gig0/0/0/1 exit interface Loopback0 exit Commit
XR5 hostname XR5 ! interface gig0/0/0/0 ip address 192.1.45.5 255.255.255.0 no shut ! interface loopback0 ip address 10.1.1.5 255.255.255.255 ! router ospf 1 router-id 0.0.0.5 area 0 interface gig0/0/0/0 exit interface Loopback0 exit commit	

Task 2

Configure LDP on all AS 100 routers for MPLS Unicast Routing using Loopback 0 as the Router-id.

XR1
 	
mpls ldp
 router-id 10.1.1.1
 interface gig0/0/0/0
  exit
!
commit
	XR2
	
	mpls ldp
	 router-id 10.1.1.2
	 interface gig0/0/0/0
	  exit
	 interface gig0/0/0/1
	  exit
	!
	commit

XR3
 	
mpls ldp
 router-id 10.1.1.3
 interface gig0/0/0/0
  exit
 interface gig0/0/0/1
  exit
!
commit
	XR4
	
	mpls ldp
	 router-id 10.1.1.4
	 interface gig0/0/0/0
	  exit
	 interface gig0/0/0/1
	  exit
	!
	Commit

XR5 mpls ldp router-id 10.1.1.5 interface gig0/0/0/0 exit ! commit	

Task 3

R1 & R5 are PE Routers. R3 is the RR for VPNv4 routes. Configure MP-iBGP between R1-R3 & R1-R5.

<table>
<tr>
<td>XR1

router bgp 100
 address-family vpnv4 unicast
  exit
 neighbor 10.1.1.3
  remote-as 100
  update-source loopback0
  address-family vpnv4 unicast
 exit
commit</td>
<td>XR5

router bgp 100
 address-family vpnv4 unicast
  exit
 neighbor 10.1.1.3
  remote-as 100
  update-source loopback0
  address-family vpnv4 unicast
 exit
commit</td>
</tr>
<tr>
<td>XR3

router bgp 100
 address-family vpnv4 unicast
  exit
 neighbor-group MP-IBGP
  remote-as 100
  update-source loopback0
  address-family vpnv4 unicast
   route-reflector-client
   exit
 exit
 neighbor 10.1.1.1
  use neighbor-group MP-IBGP
  exit
 neighbor 10.1.1.5
  use neighbor-group MP-IBGP
  exit
!
commit</td>
<td></td>
</tr>
</table>

Task 4

Configure the PE Routers with the VRF for Cust-A using a RD & RT of 100:1. Configure the PE-CE Interface based on the diagram. Run BGP as the PE-CE Routing Protocol. CE's will be in AS 65001 (R1) & AS 65002 (R2) respectively.

XR1	XR5

vrf Cust-A
 address-family ipv4 unicast
  import route-target
    100:1
    exit
  export route-target
   100:1
   exit
commit
!
Interface Gig0/0/0/1
 vrf Cust-A
 ip address 192.168.10.1 255.255.255.0
 no shut
commit
!
route-policy PASSALL
 pass
!
router bgp 100
 address-family ipv4 unicast
  exit
vrf Cust-A
 rd 100:1
 address-family ipv4 unicast
  exit
 neighbor 192.168.10.11
  remote-as 65001
  address-family ipv4 unicast
   route-policy PASSALL in
   route-policy PASSALL out
   exit
 exit
exit
exit
commit
	vrf Cust-A
	 address-family ipv4 unicast
	  import route-target
	    100:1
	    exit
	  export route-target
	   100:1
	   exit
	commit
	!
	Interface Gig0/0/0/1
	 vrf Cust-A
	 ip address 192.168.20.5 255.255.255.0
	 no shut
	commit
	!
	route-policy PASSALL
	 pass
	!
	router bgp 100
	 address-family ipv4 unicast
	  exit
	vrf Cust-A
	 rd 100:1
	 address-family ipv4 unicast
	  exit
	 neighbor 192.168.20.22
	  remote-as 65002
	  address-family ipv4 unicast
	   route-policy PASSALL in
	   route-policy PASSALL out
	   exit
	 exit
	exit
	exit
	commit

Task 5

Configure the CE Routers for PE-CE Routing. Run BGP as the PE-CE Routing Protocol. CE's will be in AS 65001 (R1) & AS 65002 (R2) respectively. Advertise the loopback network on the CE Routers into BGP.

R1	R2
Router bgp 65001 neighbor 192.168.10.1 remote-as 100 network 10.11.11.0 mask 255.255.255.0	Router bgp 65002 neighbor 192.168.20.5 remote-as 100 network 10.22.22.0 mask 255.255.255.0

Verification:

- Verify the connectivity between the 2 sites. R1 & R2 should be able to reach each other's Loopbacks networks.

Lab 2 – Configuring Segment Routing using OSPF

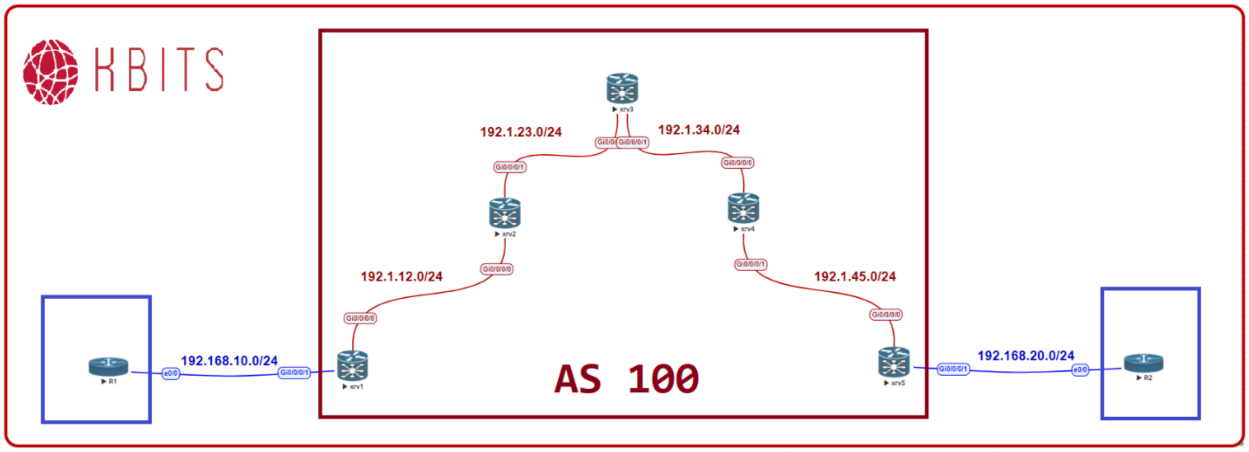

Task 1

Configure the SR label blocks on routers XR1 – XR5.

XR1 segment-routing global-block 16000 23999 commit
XR2 segment-routing global-block 16000 23999 commit
XR3 segment-routing global-block 16000 23999 commit
XR4 segment-routing global-block 16000 23999 commit
XR5 segment-routing global-block 16000 23999 commit

Task 2

Configure OSPF for Segment Routing on Routers XR1 – XR5. Use the loopback0 interface for the Prefix-index. Set the Prefix index based to router#.

XR1

```
router ospf 1
 segment-routing mpls
 area 0
  interface loopback0
  prefix-sid index 1
commit
```

XR2

```
router ospf 1
 segment-routing mpls
 area 0
  interface loopback0
  prefix-sid index 2
commit
```

XR3

```
router ospf 1
 segment-routing mpls
 area 0
  interface loopback0
  prefix-sid index 3
commit
```

XR4

```
router ospf 1
 segment-routing mpls
 area 0
  interface loopback0
  prefix-sid index 4
commit
```

XR5

```
router ospf 1
 segment-routing mpls
 area 0
  interface loopback0
  prefix-sid index 5
commit
```

Verification:

- Use the Traceroute command on R1 to verify the connectivity from R1 Loopback 0 to R2 Loopback0.
- What range of labels are being used in the SP Core?
- It should be using the Non-SR block (24XXX).

Task 3

Configure the core routers to prefer SR over LDP for MPLS Unicast Routing.

XR1

```
router ospf 1
 segment-routing sr-prefer
!
commit
```

XR2

```
router ospf 1
 segment-routing sr-prefer
!
commit
```

XR3

```
router ospf 1
 segment-routing sr-prefer
!
commit
```

XR4

```
router ospf 1
 segment-routing sr-prefer
!
commit
```

XR5

```
router ospf 1
 segment-routing sr-prefer
!
commit
```

Verification:

Use the Traceroute command on R1 to verify the connectivity from R1 Loopback 0 to R2 Loopback0. What range of labels are being used in the SP Core now?

Lab 3 – Configuring Segment Routing using IS-IS

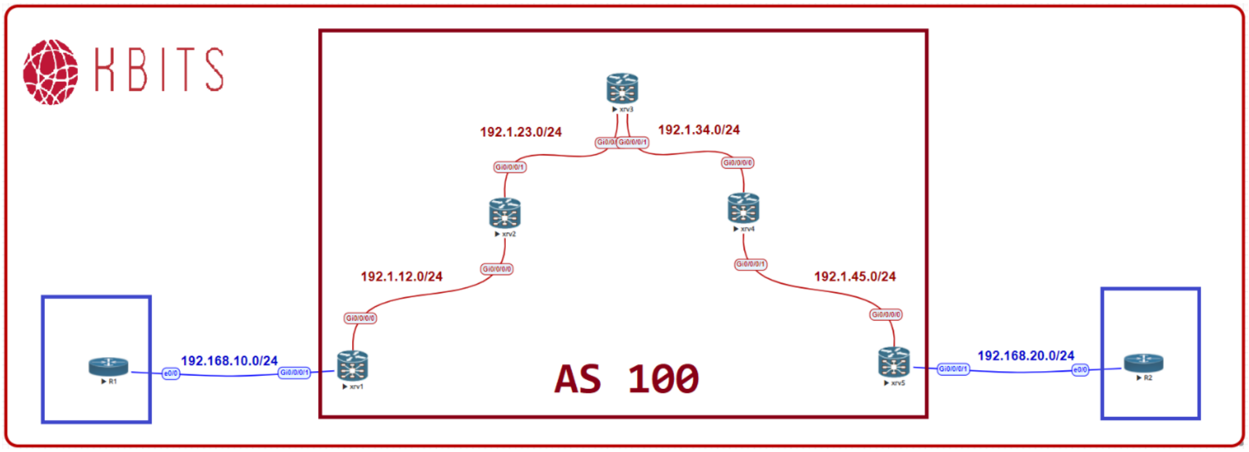

Task 1

Remove OSPF as the Routing Protocol between XR1 – XR5.

XR1 No router ospf 1 Commit
XR2 No router ospf 1 Commit
XR3 No router ospf 1 Commit
XR4 No router ospf 1 Commit
XR5 No router ospf 1 Commit

Task 2

Configure IS-IS as the routing protocol between XR1 – XR5. Use 49.0000 as the area ID with XXXX.XXXX.XXXX as the system-id, where X is the Router #. Enable metric-style wide. Configure the routes as Level-1 routers.

XR1

```
router isis 1
 is-type level-1
 net 49.0000.1111.1111.1111.00
 address-family ipv4 unicast
  metric-style wide
  exit
 interface gig0/0/0/0
  address-family ipv4 unicast
  exit
 interface Loopback0
  address-family ipv4 unicast
  exit
commit
```

XR2

```
router isis 1
 is-type level-1
 net 49.0000.2222.2222.2222.00
 address-family ipv4 unicast
  metric-style wide
  exit
 interface gig0/0/0/0
  address-family ipv4 unicast
  exit
 interface gig0/0/0/1
  address-family ipv4 unicast
  exit
 interface Loopback0
  address-family ipv4 unicast
  exit
commit
```

XR3

```
router isis 1
 is-type level-1
 net 49.0000.3333.3333.3333.00
 address-family ipv4 unicast
```

metric-style wide exit interface gig0/0/0/0 address-family ipv4 unicast exit interface gig0/0/0/1 address-family ipv4 unicast exit interface Loopback0 address-family ipv4 unicast exit commit
XR4 router isis 1 is-type level-1 net 49.0000.4444.4444.4444.00 address-family ipv4 unicast metric-style wide exit interface gig0/0/0/0 address-family ipv4 unicast exit interface gig0/0/0/1 address-family ipv4 unicast exit interface Loopback0 address-family ipv4 unicast exit commit
XR5 router isis 1 is-type level-1 net 49.0000.5555.5555.5555.00 address-family ipv4 unicast metric-style wide exit interface gig0/0/0/0 address-family ipv4 unicast exit interface Loopback0 address-family ipv4 unicast exit

```
commit
```

Task 2

Enable IS-IS for Segment Routing. Use the loopback0 interface for the Prefix-index. Set the Prefix index based to router#. Configure SR as the preferred labeling method.

XR1

```
router isis 1
 address-family ipv4 unicast
  segment-routing mpls sr-prefer
  exit
 interface loopback0
  address-family ipv4 unicast
   prefix-sid index 1
   exit
  exit
 exit
commit
```

XR2

```
router isis 1
 address-family ipv4 unicast
  segment-routing mpls sr-prefer
  exit
 interface loopback0
  address-family ipv4 unicast
   prefix-sid index 2
   exit
  exit
 exit
commit
```

XR3

```
router isis 1
 address-family ipv4 unicast
  segment-routing mpls sr-prefer
  exit
 interface loopback0
  address-family ipv4 unicast
   prefix-sid index 3
   exit
  exit
```

```
 exit
commit
```

XR4

```
router isis 1
 address-family ipv4 unicast
  segment-routing mpls sr-prefer
  exit
 interface loopback0
  address-family ipv4 unicast
   prefix-sid index 4
   exit
  exit
 exit
commit
```

XR5

```
router isis 1
 address-family ipv4 unicast
  segment-routing mpls sr-prefer
  exit
 interface loopback0
  address-family ipv4 unicast
   prefix-sid index 5
   exit
  exit
 exit
commit
```

Verification:

- Use the Traceroute command on R1 to verify the connectivity from R1 Loopback 0 to R2 Loopback0. What range of labels are being used in the SP Core?

Lab 4 – Configuring SR-LDP Mapping Server

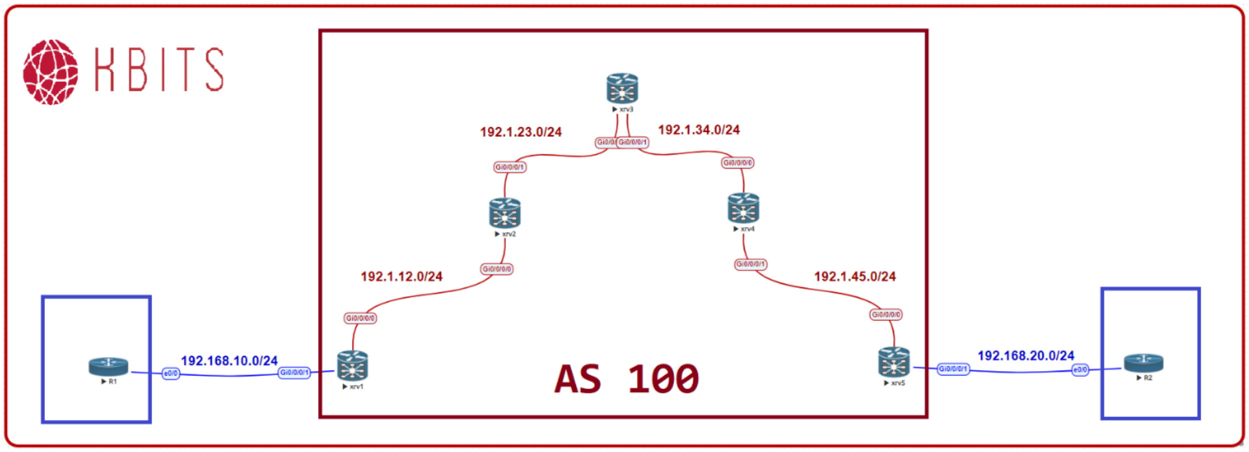

Task 1

De-configure Segment Routing on XR1 & XR2. It will going to revert back to LDP for MPLS Unicast Routing.

XR1

```
router isis 1
 address-family ipv4 unicast
  no segment-routing mpls sr-prefer
  exit
 interface loopback 0
  address-family ipv4 unicast
   no prefix-sid index 1
   exit
  exit
 exit
!
Commit
```

XR2

```
router isis 1
 address-family ipv4 unicast
  no segment-routing mpls sr-prefer
  exit
 interface loopback 0
  address-family ipv4 unicast
   no prefix-sid index 2
```

```
   exit
  exit
 exit
!
Commit
```

XR3

```
No router ospf 1
Commit
```

XR4

```
No router ospf 1
Commit
```

XR5

```
No router ospf 1
Commit
```

Task 2

De-configure LDP on XR4 & XR5. Also, disable it on the link between XR3 & XR4. It should continue to use SR for MPLS Unicast Routing.

XR3

```
Mpls ldp
 No interface Gig0/0/0/1
 Exit
!
commit
```

XR4

```
No Mpls ldp
!
Commit
```

XR5

```
No mpls ldp
!
Commit
```

Task 3

Configure XR3 as the SR-LDP Mapping server. It will run LDP with XR1 & XR2. It will run SR with XR4 & XR5. Configure it to advertise the labels for XR1 & XR2 loopbacks using a prefix-sid map.

XR3

```
segment-routing
 mapping-server
 prefix-sid-map address-family ipv4
   10.1.1.1/32 100 range 3
!
router isis 1
 address-family ipv4 unicast
  segment-routing prefix-sid-map advertise-local
 commit
```

Verification:

- Use the Traceroute command on R1 to verify the connectivity from R1 Loopback 0 to R2 Loopback0. What range of labels are being used in the SP Core?
- You should notice the label range changing after XR3.

