
South Pole Telescope Computing on OSG

OSG All Hands Meeting @ SDSC

March 6, 2017

Benedikt Riedel
Judith Stephen
University of Chicago



Open Science Grid




Overview

- Experiment & Science
- Computing Requirements
- New Infrastructure Deployment
- OSG Usage
- Future Plans

SPT Experiment

- Microwave-Millimeter telescope located at the South Pole
- Upgrades in 2016/2017:
 - SPT 3G - Third generation detectors installed - Data taking about to start
 - 10x detectors = 10x data rate
 - Added detection channel
 - New compute and storage hardware at South Pole and UChicago



The SPT-3G Collaboration (Feb. 2016)

~70 scientists, across ~20+ institutions

Funded By:



U.S. DEPARTMENT OF
ENERGY




Science

- Cosmic Microwave Background: CMB power spectrum, CMB Lensing, CMB Polarization, Epoch of Reionization
- Galaxy Clusters: Cluster Cosmology, Cluster Catalog
- Galaxy Evolution: High- z strongly-lensed star-forming galaxies, Catalog of mm sources

Computing Requirements

- Particle Physics/Astronomy-style computing: Processing exposures of the sky, MC generation of comparison skies
- Data:
 - South Pole: 450 TB storage for data
 - North: ~100 GB/day over satellite, ~200 TB/year carried back from pole
- CPU:
 - South Pole: ~0.5M CPUhours/year
 - North: ~30M CPUhours/year, 1000ish jobs requiring 70 - 100 GB RAM
- Main Support Tasks:
 - Setup of new compute infrastructure at South Pole and UChicago
 - OSG access and support for data processing and Monte Carlo production

