
XENON1T Update

LHCOPN-LHCONE Meeting @ Brookhaven

April 4, 2017



Open Science Grid

Benedikt Riedel
Suchandra Thapa
Rob Gardner
Evan Shockley
Luca Grandi (PI)
University of Chicago

Patrick de Perio
Columbia University
Fabio Valerio Massoli
University of Bologna
Jialing Fei, UCSD

Boris Bauermeister
Stockholms Universitet



Acknowledgements

- Slides adapted from Benedikt Riedel's Xenon1T computing report at the OSG All Hands Meeting @ UCSD / SDSC last month

Requirements

- Particle physics computing, i.e. data transfer to various sites, event reconstruction, Monte Carlo production
- Data: ~1.5 TB/day from LNGS (INFN)
- CPU: ~30 MCPUhours/year
- Main issues for LHCONE:
 - Data transfer from LNGS to various sites, and between sites
 - Access and support for event reconstruction and Monte Carlo production on OSG and EGI

Data Management

- Setup a distributed data management system using "Rucio" (data management system developed and used by ATLAS) and FTS at BNL
 - Rucio: File catalog service, subscription "rule" model for data placement and deletion, etc.
 - FTS: Reliable file transfer
- Currently 6 storage endpoints defined in the US, Europe and Israel



Xenon1T Rucio Status



IN2P3 RSE Status

#813 2 hours ago

IN2P3 RSE Transfer

#814 57 minutes ago

LNGS RSE Status

#2424 2 hours ago

LNGS RSE Transfer

#1004 1 hour ago

Nikhef RSE Status

#2420 2 hours ago

Nikhef RSE Transfer

#1001 2 hours ago

RCC RSE Status

#2136 1 hour ago

RCC RSE Transfer

#508 2 hours ago

Rucio Central Service

#2428 1 hour ago

Stash RSE Status

#2419 1 hour ago

Stash RSE Transfer

#993 1 hour ago

Weizmann RSE Status

#344 2 hours ago

Weizmann RSE Transfer

#346 1 hour ago

Monitoring Automated FTS Transfers

Generated at 6:23:41 PM (hs03.usatlas.bnl.gov) Overview Jobs Optimizer Error reasons Statistics Configuration Job id

Filter destination SE xenon.biggrid.nl Source storage Destination storage 12 hours Apply Reset

Overview

Showing 1 to 3 out of 3 from the last 12 hours

First Previous 1 Next Last

source

destination

	VO	Submitted	Active	Staging	S.Active	Finished	Failed	Cancel	Rate (last 1h)	VO Thr.	
+ srm://tbn18.nikhef.nl	gsiftp://gridftp.grid.uchicago.edu	xenon.k	32382	12	-	-	11333	43	5	99.62 %	31.14 MB/s
+ gsiftp://ccdcac1022.in2p3.fr	gsiftp://gridftp.grid.uchicago.edu	xenon.k	3215	12	-	-	1498	381	174	79.72 %	1.61 MB/s
+ gsiftp://xelt-datamanager.lngs.infn.it	srm://tbn18.nikhef.nl	xenon.k	-	-	-	-	348	1	-	99.71 %	-
			35597	24	0	0	13179	425	179	96.88 %	-

First Previous 1 Next Last

- Bad shape** There are submitted but no active, less than 3 active with more than 3 submitted, or a failure rate $\geq 20\%$
- Underused** Less than three actives, but no submitted waiting.
- Good shape** Success rate $\geq 90\%$, or more than three actives with a failure rate $< 20\%$.
- Nothing special** No active, no submitted, success rate between 80% and 90%.

Transfer rate





Thanks to BNL for lending use of FTS instance

Monitoring Automated FTS Transfers

Generated at 6:30:03 PM (rs03.usatlas.bnl.gov) Overview Jobs Optimizer Error reasons Statistics Configuration Job id

- All - Source storage Destination storage 12 hours Apply Reset

Details for srm://tbn18.nikhef.nl → gsiftp://gridftp.grid.uchicago.edu



First Previous 1 2 3 4 5 6 7 8 9 10 11 12 13 Next Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2017-03-02T00:29:49	11	12	32382	100.00%	25.404 MB/s	63.950 MB/s	-1	Good link efficiency, throughput deterioration
2017-03-02T00:28:44	12	13	32389	100.00%	40.042 MB/s	68.233 MB/s	-1	Good link efficiency, throughput deterioration
2017-03-02T00:27:38	13	14	32399	100.00%	53.257 MB/s	71.365 MB/s	-1	Good link efficiency, throughput deterioration
2017-03-02T00:26:34	14	11	32410	100.00%	42.797 MB/s	73.377 MB/s	0	Good link efficiency, throughput deterioration, avg. filesize decreasing

