https://e-vpn.io/fal

EVPN Introduction & Principles

Jiri Chaloupka – Cisco Technical Marketing Engineer

04/2020

Objectives

- Short Technical session (Flood & Learn)
  - No Fee
  - No Registration (Let's see if we will not overload meeting;))
- Networking topics with focus on Service Provider(SP) and SP Data Center technologies

https://e-vpn.io/fal

MPLS Transport & BGP Service

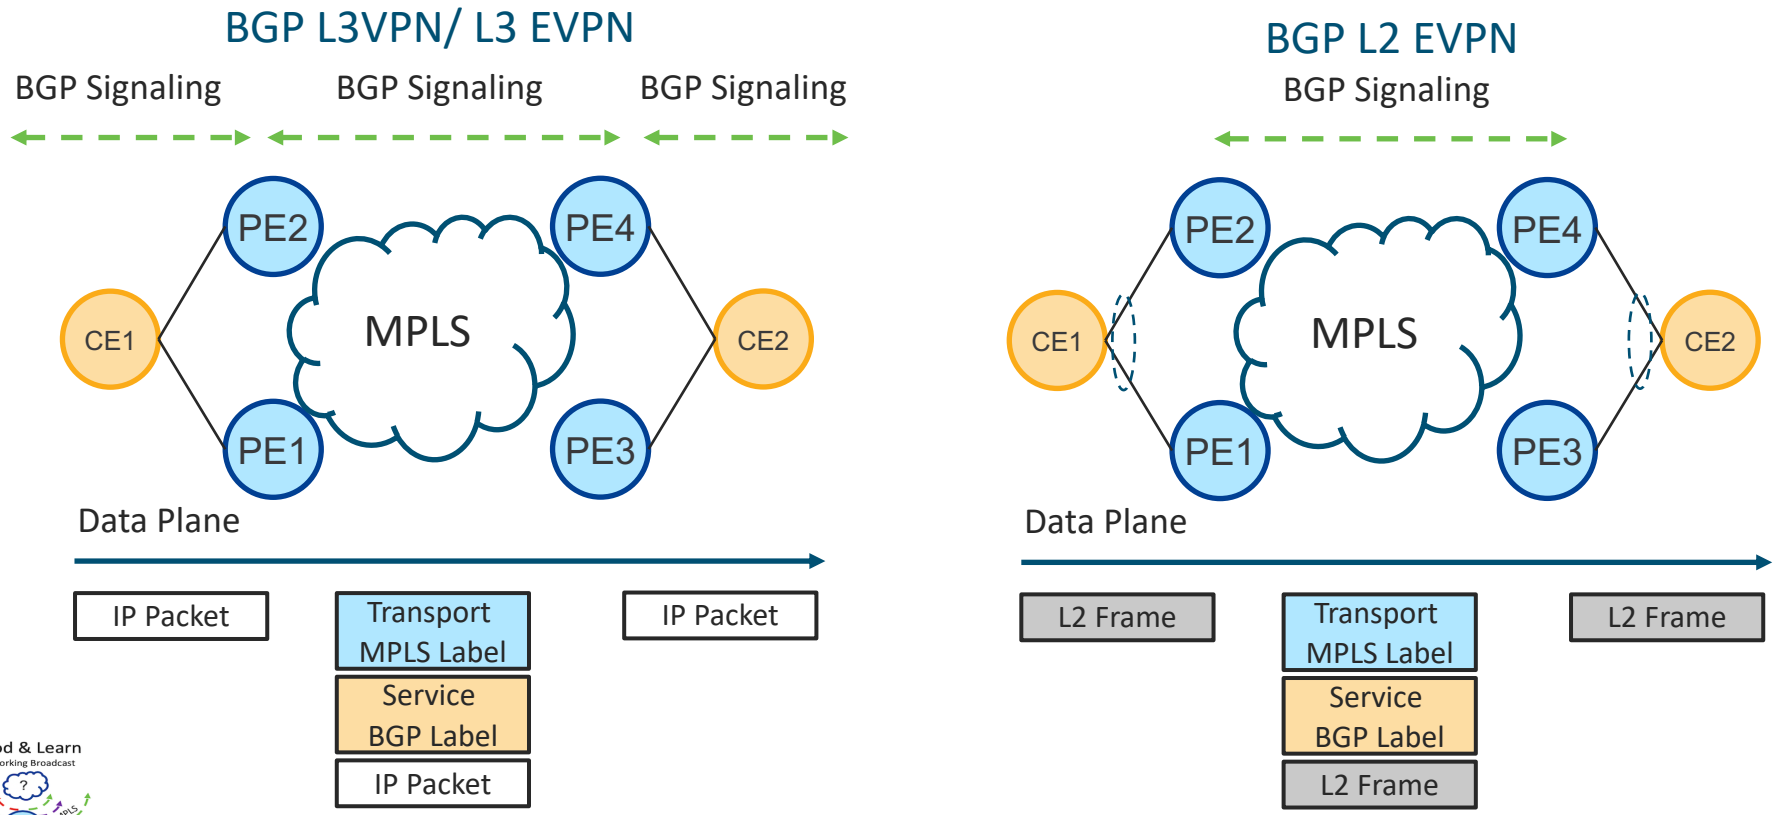

Flood & Learn
Networking Broadcast

https://e-vpn.io/fal

EVPN – Why?
Next-Generation Solutions for L2VPN
Solving VPLS challenges for per-flow Redundancy

- Existing VPLS solutions do not offer an All-Active per-flow redundancy
- Looping of Traffic Flooded from PE
- Duplicate Frames from Floods from the Core
- MAC Flip-Flopping over Pseudowire
  - E.g. Port-Channel Load-Balancing does not produce a consistent hash-value for a frame with the same source MAC (e.g. non MAC based Hash-Schemes)

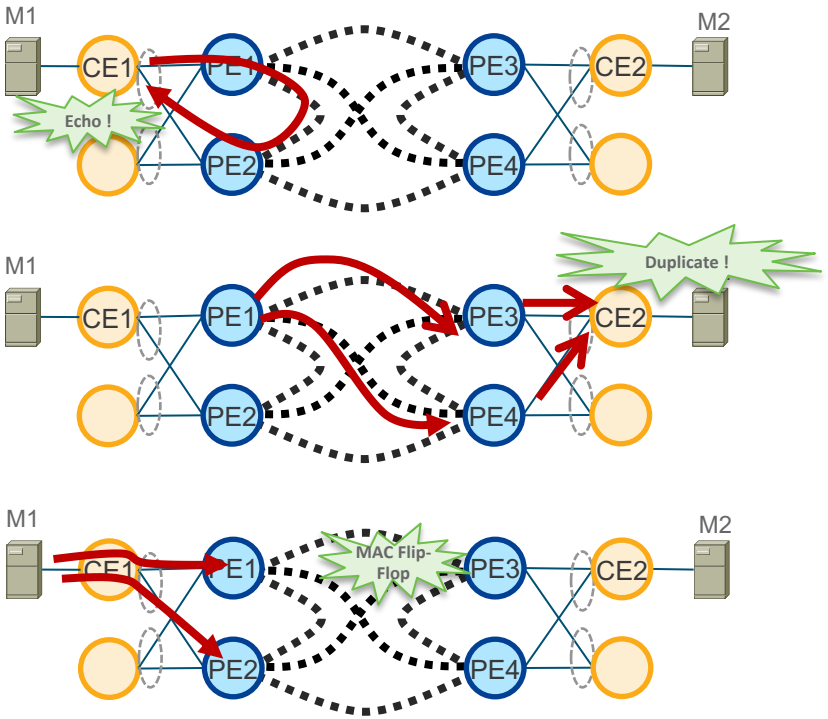

Flood & Learn
Networking Broadcast

https://e-vpn.io/fal

Concepts

EVPN Instance (EVI)

