

The HDF5 “Value Proposition” for the Fusion Data Lifecycle

Gerd Heber and Quincey Koziol, The HDF Group, {gheber, koziol}@hdfgroup.org

It Takes a Village to Publish a Paper

Building a well-designed data standard that incorporates the needs of a science community has a long-lasting value to that community (and beyond). It vastly outweighs the momentary benefits of particular hardware or software choices at any individual experimental site - the science data lifecycle involves more than just “speeds & feeds” during production. Creating a standard that captures the necessary metadata required to characterize experimental and simulation data, while accommodating future expansion and providing flexibility for the special needs of individual researchers is a challenging, but worthwhile endeavor.

Community data standards have taken root in many domains, giving researchers the ability to collaborate on larger science projects than previously possible. For example, the earth science community data standard, HDF-EOS [1], was defined in 1999, giving researchers a clearly defined standard and an open source software library for accessing that data. Having a single data standard allows researchers to track climate trends, create daily forecasts, and validate climate simulations against observational data, all without worrying about data exchange issues. A single standard also allows the community to reach a critical mass that entices commercial and open source data analysis tools to support the standard, reducing the effort spent within the community in writing one-off tools for specialty data formats.

If there is one problem that stands out for a fusion data standard, it is that of *continuous data integration*. For example, the ITER [2] experiment will operate in a pulsed mode, but data collection never really ceases and is practically continuous. Different kinds of data, which are produced nearly simultaneously and at different rates across many systems, instruments and platforms, need to be integrated. This integration has to accommodate the taxing performance requirements during acquisition and the access and dissemination needs of a community anxious to get its hands on the “hottest data.” Finally, given the expected lifetime of the facility, data needs to be integrated through time against the backdrop of ever-changing instruments, processor architectures, storage media, languages, compilers, etc. [3] Many of these challenges can be addressed by using HDF5 to gather and store ITER data.

How HDF5 Can Help Tackle the Continuous Data Integration Challenge

HDF5 [4] is a free and open source software-based storage solution featuring:

- A published, portable, self-describing family of file formats to represent a small set of highly customizable data primitives, which comes with a backward/forward compatibility guarantee covering both the programming API and file format
- An access library with APIs for C, C++, Java, and Fortran, and community developed language bindings including Python, R, Perl, and an extensive collection of FOSS tools and full commercial support from IDL, Matlab, LabView and others
- Established worldwide use across all sectors (industry, government, research) based on high quality source code and documentation developed, maintained and supported since 1996, and owned by a mission-driven, not-for-profit company since 2006

Obviously, HDF5 is not a technology stack for fusion science, but the potential benefits of HDF5 to the community can be demonstrated at practically any stage in the data lifecycle. Here we focus on three areas: data acquisition, storage and access.

During data acquisition, HDF5 performance is typically on par with the peak performance of the underlying storage layer, despite the addition of HDF5's metadata to the I/O stream. For time series, the *fast append* feature is of interest, which has been optimized for accessing datasets with one unlimited and an arbitrary number of fixed dimensions. HDF5 currently does not support multiple simultaneous writers to a single file. However, a novel, more efficient option is to let independent processes, e.g., detectors, write to datasets in separate files and compose a global view using a *Virtual Dataset* (VDS). These high-speed acquisition modes can be combined with the *Single-Writer/Multiple-Readers* (SWMR) capability to perform online analysis in a robust manner and without a need for interprocess communication or file locking.

When storing data, the rich, portable metadata capabilities including directed graph structures (e.g., hierarchies), complex attributes, and inter-object references make HDF5 a superior choice for maintaining the bond between data and metadata at the lowest level. These capabilities protect against bit rot and aid recoverability in the event of data corruption in derived metadata and cataloging layers. In addition, data integrity within HDF5 files can be established through multi-layered checksumming from file internal metadata to whole files.

Despite rapid changes in storage technologies, HDF5 acts as a differential to minimize the impact on applications. With the HDF5 *Virtual Object Layer* (VOL), which might be described as a *storage adapter plugin API for the HDF5 library*. With an HDF5 VOL plugin, the storage location (local, remote) and layout (native HDF5 file, cloud storage, etc.) become transparent, since no modifications of applications using the standard HDF5 API are required. There are now two remote access capabilities available for HDF5, which it traditionally did not provide: NETVOL is a client/server package for remotely accessing HDF5 data files, and h5serv [5] is a RESTful API for HDF5.

Without community HDF5 is empty and community without HDF5 is destitute. (A. Einstein)
Community involvement is an essential part of the HDF Group's mission: It is vital to sustaining the business and is our brain trust when making decisions about changes to HDF5, setting priorities, and adding new features. Regarding the data integration challenge, we believe that the fusion community could provide valuable insights into challenges they face with HDF5 and guidance in the areas of data provenance, replication, and indexing. We've heard from other communities, but balance is key for *building things right and building the right things*; we are eager to engage the fusion community to create a community data standard that will meet their needs for the grand science experiments planned in the years ahead and to help define the roadmap for HDF5.

- [1] <http://earthobservatory.nasa.gov/Features/HDFEOS/>
- [2] <http://www.iter.org>
- [3] Layne, R., and Capel, A., and Cook, N. and Wheatley, M. *Long Term Preservation of Scientific Data: Lessons from JET and Other Domains*. EFDA-JET-CP(11) 03/02
- [4] <http://hdfgroup.org/HDF5/>
- [5] <https://github.com/HDFGroup/h5serv>