BigPanDA on HPC's

Kaushik De

Univ. of Texas at Arlington

MAGIC Meeting

March 4, 2015

PanDA Overview

- PanDA workload management system was developed for the ATLAS experiment at the Large Hadron Collider
  - Hundreds of petabytes of data per year, thousands of users worldwide, many dozens of complex applications…
  - Leading to 398 scientific publications – and growing daily
  - Discovery of the Higgs boson, search for dark matter…
- A new approach to distributed computing
  - A huge hierarchy of computing centers working together
  - Main challenge – how to provide efficient automated performance
  - Auxiliary challenge – make resources easily accessible to all users

Core Ideas in PanDA

- Make hundreds of distributed sites appear as local
  - Provide a central queue for users – similar to local batch systems
- Reduce site related errors and reduce latency
  - Build a pilot job system – late transfer of user payloads
  - Crucial for distributed infrastructure maintained by local experts
- Hide middleware while supporting diversity and evolution
  - PanDA interacts with middleware – users see high level workflow
- Hide variations in infrastructure
  - PanDA presents uniform 'job' slots to user (with minimal sub-types)
  - Easy to integrate grid sites, clouds, HPC sites …
- Production and Analysis users see same PanDA system
  - Same set of distributed resources available to all users
  - Highly flexible – instantaneous control of global priorities by experiment

Paradigm Shift in HEP Computing

- New Ideas from PanDA
  - Distributed resources are seamlessly integrated
  - All users have access to resources worldwide through a single submission system
  - Uniform fair share, priorities and policies allow efficient management of resources
  - Automation, error handling, and other features in PanDA improve user experience
  - All users have access to same resources
- Old HEP paradigm
  - Distributed resources are independent entities
  - Groups of users utilize specific resources (whether locally or remotely)
  - Fair shares, priorities and policies are managed locally, for each resource
  - Uneven user experience at different sites, based on local support and experience
  - Privileged users have access to special resources

