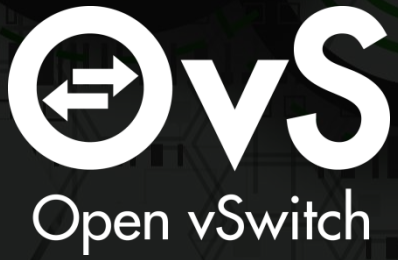

Ram.Vepa|Venu.Iyer@oracle.com

Oracle Corporation

Solaris OVS and SR-IOV

Agenda

- Solaris OVS Status
- Solaris H/W offloads: Primer
- Solaris cloud use cases
- SR-IOV and OVS: Benefits & Challenges
- Solaris SR-IOV and OVS
- Solaris OVS extension

Solaris OVS Status

- Solaris OVS (part of the next Solaris release) is based on 2.3.2
- Solaris delivers its netdev and dpif modules.
- Datapath in the kernel leverages Solaris infrastructure components such as VNIC and User Flows.
- Thanks to the community for the documentation and development discussions which helped the port!
- Solaris OVS port to 2.6 in progress
  - Netlink support on Solaris
  - OVN plug-in, as one of Solaris OpenStack Neutron ML2s.
- Solaris is available to the community developers as VM in Oracle Cloud, or via a Solaris BareMetal Loaner Program (new Sonoma-based servers).

Solaris GLDv3 API : Introduction

- Solaris GLDV3 MAC provider API for drivers to register with Solaris stack.
- As part of the registration, Solaris queries for driver capabilities and does initializations, if needed.
- As part of this capability exchange, Solaris stack knows if the device supports:
  - H/W VNICs
  - H/W Flows
    - L4-L7, to support multi-tenant enterprise apps.
  - SLA
    - Priority, limits, shares.
  - Actions
    - To enable application-centric network policies.

