

HDF Group report to LLNL
November 2012
Quincey Koziol

Summary:

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

- Metadata Aggregation (104.0 hours)
- Test HDF5 release on LLNL machines (6.3 hours)
- Project Management Tasks (0.5 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 **110.8** hours.

Items of Note:

Vailin’s continued her work on implementing the metadata aggregation feature and is still making good progress. We continue to estimate 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.

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:

- *none*

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*** (104.0 hours)
 - Design discussions
 - Implementing/refactoring code+tests for feature
 - Refactor implementation somewhat, after design discussions
 - Create superblock header message for tracking aggregation settings persistently
- ***Test HDF5 releases on LLNL machines, Albert Cheng, Quincey Koziol*** (6.3 hours)
 - Second round of HDF5 1.8.10 release testing on udawn & Aztec.
 - Fix bug in h5stat test script.
- ***Project Management Tasks, Quincey Koziol*** (0.5 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. 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.