Open Networking (white box) in the Enterprise

One Year in. Lessons learned, and way ahead.

Why am I here?

- To share Open Networking experiences from an enterprise perspective (non hyperscale)

- Matt Turner Bio
  - CCIE 16857 (Emeritus) Routing and Switching
  - Data Center Network Manager at Qualcomm Inc.
- Qualcomm Network Bio
  - 30+ data centers (~850 switches, spine/leaf topologies)
  - Many LAN & LAB switches (~2700)
  - Dedicated “NetDevOps” team ☺

What is Open Networking?

- Disaggregation, White Box, VNF’s, controllers, ONF?
  - Depends who you’re talking to.
  - For Qualcomm, Open Networking is White/Brite Box+ONIE+Software
- ONIE = Open Network Install Environment (OCP open source initiative)
  - Cumulus
  - Big Switch Monitoring Fabric
  - OpenSwitch (OPX)
  - SoNic
  - JunOS

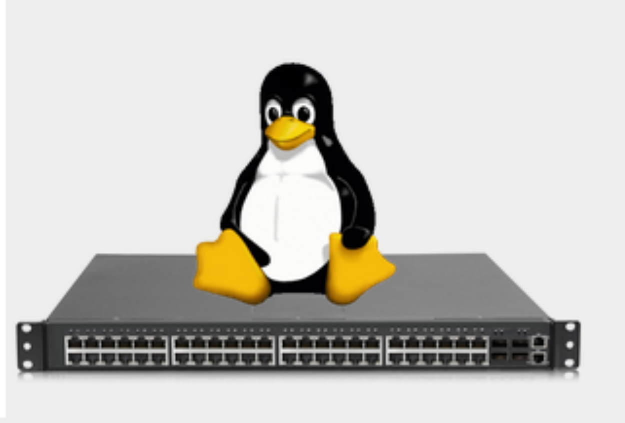

Why Open Networking

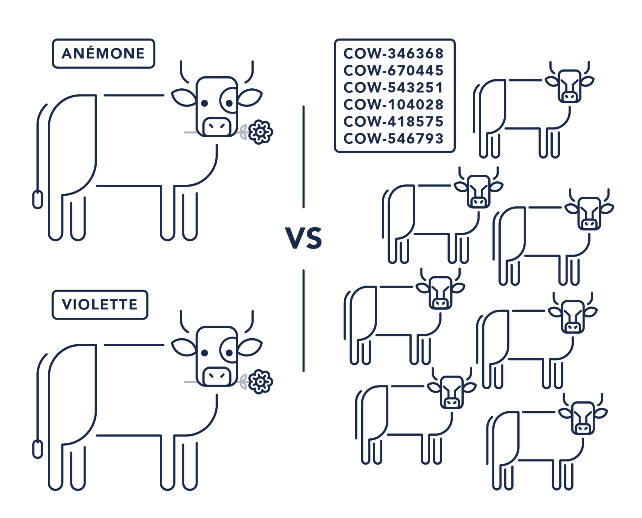

- $uper exciting!
  - Roughly 33% the cost of traditional networking
    - (discounted rate)
- Disaggregation allows flexibility
  - Big Switch BMF and Cumulus today, tomorrow?
- Linux is easier to automate than Cisco/Arista/Junos/etc
  - Ansible/Chef/Puppet built for Linux, adapted for networking
- Great way to transition from **pets to cattle** approach for network switch provisioning and MGMT
- Open Linux platform (install collectd if you like…)

Lots of Lab Testing and Evaluation…

- Decided on Cumulus for networking, Big Switch Monitoring Fabric
- Cool network features
  - BGP/OSPF Unnumbered (IPv6 link local peering)
  - BGP Redistribute Neighbor (redistribute ARP table into BGP /32 routes)
  - Cumulus NCLU (meh… for some, CLI alternative for others)
- Cool monitoring fabric features
  - OpenFlow (behind the scenes) controller based
  - ZTP/DHCP capable

Which Hardware?

- Accton/Edgecore hardware initially, Dell hardware now
  - Snap in rack rails!
- Common hardware on vendor HCL's
- Keep spares in stock vs purchasing hardware support
- Support for many brands of optics and cables
- Same chips, CPU as traditional vendors
  - Broadcom ASICs, Intel or AMD CPU, etc.

