

BMP

BGP Monitoring Protocol
GROW WG

IETF 107

Mar. 23-27, 2020

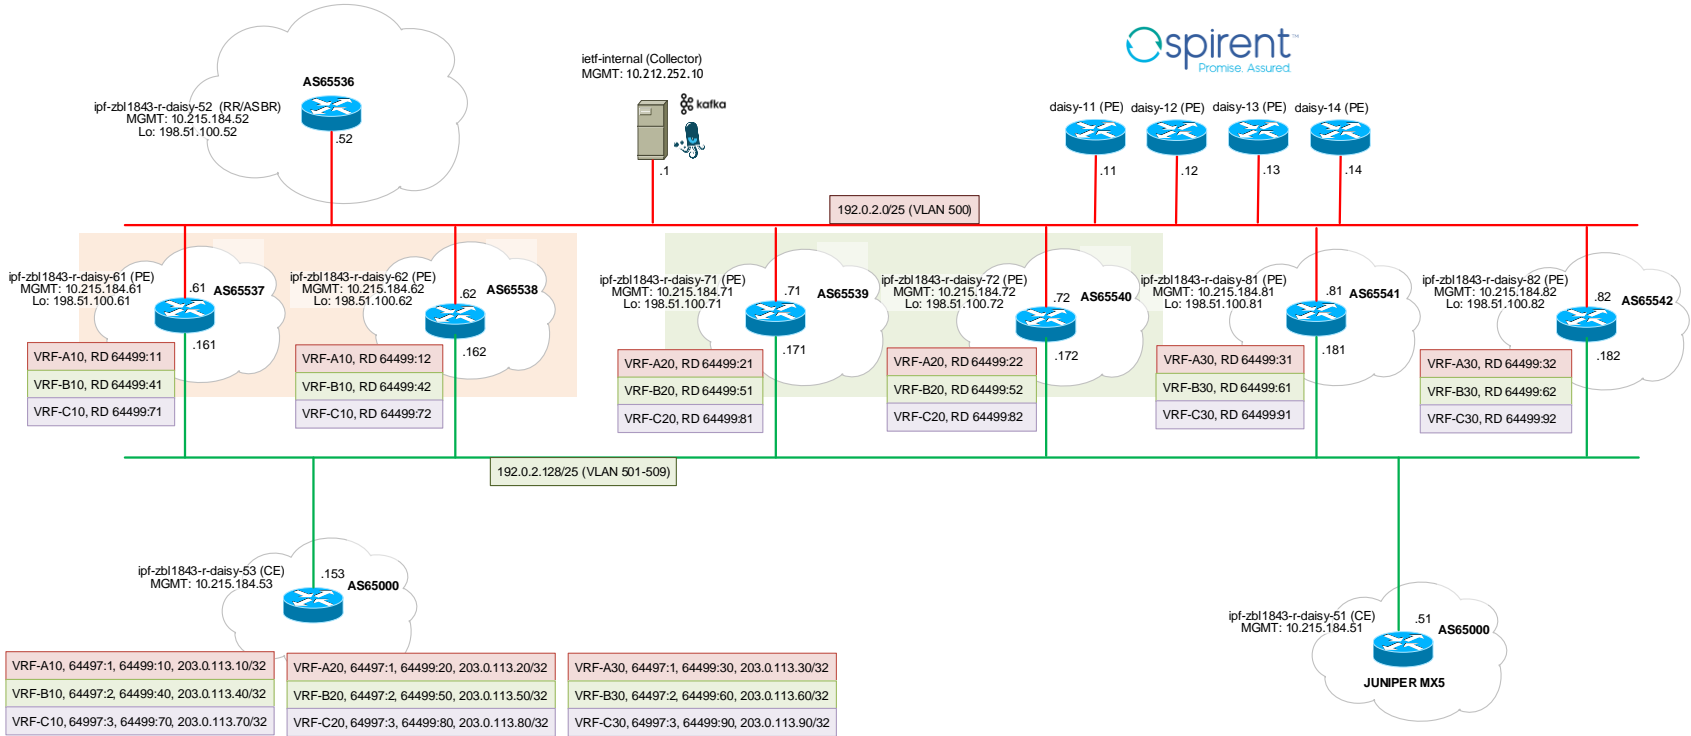
Virtual Hackathon



Hackathon Plan

- Test Interoperability between router and data-collection for initiation, peer up/down, statistics and route monitoring including message type extensions defined in
 - draft-ietf-grow-bmp-local-rib (BGP Local RIB)
 - RFC 8671 (BGP Adj-RIB Out)
 - draft-lucente-bmp-tlv (TLV support for BMP Route Monitoring and Peer Down Messages)
 - draft-cppy-grow-bmp-path-marking-tlv (Path Marking TLV)
 - draft-xu-grow-bmp-route-policy-attr-trace (BGP Route Policy and Attribute Trace)
- Test performance impact of BMP on router CPU/Memory resources and BGP route propagation.
- By using Wireshark BMP dissector, [pmacct](#) pmbmpd, Apache [Kafka](#), Apache [Druid](#) and [Pivot](#).

Topology - Test Setup



BMP initiation Message Type

The screenshot displays a dashboard for 'BMP Type 4 - Init'. The left sidebar lists dimensions: Time, Bmp Init Info String, Bmp Init Info Sysdescr, Bmp Init Info Sysname, Bmp Msg Type, Bmp Router, Bmp Router Port, Event Timestamp, Event Type, Seq, and Writer Id. The main area shows a table with the following data:

Bmp Router, Bmp Init Info Sysname	Number of Events ↓
Overall	16
192.0.2.52	2
ipf-zbl1843-r-daisy-52	2
