

Workgroup: BESS Working Group
Internet-Draft:
draft-burdet-bess-evpn-fast-reroute-05
Published: 13 March 2023
Intended Status: Standards Track
Expires: 14 September 2023
Authors: LA. Burdet, Ed. P. Brissette T. Miyasaka
 Cisco Cisco KDDI Corporation
 J. Rabadan
 Nokia

EVPN Fast Reroute

Abstract

This document summarises EVPN convergence mechanisms and specifies procedures for EVPN networks to achieve fast and scale-independent convergence.

Requirements Language

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.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 14 September 2023.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
- [3. Requirements](#)
- [4. Solution](#)
 - [4.1. Pre-selection of Backup Path](#)
 - [4.2. Failure Detection and Traffic Restoration](#)
 - [4.2.1. Simultaneous Failures in ES](#)
 - [4.2.2. Successive and Cascading Failures in ES](#)
- [5. Redirect Labels: Forwarding Behaviors](#)
 - [5.1. Bypassing DF-Election Behavior](#)
 - [5.2. Terminal Disposition Behavior](#)
- [6. Controlled Recovery Sequence](#)
- [7. Transport Underlay](#)
 - [7.1. NVO Tunnels](#)
 - [7.1.1. Ignoring Local Bias Behavior](#)
 - [7.2. Segment Routing v6](#)
 - [7.2.1. End.DT2U.Reroute : End.DT2U with Fast Reroute](#)
 - [7.2.2. End.DX2.Reroute : End.DX2 with Fast Reroute](#)
 - [7.2.3. Conflicting Endpoint Behaviors](#)
 - [7.3. Inter-AS Option B](#)
- [8. BGP Extensions](#)
- [9. Security Considerations](#)
- [10. Acknowledgements](#)
- [11. IANA Considerations](#)
- [12. References](#)
 - [12.1. Normative References](#)
 - [12.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

EVPN convergence and failure recovery methods from different types of network failures is described in [Section 17](#) of [[I-D.ietf-bess-rfc7432bis](#)]. Similarly for EVPN-VPWS, the end of [Section 5](#) of [[RFC8214](#)] briefly evokes an egress link protection mechanism.

The fundamentals of EVPN convergence rely on a mass-withdraw technique of the Ethernet A-D per ES route to unresolve all the associated forwarding paths ([Section 9.2.2](#) of [\[I-D.ietf-bess-rfc7432bis\]](#) 'Route Resolution'). The mass-withdraw grouping approach results in suitable EVPN convergence at lower scale, but is not sufficient to meet stricter convergence requirements, often sub-second. Other control-plane enhancements such as route-prioritisation ([\[I-D.ietf-bess-rfc7432bis\]](#)) help further but still provide no guarantees.

EVPN convergence using only control-plane approaches is constrained by BGP route propagation delays, routes processing times in software and hardware programming. These are additionally often performed sequentially and linearly given the potential large scale of EVPN routes present in control plane.

This document presents a mechanism for fast reroute to minimise packet loss in the case of a link failure using EVPN redirect labels (ERLs) with special forwarding behaviors. Multiple-failures where loops may occur are addressed, as are cascading failures. A mechanism for distributing redirect labels (ERLs) alongside EVPN service labels (ESLs) is shown.

The main objective is to achieve fast convergence in EVPN networks without relying on control plane actions. The procedures in this document apply to the following EVPN services: EVPN [\[I-D.ietf-bess-rfc7432bis\]](#), EVPN-VPWS [\[RFC8214\]](#), EVPN Inter-Subnet Forwarding [\[RFC9135\]](#) and EVPN IP-VRF-to-IP-VRF models as in [Section 4.4](#) of [\[RFC9136\]](#). All the EVPN Multi-Homing modes are included.

2. Terminology

Some of the terminology in this document is borrowed from [\[RFC8679\]](#) for consistency across fast reroute frameworks.

The term 'label' when used in this document, especially when referring to ERL and ESL (below) indicates an MPLS label, a VNI (VXLAN Network Identifier) or a Segment Routing IPv6 SID, depending on the transport being used.

CE: Customer Edge device, e.g., a host, router, or switch.

PE: Provider Edge device.

Ethernet Segment (ES): A set of ethernet links connected to one or more PEs.

Ethernet Segment Identifier (ESI): A unique non-zero identifier that identifies an Ethernet segment.

Egress link:

Specific Ethernet link connecting a given PE-CE, which forms part of an Ethernet Segment.

Single-Active Redundancy Mode: When only a single PE, among all the PEs attached to an Ethernet segment, is allowed to forward traffic to/from that Ethernet segment for a given VLAN, then the Ethernet segment is defined to be operating in Single-Active redundancy mode.

All-Active Redundancy Mode: When all PEs attached to an Ethernet segment are allowed to forward known unicast traffic to/from that Ethernet segment for a given VLAN, then the Ethernet segment is defined to be operating in All-Active redundancy mode.