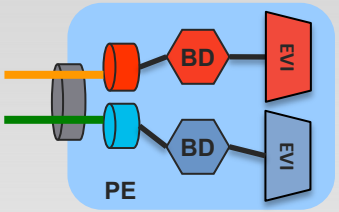

- **EVI identifies a VPN in the network**
- **Encompass one or more bridge-domains, depending on service interface type**
  - Port-based
  - VLAN-based (shown above)
  - VLAN-bundling

Ethernet Segment

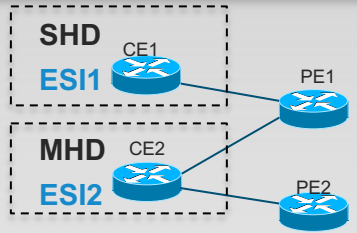

- **Represents a 'site' connected to one or more PEs**
- **Uniquely identified by a 10-byte global Ethernet Segment Identifier (ESI)**
- **Could be a single device or an entire network**
  - Single-Homed Device (SHD)
  - Multi-Homed Device (MHD)
  - Single-Homed Network (SHN)
  - Multi-Homed Network (MHN)

BGP Routes

Route Types
[1] Ethernet Auto-Discovery (AD) Route
[2] MAC/IP Advertisement Route
[3] Inclusive Multicast Route
[4] Ethernet Segment Route
[5] IP Prefix Advertisement Route

- **New SAFI [70]**
- **Routes serve control plane purposes, including:**
  - MAC address reachability
  - MAC mass withdrawal
  - Split-Horizon label adv.
  - Aliasing
  - Multicast endpoint discovery
  - Redundancy group discovery
  - Designated forwarder election
  - IP address reachability
  - L2/L3 Integration

BGP Route Attributes

Extended Communities
ESI MPLS Label
ES-Import
MAC Mobility
Default Gateway
Encapsulation

- **New BGP extended communities defined**
- **Expand information carried in BGP routes, including:**
  - MAC address moves
  - Redundancy mode
  - MAC / IP bindings of a GW
  - Split-horizon label encoding
  - Data plane Encapsulation

EVPN - load-balancing modes

All-Active (per flow)

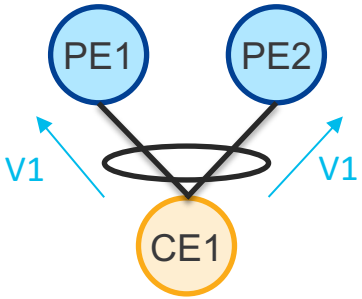

Single LAG at the CE
VLAN goes to both PE
Traffic hashed per flow
Benefits: Bandwidth, Convergence

Single-Active (per VLAN)

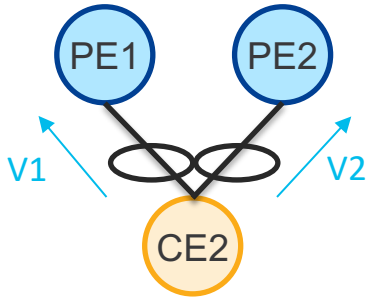

Multiple LAGs at the CE
VLAN active on single PE
Traffic hashed per VLAN
Benefits: Billing, Policing

Port-Active (per port)

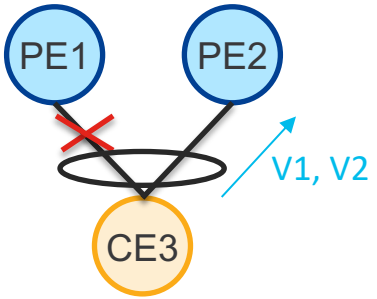

Single LAGs at the CE
Port active on single PE
Traffic hashed per port
Benefits: Protocol Simplification

All-Active Multi-Homed EVPN Access

EVPN - Ethernet-Segment for Multi-Homing

L1 and L2 (L3 and L4) have to know if they multi-home same broadcast domain

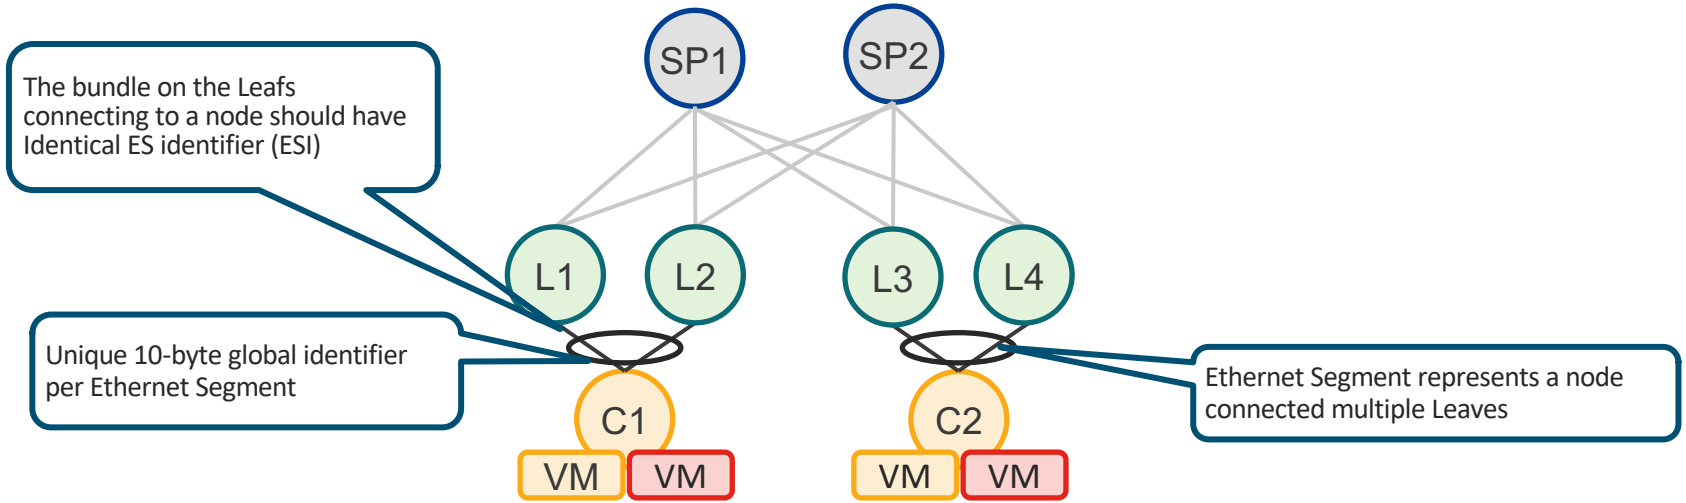

Flood & Learn
Networking Broadcast

https://e-vpn.io/fal

EVPN - Ethernet VPN

MAC address advertisement and MAC address table synchronization

Leaves run Multi-Protocol BGP to advertise & learn MAC addresses over the Network
MAC addresses are advertised to rest of Leaves
L3/4 – Learn MAC address advertised by L1
L2 – uses MAC address advertised by L1 to synchronize MAC address table
-> L2 forwards MAC via local ETH interface represented by same Ethernet Segment between L1 and L2

