



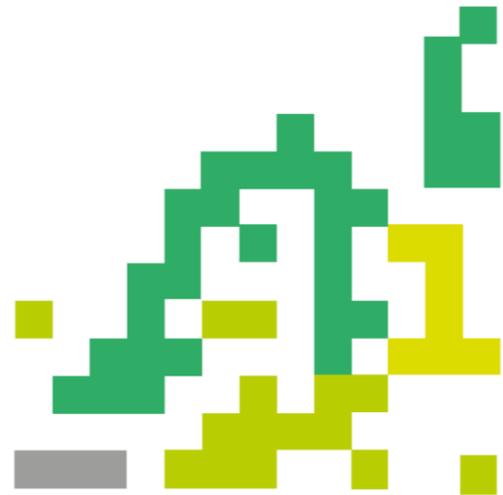
BGP EVPN - One Control Plane to Rule Them All

CSNOG 2026 @ Zlín

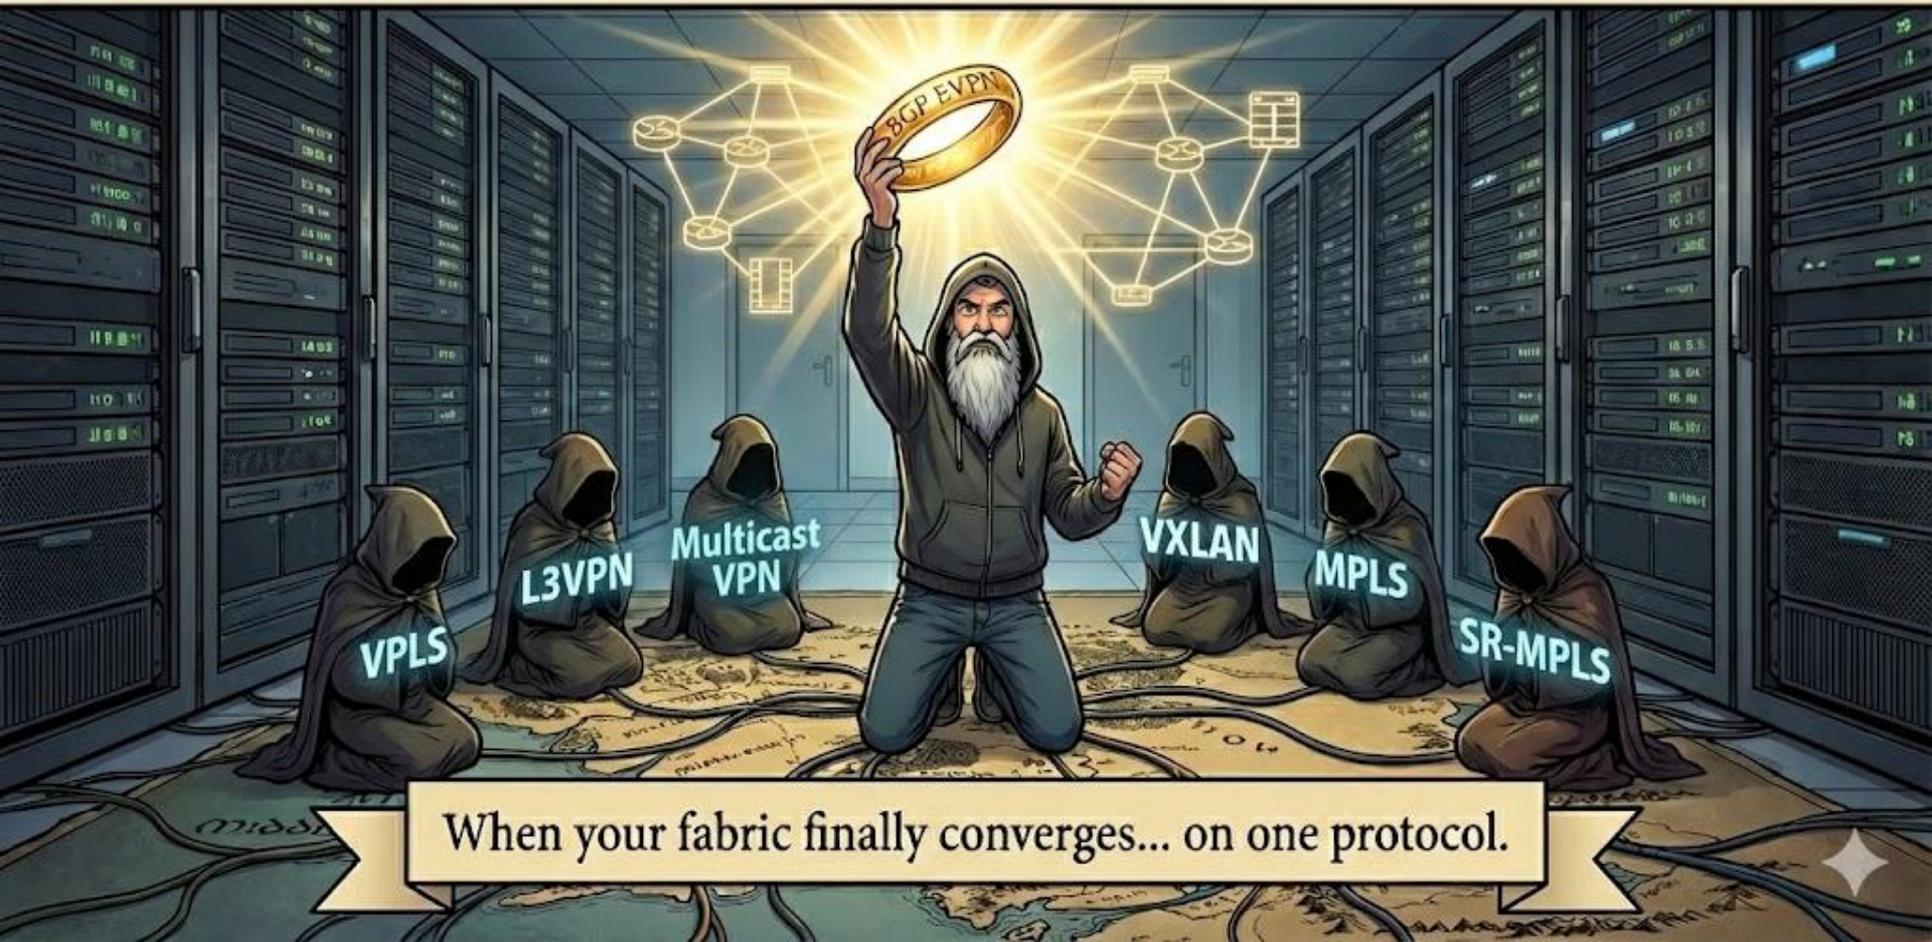
Radim Roška

22.1.2026

ALTEPRO



BGP EVPN – One Control Plane to Rule Them All



When your fabric finally converges... on one protocol.

BGP EVPN - history - context

- **2006/2007** - VPLS Martini (LDP) / Kompella (BGP AD)
- **2013** - Requirements for a VPLS successor (EVPN): RFC 7209
- **2015** - „EVPN core“ standard: RFC 7432 (BGP MPLS-Based Ethernet VPN)
- **2018** - EVPN as overlay control-plane: RFC 8365 (*EVPN for NVO3 overlays*) - Describes EVPN over VXLAN/NVGRE/MPLSoGRE (practical basis for EVPN-VXLAN fabrics)

- **2026** - Have you heard about BGP EVPN beyond the data center?

