

US ATLAS SWT2 SITE REPORT

HORST SEVERINI, JOEL SNOW, PATRICK MCGUIGAN, MARK SOSEBEE

HEPIX MEETING SPRING 2016

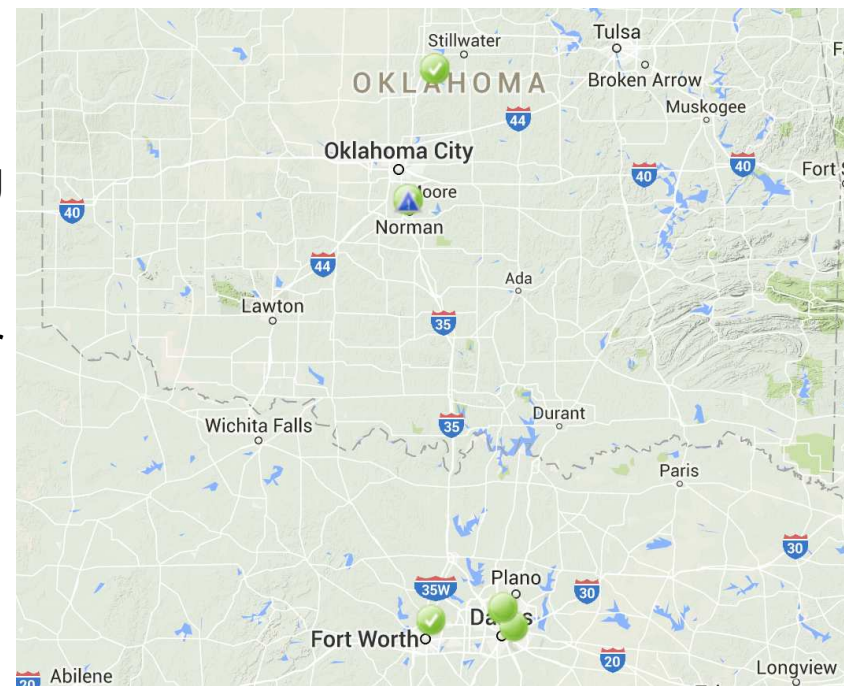
APRIL 2016

Outline

- Introduction
- Computing and Storage Hardware
- Network
- Upgrade Deployment
- Opportunistic Cycles

US ATLAS SWT2 Center

- **University of Oklahoma**
 - Oklahoma Center for High Energy Physics (OCHEP)
 - 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



OU_OCHEP_SWT2 Tier 2 Hardware

- 72 Nodes (844 Slots) – 2 GB RAM per Slot
- 8 Support Nodes (5 head, 3 storage)
- 500 TB of usable Lustre storage (DDN 9900)
- ROCKS 6.1 (RHEL 6.4), HTCondor, OSG 3.2
- Additionally, 20 new R430s and 7 R630s for OSCER cluster
(more below)



Lucille and OSCER Hardware

- **Lucille**
 - 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 6.7), HTCondor, OSG 3.3
- **OSCER Schooner Hardware (mostly opportunistic)**
 - 509 Nodes (10,368 Slots) – 2-4 GB RAM per Slot
 - 23 Nodes (600 Slots) dedicated for ATLAS
 - SALT (CentOS 7.2), SLURM, OSG 3.3



UTA Tier 2 Hardware

- **SWT2_CPB**
 - 355 Nodes (4488 Slots) – 2-3 GB RAM per Slot
 - 38 Support Nodes (13 head, 25 storage)
 - 2700 TB of usable xrootd storage (MD1200/MD3660i)
 - ROCKS 6.1 (SL 6.3), PBS/Torque, OSG 3.2
- **UTA_SWT2**
 - 157 Nodes (1496 Slots) – 2 GB RAM per Slot
 - 11 Support Nodes (6 head, 5 storage)
 - 120 TB of usable xrootd storage (MD1000/MD1200)
 - ROCKS 6.1 (SL 6.3), PBS/Torque, OSG 3.2



Network

- OU connected at 20 Gbps to I2 and ESnet via OneNet
- LU connected at 10 Gbps to OneNet
- UTA connected at 10 Gbps to LEARN
- 100 Gbps line to OneNet already live on OU campus in 1PP
- After cluster move to 4 PP, will have 40 Gbps available on DMZ
- New OSCER cluster Schooner on that same DMZ network
- Locally, can get over 2 GB/s throughput from Lustre servers to compute nodes
- UTA also plans DMZ connection at 20 Gbps



Equipment Deployment

- Have 17 R430 and 7 R630 compute nodes deployed in Schooner as dedicated 'Condo' nodes
- Have ordered 3 more R430s from first part of FY16 funds
- Will keep running older nodes until they die
- Space, Power, Cooling not an issue
- LU submitted MRI for new cluster
- UTA ordered 30 more R630s with final FY15 funds



Middleware, Dynamic Queue Deployment

- **OU_OCHEP_SWT2:** upgrade to Rocks 6.2, RHEL6.7, OSG 3.3 during next downtime
- **OU_OSCER_ATLAS:** OSG 3.3 deployed on CentOS7, testing
 - Initially, some issues with SLURM jobmanager, SELinux, and Gratia accounting
 - Now solved, continue testing
- **Lucille_CE:** OSG 3.3, fully updated
 - HTCondor-CE with whole node dynamically partitionable slots
 - Will be on same Oklahoma Friction Free Network (OFFN) DMZ as OSCER and OCHEP
- **UTA** also about to upgrade to Rocks 6.2, SL 6.7, OSG 3.3



Opportunistic Cycles

- 10,368 cores on Schooner
- Expect good usage as 'fair share OU user'
- Additionally, expect significant preemptable usage of nodes dedicated to weather forecasting 2h every day (therefore 22h of potential use)
- UTA: I/UCRC Center in Energy Efficient Systems
 - Received hardware donation from Yahoo: 450+ servers
 - Working with with CSE to setup cloud based computing services
 - Looking at making this a opportunistic production site

