flood-learn.io

How to Build Layer3 Data Center Interconnect (DCI)

Jiri Chaloupka – Cisco Technical Marketing Engineer

04/2020

Objectives

- Purpose of EVPN
  - Layer2 vs Layer3 Role
- Importance of Layer3 Data Center Interconnect (DCI)
  - Why service termination (VRF) on DCI?
- BGP Layer3 Control Plane Options
  - VPNv4/6 and EVPN
- Layer3 Interconnect Data Plane Options
  - MPLS and IP (VXLAN and SRv6)

Distributed vs Centralized Routing

Layer2 Bridging mandatory between Leaves only

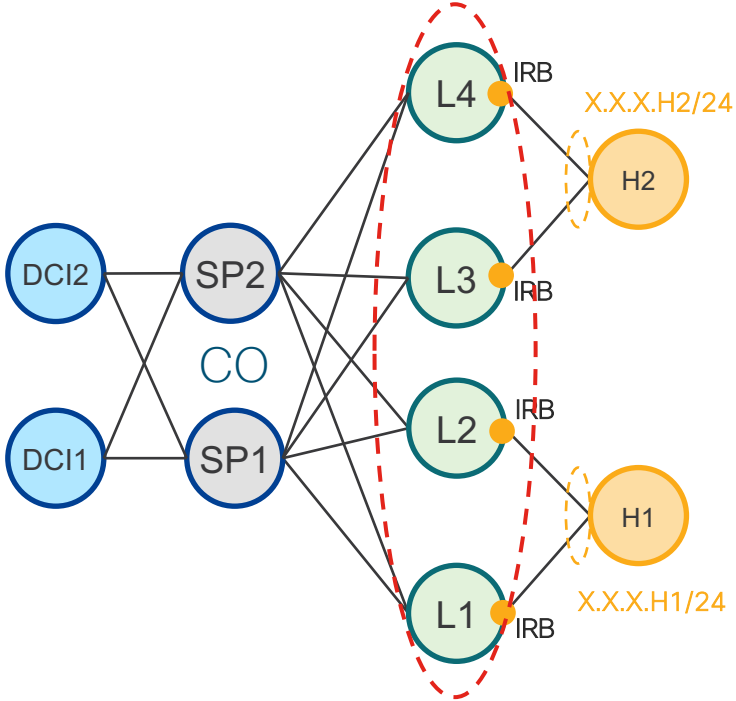

- Optimized forwarding of east-west traffic
  ARP/MAC state localized to Leafs
  Helps with horizontal scaling of DC

Layer2 Bridging mandatory between Leaves and DCI

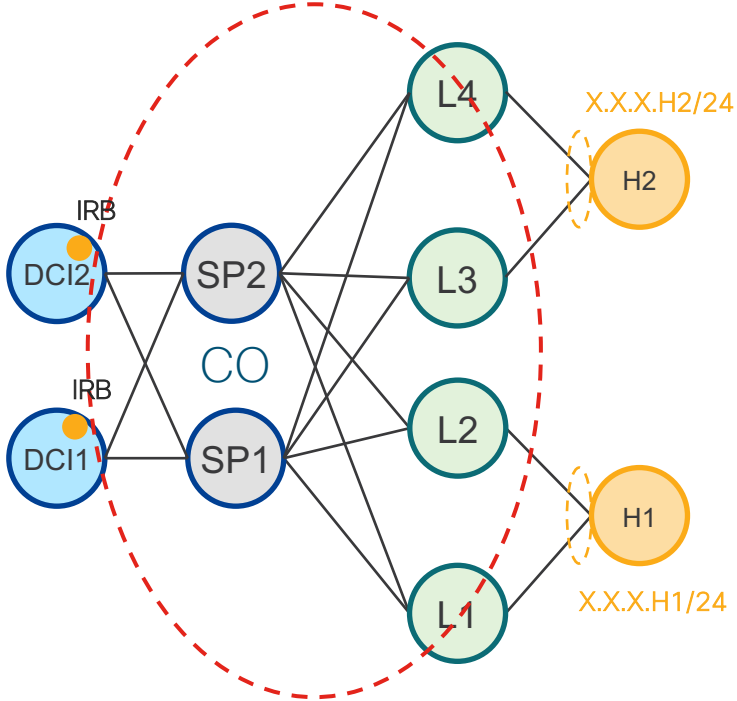

- All east<->west routed traffic traverses to centralized gateways
- Centralized gateways have full ARP/MAC state in the DCI
- Scale challenge

BGP Layer3 Interconnect
Principles

- DCI/BL provides Layer3 Interconnect
- DCI/BL participates in **L3 Routing**, but **not in Layer2 Bridging**
- **DCI/BL summarization is required/recommended**

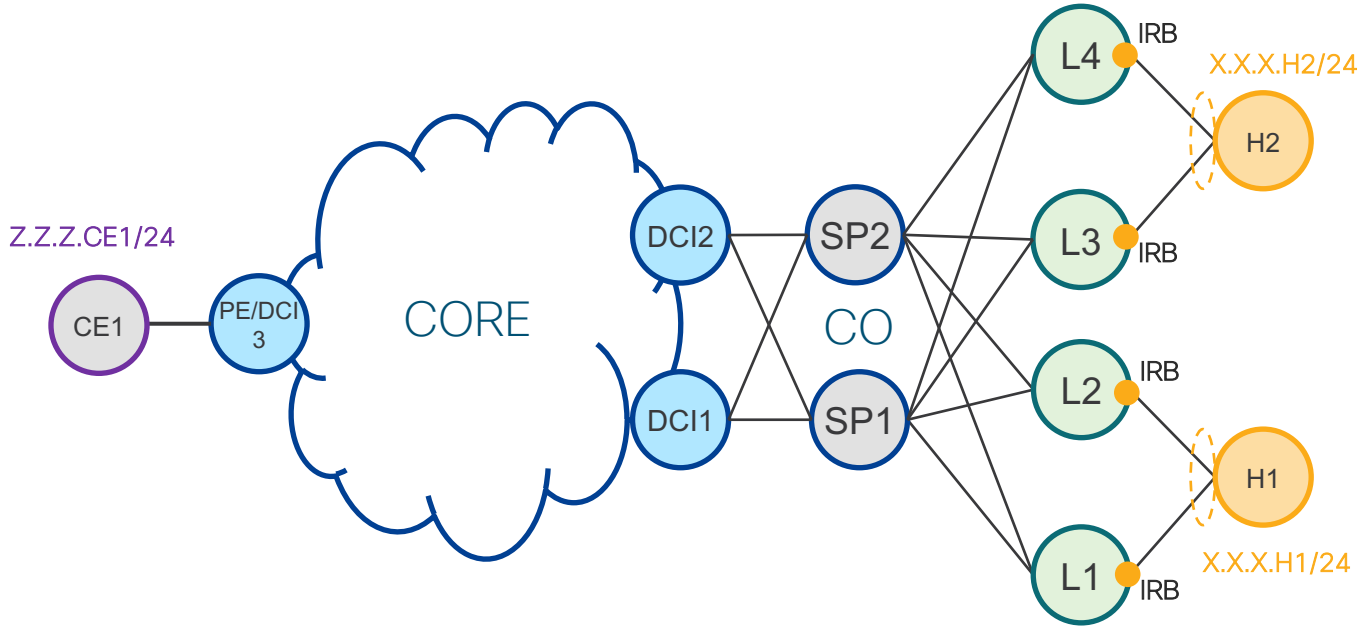

BGP Layer3 Interconnect
Principles

- DCI/BL provides Layer3 Interconnect
- DCI/BL participates in L3 Routing, but not in Layer2 Bridging
- DCI/BL summarization is required/recommended