PanDA Scale

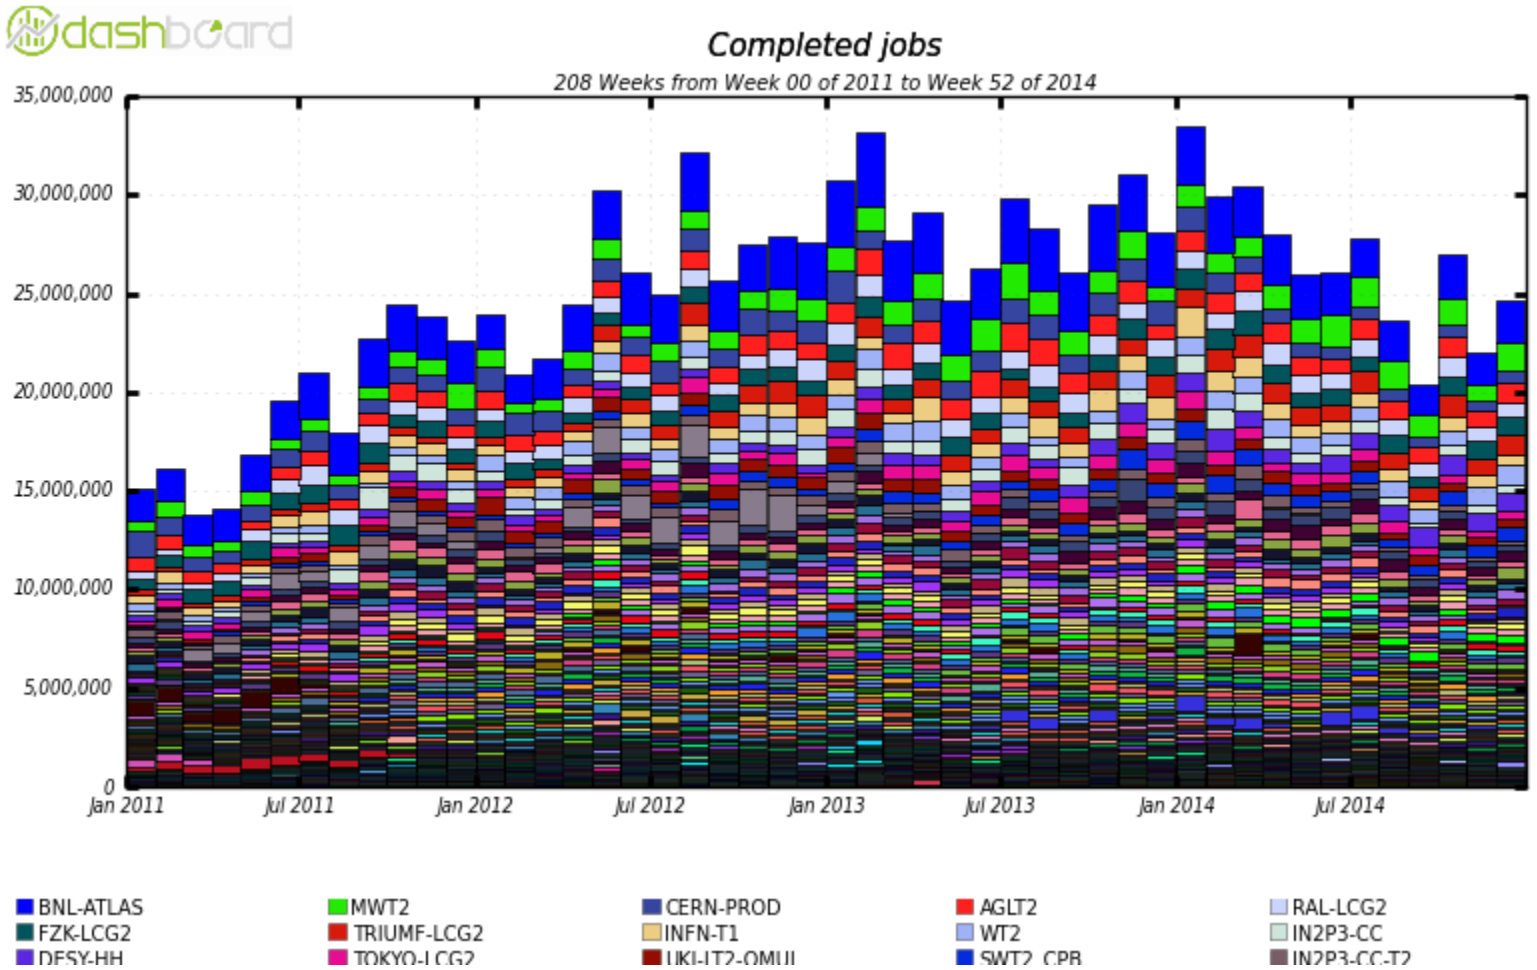

Current scale – 25M jobs completed every month at >hundred sites

First exascale system in HEP – 1.2 Exabytes processed in 2013

CPU Consumption

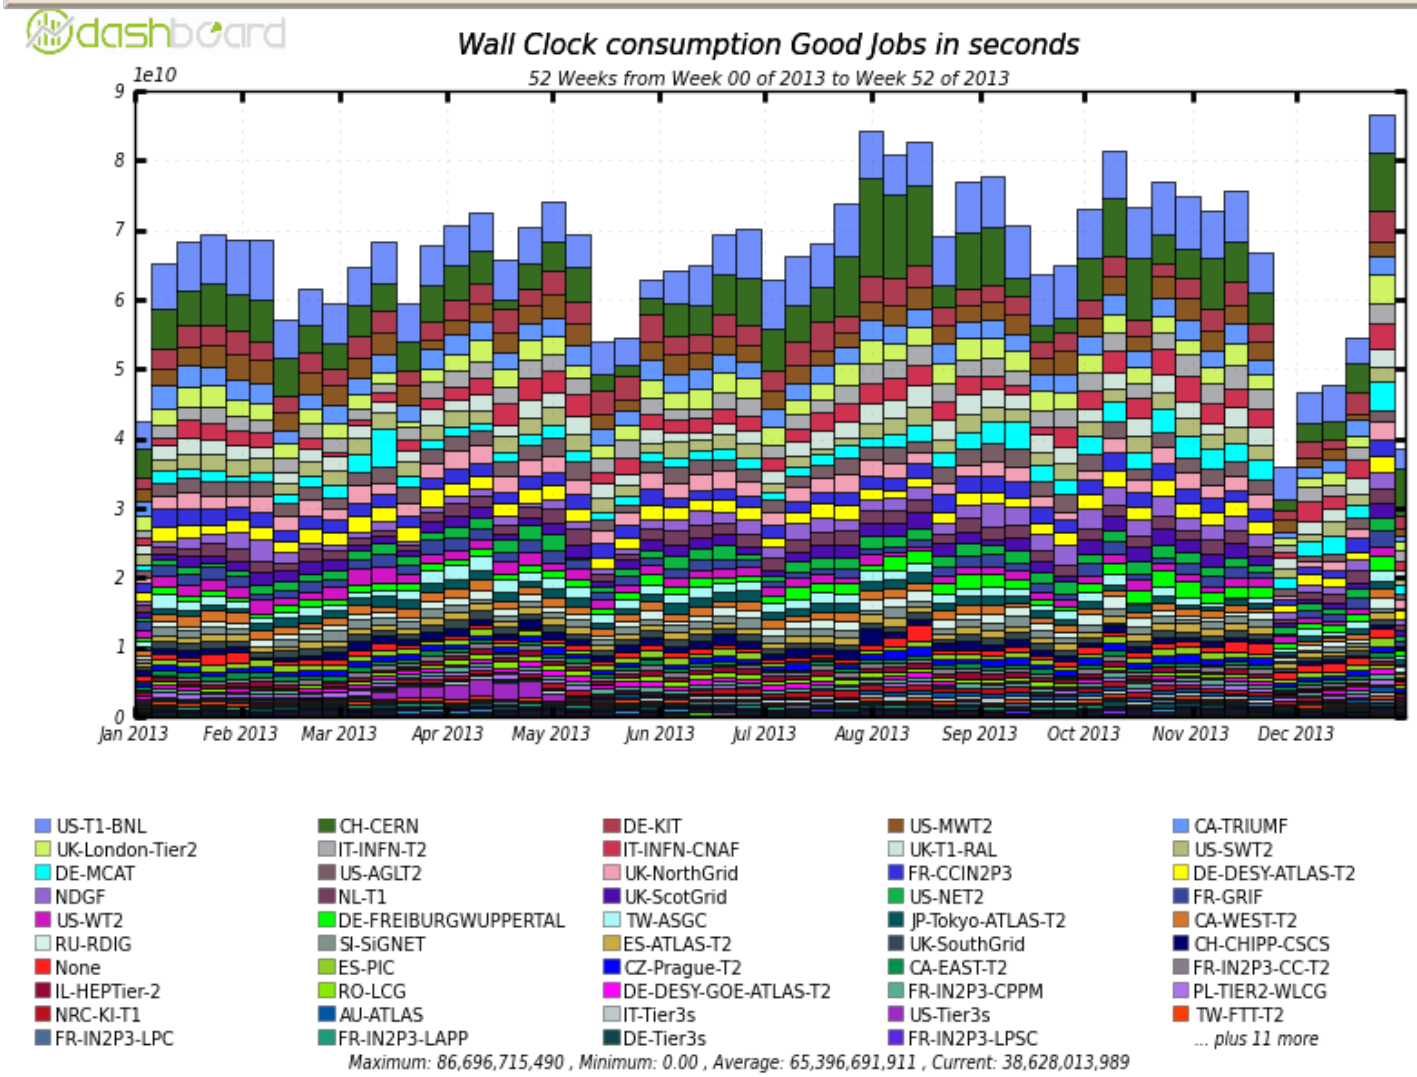

During Run 1, 2013, per week, aggregated by federation

PanDA Facts

- PanDA – Production and Distributed Analysis System
  - Deployed on WLCG infrastructure
  - Standards based implementation
    - REST framework – HTTP/S
    - Oracle or MySQL backends
    - CondorG based pilot factories
    - Python packages available from SVN and GitHub
    - Command-line and GUI/Web interfaces
- Reference
  - https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA

Next Generation "Big PanDA"

