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

Summary:

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

- Support starting core VFD from file image in memory (31.1 hours)
- Parallel performance benchmark tool (17.3 hours)
- Port and test HDF5 1.8.x releases on LLNL machines (13.8 hours)
- Support "single chunk" indexing method for chunked datasets (13.7 hours)
- Misc. Admin Tasks (8.8 hours)
- Metadata Aggregation and Page Buffering (7.4 hours)

The **total number of hours** worked is **92.1** hours.

New tasks:

During this time period the following tasks were begun:

- ***Parallel performance benchmark tool***
  - Design and implement a "next generation" parallel I/O performance benchmark tool, to measure POSIX, MPI-I/O and HDF5 performance (at least). Details at: https://visitbugs.ornl.gov/projects/hpc-hdf5/wiki/Parallel_Benchmarking

- ***Metadata Aggregation***
  - Aggregate metadata within HDF5 files into large, page-aligned blocks within the file, with a persistent mechanism.

- ***Page Buffering***
  - Page buffering will provide an additional buffering mechanism within HDF5, to allow aggregated metadata to be accessed in the file in larger I/O operations.

Completed tasks:

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

- ***Support starting core VFD from file image in memory***

- Finished RFC for feature and circulated for final comments before implementation begins.

- ***Support "single chunk" indexing method for chunked datasets***
  - Finished regression tests, now cleaning up to prepare for code review.

- ***Port and test HDF5 on LLNL machines***
  - Improved SILO daily testing to show failures in Subject line of status emails.

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*** (31.1 hours)
  - Design discussions.
  - Several review/revise iterations on RFC.
  - Sent final draft of RFC out for broad review.

- ***Parallel performance benchmark tool, Ruth Aydt*** (17.3 hours)
  - Reviewed project notes and other reference materials on wiki, highlighting things related to the benchmarking project.
  - Downloaded silo 4.8 tarball and unpacked. Looked briefly at the ioperf and pmpio codes.
  - Asked some questions to help clarify purpose(s) of benchmark effort at telecon 7/20 and also sent email w/ word doc of questions re: terms, etc. in ref materials reviewed. This "conversation" continues on the project wiki.
  - Listened to Mark's SILO podcast at his suggestion (very helpful) and also looked at Silo User's manual.
  - Getting a better handle on what is there, what terms mean (and how they are overloaded). More questions & conversation, notes are on wiki @ https://visitbugs.ornl.gov/projects/hpc-hdf5/wiki/Parallel_Benchmarking

- ***Port and test HDF5 on LLNL machines, Albert Cheng*** (13.8 hours)
  - Moving daily tests from Zeus to Sierra (and explored using Aztec, which won't work)
  - Improved SILO daily testing to show failures in Subject line of status emails

- ***Support "single chunk" indexing method for chunked datasets, Vailin Choi*** (13.7 hours)
    - Debug misc. failures.
    - Cleaning up code, for review

- ***Miscellaneous Admin Tasks, Ruth Aydt, Quincey Koziol, Albert Cheng, John Mainzer*** (8.8 hours)
    - Set up user accounts (at LLNL and INL)
    - Planning and reporting activities.
    - User discussions, status telecons & e-mail.
    - Make snapshots, etc.

- ***Metadata Aggregation and Page Buffering, John Mainzer, Quincey Koziol*** (7.4 hours)
    - Begin design discussions and start putting use cases and requirements down.

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:
    - Metadata aggregation and Page buffering design
    - "stackable" VFD design and implementation
    - Design VFDs to enable poor man's parallel I/O
- Ruth Aydt:
    - Parallel performance benchmarking tool
- Jacob Gruber
    - Enable starting "core" VFD from existing buffer: implement "load image" feature, from RFC

Ongoing tasks for next reporting period:

- ***Enable starting "core" VFD from existing buffer, John Mainzer, Jacob Gruber***
    - Update RFC for adding feature to library, based on any feedback received.
    - Implement feature.

- ***Single Chunk Index Method for Chunked Datasets, Vailin Choi***
    - Implement feature.

- ***Parallel performance benchmark tool, Ruth Aydt***
    - 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***
    - Gather requirements, use cases and goals of project
    - Write RFC describing new feature
    - 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.