Building Blocks for Success

- ONIE, Zero Touch Provisioning (ZTP)
  - ONIE boot, ZTP using DHCP options and default URL (114)
- Git, GitHub
  - Version control for ZTP, operations playbooks, global switch configurations
- Jenkins
  - CI/CD platform for centralized Ansible controller
  - Splunk logging, RBAC, store credentials, cron, GUI!
- Ansible (or Chef, Puppet, Salt)
  - We prefer Ansible for use with legacy vendor hardware/OS (agentless)

Framework – GitHub/Jenkins/Ansible

- Initially deployed for Open Networking (Cumulus)
- Playbooks stored in GitHub for version control, change MGMT, and code/peer review
- Playbooks run from Jenkins for centralization, security, auditing, logs, etc. (logs all jobs and results to Splunk)
- Ansible and associated plugins/modules installed on Jenkins server

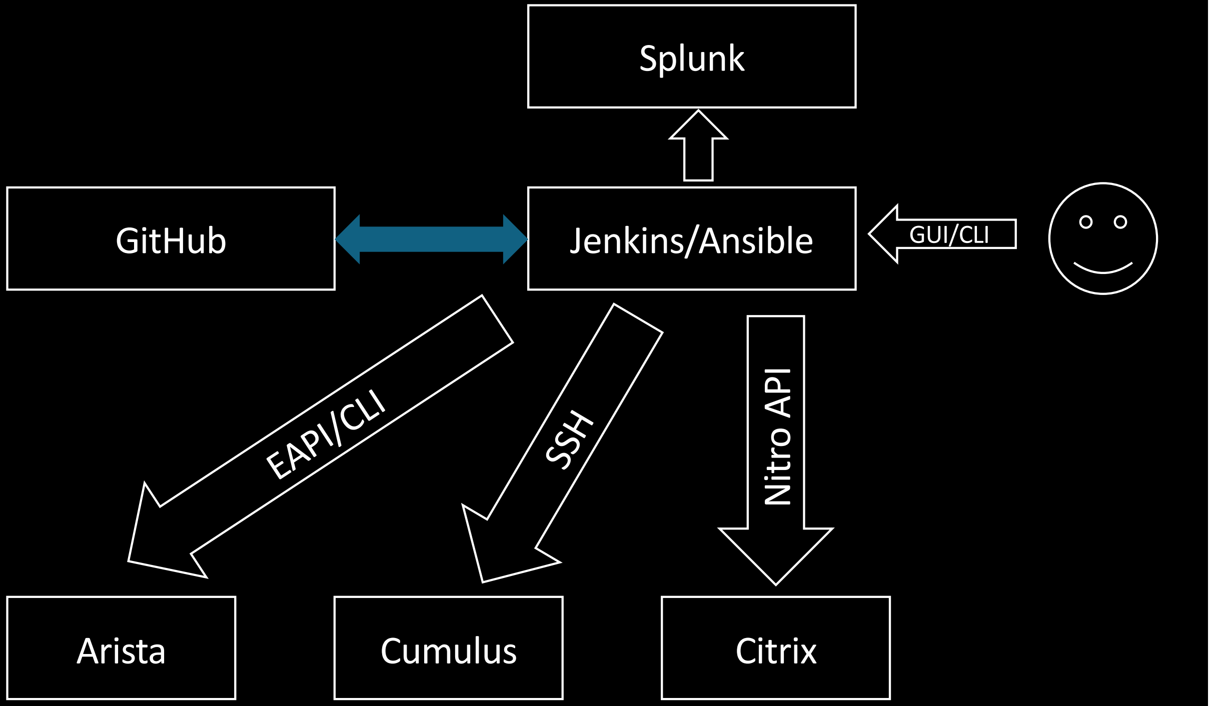

What We Automate

- Almost everything…
- ZTP for bring up
    - DHCP MAC reservation, DHCP default URL for image load
- Ansible for initial configuration
- API for user self service (rack and stack team, server/storage admins)
    - Add/change VLANs for access ports
    - Create MLAG
    - Add/change VLANs for existing MLAG ports
- Ansible for weekly global configuration compliance (declarative, no audit needed)
    - E.g. NTP servers shall be x, y, z

Do Automation Day One!

