

DATA LIB

Exascale applications generate massive amounts of data that need to be analyzed and stored to achieve their science goals. The speed at which the data can be written to the storage system is a critical factor in achieving these goals. As exascale architectures become more complex, with multiple compute nodes and accelerators and heterogenous memory systems, the storage technologies must evolve to support these architectural features. The DataLib project is focused on three distinct and critical aspects of successful storage and I/O technologies for exascale applications: enhancing and enabling traditional I/O libraries on pre-exascale and exascale architectures; establishing a nascent paradigm of data services specialized for exascale codes; and working closely with Facilities to ensure the successful deployment of their tools.

The ability to efficiently store data to the file system is a key requirement for all scientific applications. The DataLib project is providing both standards-based and custom storage and I/O solutions for exascale applications on upcoming platforms. The primary goals of this effort are to enable users of the HDF5 standard to achieve the levels of performance seen from custom codes and tools, facilitate the productization and porting of data services and I/O middleware using Mochi technologies, and continue to support application and Facility interactions using DataLib technologies.

HDF5 is the most popular high-level API for interacting with the storage system on high-

performance computers. The DataLib team is undertaking a systematic software development activity to deliver an HDF5 API implementation that achieves the highest possible performance on exascale platforms. By adopting the HDF5 API, the team is able to support the I/O needs of all the exascale applications already using this standard.

The Mochi software tool is a building block for user-level distributed data services that addresses performance, programmability, and portability. The Mochi framework components are being used by multiple exascale library and application developers, and the team is engaging with them to customize data services for their needs.

Progress to date

- The DataLib team has improved the ability to understand exascale application I/O performance using their Darshan tool, improved the performance of I/O for codes using their PnetCDF and ROMIO technologies, implemented new capabilities for storing intermediate data on burst buffers and for building custom data services, developed and refined packaging and testing of DataLib software, and supported the use of Mochi technologies in other ECP ST and vendor products.

**PI: Rob Ross, Argonne National
Laboratory**

**Collaborators: Argonne National
Laboratory, Los Alamos National
Laboratory, Northwestern University**