

HDF Group report to LLNL
June 1 – June 30, 2011
Quincey Koziol

Summary:

During the period from June 1 to June 30, 2011 the HDF Group worked on the following tasks:

- Support starting core VFD from file image in memory (63.3 hours)
- Support “single chunk” indexing method for chunked datasets (35.3 hours)
- Port and test HDF5 1.8.x releases on LLNL machines (18.8 hours)
- Misc. Admin Tasks (11.2 hours)

The **total number of hours** worked is **128.6** 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 “single chunk” indexing method for chunked datasets***
 - Finished implementation in library, now verifying with new regression tests.
- ***Port and test HDF5 on LLNL machines***
 - Now have automated daily testing working correctly for the HDF5 1.8 branch and trunk on dawndev, including use of zlib and szlib.

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, Quincey Koziol*** (63.3 hours)
 - Design discussions.
 - Several review/revise iterations on RFC.
- ***Support “single chunk” indexing method for chunked datasets, Vailin Choi*** (35.3 hours)
 - Implement changes to HDF5 library.
 - Add new tests to regression test suite to cover new feature.
 - Debug misc. failures.
- ***Port and test HDF5 on LLNL machines, Albert Cheng*** (18.8 hours)
 - Set up szlib & zlib on dawndev, for use w/HDF5 installations.
 - Installed public-use HDF5 binaries (in /usr/global/tools/hdf5dev).
 - Setting up automated daily testing w/crontab on dawndev.
 - Track daily test failures and investigate as necessary.
- ***Miscellaneous Admin Tasks, Ruth Ayt, Quincey Koziol, Albert Cheng*** (11.2 hours)
 - Set up user accounts (at LLNL and INL)
 - Planning and reporting activities.
 - User discussions, status telecons & e-mail.
 - Make snapshots, etc.

Current Projects for People:

- Quincey Koziol:
 - Design & architecture guidance
 - Project management
- Albert Cheng:
 - Port and test HDF5 on LLNL machines
- Vailin Choi:
 - Adding “single chunk” chunked dataset indexing method
- John Mainzer:
 - Enable starting “core” VFD from file image buffer
 - “stackable” VFD design and implementation
 - Design VFDs to enable poor man’s parallel I/O

Ongoing tasks for next reporting period:

- ***Enable starting “core” VFD from existing buffer, John Mainzer***
 - Write and circulate RFC for adding feature to library.

- Implement feature.
- ***Single Chunk Index Method for Chunked Datasets, Vailin Choi***
 - Scope effort for adding feature to library.
 - Implement feature.
- ***Port and test HDF5 on LLNL machines, Albert Cheng***
 - Stand up daily testing on LLNL machines.
 - Investigate and add tests for “poor man’s parallel” I/O to HDF5 regression test suite.

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.