Layer2 Bridging Required over Leaves

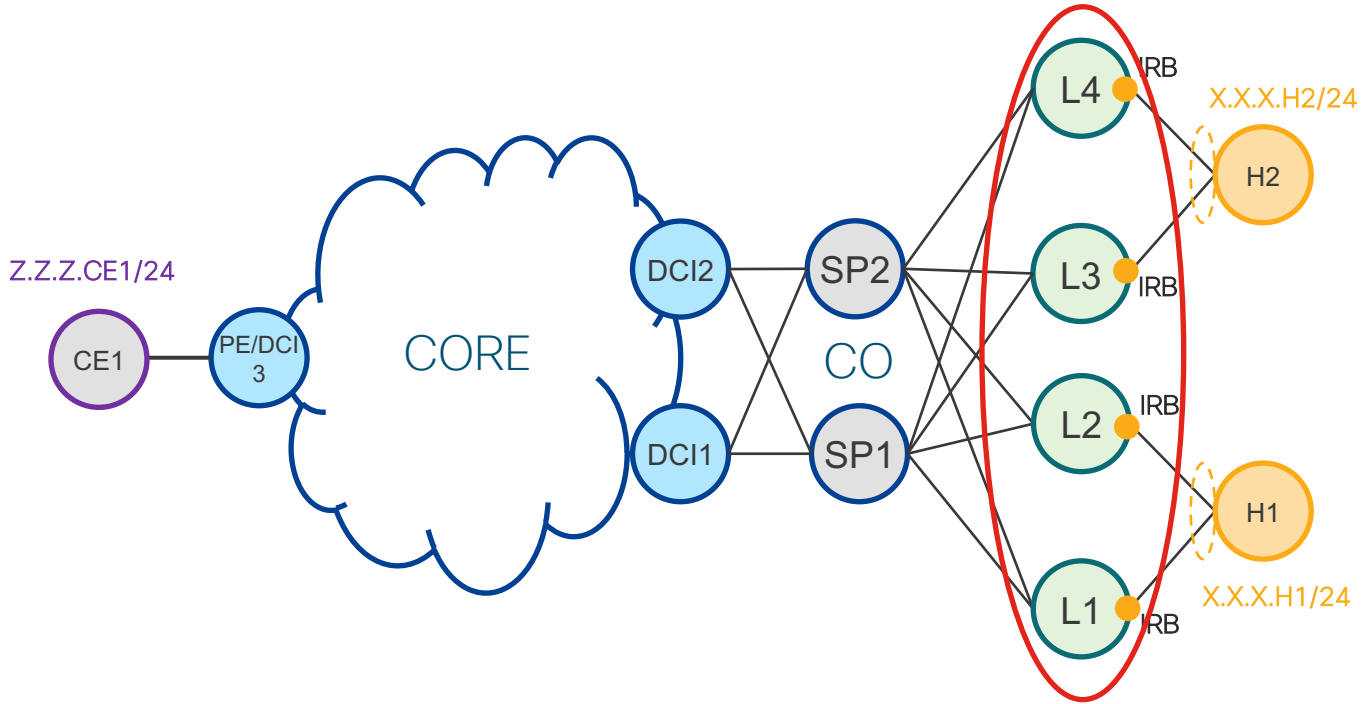

BGP Layer3 Interconnect

DCI/BL Summarization

Host-Routes are not required outside CO/DC

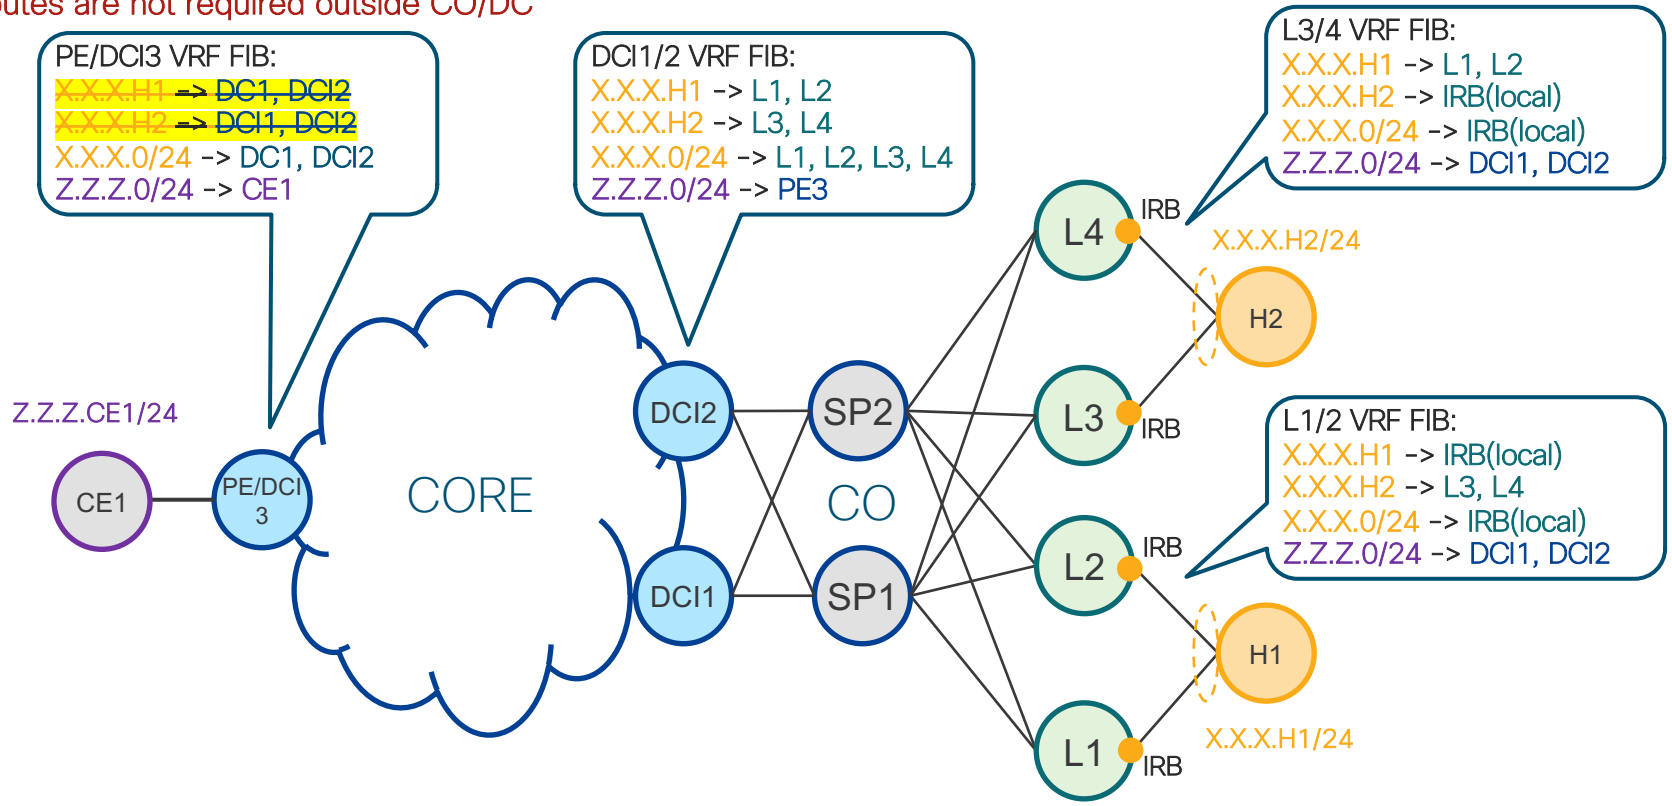

BGP Layer3 Interconnect

Control Plane

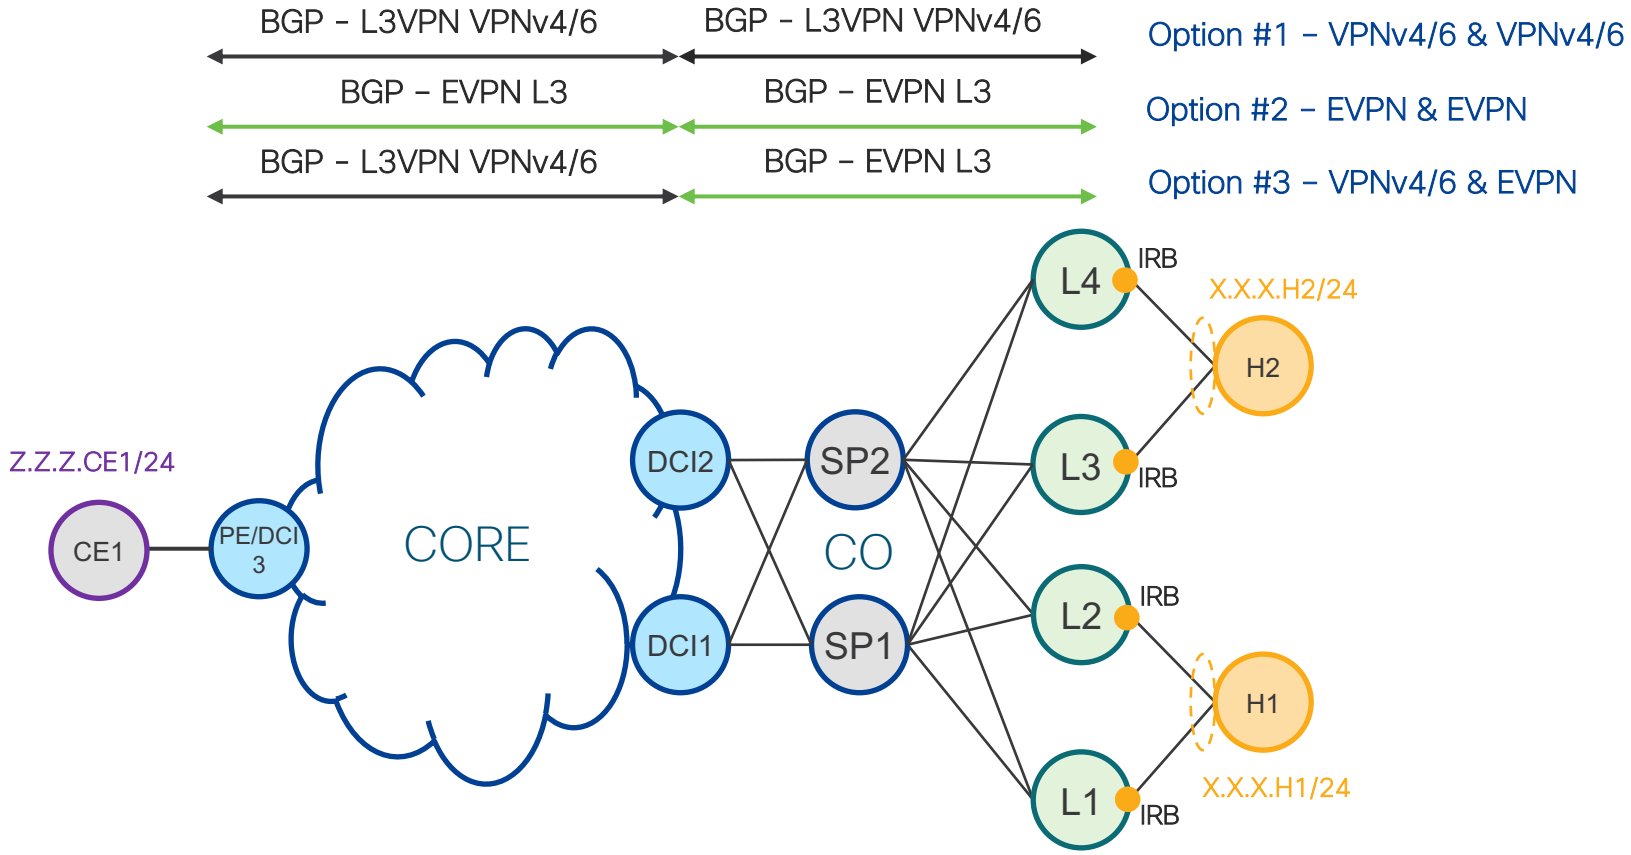

BGP Layer3 Interconnect

Option #1 – VPNv4/6 & VPNv4/6

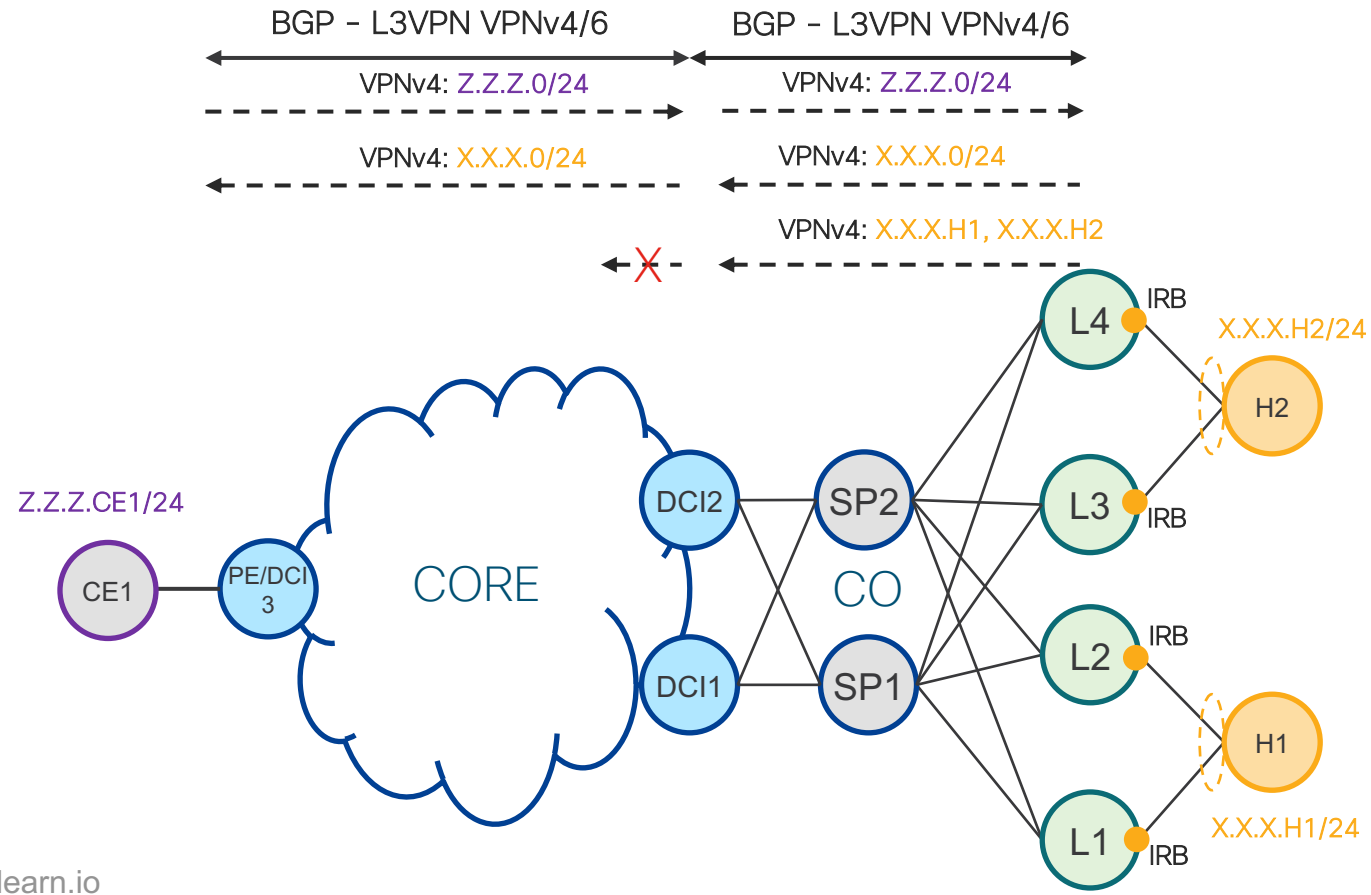

BGP Layer3 Interconnect

Option #2 – EVPN & EVPN

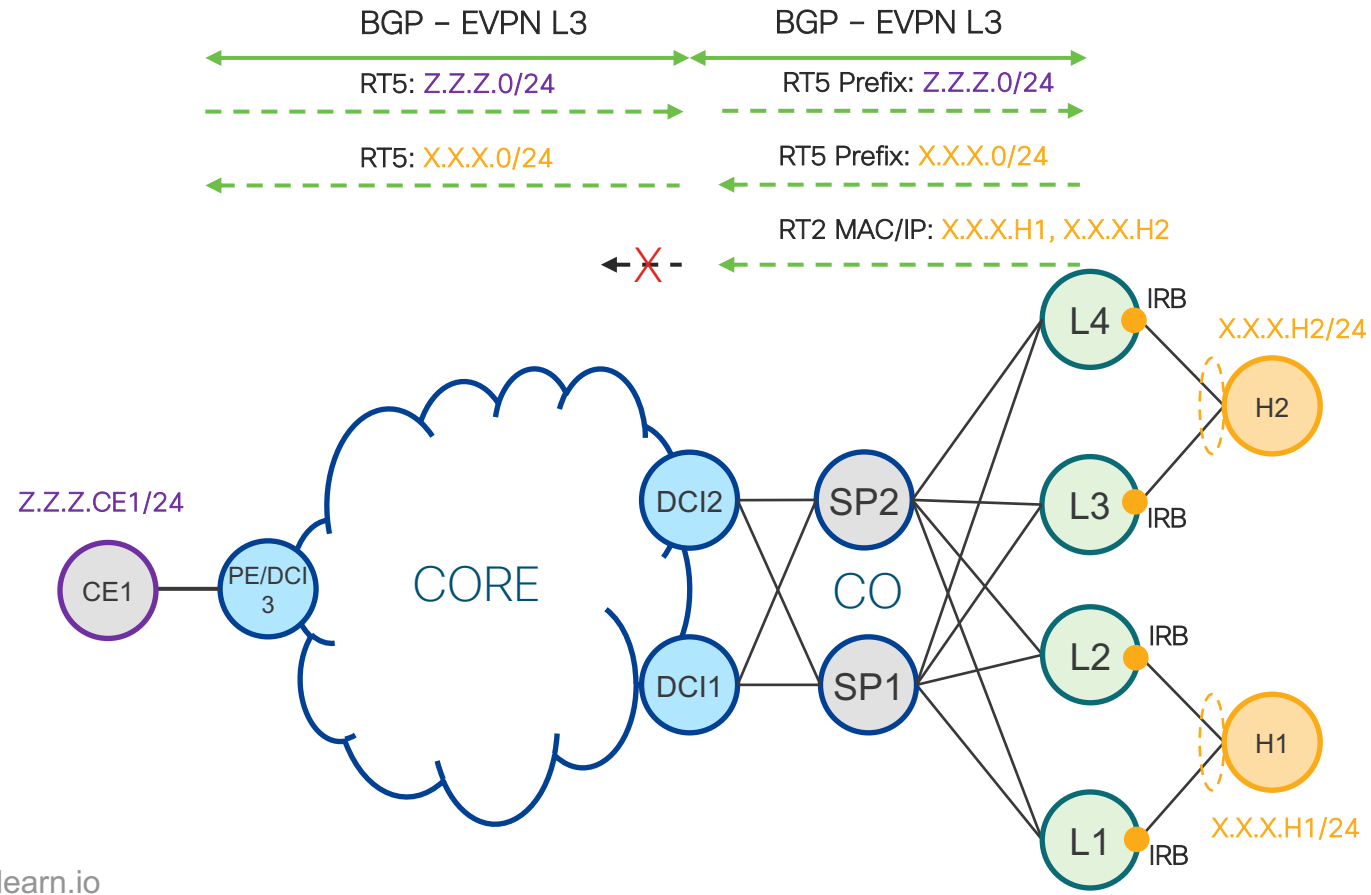

BGP Layer3 Interconnect