Zero Touch Provisioning

```
subnet 192.168.0.0 netmask 255.255.255.0 {
  range 192.168.0.20 192.168.0.200;
  option domain-name-servers 192.168.0.2;
  option routers 192.168.0.3;
  option default-url = "http://10.0.0.10/customer-abc-onie-installer";
```

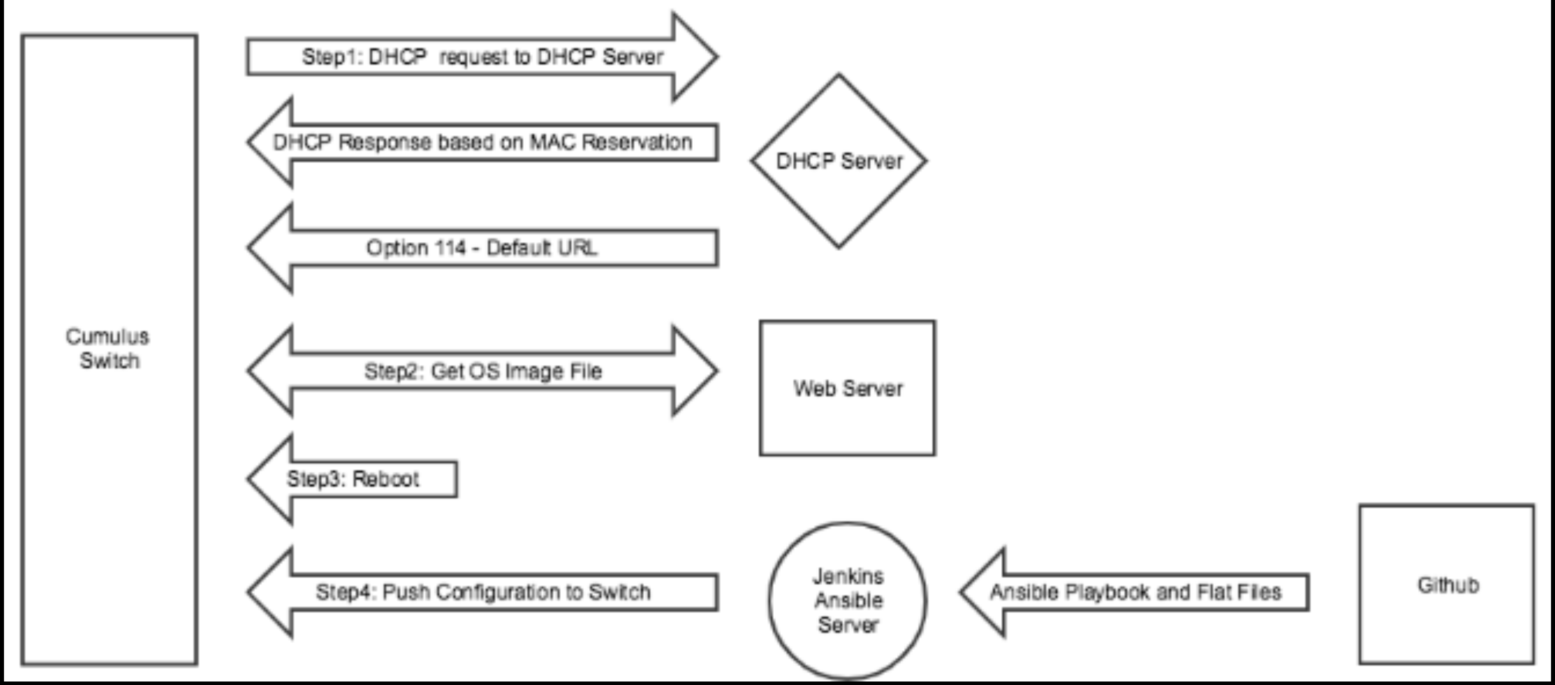
Cumulus Switch
Step1: DHCP request to DHCP Server
DHCP Response based on MAC Reservation
DHCP Server
Option 114 - Default URL
Step2: Get OS Image File
Web Server
Step3: Reboot
Step4: Push Configuration to Switch
Jenkins Ansible Server
Ansible Playbook and Flat Files
Github

