IETF84
BGP Flow-Spec Ext Community for Traffic Redirect to IP Next Hop

draft-simpson-idr-flowspec-redirect-01

J. Uttaro	uttaro@att.com
M. Texier	mtexier@arbor.net
P. Mohapatra	pmohapat@cisco.com
D. Smith	djsmith@cisco.com
W. Henderickx	wim.henderickx@alcatel-lucent.com
A. Simpson	adam.simpson@alcatel-lucent.com

Motivation

- When Flow-spec is used as a DDoS mitigation tool the ability to override the IP forwarding next-hop and redirect traffic towards an alternate destination is very useful
  - For example towards a scrubbing appliance
- Unfortunately RFC 5575 provides only a ‘redirect to VRF’ action for this purpose
  - Extended community 0x8008, encodes a route target value
  - It may not be convenient to configure a VRF on every redirecting node, especially if the VRFs needs to be part of an L3 VPN to send traffic across the core
- New I-D proposes a new flow-spec extended community, ‘redirect to IP next-hop’, which has fewer prerequisites and is simpler to use

New Extended Community

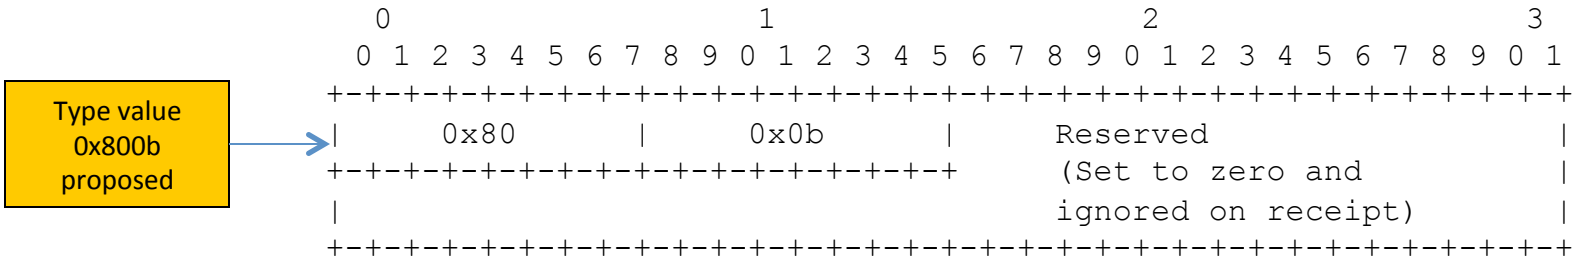

- New ext comm. plus flow-spec NLRI creates a traffic filtering rule that forwards matched packets towards the IPv4 or IPv6 address encoded in the ‘Network Address of Next-Hop’ field of the MP_REACH_NLRI
  - Router looks up the MP_REACH_NLRI next-hop in the IP FIB, finds next-hop interface/tunnel X, and forwards matched packets to X
  - X can be any type of interface/tunnel: Ethernet VLAN, GRE tunnel, LDP LSP, RSVP LSP, BGP-3107 LSP, etc.
- May appear with other flow-spec extended communities in the same Update
  - But ‘redirect to IP’ ignored if ‘redirect to VRF’ is also present

Inter-AS Considerations

- New extended community is transitive across AS
  - Next-hop address in MP_REACH_NLRI may or may not be reset when the flow-spec route is advertised to an EBGP peer; policy decision
- New validation procedure should be applied by default to 'redirect to IP' ext comm. received from EBGP peer
  - Discard the ext comm. (do not propagate further) if the last AS in the AS[4]_PATH of the longest prefix match for MP_REACH_NLRI next-hop address does not match the ASN of the EBGP peer
  - This check is in addition to the basic checks (re: NLRI with a destination prefix subcomponent) described in RFC 5575 and amended in draft-ietf-idr-bgp-flowspec-oid
  - Must be possible to disable the new check

Feedback?

- General comments
- WG adoption
- Which community combinations should be allowed & described?
  - Use cases for redirect-to-IP AND redirect-to-VRF?
- Should the new validation check result in discarding the entire flow-spec route and not just the redirect to IP ext comm.?