Cisco ASR 903 Router Design and Deployment Guide

Version 1.0

Last Updated: September 11, 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-30526-01

CONTENTS

Introduction

This document provides design guidance, and detailed configurations for deployment of the Cisco ASR 903 in service provider carrier Ethernet and mobile backhaul network. The design guide examples are created to help people who work on Cisco ASR 903 configuration.

Chapter 1 Cisco ASR 903 Chassis

The Cisco ASR 903 Router has the following hardware features:

- 3-RU modular chassis designed for installation in a 300 mm European Telecommunications Standards Institute (ETSI) cabinet
- Up to six interface modules
- Up to two Route Switch Processors (RSP)
- Up to two DC or AC power supply units
- One fan tray

```
Router# show platform
Chassis type: ASR-903
Slot      Type                State            Insert time (ago)
--------- ------------------- ---------------- -----------------
 0/0      A900-IMA8T          ok                 3d17h        ! bottom slot, 8 CU
 0/1      A900-IMA8S          ok                 3d17h        ! 8 SFP ports IM
 0/2      A900-IMA1X          ok                 3d17h        ! 10GE IM
 0/3      A900-IMA1X          ok                 3d17h
 0/4      A900-IMA16D         ok                 3d17h        ! 16 T1/E1 IM
 0/5      A900-IMA4OS         ok                 3d17h        ! top slot, OC-3 IM
R0        A903-RSP1A-55       ok, active         3d17h        ! RSP R0, bottom RSP
R1        A903-RSP1A-55       ok, standby        3d17h        ! RSP R1, top RSP
F0                             ok, active        3d17h
F1                             ok, standby       3d17h
P0        A900-PWR550-D       ok                 3d17h         ! power supply
P1        A900-PWR550-D       ok                 3d17h
P2        A903-FAN             ok                3d17h        ! fan tray

Slot      CPLD Version        Firmware Version
--------- ------------------- ----------------------------------------
R0        11102133            15.3(2r)S
R1        11102133            15.3(2r)S
F0        11102133            15.3(2r)S
F1        11102133            15.3(2r)S
```

Chapter 2 Layer 2 Solution

Ethernet Flow Point (EFP)

An Ethernet flow point (EFP) service instance is a logical interface that connects a bridge domain to a physical port or to an EtherChannel group. An incoming frame is matched against EFP matching criteria on the interface, learned on the matching EFP, and forwarded to one or more EFPs in the bridge domain. If there are no matching EFPs, the frame is dropped. EFP is the building block for Ethernet Virtual Circuit, bridged EVC hop-by-hop, EoMPLS or VPLS. An EFP can be viewed as switch access port but more powerful as it can manipulate the frames by removing and imposing vlan tags. The following EFP matches incoming dot1q vlan 10, removes the vlan tag and forwards into bridge-domain 10. On egress direction, a frame will be imposed vlan 10 before sending out on the wire.

```
interface GigabitEthernet0/2/0
 service instance 2 ethernet
  encapsulation dot1q 10
  rewrite ingress tag pop 1 symmetric
  bridge-domain 10
```

Alternatively with new CLI, the above can also be configured as below. Note that the bridge-domain does not have to be same as vlan tag and it has been purposely changed to 30 in the following example.

```
bridge-domain 30
 member GigabitEthernet0/2/0 service-instance 2
interface GigabitEthernet0/2/0
 service instance 2 ethernet
  encapsulation dot1q 10
  rewrite ingress tag pop 1 symmetric
```

Table 1. Encapsulation types

Encapsulation	Description
encapsulation untagged	Native Packets
encapsulation default	Catch-all Encap.
encasulationp priority-tagged	Dot1p Packets.
encapsulation dot1q any	All packets with one or more vlans
encapsulation dot1q *vlan*	Exact Outermost match.
encapsulation dot1q *vlan* **(second-)dot1q** *vlan2*	Exact Outermost two tags.
encapsulation dot1q *range*	Ranged Outermost match.
encapsulation dot1q *range* **(second-)dot1q** *vlan2*	Ranged Outermost match, exact second tag
encapsulation dot1q *range* **(second-)dot1q** *range*	Ranged Outermost two tags.
encapsulation dot1q *vlan* **(second-)dot1q** *range*	Exact Outermost match, ranged second tag
encapsulation dot1q *vlan* **cos** *cos*	Exact Outermost vlan, cos match.
encapsulation dot1q *range* **cos** *cos* **(second-)dot1q** *range* **cos** *cos*	Ranged combinations of outermost vlan and second tag

The following rewrite can be implemented on Cisco ASR 903.

- Pop 1, remove one vlan tag
- Pop 2, remove two vlan tags
- Push 1, impose one vlan tag

  With the above push and pop function, vlan 1 to 1, 1 to 2, 2 to 1, and 2 to 2 mapping can be achieved.

Bridge Domain Interface with EFP, Sub-interface

The Cisco ASR 903 router does *not* support sub-interfaces as on the other Cisco platforms but a workaround can be used by pairing an Ethernet flow point (EFP) with a bridge domain interface (BDI).

- BDI is similar to SVI on the switch.
- At least one EFP must be associated with BDI to have line protocol up.
- EFP must make the packet native, one tag or two tags popping may be needed.
- Bridge-domain interface supports routing.
- Symmetric keyworld is required for push or pop operation.

  The following shows an example how to create "sub-interface" off the interface of Gigabit Ethernet 0/2/0 on Cisco ASR 903 router. The EFP will match encapsulation with vlan tag 10, pop the dot1q header in ingress direction, forward the packet into bridge-domain 10 (think BD as Vlan domain as on the switch).

```
interface Gi0/2/0
  ! create an EFP instance
   service instance 1 ethernet
     encapsulation dot1q 10
     rewrite ingress tag pop 1 symmetric     ! pop is required to make packet native
     bridge-domain 10
 interface BDI10                             ! BDI 10 is associated with bridge
   ip address 10.10.10.1 255.255.255.0       ! domain 10
```

The above example is equivalent to the sub-interface on the other platforms. Note that **Vlan** command is not required for create vlan to make BDI active on the Cisco ASR 903 router.

```
interface gigabitEthernet0/0/0/1.10
    encapsulation dot1q 10
    ip address 10.10.10.1 255.255.255.0
```

Ethernet Virtual Circuit using Regular EFPs

Here is a design example to provide end-to-end connectivity between CEs by using Ethernet bridging. For illustration purpose, two subnets of A and B, are created to demonstrate connectivity. On CE1 router, vlan tag 10 and 11 are used; vlan tag 15 and 16 are used on CE2 router.

Figure 1. Ethernet Virtual Circuit with Regular EFPs

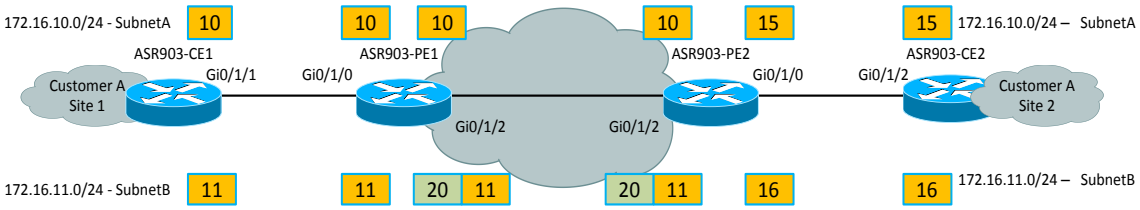

ASR903-CE1 Configuration

```
bridge-domain 10
 member GigabitEthernet0/1/1 service-instance 10
!
bridge-domain 11
 member GigabitEthernet0/1/1 service-instance 11
!
interface GigabitEthernet0/1/1
 no ip address
 negotiation auto
 service instance 10 ethernet             ! create service instance
  encapsulation dot1q 10
  rewrite ingress tag pop 1 symmetric     ! to use BDI, pop is required.
 !                                        ! otherwise rewrite may not be needed.
 service instance 11 ethernet
  encapsulation dot1q 11
  rewrite ingress tag pop 1 symmetric
!
interface BDI10                           ! subnet A
 ip address 172.16.10.1 255.255.255.0
!
interface BDI11                                 ! subnet B
 ip address 172.16.11.1 255.255.255.0
```

ASR903-PE1 Configuration

```
bridge-domain 10                                ! bridge-domain 10
 member GigabitEthernet0/1/0 service-instance 10 ! vlan tag 10 for CE subnet A
 member GigabitEthernet0/1/2 service-instance 10
!
bridge-domain 11
 member GigabitEthernet0/1/0 service-instance 11 ! vlan tag 11 for CE1 subnet B
 member GigabitEthernet0/1/2 service-instance 11
!
interface GigabitEthernet0/1/0
 description interface to CE1
 no ip address
 negotiation auto
 service instance 10 ethernet
```

```
  encapsulation dot1q 10
 !
 service instance 11 ethernet
  encapsulation dot1q 11
!
interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance 10 ethernet
  encapsulation dot1q 10                          ! Simply bridging the traffic
 !
 service instance 11 ethernet
  encapsulation dot1q 20 second-dot1q 11            ! add S-tag 20 on top of 11
  rewrite ingress tag pop 1 symmetric
```

ASR903-PE2 Configuration

```
bridge-domain 10
 member GigabitEthernet0/1/0 service-instance 10    ! vlan tag 15 facing CE2
 member GigabitEthernet0/1/2 service-instance 10    ! vlan tag 10 with PE
!
bridge-domain 11
 member GigabitEthernet0/1/0 service-instance 11    ! vlan tag 16 facing CE2
 member GigabitEthernet0/1/2 service-instance 11    ! vlan tag 20 & 11 with PE
!
interface GigabitEthernet0/1/0
 ! description interface to CE2
 no ip address
 negotiation auto
 service instance 10 ethernet
  encapsulation dot1q 15                            ! vlan tag 15 for CE subnet A
  rewrite ingress tag pop 1 symmetric
 !
 service instance 11 ethernet
  encapsulation dot1q 16                            ! vlan tag 16 for CE subnet B
  rewrite ingress tag pop 1 symmetric

interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance 10 ethernet
  encapsulation dot1q 10
  rewrite ingress tag pop 1 symmetric
 !
 service instance 11 ethernet
  encapsulation dot1q 20 second-dot1q 11     ! vlan tag 20 & 11
  rewrite ingress tag pop 2 symmetric
```

ASR903-CE2 Configuration

```
interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance 10 ethernet
  encapsulation dot1q 15
  rewrite ingress tag pop 1 symmetric
  bridge-domain 10                          ! use the legacy command
 !
 service instance 11 ethernet
  encapsulation dot1q 16
  rewrite ingress tag pop 1 symmetric
  bridge-domain 11                          ! use the legacy command
!
interface BDI10
 ip address 172.16.10.2 255.255.255.0
!
interface BDI11
 ip address 172.16.11.2 255.255.255.0
```

Verifying Connectivity

```
CE1# ping 172.16.10.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/2 ms
CE1# ping 172.16.11.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms

Router# show bridge-domain 10
Bridge-domain 10 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
Maximum address limit: 256000
    GigabitEthernet0/1/0 service instance 10
    GigabitEthernet0/1/2 service instance 10

  Nile Mac Address Entries
      BD    mac addr        type     ports
   ------------------------------------------------------------------------------
   10    c8f9.f98d.103f  DYNAMIC  Gi0/1/2.Efp10
   10    c8f9.f98d.503f  DYNAMIC  Gi0/1/0.Efp10

Router# show bridge-domain 11
Bridge-domain 11 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
```

```
Maximum address limit: 256000
    GigabitEthernet0/1/0 service instance 11
    GigabitEthernet0/1/2 service instance 11
     Nile Mac Address Entries
     BD    mac addr        type     ports
   ------------------------------------------------------------------------------
    11    c8f9.f98d.103f  DYNAMIC  Gi0/1/2.Efp11
    11    c8f9.f98d.503f  DYNAMIC  Gi0/1/0.Efp11
```

Ethernet Virtual Circuit using Trunk EFP

Trunk EFP can be used to group two or more services instances into one, which can greatly reduce the configuration burden for large network deployment. Unlike the previous example, only one service instances is created to bridging the traffic between PEs. In trunk EFP configuration, the second S-Tag will be imposed to the frame on the trunk link.

Figure 2. Ethernet Virtual Circuit with Trunk EFP

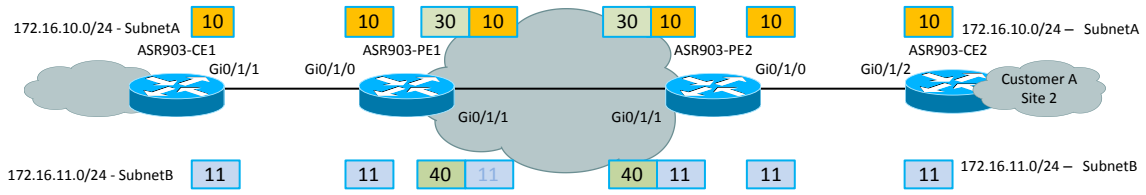

ASR903-CE1 Configuration

```
bridge-domain 10
 member GigabitEthernet0/1/1 service-instance 10
!
bridge-domain 11
 member GigabitEthernet0/1/1 service-instance 11
!
interface GigabitEthernet0/1/1
 no ip address
 negotiation auto
 service instance 10 ethernet                    ! create service instances
  encapsulation dot1q 10
  rewrite ingress tag pop 1 symmetric          ! to use BDI, pop is required.
 !                                             ! otherwise rewrite may not be needed.
 service instance 11 ethernet
  encapsulation dot1q 11
  rewrite ingress tag pop 1 symmetric
!
```

```
interface BDI10                                ! subnet A
 ip address 172.16.10.1 255.255.255.0
!
interface BDI11
 ip address 172.16.11.1 255.255.255.0          ! subnet B
```

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 service instance 10 ethernet
  encapsulation dot1q 10
  bridge-domain 30                      !  bridge-domain will be S-Tag on trunk port
 !
 service instance 11 ethernet
  encapsulation dot1q 11
  bridge-domain 40                      !  bridge-domain will be S-Tag on trunk port

interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance trunk 1 ethernet
  encapsulation dot1q 30-40               ! in this case, we only need vlan 30 and 40
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation        ! outer vlan tag will be used as bridge-
domain
```

ASR903-PE2 Configuration

```
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 service instance 10 ethernet
  encapsulation dot1q 10
  bridge-domain 30                      !  bridge-domain will be S-Tag on trunk port
 !
 service instance 11 ethernet
  encapsulation dot1q 11
  bridge-domain 40                      !  bridge-domain will be S-Tag on trunk port
interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance trunk 1 ethernet
  encapsulation dot1q 30-40               ! in this case, we only need vlan 30 and 40
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation        ! outer vlan tag is used as bridge-domain
```

ASR903-CE2 Configuration

```
interface GigabitEthernet0/1/2
```

```
 no ip address
 negotiation auto
 service instance 10 ethernet
  encapsulation dot1q 10
  rewrite ingress tag pop 1 symmetric
  bridge-domain 10
 !
 service instance 11 ethernet
  encapsulation dot1q 11
  rewrite ingress tag pop 1 symmetric
  bridge-domain 11
!
interface BDI10
 ip address 172.16.10.2 255.255.255.0
!
interface BDI11
 ip address 172.16.11.2 255.255.255.0
```

Verifying the Configuration

```
Router#  ping 172.16.10.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms

Router#  ping 172.16.11.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms


Router# show bridge-domain 30
Bridge-domain 30 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
Maximum address limit: 256000
    GigabitEthernet0/1/0 service instance 10
    GigabitEthernet0/1/2 service instance 1

   Nile Mac Address Entries

   BD    mac addr        type     ports
   ------------------------------------------------------------------------------
   30    c8f9.f98d.103f  DYNAMIC  Gi0/1/2.tefp1
   30    c8f9.f98d.503f  DYNAMIC  Gi0/1/0.Efp10
```

Layer 2 Protocol Peering

To peer with a neighbor on a port that has an EFP service instance configured, you need to configure **l2protocol peer** under the service instance.

- Layer 2 protocols peering enables protocol such as CDP, UDLD, LLDP, MSTP, LACP, and DTP on a port which has EFP configured to work with neighbor.
- CDP, MSTP, LLDP, ELMI and LACP are currently supported on ASR903.
- Layer 2 protocols use untagged frames.

Choose a bridge domain number not being used by other EFPs to have CPU process protocol packets.

```
interface GigabitEthernet port number
  service instance  instance id  ethernet
    encapsulation untagged                        !  matches untagged frame
    bridge-domain {bd number}
    l2protocol peer  [ l2protocol options ]
```

The following example enables LACP peering.

```
interface port-channel1
  service instance 1 ethernet
    encapsulation untagged
    bridge-domain  10
    l2protocol peer lacp
```

Important: The **lacp** keyword is optional in the last command, without it, it enables all the layer 2 protocols.

Spanning-Tree Configuration

MSTP is supported under EFP bridge domains. MSTP is recommended to be enabled to prevent loop, and it is disabled by default. RSTP is supported for MST instance 0 only.

- Per-vlan STP is **not** supported on ASR903.
- All incoming VLANs (outer-most or single) must belong to the same MST instance or loops could occur.
- Backup EFPs must be mapped to the same MST instance as active EFPs.
- **Spanning-tree mode mst** is to enable MSTP for EFPs and all the EFPs are default to instance 0 unless specified otherwise.
- L2 protocol peering is needed for devices to use MSTP with neighbor device.

Now another trunk link is added between two PE routers, and spanning-tree should be enabled to prevent loop.

Figure 3. Spanning Tree

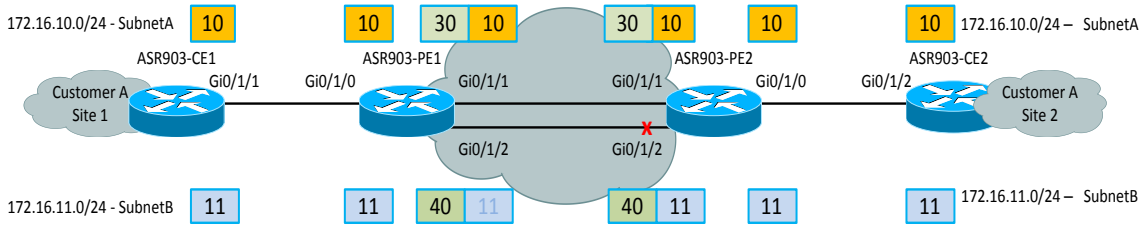

ASR903-PE1 and ASR903-PE2 Configuration

```
spanning-tree mode mst                          ! enable MST on the ASR903
interface GigabitEthernet0/1/1
 no ip address
 negotiation auto
 service instance trunk 1 ethernet
  encapsulation dot1q 30-40
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
 service instance 2 ethernet
  encapsulation untagged                           ! enable l2protocol peering
  l2protocol peer
  bridge-domain 1
interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance trunk 1 ethernet
  encapsulation dot1q 30-40
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
 service instance 2 ethernet
  encapsulation untagged
  l2protocol peer                                ! enable l2protocol peering
  bridge-domain 1
```

Verify Spanning-tree Configuration

```
ASR903-PE1# show span vlan 30
MST0
  Spanning tree enabled protocol mstp
  Root ID     Priority    32768
              Address     7010.5c51.8fbb
```

```
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID  Priority    32768  (priority 32768 sys-id-ext 0)
             Address     7010.5c51.8fbb
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi0/1/1             Desg FWD 20000     128.7    P2p
Gi0/1/2             Desg FWD 20000     128.8    P2p

ASR903-PE2# show span vlan 30

MST0
  Spanning tree enabled protocol mstp
  Root ID    Priority    32768
             Address     7010.5c51.8fbb
             Cost        0
             Port        16 (GigabitEthernet0/1/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32768  (priority 32768 sys-id-ext 0)
             Address     7010.5c51.a4bb
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi0/1/1             Root FWD 20000     128.16   P2p
Gi0/1/2             Altn BLK 20000     128.17   P2p
```

Ethernet Port-channel

Cisco ASR 903 router supports layer 2 and layer 3 port-channels, and only LACP is supported for link aggregation. Here are the steps to create layer 3 port-channel.

```
interface Port-channel  channel number
 ip address x.x.x.x y.y.y.y
interface GigabitEthernet0/0/1
 no ip address
 negotiation auto
 channel-group channel-number mode active
```

For layer 2 port-channel, if there are EFPs or EFP trunk configured under port-channel, l2 protocol peering is needed under port-channel to enable LACP.

```
interface port-channel 1
  service instance 1 ethernet
   encapsulation untagged
   bridge-domain id
   l2protocol peer lacp
```

Three modes are supported on the Cisco ASR 903 router.

- LACP—**channel-group 1 mode { *active* | *passive*}**
- On—**channel-group 1**

  On mode will force member link into port-channel whether other end joins or not. On mode is not recommended as traffic can be black holed if the peering end does not join port-channel. Here is the command to force interface into port-channel.

```
interface GigabitEthernet0/0/1
 no ip address
 negotiation auto
 channel-group 5                ! using on mode
```

Load-balancing can be achieved based on MAC address or IP address of source and destination packets. Port-channel load balancing can be configured globally for all the port-channels; individual port channel load balancing command is not supported. Here are the commands to configure port-channel load balancing.

```
Configure  terminal
 port-channel load-balance-hash-algo  {option}
  dst-ip                   Destination IP
  dst-mac                Destination MAC
  src-dst-ip               Source XOR Destination IP Addr
  src-dst-mac            Source XOR Destination MAC
  src-ip                   Source IP
  src-mac                Source MAC
```

In the example below, gig0/1/1 and gig0/1/2 is member of port-channel1.

Figure 4. Ethernet Port-Channel

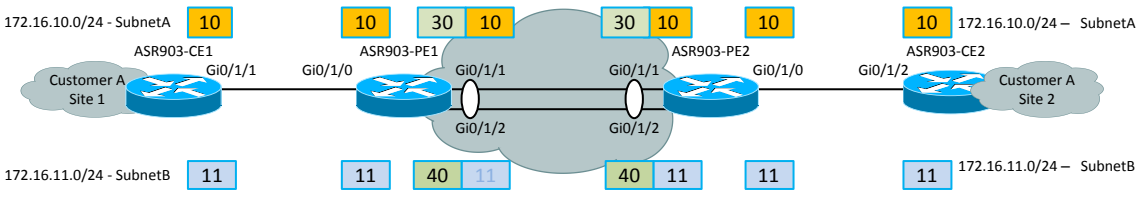

ASR903-PE1 and ASR903-PE2 Configuration

```
interface Port-channel1
 no ip address
 no negotiation auto
 service instance 1 ethernet
  encapsulation untagged
  l2protocol peer      ! l2protocol peering is needed for lacp if evc is configured
  bridge-domain 1
```

```
 !
 service instance trunk 2 ethernet
  encapsulation dot1q 30-40
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
!
interface GigabitEthernet0/1/1
 no ip address
 negotiation auto
 channel-group 1 mode active
!
interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 channel-group 1 mode active

Router# show interface port-channel 1
Port-channel1 is up, line protocol is up
  Hardware is GEChannel, address is 7010.5c51.8fc0 (bia 7010.5c51.8fc0)
  MTU 1500 bytes, BW 2000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
    No. of active members in this channel: 2
        Member 0 : GigabitEthernet0/1/1 , Full-duplex, 1000Mb/s
        Member 1 : GigabitEthernet0/1/2 , Full-duplex, 1000Mb/s
    No. of PF_JUMBO supported members in this channel : 2
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/750/0/0 (size/max/drops/flushes); Total output drops: 0


Queueing strategy: fifo
  Output queue: 0/80 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     216 packets input, 26628 bytes, 0 no buffer
     Received 5 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 206 multicast, 0 pause input
     712 packets output, 86772 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapped out

Router# show bridge-domain 40
Bridge-domain 40 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
```

```
Maximum address limit: 256000
    GigabitEthernet0/1/0 service instance 11
    Port-channel1 service instance 2

   Nile Mac Address Entries
      BD    mac addr        type     ports
   ------------------------------------------------------------------------------
   40    c8f9.f98d.103f  DYNAMIC  Po0/0/1.tefp2
   40    c8f9.f98d.503f  DYNAMIC  Gi0/1/0.Efp11
```

G.8032

The G.8032 Ethernet Ring Protection Switching implements protection switching mechanisms for Ethernet layer ring topologies. This feature uses the G.8032 Ethernet Ring Protection (ERP) protocol, defined in ITU-T G.8032, to provide protection for Ethernet traffic in a ring topology, while ensuring that no loops are within the ring at the Ethernet layer. The loops are prevented by blocking traffic on either a predetermined link or a failed link.

Figure 5. G.8032

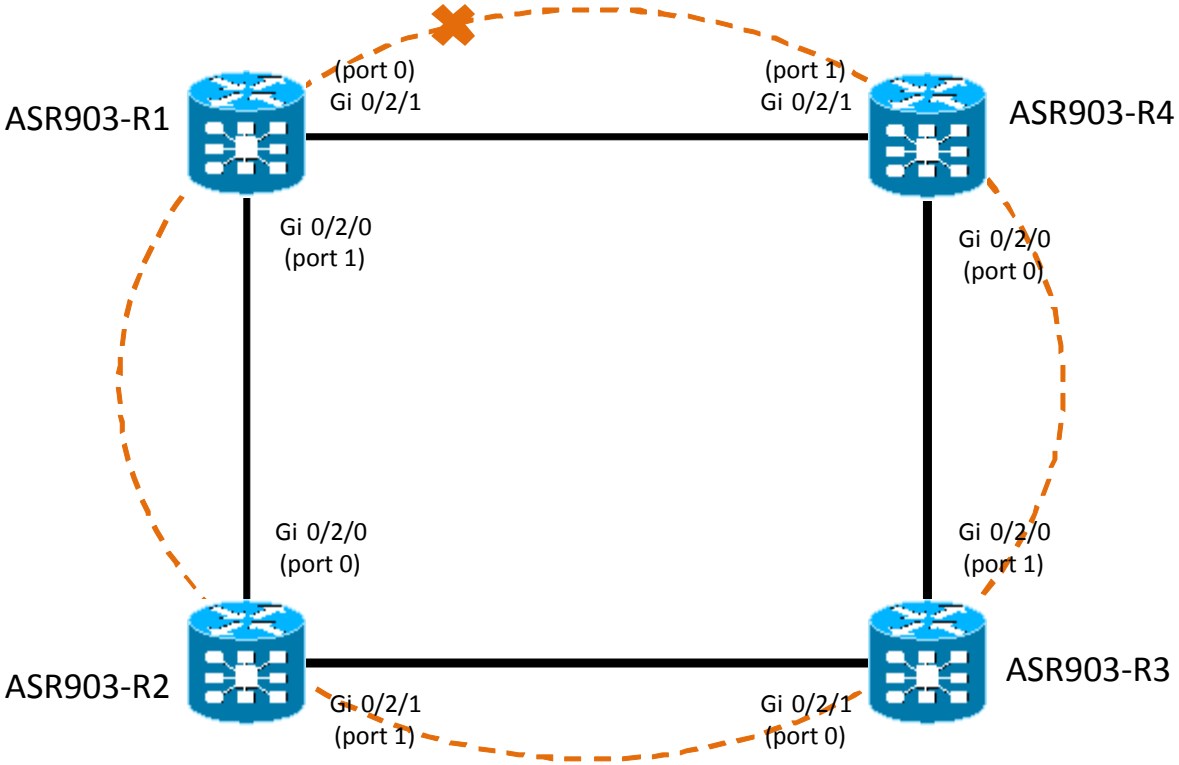

For G.8032, CFM can be configured under either an EFP or main Ethernet interface for failure detection. Following is a configuration example for CFM under EFP scenario.

ASR903-R1 Configuration:

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service evc evc_name vlan 10 direction down
  continuity-check
  continuity-check interval 3.3ms
!
ethernet evc evc_name
!
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 2
   continuity-check static rmep
   rmep mpid 1
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 1
   continuity-check static rmep
   rmep mpid 2
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
!
!
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/1
 port1 interface GigabitEthernet0/2/0
 instance 1
```

```
 profile g8032_profile
 rpl port0 owner
 inclusion-list vlan-ids 10,1000-2999
 aps-channel
  port0 service instance 1
  port1 service instance 1
 !
!
```

ASR903-R2 Configuration

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service evc evc_name vlan 10 direction down
  continuity-check
  continuity-check interval 3.3ms
!
!
ethernet evc evc_name

!

interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 1
   continuity-check static rmep
   rmep mpid 2
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
!
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 2
   continuity-check static rmep
   rmep mpid 1
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
```

```
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
!
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/0
 port1 interface GigabitEthernet0/2/1
 instance 1
  profile g8032_profile
  inclusion-list vlan-ids 10,1000-2999
  aps-channel
   port0 service instance 1
   port1 service instance 1
  !
```

ASR903-R3 Configuration

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service evc evc_name vlan 10 direction down
  continuity-check
  continuity-check interval 3.3ms
!

ethernet evc evc_name
!
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 2
   continuity-check static rmep
   rmep mpid 1
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
!
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 1
   continuity-check static rmep
   rmep mpid 2
```

```
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
!
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/1
 port1 interface GigabitEthernet0/2/0
 instance 1
  profile g8032_profile
  inclusion-list vlan-ids 10,1000-2999
  aps-channel
   port0 service instance 1
   port1 service instance 1
  !
```

ASR903-R4 Configuration

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service evc evc_name vlan 10 direction down
  continuity-check
  continuity-check interval 3.3ms
!
ethernet evc evc_name
!
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 1
   continuity-check static rmep
   rmep mpid 2
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
!
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
```

```
 service instance 1 ethernet evc_name
  encapsulation dot1q 10
  bridge-domain 10
  cfm mep domain g8032_domain mpid 2
   continuity-check static rmep
   rmep mpid 1
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/0
 port1 interface GigabitEthernet0/2/1
 instance 1
  profile g8032_profile
  rpl port1 neighbor
  inclusion-list vlan-ids 10,1000-2999
  aps-channel
   port0 service instance 1
   port1 service instance 1
```

Configuration Example: CFM under Ethernet Interface

ASR903-R1 Configuration

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service port
  continuity-check
  continuity-check interval 3.3ms
!
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 2 service g8032_service
  continuity-check static rmep
  rmep mpid 1
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 10
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
interface GigabitEthernet0/2/1
```

```
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 1 service g8032_service
  continuity-check static rmep
  rmep mpid 2
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 10
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
!
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/1
 port1 interface GigabitEthernet0/2/0
 instance 1
  profile g8032_profile
  rpl port0 owner
  inclusion-list vlan-ids 10,1000-2999
  aps-channel
   port0 service instance 1
   port1 service instance 1
  !
 !
```

ASR903-R2 Configuration

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service port
  continuity-check
  continuity-check interval 3.3ms
!

