PanDA Beyond ATLAS: Workload Management for Data Intensive Science

J. Schovancová, K. De, A. Klimentov,
T. Maeno, P. Nilsson, D. Oleynik, S. Panitkin,
A. Petrosyan, A. Vaniachine, T. Wenaus, D. Yu
on behalf of the ATLAS collaboration
Brookhaven National Laboratory, USA
University of Texas at Arlington, USA
Argonne National Laboratory, USA

HEPiX Fall 2013

BROOKHAVEN
NATIONAL LABORATORY

Introduction

- **P**roduction **an**d **D**istributed **A**nalysis System
- Highly automated, low operational manpower, integrated monitoring system
- PanDA history
  - Aug 2005: Project started
  - Sep 2005: The first prototype
  - Dec 2005: Production in US ATLAS
  - 2008: Adopted as the workload management system for the entire ATLAS collaboration
  - 2012-2013: AMS and CMS have deployed their own PanDA instances
  - 2013: LSST and ALICE also investigating PanDA integration
- Great performance during LHC Run 1
  - Data processing, simulation and analysis
  - Actively evolving to meet rapidly changing physics needs
  - Successfully managing >130 sites, ~10^5 cores, ~10^8 jobs per year, ~1500 users

PanDA System

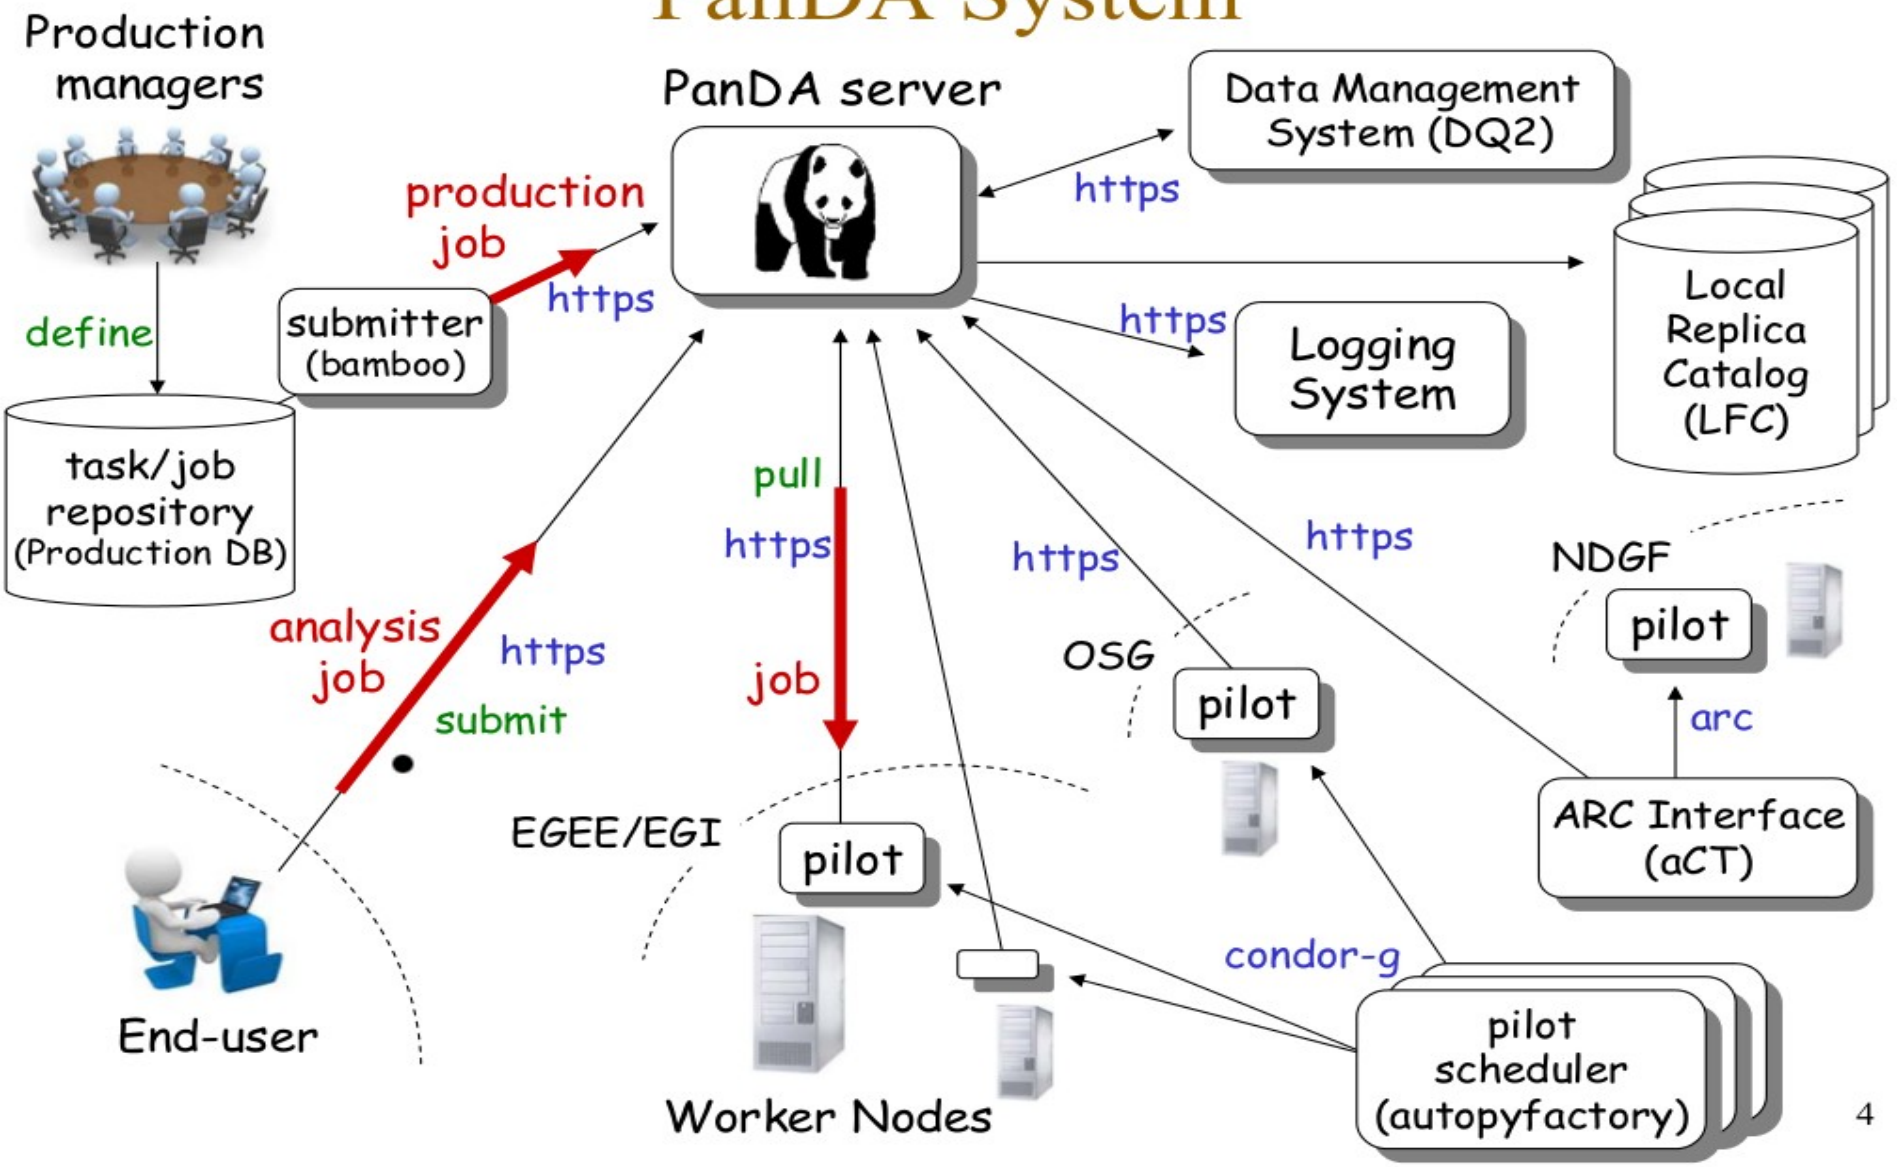

4

