



AGLT2 Site Report

Shawn McKee/University of Michigan

Dan Hayden, Philippe Laurens, Wenjing Wu


<https://indico.cern.ch/event/1377701/timetable/#20240415>

April 15, HEPix Spring 2024




AGLT2 Overview

- AGLT2 (ATLAS Great Lake Tier-2) is an LHC Tier-2 Computing Center for ATLAS, located at our UM site (University of Michigan) and MSU site (Michigan State University).
- What VO(s) we serve
 - **ATLAS** Tier2/Tier3
 - **OSG** (ligo, uscms, glow etc.)
- Resource overview
 - **17.5K cores**/226 kHS CPU,
16 PB dCache Storage
 - **2024**: Plan to add **2.4 PB**, **40 kHS**
- Resource Usage:
 - Over 92% are constantly used by ATLAS Tier-2 jobs
 - the rest is shared with UM ATLAS Tier3 and other VOs
- Resilient **100G** paths between the 2 sites and to Chicago (ESnet)




AGLT2 DC24



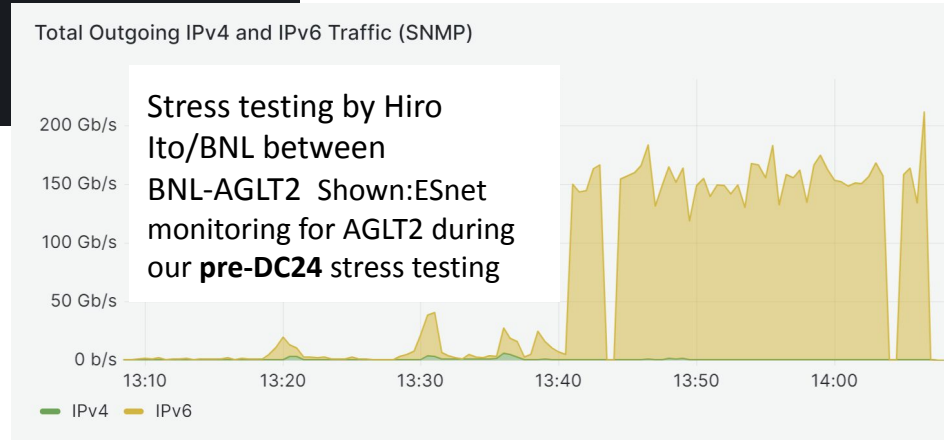
<= DC24 AGLT2 site network use

Shown are the site IN/OUT values for MSU and UM during DC24

SUMMARY: No issues for AGLT2, site was well below real limits



<= Pretesting identified (forgotten) 80G bottleneck at UM



RAID 6 VS. ZFS (1)

We want to understand if we can replace hardware RAID-6 with ZFS, removing a RAID card issues as a single point of failure.

- **Configuration**

- 2 servers with PERC H730P, 2 RAID6 over 24 disks
- 2 servers with Dell HBA330 Mini, 2 ZFS raidz2 over 24 disks, tuned with zfs recordsize=1MB compression=off atime=off xattr=sa ashift=12
- zfs-2.1.13-1(EL9)

- **We reported benchmark results with single thread IO in HEPix fall 2023**

- **work after HEPix fall 2023 (Thanks for all the inputs from HEPix fall 2023)**

- benchmark results with single thread were in favor of ZFS, but we want to test it in a production like senario:
 - With disks almost full (88% ~ 97% usage for our Tier2 storage nodes)
 - With more parallel IOs
 - Always use compression (otherwise, lose 11% usable space)
 - More reads than writes

RAID 6 VS. ZFS (2)

- **Test method**

- We dd 170 ATLAS files (size ranges from several kB to GB, total size 70GB) between Ramdisk (5.2GB/s, no bottleneck) and ZFS or RAID,
- Measure the **time** for each task (dd all the data mentioned above) Table shows ZFS vs HW RAID