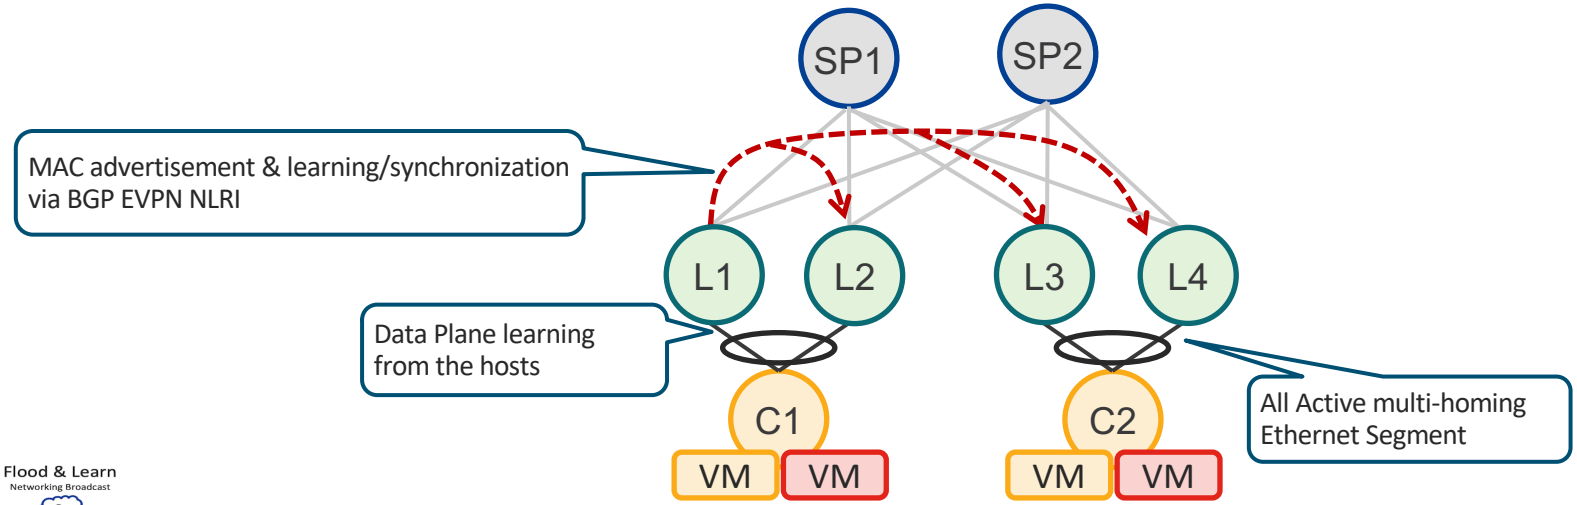

Flood & Learn
Networking Broadcast

https://e-vpn.io/fal

EVPN – BUM Ingress Replication

Two service labels per EVPN instance

BUM Label – to forward Broadcast, Unknown Unicast and Multicast

Unicast Label – to forward Unicast

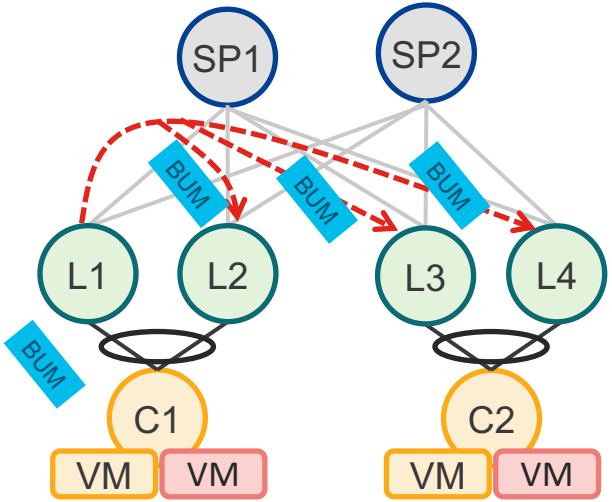

EVPN – Designated Forwarder (DF)

Challenge:

How to prevent duplicate copies of flooded traffic from being delivered to a multi-homed Ethernet Segment?

If (L3 and L4) Multi-Homing access via same Ethernet Segment -> only one of them can forward traffic to access

Same for (L1 and L2)

Why extra BUM Label?

What if Unicast Traffic is sent to L3 or L4 (not flooded)? -> DF Election applies only to BUM (from Core to Access)

DF, Redirect, Fast Re-Route (FRR), etc.

Service Label informs egress Leaf if traffic is BUM or Unicast

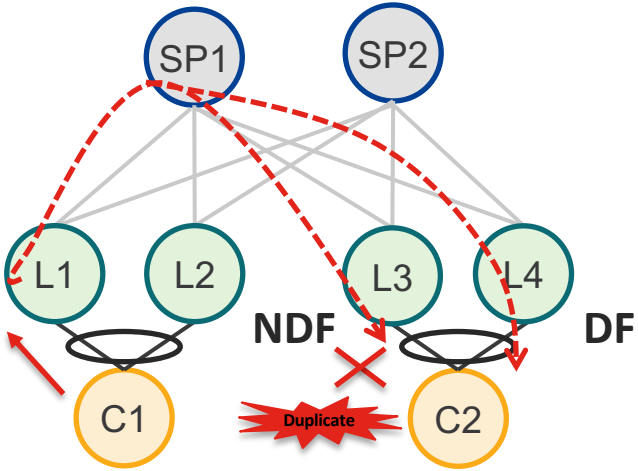

EVPN – Split Horizon

Challenge:
How to prevent flooded traffic from echoing back to a multi-homed Ethernet Segment?

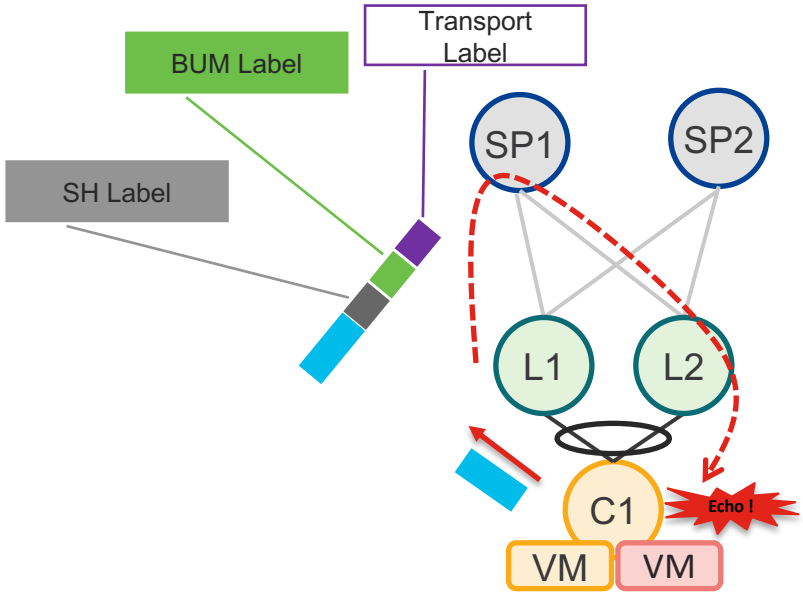

EVPN – Split Horizon

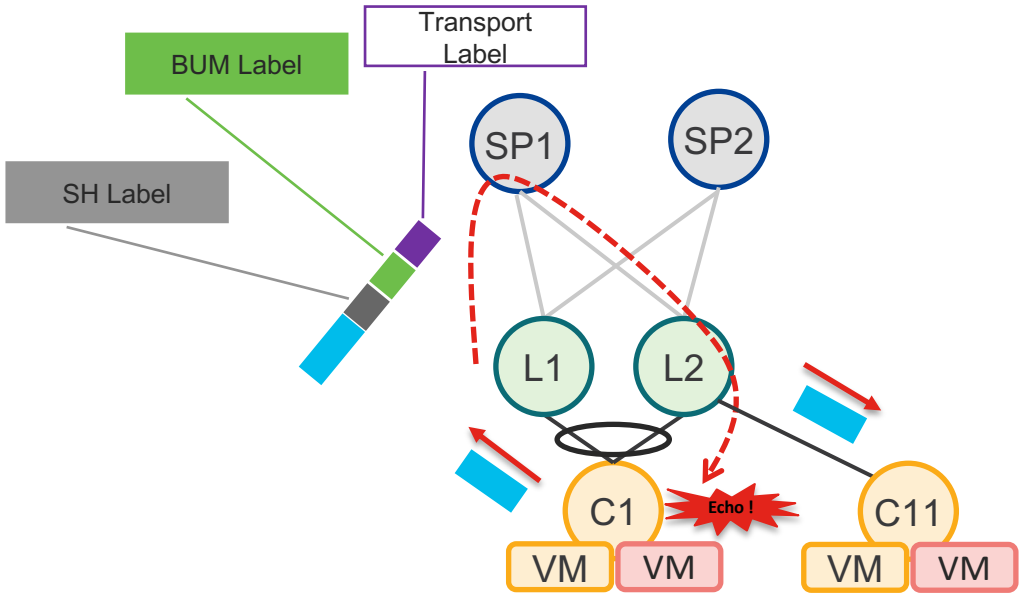

Flood & Learn
Networking Broadcast

https://e-vpn.io/fal

EVPN – MAC Mass-Withdraw

Challenge:
How to inform other Leafs of a failure affecting many MAC addresses quickly while the control-plane re-converges?