- ASCR and HEP funded project "Next Generation Workload Management and Analysis System for Big Data", 2012-2015
- Generalization of PanDA for HEP and other data-intensive sciences
- Project participants from ANL, BNL, UT Arlington
- Alexei Klimentov – Lead PI
- **WP1 (Factorizing the core):** Factorizing the core components of PanDA to enable adoption by a wide range of exascale scientific communities
- **WP2 (Extending the scope):** Evolving PanDA to support extreme scale computing clouds and Leadership Computing Facilities
- **WP3 (Leveraging intelligent networks)**: Integrating network services and real-time data access to the PanDA workflow
- **WP4 (Usability and monitoring):** Real time monitoring and visualization package for PanDA

WP2 Extending the scope PanDA @ ORNL LCF

- Special requirements for running on Titan
  - Extremely limited internet connectivity from worker nodes
  - No grid access – one-time password based authentication
  - PBS/TORQUE batch system
  - Full node jobs (16 cores)
  - Limitation on maximum number of jobs by one user
  - Parallel file system shared between nodes
  - Special data transfer nodes (high speed stage in/out)
- Many features developed through BigPanDA project to run ATLAS payload on Titan – beyond GRID's and clouds

Update on Titan activities

Example of Recent BigPanDA Work

Progress in 2014

- PanDA Pilot adopted for Titan (ORNL), Edison, Hopper (NERSC), Anselm (IT4I), work started at NRC KI ;
  - PanDA Pilot structure modified: execution on different HPC batch backends realized through plugins;
  - Introduced parameters for tuning of workload execution : limits for number of allocated nodes, minimum walltime, etc.;
- Pilot works with workload specific MPI wrapper that allows to parallelize serial workloads
- Algorithm to utilize unused resources on Titan (backfill) was introduced in Pilot and used for Titan job submissions ;
  - Also implemented local job resubmission strategy, in case job gets 'stuck' in local queue – significant decrease in wait time and increase in efficiency
- GridFTP file transfer tested between OLCF and BNL;
- Started adoption of Event Service on HPC
- Pilot for HPC branch regularly merged with production version of Pilot

ATLAS simulation on Titan

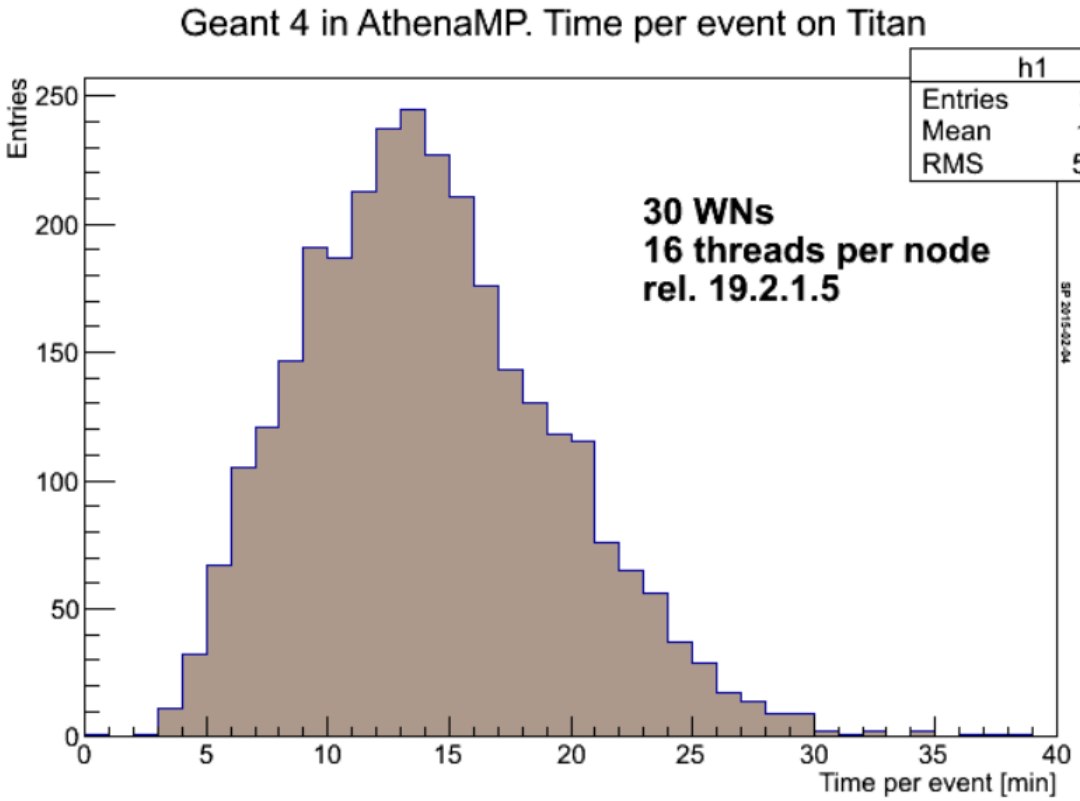

Sim_tf.py
1 input file
3000 Events
30 Worker nodes
16 worker threads per node
100 events per node