!
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 1 service g8032_service
  continuity-check static rmep
  rmep mpid 2
 service instance 1 ethernet
  encapsulation dot1q 10
```

```
  bridge-domain 10
 !
 service instance trunk 1000 ethernet

encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
!
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 2 service g8032_service
  continuity-check static rmep
  rmep mpid 1
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 10
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
!
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/0
 port1 interface GigabitEthernet0/2/1
 instance 1
  profile g8032_profile
  inclusion-list vlan-ids 10,1000-2999
  aps-channel
   port0 service instance 1
   port1 service instance 1
  !
```

ASR903-R3 Configuration

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service port
  continuity-check
  continuity-check interval 3.3ms
!
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 2 service g8032_service
  continuity-check static rmep
  rmep mpid 1
 service instance 1 ethernet
  encapsulation dot1q 10
```

```
 bridge-domain 10
!
service instance trunk 1000 ethernet
 encapsulation dot1q 1000-2999
 rewrite ingress tag pop 1 symmetric
 bridge-domain from-encapsulation
!
!
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 1 service g8032_service
  continuity-check static rmep
  rmep mpid 2
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 10
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/1
 port1 interface GigabitEthernet0/2/0
 instance 1
  profile g8032_profile
  inclusion-list vlan-ids 10,1000-2999
  aps-channel
   port0 service instance 1
   port1 service instance 1
  !
```

ASR903-R4 Configuration

```
ethernet cfm ieee
ethernet cfm domain g8032_domain level 1
 service g8032_service port
  continuity-check
  continuity-check interval 3.3ms
!
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 1 service g8032_service
```

```
  continuity-check static rmep
  rmep mpid 2
 service instance 1 ethernet

 encapsulation dot1q 10
  bridge-domain 10
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
 ethernet cfm mep domain g8032_domain mpid 2 service g8032_service
  continuity-check static rmep
  rmep mpid 1
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 10
 !
 service instance trunk 1000 ethernet
  encapsulation dot1q 1000-2999
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
ethernet ring g8032 profile g8032_profile
 timer wtr 1
!
ethernet ring g8032 g8032_ring
 port0 interface GigabitEthernet0/2/0
 port1 interface GigabitEthernet0/2/1
 instance 1
  profile g8032_profile
  rpl port1 neighbor
  inclusion-list vlan-ids 10,1000-2999
  aps-channel
   port0 service instance 1
   port1 service instance 1
  !
```

Verifying G.8032

```
ASR903-R1# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch
```

```
RingName                          Inst NodeType  NodeState     Port0    Port1
-------------------------------------------------------------------------------
g8032_ring                        1    Owner     Idle          R,B

ASR903-R1# show ethernet ring g8032 configuration

Ethernet ring g8032_ring
 Port0: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
 Port1: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
 Exclusion-list VLAN IDs:
 Open-ring: no
 Instance 1
  Description:
  Profile:     g8032_profile
  RPL:         port0 RPL Owner
  Inclusion-list VLAN IDs: 10,1000-2999
  APS channel
   Level: 7
   Port0: Service Instance 1
   Port1: Service Instance 1
  State: configuration resolved


ASR903-R1# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 2/2
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list: 10,1000-2999
         Unblock vlan list:
         REQ/ACK: 1/1
         Instance 1 is in Blocked state


ASR903-R1# show ethernet ring g8032 profile
Ethernet ring profile name: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode


ASR903-R1# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is RPL Owner node in Idle State
```

```
 Port0: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
  Status: RPL, blocked
  Remote R-APS NodeId: 0000.0000.0000, BPR: 0
 Port1: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
  Remote R-APS NodeId: 0000.0000.0000, BPR: 0
 APS Level: 7
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode

ASR903-R2# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch

RingName                        Inst NodeType  NodeState     Port0    Port1
--------------------------------------------------------------------------------
g8032_ring                      1    Normal    Idle



ASR903-R2# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 6/6
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 5/5
         Instance 1 is in Unblocked state

ASR903-R2# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is Normal Node node in Idle State
 Port0: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
  Remote R-APS NodeId: d0c2.8216.29bf, BPR: 0
 Port1: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
  Status: Non-RPL
  Remote R-APS NodeId: 0000.0000.0000, BPR: 0
```

```
 APS Level: 7
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode

ASR903-R3# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch

RingName                        Inst NodeType  NodeState     Port0    Port1
-------------------------------------------------------------------------------
g8032_ring                      1    Normal    Idle

ASR903-R3# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 2/2
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 8/8
         Instance 1 is in Unblocked state

ASR903-R3# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is Normal Node node in Idle State
 Port0: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
  Status: Non-RPL
  Remote R-APS NodeId: d0c2.8216.29bf, BPR: 0
 Port1: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
  Remote R-APS NodeId: 0000.0000.0000, BPR: 0
 APS Level: 7
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode
```

```
ASR903-R4# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch

RingName                        Inst NodeType  NodeState     Port0    Port1
---------------------------------------------------------------------------
g8032_ring                      1    Neighbor  Idle                   R,B

ASR903-R4# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 2/2
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list: 10,1000-2999
         Unblock vlan list:
         REQ/ACK: 3/3
         Instance 1 is in Blocked state

ASR903-R4# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is RPL Neighbor node in Idle State
 Port0: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
  Remote R-APS NodeId: d0c2.8216.29bf, BPR: 0
 Port1: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
  Status: RPL, blocked
  Remote R-APS NodeId: d0c2.8216.29bf, BPR: 0
 APS Level: 7
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode
```

Protection Switching

```
ASR903-R3# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ASR903-R3(config)# interface gigabitEtherenet0/2/1
ASR903-R3(config-if)# shut
ASR903-R3(config-if)#
*Jun 12 19:25:20.949: %ERP_G8032-6-STATE_CHANGED: Ethernet ring g8032_ring instance 1
changed state to Protection
*Jun 12 19:25:22.861: %LINK-5-CHANGED: Interface GigabitEthernet0/2/1, changed state
to administratively down
*Jun 12 19:25:22.882: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/2/1, changed state to down
ASR903-R3(config-if)# end
ASR903-R1#
*Jun 12 19:25:45.098: %ERP_G8032-6-STATE_CHANGED: Ethernet ring g8032_ring instance 1
changed state to Protection
ASR903-R1#

ASR903-R1# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch
RingName                        Inst NodeType  NodeState     Port0    Port1
--------------------------------------------------------------------------------
g8032_ring                      1    Owner     Protection    R

ASR903-R1# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 2/2
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 2/2
         Instance 1 is in Unblocked state

ASR903-R1# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is RPL Owner node in Protection State
 Port0: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
```

```
  Status: RPL
  Remote R-APS NodeId: d0c2.8216.1fbf, BPR: 0
 Port1: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
  Remote R-APS NodeId: d0c2.8216.20bf, BPR: 1
 APS Level: 7
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode

ASR903-R2#
*Jun 12 19:25:13.000: %ERP_G8032-6-STATE_CHANGED: Ethernet ring g8032_ring instance 1
changed state to Protection
*Jun 12 19:25:14.956: %LINK-3-UPDOWN: Interface GigabitEthernet0/2/1, changed state
to down
*Jun 12 19:25:14.980: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/2/1, changed state to down

ASR903-R2# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch

RingName                        Inst NodeType  NodeState     Port0    Port1
-------------------------------------------------------------------------------
g8032_ring                      1    Normal    Protection             F,B


ASR903-R2# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 6/6
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list:
         REQ/ACK: 6/6
         Instance 1 is in Blocked state

ASR903-R2# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is Normal Node node in Protection State
 Port0: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
```

```
  Remote R-APS NodeId: d0c2.8216.1fbf, BPR: 0
 Port1: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
  Status: Non-RPL, faulty, blocked
  Remote R-APS NodeId: 0000.0000.0000, BPR: 0
 APS Level: 7
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode

ASR903-R3# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch

RingName                        Inst NodeType  NodeState     Port0    Port1
--------------------------------------------------------------------------------
g8032_ring                      1    Normal    Protection    F,B

ASR903-R3# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 2/2
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list: 10,1000-2999
         Unblock vlan list:
         REQ/ACK: 9/9
         Instance 1 is in Blocked state

ASR903-R3# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is Normal Node node in Protection State
 Port0: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
  Status: Non-RPL, faulty, blocked
  Remote R-APS NodeId: 0000.0000.0000, BPR: 0
 Port1: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
  Remote R-APS NodeId: d0c2.8216.20bf, BPR: 1
 APS Level: 7
```

```
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode

ASR903-R4#
*Jun 12 19:19:54.816: %ERP_G8032-6-STATE_CHANGED: Ethernet ring g8032_ring instance 1
changed state to Protection

ASR903-R4# show ethernet ring g8032 brief
R: Interface is the RPL-link
F: Interface is faulty
B: Interface is blocked
FS: Local forced switch
MS: Local manual switch

RingName                        Inst NodeType  NodeState     Port0    Port1
-------------------------------------------------------------------------------
g8032_ring                      1    Neighbor  Protection             R


ASR903-R4# show ethernet ring g8032 port status
 Port: GigabitEthernet0/2/0
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 2/2
         Instance 1 is in Unblocked state

 Port: GigabitEthernet0/2/1
 Ring: g8032_ring
         Block vlan list:
         Unblock vlan list: 10,1000-2999
         REQ/ACK: 4/4
         Instance 1 is in Unblocked state

ASR903-R4# show ethernet ring g8032 status
Ethernet ring g8032_ring instance 1 is RPL Neighbor node in Protection State
 Port0: GigabitEthernet0/2/0 (Monitor: GigabitEthernet0/2/0)
  APS-Channel: GigabitEthernet0/2/0
  Status: Non-RPL
  Remote R-APS NodeId: d0c2.8216.1fbf, BPR: 0
 Port1: GigabitEthernet0/2/1 (Monitor: GigabitEthernet0/2/1)
  APS-Channel: GigabitEthernet0/2/1
  Status: RPL
  Remote R-APS NodeId: d0c2.8216.20bf, BPR: 1
 APS Level: 7
 Profile: g8032_profile
  WTR interval: 1 minutes
  Guard interval: 500 milliseconds
  HoldOffTimer: 0 seconds
  Revertive mode
```

REP

The Resilient Ethernet Protocol (REP) is a Cisco proprietary protocol that provides an alternative to the Spanning Tree Protocol (STP). REP provides a way to control network loops, handle link failures, and improve convergence time. It controls a group of ports connected in a segment, ensures that the segment does not create any bridging loops, and responds to link failures within the segment. With REP at least one port is always blocked in any given segment, that is, the alternate port. The blocked port helps ensure that the traffic within the segment is loop-free by requiring traffic flow to exit only one of the edge ports, and not both. REP provides a basis for constructing complex networks and supports VLAN load balancing.

Figure 6. REP

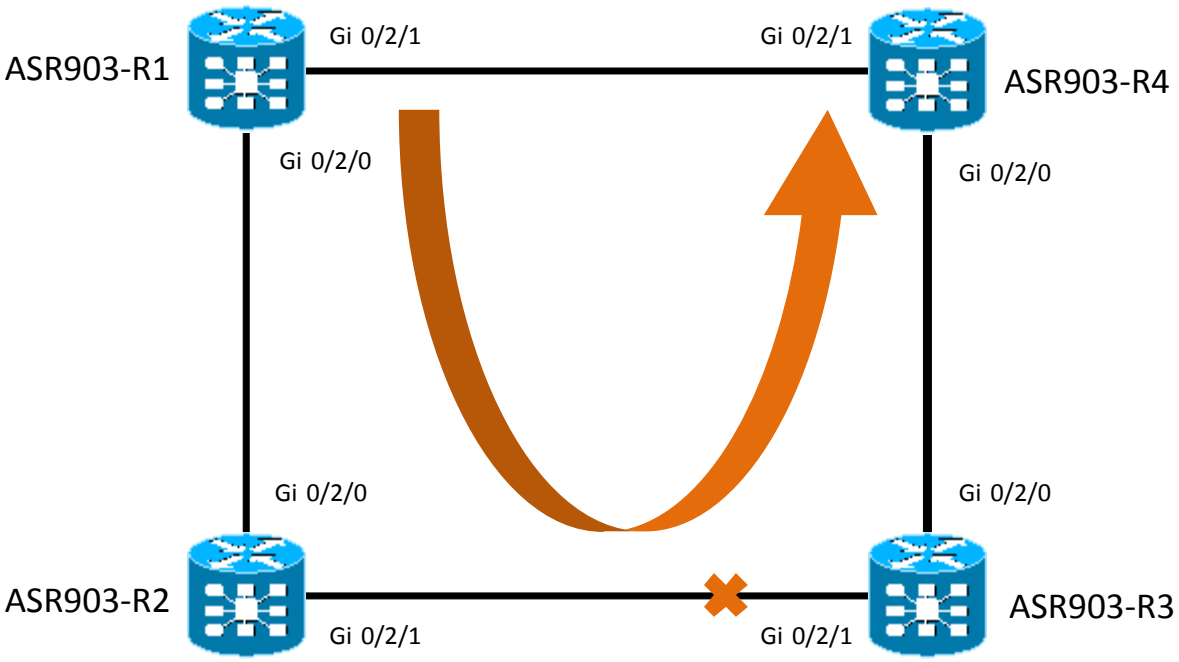

ASR903-R1 Configuration

```
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 rep segment 1 edge primary
 rep stcn stp
 rep block port preferred vlan 1-4094
 service instance trunk 1 ethernet
  encapsulation dot1q 1-4094
  rewrite ingress tag pop 1 symmetric
```

```
 bridge-domain from-encapsulation
!
```

ASR903-R2 Configuration

```
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 rep segment 1
 service instance trunk 1 ethernet
  encapsulation dot1q 1-4094
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !

interface GigabitEthernet0/2/1
 no ip address

 negotiation auto
 rep segment 1
 service instance trunk 1 ethernet
  encapsulation dot1q 1-4094
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
```

ASR903-R3 Configuration

```
interface GigabitEthernet0/2/1
 no ip address
 negotiation auto
 rep segment 1 preferred
 service instance trunk 1 ethernet
  encapsulation dot1q 1-4094
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 rep segment 1
 service instance trunk 1 ethernet
  encapsulation dot1q 1-4094
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
```

ASR903-R4 Configuration

```
interface GigabitEthernet0/2/0
 no ip address
 negotiation auto
 rep segment 1 edge
 rep stcn stp
 service instance trunk 1 ethernet
  encapsulation dot1q 1-4094
  rewrite ingress tag pop 1 symmetric
  bridge-domain from-encapsulation
 !
```

Verifying REP

```
ASR903-R1# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Alt
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Open
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open

ASR903-R2# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Alt
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Open
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open

ASR903-R3# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Alt
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Open
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open
```

```
ASR903-R4# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Alt
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Open
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open
```

Link Failover

```
ASR903-R3# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ASR903-R3(config)#int gi0/2/1
ASR903-R3(config-if)# shutdown

*Jun 10 17:11:36.053: %REP-4-LINKSTATUS: GigabitEthernet0/2/1 (segment 1) is non-
operational due to port down
*Jun 10 17:11:38.049: %LINK-5-CHANGED: Interface GigabitEthernet0/2/1, changed state
to administratively down
*Jun 10 17:11:38.069: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/2/1, changed state to down


ASR903-R3# show rep topology
REP Segment 1
Warning: REP detects a segment failure, topology may be incomplete

BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R3           Gi0/2/1         Fail
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open


ASR903-R1# show rep topology
REP Segment 1
Warning: REP detects a segment failure, topology may be incomplete

BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Sec  Open
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Fail
```

```
ASR903-R2# show rep topology
REP Segment 1
Warning: REP detects a segment failure, topology may be incomplete

BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R2           Gi0/2/1         Fail
ASR903-R2           Gi0/2/0         Open
ASR903-R1           Gi0/2/0    Sec  Open


ASR903-R3# show rep topology
REP Segment 1
Warning: REP detects a segment failure, topology may be incomplete

BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R3           Gi0/2/1         Fail
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open

ASR903-R4# show rep topology
REP Segment 1
Warning: REP detects a segment failure, topology may be incomplete

BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R4           Gi0/2/0    Sec  Open
ASR903-R3           Gi0/2/0         Open
ASR903-R3           Gi0/2/1         Fail

Link Recovery
ASR903-R3(config-if)# no shutdown
ASR903-R3(config-if)# end
ASR903-R3#
*Jun 10 17:14:12.329: %SYS-5-CONFIG_I: Configured from console by console
*Jun 10 17:14:13.376: %LINK-3-UPDOWN: Interface GigabitEthernet0/2/1, changed state
to up
*Jun 10 17:14:13.396: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/2/1, changed state to up


ASR903-R1# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Alt
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Open
```

```
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open
ASR903-R2#
*Jun 10 17:14:05.912: %LINK-3-UPDOWN: Interface GigabitEthernet0/2/1, changed state
to up
*Jun 10 17:14:05.934: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/2/1, changed state to up
*Jun 10 17:14:13.163: %REP-4-LINKSTATUS: GigabitEthernet0/2/1 (segment 1) is
operational
ASR903-R2#

ASR903-R2# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Alt
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Open
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open

ASR903-R3# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Open
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Alt
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open

ASR903-R4# show rep topology
REP Segment 1
BridgeName       PortName   Edge Role
---------------- ---------- ---- ----
ASR903-R1           Gi0/2/0    Pri  Alt
ASR903-R2           Gi0/2/0         Open
ASR903-R2           Gi0/2/1         Open
ASR903-R3           Gi0/2/1         Open
ASR903-R3           Gi0/2/0         Open
ASR903-R4           Gi0/2/0    Sec  Open
```

Chapter 3 Basic MPLS VPN Configuration

The following diagram will be used to demonstrate VPN configuration. Two customers A and B from three sites are connected to PEs. On site 2, customer A and B are connected to the same router CE2.

Figure 7. MPLS

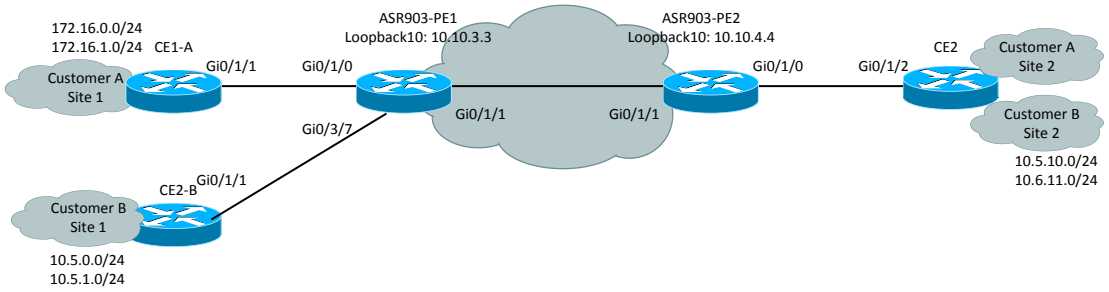

Configuration of OSPF and BFD

OSPF is chosen as IGP protocol and BFD is used to detect link failure for fast convergence. ASR 903 supports both software and hardware based BFD sessions. When a BFD session is created, depending on the hardware resources and nature of BFD session, it can be offloaded to hardware. ASR 903 supports maximum 511 no echo sessions, 255 echo sessions or combination of both in hardware, and 64 x200 ms BFD sessions in software. LDP is also enabled on both PE routers.

ASR903-PE1 Configuration

```
interface Loopback10
 ip address 10.10.3.3 255.255.255.255
interface GigabitEthernet0/1/1
 ip address 10.10.34.3 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
 mpls ip
 bfd interval 50 min_rx 50 multiplier 3
 no bfd echo                          ! use no echo function
mpls ldp router-id loopback10 force
!
router ospf 100
 router-id 10.10.3.3
```

```
 network 10.10.0.0 0.0.255.255 area 0
 bfd all-interfaces                        ! enable BFD
```

ASR903-PE2 Configuration

```
interface Loopback10
 ip address 10.10.4.4 255.255.255.255
interface GigabitEthernet0/1/1
 ip address 10.10.34.4 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 ip ospf bfd                             ! enable BFD under interface
 negotiation auto
 mpls ip
 bfd interval 50 min_rx 50 multiplier 3
 no bfd echo
mpls ldp router-id loopback10 force
!
router ospf 100
 router-id 10.10.4.4
 network 10.10.0.0 0.0.255.255 area 0
!
```

Verifying the Configuration

```
ASR903-PE2# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
10.10.3.3       1   FULL/DR         00:00:02    10.10.34.3    GigabitEthernet0/1/1

ASR903-PE2# show bfd neighbors details
IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
10.10.34.3                             1/1          Up        Up        Gi0/1/1
Session state is UP and not using echo function.
Session Host: Hardware
OurAddr: 10.10.34.4
Handle: 1
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 50000, MinRxInt: 50000, Multiplier: 3
Received MinRxInt: 50000, Received Multiplier: 3
Holddown (hits): 0(0), Hello (hits): 50(0)
Rx Count: 43858
Tx Count: 43860
Elapsed time watermarks: 0 0 (last: 0)
Registered protocols: OSPF CEF
Uptime: 00:28:22
Last packet: Version: 1                - Diagnostic: 0
             State bit: Up             - Demand bit: 0
             Poll bit: 0               - Final bit: 0
             C bit: 1
```

```
Multiplier: 3               - Length: 24
My Discr.: 1                - Your Discr.: 1
Min tx interval: 50000      - Min rx interval: 50000
Min Echo interval: 0
```

In Cisco IOS XE Release 3.9S, echo function is supported. BFD timer on the Cisco ASR 903 router can be as low as 3.3ms and timer lower than 50ms can be set by using BFD template.

```
bfd-template single-hop OSPF-BFD          ! created BFD template name OSPF-BFD
 interval min-tx 10 min-rx 10 multiplier 3    ! set BFD timer to 10ms
interface GigabitEthernet0/1/1
 bfd template OSPF-BFD                    ! apply template under the interface
```

Configuration of BGP and VRF on PE Routers

Configuration BGP and VRF are the next steps in MPLS VPN deployment. BGP routing is to ensure that VPNv4 routes can be transported across service provider backbone using MP-iBGP. Virtual Routing and Forwarding(VRF) is to put different customers into separate VRF instances to provide VPN services.

Table 2. Routing Protocols used between the CEs and PEs

Site	PE-CE Routing Protocol	Route Distinguisher and Route Target
CustomerA Site 1	OSPF	100:41
CustomerB Site 1	ISIS	100:42
CustomerA Site 2	EBGP	100:41
CustomerB Site 2	OSPF	100:42

ASR903-PE1 Configuration

```
ip vrf CustomerA                              ! define vrf CustomerA
 rd 100:41
 route-target export 100:41
 route-target import 100:41
!
ip vrf CustomerB                              ! define vrf CustomerB
 rd 100:42
 route-target export 100:42
 route-target import 100:42
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.4.4 remote-as 100
 neighbor 10.10.4.4 update-source Loopback10
 !
 address-family vpnv4                        ! VPN neighbor with ASR903-PE2, M-iBGP
  neighbor 10.10.4.4 activate
  neighbor 10.10.4.4 send-community extended
```

```
 exit-address-family
```

ASR903-PE2 Configuration

```
ip vrf CustomerA
 rd 100:41
 route-target export 100:41
 route-target import 100:41
!
ip vrf CustomerB
 rd 100:42
 route-target export 100:42
 route-target import 100:42
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.3.3 remote-as 100                ! peering with ASR903-PE1
 neighbor 10.10.3.3 update-source Loopback10
 !
 address-family vpnv4                          ! transport VPN routes
  neighbor 10.10.3.3 activate
  neighbor 10.10.3.3 send-community extended
 exit-address-family
```

CE Related Configuration on PE Routers

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
 description To CustomerA Site 1
 ip vrf forwarding CustomerA               ! put the interface under vrf CustomerA
 ip address 192.168.13.3 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
interface GigabitEthernet0/1/7
 description To CustomerB Site 1
 ip vrf forwarding CustomerB               ! put the interface under vrf CustomerA
 ip address 192.168.14.1 255.255.255.248
 ip router isis CustB
 negotiation auto
!
router ospf 41 vrf CustomerA               ! use OSPF between PE and CE
 router-id 192.168.13.3
 redistribute bgp 100 subnets                ! redistribute VPNv4 routes from other
PEs
 network 192.168.13.3 0.0.0.0 area 41
!
```

```
router isis CustB                          ! use ISIS between PE and CE
 vrf CustomerB
 net 10.0001.0042.0042.00
 is-type level-2-only
 metric-style wide
 redistribute bgp 100                      ! redistribute VPNv4 routes from other PEs
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.4.4 remote-as 100
 neighbor 10.10.4.4 update-source Loopback10
 !
 address-family ipv4 vrf CustomerA           ! distribute VRF routes into M-iBGP
  redistribute ospf 41 match internal external 1 external 2
 exit-address-family
 !
 address-family ipv4 vrf CustomerB           ! distribute VRF routes into M-iBGP
  redistribute isis CustB level-2
 exit-address-family
```

ASR903-PE2 Configuration

```
! use BDI(SVI) interfaces on PE2
interface GigabitEthernet0/1/0
 service instance 1 ethernet
  encapsulation dot1q 41
  rewrite ingress tag pop 1 symmetric
 !
 service instance 2 ethernet
  encapsulation dot1q 42
  rewrite ingress tag pop 1 symmetric

bridge-domain 41
 member GigabitEthernet0/1/0 service-instance 1     ! associate BD with service
instance
!
bridge-domain 42
 member GigabitEthernet0/1/0 service-instance 2  ! associate BD with service instance
interface BDI41                               ! virtual BDI interface for CustomerA
 ip vrf forwarding CustomerA
 ip address 192.168.40.1 255.255.255.248
!
interface BDI42                               ! virtual BDI interface for CustomerB
 ip vrf forwarding CustomerB
 ip address 192.168.41.1 255.255.255.248
!
router ospf 42 vrf CustomerB
 redistribute bgp 100 subnets
 network 192.168.41.0 0.0.0.255 area 42
```

```
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.3.3 remote-as 100                ! peering with ASR903-PE1
 neighbor 10.10.3.3 update-source Loopback10
 !
 address-family ipv4 vrf CustomerA               ! eBGP between PE and CE2 for
customerA
  neighbor 192.168.40.2 remote-as 64101
  neighbor 192.168.40.2 activate
 exit-address-family
 !
 address-family ipv4 vrf CustomerB               ! distribute OSPF into M-iBGP
  redistribute ospf 42 match internal external 1 external 2
 exit-address-family
```

Alternatively, the following configuration method can be used to define VRF, and those configurations apply to 6vPE too. This is the preferred method as it is easy to make it dual stack to support IPv6 as well.