Lab 5 – Configuring Segment Routing using BGP

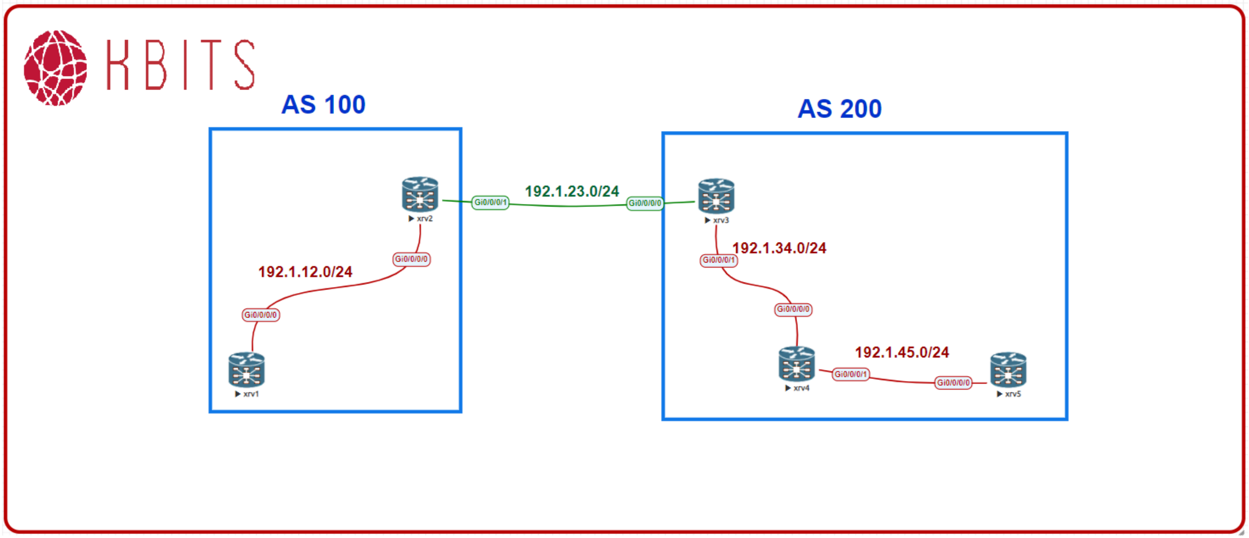

Interface IP Address Configuration

XR1

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.1	255.255.255.255
Gig0/0/0/0	192.1.12.1	255.255.255.0

XR2

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.2	255.255.255.255
Gig0/0/0/0	192.1.12.2	255.255.255.0
Gig0/0/0/1	192.1.23.2	255.255.255.0

XR3

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.3	255.255.255.255
G 0/0/0/0	192.1.23.3	255.255.255.0
G 0/0/0/1	192.1.34.3	255.255.255.0

XR4

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.4	255.255.255.255
G 0/0/0/0	192.1.34.4	255.255.255.0
G 0/0/0/1	192.1.45.4	255.255.255.0

XR5

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.5	255.255.255.255
G 0/0/0/0	192.1.45.5	255.255.255.0

AS 100

Task 1

Configure IP Addresses on XR1 & XR2 based on the above table. Run OSPF as the IGP for AS 100. Run Segment Routing for MPLS Unicast Routing using OSPF.

XR1

```
hostname XR1
!
interface gig0/0/0/0
 ip address 192.1.12.1 255.255.255.0
 no shut
!
interface loopback0
 ip address 10.1.1.1 255.255.255.255
!
segment-routing global-block 16000 23999
!
router ospf 1
 router-id 10.1.1.1
 segment-routing mpls
 area 0
  interface gig0/0/0/0
   exit
  interface loopback0
   prefix-sid index 1
   exit
  exit
 exit
```

```
commit
```

XR2

```
hostname XR2
!
Interface gig0/0/0/0
 ip address 192.1.12.2 255.255.255.0
 no shut
!
Interface gig0/0/0/1
 ip address 192.1.23.2 255.255.255.0
 no shut
!
Interface loopback0
 ip address 10.1.1.2 255.255.255.255
 exit
!
commit
!
segment-routing global-block 16000 23999
!
router ospf 1
 router-id 10.1.1.2
 segment-routing mpls
 area 0
  interface gig0/0/0/0
   exit
  interface gig0/0/0/1
   exit
  interface loopback0
   prefix-sid index 2
   exit
  exit
 exit
commit
```

AS 200

Task 1

Configure IP Addresses on R3, R4 & R5 based on the above table. Run OSPF as the IGP for AS 200. Run Segment Routing for MPLS Unicast Routing using OSPF.

XR3

```
hostname XR3
!
Interface gig0/0/0/1
 ip address 192.1.34.3 255.255.255.0
 no shut
!
Interface loopback0
 ip address 10.1.1.3 255.255.255.255
 exit
!
commit
!
segment-routing global-block 16000 23999
!
router ospf 1
 router-id 10.1.1.3
 segment-routing mpls
 area 0
  interface gig0/0/0/1
   exit
  interface loopback0
   prefix-sid index 3
   exit
  exit
 exit
commit
```

XR4

```
hostname XR4
!
Interface gig0/0/0/0
 ip address 192.1.34.4 255.255.255.0
 no shut
!
Interface gig0/0/0/1
```

```
 ip address 192.1.45.4 255.255.255.0
 no shut
!
Interface loopback0
 ip address 10.1.1.4 255.255.255.255
 exit
!
commit
!
segment-routing global-block 16000 23999
!
router ospf 1
 router-id 10.1.1.4
 segment-routing mpls
 area 0
  interface gig0/0/0/0
   exit
  interface gig0/0/0/1
   exit
  interface loopback0
   prefix-sid index 4
   exit
  exit
 exit
commit
```

XR5

```
hostname XR5
!
Interface gig0/0/0/0
 ip address 192.1.45.5 255.255.255.0
 no shut
!
Interface loopback0
 ip address 10.1.1.5 255.255.255.255
 exit
!
commit
!
segment-routing global-block 16000 23999
!
router ospf 1
 router-id 10.1.1.5
 segment-routing mpls
```

```
 area 0
  interface gig0/0/0/0
   exit
  interface loopback0
   prefix-sid index 5
   exit
  exit
 exit
commit
```

BGP – AS 100 & AS 200

Task 1

Configure IP Addresses on the external link connecting R2 & R3 based on the above table.

XR2

```
Interface gig0/0/0/1
 ip address 192.1.23.2 255.255.255.0
 no shut
!
commit
```

XR3

```
Interface gig0/0/0/0
 ip address 192.1.23.3 255.255.255.0
 no shut
!
commit
```

Task 2

Configure the Route policies to be used for BGP SR Label allocation and eBGP routes.

XR2 route-policy PASSALL pass end-policy commit ! route-policy SID($SID) set label-index $SID pass end-policy ! commit
XR3 route-policy PASSALL pass end-policy commit ! route-policy SID($SID) set label-index $SID pass end-policy ! commit

Task 3

Configure an Interface Static Route on the link between the ASBRs. This is required to provide labelled unicast on the external link.

XR1
 	
router static
 address-family ipv4 unicast
  192.1.23.3/32 gig0/0/0/1
  Exit
!
commit
	XR3
	
	router static
	 address-family ipv4 unicast
	  192.1.23.2/32 gig0/0/0/0
	  exit
	!
	commit

Task 4

Configure BGP between the ASBRs. Enable it to advertise the XR1 & XR2 routes with labels to each other. Redistribute the BGP learnt routes into OSPF.

XR2

```
router bgp 100
 address-family ipv4 unicast
  network 10.1.1.1/32 route-policy SID(1)
  network 10.1.1.2/32 route-policy SID(2)
  allocate-label all
  exit
!
 neighbor 192.1.23.3
  remote-as 200
  address-family ipv4 labeled-unicast
   route-policy PASSALL out
   route-policy PASSALL in
   exit
  exit
 exit
!
commit
!
router ospf 1
 redistribute bgp 100
 exit
commit
```

XR3

```
router bgp 200
 address-family ipv4 unicast
  network 10.1.1.3/32 route-policy SID(3)
  network 10.1.1.4/32 route-policy SID(4)
  network 10.1.1.5/32 route-policy SID(5)
  allocate-label all
  exit
!
 neighbor 192.1.23.2
  remote-as 100
  address-family ipv4 labeled-unicast
   route-policy PASSALL out
   route-policy PASSALL in
   exit
```

```
  exit
 exit
!
commit
!
router ospf 1
 redistribute bgp 200
 exit
commit
```

Verification:

- Reload XR2 & XR3. It is required for an Eve-NG setup to initialize BGP SR Labels.

- Verify the connectivity between the 2 sites. R1 & R5 should be able to reach each other's Loopbacks networks using labels.

Lab 6 – Configuring Large Scale MPLS / Unified MPLS

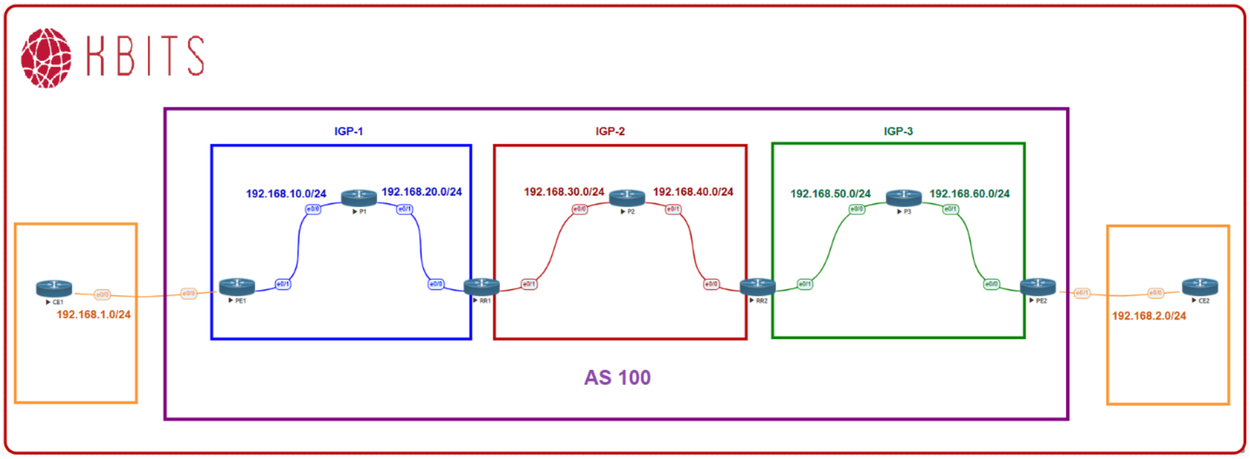

Interface IP Address Configuration

PE1

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.1	255.255.255.255
E 0/0	192.168.1.1	255.255.255.0
E 0/1	192.168.10.1	255.255.255.0

P1

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.2	255.255.255.255
E 0/0	192.168.10.2	255.255.255.0
E 0/1	192.168.20.2	255.255.255.0

RR1

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.3	255.255.255.255
E 0/0	192.168.20.3	255.255.255.0
E 0/1	192.168.30.3	255.255.255.0

P2

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.4	255.255.255.255
E 0/0	192.168.30.4	255.255.255.0
E 0/1	192.168.40.4	255.255.255.0

RR2

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.5	255.255.255.255
E 0/0	192.168.40.5	255.255.255.0
E 0/1	192.168.50.5	255.255.255.0

P3

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.6	255.255.255.255
E 0/0	192.168.50.6	255.255.255.0
E 0/1	192.168.60.6	255.255.255.0

PE2

Interface	IP Address	Subnet Mask
Loopback 0	10.1.1.7	255.255.255.255
E 0/0	192.168.60.7	255.255.255.0
E 0/1	192.168.2.7	255.255.255.0

CE1

Interface	IP Address	Subnet Mask
Loopback 0	10.11.11.11	255.255.255.255
E 0/0	192.168.1.11	255.255.255.0

CE2

Interface	IP Address	Subnet Mask
Loopback 0	10.22.22.22	255.255.255.255
E 0/0	192.168.2.22	255.255.255.0

IGP-1

Task 1

Configure IP Addresses on PE1, P1 & RR1 in IGP-1 based on the above table. Run OSPF as the IGP for this part of the network. Configure LDP on the interfaces within this domain.

PE1 hostname PE1 ! interface E 0/1 ip address 192.168.10.1 255.255.255.0 no shut ! interface loopback0 ip address 10.1.1.1 255.255.255.255 ! router ospf 1 network 10.1.1.0 0.0.0.255 area 0 network 192.168.10.0 0.0.0.255 area 0 ! mpls ldp router-id loopback0 ! Interface E 0/1 mpls ip
P1 hostname P1 ! interface E 0/0 ip address 192.168.10.2 255.255.255.0 no shut ! interface E 0/1 ip address 192.168.20.2 255.255.255.0 no shut ! interface loopback0 ip address 10.1.1.2 255.255.255.255 ! router ospf 1 network 10.1.1.0 0.0.0.255 area 0 network 192.168.10.0 0.0.0.255 area 0

```
 network 192.168.20.0 0.0.0.255 area 0
!
mpls ldp router-id loopback0
!
Interface E 0/1
 mpls ip
!
Interface E 0/0
 mpls ip
```

RR1

```
hostname RR1
!
interface E 0/0
 ip address 192.168.20.3 255.255.255.0
 no shut
!
interface loopback0
 ip address 10.1.1.3 255.255.255.255
!
router ospf 1
 network 10.1.1.0 0.0.0.255 area 0
 network 192.168.20.0 0.0.0.255 area 0
!
mpls ldp router-id loopback0
!
Interface E 0/0
 mpls ip
```

IGP-2

Task 1

Configure IP Addresses on RR1, P2 & RR2 in IGP-2 based on the above table. Run IS-IS as the IGP for this part of the network. Configure LDP on the interfaces within this domain.

RR1

```
interface E 0/1
 ip address 192.168.30.3 255.255.255.0
 no shut
!
router isis
 net 49.0000.3333.3333.3333.00
 is-type level-2
 metric-style wide
!
Interface Loopback0
 ip router isis
!
Interface E 0/1
 ip router isis
 mpls ip
```

P2

```
hostname P2
!
interface E 0/0
 ip address 192.168.30.4 255.255.255.0
 no shut
!
interface E 0/1
 ip address 192.168.40.4 255.255.255.0
 no shut
!
interface loopback0
 ip address 10.1.1.4 255.255.255.255
!
router isis
 net 49.0000.4444.4444.4444.00
 is-type level-2
 metric-style wide
!
```

```
mpls ldp router-id loopback0
!
Interface E 0/0
 ip router isis
 mpls ip
!
Interface E 0/1
 ip router isis
 mpls ip
!
Interface Loopback0
 ip router isis
```

RR2

```
hostname RR2
!
interface E 0/0
 ip address 192.168.40.5 255.255.255.0
 no shut
!
interface loopback0
 ip address 10.1.1.5 255.255.255.255
!
router isis
 net 49.0000.5555.5555.5555.00
 is-type level-2
 metric-style wide
!
mpls ldp router-id loopback0
!
Interface E 0/0
 ip router isis
 mpls ip
!
Interface Loopback0
 ip router isis
```

IGP-3

Task 1

Configure IP Addresses on RR2, P3 & PE3 in IGP-3 based on the above table. Run EIGRP as the IGP for this part of the network. Configure LDP on the interfaces within this domain.

RR2

```
interface E 0/1
 ip address 192.168.50.5 255.255.255.0
 no shut
!
router eigrp 100
 net 192.168.50.0
 net 10.1.1.0 0.0.0.255
!
Interface E 0/1
 mpls ip
```

P3

```
hostname P3
!
interface E 0/0
 ip address 192.168.50.6 255.255.255.0
 no shut
!
interface E 0/1
 ip address 192.168.60.6 255.255.255.0
 no shut
!
interface loopback0
 ip address 10.1.1.6 255.255.255.255
!
router eigrp 100
 net 192.168.50.0
 net 192.168.60.0
 net 10.1.1.0 0.0.0.255
!
mpls ldp router-id loopback0
!
Interface E 0/0
 mpls ip
!
```

Interface E 0/1 mpls ip
PE2 hostname PE2 ! interface E 0/0 ip address 192.168.60.7 255.255.255.0 no shut ! interface loopback0 ip address 10.1.1.7 255.255.255.255 ! router eigrp 100 net 192.168.60.0 net 10.1.1.0 0.0.0.255 ! mpls ldp router-id loopback0 ! Interface E 0/0 mpls ip

BGP Labelled Unicast

Task 1

Configure BGP to exchange the PE Loopback routers. The BGP relationship will be between PE1 – RR1 – RR2 – PE2. The routes advertised should be labelled. The RR's will reflect the routes.

PE1 router bgp 100 network 10.1.1.1 mask 255.255.255.255 neighbor 10.1.1.3 remote-as 100 neighbor 10.1.1.3 update-source lo0 neighbor 10.1.1.3 send-label
RR1 router bgp 100 neighbor 10.1.1.1 remote-as 100 neighbor 10.1.1.1 update-source lo0 neighbor 10.1.1.1 route-reflector-client neighbor 10.1.1.1 next-hop-self all neighbor 10.1.1.1 send-label neighbor 10.1.1.5 remote-as 100 neighbor 10.1.1.5 update-source lo0 neighbor 10.1.1.5 route-reflector-client neighbor 10.1.1.5 next-hop-self all neighbor 10.1.1.5 send-label
RR2 router bgp 100 neighbor 10.1.1.3 remote-as 100 neighbor 10.1.1.3 update-source lo0 neighbor 10.1.1.3 route-reflector-client neighbor 10.1.1.3 next-hop-self all neighbor 10.1.1.3 send-label neighbor 10.1.1.7 remote-as 100 neighbor 10.1.1.7 update-source lo0 neighbor 10.1.1.7 route-reflector-client neighbor 10.1.1.7 next-hop-self all neighbor 10.1.1.7 send-label
PE2 router bgp 100 network 10.1.1.7 mask 255.255.255.255

```
 neighbor 10.1.1.5 remote-as 100
 neighbor 10.1.1.5 update-source lo0
 neighbor 10.1.1.5 send-label
```

Verification:

- Make sure that you can ping PE2 from PE1 and vice versa.

Task 2

Configure a MP-iBGP relationship between PE1 & PE2. Make sure to only exchange VPNv4 routes by blocking all IPv4 Unicast routes from each other.

