

draft-rabadan-sajassi-bess-evpn-ipvpn-interworking-01

Jorge Rabadan (Nokia)

Ali Sajassi (Cisco)

Eric Rosen (Juniper)

John Drake (Juniper)

Wen Lin (Juniper)

Jim Uttaro (AT&T)

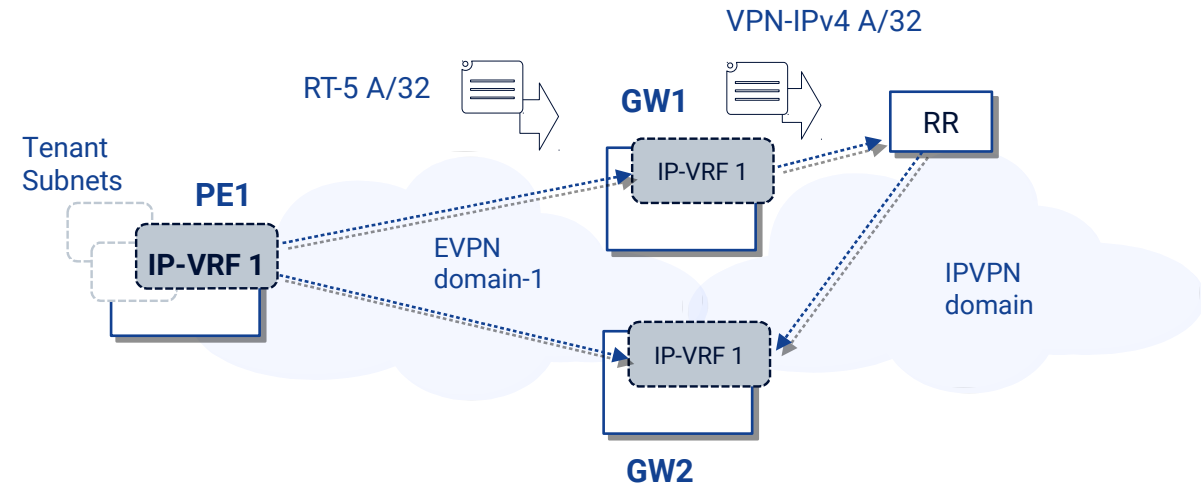
Adam Simpson (Nokia)

IETF102, Jul 2018

Montreal

Refreshing the objective

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



This document addresses the interworking aspects between ISF-SAFIs when one of them is EVPN:

