AOS-CX 10.10 Update

aruba
a Hewlett Packard Enterprise company

# Multicast NSF


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

# Overview Multicast PIM NSF (Non-Stop Forwarding)

PIM NSF support ensures lossless failover for multicast traffic in case of events like:

- PIM module(daemon) Restart/Redundancy Switchover
- VSF Switchover
- ISSU

Multicast routes programmed in Hardware are not removed during these events, and the PIM module will reconstruct its internal cache, and users will not see any disruption to existing streams.

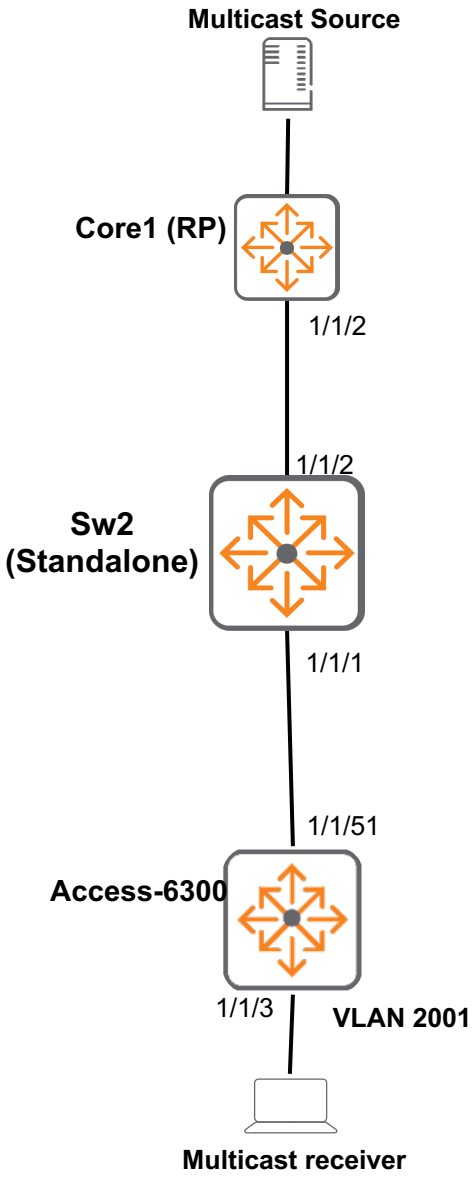

# Supported Platforms

- 6200, 6300, 6400, 8320, 8325, 8360, 8400 and 10000

# Use Cases

Deployment options and Solutions

# PIM Module (Daemon) Restart/Redundancy (Management Module)

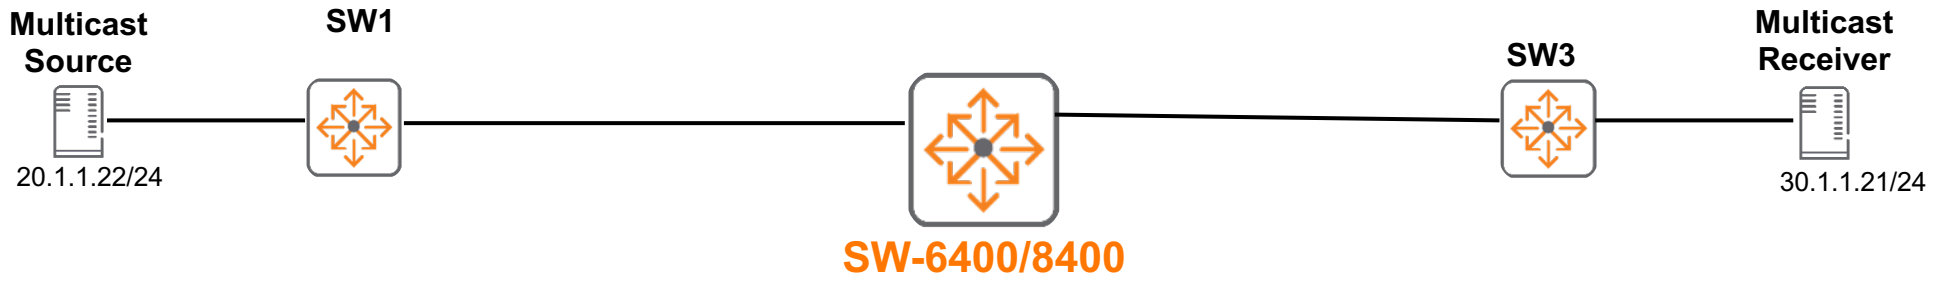


**Management Module (MM) Failover**

If the PIM daemon restarts in a standalone switch, it will NOT interrupt multicast traffic flow for the existing Mroute since multicast routes are programmed in hardware.