Why EVPN Exists: Fixing the VPLS Era Problems

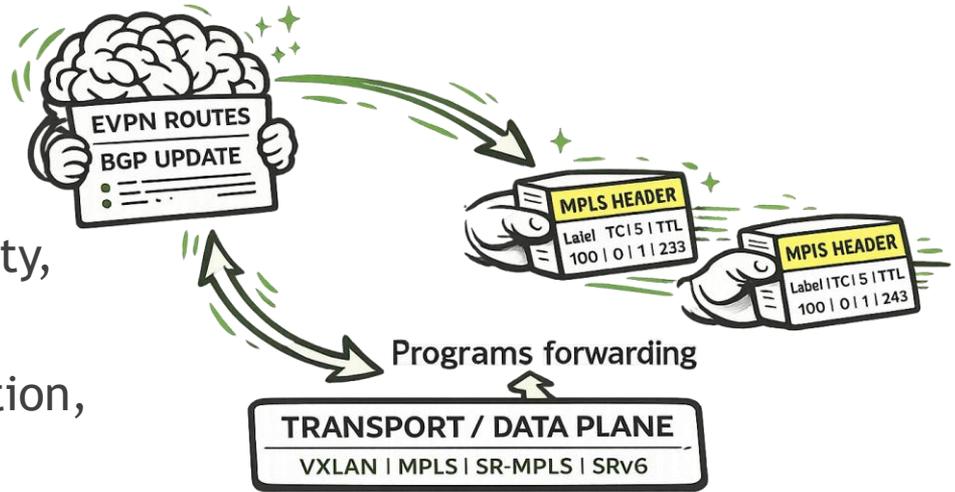
Flood & Learn → Advertise & Program

- Flood & learn for unknown unicast / ARP / ND (BUM pressure)
- MAC churn & large tables → operational pain at scale
- Multi-homing is awkward (MC-LAG, STP-ish thinking, split-horizon tricks)
- Convergence depends on learning/aging + underlay events
- Scale limits (e.g., PW mesh / signaling complexity)
- MAC/IP reachability advertised in BGP (less flood, more knowledge)
- Integrated multi-homing (ESI/DF) → active/active without hacks
- ARP/ND suppression (less BUM)
- Better policy & isolation (RD/RT model, consistent with L3VPN ops)
- Cleaner troubleshooting: “Do I see the EVPN routes?” → then dataplane



So EVPN ≠ VXLAN

- BGP EVPN = control plane
- VXLAN / MPLS / SR = data plane (transport)
- EVPN tells WHAT + WHERE (identity, reachability, redundancy)
- Transport solves HOW (encapsulation, forwarding)



How EVPN Works

1. BGP EVPN updates carrying service info

- EVPN NLRI = Route Type + RD + payload
- RT (Route Target) = who imports it

2. Working principle

Learn (MAC,IP) ~ “data plane learning”

→ Advertise (BGP EVPN)

→ Programs forwarding - “control plane based”

→ Forward (VXLAN/MPLS/SR)

