

HDF Group report to LLNL
August 2012
Quincey Koziol

Summary:

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

- Metadata Aggregation (78.4 hours)
- Page Buffering (67.0 hours)
- User Support (4.7 hours)
- Project Management Tasks (1.7 hours)
- Investigate and correct issues reported by Klocwork (1.1 hours)
- Support starting core VFD from file image in memory (0.0 hours)
- Support “single chunk” indexing method for chunked datasets (0.0 hours)

The **total number of hours** worked is **152.9** hours.

Items of Note:

Vailin's has concluded her task on another project and she has begun work on the metadata aggregation task in August.

To facilitate progress on addressing issues reported by Klocwork, we've set up a branch in our subversion repository focused on testing and changes to HDF5 that address Klocwork issues. This will enable off-site collaborators, like Mark, to more actively participate in working together with HDF Group developers on this task.

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, Quincey Koziol*** (78.4 hours)
 - Design discussions about how to implement the aggregation feature
 - Write/review/revise RFC describing feature
 - Estimate work for feature and create project plan
 - Begin work on implementing feature
- ***Page Buffering, Jacob Gruber*** (67.0 hours)
 - Wrapping up “null” VFD work
 - Prototyping page buffering feature
- ***User Support, Quincey Koziol*** (4.7 hours)
 - Investigate odd HDF5 behavior that Mark reported (several feature improvements to HDF5 filed as a result)
 - Investigate MPI hanging issue reported by Richard Hedges (appears to be in the MPI implementation, not HDF5)
- ***Project Management Tasks, Quincey Koziol*** (1.7 hours)
 - Set up user accounts
 - Planning and reporting activities.
 - User discussions, status telecons & e-mail.
 - Make snapshots, etc.
- ***Investigate and correct issues reported by Klocwork, Quincey Koziol*** (1.1 hours)
 - Reviewed Mark’s changes and merged them into the main Hdf5 repository
- ***Support starting core VFD from file image in memory, -*** (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

- Albert Cheng:
 - Test HDF5 releases on HPC machines
- Mark Evans/Frank Baker:
 - Update HDF5 documentation for new feature(s)
- Neil Fortner:
 - Help w/code reviews & project management, supervising Jacob
- Vailin Choi:
 - Metadata aggregation design & implementation
 - “Single chunk” chunked dataset indexing method
- Jacob Gruber
 - Page buffering design & implementation

Ongoing tasks for next reporting period:

- ***Metadata Aggregation, Vailin Choi***
 - Revise RFC describing new feature as needed
 - Implement feature.
- ***Page Buffering, Jacob Gruber***
 - 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.