Solaris H/W Offloads: Introduction

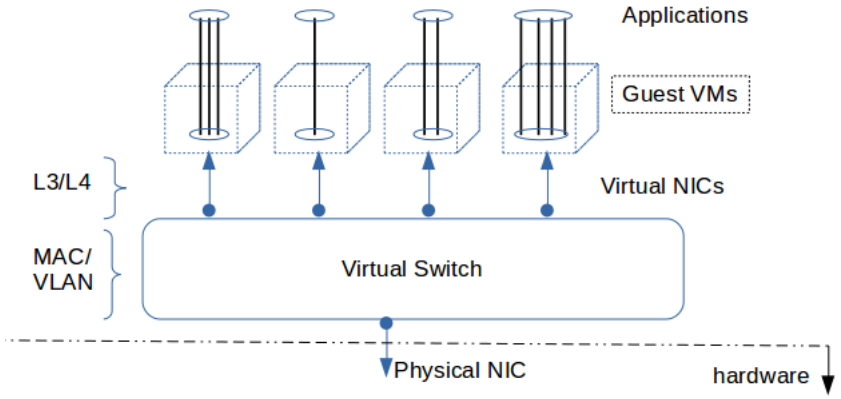

L2-L4 classification in S/W

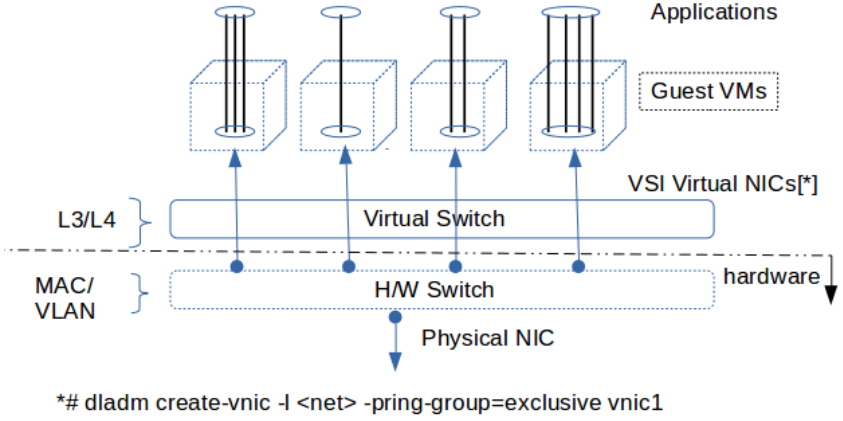

L2 classification in H/W; L3-4 in S/W

Solaris H/W offloads: Introduction

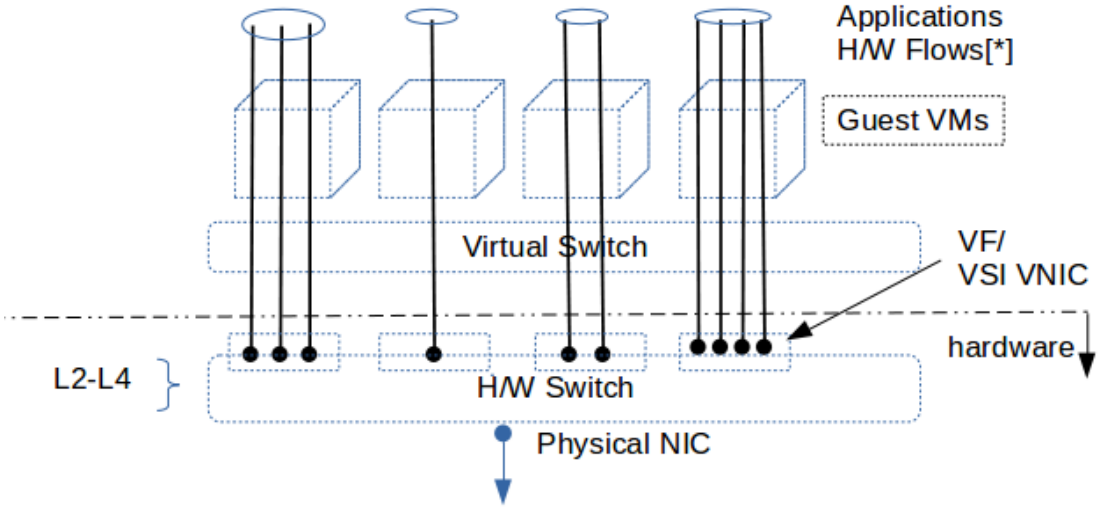

L2-4 in H/W

Oracle Cloud

- Solaris VM in the Cloud is available now!
- New Large Oracle Public Cloud (OPC) was announced in OOW’16.
- PaaS services, Enterprise Application-centric
  - Use OVS as a mechanism to deliver monitoring, SLA and networking policies to Enterprise Apps.
  - With the highest datapath performance possible
  - With OpenStack as an orchestration mechanism.