EVPN MP-BGP Route Types

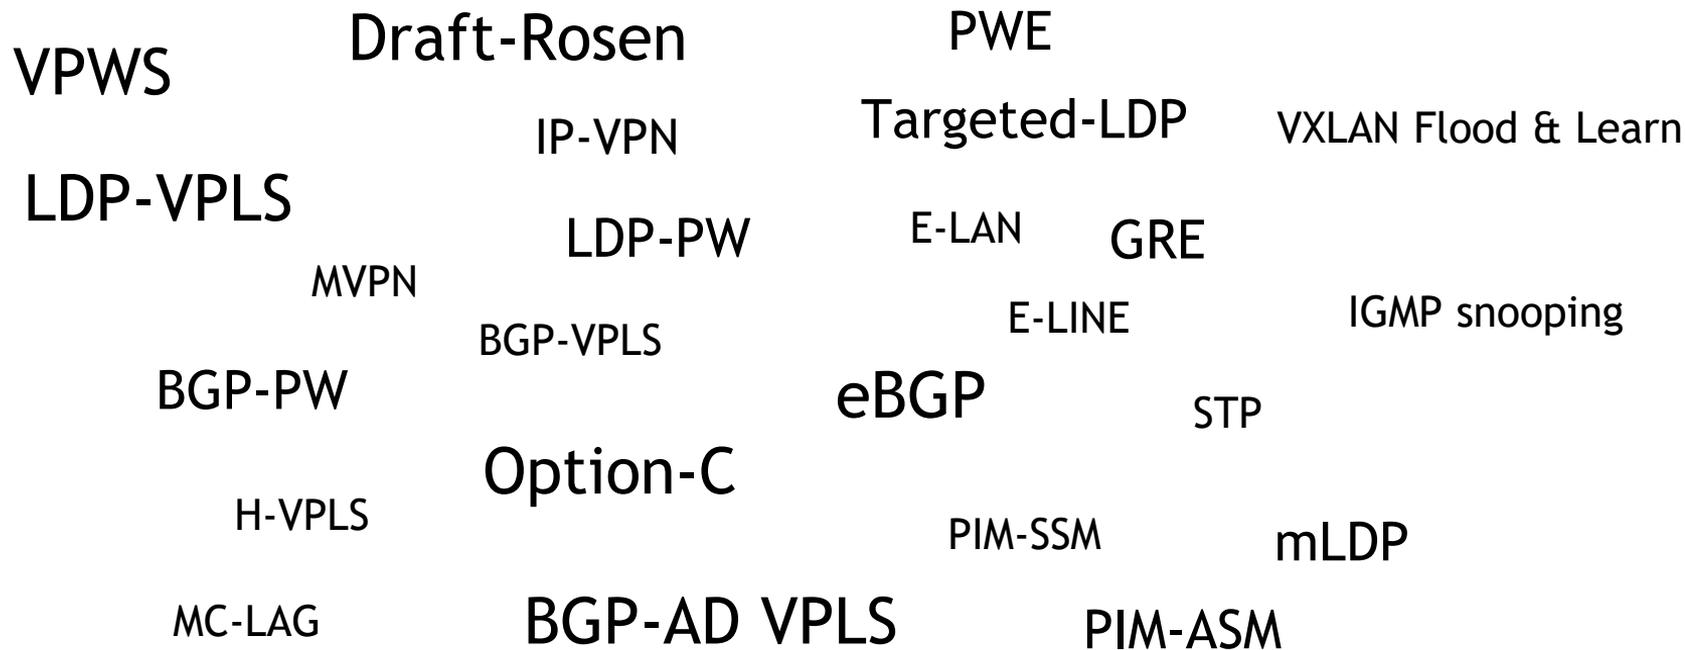
Used in multi-homed deployments

- Type 1 routes: **Ethernet A-D Routes** (A-D = Auto Discovery)
 - Are used to announce the reachability of multi-homed Ethernet segments
 - Mass withdrawal => fast convergence
 - Aliasing
- Type 4 routes: **Ethernet Segment Routes**
 - Used in multi-homed topologies to **discover remote PEs/VTEPs** on the same shared Ethernet segment and **elect a DF**

