

The Scientific Data Management Center

Arie Shoshani

Lawrence Berkeley National Laboratory

<http://sdmcenter.lbl.gov>

Summary

With the increasing volume and complexity of data produced by ultra-scale simulations and high-throughput experiments, understanding the science is largely hampered by the lack of comprehensive, end-to-end data management solutions ranging from initial data acquisition to final analysis and visualization. The Scientific Data Management (SDM) Center is bringing a set of advanced data management technologies to DOE scientists in various application domains including astrophysics, climate, fusion, and biology. Equally important, it has established collaborations with these scientists to better understand their science as well as their forthcoming data management and data analytics challenges. The SDM center has provided advanced data management technologies to DOE domain scientists in the areas of storage efficient access, data mining and analysis, and scientific process automation.

Managing scientific data has been identified as one of the most important emerging needs by the scientific community because of the sheer volume and increasing complexity of data being collected. Effectively generating, managing, and analyzing this information requires a comprehensive, end-to-end approach to data management that encompasses all of the stages from the initial data acquisition to the final analysis of the data. Based on the community input, we have identified three significant requirements. First, *more efficient access to storage systems is needed*. In particular, parallel file system improvements are needed to write and read large volumes of data without slowing a simulation, analysis, or visualization engine. These processes are complicated by the fact that scientific data are structured differently for specific application domains, and are stored in specialized file formats. Second, scientists require technologies to facilitate better understanding of their data, in particular *the ability to effectively perform complex data*

analysis and searches over large data sets. Specialized feature discovery and statistical analysis techniques are needed before the data can be understood or visualized. To facilitate efficient access it is necessary to keep track of the location of the datasets, effectively manage storage resources, and efficiently select subsets of the data. Finally, generating the data, collecting and storing the results, data post-processing, and analysis of results is a tedious, fragmented process. Tools for *automation of this process in a robust, tractable, and recoverable fashion are required to enhance scientific exploration*.

Accordingly, we have organized our activities in three layers that abstract the end-to-end data flow described above. We labeled the layers as Storage Efficient Access (SEA), Data Mining and Analysis (DMA), and Scientific Process Automation (SPA). The SEA layer is immediately on top of hardware, operating systems, file systems, and mass storage systems, and provides parallel data access technology and


transparent access to archival storage. The DMA layer, which builds on the functionality of the SEA layer, consists of indexing, feature selection, and parallel statistical analysis technology. The SPA layer, which is on top of the DMA layer, provides the ability to compose scientific workflows from the components in the DMA layer as well as application specific modules.

Over the last several years, the technologies supported by the SDM center have been deployed for a variety of application domains. Some of the most notable achievements are:

- More than a tenfold speedup in writing and reading netCDF files has been achieved by developing MPI-IO based Parallel netCDF software being utilized by astrophysics, climate, and Parallel VTK.
- Scalability to thousands of processors has been achieved by sophisticated usage of parallel and collective I/O techniques in Blast, a widely used bioinformatics sequence analysis tool. The resulting mpiBlast-pio open source library has overcome a hundred node scalability bottleneck of mpiBlast.
- An improved version of PVFS is now

offered by cluster vendors, including Dell, Atipa, and Platform, and PVFS is the only freely available parallel file system supported on IBM's BlueGene/L.

- A method for the correct classification of orbits in puncture plots from the National Compact Stellarator eXperiment (NCSX) at PPPL was developed by converting the data into polar coordinates and fitting second-order polynomials to the data points.
- A new bitmap indexing method has enabled efficient search over billions of collisions (events) in High Energy Physics, and is being applied to combustion, astrophysics, and visualization domains. It achieves more than a tenfold speedup in generating regions and tracking them over time.
- The development of a Parallel R, an open source parallel version of the popular statistical package R. These are being applied to climate, GIS, and mass spec proteomics applications.
- An easy-to-use GUI-based software, called ProRata, has provided DOE GTL researchers with robust quantification of protein abundance in high-throughput shotgun proteomics data.
- A scientific workflow management and execution system (called Kepler) has been developed and deployed within multiple scientific domains, including genomics and astrophysics.



For further information on this subject contact:

Dr. Arie Shoshani
Lawrence Berkeley National
Laboratory
Tel: (510) 486-5171
Email: shoshani@lbl.gov