



SMU ATLAS T3g

Justin Ross

ross@smu.edu

June 9, 2010

Experiences

- If you don't know the physics understanding the computing can be entertaining.
- Physics Computing is unique.
 - Math, Biology, Statistics, Engineering all run on the same cluster with Physics.
- This site is the university compute cluster but was built around Physics as the main user.

Hardware

- Worker Nodes
 - 80 Dell C6100 with 2 Nehalem E5540, 48 GB RAM, and 500 G HD
 - 27 Dell M610 with 2 Nehalem E5540, 48 GB RAM, and 250 G HD
- NFS (trouble at this size)
 - 42 TB (RAW) Nexsan SATABeast

Hardware Continued

- Lustre
 - MDS
 - Dell R710 with 24 GB RAM
 - Dell MD3000 for MDT
 - OSS
 - 4 x 84 TB (RAW) Nexsan SATABeast.
 - Each SATABeast has 2 Dell R710 with 24 GB RAM.
 - Each R710 has 4 1 Gb nics all bonded. I actually see line rate speeds, 3.9 Gb.
 - Have to use Broadcom drivers, not SL 5.4 standard drivers.
 - Small files are an issue. Biology has millions of small files. They use NFS.

Network

- Currently - Dell 8024F (24 10 Gb ports)
- 3 Dell 6048 with 4 10 Gb uplinks.
- No network issues internal to cluster.
- Administered internally to cluster.
- Each node (8 cores) gets a 1 Gb uplink.
- Campus connection to the world is a bottleneck. Campus has a packet shaper limited to 600 Mb. Working to bypass packet shaper for ATLAS data. This would give us about 1 Gb with no additional cost.

OS

- I install all machines with PXE/kickstart. Typical worker node takes 10 - 15 minutes to setup. If there is an issue with a node I reinstall it. It isn't worth debugging unless it is a recurring issue.
- Small config changes and patching are done with a parallel expect script.
- Nodes can be installed in parallel.
- SL 5.4 64bit is running on all machines.
- Any machine that is going to have heavy network usage can't use the default network drivers.

Setup

- OSG CE – SMU_HPC
 - Helpful. Allows remote use and I don't have to deal with accounts.
- OSG SE – smuosgse.physics.smu.edu
 - Data subscriptions are nice. Not much overhead for me. A few updates every once in a while.
 - Thanks to Doug, Rik and Hiro for setup help.
- GUMS

Condor


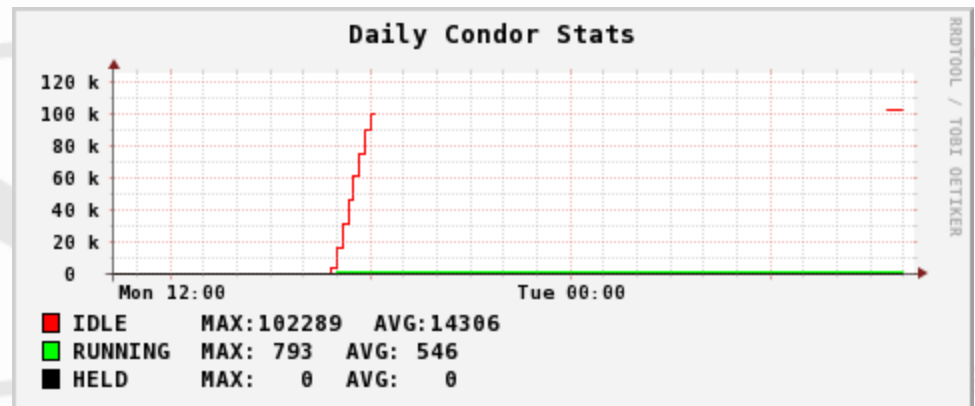
- Condor is the batch system on the cluster.
- Condor works very well for Physics/Biology. We are starting to bring in some infiniband connected nodes for parallel code in Math and Engineering. Condor doesn't seem well suited. We are going to investigate PBS. At first we will run both PBS and Condor.

CVMFS

- CVMFS provides Athena, Database releases.
- All nodes run CVMFS and point to local proxy.
- First person to use a particular piece of data or version of Athena isn't that happy. Second person doesn't see any slowdown.