Port-Active Redundancy Mode: When only a single PE, among all the PEs attached to an Ethernet segment, is allowed to forward traffic to/from that Ethernet segment for the entire interface (all VLANs), then the Ethernet segment is defined to be operating in Port-Active redundancy mode.

Single-Flow-Active Redundancy Mode: When all PEs attached to an Ethernet segment are allowed to forward known unicast traffic to/from that Ethernet segment for a given VLAN, but only one does based on receiving a traffic flow from the access for that VLAN, then the Ethernet segment is defined to be operating in Single-Flow-Active redundancy mode.

DF-Election: Designated Forwarder election, as in [\[I-D.ietf-bess-rfc7432bis\]](#) and [\[RFC8584\]](#).

DF: Designated Forwarder.

Backup-DF (BDF): Backup-Designated Forwarder.

Non-DF (NDF): Non-Designated Forwarder.

AC: Attachment Circuit.

ERL: EVPN redirect label, as described in this document.

ESL: EVPN service label, as in [\[I-D.ietf-bess-rfc7432bis\]](#), [\[RFC8214\]](#), [\[RFC9135\]](#) and [\[RFC9136\]](#).

FRR: Fast Re-Route.

3. Requirements

1. EVPN multihoming is often described as 2 peering PEs. The solution MUST be generic enough to apply multiple peering PE and no artificial limit imposed on the number of peering PEs.
2. The solution MUST apply to all EVPN load-balancing modes.
3. The solution MUST be robust enough to tolerate failures of the same ES at multiple PEs. Simultaneous as well as cascading failures on the same ES must be addressed.
4. The solution MUST support EVPN [[I-D.ietf-bess-rfc7432bis](#)], EVPN-VPWS [[RFC8214](#)], EVPN Inter-Subnet Forwarding [[RFC9135](#)] and EVPN IP-VRF-to-IP-VRF models as in [Section 4.4](#) of [[RFC9136](#)].
5. An implementation of this document SHOULD support one, or many, of the above-listed services.
6. The solution SHOULD meet stringent requirements for traffic loss of EVPN services.
7. The solution MUST allow redirected-traffic to bypass port blocking states resulting from DF-Election (BDF or NDF).
8. The solution MUST be scale-independent and agnostic of EVPN route types, scale or choice of underlay.
9. The solution MUST address egress link (PE-CE link) failures.
10. The solution MUST be loop-free, and once-redirected traffic MUST never be repeatedly redirected.
11. The solution MUST NOT rely on pushing an additional label onto the label stack, or on the definition of a special-purpose label (underlay-specific to MPLS)

4. Solution

Fast convergence in EVPN networks is achieved using a combined approach to minimising traffic loss:

*Local failure detection and restoration of traffic flows in minimal time using a pre-computed redirect path;

*Restoration of optimal traffic paths, and reconvergence of EVPN control plane with EVPN mass withdraw.

The solution presented in this document addresses the local failure detection and restoration, without impeding on or impacting existing EVPN control plane convergence mechanisms.

Consider the following EVPN topology where PE1 and PE2 are multihoming PE's on a shared ES, ESI1. EVPN (known unicast) or EVPN-VPWS traffic from CE1 to CE2 is sent to PE1 and PE2 using EVPN service labels ESL1 and/or ESL2 (depending on load-balancing mode of the ESI1 interfaces).

