AOS-CX 10.10 Update

aruba
a Hewlett Packard Enterprise company

PIM Multicast Boundary

Presenters
- Rahim Raoufi
- Daryl Wan

Agenda

1. Overview
2. Use Cases
3. Details and Caveats
4. Configuration
5. Best Practices
6. Troubleshooting
7. Demo
8. Additional Resources

Overview

PIM Multicast Boundary Overview

Allows better control over the network and prevents unwanted Multicast traffic flow across boundaries based on policies defined by the administrator.

The administrator should use the following on the boundary switches to achieve the administrative boundaries.

- pim-sparse bsr-boundary
- multicast boundary access-list

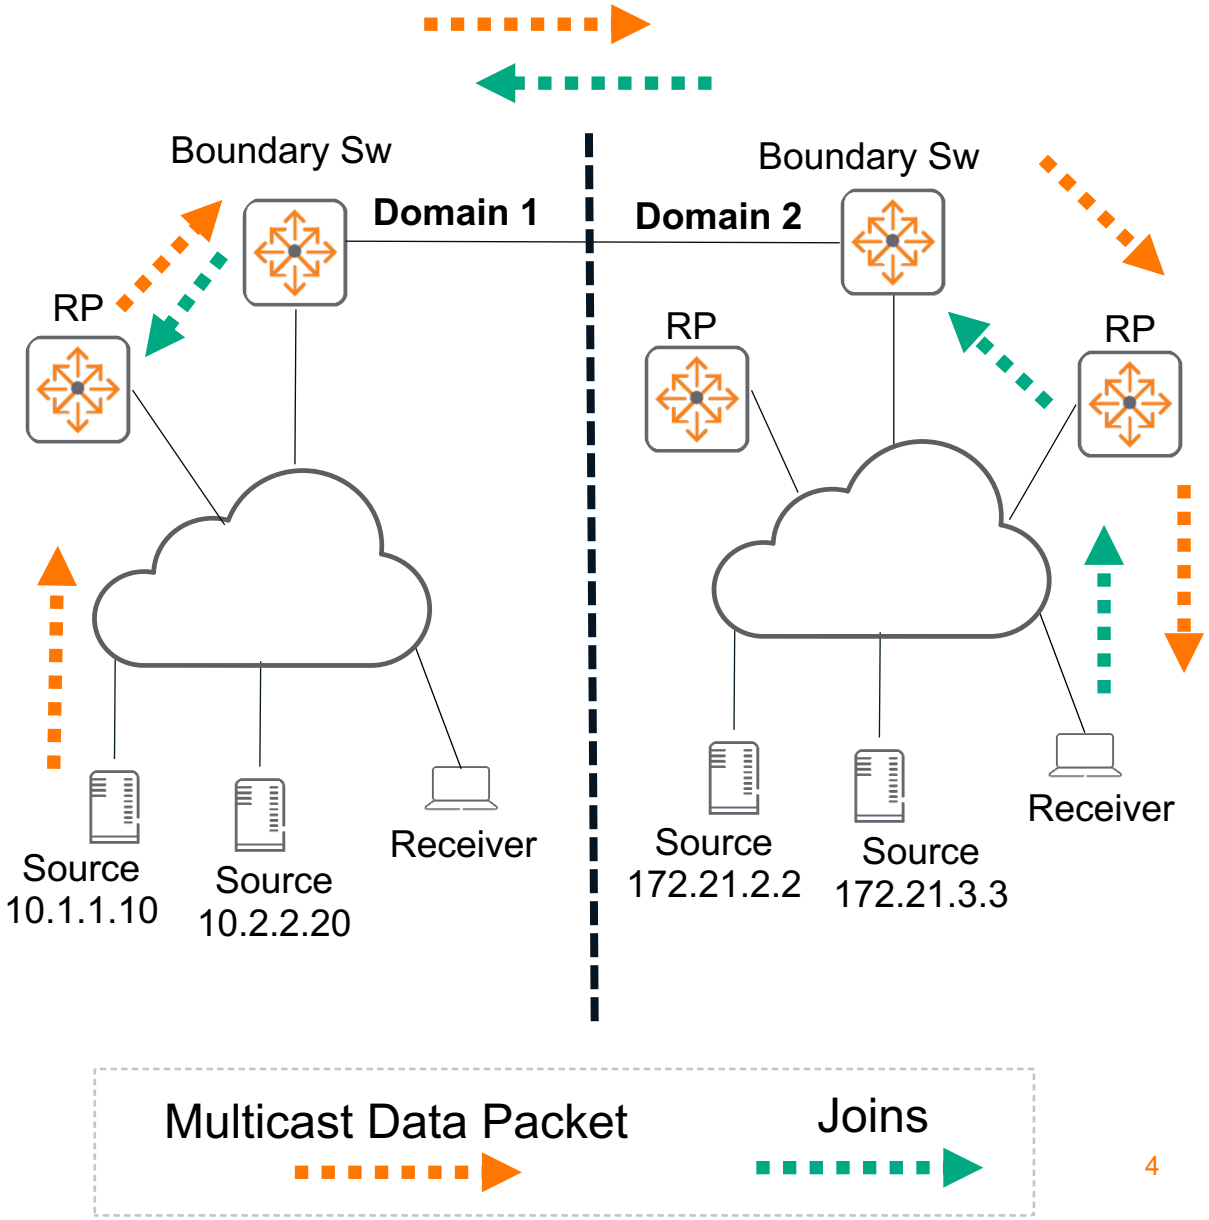

PIM Multicast Boundary Overview

BSR-boundary

- This feature helps to create multiple PIM domains by filtering **PIM BSM messages**.
- PIM RP advertisements and registrations are NOT exchanged across domains. So, each domain will elect its **own RP.**

Boundary access-list

- To restrict or allow specific multicast data traffic across domains and prevents routing of multicast traffic on the boundary interface, specify ACL filters.
- **Multicast data traffic** boundary is created with the help of boundary **ACLs**, and can filter:
  - PIM control packets like Joins/Prunes/Asserts
  - IGMP membership reports
  - Boundary ACL can match on:
    - *, G flows
    - S, G flows
    - Subnets
    - Protocols (IGMP/PIM)
- 128 ACEs are supported per ACL.

Multicast boundary can be configured on the following **interfaces:**

- ROP
- L3 LAG
- Point-to-Point SVI
- Sub-Interface

Supported Platforms

- 6300,6400,8320, 8325, 8360,8400 and 10000

Use Cases

Deployment options and Solutions

MSDP Review

MSDP is used in two deployment scenarios:

- Between PIM-SM domains (**Inter-domain scenario**)
- Within a PIM-SM domain (**Intra-domain scenario**)

