AGLT2 Site Report

Shawn Mckee/University of Michigan
Dan Hayden, Philippe Laurens, Wenjing Wu
https://indico.cern.ch/event/1222948/contributions/5316446/
March 27, HEPix Spring 2023

ATLAS Great Lakes Tier 2
UNIVERSITY OF MICHIGAN AGLT2 MICHIGAN STATE UNIVERSITY

AGLT2 Overview

- AGLT2(ATLAS Great Lake Tier-2) is a distributed LHC Tier-2 for ATLAS spanning between UM(University of Michigan)and MSU(Michigan State University).
  - What VO(s) we serve
    - ATLAS Tier2/Tier3
    - OSG (ligo, uscms, glow etc.)
  - Resource overview
    - **Now**: 225.0 kHS CPU, 12.0 PB
    - **May**: 226.5 kHS CPU, 16.9 PB
    - *Retiring lots of old equipment*
  - Resource Usage:
    - Over 98% are constantly used by ATLAS Tier-2 jobs

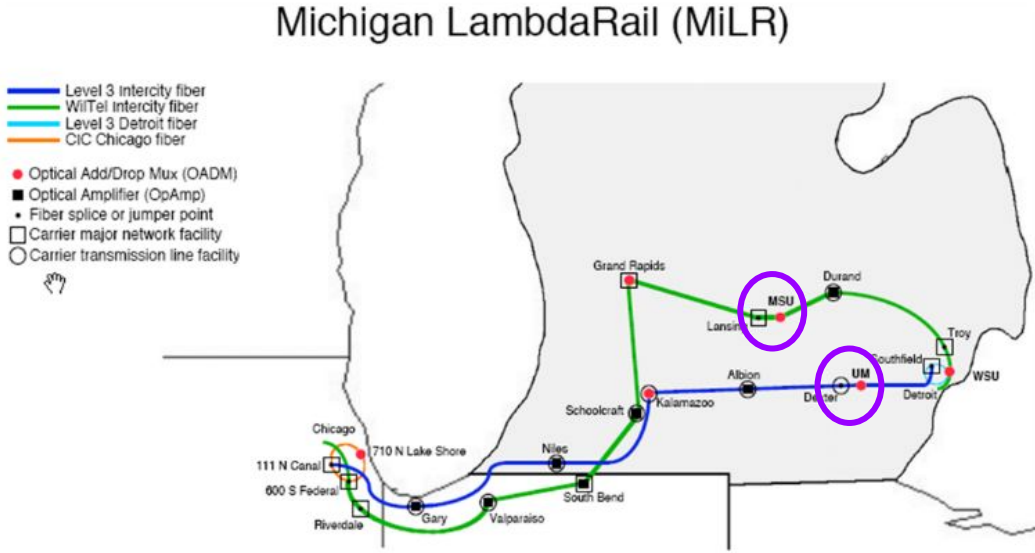

AGLT2 Site & Network Upgrades

- **Networking and site upgrades**
  - Established separate **multipath 100G MSU-UM inter-site connectivity**
  - We have completely (and very successfully) replaced our UM LAN and UM/MSU WAN network equipments starting June 2021 and finished WAN and routing changes in March 2022
  - **MSU site migration** to the MSU new data center completed in Fall 2021, which provided
    - 12x33kW racks with dual/true redundant power, multipath 100G WAN and dual/redundant data switches with 25 Gbps ports in each rack, also with room for expansion.
    - including new rack network devices, optics, cabling, configuration, at no cost to AGLT2.
  - Upgraded UM networking gear,cabling and PDU
    - Separate new data (100 Gbps/port) and mgmt switches (1 Gbps/port) for each Rack
    - Multipath 100G WAN
    - Replaced all DAC cables with **AOC**, RJ45 with slim RJ45 to reduce cable spaces in Racks.
    - Upgraded all racks with new Smart PDUs with individual socket meter and control

AGLT2 Network

- Internal ports
  - Dual 10Gb (96% work nodes), Dual 1 Gb(4% work nodes), Dual 25Gb (storage nodes), 40Gb(ELK cluster), 100Gb (VMware, perfSONAR)
- Different VLANs for management (IDRAC, PDU, ISCSI, switch OOB), private IP, public IP etc.