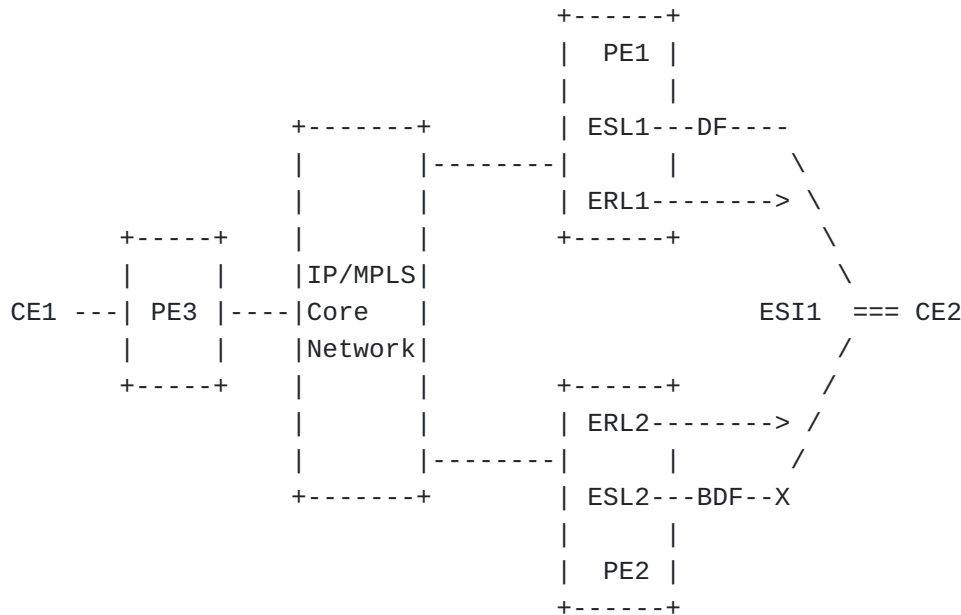


Figure 1: EVPN Multihoming with service and redirect labels

Alongside the service labels ESL1 and ESL2, two redirect labels ERL1 and ERL2 are allocated with special forwarding behaviors, as detailed in [Section 5](#). Fast-reroute and use of the ERLs is shown in [Section 4.2](#)

4.1. Pre-selection of Backup Path

EVPN DF-Election lends itself well to the selection of a pre-computed path amongst any given number of peering PE's by providing a DF-Elected and BDF-Elected node at the <EVI, ESI> granularity ([RFC8584] and [I-D.ietf-bess-rfc7432bis]).

In All-active mode, all PE's in the Ethernet Segment are actively forwarding known unicast traffic to the CE. In Single-active and Port-Active modes, only a single PE in the Ethernet Segment is actively forwarding known unicast traffic to the CE: the DF-Elected PE. The BDF-Elected PE is next to be elected in the redundancy group and is already known. In Single-flow-active mode ([I-D.ietf-bess-evpn-l2gw-proto]), only a single PE in the Ethernet

Segment is actively forwarding known unicast to the CE for a given flow: the PE which initially received that flow from the Ethernet-Segment. The backup PE is the multihoming peer in the redundancy group, referred to as "BDF" for consistency with other redundancy modes.

For consistency across PEs and load-balancing modes, the backup path selected should be in order of {DF, BDF, NDF1, NDF2, ...}. The DF-Elected PE selects the next-best BDF-Elected as backup and all BDF- and NDF-Elected nodes select the best DF-Elected for the protection of their egress links.

- *PE1 (DF) selects PE2 as BDF,

- *PE1 (DF) uses the ERL2 label signaled by PE2 to redirect the traffic of its failed local AC connected to CE2,

- *PE2 (BDF) uses the ERL1 label signaled by PE1 to redirect the traffic of its failed local AC connected to CE2,

- *PE..n (NDF) use the ERL2 label signaled by PE2 to redirect the traffic of their failed local AC connected to CE2.

The use of PE2's ERL2 as redirect label applies to local failures in all load-balancing modes at PE1.

The number of peering PEs is not limited by existing DF-Election algorithms. A solution based on DF-Election supports subsequent redirection upon multiple cascading failures, once a new DF-Election has occurred. Pre-selection of a backup path is supported by all current DF-Election algorithms, and more generally by all algorithms supporting BDF-Election, as recommended in ([[I-D.ietf-bess-rfc7432bis](#)]).

4.2. Failure Detection and Traffic Restoration

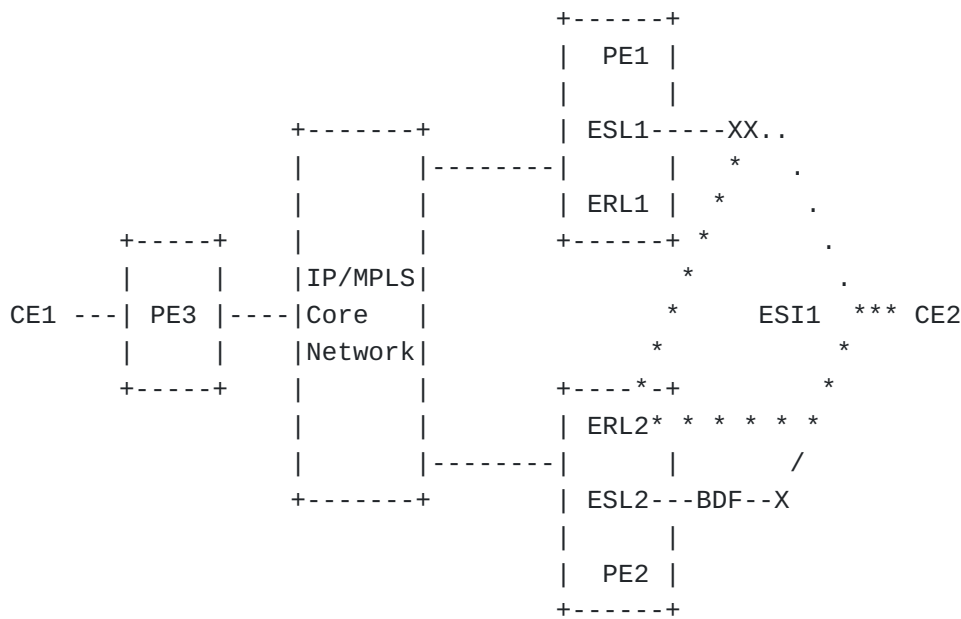


Figure 2: EVPN Multihoming failure scenario

The procedures for forwarding known unicast packets received from a remote PE on the local redirect label follow [Section 13.2.2](#) of [\[I-D.ietf-bess-rfc7432bis\]](#) for known unicast traffic. Since the CE next-hop forwarding information reflects the current BDF state of the AC, additional steps to bypass blocking state and preventing another re-direction are applied, as described further in this document.

