

CISCO *Live!*

Let's go

#CiscoLiveAPJC



The bridge to possible

# A Day in the Life of a Packet

VXLAN BGP EVPN Fabrics

Lukas Krattiger, Cisco Fellow

@CCIE21921

BRKDCN-3966

CISCO *Live!*

#CiscoLiveAPJC



# Cisco Webex App


## Questions?

Use Cisco Webex App to chat with the speaker after the session

## How

- 1 Find this session in the Cisco Live Mobile App
- 2 Click “Join the Discussion”
- 3 Install the Webex App or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until December 22, 2023.



<https://cislive.ciscoevents.com/cislivebot/#BRKDCN-3966>

# Introduction

- We are going to cover all different kind of Packet Walks
- Route, Bridge, BUM and Silent Host Discovery
- A brief intro to VXLAN with EVPN
  - Sorry, not a VXLAN or VXLAN EVPN Intro Session

# Agenda

- Introduction to VXLAN EVPN
- Layer-3 Packet Walk
- Layer-2 Packet Walk
- BUM Packet Walk
- Silent Host Discovery
- Conclusion

# Introduction

# What is VXLAN?

- Standards based Encapsulation
- RFC 7348
- Uses UDP-Encapsulation
- Transport Independent
- Layer-3 Transport (Underlay)
- Flexible Namespace
- 24-bit field (VNID) provides ~16M unique identifier
- Allows Segmentations


# What is EVPN?

- Standards based Control-Plane
- RFC 8365 (and RFC 7432)
- Uses Multiprotocol BGP
- Uses Various Data-Planes
- VXLAN (EVPN-Overlay), MPLS, Provider Backbone (PBB)
- Many Use-Cases Covered
- Bridging, MAC Mobility, First-Hop & Prefix Routing, Multi-Tenancy (VPN)




# It all starts with a Network Device

## The Dating Network - When Control- meets Data-Plane




# It all starts with a Network Device

## The Dating Network - When Control- meets Data-Plane




# It all starts with a Network Device

## The Dating Network - When Control- meets Data-Plane



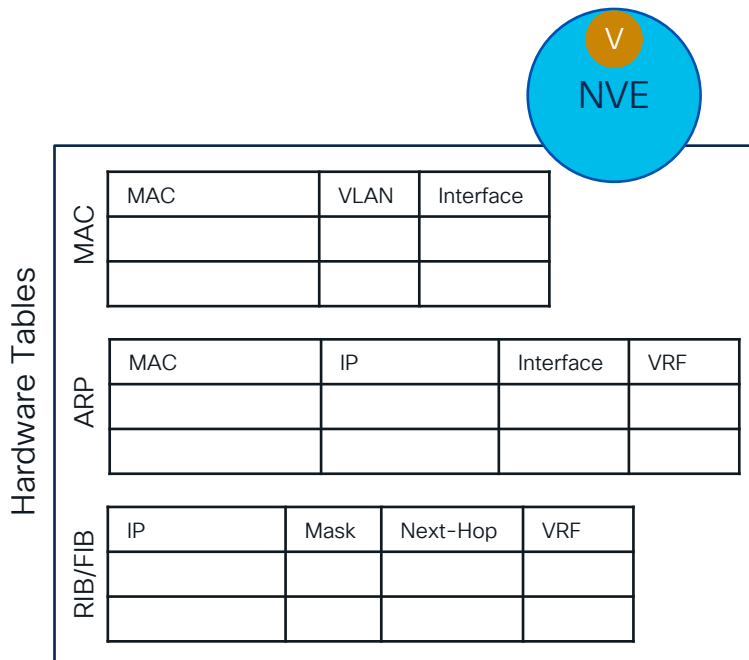
# Making the Network Device an NVE

## The Dating Network - When Control- meets Data-Plane



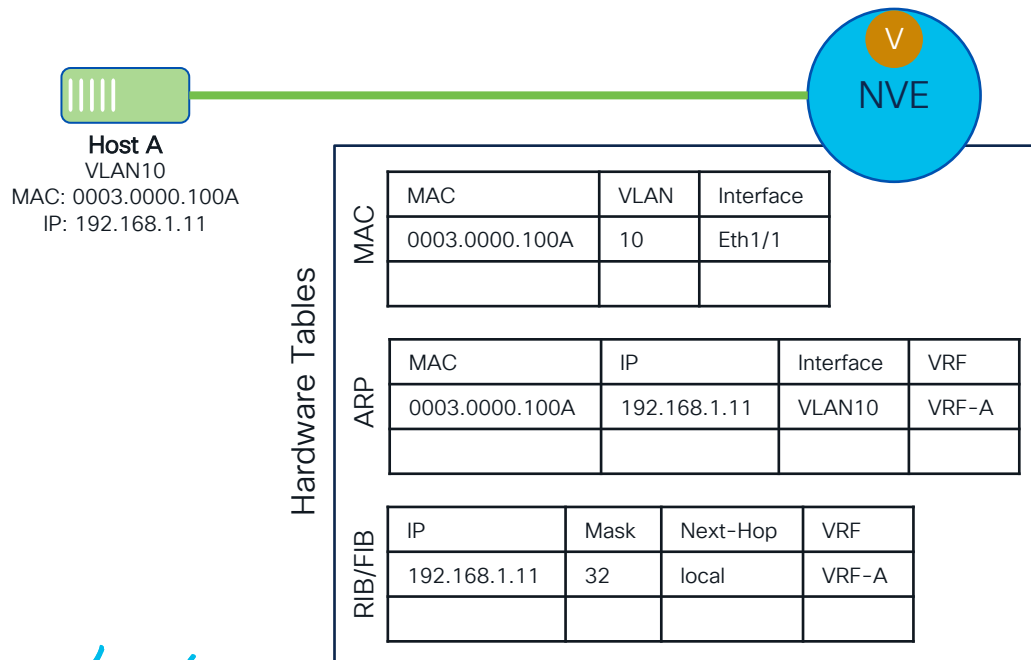
# The NVE and Some Important Table

## The Dating Network - When Control- meets Data-Plane



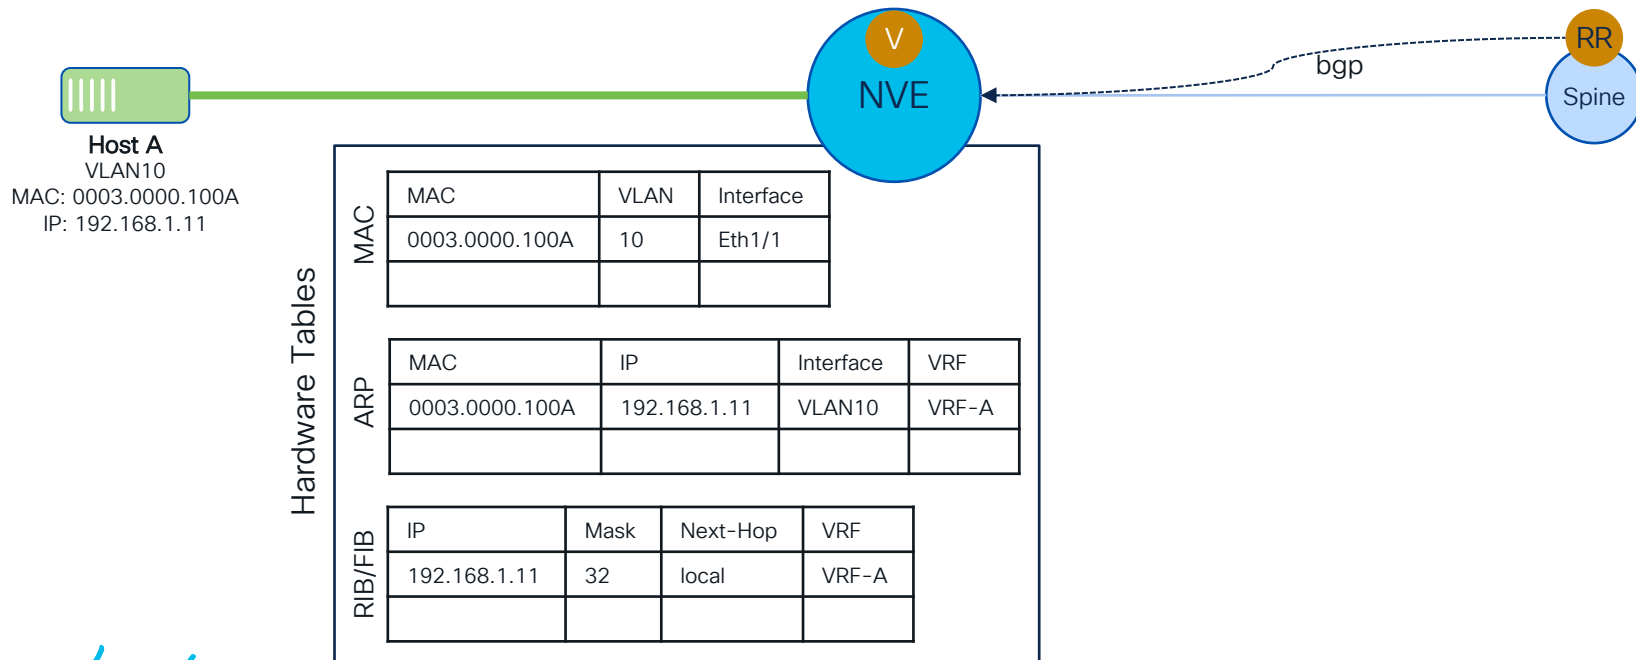
# Local Learning on the NVE

## The Dating Network - When Control- meets Data-Plane



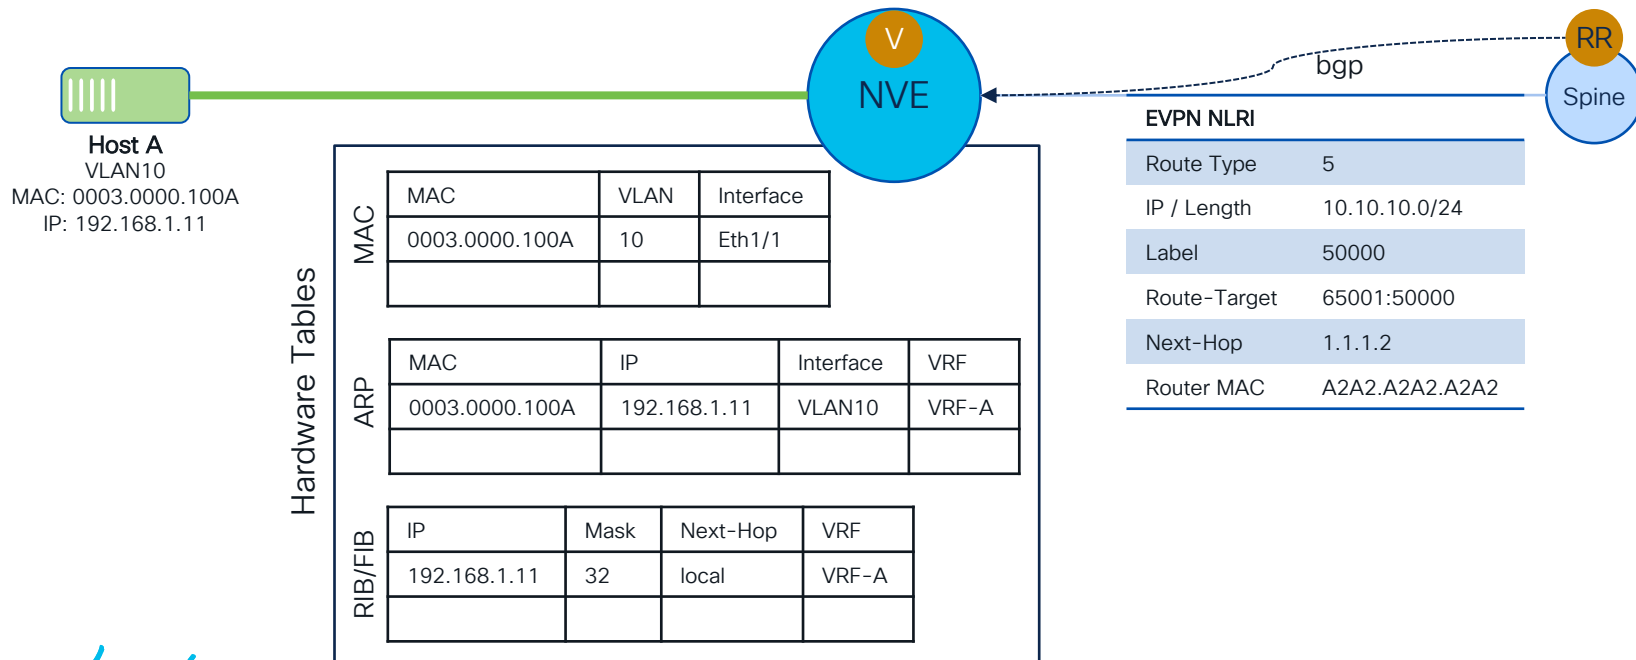
# Remote Learning from other NVE

## The Dating Network - When Control- meets Data-Plane



# Remote Learning from other NVE

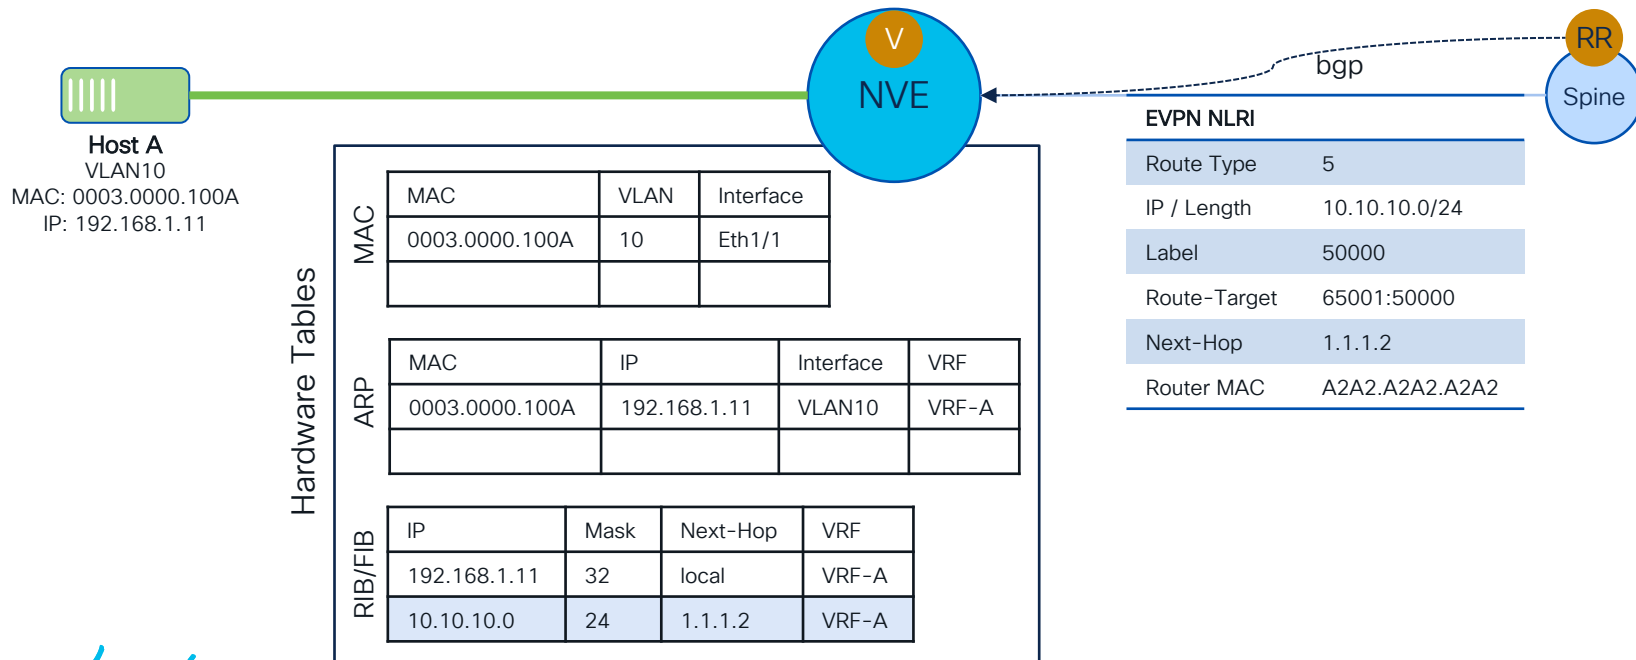
## The Dating Network - When Control- meets Data-Plane





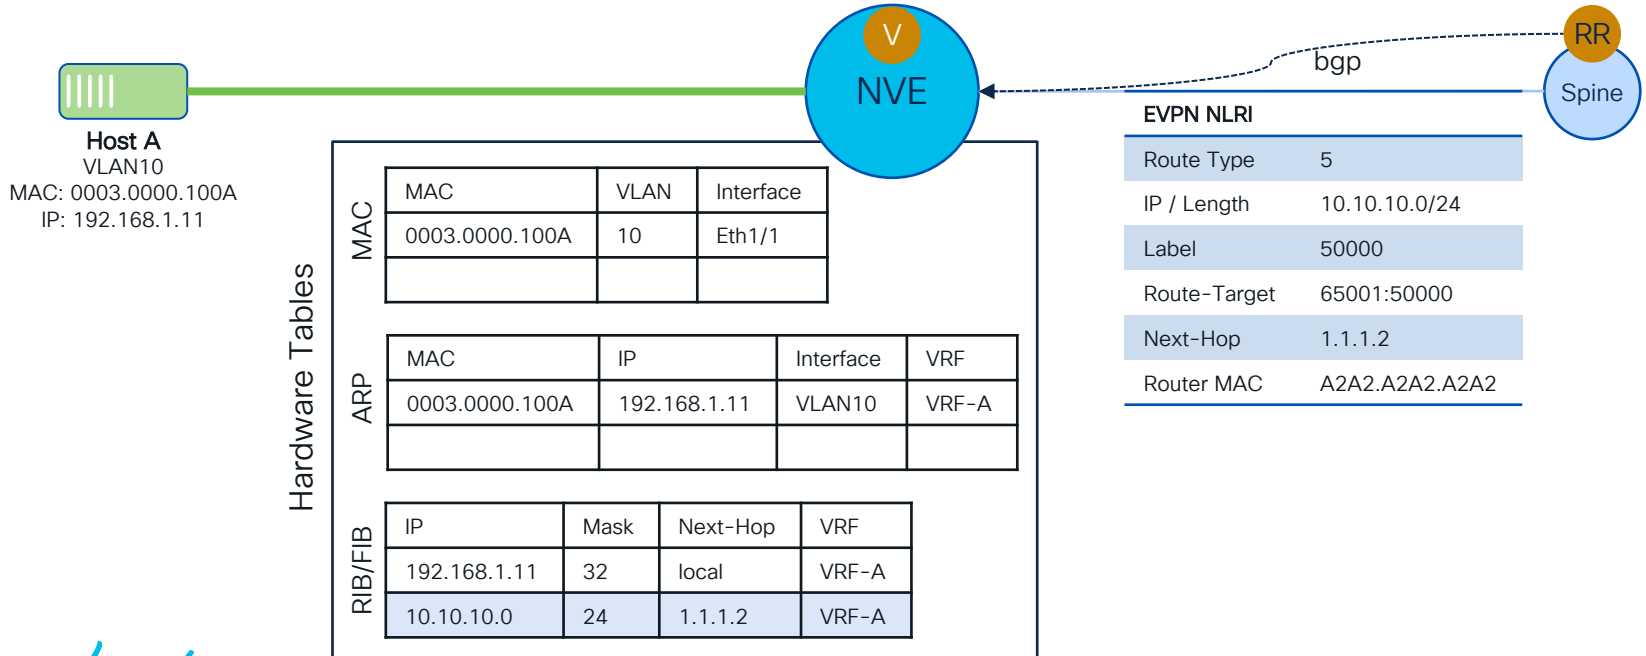

# Remote Learning from other NVE

## The Dating Network - When Control- meets Data-Plane



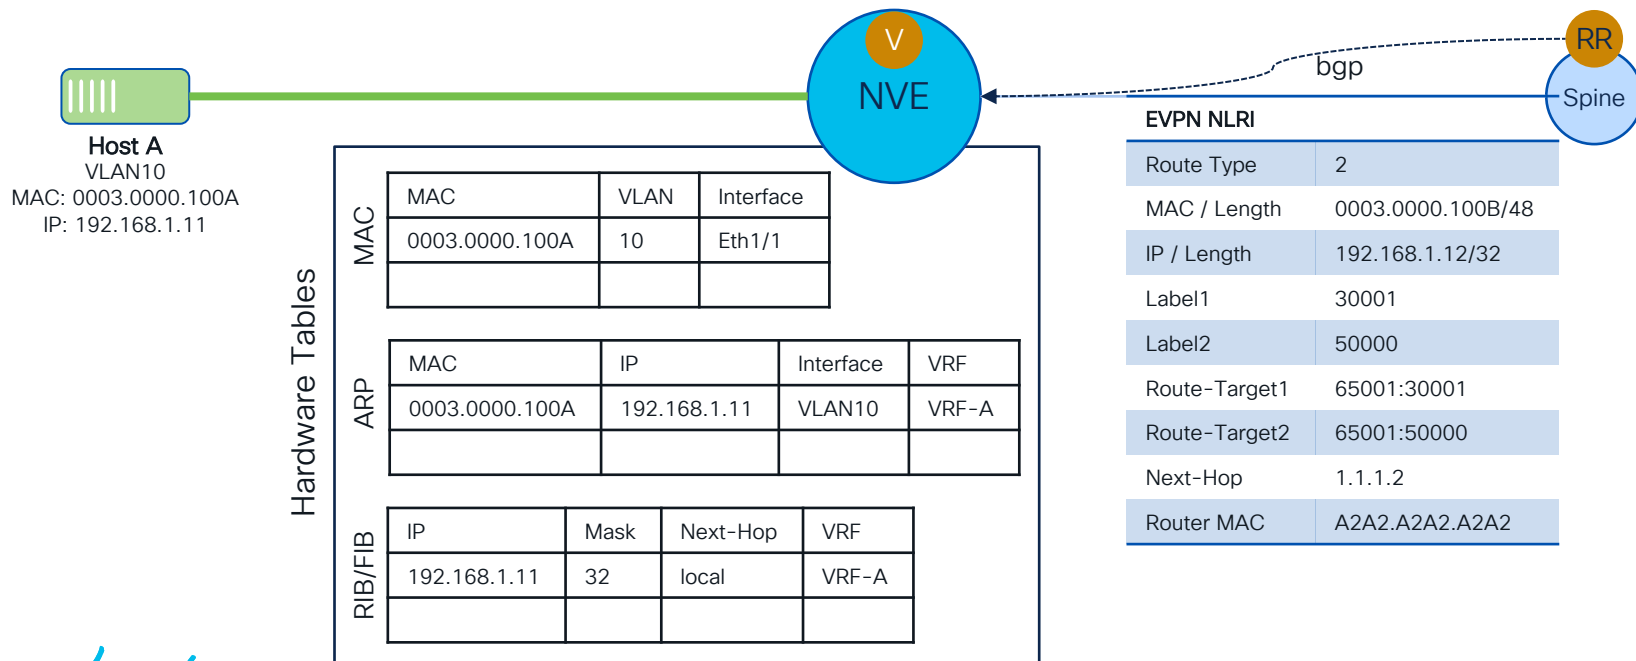
# Routing between NVE (based on VXLAN EVPN)

