

BGP FlowSpec extensions
for
Routing Policy Distribution(RPD)
draft-li-idr-flowspec-rpd-00

Liang Ou, Yujia Luo

Robin Li, Vincent Zhuang, Eric Wu

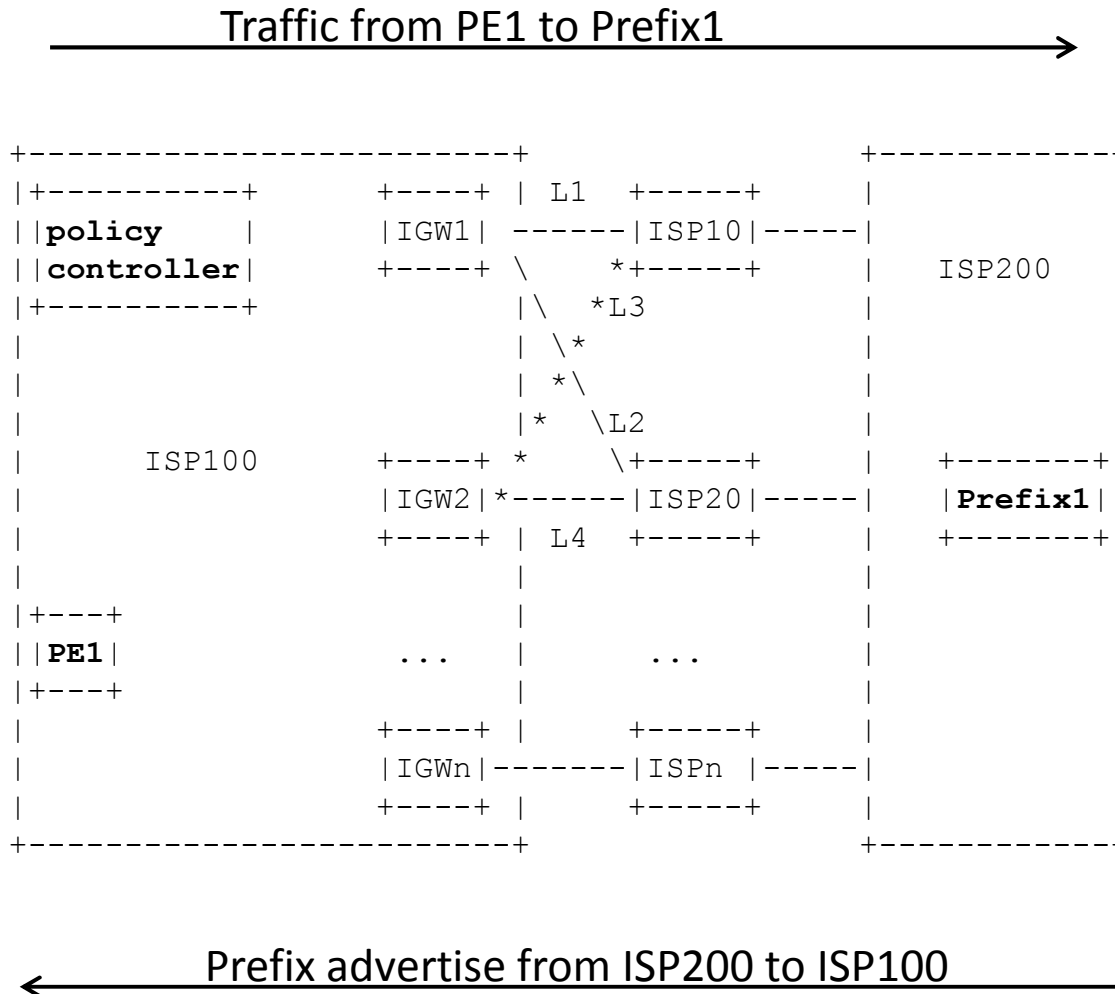
IETF93, Prague

Motivation

- Provider's motivation for traffic adjustment:
 - Link congestion or overload;
 - Packet delay, loss or corrupted;
 - Prefer lower price;