Consider the EVPN multihoming topology in [Figure 1](#), and a traffic flow from CE1 to CE2 which is currently using EVPN service label ESL1 and forwarded through the core arriving at PE1. When the local AC representing the <EVI,ESI> pair is protected using the fast-reroute solution, the pre-computed backup path's redirect label (i.e. ERL2 from BDF-Elected PE2) is installed against the AC.

Under normal conditions, PE1 disposition using ESL1 will result in forwarding the packet to the CE by selecting the local AC associated with the EVPN service label ([\[RFC8214\]](#), [\[I-D.ietf-bess-rfc7432bis\]](#)). When this local AC is in failed state, the fast-reroute solution at PE1 will begin rerouting packets using the BDF-Elected peer's nexthop and ERL2. ERL2 is chosen for redirected traffic and not ESL2 to prevent loops and overcome DF-Election timing as described in [Sections 5.2](#) and [5.1](#) respectively.

4.2.1. Simultaneous Failures in ES

In EVPN multihoming where the CE connects to peering PEs through link aggregation (LAG), a single LAG failure at the CE may manifest as multiple ES failures at all peering PEs simultaneously.

As all peering PEs would enable simultaneously the fast-reroute mechanism, redirection would be permanent causing a traffic storm or until TTL expires.

Once-redirected traffic may not be redirected again, according to the terminal nature of ERLs described in [Section 5.2](#)

4.2.2. Successive and Cascading Failures in ES

Trying to support cascading failures by redirecting once-redirected traffic is substantially equivalent to simultaneous failures above.

Once-redirected traffic may not be redirected again, according to the terminal nature of ERLs described in [Section 5.2](#) and loss is to be expected until EVPN control plane reconverges for double-failure scenarios.

In a scenario with 3 peering PEs (PE1-DF, PE2-BDF, PE3-NDF) where PE1 fails, followed by a PE2 failure before control-plane reconvergence, there is no reroute of traffic towards PE3 because the reroute-label is terminal.

In such rapid-succession failures, it is expected that control plane must first correct for the initial failure and DF-Elect PE2 as new-DF and PE3 as the new-BDF. PE2 to PE3 redirection would then begin, unless control-plane is rapid enough to correct directly, and elect PE3 new-DF.

5. Redirect Labels: Forwarding Behaviors

The EVPN redirect labels MUST be downstream assigned, and it is directly associated with the <EVI,ESI> AC being egress protected. The special forwarding characteristics and use of an EVPN redirect label (ERL) described below, are a matter of local significance only to the advertising PE (which is also the disposition PE).

Special behaviors to the ERLs do not affect any other PEs or transit P nodes. There are no extra labels appended to the label stack in the IP/MPLS network and the ERL appears to label-switching transit nodes as would any other EVPN service label. Since they appear as

EVPN service labels, ERL labels do not have any impact on Flow-Label or Control-Word procedures in [[I-D.ietf-bess-rfc7432bis](#)].

- *Traffic redirection and use of reroute labels may create routing loops upon multiple failures. Such loops are detrimental to the network and may cause congestion between protected PEs.

- *Local restoration and redirection is meant to occur much faster than control-plane operations, meaning redirected packets may arrive at the BDF PE long before a DF-Election operation unblocks the egress link.

Two special forwarding characteristics and behaviors of EVPN redirect labels are described below to mitigate these issues.

5.1. Bypassing DF-Election Behavior

Local detection and restoration at DF-Elected PE1 will begin rapidly redirecting traffic onto the backup path selected (PE2). Redirected packets will arrive at the Backup-DF port much faster than control plane DF-Election at the Backup-DF peer is capable of unblocking its local egress link for the shared ES (ESI1). All redirected traffic would drop at Backup-DF and no net reduction in traffic loss is achieved.

Traffic restoration remains dependant upon ES route or Ethernet A-D per ES/EVI routes withdrawal for a DF-Election operation and for PE1 to assume the traffic forwarding role. This is especially important in single-active load-balancing mode where known unicast traffic is blocked.

