

OpenStack and OVN

What's New with OVS 2.7

OpenStack Summit -- Boston 2017


Russell Bryant (@russellbryant)

Justin Pettit (@Justin_D_Pettit)


Ben Pfaff (@Ben_Pfaff)

Virtual Networking Overview

Provides a logical network abstraction on top of a physical network



Physical



Logical

What is OVN?

- Virtual networking for Open vSwitch (OVS)
- Developed within the OVS project
- Linux Foundation Collaborative Project
- License under the Apache license
- OVS/OVN releases every six months

OVN Feature Overview

- Manages overlays and physical network connectivity
- Flexible security policies (ACLs)
- Distributed L3 routing, IPv4 and IPv6
- Native support for NAT, load-balancing, DHCP
- Works with Linux, DPDK, and Hyper-V
- L2 and L3 gateways
- Designed to be integrated into another system
 - OpenStack, Kubernetes, Docker, Mesos, oVirt

Who's Using OVN?

- Existing Project Integrations

- OpenStack <https://github.com/openstack/networking-ovn>
- Quilt project <https://github.com/quilt/quilt>
- Kubernetes <https://github.com/openvswitch/ovn-kubernetes>
- Docker networking <https://github.com/openvswitch/ovs>
- oVirt <https://www.ovirt.org>

- Products

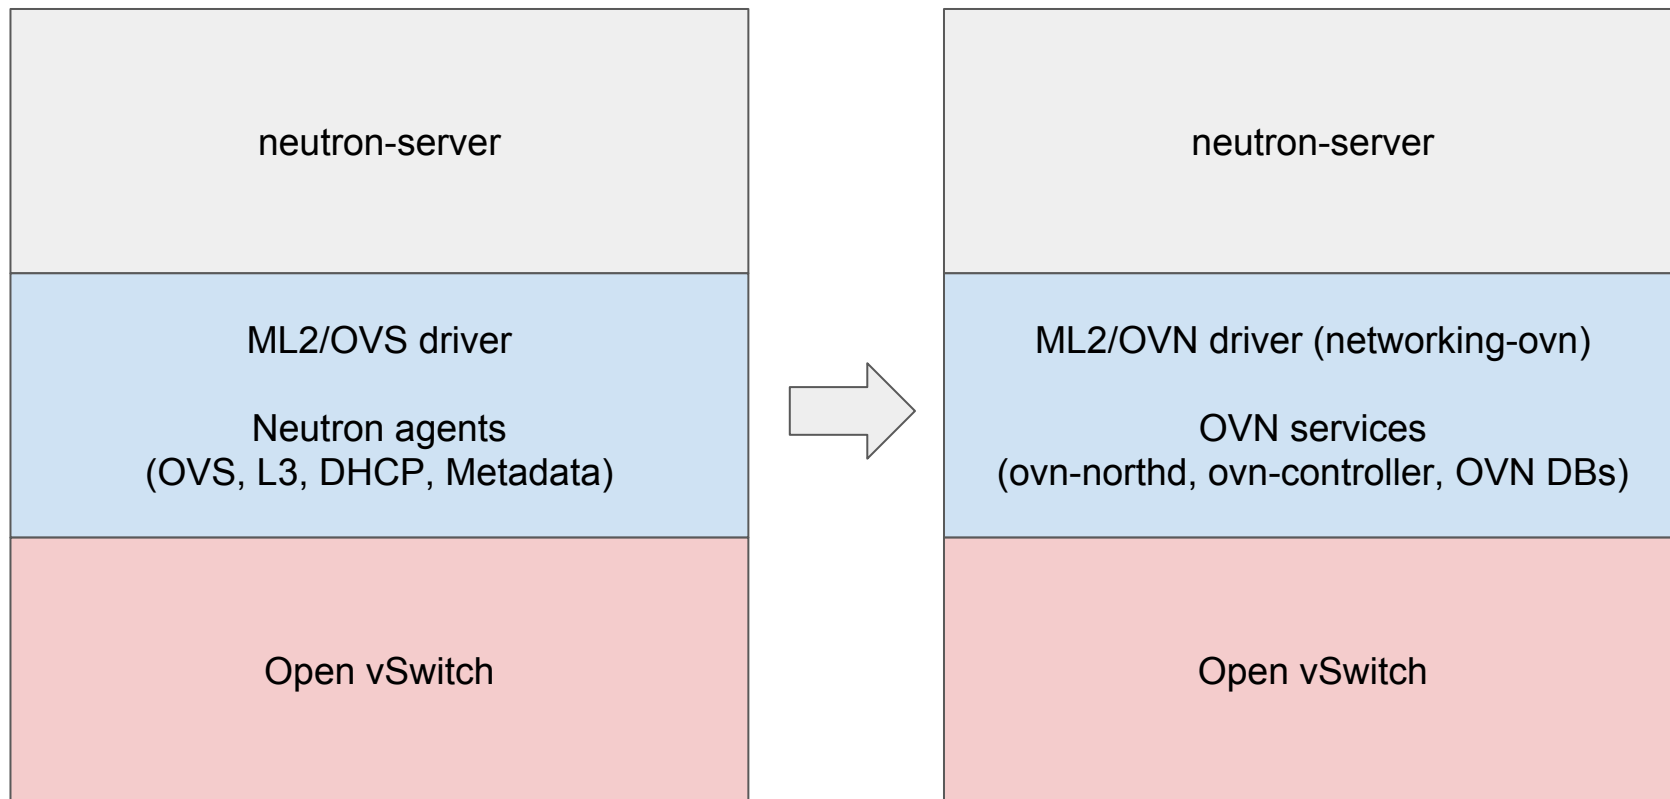
- Red Hat Virtualization (RHV) - OVN available as Tech Preview
- Red Hat OpenStack and OpenShift support on the roadmap

