



# **ATLAS Great Lakes Tier-2 Site Report**

**Ben Meekhof**

**USATLAS Tier2/Tier3 Workshop**

**SLAC, Menlo Park, CA**

**November 30<sup>th</sup>, 2007**

# AGLT2 Capacity Overview



- Dedicated hardware for AGLT2 ran first jobs in April, 2007
- Cycles from mix of dual quad-core Xeons and dual dual-core Opterons
  - Currently have ~683 available job slots, each with 2GB of RAM (UM site all T3/T2)
- Aglt2.org address space includes MSU and UM sites - MSU is nearly online with their recent purchase:
  - Processing of AGLT2 jobs at MSU site ready by mid-December.
  - Panda submission will transparently distribute across both physical sites.
- Fault-tolerant disk storage capacity: 86TB online now, +160TB ready.
- 120 TB raw disk now + 54 TB soon at MSU (utilized via dCache).
- Network connectivity via dual 10Gb links over MiLR
- With MSU site online, total Si2K for AGLT2 about 2,355,132. Total storage dedicated to Tier2 nearly 500TB.

# UM Cluster Composition

## CPU Details



- **Compute nodes**
  - 79 dual quad-core 2.66Ghz Xeons in Dell PE1950, 16GB RAM
    - 48 configured with 2 x 750GB SATA disks – striped into 1.2TB RAID0 volume, used as dCache pool.
    - 51 configured with 2 x 750GB SATA disks – JBOD, each disk a dCache pool.
  - 10 dual dual-core Opteron 285, 8GB RAM, 4 250GB disks
    - Configured with 3 disks for dCache pools, remainder as system/tmp space.
    - All with Tier-3 priority, but....
      - Per MOU, Tier-2 uses half of this resource
  - Prioritized access handled by condor configuration
  - 1,412,236 possible Si2K dedicated to Tier2
  - Dual Gb NICs (1 private, 1 public)
  - New compute nodes: Good pricing - Dell 1950 dual quad-core Intel systems about \$3,600. Purchased 52 at UM, 54 at MSU.

# UM Cluster Composition

## Storage Details



- **Dedicated Tier2 Disk Storage**
  - ~120TB raw dCache + 40TB RAID50 dCache write pools.
  - 46TB served via NFS, RAID5|6 volumes.
  - 160TB RAID space ready for allocation as needed (using for tests while we can).
  - 7.5TB for muon calibration files
  - Dual NICs (1Gb private, 10Gb public)
  - New storage node: Dell PE2950 with four MD1000 shelves (see supplemental diagram). Delivers **40TB RAID50** for just under \$25,000. Purchased 6 at UM, 5 at MSU. Also purchased 3 PE2950 for ~\$7000 with no MD1000 for use as new headnodes.
- **Tier3 Storage**
  - One 40TB PE2950 RAID50 storage node, NFS exported.
  - Various smaller RAID6 volumes, 43TB.
  - 83TB total available via NFS mounts.

# New Storage Node



## AGL Tier-2 Storage Node

10 GE (Myricom PCI-e x8, CX-4 copper)

Dell 2950, Dual quad-core E5335 (2.0 GHz), 16GB (667MHz), two 250GB SATA-II drives, DRAC5, 2 Perc 5/E (PCI-e x8/x4, x2 port SAS RAID controller)



System configuration uses hardware RAID50 (two sets of two striped RAID5 arrays) for 42TB total (21TB / RAID50)

SAS x4 Cable (PCI-e x4)

Each SAS cable plugs into a Perc 5/E port (two cards, four ports)

SAS x4 Cable (PCI-e x4)

System cost \$26K



Dell MD1000, 15 x 750GB SATA-II drives  
RAID5 Configuration, 10.5TB



Dell MD1000, 15 x 750GB SATA-II drives  
RAID5 Configuration, 10.5TB

SAS x4 Cable (PCI-e x4)

SAS x4 Cable (PCI-e x4)