Genuine requirements

□ Outbound traffic control



□ EBGP peering:

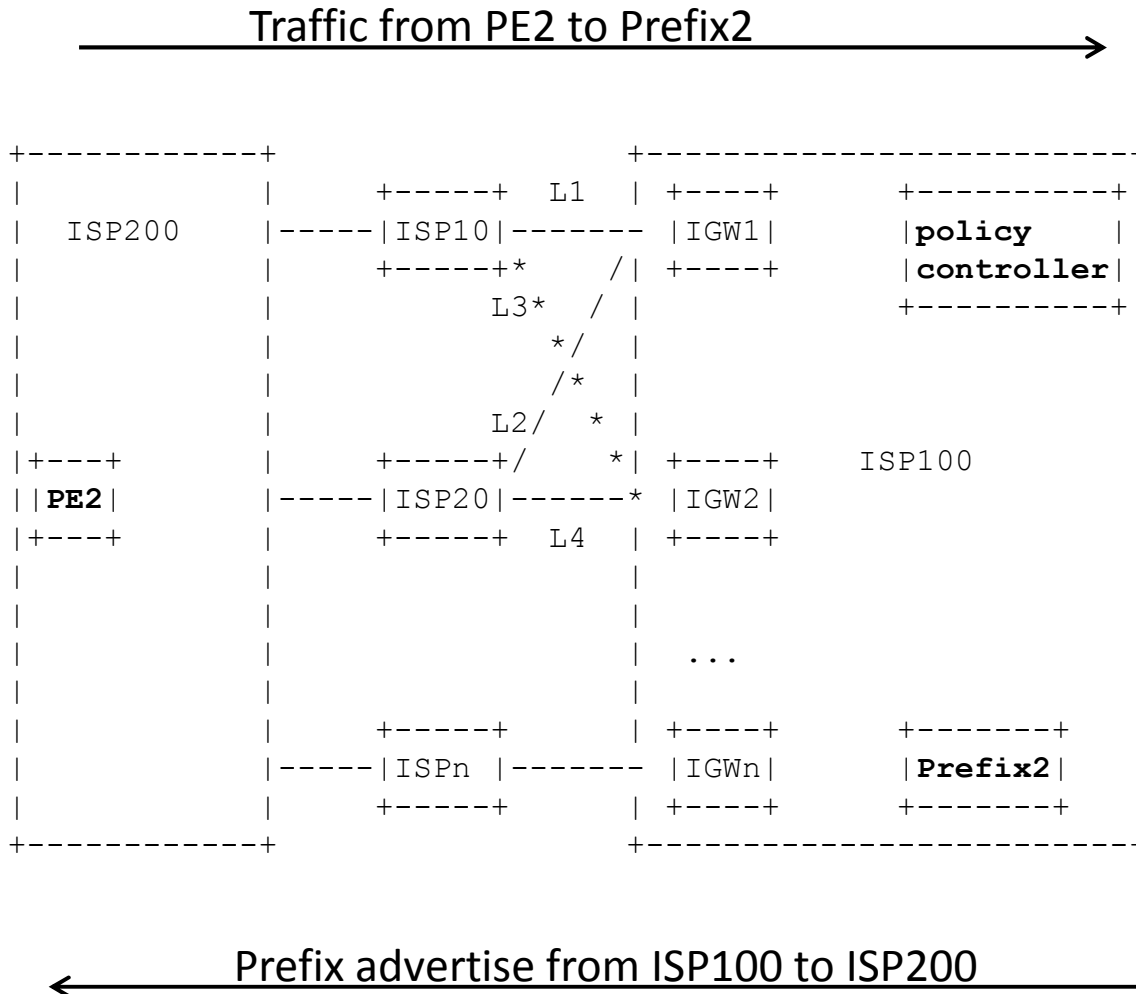
- IGW1---L1---ISP10
- IGW1---L2---ISP20
- IGW2---L3---ISP10
- IGW2---L4---ISP20

