Let’s go

The bridge to possible

Demystifying IP multicast in SD-Access

Marcin Hamroz , Principal Architect
Jarosław Gawron, Principal Engineer

What is this session (not) about?

- This session is not purely about **IP multicast** ✗
- This session is not purely about **Cisco SD-Access** ✗
- This session is about **IP multicast** in **Cisco SD-Access** ✓

  Head-end replication ✓

  Native multicast ✓

  Layer 2 flooding ✓

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

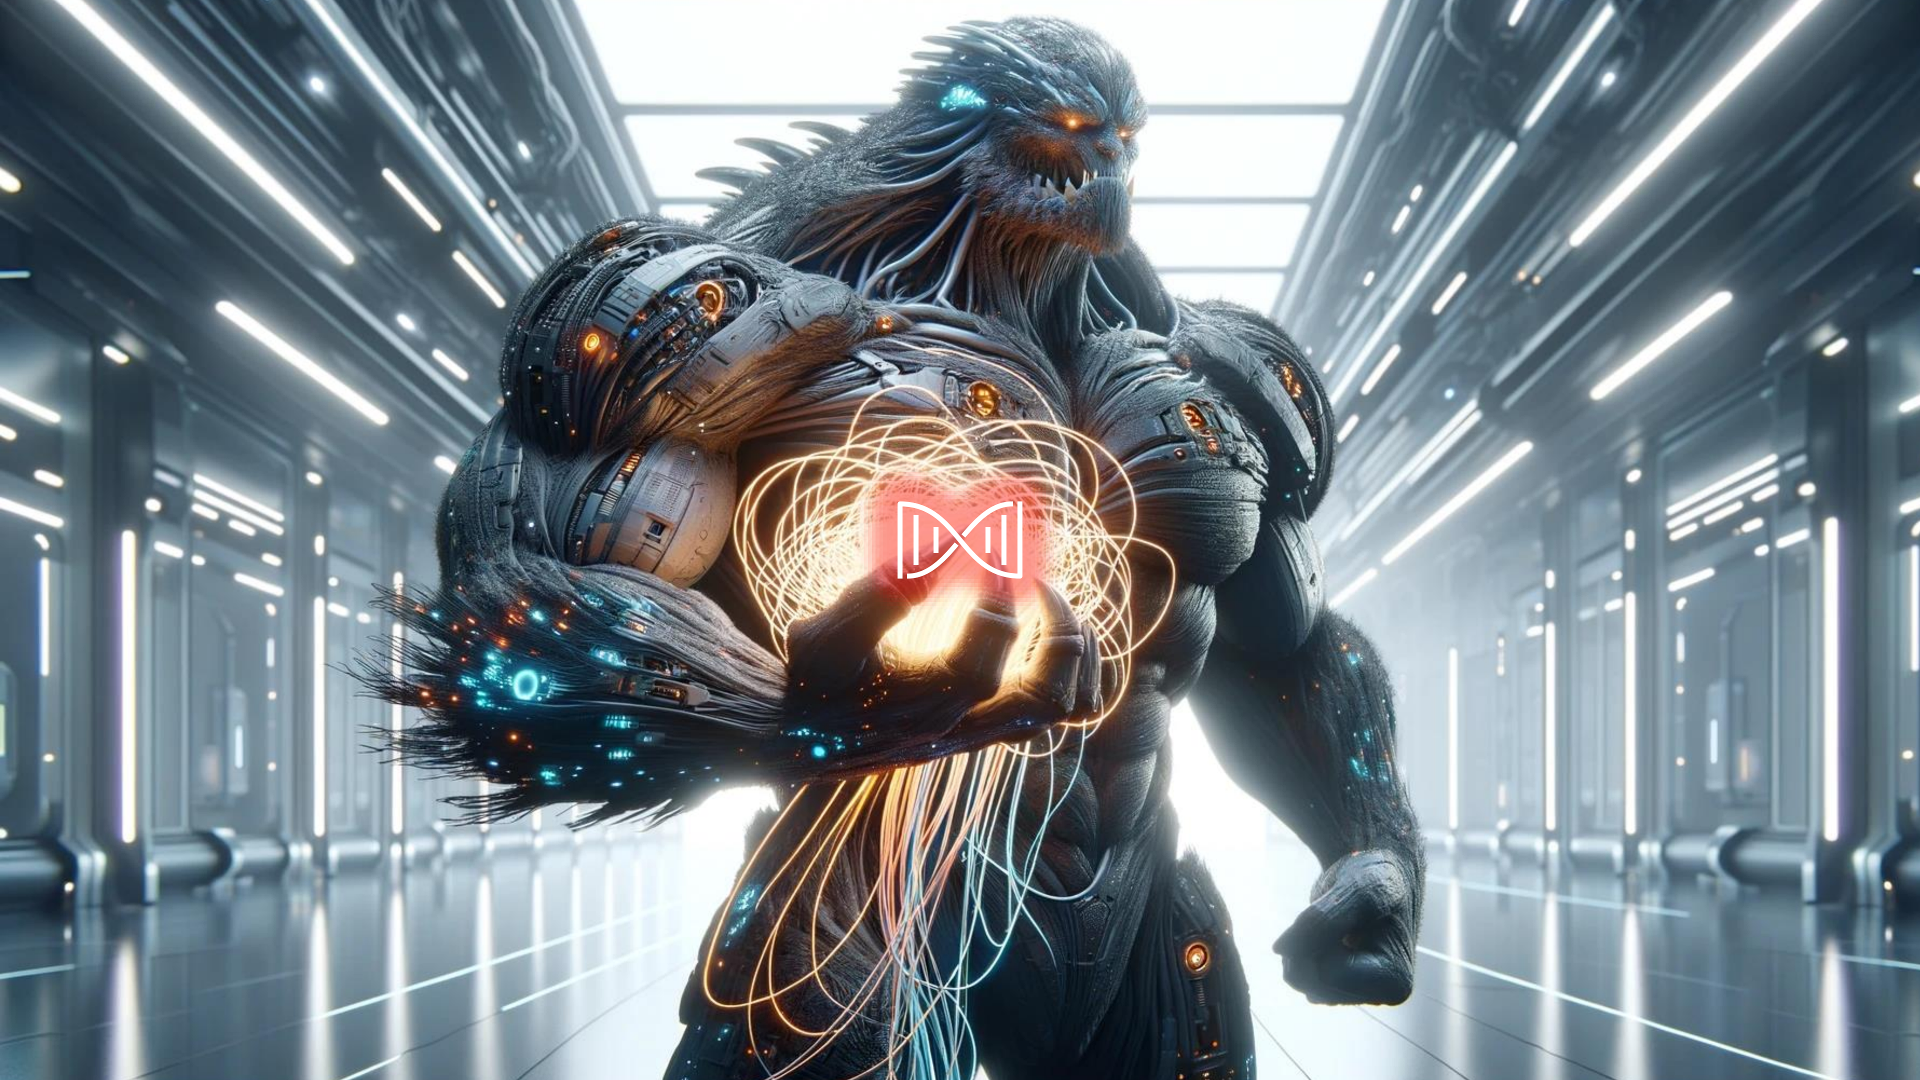

Marcin Hamróz

Principal Architect

- Part of Professional Services team
- Coined Cisco in 2012
- Based out of Cisco Krakow
- Focused on Software Defined Access & SD-WAN
- CCIE R&S / SP
- Father of three
- Fan of aviation

Jarosław (Jaro) Gawron

Principal Engineer

- In TAC from 2012
- Based out of Cisco Krakow
- Focused on Software Defined Access & Catalyst Platforms
- CCIE R&S / SP
- Father of three
- Fan of StarTrek and sailing

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- Introduction
- IP multicast recap
- Head-end replication vs Native multicast
- SD-Access multicast–deployment models
- Layer 2 flooding
- Multicast in SDA Transit
- Design considerations
- Summary

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

IP multicast recap

ASM, SSM, RP, MSDP

Multicast Routing Terminology

Source – device originating multicast traffic

Receiver – device requesting multicast traffic

First-Hop Router (FHR) – device attached to source network segment

Last-Hop Router (LHR) – device attached to receiver network segment

Multicast Device – device enabled for multicast traffic

Randezvous Point (RP) – a root for a shared tree

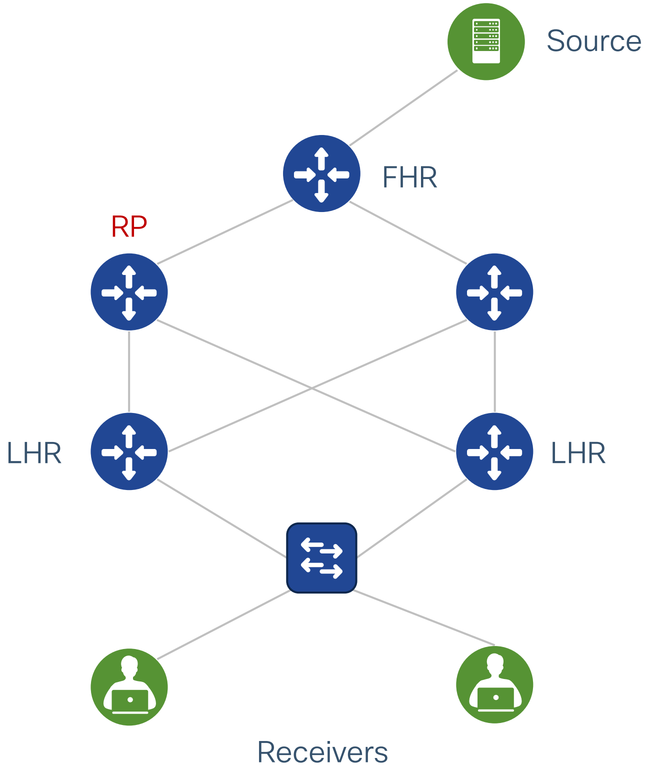

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Any-Source Multicast

- Used in scenarios where receivers do not know the sources sending to a multicast group.
- ASM is the only option in IGMP version 1 and 2. It is also supported in IGMP version 3.
- Multicast devices must learn which sources are sending to multicast group in order to forward packets to receivers.
- In ASM we need a Rendezvous Point(s)!

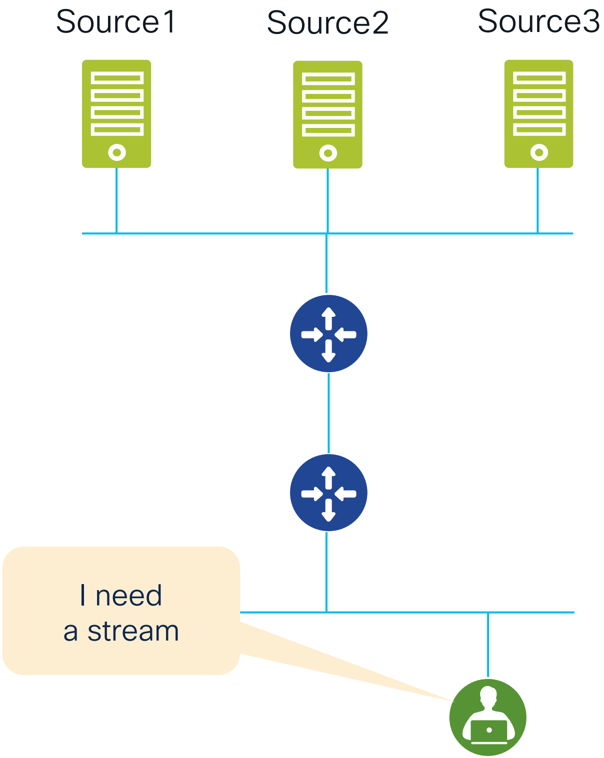

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

ASM basic workflow

- Source starts sending traffic to a multicast group. It reaches FHR first.
- FHR sends a PIM Register unicast packet encapsulated in PIM Tunnel to RP.
- At this point the multicast traffic is being sent in unicast tunneling to RP.
- What happens next depends if receivers requested a multicast stream. If there are no receivers yet, RP sends PIM Register Stop message up to FHR and waits.

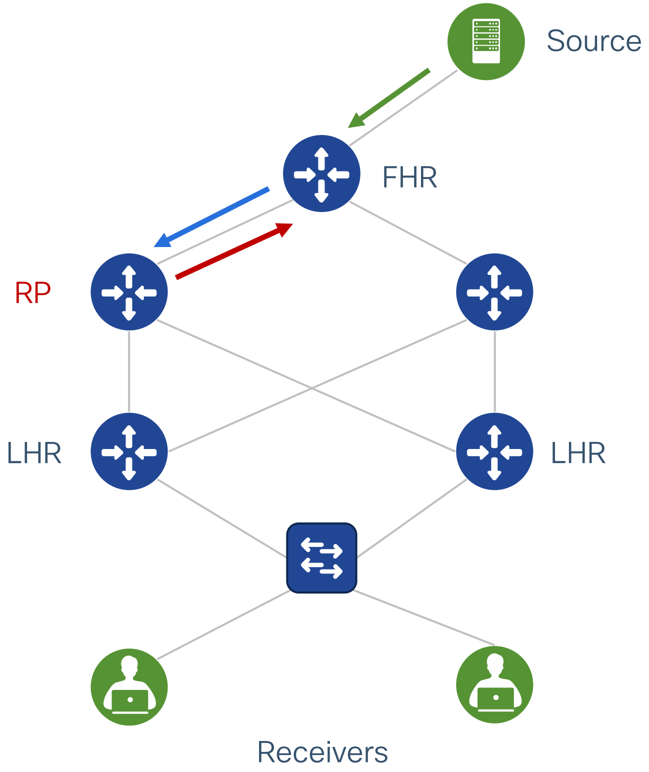

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

ASM basic workflow

- Receivers request multicast stream by sending IGMP Join messages to the segment
- Designated multicast router for this segment (DR) sends PIM Join (*,G) to RP.
- RP sends a PIM Join message to FHR to request a stream. FHR adds to OIL interface facing RP and forwards traffic.
- Shared Tree (RPT) is now ready, so multicast stream can be forwarded down to receivers.

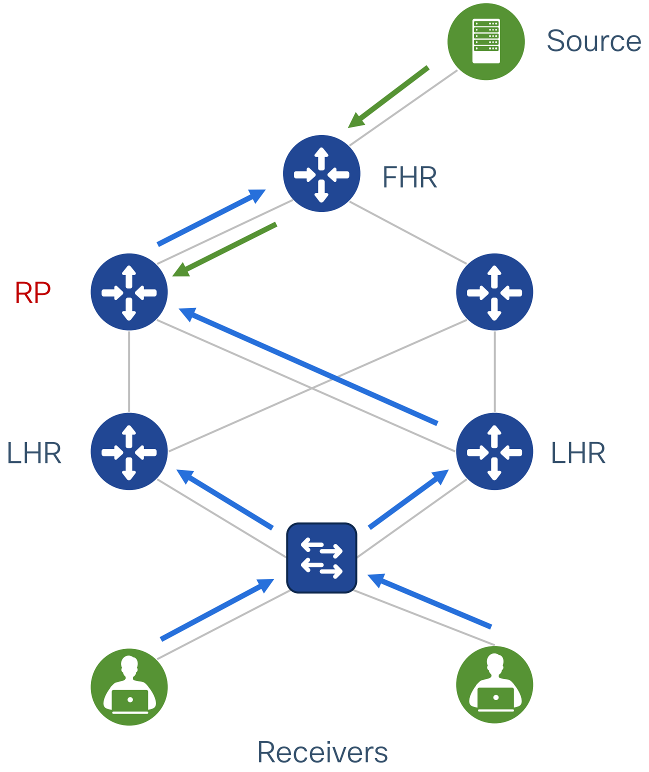

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

ASM basic workflow

Multicast traffic forwarded through RPT

- Multicast traffic flows down through RPT following OIL on all multicast devices.
- Once multicast stream hits LHR, it learns about multicast source!
- While traffic flows to receivers, LHR now starts building a separate PIM Join (S,G) directly to the source.

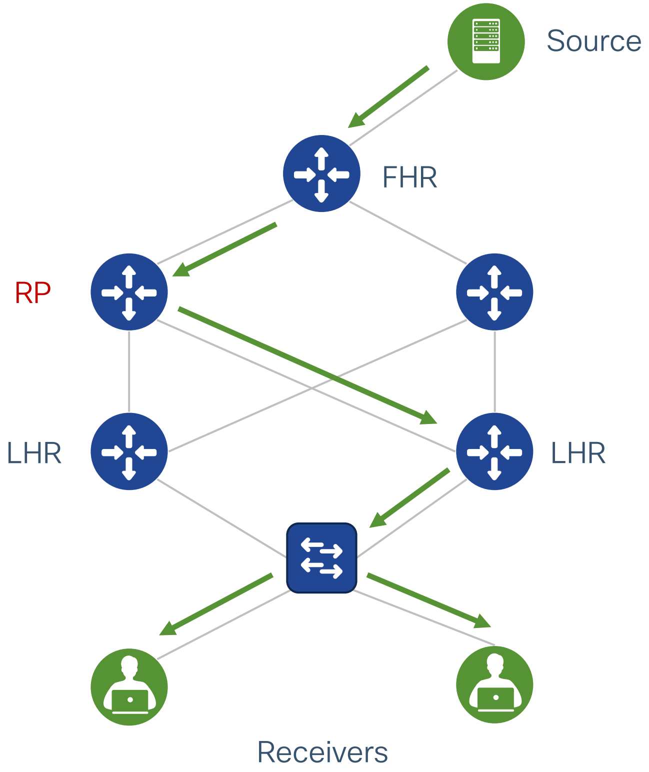

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

ASM basic workflow

Building Shortest Path Tree (SPT)

- LHR sends new PIM Join towards multicast source.
- FHR adds interface towards LHR to OIL and traffic starts flowing down OIL to LHR.
- LHR now has two multicast streams…

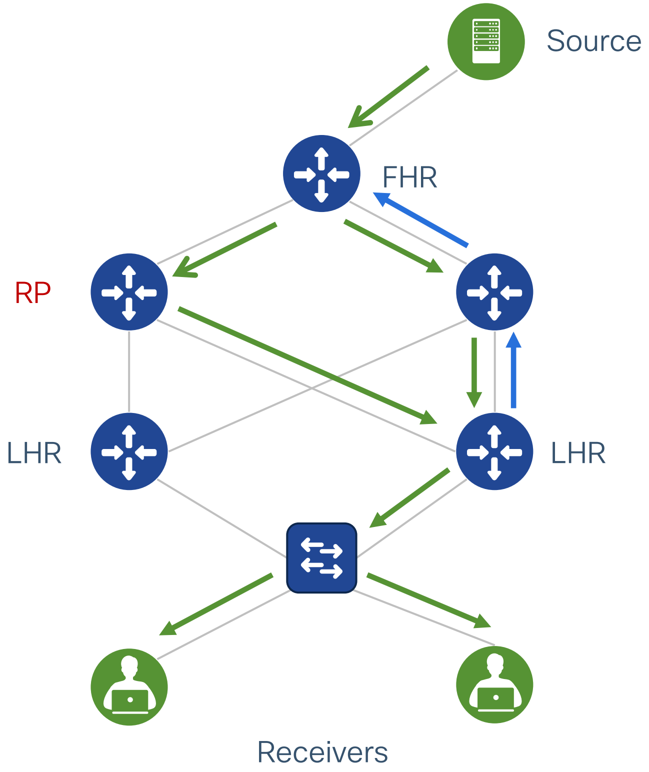

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

ASM basic workflow

Switching to Shortest Path Tree (SPT) !

- LHR sends a PIM Prune message to the RP for the (*,G) entry.
- RP removes the interface facing LHR from OIL and stops delivering traffic.
- If there are no other OIL built for that (S,G) then the RP will prune itself.
- We’ve got only SPT left.

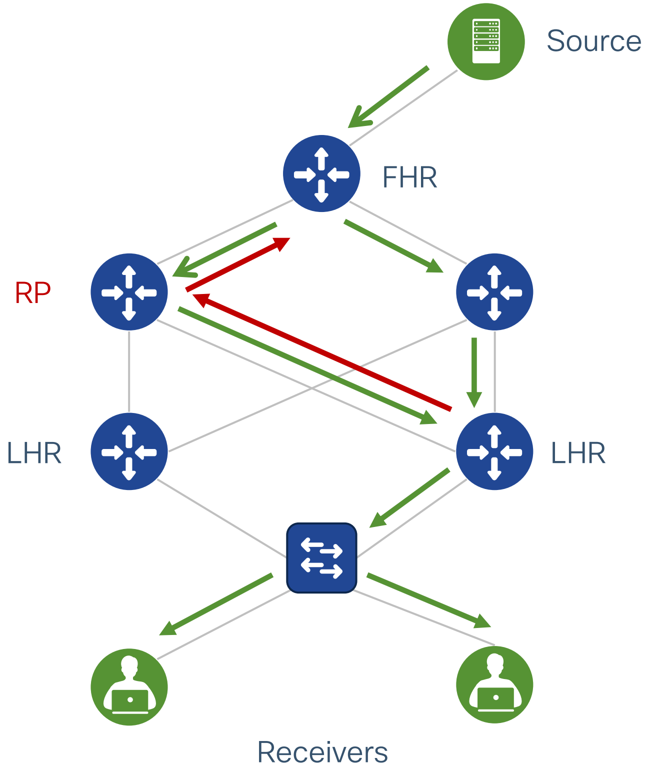

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

A closer look on the RP

Purpose:

- Helps to build SPT between a Source and Receivers.

Problems:

- How do all multicast devices agree on which one is the RP?
- If the RP fails in ASM, multicast traffic will fail unless already on SPT. How can we provide redundancy?

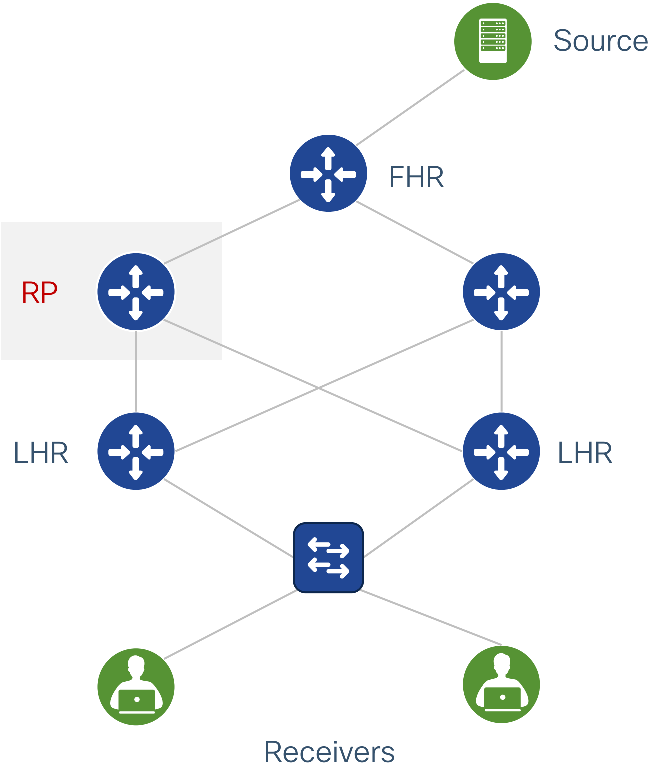

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

A closer look on the RP

Three ways to solve both problems

AutoRP (kind of old way)

- uses concept of Mapping Agent and Candidate RPs
- two dedicated multicast group used (224.0.1.39, 224.0.1.40)

BSR (better way)

- uses concept of Candidate BSR and Candidate RP
- uses All PIM Routers multicast group (224.0.0.13)

Anycast RP (smart approach)

- advertise same RP IP address from multiple devices
- all multicast routers knows RP via any method (Static, BSR, AutoRP)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

ASM with Anycast RP

How it really works?

- Source starts sending multicast. FHR sends a PIM Register unicast message encapsulated in PIM tunnel to one of RPs.
- Receivers request multicast stream by sending IGMP Join Message to the segment.
- DR sends PIM Join (*,G) to RP based on routing table / load-balancing algorithm.
- What if RP node which received PIM Join (*,G) doesn't have a knowledge about the source !?

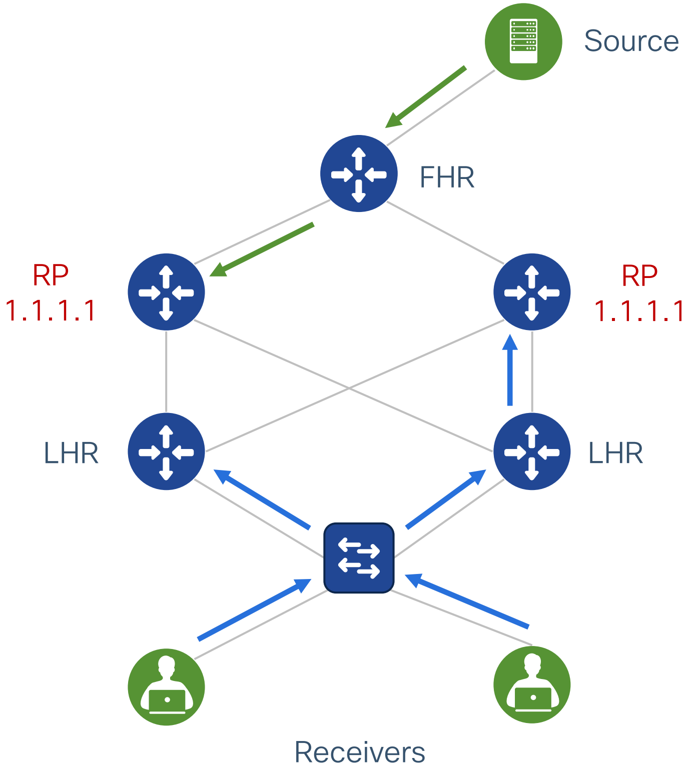

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

ASM with Anycast RP

Introducing Multicast Source Discovery Protocol (MSDP)

- Uses unique interfaces to exchange messages between Anycast RPs.
- When any RP receives PIM Register, sends MSDP SA message to the peer.
- MSDP Source Active message contains the IP of source and group address, if another RP has active PIM Joins and OIL for this group, it triggers that RP to build PIM Join to source.