```
vrf definition CustomerA                         ! define the VRF
 rd 100:41
 !
 address-family ipv4
  route-target export 100:41
  route-target import 100:41
 exit-address-family
interface GigabitEthernet0/1/0
 description To CustomerA Site 1
 vrf forwarding CustomerA                     ! put the interface under vrf CustomerA
```

Configuration of CE Routers

CE1-A Configuration

```
interface GigabitEthernet0/1/1
 description to PE-1 interface
 ip address 192.168.13.1 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
router ospf 41                             ! use OSPF as PC-CE protocol
 network 172.16.0.0 0.0.255.255 area 41
 network 192.168.13.0 0.0.0.255 area 41
interface Loopback110                      ! routes to advertise out from CE1-A
 ip address 172.16.0.1 255.255.255.0
 ip ospf network point-to-point
interface Loopback111                      ! routes to advertise out from CE1-A
 ip address 172.16.1.1 255.255.255.0
 ip ospf network point-to-point
```

CE-2 Configuration

```
! use Bridge-domain and BDI to create two "sub-interfaces" on CE-2, one for customerA
and
! another for customerB

bridge-domain 41
 member GigabitEthernet0/1/2 service-instance 1
!
bridge-domain 42
 member GigabitEthernet0/1/2 service-instance 2
interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance 1 ethernet
  encapsulation dot1q 41
  rewrite ingress tag pop 1 symmetric
 !
 service instance 2 ethernet
  encapsulation dot1q 42
  rewrite ingress tag pop 1 symmetric

! create two BDI interfaces

interface BDI41
 ip address 192.168.40.2 255.255.255.248
!

interface BDI42
 ip address 192.168.41.2 255.255.255.248
interface Loopback90
 ip address 172.16.10.1 255.255.255.0
!

interface Loopback91
 ip address 172.16.11.1 255.255.255.0
!

interface Loopback100
 ip address 10.5.10.1 255.255.255.0
 ip ospf network point-to-point
!

interface Loopback101
 ip address 10.5.11.1 255.255.255.0
 ip ospf network point-to-point
!
```

```
! OSPF is used for customerA
router ospf 42
 network 10.5.10.0 0.0.0.255 area 42
 network 10.5.11.0 0.0.0.255 area 42
 network 192.168.41.0 0.0.0.255 area 42
! eBGP is used for customerB
router bgp 64101
 bgp log-neighbor-changes
 neighbor 192.168.40.1 remote-as 100
 !
 address-family ipv4
  network 172.16.10.0 mask 255.255.255.0        ! advertise those routes through
ebgp
  network 172.16.11.0 mask 255.255.255.0
  neighbor 192.168.40.1 activate
```

Verification of PE and CE Routers

Verifying BGP Session establishment between two PE Routers

```
ASR903-PE2# show ip bgp summary
BGP router identifier 10.10.100.100, local AS number 100
BGP table version is 1, main routing table version 1

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
10.10.3.3       4          100    2041    2038        1    0    0   1d06h           0
```

Verifying VPN Routes Learnt from Local CE and Remote PE Routers

```
ASR903-PE2#show ip bgp vpnv4 vrf CustomerA
BGP table version is 25, local router ID is 10.10.100.100

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,

              x best-external, a additional-path, c RIB-compressed,

Origin codes: i - IGP, e - EGP, ? - incomplete

RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 100:41 (default for vrf CustomerA)
 *>i 172.16.0.0/24    10.10.3.3           2    100      0 ?       ! routes from PE1
 *>i 172.16.1.0/24    10.10.3.3           2    100      0 ?       ! routes from PE1
 *>  172.16.10.0/24   192.168.40.2        0             0 64101 I ! routes from local
 *>  172.16.11.0/24   192.168.40.2        0             0 64101 I ! CE
 *>i 192.168.13.0/29  10.10.3.3           0    100      0 ?
```

```
ASR903-PE1#show ip bgp vpnv4 all
BGP table version is 27, local router ID is 10.10.3.3
     Network          Next Hop            Metric  LocPrf Weight Path
Route Distinguisher: 100:41 (default for vrf CustomerA)
 *>  172.16.0.0/24    192.168.13.1             2         32768 ?
 *>  172.16.1.0/24    192.168.13.1             2         32768 ?
 *>i 172.16.10.0/24   10.10.4.4               0    100      0 64101 I  ! routes from
 *>i 172.16.11.0/24   10.10.4.4               0    100      0 64101 I  ! PE2
 *>  192.168.13.0/29  0.0.0.0                  0         32768 ?
Route Distinguisher: 100:42 (default for vrf CustomerB)
 *>  10.5.0.0/24      192.168.14.2            20         32768 ?
 *>  10.5.1.0/24      192.168.14.2            20         32768 ?
 *>i 10.5.10.0/24     10.10.4.4                2    100      0 ?    ! routes from PE2
 *>i 10.5.11.0/24     10.10.4.4                2    100      0 ?    ! routes from PE2
 *>i 192.168.41.0/29  10.10.4.4               0    100      0 ?
```

Verifying the Routing Tables on CE Devices

```
CE1-A# show ip route
      172.16.0.0/16 is variably subnetted, 6 subnets, 2 masks
C        172.16.0.0/24 is directly connected, Loopback110     ! local routes
L        172.16.0.1/32 is directly connected, Loopback110
C        172.16.1.0/24 is directly connected, Loopback111
L        172.16.1.1/32 is directly connected, Loopback111
O E2     172.16.10.0/24                              ! routes from CustomerA site 2
           [110/1] via 192.168.13.3, 04:46:18, GigabitEthernet0/1/1
O E2     172.16.11.0/24                              ! routes from CustomerA site 2
           [110/1] via 192.168.13.3, 04:46:18, GigabitEthernet0/1/1
      192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.13.0/29 is directly connected, GigabitEthernet0/1/1
L        192.168.13.1/32 is directly connected, GigabitEthernet0/1/1

CE2# show ip route
    10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E2     10.5.0.0/24 [110/1] via 192.168.41.1, 04:50:16, BDI42! routes from CustomerB
site 1
O E2     10.5.1.0/24 [110/1] via 192.168.41.1, 04:43:23, BDI42! routes from CustomerB
site 1
C        10.5.10.0/24 is directly connected, Loopback100
L        10.5.10.1/32 is directly connected, Loopback100
C        10.5.11.0/24 is directly connected, Loopback101
L        10.5.11.1/32 is directly connected, Loopback101
      172.16.0.0/16 is variably subnetted, 6 subnets, 2 masks
B        172.16.0.0/24 [20/0] via 192.168.40.1, 1d00h ! routes from CustomerA site 1
B        172.16.1.0/24 [20/0] via 192.168.40.1, 1d00h ! routes from CustomerA site 1
C        172.16.10.0/24 is directly connected, Loopback90
```

```
L        172.16.10.1/32 is directly connected, Loopback90
C        172.16.11.0/24 is directly connected, Loopback91
L        172.16.11.1/32 is directly connected, Loopback91
      192.168.13.0/29 is subnetted, 1 subnets
B        192.168.13.0 [20/0] via 192.168.40.1, 1d00h

      192.168.40.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.40.0/29 is directly connected, BDI41
L        192.168.40.2/32 is directly connected, BDI41
      192.168.41.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.41.0/29 is directly connected, BDI42
L        192.168.41.2/32 is directly connected, BDI42

CE1-A#  ping 172.16.10.1 source 172.16.0.1    ! connectivity between customerA sites
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.0.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms

CE2# ping 10.5.0.1 source 10.5.10.1           ! connectivity between customerB sites
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.5.0.1, timeout is 2 seconds:
Packet sent with a source address of 10.5.10.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Chapter 4 Ethernet over MPLS

In Ethernet over MPLS environment, Ethernet frames are exchanged between customer sites using SP backbone as the medium of transport.

Figure 8. Ethernet over MPLS

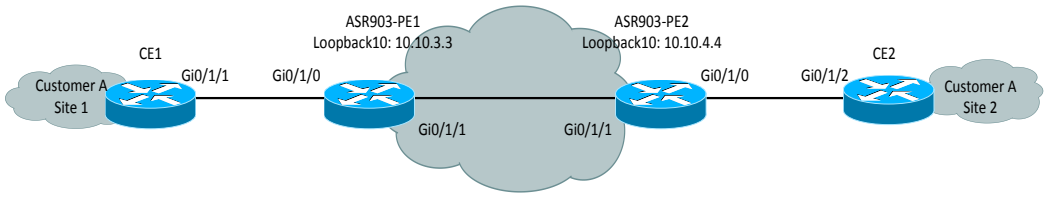

Ethernet over MPLS using Ethernet Flow Point

ASR903 only supports service instance based pseudo-wire.

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
service instance 2 ethernet
  encapsulation dot1q 200
  xconnect 10.10.4.4 200 encapsulation mpls         ! create PW
```

ASR903-PE2 Configuration

```
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 2 ethernet
  encapsulation dot1q 200
  xconnect 10.10.3.3 200 encapsulation mpls         ! create PW, VCID must be same
```

CE1 Configuration

```
interface BDI200
```

```
 ip address 192.168.200.1 255.255.255.0
interface GigabitEthernet0/1/1
 no ip address
 negotiation auto
 service instance 200 ethernet
  encapsulation dot1q 200
  rewrite ingress tag pop 1 symmetric
  bridge-domain 200
```

CE2 Configuration

```
interface BDI200
 ip address 192.168.200.2 255.255.255.0
interface GigabitEthernet0/1/2
 no ip address
 negotiation auto
 service instance 200 ethernet
  encapsulation dot1q 200
  rewrite ingress tag pop 1 symmetric
  bridge-domain 200
```

The following options are also supported for EoMPLS:

1. Remove a vlan tag
2. Same PW used for a range of Vlans
3. Push a Vlan tag

Removing the Vlan Tag before Transmitting on the PW

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
service instance 2 ethernet
  encapsulation dot1q 200
  rewrite ingress tag pop 1 symmetric           ! pop the vlan header
  xconnect 10.10.4.4 200 encapsulation mpls
```

ASR903-PE2 Configuration

```
interface GigabitEthernet0/1/0
service instance 2 ethernet
  encapsulation dot1q 300                    ! notice different vlan tag
  rewrite ingress tag pop 1 symmetric        ! pop the vlan header
  xconnect 10.10.4.4 200 encapsulation mpls
```

Same PW used for a Range of VLANs

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
service instance 2 ethernet
  encapsulation dot1q 200-300
  xconnect 10.10.4.4 200 encapsulation mpls
```

ASR903-PE2 Configuration

```
interface GigabitEthernet0/1/0
service instance 2 ethernet
  encapsulation dot1q 200-300
  xconnect 10.10.4.4 200 encapsulation mpls
```

Push a Vlan Tag

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
service instance 2 ethernet
  encapsulation dot1q 200-300
  rewrite ingress tag push dot1q 300 symmetric            ! push a vlan tag
  xconnect 10.10.4.4 200 encapsulation mpls
```

ASR903-PE2 Configuration

```
interface GigabitEthernet0/1/0
service instance 2 ethernet
  encapsulation dot1q 200-300                             ! push a vlan tag
  rewrite ingress tag push dot1q 300 symmetric
  xconnect 10.10.4.4 200 encapsulation mpls
```

New l2vpn Commands

There are a new set of commands available to create EoMPLS as illustrated in PE configuration.

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 1 ethernet
```

```
  encapsulation dot1q 200-300
  rewrite ingress tag push dot1q 300 symmetric
l2vpn xconnect context PW200                 ! use l2vpn context

 member GigabitEthernet0/1/0 service-instance 1
 member 10.10.4.4 200 encapsulation mpls
```

ASR903-PE2 Configuration

```
! create a pesudowire interface
interface pseudowire200
 encapsulation mpls
 neighbor 10.10.3.3 200
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 1 ethernet
  encapsulation dot1q 200-300
  rewrite ingress tag push dot1q 300 symmetric
l2vpn xconnect context PW200                 ! use l2vpn context
 member GigabitEthernet0/1/0 service-instance 1
 member pseudowire200
```

Port mode by using Encapsulation Default

Important: Port mode is *not* supported on ASR903 but encapsulation default can be used as a workaround.

ASR903-PE1 Configuration

```
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 1 ethernet
  encapsulation default                  ! this encapsulation catches all
  xconnect 10.10.4.4 100 encapsulation mpls
```

ASR903-PE2 Configuration

```
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 1 ethernet
  encapsulation default
  xconnect 10.10.3.3 100 encapsulation mpls
```

CE1 Configuration

```
interface GigabitEthernet0/1/1
ip address 192.168.13.1 255.255.255.248
```

CE2 Configuration

```
interface GigabitEthernet0/1/1
ip address 192.168.13.2 255.255.255.248
```

Verify the L2VPN is up and running and CE1 can reach CE2, vice versa.

```
ASR903-PE2# show mpls l2transport vc
Local intf     Local circuit              Dest address    VC ID      Status
-------------  -------------------------- --------------- ---------- ----------
Gi0/1/0        Ethernet:1                 10.10.3.3       100         UP

CE1# ping 192.168.13.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.13.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Multi-segment Pseudowire

Multi-segment pseudowire extends the pseudo wire into multiple autonomous systems. In the case below, PW between PE1 and PE2 have two segments.

Figure 9. Multi-segment Pseudowire

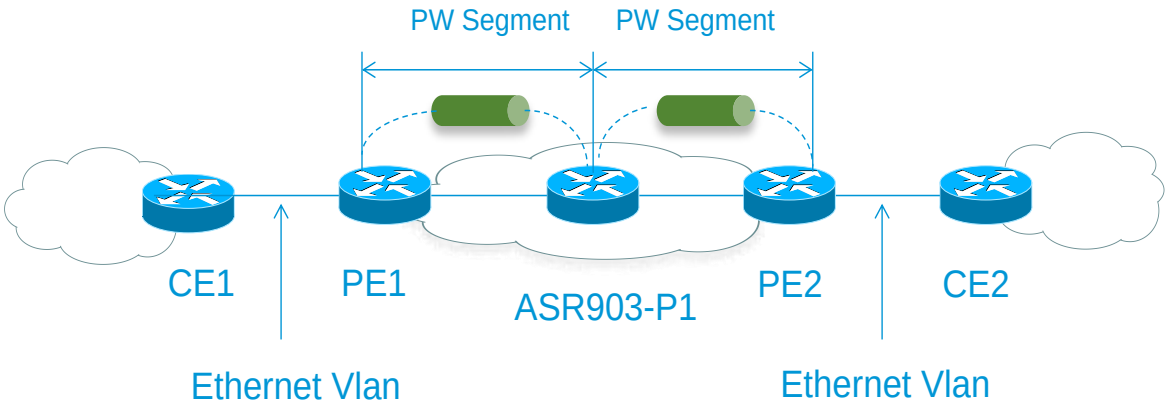

ASR903-PE1 Configuration (10.10.3.3)

```
interface GigabitEthernet0/0/4
 service instance 2 ethernet
  encapsulation dot1q 100 second-dot1q 10
  rewrite ingress tag pop 1 symmetric
  xconnect 10.10.2.2 1100 encapsulation mpls
```

ASR903-PE2 Configuration (10.10.4.4)

```
interface GigabitEthernet0/0/7
 service instance 2 ethernet
  encapsulation dot1q 100 second-dot1q 10
  rewrite ingress tag pop 1 symmetric
  xconnect 10.10.2.2 1100 encapsulation mpls
```

ASR903-P1 Configuration (10.10.2.2)

```
l2vpn xconnect context 1100
 member 10.10.3.3 1100 encapsulation mpls
 member 10.10.4.4 1100 encapsulation mpls
```

Verifying Pseudowire

```
ASR903-P1# show mpls l2transport vc
Local intf         Local circuit            Dest address    VC ID     Status
-------------  -------------------------- ----------------- -------- ----------
pw100010           10.10.4.4 1100             10.10.3.3      1100       UP
pw100009           10.10.3.3 1100             10.10.4.4      1100       UP
```

Ethernet OAM

CFM, 802.3ah(ethernet in the first mile), Ethernet Local Management Interface (LMI) are supported on the ASR903.

Connectivity Fault Management

Ethernet connectivity fault management (CFM) is an end-to-end per-service-instance Ethernet layer OAM protocol that includes proactive connectivity monitoring, fault verification, and fault isolation.

ASR903-PE1 Configuration

```
ethernet cfm ieee
ethernet cfm global                          ! enable CFM on the router
ethernet cfm domain PE1-2 level 6            ! define domain PE1-2
 service EVC-PE-200 evc evc-200
  continuity-check
  continuity-check interval 1s
!
ethernet cfm logging
ethernet evc evc-200
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 200 ethernet evc-200
  encapsulation dot1q 200-300
  cfm mep domain PE1-2 mpid 1200             ! created MEP
l2vpn xconnect context PW200
 member GigabitEthernet0/1/0 service-instance 200
 member 10.10.4.4 200 encapsulation mpls
```

ASR903-PE2 Configuration

```
ethernet cfm ieee
ethernet cfm global                          ! enable CFM on the router
ethernet cfm domain PE1-2 level 6
 service EVC-PE-200 evc evc-200
  continuity-check
  continuity-check interval 1s
!
ethernet cfm logging
ethernet evc evc-200
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
```

```
 service instance 200 ethernet evc-200
  encapsulation dot1q 200-300
  cfm mep domain PE1-2 mpid 1201          ! mpid must be different from remote end
interface pseudowire200
 encapsulation mpls
 neighbor 10.10.3.3 200
l2vpn xconnect context PW200
 member GigabitEthernet0/1/0 service-instance 200
 member pseudowire200
```

Verify OAM

```
ASR903-PE2# show ethernet cfm maintenance-points local
Local MEPs:
--------------------------------------------------------------------------------
MPID Domain Name                                Lvl   MacAddress     Type  CC
Ofld Domain Id                                  Dir   Port           Id
     MA Name                                          SrvcInst       Source
     EVC name
--------------------------------------------------------------------------------
1201 PE1-2                                      6     7010.5c51.a4bf XCON  Y
No   PE1-2                                      Up    Gi0/1/0        N/A
     EVC-PE-200                                       200            Static
     evc-200

Total Local MEPs: 1

ASR903-PE2# show ethernet cfm maintenance-points remote
--------------------------------------------------------------------------------
MPID  Domain Name                                MacAddress          IfSt  PtSt
 Lvl  Domain ID                                  Ingress
 RDI  MA Name                                    Type Id             SrvcInst
      EVC Name                                                       Age
      Local MEP Info
--------------------------------------------------------------------------------
1200  PE1-2                                      7010.5c51.8fbf       Up     Up
 6    PE1-2                                      Gi0/1/0:(10.10.3.3, 200)
 -    EVC-PE-200                                 XCON N/A            200
      evc-200                                                        0s
      MPID: 1201 Domain: PE1-2 MA: EVC-PE-200

Total Remote MEPs: 1

ASR903-PE2# ping ethernet mpid 1200 domain PE1-2 service EVC-PE-200 source 1201
Type escape sequence to abort.
Sending 5 Ethernet CFM loopback messages to 7010.5c51.8fbf, timeout is 5
seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/4 ms

ASR903-PE2# traceroute ethernet mpid 1200 domain PE1-2 service EVC-PE-200  source
1201
Type escape sequence to abort. TTL 64. Linktrace Timeout is 5 seconds
```

```
Tracing the route to 7010.5c51.8fbf on Domain PE1-2, Level 6, service EVC-PE-200, evc
evc-200
Traceroute sent via Gi0/1/0:(10.10.3.3, 200), path found via MPDB

B = Intermediary Bridge
! = Target Destination
* = Per hop Timeout
-------------------------------------------------------------------------------
                             MAC      Ingress        Ingr Action  Relay Action
  Hops   Host             Forwarded   Egress         Egr Action   Previous Hop
-------------------------------------------------------------------------------
B 1                   7010.5c51.a4bf Gi0/1/0        IngOk        RlyMPDB
                      Forwarded
! 2                   7010.5c51.8fbf                             RlyHit:MEP
                      Not Forwarded                              7010.5c51.a4bf
```

Y.1731 one way delay measurement, PE1 initiates the packets and PE2 receive the packets. To get the accurate one way delay measurement, it is recommended that those devices are time synchronized through either PTP or NTP, preferably PTP.

ASR903-PE2 Configuration

```
ip sla 100
 ethernet y1731 delay receive 1DM domain PE1-2 evc evc-200 cos 0 mpid 1201
ip sla schedule 100 start-time pending              ! start receive end first
```

ASR903-PE1 Configuration

```
ip sla 100
 ethernet y1731 delay 1DM domain PE1-2 evc evc-200 mpid 1201 cos 0 source mpid 1200
  aggregate interval 30
ip sla schedule 100 start-time now
```

Y.1731 two way delay measurement, here is an example to configure DMM on router PE1.

ASR903-PE1 Configuration

```
ip sla 110
 ethernet y1731 delay DMM domain PE1-2 evc evc-200 mpid 1201 cos 0 source mpid 1200
  aggregate interval 30
ip sla schedule 110 life forever start-time now

ASR903-PE1# show ip sla statistics 110 (detail)
IPSLAs Latest Operation Statistics

IPSLA operation id: 110
Delay Statistics for Y1731 Operation 110
Type of operation: Y1731 Delay Measurement
Latest operation start time: *02:24:08.354 UTC Sat Aug 10 2013
```

```
Latest operation return code: OK
Distribution Statistics:

Interval
 Start time:  *02:24:08.354 UTC Sat Aug 10 2013
 Elapsed time: 29 seconds
 Number of measurements initiated: 24
 Number of measurements completed: 24
 Flag: OK
```

Chapter 5 Virtual Private LAN Service - VPLS

VPLS enables enterprises to link together their Ethernet-based LANs from multiple sites via the infrastructure provided by their service provider. From the enterprise perspective, the service provider's public network looks like one giant Ethernet LAN. For the service provider, VPLS provides an opportunity to deploy another revenue-generating service on top of their existing network without major capital expenditures. Operators can extend the operational life of equipment in their network.

Virtual Private LAN Services (VPLS) uses the provider core to join multiple attachment circuits together to simulate a virtual bridge that connects the multiple attachment circuits together. From a customer point of view, there is no topology for VPLS. All of the CE devices appear to connect to a logical bridge emulated by the provider core (see the figure below).

Figure 10. VPLS

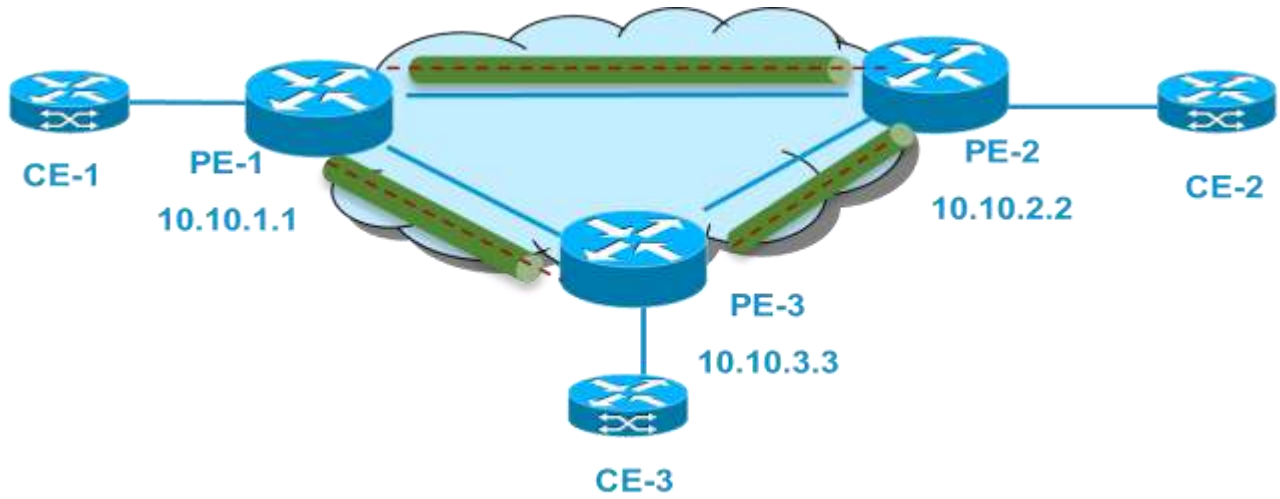

Manually Provisioned VPLS

PE-1 Configuration

```
! use the legacy l2 VPLS configuration commands
l2 vfi VPLS30 manual
 vpn id 30
 bridge-domain 30
 neighbor 10.10.3.3 encapsulation mpls
```

```
 neighbor 10.10.2.2 encapsulation mpls
! Interface to CE-1
interface GigabitEthernet0/0/7
  service instance 2 ethernet evc30
    encapsulation dot1q 30

bridge-domain 30
```

PE-2 Configuration

```
l2vpn vfi context VPLS30                  ! new l2vpn commands to VFI context
 vpn id 30
 member 10.10.1.1 encapsulation mpls
 member 10.10.3.3 encapsulation mpls
bridge-domain 30
 member GigabitEthernet0/0/7 service-instance 2
 member vfi VPLS30
interface GigabitEthernet0/0/7
  service instance 2 ethernet evc30
    encapsulation dot1q 30
```

PE-3 Configuration

```
l2vpn vfi context VPLS30
 vpn id 30
 member 10.10.1.1 encapsulation mpls
 member 10.10.2.2 encapsulation mpls
bridge-domain 30
 member GigabitEthernet0/0/7 service-instance 2
 member vfi VPLS30
interface GigabitEthernet0/0/7
  service instance 2 ethernet evc30
encapsulation dot1q 30
```

Verifying VC Neighbors

```
PE-1# show vfi
Legend: RT=Route-target, S=Split-horizon, Y=Yes, N=No
VFI name: VPLS30, state: up, type: multipoint, signaling: LDP
  VPN ID: 30
  Bridge-Domain 30 attachment circuits:
  Neighbors connected via pseudowires:
  Peer Address        VC ID        S
  10.10.2.2           30           Y
  10.10.3.3           30           Y
```

VPLS – BGP Auto-discovery and LDP Signaling

VPLS can rely on BGP protocol for discovering and provisioning PEs, which reduces VPN configuration and errors associated with configuration.

PE-1 configuration

```
 l2 vfi VPLS35 autodiscovery              ! use legacy commands
    vpn id 35
bridge-domain 35
interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 35 ethernet
  encapsulation dot1q 35
bridge-domain 35

router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.2.2 remote-as 100
 neighbor 10.10.2.2 update-source Loopback10
 neighbor 10.10.3.3 remote-as 100
 neighbor 10.10.3.3 update-source Loopback10
 !
 address-family l2vpn vpls
  neighbor 10.10.2.2 activate
  neighbor 10.10.2.2 send-community extended
  neighbor 10.10.3.3 activate
  neighbor 10.10.3.3 send-community extended
```

PE-2 Configuration

```
l2vpn vfi context VPLS35           ! use new l2vpn commands
 vpn id 35
 autodiscovery bgp signaling ldp    ! use bgp to discover VPLS neighbor and LDP to
negotiate
                             ! inner VC label
bridge-domain 35                    ! associate Attachement circuit with VPLS VFI
 member GigabitEthernet0/1/0 service-instance 35
 member vfi VPLS35

interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 35 ethernet
  encapsulation dot1q 35
```

```
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.1.1 remote-as 100
 neighbor 10.10.1.1 update-source Loopback10
 neighbor 10.10.3.3 remote-as 100
 neighbor 10.10.3.3 update-source Loopback10
 !
 address-family l2vpn vpls
  neighbor 10.10.1.1 activate
  neighbor 10.10.1.1 send-community extended
  neighbor 10.10.3.3 activate
  neighbor 10.10.3.3 send-community extended
```

PE-3 Configuration

```
l2vpn vfi context VPLS35           ! use new l2vpn commands
 vpn id 35
 autodiscovery bgp signaling ldp   ! use bgp to discover VPLS neighbor and LDP to
negotiate
                                   ! inner VC label
bridge-domain 35                   ! associate Attachement circuit with VPLS VFI
 member GigabitEthernet0/1/0 service-instance 35
 member vfi VPLS35

interface GigabitEthernet0/1/0
 no ip address
 negotiation auto
 no keepalive
 service instance 35 ethernet
  encapsulation dot1q 35
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.1.1 remote-as 100
 neighbor 10.10.1.1 update-source Loopback10
 neighbor 10.10.2.2 remote-as 100
 neighbor 10.10.2.2 update-source Loopback10
 !
 address-family l2vpn vpls
  neighbor 10.10.1.1 activate
  neighbor 10.10.1.1 send-community extended
  neighbor 10.10.2.2 activate
  neighbor 10.10.2.2 send-community extended
```

Verifying VPLS Neighbors

```
PE-1# show vfi name VPLS35
Legend: RT=Route-target, S=Split-horizon, Y=Yes, N=No
VFI name: VPLS35, state: up, type: multipoint, signaling: LDP
  VPN ID: 35, VPLS-ID: 100:35
  RD: 100:35, RT: 100:35
  Bridge-Domain 35 attachment circuits:
```

```
Neighbors connected via pseudowires:
  Peer Address     VC ID        Discovered Router ID    S
  10.10.2.2         35               10.10.2.2         Y
  10.10.3.3         35               10.10.3.3         Y

PE-1# show mpls l2transport vc 35
Local intf           Local circuit           Dest address    VC ID      Status
------------------  ---------------------- ---------------------- ---------- -----
-----
VFI VPLS35          vfi                        10.10.2.2      35         UP
VFI VPLS35          vfi                        10.10.3.3      35         UP
```

VPLS – BGP Auto-discovery and Signaling

The VPLS control plane is used for auto-discovery and signaling. Auto-discovery involves locating all provider edge (PE) devices that participate in a particular VPLS instance. Signaling is accomplished by configuring pseudo-wires for a VPLS instance. In the past, LDP was used for signaling and BGP was used for auto-discovery (RFC 6074). Now, the VPLS BGP Signaling L2VPN feature supports RFC 4761 by simplifying the auto-discovery and signaling of all known PE devices in a VPLS instance by using BGP for both functions.

