Avoiding Traffic Black-Holes for Route Aggregation in IS-IS

draft-chen-isis-black-hole-avoid-02

Zhe Chen, Xiaohu Xu, Dean Cheng

Huawei

2018.3

Changes since Version -00

- Added Dean Cheng as co-author
- Introduce INFINITE cost, thus reusing existing TLV instead of creating a new one

Problem Description

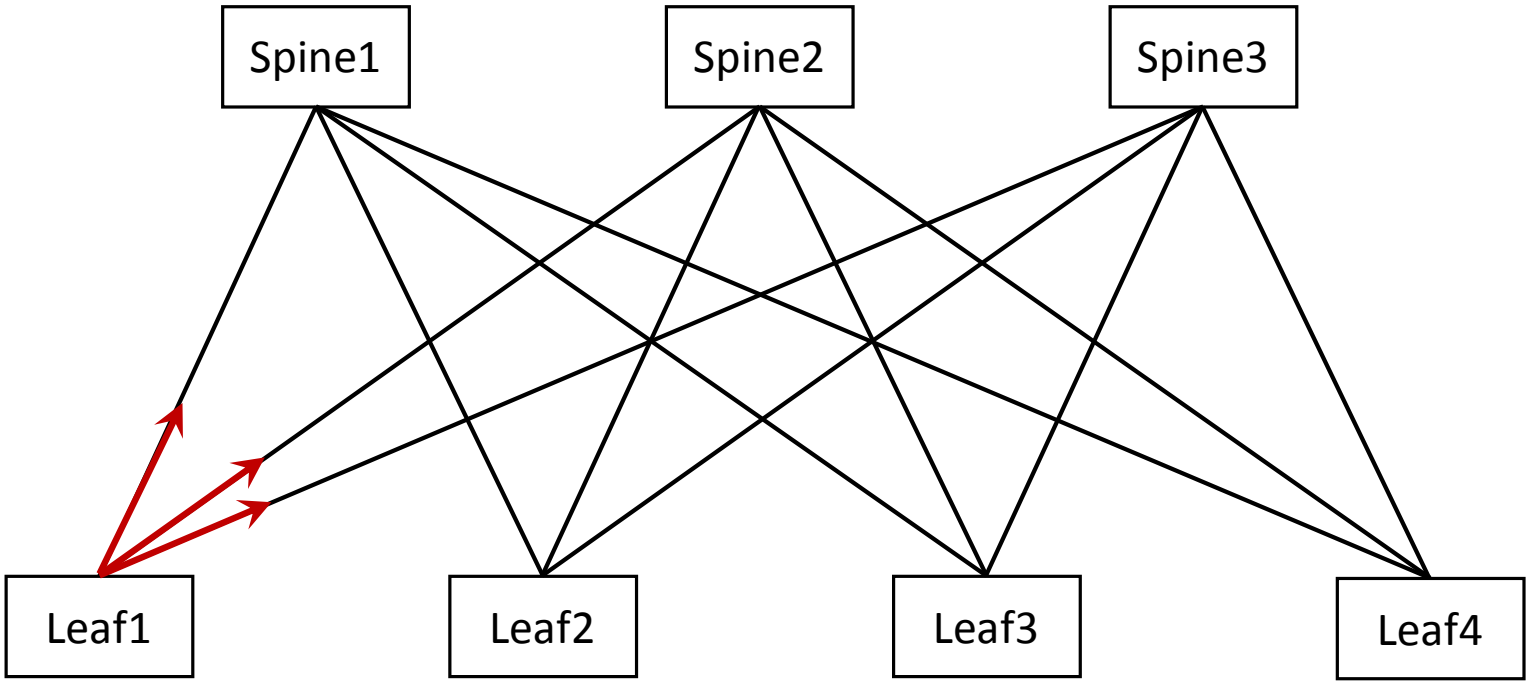

Dst	Nexthop
0.0.0.0/0	Spine1 Spine2 Spine3

PrefixA
PrefixB

Problem Description

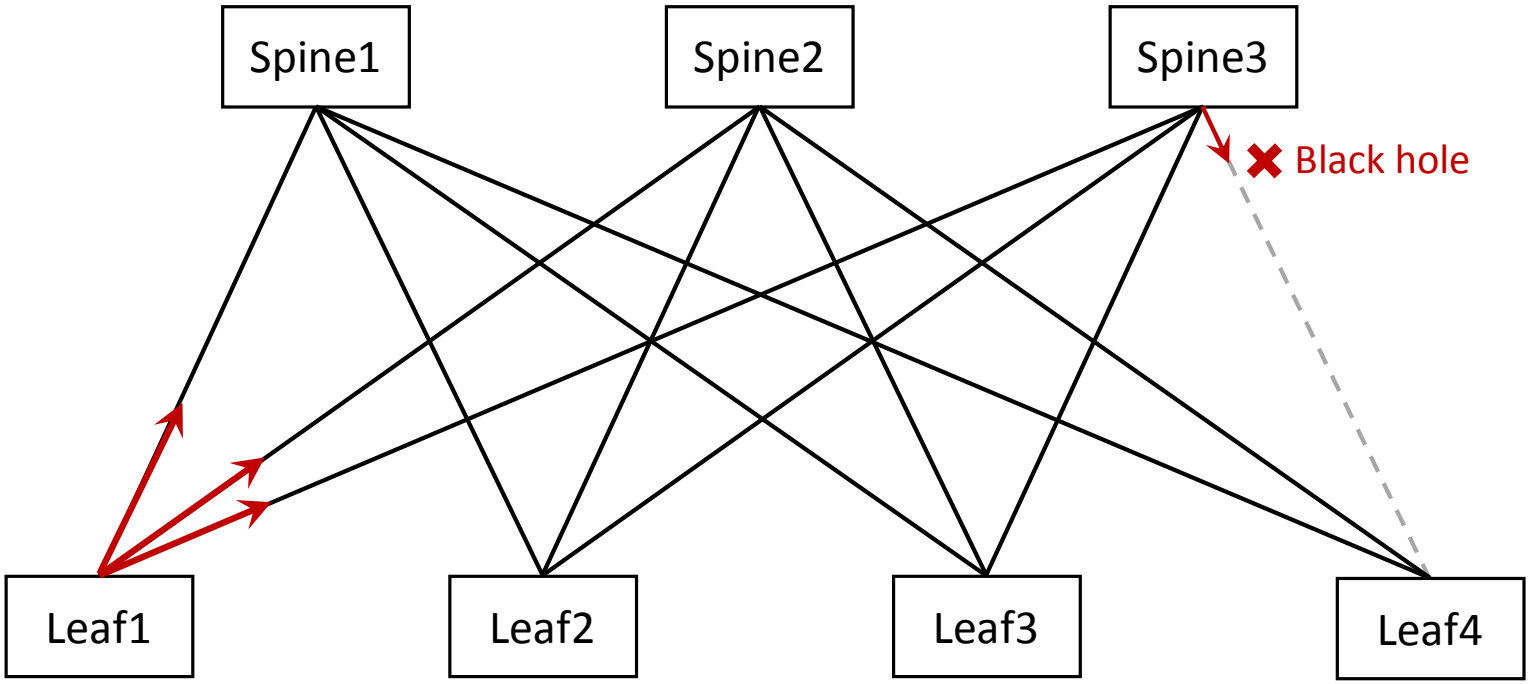

