

BGP FlowSpec extensions for Routing Policy Distribution(RPD) draft-li-idr-flowspec-rpd-02

Zhenbin Li (lizhenbin@huawei.com)

Liang Ou (oul@gsta.com)

Yujia Luo (luoyuj@gsta.com)

Sujian Lu (jasonlu@tencent.com)

Shunwan Zhuang (zhuangshunwan@huawei.com)

Nan Wu (eric.wu@huawei.com)

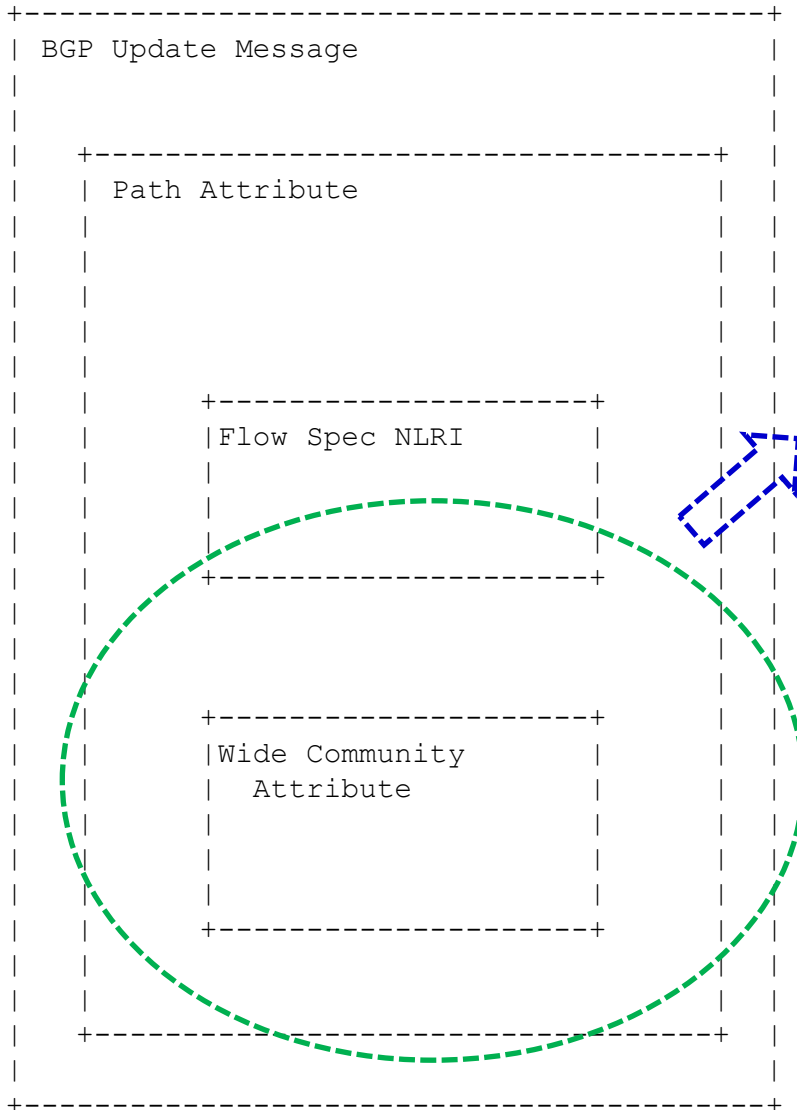
IETF97, Seoul

History & Changed from 01 version

- ❑ History: Presented at the 93th, 94th IETF, idr-interim meeting(2015-12-9)

- ❑ Changed from 01 version:
 - Refine use cases , some related use cases are recorded in the document [draft-luo-grow-ts-use-cases].
 - <https://tools.ietf.org/id/draft-luo-grow-ts-use-cases-01.txt>
 - Implemented a prototype, some brief implementation information is documented in [draft-wu-idr-flowspec-rpd-impl].
 - <https://tools.ietf.org/id/draft-wu-idr-flowspec-rpd-impl-00.txt>
 - The next version will provides more information.

