

Multicast VPN using BIER

IETF 91, Honolulu

<http://www.ietf.org/id/draft-rosen-l3vpn-mvpn-bier-01.txt>

Eric Rosen

Juniper Networks

Mahesh Sivakumar

Cisco Systems

Ijsbrand Wijnands

Cisco Systems

Sam K Aldrin

Huawei Technologies

Andrew Dolganow

Alcatel-Lucent

Tony Przygienda

Ericsson

Outline

- Introduction
- P-tunnels and PMSI
- Explicit Tracking
- Data Plane

Introduction

- **Multicast Virtual Private Network (MVPN)** carries multicast traffic over a Service Provider (SP) backbone network requires multicast tunnels to carry them
- **Bit Index Explicit Replication (BIER)** is an architecture to provide optimal multicast forwarding through a “multicast domain”
- This draft specifies the protocol and procedures to allow MVPN to use BIER to carry multicast traffic over an SP backbone network

Outline

- Overview
- P-tunnels and PMSI
- Explicit Tracking
- Data Plane

P-tunnels

- [RFC-6513] and [RFC-6514] specify protocols and procedures for carrying customer multicast traffic through an SP - backbone network
 - Multicast provider tunnels (**P-tunnels**) are created to carry multicast traffic
 - Specification allows for several P-tunnel technologies, including
 - MLDP (Multicast Label Distribution Protocol)
 - P2MP-TE (Point-to-MultiPoint Traffic Engineering)
 - IR (Ingress Replication)
 - **BIER** now added as an additional P-tunnel technology

Terminology


- **BFR** – Bit-Forwarding Router
- **BFIR/BFER** – Bit-Forwarding Ingress/Egress Router
- **MVPN** – Multicast Virtual Private Network
- **P-Tunnel** – tunnel through the SP network; transport multicast data
- **C-S, C-G** – multicast source, group used by a VPN customer
- **C-Flow** – A customer multicast flow
- **I-PMSI AD route** – Inclusive Provider Multicast Service Interface Auto Discovery route, carried in a BGP update message to advertise the default MVPN p-tunnel
- **S-PMSI AD route** – Selective PMSI AD route; carried in BGP update messages to bind c-flows to particular p-tunnels
- **PTA** – PMSI Tunnel Attribute; BGP attribute carried to identify a particular p-tunnel. Also carries information to multiplex/de-multiplex c-flows when traffic from multiple VPNs are carried in a single p-tunnel

BIER P-tunnel

- No explicit multicast tunnel building with BIER
 - Can however be modeled as an implicit P2MP tunnel through a **BIER domain** (from a **BFIR** to all the **BFERs**)
- BIER “**tunnel**” not specific to any VPN
 - Aggregates traffic from multiple MVPNs
 - Control plane will leverage existing procedures
- BIER carries traffic within one BIER domain
 - MVPN allows P-tunnels to be “**segmented**” at “**border routers**” (such as ABRs or ASBRs)
 - BIER domain will correspond to P-tunnel segment
 - BIER domain coextensive with IGP network or area
 - **Area Border Routers (ABRs)** and/or **Autonomous System Border Routers (ASBRs)** need to be capable of acting as BFIRs and BFERs

PMSI Tunnel Attribute (PTA)

- Identifies the P-tunnel to which one or more flows are bound



- Tunnel Type**
 - New tunnel type for “BIER”
- Tunnel Identifier**
 - BFR-Prefix of the originator of the route carrying the PTA
- Leaf-Info-Required bit** (in Flags octet)
 - Set in S-PMSI A-D route and clear in I-PMSI A-D route

Upstream Assigned Label

- **MPLS Label**

- Upstream-assigned by the router originating the PMSI route to which PTA is attached
- MPLS label in two x-PMSIs ($x \rightarrow I$ or S) originated by an ingress PE for BIER **MUST** be different if
 - They carry a different set of Route Targets (RTs)
 - Ingress PE supports extranet and routes are from different VRFs
 - Ingress PE supports “Extranet separation” and only one of the routes carry the “Extranet Separation” EC
- Thus at a given egress PE, two packets with same upstream-assigned label will be *from* the same ingress VRF (of ingress PE), and destined *to* the same set of egress VRFs on the egress PE

Outline

- Overview
- P-tunnels and PMSI
- **Explicit Tracking**
- Data Plane

Explicit Tracking

- For BFIRs to determine the set of BFERs to which packets are to be delivered, explicit tracking, as specified in [RFC6513] and [RFC6514] will be used
 - BFIR originates S-PMSI A-D route (for each C-flow) with LIR bit set in the PTA
 - Per [RFC6513], BFERs having interested receivers for the C-flow will respond with a Leaf A-D route
 - The BFIR will match the received Leaf A-Ds to the originated S-PMSI to determine the receivers
 - If a matching I or (*,*) S-PMSI exists for the C-flow, explicit tracking will be done via S-PMSI for the C-flow which includes a PTA specifying “no tunnel type”

Outline

- Overview
- P-tunnels and PMSI
- Explicit Tracking
- Data Plane

Transmitting an MVPN data packet

- Ingress PE finds matching x-PMSI route (following rules of [RFC6625]) with PTA specifying “BIER”
- MPLS label from matching route is pushed on to the packet label stack and forwarded according to procedures of [BIER_ARCH] and [BIER_ENCAPS]
- Since the payload is an MPLS packet with upstream assigned label, the **I bit** is set and the **BFIR-ID** MUST be included in the BIER header

Receiving an MVPN data packet

- A BFER receiving an bier-encapsulated MVPN data packet sends the below information to the multicast flow layer
 - the “BFR-prefix” corresponding to the BIER-ID in the packet BIER-header
 - the “payload”, an MPLS packet with an upstream assigned top label
- The BFR-prefix provides the context in which the upstream assigned label is interpreted

Questions?

References

[BIER_ARCH]

Wijnands, IJ., “Multicast using Bit Index Explicit Replication Architecture”, internet-draft, draft-wijnands-bier-architecture-00, September 2014.

[BIER_ENCAPS]

Wijnands, IJ., “Multicast encapsulation using Bit Index Explicit Replication Architecture”, internet-draft, draft-wijnands-mpls-bier-encapsulation-00, September 2014.

[RFC6513]

Rosen, E. and R. Aggarwal, “Multicast in MPLS/BGP IP VPNs”, RFC 6513, February 2012

[RFC6514]

Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, “BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs”, RFC 6514, February 2012

[RFC6625]

Rosen, E., Rekhter, Y., Hendrickx, W., and R. Qiu, “Wildcards in Multicast VPN Auto-Discovery Routes”, RFC 6625, May 2012.