

Fully Adaptive Routing Ethernet using BGP

Draft-xu-idr-fare-04


Xiaohu Xu@China Mobile
Shraddha Hegde@HPE
Keyur Patel@Arrcus
Zongying He@Broadcom
Junjie Wang@Centec
Hongyi Huang@Huawei

Qingliang Zhang@H3C
Hang Wu@Ruijie Networks
Yadong Liu@Tencent
Yinben Xia@Tencent
Peilong Wang@Baidu
Tiezheng Li@IEIT SYSTEMS

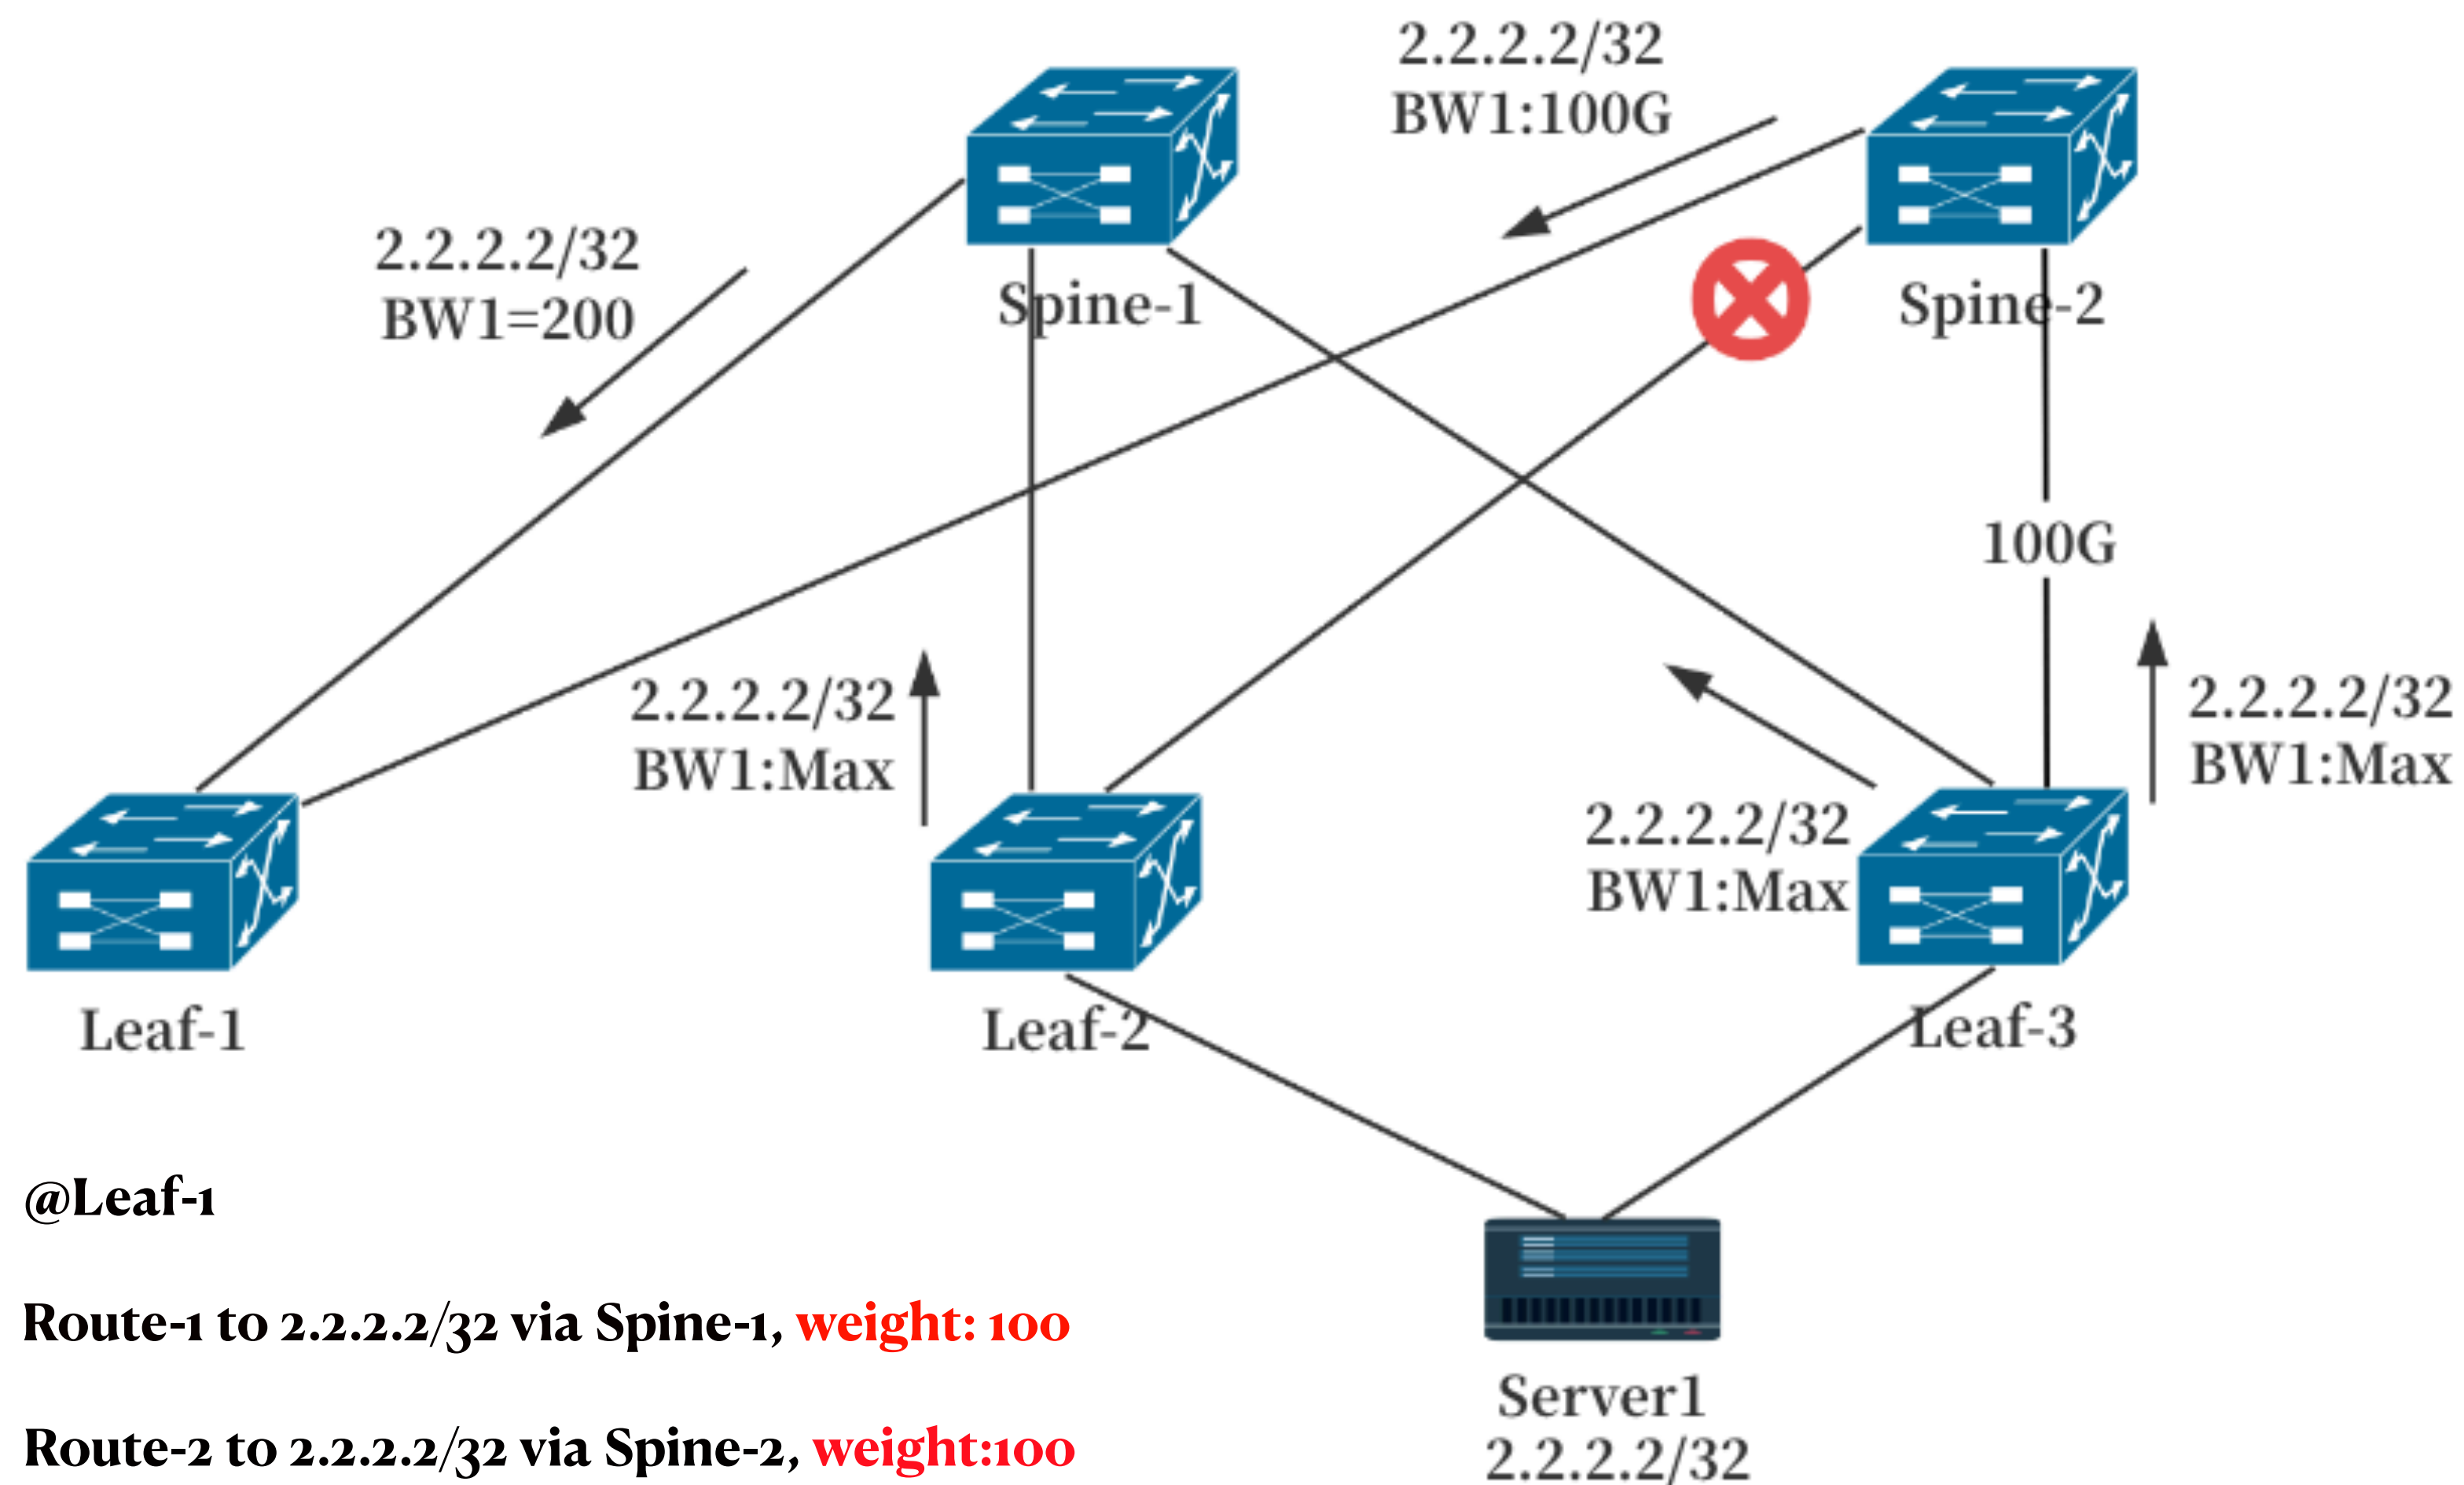
IETF125, Shenzhen

Changes Since -03


- Add a new co-author: Keyur Patel@Arrcus.
- Add a description of path bandwidth computation in multi-homing scenarios.



