

# SRv6 Operation Step by Step



# Agenda

- **L3VPN/EVPN L3VPN over SRv6 BE**
- **EVPN VPWS over SRv6-BE**

# End.DT4: Decapsulation and Specific IPv4 Table Lookup

## Processing Logic:

*IF NH=SRH and SL > 0*

drop the packet

*ELSE IF ENH = 4*

pop the (outer) IPv6 header and its extension headers

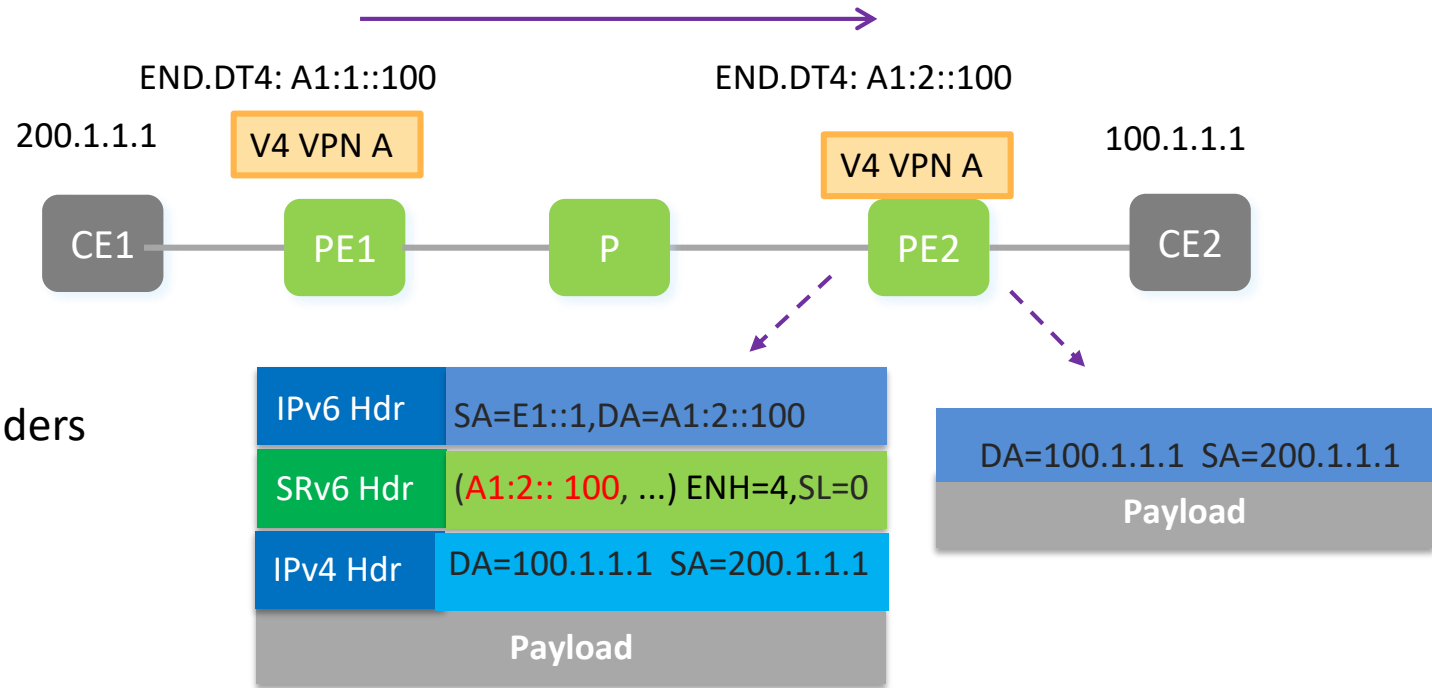
lookup the exposed inner IPv4 DA in IPv4 table T

forward via the matched table entry

*ELSE*

drop the packet

**Usage:** L3VPN use-case where a FIB lookup in a specific tenant table at the egress PE is required.



## Reminder:

The End.DT4 SID must always be the last SID, or it can be the destination Address of an IPv6 packet with no SRH header.

ENH (effective next header) 4 refers to IPv4 encapsulation as defined by IANA allocation for Internet Protocol Numbers

# Basic BGP Configuration

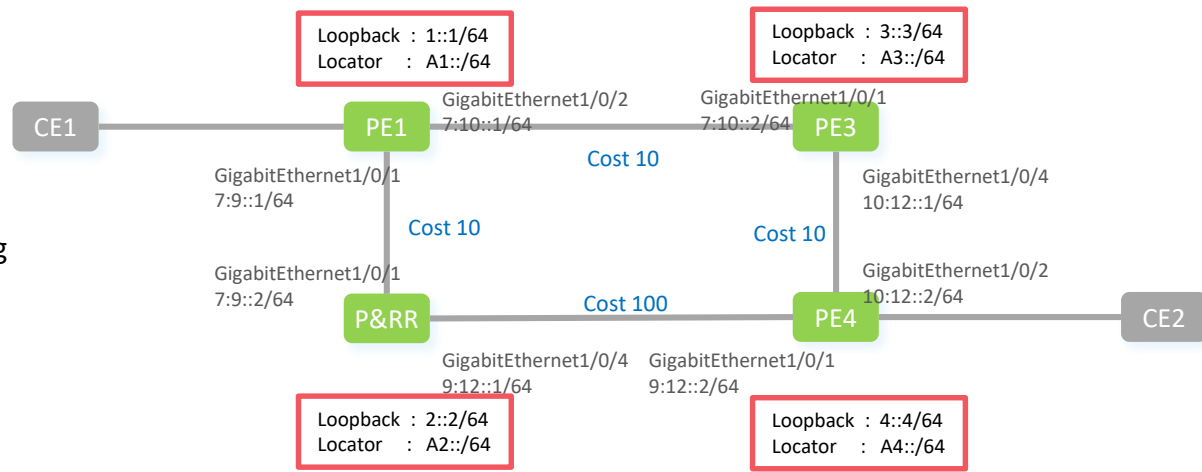
## BGP with VPNv4 and EVPN address-family

```
[~PE1-bgp]display this
#
bgp 100
peer 2::2 as-number 100
peer 2::2 connect-interface LoopBack0
#
ipv4-family vpnv4
policy vpn-target
peer 2::2 enable
peer 2::2 prefix-sid
#
l2vpn-family evpn
policy vpn-target
peer 2::2 enable
peer 2::2 advertise encap-type srv6
#
```

