

BGP Extensions for BIER

draft-xu-idr-bier-extensions-01

Xiaohu Xu (Huawei)

Mach Chen (Huawei)

Keyur Patel (Cisco)

IJsbrand Wijnands (Cisco)

Tony Przygienda (Ericsson)

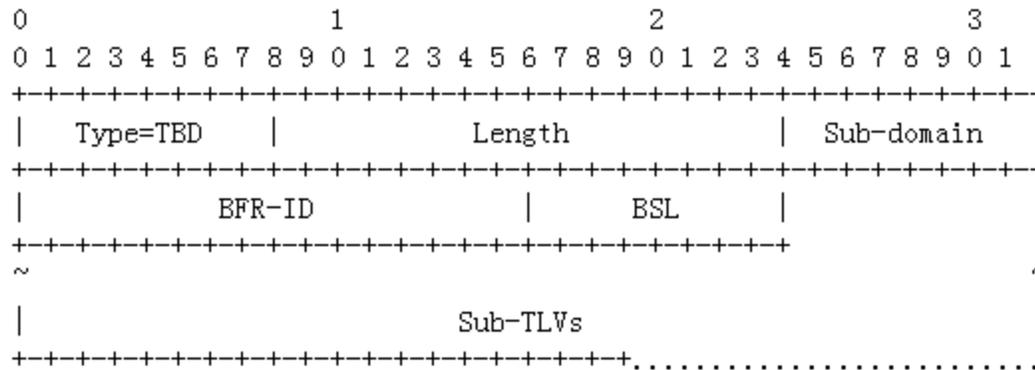
IETF92, Dallas

Motivation

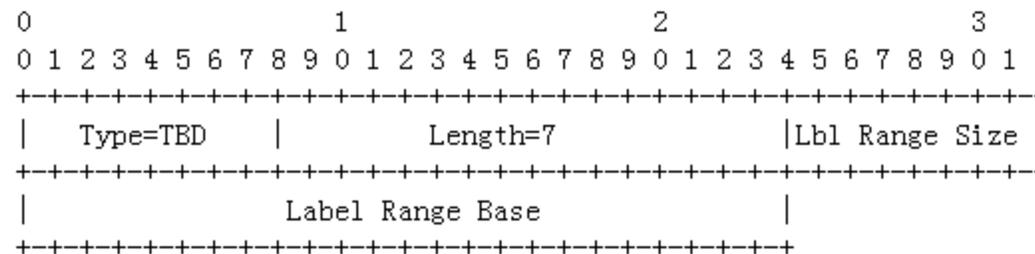
- **Bit Index Explicit Replication (BIER) is applicable in multi-tenant data center network environments for efficient delivery of BUM traffic while eliminating the need for maintaining multicast states in the underlay[I-D.kumar-bier-use-cases].**
- **BGP instead of IGP is used as an underlay in some large multi-tenant data center network environments [I-D.ietf-rtgwg-bgp-routing-large-dc].**
- **This document describes BGP extensions 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.**

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 TLV as shown below:



- MPLS-BIER Encapsulation sub-TLV is a sub-TLV of the BIER TLV encoding the MPLS-BIER specific information.



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 routable addresses (e.g., a loopback address) is contained in the NLRI.
 - The routable address contained in the NLRI is **RECOMMENDED** to be the one used for establishing BGP sessions.

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

- **Comments?**