Parallel IOs	zfs no compression vs. raid				zfs lz4 vs. raid			
	2	10	20	40	2	10	20	40
Disk Empty / Write Time	0.39%	-9.71%	-13.42%	-17.70%	-9.02%	-19.43%	-22.78%	-25.53%
Disk Empty / Read Time	0.95%	15.45%	40.46%	59.00%	20.38%	20.92%	44.67%	68.68%
Disk 88% full / Write Time	8.91%	1.04%	-5.90%	-11.16%	8.53%	-6.25%	-13.67%	-19.37%
Disk 88% full / Read Time	6.86%	44.55%	78.17%	87.27%	52.45%	51.85%	90.35%	135.46%
Disk 97% full / Write Time	8.08%	-8.85%	-9.17%	-7.03%	26.54%	2.18%	12.35%	11.93%
Disk 97% full / Read Time	17.10%	50.57%	76.71%	86.95%	61.66%	53.46%	94.60%	136.99%

- **Conclusion**

- when the used disk space increases, zfs read/write performance degrades significantly, RAID performance stays stable
- when number of parallel IOs increases, zfs write performance increases, read decreases.
- **In a production similar environment (97% disk full, >40 IOs, more Read than Write), RAID is a much better option.**

Transition to RHEL9

EOL for RHEL7 and clones in June 2024 means we need a new OS

- **Choice of OS after CentOS 7-> RHEL9.3**
 - Both **MSU** and **UM** have licenses for RHEL through the University and satellite servers are hosted by IT
 - RHEL9+ gives a modern kernel, compilers, and improvements for data transfers and long lifetime
 - Skipping RHEL8 and clones for longer term solution
- **Challenges from transitioning from cobbler to the UM Satellite Server (version 6.12 on EL8)**
 - The AGLT2 network is not routed to access the UM College Satellite server
 - The UM College Satellite server is not set up support PXE booting
- **Solution**
 - We deployed a capsule server (~Foreman SmartProxy) as a proxy between the AGLT2 network and the UM Satellite server
 - Settled on using UEFI HTTPboot as the bootstrap mechanism (lots of little bugs and quirks to deal with)
 - Requires BIOS change to UEFI and enabling HTTPboot on the correct NIC but works well
- **Lots to do to improve the build and leverage the various Satellite features**
 - **Ansible** or **Puppet** are potential options for us, if we choose to move off **CFengine**
 - Lots of security, auditing, monitoring and version control tools to explore in Satellite

Progress on EL9 transition

Goal is to migrate most of servers to RHEL 9 by the end of June 2024.

- **Satellite server**

- Enabled UEFI boot/dhcp server on the capture server
- Enabled lACP fallback, so either interface of a bonded link can be used for provision
- Improves the kickstart/snippets to support different flavor of nodes
- We built different type of nodes for testing purpose, **no RHEL 9 in production yet**

- **Ansible**

- Server: Ansible 7.7 on RHEL9
- Uses the ansible pull architecture, subversion as repository
- Working on rewriting all configurations in Ansible (we use cfengine for EL7)
 - have finished playbooks for common tasks (install and configure network, ssh,cvmfs, autofs, afs, chrony, user/group info, customized package installation)
 - built and configured HTCondor (condor 23.06)work node, it works with condor (10.0.7) head node in EL7.

Network Security

AGLT2 has been working with the WLCG SOC effort to help secure our networks while maintaining performance and participated in the WLCG SOC Hackathon, March 2024.

All components are now installed at AGLT2, **2 capture nodes at UM**, **one at MSU**, **MISP** configured and updated and **Opensearch** running (6 data nodes, 3 master/dashboards)