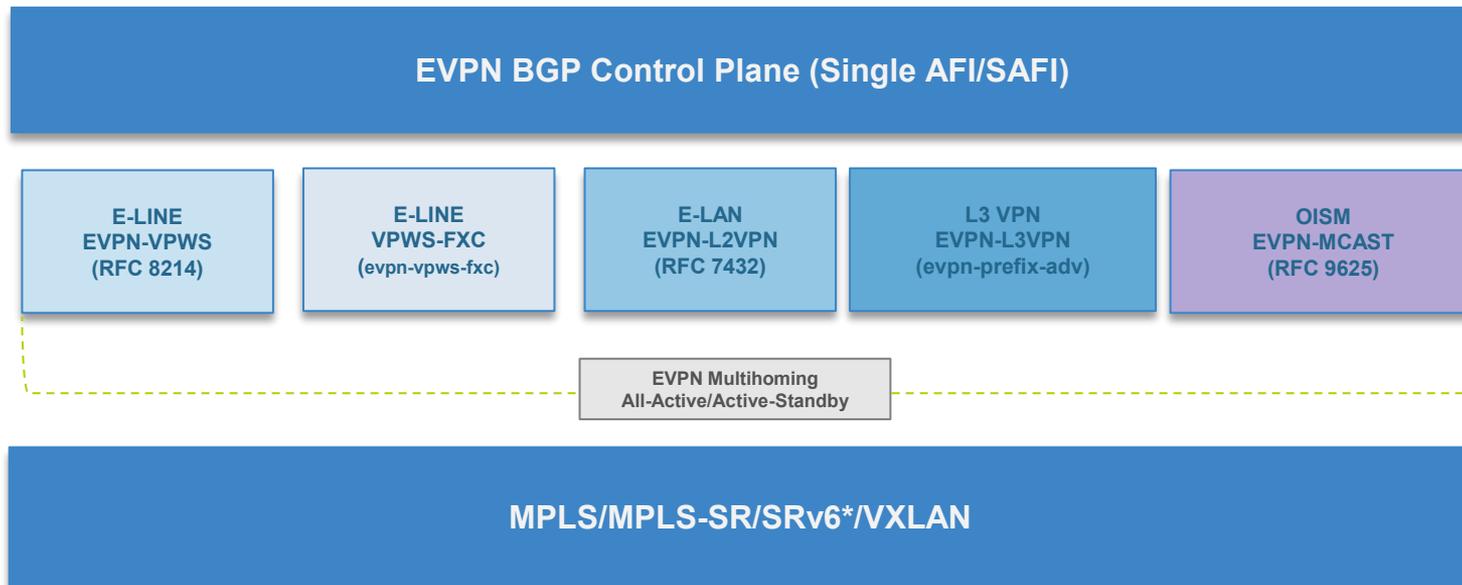
Multicast Optimization

- RT 6, 7, 8: IGMP/MLD State Sync (L2 efficiency)
- RT 9, 10, 11: Transport Efficiency (Ingress Rep. vs Underlay Multicast)