PE1

```
ip prefix-list DENYALL deny 0.0.0.0/0 le 32
!
router bgp 100
 neighbor 10.1.1.7 remote-as 100
 neighbor 10.1.1.7 update-source lo0
 neighbor 10.1.1.7 prefix-list DENYALL out
!
 address-family vpnv4
  neighbor 10.1.1.7 activate
```

PE2

```
ip prefix-list DENYALL deny 0.0.0.0/0 le 32
!
router bgp 100
 neighbor 10.1.1.1 remote-as 100
 neighbor 10.1.1.1 update-source lo0
 neighbor 10.1.1.1 prefix-list DENYALL out
!
 address-family vpnv4
  neighbor 10.1.1.1 activate
```

Task 3

Configure PE-CE Routing to connect the PE routers to the customer routers. Advertise the Loopback interfaces on the CE Routers. Configure the CE Routers in AS 65001 & 65002 respectively.

PE1
 	
vrf definition CUST-A
 rd 100:1
 address-family ipv4
  route-target both 100:1
!
Interface E 0/0
 vrf forwarding CUST-A
 ip add 192.168.1.1 255.255.255.0
 no shut
!
router bgp 100
 address-family ipv4 vrf CUST-A
  neighbor 192.168.1.11 remote-as 65001

PE2
 	
vrf definition CUST-A
 rd 100:1
 address-family ipv4
  route-target both 100:1
!
Interface E 0/1
 vrf forwarding CUST-A
 ip add 192.168.2.7 255.255.255.0
 no shut
!
router bgp 100
 address-family ipv4 vrf CUST-A
  neighbor 192.168.2.22 remote-as 65002

CE1 Interface E 0/0 Ip address 192.168.1.11 255.255.255.0 No shut ! Interface Loopback0 Ip address 10.11.11.11 255.255.255.0 !

```
Router bgp 65001
 Neighbor 192.168.1.1 remote-as 100
 Network 10.11.11.0 mask 255.255.255.0
```

CE2

```
Interface E 0/0
 Ip address 192.168.2.22 255.255.255.0
 No shut
!
Interface Loopback0
 Ip address 10.22.22.22 255.255.255.0
!
Router bgp 65002
 Neighbor 192.168.2.7 remote-as 100
 Network 10.22.22.0 mask 255.255.255.0
```

Verification:

- Make sure that you can ping CE2 Loopback0 from CE1 and vice versa.

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring Multicast-VPN (M-VPN)

Lab 1 – Configuring SP Core – Intra-AS MPLS VPN & Multicast-Routing

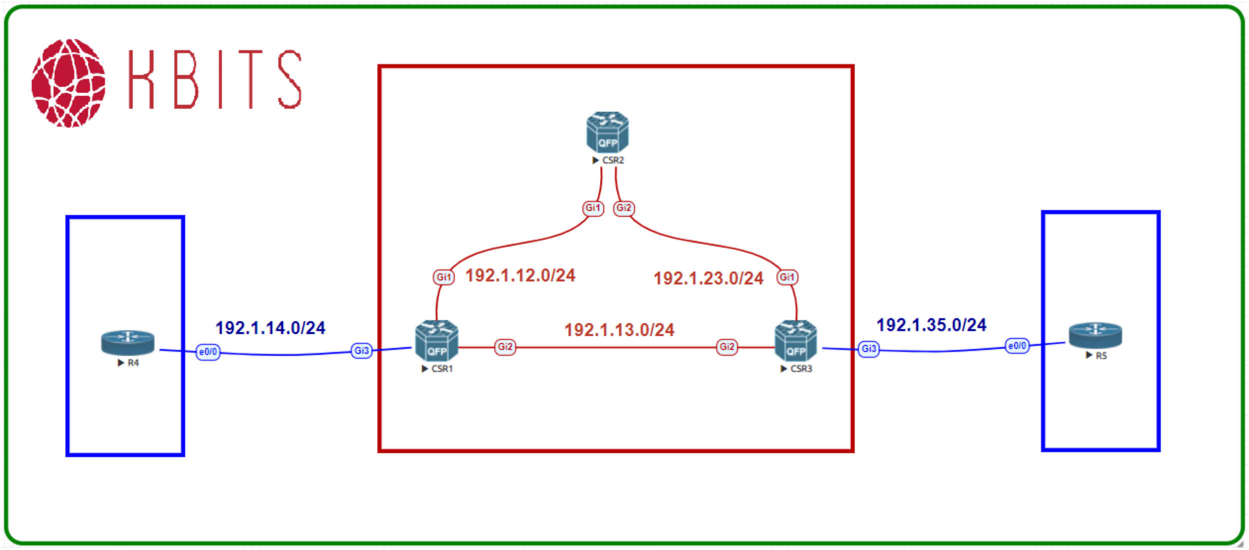

Interface IP Address Configuration

CSR1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
Gig1	192.1.12.1	255.255.255.0
Gig2	192.1.13.1	255.255.255.0
Gig3	192.1.14.1	255.255.255.0

CSR2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
Gig1	192.1.12.2	255.255.255.0
Gig2	192.1.23.2	255.255.255.0

CSR3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
Gig1	192.1.23.3	255.255.255.0
Gig2	192.1.13.3	255.255.255.0
Gig3	192.1.35.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	10.4.4.4	255.255.255.0
E 0/0	192.1.14.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	10.5.5.5	255.255.255.0
E 0/0	192.1.35.5	255.255.255.0

Task 1

Configure OSPF between all the SP routers (CSR1, CSR2 & CSR3). Use x.x.x.x as the router-id, where x is the Router number. Advertise all interal links in OSPF.

CSR1 Router ospf 1 Router-id 1.1.1.1 Network 1.1.1.1 0.0.0.0 area 0 Network 192.1.12.1 0.0.0.0 area 0	**CSR2** Router ospf 1 Router-id 2.2.2.2 Network 2.2.2.2 0.0.0.0 area 0 Network 192.1.12.2 0.0.0.0 area 0 Network 192.1.23.2 0.0.0.0 area 0
CSR3 Router ospf 1 Router-id 3.3.3.3 Network 3.3.3.3 0.0.0.0 area 0 Network 192.1.23.3 0.0.0.0 area 0 Network 192.1.34.3 0.0.0.0 area 0	**CSR4** Router ospf 1 Router-id 4.4.4.4 Network 4.4.4.4 0.0.0.0 area 0 Network 192.1.34.4 0.0.0.0 area 0

Task 2

Configure MPLS on all the physical links in the SP Network. The LDP neighbour relationships should be formed based on the most reliable interface.

CSR1 mpls ldp router-id loopback0 ! Interface Gig1 mpls ip ! Interface Gig2 mpls ip	**CSR2** mpls ldp router-id loopback0 ! Interface Gig1 mpls ip ! Interface Gig2 mpls ip
CSR3 mpls ldp router-id loopback0 ! Interface Gig1 mpls ip ! Interface Gig2 mpls ip	

Task 3

Configure MP-iBGP between CSR1-CSR2-CSR3 such that CSR2 is configured as the RR to progagate VPNv4 routers.

CSR1
 	
router bgp 100
 neighbor 2.2.2.2 remote-as 100
 neighbor 2.2.2.2 update-source lo0
!
 address-family vpnv4
  neighbor 2.2.2.2 activate
	CSR2
	
	router bgp 100
	 neighbor IBGP peer-group
	 neighbor IBGP remote-as 100
	 neighbor IBGP update-source lo0
	 neighbor 1.1.1.1 peer-group IBGP
	 neighbor 3.3.3.3 peer-group IBGP
	!
	 address-family vpnv4
	  neighbor IBGP route-reflector-client
	  neighbor 1.1.1.1 activate
	  neighbor 3.3.3.3 activate

CSR3 router bgp 100 neighbor 2.2.2.2 remote-as 100	

neighbor 2.2.2.2 update-source lo0 ! address-family vpnv4 neighbor 2.2.2.2 activate	

Task 4

Configure a VRF called Cust-A on the PE Routers (CSR-1 & CSR-3). Use the RD & RT values as 100:1. Configure the interfaces between the PE-CE based on the diagram. Configure BGP as the PE-CE Routing protocol. R4 will be in AS 65004 and R5 will be in AS 65005.

CSR1	**CSR3**
vrf definition Cust-A rd 100:1 address-family ipv4 route-target both 100:1 ! Interface Gig3 vrf forwarding Cust-A ip address 192.1.14.1 255.255.255.0 no shut ! router bgp 100 address-family ipv4 vrf Cust-A neighbor 192.1.14.4 remote-as 65004	vrf definition Cust-A rd 100:1 address-family ipv4 route-target both 100:1 ! Interface Gig3 vrf forwarding Cust-A ip address 192.1.35.3 255.255.255.0 no shut ! router bgp 100 address-family ipv4 vrf Cust-A neighbor 192.1.35.5 remote-as 65005

Task 5

Configure the CE Routers (R4 & R5). Configure BGP as the PE-CE Routing protocol. R4 will be in AS 65004 and R5 will be in AS 65005. Advertise loopback0 interface of R4 & R5 in BGP.

R4	**R5**
Interface E 0/0 ip address 192.1.14.4 255.255.255.0 no shut ! Interface loopback0 ip address 10.4.4.4 255.255.255.0 ! router bgp 65004 neighbor 192.1.14.1 remote-as 100 network 10.4.4.0 mask 255.255.255.0	Interface E 0/0 ip address 192.1.35.5 255.255.255.0 no shut ! Interface loopback0 ip address 10.5.5.5 255.255.255.0 ! router bgp 65005 neighbor 192.1.35.3 remote-as 100 network 10.5.5.0 mask 255.255.255.0

Lab 2 – Configuring Multicast VPN (M-VPN) Using Static RP

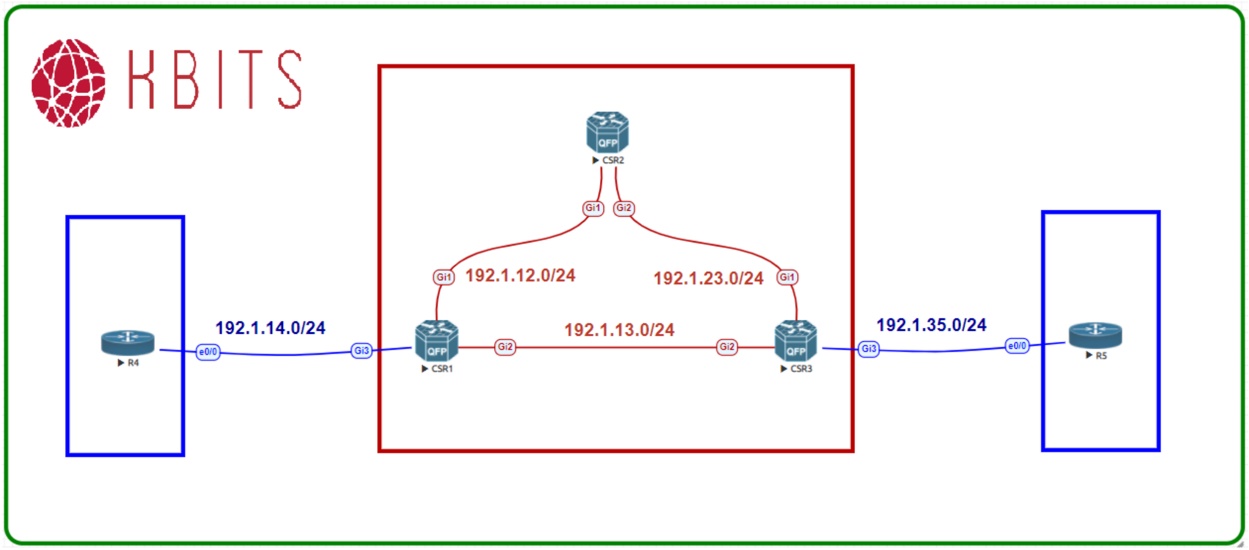

Task 1

Configure Multicast routing on all SP routers using PIM Sparse-dense-mode. Enable Multicast routing for the CUST-A VRF on the PE routers.

CSR1 Ip multicast-routing distributed ! Interface Gig1 ip pim sparse-dense-mode ! Interface Gig2 ip pim sparse-dense-mode ! Interface loop0 ip pim sparse-dense-mode	**CSR2** Ip multicast-routing distributed ! Interface Gig1 ip pim sparse-dense-mode ! Interface Gig2 ip pim sparse-dense-mode ! Interface loop0 ip pim sparse-dense-mode
CSR3 Ip multicast-routing distributed ! Interface Gig1 ip pim sparse-dense-mode !	

Interface Gig2 ip pim sparse-dense-mode ! Interface loop0 ip pim sparse-dense-mode	

Task 2

Configure Multicast routing on the Customer Routers (R4 & R5) using PIM Sparse-dense-mode. Configure 10.4.4.4 as the RP-Address. Have the loopback interface join the 224.45.45.45 multicast group.

R4	**R5**
ip multicast-routing ! int E 0/0 ip pim sparse-dense-mode ! int Loopback0 ip pim sparse-dense-mode ip igmp join-group 224.45.45.45 ! ip pim rp-address 10.4.4.4	ip multicast-routing ! int E 0/0 ip pim sparse-dense-mode ! int Loopback0 ip pim sparse-dense-mode ip igmp join-group 224.45.45.45 ! ip pim rp-address 10.4.4.4

Task 3

Configure VRF Cust-A for Multicast-Routing on the PE Routers. Use the 239.1.1.1 for the Default MDT group. Use 232.1.1.0/24 for the MDT Data Group.

CSR1
 	
ip multicast-routing vrf Cust-A distributed
!
vrf definition Cust-A
 address-family ipv4
  mdt default 239.1.1.1
  mdt data 232.1.1.0 0.0.0.255
!
interface Gig3
 ip pim sparse-mode
!
ip pim vrf Cust-A rp-address 10.4.4.4

CSR3
 	
ip multicast-routing vrf Cust-A distributed
!
vrf definition Cust-A
 address-family ipv4
  mdt default 239.1.1.1
  mdt data 232.1.1.0 0.0.0.255
!
interface Gig3
 ip pim sparse-mode
!
ip pim vrf Cust-A rp-address 10.4.4.4

Task 4

Ping 224.45.45.45 from R4 & R5. You should receive replies from R5 & R6. Verify the mroute Table entries

R5 Ping 224.45.45.45
R6 Ping 224.45.45.45

Lab 3– Configuring Multicast MPLS VPN (M-VPN) Using Auto-RP

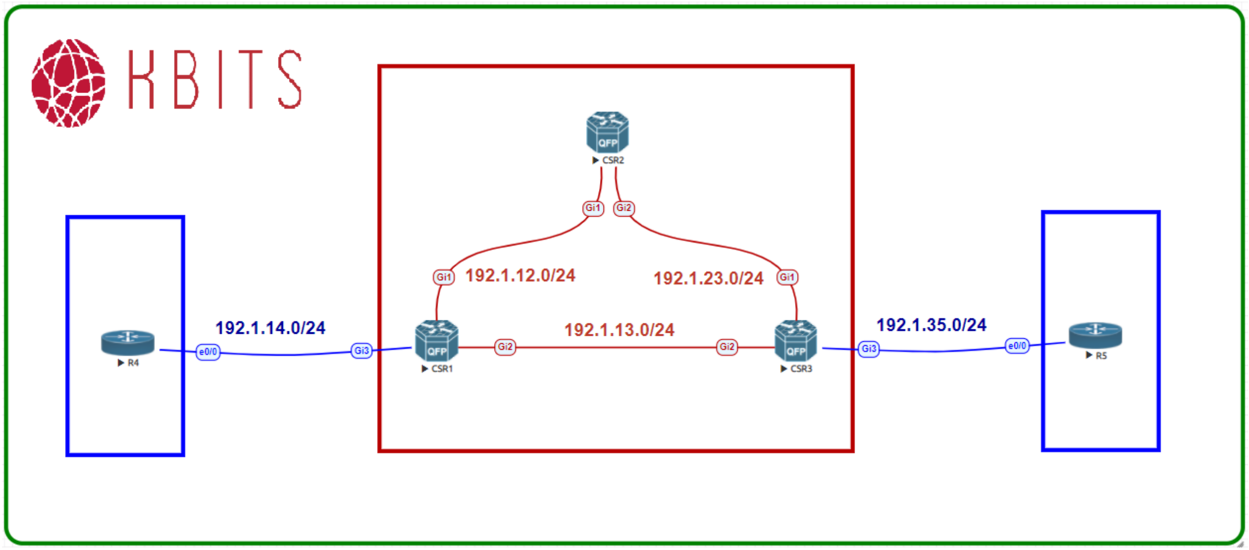

Task 1

As we will be configuring Auto-RP on R4, take out the static assignment of the RP on CSR1, R4, R5 & CSR3 Routers

CSR1 No ip pim vrf Cust-A rp-address 10.4.4.4	**CSR3** No ip pim vrf Cust-A rp-address 10.4.4.4
R4 No ip pim rp-address 10.4.4.4	**R5** No ip pim rp-address 10.4.4.4

Task 2

Configure R4 as the RP-Candidate and the Mapping agent. Use Loopback 0 as the source address with a scope of 10 and an interval of 10 secs.

