Workshop on Central Computing Support for Photon Sciences

Monday, 24 September 2018 - Monday, 24 September 2018

BNL, Physics

Book of Abstracts

Contents

Computing and Storage / **2**

Bursty Data Analytics on HPC

Author: David Skinner[1]

[1] *LBNL*

Corresponding Author: deskinner@lbl.gov

Abstract: We present a general overview of the challenges and opportunities in marshaling computational intensity around bursts of data generated in a (mostly) scheduled manner. Where computation "fits" in the data analytic pipeline (between detector and actionable knowledge) is an important architectural concern for advanced instruments with bursty data. Design boundary conditions include instrument duty-cycle, experiment predictability, the stubborn constancy of the speed of light, and a variety of data reduction opportunities and constraints. NERSC systems aim to capably capture the most intense computational peaks in these workflows. Opportunistically upstreaming computation in the analytic pipeline has significant promise in mitigating the "data deluge" through HPC-informed DAQ design, using NERSC systems to develop algorithms which can be back-ported to the DAQ system. Examples from LCLS and NCEM are presented with the intent of gathering future needs of DAQ designers.

Authorization & Authentication Infrastructure / **3**

Automating Inter-facility Science with Fine Grained Authorization

Author: Mark Day[1]

[1] *LBNL*

Corresponding Author: mrday@lbl.gov

Abstract: Inter-facility workflows by their nature cross facility boundaries thereby implying attention to how users are authenticated at each facility and how their workflow steps are authorized. A variety of approaches can be used to make these boundary crossings less intensive in terms of human effort. We suggest fine grained authorizations as a means to automation by forming a minimal set of inter-operational controls which abide the policies of both facilities. A spectrum of authorized actions are examined from read-only access, posting of future intents, to full access are considered in the context of photon science data analysis.

Computing and Storage / **4**

High Throughput Computing Challanges

Author: William Strecker-Kellogg[1]

[1] *BNL*

Corresponding Author: willsk@bnl.gov

Batch processing with (embarrassingly) parallel workloads affords numerous possible architectures all with their own tradeoffs and challenges. We will discuss the various architectures we've employed at the RACF/SDCC to tackle those challanges and our experiences with them

Archival Tape Storage / **5**

Archival Storage for a Scientific Research Environment

Author: David Yu[1]

[1] *Brookhaven National Lab*

Corresponding Author: david.yu@bnl.gov

As the precision, energy, and output of scientific instruments such as particle colliders increase, so does the volume of data generated from science experiments. With this data has increasing rapidly, there is a serious need to keep the data in storage that is reliable and cost effective. Disk storage is efficient – ideal for frequently accessed data – but is often very costly, and is not a good solution for long-term archiving, especially when data becomes less active.

Cold storage, such as tape storage, has been an ideal solution for long-term data preservation, due to being cost-effective, environmentally friendly and having a long lifespan. Tape technology has been improved in both capacity and performance over the recent decades. Therefore, tape technology has been playing a very important role in managing the exponential growth of scientific data. Tape systems are great for archiving, due to the scalability and high sequential writing speed. However, accessing files from massive amounts of tapes usually is a major challenge for the tape storage system.

In BNL, we have implemented a high throughput active archive system currently stored near 150 PB of scientific data and serving scientists from multiple collaborations worldwide. The implementation concept is based on the most cost-effective and energy-efficient (green) memory model available today.

In this presentation, we will describe the concept of our archival storage and the underlying tape storage complex, as well as the challenges we are facing for future scientific data.

Computing and Storage / **6**

Supporting users for the long term

Author: Andrew Richards[1]

[1] *Diamond Light Source Ltd.*

Corresponding Author: andrew.j.richards@diamond.ac.uk

The historical implementation of scientific computing infrastructure within the Diamond Light Source has been to support the initial stages of data acquisition, from immediate sample validation, to the first stages of data processing. For many users, once the visit period had completed there was little if any continued interaction with the data stored at Diamond.

As data volumes increase it has become increasingly not viable for users to take the data home with them, or to transfer the data to other compute facilities to further proceed with data analysis. As a result Diamond is now working to address how best to separate the data collection and analysis associated with the actual visit, and the continued requirement to post-process data at some point in time after the visit has concluded. This is leading to investigations of offsite, cloud-like, resources and to work closer with activities such as the developing IRIS e-Infrastructure, a compute and storage platform across multiple sites to support STFC funded facilities.
As Diamond has stored all data produced over the last 11 years of operation, it is also now looking at how any future data archive should be provisioned to potentially enable open access to future data sets. Access to data, the implications on software for analysis, implications on infrastructure from storage to network and where best to locate the data for future post-processing compute requirements are all current topics being investigated in order to provide a facility that supports its users for the long term, and not just during their visit.

