
XSEDE Integration

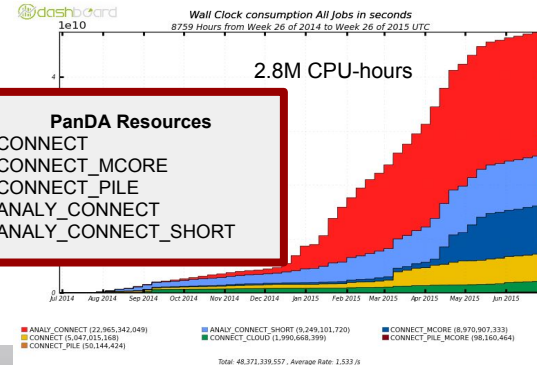


Rob Gardner

US ATLAS Physics, Software and Computing Technical Planning Meeting,
Chicago, August 1, 2016

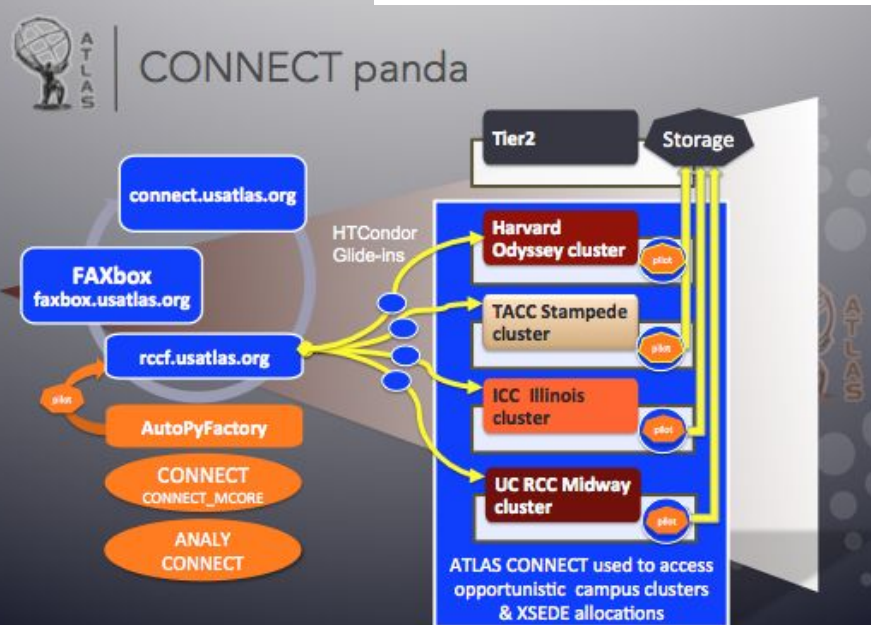
XSEDE and ATLAS activity

Stampede (2014-15)



Panda – TACC/Stampede Integration

- We have 250k hour XSEDE allocation on TACC Stampede
- Use ATLAS Connect to attach clusters to Panda
 - Simple SSH login to Stampede SLURM submit node
 - Submission of HTCondor glide-in's which receive ATLAS pilots; one stampede slot = 16 single core pilots or 2x8 core MCore jobs
 - Squid caches and outputs to a Rucio storage element at Midwest Tier2
 - Wide area federated storage (FAX) for input data access
 - **Minimize Stampede admin labor:**
 - Panda queues maintained by MWT2 admins
- Long standing software access issues **now resolved:**
 - ATLAS releases (incl. nightlies), conditions data, WLCG platform dependencies installed to local disk via "Stratum-R" servers at UIUC and Chicago (no FUSE kernel modules needed on compute nodes)
 - HammerCloud validations passing for all releases
- Scale testing this week (focus on MCore production)
- Stratum-R a general solution for CVMFS → shared NFS disk, typical for campus and HPC clusters

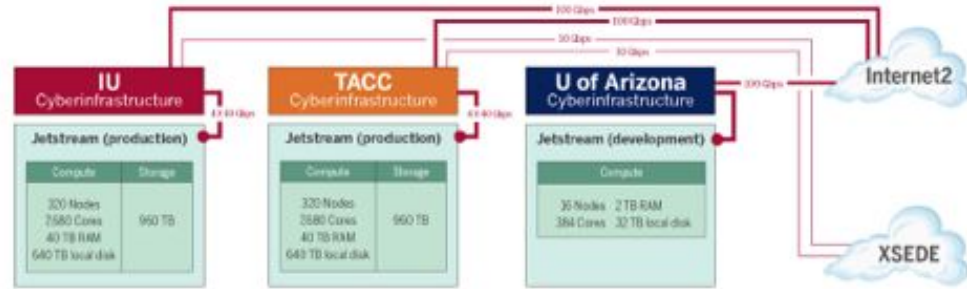


TACC → LHCONE in progress

XSEDE and ATLAS activity

- XRAC proposal Oct '15
 - Onysi (PI), Gardner, Lesny
 - “Cloud Computing on Jetstream for the ATLAS Experiment at the Large Hadron Collider”
- Jetstream VM resources
- Leverages ATLAS Connect & UT Cloud Tier3
- 2.1M SU (PHY140033)

