

# US ATLAS SWT2 SITE REPORT

HORST SEVERINI, CHRIS WALKER, PATRICK MCGUIGAN, MARK SOSEBEE

HEPIX MEETING FALL 2019

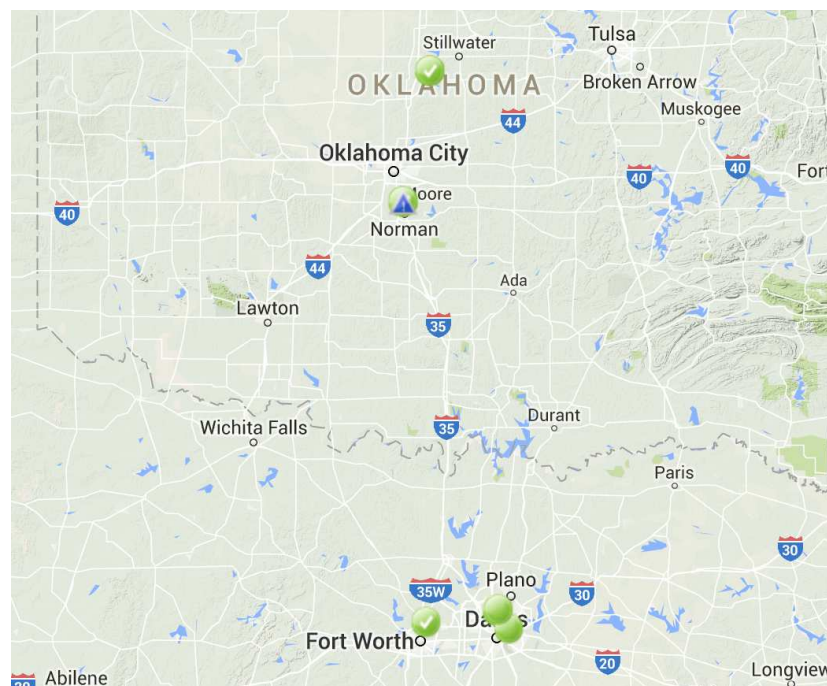
OCTOBER 2019

## Outline

- Introduction
- Computing and Storage Hardware
- Network
- Ongoing Work and Plans
- Summary and Conclusions

## US ATLAS SWT2 Center

- University of Oklahoma
  - Oklahoma Center for High Energy Physics (OCHEP) – just retired
  - OU Supercomputing Center for Education and Research (OSCER)
- University of Texas Arlington
  - Chemistry and Physics Building (CPB)
  - Arlington Regional Data Center (ARDC) in Fort Worth
- Langston University
  - Lucille (being upgraded)



## OU\_OCHEP\_SWT2 Tier 2 Hardware

---

- 72 Nodes (844 Slots) – 2 GB RAM per Slot
- 5 Support Nodes
- ROCKS 6 (RHEL 6), HTCondor, OSG 3
- Very old hardware, just retired
- Also, OUHEP\_OSG Tier3 HTCondor Desktop Cluster
  - 400 cores, 400 TB usable XFS/NFS storage
  - Used for local Athena, root, and theory calculations
  - 5 node OSG Integration Testbed cluster



## OU\_OSCER\_ATLAS Tier 2 Hardware

- 52 Nodes (2400 Slots) – 2 GB RAM per Slot
- 10 Support Nodes (1 GK, 1 DTN/Proxy, 1 redirector, 7 storage)
- 700 TB of usable xrootd storage (7 T630s with 16 8 TB drives, RAID6, xfs)
- SALT (CentOS 7.3), SLURM 17.11, OSG 3.4, XRootD 4.10
- Part of generic OSCER HPC cluster
- Rest of OSCER Schooner Hardware
  - 600 Nodes (about 18k Slots) – 2-4 GB RAM per Slot total
  - 300 Nodes (about 10k Slots) – 2-4 GB RAM per Slot public
  - Opportunistically available for ATLAS production

## Lucille Hardware

---

- 30 Nodes (960 Slots) – 4 GB RAM per Slot
- 5 Support Nodes (2 head, 3 storage)
- 110 TB of usable XFS storage (MD1200/MD3200)
- ROCKS 6.1 (CentOS 6.9), HTCondor, OSG 3.3
- Currently being upgraded to OpenHPC, CentOS 7, SLURM
- Cluster also has some GPU nodes that have been used for other OSG testing

## UTA Tier 2 Hardware

---

- **SWT2\_CPB**

- 458 Nodes (9652 Slots) – 2-3 GB RAM per Slot
- 44 Support Nodes (13 head, 31 storage)
- 5500 TB of usable xrootd storage (MD1000/MD1200/MD3X60)
- ROCKS 7.0 (CentOS 7.6), SLURM, OSG 3.4

- **UTA\_SWT2**

- 108 Nodes (2328 Slots) – 2-3 GB RAM per Slot
- 16 Support Nodes (10 head, 6 storage)
- 325 TB of usable xrootd storage (MD1000/MD1200)
- ROCKS 7.0 (CentOS 7.6), SLURM, OSG 3.4

## Network

---

- OU connected at 100 Gbps to I2 and ESnet via OneNet
- LU connected at 10 Gbps to OneNet
- UTA peers with LEARN on Campus at 100 Gbps
- OSCER connected at dual 10 Gbps to OneNet
- Internally, ipv6 ready; working on external connectivity
- Working on 40 Gbps DMZ for OSCER cluster
  - Migrate DTN from central OU network to DMZ
  - Put that subnet onto LHCONE
  - perfSonar node already on that subnet

## Equipment Deployment

---

- Have 45 R430 and 7 R630 compute nodes deployed in Schooner as dedicated 'Condo' nodes, as well as 7 T630 storage nodes
- Have ordered and received 4 R440s with Cascade Lake CPUs (80 slots/node, 192 GB of RAM), to be deployed shortly
- Old OCHEP cluster just retired
- Space, Power, Cooling not an issue
- UTA is currently deploying 1.2 PB of storage (to replace older storage) and replacing 10 older compute nodes with 10 R440 (56 slots/machine)



## Opportunistic Cycles, Containers

---

- Opportunistic Cycles

- 10k cores on Schooner
- Running Event Service jobs on new hep\_killable queue
- Getting very good usage as 'fair share OU user' – up to 3k at times

- Containers

- Singularity 2.6.1 installed on all OSCER nodes
- Running all dedicated ATLAS jobs in Containers now
- Ready to switch over Opportunistic Jobs to Containers now as well



## XRootD Testing

---

- 700 TB OU XRootD storage very stable and performant
- Part of Third Party Copy (TPC) testbed
- Able to transfer 800+ MB/s, which is probably close to current available OSCER core network limit
- Recently switched to pre-assigned space groups (available in v4.10), to make space reporting easier
- Next: enable HTTP for XRootD
- Looking forward to network improvement once on DMZ / LHCONE

## Summary and Conclusions

- SWT2 active and successful ATLAS Tier2 Center
- Consistently in Top 3 in the US, Top 5 in the World
- Very active in ATLAS/Panda Testing and Development
- Also OSG and XRootD Integration Testing and Deployment

