



# The evolution of Open vSwitch integration for OpenStack

Daniel Alvarez  
Numan Siddique

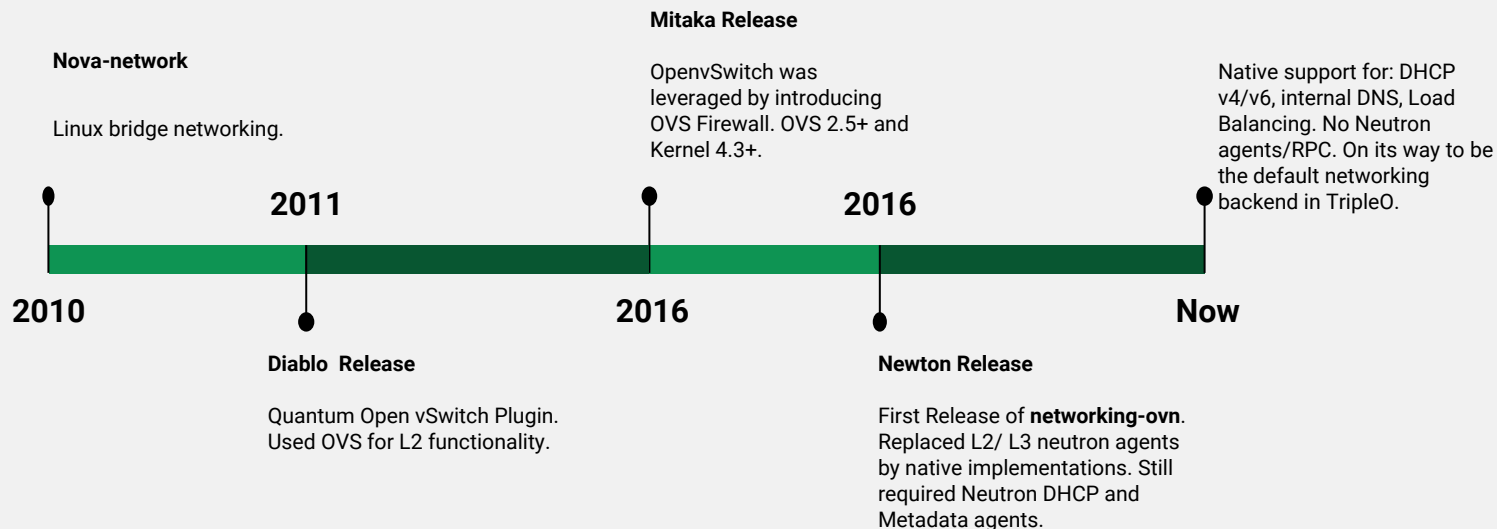
[dalvarez@redhat.com](mailto:dalvarez@redhat.com)  
[nusiddiq@redhat.com](mailto:nusiddiq@redhat.com)  
[@numansiddique](https://twitter.com/numansiddique)

13 Nov 2018

# Agenda

- History of OpenvSwitch in OpenStack
- OVN architecture overview
- ML2/OVN vs ML2/OVS
  - Features
  - Performance
- What's next?