PE-1 configuration

```
l2vpn vfi context VPLS30
 vpn id 30
 autodiscovery bgp signaling bgp
   ve id 32                                ! VE ID must unique in the same VPLS domain
 bridge-domain 30
   member GigabitEthernet0/0/4 service-instance 2
   member vfi VPLS30
router bgp 100
neighbor 10.10.2.2 remote-as 100
 neighbor 10.10.2.2 update-source Loopback10
 neighbor 10.10.3.3 remote-as 100
 neighbor 10.10.3.3 update-source Loopback10
 address-family l2vpn vpls
  neighbor 10.10.2.2 activate
  neighbor 10.10.2.2 send-community extended
  neighbor 10.10.2.2 suppress-signaling-protocol ldp          ! suppress LDP to use
BGP
  neighbor 10.10.3.3 activate
  neighbor 10.10.3.3 send-community extended
  neighbor 10.10.3.3 suppress-signaling-protocol ldp
```

PE-2 configuration

```
l2vpn vfi context VPLS30
 vpn id 30
 autodiscovery bgp signaling bgp
   ve id 34
 bridge-domain 30
   member GigabitEthernet0/0/4 service-instance 2
   member vfi VPLS30
router bgp 100
 neighbor 10.10.1.1 remote-as 100
 neighbor 10.10.1.1 update-source Loopback10
 neighbor 10.10.3.3 remote-as 100
 neighbor 10.10.3.3 update-source Loopback10
 !
 address-family l2vpn vpls
  neighbor 10.10.1.1 activate
  neighbor 10.10.1.1 send-community extended
  neighbor 10.10.1.1 suppress-signaling-protocol ldp
  neighbor 10.10.3.3 activate
  neighbor 10.10.3.3 send-community extended
  neighbor 10.10.3.3 suppress-signaling-protocol ldp
```

PE-3 configuration

```
l2vpn vfi context VPLS30
 vpn id 30
 autodiscovery bgp signaling bgp
  ve id 33  ç
bridge-domain 30
 member GigabitEthernet0/0/7 service-instance 2
 member vfi VPLS30
router bgp 100
 neighbor 10.10.1.1 remote-as 100
 neighbor 10.10.1.1 update-source Loopback10
 neighbor 10.10.2.2 remote-as 100
 neighbor 10.10.2.2 update-source Loopback10
 !
 address-family l2vpn vpls
  neighbor 10.10.1.1 activate
  neighbor 10.10.1.1 send-community extended
  neighbor 10.10.2.2 suppress-signaling-protocol ldp
  neighbor 10.10.2.2 activate
  neighbor 10.10.2.2 send-community extended
  neighbor 10.10.2.2 suppress-signaling-protocol ldp
```

Verifying VPLS

```
PE-1# show vfi name VPLS30
Legend: RT=Route-target, S=Split-horizon, Y=Yes, N=No
VFI name: VPLS30, state: up, type: multipoint, signaling: BGP
```

```
VPN ID: 30, VE-ID: 32, VE-SIZE: 10
  RD: 100:30, RT: 100:30
  Bridge-Domain 30 attachment circuits:
  Neighbors connected via pseudowires:
  Interface           peer Address     VE-ID  Local Label  Remote Label    S
  pseudowire100023    10.10.2.2        32          30            33         Y
  pseudowire100024    10.10.3.3        33          31            35         Y

PE-1# show mpls l2transport vc 30
Local intf     Local circuit              Dest address    VC ID    Status
-------------  ------------------------ ---------------------- ---------- ----------
VFI VPLS30     vfi                        10.10.2.2        30         UP
VFI VPLS30     vfi                        10.10.3.3        30         UP
```

VPLS Hub-Spoke Configuration

VPLS use split-horizon to avoid loop so packets received in one VC are forwarded only to attachment circuits, not to other VCs. However in the hub-spoke scenario, split-horizon prevents spoke sites communicating with each other. In this case, it can be disabled.

Figure 11. VPLS Hub-Spoke

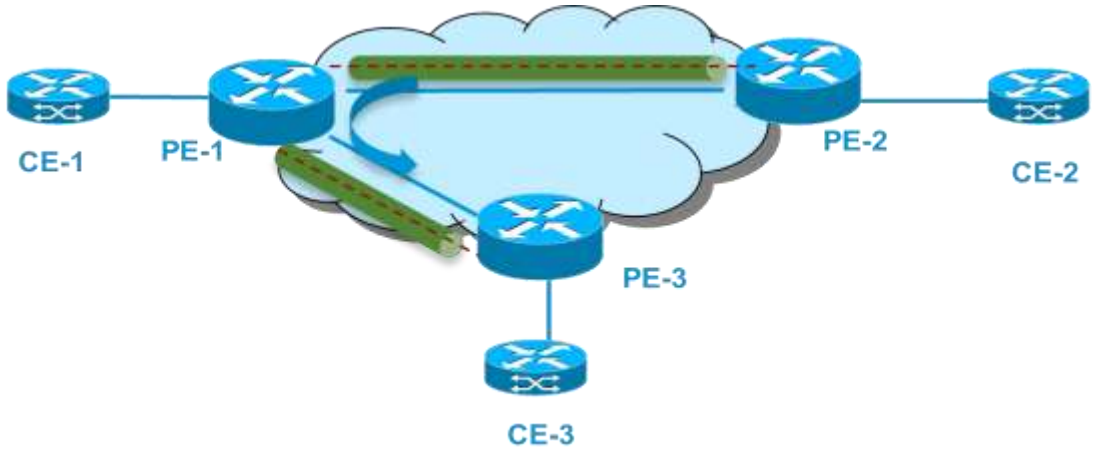

PE-1 Configuration

```
l2 vfi VPLS30 manual
 vpn id 30
 bridge-domain 30
 neighbor 10.10.3.3 encapsulation mpls  no-split-horizon
 neighbor 10.10.2.2 encapsulation mpls  no-split-horizon
! Interface to CE-1
```

```
interface GigabitEthernet0/0/7
  service instance 2 ethernet evc30
    encapsulation dot1q 30
    bridge-domain 30
```

H-VPLS with Ethernet Q-in-Q Access

In a flat or non-hierarchical VPLS configuration, a full mesh of pseudowires (PWs) is needed between all PE nodes. Hierarchical Virtual Private LAN Service (H-VPLS) reduces both signaling and replication overhead by using a combination of full-mesh and hub-and-spoke configurations.

Figure 12. H-VPLS with Ethernet Q-in-Q Access

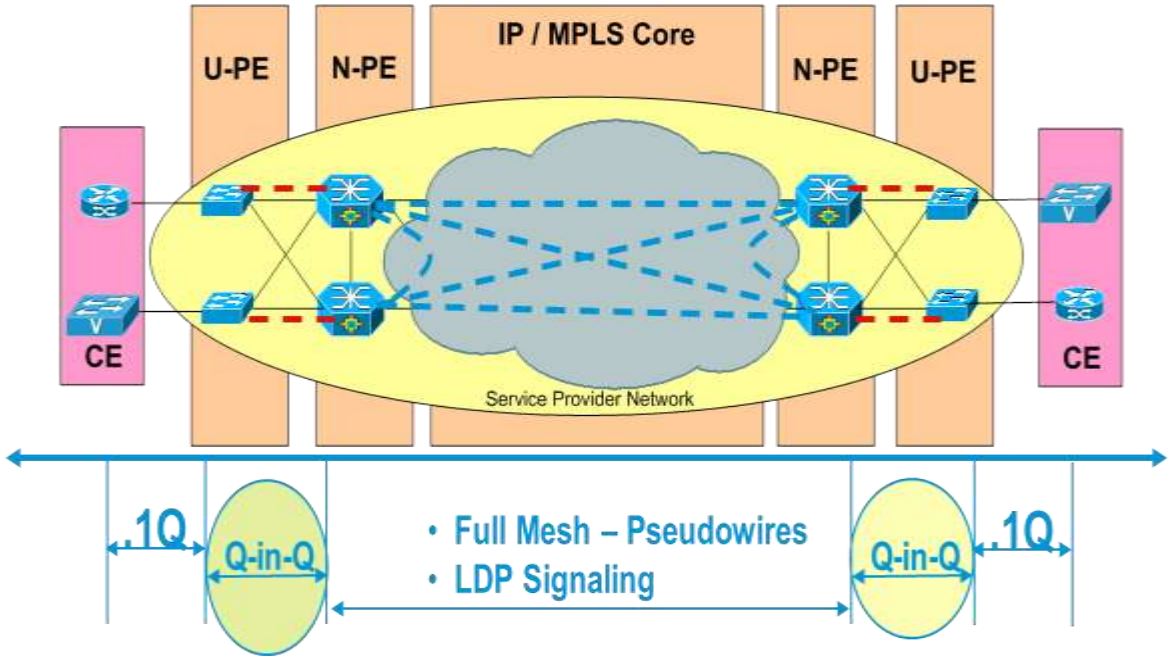

In the example below, U-PE device will add S-Tag 30 on top of customer vlan 10. N-PE device is only aware of S-TAG 30.

U-PE(ASR903)

```
! To CE interface, C- tag 10
interface GigabitEthernet0/0/2
 service instance 1 ethernet
  encapsulation dot1q 10
  bridge-domain 10
!  Interface to N-PE
!  S-TAG 30 is added on top of C-tag 10
interface GigabitEthernet0/0/1
 service instance 3 ethernet
  encapsulation dot1q 30 second-dot1q 10
```

```
 rewrite ingress tag pop 1 symmetric
 bridge-domain 10
```

N-PE (ASR 903)

```
! Interface to U-PE device
! N-PE device does not have to be c-tag aware
!
interface GigabitEthernet0/0/4
 service instance 2 ethernet
  encapsulation dot1q 30                    ! N-PE device is C-tag agnostic
  bridge-domain 30
 l2 vfi VPLS30 manual
 vpn id 30
 bridge-domain 30
 neighbor 10.10.3.3 encapsulation mpls
 neighbor 10.10.2.2 encapsulation mpls
```

H-VPLS with MPLS Access

Figure 13. H-VPLS with MPLS Access

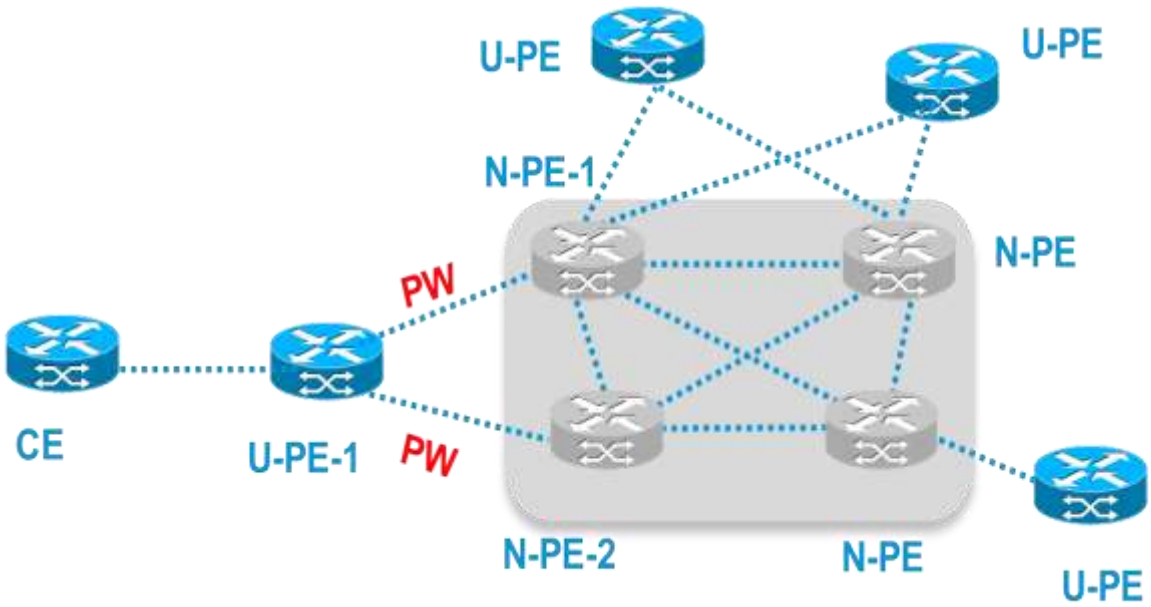

For H-VPLS with MPLS access based on pseudowire redundancy, the MPLS network has pseudowires to the VPLS core N-PE routers.

As shown in figure above, one pseudowire transports data between the U-PE router and its peer N-PE routers. When a failure occurs along the path of the U-PE router, the backup pseudowire and the redundant N-PE router become active and start transporting data.

U-PE-1(ASR-901)

```
interface GigabitEthernet0/2
 ! descripton interface to CE
 service instance 1 ethernet
   encapsulation dot1q 30
   xconnect 10.10.3.3 1300 encapsulation mpls         ! To N-PE-1
   backup peer 10.10.4.4 1300                          ! To N-PE-2
   mtu 1500
```

N-PE-1(ASR903)

```
! Create VPLS neighbors
l2vpn vfi context VPLS30
 vpn id 30
 member 10.10.4.4 encapsulation mpls               ! to other N-PE
 member 10.10.2.2 encapsulation mpls               ! to other N-PE
 bridge-domain 30
 member GigabitEthernet0/0/7 service-instance 2     ! this is local attachement
circuit
 member vfi VPLS30
 ! To MPLS access u-PE
 member 10.10.10.10 1300 encapsulation mpls

interface GigabitEthernet0/0/7
 service instance 2 ethernet
  encapsulation dot1q 30
```

Chapter 6 TDM

The Cisco ASR 903 router supports 16 T1/E1 and 4 port OC-3/1 port OC-12 combo card. HDLC, PPP, MLPPP, ATM and CEM are supported on the Cisco ASR 903 router.

Optical Channel Mapping and Related Configuration

The following configuration and diagram show how to configure channelized OC-3/12 down to T1/E1 level.

Figure 14. OC3 E1 Mapping at DS1 Level In SDH

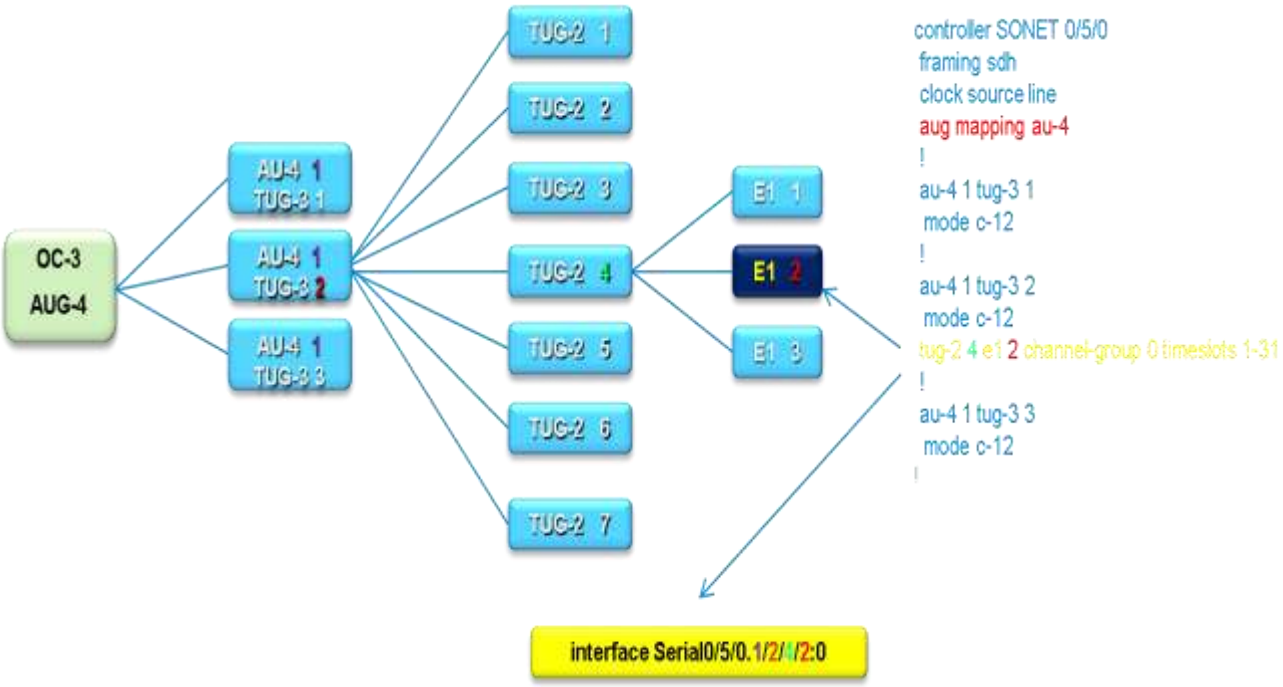

Figure 15. OC-3 T1 Mapping at DS1 Level in SDH

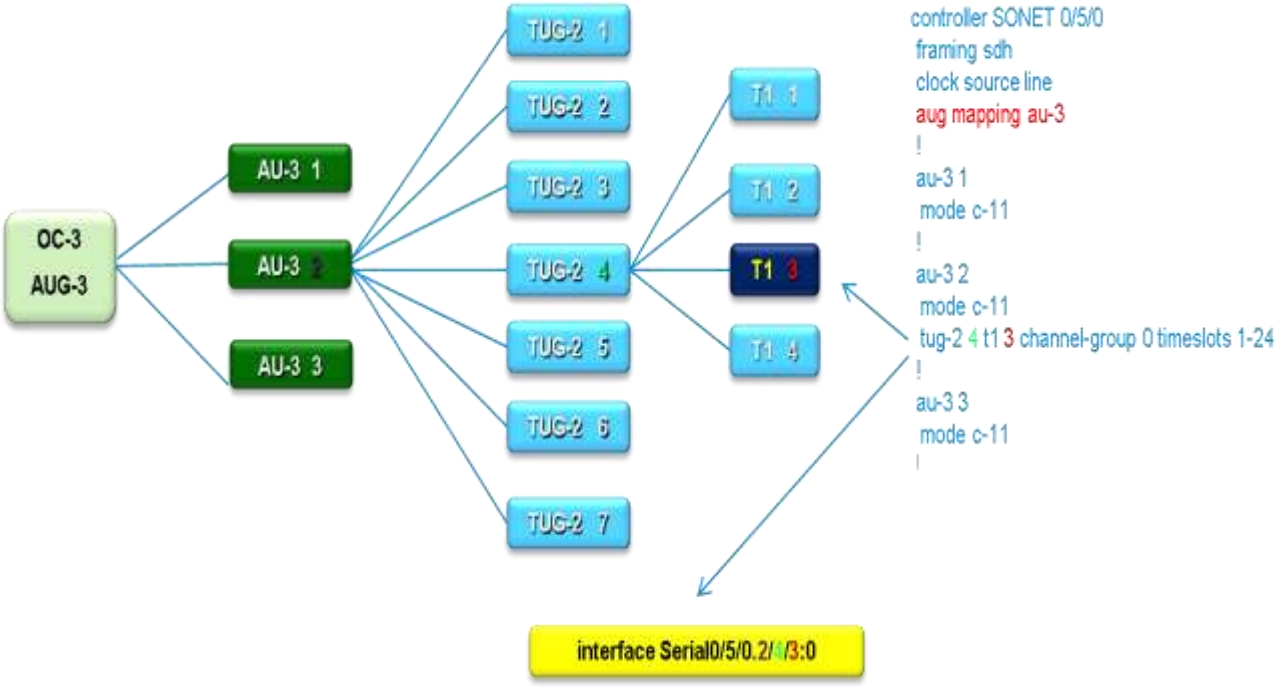

Figure 16. OC-3 T1 Mapping at DS1 Level in SONET

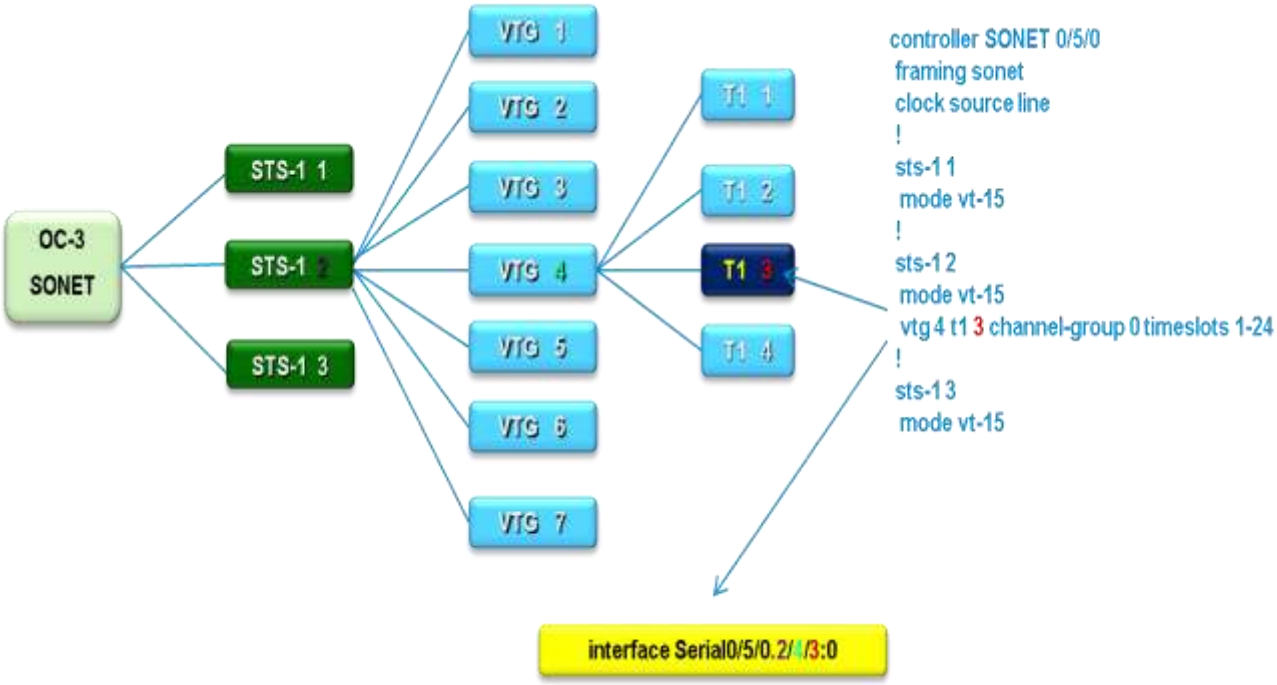

Figure 17. OC-12 E1 Mapping at DS1 Level in SDH

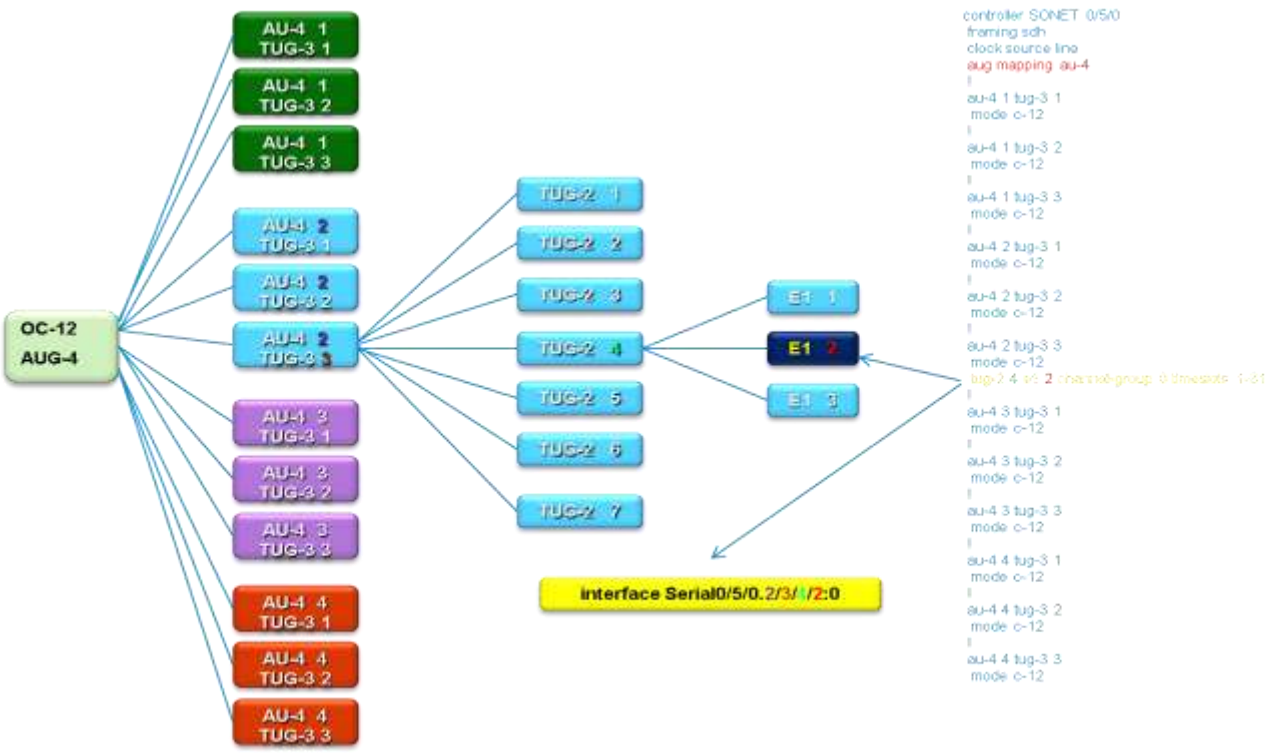

Figure 18. OC-12 T1 Mapping at DS1 Level in SDH

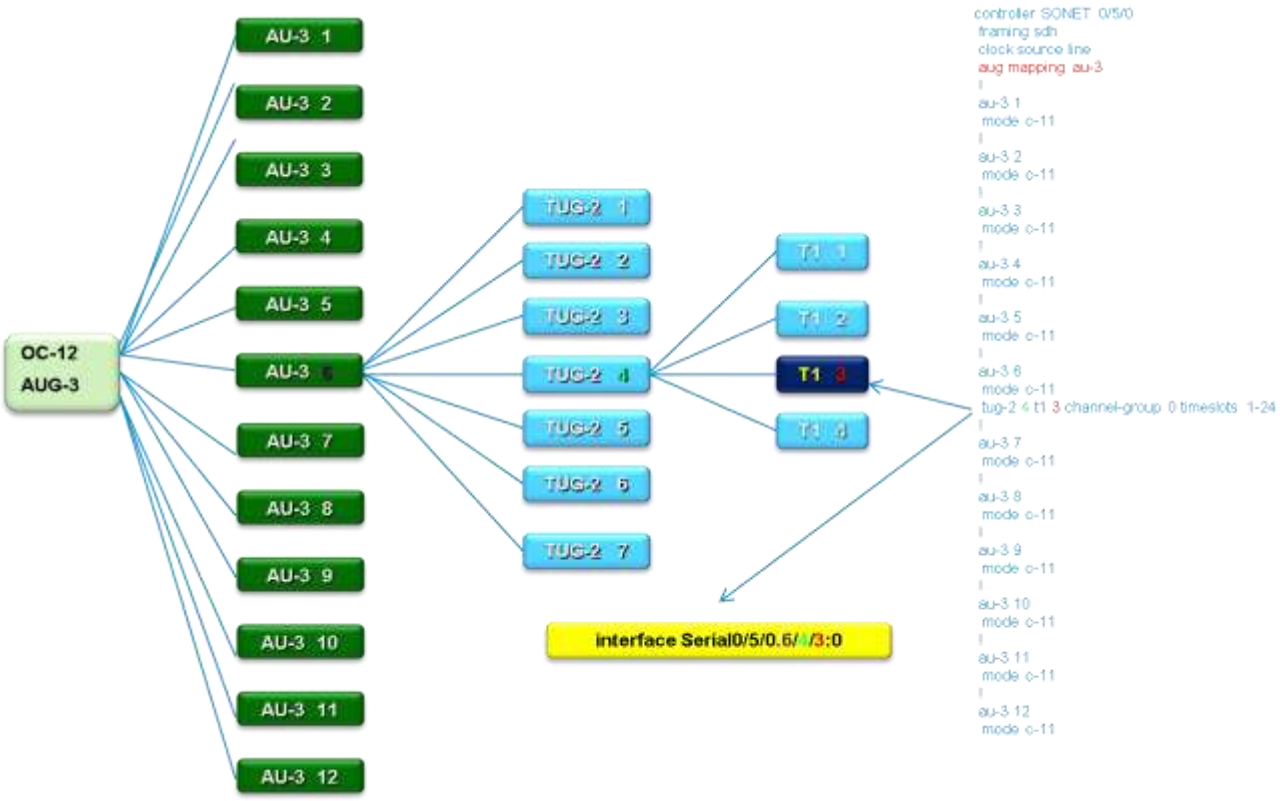

Figure 19. OC-12 T1 Mapping at DS1 Level in SONET

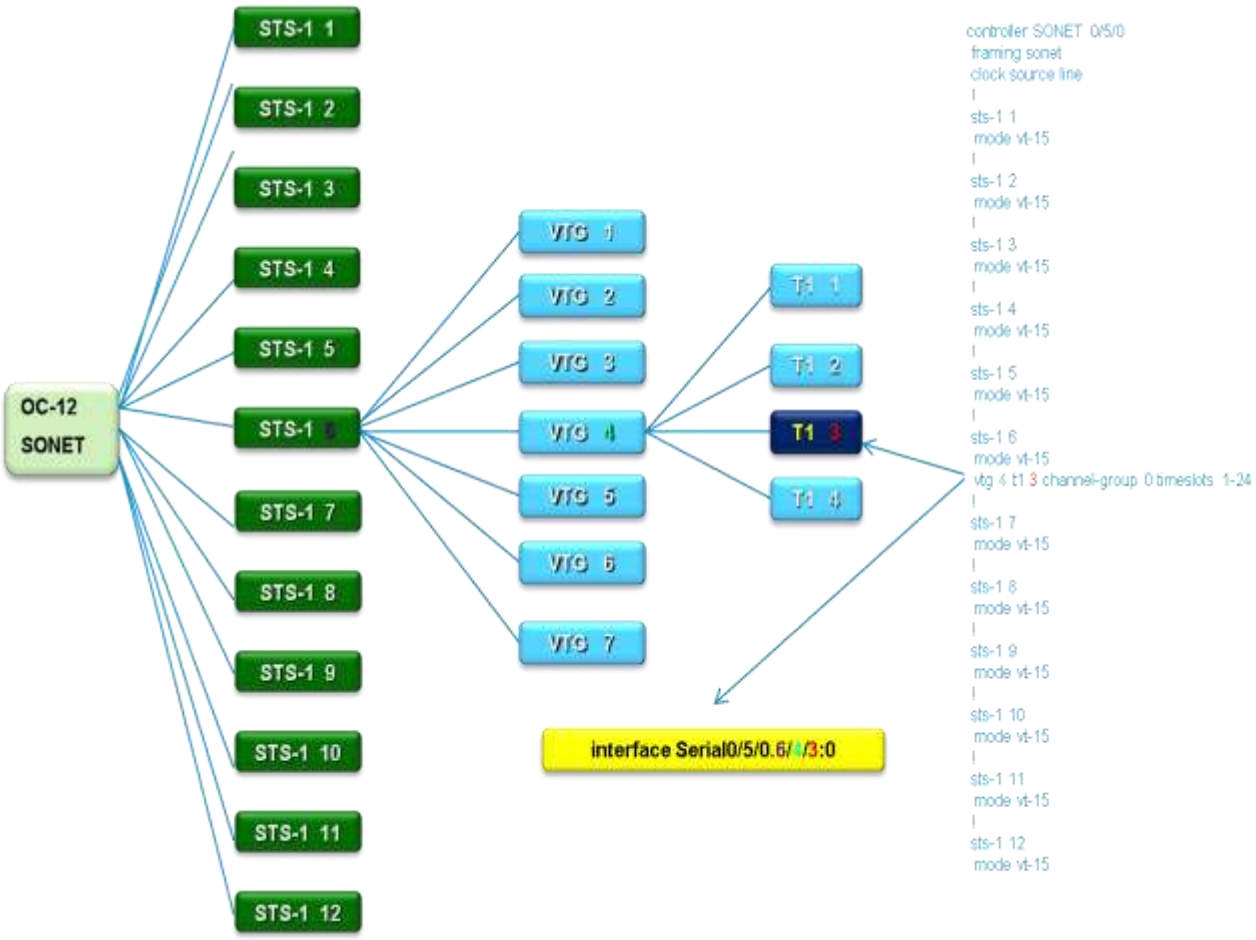

Circuit Emulation Service over Packet-Switched Network (CESoPSN)