Dell MD1000, 15 x 750GB SATA-II drives  
RAID5 Configuration, 10.5TB




Dell MD1000, 15 x 750GB SATA-II drives  
RAID5 Configuration, 10.5TB

# UM Rack Layout ~ November 2007



  
 May  
 2007  
  
  
 Oct.  
 2007

# UM Rack Layout ~ November 2007



Oct. 2007

# UM Installation





# MSU Cluster Composition



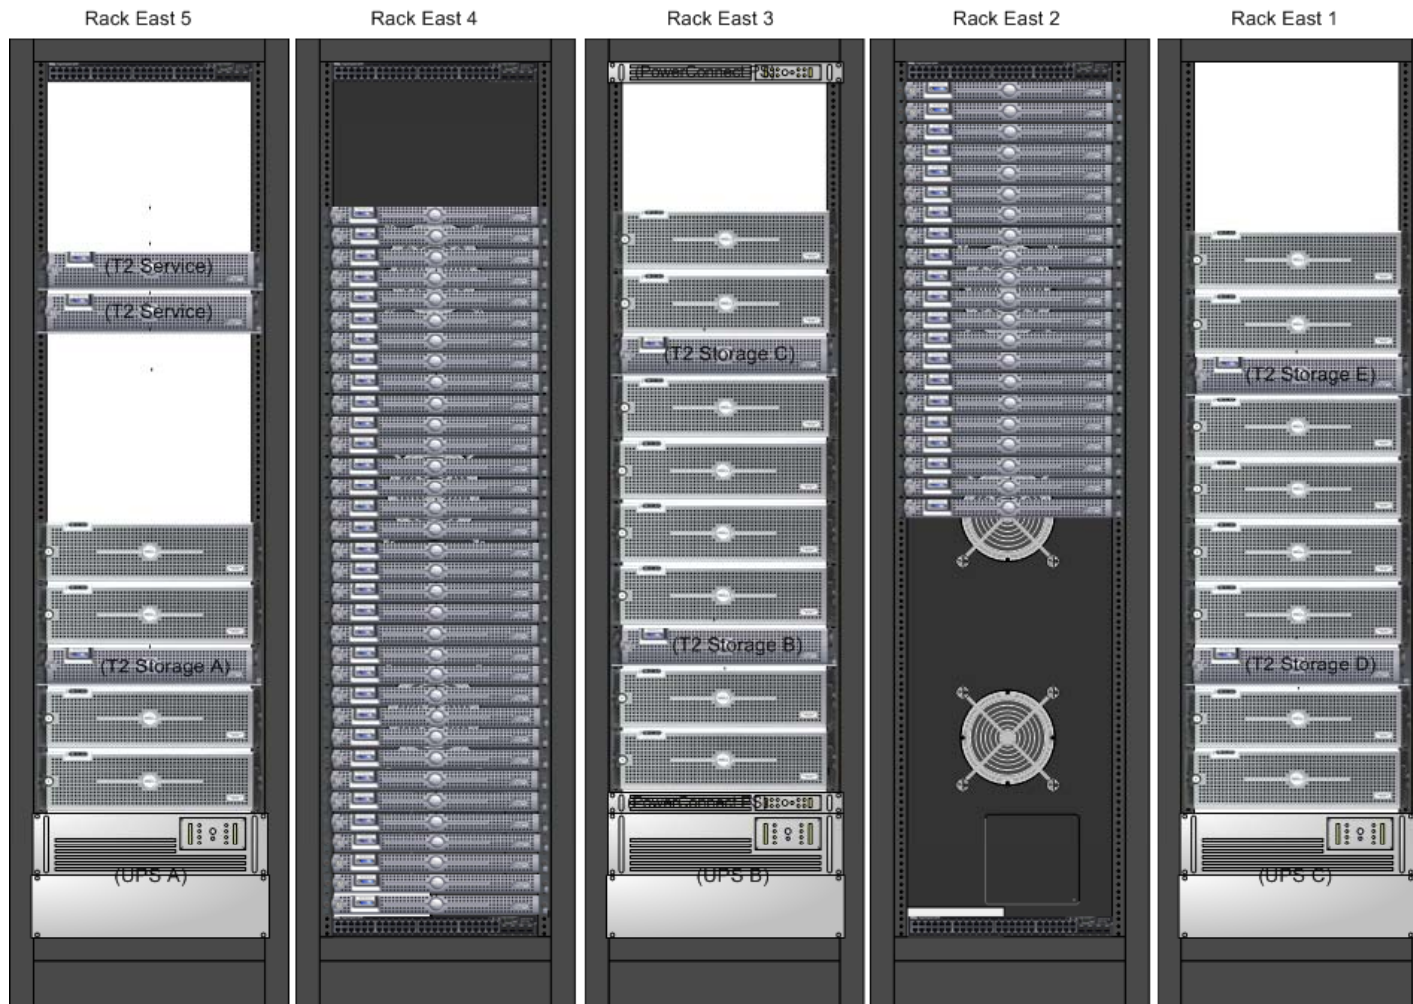
- UM and MSU submit bid request jointly, obtained same pricing.
- **Compute Nodes**
  - 54 PE1950 compute nodes 16GB RAM, Dual quad-core Xeon 5355 (2.67Ghz), 2 x 750GB SATA HD.
  - No RAID, drives allocated with 2 dCache partitions and job /tmp space.
- **Total SPECint\_2000 at MSU = 940,896 Si2k.**
- **Storage Nodes**
  - 5 Dell PE2950 storage nodes with MD1000 disk units and 225TB of raw space for allocation in AGLT2.
  - Details of allocation, dCache vs NFS, not yet determined but will match UM
- **Physical infrastructure done, 10G external network connected to world now.**  
**Compute nodes allocated (Rocks cluster tools used for allocation, as at UM).**

# MSU Site



- Power: 150 KW Liebert PDU
- Cooling: 30 Ton Liebert Airhandler
  - additional unit to be added early next year
- Safety: Vesda smoke system
- Size: 400 Sq. Feet, 20 Racks.
- Despite enormous amounts of frustration and investment of time/money, everything has come through and MSU has a very nice facility. Many congratulations to MSU people on overcoming the challenges renovating the space - we are installed, networked, and ready to turn on!