## The Dating Network - When Control- meets Data-Plane



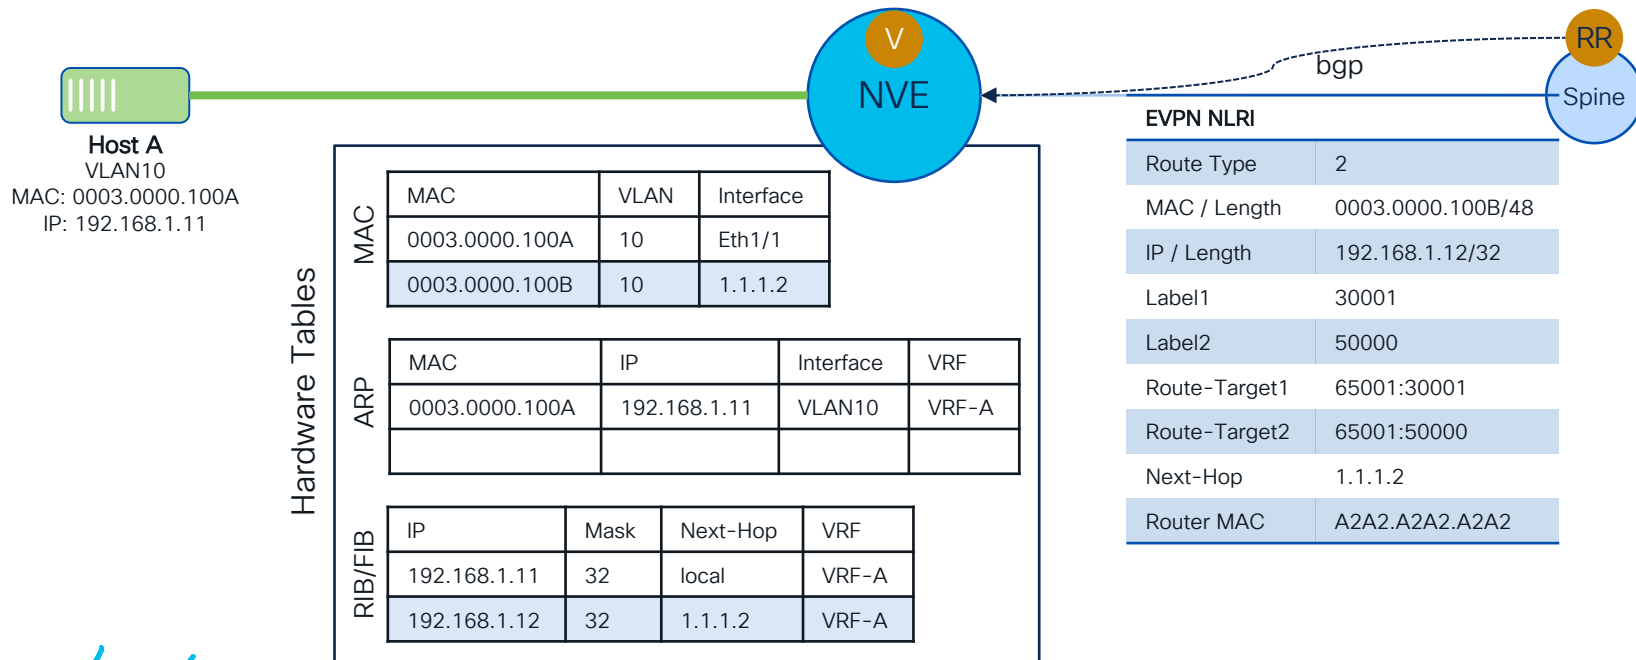
# Remote Learning from other NVE

## The Dating Network - When Control- meets Data-Plane



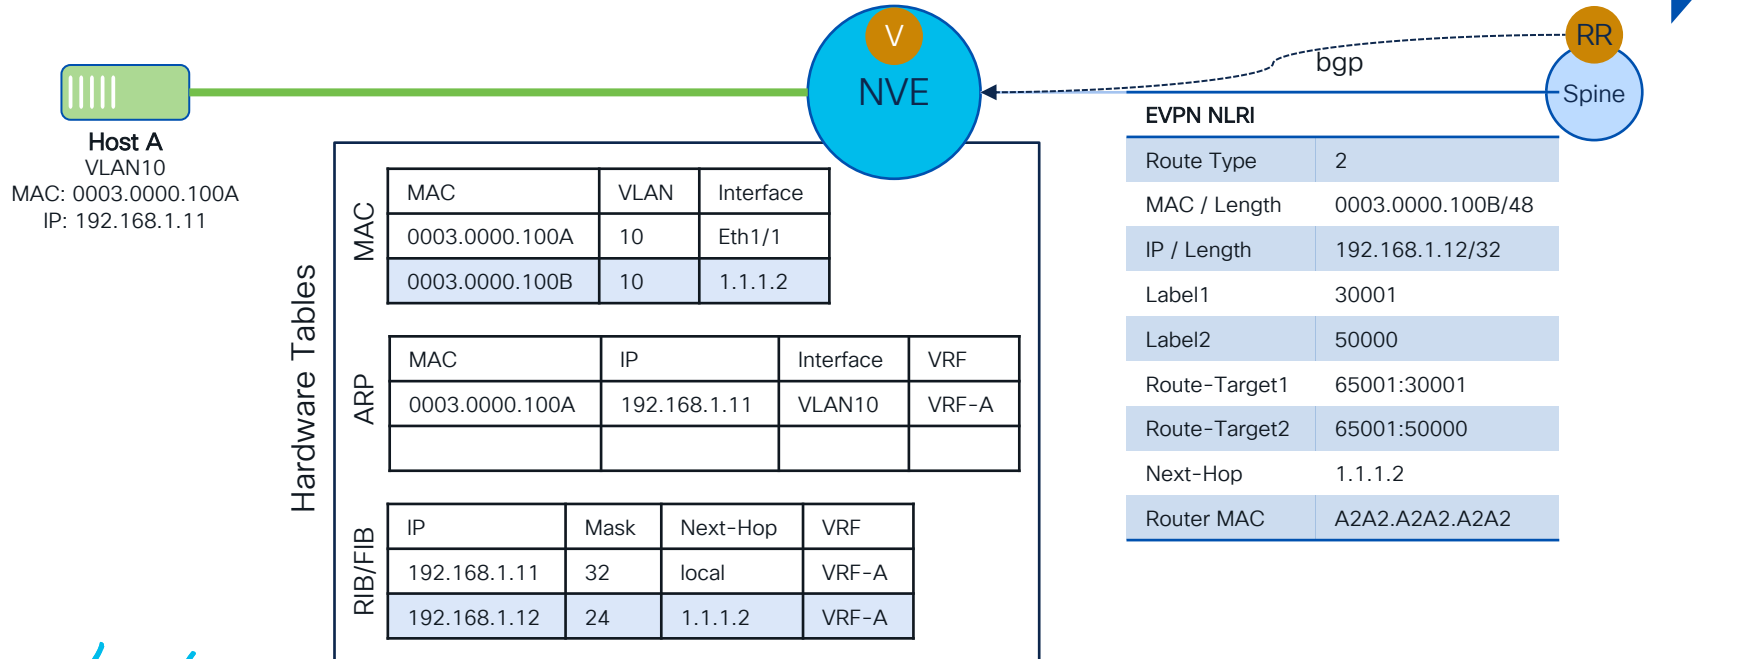

# Remote Learning from other NVE

## The Dating Network - When Control- meets Data-Plane



# Bridging between NVE (based on VXLAN EVPN)


## The Dating Network - When Control- meets Data-Plane



# Packet Walk: Layer-3 – Host to External World

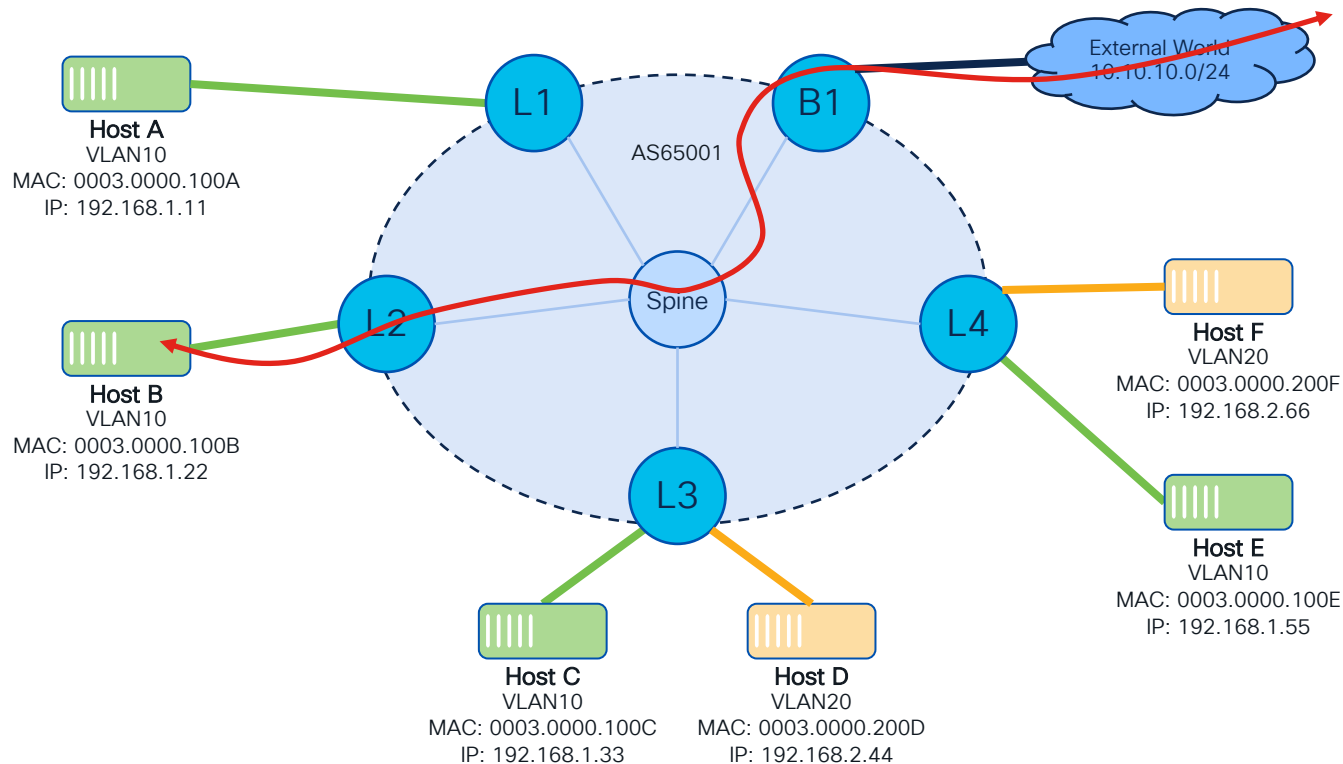
# Topology Overview

## Layer-3 Packet Walk



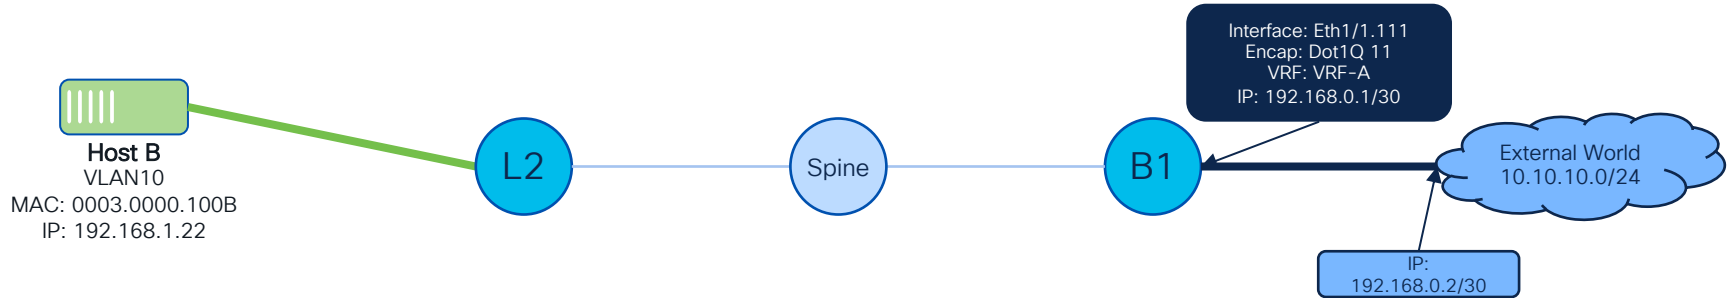
# Topology Overview

## Layer-3 Packet Walk

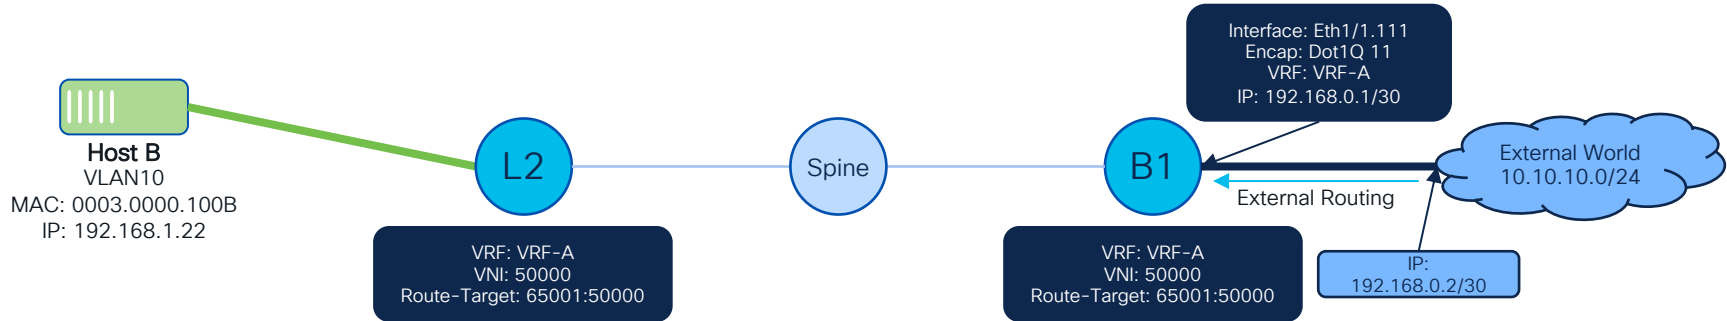




# Learning: External World to Leaf2

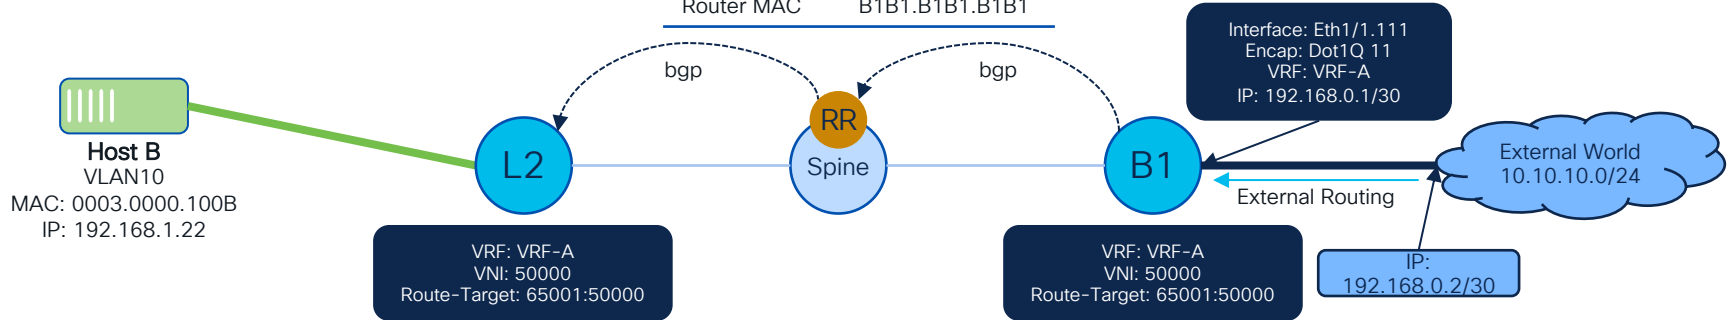


# Learning: External World to Leaf2

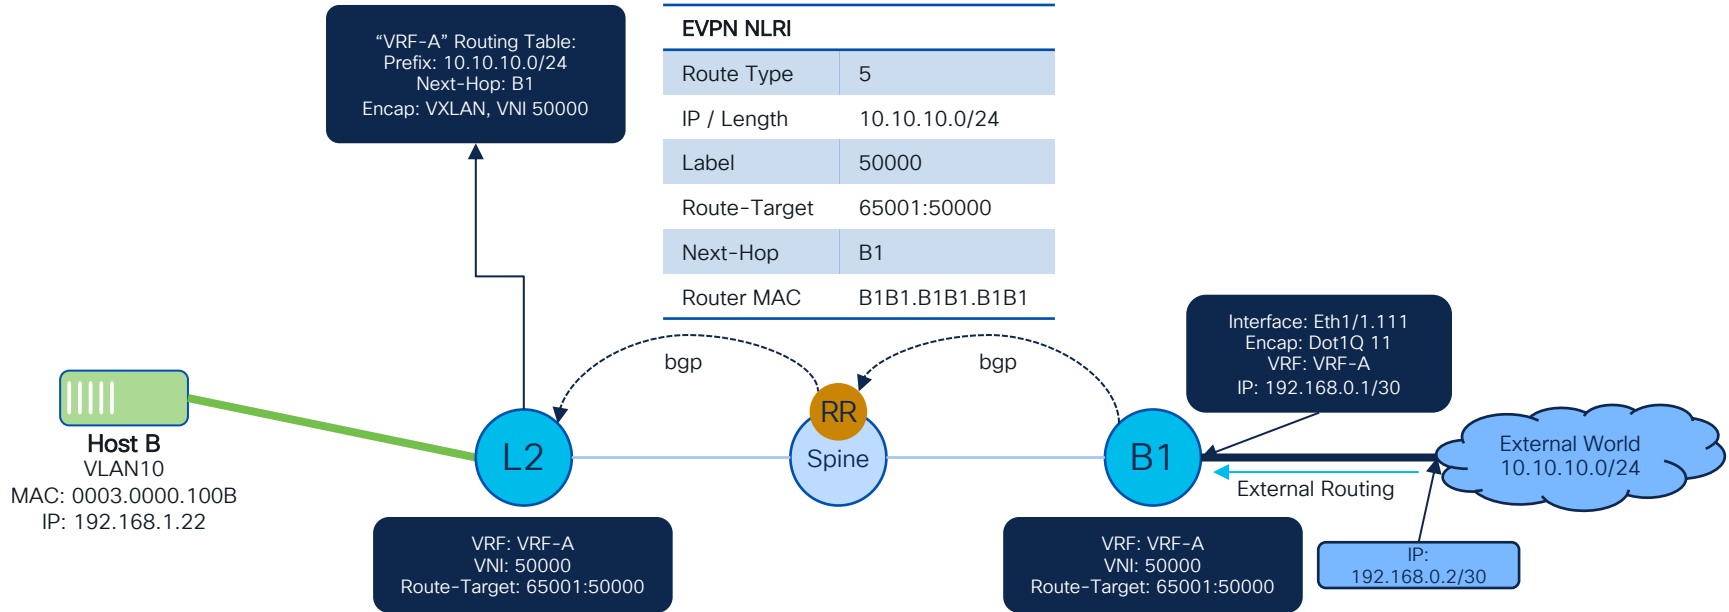


# Learning: External World to Leaf2

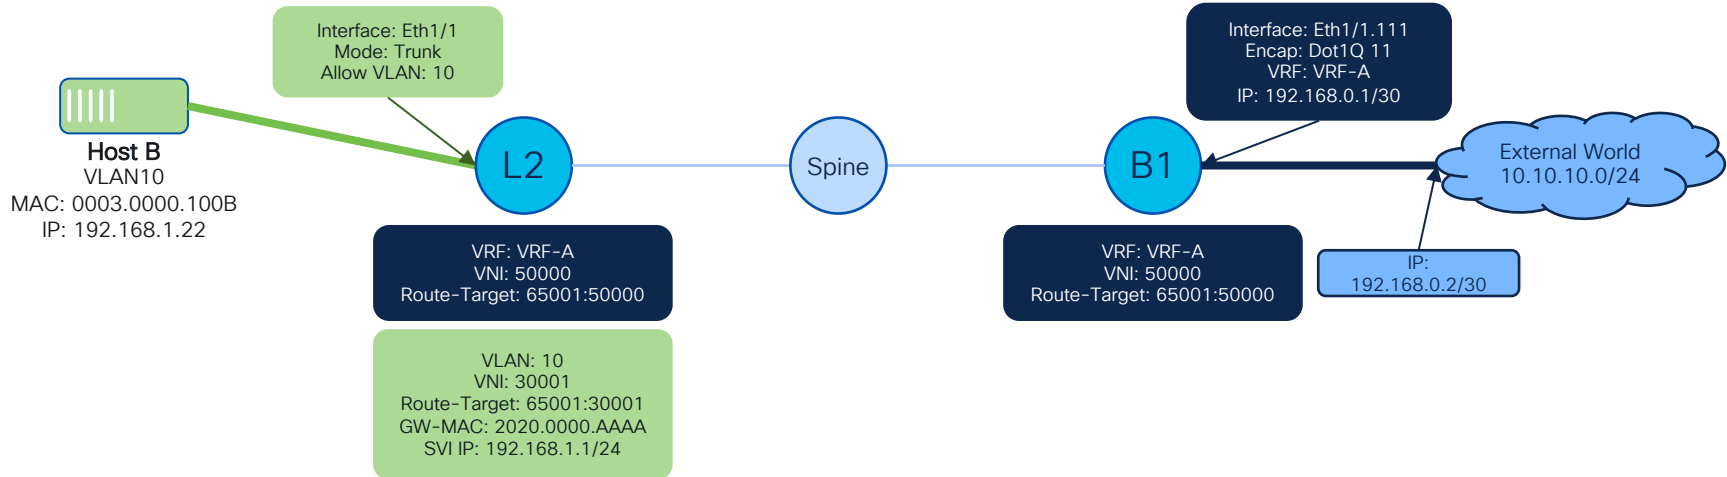
EVPN NLRI	
Route Type	5
IP / Length	10.10.10.0/24
Label	50000
Route-Target	65001:50000
Next-Hop	B1
Router MAC	B1B1.B1B1.B1B1



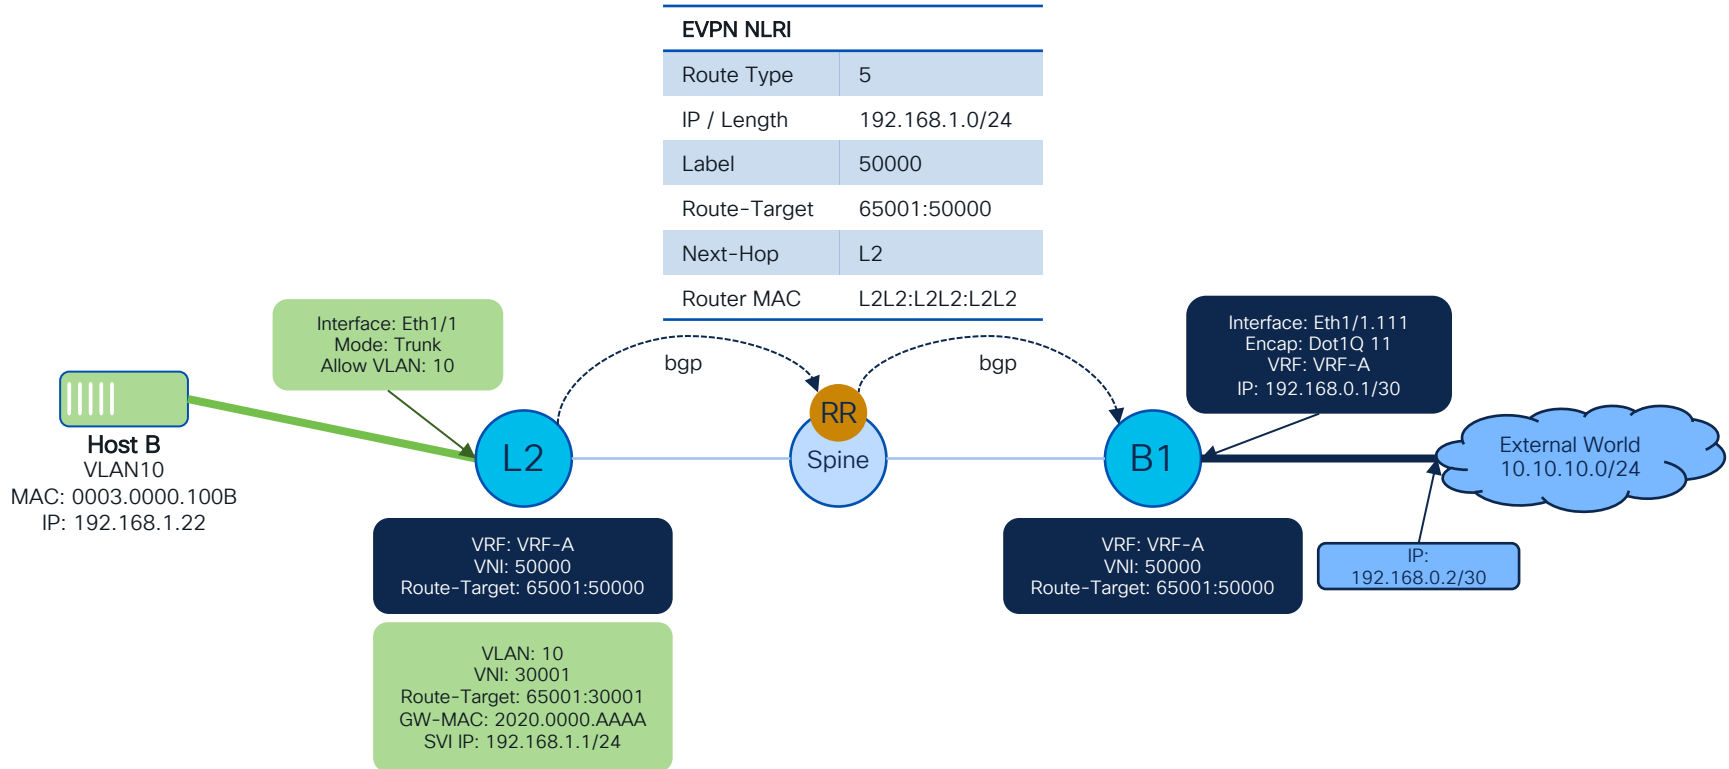
# Learning: External World to Leaf2



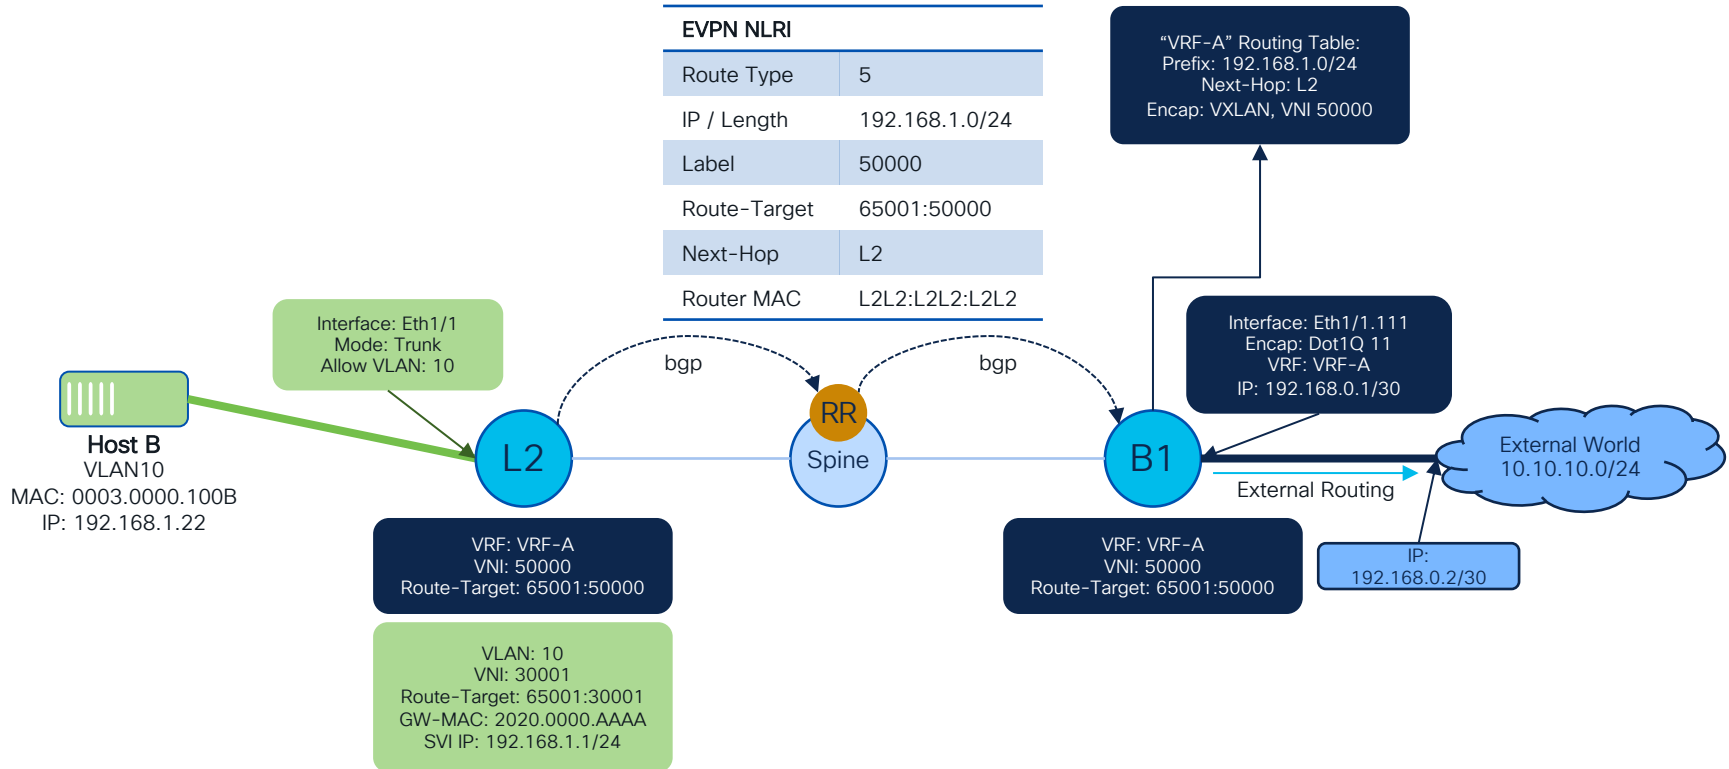
# Learning: HostB to External World



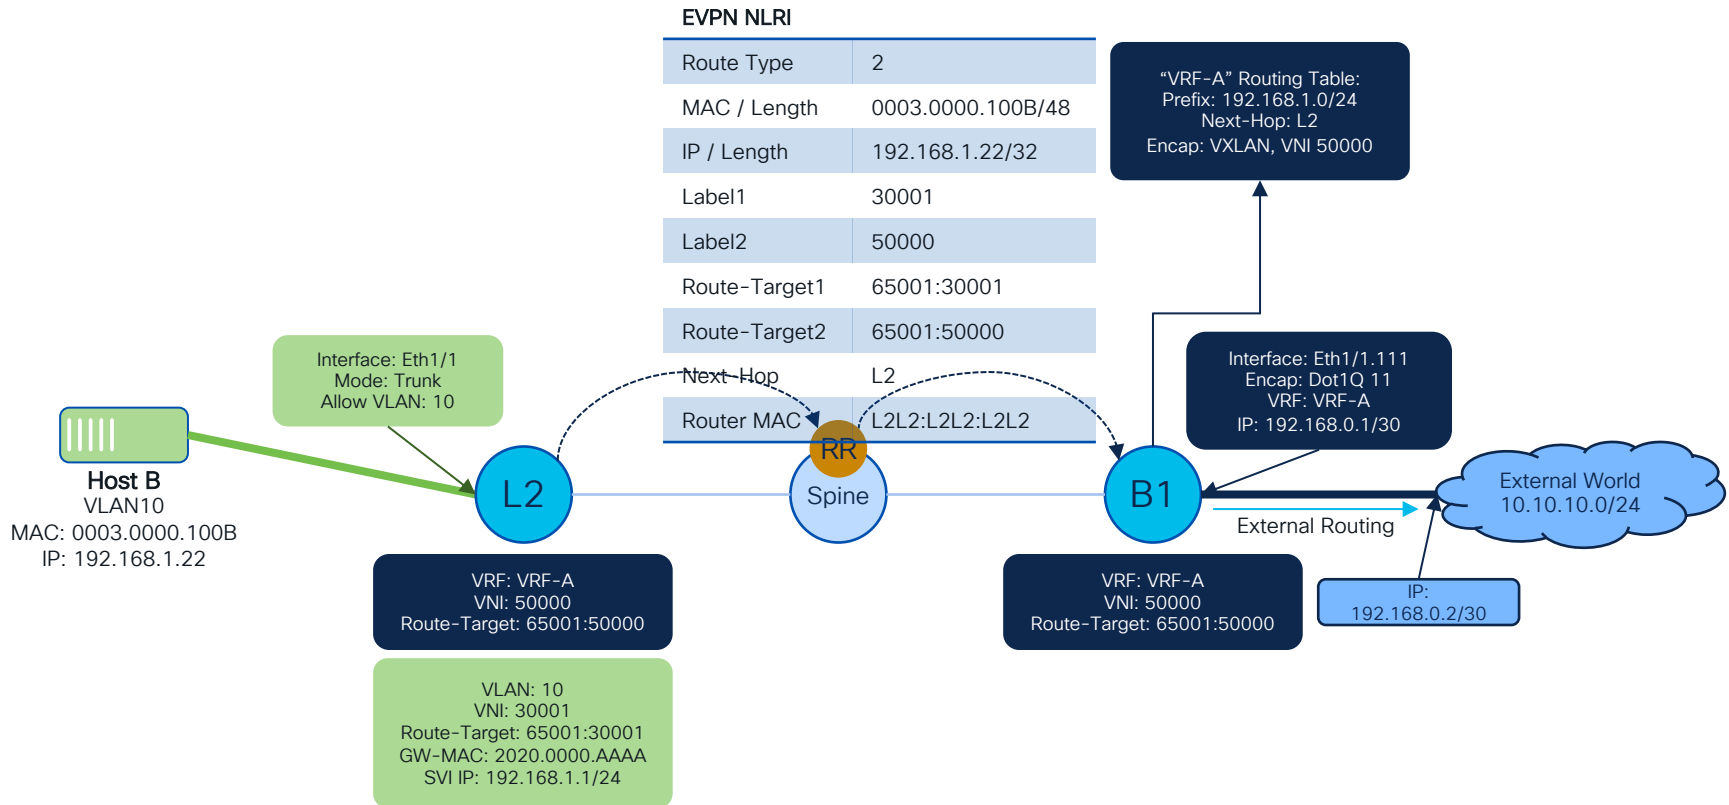
# Learning: HostB to External World



# Learning: HostB to External World

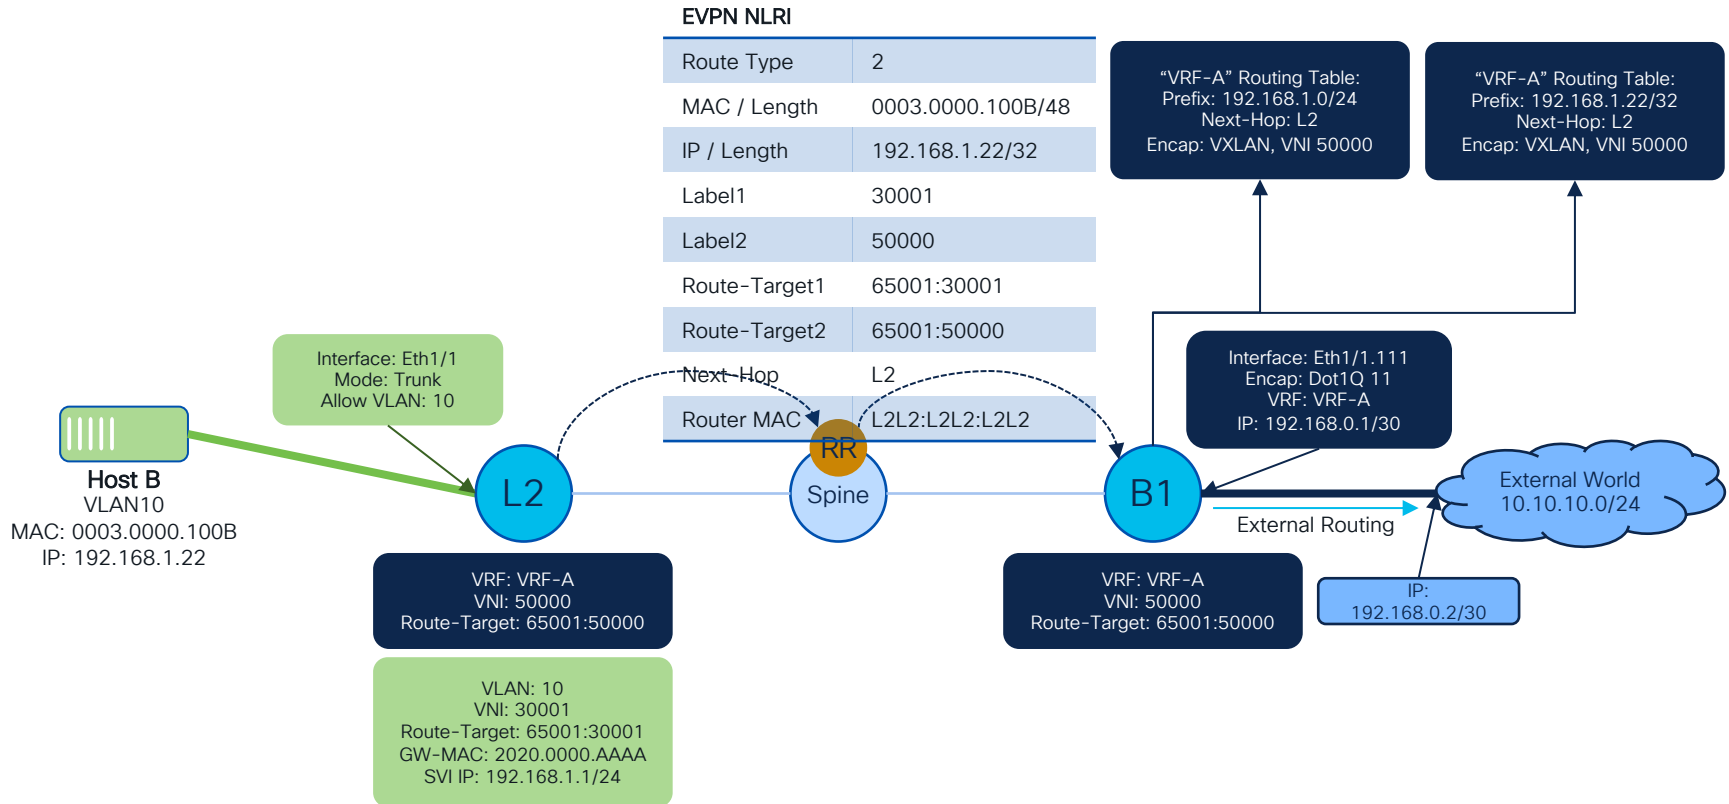


# Learning: HostB to External World

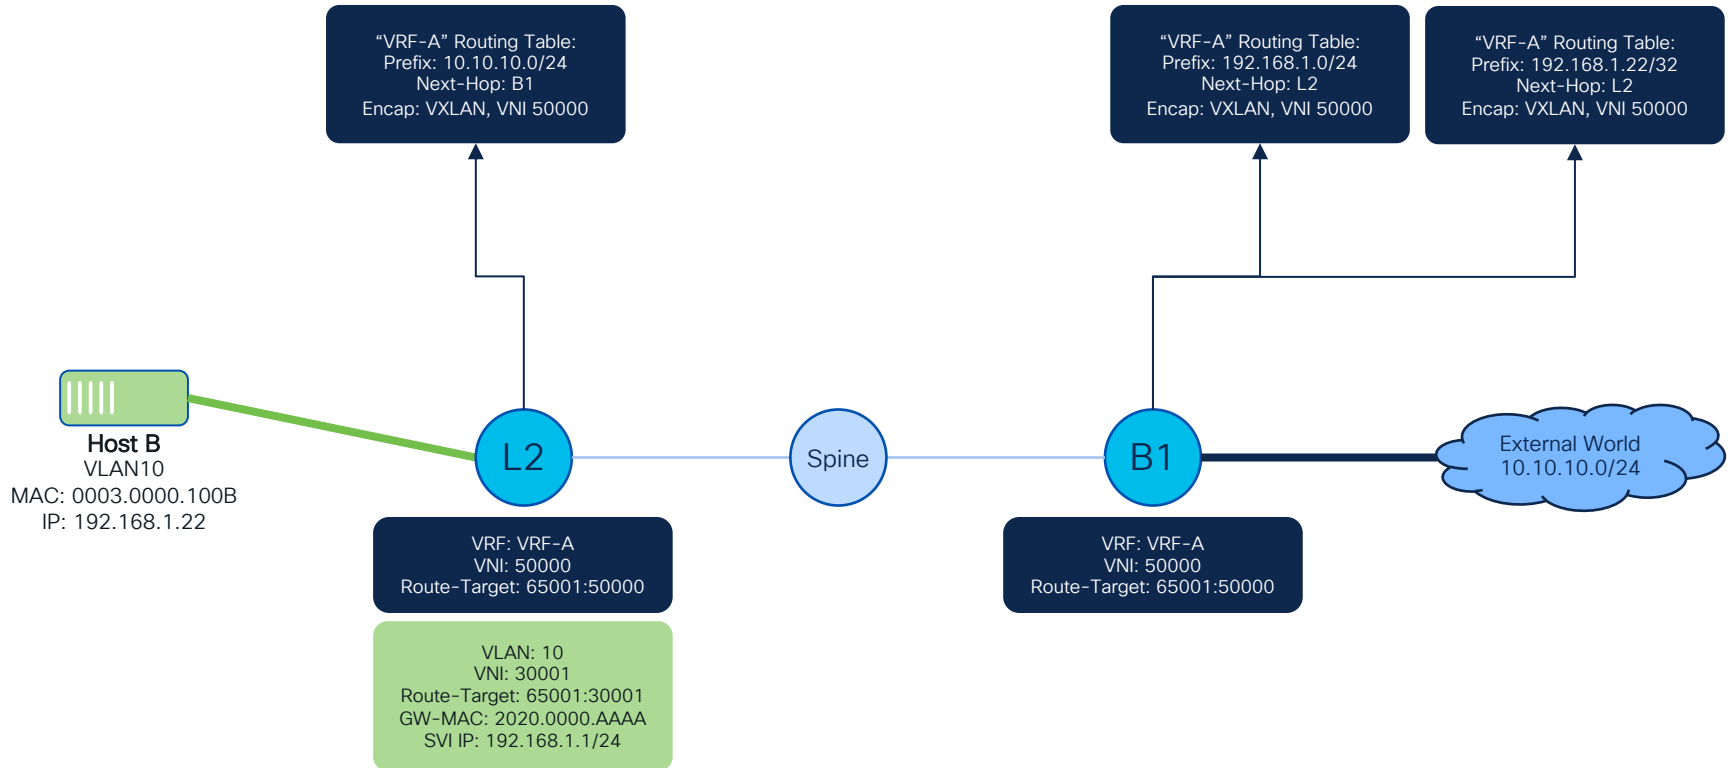




# Learning: HostB to External World




# Overview: Forwarding Tables



# HostB to External World


SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100B	2020.0000.AAAA	10	192.168.1.22	10.10.10.77	