A switch with dual management modules multicast traffic continues to flow if active MM fails.

- Switch DB is synchronized and multicast routes are programmed in hardware.
- Mroutes are not removed during these events, and the PIM module on the newly active MM will reconstruct its internal cache, and users will not see any disruption to existing streams.

# PIM module Switchover/VSF

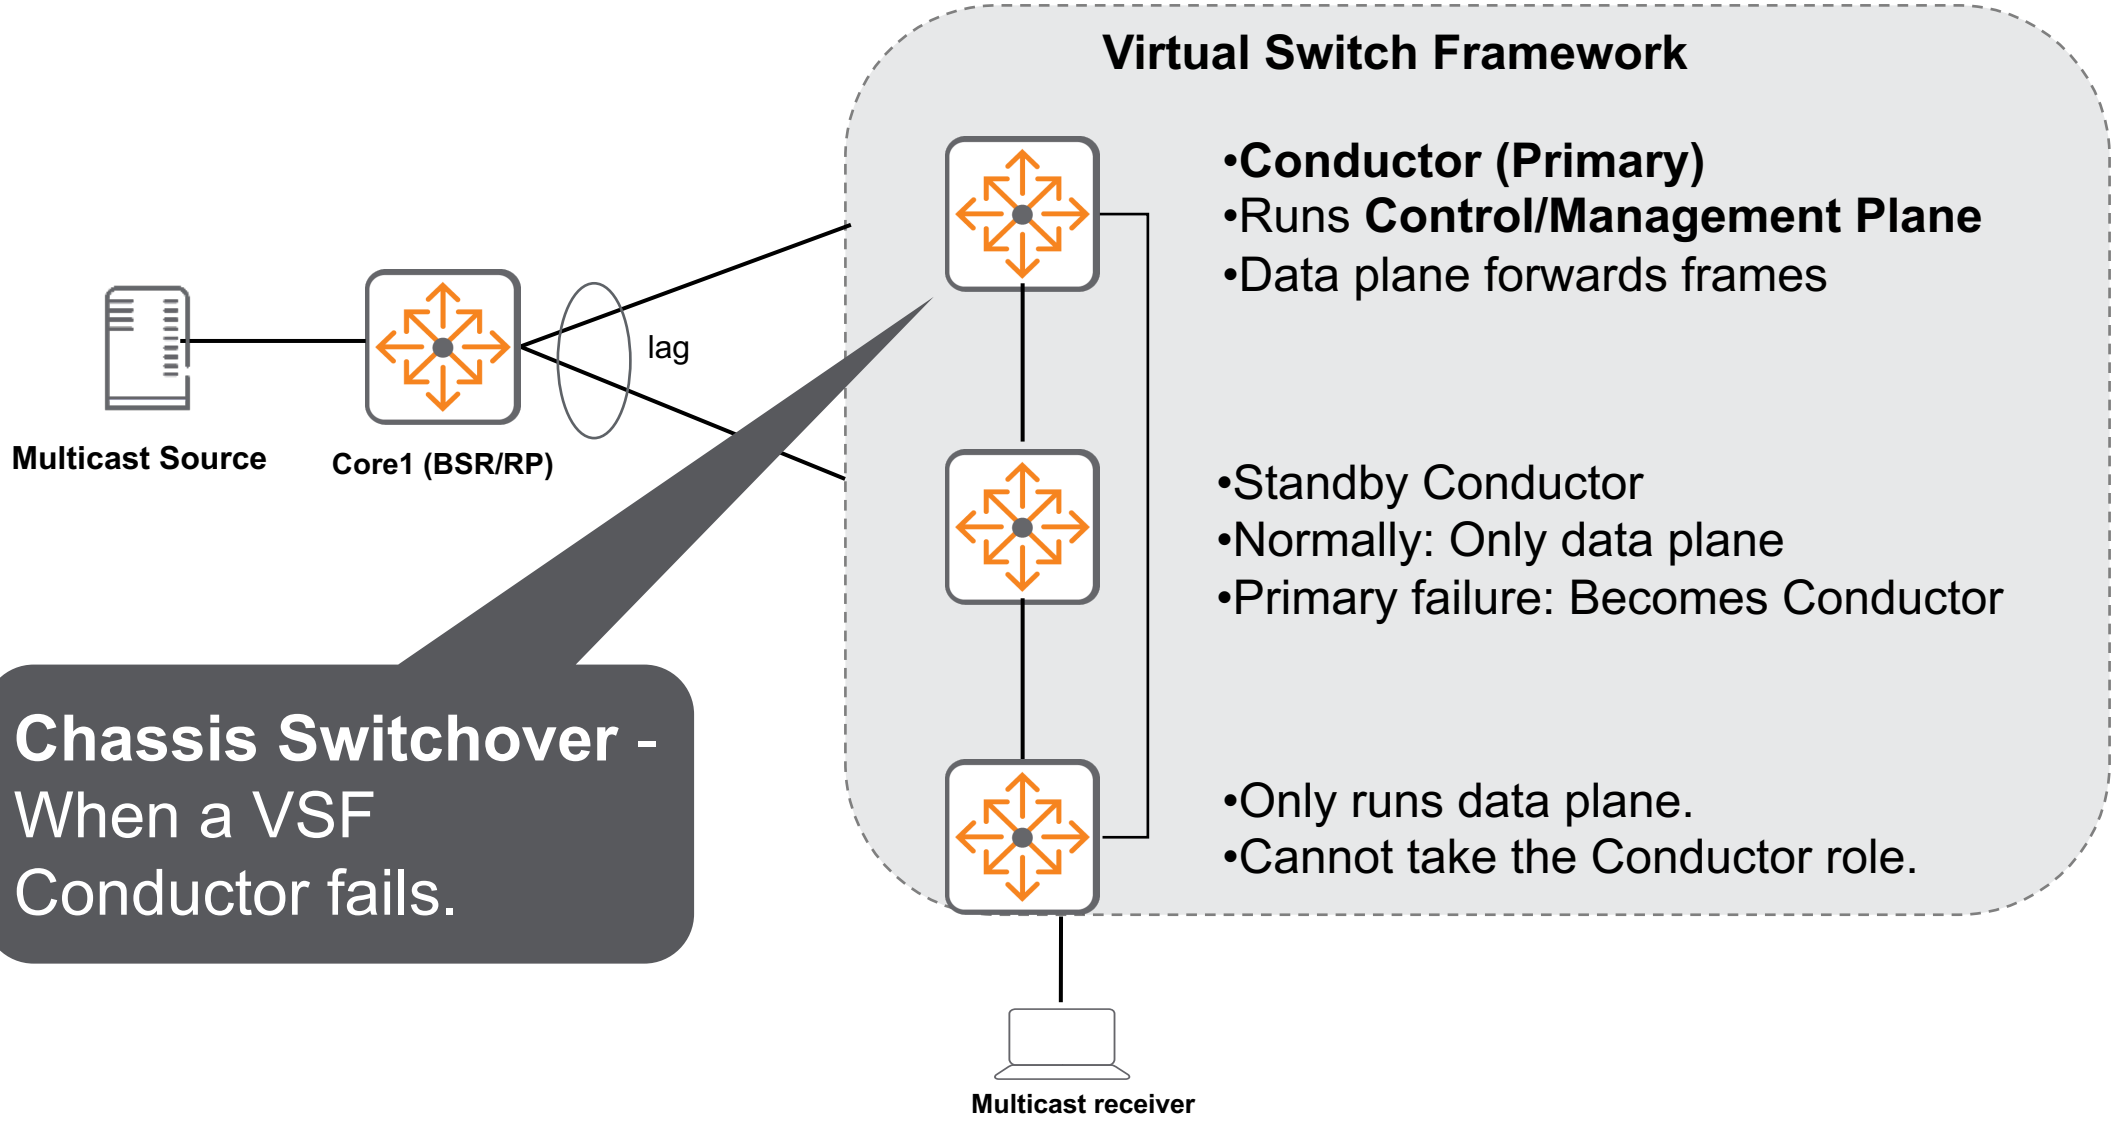

# PIM module Switchover/ISSU

ISSU - When a user triggers ISSU

From the PIM perspective, both ISSU and Switchover are similar to a daemon restart

PIM NSF is local to a device, and it does not exchange any information with VSX peers