Neighbor with RR using IPv6 address

Prefix-SID extension

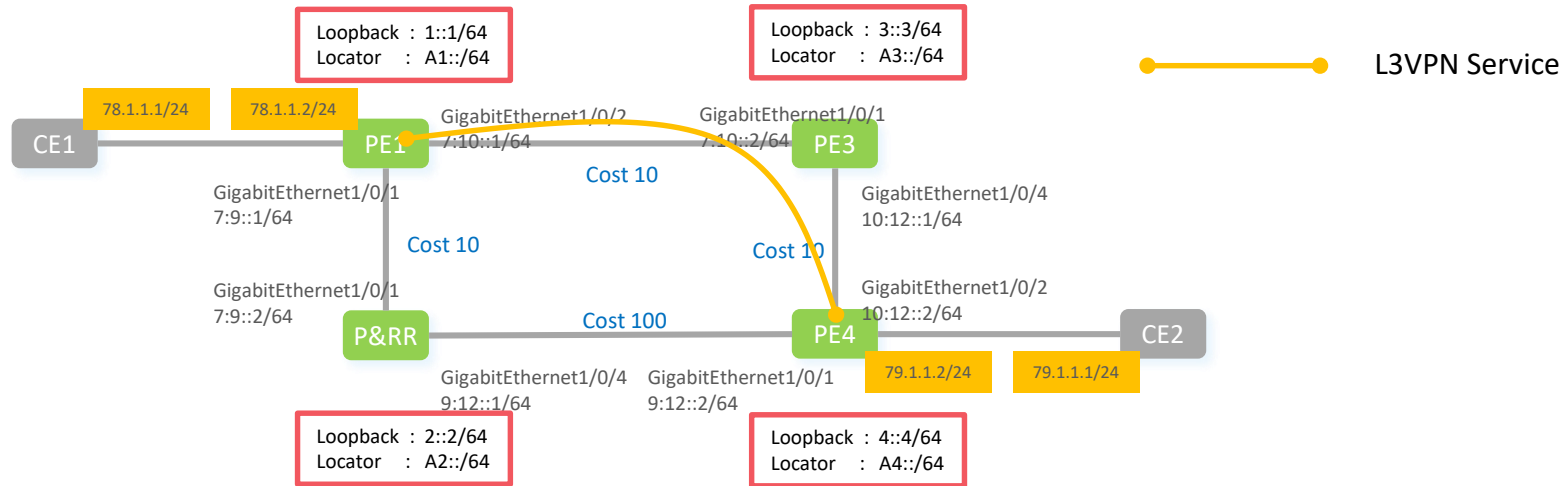
Prefix-SID extension and SRv6 encapsulation



## BGP Configuration on RR

```
[~P&RR-bgp]display this
#
bgp 100
peer 1::1 as-number 100
peer 1::1 connect-interface LoopBack0
peer 3::3 as-number 100
peer 3::3 connect-interface LoopBack0
peer 4::4 as-number 100
peer 4::4 connect-interface LoopBack0
#
ipv4-family vpnv4
undo policy vpn-target
peer 1::1 enable
peer 1::1 reflect-client
peer 1::1 prefix-sid
peer 3::3 enable
peer 3::3 reflect-client
peer 3::3 prefix-sid
#
l2vpn-family evpn
peer 4::4 enable
peer 4::4 reflect-client
peer 4::4 prefix-sid
#
l2vpn-family evpn
undo policy vpn-target
peer 1::1 enable
peer 1::1 reflect-client
peer 1::1 advertise encap-type srv6
peer 3::3 enable
peer 3::3 reflect-client
peer 3::3 advertise encap-type srv6
peer 4::4 enable
peer 4::4 reflect-client
peer 4::4 advertise encap-type srv6
#
```

# L3VPN over SRv6: End.DT4



## Traditional L3VPN and AC interface Without any change

```
[~PE1-vpn-instance-srv6_vpn1]display this
#
ip vpn-instance srv6_vpn1
ipv4-family
route-distinguisher 100:1
vpn-target 100:1 export-extcommunity
vpn-target 100:1 import-extcommunity
#
[~PE1-GigabitEthernet1/0/0.1]display this
#
interface GigabitEthernet1/0/0.1
vlan-type dot1q 1
ip binding vpn-instance srv6_vpn1
ip address 78.1.1.2 255.255.255.0
#
```

## SRv6 for L3VPN Service

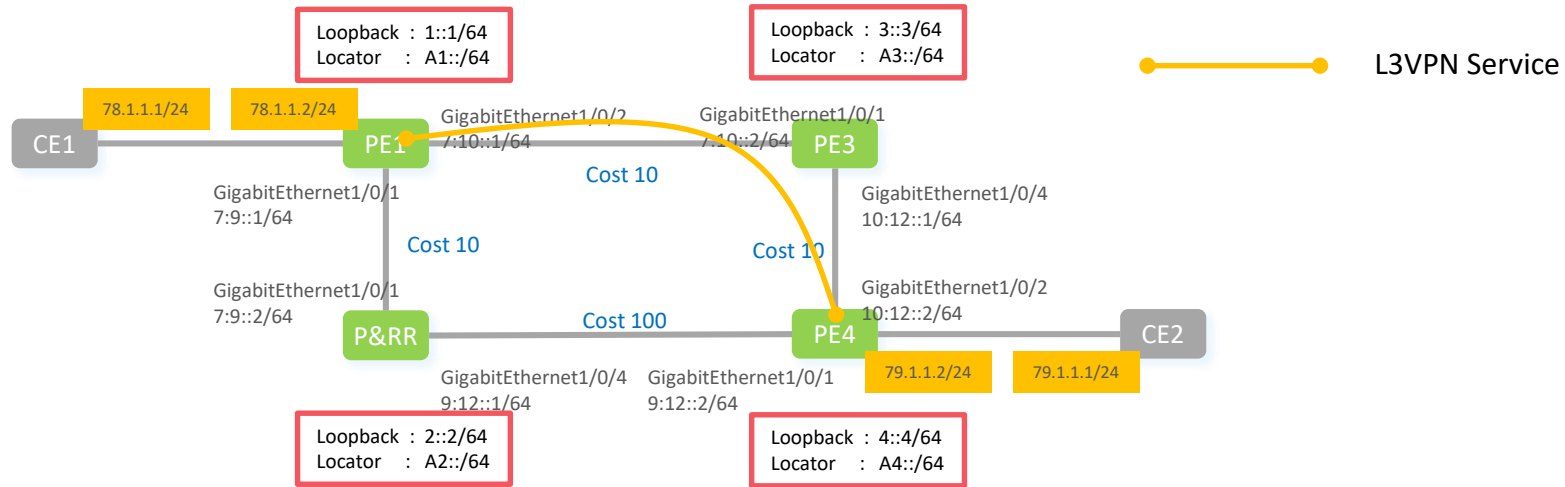
```
[~PE1-bgp-srv6_vpn1]display this
#
ipv4-family vpn-instance srv6_vpn1
import-route direct
segment-routing ipv6 locator SRv6_locator
segment-routing ipv6 best-effort
peer 78.1.1.1 as-number 65002
#
Optional
[~PE1-segment-routing-ipv6-locator]display this
locator SRv6_locator ipv6-prefix A1:: 64 static 32
opcode ::80 end-dt4 vpn-instance srv6_vpn1
#
```

