ANF SYSTEMS ARCHITECTURE

Antelope Users Group 2012
University of Nevada, Reno, NV

WHERE WE WERE

- Mid-2005
  - Single Sun V240 with 3.5 TB storage
  - SIngle Linux server for Web site
  - Single Datacenter

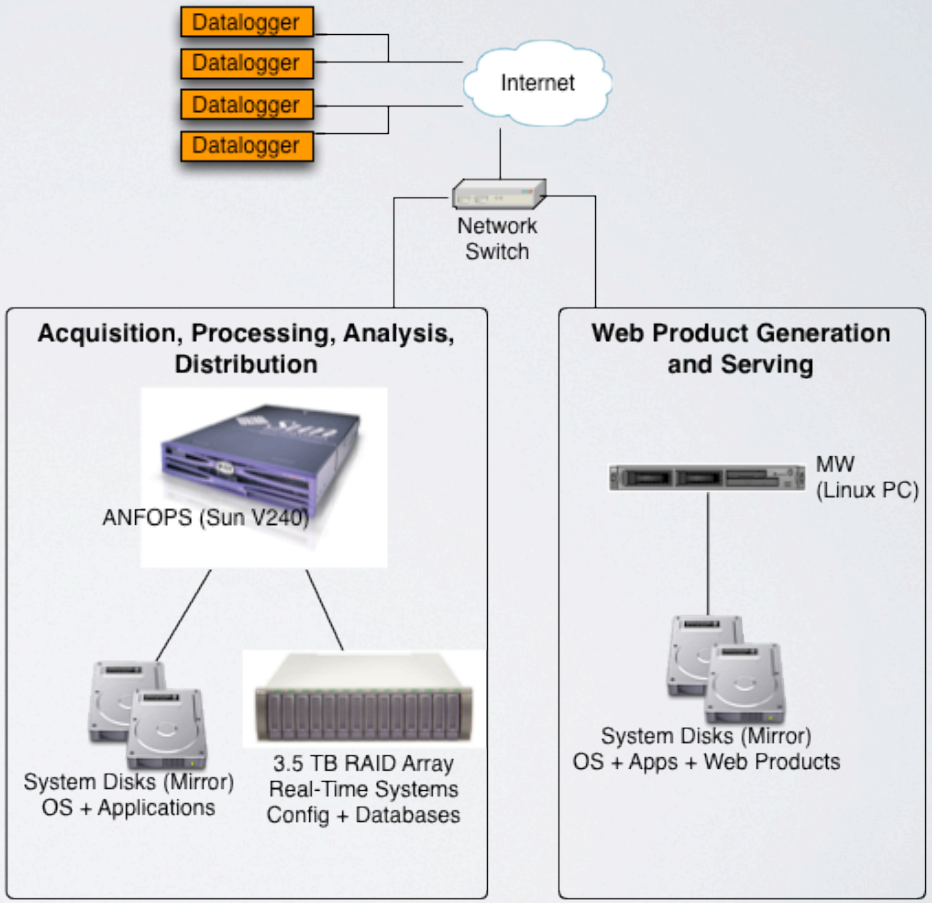

INFRASTRUCTURE AT MAIN DATACENTER

- 3 Dell R710 x86_64 servers,
  - 192 GB RAM, 512GB mirrored root disks, 2 port Fibre Channel HBA
- Redundant Cisco Switches
- Cisco VPN Gateway, Serial Console Server for Mgmt.
- 3 Apple Xserves
- Redundant QLogic SAN Switches
- Compellent Storage System for Block Storage
- Shared data via NFS
- Legacy Sun Hardware - 3 T5220 systems, 1 T2000 (for network monitoring)

VIRTUAL MACHINES AND ZONES

- Live on top of our physical infrastructure
- Primary OS disk and host specific data volumes on SAN storage
- Can be moved between physical servers with varying degrees of ease. VMware is automatic, Zones manual
- 15-20 CentOS Linux VMs (varies)
- ~15 Solaris Zones

BACKUP DATACENTERS

- Main DR site at IGPP
  - 2 Sun T2000 servers
  - 3 Nexsan SATABeast Storage Arrays
  - 1 AC&NC JetStor Array
  - Redundant LAN and SAN
  - Total ~160 TB backup storage
- Scorched earth site at IRIS DMC
  - Single Sun T2000 with 768 GB RAID5 internal disks