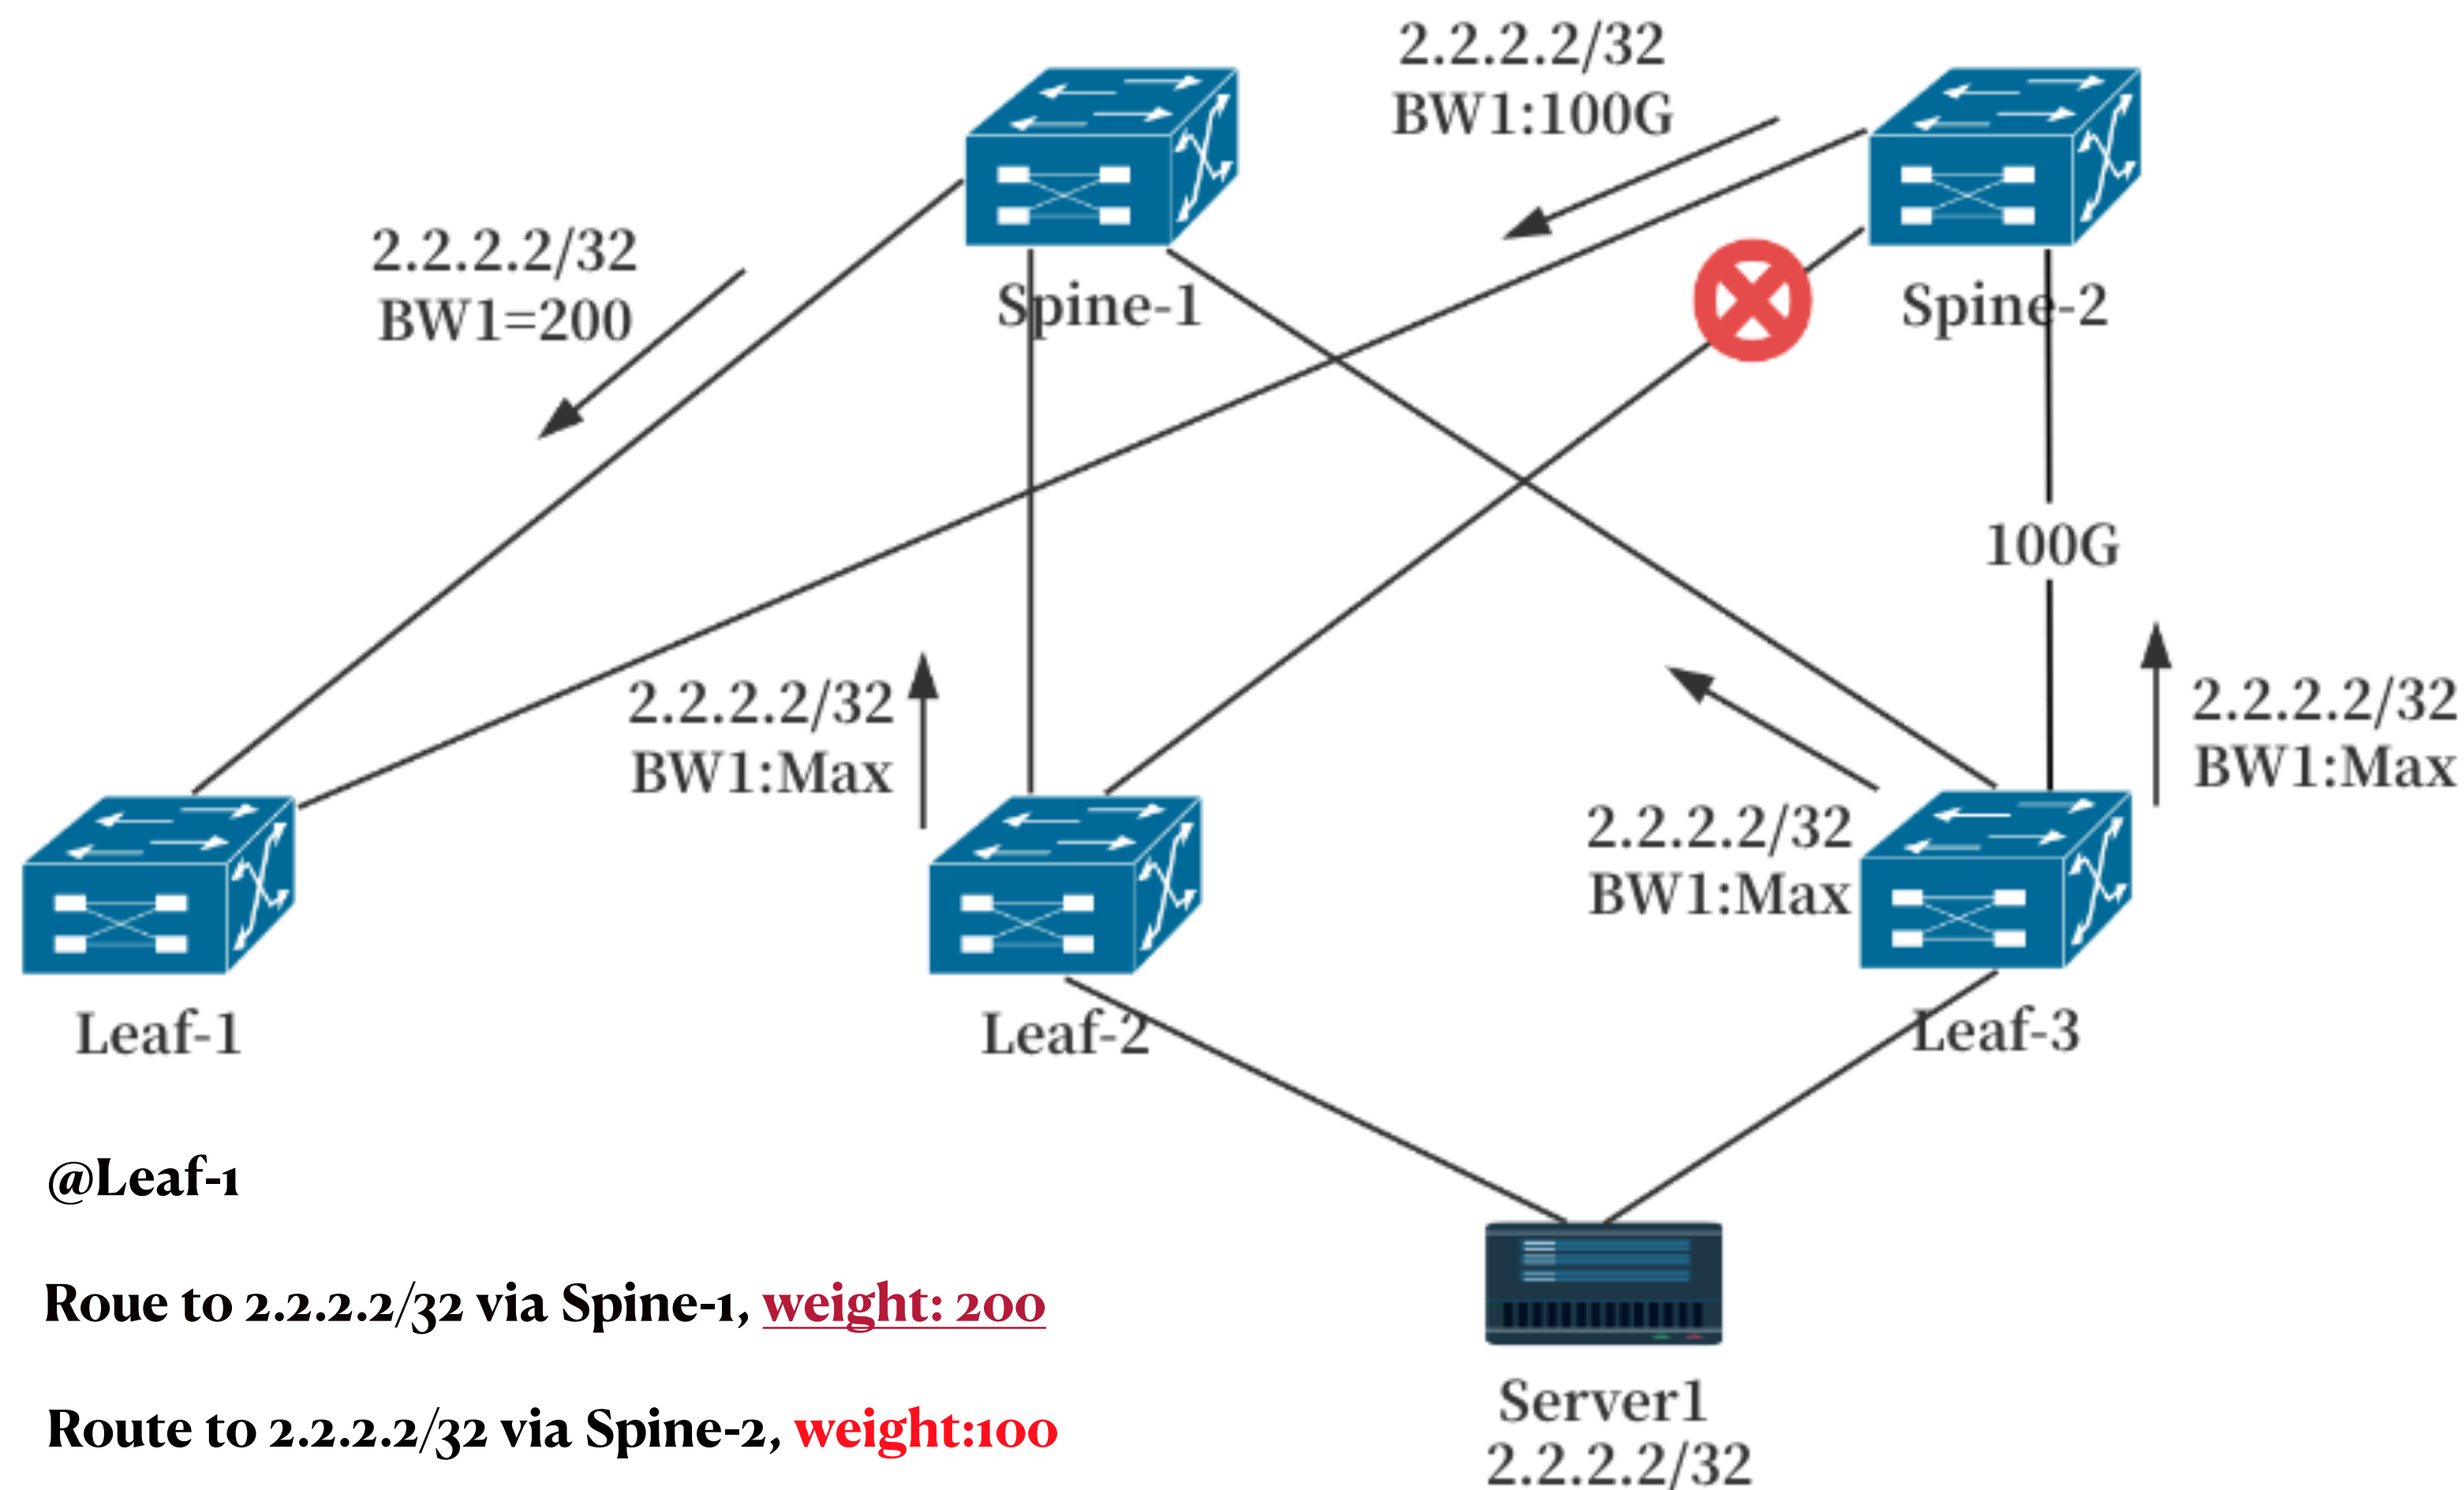
Previous: No Special Consideration for Multihoming F[®]



