Troubleshoot Different ESI Load Balancing Modes

Contents

Introduction

This document describes different Ethernet Segment Identifier (ESI) load balancing modes which is essential for optimizing traffic distribution and maintaining network reliability.

The different load balancing modes include

- Port-Active
- Single-Active
- All-Active

- Single-Flow Active

Prerequisites

Requirements

Cisco recommends that you have basic knowledge of :

- Multiprotocol Label Switching (MPLS).
- Layer 2 Virtual Private Network (L2VPN).
- Ethernet Virtual Private Network (EVPN).

Components Used

- The information in this document is based on Device: Aggregation Services Router 9000 (ASR9K).
- The information in this document was created from the devices in a specific lab environment. All the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Network Topology

Using Figure 1 Network Topology diagram to illustrate the working of different load balancing modes.

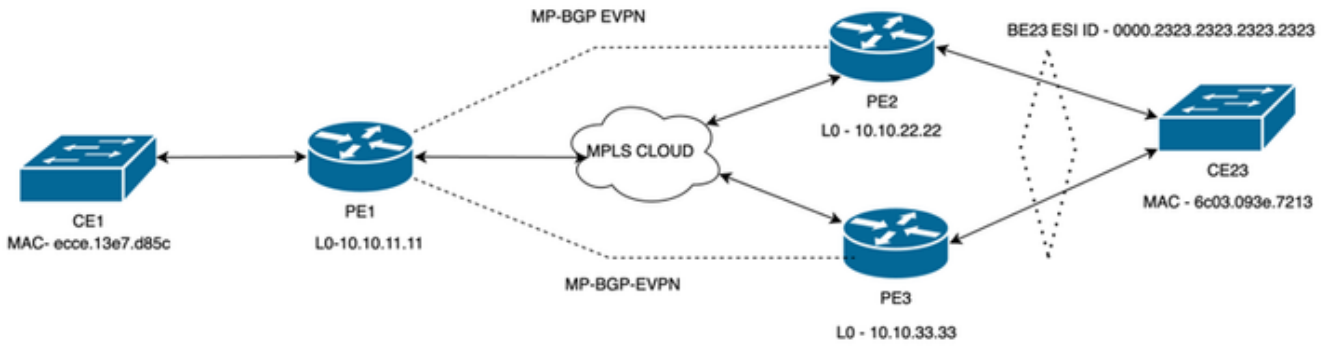

Figure 1 Network Topology

Port Active

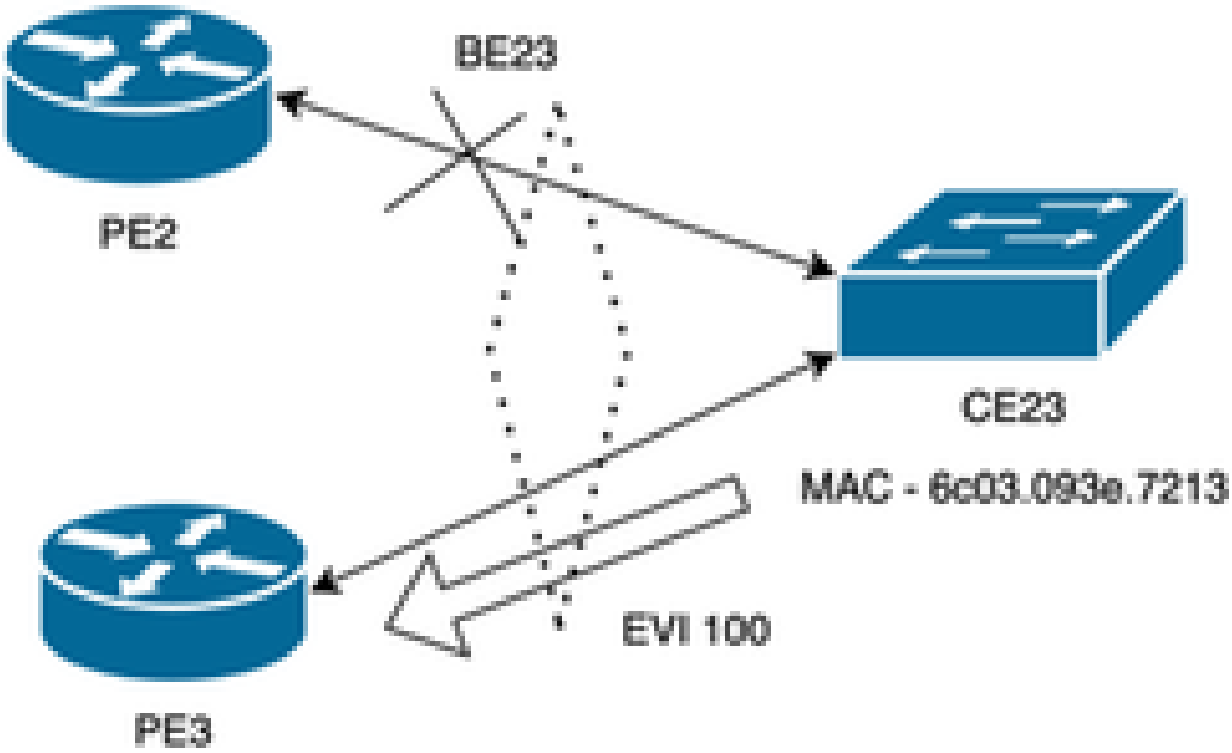

Figure 2 Port-Active Redundancy Mode

Figure 2 Port active Redundancy Mode describes a redundancy mode in which active/standby redundancy is configured at the interface level. In this setup, traffic for an Ethernet Virtual Instance (EVI), such as EVI 100, is forwarded only through the active interface (in this case, interface connected to Provider Edge -PE3). All traffic destined for EVI 100 is hashed onto the active router (PE3), while the standby router(s) (PE2) remains ready to take over in case the active interface fails.

Configuration

<#root>

```
evpn
 interface Bundle-Ether23
  ethernet-segment
   identifier type 0 00.23.23.23.23.23.23.23.23
```

load-balancing-mode port-active >> configuration required to enable this mode

Verification Of Port-Active ESI

Interface status

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show interfaces BE23
Bundle-Ether23 is down, line protocol is down
```

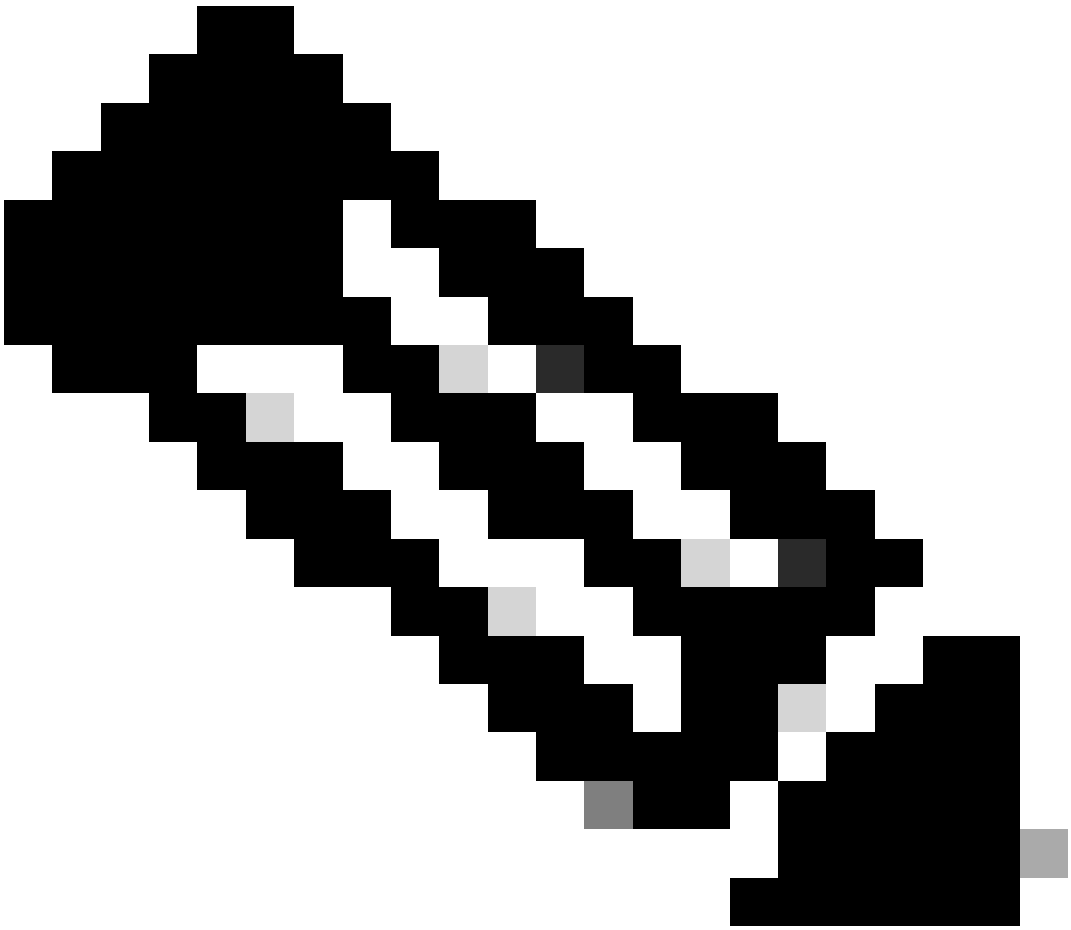

Note: On the standby PE (PE2) the Bundle ether (BE) interface status is down.

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show interfaces BE23
Bundle-Ether23 is up, line protocol is up
```

Bundle status

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show bundle bundle-ether 23
Bundle-Ether23
  Status:                                  LACP OOS (out of service)
<snip>
Port                  Device           State        Port ID         B/W, kbps
  --------------------  ---------------  -----------  --------------  ----------
  Hu0/1/0/7             Local
```

Standby

```
      0x8000, 0x0001   100000000
      Link is in standby due to bundle out of service state
```

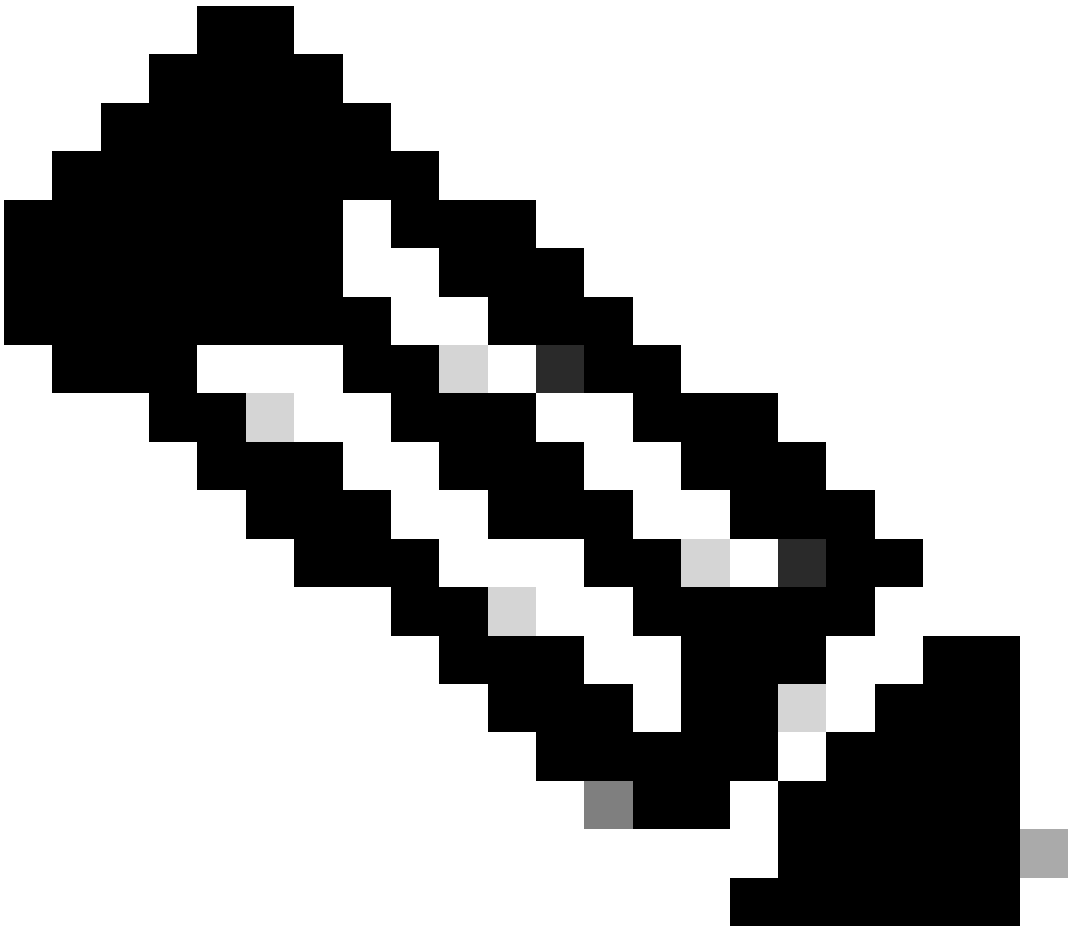

Note:

- Interface BE23 on PE2 is down due to “bundle out of service state” as this PE2 is in standby state once the BE23 on PE3(Active PE) goes down BE23 on PE2 becomes active .
- By default Bundle status is "Out of service" but you can set the access signal to down with explicit configuration "**access-signal bundle-down**" command under ethernet-segment.

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show bundle bundle-ether 23
Bundle-Ether23
  Status:                             Up
<snip>
  Port                Device           State        Port ID         B/W, kbps
  ------------------  ---------------  -----------  --------------  ----------
  Hu0/1/0/7           Local            Active       0x8000, 0x0001   100000000
```

```
      Link is Active
```

The active interface on which all the traffic must be hashed from the Customer Edge (CE) 23.

