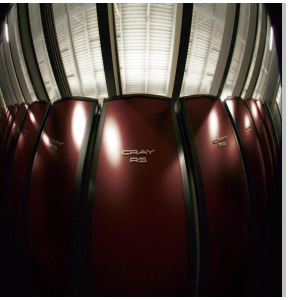

Metadata Management to Support Scientific Inquiry

Jay Lofstead

Sandia National Laboratories

Exceptional service in the national interest

HUG 23 8-17-2023

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Attributes vs. Metadata

Sandia National Laboratories

- Metadata definition: Any data about the data to aid scientific inquiry

- HDF5 attributes attached to the hierarchy
- Metadata may be
  - HDF5 attribute
  - whole variable
  - Ideally: attached to a variable region, timestep, run or even campaign

- Question:

  How do we leverage Rich Metadata to aid scientific inquiry?

3 Existing Tool Generations

Sandia National Laboratories

Generations refer to classes and development order rather than obsolescence!

- 1st generation: file level metadata, domain specific tools
- 2nd generation: raw data level indexing
- 3rd generation: data-level tagging

Provenance-related tools are a completely separate discussion (exercise left to the reader)

First Generation Tools (Storage Sys)

Sandia National Laboratories

- POSIX extended attributes
  - Implementation and limitations varies by file system
  - File system may have to be mounted differently to enable them
  - Key-value pairs
  - Limited to block size for portability (1024 bytes for maximum compatibility)
  - Tools like `tar` may require special flags to save, if supported at all
- HPSS
  - Implements POSIX extended attributes
  - Stores all metadata in a RDBMS outside the tape archive for performant queries
  - Can export data for use in systems like Starfish (see next)

First Generation Tools (Abstracted)

Sandia National Laboratories

- Starfish (https://starfishstorage.com/)
  - File level arbitrary metadata
  - Strong query capabilities
  - Integration with HPSS (possible, if not standard)
  - Works across storage systems
  - Depends on files being in a fixed location or location updated as files move in storage
- JAMO – Joint Genome Institute (JGI) Archive and Metadata Organizer
  - One example of a domain specific tool
  - Repository for data, metadata, and provenance with full query API
  - Standard + User defined Templates for metadata
  - Can tag at a more detailed level as well

Second Generation Tools (indexing)

Sandia National Laboratories

Raw data indexing

- Exact value or binning

- FastBit
  - Make Bloom-filter style index on data.
  - Bins for value ranges with a 1/0 to indicate present/absent
  - Query capabilities: Does a value exist in this file/dataset?
    - Yes/no
    - No ability to know where

Second Generation Tools (custom db)

Sandia National Laboratories

- SciDB
  - Built for astronomy images and other 2-D data sets originally
  - Multi-dimensional array data model
  - Indexing capabilities
  - SQL front-end
  - Query for data values or thresholds

Second Generation Tools (IO libs)

Sandia National Laboratories

All offer an attribute capability, but it varies a bit

- HDF5
  - Attach to any part of the hierarchy
  - Walk tree to search
- NetCDF
  - Attach to a variable for things like units
  - Flat hierarchy makes it a linear search
- ADIOS
  - Attach to any part of the hierarchy
  - Separate attribute index for fast searching

Third Generation Tools (features)

Sandia National Laboratories

- Associate tags data/parts of data
  - Region/var/timestep/run
  - Bounding box, simple tag
- Solutions use key-value and RDBMS approaches for different tradeoffs
- Key-value
  - Flat hierarchy with all encoding in key name
  - O(1) to get next item adjacent in sorted keys
  - O(n) to find something that is not exactly matching key

Third Generation Tools (key-value)

Sandia National Laboratories

- SoMeta (LBL)
  - Developed as a way to see how to use a KV store to optimize data access for object stores
  - Encode tag information into key and use value for data location
  - Key searching speed based on key storage system
- TagIt (ORNL)
  - Integrated into storage system
  - Assuming distributed, shared nothing storage
  - Works like a distributed DB index placing index next to data for faster access

Third Generation Tools (RDBMS)