CURRENT INFRASTRUCTURE

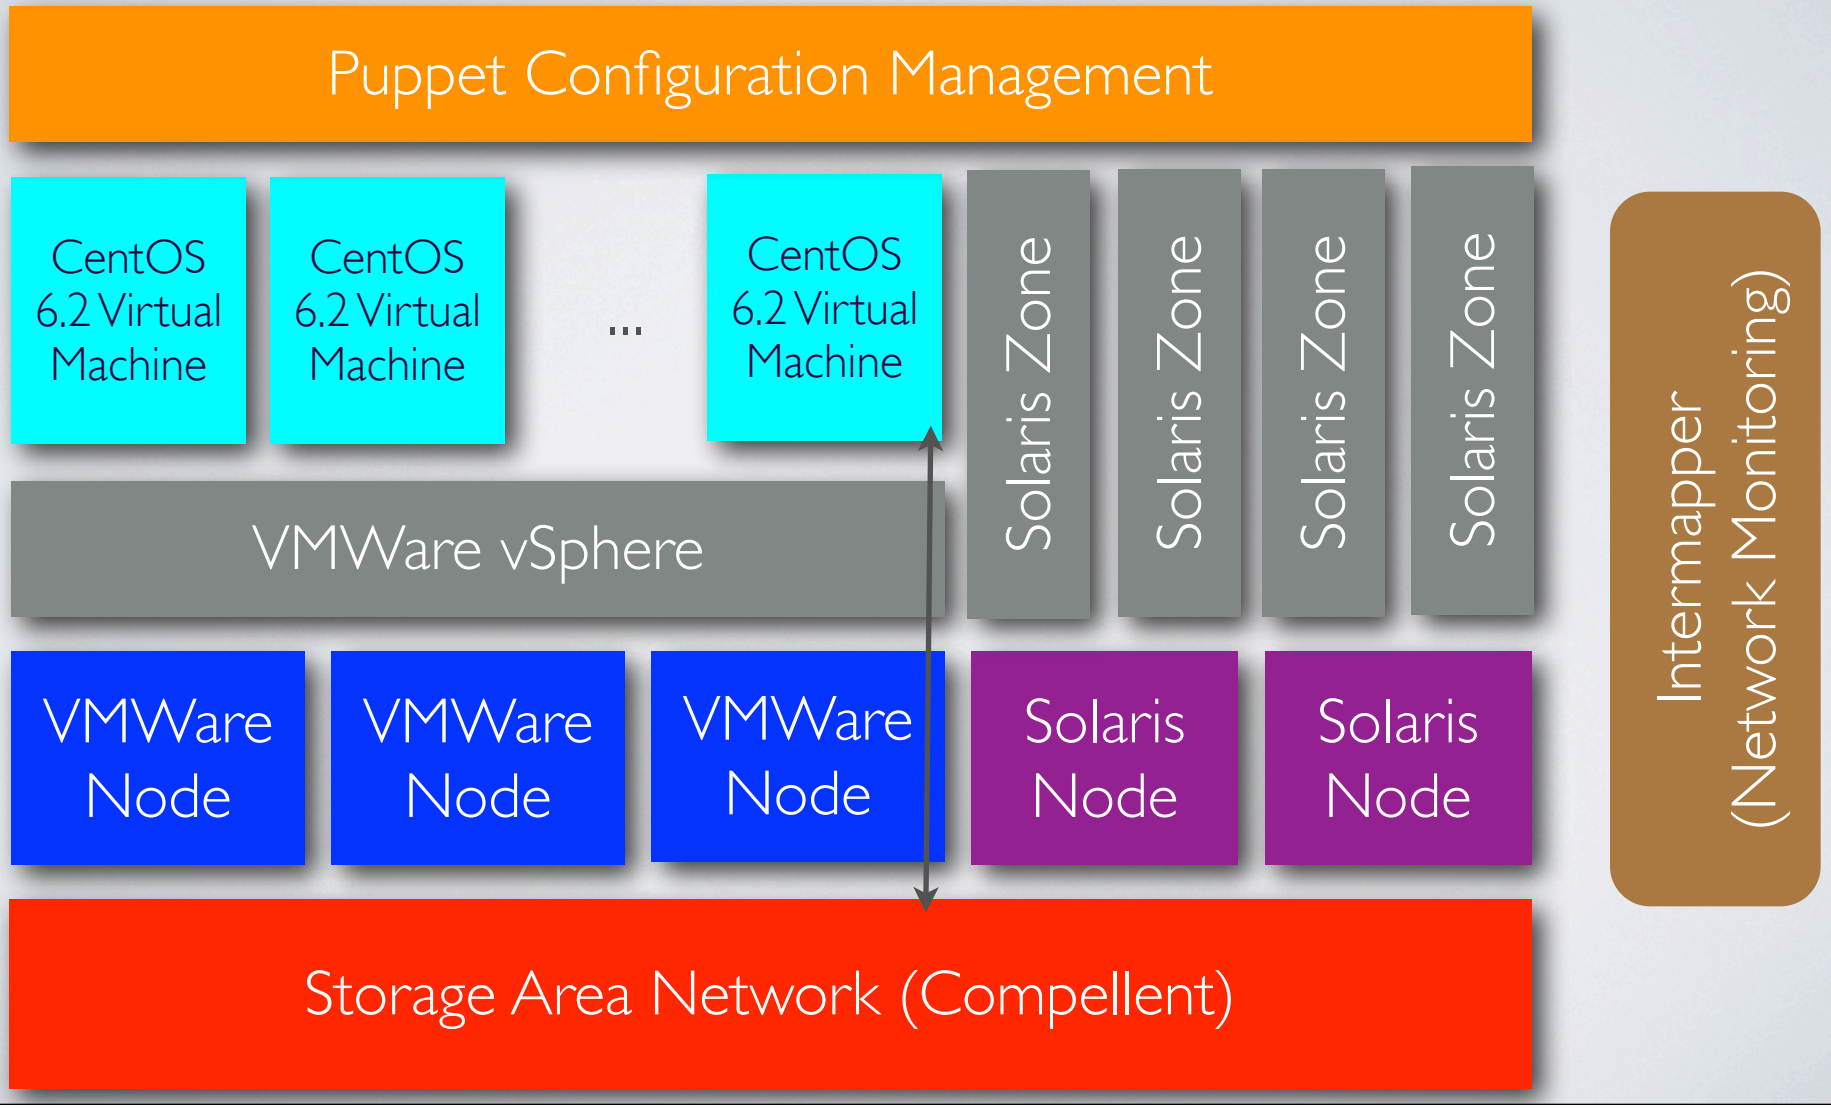

KEY ARCHITECTURE HIGHLIGHTS

- Virtualization
- SAN Storage
- Redundant Network and SAN Connections
- Configuration Management

WHY VIRTUALIZATION?

- Keep intensive processing jobs from interfering with real-time data collection
- Separation of core systems from analyst work and testing
- Easier hardware maintenance
- Easily create clones of existing systems for testing