Solaris cloud use case

Predictable Performance of Multi tenant Applications

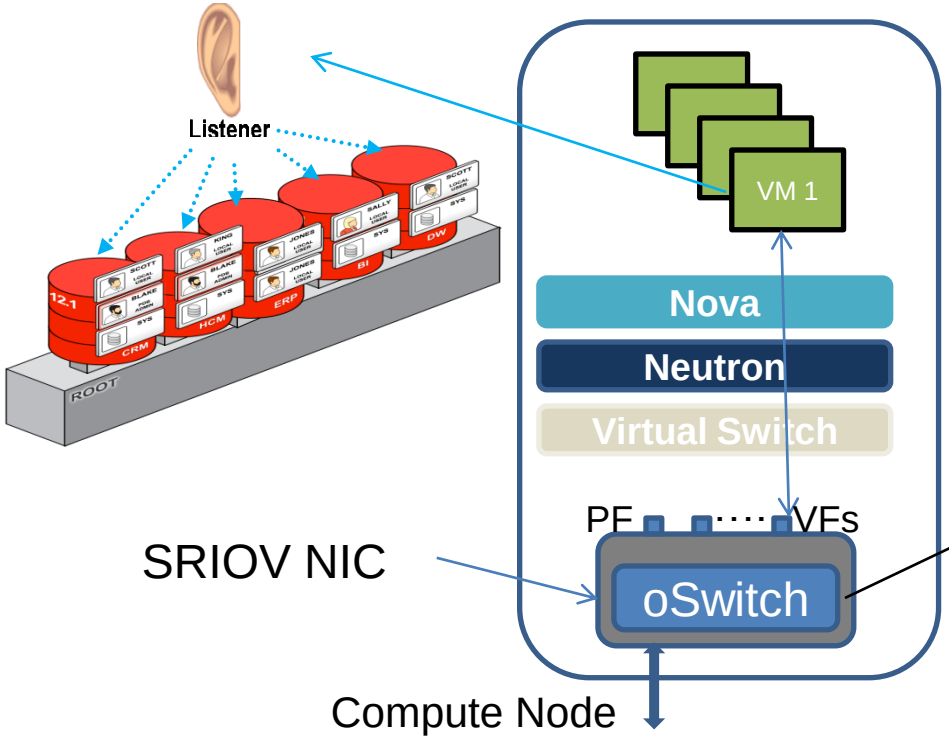

Offloaded rules

	L2-L7 filters with actions • VTEP • ACL & Security groups • SLAs per flow • S/D NAT • Metadata for SFC • Open Flow stats

SR-IOV and OVS: Benefits

- SR-IOV and H/W based VNICs
  - Isolation; uses hardware resources that are bound to the VF / H/W VNIC.
  - Performance; datapath doesn’t go via the host (para-virtualized)
  - Preferred for high traffic Virtual Machines, e.g. routers, that are L3-centric.

- Hardware Flows (L3+)
  - Isolation: Flows have dedicated resources.
  - Performance: Classification and action offload to the H/W

SR-IOV and OVS: Challenges

- Challenges with Open Flow flows in H/W
  - NIC (H/W Switch) should support classification attributes: L2-L4.
  - NIC should support all the actions, esp. tunneling, NAT, conntrack etc.

- Challenges with OVS and SR-IOV
  - NICs typically support a subset of OVS flow attributes and actions.
  - An all-or-nothing offload will not be optimal.
  - SmartNIC does offload OVS to h/w and could be an option, but cost may be a factor.

SR-IOV and OVS: Alternatives

- External [OVS] switch (TOR)
  - e.g. VTEP could be on the external switch.

- DPDK
  - Addresses the performance issues
  - Relies on the datapath going to the primary domain or host, which becomes an issue where the primary is not configured for that purpose (bandwidth, CPU etc.)
  - Isolation is still an issue; a guest under attack also impacts the primary
  - Performance can be optimized if we can drop packets in the h/w.

- OpenNFP and other accelerations.

Solaris SR-IOV and OVS

E.g: Src Port = OF#2 (a), Output: OF# 3 (e)
Action 1: xyz (s/w) (b, c), Action 2: BW limit (h/w) (d)

