BESS Workgroup
Internet-Draft
Intended status: Standards Track
Expires: 15 November 2025

J. Rabadan, Ed.
S. Sathappan
V. Prabhu
Nokia
W. Lin
Juniper
P. Brissette
Cisco Systems
14 May 2025

Ethernet VPN Virtual Private Wire Services Gateway Solution
draft-ietf-bess-evpn-vpws-gateway-00

Abstract

Ethernet Virtual Private Network Virtual Private Wire Services (EVPN VPWS) need to be deployed in high scale multi-domain networks, where each domain can use a different transport technology, such as MPLS, VXLAN or Segment Routing with MPLS or IPv6 Segment Identifiers (SIDs). While transport interworking solutions on border routers spare the border routers from having to process service routes, they do not always meet the multi-homing, redundancy, and operational requirements, or provide the isolation that each domain requires. This document analyzes the scenarios in which an interconnect solution for EVPN VPWS using EVPN Domain Gateways is needed, and adds the required extensions to support it.

Table of Contents

1. Introduction

Ethernet VPN Virtual Private Wire Services (EVPN VPWS) [RFC8214] need to be deployed in high scale multi-domain networks, where each domain can use a different transport technology, such as MPLS, VXLAN or Segment Routing with MPLS or IPv6 Segment Identifiers (SIDs). While the so-call transport interworking solutions on border routers spare the border routers from having to process service routes, they do not always meet the multi-homing, redundancy, and operational requirements, or provide the isolation that each domain requires. This document analyzes the scenarios in which an interconnect solution for EVPN VPWS using EVPN Domain Gateways is needed, and adds

the required extensions to support it.

1.1. Terminology

This section summarizes the terminology that is used throughout the rest of the document.

* BR: Border Router, router that provides connectivity between domains, typically an Area Border Router (ABR) or Autonomous System Border Router (ASBR).

* BUM: Broadcast, Unknown unicast and Multicast traffic.

* Domain: in this document Domain and EVPN Domain are used interchangeably.

* E-PE: Egress Provider Edge router.

* ES and ESI: Ethernet Segment and Ethernet Segment Identifier, as defined in [I-D.ietf-bess-rfc7432bis].

* EVPN Domain and EVPN Domain Gateway: two PEs are in the same EVPN Domain if they are attached to the same service and the packets between them do not require a data path lookup of the inner frame in any intermediate router. An EVPN Domain is typically a group of PE, P and Border Routers that belong to the same IGP instance or BGP domain. EVPN services are instantiated on the PEs and Border Routers, which are referred to as EVPN Domain Gateways in this document. An EVPN Domain Gateway connects two or more EVPN Domains and is configured with multiple Domain identifiers (EVPN Domain-ID) in the VPWS that connects those EVPN Domains. Each EVPN Domain-ID representing an EVPN Domain. Another definition of EVPN Domain Gateway is a Border Router that implements the Service Interworking procedures described in this document.

* I-ES and I-ESI: Interconnect Ethernet Segment and Interconnect Ethernet Segment Identifier. An I-ES is defined for multihoming to the domains to which a Service Gateway is attached [RFC9014].

* I-PE: Ingress Provider Edge router.

* NVO: Network Virtualization Over Layer-3 tunnels.

1.2. EVPN Interconnect Options

This section describes the EVPN [I-D.ietf-bess-rfc7432bis] high level interconnect options and discusses their applicability to EVPN VPWS.