ONIE Boot – ZTP

```
Info: Mounting ONIE-BOOT on /mnt/onie-boot ...
Info: Mounting EFI System on /boot/efi ...
Info: Using eth0 MAC address: 3c:2c:30:38:ed:00
Info: eth0:  Checking link... up.
Info: Trying DHCPv4 on interface: eth0
ONIE: Using DHCPv4 addr: eth0: 10.1.19.221 / 255.255.255.224
<SNIP>

Please press Enter to activate this console. Info: eth0:  Checking link... up.
Info: Trying DHCPv4 on interface: eth0
ONIE: Using DHCPv4 addr: eth0: 10.1.19.221 / 255.255.255.224
ONIE: Starting ONIE Service Discovery
Info: Fetching http://10.43.255.182/cumulus/cumulus-linux-3.7.0-bcm-amd64.bin ...
[   21.497593] random: crng init done
ONIE: Executing installer: http://10.43.255.182/cumulus/cumulus-linux-3.7.0-bcm-amd64.bin
Verifying image checksum ...OK.
Preparing image archive ... OK.
<SNIP>

Please reboot to start installing OS.
ONIE: NOS install successful: http://10.43.255.182/cumulus/cumulus-linux-3.7.0-bcm-amd64.bin
ONIE: Rebooting...
```

Framework

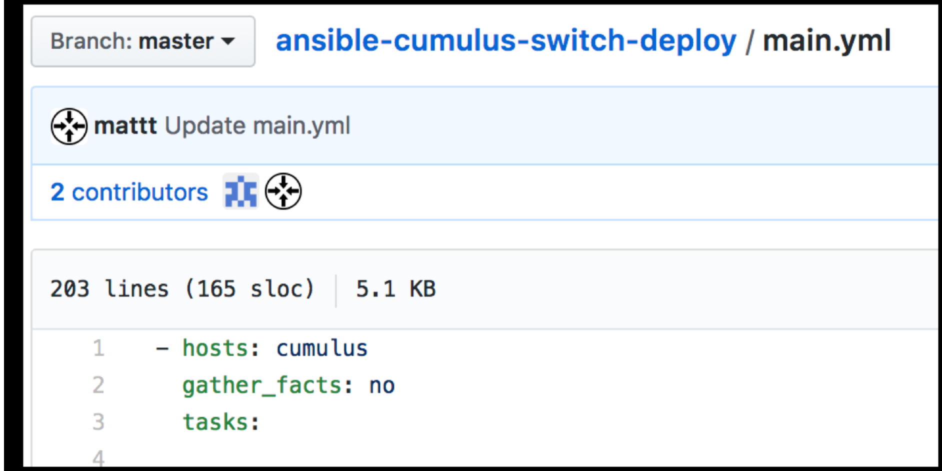

Branch: master ▾ **ansible-cumulus-switch-deploy / main.yml**

mattt Update main.yml

2 contributors

```
203 lines (165 sloc)  |  5.1 KB
```

```yaml
- hosts: cumulus
  gather_facts: no
  tasks:
```