192.0.2.53	2
ipf-zbl1843-r-daisy-53	2
192.0.2.61	2
ipf-zbl1843-r-daisy-61	2
192.0.2.62	2
ipf-zbl1843-r-daisy-62	2
192.0.2.71	2
ipf-zbl1843-r-daisy-71	2
192.0.2.72	2
ipf-zbl1843-r-daisy-72	2
192.0.2.81	2
ipf-zbl1843-r-daisy-81	2
192.0.2.82	2
ipf-zbl1843-r-daisy-82	2

*BMP router IP address
grouped by BMP router name*

BMP peer_up Message Type

Bmp Router, Peer Ip	Count ↓
Overall	134
192.0.2.61	42
0.0.0.0	18
192.0.2.153	12
2001:db8:1::153	12
192.0.2.53	24
192.0.2.161	12
2001:db8:1::161	12
192.0.2.62	12
0.0.0.0	12
192.0.2.71	12
0.0.0.0	12
192.0.2.72	12
0.0.0.0	12
192.0.2.81	12
0.0.0.0	12
192.0.2.82	12
0.0.0.0	12
192.0.2.52	8
198.51.100.61	8

*Peer IP addresses grouped by
BMP router*

Bmp Router, Peer Ip, Rd	Count ↓
Overall	42
192.0.2.61	42
0.0.0.0	18
0:64499:11	6
0:64499:41	6
0:64499:71	6
192.0.2.153	12
0:64499:11	4
0:64499:41	4
0:64499:71	4
2001:db8:1::153	12
0:64499:11	4
0:64499:41	4
0:64499:71	4

*Peer IP addresses of BMP router 192.0.2.61 grouped by
BGP route-distinguisher*

BMP stats Message Type

Dimensions	Number of Events
Overall	200.10 k
0:64499:11	66.81 k
192.0.2.153	31.95 k
2001:db8:1::153	31.10 k
0.0.0.0	3.76 k
0:64499:41	66.73 k
192.0.2.153	31.95 k
2001:db8:1::153	31.02 k
0.0.0.0	3.76 k
0:64499:71	66.55 k
192.0.2.153	31.80 k
2001:db8:1::153	30.99 k
0.0.0.0	3.76 k

