ribbon

Technical White Paper

Version 1.1 · July 2026

EVPN Services over MPLS

Technology Overview and Deployment on Ribbon NPT Platforms

Your Path to Autonomonus Networks

Copyright and Disclaimer

This document is provided for informational purposes only and is subject to change without notice. It does not constitute a commitment on the part of Ribbon. Only a definitive contract executed by the authorized representatives of Ribbon and the purchaser shall be binding on the parties.

While reasonable efforts have been made in the preparation of this document to ensure its accuracy, Ribbon assumes no liability resulting from technical or editorial errors or omissions, or from any damages whatsoever resulting from the furnishing, performance, or use of the information contained herein. Ribbon reserves the right to make changes to this document and to Ribbon products without notice in its sole discretion.

Version	Date	Description
1.0	July 2026	Initial release
1.1	July 2026	Revised introduction; expanded Type-5 route description; VLAN mode use-case guidance; theory and lab parts; multihoming mode selection guidance; EVPN vs VPLS comparison added; spreadsheet results; strengthened conclusion

Contents

Introduction....4
Evolution of Carrier Ethernet Services....4
Part I — EVPN Technology Essentials....5
EVPN Fundamentals....5
Key Concept....5
MAC Learning in EVPN....6
EVPN Route Types....6
Type-1 — Ethernet Auto-Discovery....6
Type-2 — MAC/IP Advertisement....7
Type-3 — Inclusive Multicast Ethernet Tag....7
Type-4 — Ethernet Segment....7
Type-5 — IP Prefix....7
VLAN Service Modes....8
Which Mode for Which Use Case....9
Part II — EVPN in Practice: Lab on NPT Platforms....10
Lab Deployment on NPT Platforms....10
Topology and Scenario....10
Choosing a Multihoming Mode....11
Aggregate Interfaces....11
EVI, L3 VRF, and IRB....13
BGP Configuration....15
Ethernet Segments....15
Service Validation....17
Failure Emulation....21
EVPN and VPLS: Feature Comparison....22
Conclusion....23

Introduction

Ethernet services remain the backbone of enterprise connectivity, mobile transport, and wholesale offerings — but the expectations placed on them have changed dramatically. Customers demand that every purchased link carries traffic, that a failure anywhere in the network is invisible to their applications, and that Layer 2 and Layer 3 services can be combined without operational gymnastics. Networks built on legacy VPLS struggle to meet these expectations: paid-for standby links sit idle, a single link flap can trigger network-wide MAC re-learning, and every new service adds manual provisioning effort.

Ethernet VPN (EVPN) changes the economics and the operational model of Ethernet services. By moving MAC learning from data-plane flooding into the BGP control plane, EVPN unlocks capabilities that VPLS was never able to deliver: true all-active multihoming, sub-second convergence, built-in auto-discovery, and seamless integration of bridging and routing in a single service. It is the technology the industry has converged on — in service provider metro networks and data centers alike.

This white paper is a practical guide to EVPN on Ribbon's NPT platforms. It is organized into two parts. Part I explains the technology: the EVPN building blocks, BGP route types 1 through 5, and the three VLAN service modes, with guidance on when to use each. Part II proves it in practice: a complete EVPN deployment in a seven-node NPT lab — configuration, validation, and measured failure convergence — followed by a direct feature comparison of EVPN and VPLS. The paper is written for network architects and engineers who plan, deploy, or operate Carrier Ethernet services.

Evolution of Carrier Ethernet Services

Ethernet has come a long way since the first research was published by Robert Metcalfe in 1973: from a shared laboratory medium to the dominant technology for carrier-grade Layer 2 services. To extend Ethernet across provider networks, the industry adopted Virtual Private LAN Service (VPLS), available in two flavors: LDP-signaled VPLS (RFC 4762), which is simple to configure but offers no auto-discovery, and BGP-signaled VPLS (RFC 4761), which adds automatic neighbor discovery through the BGP control plane.

Years of operating VPLS at scale exposed several structural limitations:

- Multihomed sites cannot use all attachment links at the same time — one link carries the traffic while the others stay idle.
- MAC addresses are learned in the data plane by flooding, which increases Broadcast, Unknown unicast, and Multicast (BUM) traffic across the network.
- Layer 3 capabilities are limited, complicating designs that combine bridging with routed access to external networks.

These limitations prompted the industry to design a successor rather than another patch. Ethernet VPN (EVPN, RFC 7432) keeps the familiar service model of VPLS but replaces its weakest element: MAC addresses are advertised between PE routers in the BGP control plane instead of being flooded in the data plane. This single architectural change enables all-active multihoming, fast convergence, integrated Layer 2 and Layer 3 services, and much better control over BUM traffic. A detailed feature-by-feature comparison of the two technologies is provided near the end of this document.

Part I EVPN Technology Essentials

This part explains how EVPN works: the core concepts and terminology, the BGP route types that form its control plane, and the VLAN service modes with guidance on selecting the right one. No lab equipment is required to follow it.

EVPN Fundamentals

Key Concepts

A typical EVPN service connects customer sites across an MPLS core, as shown in Figure 1. The following terms are used throughout this document:

- **EVI (EVPN Instance)** — the routing and forwarding instance of an EVPN service on a PE, also known as a MAC-VRF (RFC 7432). It holds the MAC address table of the service.
- **Ethernet Segment (ES) and ESI** — the set of links that connects a customer site to one or more PEs. When a site is connected to two or more PEs, the segment is identified by a non-zero Ethernet Segment Identifier (ESI) shared by those PEs.
- **Multihoming modes** — a CE may be single-homed (one PE) or multihomed. On NPT platforms a multihomed segment can operate in single-active, port-active, or active-active (all-active) mode, defining how many links forward traffic at the same time.
- **Designated Forwarder (DF)** — the PE elected to forward BUM traffic arriving from the MPLS core towards a multihomed segment. The election prevents duplicate delivery of the same frame.
- **Control plane and data plane separation** — local MAC addresses are learned from the attachment circuit in the data plane; remote MAC addresses are learned from BGP advertisements in the control plane.