Specify the locator and End.DT4  
SID will be auto-generated

Specify using VPN SID for  
recursive lookup and SRv6  
Best Effort forwarding

Specify static End.DT4 SID for  
vpn-instance in locator, and  
VPN will prefer the static SID

# Check End.DT4 Function



## Display Local-SID forwarding table for End.DT4

```
[~PE1-segment-routing-ipv6-locator]display segment-routing ipv6 local-sid end-dt4 forwarding

My Local-SID End.DT4 Forwarding Table
-----
SID       : A1::1:0:9B/128      FuncType : End.DT4
VPN Name  : srv6_vpn1          VPN ID   : 2
LocatorName: SRV6_locator     LocatorID: 1

SID       : A1::1:0:9C/128      FuncType : End.DT4
VPN Name  : srv6_vpn2          VPN ID   : 5
LocatorName: SRV6_locator     LocatorID: 1

Total SID(s): 2
```

Using dynamic generated SID A1::1:0:9B for srv6\_vpn1

```
[~PE1-segment-routing-ipv6-locator]opcode ::80 end-dt4 vpn-instance srv6_vpn1
[*PE1-segment-routing-ipv6-locator]comm
[~PE1-segment-routing-ipv6-locator]display segment-routing ipv6 local-sid end-dt4 forwarding

My Local-SID End.DT4 Forwarding Table
-----
SID       : A1::80/128         FuncType : End.DT4
VPN Name  : srv6_vpn1          VPN ID   : 2
LocatorName: SRV6_locator     LocatorID: 1

SID       : A1::1:0:9C/128      FuncType : End.DT4
VPN Name  : srv6_vpn2          VPN ID   : 5
LocatorName: SRV6_locator     LocatorID: 1

Total SID(s): 2
```

After static configure SID, using SID A1::80 for srv6\_vpn1

# Check L3VPN Routing on Remote PE

Display BGP routing table for CE1 address on remote PE4

```
[~PE4]display bgp vpnv4 vpn-instance srv6_vpn1 routing-table 78.1.1.0

BGP local router ID : 8.37.112.122
Local AS number : 100

VPN-Instance srv6_vpn1, Router ID 8.37.112.122:
Paths: 1 available, 1 best, 1 select, 0 best-external, 0 add-path
BGP routing table entry information of 78.1.1.0/24:
Route Distinguisher: 100:1
Remote-Cross route
Label information (Received/Applied): 3/NULL
From: 2::2 (8.37.112.119)
Route Duration: 0d00h00m09s
Relay IP Nexthop: FE80::82B5:75FF:FE4C:326D
Relay IP Out-Interface: GigabitEthernet1/0/2
Relay Tunnel out-Interface:
Original nexthop: 1::1
Qos information : 0x0
EXT-Community: RT <100 : 1>
Prefix-sid: A1::80
AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255, IGP cost 20
Originator: 8.37.112.117
Cluster list: 8.37.112.119
Advertised to such 1 peers:
 79.1.1.1
```

End.DT4 SID allocated on PE1

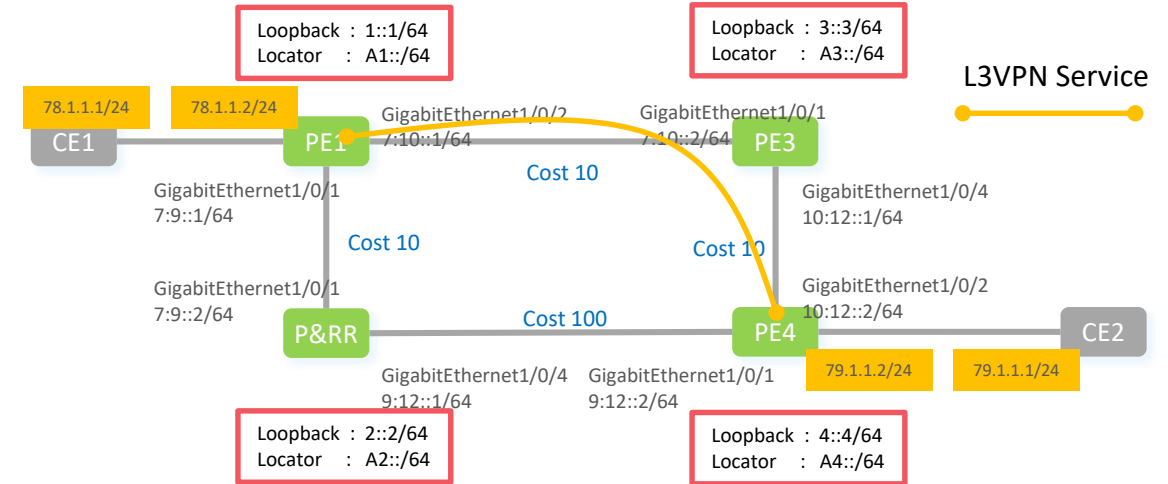
Display IP routing table for CE1 address on remote PE4

```
[~PE4]display ip routing-table vpn-instance srv6_vpn1 78.1.1.0 v
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black hole route
-----
Routing Table : srv6_vpn1
Summary Count : 1

Destination: 78.1.1.0/24
Protocol: IBGP
Preference: 255
NextHop: A1::80
State: Active Adv Retied
Tag: 0
Label: 3
IndirectID: 0x100026B
RelayNextHop: FE80::82B5:75FF:FE4C:326D
TunnelID: 0x0

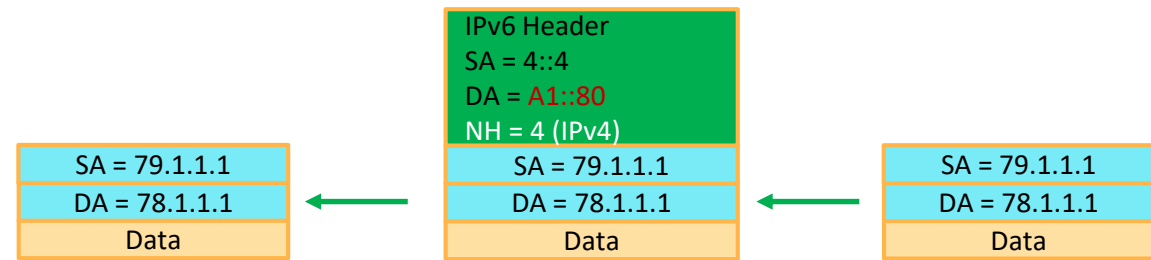
Process ID: 0
Cost: 0
Neighbour: 2::2
Age: 00h00m05s
Priority: low
QoSInfo: 0x0
Instance:
Interface: GigabitEthernet1/0/2
Flags: RD
```