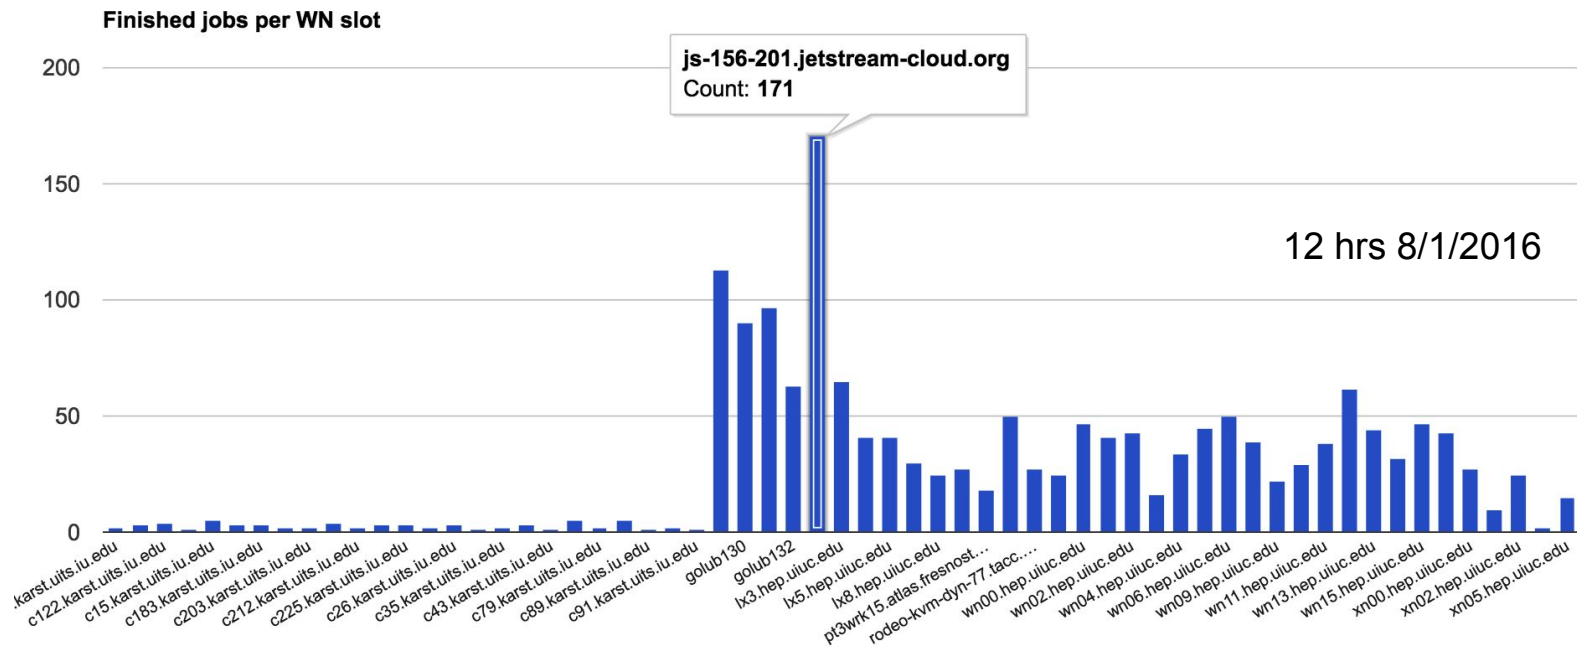
Jetstream System Overview



- Geographically distributed cloud; 0.5 PetaFLOPS
- Globus for large-scale file transfer, authentication

Commissioning Jetstream

<http://bigpanda.cern.ch/wns/CONNECT/>



Discussion, Summary

- Two major XSEDE resources integrated with the ATLAS production system
 - Automated. All releases. Any data at MWT2.
- Next logical resource would be Comet at SDSC
- Next steps?
 - Can provide as service to individual US ATLAS PI's
 - Special XRAC allocation given NSB review?

Extra slides (ATLAS Connect - Panda)

