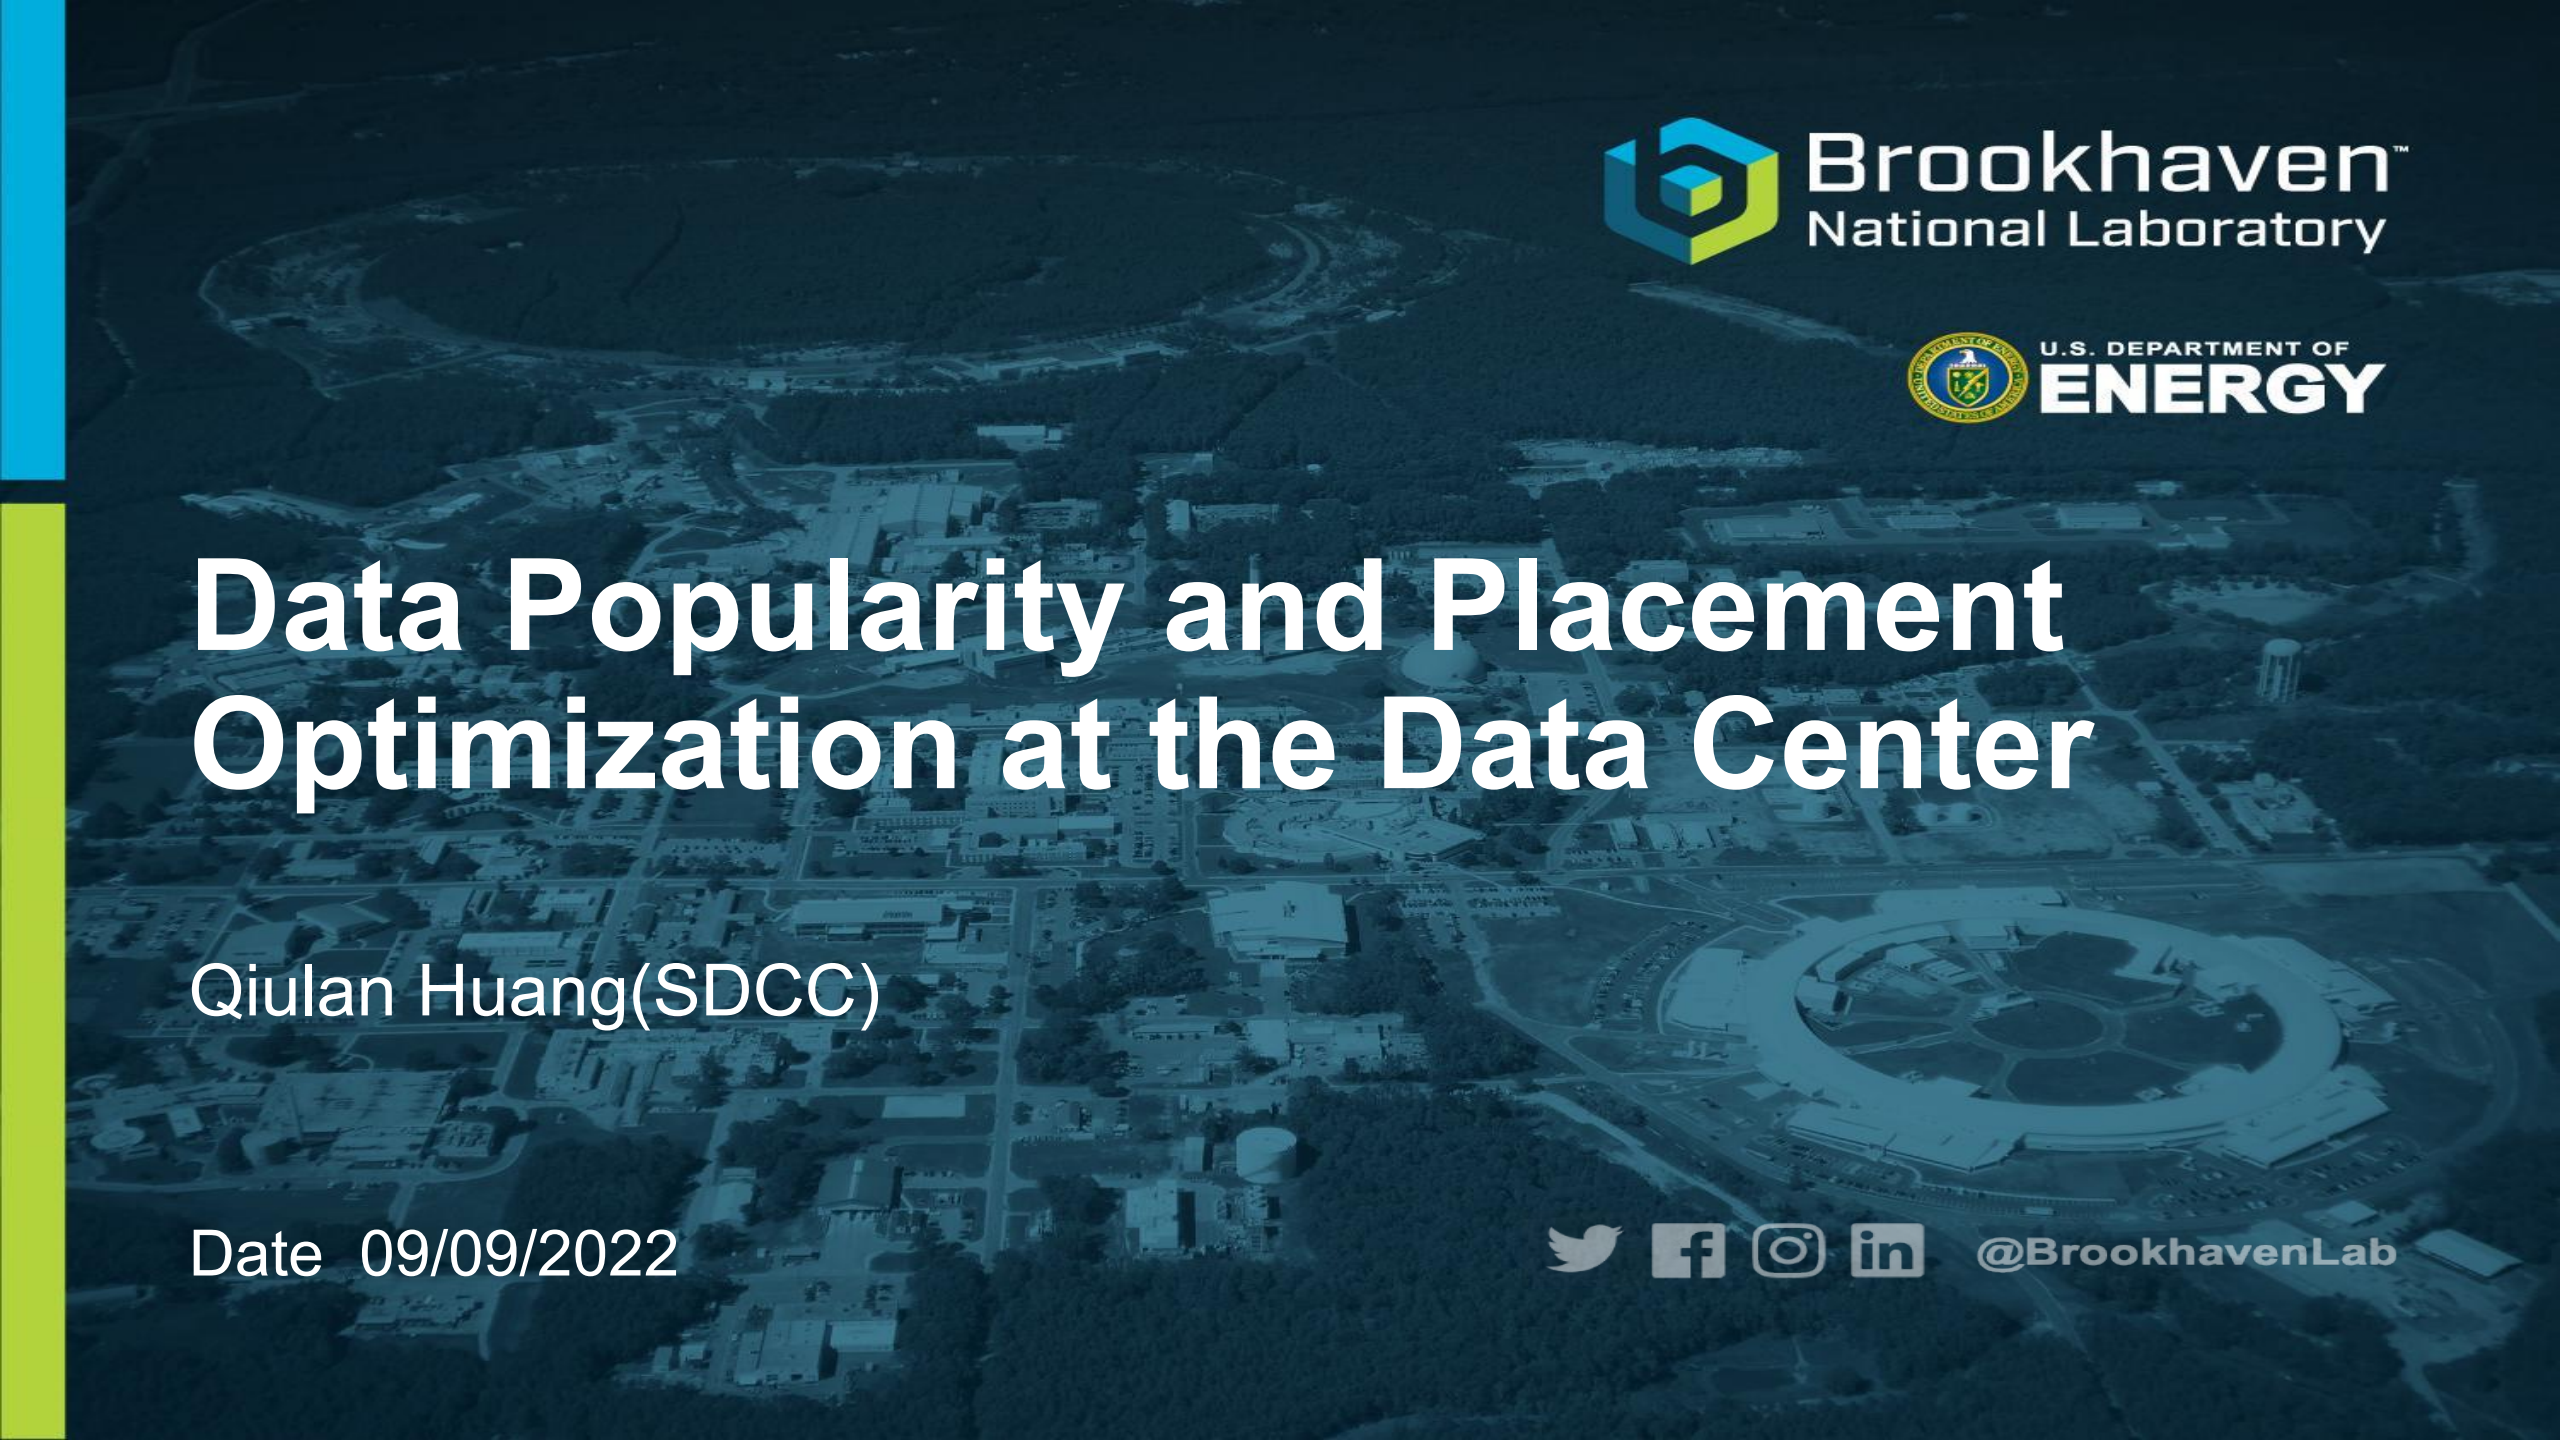

Data Popularity and Placement Optimization at the Data Center

Qiulan Huang(SDCC)