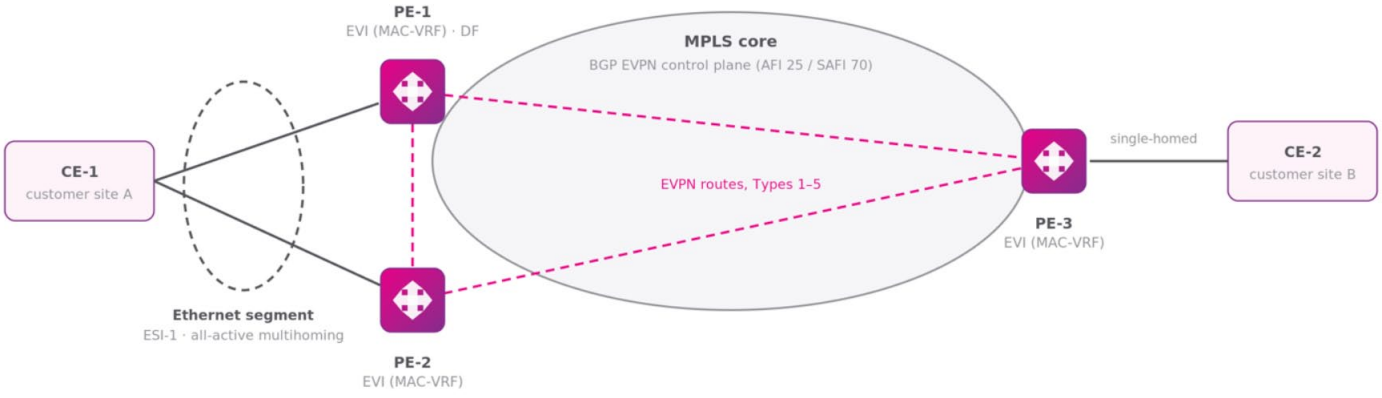

Figure 1 — EVPN reference model — multihomed and single-homed sites over an MPLS core

MAC Learning in EVPN

Consider a station behind PE-3 sending its first frame to a station behind the multihomed segment ESI-1. Because the destination is unknown, the station first issues a broadcast ARP request. PE-3 floods this broadcast only to the PEs that have advertised readiness to receive BUM traffic for this broadcast domain (route Type-3), and only the elected DF forwards it to the customer site. The ARP reply returns as unicast; each PE learns the local MAC address from its attachment circuit, advertises it in BGP (route Type-2), and every other PE installs the MAC with its MPLS forwarding label. From this moment traffic flows as labeled unicast — no further flooding is required.

This is the fundamental difference from VPLS, where every unknown destination is repeatedly discovered by flooding. In EVPN the network floods once, learns in the control plane, and then forwards deterministically.

EVPN Route Types

EVPN defines several BGP route types (NLRIs) carried in the L2VPN EVPN address family (AFI 25, SAFI 70). Types 1–5 are the foundation of every EVPN deployment; types 6–8 extend EVPN with multicast optimizations (RFC 9251), and types 9–11 are being standardized as part of the BUM procedure updates draft.

Type	Name	Defined In	Purpose
1	Ethernet Auto-Discovery	RFC 7432	Multihoming: fast convergence, split horizon, aliasing
2	MAC/IP Advertisement	RFC 7432 / RFC 9135	MAC (and optional IP) reachability
3	Inclusive Multicast Ethernet Tag	RFC 7432	BUM traffic delivery tree
4	Ethernet Segment	RFC 7432	ES discovery and DF election
5	IP Prefix	RFC 9136	IP prefix reachability, inter-subnet routing
6-8	Multicast (SMET, Join/Leave Synch)	RFC 9251	IGMP/MLD proxy and multicast optimization
9-11	Per-Region I-PMSI, S-PMSI, Leaf A-D	IETF draft	BUM procedure updates (in standardization)

The following sections describe route types 1 through 5 in ascending order.

Type-1 — Ethernet Auto-Discovery

A Type-1 route is announced only when the PE participates in a multihomed Ethernet segment (non-zero ESI). It serves three purposes:

- **Fast convergence (mass withdrawal)** — when a PE-CE link fails, the PE withdraws a single Type-1 route instead of thousands of individual Type-2 routes. Remote PEs immediately switch all affected MAC addresses to the remaining next hop.
- **Split horizon** — the route carries an ESI label used to make sure that BUM traffic originating from a segment is never sent back to the same segment through another PE, preventing Layer 2 loops.
- **Aliasing** — in all-active mode, remote PEs can load-balance traffic to all PEs attached to the segment, even to those that have not advertised a given MAC address themselves.

Type-2 — MAC/IP Advertisement

The Type-2 route is the main building block of an EVPN service: it advertises a MAC address (optionally together with an IP address) learned on an attachment circuit. The route is generated only after the MAC address has been learned locally in the data plane.

Mandatory fields include the Route Distinguisher (8 octets), Ethernet Segment Identifier (10 octets), Ethernet Tag ID (4 octets), MAC address length (1 octet), MAC address (6 octets), and MPLS Label1 (3 octets). Optional fields carry the IP address and its length, and a second label (Label2) used for inter-subnet routing when the route also carries IP information (symmetric IRB, RFC 9135).

Type-3 — Inclusive Multicast Ethernet Tag

The Type-3 route announces that a PE is ready to receive BUM traffic for a given broadcast domain, identified by the Ethernet Tag ID. It is present in every EVPN topology except EVPN-VPWS and appears immediately after the EVPN instance is configured, even before any customer traffic exists. A PE floods broadcast frames only to the PEs from which it has received a Type-3 route; with ingress replication, the advertised inclusive multicast label is used to build the replication list.

Type-4 — Ethernet Segment

Like Type-1, the Type-4 route is announced only by PEs that are part of a multihomed segment. It allows the PEs sharing the same ESI to discover each other and to elect the Designated Forwarder. The default election algorithm defined in RFC 7432 is modulo-based VLAN carving, which distributes the DF role for different VLANs across the PEs of the segment. NPT platforms additionally support preference-based DF election, which gives the operator direct control over which PE becomes the DF — this method is used in the lab example later in this document.

Type-5 — IP Prefix

The Type-5 route, defined in RFC 9136, carries IPv4 or IPv6 prefixes in the EVPN address family. While a Type-2 route can advertise a single host IP together with a MAC address, Type-5 advertises whole subnets and is fully decoupled from any specific MAC address or Ethernet segment. The route carries the Route Distinguisher, an optional ESI, the Ethernet Tag ID, the IP prefix with its length, an optional gateway IP address, and an MPLS label pointing to the L3 context (IP-VRF) of the advertising PE.

