



HDF5

Past, Present, and Future

Quincey Koziol

Ruth Aydt

The HDF Group

{koziol, aydt}@hdfgroup.org



Outline

- HDF5 in the past and present
 - What is HDF5 and how is it used?
 - What is The HDF Group?
- What is scheduled for next major release?
- What is further out on the horizon?



What is HDF5?

- **A versatile data model** that can represent very complex data objects and a wide variety of metadata.
- **A completely portable file format** with no limit on the number or size of data objects stored.
- **An open source software library** that runs on a wide range of computational platforms, from cell phones to massively parallel systems, and implements a high-level API with C, C++, Fortran 90, and Java interfaces.
- **A rich set of integrated performance features** that allow for access time and storage space optimizations.
- **Tools and applications** for managing, manipulating, viewing, and analyzing the data in the collection.



Why 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.
 - Very large datasets, very fast access requirements, or very complex datasets.
- Accessible data:
 - Data accessed by applications written in different programming languages on a wide variety of computational platforms.
 - Data shared within and across disciplines.
 - Data processed/visualized/analyzed by commercial and open-source tools that understand HDF5.



Who uses HDF5?

- Some examples of HDF5 users
 - Astrophysics
 - Astronomers
 - NASA Earth Science Enterprise
 - DOE Labs
 - Supercomputing centers in US, Europe, and Asia
 - Financial Institutions
 - NOAA
 - Manufacturing industries
 - Many others
- For a more detailed list, visit
 - <http://www.hdfgroup.org/HDF5/users5.html>



What is The HDF Group?



The HDF Group Mission

The HDF Group facilitates discovery, now and in the future, through its stewardship of HDF technologies and support of HDF users.