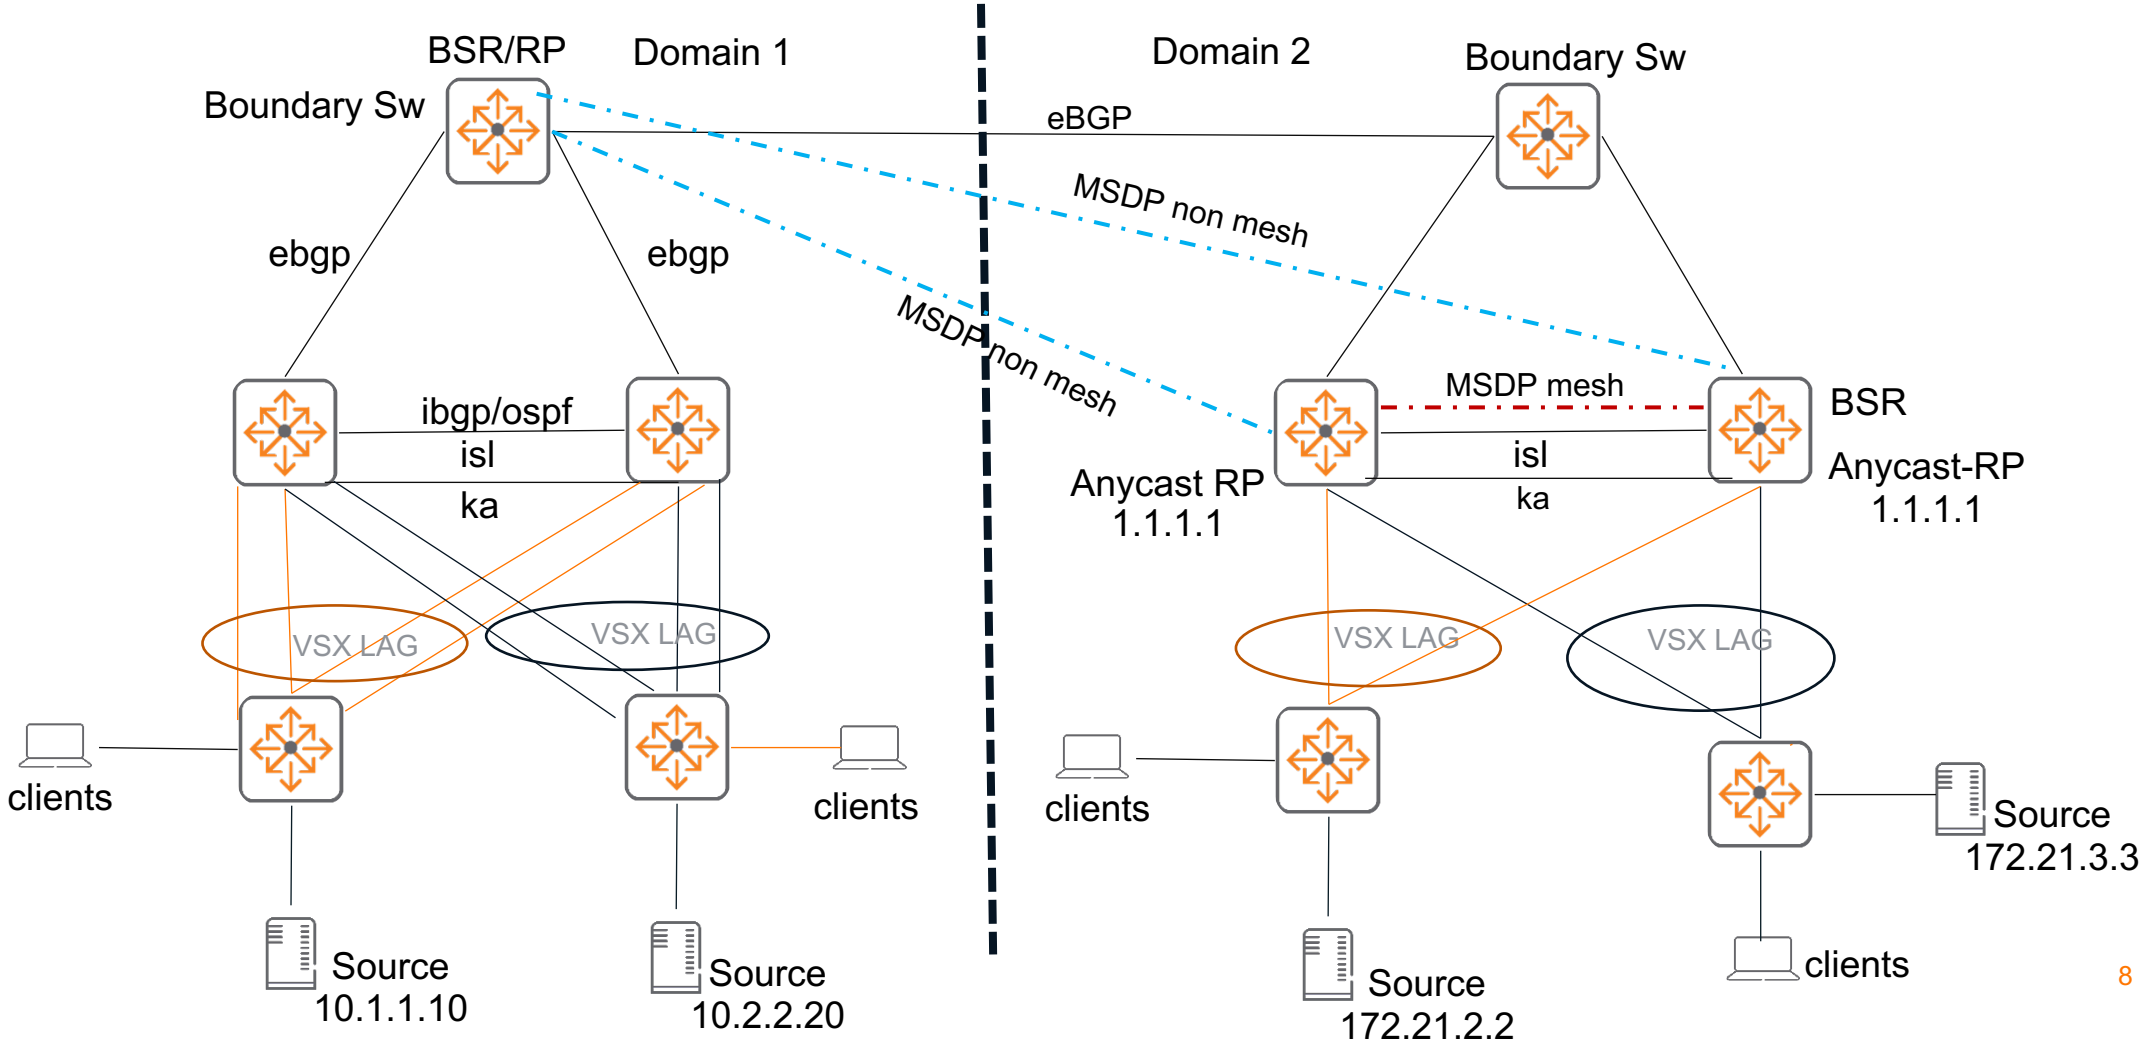

Combined MSDP and Boundary Operations

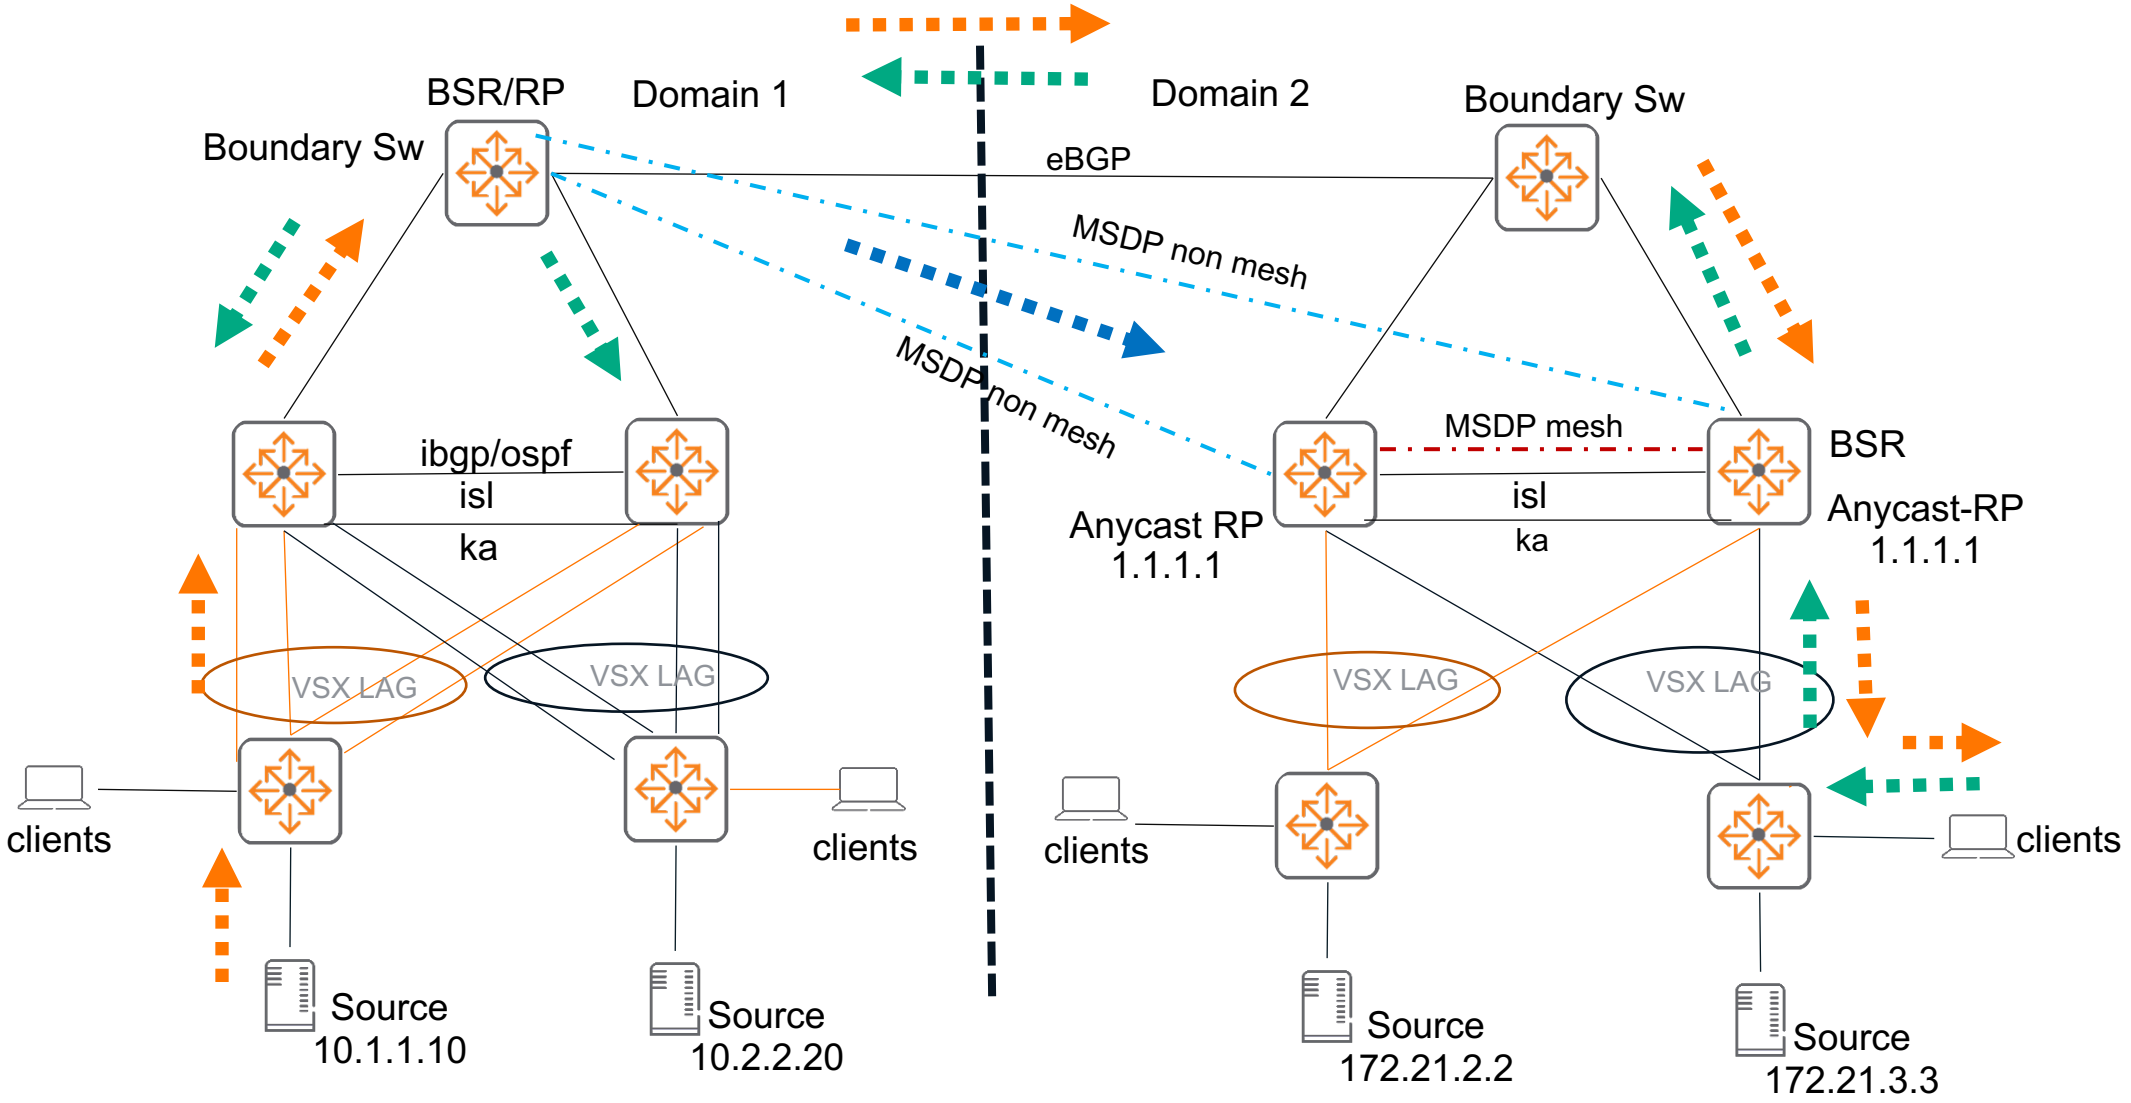

Details and Caveats

Solution 1 – Boundary on VSX with ROP interface

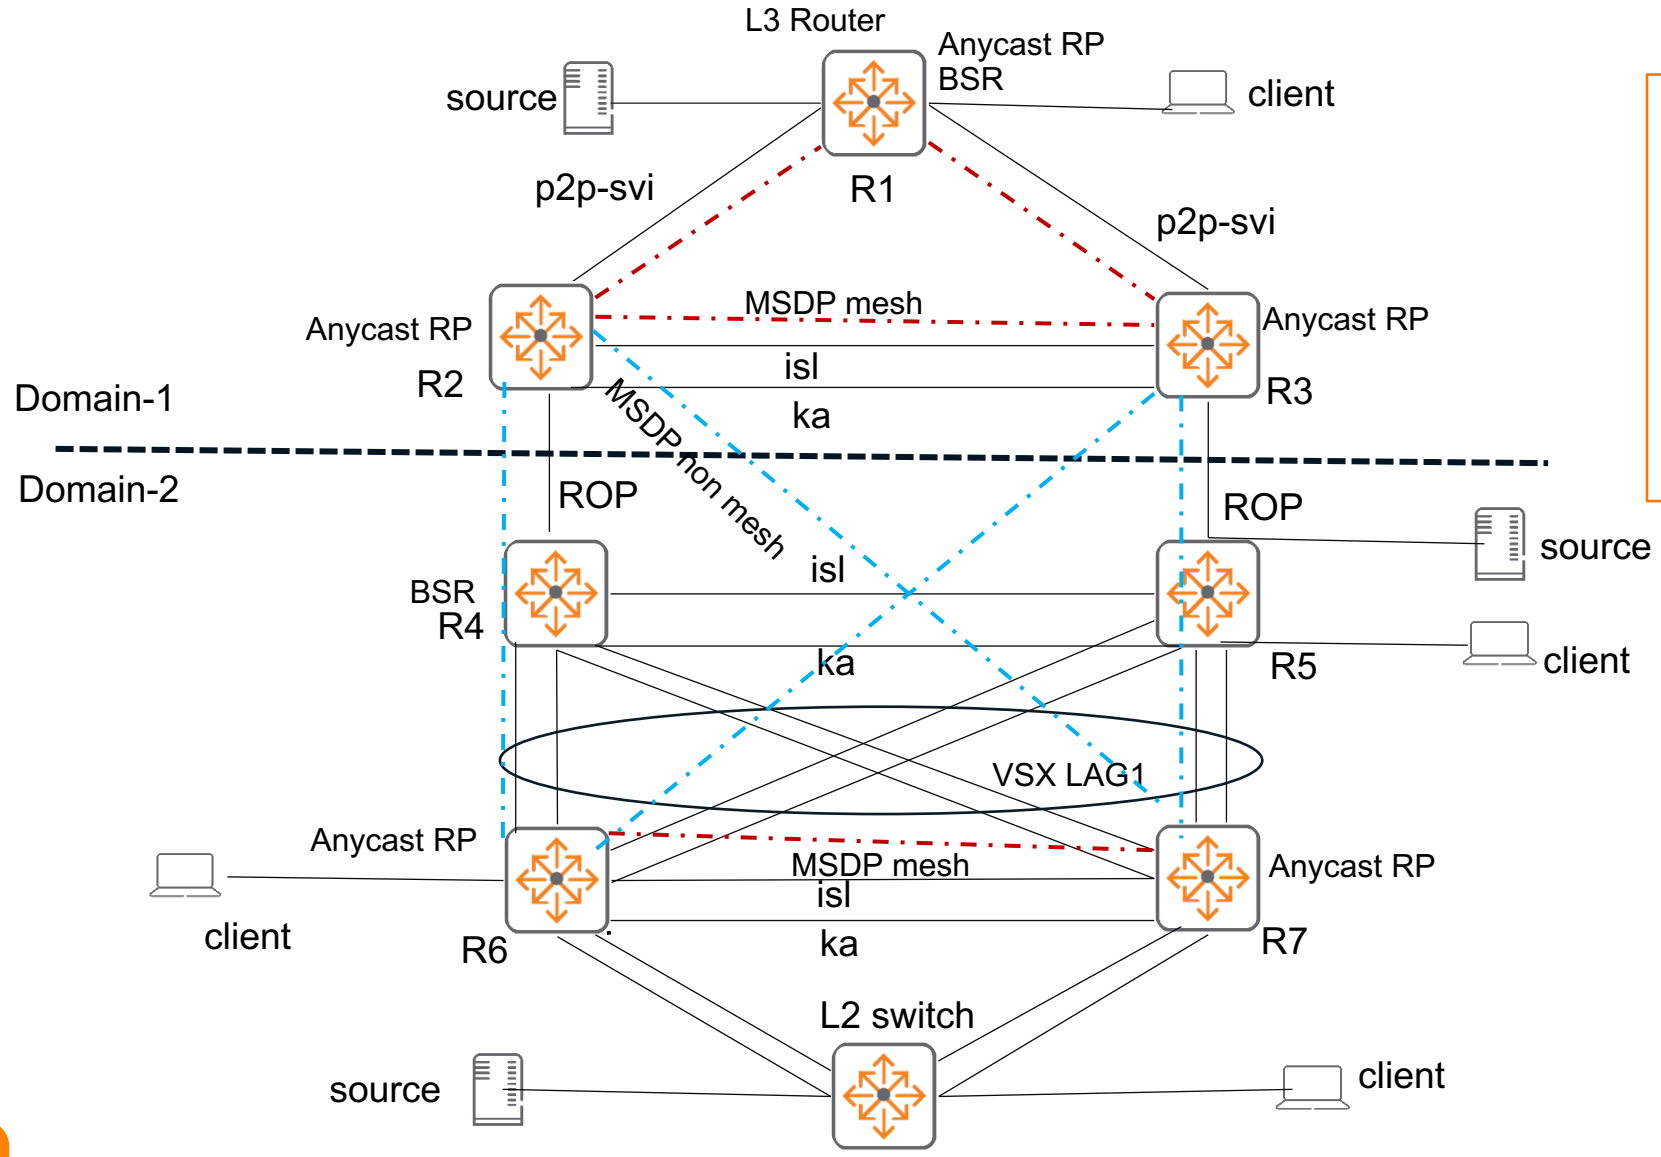

The multicast boundary can be configured on the following **interfaces:**

- ROP
- L3 LAG
- Point-to-Point SVI
- Sub-Interface

Solution 2 – Interop with PIM-SSM

PIM-SSM works across the boundary. PIM-SSM flows don't need MSDP to work.

If you also have non-SSM flows in the network and SSM flows, you still need MSDP for non-SSM flows to work. This solution validates both SM and SSM flows.

In the case of a pure PIM-SSM network, there is no RP configured, so the BSR-boundary command is not required; only boundary ACL is needed to filter the flows in such networks.

