ABSTRACT

NEXT GENERATION WORKLOAD MANAGEMENT AND ANALYSIS SYSTEM FOR BIG DATA

Lead Institution: Brookhaven National Laboratory
Lead PI: Alexei Klimentov
Principal Investigators:
Brookhaven National Laboratory:
Group Leader, Physics Application Software: Alexei Klimentov (alexei.klimentov@cern.ch)
Senior Application Software Engineer: Sergey Panitkin (panitkin@bnl.gov)
Senior Physicist: Torre Wenaus (wenaus@gmail.com)
Group Leader, Computational Science Center: Dantong Yu (dtyu@bnl.gov)
Argonne National Laboratory:
Software Engineer: Alexandre Vaniachine (Sacha.Vaniachine@cern.ch)
The University of Texas at Arlington:
Professor of Physics: Kaushik De (kaushik@uta.edu)
Associate Professor of Computer Science Engineering: Gergely Zaruba (zaruba@uta.edu)

One of the largest scientific collaborations ever assembled, the ATLAS experiment at the Large Hadron Collider (LHC), is designed to explore the fundamental properties of matter for the next decade. The origin of mass, sources of dark matter and dark energy, higher symmetries of nature, extra dimensions of space-time, and other mysteries of nature will be studied while pushing the boundaries of our knowledge of the Standard Model of physics. In operation since 2009, the ATLAS experiment already manages over fifty petabytes of data at more than a hundred computer centers worldwide. Thousands of physicists analyze tens of millions of collisions from the LHC daily, leading to weekly publications of new results in peer-reviewed journals. Over the next few years, ATLAS data will grow to exabyte scale, providing an active incubator for data intensive scientific challenges.

An important foundation underlying the impressive success of ATLAS data processing and analysis is the Production and Distributed Analysis (PanDA) workload management system. Proposed and developed seven years ago in the US by the stakeholders in this proposal, and funded by the DoE and NSF, PanDA has been adopted by the more than three thousand physicists in the international ATLAS community spanning over forty countries. PanDA is deployed on all Worldwide LHC Computing Grid (WLCG) resources, comprising over a hundred thousand batch slots and over fifty petabytes of storage, and of late is expanding into research and commercial cloud resources. PanDA is a highly successful model of U.S. technological innovation in big data sciences.

PanDA was designed specifically for the ATLAS experiment at the LHC and has proved to be highly successful in meeting all the distributed computing needs of the experiment. However, the core design of PanDA is not experiment specific. The PanDA workload management system is capable of meeting the needs of other data intensive scientific applications. PanDA has been used to run the CHARMM protein folding code on the Open Science Grid (OSG). Several big data experiments such as the Alpha-Magnetic Spectrometer (AMS, an astro-particle experiment on the International Space Station) and the Compact Muon Solenoid (CMS, a LHC experiment) are considering the adoption of PanDA. We propose here a program to develop a generic version of PanDA which can be easily used by many data intensive sciences. With a modest investment of effort, we can enable easy adoption of PanDA by others. We propose generalizing PanDA as a meta-application, providing location transparency of processing and data access, for High Energy Physics, other data-intensive sciences, and a wider exascale community. The proponents of this proposal consist of the core PanDA development team at Brookhaven National Laboratory and the University of Texas at Arlington, who proposed and developed PanDA from a new concept to one of the most successful and comprehensive distributed software for big data.