# HostB to External World

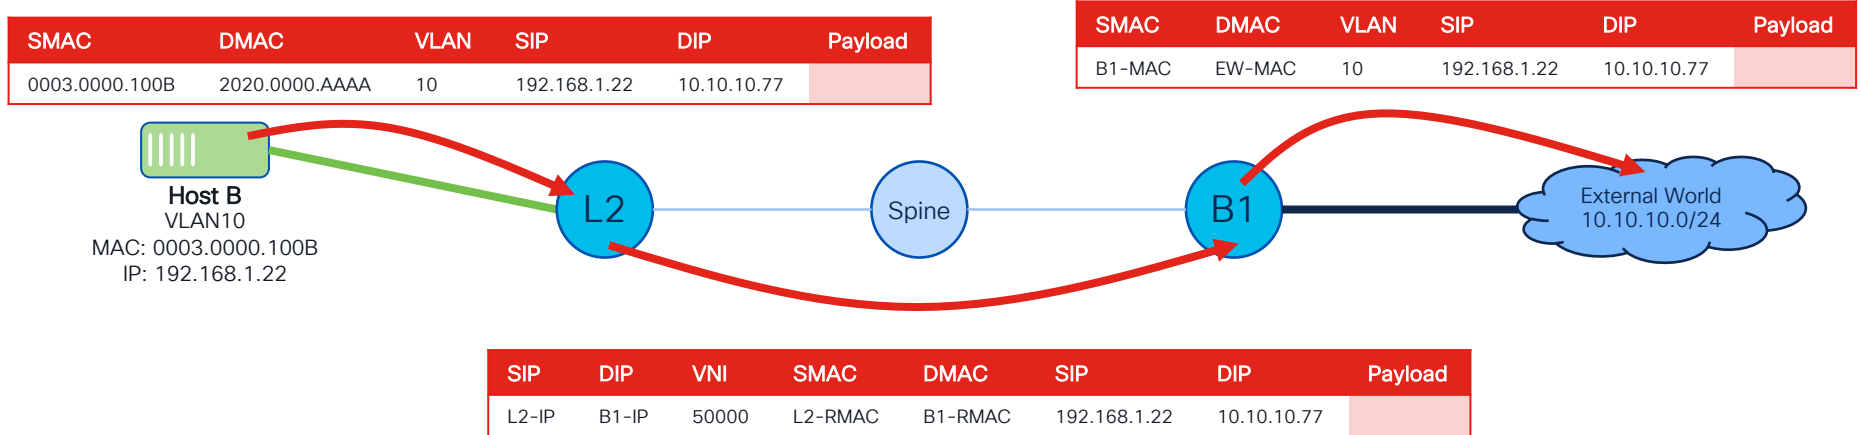
SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100B	2020.0000.AAAA	10	192.168.1.22	10.10.10.77	

Host B  
VLAN10  
MAC: 0003.0000.100B  
IP: 192.168.1.22

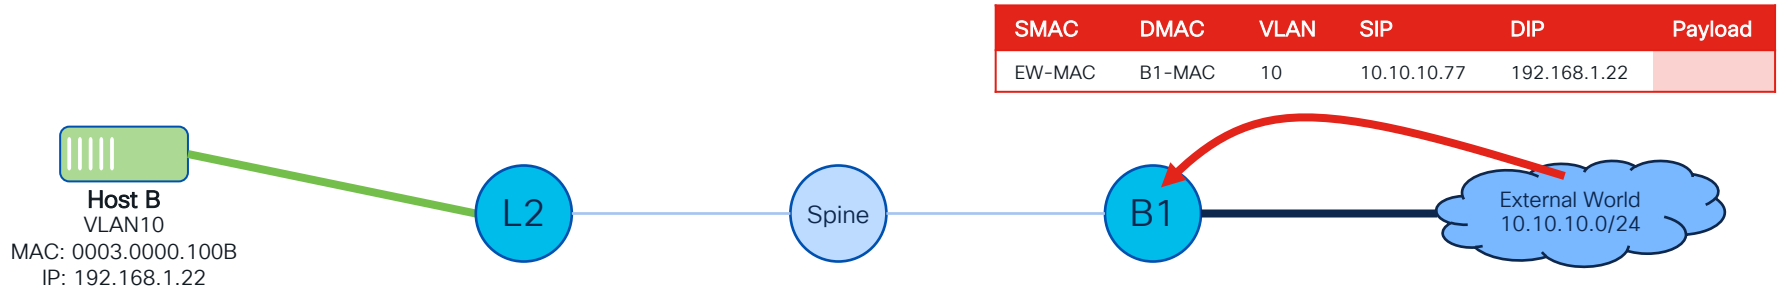


SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L2-IP	B1-IP	50000	L2-RMAC	B1-RMAC	192.168.1.22	10.10.10.77	

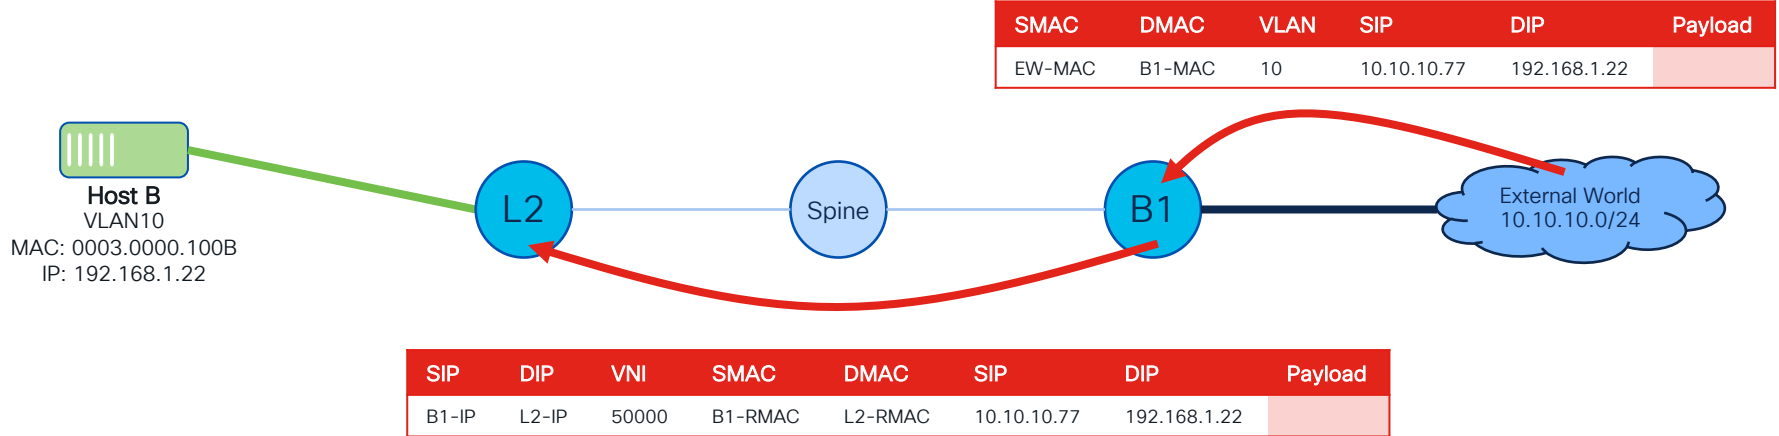
# HostB to External World



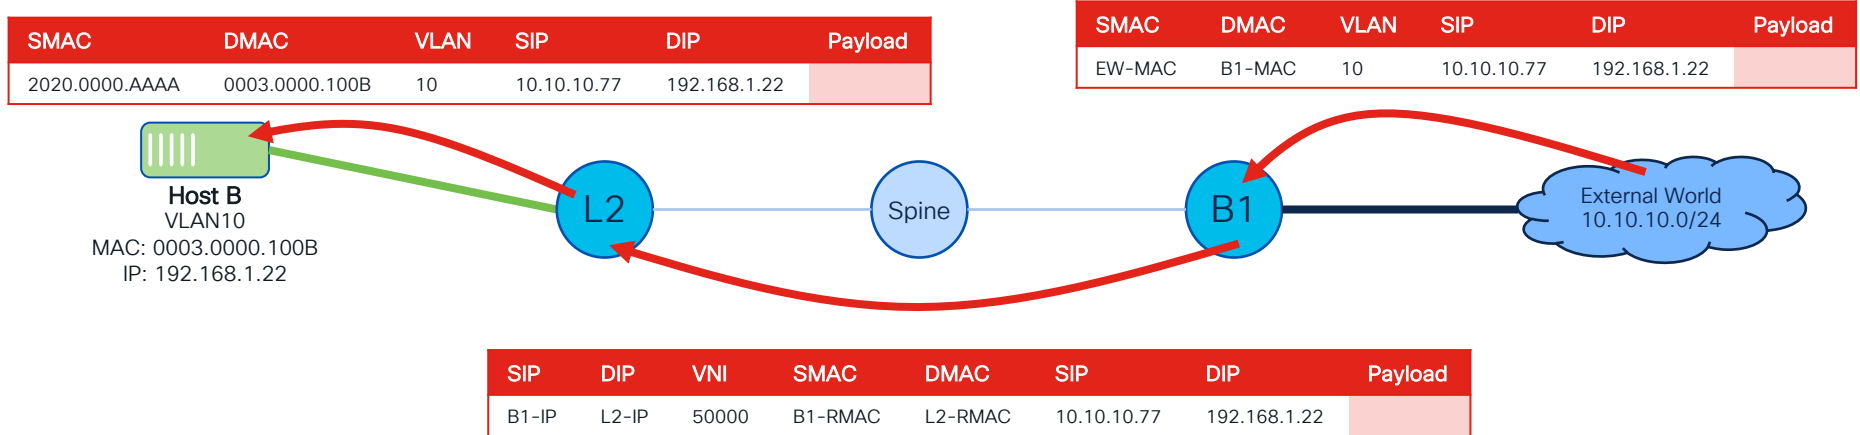
# External World to HostB



# External World to HostB



# External World to HostB





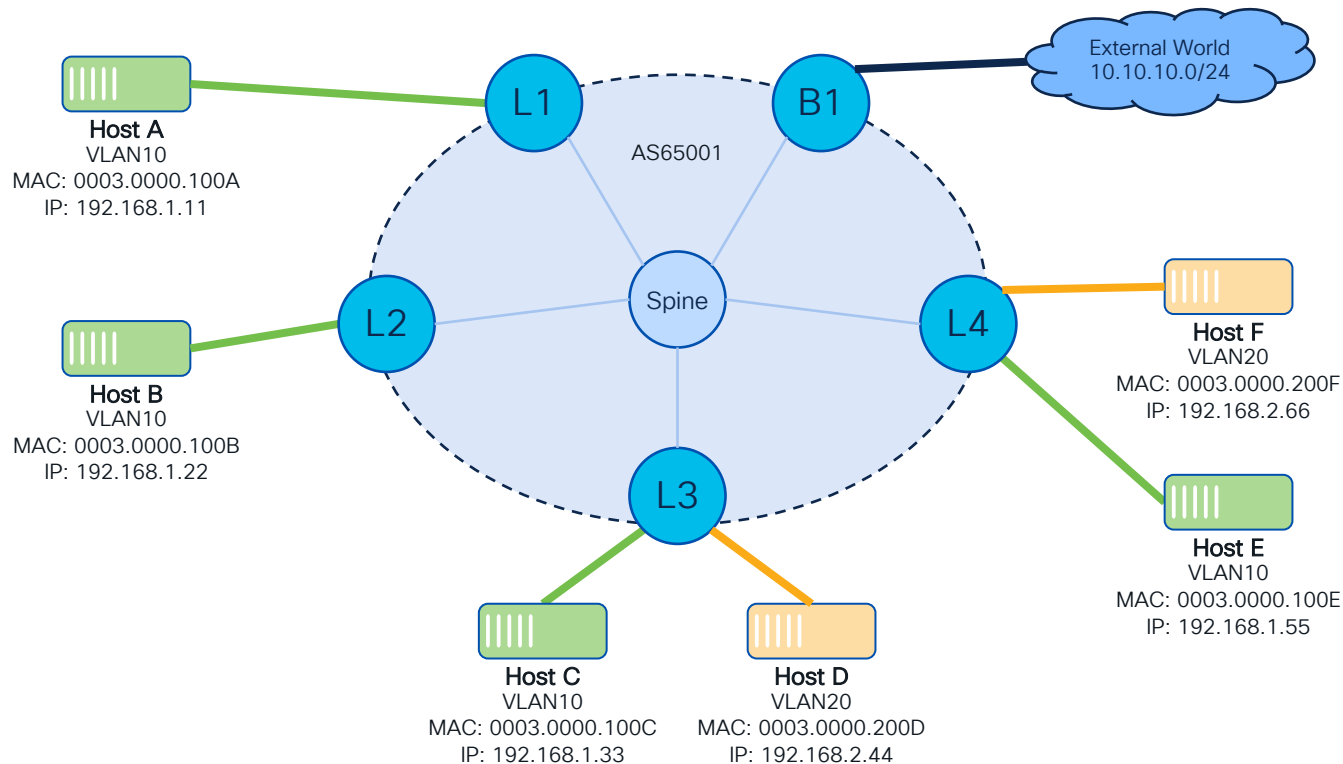
# How to Talk to the Rest of the World - External Connectivity for VXLAN EVPN Fabrics

BRKDCN-2267

# Packet Walk: Layer-3 – Host to Host

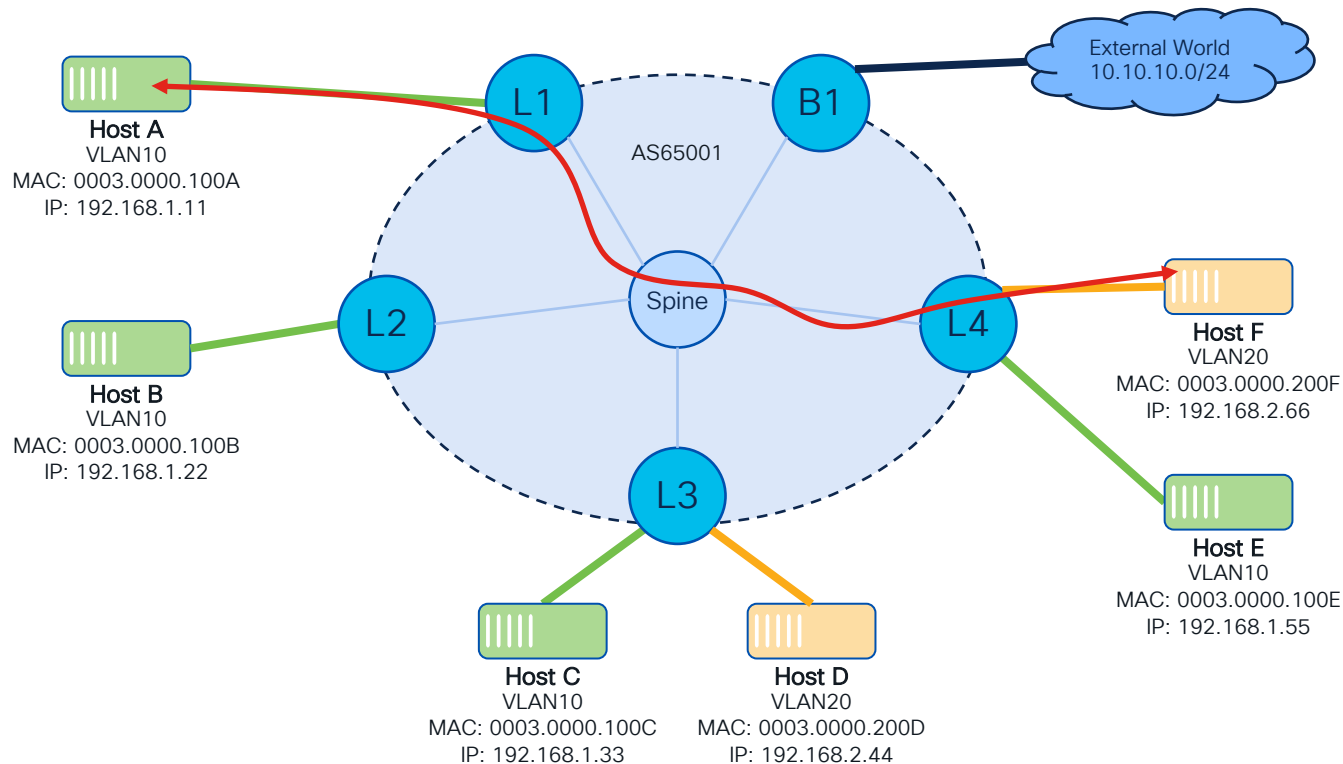
# Topology Overview

## Layer-3 Packet Walk




# Topology Overview

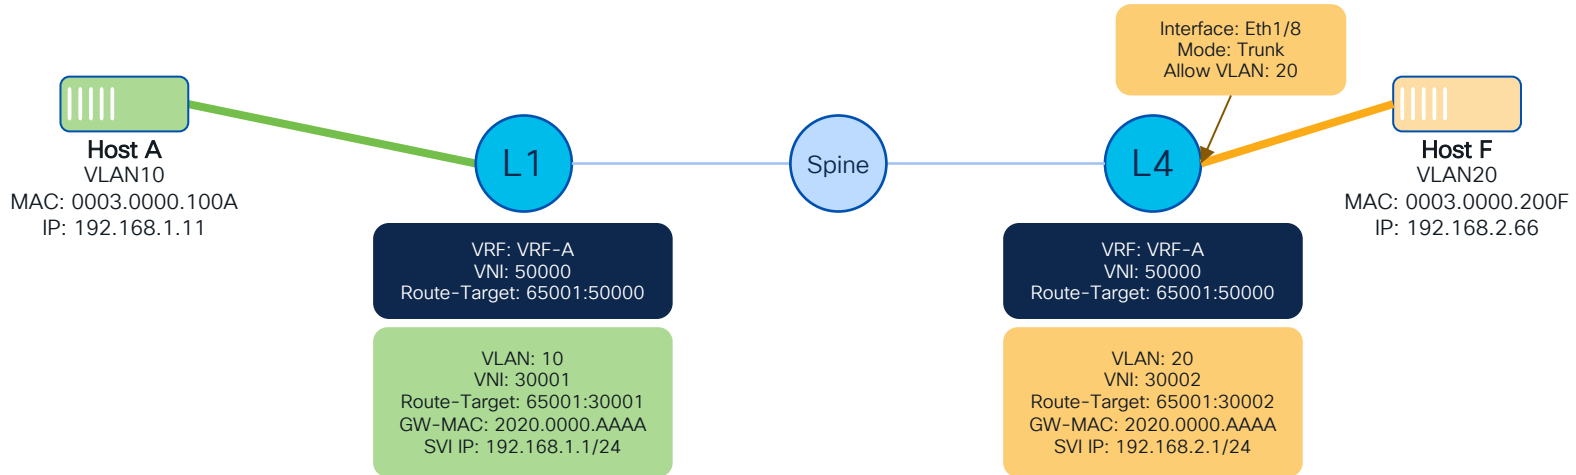
## Layer-3 Packet Walk



# Learning: HostF to Leaf1

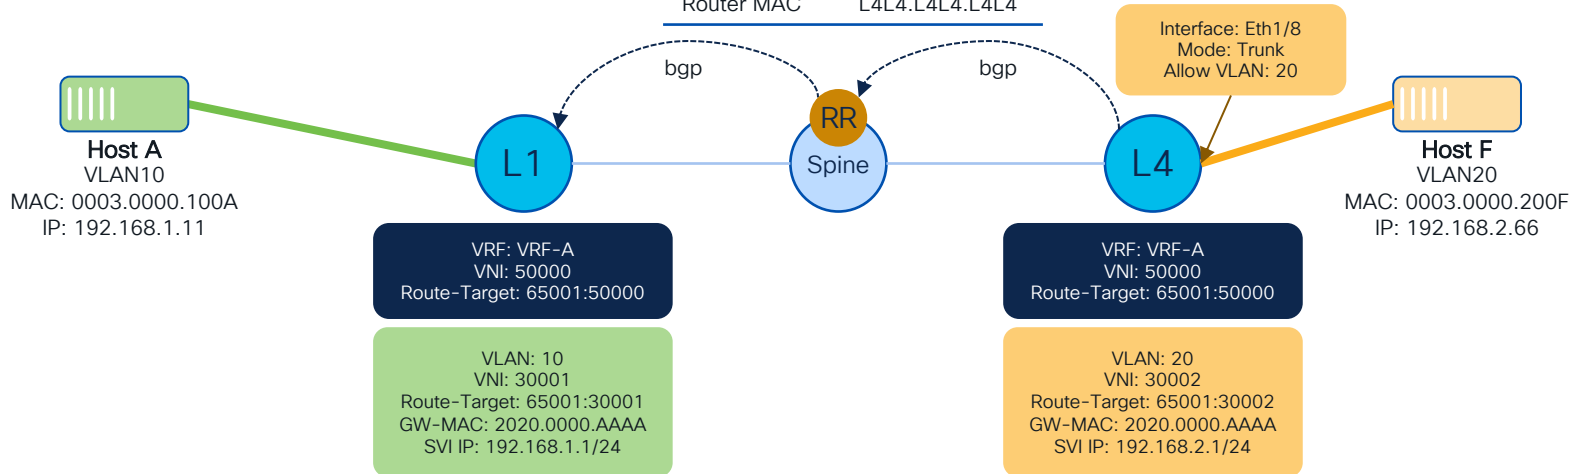


# Learning: HostF to Leaf1

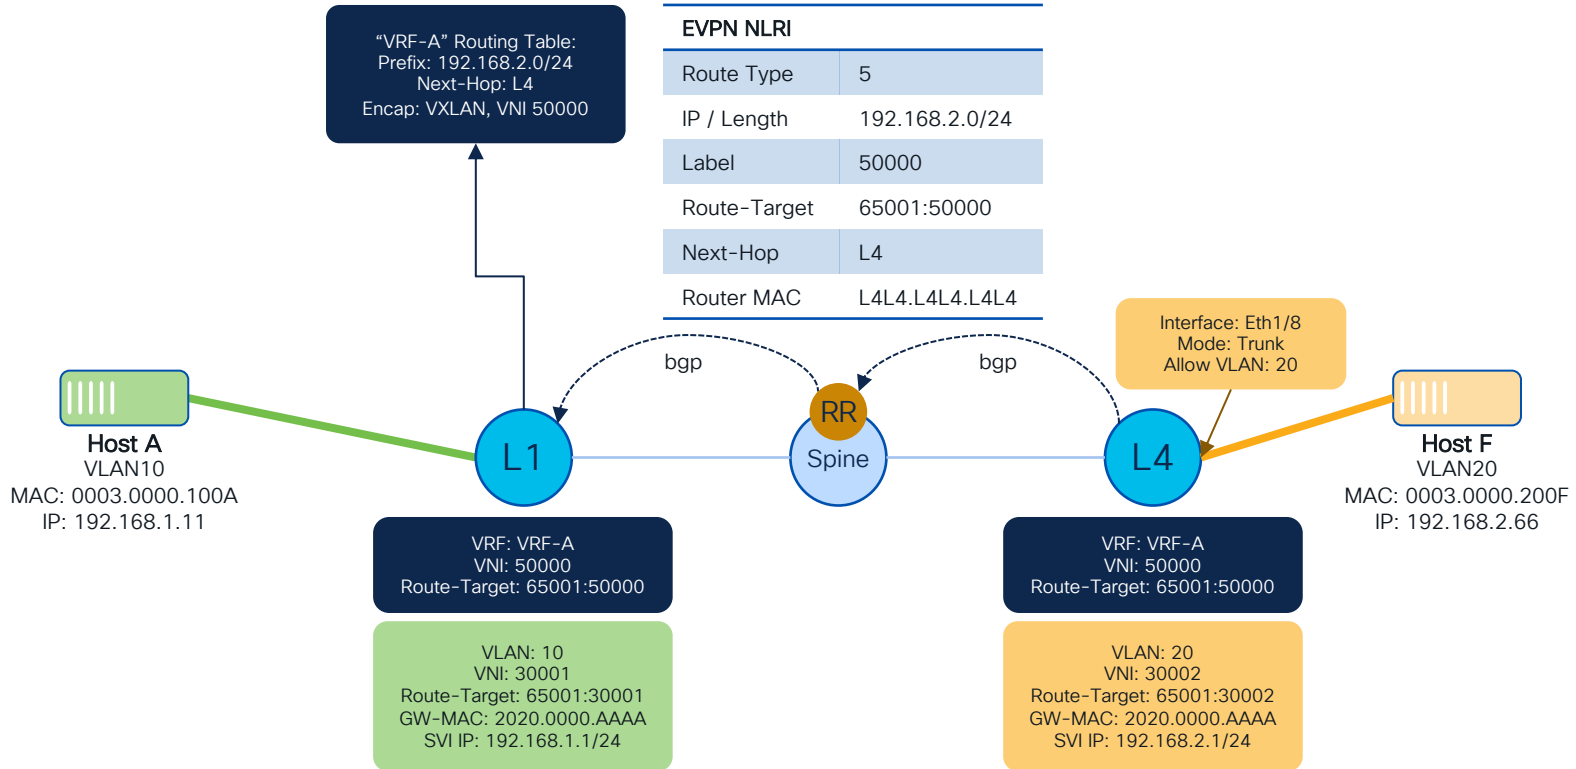


# Learning: HostF to Leaf1

EVPN NLRI	
Route Type	5
IP / Length	192.168.2.0/24
Label	50000
Route-Target	65001:50000
Next-Hop	L4
Router MAC	L4L4.L4L4.L4L4