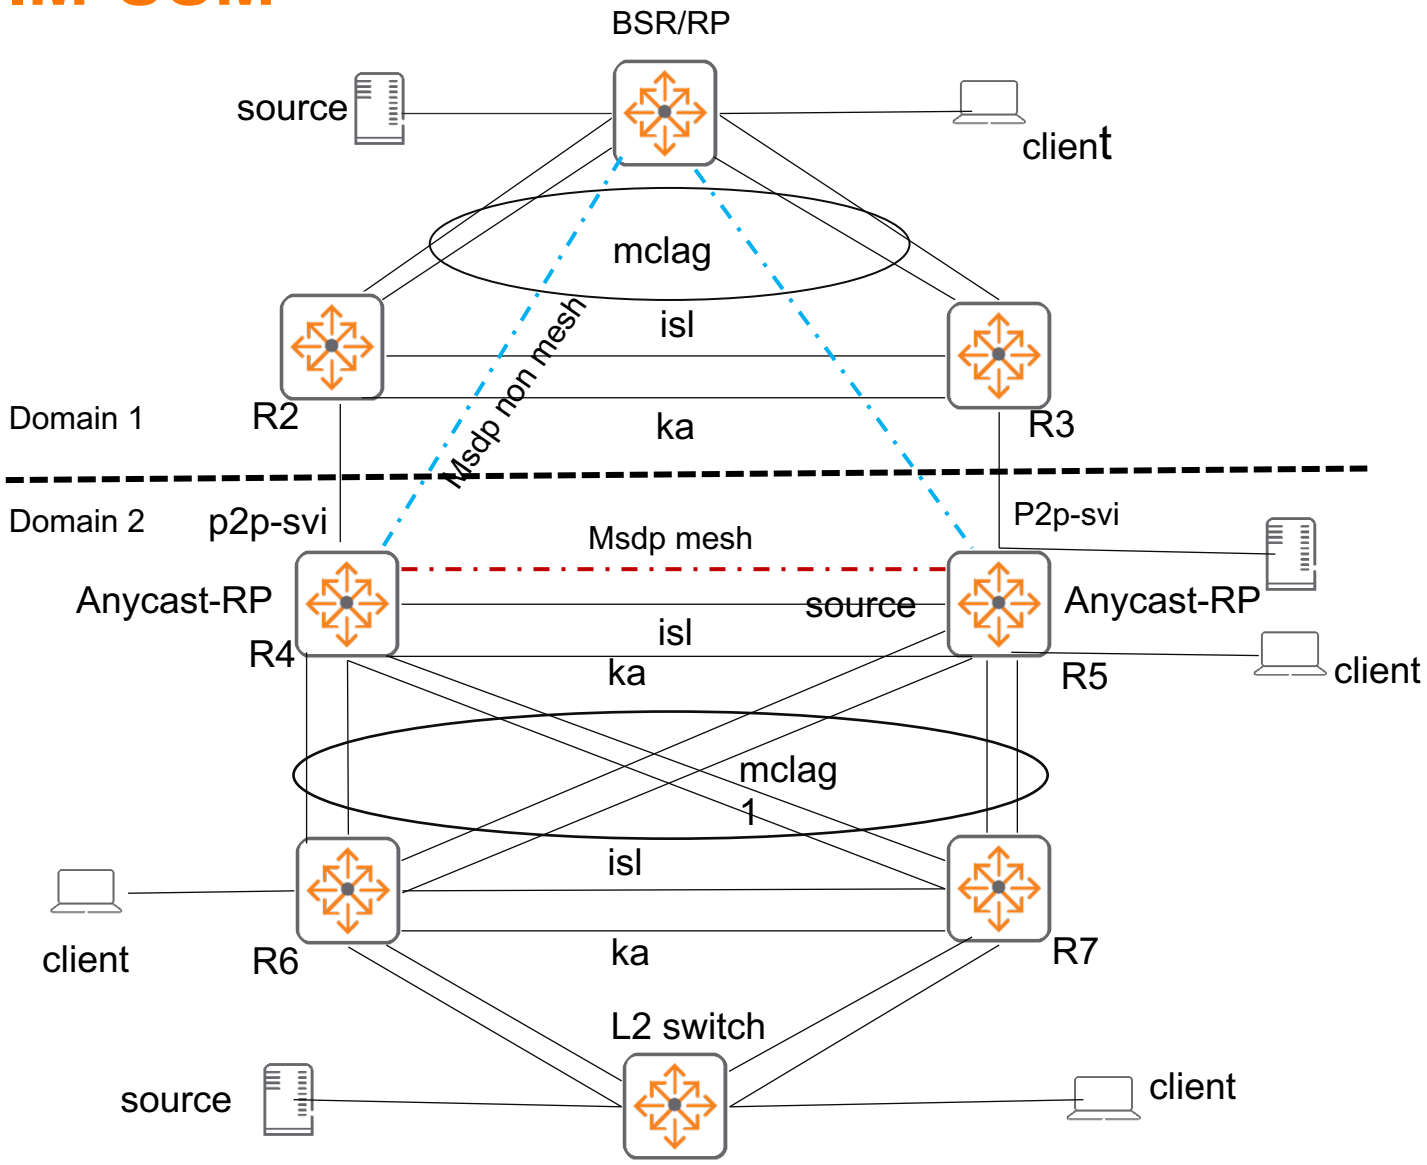

Solution 3 Combining with VXLAN

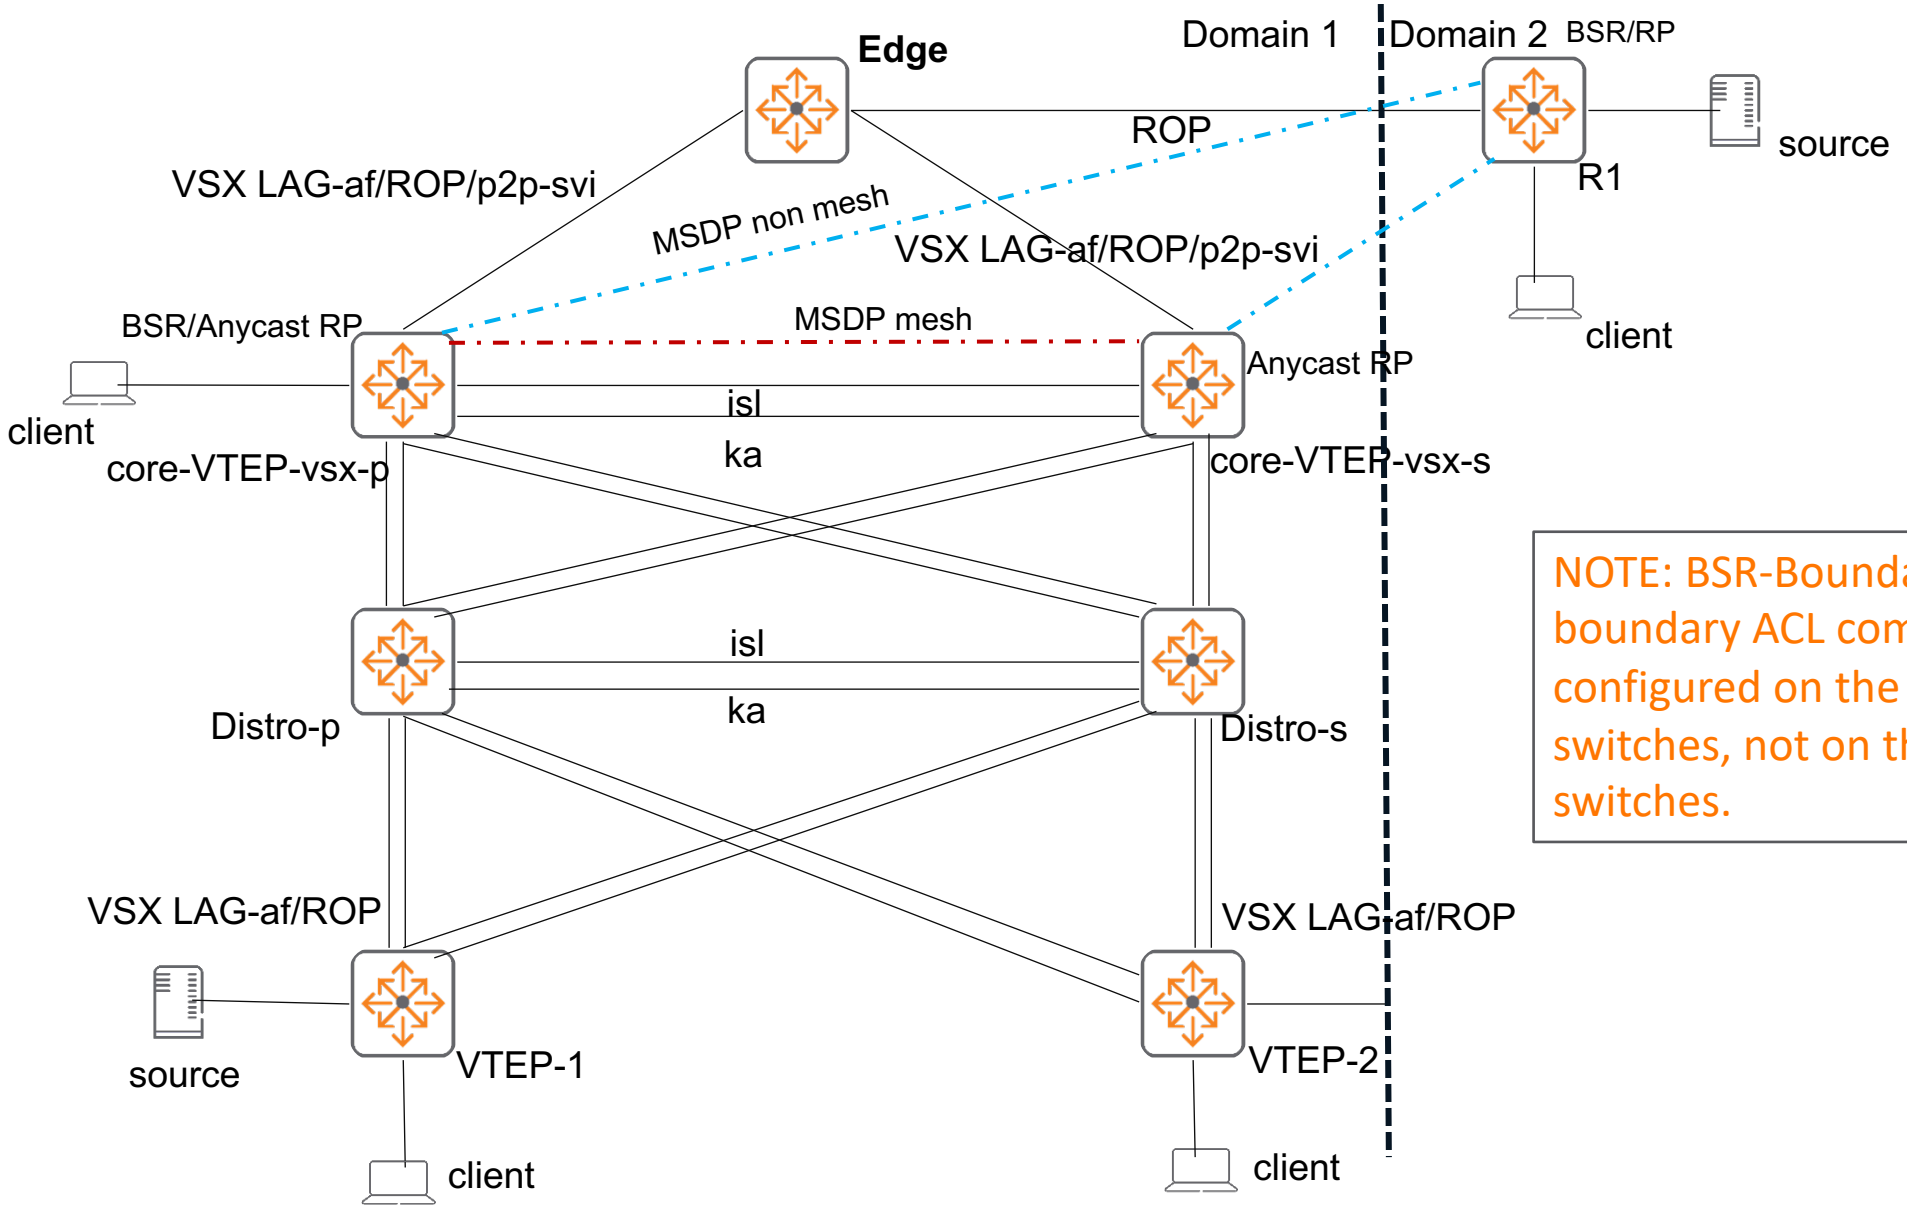

NOTE: BSR-Boundary and boundary ACL commands are configured on the edge routing switches, not on the VTEP routing switches.

Caveats

Unsupported Platforms

- IPv6 PIM boundary is not supported
- PIM boundary interop with multicast-over-VXLAN is not supported

Configuration

New CLIs

ip pim-sparse bsr-boundary

ip multicast boundary access-list

show ip multicast boundary interface

PIM BSR boundary, Candidate-RP advertisements are not exchanged across domains. So, each domain will elect its RP.

Usage

```
switch# configure terminal
switch(config)#interface <name >
switch(config-if)#[no] ip pim-sparse bsr-boundary
switch(config-if)# [no] ip multicast boundary access-list <acl_name>
```

To filter multicast data traffic
PIM control packets like Joins/Prunes/Asserts
IGMP membership reports
Boundary ACL can match on:
- *,G flows
- S,G flows
- Subnets
- Protocols (IGMP/PIM)

Verification

```
switch(config)#show ip multicast boundary interface <name>
```

Configuration example

Deny all sources for groups in range 239.0.0.0/8, Permit other streams:

```
access-list ip boundary1
    10 deny any any 239.0.0.0/255.0.0.0
    20 permit any any 224.0.0.0/240.0.0.0
interface lag 1
    ip address 40.1.1.1/24
    ip pim-sparse enable
    ip pim-sparse bsr-boundary
    ip multicast boundary access-list boundary1
```

NOTE: The ACL applied is applicable for both inbound/outbound direction

Configuration example

Permit specific S, G traffic and implicitly deny all other traffic.

```
access-list ip boundary2
    10 permit any 192.168.1.1 225.1.1.0/255.255.255.0
    20 permit any 172.168.1.1 239.1.1.0/255.255.255.0
interface 1/1/1
    ip address 40.1.1.1/24
    ip pim-sparse enable
    ip pim-sparse bsr-boundary
    ip multicast boundary access-list boundary2
```

NOTE: The ACL applied is applicable for both inbound/outbound direction

Multicast Source Directly Connected to VSX and Anycast-RP Switches

anycast-rp source-directly-connected

- When a source is directly connected to VSX via L2-VSX-LAG and VSX peers are configured as anycast-RP, the NEW CLI needs to be configured under router pim context on both VSX peers.