1. Service Interworking solution:

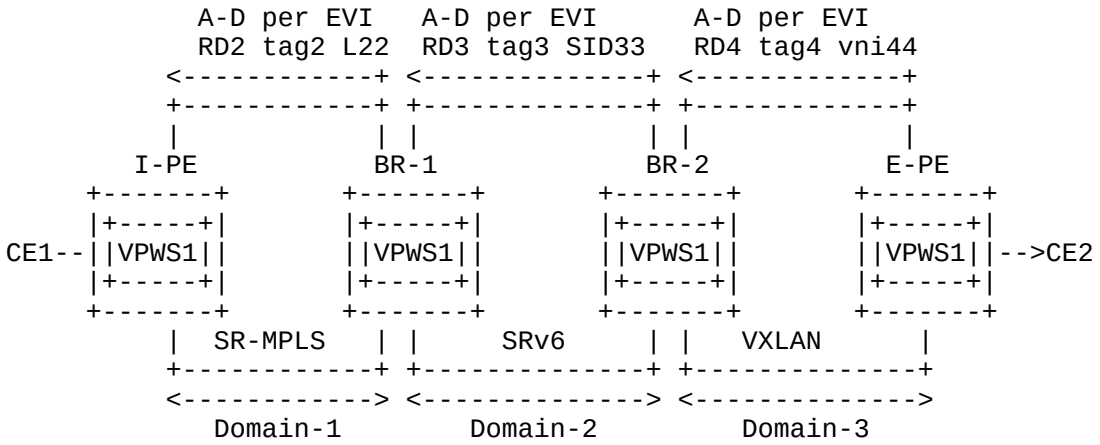

Figure 1: Service Interworking Interconnect

[RFC9014] section 4 describes an end-to-end EVPN interconnect solution using EVPN Domain Gateways, or simply Gateways. The Gateways provide connectivity across EVPN Domains, where those Domains can use MPLS tunnels, NVO3 tunnels (e.g., VXLAN) or Segment Routing tunnels. Procedures are extrapolated to SRv6 domains too. The Gateways provide independence in terms of the Route Targets and Route Distinguishers used in each Domain, or the type of multicast tree used for BUM traffic in each domain, while keeping the key EVPN properties end-to-end, such as MAC mobility, MAC protection or ARP suppression. The Gateways also provide all-active and single-active multi-homing redundancy by extending the concept of the multi-homing Ethernet Segment for interconnect domains (I-ES). In this document, we refer to this solution as the Service Interworking option, and the Border Routers play the role of EVPN Domain Gateways. Since [RFC9014] section 4 only describes the solution for EVPN multi-point services, this document extends the procedures to support EVPN VPWS services with the required extensions. Figure 1 illustrates the Service Interworking solution across domains of different transport encapsulations when applied to EVPN VPWS services.

2. Inter-domain Option-B solution:

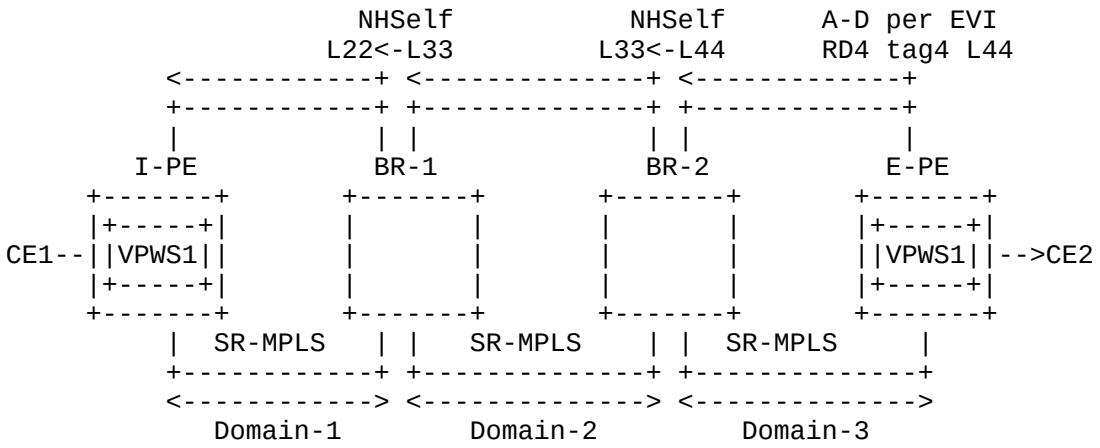

Figure 2: Inter-domain Option-B

[RFC8365] section 10 provides an alternative interconnect solution for EVPN services by using Border Routers that re-write the EVPN BGP next hops and program a swap operation of the VNIs or MPLS labels (depending on whether the encapsulation is NVO3-based or MPLS-based). This solution does not require the instantiation of Services on the Border Routers that perform a lookup on the inner destination MAC (as it is the case in [RFC9014]), however the solution is limited to the interconnect of domains of the same encapsulation. In addition, the solution does not support per-ES mass withdraw of the EVPN MAC/IP Advertisement routes, as described in [RFC8365]. In this document we refer to this solution as Inter-domain Option-B. Figure 2 illustrates this model applied to EVPN VPWS, where all three domains use the same encapsulation, and no service instantiation occurs on the Border Routers.