This decoupling is exactly what makes Type-5 valuable. A subnet can be advertised even when no host from it has been learned yet; the prefix survives host moves and MAC changes; and summarization keeps the routing table small — one Type-5 route can replace thousands of host-level Type-2 entries. Typical applications include:

- **Data Center Interconnect (DCI)** — Layer 2 domains in different data centers are stitched together, while Type- provides the routed connectivity between subnets that must not be bridged.
- **Inter-subnet routing at scale** — combined with anycast IRB gateways (as in the lab in Part II), Type-5 advertises the subnets of each bridge domain into a shared L3 VRF, enabling routing between EVPN services without a separate L3VPN overlay.
- **External and default routes** — prefixes learned from outside the fabric (including a default route to a border router or firewall) are injected into the EVPN domain as Type-5 routes, giving every PE a consistent exit path.

In short, Type-2 answers the question “where is this host?”, while Type-5 answers “where is this network?” — and a complete EVPN design normally uses both together.

VLAN Service Modes

RFC 7432 defines three ways to map customer VLANs to an EVPN instance, and all three are fully supported on NPT platforms: VLAN-based, VLAN-bundle, and VLAN-aware (Figure 2).

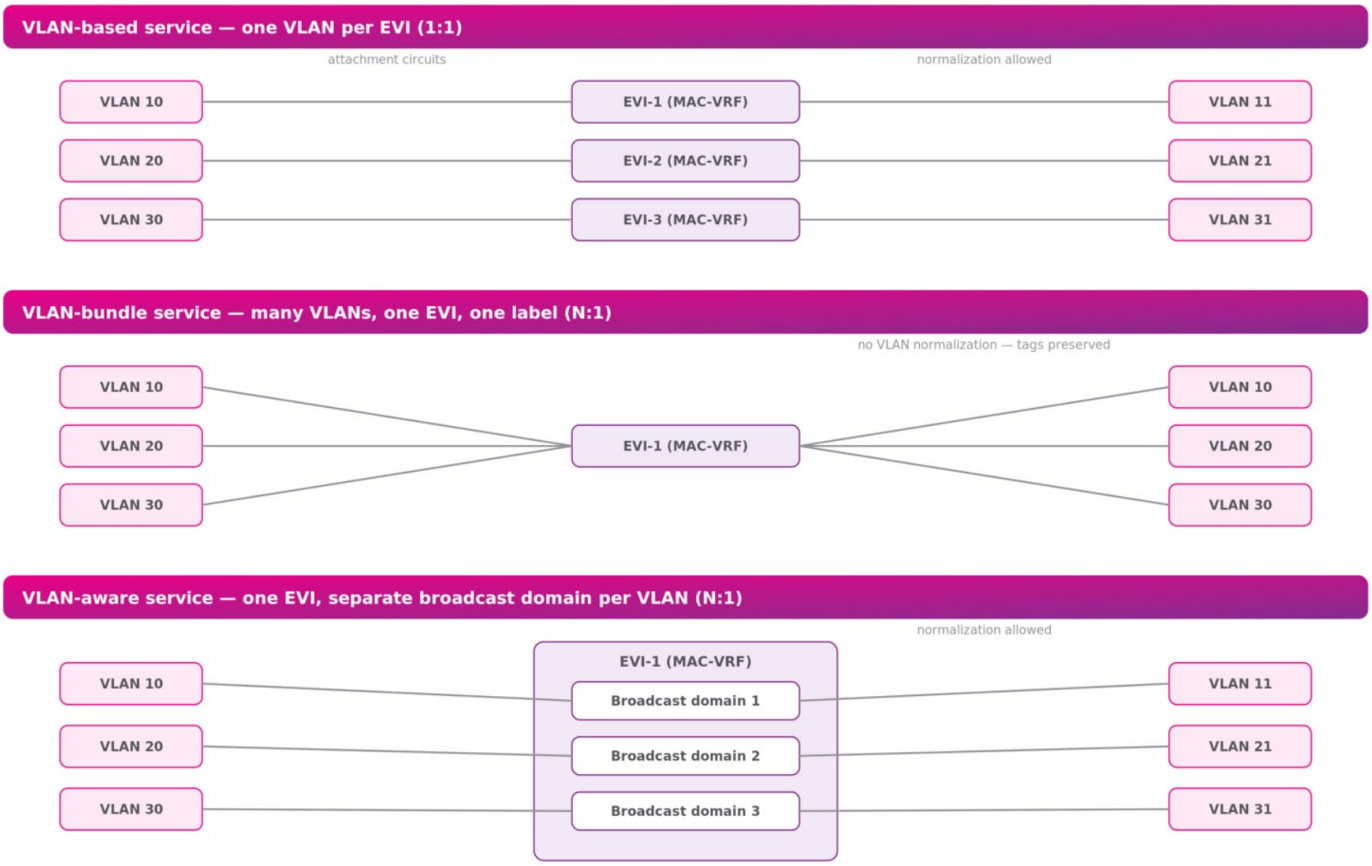

Figure 2 — EVPN VLAN service modes

- **VLAN-based** — one VLAN maps to one EVI. The simplest model: per-VLAN MPLS labels, VLAN normalization on the interface, and optimal per-VLAN BUM distribution, at the cost of more label and instance resources.
- **VLAN-bundle** — many VLANs map to one EVI sharing a single broadcast domain and a single label. The best scaling option, but VLAN tags must be preserved end to end (no normalization) and BUM traffic of one VLAN reaches all VLANs of the bundle.
- **VLAN-aware** — many VLANs map to one EVI, but each VLAN keeps its own broadcast domain inside the instance. Combines N:1 scaling with per-VLAN labels, normalization, and optimal BUM distribution.