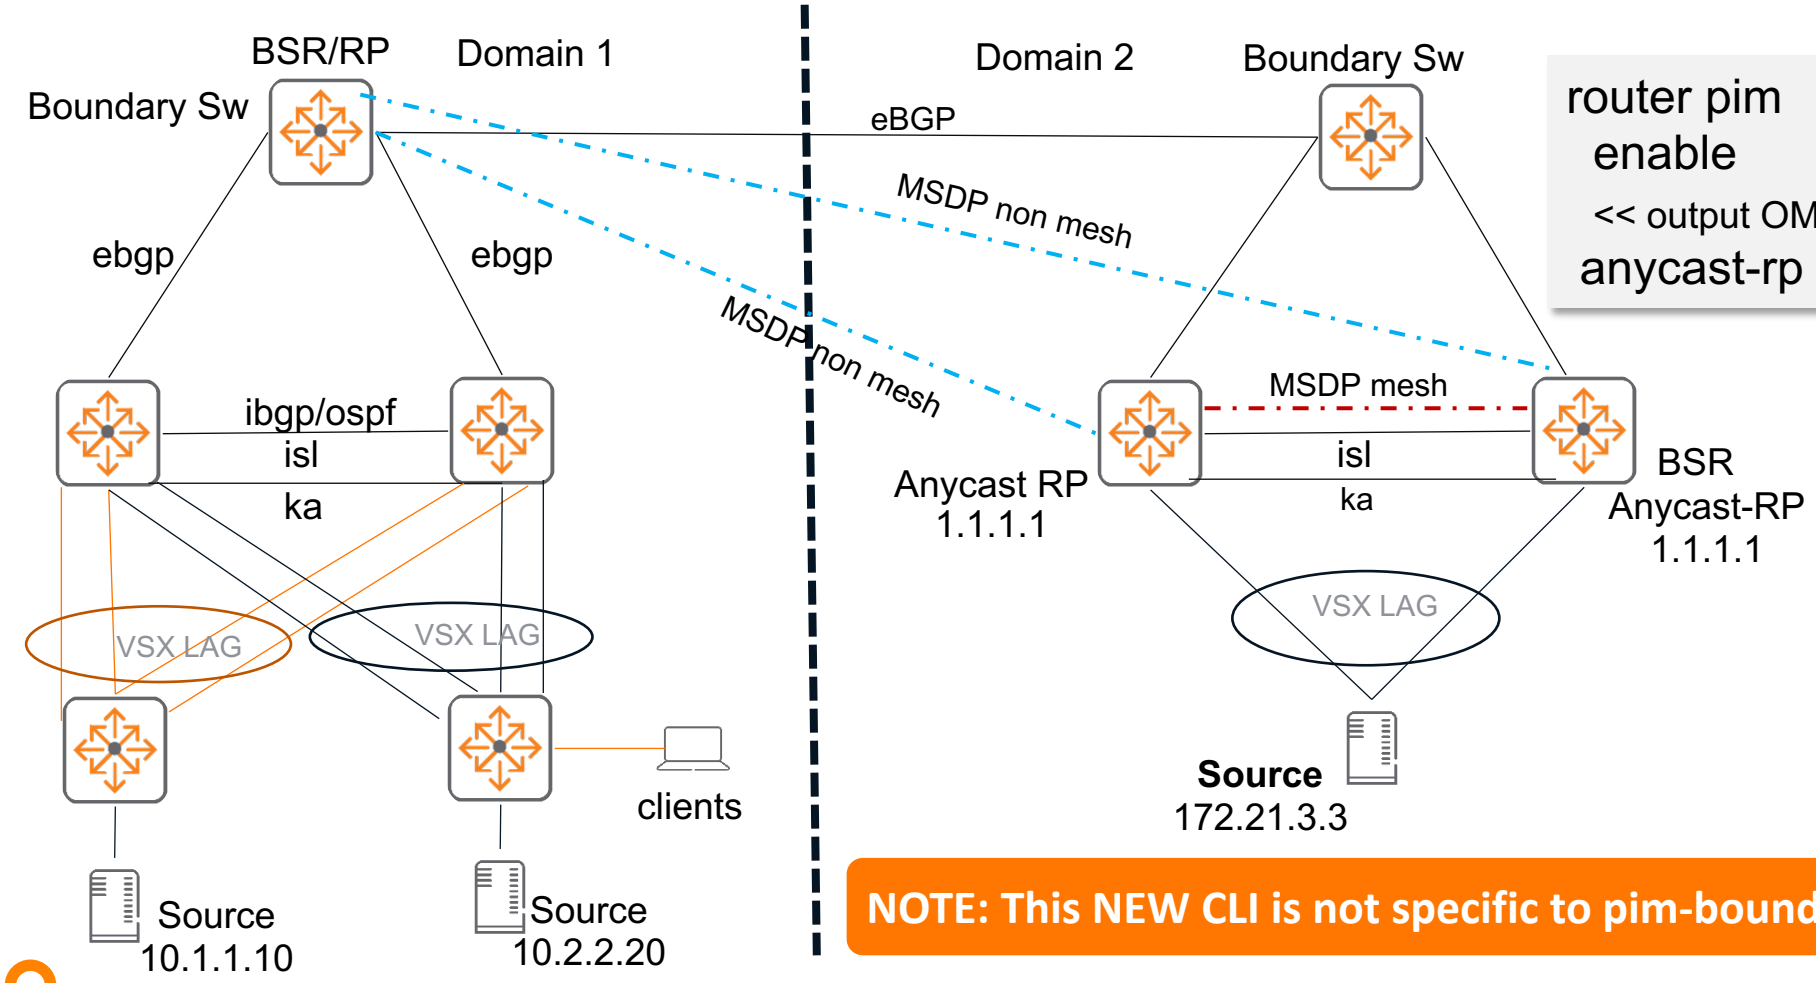


```
router pim
  enable
  << output OMITTED>>
  anycast-rp source-directly-connected
```

NOTE: This NEW CLI is not specific to pim-boundary but MSDP requirement.

Best Practices

Best Practices

- Always use both BSR boundary and multicast boundary ACL to avoid RPs across domains (unless it's an SSM only network or a single domain)

- The boundary ACLs need to be reviewed carefully by the admin as the ACL will define the traffic behavior. Note that there will be any implicit deny all in the end. So review and decide on the policies.

Troubleshooting

PIM boundary Troubleshooting

- Have a topology diagram ready
- Ensure IP interface details are included
- Check physical cabling and generate “show tech” when opening a TAC case
- Check network: Using show commands (`show ip multicast boundary interface <name>`), ensure directly connected PIM neighbors have connectivity by using ping and traceroute between loopbacks and interfaces, fix any issues found

Recommended troubleshooting flow

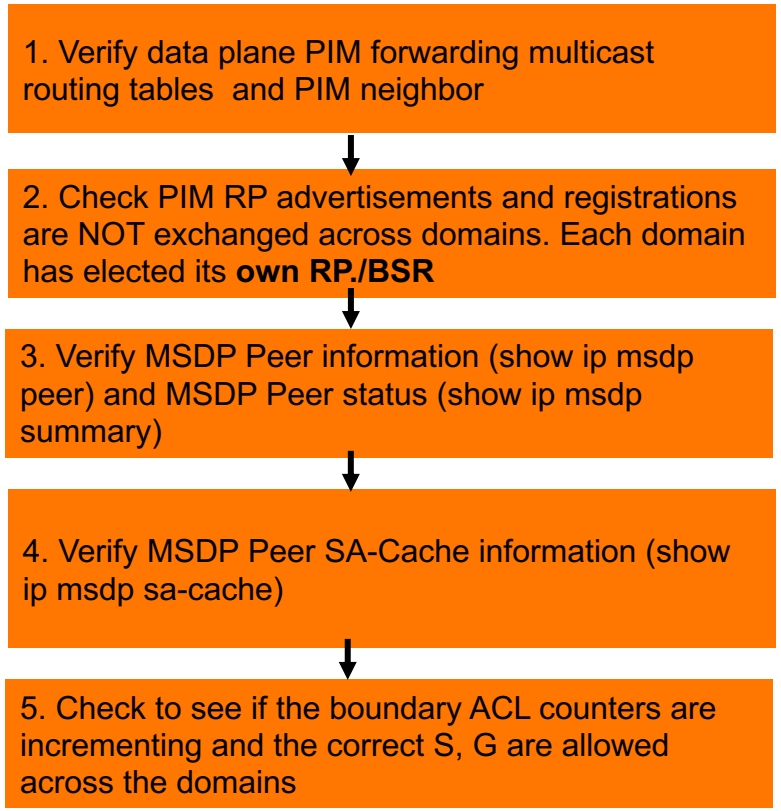

1. Verify data plane PIM forwarding multicast routing tables and PIM neighbor

```
SW2# show ip pim neighbor
PIM Neighbor

VRF                    : default
Total number of neighbors : 2

IP Address             : 192.168.100.0
Interface              : 1/1/3
Up Time (HH:MM:SS)     : 11 days 21:15:19
Expire Time (HH:MM:SS) : 00:01:39
DR Priority            : 1
Hold Time (HH:MM:SS)   : 00:01:45

IP Address             : 192.168.100.3
Interface              : 1/1/1
Up Time (HH:MM:SS)     : 1 days 06:44:33
Expire Time (HH:MM:SS) : 00:01:37
DR Priority            : 1
Hold Time (HH:MM:SS)   : 00:01:45
```

