

BMP

BGP Monitoring Protocol
GROW WG

IETF 106

Nov. 16-17, 2019

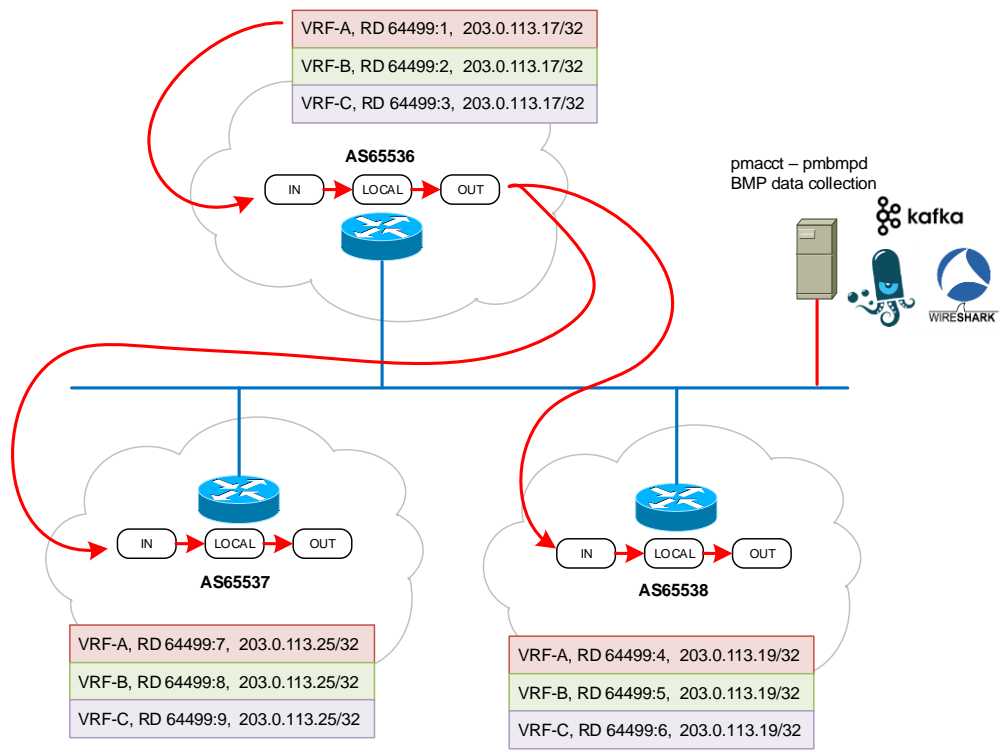
Singapore



Hackathon Plan

- Test Interoperability between router and data-collection for route monitoring and peer up/down message type extensions defined in
 - draft-ietf-grow-bmp-local-rib (BGP Local RIB)
 - RFC 8671 (BGP Adj-RIB Out)
- By using Wireshark BMP dissector and pmacct pmbmpd Apache AVRO messages

Topology - Test Setup



Gaps identified (1)

- BGP withdraw in Adj-RIB In does show in packet capturing but not in pmbmpd Kafka AVRO message.
- RFC 4271 doesn't specify BGP next-hop attribute for local originated routes in BGP local RIB. Propose to add recommendation in draft-ietf-grow-bmp-local-rib BMP (127.0.0.1 vs. 0.0.0.0?)

```
subsequent address family identifier (>AFI): Labeled VPN UNICAST (128)
  Next hop network address (12 bytes)
    Next Hop: Empty Label Stack RD=0:0 IPv4=127.0.0.1
    Number of Subnetwork points of attachment (SNPA): 0
  Network layer reachability information (16 bytes)
-----
```

0060	c5 f9 00 00 00 00 ff ff	ff ff ff ff ff ff ff ff
0070	ff ff ff ff ff ff 00 73	02 00 00 00 5c 40 01 01s \@ ..
0080	00 40 02 00 80 04 04 00	00 00 00 40 05 04 00 00	..@..... ..@.....
0090	40 10 c0 08 14 fb f0 01	2b fb f0 03 e9 fb f0 04	@..... +.....
00a0	09 fb f1 00 01 fb f3 00	11 c0 10 08 00 02 fb f1

Gaps identified (2)

- When BMP Adj-RIB Out and/or post policy is configured on router, one or multiple BMP peer up messages with different peer header (O and L bit set) are seen in packet capture.
- Due to this ambiguity, Wireshark BMP dissector shows peer up notifications under update message and pmbmpd duplicated peer up message types.
- We like to collect at GROW WG how other vendors implemented peer up.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	192.0.2.17	192.0.2.1	BGP	212	OPEN Message, OPEN Message
2	0.001141	192.0.2.19	192.0.2.1	BGP	212	OPEN Message, OPEN Message
3	0.002082	192.0.2.17	192.0.2.1	BGP	606	UPDATE Message


```
> Frame 3: 606 bytes on wire (4848 bits), 606 bytes captured (4848 bits)
> Ethernet II, Src: HuaweiTe_90:ff:ec (d0:c6:5b:90:ff:ec), Dst: Vmware_1b:75:4d (00:0c:29:1b:75:4d)
> Internet Protocol Version 4, Src: 192.0.2.17, Dst: 192.0.2.1
> Transmission Control Protocol, Src Port: 51354, Dst Port: 1790, Seq: 159, Ack: 1, Len: 552
> BGP Monitoring Protocol, Type Peer Up Notification
> BGP Monitoring Protocol, Type Peer Up Notification
> BGP Monitoring Protocol, Type Peer Up Notification
> BGP Monitoring Protocol, Type Route Monitoring
```

What we learned

- Good
 - Dedicated BMP Hackathon lab environment's (Internet & VPN) are working fine.
- Bad
 - Tests take longer than expected. Need automated test environment (CI, continuous integration).

Next Step for 107 Hackathon

- Understand different BMP implementations and support found gaps before next hackathon.
- Define BGP RIB state changes and their intended results.
- Develop test script which performs BGP configuration changes and validate results at data collection.
- Import BMP metrics with schema into TSDB and make it accessible with Web UI.
- Implement the following drafts for router and data collection
 - draft-ietf-grow-bmp-tlv
 - draft-cppy-grow-bmp-path-marking-tlv
 - draft-xu-grow-bmp-route-policy-attr-trace

Thanks to...

- Yunan Gu
- Alexis La Goutte
- Binyang Huang (remote)
- Camilo Cardona
- Christian Kuster
- Matthias Arnold
- Paolo Lucente
- Thomas Graf

