Next Generation Workload Management System For Big Data on Heterogeneous Distributed Computing

Next Generation Workload Management System For Big Data on Heterogeneous Distributed Computing

A Klimentov[1], P Buncic[2], K De[3], S Jha[4], T Maeno[1], R Mount[5], P.Nilsson[1], D Oleynik[3], S Panitkin[1], A Petrosyan[3], R J Porter[6], K F Read[7], A Vaniachine[8], J C Wells[7] and T Wenaus[1]

[1] Brookhaven National Laboratory, Upton, NY, USA
[2] CERN, Geneva, Switzerland
[3] University Texas at Arlington, Arlington, TX, USA
[4] Rutgers University, Piscataway, NJ, USA
[5] SLAC National Accelerator Laboratory, Menlo Park, CA, USA
[6] Lawrence Berkeley National Laboratory, CA, USA
[7] Oak Ridge National Laboratory, Oak Ridge, TN, USA
[8] Argonne National Laboratory Lemont, IL, USA

E-mail: aak@bnl.gov

Abstract. The Large Hadron Collider (LHC), operating at the international CERN Laboratory in Geneva, Switzerland, is leading Big Data driven scientific explorations. Experiments at the LHC explore the fundamental nature of matter and the basic forces that shape our universe, and were recently credited for the discovery of a Higgs boson. ATLAS and ALICE are the largest collaborations ever assembled in the sciences and are at the forefront of research at the LHC. To address an unprecedented multi-petabyte data processing challenge, both experiments rely on a heterogeneous distributed computational infrastructure. The ATLAS experiment uses PanDA (Production and Data Analysis) Workload Management System (WMS) for managing the workflow for all data processing on hundreds of data centers. Through PanDA, ATLAS physicists see a single computing facility that enables rapid scientific breakthroughs for the experiment, even though the data centers are physically scattered all over the world. The scale is demonstrated by the following numbers: PanDA manages $O(10^2)$ sites, $O(10^5)$ cores, $O(10^8)$ jobs per year, $O(10^3)$ users, and ATLAS data volume is $O(10^{17})$ bytes. In 2013 we started an ambitious program to expand PanDA to all available computing resources, including opportunistic use of commercial and academic clouds and Leadership Computing Facilities (LCF). The project titled 'Next Generation Workload Management and Analysis System for Big Data' (BigPanDA) is funded by DOE ASCR and HEP. Extending PanDA to clouds and LCF presents new challenges in managing heterogeneity and supporting workflow. The BigPanDA project is underway to setup and tailor PanDA at the Oak Ridge Leadership Computing Facility (OLCF) and at the National Research Center "Kurchatov Institute" together with ALICE distributed computing and ORNL computing professionals. Our approach to integration of HPC platforms at the OLCF and elsewhere is to reuse, as much as possible, existing components of the PanDA system. We will present our current accomplishments with running the PanDA WMS at OLCF and other supercomputers and demonstrate our ability to use PanDA as a portal independent of the computing facilities infrastructure for High Energy and Nuclear Physics as well as other data-intensive science applications.