T. Maeno et al.: *Evolution of the ATLAS PanDA Workload Management System for Exascale Computational Science*, CHEP2013

Completed ATLAS jobs Jan 2010 – Oct 2013

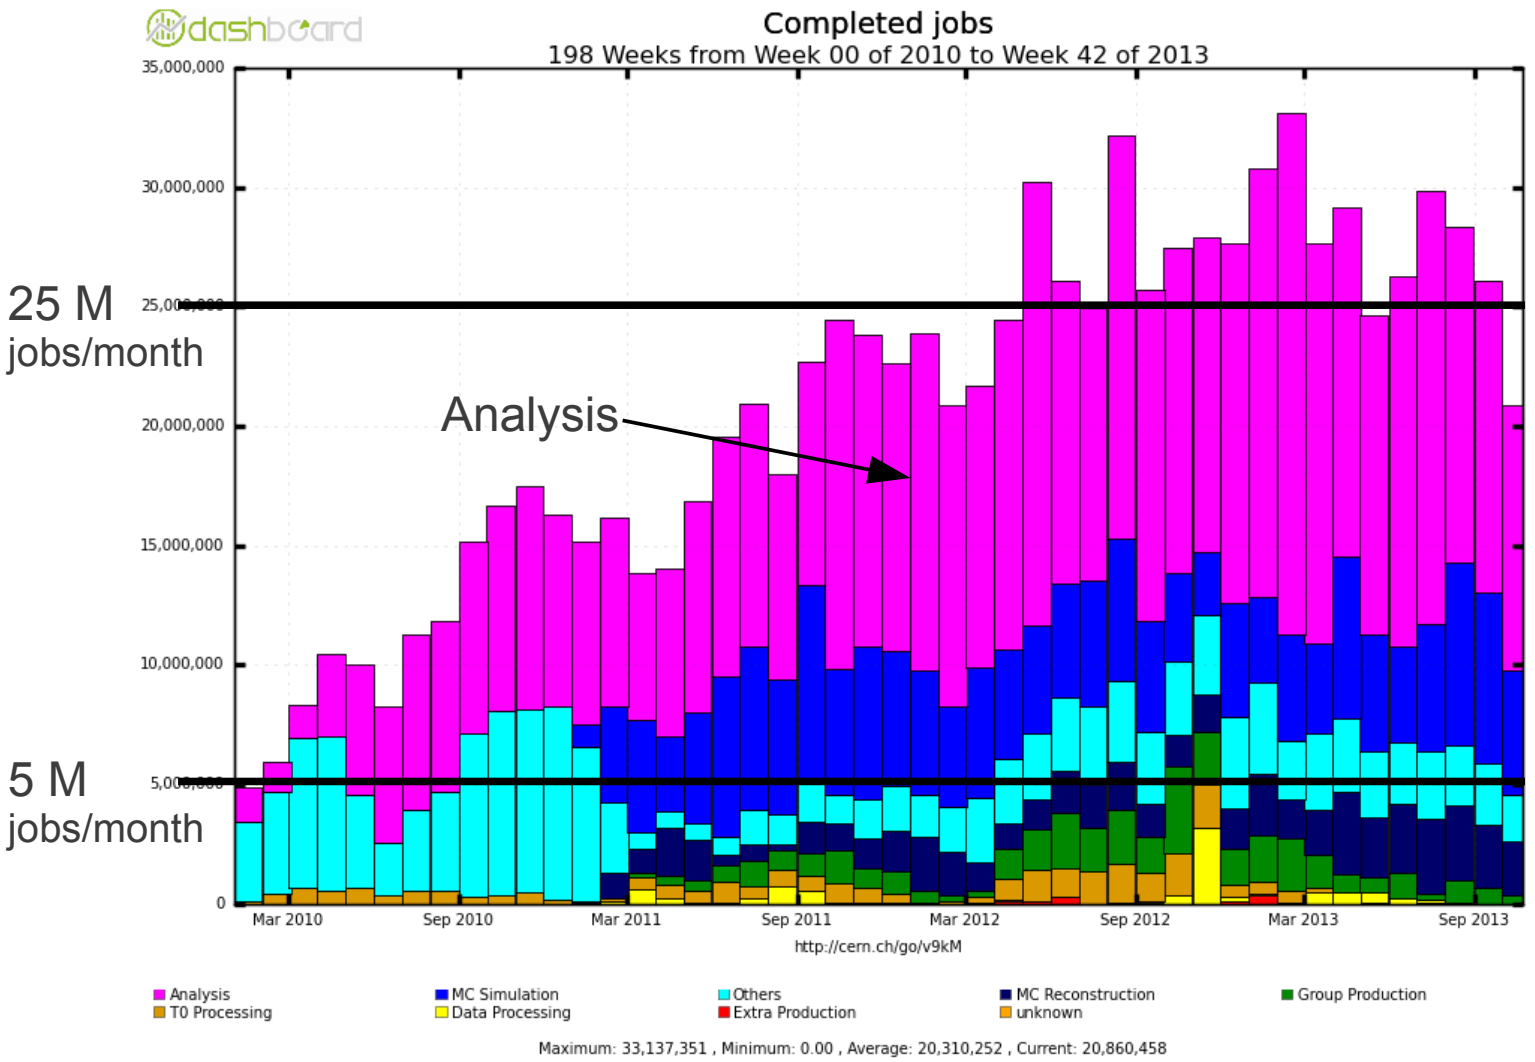

BigPanDA

- Evolution of PanDA for Advanced Scientific Computing
  - interest in PanDA by other big data sciences
- Proposal **Next Generation Workload Management and Analysis System for BigData** submitted to ASCR DOE in April 2012
  - Generalize PanDA as meta application, providing location transparency of processing and data management for HEP and other data-intensive sciences, and a wider exascale community
  - DOE ASCR and HEP funded the project since Sep 2012
  - 3 dimensions:
    - ➔ Make PanDA available beyond ATLAS and HEP
    - ➔ Extend beyond Grid: Leadership Computing Facilities, Clouds, University clusters
    - ➔ Integrate network as a resource in the workload management

BigPanDA Work Plan

- 3 years plan
  - Year 1: Set up the collaboration, define algorithms and metrics
  - Year 2: Prototype and implement
  - Year 3: Production and operations
- 4 work packages
  - WP1: Factorizing the core
  - WP2: Extending the scope
  - WP3: Leveraging intelligent networks
  - WP4: Usability and monitoring