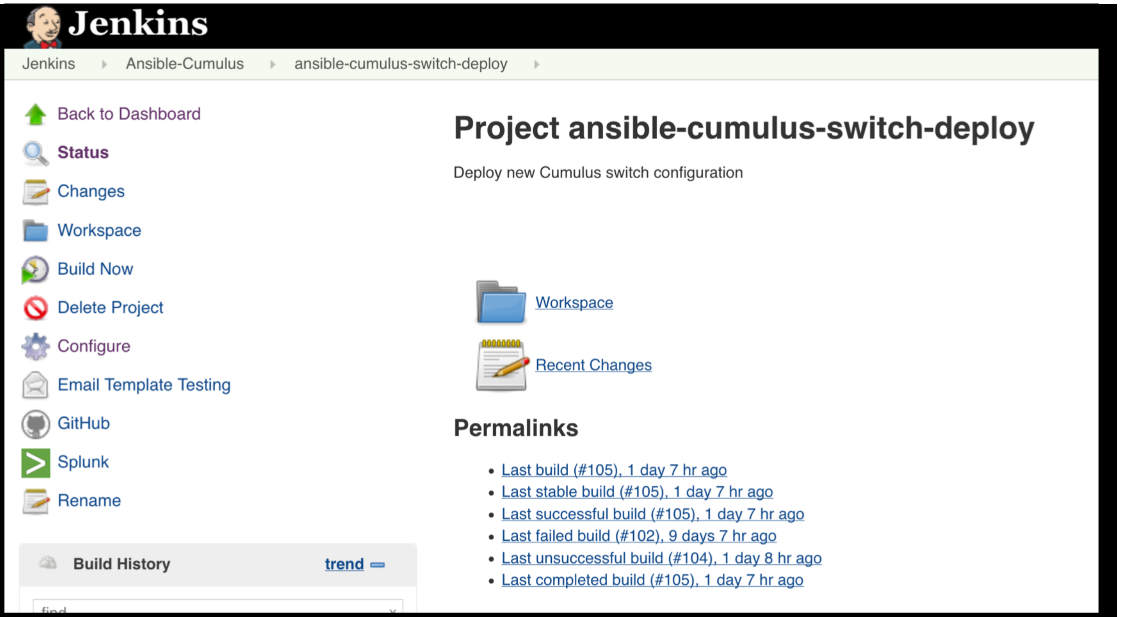

Project ansible-cumulus-switch-deploy

Deploy new Cumulus switch configuration

Workspace

Recent Changes

Permalinks

- Last build (#105), 1 day 7 hr ago
- Last stable build (#105), 1 day 7 hr ago
- Last successful build (#105), 1 day 7 hr ago
- Last failed build (#102), 9 days 7 hr ago
- Last unsuccessful build (#104), 1 day 8 hr ago
- Last completed build (#105), 1 day 7 hr ago

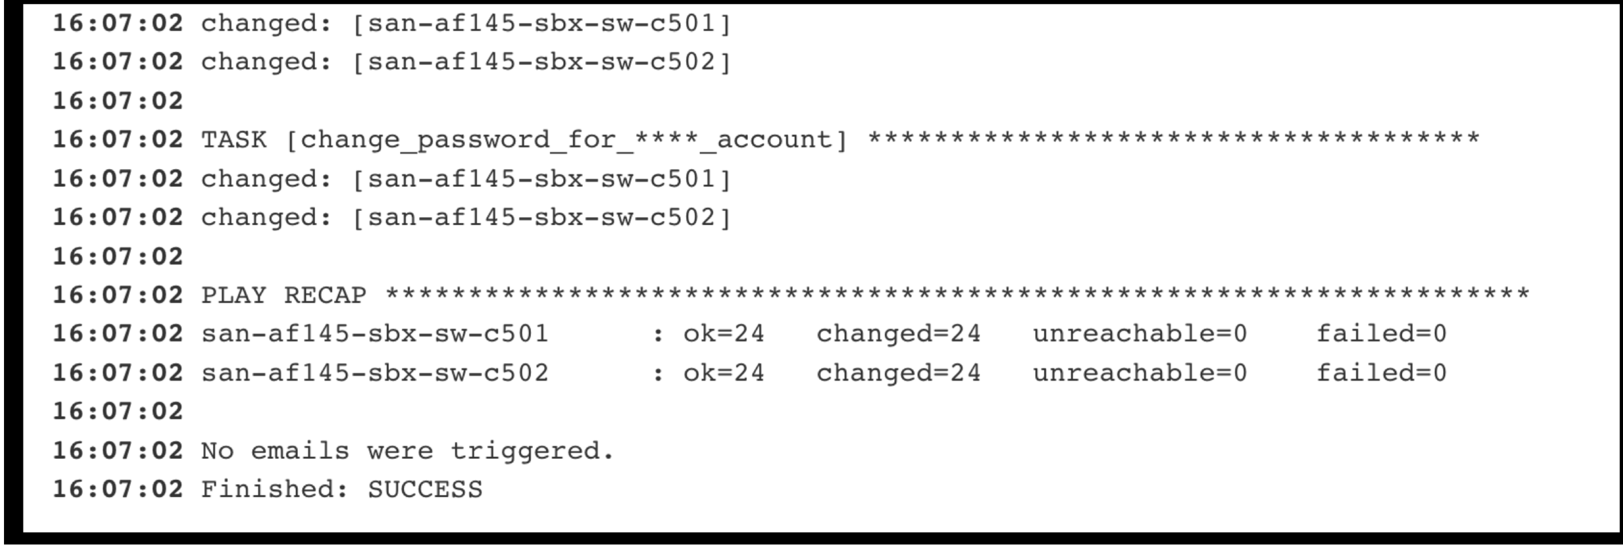

```
16:07:02 changed: [san-af145-sbx-sw-c501]
16:07:02 changed: [san-af145-sbx-sw-c502]
16:07:02
16:07:02 TASK [change_password_for_****_account] *************************************
16:07:02 changed: [san-af145-sbx-sw-c501]
16:07:02 changed: [san-af145-sbx-sw-c502]
16:07:02
16:07:02 PLAY RECAP ***********************************************************************
16:07:02 san-af145-sbx-sw-c501       : ok=24    changed=24    unreachable=0    failed=0
16:07:02 san-af145-sbx-sw-c502       : ok=24    changed=24    unreachable=0    failed=0
16:07:02
16:07:02 No emails were triggered.
16:07:02 Finished: SUCCESS
```