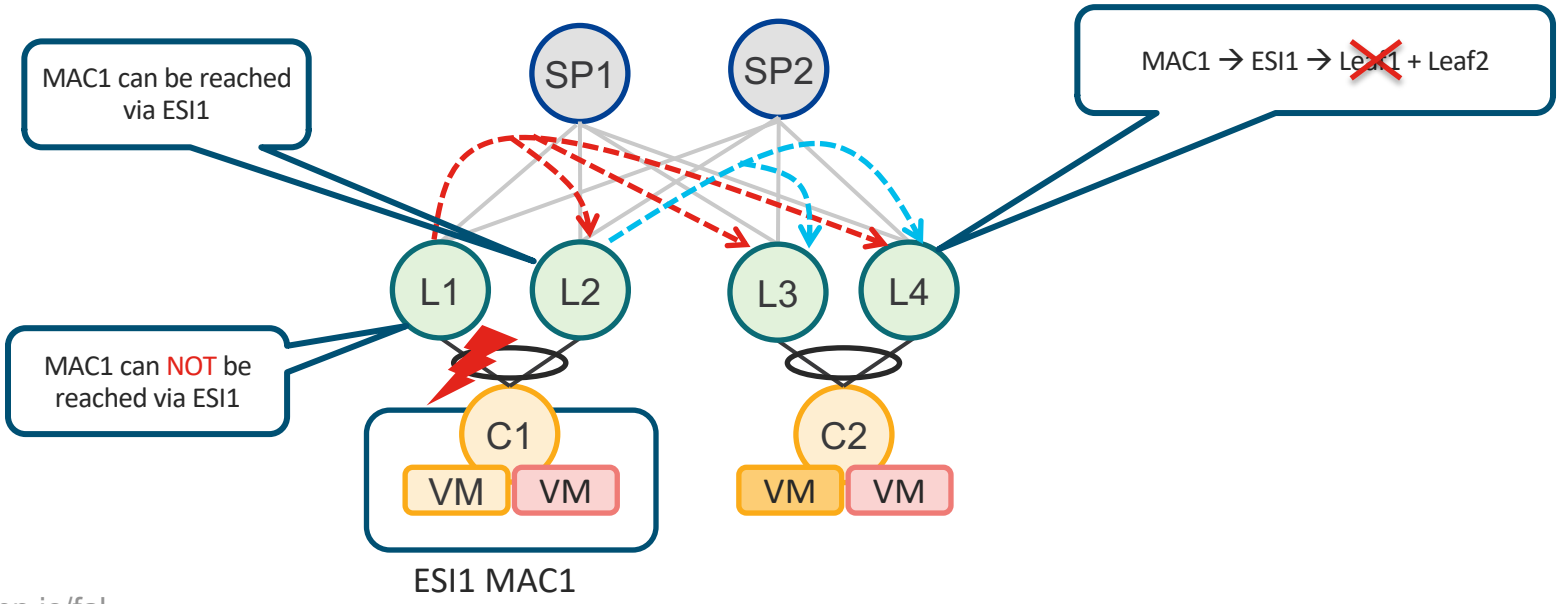

R36, R37, R38, R39 - EVPN Startup
R36 - Example

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**

Service Carving: 100 modulo 2 = 0
R36 is DF for EVI-100

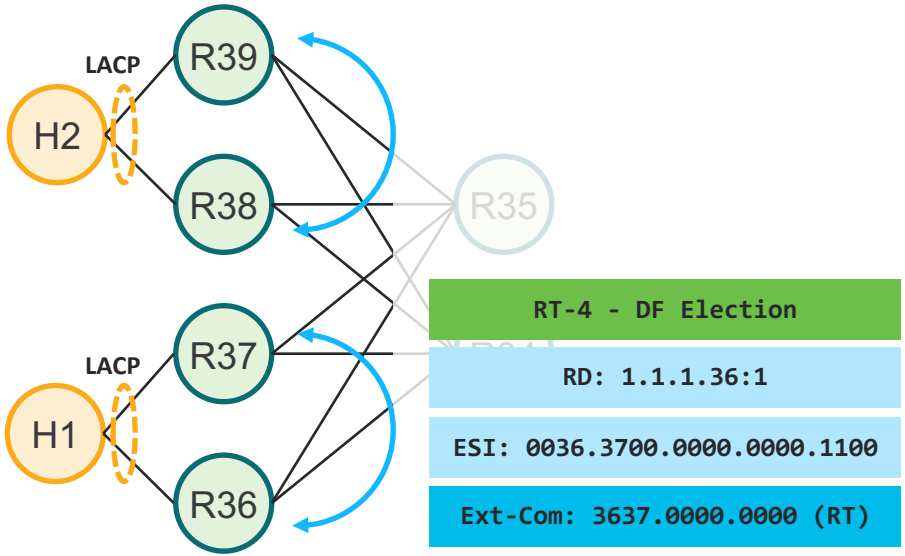

R36, R37, R38, R39 - EVPN Startup
R36 - Example

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**

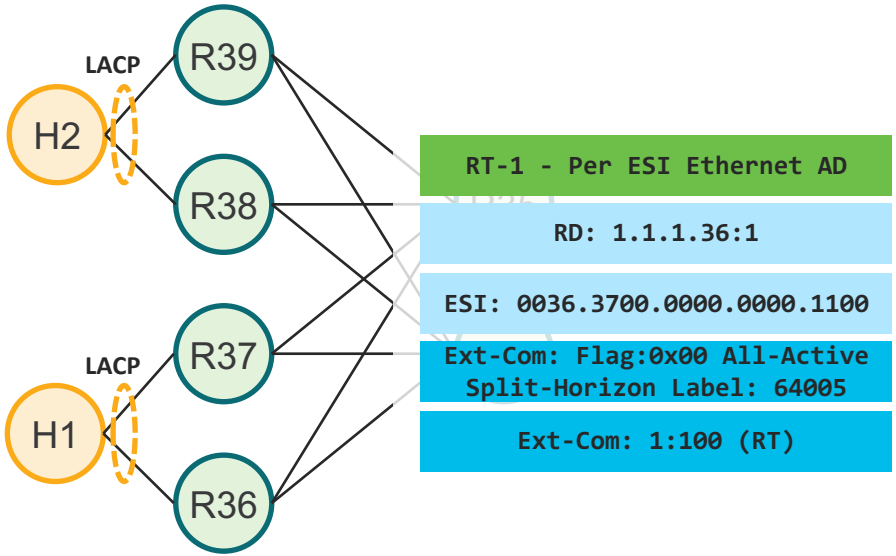

R36, R37, R38, R39 - EVPN Startup
R36 - Example

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**

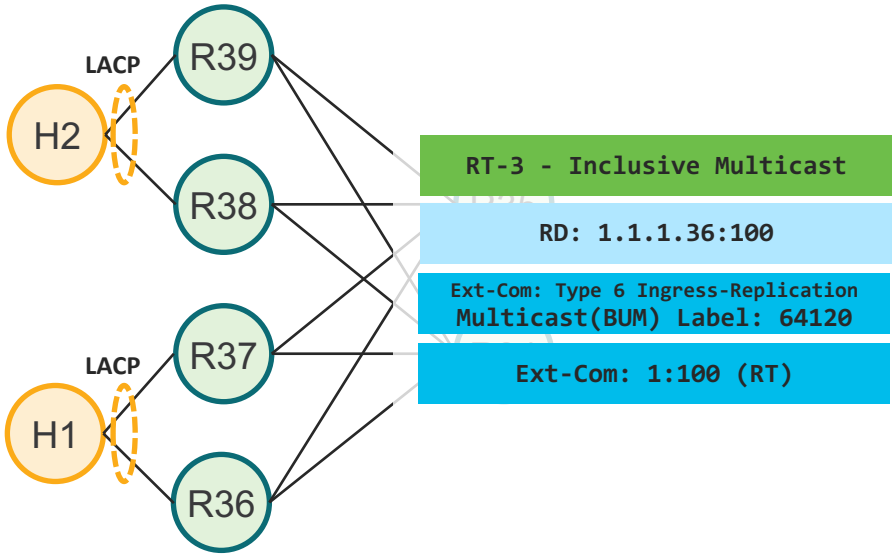