Average time per event ~14 min. Broad distribution from ~1 to ~40 minutes

Optimization of Initialization Time

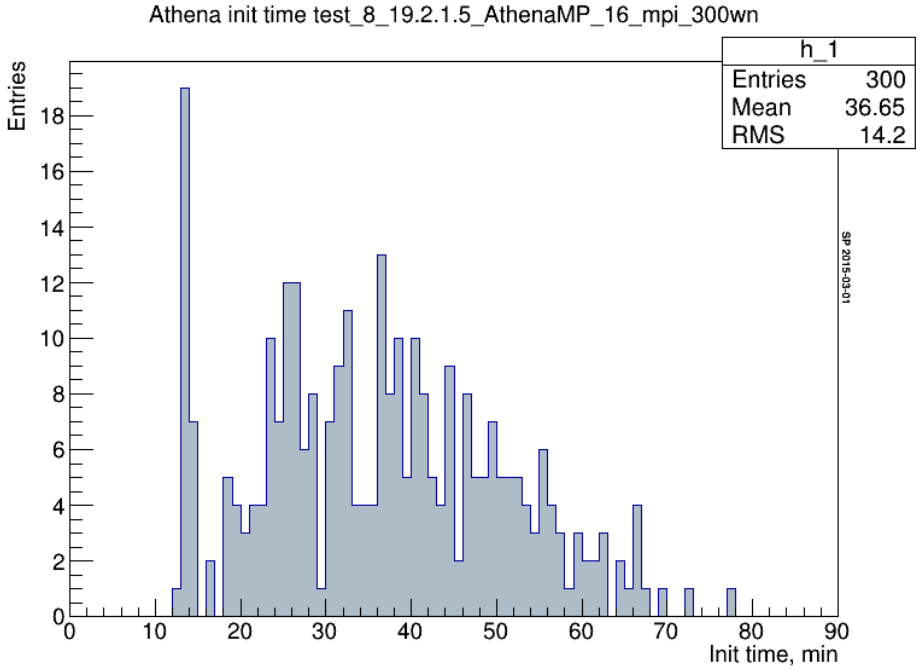

300 AthenaMP Jobs
On Titan

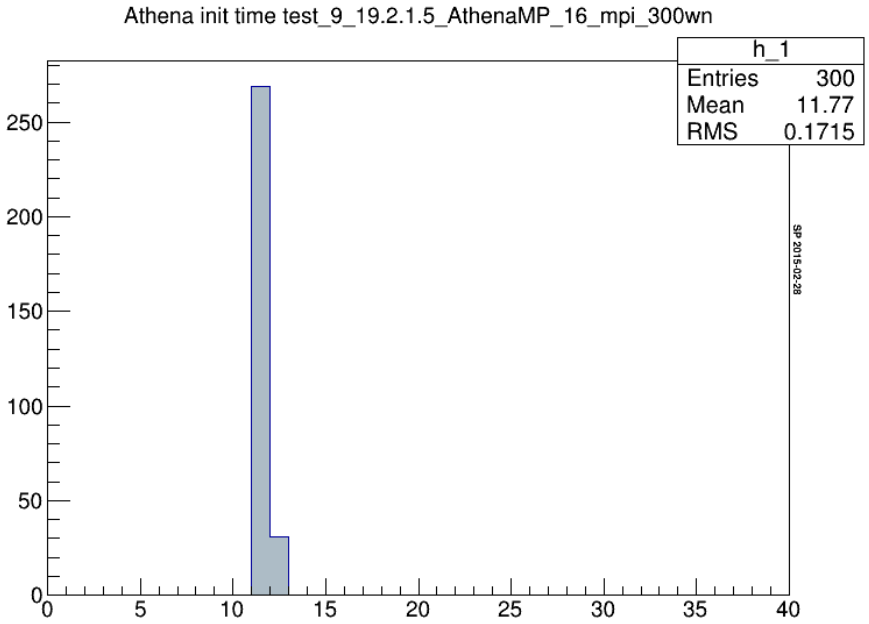

Why PanDA on HPC?

The Strategy – Requests

CPU:

- Stress the value of beyond pledge resources and give an approximate HS06 goal (~ twice the pledged resources)
- No change (wrt previous submission) to the requested pledged resources.

Storage:

- Disk: No change (wrt previous submission) to the requested pledged resources.
- Tape: Large increase in requested pledged resources in 2016 and 2017
- State value of beyond-pledge tape resources if they can be made available in 2015 (likely at some sites)

Rising CPU Needs at LHC

The Request

	2015 C-RSG	2016 C-RSG CERN–RRB–2014–079	2016 ATLAS (Revised)	2017 ATLAS (August 2014)	2017 ATLAS (Revised)
T0 CPU	205	257	257	273	273
T1 CPU	450	540	540	691	691
T2 CPU	520	608	608	732	732
T0 Disk	14	17	17	19	19
T1 Disk	36	47	47	58	58
T2 Disk	53	72	72	98	98
T0 Tape	33	42	42	54	54
T1 Tape	65	84	**116**	108	**208**

Other PanDA HPC Efforts

- Many efforts in addition to BigPanDA
- All teams are working together – common goals
- Major effort underway at NERSC
- Huge allocation at ALCF – MIRA
- Dozens of European and Asian HPC's

Example: Event Service

Event Service (ES)

- Change **job granularity from files to events**
  - Only deliver those events to the worker node, which will be processed by the payload job
  - Don't stage in entire input files
- Our approach to running ATLAS software on **opportunistic resources**
- The concept of the Event Service has been presented multiple times at various ATLAS meetings
- Event Service **wiki page** with a detailed description of the concept, workflow, development plans and more

  https://twiki.cern.ch/twiki/bin/view/PanDA/EventServer

- **First use-case: Geant4 Simulation**

V. Tsulaia Dec-3, 2014

Event Service at HPC

Yoda: Event Service @ HPC

- We cannot run the **conventional ES** on most HPC platforms (no internet connection from HPC compute nodes)
- For such machines ES had to be implemented in a different way

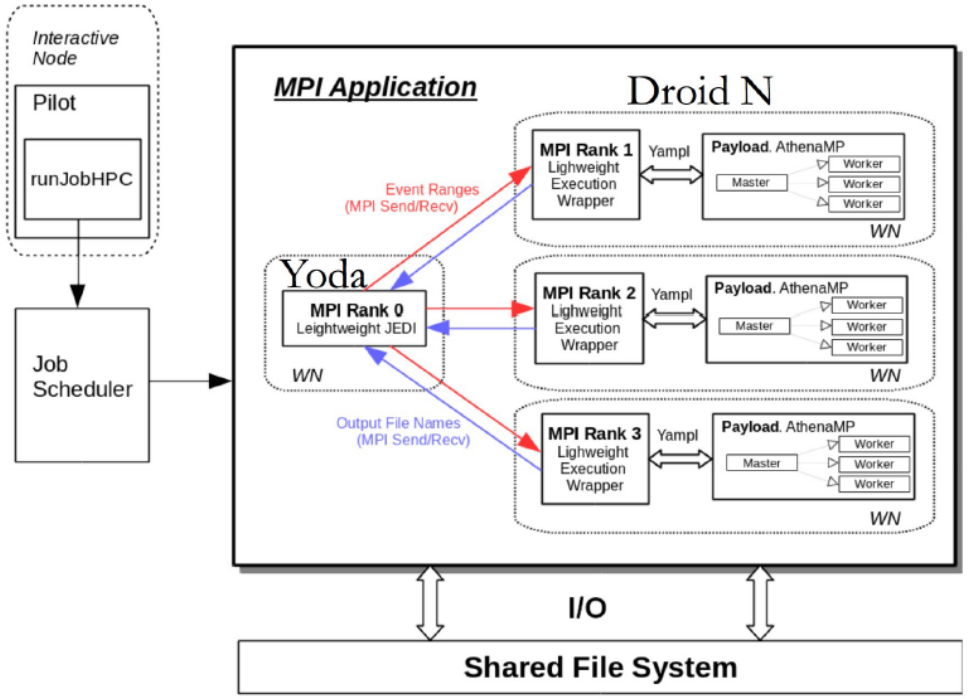

- **MPI application**
- **Reuse conventional ES code wherever possible**
- **Rank 0 (Yoda, master).** Distributes workload between slave ranks
- **Fine grained workload:** individual events or event ranges
- **Rank N (Droid, slave).** Processes assigned workload, saves output to the shared file system, asks for the next workload ...

