



HDF5: State of the Union

Quincey Koziol, Elena Pourmal
{koziol,epourmal}@hdfgroup.org

The HDF Group

SC10 – HDF5 BOF

November 18, 2010



Data Challenges Addressed by HDF5

- Organize complex collections of data
- Efficient and scalable data storage and access
- Integrates a wide variety of types of data
- The evolution of data technologies
- Long term preservation of data



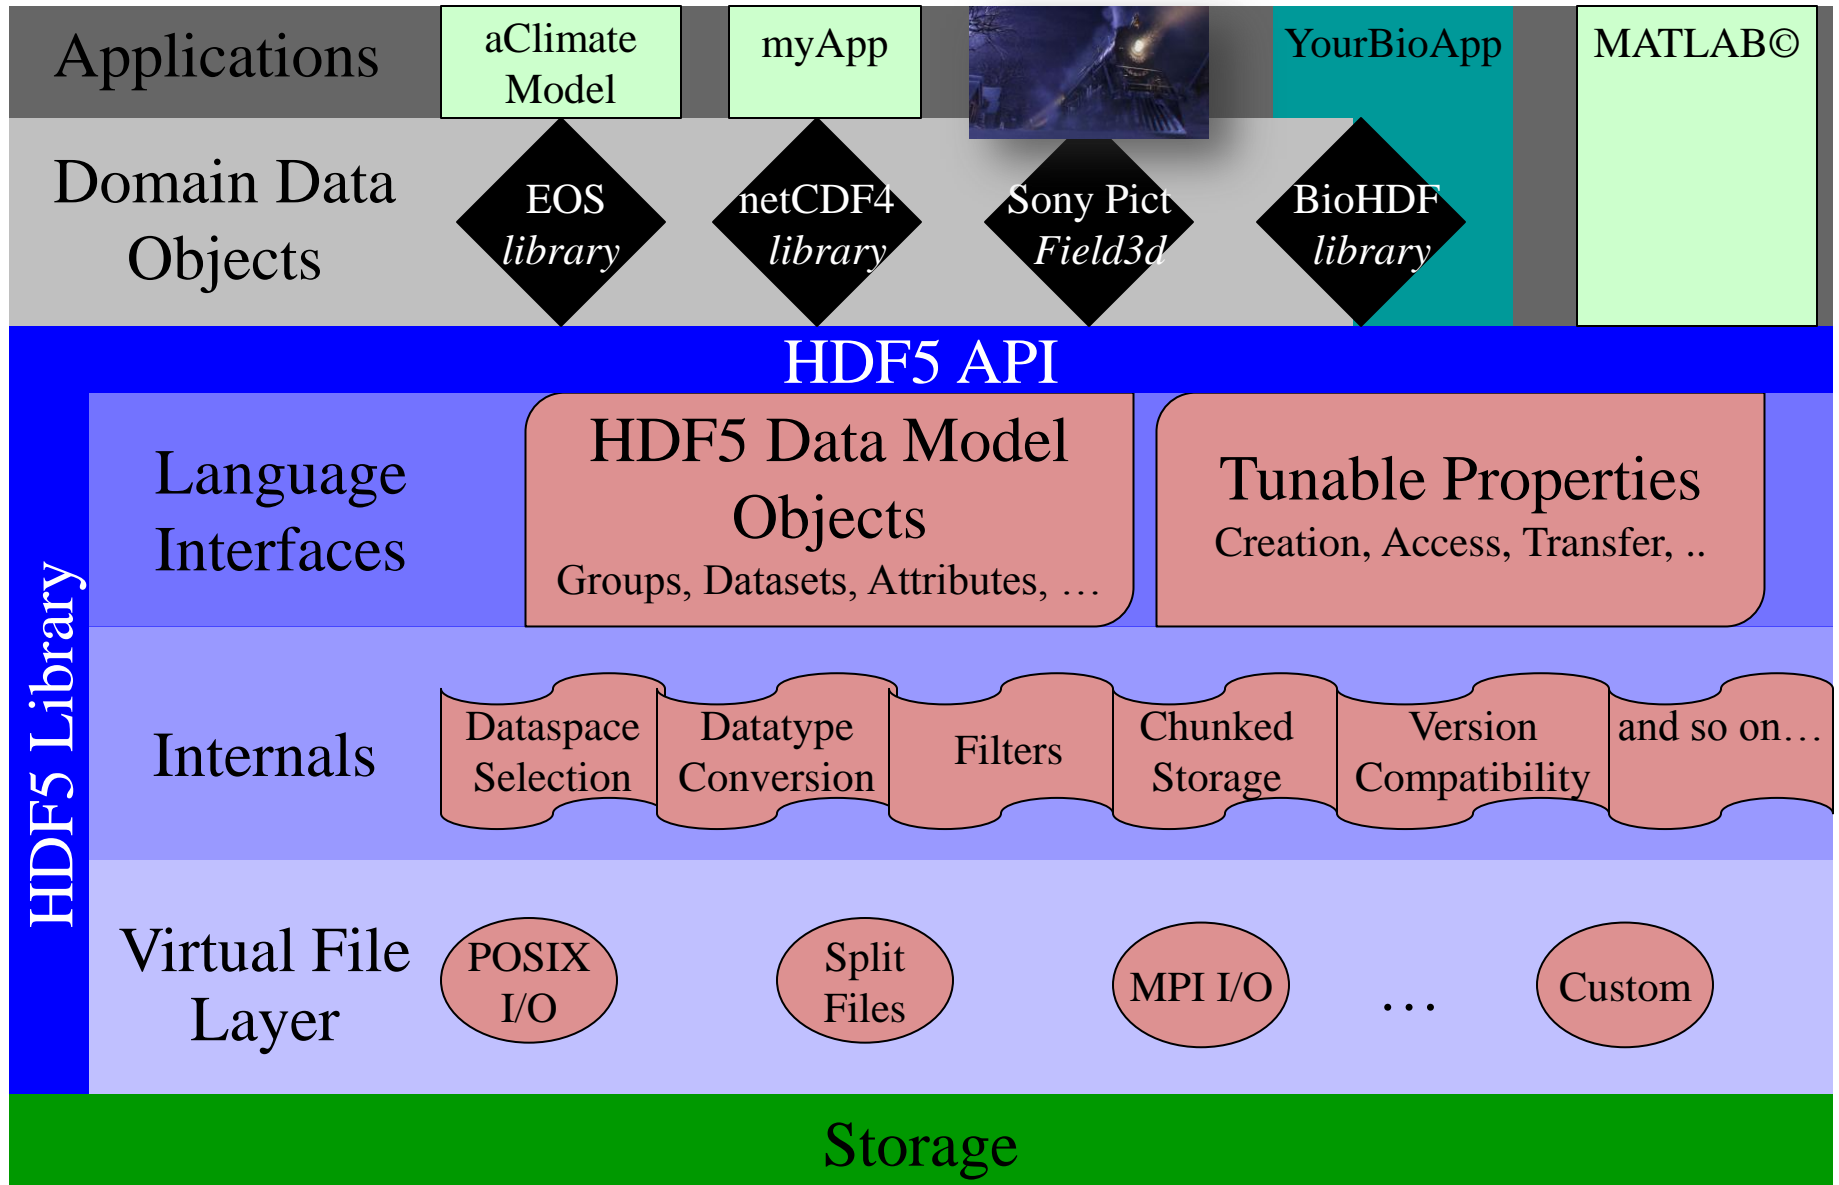
- **HDF5 Abstract Data Model**
 - Defines the “building blocks” for data organization and specification
 - Files, Groups, Links, Datasets, Attributes, Datatypes, Dataspaces