R4 ip pim send-rp-announce Loopback0 scope 10 interval 10 ip pim send-rp-discovery Loopback0 scope 10

Task 3

Verify that the CE routers are notified of the RP Mapping using the show ip pim rp mapping command

R4 **sh ip pim rp mapping** PIM Group-to-RP Mappings Group(s) 224.0.0.0/4 RP **10.4.4.4** (?), v2v1 Info source: **10.4.4.4** (?), elected via **Auto-RP** Uptime: 00:11:06, expires: 00:02:44
R5 **sh ip pim rp mapping** PIM Group-to-RP Mappings Group(s) 224.0.0.0/4 RP **10.4.4.4** (?), v2v1 Info source: **10.4.4.4** (?), elected via **Auto-RP** Uptime: 00:12:25, expires: 00:02:28

Lab 4 – Configuring Multicast MPLS VPN using mLDP

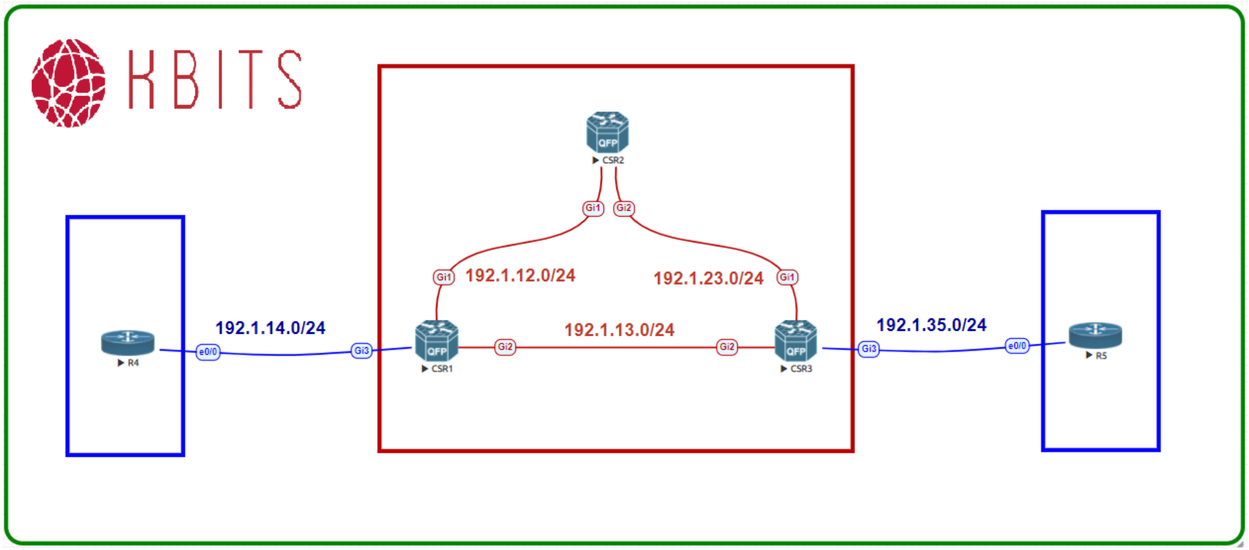

Task 1

De-Configure MDT default & data under the VRF Cust-A on the PE Routers (CSR1 & CSR3).

CSR1	CSR3

vrf definition CUST-A
 address-family ipv4
  no mdt default 239.1.1.1
  no mdt data 232.1.1.0 0.0.0.255
	vrf definition CUST-A
	 address-family ipv4
	  no mdt default 239.1.1.1
	  no mdt data 232.1.1.0 0.0.0.255

Task 2

Configure PE Routers to setup a mLDP neighbor for VRF Cust-A towards each other. Use a VPN ID of 100:1 with a MLDP Label of 200.

CSR1	CSR3

vrf definition CUST-A
!
 vpn id 100:1
  address-family ipv4
   mdt default mpls mldp 3.3.3.3
   mdt data mpls mldp 200
	vrf definition CUST-A
	!
	 vpn id 100:1
	 address-family ipv4
	  mdt default mpls mldp 1.1.1.1
	  mdt data mpls mldp 200

Task 3

Verify that the other routers are notified of the RP Mapping using the show ip pim rp mapping command

R4 **sh ip pim rp mapping** PIM Group-to-RP Mappings Group(s) 224.0.0.0/4 RP **10.4.4.4** (?), v2v1 Info source: **10.4.4.4** (?), elected via **Auto-RP** Uptime: 00:11:42, expires: 00:02:11
R5 **sh ip pim rp mapping** PIM Group-to-RP Mappings Group(s) 224.0.0.0/4 RP **10.4.4.4** (?), v2v1 Info source: **10.4.4.4** (?), elected via **Auto-RP** Uptime: 00:12:25, expires: 00:02:28

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring MPLS Traffic Engineering

Lab 1 – Configuring SP Core Network with MPLS Unicast Routing

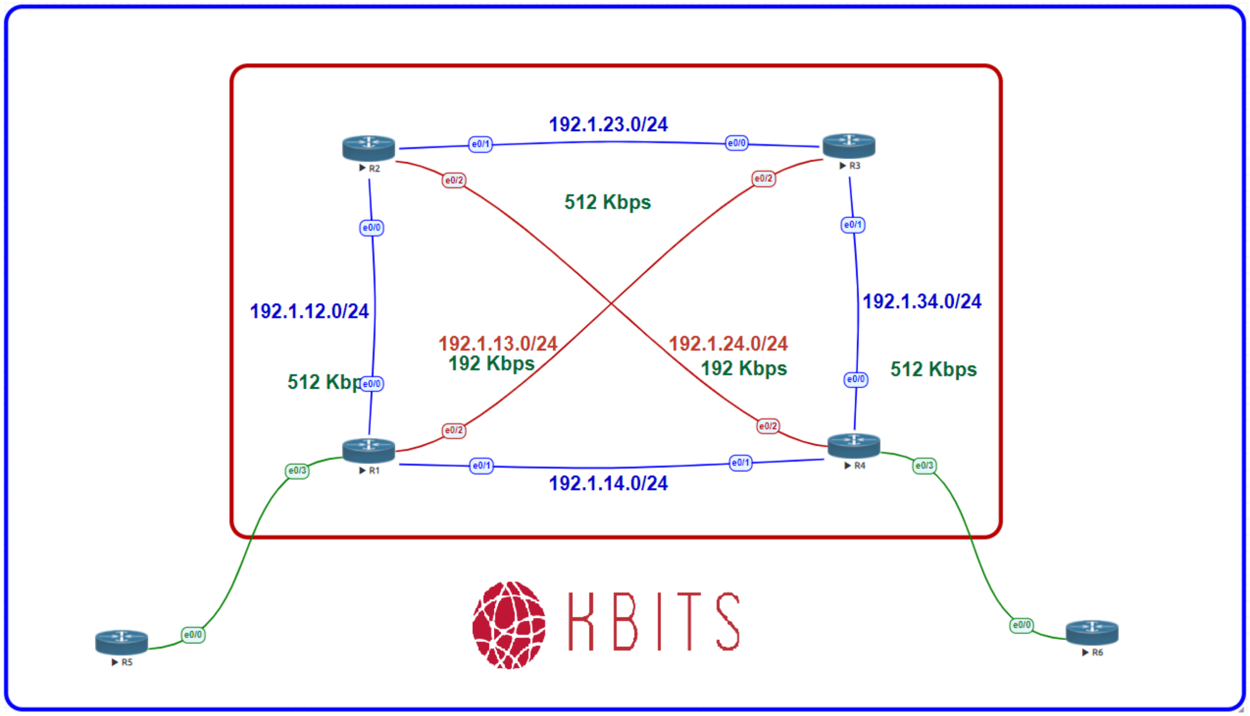

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
E 0/0	192.1.12.1	255.255.255.0
E 0/1	192.1.14.1	255.255.255.0
E 0/2	192.1.13.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
E 0/0	192.1.12.2	255.255.255.0
E 0/1	192.1.23.2	255.255.255.0
E 0/2	192.1.24.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
E 0/0	192.1.23.3	255.255.255.0
E 0/1	192.1.34.3	255.255.255.0
E 0/2	192.1.13.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
E 0/0	192.1.34.4	255.255.255.0
E 0/1	192.1.14.4	255.255.255.0
E 0/2	192.1.24.4	255.255.255.0

Task 1

Configure OSPF between all the SP routers (R1, R2, R3, R4). Use 0.0.0.X as the router-id, where x is the Router #. Advertise all links in OSPF.

R1 router ospf 1 router-id 0.0.0.1 network 192.1.12.0 0.0.0.255 area 0 network 192.1.13.0 0.0.0.255 area 0 network 192.1.14.0 0.0.0.255 area 0 network 1.0.0.0 0.255.255.255 area 0	**R2** router ospf 1 router-id 0.0.0.2 network 192.1.12.0 0.0.0.255 area 0 network 192.1.23.0 0.0.0.255 area 0 network 192.1.24.0 0.0.0.255 area 0 network 2.0.0.0 0.255.255.255 area 0
R3 router ospf 1 router-id 0.0.0.3 network 192.1.13.0 0.0.0.255 area 0 network 192.1.23.0 0.0.0.255 area 0 network 192.1.34.0 0.0.0.255 area 0 network 3.0.0.0 0.255.255.255 area 0	**R4** router ospf 1 router-id 0.0.0.4 network 192.1.14.0 0.0.0.255 area 0 network 192.1.24.0 0.0.0.255 area 0 network 192.1.34.0 0.0.0.255 area 0 network 4.0.0.0 0.255.255.255 area 0

Task 2

Configure MPLS on all the physical links in the SP Network. The LDP neighbour relationships should be formed based on the most reliable interface.

R1	R2
Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip ! Interface E 0/2 mpls ip	Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip ! Interface E 0/2 mpls ip
R3 Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip ! Interface E 0/2 mpls ip	**R4** Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip ! Interface E 0/2 mpls ip

Lab 2 – Configuring MPLS TE - Static Tunnels

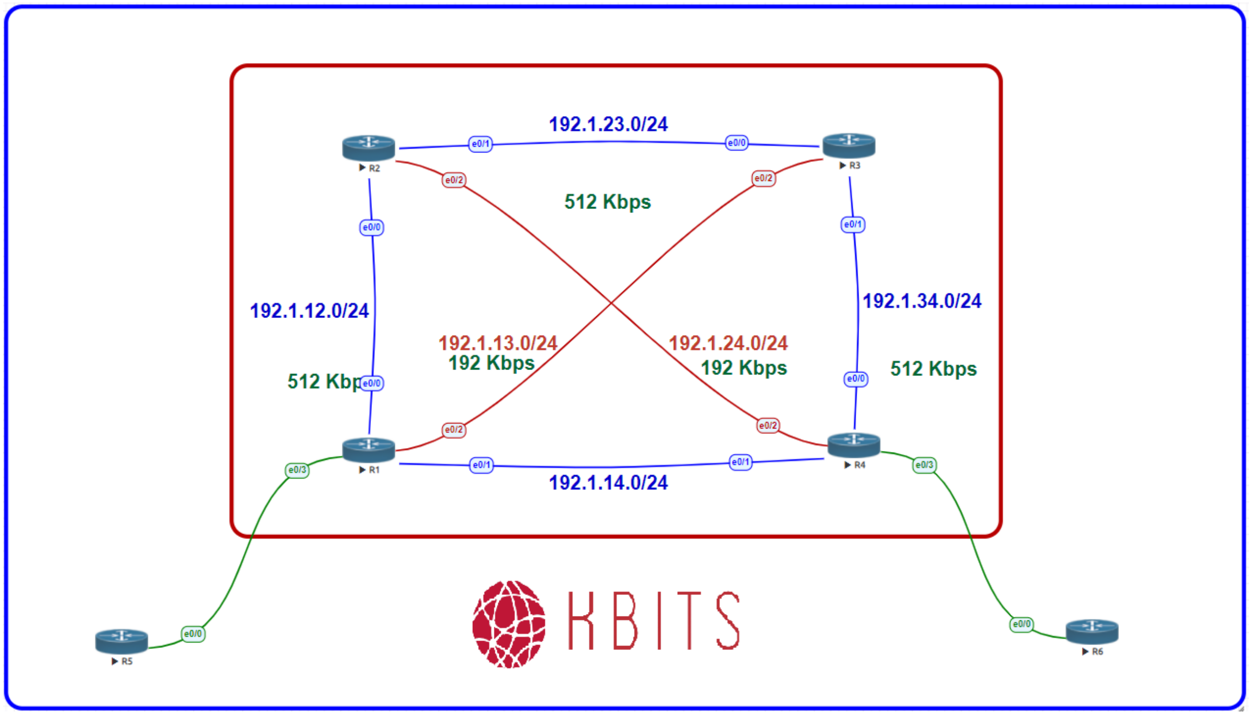

Task 1

Enable MPLS Traffic Engineering on all the Core Routers. OSPF Area 0 should also be enabled for MPLS TE. Use Loopback 0 as the router-id.

R1 Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0	**R2** Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0
R3 Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0	**R4** Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0

Task 2

Enable MPLS Traffic Engineering on all the SP Interfaces.

R1 Interface E 0/0 mpls traffic-eng Tunnels ! Interface E 0/1 mpls traffic-eng Tunnels ! Interface E 0/2 mpls traffic-eng Tunnels	**R2** Interface E 0/0 mpls traffic-eng Tunnels ! Interface E 0/1 mpls traffic-eng Tunnels ! Interface E 0/2 mpls traffic-eng Tunnels
R3 Interface E 0/0 mpls traffic-eng Tunnels ! Interface E 0/1 mpls traffic-eng Tunnels ! Interface E 0/2 mpls traffic-eng Tunnels	**R4** Interface E 0/0 mpls traffic-eng Tunnels ! Interface E 0/1 mpls traffic-eng Tunnels ! Interface E 0/2 mpls traffic-eng Tunnels

Task 3

Enable RSVP Bandwidth reservation on all Interfaces on all SP Routers. Configure a reservation of 512 on all E 0/0 & E 0/1 ports. Configure a reservation of 192 on all E 0/2 Interfaces.

R1 Interface E 0/0 ip rsvp bandwidth 512 ! Interface E 0/1 ip rsvp bandwidth 512 ! Interface E 0/2 ip rsvp bandwidth 192	**R2** Interface E 0/0 ip rsvp bandwidth 512 ! Interface E 0/1 ip rsvp bandwidth 512 ! Interface E 0/2 ip rsvp bandwidth 192
R3 Interface E 0/0 ip rsvp bandwidth 512 !	**R4** Interface E 0/0 ip rsvp bandwidth 512 !

Interface E 0/1 ip rsvp bandwidth 512 ! Interface E 0/2 ip rsvp bandwidth 192	Interface E 0/1 ip rsvp bandwidth 512 ! Interface E 0/2 ip rsvp bandwidth 192

Task 4

Configure a MPLS TE Tunnel to have R1 use the R1-> R2 -> R4 path for Traffic destined to 4.4.4.4. This tunnel should have a bandwidth reservation requirement of 128 kbps.

R1

```
ip explicit-path name R1-R4 enable
 next-address 2.2.2.2
 next-address 4.4.4.4
!
Interface Tunnel 14
 ip unnumbered Loopback0
 tunnel destination 4.4.4.4
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 4 4
 tunnel mpls traffic-eng bandwidth 128
 tunnel mpls traffic-eng path-option 1 explicit name R1-R4
```

Notes: Make sure the Tunnel comes up and it is used in the routing table as the next hop for the 4.4.4.4 network.

Task 5

Configure a MPLS TE Tunnel to have R2 use the R2-> R1 path for Traffic destined to 1.1.1.1. This tunnel should have a bandwidth reservation requirement of 400 kbps.

R2

```
ip explicit-path name R2-R1 enable
 next-address 1.1.1.1
!
Interface Tunnel 21
 ip unnumbered Loopback0
 tunnel destination 1.1.1.1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 400
 tunnel mpls traffic-eng path-option 1 explicit name R2-R1
```

```
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng autoroute announce
```

Notes: Make sure the Tunnel comes up and it is used in the routing table as the next hop for the 1.1.1.1 network.

Task 6

Configure a MPLS TE Tunnel to have R4 use the R4-> R3 -> R1 path for Traffic destined to 1.1.1.1. This tunnel should have a bandwidth reservation requirement of 128 kbps.

R4

```
ip explicit-path name R4-R1 enable
 next-address 3.3.3.3
 next-address 1.1.1.1
!
Interface Tunnel 41
 ip unnumbered Loopback0
 tunnel destination 1.1.1.1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 128
 tunnel mpls traffic-eng path-option 1 explicit name R4-R1
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng autoroute announce
```

Notes: Make sure the Tunnel comes up and it is used in the routing table as the next hop for the 1.1.1.1 network.

Task 7

Configure a MPLS TE Tunnel to have R1 use the R1-> R2 -> R3 path for Traffic destined to 1.1.1.1. This tunnel should have a bandwidth reservation requirement of 128 kbps. This is a high priority tunnel.

R1

```
ip explicit-path name R1-R3 enable
 next-address 2.2.2.2
 next-address 3.3.3.3
!
Interface Tunnel 13
 ip unnumbered Loopback0
 tunnel destination 3.3.3.3
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 400
 tunnel mpls traffic-eng path-option 1 explicit name R1-R3
```

tunnel mpls traffic-eng priority 2 2 tunnel mpls traffic-eng autoroute announce
Notes: Make sure the Tunnel comes up and it is used in the routing table as the next hop for the 3.3.3.3 network.

Lab 3 – Configuring MPLS TE - Dynamic Tunnels

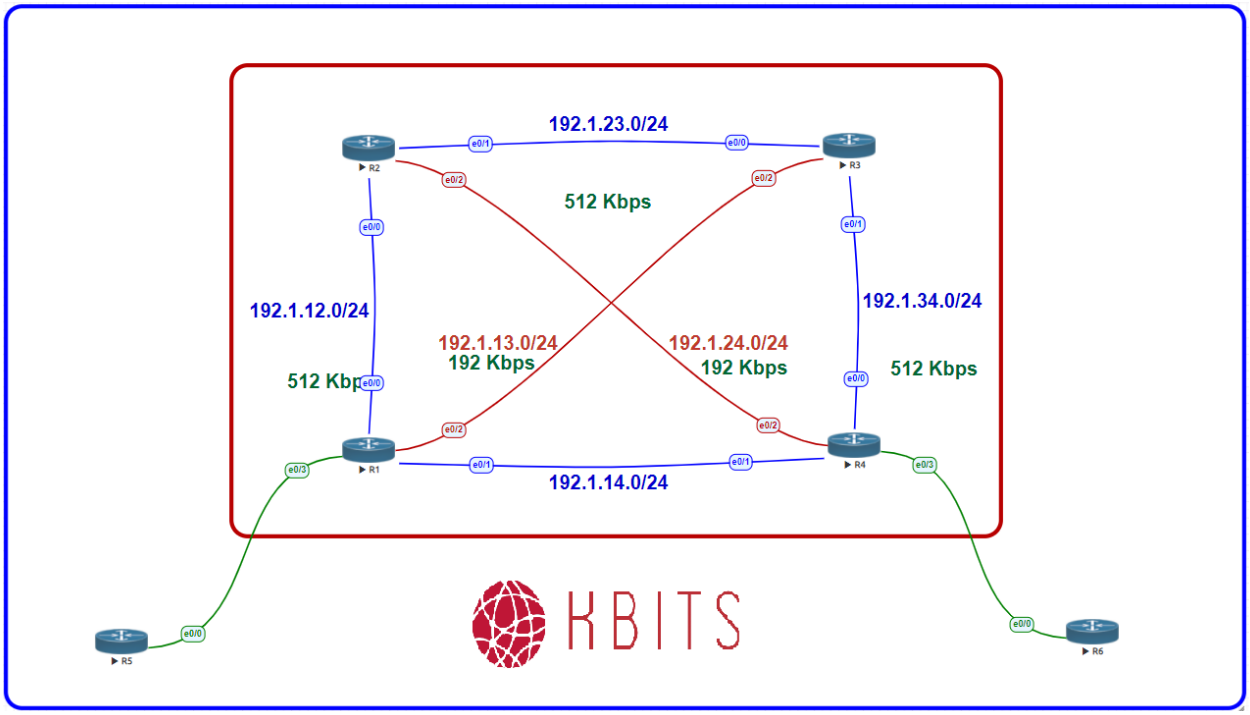

Task 1

Delete the Tunnels created in Lab 2. We will be creating tunnels dynamically based on the available bandwidth the routing protocol metrics.