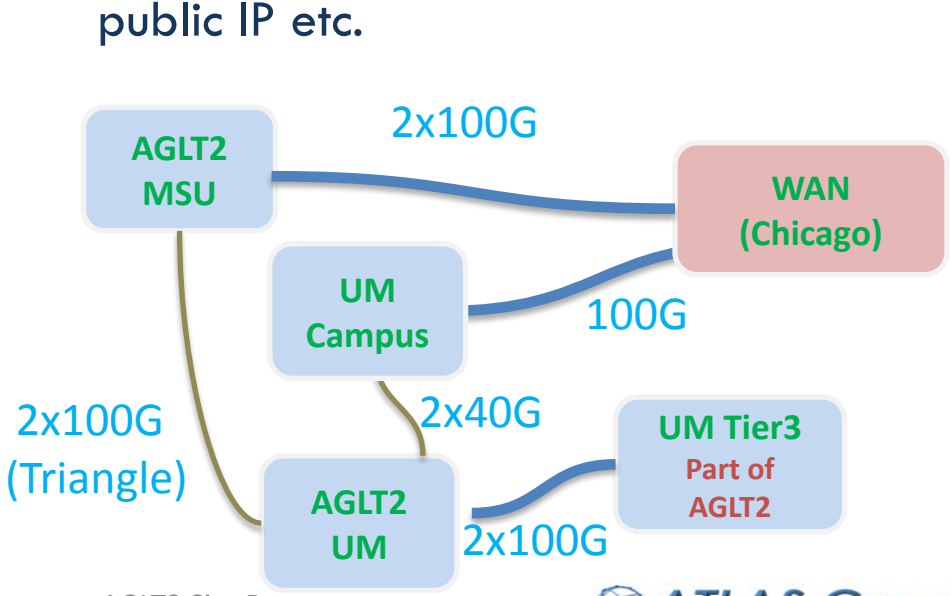

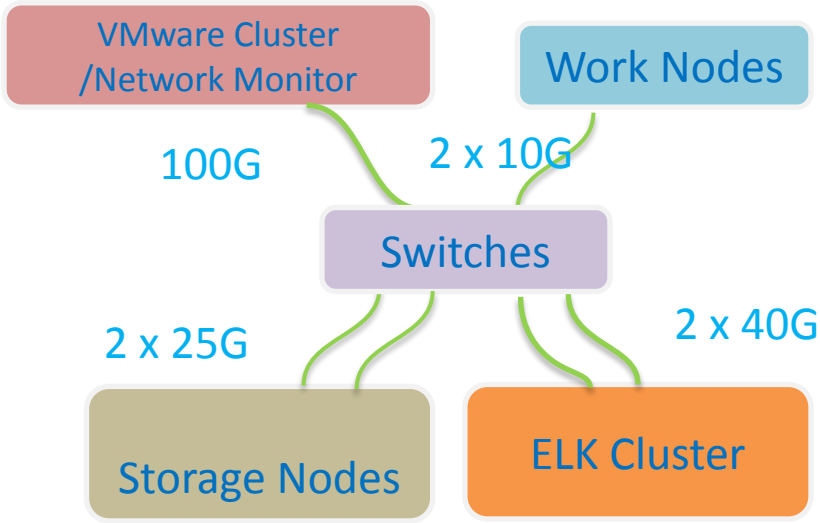

AGLT2 Storage

- **dCache (serve ATLAS Tier2)**
  - Total capacity **12PB**
  - 2 head nodes (@UM, each has postgresql database), with another 2 slave nodes for postgresql hot standby (all postgresql are on ZFS)
  - 6 door nodes (3@UM, 3@MSU)
  - 39 pool nodes (17 @MSU, 22 @UM)
  - Bonded Ethernet for pool nodes with 50Gbps.
- **Lustre (serves UM ATLAS Tier3, mounted to all the UM cluster nodes)**
  - Total capacity **1.7PB**
  - Version: Server 2.12.3, Client 2.12.5
  - ZFS 0.7.11 on the OSS, ldiskfs on MGS/MGT
  - 2x10G or 2x25G bonded Ethernet on each OSS
  - 51% storage capacity are on new hardware (R740xD2)
  - New testbed on CentOS Stream 8, 2.15.1

AGLT2 HTCondor Cluster

- 3 gatekeepers for ATLAS to increase resiliency
- 357 Physical nodes, with cores range in (24,32,40,56,64,96), RAM per core ranges in (2,3,5,6)GB
    - Generations of work nodes: Dell R620->R630->C6420->R6525
- **18064** logical cores, total of **225 kHS06**, average 12.46 HEPSCORE/core
- 2GB ~6.3 GB RAM/core, 1000 job slots for High Memory Queue (6 GB/core)
- 14GB ~52 GB Disk/core, supports Merge Queue with higher disk requirement.
- All slots dynamical partitioning, and Cluster Utilization rate over 97.9% (including downtime for hardware and software upgrade)
- Continue to run ATLAS@home/BOINC backfilling jobs (since March 2019).