Peer IP addresses of BMP router 192.0.2.61 grouped by BGP route-distinguisher

Counter Type Str	Number of Events
Overall	31.10 k
Number of routes in per-AFI/SAFI Adj-RIB-In	5.98 k
Number of routes in per-AFI/SAFI Adj-RIB-Out Post-Policy	5.98 k
Number of routes in per-AFI/SAFI Loc-RIB	5.98 k
Number of (known) duplicate prefix advertisements	1.20 k
Number of (known) duplicate withdraws	1.20 k
Number of prefixes rejected by inbound policy	1.20 k
Number of routes in Adj-RIBs-in	1.20 k
Number of routes in Adj-RIBs-Out Post-Policy	1.20 k
Number of routes in Loc-RIB	1.20 k
Number of updates invalidated due to AS_CONFED loop	1.20 k
Number of updates invalidated due to AS_PATH loop	1.20 k
Number of updates invalidated due to CLUSTER_LIST loop	1.20 k
Number of updates invalidated due to ORIGINATOR_ID	1.20 k
Number of updates subjected to treat-as-withdraw	1.20 k

Peer statistics of BMP router 192.0.2.61, BGP route-distinguisher 0:62459:184549376 and Peer IP 2001:db8:1::153

BMP route_monitoring Message Type

BMP Type 0 - Routemonitor

FILTER Latest 6 hours Bmp Router: 192.0.2.52

SHOW Rd

Rd	Number of Events ↓
Overall	432
0:64499:11	28
0:64499:12	28
0:64499:21	28
0:64499:22	28
0:64499:31	28
0:64499:32	28
0:64499:41	28
0:64499:42	28
0:64499:51	28
0:64499:52	28
0:64499:61	28
0:64499:62	28
0:64499:81	28
0:64499:82	28
0:64499:71	20
0:64499:72	20

*BGP route-distinguishers of
BMP router 192.0.2.52*

BMP Type 0 - Routemonitor

FILTER Latest 6 hours Bmp Router: 192.0.2.52 Rd: 0:64499:11

SHOW Ip Prefix Peer Ip

Ip Prefix, Peer Ip	Number of Events ↓
Overall	28
203.0.113.10/32	28
198.51.100.61	8
198.51.100.62	4
198.51.100.71	4
198.51.100.72	4
198.51.100.81	4
198.51.100.82	4

*BGP Prefixes with BGP route-distinguisher
0:64499:11 grouped by Peering IP Address*

Pmacct -pmbmpd

Achievements

- First-pass implementation of draft-xu-grow-bmp-route-policy-attr-trace
- First-pass implementation of draft-lucente-grow-bmp-tlv-ebit (decoding of PEN), applied to draft-cppy-grow-bmp-path-marking-tlv
- Honoring of Peer Distinguisher in per-peer message headers for peer_up, peer_down and statistics

Gaps Identified

- Honoring of Peer Distinguisher in per-peer message headers for peer_up, peer_down, statistics and route_monitoring

Next Steps

- Honoring of Peer Distinguisher in per-peer message headers for route_monitoring
- Decode path-status attribute in draft-xu-grow-bmp-route-policy-attr-trace

<https://github.com/pmacct/pmacct/>

Swisscom – lab environment

Achievements

- Lab topology increased to 9 routers
- Netconf-yang orchestrated
- Spirent test center added for BGP route generation
- Apache Druid timeseries DB and Pivot UI for metric validation

Gaps Identified

- CPU and memory monitoring for stress testing

Next Steps

- Implement CPU and memory monitoring with yang push

Huawei - VRP

Achievements

- Supporting draft-xu-grow-bmp-route-policy-attr-trace-03
- Supporting path status of draft-cppy-grow-bmp-path-marking-tlv-03
- Wireshark BMP dissector for packet validation
- Stress tests showing CPU usage increase but no BGP propagation delay

Gaps Identified

- Repeated report of route trace attribute

Next Steps

- Supporting reason code of draft-cppy-grow-bmp-path-marking-tlv-03

What we learned

- Good
 - Metrics were easily discoverable with Pivot UI.
Makes it a breeze to validate
- Bad
 - Need to improve testing setup to validate BMP stress test better

Thanks to...

- Yunan Gu
- Binyang Huang
- Paolo Lucente
- Matthias Arnold
- Thomas Graf

...[Imply](#) and Swisscom Time Analytics Platform team for providing us the big data environment.

