

PIM Source Specific Multicast (SSM)

Presenter:

Daryl Wan



Agenda

1 Overview

2 Use Cases

3 Details and Caveats

4 Configuration

5 Best Practices

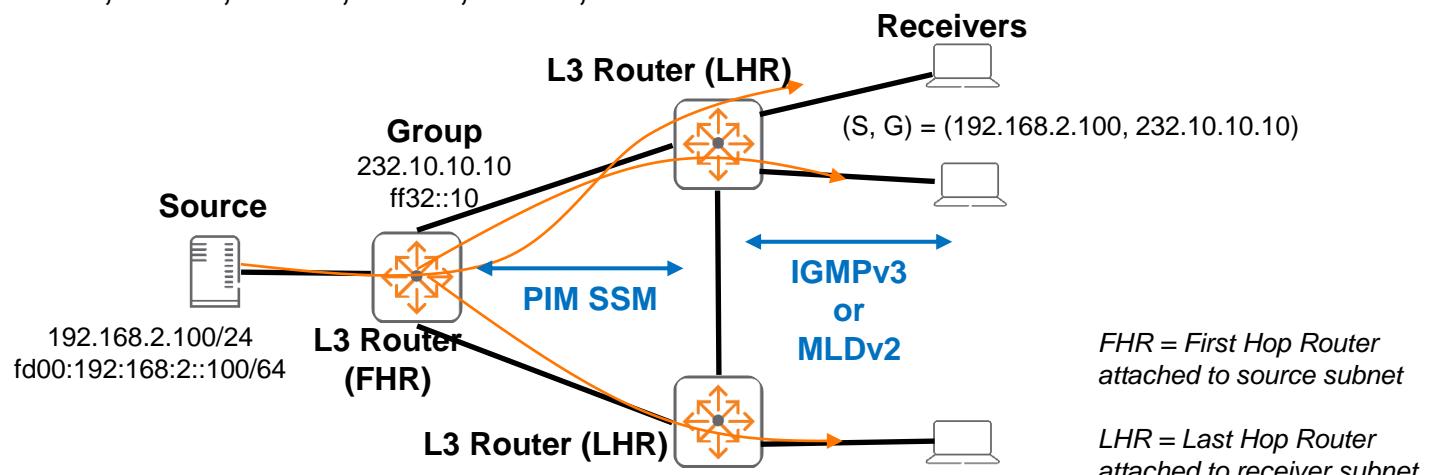
6 Troubleshooting

7 Demo

Overview

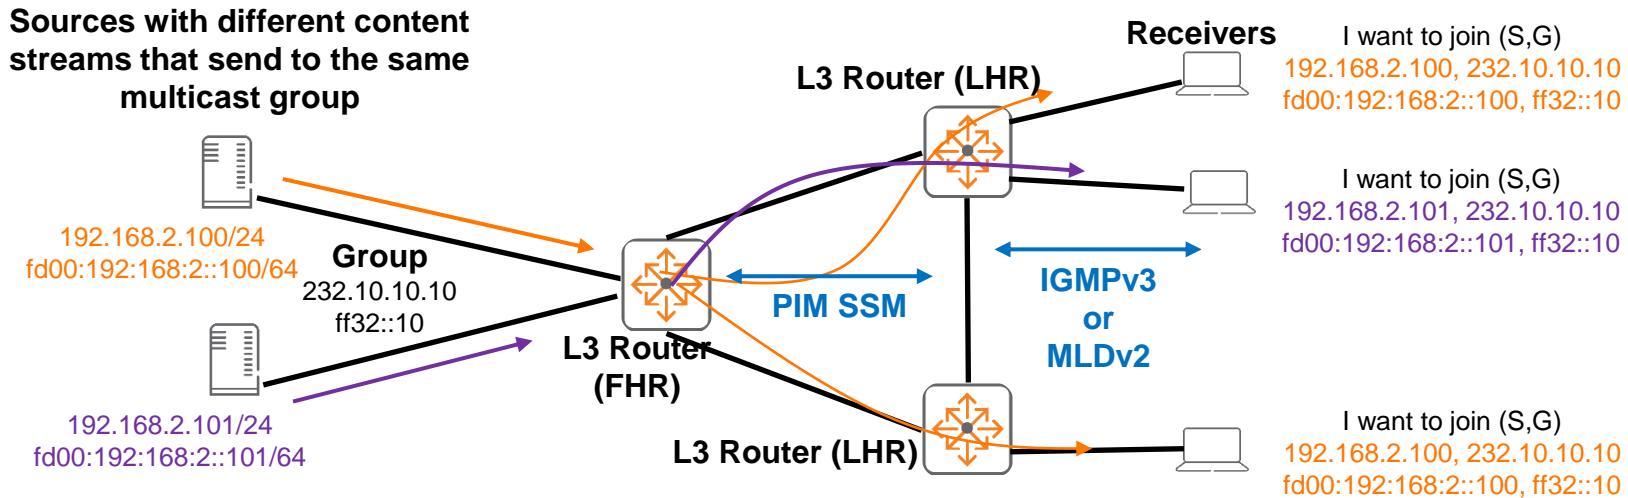
PIM SSM(Source Specific Multicast) Overview

- 10.9 adds PIM SSM support for non VXLAN deployments
- Allows multicast receivers to request for a multicast stream from a specific source
 - Only (S, G) is used, (*, G) is not used
 - PIM SSM suitable for one source stream to many receivers model/applications
- PIM SSM no longer requires rendezvous points (RPs)
- Both IPv4 and IPv6 SSM are supported
- Requires IGMPv3 (IPv4) or MLDv2 (IPv6) between receivers and LHRs to request for specific multicast source
- Supported platforms:
 - 6200, 6300, 6400, 8320, 8325, 8360, 8400, CX10000