The Service Control-Plane Zoo

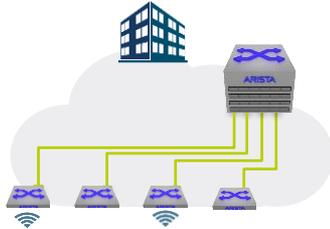


EVPN provides a flexible foundation for the broad range of services



Enterprise Routing - Network as a Platform

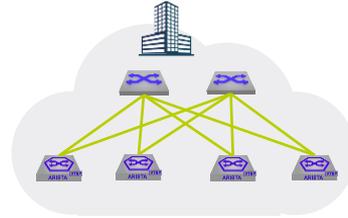
Enterprise Campus



➔ EVPN

➔ VXLAN

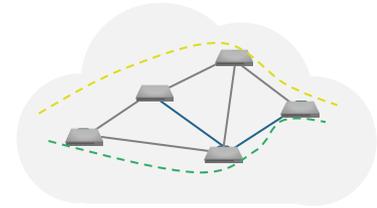
Datacenter



➔ EVPN

➔ VXLAN

WAN/Backbone



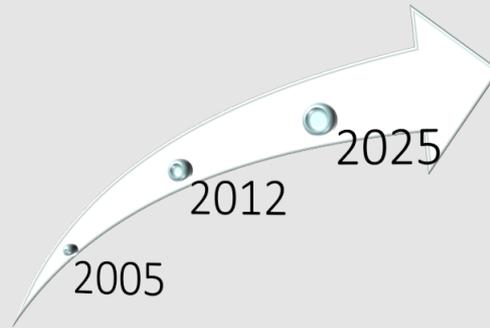
➔ EVPN
➔ IP-VPN
➔ MVPN

➔ VXLAN
➔ MPLS
➔ MPLS-SR
➔ SRv6

Reference case - Rowanet - EVPN + VXLAN

Network history

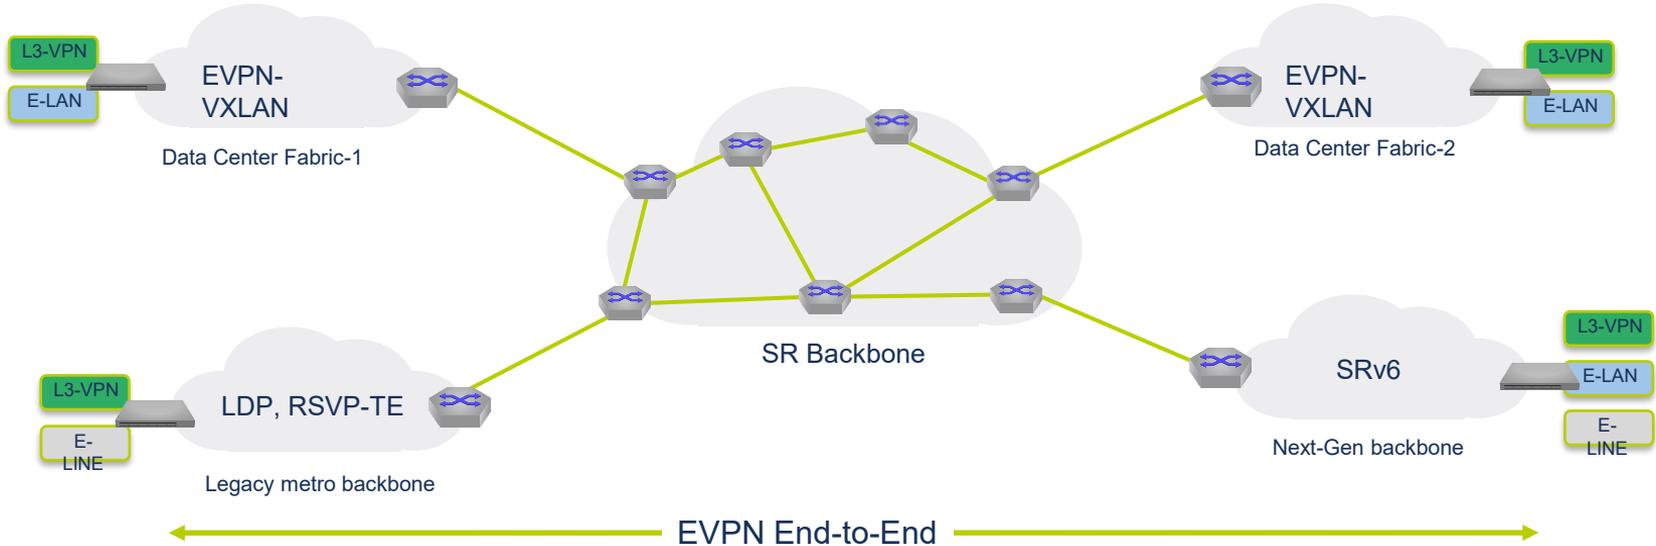
- **2005 – ROWANet I.**
 - 1x central device CISCO Catalyst 6500, 7x PE CISCO Catalyst 3750ME
 - VLAN 802.1Q, backbone lines speed 100Mb/s or 1Gb/s
 - Star topology, transmission technology CWDM
- **2012 - ROWANet II.**
 - 2x central device HP A7500
 - 12x PE HP A5800, 70x CE CISCO SG300
 - MPLS, backbone lines speed 10Gb/s or 1Gb/s
 - Star topology, transmission technology CWDM
- **2025 - ROWANet III.**
 - ARISTA 7280 – 2x CORE , 28xPE, 207x CE CISCO boxes
 - backbone lines speed 100Gb/s
 - EVPN + VXLAN, ring topology, transmission technology DWDM



EVPN rules, but what about data plane?