Date 09/09/2022

@BrookhavenLab

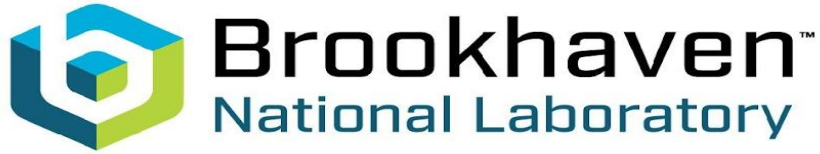

Data Popularity and Placement Optimization at the Data Center

Qiulan Huang(SDCC)

Date 09/09/2022

Major challenges ahead of us will bring the Scientific Data and Computing Center (SDCC) at the Exabyte-scale

- Current storage status
  - Disk storage: ~104 PB (fast storage)
  - Tape storage: ~183.5PB(slow storage)
- Scientific data keeps increasing in an order of magnitude over the next 5 years
- Increasing complex as new bleeding edge technologies are being deployed in hardware
  - E.g., Flash storage, SSD, etc.

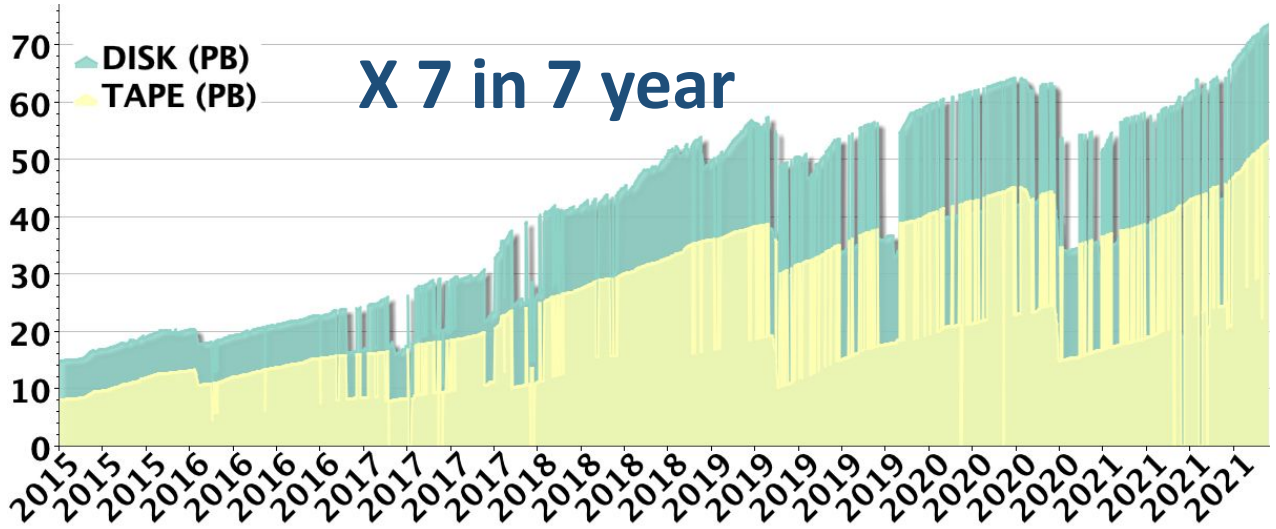

LHC experiment(ATLAS) data at BNL / SDCC

Brookhaven National Laboratory

The SDCC optimization problem

- Different storage “classes” have various cost and quality of services. E.g., slow storage and cheaper for archiving vs. fast and expensive storage for quick and efficient data access
- Increasing requirements for mass storage system like capacity, reliability, scalability, and cost efficient
- Problems with data classification, placement, and migration
  - Data often has no clear boundaries. e.g., which data should be in fast storage? Which data should be in slow storage?
  - Unreasonable data placement. e.g., unused data stored on the expensive storage
  - Migrate data manually with simple rules defined by Data Management Admin
    - Based on human knowledge,decisions and simple usage statistics

Well-known Problem in storage

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

⇒How to manage data in a cost efficient manner while enabling effective data access and processing at scale