RPD Mechanism in Summary



Processing steps:

1. Effective on which routes → Filtered by Flowspec NLRI
2. Effective on which peers → Filtered by **Wide Community Attribute**
3. Take the action in **Wide Community Attribute**

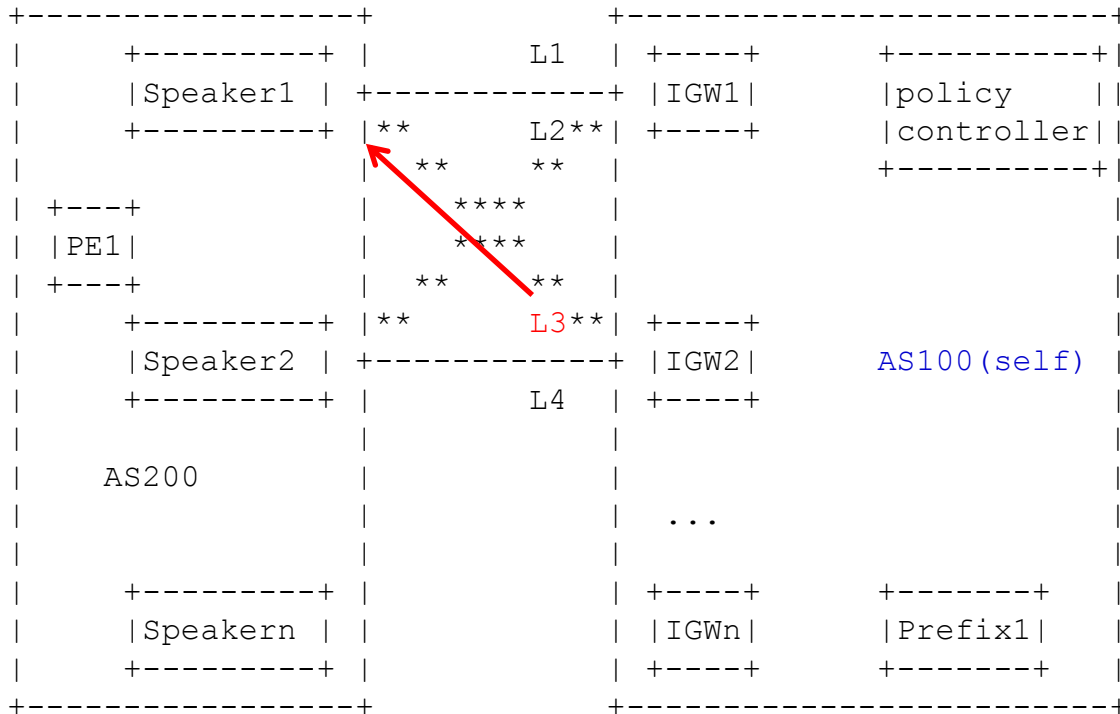
The routing policy can be mapped into the Wide BGP Community Attribute exactly.

Application (1)

□ Inbound traffic control

Traffic from PE1 to Prefix1

----->



Prefix1 advertises from AS100 to AS200

<-----

□ EBGP peering:

- Speaker1---L1---IGW1
- Speaker2---L2---IGW1
- Speaker1---L3---IGW2
- Speaker2---L4---IGW2

□ Requirement:

- Administration only on AS100
- Traffic enter AS100 through L3

Encoding Example (1)

