

Virtual Hub & Spoke with BGP EVPNs

draft-keyupate-evpn-virtual-hub-00

Keyur Patel, Ali Sajassi, John Drake, Wim Henderickx

IETF 94, November 2015, Yokohama, Japan

Motivation

- EVPN is becoming a pervasive technology for overlay networks in MSDCs & SPDCs
- EVPN route scale in DCs is typically achieved by
 - Leaf nodes learn and store local DC routes (MAC and IP addresses)
 - Any remote DC routes are stored on Gateway nodes
 - Leaf nodes are installed with default route pointing towards Gateway nodes
- Draft further optimizes route learning such that
 - Leaf nodes learn and store routes (MAC and IP addresses) for local sites only
 - Leaf nodes do not store routes advertised by other Leaf nodes residing in the same DC

Solution

- RFC7024 defines rules for Virtual Hub & Spoke for BGP L3VPNs
- Draft extends RFC7024 for EVPN address family
 - Rules for generating and processing of unknown MAC Route
 - Modifications to Aliasing, Split Horizon and Mass Withdraws
 - Forwarding considerations for Bridging mode and IRB mode
 - Rules for ARP Suppression

Solution – Rules for Unknown MAC Route

- Unknown MAC Route is equivalent to VPN-IP default route in RFC7024
- V-Hub exports Unknown MAC Route for a given EVI for set of its V-Spokes
 - RD makes route unique per V-Hub
- Rules to originate the Unknown MAC Route is same as VPN-IP default route in RFC7024
 - Unknown MAC Route is announced using MAC/IP advertisement route
- V-Spoke imports the route unless route-target mismatch happens
- Rules to import the Unknown MAC Route is same as VPN-IP default route in RFC7024
 - In absence of more specific EVPN-MAC route, unknown MAC Route is used in MAC-Table/ARP/ND Cache
 - Forces the traffic to be forwarded to V-HUB from a V-Spoke

Solution – Modifications to Aliasing, Split Horizon and Mass Withdrawals

- V-Spokes usually do NOT have EVPN MAC Routes from other V-Spokes
 - Ethernet A-D routes of V-Spokes are only announced to V-HUB
 - V-Hub only announces Ethernet A-D per EVI route when it announces at least 1 or more specific MAC/IP route to a V-Spoke. Otherwise there is no point in propagating Ethernet A-D per EVI routes between spokes by default.
 - Same rule applies to Ethernet A-D per ES routes

Solution – Forwarding Considerations

- IP Forwarding
 - Rules of RFC7024 apply for IP only forwarding
- MAC only forwarding (Bridging Mode)
 - MPLS label of Unknown MAC route identifies the MAC-VRF of V-Hub
 - V-Hub receiving data packets for such a label pops the label and performs a lookup in MAC-VRF to determine further disposition of the packet
 - Data packets for more specific routes are forwarded towards their nexthops
- IRB
 - If the host destination address of the packet is an IRB interface then V-Spoke does an IP lookup in MAC-VRF
 - If the IP address is not found then V-Spoke MAY forward the packet to V-HUB using Default VPN-IP Route
 - If the host destination address of the packet is a MAC address then V-Spoke does an MAC lookup in MAC-VRF
 - If the MAC address is not found then V-Spoke MAY forward the packet to V-HUB using unknown MAC Route

Solution – ARP Suppression

- RFC7432 defines procedures to terminate ARP/ND request, translate it into BGP updates and then use the information to generate ARP/ND responses (AKA ARP Suppression)
- ARP suppression needs to be performed by V-Hub and V-Spoke:
 - V-Spoke terminates gratuitous ARP/ND request, stores the source MAC/IP pair in ARP/ND Cache and announce it in EVPN SAFI to V-Hubs (if created for the first time)
 - V-Spoke responds to the received ARP/ND request if it find a match in its local ARP/ND cache. Otherwise the request is unicasted towards one of V-Hubs
 - V-Hub receives these routes from V-Spoke and stores it in its ARP/ND Cache table.
 - V-Hub responds to any received unicast ARP/ND requests. Optionally V-Hub MAY announce the matching route in EVPN SAFI as well (as a matter of policy)

Questions?