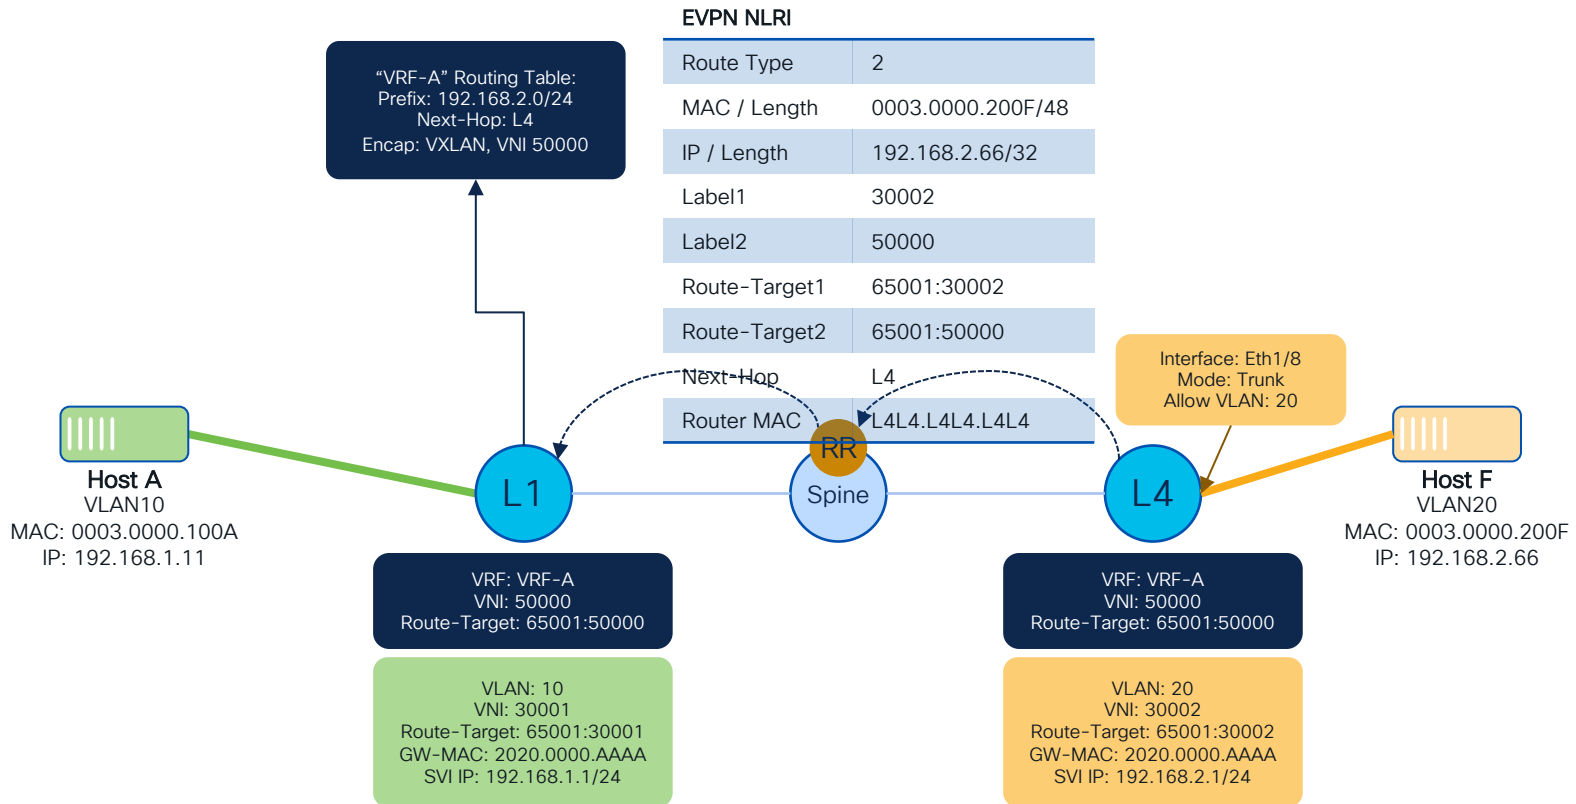


# Learning: HostF to Leaf1

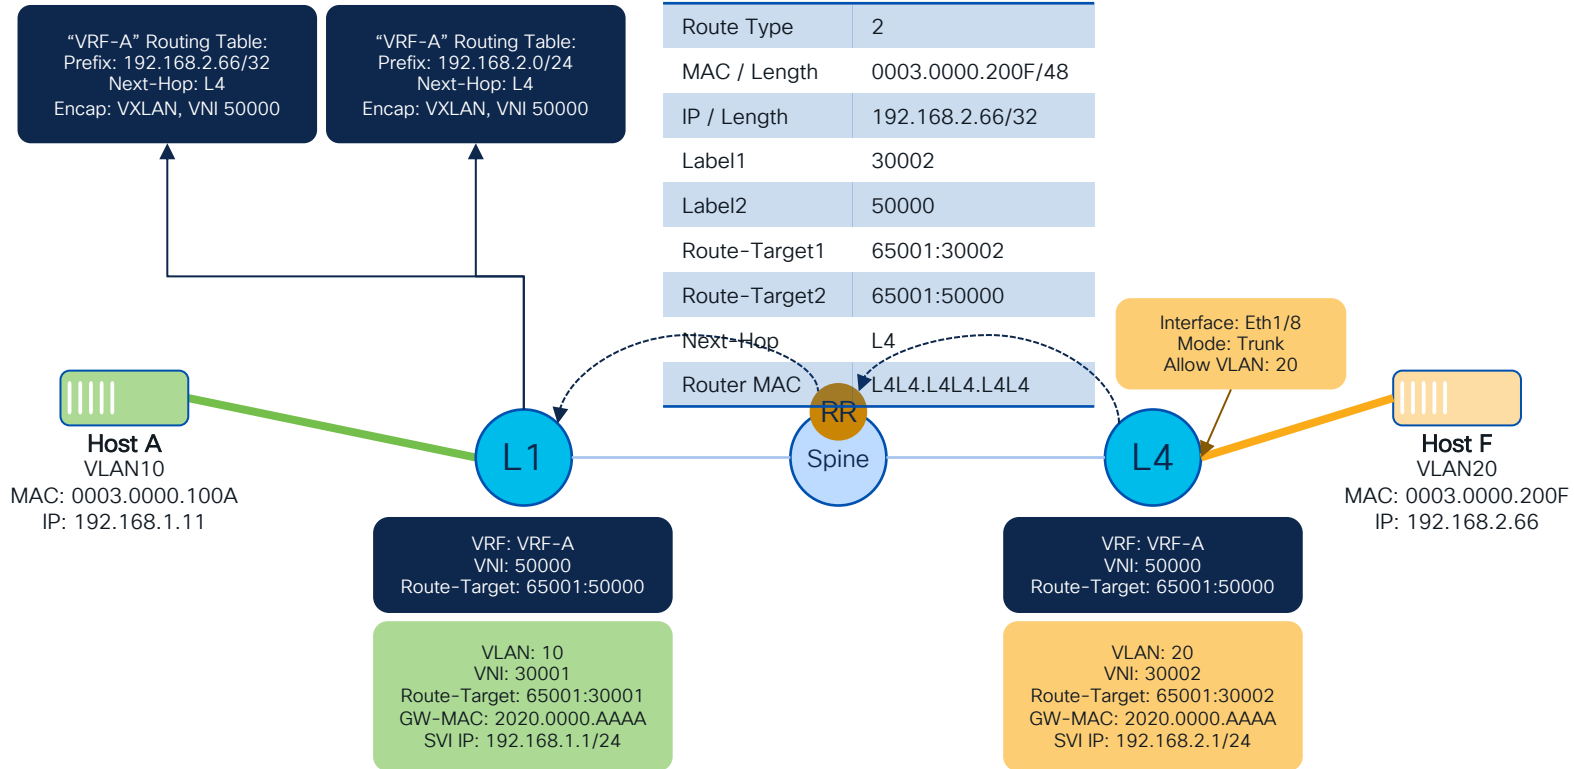




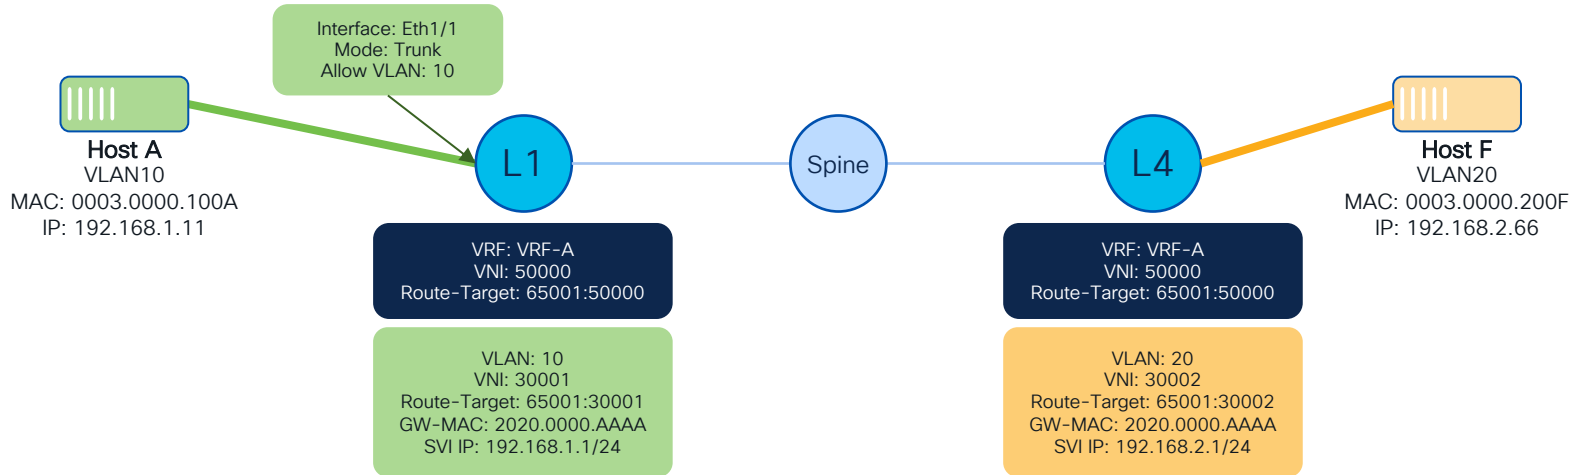
# Learning: HostF to Leaf1



# Learning: HostF to Leaf1

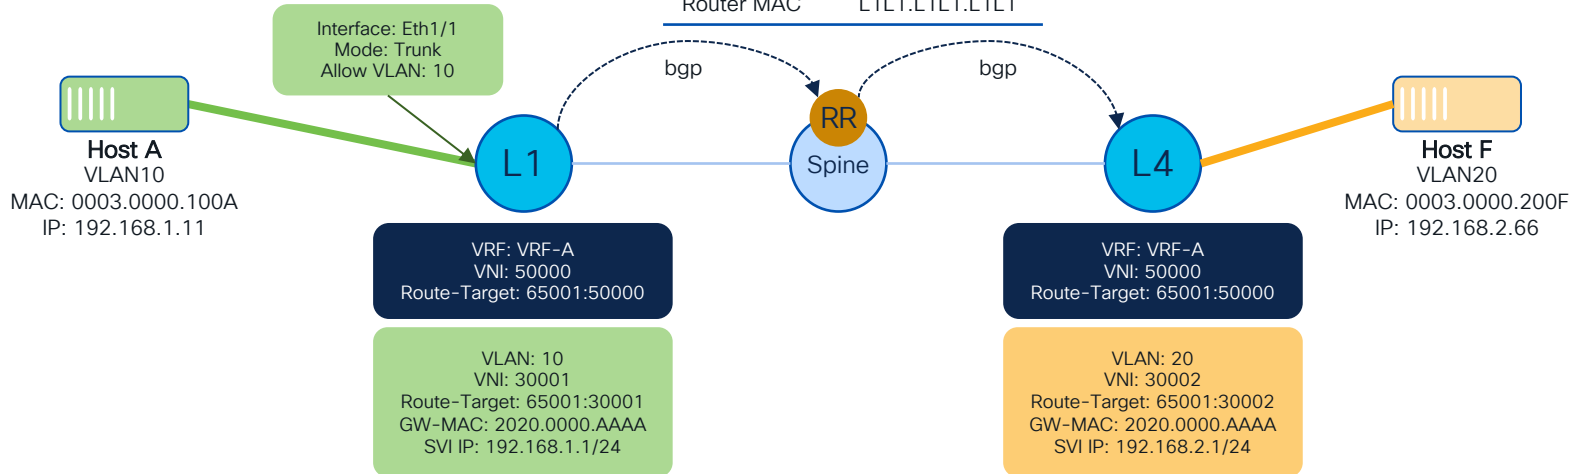


# Learning: HostA to Leaf4

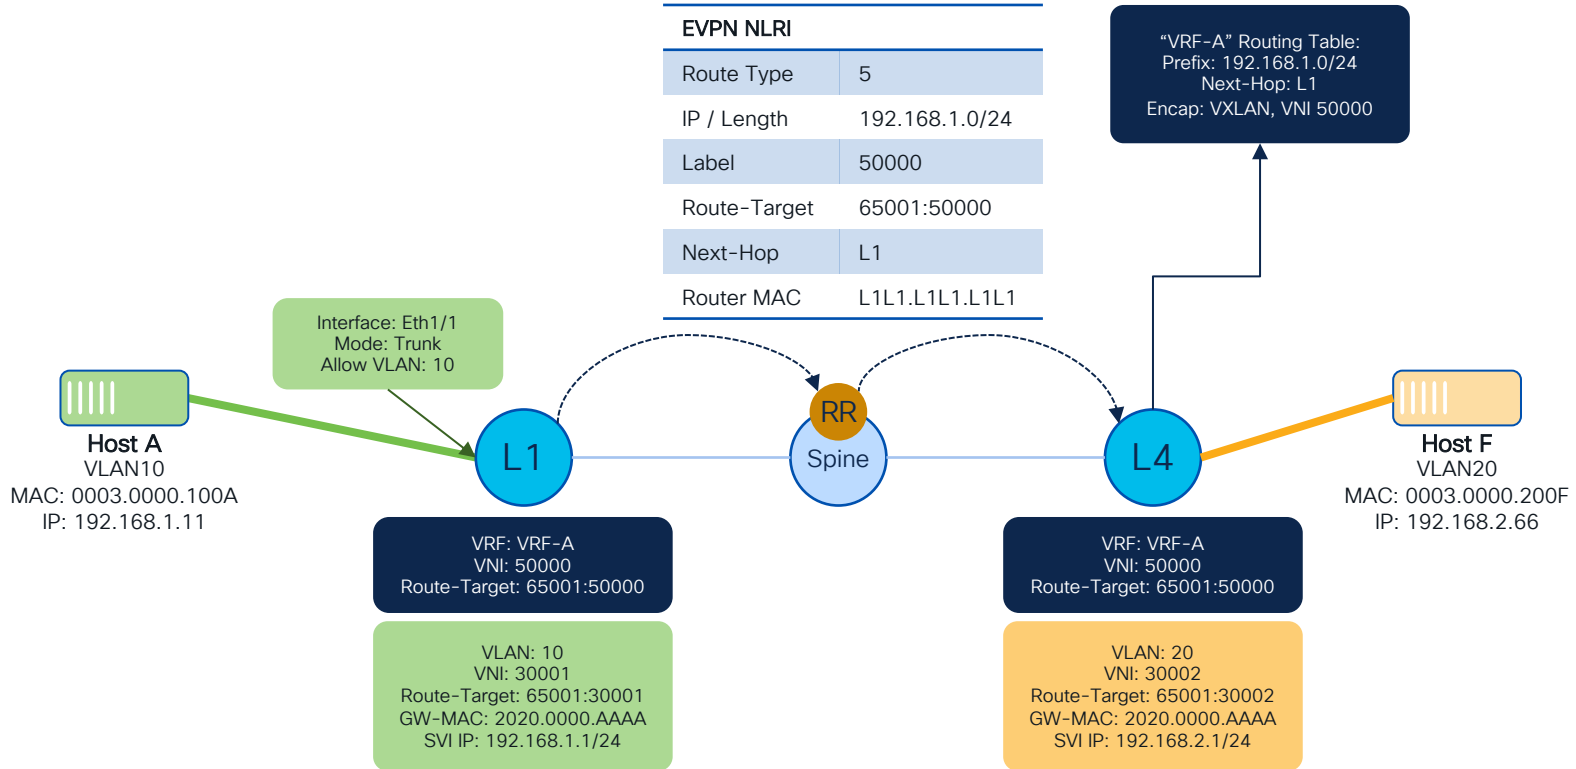


# Learning: HostA to Leaf4

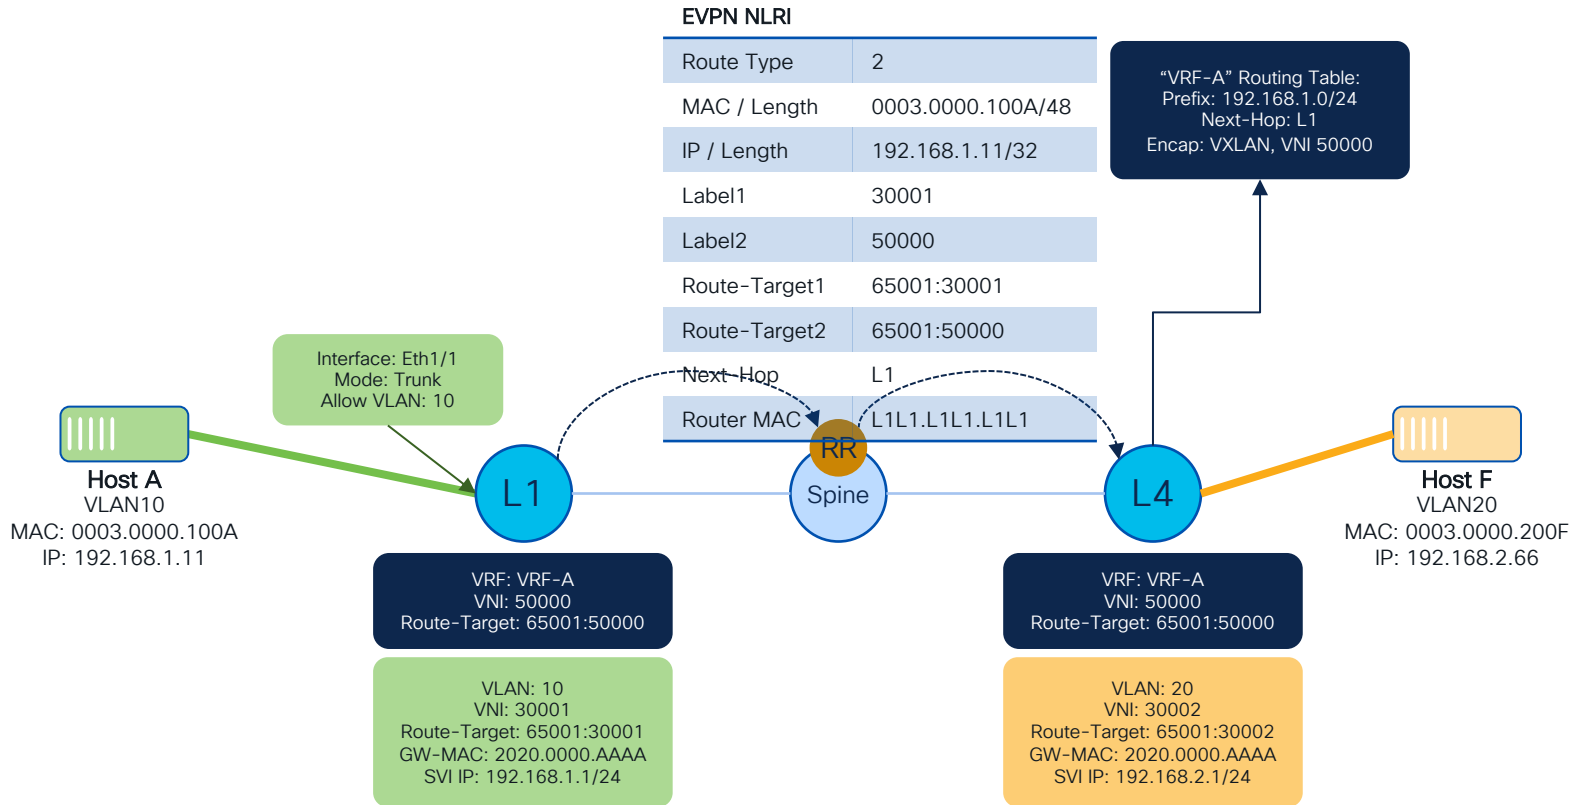
EVPN NLRI	
Route Type	5
IP / Length	192.168.1.0/24
Label	50000
Route-Target	65001:50000
Next-Hop	L1
Router MAC	L1L1.L1L1.L1L1



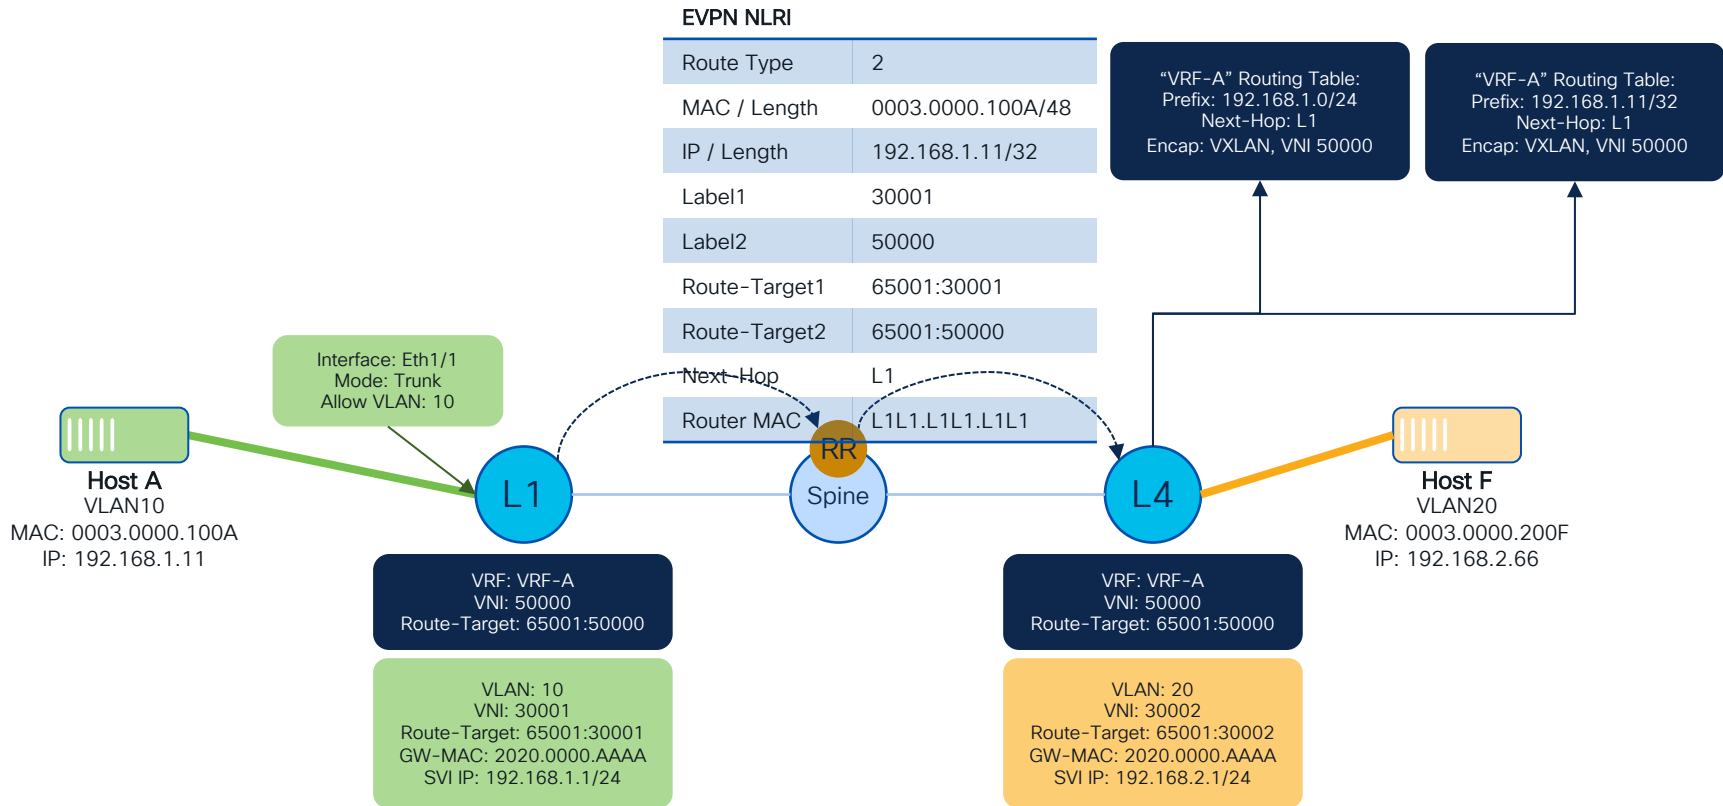
# Learning: HostA to Leaf4



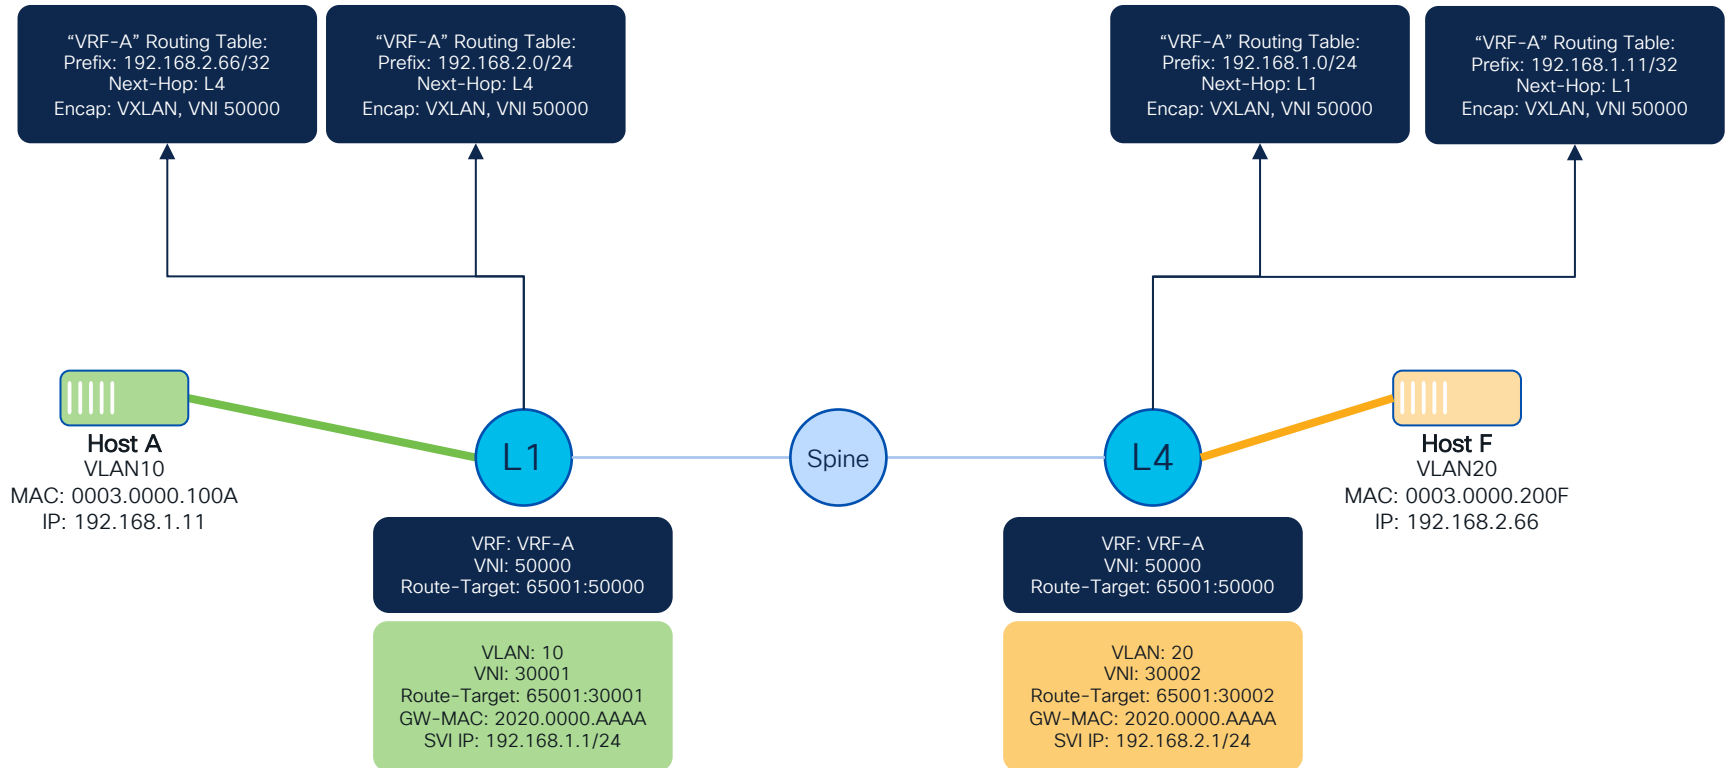
# Learning: HostA to Leaf4



# Learning: HostA to Leaf4




# Overview: Forwarding Tables





# HostA to HostF


SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.66	



# HostA to HostF

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.66	

Host A  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11




Host F  
VLAN20  
MAC: 0003.0000.200F  
IP: 192.168.2.66

SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	L4-IP	50000	L1-RMAC	L4-RMAC	192.168.1.11	192.168.2.66	

# HostA to HostF

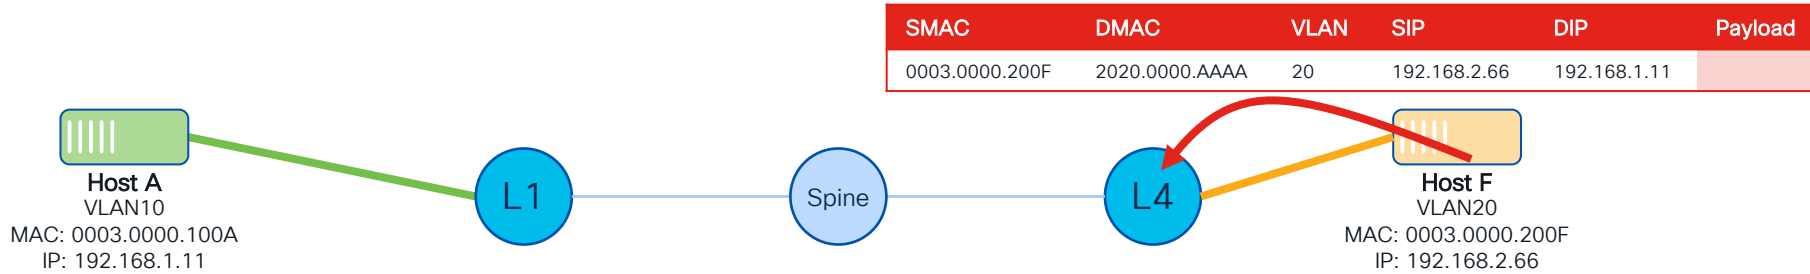
SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.66	

SMAC	DMAC	VLAN	SIP	DIP	Payload
2020.0000.AAAA	0003.0000.200F	20	192.168.1.11	192.168.2.66	

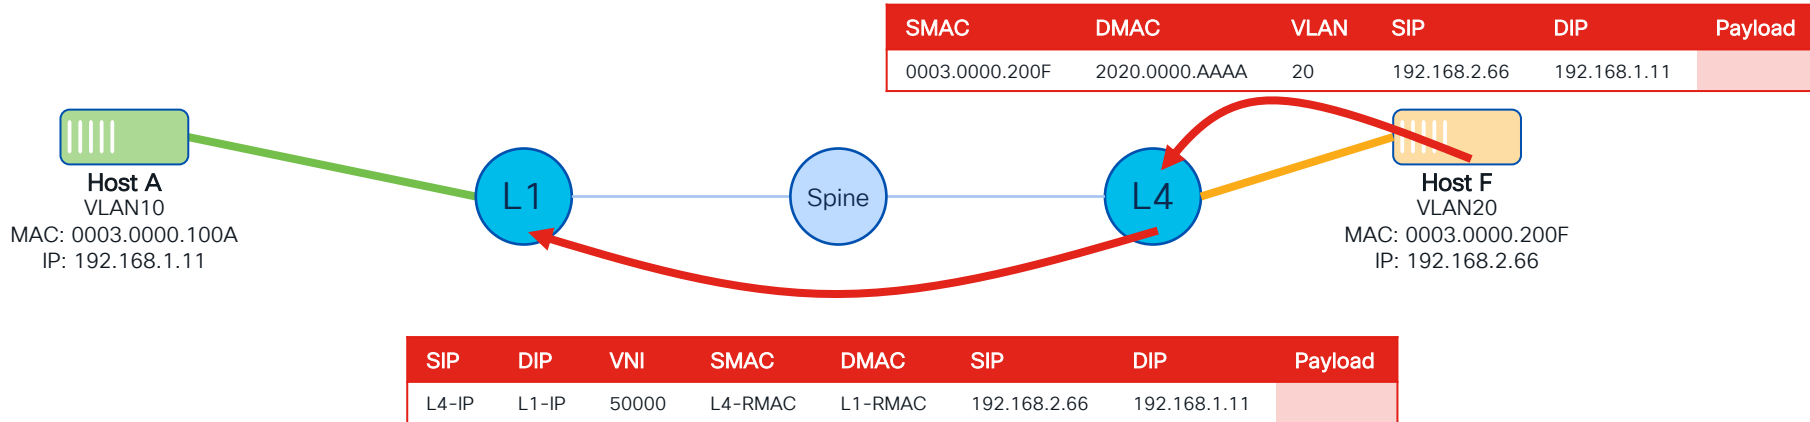


SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	L4-IP	50000	L1-RMAC	L4-RMAC	192.168.1.11	192.168.2.66	