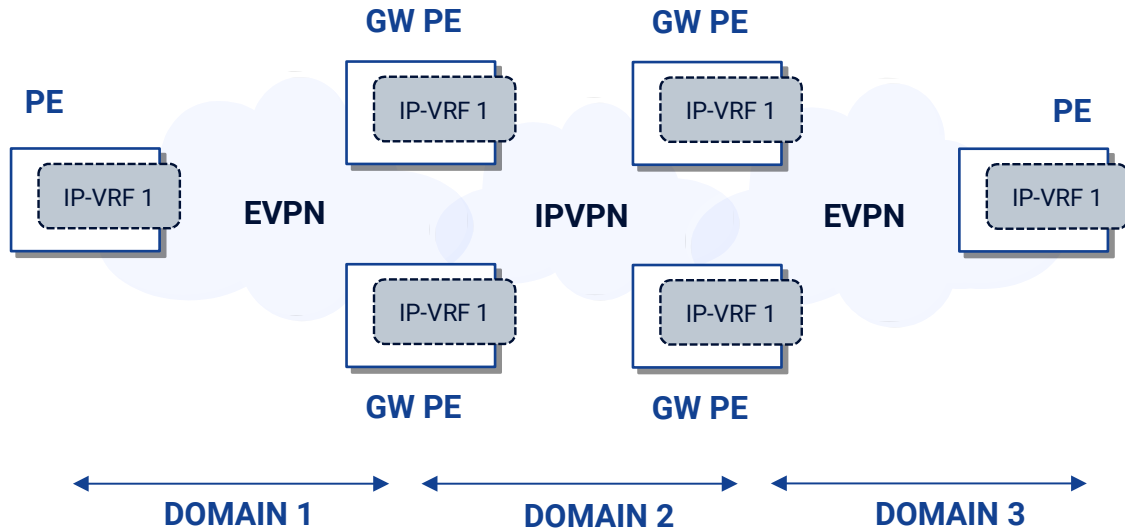
- BGP route selection

- Loop Prevention

- Path Attribute Propagation

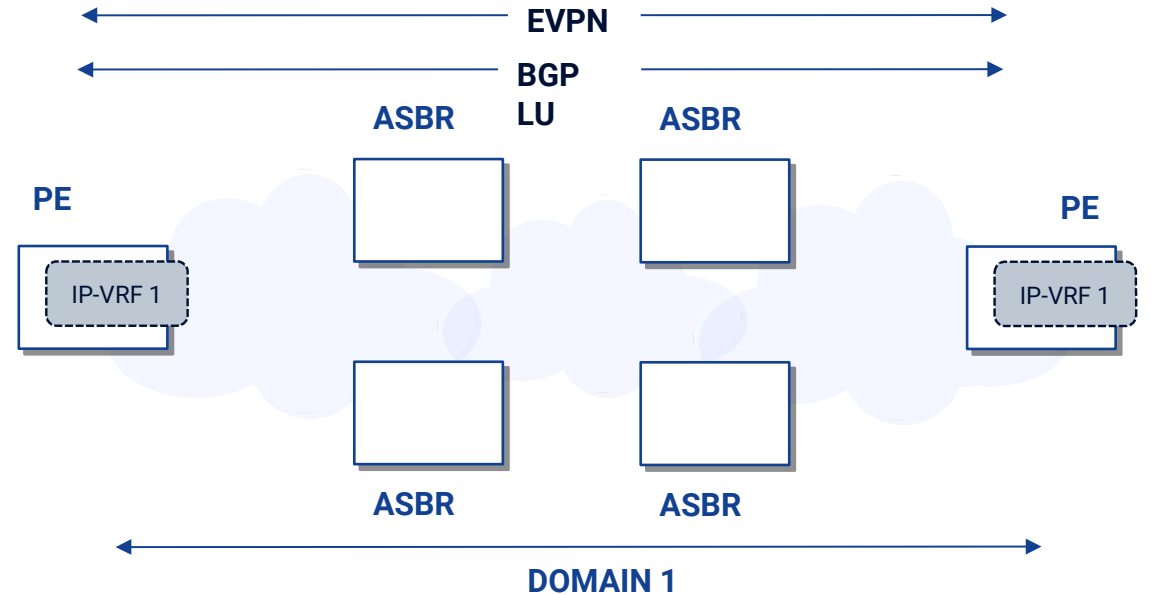
- Route Aggregation (new in rev 01)**

How can a DOMAIN be defined?