Wherever ISSU is triggered, PIM NSF operations also get triggered

During NSF, PIM construct its internal cache from its neighbors

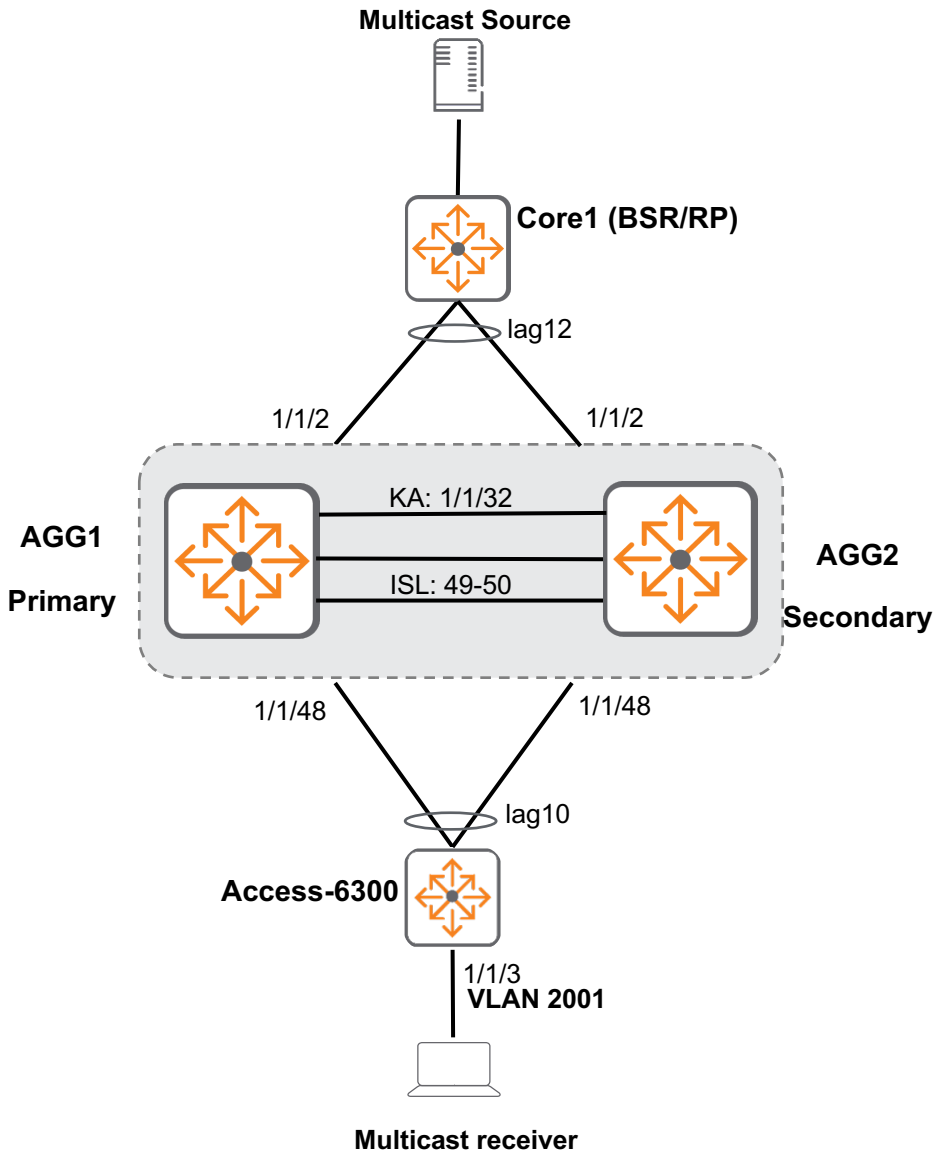

# Details and Caveats

# PIM NSF Mode

PIM NSF mode is activated by the events like PIM module restart/redundancy switchover/VSF switchover/ISSU. When PIM NSF mode is activated, PIM module triggers a restart timer and internally starts rebuilding the protocol states.

When the restart timer is running:

1. **Existing Mroutes are not impacted and traffic flow is uninterrupted.**

2. Any updates to existing Mroutes like an add or delete of outgoing interfaces are not updated to DB immediately. This will be updated once the NSF timer expires.

3. If any flow that is newly arrived in any interface during this time, will be updated to DB and installed in hardware immediately, and if there is any further update such as outgoing interface list update or deletion to the newly arrived stream would take place immediately.