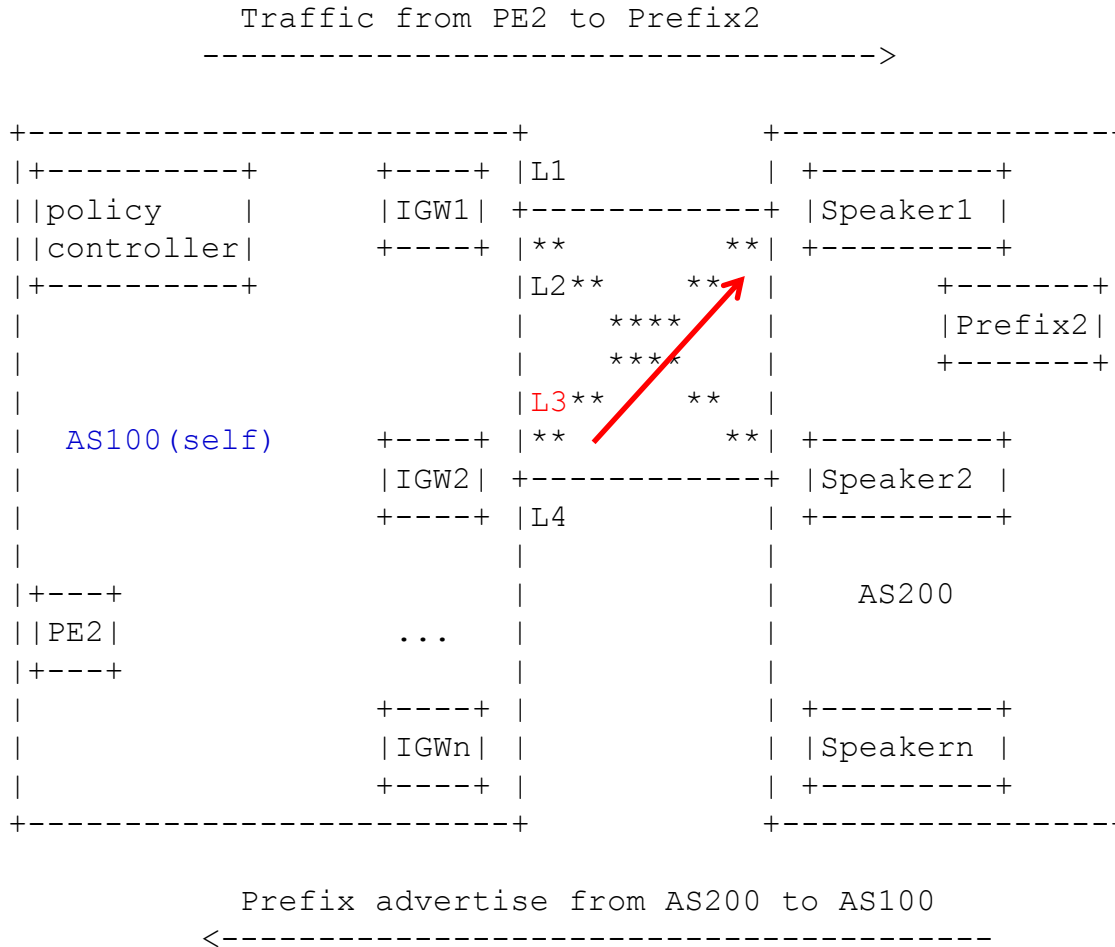
❑ Inbound Traffic Control encoding example



- ❑ As required in the case, traffic from PE1 to Prefix1 need to enter through L3, so IGWs except IGW2 should prepend ASN list to Prefix1 when populating to AS100.
- ❑ As shown in left figure, community "PREPEND N TIMES TO AS" and "Exclude Target(s) TLV" are be used.

Application (2)

□ Outbound traffic control



□ EBGP peering:

- IGW1---L1---Speaker1
- IGW1---L2---Speaker2
- IGW2---L3---Speaker1
- IGW2---L4---Speaker2

□ Requirement:

- Administration only on AS100
- Traffic exit through L3

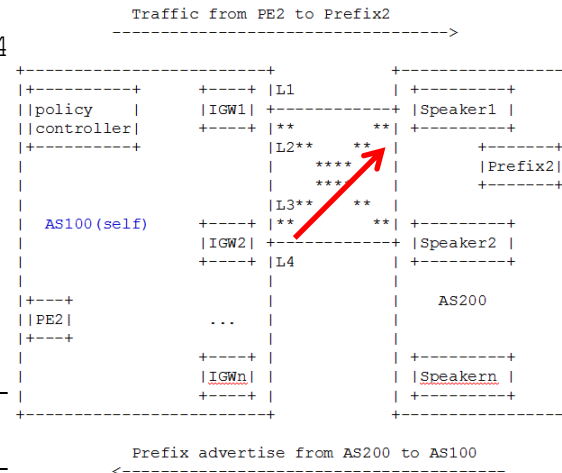
Encoding Example (2)

Outbound Traffic Control encoding example

```

0
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Container Type 1 (1) |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| 1 0 0 0 0 0 0 0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Hop Count: 0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Length: 36 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Community: LOCAL PREFERENCE
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Own ASN | 100 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Context ASN# | 100 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| TargetTLV(1) | Length: 11 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| IPv4Neig(TBD) | Length: 8 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Local Speaker | #IGW2 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Remote Speaker | #Speaker1 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Param TLV (3) | Length: 7 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Integer (4) | Length: 4 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Increment # | 100 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```



- EBGP peering:
 - IGW1---L1---Speaker1
 - IGW1---L2---Speaker2
 - IGW2---L3---Speaker1
 - IGW2---L4---Speaker2

- Requirement:
 - Administration only on AS100
 - Traffic exit through L3

- As required in the case, traffic from PE2 to Prefix2 need to exit through L3, so IGWs should prefer the route from IGW2 to Speaker1.
- As shown in left figure, community "LOCAL PREFERENCE" and "Target(s) TLV" are be used.

Next step

- ❑ Collect feedback and comments.
- ❑ Refine this draft according to comments.
- ❑ Adding new use cases from operators.
- ❑ Ask for early allocation for Wide Community Attribute