draft-saum-evpn-lsp-ping-extension

Saumya Dikshit

Agenda

- Introduction
- Problem Statement and Use-Case
- Solution
- Inter-operability and Backward Compatibility

Introduction

This draft addresses the limitation(s) in current standards

- https://datatracker.ietf.org/doc/html/draft-jain-bess-evpn-lsp-ping

Requirements spelled out in the problem section are not resolved by current set of standards

Fallout of comments provided as part of reviewing “draft-jain-bess-evpn-lsp-ping”.

The authors of “draft-jain-bess-evpn-lsp-ping” advised

- new draft instead of adding the new requirements and corresponding solution in “draft-jain-bess-evpn-lsp-ping”

Requirements

Ease of usage

- EVPN NLRI key is long and complex
- Exact prefix key not top-of-mind for an operator.
- Attributes like RD, RT, ESI, ESI are required along host credentials are combined to be treated as long string index.

Validation type

- Legacy OAM pings include both control plane and data plane validation
- Routing convergence may lead to delayed or no response from destination due to a churn and/or source application may bail-out/time-out before the response arrives
- Admin access to remote device not-in-place and rest-support is questionable
  - OAM ping can be a good way to obtain Control plane data (RIB values of the protocol)

OAM reachability to liaison VRFs

- The state of VRF:
  - Working Configuration: VRF is operationally and administratively UP and WORKING
  - Network Reachability, that is, VRF is reachable via/from remote fabric devices (Vteps or LSR or LER)
- Existing OAM toolset is not armed-enough to address the following:
  - If there is no route leaked into the VRF, the hosting device MAY not form a tunnel with any device across the fabric.
  - Hence an OAM reachability check to VRF is not possible with current set of standard toolset

Requirements: Example

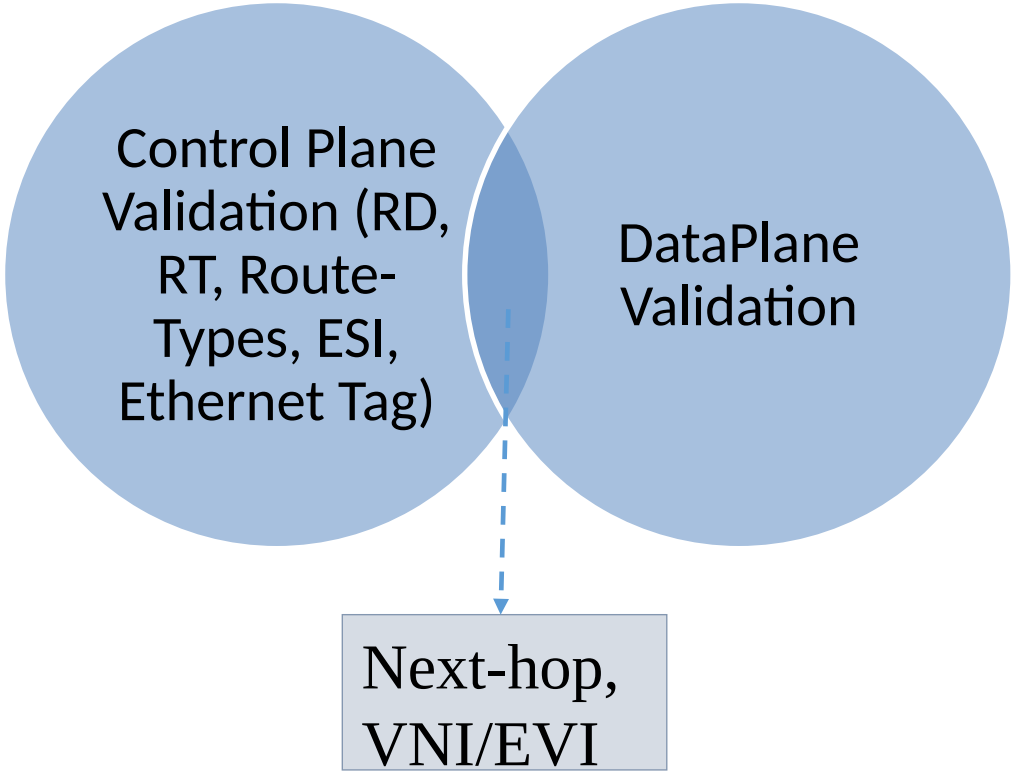

Venn diagram gives an apt description about

- the parameters which are only control plane specific and
- may not requirement validation when being asked in “Only Dataplane” mode

