IETF 101 – London
Mar 2018
LSR Working Group

draft-ketant-ospf-reverse-metric-00

Ketan Talaulikar (ketant@cisco.com)
Peter Psenak (ppsenak@cisco.com)
Hugh Johnston (hugh_johnston@labs.att.com)

What does this draft propose?

- Enables an OSPFv2/v3 router to signal its neighbour the “reverse-metric” that the neighbour should use on the link toward itself

- R1 does reverse-metric signalling towards its neighbour R2 via link-local mechanism
- The neighbour R2 modifies the metric in its Router-LSA for its link to R1
- Only R1 and R2 need to support this draft; other routers are not involved and will start using the updated metric on the link from R2 → R1

Use-case 1

- Many OSPF deployments use “auto-cost” based on reference bandwidth
- Metric are symmetrically setup in both directions on a link based on bandwidth
- Link capacity is managed at layer-2 as bundle members including over layer-2 aggregation networks

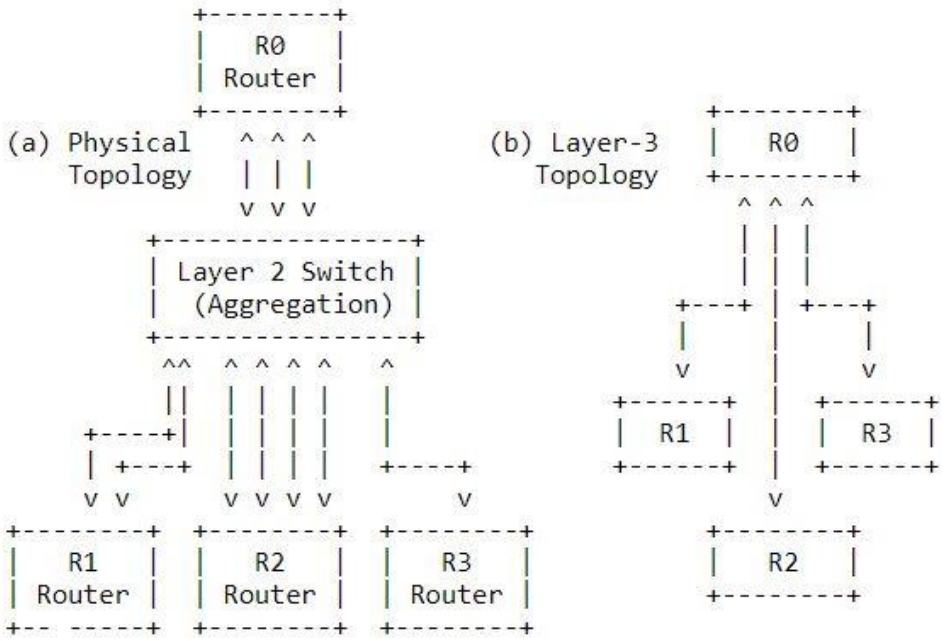

- Link down or capacity degradation in one direction is not detected and reflected in OSPF metric in the other direction
- Reverse-Metric mechanism allows the router detecting the capacity degradation to signal metric for the link to the neighbor

Use-case 2

- AGGR routers are aggregating traffic from multiple downstream routers towards the core
- Based on the uplink bandwidth or congestion considerations, AGGR1 may want to offload upstream traffic coming to it from R1 to its pair AGGR2

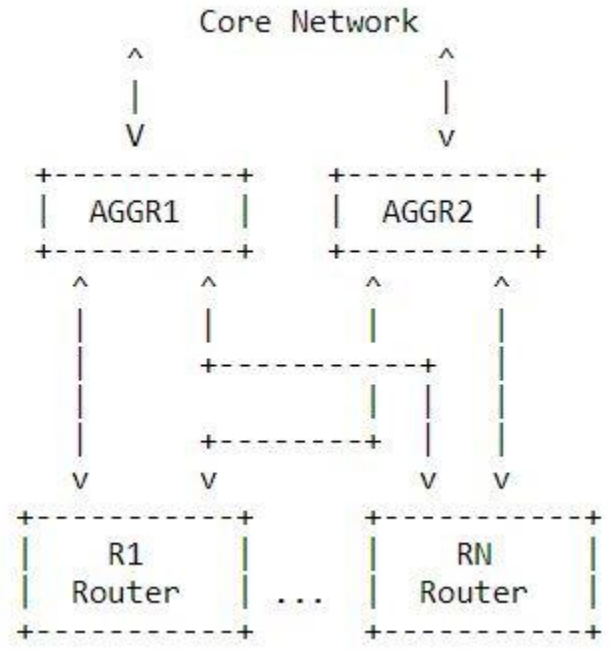

- Reverse-Metric mechanism allows AGGR1 to signal R1 to increase its metric towards AGGR1 such that the path via AGGR2 becomes the preferred path.

How is this done?

- New Reverse Metric TLV is introduced for Link-Local-Signalling

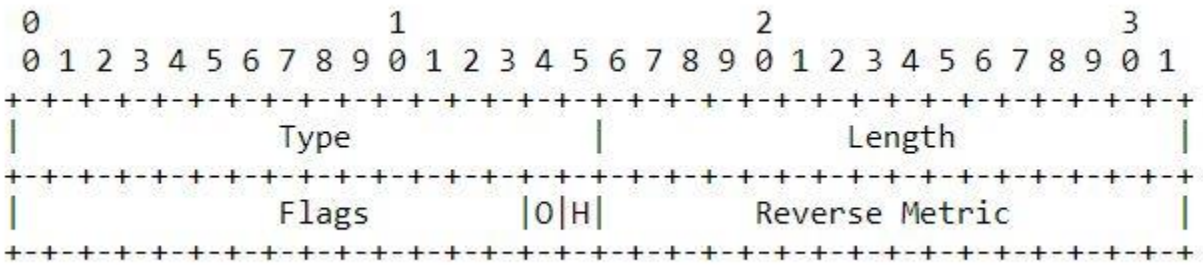

- O-bit : reverse metric value is an offset to be added to existing original metric by receiver
- H-bit : the absolute value of reverse metric is to be used only when larger than the existing original metric by receiver
- TLV is included in the Hello message while the reverse-metric value is to be signalled

Other aspects …

- This draft applies to point-to-point and point-to-multipoint links only.
- RFC 8042 (OSPF Two Part Metric) needs to be used for broadcast and NBMA
- Unlike Two Part Metric, this mechanism local to the link involved and has no backward compatibility challenges
- The mechanism to determine the value of reverse-metric to be signalled and to trigger this signalling is outside the scope of the draft and implementation specific aspect (e.g. via management station or controller, or other detection mechanism on the router, etc.)

Comparison with draft-ietf-isis-reverse-metric

- ISIS draft proposes to handle multiple scenarios using reverse-metric which are done differently by existing OSPF documents
  - RFC 8042 – OSPF two part metric : handles the LAN scenarios
  - draft-ietf-ospf-link-overload : handles link-overload signalling
  - this draft handles the other generic use-cases where a router needs to locally signal a reverse-metric to its neighbour
- OSPF reverse-metric draft allows reverse-metric to be signalled as an absolute value OR as an incremental offset

Next Steps …

- Solicit WG review and comments/inputs/feedback