



THE STATE OF THE OVN

Baking tasty virtual networks.

Russell Bryant - @russellbryant

November 16, 2017

MY PERSPECTIVE

- OVN Contributor and Committer
- OpenStack contributor and committer
- Red Hat CTO Office

THE OVN JOURNEY FOR OPENSTACK

WHAT PROBLEMS DID WE WANT TO SOLVE?

Why OVN?

- I started with a focus on OpenStack as the consumer.
- Can we do something for OpenStack that ...
 - Meets functional requirements
 - Ideally zero-to-minimal dependency / complexity addition
 - Is easy enough to integrate with
 - Is reusable outside of OpenStack

- And off we went, starting early 2015 ...

L2 VIRTUAL NETWORKS

- A Neutron network == OVN Logical Switch
- From zero to ping in 6 weeks.
- Live demo at OpenStack Summit in Vancouver (Spring 2015)
 - Bonus: included container-in-vm modeling this early!
- What is different?
 - All OpenFlow, no NORMAL action used.
 - Neutron OVS agent replaced by ovn-controller.
 - Local ARP responder is automatic, not configured and optional. (I2pop in OpenStack)

SECURITY GROUPS

OVN ACLs

- OpenStack Security Groups == OVN ACLs
- A distributed firewall.
- Implemented by Fall 2015
- Built on top of OVS conntrack integration.
 - More efficient than earlier OpenStack iptables usage, but similar to OVS firewall driver.
- OVN L2/L3/L4 ACLs are much more flexible than security groups
 - Great for re-use by different security models in other systems!
- More recent value add: ACL logging!

L3 LOGICAL ROUTERS (v4 and v6)

East/West routing

- Neutron router == OVN Logical Router
- Implemented by Fall 2015
- Distributed, not something configured and optional.
- Implemented using OpenFlow
 - ovn-controller starts replacing job of Neutron L3 agent
 - Flow caching provided a performance boost here
 - Works with OVS-DPDK

L2 GATEWAYS

- OVN added hardware_vtep L2 gateway support by Fall 2015.
- Interesting to OpenStack, but never really a focus ...
 - Early POC integration made available
 - Still have not worked on networking-l2gw API integration
- OVN later added software (ovn-controller) based L2 gateway support.

OVS 2.5 - February 2016

First experimental release of OVN

PROVIDER NETWORKS

- A Neutron network can be a virtual network, or a mapping to a physical one.
 - OVN Logical Switch with attached “localnet” port
- Can map Flat or VLAN networks into OVN.
- Can attach ports directly to these networks
 - OVN still useful for managing security policies here.
- Define gateways (L2 or L3) from virtual networks to these physical networks.

L3 GATEWAYS

- L3 gateways between logical and physical networks.
- NAT supported using OVS conntrack integration
 - OpenStack: both SNAT and floating IPs
- Like everything else, defined in OVN pipelines and implemented with OpenFlow

LOAD BALANCING

- OVN added distributed L4 load balancing
- Not a feature parity issue for OpenStack, but valuable new thing!
- OpenStack patches just now in progress to make use of this
 - Without it, OpenStack LBaaS falls back to service VMs running haproxy

OVN TRACING

ovn-trace

- Allows you to trace a sample packet through OVN's logical pipelines to determine how it would be processed
- Output at varying levels of detail.
- Super helpful in development, and for learning the system.
- Also helps debugging by starting at tracing at a higher layer
 - Can more quickly identify if observed behavior is also what OVN expects based on current configuration
- A really nice value add!

NATIVE DHCP (v4 and v6)

- OVN can intercept DHCP requests from logical ports and respond to them to do VM bootstrapping
- For OpenStack
 - Replaces Neutron DHCP agent, which managed dnsmasq processes to do this instead
 - An improvement: one less agent and external dependency
 - Removed need for potentially *many* dnsmasq processes

OVS 2.6 - September 2016

First Supported Release of OVN

- Major pieces in place
- Start to do more polish and improvements

OVS 2.7 - February 2017

- Simplified modeling of routers that have both distributed (east/west) and centralized (NAT/gateway) roles
 - Big simplification for OpenStack OVN integration
- QoS (bandwidth limits, DSCP marking) added - used by OpenStack
- Source IP based routing policy support for L3 gateways
- Basic IPAM sees some minor enhancements
- ovn-trace enhancements (OpenFlow, DHCP support)
- SSL config enhanced / simplified
- performance testing (control and data plane)

OVS 2.8 - August, 2017

The OVN is really cooking!

- L3 Gateway High Availability
- Native internal DNS support
- Basic IPAM sees enhancements
- ovn-trace improvements (supports ACLs)
 - “Why can’t I reach my VM?! Oh, because an ACL drops it.”
- OVSDB RBAC allows to start locking down ovn-controller db access
- ACL Logging mentioned earlier was added here

L3 GATEWAY HIGH AVAILABILITY

- <http://docs.openvswitch.org/en/latest/topics/high-availability/>
- Define a prioritized list of chassis (hosts) for a gateway
- Each node that may send traffic to that gateway monitors connectivity to each gateway host using BFD
- Traffic sent to highest priority reachable gateway host based on BFD status
- Gateway hosts also monitor each other. A lower priority gateway host will take over if all higher priority gateway hosts are not reachable and it's still able to

WHAT HAPPENS NEXT

- OVSDB clustering - active/active HA
 - Currently support active/passive HA with standby replicas
- Closed a few lingering IPv6 gaps
- ACL logging performance improvements
- More control plane performance testing and enhancements
- OVN tunnel encryption
- Port mirroring
- Service Insertion (SFC)

BACK TO THE OPENSTACK PERSPECTIVE

- I'm feeling very satisfied with what we have as covering all of our base requirements.
- It's reusable, and offers benefits over the past iteration of OVS integration for OpenStack.
- The next most important things are being worked on.
- Most new features provide new value to OpenStack vs catching up
- THANK YOU to everyone who has helped make this happen.

ON REUSABILITY

- ovn-kubernetes
- Docker
- Mesos (experimental, at least)
 - Container integration talk up next!
- Nutanix
 - The talk after that!
- oVirt
- We'd love to talk to more users!

WHERE TO LEARN MORE

- OVN
 - <http://docs.openvswitch.org/>
 - OVS discuss and dev mailing lists
 - #openvswitch on Freenode IRC network
- OpenStack+OVN
 - <http://docs.openstack.org/developer/networking-ovn/>
 - openstack-dev mailing list
 - #openstack-neutron on Freenode



THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHatNews



youtube.com/user/RedHatVideos