

IGP Fast Reroute

Cengiz Alaettinoglu
cengiz@packetdesign.com

Alex Zinin
zinin@psg.com

Slides: <http://www.packetdesign.com/publications>

Convergence Framework

- Link failure detection
 - Hardware detection (preferred) or hello protocols
 - Damping at recovery of flapping links
- Reaction to failure
 - Switch to precomputed detours, MPLS or IGP
 - Generate an LSP
 - Trigger an SPF
- Propagation of the news
 - LSP gets flooded across the network
- Convergence to new routes
 - SPF completes and new routes are installed

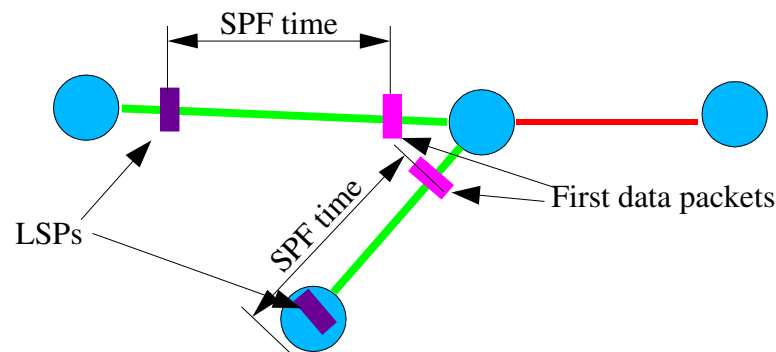


IGP Convergence Times

- Vendors have achieved sub-second convergence
 - Filsfils, <http://www.nanog.org/mtg-0206/filsfils.html>
 - Katz, <http://www.nanog.org/mtg-0206/katz.html>
- And working on improving several aspects of it
 - draft-kompella-rag-plp-00.txt
 - draft-retana-marp-01.txt
 - ...

How fast can IGP convergence be?

- *Few milliseconds is possible in theory*
 - *sub-100 milliseconds can be expected in routers*
- Single SPF time + network diameter x scheduling delays
 - Incremental SPF in tens to hundreds of microseconds
 - Requires LSP is transmitted before data packets and SPF
 - LSP is ahead of any data packet by SPF time
- Requires fast route install
- Requires small CPU scheduling delays
 - *Really good implementation*



Road Blocks

- Fast hardware failure detection
 - Works today for at least POS & DWDM interfaces
- Prioritizing control traffic transmission over data
 - Deployment issues remain
- Fast route install
 - Vendors are are working on two level FIB lookup, first one for BGP and then for IGP
- Small scheduling delays
 - *This is hard and will take time*

Reactive Solutions for the Interim


- MPLS fast reroute
 - Some ISPs have reported success
 - Rao et. al. <http://www.nanog.org/mtg-0210/shankar.html>
 - Protection against single link failure
 - Protects core to core links
 - Due to MPLS full mesh complexity
- IGP fast reroute
 - Can we do the same for IGP, i.e. switch next hops before SPF?

IGP Fast Reroute

- SPF computes a set of next hops
 - Goal is to have at least two next hops
 - *Shortest* and if necessary *feasible* next hops
- Install two classes of next hops in line cards
 - Primary is always shortest
 - Secondary is either shortest (ECMP) or best feasible next hop
- When the primary fails, switch over to the secondary
 - This switching is done in the line cards
 - The secondary is only used while SPF is being computed

Feasible Next Hops

- A neighbor whose distance is less than yours
 - ISIS: downstream router
 - Section 7.2.6.2, see Parker/Perlman/Katz' emails on ISIS list on May 2001
 - EIGRP: feasible successor :-)
- A neighbor whose path does not contain you, i.e. loop free
 - Relaxes the shorter distance requirement