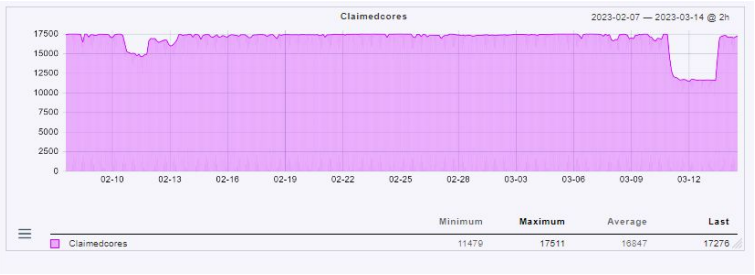

Cluster backfilling (1)

- Backfilling is to run 2 set of jobs on the cluster simultaneously, and they are controlled by 2 batch systems (HTCondor/BOINC) and different users.
- Grid jobs are controlled by HTCondor, having NICE number 0
- Backfilling jobs are from the ATLAS@home project, controlled by BOINC, having NICE number 19
- cgroup is used to control the CPU usage by each set of jobs
- Backfilling jobs are only run when the Condor jobs do not need CPU cycles.
- Backfilling jobs appear as a separate ATLAS PanDA Queue AGLT2_BOINC (jobs slots is not factorized with cores/job)
- Biggest site contributing to ATLAS@home, Scavenged CPU time from AGLT2 (2313 CPU days per day in the recent 100 days) is equivalent to a site with 3K Cores.

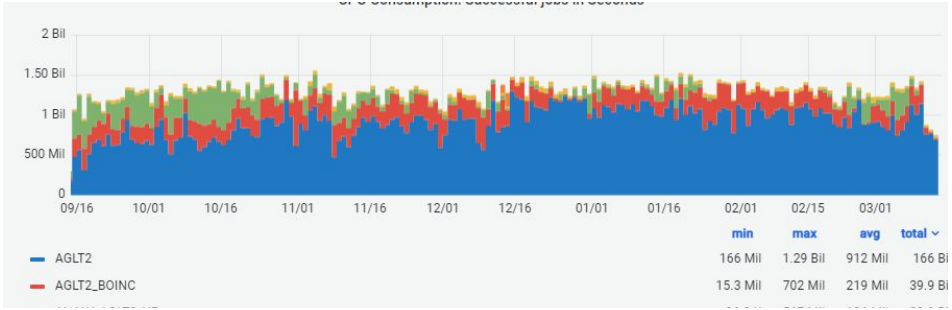

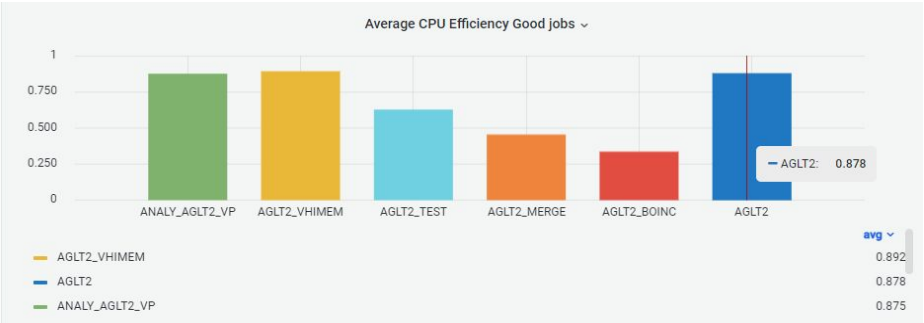

Cluster backfilling (2)

- BOINC optimization
  - configuring BOINC as a service and put it under the system.slice cgroup can reduce the CPU Efficiency loss for HTCondor jobs by 5%
  - having BOINC jobs use 50% of the cores (instead of 100%)can further reduce the CPU Efficiency loss for HTCondor jobs by another 5%.
  - After the optimization, Avg CPU Eff. HTCondor Jobs, 0.878 (vs. 0.92 in other USATLAS sites), BOINC Jobs, 0.336
- Harvest from running BOINC jobs.
  - It increases the CPU Utilization of the cluster, taking the recent 100 days for example, the CPU Utilization reaches **97%** combining both Grid and BOINC jobs (82% for HTCondor, and 15% for BOINC jobs)
  - **Fills the cluster during site downtime/cluster draining** (HTCondor update)/grid service or network issues as BOINC jobs requires only the work node itself and intermittent network access.