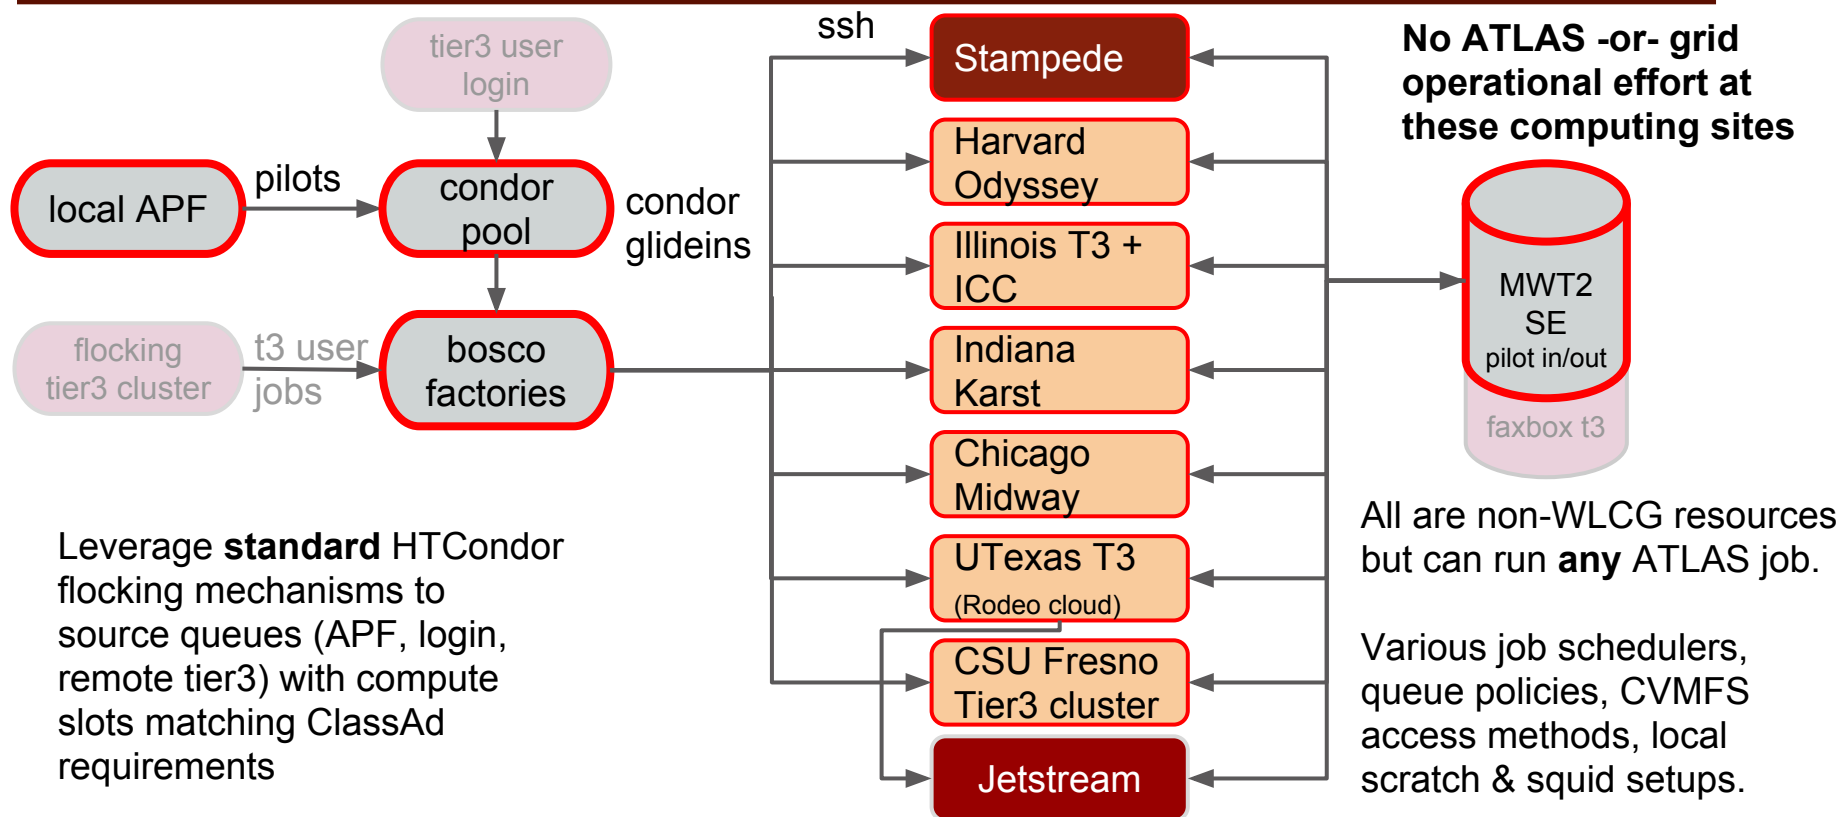
Distributed resource targets

- Goal is to make the connection lightweight
 - Local site provides only a user account, and optionally a frontier-squid server
- MWT2 hosts the ATLAS required services
 - CVMFS, input datasets, output storage
- Priorities and preemption policies differ
- Continuously run, but with variable availability
- Operated by MWT2 staff

ATLAS Connect & ADC

Panda

tier3 users



Technical Challenges

- Software access
 - Multiple approaches taken to deliver CVMFS to the site
 - “nfsCVMS” - client on server, export to workers
 - “portableCVMFS” - comes with job, user loads fuse modules dynamically
 - “replicaCVMFS” - replicate CVMFS tree to ‘local’ server, followed by rsync to target system disk
 - requires sys admin provide /cvmfs sym link to project disk
- Data delivery is via WAN