

BGP Extended Community for Identifying the Target Node

draft-dong-idr-node-target-ext-comm-00

Jie Dong, Shunwan Zhuang @ Huawei

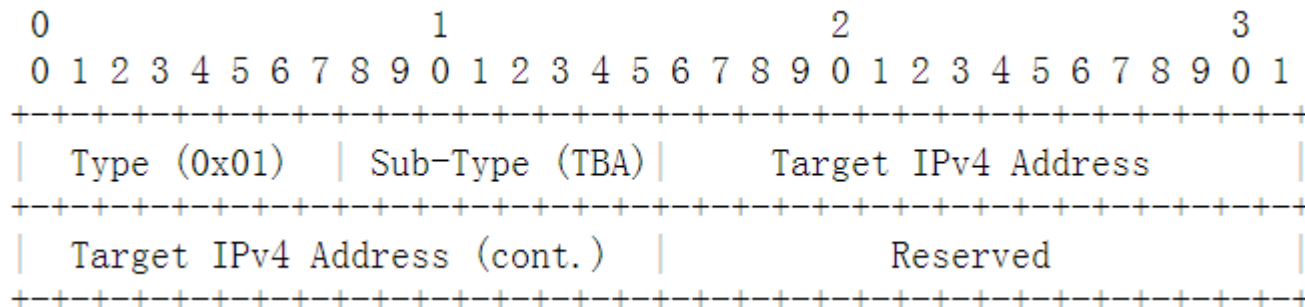
Gunter Van de Velde @ Nokia

Motivation

- BGP has been used for the distribution of various routing and policy information
 - In some cases the intended target may be one or several particular BGP nodes
 - One example is BGP FlowSPec
- BGP does not have the mechanism of designating the receiving nodes
 - Was designed for P2MP distribution
 - Route Target (RT) is used for matching VPN routes to VRFs
- A general mechanism is needed to control the distribution of BGP information to particular nodes

Proposed Solution

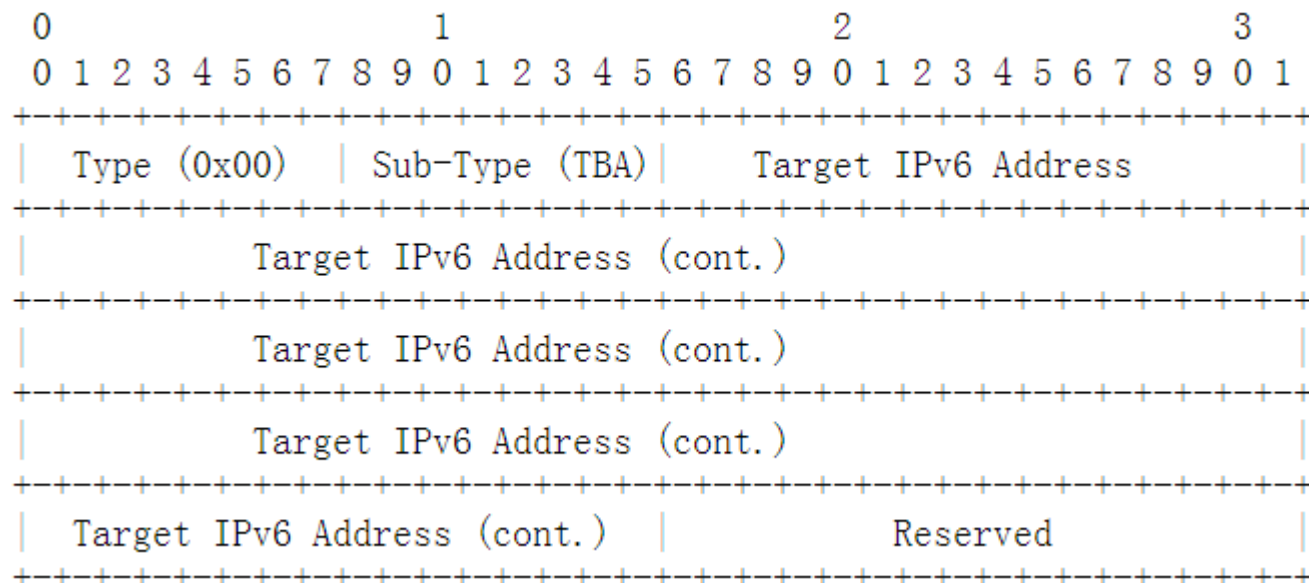
- New BGP extended community to carry the target node information
 - IPv4 Node Target extended community



- One or more IPv4 Node Target extended communities may be carried in BGP Update

Proposed Solution (Cont.)

- New IPv6 address specific extended community to carry the target node information
 - IPv6 Node Target extended community



- One or more IPv6 Node Target extended communities may be carried in BGP Update

Procedures for Intra-AS Scenario

- Originator
 - Encode the IPv4/v6 Node Target extended communities with the target nodes' IPv4/v6 addresses
- Receiving nodes
 - Non-RR speaker: check the node target extended communities in the received Update
 - If any of them match with the local address, the routes are valid for local use
 - Otherwise, discard the Update message
 - Route reflector: in addition to checking for local match, also responsible for the further distribution to the clients

Comments and Discussion

- Is IPv6 Node Target necessary? The control plane may still be IPv4 based
 - Do we need to support pure IPv6 (data plane and control plane)?
- Restrictions on the target IP addresses, suggest to only use BGP Router-ID (4-octet)
 - Can BGP Router-ID be used in pure IPv6 scenario?
 - How about the host IP addresses (IPv4 or IPv6)?
- Transitive or non-transitive extended community
 - Depends on whether we need to cover the inter-AS scenario

Next Steps

- Collect feedbacks on the problem space
 - Is this a valid problem?
 - Do we need to cover both IPv4 and IPv6?
 - Do we need to cover both intra-AS and Inter-AS?
- Discuss and polish the solution
 - Contributions are welcome!