3. Transport Interworking solution:

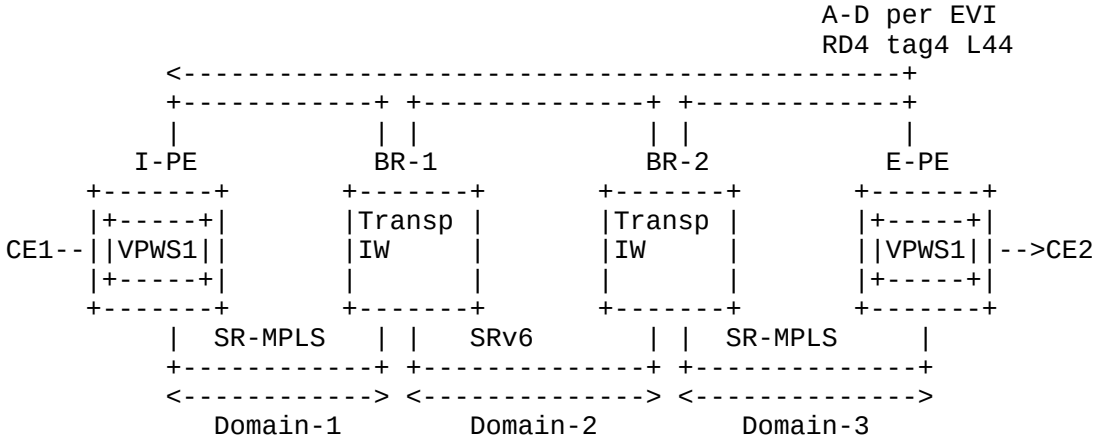

Figure 3: Transport Interworking option

Other proposals are currently being investigated, in the context of SRv6 to MPLS interworking, e.g., [I-D.ietf-spring-srv6-mpls-interworking]. In these solutions, the Border Routers do not change the EVPN BGP next hops, or process EVPN routes for that matter. The Border Routers provide stitching between MPLS and SRv6 tunnels. In this case, the solution allows the interconnect of domains of different encapsulation, as long as the ingress and egress PEs support the same encapsulation. A variation of this solution is the Inter-domain Option-C solution, where a BGP LU (Label Unicast) tunnel provides the stitching across the domains, as long as all the domains use the same encapsulation. In this document, we refer to this solution as Transport Interworking option. Figure 3 illustrates this model when applied to EVPN VPWS, where I-PE and E-PE are attached to domains of the same encapsulation. Intermediate domains - such as Domain-2 - may use encapsulations different from those in the ingress and egress domains. However, the EVPN route remains unchanged and is not processed by the Border Routers.

1.3. When is the Service Interworking Solution Required for EVPN VPWS

The three interconnect solutions described in Section 1.2 are valid, however, this section describes the requirements that make the Service Interworking solution needed. Those requirements are:

a. Per-domain EVPN Multi-Homing

The Service Interworking solution allows the use of different Ethernet Segment Identifiers (ESI) per domain, as well as the implementation of the aliasing and backup procedures on a per-domain basis. The use of different ESIs per domain may help guarantee the uniqueness of the ESI when different domains independently managed and operated are interconnected. The implementation of independent aliasing and backup procedures per domain, spares the need for propagation of the EVPN A-D per ES routes by the Border Routers (which are EVPN Domain Gateways in the Service Interworking solution). These A-D per ES routes are consumed within the domain, which results in a significant reduction of the number of routes that the ingress PEs need to process. Another consequence of the processing of A-D per ES routes per domain, is a faster convergence in case of ES PE or link failure, since A-D per ES routes are no longer propagated by all the Border Routers along the domains, but processed by the Border Routers of the originating domain. Per-domain EVPN Multi-Homing procedures are not possible in the Inter-domain Option-B or Transport Interworking solutions.

b. Per-ES Mass Withdrawal

In order to benefit from the per-ES mass withdrawal property of EVPN Multi-Homing, the received BGP next hops of the selected EVPN A-D per EVI and A-D per ES routes need to match on a PE. This cannot be guaranteed in an Inter-domain Option-B solution, as described in [RFC8365] section 10.2.2. However, it is always ensured in both the Service Interworking and Transport Interworking solutions.

c. Per-domain Route Distinguishers (RDs) and Route Targets (RTs)

In case of merge of domains coming from different administrative entities, the uniqueness of RDs and RTs across domains for the same service is not guaranteed. Hence the re-write of RD/RTs at the Border Routers may be required. If that is the case, the Service Interworking solution provides the support for re-writing RD/RTs. The Inter-domain Option-B may allow re-writing RD/RTs, however, it is not considered a common practice. The Transport Interworking solution does not support the translation of RD/RTs.

d. Ethernet Tag IDs per domain

Similar to per-domain RDs and RTs, re-writing of Ethernet Tag IDs used in the A-D per EVI routes may be needed in case of interconnecting domains that belong to different administrative entities. This can be only supported by a Service Interworking solution.

e. Control Word, Flow Label and MTU (Maximum Transfer Unit) signaling per domain

   As described in [I-D.ietf-bess-rfc7432bis], the use of Control Word and Flow Label, as well as the MTU are signaled in the EVPN Layer 2 Attributes extended community along with the A-D per EVI routes. The signaling and use of Control Word is recommended in those domains where P routers can get confused when hashing based on the tunneled EVPN packet payload, but the Control Word may not be needed in some domains. Similarly, the Flow Label introduces an additional level of entropy in EVPN encapsulated packets, that may be needed in some domains but adding unnecessary extra overhead in other domains. Different MTUs may be supported in different domains, due to the domains running on different physical media. A Service Interworking model allows the signaling and use of Control Word, Flow Label, and Layer-2 MTU on a per domain basis. This is not the case in the other two models analyzed in this document.

