

BGP Route Policy and Attribute Trace Using BMP

draft-xu-grow-bmp-route-policy-attr-trace-00


Feng Xu, Tencent

Yunan Gu, Shunwan Zhuang, Zhenbin Li, Huawei

Mar. 24, 2019

Motivations

- Route monitoring
 - BMP, RouteViews, RIPE RIS...
- Policy monitoring
 - Netconf, gRPC, CLI...
- Correlated route and policy monitoring?
 - Coming soon...
- In a nutshell
 - Each route policy processing is recorded as an event
- Use cases
 - Route policy validation
 - Root cause analysis

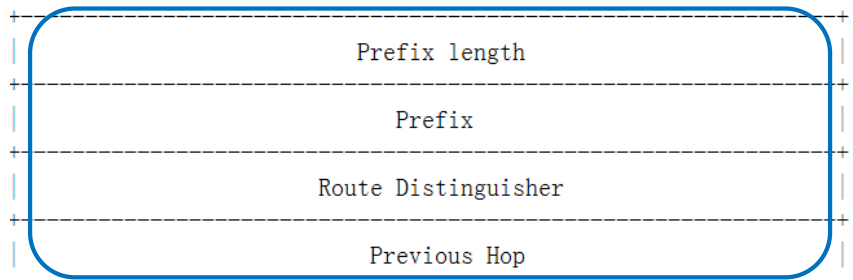


Prefix	Route event	Route policy	Time stamp	Next hop	Cost
172.17.0.0/16	1	ISIS: Route-policy r1 : permit/permit : cost 100	xx:xx:xx	192.168.2.2/24	100
	2	RM: Route-policy r2 : permit/deny : next-hop	xx:xx:xx	192.168.1.1/24	100
	3	RM: Route-policy r3 : permit/deny : cost 200	xx:xx:xx	192.168.1.1/24	200

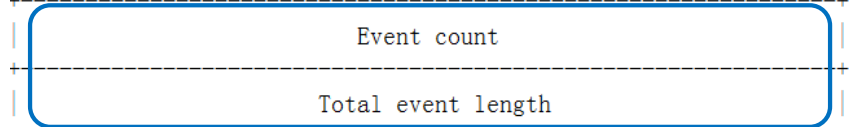
BMP Extensions

- Why using BMP?
 - BGP routes and policy
- Why not extending BMP local rib?
 - A prefix is not necessarily added to the local-rib
 - One prefix maps to multiple events, the record is event based
- How?
 - BMP Common header
 - A new BMP message type
 - Type = TBD: Route Policy and Attribute Trace Message
 - No Per Peer Header
 - Message format