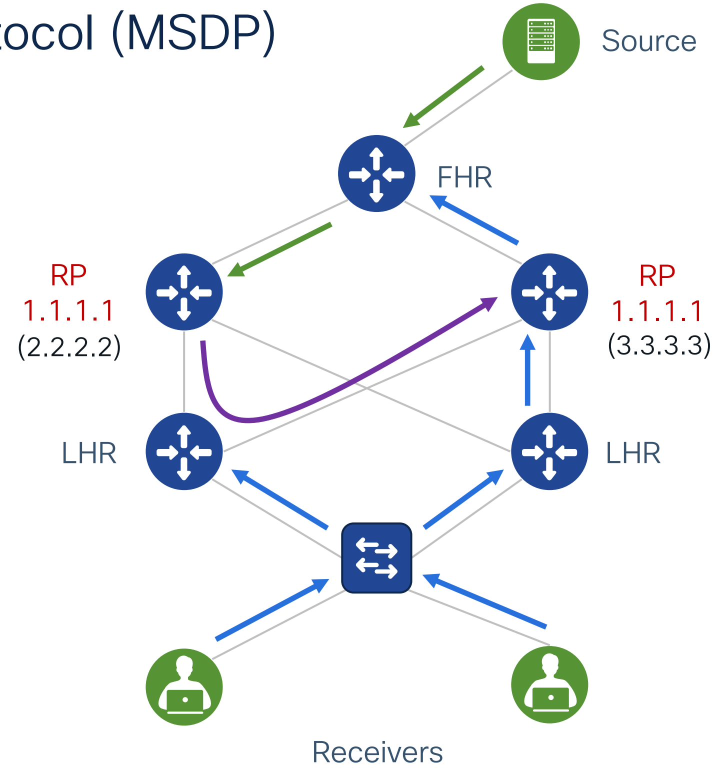

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Source-Specific Multicast

- A solution that allows receivers to point the source from which they want to receive traffic.
- SSM requires receivers to know sources and uses IGMPv3 Join message to request source and group pair.
- Uses reserved multicast group address range 232.0.0.0/8. *
- In SSM we don't need a Rendezvous Point(s)!

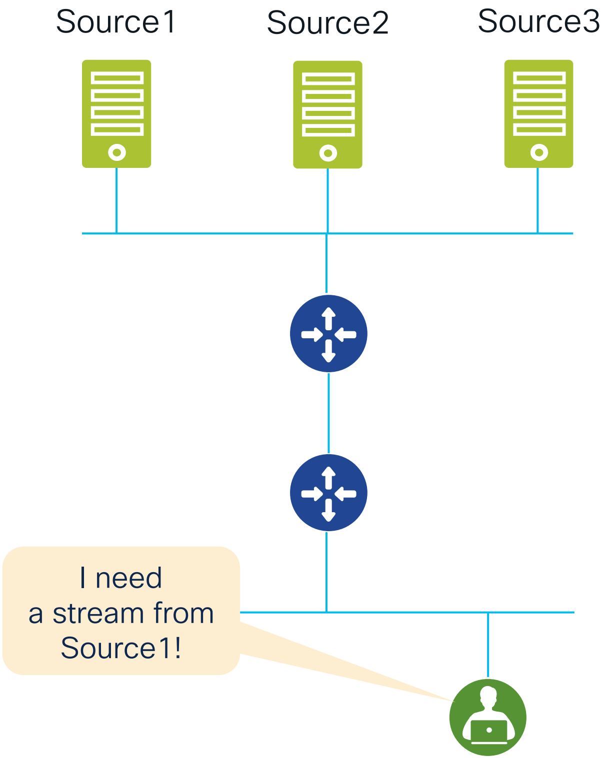

* Can be customized in a configuration

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

SSM basic workflow

Source starts sending multicast stream

- When a source sends multicast traffic to a group it hits FHR first and nothing happens.
- Receivers signal interest in multicast group by sending IGMPv3 Join (S,G) message.
- DR starts building SPT by sending PIM Join (S,G) towards the FHR.
- Multicast traffic flows down through SPT following OIL.

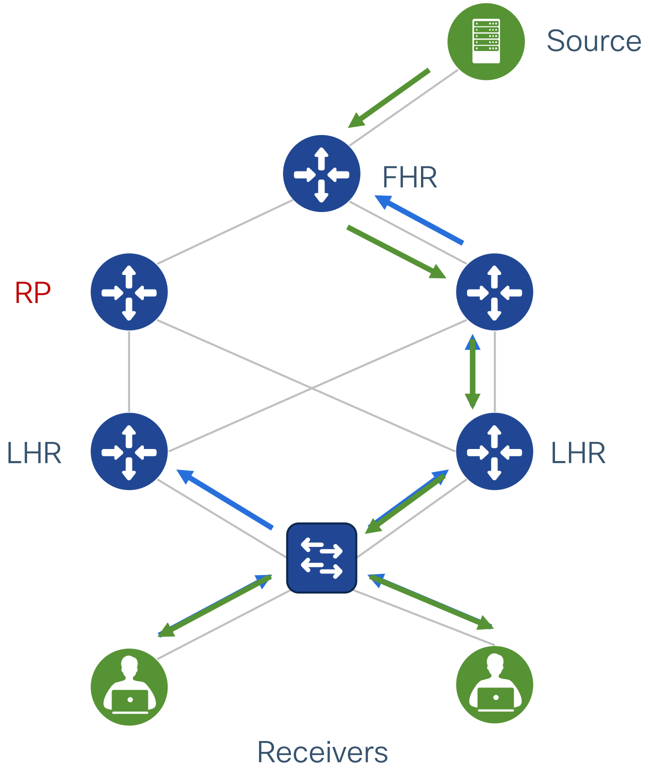

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication vs Native multicast

SD-Access multicast

Supported deployment modes (in a nutshell)

Head-end replication

- Forwarding in overlay
- Multicast over unicast scenario
- Supported from the beginning
- 1000 multicast groups* supported in overlay

Native multicast

- Forwarding in underlay
- Multicast over multicast scenario
- Supported from Cisco DNA Center 1.2.5 and IOS-XE 16.9.1s
- 1000 multicast groups* supported in underlay (mapping starts with **232.0.0.1** and ends with **232.0.3.232**)

(*) platform dependency

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication

Supported modes (overlay):	ASM, SSM
RP placement (ASM, overlay):	Inside or outside the fabric
Multicast source placement:	Inside or outside the fabric
RP redundancy (ASM, overlay):	MSDP
Configuration:	Fully automated by Cisco DNA Center

General rule:

Multicast packets are encapsulated in VXLAN and forwarded as unicast towards each Edge Node separately.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast

Supported modes (overlay):	ASM, SSM
Supported modes (underlay):	SSM
RP placement (ASM, overlay):	Inside or outside the fabric
Multicast source placement:	Inside or outside the fabric
RP redundancy (ASM, overlay):	MSDP
Configuration:	Fully automated by Cisco DNA Center only when LAN Automation feature used to configure underlay

General rule:

Multicast packets are encapsulated in VXLAN as multicast and forwarded to Edge Nodes through underlay multicast tree.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Head-end replication with ASM (config)

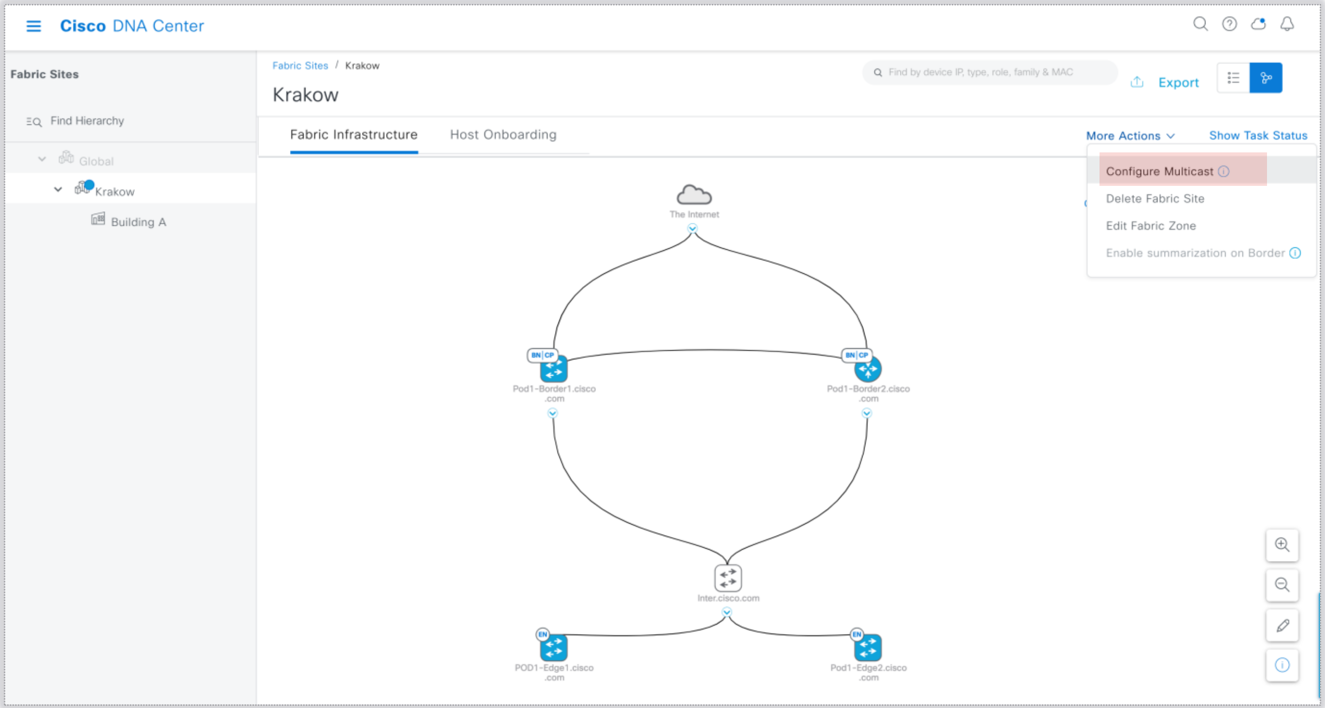

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM (config)

Selection of replication mode

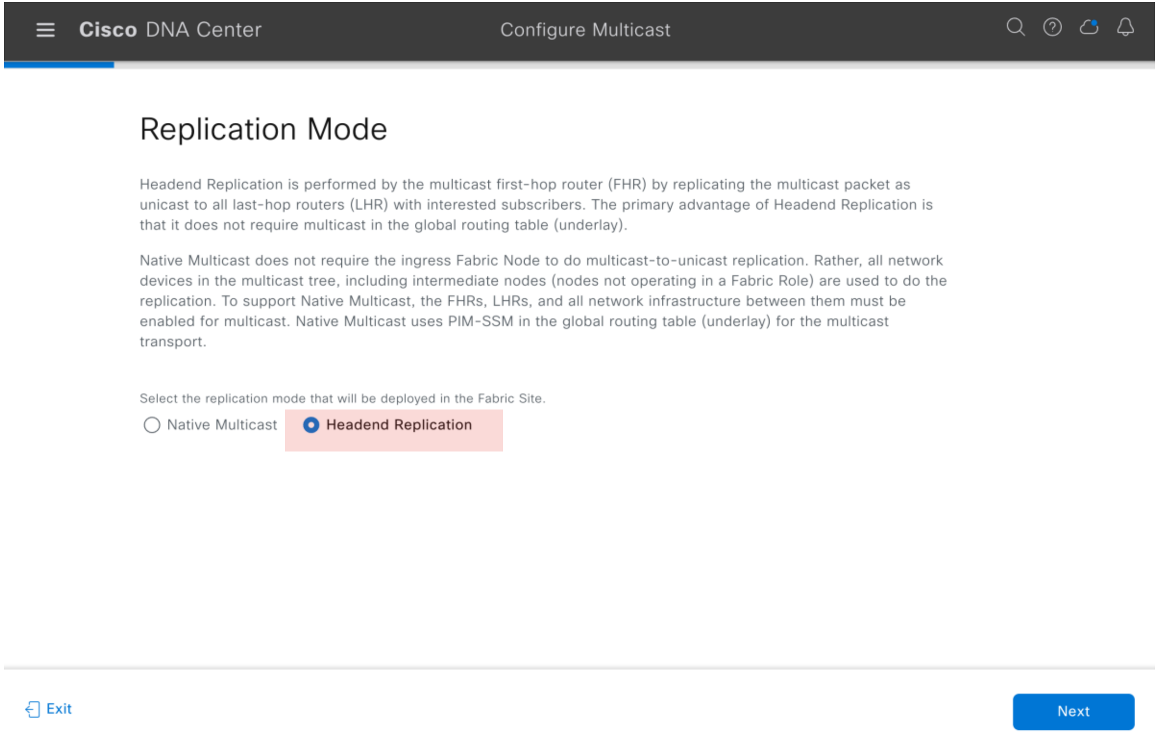

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM (config)

Selection of Virtual Network

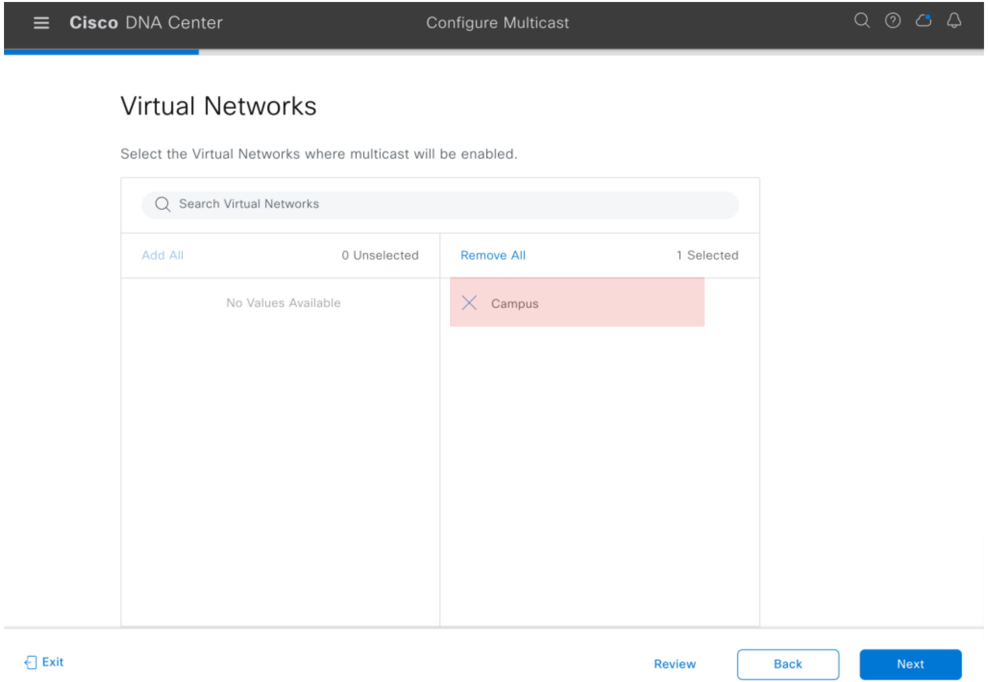

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM (config)

Multicast pool mapping

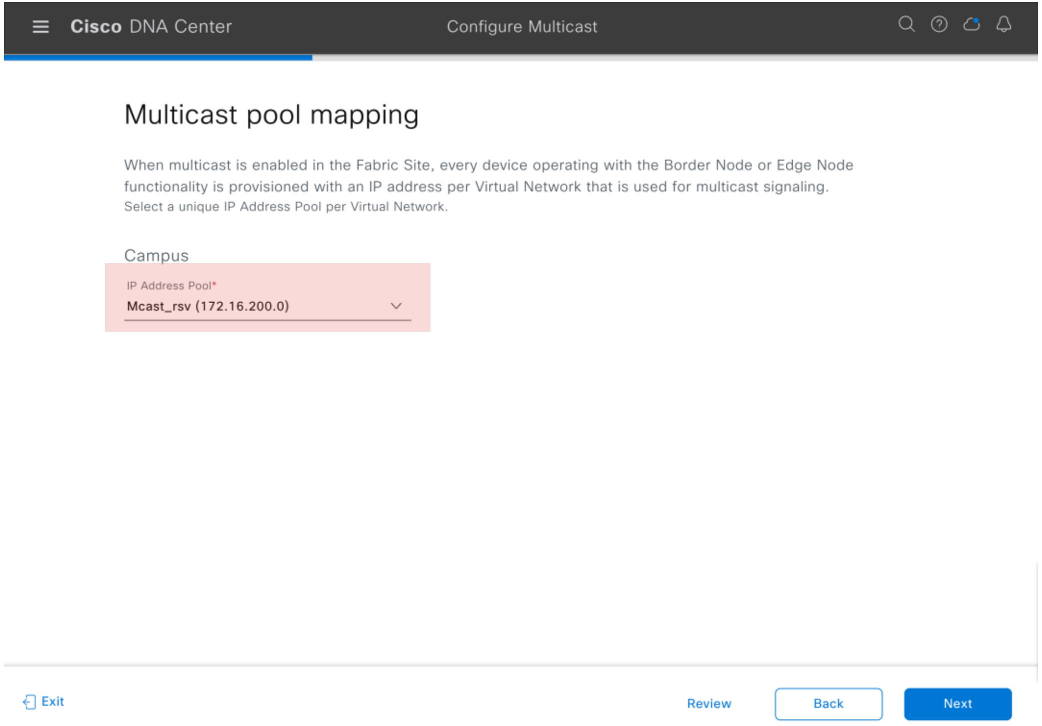

TIP

Configure one big IP pool in a global level, reserve pools per VN here

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

How big should be multicast IP pool?

- Every non-RP fabric device needs just one IP address.
- RP fabric devices need one common IP address (AnycastRP) and one IP address per device (MSDP), so three IP addresses in total.

You cannot use the same IP pool for multiple VNs. Each VN requires dedicated IP Pool!

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Selection of multicast mode

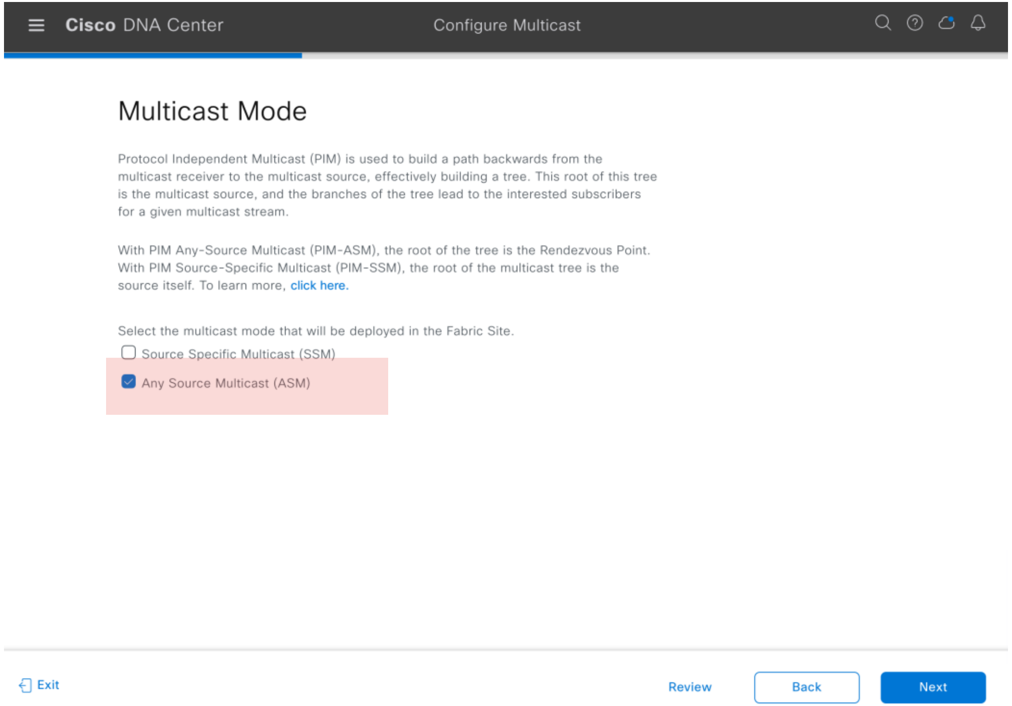

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (config)

RP Mapping

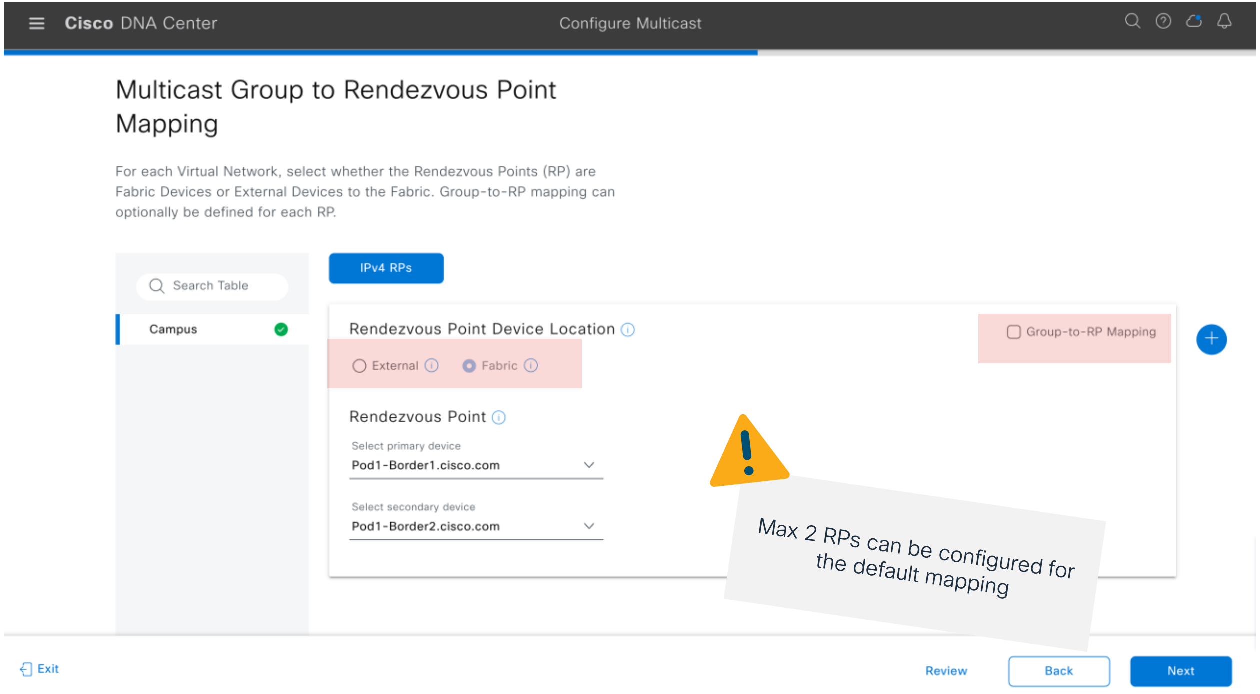

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM (configuration)

Final state

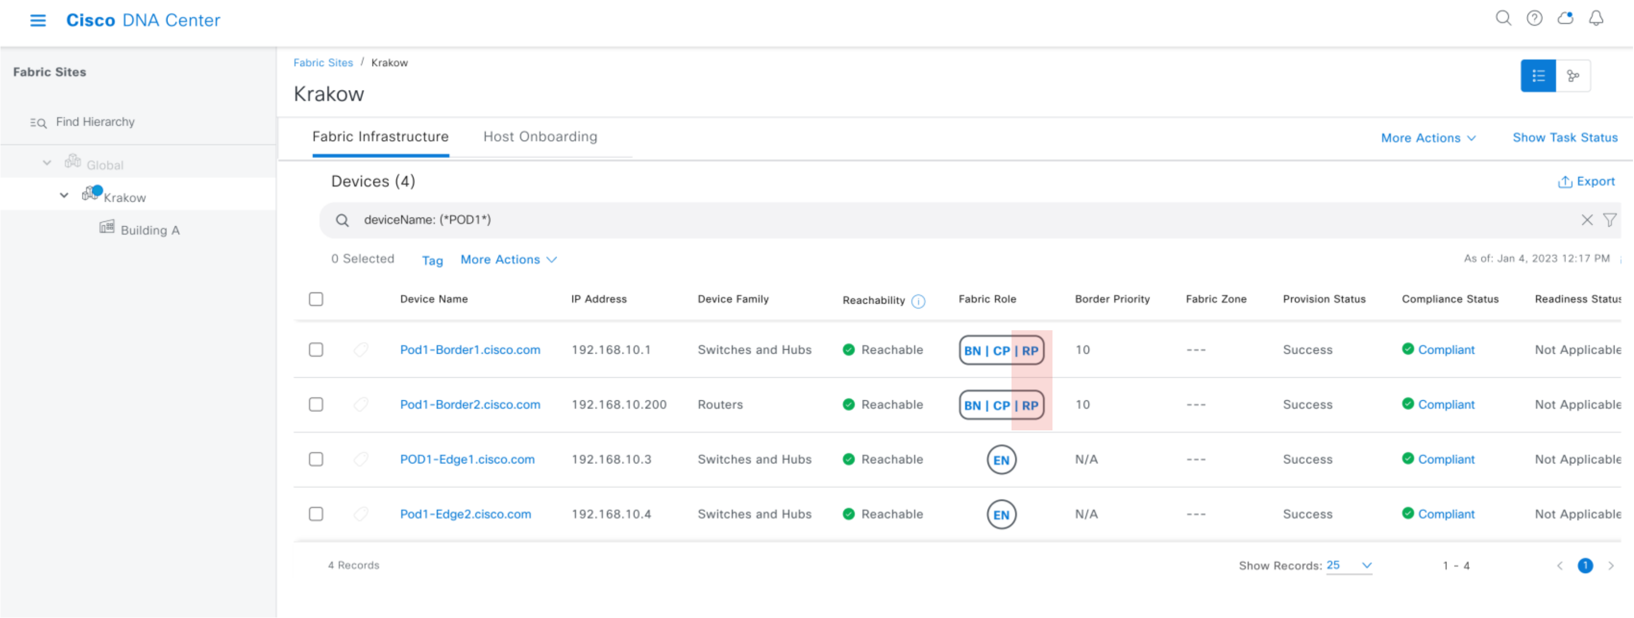

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

A closer look. Initial status.

INFO
RP advertisement in BGP automated by DNA-C!