Ethernet Segment

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
  ES to BGP Gates   : Ready
  ES to L2FIB Gates : Ready
  Main port         :
     Interface name : Bundle-Ether23
     Interface MAC  : 08ec.f50e.6af6
     IfHandle       : 0x040001a0
     State          : Standby
     Redundancy     : Not Defined
  ESI type          : 0
     Value          : 0000.2323.2323.2323.2323
  ES Import RT      : 0023.2323.2323 (from ESI)
  Source MAC        : 0000.0000.0000 (N/A)
  Topology          :
     Operational    : MH
     Configured     : Port-Active
  <snip>
  Service Carving Results:
     Forwarders     : 2
     Elected        : 0

Not Elected    : 2

            EVI NE  :      100,     200

  <snip>
  Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.33.33

Access signal mode: Bundle OOS
```

PE2 - Non-Designated Forwarder router implements a directional blocking for all the traffic coming from and towards the CE23.

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
  Main port         :
     Interface name : Bundle-Ether23
     Interface MAC  : 08ec.f52e.55b5
     IfHandle       : 0x000001a0
     State          : Up
     Redundancy     : Not Defined
<snip>
  Topology          :
     Operational    : MH
     Configured     : Port-Active
  <snip>
  Service Carving Results:
     Forwarders     : 2


Elected        : 2


            EVI E   :      100,     200


     Not Elected    : 0
  <snip>
  HRW Reset timer   : 5 sec [not running]
  Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.22.22
  Access signal mode: Bundle OOS (Default)
```

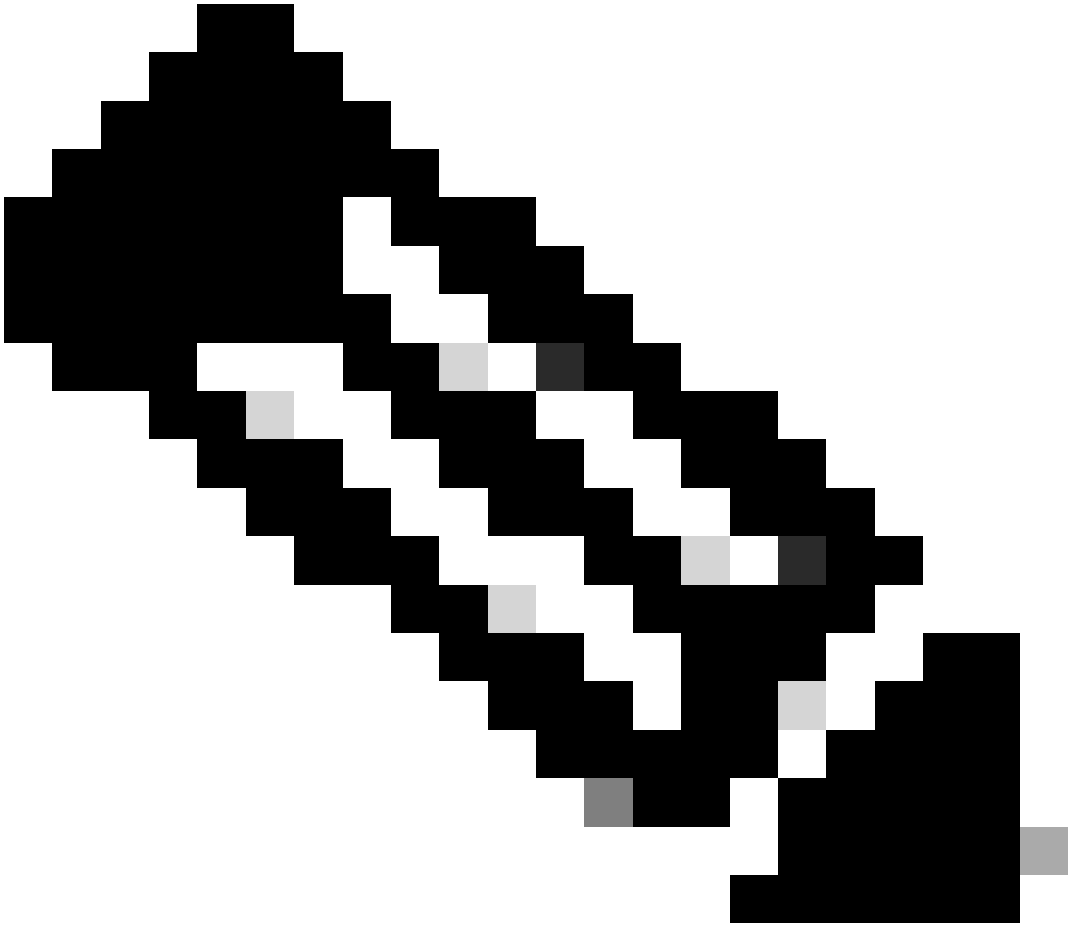

Note: PE3 is elected as the Designated Forwarder(DF) for this Ethernet Segment , all the traffic towards the CE23 is sent and received via PE3.

Routes advertised to remote router PE1 from PE2 and PE3

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show bgp l2vpn evpn
Route Distinguisher: 10.10.11.11:100 (default for vrf EVPN-ELAN-1)
*>i[1][0000.2323.2323.2323.2323][0]/120 >>>>>>>>>>>>>>>>>>>>>>>>> per ESI route
                     10.10.22.22                   100      0 i
* i                   10.10.33.33                   100      0 i
*>i[1][0000.2323.2323.2323.2323][4294967295]/120 >>>>>>>>>>>>>>>> per EVI route
                     10.10.22.22                   100      0 i
* i                   10.10.33.33                   100      0 i
*>i
```

[2][0][48][6c03.093e.7213][0]/104 >>>>>>>>>>>>>>>>>>>>>>>>>>>>>> Route-Type 2 for the CE23 mac address is

```
                     10.10.33.33                   100      0 i
```

```
*> [2][0][48][ecce.13e7.d85c][0]/104
                      0.0.0.0                                0 i
```

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show evpn internal-label vpn-id 100 detail
VPN-ID     Encap  Ethernet Segment Id         EtherTag     Label
---------- ------ -------------------------- ----------   --------
100        MPLS   0000.2323.2323.2323.2323    0            24010
   Multi-paths resolved: TRUE (Remote single-active)
   Multi-paths Internal label: 24010
    MAC              10.10.33.33                            24001
    EAD/ES           10.10.22.22                              0
                     10.10.33.33                              0
    EAD/EVI          10.10.22.22                            24001
                     10.10.33.33                            24001
   Summary pathlist:

 0x02000002 (P) 10.10.33.33                                 24001


 0x00000000 (B) 10.10.22.22                                 24001
```

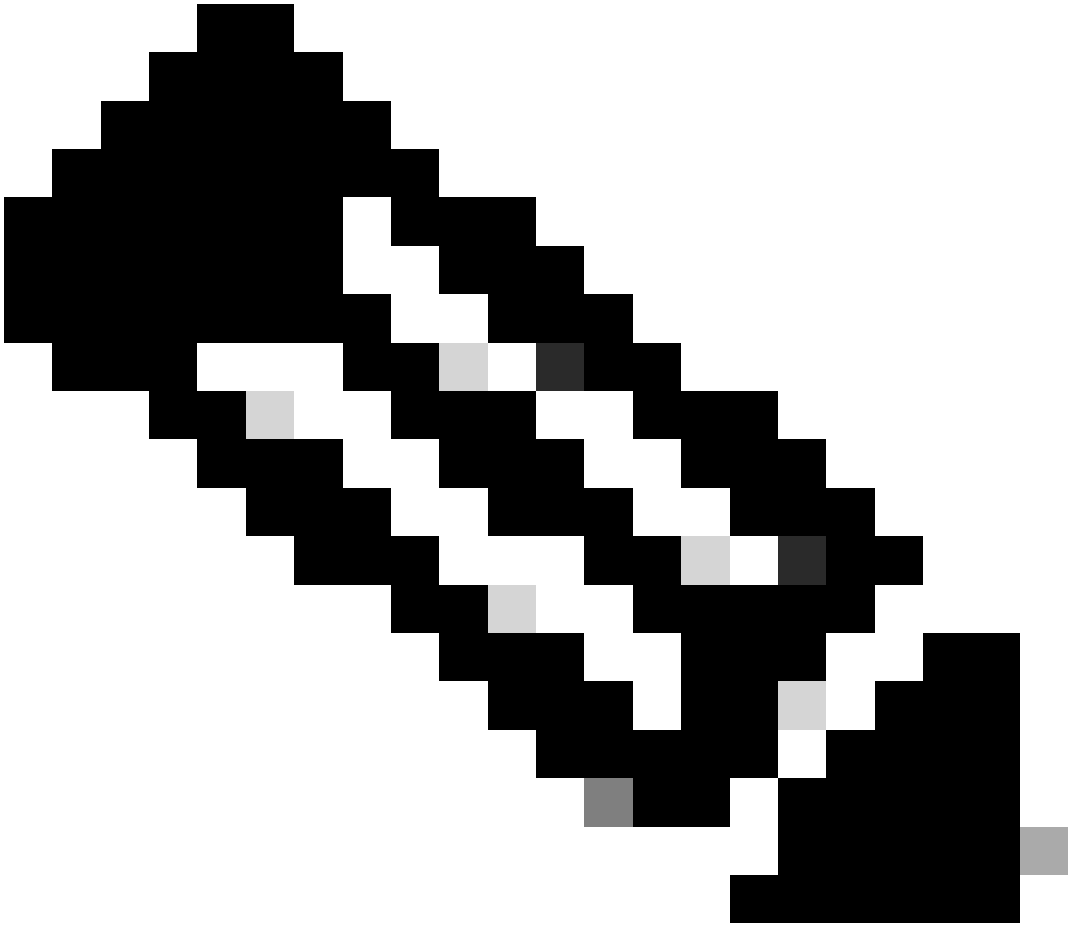

Note: The Route-Type 2 of CE23 Is only advertised by PE3 hence PE3 is set as primary for this Ethernet Segment and PE2 as secondary .

Mac address learning

L2vpn

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show l2vpn forwarding bridge-domain EVPN-ELAN-1:EVPN-ELAN-1 mac-address lo
Mac Address    Type    Learned from/Filtered on    LC learned Resync Age/Last Change Mapped to
-------------- ------- -------------------------- ---------- ---------------------- --------------
ecce.13e7.d85c dynamic Hu0/1/0/2.100              N/A        24 Dec 08:58:17        N/A

6c03.093e.7213 EVPN    BD id: 1                   N/A        N/A                    N/A
```

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show l2vpn forwarding bridge-domain  EVPN-ELAN:EVPN-ELAN mac-address loca
```

```
Mac Address    Type    Learned from/Filtered on    LC learned Resync Age/Last Change Mapped to
-------------- ------- --------------------------- ---------- ---------------------- --------------

6c03.093e.7213 EVPN    BD id: 0                    N/A        N/A                    N/A

ecce.13e7.d85c EVPN    BD id: 0                    N/A        N/A                    N/A
```

PE2 learns both the mac addresses of CE1 and CE23 as EVPN routes from PE1 and PE3 .