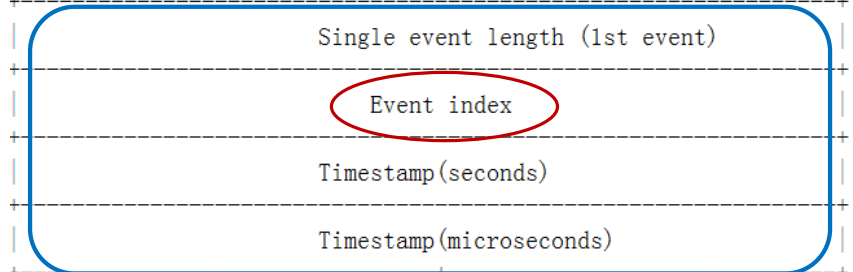
Prefix information




Event count





Single event information



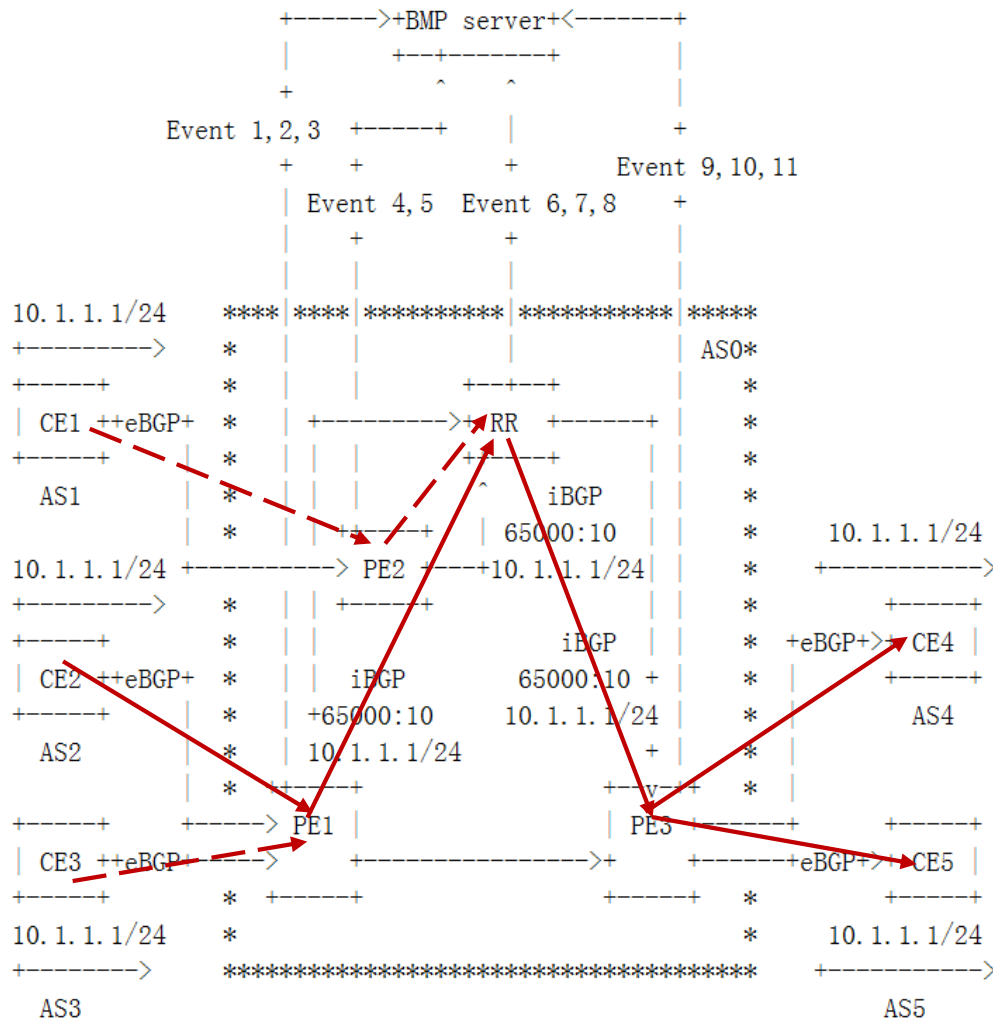
Route policy information



Route attribute information



Route Footprint Recovery Example



- PE1
 - Event 1: 10.1.1.1/24 inbound from CE2
 - Event 2: 10.1.1.1/24 inbound from CE3
 - Event 3: 10.1.1.1/24 from CE2 outbound to RR
- PE2
 - Event 4: 10.1.1.1/24 inbound from CE1
 - Event 5: 10.1.1.1/24 from CE1 outbound to RR
- RR
 - Event 6: 10.1.1.1/24 inbound from PE2
 - Event 7: 10.1.1.1/24 inbound from PE1
 - Event 8: 10.1.1.1/24 from PE1 outbound to PE3
- PE3
 - Event 9: 10.1.1.1/24 inbound from RR
 - Event 10: 10.1.1.1/24 from RR outbound to CE4
 - Event 11: 10.1.1.1/24 from RR outbound to CE5

Event 1,2,3

RD: 65000:10	
Prefix: 10.1.1.1/24	
Event 1	
Timestamp 1	
Policy ID: WC1, node 101	Inbound policy
Peer ID: CE2	
Peer AS: AS2	
VRF/Table name: VRF 1	
AFI: IPv4	
Previous Hop: CE2	
Event 2	
Timestamp 2	
Policy ID: WC1, node 102	Inbound policy
Peer ID: CE3	
Peer AS: AS3	
VRF/Table name: VRF 1	
AFI: IPv4	
Previous Hop: CE3	
Event 3	
Timestamp 3	
Policy ID: RR1, node 103	Outbound policy
Peer ID: RR	
Peer AS: AS0	
VRF/Table name: VRF 1	
AFI: VPNv4	
Previous Hop: CE2	