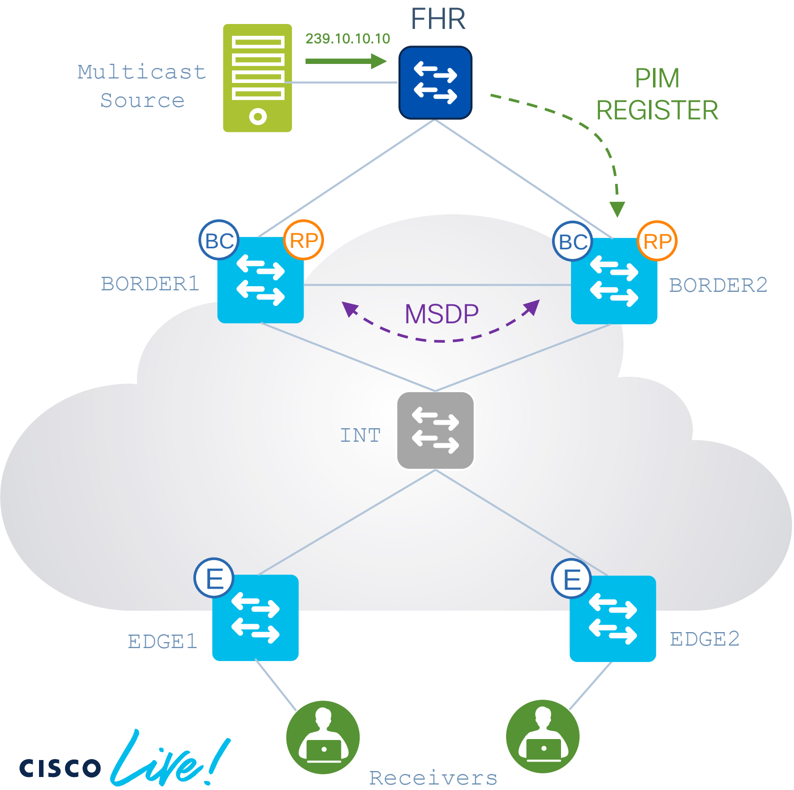

1. Receivers are connected to Edge Nodes (LHRs). Multicast Source is located outside the fabric.
2. Rendezvous Points (RPs) are present in overlay as a part of endpoint IP space.
3. PIM-SM is enabled in overlay between fabric devices and outside the fabric within VN/VRF.
4. Source starts to stream multicast traffic. FHR receives multicast stream and triggers any-source PIM Register towards RP (based on routing protocol).
5. RP updates a peer RP about the source via MSDP.
6. If there are no receivers yet, RP sends PIM Register Stop message up to FHR and waits.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

A closer look. Clients request multicast stream.

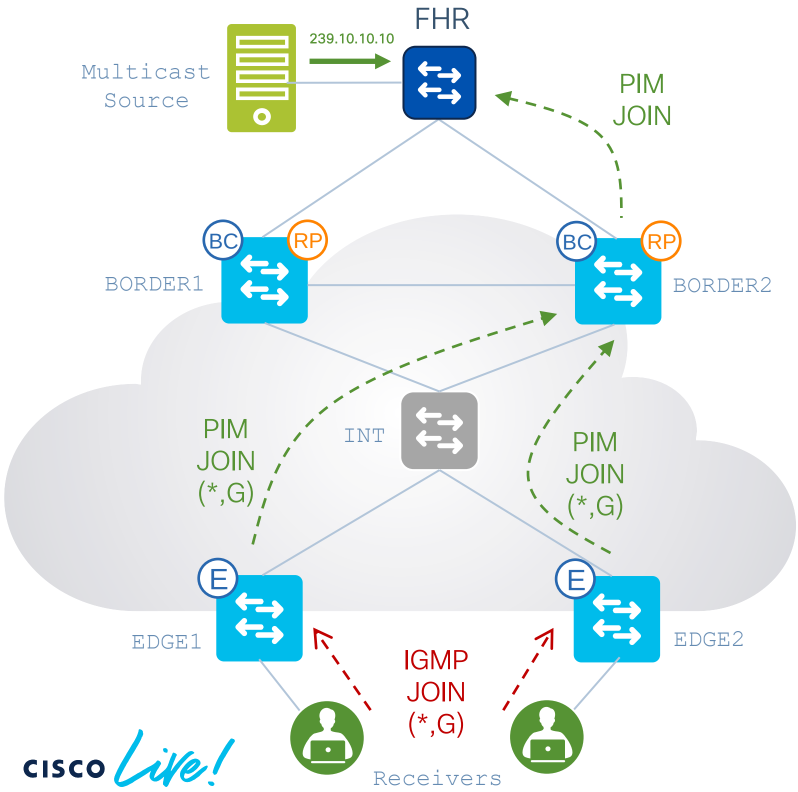

1. Clients send IGMP Join for any-source multicast group (*,G).

2. The Edge Node receives IGMP Join on SVI and triggers new PIM Join towards RP. Since RPs have same IP address (anycast), the hashing algorithm will be used to forward the packet (LISP next-hop towards RP).

3. The fabric RP now has both (S,G) knowledge for the multicast group and triggers PIM Join to FHR.

4. FHR adds interface facing RP to OIL and starts forwarding multicast traffic.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

A closer look. Multicast stream forwarded to receivers.

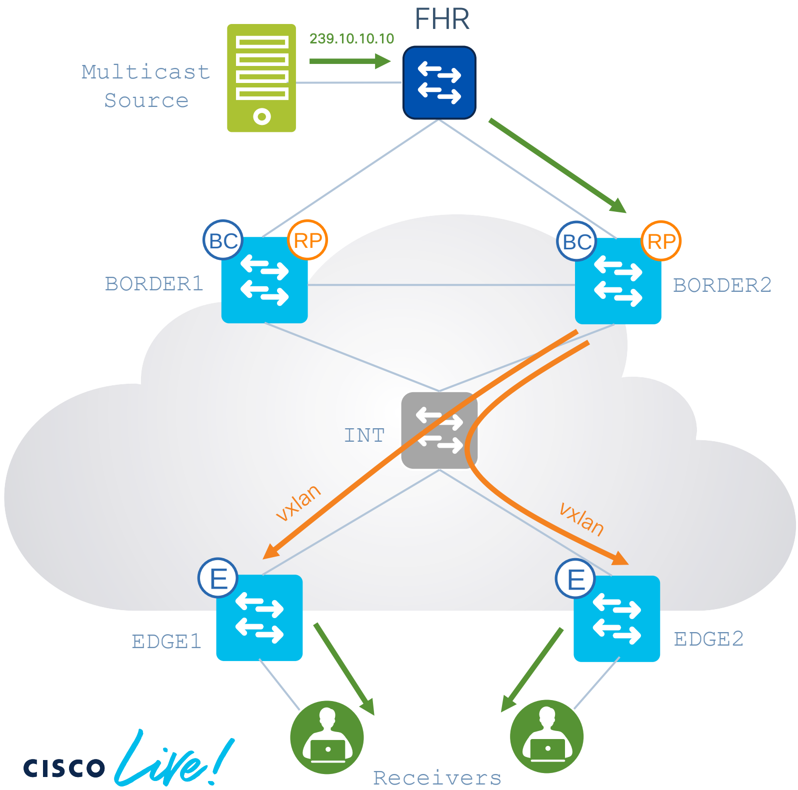

1. The Border Node (RP) receives a multicast packet from FHR.
2. The Border Node creates a copy of the original multicast packet for each LHR (Edge Node), encapsulate it in VXLAN and then unicasts it.
3. Each LHR receives VXLAN packet, decapsulates it and sends original multicast packet towards the client.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Let's get some captures!

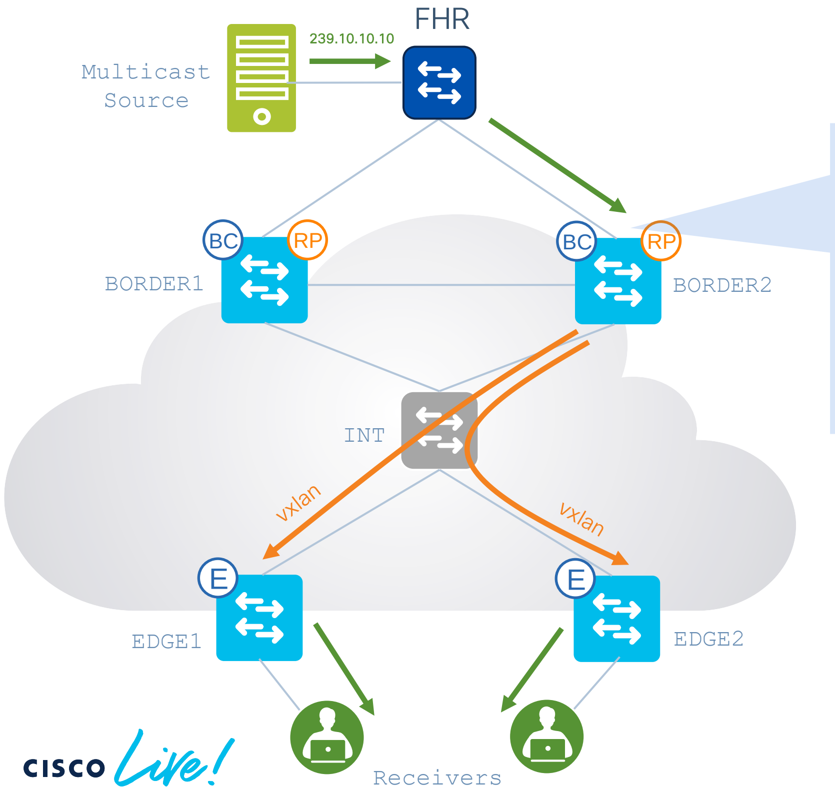

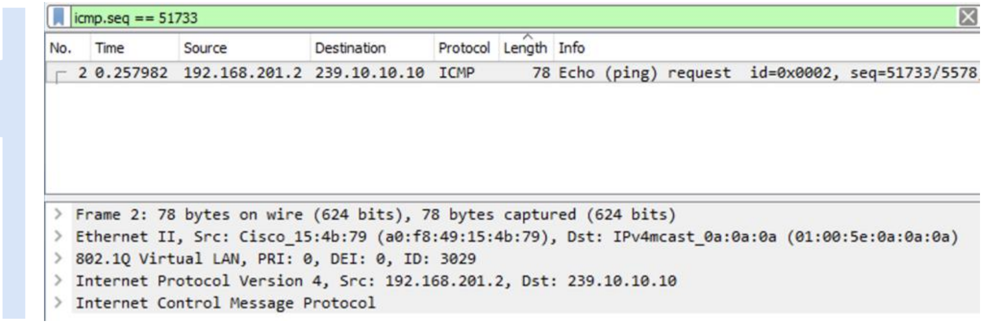

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Let’s get some captures!

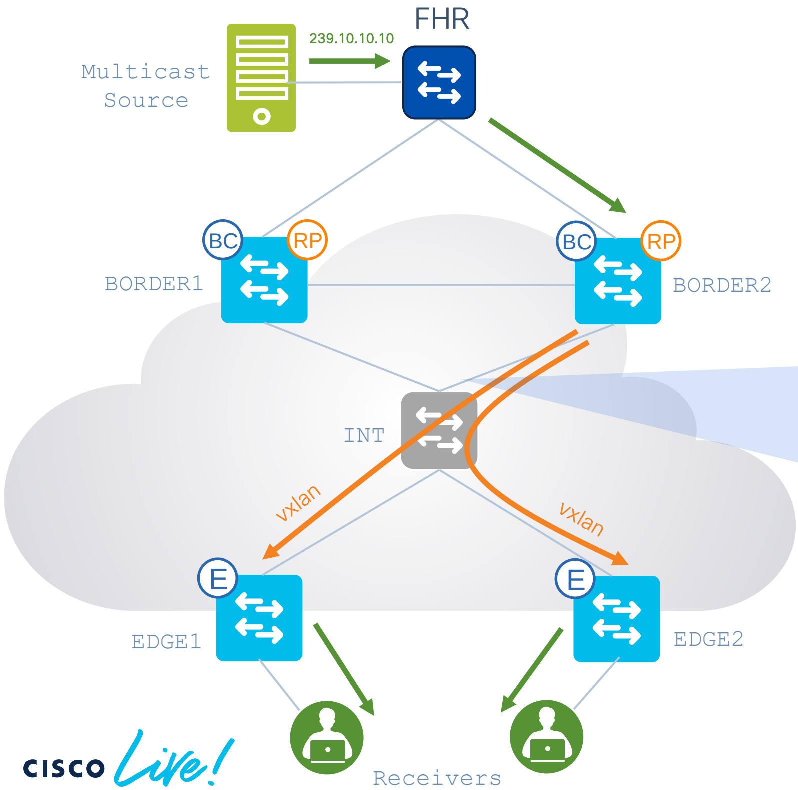

icmp.seq == 51733

No.	Time	Source	Destination	Protocol	Length	Info
…	2.546428	192.168.201.2	239.10.10.10	ICMP	124	Echo (ping) request id=0x0002, seq=51733/5578
…	2.546548	192.168.201.2	239.10.10.10	ICMP	124	Echo (ping) request id=0x0002, seq=51733/5578

```
> Frame 13: 124 bytes on wire (992 bits), 124 bytes captured (992 bits) on interface 0
> Ethernet II, Src: Cisco_d4:8a:07 (00:6b:f1:d4:8a:07), Dst: Cisco_38:b3:74 (00:42:5a:38:b3:74)
> Internet Protocol Version 4, Src: 192.168.10.200, Dst: 192.168.10.4
> User Datagram Protocol, Src Port: 60842, Dst Port: 4789
v Virtual eXtensible Local Area Network
  > Flags: 0x8840, GBP Extension, Don't Learn, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 4099
    Reserved: 0
> Ethernet II, Src: Cisco_d4:8a:07 (00:6b:f1:d4:8a:07), Dst: IPv4mcast_0a:0a:0a (01:00:5e:0a:0a:0a)
> Internet Protocol Version 4, Src: 192.168.201.2, Dst: 239.10.10.10
> Internet Control Message Protocol
```

What if receivers connected to 20 x Edge Nodes !?

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Let’s get some captures!

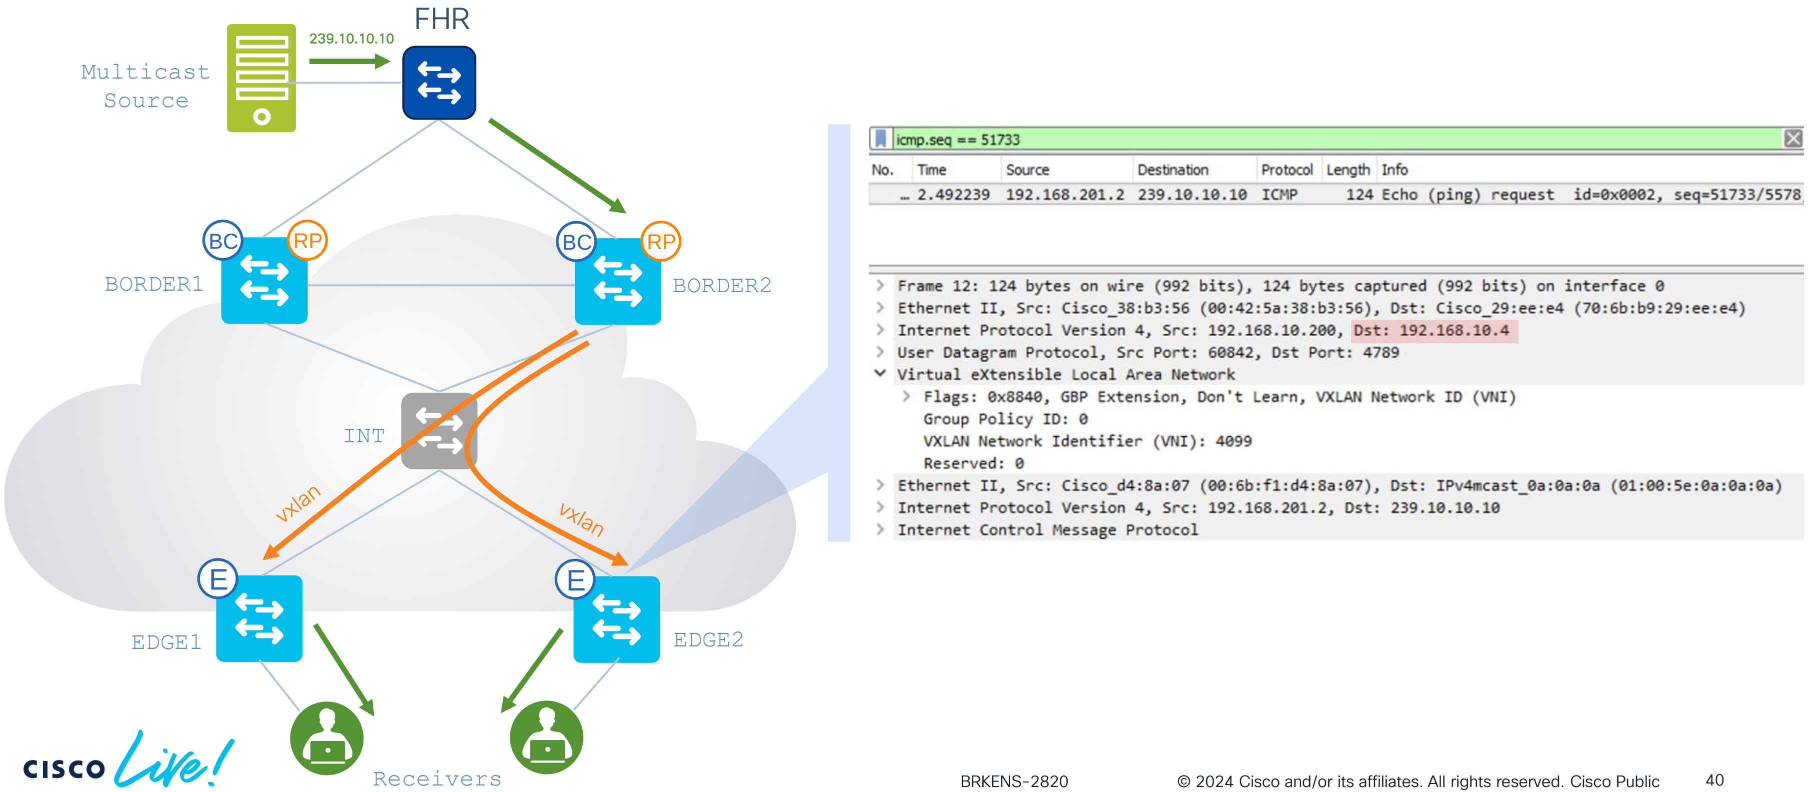

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Final state.

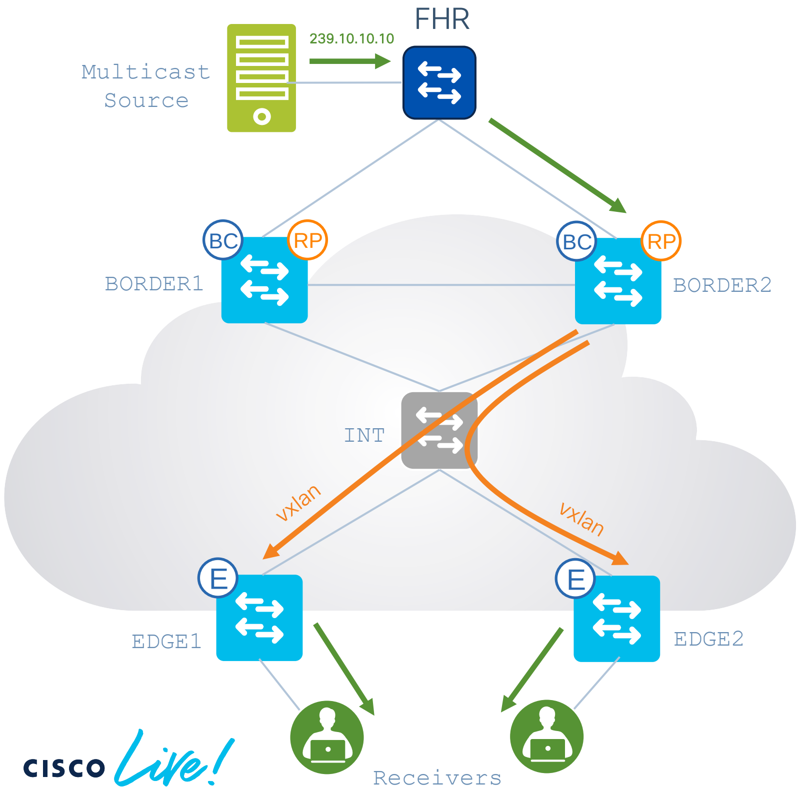

1. Multicast traffic forwarded to receivers through RPT.

2. What about switching to SPT?

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

What if more than two border nodes?

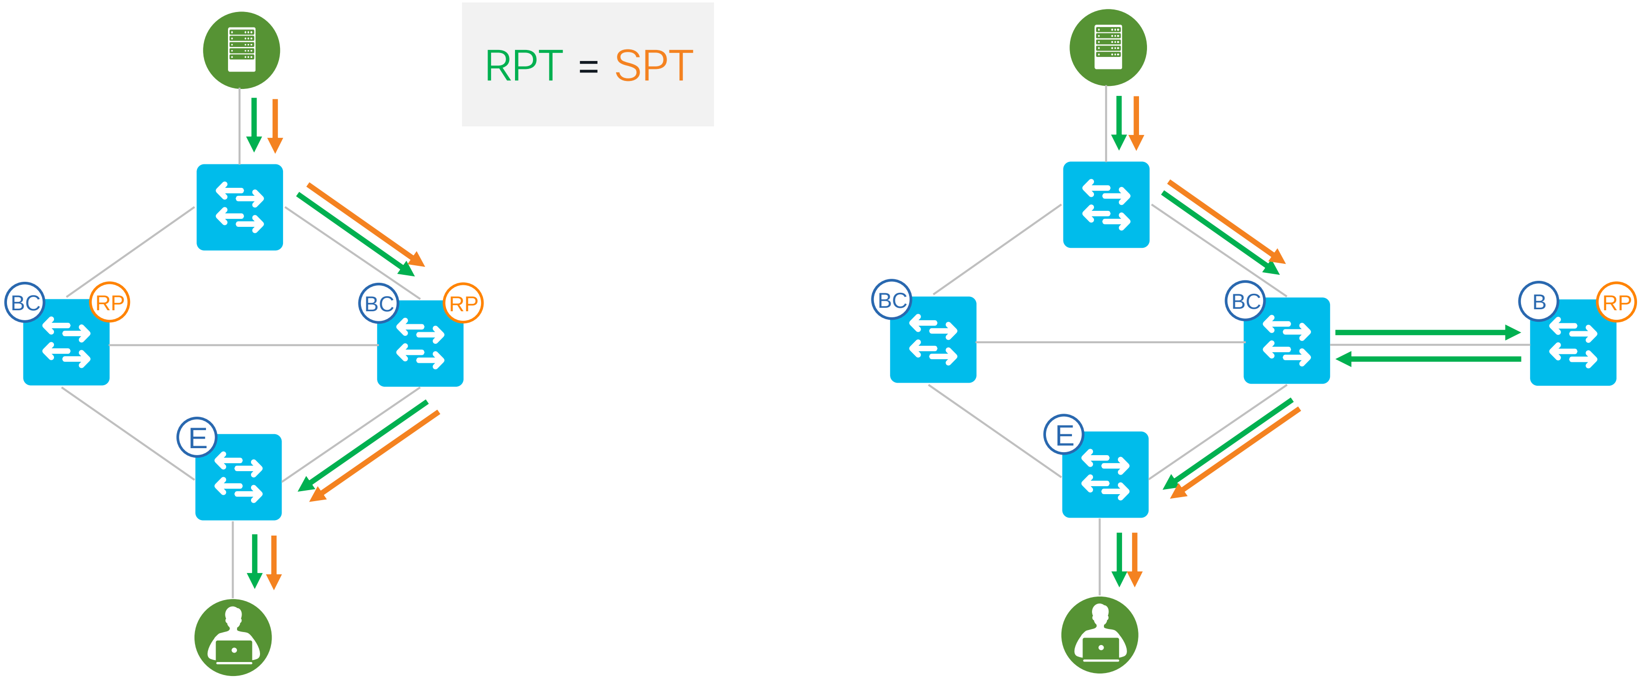

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

More than two border nodes.

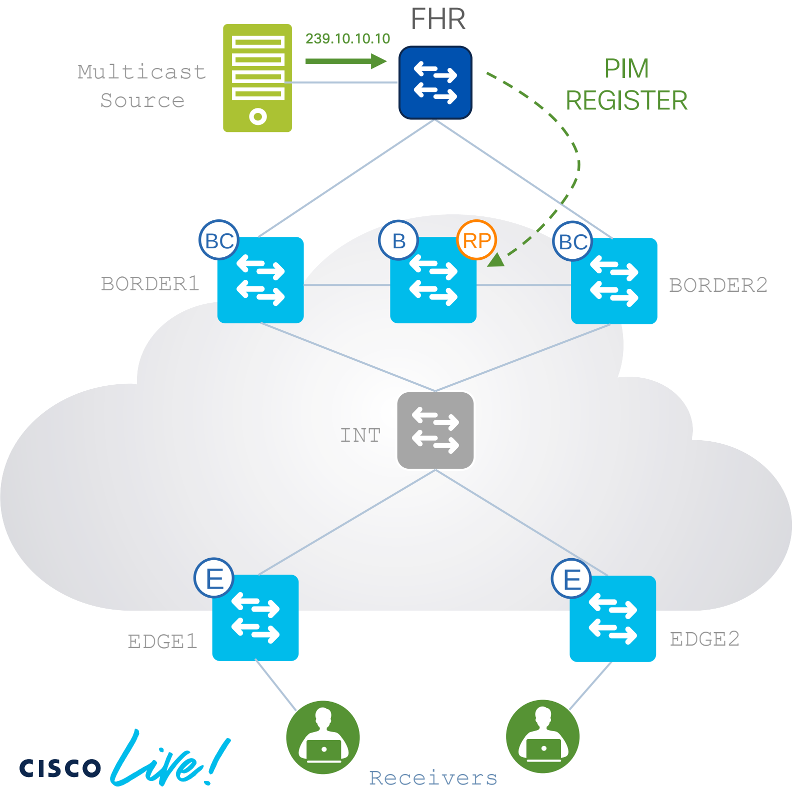

1. Receivers are connected to Edge Nodes (LHR). Multicast source is located outside the fabric.
2. Rendezvous Point is present in overlay as a part of endpoint IP space.
3. PIM-SM is enabled in overlay between fabric devices and outside the fabric within VN/VRF.
4. Source starts to stream multicast traffic. FHR receives multicast stream and triggers any-source PIM Register towards RP.
5. If there are no receivers yet, RP sends PIM Register Stop message up to FHR and waits.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

More than two border nodes.