4. Any inactive Mroute entries that get timed out will be removed from the Mroute table once the restart timer expires.

5. When the restart timer expires:

   a) **Mark operation** (where new entries present in the PIM protocol cache are inserted into DB) and

   b) **Sweep operation** (where the stale entries which are not present in the PIM protocol cache are deleted from DB) is performed to remove the stale entries from Mroute/next-hop tables in DB.

   c) Upon Mark and Sweep operation completion, normal processing of Mroutes is resumed.

# PIM NSF Mode

**NSF Inactive state**

```
SW2# show ip pim
PIM Global Parameters
VRF                                 :  default
PIM Status                          :  Enabled
PIM SSM Status                      :  Disabled
PIM SSM Range ACL                   :  Not Configured
Join/Prune Interval (sec)           :  60SPT
Threshold                           :  Enabled
State Refresh Interval (sec)        :  60
PIM NSF Status                      :  Inactive
```

**NSF Active state**

```
SW2# show ip pim
PIM Global Parameters
VRF                                 :  default
PIM Status                          :  Enabled
PIM SSM Status                      :  Disabled
PIM SSM Range ACL                   :  Not Configured
Join/Prune Interval (sec)           :  60SPT
Threshold                           :  Enabled
State Refresh Interval (sec)        :  60
PIM NSF Status                      :  Active
PIM NSF Time Remaining (HH:MM:SS)   :  00:03:25
```

PIM NSF is enabled by default

PIM NSF mode is activated by the events

The default restart timer value is **210 seconds**. When the restart timer is running, all the multicast routes which are present will not be changed. Currently, the timer is not configurable but in future code, it will.

# PIM NSF Support

With PIM SM, PIM SSM and PIM DM protocols.

Supports IPv4 and IPV6

# Caveats

- With PIM-SM, the solution only works on non-RP devices and non-source-connected DR routers.
- It is supported on other routers like last hop/intermediate routers and BSR routers.
- If there's a PIM module crash during Graceful shutdown process, multicast traffic outage is expected as PIM NSF is not triggered in this scenario.
- Updates to existing Neighbors/BSR/RP-Set information including timers are not published while timer is active.
- While PIM NSF mode is active during NSF timer (210 sec):
  - Modifications to existing mroutes, will not be allowed
  - If you get a leave message for an existing multicast group during that time, it will not be processed and traffic will keep flowing till the 210 seconds timer expires.
  - This is to ensure that any existing traffic flow is unaffected during the restart timer.

**NOTE:** In the case of PIM SSM and PIM DM, the solution works for all the devices.

# Configuration

# Commands

PIM NSF is enabled by default

There is no option to disable it

**Verification**

```
switch# show ip mroute
Switch# show ip pim
Switch# show events -d pimd
```

# Best Practices

# Best Practices

– NONE because it is enabled by default

# 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, ensure the PIM routing setup is correct and the outgoing interface list acquired mode, fix any issues found
- Ensure mroute is intact and forwarding when NSF is in active mode

