

Multi-Chassis Ring Layer 2 with Enhanced Subscriber Management

In This Chapter

This section provides information about MC-Ring Layer 2 with Enhanced Subscriber Management (ESM).

Topics in this section include:

- [Applicability on page 2546](#)
- [Summary on page 2547](#)
- [Overview on page 2548](#)
- [Configuration on page 2553](#)
- [Conclusion on page 2577](#)

Applicability

These configuration notes are applicable to all of the 7450, 7750 and 7710 SR series and was tested on Release 7.0R5. The 7750 SR-c4 is supported from 8.0R4 and higher.

MC-Ring L2 with Enhanced Subscriber Management (ESM) was introduced in SR OS 6.0. There are no other pre-requisites for the 7450, 7750 and 7710 SR this configuration.

Summary

Multi-Chassis Ring (MC-ring) is an extension for dual homing support in TPSDA (Triple Play Service Delivery Architecture) networks based on Layer 2-CO (Layer 2 Central Office) model. The extension addresses networks where multiple access nodes (for example, DSLAMs, GPON OLT) are connected in a single ring.

MC-ring Layer 2 ESM is considered an extension or evolution for ring topologies of the Multi-Chassis LAG dual-homing solution used for directly connected access nodes.

MC-ring Layer 2 CO is documented in the Triple Play Enhanced Subscriber Management / Dual Homing section of the Triple Play Guide.

Overview

MC-Ring

Figure 391 shows a typical ring-based configuration of network model based on the Layer 2 CO model.

Individual rings of access nodes are aggregated at the Broadband Service Aggregator (BSA) level in one (or multiple) Virtual Private LAN Service (VPLS) service(s). At higher aggregation levels, Broadband Service Router (BSR) individual BSAs are connected to Layer 3 interface (Internet Enhanced Service (IES) or Virtual Private Routed Network (VPRN)) by means of spoke-SDP (Service Distribution Point) termination. Every Layer 3 interface at the BSR level aggregates all subscribers in one or more subnets. ESM is performed in the BSAs.

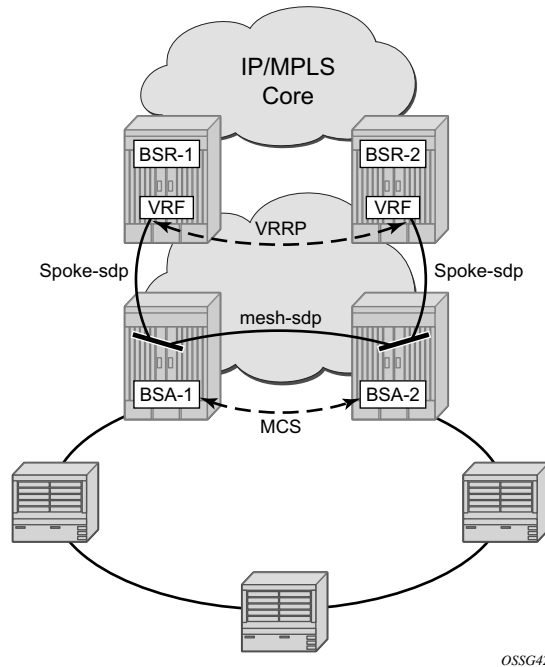


Figure 391: MC-Ring Layer 2 CO Dual Homing

The key functional components of the MC-Ring Layer 2 CO dual-homing redundancy solution are:

1. Mirroring of the subscriber state between the two BSAs performing subscriber management using the multi-chassis synchronization (MCS) protocol (BSA-1 and BSA-2 in Figure 391).
2. Ring control between the two BSAs, using the following mechanisms:

- In-band bi-directional forwarding detection (BFD) between the BSAs over the ring to monitor ring integrity and detect failures. A BFD session between BSA-1 and BSA-2 runs through the access ring using a dedicated IES/VPRN interface configured on both BSAs. This connection uses a separate VLAN throughout the ring (access nodes provide transparent bridging for this VLAN).
 - Out-of-band communication between the BSA nodes to exchange information about the reachability of individual access nodes, and to verify the configuration consistency of the ring. The information on configuration is synchronized through MCS (this use of MCS is related to, but independent to, MCS for subscriber-state synchronization mentioned above).
3. Ring Node Connectivity Verification (RNCV). Each BSA uses RNCV to detect the reachability of individual Access Nodes. It is used after a ring failure to determine which BSA should handle the traffic of each Access Node. RNCV uses ARP requests to “ping” individual ANs, which must be configured with an IP address for this purpose.
 4. Per-subscriber attribute (intermediate destination ID) for the BSA to know the Access Node each subscriber is connected to (assigned through RADIUS or DHCP/Python).
 5. VRRP in the BSRs to provide a redundant default gateway to the CPEs/Home Gateways. BSRs do not perform any subscriber management functions.

The operation of dual homing at the BSA level will be covered based on two underlying mechanisms.

MC-Ring Layer 2 CO Operation

To describe the functional behavior and operation of the dual-homing concept in a ring, it is best to sub-divide the description into the following three parts:

- Steady-state operation with ring fully closed
- Transition to ring-broken state
- Transition from ring-broken to steady state operation

Steady-State Operation of Dual-Homed Ring

Figure 392 illustrates the detailed operation of the dual-homed ring. The steady-state is achieved under following conditions:

- Both nodes are configured in a consistent way
- The MCS peering relation is up
- In-Band Ring Control Connection (IB-RCC) is in an operationally UP state. Note that this connection is set-up using BFD session between IP interfaces on BSA-1 and BSA-2