To mitigate this, the redirect labels allocated must carry a special attribute in the local forwarding and decapsulation chain: for traffic received on the ERL when the AC is up, an override to the DF-Election is applied and traffic from the ERL will bypass the local Backup-DF blocking state. Once EVPN control plane reconverges, traffic from the ERL will cease and the optimal forwarding path based on ESLs will resume.

The EVPN redirect label MUST carry a context locally, such that from disposition to egress redirected packets are allowed to bypass the Backup-DF blocking state that would otherwise drop. Similarly, this may open the gate to the traffic in the reverse direction.

In Port-Active mode, the Backup-DF interface may signal Out-of-Service but remain in Up/Backup state: to support EVPN Fast Reroute, the CE must be able to receive traffic from an OOS LAG link.

5.2. Terminal Disposition Behavior

The reroute scheme is susceptible to loops and persistent redirects between peering PEs which have setup FRR redirection. Consider the scenario where both CE-facing interfaces fail simultaneously, fast reroute will be activated at both PE1 and PE2 effectively bouncing a redirected packet between the two PEs indefinitely (or until the TTL expires) causing a traffic storm.

To prevent this, a distinction is made between 'regular' EVPN service labels for disposition (i.e. known unicast EVI label or EVPN-VPWS label) and reroute labels with terminal disposition.

At the redirecting PE2, we consider the case of ESL2 vs. ERL2 , where both are locally allocated and provided in EVPN routes (downstream allocation) to BGP peers:

1. EVPN Service label, ESL2:

- *Regular MAC-lookup or traffic forwarding occurs towards the access AC.

- *If the AC is up, traffic will exit the interface, subject to local blocking state on the AC from DF-Election.

- *If the AC is down and fast-reroute procedures are enabled, traffic may be re-encapsulated using BDF peer's redirect label ERL1 (if received).

2. EVPN Reroute label, ERL2:

- *Regular MAC-lookup or traffic forwarding occurs towards the access AC.

- *If the AC is up, traffic will apply an override to DF-Election and bypass the local blocking state on the AC.

- *If the AC is down, traffic is dropped. No reroute must occur of once-rerouted traffic. Redirecting towards peer's redirect label ERL1 is explicitly prevented.

The ERL acts like a local cross-connect by providing a direct channel from disposition to the AC. ERLs are terminal-disposition and prevents once-redirection packets from being redirected again. With this forwarding attribute on ERLs, known only locally to the downstream-allocating PE, redirection is achieved without growing the label stack with another special purpose label.

6. Controlled Recovery Sequence

Fast reroute mechanisms such as the one described in this document generally provide a way to preserve traffic flows at failure time. Use of fast reroute in EVPN, however, permits setting up a controlled recovery sequence to shorten the period of loss between an interface coming up and the EVPN DF-Election procedures and default timers for peer discovery.

The benefit of a controlled recovery sequence is amplified when used in conjunction with [[I-D.ietf-bess-evpn-fast-df-recovery](#)] (synchronised DF-Election)>

7. Transport Underlay

The solution is agnostic to transport underlays, for instance similar behavior is carried forward for NVO tunnels (VXLAN) and SRv6.

7.1. NVO Tunnels

The rerouting procedures and behaviors in this document apply as well for [[RFC8365](#)] NVO tunnels.

For MPLS-based NVO tunnels, i.e. MPLSoGRE, MPLSoUDP, etc., no additional behaviors are required.

For non-MPLS NVO tunnels, the labels are 24-bit VNIs, not downstream assigned and usually global, i.e. same value for all the PEs attached to the BD. In this case, the rerouting mechanisms described in this document would not work without some additional behaviors: the rerouting mechanism needs to avoid local-bias split-horizon filtering upon reception of the redirected packets. For non-MPLS NVO tunnels, an additional identifier is advertised in Ethernet A-D per EVI routes to enable EVPN Fast Reroute.

7.1.1. Ignoring Local Bias Behavior

Non-MPLS NVO tunnel encapsulations may use local-bias procedures instead of ES label-based split-horizon (for EVPN multihoming). This means that, e.g. when PE1 sends redirected traffic to multihoming peer PE2 with the ERL VNI, PE2 will drop the packets due to the filtering based on the tunnel source IP. To support non-MPLS NVO tunnels such as VXLAN, PE2 in the example above needs to bypass the source IP based filtering if the VNI identifies a local redirection instance. The split-horizon filtering would be based on source-IP + FRR-VNI, as opposed to source-IP only. Since the VNI is global and not e.g. downstream-assigned, a VNI must be allocated per ES,EVI for the rerouting mechanisms described in this document to apply.

7.2. Segment Routing v6

Ethernet A-D per EVI routes are advertised along with the Service SID used for End.DX2 or End.DT2U behaviors [Section 6.1.2](#) of [\[RFC9252\]](#). These advertisements correspond to the ESL behavior in this document (EVPN Service SID). An additional EVPN Redirect SID is advertised in Ethernet A-D per EVI routes to enable EVPN Fast Reroute, with one of 2 new behaviors.