Requirements : Multi-fabric Topology

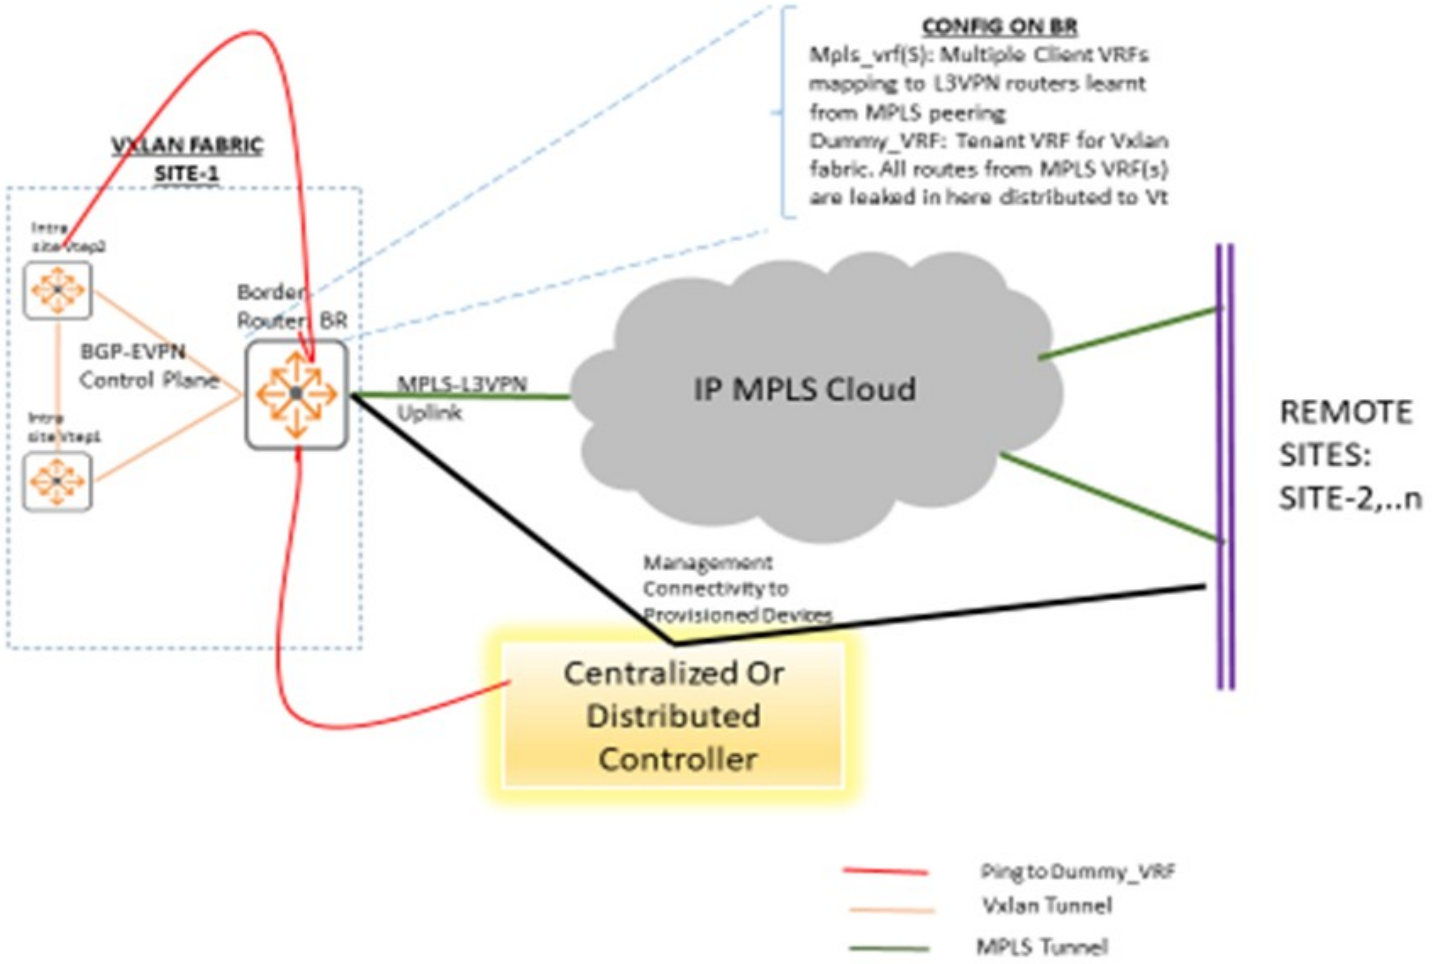

Solution Overview

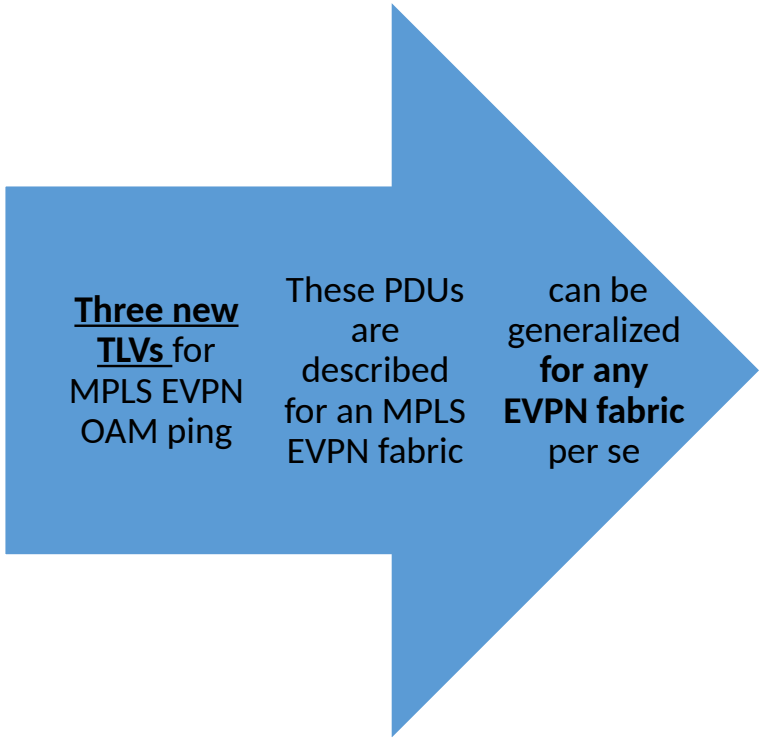

Wild Card List TLV

- Don’t care option

Validation TLV

- mode of validation (Control-plane, data-plane, both)

EVI Sub Tlv

- liaison vrf information.

Solution : Wild Card TLV

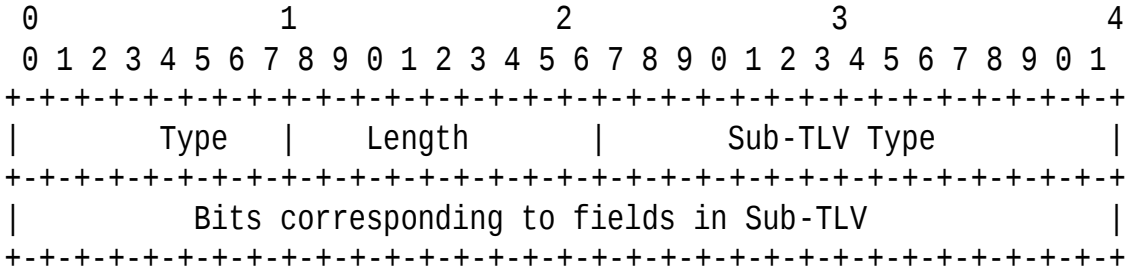

Wild-Card List TLV	
Field	**Description**
Type	Type field can be newly defined as a proprietary one.
Length	length of the TLV
Sub-TLV Type	Sub-TLV type value as defined in https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-lsp-ping-05.txt
Bitmap for fields inside Sub-TLV	The bit-map defines which field(s) in the "Sub-TLV type" is carried as wild card. The bitmap for fields is very specific to the sub-tlv. The assumption is that there are no more than 32 unique fields inside a sub-tlv. For example, in EVPN MAC Sub-TLV, https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-lsp-ping-05.txt#section-4.1, the RD is to set as wild card, then the Sub-TLV-Type carries a value 2 (defined in https://datatracker.ietf.org/doc/html/rfc7432#section-20), and bitmap has 1st bit set indicating the 1st field of the TLV is RD.

