


Smart Network Management

AOL's Take



Jay Moran
Distinguished Engineer,
AOL Technology Operations

Problems AOL Faces

~500 routers and ~6000 “switches” (~40k physicals, 15k++ VMs)

Scale and speed of deployment for changes; customer dependent

Ongoing & timely validation of network state

Extensibility and/or doesn't do specifically what we want

Integration costs/time of COTS sometimes equal to home grown

Multiple network vendors - Juniper, Brocade, Arista, Dell, HP, Cisco, and potentially L4-7 with Netscaler, A10, and Brocade ADX

Combination of human at the CLI and system driven configuration

Ops org (1st Tier NOC and 2nd Tier TAC) separate from Eng & Arch

Central Infrastructure Development org separate



SIMIAN – our in-house custom OSS-lite platform

Actually twiddles the bits on the devices

In House developed over the past few years

Originally designed to manage the rather complex ACL set

Based on Twisted Framework for Python

Moving all of our various Perl/Python/shell scripts to it

Multithreaded, event driven

Driver based for different NOS'

Greatly eases configuring all of our devices

We need common vendor interfaces; whatever... just do it!

Ties to the rest of our tools (ORB, AMIS, Remedy, etc.)



CMDB & Asset Management

ORB is AOL's internally developed CMDB

- All of AOL Tech Ops – Systems, Apps, Networks, etc.
- Integration at the project level across technologies
- Base information store; other infrastructure apps use & update
- Easy query/integration MySQL over HTTP (ATOMICS)
- Extended for our private cloud VMs; public cloud no problem
- Key to success at scale; damn cool to get exactly what you need

AMIS is AOL's COTS Asset Management System

- COTS works, but can be a pain
- Accepted by financial teams as asset records
- Should likely be replaced by ORB (personal opinion, finance disagrees)



Off the Shelf Configuration Components

RANCID

- Pretty straightforward RANCID implementation
- Generally “last resort” for config backups
- Most of config stored in CMDB type data sets elsewhere

JunOS Ops & Commit Scripts

- Various “known value” checking
- Automation of replicated information
- Using more, lots of ideas, implementing as we can, “op bgpstatus” example

Syslogd and Splunk

- Central syslogging is key to event correlation and other data mining
- Syslog at 60,000+ device scale is actually a pretty hard task, done right
- For log processing, we do have custom components, viewing is Splunk



Audit Tools

Arbor for NetFlow – planning, visibility, and security uses

Cariden for planning, TE, and event simulation

Cyclops & BGPlay for outside views

Daily/Weekly Reports for actioning

In house developed metrics collection – Argus

- 10M SNMP metrics, ~40M Argus metrics... EVERY minute
- SNMP, custom collection scripts, app events, log4j, etc.

In house developed metrics visualization - TOGA

- Argus' time series metrics overlaid, amazing correlation
- Lots of ex-AOLers call back wishing they had ORB/Argus/TOGA again

Hashed multi-server open source packet capture capabilities

- Brocade MLXe “10GE hub” and hashing capabilities
- Cheap, industry standard servers, multi-10Gbs, TBs of storage



Takeaways from AOL and my experience...

If your technology stack = revenue, invest in tooling!

No silos, leverage common tooling (CMDB, Metrics, etc) across technologies for benefits; do not recreate the wheel!

Leverage common frameworks for device drivers, etc.

Visibility across techs hard, worse at scale; but data can provide greater planning/decision making and aid in troubleshooting

At scale a lot of the pricing models would bankrupt you for COTS

As with all business decisions, use COTS when there is a clear ROI

Invest in network engineers who can help you automate;
all should be taught when to recognize helpful automation

