LFA Manageability

draft-litkowski-rtgwg-lfa-manageability-00

Stephane Litkowski	Orange
Bruno Decraene	Orange
Clarence Filsfils	Cisco Systems
Kamran Raza	Cisco Systems

Agenda

- Goals
- Some operational issues with LFA selection
- A call for policy based LFA selection
- Some additional operational aspects

Goals

- Discuss LFA management
  - Missing point highlighted during IESG review of RFC 6571(LFA applicability in SP networks)
  - Asked by rtgwg chair (Alia 2012/01/18)
- Provides feedback following LFA deployment
- Highlights some limitations
- Call for some improvements

Issue 1: PE used to protect a P

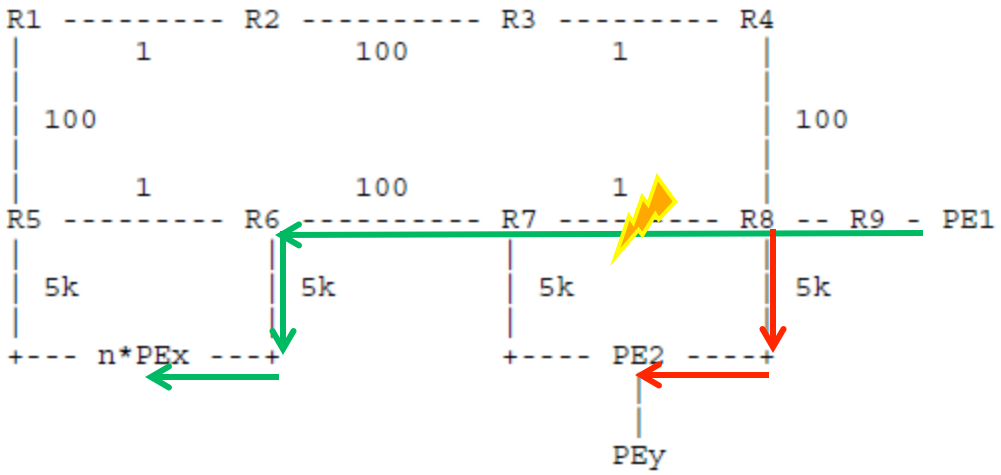

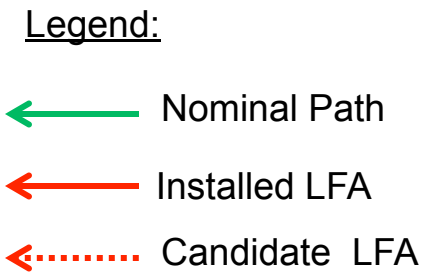

- Edge node / edge link used as a protection for a core link / node
  - because no other LFA was available
- Routing policy issue: edge node used to route core traffic.
- (link) capacity issue:
  - PE2 was not impacted by the original failure but becomes congested following LFA activation.

Issue 2: PE selected as best LFA to protect a P

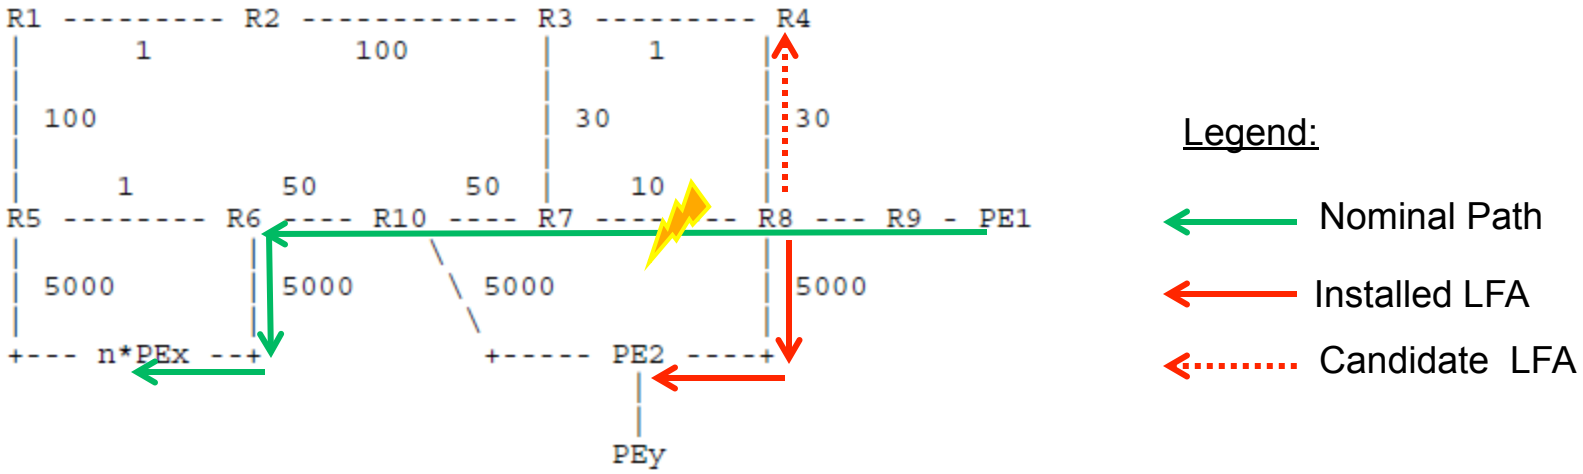

- Edge node / edge link used as a protection for a core link / node
  - because PE2 is node protecting while R4 is link protecting
- Routing policy issue: edge node used to route core traffic.
- (link) capacity issue.