Use Cases

PIM SSM Use Case

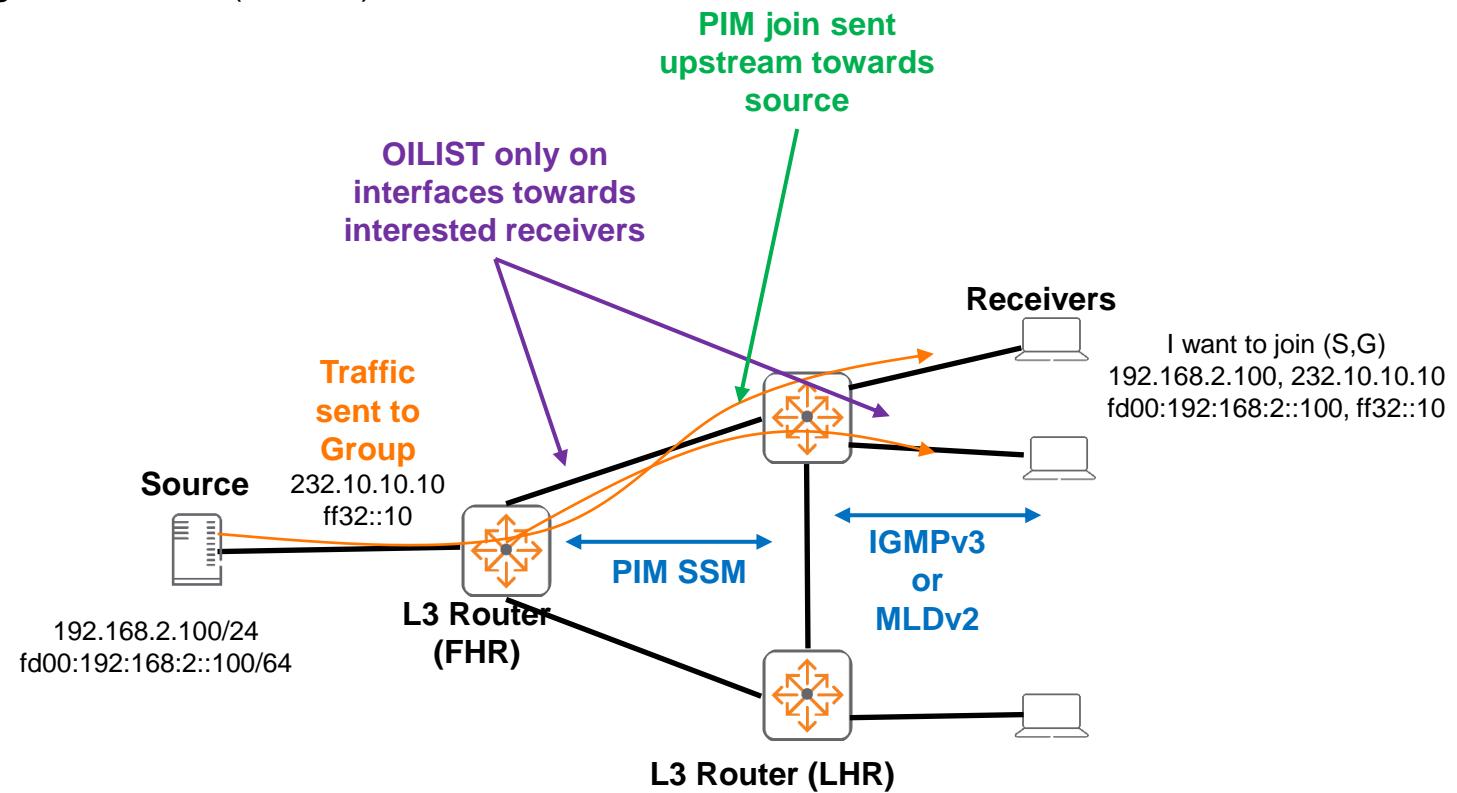


- Allows the same multicast group to be used by different sources without interference
- Or when the source IP is known
- Receivers typically join their interested stream via a web page

Details

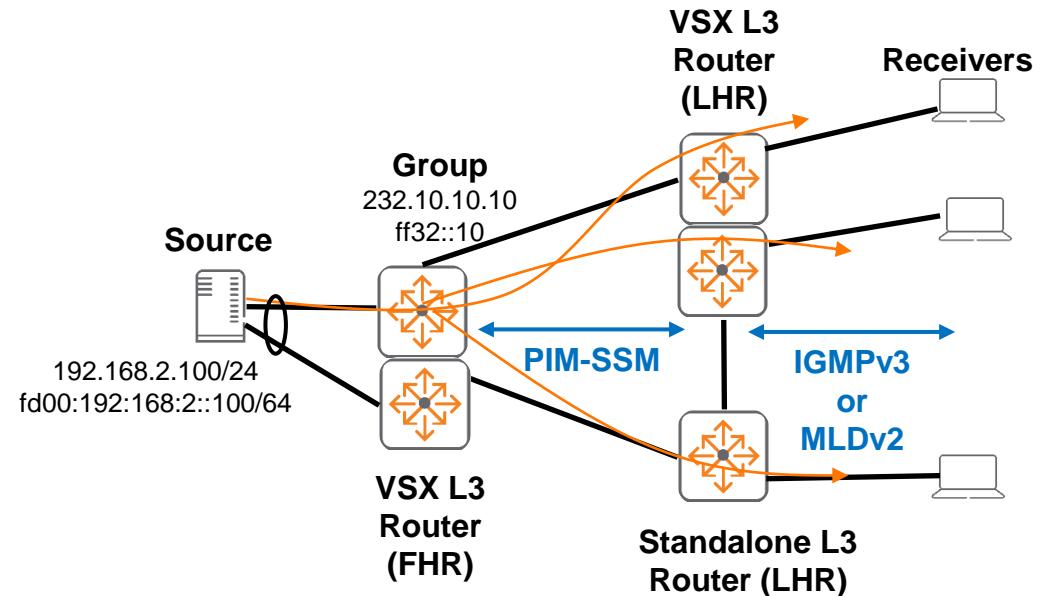
Multicast Routing and Traffic Forwarding