R1 No Interface Tunnel 14 No Interface Tunnel 13	**R2** No Interface Tunnel 21
R4 No Interface Tunnel 41	

Task 2

Configure a dynamic MPLS TE Tunnel on R1 toward R4 with a bandwidth requirement of 128 kbps.

R1 Interface Tunnel 14 ip unnumbered Loopback0

```
 tunnel destination 4.4.4.4
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 128
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng autoroute announce
```

Task 3

Configure a dynamic MPLS TE Tunnel on R2 toward R1 with a bandwidth requirement of 400 kbps.

R2

```
Interface Tunnel 21
 ip unnumbered Loopback0
 tunnel destination 1.1.1.1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 400
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng autoroute announce
```

Task 4

Configure a dynamic MPLS TE Tunnel on R1 toward R2 with a bandwidth requirement of 400 kbps.

R1

```
Interface Tunnel 12
 ip unnumbered Loopback0
 tunnel destination 2.2.2.2
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 400
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng priority 2 2
 tunnel mpls traffic-eng autoroute announce
```

Task 6

Verify which tunnels were created. Also, verify the path that is being taken by the Tunnels.

R1 show mpls traffic-eng tunnel	**R2** show mpls traffic-eng tunnel

Lab 4 – Configuring MPLS TE - Dynamic Tunnels with Pre-emption

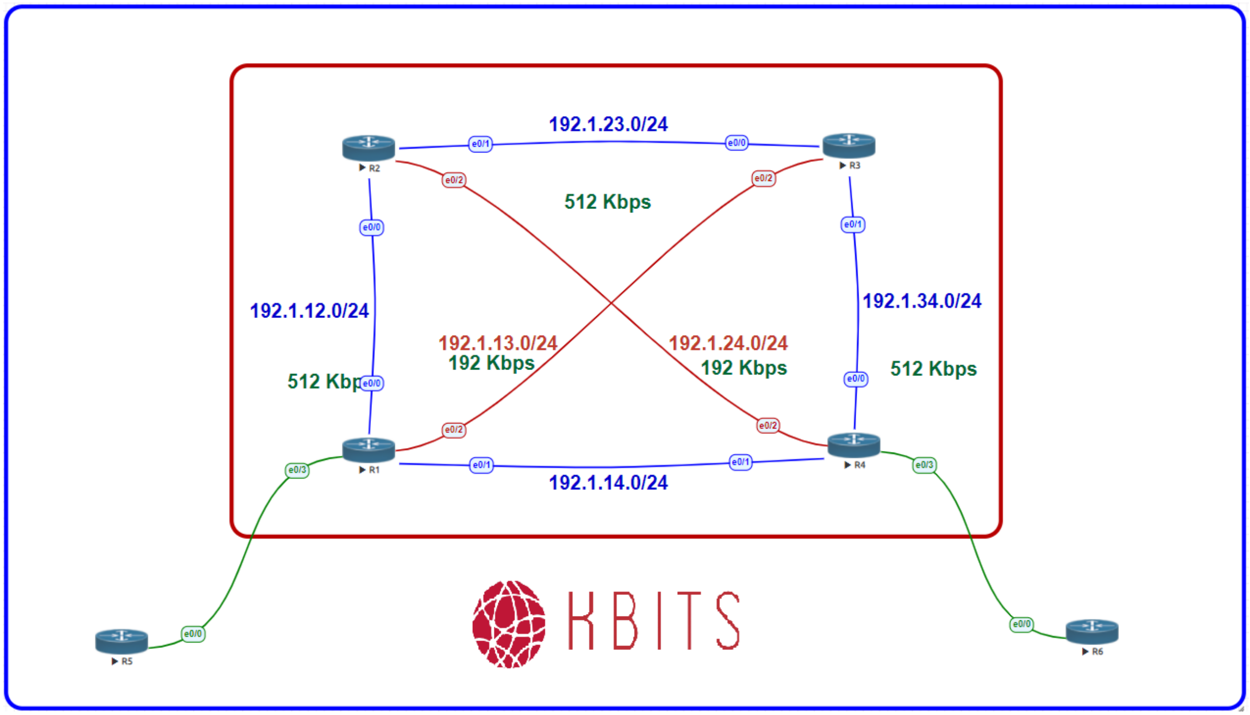

Task 5

Configure a dynamic MPLS TE Tunnel on R1 toward R3 with a bandwidth requirement of 400 kbps.

R1

```
Interface Tunnel 13
 ip unnumbered Loopback0
 tunnel destination 3.3.3.3
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 300
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng autoroute announce
```

Task 2

Verify if the tunnel came up by checking the routing table for reachability to 3.3.3.3.

R1 Show IP route **Note :** The tunnel didn't come up as the all the bandwidth was already reserved by the other Tunnels.

Task 3

Change the mpls traffic-eng priority for the Tunnel on R1 towards 3.3.3.3 to 2 for both Setup and Hold. Check to see if the Tunnel Came up.

R1 Interfae Tunnel13 Tunnel mpls traffic-eng priority 2 2 Show IP route **Note :** The tunnel did come up as the priority for this tunnel was better than the previous tunnels.

Lab 5 – Configuring MPLS TE with IS-IS

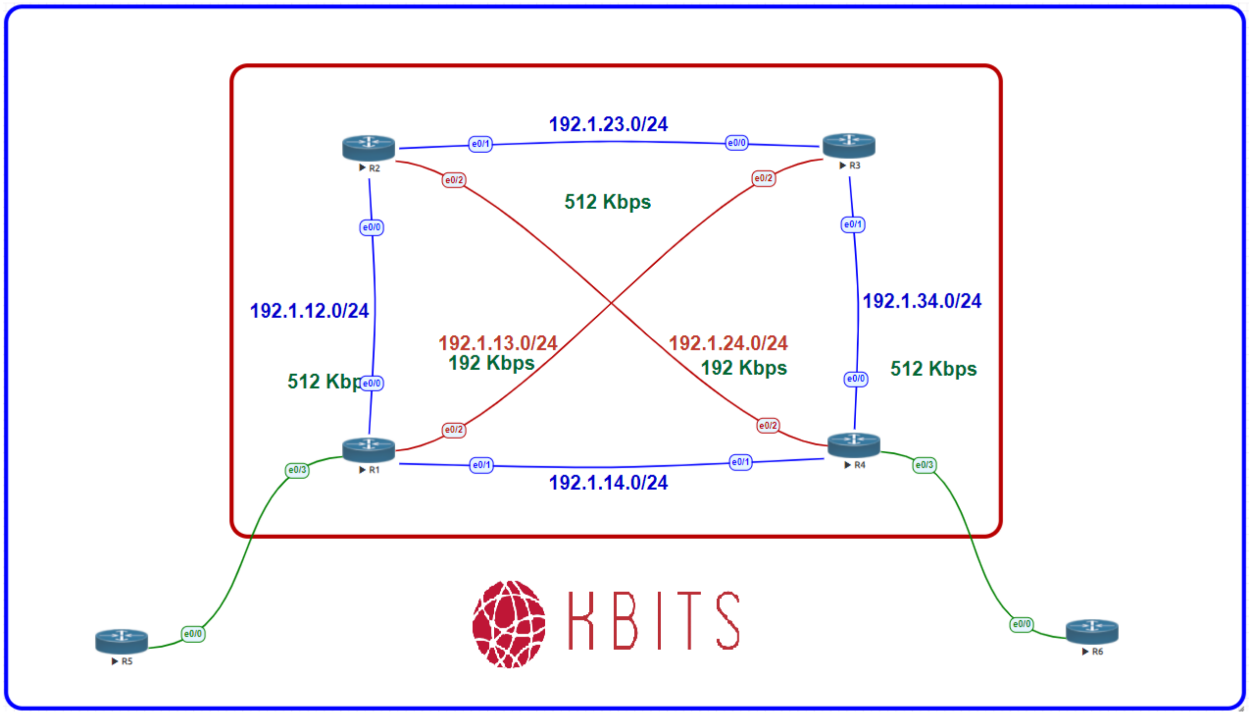

Task 1

Disable OSPF on all 4 Core Routers.

R1 No router ospf 1	**R2** No router ospf 1
R3 No router ospf 1	**R4** No router ospf 1

Task 2

Configure IS-IS on all 4 routers in a single Area 49.0000. Use XXXX.XXX.XXXX as the System ID. Advertise all the Loopbacks in IS-IS. Make sure that the Routers only establish L1 Adjacencies with each other. Enable Wide style metric to accomodate MPLS TE.

R1 Router isis Net 49.0000.1111.1111.1111.00 Is-type level-1 Metric-style wide	**R2** Router isis Net 49.0000.2222.2222.2222.00 Is-type level-1 Metric-style wide

! Int Loopback0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis Int E 0/2 Ip router isis	! Int Loopback0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis Int E 0/2 Ip router isis
R3 Router isis Net 49.0000.3333.3333.3333.00 Is-type level-1 Metric-style wide ! Int Loopback0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis Int E 0/2 Ip router isis	**R4** Router isis Net 49.0000.4444.4444.4444.00 Is-type level-1 Metric-style wide ! Int Loopback0 Ip router isis Int E 0/0 Ip router isis Int E 0/1 Ip router isis Int E 0/2 Ip router isis

Task 3

Enable MPLS Traffic Engineering IS-IS for all Level-1 neighbors. Use Loopback 0 as the router-id.

R1 router isis Mpls traffic-eng router-id Loopback0 Mpls traffic-eng level-1	**R2** router isis Mpls traffic-eng router-id Loopback0 Mpls traffic-eng level-1
R3 router isis Mpls traffic-eng router-id Loopback0 Mpls traffic-eng level-1	**R4** router isis Mpls traffic-eng router-id Loopback0 Mpls traffic-eng level-1

Task 4

Verify the status of the tunnels on R1 & R2.

R1	**R2**
Show IP route ! Sh mpls traffic-eng tunnel brief	Show IP route ! Sh mpls traffic-eng tunnel brief

Lab 6 – Configuring MPLS TE - Link Protection with Fast Re-Route (FRR)

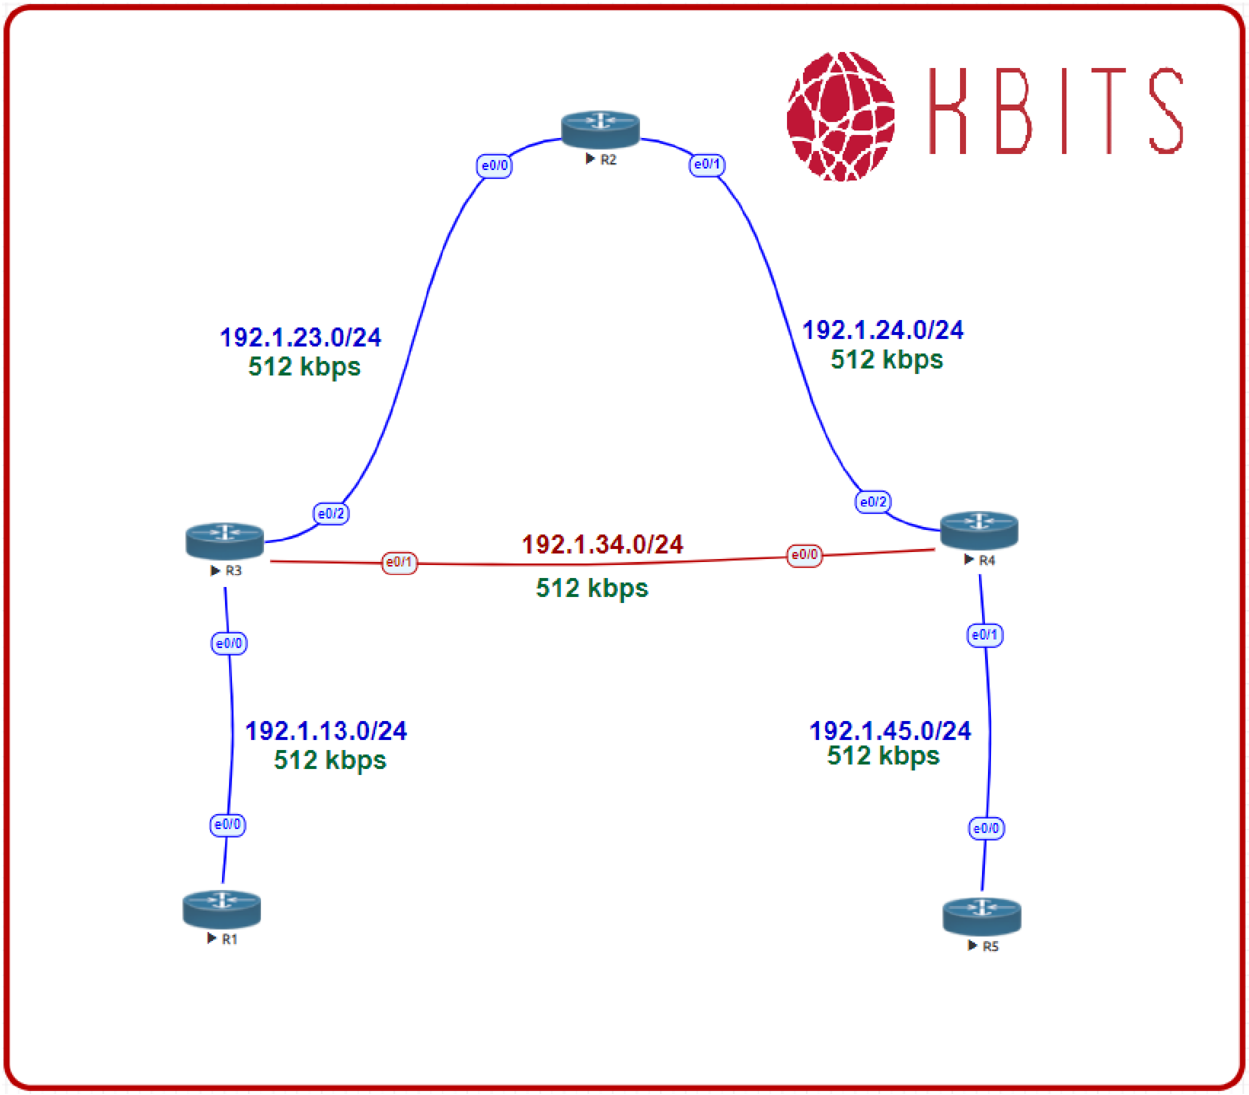

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
E 0/0	192.1.13.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
E 0/0	192.1.23.2	255.255.255.0
E 0/1	192.1.24.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
E 0/0	192.1.13.3	255.255.255.0
E 0/1	192.1.34.3	255.255.255.0
E 0/2	192.1.23.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
E 0/0	192.1.34.4	255.255.255.0
E 0/1	192.1.45.4	255.255.255.0
E 0/2	192.1.24.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	5.5.5.5	255.255.255.0
E 0/0	192.1.45.5	255.255.255.0

Task 1

Configure OSPF between all the SP routers (R1, R2, R3, R4 & R5). Use x.x.x.x as the router-id, where x is the Router number. Advertise all links in OSPF.

R1 Router ospf 1 Router-id 1.1.1.1 Network 1.1.1.1 0.0.0.0 area 0 Network 192.1.13.1 0.0.0.0 area 0	**R2** Router ospf 1 Router-id 2.2.2.2 Network 2.2.2.2 0.0.0.0 area 0 Network 192.1.23.2 0.0.0.0 area 0 Network 192.1.24.2 0.0.0.0 area 0
R3 Router ospf 1	**R4** Router ospf 1

Router-id 3.3.3.3 Network 3.3.3.3 0.0.0.0 area 0 Network 192.1.13.3 0.0.0.0 area 0 Network 192.1.23.3 0.0.0.0 area 0 Network 192.1.34.3 0.0.0.0 area 0	Router-id 4.4.4.4 Network 4.4.4.4 0.0.0.0 area 0 Network 192.1.24.4 0.0.0.0 area 0 Network 192.1.34.4 0.0.0.0 area 0 Network 192.1.45.4 0.0.0.0 area 0
R5 Router ospf 1 Router-id 5.5.5.5 Network 5.5.5.5 0.0.0.0 area 0 Network 192.1.45.5 0.0.0.0 area 0	

Task 2

Configure MPLS on all the physical links in the SP Network. The LDP neighbour relationships should be formed based on the most reliable interface.

R1 Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip	**R2** Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip
R3 Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip ! Interface E 0/2 Mpls ip	**R4** Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip ! Interface E 0/1 Mpls ip ! Interface E 0/2 Mpls ip
R5 Mpls ldp router-id Loopback0 ! Interface E 0/0 Mpls ip	

Task 3

Enable MPLS Traffic Engineering on all the Core Routers. OSPF Area 0 should also be enabled for MPLS TE. Use Loopback 0 as the router-id.

R1 Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0	**R2** Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0
R3 Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0	**R4** Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0
R5 Mpls traffic-eng tunnels ! router ospf 1 Mpls traffic-eng router-id Loopback0 Mpls traffic-eng area 0	

Task 4

Enable RSVP Bandwidth reservation & MPLS Traffic Engineering on all SP Interfaces on all SP Routers. Configure a reservation of 512 on all Fastethernet ports. Configure a reservation of 192 on all Serial ports.

