

draft-litkowski-rtgwg-uloop-delay-03

Stephane Litkowski, Orange

Bruno Decraene, Orange

Clarence Filsfils, Cisco Systems

Pierre Francois, IMDEA Networks Institute

IETF 89, RTGWWG WG

Diffs


- Motivation of the solution
- More precise definitions (local event, ...)
- More precise focus (link down and link up)
- Previous works

Microloop prevention by introducing a local convergence delay

- Microloop
 - transient inconsistency among routers' FIB
 - during a convergence event
- Local convergence delay
 - Avoid inconsistencies
 - In the direct neighborhood of the node responsible of the convergence
 - by controlling when the node updates its FIB


Motivation

- Benefits of FRR is negated by microloops
 - if PLR directly goes into FIB update after FRR activation
- Similar issues when the link comes back up
- Microloops are bad:
 - Packet drops (TTL)
 - Link saturation
(more drops, even for “unaffected” destinations)




Link down

- Link down $A \rightarrow D$, A does FRR using RLFA to C
- A
 - Floods LSP
 - Delays its FIB update
- Other nodes (B) update first



Link up

- Link A-D comes back up
 - A floods LSP with largest metric (2 way CC)
 - A updates its FIB, (SPF using configured metric)
 - A floods LSP with configured metric
- A updates its FIB before other nodes (B)



Uloop-delay, oFIB, and PLSN

- oFIB
 - Orders FIB update among all routers of the net
 - Requires interoperability and full support
 - uloop-delay is a subset:
 - only deal with local inconsistencies
- PLSN
 - Delays FIB update while transiently use loopfree neighbours
 - uloop-delay is different:
 - Local only
 - Traffic does not transiently follow non shortest-paths

Possible future items

- Metric reconfiguration (easy)
- Node event
- LAN's