VPN SID as next-hop for route-recursive

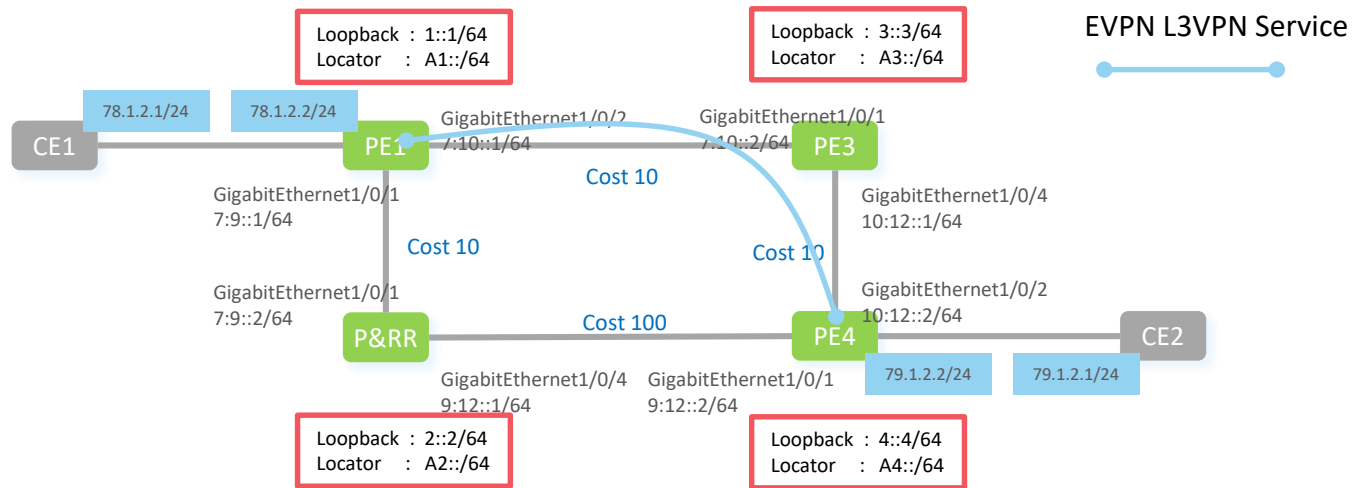


```
<CE2>ping -a 79.1.1.1 78.1.1.1
PING 78.1.1.1: 56 data bytes, press CTRL_C to break
Reply from 78.1.1.1: bytes=56 Sequence=1 ttl=253 time=2 ms
Reply from 78.1.1.1: bytes=56 Sequence=2 ttl=253 time=1 ms
Reply from 78.1.1.1: bytes=56 Sequence=3 ttl=253 time=1 ms
Reply from 78.1.1.1: bytes=56 Sequence=4 ttl=253 time=1 ms
Reply from 78.1.1.1: bytes=56 Sequence=5 ttl=253 time=1 ms

--- 78.1.1.1 ping statistics ---
5 packet(s) transmitted
5 packet(s) received
0.00% packet loss
round-trip min/avg/max = 1/1/2 ms
```



# SRv6 Configuration Example – EVPN L3VPN



## EVPN L3VPN and AC interface

```
[~PE1-vpn-instance-srv6_vpn2]display this
#
ip vpn-instance srv6_vpn2
ipv4-family
route-distinguisher 100:2
vpn-target 100:2 export-extcommunity evpn
vpn-target 100:2 import-extcommunity evpn
#
// Same AC Configuration
[~PE1-GigabitEthernet1/0/0.2]display this
#
interface GigabitEthernet1/0/0.2
vlan-type dot1q 2
ip binding vpn-instance srv6_vpn2
ip address 78.1.2.2 255.255.255.0
#
```

VPN-target for EVPN L3VPN

## SRv6 for EVPN L3VPN Service

```
[~PE1-bgp-srv6_vpn1]display this
#
ipv4-family vpn-instance srv6_vpn2
import-route direct
advertise l2vpn evpn
segment-routing ipv6 locator SRv6 locator evpn
segment-routing ipv6 best-effort evpn
peer 78.1.2.1 as-number 65002
#
Optional
[~PE1-segment-routing-ipv6-locator]display this
locator SRv6_locator ipv6-prefix A1:: 64 static 32
opcode ::80 end-dt4 vpn-instance srv6_vpn1
opcode ::81 end-dt4 vpn-instance srv6_vpn2 evpn
#
```

Advertise routes with EVPN L3VPN type-5 IP Prefix

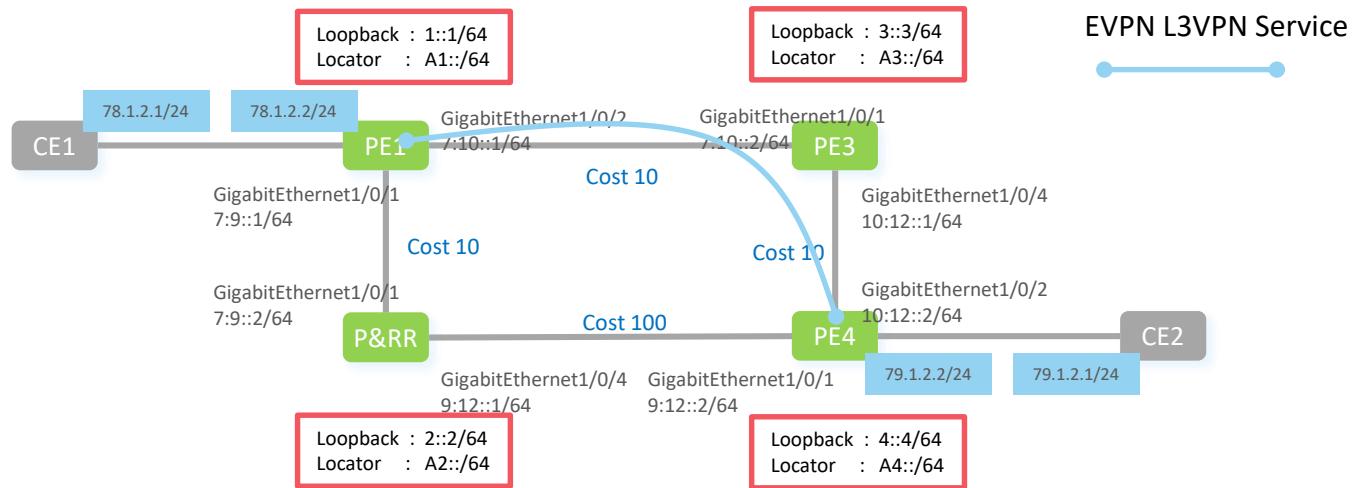
Specify the locator and End.DT4 SID will be auto-generated

