

HDF Group report to LLNL  
February 2013  
Quincey Koziol

Summary:

During the month of February, 2013 the HDF Group worked on the following tasks:

- Metadata Aggregation (113.6 hours)
- Project Management Tasks (0.4 hours)
- Test HDF5 release on LLNL machines (0.0 hours)
- Support starting core VFD from file image in memory (0.0 hours)
- User Support (0.0 hours)
- Investigate and correct issues reported by Klocwork (0.0 hours)
- Page Buffering (0.0 hours)
- Support “single chunk” indexing method for chunked datasets (0.0 hours)

The **total number of hours** worked is **114.0** hours.

Items of Note:

Vailin has finished the full implementation of the metadata aggregation feature, with backward/forward compatibility tests and support for it in the tools. She’s now returning to the core implementation and making revisions after review & discussion.

New tasks:

During this time period the following tasks were begun:

- *none*

Completed tasks:

During this time period the following tasks or sub-tasks were completed:

- ***Metadata Aggregation***
  - Finished backward/forward compatibility testing
  - Updated tools to support new feature (h5stat, h5dump, h5repack)
  - Updated documentation with changes for new feature

Deferred tasks:

During this time period the following tasks or sub-tasks were deferred:

- *none*

### Tasks in progress:

During this period of time The HDF Group worked on the following tasks:

- ***Metadata Aggregation, Vailin Choi, Quincey Koziol, Frank Baker*** (113.6 hours)
  - Design discussions
  - Finished backward/forward compatibility testing
  - Updated tools to support new feature: h5stat, h5dump, h5repack
  - Updated documentation for file free space usage, new API routines, file format change, and changes to tools.
  - Revising approach to setting property for enabling feature, to simplify it and make it more logical for users. (And propagate change to tests, etc)
- ***Test HDF5 releases on LLNL machines, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***Project Management Tasks, Quincey Koziol*** (0.4 hours)
  - Set up user accounts
  - Planning and reporting activities.
  - User discussions, status telecons & e-mail.
  - Make snapshots, etc.
- ***Support starting core VFD from file image in memory, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***User Support, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***Page Buffering, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***Investigate and correct issues reported by Klocwork, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***Support “single chunk” indexing method for chunked datasets, -*** (0.0 hours)
  - *Nothing to report during this time period*

### Current Projects for People:

- Quincey Koziol:
  - Design & architecture guidance

- Project management
- Vailin Choi:
  - Metadata aggregation design & implementation
  - Page buffering design & implementation
  - “Single chunk” chunked dataset indexing method
- Albert Cheng:
  - Test HDF5 releases on HPC machines
- Mark Evans/Frank Baker:
  - Update HDF5 documentation for new feature(s)

Ongoing tasks for next reporting period:

- ***Metadata Aggregation, Vailin Choi***
  - Revise RFC describing new feature as needed
  - Implement feature.
- ***Page Buffering, Vailin Choi***
  - Write RFC describing new feature
  - Implement feature.
- ***Single Chunk Index Method for Chunked Datasets, Vailin Choi***
  - Second review.
  - Check in to subversion.
- ***Test HDF5 releases on LLNL machines, Albert Cheng***
  - Release testing on LLNL machines (in May & November).
- ***Investigate and correct issues reported by Klocwork tool, Quincey Koziol***
  - Investigate issues reported by Klocwork and correct them.

Deferred/Future tasks:

- ***Scope effort for implementing “stackable” VFDs***
  - Discuss feature and write RFC for allowing VFDs to be “stacked” on top of each other.
- ***Design VFDs to enable poor man’s parallel I/O***
  - Discuss feature and write RFC for VFDs that can improve “Poor Man’s Parallel” I/O on HPC systems.