

HDF Group report to LLNL  
October 2012  
Quincey Koziol

Summary:

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

- Metadata Aggregation (119.9 hours)
- Test HDF5 release on LLNL machines (17.8 hours)
- Project Management Tasks (1.0 hours)
- Support starting core VFD from file image in memory (0.9 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 **139.6** hours.

Items of Note:

Vailin’s continued her work on implementing the metadata aggregation feature and is making good progress. We’re estimating that she’ll be done around the end of the year and that we’ll have enough funding to pursue the “page buffering” feature, which complements it.

Unfortunately, Jacob’s gotten very busy with his schoolwork and doesn’t have the time to continue working for us any longer, so won’t be included in future projects.

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:

- ***Test HDF5 releases on LLNL machines***
  - HDF5 1.8.10 release testing on udawn & Aztec.

## 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*** (119.9 hours)
  - Implementing/refactoring code+tests for feature
  - Design discussions
- ***Test HDF5 releases on LLNL machines, Albert Cheng*** (17.8 hours)
  - HDF5 1.8.10 release testing on udawn & Aztec.
- ***Project Management Tasks, Quincey Koziol, Albert Cheng*** (1.0 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, Frank Baker*** (0.9 hours)
  - Updates to reference manual entries for new routines.
- ***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.