DOMAIN IDs

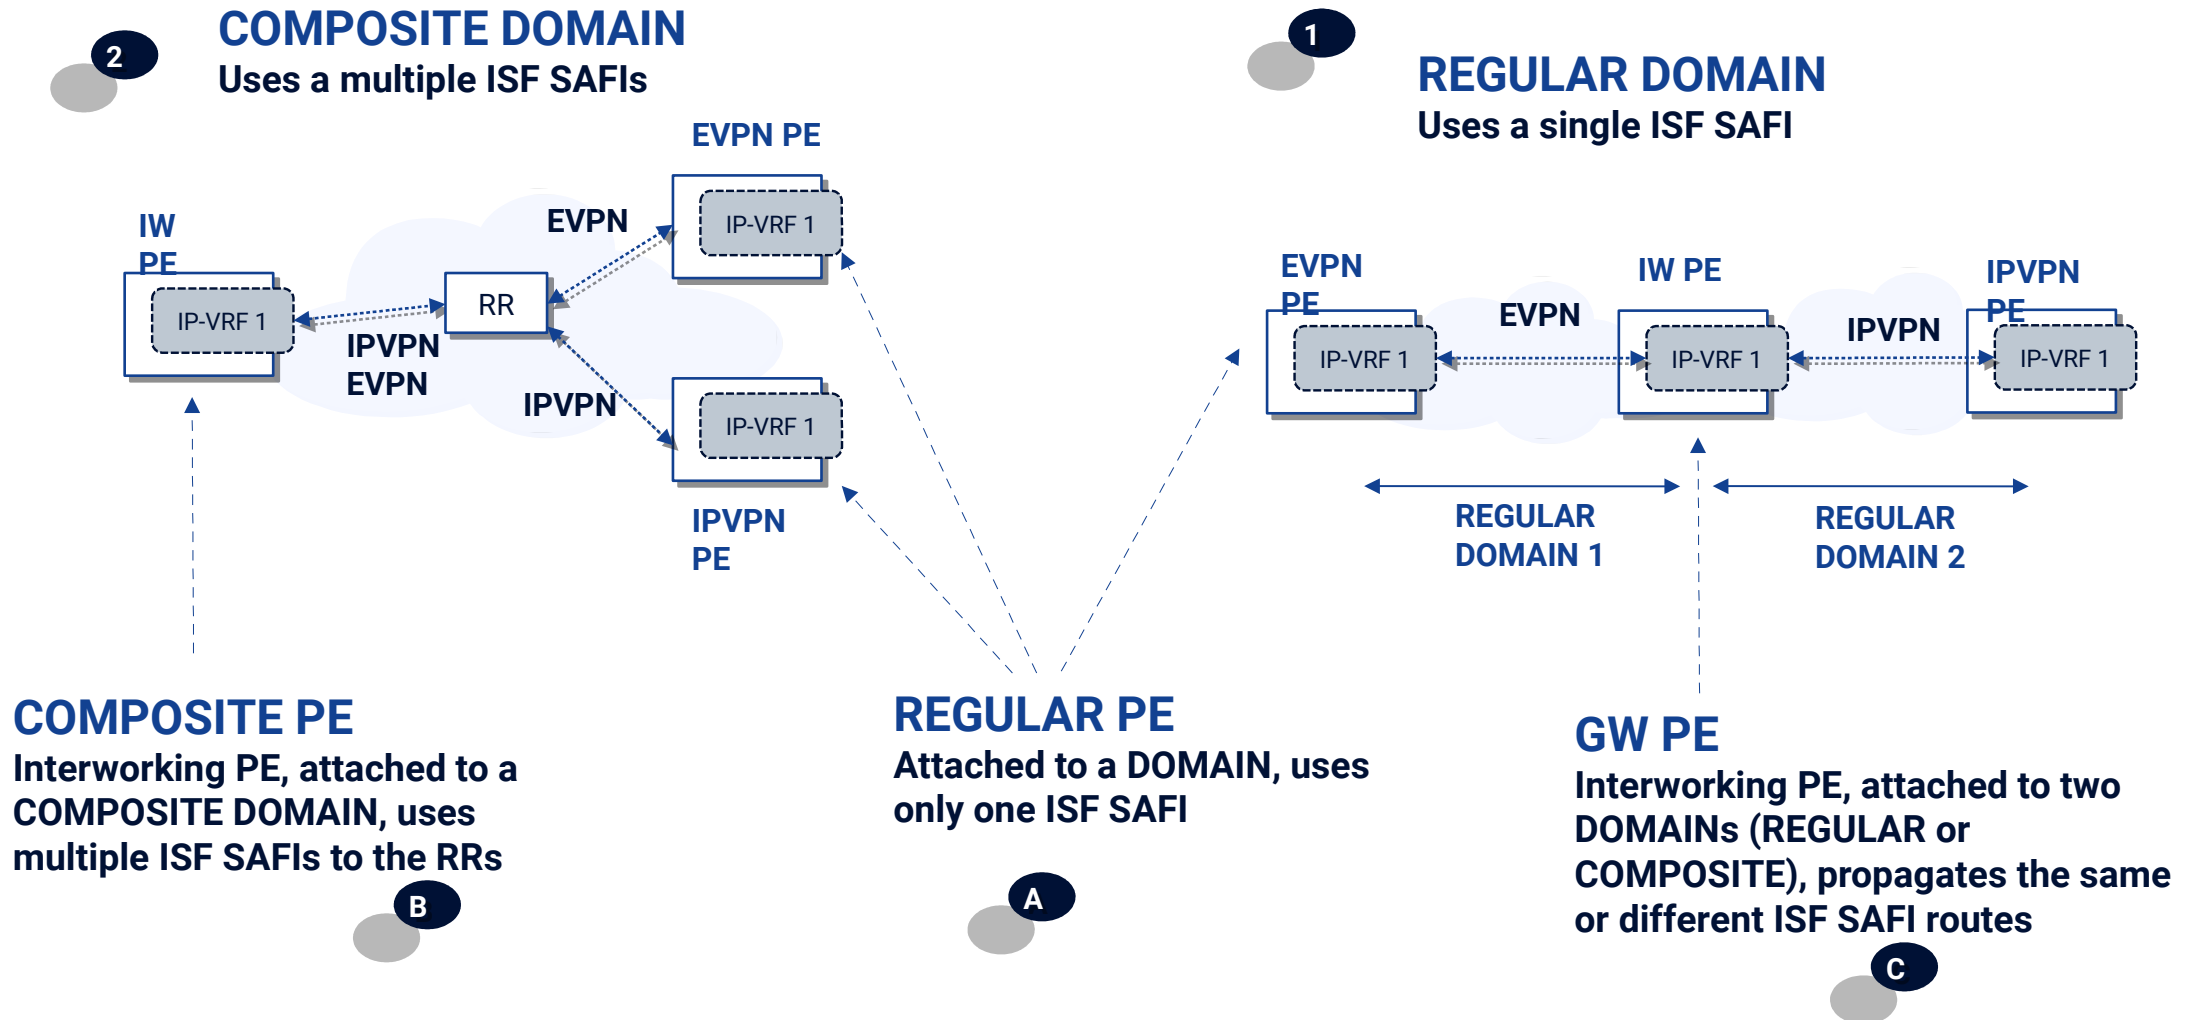
A GW is always configured with multiple Domain_IDs