Feature	VLAN-based	VLAN-bundle	VLAN-aware
VLAN to EVI mapping	1:1	N:1	N:1
Label allocation	Per VLAN	Per EVI	Per VLAN (per BD)
VLAN normalization	Available	Not available	Available
BUM distribution	Optimal (per VLAN)	Whole bundle	Optimal (per BD)
Resource usage	Highest	Lowest	Moderate

Which Mode for Which Use Case

The three modes are not competitors — each has a natural fit. The table below maps common service scenarios to the recommended mode:

Use case	VLAN-based	VLAN-bundle	VLAN-aware
A few premium services with strict per-service isolation and SLA	Best fit	Not recommended	Suitable
Transparent transport of many customer VLANs (wholesale, C-VLAN pass-through)	Wasteful	Best fit	Suitable
Many VLANs per customer with per-VLAN control and normalization	Wasteful at scale	No normalization	Best fit
Minimal label and instance consumption on the PE	Highest usage	Best fit	Moderate
Per-VLAN QoS and controlled BUM distribution	Best fit	Not available	Best fit
Simple migration of an existing per-VLAN VPLS design	Best fit	Suitable	Suitable

As a rule of thumb: choose VLAN-based for a small number of high-value services, VLAN-bundle when the goal is bulk transparent transport with minimum resources, and VLAN-aware when one customer brings many VLANs that still need individual treatment.

Part II EVPN in Practice: Lab on NPT Platforms

This part demonstrates a complete EVPN deployment on Ribbon NPT hardware: topology and design choices, full device configuration, control-plane and data-plane validation, and measured convergence under three failure scenarios. All outputs are taken from the live lab..

Lab Deployment on NPT Platforms

Topology and Scenario

The lab uses the seven-node NPT test network shown in Figure 3, with Segment Routing MPLS transport and an iBGP full mesh (AS 100) running the L2VPN EVPN address family. A single EVPN service demonstrates all three attachment modes at once:

- **VM-1 (192.168.101.101)** is attached through an access switch that is dual-homed to R1 and R2 — Ethernet segment ESI-NPT-1-2, bundle Be3, active-active mode.
- **VM-2 (192.168.101.102)** is attached through an access switch dual-homed to R6 and R7 — Ethernet segment ESI-NPT-6-7, bundle Be2, port-active mode.
- **VM-5 (192.168.201.105)** is attached to a single port on R3 — single-homed, no Ethernet segment required.

The EVIs on R1, R2, R6, and R7 belong to the same Layer 2 bridge domain (subnet 192.168.101.0/24); each of these routers also hosts an IRB interface with the same anycast gateway address 192.168.101.254. The EVI on R3 is a different bridge domain with its own subnet, 192.168.201.0/24, and gateway 192.168.201.254. Traffic between the two subnets is routed through the IRB interfaces, which are members of a common L3 VRF.

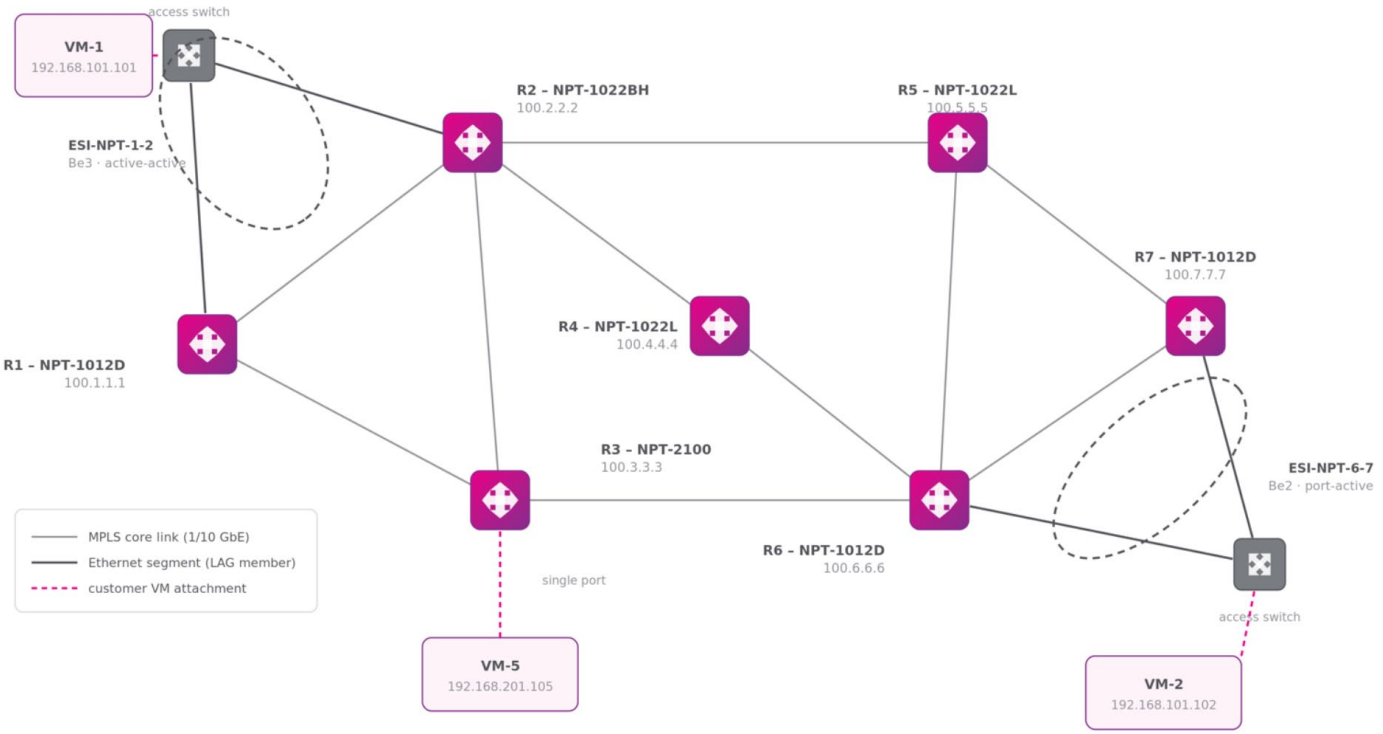

Figure 3 — Lab topology — one EVPN service with active-active, port-active, and single-port attachment

Choosing a Multihoming Mode

The three attachment modes used in the lab were not chosen at random — each answers a different set of requirements, and the same reasoning applies to production designs:

- **Active-active (all-active)** — choose it when the goal is to use every purchased link and get the fastest possible failover. Both PE-CE links forward traffic simultaneously, remote PEs load-balance flows to both PEs (aliasing), and a link failure removes only half of the capacity with no reconvergence of the CE side. It requires a CE or access switch capable of running a single LACP bundle towards two different PEs, and it is the mode that VPLS fundamentally cannot offer — which is why it is demonstrated on R1/R2.
- **Port-active** — choose it when deterministic traffic placement matters more than bandwidth: one port is active, the other is hot standby, and the operator controls which side is active through DF preference (as configured on R6/R7 with preferences 16000/15000). This suits access devices with limited LACP capabilities, asymmetric core capacity, or services where per-flow load balancing between PEs is undesirable — while still delivering about one-second failover (or with new EVPN fast reroute feature even significant faster).
- **Single-active** — choose it when per-service (per-VLAN) load sharing is wanted across the two PEs: each broadcast domain is active on one PE and standby on the other, so different services use different links. It is a natural middle ground for multi-service access rings.
- **Single-homed (single port)** — choose it where redundancy is not justified: lab or test attachments, low-priority sites, or CE devices with a single uplink. No Ethernet segment is configured at all, which keeps the configuration minimal — demonstrated on R3.

Aggregate Interfaces

The multihomed attachment circuits are LACP bundles. The bundle configuration must match on both PEs of the segment and on the access switch. On R1 and R2 the bundle Be1/Be3 is configured as follows:

R1 — aggregate interface configuration

```
[edit]
admin@R1# show interfaces bundle Be1
admin-actor-key 2;
links {
    et-xsa/3 {
        hold-off-timer 0;
        lacp-options {
            mode active;
            time-out short;
        }
        master;
    }
}
min-active-links 1;
description bundle1;
system-mac 00:00:00:12:12:12;
system-priority 1;

[edit]
admin@R1#
```

R2 — aggregate interface configuration

```
[edit]
admin@R2# show interfaces bundle
Be1 {
    admin-actor-key 2;
    links {
        ge-xsa/6 {
            hold-off-timer 0;
            lacp-options {
                mode active;
                time-out short;
            }
            master;
        }
    }
    min-active-links 1;
    description bundle1;
    system-mac 00:00:00:12:12:12;
    system-priority 1;
}

[edit]
admin@R2#
```

Before configuring the service, verify that the bundle members are up and in the distributing state:

R2 — LACP distributing state check

```
[edit]
admin@R2# run show interfaces bundle Be1 | match dist
    Partner Admin State Distributing    :    no
    Partner Oper State Distributing     :    yes
    Actor Oper State Distributing       :    yes
    Actor Admin State Distributing      :    no

[edit]
admin@R1# run show interfaces bundle Be1 | match dist
    Partner Admin State Distributing    :    no
    Partner Oper State Distributing     :    yes
    Actor Oper State Distributing       :    yes
    Actor Admin State Distributing      :    no
```

EVI, L3 VRF, and IRB

Each PE requires an EVI (MAC-VRF) for the Layer 2 service and an L3 VRF that includes the IRB interface for routing between subnets. Note the attachment circuit with VLAN translation and the anycast MAC address shared by all gateways of the bridge domain:

R1 — EVI, L3 VRF, IRB, and attachment circuit

```
[edit]
admin@R1# show routing-instances
vrf evpn-1-2 {
    route-distinguisher 100.1.1.1:122;
    interface irb.12;
    route-target {
        both 64512:1;
    }
}
evi EVI-NPT-12 {
    interface et-xsa/3.100;
    route-distinguisher 100.1.1.1:12;
    route-target {

       both 64512:123;
    }
    routing-interface irb.12;
}

[edit]
admin@R1# show interfaces irb
unit 12 {
    family inet {
        address 192.168.101.254/255.255.255.0;
    }
    anycast-mac 00:00:00:01:02:03;
}

[edit]
admin@R1# show interfaces et-xsa/3
type {
    vlan-tagged;
}
unit 100 {
    family vpls;
    vlan-id-list 100;
    vlan-translation {
        egress-all-to-one {
            outer-vlan {
                vlan-id 100;
            }
        }
        ingress-all-to-one {
            strip;
        }
    }
}
```

R2 — IRB with the same anycast gateway address

```
[edit]
admin@R2# show interfaces irb
unit 12 {
    family inet {
        address 192.168.101.254/255.255.255.0;
    }
    anycast-mac 00:00:00:01:02:03;
}
```

R3 serves the second bridge domain, so its IRB address belongs to the 192.168.201.0/24 subnet:

R3 — EVI, L3 VRF, and IRB for the second subnet

```
[edit]
admin@R3# show routing-instances
vrf SPDC {
    route-distinguisher 100.3.3.3:1;
    interface irb.33;
    route-target {
        both 64512:1;
    }
}
evi EVI-NPT-3 {
    interface et-bs/26.0;
    route-distinguisher 100.3.3.3:64512;
    route-target {
        both 64512:123;
    }
    routing-interface irb.33;
}

[edit]
admin@R3# show interfaces irb
unit 33 {
    family inet {
        address 192.168.201.254/255.255.255.0;
    }
    anycast-mac 00:00:00:01:02:04;
}
```

BGP Configuration

All PEs run iBGP sessions with the L2VPN EVPN address family enabled. MPLS transport must be operational between all devices:

R2 — BGP configuration with the l2vpn evpn family

```
[edit]
admin@R2# show protocols bgp
autonomous-system 100;
enable;
family inet-vpn unicast {
    enable;
}
family l2vpn evpn {
    enable;
}
group ibgp {
    community-type both;
    export-default accept-all;
    family inet-vpn unicast {
        enable;
    }
    import-default accept-all;
    neighbor 100.1.1.1 {
        family l2vpn evpn {
            enable;
        }
        peer-as 100;
    }
    neighbor 100.3.3.3 { ... }
    neighbor 100.4.4.4 { ... }
    neighbor 100.5.5.5 { ... }
    neighbor 100.6.6.6 { ... }
    neighbor 100.7.7.7 { ... }
}
```

Ethernet Segments

Ethernet segment configuration is required on every PE that is part of a multihomed segment. R1 and R2 share segment ESI-NPT-1-2 in active-active mode (mh-type a_a):