V. Tsulaia Dec-3, 2014

Edison tests @ NERSC

"Killable queue" test

- A backfilling-like test
  - Submit a bunch of large jobs to the killable queue either before the scheduled maintenance time, or before a large reservation created specifically for this test
- Yoda jobs run in the killable queue until they get terminated by the batch scheduler

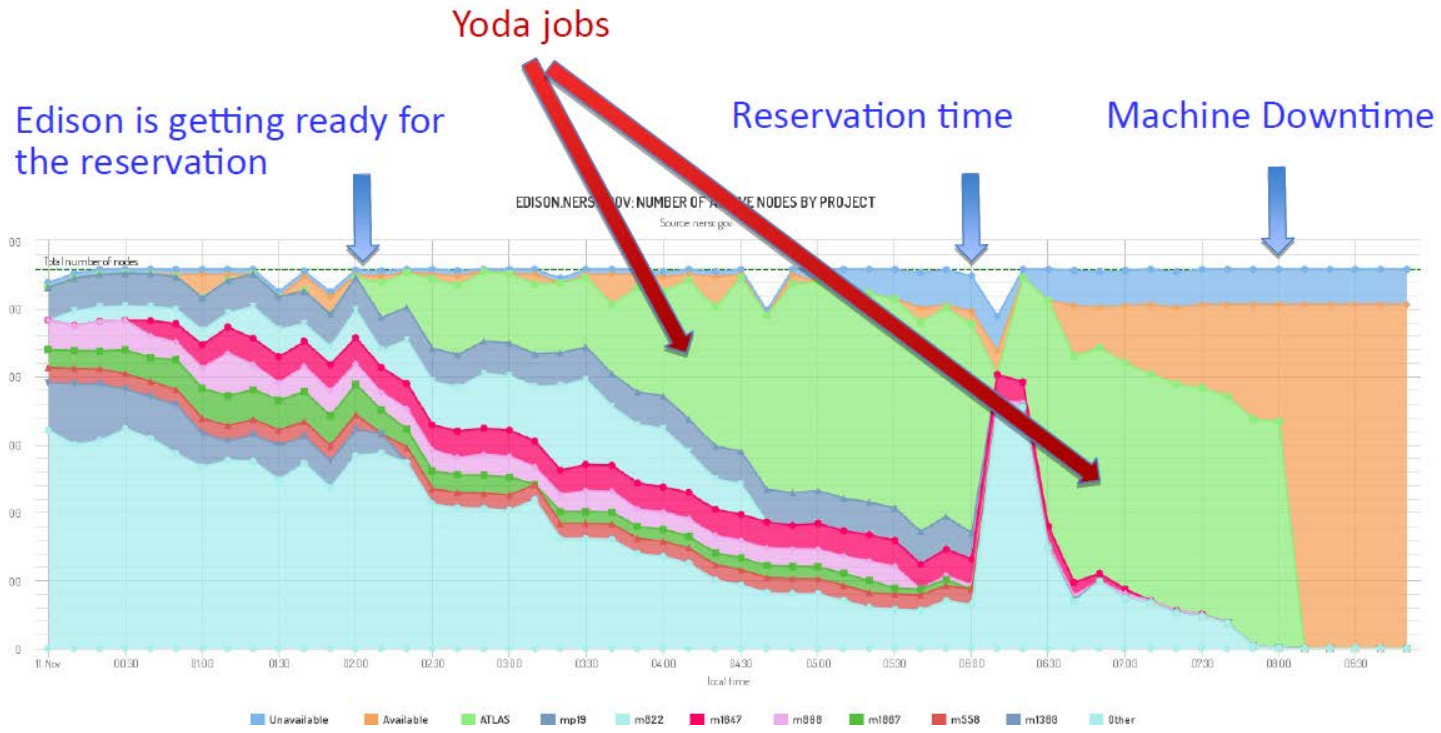

Big Push at ALCF

Argonne NATIONAL LABORATORY

Update from Argonne

Taylor Childers (Argonne)

Tom LeCompte (Argonne)

Tom Uram (Argonne)

Doug Benjamin (Duke)