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

VRF : default
Total number of entries : 1

Group Address          : 239.1.1.1
Source Address         : 20.1.1.22
SSM Mroute             : False
Neighbor               : 192.168.100.2
Uptime                 : 1 days 06:47:13
State                  : route
Incoming interface     : 1/1/50
Outgoing Interface List :
Interface      State
-----------    ----------
vlan30         forwarding
```

2. Check PIM RP advertisements and registrations are NOT exchanged across domains. Each domain has elected its own RP/BSR

SW1# show ip pim rp-set

VRF: default

Status and Counters - PIM-SM Learned RP-Set Information

Group Address	Group Mask	RP Address	Hold Time	Expire Time
224.0.0.0	240.0.0.0	1.1.1.1	150	95

SW1# show ip pim bsr

Status and Counters- PIM-SM Bootstrap Router Information

VRF : default
E-BSR Address : 1.1.1.1
E-BSR Priority : 0
E-BSR Hash Mask Length : 30
E-BSR Up Time : 12 mins 45 secs
Next Bootstrap Message : 15 secs

C-BSR Admin Status : This system is a Candidate-BSR
C-BSR Address : 1.1.1.1
C-BSR Priority : 0
C-BSR Hash Mask Length : 30
C-BSR Message Interval : 60
C-BSR Source IP Interface : loopback0

C-RP Admin Status : This system is a Candidate-RP
C-RP Address : 1.1.1.1
C-RP Hold Time : 150
C-RP Advertise Period : 60
C-RP Priority : 192
C-RP Source IP Interface : loopback0

Group Address	Group Mask
224.0.0.0	240.0.0.0

3. Verify MSDP Peer information and status

```
SW1# show ip msdp summary
VRF: default

MSDP Peer Status Summary
Peer address     State    Uptime(Downtime)     Reset Count     SA Count
3.3.3.3          up       2h 23m               0               145


SW1# show ip msdp peer 3.3.3.3

VRF: default

MSDP Peer: 3.3.3.3
Connection status
State: up  Resets: 0  Connection Source: loopback0(1.1.1.1)
Uptime(Downtime): 2h 23m    SA Messages sent: 3
SA's learned from this peer: 145
Peer Keepalive interval: 60
Peer Hold time: 75
Peer Connection Retry interval: 30
```

4. Verify MSDP Peer SA-Cache information

```
SW1# show ip msdp sa-cache

VRF: default
Total entries: 1

(20.1.1.22, 239.1.1.1)  RP: 3.3.3.3  Peer: 3.3.3.3
```

5. Check to see if the boundary ACL counters are incrementing and the correct S, G are allowed across the domains

```
SW2(config)# show ip multicast boundary interface 1/1/1
IP Multicast Boundary Configurations
------------------------------------
access-list ip boundary
    10 permit any 20.1.1.22 239.1.1.0/255.255.255.0
    20 deny any any 239.0.0.0/255.0.0.0 count
    30 permit any any 224.0.0.0/240.0.0.0 count

IP Multicast Boundary Rx packet drop counters
-------------------------------------------
PIM Joins/Prunes            7
PIM BSM                     24
PIM C-RP Advertisements     0
PIM Asserts                 0
Multicast Data Packets      0
IGMP Joins                  0

Displays IP Multicast boundary ACL configurations and counters for the given
interface.
```

Demo

PIM Multicast Boundary Demo

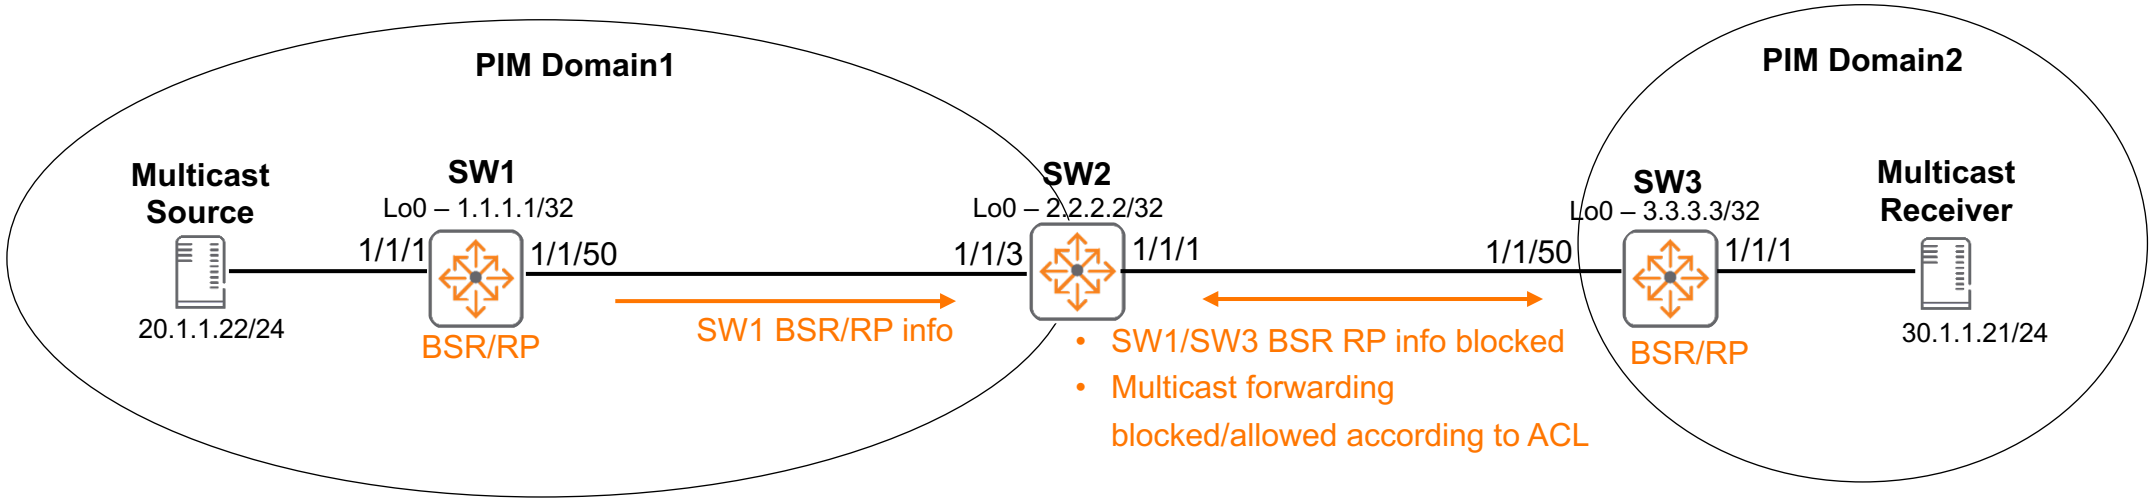

- Show before and after boundary commands are applied on SW2
- ACL is applicable to both inbound and outbound directions

```
SW2#
access-list ip boundary
    10 permit any 20.1.1.22 239.1.1.0/255.255.255.0
    20 deny any any 239.0.0.0/255.0.0.0
    30 permit any any 224.0.0.0/240.0.0.0
interface 1/1/1
    ip address 192.168.100.2/31
    ip pim-sparse enable
    ip pim-sparse bsr-boundary
    ip multicast boundary access-list boundary
```

Resources

Feature/Solution References

- Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)
  - **https://datatracker.ietf.org/doc/html/rfc5059**

- Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
  - **https://datatracker.ietf.org/doc/html/rfc3446**

- MSDP
  - https://en.wikipedia.org/wiki/Multicast_Source_Discovery_Protocol
  - *"RFC3618 - Multicast Source Discovery Protocol (MSDP)"*
  - *"RFC4611 - Multicast Source Discovery Protocol (MSDP) Deployment Scenarios"*
  - https://techhub.hpe.com/eginfolib/networking/docs/switches/7500/5200-1936a_ip-multi_cg/content/495505225.htm

- Multicast address
  - https://en.wikipedia.org/wiki/Multicast_address

aruba
a Hewlett Packard Enterprise company

Thank you

fardin.rah.raoufi@hpe.com
daryl.wan@hpe.com