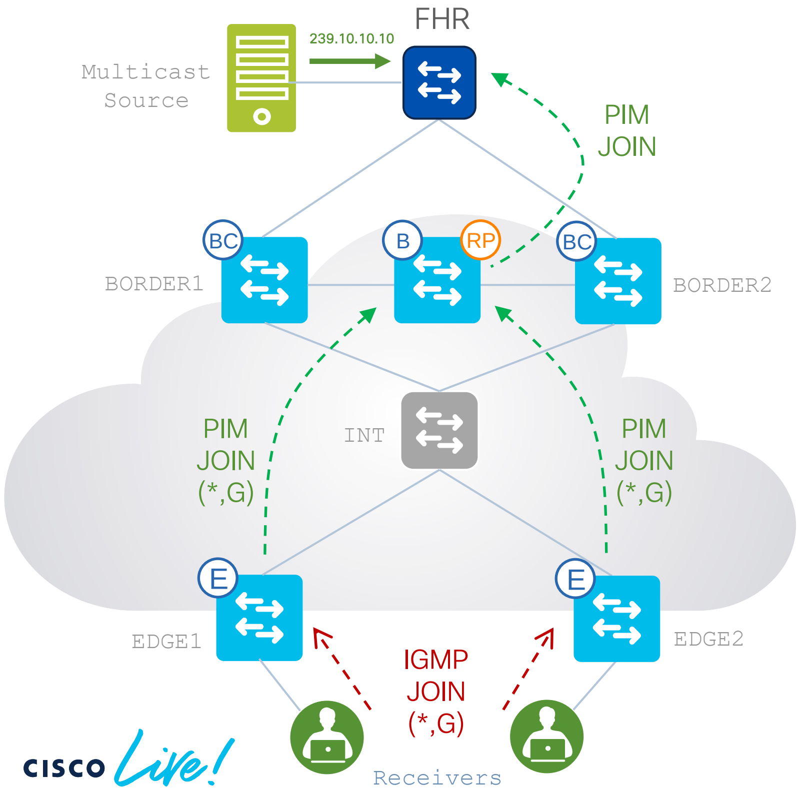

1. Clients send IGMP Join for any-source multicast group (*,G).
2. The Edge Node receives IGMP Join on SVI and triggers new PIM Join towards RP.
3. The fabric RP now has both (S,G) knowledge for the multicast group. The PIM Join is triggered from RP to FHR.
4. FHR adds interface facing RP to Outgoing Interface List (OIL) and starts forwarding multicast stream.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

More than two border nodes.

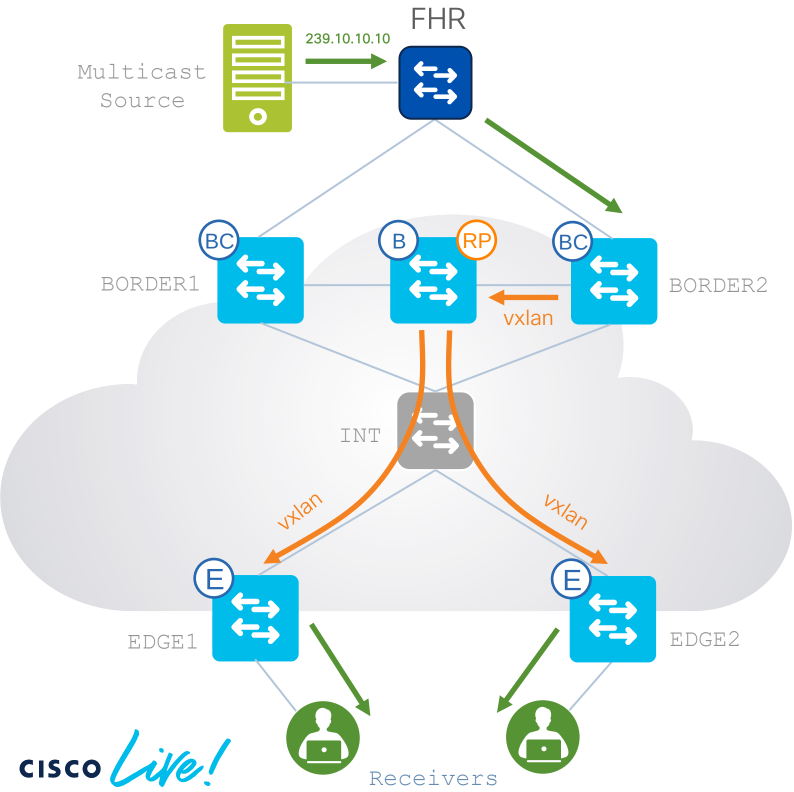

1. Once Border Nodes receives multicast traffic, it will unicast it over the VXLAN to fabric RP.
2. Then, the fabric RP will unicast traffic to Edge Nodes (LHRs) over the VXLAN.
3. Edge Nodes receive VXLAN packets, decapsulate it and send original IP multicast packets towards clients.
4. Classic scenario, multicast traffic forwarded through RPT.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

More than two border nodes.

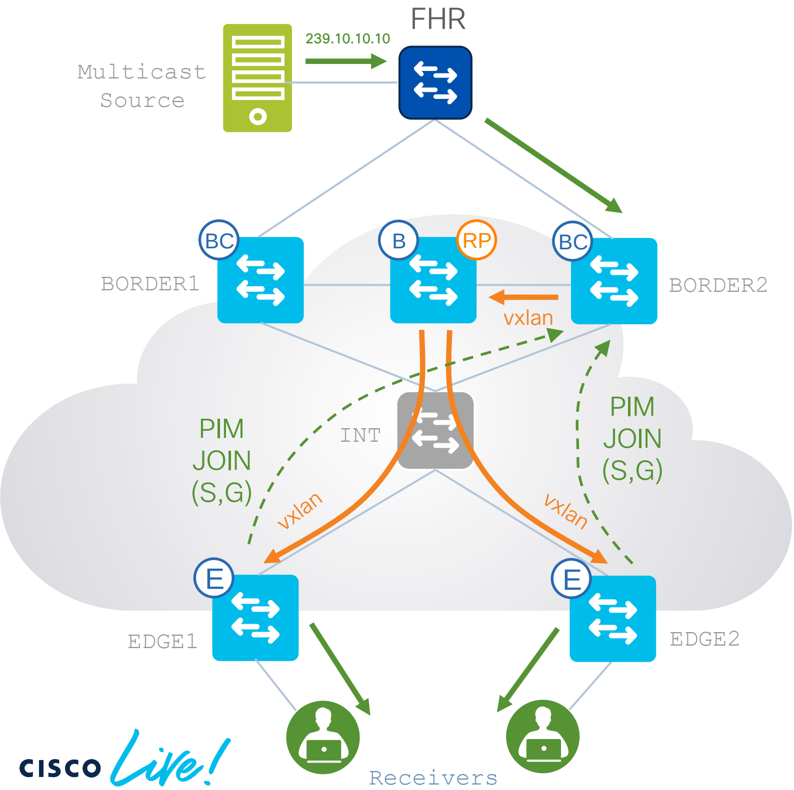

1. Once the first multicast packet arrives on the Edge Node, the shortest path tree (SPT) switchover occurs, which triggers a new PIM Join (S,G) directly to Border Node.

2. The Border Node now knows which Edge Nodes have clients attached based on the received PIM Join message.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

More than two border nodes.

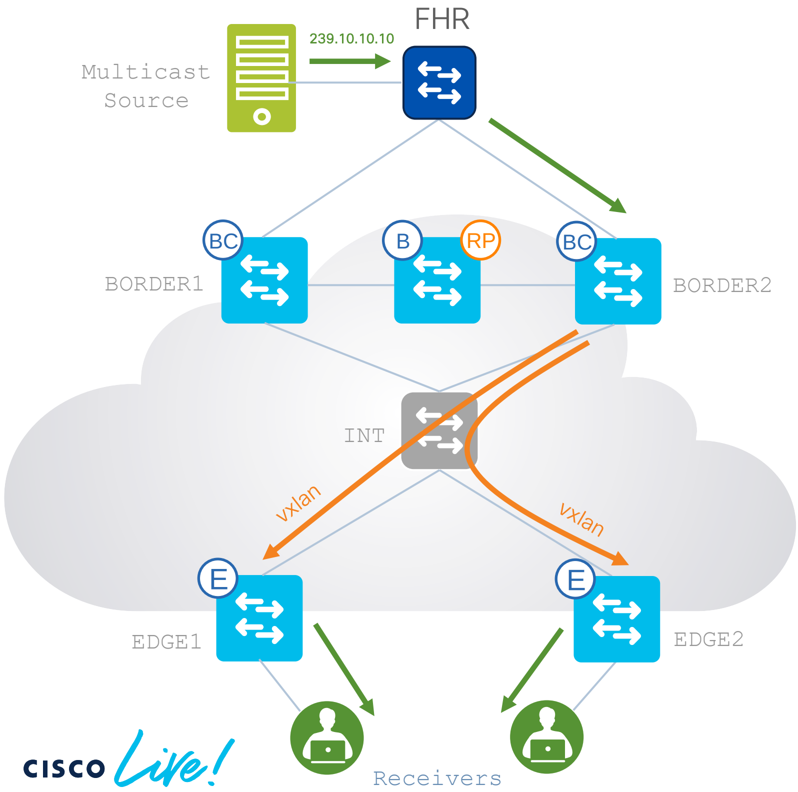

1. The Border Node creates a copy of the original multicast packet for each LHR (Edge Node), encapsulate it in VXLAN and then unicasts it.

2. Each LHR receives VXLAN packet, decapsulate it and sends original multicast packet towards the client.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Source and receivers behind Edge Nodes.

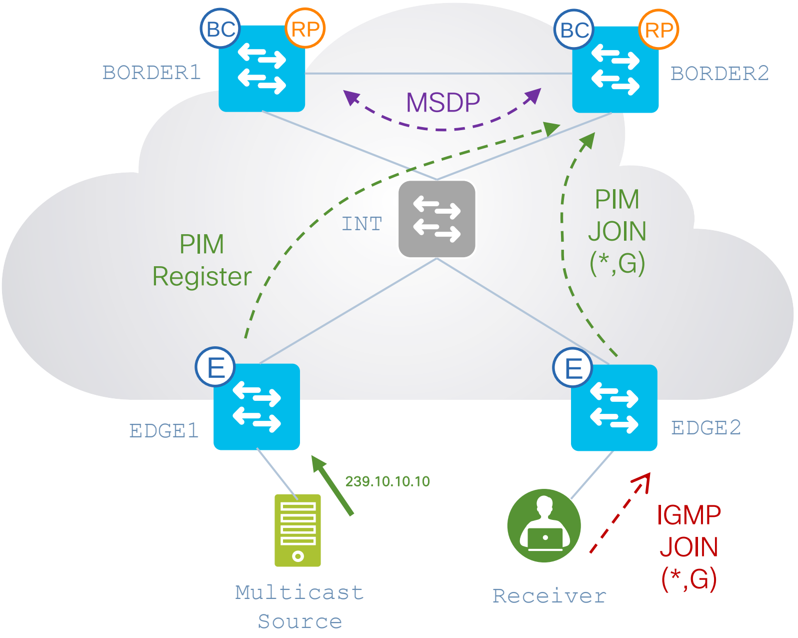

1. Receiver and source are connected to Edge Nodes (FHR & LHR).
2. Rendezvous Points (RPs) are present in overlay as a part of endpoint IP space.
3. PIM-SM is enabled in overlay between fabric devices.
4. Source starts to stream multicast traffic. FHR Edge Node receives multicast stream and triggers any-source PIM Register towards RP.
5. Receiver request multicast traffic by sending IGMP Join (*,G). LHR Edge Node triggers PIM Join (*,G) towards RP.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Source and receivers behind Edge Nodes.

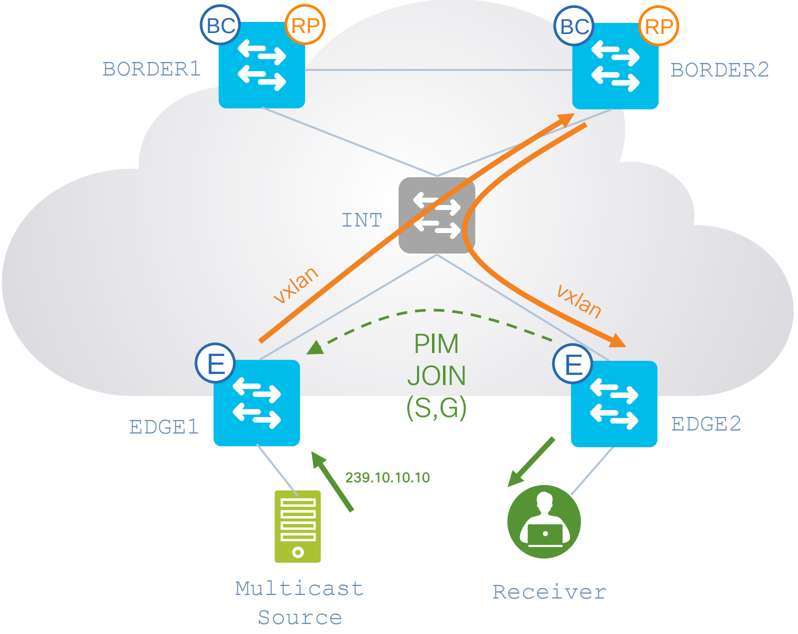

1. FHR Edge Node encapsulates multicast traffic in vxlan and forw towards RP.

2. RP decapsulate multicast traffic and forward it down to LHR Edge Node following RPT.

3. Once the first multicast packet arrives on the LHR Edge Node, the shortest path tree (SPT) switchover occurs, which triggers a new PIM Join directly to FHR Edge Node.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with ASM

Source and receivers behind Edge Nodes.

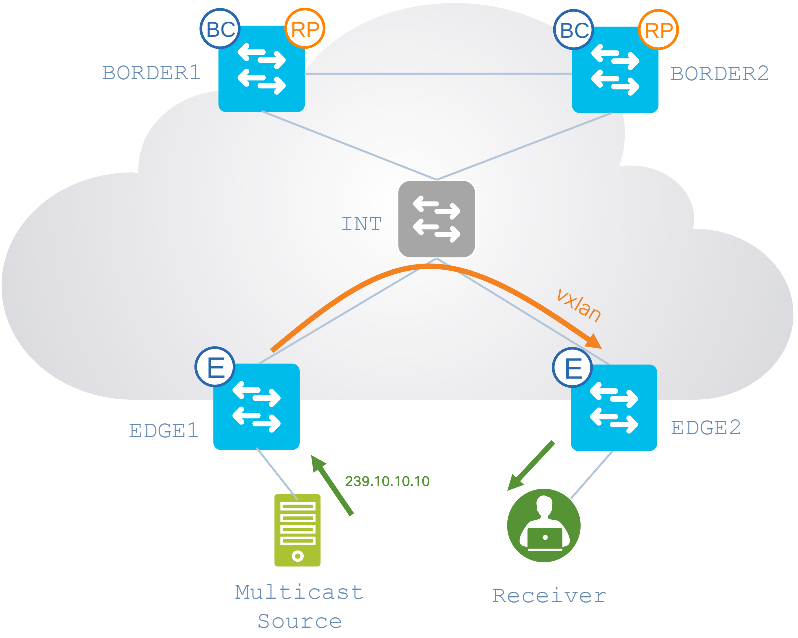

1. The FHR Edge Node creates a copy of the original multicast packet for each LHR, encapsulate it in VXLAN and then unicasts it.

2. The LHR Edge Node receives VXLAN packet, decapsulate it and sends original multicast packet towards the client.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head-end replication with SSM

Head-end replication with SSM (configuration)

- Configuration steps are exactly the same as for Head-end replication with ASM.
- As a result of provisioning process the **ip pim vrf <vrf-name> ssm default** command is pushed to all fabric devices which enables SSM in the overlay.
- When host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S, RP is not contacted in this process by the receiver.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

DEMO #1:

Head-end replication with ASM

https://youtu.be/O7oQuge5vDA

Lab setup used for demo

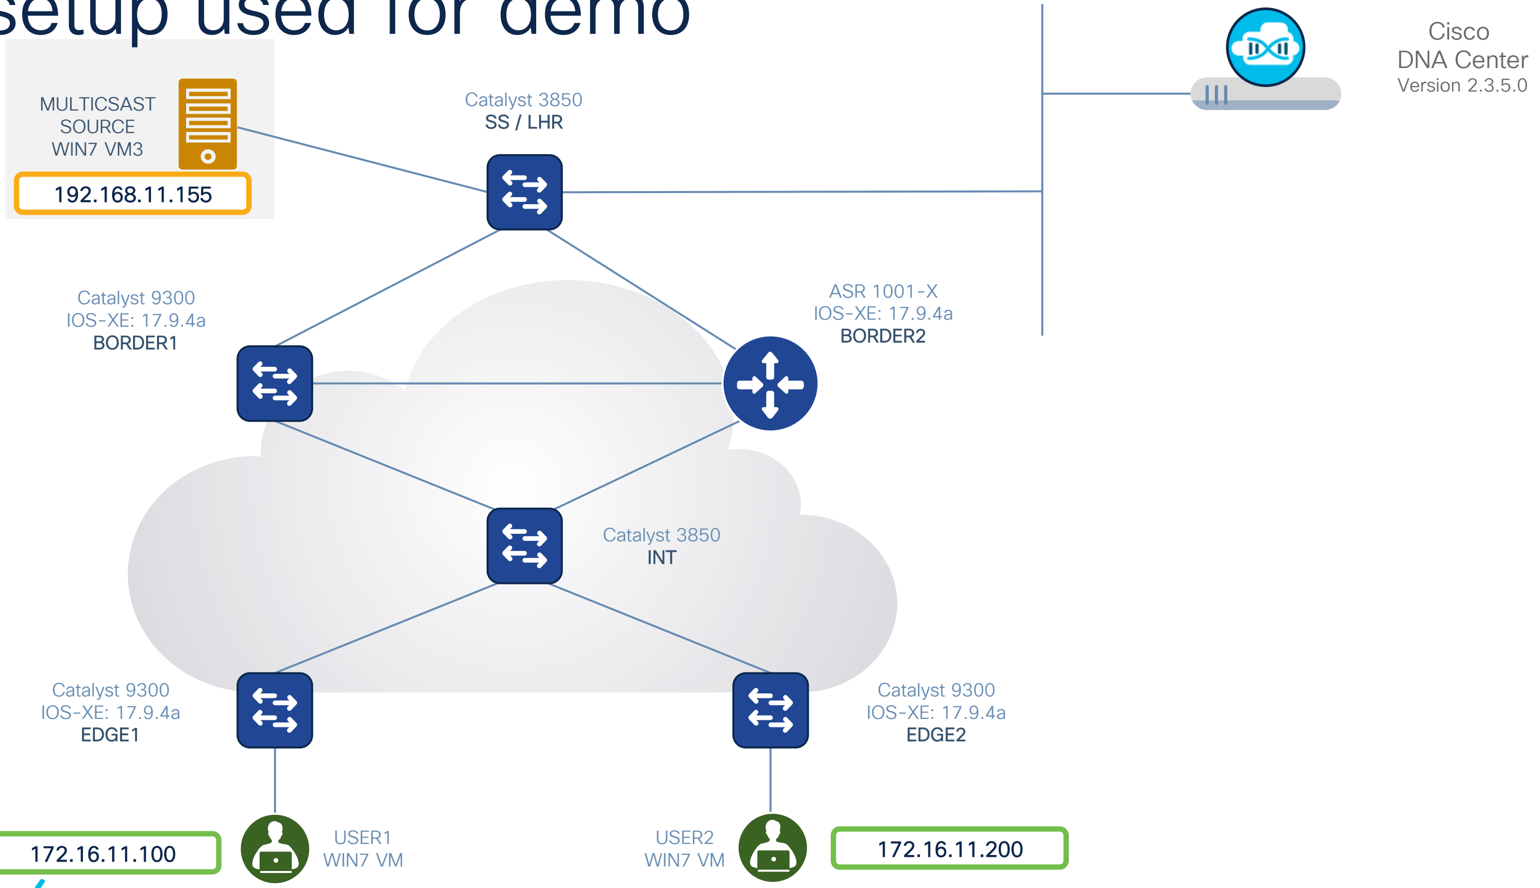

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

Native multicast with ASM (underlay config)

By LAN Automation.

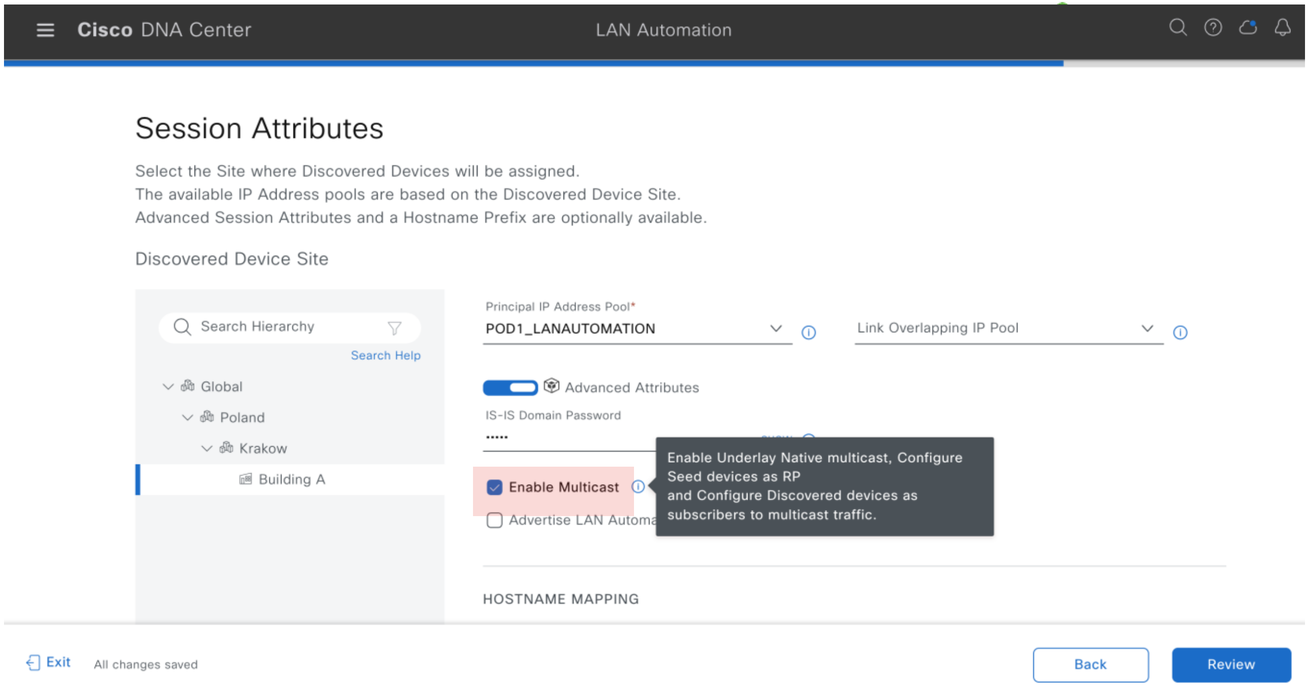

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (underlay config)

Manual approach (CLI), minimal configuration.

All of the underlay devices (including Intermediate Nodes)

```
ip multicast-routing
ip pim ssm default
```

