

Making Data Useful Through Metadata Techniques

Jay Lofstead



*Exceptional
service
in the
national
interest*

CHEOPS

08 May 2023



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.



Problem Overview

- Bits in a file are meaningless....unless you know:
 1. The file format/organization
 2. The application type/specific application that generated it

- For science, it can still be meaningless....unless you know:
 3. What configuration that application used
 4. What OS version, compiler, and TPLs were used
 5. Environment variables that affected the run
 6. And many others

Problem Overview

- Partial solutions:
 - Standard file formats and libraries
 - Provenance systems
 - Source repo
 - DevOps scripts and logs
 - Container construction manifests

- Biggest remaining problem:
 - Is there anything interesting in the data to look at?

Solution Space

Specific problem definition:

As an analyst evaluating simulation and observation data, what data properties exist within what file (object) within a dataset that may suit my explorations?

3 Existing Tool Generations

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)



- 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)



- 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)



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)



- 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)

All offer an attribute capability, but it varies a bit

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

Third Generation Tools (features)

- Associate tags data/parts of data
 - Region/var/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)



- 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)

- 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

- 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

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

Applying Metadata in the Large

- National Science Data Fabric (NSDF) pilot project
 - Valerio Pascucci (Utah), Michela Taufer (UTK) PIs
- Materials Commons 2.0
- FAIR data principles focused (another separate discussion)



Catalyze the creation of a globally connected, democratized infrastructure in which scientific investigation and training are unhindered by the limitations of dealing with extreme amounts of data.

Democratize access to large-scale scientific data by developing production-grade, scalable solutions to data storage, movement, and processing that can be deployed on commodity hardware, cloud computing, and HPC resources.

(courtesy of Valerio Pascucci)

National Science Data Democratization Consortium



THE UNIVERSITY OF TENNESSEE
KNOXVILLE




UNIVERSITY OF MICHIGAN



JOHNS HOPKINS UNIVERSITY

SDSC

SAN DIEGO SUPERCOMPUTER CENTER



Protocol Labs Research



DoubleCloud

Seal Storage Technology
Web3 Cloud Storage

WHO WE ARE
We're cloud storage and blockchain experts with over 100 years of experience in enterprise data storage from Seagate, Oracle, Cisco, and more. By seamlessly stewarding our clients into decentralized cloud storage, we're making Web3 an accessible reality for universities, research institutes, enterprises, and Web3 firms alike.

WHY SEAL
Seal provides sustainable, immutable, and affordable data storage.

- DATA RETRIEVAL**
Access data in hours vs days
- COST EFFECTIVE**
Up to 80% less than competitors
- SECURE**
Tamper proof and verifiable

(courtesy of Valerio Pascucci)

NSDF + Metadata Management

- Generation 5?
 - Metadata management for multiple applications in distributed store

- Issues addressed:
 1. What part of the rich, but large metadata can we save to help for data selection?
 2. How do we distribute the metadata system as a local cache for large, distributed data storage system?
 3. How do we deal with different kinds of storage systems, file formats, data organizations, and other complexities?

Thank you



- gflofst@sandia.gov