VIRTUALIZED ANTELOPE SYSTEMS

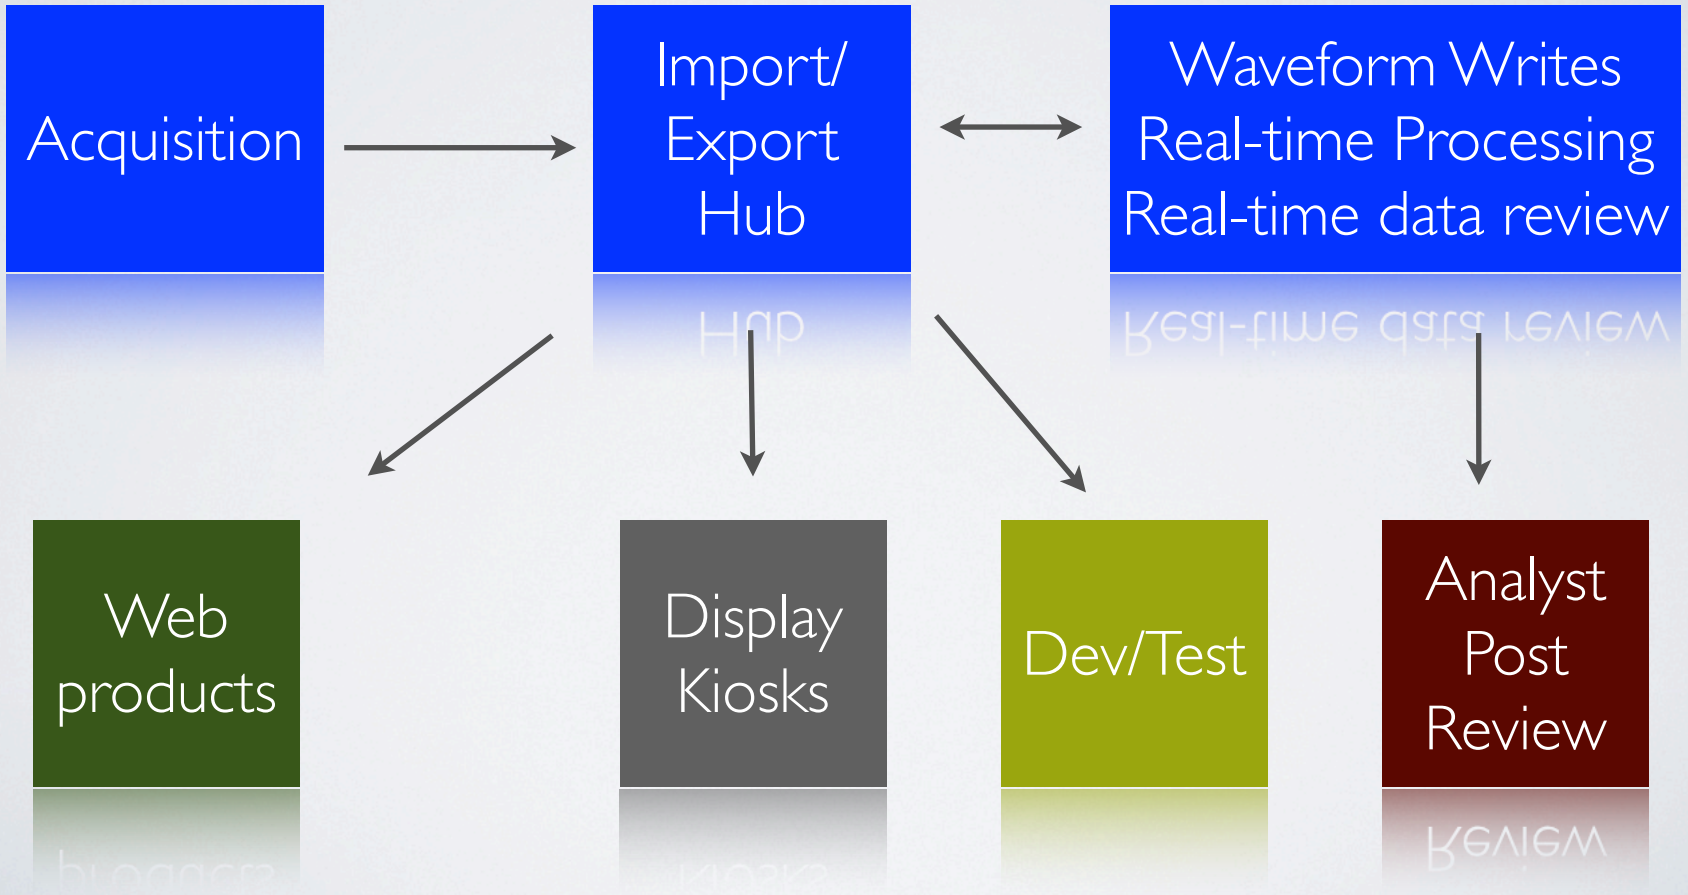

INTERMAPPER

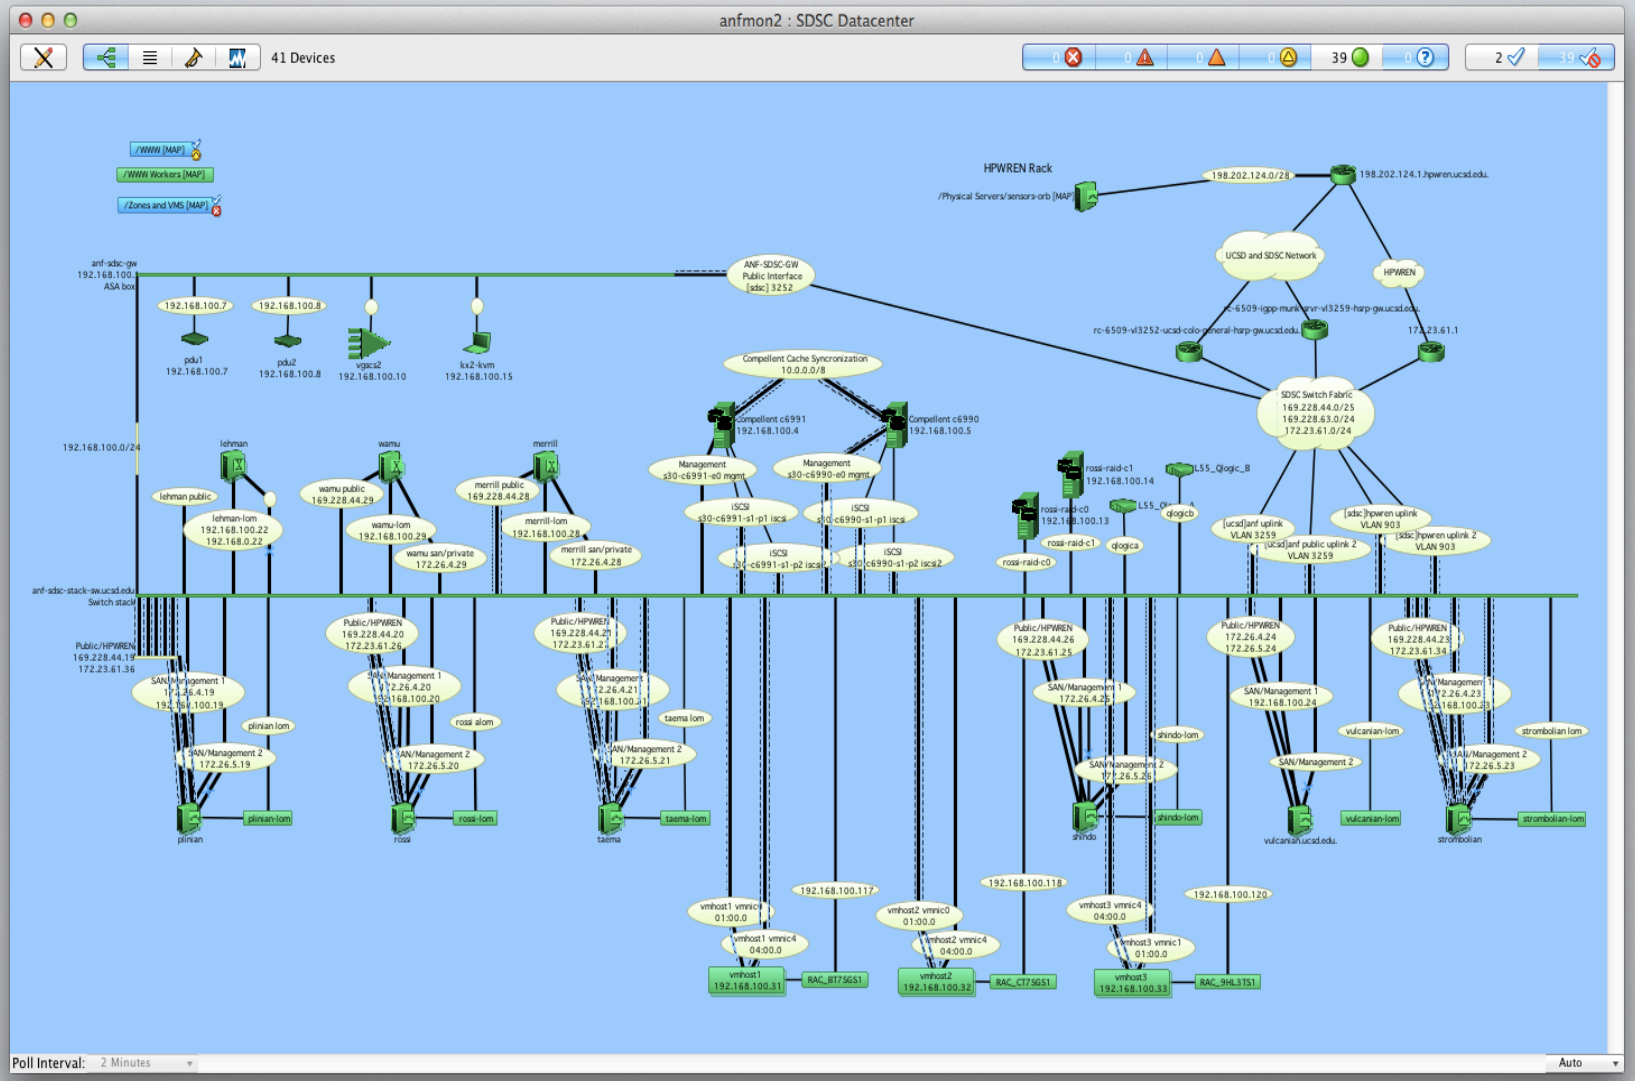

SAN STORAGE

- Compellent system at primary Datacenter
  - 110TB primary storage
  - no more manually trying to move data to fast disk or slow disk depending on workload
- All of our old disparate storage arrays pressed into service as offsite replication target for Disaster Recovery
- DR data synced via ZFS snapshots

WHY COMPELLENT?

- Antelope and our local extensions generate a lot of different storage workloads
  - Waveform data is write biased, lots of small writes
    - Once it's written, it's typically read in large sequential chunks