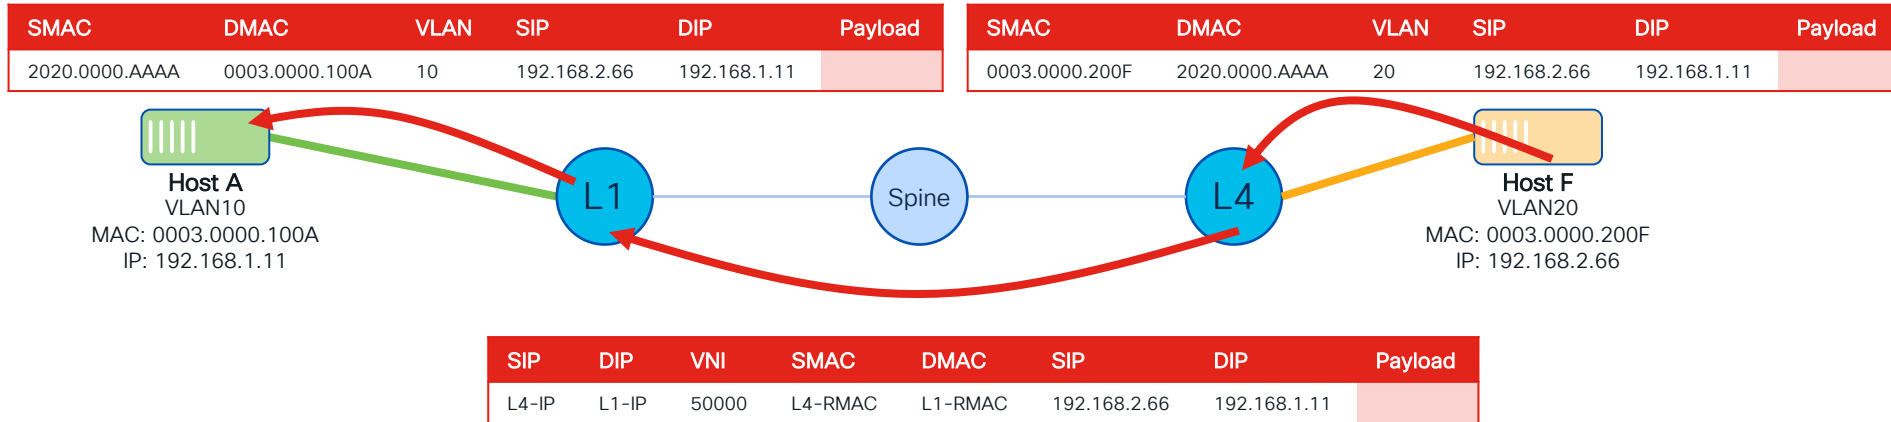
# HostF to HostA



# HostF to HostA



# HostF to HostA




# Packet Walk: Layer-2



# Topology Overview


## Layer-2 Packet Walk






# Topology Overview

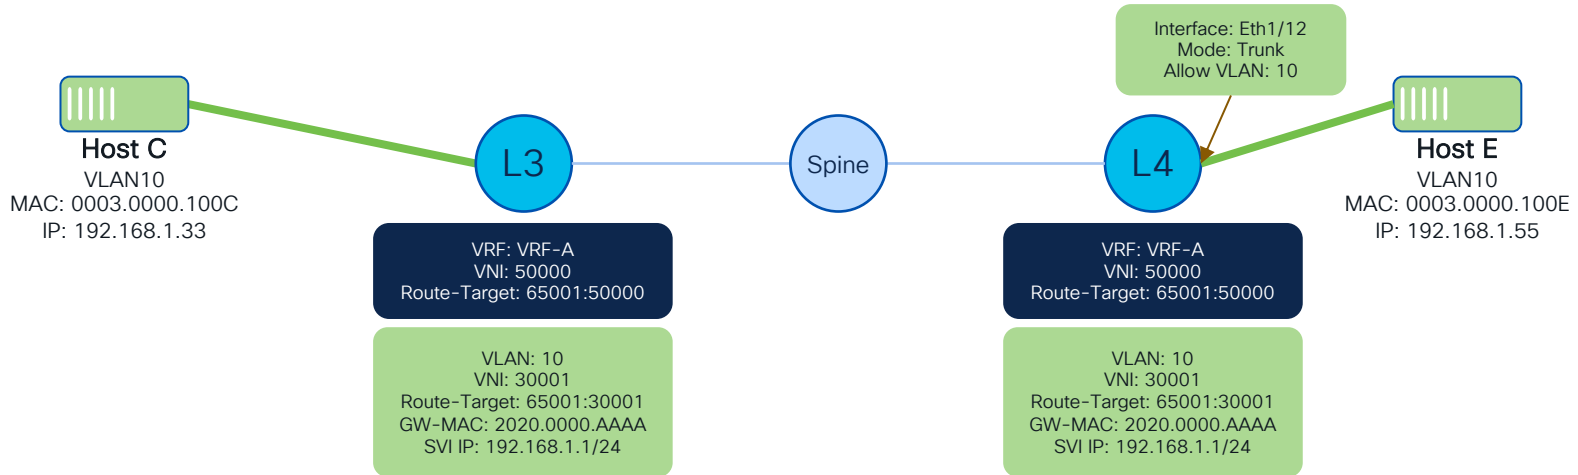
## Layer-2 Packet Walk



# Learning: HostE to Leaf3

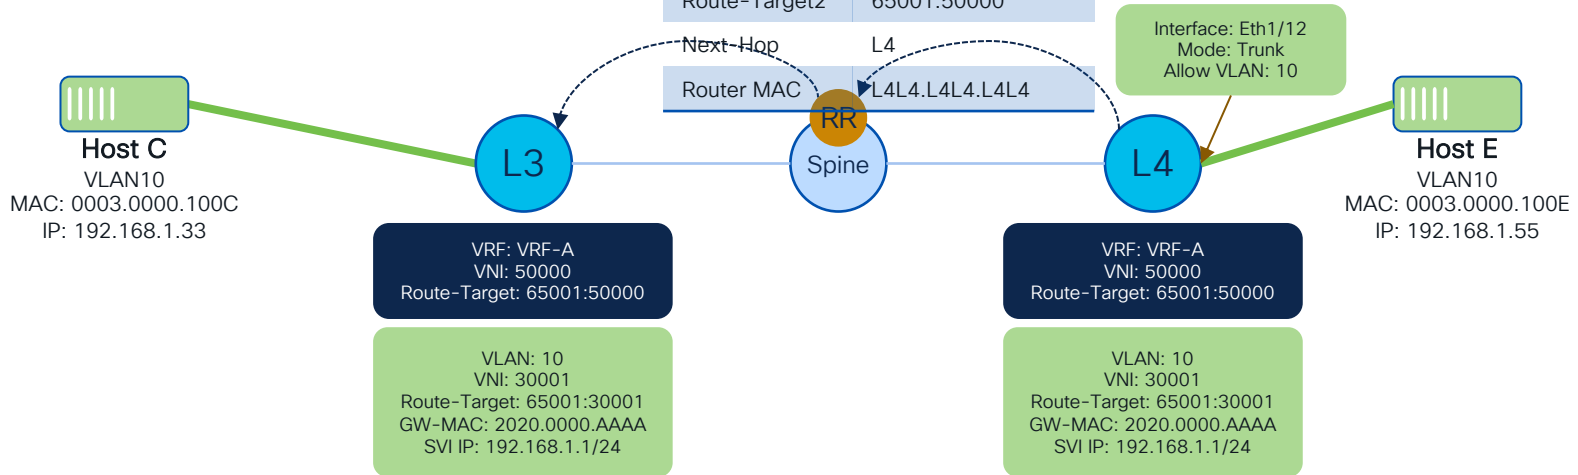


# Learning: HostE to Leaf3

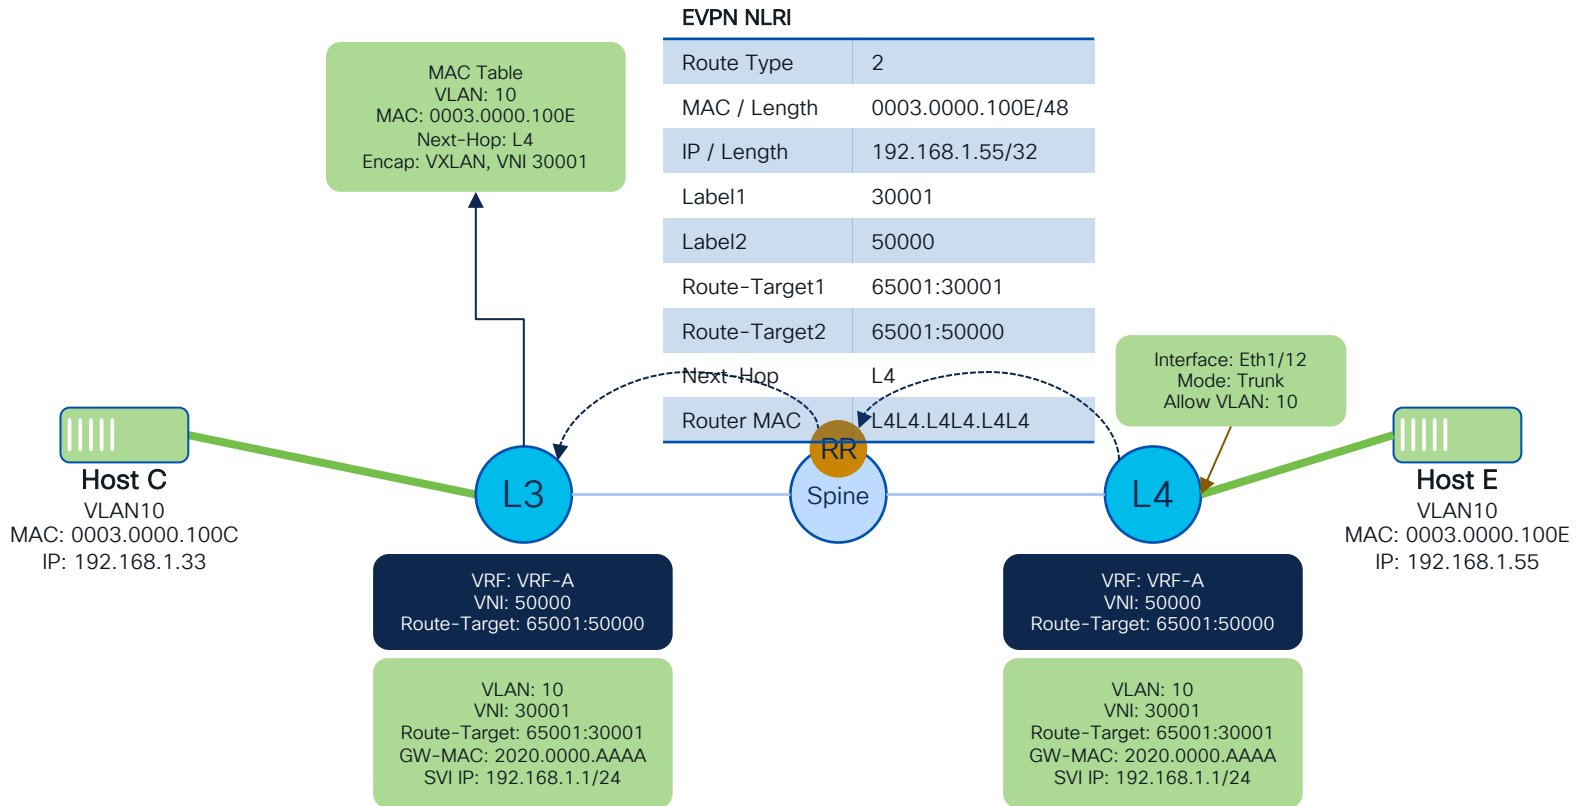


# Learning: HostE to Leaf3

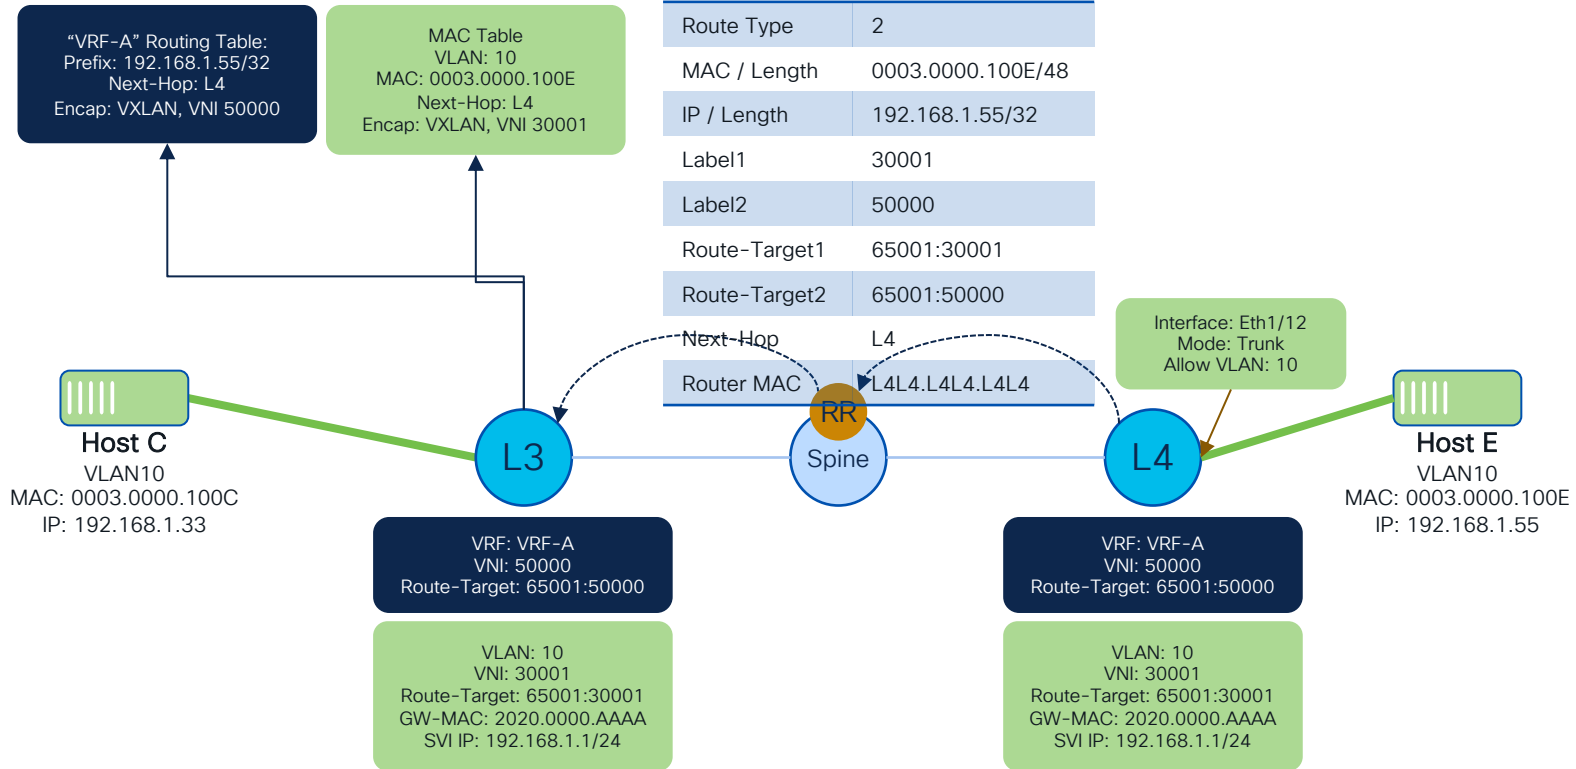
EVPN NLRI	
Route Type	2
MAC / Length	0003.0000.100E/48
IP / Length	192.168.1.55/32
Label1	30001
Label2	50000
Route-Target1	65001:30001
Route-Target2	65001:50000
Next-Hop	L4
Router MAC	L4L4.L4L4.L4L4



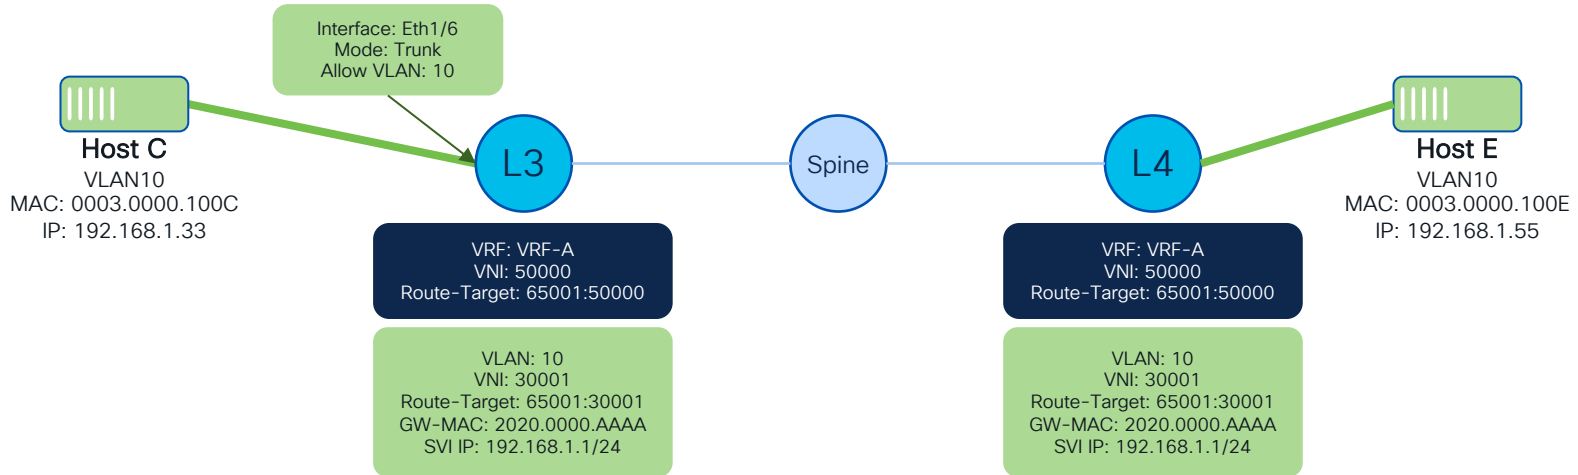
# Learning: HostE to Leaf3



# Learning: HostE to Leaf3

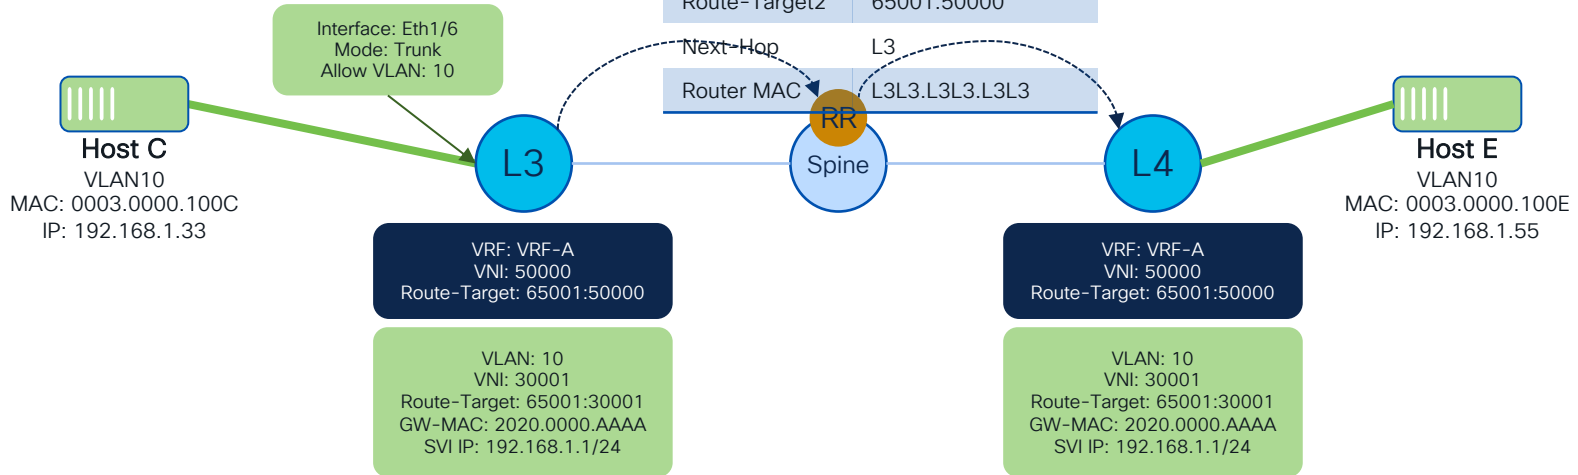


# Learning: HostC to Leaf4



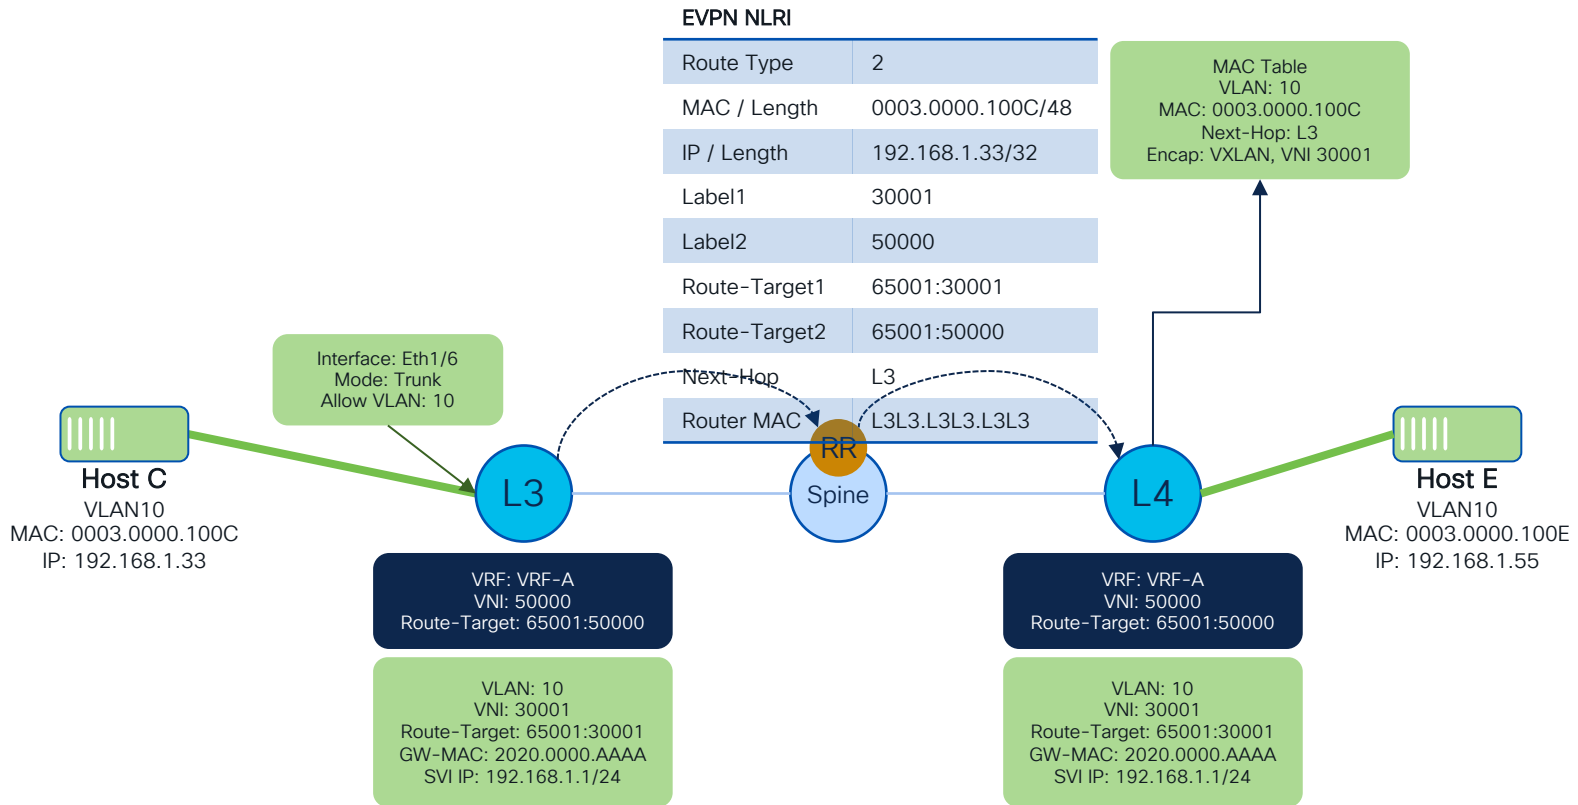
# Learning: HostC to Leaf4

EVPN NLRI	
Route Type	2
MAC / Length	0003.0000.100C/48
IP / Length	192.168.1.33/32
Label1	30001
Label2	50000
Route-Target1	65001:30001
Route-Target2	65001:50000
Next-Hop	L3
Router MAC	L3L3.L3L3.L3L3

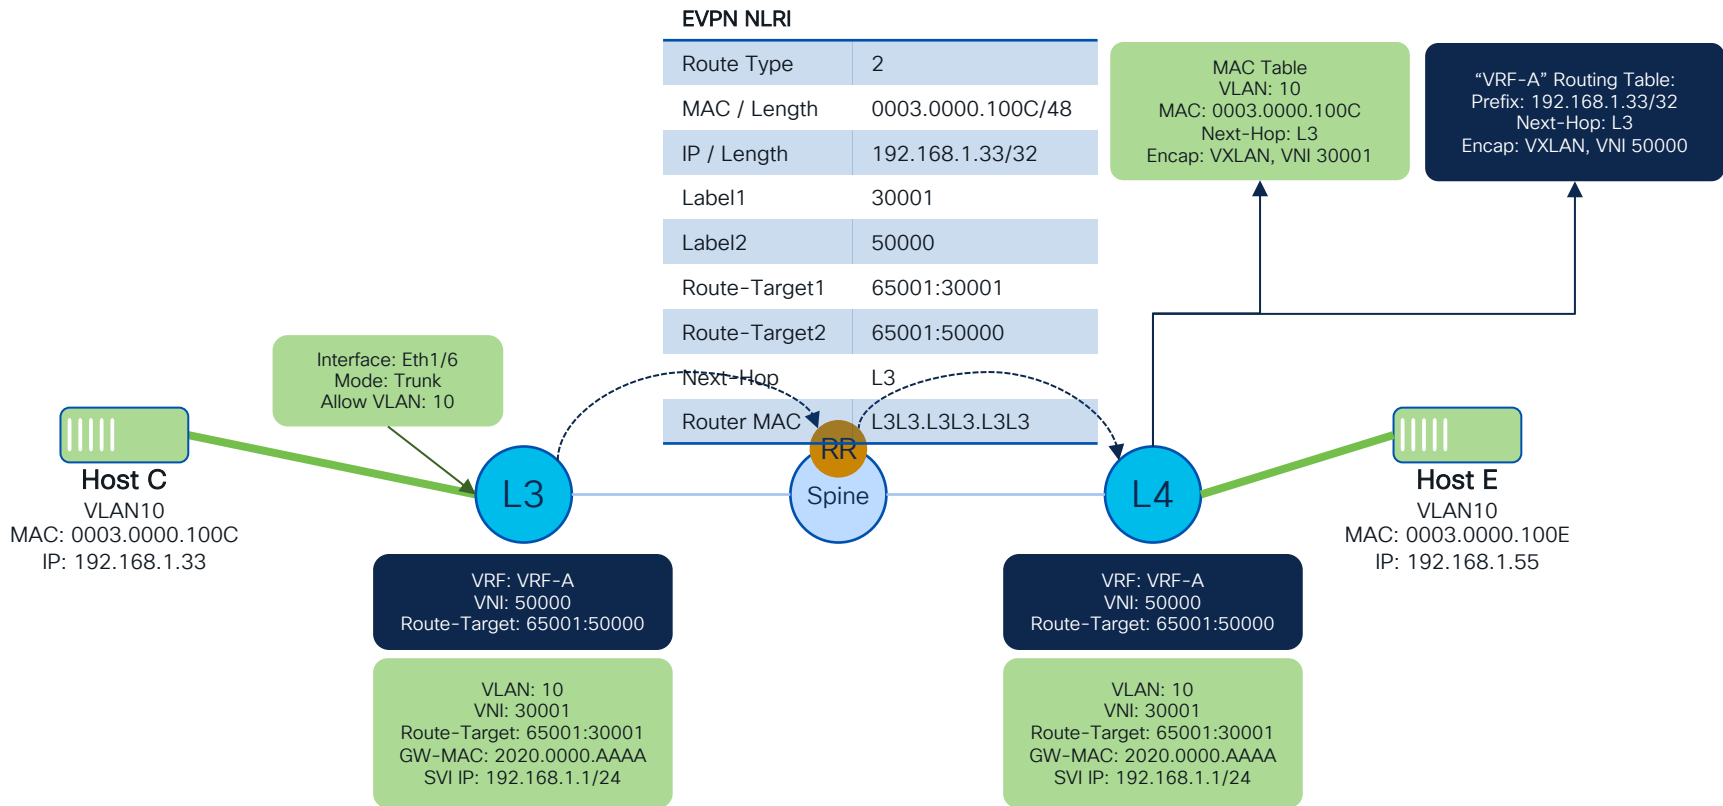




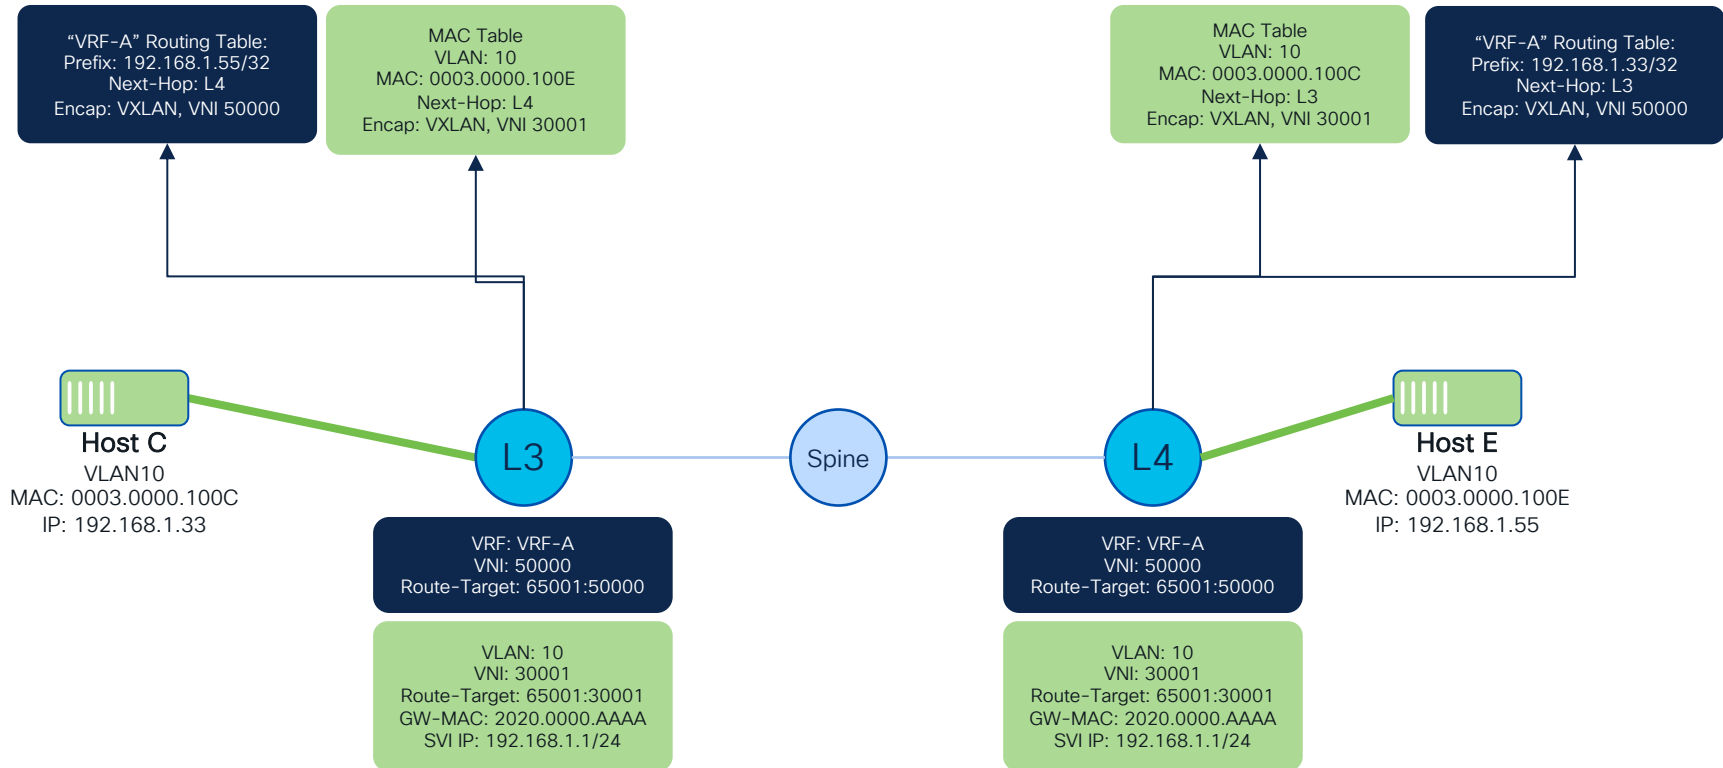
# Learning: HostC to Leaf4



# Learning: HostC to Leaf4




# Forwarding Tables




# HostC to HostE


SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100C	0003.0000.100E	10	192.168.1.33	192.168.1.55	