Option #3 – VPNv4/6 & EVPN

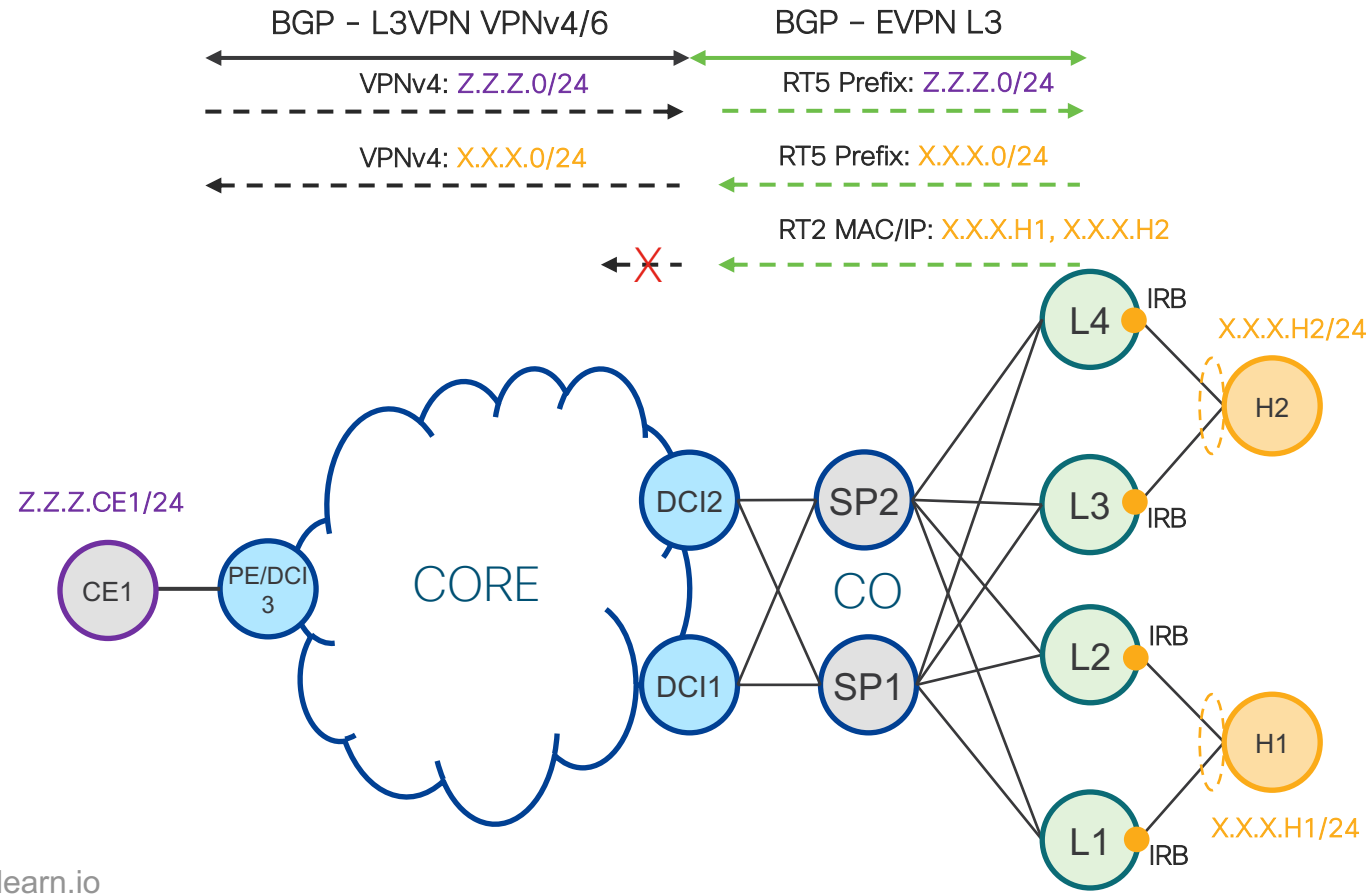

BGP Layer3 Interconnect
Control Plane Options Highlight

- Option #1 – VPNv4/6 & VPNv4/6
  - + VPNv4/6 Industry proofed solution for Layer3 VPN
  - + DCI doesn't need to understand BGP EVPN AF
  - – Leaf has to peer with Route-Reflector via both BGP EVPN and VPNv4/6 AF
    - EVPN AF to support L2 stretch (MAC advertisement) across DC/CO between Leaves
    - EVPN AF to sync ARP/ND for Multi-Homed All-Active
  - – DC/CO Route-Reflector has to support both BGP EVPN and VPNv4/6 AF
  - – Leaf has to advertise VM Host-Routes via VPNv4/6

- Option #2 – EVPN & EVPN
  - + Single BGP Address Family End-To-End in Network
  - – Existing L3 VPNv4/6 services has to to migrated to L3 EVPN
    - No technical benefit to migrate existing L3 VPNv4/6 to L3 EVPN

- Option #3 – VPNv4/6 & EVPN
  - + Recommended solution which benefits from both Options #1 and #2
  - + New DC/CO – Leaf, Route-Reflector use single BGP AF EVPN
  - + Existing L3 VPNv4/6 services stay untouched