□ Requirement:

- Exit through L3

Genuine requirements

□ Inbound traffic control



□ EBGP peering:

- IGW1---L1---ISP10
- IGW1---L2---ISP20
- IGW2---L3---ISP10
- IGW2---L4---ISP20

□ Requirement:

- Enter through L3

Solution requirements

- Minimal configuration burden, less manually;
- Dynamic provisioning instead of static one;
- Safe to undo or rollback without damage;
- Easy to maintain with reasonable complexity;
- High performance mechanism with small delay;
- Backward compatible with less software upgrade;


Proposed solution

□ Routing Policy Distribution (RPD)

- Taking effect on control plane
- Impact decision on remote site

□ RPD protocol: BGP Flowspec


- Filtering rule: destination for prefix1/prefix2
- Action: R-bit introduced, more info carried in new attribute



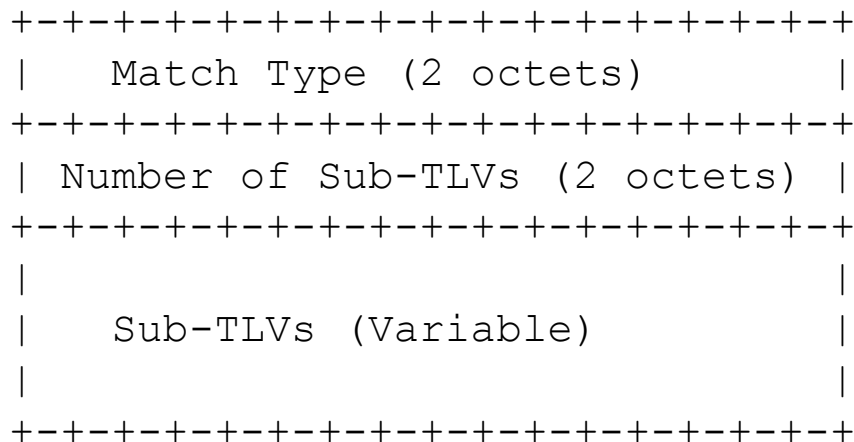
Proposed solution

□ BGP Policy Attribute


• Attribute structure



• Match field



• Action field



• Action type 1: Route-Preference

• Action type 2: Route-Prepend-AS

□ Match type

• Value 0: Positive match

• Value 1: Negative match

□ Sub-TLVs

• Type 1: IPv4 Neighbor

• Type 2: IPv6 Neighbor

• Type 3: ASN list

Proposed solution

❑ Outbound traffic control

- Match type: permit
- IPv4 neighbor sub-TLV:
 - ✓ Local BGP Speaker IGW2
 - ✓ Remote BGP Peer ISP10
- Action type: Route-Preference

❑ Inbound traffic control

- Match type: deny
- IPv4 neighbor sub-TLV:
 - ✓ Local BGP Speaker IGW2
 - ✓ Remote BGP Peer ISP10
- Action type: Route-Prepend-AS
- Action value: Prepend-AS five times

Comparing to traditional routing policy


❑ Drawbacks using traditional routing policy:

- Manual configuration burden and mistakes.
- Complexity increased and difficulty to maintain.
- Static configuration versus dynamic requirements.

Comparing to traditional BGP-Flowspec

□ Current BGP-Flowspec limitation

- Can't affect AS external device's decision.




- On IGW1, L2 is preferred previously.
- Flowspec can redirect to L1
- CE prefers IGW3, which makes IGW1's redirection useless.

Comparing to traditional BGP-Flowspec

❑ Current BGP-FS limitation

- Changing decision in forwarding plane may introduce loop.



- IGW2 redirect from M2 to M3;
- M3 may still prefer IGW2's nexthop
- Loop can happen between M3 and IGW2.

Next step

- ❑ Collect feedback and comments.
- ❑ Refine this draft according to comments.