1. Introduction

The largest scientific instrument in the world, the Large Hadron Collider, is operating at the CERN Laboratory in Geneva, Switzerland [1]. The ATLAS [2], ALICE [3] and other LHC experiments explore the fundamental nature of matter and the basic forces that shape our universe. To address an unprecedented multi-petabyte data processing challenge, the LHC experiments rely on the computational grids infrastructure deployed in the framework of the Worldwide LHC Computing Grid (WLCG) [4]. WLCG is by far the largest academic distributed computing environment in the world and the levels of use are ground-breaking. Thanks to the outstanding LHC performance in 2010-2013, ATLAS manages over 160 petabytes of data, submitting up to two million jobs per day on more than 100 sites (figure 1). Following the massive data processing on the Grid, more than 8000 scientists analyze LHC data in search of new discoveries. ATLAS leads the WLCG usage in the number of jobs, processed data volume, and in core-hours.
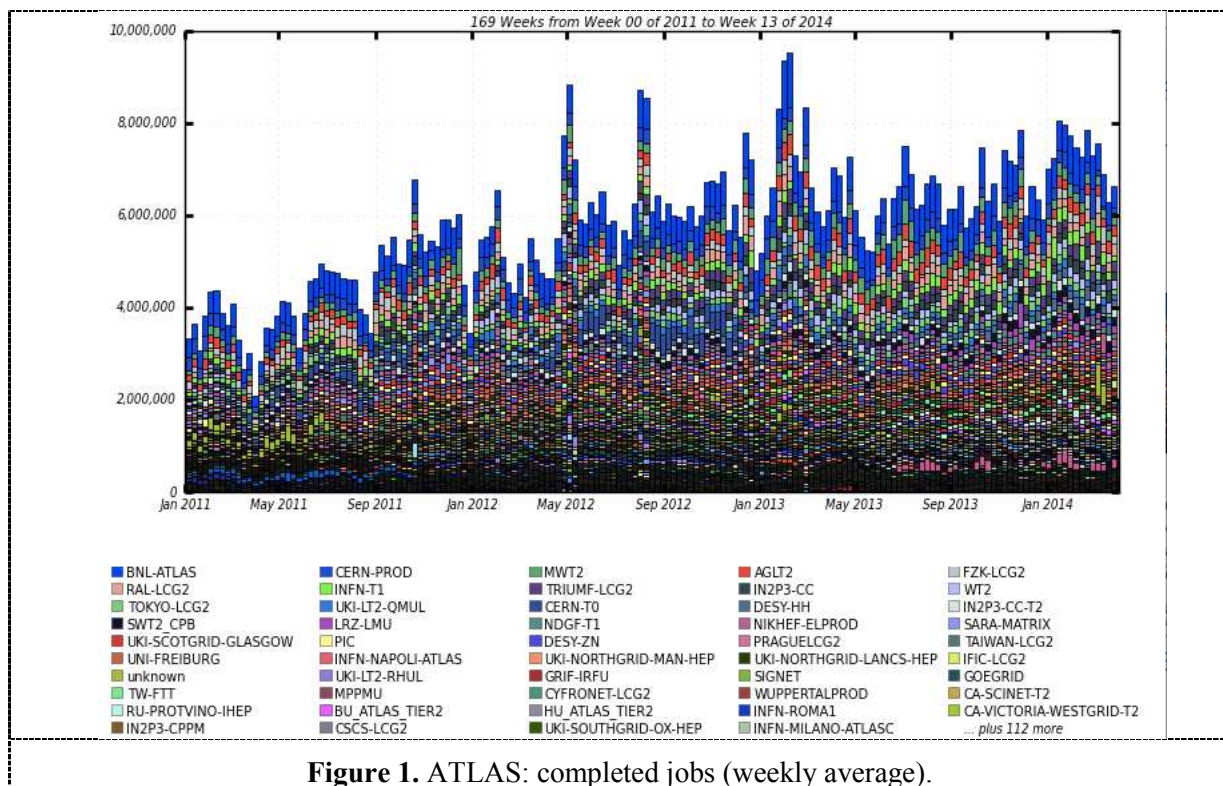

Figure 1. ATLAS: completed jobs (weekly average).

The challenges posed by such mega-science experiments are numerous and not limited to the unprecedented size of the data. Exascale data is often highly distributed and accessed by large international collaborations. A sophisticated Workload Management System (WMS) is needed to manage the distribution and processing of such data. One of the most successful WMS developed in the U.S. is PanDA (acronym for Production and Distributed Analysis System) [5], used by thousands of physicists in the ATLAS experiment at the Large Hadron Collider.

PanDA delivers transparency of data and processing in a distributed computing environment to ATLAS physicists. It provides execution environments for a wide range of experimental applications, automates centralized data production and processing, enables analysis activity of physics groups, supports custom workflow of individual physicists, provides a unified view of distributed worldwide resources, presents status and history of workflow through an integrated monitoring system, archives and curates all workflow, manages distribution of data as needed for processing or physicist access. The rich menu of features provided, coupled with support for heterogeneous computing environments, makes PanDA ideally suited for data intensive sciences.

PanDA has a highly scalable and flexible architecture. Scalability has been demonstrated in ATLAS through the rapid increase in usage over the past three years. PanDA was designed to have the flexibility to adapt to emerging computing technologies in processing, storage, networking as well as the underlying software stack (middleware and file management). This flexibility has also been successfully demonstrated through the past five years of evolving technologies adapted by computing centres in ATLAS, which span many continents and yet are seamlessly integrated into PanDA. This proven scalability and flexibility makes PanDA ideally suited for adoption by future megascience projects.