Computing and Storage / **7**

Data Analysis as a Service at STFC

Author: Ian Collier[1]

[1] *STFC Rutherford Appleton Laboratory*

Corresponding Author: ian.collier@stfc.ac.uk

In order to better support the scientists using STFC Facilities (including The Diamond LightSource, the ISIS neutron source and CLF laser Facility) STFC's Scientific Computing Department has been developing the Data Analysis as a Service platform.
DAaaS brings together the facilities data already archived by SCD with a flexible, extensible platform to deploy scientific workflows on STFC's OpenStack based cloud platform.
We describe the challenges and potential benefits.

Authorization & Authentication Infrastructure / **8**

A modern approach to SSO at the RACF/SDCC

Author: Jamal Irving[1]

[1] *BNL*

Corresponding Author: jamal@bnl.gov

Authentication and authorization is an integral part of any organization to maintain integrity and control over shared resources. Historically, RACF/SDCC projects were responsible for maintaining their own user accounts and IT resources to perform the experiments the project received funding for. We will discuss our current authentication and authorization architecture that we use at the RACF/SDCC and the steps we are taking to modernize our authorization and authentication stack to facilitate single sign on and two factor authentication from within the facility and beyond.

Computing and Storage / **9**

Planning for the LCLS-II Data System: Requirements, Benchmarks and Design

Author: Amedeo Perazzo[1]

[1] *SLAC National Accelerator Laboratory*

Corresponding Author: perazzo@slac.stanford.edu

LCLS and SLAC have done extensive analysis to determine the facility computing needs over the next decade. Based on the set of experiments currently planned and based on today's understanding of the computing and data requirements we estimated the computing demand from LCLS-II. This presentation describes the methodology adopted for deriving the computing rates, throughput and storage for LCLS-II, and how this methodology drives the design of our future data system.

Wide Area Network / **10**

ESnet WAN Service & Support

Author: Michael O'Connor[1]

[1] *Esnet*

Corresponding Author: moc@es.net

ESnet is the DOE's High Performance Network (HPN). It is viewed as a scientific user facility and in many ways an instrument to accelerate research and discovery, with a history of alignment with national laboratory enterprise network organizations to accomplish facility missions and objectives, beyond the campus perimeter to collaborations scaling both nationally and globally. A discussion of network architectures supporting data transfer, such as Science DMZ and Research and Education Internet Exchanges.

Wide Area Network / 11

Network Architecture and Operations of the RACF/SDCC Facility

Author: Alexandr Zaytsev[1]

[1] *Brookhaven National Laboratory (BNL)*

Corresponding Author: alezayt@bnl.gov

This talk gives a summary of the current structure of network systems deployed in B515 based RACF/SDCC facility, emphasizing of the architecture of the Science Zone and its expected evolution in 2019-2023 period during which the new B725 based datacenter to be constructed under umbrella of the BNL Computing Facility Revitalization (CFR) project is expected to enter the fully operational state. The main challenges and shifts in network technology anticipated in this timeframe are discussed. The model of how other BNL user facilities can access the RACF/SDCC resources and use RACF/SDCC Facility as a bridge to external network resources is outlined.

Computing and Storage / 12

Computing&Storage for on-site experiments - Petra3, FLASH and EuXFEL

Author: Martin Gasthuber[1]

[1] *DESY*

Corresponding Author: martin.gasthuber@desy.de

This presentation will give a short overview of the current setup and operation for the DAQ and offline analysis part. The second half will focus on current activities/improvements for the storage and online data analysis services.

Software Support / 13

BNL Box

Author: Hironori Ito[1]

[1] *Brookhaven National Laboratory*

Corresponding Author: hito@bnl.gov

An update on the latest developments towards production use of BNL Box.

Computing and Storage / 14

NSLS-II status and computing challenges

Authors: Stuart Campbell[1]; Richard Farnsworth[2]

[1] *NSLS-II*

[2] *NSLS-II*

Corresponding Authors: campbell@bnl.gov, rif@bnl.gov

A short update of the Status of the NSLS-II computing situation, its challenges and futures. Including data retention, intentions, remote access requirements, remote control now and further, Real time cluster needs, experimental support and post experiment analysis.
Some discussion of the things that work well and some things that don't.

Software Support / 15

Ubiquitous Big Data: Supporting the Proliferation of Big Data Experiments from the Data Center

Author: Shigeki Misawa[1]

[1] *BNL*

Corresponding Author: misawa@bnl.gov

Advances in electronics has resulted in the explosion of scientific instruments that are capable of generating "Big Data". Traditionally, there were few big data experiments in operation at any given instant. These experiments were typically large endeavors with the financial, infrastructure, computing, and personnel resources to manage the data volumes generated by the experiment. With the proliferation of big data scientific instruments, next generation Big Data experiments will typically be smaller scale operations with limited resources. Individually, these next generation experiments may not have the resources to handle the data volumes. However, collectively, they may be able to support the necessary computing and storage infrastructure. This talk discusses how a central data center can help provide these resources.