Dashboard


x1t-rucio-test-dashboard

New Add Save Open Share Options ⌕ March 5th 2017, 23:54:20.227 to April 2nd 2017, 22:34:58.023



x1t-rucio



Xenon 1T Rucio Successful Transfers Total Throughput per Hour



Xenon 1T Rucio Successful Transfers Source RSEs per Day



Major Improvements

- A primary need is to quickly process data from the detector using dedicated compute resources at UChicago
- We did three things to speed things up:
 - Transitioned from scripted scp to FTS3-managed gridftp
 - Rucio managed replication first to Nikhef, then Chicago
 - Moved their data endpoint at RCC-Chicago into the SciDMZ, peered with LHCONE

Managed XrootD storage infrastructure

- Number of datasets: 6384
- Number of files: 417218
- Total disk storage used: 407 TB
- 175 TB on tape at PDC at KTH (Sweden)
- 6 Rucio endpoints:
 - LNGS (Italy), IN2P3 (Lyon), Nikhef, OSG (Chicago), RCC (Chicago), Weizmann (Israel)
- XRootD cache at UCSD for access to Comet









Next steps

- LNGS peering with LHCONE
 - Process formally started last week
- Improve FTS configuration between major endpoints (LNGS, MWT2, NIKHEF, and CCIN2P3)

Extra



~130 scientists from 22 institutions

 Columbia
 RPI
 Nikhef
 Stockholm University
 Münster
 Mainz
 Chicago
 UCLA
 UC San Diego
 UCSD
 Rice
 Purdue
 Coimbra
 Subatech
 LPNHE
 Bologna
 LNGS
 INFN Torino
 MPIK
 Freiburg
 Zurich
 NYU | ABU DHABI
 NYUAD
 Weizmann



Laboratori Nazionali
del Gran Sasso
(LNGS), Italy

XENON1T




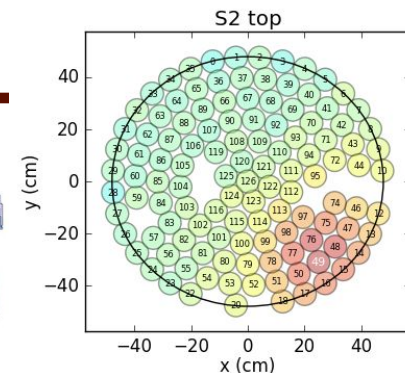

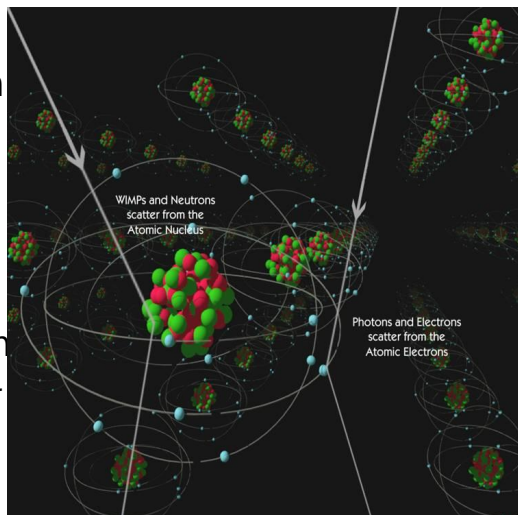



- 1st ton-scale experiment
- 3.2t of LXe, 2t in TPC
- All systems commissioned since Fall 2016
- Calibration and science data taking now ongoing