PanDA was conceived in 2005. At the time, a variety of workload management systems were deployed in ATLAS, separately for different applications, using distinct systems for the chaotic workflows from physicists and the organized workflows from central production. PanDA emerged as the best system and was adopted as the default and single system for ATLAS before the LHC started operating in 2009. Today, PanDA has grown to support all distributed workflows in ATLAS, and enjoys a huge user base with worldwide support.

PanDA was adapted and successfully used for molecular dynamics simulations of protein folding using the CHARMM [6] molecular modeling software, leading to a publication describing the implementation as general, flexible, easily modifiable for use with other molecular dynamics programs and other grids and automated in terms of job submission, monitoring, and resubmission. In 2012, the Alpha Magnetic Spectrometer (AMS) experiment [7] began using PanDA for Monte Carlo simulation and data processing. AMS set up the PanDA infrastructure at the largest academic computing center of Asia-Pacific region (ASGC). The Large Synoptic Survey Telescope collaboration [8] has conducted the first Monte Carlo simulation in 2013-2014 using a PanDA instance installed on the Amazon EC2 cloud.

Interest in PanDA by other big data sciences provided the primary motivation to make a proposal titled "Next Generation Workload Management and Analysis System for Big Data." The idea was to generalize PanDA as meta-application, providing location transparency of processing and data management, for High Energy Physics (HEP) community and other data-intensive sciences, and the wider exascale community. The DOE Advanced Scientific Computing Research office (ASCR) awarded a grant to expand PanDA beyond HEP, to add network awareness to PanDA, and to expand PanDA to Leadership Computing Facilities and cloud computing (BigPanDA) [9].

2. Next Generation Workload Management and Analysis System for Big Data

The project was started at the end of 2013 and is jointly funded by the DOE ASCR and HEP offices. The work proposed within the project will enable the use of PanDA by new scientific collaborations and communities as a means of leveraging extreme scale computing resources with a low barrier of entry. Particular attention is given to enabling the computing infrastructures supported by the DOE that were not initially supported by PanDA such as the Leadership Computing Facilities (LCF).

Support for Leadership Computing Facilities will expand the potential user community for PanDA and will also, even in the near term, benefit ATLAS and ALICE by running jobs on Leadership Computing Facilities with direct access to the data hosted by US National Laboratories, dynamically acquiring supplementary CPU resources when needed. Extending PanDA beyond the Grid and integrating distributed computing with cloud computing (or grids of clouds) will further expand the potential user community and the resources available to them, for example making it possible for non-LHC and non-HEP experiments such as AMS and LSST to use PanDA.

In the longer term, exascale facilities and cloud platforms could become a real alternative to very large data centers owned and managed by the scientific community. This work would create in PanDA a WMS system enabling users to bridge the near and longer terms smoothly, with minimal effort, and preserve the same global system view throughout, by leveraging PanDA support for both present and emerging computing infrastructures. ATLAS and ALICE would certainly lead the way in utilizing exascale platforms this way, and we would seek through this work to make it an easy path for others to follow as well.

There are three dimensions for the overall workload system evolution. First, we should make PanDA available beyond the LHC and HEP. Second, we should extend PanDA beyond the Grid. Third, we should integrate network services as a resource in workload management. The following four work packages have been identified:

- *Factorizing the core:* Factorizing the core components of PanDA to enable adoption by a wide range of exascale scientific communities.
- *Extending the scope:* Evolving PanDA to support extreme scale computing clouds and Leadership Computing Facilities.
- *Leveraging intelligent networks:* Integrating network services and real-time data access to the PanDA workflow.
- *Usability and monitoring:* Real time monitoring and visualization package for PanDA

We discuss below in more detail how to evolve PanDA to support computing clouds and HPC.

3. BigPanDA to support extreme scale computing clouds and Leadership Computing Facilities

3.1. Scale of Needs
The ATLAS experiment uses a geographically distributed grid of approximately 130,000 cores continuously, (over 1000 million core-hours per year) to simulate and analyze its data. After the early success in discovering a new particle consistent with the long awaited Higgs boson [10], ATLAS is preparing for precision measurements and further discoveries that will be made possible by much higher LHC collision rates after early 2015. The need for simulation and analysis would overwhelm the expected capacity of WLCG computing facilities unless the range and precision of physics studies were to be curtailed.