# MSU Rack Layout – November 2007




**East Row - MSU BPS Room 1221 - AGLT2 Only**

24-Sep-2007

# MSU Installation



# MSU Installation

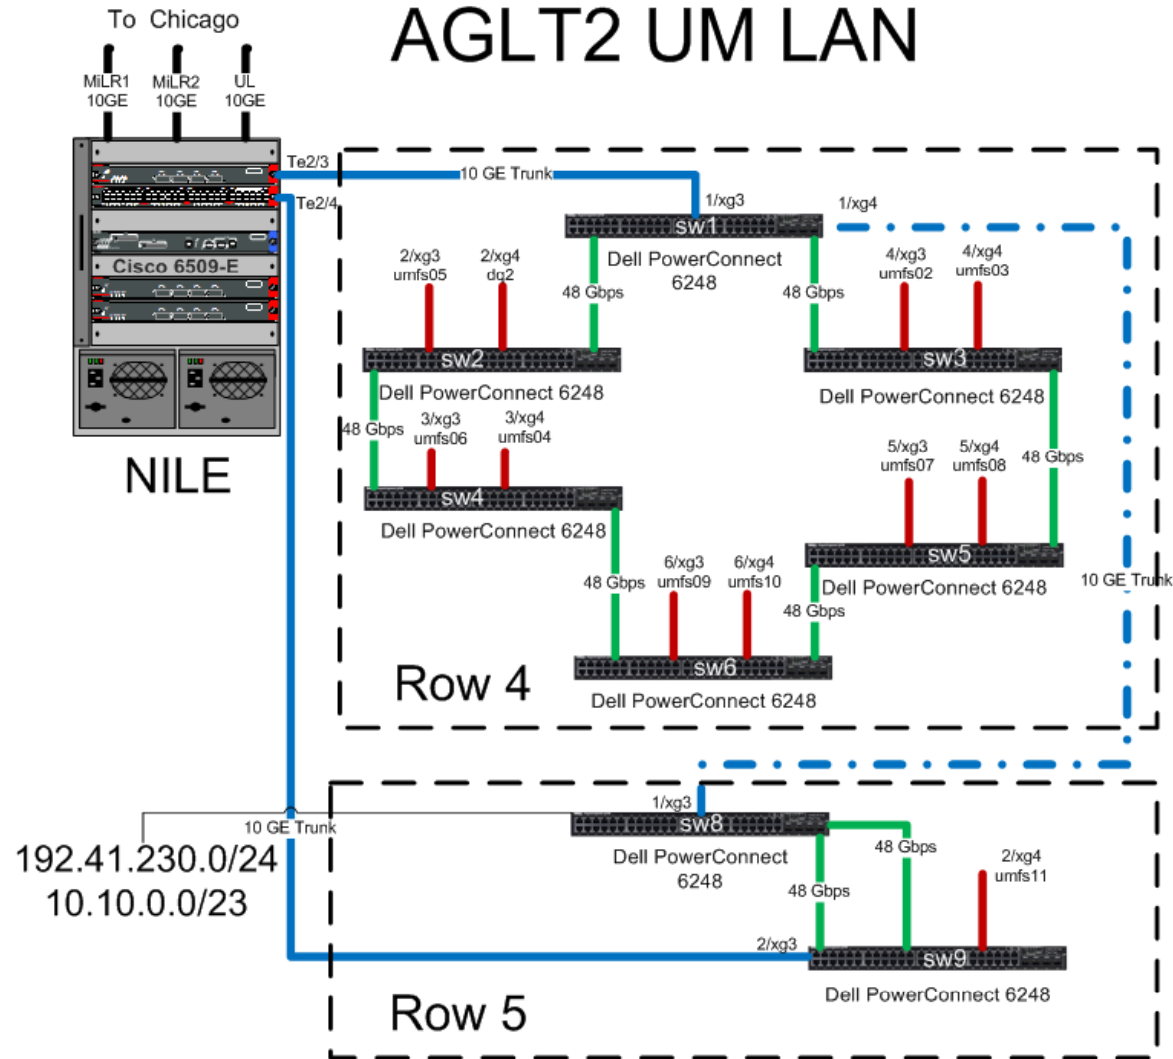


# Network Equipment



- Network switches: Dell PowerConnect 6248 - stacking modules (48Gbps) 48 copper 10/100/1000 ports and two 10 gig module slots.
- UM has 8 switches in stacks of 2 and 6 units. MSU purchased 6 units, 5 stacked and 1 spare. Pricing under \$3000 per unit, including stacking/uplink modules and CX cables.
- Reliability problems at UM site. Switch has many software bugs exacerbated by large stacks and complex VLAN arrangements. Some hardware failures (2 of 8 units at UM since April 07).
- Dell support not prepared for complex configurations, and often unreceptive and/or unwilling to acknowledge that the switch is not behaving as documented.

# Network Architecture



# Remote Site Access



- Raritan KVM for console access to non-Dell machines.
- DRAC access to Dell consoles – unfortunately IE only.
- DRAC Virtual Media – remote installations, works simply and reliably.
- Serial Console IP servers for access to switches.
- APC power infrastructure, remotely switch outlets.



# FY08 Plans for AGLT2



- September 2007 UM/MSU combined spent around \$625K on storage/compute:
  - 11 storage nodes ~ **460 TB**, 106 compute nodes for ~ **1635 kSI2K** additional CPU
- UM/MSU - tentatively planning to purchase in Spring, around \$400K-\$500K
  - If we spend the equipment money for the **whole next fiscal year** AND use some of our local leveraged funds
  - Anticipate being able to acquire similar capability to the last round of purchases
- MSU – capable of significant expansion because of augmented cooling and space beyond original plans