Figure 20. CESoPN

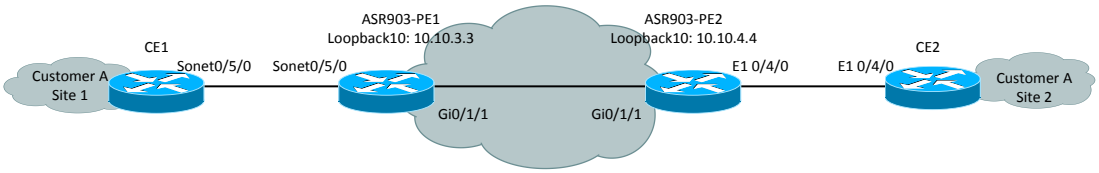

ASR903-PE1 Configuration

```
controller SONET 0/5/0
 framing sdh
 clock source internal
 aug mapping au-4
 !
 au-4 1 tug-3 1
  mode c-12
  tug-2 1 e1 1 cem-group 10 timeslots 1-31          ! CES for the whole E1 circuit
interface CEM0/5/0
 no ip address
 cem 10
  xconnect 10.10.4.4 101 encapsulation mpls
```

ASR903-PE2 Configuration

```
controller E1 0/4/0
 clock source internal
 cem-group 10 timeslots 1-31
!
interface CEM0/4/0
 no ip address
 cem 10
  xconnect 10.10.3.3 101 encapsulation mpls
```

CE1 Configuration (ASR903)

```
controller SONET 0/5/0
 framing sdh
 clock source line
 aug mapping au-4
 !
 au-4 1 tug-3 1
  mode c-12
  tug-2 1 e1 1 channel-group 0 timeslots 1-31       ! the CEM circuit
```

```
  tug-2 1 e1 2 channel-group 0 timeslots 1-31
interface Serial0/5/0.1/1/1/1:0
 ip address 100.100.20.1 255.255.255.248
 encapsulation ppp
```

CE2 Configuration (ASR 903)

```
controller E1 0/4/0
 channel-group 0 timeslots 1-31
interface Serial0/4/0:0
 ip address 100.100.20.2 255.255.255.248
 encapsulation ppp
```

Verifying CEM

```
ASR903-PE1# show mpls l2transport vc

Local intf     Local circuit              Dest address    VC ID      Status
-------------  -------------------------- --------------- ---------- ----------
CE0/5/0        CESoPSN Basic 10           10.10.4.4     101        UP

CE1# ping 100.100.20.2                 ! ping CE2 ip address
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.20.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 7/9/15 ms
```

Structure-Agnostic TDM over Packet (SAToP)

Figure 21. SAToP

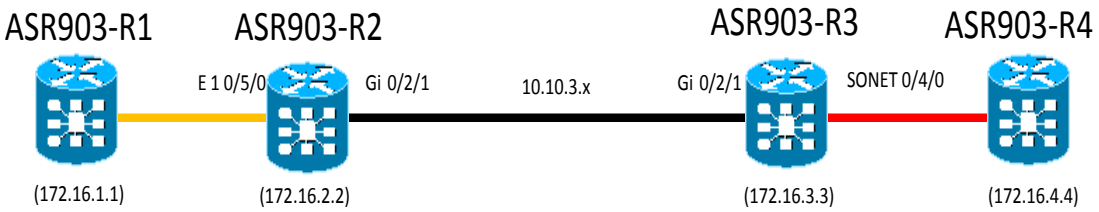

ASR903-R2 Configuration

```
!
controller E1 0/5/0
 framing unframed
 cem-group 0 unframed
!
interface Loopback0
 ip address 172.16.2.2 255.255.255.255
!
interface GigabitEthernet0/2/1
 ip address 10.10.3.1 255.255.255.0
 negotiation auto
 mpls ip
!
interface CEM0/5/0
 no ip address
 cem 0
  xconnect 172.16.3.3 1001 encapsulation mpls
 !
!
router ospf 1
 router-id 172.16.2.2
 network 172.16.2.2 0.0.0.0 area 0
 network 10.10.3.0 0.0.0.255 area 0
!
mpls ldp router-id Loopback0 force
!
```

ASR903-R3 Configuration:

```
!
controller SONET 0/4/0
 framing sdh
 clock source line
 aug mapping au-4
 !
 au-4 1 tug-3 1
  mode c-12
  tug-2 1 e1 1 cem-group 0 unframed
  tug-2 1 e1 1 framing unframed
 !
 au-4 1 tug-3 2
  mode c-12
 !
 au-4 1 tug-3 3
  mode c-12
!
interface Loopback0
 ip address 172.16.3.3 255.255.255.255
!
interface GigabitEthernet0/2/1
 ip address 10.10.3.2 255.255.255.0
```

```
 negotiation auto
 mpls ip
!
interface CEM0/4/0
 no ip address
 cem 0
  xconnect 172.16.2.2 1001 encapsulation mpls
 !
!
router ospf 1
 router-id 172.16.3.3
 network 172.16.3.3 0.0.0.0 area 0
 network 10.10.3.0 0.0.0.255 area 0
!
mpls ldp router-id Loopback0 force
!
```

Verifying TDM PW

```
ASR903-R2# show xconnect all detail
Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
  UP=Up       DN=Down            AD=Admin Down      IA=Inactive
  SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware

XC ST  Segment 1                         S1 Segment 2                          S2
------+---------------------------------+--+---------------------------------+--
UP pri   ac CE0/5/0:0(SATOP E1)          UP mpls 172.16.3.3:1001                 UP
            Interworking: L2L                    Local VC label 24
                                                 Remote VC label 22


ASR903-R2# show mpls l2transport vc 1001

Local intf     Local circuit              Dest address    VC ID      Status
-------------  -------------------------- --------------- ---------- ----------
CE0/5/0        SATOP E1 0                 172.16.3.3          1001          UP


ASR903-R2# show mpls l2transport vc 1001 detail
Local interface: CE0/5/0 up, line protocol up, SATOP E1 0 up
  Destination address: 172.16.3.3, VC ID: 1001, VC status: up
    Output interface: Gi0/2/1, imposed label stack {0 22}
    Preferred path: not configured
    Default path: active
    Next hop: 10.10.3.2
  Create time: 01:10:31, last status change time: 00:51:21
    Last label FSM state change time: 00:57:43
  Signaling protocol: LDP, peer 172.16.3.3:0 up
```

```
    Targeted Hello: 172.16.2.2(LDP Id) -> 172.16.3.3, LDP is UP
    Graceful restart: not configured and not enabled
    Non stop routing: not configured and not enabled
    Status TLV support (local/remote)   : enabled/supported
      LDP route watch                   : enabled
      Label/status state machine        : established, LruRru
      Last local dataplane   status rcvd: No fault
      Last BFD dataplane     status rcvd: Not sent
      Last BFD peer monitor  status rcvd: No fault
      Last local AC  circuit status rcvd: No fault
      Last local AC  circuit status sent: No fault
      Last local PW i/f circ status rcvd: No fault
      Last local LDP TLV     status sent: No fault
      Last remote LDP TLV    status rcvd: No fault
      Last remote LDP ADJ    status rcvd: No fault
    MPLS VC labels: local 24, remote 22
    Group ID: local 0, remote 0
    MTU: local 0, remote 0
    Remote interface description:
  Sequencing: receive disabled, send disabled
  Control Word: On (configured: autosense)
  SSO Descriptor: 172.16.3.3/1001, local label: 24
  Dataplane:
    SSM segment/switch IDs: 4098/8193 (used), PWID: 2
  VC statistics:
    transit packet totals: receive 3243829, send 3455361
    transit byte totals:   receive 882321488, send 912215304
transit packet drops:  receive 0, seq error 0, send 0


ASR903-R2# show controllers e1 0/5/0 brief
E1 0/5/0 is up.
  Applique type is A900-IMA16D
  Cablelength is Unknown
  No alarms detected.
  alarm-trigger is not set
  Framing is unframed, Line Code is HDB3, Clock Source is Line.
  Data in current interval (340 seconds elapsed):
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs, 0 Unavail Secs
     0 Near-end path failures, 0 Far-end path failures, 0 SEF/AIS Secs
  Total Data (last 24 hours)
     0 Line Code Violations, 0 Path Code Violations,
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins,
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs, 0 Unavail Secs
     0 Near-end path failures, 0 Far-end path failures, 0 SEF/AIS Secs
```

```
ASR903-R3#show xconnect all detail
Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
  UP=Up       DN=Down            AD=Admin Down      IA=Inactive
  SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware

XC ST        Segment 1                   S1 Segment 2                          S2
------+-----------------------------------+--+-----------------------------------+--
UP pri   ac CE0/4/0:0(SATOP E1)          UP   mpls 172.16.2.2:1001            UP
            Interworking: L2L                    Local VC label 22
                                                 Remote VC label 24

ASR903-R3#show mpls l2 vc 1001 detail
Local interface: CE0/4/0 up, line protocol up, SATOP E1 0 up
  Destination address: 172.16.2.2, VC ID: 1001, VC status: up
    Output interface: Gi0/2/1, imposed label stack {0 24}
    Preferred path: not configured
    Default path: active
    Next hop: 10.10.3.1
  Create time: 01:00:24, last status change time: 00:54:02
    Last label FSM state change time: 01:00:24
  Signaling protocol: LDP, peer 172.16.2.2:0 up
    Targeted Hello: 172.16.3.3(LDP Id) -> 172.16.2.2, LDP is UP
    Graceful restart: not configured and not enabled
    Non stop routing: not configured and not enabled
    Status TLV support (local/remote)   : enabled/supported
      LDP route watch                   : enabled
      Label/status state machine        : established, LruRru
      Last local dataplane   status rcvd: No fault
      Last BFD dataplane     status rcvd: Not sent
      Last BFD peer monitor  status rcvd: No fault
      Last local AC  circuit status rcvd: No fault
      Last local AC  circuit status sent: No fault
      Last local PW i/f circ status rcvd: No fault
      Last local LDP TLV     status sent: No fault
      Last remote LDP TLV    status rcvd: No fault
      Last remote LDP ADJ    status rcvd: No fault
    MPLS VC labels: local 22, remote 24
    Group ID: local 0, remote 0
    MTU: local 0, remote 0
    Remote interface description:
  Sequencing: receive disabled, send disabled
  Control Word: On (configured: autosense)
  SSO Descriptor: 172.16.2.2/1001, local label: 22
  Dataplane:
    SSM segment/switch IDs: 8194/4096 (used), PWID: 1
  VC statistics:
    transit packet totals: receive 3725129, send 3412605
    transit byte totals:   receive 1014481148, send 900927720
    transit packet drops:  receive 0, seq error 0, send 0
```

```
ASR903-R3# show controllers sonet 0/4/0
SONET 0/4/0 is up.
  Hardware is A900-IMA4OS

 Applique type is Channelized Sonet/SDH
 Clock Source is Line, AUG mapping is AU4.
Medium info:
  Type: SDH, Line Coding: NRZ,
 Regenerator Section:
  LOS = 1          LOF = 1                           BIP(B1) = 0

SONET/SDH Section Tables
  INTERVAL       CV    ES   SES  SEFS
  22:48-22:57     0     0     0     0
  22:33-22:48     0     0     0     0
Total of Data in Current and Previous Intervals
  22:33-22:57     0     0     0     0


Multiplex Section:
  AIS = 1          RDI = 0          REI = 0          BIP(B2) = 0
Active Defects: None
Detected Alarms: None
Asserted/Active Alarms: None
Alarm reporting enabled for: SLOS SLOF SF B1-TCA B2-TCA
BER thresholds:  SF = 10e-3  SD = 10e-6
TCA thresholds:  B1 = 10e-6  B2 = 10e-6
Rx: S1S0 = 00
    K1 = 00,   K2 = 00
    J0 = 01
    RX S1 = 00

Tx: S1S0 = 02
    K1 = 00,   K2 = 00
    J0 = 01

SONET/SDH Line Tables
  INTERVAL       CV    ES   SES   UAS
  22:48-22:57     0     0     0     0
  22:33-22:48     0     0     0     0
Total of Data in Current and Previous Intervals
  22:33-22:57     0     0     0     0


High Order Path:

PATH 1:
  AIS = 1          RDI = 0          REI = 0          BIP(B3) = 0
  LOP = 0          PSE = 0          NSE = 0          NEWPTR = 0
  LOM = 0          PLM = 1          UNEQ = 0
Active Defects: None
Detected Alarms: None
Asserted/Active Alarms: None
```

```
Alarm reporting enabled for: PLOP LOM B3-TCA

TCA threshold:  B3 = 10e-6
Rx: C2 = 02
Tx: C2 = 02

PATH TRACE BUFFER : STABLE
CRC-7: 0x9E OK

  52 75 64 79 2D 34 20 30 2F 34 2F 30 2E 31 20     ASR903-R4 0/4/0.1

SONET/SDH Path Tables
  INTERVAL       CV    ES   SES   UAS
  22:48-22:57     0     0     0     0
  22:33-22:48     0     0     0     0
Total of Data in Current and Previous Intervals
  22:33-22:57     0     0     0     0


 AU-4 1, TUG-3 1, TUG-2 1, E1 1 (C-12 1/1/1/1) is up
 VT Receiver has no alarm.
  No alarms detected.
  Framing is unframed, Clock Source is Internal
  Data in current interval (270 seconds elapsed):
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     0 Unavail Secs, 0 Stuffed Secs
  Data in Interval 1:
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     0 Unavail Secs, 0 Stuffed Secs
  Data in Interval 2:
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     0 Unavail Secs, 0 Stuffed Secs
  Data in Interval 3:
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     0 Unavail Secs, 0 Stuffed Secs
  Data in Interval 4:
     0 Line Code Violations, 0 Path Code Violations
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     0 Unavail Secs, 0 Stuffed Secs
  Total Data (last 4 15 minute intervals):
```

```
     0 Line Code Violations,0 Path Code Violations,
     0 Slip Secs, 0 Fr Loss Secs, 0 Line Err Secs, 0 Degraded Mins,
     0 Errored Secs, 0 Bursty Err Secs, 0 Severely Err Secs
     0 Unavail Secs, 0 Stuffed Secs

 AU-4 1, TUG-3 1, TUG-2 1, E1 2 (C-12 1/1/1/2)
    Not configured.
.........................
```

POS Configuration

Here are the configurations for back-to-back POS configuration between ASR903-R1 and ASR903-R3.

OC-3 SDH POS

ASR903-R1 configuration

```
controller SONET 0/5/0
 framing sdh
 clock source line
 aug mapping au-4
 !
 au-4 1 pos                  ! create POS interface
interface POS0/5/0.1
 ip address 10.10.100.1 255.255.255.248
 encapsulation ppp           ! using encapsulation PPP, default is HDLC
```

ASR903-R3 configuration

```
controller SONET 0/5/0
 framing sdh
 clock source internal
 aug mapping au-4
 !
 au-4 1 pos

interface POS0/5/0.1
 ip address 10.10.100.3 255.255.255.248
 encapsulation ppp
```

Verifying the Configuration

```
ASR903-R3# show interface po0/5/0.1
POS0/5/0.1 is up, line protocol is up
  Hardware is SPA_4xOC3_1xOC12
  Internet address is 10.10.100.3/29
  MTU 4470 bytes, BW 155000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation PPP, LCP Open
```

```
 Open: IPCP, crc 16, loopback not set
  Keepalive set (10 sec)
  Scramble disabled
  Last input 00:00:04, output 00:00:04, output hang never
  Last clearing of "show interface" counters 00:06:07
  Input queue: 0/375/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     89 packets input, 4378 bytes
     Received 0 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     85 packets output, 4040 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
  Non-inverted data

ASR903-R3# ping 10.10.100.1 re 10
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 10.10.100.1, timeout is 2 seconds:
!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 1/1/2 ms
```

OC-3 SONET POS

ASR903-R1 Configuration

```
controller SONET 0/5/0
 framing sonet
 clock source line
 !
 sts-1 1 - 3 pos                ! create POS interface

interface POS0/5/0.1
 ip address 10.10.100.1 255.255.255.248
 encapsulation ppp           ! using encapsulation PPP, default is HDLC
```

ASR903-R3 Configuration

```
controller SONET 0/5/0
 framing sonet
 clock source internal
 !
```

```
 sts-1 1 - 3 pos

interface POS0/5/0.1
 ip address 10.10.100.3 255.255.255.248
 encapsulation ppp
```

Similar configuration can be for OC-12.

OC-12 SDH POS

Optical combo card can also be configured as 1 OC-12 port by using the following command.

card type oc12 0 5

ASR903-R1 Configuration

```
controller SONET 0/5/0
 framing sdh
 clock source line
 aug mapping au-4
 !
 au-4 1 - 4 pos              ! create POS interface
interface POS0/5/0.1
 ip address 10.10.100.1 255.255.255.248
 encapsulation ppp           ! using encapsulation PPP, default is HDLC
```

ASR903-R3 Configuration

```
controller SONET 0/5/0
 framing sdh
 clock source internal
 aug mapping au-4
 !
 au-4 1 - 4 pos
interface POS0/5/0.1
 ip address 10.10.100.3 255.255.255.248
 encapsulation ppp
```

OC-12 SONET POS

Optical combo card can also be configured as 1 OC-12 port by using the following command.

card type oc12 0 5

ASR903-R1 Configuration

```
controller SONET 0/5/0
 framing sonet
 clock source line
 !
 sts-1 1 - 12 pos               ! create POS interface
```

```
interface POS0/5/0.1
 ip address 10.10.100.1 255.255.255.248
 encapsulation ppp              ! using encapsulation PPP, default is HDLC
```

ASR903-R3 Configuration

```
controller SONET 0/5/0
 framing sonet
 clock source internal
 !
 sts-1 1 - 12 pos
interface POS0/5/0.1
 ip address 10.10.100.3 255.255.255.248
 encapsulation ppp
```

Chapter 7 MPLS TE Configuration

MPLS traffic engineering (MPLS TE) is to use a traffic engineered label switched path (TE LSP or tunnel) to forward the traffic across network by taking into account of a set of constraints, network topology and resources available with the objective of make efficient use of the network. MPLS TE is implemented by creating tunnel interface on the head-end router and then established the label switched path from hop to hop.

Here are the well-known engineering attributes which determine the desired LSP path.

- Destination: the sources of the TE LSP, the head-end router needs to know where the LSP terminates.
- Bandwidth: bandwidth requested by the TE LSP.
- Affinities: it can be viewed as coloring scheme for the links and it is desirable sometime to ensure a TE LSP to traverse links of specified colors.
- Preemption: seven levels of priority allow high-priority TE LSP to preempt lower priority TE LSPs in the situation of resource contention such as lack of bandwidth resources on a link.
- Protection by fast reroute: a way to quickly re-route traffic to a pre-signaled back tunnel with tens of milliseconds when there is any link or node failure.
- Optimized Metric: MPLS TE uses either IGP or TE metrics to pick the shortest path for a tunnel that satisfies specified constraints.

ISIS and OSPF IGP extensions are used for path calculation by also putting available bandwidth on each link into consideration. As TE LSP are setup and torn down, the amount of reserved bandwidth varies on each link and is reflected by the IGP. Once the TE LSP path is computed, RSVP TE uses RSVP messages to set up, maintain, signal an error condition, and tear down a TE LSP.

Below is the topology we are going to use to demonstrate MPLS TE configuration.

Figure 22. MPLS-TE

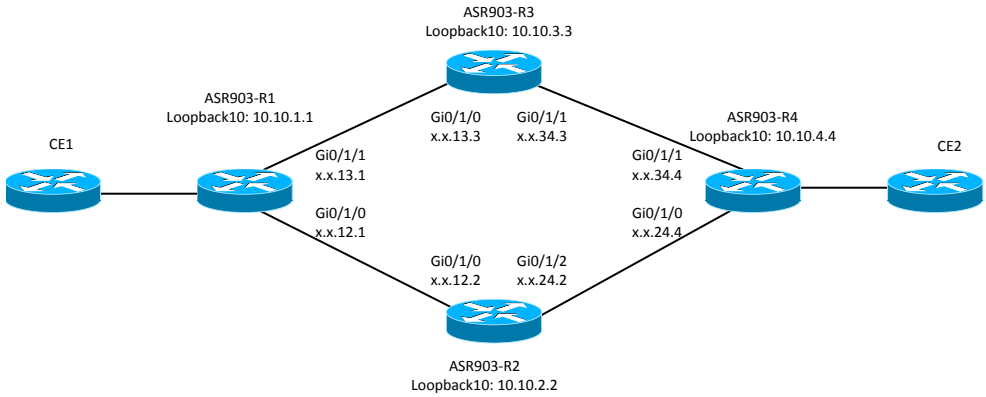

Basic Steps to Configure Traffic Engineering

ASR903-R1 Configuration

```
interface GigabitEthernet0/1/0
 ip address 10.10.12.1 255.255.255.248
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/1/1
 ip address 10.10.13.1 255.255.255.248
 negotiation auto
 mpls ip
!
router ospf 100
 router-id 10.10.1.1
 network 10.10.1.1 0.0.0.0 area 0
 network 10.10.12.1 0.0.0.0 area 0
 network 10.10.13.1 0.0.0.0 area 0
```

Once connectivity and MPLS configuration are verified, the next step is to enable TE and configure IGP for MPLS TE support.

```
mpls traffic-eng tunnels                  ! Enable MPLS TE globally and under the
interfaces.
interface GigabitEthernet0/1/1
 mpls traffic-eng tunnels
interface GigabitEthernet0/1/0
 mpls traffic-eng tunnels

router ospf 100                           ! Configure IGP for TE support
 mpls traffic-eng router-id Loopback10    ! Specify TE router ID
 mpls traffic-eng area 0                  ! Enable TE in area 0
```

Next, specify the bandwidth can be used for traffic engineering. This is the total bandwidth which can be used for TE under the configured interface. The bandwidth can be lower or higher than actual bandwidth and care should be taken to make sure other traffic has enough bandwidth. The following command allocates 40M bandwidth under interface gig0/1/1 and 30M under gi0/1/0 for traffic engineering.

```
interface GigabitEthernet0/1/1
 ip rsvp bandwidth 40000                  ! Reserved bandwidth in kbps
interface GigabitEthernet0/1/0
 ip rsvp bandwidth 30000                  ! Reserved 30M under another interface
```

ASR903-R1 Configuration

```
interface GigabitEthernet0/1/0
 ip address 10.10.12.1 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 ip rsvp bandwidth 30000
!
interface GigabitEthernet0/1/1
 ip address 10.10.13.1 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 ip rsvp bandwidth 40000
!
mpls label protocol ldp
mpls ldp explicit-null
mpls ldp session protection
mpls ldp igp sync holddown 200
mpls ldp discovery targeted-hello accept
mpls traffic-eng tunnels
mpls ldp router-id Loopback10 force
!
router ospf 100
 router-id 10.10.1.1
 timers throttle spf 5 50 200
 timers throttle lsa 5 50 200
 timers lsa arrival 10
 timers pacing flood 5
 network 10.10.1.1 0.0.0.0 area 0
 network 10.10.12.1 0.0.0.0 area 0
 network 10.10.13.1 0.0.0.0 area 0
 mpls traffic-eng router-id Loopback10
 mpls traffic-eng area 0
```

ASR903-R2 Configuration

```
interface GigabitEthernet0/1/2
 ip address 10.10.24.2 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 ip rsvp bandwidth 30000
!
interface GigabitEthernet0/1/0
 ip address 10.10.12.2 255.255.255.248
```

```
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 ip rsvp bandwidth 30000
!
mpls label protocol ldp
mpls ldp explicit-null
mpls ldp session protection
mpls ldp igp sync holddown 200
mpls ldp discovery targeted-hello accept
mpls traffic-eng tunnels
mpls ldp router-id Loopback10 force

router ospf 100
 router-id 10.10.2.2
 network 10.10.0.0 0.0.255.255 area 0
 mpls traffic-eng router-id Loopback10
 mpls traffic-eng area 0
!
```

Similar configurations are configured on ASR903-R3 and ASR903-R4. For simplicity, their configurations are not shown in this document.

Dynamic TE Tunnel

With the steps accomplished above, TE tunnel can be created. Here is an example to create a dynamic TE tunnel from ASR903-R1 to ASR903-R4.

Figure 23. Dynamic TE Tunnel

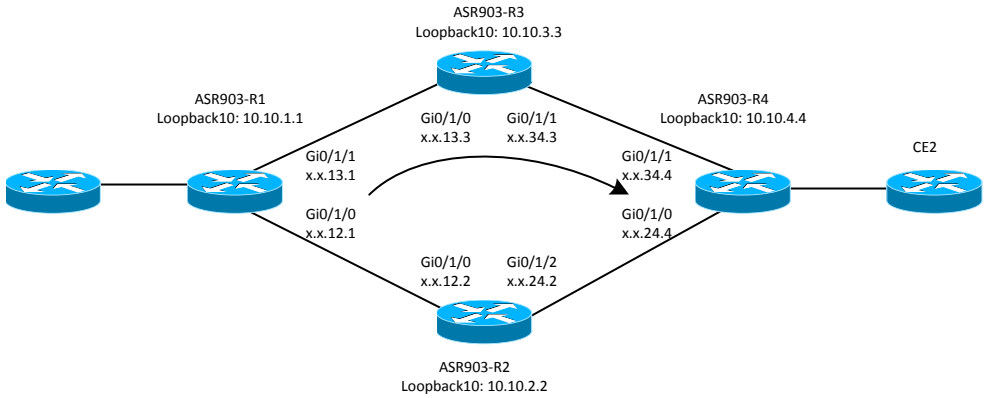

ASR903-R1 Configuration

```
interface Tunnel10
 ip unnumbered Loopback10
 tunnel source Loopback10                    ! specify tunnel source
 tunnel mode mpls traffic-eng
 tunnel destination 10.10.4.4                ! specify tunnel destination
 tunnel mpls traffic-eng priority 5 5           ! priority for the tunnel
 tunnel mpls traffic-eng bandwidth 1000         ! bandwidth requested by this TE
tunnel
 tunnel mpls traffic-eng path-option 3 dynamic      ! dynamic tunnel

ASR903-R1#show mpls traffic-eng tunnels

P2P TUNNELS/LSPs:

Name: ASR903-R1_t10                        (Tunnel10) Destination: 10.10.4.4
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 3, type dynamic (Basis for Setup, path weight 2)

  Config Parameters:
    Bandwidth: 1000     kbps (Global)  Priority: 5  5   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute: enabled  LockDown: disabled Loadshare: 1000 [2000000] bw-based
    auto-bw: disabled
  Active Path Option Parameters:
    State: dynamic path option 3 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled


  InLabel  :  -
  OutLabel : GigabitEthernet0/1/1, 27
  Next Hop : 10.10.13.3
  RSVP Signalling Info:
       Src 10.10.1.1, Dst 10.10.4.4, Tun_Id 10, Tun_Instance 3
    RSVP Path Info:
      My Address: 10.10.13.1
      Explicit Route: 10.10.13.3 10.10.34.3 10.10.34.4 10.10.4.4
      Record   Route:   NONE
      Tspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits

ASR903-R1#show ip route 10.10.4.4
Routing entry for 10.10.4.4/32
  Known via "ospf 100", distance 110, metric 3, type intra area
  Last update from 10.10.13.3 on GigabitEthernet0/1/1, 00:00:11 ago
  Routing Descriptor Blocks:
    10.10.13.3, from 10.10.4.4, 00:00:11 ago, via GigabitEthernet0/1/1
      Route metric is 3, traffic share count is 1
  * 10.10.12.2, from 10.10.4.4, 00:00:11 ago, via GigabitEthernet0/1/0
```

```
      Route metric is 3, traffic share count is 1
```

Now announce tunnel interface for use by IGP.

```
ASR903-R1# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ASR903-R1(config)# interface tun10
ASR903-R1(config-if)# tunnel mpls traffic-eng autoroute announce
ASR903-R1(config-if)#end
ASR903-R1# show ip route 10.10.4.4
Routing entry for 10.10.4.4/32
  Known via "ospf 100", distance 110, metric 3, type intra area
  Last update from 10.10.4.4 on Tunnel10, 00:00:07 ago
  Routing Descriptor Blocks:
  * 10.10.4.4, from 10.10.4.4, 00:00:07 ago, via Tunnel10
      Route metric is 3, traffic share count is 1
```

Explicit Path Configuration

Explicit LSP path can be specified as well for TE tunnel.