FUNCTIONALITY

- Carries the information regarding the fields (TLVs or sub TLVs),
  - That need to be ignored while processing in mpls lsp ping PDU at the OAM PDU destination

EXAMPLE

- Send Side
  - if an OAM ping to a prefix does not requires any RD (Route-Distinguisher) validation,
    - then RD value, to be carried in IP prefix TLV; can be indicated as wild-card (dont care).
- Receive Side (target device)
  - The control-plane validation of the lsp-ping should ignore the RD value in the TLV,
  - and respond back as success
    - even if there is atleast one NLRI which complies with other attributes (not set as wild card).

Solution : Validation Scope TLV

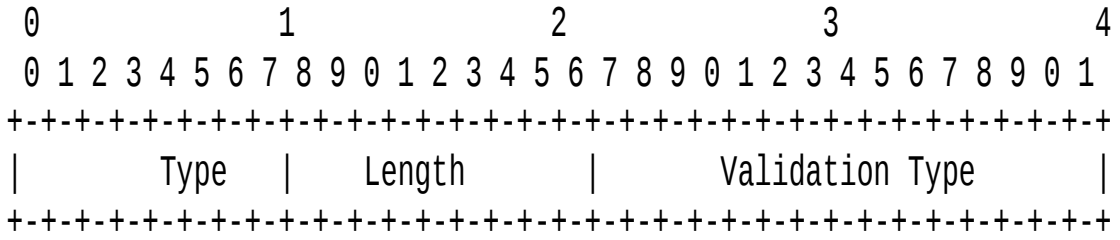

Validation TLV	
Field	**Description**
Type	Type field can be newly defined as a proprietary one.
Length	Length of the TLV
Validation type	Three values for the validation as of now: 0 - Both Control plane and Dataplane Validation (DEFAULT) 1 – Only Control plane Validation 2 – Only Data plane Validation

- Validation-type to be done for the OAM mpls ping
- Dataplane Validation:
  - FIB (forwarding information base) or routing/switching/bridging table
- Control Plane Validation:
  - Protocol RIB parameters
  - CPU intensive and can impact the control plane processing
- Both Control plane and Dataplane Validation:
  - Sanitize the network in a new-installation or post/pre upgrades
  - network is in steady state and routers/switches in contention are not experiencing protocol churns.

Solution : EVI Sub TLV

EVI sub-tlv:

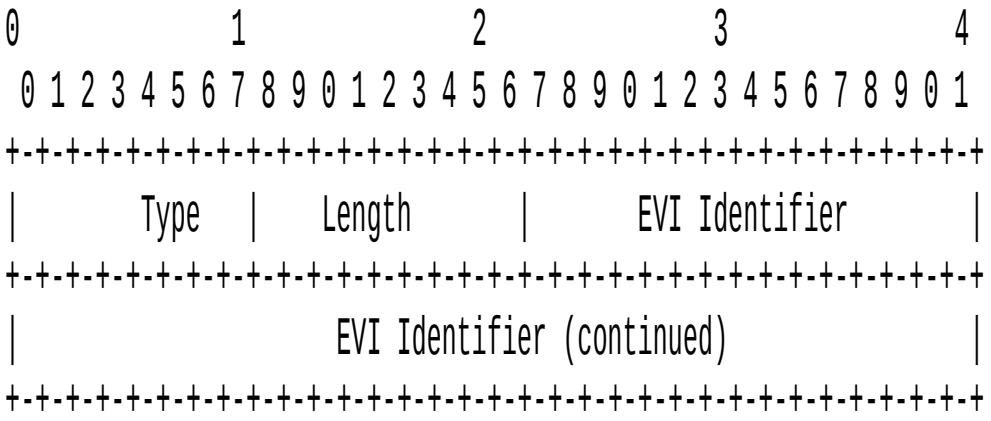

Field Description:

- *Type*: 1 octet: Type field can be newly defined as a proprietary one.
- *Length*: 1 octet: Defines the length of the Value field
- *Value*: EVI identifier and depending on the length field being carried. The EVI can be an MPLS label or VNI in case of Vxlan.

EVI (Virtual Network Identifier) information, thus ensuring that following tools and/or action-sets can be supported:

- Ping or path tracing to check the configuration of an EVI on a remote Vtep
- Ping to check VRF configuration (mapped to an EVI) on remote Vtep,
  - even though no layer-3 configuration is enable against that VRF
- Ping to check VRF configuration (mapped to an EVI) on remote Vtep,
  - For which EVPN tunnel not been provisioned yet.

Further Actions Requested

- Requesting for WG Adoption

Or

- Merging into existing mpls-evpn-lsp-ping draft