

OSPF Two-part Metrics

Jeffrey Zhang

Juniper Networks

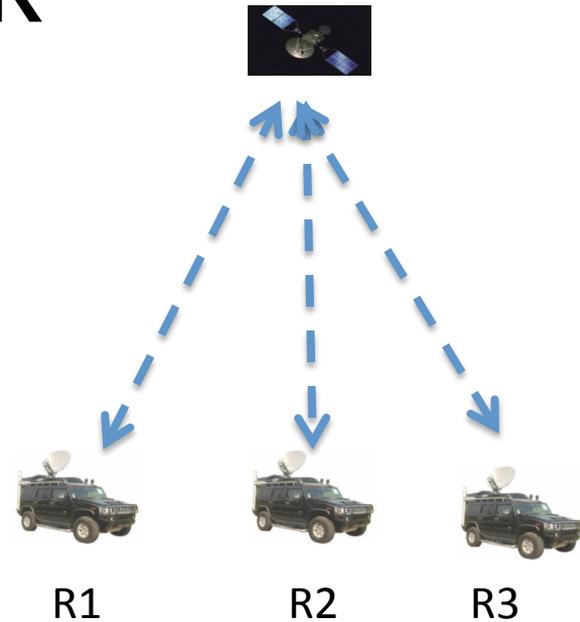
90th IETF, Toronto

Summary

- -00 revision presented in 88th IETF, Vancouver
 - Problem statement and proposed solution
- -01 revision addressed compatibility issue and filled in missing details
 - Presented in 89th IETF, London
 - With two options of encoding from-network cost and ensuring backward compatibility
- This presentation:
 - Reviews general idea
 - Presents the chosen option and details in -02 revision
 - Request WG adoption

An example network

- Satellite based, with fixed/mobile routers
 - A true broadcast network, not MANET
- Different costs between different pairs
 - Conventionally modeled as p2mp network



R1->R2: 10

R1->R3: 20

R2->R1: 15

R2->R3: 25

R3->R1: 20

R3->R2: 25

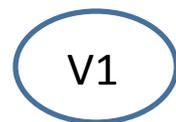
RFC 6845 Hybrid Interface

- Treat as broadcast for Hello, adjacency and database synchronization purpose
- Treat as p2mp to advertise different costs for different neighbors
 - Each Router LSA has N-1 links for the interface
 - One for each neighbor
 - Change in one router's communication capability causes all routers to update their Router LSAs
 - unbearable flooding in a large network with routers constantly moving around

Proposed solution

- Model as broadcast network
 - Do not advertise p2p links
- Break router to router cost to two parts: to/from-network
 - Advertise from-network cost in a stub-link for the network
- SPF calculation to consider both to- and from-network costs

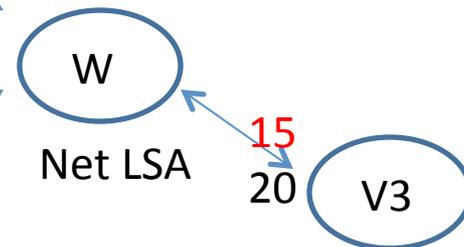
Router LSA for R1



Black: to-net
Red: from-net



Router LSA for R2



Router LSA for R3

Router-Router Cost calculation:

$$R1 \rightarrow R2: 10 + 10 = 20$$

$$R1 \rightarrow R3: 10 + 15 = 25$$

$$R2 \rightarrow R1: 10 + 5 = 15$$

$$R2 \rightarrow R3: 10 + 15 = 25$$

$$R3 \rightarrow R1: 20 + 5 = 25$$

$$R3 \rightarrow R2: 20 + 10 = 30$$

Advantages

- Full benefit of broadcast network model
 - Hello, Adjacency, Synchronization
 - Reduced LSA size ($2N$ vs. N^2)
 - One link in each Router LSA, N link in the Net LSA
 - Vs. $(N-1)$ link in each Router LSA
- Reduced update frequency
 - When one router's communication capability changes, only its own Router LSA needs update (for OSPFv2)
 - For OSPFv3, that router's Intra-Area-Prefix-LSA also needs update
- Generally applicable to ANY broadcast network
 - As long as cost can be logically broken into two parts
 - Not a drastic change in concept from RFC 2328:
 - Section 2.2: "Edges that are not marked with a cost have a cost of zero (these are network-to-router links)"
 - Section A.4.3: "The distance from the network to all attached routers is zero"

Encoding from-network cost: In a stub link (OSPFv2)

- Thanks to Acee for the idea
- In addition to the type-2 transit link, advertise a type-3 stub link to the network, with cost set to the from-network cost
- During SPF calculation, the presence of both a transit and stub link to the network indicates the network is using two-part metric, and router-router costs are calculated accordingly
- The stub link is skipped in the second stage of SPF calculation (for stub networks)

Encoding from-network cost: OSPFv3

- For OSPFv3, there is no stub-link in Router LSA but there are separate Intra-Area-Prefix-LSAs
 - advertised by DR for transit networks and by everyone for stub networks
- If a DR includes a transit link in its Router-LSA for a two-part-metric network, then for the associated prefix in its Intra-Area-Prefix-LSA, the metric is set to its from-network cost (vs. 0)
- For a non-DR, include the associated prefix in its Intra-Area-Prefix-LSA
 - Metric set to its from-network-cost
 - Reference set to the Transit Link advertised by the DR
 - Referenced LS Type set to 0x2002
 - Referenced LS ID sets to DR's InterfaceID
 - Referenced Advertising Router set to DR's Router ID

An optional optimization for OSPFv3

- Encode all from-network costs in DR's Intra-Area-Prefix-LSA
 - If the underlying network already has a mechanism to communicate those costs to DR
 - Non-DR's from-network costs encoded in additional prefix entries following the prefix entry that the DR normally includes already
 - Referenced Advertising Router set to those Non-DR
- This further reduces flooding in affect-all events
 - e.g. large-area rainstorm affecting a satellite based network

SPF calculation

- When a network LSA is examined, for each listed neighbor, check if there is advertised from-network cost
 - OSPFv2: check for the stub-link for the transit network in the neighbor's Router LSA
 - OSPFv3: check for corresponding prefixes in DR's Intra-Area-Prefix-LSA first. If not found, check the neighbor's Intra-Area-Prefix-LSA
- The stub-link or prefixes described above are ignored when calculating routes to stub networks

Compatibility

- All routers supporting this capability must advertise Router Information (RI) LSA with a newly assigned bit set in *Router Information Capabilities TLV*
- All routers must disable two-part metric when detecting the presence of a reachable Router LSA w/o a companion RI LSA w/ the bit set:
 - Remove its stub link for a transit network
 - Recalculate routes w/o considering the from-network cost

Next steps

- Draft is done – no further update expected
 - Except for potential changes warranted from the WG review
- Seek WG adoption