ASR903-R1 Configuration

```
ip explicit-path name R1-R2-R4-1 enable
 next-address 10.10.12.2
 next-address 10.10.24.4
interface Tunnel11
 ip unnumbered Loopback10
 tunnel source Loopback10
 tunnel mode mpls traffic-eng
 tunnel destination 10.10.4.4
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 5 5
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 3 explicit name R1-R2-R4-1

ASR903-R1# show mpls traffic-eng tunnels tunnel 11
Name: ASR903-R1_t11                       (Tunnel11) Destination: 10.10.4.4
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 3, type explicit R1-R2-R4-1 (Basis for Setup, path weight 2)

  Config Parameters:
    Bandwidth: 1000     kbps (Global)  Priority: 5  5   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute: enabled  LockDown: disabled Loadshare: 1000 [2000000] bw-based
    auto-bw: disabled
```

```
Active Path Option Parameters:
  State: explicit path option 3 is active
  BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled


InLabel  :  -
OutLabel : GigabitEthernet0/1/0, 26
Next Hop : 10.10.12.2
RSVP Signalling Info:
     Src 10.10.1.1, Dst 10.10.4.4, Tun_Id 11, Tun_Instance 44
  RSVP Path Info:
    My Address: 10.10.12.1
    Explicit Route: 10.10.12.2 10.10.24.2 10.10.24.4 10.10.4.4
    Record   Route:   NONE
    Tspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits
  RSVP Resv Info:
    Record   Route:   NONE
    Fspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits
Shortest Unconstrained Path Info:
  Path Weight: 2 (TE)
  Explicit Route: 10.10.13.1 10.10.13.3 10.10.34.3 10.10.34.4
                  10.10.4.4
```

Explicit Path by using Affinity

Affinity and mask can be used to color the paths. Once those links are colored, an attribute flag can be used to designate which path LSP the tunnel uses.

Figure 24. Explicit Path

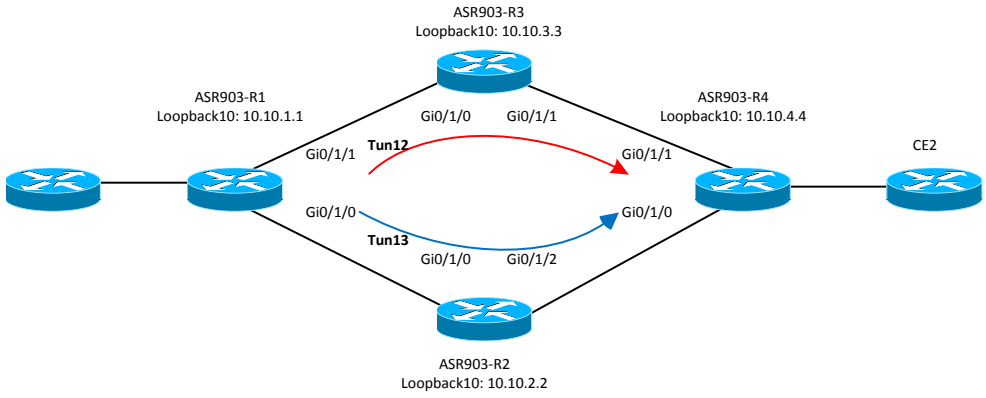

Creating Tunnel12 using Red Path

ASR903-R1 Configuration

```
interface Tunnel12
 ip unnumbered Loopback10
 tunnel source Loopback10
 tunnel mode mpls traffic-eng
 tunnel destination 10.10.4.4
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng bandwidth 500
 tunnel mpls traffic-eng affinity 0x1 mask 0x1
 tunnel mpls traffic-eng path-option 3 dynamic
interface GigabitEthernet0/1/1
 mpls traffic-eng attribute-flags 0x1
```

ASR903-R3 Configuration

```
interface GigabitEthernet0/1/1
 mpls traffic-eng attribute-flags 0x1
!
interface GigabitEthernet0/1/1
 mpls traffic-eng attribute-flags 0x1
```

ASR903-R4 Configuration

```
interface GigabitEthernet0/1/1
 mpls traffic-eng attribute-flags 0x1
```

Verify TE tunnel and Tunnel Path

```
ASR903-R1# show mpls traffic-eng tunnels tunnel 12

Name: ASR903-R1_t12                       (Tunnel12) Destination: 10.10.4.4
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 3, type dynamic (Basis for Setup, path weight 2)

  Config Parameters:
    Bandwidth: 500      kbps (Global)  Priority: 3  3   Affinity: 0x1/0x1
    Metric Type: TE (default)
    AutoRoute: enabled  LockDown: disabled Loadshare: 500 [4000000] bw-based
    auto-bw: disabled
  Active Path Option Parameters:
    State: dynamic path option 3 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled
```

```
  InLabel  :  -
  OutLabel : GigabitEthernet0/1/1, 25
  Next Hop : 10.10.13.3
  RSVP Signalling Info:
       Src 10.10.1.1, Dst 10.10.4.4, Tun_Id 12, Tun_Instance 14
    RSVP Path Info:
      My Address: 10.10.13.1
      Explicit Route: 10.10.13.3 10.10.34.3 10.10.34.4 10.10.4.4
      Record   Route:   NONE
      Tspec: ave rate=500 kbits, burst=1000 bytes, peak rate=500 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=500 kbits, burst=1000 bytes, peak rate=500 kbits
  Shortest Unconstrained Path Info:
    Path Weight: 2 (TE)
    Explicit Route: 10.10.13.1 10.10.13.3 10.10.34.3 10.10.34.4
                    10.10.4.4
  History:
    Tunnel:
      Time since created: 11 hours, 23 minutes
      Time since path change: 11 hours, 15 minutes
      Number of LSP IDs (Tun_Instances) used: 14
    Current LSP: [ID: 14]
      Uptime: 11 hours, 15 minutes
    Prior LSP: [ID: 1]
      ID: path option unknown
      Removal Trigger: configuration changed
```

Creating Tunnel13 using Blue Path

ASR903-R1 Configuration

```
interface Tunnel13
 ip unnumbered Loopback10
 tunnel source Loopback10
 tunnel mode mpls traffic-eng
 tunnel destination 10.10.4.4
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng bandwidth 500
 tunnel mpls traffic-eng affinity 0x2 mask 0x2
 tunnel mpls traffic-eng path-option 3 dynamic
interface GigabitEthernet0/1/0
 mpls traffic-eng attribute-flags 0x2
ASR903-R2 Configuration
interface GigabitEthernet0/1/0
 mpls traffic-eng attribute-flags 0x2
!
interface GigabitEthernet0/1/2
 mpls traffic-eng attribute-flags 0x2
```

```
ASR903-R4 Configuration
interface GigabitEthernet0/1/0
 mpls traffic-eng attribute-flags 0x2
```

Traffic Engineering Fast Reroute

In the previous example, tunnel12 goes through ASR903-R3 and tunnel13 goes through ASR903-R2. Tunnel 13 is to be set up as backup path for fast reroute.

ASR903-R1 Configuration

```
interface GigabitEthernet0/1/1
 ip address 10.10.13.1 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 mpls traffic-eng backup-path Tunnel13
 mpls traffic-eng attribute-flags 0x1
 ip rsvp bandwidth 40000
!
interface Tunnel12
 ip unnumbered Loopback10
 tunnel source Loopback10
 tunnel mode mpls traffic-eng
 tunnel destination 10.10.4.4
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng bandwidth 500
 tunnel mpls traffic-eng affinity 0x1 mask 0x1
 tunnel mpls traffic-eng path-option 3 dynamic
 tunnel mpls traffic-eng fast-reroute
!
interface Tunnel13
 ip unnumbered Loopback10
 tunnel source Loopback10
 tunnel mode mpls traffic-eng
 tunnel destination 10.10.4.4
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng bandwidth 500
 tunnel mpls traffic-eng affinity 0x2 mask 0x2
 tunnel mpls traffic-eng path-option 3 dynamic
```

Verifying Fast Reroute Configuration

```
ASR903-R1# show mpls traffic-eng fast-reroute database detail
FRR Database Summary:
  Protected interfaces    : 1
  Protected LSPs/Sub-LSPs : 1
  Backup tunnels           : 1
  Active interfaces       : 0
  FRR Active tunnels      : 0

P2P LSPs:
 Tun ID: 12, LSP ID: 26, Source: 10.10.1.1
 Destination: 10.10.4.4
  State        : ready
  InLabel      : Tunnel Head
  OutLabel     : Gi0/1/1:26
  FRR OutLabel : Tu13:implicit-null

ASR903-R1# show mpls traffic-eng tunnels tunnel 12 protection
ASR903-R1_t12
  LSP Head, Tunnel12, Admin: up, Oper: up
  Src 10.10.1.1, Dest 10.10.4.4, Instance 26
  Fast Reroute Protection: Requested
    Outbound: FRR Ready
      Backup Tu13 to LSP nnhop
        Tu13: out i/f: Gi0/1/0, label: 17
      LSP signalling info:
        Original: out i/f: Gi0/1/1, label: 26, nhop: 10.10.13.3
                  nnhop: 10.10.4.4, nnhop rtr id: 10.10.4.4
        With FRR: out i/f: Tu13, label: implicit-null
      LSP bw: 500 kbps, Backup level: any-unlim, type: any pool
  Path Protection: None
After the interface to ASR903-R1 on the ASR903-R3 router is shutdown, fast reroute
became active.

ASR903-R1# show mpls traffic-eng tunnels tunnel 12 protection
ASR903-R1_t12
  LSP Head, Tunnel12, Admin: up, Oper: up
  Src 10.10.1.1, Dest 10.10.4.4, Instance 26
  Fast Reroute Protection: Requested
    Outbound: FRR Active
      Backup Tu13 to LSP nnhop
        Tu13: out i/f: Gi0/1/0, label: 17
      LSP signalling info:
        Original: out i/f: Gi0/1/1, label: 26, nhop: 10.10.4.4
                  nnhop: 10.10.4.4, nnhop rtr id: 10.10.4.4
        With FRR: out i/f: Tu13, label: implicit-null
      LSP bw: 500 kbps, Backup level: any-unlim, type: any pool
  Path Protection: None
```

BFD Triggered TE Fast Reroute

ASR903-R1 Configuration

```
interface GigabitEthernet0/1/1
 ip address 10.10.34.3 255.255.255.248
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 mpls traffic-eng attribute-flags 0x1
 bfd interval 50 min_rx 50 multiplier 3
 ip rsvp bandwidth 40000
 ip rsvp signalling hello bfd
```

ASR903-R3 Configuration

```
interface GigabitEthernet0/1/0
 ip address 10.10.13.3 255.255.255.248
 ip ospf dead-interval 3
 ip ospf hello-interval 1
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 mpls traffic-eng attribute-flags 0x1
 bfd interval 50 min_rx 50 multiplier 3
 ip rsvp bandwidth 40000
 ip rsvp signalling hello bfd
!
interface GigabitEthernet0/1/1
 ip address 10.10.34.3 255.255.255.248
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 mpls traffic-eng attribute-flags 0x1
 bfd interval 50 min_rx 50 multiplier 3
 ip rsvp bandwidth 40000
 ip rsvp signalling hello bfd
```

ASR903-R4 Configuration

```
interface GigabitEthernet0/1/1
 ip address 10.10.34.4 255.255.255.248
 negotiation auto
 mpls ip
 mpls traffic-eng tunnels
 mpls traffic-eng attribute-flags 0x1
 bfd interval 50 min_rx 50 multiplier 3
 ip rsvp bandwidth 100000
 ip rsvp signalling hello bfd
```

Verifying BFD Triggered Fast Reroute

```
ASR903-R1# show bfd neighbors details

IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
10.10.13.3                             1/1           Up        Up        Gi0/1/1
Session state is UP and using echo function with 50 ms interval.
Session Host: Hardware - session negotiated with platform adjusted timer values.
              MinTxInt - configured: 50000       adjusted: 1000000
OurAddr: 10.10.13.1
Handle: 1
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 1000000, MinRxInt: 1000000, Multiplier: 3
Received MinRxInt: 1000000, Received Multiplier: 3
Holddown (hits): 0(0), Hello (hits): 1000(0)
Rx Count: 53031
Tx Count: 53023
Elapsed time watermarks: 0 0 (last: 0)
Registered protocols: FRR CEF
Uptime: 00:32:44
Last packet: Version: 1                  - Diagnostic: 0
             State bit: Up               - Demand bit: 0
             Poll bit: 0                 - Final bit: 0
             C bit: 1
             Multiplier: 3               - Length: 24
             My Discr.: 1                - Your Discr.: 1
             Min tx interval: 1000000    - Min rx interval: 1000000
             Min Echo interval: 50000
```

Chapter 8 LFA & Remote LFA (rLFA)

When a link or a router fails, distributed routing algorithms compute new routes that take into account the failure. The time taken for computation is called routing transition. Until the transition is complete and all routers are converged on a common view of the network, the connectivity between the source and destination pairs is interrupted. IPv4 Loop-Free Alternate (LFA) Fast Reroute (FRR) feature can reduce the routing transition time to less than 50 milliseconds using a pre-computed alternate next hop. When a router is notified of a link failure, the router immediately switches over to the repair path to reduce traffic loss.

IPv4 LFA FRR supports the pre-computation of repair paths. The repair path computation can be done by IS-IS or OSPF routing protocol, and the resulting repair paths are sent to the RIB.

Figure 25. LFA

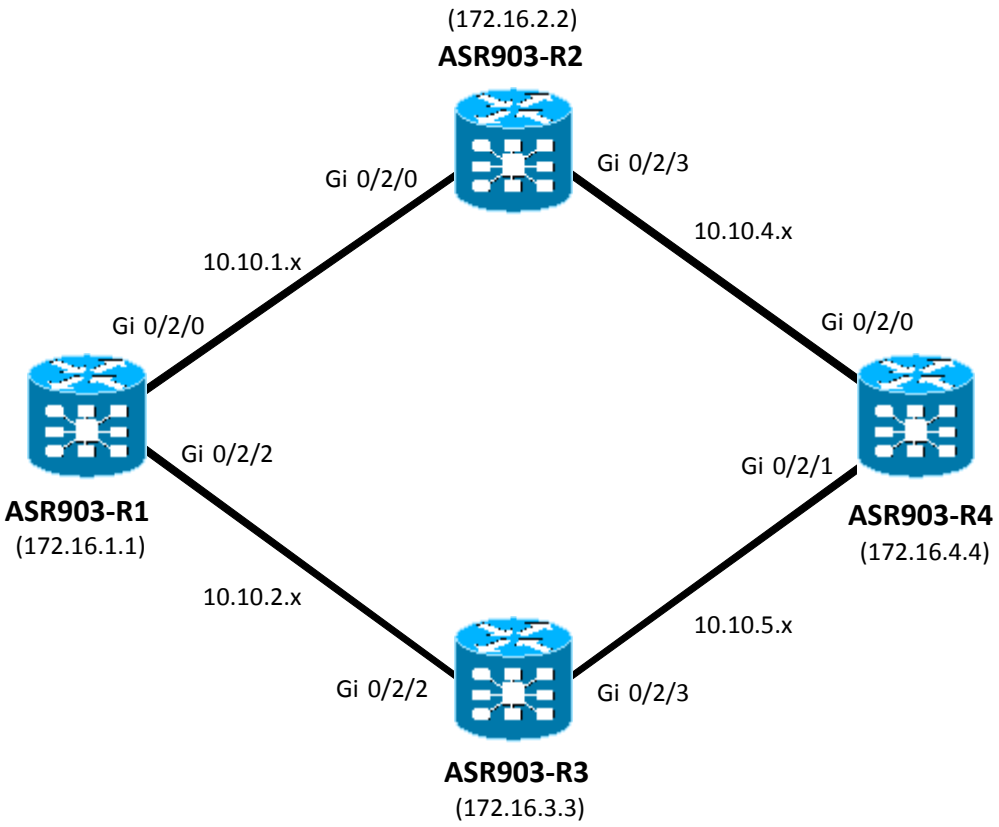

Configuration for LFA (with OSPF)

ASR903-R1 Configuration

```
router ospf 1
 router-id 172.16.1.1
 fast-reroute per-prefix enable prefix-priority low
 network 172.16.1.1 0.0.0.0 area 0
 network 10.10.1.0 0.0.0.255 area 0
 network 10.10.2.0 0.0.0.255 area 0
```

Configuration for LFA (with ISIS)

ASR903-R1 Configuration

```
interface GigabitEthernet0/2/0
 ip address 10.10.1.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
interface GigabitEthernet0/2/2
 ip address 10.10.2.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
router isis
 net 49.0001.0000.0000.0001.00
 fast-reroute per-prefix level-1 all
 fast-reroute per-prefix level-2 all
passive-interface Loopback0
```

Verifying LFA

```
ASR903-R1# show ip route 172.16.4.4
Routing entry for 172.16.4.4/32
  Known via "ospf 1", distance 110, metric 3, type intra area
  Last update from 10.10.1.2 on GigabitEthernet0/2/0, 00:00:06 ago
  Routing Descriptor Blocks:
    10.10.2.2, from 172.16.4.4, 00:00:06 ago, via GigabitEthernet0/2/2
      Route metric is 3, traffic share count is 1
      Repair Path: 10.10.1.2, via GigabitEthernet0/2/0
  * 10.10.1.2, from 172.16.4.4, 00:00:06 ago, via GigabitEthernet0/2/0
      Route metric is 3, traffic share count is 1
      Repair Path: 10.10.2.2, via GigabitEthernet0/2/2
```

```
ASR903-R1#show ip cef 172.16.4.4 de
172.16.4.4/32, epoch 2, per-destination sharing
  local label info: global/30
  nexthop 10.10.1.2 GigabitEthernet0/2/0 label [16|18]
    repair: attached-nexthop 10.10.2.2 GigabitEthernet0/2/2
  nexthop 10.10.2.2 GigabitEthernet0/2/2 label [18|16]
repair: attached-nexthop 10.10.1.2 GigabitEthernet0/2/0
```

Configuration for rLFA (with OSPF)

ASR903-R1 Configuration

```
mpls label protocol ldp
mpls ldp explicit-null
mpls ldp session protection
mpls ldp router-id Loopback0 force
!
router ospf 1
 router-id 172.16.1.1
 fast-reroute per-prefix enable prefix-priority low
 fast-reroute per-prefix remote-lfa tunnel mpls-ldp
 network 172.16.1.1 0.0.0.0 area 0
 network 10.10.1.0 0.0.0.255 area 0
 network 10.10.2.0 0.0.0.255 area 0
```

ASR903-R4 Configuration

```
!
mpls ldp discovery targeted-hello accept
mpls ldp router-id Loopback0 force
!
router ospf 1
 router-id 172.16.4.4
 network 172.16.4.4 0.0.0.0 area 0
 network 10.10.4.0 0.0.0.255 area 0
 network 10.10.5.0 0.0.0.255 area 0
```

Configuration for rLFA (with ISIS)

ASR903-R1 Configuration

```
!
mpls label protocol ldp
mpls ldp explicit-null
mpls ldp session protection
mpls ldp router-id Loopback0 force
```

```
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.255
!
interface GigabitEthernet0/2/0
 ip address 10.10.1.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
interface GigabitEthernet0/2/2
 ip address 10.10.2.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
router isis
 net 49.0001.0000.0000.0001.00
 fast-reroute per-prefix level-1 all
 fast-reroute per-prefix level-2 all
 fast-reroute remote-lfa level-1 mpls-ldp
 fast-reroute remote-lfa level-2 mpls-ldp
 passive-interface Loopback0
!
```

ASR903-R2 Configuration

```
!
mpls ldp explicit-null
mpls ldp router-id Loopback0 force
!
interface GigabitEthernet0/2/0
 ip address 10.10.1.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
interface GigabitEthernet0/2/3
 ip address 10.10.4.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
router isis
 net 49.0001.0000.0000.0002.00
 passive-interface Loopback0
!
```

ASR903-R3 Configuration

```
mpls ldp explicit-null
mpls ldp router-id Loopback0 force
!
interface GigabitEthernet0/2/2
 ip address 10.10.2.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
interface GigabitEthernet0/2/3
 ip address 10.10.5.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
router isis
 net 49.0001.0000.0000.0003.00
 passive-interface Loopback0
!
```

ASR903-R4 Configuration

```
!
mpls ldp explicit-null
mpls ldp session protection
mpls ldp discovery targeted-hello accept
mpls ldp router-id Loopback0 force
!
interface GigabitEthernet0/2/0
 ip address 10.10.4.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
interface GigabitEthernet0/2/1
 ip address 10.10.5.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis network point-to-point
!
router isis
 net 49.0001.0000.0000.0004.00
 passive-interface Loopback0
```

Verifying rLFA

```
ASR903-R1# show ip ospf fast-reroute remote-lfa tunnels

OSPF Router with ID (172.16.1.1) (Process ID 1)
Area with ID (0)

Base Topology (MTID 0)
Interface MPLS-Remote-Lfa1
  Tunnel type: MPLS-LDP
  Tailend router ID: 172.16.4.4
  Termination IP address: 172.16.4.4
  Outgoing interface: GigabitEthernet0/2/0
  First hop gateway: 10.10.1.2
  Tunnel metric: 2
  Protects:
    10.10.2.2 GigabitEthernet0/2/2, total metric 3

Interface MPLS-Remote-Lfa2
  Tunnel type: MPLS-LDP
  Tailend router ID: 172.16.4.4
  Termination IP address: 172.16.4.4
  Outgoing interface: GigabitEthernet0/2/2
  First hop gateway: 10.10.2.2
  Tunnel metric: 2
  Protects:
    10.10.1.2 GigabitEthernet0/2/0, total metric 3

ASR903-R1# show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
17         explicit-n 10.10.5.0/24     0             Gi0/2/2    10.10.2.2
21         explicit-n 172.16.3.3/32    0             Gi0/2/2    10.10.2.2
25         explicit-n 172.16.2.2/32    0             Gi0/2/0    10.10.1.2
27         explicit-n 10.10.4.0/24     0             Gi0/2/0    10.10.1.2
30         16         172.16.4.4/32    144           Gi0/2/0    10.10.1.2
           18         172.16.4.4/32    0             Gi0/2/2    10.10.2.2

ASR903-R1# show mpls ldp bindings
  lib entry: 172.16.1.1/32, rev 2
        local binding:  label: imp-null
        remote binding: lsr: 172.16.2.2:0, label: 19
        remote binding: lsr: 172.16.3.3:0, label: 28
        remote binding: lsr: 172.16.4.4:0, label: 17
  lib entry: 172.16.2.2/32, rev 195
        local binding:  label: 25
        remote binding: lsr: 172.16.4.4:0, label: 29
        remote binding: lsr: 172.16.3.3:0, label: 31
        remote binding: lsr: 172.16.2.2:0, label: exp-null
  lib entry: 172.16.3.3/32, rev 262
        local binding:  label: 21
        remote binding: lsr: 172.16.2.2:0, label: 17
        remote binding: lsr: 172.16.4.4:0, label: 16
```

```
        remote binding: lsr: 172.16.3.3:0, label: exp-null
  lib entry: 172.16.4.4/32, rev 264
        local binding:  label: 30
        remote binding: lsr: 172.16.2.2:0, label: 16
        remote binding: lsr: 172.16.3.3:0, label: 18
        remote binding: lsr: 172.16.4.4:0, label: exp-null
  lib entry: 10.10.1.0/24, rev 10
        local binding:  label: imp-null
        remote binding: lsr: 172.16.4.4:0, label: 20
        remote binding: lsr: 172.16.3.3:0, label: 23
        remote binding: lsr: 172.16.2.2:0, label: exp-null
  lib entry: 10.10.2.0/24, rev 242
        local binding:  label: imp-null
        remote binding: lsr: 172.16.2.2:0, label: 22
        remote binding: lsr: 172.16.4.4:0, label: 26
        remote binding: lsr: 172.16.3.3:0, label: exp-null
  lib entry: 10.10.4.0/24, rev 257
        local binding:  label: 27
        remote binding: lsr: 172.16.4.4:0, label: exp-null
        remote binding: lsr: 172.16.3.3:0, label: 27
        remote binding: lsr: 172.16.2.2:0, label: exp-null
  lib entry: 10.10.5.0/24, rev 263
        local binding:  label: 17
        remote binding: lsr: 172.16.2.2:0, label: 23
        remote binding: lsr: 172.16.4.4:0, label: exp-null
        remote binding: lsr: 172.16.3.3:0, label: exp-null

ASR903-R1# show ip route 172.16.3.3
Routing entry for 172.16.3.3/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.10.2.2 on GigabitEthernet0/2/2, 00:24:33 ago
  Routing Descriptor Blocks:
  * 10.10.2.2, from 172.16.3.3, 00:24:33 ago, via GigabitEthernet0/2/2
      Route metric is 2, traffic share count is 1
      Repair Path: 172.16.4.4, via MPLS-Remote-Lfa1

ASR903-R1# show ip route 172.16.2.2
Routing entry for 172.16.2.2/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.10.1.2 on GigabitEthernet0/2/0, 00:30:30 ago
  Routing Descriptor Blocks:
  * 10.10.1.2, from 172.16.2.2, 00:30:30 ago, via GigabitEthernet0/2/0
      Route metric is 2, traffic share count is 1
      Repair Path: 172.16.4.4, via MPLS-Remote-Lfa2
```

Chapter 9 BGP PIC & RFC3107

The BGP PIC for IP and MPLS-VPN feature improves BGP convergence after a network failure. This convergence is applicable to both core and edge failures and can be used in both IP and MPLS networks. The BGP PIC for IP and MPLS-VPN feature creates and stores a backup/alternate path in the routing information base (RIB), forwarding information base (FIB), and Cisco Express Forwarding (CEF) so that when a failure is detected, the backup/alternate path can immediately take over, thus enabling fast failover.

RFC 3107 specifies how to add label information to multiprotocol BGP address families using a subsequent address family identifier (SAFI). The Cisco IOS implementation of MPLS uses RFC 3107 to provide support for sending IPv4 routes with a label. VPNv4 routes implicitly have a label associated with each route.

