



Fast notifications

draft-lu-fast-notification-framework-00.txt

IETF79 (Beijing) – Nov 7-12, 2010

Wenhu Lu, Albert Tian, Sriganesh Kini

Motivation

- › Fast convergence is important topic in its own right
 - FRR can not replace it (coverage/complexity issues etc)
- › Convergence time is dependent on several factors: detection, propagation, SPF, RIB/FIB
 - Need to attack from all fronts
 - Onion peeling effect: Improvements in one area result in other areas becoming more significant
 - Improvements in hardware, forwarding-architecture and network-design are making FIB installation delays less of an issue.
- › This work is trying to reduce the propagation delay
 - Hop by hop control-plane processing introduces non-determinism
 - Unix like non-real-time OS is getting popular in the control plane, and it does not provide guarantee under load

Motivation (contd)

- › A framework for fast dissemination of control protocol information needs to be examined for potential gains
- › Fast convergence is the first user of such a framework
- › Other application could also use it
 - E.g. it maybe useful to trigger other recovery mechanisms
- › Can be used by many potential applications
- › Can be supported by many different mechanisms

Layered Architecture

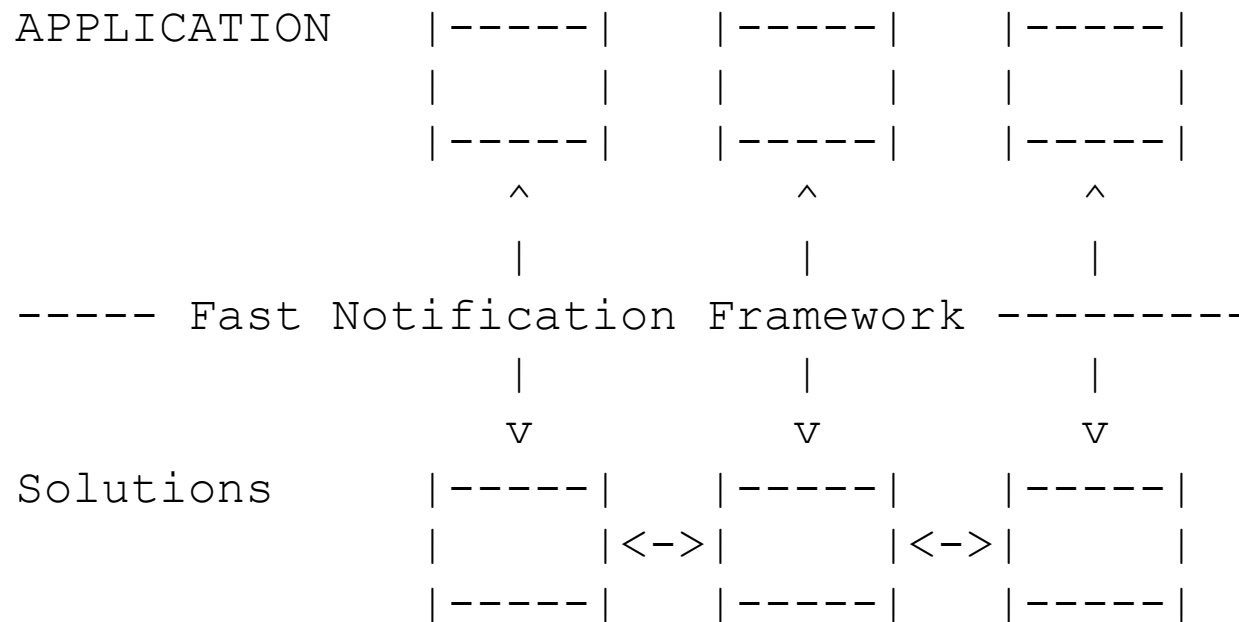


Figure 1: Fast Notification Architecture

FN using OSPF

- › LSA can (re-)used to convey the fast-notification
- › Pure optimization on top of existing slow flooding procedure so that flooding and convergence correctness is guaranteed

zN
§ @*«~@~>
k0U000Yb
aAaaCcCcC
NnNnOoOoO
ZzZzZzZzZz
π
GGGGllllkk
TTT00000
XΨΥΑΕΗΓ
НОПРСТУФ
ОПРСТУФХ
ёёёvгГёёё

Design choices

- › Useful to provide right tradeoff between information duplication and reliability
- › Retain existing slow (hop-by-hop control-plane) flooding when doing **FN** so that any errors in **FN** flooding are recovered
- › **FN** occupies a sequence number at the originator
- › FN can trigger slow flooding
- › FN is used for bad news (could also be used for good news)

Choice of flooding paths - Maximal Redundant Trees

- › Many different ways to choose flooding paths
- › <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

z[]
§ @*~@~>
k0000Yb
aAaaCcCcC
NpNn000eE
ZzZzZz fSfS
π
GGGGllllkk
TTT00000
xψÿAεηι
НОПРСТУФ
ОПРСТУФХ
ёёenvrГёæ

FN using ISIS

- › Similar design can apply to ISIS

z]`
\$ @*~@~>
k0000Yb
aAaaCcCcC
NnNnOoOeE
ZzZzZz f\$§~
π
GGGGllllkk
TTT00000
xψÿαєи
нОПРСТУФ
оПРСТУФХ
ёёөvГГеё

› Acknowledgements

– Alia Atlas, Dave Allan, Acee Lindem, Joel Halpern, Jeff Tantsura, Mike Shand, Olivier Bonaventure, Jakob Heitz, etc

› Comments welcome