draft-saumthimma-evpn-ip-binding-sync

Saumya Dikshit

Agenda

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

Introduction

In Campus networks of Colleges, Branch office, Headquarters and Data Center networks,

- network extend Layer-2 across sites and geographies over WAN, Wide Area Network, via well defined overlay fabric (including EVPN)

'Client IP Binding database' is significant in last-mile

- Access Switches, build trusted database of L3 clients by snooping into DHCP/ND packets in the client VLAN.
- Essential for securing a Campus network

Enterprise networks are extended across geographies

- Synchronization required across sites/fabrics and/or across geographies

For Overlay interconnects:

- BGP control plane constructs can be leveraged
- BGP is a typical control plane protocol configured to communicate across network boundries.¶

Requirements

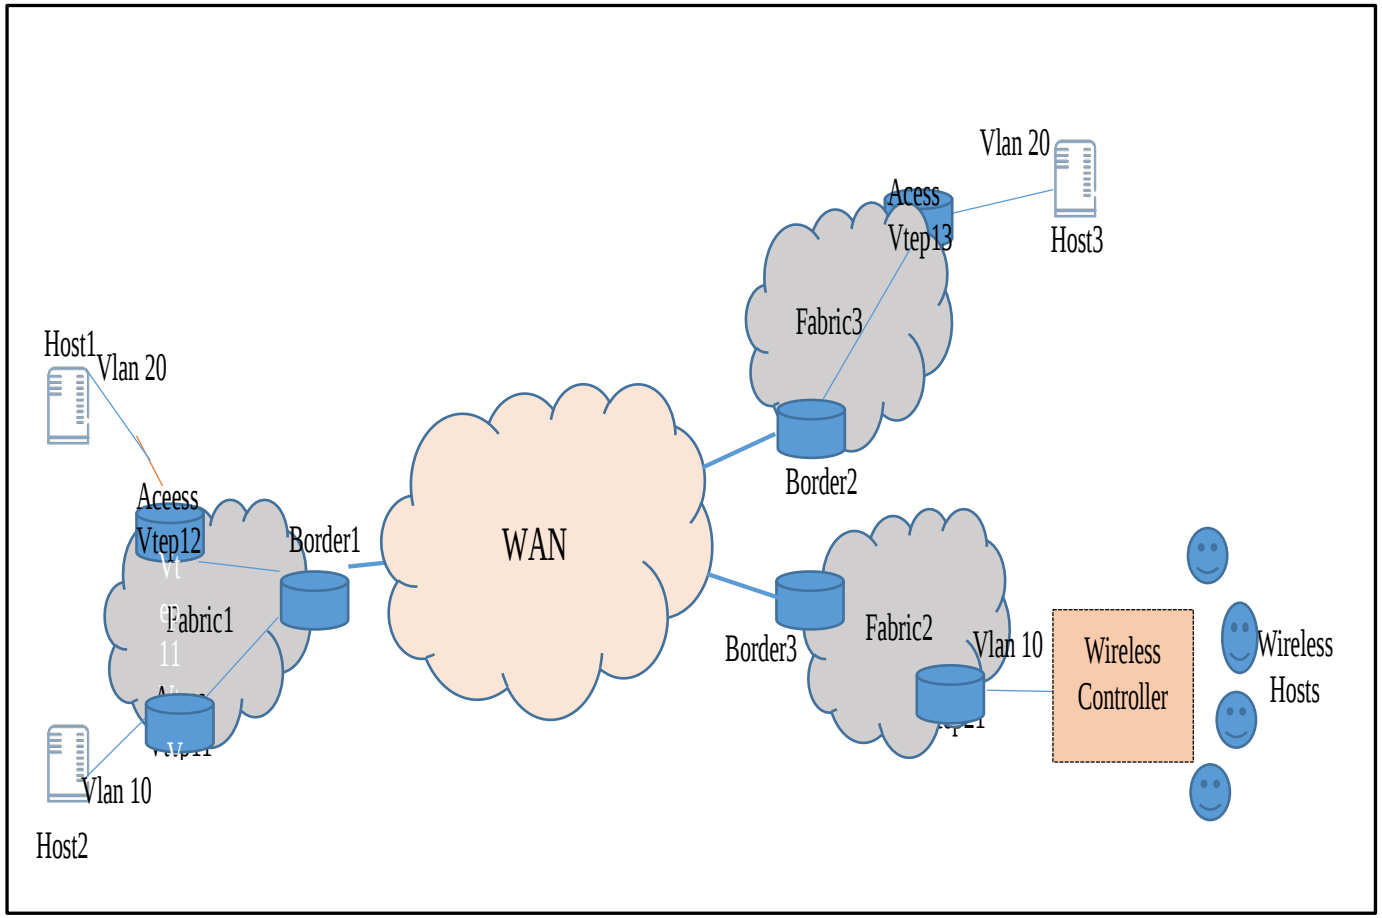

- Client IP Binding Database" is local to client connected to the switches within the fabric.
- Without explicit solution in place, these details are not known to other switches within and outside the fabric, unless and until explicitly communicated.
  - the Access-Vtep11 and Access-Vtep12, though in the same fabric, can snoop in to DHCP packets originating/destined from the locally attached hosts, i.e., Host1 and Host2 respectively,
    - but not the other Access's hosts. These packets are typically unicasted to/from the DHCP server located in a centralized location
  - The **packets are not leaked to remote fabrics** as well.
    - The DHCP packets generated from Host1 or Host2 will not be leaked to remote fabrics (Fabric2 or Fabric3)
    - hence cannot be snooped in by remote Access devices Access-Vtep12 and Access-Vtep13.

