

# **BGP Extensions for BIER**

## **draft-xu-idr-bier-extensions**

**Xiaohu Xu (Huawei)**

**Mach Chen (Huawei)**

**Keyur Patel (Cisco)**

**IJsbrand Wijnands (Cisco)**

**Tony Przygienda (Ericsson)**


**IETF93, Prague**

# Motivation

- **BIER is applicable in multi-tenant data center networks for efficient delivery of BUM traffic while eliminating the need for maintaining multicast states in the underlay[I-D.ietf-bier-use-cases].**
- **Since BGP has been used as a underlay routing protocol in many large multi-tenant data center networks [I-D.ietf-rtgwg-bgp-routing-large-dc], it seems necessary to extend BGP for advertising the BIER-specific information.**
  - **A new optional, transitive path attribute, referred to as the BIER attribute, can be attached to a BGP UPDATE message by the originator so as to indicate the BIER-specific information of a particular BFR which is identified by the /32 or /128 address prefix contained in the NLRI. In other words, if the BIER path attribute is present, the NLRI is treated by BIER as a "BFR-prefix".**

# BIER Path Attribute

- The attribute type code for the BIER Attribute is TBD. The value field of the BIER Attribute contains one or more BIER TLVs .



- MP ~ ER TLV  
enc | Sub-TLVs |

Figure 1:BIER TLV

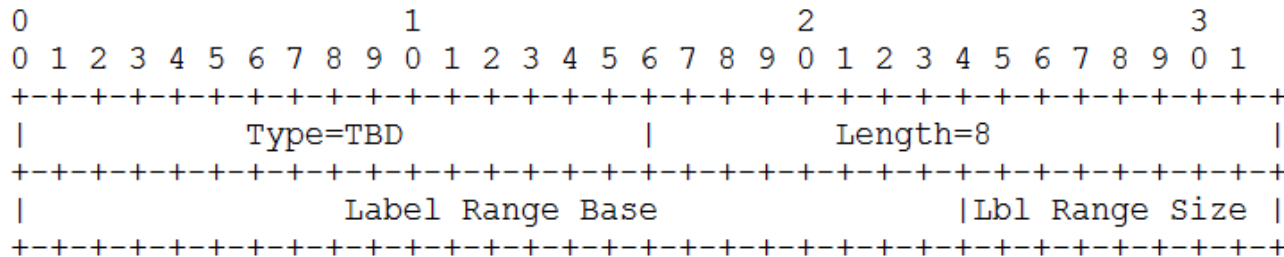


Figure 2:BIER MPLS Encapsulation sub-TLV

# Originating BIER Attribute

- An implementation that supports the BIER attribute **MUST** support a policy to enable or disable the creation of the BIER attribute and its attachment to specific BGP routes.
- An implementation **MAY** disable the creation of the BIER attribute unless explicitly configured to do so otherwise.
- A BGP speaker **MUST** only attach the locally created BIER attribute to a BGP UPDATE message in which at least one of its BFR-prefixes is contained in the NLRI.

# Restrictions on Sending/Receiving

- **An implementation that supports the BIER attribute MUST support a per-EBGP-session policy, that indicates whether the attribute is enabled or disabled for use on that session.**
- **The BIER attribute MUST NOT be sent on any EBGP peers for which the session policy is not configured.**
  - **If an BIER attribute is received on a BGP session for which session policy is not configured, then the received attribute MUST be treated exactly as if it were an unrecognised non-transitive attribute. That is, “it MUST be quietly ignored and not passed along to other BGP peers“.**
- **To prevent the BIER attribute from “leaking out” of an BIER domain, each BGP router on the BIER domain MUST support an outbound route announcement policy. Such a policy MUST be disabled on each EBGP session by default unless explicitly configured.**

# Deployment Considerations

- It's assumed by this document that the BIER domain is aligned with the Administrative Domain (AD) which are composed of multiple ASes (either private or public ASes).
  - Use of the BIER attribute in other scenarios is outside the scope of this document.
- Since the BIER attribute is an optional, transitive path attribute, a non-BFR BGP speakers could still advertise the received route with a BIER attribute.
  - This is desirable in the incremental deployment scenario where a BGP speaker could tunnel a BIER packet or the payload of a BIER packet to a BFER directly if the BGP next-hop of the route for that BFER is a non-BFR.
- A BGP speaker is allowed to tunnel a BIER packet to the BGP next-hop if these two BFR-capable BGP neighbors are not directly connected (e.g., multi-hop EBGP) .

# Next Steps

- **WG adoption?**