- Dependent on unicast routing, ensure receiver has L3 connectivity to source
- IGMPv3 or MLDv2 joins are sent from receiver to source connected LHR
- PIM joins are sent upstream towards multicast source/FHR
- Multicast traffic is only forwarded on L3 interfaces towards interested receivers
 - Mroute Outgoing interface list (OILIST) will state this



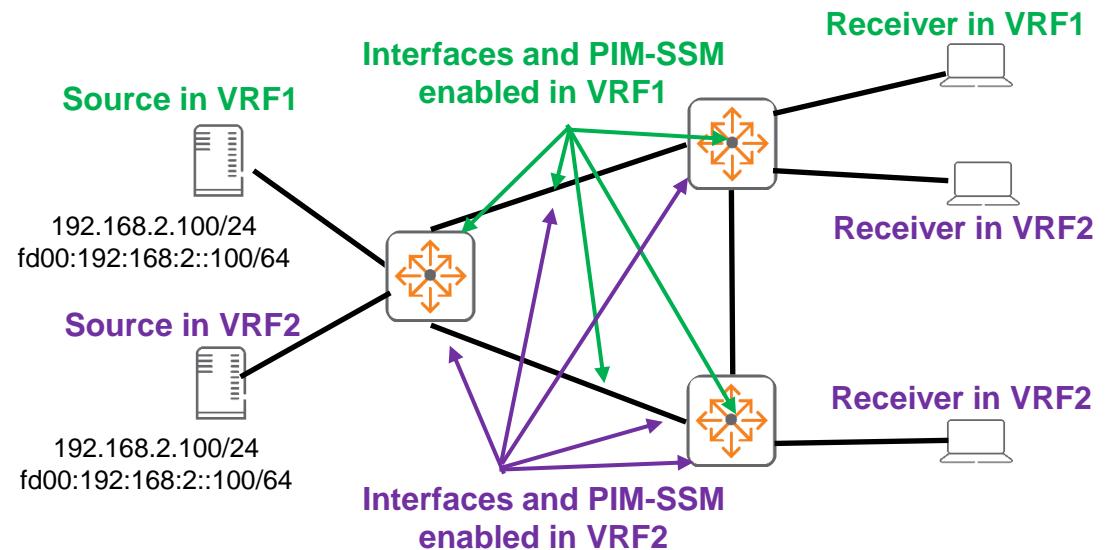
VSX Support

- VSX is supported on both
 - Switches connected to sources (FHR)
 - Switches connected to receivers (LHR)



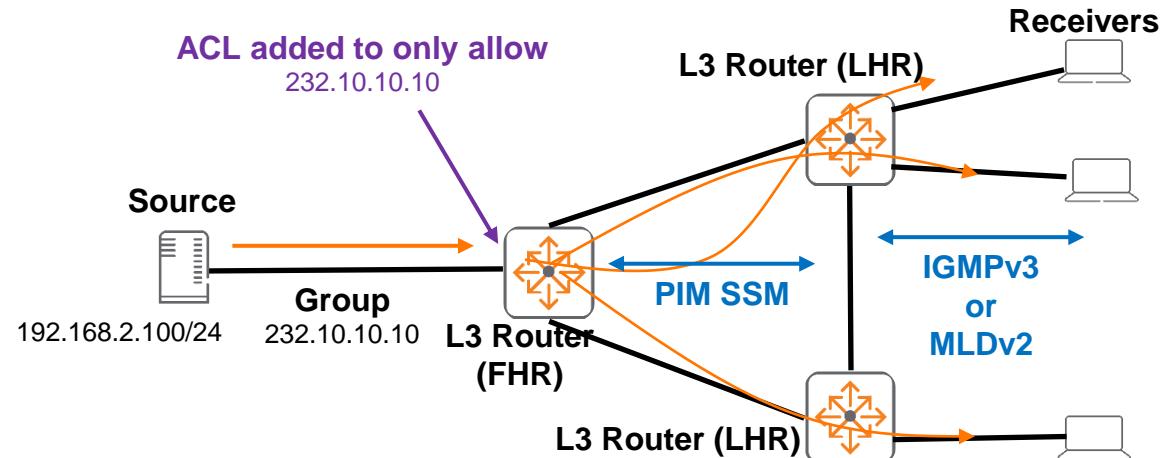
VRF Support

- VRFs are supported if required for multi-tenancy
- VRFs with PIM have to be enabled on every link/router between source and receivers
- Sub-interfaces between AOS-CX switches can also be used



PIM SSM ACL support

- Default IPv4 SSM address range = 232.0.0.0/8
- Default IPv6 SSM address range = FF3x::/32
- ACLs can be used together with PIM SSM
 - To allow desired groups to be forwarded
 - To block certain groups from being forwarded
 - To modify the default ssm range to a user defined range
 - e.g. 233.0.0.0/8 or FF4x::/32



PIM SSM and MSDP interaction

- For groups within the SSM range, no MSDP Source-Active (SA) messages within the SSM range will be accepted, generated, or forwarded
 - Applicable for default SSM range and when SSM ACL is modified