  - Support for up to 13 more racks of similar equipment with planned cooling expansions

# Software Access



- **pacman** kit installation to afs space
  - Mirror for Tier-3 access elsewhere
- **OSG** installed to AFS where read-only access required
  - OSG 0.6.0 with VDT 1.6.1i
  - Soft links back to local disk for configs and needed write access
  - Gatekeepers upgraded to OSG 0.8.0, VDT 1.8.1b via local installation.
- **Gatekeeper** is `gate01.aglt2.org`
  - Condor v6.9.4 collector located on cluster head node `umopt1`
  - `gate02.aglt2.org` is backup gatekeeper
- **OS** is **SLC4.4** or **4.5**
  - Kernel is customized for AGLT2 and UltraLight - 2.6.20-20UL3smp
  - 64-bit kernel with 32-bit compatibility libraries for code

# Resource Allocation



- We use Condor 6.9.4 job scheduler. Policy on Tier-2 is:
  - Highest priority for usatlas2 (admin, software install)
  - Next highest priority for usatlas1 (ATLAS production)
  - Next highest priority for usatlas3 (US ATLAS users)
  - Lowest priority for OSG (including usatlas4)
- The Tier-3 is “integrated” with the Tier-2 provisioning (ROCKS) but has different priorities reflecting 50% contribution to the Tier-2.
- If either Tier-2 or Tier-3 is idle, the other expands to fill
  - Equilibrium when all is busy at about 80% of job slots to Tier-2
- OSG/usatlas4 limited to, at most, 16 cores
- This only applies to UM – MSU is still planning Tier3 structure

# AGLT2 Analysis Queues



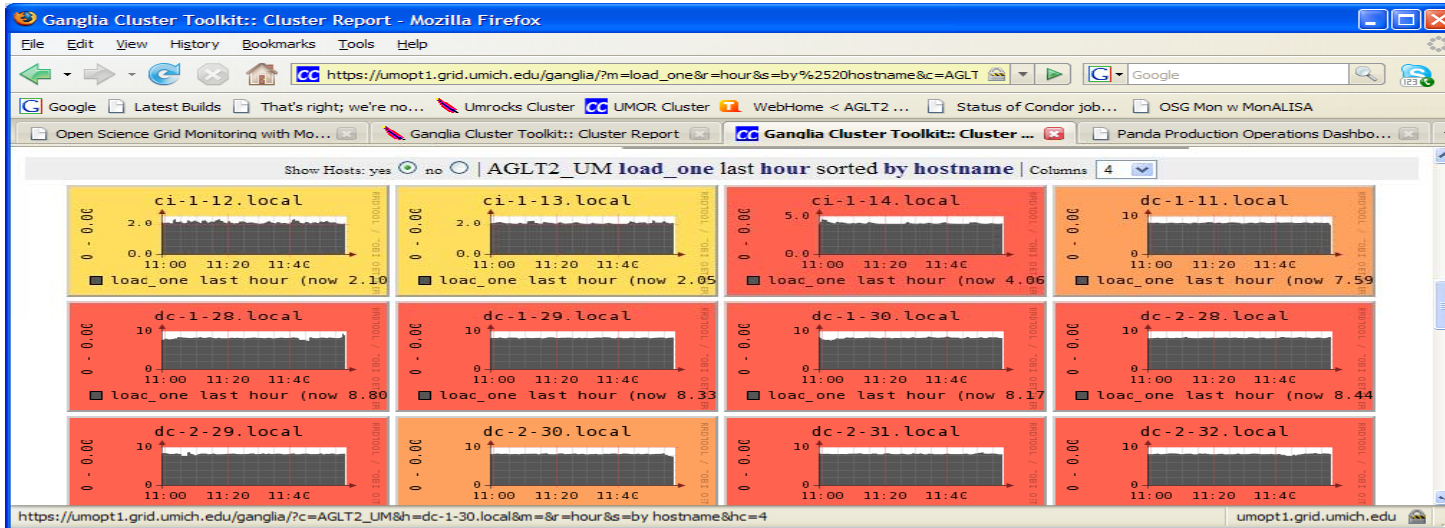
- Analysis queue is setup and tested, generally idle.
  - One dedicated 8-CPU Dell node.
  - Bob Ball is working with other T2 sites to help setup.
  - OU (Condor) and BU (PBS) volunteered first.
- All analysis queue sites are waiting for full set of AOD and NTUPLES needed for analysis jobs.