Specify using VPN SID for recursive lookup and SRv6 Best Effort forwarding

Specify static End.DT4 SID for vpn-instance in locator, and EVPN L3VPN will prefer the static SID



# Check End.DT4 Function



## Display Local-SID forwarding table for End.DT4

```
[~PE1-segment-routing-ipv6-locator]display segment-routing ipv6 local-sid end-dt4 forwarding
My Local-SID End.DT4 Forwarding Table
-----
SID      : A1::80/128          FuncType : End.DT4
VPN Name : srv6_vpn1         VPN ID   : 2
LocatorName: SRv6_locator   LocatorID: 1

SID      : A1::1:0:9C/128    FuncType : End.DT4
VPN Name : srv6_vpn2         VPN ID   : 5
LocatorName: SRv6_locator   LocatorID: 1

Total SID(s): 2
```

Using dynamic generated SID A1::1:0:9C for EVPN L3VPN srv6\_vpn2

```
[~PE1-segment-routing-ipv6-locator]opcode ::81 end-dt4 vpn-instance srv6_vpn2 evpn
[*PE1-segment-routing-ipv6-locator]comm
[~PE1-segment-routing-ipv6-locator]display segment-routing ipv6 local-sid end-dt4 forwarding
My Local-SID End.DT4 Forwarding Table
-----
SID      : A1::80/128          FuncType : End.DT4
VPN Name : srv6_vpn1         VPN ID   : 2
LocatorName: SRv6_locator   LocatorID: 1

SID      : A1::81/128        FuncType : End.DT4
VPN Name : srv6_vpn2         VPN ID   : 5
LocatorName: SRv6_locator   LocatorID: 1

Total SID(s): 2
```

After static configure SID, using SID A1::81 for EVPN L3VPN srv6\_vpn2

# Check EVPN L3VPN Routing on Remote PE

Display BGP routing table for CE1 address on remote PE4

```
[~PE4]display bgp vpnv4 vpn-instance srv6_vpn2 routing-table 78.1.2.0
BGP local router ID : 8.37.112.122
Local AS number : 100

VPN-Instance srv6_vpn2, Router ID 8.37.112.122:
Paths: 1 available, 1 best, 1 select, 0 best-external, 0 add-path
BGP routing table entry information of 78.1.2.0/24:
Route Distinguisher: 100:2
Remote-Cross route
Evpn route: Type 5, ip-prefix
Label information (Received/Applied): 3/NULL
From: 2::2 (8.37.112.119)
Route Duration: 0d00h06m54s
Relay IP Nexthop: FE80::82B5:75FF:FE4C:326D
Relay IP Out-Interface: GigabitEthernet1/0/2
Relay Tunnel Out-Interface:
Original nexthop: 1::1
Qos information : 0x0
Ext-Community: RT <100 : 2>
Prefix-sid: A1::81
AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255, IGP cost 20
Originator: 8.37.112.117
Cluster list: 8.37.112.119
Advertised to such 1 peers:
79.1.2.1
```

EVPN Type-5 IP Prefix Route

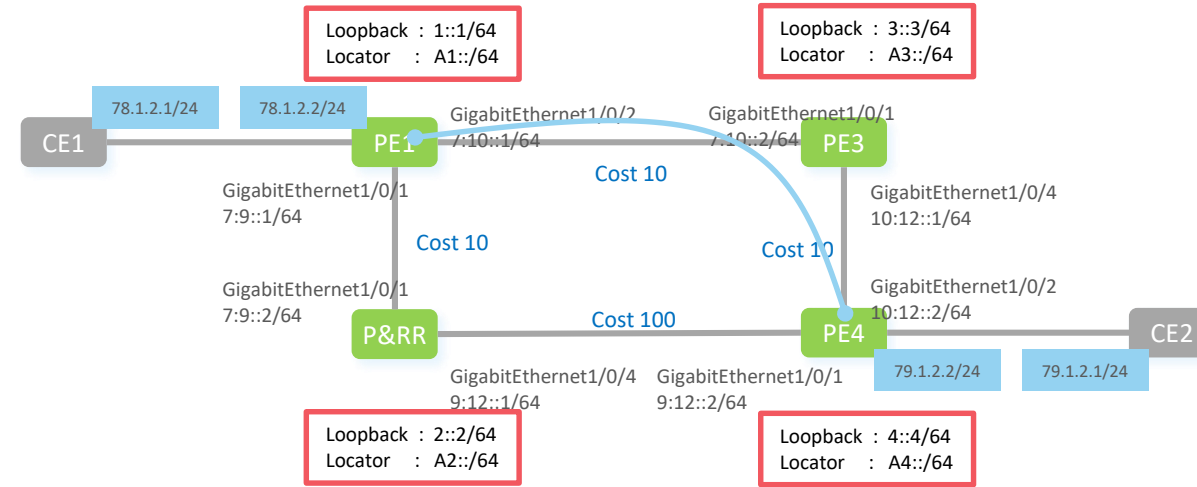
End.DT4 SID allocated on PE1

Display IP routing table for CE1 address on remote PE4

```
[~PE4]display ip routing-table vpn-instance srv6_vpn2 78.1.2.0 v
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black hole route
-----
Routing Table : srv6_vpn2
Summary Count : 1

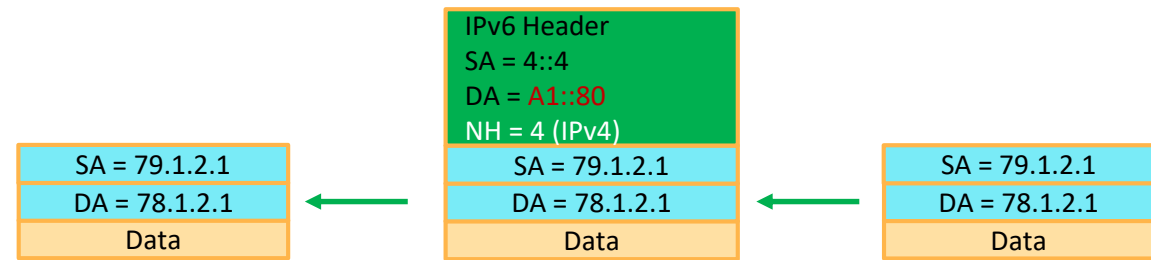
Destination: 78.1.2.0/24
Protocol: IBGP          Process ID: 0
Preference: 255         Cost: 0
NextHop: A1::81         Neighbour: 2::2
State: Active Adv Relyd Age: 00h09m39s
Tag: 0                  Priority: low
Label: NULL             QoSInfo: 0x0
IndirectID: 0x1000274   Instance:
RelayNextHop: FE80::82B5:75FF:FE4C:326D Interface: GigabitEthernet1/0/2
TunnelID: 0x0          Flags: RD
```

