10 YEARS OF OPEN VSWITCH SUCCESS AND FAILURE

BEN PFAFF

WHAT IS OPEN VSWITCH?

Open source project

- Software switch
- Portable
- Programmable
- Fast

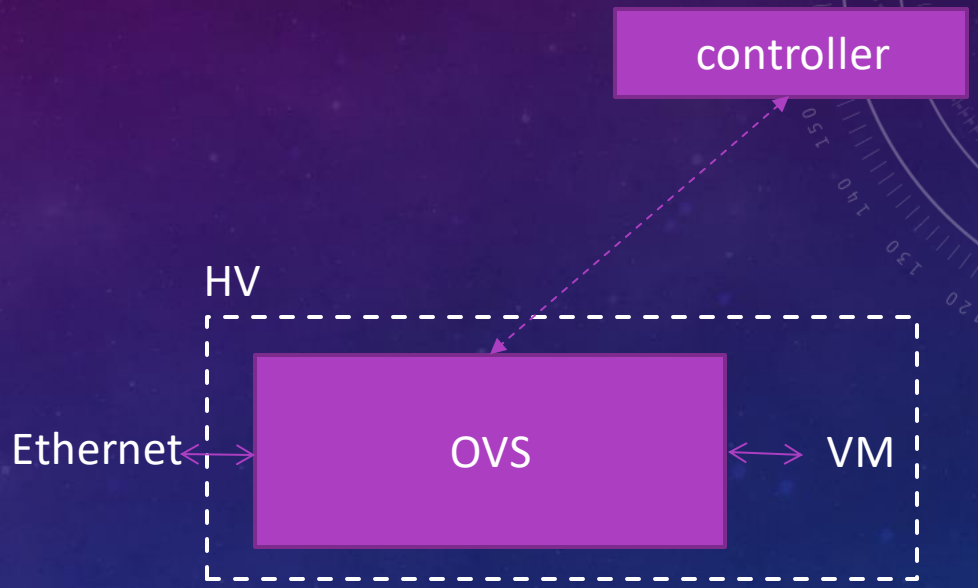

OPEN VSWITCH SUCCESS

- Used in NSX on KVM, Windows, and public cloud
- Widely used in Xen, KVM, OpenStack
- Incumbent targeted by new projects
- Over 5,000 academic citations

LET'S CELEBRATE!

- 2018: ACM SIGCOMM SOSR Software Systems Award.
- 2015: NSDI "Best Paper" Award for "The Design and Implementation of Open vSwitch".

KEYS TO SUCCESS

PEOPLE, TECHNOLOGY, AND ENVIRONMENT

ENVIRONMENT

- Open source virtualization was exploding, but there was no good virtual switch.
- Open vSwitch was in the right place at the right time. It filled a gap.
- Nicira founders had relevant academic and business connections.
- Nicira founders provided a valuable vision.

PEOPLE

- OVS had the best people: Justin Pettit, Jesse Gross, Ethan Jackson.

  Martin Casado: "You end up putting basically your best development resources... on the open source side of the house, which from a company is kind of difficult to do, because you basically want that best developer to work on the core product on the inside." (on "a16 Podcast: Monetizing Open Source", April 10, 2017.)

- Nicira had PhD-type innovators:
  - Links to academics.
  - Willing and able to do library research and recognize where it was valuable.
  - Able to step back and consider larger problems.
- Kind (by and large).

TECHNOLOGY

- We didn't know what Nicira needed, so we built OVS to do anything.
  - No one would want NSX-specific agents except NSX customers.
- Lots of open source projects are aimed at one company's use case. OVS was not.
- We did not have to reject user feedback and patches that made OVS better for use cases outside of "ours." The criteria are simpler: "Will this make OVS more general purpose?"

FAILURES

RELATIONSHIPS!

OPENFLOW

OpenFlow is the protocol used to program Open vSwitch. Nicira invented it and then worked through standardizing it.

Lessons:

- Don't assume SDOs will fix anything.
- SDOs will happily ignore reality.
- SDOs are just collections of individual people and those particular people matter a great deal.
- SDOs can hijack or kill your project if you don't fight for it.

Status: OpenFlow is an unmaintained standard.

HARDWARE

Open vSwitch was supposed to support hardware switches from the beginning, but:

- Most switch ASIC companies are as closed as possible.
- Most switch vendors see open source software as read-only.
- SDOs cater to both of the above behaviors.

We tried numerous times to launch hardware projects based on OVS. All of them failed.

- We hoped OVS and OpenFlow would encourage ASIC and switch vendors to open up. They did not.
- Switch vendors often use OVS as OpenFlow agent. They don't usually contribute back but do complain.
- OVS comes with a simple API for VXLAN configuration for hardware switches. It isn't OVS, though.

Status: OVS hardware support is fractured across vendor implementations.

KERNEL

OVS has a kernel module that requires development and maintenance.

The relationship between the OVS team and the kernel network stack developers has been rocky at times.

The core kernel network developers are concerned about losing control of networking on Linux.

OVS developers have been focused on OVS and not necessarily on the kernel.

Status: OVS developers continue to try to be good kernel citizens.

COMMUNITY DECENTRALIZATION

OVS was originally just Nicira. We encouraged outside contributions but received few.

- In 2011, ~97% of commits were internal.
- In 2017, ~50% of commits were internal.

Success?

2017 Commits by Contributor:

1825 total
904 VMware
299 Red Hat
140 Cloudbase Solutions
123 Intel
74 Mellanox
56 Samsung
32 erig.me
31 opencloud.tech
31 Ericsson
17 eBay
14 Huawei
10 ZTE
9 DT Dream
7 Wind River
78 other

COMMUNITY DECENTRALIZATION

OVS was originally just Nicira. We encouraged outside contributions but received few.

- In 2011, ~97% of commits were internal.
- In 2017, ~50% of commits were internal.

Success? Maybe:

- 392 (21%) of the 2017 commits were mine.
- 6944 (40%) of all commits are mine.
- 9757 (57%) of all commits were committed by me.

Delegation is hard and I am only learning slowly.

2017 Commits by Contributor:

1825 total
904 VMware
299 Red Hat
140 Cloudbase Solutions
123 Intel
74 Mellanox
56 Samsung
32 erig.me
31 opencloud.tech
31 Ericsson
17 eBay
14 Huawei
10 ZTE
9 DT Dream
7 Wind River
78 other

FUTURE

MOVING UP THE STACK

NETWORK VIRTUALIZATION: OVN

OVN, the Open Virtual Network, is a subproject of Open vSwitch for network virtualization.

It is starting to get adoption in some areas and by some companies.

VMware is considering using it in at least one product.

The OVS project has been doing a poor job of promoting it. We are working on that,

and on making it easier to install.

CONTAINERS

Everyone is talking about Kubernetes these days.

OVN integrates with Kubernetes and supports all Kubernetes networking features, but it isn't well known;

Calico has taken the lead.

Container folks don't care much about networking; do the power and scale of OVN matter to them?

Does Open vSwitch matter in a container world?

SERVICE PROXY

Envoy and linkerd are "service proxies" that act as application-layer L7 switches.

They are exploding in popularity in the container world.

When everything happens at L7, is there value in an L2/L3/L4 virtual switch?

Should OVS attempt to move to a higher layer?

Should we start a new project that acts at this layer?

QUESTIONS?