



The HDF Group



SCIENCE TEAM SUPPORT FOR HDF5 ON BLUE WATERS

PRINCIPAL INVESTIGATOR:

Quincey Koziol (koziol@hdfgroup.org)

TEAM LEAD AND HDF APPLICATION SPECIALIST:

Scot Breitenfeld (brtnfld@hdfgroup.org)

PARALLEL HDF DEVELOPMENT:

Mohamad Chaarawi (chaarawi@hdfgroup.org)

THE HDF GROUP HOME PAGE:

<https://www.hdfgroup.org>

HDF BLUE WATERS PROJECT SUMMARY PAGE:

<https://ncsa-bw.atlassian.net/wiki/display/HDFBW>



Why use HDF5?

CHALLENGING DATA

- Application data that pushes the limits of what can be addressed by traditional database systems, XML documents, or in-house data formats.

SOFTWARE SOLUTIONS

- For very large datasets, very fast access requirements, or very complex datasets.
- To easily share data across a wide variety of computational platforms using applications written in different programming languages.
- That take advantage of the many open-source and commercial tools that understand HDF5.
- Enabling long-term preservation of data.



Recent HPC Success Story

SIMULATION

- I/O Kernel of a DOE plasma physics application
- ~10 Trillion particles
- Using 1 GB stripe size and 160 Lustre OSTs
- Running on **298,048** cores

PERFORMANCE RESULTS ON BLUE WATERS

- Achieved 52 GB/s
- ~50% of the peak performance



HDF5 1.10 Planned Features: HPC

IMPROVE USE OF MPI AND PARALLEL FILE SYSTEMS

- Remove 'truncate' operation on file close, etc.

REDUCE # OF I/O ACCESSSES FOR METADATA ACCESS

- Collective Read/Write of metadata

MULTI-DATASET COLLECTIVE I/O

SUPPORT FOR COMPRESSION IN PARALLEL

- COLLECTIVE ACCESS MODE ONLY

POSSIBLY SUPPORT SINGLE-WRITE/MULTIPLE-READER (SWMR) ACCESS IN PARALLEL



Services Provided

WE ARE HERE TO ASSIST YOU BY:

- Performing assessments of your application's needs
 - Audit of your application's current HDF5 use
 - Planning and assessments of potential improvements
- Providing expedited resolution of HDF5 defects and performance requirements
- Supporting development efforts based on specific applications needs
 - Help in introducing HDF5 into your application