PIM SSM Caveats

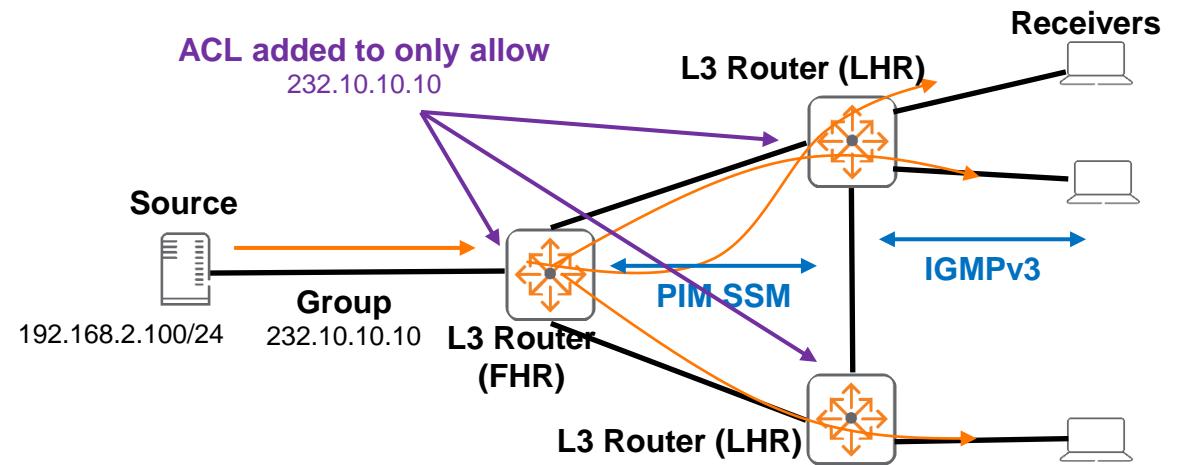
- It's not mandatory to only utilize PIM SSM
 - A parallel PIM SM deployment with BSR/RP can be used for non SSM address ranges
 - e.g. both PIM SSM supporting 232.0.0.0/8 and PIM SM supporting 239.0.0.0/8 with BSR/RP can be used on a network
- On VSX switches
 - PIM active-active should be enabled when clients are connected via downstream VSX LAG for faster convergence
 - Applicable to both PIM SSM and PIM SM
- PIM SSM ACL
 - Only "multicast group" parameter in ACL is supported
 - Need to add default SSM range into ACL if those are desired
- Requires IGMPv3 (IPv4) or MLDv2 (IPv6) between receivers and LHRs to request for specific multicast source
- PIM-SSM mapping is not supported today (planned for future AOS-CX release)
- PIM-SSM mapping will support receivers that are not IGMPv3 or MLDv2 capable

```
access-list ip ssm
    20 permit any any 239.10.10.10
    30 permit any any 232.0.0.0/255.0.0.0
```

Configuration

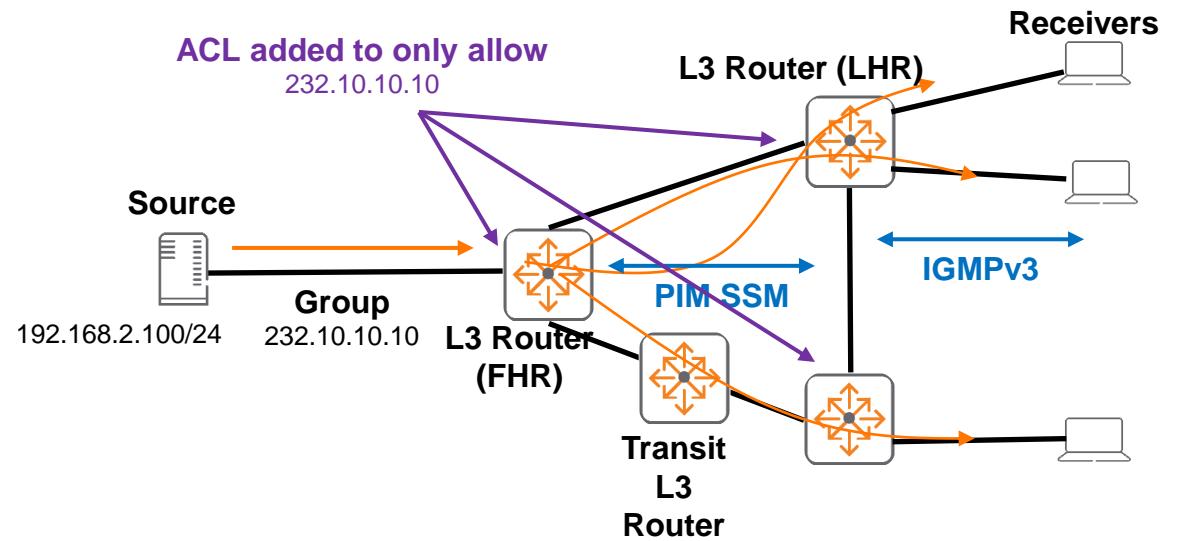
PIM SSM IPv4 Configuration – FHR example

```
vlan 1,20
!
access-list ip ssm
    20 permit any any 232.10.10.10
!
interface 1/1/49
    no shutdown
    description Downlink to LHR
    ip address 192.168.3.0/31
    ip ospf 1 area 0.0.0.0
    ip pim-sparse enable
!
interface vlan 20
    description Source subnet
    ip address 192.168.2.1/24
    ip ospf 1 area 0.0.0.0
    ip igmp enable
    ip pim-sparse enable
!
router ospf 1
    area 0.0.0.0
!
router pim
    enable
    pim-ssm
    pim-ssm range-access-list ssm
```