OVN Project Releases

- OVS Version 2.7
 - QoS via egress shaping
 - DSCP marking
 - IP source-based routing
- OVS Version 2.8
 - Native DNS support
 - IPAM IPv6 support
 - Database clustering (in progress)
 - ACL logging (in progress)
- Ocata Release
 - Support for OVN L3 gateways with NAT
 - Native DHCPv6 support
 - Mixed DPDK and kernel datapath environment
 - Many new tests
- Pike Release (Completed)
 - Simplified L3 gateway support
 - SSL support for OVN DBs
 - CI with OVS master and latest release
 - Many new tests
- Pike Release (In Progress)
 - Support for OpenStack Metadata API
 - L3 Gateway high-availability
 - Kolla container images
 - Improvements to TripleO Support
 - Ansible-based ML2/OVS to OVN migration

OpenStack Integration

How OVN fits into Neutron



openstack/networking-ovn

- Implements the OVN backend for Neutron
- Implements various Neutron interfaces
 - ML2 mechanism driver
 - L3 service plugin
 - QoS notification driver
 - Trunk driver (vlan-aware-vms)
- Configures OVN using OVSDDB protocol via python-ovs library

Migration to OVN

- Proof of concept Ansible playbook available to migrate from ML2/OVS
- Requires use of OVS Firewall with ML2/OVS first
- In-place upgrade without any VM migrations
- Attempts to minimize data plane downtime
 - Measured as ~10 seconds in the migration of a 10 node cloud
- Open question: do we need to build a no-downtime rolling migration?
 - Much more difficult than the in-place migration with minimal downtime

OpenStack Deployment with OVN

- DevStack
 - For development and basic testing only
 - <https://docs.openstack.org/developer/networking-ovn/testing.html>
- TripleO
 - tripleo-heat-templates/environments/neutron-ml2-ovn.yaml

Performance

Differences from ML2/OVS Control Plane

- ML2/OVS
 - RPC over message queues
 - Neutron agents
 - Python services
- OVN
 - Database-driven
 - OVN replaces Neutron agents
 - C binaries