  - Orbservers can be read or write biased depending on their use
  - RRDs really write biased, horribly inefficient

WHY COMPELLENT?

- Existing storage couldn't keep up with write biased workload.
- We bought faster disk but it couldn't fit all of our data there
- Copying data from one volume to another is slow, adds even more workload to overtaxed storage
- We have core systems requirements, but we're always conducting experiments with data. Hard to predict what will be the most popular data for the day.

COMPELLENT MANAGEMENT

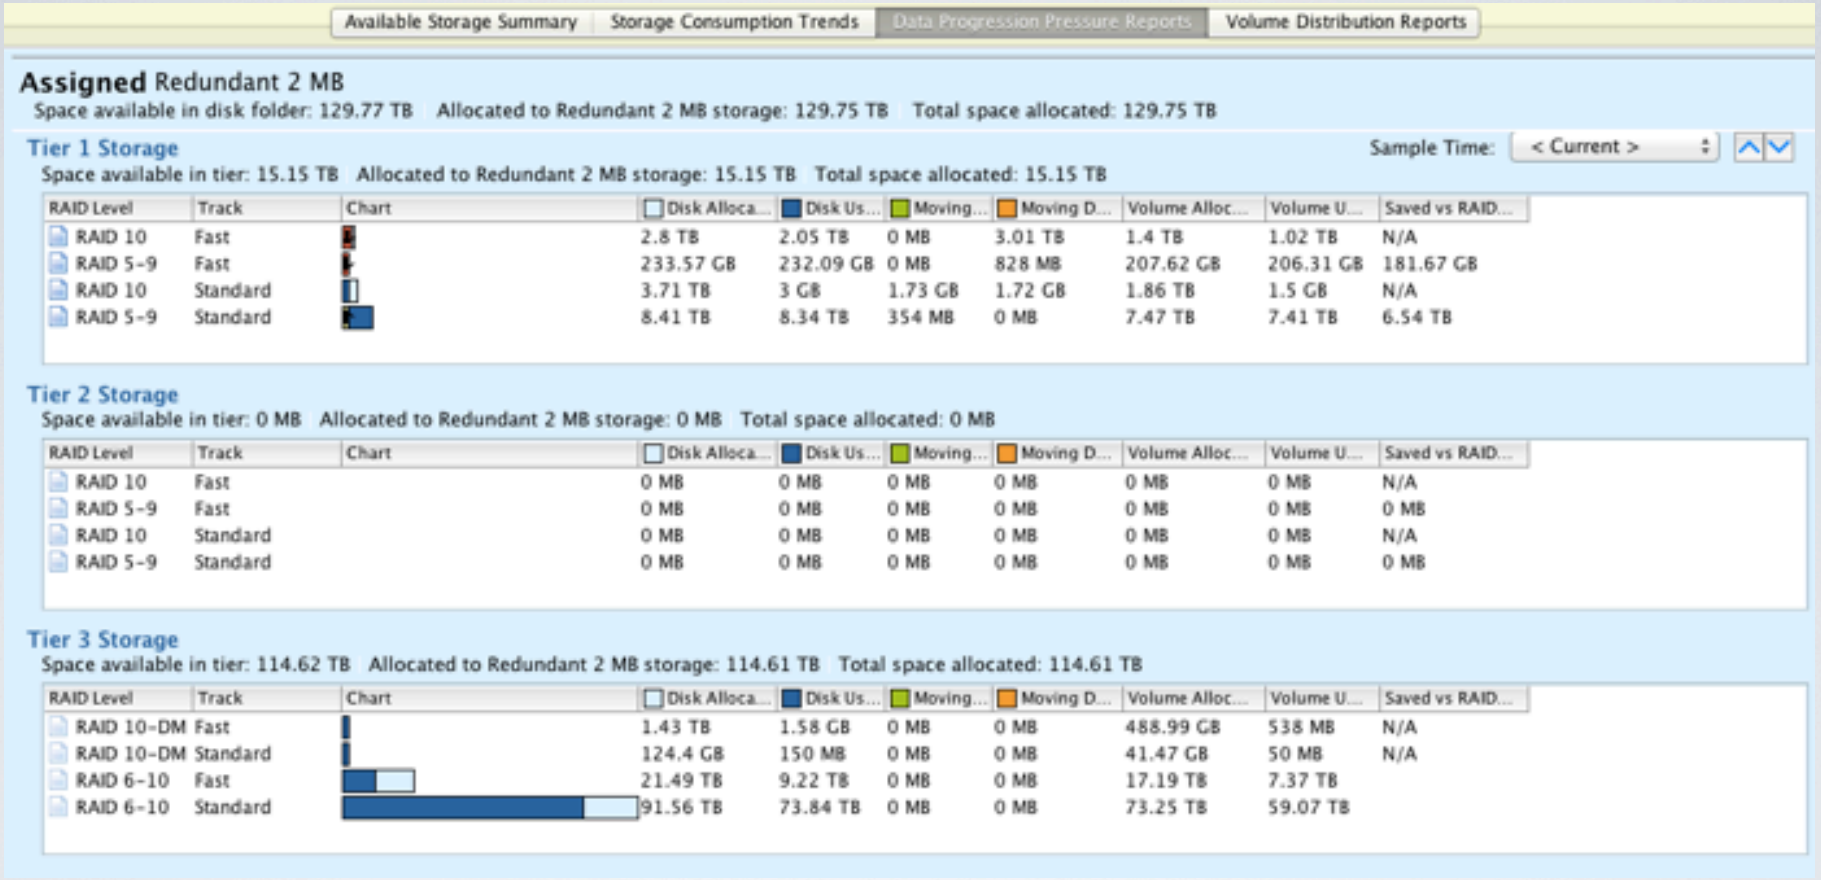

Writes go into top tier fast disk, trickle down to slower storage automatically

Least used data trickles to lowest tier of cheap slow disk.