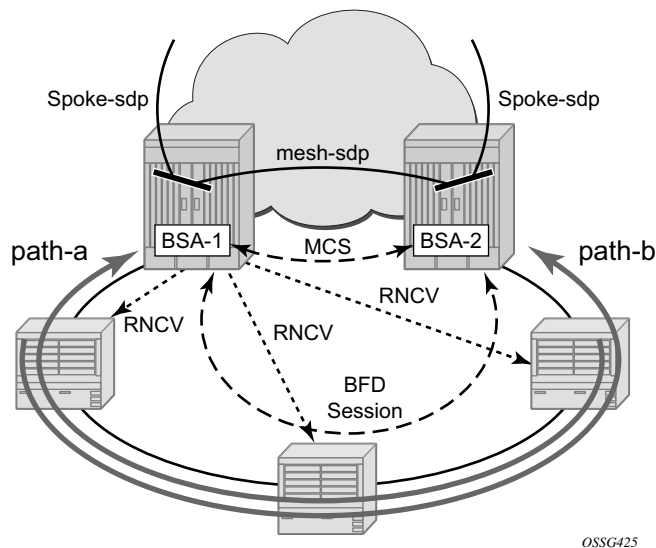


Figure 392: Dual homing Under Steady-State Condition

Under the above conditions, the ring is fully closed and every access node (the **ring-node**) has two possible paths toward the VPLS core. Figure 392 refers to them as **path-a** and **path-b**. In order to avoid the loop created by the ring, only one of the paths may be used by any given ring-node for any given VLAN. The assignment of the individual VLANs to path-a or path-b, respectively, has

to be provisioned on both BSAs in a consistent manner. The BSA with a lower IP address in the interface used for BFD will be the master for the VLANs associated with path-a and standby for the VLANs associated with path-b.

The following summarizes the behavior of the master and standby BSA for a given path (a or b) and the VLANs configured for that path:

Master BSA for the VLANs/SAPs associated with a given path (a or b):

- SAPs associated with the path are operationally up and FDB entries for sub-hosts point to the corresponding SAP
- Master BSA for a path performs RNCV checks to all ring nodes. RNCV failures trigger an alarm but do **not** trigger a switchover
- The ARP reply agent replies to ARP requests for subscriber-hosts associated with ring-nodes for which the BSA is master

Standby BSA:

- All SAPs associated with the path for which the BSA is standby will be operationally down and all FDB entries of subscriber hosts associated with those SAPs will point towards an SDP connecting to the master BSA.

Broken-Ring Operation and the Transition to this State

[Figure 393](#) illustrates the scenario with the broken ring (link failure or ring-node failure). This state is reached under following conditions:

- Both nodes are configured in a consistent way
- The MCS peering relation is up
- IB-RCC is in operationally down state

In this situation, every ring node has only one access path towards the VPLS core and hence the path-a and path-b notion has no real meaning in this situation.

From a functional point of view, both BSAs are now the master for the ring-nodes they can reach. For all hosts behind unreachable ring-nodes, the corresponding subscriber host FDB (Forwarding DataBase) entries will point to SDP pointing to the other head-end ring PE.

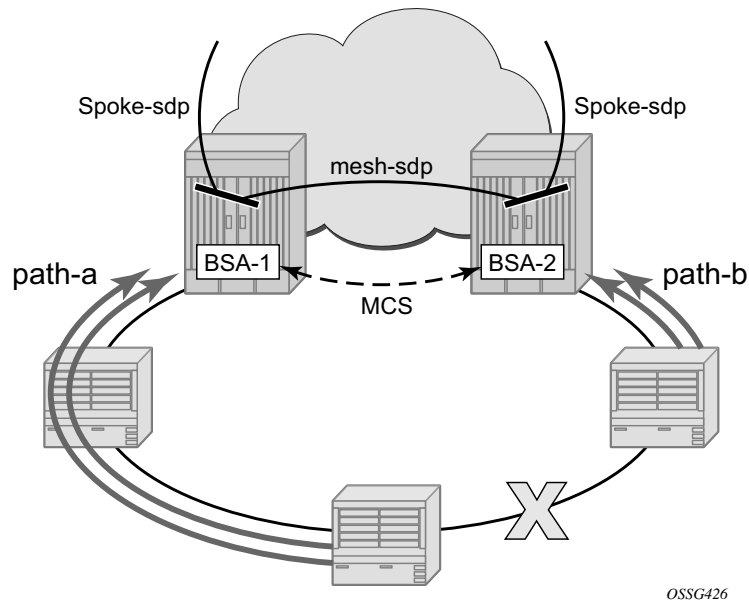


Figure 393: Broken Ring State

Transition from Broken to Closed Ring State

MC-ring operates in a revertive mode. This means that whenever the ring connectivity is restored, the BSA with the lower IP address in the IB-RCC communication channel will become master of path-a and slave for ring-b with vice versa operations on the other BSA.

The actions in this case are straightforward. After restoration of BFD session, the master functionality is assumed by respective BSAs. The FDB tables are updated according to the master/standby role of the given BSA and FDB population messages will be sent accordingly.

Configuration




Figure 394: Network Topology

The network topology is displayed in [Figure 394](#). The setup consists of two BSA nodes (PE-1 and PE-2), two BSR nodes (PE-3, PE-4) and another PE router (PE-5). A setup with one BSR node and two BSA nodes can also be used, but in this example, two (2) BSR nodes were used to show the typical Layer 2-CO setup with VRRP and Protocol Independent Multicast (PIM) on the BSRs. The access ring consists of four CE nodes.

Base Topology

This section assumes that the following base configuration has been implemented on the PE:

- Cards, MDAs and ports configured
- Interfaces configured
- IGP (Interior Gateway Protocol) configured and converged LDP (Label Distribution Protocol) configured on the interfaces between PE-3-PE-1, PE-1-PE-2 and PE-2-PE-4
- T-LDP (Targeted-LDP) configured on PE-1, PE-2, PE-3, PE-4
- SDPs configured between PE-3-PE-1, PE-1-PE-2 and PE-2-PE-4

Note that you can choose between OSPF and IS-IS as the IGP. Both LDP and RSVP (Resource Reservation Protocol) can be used for signaling the transport MPLS labels. Alternatively, GRE (Generic Routing Encapsulation) can be used for the transport tunnels. In this example, OSPF is configured and LDP SDPs are used.