- When a leaf node, such as Leaf1, receives multiple equal-cost routes for that prefix from spine nodes (e.g., Spine-1 and Spine-2), for each route, it will determine the minimum value between the bandwidth of the link towards the advertising node and the value of the path bandwidth extended community carried in the received route, and then use that minimum bandwidth value as a weight value for that route when performing weighted ECMP load-balancing.



Now: Special Consideration for Multihoming^{TF}



- When a given IP prefix is multi-homed to multiple leaf nodes (e.g., two leaf nodes), the value of the path bandwidth extended community carried in the received route SHOULD be divided by the number of multi-homed leaf nodes (e.g., two) before determining the minimum value.
- Alternatively, the bandwidth of the link towards the advertising node SHOULD be multiplied by the number of multi-homed leaf nodes (e.g., two) before determining the minimum value.

Implementation Status

- The FARE protocol has been implemented by at least 2 cloud providers and 5 vendors, including **China Mobile Cloud, Baidu Cloud, Ruijie, IEIT System, H3C, Accton, and EdgeCore.**
- It has also been successfully adapted to more than four mainstream switch chips, including but not limited to:
 - Broadcom Tomahawk 5 (51.2T);
 - Nvidia Spectrum-4 (51.2T);
 - Huawei HiSilicon SD58XX(51.2T);
 - Centec Arctic (25.6T&12.8T);

Next Step

- The document is ready for the WG adoption call.