

# BGP-SPF Flooding Reduction

draft-chen-lsvr-flood-reduction-01

Huaimo Chen

Gyan Mishra

Aijun Wang

Yisong Liu

Haibo Wang

Yanhe Fan

IETF 113

# Overview

Thanks people below for their comments and suggestions

- Acee Lindem
- Keyur Patel
- Donald E. Eastlake

## Updates to Previous Version

- Simplify it by removing RR groups
- Some Editorial Changes

# Revised Flooding Procedure in RR Model

- ✓ BGP SPF speaker/node sends its Link NLRI to some such as one of RRs.
- ✓ After receiving it, the RR sends the NLRI to the other BGP SPF speakers.

For example,

A sends its link state to one RR RR1, not to RR2.

After receiving the link state from A, RR1 sends it to B and C.




Figure 2. BGP-SPF Domain with two RRs

B and C receives only one copy of the same NLRI

Comparing to normal flooding in RR model, revised flooding reduced the amount of flooding by half.

# Revised Flooding Procedure in Node Connections Model

Similar to the one in ietf-lsr-dynamic-flooding:

- ✓ Each node has a flooding topology (FT).
  - In an option, FT is computed in a distributed mode, where every BGP SPF speaker computes a FT for the network using a same algorithm.
  - In another, FT is computed in a centralized mode, where one BGP SPF speaker elected as a leader computes a FT and advertises FT to every BGP SPF speaker. For a new FT computed, **only changes are advertised**.
- ✓ Each node sends link NLRI to its peers on FT (i.e., are connected by the links on the FT).




Figure 2a. A Flooding Topology




Figure 3a. Advertise NLRI Using Flooding Topology

For example, Figure 3a shows a flooding flow of a link NLRI.

A sends the NLRI to its peers B and D. B and D are peers of node A and on the FT. A does not send the NLRI to its peer C since C is not connected to A on the FT.

After receiving it from A, B sends it to C; D sends it to C.

The number of NLRIs flooded in revised flooding is much less than that in normal flooding.

# Next Step

Comments

Adoption