If something suddenly becomes "hot" it will move up to higher tiers

VMWARE AND LINUX MIGRATION

- VMware cluster
  - 3 servers running at 2/3rds capacity, can survive outage of one server with no degradation of performance
- Core TA acquisition, processing, distribution migrated to Linux VMs(CentOS 6.2)
- 110 TB of Compellent storage
- Most Legacy Solaris Systems are being phased out
- TA acquisition, import/export and analyst processing all on Linux
- Web processing partially migrated

MIGRATION ADVANTAGES AND BENEFITS

- Commodity Intel hardware price/performance significantly better than Oracle (formerly Sun) SPARC, especially post Oracle
- Ongoing software support looks better long term on Linux
  - especially with regards to Antelope
  - Other open source components used in our "stack" have been harder to keep up to date on Solaris
- CPU speeds significantly faster
  - No need to differentiate workloads between floating point and integer
  - Catching up from a maintenance outage is significantly faster
  - Real-time database processing by analysts is significantly faster
    - startup time for dbloc2 went from 10 minutes down to 30 seconds

FURTHER MIGRATION ADVANTAGES AND BENEFITS

- Easier to spin up new virtual machines
- duplicate most of environment for testing
- scaling - add new processing nodes as needed, configured like existing ones
- Made easier by our use of Puppet
  - clone a fresh base system, assign role, puppet does the rest. Installs packages, starts/stops services, etc
  - Infrastructure represented in code. Software, NFS permissions, allowed users, switch configurations.