7.2.1. End.DT2U.Reroute : End.DT2U with Fast Reroute

The "End.DT2U with Fast Reroute" behavior ("End.DT2U.Reroute" for short) is a variant of the End.DT2U behavior.

The End.DT2U.Reroute behavior is defined for the fast-reroute application between two EVPN multi-homing peers, and extends the base End.DT2U behavior. This behavior takes an optional Fast Reroute argument: "Arg.FR2". This argument provides a local mapping to Attachment Circuit (EVI/ESI) for the received traffic, which also implements the forwarding behaviors in [Section 5](#).

Any SID instance of this behavior may be used in two ways:

1. by ingress PEs not performing any reroute (such as PE3 in [Figure 1](#)) by setting the Arg.FR2 argument as zero for handling at an egress PE that is the same as End.DT2U
2. by peering PEs performing redirection (such as PE1 in [Figure 2](#)), by setting the argument Arg.FR2 with a non-zero value for the reroute handling in addition to the End.DT2U functionality

Thus, the SID entry for this behavior when instantiated in the FIB performs the disposition of both base L2 Table traffic (i.e., the base End.DT2U behavior) traffic as well as rerouted traffic (i.e., the End.DT2U+Arg.FR2 handling). End.DT2U processing is as in [Section 4.11](#) of [\[RFC8986\]](#).

When processing the Upper-Layer header of a packet matching a FIB entry locally instantiated as an End.DT2U.Reroute SID, N does the following:

```

S01. If (Upper-Layer header type == 143(Ethernet) ) {
S02.   Remove the outer IPv6 header with all its extension headers
S03.   If (Arg.FR2 is 0) {
S04.     Process as per Section 4.11 of [RFC8986] (End.DT2U)
S05.   } Else {
S06.     Lookup the egress interface L2 OIF I for Arg.FR2
S07.     If (L2 OIF interface I is down) {
S08.       Drop the Ethernet frame
S09.     } Else {
S10.       Forward the Ethernet frame to the OIF I
           bypassing any EVPN DF-Election blocking state
S11.     }
S12. } Else {
S13.   Process as per Section 4.1.1 of [RFC8986]
S14. }

```

To maintain backwards-compatibility, both End.DT2U.Reroute and End.DT2U Behavior SIDs MAY be advertised together whereby legacy receivers ignore the SRv6 SID of unknown behavior End.DT2U.Reroute.

The SRv6 L2 Service TLV in this case will carry two SRv6 SID Information sub-TLVs:

- *the first one with the base End.DT2U behavior and

- *the second one with the End.DT2U.Reroute behavior variant.

The second one will have a non-zero Arg length (AL) and convey Arg.FR2 embedded in the advertised SID

When advertised alongside an End.DT2U EVPN Service SID, the End.DT2U.Reroute EVPN Reroute SID MUST be identical to the End.DT2U except for the inclusion of an Argument Arg.FR2. Both SRv6 SIDs can use transposition since the function MUST be identical between the 2 SIDs. A receiver unable to validate the applicability of arguments for SRv6 Endpoint Behaviors that are unknown to it MUST ignore the End.DT2U.Reroute SID ([Section 3.2.1](#) of [\[RFC9252\]](#)).

Following is an example representation of the BGP Prefix-SID Attribute encoding in this case for a 16-bit argument Arg.FR2 (0xaaaa):

```

BGP Prefix SID Attr:
  SRv6 L2 Service TLV:
    SRv6 SID Information sub-TLV:
      SID: 2001:db8:b:1:fb1::
      Behavior: End.DT2U
    SRv6 SID Structure sub-sub-TLV:
      LBL: 48, LNL: 16, FL: 16, AL: 0, TPOS-L: 0, TPOS-O: 0
  SRv6 SID Information sub-TLV:
    SID: 2001:db8:b:1:fb1::aaaa::
    Behavior: End.DT2U.Reroute
    SRv6 SID Structure sub-sub-TLV:
      LBL: 48, LNL: 16, FL: 16, AL: 16, TPOS-L: 0, TPOS-O: 0

```

Figure 3: EVPN Route Type 1 with dual End.DT2U SIDs

When both End.DT2U.Reroute and End.DT2U are advertised, the ingress PE not performing reroute MUST use the End.DT2U as the EVPN Service SID.

7.2.2. End.DX2.Reroute : End.DX2 with Fast Reroute

The "End.DX2 with Fast Reroute" behavior ("End.DX2.Reroute" for short) is a variant of the End.DX2 behavior.

The text in this section mirrors that of [Section 7.2.1](#) (End.DT2U.Reroute) and is included for completeness' sake.