AGLT2 Site CPU Utilization(16000 CPU cores in 100 days)
Wed Nov 23 23:00:00 2022 to Fri Mar 3 23:00:00 2023

	cpu_eff	cpu_util	wall_util
BOINC	0.34	0.15	0.45
Grid	0.88	0.82	0.95
All	0.70	0.97	1.40

Software and Technology Details

- **AGLT2 runs a number of software packages required for an ATLAS site:**
  - OSG 3.6/HTCondor-CE 5.1.5/HTCondor 9.0.17
  - dCache 8.2.13
- **VMware cluster**
  - To host and manage critical services we rely upon VMware, which provide high availability and supports live migration of services to allow hardware, firmware and software updates
  - VMware 7, TrueNAS
- **Storage**
  - Lustre(1.7PB, 2.12.8 on CentOS 7 and 2.15.1 on CentOS 8 stream)
  - NFS(0.5PB), AFS(1.8.7) and have collaborative access to Ceph on OSiRIS (12 PB).
- **Site Monitoring**
  - We have a combination of custom built monitoring tools, along with CheckMK(2.0.0-p6), Elasticsearch(7.17), Zeek(4.2.0), Elastiflow(5.3.4) and NetDisco to provide required management and operations visibility.
- **Tape Backup**: Amanda 3.5.1 on CentOS 7 (most recent version: 3.5.2 on CentOS 7)
- **Provision**: Cobbler 2.8.5 on CentOS 7 (most recent version: 3.3.3 on CentOS Stream 8)
- **Configuration**: CFEngine 3.7.2 on CentOS 7 and 3.12.4 on CentOS 8 Stream (most recent version : 3.20.0 on CentOS 8 stream)

Software and Technology Evolution

Next generation of Operating System

- Currently our base operating system is CentOS 7.9 but we would like to migrate to either RHEL 9 or a RHEL 9 compatible OS (Rocky 9, Almalinux 9, CentOS Streams 9).
- We succeed in adding CentOS 8 Stream in cobbler to build different flavors of nodes and make it work with CFEngine to configure the nodes.
- We are experimenting migrating different software systems to CentOS 8 Stream to prepare for the transition to 9.

software	Status	Notes
lustre server	finished	2.15.1 installed on CC8 Stream
lustre client	finished	2.15.1 installed on CC8 Stream
Amanda Server	In progress	No rpm, src tarball does not build on CC8 Stream
Amanda client	In progress	No rpm, src tarball does not build on CC8 Stream
openafs server	finished	1.8.8 built from src rpm, in umatlas repository
openafs client	finished	1.8.8 built from src rpm, in umatlas repository
CFEngine Server	Not started	Need to install 3.12.4 and test existing code
CFEngine client	finished	3.12.4 in umatlas repo, working with CFEngine 3.7.2 server and the majority of the existing code are fixed to work with the 3.12.4 client
OSG software	finished	OSG 3.6 installed on CC8 Stream
HTCondor client	In progress	9.0.17 installed on CC8 Stream, problem with starting jobs
HTCondor Server	Not started	Need to install 9.0.17 on CC8 Stream
Cobber server	In progress	Install cobbler 3.x on CC8 Stream
checkmk	In progress	Client rpm is ready for CC8 Stream

Network Security

AGLT2 has been working with the WLCG SOC effort to help secure our networks while maintaining performance

Our original network had a Zeek+MISP+Elasticsearch setup for dual 40G. Cost to set up was about $2K plus repurposing an R630

Our new network is **4x100G**
We have purchased two "network capture" nodes (Del R7525) each with two Bluefield-2 NICs (each 2x100G)
Have a milestone for April 2023 to get it into production…