# HostC to HostE

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100C	0003.0000.100E	10	192.168.1.33	192.168.1.55	

  
**Host C**  
VLAN10  
MAC: 0003.0000.100C  
IP: 192.168.1.33




  
**Host E**  
VLAN10  
MAC: 0003.0000.100E  
IP: 192.168.1.55

SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L3-IP	L4-IP	30001	0003.0000.100C	0003.0000.100E	192.168.1.33	192.168.1.55	

# HostC to HostE

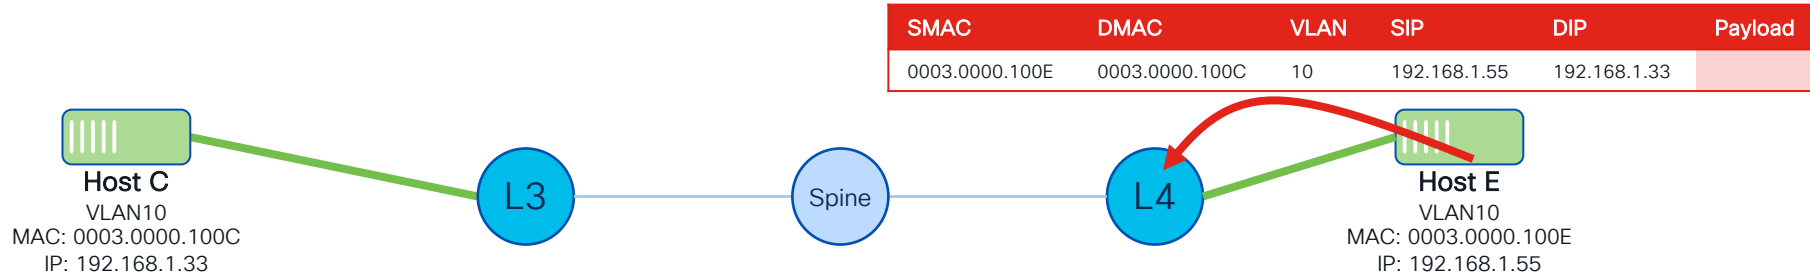
SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100C	0003.0000.100E	10	192.168.1.33	192.168.1.55	

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100C	0003.0000.100E	10	192.168.1.33	192.168.1.55	

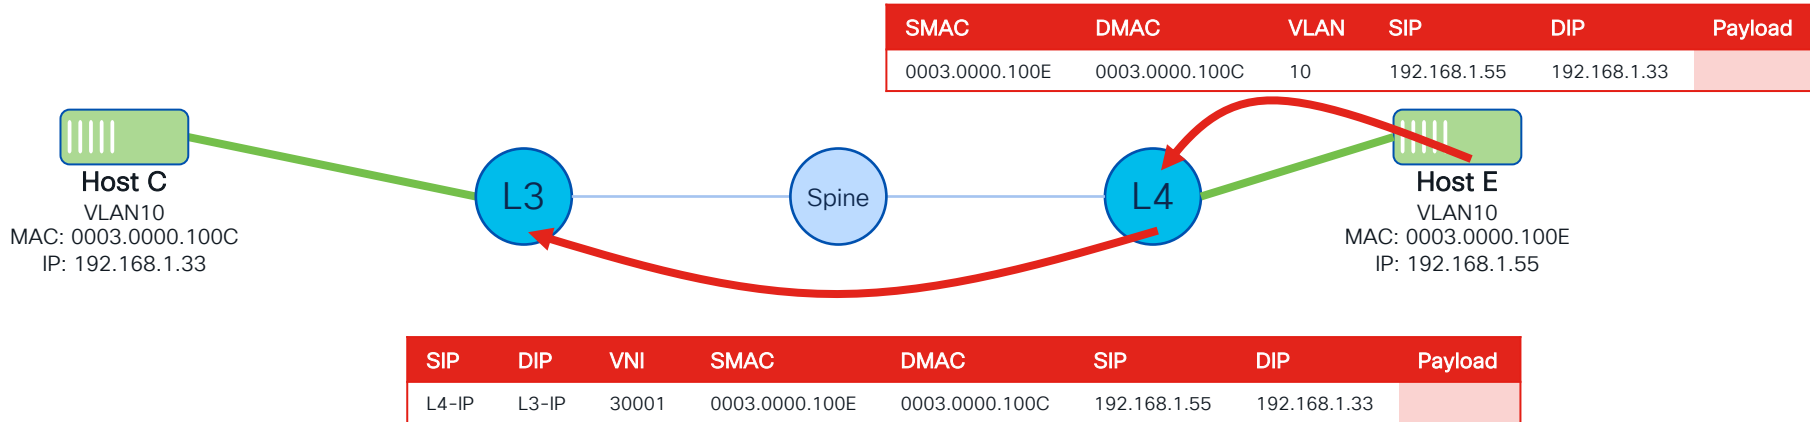


SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L3-IP	L4-IP	30001	0003.0000.100C	0003.0000.100E	192.168.1.33	192.168.1.55	

# HostE to HostC



# HostE to HostC








# HostE to HostC

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100E	0003.0000.100C	10	192.168.1.55	192.168.1.33	

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100E	0003.0000.100C	10	192.168.1.55	192.168.1.33	

  
**Host C**  
 VLAN10  
 MAC: 0003.0000.100C  
 IP: 192.168.1.33



  
**Host E**  
 VLAN10  
 MAC: 0003.0000.100E  
 IP: 192.168.1.55

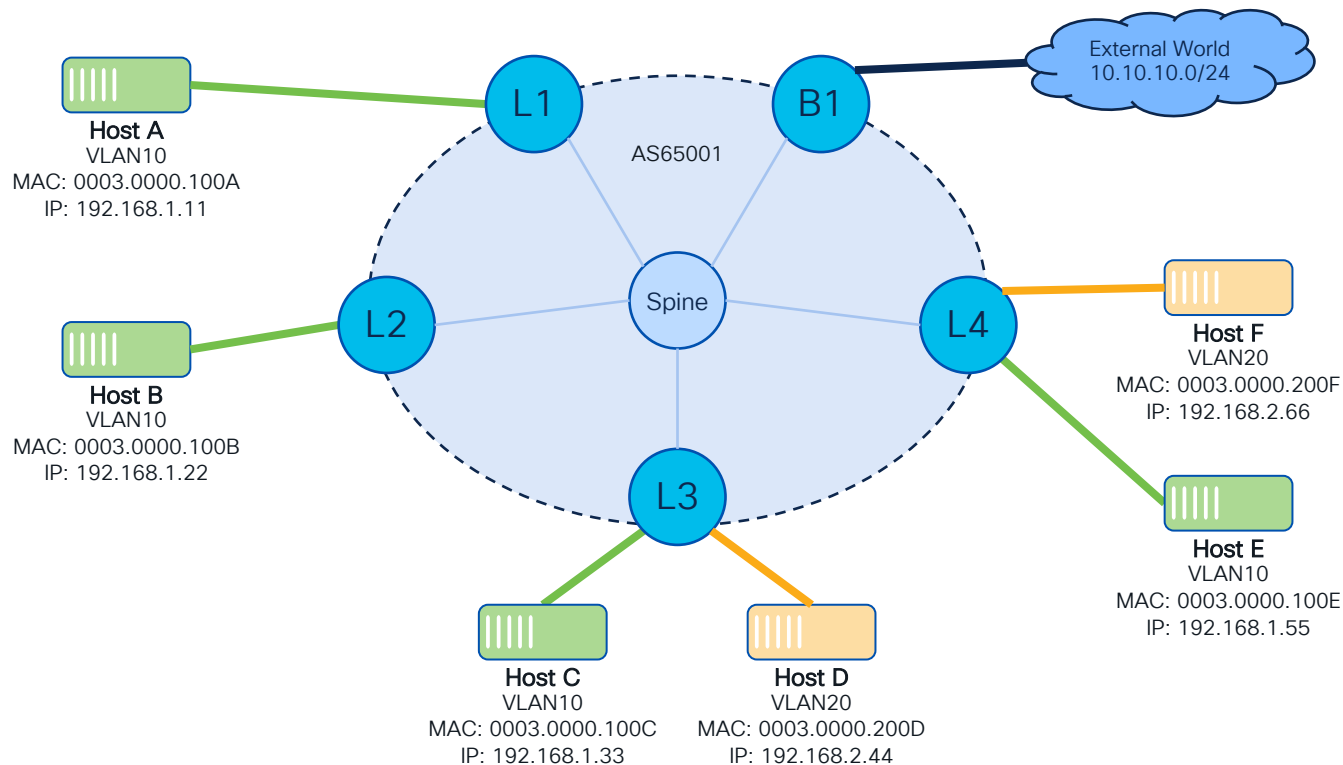
SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L4-IP	L3-IP	30001	0003.0000.100E	0003.0000.100C	192.168.1.55	192.168.1.33	

# Packet Walk: BUM



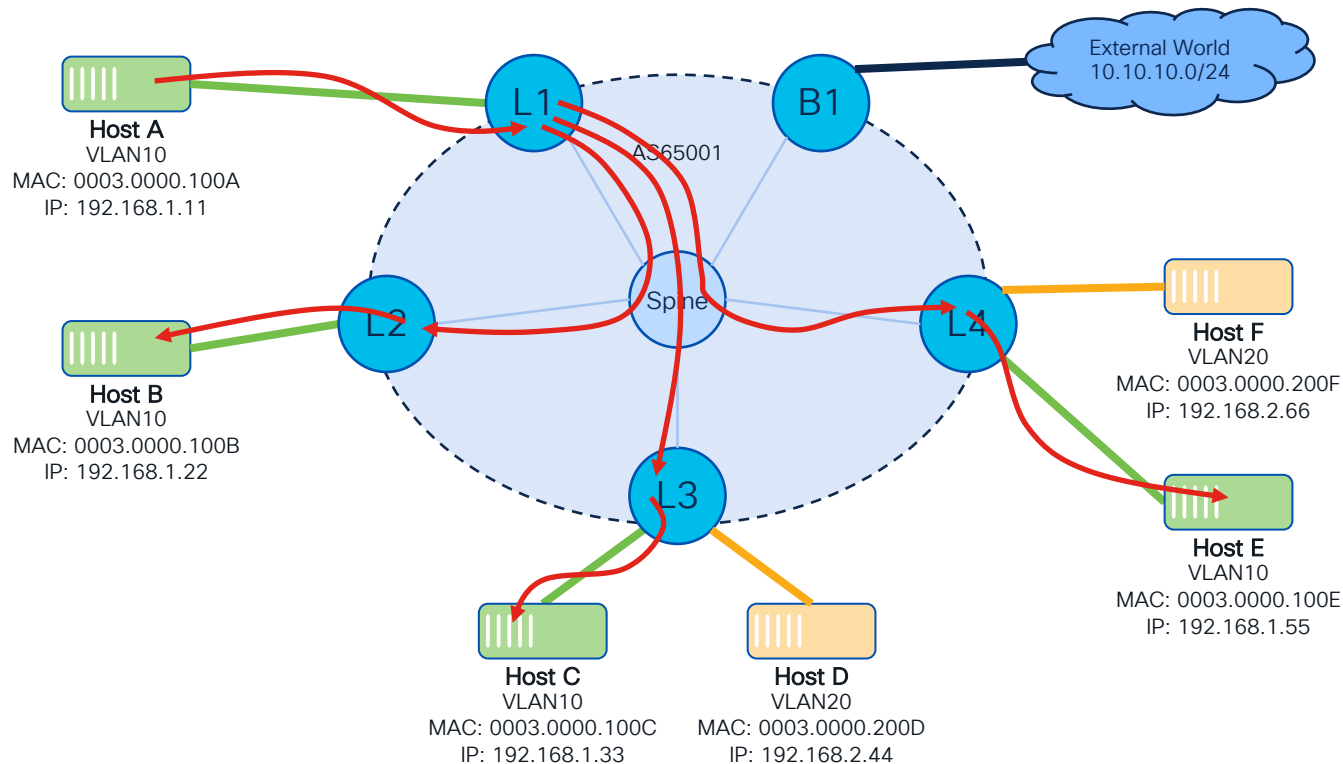
# Topology Overview

## BUM Packet Walk

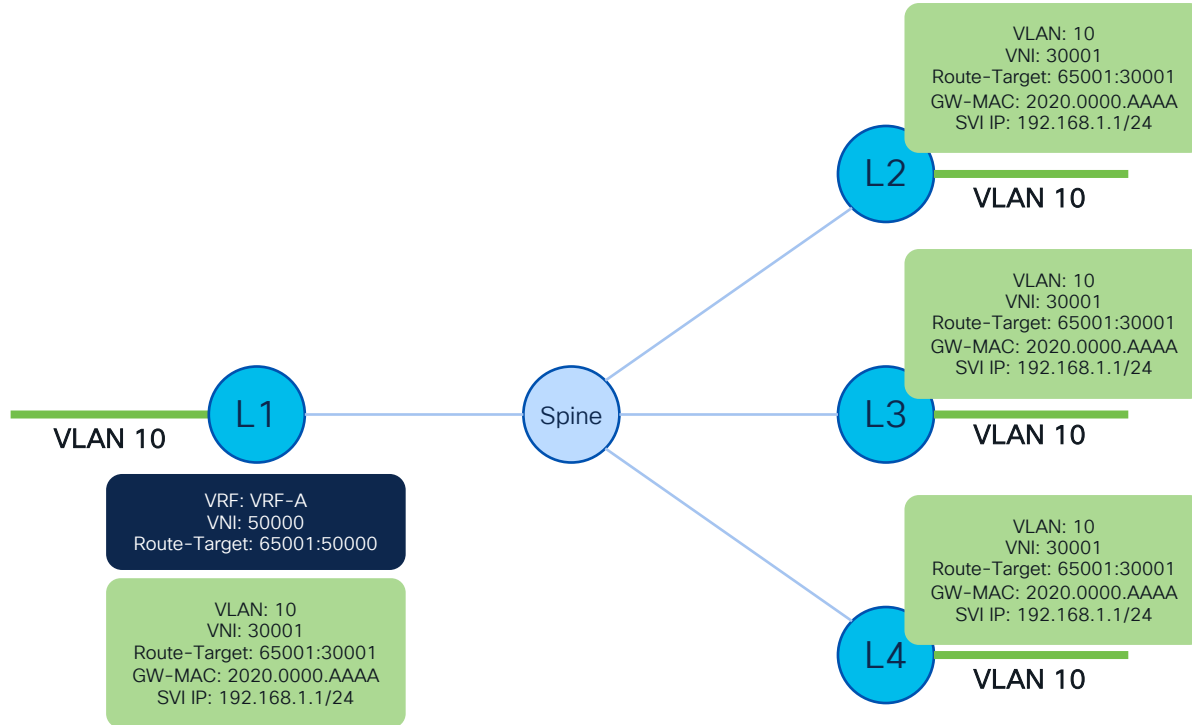


# Topology Overview

## BUM Packet Walk (Ingress Replication)

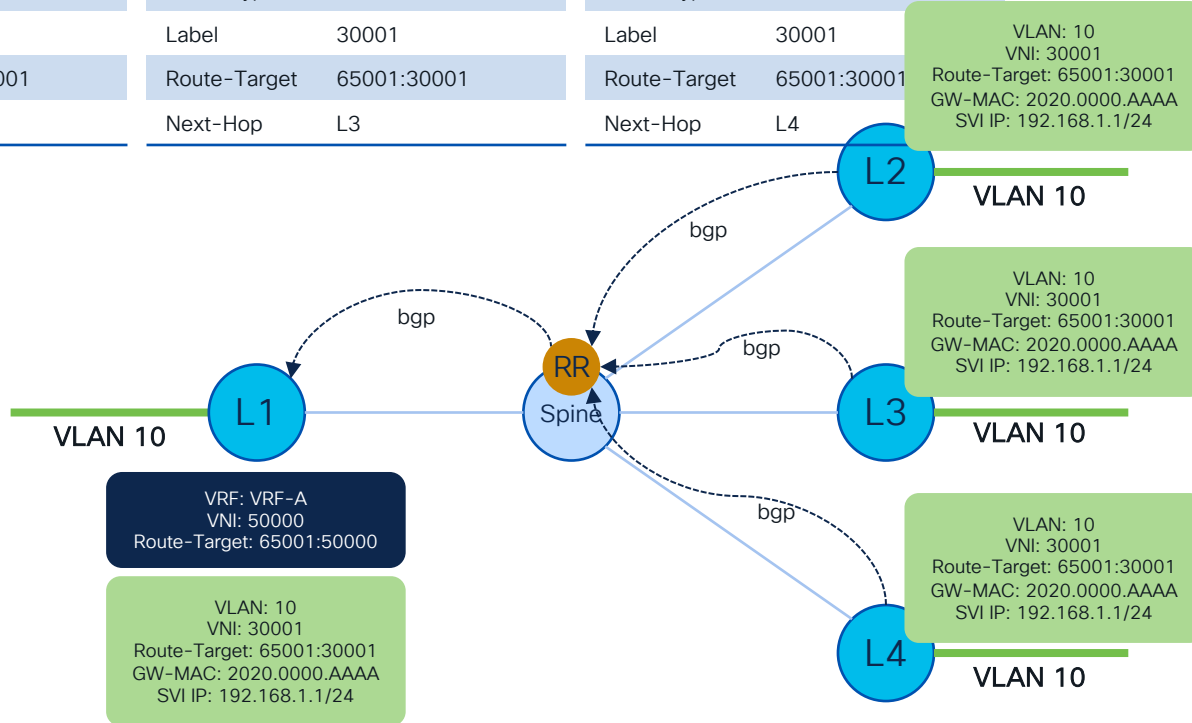


# Learning: VNI 30001 Participation (IR)



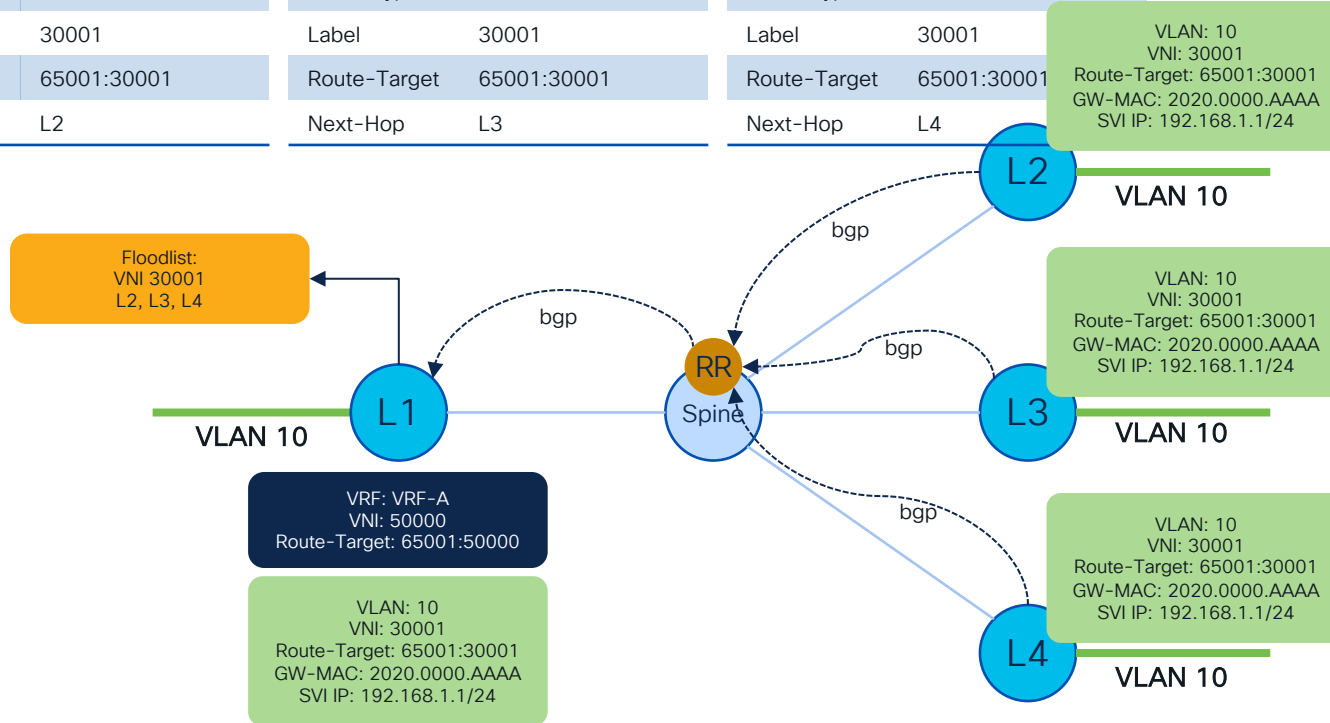
# Learning: VNI 30001 Participation (IR)

EVPN NLRI		EVPN NLRI		EVPN NLRI	
Route Type	3	Route Type	3	Route Type	3
Label	30001	Label	30001	Label	30001
Route-Target	65001:30001	Route-Target	65001:30001	Route-Target	65001:30001
Next-Hop	L2	Next-Hop	L3	Next-Hop	L4

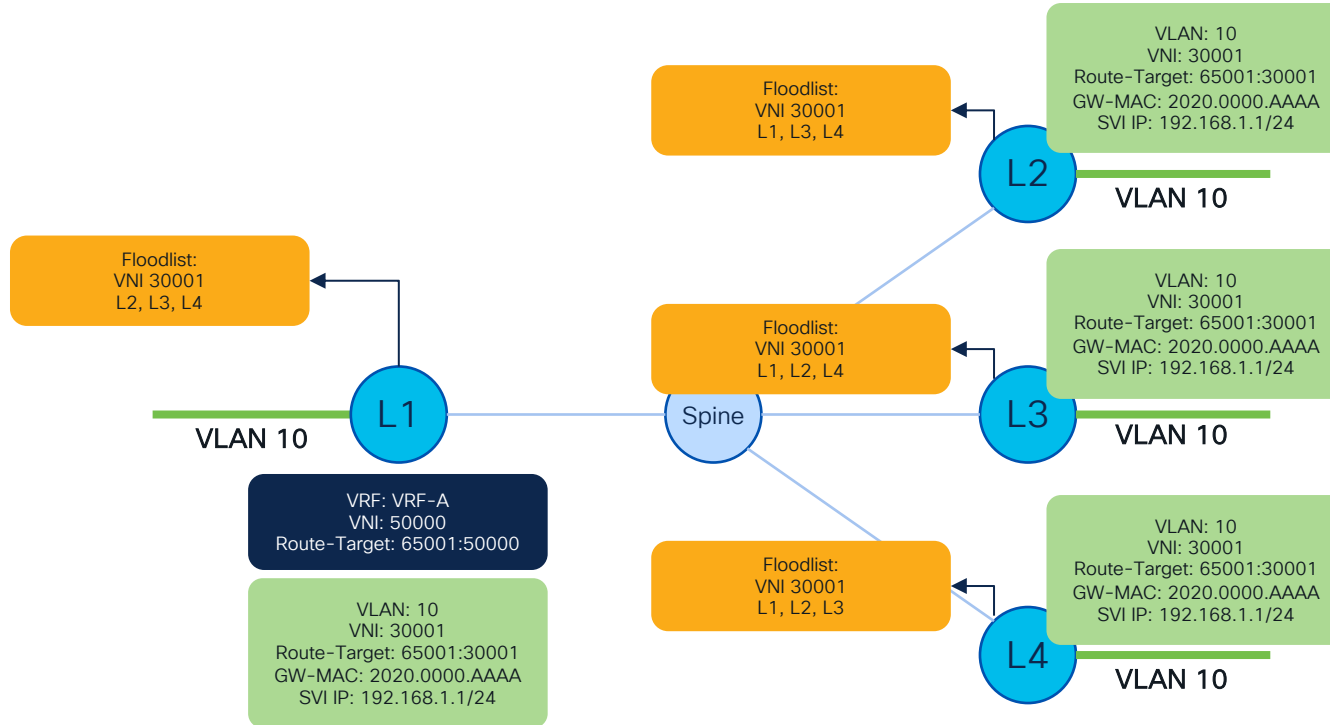


# Learning: VNI 30001 Participation (IR)

EVPN NLRI		EVPN NLRI		EVPN NLRI	
Route Type	3	Route Type	3	Route Type	3
Label	30001	Label	30001	Label	30001
Route-Target	65001:30001	Route-Target	65001:30001	Route-Target	65001:30001
Next-Hop	L2	Next-Hop	L3	Next-Hop	L4



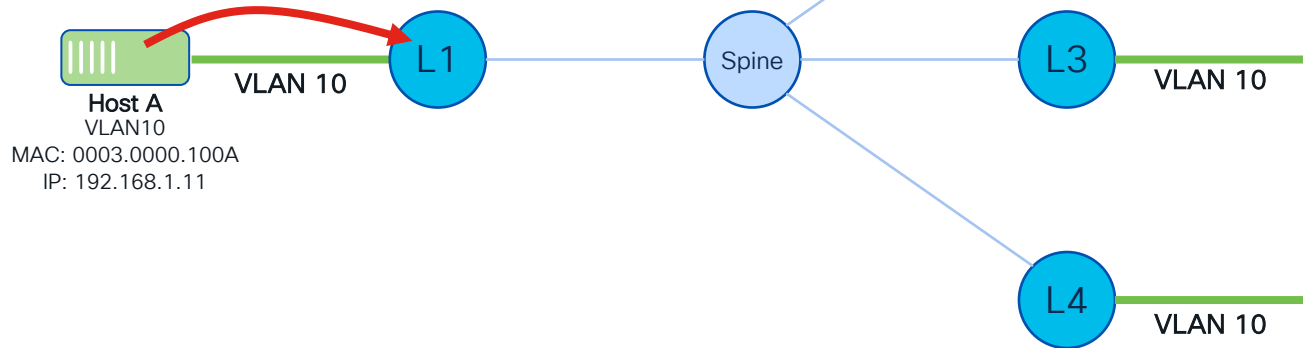
# Forwarding Tables (IR)



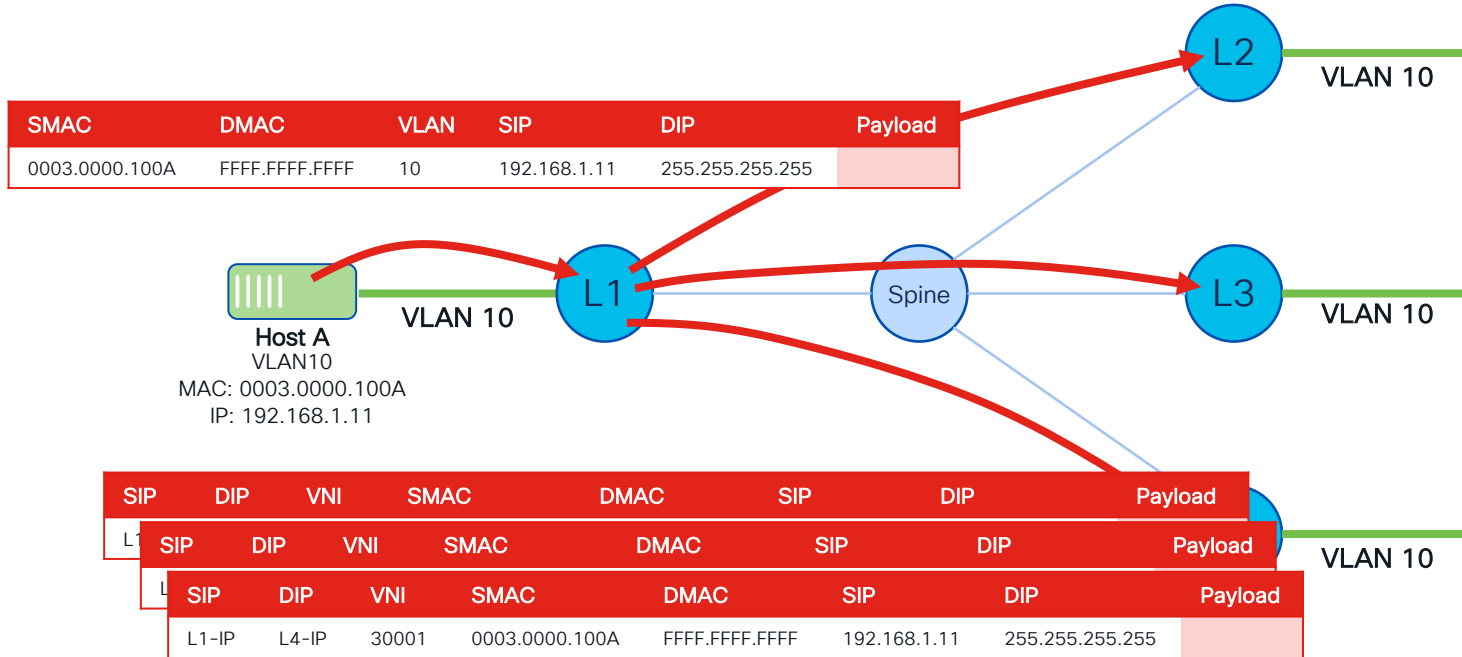


# HostA sends BUM (IR)

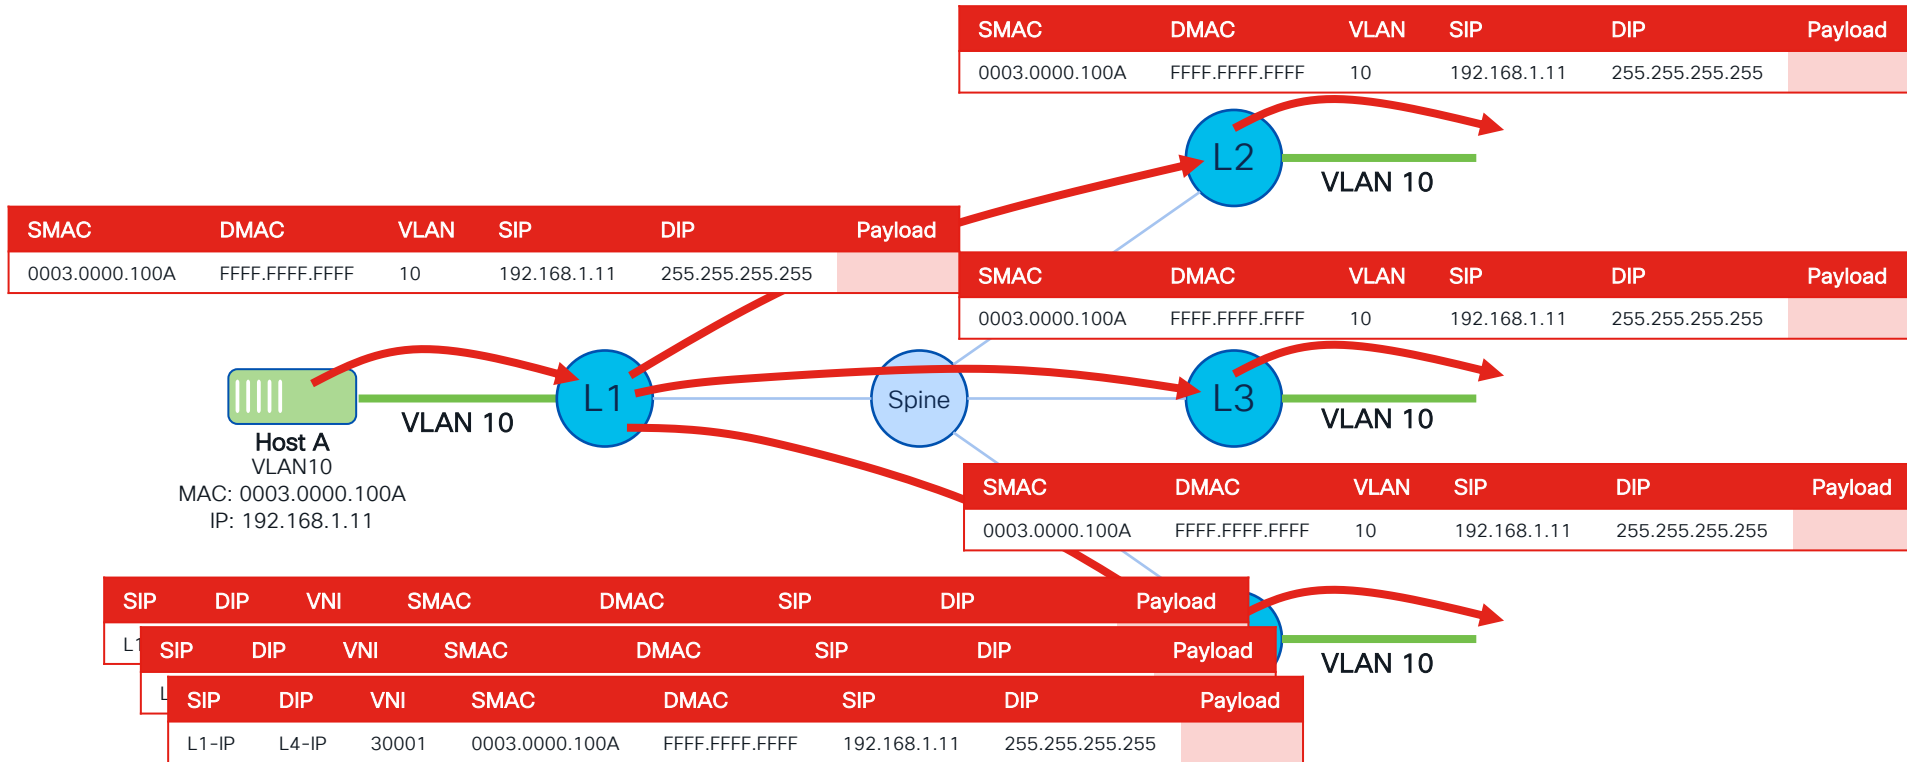
SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	



