Optimizing Storage Utilization in the HL-LHC Era

Qiulan Huang, Vincent Garonne
Scientific Data and Computing Center (SDCC)
Brookhaven National Laboratory

US ATLAS Computing Facilities Meeting
11/30/2022

@BrookhavenLab

LDRD project

- Project title:Data popularity, placement optimization and storage usage effectiveness at Data Center
- Investigators: Qiulan Huang(PI, SDCC), Vincent Garonne(SDCC),Ai Kagawa (CSI, Computational Science Initiative), Xin Dai (CSI)
  - AI/ML team(CSI): Develop the state-of-the-art data usage prediction model
  - Storage team(SDCC): Quality Assurance & Control of training data & performance monitoring
- Term: 2022- 2024

Current Storage at SDCC

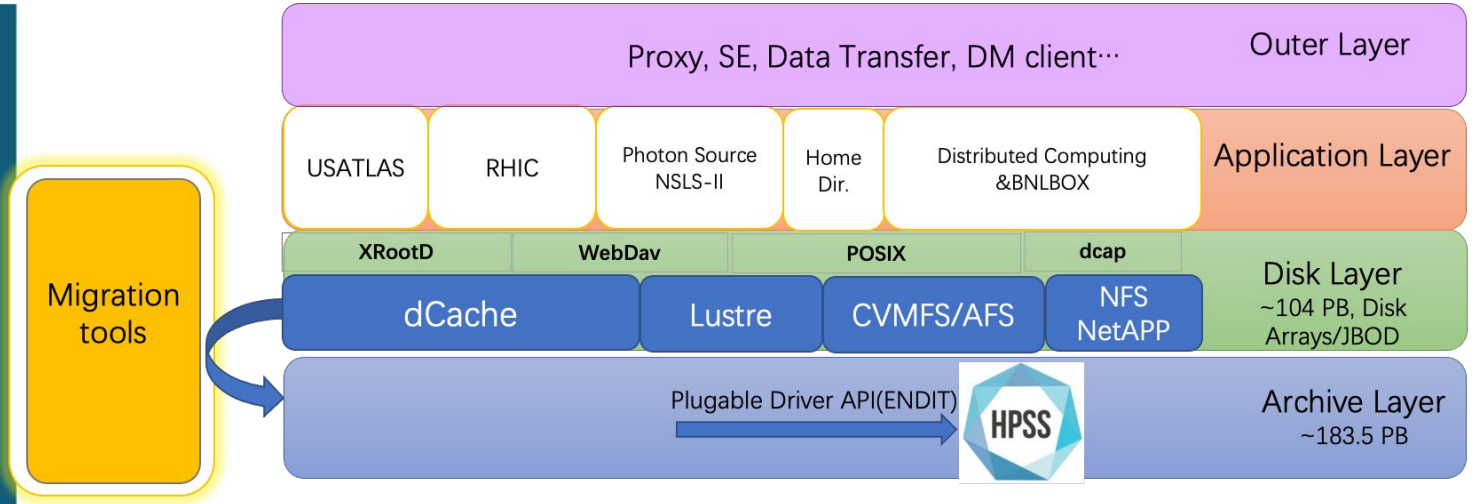

Storage "classes" (DISK,TAPE, FLASH,etc) have various cost and quality of services, E.g.,

- **slow and cheaper storage for archiving vs. fast and expensive storage for quick and efficient data access**

Two storage classes : Disk and Tape

- **Disk storage: ~104 PB (fast storage)**
- **Tape storage: ~183.5PB(slow storage)**

- **Several millions of files created/deleted/transferred per day**
- **Peak traffic at 100GB/s**
- **Data accessed by millions of jobs per day**
- **Hundred of thousands of storage devices**

Brookhaven National Laboratory

Data temperature of ATLAS

~50% files are frozen, **~22%** of total size(~8PB) is frozen
Cannot be solved by adding more storage but by better use of storage resources

★ Plot done more than one year ago on all dCache replica categories (cache, precious, etc) without curation