All of the L3 interfaces

```
 ip pim sparse-mode
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (underlay config)

Manual approach (Templated Editor), minimal configuration.

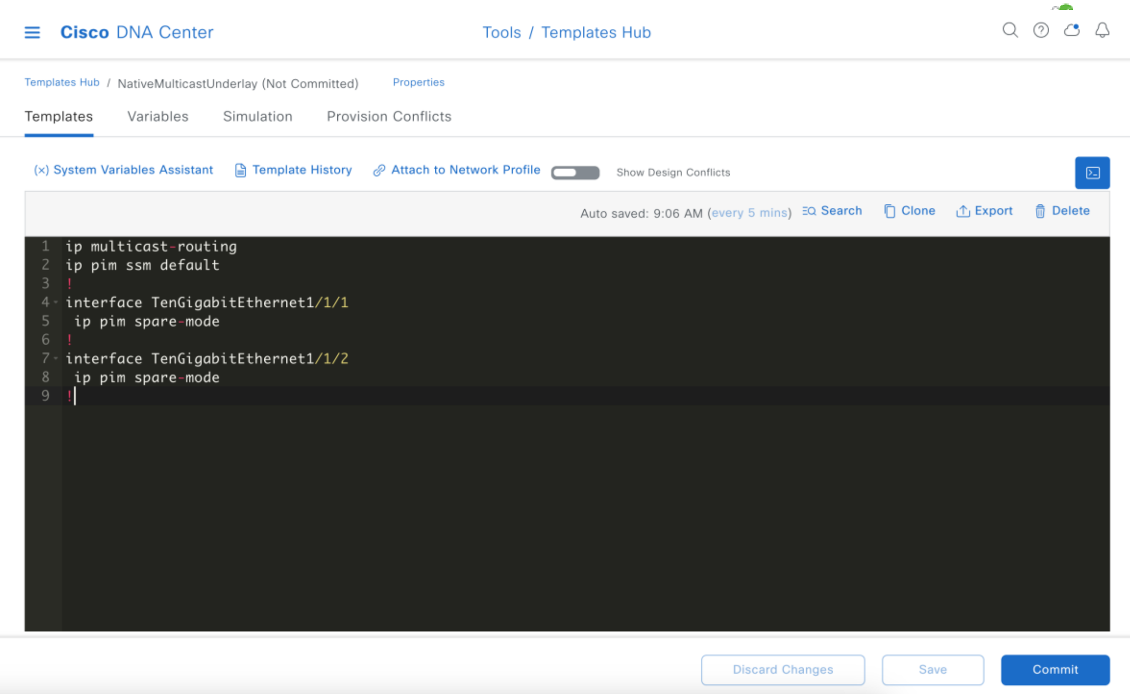

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (config)

New Workflow

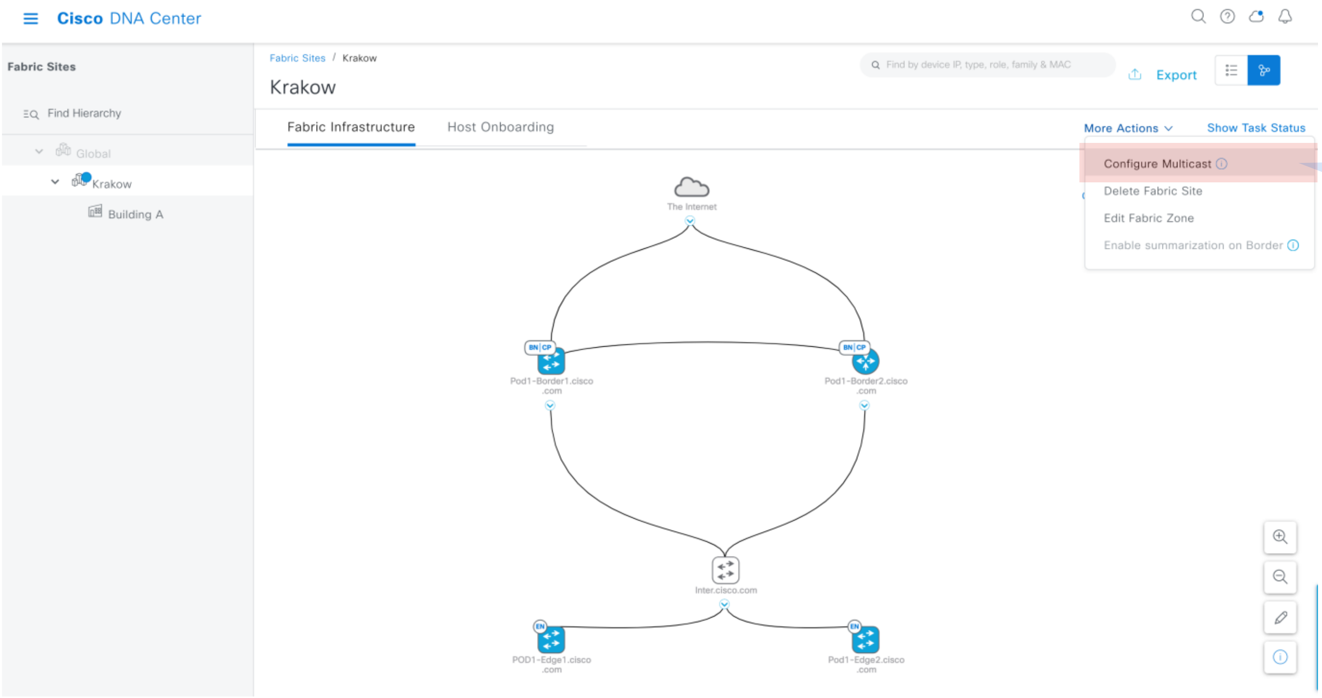

New GUI workflow introduced in 2.3.5.X

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (config)

Selection of replication mode

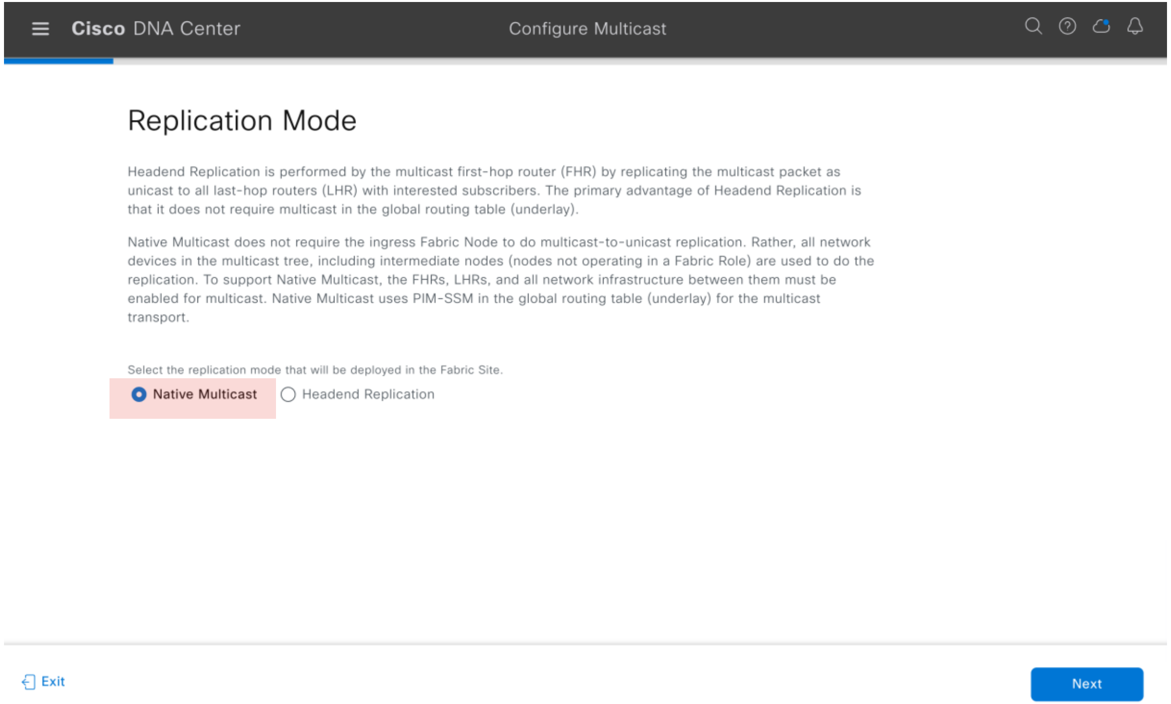

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (config)

Selection of Multicast Mode

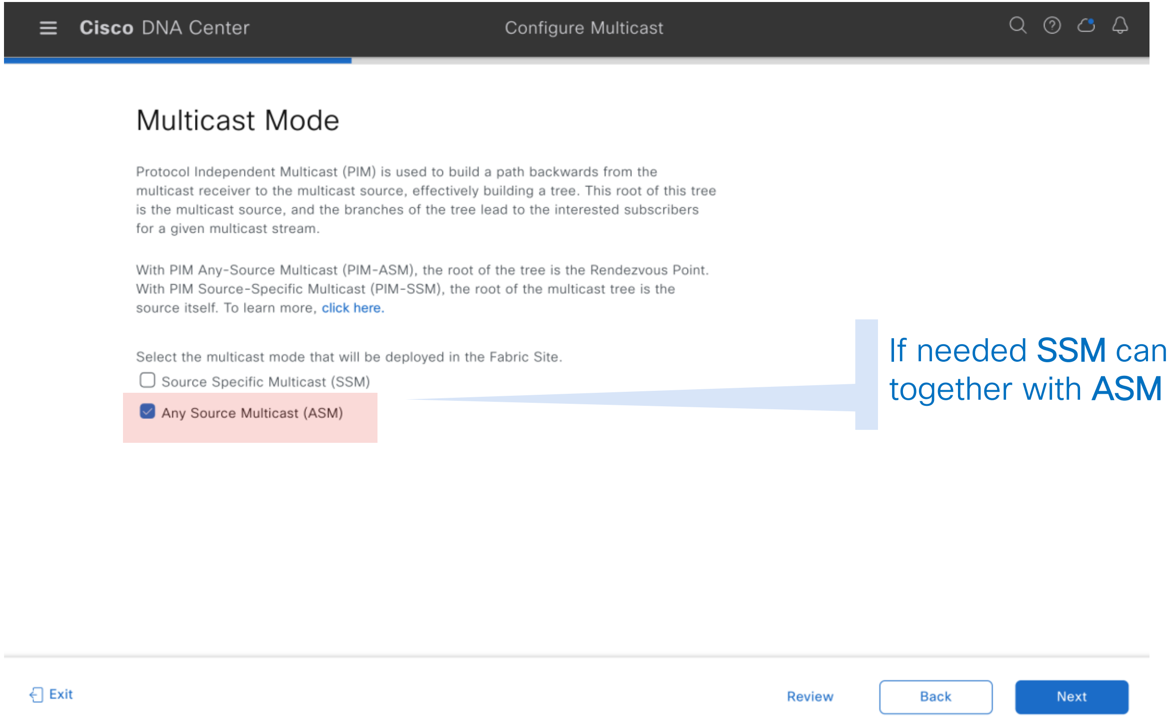

If needed **SSM** can be enabled together with **ASM**

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (config)

RP Mapping

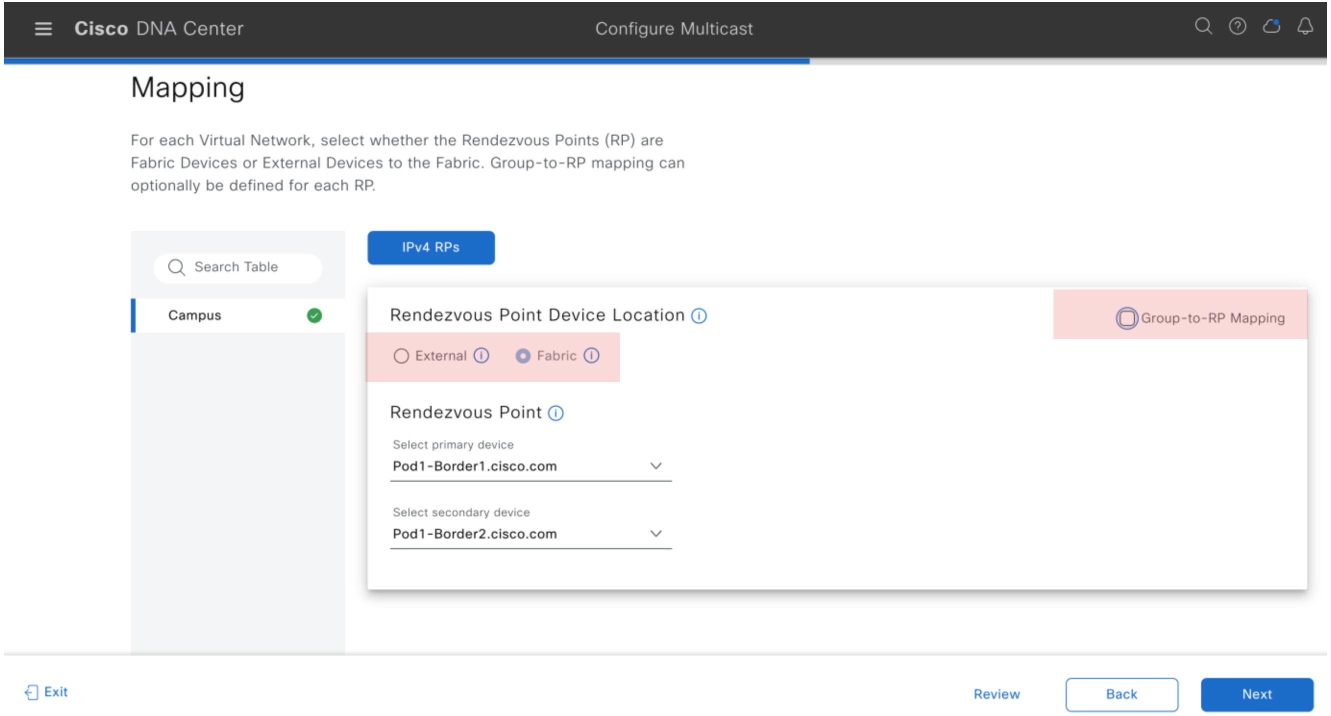

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (config)

RP Mapping (Group-to-RP Mapping)

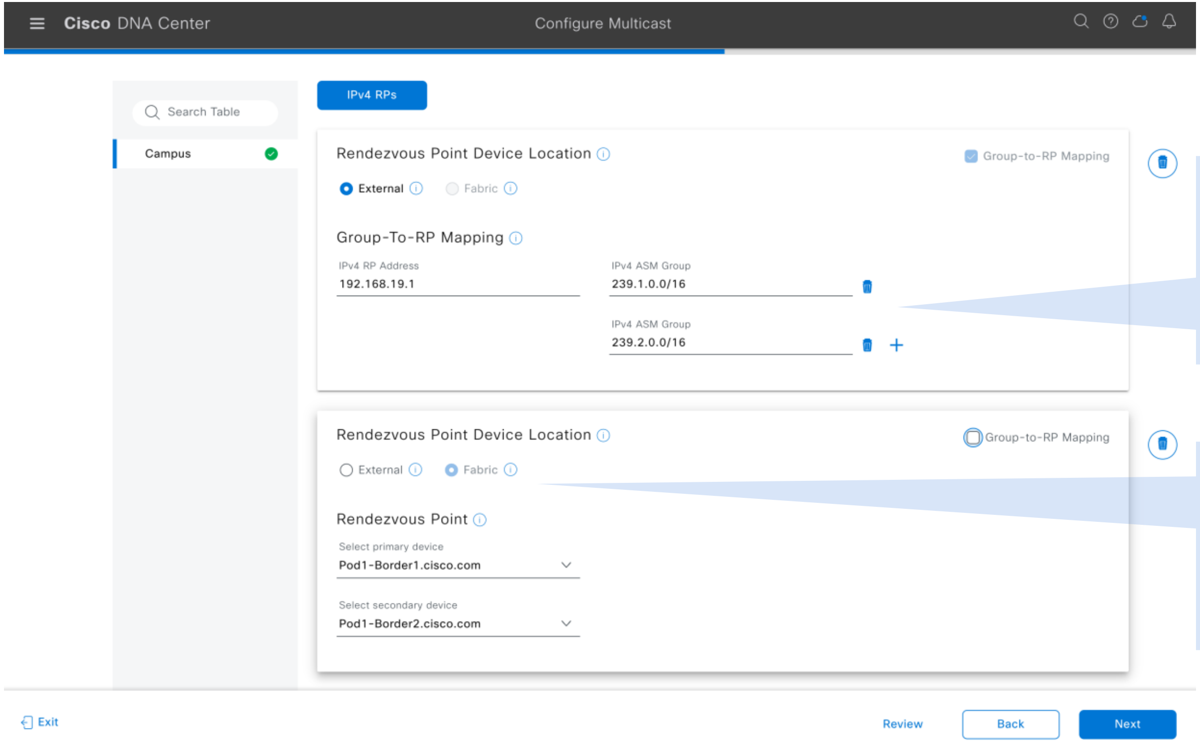

External RP for only 2 selected ranges

- `239.1.0.0 0.0.255.255`
- `239.2.0.0 0.0.255.255`

Internal RP for all others

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM (config)

RP Mapping (Group-to-RP Mapping)

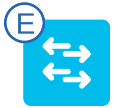


```
<snip>
!
ip pim vrf Campus rp-address 192.168.19.1 ASM_ACL_IPV4_Campus_192.168.19.1
ip pim vrf Campus rp-address 172.16.200.1

ip access-list standard ASM_ACL_IPV4_Campus_192.168.19.1
 10 permit 239.1.0.0 0.0.255.255
 20 permit 239.2.0.0 0.0.255.255
!
<snip>
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

A closer look. Initial status (underlay).

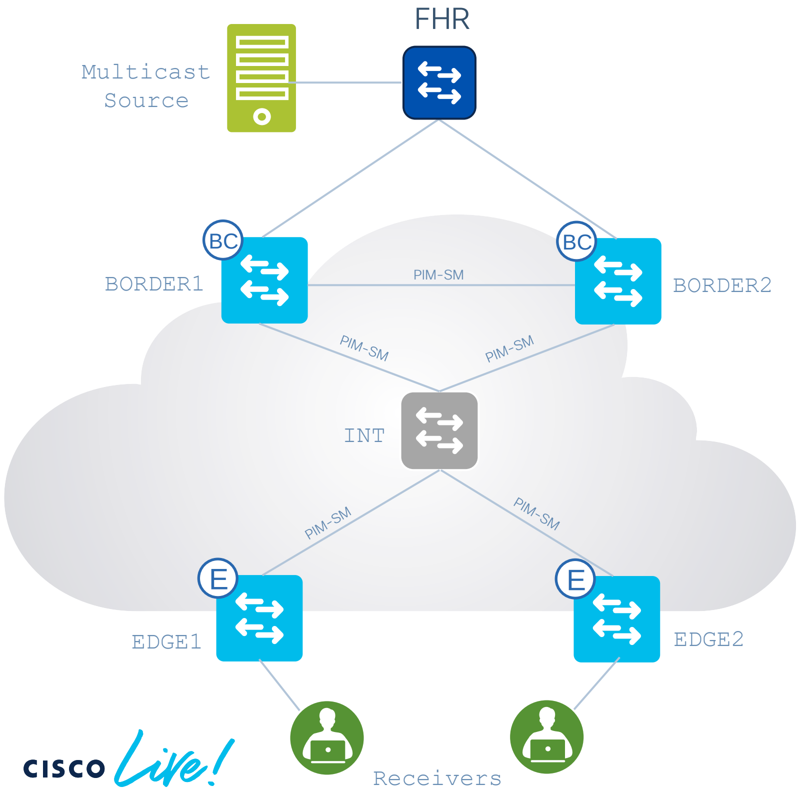

1. Multicast routing is enabled on all underlay devices.
2. SSM is enabled on all underlay devices.
3. PIM-SM is enabled on all L3 interfaces on all underlay devices.
4. Native multicast is enabled for particular site in Cisco DNA Center GUI.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

A closer look. Initial status (overlay).

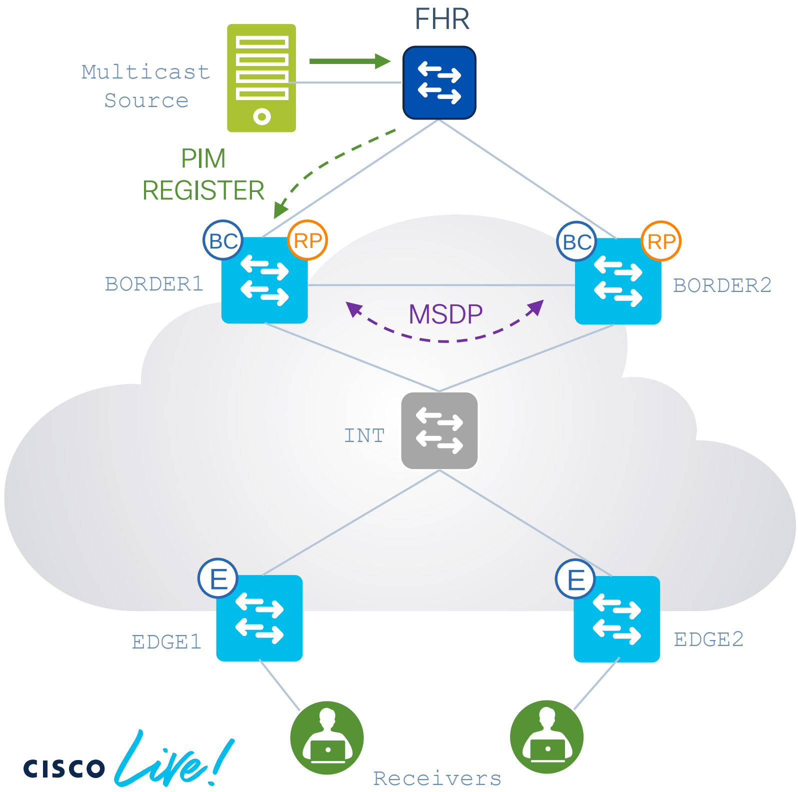

1. Receivers are connected to Edge Nodes (LHRs). Multicast Source is located outside the fabric.
2. Rendezvous Points (RPs) are present in overlay as a part of endpoint IP space.
3. PIM-SM is enabled in overlay between fabric devices and outside the fabric within VN/VRF.
4. Source starts to stream multicast traffic. FHR receives multicast stream and triggers any-source PIM Register towards RP.
5. RP updates a peer RP about the source via MSDP.
6. If there are no receivers yet, RP sends PIM Register Stop message up to FHR and waits.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

A closer look. Clients request multicast traffic.

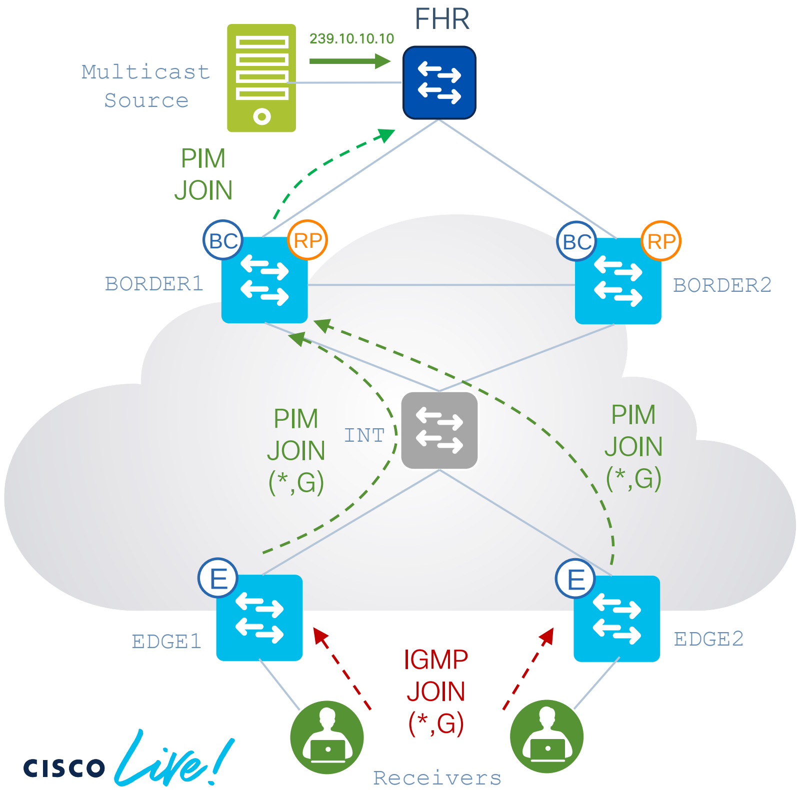

1. Clients send IGMP Join for any-source multicast group (*,G).
2. The Edge Node receives IGMP Join on SVI and triggers new PIM Join towards RP. Since RPs have same IP address (anycast), the hashing algorithm will be used to forward the packet (LISP next-hop towards RP).
3. The fabric RP now has both (S,G) knowledge for the multicast group and triggers PIM Join to FHR.
4. FHR adds interface facing RP to Outgoing Interface List (OIL) and starts forwarding multicast traffic.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

A closer look. Multicast stream forwarded to receivers.

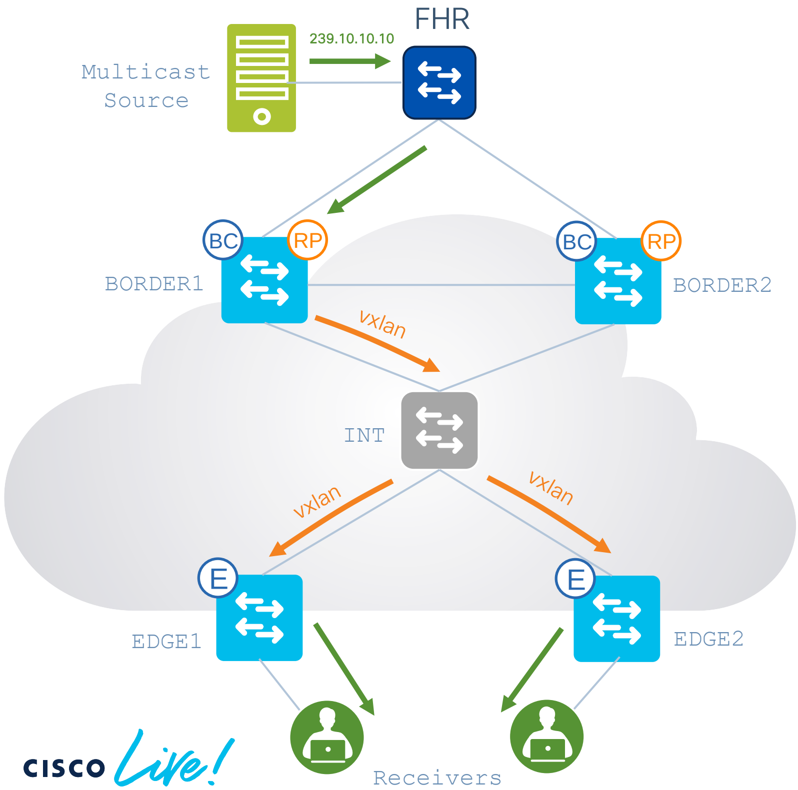

1. The Border Node receives a multicast packet from FHR.
2. The Border Node copies original multicast packet to underlay multicast SSM tree (232.0.0.0/8) and sends it towards Intermediate Node. As a destination IP address in VXLAN header, the multicast group address is used.
3. Intermediate Node replicates original multicast packet and forward one copy to each of the Edge Nodes where receivers are connected.
4. Edge Nodes (LHRs) decapsulate VXLAN packet, put it back to multicast tree in respective VN and forward it directly to receivers.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

Let’s get some captures!