BGP Layer3 Interconnect

End-To-End MPLS Data Plane

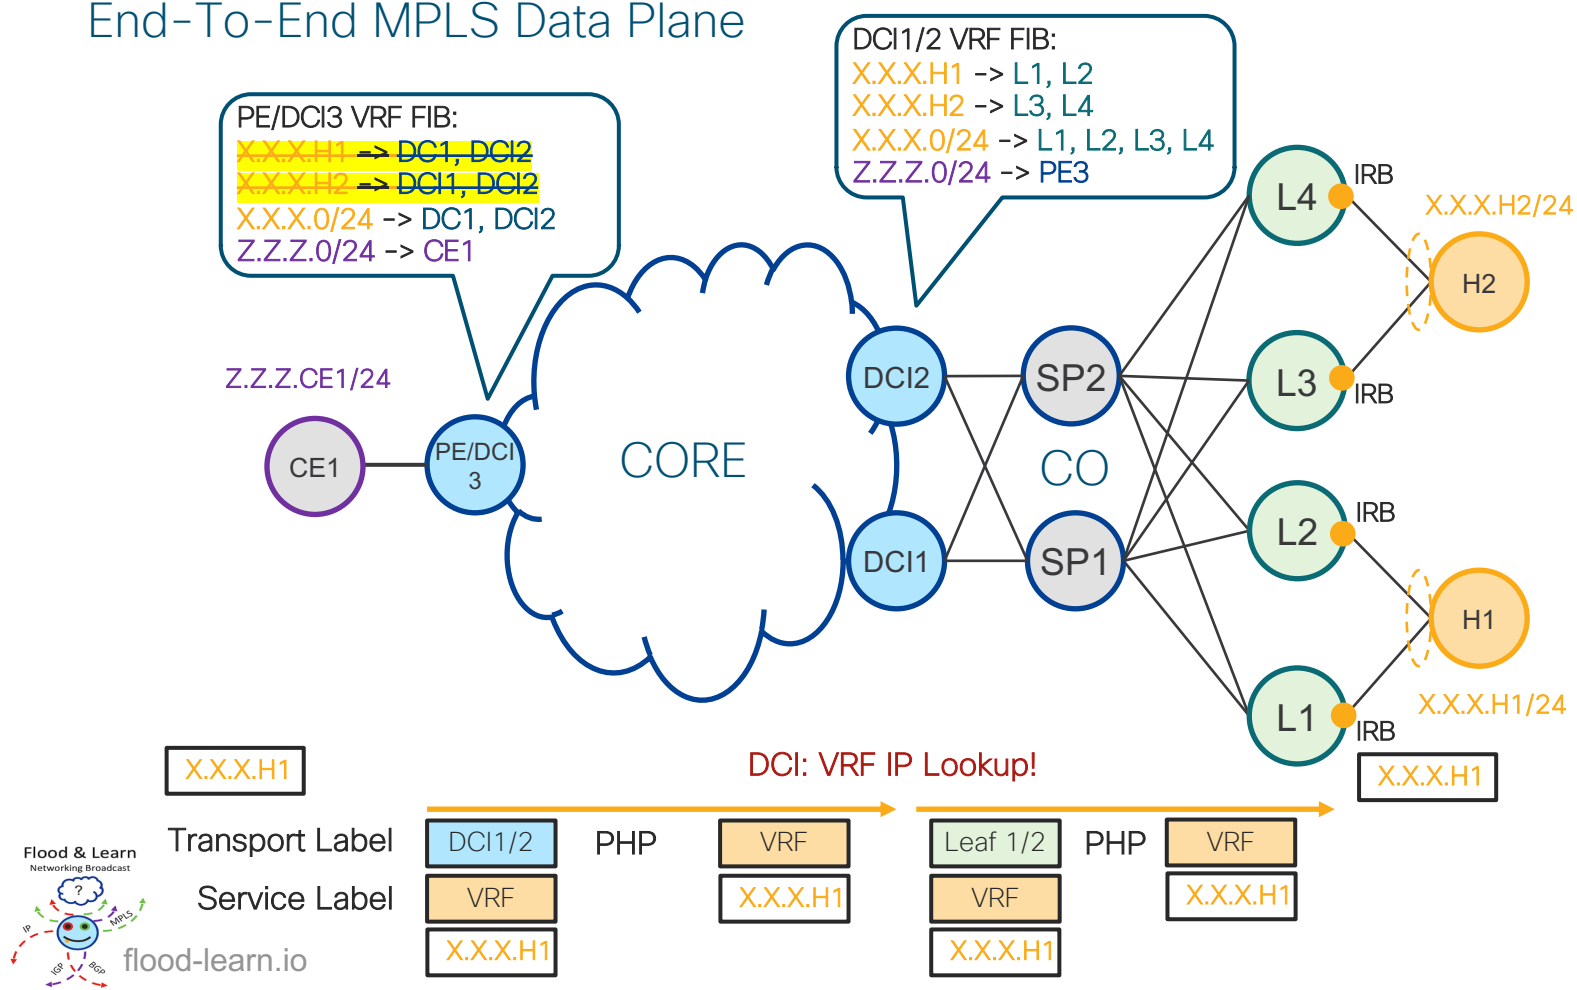

Flood & Learn
Networking Broadcast

flood-learn.io

BGP Layer3 Interconnect

End-To-End MPLS Data Plane

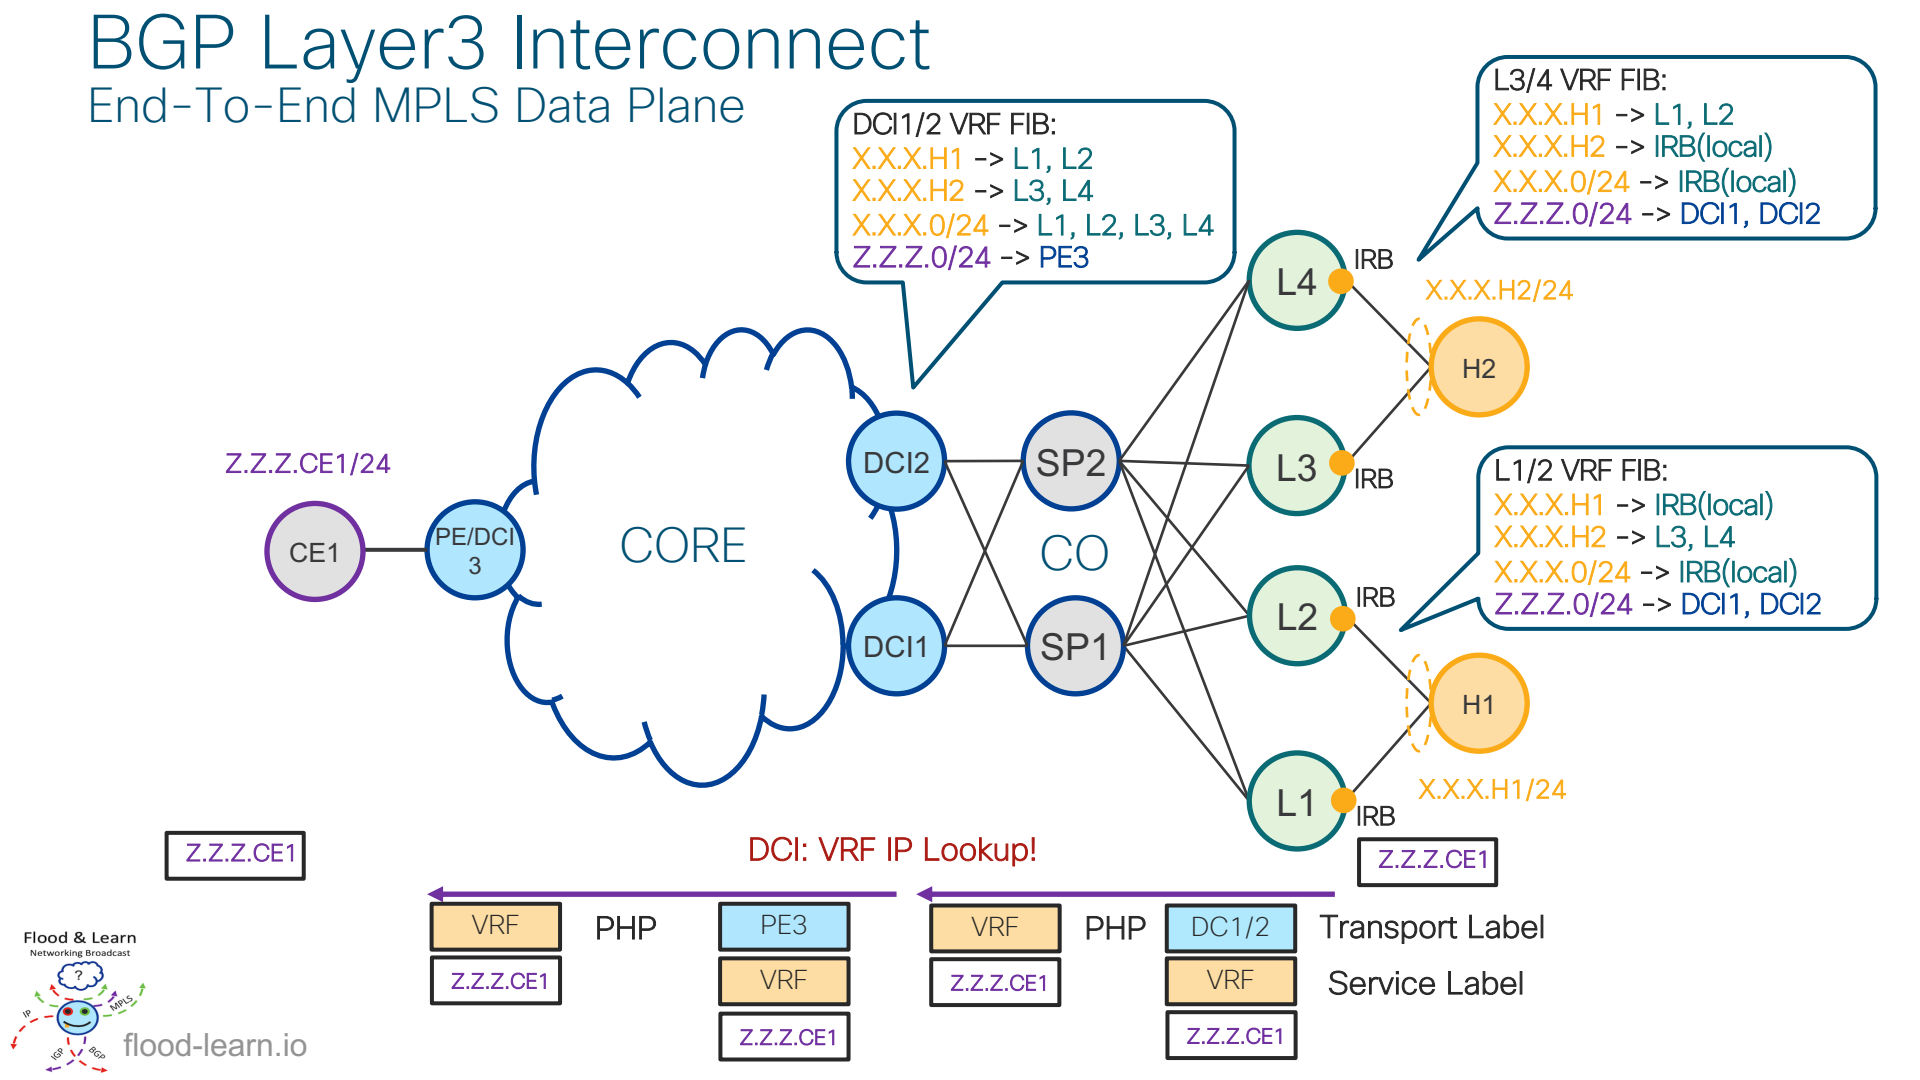

BGP Layer3 Interconnect

CO - EVPN VXLAN Data Plane

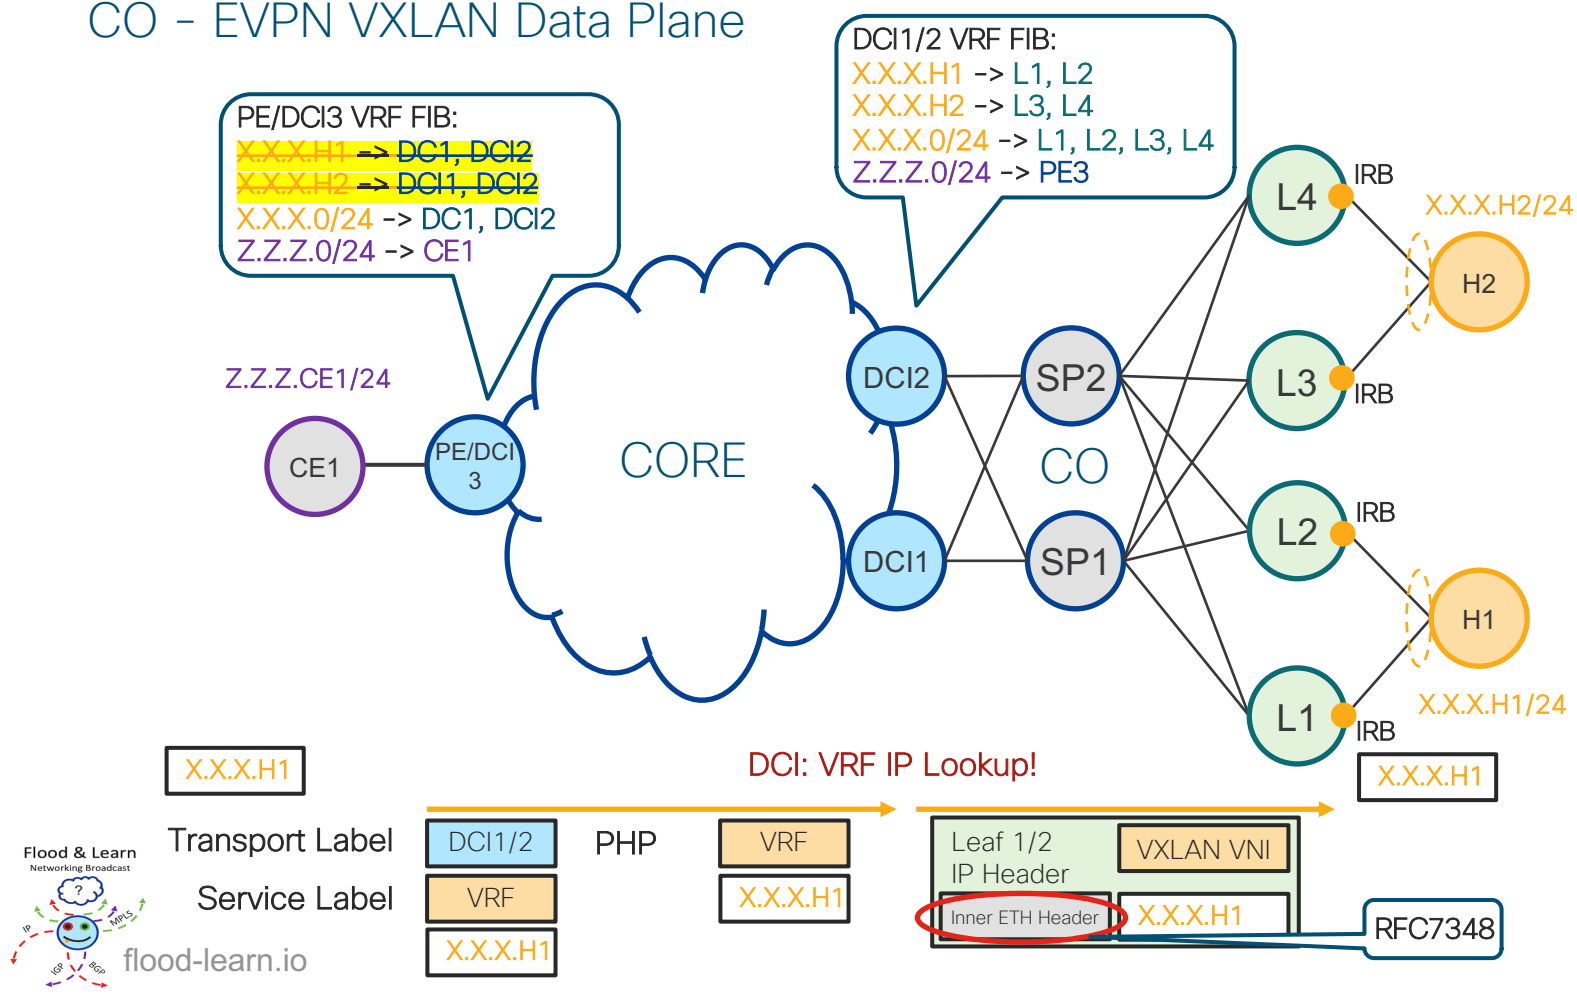

Flood & Learn
Networking Broadcast
flood-learn.io

BGP Layer3 Interconnect

CO - EVPN VXLAN Data Plane

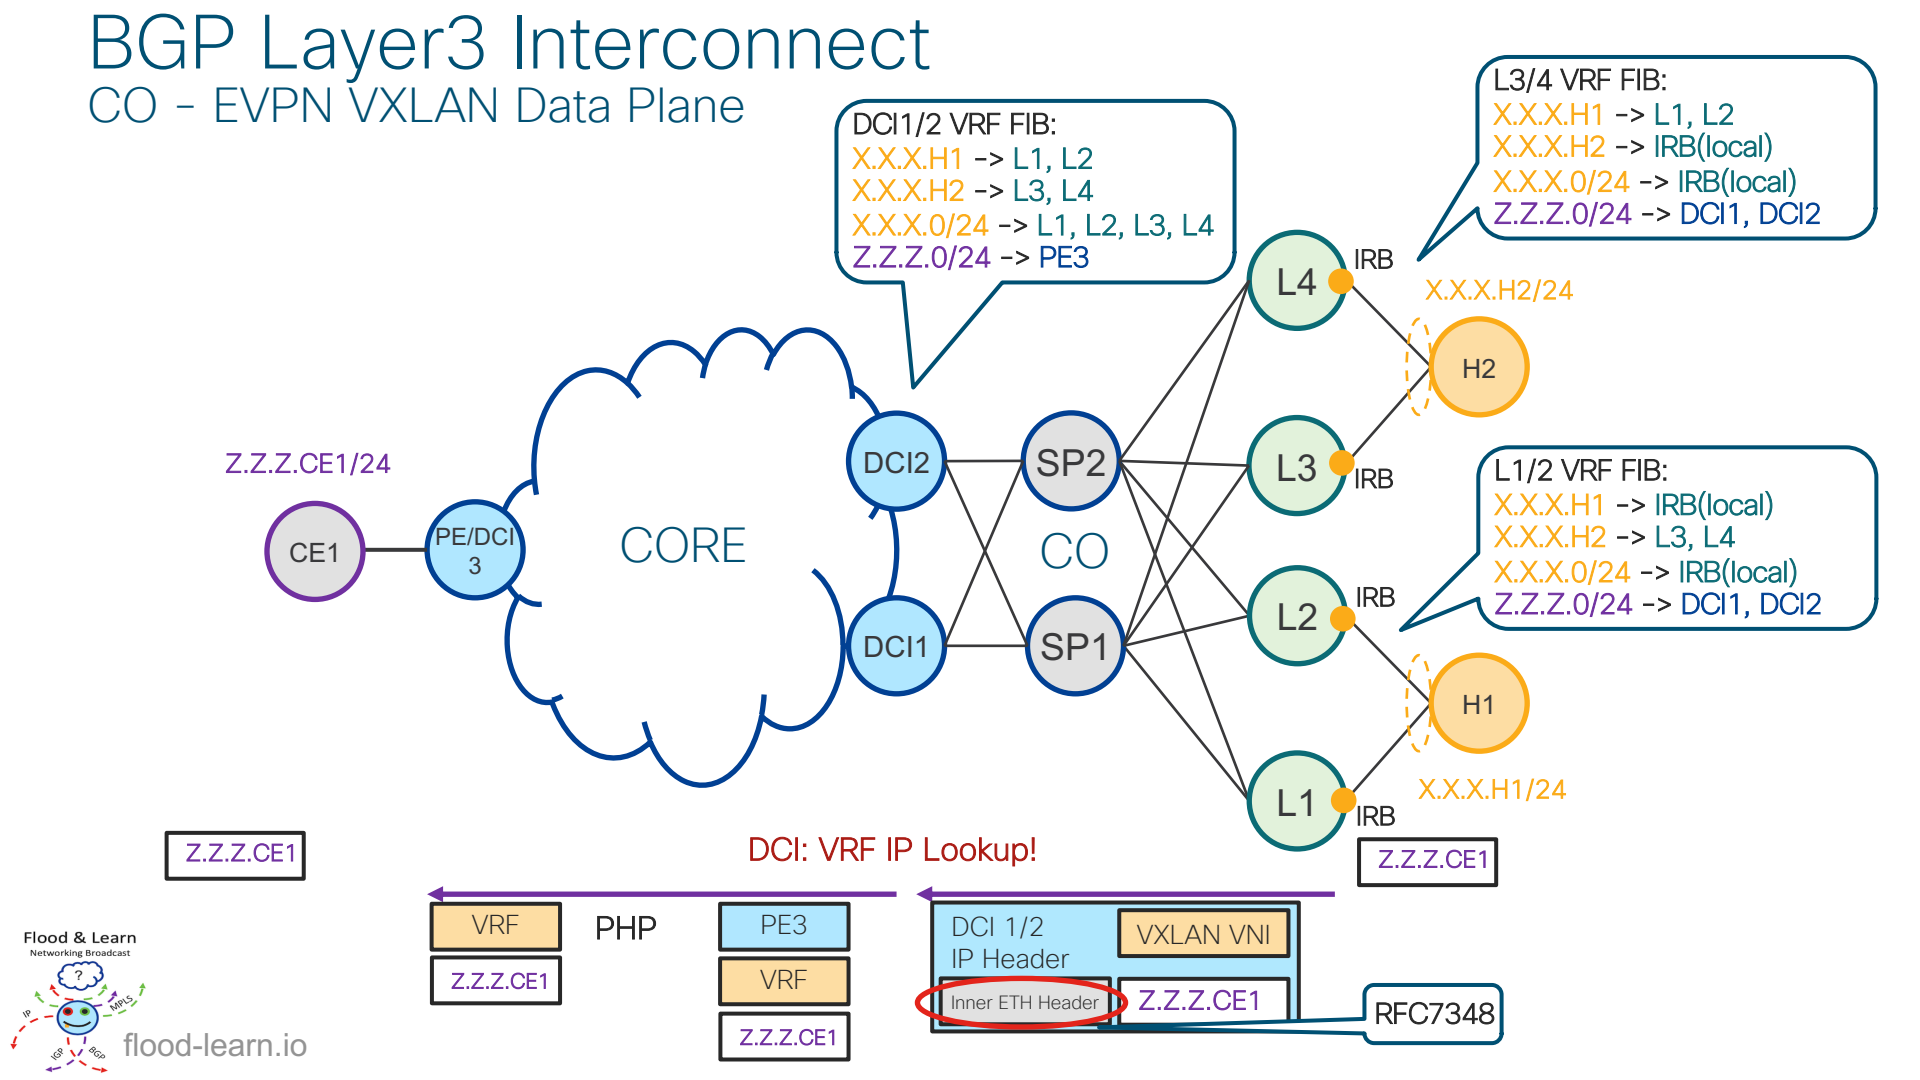

Flood & Learn
Networking Broadcast

flood-learn.io

BGP Layer3 Interconnect

End-To-End SRv6 Data Plane

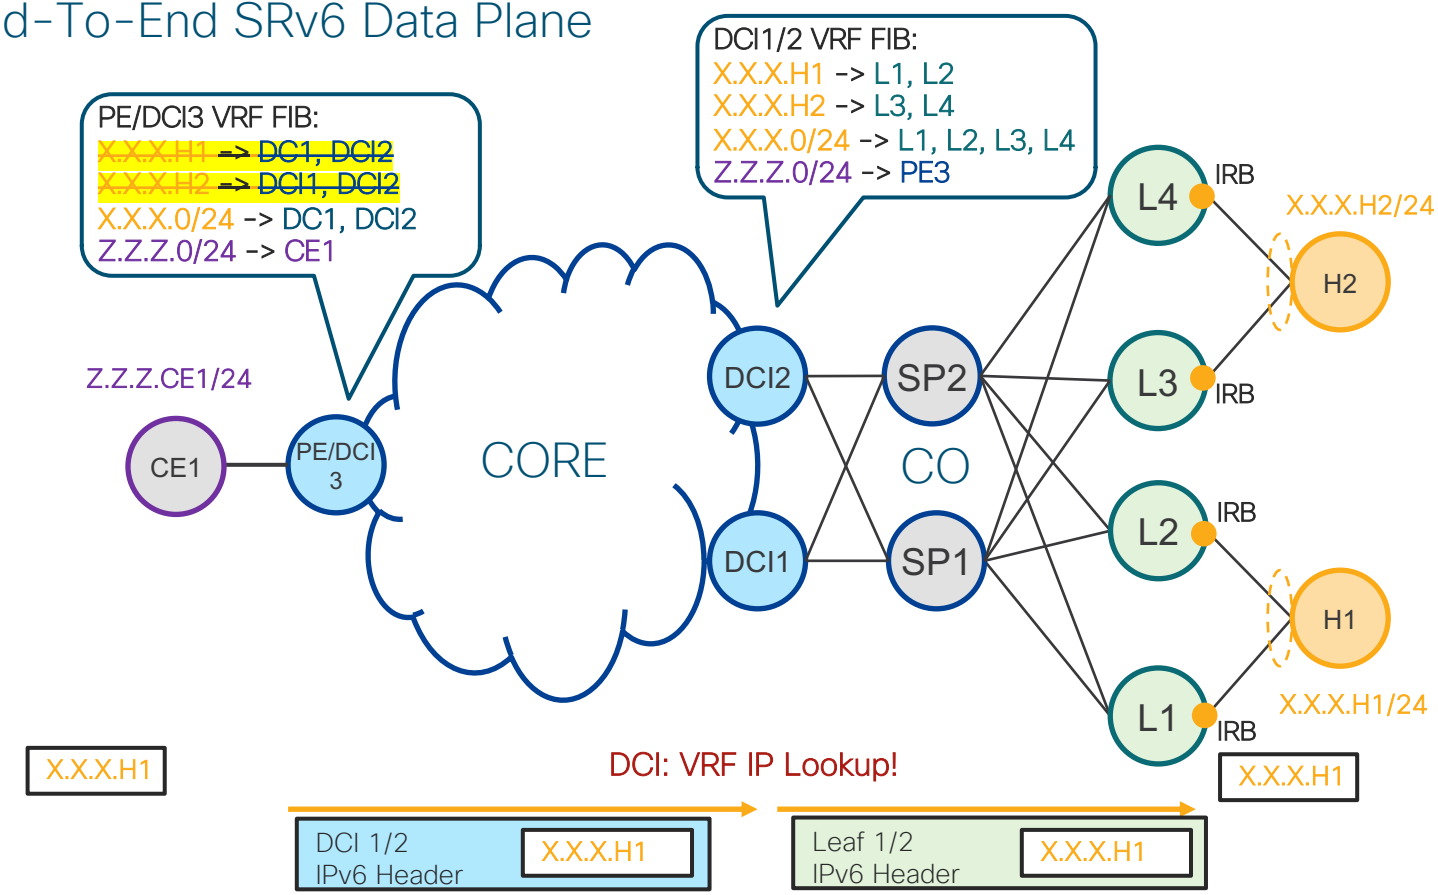

BGP Layer3 Interconnect