# Local Monitors



- Environmental
  - APC for power, temp, humidity
- Ganglia for overall, fast view of performance
  - Perl script for details of running processes and immediate log access
- Hardware health and performance via snmp/cacti
- OS logs to central, searchable syslog-ng server
- Email notification of most “out of range” violations
- Following is a sampling of available plots

# Ganglia and Perl Job Monitors



Status of Condor jobs for - Mozilla Firefox

http://gate01.aglt2.org/Monitoring/condor\_jobs.html

### Status of Condor jobs on AGLT2

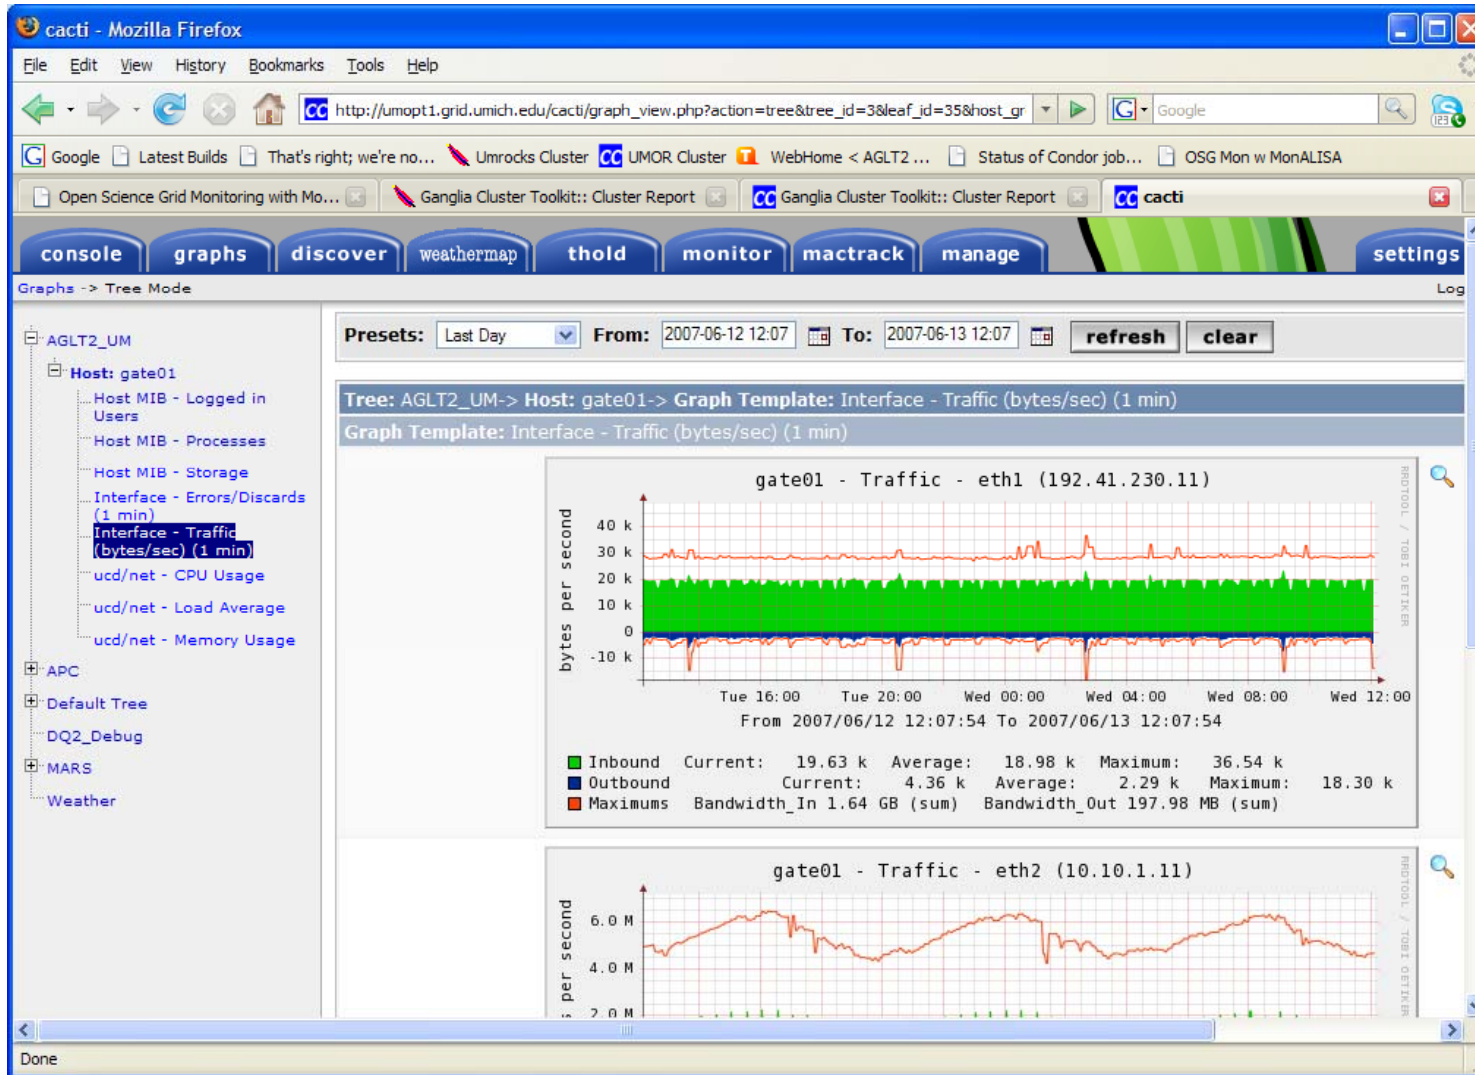
Data collected at Wed Jun 13 12:00:59 2007 (updated every 5 minutes)