Dst	Nexthop
0.0.0.0/0	Spine1 Spine2 Spine3

Solution Overview

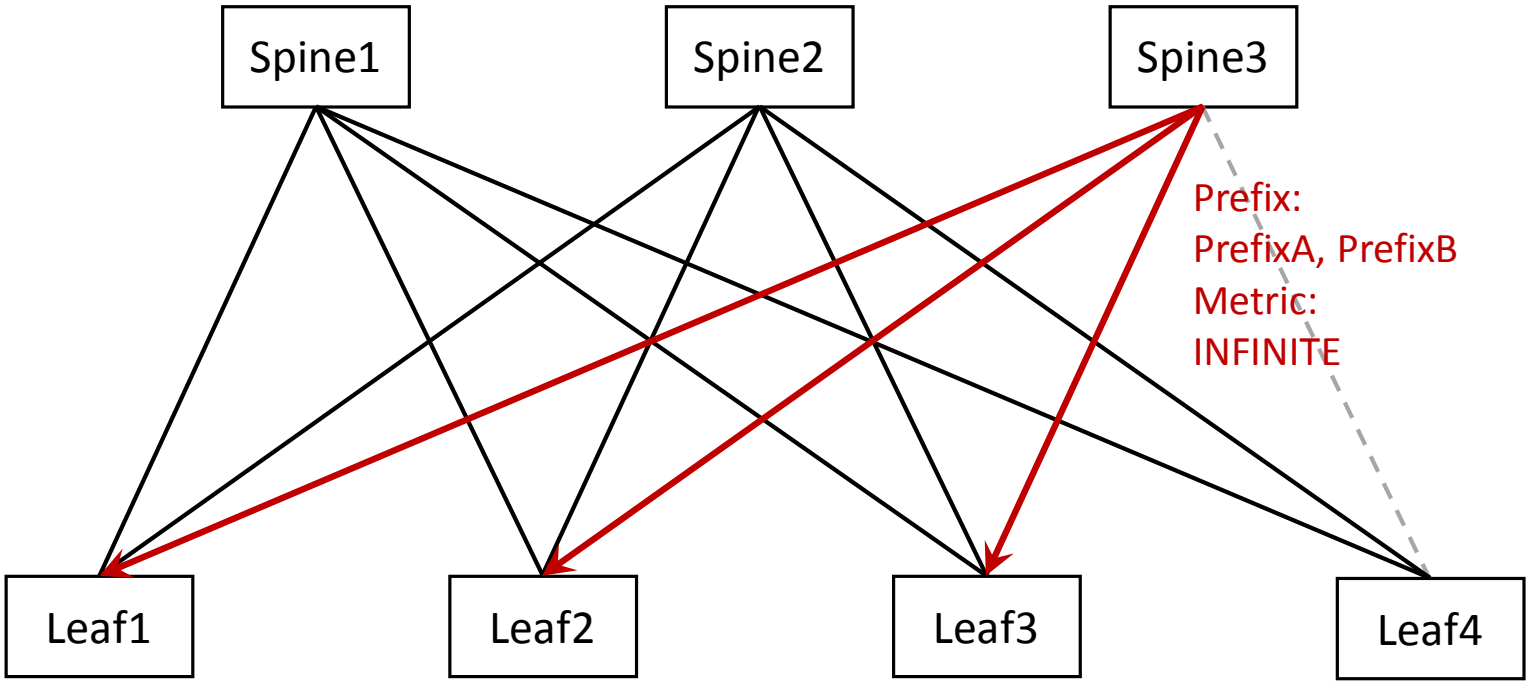

Dst	Nexthop
0.0.0.0/0	Spine1 Spine2 Spine3

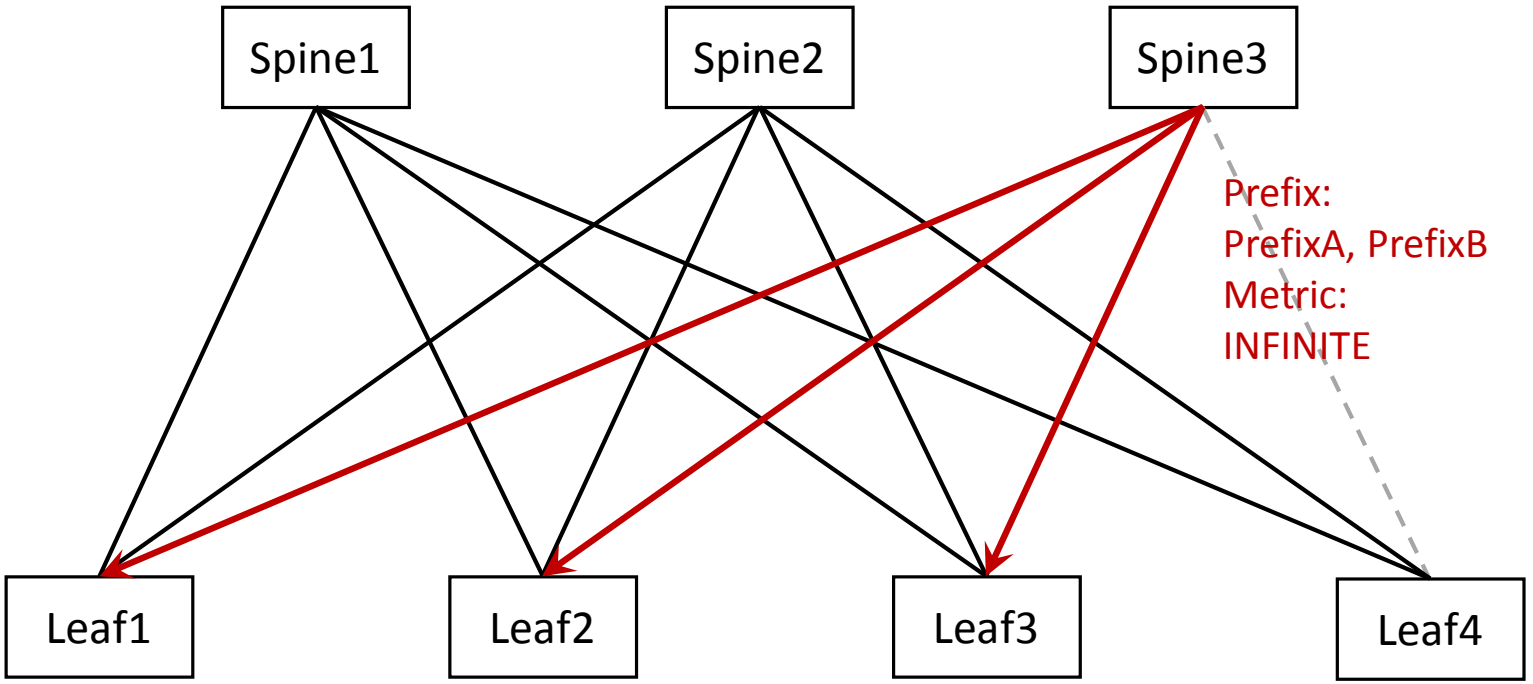

Dst	Nexthop
0.0.0.0/0	Spine1 Spine2 Spine3

Dst	Nexthop
0.0.0.0/0	Spine1 Spine2 Spine3
PrefixA	Spine1 Spine2
PrefixB	Spine1 Spine2

Details

- When link failure happens between a Spine node and a Leaf node, the Spine node SHOULD
  - 1) encode all prefixes attached to the Leaf node into the IP Reachability TLV,
  - 2) set the cost of the prefixes to be INFINITE,
  - 3) append the IP Reachability TLV to the IS-IS LSP, and
  - 4) send the LSP to every other Leaf node it connects to.
- The Leaf node SHOULD:
  - 1) install each of the prefixes into its routing table,
  - 2) set the next hop to be an ECMP group including all Spine nodes it connects to, except the one who advertises the prefix.

Next Steps

- We need more reviews and comments.