Commenting on behalf of the HEP/NP data intensive, high throughput computing programs in the BNL Physics Department.

Name, credentials, organization

Torre Wenaus, Staff scientist, BNL Physics Department
Distributed Computing Co-Coordinator for the ATLAS experiment at the LHC
Software and Computing Manager for the US ATLAS Operations Program
PanDA Distributed Workload Management System Project Co-Leader and Developer

Contact information

wenaus@bnl.gov
631-344-4755

Experience working with big data and your role in the big data innovation ecosystem

RHIC ushered in high throughput, data intensive computing as a BNL capability 15 years ago, and LHC ATLAS computing at BNL has built on that over the past decade. The LHC is the largest scientific enterprise worldwide in data intensive computing. BNL's ATLAS Tier 1 Center is the largest LHC computing center outside CERN. The BNL Tier 1 is part of the RHIC and ATLAS Computing Facility (RACF), a unified facility delivering leveraged, cost-effective computing to both programs, as well as many others in BNL Physics. BNL also provides the PanDA (Production and Distributed Analysis) system that orchestrates ATLAS' data intensive computing at RACF and around the world. These capabilities make BNL one of the largest Big Data / High Throughput Computing resources and expertise pools in US science.

RACF consists of about 30 scientists and IT professionals; the PAS group consists of about 10. RACF runs over 40,000 concurrent jobs in steady state. Storage is O(100PB), about 25% of that on disk. Networking is at Terabit scale. Monthly averaged transfer rates are up to 800 MB/s between BNL and the rest of ATLAS; peaks are up to five times higher.

PanDA runs 150-200k jobs concurrently at well over 100 sites around the world, close to a million jobs a day, with about 1400 users. The system has reached Exascale in data volume processed: about 1.3 Exabytes were processed by PanDA in 2013. The ATLAS data set emerging from the first LHC run is currently about 150 PB; Run 2 (2015-2018) will accrue about six times the raw data volume of Run 1.

In its ATLAS and RHIC computing BNL is meeting challenges not only in terms of capability and scale but also in the cost effectiveness needed, in the prevailing environment of flat-flat budgets, with data processing demands growing by an order of magnitude over the next decade. BNL is a leader in developing cutting edge cost effective solutions. Examples include highly distributed federated storage serving scientists at dozens of sites across the US and globally; systems to leverage commercial clouds in a cost-effective and transparent way for data-intensive

processing, interfaced to peta-scale storage; intelligent, high performance networking via ESnet, and extending intelligent network awareness to the application level; innovations in fine-grained processing workflows to utilize transient opportunistic resources and powerful networking fully, economizing on storage; and highly scalable and accessible data storage via object stores. We are extending data intensive computing to new platforms and new resources opportunistically -- high performance computing (HPCs) as well as clouds. Efficient use of storage at massive scales is an objective common to many of these initiatives – data storage is the largest cost component for data intensive computing today.

Comments and suggestions on the Visions and Priority Actions document

We find the high impact frontier of big data computing to be in exploiting the synergies between 1) powerful intelligent networks, 2) intelligent data-aware distributed workload management, and 3) highly scalable and accessible data stores to build agile, cost-efficient data intensive computing capabilities that can enable a scientific enterprise to leverage widely disparate distributed resources in a manageable way, making possible scales of computation well beyond what conventional approaches allow. R&D and other investments in these areas could be game-changing in giving increasingly data intensive scientific domains an accessible on-ramp to massively scaled Big Data computing.

Why you feel your contribution/ideas should be included in the strategic plan

In BNL's HEP/NP computing program we've been experts at data intensive high throughput computing for 15 years, thanks to RHIC and ATLAS. ATLAS' unprecedented scale and BNL's leading roles in computing place BNL in the vanguard of large scale data intensive scientific computing, in the US and globally. While other science domains are relatively new to or only approaching truly Big Data computing, ATLAS' early ascent has made BNL an expert today. BNL leverages this capability to benefit the wider HEP, NP and scientific computing communities, and is also seeking to expand to new fields such as BES and biology, and identify new ways in which its capabilities can support and advance scientific computing in the US.