R1 int E 0/0 mpls traffic-eng tunnel ip rsvp bandwidth 512	**R2** int E 0/0 mpls traffic-eng tunnel ip rsvp bandwidth 512 int E 0/1 mpls traffic-eng tunnel ip rsvp bandwidth 512
R3 int E 0/0 mpls traffic-eng tunnel ip rsvp bandwidth 512	**R4** int E 0/0 mpls traffic-eng tunnel ip rsvp bandwidth 512

int E 0/1 mpls traffic-eng tunnel ip rsvp bandwidth 512 int E 0/2 mpls traffic-eng tunnel ip rsvp bandwidth 192	int E 0/1 mpls traffic-eng tunnel ip rsvp bandwidth 512 int E 0/2 mpls traffic-eng tunnel ip rsvp bandwidth 192
R5 int E 0/0 mpls traffic-eng tunnel ip rsvp bandwidth 512	

Task 5

Configure a MPLS TE Tunnel to have R1 use the R1-> R3 -> R4 -> R5 path for Traffic destined to 5.5.5.5. This tunnel should have a bandwidth reservation requirement of 128 kbps. Enable Link Protection using the Fast Re-route feature.

R1 Interface Tunnel 1 ip unnumbered Loopback0 tunnel destination 5.5.5.5 tunnel mode mpls traffic-eng tunnel mpls traffic-eng autoroute announce tunnel mpls traffic-eng priority 4 4 tunnel mpls traffic-eng bandwidth 128 tunnel mpls traffic-eng path-option 1 explicit name R1-R5 tunnel mpls traffic-eng fast-reroute ! ip explicit-path name R1-R5 enable next-address 3.3.3.3 next-address 4.4.4.4 next-address 5.5.5.5
Notes: Make sure the Tunnel comes up and it is used in the routing table as the next hop for the 5.5.5.5 network.

Task 6

Configure a MPLS TE Tunnel on R3 to provide link protection for the R3-R4 link by using a backup path via R2 to R4. This tunnel should also have a bandwidth reservation requirement of 128 kbps. This tunnel should be used only as a backup in case of R3-R4 link going down. Configure this tunnel as the back link on the Protected Link.

R3

```
Interface Tunnel 1
 ip unnumbered Loopback0
 tunnel destination 4.4.4.4
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 4 4
 tunnel mpls traffic-eng bandwidth 128
 tunnel mpls traffic-eng path-option 1 explicit name R3-R4
!
ip explicit-path name R3-R4 enable
 next-address 2.2.2.2
 next-address 4.4.4.4
!
Interface E 0/1
 Mpls traffic-eng backup tunnel 1
```

Note: Use the **show mpls traffic-eng fast-reroute database** to check whether the FRR interface is being used or is ready for Use. Under normal circumstances, it will be in Ready state. When the protected link is down, it will do into Active state.

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring MPLS QoS

Lab 1 – Configuring Intra-AS MPLS VPN

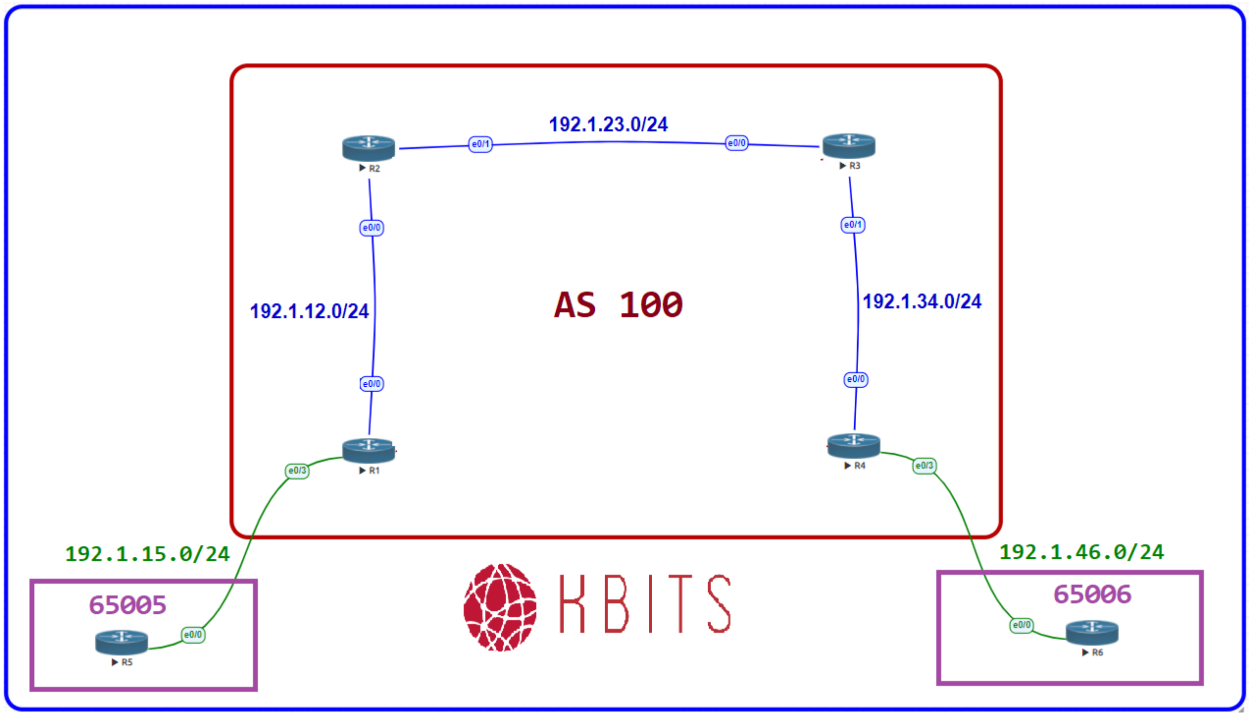

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
E 0/0	192.1.12.1	255.255.255.0
E 0/3	192.1.15.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
E 0/0	192.1.12.2	255.255.255.0
E 0/1	192.1.23.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
E 0/0	192.1.23.3	255.255.255.0
E 0/1	192.1.34.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	4.4.4.4	255.255.255.255
E 0/0	192.1.34.4	255.255.255.0
E 0/3	192.1.46.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	10.5.5.5	255.255.255.255
E 0/0	192.1.15.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	10.6.6.6	255.255.255.255
E 0/0	192.1.46.6	255.255.255.0

Task 1

Configure OSPF between all the SP routers (R1, R2, R3, R4). Use 0.0.0.X as the router-id, where x is the Router #. Advertise all links in OSPF.

R1 router ospf 1 router-id 0.0.0.1 network 192.1.12.0 0.0.0.255 area 0 network 1.0.0.0 0.255.255.255 area 0	**R2** router ospf 1 router-id 0.0.0.2 network 192.1.12.0 0.0.0.255 area 0 network 192.1.23.0 0.0.0.255 area 0 network 2.0.0.0 0.255.255.255 area 0
R3 router ospf 1 router-id 0.0.0.3 network 192.1.13.0 0.0.0.255 area 0 network 192.1.23.0 0.0.0.255 area 0 network 3.0.0.0 0.255.255.255 area 0	**R4** router ospf 1 router-id 0.0.0.4 network 192.1.14.0 0.0.0.255 area 0 network 4.0.0.0 0.255.255.255 area 0

Task 2

Configure MPLS on all the physical links in the SP Network. The LDP neighbour relationships should be formed based on the most reliable interface.

R1 Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip	**R2** Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip
R3 Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip ! Interface E 0/1 mpls ip	**R4** Mpls ldp router-id Loopback 0 ! Interface E 0/0 mpls ip

Task 3

Configure MP-iBGP between the PE Routers (R1 & R4).

R1
 	
router bgp 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 4.4.4.4 update-source lo0
!
 address-family vpnv4
  neighbor 4.4.4.4 activate
	R4
	
	router bgp 100
	 neighbor 1.1.1.1 remote-as 100
	 neighbor 1.1.1.1 update-source lo0
	!
	 address-family vpnv4
	  neighbor 1.1.1.1 activate

Task 4

Configure a VRF called Cust-A on the PE Routers (R1 & R4). Use the RD & RT values as 100:1. Configure the interfaces between the PE-CE based on the diagram. Configure BGP as the PE-CE Routing protocol. R5 will be in AS 65005 and R6 will be in AS 65006.

R1

```
vrf definition Cust-A
 rd 100:1
 address-family ipv4
  route-target both 100:1
!
Interface E 0/3
 vrf forwarding Cust-A
 ip address 192.1.15.1 255.255.255.0
 no shut
!
router bgp 100
 address-family ipv4 vrf Cust-A
   neighbor 192.1.15.5 remote-as 65005
```

R4

```
vrf definition Cust-A
 rd 100:1
 address-family ipv4
  route-target both 100:1
!
Interface E 0/3
 vrf forwarding Cust-A
 ip address 192.1.46.4 255.255.255.0
 no shut
!
router bgp 100
 address-family ipv4 vrf Cust-A
   neighbor 192.1.46.6 remote-as 65006
```

Task 5

Configure the CE Routers (R5 & R6). Configure BGP as the PE-CE Routing protocol. R5 will be in AS 65005 and R6 will be in AS 65006. Advertise loopback0 interface of R5 & R6 in BGP.

R5

```
Interface E 0/0
 ip address 192.1.15.5 255.255.255.0
 no shut
!
Interface loopback0
 ip address 10.5.5.5 255.255.255.0
!
router bgp 65005
 neighbor 192.1.15.1 remote-as 100
 network 10.5.5.0 mask 255.255.255.0
```

R6

```
Interface E 0/0
 ip address 192.1.46.6 255.255.255.0
 no shut
!
Interface loopback0
 ip address 10.6.6.6 255.255.255.0
!
router bgp 65006
 neighbor 192.1.46.6 remote-as 100
 network 10.6.6.0 mask 255.255.255.0
```

Lab 2 – Configuring MPLS QoS – Uniform Mode

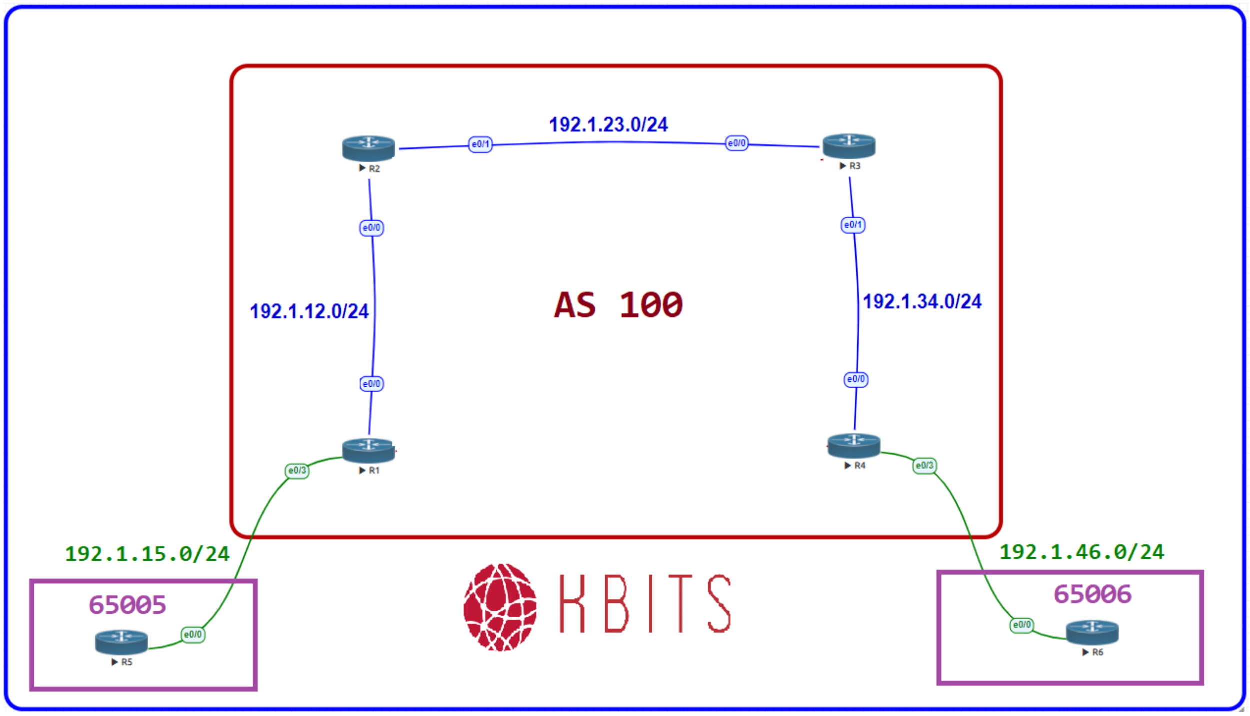

Agreement between Customer & SP

- All traffic with a PREC of 5 should be prioritized to 1 mbps
- All traffic with a PREC of 3 should be policied to 256kbps

Task 1

Configure the CE router (R5) mark Telnet traffic with a Precedence of 5 so that the SP can prioritize the traffic. Also, configure the router to mark ICMP traffic with a precedence of 3 so that the SP can police the traffic.

R5

```
access-list 101 permit tcp any any eq 23
access-list 102 permit icmp any any
!
class-map CM-TELNET
 match access-group 101
class-map CM-ICMP
 match access-group 102
```

```
!
policy-map PM-QOS
 class CM-TELNET
   priority percent 10
   set ip precedence 5
 class CM-ICMP
   police 256000
   set ip precedence 3
!
Interface E 0/0
 service-policy output PM-QOS
```

Task 2

Configure the PE router (R1) to copy the appropriate CE Markings into the top label.

```
R1

class-map IP-PREC-3
 match ip precedence 3
!
class-map IP-PREC-5
 match ip precedence 5
!
policy-map SET-MPLS-EXP
 class IP-PREC-3
   set mpls experimental imposition 3
 class IP-PREC-5
   set mpls experimental imposition 5
!
Interface E 0/3
 service-policy input SET-MPLS-EXP
```

Task 3

Configure the PE router (R1) to use the markings copied to the top label in the previous task to fulfil the Customer SLA.

```
R1

class-map MPLS-PREC-3
 match mpls experimental topmost 3
!
class-map MPLS-PREC-5
 match mpls experimental topmost 5
```

```
!
policy-map QOS-TO-R2
 class MPLS-PREC-3
   police 256000
 class MPLS-PREC-5
   priority percent 10
!
Interface E 0/0
 service-policy output QOS-TO-R2
```

Task 4

The SP Core network uses Precedence 4 for Policing & Precedence 7 for Priority. Configure the P Router (R2) to change the topmost marking to adhere to the SP QoS Policy.

R2

```
class-map MPLS-PREC-3
 match mpls experimental topmost 3
!
class-map MPLS-PREC-5
 match mpls experimental topmost 5
!
policy-map SET-MPLS-EXP
 class MPLS-PREC-3
   set mpls experimental topmost 4
 class MPLS-PREC-5
   set mpls experimental topmost 7
!
Interface E 0/0
 service-policy input SET-MPLS-EXP
```

Task 5

Configure the PE router (R2) to use the markings to set the SP QoS Policy.

R2

```
class-map MPLS-PREC-4
 match mpls experimental topmost 4
!
class-map MPLS-PREC-7
 match mpls experimental topmost 7
!
policy-map QOS-TO-R3
```

```
 class MPLS-PREC-4
   police 256000
 class MPLS-PREC-7
   priority percent 10
!
Interface E 0/1
 service-policy output QOS-TO-R3
```

Task 6

R3 will perform the PHP for R4 routes. Copy the Top Label marking so that they can be used in the QoS Policy between R3 & R4.

R3

```
class-map MPLS-PREC-4
 match mpls experimental topmost 4
!
class-map MPLS-PREC-7
 match mpls experimental topmost 7
!
policy-map SET-QOS-GROUP
 class MPLS-PREC-4
   set qos-group mpls experimental topmost
 class MPLS-PREC-7
   set qos-group mpls experimental topmost
!
Interface E 0/0
 service-policy input SET-QOS-GROUP
```

Task 7

Impose the marking copied from the top label to the next label and configure the QoS Policy on the link between R3 & R4.

R3

```
class-map QOS-PREC-4
 match qos-group 4
!
class-map QOS-PREC-7
 match qos-group 7
!
policy-map USE-QOS-GROUP
 class QOS-PREC-4
   police 256000
```

```
   set mpls experimental topmost qos-group
 class QOS-PREC-7
   priority percent 10
   set mpls experimental topmost qos-group
!
Interface E 0/1
 service-policy output USE-QOS-GROUP
```

Task 8

Configure MPLS QoS – Uniform mode on the PE-CE leg (R4-R6). Copy the SP marking onto the Customer Packet and use the marking on the QoS Policy on this link.