Day Two Automation – Self Service Tools

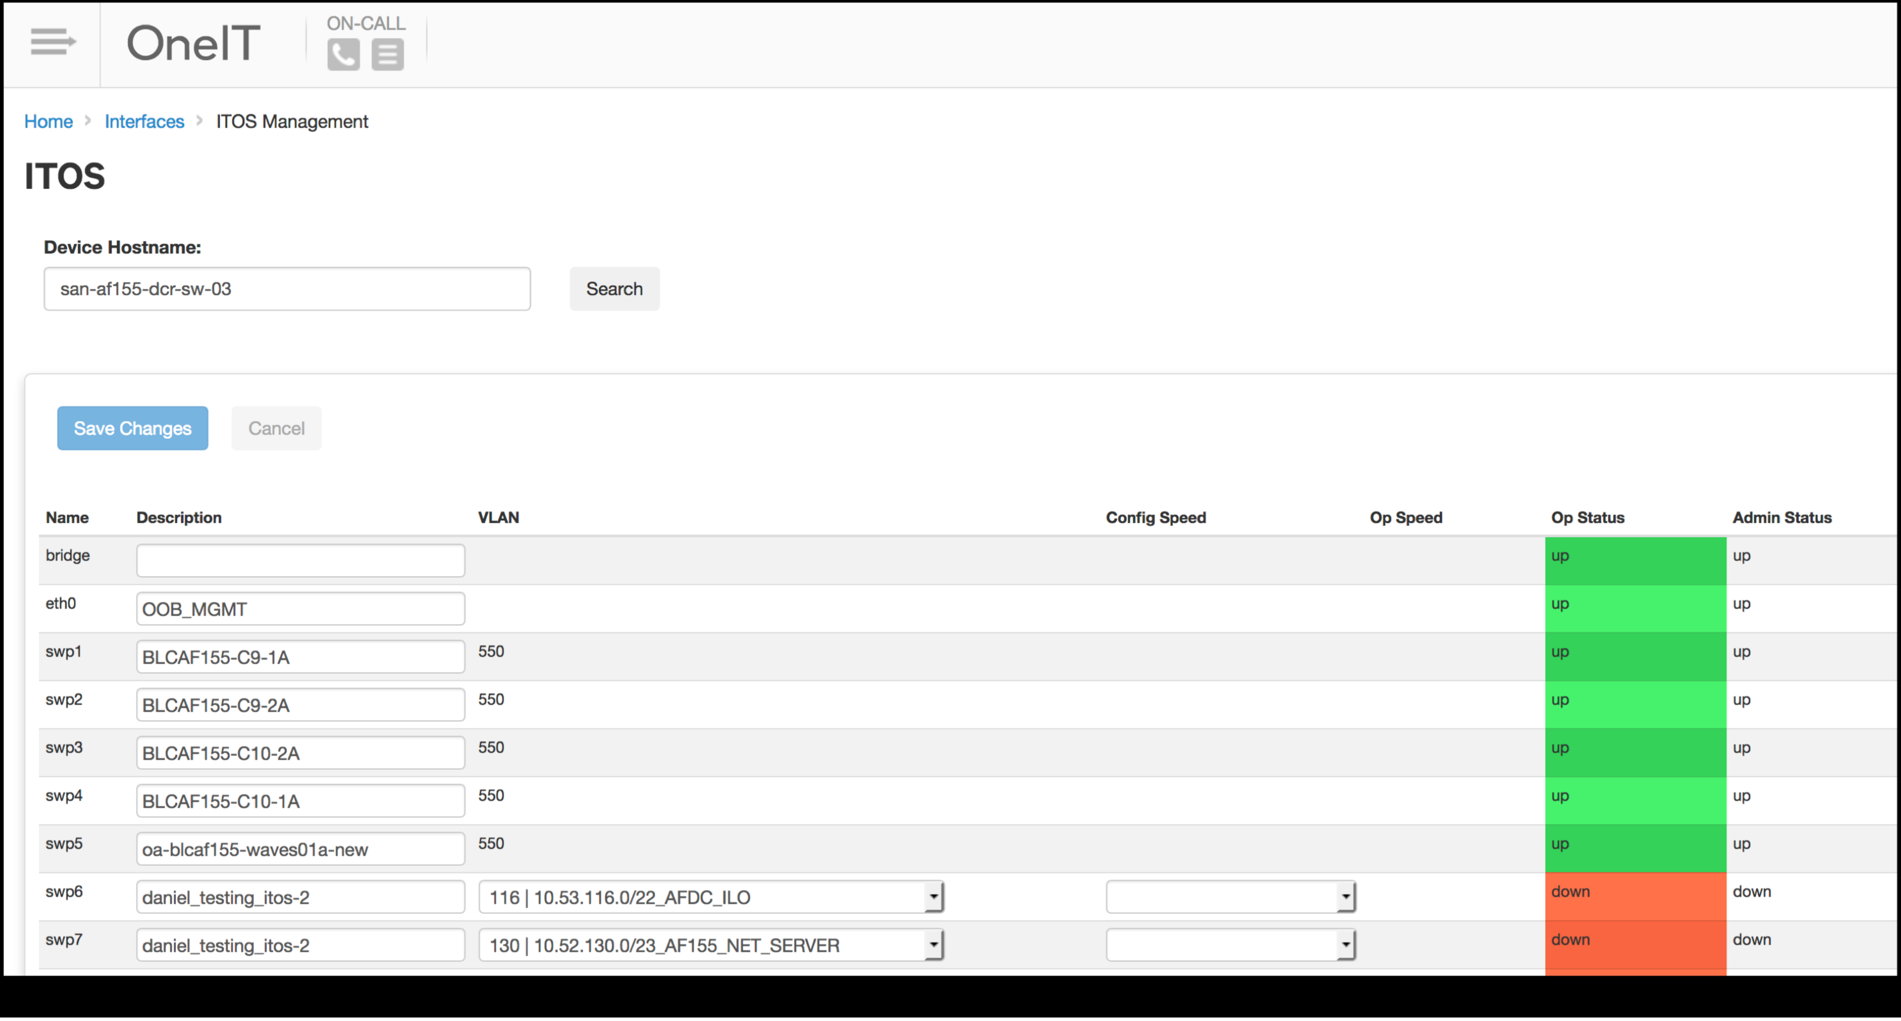

Home › Interfaces › ITOS Management

ITOS

Device Hostname: san-af155-dcr-sw-03

Name	Description	VLAN	Config Speed	Op Speed	Op Status	Admin Status
bridge					up	up
eth0	OOB_MGMT				up	up
swp1	BLCAF155-C9-1A	550			up	up
swp2	BLCAF155-C9-2A	550			up	up
swp3	BLCAF155-C10-2A	550			up	up
swp4	BLCAF155-C10-1A	550			up	up
swp5	oa-blcaf155-waves01a-new	550			up	up
swp6	daniel_testing_itos-2	116 | 10.53.116.0/22_AFDC_ILO			down	down
swp7	daniel_testing_itos-2	130 | 10.52.130.0/23_AF155_NET_SERVER			down	down