Issue 3: low bandwidth link used

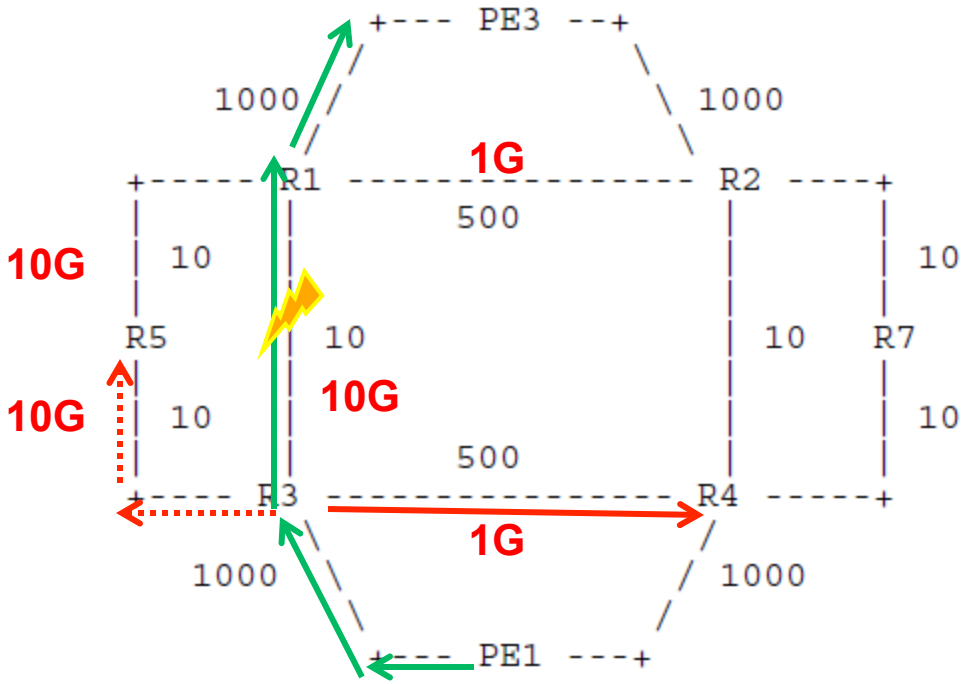

Legend:

Nominal Path

Installed LFA

Candidate LFA

- Low bandwidth link used as a protection for high bandwidth link
  - because R4 is node protecting while R5 is link protecting

(link) capacity issue.

Issue 4: high cost/delay link selected as LFA

- 4 neighbors are candidate LFA, for the failure of link CORE1-CORE2.
- PE2 is selected as best LFA and installed, while it's an oversea PE.
- CORE3 would be the preferred choice.

Selected & installed

Link protected	Destination	Alternate	Type/ metric
CORE1 -> CORE2	PE1	PE2	node protect /260000
		PE3	node protect /270000
		PE4	node protect /280000
		CORE3	Link protect /200000

Oversea PEs …

Meant to be the best LFA

Calling for a policy based LFA selection

- Current tie-breakers for selecting the LFA are not flexible enough to accommodate for all cases.
- Calling for a policy based LFA selection, controlled by the SP according to local constraints
- Multiple criteria expected:
  - Level of protection: node, link, srlg, local srlg
  - Type of LFA: primary, downstream, LFA
  - IGP metric to destination
  - Link coloring / node coloring (e.g. core, edge, core&edge)
  - Link info (a la TE): affinity, speed, available bandwidth, delay
  - Connectivity toward the Merge Point: link, tunnel (rLFA), TE-tunnel
    - computes rLFA even if LFA exist as rLFA may be preferred.
- Applied per protected interface or set of destinations.
- More details in §3.2 “Policy based LFA selection”

Example of a policy using link coloring

- Marking:
  - PE links as RED
  - 10G CORE links as BLUE
  - 1G CORE links as YELLOW
- LFA Policy:
  - Include BLUE, preference 200
  - Include YELLOW, preference 100
  - Exclude RED
- Result:
  - assuming all routers are candidate LFA
  - P2 is selected as best (non PE, 10G interface):
  - PE links not used to protect core links
  - 10G links preferred over 1G links

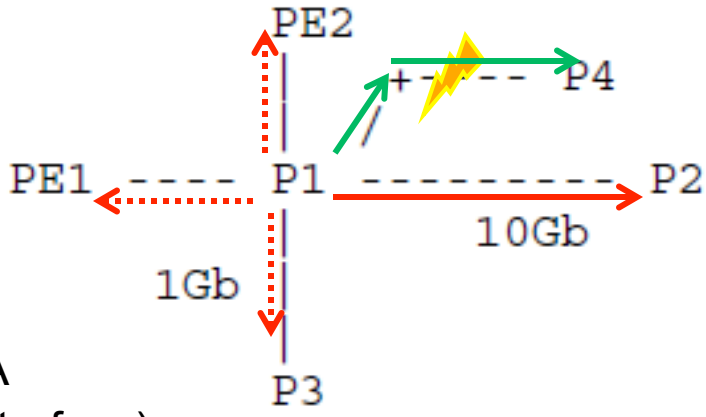

One more thing…

- LFA activation granularity
  - per address-family, per routing context, per interface, possibly per prefix.
- Controlling LFA computations
  - a la SPF delay / back off algorithm
  - abort LFA computation if an IGP SPF is scheduled
- Checking coverage
  - show coverage per IGP domain (area/level, topology, instance, virtual router), per protected link, possibly per prefix priority group
  - show non protected prefix, possibly per prefix priority group
    - providing the reason (e.g. rejected by policy),
  - alert/log if coverage falls below a threshold
- Checking LFA selection
  - show installed LFA & candidates LFA
  - per prefix, per interface.
  - provides the reason for selecting the LFA

Next steps

- Some first comments received
  - -01 being edited
- Soliciting more comments
  - Additional cases/issue found during deployment
  - Improvements
  - ….

thank you