R1 and R2 — Ethernet segment, active-active

```
[edit]
admin@R1# show ethernet-segment
ESI-NPT-1-2 {
    esi {
        import-target 12:12:12:12:12:12;
        type Arbitrary;
        value 12:12:12:12:12:12:12:12:12:12;
    }
    mh-type a_a;
    supporting-interface et-xsa/3;
}
[edit]
admin@R2# show ethernet-segment
ESI-NPT-1-2 {
    esi {
        import-target 12:12:12:12:12:12;
        type Arbitrary;
        value 12:12:12:12:12:12:12:12:12:12;
    }
    mh-type a_a;
    supporting-interface ge-xsa/6;
}
```

R6 and R7 share segment ESI-NPT-6-7 in port-active mode (mh-type p_a). Preference-based DF election is configured: R7 has the higher preference (16000) and becomes the active PE:

R6 and R7 — Ethernet segment, port-active with DF preference

```
[edit]
admin@R6# show ethernet-segment
ESI-NPT-6-7 {
    df-elect-method Preference-based {
        preference 15000;
        preemption-enable;
    }
    esi {
        import-target 67:67:67:67:67:67;
        type Arbitrary;
        value 67:67:67:67:67:67:67:67:67:67;
    }
    mh-type p_a;
    supporting-interface et-xsa/3;
}

admin@R7# show ethernet-segment
ESI-NPT-6-7 {
    df-elect-method Preference-based {
        preference 16000;
        preemption-enable;
    }
    esi {
        import-target 67:67:67:67:67:67;
        type Arbitrary;
        value 67:67:67:67:67:67:67:67:67:67;
    }
    mh-type p_a;
    supporting-interface et-xsa/3;
}
```

R3 attaches its CE over a single port, so no Ethernet segment is defined at all:

R3 — no Ethernet segment required for single-homed attachment

```
[edit]
admin@R3# show ethernet-segment


[edit]
admin@R3#
```

Service Validation

Validation starts with the control plane: all BGP sessions must be established, and the peers must negotiate the l2vpn-evpn address family:

R1 — BGP session state

```
[edit]
admin@R1# run show bgp neighbor

Instance: master
Number of Peers: 7
 Peer                         AS       Local                        AS
     State
192.168.133.232                  100      172.16.102.23                    100
     ESTABLISHED
100.2.2.2                     100      100.1.1.1                    100
     ESTABLISHED
100.3.3.3                     100      100.1.1.1                    100
     ESTABLISHED
100.4.4.4                     100      100.1.1.1                    100
     ESTABLISHED
100.5.5.5                     100      100.1.1.1                    100
     ESTABLISHED
100.6.6.6                     100      100.1.1.1                    100
     ESTABLISHED
100.7.7.7                     100      100.1.1.1                    100
     ESTABLISHED
```

R1 — negotiated capabilities for the l2vpn-evpn family

```
admin@R1# run show bgp neighbor 100.6.6.6 detail

Instance: master
Number of Peers: 7

  Peer: 100.6.6.6+179  AS: 100  Local: 100.1.1.1+36498  AS: 100
  Type: internal  State: ESTABLISHED  Uptime: 3h:26m:42s, BGP ID: 100.6.6.6
  Hold Time : 210, RR Client: no
  Neighbor Last Error:
  Supported capabilities:
   Capability       Supported Mode              Famliy List
   GRACEFUL RESTART  yes       BILATERAL          l2vpn-evpn vpn-inet-unicast
   MPBGP            yes                    l2vpn-evpn vpn-inet-unicast
   ROUTE-REFRESH     yes
   ASN32            yes
   ADD-PATH          no
  Add Path Mode:
   ... (output truncated)
```

Next, verify the EVPN routes themselves. The output below shows Type-2 (MAC/IP) advertisements sent by R2. Three kinds of Type-2 routes are visible: a MAC-only route (marked no-ip, prefix length /33), a MAC-with-IP route for the same customer MAC (/40), and the anycast gateway MAC/IP route advertised with a zero ESI:

R2 — Type-2 MAC/IP advertisements towards the EVPN peers

```
[edit]
admin@R2# run show bgp route out family l2vpn-evpn evpn-prefix-specific route-type mac_ip

Instance: master
  54 routes

  Prefix          Family          Origin    Local Pref    Neighbor          Next Hop
  100.2.2.2:12--Arbitrary:12:12:12:12:12:12:12:12:12--0--52:54:00:c1:76:bd--no-ip--524290/33 MAC/IP
                  l2vpn-evpn      incomplete 100          100.1.1.1         100.2.2.2
  100.2.2.2:12--Arbitrary:00:00:00:00:00:00:00:00:00--0--00:00:00:01:02:03--192.168.101.254--524290/40 MAC/IP
                  l2vpn-evpn      incomplete 100          100.1.1.1         100.2.2.2
  100.2.2.2:12--Arbitrary:12:12:12:12:12:12:12:12:12--0--52:54:00:c1:76:bd--192.168.101.102--524290/40 MAC/IP
                  l2vpn-evpn      incomplete 100          100.1.1.1         100.2.2.2
  100.2.2.2:12--Arbitrary:12:12:12:12:12:12:12:12:12--0--52:54:00:c1:76:bd--no-ip--524290/33 MAC/IP
                  l2vpn-evpn      incomplete 100          100.3.3.3         100.2.2.2
  100.2.2.2:12--Arbitrary:00:00:00:00:00:00:00:00:00--0--00:00:00:01:02:03--192.168.101.254--524290/40 MAC/IP
                  l2vpn-evpn      incomplete 100          100.3.3.3         100.2.2.2
   ... (output truncated)
```

With the control plane converged, the VMs in the same subnet reach each other directly through the EVPN bridge domain:

VM-2 — reachability check inside the bridge domain

```
[root@vm2 ~]# ping 192.168.101.101
PING 192.168.101.101 (192.168.101.101) 56(84) bytes of data.
64 bytes from 192.168.101.101: icmp_seq=1 ttl=64 time=0.157 ms
64 bytes from 192.168.101.101: icmp_seq=2 ttl=64 time=0.162 ms
^C
--- 192.168.101.101 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.157/0.159/0.162/0.012 ms
[root@vm2 ~]# ping 192.168.101.102
PING 192.168.101.102 (192.168.101.102) 56(84) bytes of data.
64 bytes from 192.168.101.102: icmp_seq=1 ttl=64 time=0.016 ms
64 bytes from 192.168.101.102: icmp_seq=2 ttl=64 time=0.023 ms
^C
--- 192.168.101.102 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.016/0.019/0.023/0.005 ms
```