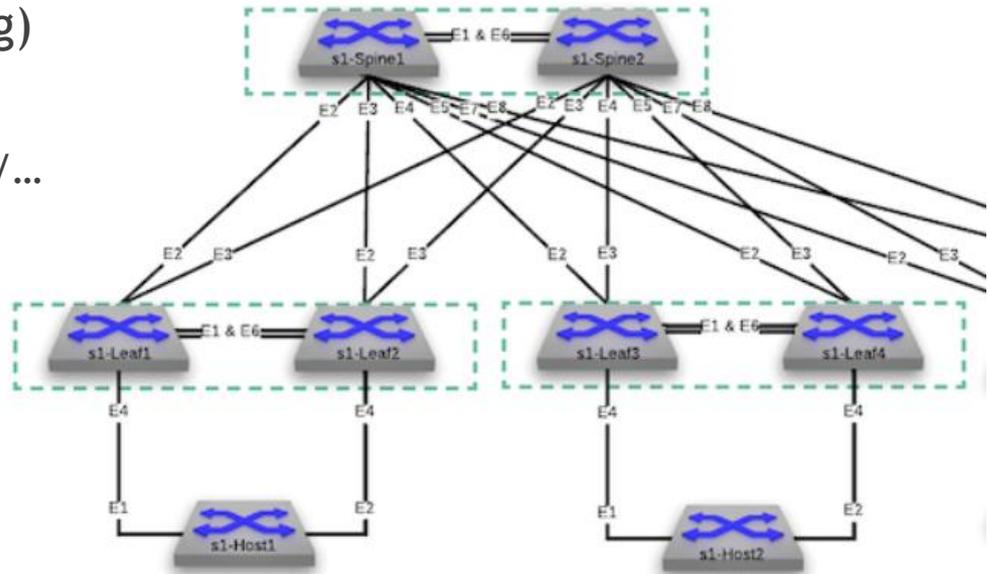
DECISION DRIVERS	EVPN/VXLAN	EVPN/MPLS	EVPN/SR-MPLS	EVPN/SRv6
Ethernet-centric fabric (mainly Campus/DC overlay)	✓	!	!	!
WAN services & provider-style operations (InterAS, BGP-LU, historically MPLS based)	!	✓	✓	!
Greenfield simplicity	✓✓ (IP underlay)	!	✓	!
Traffic Engineering / path steering	✗	! (needs RSVP-TE)	✓ (native)	✓
Fast reroute / deterministic protection	! (BFD + IP FRR)	✓ (RSVP-TE FRR)	✓ (TI-LFA)	✓
“Wire-like” services (VPWS / PW) (L2 transparency)	!	✓	✓	!
Non-Ethernet legacy payloads (PWE3: TDM/ATM/etc.)	✗	✓	✓	✗
OAM maturity / troubleshooting ergonomics (LSP ping/traceroute, MPLS OAM, TWAMP,...)	!	✓	✓	!
Encap overhead / MTU sensitivity across “complicated” paths	! ~ 54B	✓ ~ 8-12B typical	✓ ~8-12B typical	! ~40B+ x 8B

Why Do We Need EVPN Gateways?



DEMO - EVPN / VXLAN

- Same principles apply to MPLS/SR...
- ARISTA CVP - creating the services & see the CLI difference
- *shows* in CLI (seeing is believing)
 - show bgp evpn
 - show bgp evpn route type mac-ip/...
 - show vxlan address-table
 - show mac address



Conclusion

1. EVPN is **the** control plane for L2/L3/MC services. Transport is a choice
2. RT2+RT5 are the base; RT1/4 = multihoming; RT3 = BUM
3. Same mental model across campus/DC/WAN

From “Flood & Pray” → “Advertise & Program”

Conclusion



*ONE PROTOCOL TO RULE THEM ALL. ONE TO FIND THEM.
ONE TO BRING THEM ALL, INTO FABRIC BIND THEM.*



Thank you
www.altepro.cz
Radim Roška



ALTEPRO solutions a.s.
Na Maninách 1092/20, Praha 7



+420 220 611 111
info@altepro.cz