From the usage information above, it becomes clear that HPC contributions of the order of 10 million or more core hours per year become important and valuable. ATLAS computing can also be a close-to-ideal "crack-filling" application. The ATLAS production management system is being upgraded to make it aware of dynamically changing resources, and thus able to exploit groups of processors that become available for relatively short times. Storage capacity on the petabyte scale will be an objective of the work package, adding extra resources such as computing clouds and Leadership Computing Facilities to those supported by PanDA. Extending PanDA beyond the Grid will further expand the potential user community and the resources available to them.

3.2. Challenges
The major issues and challenges, which must be addressed to attain excellent scientific productivity on advanced computer architectures, are described below.

3.2.1. Execution efficiency. In common with most of HEP, ATLAS code typically executes around 0.8 instructions per clock cycle, even though the (Intel, AMD) hardware it executes on can perform more than 10 instructions/cycle and achieves between 2 and 3 instructions per cycle for much scientific code. The reasons for this "bad" performance are the many conditional branches and lack of repetitious activities, both due to the complex nature of the physics and devices being simulated or analyzed.

3.2.2. Parallelism. In its current form, ATLAS simulation and analysis is highly amenable to trivial high-level parallelism: events (i.e. an initial collision and the subsequent response of the detector to all the collision products) can be independently simulated or analyzed on thousands of cores. Current ATLAS simulation and analysis makes heavy use of this parallelism. ATLAS codes also have the potential for massive low-level parallelism that may be a good match for vector and GPU hardware. To expose this parallelism, the execution architecture must be transformed to present many similar (but unrelated) operations for simultaneous execution. The level of parallelism that can present "close-to-SIMD work packages" for efficient execution on many-core CPUs or GPUs will benefit from the

high bandwidth communications network, typical of an HPC facility, to assemble work and distribute the results.

3.2.3. Production environment. Major HEP simulation and analysis, such as that of ATLAS, requires sophisticated systems to control the flow of data and tasks to computing resources, to monitor workflow and to assure automated recovery from most "environmental" errors. To make productive use of HPC, interfacing to such production environments is essential.

3.2.4. Living code. There are 6 million lines of active code in the ATLAS SVN repository. This is living code, being continuously improved as our understanding of the detector and of physics develops. The code must remain intelligible and maintainable. While small amounts of code can certainly be optimized for particular target architectures, the majority of the code must remain generic.
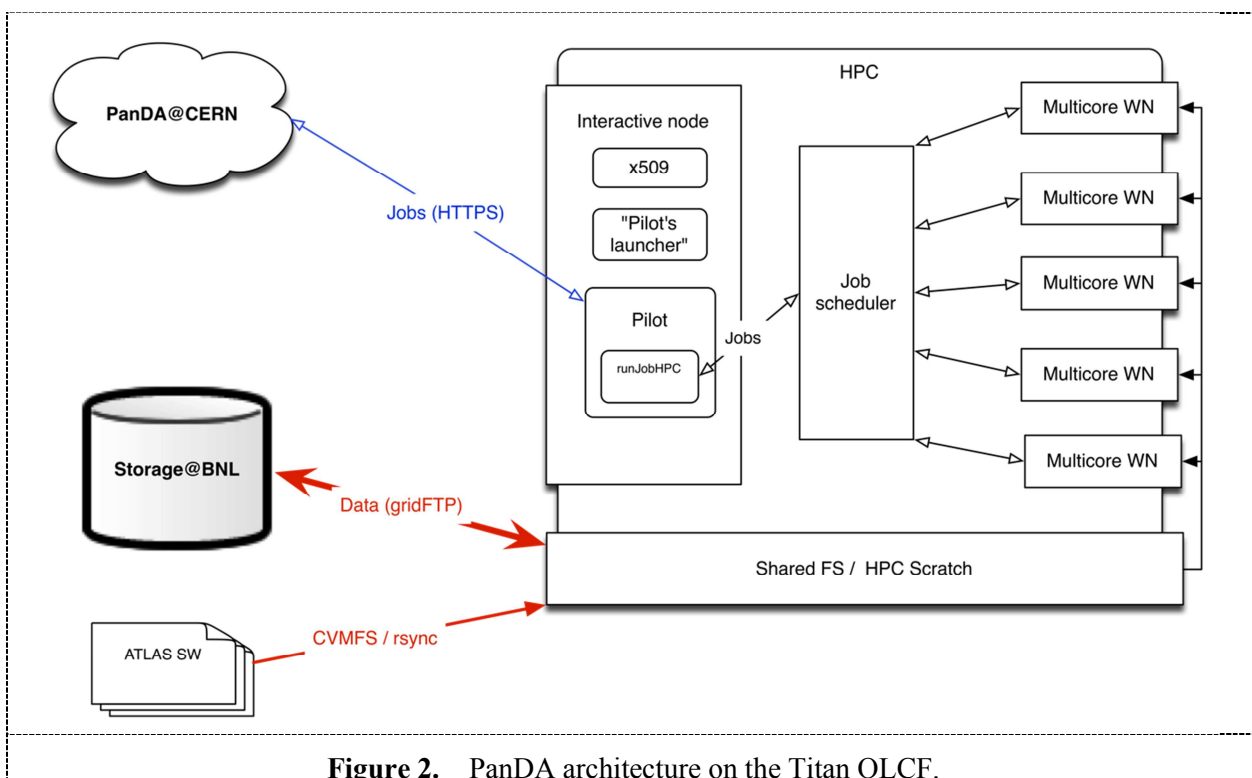

