



AGLT2 Site Report

Shawn McKee/University of Michigan
Chip Brock, Dan Hayden, Philippe Laurens, Mike Nila,
Wenjing Wu

HEPiX Fall Meeting Oct 15, 2019 / Amsterdam



AGLT2 Numbers

- 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
 - 12045 logical cores
 - by quota, Tier-3 uses 8%, Tier-2 uses 92%
 - Total of **128 kHS06**, Average **10.69** HS06/slot
 - **6.9 Petabytes** of storage (but some is very old!)
- Tier-2 services virtualized in VMware 6.7 (MSU site still 5.5)
- 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 or 2 x 100Gb bonded links
- 40Gb link between Tier-2 and Tier-3 physical locations

Personnel Updates

- Our long-time grad student at MSU, Forrest Phillips is handing off to our new grad student Carlos Buxó-Vázquez
- We have three undergraduates who started working with us at UM: Asad Siddiqui, Elizabeth Bott and Jaroslaw Kawa
- We have added a new faculty member at MSU: we welcome **Dan Hayden** who will be taking over as AGLT2 Co-PI
- Which brings up a less happy personnel update...
- Founding Co-PI **Chip Brock** / MSU is transitioning to a teaching focus, heading to sunny Arizona for Winter Semesters starting **January 2020**
 - We will miss Chip's leadership, guidance and advice
 - Thanks for almost 14 years of involvement with AGLT2
 - **We wish him all the best**

Tier-2 Hardware Additions

- Our focus for new hardware acquisition for AGLT2 for this year was **storage**
 - ATLAS and USATLAS are pushing to replace unreliable, out-of-warranty storage as quickly as possible
- We choose a Dell R740xd2 server as our replacement host
 - 96G, Dual 25G, 24x12TB NLSAS disks, 240TB usable(R6)
 - Acquiring 6 servers this round (**1440TB** usable)
 - Will allow us to replace all our storage older than 2012 (**2008:201TB, 2009:226TB, 2010:557TB, 2011:406TB**)

Tier-3 Compute Additions

- We acquired a **GPU server** to support ATLAS ML work
 - Two Gold 6136 3.0GHz 12C/24T, 96G, 2xNVIDIA V100 GPUs, Dual 25G
- Currently the Tier-3 has **728** job slots within the larger Tier-2
 - **728: In 2019 we are retiring 560 job slots and purchasing 480**
- **Buying two types of compute servers (will be part of AGLT2)**
 - Buying 4 Formulus Black servers, ~\$8K/server
 - Has Two E5-2697v4 2.3G (**18C/36T** core), **512GB** (8x 64GB 2400), **Two NVME 3.2TB** storage, 2.5 year warranty, One year software
 - Testing out the Forsa software with ATLAS jobs, evaluating impact
 - Purchase 3x6420 servers from Dell with newer processors Gold 6242 2.8G, **16C/32T, 192G** 2933 mem, 5-year warr.

Software Updates Since Last Mtg

- dCache updated to 5.2.6 & Postgresql 10.5->11.5
- HTCondor update to version 8.6.13
- OSG CE updated to 3.4.22
- UM Site VMware updated to 6.7
- Various switch firmware updates applied
- Monitoring updates: OMD/check_mk to 1.5.0p12
- Testing out Forsa software with ATLAS jobs

Lustre at AGLT2

- Lustre at AGLT2 is used for “local” users (Tier-3)
 - About 1.5 PB of space running on SL7/Lustre 2.10.4/ZFS 0.7.11 and “old” hardware. Clients are 2.10.6 / 2.12.2
 - ZFS 0.7.11 on the OSS, ldiskfs on MGS/MGT
- 20Gb/s bonded Ethernet on each OSS
- Has been working well but we are seeing increased disk failures from very old disks (2TB and 3TB disks)
- We are in the process of getting 4 new servers (240TB each) that will allow us to retire our old/problematic servers
 - One R740xd2 server has arrived, 3 more ordered

Networking Updates

As part of our hardware refresh for the Tier-3, we are adding a new network switch Dell 5032F-ON