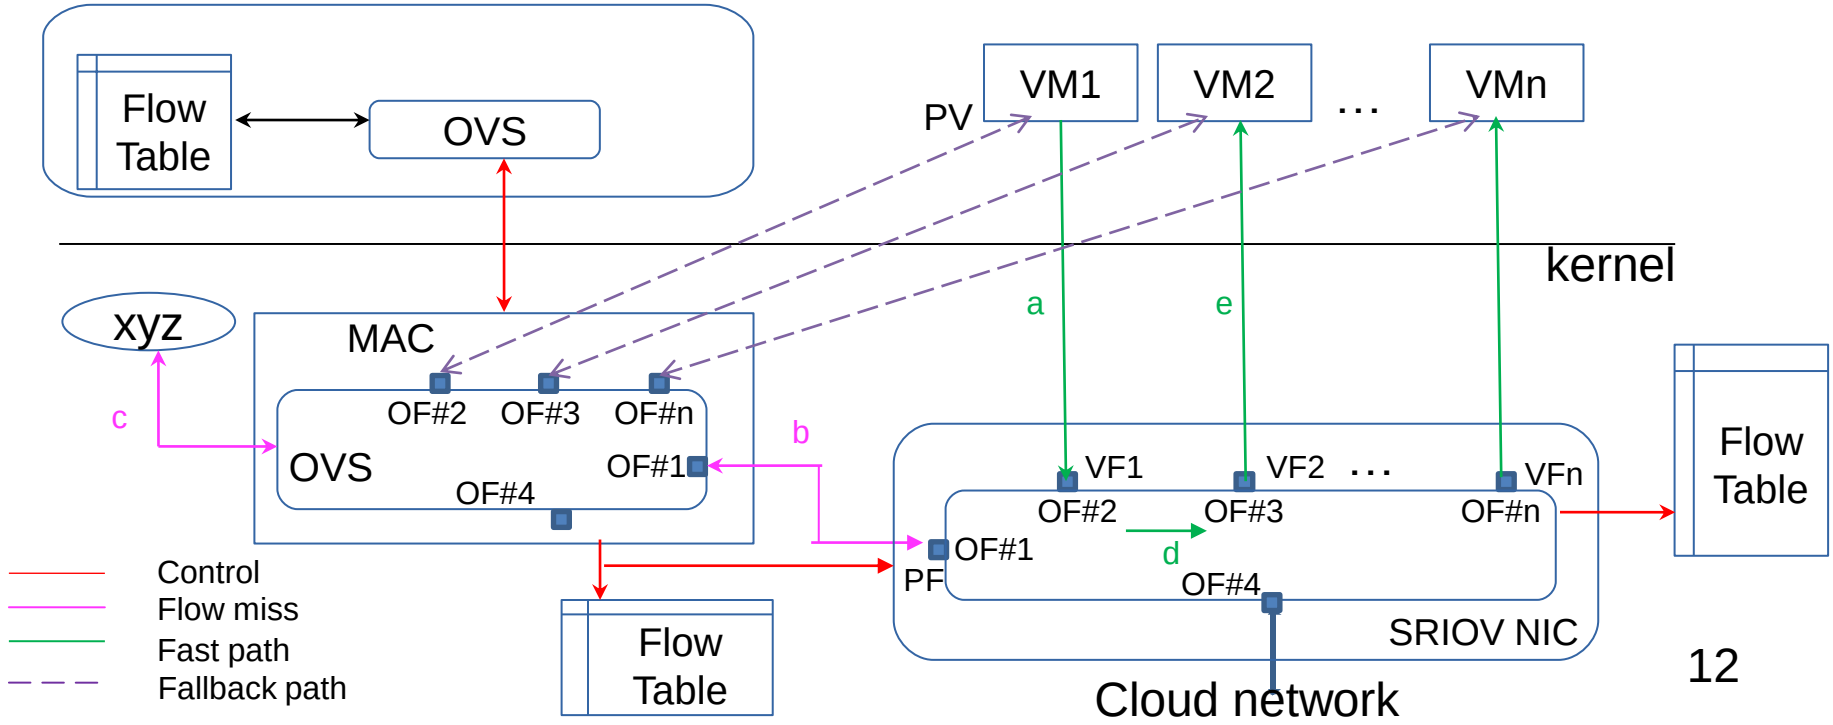

Solaris OVS potential enhancements

- Solaris OVS metadata extension:
  - LSO, Checksum
    - Need to pass packets to OVS with some metadata, which can be sent back along with the packet
  - Application ID
    - Need to include some metadata with the packet to enable end to end classification:
      - On the host via metadata
      - On the wire using Overlay.
- Debugging tool
- Engage with the community to push Solaris port upstream.