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show l2vpn forwarding bridge-domain  EVPN-ELAN:EVPN-ELAN mac-address locat
Mac Address    Type    Learned from/Filtered on    LC learned Resync Age/Last Change Mapped to
-------------- ------- --------------------------- ---------- ---------------------- --------------
ecce.13e7.d85c EVPN    BD id: 0                    N/A        N/A                    N/A

6c03.093e.7213 dynamic BE23.100                    N/A        24 Dec 07:26:58        N/A
```

L2rib

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show l2route evpn mac all
Topo ID  Mac Address    Producer    Next Hop(s)
-------- -------------- ----------- ----------------------------------------
1

6c03.093e.7213 L2VPN       24010/I/ME, N/A

1        ecce.13e7.d85c LOCAL       HundredGigE0/1/0/2.100, N/A
```

EVPN:

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show evpn evi vpn-id 100  mac
VPN-ID     Encap      MAC address    IP address                              Nexthop
---------- ---------- -------------- --------------------------------------- ------------------------

100        MPLS       6c03.093e.7213 ::                                      10.10.33.33

100        MPLS       ecce.13e7.d85c ::                                      HundredGigE0/1/0/2.100
```

BGP:

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show bgp l2vpn evpn rd 10.10.11.11:100 [2][0][48][6c03.093e.7213][0]/104
Local
```

```
    10.10.33.33 (metric 10) from 10.10.33.33 (10.10.33.33)
      Received Label 24001
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported, rib-ins
      Received Path ID 0, Local Path ID 1, version 1321
      Extended community:
```

SoO:10.10.33.33:100

```
 0x060e:0000.0000.0064 RT:100:100
      EVPN ESI: 0000.2323.2323.2323.2323
      Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.10.33.33:100
```

In this mode when traffic from CE1 has to be sent to CE23, from PE1 it is only aliased to PE3 as we receive Route-Type 2 only from PE3 . When BE23 goes down on PE3, the traffic is shifted to PE2 .

Single-Active

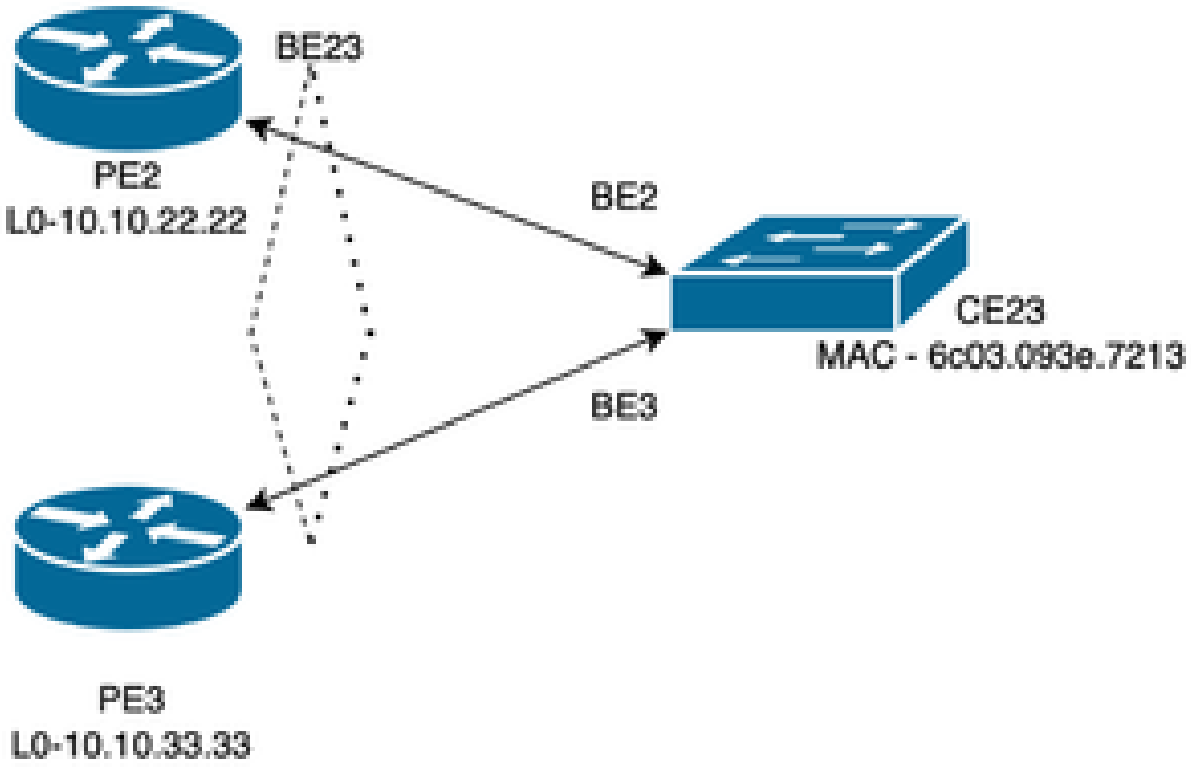

Figure 3 Single-Active Redundancy Mode

As seen in the Figure 3 Single-Active Redundancy Mode , In this redundancy mode, both PE access interfaces remain active, Each link towards the PE is assigned a unique Ethernet Bundle interface on the CE23, with VLANs 100 and 200 permitted on both the interfaces . Since these links belong to separate Ethernet Bundles, CE23 initially floods traffic to both PEs. However, only the Designated Forwarder (DF) of the Ethernet Segment forwards traffic toward the core. Consequently, the CE maintains a single Ethernet Bundle interface in its forwarding table, ensuring per-VLAN single-active multi-homing.

Configuration

<#root>

```
interface Bundle-Ether23
  ethernet-segment
   identifier type 0 00.23.23.23.23.23.23.23.23
```

load-balancing-mode single-active >> configuration required to enable this mode

Verification Of Single-Active ESI

Interface status

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show interfaces bundle-ether 23
Bundle-Ether23 is up, line protocol is up

RP/0/RSP0/CPU0:ASR9910-3-PE3#show interfaces bundle-ether 23
Bundle-Ether23 is up, line protocol is up
```

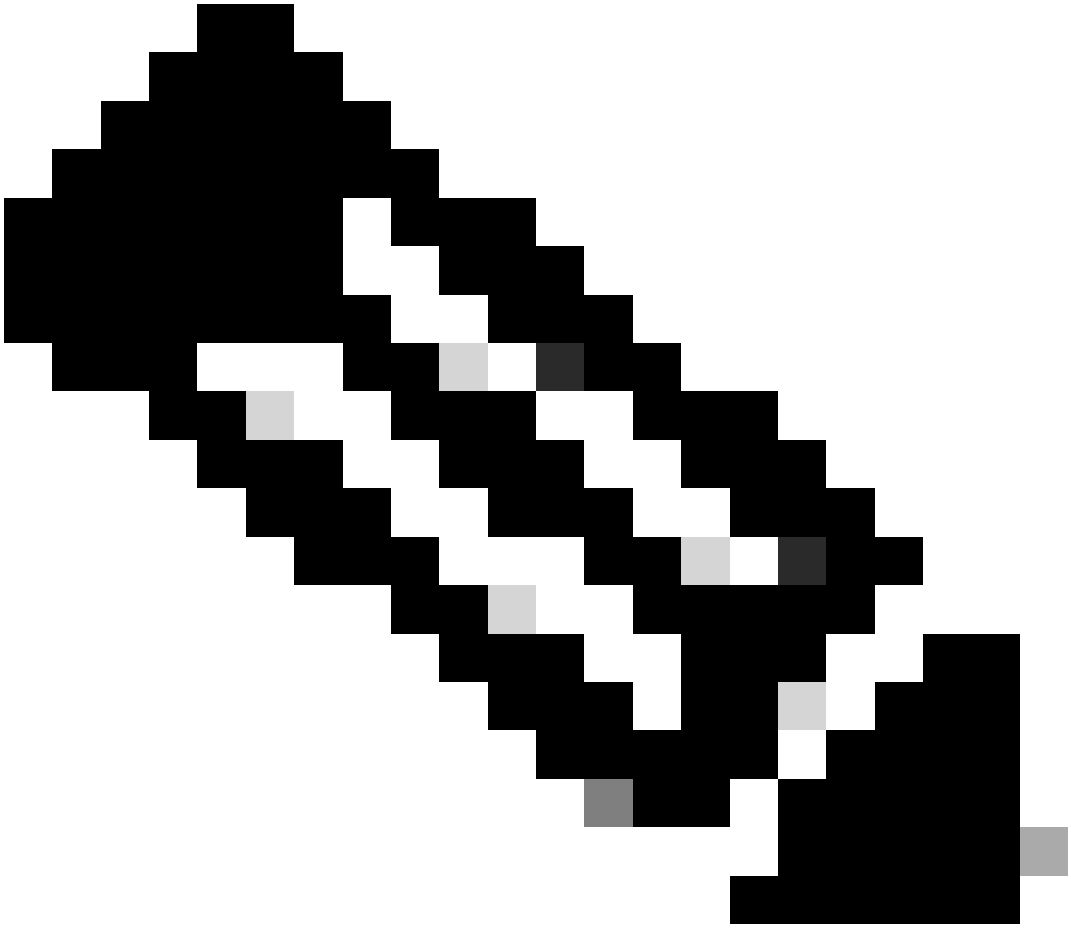

Note: The physical status and the bundle status of BE23 on PE2 and PE3 are up and active.

Ethernet Segment

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
  <snip>
  Topology          :
     Operational    : MH, Single-active
     Configured     : Single-active (AApS)
  <snip>
  Service Carving Results:
     Forwarders     : 2

Elected        : 1

            EVI E   :       200

     Not Elected    : 1
  <snip>
  Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.33.33
  Access signal mode: Bundle OOS

RP/0/RSP0/CPU0:ASR9910-3-PE3#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
<snip>
Topology          :
     Operational    : MH, Single-active
     Configured     : Single-active (AApS)
<snip>
Service Carving Results:
     Forwarders     : 2

Elected        : 1

             EVI E  :      100

     Not Elected    : 1
<snip>
  Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.22.22
```

```
  Access signal mode: Bundle OOS (Default)
```

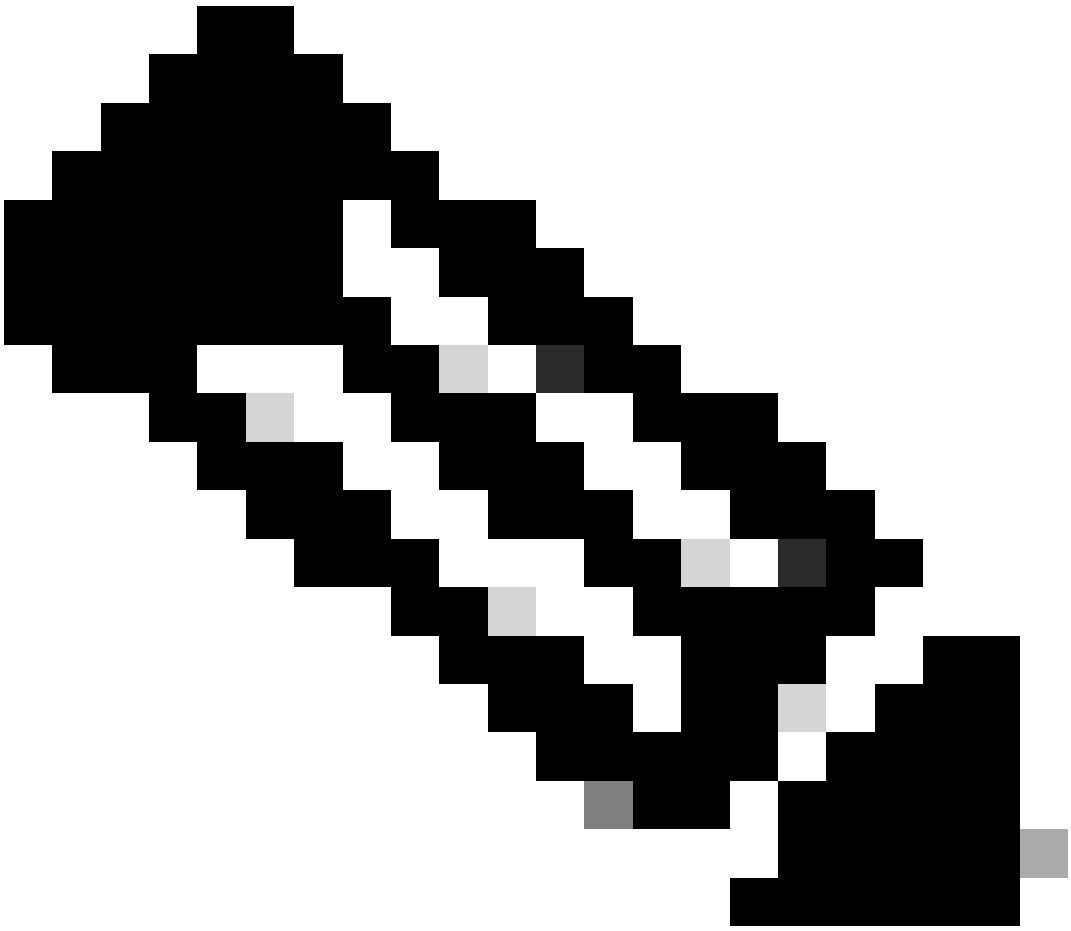

Note:

- PE2 is elected as DF for EVI 200 .
- PE3 is elected as DF for EVI 100 .
- For EVI 100 and EVI 200 the remote PE (PE1) can send and receive traffic only via PE3 and PE2 respectively.