U.S. DEPARTMENT OF ENERGY

Large Allocation

2014 ALCC Award

Also requested 50M CPU Hours in 2015

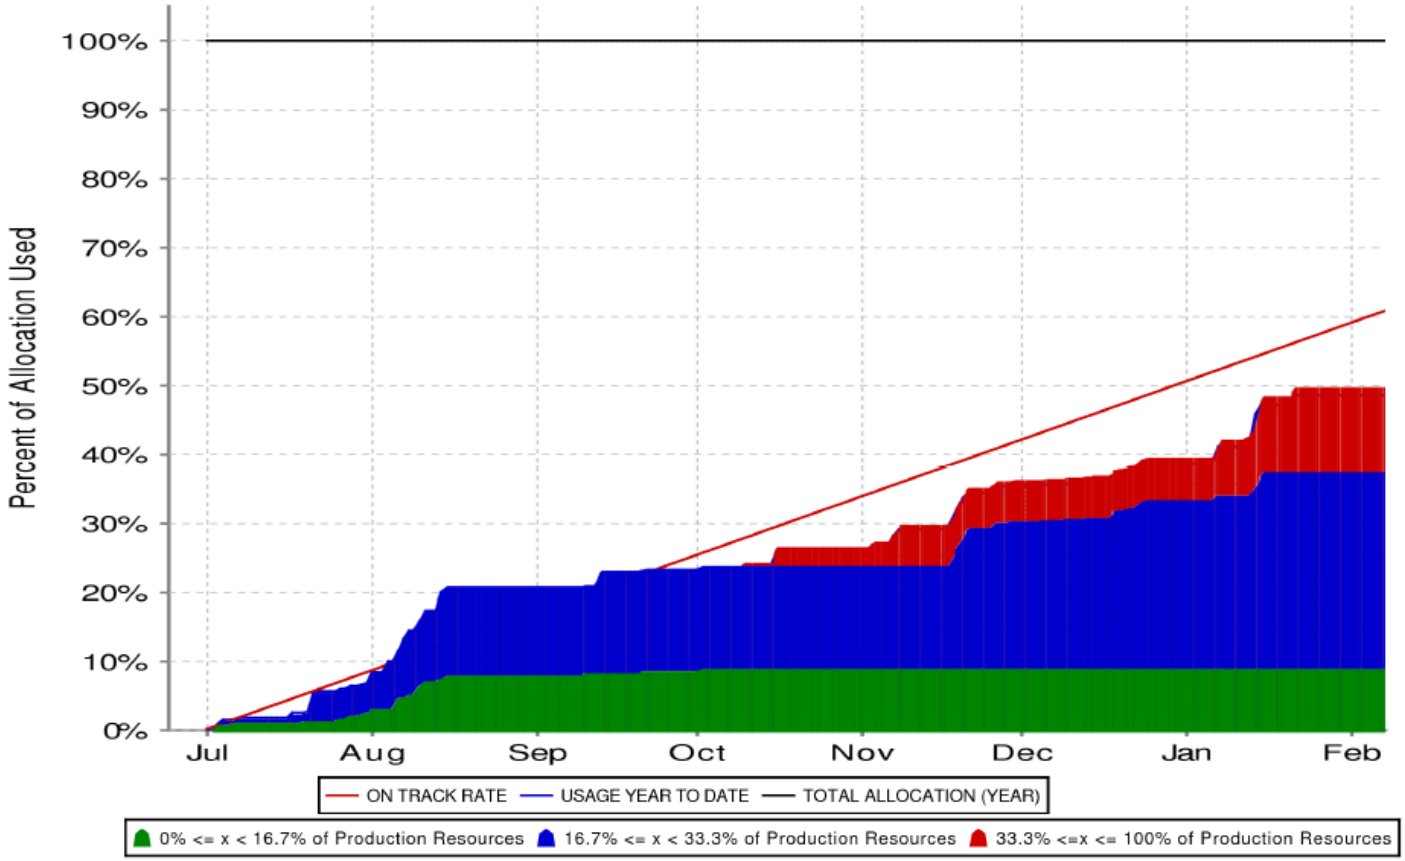

PanDA at ALCF

ARGO Integration with PanDA

ARGO submits each step as a job to Balsam.

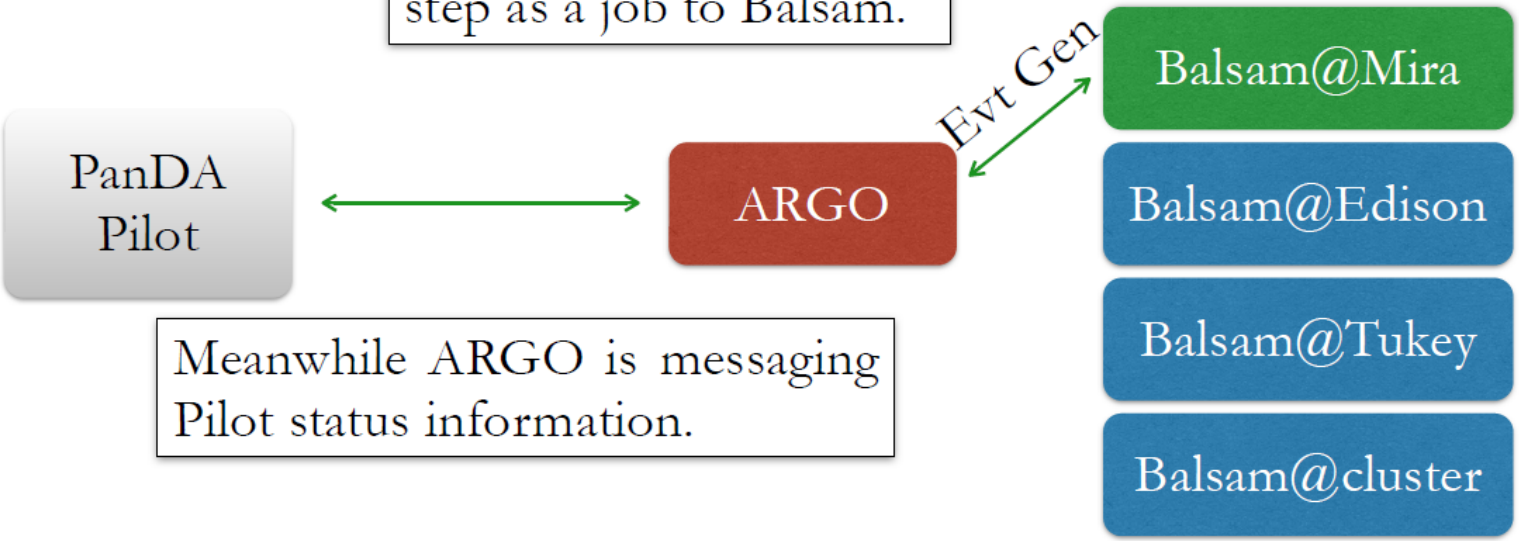

Meanwhile ARGO is messaging Pilot status information.