Figure 2.   PanDA architecture on the Titan OLCF.

3.3. Extending PanDA to Leadership Computing Facilities

Unlike the Grid, HPC centers are primarily built to execute parallel payloads and there are usually significant limitations: worker node set up is fixed, outbound network connections are not possible from the worker nodes, limited RAM per core, and a specialized Operating System. Standard Workload Management Systems (like PanDA) are not well-adapted to submit payloads to HPC resources. Together with colleagues at the OLCF, we adapted PanDA for the Titan LCF, trying to reuse existing PanDA components and workflow as much as possible. The PanDA connection layer runs on front-end nodes in user space. There is a predefined host to communicate with CERN from the OLCF, with connections initiated from the front-end nodes. We use the SAGA framework [10] (a Simple API for Grid Applications) as the local batch interface.

The PanDA architecture on Titan is shown in figure 2. The PanDA Pilot, responsible for payload submission, runs on a Titan interactive node, communicating with local batch scheduler to manage jobs on Titan. Outputs are transferred to the Brookhaven National Laboratory Tier-1 Grid center. The initial demonstration was done using Monte Carlo event generators as a payload. Event generators are

mostly computational having stand-alone code with a small amount of stage-in/out data. For LHC experiments 10-15% of all CPU resources are utilized to run event generators.

It is important to define the mode of payload execution. One of possible scenarios would be to run HEP applications in backfill mode. Since HPC centers are geared toward large-scale jobs, about 10% of capacity of a typical HPC machine is unused due to a mismatch between job sizes and available resources. A rough estimation of a potential resource is 300 M CPU hours per year, primarily available via short jobs. PanDA can be a vehicle to harvest Titan resources opportunistically. Taking the above considerations into account, we have adapted the PanDA Pilot algorithm to use backfill information and to submit jobs in backfill mode as follows:

- Pilot periodically queries MOAB scheduler about available resources.
- Scheduler returns information about available (unscheduled) nodes and interval of availability.
- Pilot chooses the largest available block of nodes and submits jobs taken into account Titan's scheduling policy limitations.