We have 813 Condor jobs (239 running, 574 idle) on AGLT2\_Condor on UMOP1 on AGLT2 (Total Slots 299)  
 CondorVersion: 6.8.5 May 17 2007 for X86\_64-LINUX\_RHEL3

Condor Jobs Status

ID	Submitter	Proxy	StartDate	Host	Log	Sys	User	ImSize	Status	RunCat	NEvents	Type	CMD(l
10000.0	daits	daits	Jun 13 11:42:22	vm4 on c-2-32	Log	0	0	220000	2	1	31	Job(5)	ZoG010 2lep 0C
10001.0	daits	daits	Jun 13 11:43:38	vm7 on c-1-28	Log	0	0	220000	2	1	26	Job(5)	ZoG010 2lep 0C

# Cacti



# Syslog-ng



php-syslog-ng 2.8: REGULAR RESULTS - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://atgrid.grid.umich.edu/php-syslog-ng/index.php?table=logs&excludeHost=1&host2=&exclud

Google Latest Builds That's right; we're no... Umrocks Cluster CC UMOR Cluster WebHome < AGLT2 ... Status of Condor job... OSG Mon w MonALISA

Open Science Grid Monitoring with Mo... Ganglia Cluster Toolkit:: Cluster Report CC Ganglia Cluster Toolkit:: Cluster Report php-syslog-ng 2.8: REGULAR RES...

php-syslog-ng Wednesday June 13th, 2007 - 12:14:28  
Network Syslog Monitor Your IP: 141.211.99.16

Logout Search Config Help About

Use this link to reference this query directly: QUERY

BACK TO SEARCH  
Number of Entries Found: 162

DEBUG INFO NOTICE WARNING ERROR CRIT ALERT EMERG SEVERITY LEGEND

The SQL query: SELECT SQL\_CALC\_FOUND\_ROWS \* FROM logs WHERE priority in ('alert','emerg') ORDE