BUM Forwarding

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**

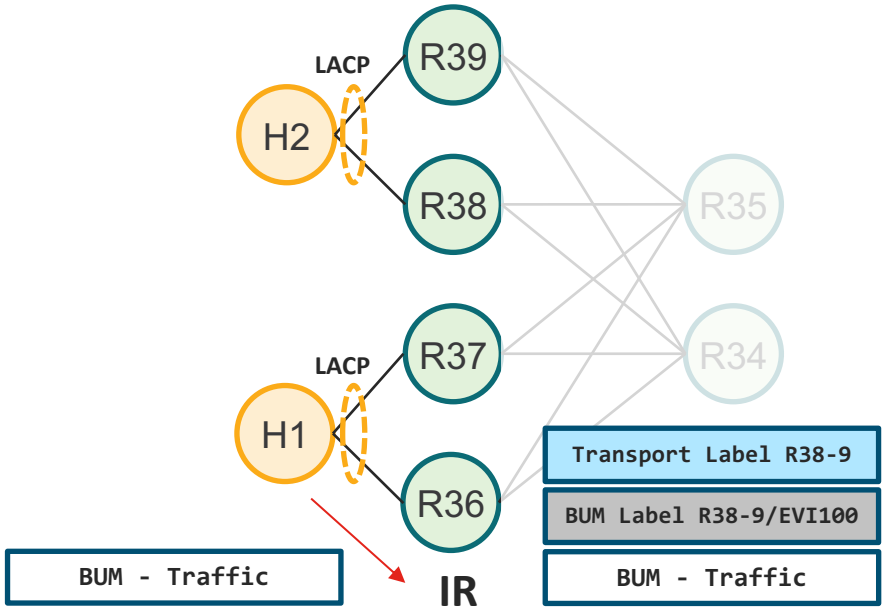

BUM Forwarding

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**

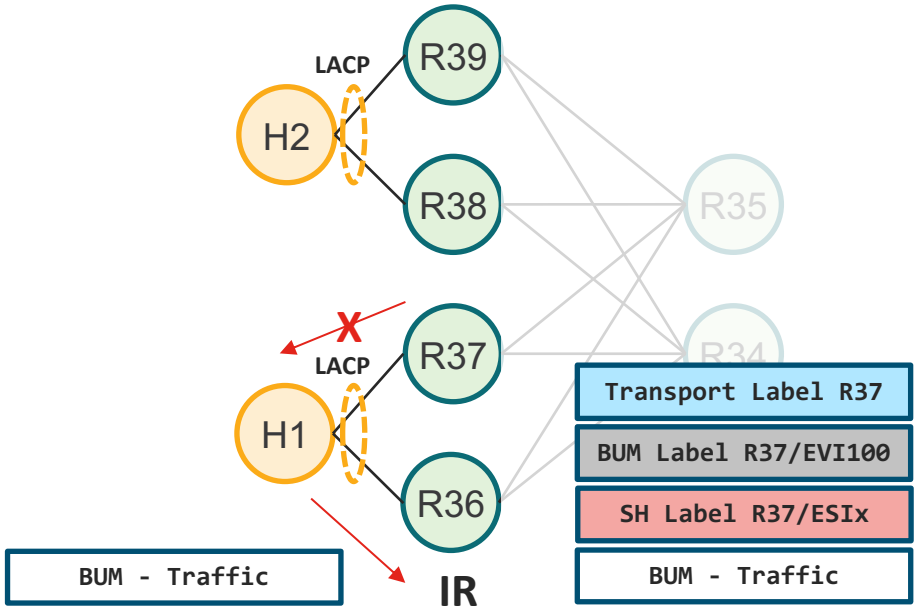

R36, R37, R38, R39 - EVPN Startup
R36 - Example

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**
4. **RT2: MAC Advertisement**

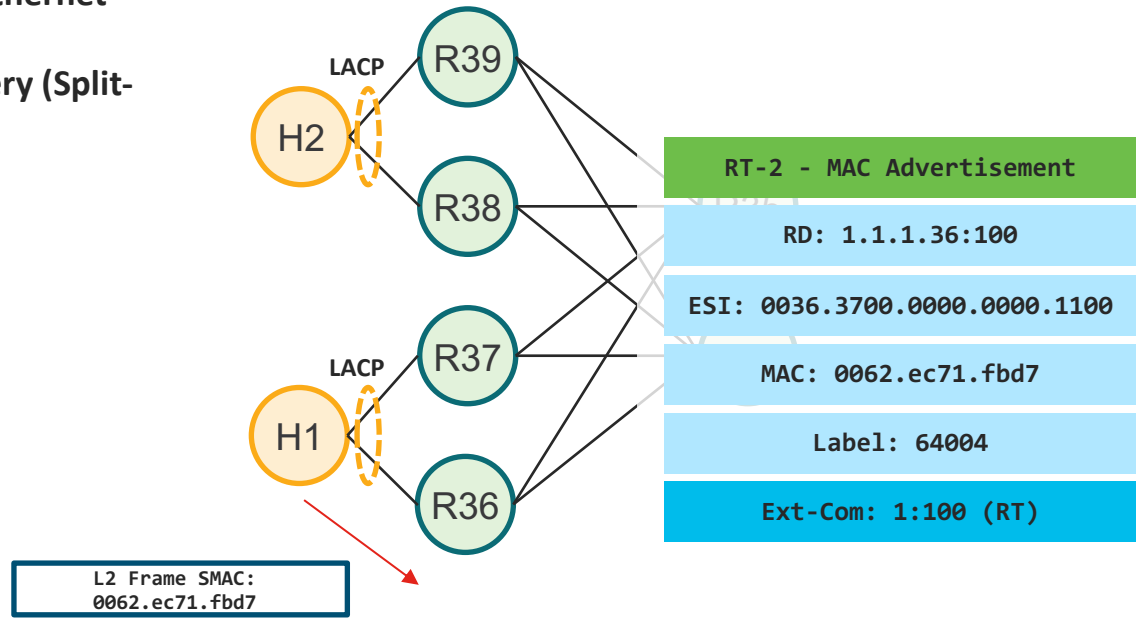

Unicast Forwarding

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**
4. **RT2: MAC Advertisement**