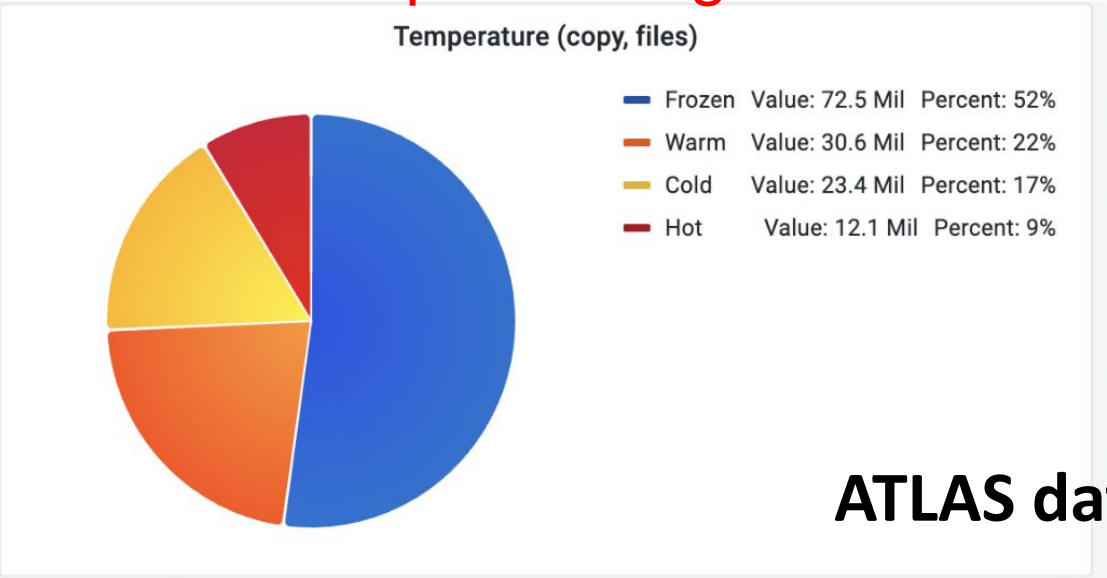

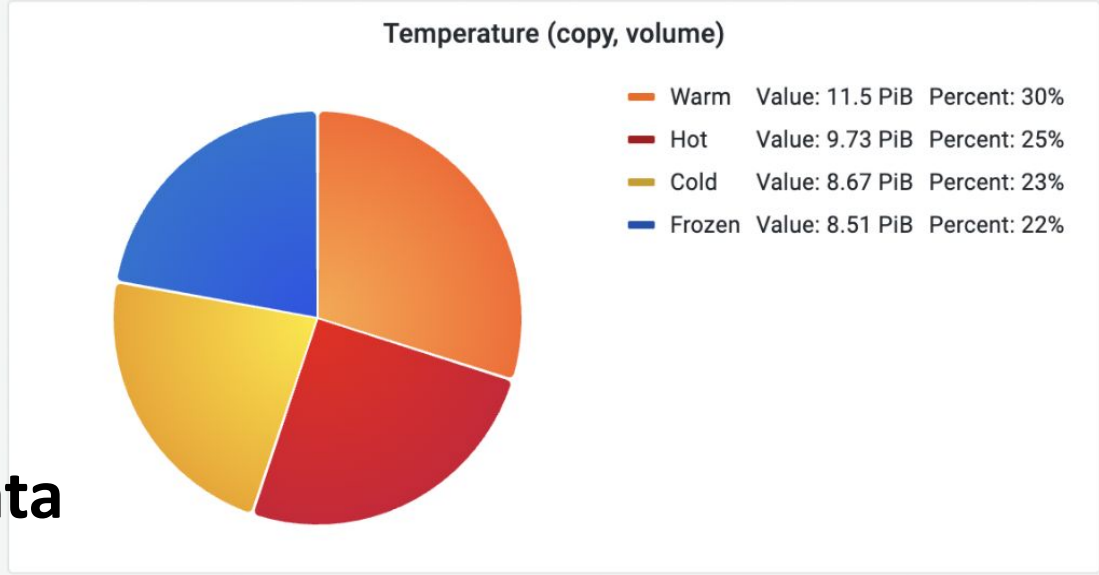

ATLAS data

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

Brookhaven National Laboratory

AI/ML For Storage Optimization

Limitations of current AI/ML storage optimization solutions:

- Non-scalable (designed for small storage)
- No support for multi-tier storage architecture (two-tier fast-slow partition)
- Inaccurate for long-term prediction (low hit-ratio)

Undergoing efforts:

- Develop the state-of-the-art data usage prediction model
  - Scalable, versatile and accurate long-term predictive model
- Quality Assurance & Control of training data & performance monitoring
  - Strong data mgt expertise to validate solutions in production with incremental delivery strategy

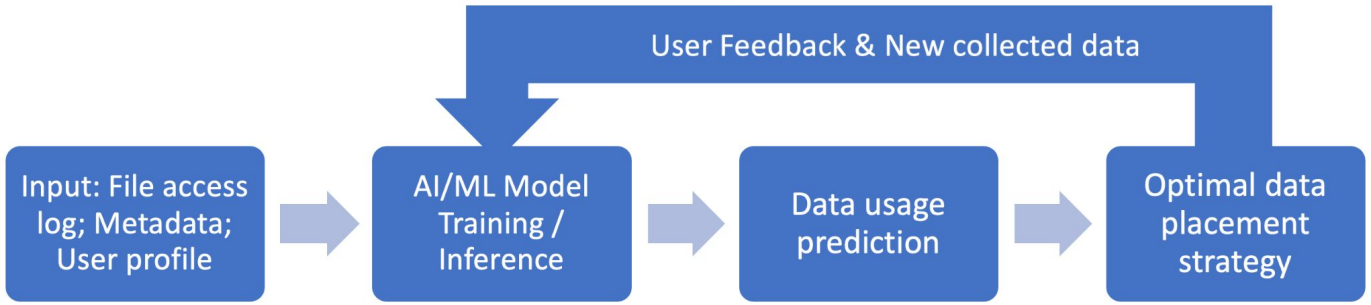

Brookhaven National Laboratory

Summary

- This project is supported by LDRD funding to develop big data analysis facilities and support data at extreme scale for various programs in NP & HEP for the current and long term.
- A precise AI/ML prediction model will be developed to forecast the future usage of the data for optimal data placement, *which will be the core of our development of novel, automated extreme-scale data pipeline paradigms*
- It will bring great benefits both to scientists and BNL, e.g. save $$ for the long term
    - Manage data in a cost efficient manner while enabling effective data access and processing at scale will be achieved

