

HDF Group report to LLNL  
February 2012  
Quincey Koziol

Summary:

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

- Metadata Aggregation (92.1 hours)
- Support starting core VFD from file image in memory (42.1 hours)
- Page Buffering (31.6 hours)
- Misc. Admin Tasks (3.8 hours)
- Port and test HDF5 on LLNL machines (0.0 hours)
- Parallel performance benchmark tool (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 **169.6** 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:

- ***Metadata Aggregation***
  - Finished next revision of RFC for feature.
- ***Support starting core VFD from file image in memory***
  - Finished documentation for 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, John Mainzer, Frank Baker*** (92.1 hours)
  - Revised and updated RFC for feature:
    - Investigate & document free space manager design
    - Update overall design
    - Send to Quincey for review
- ***Support starting core VFD from file image in memory, Mark Evans, John Mainzer, Frank Baker*** (42.1 hours)
  - Update documentation for new routines and mode of operation.
- ***Page Buffering, Jacob Gruber, John Mainzer, Quincey Koziol*** (31.6 hours)
  - Updated RFC for “NULL” VFD
  - Coding & testing “NULL” VFD
  - Review/revise ‘sb\_verify’ RFC
  - Update ‘sb\_verify’ implementation
- ***Miscellaneous Admin Tasks, Quincey Koziol, Albert Cheng, Jacob Gruber*** (3.8 hours)
  - Set up user accounts
  - Planning and reporting activities.
  - User discussions, status telecons & e-mail.
  - Make snapshots, etc.
- ***Port and test HDF5 on HPC machines, -*** (0.0 hours)
  - *Nothing to report during this time period*
- ***Parallel performance benchmark tool, -*** (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
- 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
- John Mainzer:
  - 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
- Ruth Aydt:
  - Parallel performance benchmarking tool
- Jacob Gruber
  - Prototype page buffering implementation

Ongoing tasks for next reporting period:

- ***Enable starting “core” VFD from file image, John Mainzer, Christian Chilan***
  - Final code review.
  - Merge to 1.8 release branch.
  - Update documentation to describe new feature.
- ***Single Chunk Index Method for Chunked Datasets, Vailin Choi***
  - Second review.
- ***Parallel performance benchmark tool, Ruth Aydt***
  - *On hold, pending resolution of funding for next year*
  - Gather requirements, use cases and goals of project
  - Write RFC describing new tool
  - Implement tool.
- ***Metadata Aggregation, John Mainzer***
  - Gather requirements, use cases and goals of project
  - Write RFC describing new feature
  - Implement feature.
- ***Page Buffering, John Mainzer, Jacob Gruber***
  - Gather requirements, use cases and goals of project
  - Write RFC describing new feature
  - Implement feature.

- ***Port and test HDF5 on LLNL machines, Albert Cheng***
  - Maintain daily testing on LLNL machines.
  - Stand up daily testing on INL machines.
  - Stand up Silo testing on HDF Group and LLNL machines.
  - Investigate and add tests for “poor man’s parallel” I/O to HDF5 regression test suite.
- ***Investigate and correct issues reported by Klocwork tool, John Mainzer***
  - 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.