SEQ	HOST	FACILITY	DATE TIME	MESSAGE
2852446	mars03.cern.ch	auth-alert	2007-06-13 14:45:50	sshd(pam_unix)[18978]: check pass; user unknown
2852450	mars04.cern.ch	auth-alert	2007-06-13 14:45:50	sshd(pam_unix)[595]: check pass; user unknown
2904053	linat09.grid.umich.edu	daemon-alert	2007-06-13 10:08:40	2 * watchdog[2457]: shutting down the system because of error -1
2325516	linat09.grid.umich.edu	daemon-alert	2007-06-12 18:09:07	watchdog[2463]: shutting down the system because of error -1
2176737	linat05.grid.umich.edu	auth-alert	2007-06-12 14:04:12	4 * vsftpd(pam_unix)[20783]: check pass; user unknown
1749108	mars05.cern.ch	auth-alert	2007-06-12 08:22:27	sshd(pam_unix)[10683]: check pass; user unknown
1749113	mars04.cern.ch	auth-alert	2007-06-12 08:22:27	sshd(pam_unix)[4239]: check pass; user unknown
1599533	linat06.grid.umich.edu	auth-alert	2007-06-11 22:12:41	sshd(pam_unix)[29011]: check pass; user unknown
1599525	gate02.grid.umich.edu	auth-alert	2007-06-11 22:12:38	sshd(pam_unix)[12412]: check pass; user unknown
1599507	umfs01.grid.umich.edu	auth-alert	2007-06-11 22:12:34	sshd(pam_unix)[15918]: check pass; user unknown
1599488	linat11.grid.umich.edu	auth-alert	2007-06-11 22:12:32	sshd(pam_unix)[22593]: check pass; user unknown
1599494	linat08.grid.umich.edu	auth-alert	2007-06-11 22:12:32	sshd(pam_unix)[27511]: check pass; user unknown
1599463	linat05.grid.umich.edu	auth-alert	2007-06-11 22:12:31	sshd(pam_unix)[15745]: check pass; user unknown
1599317	linat06.grid.umich.edu	auth-alert	2007-06-11 22:12:11	sshd(pam_unix)[29003]: check pass; user unknown
1599285	gate02.grid.umich.edu	auth-alert	2007-06-11 22:12:10	sshd(pam_unix)[12400]: check pass; user unknown
1599277	umfs01.grid.umich.edu	auth-alert	2007-06-11 22:12:09	sshd(pam_unix)[15908]: check pass; user unknown
1599259	linat08.grid.umich.edu	auth-alert	2007-06-11 22:12:07	sshd(pam_unix)[27503]: check pass; user unknown
1599237	linat11.grid.umich.edu	auth-alert	2007-06-11 22:12:06	sshd(pam_unix)[22583]: check pass; user unknown
1599244	linat01.grid.umich.edu	auth-alert	2007-06-11 22:12:06	sshd(pam_unix)[748]: check pass; user unknown
1599197	gate02.grid.umich.edu	auth-alert	2007-06-11 22:12:04	sshd(pam_unix)[12386]: check pass; user unknown
1599206	linat05.grid.umich.edu	auth-alert	2007-06-11 22:12:04	sshd(pam_unix)[15735]: check pass; user unknown
1599166	linat06.grid.umich.edu	auth-alert	2007-06-11 22:12:03	sshd(pam_unix)[28957]: check pass; user unknown
1599175	linat07.grid.umich.edu	auth-alert	2007-06-11 22:12:03	sshd(pam_unix)[4423]: check pass; user unknown
1599181	umfs01.arid.umich.edu	auth-alert	2007-06-11 22:12:03	sshd(pam_unix)[15906]: check pass; user unknown