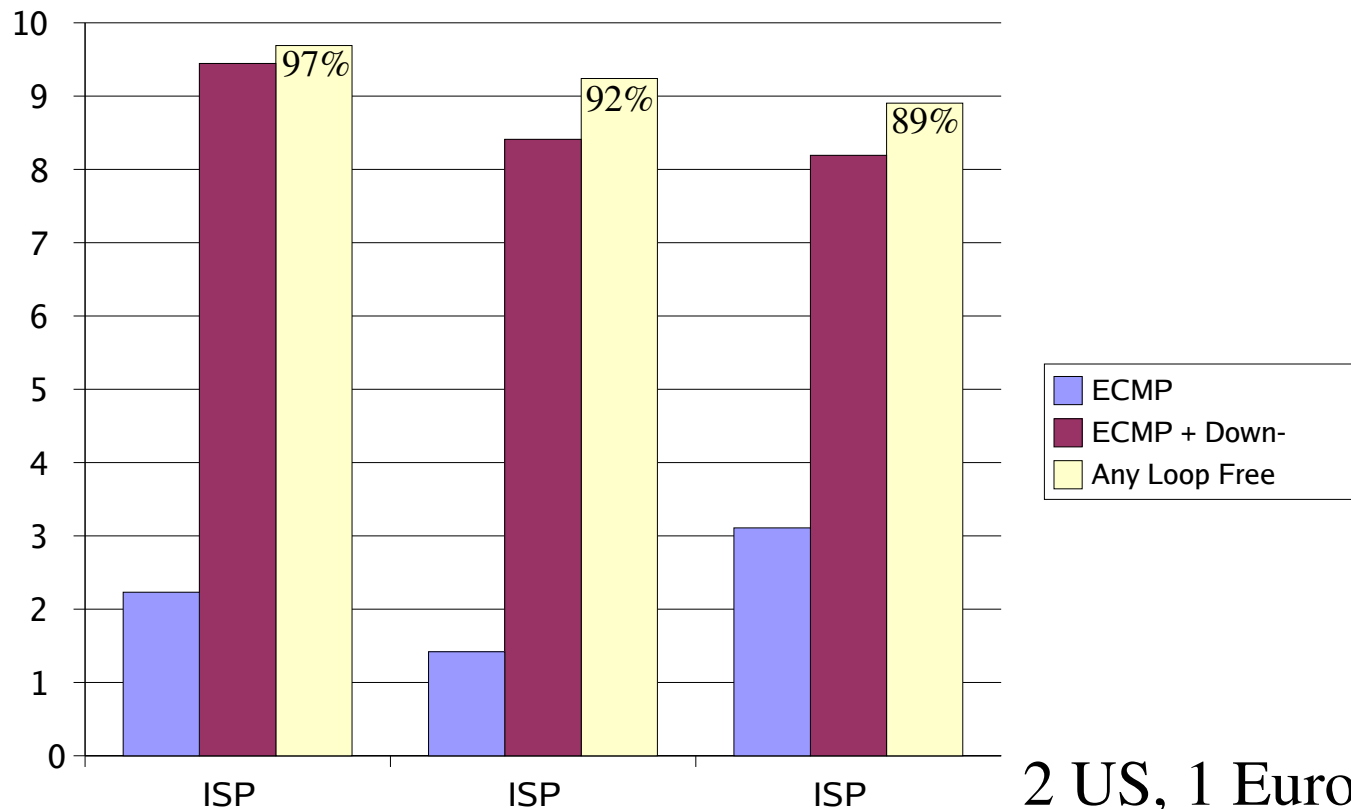


Are Feasible Next Hops *Feasible*?

- Most ISP networks have redundancy built-in
 - Border routers connected to multiple core routers
 - *Send to the other core router when next hop one fails*
 - Multiple connections across POPs
 - Hence, we expect to see feasible next hops

Availability of Feasible Next Hops

% of Source-Sink Pairs w/ Feasible Next Hops



2 US, 1 European, Tier 1 ISPs

- This is not % of links protected
→ That should be higher

Fast Switching in Line Cards

- Like MPLS FRR, trigger changes on line cards directly
 - Except for label swapping, pushing, popping, rewriting
- Two level FIB is key for scaling

Computing Feasible Paths

- Brute force: maintain SPF trees of neighbors
 - Resource requirement is same as cSPF for MPLS FRR
 - Incremental SPF helps a lot here
 - Extra SPFs can be computed lower priority

Conclusion

- IGP convergence can be in few to tens of milliseconds
 - But it is hard to implement, will take time
- IGP fast reroute can provide relief
 - For ISPs who are not deploying MPLS fast reroute
 - For ISPs who are deploying MPLS but not edge-to-edge
- Question to the group and ADs:
 - Should IETF work on this?

Acknowledgments

- We would like to thank our ISP friends and Van Jacobson for comments and for letting us use their topologies for analysis.