Obstacles to Overcome

- “Where’s my config-t?”
- Upper MGMT directors are/were CCIE’s, “Who do I call for support?”
- Legacy Network Management and Monitoring Tools
  - RSA/ACS – challenging to set up at first
  - SNMP – mostly works
  - Config Repo (HPNA Opsware for Cisco/Arista, GitHub/Jenkins for Cumulus)
- Change in mindset from a single config file, to Linux “net-sysadmin”
  - IMO this evolution needs to occur anyway for OpenStack, K8s, etc.. (Linux networking)

Non-Critical and Simple Deployments First

- OoB Data Center Network (switch mgmt.)- copper
- OoB Server Network (iLO/DRAC/MGMT) - copper
- Lab/Test/Dev Environments – fiber and copper
- LAN Access – copper PoE for fun and testing (works fine)
- Simple Critical Environments - HPC-LSF Top of Rack
  - Only requires BGP, LACP, MLAG
  - 80-96 servers per rack
  - QSFP Twinax cables to 4x25G SFP+

Test Network

- Have at least one…
- Vagrant/VirtualBox works well for us
  - Pre-canned topologies, stored in GitHub/GitLab
  - Great for learning, testing, planning for changes, developing automation
- Physical lab setup for optics, monitoring, etc. testing

Lessons Learned

- Adoption can be tough for seasoned network engineers
  - Need to learn Linux, Git/GitHub version control, CI/CD tools like Jenkins
  - Should learn Ansible/Puppet/Chef
  - Need to let go of the “config t”
- Linux experience very beneficial
- Automation required, day one
- Cattle instead of pets mindset
- Switch VM’s are great learning and testing tools
- https://github.com/mattincarlsbad

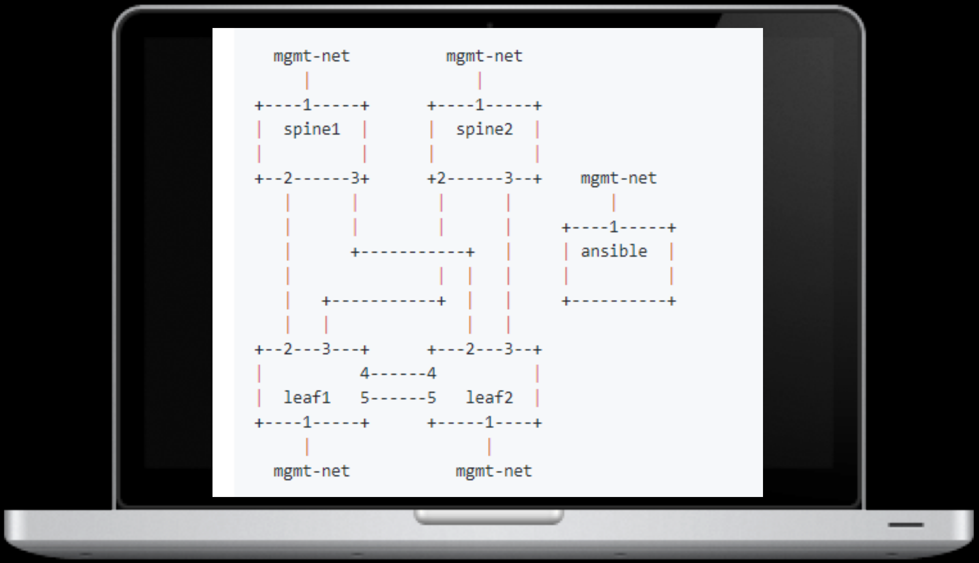

```
  mgmt-net          mgmt-net
     |                 |
+----1-----+      +----1-----+
|  spine1  |      |  spine2  |
|          |      |          |
+--2------3+      +2------3--+    mgmt-net
   |      |        |      |          |
   |      |        |      |     +----1-----+
   |      +-----------+   |     | ansible  |
   |               |  |   |     |          |
   |   +-----------+  |   |     +----------+
   |   |              |   |
+--2---3---+      +---2---3--+
|          4------4          |
|  leaf1   5------5   leaf2  |
+----1-----+      +-----1----+
     |                  |
  mgmt-net           mgmt-net
```

Conclusion

- Enterprises can:
  - Deploy and run white box switches
  - Save money by doing so
  - Usher in the new era of Linux networking

- As long as they…
  - Start in the lab
  - Start small
  - Don’t expect “config t”
  - Keep an open mind

Questions?

Pets vs Cattle…