7750 SR routers are used as ring-nodes to simulate Access Nodes. Ring-nodes can be any L2 device that support VLANs and have the ability to connect an IP interface to one VLAN (required for RNCV).

NTP Configuration

Time must be the same on the redundant NSAs (PE-1 and PE-2); otherwise, the lease times will be different on both nodes. NTP can be used:

```
configure system time
  ntp
    server 10.30.30.30 prefer
    no shutdown
  exit
exit all
```

MC-Ring Configuration

Configure a VPLS service on the CE routers for the In-Band Ring Control connection (BFD packets between PE-1 and PE-2 traversing the ring). This VPLS service will also be used for RNCV. Note that an Epipe service can also be used in case the service is only used for BFD. In that case, a separate service is required for the RNCV.

In this example, QinQ encapsulation is used and BFD and RNCV packets will use VLAN tag 1. Note that dot1q encapsulation can also be used.

The configuration of CE-1 is shown below. A similar configuration is required on the other CE routers.

```

configure
  port 1/1/1
    ethernet
      mode access
      encap-type qinq
    exit
  no shutdown
exit
port 1/1/2
  ethernet
    mode access
    encap-type qinq
  exit
  no shutdown
exit
service
  vpls 1 customer 1 create
    interface "lo1" create
      address 172.16.0.1/24
    exit
    sap 1/1/1:1.0 create
    exit
    sap 1/1/2:1.0 create
    exit
  no shutdown
  exit
exit all

```

An interface **lo1** is created in the VPLS. This interface will be used for RNCV.

On the BSA nodes (PE-1 and PE-2), configure an IES services that will originate BFD and RNCV messages. On PE-1:

```

configure
  port 1/1/1
    ethernet
      mode access
      encap-type qinq
    exit
  no shutdown
exit
service
  ies 1 customer 1 create
    interface "bfd-rncv1" create
      address 172.16.0.251/24
      bfd 100 receive 100 multiplier 3
      sap 1/1/1:1.0 create
    exit
  exit
  no shutdown
  exit
exit all

```

MC-Ring Configuration

On PE-2:

```
configure
  port 1/1/2
    ethernet
      mode access
      encap-type qinq
    exit
  no shutdown
exit
service
  ies 1 customer 1 create
  interface "bfd-rncv1" create
    address 172.16.0.252/24
    bfd 100 receive 100 multiplier 3
    sap 1/1/2:1.0 create
  exit
  exit
  no shutdown
exit
exit all
```

In-Band Ring Control Connection BFD is always originated on an IES or VPRN service on the BSA nodes. RNCV messages can be originated from the same IES/VPRN service or from a separate service.

Configure Multi-Chassis Synchronization (MCS) on the BSA nodes. Enable MCS for **igmp-snooping**, **mc-ring** and **sub-mgmt**. The configuration of PE-1 is shown below. The configuration of PE-2 is similar.

```
configure redundancy multi-chassis peer 192.0.2.2 create
  sync
    igmp-snooping
    mc-ring
    sub-mgmt
    port 1/1/1 sync-tag "l2ring1" create
  exit
  no shutdown
exit
no shutdown
exit all
```

Add the MC-ring configuration on the BSA nodes and link the Ring Control Connection BFD session to the IES service created before.

```
configure redundancy multi-chassis peer 192.0.2.2 create
  mc-ring
    ring "l2ring1" create
    in-band-control-path
      service-id 1
      interface "bfd-rncv1"
      dst-ip 172.16.0.252
    exit
  no shutdown
exit
```

```

        exit
        no shutdown
exit all

```

Note that the ring name is exactly the same as the sync-tag already configured.

At this moment, the MC-ring should be up:

```

A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 detail
=====
Multi-Chassis MC-Ring Detailed Information
=====
Peer           : 192.0.2.2
Ring Type      : Layer 2
Sync Tag       : l2ring1
Port ID        : 1/1/1
Admin State    : inService
Oper State     : connected
Admin Change   : 11/05/2009 21:17:54
Oper Change    : 11/05/2009 21:17:54
Failure Reason : None
Control B Path : No
-----
In Band Control Path
-----
Service ID     : 1
Interface Name : bfd-rncv1
Oper State     : connected
Dest IP        : 172.16.0.252
Src IP         : 172.16.0.251
...

```

Next, configure under MCS the ring nodes on PE-1 and PE-2. This configuration will be used for the RNCV. The ring node configuration for CE-1 is shown below:

```

configure redundancy multi-chassis peer 192.0.2.2 mc-ring ring l2ring1
    ring-node "CE-1" create
        connectivity-verify
        dst-ip 172.16.0.1
        interval 1
        service-id 1
        vlan 1
        no shutdown
    exit
exit
exit all

```

MC-Ring Configuration

Note that the **interval** parameter is the interval used to check the reachability of the CE nodes (configurable from 1 to 6000 minutes). A BFD failure will also be a trigger for a reachability check.

The configuration on PE-2 is identical and the configuration of the other ring nodes is similar..

Configure VLAN 3 to take path-b (default is path-a).

```
configure redundancy multi-chassis peer 192.0.2.2 mc-ring ring l2ring1
    path-b
        range 3-3
    exit
exit all
```

MC-Ring Verification

Verify that MCS is up and running:

```
A:PE-1# show redundancy multi-chassis all
=====
Multi-Chassis Peers
=====
Peer IP          Src IP          Auth          Peer Admin    MC-Ring Oper MC-EP Adm
MCS Admin       MCS Oper       MCS State    MC-LAG Adm    MC-LAG Oper
-----
192.0.2.2       192.0.2.1     None         Enabled       inService    --
Enabled         Enabled       inSync       Disabled      Disabled
=====
```

The output shows that MCS is administrative and operational up and that both peers are synchronized. MC-ring is operational in service.