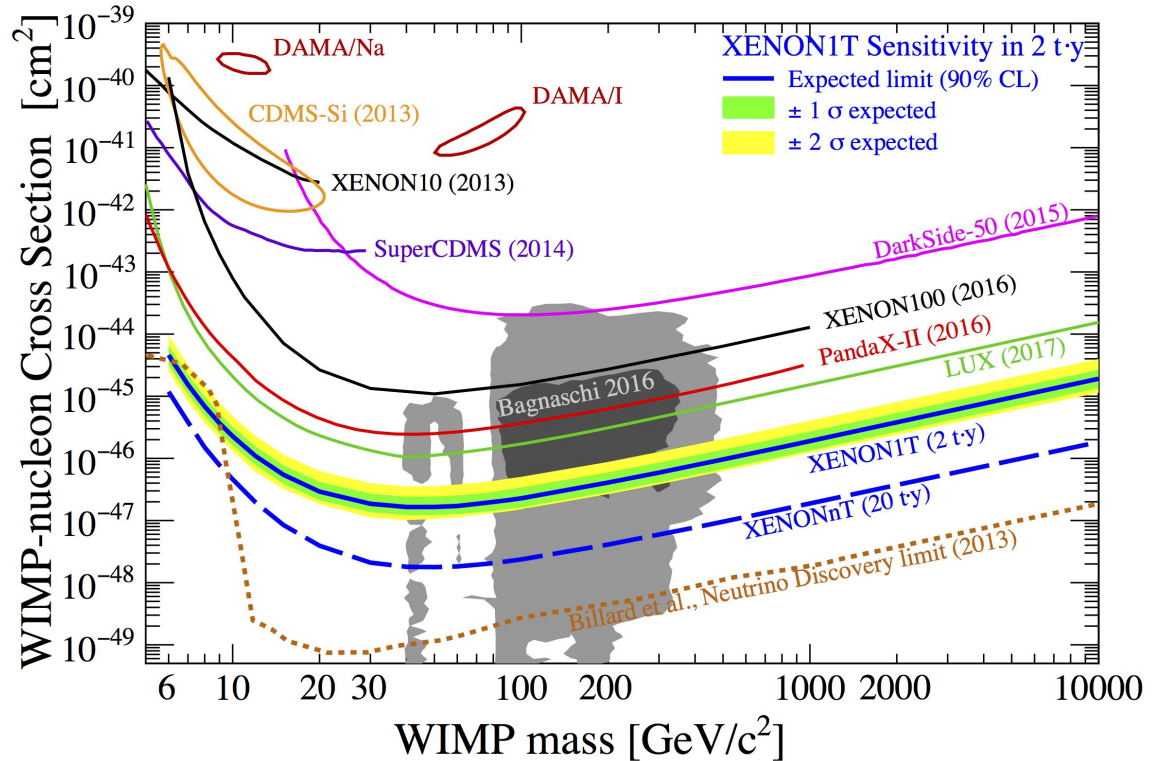
XENON1T Science

- No idea what 95.4% of Universe is made of
- Looking for Dark Matter candidate - Weakly Interactive Massive Particles (WIMP) - through nuclear recoil on liquid Xenon
- WIMP nuclear recoil produces charge and time signature in the detector readout that can easily distinguish nuclear recoil from other interactions



Projected WIMP Limits

- Projected to be sensitive in theoretical space
- Preparation for future larger detectors (XENONnT)