The End.DX2.Reroute behavior is defined for the fast-reroute application between two EVPN multi-homing peers, and extends the base End.DX2 behavior. This behavior takes an optional Fast Reroute argument: "Arg.FR2". This argument provides a local mapping to Attachment Circuit (EVI/ESI) for the received traffic, which also implements the forwarding behaviors in [Section 5](#).

Any SID instance of this behavior may be used in two ways:

1. by ingress PEs not performing any reroute (such as PE3 in [Figure 1](#)) by setting the Arg.FR2 argument as zero for handling at an egress PE that is the same as End.DX2
2. by peering PEs performing redirection (such as PE1 in [Figure 2](#)), by setting the argument Arg.FR2 with a non-zero value for the reroute handling in addition to the End.DX2 functionality

Thus, the SID entry for this behavior when instantiated in the FIB performs the disposition of both base L2 Table traffic (i.e., the base End.DX2 behavior) traffic as well as rerouted traffic (i.e., the End.DX2+Arg.FR2 handling). End.DX2 processing is as in [Section 4.9](#) of [[RFC8986](#)].

When processing the Upper-Layer header of a packet matching a FIB entry locally instantiated as an End.DX2.Reroute SID, N does the following:

```
S01. If (Upper-Layer header type == 143(Ethernet) ) {
S02.   Remove the outer IPv6 header with all its extension headers
S03.   If (Arg.FR2 is 0) {
S04.     Process as per Section 4.9 of [RFC8986] (End.DX2)
S05.   } Else {
S06.     Lookup the egress interface L2 OIF I for Arg.FR2
S07.     If (L2 OIF interface I is down) {
S08.       Drop the Ethernet frame
S09.     } Else {
S10.       Forward the Ethernet frame to the OIF I
           bypassing any EVPN DF-Election blocking state
S11.     }
S12. } Else {
S13.   Process as per Section 4.1.1 of [RFC8986]
S14. }
```

To maintain backwards-compatibility, both End.DX2.Reroute and End.DX2 Behavior SIDs MAY be advertised together whereby legacy receivers ignore the SRv6 SID of unknown behavior End.DX2.Reroute.

The SRv6 L2 Service TLV in this case will carry two SRv6 SID Information sub-TLVs:

- *the first one with the base End.DX2 behavior and

- *the second one with the End.DX2.Reroute behavior variant.
The second one will have a non-zero Arg length (AL) and convey Arg.FR2 embedded in the advertised SID

When advertised alongside an End.DX2 EVPN Service SID, the End.DX2.Reroute EVPN Reroute SID MUST be identical to the End.DX2 except for the inclusion of an Argument Arg.FR2. Both SRv6 SIDs can use transposition since the function MUST be identical between the 2 SIDs. A receiver unable to validate the applicability of arguments for SRv6 Endpoint Behaviors that are unknown to it MUST ignore the End.DX2.Reroute SID ([Section 3.2.1](#) of [\[RFC9252\]](#)).

Following is an example representation of the BGP Prefix-SID Attribute encoding in this case for a 16-bit argument Arg.FR2 (0xaaaa):


```

BGP Prefix SID Attr:
  SRv6 L2 Service TLV:
    SRv6 SID Information sub-TLV:
      SID: 2001:db8:b:1:fb1::
      Behavior: End.DX2
      SRv6 SID Structure sub-sub-TLV:
        LBL: 48, LNL: 16, FL: 16, AL: 0, TPOS-L: 0, TPOS-O: 0
    SRv6 SID Information sub-TLV:
      SID: 2001:db8:b:1:fb1:aaaa::
      Behavior: End.DX2.Reroute
      SRv6 SID Structure sub-sub-TLV:
        LBL: 48, LNL: 16, FL: 16, AL: 16, TPOS-L: 0, TPOS-O: 0

```

Figure 4: EVPN Route Type 1 with dual End.DX2 SIDs

When both End.DX2.Reroute and End.DX2 are advertised, the ingress PE not performing reroute MUST use the End.DX2 as the EVPN Service SID.

7.2.3. Conflicting Endpoint Behaviors

End.DT2U.Reroute and End.DX2.Reroute are variants of their respective base behaviours and when two SIDs are advertised together in an Ethernet A-D per EVI route, the variant advertised MUST be the same as base behaviour.

In other words, advertisement of an End.DT2U.Reroute variant alongside an End.DX2 base is unusable and SHALL be discarded by receivers, and similarly an End.DX2.Reroute variant advertised alongside an End.DT2U base SHALL be discarded by receivers.

7.3. Inter-AS Option B

EVPN multi-homing peers in different AS are rather an exception. In Inter-AS Option B or inter-domain scenarios, the ASBR/ABR and BGP route-reflectors with nexthop-self procedures are extended:

- *Prior to this spec the ABR/ASBR receives the Ethernet A-D per EVI route, programs a label swap operation and redistributes the route with a new allocated label in the NLRI's label field.

