The Blurry Line Between Exascale and Big Data

E. Wes Bethel, Lawrence Berkeley National Laboratory

Background

Over the past few decades, the US Department of Energy (DOE) has invested substantial resources in computing infrastructure and in the R&D programs needed to enable use of high-end computing infrastructure for use by its scientific research community. The facilities large-scale computing centers located at national laboratories along with a leading-edge production network (ESnet). The R&D programs have included a diverse portfolio of projects targeting cultivation of key technology (e.g., applied mathematics, performance optimization, programming models, and more), and applying these technologies to science projects.

Insofar as facilities and programmatic planning goes, one could characterize the vast majority of requirements-gathering efforts in past years as focusing on topics like what are the science challenges facing DOE science programs that could be addressed with a peta or exascale class machine, along with articulating the CS challenges needed to realize those objectives. A handful of recent workshops have focused on data-centric issues. Requirements and challenges are usually gathered in an *ad hoc* fashion, in the form of workshops organized by one or more DOE offices. See, for example, workshop reports for DOE's Office of Advanced Scientific Computing[1].

As with other agencies and organizations, there is increasing awareness of the impact, opportunities, and challenges associated with the fact we live in world increasingly dominated by data-centric challenges. The "data problems" are faced by essentially all science projects, ranging from those that are primarily computational- and simulation-based, to those that are experimental and/or observational based. Some elements of DOE's current Exascale Research planning process are focusing on data issues, though those are limited to a focus on data produced by simulations on the exascale platform as well as use of the exascale platform for the processing of experimental and observational data (EOD) that is resident on persistent storage accessible to the exascale platform.

Position Statement

My position statement can be articulated as: *There is a blurry line between "big data" and HPC concerns.*

The idea of "big data" can be succinctly summarized as follows: there is a wealth of information and knowledge waiting to be discovered in collections of data. This idea is not new, by any means, yet we are more acutely aware of this fact in the present due to the explosive growth in data that results from an evolving technological landscape.

By HPC concerns, I'm primarily focusing on the combination of facilities and R&D programs that use them and provide technology that enable their use, and the increasingly larger and more complex collections of data that are produced by computational projects on those platforms.

[1] Office of Advanced Scientific Computing Workshop U. S. Department of Energy. ASCR Program Documents. `http://science.energy.gov/ascr/news-and-resources/program-documents/`, 2015. [Online; last accessed Sun Jan 18 08:36:13 PST 2015].

While there are clear differences in the needs of computational and EOD projects, the "data challenges" facing both types of projects are largely the same. Both produce ever-larger and quantities of data. Both are faced with a widening gap between the ability to produce/collect data and the ability to manage and gain understanding from it. Both can benefit from a movement towards a more unified, standardized infrastructure and ecosystem for managing and working with data; ad hoc, single-investigator solutions are simply not a sustainable approach for a large program.

Issues to Consider Moving Forward

The requirements gathering process. Shaping program focus areas and priorities requires understanding the needs of several different segments of the population. The science community will articulate what new science they wish to perform using HPC/data technologies. The CS R&D community, as well as industry, will provide input as to potential approaches for meeting those needs. The requirements gathering process should include a focus on both computational and data issues.

Evolution of the HPC programs. In the past, HPC centers focused their attention on the collection of software, the "software stack", run on their machines. Procurements were engineered to provision a resource targeted at servicing the anticipated workload. This approach has worked well for many years. Going forward, the definition of "workload" needs to expand to include data-centric needs in addition to the traditional computational workload. One challenge is that these data-centric needs are likely more diverse and difficult to quantify than computational-only workloads. Concurrently, it should be recognized that an increasing portion of the workload will be run in a distributed fashion, with some processing happening close to where data is first collected, then moved to an HPC center for creating large collections, community access, more advanced and extensive processing, and so forth.

Defining the "data-centric" program. Given an evolving workload for HPC programs and the need to characterize the data-centric workloads, the process of defining the needs and technical approaches for a data program should be on "equal footing" with computational requirements and workload assessments. Given the increasingly blurry line between HPC and data-centric concerns, It may be more productive, albeit more of a burden on the community, to combine the assessment of computational and data needs, challenges, and opportunities into a unified process.

Towards a unified data/HPC program. In the HPC community, there has emerged a set of reasonably stable "standards" and infrastructure that have enabled long-term, sustained growth of computational science efforts. For example, at the dawn of the MPP era in the early 1990s, code teams began to migrate from serial Fortran and a focus shared-memory parallelism and vectorization to a distributed-memory programming/execution model. Codes written in MPI were able to largely survive technology evolution for the next two decades, with recent advances focusing on so-called hybrid parallelism, where MPI accommodates distributed-memory programming/runtime issues, and other technologies (e.g., OpenMP, etc.) accommodate shared-memory parallel issues. HPC R&D programs over the years have largely focused on cultivating and managing the technology pipeline to facilitate the creation of codes that can run effectively on HPC platforms. The new issue is to identify the equivalents in the data space. There are traditional concepts, from an HPC sense, that are reasonably well understood: parallel I/O libraries, and standalone applications for visualization and analysis. However, the issues facing data-centric science are much deeper and far-reaching than can be addressed only with I/O libraries and a handful of off-the-shelf applications. The R&D programs in data should cultivate growth for the diverse technologies needed to enable data-centric science to effectively use HPC resources.