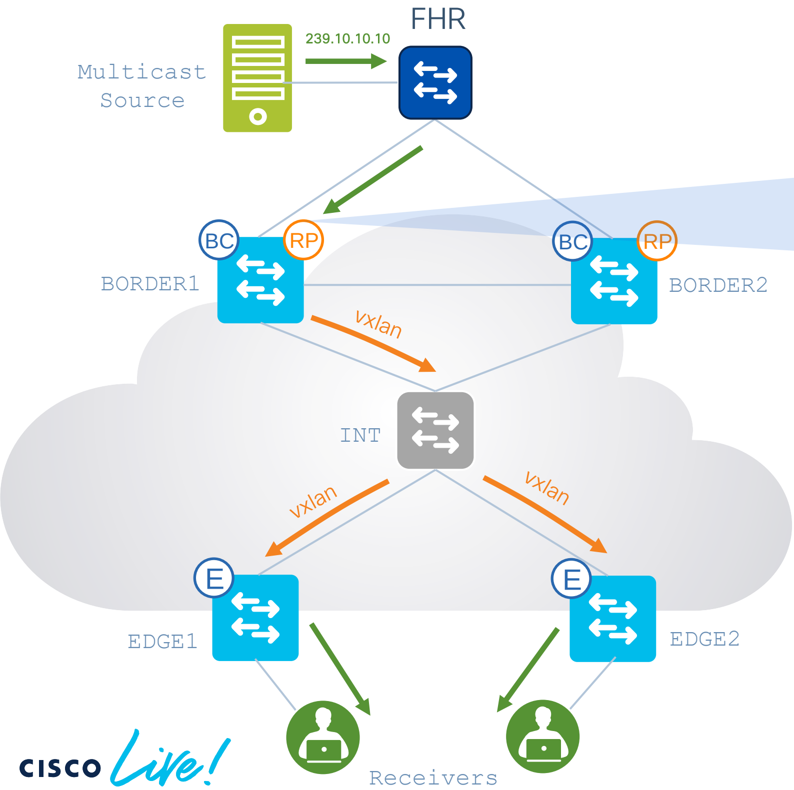


```
icmp.seq == 55082

No.  Time      Source         Destination   Protocol  Length Info
  2 0.183978  192.168.201.2  239.10.10.10  ICMP          78 Echo (ping) request  id=0x0002

> Frame 2: 78 bytes on wire (624 bits), 78 bytes captured (624 bits) on interface 0
> Ethernet II, Src: Cisco_15:4b:64 (a0:f8:49:15:4b:64), Dst: IPv4mcast_0a:0a:0a (01:00:5e:0a:0a:0a)
> 802.1Q Virtual LAN, PRI: 0, DEI: 0, ID: 3030
> Internet Protocol Version 4, Src: 192.168.201.2, Dst: 239.10.10.10
> Internet Control Message Protocol
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

Let’s get some captures!

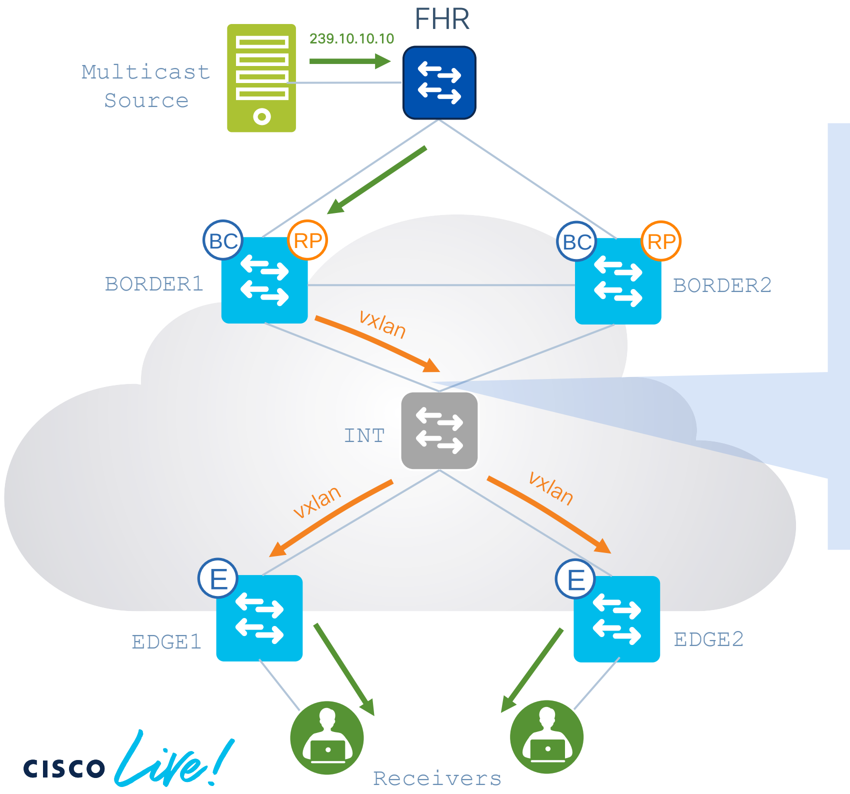

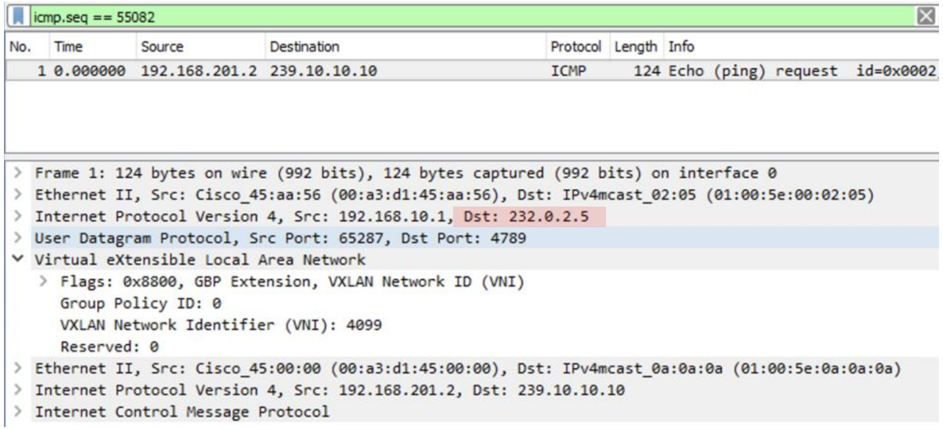

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with ASM

Let's get some captures!

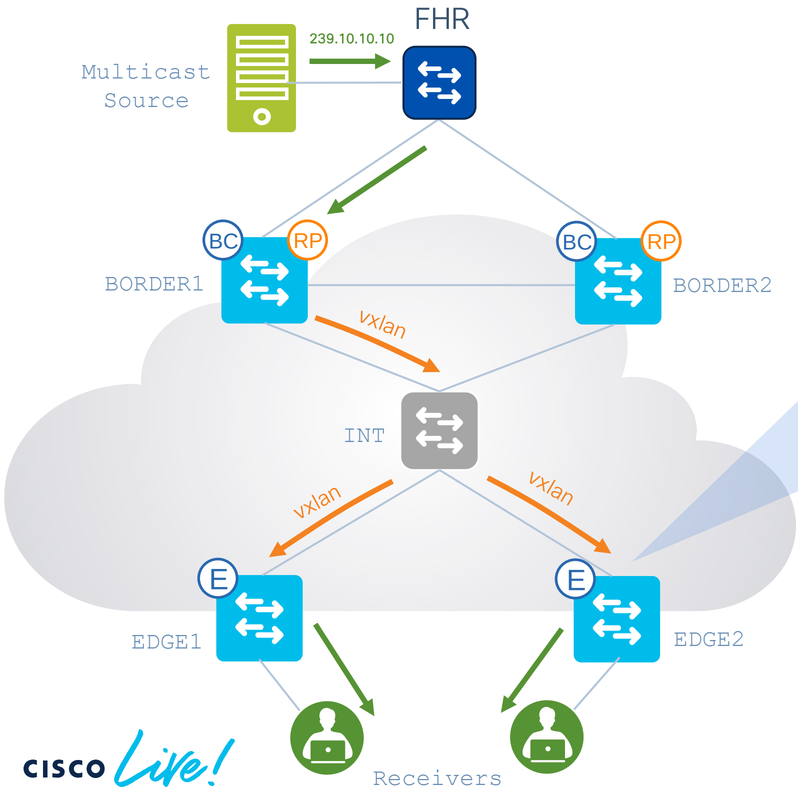

icmp.seq == 55082

No.	Time	Source	Destination	Protocol	Length	Info
2	0.100809	192.168.201.2	239.10.10.10	ICMP	124	Echo (ping) request id=0x0002

```
> Frame 2: 124 bytes on wire (992 bits), 124 bytes captured (992 bits) on interface 0
> Ethernet II, Src: Cisco_38:b3:56 (00:42:5a:38:b3:56), Dst: IPv4mcast_02:05 (01:00:5e:00:02:05)
> Internet Protocol Version 4, Src: 192.168.10.1, Dst: 232.0.2.5
> User Datagram Protocol, Src Port: 65287, Dst Port: 4789
v Virtual eXtensible Local Area Network
  > Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 4099
    Reserved: 0
> Ethernet II, Src: Cisco_45:00:00 (00:a3:d1:45:00:00), Dst: IPv4mcast_0a:0a:0a (01:00:5e:0a:0a:0a)
> Internet Protocol Version 4, Src: 192.168.201.2, Dst: 239.10.10.10
> Internet Control Message Protocol
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native multicast with SSM

Native multicast with SSM (configuration)

- Configuration steps are exactly the same as for native multicast with ASM in overlay.
- The PIM SSM in underlay is still required to be configured (either by LAN Automation feature or manually).
- When host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S, RP is not contacted in this process by the receiver.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Other Use Cases

Fabric Enabled Wireless – basic workflow

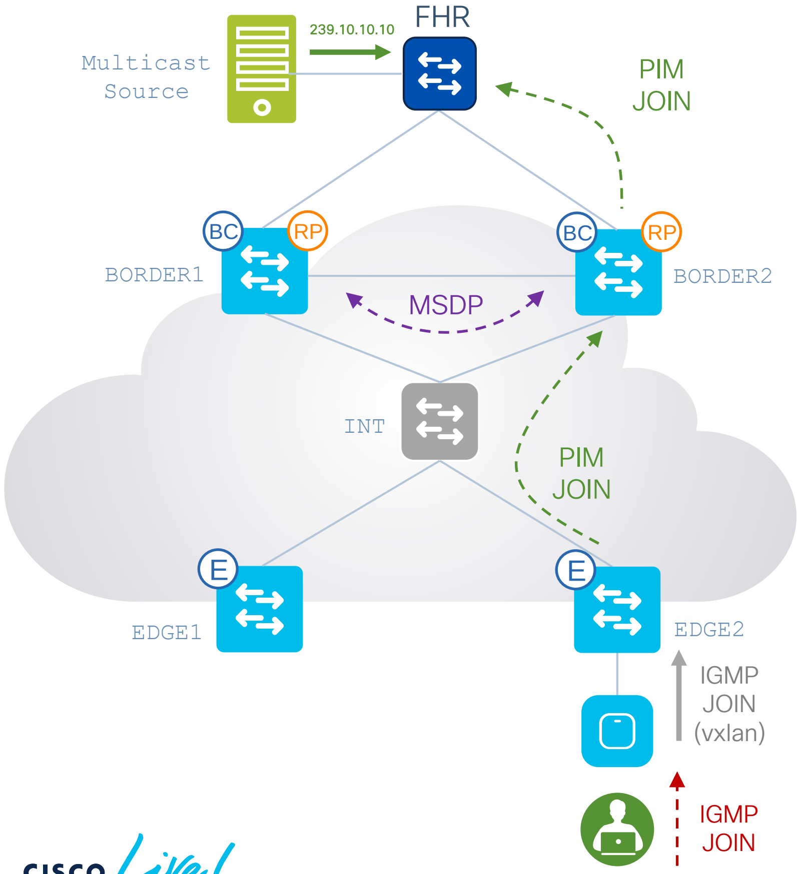

1. Let's assume source is already registered to RP and we use head-end replication with ASM.

2. Client requests multicast traffic by sending IGMP Join.

3. AP encapsulate it in VXLAN and sends over to Edge Node via AccessTunnel interface.

4. Edge Node receives it and triggers PIM Join towards RP.

5. RP triggers PIM Join towards FHR to request multicast traffic.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Fabric Enabled Wireless – basic workflow

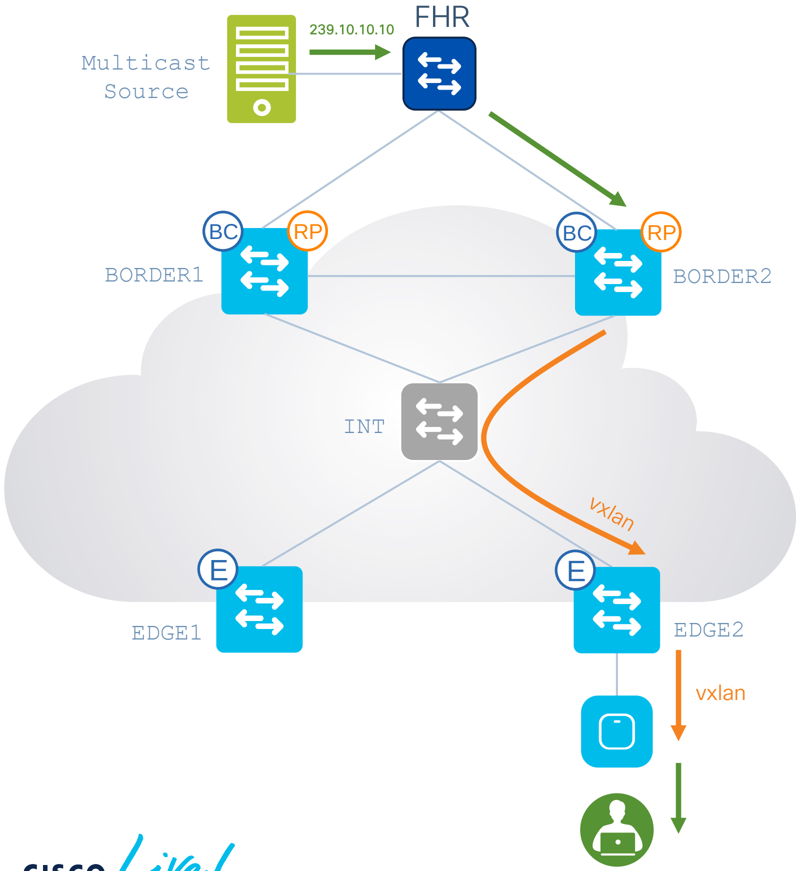

1. RP receives IP multicast packet and creates a copy for each Edge Node where APs with wireless clients requesting multicast stream are connected.

2. RP encapsulate it in VXLAN and send it over towards the Edge Node.

3. Edge Node decapsulate the traffic and encapsulate it again in VXLAN in order to send it through AccessTunnel interface.

4. The AP removes the VXLAN header and send the original IP multicast packet into the air towards wireless clients.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Extended Node – basic workflow

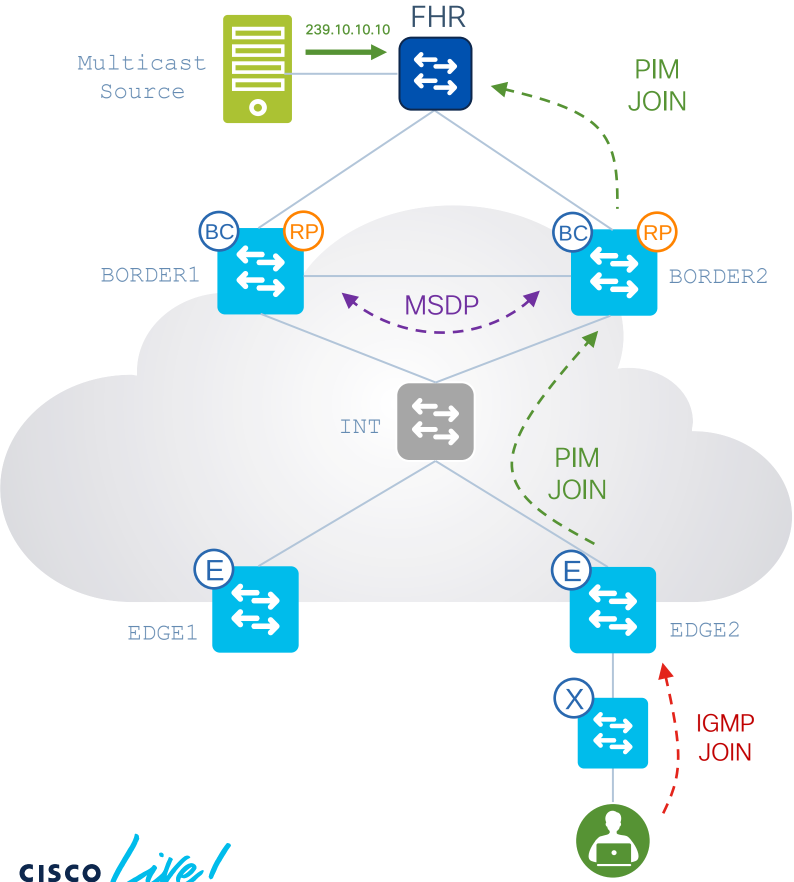

1. Let's assume source is already registered to RP and we use head-end replication with ASM.
2. Client requests multicast traffic by sending IGMP Join.
3. Extended Node as L2 device forwards it towards the querier (Edge Node).
4. Edge Node receives it and triggers PIM Join towards RP.
5. RP triggers PIM Join towards FHR to request multicast traffic.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Extended Node – basic workflow

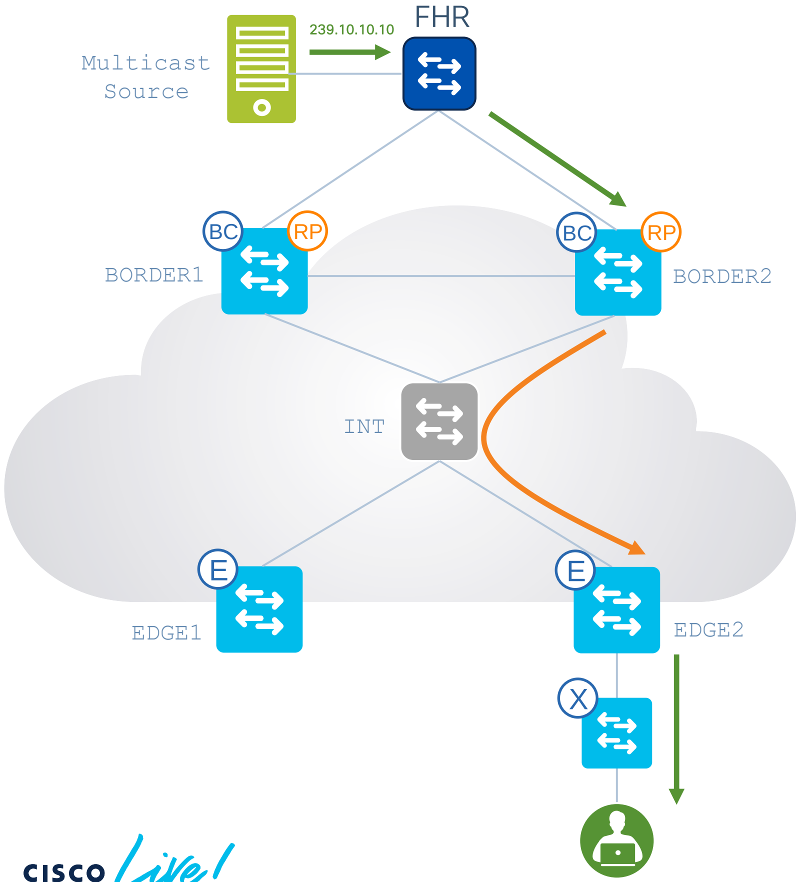

1. RP receives IP multicast packet and creates a copy for each Edge Node where clients requesting multicast stream are connected (no matter if directly or through Extended Node).

2. RP encapsulate it in VXLAN and send it over towards the Edge Node.

3. Edge Node decapsulate the traffic and send it natively towards Extended Node via Port-channel interface (trunk).

4. Extended Node checks IGMP snooping table and replicate IP multicast packet to interfaces where clients are connected.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

(aka selective flooding)

Layer 2 flooding in SD-Access

Why and how?

BUM = Broadcast, Unknown Unicast and Multicast

- By default BUM traffic is not flooded across the SD-Access fabric. The Layer 2 unicast is handled by LISP (MAC address mapped to RLOC).
- With the help of multicast in the underlay we can get BUM traffic flood across SD-Access fabric to be forwarded to all devices with hosts in respected subnet.
- Incoming BUM traffic for a given VN is encapsulated in VXLAN, and then sent with {Source IP = RLOC, Destination IP = Underlay Multicast Group} as the outer IP addresses.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding in SD-Access

Mapping details

INFO

The number of supported IP Pools depends on the DNAC hardware specifications

- In the SD-Access fabric today we support max 1000 IP subnets

- **Prior DNAC 1.3.3.X**

The associated multicast group range associated with Layer 2 flooding is in the range of 239.0.0.1 to 239.0.1.246

Vlan 1021 is assigned to 239.0.0.1
Vlan 1022 is assigned to 239.0.0.2
Vlan 1023 is assigned to 239.0.0.3
…
Vlan 1521 is assigned to 239.0.1.246

- **Post DNAC 1.3.3.X**

The associated multicast group associated with Layer 2 flooding is hardcoded to 239.0.17.1

Vlan 1021 is assigned to 239.0.17.1
Vlan 1022 is assigned to 239.0.17.1
Vlan 1023 is assigned to 239.0.17.1
…
Vlan 2021 is assigned to 239.0.17.1

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding vs Native multicast

A bit of confusion...

Native multicast requires PIM SSM to be configured in the underlay for basic operations.

Layer 2 flooding requires PIM ASM to be configured in the underlay for basic operations.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

A good news is… LAN Automation!

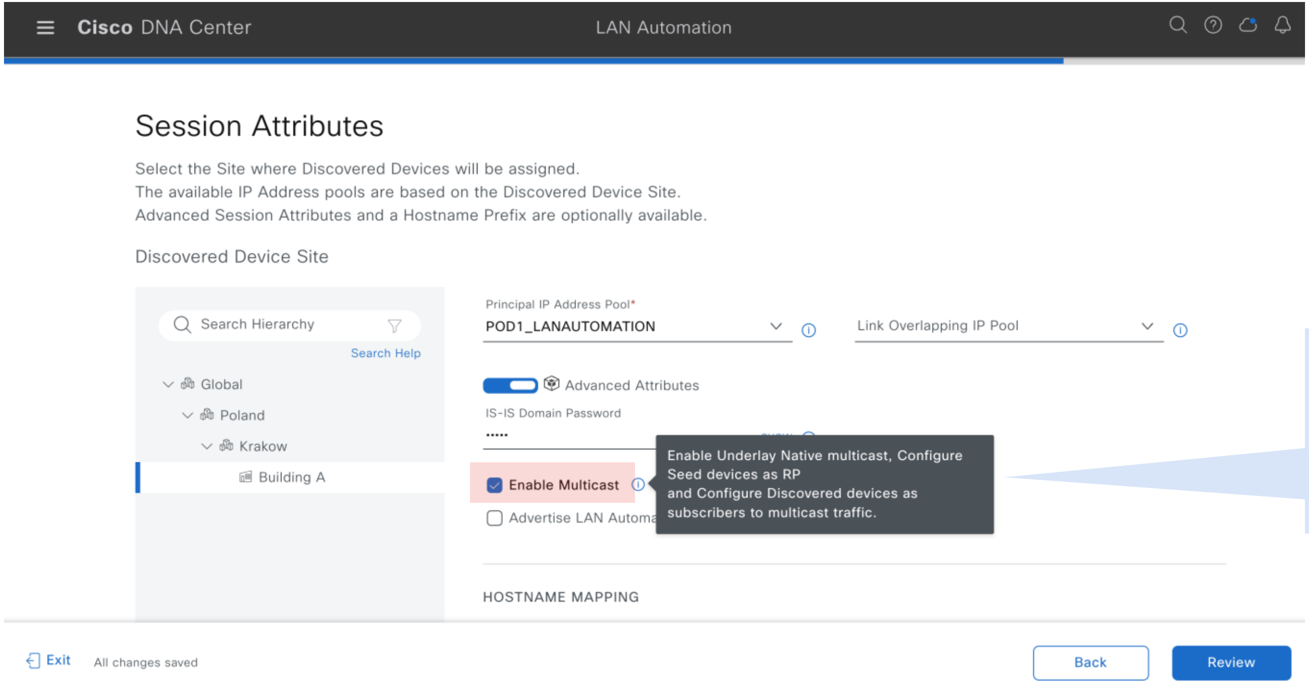

With this option enabled we cover both scenarios:

- Native multicast
- L2 flooding

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

Manual approach (CLI), minimal configuration.

RP devices (typically redundant Border Nodes)

```
interface Loopback60000
 ip address <loopback60000-ip-address> 255.255.255.255
 ip router isis
 ip pim sparse-mode
!
ip multicast-routing
ip pim rp-address <loopback60000-ip-address>
ip pim register-source Loopback60000
!
ip msdp peer <loopback0-other-rp> connect-source Loopback0
```

Interface Loopback60000 (or other chosen) needs to be belong to underlay domain.

Non-RP underlay devices

```
ip multicast-routing
ip pim rp-address <rp-address>
ip pim register-source Loopback0
```

All of the L3 interfaces

```
 ip pim sparse-mode
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

Manual approach (Template Hub/Editor), minimal configuration.

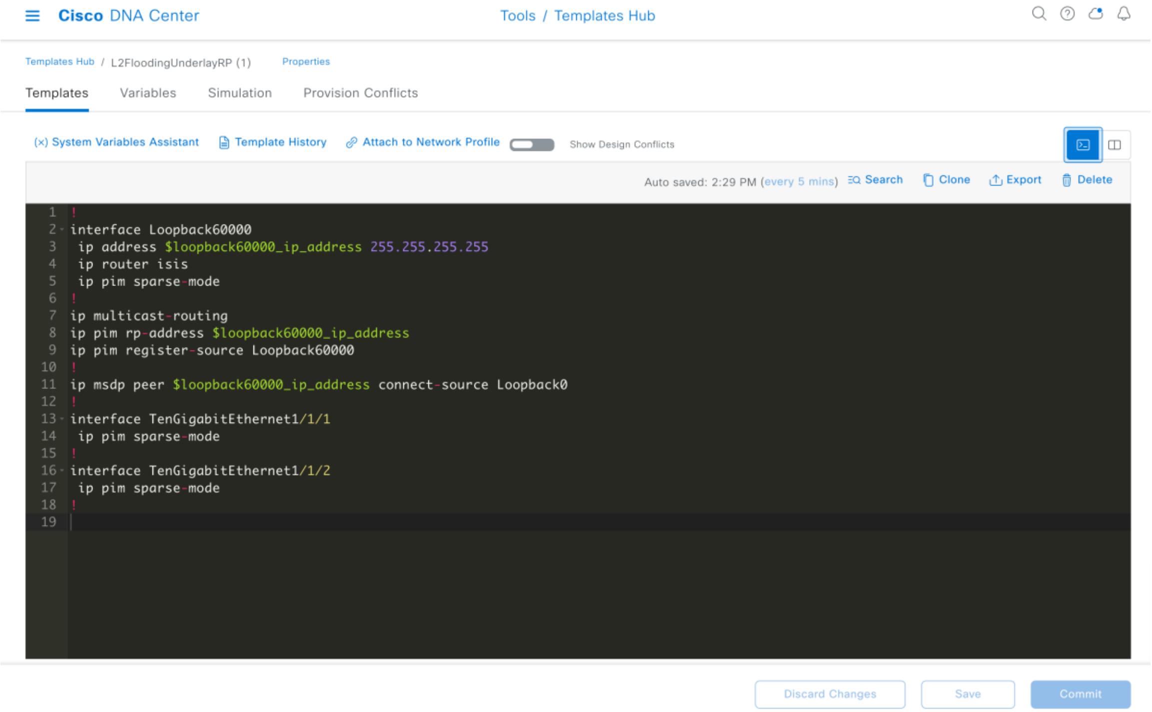

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

Enabling the feature for the subnet.

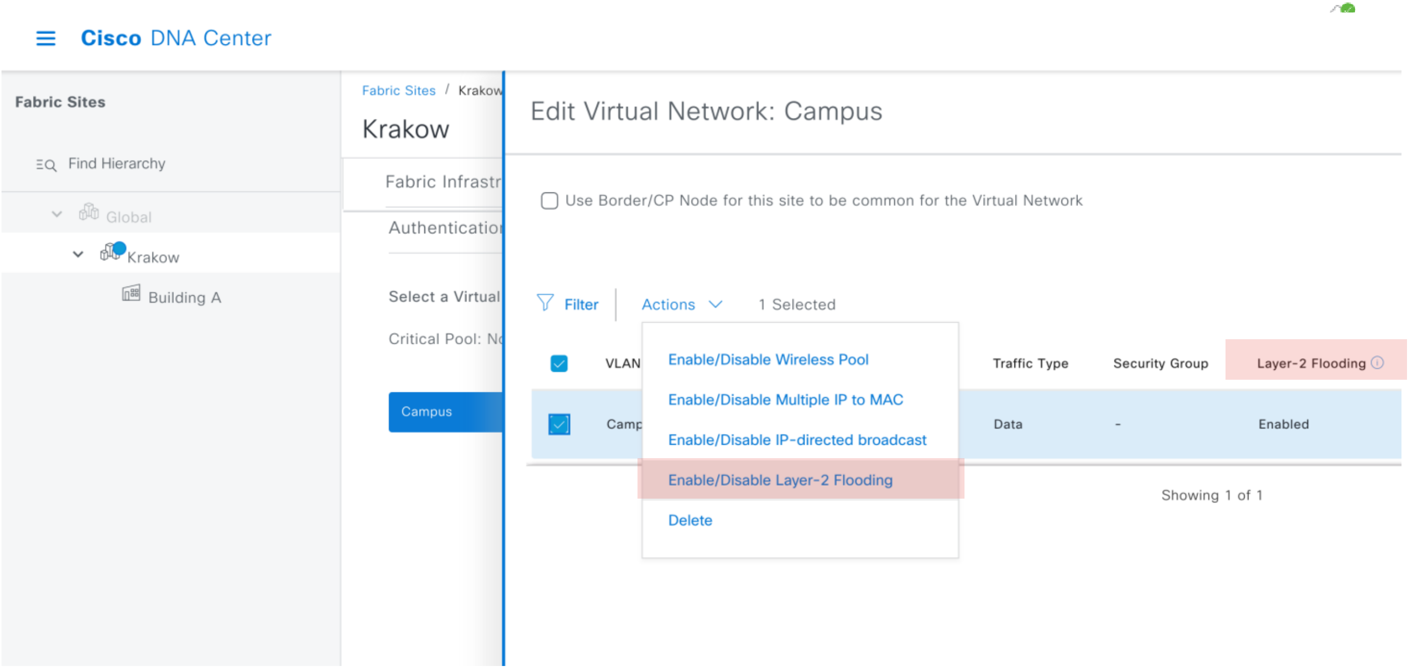

INFO

No need to re-provision fabric devices.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

What happens in background?

LISP configuration is updated on all of the fabric devices with hosts attached to the subnet where L2 flooding enabled (including border nodes if Layer 2 border handoff enabled)

```
Edge1#show run | sec lisp
(snipped)
 !
 instance-id 8188
  remote-rloc-probe on-route-change
  service ethernet
   eid-table vlan 200
   broadcast-underlay 239.0.17.1
   flood arp-nd
   flood unknown-unicast
   database-mapping mac locator-set rloc_671c2347-7a35-4597-a27e-d3c4bdc72ba1
   exit-service-ethernet
  !
  exit-instance-id
 !
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

A closer look. Initial status.