Routes advertised to PE1 by PE2 and PE3

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show bgp l2vpn evpn
Route Distinguisher: 10.10.11.11:
```

100

```
 (default for vrf EVPN-ELAN-1)
```

```
*>i[1][0000.2323.2323.2323.2323][0]/120>>>>>>>>>>>>>>>>>>>>>>>>>>> Per EVI route
                     10.10.22.22                   100      0 i
* i                   10.10.33.33                   100      0 i
*>i[1][0000.2323.2323.2323.2323][4294967295]/120>>>>>>>>>>>>>>>>> per ESI route
                     10.10.22.22                   100      0 i
* i                   10.10.33.33                   100      0 i
*>i
```

[2][0][48][6c03.093e.7213][0]/104>>>>>>>>>>>>>>>>>>>>>>>>>>>>>> Route-type 2 Advertised only by PE3 for

10.10.33.33 100 0 i

In Single active load balancing mode utilizes both the PE1 and PE2 access interface , EVI 100 traffic hashed to PE3 and EVI 200 traffic hashed to PE2 . CE23 mac address on EVI 100 is reachable only via the PE3 , as this PE3 Is the Designated Forwarder for EVI 100 which announcing the associated Mac/IP route, PE3 is referred to as the primary PE on PE1.

The mac address learning and propagation through various components remains the same as port active redundancy mode.

All- Active

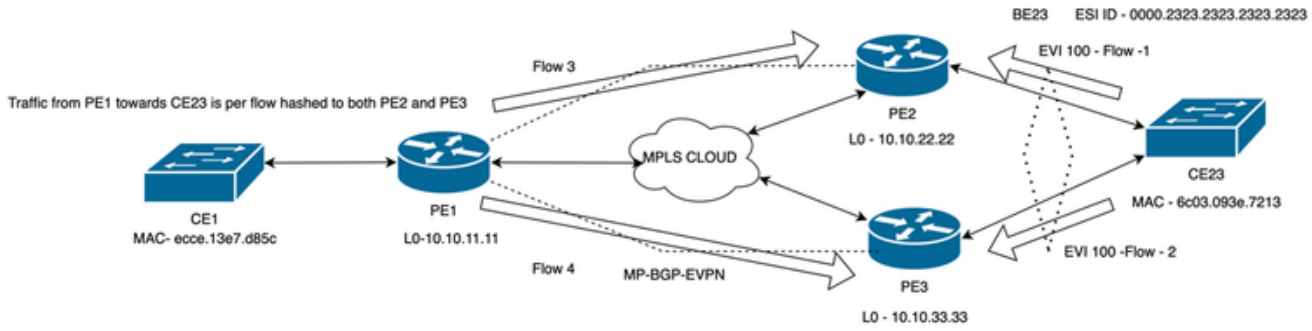

Figure 4 All- Active Redundancy Mode

As illustrated in the Figure 4 All- Active Redundancy Mode , in all-active redundancy mode for an EVI both PE2 and PE3 can send and receive known unicast traffic, the traffic is hashed per flow on both the PEs.

Configuration

```
Evpn
interface Bundle-Ether23
  ethernet-segment
   identifier type 0 00.23.23.23.23.23.23.23.23
```

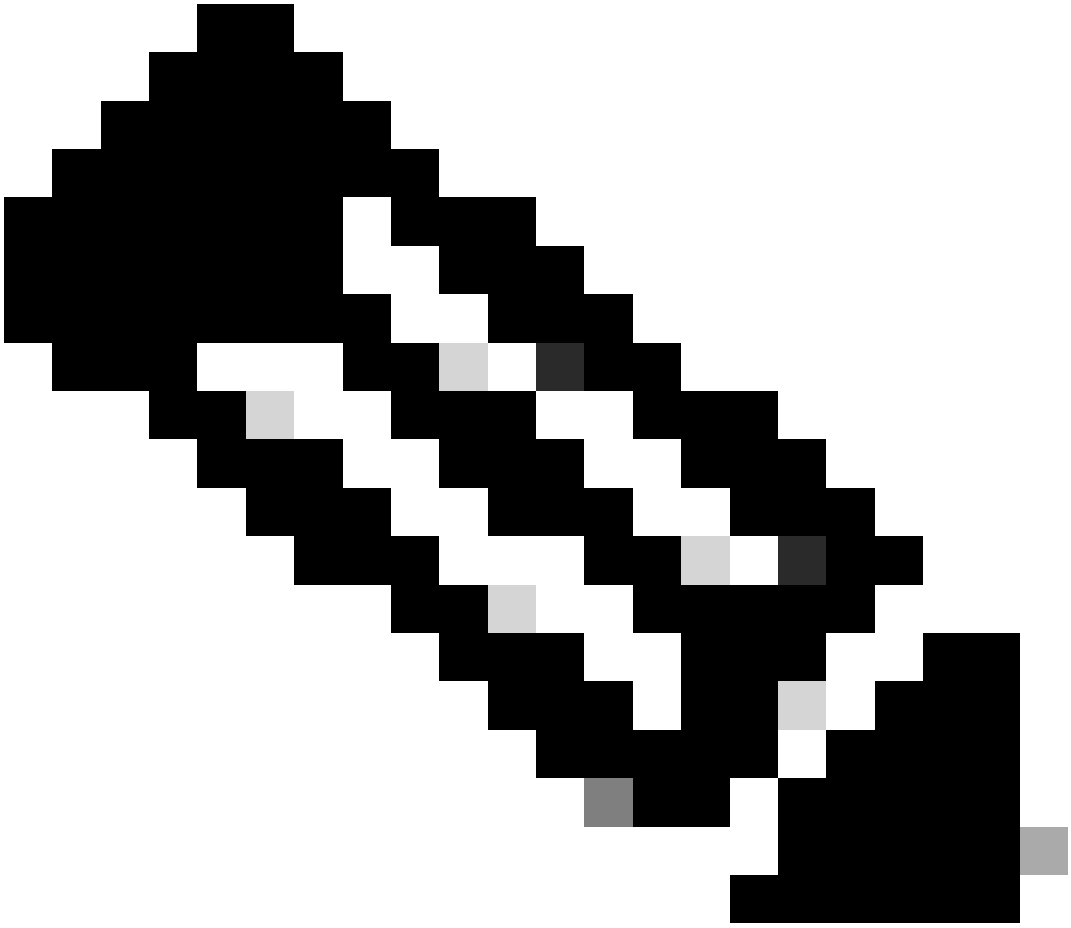

Note: The default mode is all-active .

Verification Of All-Active ESI

Interface status

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show interfaces Be23
Bundle-Ether23 is up, line protocol is up

RP/0/RSP0/CPU0:ASR9910-3-PE3#show interfaces BE23
Bundle-Ether23 is up, line protocol is up
```

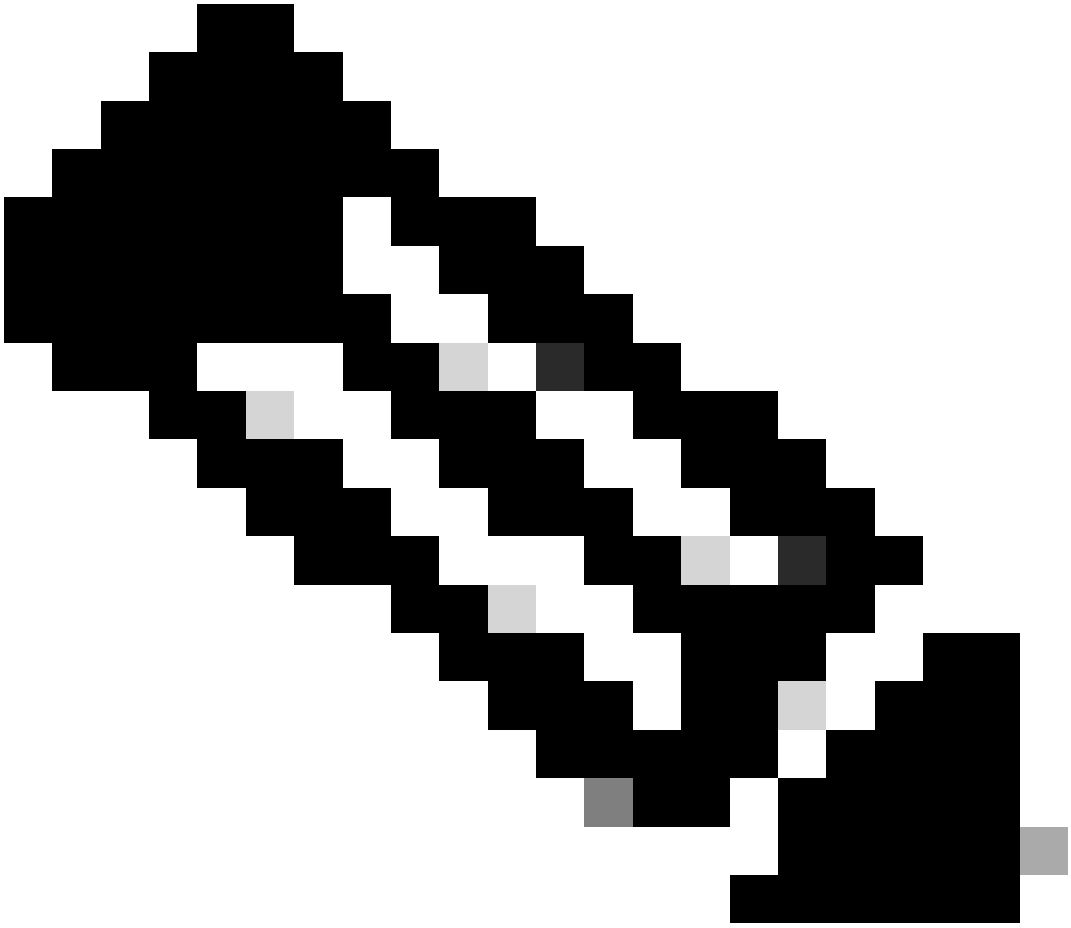

Note: The physical status and the bundle status of BE23 on PE2 and PE3 are up and active .

