


Open vSwitch Offload

Simon Horman

DXDD Europe

Utrecht, 8th June 2017

- What is Open vSwitch?
- Offload Models
- TC-Flower Based Offload

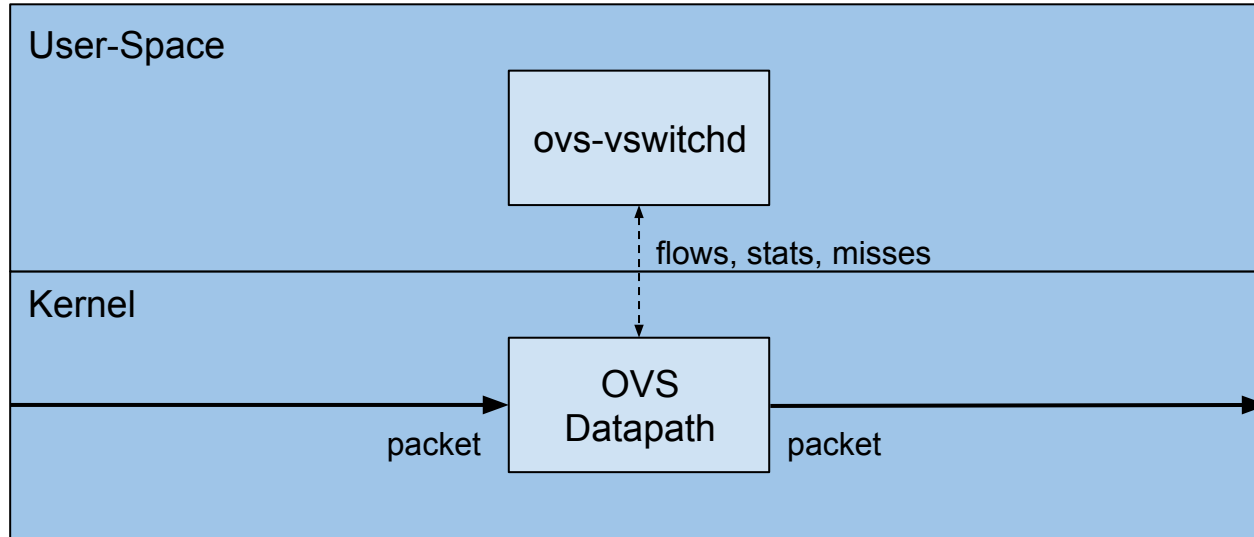



What is Open vSwitch

What is Open vSwitch (OvS)

- Fully featured virtual switch
- Provides match-action scheme
- Configuration via extended OpenFlow and OVSDDB
- Accelerated forwarding of cached flows
- Mega Flows to enhance caching behaviour
- Tunnel termination and origination
- Platform dependent extensions: QoS, Rate-Limiting