- **HDF5 Software**
 - Tools
 - Language Interfaces
 - HDF5 Library

- **HDF5 Binary File Format**
 - Bit-level organization of HDF5 file
 - Defined by HDF5 File Format Specification



HDF5 API and Applications





Topics

What's up with The HDF Group?

Library Update

Tools update

HDF Java Products

Library development in the works

Other activities



The HDF Group

- Established in 1988
 - 18 years at University of Illinois' National Center for Supercomputing Applications
 - 5 years as independent non-profit company, "The HDF Group"
- The HDF Group owns HDF4 and HDF5
 - Basic HDF4 and HDF5 formats, libraries, and tools are open and free
- Currently employ 33 FTEs



The HDF Group Mission

To ensure long-term accessibility of HDF data through sustainable development and support of HDF technologies.



Goals of The HDF Group

- Maintain and evolve HDF for sponsors and communities that depend on it
- Provide support to the HDF communities through consulting, training, tuning, development, research
- Sustain the company for the long term to assure data access over time



The HDF Group Services

- Helpdesk and Mailing Lists
 - Available to all users as a first level of support: help@hdfgroup.org
- Priority Support
 - Rapid issue resolution and advice
- Consulting
 - Needs assessment, troubleshooting, design reviews, etc.
- Training
 - Tutorials and hands-on practical experience
- Enterprise Support
 - Coordinating HDF activities across departments
- Special Projects
 - Adapting customer applications to HDF
 - New features and tools
 - Research and Development




Members of the HDF support community

- NASA – Earth Observing System
- NOAA/NASA/Riverside Tech – NPOESS
- A large financial institution
- DOE – projects w/LBNL & PNNL, ANL & ORNL
- Lawrence Livermore National Lab
- Army Geospatial Center
- NIH/Geospiza (bio software company)
- Lawrence Berkeley National Lab
- University of Illinois/NCSA
- Sandia National Lab
- A leading U.S. aerospace company
- Projects for petroleum industry, vehicle testing, weapons research, others
- “In kind” support



Income Profile – past 12 months



Total income approximately \$3.4 million



New Directions We're Taking

- Renewed and Re-energized effort in HPC
 - ~\$800K/year in new DOE grants and contracts, focused on HPC improvements to HDF5
- Exploring applications of HDF5 in the Bioinformatics field
 - Small current grant and multiple new grants submitted for Bio applications of HDF5
- Staffing up to explore synthesis of HDF5 and traditional database technology
 - Hired Dr. Gerd Heber to lead effort



Topics

What's up with The HDF Group?

Library Update

Tools Update


HDF Java Products

Library development in the works

Other activities




Repository Statistics





Repository Statistics

HDF5 Library Source Code, Lines of Code by Release





Where We've Been

- Release 1.0
 - First “prototype” release in Oct, 1997
 - Incorporated core data model: datatypes, dataspace & datasets and groups
 - Parallel support added in r1.0.1, in Jan, 1999
- Release 1.2.0 - Oct, 1999
 - Added support for bitfield, opaque, enumeration, variable-length and reference datatypes.
 - Added new ‘h5toh4’ tool
 - Lots of polishing
 - Performance optimizations



Where We've Been

- Release 1.4.0 - Feb, 2001
 - Added Virtual File Driver (VFD) API layer, with many drivers
 - Added 'h4toh5', h5cc tools, XML output to h5dump
 - Added array datatype
 - F90 & C++ API wrappers
 - Performance optimizations
- Release 1.6.0 - July, 2003
 - Generic Property API
 - Compact dataset storage
 - Added 'h5diff', 'h5repack', 'h5jam', 'h5import' tools
 - Performance optimizations




Where We're At Now

- Release 1.8.0 - Feb, 2008
 - Features to support netCDF-4
 - Creation order indexing on links and attributes
 - Integer-to-floating point conversion support
 - NULL dataspace
 - More efficient group storage
 - External Links
 - New H5L (links) and H5O (objects) APIs
 - Shared Object Header Messages
 - Unicode-8 Support
 - Anonymous object creation
 - New tools: 'h5mkgrp', 'h5stat', 'h5copy'
 - CMake build support
 - Performance optimizations



Performance Regression Test Suite





Software Engineering in HDF5

- We spend *considerable* effort in ensuring we produce very high quality code for HDF5.
- Current efforts:
 - Correctness regression testing
 - Nightly testing of >60 configurations on >20 machines
 - Performance regression testing
 - Applying static code analysis - Coverity
 - Memory leak detection – valgrind
 - Code coverage – *coming soon*



- New library features
 - Added support for filtering densely stored groups.
- Tool features:
 - h5dump: Added the new packed bits feature which prints packed bits stored in an integer dataset.
 - h5diff: Added new flag --follow-symlinks.
- Bugs fixed:
 - Fixed various "strict aliasing" problems.
 - Fixed a couple of file corruption bugs.



HDF5 1.8.6 minor release (Nov 10)

- Build:
 - CMake builds supported
- Library:
 - Added support for thread safety on Windows using the Windows threads library.
 - Improved performance of I/O on datasets with the same shape, but different rank.
 - Improved performance of the chunk cache by avoiding unnecessary B-tree lookups of chunks already in cache.
- Parallel I/O:
 - Improved performance of metadata I/O by changing the default algorithm to perform I/O from all processes (instead of just process 0) when using parallel I/O drivers.
- Bugs fixed:
 - Fixed many memory issues that valgrind exposed.
- Tools:
 - h5diff: Added a new flag: --exclude-path.
 - h5ls: Added a new flag --follow-symlinks.