WP1: Factorizing the core

- Factorize the core components of PanDA to enable adoption by wide range of data-intensive scientific communities
- Package core components of PanDA in an experiment-neutral package
  - General components, customizable layers
  - Experiment-specific layers in plugins and configuration files
- Provide advanced features with sensible defaults

WP2: Extending the scope

- Evolve PanDA to support extreme scale computing clouds and Leadership Computing Facilities
- Adding extra resources
- Expansion to the resources available to potential user community

WP3: Leveraging intelligent networks

- Research efforts in dynamic network provisioning, quality of service and traffic management
- Integration of network services and real-time data access to the PanDA workflow
  - Integrate the services within existing and evolving infrastructures
  - Automated discovery and usage of such services transparently to the scientists

WP4: Usability and monitoring

- Special effort to factorize and generalize PanDA monitoring
- Design a generic monitor skeleton
  - Experiment-specific views as customization
- Provide generic components and APIs
  - For user communities to easily implement and customize their monitoring views and workflow visualizations
- Documentation, tutorials

PanDA server and monitor

PanDA server

- Multiple DB backend support
  - Oracle and MySQL
- VO–independent instance on Amazon EC2
  - MySQL DB backend
  - Serve the non-LHC experiments in OSG
- Refactoring to decompose experiment-specific code to plugins is well underway

PanDA monitor

- Refactoring to separate common skeleton and experiment-specific plugins

Pilot

- Already being refactored in context of the Common Analysis Framework project
- Experiment-specific plugins
  - payload configuration
  - input/output data management
  - site configuration information
- Incremental updates to avoid affecting current production

P. Nilsson et al.: *Next Generation PanDA Pilot for ATLAS and Other Experiments*, CHEP2013

Use of Google Compute Engine in PanDA

- Common ATLAS project with Google
  - ~5M CPU hours, 4000 cores for ~2 months, processed 457k jobs
- Resources organized as HTCondor-based PanDA queue
- Transparent inclusion of cloud resources into ATLAS grid
- Delivered to ATLAS as a **production resource**, not as an R&D platform