The following output shows more detailed information about the configured ring:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 detail
=====
Multi-Chassis MC-Ring Detailed Information
=====
Peer           : 192.0.2.2
Ring Type      : Layer 2
Sync Tag       : l2ring1
Port ID        : 1/1/1
Admin State    : inService
Oper State     : connected
Admin Change   : 11/05/2009 21:22:29
Oper Change    : 11/05/2009 21:22:29
Failure Reason : None
Control B Path : No
-----
In Band Control Path
```

```

-----
Service ID       : 1
Interface Name   : bfd-rncv1
Oper State      : connected
Dest IP         : 172.16.0.252
Src IP          : 172.16.0.251
Debounce State  : inService
Debounce Max    : 10 s
Debounce Guard  : 60 s
-----

```

```

-----
VLAN Managed Range
-----

```

```

full range
-----

```

```

VLAN Map B Path Provisioned
-----

```

```

range 3-3
-----

```

```

VLAN Map Excluded Path Provisioned
-----

```

```

no range
-----

```

```

VLAN Map B Path Operational
-----

```

```

range 3-3
-----

```

```

VLAN Map Excluded Path Operational
-----

```

```

no range
-----

```

The output above shows that the ring **l2ring1** on port 1/1/1 is in service and connected. Ring-node Connectivity Check (BFD) is running on **interface bfd-rncv1** in service **1**. All VLANs on port 1/1/1 use path-a (default) except VLAN 3, which uses path-b.

The BSA peer with the lower IP address (the master) will put the SAPs configured for path-a in operational up state while the SAPs configured in path-b will be put in operational down state. The other BSA peer will do the reverse. This is done to prevent loops in the ring.

Note that no VLANs are configured to be excluded from the paths. If a VLAN is configured to be excluded from the paths, the respective SAPs of this VLAN on both BSA nodes will be operationally up.

Check which ring-nodes are configured and connected:

```

A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node
=====

```

```

MC-Ring Node entries
=====

```

Name	Loc Oper St.	Failure Reason
In Use	Rem Oper St.	
CE-1	connected	None
No	notTested	
CE-2	connected	None
No	notTested	
CE-3	connected	None

MC-Ring Configuration

```
      No                notTested
CE-4   connected        None
      No                notTested
```

```
-----
No. of MC-Ring Node entries: 4
=====
```

```
A:PE-1#
```

The output above shows that four ring-nodes are connected. Only the master will send RNCV messages to the ring-nodes. As soon as a ring failure occurs, the BFD session goes down and both BSA nodes send out RNCV messages to see which ring-nodes are connected. Note that when the reachability of the CE nodes is determined, the RNCV messages will be sent at a continuous interval (see above).

More detail about each ring-node can be provided with following command:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node CE-1
detail
```

```
=====
Multi-Chassis MC-Ring Node Detailed Information
=====
```

```
Peer           : 192.0.2.2
Sync Tag       : l2ring1
Node Name      : CE-1
Oper State Loc : connected
Oper State Rem : notTested
In Use        : False
Admin Change   : 11/05/2009 21:21:30
Oper Change    : 11/05/2009 21:22:32
Failure Reason : None
```

```
-----
Ring Node Connectivity Verification
=====
```

```
Admin State    : inService
Service ID     : 1
Encap Value    : 1
Dest IP        : 172.16.0.1
Src IP         : None
Interval       : 1 minutes
Src MAC        : None
```

```
=====
A:PE-1#
```


Unicast Services Configuration

Figure 395 shows the logical setup of the services that will be created. Note that a mesh SDP is used between PE-1 and PE-2. A spoke SDP could also be used.

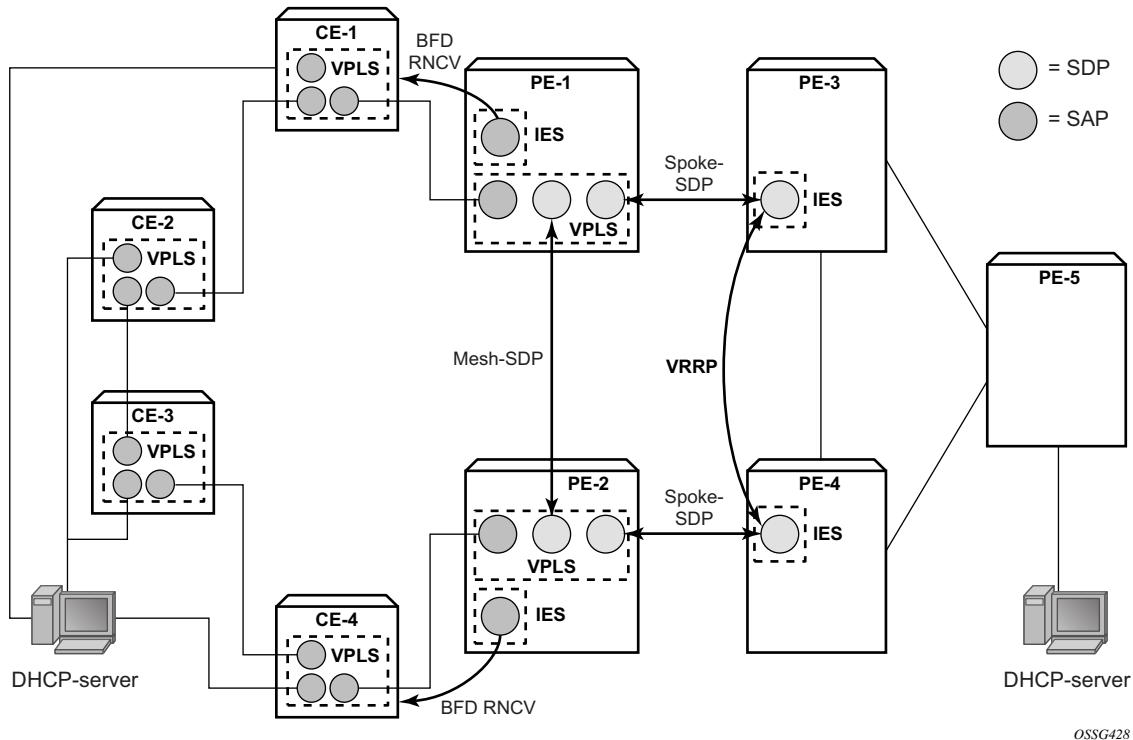


Figure 395: Unicast Services — Logical Topology

In this setup, two unicast services (VPLS 2 and VPLS 3) are created. VPLS 2 uses path-a and VPLS 3 uses path-b.