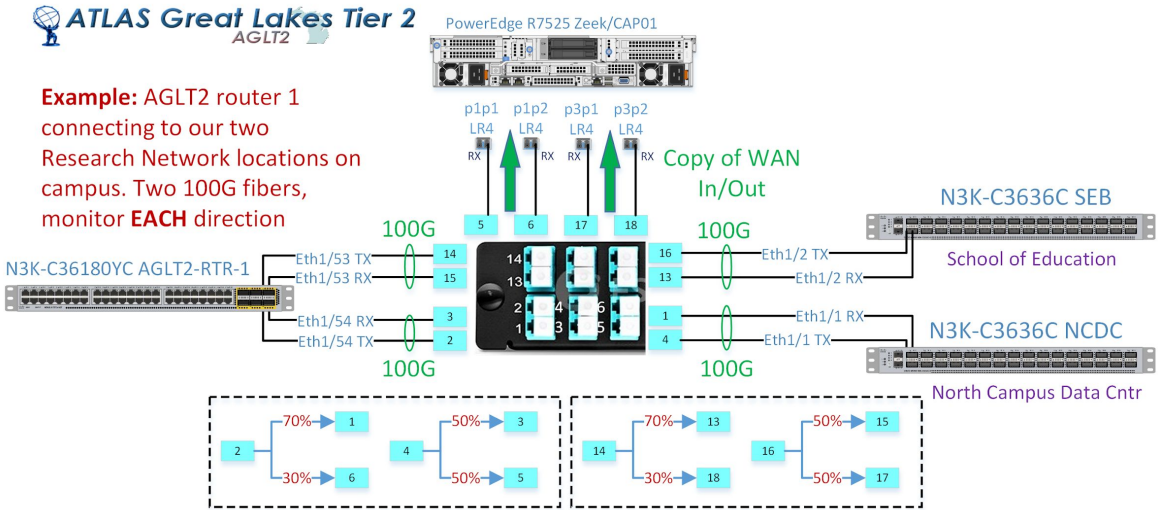

Enabling PTP (1 / 2)

For about $1500, AGLT2 added dual GPS clocks to enable PTP

- Challenge is the antenna; ideally switch support
- PTP provides < 1 microsecond time accuracy
- Makes perfSONAR latency much more powerful **BUT** needs pS mods
  To do: PTP Clients (NTP ~20 u-secs)

Enabling PTP (2 / 2)

Start of Nov, we got a our antenna installed on the roof of Physics

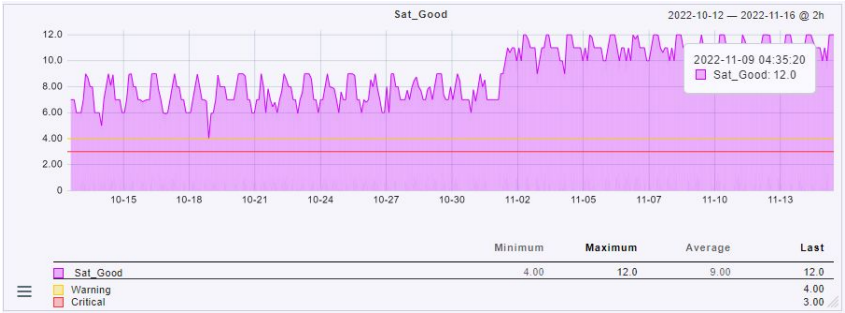

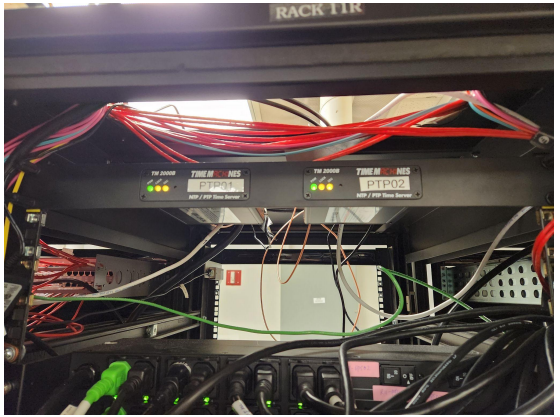

Summary

- Updates of OS, software, firmware and security patches are applied in a timely way to keep AGLT2 current
- Both Data centers had big upgrades to a next generation infrastructure.
- **FUTURE**: Enabling PTP across our systems, EL9 OS, WLCG SOC implementation and new hardware.

Questions ?