

Loc-RIB and Adj-RIB-Out Extensions to BMP

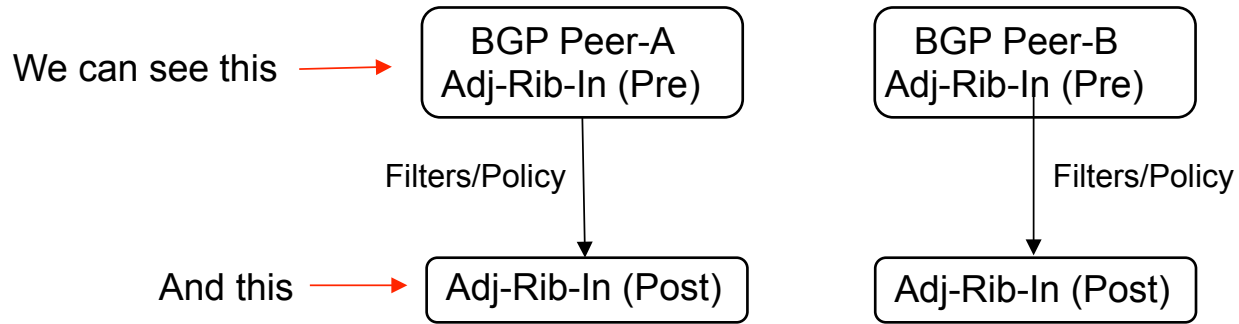
Serpil Bayraktar/Tim Evens Cisco, Paolo Lucente, NTT