f. Heterogeneous Encapsulations

   Interconnecting domains that use different encapsulations (e.g., VXLAN, SRv6, MPLS, SR-MPLS, etc.) is a common requirement. This becomes important in case the domains have different platform features, or migrations to new encapsulations or transport types are needed. In the Service Interworking model the EVPN routes are generated and consumed at every Border Router (which is an EVPN Domain Gateway), hence the encapsulation indicated along with the route can be advertised independently at each Border Router. That is not the case in the models 2 and 3 in Section 1.2. The Inter-domain Option-B model requires the same encapsulation in each of the domains the Border Router connects, whereas the Transport Interworking model requires that at least the ingress and egress domains have the same encapsulation.

g. Per-domain EVPN Service OAM

   [RFC9062] defines the Service OAM requirements for EVPN services. When applied to the Interconnect solutions, the three solutions in Section 1.2 allow for the use of MEPs and MIPs on the ingress and egress PEs, but only the Service Interworking solution supports MEPs and MIPs on the Border Routers. In other words, per-domain EVPN Service OAM is only supported in the Service Interworking option.

The above requirements and their support across the Interconnect solutions are summarized in Table 1.

Requirement	Service Interworking	Inter-domain Option-B	Transport Interworking
Per-domain EVPN Multi-Homing	Yes	No	No
Per-ES Mass Withdrawal	Yes	No	Yes
Per-domain RD/ RTs	Yes	Yes*	No
Ethernet Tag IDs per domain	Yes	No	No
Control Word, Flow Label and MTU signaling per domain	Yes	No	No
Heterogeneous encapsulations	Yes	No	Yes**
Per-domain EVPN Service OAM	Yes	No	No

Table 1: EVPN VPWS Interconnect Options Comparison

* Although possible, it is unusual to re-write RD/RTs in the Inter-domain Option-B solution

** Supported only when the ingress and egress domains are of the same encapsulation

1.4. Service Gateway Extensions for EVPN VPWS

The rest of the document specifies the extensions required for the EVPN Domain Gateways to implement the Service Interworking solution to deploy end-to-end EVPN VPWS services. In a nutshell, the AD per EVI routes advertised by the E-PE are redistributed across domains and delivered to the I-PE, while ES and A-D per ES routes advertised by E-PEs are not redistributed by the EVPN Domain Gateways. In addition, this document defines how Gateway redundancy works using either an Anycast Gateway solution, or by extending the I-ES concept

already defined for multi-point EVPN services in [RFC9014].

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Service Interworking procedures for EVPN VPWS

This section describes the EVPN VPWS extensions on the EVPN Domain Gateways (or simply Gateways) to support the Service Interworking model. An EVPN Domain Gateway in this context is a Border Router that connects EVPN Domains and implements the Service Interworking model of Section 1.2. Section 3.1 specifies the Gateway rules to redistribute EVPN routes. When redundant Gateways attached to two or more EVPN Domains are deployed, there are two redundancy mechanisms that can be used. Section 3.2 describes a redundancy method that we refer to as "Anycast" and is based on the redundant Gateways behaving as a single system for the remote PEs. Section 3.3 describes the redundancy based on I-ES, as an extension of the I-ES procedures specified in [RFC9014], only for EVPN VPWS services. The Anycast redundancy does not require the use of I-ES and supports single-active multi-homing connectivity, but it will not support all-active, aliasing, backup, or mass withdraw features that are supported along with the use of I-ES and EVPN Multi-Homing.

3.1. Redistribution of EVPN Routes Across Domains

The EVPN Domain Gateways MUST establish separate BGP sessions for sending/receiving EVPN routes to/from each different Domain to which they are attached. We refer to redistribution of an EVPN route as the set of procedures on the Gateway that include receiving and processing the EVPN route from the source domain, programming the corresponding forwarding path, and re-advertising the route to a different domain (the next destination domain).

The reception and processing of EVPN routes for an EVPN VPWS service follows [RFC8214]. If the D-PATH attribute is contained in the EVPN A-D per EVI route, loop detection and best path selection follows [I-D.ietf-bess-evpn-dpath]. The Gateway imports the valid best EVPN A-D per EVI route required for an Ethernet Tag ID based on the matching import Route Target and the best path selection described in [I-D.ietf-bess-rfc7432bis], section 7.13.2. If a non-zero ESI is included in the route, the [RFC8214] procedures for aliasing, backup, and mass withdraw are followed on the Gateway. Note that the best

path selection for A-D per EVI routes (with non-zero ESI) in [I-D.ietf-bess-rfc7432bis] section 7.13.2 also influences how the Gateway adds primary or back-up next hops to the created Ethernet Segment destinations. As an example, suppose a Gateway receives "m" A-D per EVI routes for ESI "x" and Ethernet Tag ID "y" (all of them with different Route Distinguishers and the flag P set) but supports only "n" paths in the Aliasing list for ESI "x" (with m>n). In this case, the Gateway orders the "m" routes following the best path selection in [I-D.ietf-bess-rfc7432bis] section 7.13.2, and selects the "n" top routes of the ordered list.

