

HDF Group report to LLNL  
April 2012  
Quincey Koziol

Summary:

During the month of April, 2012 the HDF Group worked on the following tasks:

- Support starting core VFD from file image in memory (20.3 hours)
- Port and test HDF5 on LLNL machines (8.0 hours)
- Misc. Admin Tasks (1.8 hours)
- Metadata Aggregation (0.0 hours)
- Page Buffering (0.0 hours)
- Support “single chunk” indexing method for chunked datasets (0.0 hours)
- Collaborations w/LANL developers about stackable VFD ideas (0.0 hours)
- Investigate and correct issues reported by Klocwork tool (0.0 hours)

The **total number of hours** worked is **30.1** hours.

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:

- ***Support starting core VFD from file image in memory***
  - Finished documentation and merged code into repositories, for inclusion in the 1.8.9 release.

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:

- ***Support starting core VFD from file image in memory, Frank Baker, Quincey Koziol*** (20.3 hours)
  - Finished documentation for new routines and mode of operation.
  - Merged feature to trunk and 1.8 release branch.
- ***Port and test HDF5 on HPC machines, Albert Cheng*** (8.0 hours)
  - HDF5 1.8.9 release testing on LLNL systems
  - Work on Silo testing issues
- ***Miscellaneous Admin Tasks, Quincey Koziol, Albert Cheng*** (1.8 hours)
  - Set up user accounts
  - Planning and reporting activities.
  - User discussions, status telecons & e-mail.
  - Make snapshots, etc.
- ***Page Buffering, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***Metadata Aggregation, -*** (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*
- ***Collaborations w/LANL developers about stackable VFD ideas, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***Investigate and correct issues reported by Klocwork tool, -*** (0.0 hours)
  - *Nothing to report during this time period*

### Current Projects for People:

- Quincey Koziol:
  - Design & architecture guidance
  - Project management
  - Metadata aggregation and Page buffering design
  - “stackable” VFD design, implementation and collaborations w/LANL developers
  - Investigate issues reported by Klocwork tool
  - Design VFDs to enable poor man’s parallel I/O
- Albert Cheng:
  - Port and test HDF5 on HPC machines
- Mark Evans/Frank Baker:
  - Update HDF5 documentation for new feature(s)

- Vailin Choi:
  - “Single chunk” chunked dataset indexing method
- Jacob Gruber
  - Prototype page buffering implementation

#### Ongoing tasks for next reporting period:

- ***Metadata Aggregation, Vailin Choi***
  - Update RFC describing new feature
  - Implement feature.
- ***Page Buffering, Jacob Gruber***
  - Gather requirements, use cases and goals of project
  - Write RFC describing new feature
  - Implement feature.
- ***Single Chunk Index Method for Chunked Datasets, Vailin Choi***
  - Second review.
  - Check in to subversion.
- ***Port and test HDF5 on LLNL machines, Albert Cheng***
  - Release testing on LLNL machines.
- ***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.