March 19, 2018

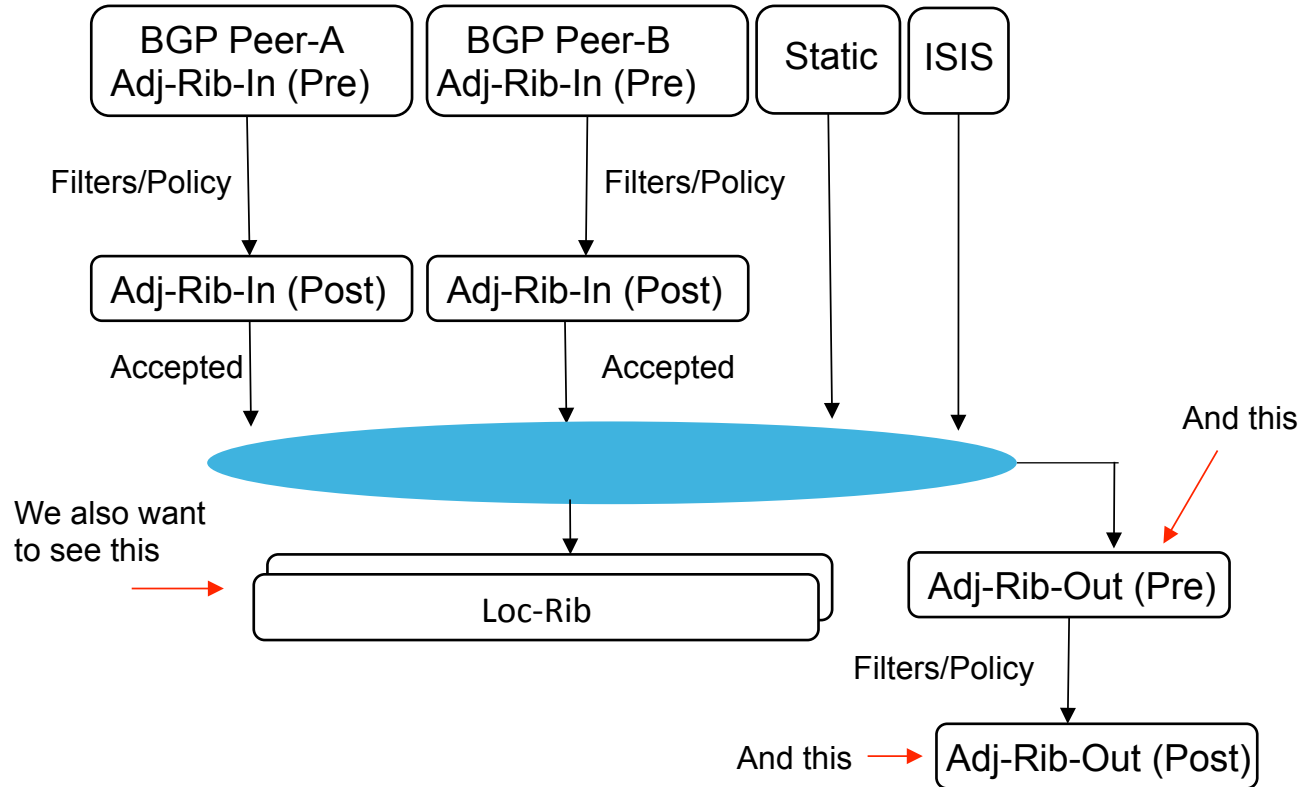
IETF, GROW Working Group Presentation

Current State

- RFC 7854 defines access to Adj-RIB-In only



Proposal: Extend access to Adj-RIB-Out and Loc-RIB



draft-ietf-grow-bmp-loc-rib: 00 -> 01 (1/3)

4.1. Peer Type

A new peer type is defined for Loc-RIB to distinguish that it represents Loc-RIB with or without RD and local instances. Section 4.2 [RFC7854] defines a Local Instance Peer type, which is for the case of non-RD peers that have an instance identifier.

This document defines the following new peer type:

- o Peer Type = TBD: Loc-RIB Instance Peer

Clarification only

draft-ietf-grow-bmp-loc-rib: 00 -> 01 (2/3)

5. Loc-RIB Monitoring

Loc-RIB contains all routes from BGP peers as well as any and all routes redistributed or otherwise locally originated. In this context, only the BGP instance Loc-RIB is included. Routes from other routing protocols that have not been redistributed, originated by or into BGP, or received via Adj-RIB-In are not considered.

Loc-RIB in this context does not attempt to maintain a pre-policy and post-policy representation. Loc-RIB is the selected and used routes, which is equivalent to post-policy.

For example, VRF "Blue" imports several targets but filters out specific routes. The end result of VRF "Blue" Loc-RIB is conveyed. Even though the import is filtered, the result is complete for VRF "Blue" Loc-RIB. The F flag is not set in this case since the Loc-RIB is complete and not filtered to the BMP receiver.

Clarification only

draft-ietf-grow-bmp-loc-rib: 00 -> 01 (3/3)

5.3. Peer Down Notification

Peer down notification SHOULD follow the section 4.9 [RFC7854] reason 2.

The VRF/Table Name informational TLV SHOULD be included if it was in the Peer UP.

Trying to reduce state

draft-ietf-grow-bmp-adj-rib-out: 00 -> 01 (1/3)

5.1. Post-Policy

The primary use-case in monitoring Adj-RIB-Out is to monitor the updates transmitted to the BGP peer after outbound policy has been applied. These updates reflect the result after modifications and filters have been applied (e.g. Adj-RIB-Out Post-Policy). **Some attributes are set when the BGP message is transmitted, such as next-hop. Adj-RIB-Out Post-Policy MUST convey what is actually transmitted to the peer, next-hop and any attribute set during transmission should also be set and transmitted to the BMP receiver.**

The L flag MUST be set to 1 ~~in this case~~ to indicate post-policy.

Clarification only

draft-ietf-grow-bmp-adj-rib-out: 00 -> 01 (2/3)

5.2. Pre-Policy

As with Adj-RIB-In policy validation, there are use-cases that pre-policy Adj-RIB-Out is used to validate and audit outbound policies. For example, a comparison between pre-policy and post-policy can be used to validate the outbound policy.