The services use a mesh SDP between PE-1 and PE-2 and a spoke-SDP between PE-1/PE-3 and another spoke SDP between PE-2/PE-4. On PE-3 and PE-4 a spoke SDP terminated IES is configured with a VRRP default gateway address. The VRRP packets are switched through the BSAs.

IGMP snooping and ESM are configured on both services. ESM is required since the BSA node must know which ring-node each subscriber is connected to. In this setup, the intermediate destination identifier (int_dest_id) will be returned by Option 254 in x Dynamic Host Control Protocol (DHCP). ESM configuration details are outside the scope of this document. Refer to the appropriate platform OS Triple Play Guide.

Unicast Services Configuration

The following output shows the configuration of VPLS 2 on PE-1:

```
configure service
  vpls 2 customer 1 create
    description "VLAN_2"
    split-horizon-group "RSHG" residential-group create
    exit
    sap 1/1/1:2.* split-horizon-group "RSHG" create
      dhcp
        snoop
        lease-populate 10
        no shutdown
      exit
      anti-spoof ip-mac
      sub-sla-mgmt
        def-sub-profile "initial"
        def-sla-profile "initial"
        sub-ident-policy "speedy"
        multi-sub-sap 10
        no shutdown
      exit
    exit
  mesh-sdp 12:2 create
    dhcp
      snoop
    exit
  spoke-sdp 13:2 create
    dhcp
      snoop
    exit
  exit
  no shutdown
exit all
```

The configuration of VPLS 3 is similar. The configuration of VPLS 2 and 3 are similar on PE-2.

The output below shows the subscriber management configuration on PE-1. Similar configuration is required on PE-2.

```
configure subscriber-mgmt
  sla-profile "initial" create
  exit
  sub-profile "initial" create
  exit
  sub-ident-policy "speedy" create
    strings-from-option 254
  exit
exit all
```

Also, configure VLAN 2 and 3 on all the ring-nodes. This configuration is straightforward. Below is an example of the VLAN 2 configuration on CE-1. In this example, the client is connect to port 1/1/3 and should send packets with an outer tag of 2:

```
configure service
  vpls 2 customer 1 create
    sap 1/1/1:2.* create
    exit
    sap 1/1/2:2.* create
    exit
    sap 1/1/3:2.* create
    exit
    no shutdown
  exit
exit all
```

In the example, VLAN 2.* and 3.* is used to allow for transparent transport of the customer VLAN.

The IES service where VPLS 2 terminates looks like this on BSR PE-3:

```
configure service
  ies 2 customer 1 create
    interface "VLAN_2" create
      address 10.0.2.3/24
      dhcp
        server 10.10.10.10
        trusted
        no shutdown
      exit
      ip-mtu 1500
      vrrp 2
        backup 10.0.2.254
        ping-reply
      exit
      local-proxy-arp
      spoke-sdp 31:2 create
      exit
    exit
    no shutdown
  exit
exit all
```

And on PE-4:

```
configure service
  ies 2 customer 1 create
    interface "VLAN_2" create
      address 10.0.2.4/24
      dhcp
        server 10.10.10.10
        trusted
        no shutdown
      exit
      ip-mtu 1500
      vrrp 2
```

Unicast Services Configuration

```
        backup 10.0.2.254
        ping-reply
    exit
    local-proxy-arp
    spoke-sdp 42:2 create
    exit
exit
no shutdown
exit
exit all
```

Notice that the ip-mtu must be set to match the vc-mtu signaled by the other side of the spoke-SDP. Otherwise, the service will be operationally down with a ServiceMTUMismatch.

Note also in the configuration that DHCP relay is done by configuring a DHCP server under the IES interface.

The configuration of IES 3 on PE-3 and PE-4 is similar..

On PE-5 an IES interface is configured to the DHCP-server. The interface is configured with a Dot1Q encapsulated port because this port will be also be used for an interface to the multicast server.

```
configure service
    ies 2 customer 1 create
        interface "dhcp-server" create
            address 192.168.6.1/30
            sap 1/1/3:2 create
        exit
    exit
no shutdown
exit
exit all
```

Unicast Services Verification

Request an IP address on VLAN 2 on CE-1. On BSA routers PE-1 and PE-2 following DHCP info can be checked:

```
A:PE-1# show service id 2 dhcp lease-state
=====
DHCP lease state table, service 2
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
                  LifeTime        Origin          Stdbby
-----
10.0.2.107      00:00:64:01:01:02 1/1/1:2.*      09d07h38m  DHCP
-----
Number of lease states : 1
=====
```

ESM show commands can be used to obtain the subscriber identity, IP address, MAC address and SLA-profile:

```
A:PE-1# show service active-subscribers
=====
Active Subscribers
=====
Subscriber subscriber_1_vlan_2 (initial)
-----
(1) SLA Profile Instance sap:1/1/1:2.* - sla:initial
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
10.0.2.107      00:00:64:01:01:02 N/A      DHCP
-----
Number of active subscribers : 1
=====
```