The interest to run LHC experiment payloads using the scenario described above was expressed by several HPC centers: Czech IT4i, Swiss CSCS (Piz Daint), Taipei ASGC, and others.
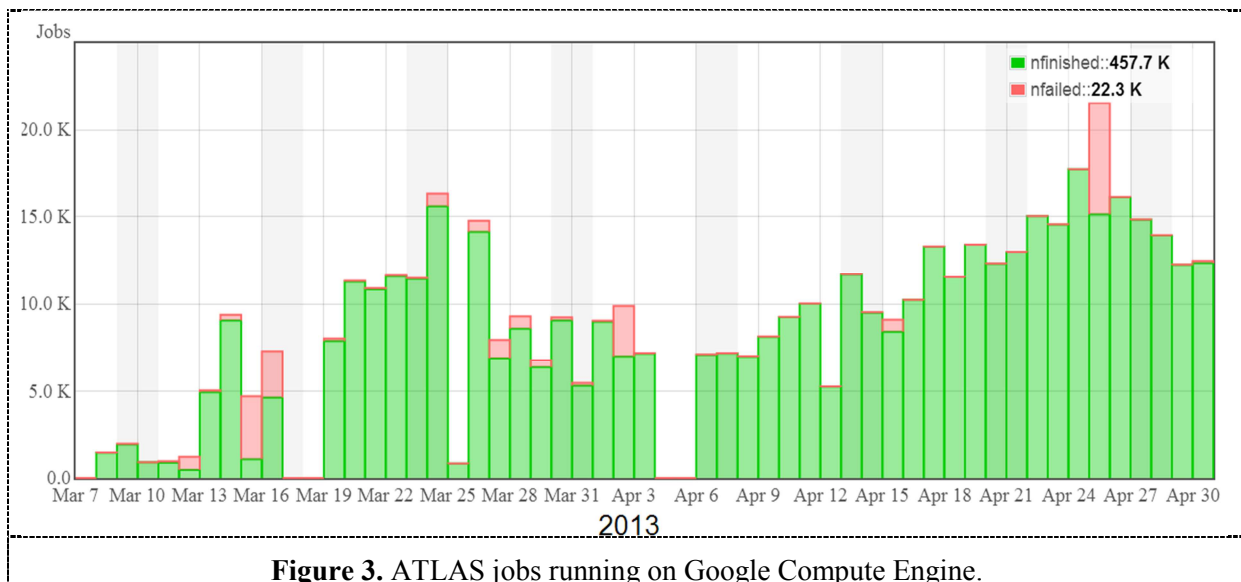

Figure 3. ATLAS jobs running on Google Compute Engine.

3.4. Extending PanDA to Clouds

The demand on computing resources to accommodate physics calculation requirements increases inexorably. Commercial and research clouds are one of the venues we explored over the past two years. ATLAS was invited to participate in the Google Compute Cloud preview. After an initial period, Google allocated 5 M core hours (using 4 k cores) for two months. The idea was to run a realistic physics application to gain experience and to understand a long-term stability while running cloud facilities similar in size to a typical Tier-2 Grid center. We ran ATLAS Monte Carlo production for 10 weeks (figure 3). We reached a throughput of 15 k jobs per day, with 214 M events have been generated and processed. A similar test has been conducted on Amazon EC2.

In broad terms, commercial clouds are more expensive than Grid Computing Centers, but clouds resources may be used to advantage by scientific collaborations during peak periods. Figure 4 shows the variation in the number of Grid jobs waiting because of lack of resources, with peaks corresponding to the activities before major physics conferences.

Figure 4. ATLAS: Waiting Jobs (weekly average).
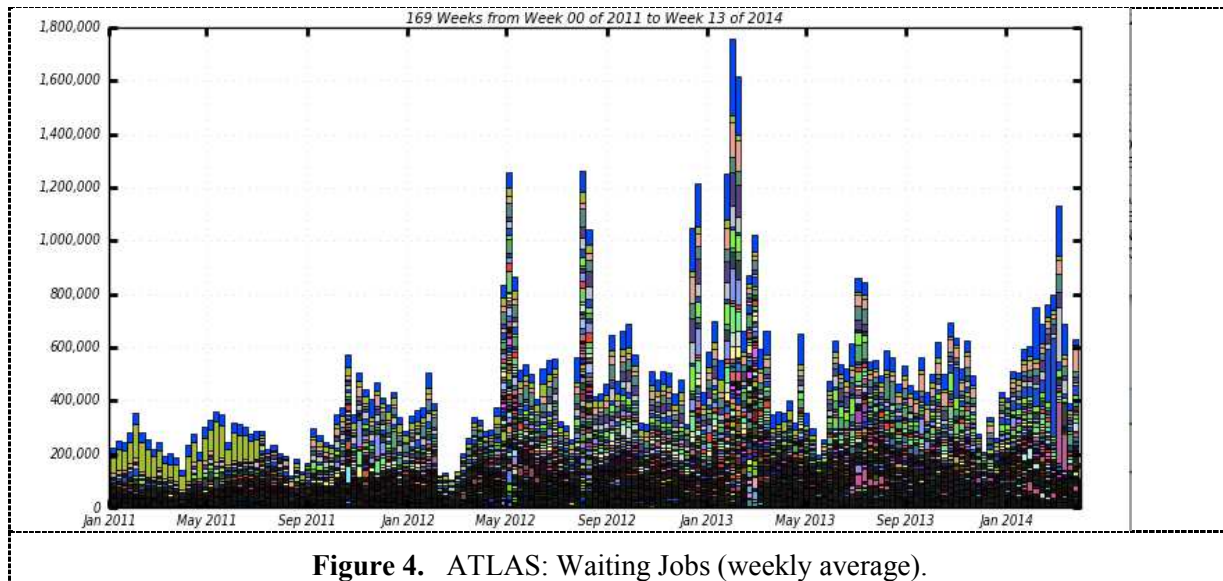
4. Conclusions