At the South Pole

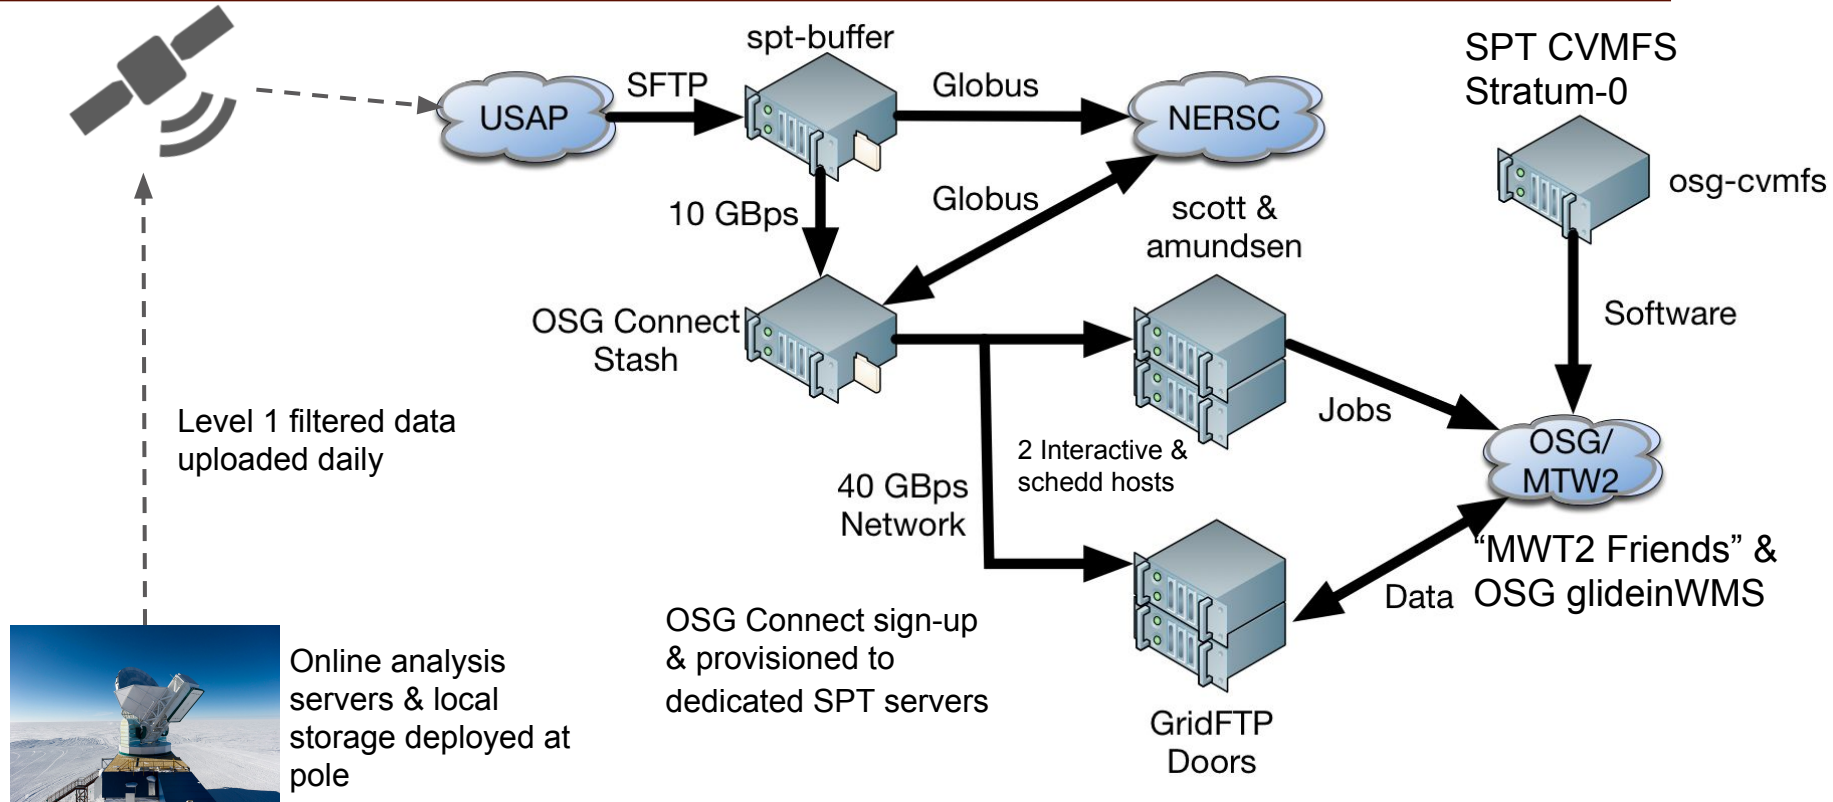


New Infrastructure Deployment - South Pole



- New Hardware in red (no Fry's down the road edition):
 - 4x Dell R730s:
 - 2x R730 for analysis work (HTCondor pool)
 - 1x R730 as hypervisor
 - 1x R730 hot spare
 - 2x Dell R330s: Storage controller + backup
 - 2x Dell MD1280s:
 - Primary Copy: ZFS pool, 42x 8 TB, NFS mounted to all R730s
 - Secondary Copy: JBOD, 28x 8 TB
 - 2x UPSes, 6x PDUs
- Old hardware in green - Part of analysis HTCondor pool


New Infrastructure Deployment - UChicago



New Infrastructure Deployment - UChicago


● Data Infrastructure:

- spt-buffer:
 - VM with 4 TB disk attached
 - Total 40 day buffer
- /spt
 - Part of OSG Stash (CephFS)
 - rsync'ed from spt-buffer
- Backup from spt-buffer to NERSC HPSS, partially automated



New Infrastructure for Software & Analysis

- Analysis Servers
 - "scott" & "amundsen"
 - Dell R630s
 - Login, interactive, schedd
 - 10Gbps
- /spt from OSG Stash
- 2x gridftp with bonded 2x10 Gbps
- osg-cvmfs - SPT
- OASIS Sratum-0