# HostA sends BUM (IR)

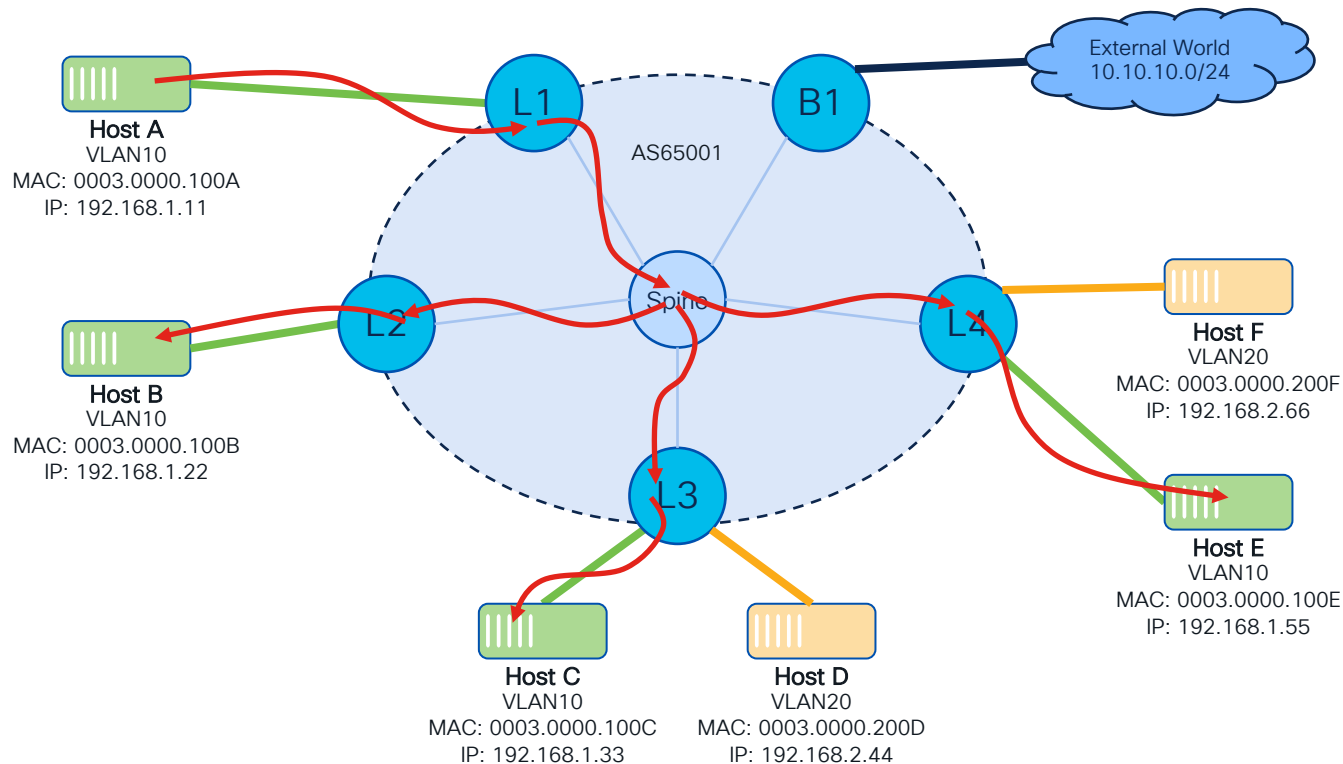


# HostA sends BUM (IR)

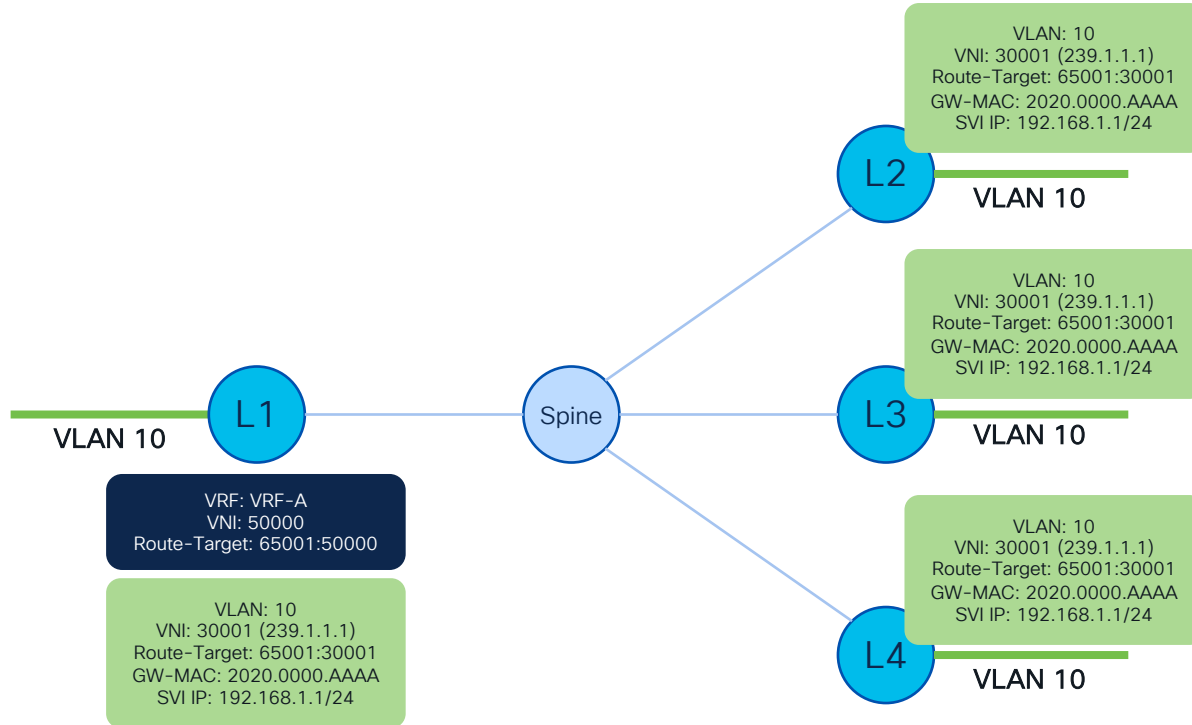


# Topology Overview

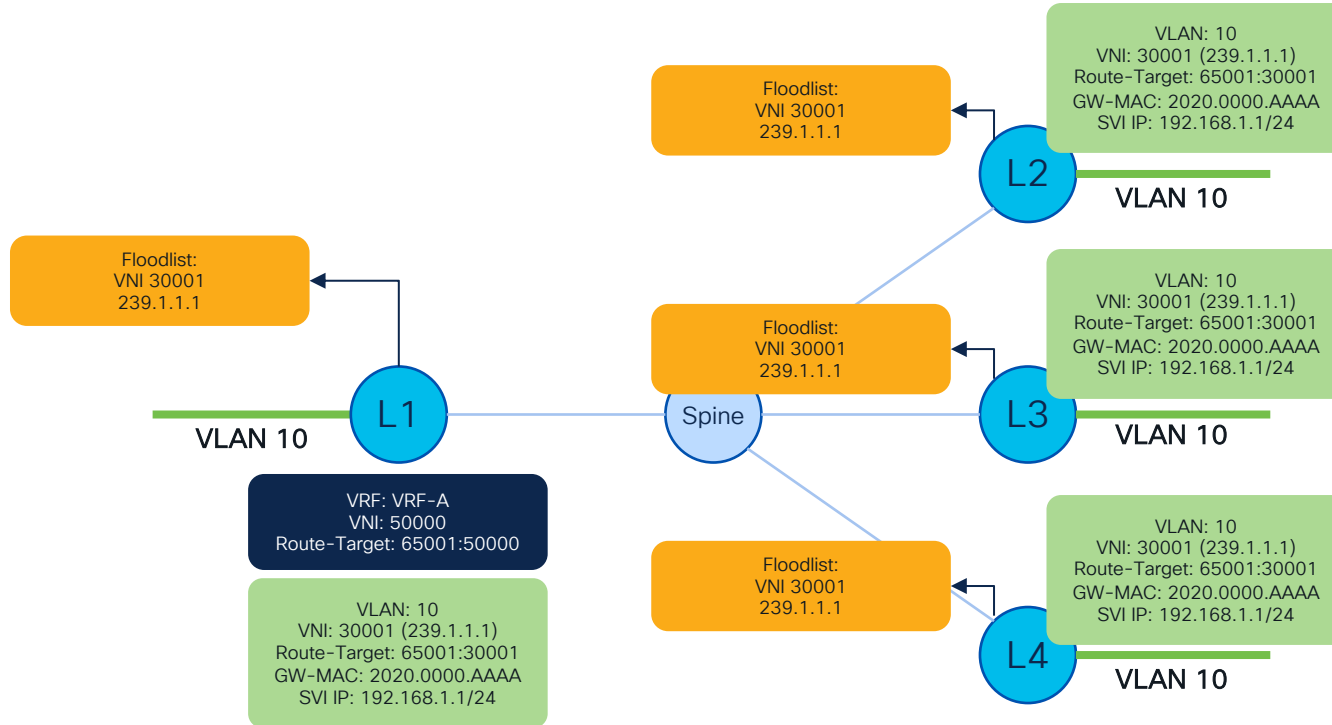
## BUM Packet Walk (Multicast)



# Learning: VNI 30001 Participation (MCAST)




# Forwarding Tables (MCAST)



# HostA sends BUM (MCAST)

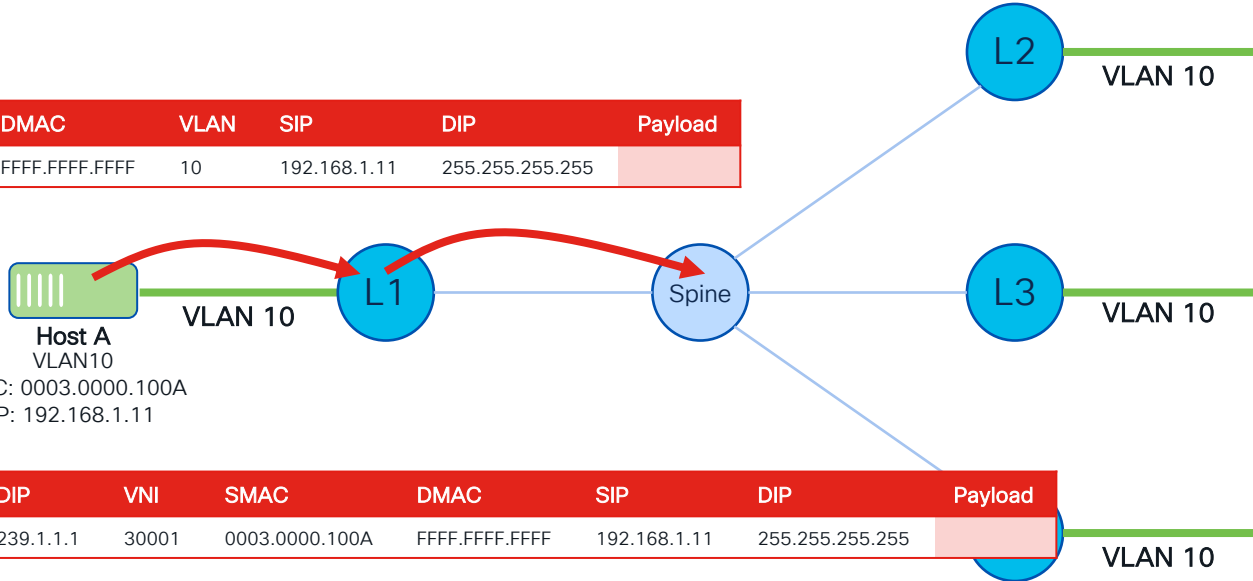
SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	



# HostA sends BUM (MCAST)

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	

Host A  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11



SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	239.1.1.1	30001	0003.0000.100A	FFFF.FFFF.FFFF	192.168.1.11	255.255.255.255	

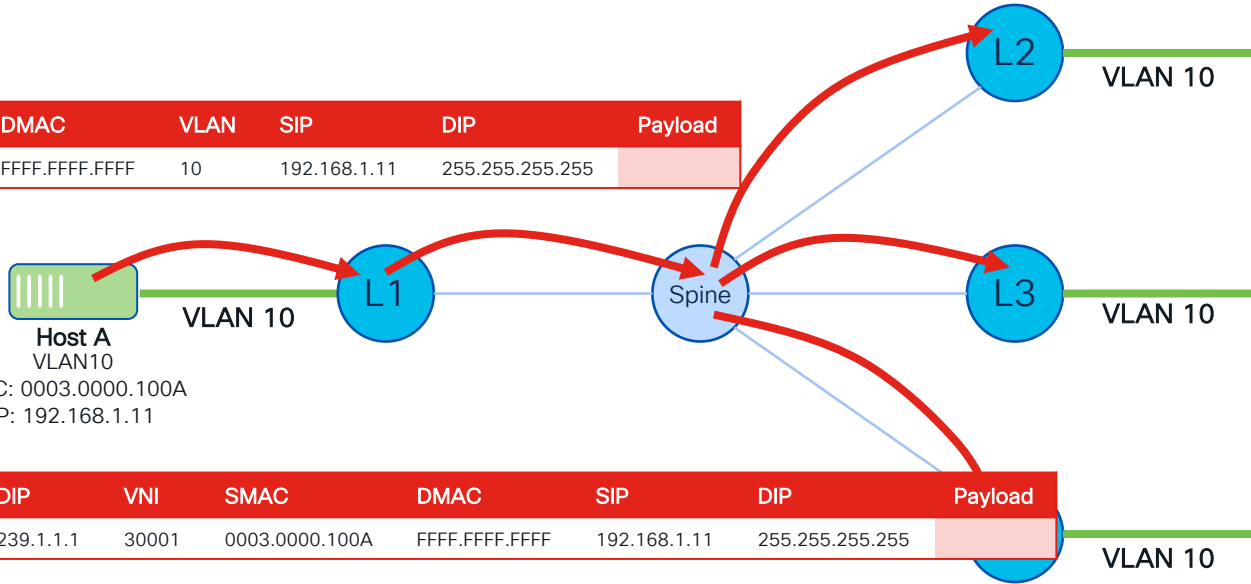


# HostA sends BUM (MCAST)

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	

Host A  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11

SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	239.1.1.1	30001	0003.0000.100A	FFFF.FFFF.FFFF	192.168.1.11	255.255.255.255	



# HostA sends BUM (MCAST)

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	

Host A  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11

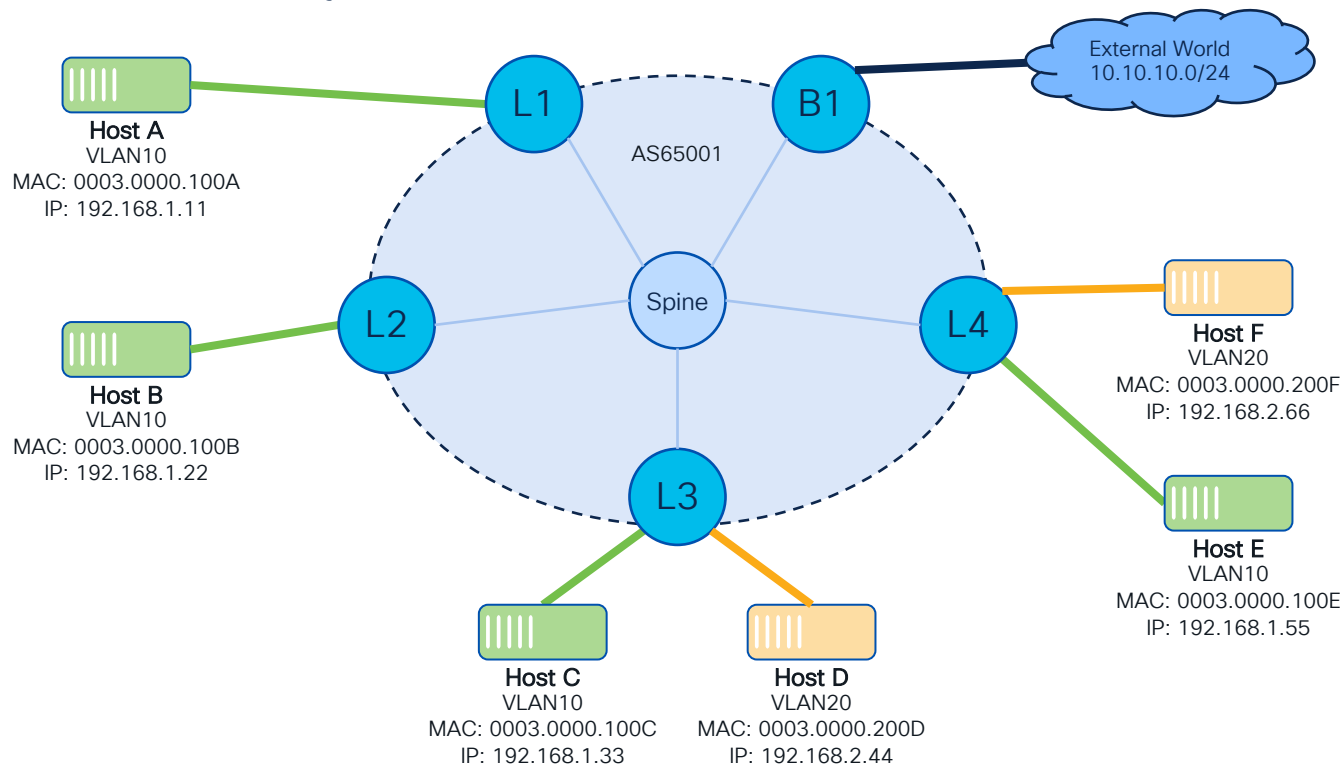
SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	FFFF.FFFF.FFFF	10	192.168.1.11	255.255.255.255	

SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	239.1.1.1	30001	0003.0000.100A	FFFF.FFFF.FFFF	192.168.1.11	255.255.255.255	