If an A-D per EVI route for a service is successfully imported and processed, forwarding state is programmed in the data path using the MPLS label, VNI or SRv6 SID that was received in the EVPN A-D per EVI route. In addition, depending on the encapsulation of the route's next destination domain, the router allocates a new MPLS label, VNI or SRv6 SID and programs a data path switching operation between the identifiers of the source and next destination domains. Immediately after, the Gateway re-advertises the route to the BGP speaker in the next domain. The source domain refers to the domain from which the Gateway receives the route, while the next domain is the EVPN domain where the Gateway redistributes the route. The following considerations apply to the redistributed EVPN A-D per EVI routes:

a. The redistributed A-D per EVI route MUST carry a different RD than the source A-D per EVI route did. This ensures that, in case of redundant Gateways, there is full path visibility in the next domain where the route is advertised.

b. The redistributed route MAY carry the same set of Route Targets as the source route did, if the source and next destination domains use different encapsulations, however translation or re-write of Route Targets SHOULD be supported in this case. In case the source and next destination domains use the same encapsulation, the Gateway MUST use either different import Route Targets in the two domains, or use different Ethernet Tag IDs to create forwarding state in the two domains. This ensures the Gateway does not loop packets back to the source domain and the redistributed routes are not leaked back to the source domain.

c. The ESI of the redistributed route MUST be set to zero or the value of the I-ESI defined in the Gateway (if any).

d. The Ethernet Tag ID of the redistributed route MAY have the same value as the source route. Translation of the Ethernet Tag IDs SHOULD be supported though.

e. The EVPN Layer 2 Attributes extended community is regenerated for the redistributed route. The value of the P and B flags are set based on the Gateway's I-ES and MUST NOT be propagated from the source route. The Control Word, Flow Label flags, as well as the MTU, MAY be set to different values from the source A-D route. The M and V flags [RFC9744] of the redistributed route MUST be copied from the M and V flag values of the source route.

f. The encapsulation specific attributes of the redistributed route are regenerated based on the encapsulation of the next domain. That includes the encoding of the A-D per EVI route NLRI as specified in [RFC8214] or [RFC8365], or the addition of the SRv6 Services TLV as in [RFC9252].

g. The redistributed route SHOULD carry the Communities, Extended Communities, Large Communities and Wide Communities of the source route.

   * The source route in this context is the best A-D per EVI route for the Ethernet Tag ID, as per the best path selection in [I-D.ietf-bess-rfc7432bis] section 7.13.2, irrespective of the ESI being zero or non-zero.

   * Exceptions to the propagation rule are Route Targets (which are reoriginated), EVPN Extended Communities and BGP Encapsulation Extended Communities [RFC9012]. EVPN Extended Communities and BGP Encapsulation Extended Communities MUST NOT be propagated across domains.

h. The redistributed A-D per EVI route MUST update the D-PATH attribute of the received route, or add the D-PATH attribute if the received route did not contain a D-PATH [I-D.ietf-bess-evpn-dpath].

EVPN VPWS services also make use of multi-homing routes, that is, EVPN A-D per ES routes and Ethernet Segment routes. These multi-homing routes are processed in the Gateway as in [RFC8214]. The A-D per ES and Ethernet Segment routes are only processed in the context of the domain they are received, and they MUST NOT be redistributed to any other domain. A-D per ES and Ethernet Segment routes may be originated at the Gateway though, if the Gateway is attached to an I-ES, as described in Section 3.3.

The procedures on the EVPN Domain Gateways described in this document are compatible with PEs that implement either the default Flexible Crossconnect (FXC) mode or the VLAN-Signaled Flexible Crossconnect mode described in [RFC9744].

3.2. EVPN Domain Anycast Gateways for redundancy

The Anycast Service Gateway redundancy is specified as follows:

a. All the Anycast Gateways attached to the same two domains MUST redistribute the EVPN A-D per EVI routes between domains as per Section 3.1 with the following considerations:

   * No I-ES is used on the Gateways, therefore the ESI value MUST be set to zero when redistributing EVPN A-D per EVI routes.

   * All the redundant Gateways can set the same (or different) Ethernet Tag ID in the redistributed A-D per EVI route.

b. All Anycast Gateways MUST process the received D-PATH attribute and update the D-PATH (with the source domain-id) when redistributing the A-D per EVI route to the next domain. The D-PATH attribute will avoid control plane loops.