The DF election result — an outcome of the Type-4 routes exchanged between R6 and R7 — confirms that R7 is the Active-DF for the port-active segment thanks to its higher preference. The ESI label shown in the same output is carried in the Type-1 route and used for split-horizon filtering:

R6 and R7 — DF election result and ESI label

```
admin@R7> show ethernet-segment detail
ethernet-segment: ESI-NPT-6-7
   Supporting I/F: et-xsa/3, Oper: up
   ESI Type: Arbitrary, ESI Value: 67:67:67:67:67:67:67:67:67:67, ESI label: 524306
   DF Election Method: Preference-based, Actual DF Election Method: Preference-based
   DF Preference, Configured: 16000, Operational: 16000
   DF Election Result: Active-DF
   DF Preemption Mode, Configured: true, Operational: true
   MH Type: Port-Active, Actual MH Type:: Port-Active
   Actual ES Import Target: 67:67:67:67:67:67
   MH ES members: 100.6.6.6, *100.7.7.7
   RT List: 1

admin@R6> show ethernet-segment detail
ethernet-segment: ESI-NPT-6-7
   Supporting I/F: et-xsa/3, Oper: up
   ESI Type: Arbitrary, ESI Value: 67:67:67:67:67:67:67:67:67:67, ESI label: 524300
   DF Election Method: Preference-based, Actual DF Election Method: Preference-based
   DF Preference, Configured: 15000, Operational: 15000
   DF Election Result: Backup-DF
   DF Preemption Mode, Configured: true, Operational: true
   MH Type: Port-Active, Actual MH Type:: Port-Active
   Actual ES Import Target: 67:67:67:67:67:67
   MH ES members: *100.6.6.6, 100.7.7.7
   RT List: 1
```

The MAC forwarding table shows the data-plane result of the Type-2 advertisements: the local customer MAC points to the attachment circuit, while remote MACs point to Segment Routing LSPs with the labels received from BGP:

R1 and R2 — EVI MAC forwarding table

```
[edit]
admin@R1# run show routing-instances evi forwarding-table
- = Frozen MAC
Instance: EVI-NPT-12
  Destination        Type    Aging(S)   NextHop                                  Label
     Protection
  00:00:00:01:02:04  static   N/A        100.3.3.3\lsp-sr:algo-0:NA                524297
     Primary
  52:54:00:7f:fc:a6  static   N/A        100.7.7.7\lsp-sr:algo-0:NA                524303
     Primary
  52:54:00:c1:76:bd  static   300        et-xsa/3.100                              N/A
     Primary
  Total: 3 entries

[edit]
admin@R2# run show routing-instances evi forwarding-table
- = Frozen MAC
Instance: EVI-NPT-12
```

```
Destination        Type     Aging(S)    NextHop                                  Label
   Protection
00:00:00:01:02:04  static   N/A         100.3.3.3\lsp-sr:algo-0:NA                524297
   Primary
52:54:00:7f:fc:a6  static   N/A         100.7.7.7\lsp-sr:algo-0:NA                524303
   Primary
52:54:00:c1:76:bd  static   300         ge-xsa/6.100                              N/A
   Primary
Total: 3 entries
```

Finally, three MPLS labels identify the service traffic, each corresponding to a different EVPN mechanism: the IR (inclusive replication) label from the Type-3 route carries BUM traffic, the EVI VPN label carries known unicast, and the VRF EVPN label is used for routed traffic between the subnets:

R1 — EVI and VRF label allocation

```
[edit]
admin@R1# run show routing-instances evi detail
Instance-type : MACVRF
  Instance: EVI-NPT-12, Oper State: Up
    Route Distinguisher Configured: manual, Actual: 100.1.1.1:12
    Label Allocation Method: PER-MACVRF
    VPN Label: 524289, IR Label 319488
    Interfaces:
      Interface       Admin Oper  IsDF
      et-xsa/3.100    Up    Up    Yes
    IRBs:
      IRB       Admin    Oper    Anycast MAC Address
      irb.12    Up       Up      Explicitly configured

[edit]
admin@R1# run show routing-instances vrf
Instance-type : VRF
  Instance: evpn-1-2,  Oper State: up,  Is Management:None
    Route Distinguisher Configured: manual, Actual: 100.1.1.1:122
    Max Routes: 0, Threshold: 90, Current Routes: 0
    Label Allocation Method: PER-VRF
    VPN Label: 524288, EVPN Label 524290
   Interfaces:
     Interface                    Admin     Oper
     irb.12                       up        up
Instance-type : VRF
  Instance: master
```

Failure Emulation

Three failure scenarios were emulated to measure service convergence (Figure 4):

- **Case 1** — with traffic running from VM-2 to VM-1, the active link between the access switch and R2 fails. This tests convergence of the active-active segment: R2 withdraws its Type-1 route (mass withdrawal), and remote PEs immediately redirect all MAC addresses to R1.
- **Case 2** — with traffic running from VM-1 to VM-2, the active link between the access switch and R7 fails. This tests the port-active segment: the DF role fails over to R6 based on the Type-4 election.
- **Case 3** — a core link on the active traffic path fails while Ti-LFA (Topology-Independent Loop-Free Alternate) protection is enabled on the Segment Routing transport.