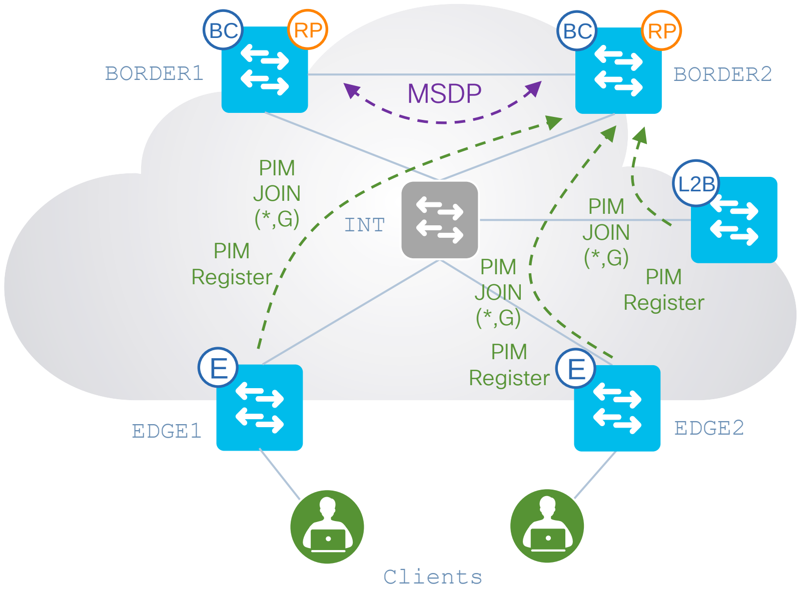

1. Rendezvous Points (RPs) are present in the underlay (AnycastRP).
2. PIM-SM is enabled in underlay between fabric devices.
3. MSDP is configured for exchanging information about sources.
4. L2 flooding is enabled for the subnet in Cisco DNA Center.
5. Once L2 flooding enabled, Edge Nodes (and L2 Borders if configured) send PIM Join (*,G) to RP declaring their own Loopback0 as an IGMP reporter.
6. Same time Edge Nodes (and L2 Borders if configured) send PIM Register message to RP registering the source (itself).

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

A closer look. Switching to SPT.

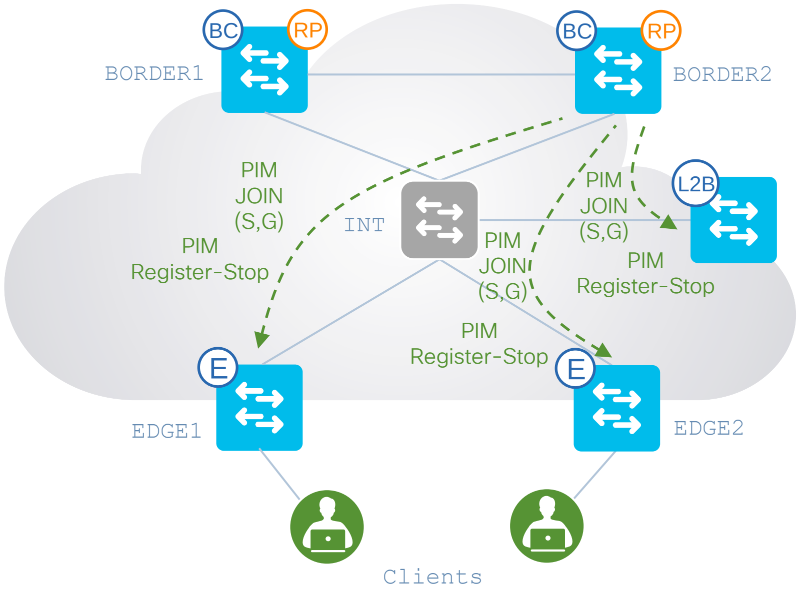

1. RP replies with PIM Register-stop message.
2. Then the PIM Join is sent (S,G) down towards Edge Nodes (and L2 Borders if configured).
3. The OIL is updated accordingly and SPT is built.
4. In the last step RP prunes itself (if it's not a L2 Border).

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

A closer look. Final stage.

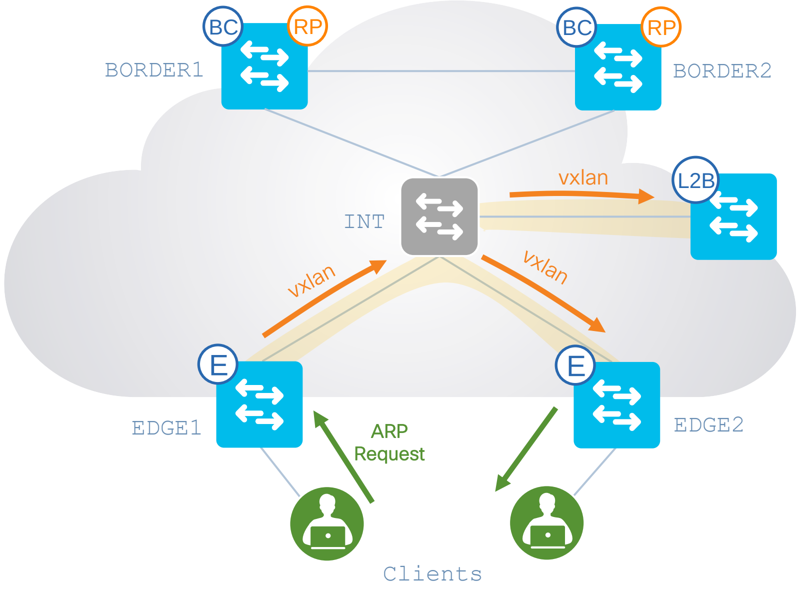

1. SPT is pre-built for a given IP pool even if there is no actual BUM traffic.

2. Once the BUM traffic (e.g. ARP Request) hits one device, it is intercepted and sent over via a dedicated multicast group in the underlay.

3. Traffic is encapsulated in VXLAN and the multicast group assigned to the given IP subnet is used as a destination address.

4. The underlay device (INT) is responsible for replicating traffic as needed (regular multicast operations).

5. All Edge Nodes and L2 Borders receive traffic originated by EDGE1 which can be forwarded to clients.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

Let’s get some captures!

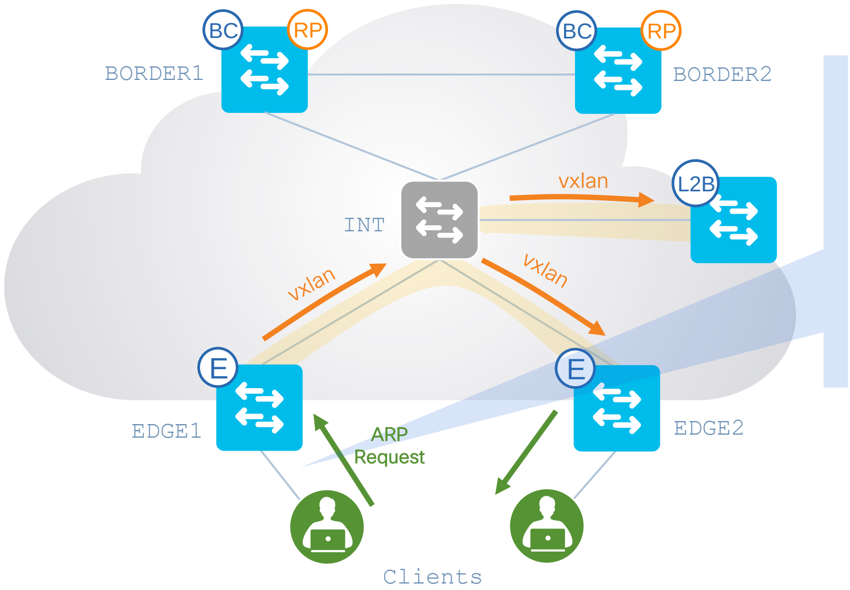

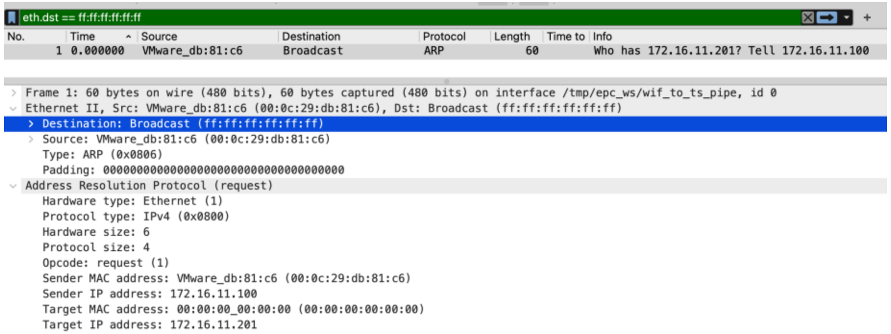

```
eth.dst == ff:ff:ff:ff:ff:ff
No.  Time      Source          Destination  Protocol  Length  Time to  Info
  1  0.000000  VMware_db:81:c6  Broadcast    ARP       60              Who has 172.16.11.201? Tell 172.16.11.100

> Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface /tmp/epc_ws/wif_to_ts_pipe, id 0
v Ethernet II, Src: VMware_db:81:c6 (00:0c:29:db:81:c6), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
  > Destination: Broadcast (ff:ff:ff:ff:ff:ff)
  > Source: VMware_db:81:c6 (00:0c:29:db:81:c6)
    Type: ARP (0x0806)
    Padding: 000000000000000000000000000000000000
v Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: VMware_db:81:c6 (00:0c:29:db:81:c6)
    Sender IP address: 172.16.11.100
    Target MAC address: 00:00:00_00:00:00 (00:00:00:00:00:00)
    Target IP address: 172.16.11.201
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

Let's get some captures!

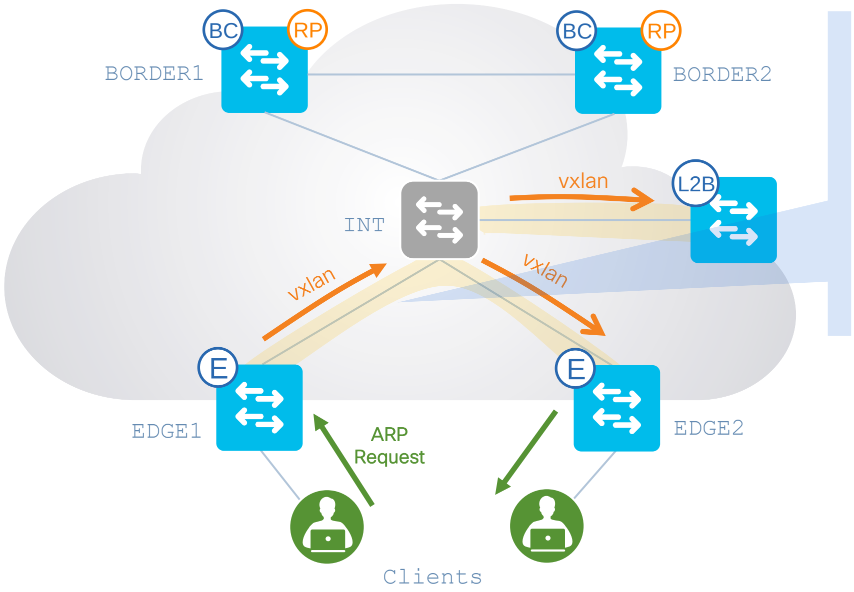

eth.dst == ff:ff:ff:ff:ff:ff

No.	Time	Source	Destination	Protocol	Length	Time to	Info
25	2.765816	VMware_db:81:c6	Broadcast	ARP	110	253	Who has 172.16.11.201? Tell 172.16.11.100

```
> Frame 25: 110 bytes on wire (880 bits), 110 bytes captured (880 bits) on interface /tmp/epc_ws/wif_to_ts_pipe, id 0
v Ethernet II, Src: Cisco_38:b3:56 (00:42:5a:38:b3:56), Dst: IPv4mcast_11:01 (01:00:5e:00:11:01)
    > Destination: IPv4mcast_11:01 (01:00:5e:00:11:01)
    > Source: Cisco_38:b3:56 (00:42:5a:38:b3:56)
      Type: IPv4 (0x0800)
> Internet Protocol Version 4, Src: 192.168.10.3, Dst: 239.0.17.1
> User Datagram Protocol, Src Port: 65485, Dst Port: 4789
v Virtual eXtensible Local Area Network
    > Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
      Group Policy ID: 0
      VXLAN Network Identifier (VNI): 8188
      Reserved: 0
> Ethernet II, Src: VMware_db:81:c6 (00:0c:29:db:81:c6), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
> Address Resolution Protocol (request)
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 2 flooding

Let's get some captures!

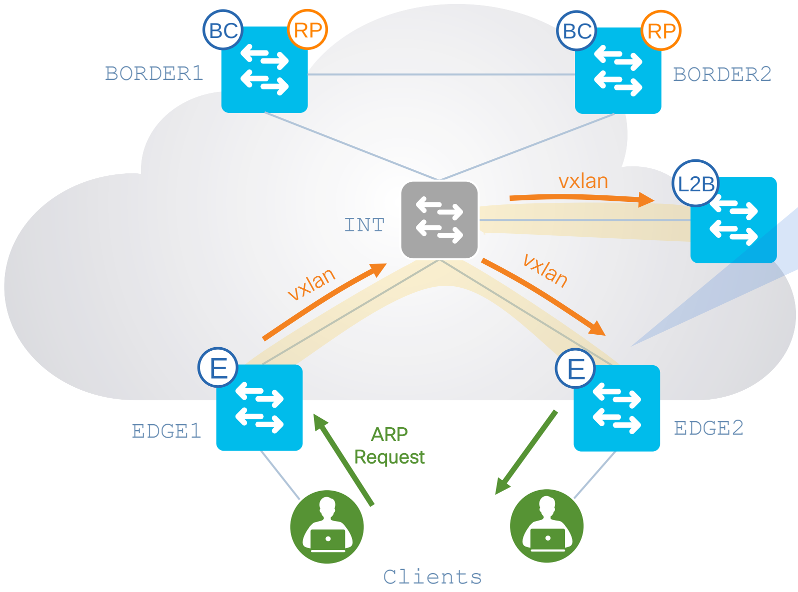

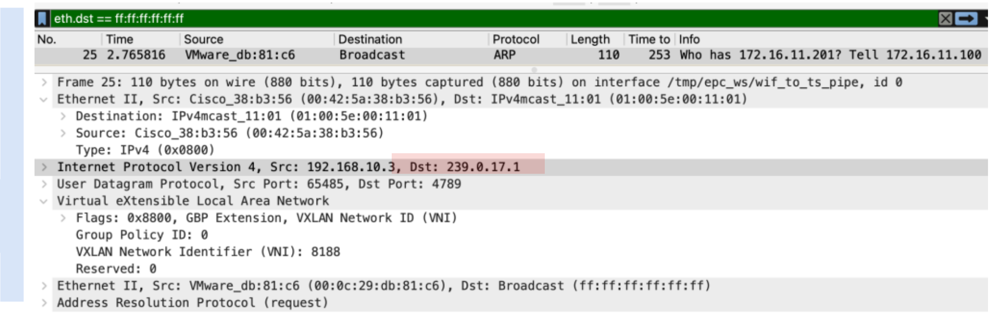

```
eth.dst == ff:ff:ff:ff:ff:ff
No.  Time      Source          Destination  Protocol  Length  Time to  Info
 25  2.765816  VMware_db:81:c6  Broadcast    ARP       110     253      Who has 172.16.11.201? Tell 172.16.11.100

> Frame 25: 110 bytes on wire (880 bits), 110 bytes captured (880 bits) on interface /tmp/epc_ws/wif_to_ts_pipe, id 0
v Ethernet II, Src: Cisco_38:b3:56 (00:42:5a:38:b3:56), Dst: IPv4mcast_11:01 (01:00:5e:00:11:01)
  > Destination: IPv4mcast_11:01 (01:00:5e:00:11:01)
  > Source: Cisco_38:b3:56 (00:42:5a:38:b3:56)
    Type: IPv4 (0x0800)
> Internet Protocol Version 4, Src: 192.168.10.3, Dst: 239.0.17.1
> User Datagram Protocol, Src Port: 65485, Dst Port: 4789
v Virtual eXtensible Local Area Network
  > Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 8188
    Reserved: 0
> Ethernet II, Src: VMware_db:81:c6 (00:0c:29:db:81:c6), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
> Address Resolution Protocol (request)
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

L2 Flooding enhancement

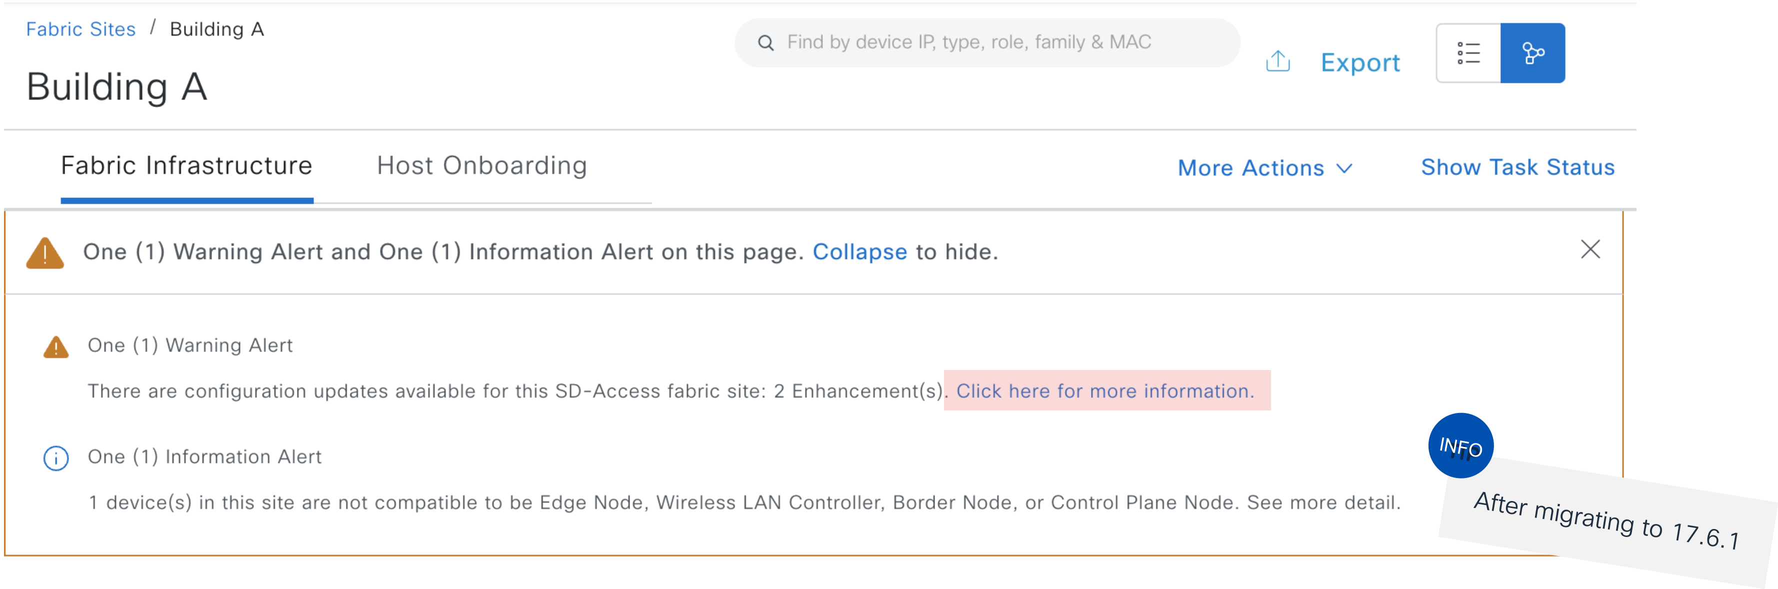

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

L2 Flooding enhancements

What is actually means ?

Fabric Configuration Updates

2 Enhancement(s)

⌄ ⚠ Enhancements (2) Apply All

Layer 2 Flooding Access Control List

This release provides an updated Layer 2 Flooding behavior in Cisco SD-Access for IOS XE-based switches by enabling an Access Control List (ACL) on the Layer 2 LISP sub-interfaces. This ACL is provisioned if Layer 2 Flooding is enabled, and it is applied on all devices operating with the Edge Node function and Border Nodes that have a Layer 2 Handoff. To enable this ACL support, switches operating in fabric roles must be upgraded to IOS XE 17.6.x or later. Once these changes are incorporated, you will not be able to add new switches in fabric roles which do not meet this minimum IOS XE version requirement.

Layer 2 Flooding Update

When Layer 2 Flooding is enabled for a given address pool, broadcast, link-local multicast, and ARP frames are flooded. This release provides an updated Layer 2 Flooding behavior in Cisco SD-Access for IOS XE-based switches to further enable support for 'silent hosts.' This is provided by flooding unknown-unicast and IPv6 Neighbor Discovery frames. To enable this updated flooding behavior, switches operating in fabric roles must be upgraded to IOS XE 16.12.2 or later. Once these changes are incorporated, you will not be able to add new switches in fabric roles which do not meet this minimum IOS XE version requirement.