Problem and Requirements

Client IP Binding Database" is local to client connected to the switches within the fabric.	Example (Diagram Reference)	Data packets **NOTleaked to remote fabrics**
• Without explicit solution in place • CLIENT details are not known to other switches within and outside the fabric, unless and until explicitly communicated.	• Access-Vtep11 and Access-Vtep12, though in the same fabric, can snoop in to DHCP packets originating/destined from the locally attached hosts, i.e., Host1 and Host2 respectively, • but not the other Access's hosts. These packets are typically unicasted to/from the DHCP server located in a centralized location	• The DHCP packets generated from Host1 or Host2 will not be leaked to remote fabrics (Fabric2 or Fabric3) • hence cannot be snooped in by remote Access devices Access-Vtep12 and Access-Vtep13.

Proposed Solution cont…

This document proposes a solution for synchronizing the IP-Bindings between the Access-Switches across EVPN provisioned overlays:

- By leveraging the BGP (EVPN) control plane constructs called Route-Type-2 NLRIs (EVPN NLRI: MAC Advertisement Route).
- A new extended community, indicating
  - the IP-Binding-synchronization,
  - to be carried, in BGP-Update message carrying the Route-Type-2 NLRI (Network Layer Reachability Information).
- As EVPN is the at the core of fabric deployments to support multi-homing and multi-tenancy,
  - This control plane extension alleviates the Synchronization of IPs which is specific to tenants and
  - applicable to all/selective hosts attached to the Access-Vtep. T
  - The host can be attached to more than one Vtep (indicating multi-homing) over a segment (called EVI). This Bullet is specific to BGP-EVPN Control plane.
- EVPN is a goto solution for Vxlan fabric extension across the WAN
  - as its also easy to realize the native-transport (underlay) encapsulation as IP based.