Any questions?
Thank you

Backup

Current Storage at SDCC

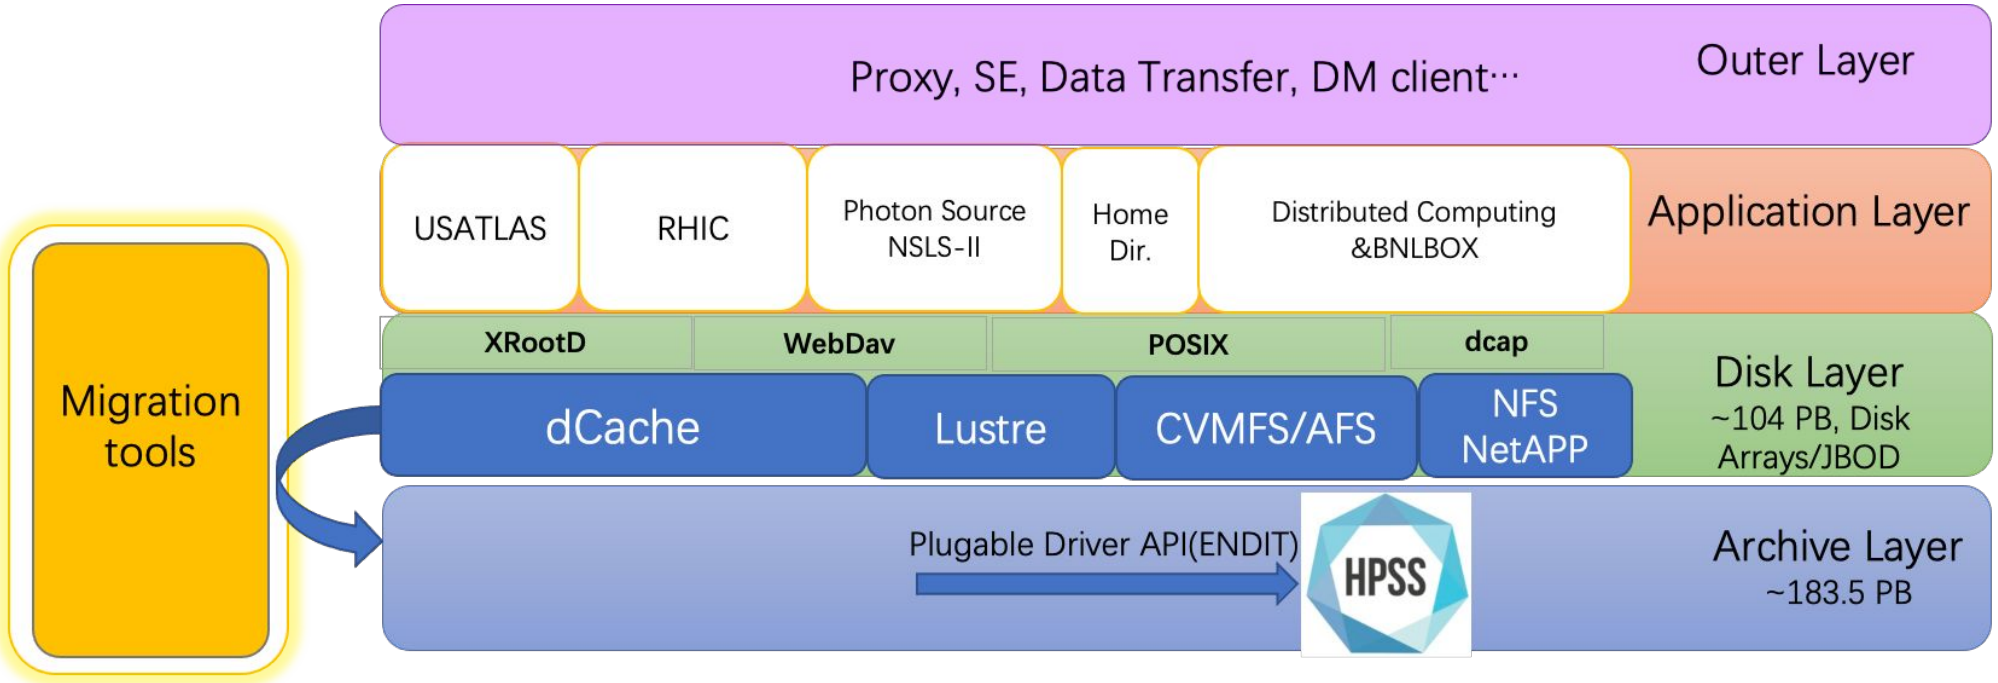

- **Disk storage: ~104 PB (fast storage)**
- **Tape storage: ~183.5PB(slow storage)**

