

OSPF Two-part Metrics

Jeffrey Zhang

Juniper Networks

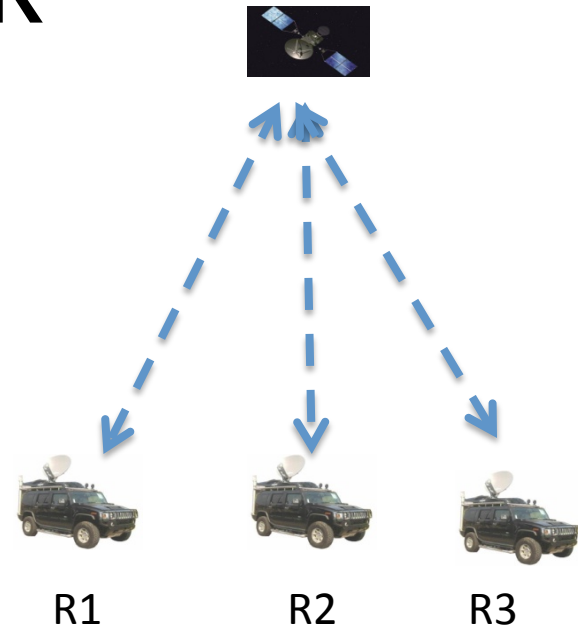
89th IETF, Landon

Summary

- -00 revision presented in 88th IETF, Vancouver
- -01 revision addresses compatibility issue and finishes missing details
- This presentation:
 - Reviews general idea
 - Presents updates in -01
 - Presents an alternative encoding for the from-network cost

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

Observation 1

- If one router's communication capability changes, all costs change in a similar fashion:
 - other routers' cost to it
 - its cost to all other routers



R1 R2 R3

R1->R2: 10
R1->R3: 20 + 150

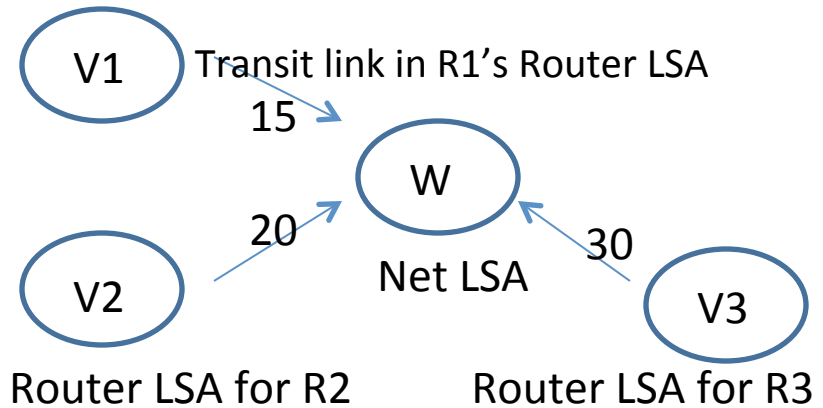
R2->R1: 15
R2->R3: 25 + 150

R3->R1: 20 + 100
R3->R2: 25 + 100

Observation 2

- Network LSA does not have costs associated with listed routers
 - It is assumed that a router's cost to all neighbors are the same
 - encoded in the transit link in Router LSA
- Note that different routers can still encode different costs in the transit link of their own Router LSA for the same network

Router LSA for R1



SPF calculation result:

R1->R2,R3: cost 15



R2->R1,R3: cost 20

R3->R1,R2: cost 30

Proposed solution

- Model as broadcast network
 - Use Router and Net LSAs instead of advertising p2p links
- Break router to router cost to two parts: to/from-network
 - Advertise both in the transit link (of a different type X) in Router LSA
 - Encode from-network cost as an MT cost
 - Network LSA as is
- SPF calculation to consider both to- and from-network costs

Router LSA for R1




Router LSA for R2



Net LSA

Black: to-net
Red: from-net



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
- 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:

Option 2 - encode in a stub link

- 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)

Compatibility: Encoding Option 1

- New Link Type X not backward compatible
- Routers not supporting this must not form adjacencies with those that do support it, in an area with two-part metric networks
 - Use LLS EOF-TLV for this purpose

Bit	Name	Reference
0x00000001	LSDB Resynchronization (LR)	[RFC4811]
0x00000002	Restart Signal (RS-bit)	[RFC4812]
0x00000004	Two-Part Metric (TM-bit)	[this proposal]

Compatibility: Encoding Option 2

- All routers supporting this capability must advertise Router Information (RI) LSA with a newly assigned bit set in *Router Informational Capabilities TLV*
- All routers 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

Plan

- Update draft with new encoding scheme
- Seek review, comments, and WG adoption