

IP Fast Reroute for AI/ML Fabrics

draft-clad-rtgwg-ipfrr-aiml

Francois Clad, Clarence Filsfils, Roy Jiang, Dennis Cai
IETF 125, Shenzhen
14-20 March 2026




AI Backend Networks (Reminder)

- ~~Scale-Up~~ — Leave it out for now
- Scale-Out
 - **Topology:** typically, 2- or 3-tier folded Clos (sometimes multi-plane and/or multi-rail)
 - **Routing protocol:** (e)BGP
 - **Resiliency mechanisms:** ECMP
- Scale-Across
 - **Topology:** arbitrary
 - **Routing protocol:** IS-IS, OSPF, (e)BGP
 - **Resiliency:** ECMP, LFA, TI-LFA, etc.

Limitation of Existing Mechanisms

- CPU-activated resiliency mechanisms
 - Multiple orders of magnitude too slow for AI backend
- Local-only protection
 - ECMP is not always available
 - TI-LFA may hairpin traffic
- No FRR-like mechanism for capacity/quality degradation
 - Partial link (e.g., bundle member) or path (e.g., spine-to-super-spine in 3T folded Clos) failures can create congestions



Requirements

- Hardware-accelerated protection activation
- Hardware-accelerated network notifications
- Complete topology visibility
- Quality-aware and remote protection

Hardware-Accelerated Protection Activation

- Leverage extended NPU capabilities to activate protection without CPU intervention
- Same hardware-accelerated logic should apply to:
 - Local failure / quality degradation
 - Network notification of remote event

Hardware-Accelerated Network Notifications

- Extensively discussed in FANTEL context
- Problem statement in draft-ietf-rtgwg-net-notif-ps
- Standardization effort starting at IETF such as draft-camarillo-rtgwg-lsn

Complete Topology Visibility

- In BGP-only fabrics, nodes don't have any visibility beyond their direct neighbors
- However, complete topology visibility is needed for
 - LFA, TI-LFA path calculation
 - Any form of remote protection
- Easily achieved by enabling BGP-LS
(draft-ietf-idr-bgp-ls-bgp-only-fabric)

Quality-Aware Remote Protection

- React to remote failure or performance degradation
- Protection can consist in
 - Adjustments to ECMP weights
 - Use of repair path(s) outside the ECMP set
 - A combination of both
- One solution described in draft-clad-rtgwg-efficient-remote-protection

Next Steps

- Seek feedback from the working group