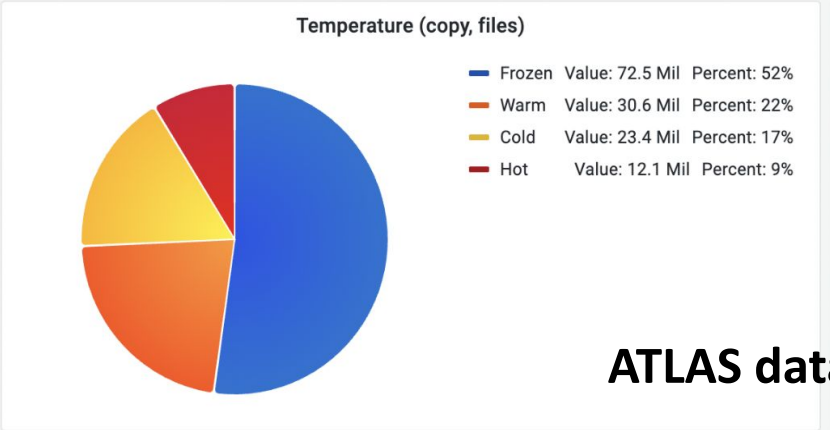

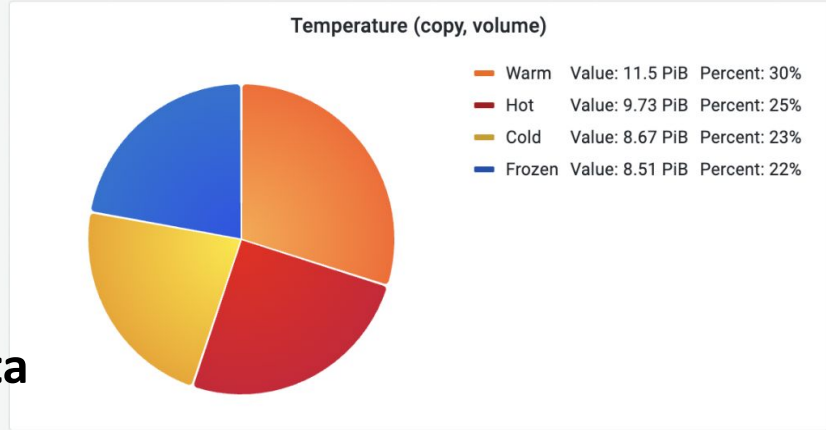

ATLAS data

Hot: Last access in the last month
Cold: Last access between 6 months and one year

Warm: Last access in the last 6 months
Frozen: Not accessed in the last year

⇒How to map data category with activities and storages classes?

⇒How to provide a better accurate view for more experiments and technology: dcache, lustre (robinhood), etc

AI/ML For Storage Optimization

Motivation

- In the current multi-tier storage "class" system at the Data Center
  - Unused data is stored on expensive storage
  - Fast IO storage is not currently used effectively

Goals

- Design an efficient monitoring platform to collect the relevant information from various distributed data sources
- Develop an optimal data management system for the data center to maximize usable space while minimizing access latency, within budget, hardware, and compliance constraints
  - Heavy use of storage, metadata and data popularity information
  - Detect early failures and pathological usage pattern
  - Develop a precise AI/ML prediction model to possibly forecast the future usage of the data
  - Orchestration of data for optimal movement and placement