The following command gives more information about a specific subscriber:

```
A:PE-1# show service active-subscribers subscriber subscriber_1_vlan_2 detail
=====
Active Subscribers
=====
Subscriber subscriber_1_vlan_2 (initial)
-----
I. Sched. Policy : N/A
E. Sched. Policy : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A
Rad. Acct. Pol.   : N/A
Dupl. Acct. Pol.  : N/A
E. Agg Rate Limit: Max
Collect Stats     : Disabled
```

Unicast Services Configuration

```
ANCP Pol.          : N/A
HostTrk Pol.       : N/A
Sub. ANCP-String   : "subscriber_1_vlan_2"
Sub. Int Dest Id   : "CE-1"
Host Trk Rate Adj  : N/A
```

```
-----
(1) SLA Profile Instance
    - sap:1/1/1:2.* (VPLS 2)
    - sla:initial
-----
```

```
-----
Description        : (Not Specified)
Host Limit         : No Limit
Ingress Qos-Policy : 1                      Egress Qos-Policy : 1
Ingress Queuing Type : Service-queuing
Ingress Filter-Id  : N/A                    Egress Filter-Id  : N/A
Ingress Report-Rate : N/A
Egress Report-Rate : N/A
Egress Remarking   : from Sap Qos
Credit Control Pol. : N/A
-----
```

```
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
10.0.2.107     00:00:64:01:01:02 N/A           DHCP
-----
```

The output above gives more details about which ring-node the customer is connected to (Sub. Int Dest Id : CE-1), which QoS policies are applied, statistics of each queue,

MCS Verification

Check if the two redundant peers are in sync and check the detailed MCS info:

```
A:PE-1# show redundancy multi-chassis sync peer 192.0.2.2
```

```
=====
Multi-chassis Peer Table
=====
```

```
Peer
```

```
-----
Peer IP Address      : 192.0.2.2
Description          : (Not Specified)
Authentication       : Disabled
Source IP Address    : 192.0.2.1
Admin State          : Enabled
-----
```

```
Sync-status
```

```
-----
Client Applications  : IGMP Snooping SUBMGMT RING
Sync Admin State     : Up
Sync Oper State      : Up
DB Sync State        : inSync
Num Entries          : 11
Lcl Deleted Entries  : 0
Alarm Entries        : 0
Rem Num Entries      : 11
Rem Lcl Deleted Entries : 0
-----
```

```

Rem Alarm Entries      : 0
=====
MCS Application Stats
=====
Application           : igmp
Num Entries           : 0
Lcl Deleted Entries   : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : igmpSnooping
Num Entries           : 0
Lcl Deleted Entries   : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : subMgmt
Num Entries           : 1
Lcl Deleted Entries   : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : srrp
Num Entries           : 0
Lcl Deleted Entries   : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : mcRing
Num Entries           : 10
Lcl Deleted Entries   : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 10
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : mldSnooping
Num Entries           : 0
Lcl Deleted Entries   : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : dhcpServer

```

Unicast Services Configuration

```
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
-----
Application          : subHostTrk
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
-----
=====
A:PE-1#
```

The output shows that both MCS peers are in sync and that entries are populated for MC-ring and Subscriber Management (DHCP lease states).

Notice that the lease states are also populated on PE-2:

```
A:PE-2# show service id 2 dhcp lease-state
=====
DHCP lease state table, service 2
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining      Lease      MC
                LifeTime         Origin          Stdby
-----
10.0.2.107      00:00:64:01:01:02 1/1/2:2.*      09d06h16m     DHCP      Yes
-----
Number of lease states : 1
=====
A:PE-2#
```

The output is similar to the output on PE-1 except that on PE-2 the flag MC-Stdby is set to yes, which implies that this node is in standby mode for this VLAN.

This can be verified by looking at the status of the SAP on PE-1 and PE-2. On PE-1 the SAP is operationally up:

```
A:PE-1# show service id 2 sap 1/1/1:2.*
=====
Service Access Points(SAP)
=====
Service Id      : 2
SAP             : 1/1/1:2.*          Encap           : qinq
QinQ Dot1p     : Default
Description     : (Not Specified)
Admin State     : Up                Oper State      : Up
Flags           : None
Multi Svc Site : None
Last Status Change : 11/06/2009 17:17:26
```



```
Last Mgmt Change   : 11/04/2009 22:51:15
```

```
=====
```

On PE-2 the situation is different:

```
A:PE-2# show service id 2 sap 1/1/2:2.*
```

```
=====
```

Service Access Points(SAP)

```
=====
```

Service Id	: 2		
SAP	: 1/1/2:2.*	Encap	: qinq
QinQ Dot1p	: Default		
Description	: (Not Specified)		
Admin State	: Up	Oper State	: Down
Flags	: StandByForMcRing		
Multi Svc Site	: None		
Last Status Change	: 11/06/2009 18:07:30		
Last Mgmt Change	: 11/04/2009 23:43:16		

```
=====
```

Notice that the SAP on PE-2 is operationally down and that a flag is set: StandByForMcRing.

The situation is reversed for VLAN 3 since it was configured for path-b; here the SAP on PE-1 is operationally down and the SAP on PE-2 is operationally up.

Ring Failure Verification

In case a ring failure occurs (either ring link failure or ring-node failure), the IB-RCC BFD session between PE-1 and PE-2 will go down and both nodes will put the SAP in operational up state.

Break the link between CE-2 and CE-1.

Observe that the ring is now in a **broken** state:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2
```

```
=====
```

MC-Ring entries

```
=====
```

Sync Tag	Oper State	Failure Reason
l2ring1	broken	None

```
=====
```

```
No. of MC-Ring entries: 1
```

```
=====
```

Unicast Services Configuration

The following command shows the BSA to ring nodes connections:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node
=====
MC-Ring Node entries
=====
Name                               Loc Oper St.      Failure Reason
  In Use                            Rem Oper St.
-----
CE-1                               connected         None
  Yes                               disconnected
CE-2                               disconnected       None
  No                                connected
CE-3                               disconnected       None
  No                                connected
CE-4                               disconnected       None
  No                                connected
-----
No. of MC-Ring Node entries: 4
=====
```

The output shows that CE-1 is connected to PE-1. CE-2, CE-3 and CE4 are connected to PE-2.

