

Routing in Fat Trees (RIFT) Update draft-przygienda-rift-02

Tony Przygienda, Juniper
Alankar Sharma, ComCast
Alia Atlas, Juniper
John Drake, Juniper

Progress since -01

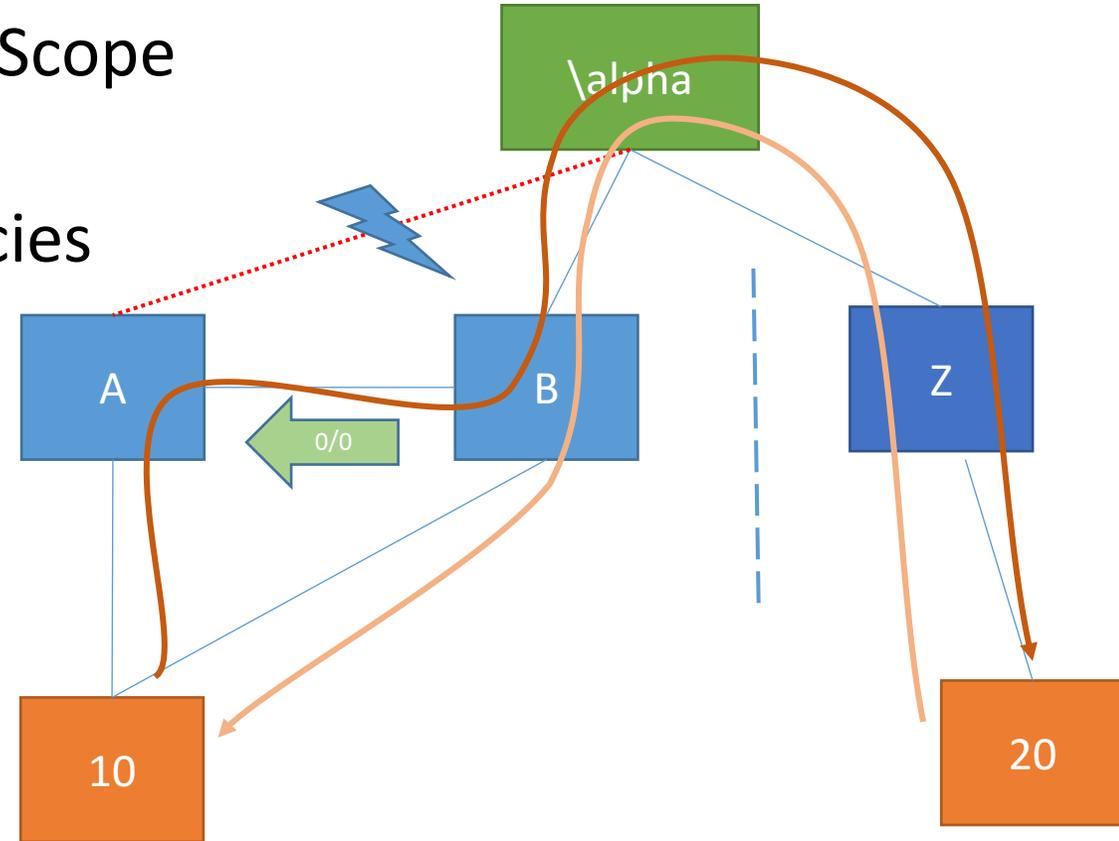
- New Co-Author
- Good Amount of Industry Input and Real ;-) Experience Gained
- Significant Clarifications/Extensions on the Specification
 - Requirements Extended
 - Optional E-W Link Procedures Added
 - Internal Node Reachability Considerations Added
 - More Explicit Description of Flooding Scopes Added
 - Description of Default Origination Procedures Added
 - BFD Interactions Added
 - Encoding Improved for More Efficient Implementations

Requirements Extension

- RIFT Distinguishes Parallel, Non-Uniquely Addressed and Unnumbered Links. LIEs Are Exchanges over All Links, i.e. Native IP “LAG” is Possible
 - Flooding Optimization over Parallel Links is Left Unspecified
- MC-LAG Considered Non-Requirement = IP Multi-Homing
- No Special Support for Forwarding Adjacencies and so on is Necessary
- An Optional Mode Allowing Basically a Near Zero Touch Configuration of a Whole DC is Desired

Optional E-W Procedures

- E-W Flooding Rules Same as Flooding Scope Towards South Peers
- During SPF a Node Uses E-W Adjacencies IIF
 - It Has No Northbound Itself AND
 - Neighbor Has Northbound Adjacencies
- One-Hop Split-Horizon Black-Hole Healing
- E-W NOT Used Unless Failures Occur



Internal Fabric Node Reachability Considerations

- New Section Added
- Nodes Can Inject Loopback Into N-Prefix TIEs for Reachability “From the North” or “From Everywhere under Normal Conditions”
- Nodes Can Inject Loopbacks Always or on North Connectivity Failures Into S-PGP TIEs for Reachability “From the South”

Precise S-Prefix Default Origination Procedures

(Spoiler Warning: Not Entirely Trivial)

Node X That Is

- NOT Overloaded AND
- Has East-West OR South Adjacencies

Originates S-Prefix TIE Default IIF

- All Other Nodes @ Same Level Are Overloaded OR Have No N-Adjacencies OR
- X Has Computed Default Reachability in N-SPF

BFD Interactions

- 3-way Hello MAY bring up BFD Sessions
- BFD After Being Up Can Drive the 3-way Hello Down
- RIFT Link IDs Are Aligned Size-Wise with BFD Discriminators so They MAY be Same
- BFD MAY Run on Any Subset of RIFT Links

Encoding Improvements

- LIE Matches Up LinkIDs on Both Sides
- Several Encoding Elements Became <maps> Instead of <sets> for Fast Searches

THANK YOU FOR YOUR ATTENTION