VPN SID as next-hop for route-recursive



```
<CE2>ping -a 79.1.2.1 78.1.2.1
PING 78.1.2.1: 56 data bytes, press CTRL_C to break
Reply from 78.1.2.1: bytes=56 Sequence=1 ttl=253 time=1 ms
Reply from 78.1.2.1: bytes=56 Sequence=2 ttl=253 time=1 ms
Reply from 78.1.2.1: bytes=56 Sequence=3 ttl=253 time=1 ms
Reply from 78.1.2.1: bytes=56 Sequence=4 ttl=253 time=1 ms
Reply from 78.1.2.1: bytes=56 Sequence=5 ttl=253 time=1 ms

--- 78.1.2.1 ping statistics ---
5 packet(s) transmitted
5 packet(s) received
0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```



# Agenda

- L3VPN/EVPN L3VPN over SRv6 BE
- **EVPN VPWS over SRv6-BE**

# End.DX2: Decapsulation and L2 cross-connect to OIF

## Processing Logic:

*IF* NH=SRH and SL > 0

drop the packet

*ELSE IF* ENH = 59

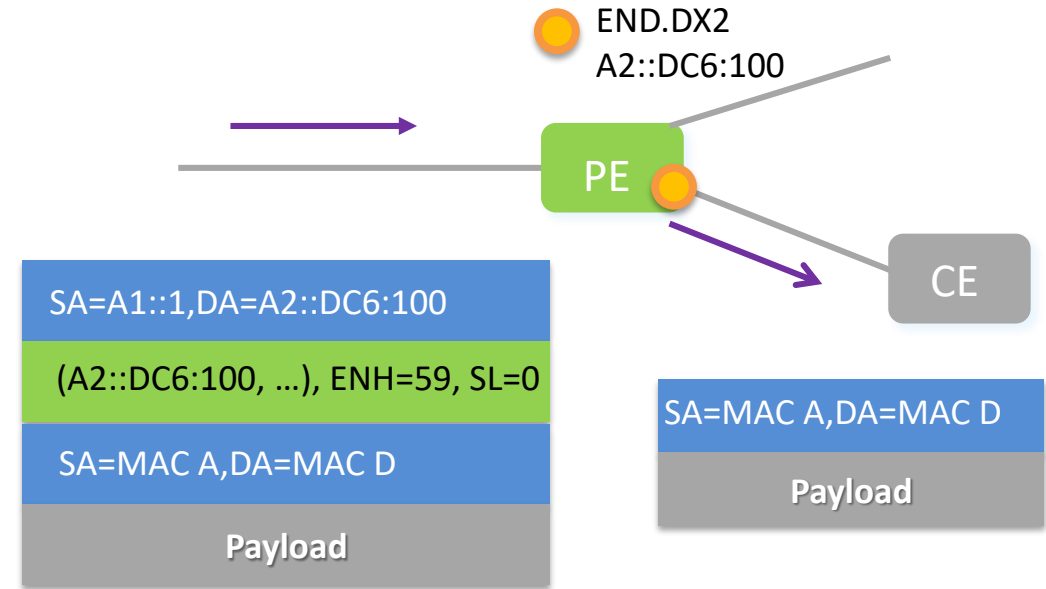
pop the (outer) IPv6 header and its extension headers

forward the resulting frame to OIF bound to the SID S

*ELSE*

drop the packet

**Usage:** L2VPN/EVPN VPWS use-case.



## Reminder:

An End.DX2 SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header

# Single-Homed EVPN VPWS over SRv6 BE

## Note:

BGP EVPN address-family has already deployed in [page 17](#)

### EVPN VPWS and AC interface

```
[~PE1-vpws-evpn-instance-srv6_vpws]display this
```

```
#
evpn vpn-instance srv6_vpws vpws
route-distinguisher 100:2
```

```
segment-routing ipv6 best-effort
```

```
vpn-target 100:2 export-extcommunity
vpn-target 100:2 import-extcommunity
```

```
#
[~PE1-evpl-srv6-1]display this
```

```
#
evpl instance 1 srv6-mode
```

```
evpn binding vpn-instance srv6_vpws // Binding evpl instance to evpn
instance, multi evpl instance could be bind to one evpn instance
local-service-id 100 remote-service-id 200 // Specify local and remote
AC-ID, on the other PE, configuration should be reversed
```

```
segment-routing ipv6 locator SRv6_locator
```

```
#
[~PE1-GigabitEthernet1/0/0.100]display this
```

```
#
interface GigabitEthernet1/0/0.100 mode I2
encapsulation dot1q vid 100
rewrite pop single
evpl instance 1 // Binding AC to evpl instance
#
```

Specify using VPN SID for recursive lookup and SRv6 Best Effort forwarding

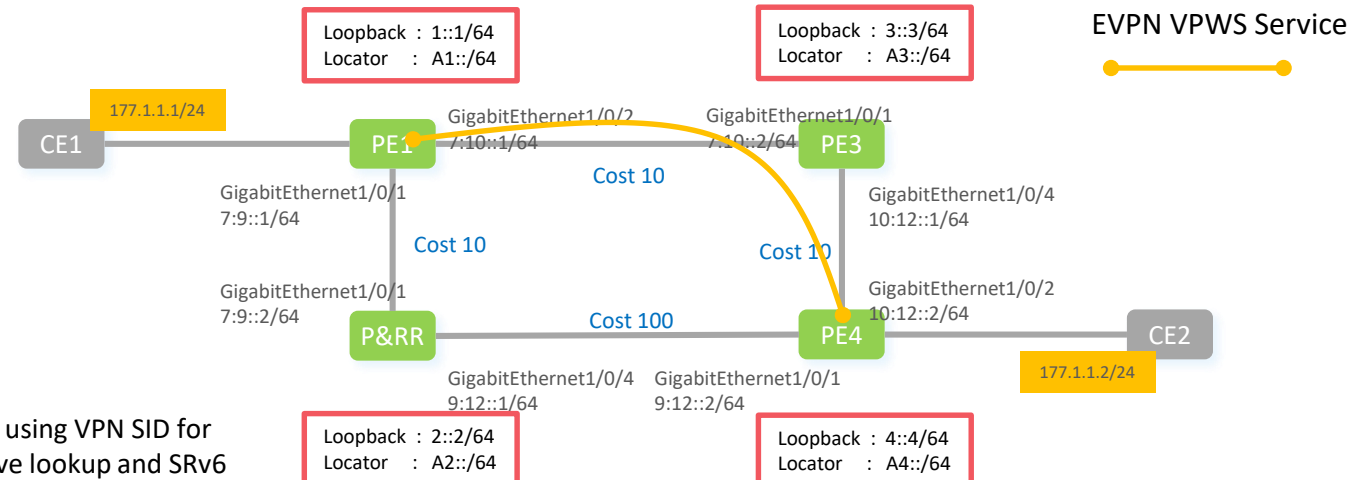
Specify evpl instance in SRv6 mode

Specify the locator and End.DX2 SID will be auto-generated

### Optional