VMWARE

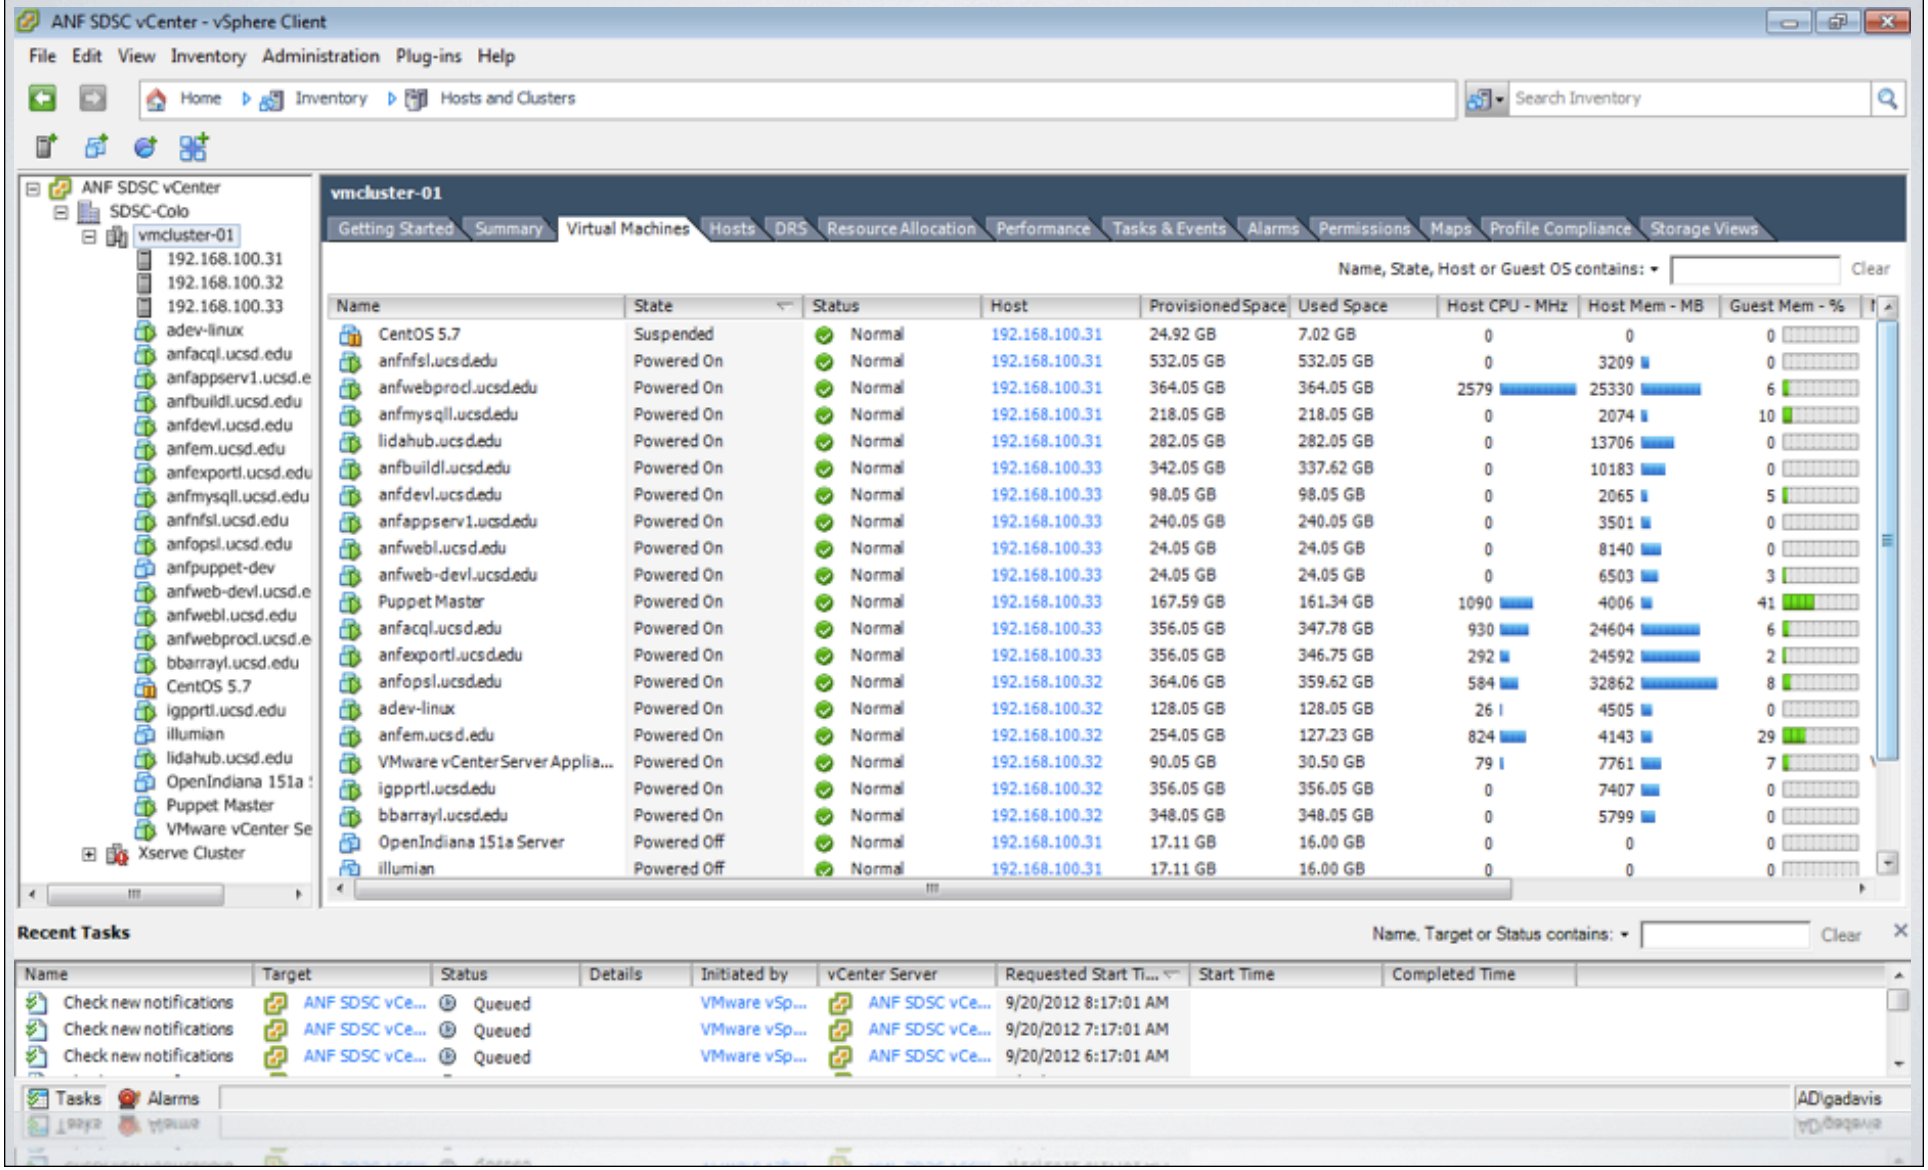
ANF SDSC vCenter - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home Inventory Hosts and Clusters
Search Inventory
ANF SDSC vCenter
SDSC-Colo
vmcluster-01
192.168.100.31
192.168.100.32
192.168.100.33
adev-linux
anfacql.ucsd.edu
anfappserv1.ucsd.e
anfbuildl.ucsd.edu
anfdevl.ucsd.edu
anfem.ucsd.edu
anfexportl.ucsd.edu
anfmysqll.ucsd.edu
anfnfsl.ucsd.edu
anfopsl.ucsd.edu
anfpuppet-dev
anfweb-devl.ucsd.e
anfwebl.ucsd.edu
anfwebprocl.ucsd.e
bbarrayl.ucsd.edu
CentOS 5.7
igpprtl.ucsd.edu
illumian
lidahub.ucsd.edu
OpenIndiana 151a
Puppet Master
VMware vCenter Se
Xserve Cluster
vmcluster-01
Getting Started Summary Virtual Machines Hosts DRS Resource Allocation Performance Tasks & Events Alarms Permissions Maps Profile Compliance Storage Views
Name, State, Host or Guest OS contains:
Clear
Name State Status Host Provisioned Space Used Space Host CPU - MHz Host Mem - MB Guest Mem - %
CentOS 5.7 Suspended Normal 192.168.100.31 24.92 GB 7.02 GB 0 0 0
anfnfsl.ucsd.edu Powered On Normal 192.168.100.31 532.05 GB 532.05 GB 0 3209 0
anfwebprocl.ucsd.edu Powered On Normal 192.168.100.31 364.05 GB 364.05 GB 2579 25330 6
anfmysqll.ucsd.edu Powered On Normal 192.168.100.31 218.05 GB 218.05 GB 0 2074 10
lidahub.ucsd.edu Powered On Normal 192.168.100.31 282.05 GB 282.05 GB 0 13706 0
anfbuildl.ucsd.edu Powered On Normal 192.168.100.33 342.05 GB 337.62 GB 0 10183 0
anfdevl.ucsd.edu Powered On Normal 192.168.100.33 98.05 GB 98.05 GB 0 2065 5
anfappserv1.ucsd.edu Powered On Normal 192.168.100.33 240.05 GB 240.05 GB 0 3501 0
anfwebl.ucsd.edu Powered On Normal 192.168.100.33 24.05 GB 24.05 GB 0 8140 0
anfweb-devl.ucsd.edu Powered On Normal 192.168.100.33 24.05 GB 24.05 GB 0 6503 3
Puppet Master Powered On Normal 192.168.100.33 167.59 GB 161.34 GB 1090 4006 41
anfacql.ucsd.edu Powered On Normal 192.168.100.33 356.05 GB 347.78 GB 930 24604 6
anfexportl.ucsd.edu Powered On Normal 192.168.100.33 356.05 GB 346.75 GB 292 24592 2
anfopsl.ucsd.edu Powered On Normal 192.168.100.32 364.06 GB 359.62 GB 584 32862 8
adev-linux Powered On Normal 192.168.100.32 128.05 GB 128.05 GB 26 4505 0
anfem.ucsd.edu Powered On Normal 192.168.100.32 254.05 GB 127.23 GB 824 4143 29
VMware vCenter Server Applia... Powered On Normal 192.168.100.32 90.05 GB 30.50 GB 79 7761 7
igpprtl.ucsd.edu Powered On Normal 192.168.100.32 356.05 GB 356.05 GB 0 7407 0
bbarrayl.ucsd.edu Powered On Normal 192.168.100.32 348.05 GB 348.05 GB 0 5799 0
OpenIndiana 151a Server Powered Off Normal 192.168.100.31 17.11 GB 16.00 GB 0 0 0
illumian Powered Off Normal 192.168.100.31 17.11 GB 16.00 GB 0 0 0
Recent Tasks
Name, Target or Status contains:
Clear
Name Target Status Details Initiated by vCenter Server Requested Start Ti... Start Time Completed Time
Check new notifications ANF SDSC vCe... Queued VMware vSp... ANF SDSC vCe... 9/20/2012 8:17:01 AM
Check new notifications ANF SDSC vCe... Queued VMware vSp... ANF SDSC vCe... 9/20/2012 7:17:01 AM
Check new notifications ANF SDSC vCe... Queued VMware vSp... ANF SDSC vCe... 9/20/2012 6:17:01 AM
Tasks Alarms
ADigadavis

VMWARE

- Three physical systems running ESXi 5.1 on bare metal
- Virtual machine that works as a supervisor, command dispatcher
- Automatically monitors load, migrates VMs to less loaded physical nodes as needed, transparent to users.
- Running at 2/3rds capacity so one server can be down for maintenance

