

HDF Group report to LLNL
December 2011
Quincey Koziol

Summary:

During the month of December, 2011 the HDF Group worked on the following tasks:

- Support starting core VFD from file image in memory (133.3 hours)
- Page Buffering (16.0 hours)
- Port and test HDF5 on LLNL machines (5.5 hours)
- Misc. Admin Tasks (4.5 hours)
- Metadata Aggregation (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 **159.3** 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:

- ***Support starting core VFD from file image in memory***
 - Wrap up development 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:

- ***Support starting core VFD from file image in memory, John Mainzer, Christian Chilan, Quincey Koziol*** (133.3 hours)
 - Minor revisions on RFC
 - Wrapped up testing for high-level API routines for feature
 - Code review and revision cycle
 - Design discussions
 - Push code through daily testing regime, tracking down and fixing several errors on the way
- ***Page Buffering, Jacob Gruber*** (16.0 hours)
 - Building skeleton for page buffering VFD, by creating “NULL” VFD
- ***Port and test HDF5 on HPC machines, Albert Cheng*** (5.5 hours)
 - Working on Silo weekly testing with FTP fetch of Silo tarball, on koala @ THG and aztec @ LLNL
- ***Miscellaneous Admin Tasks, Quincey Koziol, Albert Cheng, John Mainzer*** (4.5 hours)
 - Set up user accounts
 - Planning and reporting activities.
 - User discussions, status telecons & e-mail.
 - Make snapshots, etc.
- ***Metadata Aggregation, -*** (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
- 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***
 - Deliver snapshot.
 - 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.