Proposed Solution cont…

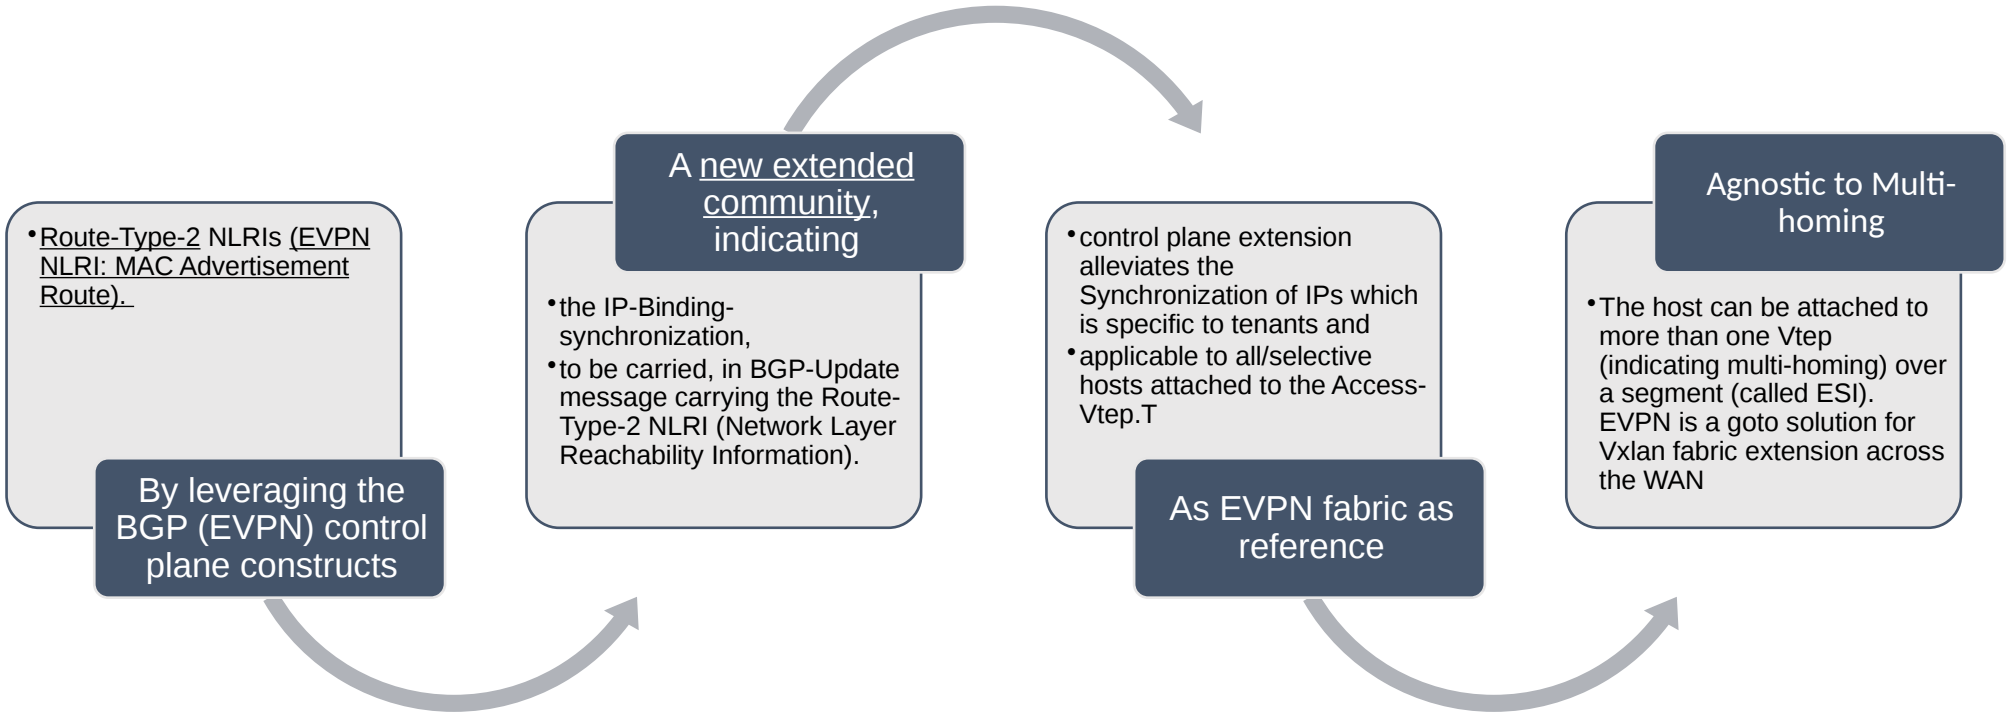

Proposed Solution

FORMAT: “Client IP Binding Sync” Extended Communities

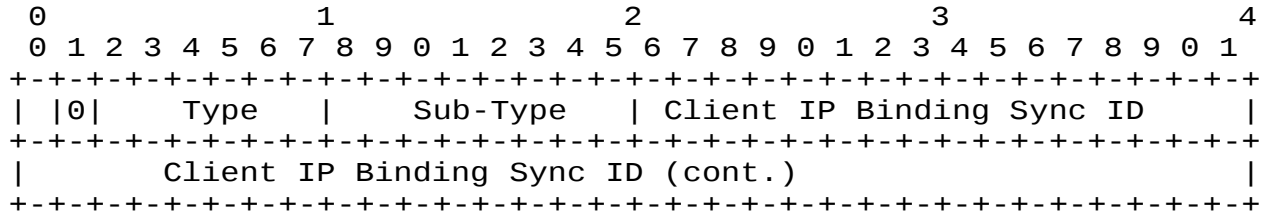

Client IP Binding Sync Extended Communities		
Field	**Size**	**Description**
Type	1 Octet	Type field can be newly defined as a proprietary one. It’s a transitive Attribute, hence 1-bit in this octet is “0”.
Sub-Type	1 Octet	The value of sub-type will be allocated as per https://datatracker.ietf.org/doc/html/rfc4360#section-7 For experimental purposes, it can be used as any unallocated value within the range which can be defined as per: https://datatracker.ietf.org/doc/html/rfc7153#section-5.2 This field when set to appropriate value indicates to the receiver about processing the NLRI for IP-Binding synchronization.
Client IP Binding Sync ID	6 Octets	This field carries the Identifier configured to categorize the IP-Binding instance. The field is applicable to all data carried in MAC/IP NLRIs in the same update message. The instance is configurable and value is implementation dependent.

“Client-IP Binding Sync” Extended Communities is defined.
optional attribute and also transitive

- Encapsulated along with BGP-update message with “MAC/IP” bindings (with “EVPN NLRI: MAC Advertisement Route”/“Route-Type-2”), indicates the following to the receiving BGP-peer:
- MAC/IP data to be leveraged for Client IP Binding synchronization.
- The data is to be handed over to the Security entity for validating the IP allocation
- The handover procedure is implementation specific and outside the purview of this invention.

Applicability

Generic solution for all Fabrics

- **for any BGP provisioned overlay network like VPLS, VPWS**, L3-VPN etc.
- The applicability is for various fabrics, like Vxlan, MPLS, NV-GRE, GPE etc

leveraged for **other overlay provisioning protocols** like LDP (Label Distribution Protocol),

- by aligning the Path attribute information in context of target protocol.

Further Actions and Requests

- Review request by BESS WG member.