DOMAIN LIMITS

Two PEs are in the same Domain if there are no (tenant) IP lookups in intermediate routers.
A GW is always configured with multiple Domain_IDs

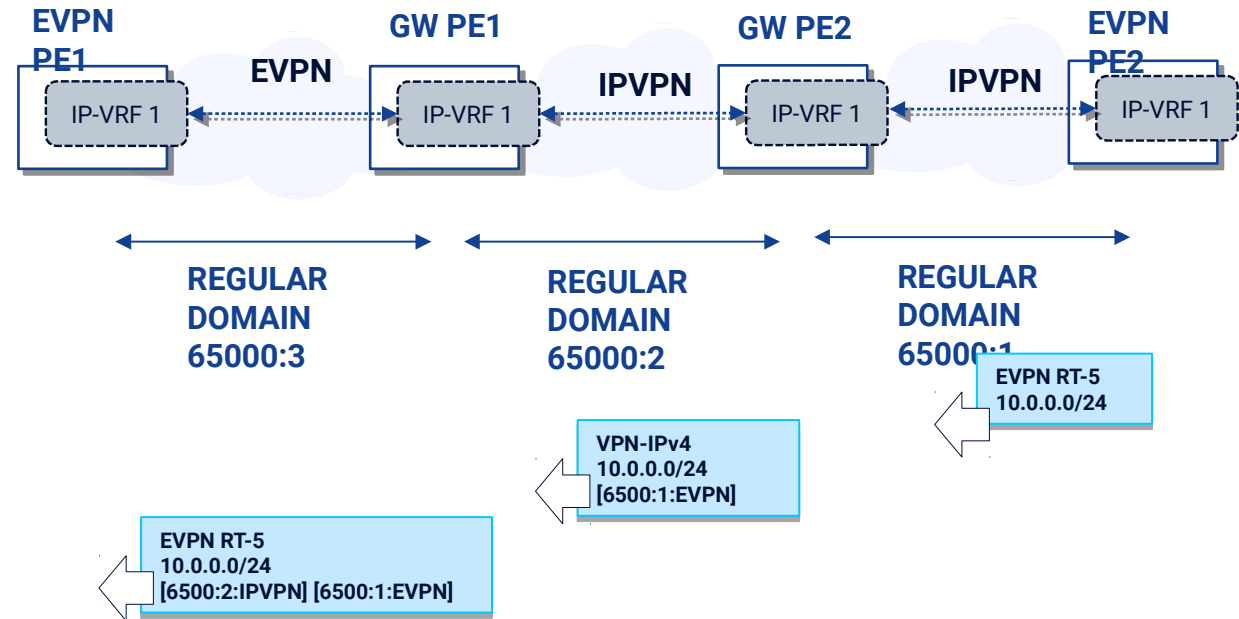
Building blocks: DOMAIN types and PE types