- *To implement the procedures in this document, the ABR/ASBR needs to allocate two downstream labels for each Ethernet-A-D per EVI route: one for the NLRI's label (ERL) and another one for the ESI Label Extended Community label (ESL). A label swap operation is programmed for both ERL and ESL labels.

8. BGP Extensions

While this document describes a new behavior, there are no new BGP extensions required to advertise the redirect label(s) used for EVPN egress link protection. The ESI Label Extended Community defined in

[Section 7.5](#) of [[I-D.ietf-bess-rfc7432bis](#)] may be advertised along with Ethernet A-D routes:

*When advertised with an Ethernet A-D per ES route, it enables split-horizon procedures for multihomed sites as described in [Section 8.3](#) of [[I-D.ietf-bess-rfc7432bis](#)];

*When advertised with an Ethernet A-D per EVI route, it enables link protection and fast-reroute procedures for multihomed sites as described in this document. The label value represents the per-<EVI,ESI> EVPN redirect label (ERL). The Flags field SHOULD NOT be set and MUST be ignored.

Prior to this document, advertising the ESI Label Extended Community along with an Ethernet A-D per EVI route (Ethertag different than MAX-ET) was undefined, and presumably ignored.

Remote PEs SHOULD NOT use the ERLs as a substitution for ESLs in route resolution, and is especially not to be confused with the aliasing and backup path ESL as described and used in [Section 8.4](#) of [[I-D.ietf-bess-rfc7432bis](#)].

9. Security Considerations

The mechanisms in this document use the EVPN control plane as defined in [[I-D.ietf-bess-rfc7432bis](#)] and [[RFC8214](#)], and the security considerations described therein are equally applicable. Reroute labels redistributed in EVPN control plane are meant for consumption by the peering PE in a same ES. It is, however, visible in the EVPN control plane to remote peers. Care shall be taken when installing reroute labels, since their use may result in bypassing DF-Election procedures and lead to duplicate traffic at CEs if incorrectly installed.

10. Acknowledgements

Authors would like to thank Ketan Talaulikar for his review of SRv6 procedures in this document.

11. IANA Considerations

This document introduces two new Endpoint behaviors. This document requests IANA assign a two new values and update the "SRv6 Endpoint Behaviors" subregistry under the top-level "Segment Routing" registry as follows:

Value	Hex	Endpoint Behavior	Reference
TBD	TBD	End.DT2U.Reroute	This document
TBD	TBD	End.DX2.Reroute	This document

Table 1: SRv6 Endpoint Behaviors Subregistry

12. References

12.1. Normative References

- [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-06, 5 January 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-rfc7432bis-06>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8584] Rabadan, J., Ed., Mohanty, S., Ed., Sajassi, A., Drake, J., Nagaraj, K., and S. Sathappan, "Framework for Ethernet VPN Designated Forwarder Election Extensibility", RFC 8584, DOI 10.17487/RFC8584, April 2019, <<https://www.rfc-editor.org/info/rfc8584>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

12.2. Informative References

- [I-D.ietf-bess-evpn-fast-df-recovery] Brissette, P., Sajassi, A., Burdet, L. A., Drake, J., and J. Rabadan, "Fast Recovery for EVPN Designated Forwarder Election", Work in Progress, Internet-Draft,

draft-ietf-bess-evpn-fast-df-recovery-06, 24 August 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-fast-df-recovery-06>>.

[I-D.ietf-bess-evpn-l2gw-proto] Brissette, P., Sajassi, A., Burdet, L. A., and D. Voyer, "EVPN Multi-Homing Mechanism for Layer-2 Gateway Protocols", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-l2gw-proto-02, 24 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-l2gw-proto-02>>.

[RFC8679] Shen, Y., Jeganathan, M., Decraene, B., Gredler, H., Michel, C., and H. Chen, "MPLS Egress Protection Framework", RFC 8679, DOI 10.17487/RFC8679, December 2019, <<https://www.rfc-editor.org/info/rfc8679>>.

[RFC9135] Sajassi, A., Salam, S., Thoria, S., Drake, J., and J. Rabadan, "Integrated Routing and Bridging in Ethernet VPN (EVPN)", RFC 9135, DOI 10.17487/RFC9135, October 2021, <<https://www.rfc-editor.org/info/rfc9135>>.

[RFC9136] Rabadan, J., Ed., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in Ethernet VPN (EVPN)", RFC 9136, DOI 10.17487/RFC9136, October 2021, <<https://www.rfc-editor.org/info/rfc9136>>.

[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>>.

Authors' Addresses

Luc Andre Burdet (editor)
Cisco

Email: lburdet@cisco.com

Patrice Brissette
Cisco

Email: pbrisset@cisco.com

Takuya Miyasaka
KDDI Corporation

Email: ta-miyasaka@kddi.com

Jorge Rabadan

Nokia

Email: jorge.rabadan@nokia.com