Done



# Performance



Analysis job summary, last 24 hours (Details: errors nodes)															
<a href="#">US Region</a>	952	25589	10-28 09:40	1	598	0	1027	1956	244	3749	17816	197	1%	0%	1%
<a href="#">CA Region</a>	816	7110	10-28 09:40	0	228	0	3862	1311	157	504	539	509	49%	18%	31%
<a href="#">FR Region</a>	145	4233	10-28 08:46	0	0	0	2529	0	75	1533	0	96	100%	0%	100%
<a href="#">UK Region</a>	989	3655	10-28 09:40	0	0	0	38	51	85	2587	432	462	52%	0%	52%
US Sites	Nodes	Jobs	Latest	defined	assigned	waiting	activated	running	holding	transferring	finished	failed	tot	trf	other
Site Name	952	25589	10-28 09:40	1	598	0	1027	1956	244	3749	17816	197	1%	0%	1%
<a href="#">AGLT2</a>	73	6721	10-28 09:39	0	<a href="#">29</a>	0	0	<a href="#">167</a>	<a href="#">5</a>	<a href="#">1029</a>	<a href="#">5469</a>	<a href="#">22</a>	0%	0%	0%
<a href="#">BNL ATLAS_1</a>	173	2331	10-28 09:40	<a href="#">1</a>	<a href="#">440</a>	0	<a href="#">9</a>	<a href="#">479</a>	<a href="#">62</a>	0	<a href="#">1281</a>	<a href="#">59</a>	4%	0%	4%
<a href="#">BNL ATLAS_2</a>	0	0		0	0	0	0	0	0	0	0	0			
<a href="#">BU ATLAS_Tier2</a>	26	715	10-28 09:39	0	0	0	<a href="#">112</a>	<a href="#">61</a>	<a href="#">1</a>	<a href="#">156</a>	<a href="#">385</a>	0	0%	0%	0%
<a href="#">BU ATLAS_Tier2o</a>	80	2448	10-28 09:40	0	<a href="#">1</a>	0	<a href="#">118</a>	<a href="#">99</a>	<a href="#">1</a>	<a href="#">316</a>	<a href="#">1912</a>	<a href="#">1</a>	0%	0%	0%
<a href="#">IU_OSG</a>	17	631	10-28 09:40	0	0	0	<a href="#">52</a>	<a href="#">64</a>	<a href="#">1</a>	<a href="#">109</a>	<a href="#">369</a>	<a href="#">36</a>	9%	0%	9%
<a href="#">LTU CCT</a>	0	0		0	0	0	0	0	0	0	0	0			
<a href="#">MWT2_IU</a>	41	2005	10-28 09:38	0	<a href="#">25</a>	0	0	<a href="#">5</a>	0	<a href="#">128</a>	<a href="#">1844</a>	<a href="#">3</a>	0%	0%	0%
<a href="#">MWT2_UC</a>	56	3470	10-28 09:40	0	<a href="#">80</a>	0	<a href="#">10</a>	<a href="#">208</a>	<a href="#">79</a>	<a href="#">625</a>	<a href="#">2407</a>	<a href="#">60</a>	2%	0%	2%
<a href="#">OU_OCHEP_SWT2</a>	0	0		0	0	0	0	0	0	0	0	0			
<a href="#">OU_OSCER_ATLAS</a>	0	0		0	0	0	0	0	0	0	0	0			
<a href="#">SLACXRD</a>	286	4659	10-28 09:40	0	<a href="#">23</a>	0	<a href="#">148</a>	<a href="#">485</a>	<a href="#">21</a>	<a href="#">743</a>	<a href="#">3234</a>	<a href="#">5</a>	0%	0%	0%
<a href="#">UC ATLAS_MWT2</a>	0	0		0	0	0	0	0	0	0	0	0			
<a href="#">UC_Teraport</a>	2	29	10-27 22:39	0	0	0	<a href="#">21</a>	0	<a href="#">8</a>	0	0	0			
<a href="#">UTA-DPCC</a>	42	1010	10-28 09:40	0	0	0	<a href="#">260</a>	<a href="#">79</a>	<a href="#">53</a>	<a href="#">132</a>	<a href="#">482</a>	<a href="#">4</a>	1%	0%	1%
<a href="#">UTA_SWT2</a>	156	1570	10-28 09:40	0	0	0	<a href="#">297</a>	<a href="#">309</a>	<a href="#">13</a>	<a href="#">511</a>	<a href="#">433</a>	<a href="#">7</a>	2%	0%	2%
<a href="#">UTD-HEP</a>	0	0		0	0	0	0	0	0	0	0	0			

# Problems



- **Non ATLAS-wide AGLT2 outages from variety of sources**
  - Gate-keeper, head node crashes under heavy load - planning to move both systems to more reliable equipment
  - Network outages – equipment issues.
  - NFS server crashes - Disruptive, but not fatal. We seem to have no or few of these anymore on our newer systems with the latest kernel and nfs utilities.
- **Software configuration**
  - Recently upgraded OSG/VDT on our gatekeeper, site operations disrupted most of a day. Upgrade was successful, and we were back online the same day.
- **Accounting issues due to broken Gratia, fixed October 2007 onwards.**

# ATLAS Memory Limits



- We have seen jobs allocating 2.5GB or more of RAM.
- Jobs like this cause nodes to use swap RAM, which can hugely slow the node down and cause loads to rise excessively
- Some sites configure with no swap – nodes crash if a job allocates more than was planned for each job slot.
- It is possible to have the scheduler (condor) kill jobs that exceed memory limits
- An ATLAS-wide policy regarding this would be a good idea.

# Summary



- AGLT2 is in good shape and will only get better.
- Accounting issues resolved as of Oct 2007.
- Dell switch problems responsible for some outages. Servers generally very reliable, though.
- OSG/VDT is up to date on the gatekeepers.
- Lots of monitors allow quick response to problems. Flexible, cross-platform/browser remote access tools are essential.
- Expansion scheduled at MSU site is very close to being finished and processing jobs. Lots of additional expansion planned.
- Review presentation by S. Mckee for more info about I/O performance tests performed at AGLT2.

# Supplementary Slides



# Performance – AGLT2



Jobs status for ATLAS VO