Domain Path Attribute (D-PATH)

A new attribute for loop protection and ISF SAFI visibility

- Length field followed by a sequence of Domain segments, where each Domain segment is represented by <DOMAIN-ID:ISF_SAFI_TYPE>
- D-PATH used for Loop Prevention: any PE that imports a Prefix route MUST drop the route if its D-PATH contains a <DOMAIN- ID:ISF_SAFI_TYPE> segment, where DOMAIN-ID matches a local DOMAIN-ID



BGP Path Propagation Across ISF-SAFIs

Two modes of operation

- No Propagation mode (default)
- Uniform Propagation mode

Uniform Propagation:

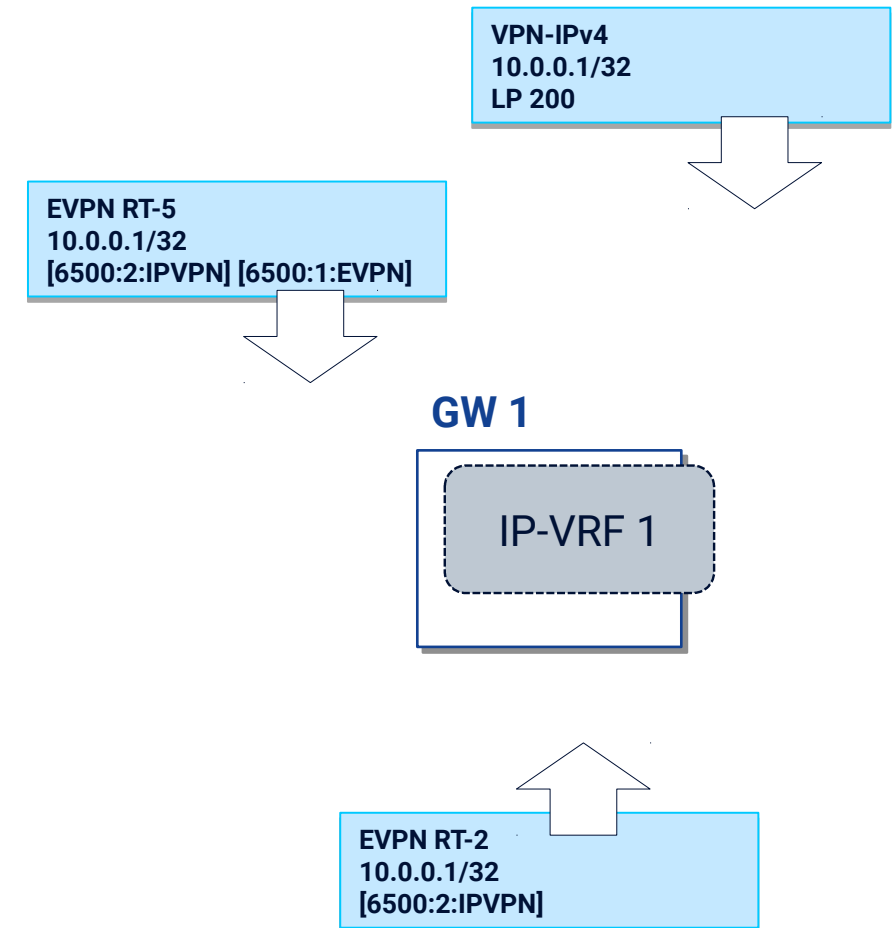
- AS_PATH
- D-PATH
- IBGP-only Path Attributes: LOCAL_PREF, ORIGINATOR_ID, CLUSTER_ID
- MED
- AIGP
- Communities, (non-EVPN) Extended Communities and Large Communities