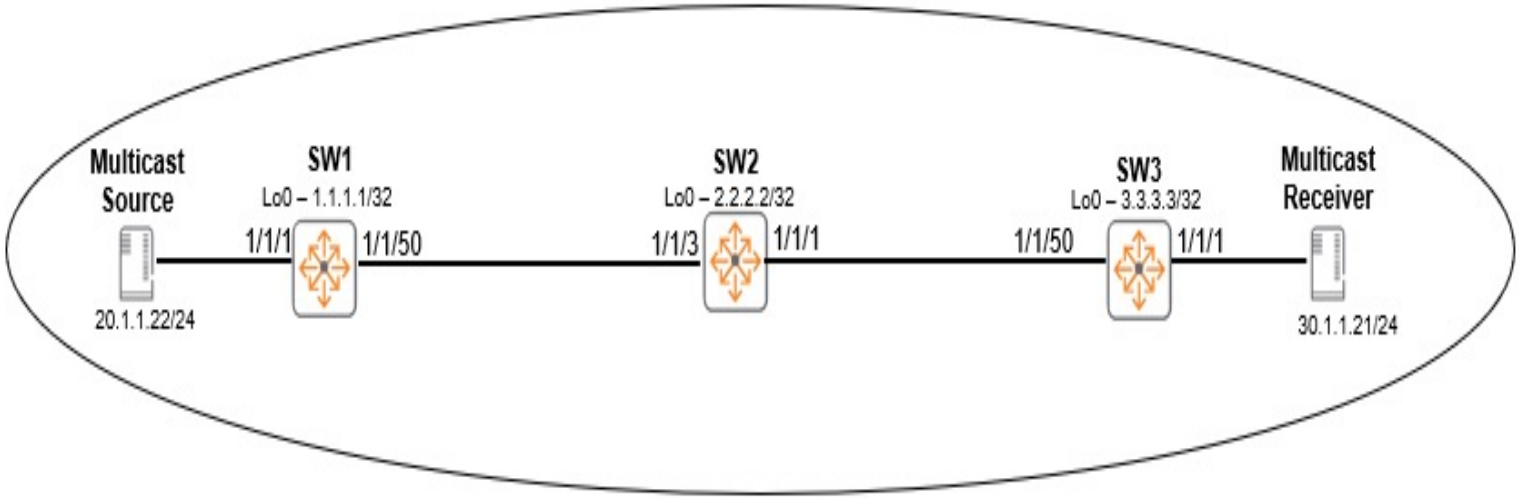


**Recommended troubleshooting flow**

1. Verify PIM neighbors up/operational and PIM routing tables are correct between the multicast source and receiver.
2. Verify data plane PIM NSF status (active or inactive), and time remaining counters (show ip pim)
3. Check to see if the mroute is still intact and forwarding, state “route” with uptime counters
4. Check events log for pimd NSF
5. You can use Shell command to test NSF

# 1.Verify PIM neighbors up/operational and PIM routing tables are correct between the multicast source and receiver.

```
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               : 192.168.100.0
Uptime                 : 1 days 06:47:13
State                  : route
Incoming interface     : 1/1/3
Outgoing Interface List :
Interface      State
----------     ----------
1/1/1          forwarding
```

# 2. Verify data plane PIM NSF status (active or inactive), and time remaining counters

**NSF Inactive state**

```
SW2# show ip pim
PIM Global Parameters
VRF                                 :  default
PIM Status                          :  Enabled
PIM SSM Status                      :  Disabled
PIM SSM Range ACL                   :  Not Configured
Join/Prune Interval (sec)           :  60SPT
Threshold                           :  Enabled
State Refresh Interval (sec)        :  60
PIM NSF Status                      :  Inactive
```

**NSF Active state**

```
SW2# show ip pim
PIM Global Parameters
VRF                                 :  default
PIM Status                          :  Enabled
PIM SSM Status                      :  Disabled
PIM SSM Range ACL                   :  Not Configured
Join/Prune Interval (sec)           :  60SPT
Threshold                           :  Enabled
State Refresh Interval (sec)        :  60
PIM NSF Status                      :  Active
PIM NSF Time Remaining (HH:MM:SS)   :  00:03:14
```

# 3. Check to see if the mroute is still intact and forwarding, state “route” with uptime counters

```
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               : 192.168.100.0
Uptime                 : 1 days 06:47:13
State                  : route
Incoming interface     : 1/1/3
Outgoing Interface List :
Interface      State
----------     ----------
1/1/1          forwarding
```

# Check Event Logs (4. Check events log for pimd NSF)

**Verification**

```
SW2# show events -d pimd | inc NSF

SW2 pimd[12694]: Event|5133|LOG_INFO|AMM|1/1|PIM NSF timer activated.
SW2 pimd[12694]: Event|5133|LOG_INFO|AMM|1/1|PIM NSF timer complete, resuming normal operation.
```

# Testing Restarted PIM Daemon (5. You can use Shell command to test NSF)

**NSF Inactive state**

```
SW2# show ip pim
PIM Global Parameters
VRF                                 :  default
PIM Status                          :  Enabled
PIM SSM Status                      :  Disabled
PIM SSM Range ACL                   :  Not Configured
Join/Prune Interval (sec)           :  60SPT
Threshold                           :  Enabled
State Refresh Interval (sec)        :  60
PIM NSF Status                      :  Inactive
```

```
SW2# start-shell
SW2:~$ sudo bash
SW2:/home/admin# systemctl restart pim
SW2:/home/admin# exit
SW2:~$ exit
```

