



Mechanisms for Optimal LAG/ECMP Component Link Utilization in Networks

draft-krishnan-large-flow-load-balancing-04

IETF 86

Ram Krishnan, Sanjay Khanna (Brocade Communications)

Lucy Yong (Huawei)

Anoop Ghanwani (Dell)

Ning So (Tata Communications)

Bhumip Khasnabish (ZTE)



OVERVIEW

- Mechanisms for optimal utilization of LAGs
- Based on long-lived large flow identification
 - Flows that exceed a certain BW over a pre-defined interval; e.g. 5% of link BW over 1 sec
- Two mechanisms for identification of flows
 - Sampling techniques
 - Inline, line rate hardware-based technique
- Two mechanisms for rebalancing of flows
 - Move long-lived large flows, keep small flows as is
 - Move small flows, keep long-lived large flows as is

ADDITIONS/CHANGES SINCE IETF 85

- Positioning
 - Applicable to all wire-line networks – not limited to SP backbone networks
- Mechanisms
 - Move small flows (flows other than the identified long-lived large flows)
 - Adjusting the LAG hash table to account for the long-lived large flows
- Background information
 - Simulation Results
 - Detailed explanation of algorithm for automatic hardware identification
- Editorial cleanup



NEXT STEPS

- Adopt as a work item in OPSAWG
 - Individuals from both operators and vendors have expressed interest and provided feedback