Figure 26. BGP PIC for IPand MPLS-VPN

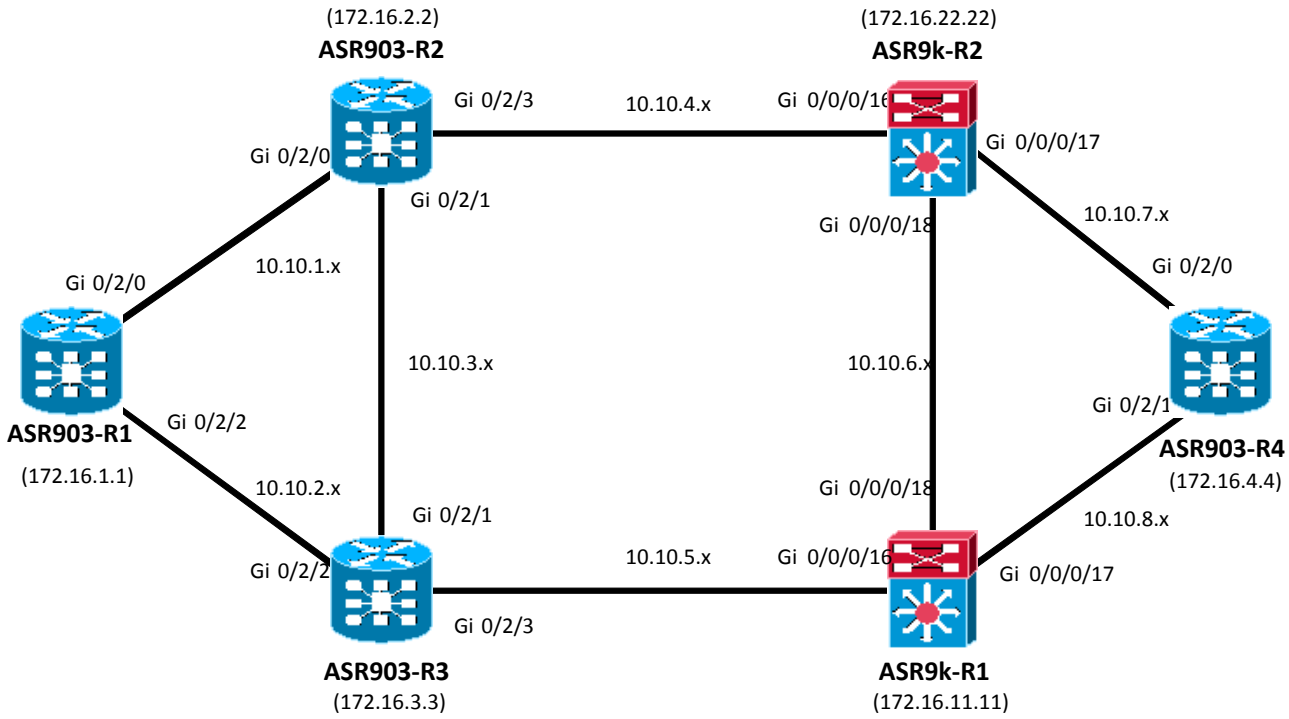

Configuration for BGP PIC (with OSPF)

ASR903-R1 Configuration

```
interface Loopback0
 ip address 172.16.1.1 255.255.255.255
!
interface GigabitEthernet0/2/0
 ip address 10.10.1.1 255.255.255.0
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/2
 ip address 10.10.2.1 255.255.255.0
 negotiation auto
 mpls ip
!
router ospf 1
 router-id 172.16.1.1
 network 172.16.1.1 0.0.0.0 area 10
 network 10.10.1.0 0.0.0.255 area 10
 network 10.10.2.0 0.0.0.255 area 10
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 172.16.2.2 remote-as 100
 neighbor 172.16.2.2 update-source Loopback0
 neighbor 172.16.3.3 remote-as 100
 neighbor 172.16.3.3 update-source Loopback0
 !
 address-family ipv4
  bgp additional-paths install
  network 172.16.1.1 mask 255.255.255.255
  neighbor 172.16.2.2 activate
  neighbor 172.16.2.2 send-label
  neighbor 172.16.3.3 activate
  neighbor 172.16.3.3 send-label
 exit-address-family
!
mpls ldp router-id Loopback0
```

ASR903-R2 Configuration

```
interface Loopback0
 ip address 172.16.2.2 255.255.255.255
!
interface GigabitEthernet0/2/0
 ip address 10.10.1.2 255.255.255.0
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/1
 ip address 10.10.3.1 255.255.255.0
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/3
 ip address 10.10.4.1 255.255.255.0
 negotiation auto
 mpls ip
!
router ospf 1
 router-id 172.16.2.2
 area 10 filter-list prefix DENY in
 area 10 filter-list prefix PERMIT out
 network 172.16.2.2 0.0.0.0 area 10
 network 10.10.1.0 0.0.0.255 area 10
 network 10.10.3.0 0.0.0.255 area 10
 network 10.10.4.0 0.0.0.255 area 0
!
router bgp 100
 bgp cluster-id 172.16.2.2
 bgp log-neighbor-changes
 neighbor 172.16.1.1 remote-as 100
 neighbor 172.16.1.1 update-source Loopback0
 neighbor 172.16.3.3 remote-as 100
 neighbor 172.16.3.3 update-source Loopback0
 neighbor 172.16.22.22 remote-as 100
 neighbor 172.16.22.22 update-source Loopback0
 !
 address-family ipv4
  neighbor 172.16.1.1 activate
  neighbor 172.16.1.1 route-reflector-client
  neighbor 172.16.1.1 next-hop-self all
  neighbor 172.16.1.1 send-label
  neighbor 172.16.3.3 activate
  neighbor 172.16.22.22 activate
  neighbor 172.16.22.22 next-hop-self all
  neighbor 172.16.22.22 send-label
 exit-address-family
!
```

```
!
ip prefix-list DENY seq 5 deny 0.0.0.0/0 ge 1
!
ip prefix-list PERMIT seq 5 permit 172.16.2.2/32
!
mpls ldp router-id Loopback0
```

ASR903-R3 Configuration

```
interface Loopback0
 ip address 172.16.3.3 255.255.255.255
!
interface GigabitEthernet0/2/1
 ip address 10.10.3.2 255.255.255.0
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/2
 ip address 10.10.2.2 255.255.255.0
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/3
 ip address 10.10.5.1 255.255.255.0
 negotiation auto
 mpls ip
!
router ospf 1
 router-id 172.16.3.3
 area 10 filter-list prefix DENY in
 area 10 filter-list prefix PERMIT out
 network 172.16.3.3 0.0.0.0 area 10
 network 10.10.2.0 0.0.0.255 area 10
 network 10.10.3.0 0.0.0.255 area 10
 network 10.10.5.0 0.0.0.255 area 0
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 172.16.1.1 remote-as 100
 neighbor 172.16.1.1 update-source Loopback0
 neighbor 172.16.2.2 remote-as 100
 neighbor 172.16.2.2 update-source Loopback0
 neighbor 172.16.11.11 remote-as 100
 neighbor 172.16.11.11 update-source Loopback0
 !
 address-family ipv4
  neighbor 172.16.1.1 activate
  neighbor 172.16.1.1 route-reflector-client
  neighbor 172.16.1.1 next-hop-self all
```

```
  neighbor 172.16.1.1 send-label
  neighbor 172.16.2.2 activate
  neighbor 172.16.11.11 activate
  neighbor 172.16.11.11 next-hop-self all
  neighbor 172.16.11.11 send-label
 exit-address-family
!
ip prefix-list DENY seq 5 deny 0.0.0.0/0 ge 1
!
ip prefix-list PERMIT seq 5 permit 172.16.3.3/32
!
mpls ldp router-id Loopback0
```

ASR903-R4 Configuration

```
interface Loopback0
 ip address 172.16.4.4 255.255.255.255
!
interface GigabitEthernet0/2/0
 ip address 10.10.7.2 255.255.255.0
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/1
 ip address 10.10.8.2 255.255.255.0
 negotiation auto
 mpls ip
!
router ospf 1
 router-id 172.16.4.4
 network 172.16.4.4 0.0.0.0 area 20
 network 10.10.7.0 0.0.0.255 area 20
 network 10.10.8.0 0.0.0.255 area 20
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 172.16.11.11 remote-as 100
 neighbor 172.16.11.11 update-source Loopback0
 neighbor 172.16.22.22 remote-as 100
 neighbor 172.16.22.22 update-source Loopback0
 !
 address-family ipv4
  bgp additional-paths install
  network 172.16.4.4 mask 255.255.255.255
  network 10.10.7.0 mask 255.255.255.0
  neighbor 172.16.11.11 activate
  neighbor 172.16.11.11 send-label
  neighbor 172.16.22.22 activate
  neighbor 172.16.22.22 send-label
 exit-address-family
!
mpls ldp router-id Loopback0
!
```

ASR9k-1 Configuration

```
!
interface GigabitEthernet0/0/0/16
 ipv4 address 10.10.5.2 255.255.255.0
 negotiation auto
!
interface GigabitEthernet0/0/0/17
 ipv4 address 10.10.8.1 255.255.255.0
 negotiation auto
!
interface GigabitEthernet0/0/0/18
 ipv4 address 10.10.6.2 255.255.255.0
 negotiation auto
!
router ospf 10
 router-id 172.16.11.11
 area 0
  interface GigabitEthernet0/0/0/16
  !
 !
 area 20
  route-policy DENY in
  route-policy PERMIT out
  interface Loopback10
  !
  interface GigabitEthernet0/0/0/17
  !
  interface GigabitEthernet0/0/0/18
  !
 !
!
router bgp 100
 ibgp policy out enforce-modifications
 address-family ipv4 unicast
  advertise best-external
  allocate-label all
 !
 neighbor 172.16.3.3
  remote-as 100
  update-source Loopback10
  address-family ipv4 unicast
   next-hop-self
  !
  address-family ipv4 labeled-unicast
   next-hop-self
  !
 !
 neighbor 172.16.4.4
```

```
  remote-as 100
  update-source Loopback10
  address-family ipv4 unicast
   route-reflector-client
   next-hop-self
  !
  address-family ipv4 labeled-unicast
   route-reflector-client
   next-hop-self
  !
 !
 neighbor 172.16.22.22
  remote-as 100
  update-source Loopback10
  address-family ipv4 unicast
  !
 !
!
route-policy DENY
  drop
end-policy
!
route-policy PERMIT
  if destination in (172.16.11.11/32) then
    pass
  else
    drop
  endif
end-policy
!
```

ASR9k-2 Configuration

```
!
interface GigabitEthernet0/0/0/16
 ipv4 address 10.10.4.2 255.255.255.0
 negotiation auto
!
interface GigabitEthernet0/0/0/17
 ipv4 address 10.10.7.1 255.255.255.0
 negotiation auto
!
interface GigabitEthernet0/0/0/18
 ipv4 address 10.10.6.1 255.255.255.0
 negotiation auto
!
router ospf 10
 router-id 172.16.22.22
 area 0
  interface GigabitEthernet0/0/0/16
  !
 !
 area 20
```

```
  route-policy DENY in
  route-policy PERMIT out
  interface Loopback10
  !
  interface GigabitEthernet0/0/0/17
  !
  interface GigabitEthernet0/0/0/18
  !
 !
!
router bgp 100
 ibgp policy out enforce-modifications
 address-family ipv4 unicast
  advertise best-external
  allocate-label all
 !
 neighbor 172.16.2.2
  remote-as 100
  update-source Loopback10
  address-family ipv4 unicast
   next-hop-self
  !
  address-family ipv4 labeled-unicast
   next-hop-self
  !
 !
 neighbor 172.16.4.4
  remote-as 100
  update-source Loopback10
  address-family ipv4 unicast
   route-reflector-client
   next-hop-self
  !
  address-family ipv4 labeled-unicast
   route-reflector-client
   next-hop-self
  !
 !
 neighbor 172.16.11.11
  remote-as 100
  update-source Loopback10
  address-family ipv4 unicast
  !
 !
!
route-policy DENY
  drop
end-policy
!
```

```
route-policy PERMIT
  if destination in (172.16.22.22/32) then
    pass
  else
    drop
  endif
end-policy
!
```

Verifying BGP PIC & RFC3107

```
ASR903-R1# show ip bgp
BGP table version is 270, local router ID is 172.16.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  172.16.1.1/32    0.0.0.0                  0         32768   i
 *bi 172.16.4.4/32    172.16.3.3               0    100      0   i
 *>i                  172.16.2.2               0    100      0   i
 *bi 10.10.7.0/24     172.16.3.3               0    100      0   i
 *>i                  172.16.2.2               0    100      0   i

ASR903-R1# show ip bgp 172.16.4.4
BGP routing table entry for 172.16.4.4/32, version 269
Paths: (2 available, best #2, table default)
  Additional-path-install
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.3.3 (metric 2) from 172.16.3.3 (172.16.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
      Originator: 172.16.4.4, Cluster list: 172.16.3.3, 203.0.113.101
      mpls labels in/out nolabel/28
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 3
  Local
    172.16.2.2 (metric 2) from 172.16.2.2 (172.16.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 172.16.4.4, Cluster list: 172.16.2.2, 203.0.113.102
      mpls labels in/out nolabel/28
      rx pathid: 0, tx pathid: 0x0

ASR903-R1# show ip bgp label
   Network          Next Hop      In label/Out label
   172.16.1.1/32     0.0.0.0         imp-null/nolabel
   172.16.4.4/32    172.16.3.3         nolabel/28
                    172.16.2.2         nolabel/28
   10.10.7.0/24     172.16.3.3         nolabel/30
```

```
                    172.16.2.2         nolabel/30
ASR903-R1# ping 172.16.4.4 source 172.16.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.4, timeout is 2 seconds:
Packet sent with a source address of 172.16.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms


ASR903-R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
C        172.16.1.1 is directly connected, Loopback0
      2.0.0.0/32 is subnetted, 1 subnets
O        172.16.2.2 [110/2] via 10.10.1.2, 1w0d, GigabitEthernet0/2/0
      3.0.0.0/32 is subnetted, 1 subnets
O        172.16.3.3 [110/2] via 10.10.2.2, 2d21h, GigabitEthernet0/2/2
      4.0.0.0/32 is subnetted, 1 subnets
B        172.16.4.4 [200/0] via 172.16.2.2, 1d20h
      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
C        10.10.1.0/24 is directly connected, GigabitEthernet0/2/0
L        10.10.1.1/32 is directly connected, GigabitEthernet0/2/0
C        10.10.2.0/24 is directly connected, GigabitEthernet0/2/2
L        10.10.2.1/32 is directly connected, GigabitEthernet0/2/2
O        10.10.3.0/24 [110/2] via 10.10.2.2, 2d21h, GigabitEthernet0/2/2
                      [110/2] via 10.10.1.2, 2d21h, GigabitEthernet0/2/0
B        10.10.7.0/24 [200/0] via 172.16.2.2, 1d20h


ASR903-R1# show ip route 172.16.4.4
Routing entry for 172.16.4.4/32
  Known via "bgp 100", distance 200, metric 0, type internal
  Last update from 172.16.2.2 1d20h ago
  Routing Descriptor Blocks:
  * 172.16.2.2, from 172.16.2.2, 1d20h ago
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: 28
```

```
ASR903-R1# show ip cef 172.16.4.4 detail
172.16.4.4/32, epoch 2, flags rib defined all labels
  recursive via 172.16.2.2 label 28
    nexthop 10.10.1.2 GigabitEthernet0/2/0
  recursive via 172.16.3.3 label 28, repair
    nexthop 10.10.2.2 GigabitEthernet0/2/2

ASR903-R4# show ip bgp
BGP table version is 92, local router ID is 172.16.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 172.16.1.1/32    172.16.11.11             0    100      0 i
 *bi                  172.16.22.22             0    100      0 i
 *>  172.16.4.4/32     0.0.0.0                  0         32768 i
 *>  10.10.7.0/24      0.0.0.0                  0         32768 i

ASR903-R4# show ip bgp 172.16.1.1
BGP routing table entry for 172.16.1.1/32, version 92
Paths: (2 available, best #1, table default)
  Additional-path-install
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.11.11 (metric 2) from 172.16.11.11 (203.0.113.101)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 172.16.1.1, Cluster list: 203.0.113.101, 172.16.3.3
      mpls labels in/out nolabel/17248
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  Local
    172.16.22.22 (metric 2) from 172.16.22.22 (203.0.113.102)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
      Originator: 172.16.1.1, Cluster list: 203.0.113.102, 172.16.2.2
      mpls labels in/out nolabel/16001
      rx pathid: 0, tx pathid: 0

ASR903-R4# show ip bgp labels
   Network          Next Hop      In label/Out label
   172.16.1.1/32    172.16.11.11    nolabel/17248
                    172.16.22.22    nolabel/16001
   172.16.4.4/32    0.0.0.0         imp-null/nolabel
   10.10.7.0/24     0.0.0.0         imp-null/nolabel

ASR903-R4# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
B        172.16.1.1 [200/0] via 172.16.11.11, 1d20h
      4.0.0.0/32 is subnetted, 1 subnets
C        172.16.4.4 is directly connected, Loopback0
      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O        10.10.6.0/24 [110/2] via 10.10.7.1, 1w0d, GigabitEthernet0/2/0
C        10.10.7.0/24 is directly connected, GigabitEthernet0/2/0
L        10.10.7.2/32 is directly connected, GigabitEthernet0/2/0
C        10.10.8.0/24 is directly connected, GigabitEthernet0/2/1
L        10.10.8.2/32 is directly connected, GigabitEthernet0/2/1
      11.0.0.0/32 is subnetted, 1 subnets
O        172.16.11.11 [110/2] via 10.10.8.1, 1w0d, GigabitEthernet0/2/1
      22.0.0.0/32 is subnetted, 1 subnets
O        172.16.22.22 [110/2] via 10.10.7.1, 1w0d, GigabitEthernet0/2/0

ASR903-R4# show ip route 172.16.1.1
Routing entry for 172.16.1.1/32
  Known via "bgp 100", distance 200, metric 0, type internal
  Last update from 172.16.11.11 1d20h ago
  Routing Descriptor Blocks:
  * 172.16.11.11, from 172.16.11.11, 1d20h ago
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: 17248

ASR903-R4# show ip cef 172.16.1.1 detail
172.16.1.1/32, epoch 2, flags rib defined all labels
  recursive via 172.16.11.11 label 17248
    nexthop 10.10.8.1 GigabitEthernet0/2/1
  recursive via 172.16.22.22 label 16001, repair
    nexthop 10.10.7.1 GigabitEthernet0/2/0

ASR903-R4# ping 172.16.1.1 source 172.16.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Configuration for BGP PIC & RFC3107 (with ISIS)

ASR903-R1 Configuration

```
interface GigabitEthernet0/2/0
 ip address 10.10.1.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/2
 ip address 10.10.2.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
!
router isis
 net 49.0001.0000.0000.0001.00
 is-type level-1
 passive-interface Loopback0
!
```

ASR903-R2 Configuration

```
interface GigabitEthernet0/2/1
 ip address 10.10.3.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis circuit-type level-2-only
!
interface GigabitEthernet0/2/3
 ip address 10.10.4.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis circuit-type level-2-only
!
interface GigabitEthernet0/2/0
 ip address 10.10.1.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis circuit-type level-1
!
router isis
 net 49.0001.0000.0000.0002.00
 passive-interface Loopback0
!
```

ASR903-R3 Configuration

```
interface GigabitEthernet0/2/2
 ip address 10.10.2.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis circuit-type level-1
!
interface GigabitEthernet0/2/1
 ip address 10.10.3.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis circuit-type level-2-only
!
interface GigabitEthernet0/2/3
 ip address 10.10.5.1 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
 isis circuit-type level-2-only
!
router isis
 net 49.0001.0000.0000.0003.00
 passive-interface Loopback0
!
```

ASR903-R4 Configuration

```
interface GigabitEthernet0/2/0
 ip address 10.10.7.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
!
interface GigabitEthernet0/2/1
 ip address 10.10.8.2 255.255.255.0
 ip router isis
 negotiation auto
 mpls ip
!
router isis
 net 49.0001.0000.0000.0004.00
 is-type level-1
 passive-interface Loopback0
!
```

ASR9K-1 Configuration

```
router isis 1
 net 49.0001.0000.0000.0011.00
 interface Loopback10
  passive
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/16
  circuit-type level-2-only
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/17
  circuit-type level-1
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/18
  circuit-type level-2-only
  address-family ipv4 unicast
  !
```

ASR9K-2 Configuration

```
router isis 1
 net 49.0001.0000.0000.0022.00
 interface Loopback10
  passive
  address-family ipv4 unicast
  !
!
 interface GigabitEthernet0/0/0/16
  circuit-type level-2-only
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/17
  circuit-type level-1
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/18
  circuit-type level-2-only
  address-family ipv4 unicast
  !
 !
```

Chapter 10 QoS

Cisco ASR 903 supports three level hierarchical QoS, port, vlan and class level. Classification, marking, policing, shaping and LLQ are supported. At class level, you can have up to 8 classes including class-default under a policy-map. Three-level policy can only be applied to a physical port on the router. A three-level policy consists of:

- Topmost policy: class-default
- Middle policy: match vlan
- Lowest policy: match qos-group/match prec/match cos/match dscp

For example,

```
! class-level
class-map match-any P1
 match dscp ef
class-map match-any C3
 match dscp af31
class-map match-any C2
 match dscp af41
class-map match-any C1
 match dscp af11  af12
policy-map GRAND_CHILD          ! can have up to 8 classes at class-level
  class P1
    priority
    police cir percent 10
  class C1
   bandwidth remaining percent 10
  class C2
   bandwidth remaining percent 30
  class C3
   bandwidth remaining percent 40
  class class-default
! vlan level
class-map match-all vlan10
  match vlan 10
class-map match-all vlan20
  match vlan 20
class-map match-all vlan30
  match vlan 30
!
policy-map CHILD
  class vlan10
    shape average 30000000
    service-policy GRAND_CHILD
  class vlan20
    shape average 40000000
```

```
    service-policy GRAND_CHILD
  class vlan30
    shape average 50000000
    service-policy GRAND_CHILD
! port level
policy-map PARENT
  class class-default
    shape average 200000000
    service-policy CHILD

interface GigabitEthernet0/1/3
 no ip address
 negotiation auto
 service-policy output PARENT
 service instance 10 ethernet
  encapsulation dot1q 10
  rewrite ingress tag pop 1 symmetric
  bridge-domain 10
 !
 service instance 20 ethernet
  encapsulation dot1q 20
  rewrite ingress tag pop 1 symmetric
  bridge-domain 20
 !
 service instance 30 ethernet
  encapsulation dot1q 30
  rewrite ingress tag pop 1 symmetric
  bridge-domain 30
```

Not all the policy has to be three-level policy and here are some other choices.

Two-Level Policy

- Topmost policy: match vlan
- Lowest policy: match qos-group/match prec/match cos/match dscp

Two-Level Policy

- Topmost policy: class-default
- Lowest policy: match vlan

Two-Level Policy

- Topmost policy: class-default
- Lowest policy: match mpls experimental topmost

Flat policy: match ip dscp

Flat policy: match vlan inner

Flat policy: class-default

As a best practice, it is recommended to mark ingress packets into internal qos-group and then take action based on those qos-groups. Qos-groups are internal entities on the router and there are 100 usable qos-groups.

Ingress

```
class-map match-any NA4-1Q-VoIP
 match cos 5
 match mpls exp topmost 5
 match ip prec 5

class-map match-any NA4-1Q-Video
 match cos 4
 match mpls exp topmost 4
 match ip prec 4

policy-map NA4-1Q-Ingress
 class NA4-1Q-VoIP
   set qos-group 5
 class NA4-1Q-Video
   set qos-group 4
```

Egress

```
class-map match-all match-qos5
 match qos-group 5
class-map match-all match-qos4
 match qos-group 4



policy-map NA4-1Q-Egress
 class match-qos5
    priority
 class match-qos4
bandwidth remaining percent 15
```

For more information on configuring QoS on the Cisco ASR 903 router, see the *Quality of Service Configuration Guide for Cisco ASR 903 Router*.

Chapter 11 Clocking

ASR903 supports both Sync-E and Precision Time Protocol 1588v2. Here is the front panel of ASR903 RSP. BITS and 10MHz inputs can be used as external clocking sources. 1PPS and ToD can be used for 1588 PTP.

Figure 27. Cisco ASR903-RSP

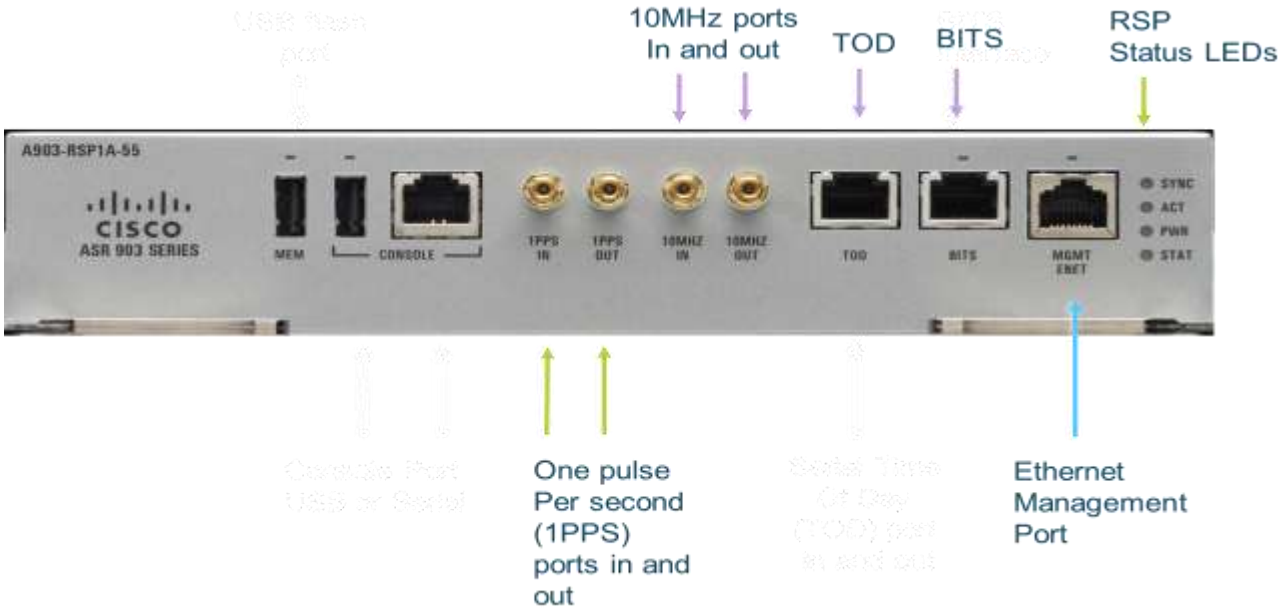

Table 3. Framing Modes for BITS Port on a Cisco ASR 903 Series Router

BITS or SSU Port Support Matrix	Framing Modes Supported	SSM or QL Support	Tx Port	Rx Port
T1	T1 ESF	Yes	Yes	Yes
T1	T1 SF	No	Yes	Yes
E1	E1 CRC4	Yes	Yes	Yes
E1	E1 FAS	No	Yes	Yes
2048 kHz	2048 kHz	No	Yes	Yes

Table 4. Clock Quality Level Matrix

	G.781		ESMC		PTP Clock Class
	Option I	Option II	Option I	Option II	
0001		QL-PRS			80
0000		QL-STU			82
0010	QL-PRC				84
0111		QL-ST2			86
0011					88
0100	QL-SSU-A	QL-TNC			90
0101					92
0110					94
1000	QL-SSU-B				96
1001					98
1101		QL-ST3E			100
1010		QL-ST3		QL-EEC2	102
1011	QL-SEC		QL-EEC1		104
1100		QL-SMC			106
1110		QL-PROV			108
1111	QL-DNU	QL-DUS			110

When the best available clock is to be selected, there are two modes for consideration. In quality disabled mode, only clock availability and priority matter, and quality level is ignored. In quality mode, quality level is used a selection criteria too.

Quality Disabled Mode:

- Clock Availability ! clock must be available and error & alarm free
- Local Priority

Quality Enabled Mode:

- Clock Availability ! clock must be available and error & alarm free
- Quality Level ! next is to check quality level
- Local Priority ! last is to look at the priority

Sync-E Configuration

Synchronous Ethernet is an extension of Ethernet designed to provide the reliability found in traditional SONET/SDH and T1/E1 networks to Ethernet packet networks by incorporating clock synchronization features. It supports the Synchronization Status Message (SSM) and Ethernet Synchronization Message Channel (ESMC) for synchronous Ethernet clock synchronization.

The following are available clocking sources on the ASR903.

From the Interfaces of IM (Max 2 per IM, up to 4 in total)

- T1/E1 Controllers of 16x T1/E1 TDM IM (SSM not supported yet)
- OC3/OC12 Controller (SSM not supported yet)
- Gig interface of 8 SFP IM
- Gig interface of 1x 10 Gig IM
- Gig interface on 8 Cu IM

From BITS and 10Mhz from RSP

- BITS interface on Active RSP(SSM supported)
- BITS interface on Standby RSP(SSM supported)
- 10 Mhz on Active RSP

Here are the commands to specify clocking sources.

Elect Interface as Clocking Candidates

```
network-clock input-source priority   interface  interface
network-clock input-source 10 interface GigabitEthernet 0/0/1
```

Elect 10MHz and BITS as Input Sources