Route Selection

Between EVPN and other ISF-SAFIs

Removes from consideration the routes following the rules and the order defined in [RFC4271], with the following exceptions and in the following order:

1. After removing from consideration lowest Local Preference, any routes that do not have the shortest D-PATH are also removed from consideration. Routes with no D-PATH are considered to have a zero-length D-PATH.
2. Then regular [RFC4271] selection criteria is followed.
3. If at least one route still under consideration is an RT-2 route, remove from consideration any RT-5 routes.
4. Steps 1-3 could possibly leave Equal Cost Multi-Path (ECMP) between IP and EVPN paths.
 - A. By default, the EVPN path is considered
 - B. If ECMP across ISF SAFIs is enabled by policy, and an "IP path" and an "EVPN path" remain at the end of step 3, both path types will be used.



EVPN and IPVPN Interworking

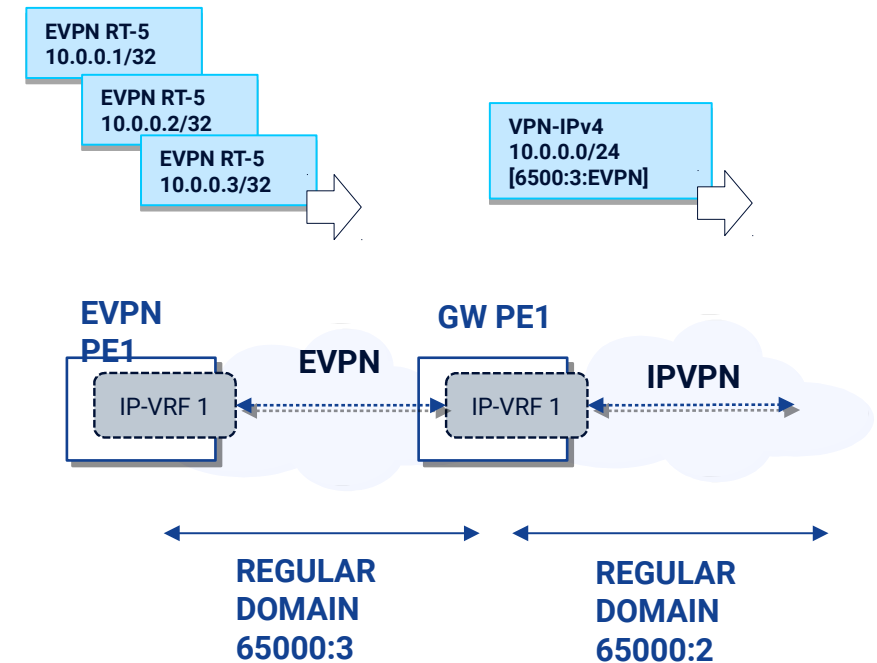
Changes in rev 01

- Terminology section clarified and general cleaning (credits to John)
- Some procedures are clarified
- Route Aggregation section added

Aggregation of Routes to minimize the impact of host routes in DCs

And Path Attribute Propagation

- Combination of the characteristics of multiple ISF routes of the same ISF SAFI in such way that a single aggregate ISF route of a different ISF SAFI can be propagated.
- Done ONLY by a gateway PE.
- If Uniform Propagation mode is used:
 - AS_PATH aggregated based on [RFC4271].
 - ISF routes that have different attributes of the following type codes MUST NOT be aggregated: D-PATH, LOCAL_PREF, ORIGINATOR_ID, CLUSTER_ID, MED or AIGP.
 - Community, Extended Community and Large Community attributes of the aggregate ISF route MUST contain all the Communities/Extended Communities/Large Communities from all of the aggregated ISF routes.



Next steps

The Authors would like to request Working Group Adoption

Thank you