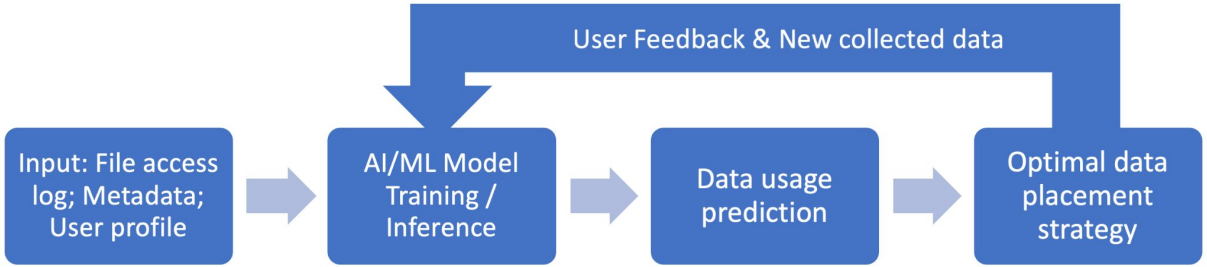

Brookhaven National Laboratory

What benefits of the LDRD

- Help in making important business decisions, policies and lead to more cost effective operations
  - Implement the automatic migration of data between different levels of storage
  - Measurable objectives, e.g., min(access time) and max(total usable space) on a same budget
- Define optimal data placement decisions
  - Data distribution in a reasonable tiered storage and provide good IO performance for scientists

Novel Storage at SDCC

The LDRD contributions are shown in the red dashed boxes

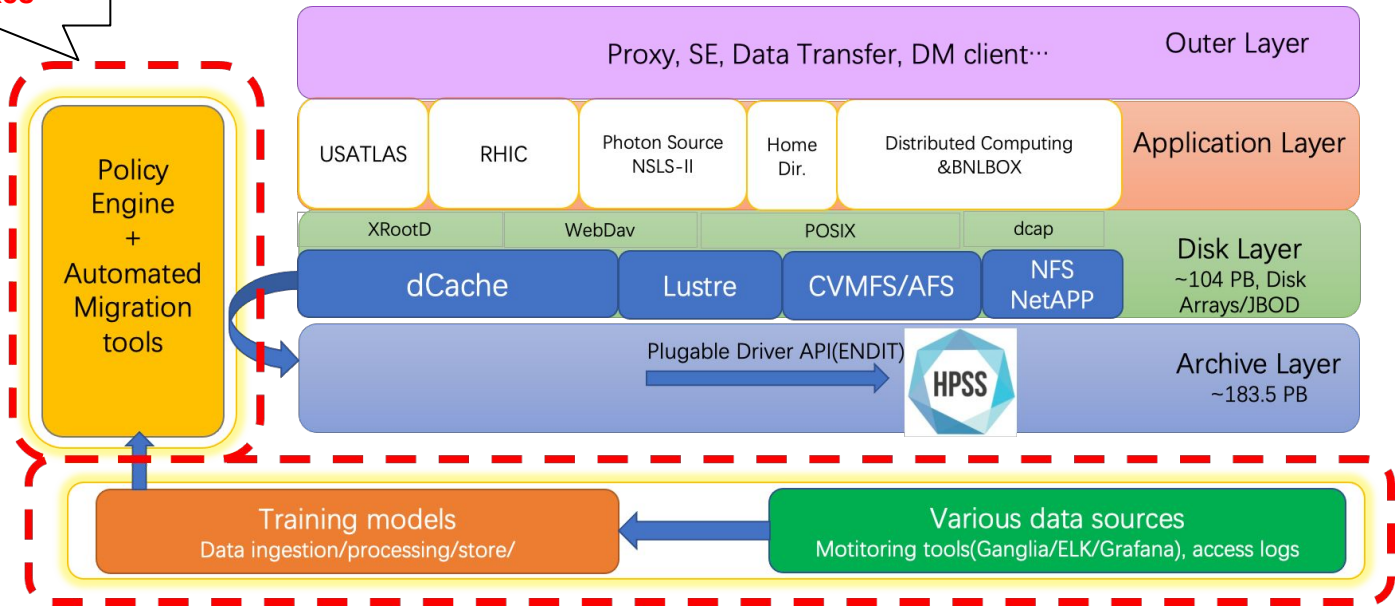

- **Introduce AI/ML combined with our big data, conduct extensive training, and build data prediction model**
- **Define the policy engine and perform automated migration actions transparently**

Status

- Collect and analyze data samples is undergoing
  - dCache billing, chimeradb, access logs,etc
  - [File id],[File size],[access_time],[disk usage],etc
- Machine learning analyzes all of the collected data samples
  - Start from a simple definition to predict near future(like 30 days) data usage
- Biweekly meetings are scheduled