System Monitoring/Assumptions

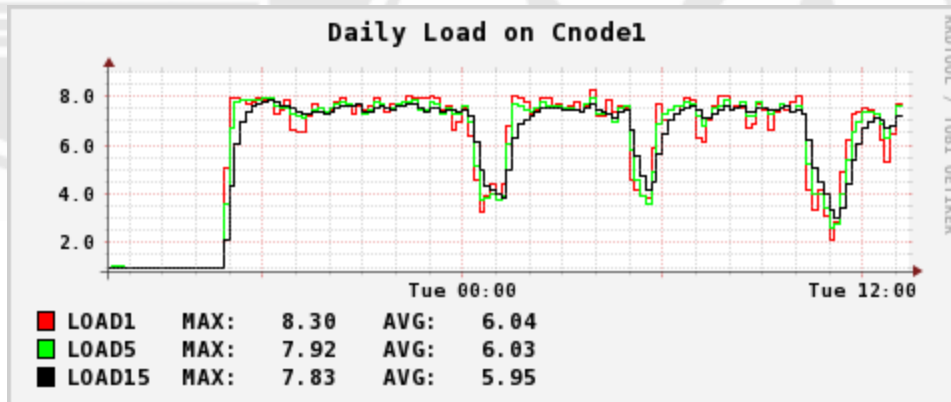
- Monitoring done with Nagios
- Usage
- Environmental
- Cluster originally had 216 cores. Never thought the queue would have over 100,000 jobs. Now with 856 cores and the ability to run 14,500 jobs per day the queue does get over 100,000 jobs. RRD max was set to 100,000.



More Monitoring

apcpdu1	PING		OK	06-08-2010 13:25:56	15d 21h 55m 19s	1/3	PING OK - Packet loss = 0%, RTA = 3.85 ms
	SNMP_PDU_LOAD_1		OK	06-08-2010 13:26:58	15d 21h 54m 4s	1/3	OK - PDU Load: 134
	SNMP_PDU_LOAD_2		OK	06-08-2010 13:28:15	15d 21h 52m 49s	1/3	OK - PDU Load: 124
	SNMP_PDU_LOAD_3		OK	06-08-2010 13:29:28	15d 21h 51m 34s	1/3	OK - PDU Load: 138

Host ↑↓	Service ↑↓	Status ↑↓	Last Check ↑↓	Duration ↑↓	Attempt ↑↓	Status Information	
cnode1	CPU		OK	06-08-2010 13:19:07	1d 8h 0m 57s	1/3	CPU STATISTICS OK : user=0.00% system=0.37% iowait=0.00% idle=50.03% nice=49.59% steal=0.00%
	DISK_BOOT		OK	06-08-2010 13:10:18	1d 7h 59m 46s	1/3	DISK OK - free space: /boot 446 MB (95% inode=99%):
	DISK_ROOT		OK	06-08-2010 13:11:33	1d 8h 8m 31s	1/3	DISK OK - free space: / 411652 MB (99% inode=99%):
	ETH1_TRAFFIC		OK	06-08-2010 13:14:38	1d 8h 5m 26s	1/3	eth1 RX: 128688011192 TX: 291392516372
	LOAD		OK	06-08-2010 13:14:07	1d 7h 35m 57s	1/3	OK - load average: 7.99, 7.82, 7.36
	MPD		OK	06-08-2010 13:14:49	1d 8h 5m 15s	1/3	PROCS OK: 1 process with args 'gird/software/mpich2/bin/mpd'
	PING		OK	06-08-2010 13:15:36	13d 9h 54m 32s	1/3	PING OK - Packet loss = 0%, RTA = 0.16 ms
	SSH		OK	06-08-2010 13:17:55	1d 8h 2m 9s	1/3	SSH OK - OpenSSH_4.3 (protocol 2.0)
	SWAP		OK	06-08-2010 13:19:06	1d 8h 0m 58s	1/3	SWAP OK - 100% free (24430 MB out of 24575 MB)



Issues

- Users. If it will cause an issue a user will figure out how to do it. This cluster has users that login and submit jobs directly to condor. This is necessary for non physics users. Pathena will help with physics users.
- NFS scaling. Only an issue with two users.
- Panda – I would like it. Going to work on fitting it into our model. Need user submitted pilots.