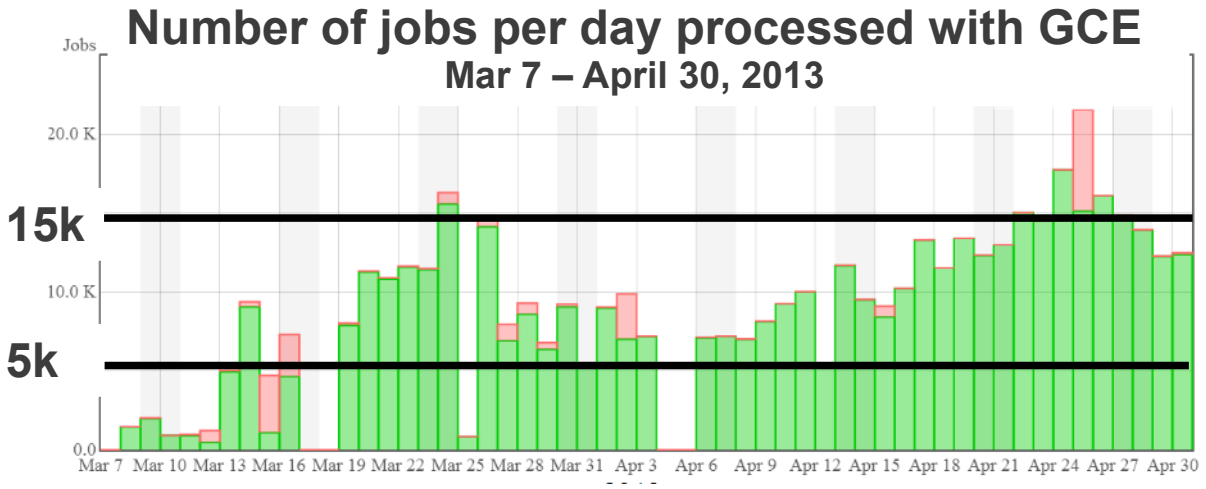

S. Panitkin et al.: *ATLAS Cloud Computing R&D*, CHEP2013

PanDA and HPC

- In collaboration with Oak Ridge National Laboratory Leadership Computing Facility
  - Gain experience with all aspects of platform and workload
    - Job submission, output storage and transfers, monitoring, security, etc.
  - Develop pilot/agent model for Titan
  - Focus: ATLAS Geant4 simulation

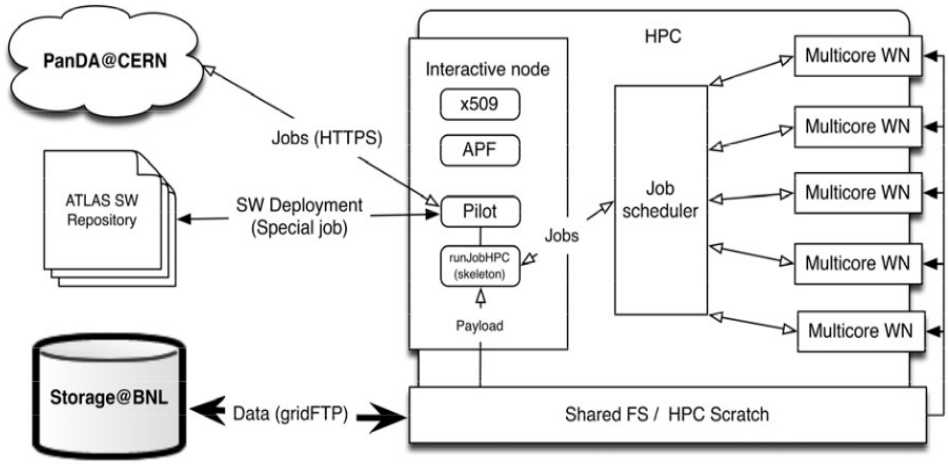

Leveraging Intelligent Network

- Synchronized with two other efforts
  - Integration of Federated XrootD for ATLAS (FAX) with PanDA, and ANSE project
- Three layered SW architecture
  - Collector: network performance information from various sources: perfSonar, ATLAS Site Status Board, FAX, etc.
  - AGIS (ATLAS Grid Information System): information pool
  - Calculator: weights to be taken into account for site selection by brokerage
- PanDA brokerage site selection algorithm being enriched to consider FAX cost matrix

