



Contribution ID: 367

Type: **not specified**

## Reducing the Overhead of Network Virtualization: Software and Hardware Solutions

*Wednesday, 18 September 2024 15:00 (30 minutes)*

Virtualization comes with overhead, and networking is no exception. In a typical virtualized scenario, the traffic traverses the network stacks of both the guest and the host. This presentation will introduce software and hardware solutions to minimize this overhead.

One of the challenges a virtualized network stack will face is the consistency of policy to assign received packets to queues. The use of multiqueues is essential for scalability, but inconsistent packet assignment policies of the guest and the host result in unnecessary synchronization among processors, hurting scalability. The current virtio specification allows negotiating the use of a packet assignment policy called RSS. We will present an RSS implementation utilizing eBPF in QEMU. We will also propose to add a new feature to the tuntap device to pass hash values calculated for RSS to the guest via vhost-net so that the guest can reuse them to assign packets to internal queues, reducing the duplication of work further. These kernel mechanisms are potentially useful to implement other packet assignment policies the virtio spec may gain in the future.

While software optimizations are important for general setups, dedicated network virtualization hardware is the best choice for optimized virtualization environments. Some NICs support virtualization based on the PCI Express SR-IOV specification. We have implemented an emulation of Intel's NIC that supports SR-IOV in QEMU to ease testing SR-IOV. In addition, we are also proposing to add SR-IOV support to QEMU's virtio implementation as a preparation for applying it to vDPA.

**Primary author:** ODAKI, Akihiko (Daynix Computing Ltd.)

**Presenter:** ODAKI, Akihiko (Daynix Computing Ltd.)

**Session Classification:** Networking Track

**Track Classification:** Networking Track