Contribution ID: **145** Type: **not specified**

Scaling Linux Bridge Forwarding Database

Tuesday, November 13, 2018 12:10 PM (35 minutes)

Linux bridge is deployed on Hosts, Hypervisors, Container OS's and in most recent years on data center switches. It is complete in its feature set with forwarding, learning, proxy and snooping functions. It can bridge Layer-2 domains between VM's, Containers, Racks, POD's and between data centers as seen with Ethernet-Virtual Private networks [1, 2]. With Linux bridge deployments moving up the rack, it is now bridging Larger Layer-2 domains bringing in scale challenges. The bridge forwarding database can scale to thousands of entries on a data center switch with hardware acceleration support.

In this paper we discuss performance and operational challenges with large scale bridge fdb database and solutions to address them. We will discuss solutions like fdb dst port failover for faster convergence, faster API for fdb updates from control plane and reducing number of fdb dst ports with Light weight tunnel endpoints for bridging over a tunneling solution (eg vxlan).

Most solutions though discussed around the below deployment scenarios are generic and can be applied to all bridge use-cases:

- Multi-chassis link aggregation scenarios where Linux bridge is part of the active-active switch redundancy solution
- Ethernet VPN solutions where Linux bridge forwarding database is extended to reach Layer-2 domains over a network overlay like VxLAN

[1] https://tools.ietf.org/html/draft-ietf-bess-evpn-overlay-11
[2] https://www.netdevconf.org/2.2/slides/prabhu-linuxbridge-tutorial.pdf

Presenters: PRABHU, Roopa (Cumulus Networks); ALEKSANDROV, Nikolay (Cumulus Networks)

Session Classification: Networking Track