For Example:
1. Generator Integration Step
2. Generator Event Generation Step

Impressive – Full Throttle

Learning to grow

Argonne NATIONAL LABORATORY
Leadership Computing Facility

Mira Activity

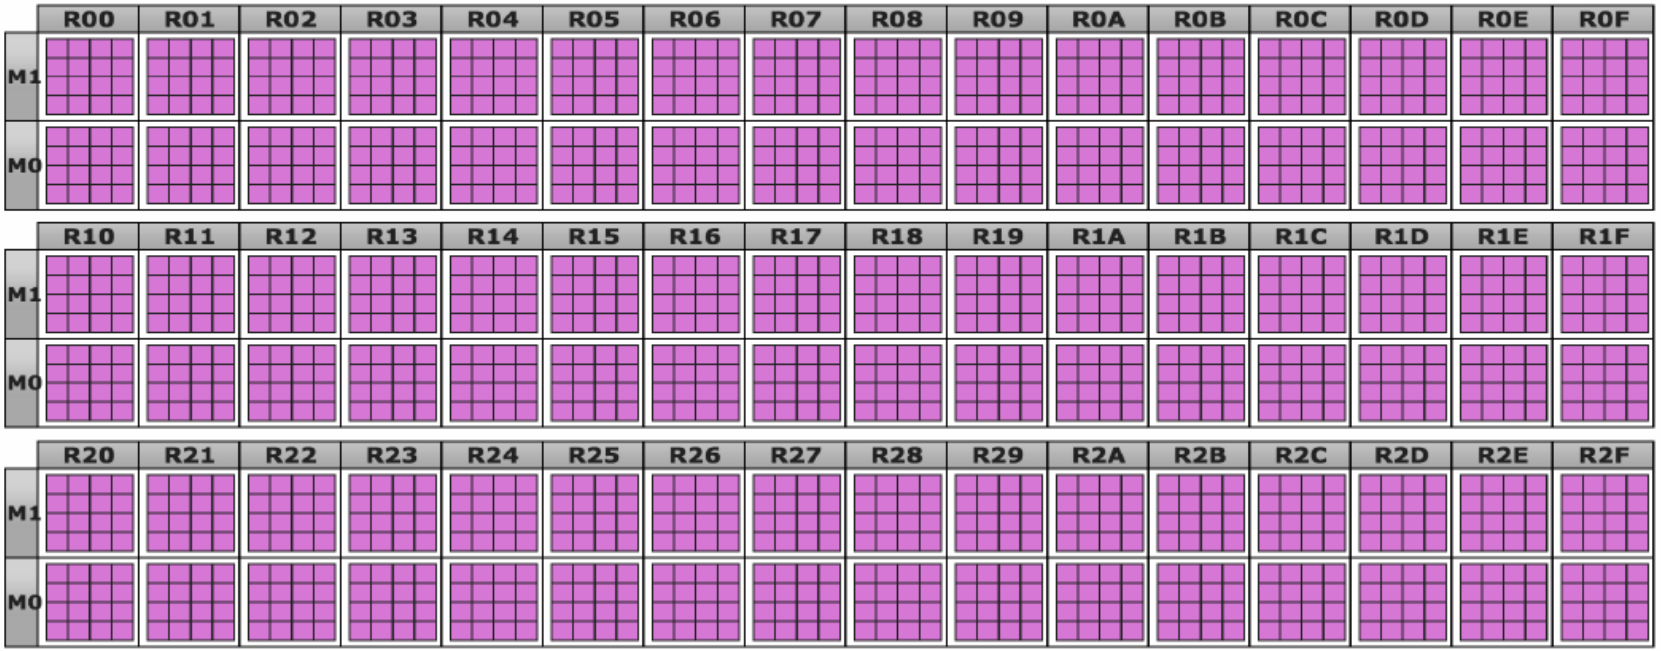

- Here we are running Alpgen in a Mira-sized (49,152 node) partition (32 ranks per node)
- 1,572,864 ranks

HPC's in Europe

Some HPC Resources in Europe

SuperMUC
LRZ, Germany

Hydra
RZG, Germany

ARCHER
RCUK, UK

Piz Daint
CSCS, Switzerland

Using Nordugrid ARC with PanDA

ARC-CE remote job submission (3)