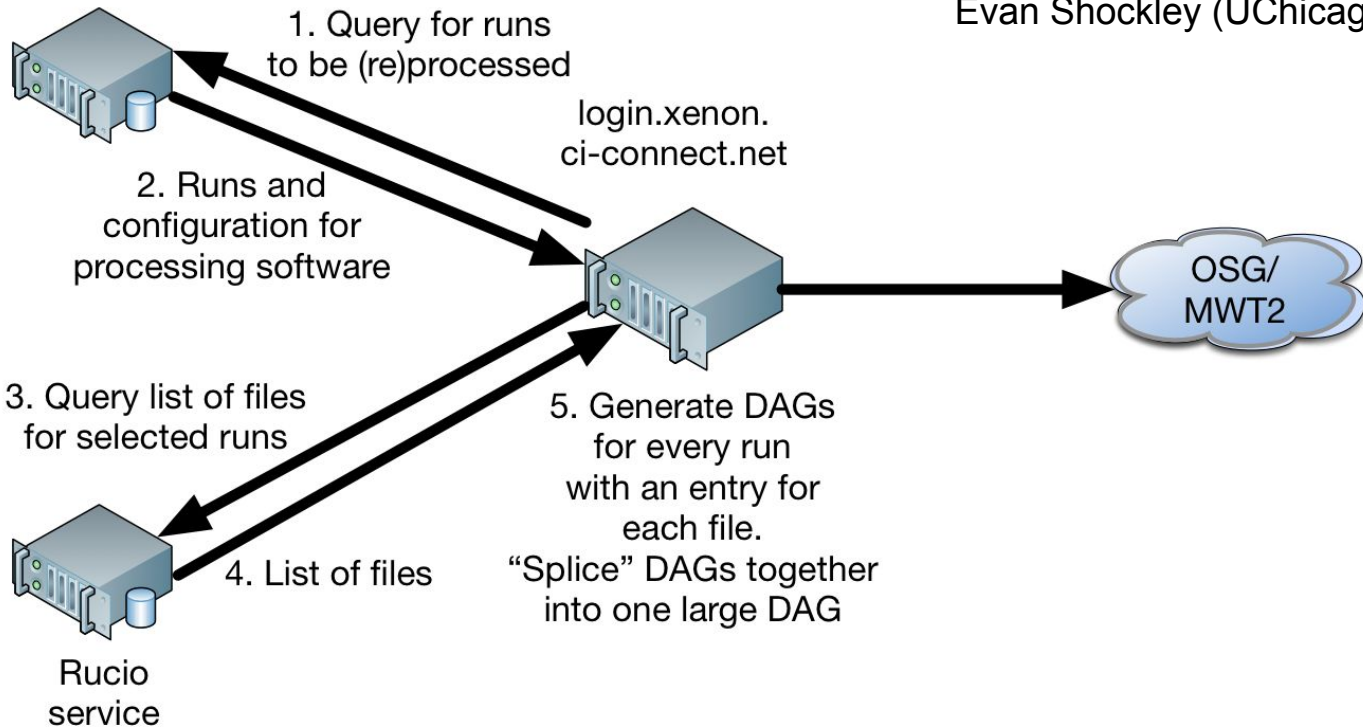
Xenon1T Infrastructure - UChicago

- Dedicated submit host - `login.xenon.ci-connect.net`
- MWT2 Friends - Special opportunistic priority on MWT2 compute resources
- 2x GridFTP doors for OSG Stash, shared with SPT
- CVMFS Stratum 0 - CVMFS deployment automated through DeployHQ targeting for new releases and changes to Xenon Git repository
- Backup Jupyterhub for analysis workshop