```
[~PE1-segment-routing-ipv6-locator]display this
locator SRv6_locator ipv6-prefix A1:: 64 static 32
opcode ::82 end-dx2 evpl-instance 1
#
```

Specify static End.DX2 SID for evpl instance



# Check End.DX2 Function and Local EVPN A-D Route

Display End.DX2 forwarding table on PE1

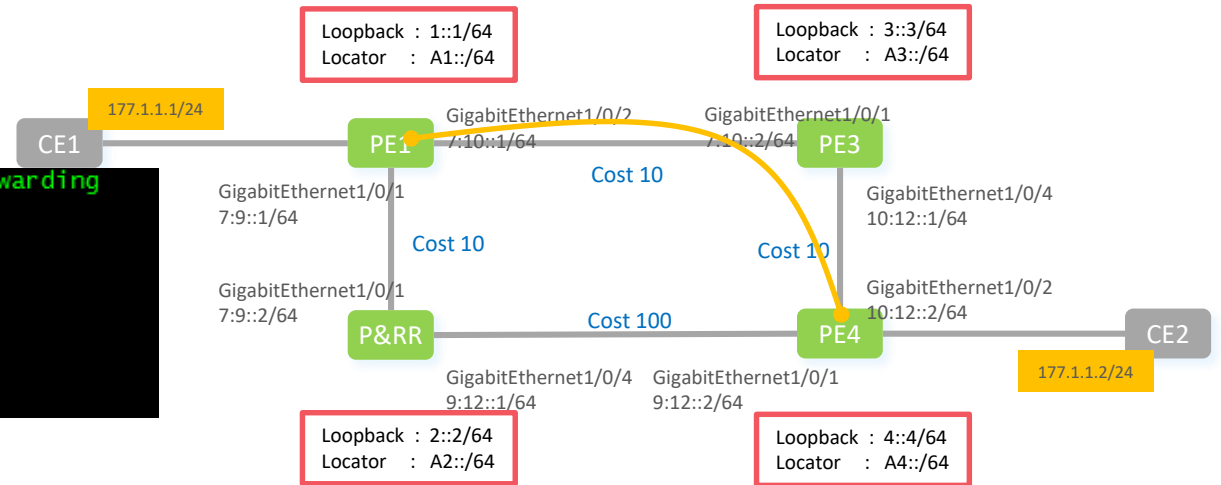
```
[~PE1-segment-routing-ipv6]display segment-routing ipv6 local-sid end-dx2 forwarding
My Local-SID End.DX2 Forwarding Table
-----
SID       : A1::82/128  → Static configured      FuncType : End.DX2
EVPL ID   : 1          → End.DX2 SID A1::82
LocatorName: SRV6_locator      LocatorID: 1
Total SID(s): 1
```

Display local EVPN A-D route for evpl instance 1 on PE1

```
[~PE1]display bgp evpn vpn-instance srv6_vpws routing-table ad-route 0000.0000.0000.0000.0000:100

BGP local router ID : 8.37.112.117
Local AS number : 100

EVPN-Instance srv6_vpws:
Number of A-D Routes: 1
BGP routing table entry information of 0000.0000.0000.0000.0000:100:
Route Distinguisher: 100:2
Aggregated route.
Route Duration: 0d00h04m00s
Relay IP Nexthop: 0.0.0.0
Relay IP Out-Interface:
Original nexthop: 127.0.0.1
Qos information : 0x0
Ext-Community: EVPN L2 Attributes <MTU:1500 C:0 P:1 B:0>
AS-path Nil, origin incomplete, pref-val 0, valid, local, best, select, pre 255
Prefix-sid: A1::82
Route Type: 1 (Ethernet Auto-Discovery (A-D) route)
ESI: 0000.0000.0000.0000.0000, Ethernet Tag ID: 100
Not advertised to any peer yet
```



# Check Remote EVPN A-D Route and EVPL Instance

Display all EVPN A-D route for evpl instance 1 on PE4

```
[~PE4]display bgp evpn vpn-instance srv6_vpws routing-table ad-route
BGP Local router ID is 8.37.112.122
Status codes: * - valid, > - best, d - damped, x - best external, a - add path,
              h - history, i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

EVPN-Instance srv6_vpws:
Number of A-D Routes: 2
 *>i  0000.0000.0000.0000.0000 100 Remote AC-ID on PE4      NextHop
 *>   0000.0000.0000.0000.0000 200 Local AC-ID on PE4     127.0.0.1
```



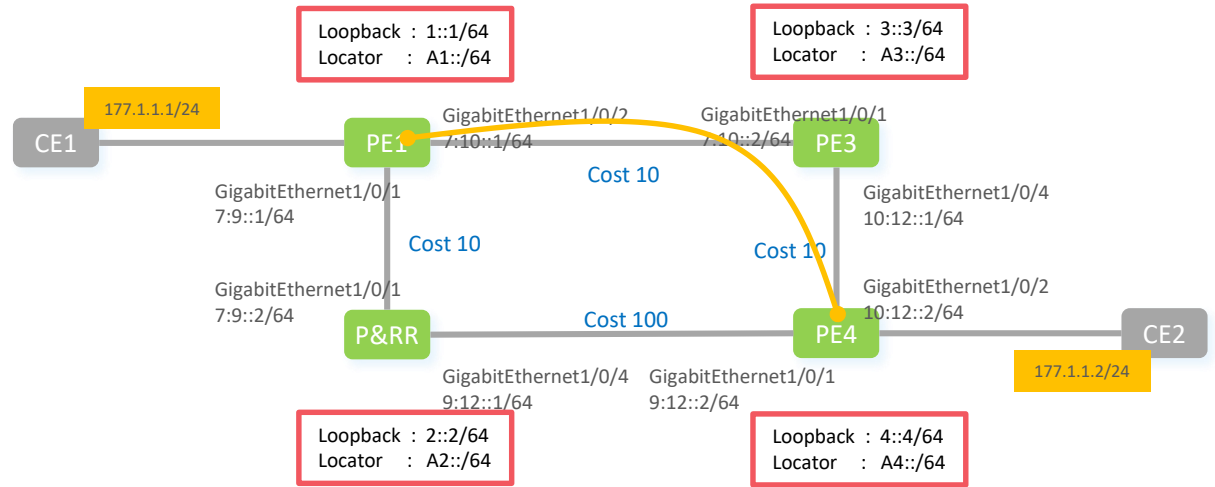
Display remote EVPN A-D route for evpl instance 1 on PE4