It is now accepted that High Energy Physics and High Energy Nuclear Physics applications must undergo a transformation to be able to make good use of current and future computer architectures. Only by using hardware efficiently will the experiments in those (and other) communities be able to achieve the massive computing throughput that will be needed in the next decade. HPC facilities are a particular target because of the massive computing power they can bring to bear, and because they provide the lowest possible communications overhead for massively decomposed tasks.

Fully optimized use of HPC facilities is a long-term goal, certainly requiring work of five or more years. Early access to HPC facilities will provide an important stimulus to this work. Access to facilities coupled with collaborative help in the transformation of HEP code would be a major scientific contribution to the physics discoveries of the next ten years.

The PanDA system played a key role during LHC Run 1 (2009-2013) data reprocessing, simulation and analysis. The challenge of how to process and analyze the data and produce timely physics results was substantial, but at the end resulted in a great success.

The interest in PanDA by other big data sciences provided the primary motivation to generalize the PanDA system. DOE ASCR gave us a great opportunity to start the BigPanDA project.

- The work on extending PanDA to Leadership Computing Facilities has started. PanDA has been successfully ported to OLCF Titan.
- Large-scale PanDA deployments on commercial clouds are already producing valuable results.

The work reported here will enable the use of PanDA by new scientific collaborations and communities as a means of leveraging extreme scale computing resources with a low barrier of entry. The technology base provided by BigPanDA will enhance the usage of a variety of high-performance computing resources available to basic research.

Acknowledgments

References
[1] LHC – The Large Hadron Collider, URL http://lhc.web.cern.ch
[2] The ATLAS Collaboration 2008 The ATLAS experiment at the CERN Large Hadron Collider *J. Inst.* **3** S08003
[3] The ALICE Collaboration 1995 *Technical proposal for A Large Ion Collider Experiment at the CERN LHC* LHCC, December 1995, ISBN 92-9083-077-8
[4] WLCG Worldwide LHC Computing Grid, URL http://lcg.web.cern.ch
[5] Maeno T *et al.* 2008 PanDA: Distributed Production and Distributed Analysis System for ATLAS *J. Phys.: Conf. Ser.* **119** (6) 062036
[6] Damjanović A, Miller B T, Wenaus T J, Maksimović P, García-Moreno E B, Brooks B R 2008 Open science grid study of the coupling between conformation and water content in the interior of a protein, *J Chem Inf Model* Oct; **48** (10) 2021-9 doi:10.1021/ci800263c
[7] Battiston R 2008 The antimatter spectrometer (AMS-02): A particle physics detector in space *Nucl. Instrum. Methods Phys. Res., Sect. A* **588**, 227-234
[8] Ivezic Z *et al.* 2008 LSST: from Science Drivers to Reference Design and Anticipated Data Products, *Preprint* astro-ph arXiv:0805.2366
[9] Next Generation Workload Management and Analysis System for Big Data, URL http://science.energy.gov/~/media/ascr/pdf/research/nextgen/KlimentovAlexei_BNL_abstract.pdf
[10] ATLAS Collaboration 2012 Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC *Phys. Lett. B* **716** 1-29
[11] The SAGA Project, URL http://www.saga-project.org