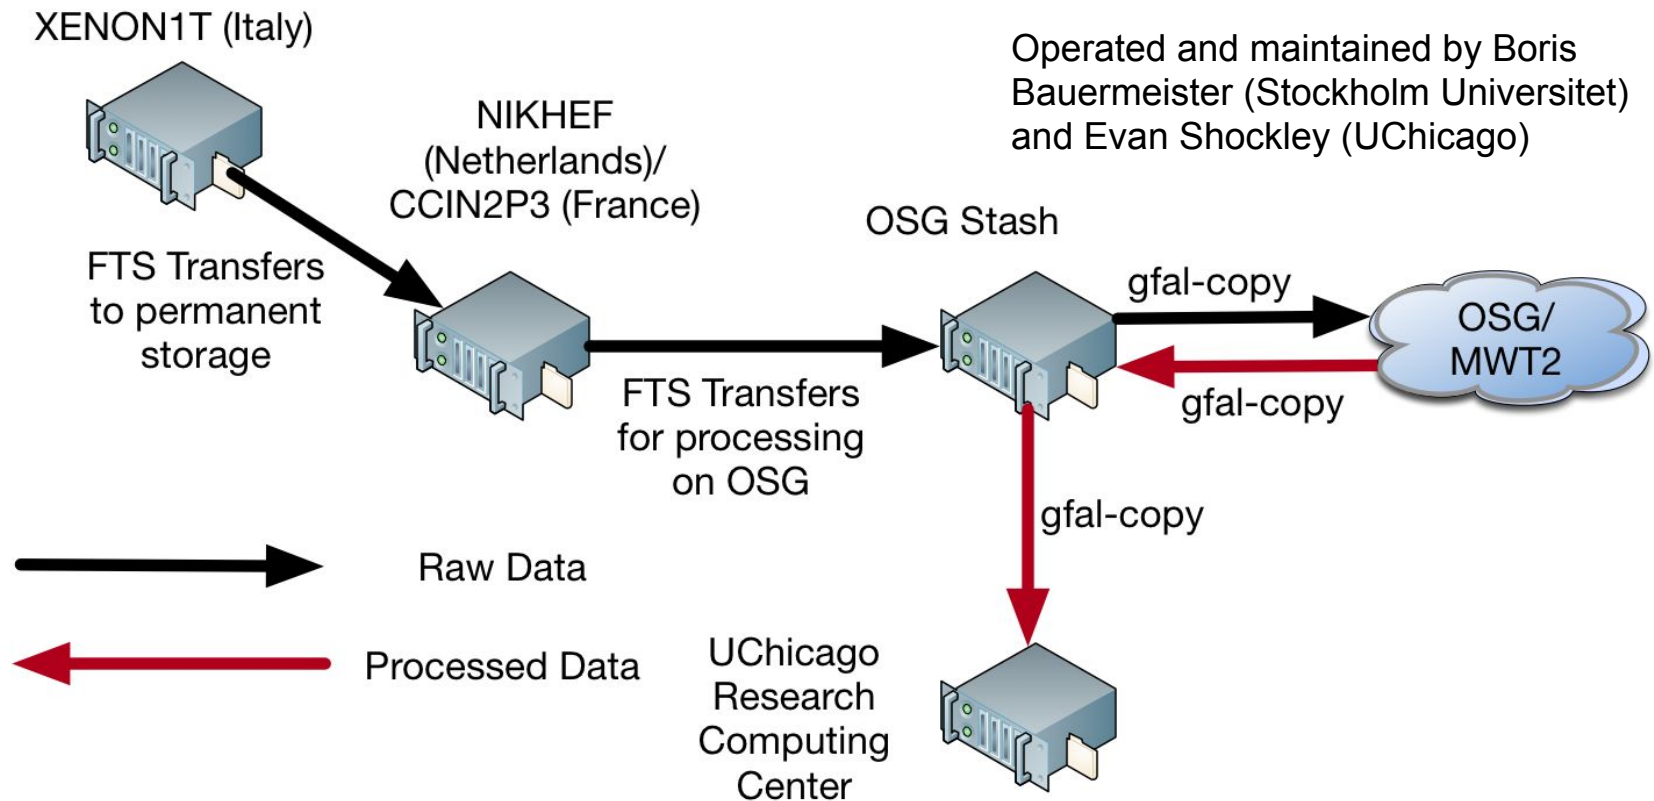
Data Processing Workflow

XENON1T
Runs
Database

Operated and maintained by
Evan Shockley (UChicago)


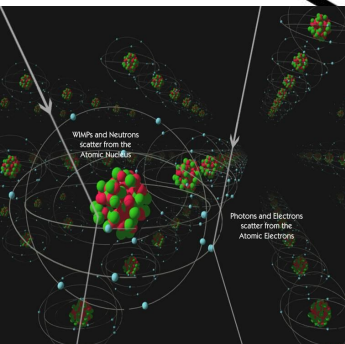


Data Flow for Processing




MonteCarlo Workflow

GEANT4
Event Generation





Generate Detector
Response (S1 and S2)

Waveform Generation



Reconstruction

Analysis ready output



Operated and
maintained by Patrick
de Perio (Columbia),
Fabio Valerio Massoli
(Bologna), and Jialing
Fei (UCSD)

MonteCarlo Workflow

- Implemented using **Pegasus** workflow manager to handle running jobs and errors
- Submit to OSG or EGI sites using a switch at runtime

```
mc_process.py --flavor G4 --config optPhot --source-macro  
run_optPhot_fullvolume.mac --batch-size 10000 --events 1000  
--mc-version v0.1.7 --pax-version v6.2.1 --grid-type osg
```

