



OSPF FAST NOTIFICATIONS

IETF79 (BEIJING) – NOV 7-12, 2010

SRIGANESH KINI, WENHU LU, ALBERT TIAN

MOTIVATION

- › Convergence time is dependent on several factors, some of which are orthogonal to others
- › Improvements in one area result in other areas becoming (relatively) more significant contributors to convergence delays
 - Improvements in hardware, forwarding-architecture and network-design are making FIB installation delays lesser of an issue.
- › Hop by hop control-plane processing introduces non-determinism

MOTIVATION (CONTD)

- › Area wide multicasting of control protocol information needs to be examined for potential gains
- › A fast, network-wide (scoped) notification mechanism is useful to trigger recovery mechanisms

DESIGN CONSIDERATIONS

- › Design a fast-notification (**FN**) mechanism that can be used by any potential application
- › Use a widely deployed protocol
- › Use multicasting as much as possible

FN USING OSPF – DESIGN CONSIDERATIONS

- › LSA can (re-)used to convey the fast-notification

- › Lacks (area-wide) multicasting of control information (LSAs) that does not involve hop-by-hop control plane processing

- › Problem can be split into two parts
 - Choice of path to multicast **FN**
 - Ensuring that LSA carrying **FN** conforms to OSPF flooding procedures

DESIGN CHOICES

- › Provide a mechanism to choose the type of flooding paths. Useful to provide right tradeoff between information duplication and reliability
- › Define a multicast address to deliver **FN**
- › Retain existing (hop-by-hop control-plane) flooding when doing **FN** so that any errors in **FN** flooding are recovered.

CHOICE OF FLOODING PATH - SPANNING TREE

- › Simple to compute
- › Most topologies/failures FN reaches all routers
- › For some specific cases of types of topologies/failure/spanning-tree the FN may not reach some routers

CHOICE OF FLOODING PATHS - MAXIMAL REDUNDANT TREES

- › <http://www.ieee-iscc.org/2009/program.html>
 - “On finding maximally redundant trees in strictly linear time”
 - › Andras Csaszar, et al

- › Can cover all failure conditions

ENSURING CONSISTENCY WITH LSA FLOODING

- › Generate FN along the path to flood FN quickly (preferably in hardware)
- › Use the same LSA comparison algorithm but treat copy of LSA received by FN as a temp to do any processing earlier

- › Acknowledgements

- Atlas Alia, Dave Allan, Acee Lindem, Joel Halpern

- › Comments welcome