Where We'll Be Soon

- Release 1.10 - Overview
 - May make it into 2011, but no guarantee
 - Stopped adding major features, fleshing out our current efforts now
 - Major Efforts:
 - Improved scalability of chunked dataset access
 - Single-Writer/Multiple Reader (SWMR) Access
 - Improved fault tolerance



Where We'll Be Soon

- Release 1.10 - Details
 - New chunked dataset indexing methods
 - Single-Writer/Multiple-Reader (SWMR) Access
 - Improved Fault Tolerance
 - Journalled Metadata Writing
 - Ordered Updates
 - Persistent file free space tracking
 - Basic support for asynchronous I/O
 - Expanded Virtual File Driver (VFD) interface
 - Lazy metadata writes (in serial)
 - F2003 Support
 - Compressed group information
 - Open file caching for external links
 - Performance optimizations




Where We Might Get To

- Release 1.10 - Maybe?
 - Full C99 type support (long double, complex, boolean types, etc)
 - First steps to improved multi-threaded concurrency
 - High-level “HPC” API
 - Support for runtime file format limits
 - Improved variable-length datatype storage
 - Re-write h5diff tool to perform better and have a more intuitive behavior

What's up with The HDF Group?

Library Update

Tools update



HDF Java Products

Library development in the works

Other activities



Tool activities in the works

- New tool – h5watch
 - Display changes to a dataset, metadata and raw data
- Improved code quality and testing
- Tools library: general purpose APIs for tools
 - Tools library currently only for our developers
 - Want to make it public so that people can use it in their products



What's up with The HDF Group?

Library Update

Tools update

HDF Java Products

Library development in the works

Other activities



HDF-Java 2.7 is on the way

- Full support for [most] features in 1.8.x C library
- Many bug fixes and much better regression tests
- Performance improvements

What's up with The HDF Group?

Library Update

Tools update

HDF Java Products

Library development in the works

Other activities





What Next?

- Release 1.12 - Overview
 - Special focus on HPC features & performance
 - Improvements to data model
 - “Virtual Object Layer”



What Next?

- Release 1.12 - Details
 - Special focus on HPC features & performance
 - Remove “all collective” metadata operation requirement
 - Use “staging area” to buffer and analyze data
 - Ordered updates in parallel
 - “Append-only” optimizations
 - “Embarrassingly parallel” optimizations
 - Asynchronous metadata and raw data I/O in parallel
 - Indexing on both raw data and metadata, along with query language
 - Autotune to parallel file system
 - Support MPI fault tolerance



What Next?

- Release 1.12 - Details
 - Improvements to data model:
 - Shared dataspace
 - Attributes on dataspace and datatypes
 - “Virtual Object Layer”
 - HDF5 without a file?
 - Introduce new layer in library focused on abstract data model
 - Working with Carlos Maltzahn @ UCSC to explore merging HDF5 with Ceph file system



What Next?

- Release 1.12 – Details
 - Further along with multi-threaded concurrency
 - Allow component values of datasets with composite datatypes to be stored in non-interleaved form
 - Port HDF5 to the NMI Build and Test facility
 - More, certainly...

What's up with The HDF Group?

Library Update

Tools update

HDF Java Products

Library development in the works

Other activities




GES DISC

Goddard Earth Sciences Data and Information Services Center



geospiza





NIH STTR with Geospiza, Seattle WA

BIOHDF : TOWARD SCALABLE BIOINFORMATICS INFRASTRUCTURES



Next Generation DNA Sequencing

“Transforms today’s biology”

“Democratizing genomics”

“Changing the landscape”

“Genome center in a mail room”

“The beginning of the end for microarrays”

NGS is Powerful






- ***Goal: Move bioinformatics problems from organizing and structuring data to asking questions and visualizing data***
 - Develop data models and tools to work with NGS data in HDF5
 - Create HDF5 domain-specific extensions and library modules to support the unique aspects of NGS data → BioHDF
 - Integrate BioHDF technologies into Geospiza products
- **Deliver core BioHDF technologies to the community as open-source software**