Notice that the SAPs on both PE-1 and PE-2 are now operationally up. This can be checked with the **show service id 2 sap 1/1/1:2.*** command on PE1 and the **show service id 2 sap 1/1/2:2.*** command on PE2.

Following the ring failure, the BSA nodes send a MAC address withdrawal message to all the SDP peers configured in the VPLS.

Multicast Service Configuration

In general, BTV (Broadcast TV) services are delivered to the DSLAM on a different VLAN using a different VPLS service. The DSLAM can send/relay IGMP group membership messages if it wants to receive a multicast stream. At the BSA level (PE-1 and PE-2), IGMP snooping is configured. At the BSR level (PE-3, PE-4) and the PE where the video source is located, PIM is configured and IGMP is configured on the IES service facing the BSA ring. The BSA ring consists of concatenated spoke SDPs. The spoke SDP ring is not closed between PE-3 and PE-4 to avoid a loop. [Figure 396](#) shows the logical setup for the multicast service with a MC-ring.

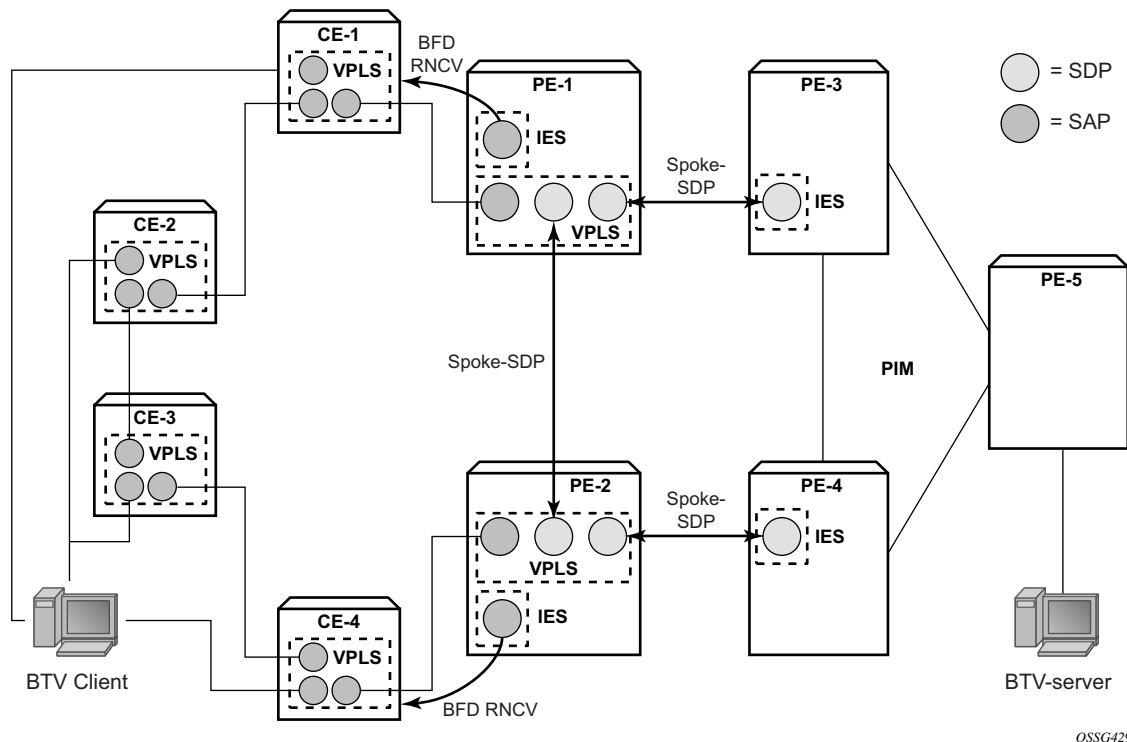


Figure 396: Multicast Service — Logical Setup

Configure an interface to the multicast server on PE-5:

```
configure service
  ies 3 customer 1 create
    interface "mcast-server" create
      address 192.168.7.1/30
      sap 1/1/3:3 create
    exit
  exit
no shutdown
```

Multicast Service Configuration

```
        exit
exit all
```

Configure a VPLS service on PE-1 and PE-2. The configuration on PE-1 is shown below. The VPLS configuration on PE-2 is similar.

```
configure service
  vpls 4 customer 1 create
    description "Multicast VLAN"
    igmp-snooping
      no shutdown
    exit
    sap 1/1/1:4.* create
    exit
    spoke-sdp 12:4 create
      igmp-snooping
        mrouter-port
      exit
    exit
    spoke-sdp 13:4 create
      igmp-snooping
        mrouter-port
      exit
    exit
    no shutdown
  exit
exit all
```

Notice that igmp-snooping has been enabled and that the spoke SDPs are configured as mrouter-ports in order to forward the IGMP joins to the BSRs.

Configure a spoke SDP terminated IES service on PE-3:

```
configure service
  ies 4 customer 1 create
    interface "btv-dst" create
      address 10.0.4.3/24
      ip-mtu 1500
      spoke-sdp 31:4 create
    exit
  exit
  no shutdown
exit
exit all
```

On PE-4:

```
configure service
  ies 4 customer 1 create
    interface "btv-dst" create
      address 10.0.4.4/24
      ip-mtu 1500
      spoke-sdp 42:4 create
    exit
  exit
  no shutdown
```

```

        exit
    exit all

```

Notice that also here the **ip-mtu** must be configured to bring the service up. The **ip-mtu** is required to have an MTU match on the spoke SDP between the IES service and the VPLS service.