Ethernet Segment

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
<snip>
  Topology          :

Operational    : MH, All-active

     Configured     : All-active (AApF) (default)
  Service Carving   : Auto-selection
<snip>
```

```
Service Carving Results:
     Forwarders     : 2


 Elected        : 2


            EVI E   :      100,     200


     Not Elected    : 0
<snip>
Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.33.33
  Access signal mode: Bundle OOS

RP/0/RSP0/CPU0:ASR9910-3-PE3#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
<snip>
Topology          :


Operational    : MH, All-active


     Configured     : All-active (AApF) (default)
  Service Carving   : Auto-selection
<snip>
Service Carving Results:
     Forwarders     : 2
     Elected        : 0


 Not Elected    : 2


            EVI NE  :      100,     200


<snip>
Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.22.22
  Access signal mode: Bundle OOS (Default)
```

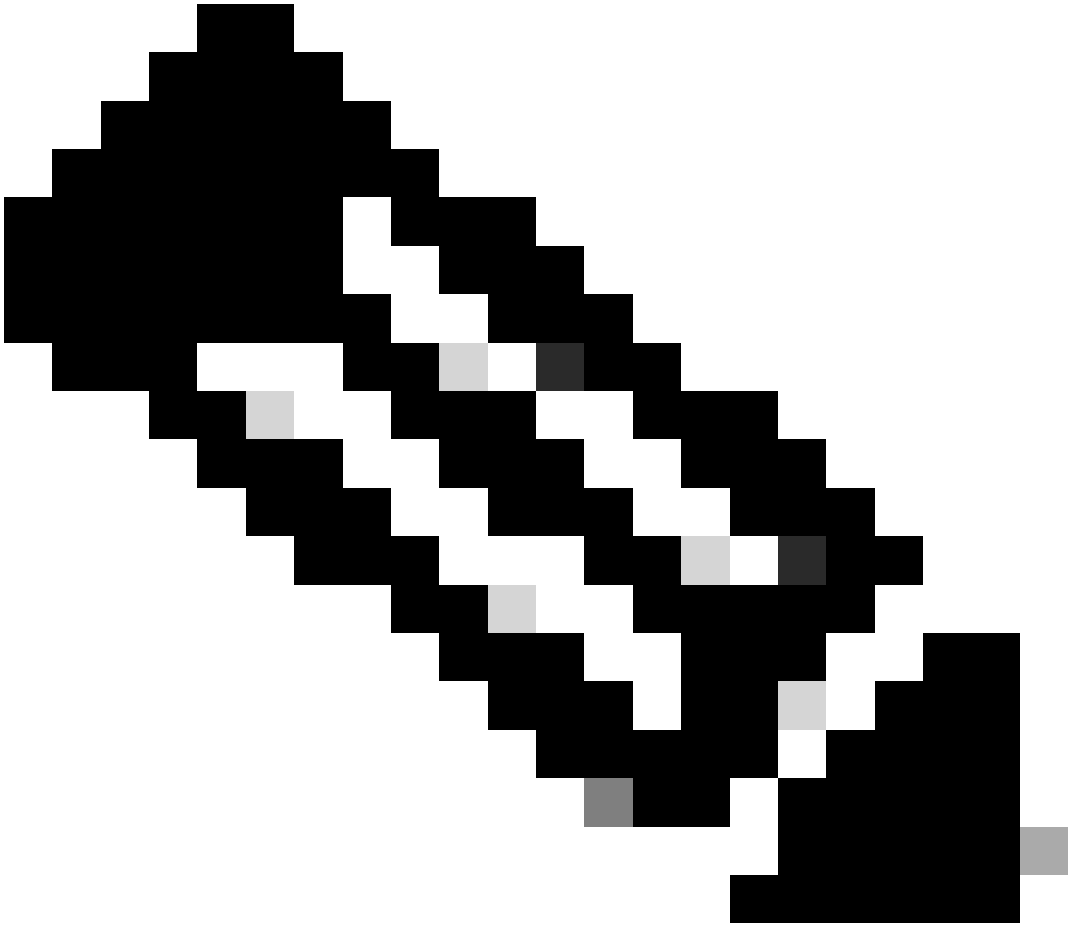

Note:

- The Designated Forwarder influences only the Broadcast, Unknown unicast and Multicast (BUM) traffic sent from core towards access device ,that is ,when BUM traffic from PE1 is sent to both PE2 and PE3 only PE2 forwards this traffic towards the CE23 .
- Whereas BUM traffic from CE23 can be hashed to both PE2 and PE3 which is forwarded to the remote PE1.
- for details on the Split-Horizon Mechanism and the SHG label , refer to this : https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/layer-2-vpns/222561-configure-evpn-route-types-and-their-fun.html#toc-hId--1100150104

Routes advertised to PE1 by PE2 and PE3

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show bgp l2vpn evpn
Route Distinguisher: 10.10.11.11:100 (default for vrf EVPN-ELAN-1)
*>i[1][0000.2323.2323.2323.2323][0]/120 >>>>>>>>>>>>>>>>>>>>>> Per EVI
                     10.10.22.22                   100      0 i
* i                   10.10.33.33                   100      0 i
```

```
*>i[1][0000.2323.2323.2323.2323][4294967295]/120>>>>>>>>>>>>>> Per ESI
                     10.10.22.22                    100      0 i
* i                   10.10.33.33                    100      0 i
```

***>i[2][0][48][6c03.093e.7213][0]/104 >>>>>>>>>>>>>>>>>>>>>>>>> Route-Type2 is advertised by both PE2 and**

10.10.22.22 100 0 i

*** i 10.10.33.33 100 0 i**

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show evpn internal-label detail
VPN-ID     Encap  Ethernet Segment Id         EtherTag     Label
---------- ------ -------------------------- ----------   --------
100        MPLS   0000.2323.2323.2323.2323    0            24010
   Multi-paths resolved: TRUE (Remote all-active)
   Multi-paths Internal label: 24010
    MAC         10.10.22.22                               24001
                10.10.33.33                               24001
    EAD/ES      10.10.22.22                               0
                10.10.33.33                               0
    EAD/EVI     10.10.22.22                               24001
                10.10.33.33                               24001

   Summary pathlist:

 0x02000001 (P) 10.10.22.22                               24001

 0x02000002 (P) 10.10.33.33                               24001
```

From the output we can observe that to reach CE23 on PE1 , PE2 and PE3 are consider as primary path and the traffic is hashed to both PE2 and PE3

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show mpls forwarding labels 24010 detail
Local  Outgoing    Prefix            Outgoing     Next Hop        Bytes
Label  Label       or ID             Interface                    Switched
------ ----------- ----------------- ------------ --------------- ------------
24010  24001       EVPN:100                       10.10.22.22     0
     Updated: Dec 29 12:19:20.649
     Version: 333, Priority: 3
     Label Stack (Top -> Bottom): { 24001 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
     MAC/Encaps: 0/4, MTU: 0
     Packets Switched: 0
       24001       EVPN:100                       10.10.33.33     0
     Updated: Dec 29 12:19:20.649
     Version: 333, Priority: 3
     Label Stack (Top -> Bottom): { 24001 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
     MAC/Encaps: 0/4, MTU: 0
     Packets Switched: 0
```

Mac address synchronization

L2vpn :

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show l2vpn forwarding bridge-domain  EVPN-ELAN:EVPN-ELAN mac-address loca
Mac Address    Type    Learned from/Filtered on    LC learned Resync Age/Last Change Mapped to
-------------- ------- -------------------------- ---------- --------------------- --------------
ecce.13e7.d85c EVPN    BD id: 0                   N/A        N/A                   N/A

6c03.093e.7213 static  BE23.100                   N/A        N/A                   N/A
```

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show l2vpn forwarding bridge-domain  EVPN-ELAN:EVPN-ELAN mac-address locat
Mac Address    Type    Learned from/Filtered on    LC learned Resync Age/Last Change Mapped to
-------------- ------- -------------------------- ---------- --------------------- --------------
ecce.13e7.d85c EVPN    BD id: 0                   N/A        N/A                   N/A

6c03.093e.7213 dynamic BE23.100                   N/A        30 Dec 06:49:50       N/A
```

From CE23 only 1 flow was initiated it was hashed to PE3 hence the learning is **dynamic on PE3** , on PE2 the mac address was synchronized for this ESI the entry is **static on PE2**.

L2RIB :

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show l2route evpn mac all
Topo ID  Mac Address    Producer    Next Hop(s)
-------- -------------- ----------- --------------------------------------
0        6c03.093e.7213 L2VPN       Bundle-Ether23.100, N/A
```

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show l2route evpn mac all
Topo ID  Mac Address    Producer    Next Hop(s)
-------- -------------- ----------- --------------------------------------
0        6c03.093e.7213 LOCAL       Bundle-Ether23.100, N/A
```

EVPN:

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show evpn evi vpn-id 100  mac
VPN-ID     Encap      MAC address    IP address                              Nexthop
---------- ---------- -------------- --------------------------------------- ------------------------

100        MPLS       6c03.093e.7213 ::                                      10.10.33.33
```

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show evpn evi vpn-id 100  mac
VPN-ID     Encap      MAC address    IP address                                Nexthop
---------- ---------- -------------- ----------------------------------------- ------------------------

100        MPLS       6c03.093e.7213 ::                                        Bundle-Ether23.100
```

This output at the EVPN component level clearly shows that PE3 dynamically learnt the mac address of CE23 and advertised Route-Type 2 towards PE2 and PE1 .

PE2 imported this Route-Type2 with the sync flag set and readvertises the Route-Type 2 for the mac address of CE23 with PE2 has the Next-Hop.

L2RIB traces from PE2 which shows this behavior :

<#root>

```
[12/29/24 11:14:07.763 UTC 225f 7879] Received MAC ROUTE msg: addr: (0, 6c03.093e.7213) vni: 0 admin_di

flags: S

 soo: 0 dg_count: 0 res: 0 esi: (F) >>> sync flag is set
[12/29/24 11:14:07.763 UTC 2262 7879] (0,6c03.093e.7213,9):Updated recv attrs seq:0 flags:S NH-type:2 N
[12/29/24 11:14:07.763 UTC 2264 7879] (0,6c03.093e.7213,9):set route flags: BEST
[12/29/24 11:14:07.763 UTC 2265 7879] (0,6c03.093e.7213,9):

MAC route created seq num:0 flags:BS (Rcv) rcv_seq:0 rcv_flags:S slot_id:0

 >>>>>>>>> the mac route is created with PE2 has the NH
[12/29/24 11:14:07.763 UTC 2266 7879] (0,6c03.093e.7213,9):MAC route created BR:9 Dup Moves:0 NH-type:2
[12/29/24 11:14:07.765 UTC 2267 7879] (0,6c03.093e.7213,9):Encoding MAC BR (ADD) Client = 3 BR = 0x5648
```

BGP:

<#root>

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show bgp l2vpn evpn rd 10.10.11.11:100 [2][0][48][6c03.093e.7213][0]/104
Local
    10.10.22.22 (metric 10) from 10.10.22.22 (10.10.22.22)
      Received Label 24001
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported, rib-ins
      Received Path ID 0, Local Path ID 1, version 1699
      Extended community:

SoO:10.10.33.33:100

 0x060e:0000.0000.0064 RT:100:100
      EVPN ESI: 0000.2323.2323.2323.2323
      Source AFI: L2VPN EVPN, Source VRF: default, Source

Route Distinguisher: 10.10.22.22:100


  Path #2: Received by speaker 0
  Not advertised to any peer
  Local
```

```
    10.10.33.33 (metric 10) from 10.10.33.33 (10.10.33.33)
      Received Label 24001
      Origin IGP, localpref 100, valid, internal, import-candidate, imported, rib-install
      Received Path ID 0, Local Path ID 0, version 0
      Extended community:
```

SoO:10.10.33.33:100