Depending on BGP peering session type (IBGP, IBGP route reflector client, EBGP) the candidate routes that make up the Pre-Policy Adj-RIB-Out do not contain all local-rib routes. Pre-Policy Adj-RIB-Out conveys only routes that are available based on the peering type. Post-Policy represents the filtered/changed routes from the available routes.

Some attributes are set only during transmission of the BGP message, e.g. Post-Policy. It is common that next-hop may be null, loopback, or similar during this phase. All mandatory attributes, such as next-hop, MUST be either ZERO or have an empty length if they are unknown at the Pre-Policy phase. The BMP receiver will treat zero or empty mandatory attributes as self originated.

The L flag MUST be set to 0 ~~in~~
~~this case~~ to indicate pre-policy.

Clarification only

draft-ietf-grow-bmp-adj-rib-out: 00 -> 01 (3/3)

6.3.1. Peer Up Information

The following peer UP information TLV types are added:

- o Type = TBD: Admin Label. The Information field contains a free-form UTF-8 string whose length is given by the Information Length field. The value is administratively assigned. There is no requirement to terminate the string with null or any other character.

Multiple admin labels can be included in the Peer UP. When multiple admin labels are included the BMP receiver MUST preserve the order.

The TLV is optional.

New optional element

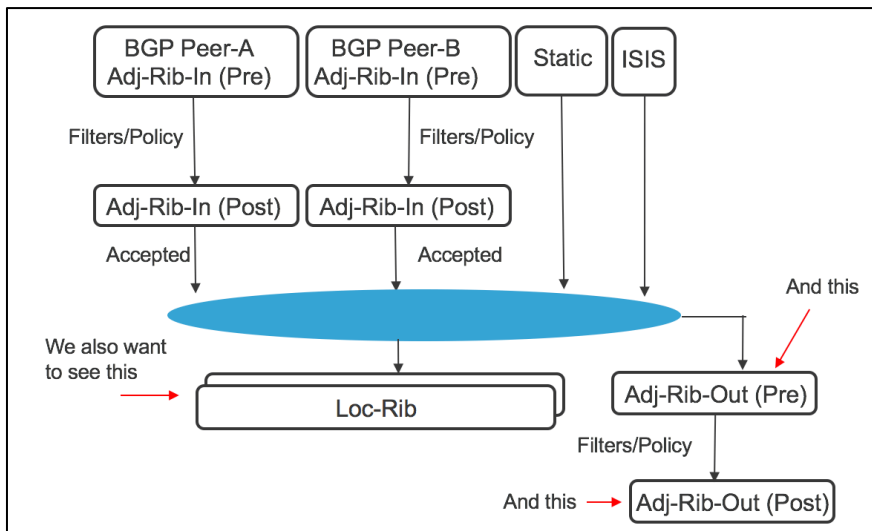
Loc-RIB and Adj-RIB-Out Extensions to BMP

`exit(0)`

Loc-RIB and Adj-RIB-Out Extensions to BMP

Backup Slides

Use Cases



Local-RIB

- Monitor routes selected and used by the router
- ECMP
- Correlation with netflow/ipfix
- Next hop preservation
- Monitor locally originated and BGP routes without requiring peering
- Policy verification

Adj-RIB-Out

- Monitor routes advertised to peers
- Policy verification

Extensions to RFC 7854

Local-RIB

- New Peer Type 3
 - Supports backwards capability
- Flags are specific to Peer Type
- New F Flag: indicates filtered
- Using existing Stats
- Replaces “Locally Originated Routes” specified in Section 8.2
- Peer Up Info TLV Type 3 : vrf name

Adj-RIB-Out

- New per peer header O flag to indicate Adj-RIB-In or Adj-RIB-Out
- Added statistics Types 14-17
 - 14 = Adj-RIB-Out pre-policy
 - 15 = Adj-RIB-Out post-policy
 - 16 = Adj-RIB-Out pre-policy per AFI/SAFI
 - 17 = Adj-RIB-Out post-policy per AFI/SAFI