R4

Copy down the MPLS EXP bits

```
class-map MPLS-PREC-4
  match mpls experimental topmost 4
!
class-map MPLS-PREC-7
  match mpls experimental topmost 7
!
policy-map SET-QOS-GROUP
 class MPLS-PREC-4
   set qos-group mpls experimental topmost
 class MPLS-PREC-7
   set qos-group mpls experimental topmost
!
Interface E 0/0
 service-policy input SET-QOS-GROUP
```

Push down the MPLS EXP value to the IP Packet

```
class-map QOS-PREC-4
 match qos-group 4
!
class-map QOS-PREC-7
 match qos-group 7
!
policy-map USE-QOS-GROUP
 class QOS-PREC-4
   police 128000
   set precedence qos-group
```

```
 class QOS-PREC-7
   priority percent 10
   set precedence qos-group
!
Interface E 0/3
 service-policy output USE-QOS-GROUP
```

Lab 3 – Configuring MPLS QoS – Long Pipe Mode

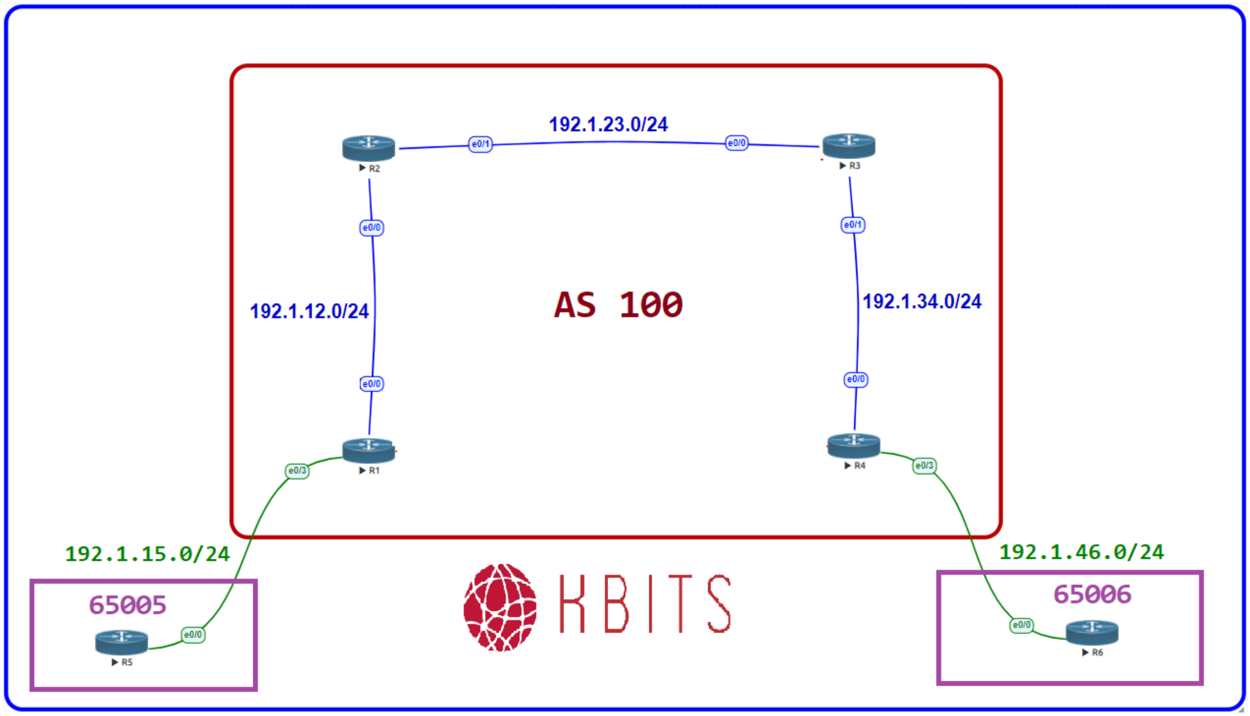

Agreement between Customer & SP

- All traffic with a PREC of 5 should be prioritized to 1 mbps
- All traffic with a PREC of 3 should be policied to 256kbps

Task 1

You will be configuring the MPLS QoS – Long Pipe mode on the PE-CE leg (R4-R6). De-configure the Class-map & Policy-maps from the previous labs on R4.

R4

```
Interface E 0/0
 No service-policy input SET-QOS-GROUP
!
Interface E 0/3
 No service-policy output USE-QOS-GROUP
!
No policy-map SET-QOS-GROUP
No policy-map USE-QOS-GROUP
```

```
!
No class-map MPLS-PREC-4
No class-map MPLS-PREC-7
No class-map QOS-PREC-4
No class-map QOS-PREC-7
```

Task 2

Configure MPLS QoS – Long Pipe mode on the PE-CE leg (R4-R6). Copy the SP marking in the QoS-Group. Do not copy it to the Customer Packet. Use the SP marking on the QoS Policy on this link.

R4

Copy down the MPLS EXP bits

```
class-map MPLS-PREC-4
  match mpls experimental topmost 4
!
class-map MPLS-PREC-7
  match mpls experimental topmost 7
policy-map SET-QOS-GROUP
 class MPLS-PREC-4
   set qos-group mpls experimental topmost
   set discard-class 4
 class MPLS-PREC-7
   set qos-group mpls experimental topmost
   set discard-class 7
!
Interface E 0/0
 service-policy input SET-QOS-GROUP
```

Use the QOS Group for QoS. Don't Set it on the IP Packet

```
class-map QOS-PREC-4
 match qos-group 4
!
class-map QOS-PREC-7
 match qos-group 7
!
policy-map USE-QOS-GROUP
 class QOS-PREC-4
   police 128000
 class QOS-PREC-7
```

```
   priority percent 10
!
Interface E 0/3
 service-policy output USE-QOS-GROUP
```

Lab 4 – Configuring MPLS QoS – Short Pipe Mode

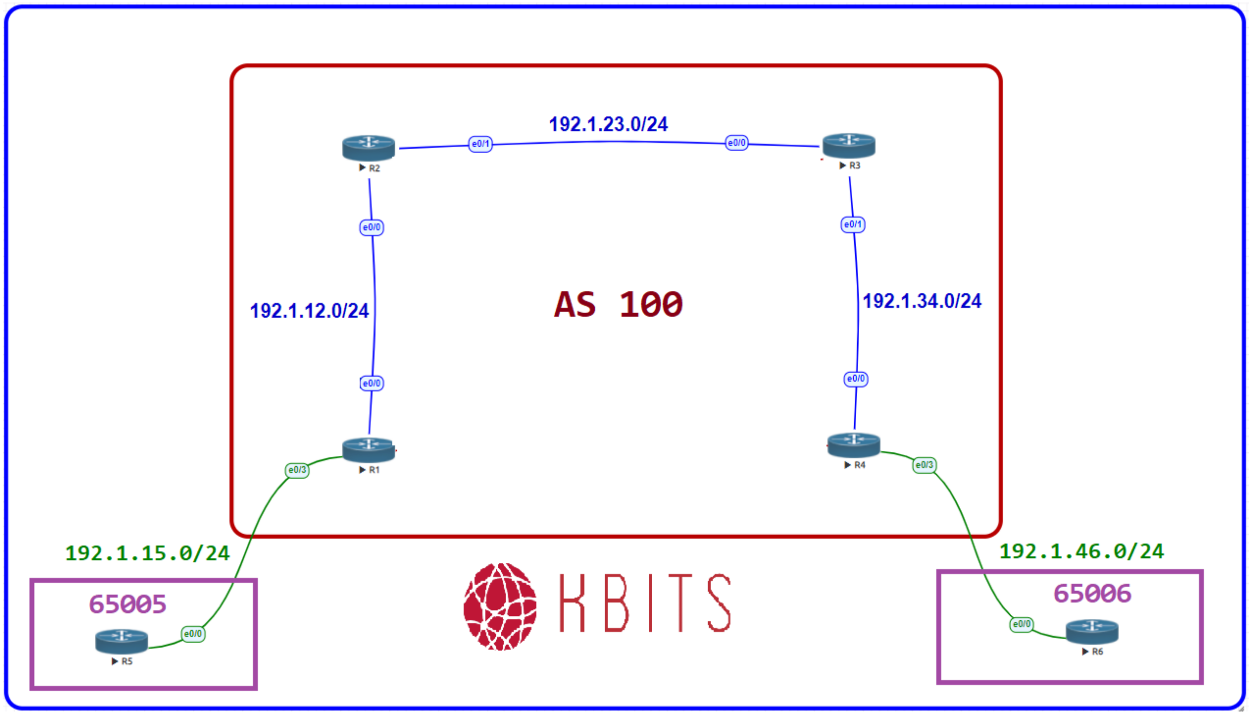

Agreement between Customer & SP

- All traffic with a PREC of 5 should be prioritized to 1 mbps
- All traffic with a PREC of 3 should be policied to 256kbps

Task 1

You will be configuring the MPLS QoS – Short Pipe mode on the PE-CE leg (R4-R6). De-configure the Class-map & Policy-maps from the previous labs on R4.

R4

```
Interface E 0/0
 No service-policy input SET-QOS-GROUP
!
Interface E 0/3
 No service-policy output USE-QOS-GROUP
!
No policy-map SET-QOS-GROUP
No policy-map USE-QOS-GROUP
```

```
!
No class-map MPLS-PREC-7
No class-map QOS-PREC-4
No class-map QOS-PREC-7
```

Task 2

Configure MPLS QoS – Short Pipe mode on the PE-CE leg (R4-R6). There is not need to work on the SP Markings. Use the Customer markings on the QoS Policy on this link.

R4

```
class-map QOS-PREC-3
 match precedence 3
!
class-map QOS-PREC-5
 match precedence 5
!
policy-map USE-QOS-GROUP
 class QOS-PREC-3
   police 256000
 class QOS-PREC-5
   priority percent 10
!
Interface E 0/3
 service-policy output USE-QOS-GROUP
```

CCIE Service Provider Workbook

Authored By:

Khawar Butt
Hepta CCIE # 12353
CCDE # 20110020

Configuring Layer 2 MPLS VPNs

Lab 1 – Implementing AToM – Ethernet VLAN over MPLS

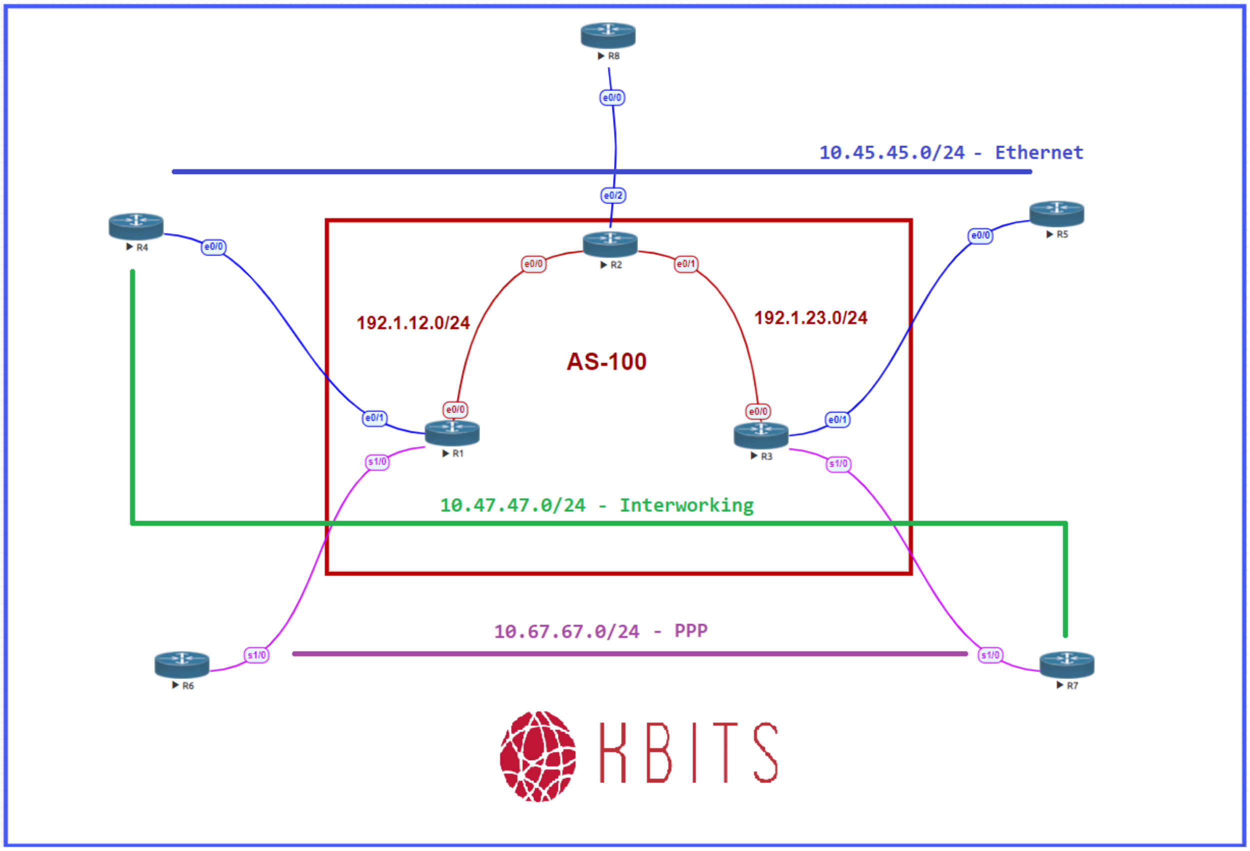

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
F 0/0.12	192.1.12.1	255.255.255.0

R2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
F 0/0.12	192.1.12.2	255.255.255.0
F 0/0.23	192.1.23.2	255.255.255.0

R3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
F 0/0.23	192.1.23.3	255.255.255.0

Task 1

Configure OSPF between all the SP routers in AS 00 (R1, R2, R3). Hard code the Router-ID's. Advertise all links in OSPF except the Links between CE-PE's.

R1 Router ospf 1 Router-id 1.1.1.1 Network 1.1.1.1 0.0.0.0 area 0 Network 192.1.12.1 0.0.0.0 area 0	**R2** Router ospf 1 Router-id 2.2.2.2 Network 2.2.2.2 0.0.0.0 area 0 Network 192.1.12.2 0.0.0.0 area 0 Network 192.1.23.2 0.0.0.0 area 0
R3 Router ospf 1 Router-id 3.3.3.3 Network 3.3.3.3 0.0.0.0 area 0 Network 192.1.23.3 0.0.0.0 area 0	

Task 2

Configure MPLS on all the physical links in the SP AS Network. Use LDP to distribute labels. The LDP neighbour relationships should be formed based on the most reliable interface.

<table>
<tr><td>R1

Mpls ldp router-id Loopback0
!
Mpls ip
!
Interface E 0/0
 Mpls ip</td><td>R2

Mpls ldp router-id Loopback0
!
Mpls ip
!
Interface E 0/0
 Mpls ip
!
Interface E 0/1
 Mpls ip</td></tr>
<tr><td>R3

Mpls ldp router-id Loopback0
!
Mpls ip
!
Interface E 0/0
 Mpls ip</td><td></td></tr>
</table>

Task 3

You need to connect R4 to R5 using Ethernet over MPLS. Configure R4 and R5 to be part of VLAN 45. Configure them with an IP address of 10.45.45.X/24, where X is the 4 for R4 and 5 R5. Router EIGRP in AS 45 between them. Advertise loopback interfaces on R4(10.4.4.4/24) & R5 (10.5.5.5/24) to each other.

<table>
<tr><td>R4

Interface loopback0
 Ip address 10.4.4.4 255.255.255.0
!
Interface E 0/0
 No shut
!
Interface E 0/0.45
 Encapsulation dot1q 45
 Ip address 10.45.45.4 255.255.255.0
!</td><td>R5

Interface loopback0
 Ip address 10.5.5.5 255.255.255.0
!
Interface E 0/0
 No shut
!
Interface E 0/0.45
 Encapsulation dot1q 45
 Ip address 10.45.45.5 255.255.255.0
!</td></tr>
</table>

router eigrp 45 Network 10.0.0.0	router eigrp 45 Network 10.0.0.0

Task 4

Configure the PE Routers R1 and R3 to connect CE routers R4 & R5 to each other. Use 145 as the PVC. Use VLAN 45 on the Sub-interfaces.

R1 Interface E 0/1.45   encapsulation dot1Q 45 xconnect 3.3.3.3 145 encapsulation mpls
R3 Interface E 0/1.45   encapsulation dot1Q 45 xconnect 1.1.1.1 145 encapsulation mpls

Lab 2 – Implementing AToM – Ethernet over MPLS

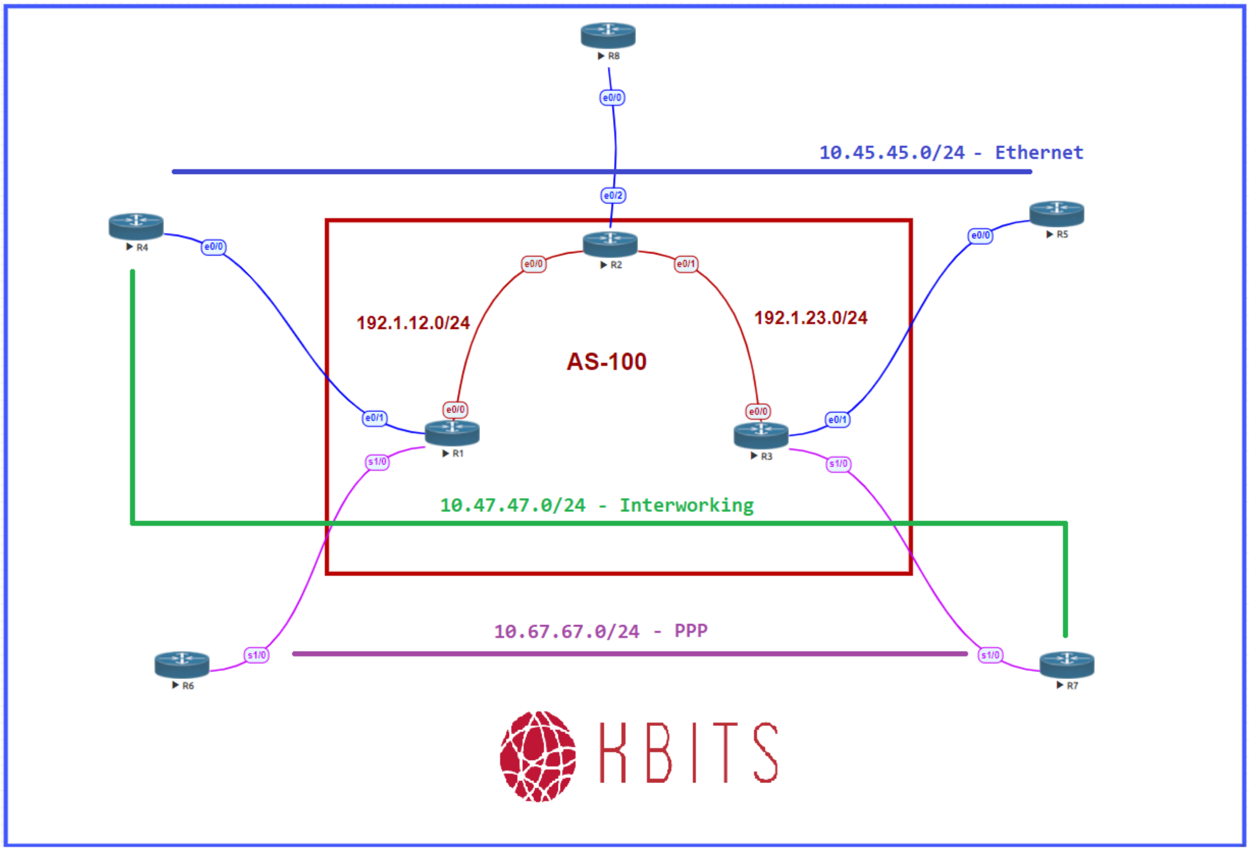

Task 1

Default the PE-CE Interfaces on R1, R3, R4 & R5 routers configured in the previous lab.

