

EVPN Interworking with IPVPN

draft-ietf-bess-evpn-ipvpn-interworking-08

Jorge Rabadan (Nokia)

Ali Sajassi (Cisco)

Eric Rosen (Juniper)

John Drake (Juniper)

Wen Lin (Juniper)

Jim Uttaro

Adam Simpson (Nokia)

IETF117, Jul 2023

San Francisco

Agenda

draft-ietf-bess-evpn-ipvpn-interworking

Introduction and Document History

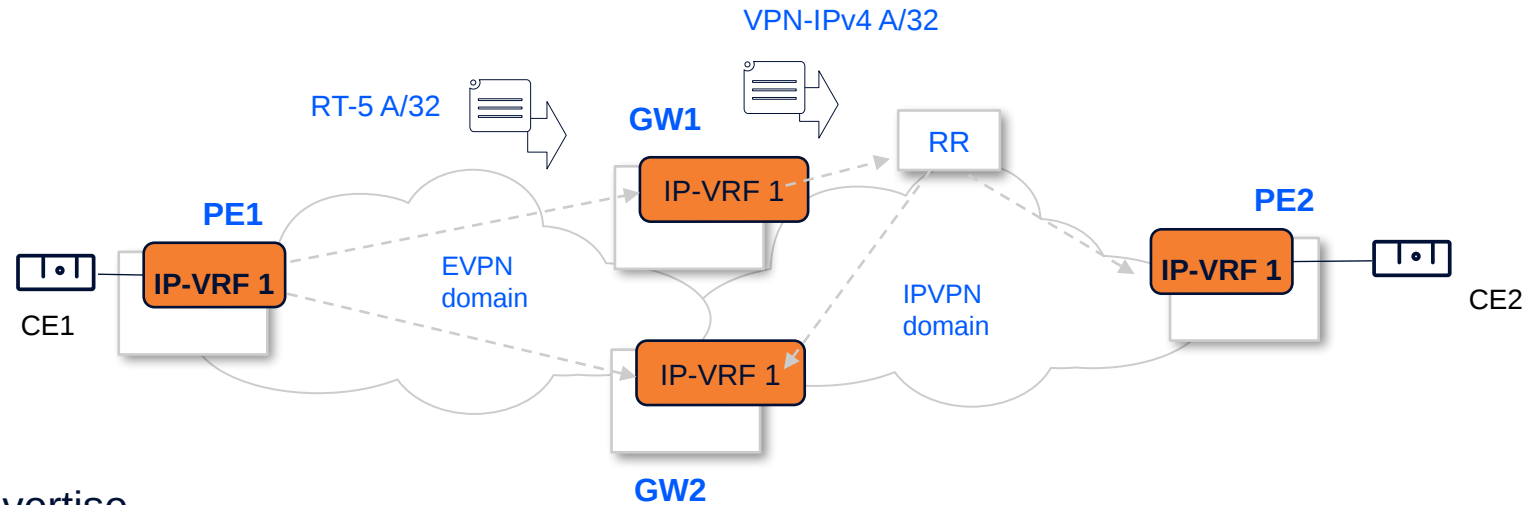
Changes in rev 08

Next Steps

What is this document about

Main Concepts

A VPN may span **multiple domains** where different Inter-Subnet-Forwarding SAFIs (ISF SAFIs) may provide inter-subnet connectivity.



Interworking PEs / Gateways receive and advertise IP Prefixes in different ISF SAFIs and therefore there is a need to specify certain functions across ISF SAFIs, such as:

- BGP path selection
- Loop Prevention
- Path Attribute Propagation
- Route Aggregation

The document defines D-PATH as a path attribute with value 36 (IANA allocated) as optional transitive

- A sequence of domain segments (segment length, value)
- Where value is a sequence of domains with format <DOMAIN-ID:ISF_SAFI_TYPE>