```
Edge1#show run
(snipped)
!
interface L2LISP0
 ip access-group SDA-FABRIC-LISP in
 ip access-group SDA-FABRIC-LISP out
!
(snipped)
ip access-list extended SDA-FABRIC-LISP
 10 deny ip any host 224.0.0.22
 20 deny ip any host 224.0.0.13
 30 deny ip any host 224.0.0.1
 40 permit ip any any
!
(snipped)
```

INFO

224.0.22 - IGMP
224.0.13 - All PIM Routers
224.0.0.1 - All Systems on this Subnet

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multicast enhancements

New multicast features introduced in DNAC 2.3.5

- UI Workflow that allows continuing using multicast feature per Virtual Network (VN).
A user will now be able to configure both ASM and SSM together.
It will be presented with SSM configuration first, followed by ASM configuration.

- The ability to map different ASM groups to different RP addresses per L3VN.
UI will provide the functionality to configure any number of groups to external RP mappings.

- ASM configuration screens are changed to configure external or internal RP
per ASM group range(s). The number of maximum internal RPs continues to be 2.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multicast in SDA Transit

Multicast in SDA Transit

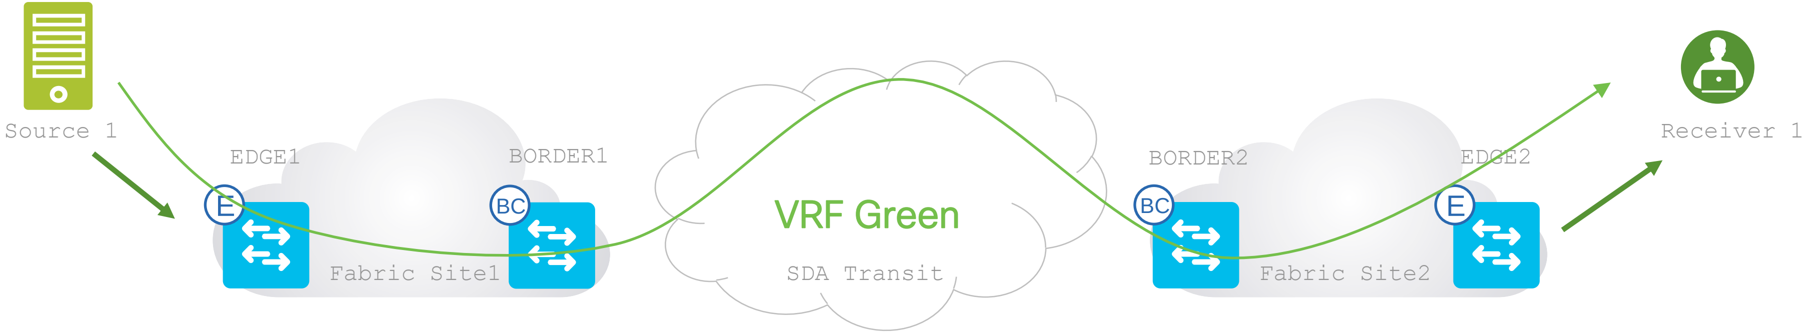

Fabric Site #1	SDA Transit	Fabric Site #2
vrf GREEN 192.168.10.1 -> 224.100.100.100 Underlay: 192.168.10.1 -> 232.0.0.100	Overlay: 192.168.10.1 -> 224.100.100.100 Underlay : 192.168.10.1 -> 232.0.0.100	vrf GREEN 192.168.10.1 -> 224.100.100.100 Underlay : 192.168.10.1 -> 232.0.0.100
`ip pim lisp core-group-range 232.0.0.1 1000`	`ip pim lisp core-group-range 232.0.0.1 1000`	`ip pim lisp core-group-range 232.0.0.1 1000`

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multicast in SDA Transit

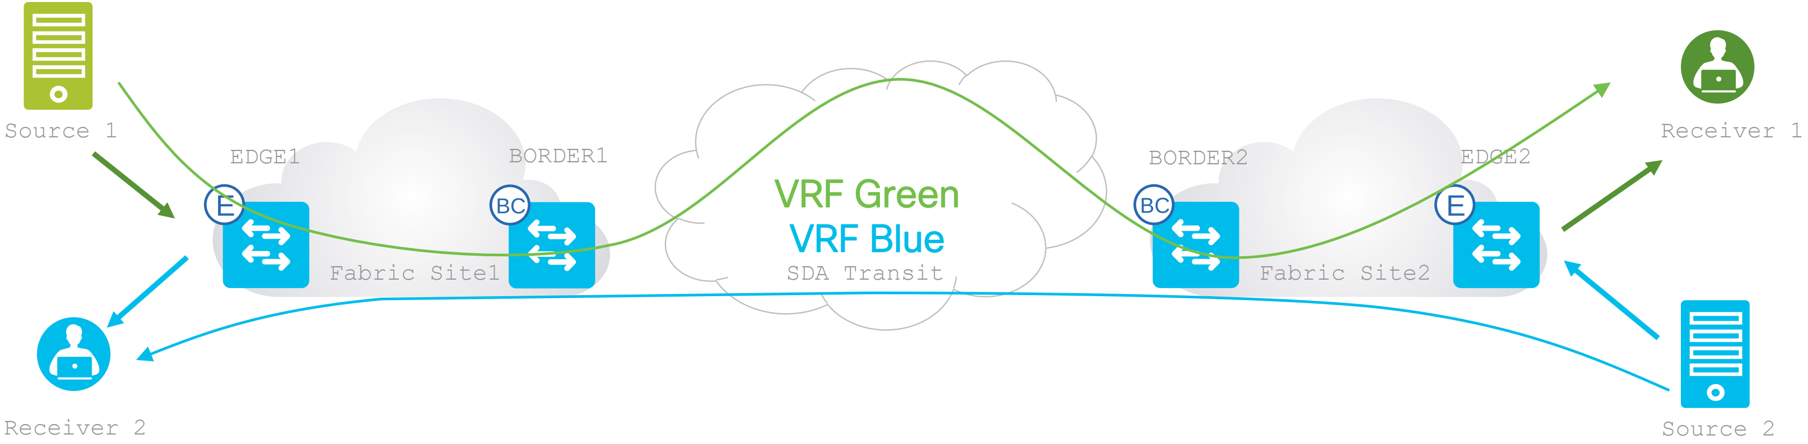

In case of introduction of a new source and receiver spread across fabric sites in different/same VRF, and using the same overlay multicast group, a multicast packet loop might be formed between the borders due to the derivation of the same underlay group for the overlay multicast

- For VRF Green, Border1 would decapsulate the traffic from Edge1 and encapsulate to Border2
- For VRF Blue, Border2 would decapsulate the traffic from Edge2 and encapsulate to Border1

Since Border1 and Border2 derives the same multicast group, results in both pointing to each other in multicast routing table for decapsulation and encapsulation.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multicast in SDA Transit – support in 2.3.5

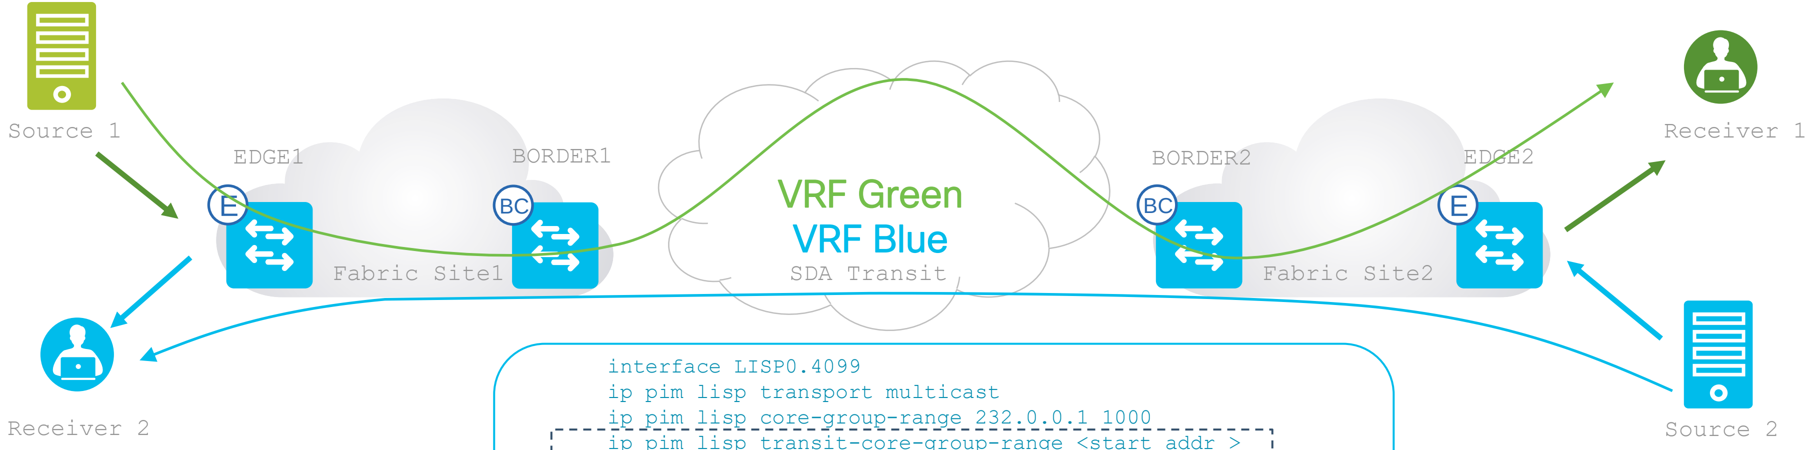


```
interface LISP0.4099
ip pim lisp transport multicast
ip pim lisp core-group-range 232.0.0.1 1000
ip pim lisp transit-core-group-range <start_addr >
<num_of_groups>

interface LISP0.4100
ip pim lisp transport multicast
ip pim lisp core-group-range 232.0.0.1 1000
ip pim lisp transit-core-group-range <start_addr >
<num_of_groups>
```

In DNAC 2.3.5 a new multicast group address is introduced that is only used for forwarding between the border nodes in SDA transit

The multicast address for the underlay within a fabric site doesn’t overlap with the address used for forwarding in SDA transit

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multicast in SDA Transit

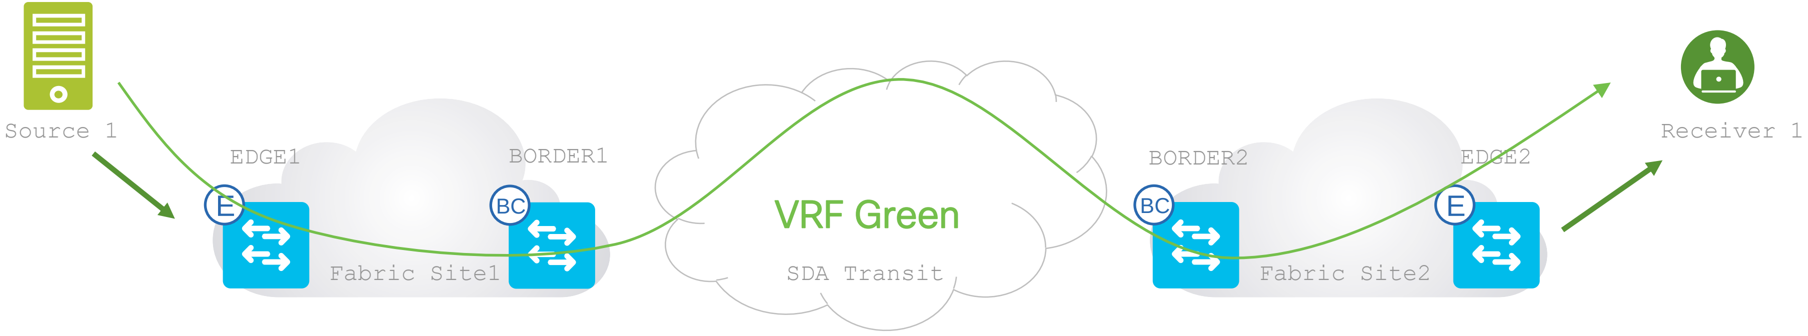

Fabric Site #1	SDA Transit	Fabric Site #2
vrf GREEN 192.168.10.1 -> 224.100.100.100 Underlay: 192.168.10.1 -> 232.0.0.100	Overlay. 192.168.10.1 -> 224.100.100.100 Underlay : 192.168.10.1 -> 232.11.0.100	vrf GREEN 192.168.10.1 -> 224.100.100.100 Underlay : 192.168.10.1 -> 232.0.0.100
`ip pim lisp core-group-range 232.0.0.1 1000`	`ip pim lisp transit-core-group-range 232.11.0.1 1000`	`ip pim lisp core-group-range 232.0.0.1 1000`

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

DEMO #2:

Multicast in SDA Transit – Demo

https://youtu.be/SgxMv7DTK7c

LAB topology

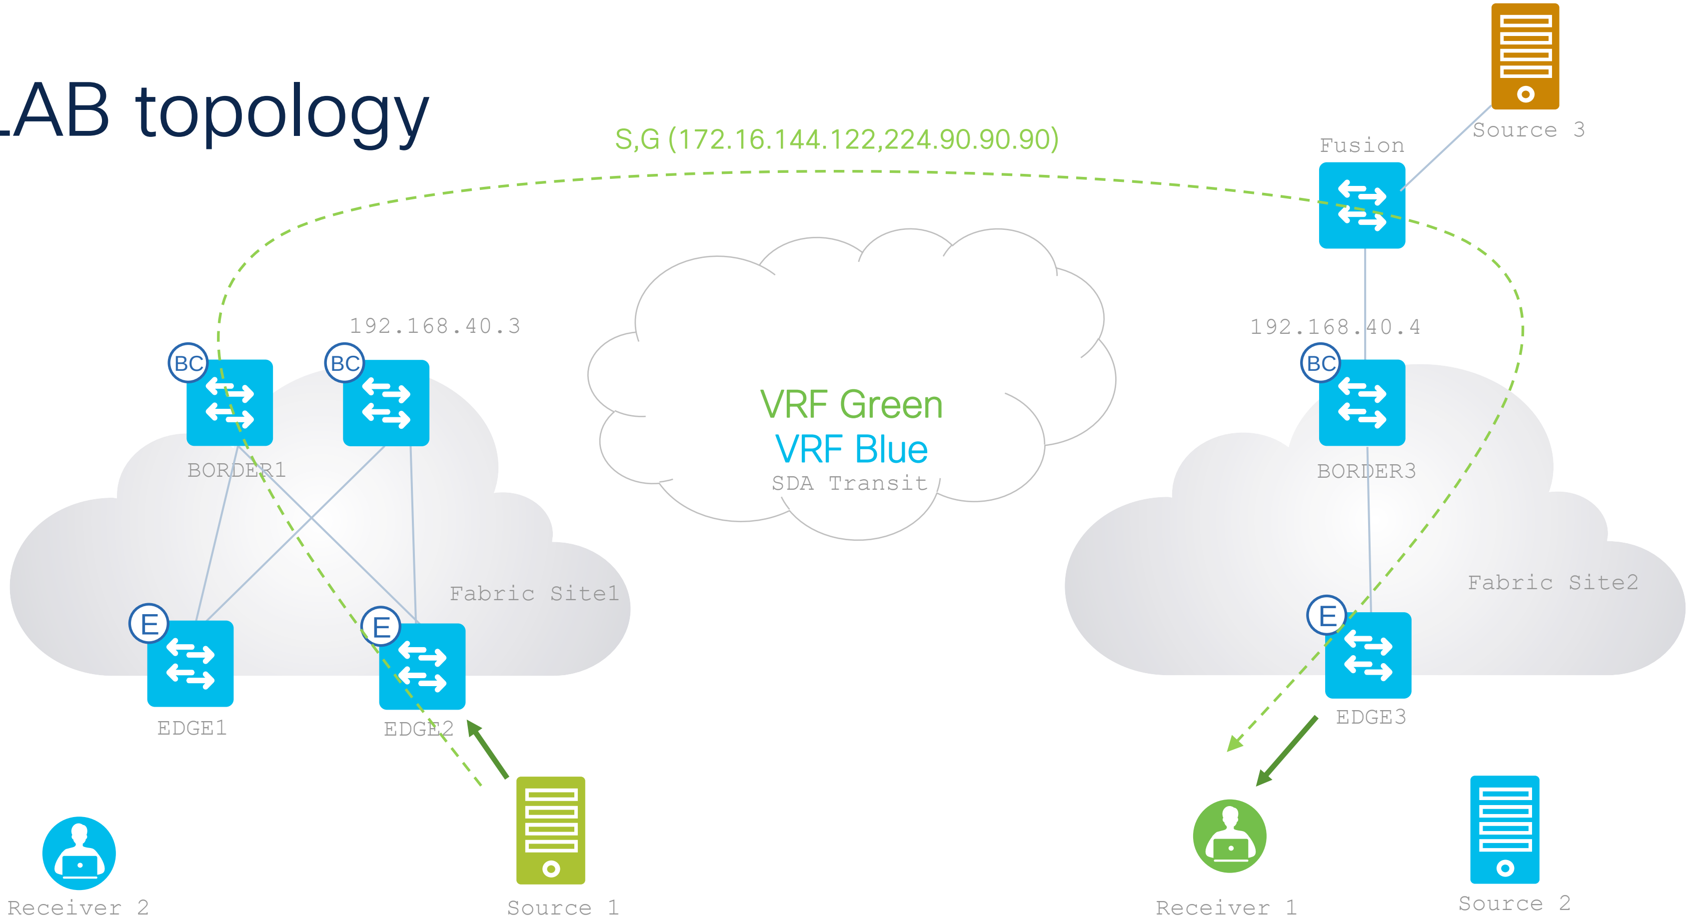

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco DNA Center

Welcome, admin

Explore

Some of your license compliance requirements have not been met. Learn more.

Learn about new capabilities in this release on the Cisco DNA Center YouTube Channel.

Stay up to date with your network and Cisco DNA Center through our insight email

Receive announcements, network highlights, weekly snapshots, and executive summaries all neatly packaged in a single email.

Insights

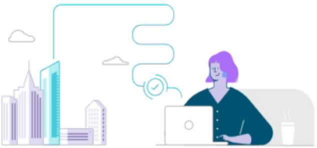

Assurance Summary

Health

Healthy as of Jan 18, 2023 5:20 PM

100%	--%	100%
Network Devices	Wireless Clients	Wired Clients

View Details

Critical Issues

Last 24 Hours

0	3
P1	P2

View Details

Trends and Insights

Last 30 Days

--	--
AP Performance Advisories	Trend Deviations

View Details

Network Snapshot

Sites

As of Jan 18, 2023 7:10 PM

8

DNS Servers : 1
NTP Servers : 0

Network Devices

As of Jan 18, 2023 7:10 PM

8

Unclaimed: 0
Unprovisioned: 1
Unreachable: 0

Application QoS Policies

As of Jan 18, 2023 7:12 PM

0

Successful Deploys: 0
Errored Deploys: 0
Stale Policies: 0

Design Considerations

RP(s) already exists in the network

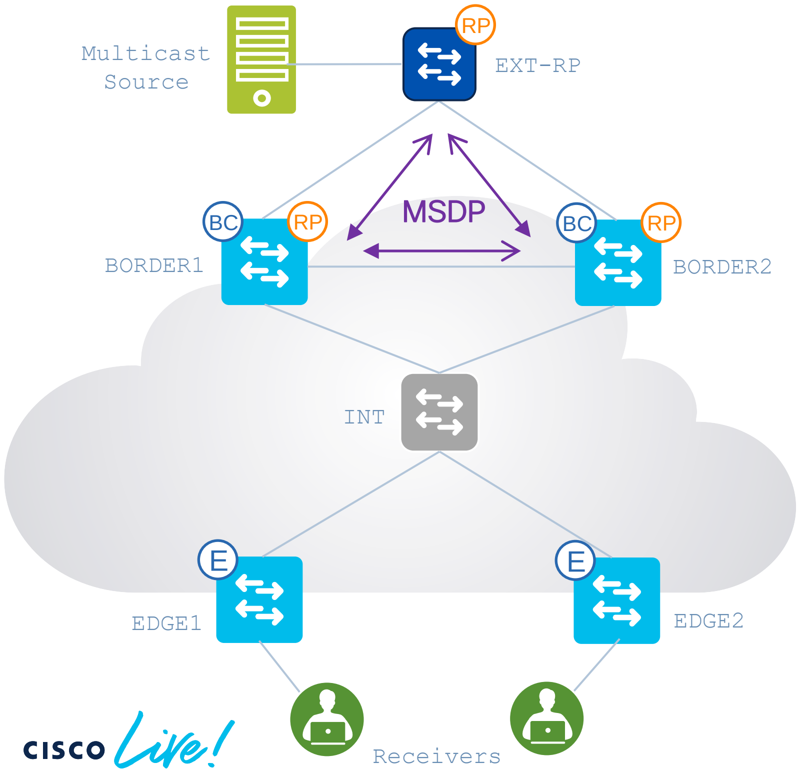

INFO

Support for the total number of external RP addresses beyond current maximum of 2

It is quite common that RP already exists in the network and reconfiguring all existing devices to use fabric RP instead is not a good practice (not optimal).

In that case MSDP can be used to transfer information about multicast source towards fabric RPs.

It requires manual configuration (it can be automated via Templated Editor in DNA-C).

MSDP needs to be configured per Virtual Network (or GRT-to-VN if EXT-RP belongs to GRT).

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

RP(s) already exists in the network

MSDP - configuration example.

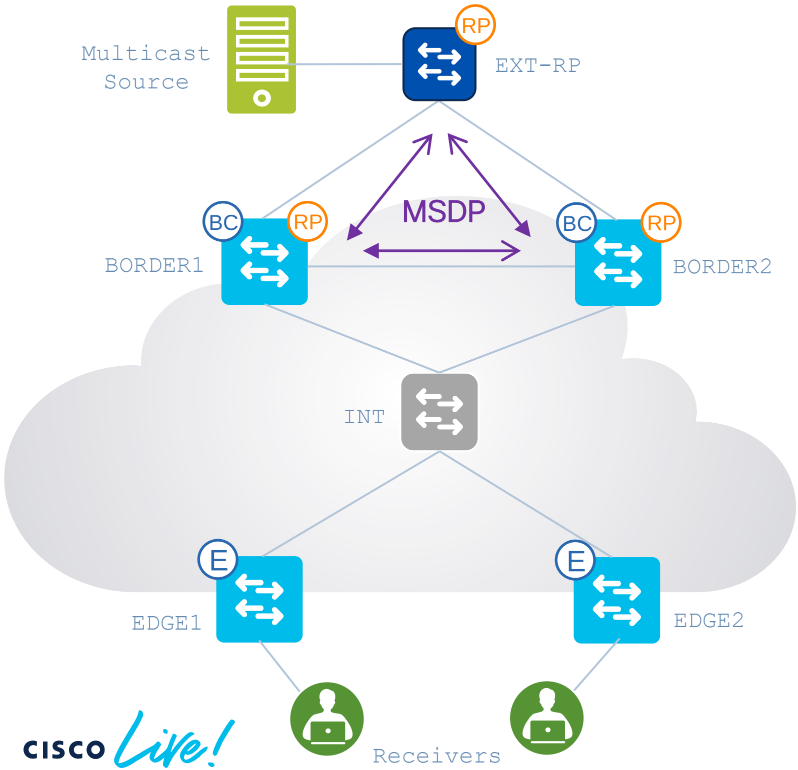

EXT-RP

```
interface LoopbackX
 ip address 1.1.1.1 255.255.255.255
 ip pim sparse-mode
!
ip msdp peer <Border1-MSDP-Loopback-IP> connect-source Loopback0
ip msdp peer <Border2-MSDP-Loopback-IP> connect-source Loopback0
ip msdp cache-sa-state
ip msdp originator-id LoopbackX
```

BORDER1 (FABRIC-RP)

```
ip msdp vrf <VN-name> peer 1.1.1.1 connect-source <Border1-MSDP-Loopback-IP>
```

BORDER2 (FABRIC-RP)

```
ip msdp vrf <VN-name> peer 1.1.1.1 connect-source <Border2-MSDP-Loopback-IP>
```

Take care of necessary routing and leaking

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Clients located in multiple Virtual Networks

A design challenge.

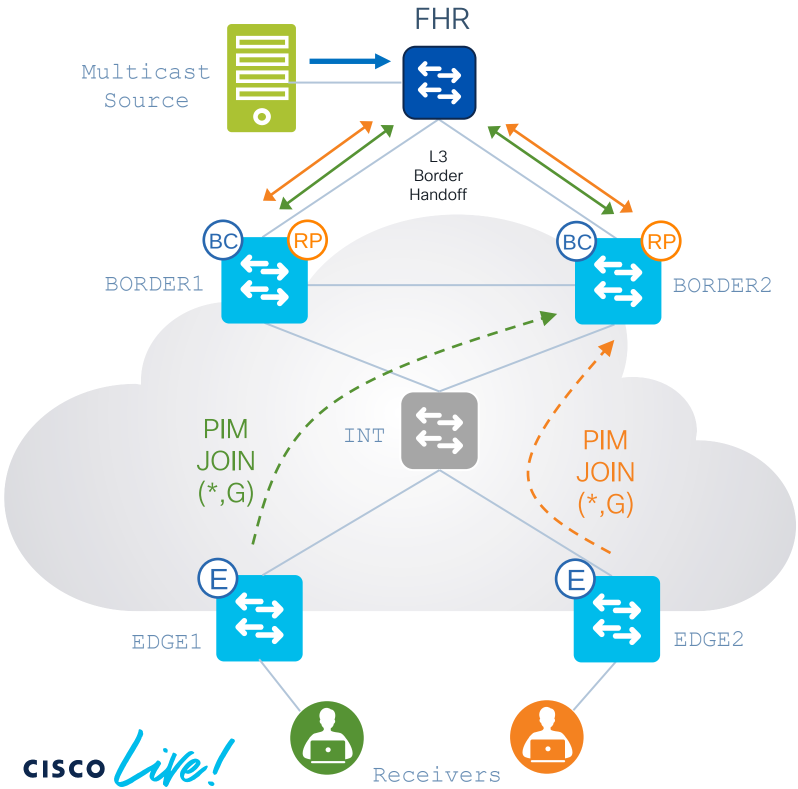

VN1 - VN2 - VRF3

1. Client1 located in VN1, Client2 located in VN2.
2. Border Nodes are configured as RPs for both VNs.
3. Layer 3 Border Handoff (eBGP) is configured towards FHR.
4. Both clients request multicast traffic from the same group.
5. Edge Nodes send PIM Join message towards RP in respective Virtual Network.

How to make FHR to register multicast source in two different Virtual Networks (VRFs) ?

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multicast VPN Extranet feature

A possible solution.

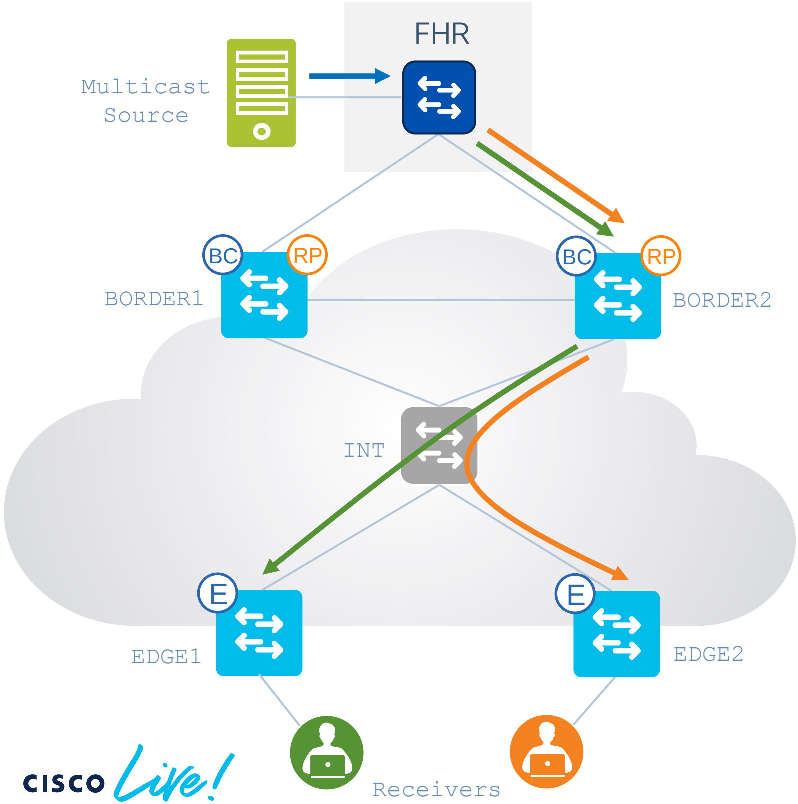

VN1 - VN2 - VRF3

It allows to distribute IP multicast traffic originated from one VRF to other VRFs (aka multicast leaking).

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Summary / Call to action

IP multicast in SD-Access

Key use cases.

Head-end replication / Native multicast

- CCTV
- Hospitality TV
- University Campus Multimedia
- Medical devices

Layer 2 flooding

- Migration (with use of L2 Border)
- Silent hosts
- Wake on LAN (WoL)*
- BMS

* SDA Transit not supported in DNAC 2.3.5.0

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Key takeaways

- Head-end replication can be used in both ASM and SSM mode in the overlay.
- Native multicast can be used in both ASM and SSM mode in the overlay.
- Native multicast requires SSM to be configured and enabled in the underlay.
- L2 flooding requires ASM to be configured and enabled in the underlay.
- ASM/SSM in the underlay can be automated via LAN Auto.
- Multicast over SDA Transit fully supported in DNAC 2.3.5 (*)

(*) Multicast over Pub/Sub SD-Access Transit is supported with IOS XE 17.10 and Cisco DNA Center 2.3.5.x. Multicast routing over LISP/BGP SD-Access Transit is unsupported.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

The bridge to possible

Thank you

cisco Live!

Let’s go