R1	**R3**
Default interface E 0/1 Default interface E 0/1.45	Default interface E 0/1 Default interface E 0/1.45
R4	**R5**
Default interface E 0/0 Default interface E 0/0.45	Default interface E 0/0 Default interface E 0/0.45

Task 2

You need to connect R4 to R5 using Ethernet over MPLS. Configure R4 and R5 with an IP address of 10.45.45.X/24, where X is the 4 for R4 and 5 R5 on the Physical PE-CE Interface.

R4	**R5**
Interface E 0/0 Ip address 10.45.45.4 255.255.255.0 No shut	Interface E 0/0 Ip address 10.45.45.5 255.255.255.0 No shut

Task 4

Configure the PE Routers R1 and R3 to connect CE routers R4 & R5 to each other. Use 145 as the PVC.

R1
Interface E 0/1 xconnect 3.3.3.3 145 encapsulation mpls
R3 Interface E 0/1 xconnect 1.1.1.1 145 encapsulation mpls

Lab 3 – Implementing AToM – PPP over MPLS

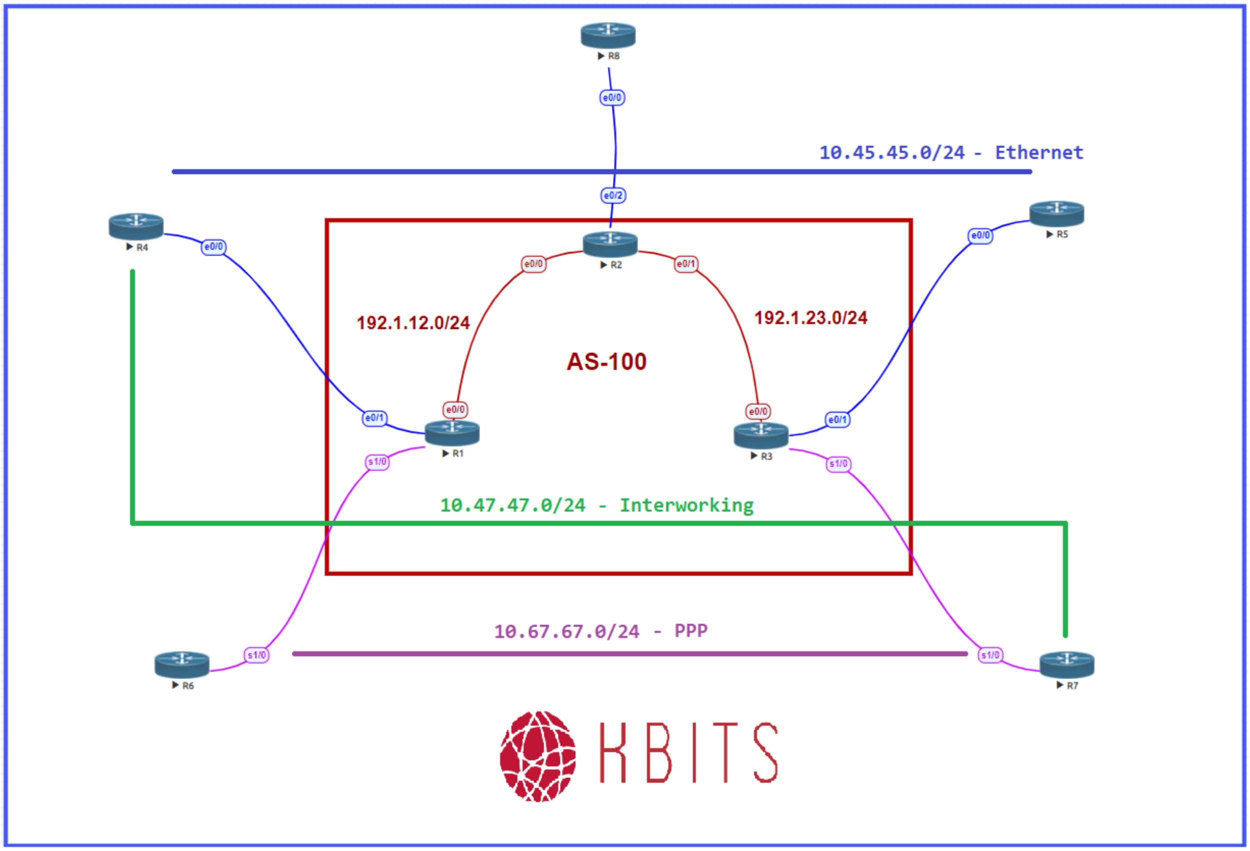

Task 1

You need to connect R6 to R7 using PPP over MPLS. Configure R6 and R7 with an IP address of 10.67.67.X/24 on the S 1/0 interface, where X is the 6 for R6 and 7 R7. Router EIGRP in AS 67 between them. Advertise loopback interfaces on R6(10.6.6.6/24) & R7 (10.7.7.7/24) to each other.
.

R6	**R7**
Interface Loopback0 Ip address 10.6.6.6 255.255.255.0 ! Interface S 1/0 Ip address 10.67.67.6 255.255.255.0	Interface Loopback0 Ip address 10.7.7.7 255.255.255.0 ! Interface S 1/0 Ip address 10.67.67.7 255.255.255.0

Encapsulation ppp No shut ! Rotuer eigrp 67 Network 10.0.0.0	Encapsulation ppp No shut ! Rotuer eigrp 67 Network 10.0.0.0

Task 2

Configure the PE Routers R1 and R3 to connect CE routers R6 & R7 to each other. Use 167 as the PVC.

R1 Interface S 1/0 Encapsulation ppp xconnect 3.3.3.3 167 encapsulation mpls no shut
R3 Interface S 1/0 Encapsulation ppp xconnect 1.1.1.1 167 encapsulation mpls no shut

Lab 4 – Implementing AToM – Interworking Ethernet & PPP

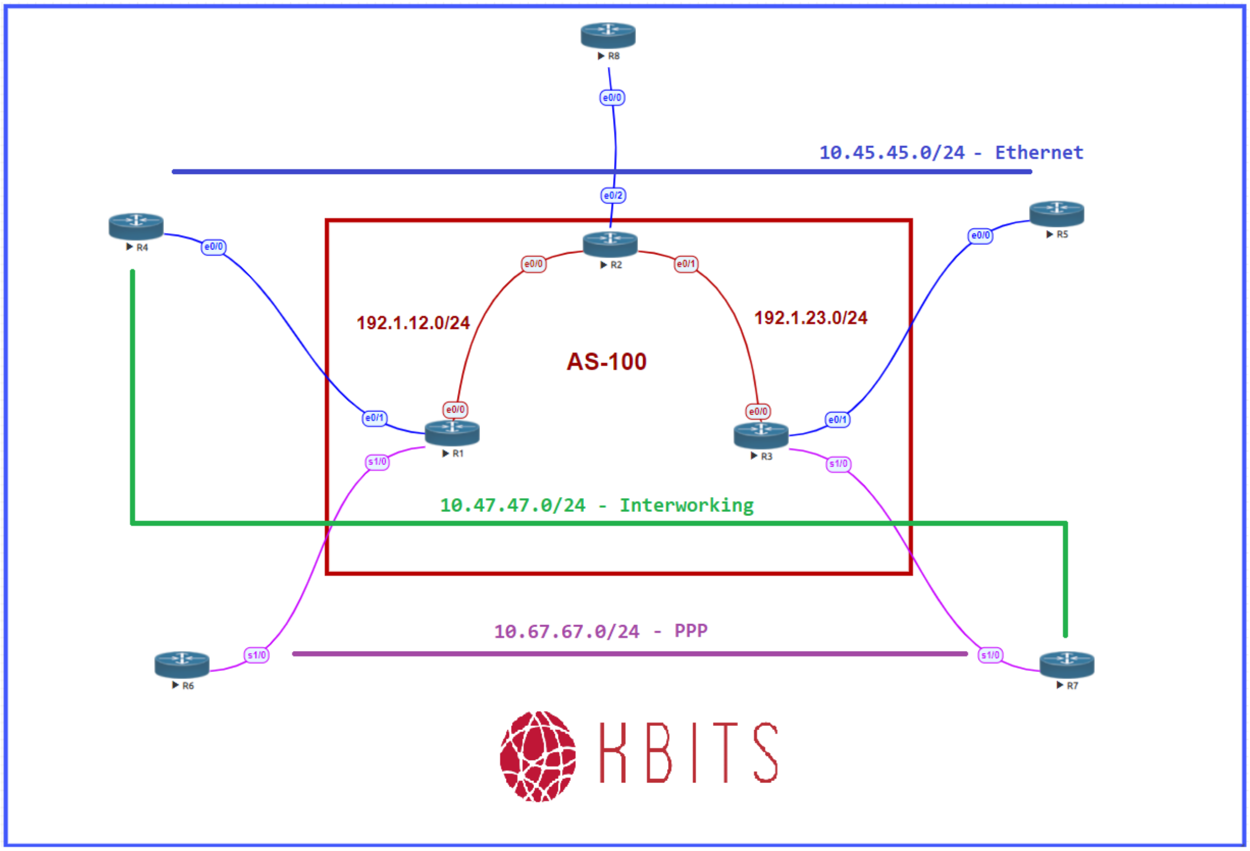

Task 1

Default the PE-CE Interfaces on R1, R3, R4, R5, R6 & R7 routers configured in the previous lab.

R1 Default interface E 0/1 Default interface S 1/0	**R3** Default interface E 0/1 Default interface S 1/0
R4 Default interface E 0/0	**R5** Default interface E 0/0
R6 Default interface S 1/0	**R7** Default interface S 1/0

Task 2

You need to connect R4 to R7 using Ethernet/PPP over MPLS. Configure R4 with an IP Address of 10.47.47.4/24 on the E 0/0 Interface. Configure R7 with an IP address of 10.47.47.7/24 on the S 1/0 interface. Configure it with an encapsulation of PPP.

R4	R7
Interface E 0/0 Ip address 10.47.47.4 255.255.255.0 No shut	Interface S 1/0 Ip address 10.47.47.7 255.255.255.0 encapsulation No shut

Task 3

Configure the PE Routers R1 and R3 to connect CE routers R4 & R7 to each other using Interworking MPLS. Use 147 as the PVC ID.

R1

```
pseudowire-class IW-R4-R7
 encapsulation mpls
 interworking ip
!
Interface E0/1
 xconnect 3.3.3.3 147 pw-class IW-R4-R7
 no shut
```

R3

```
pseudowire-class IW-R4-R7
 encapsulation mpls
 interworking ip
!
Interface S1/0
 encap ppp
 xconnect 1.1.1.1 147 pw-class IW-R4-R7
 no shut
```

Lab 5 – Configuring VPLS

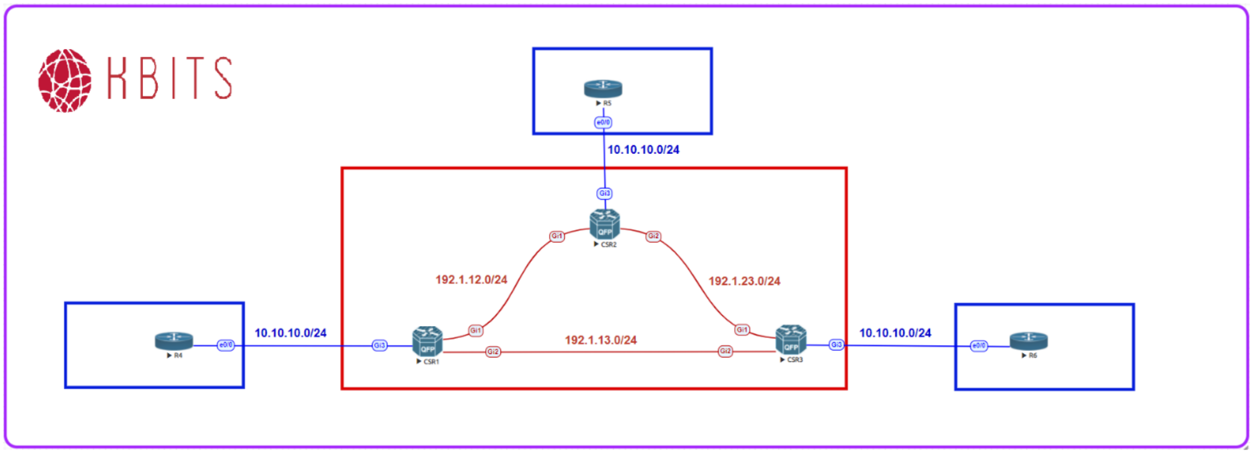

Interface IP Address Configuration

CSR1

Interface	IP Address	Subnet Mask
Loopback 0	1.1.1.1	255.255.255.255
Gig1	192.1.12.1	255.255.255.0
Gig2	192.1.13.1	255.255.255.0

CSR2

Interface	IP Address	Subnet Mask
Loopback 0	2.2.2.2	255.255.255.255
Gig1	192.1.12.2	255.255.255.0
Gig2	192.1.23.2	255.255.255.0

CSR3

Interface	IP Address	Subnet Mask
Loopback 0	3.3.3.3	255.255.255.255
Gig1	192.1.23.3	255.255.255.0
Gig2	192.1.13.3	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Loopback 0	10.4.4.4	255.255.255.0
E 0/0.1	10.10.10.4	255.255.255.0

R5

Interface	IP Address	Subnet Mask
Loopback 0	10.5.5.5	255.255.255.0
E 0/0.1	10.10.10.5	255.255.255.0

R6

Interface	IP Address	Subnet Mask
Loopback 0	10.6.6.6	255.255.255.0
E 0/0.1	10.10.10.6	255.255.255.0

Task 1

Configure the SP routers (CSR1, CSR2 & CSR3) with Interface IP's based on the diagram. Configure OSPF with the core. Enable LDP on the SP Core Interfaces.

CSR1

```
Interface Gig1
 ip address 192.1.12.1 255.255.255.0
 mpls ip
 no shut
!
Interface Gig2
 ip address 192.1.13.1 255.255.255.0
 mpls ip
 no shut
!
Interface Loopback10
 ip address 1.1.1.1 255.255.255.255
!
mpls ldp router-id Loopback10
!
router ospf 1
 router-id 0.0.0.1
 network 192.1.12.0 0.0.0.255 area 0
 network 192.1.13.0 0.0.0.255 area 0
 network 1.0.0.0 0.255.255.255 area 0
```

CSR2

```
Interface Gig1
 ip address 192.1.12.2 255.255.255.0
 mpls ip
 no shut
!
Interface Gig2
 ip address 192.1.23.2 255.255.255.0
 mpls ip
 no shut
!
Interface Loopback10
 ip address 2.2.2.2 255.255.255.255
!
mpls ldp router-id Loopback10
!
router ospf 1
 router-id 0.0.0.2
 network 192.1.12.0 0.0.0.255 area 0
 network 192.1.23.0 0.0.0.255 area 0
 network 2.0.0.0 0.255.255.255 area
```

CSR3 Interface Gig1 ip address 192.1.23.3 255.255.255.0 mpls ip no shut ! Interface Gig2 ip address 192.1.13.3 255.255.255.0 mpls ip no shut ! Interface Loopback10 ip address 3.3.3.3 255.255.255.255 ! mpls ldp router-id Loopback10 ! router ospf 1 router-id 0.0.0.3 network 192.1.13.0 0.0.0.255 area 0 network 192.1.23.0 0.0.0.255 area 0 network 3.0.0.0 0.255.255.255 area 0	

Task 2

Connect the CE sites to each other using VPLS. Use a Bridge-Domain ID of 100. Create a VFI of CUSTA with a VPN ID os 111 and Bridge Domain of 100. Use VLAN 10 on the PE-CE Interface.

<table>
<tr>
<td>

CSR1

```
interface GigabitEthernet3
 no shut
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 100
!
l2 vfi CUSTA manual
 vpn id 111
 bridge-domain 100
 neighbor 3.3.3.3 encapsulation mpls
 neighbor 2.2.2.2 encapsulation mpls
```

</td>
<td>

CSR2

```
interface GigabitEthernet3
 no shut
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 100
!
l2 vfi CUSTA manual
 vpn id 111
 bridge-domain 100
 neighbor 1.1.1.1 encapsulation mpls
 neighbor 3.3.3.3 encapsulation mpls
```

</td>
</tr>
<tr>
<td>

CSR3

```
interface GigabitEthernet3
 no shut
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 100
!
l2 vfi CUSTA manual
 vpn id 111
 bridge-domain 100
 neighbor 10.1.1.1 encapsulation mpls
 neighbor 10.2.2.2 encapsulation mpls
```

</td>
<td></td>
</tr>
</table>

Task 3

Configure the CE sites with a Sub-interface on the E 0/0 interface using a VLAN of 10. Configure the interfaces based on the table above. Run EIGRP in AS 123 to route the loopbacks.

<table>
<tr>
<td>

R4

```
Interface E 0/0
 no shut
Interface E 0/0.1
 encapsulation dot1q 10
 ip address 10.10.10.4 255.255.255.0
```

</td>
<td>

R5

```
Interface E 0/0
 no shut
Interface E 0/0.1
 encapsulation dot1q 10
 ip address 10.10.10.5 255.255.255.0
```

</td>
</tr>
</table>

no shut ! Interface loopback 0 ip address 10.4.4.4 255.255.255.0 ! router eigrp 100 network 10.0.0.0	no shut ! Interface loopback 0 ip address 10.5.5.5 255.255.255.0 ! router eigrp 100 network 10.0.0.0
R6 Interface E 0/0 no shut Interface E 0/0.1 encapsulation dot1q 10 ip address 10.10.10.6 255.255.255.0 no shut ! Interface loopback 0 ip address 10.6.6.6 255.255.255.0 ! router eigrp 100 network 10.0.0.0	