Proposed Update to BGP Link-State SPF NLRI Selection Rules

draft-dong-lsvr-bgp-spf-selection-01

Jie Dong @Huawei

Jinqiang Chen @Huawei

Sheng Fang @Huawei

LSVR WG IETF 121 Meeting Nov. 2024

Motivation

- BGP-LS SPF is designed for link-state information distribution and SPF path calculation in MSDC scenarios
  - Leverages the mechanisms of the BGP base protocol and BGP-LS extensions
- The NLRI selection rules for all BGP-LS SPF NLRIs are defined as below:
  1. NLRI originated by directly connected BGP SPF peers are preferred
  2. The NLRI with the most recent Sequence Number TLV, i.e., highest sequence number is selected
  3. The NLRI received from the BGP SPF speaker with the numerically larger BGP Identifier is preferred
- In some cases, these rules may not be good enough for efficient route exchange and convergence
  - Redundant route advertisement may cause delayed route convergence
  - Consistent SPF computation result can be achieved with relaxed NLRI selection
- This document describes the problem scenarios, and proposes updates to the selection rules of BGP-LS SPF NLRI
  - Some problem scenarios in -00 version has been resolved in the updated BGP-SPF base document

Problem Scenario 1

- A new BGP-SPF session is established between R1 and R6
- R1 advertises the link NLRI R1-R6 to its neighbors (R2 and R3)
- R2 firstly receives the link NLRI R1-R6 from R1 directly, and advertise it further to its neighbors (R4 and R5)
- R4 receives the link NLRI R1-R6 with the same sequence number from both R3 and R2, and **prefer the one from the peer with larger BGP ID (R3)**
- R4 advertises link NLRI R1-R6 to R2
- R2 prefers the link NLRI received from the peer with larger BGP-ID (R4)
- R2 advertise the link NLRI received from R4 to R5, which is **a redundant advertisement of the same link NLRI**, as it does not impact the result of SPF computation

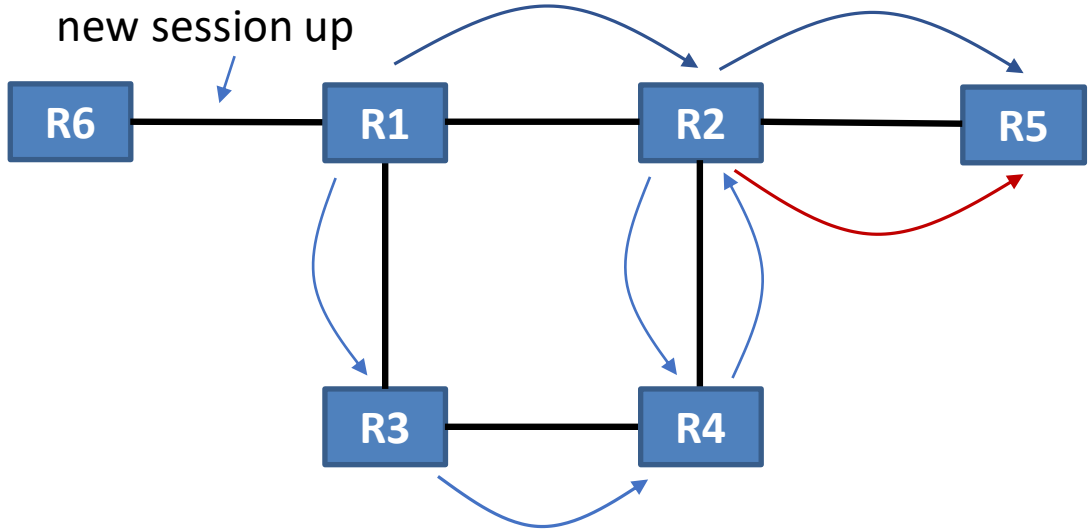

Problem Scenario 2

- There are two parallel links between R1 and R2, on each link a separate BGP session is established
- For the same NLRI with the same sequence number received from R1 via different sessions, the current NLRI selection rule cannot determine which one is the preferred route
- No matter what tie-breaking rule is used, depending on the order of the routes received from R1, R2 may need to advertise duplicated BGP-SPF routes to R4.

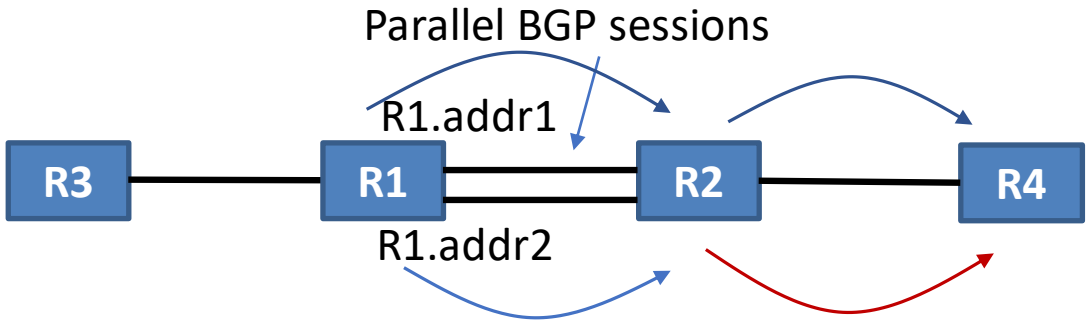

Proposed Updates to NLRI Selection Rules

- The assumption is that for BGP-SPF NRLIs with the same sequence number, the information used for SPF computation would be the same, thus additional NLRI selection and update caused due to other attributes are considered not necessary from the perspective of consistent SPF computation result
- This document proposes to update the selection rules for all BGP-LS SPF NLRI as follows
    1. NLRI originated by directly connected BGP SPF peers SHOULD be preferred.
    2. The NLRI with the most recent Sequence Number TLV, i.e., highest sequence number SHOULD be selected.
    3. NLRI received from a BGP-LS-SPF peer with the same sequence number as the one of the current selected NLRI SHOULD not be selected.
- The new rule 3 can help to solve the duplicated advertisement problem

Next Steps

- Suggest to make this an optional optimization to BGP-SPF NLRI selection rules
  - Can be useful in network scenarios where redundant advertisement is a bigger concern comparing to deterministic route selection
- Comments and feedbacks are appreciated

Thank You