As an illustration of this redundancy method, suppose all four Service Gateways in Figure 4 are configured as Anycast Service Gateways, and local and remote Ethernet Tag IDs are configured as 1, 2 and 3 on all routers in the domains 1, 2 and 3 respectively.

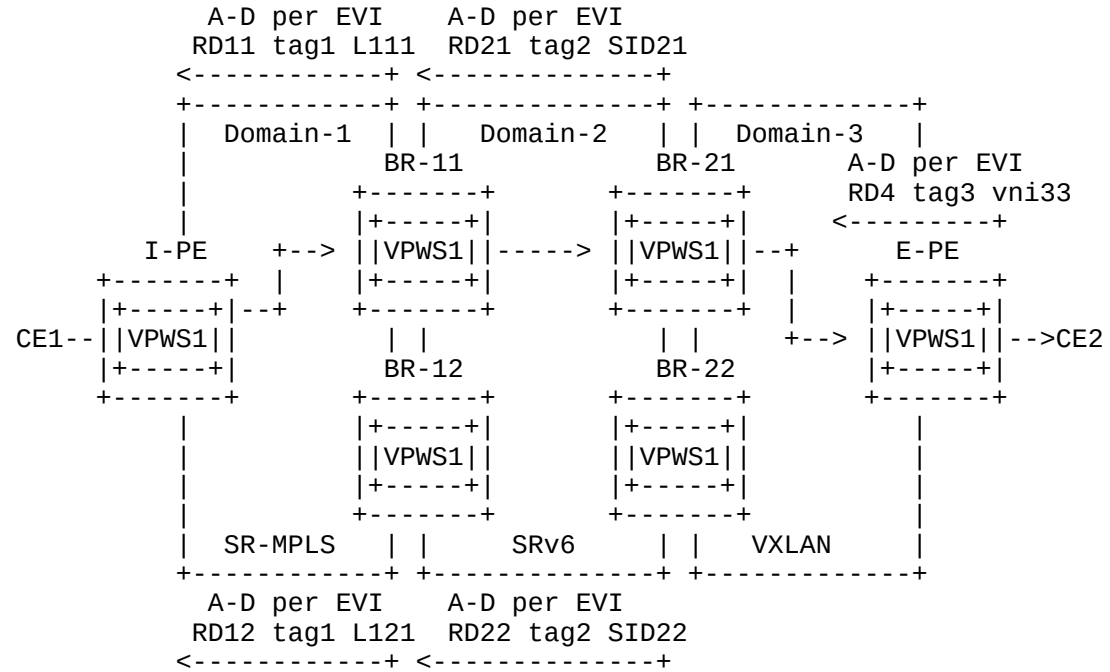

Figure 4: Anycast Redundancy

In the example in Figure 4 E-PE advertises an EVPN A-D per EVI route for Ethernet Tag ID 3. Both BR-21 and BR-22 import the route and redistribute it with Ethernet Tag ID 2 and new RD and encapsulation into domain-2. When redistributing, both BR-21 and BR-22 update (if it existed before) or insert a D-PATH attribute with the domain-id of domain-3. That prevents BR-21 and BR-22 from redistributing back into domain-3 each other's route [I-D.ietf-bess-evpn-dpath]. BR-11 and BR-12 import the routes after best path selection and perform the same process and redistribution into domain-1. I-PE will receive two routes for Ethernet Tag ID 1, from BR-11 and BR-12, and will perform best path selection for Ethernet Tag ID 1. Based on the best path selection carried out by I-PE and the BRs along the way, all flows from CE1 to CE2 will follow, e.g., I-PE, BR-11, BR-21 and E-PE. In case of failure on any of the BRs in the data path, the routers will select the alternate route for the Ethernet Tag ID. The same control plane exchange and traffic flow happen in the reverse direction, where I-PE becomes the egress PE and E-PE the ingress PE.

As illustrated in Figure 4, this model does not support per-flow load balancing (all-active multi-homing) to all the BR nodes along the way from CE to CE.

3.3. EVPN Multi-Homing for Domain Gateway Redundancy (I-ES)

EVPN Multi-Homing procedures can be used on the EVPN Domain Gateways. For that, an I-ES and its assigned I-ESI will be configured on the Gateways for multihoming. The I-ES concept is introduced in [RFC9014], and it is used in this document for EVPN VPWS services. This I-ES represents a domain to the next domain, in both directions. Therefore two or more Gateways attached to the same two domains will use the same I-ESI when advertising routes to the two domains.