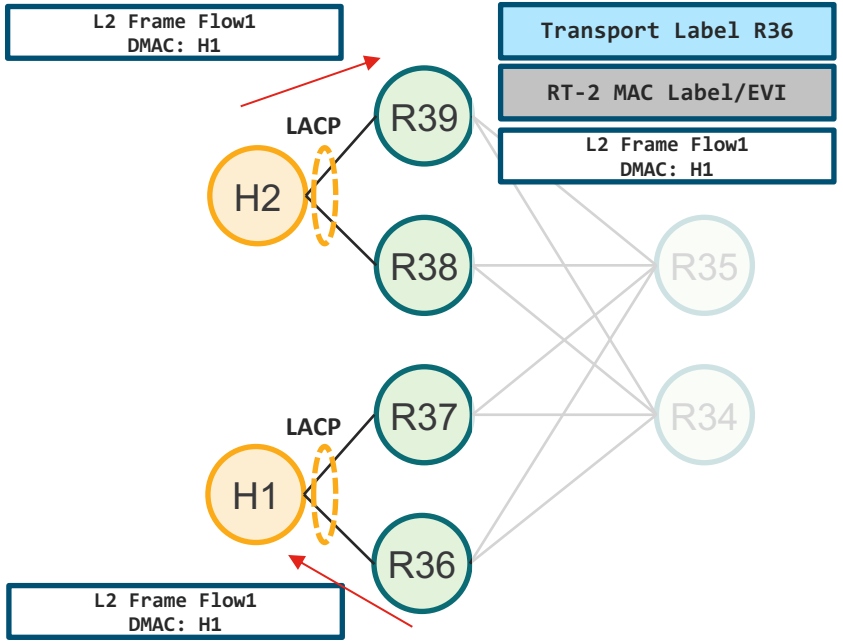

EVPN – Aliasing

Challenge:
How to load-balance traffic towards a multi-homed device across multiple Leaves when MAC addresses are learnt by only a single Leaf?

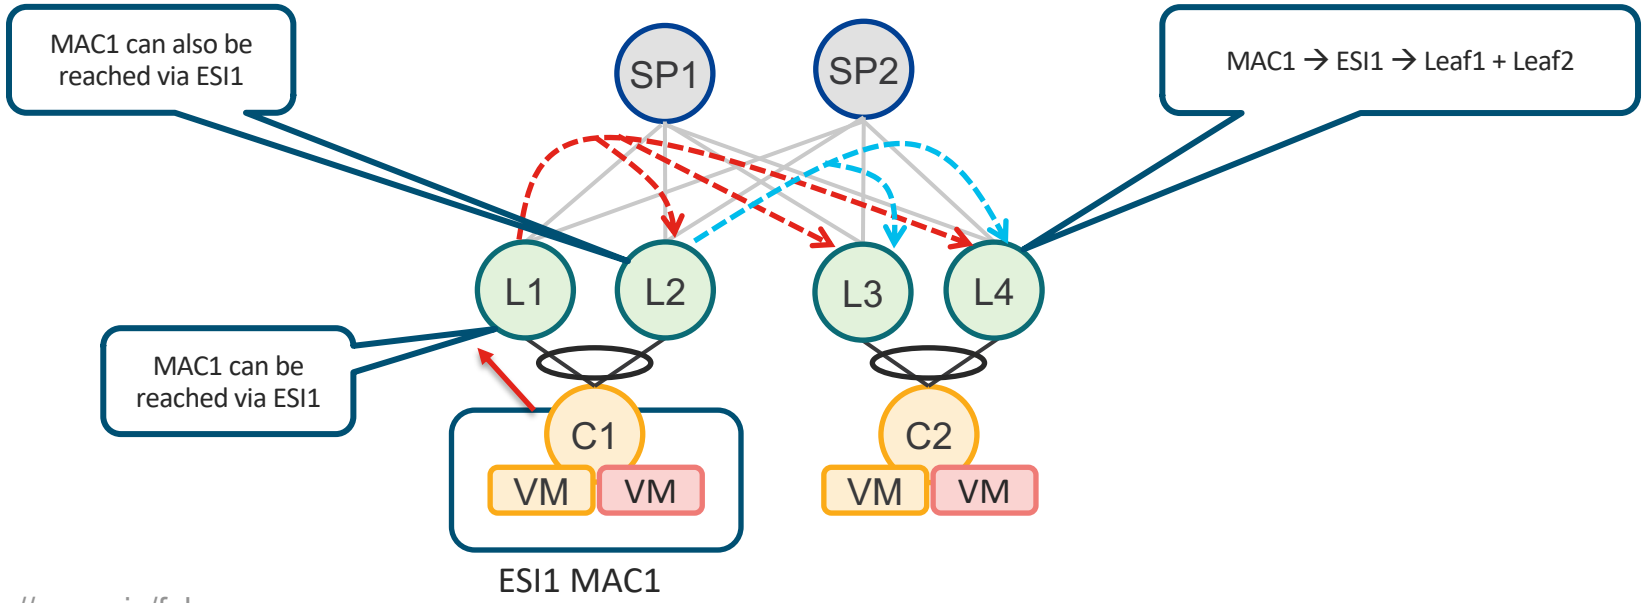

Flood & Learn
Networking Broadcast

https://e-vpn.io/fal

R36, R37, R38, R39 - EVPN Startup
R36 - Example

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**
4. **RT2: MAC Advertisement**
5. **RT1: Per EVI Ethernet Auto-Discovery**

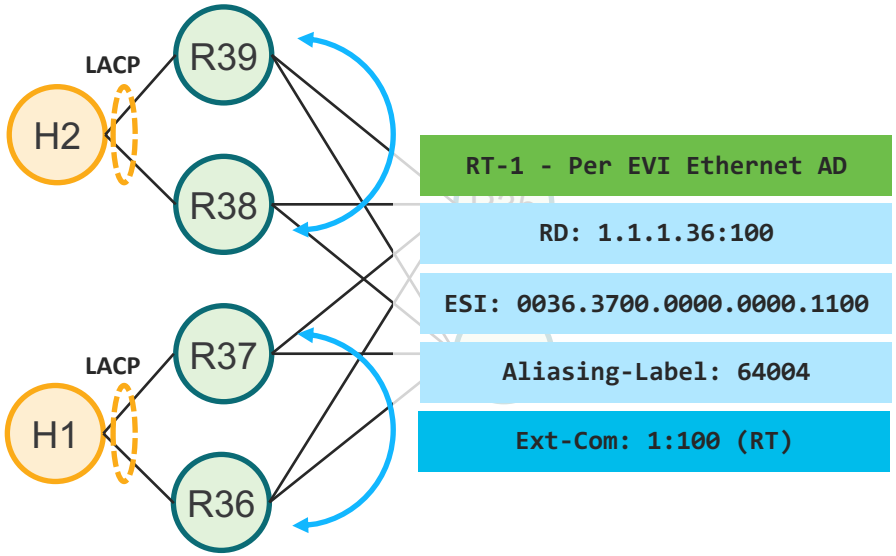

Unicast Forwarding

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**
4. **RT2: MAC Advertisement**
5. **RT1: Per EVI Ethernet Auto-Discovery**