- **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

New Storage at SDCC

The contributions of this project are shown in the red dashed boxes

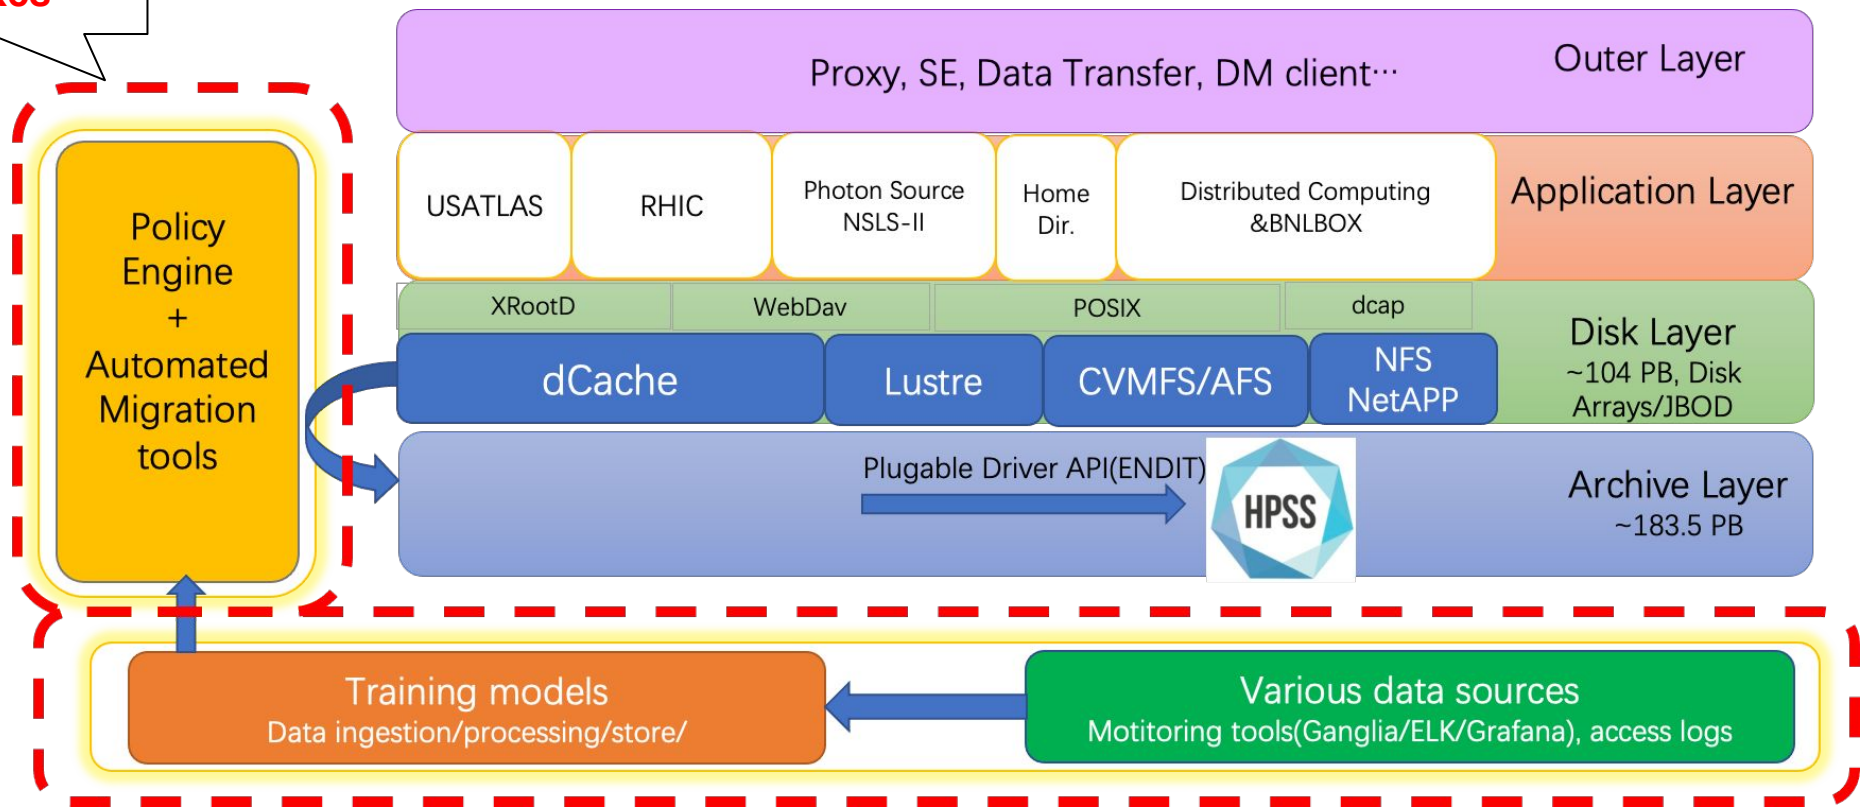

- **Introduce AI/ML combined with our big data, conduct extensive training, and build data prediction model**
- **Define the policy engine and perform automated migration actions transparently**