```
 0x060e:0000.0000.0064 RT:100:100
      EVPN ESI: 0000.2323.2323.2323.2323
      Source AFI: L2VPN EVPN, Source VRF: default,
```

Source Route Distinguisher: 10.10.33.33:100

Source of Origin Attribute shows that this Route-Type 2 was originated by PE3 , but the Route Distinguisher clearly indicates PE2 readvertised this Route-Type 2 .

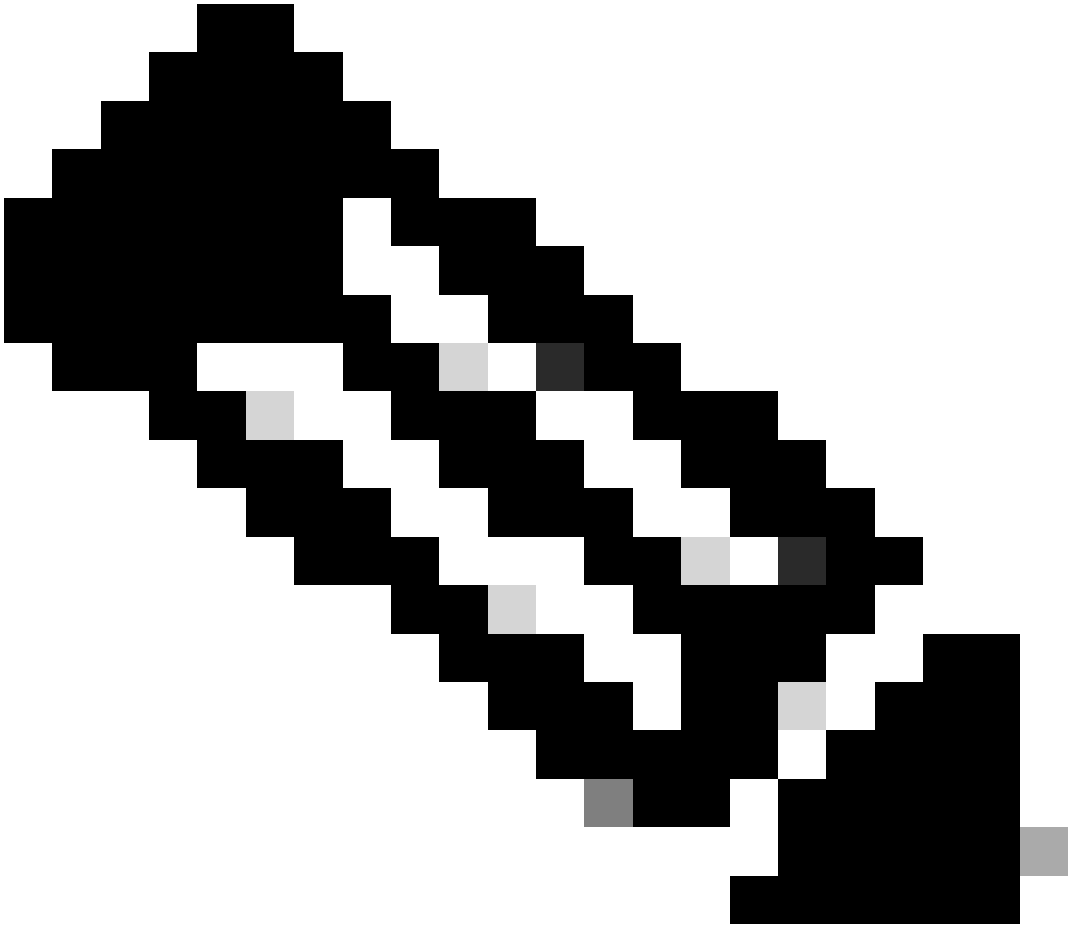

Note: When traffic is received on PE3 and PE2 from CE23 as various flows are hashed to both the links both PE2 and PE3 can dynamically learn the mac address and advertised Route-Type2 .

Single-Flow Active (SFA)

In this redundancy mode a Vlan can be active on all the PEs in the redundancy group but each unique L2 flow of that Vlan can be active on only one of the PEs in the redundancy group at a time .

Configuration

<#root>

```
Evpn
interface Bundle-Ether23
  ethernet-segment
   identifier type 0 00.23.23.23.23.23.23.23.23
```

load-balancing-mode single-flow-active >>> this command enables SFA redundancy mode

Verification Of Single-Flow Active ESI

Interface status

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show interfaces BE23
Bundle-Ether23 is up, line protocol is up
```

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show interfaces BE23
Bundle-Ether23 is up, line protocol is up
```

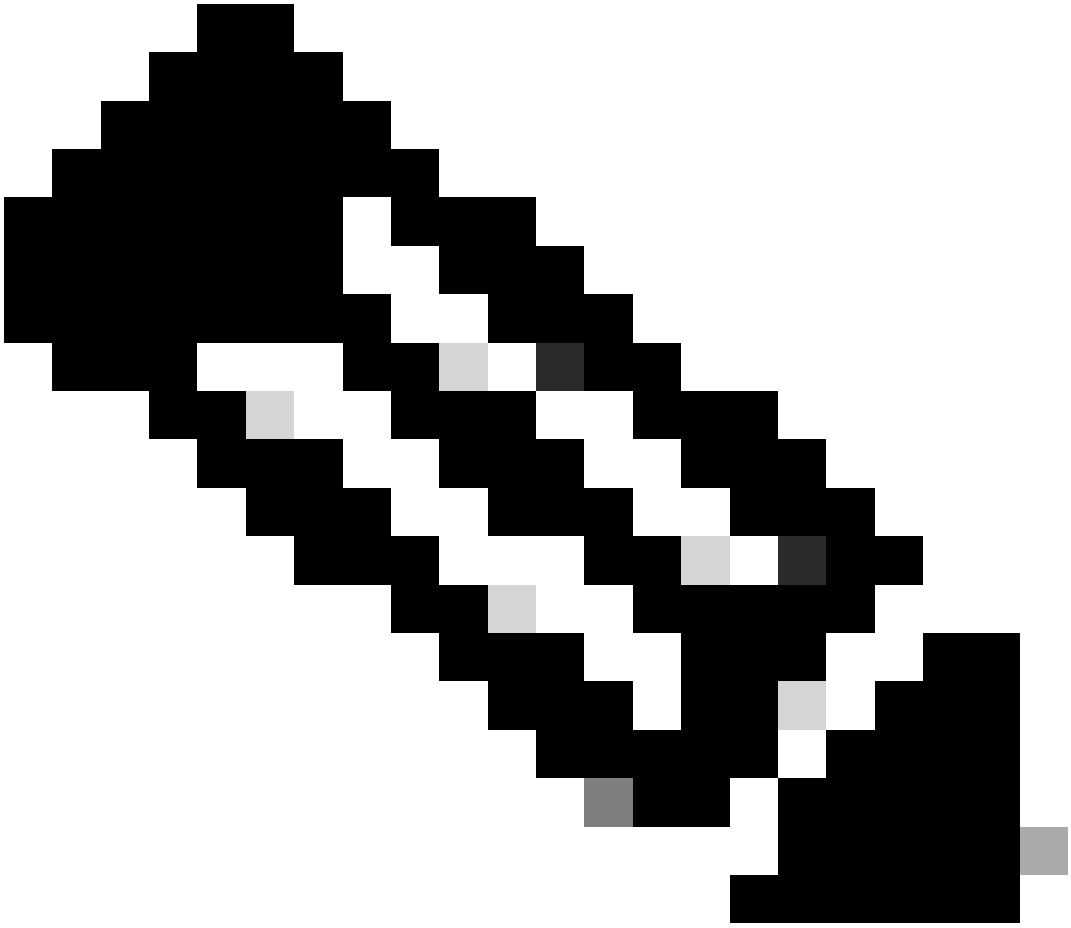

Note: The physical status and the bundle status of BE23 on PE2 and PE3 are up and active.

Ethernet Segment

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
<snip>
Topology          :

Operational     : MH, Single-flow-active

     Configured     : Single-flow-active
  Service Carving   : Auto-selection
<snip>
```

```
Service Carving Results:

Forwarders     : 2

     Elected        : 0

     Not Elected    : 0

<snip>
Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.33.33
  Access signal mode: Bundle OOS
```

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show evpn ethernet-segment carving detail
Ethernet Segment Id      Interface                         Nexthops
------------------------ --------------------------------- --------------------
0000.2323.2323.2323.2323 BE23                              10.10.22.22
                                                           10.10.33.33
<snip>
Topology          :

Operational    : MH, Single-flow-active

     Configured     : Single-flow-active
  Service Carving   : Auto-selection
<snip>
Service Carving Results:

Forwarders     : 2

     Elected        : 0

     Not Elected    : 0

<snip>
Local SHG label   : 24003
  Remote SHG labels : 1
              24003 : nexthop 10.10.22.22
  Access signal mode: Bundle OOS (Default)
```

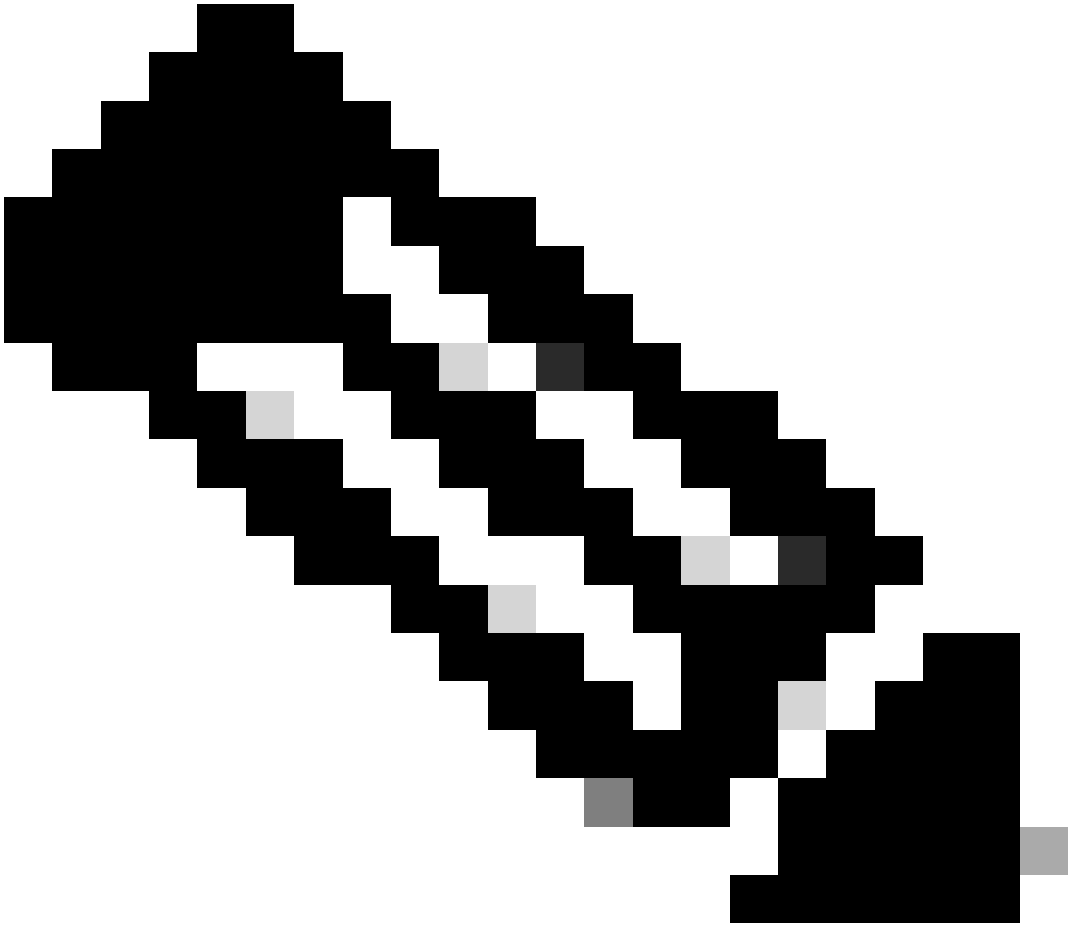

Note:

- There is no DF election in this redunacy mode .
- SFA - Split horizon is disabled .

Routes advertised to PE1 by PE2 and PE3 .

Per ESI route :

```
RP/0/RSP1/CPU0:ASR9906-1-PE1#show bgp l2vpn evpn rd 10.10.11.11:100 [1][0000.2323.2323.2323.2323][42949
  Local
    10.10.22.22 (metric 10) from 10.10.22.22 (10.10.22.22)
      Received Label 0
      Extended community: EVPN ESI Label:0x02:24003 RT:100:100 RT:200:200
      Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.10.22.22:1
  Path #2: Received by speaker 0
  Not advertised to any peer
  Local
    10.10.33.33 (metric 10) from 10.10.33.33 (10.10.33.33)
      Received Label 0
```

