A Walk through Ciena's experience with Trellis

CORD Build 2017

ciena

CORD
Central Office Re-architected as a Datacenter

Background

Started with investment with ODL

Shifted investment to ONOS 2015

Focused on CORD 2016

- Automation of physical POD build out
- Compute node network configuration

Leaf - Spine Fabric

Dynamic network sizing

One of the most interesting technologies or CORD

Resourcing not reflective of importance

Does CORD need a fabric abstraction?

Infancy

Manual configuration

No dynamic updates

Frequent fix was complete restart

Childhood

Container based service to generate configuration

- Queried ONOS for information on hosts, ports, devices
- Generated a configuration that could be pushed to ONOS

Did not address inability of SR to process dynamic updates

Required hosts to ping ONOS before generation

Solution is Complex

- It is not just the SR configuration
- Default gateway when there is both data and control networks
- https://github.com/opencord/cord/blob/master/docs/appendix_basic_config.md

Adolescence

Configuration generator not really used

- Prerequisites were an issue
- Not kept up to date with Trellis changes
- Often used to hand generate initial configuration followed by hand editing

Troubleshooting difficult

- Restarting sometimes fixes things
- Walking flow tables in ONOS
- Can’t be sure if it is OFDPA, ONOS, phase of the moon
- Configuration works on Monday not on Wednesday
- `@charles` a big help!

Early Adulthood

Better Integration with XOS

- XOS use of configuration generator

Host subnet configuration (Where is DHCP Relay?)

- Container to allocate fabric address to compute hosts in CORD, deeply flawed

More Configuration Changes

- Keeping configuration generator up to date

Stability still a Hit and Miss

Midlife

Mostly ignored SR

- Got it working
- Fixed it when it broke

Rinse and Repeat

- Kill and reset ofdpa on leaf1, leaf2, spine1, spine2
- `app deactivate org.onosproject.segmentrouting`
- `wipeout please`
- `POST nework-cfg.json` (thanks `curl`)
- `app activate org.onosproject.segmentrouting`
- Start OFDPA and connect (leaf1, leaf2, spine1, spine2) to ONOS
- Repeat on failure

Mature Adulthood (cont)

Goal is to Demonstrate Fully HA POD

- Customer (behind OLT) to Internet
- Fabric (+ Connection to Internet)
- Physical hosts
- Network switches and OLT/ONU
- Container orchestration (k8s)
- NFVs
  - vSG
  - Spanneti
  - ONOS

Would like to use fabric for control and data traffic

Mature Adulthood

Trellis under a Container Management System (k8s)

- Experimenting with reactive network programming as part of container brigade
- Network configuration generated and hand edited
- 2 x Leaf Switches, 2 x Spine Switches
- 3 Compute nodes behind each switch
- 1 OLT/ONU

Single ONOS Instance (1.10.9)

- ONOS as a Service (VIP)
- Single ONOS instance behind service (currently)
- ONOS can run on any compute node

Mature Adulthood (cont)

In Support of Fabric

- ONOS auto starts as service with single instance
- `ofconnect` service on switches to auto connect to ONOS
- Configuration Push Container
  - Periodically poll ONOS for network configuration
  - Reconcile and Update
  - Static network means no dynamic updates
- Compute node network configuration fixed
  - IP addressing
  - Static cross-fabric routes

Mature Adulthood (cont)

Still Experiencing Stability Issues

- Fabric stopped working 2 weeks ago, unexplained
- Tried restart procedure (several times) no joy

Need to dump troubleshooting data

- Flows, devices, hosts, etc.
- `@charles` is not a scalable solution

Frustration

- When it works, life is good
- When it doesn’t work, can be hard / unexplainable how to fix it
- It “should” just work

Looking Toward the Future

DHCP

- Currently compute node IP addresses are statically assigned
- Move to DHCP using DHCP relay based on leaf subnet
- Existing in current Trellis, needs to be tested

XOS (orchestration really)

- Allocate a subnet to a leaf when it comes on line
- Update DHCP server, SR config, accordingly
- Should a new switch / compute node be governed (orchestrated) as a XOS service?

Take-a-ways

Trellis represents and import network capability

Maturing over time, but not ready for production

- Combination of hardware, software, configuration, and controller issues

Need more effective troubleshooting tools

- Packet traces through ONOS flows

Configuration should be mostly automated

- Some debate on this (intentional v. discovered)
- Security issues

Mèsi Anpil