PanDA at ASGC, TW

Other Applications Support using PanDA

- Application Support in various research groups as below
  - Other HEP research groups (App: GMC(geant4))
  - Polymer Physics & Biomacromolecule physics
    - Computer simulations, Serial computing & batch submission
  - Earth science, Climate changes
    - MPI & OpenMP (App: gemb(x))

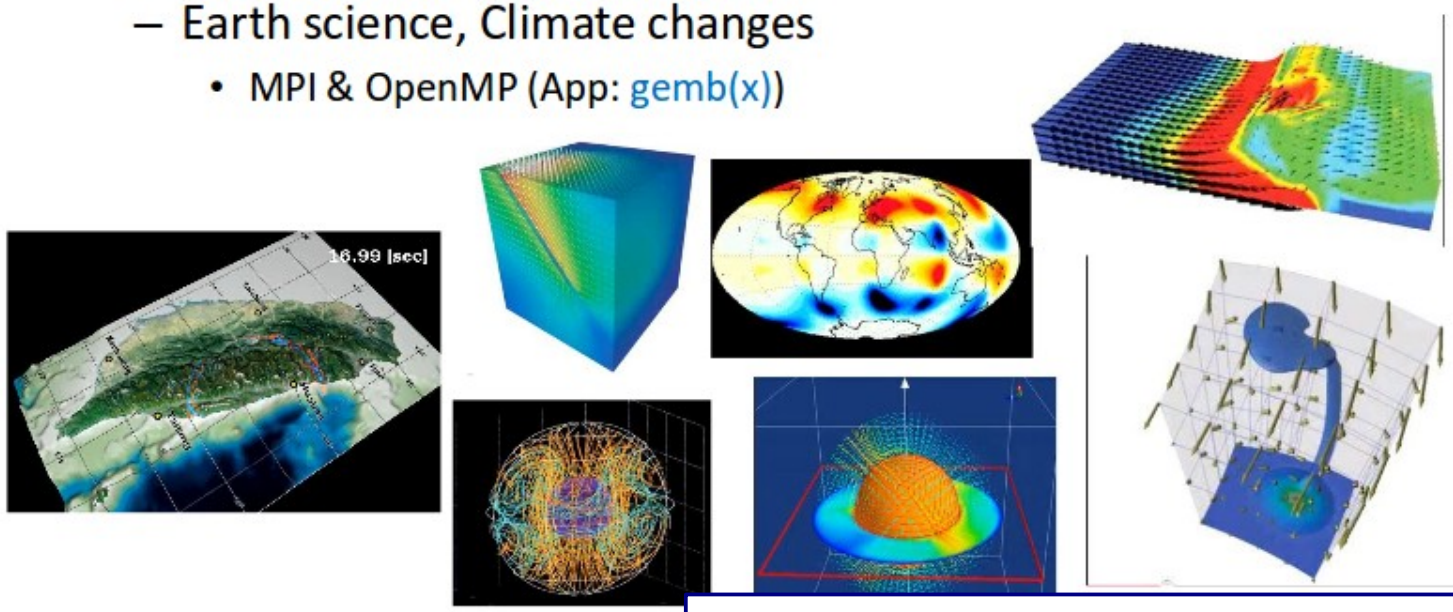

J. You: *Status report of Taiwan PanDA*,
BigPanDA workshop at CERN, 21st Oct 2013

Summary

- ✔ **The PanDA system played a key role during LHC Run 1 data processing, simulation and analysis with a great success**
  - ➢ **While actively evolving to meet rapidly changing physics needs!**
- ✔ **Interest in PanDA by other data-intensive sciences motivates generalization of the PanDA system**
- ✔ **ASCR DOE gave us a great opportunity to evolve PanDA beyond ATLAS and HEP and to start the BigPanDA project**
- ✔ **Progress in many areas: networking, VO independent PanDA instance, cloud computing, HPC**
- ✔ **Strong interest in the project from several experiments and scientific centers to have a collaborative project**