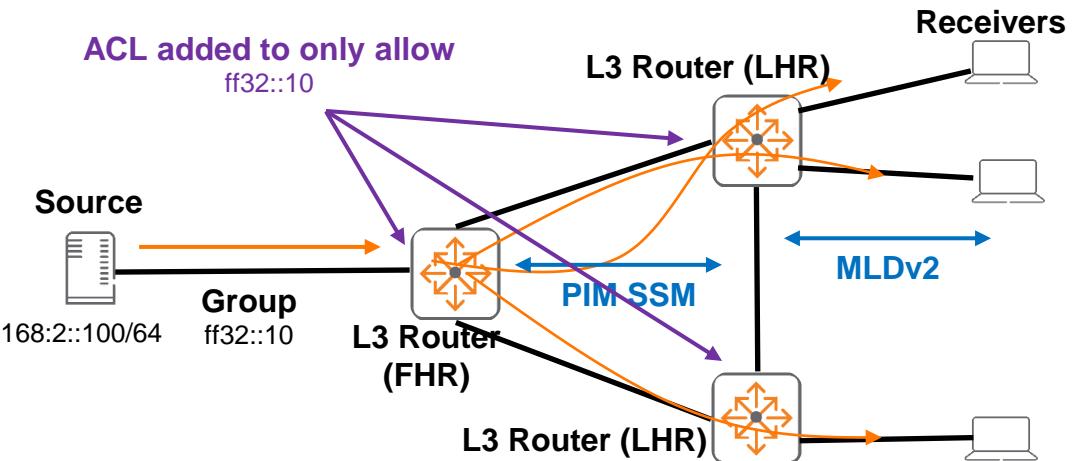
PIM SSM IPv4 Configuration – Transit router or LHR example

```
vlan 1,20
!
access-list ip ssm
    20 permit any any 232.10.10.10
1
interface 1/1/49
    no shutdown
    description Uplink to FHR
    ip address 192.168.3.1/31
    ip ospf 1 area 0.0.0.0
    ip pim-sparse enable
!
interface vlan 20
    description Receiver subnet
    ip address 20.1.1.220/24
    active-gateway ip mac 20:00:00:00:20:20
    active-gateway ip 20.1.1.254
    ip ospf 1 area 0.0.0.0
    ip igmp enable ← IGMPv3 used by default
    ip pim-sparse enable
!
router ospf 1
    area 0.0.0.0
!
router pim
    enable
    pim-ssm
    pim-ssm range-access-list ssm
```



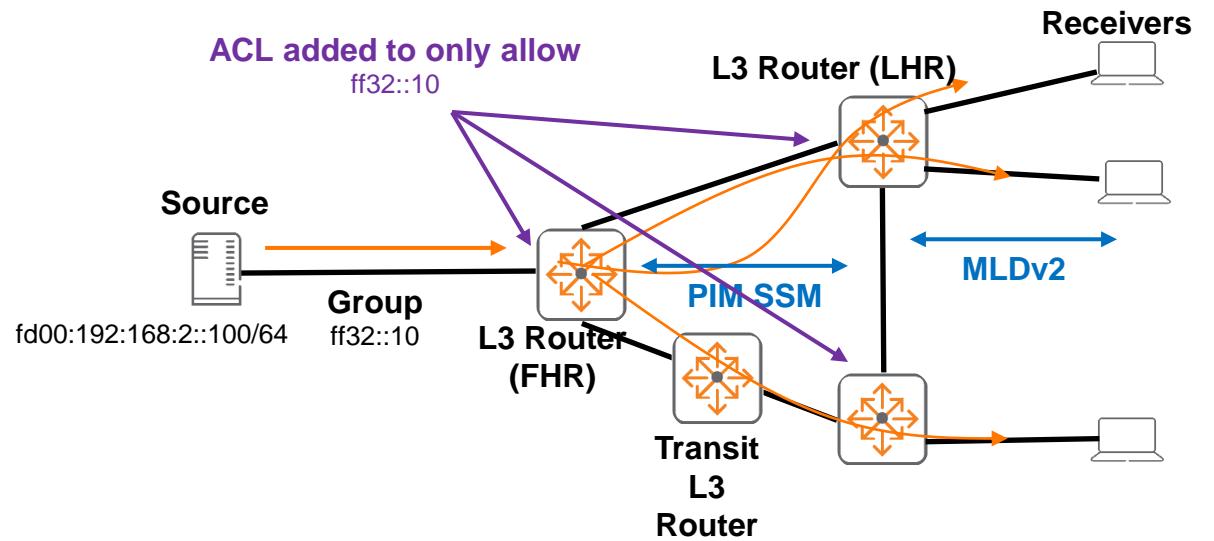
PIM SSM IPv6 Configuration – FHR example

```
vlan 1,20
!
access-list ipv6 ssm
    20 permit any any ff32::10
!
interface 1/1/49
    no shutdown
    description Downlink to LHR
    ipv6 address fd00:192:168:3::1/127
    ipv6 ospfv3 1 area 0.0.0.0
    ipv6 pim6-sparse enable
!
interface vlan20
    description Source subnet
    ipv6 address fd00:192:168:2::1/64
    ipv6 ospfv3 1 area 0.0.0.0
    ipv6 pim6-sparse enable
    ipv6 mld enable
!
router ospfv3 1
    area 0.0.0.0
!
router pim6
    enable
    pim-ssm
    pim-ssm range-access-list ssm
```