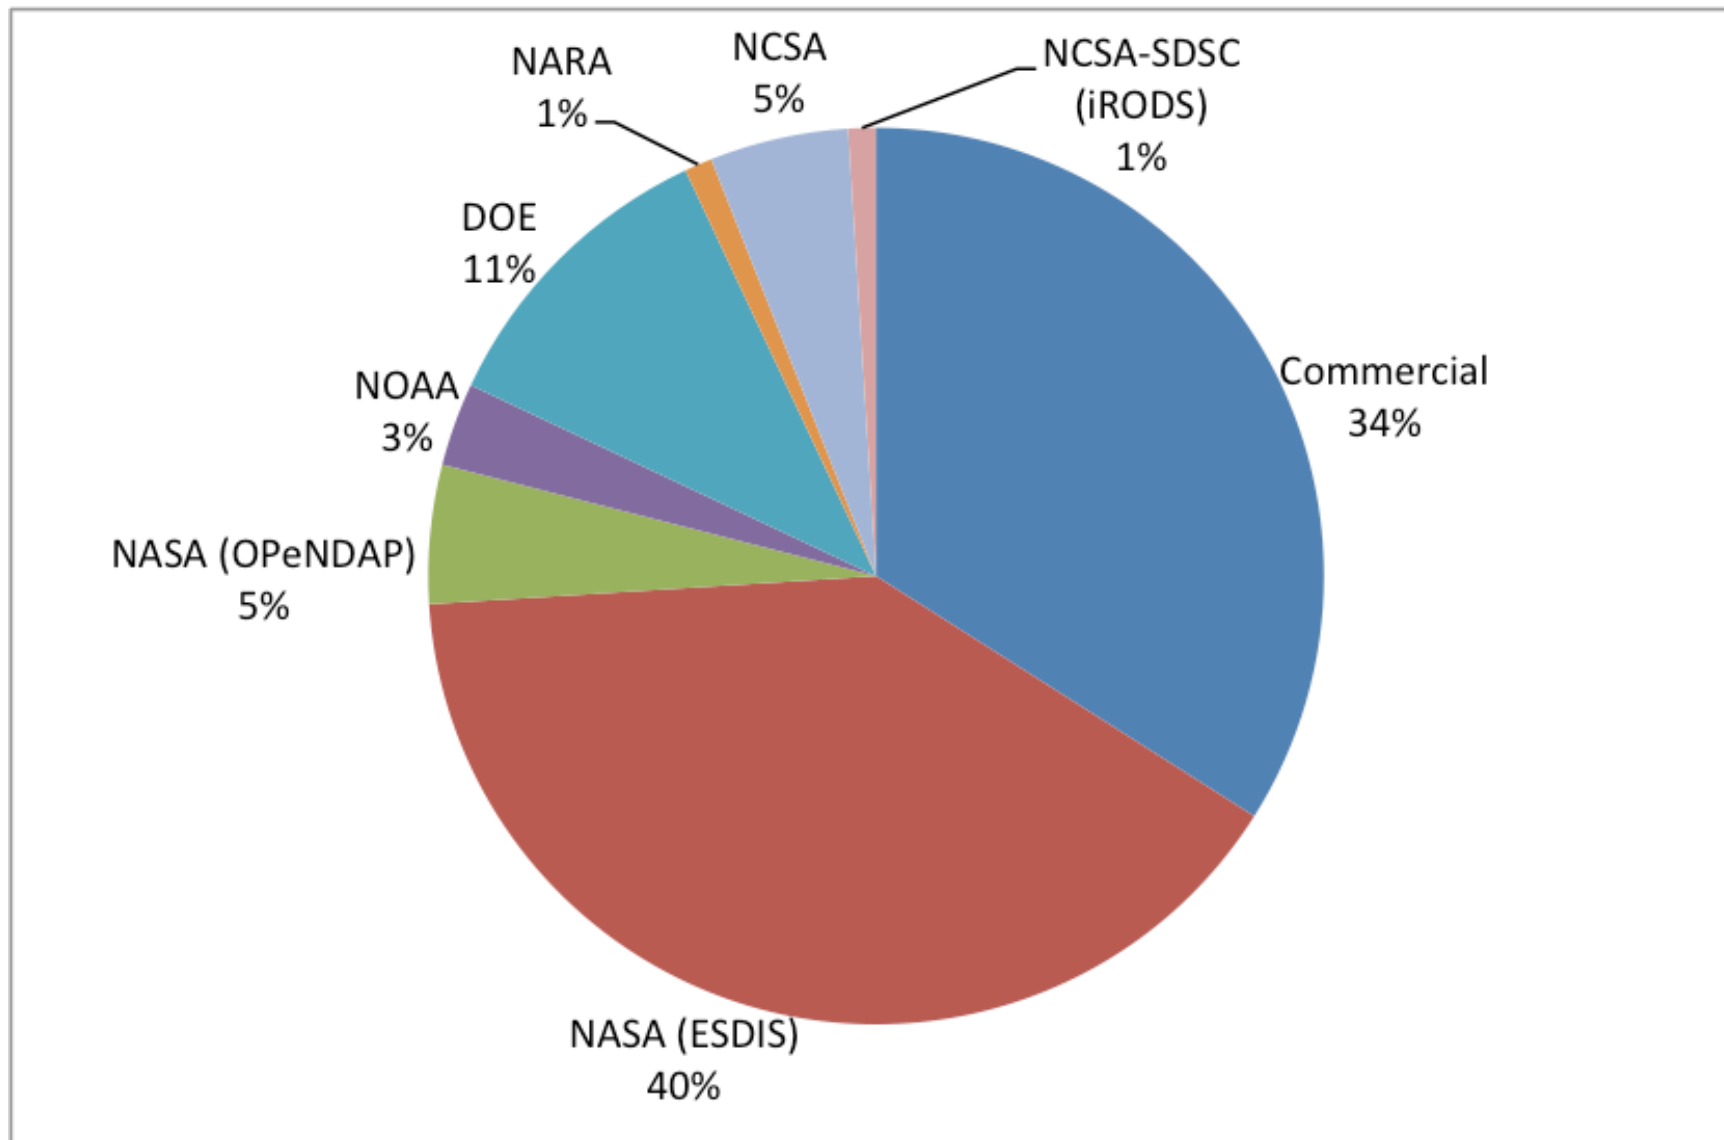


History of The HDF Group

- 18 Years at University of Illinois National Center for Supercomputing Applications
- Spun-off from University July 2006
- Non-profit
- 20+ scientific, technology, professional staff
- Intellectual property:
 - The HDF Group owns HDF4 and HDF5
 - HDF formats and libraries to remain open
 - BSD-type license



Income Profile – past 18 months





HDF5 Road Map





Plans for the next major release

- *Performance* – Improvements “behind the scenes”
- *Ease of Use* - Fortran 2003 features
- *Robustness* - Journalled HDF5 files
- More...



Performance Improvements

- Improved metadata cache.
- Rewrote file free-space manager:
 - Much faster
 - Persistent
- Added two chunk indexing methods:
Fixed & Extensible Arrays
 - Faster
 - Single-Writer/Multi-Reader access to data



Fortran 2003 Features

- New 2003 features allowed us to support
 - Any Fortran INTEGER and REAL type data in HDF5 files
 - Fortran derived types and HDF5 compound datatypes
 - Fortran enumerated types and HDF5 enumerated types
 - HDF5 APIs with callbacks



Journalized HDF5 Files

- Problem:
 - Data in an open HDF5 file susceptible to corruption in the event of an application or system crash.
 - Corruption possible if an open HDF5 file has been updated when the crash occurs.
- Initial Objective:
 - Guarantee an HDF5 file with consistent metadata can be reconstructed in the event of a crash.
 - No guarantee on state of raw data – contains whatever made it to disk prior to crash.



Journalized HDF5 Files, cont.

- Approach: Metadata Journaling
 - When an HDF5 file is opened with Metadata Journaling enabled, a companion journal file is created.
 - When an HDF5 API function that modifies metadata is completed, a transaction is recorded in the journal file before the modified metadata is written to the HDF5 file.
 - If the application crashes, a recovery program can replay the journal by applying in order all metadata writes until the end of the last completed transaction written to the journal file.



HDF5 in the Future

“It’s hard to make predictions, especially about the future” – Yogi Berra



Plans, Guesses, and Speculations

- Improve Multi-threaded Concurrency:
 - Currently thread-safe, but not concurrent
 - Start “pushing global lock down”
 - Use lock-free data structures inside library where possible
- Allow more operations to be Single-Writer / Multiple-Reader (SWMR) safe:
 - Expand range of operations beyond new “extensible array” structure.
 - Allow an MPI application to be the writer



Plans, Guesses, and Speculations

- Improve Parallel I/O Performance:
 - Better leveraging of MPI and file system features
 - Reduce # of I/O accesses for metadata access
 - Allow independent metadata creation operations
- Improve Journalled HDF5 Files:
 - Journal raw data operations
 - Allow "super-transactions" to be created by applications
 - Enable journaling for Parallel HDF5



Plans, Guesses, and Speculations

- Improve raw data chunk cache implementation
- Implement more efficient storage and I/O of variable-length data, including compression
- Work with HPC community to serve their needs:
 - Participating in MPI Forum
 - Focus on high-profile applications or “I/O cores” and remove HDF5 bottlenecks discovered



You tell us!

- Please complete and return the survey so that we can understand your needs
- Questions?