End-To-End SRv6 Data Plane

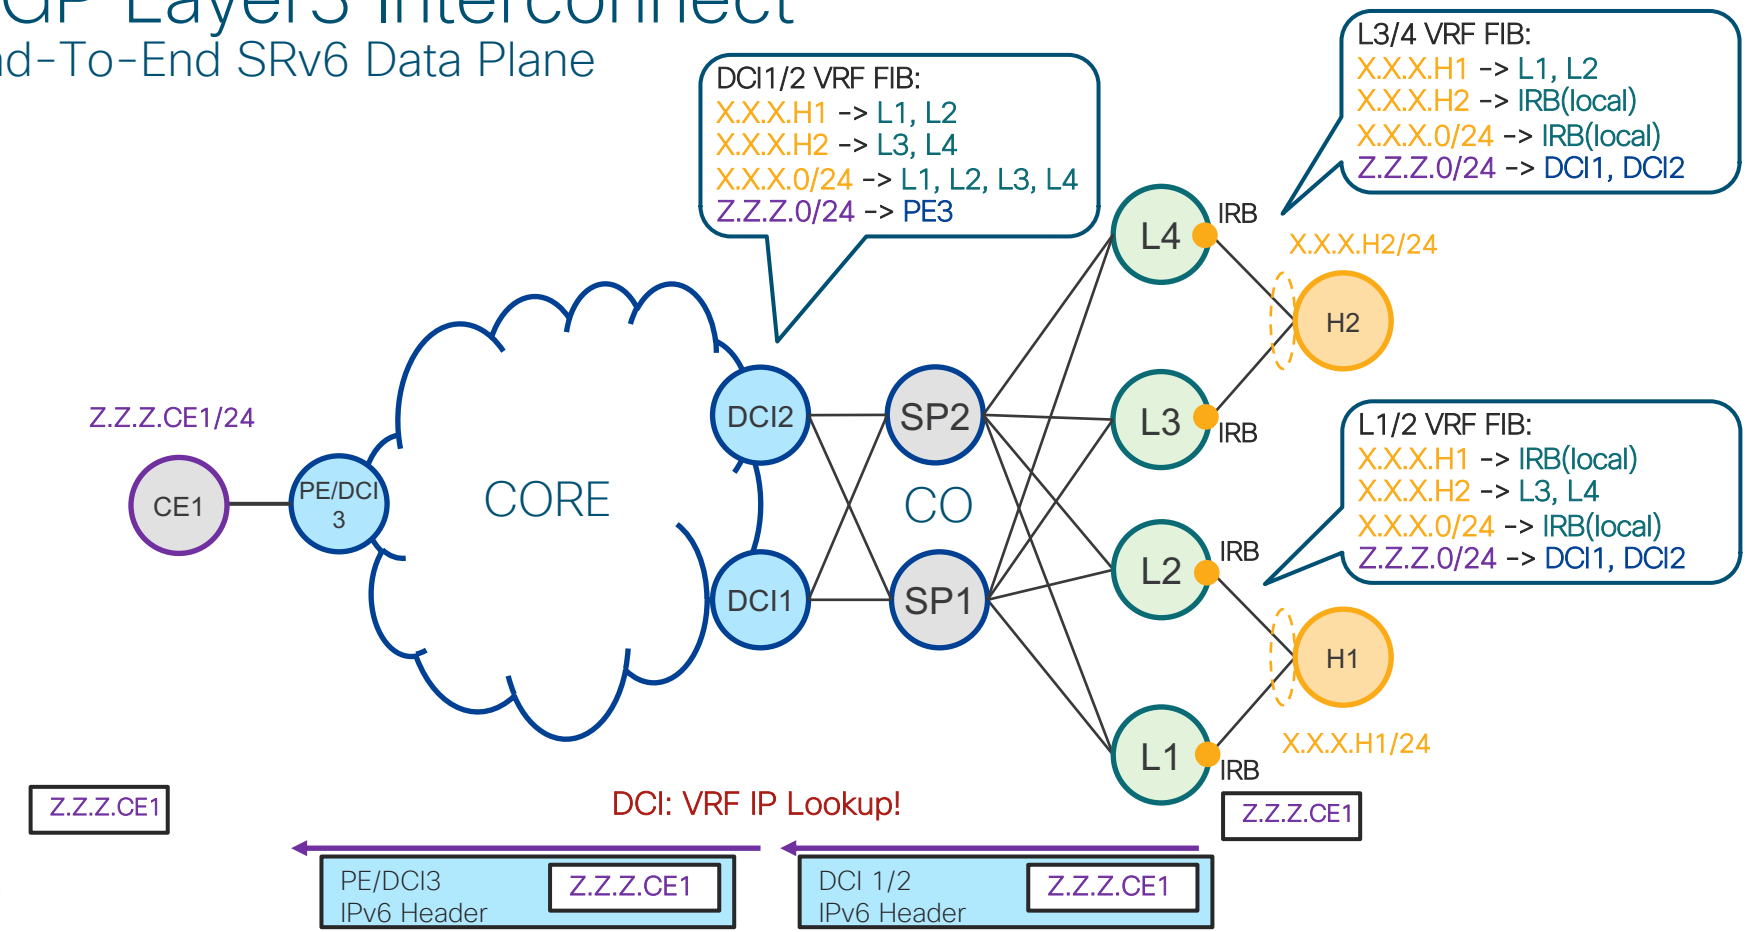

BGP Layer3 Interconnect
Data Plane Highlight - MPLS

- MPLS Data Plane

  + The packet structure is always identical, regardless of BGP VPNv4/6 or L3 EVPN Control Plane
    Less Complexity, Simple Troubleshooting
  + MPLS Load-Balancing (ECMP) by Inner IP Header Lookup
  + Segment Routing provides Traffic Engineering and Fast Re-Reroute (FRR) capability

BGP L3 EVPN or VPNv4/6 MPLS Packet

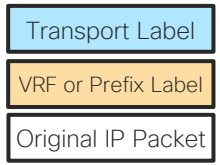

BGP Layer3 Interconnect

Data Plane Highlight - IP

- VXLAN Data Plane – RFC7348
  - EVPN Signaling only