Configure PIM on PE-3, PE-4 and PE-5. The configuration on PE-3 is shown below. The PIM configuration on PE-4 and PE-5 is similar.

```

configure router pim
    interface "system"
    exit
    interface "int-PE-3-PE-4"
        priority 10
    exit
    interface "int-PE-3-PE-5"
    exit
    interface "btv-dst"
    exit
    rp
        static
            address 192.0.2.5
            group-prefix 224.0.0.0/4
        exit
    exit
exit all

```

Note that PE-5 is statically configured as the RP. This is just an example. Different configurations can be used.

Configure IGMP on PE-3/PE-4:

```

configure router igmp interface btv-dst no shutdown

```

The service should also be configured on all ring nodes. Below, the configuration on CE-2 is shown. Similar configurations are required on the other ring-nodes.

```

configure service
    vpls 4 customer 1 create
        sap 1/1/1:4.* create
        exit
        sap 1/1/2:4.* create
        exit
        sap 1/1/3:4.* create
        exit
        no shutdown
    exit
exit all

```

Multicast Service Verification

Configure the multicast server to send one or more multicast streams. Have the BTV client connected to CE-2 send an IGMP join message for this multicast stream.

On the BSA routers (PE-1/PE-2), IGMP snooping can be checked:

```
A:PE-1# show service id 4 mfib
=====
Multicast FIB, Service 4
=====
Source Address  Group Address      Sap/Sdp Id          Svc Id  Fwd/Blk
-----
*                *                  sdp:12:4            Local    Fwd
                  sdp:13:4            Local    Fwd
*                225.1.1.1         sap:1/1/1:4.*      Local    Fwd
                  sdp:12:4            Local    Fwd
                  sdp:13:4            Local    Fwd
-----
Number of entries: 2
=====
```

On PE-3/PE-4 the IGMP groups can be checked:

```
A:PE-3# show router igmp group
=====
IGMP Groups
=====
(*,225.1.1.1)                               Up Time : 0d 00:01:03
  Fwd List  : btv-dst
-----
(*,G)/(S,G) Entries : 1
=====
```

Following command shows that the MCS peers are synchronized and that there is one IGMP entry on both peers:

```
A:PE-1# show redundancy multi-chassis sync peer 192.0.2.2 detail
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 192.0.2.2
Description          : (Not Specified)
Authentication       : Disabled
Source IP Address    : 192.0.2.1
Admin State          : Enabled
-----
Sync-status
-----
Client Applications  : IGMP Snooping SUBMGMT RING
Sync Admin State     : Up
Sync Oper State      : Up
DB Sync State        : inSync
```

```

Num Entries          : 12
Lcl Deleted Entries  : 0
Alarm Entries        : 0
Rem Num Entries      : 12
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0

```

```

=====
MCS Application Stats
=====

```

```

Application          : igmp
Num Entries          : 1
Lcl Deleted Entries  : 0
Alarm Entries        : 0

```

```

-----
Rem Num Entries      : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0

```

Notice that the MFIB on PE-2 has also been updated:

```

A:PE-2# show service id 4 mfib

```

```

=====
Multicast FIB, Service 4
=====

```

Source Address	Group Address	Sap/Sdp Id	Svc Id	Fwd/Blk
*	*	sdp:21:4	Local	Fwd
		sdp:24:4	Local	Fwd
*	225.1.1.1	sap:1/1/2:4.*	Local	Fwd
		sdp:21:4	Local	Fwd
		sdp:24:4	Local	Fwd

```

-----
Number of entries: 2

```

```

=====
A:PE-2#

```

The SAP on PE-2 is down to avoid duplicated traffic and loops:

```

A:PE-2# show service id 4 sap 1/1/2:4.*

```

```

=====
Service Access Points(SAP)
=====

```

```

Service Id          : 4
SAP                 : 1/1/2:4.*          Encap           : qinq
QinQ Dot1p         : Default
Description         : (Not Specified)
Admin State         : Up                Oper State       : Down
Flags               : StandByForMcRing
Multi Svc Site      : None
Last Status Change  : 11/06/2009 18:56:25
Last Mgmt Change    : 11/04/2009 23:43:16

```

If a failure occurs in the ring-node, the IB-RCC BFD session between PE-1 and PE-2 will go down and the SAP on both PE-1 and PE-2 will be put in operational up state.

Configuration Notes

RNCV (used for ring-node connectivity check) and BFD (used for ring control connection) can either run on the same VLAN in the same IES or VPRN service or can run on different VLAN.

MCS for IGMP or DHCP states on a MC-ring requires ESM since the BSA nodes must know which ring node a subscriber is connected to in case a ring failure occurs. The ring-node name is returned through a Python script, a RADIUS server or through a local user database. This string (`int_dest_id`) must match one of the ring nodes defined in the redundancy configuration.

Convergence time after a ring failure should be $3 * \text{BFD timer} + \text{MCS convergence time}$.
Convergence time after a BSA failure should be likewise.

Note that a debounce timer runs on the MC-ring peers. After a ring failure, the MC-ring converges immediately (after the BFD session times out) and the debounce timer is started. After the ring is fixed and the BFD session is up the MC-ring converges immediately again. If another ring failure occurs before the debounce timer expires, convergence will be slowed down by two (2) seconds. If a third ring failure occurs before the debounce timer expires, four second delays are introduced. In case of a fourth failure, an eight second delay is introduced. 200 seconds of delay is the maximum. The debounce timer can be configured under the MC-ring.

Conclusion

This section covers an extension of dual homing support in TPSDA networks based on Layer 2 CO model. The extension addresses networks where multiple access nodes (for example, DSLAMs) are connected in a single ring. The examples show the use of a ring with four access nodes in a ring. The behavior is described in normal operation and in case a failure occurs in the access ring.

Conclusion