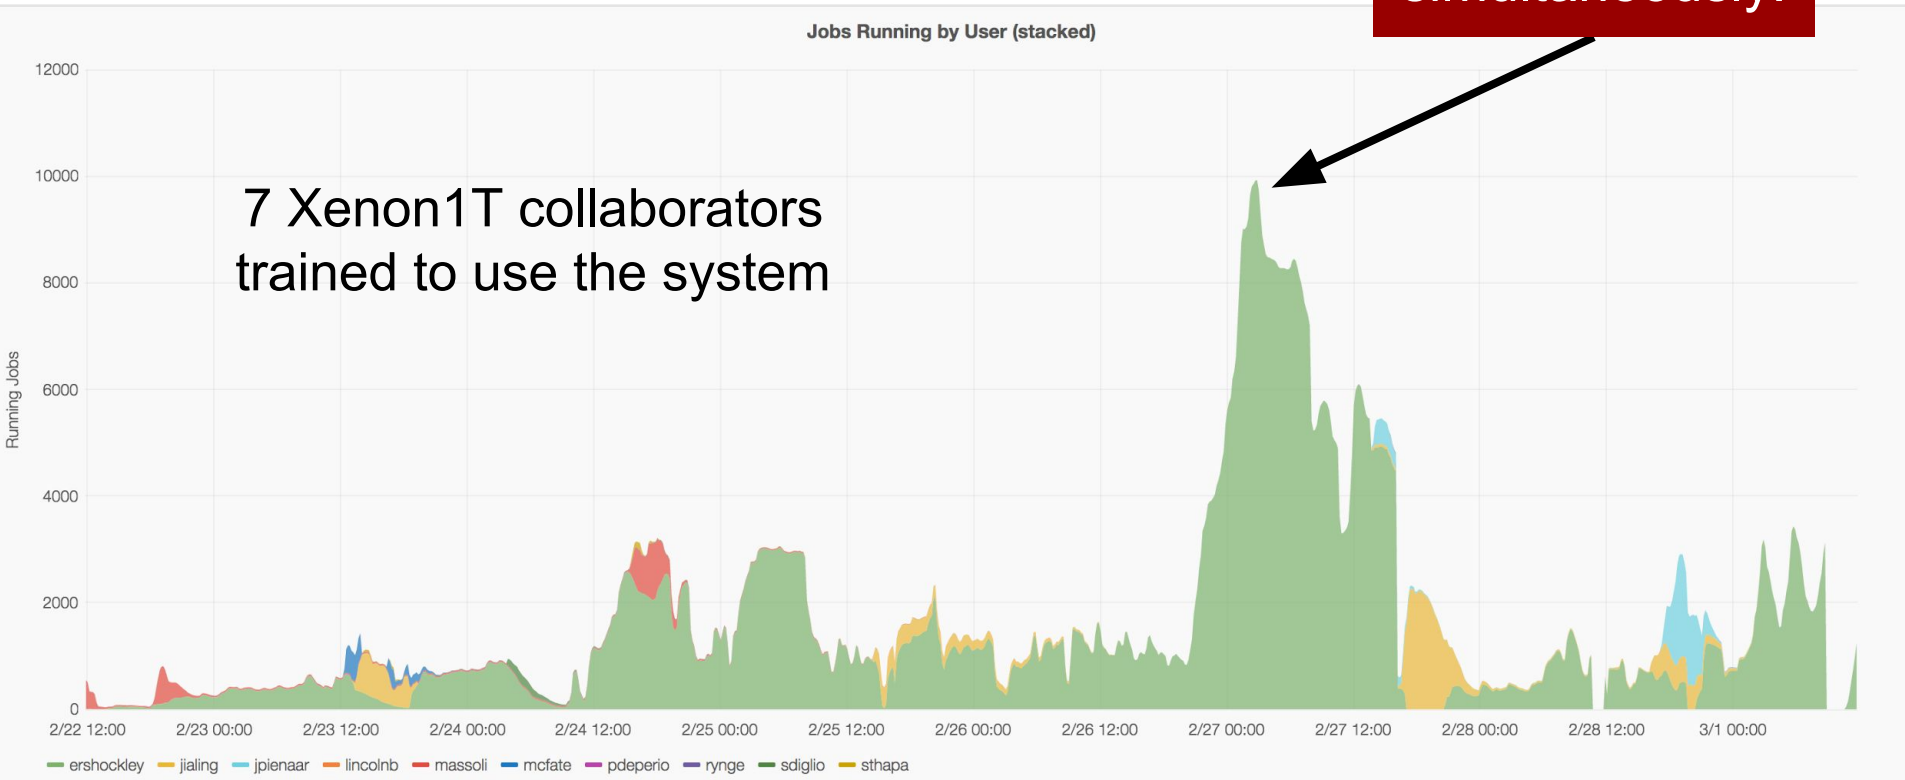
- 10+ data sets generated so far using OSG resources with production still ongoing

Status Today

- Data:
 - Rucio service has been stable modulo some peak usage periods, increased the maximum number of concurrent DB connections to resolve
 - OSG helped XENON1T become part of LHCOne
 - Transfer speeds between EU sites and US can vary greatly, under investigation
- Processing
 - Reprocessing campaign has started, expected use ~42 kCPUHours
 - Overall used 583933 CPUhours in February for reprocessing campaign and Monte Carlo production - Number 1 OSG Connect group

OSG Usage


10000 cores
simultaneously!




Comet usage with glideinWMS

- Added Comet as a dedicated target for XENON1T VO through glideinWMS
- Goes through osg-flock and a special glidein is started on the UCSD side
- Uses docker to setup environment
- Transparent to XENON1T users
- **Thanks to Edgar and Mats!**


By VO



By Facility



By Project



Xenon1T Compute Environment

- Integrated **campus cluster, OSG, EGI & HPC resources** into one submission environment
- Data management system that allows transparent access to data across continents