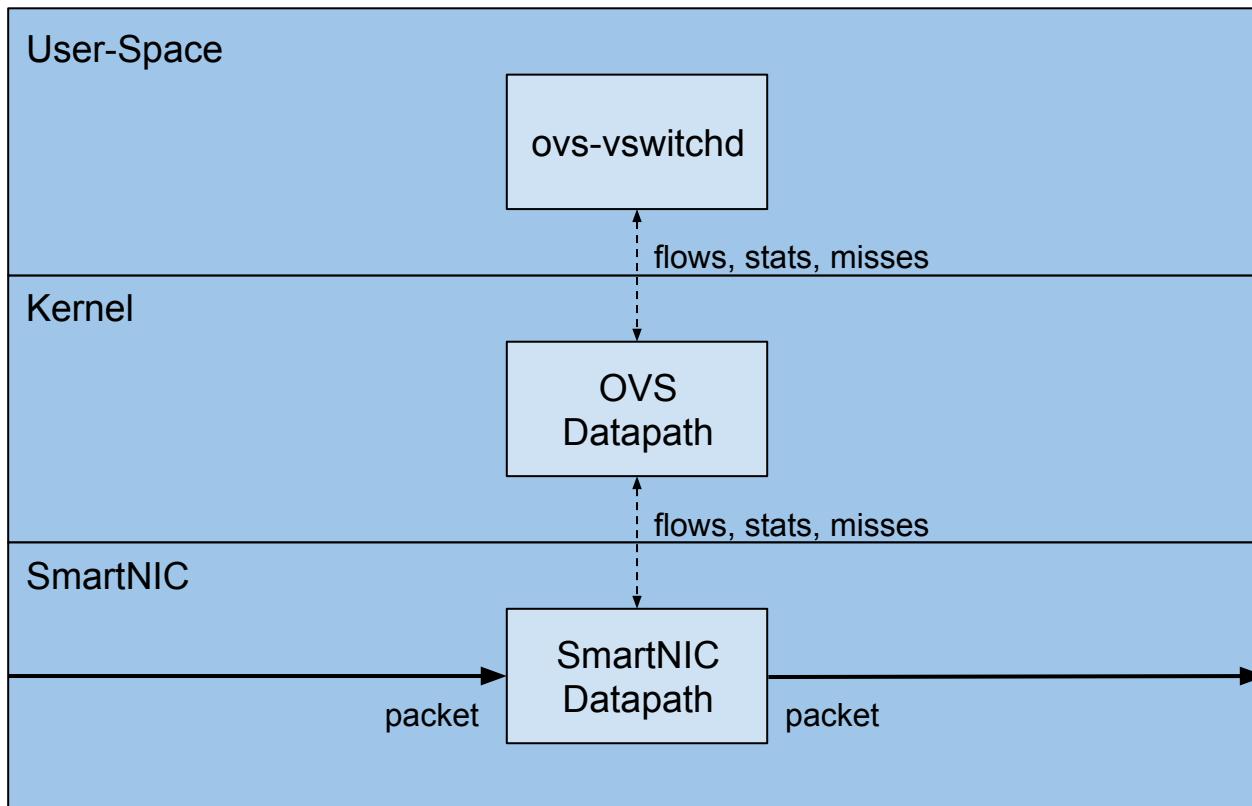
Kernel Datapath

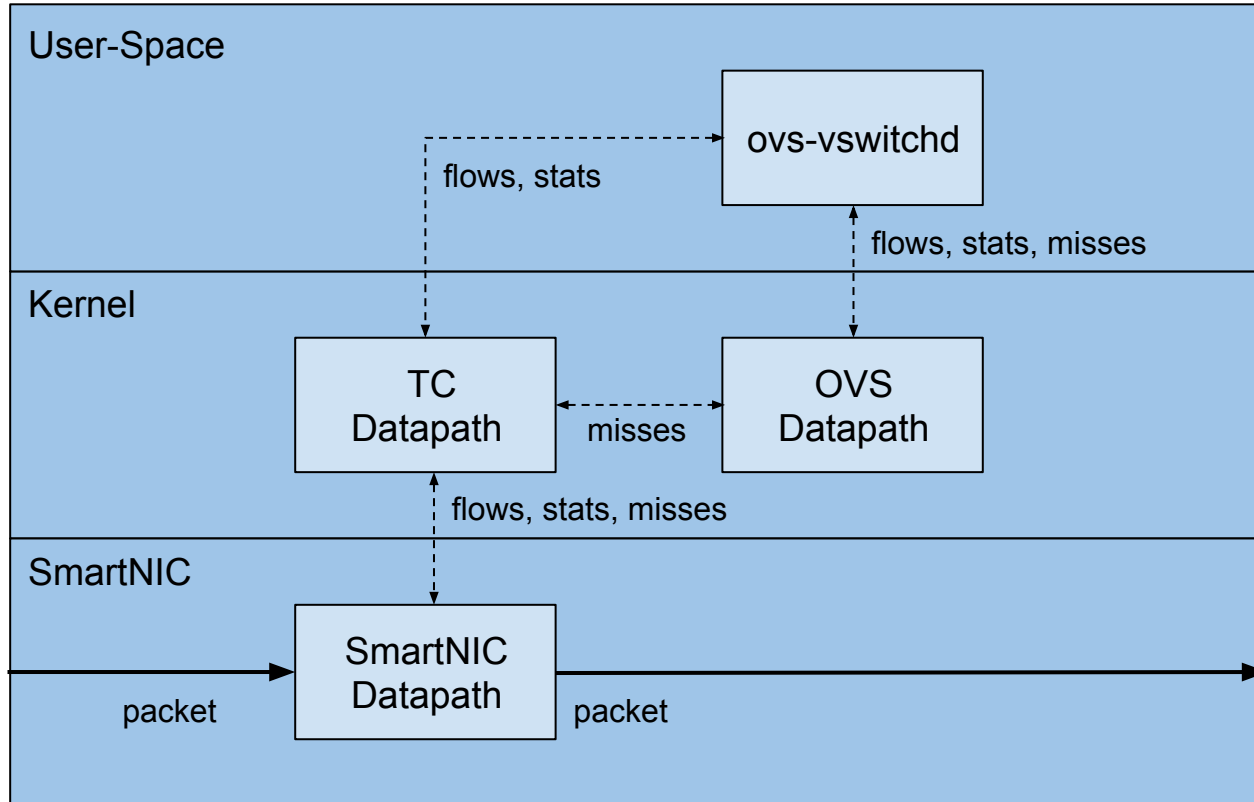





Offload Models

OVS Datapath Hooks







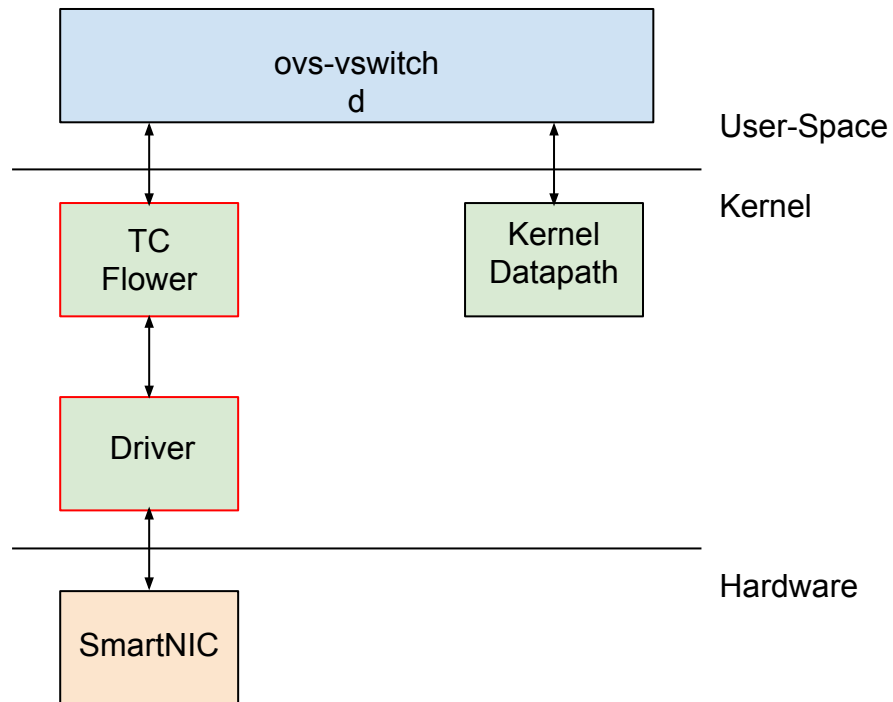
TC Flower Based Offload

- What is TC Flower?
- Overview of OvS-TC
- Example Modification
- How to Participate

- Packet classifier for Linux kernel traffic classification (TC) subsystem
- Allows match on key with a wide number of packet and metadata fields
- TC actions may be used to provide match-action behaviour similar to Ovs

- OvS Datapath
 - Single table
 - Match on in_port
 - Flows have a wide key and are disjoint
 - And therefore can be partitioned into slices
 - MegafloWS are priority independent
- TC Flower
 - Recently gained multi-table (chain) support
 - Attached to in_port
 - Flows have a wide key
 - Only one mask per priority

OvS with Offloaded TC Flower



- May flag TC filters as software-only, hardware-only
- Default is software and if available hardware

- Stage 1
 - L2, L3 & L4 matches
 - Drop, output & VLAN actions
- Stage 2
 - VXLAN
- Stage 3 - Requires further user-space enhancements
 - Set Action

- OvS-TC
 - Patches available
 - Hope to see these included in OvS v2.8
- Kernel
 - TC Flower and related actions are present
 - More features are being added, more to be added...
- NFP
 - Prototyping
 - Plan to post for inclusion in Linux kernel soon

Example TC Flower Modification

- Two Kernel Components Work Together
 - Network stack core: [flow dissector](#)
 - TC: [flower classifier](#)
- Userspace Utility for Exercising Code
 - Iproute2: [tc](#)

- Feedback on feature set
 - OvS has many features
 - OvS-TC is starting with few
 - Plans for adding new features to OvS-TC is not fixed
- Discussion and code review on mailing lists
 - Kernel: netdev@vger.kernel.org
 - Open vSwitch: dev@openvswitch.org



Thank You