- Will provide 32x100G ports.
- Each 100G can be alternately used as 40G, 4x10G, 2x50G or 4x25G
- Will allow us to upgrade the Tier-3 connection to the Tier-2 from current **40G to 100G** and connect Lustre at higher speeds


We also added a resilient routing configuration to our Tier-2 private subnets between UM and MSU. Now maintenance events on one part of our 'triangle' (UM-MSU-Chicago) is transparent for operations

Improving CPU Usage at AGLT2

- Convert all old and less resourceful work nodes to run multi-core jobs
 - In general, multi-core jobs require less memory and disk per core
- Create 2 logical pools to avoid defragmentation by dynamical partitioning
 - pool1: static partitioning, runs only multi-core jobs (87% cores)
 - pool2: dynamic partitioning, runs a mixture of multi and single core jobs (12% cores)
 - multi-core jobs can use both pools, but has less priority on pool2
 - single-core jobs can only use pool 2

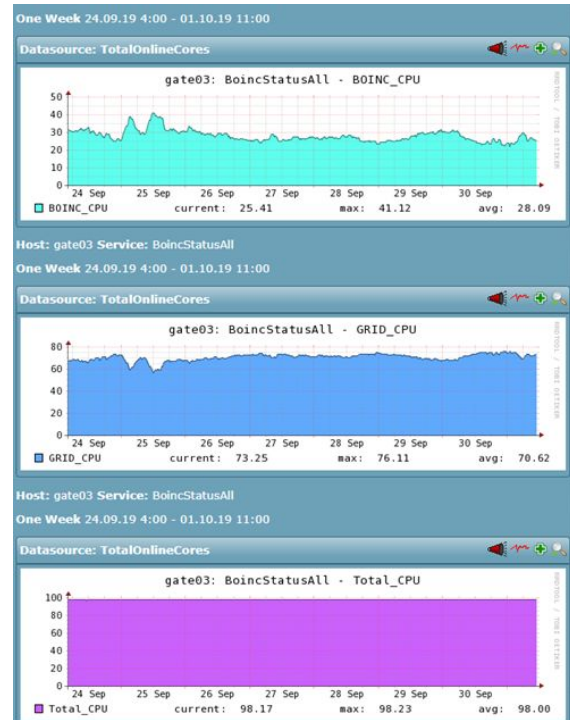
• Outcome

- The utilization(claimed_core/online_cores) is increased from 95.5 to 99.26%



BOINC Backfilling at AGLT2

- **AGLT2** started to use ATLAS@home(**BOINC**) jobs to backfill the cluster in March 2019, and it has **improved the CPU utilization** of the cluster by **over 20%**. (The average CPU utilization is 90%)
- Because the OS process schedule policy does not honor the *nice* priority for backfilling jobs, the BOINC jobs get the equal priority (cputime opportunity) as HTCondor jobs, hence the CPU Efficiency of HTCondor jobs was **reduced**.
- We started to use **cgroup** to limit the cputime usage by BOINC jobs (limit at 5% when everything else is busy), so BOINC jobs only scavenges idle CPU cycles which could not be utilized by all other jobs.



With cgroup, HTCondor jobs use 70% CPU, BOINC jobs use 28% CPU

Near Term Plans

- Participating in SC19 next month
 - Will demo use of OSiRIS caching capabilities across our 100G link
 - Demo of SLATE capabilities (containers) to deploy workflows for other demos
- Will be experimenting with Puppet and Foreman
- Completing VMware upgrade: MSU 5.5->6.7, UM on 6.7, new configuration for HA, new storage backend to prod. at MSU
- Will begin testing dCache on Ceph using AGLT2 and OSiRIS
- Testing OVS use between MWT2, KIT and AGLT2
- Have students engaged in creating AGLT2 Kibana dashboards from our central syslogging infrastructure (ELK 7.2.0)

Summary

- Monitoring stabilized, update procedures working well
- New configuration to maximize CPU use for science based upon HTCondor+BOINC+cgroups working well.
- Tier-2 services and tools are evolving. Site continues to expand and operations are smooth.
- FUTURE: SDN/NFV testing, SC19, dCache over Ceph

Questions ?