HDF-EOS5/netCDF-4 Augmentation Tool

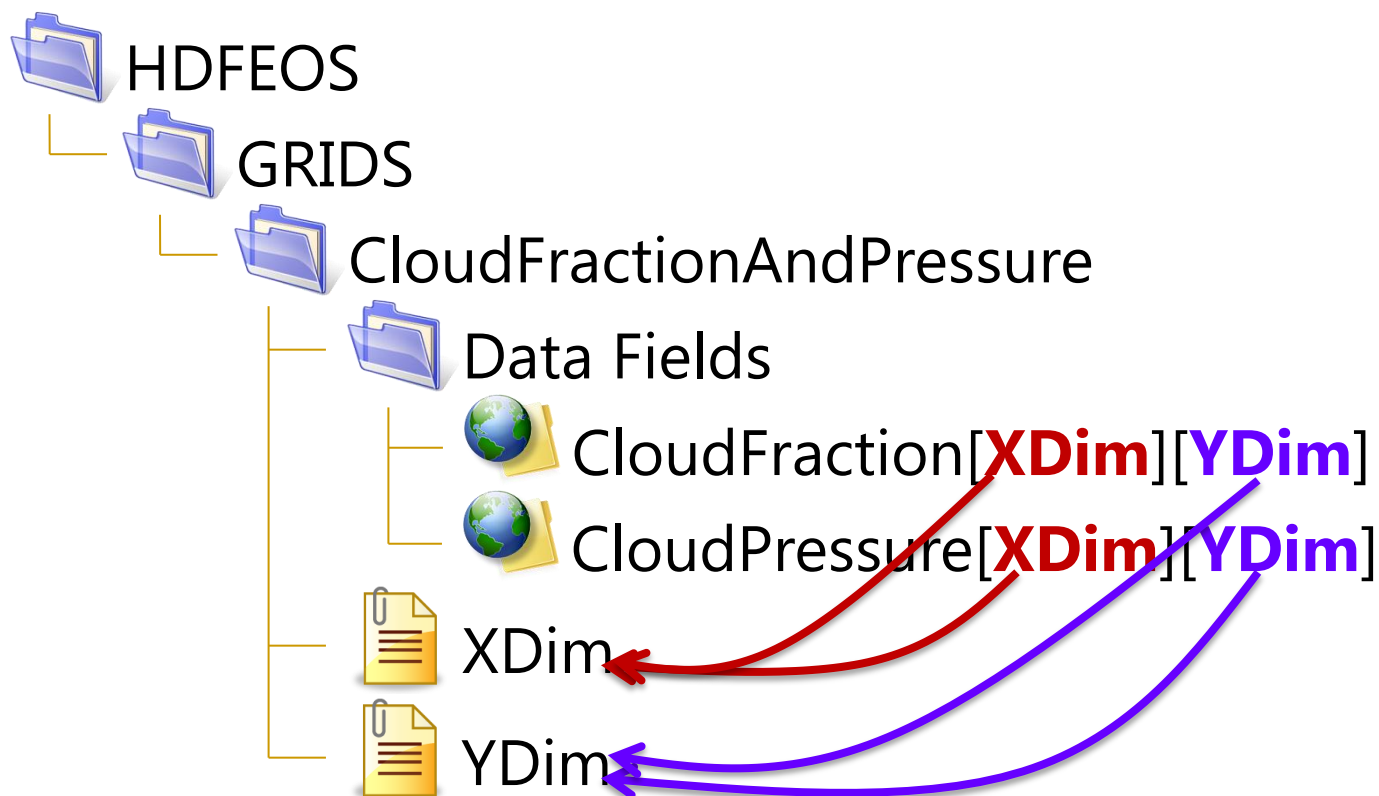
Accessing HDF-EOS5 files via netCDF-4 API

- NetCDF-4 model follows the HDF5 dimension scale model but HDF-EOS5 does not.



No HDF5 dimension scales are associated with this variable

- Provide dimensions required by netCDF-4





The logo for OPeNDAP, featuring a blue star with a red orbital path above the text 'OPeNDAP' in a blue, serif font.



- HDF5-OPeNDAP handler
 - Served OMI Swath data
- HDF4-OPeNDAP handler
 - Tested with some AIRS data and some MODIS data



Thank You!

Questions & Comments?