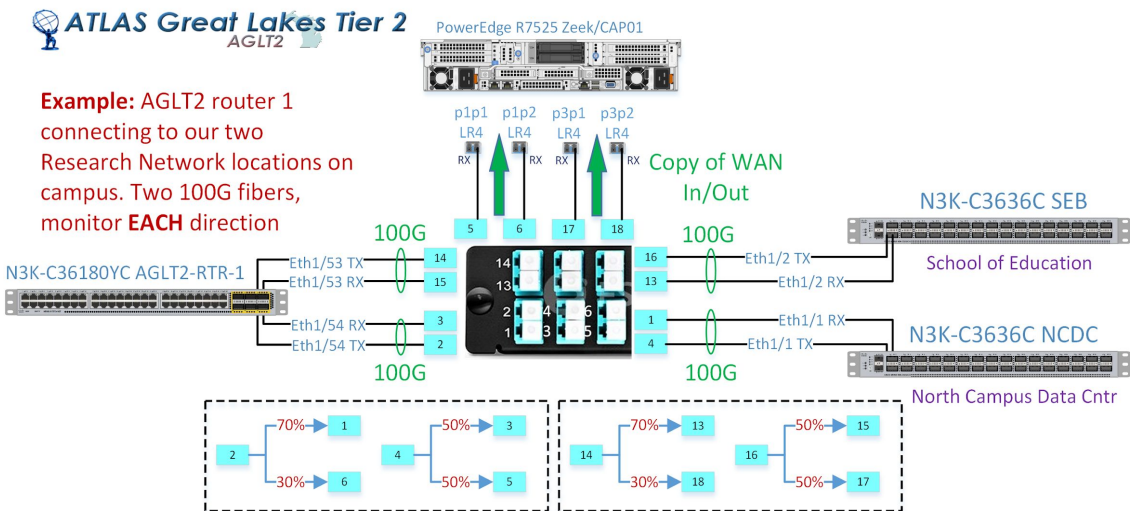
Our new UM network needs 2 capture nodes, while MSU needs 1

- Dell R7525, each with two Bluefield-2 NICs (2x100G)
- **Installed Zeek on RHEL9.3**

ElasticSearch:

- Migration to OpenSearch on RHEL9.3 completed.

Next: better integration/config with MISP and associated alerting.



CAP01/RTR-1 diagram (similar for CAP02/RTR-2)

Challenges with Bluefield-2 NICs

Part of getting our WLCG Security Operations Center fully functional requires reliable traffic capture. We choose the NVidia (Mellanox) **Bluefield-2 NICs** which combine ARM cores with a ConnectX-6.

There were a few problems when (re)installing the Mellanox OFED on RHEL 9.3. The FW updates seemed to work but didn't actually update things. Two (out of three) of our systems had no NIC assigned after reboot. We needed to use the `mlxfwmanager` to force the updates. All cards but one got the newest firmware update in place after one or two reboots; one card shows the newest FW installed but the older FW "running", even after power down and restart hours later :(

```
mlxfwmanager -d /dev/mst/mt41686_pciconf0 --online --update --force
```

For **Zeek** capture, we end up with only ONE NIC reliably capturing traffic, even though all NICs show "Link: yes" and are in promiscuous mode. **We continue to work on this...**

Summary

- Updates of OS, software, firmware and security patches are applied in a timely way to keep AGLT2 updated
- We are focused on getting our existing services and infrastructure operating on RHEL 9.x by mid-summer
- **FUTURE:** Migrate all site systems to **RHEL 9**, enable **WLCG SOC** for effective alerting and integration with other WLCG sites.

Questions ?

Bluefield-2 Card Running vs Current

```
[root@cap01 ~]# mlxfwmanager -d /dev/mst/mt41686_pciconf1 --query --online
Querying Mellanox devices firmware ...
```

Device #1:

```
Device Type:      BlueField2
Part Number:     MBF2M516A-CENO_Ax_Bx
Description:     BlueField-2 E-Series DPU 100GbE Dual-Port QSFP56; PCIe Gen4 x16; Crypto Disabled; 16GB
on-board DDR; 1GbE OOB management; FHHL
PSID:           MT_0000000560
PCI Device Name: /dev/mst/mt41686_pciconf1
Base GUID:      b8cef60300a15aae
Base MAC:       b8cef6a15aae
Versions:
  Current      Available
  FW          24.40.1000  24.40.1000
  FW (Running) 24.38.1002  N/A
  PXE          3.7.0201      3.7.0300
  UEFI         14.31.0020  14.33.0010
  UEFI Virtio blk 22.4.0010     N/A
  UEFI Virtio net 21.4.0010     N/A

Status:         Up to date
```