

# BGP Flowspec

## Redirect Load Balancing Group Community

`draft-wu-idr-flowspec-redirect-group-00`


Z. Wu, H. Wang, L. Wang, Z. Tan (*Huawei Technologies*)

IETF113, Mar 2022

Vienna

# Motivation & Problem Statement

- BGP Flow Specification allows flow routes that carry traffic policies to steer traffic.
- In some scenarios, we need to steer traffic to a load balancing group, either ECMP or UCMP.



- For the 50G traffic to R1 from R0, we may want to steer traffic to R2, R3 and R4 in a 1:2:2 ratio
- The ratio could be changing according to network running status
- Routers may be deployed in IP, MPLS or SRv6 networks
- Current set of mechanisms can hardly support neither ECMP of SRv6 tunnels nor UCMP of either types.

# Redirect Load Balancing Group Community

- New Type of BGP Wide Community
- Extension to "BGP Community Container Attribute" [draft-ietf-idr-wide-bgp-communities]
- General formats comply with "BGP Community Container Attribute"
- New Community Value: **Redirect Load Balancing Group (TBD)**, require IANA reregistration
- **MUST** contain only 1 Parameter TLV(subtype 3)
- Parameter TLV contains a list of atoms, each represents a redirection action

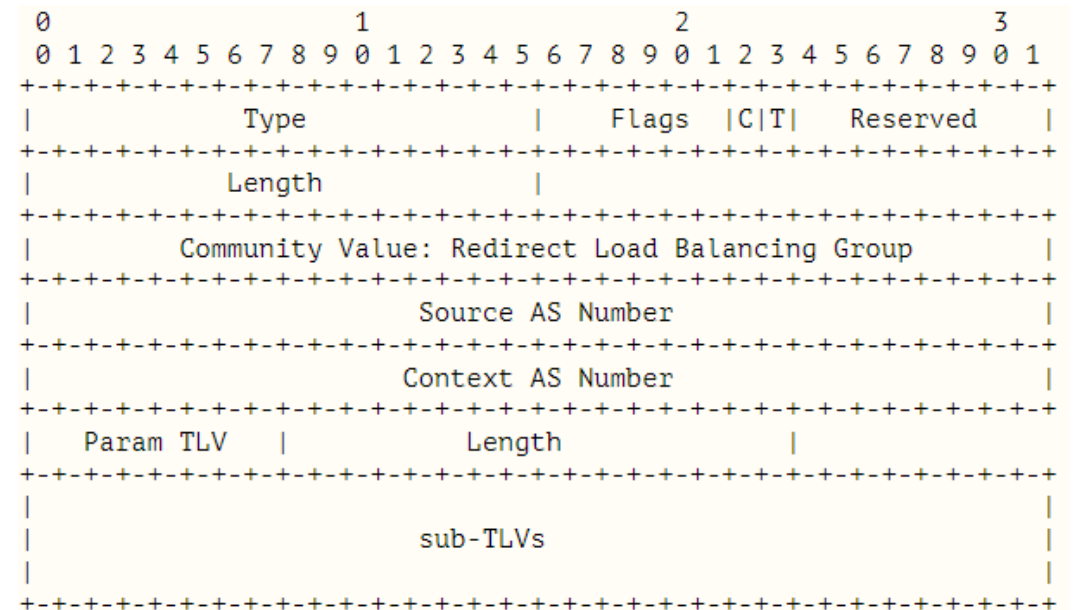


Figure 1: Redirect Load Balancing Group Community Format

# Atom (Param Sub-TLV) Formats

- Comply with wide community atom format

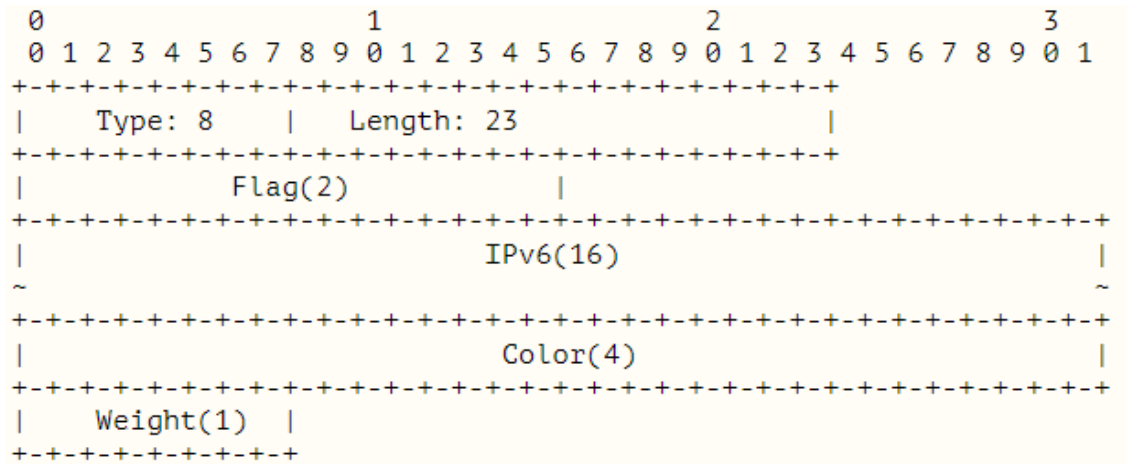


Figure 10: Atom Type 8: IPv6 Prefix with Color and Weight

- Type:** Used exclusively within Redirect Load Balancing Group Community
- Length:** fixed for each atom
- Flag:** 2 octets, reserved for future use, MUST be set to 0
- IPv6:** 16-octet IPv6 address or SRv6 tunnel Endpoint for redirection
- Color:** 4 octets, SRv6 tunnel Color for redirection
- Weight:** 1 octet, values from 1~255, load balancing weight

# Supported Atoms: redirection actions

Type 1:	IPv4 Prefix Only	(unweighted IPv4 address)
Type 2:	IPv4 Prefix with Weight	(weighted IPv4 address)
Type 3:	IPv4 Prefix with Color	(unweighted SR-TE tunnel)
Type 4:	IPv4 Prefix with Color and Weight	(weighted SR-TE tunnel)
Type 5:	IPv6 Prefix Only	(unweighted IPv6 address)
Type 6:	IPv6 Prefix with Weight	(weighted IPv6 address)
Type 7:	IPv6 Prefix with Color	(unweighted SRv6 tunnel)
Type 8:	IPv6 Prefix with Color and Weight	(weighted SRv6 tunnel)

- In principle, sub-TLVs may be combined in any mode
- Supported combinations depend on the specific implementation
- ECMP: "Redirect Group" contains more than 1 atoms, not all atoms are of a weighted type
- UCMP: "Redirect Group" contains more than 1 atoms, all atoms are of a weighted type

## **Next step**

- Welcome more comments and discussion

Thank you!