The Gateways attached to the same I-ES:

a. Advertise EVPN Ethernet Segment routes and A-D per ES routes for the I-ES. Those routes are not redistributed beyond the Domain into which they are originated.

b. Receive Ethernet Segment and A-D per ES routes from the I-ES peer(s), and use them for I-ES Designated Forwarding (DF) Election and mass withdraw respectively, as described in [RFC8214] and [I-D.ietf-bess-rfc7432bis].

c. Set the I-ESI into the EVPN A-D per EVI routes that are redistributed across domains. P and B flags are set based on the result of the DF Election [RFC8214].

d. Identify loops if the received EVPN A-D per EVI routes include a local domain-id in the D-PATH attribute. Also EVPN A-D per EVI routes that include a local ESI MUST NOT be redistributed to another domain, irrespective of the presence of the D-PATH attribute.

Figure 5 illustrates the use of I-ES or EVPN Multi-Homing procedures in EVPN Domain Gateways. In the example, BR-11 and BR-12 are attached to I-ES-1 (with ESI-1 as identifier), whereas BR-21 and BR-22 are attached to I-ES-2 (using ESI-2).

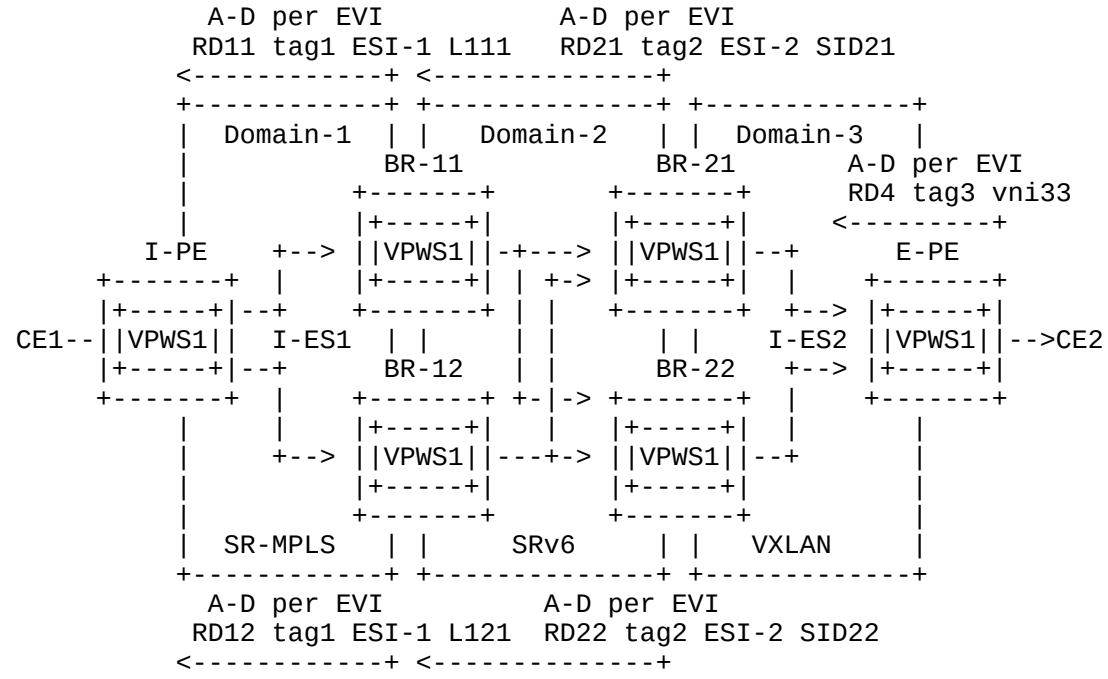

Figure 5: EVPN Multi-Homing

E-PE advertises an A-D per EVI route for tag3, that gets redistributed by BR-21/BR-22 first, and BR-11/BR-12 later, translating the Ethernet Tag ID and encapsulation in each redistribution. The BR nodes implement the EVPN Multi-Homing procedures for their own Ethernet Segment as in [RFC8214], and set the P and B flags accordingly when redistributing the A-D per EVI routes, to indicate the forwarding mode to the receiving nodes. If I-ES-1 and I-ES-2 are defined as all-active multi-homing Ethernet Segments, per-flow load balancing will be performed not only by the I-PE to the Gateways in domain-1, but also by the Gateways at each domain of the EVPN VPWS service, as depicted in Figure 5. The same control plane exchange and traffic flow happen in the reverse direction, where I-PE becomes the egress PE and E-PE the ingress PE.

I-ES-1 and I-ES-2 are independent of each other, e.g., I-ES-1 can work in single-active mode, whereas I-ES-2 uses all-active mode. If that is the case, BR-11 and BR-12 run Designated Forwarded (DF) Election and BR-11 signals P=1 and B=0 (in the EVPN Layer 2 Attributes extended community) if it is elected as DF, whereas BR-12 signals P=0 and B=1 if elected as Backup DF router. I-PE then sends all traffic to BR-11, and BR-21/BR-22 send all traffic to BR-11 in the reverse direction. Since BR-21/BR-22 work in all-active mode, they both signal P=1/B=0 to both, E-PE and BR-11/BR-12. Therefore traffic from BR-11/BR-12 is sprayed to both BR-21/BR-22, and so is traffic from E-PE.