Sandia National Laboratories

- Biomedical Image Metadata Manager (BIMM)
  - Image database with tags
  - Search for semantic features
  - Domain specific, but is really just image tagging capability
- EMPRESS
  - Embedded database(s) for tagging
  - Run, timestep, var, and hyperslab tagging
  - Flexible query API + ability to use SQL directly on database
  - Per process/node/job database granularity
  - In memory, hybrid, on persistent storage
  - Distributed shared nothing
  - Use logical locations to be file/object format independent

State of the Art Summary

Sandia National Laboratories

- Long term data archiving
  - Which data set(s) contain what I want to study?
- Short/medium term data set identification
  - What does each set contain and which ones to save?
- Provenance largely focuses on environment
  - Critical, but insufficient/inefficient/impractical/fragile

How do we look for complex data features we didn't think to annotate earlier?

Fourth Generation Tool

Sandia National Laboratories

Coeus project from US DOE ASCR funding

- Derived quantity information-based tagging
  - For a climate model, where is the pressure gradient greater than a particular value
- Problems!
  - Derived quantities take time to compute
  - Derived quantities can be as large as the original data

Query Examples

- Select data from ‘CFD run’ where magnitude (gradient) >= 10
- Select data from ‘PIC run’ where vorticity != 0
- Select data from ‘combustion campaign 2023’ where ‘ignition timestep’ < 50

Integrated Solution

- Query front end
  - ADIOS (demonstration)
- Storage management
  - (Hermes)
- Metadata Management
  - (Empress++)

- When/how to pre-stage data
- When to calculate vs. store metadata
- Focus on session rather than individual query

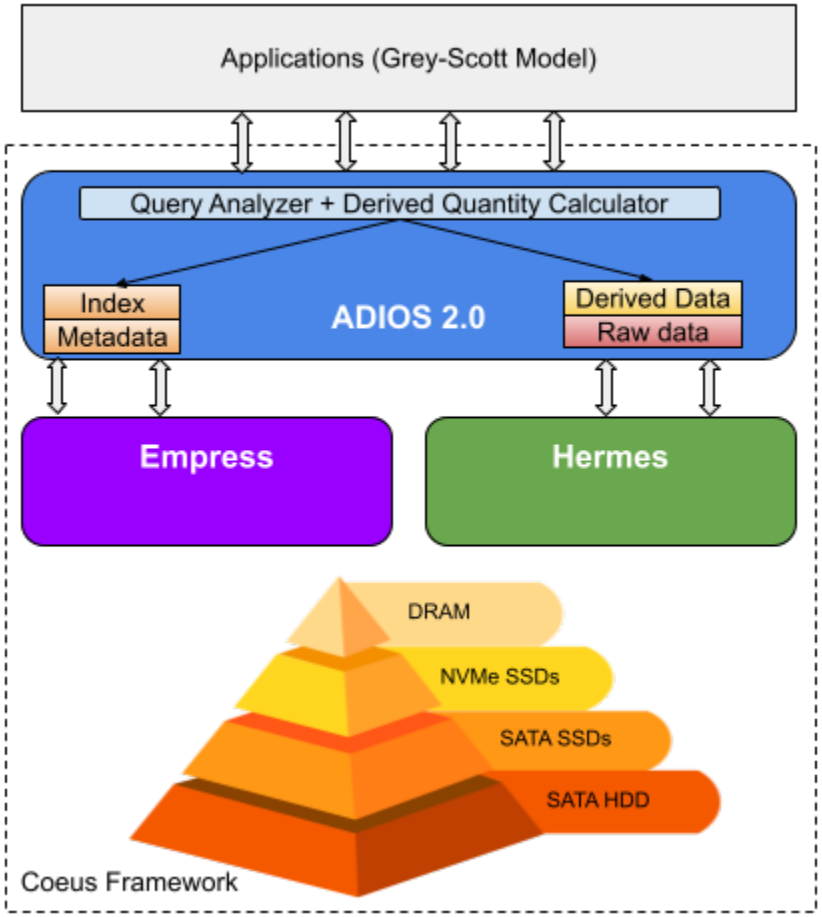

Questions

Sandia National Laboratories

- Jay Lofstead
- gflofst@sandia.gov