  - RFC7348 requires Inner Ethernet encapsulation => Unnecessary overhead for L3 Forwarding

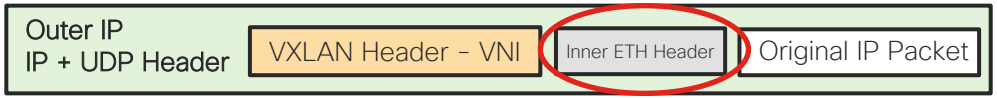

  - Inner Ethernet Header encapsulation/decapsulation typically done by Integrated Routing and Bridging (IRB) Interface
    IRB requires Bridge-Domain
    DCI doesn't participate in L2 Forwarding => Bridge-Domain (BD) requires unnecessary HW resources

  + VXLAN draft-ietf-nvo3-vxlan-gpe can simplify

- SRv6
  + Transport and Service is integrated in Outer IPv6 Header
  + The packet structure is always identical, regardless of BGP VPNv4/6 or L3 EVPN Control Plane
    Less Complexity, Simple Troubleshooting

  + Load-Balancing (ECMP) by Flow-Label in outer IPv6 header
  + Doesn't require additional header compared to VXLAN
  + Same Principles as Segment Routing MPLS
    Optional Segment Routing Header (SRH) can extend Traffic Engineering, Service Chaining and Fast Re-Reroute (FRR) capabilities

Conclusion

- Data Center Interconnect (DCI) is required for IP summarization
- EVPN is not strictly a replacement of “traditional” VPNv4/6
  - EVPN and VPNv4/6 can coexist
- Service Layer is Data Plane independent, but the right Data Plane selection decreases complexity and provides additional capabilities

EVPN - Stay Up-To-Date

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