ML2/OVS vs OVN Control Plane Comparison

- Results for Nova VM Creation time with OVN:

	ML2/OVS (Seconds)	OVN (Seconds)	% Improvement
Average	80.7	23.4	70.9%
95%	163.2	35.3	78.4%
Maximum	221.9	47.8	78.4%
Minimum	18.7	3.8	79.8%

<https://blog.russellbryant.net/2016/12/19/comparing-openstack-neutron-ml2ovs-and-ovn-control-plane/>

Data Plane Performance

- Distributed routing
 - Does not bounce traffic through namespaces
 - OVN performs route calculation then applies simple packet modification to future packets
- ACLs and NAT use native connection tracking functionality
 - Linux: Netfilter conntrack kernel module
 - DPDK: New OVS userspace connection tracker

Geneve vs VxLAN Performance

- OVN supports Geneve and STT for hypervisor connectivity
 - VxLAN doesn't provide enough metadata
- Support for Geneve-offloading on some NICs (more coming)
- Existing checksum-offloading in many NICs provides comparable Geneve performance to VxLAN-offloading

Future

Feature Development

- Database clustering (Under development)
- ACL Logging (Under development)
- Securely handle a compromised hypervisor
- Scaling improvements
- Service function chaining
- Encrypted tunnels
- Native OpenStack LBaaS support
- OpenStack support of multiple SNAT gateways on a network

Open vSwitch BPF Datapath

- BPF provides a safe, virtual sandbox in the Linux kernel (as well as other platforms)
- Potentially greater portability across kernel versions and platforms
- Insert new functionality at run-time:
 - New network and tunneling protocols
 - Push OVN-specific actions into the datapath

Break OVN out of OVS Repository?

- Projects have different trajectories
- Separate group of OVN contributors
- Considerations
 - Work needed to break things into libraries
 - Development under the “openvswitch” project?
 - Use same governance policy as OVS?
 - Continue same release cycle as OVS?

Other Resources

- OVS/OVN Repository
 - <https://github.com/openvswitch/ovs>
- OpenStack OVN Integration
 - <https://docs.openstack.org/developer/networking-ovn/>
- Kubernetes OVN Plugin
 - <https://github.com/openvswitch/ovn-kubernetes>
- OVS Orbit Podcast
 - <https://ovsorbit.org/>


Thank you for attending!

Ben Pfaff (@Ben_Pfaff)

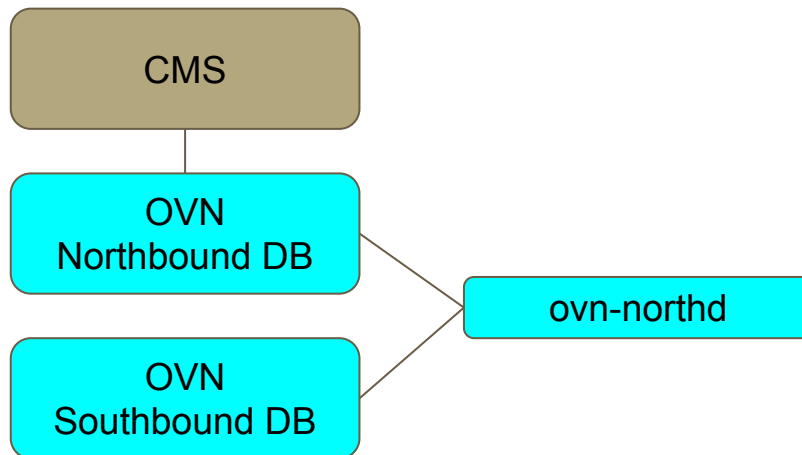
Justin Pettit (@Justin_D_Pettit)

Russell Bryant (@russellbryant)

1. Logical configuration in Northbound DB



2. ovn-northd populates Southbound logical flows



3. Hypervisors generate physical flows