```
Extended community: EVPN ESI Label:0x02:24003 RT:100:100 RT:200:200
Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.10.33.33:1
```

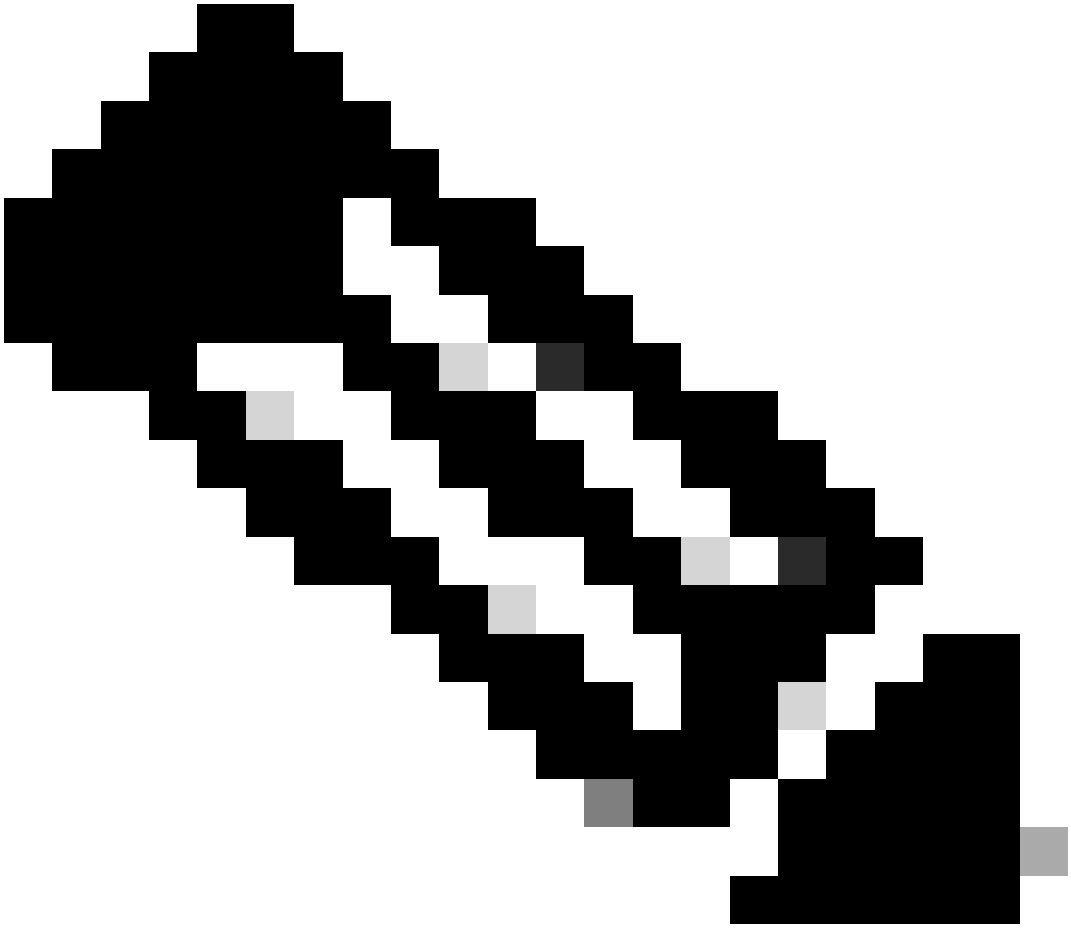

Note:

- EVPN ESI Label:0x02 – this indicates the remote PE2 and PE3 are in SFA redundancy mode .
- Once this extended community is received by the remote PE1 aliasing is disabled at Layer 2 and layer 3 .

Mac address Learning

L2vpn

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show l2vpn forwarding bridge-domain EVPN-ELAN-200:EVPN-ELAN-200 mac-addre
Mac Address    Type    Learned from/Filtered on    LC learned Resync Age/Last Change Mapped to
```

```
-------------- ------- -------------------------- ---------- ---------------------- --------------
ecce.13e7.d85c EVPN    BD id: 1                   N/A        N/A                    N/A

6c03.093e.7213 dynamic BE23.200                   N/A        30 Dec 10:05:07        N/A
```

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show l2vpn forwarding bridge-domain EVPN-ELAN-200:EVPN-ELAN-200 mac-address
Mac Address    Type    Learned from/Filtered on    LC learned Resync Age/Last Change Mapped to
-------------- ------- -------------------------- ---------- ---------------------- --------------

6c03.093e.7213 EVPN    BD id: 1                   N/A        N/A                    N/A


ecce.13e7.d85c EVPN    BD id: 1                   N/A        N/A                    N/A
```

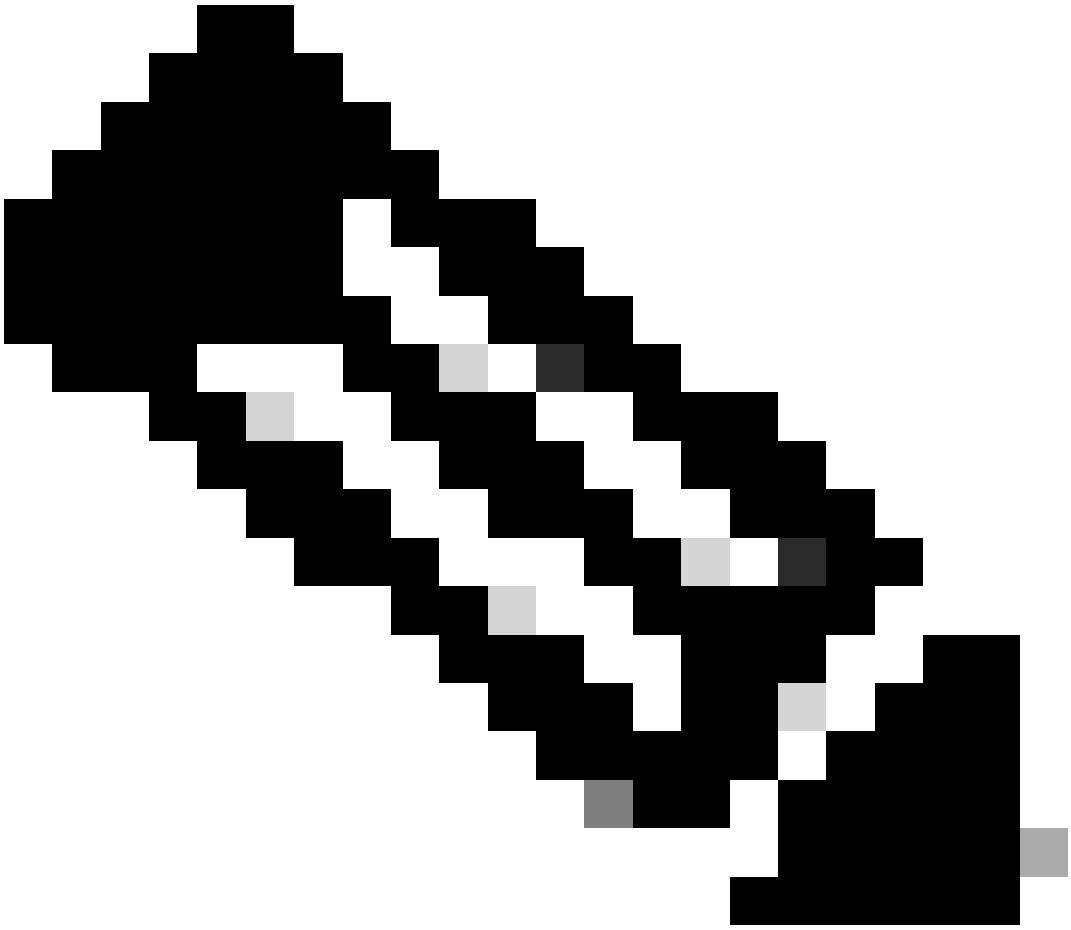

Note: On PE2 Mac is learnt dynamically and on PE3 mac is learnt via EVPN .

ARP :

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show arp vrf TEST
192.168.200.23  00:03:37   6c03.093e.7213
```

Dynamic

```
    ARPA  BVI200
```

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show arp vrf TEST
192.168.200.23  -          6c03.093e.7213
```

EVPN_SYNC

```
  ARPA  BVI200
```

L2rib:

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show l2route  evpn mac-ip all detail
1        6c03.093e.7213 192.168.200.23  L2VPN       Bundle-Ether23.200, 24006/I/ME           1        S
```

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show l2route  evpn mac-ip all detail
1        6c03.093e.7213 192.168.200.23  L2VPN       Bundle-Ether23.200, 24010/I/ME           2        B
```

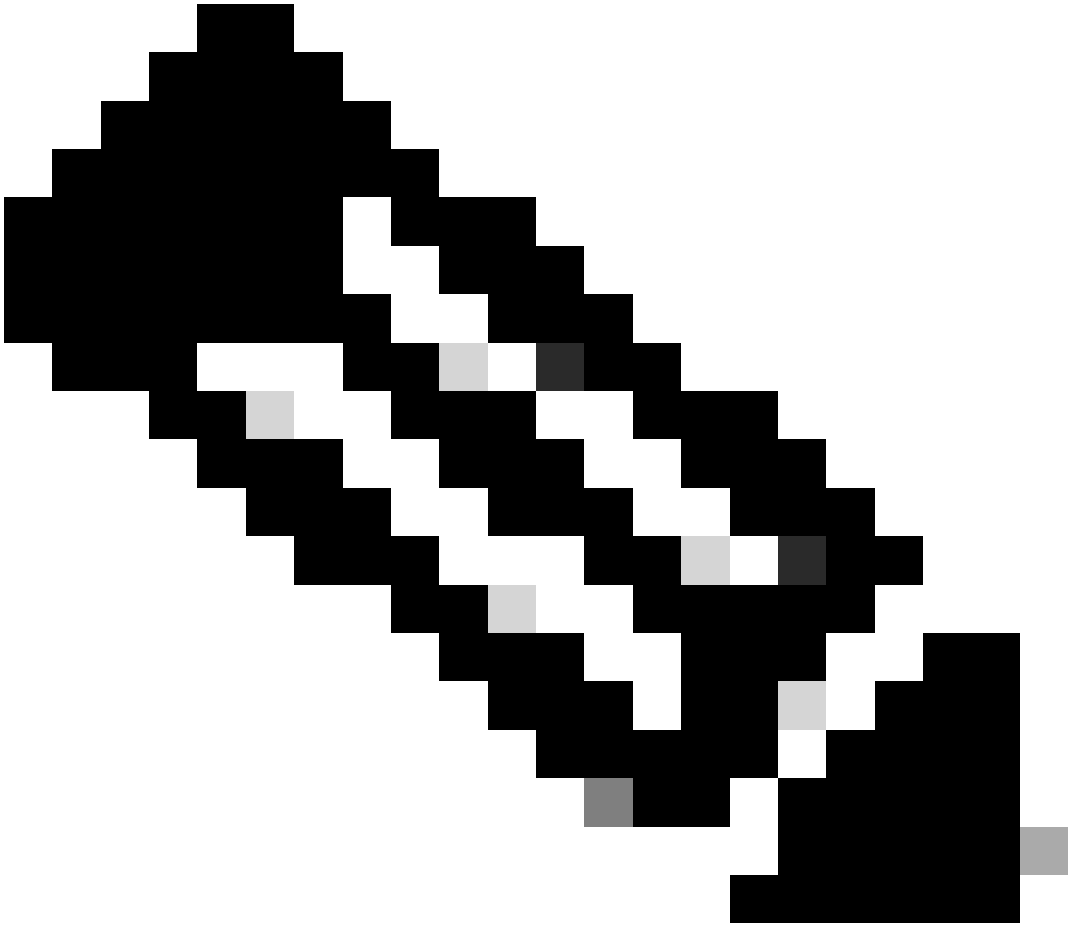

Note: BSSfa >> for this mac the best route is the remote route with sync flag set .

EVPN :

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show evpn evi vpn-id 200 mac ipv4 192.168.200.23
VPN-ID     Encap      MAC address    IP address                              Nexthop
---------- ---------- -------------- --------------------------------------- ------------------------
200        MPLS       6c03.093e.7213 192.168.200.23                          Bundle-Ether23.200
```