SYSTEMS PROVISIONING

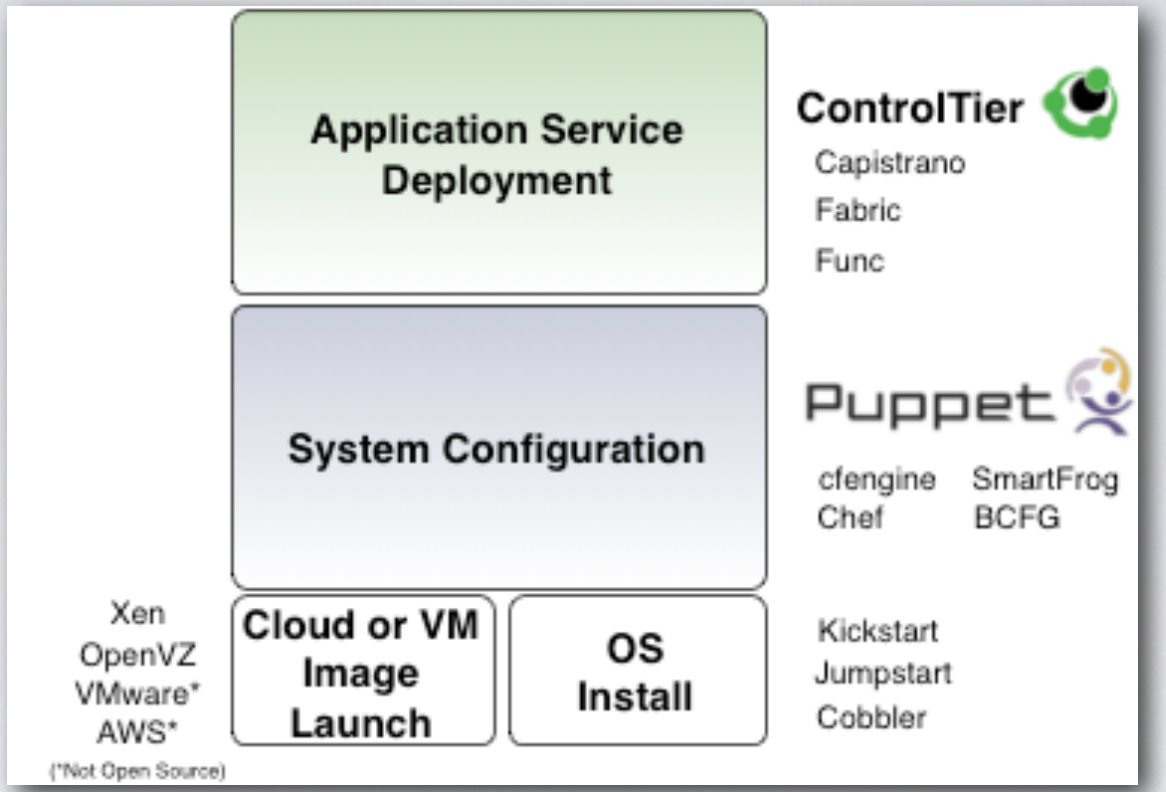

- The cool kids call it "DevOps"
- Something similar in use almost since the beginning of the project (2005), prior to that buzzword being coined
- We aren't using ControlTier but something at that tier will replace my few remaining SSH loops

CONFIGURATION MANAGEMENT WITH PUPPET

- http://puppetlabs.com
- Puppet is a configuration management tool with a declarative syntax. Similar in concept to Makefiles.
- You describe what something should look like, and the interpreter figures out the steps from current configuration to desired configuration.

INFRASTRUCTURE AS CODE. LITERALLY.