PIM SSM IPv6 Configuration – Transit router or LHR example

```
vlan 1,20
!
access-list ipv6 ssm
  20 permit any any ff32::10
!
interface 1/1/49
  no shutdown
  description Uplink to FHR
  ipv6 address fd00:192:168:3::1/127
  ipv6 ospfv3 1 area 0.0.0.0
  ipv6 pim6-sparse enable
!
interface vlan 20
  description Receiver subnet
  ipv6 address fd00:20:1:1::220/64
  active-gateway ipv6 mac 20:00:00:00:20:20
  active-gateway ipv6 fd00:20:1:1::254
  ipv6 ospfv3 1 area 0.0.0.0
  ipv6 mld enable ← MLDv2 used by default
  ipv6 pim6-sparse enable
!
router ospfv3 1
  area 0.0.0.0
!
router pim6
  enable
  pim-ssm
  pim-ssm range-access-list ssm
```



Best Practices

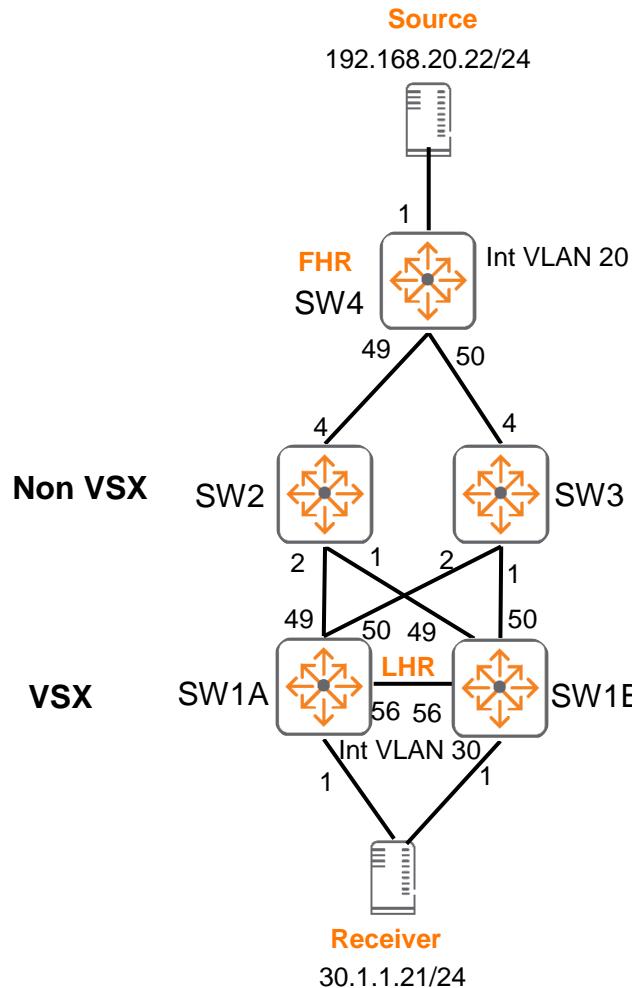
Best Practices

- Remember to
 - Enable “pim ssm” under “router pim”
 - Enable “pim sparse mode” between pim routers and on SVIs facing sources and receivers
 - Ensure default SSM address range is used by multicast source
 - ACL needs to be modified if non default is used, need to add default range into ACL if those are desired

Troubleshooting

PIM SSM Troubleshooting

- Have a topology diagram with interface and IP details ready
- Check physical cabling and generate “show tech” when opening a TAC case
- Check network: show LLDP neighbor, ensure unicast network works using ping and traceroute between loopbacks and interfaces, fix any issues found



– Recommended troubleshooting flow

1. Check unicast routing between source and receiver
2. Check PIM SSM configs are correct on LHR, transit routers and FHR
3. Check PIM neighbors are seen between PIM routers
4. Check IGMPv3 or MLDv2 joins are seen on receiver connected LHR
5. Check mroute is seen on LHR, transit routers and FHR

1. Check unicast routing between source and receiver

- In order for PIM SSM to correctly forward traffic from source towards receivers:
 - Ensure source/destination IPs are correctly advertised in the network
 - e.g. "sh ip route", test pings using between source and receiver IPs
 - Fix connectivity issues if discovered
 - This applies to both IPv4/IPv6 PIM SSM

```
C:\> Select Command Prompt

Connection-specific DNS Suffix . : 
Description . . . . . : Intel(R) 82574L Gigabit Network Connection
Physical Address . . . . . : 00-50-56-8E-FB-12
DHCP Enabled. . . . . : No
Autoconfiguration Enabled . . . . . : Yes
IPv6 Address. . . . . : fd00:30:1:1::21(Preferred)
Link-local IPv6 Address . . . . . : fe80::3481:f3bc:dea5:ca9f%13(Preferred)
IPv4 Address. . . . . : 30.1.1.21(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : fd00:30:1:1::254
                                         30.1.1.254
DHCPv6 IAID . . . . . : 100666409
DHCPv6 Client DUID. . . . . : 00-01-00-01-28-BF-3E-71-00-50-56-8E-FB-12
DNS Servers . . . . . : fec0:0:0:ffff::1%1
                                         fec0:0:0:ffff::2%1
                                         fec0:0:0:ffff::3%1
NetBIOS over Tcpip. . . . . : Enabled

C:\>ping 192.168.20.22

Pinging 192.168.20.22 with 32 bytes of data:
Reply from 192.168.20.22: bytes=32 time=49ms TTL=125
Reply from 192.168.20.22: bytes=32 time<1ms TTL=125
Reply from 192.168.20.22: bytes=32 time=5ms TTL=125
Reply from 192.168.20.22: bytes=32 time=4ms TTL=125

Ping statistics for 192.168.20.22:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 49ms, Average = 14ms

C:\>

C:\>ping fd00:192:168:20::22

Pinging fd00:192:168:20::22 with 32 bytes of data:
Reply from fd00:192:168:20::22: time=1ms
Reply from fd00:192:168:20::22: time=3ms
Reply from fd00:192:168:20::22: time=2ms
Reply from fd00:192:168:20::22: time<1ms

Ping statistics for fd00:192:168:20::22:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 3ms, Average = 1ms
```

