

Open vSwitch

Justin Pettit

NetFilter Workshop

22 June 2015

Overview

- OVS is a multi-layer switch
- Visibility (NetFlow, sFlow, SPAN/RSPAN)
- Fine-grained ACLs and QoS policies
- Port bonding, LACP, tunneling
- Centralized control through OpenFlow and OVSDDB
- Open source
- Multiple ports to physical switches

Where is it used?

- Broad support
 - Linux, FreeBSD, NetBSD, Windows, ESX
 - KVM, Xen, Docker, VirtualBox, Hyper-V,...
 - OpenStack, Cloudstack, OpenNebula,...
- Widely used
 - Most popular OpenStack networking backend
 - Default network stack in XenServer
 - 1440 hits in Google Scholar
 - Thousands of subscribers to OVS mailing lists

(Partial) List of Contributors



Who develops OVS?

- 184 contributors listed in AUTHORS
 - From 80 unique email domains
- 13 “committers”
- Commits from outside Nicira/VMware growing
 - 2012 and 2013: 19%
 - 2014: 24%
 - 2015 to date: 31%


Software-Defined Networking

- SDN is a method of using software to define networking behavior instead of having it baked into ASICs
- Benefit of having God-like knowledge of entire system as opposed to cooperatively working with neighbors
- OpenFlow is often used interchangeably with SDN, but SDN does not require OpenFlow

OpenFlow

- Idealized view of a switch's datapath
- Centralized controller configures flow table
 - Lookup based on L2-L4
 - Supports full wildcarding and priorities
 - Flows associated with actions: forward, drop, modify
 - Missed flows go to controller
- Remote visibility
 - Description of switch (supported actions, flow tables' sizes, etc.)
 - Statistics (flows, tables, ports)

Main Components




Flow Complexity

- Many different applications of SDN
 - Traffic management (Google-scale)
 - Network virtualization (Enterprises and Cloud Providers)
 - Security policies (Government Agencies)
- An application can includes hundreds of thousands of rules on an OVS instance with dozens of lookups per packet

Forwarding Components

- ovs-vswitchd (slow path)
 - Forwarding logic (Flow tables)
 - Remote configuration and visibility
- Kernel module (fast path)
 - Packet lookup, modification, and forwarding
 - Tunnel encapsulation/decapsulation



Decoupling

- Decoupling Helps
 - A number of different SDN applications have been written without requiring changes to OVS.
 - A number of new OpenFlow protocols have been added without changes to kernel
- Flow programming with slow-path/fast-path design often performs better than fixed-pipeline
- NSDI paper on design and implementation:
 - <http://openvswitch.org/support/papers/nsdi2015.pdf>

Kernel Integration

- Kernel is where performance-critical processing is done
- Looking to integrate with other components
 - Connection tracking for stateful firewalls
 - Load-balancing
 - NAT
- Reducing duplication of code

Future Plans

- Focus is on moving up the stack (OVN) and adding more stateful services
- Better integration with the kernel
- Bringing OVS to more platforms