If EVPN Multi-Homing is used in the redundant Gateways, Fast Reroute procedures as in [I-D.burdet-bess-evpn-fast-reroute] MAY be applied to speed up convergence in case one of the Gateways loses its connectivity to the adjacent domain.

The Anycast Gateway and the EVPN Multi-Homing redundancy solutions can coexist. The Gateways of the same redundancy group MUST implement the same redundancy method, but different redundancy Gateway groups MAY implement different methods. In the example, BR-11/BR-12 constitutes a redundancy group and BR-21/BR-22 constitutes a different redundancy group.

4. Security Considerations

This document describes an Interconnect solution for EVPN VPWS services based on Service Gateways. While other interconnect options for EVPN VPWS exist - as outlined in Section 1.2 - the Service Gateway solution presented here offers isolation between interconnected domains. This isolation improves scalability for the PEs within each domain and helps mitigate risks, such as the leakage of unintended routes with matching Route Targets and Ethernet Tag IDs from remote, unmanaged domains into local domain PEs. Although

Service Gateways provide an additional layer of security for the PEs within the domain, they do so at the cost of requiring EVPN route processing - unlike other interconnect options. Consequently, the security considerations from [RFC8214] also apply to the Border Routers connecting the domains.

5. IANA Considerations

None.

6. Acknowledgments

7. Contributors

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365, DOI 10.17487/RFC8365, March 2018, <https://www.rfc-editor.org/info/rfc8365>.

[I-D.ietf-bess-evpn-dpath] Rabadan, J., Sathappan, S., Gautam, M., Brissette, P., and W. Lin, "Domain Path (D-PATH) for Ethernet VPN (EVPN) Interconnect Networks", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-dpath-02, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-dpath-02>.

[I-D.ietf-bess-rfc7432bis] Sajassi, A., Burdet, L. A., Drake, J., and J. Rabadan, "BGP MPLS-Based Ethernet VPN", Work in Progress, Internet-Draft, draft-ietf-bess-rfc7432bis-12, 18 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-rfc7432bis-12>.

[RFC9014] Rabadan, J., Ed., Sathappan, S., Henderickx, W., Sajassi, A., and J. Drake, "Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks", RFC 9014, DOI 10.17487/RFC9014, May 2021, <https://www.rfc-editor.org/info/rfc9014>.

[RFC8214] Boutros, S., Sajassi, A., Salam, S., Drake, J., and J. Rabadan, "Virtual Private Wire Service Support in Ethernet VPN", RFC 8214, DOI 10.17487/RFC8214, August 2017, <https://www.rfc-editor.org/info/rfc8214>.

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <https://www.rfc-editor.org/info/rfc9252>.

[RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <https://www.rfc-editor.org/info/rfc9012>.

8.2. Informative References

[RFC9062] Salam, S., Sajassi, A., Aldrin, S., Drake, J., and D. Eastlake 3rd, "Framework and Requirements for Ethernet VPN (EVPN) Operations, Administration, and Maintenance (OAM)", RFC 9062, DOI 10.17487/RFC9062, June 2021, <https://www.rfc-editor.org/info/rfc9062>.

[I-D.ietf-spring-srv6-mpls-interworking] Agrawal, S., Filsfils, C., Voyer, D., Dawra, G., Li, Z., and S. Hegde, "SRv6 and MPLS interworking", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-mpls-interworking-00, 17 October 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-mpls-interworking-00>.

[RFC9744] Sajassi, A., Ed., Brissette, P., Uttaro, J., Drake, J., Boutros, S., and J. Rabadan, "EVPN Virtual Private Wire Service (VPWS) Flexible Cross-Connect (FXC) Service", RFC 9744, DOI 10.17487/RFC9744, March 2025, <https://www.rfc-editor.org/info/rfc9744>.

[I-D.burdet-bess-evpn-fast-reroute]
Burdet, L. A., Brissette, P., Miyasaka, T., Rabadan, J., Liu, Y., and C. Lin, "EVPN Fast Reroute", Work in Progress, Internet-Draft, draft-burdet-bess-evpn-fast-reroute-09, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-burdet-bess-evpn-fast-reroute-09>.

Authors' Addresses

J. Rabadan (editor)
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America
Email: jorge.rabadan@nokia.com

S. Sathappan
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America
Email: senthil.sathappan@nokia.com

V. Prabhu
Nokia
600 March Rd
Kanata ON K2K 2T6
Canada
Email: vinod.prabhu@nokia.com

W. Lin
Juniper
United States of America
Email: wlin@juniper.net

P. Brissette
Cisco Systems
Canada
Email: pbrisset@cisco.com