2. Check PIM SSM configs are correct on LHR, transit routers and FHR

- Refer to config section for LHR, transit and FHR sample configs

3. Check PIM neighbors are seen between PIM routers

- In order for PIM traffic to correctly flow down the multicast tree
 - PIM neighbors should be up on all interfaces between PIM routers
 - This applies to both IPv4/IPv6 PIM SSM
- IPv6 equivalent command

```
SW1A# sh ipv6 pim6 neighbor
```

```
SW1A# sh ip pim neighbor

PIM Neighbor

VRF : default
Total number of neighbors : 4

IP Address : 20.1.1.221
Interface : vlan20
Up Time (HH:MM:SS) : 4 days 06:04:09
Expire Time (HH:MM:SS) : 00:01:41
DR Priority : 1
Hold Time (HH:MM:SS) : 00:01:45
Secondary IP Addresses :
  20.1.1.254

IP Address : 30.1.1.221
Interface : vlan30
Up Time (HH:MM:SS) : 4 days 06:04:09
Expire Time (HH:MM:SS) : 00:01:42
DR Priority : 1
Hold Time (HH:MM:SS) : 00:01:45
Secondary IP Addresses :
  30.1.1.254

IP Address : 192.168.3.4
Interface : 1/1/49
Up Time (HH:MM:SS) : 4 days 06:05:00
Expire Time (HH:MM:SS) : 00:01:21
DR Priority : 1
Hold Time (HH:MM:SS) : 00:01:45

IP Address : 192.168.3.12
Interface : 1/1/50
Up Time (HH:MM:SS) : 4 days 06:04:30
Expire Time (HH:MM:SS) : 00:01:24
DR Priority : 1
Hold Time (HH:MM:SS) : 00:01:45
```

4. Check IGMPv3 or MLDv2 joins are seen on LHR

- In order for PIM traffic to correctly flow down the multicast tree
 - IGMPv3 or MLDv2 joins should be seen on LHR
 - This applies to both IPv4/IPv6 PIM SSM
- IPv6 equivalent command

```
SW1A# sh ipv6 mld groups
```

```
SW1A# sh ip igmp groups

IGMP group information for group 239.255.255.250

Interface Name      : vlan20
VRF Name            : default

Group Address       : 239.255.255.250
Last Reporter        : 20.1.1.21

          V1      V2      Sources      Sources
Vers Mode Uptime    Expires   Timer     Timer   Forwarded Blocked
----- -----
3   EXC  4d 6h 12m 3m 43s

IGMP group information for group 232.10.10.10

Interface Name      : vlan30
VRF Name            : default

Group Address      : 232.10.10.10
Last Reporter      : 30.1.1.21

          V1      V2      Sources      Sources
Vers Mode Uptime    Expires   Timer     Timer   Forwarded Blocked
----- -----
3   INC  52m 26s    3m 36s           1

Group Address      : 232.10.10.10
Source Address     : 192.168.20.22

Mode Uptime      Expire
----- -----
INC  52m 28s    3m 34s
```

4. Check IGMPv3 or MLDv2 joins are seen on LHR

- Packet captures might be required to check if receiver sends joins correctly to LHR

17 11.571125	30.1.1.21	224.0.0.22	IGMPv3	58 Membership Report / Group 232.10.10.10, new source {192.168.2...
18 11.637364	30.1.1.21	224.0.0.22	IGMPv3	58 Membership Report / Group 232.10.10.10, new source {192.168.2...
19 12.394914	192.168.20.22	232.10.10.10	UDP	1370 65492 → 5004 Len=1328
20 12.394916	192.168.20.22	232.10.10.10	UDP	1370 65492 → 5004 Len=1328
21 12.398786	192.168.20.22	232.10.10.10	UDP	1370 65492 → 5004 Len=1328
22 12.398788	192.168.20.22	232.10.10.10	UDP	1370 65492 → 5004 Len=1328
23 12.398932	192.168.20.22	232.10.10.10	UDP	1370 65492 → 5004 Len=1328
24 12.398933	192.168.20.22	232.10.10.10	UDP	1370 65492 → 5004 Len=1328


```
> Frame 17: 58 bytes on wire (464 bits), 58 bytes captured (464 bits) on interface \Device\NPF_{E417155C-6750-4480-B8C6-9C1CFC9EF982}, id 0
> Ethernet II, Src: VMware_8e:fb:12 (00:50:56:8e:fb:12), Dst: IPv4mcast_16 (01:00:5e:00:00:16)
> Internet Protocol Version 4, Src: 30.1.1.21, Dst: 224.0.0.22
▼ Internet Group Management Protocol
    [IGMP Version: 3]
    Type: Membership Report (0x22)
    Reserved: 00
    Checksum: 0x122a [correct]
    [Checksum Status: Good]
    Reserved: 0000
    Num Group Records: 1
▼ Group Record : 232.10.10.10 Allow New Sources
    Record Type: Allow New Sources (5)
    Aux Data Len: 0
    Num Src: 1
    Multicast Address: 232.10.10.10
    Source Address: 192.168.20.22
```

- Config to mirror traffic

```
mirror session 1
enable
destination interface 1/1/40
source interface 1/1/51 both
```