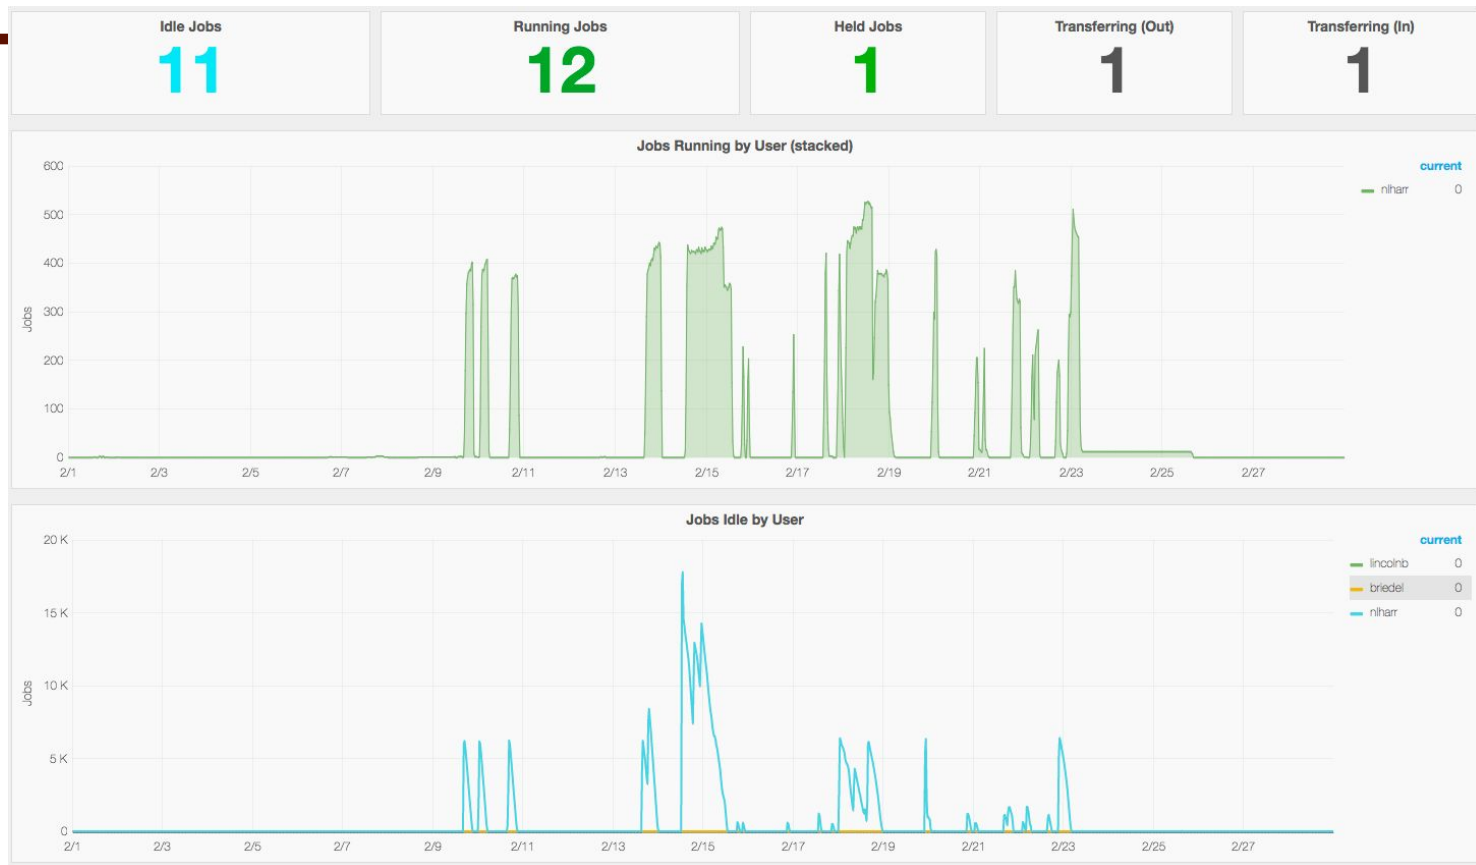
New Infrastructure for processing: OSG "ready"

- MWT2 Friends - Special opportunistic priority on MWT2 compute resources for groups working with UChicago OSG Team
- Data Movement - SPT groups with local storage pull data through Stash Globus endpoint or gridftp endpoint, depends on local security

Status Today

- Data (experiment is **live!**)
 - Ingested over 1 TB of new data into /spt
 - Ingesting data from 2nd generation experiment into Stash Object Store
- Compute:
 - 100k jobs submitted to OSG so far
 - Consistently getting 500 slots
 - Usage continuing to rise as more data comes in and users come back from South Pole

OSG Usage



Future Plans

- Backup Automation
 - Looking into automating the backup to NERSC - Globus Shared Endpoints, GridFTP-based data management, etc.
- Jupyterhub
 - Replace VNC-based workflows with JupyterHub-based workflows
 - Easier to administer and maintain
- SL7 Builds in CVMFS: Adding machine to generate SL7 builds for cvmfs

Summary

- OSG is working closely with SPT from the telescope at the pole to the grid to facilitate data handling, processing & simulation
- Online data collection and analysis facility built at the South Pole
- OSG-"ready" infrastructure deployed and ready for SPT-3G data
- Looking forward to new physics in 2017!