



Contribution ID: 71

Type: **not specified**

LAG and hardware offload to support RDMA and IO virtualized interfaces

Monday 9 September 2019 15:45 (45 minutes)

Link Aggregation (LAG) is traditionally served by bonding driver. Linux bonding driver supports all LAG modes on almost any LAN drivers - in the software. However modern hardware features like SR-IOV-based virtualization and state full offloads such as RDMA are currently not well supported by this model. One of possible options to solve that is to implement LAG functionality entirely in NIC's hardware or firmware. In our presentation we present another approach, where LAG functionality for state full offloads such as RDMA and IO virtualization is implemented mostly in software, with very limited support from existing Hardware and firmware. A concept that should make the solution more generic without complicating the HW any further.

The presentation is focused on 3 areas: implementation of active-backup mode for RDMA and virtual functions, usage of RX hash value to implement flow-based active-active mode and new active-active mode for virtual functions.

Proposed implementation of the active-backup mode for RDMA is done in RDMA and LAN drivers. An application continues using direct HW support for RDMA. LAN driver (with the help of RDMA driver) observes notifications from the bonding driver and accordingly controls low-level TX scheduling and RX rules for RDMA queues. The same mechanism can be used to transparently redirect network virtual functions from active to backup. We further explore the use of RX hash to implement active-active mode.

I agree to abide by the anti-harassment policy

Yes

I confirm that I am already registered for LPC 2019

Primary authors: Mr KASHYAP, Vivek (Intel); Ms SINGHAI JAIN, Anjali (Intel); Dr UMINSKI, Piotr (Intel)

Presenters: Mr KASHYAP, Vivek (Intel); Ms SINGHAI JAIN, Anjali (Intel); Dr UMINSKI, Piotr (Intel)

Session Classification: Networking Summit Track