

# Submission Details: symp151s1

Submitted: 2018-01-08 01:04 · Last updated: 2018-01-08 12:22

## Scientific Domain

*Physics:* 1

*Computer Science and Applied Mathematics:* 2

## Title (Maximum 20 words)

*Title (Maximum 20 words):* Opening Doors for New Data Intensive Sciences at Titan and Beyond Using BigPanDA

## Organizer Information

### Organizer 1:

*Gender:* male

*Name:* Dr. Kaushik De

*Email:* kaushik@uta.edu

*Affiliation:* University of Texas at Arlington

*Affiliation 2:*

*Country of the Primary Affiliation:* United States of America

*Primary Telephone:* +1 817 272 2813

*Primary Phone Type:* Work

*Will this person chair the minisymposium at the conference?* No

*Is this person on the reviewing committee?* No

### Organizer 2:

*Gender:* male

*Name:* Dr. Alexei Klimentov

*Email:* Alexei.Klimentov@cern.ch

*Affiliation:* Brookhaven National Laboratory

*Affiliation 2:*

*Country of the Primary Affiliation:* United States of America

*Primary Telephone:* +1 631 344 7855

*Primary Phone Type:* Work

*Will this person chair the minisymposium at the conference?* No

*Is this person on the reviewing committee?* No

### Organizer 3:

*Gender:* male

*Name:* Dr. Torre Wenaus

*Email:* wenaus@gmail.com

*Affiliation:* Brookhaven National Laboratory

*Affiliation 2:*

*Country of the Primary Affiliation:* United States of America

*Primary Telephone:* +1 631 344 4755

*Primary Phone Type:* Work

*Will this person chair the minisymposium at the conference?* Yes

*Is this person on the reviewing committee?* No

## Description for Review

*Description for Review (Maximum 800 words):*

The ATLAS experiment at CERN's Large Hadron Collider uses the Worldwide LHC Computing Grid, the WLCG, for its remote computing infrastructure. Through the workload management system PanDA, ATLAS provides seamless access to hundreds of WLCG grid and cloud based resources that are distributed worldwide, to thousands of physicists. PanDA annually processes over an exabyte of data using an average of 250,000 distributed batch slots, to enable hundreds of new scientific results from ATLAS. However, the resources available to PanDA have been insufficient to meet ATLAS simulation needs over the past few years as the volume of data from the LHC has grown. Consequently, an effort was launched to extend PanDA, called BigPanDA, to access HPC resources, funded by the US Department of Energy (DOE-ASCR). Through this successful effort, ATLAS today uses about 20 million hours monthly on the Titan supercomputer at Oak Ridge National Laboratory. Many other supercomputers have also been integrated into ATLAS computing, providing large periodic bursts of computing cycles to ATLAS. This minisymposia will explore the software and operational lessons learned from integrating HPCs with traditional grid computing resources in ATLAS, and describe recent efforts to use BigPanDA for many other scientific domains.

Many challenges were met and overcome in using HPCs for complex ATLAS workloads and workflows. New software systems were developed to integrate the different resources seamlessly. New operational models were required for distributed large scale facilities with short duration access. These software and operations innovations will be described in the talks and discussed by the expert panel in this proposed minisymposia.

The first talk will describe the software and operational lessons learned. ATLAS uses two different operational modes at Titan. Using allocations through specific experimental allocations is the traditional mode - but requires software innovations to fit the low latency requirements of experimental science. New techniques were implemented to fit the bursty nature of traditional large jobs using allocations on a leadership class machine. In the second mode, high priority work is constantly sent to Titan to backfill high priority leadership class jobs. This has resulted in impressive gains in overall utilization of Titan, while benefiting the physics objectives of ATLAS.

Following the success with ATLAS, BigPanDA is now being tested and prototyped for a wide range of non-LHC data sciences - from biology, to chemistry to other physics disciplines. Two talks will describe the use of BigPanDA at HPCs beyond ATLAS. Many groups of scientists, from small PI driven sciences, to the largest experiments, have recently started to use cycles on Titan through BigPanDA.

We propose an expert panel discussion for the final 30 minute slot in this minisymposia. It will primarily focus on future plans and new opportunities. The 3 speakers from this minisymposia will join other experts from HPC and LHC communities during the panel discussion.

This proposed minisymposia will provide an important venue to present the experience of integrating HPCs with traditional distributed systems, and to discuss the future of HPC systems for instrumental and data sciences.

**Abstract for Publishing**

*Abstract for Publishing (Maximum 200 words):*

The ATLAS experiment at CERN's Large Hadron Collider depends on the Worldwide LHC Computing Grid, the WLCG, for its remote computing infrastructure. PanDA, the workload management system used by ATLAS, annually processes over an exabyte of data using an average of 250,000 distributed batch slots, to enable hundreds of new scientific results. An effort was launched to extend PanDA, called BigPanDA, to access HPC resources, funded by the US Department of Energy (DOE-ASCR). Through this successful effort, ATLAS today uses about 20

million hours monthly on the Titan supercomputer at Oak Ridge National Laboratory. Many other supercomputers have also been integrated into ATLAS computing. This minisymposia will explore the software and operational lessons learned in integrating HPCs with traditional grid computing, and describe recent efforts to use BigPanDA for many other scientific domains. Three talks will summarize the state of the art and the future wishlist for HPC usage for current and future experiments, while a concluding expert panel discussion will focus on the future.

## **Presenter Information**

### **Presenter 1:**

*Gender:* male

*Name:* Dr. Alexei Klimentov

*Email:* Alexei.Klimentov@cern.ch

*Affiliation:* Brookhaven National Laboratory

*Affiliation 2:*

*Country of the Primary Affiliation:* United States of America

*Presentation Title:* BigPanDA Experience on Titan for the ATLAS Experiment at the LHC

### **Presenter 2:**

*Gender:* male

*Name:* Dr. Fabien Delalondre

*Email:* fabien.delalondre@epfl.ch

*Affiliation:* EPFL

*Affiliation 2:*

*Country of the Primary Affiliation:* Switzerland

*Presentation Title:* Blue Brain on Titan using BigPanDA

### **Presenter 3:**

*Gender:* male

*Name:* Mr. Danila Oleynik

*Email:* Danila.Oleynik@cern.ch

*Affiliation:* University of Texas at Arlington

*Affiliation 2:*

*Country of the Primary Affiliation:* United States of America

*Presentation Title:* Enabling Biology, Chemistry and Other Sciences on Titan through BigPanDA

### **Presenter 4:**

*Gender:* male

*Name:* Dr. Jack C. Wells

*Email:* wellsjc@ornl.gov

*Affiliation:* Oak Ridge National Laboratory

*Affiliation 2:*

*Country of the Primary Affiliation:* United States of America

*Presentation Title:* Panel Discussion: BigPanDA Experience at Oak Ridge - Learning from the LHC, Going Far Beyond