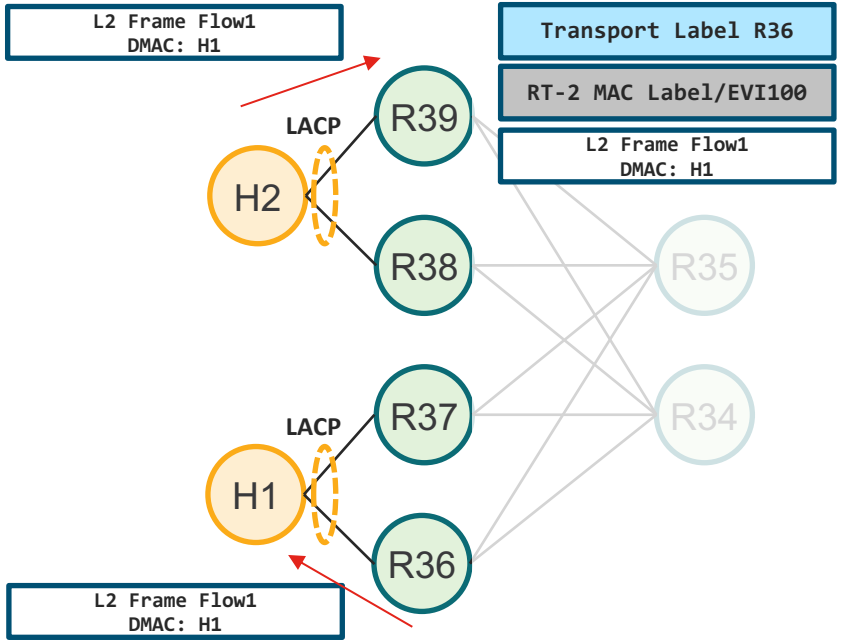

Unicast Forwarding

1. **RT4: DF Election & Multi-Homed Ethernet Segment Auto-Discovery**
2. **RT1: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)**
3. **RT3: Inclusive Multicast**
4. **RT2: MAC Advertisement**
5. **RT1: Per EVI Ethernet Auto-Discovery**

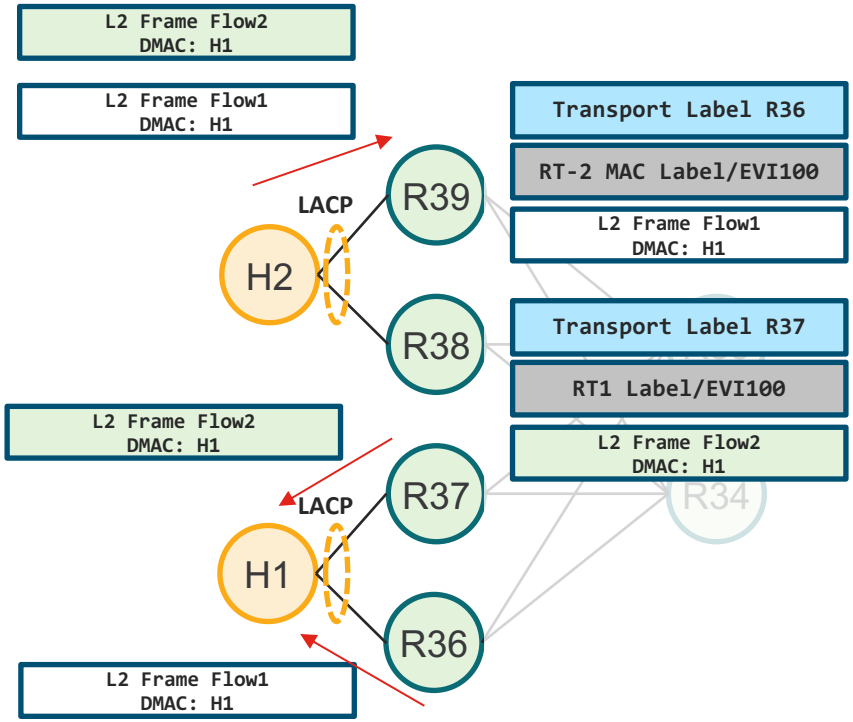

Unicast Forwarding

1. **RT4**: DF Election & Multi-Homed Ethernet Segment Auto-Discovery
2. **RT1**: Per ESI Ethernet Auto-Discovery (Split-Horizon, Mass-Withdraw)
3. **RT3**: Inclusive Multicast
4. **RT2**: MAC Advertisement
5. **RT1**: Per EVI Ethernet Auto-Dis

Per Flow Balancing via R36 and R37 - Aliasing

Per Flow Balancing via R36 and R37 - Aliasing

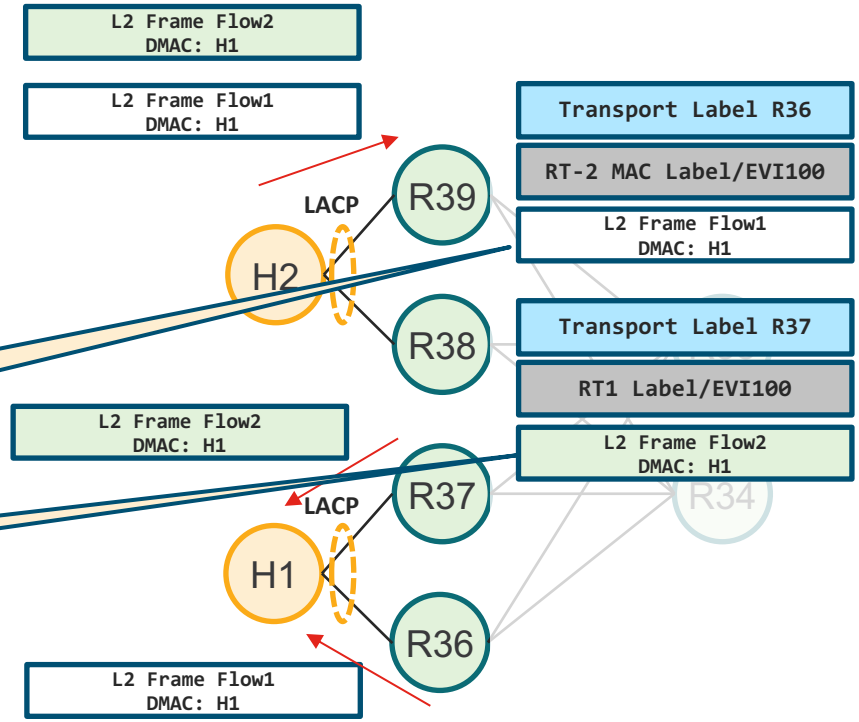

EVPN – MAC Mobility

Challenge:
How to detect the correct location of MAC after the movement of host from one Ethernet Segment to another also called “MAC move”?