```
[~PE4]display bgp evpn vpn-instance srv6_vpws routing-table ad-route 0000.0000.0000.0000:100

BGP Local router ID : 8.37.112.122
Local AS number : 100

EVPN-Instance srv6_vpws:
Number of A-D Routes: 1
BGP routing table entry information of 0000.0000.0000.0000:100:
Route Distinguisher: 100:2
Remote-Cross route
Label information (Received/Applied): 3/NULL
From: 2::2 (8.37.112.119)
Route Duration: 0d00h02m01s
Relay IP Nexthop: FE80::82B5:75FF:FE4C:326D
Relay IP Out-Interface:GigabitEthernet1/0/2
Relay Tunnel Out-Interface:
Original nexthop: 1::1
Qos information : 0x0
Ext-Community: RT <100 >2>, EVPN L2 Attributes <MTU:1500 C:0 P:1 B:0>
AS-path Nil, origin incomplete, localpref 100, pref-val 0, valid, internal, best, select, pre 255, IGP cost 20
Originator: 8.37.112.117
Cluster list: 8.37.112.119
Prefix-sid: A1::82
Route type: 1 (Ethernet Auto-Discovery (A-D) route)
ESI: 0000.0000.0000.0000.0000, Ethernet Tag ID: 100
Not advertised to any peer yet
```

End.DX2 SID allocated on PE1



Display evpl instance 1 state on PE4

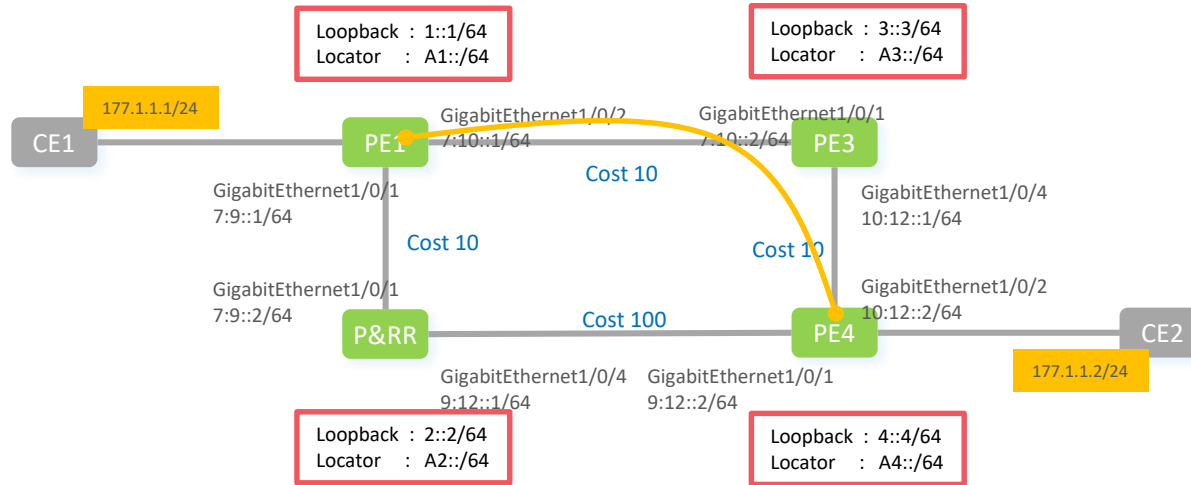
```
[~PE4]display bgp evpn evpl instance-id 1
Total EVPLs: 1      1 Up      0 Down

EVPL ID : 1
State : up
Evpl type : srv6-mode
Interface : GigabitEthernet1/0/0.100
Local MTU : 1500
Local Control word : false
Local Redundancy Mode : all-active
Local DF State : primary
Local ESI : 0000.0000.0000.0000.0000
Remote Redundancy Mode : all-active
Remote Primary DF Number : 1
Remote Backup DF Number : 0
Remote None DF Number : 0
Peer IP : 2::2
Origin Nexthop IP : 1::1
DF State : primary
Eline Role : primary
Remote MTU : 1500
Remote Control word : false
Remote ESI : 0000.0000.0000.0000.0000
Tunnel info : 1 tunnels
NO.0 Tunnel Type : Native IPV6, Tunnel ID :
Last interface UP Timestamp : 2019-4-22 16:18:42:902
Last Designated Primary Timestamp : 2019-4-22 15:38:45:710
Last Designated Backup Timestamp : --
```

Using SRv6 BE forwarding



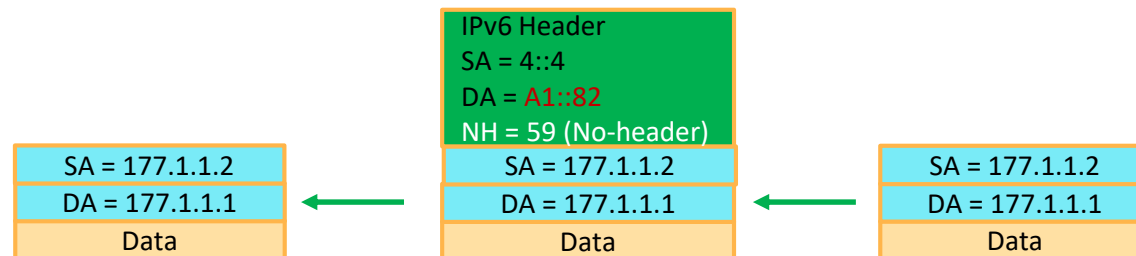
# Ping from CE2 to CE1 across EVPN VPWS



```

<CE2>ping -a 177.1.1.2 177.1.1.1
PING 177.1.1.1: 56 data bytes, press CTRL_C to break
Reply from 177.1.1.1: bytes=56 Sequence=1 ttl=255 time=2 ms
Reply from 177.1.1.1: bytes=56 Sequence=2 ttl=255 time=1 ms
Reply from 177.1.1.1: bytes=56 Sequence=3 ttl=255 time=1 ms
Reply from 177.1.1.1: bytes=56 Sequence=4 ttl=255 time=1 ms
Reply from 177.1.1.1: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 177.1.1.1 ping statistics ---
 5 packet(s) transmitted
 5 packet(s) received
 0.00% packet loss
 round-trip min/avg/max = 1/1/2 ms
    
```





# Thanks

Huawei Live