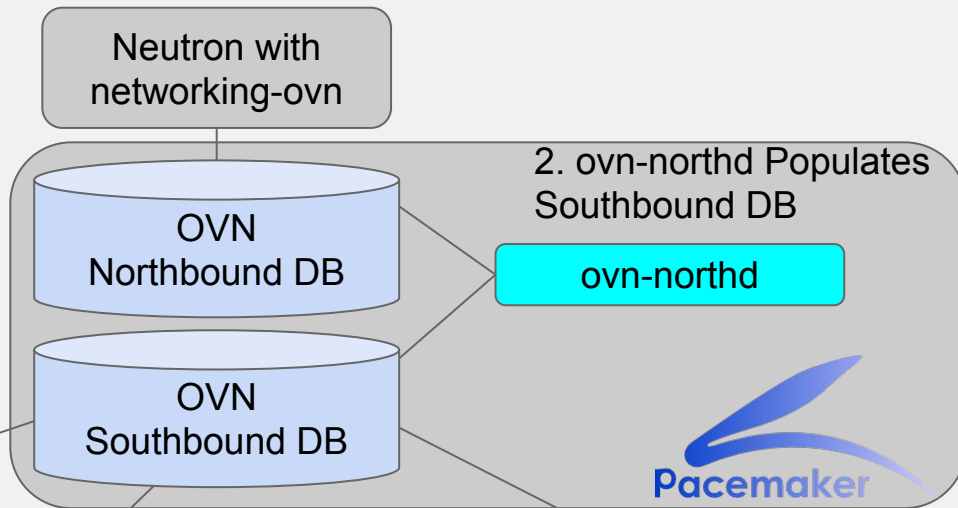
# History of OpenvSwitch in OpenStack



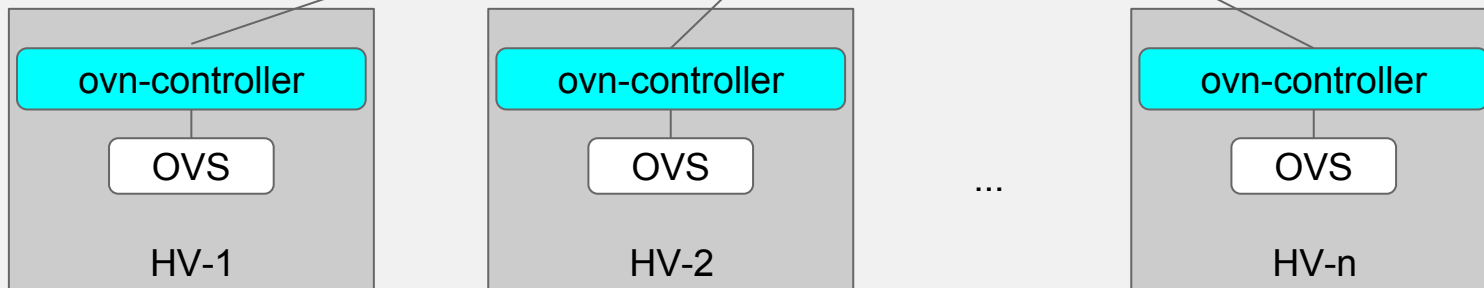
# OVN Architecture overview

# OVN Architecture

1. A Cloud Management System (Openstack in this eg.) creates Logical network components via the OVN Northbound DB



2. ovn-northd Populates Southbound DB



3. Hypervisors Generate Physical Flows

# Comparing

ML2/OVN and ML2/OVS

# Comparing ML2/OVN and ML2/OVS (I)

	<b>ML2/OVN</b>	<b>ML2/OVS</b>
<b>Native Product Compatibility</b>	OpenStack Kubernetes oVirt	OpenStack
<b>Resources/ Complexity</b>	C services/ single C agent, OVSDB protocol, smaller footprint.	Multiple python agents, rabbitmq, medium footprint
<b>L3</b>	OpenFlow based	L3-agent / Linux kernel namespaces, routing and iptables
<b>L3HA</b>	OpenFlow + BFD (Native)	L3-agent / Linux kernel namespaces + keepalived + VRRP over ha_xx network.

# Comparing ML2/OVN and ML2/OVS (II)


	<b>ML2/OVN</b>	<b>ML2/OVS</b>
<b>L3 DISTRIBUTED East/West</b>	Always (except for VLAN tenant networks)	Only with DVR, many namespaces and hops (fip-, snat-, qrouter-).
<b>L3 DISTRIBUTED North/South (FIP)</b>	OpenFlow, SNAT traffic through the networker nodes NAT using OVS connection tracking	L3-agent / Linux kernel / many namespaces and hops (fip- snat-, qrouter-) SNAT through networker nodes
<b>DHCP</b>	OpenFlow (controller action), Response from comp. local to the instances  Fully distributed HA on compute nodes.	Response from networker nodes  dhcp-agent / dnsmasq + qdhcp- namespaces



# Comparing ML2/OVN and ML2/OVS (III)

	<b>ML2/OVN</b>	<b>ML2/OVS</b>
<b>ENCAP.</b>	Geneve, VLAN <sup>3</sup>	VXLAN, GRE, VLAN
<b>Agents</b>	ovn-controller (C <sup>1</sup> + N <sup>2</sup> ) ovn-metadata-agent (C)	Neutron-l3-agent (C + N) Neutron-dhcp-agent (N) Neutron-metadata-agent (C + N) Neutron-openvswitch-agent (C + N)
<b>IPv6</b>	OpenFlow RA, RS, ND, NS handled locally in compute nodes	Neutron-l3-agent + radvd (N)
<b>L4 Load Balancing</b>	Octavia driver, handles distributed L4 Load Balancer in OpenFlow	No
<b>Internal DNS</b>	OpenFlow	Neutron-dhcp-agent + dnsmasq (N)


# Performance: Controlplane



# Performance: Dataplane


## TCP STREAM THROUGHPUT

Higher is Better




## UDP Packets

Higher is Better

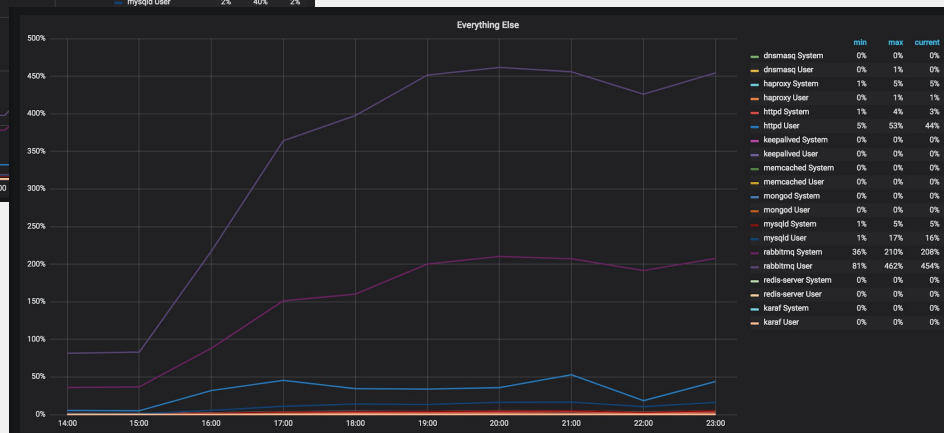


# Performance: CPU utilization



ML2/OVN

ML2/OVS



# What's next?

- ML2/OVS to ML2/OVN migration tool
- ML2/OVS parity: QoS, SG logging, ...
- Split OVN from OVS project for better agility and independence
- Adopt Raft OVSDDB clustering (A/A)
- Performance: Incremental processing for ovn-northd and ovn-controller

# Q&A