What is this document about

Domains and D-PATH

Domains and Domain-IDs

A GW is always configured with multiple Domain-IDs

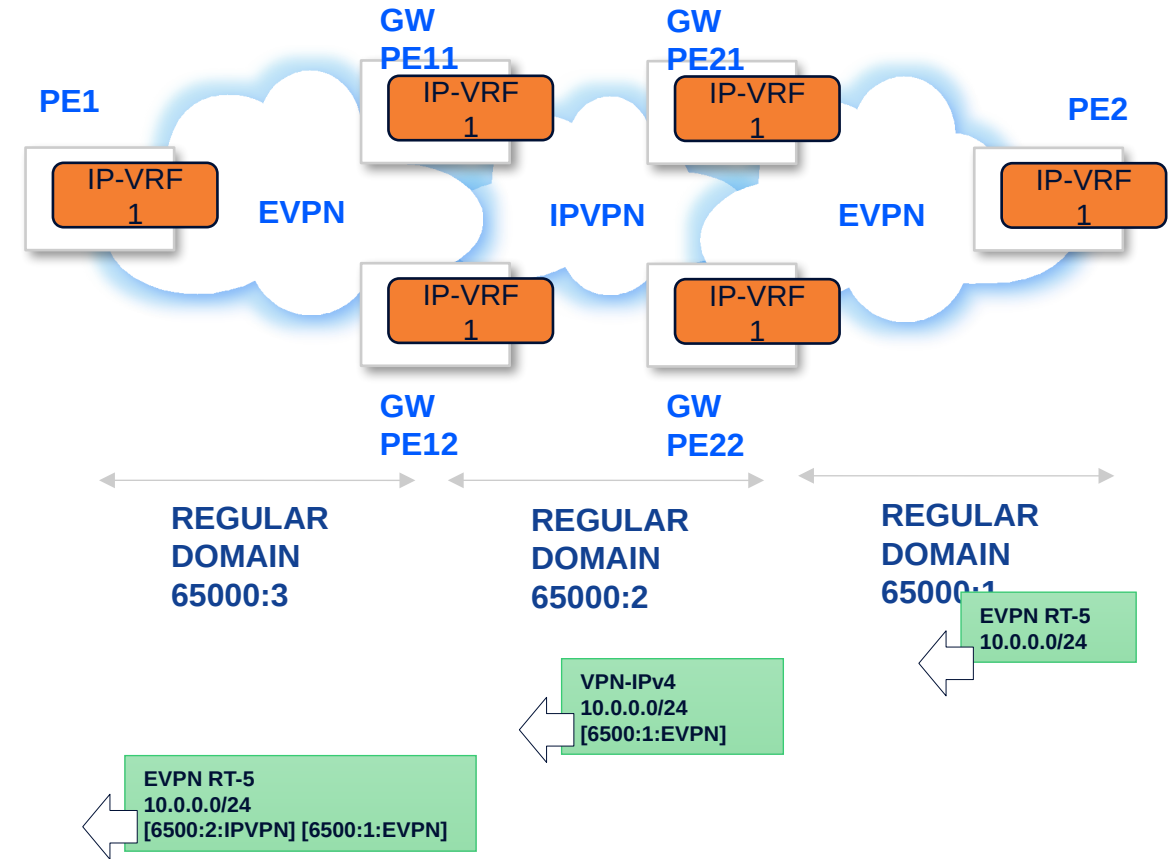
Two PEs are in the same Domain if there are no (tenant) IP lookups done in intermediate routers

D-PATH and Loop Prevention

D-PATH provides visibility of the domains and SAFIs through which the IP Prefix has gone

Modifies best path selection for IPVPN and EVPN routes (shortest D-PATH wins) and prevent loops. E.g.,:

- GW PE21 receives EVPN 10.0.0.0/24 from PE2 with D-PATH=0 and VPN-IPv4 from GW PE22 with D-PATH=(6500:1:EVPN)
- GW PE21 selects the EVPN route based on shortest D-PATH
- GWs MUST NOT re-export a route that contains a local domain-id



Document History

draft-ietf-bess-evpn-ipvpn-interworking

- Initial document draft-rabadan-sajassi-bess-evpn-ipvpn-interworking originated based on the consensus among a few vendors and operators to merge of draft-rs-bess-evpn-attr-prop and draft-sajassi-bess-evpn-ipvpn-interop
- Presented multiple times at IETF
- Adopted by the BESS Working Group in March 2019, early allocation by IANA
- Implemented by multiple vendors and interoperability proven at public events, e.g. <https://eantc.de/fileadmin/eantc/downloads/events/2022/EANTC-InteropTest2022-TestReport.pdf>
- BESS Working Group Last Call in November 2020, sent to IDR chairs and IDR WG for review
- Document reviewed by IDR chairs (big thanks to Sue and Jeff), with several issues raised and addressed by authors
- Main outstanding concerns related to D-PATH and best path selection (quoting IDR chairs):
 1. D-PATH changes fundamental route selection behavior for BGP
 2. Can result in inconsistent route selection
 3. Accidentally applicable to Internet-scoped BGP routes (attribute escape)

Changes in rev 08

Addressing concerns and other comments

1. Remove complete support for D-PATH on SAFI 1 routes, including proper normative language
 1. D-PATH MUST NOT be advertised along with routes different from IPVPN and EVPN routes, and only changes best path selection for IPVPN and EVPN routes.
 2. By default, the BGP D-PATH attribute is not advertised and MUST be explicitly enabled by configuration on the Gateway PEs.
 3. If D-PATH is received along with routes of AFI/SAFI different from the IPVPN and EVPN families, the behavior treat-as-withdraw is applied.
 4. The use of D-PATH is restricted to "walled garden" Virtual Private Networks, and the operator MUST NOT turn on the generation of D-PATH along with IPVPN and/or EVPN routes if there are CEs attached to a PE (of any domain in the Virtual Private Network) that are connected to the Internet.
 5. In addition, a Gateway PE MUST support the removal of the D-PATH attribute on import and on export, based on configuration.
2. New discussion on the Security Considerations section about "walled garden" VPNs and what happens in case of D-PATH escape
3. Addressed John Scudder's comment on mailing list

Next Steps

1. Request Feedback from WG, BESS and IDR chairs about rev 08
2. Make progress and publish ASAP

Thank you