HDF Group report to LLNL
May 1 – May 31, 2011
Quincey Koziol

Summary:

During the period from May 1 to May 31, 2011 the HDF Group worked on the following tasks:

- Port and test HDF5 1.8.x releases on LLNL machines (75.6 hours)
- Support starting core VFD from file image in memory (64.2 hours)
- Support "single chunk" indexing method for chunked datasets (19.8 hours)
- Misc. Admin Tasks (8.2 hours)

The **total number of hours** worked is **167.8** hours.

New tasks:

During this time period the following tasks were begun:

- ***Support "single chunk" indexing method for chunked datasets.***
  - Task to reduce overhead when entire dataset fits into one chunk. Instead of creating a chunk index when the dataset will only ever have a single chunk (i.e. it's fixed dimensions and the chunk dimensions are equal to the dataset dimensions), just point to the chunk directly from the dataset's object header.

- ***Support starting core VFD from file image in memory.***
  - Task to allow application to "open" an HDF5 file image in memory, primarily with the core VFD. RFC located here: http://visitbugs.ornl.gov/attachments/66/load_core_file_driver_from_image_RFC_v04.pdf

Completed tasks:

During this time period the following tasks or sub-tasks were completed:

- ***Support starting core VFD from file image in memory***
  - Finished initial draft RFC for feature and circulated with users, for feedback.

Deferred tasks:

During this time period the following tasks or sub-tasks were deferred:

- ***Stackable VFD support***
  - After discussions w/Mark Miller about the size and scope of this effort, we decided to deferring further work on it until some smaller, easier tasks were completed, as well as a stronger use-case could be made for the stackable VFD feature.

Tasks in progress:

During this period of time The HDF Group worked on the following tasks:

- ***Port and test HDF5 1.8.x releases on LLNL machines, Albert Cheng*** (75.6 hours)
  - Testing 1.8.7 release,1.8 branch, and trunk on dawndev.
  - Added support for C++ and FORTRAN in the 1.8 branch, on dawndev.
  - Set up automated Silo testing on HDF Group machines.
  - Set up automated Silo testing on Zeus cluster @ LLNL.

- ***Support starting core VFD from file image in memory, John Mainzer, Quincey Koziol*** (64.2 hours)
  - Design discussions.
  - Write & revise RFC, after more discussions.

- ***Support "single chunk" indexing method for chunked datasets, Vailin Choi*** (19.8 hours)
  - Create more detailed time estimate for task.
  - Study code and plan design.
  - Implement changes to HDF5 library.

- ***Miscellaneous Admin Tasks, Quincey Koziol, Albert Cheng*** (8.2 hours)
  - 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.