

HDF Group report to LLNL
January 2012
Quincey Koziol

Summary:

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

- Support starting core VFD from file image in memory (106.1 hours)
- Page Buffering (83.1 hours)
- Metadata Aggregation (78.0 hours)
- Visit LLNL (8.2 hours)
- Port and test HDF5 on LLNL machines (1.5 hours)
- Misc. Admin Tasks (0.3 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 **277.2** 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:

- ***Page Buffering***
 - Created RFC for new ‘sb_verify’ VFD callback.
- ***Visit LLNL***
 - Give presentation about HDF5, talk to LLNL developers about HDF5 use and discuss future development.

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, Mark Evans, John Mainzer, Frank Baker, Quincey Koziol*** (106.1 hours)
 - Minor revisions on RFC
 - Update documentation for new routines and mode of operation.
- ***Page Buffering, Jacob Gruber, John Mainzer, Quincey Koziol*** (83.1 hours)
 - Writing RFC for “NULL” VFD
 - Coding & testing “NULL” VFD
 - Determined VFD layer is missing callback routine required for correctly detecting VFD that created the file. Created RFC describing new callback (‘sb_verify’)
 - Review/revise ‘sb_verify’ RFC
- ***Metadata Aggregation, John Mainzer*** (78.0 hours)
 - Revising and updating RFC for feature
- ***Visit LLNL, Quincey Koziol*** (8.2 hours)
 - Discussions w/Mark @ LLNL, including presentations, etc.
 - Discuss Oracle support w/NIF developers in followup telecon
- ***Port and test HDF5 on HPC machines, Albert Cheng*** (1.5 hours)
 - Working on Silo weekly testing with FTP fetch of Silo tarball, on koala @ THG and aztec @ LLNL
- ***Miscellaneous Admin Tasks, Quincey Koziol*** (0.3 hours)
 - Set up user accounts
 - Planning and reporting activities.
 - User discussions, status telecons & e-mail.
 - Make snapshots, etc.
- ***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.