


SL 6 at Nebraska's CMS Tier-2



OSG All-hands Meeting
March 19, 2012
Lincoln, NE
Carl Lundstedt






- A quick look back at Nebraska's Tier-2
- Motivation to move to SL6
- Current Status of our Tier 2 w.r.t. SL 6




Nebraska's Tier-2 at the Beginning (Red 1.0)



- 256 cores in 64 homogeneous nodes
- Single headnode running NFS, gatekeeper, NAT, login and the batch scheduler
- 14 TB in storage – single FC appliance
- 3 “support” nodes (Gums, Phedex, dcache/srm)
- 640 Mbit WAN connectivity
- 2 FTE + 1 Grad Student Brian



Nebraska's Tier-2 Today



- 3100 job slots across 278 widely varied nodes
- 4 Gatekeepers, 1 dedicated batch scheduler
- 2300 TB of Hadoop storage
- 2 NFS servers
- 3 SRM endpoints
- 12 Gridftp doors
- 18 “Support” machines
- 10+ Gbit WAN connectivity
- 2 FTE + 1 Dr. Brian



Management Tools Have Scaled

- Nagios + RSV +SAM for monitoring
- Cobbler for deployment
- Remote console on almost all our nodes
- Puppet for configuration and provisioning
- Transitioned to VMs for many services
- Moved to Hadoop



“Per Job” Control Hasn't Scaled

- No longer have time to babysit the queue(s)
- Vastly different user profile than before has separated us from our users – no more emails
- Jobs now live along side our storage on the same physical node – a bad job can no longer be allowed to kill off a node.




CGroups

- Control Groups are a kernel level process aggregation that allow for better partitioning of resources – no longer can a run away job kill off a storage node
- Attempts to run the “Oracle” kernel implementing cgroups on SL 5 failed miserably – random node crashes, lost blocks in storage, sad admins
- Cgroups in SL6 was the answer to getting automated containment of 'bad' jobs



Cgroups are Supported by Condor

- Since Condor 7.7.0 cgroup support has been included in condor.
- Each condor job is put into a dedicated group for block I/O, memory, CPU.



But...but...CMSSW...

- CMS software is not SL6 ready/certified

chroot Jail

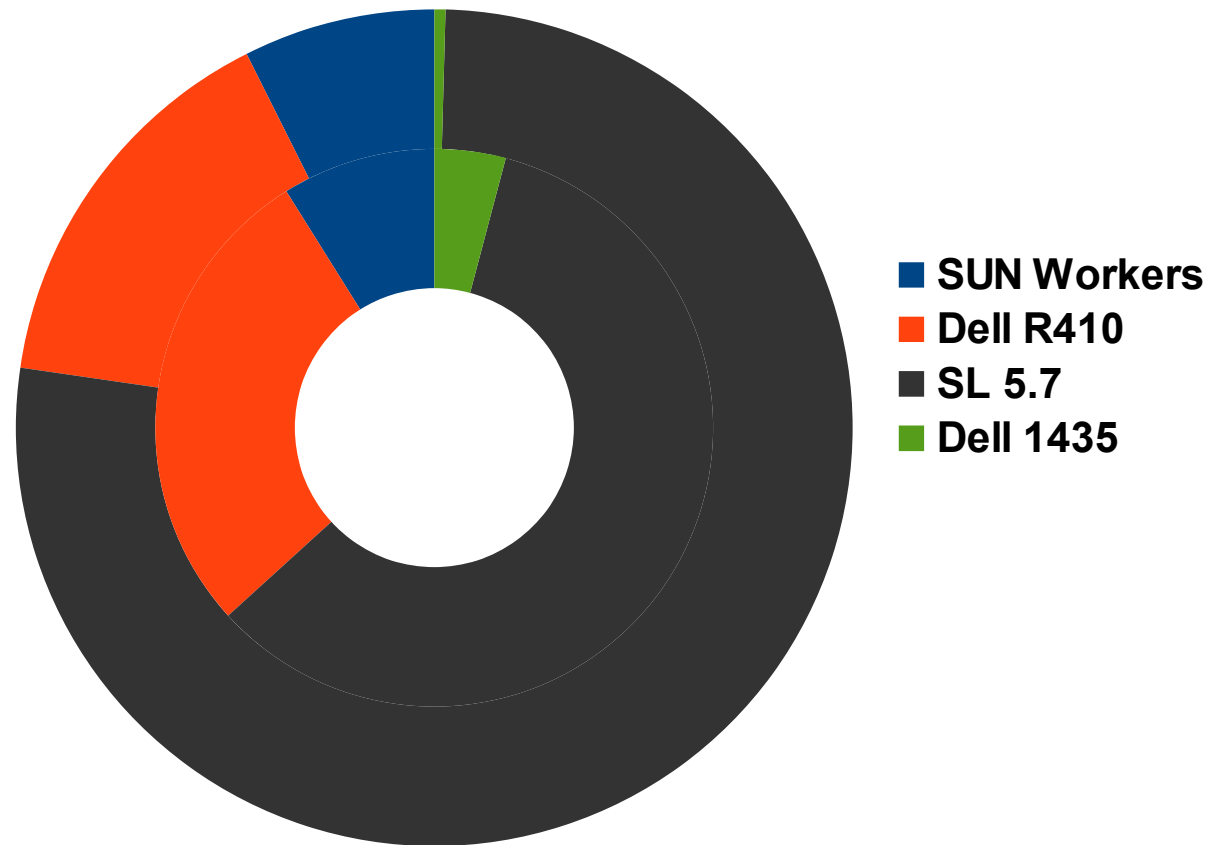

- Setup a chroot jail on the workers and confine jobs there to ensure an SL5 environment (although running with an SL6 kernel)
- /chroot/sl5-v3/root – is where we have our chroot's
- Filesystems need mounted both at the OS level and within the chroot
- Chroot configuration done via chroot-tool (by Brian Bockelman)



First Nodes to go SL6

- Our new R410s were the first to go SL6.
 - 3TB drives, GPT and SL5 didn't get along. We eventually got SL5 to work on these new nodes but only with low level formatting tricks.
- Next were Dell 1435s we inherited from Firefly
 - These nodes lacked hadoop storage, but also remote consoles. These were a natural test bed for SL6 (freebies).
- Most recent have been our 8-core SUN nodes with 3 TB of hadoop space
- Next up are our Dell 510s at 20TB and 16 slot

Transition Progress to Date



Inner Ring: Slots
Outer Ring: Storage

What remains

- Dell R710s (10 TB, 12 slot)
- Dell R510s (20 TB, 16 slot)
- Older Dell R410s (7 TB, 12 slot)
- Red's Dell 1435s (3.5 TB, 4 slot)
- SUN x4275 (21.5 TB, 12 slot)

1832 job slots (60%) , 1.8 PB of storage (80%)

Other Benefits to SL6

- Native XFS support
- Native FUSE support
- Boot support for GPT and large disks
- In the words of RedHat “More sockets, more cores, more memory”
- All our new clusters are going to SL6