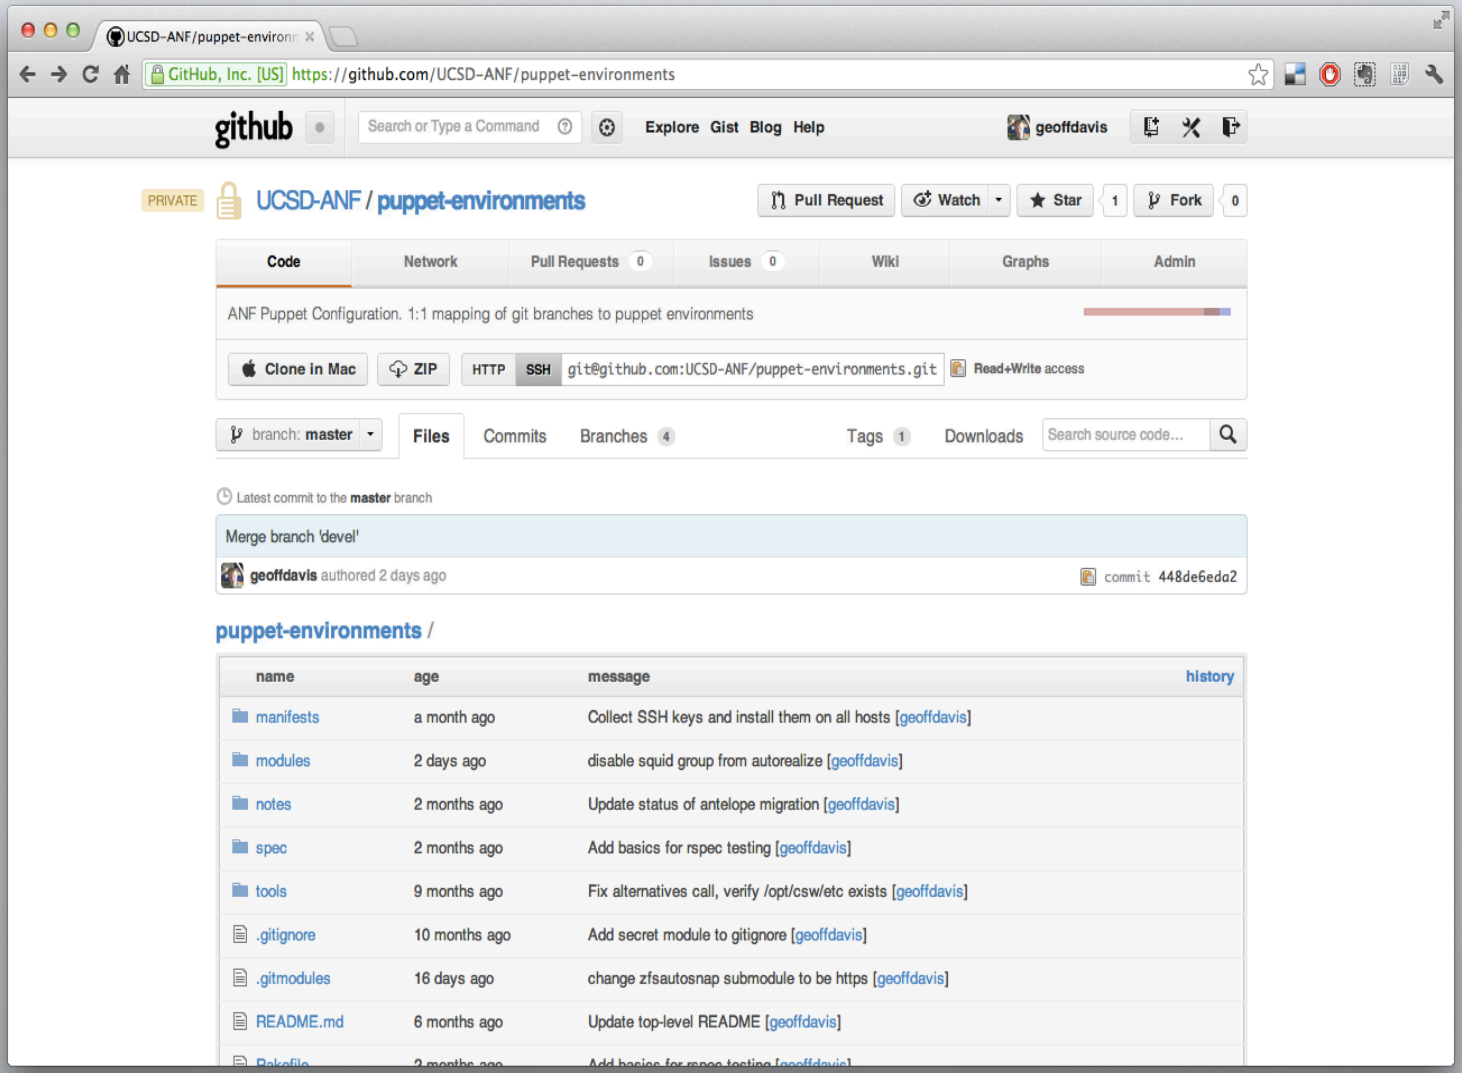

PUPPET CODE EXAMPLE

```
package { 'openssh-server':
    ensure => installed,
}

file { '/etc/ssh/sshd_config':
    source  => 'puppet:///modules/sshd/sshd_config',
    owner   => 'root',
    group   => 'root',
    mode    => '640',
    notify  => Service['sshd'], # sshd will restart whenever you edit this file.
    require => Package['openssh-server'],
}

service { 'sshd':
    ensure => running,
    enable => true,
    hasstatus => true,
    hasrestart => true,
}
```

PUPPET DASHBOARD

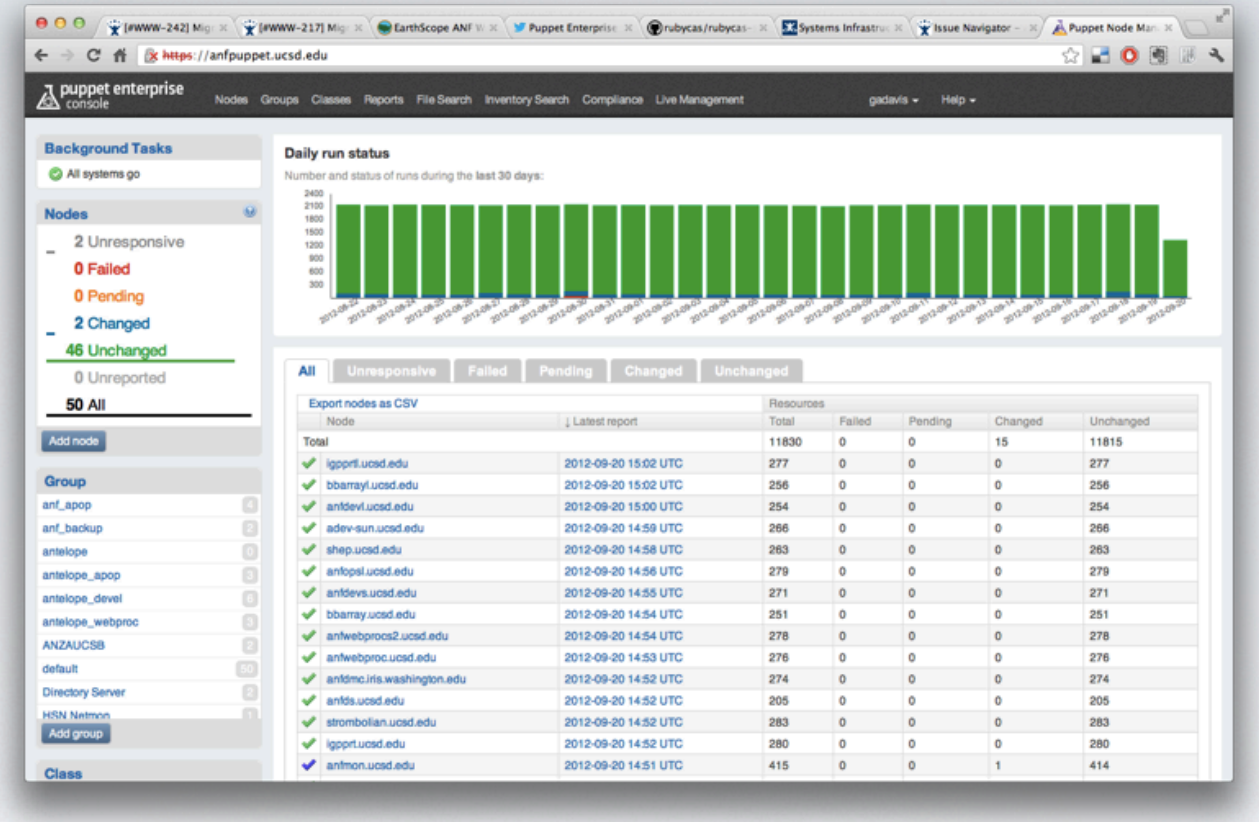

- Dashboard gives me a quick overview of changes to code.
- Can also monitor other files/services not controlled by Puppet to alert to changes.

PUPPET-ANTELOPE

- This wouldn't belong at AUG without tying it into Antelope somehow.
- https://github.com/UCSD-ANF/puppet-antelope

```
 antelope::instance { 'antelope' :
   user => 'rt',
   dirs => ['/rtsystems/foo', '/rtsystems/bar'],
}

antelope::instance { 'antelope-baz' :
   user        => 'basil',
   dirs        => '/rtsystems/baz',
   manage_fact => false, # don't create an entry in the antelope_services fact for
this instance.
}
```

ANTELOPE MANAGEMENT

- We use a lot of extensions to Antelope
  - PHP, Perl bindings, Image Magick
- Customized parameter files for ANF
- Instrument responses
- Typically running two or three releases of Antelope
- Want consistent extensions available on all of them

ANTELOPE BUILD PROCESS

- All managed by a series of Makefiles and helper scripts
  - Install Antelope
  - Copy in initial license file for the build host
  - Install $ANTELOPE overlays (localmake_config, some instrument responses)
  - Build $ANF site local tree
- Build contrib

ANTELOPE DISTRIBUTION

- Rsync based
  - We tried packages, but too much flux of contrib code, don't want to redistribute 600+ MB package for 2 MB change during incremental builds
- Deploy to development hosts first
- Stop acquisition, kick analysts off
- Deploy to production
- Deploy to backup sites

POST DISTRIBUTE TASKS

- All managed with Puppet
- System role specific parameter files (Transportable Array versus ANZA)
- Site-specific (and version specific license files)
- Install /etc/init.d/antelope with customized list of rtexec instances