```
network-clock input-source priority external  {R0 |R1 }   input source
```

Figure 28. Input Source Options

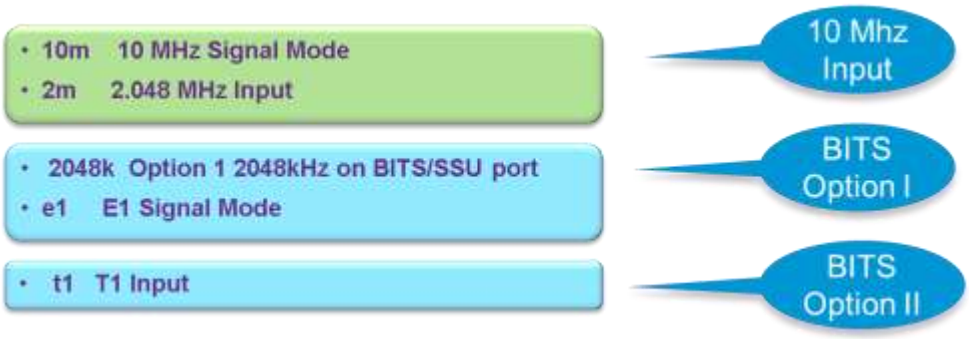

Here is the diagram for SYNC-E configuration. R1 router has its BITS port connected to external clock source and then uses SYNC-E to distribute it to other routers.

Figure 29. SYNC-E

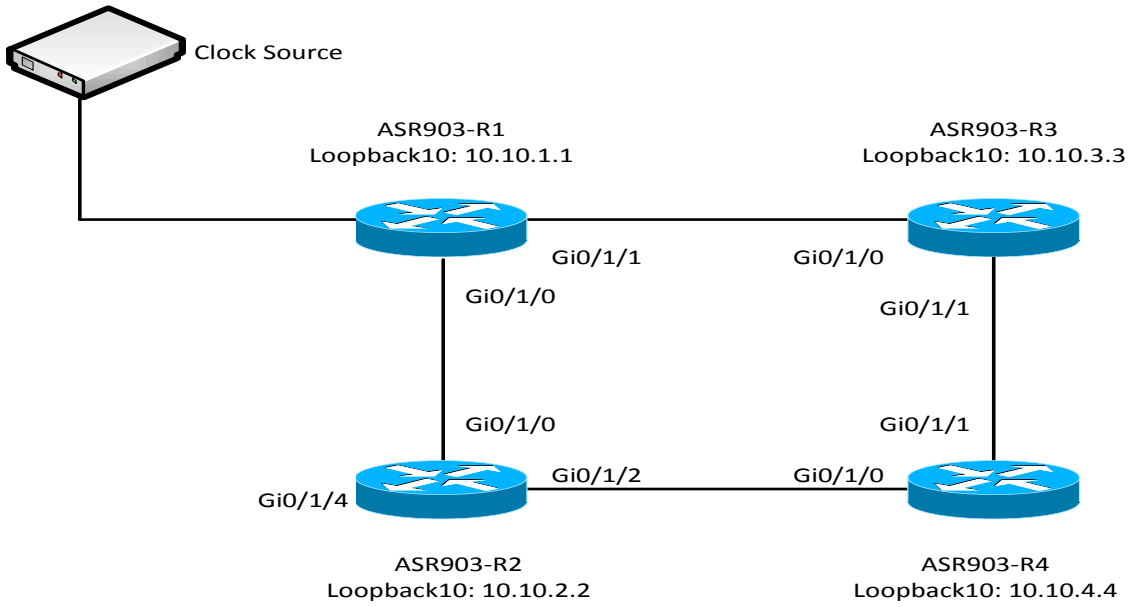

ASR903-R1 Configuration

```
interface GigabitEthernet0/1/1
  ! description to R3
  synchronous mode                                ! enable sync-e under interface
!
interface GigabitEthernet0/1/0
  description to R2
  synchronous mode

network-clock synchronization automatic                ! enable clock selection
algorithm
network-clock synchronization ssm option 1             ! to select E1 or T1 type
network-clock synchronization mode QL-enabled          ! use Qualiy enabled mode
network-clock input-source 10 External R0 e1 crc4      ! BITS clock as source
network-clock wait-to-restore 30 global
esmc process                                           ! enable ESMC process
```

ASR903-R2 Configuration

```
interface GigabitEthernet0/1/0
  ! description to R1
  synchronous mode                                ! enable sync-e under interface
!
interface GigabitEthernet0/1/2
  description to R4
  synchronous mode

network-clock synchronization automatic                ! enable clock selection
algorithm
network-clock synchronization ssm option 1             ! to select E1 or T1 type
network-clock synchronization mode QL-enabled          ! use Qualiy enabled mode
network-clock input-source 10 interface Gig0/1/0       ! elect gi0/1/0 as candidiate
                                                       ! lower priority is preferred
network-clock input-source 20 interface Gig0/1/2       ! elect gi0/1/2 as candidiate
network-clock wait-to-restore 30 global
esmc process                                           ! enable ESMC process
```

ASR903-R3 Configuration

```
interface GigabitEthernet0/1/0
  ! description to R1
  synchronous mode                                ! enable sync-e under interface
!
interface GigabitEthernet0/1/1
  description to R4
  synchronous mode

network-clock synchronization automatic                ! enable clock selection
algorithm
network-clock synchronization ssm option 1             ! to select E1 or T1 type
```

```
network-clock synchronization mode QL-enabled          ! use Qualiy enabled mode
network-clock input-source 10 interface Gig0/1/0       ! elect gi0/1/0 as candidiate
                                                       ! lower priority is preferred
network-clock input-source 20 interface Gig0/1/1       ! elect gi0/1/1 as candidiate
network-clock wait-to-restore 30 global
esmc process                                           ! enable ESMC process
```

ASR903-R4 Configuration

```
interface GigabitEthernet0/1/0
  ! description to R2
  synchronous mode                                ! enable sync-e under interface
!
interface GigabitEthernet0/1/1
  description to R3
  synchronous mode

controller E1 0/4/0
 clock source internal                            ! to use synchronized clock
!
network-clock synchronization automatic                ! enable clock selection
algorithm
network-clock synchronization ssm option 1             ! to select E1 or T1 type
network-clock synchronization mode QL-enabled          ! use Qualiy enabled mode
network-clock input-source 10 interface Gig0/1/1       ! elect gi0/1/1 as candidiate
                                                       ! lower priority is preferred
network-clock input-source 20 interface Gig0/1/0       ! elect gi0/1/0 as candidiate
network-clock wait-to-restore 30 global
esmc process                                           ! enable ESMC process
```

Verifying Clocking

```
ASR903-R1# show network-clock synchronization detail
Symbols: En - Enable, Dis - Disable, Adis - Admin Disable
NA - Not Applicable
* - Synchronization source selected
# - Synchronization source force selected
& - Synchronization source manually switched

Automatic selection process : Enable
Equipment Clock : 2048 (EEC-Option1)
Clock Mode : QL-Enable
ESMC : Enabled
SSM Option : 1
T0 : External R1 e1 crc4 120ohms                    ! clock is synced to
Hold-off (global) : 300 ms
Wait-to-restore (global) : 30 sec
Tsm Delay : 180 ms
```

```
Revertive : No
Force Switch: FALSE
Manual Switch: FALSE
Number of synchronization sources: 1
Squelch Threshold: QL-SEC
sm(netsync NETCLK_QL_ENABLE), running yes, state 1A
Last transition recorded: (begin)-> 2A (ql_mode_enable)-> 1A (src_added)-> 1A
(ql_change)-> 1A (ql_change)-> 1A (sf_change)-> 1A

Nominated Interfaces

Interface     SigType    Mode/QL   Prio   QL_IN   ESMC Tx ESMC Rx
Internal         NA       NA/Dis     251   QL-SEC     NA     NA
*External  R1 E1 CRC4 120 NA/En     10    QL-PRC     NA     NA


Interface:
-----------------------------------------------
Local Interface: Internal
Signal Type: NA
Mode: NA(Ql-enabled)
SSM Tx: DISABLED
SSM Rx: DISABLED
Priority: 251
QL Receive: QL-SEC
QL Receive Configured: -
QL Receive Overrided: -
QL Transmit: -
QL Transmit Configured: -
Hold-off: 0
Wait-to-restore: 20
Lock Out: FALSE
Signal Fail: FALSE
Alarms: FALSE
Slot Disabled: FALSE
SNMP input source index: 1
SNMP parent list index: 0

Local Interface: External R1
Signal Type: E1 CRC4 120ohms
Mode: NA(Ql-enabled)
SSM Tx: ENABLED
SSM Rx: ENABLED
Priority: 10
QL Receive: QL-PRC
QL Receive Configured: -
QL Receive Overrided: -
QL Transmit: -
QL Transmit Configured: -
Hold-off: 300
Wait-to-restore: 20
```

```
Lock Out: FALSE
Signal Fail: FALSE          ! signal fail must be false to be clock candidiate
Alarms: FALSE               ! alarms  must be false to be clock candidiate
Active Alarms : None
Slot Disabled: FALSE
SNMP input source inde

ASR903-R3# show network-clocks synchronization
Symbols:     En - Enable, Dis - Disable, Adis - Admin Disable
             NA - Not Applicable
             *  - Synchronization source selected
             #  - Synchronization source force selected
             &  - Synchronization source manually switched

Automatic selection process : Enable
Equipment Clock : 2048 (EEC-Option1)
Clock Mode : QL-Enable
ESMC : Enabled
SSM Option : 1
T0 : GigabitEthernet0/1/0
Hold-off (global) : 300 ms
Wait-to-restore (global) : 30 sec
Tsm Delay : 180 ms
Revertive : No

Nominated Interfaces

 Interface            SigType     Mode/QL      Prio  QL_IN  ESMC Tx  ESMC Rx
 Internal             NA          NA/Dis       251   QL-SEC    NA        NA
*Gi0/1/0              NA          Sync/En      10    QL-PRC    -         -
 Gi0/1/1              NA          Sync/En      20    QL-DNU    -         -


ASR903-R4# show network-clocks synchronization
Symbols:     En - Enable, Dis - Disable, Adis - Admin Disable
             NA - Not Applicable
             *  - Synchronization source selected
             #  - Synchronization source force selected
             &  - Synchronization source manually switched

Automatic selection process : Enable
Equipment Clock : 2048 (EEC-Option1)
Clock Mode : QL-Enable
ESMC : Enabled
SSM Option : 1
T0 : GigabitEthernet0/1/1
Hold-off (global) : 300 ms
Wait-to-restore (global) : 30 sec
Tsm Delay : 180 ms
```

```
Revertive : No

Nominated Interfaces

 Interface            SigType      Mode/QL      Prio  QL_IN  ESMC Tx  ESMC Rx
 Internal             NA           NA/Dis       251   QL-SEC    NA        NA
*Gi0/1/1              NA           Sync/En      10    QL-PRC    -         -
 Gi0/1/0              NA           Sync/En      20    QL-PRC    -
```

PTP Configuration

The Precision Time Protocol (PTP), as defined in the IEEE 1588 standard, synchronizes with nanosecond accuracy the real-time clocks of the devices in a network. The clocks in are organized into a master-member hierarchy. PTP identifies the switch port that is connected to a device with the most precise clock. This clock is referred to as the master clock. All the other devices on the network synchronize their clocks with the master and are referred to as members. Constantly exchanged timing messages ensure continued synchronization. 1588 PTP can be used to synchronize clocking and time.

The following modes are supported on the ASR903.

Figure 30. PTP Modes

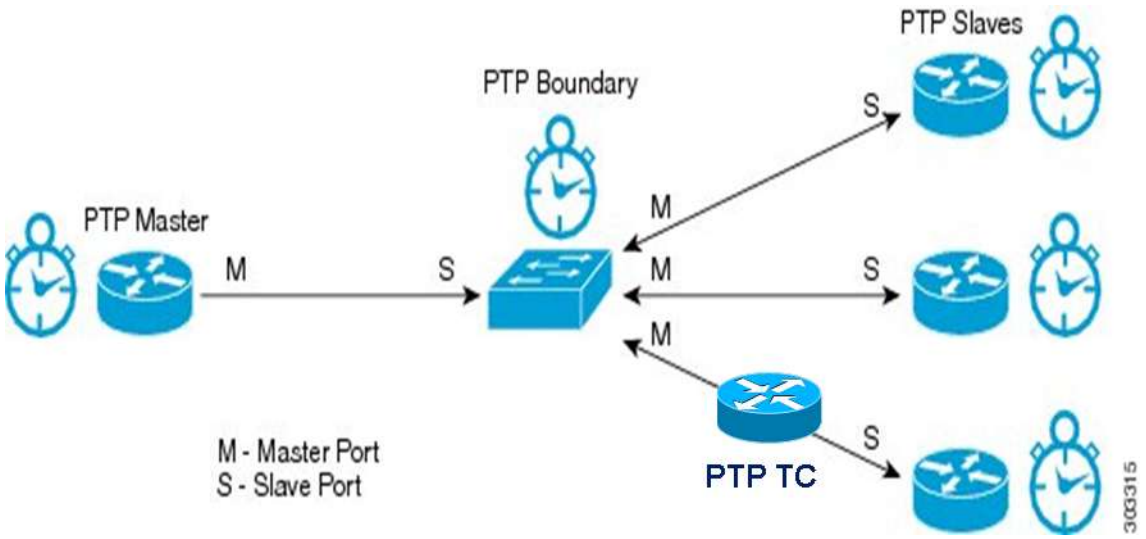

- Master Clock: one master clock port, requires 10 Mhz, 1PPS and ToD
- Slave Clock: one slave clock port
- Boundary: one master clock port and one or more slave ports.
- Hybrid: use sync-e for frequency and PTP for time/phase.
- Transparent: just time-stamp PTP packets which transit the router.

  Here is the topology to use for PTP configuration. ASR903-R1 and R2 will be configured as BC, ASR903-R3 as TC, and ASR903-R4 as slave to synchronize with R1 and R2.

Figure 31. PTP

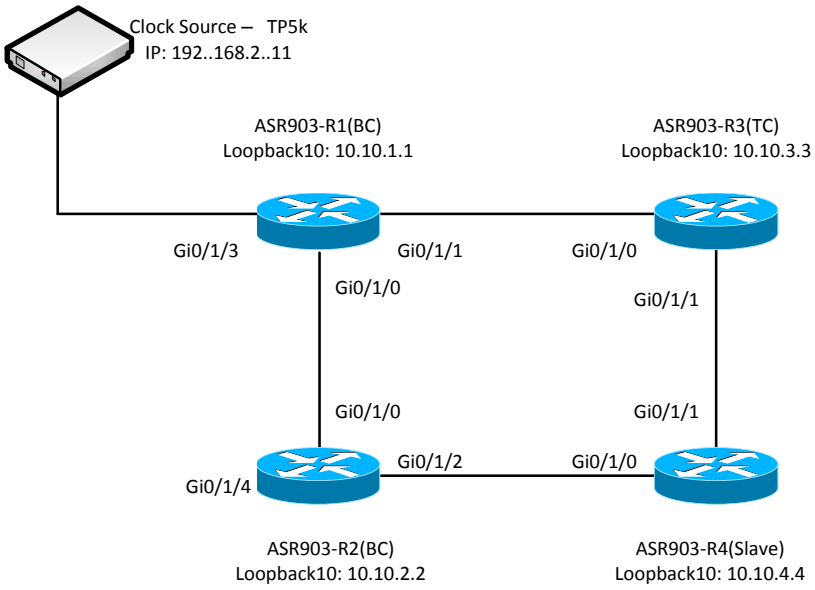

ASR903-R1 Configuration

```
interface Loopback10                        ! use seperate interfaces for master and
slave
 ip address 10.10.1.1 255.255.255.255
!
interface Loopback11
 ip address 10.10.158.1 255.255.255.255
!
network-clock synchronization automatic
network-clock synchronization mode QL-enabled
network-clock input-source 10 ptp domain 0    ! specify PTP as clock source
!
ptp clock boundary domain 0
 clock-port TPSlave slave                   ! upstream master and itself slave
  transport ipv4 unicast interface Lo10 negotiation
  clock source 192.168.2.11                 ! upstream master IP address
 clock-port Master master                   ! master for the downstream slaves
  transport ipv4 unicast interface Lo11 negotiation
```

ASR903-R2 Configuration

```
interface Loopback10                        ! use seperate interfaces for master and
slave
 ip address 10.10.2.2 255.255.255.255
!
interface Loopback11
 ip address 10.10.158.2 255.255.255.255
!
network-clock synchronization automatic
network-clock synchronization mode QL-enabled
network-clock input-source 10 ptp domain 0          ! specify PTP as clock source
!
ptp clock boundary domain 0
 clock-port TPSlave slave
  transport ipv4 unicast interface Lo10 negotiation
  clock source 10.10.158.1
 clock-port Master master
  transport ipv4 unicast interface Lo11 negotiation
```

ASR903-R3 Configuration

```
ptp clock e2e-transparent domain 0
```

ASR903-R4 Configuration

```
interface Loopback10
 ip address 10.10.4.4 255.255.255.255
!
network-clock synchronization automatic
network-clock synchronization mode QL-enabled
network-clock input-source 10 ptp domain 0
!
ptp clock ordinary domain 0
 clock-port SLAVE slave
  transport ipv4 unicast interface Lo10 negotiation
  clock source 10.10.158.1 1       ! can select multiple sources, lower number
  clock source 10.10.158.2 2       ! higher priority
```

Verifying PTP and Clock Synchronization

```
ASR903-R1# show ptp clock running domain 0

                    PTP Boundary Clock [Domain 0]

         State          Ports          Pkts sent      Pkts rcvd      Redundancy Mode

         PHASE_ALIGNED  2              424467         362306         Hot standby

                              PORT SUMMARY

PTP Master
```

```
Name                Tx Mode      Role         Transport    State        Sessions
Port Addr

TPSlave             unicast      slave        Lo10         Slave        1
192.168.2.11
Master              unicast      master       Lo11         Master       2            -


                              SESSION INFORMATION

TPSlave [Lo10] [Sessions 1]

 Peer addr           Pkts in    Pkts out   In Errs    Out Errs

 192.168.2.11        250189     84844      0          0

Master [Lo11] [Sessions 2]

 Peer addr           Pkts in    Pkts out   In Errs    Out Errs

 10.10.2.2          64635      195790     0          0
 10.10.4.4          47482      143833     0

ASR903-R1# show network-clocks synchronization
Symbols:     En - Enable, Dis - Disable, Adis - Admin Disable
             NA - Not Applicable
             *  - Synchronization source selected
             #  - Synchronization source force selected
             &  - Synchronization source manually switched

Automatic selection process : Enable
Equipment Clock : 2048 (EEC-Option1)
Clock Mode : QL-Enable
ESMC : Disabled
SSM Option : 1
T0 : Internal
Hold-off (global) : 300 ms
Wait-to-restore (global) : 30 sec
Tsm Delay : 180 ms
Revertive : No

Nominated Interfaces

 Interface            SigType     Mode/QL      Prio  QL_IN  ESMC Tx  ESMC Rx
 Internal             NA          NA/Dis       251   QL-SEC    NA        NA
 *PTP (Dom 0)         NA          NA/En        10    QL-PRC    NA        NA
```

```
ASR903-R4# show ptp  clock running domain 0

                    PTP Ordinary Clock [Domain 0]

         State           Ports          Pkts sent      Pkts rcvd      Redundancy Mode

         PHASE_ALIGNED   1              106295         321979         Hot standby

                              PORT SUMMARY

PTP Master
Name                Tx Mode      Role         Transport    State        Sessions
Port Addr

SLAVE               unicast      slave        Lo10         Slave        2
10.10.158.1


                             SESSION INFORMATION

SLAVE [Lo10] [Sessions 2]

 Peer addr          Pkts in    Pkts out   In Errs    Out Errs

 10.10.158.1      160990     53148      0          0
 10.10.158.2      160989     53147      0          0
```

Chapter 12 Cisco ASR 903 Software Upgrade

Software Boot up options

Software packages in ASR 903 are broadly categorized into two parts as below

- Consolidated package mode
- Sub package mode

Important: ISSU in consolidated mode is not supported on ASR 903

The below sections explain in details the above two modes of software and ISSU procedures.

Consolidated package mode

Consolidated image is a single software image containing a full collection of software packages. Consolidated mode provides a simplified installation and can be stored in bootflash, a TFTP server, or a network server.

Booting in Consolidated Mode

Step 1 Set the configuration register boot field to boot the system image as specified by the boot system command in the NVRAM configuration. (0x2)

```
Router# configure terminal
Router(config)# config-register 0x2
Router(config)# end
```

Step 2 Configure the NVRAM boot system command to boot a consolidated system image which has previously been copied to the system bootflash. In a dual RP system, the consolidated package must also exist in the same location on the standby RP bootflash.

```
Router# configure terminal
Router(config)# boot system bootflash:asr903-adventerprisek9.base.bin
Router(config)# end
Router#write mem
```

Step 3 Reload the router.

Important: In case a device booted in consolidate mode, you should first re-boot the router in sub package mode and perform ISSU. The procedure to boot router in sub package mode is described below.

Sub package mode

Sub-package mode is achieved when the system is initially booted with a packages.conf file. A packages.conf file is created when a consolidated package is expanded using the **request platform software packages expand** command.

Booting in Sub-Package Mode

Step 1 Set the configuration register boot field to boot the system image as specified by the boot system command in the NVRAM configuration. (0x2)

```
Router# configure terminal
Router(config)# config-register 0x2
Router(config)# end
```

Step 2 Expand the consolidated image which has previously been copied to the system bootflash. In a dual RP system, the consolidated image should exist on, and be expanded on both the active and standby RP bootflash.

It is good practice to create a separate directory for ISSU and expand the image in that directory. After expanding the consolidated package, the original package will exist as well as sub package files and a packages.conf file that can be used for booting.

Important: Do not copy the packages.conf file to a new directory after expanding the package. It is required that the packages.conf file and sub package files exist in the same directory.

```
Router# mkdir issu
Create directory filename [issu]?
Created dir bootflash:/issu
Router# copy filename directory
```

Example:

```
Router# copy asr903rsp1-adventerprisek9.upgrade.bin bootflash:/issu
Router# request platform software package expand file  bootflash:issu/asr903rsp1-
adventerprisek9.base.bin
Router# request platform software package expand file stby-bootflash:issu/asr903rsp1-
adventerprisek9.base.bin
```

Step 3 Configure the NVRAM boot system command to boot the packages.conf file.

```
Router# configure terminal
Router(config)# boot system bootflash:issu/packages.conf
Router(config)# end
Router# write mem
```

Step 4 Reload the router.

ISSU One Shot Upgrade Overview

The ISSU one shot update is the preferred method of updating software on a dual RSP ASR 903 system. The one shot update is designed to operate optimally when the router is booted in sub-package mode. In this mode, the router will perform rolling upgrades of interface module software (and firmware if applicable).

The one-shot ISSU procedure can be used to install a complete set of sub-packages using a single command. When used in sub-package mode, this command will:

Step 1 Expand the consolidated package into a complete set of subpackages.

Step 2 Install the complete set of packages on the standby RP.

Step 3 Reload the standby RP.

When the standby RP is reloaded, the Niles manger will determine whether or not a firmware upgrade is required for the Niles Handoff FPGA. (i.e. Firmware versioning is performed independently by the subpackage and is not maintained directly in the ISSU scripts or packages.conf file). If a firmware upgrade is required, the firmware will be programmed into the EEPROM and the FPGA will be restarted with the upgraded firmware. This will extend the re-boot time but will not affect packet forwarding since the process occurs independently on the standby RP. It should be noted that firmware upgrades are rare and are normally not required during most ISSU upgrades.

Step 4 Perform a rolling reload of each interface modules on the active RP by initiating software OIR of each IM. When each IM is OIRed, the IOMD software will determine whether or not a firmware upgrade is required for an IM FPGA. If a firmware upgrade is required, the firmware upgrade will be programmed into the EEPROM and the IM will be OIRed again. This will extend the outage time for each IM that requires a firmware upgrade.

Important: Firmware upgrades are relatively rare and are not required during most software upgrades.

Some IMs (presently E1/T1 modules) are intelligent with respect to the fact that they must be downloaded with software prior to becoming operational after an OIR . For these IMs, a software download is expected for every ISSU procedure. In the case that a firmware upgrade is required for these IMs, the firmware will be upgraded first. Then the IM will be OIRed a second time before becoming operational. The rolling upgrade will not proceed with subsequent IM modules until after the current IM module has become operational. The amount of delay between each IM OIR can be specified by the user using the interface-module-delay keyword option.

Step a Initiate an HA switchover.

Step b Install the remaining packages on the new standby (previously active) RP.

Step c Reload the new standby RP.

When the new standby is reloaded, it may result in a firmware upgrade of the Niles Handoff FPGA as described above.

It is important to recognize that the elapsed time to perform the one shot ISSU is variable and dependent upon several factors. The factors that affect user downtime are:

- Per IM outage times— this variable depends upon FPGA upgrade time (if required), the IM software download time, and the value of the interface-module-delay keyword
- HA switchover time — 50 msec or less
- Network topology—to reduce outage time during an ISSU upgrade. It is important to design the network with redundant route paths that allow traffic to be re-routed during individual IM outages. In this respect care must be taken to insure redundant L2 and L3 routes are properly configured.

Table 5. IM-FPGA Upgrade

IM	Reload required	Packet Loss due to RSP Switchover	Packet loss due to IM switchover	Packet loss due to IM-FPGA upgrade
Ethernet IMs	Only when IM-FPGA upgraded	50 ms	-	30 – 60 seconds
TDM IMs	Always	50 ms	30 – 60 seconds	-

ISSU Preparation

In order to successfully achieve an ISSU upgrade using this procedure. A dual RP router must have previously been booted in sub-package mode and the redundancy mode SSO must have been achieved. A consolidated upgrade image should have been copied on the bootflash to the directory that was used to boot the base image. It is not necessary to copy the upgrade image to the stby-bootflash as this will be performed as part of the one shot procedure. Care should be taken to make sure stby-bootflash can accommodate the consolidated bin file as well expanded package files. ISSU requires at least 450M space on both active and standby bootflash. This is to allow space for expansion and extraction of files during bootup.

The following example demonstrates the contents of the bootflash and stby-bootflash prior to the one shot ISSU procedure.

```
Directory of bootflash:/issu/
     6380   Apr 18 2013 22:38:05 +00:00  packages.conf
223480164  Mar 14 2000 08:17:36 +00:00  asr903rsp1-universal.03.08.01.S.153-1.S1.bin
 34371940  Apr 18 2013 16:01:00 +00:00  asr903rsp1-espbase.03.08.01.S.153-1.S1.pkg
     5646   Apr 18 2013 16:01:00 +00:00  asr903rsp1-packages-
universal.03.08.01.S.153-1.S1.conf
 25194852  Apr 18 2013 16:01:01 +00:00  asr903rsp1-rpaccess.03.08.01.S.153-1.S1.pkg
 34224484  Apr 18 2013 16:01:08 +00:00  asr903rsp1-rpbase.03.08.01.S.153-1.S1.pkg
 26745188  Apr 18 2013 16:01:15 +00:00  asr903rsp1-rpcontrol.03.08.01.S.153-1.S1.pkg
 57137508  Apr 18 2013 16:01:31 +00:00  asr903rsp1-rpios-universal.03.08.01.S.153-
1.S1.pkg
 25725284  Apr 18 2013 16:01:38 +00:00  asr903rsp1-sipbase.03.08.01.S.153-1.S1.pkg
 24547684  Apr 18 2013 16:01:45 +00:00  asr903rsp1-sipspa.03.08.01.S.153-1.S1.pkg
228995428  May 23 2013 21:10:05 +00:00  asr903rsp1-universal.03.09.00.S.153-2.S.bin

Directory of stby-bootflash:/issu/
     6380   Apr 18 2013 22:38:05 +00:00  packages.conf
223480164  Mar 14 2000 08:17:36 +00:00  asr903rsp1-universal.03.08.01.S.153-1.S1.bin
 34371940  Apr 18 2013 16:01:00 +00:00  asr903rsp1-espbase.03.08.01.S.153-1.S1.pkg
     5646   Apr 18 2013 16:01:00 +00:00  asr903rsp1-packages-
```

```
universal.03.08.01.S.153-1.S1.conf
 25194852  Apr 18 2013 16:01:01 +00:00  asr903rsp1-rpaccess.03.08.01.S.153-1.S1.pkg
 34224484  Apr 18 2013 16:01:08 +00:00  asr903rsp1-rpbase.03.08.01.S.153-1.S1.pkg
 26745188  Apr 18 2013 16:01:15 +00:00  asr903rsp1-rpcontrol.03.08.01.S.153-1.S1.pkg
 57137508  Apr 18 2013 16:01:31 +00:00  asr903rsp1-rpios-universal.03.08.01.S.153-
1.S1.pkg
 25725284  Apr 18 2013 16:01:38 +00:00  asr903rsp1-sipbase.03.08.01.S.153-1.S1.pkg
 24547684  Apr 18 2013 16:01:45 +00:00  asr903rsp1-sipspa.03.08.01.S.153-1.S1.pkg
```

The redundancy mode should be configured for SSO and the SSO state should be reached prior to starting the test.

```
Router# configure terminal
Router(config)# redundancy
Router(config)# mode ss0
Router(config)# end
*Jan 12 17:52:26.516: %RF-5-RF_TERMINAL_STATE: Terminal state reached for (SSO)
```

ISSU Procedure

Step 1 Issue the request platform software package install command on the active RP.

```
Router# request platform software package install node file bootflash:/issu/
asr903rsp1-universal.03.09.00.S.153-2.S.bin interface-module-delay 150
```

Step 2 Wait for the successive stages to complete as indicated by STAGE and SUCCESS messages that are displayed on the active RP console after each stage completes. In general interface-module-delay timer of 120-150 seconds is recommended most of the cases to make sure that one IM can fully boot up before resetting next IM.

Step 3 Wait for the original active RP to reboot. This will occur after the completion of STAGE 5.

Step 4 Connect to the new active console wait for the redundancy state to return to SSO state.

Appendix A References

ASR 903 Web site:

http://www.cisco.com/en/US/products/ps11610/index.html