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show evpn evi vpn-id 200 mac ipv4 192.168.200.23
VPN-ID     Encap      MAC address    IP address                              Nexthop
---------- ---------- -------------- --------------------------------------- ------------------------
200        MPLS       6c03.093e.7213 192.168.200.23                          10.10.22.22
```

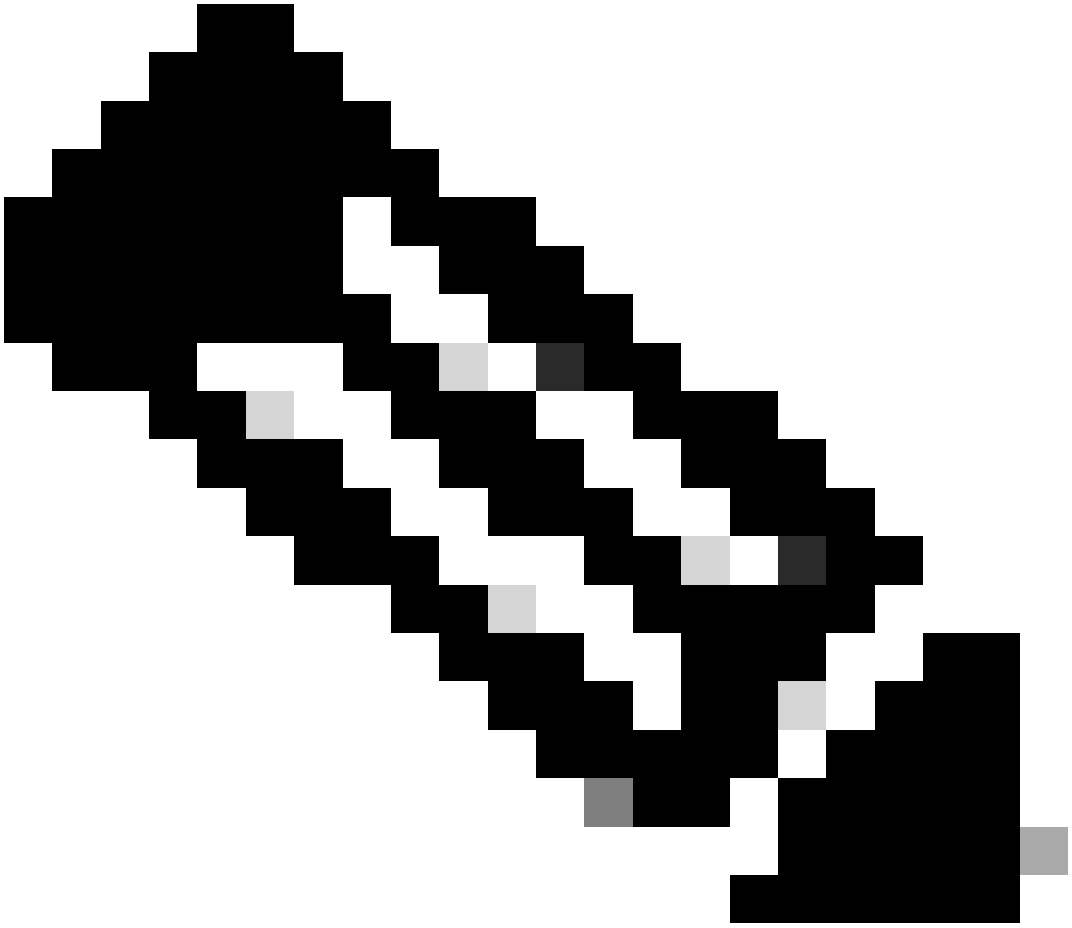

Note: From PE3 to reach 192.168.200.23 the Next-Hop is PE2 and not the local interface .

<#root>

```
RP/0/RSP1/CPU0:ASR-9904-5-PE2#show bgp l2vpn evpn rd 10.10.22.22:200  [2][0][48][6c03.093e.7213][32][19
BGP routing table entry for [2][0][48][6c03.093e.7213][32][192.168.200.23]/136, Route Distinguisher: 10
<snip>
  Local
    0.0.0.0 from 0.0.0.0 (10.10.22.22)
      Second Label 24012
      Origin IGP,
```

localpref 100

```
, valid, redistributed, best, group-best, import-candidate, rib-install
      Received Path ID 0, Local Path ID 1, version 2022
      Extended community: Flags 0xe: SoO:10.10.22.22:200 EVPN MAC Mobility:0x00:2 0x060e:0000.0000.00c:
      EVPN ESI: 0000.2323.2323.2323.2323
  Path #2: Received by speaker 0
  Not advertised to any peer
  Local
    10.10.33.33 (metric 10) from 10.10.33.33 (10.10.33.33)
      Received Label 24006, Second Label 24008
```

```
      Origin IGP,
localpref 80
, valid, internal, import-candidate, imported, rib-install
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: SoO:10.10.22.22:200 EVPN MAC Mobility:0x00:2 0x060e:0000.0000.00c8 RT:200:200
      EVPN ESI: 0000.2323.2323.2323.2323
      Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.10.33.33:200
```

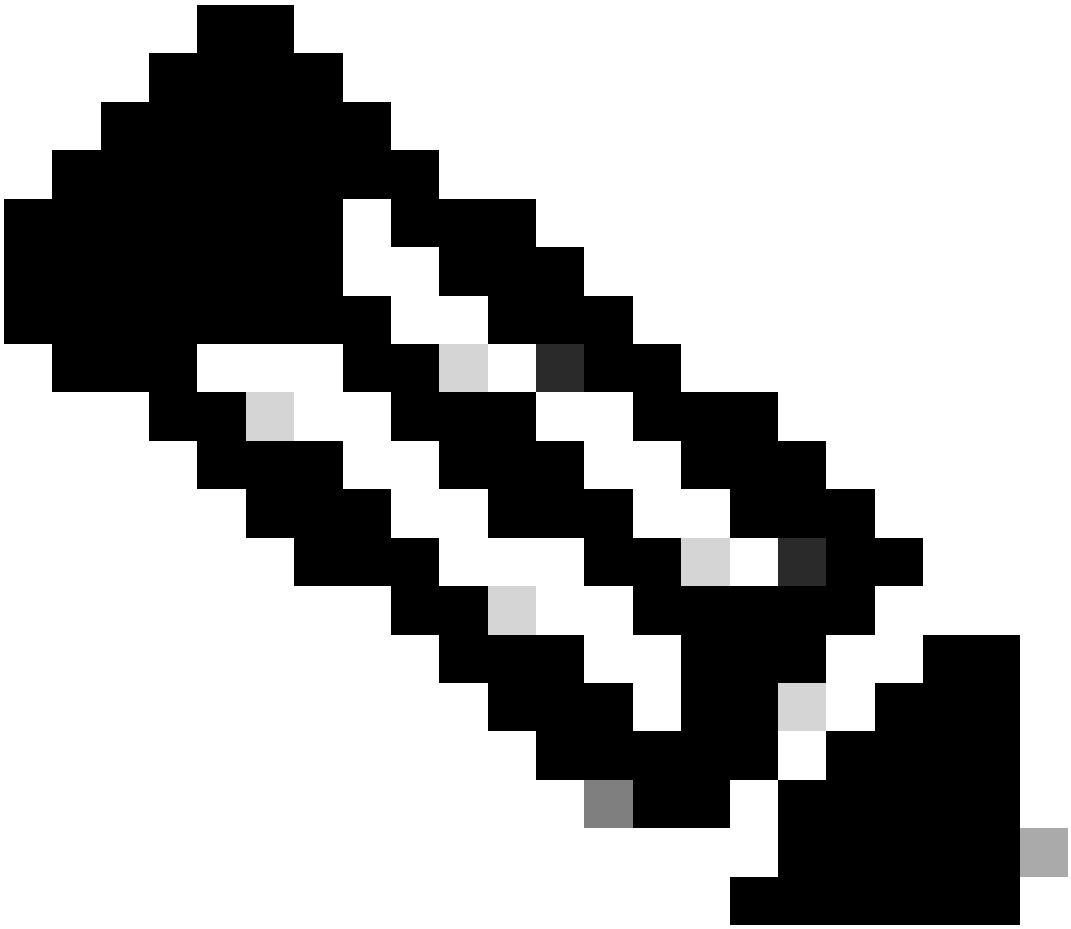

Note:

- In SFA when the mac address is learnt on PE2 it is synced with PE3 by advertising Route-type 2 , PE3 then re-advertises the Route-Type 2 for CE23.
- The remote router PE1 determines which is the primary PE for CE23 by the **Local Preference(LP) value** of the advertised Route Type 2 for CE23 ie PE2 has a LP of 100 and PE3 LP of 80 .
- On PE1 to send traffic to CE23 PE2 is preferred over PE3 .

On PE3 when the Route-Type2 for CE23 is synced and readvertised by PE3 it is advertised with Local

Preference 100 which is changed by the BGP outbound policy on PE3 this behavior can be observed in the subsequent outputs .

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3# show bgp l2vpn evpn rd  10.10.33.33:200 [2][0][48][6c03.093e.7213][32][19
<snip>
Local
    0.0.0.0 from 0.0.0.0 (10.10.33.33)
      Second Label 24008
      Origin IGP,
```

localpref 100

```
, valid, redistributed, best, group-best, import-candidate, rib-install
      Received Path ID 0, Local Path ID 1, version 1365
      Extended community: Flags 0xe: SoO:10.10.22.22:200 EVPN MAC Mobility:0x00:2 0x060e:0000.0000.00c8
      EVPN ESI: 0000.2323.2323.2323.2323
```

This Local preference is changed on PE3 when BGP applies the outbound policy before advertising the Route-Type2 to other remote PEs

<#root>

```
RP/0/RSP0/CPU0:ASR9910-3-PE3#show evpn evi vpn-id 200 mac ipv4 192.168.200.23 private
VPN-ID     Encap      MAC address    IP address                              Nexthop
---------- ---------- -------------- --------------------------------------- ------------------------
200        MPLS       6c03.093e.7213 192.168.200.23                          10.10.22.22
   Ethernet Tag                          : 0
   Multi-paths Resolved                  : True
   Multi-paths Internal label            : 24010
<snip>
   Ext Flags                             : 0x00000510 (Lcl Spec,Pref-Rib,
```

LP-80

```
,)
```

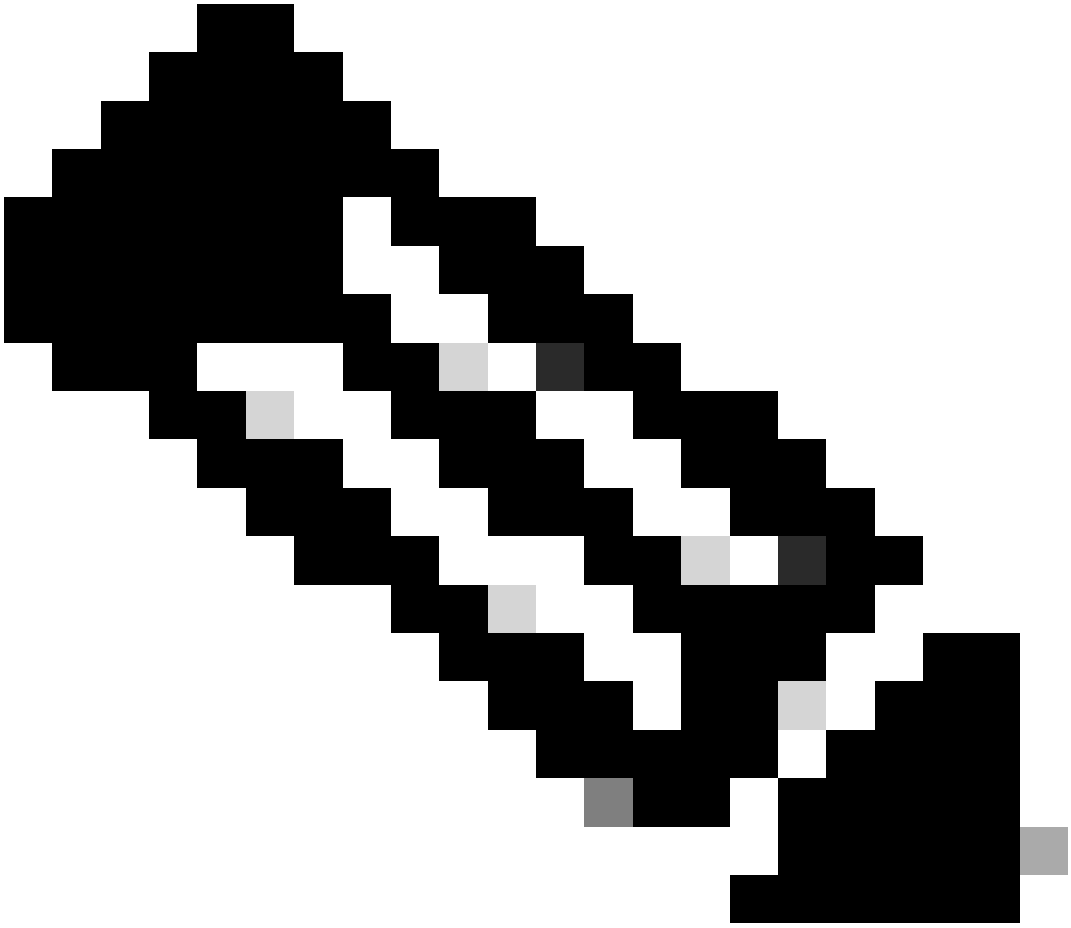

Note:

- Pref-Rib – prefer rib over local as this is a remote route .
- LP-80 – the outbound policy applied to Route-Type2 on PE3 to change the local preference .

- this document is also applicable for other XR products such as NCS 5500 , NCS 5700 .

Conclusion

EVPN Ethernet Segment Identifier (ESI) modes provide robust redundancy and traffic distribution mechanisms, ensuring high availability and efficient resource utilization in modern network architectures. By leveraging features such as Single-Active and All-Active redundancy modes, EVPN enables seamless failover, load balancing, and operational resilience. Understanding and implementing the appropriate ESI mode for specific network requirements is crucial for optimizing performance and maintaining network reliability.