# Silent Host Discovery


# Topology Overview

## Silent Host Discovery

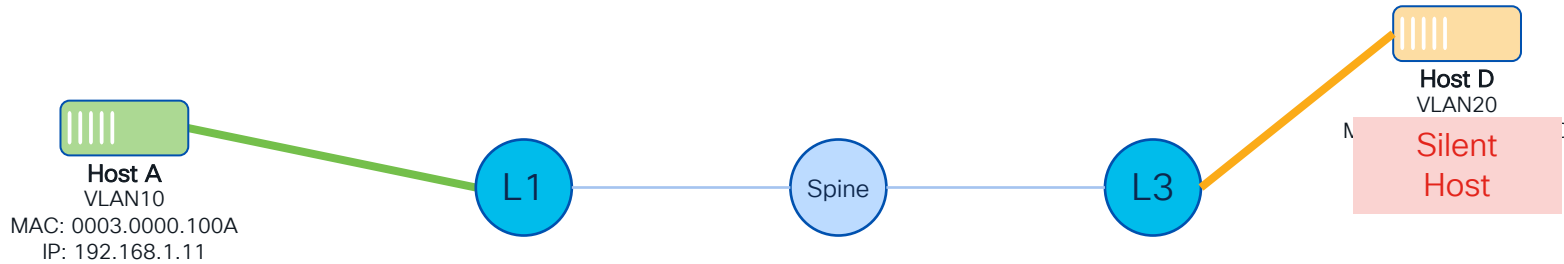


# Topology Overview

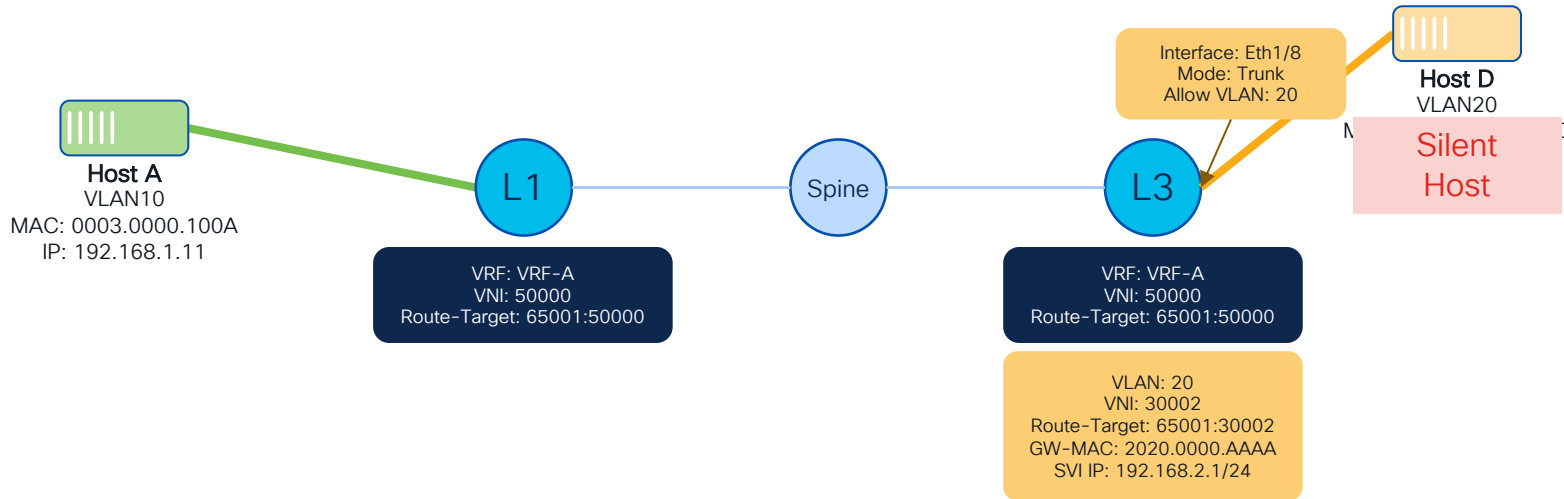
## Silent Host Discovery



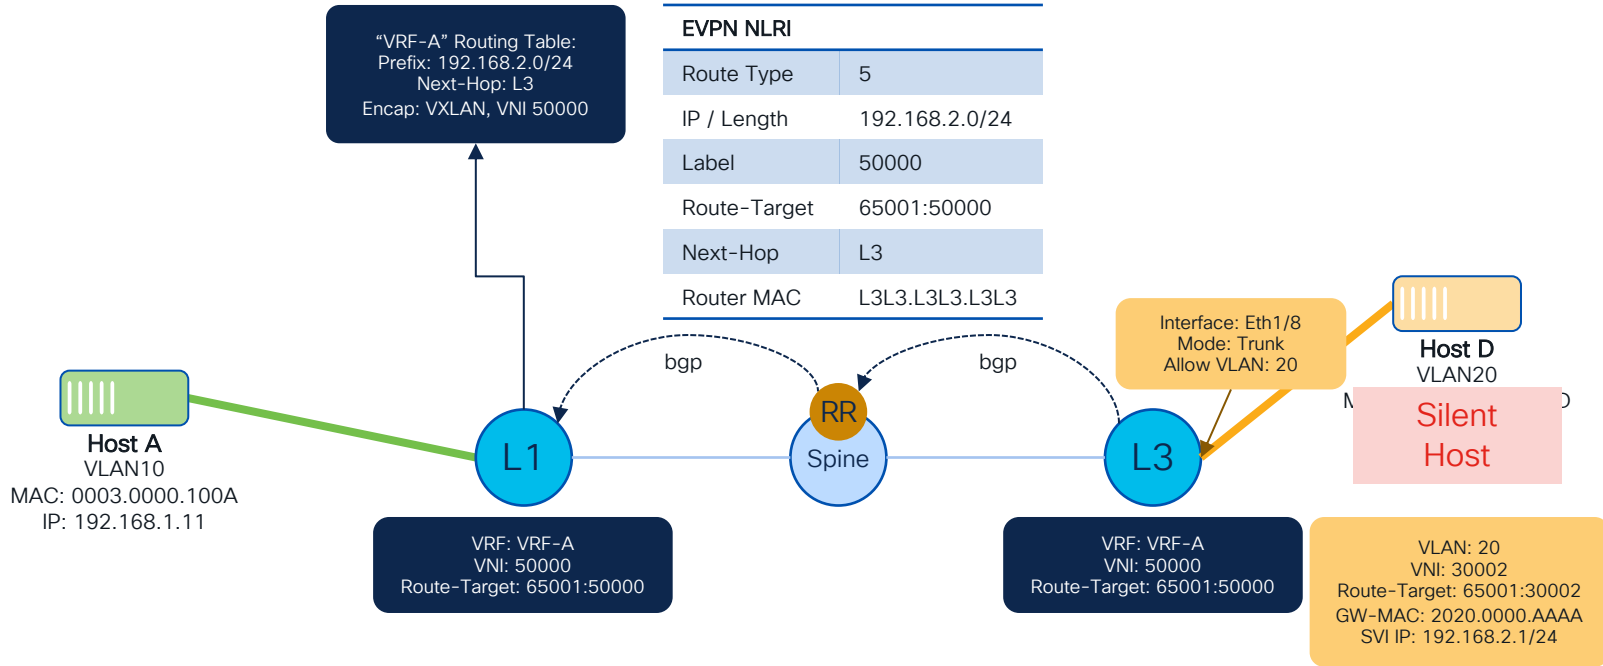
# Learning: HostD to Leaf1



# Learning: HostD to Leaf1

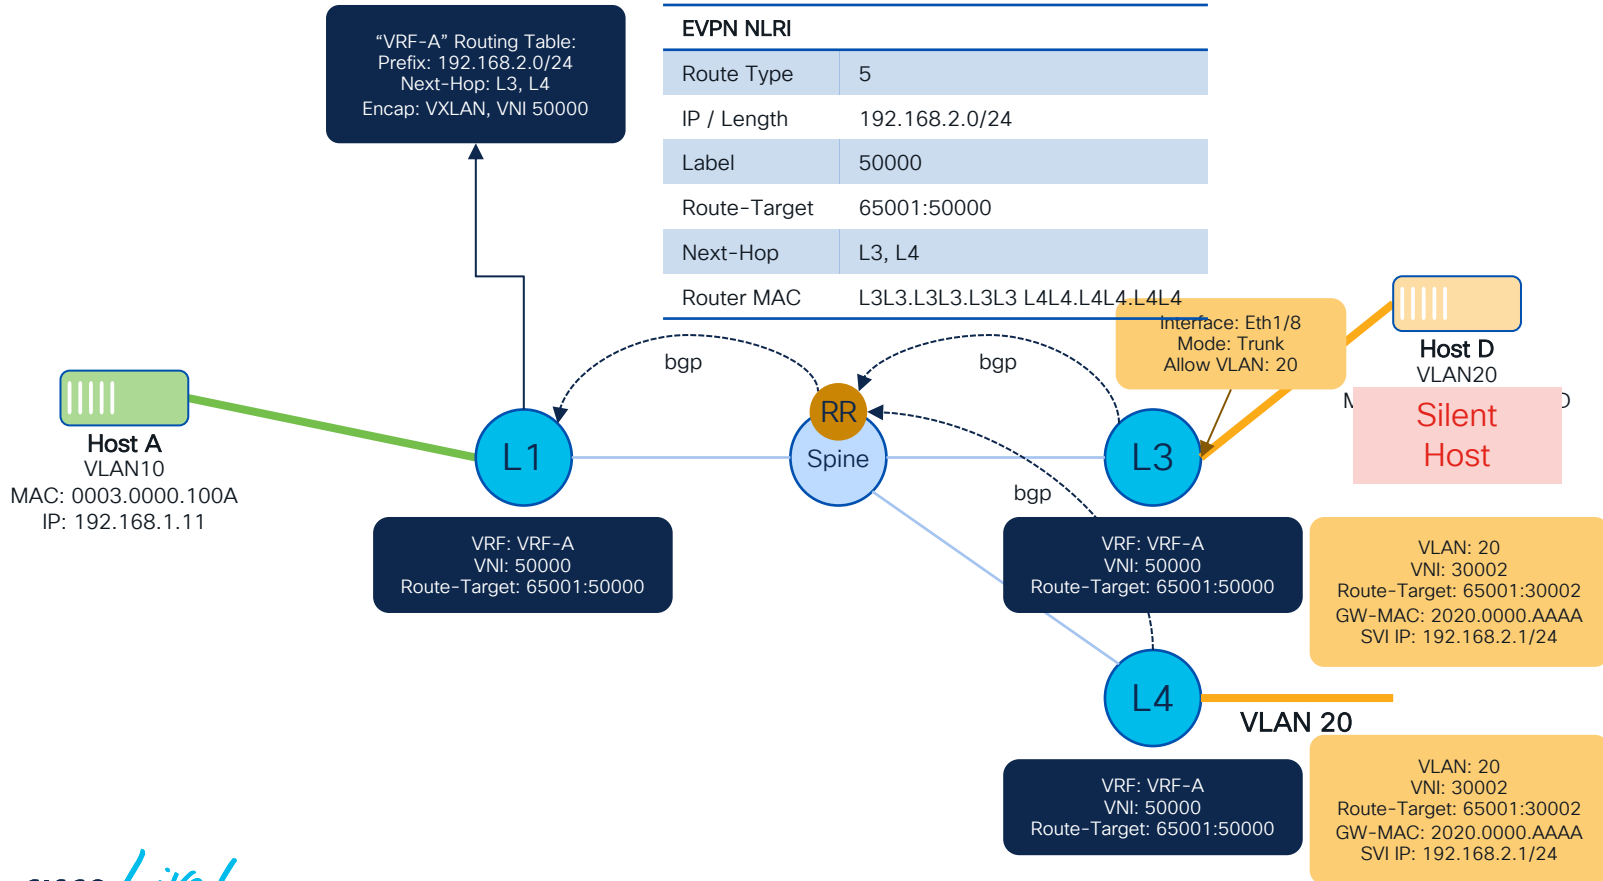


# Learning: HostD to Leaf1

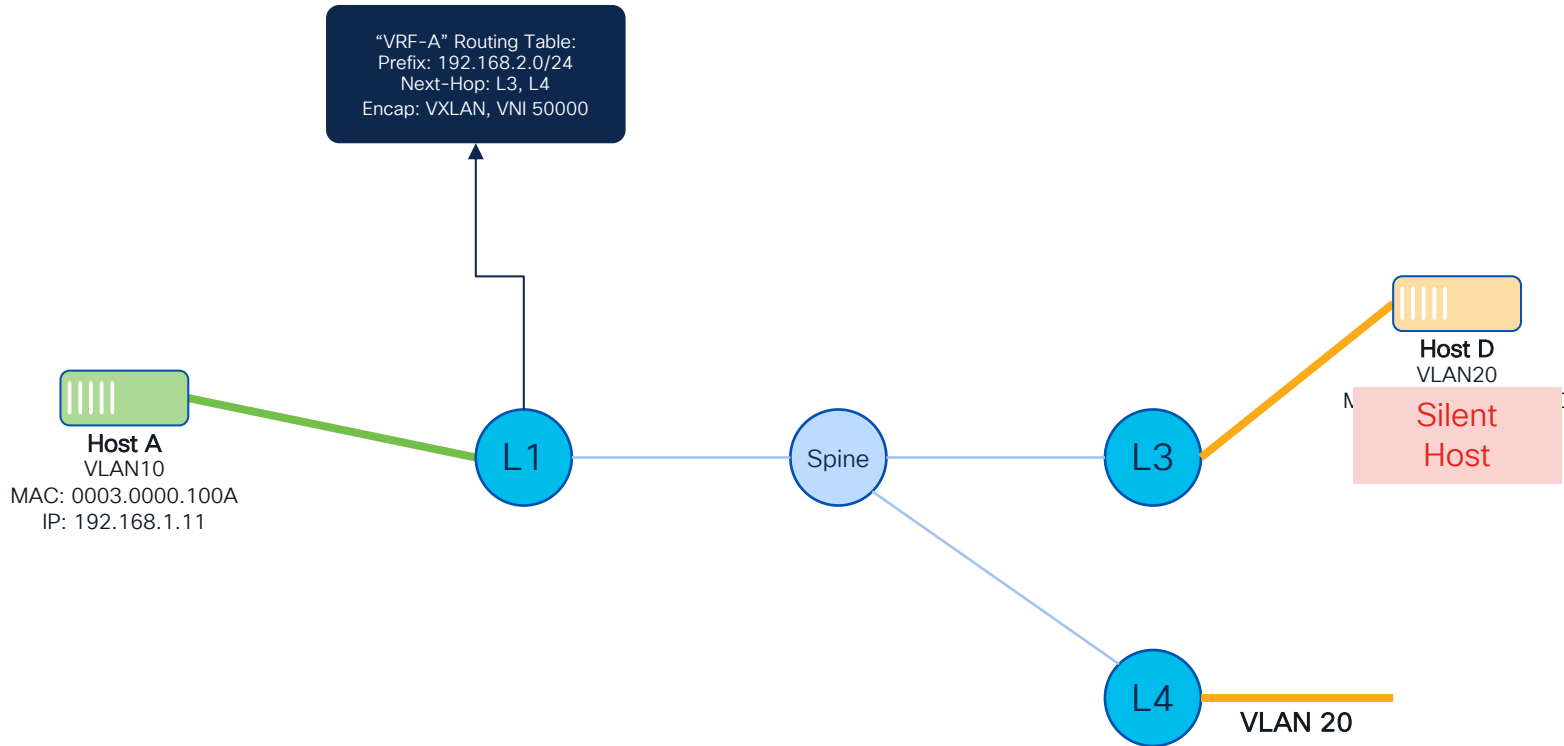




# Learning: HostD to Leaf1








# Forwarding Tables



# HostA to HostD

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.44	

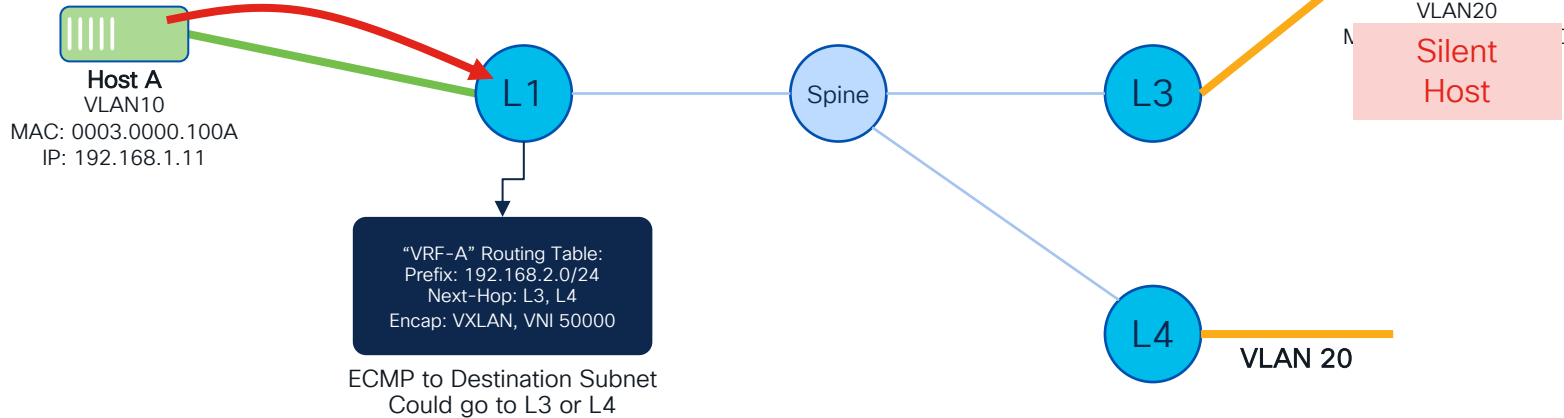
  
**Host A**  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11



VLAN 20


# HostA to HostD



SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.44	



# HostA to HostD

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.44	

  
**Host A**  
 VLAN10  
 MAC: 0003.0000.100A  
 IP: 192.168.1.11



**Host D**  
 VLAN20

Silent  
 Host



VLAN 20

SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	L4-IP	50000	L1-RMAC	L4-RMAC	192.168.1.33	192.168.2.44	

# HostA to HostD

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.44	

Host A  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11



Host D  
VLAN20


Silent  
Host



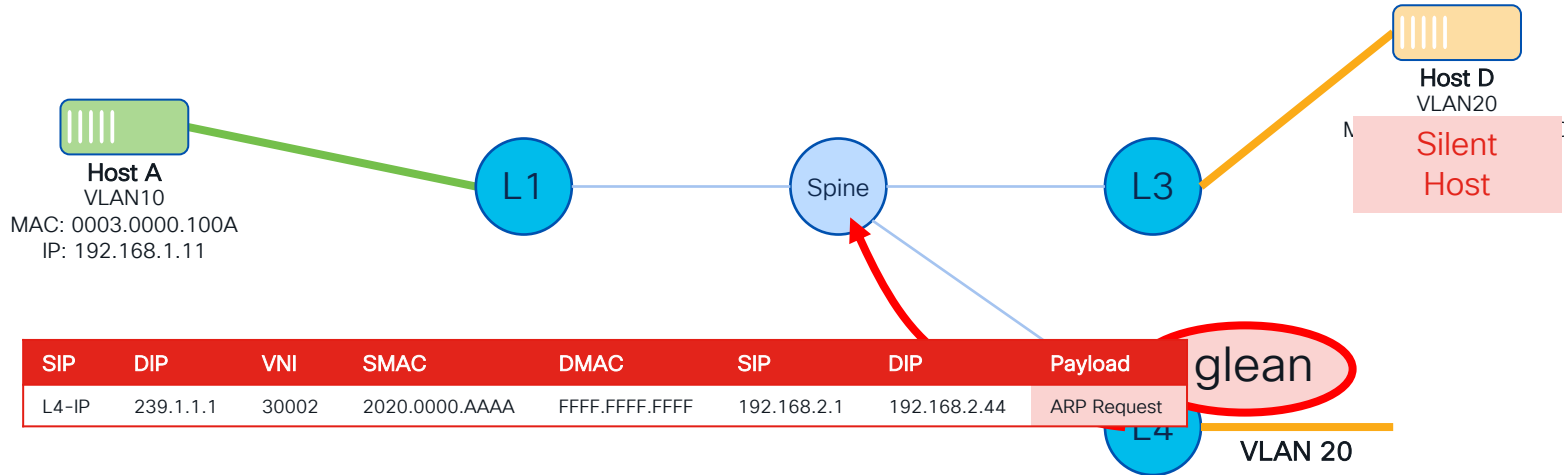
VLAN 20

SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	L4-IP	50000	L1-RMAC	L4-RMAC	192.168.1.33	192.168.2.44	

# HostA to HostD

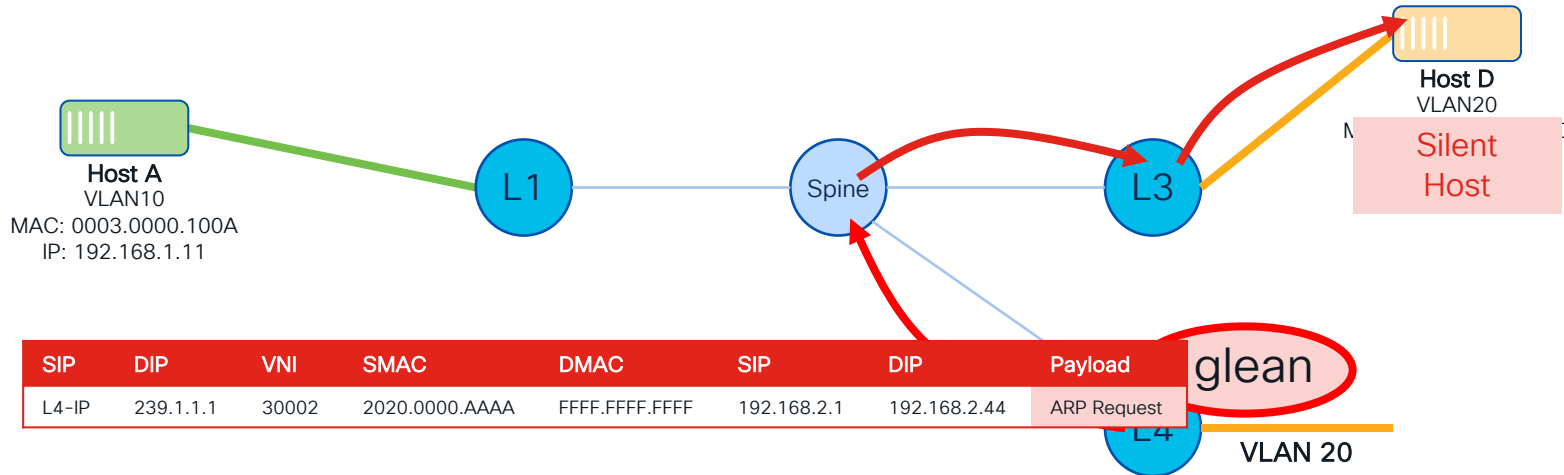


# HostA to HostD

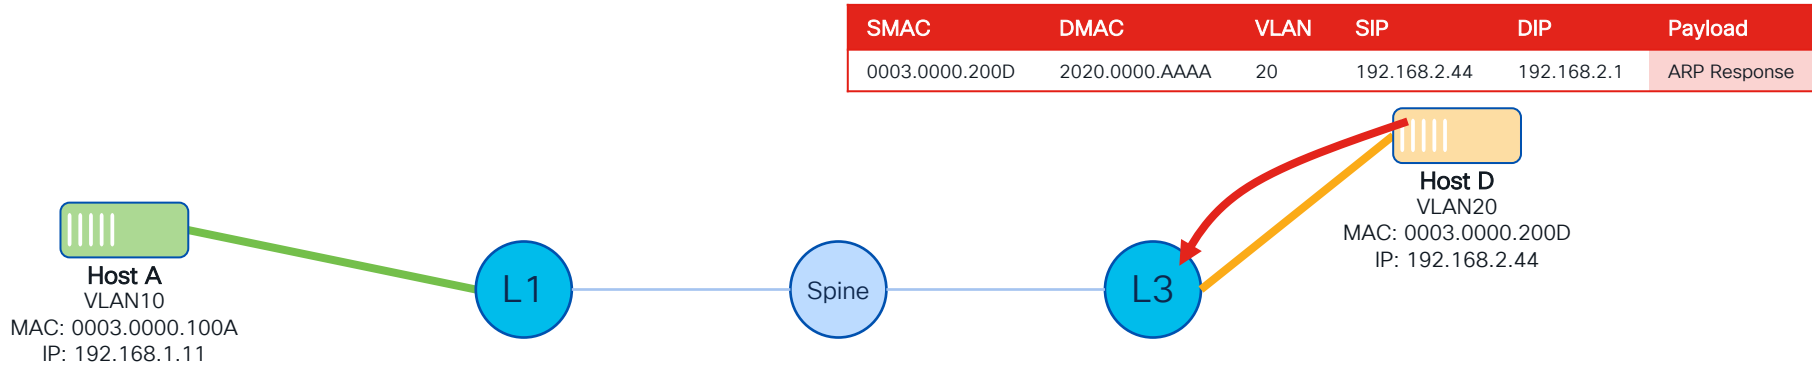




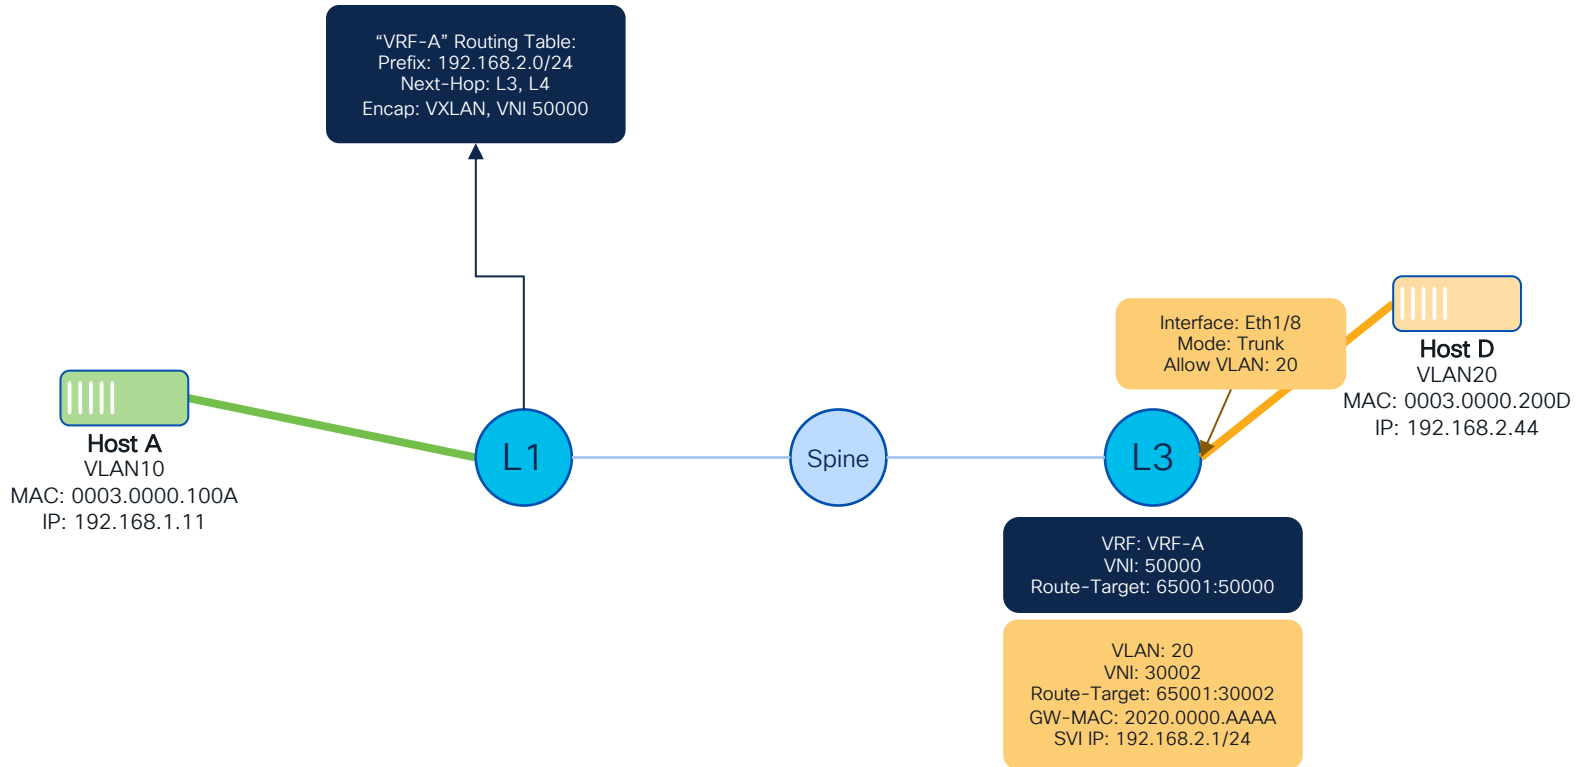
# HostA to HostD



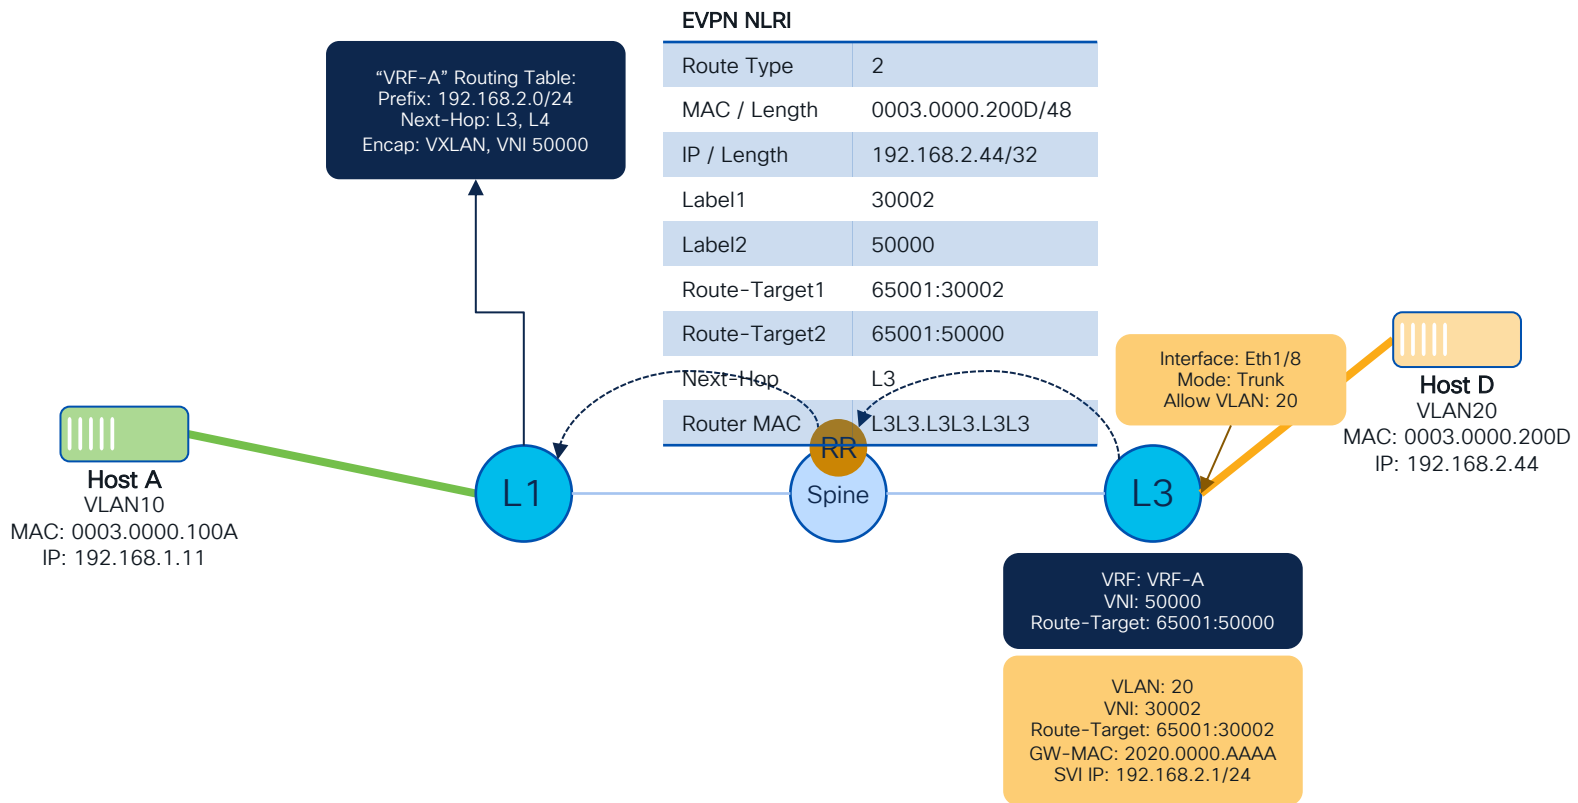
# HostA to HostD



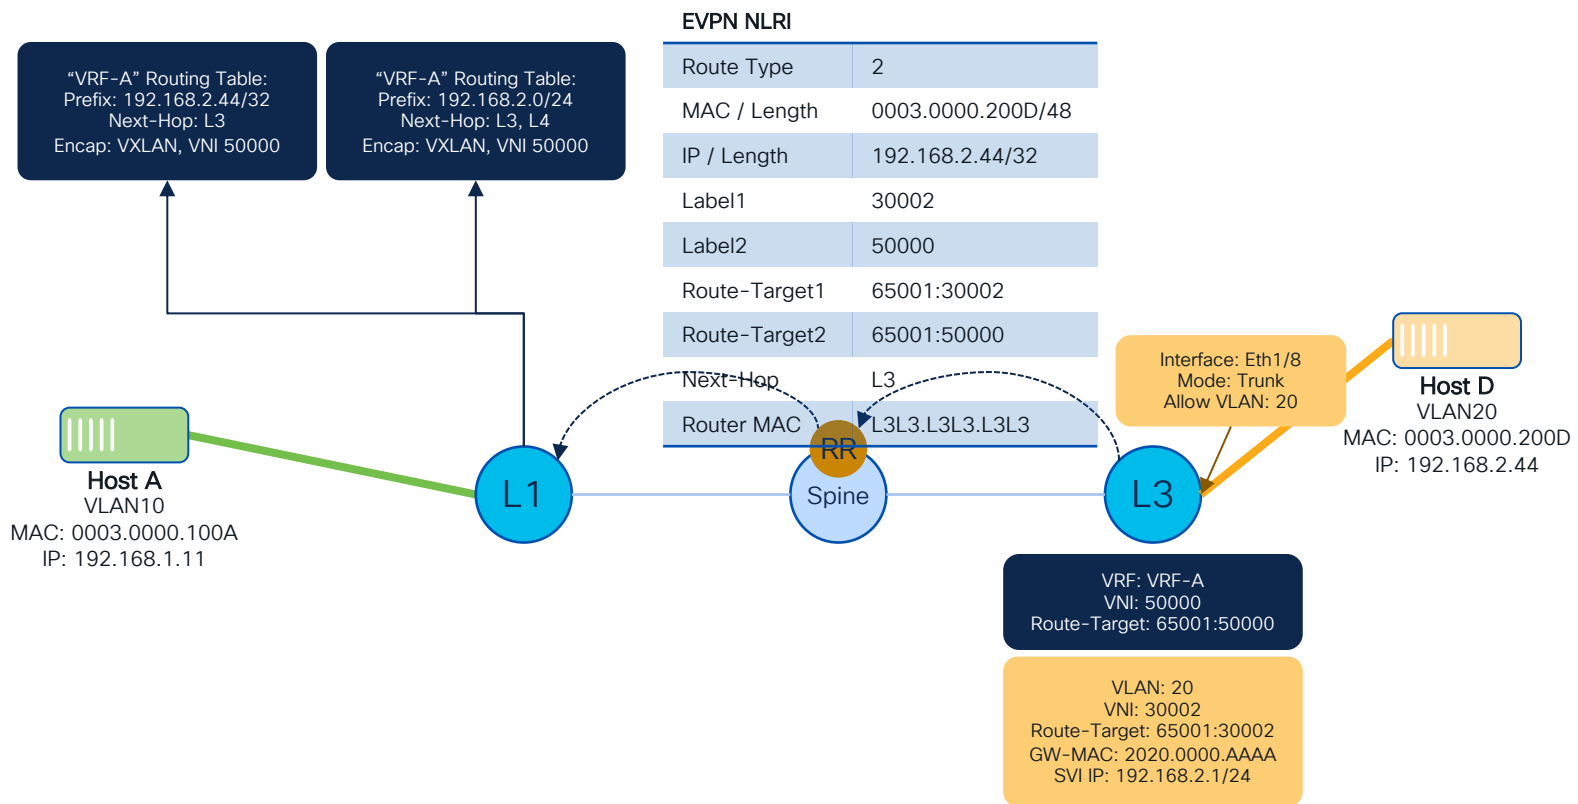
# Learning: HostD to Leaf1



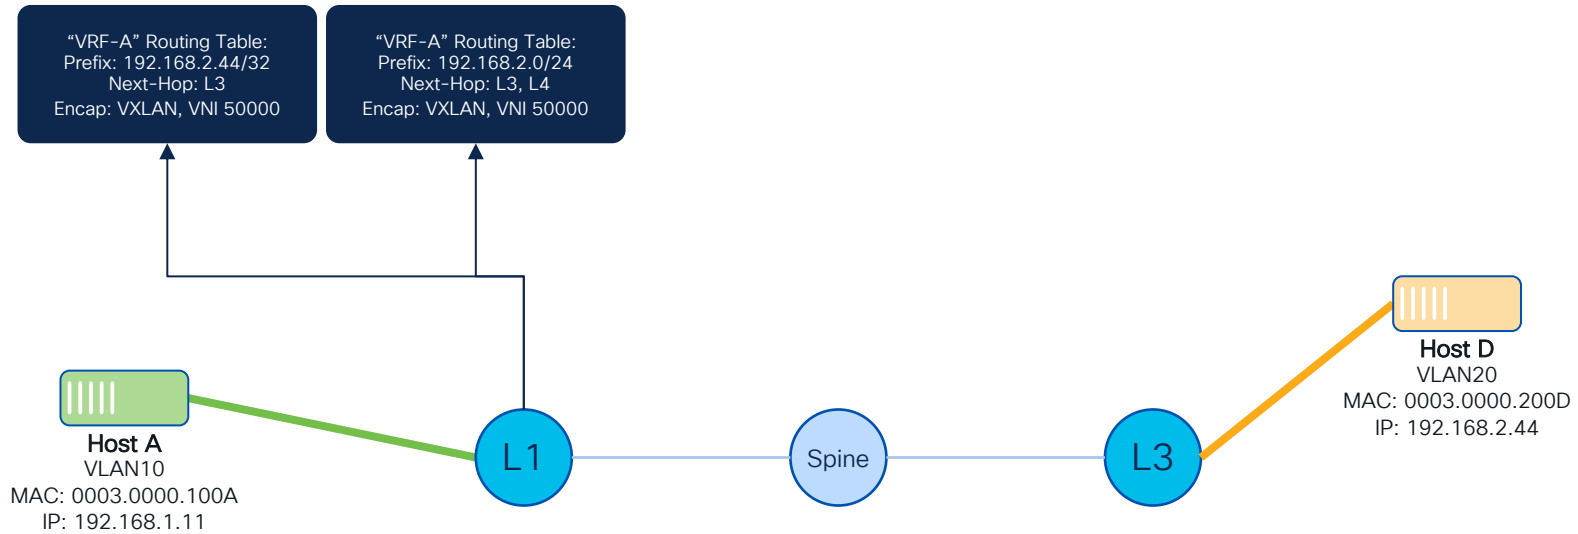
# Learning: HostD to Leaf1



# Learning: HostD to Leaf1







# Forwarding Tables




# HostA to HostD

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.44	

  
**Host A**  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11




  
**Host D**  
VLAN20  
MAC: 0003.0000.200D  
IP: 192.168.2.44

# HostA to HostD

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.44	

Host A  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11



Host D  
VLAN20  
MAC: 0003.0000.200D  
IP: 192.168.2.44


SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	L3-IP	50000	L1-RMAC	L3-RMAC	192.168.1.11	192.168.2.44	



# HostA to HostD

SMAC	DMAC	VLAN	SIP	DIP	Payload
0003.0000.100A	2020.0000.AAAA	10	192.168.1.11	192.168.2.44	

Host A  
VLAN10  
MAC: 0003.0000.100A  
IP: 192.168.1.11



Host D  
VLAN20  
MAC: 0003.0000.200D  
IP: 192.168.2.44

SMAC	DMAC	VLAN	SIP	DIP	Payload
2020.0000.AAAA	0003.0000.200D	20	192.168.1.11	192.168.2.44	

SIP	DIP	VNI	SMAC	DMAC	SIP	DIP	Payload
L1-IP	L3-IP	50000	L1-RMAC	L3-RMAC	192.168.1.11	192.168.2.44	

# Conclusion

# Conclusion

- Did you have enough Packet Walks?
- We covered
  - Host to External Network (RT-5 based routing)
  - Host to Host in different Subnet (RT-2 based routing)
  - Host to Host in same Subnet (RT-2 based bridging)
  - BUM – Broadcast, Unknown Unicast and Multicast (bridged)
    - We looked at Ingress / Head-End Replication and Multicast
    - Note: EVPN works well with BUM forwarding in Multicast (efficiency)
  - Silent Host Discovery (integrated forward and learn)

CISCO *Live!*

# Did you know?

You can have a one-on-one session with a technical expert!

Visit Meet the Expert in The HUB to meet, greet, whiteboard & gain insights about your unique questions with the best of the best.



## Meet the Expert Opening Hours:

<b>Tuesday</b>	<b>3:00pm - 7:00pm</b>
<b>Wednesday</b>	<b>11:15am - 7:00pm</b>
<b>Thursday</b>	<b>9:30am - 4:00pm</b>
<b>Friday</b>	<b>10:30am - 1:30pm</b>

# Session Surveys

We would love to know your feedback on this session!

- Complete a minimum of four session surveys and the overall event surveys to claim a Cisco Live T-Shirt



# Continue your education

CISCO *Live!*

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Expert meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at [www.CiscoLive.com/on-demand](https://www.CiscoLive.com/on-demand)




The bridge to possible

# Thank you

CISCO *Live!*

#CiscoLiveAPJC



The Cisco Live! logo features the word "CISCO" in a bold, black, sans-serif font, followed by "Live!" in a black, cursive script font. The background of the entire image is a vibrant, multi-colored abstract pattern of overlapping, wavy bands in shades of red, orange, yellow, green, and blue, radiating from a bright white center on the right side.

CISCO *Live!*

Let's go

#CiscoLiveAPJC