5. Check mroute is seen on LHR, transit routers and FHR

- Ensure expected incoming and outgoing interfaces are seen on each router

– FHR

```
SW4# sh ip mroute
IP Multicast Route Entries

VRF : default
Total number of entries : 1

Group Address          : 232.10.10.10
Source Address         : 192.168.20.22
SSM Mroute            : True
Neighbor                 :
Incoming interface     : vlan20
Outgoing Interface List:
Interface    State
-----
1/1/49      forwarding
```

– Transit router

```
SW2# sh ip mroute
IP Multicast Route Entries

VRF : default
Total number of entries : 1

Group Address          : 232.10.10.10
Source Address         : 192.168.20.22
SSM Mroute            : True
Neighbor                 : 192.168.3.0
Incoming interface     : 1/1/4
Outgoing Interface List:
Interface    State
-----
1/1/1       forwarding
```

– LHR

```
SW1B#   sh ip mroute
IP Multicast Route Entries

VRF : default
Total number of entries : 1

Group Address          : 232.10.10.10
Source Address         : 192.168.20.22
SSM Mroute            : True
Neighbor                 : 192.168.3.2
Incoming interface     : 1/1/49
Outgoing Interface List:
Interface    State
-----
vlan30      forwarding
```

- IPv6 equivalent command

```
SW1B# sh ipv6 mroute
```

5. Check mroute is seen on LHR, transit routers and FHR

- On VSX switches, only the DR will state outgoing interface in mroute

```
SW1A# sh ip mroute
IP Multicast Route Entries

VRF : default
Total number of entries : 1

Group Address          : 232.10.10.10
Source Address         : 192.168.20.22
SSM Mroute             : True
Neighbor                :
Incoming interface      : vlan30
```

```
SW1B# sh ip mroute
IP Multicast Route Entries

VRF : default
Total number of entries : 1

Group Address          : 232.10.10.10
Source Address         : 192.168.20.22
SSM Mroute             : True
Neighbor                : 192.168.3.2
Incoming interface      : 1/1/49
Outgoing Interface List :
Interface      State
-----
vlan30           forwarding
```

- Method to check which VSX switch is DR for an SVI

```
SW1A# sh ip pim int vlan 30

PIM Interfaces

VRF: default

Interface : vlan30
Neighbor count : 1
IP Address : 30.1.1.220/24
Mode       : sparse
Designated Router : 30.1.1.221
Proxy DR      : false
Hello Interval (sec) : 30
Hello Delay (sec)  : 5
Override Interval (msec) : 2500
Propagation Delay (msec) : 500
Operational DR Priority : 1
Neighbor Timeout   : 80
Lan Prune Delay    : Yes
Configured DR Priority : 1
```

```
SW1B# sh ip pim int vlan 30

PIM Interfaces

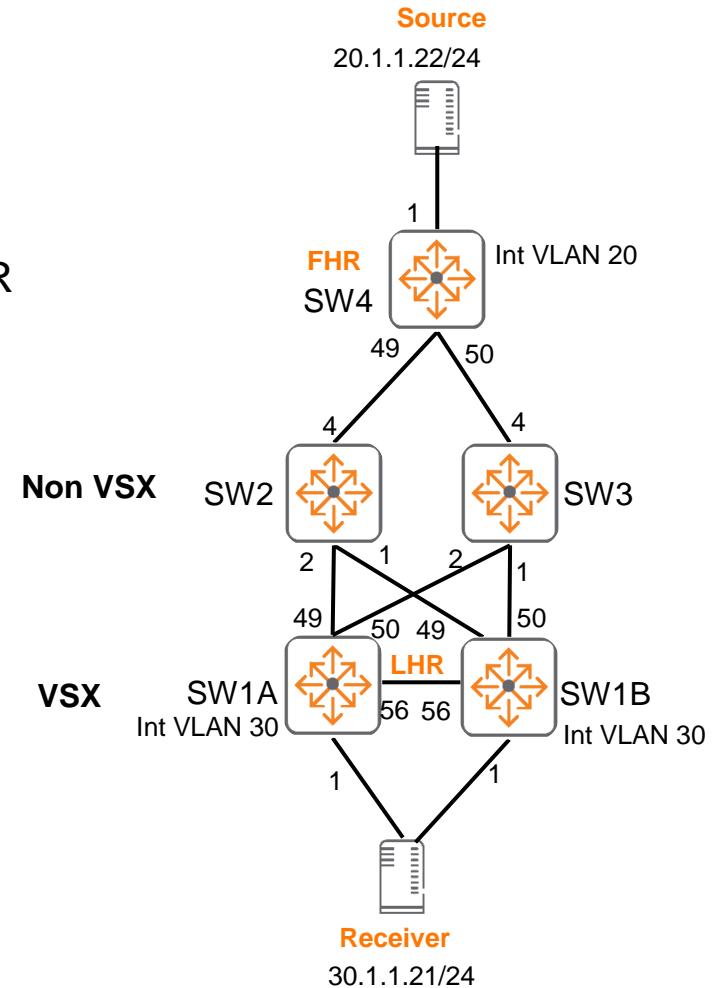
VRF: default

Interface : vlan30
Neighbor count : 1
IP Address : 30.1.1.221/24
Mode       : sparse
Designated Router : 30.1.1.221
Proxy DR      : false
Hello Interval (sec) : 30
Hello Delay (sec)  : 5
Override Interval (msec) : 2500
Propagation Delay (msec) : 500
Operational DR Priority : 1
Neighbor Timeout   : 100
Lan Prune Delay    : Yes
Configured DR Priority : 1
```

Demo

PIM SSM Demo

- Demo flow
 - Show mroute at FHR
 - Join multicast stream at receiver
 - Show IGMPv3 join at LHR
 - Show updated mroute at LHR, transit router and FHR
 - Show failover of PIM link and impact to receiver



Thank you

daryl.wan@hpe.com