**NSF Active state**

```
SW2# show ip pim
PIM Global Parameters
VRF                                 :  default
PIM Status                          :  Enabled
PIM SSM Status                      :  Disabled
PIM SSM Range ACL                   :  Not Configured
Join/Prune Interval (sec)           :  60SPT
Threshold                           :  Enabled
State Refresh Interval (sec)        :  60
PIM NSF Status                      :  Active
PIM NSF Time Remaining (HH:MM:SS)   :  00:03:14
```

# Demo

# Multicast NSF Demo

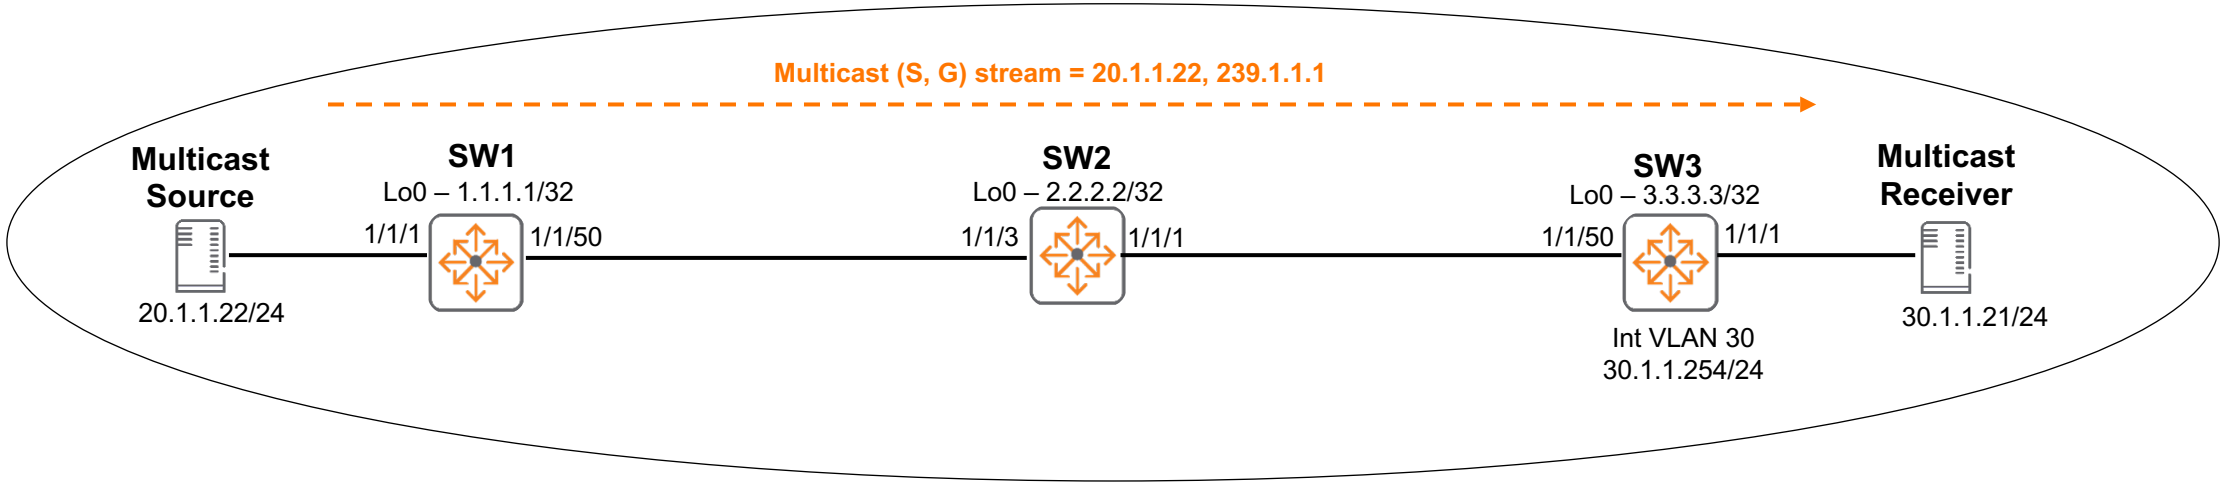


- Show multicast traffic forwarding is impacted by control plane restart on SW3 (10.9)
- Show multicast traffic forwarding is not impacted by control plane restart on SW2 (10.10) due to PIM NSF (no additional configuration required)
- PIM NSF applicable to both IPv4/IPv6 multicast traffic

aruba
a Hewlett Packard Enterprise company

# Thank you

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