



AGLT2 Site Report

Shawn McKee/University of Michigan

Bob Ball, Chip Brock, Philippe Laurens, Mike Nila

HEPiX Spring 2017 / Budapest



Site Summary

- 📄 The ATLAS Great Lake Tier-2 (AGLT2) is a distributed LHC Tier-2 for ATLAS spanning between UM/Ann Arbor and MSU/East Lansing. Roughly 50% of storage and compute at each site
 - 📄 9408 logical cores
 - 📄 MCORE slots 950 (dynamic) + 10 (static)
 - 📄 Additional 720 Tier-3 job slots usable by Tier-2
 - 📄 Average 10.21 HS06/slot
 - 📄 **6.85 Petabytes** of storage
 - 📄 Total of **96.1 kHS06**
 - 📄 Tier-2 services virtualized in VMware 5.5 (soon upgrading to 6.5)
- 📄 2x40 Gb inter-site connectivity, UM has 100G to WAN, MSU has 10G to WAN, lots of 10Gb internal ports and 20 x 40Gb ports, 32x100G/40G or 64x50G/25G ports
- 📄 High capacity storage systems have 2 x 50Gb bonded links
- 📄 40Gb link between Tier-2 and Tier-3 physical locations

Personnel Reorganization

- Over the last year we have lost two tier-2 cluster administrators, one at UM and one at MSU
- For the near-term, we are not planning on hiring a replacement
 - We have been very successful in getting good undergraduates to work on an hourly basis for us.
 - MSU added Mike Nila (Engineering-tech) 20% time
 - Looking to augment our current two undergraduate hourly students with 1-2 more to maintain “institutional” knowledge as students move on.

Hardware Additions

- ☐ Last summer we had a one-time infusion of NSF funds
 - ☐ AGLT2 received **\$398K additional funds** for hardware
- ☐ We negotiated good pricing with Dell and purchased:
 - ☐ 55 R630s; 2xE2640v4 processors, 128GB ram, 2x800 SSDs
 - ☐ 4 Storage nodes (R730xd+MD3460+MD3060e; 120x8TB)
 - ☐ 2 MD3060e shelves (each 60x8TB disk)
 - ☐ All with 5-year warranties
- ☐ Total **addition** was **24 kHS06**, 4.8 PB(raw) / **3.84 PB(r6)**
 - ☐ All equipment was added and has been operational in 2017
- ☐ We also spent remaining funds from last 5-year grant in January
 - ☐ 18 R630 nodes, 4 N2048, 1 N4032F
 - ☐ **Added 7.8 kHS06**

Software Updates Since Last Mtg

- Tier-2 VMs rebuilt to use SL7 (old SL5)
- dCache updated to 3.0.11 and Postgresql 9.5
- HTCondor now running version 8.4.11
- OSG CE updated to 3.3.21
- Various switch firmware updates applied
- Monitoring updates: OMD/check_mk to 1.2.8p18, ELK stack upgraded, custom monitoring upgraded.

Lustre at AGLT2

- We have updated our Lustre storage, using new hardware and incorporating old servers
- The new Lustre server and storage shelves were racked in the Tier-2 center last year
 - Lustre version 2.7.0 was installed and new file system created (Using ZFS)
 - Old files from /lustre/umt3 were copied to the new system
 - Old Lustre servers were then recycled to increase total storage in new file system
 - New Lustre file system is (re)mounted at the old location, /lustre/umt3
- Current AGLT2 lustre size is now 1.1 PB (retired some *old* storage)
 - We continue to have challenges with older hardware and this combination of Lustre (2.7) + ZFS (0.6.4.2)
- **Next up:** go to Lustre 2.10+ZFS v0.7 when ready

Possible Relocation at UM

- 📄 I reported last year that **AGLT2** may need to move to a new physical location at the University of Michigan
- 📄 **Reminder:** After exploring our requirements the University had the architects involved with the building renovation work around our Tier-2 center space.
- 📄 **Update:** Our college just purchased new batteries for the two APC Symmetra 80kW systems in that room.
 - 📄 We were also informed that the room will not be altered. **All good news. Hopefully we are set for at least the next 5 years.**

Future Plans

- Participating in **SC17** (simple infrastructure for 100G+)
 - Likely will demo use of **OSiRIS** object stores as part of the **ATLAS** distributed event service
- Working on integration of **OVS** (Open vSwitch) on production storage and Software Defined Networking (OpenFlow). We are experimenting with SDN in our Tier-2 and as part of LHCONE point-to-point testbed.
- Working on IPv6 dual-stack for all nodes in our Tier-2
 - Have IPv6 address block for AGLT2 (spans UM/MSU)
 - Our perfSONAR nodes are all dual-stacked and working fine.
 - **Dual-stacking our dCache system is our next step. Planned for sometime in the next month or two.**

Summary

- Tier-2 services and tools are evolving. Site continues to expand and operations are smooth.
- Monitoring stabilized, update procedures working well
- FUTURE: IPv6 for storage, SDN, SC17

Questions ?