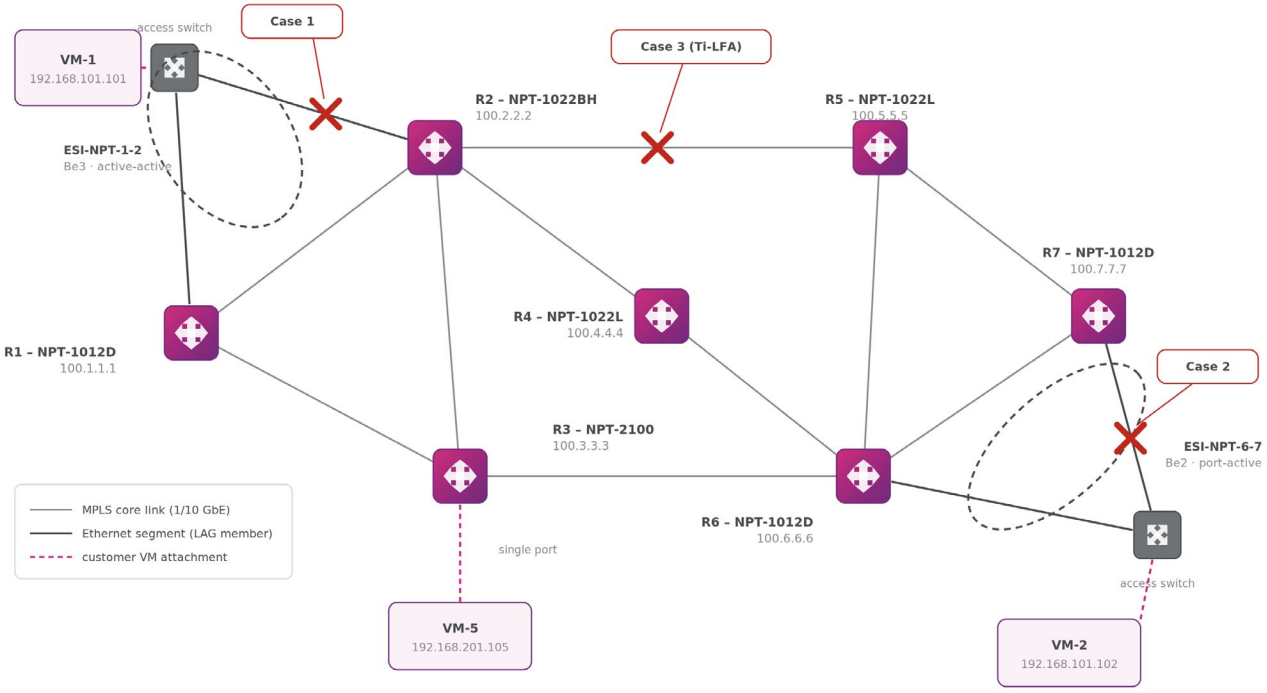

Figure 4 — Failure emulation — access link failures (Cases 1 and 2) and core link failure with Ti-LFA (Case 3)

Across repeated attempts, both access-side failures (Cases 1 and 2) converged in close to one second, driven by the EVPN mass-withdrawal and DF re-election mechanisms. The core link failure protected by Ti-LFA (Case 3) restored traffic in under 50 milliseconds, since the repair happens locally in the forwarding plane without waiting for control-plane reconvergence.

Below the table with measured results from different cases:

Use case	Min	Avg	Max
Case 1	980 ms	1350 ms	2250 ms
Case 2	966 ms	1440 ms	2330 ms
Case 3	0 ms	40 ms	75 ms

* With EVPN fast-reroute feature enabled results for cases 1 and 2 must be significantly improved.

EVPN and VPLS: Feature Comparison

The lab results make the architectural argument concrete. The table below summarizes where VPLS is limited by design and how EVPN addresses each limitation:

Feature	VPLS	EVPN
MAC learning	Data plane, by flooding frames across all pseudowires	BGP control plane (Type-2); flooding only for first contact
All-active multihoming	Not possible — one link active, others idle	Native (Type-1/Type-4); all links forward, flows balanced via aliasing
Standby link utilization	Paid-for links carry no traffic	Active-active uses every link; port-active gives controlled standby
Convergence after access failure	Slow — MAC tables flushed and re-learned by flooding	Mass withdrawal (Type-1): one BGP update repoints all MACs, ~1 s
Core protection	Depends on transport	Segment Routing with Ti-LFA — sub-50 ms local repair (lab-verified)
BUM traffic control	Flooding grows with network size	Delivered only to PEs advertising Type-3; multicast optimized by Types 6–8
Loop prevention (multihoming)	Manual designs, spanning tree at the edge	Built-in split horizon with ESI labels (Type-1)
Auto-discovery	None in LDP-VPLS; partial in BGP-VPLS	Built into the BGP EVPN address family
Integrated L2 + L3	Limited, external constructs required	IRB with anycast gateways (Type-2 + RFC 9135), IP prefixes (Type-5)
MAC mobility	No mechanism — stale entries until aging	Sequence-numbered Type-2 updates move hosts instantly
Provisioning a new PE	Full mesh of pseudowires must be extended	Configure the EVI — BGP distributes everything else

Every row of this table was exercised in the lab in Part II: control-plane learning was observed in the Type-2 advertisements, all-active and port-active multihoming were configured side by side, split horizon and DF election were verified in the show outputs, and the convergence figures were measured under real simulated failures. VPLS remains a functional technology, but each of its structural limitations requires an operational workaround; EVPN removes the limitation itself.

Conclusion

EVPN has moved from a promising standard to the default choice for Carrier Ethernet: it delivers all-active multihoming, control-plane MAC learning, fast and predictable convergence, and integrated Layer 2 and Layer 3 services within a single BGP-based framework. The technology also keeps evolving — route types 6 through 8 (RFC 9251) bring IGMP/MLD awareness into the control plane so multicast reaches only interested receivers, and route types 9 through 11, currently in standardization, further refine BUM procedures.

The lab in this document demonstrates that these benefits are available on Ribbon NPT platforms today, not in a roadmap: a single EVPN service combined active-active, port-active, and single-homed attachments with anycast IRB gateways over Segment Routing transport, converging in about one second on access failures and under 50 milliseconds on Ti-LFA-protected core failures — figures that legacy VPLS designs cannot reach.

Ribbon NPT provides a standards-based EVPN implementation that enables operators to modernize their Ethernet services, simplify day-to-day operations, and confidently migrate from legacy VPLS architectures — preserving the service model their customers rely on while removing the limitations that constrained it.

Contact Us Contact us to learn more about Ribbon solutions.

About Ribbon

Ribbon Communications (Nasdaq: RBBN) is a global provider of voice communications software, IP routing, and optical networking to mobile and wireline service providers, enterprises, critical infrastructure and defense sectors. We support our customers' Path to Autonomous Networks by leveraging the latest AIOps automation platforms and Agentic AI technologies, helping them deliver better customer experiences, reduce operational costs, and achieve sustainable growth. To learn more about Ribbon visit rbbn.com.