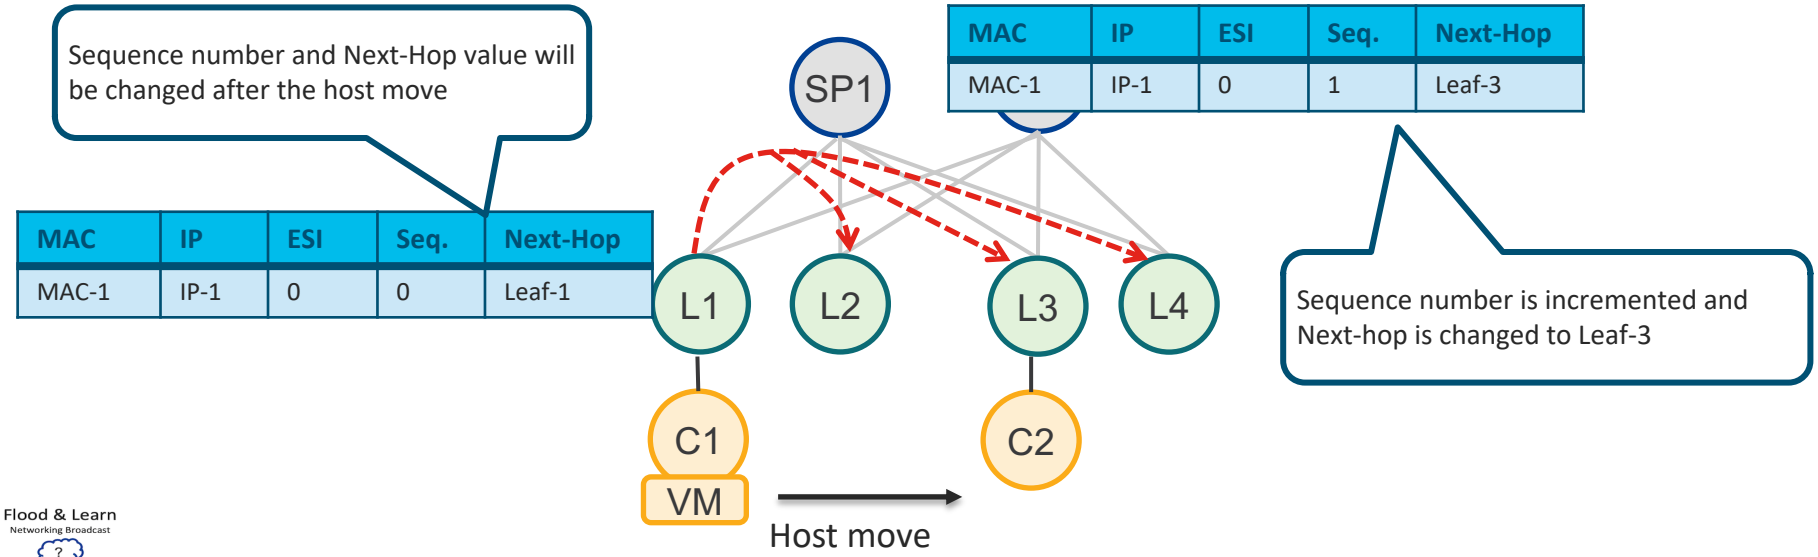

Flood & Learn
Networking Broadcast

https://e-vpn.io/fal

EVPN – Distributed Symmetric Anycast Gateway

Leaves run Multi-Protocol BGP to advertise & learn MAC + HOST IP addresses over the Network
MAC + IP addresses are advertised to rest of Leaves
L3/4 – Learn MAC + IP HOST address advertised by L1
-> L2/L3 update MAC address table + **IP Forwarding table**
L2 – uses MAC address advertised by L1 to synchronize MAC address table
-> L2 forwards MAC via local ETH interface represented by same Ethernet Segment between L1 and L2
L2 – uses MAC + IP HOST address advertised by L1 to synchronize ARP/ND information
-> L2 forwards IP via local ETH interface

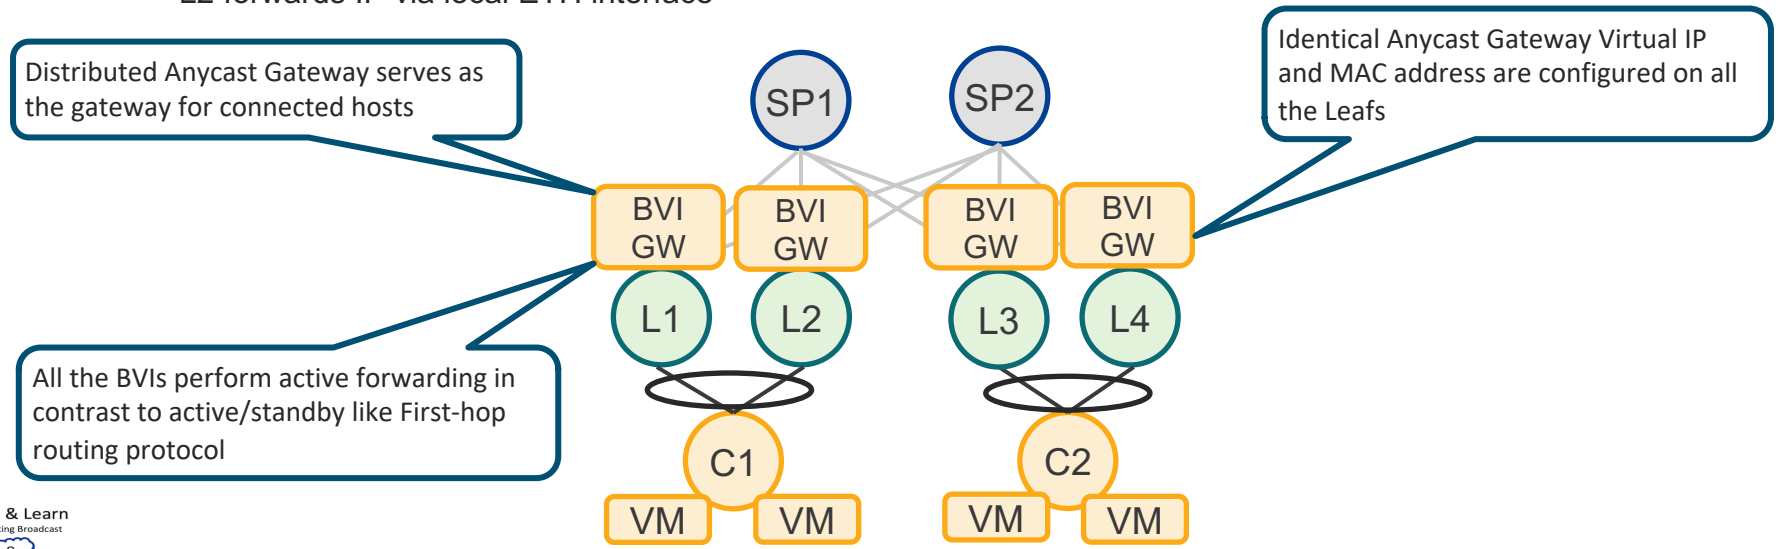

EVPN – IRB in Network Fabric

Purpose:

Optimal intra and inter-subnet connectivity with seamless workload mobility

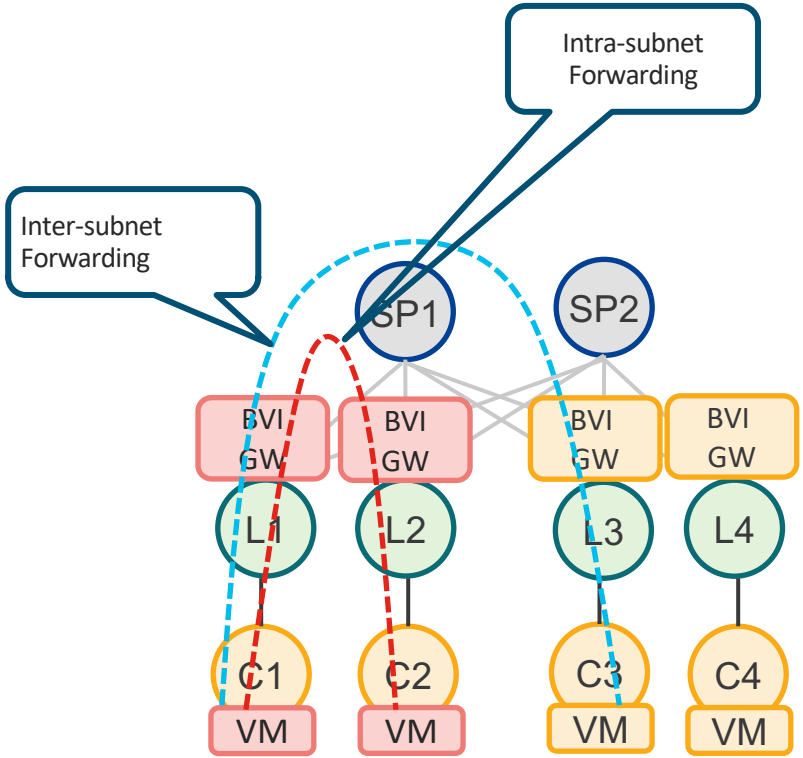

EVPN - Stay Up-To-Date

EVPN

- https://e-vpn.io/
- Upcoming “Flood & Learn” Networking Broadcast: https://e-vpn.io/fal/

https://e-vpn.io/fal