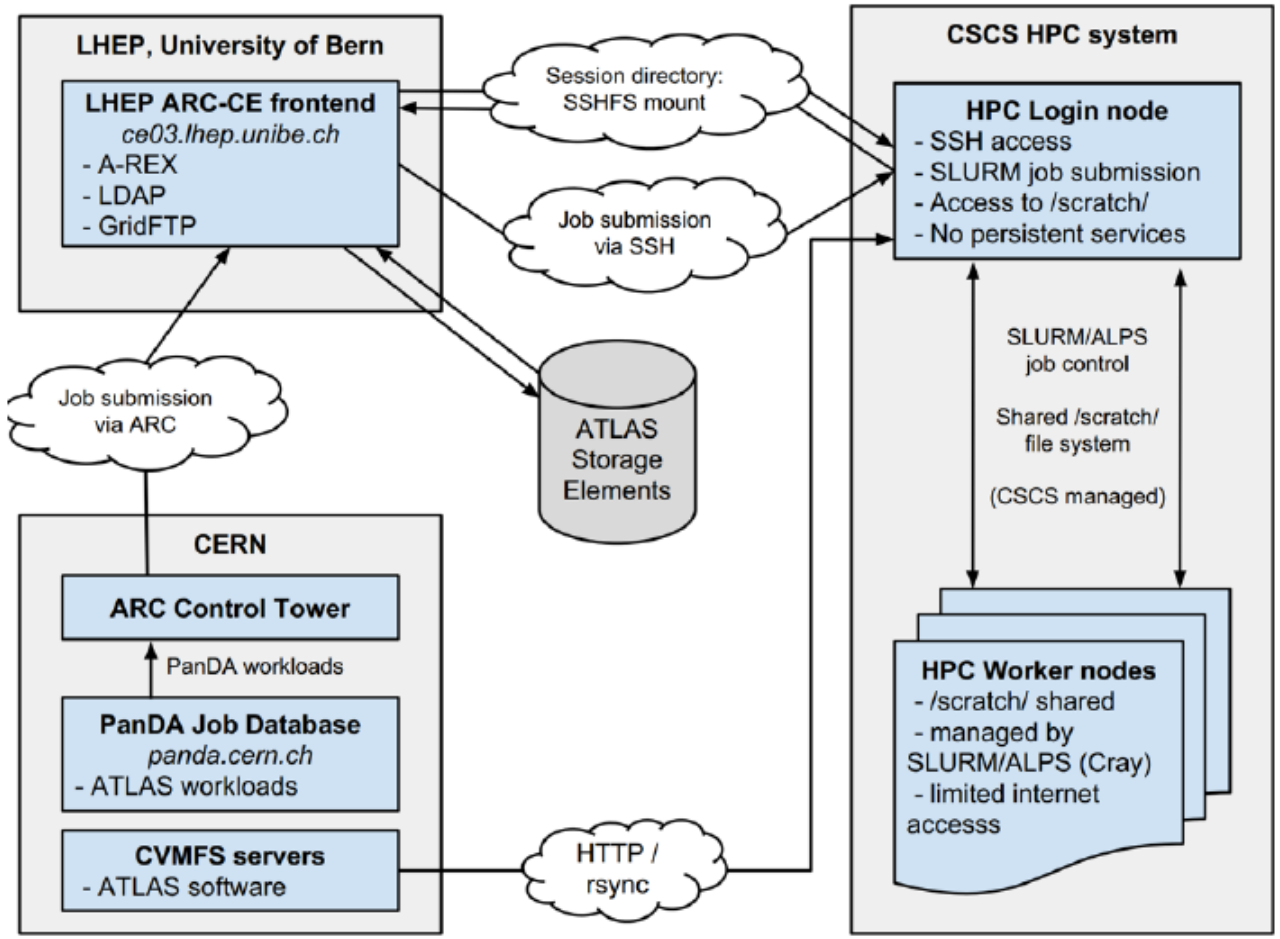

Utilizing HPC CPU's with ARC

Current Status: 3 HPC sites running

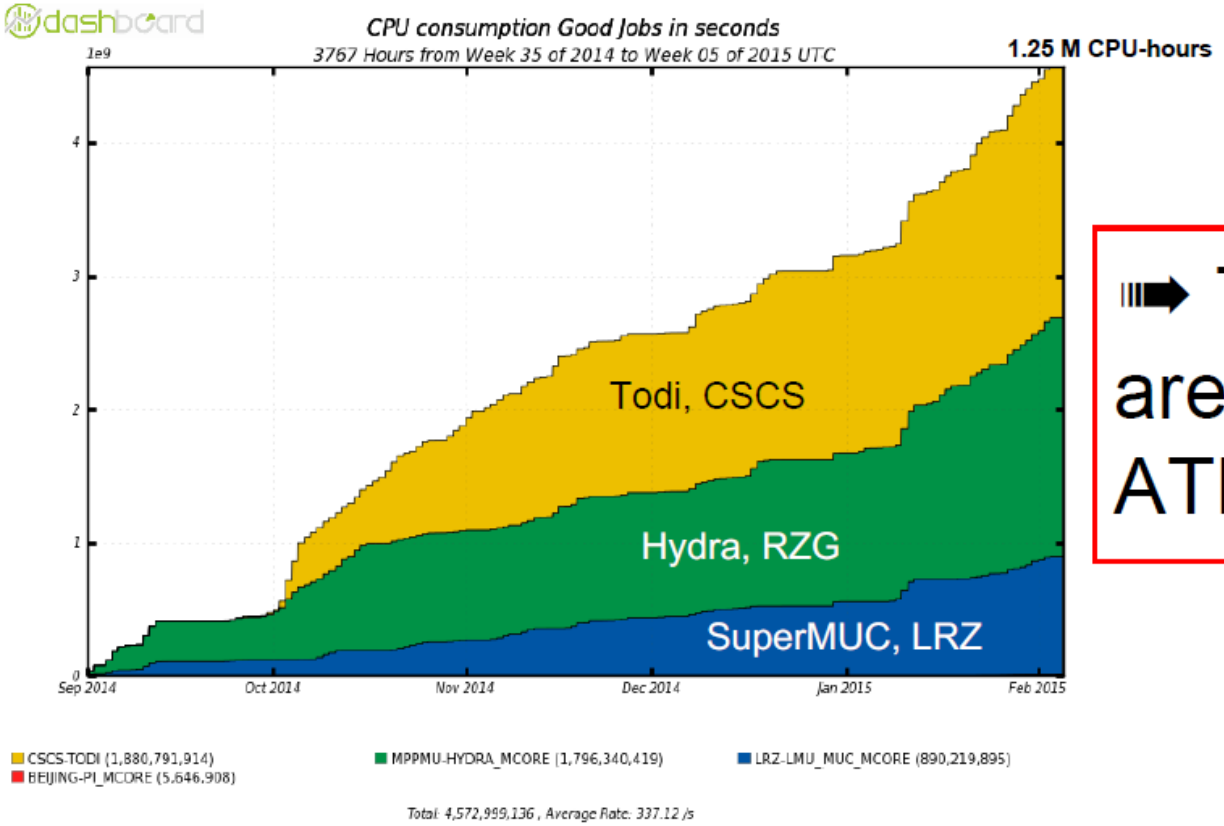

➠ These systems are up and running ATLAS production

Outlook and Plans

- **Switzerland**: Requested 50M CPU-hours on Piz Daint
  - Opportunistic backfill still an option
  - Runs on Todi show feasibility to CSCS
- **Germany**: Hydra, SuperMUC: continuous backfill
- **UK**: ARCHER: integration ongoing
- **China**: Pi and more?
  - Integration of Pi ongoing
  - Development of an ARC-CE 'backend' to Chinese 'SCGrid' (15 HPCs)

- Further software development:
  - SSH submission backend in ARC
  - Automated software provisioning, validation and tagging

The Growing PanDA EcoSystem

- ATLAS PanDA core
  - US ATLAS, CERN, UK, DE, ND, CA, Russia, OSG …
- ASCR/HEP BigPanDA
  - DoE funded project at BNL, UTA – PanDA beyond HEP, at LCF
- ANSE PanDA
  - NSF funded network project - CalTech, Michigan, Vanderbilt, UTA
- HPC and Cloud PanDA – very active
- Taiwan PanDA – AMS and other communities
- Russian NRC KI PanDA, JINR PanDA – new communities
- AliEn PanDA, LSST PanDA, other experiments
- MegaPanDA (COMPASS, ALICE, NICA) …

Resources Accessible via PanDA

WLCG
Worldwide LHC Computing Grid

NORDUGRID
Grid Solution for Wide Area Computing and Data Handling

Open Science Grid

Many Others

nectar

NERSC

НАЦИОНАЛЬНЫЙ ИССЛЕДОВАТЕЛЬСКИЙ ЦЕНТР «КУРЧАТОВСКИЙ ИНСТИТУТ»

amazon web services™ EC2

OLCF
Oak Ridge Leadership Computing Facility

FutureGrid
an XSEDE resource provider

Google Cloud Platform

About 150,000 job slots used continuously 24x7x365