Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*

DIRAC @ PNNL

MALACHI SCHRAM
ON BEHALF OF PNNL PROJECT8, IPPD2, AND CCSDI
Pacific Northwest National Laboratory

DIRAC experiments and development

- Current work:
  - Belle II (High Energy Physics)
  - Integrated End-to-End Performance Prediction and Diagnosis for Extreme Scientific Workflows (Advanced Computing)
  - Project8 (Nuclear Physics)
  - CCSDi (Complex Chemical Systems)

- Demo and evaluations:
  - SuperCDMS (High Energy Physics)
  - miniCLEAN (High Energy Physics)
  - nEXO (Nuclear Physics)

Belle II @ PNNL

- PNNL was the lead for the Distributed Data Management System operations and development

- PNNL managed several Belle II DIRAC servers:
  - DDM production slave servers (now at BNL)
  - DDM development slave servers (still online)
  - BelleDIRAC development master server
  - BelleDIRAC migrate master server
  - BelleDIRAC certification master server

- HPC resources access and containers is no longer supported

- PNNL is officially no longer working on the Belle II Distributed Computing at the end of June

IPPD: Integrated Performance Prediction & Diagnosis

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

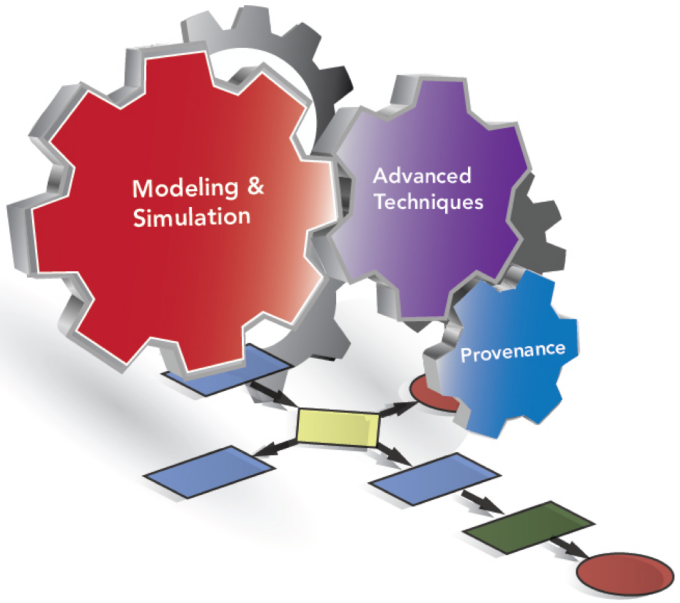

- **Aim to provide an integrated approach to the modeling of extreme scale scientific workflows**
- **Brings together researchers working on modeling / simulation / empirical analysis, workflows and domain scientists**
- **Builds upon existing research much of which has focused to date on large-scale HPC systems and applications**

BROOKHAVEN NATIONAL LABORATORY

UNIVERSITY OF CALIFORNIA SAN DIEGO

Explore in advance – Design-space exploration & Sensitivity Analyses

Optimize at run-time – Guide execution based on dynamic behavior

Key research areas in IPPD

- **Provenance**
  - Capture empirical performance information from workflows enabling baseline performance to be established; to identify and help diagnose variability; to feed simulation and modeling;
- **Modeling and Simulation**
  - Develop modeling and simulation to enable both exploration in advance of possible workflow configuration and optimizations, as well as rapid performance prediction to guide the dynamic adaptation of workflows and optimization of resource utilization.
- **Advanced Techniques**
  - Explore novel optimization techniques for workflow optimization for both processing and data organization. Leverage unit-commitment ideas (from the smart power-grid), and transparent page management and compression (for data-movement / storage reduction)
- **Workflows – *Belle II*, ACME, Project8**
  - Initially explore our developments for analyzing, optimizing and considering future design of the Belle II High Energy Physics project
  - Also extending to others including metagenomics, climate.
  - Explore new data streaming techniques within DIRAC for the Project8 experiment

IPPD's Capability Demonstrator: 'Enhanced' Belle II Workflow Execution

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Challenge
- Mitigate contention
- Consider power/energy

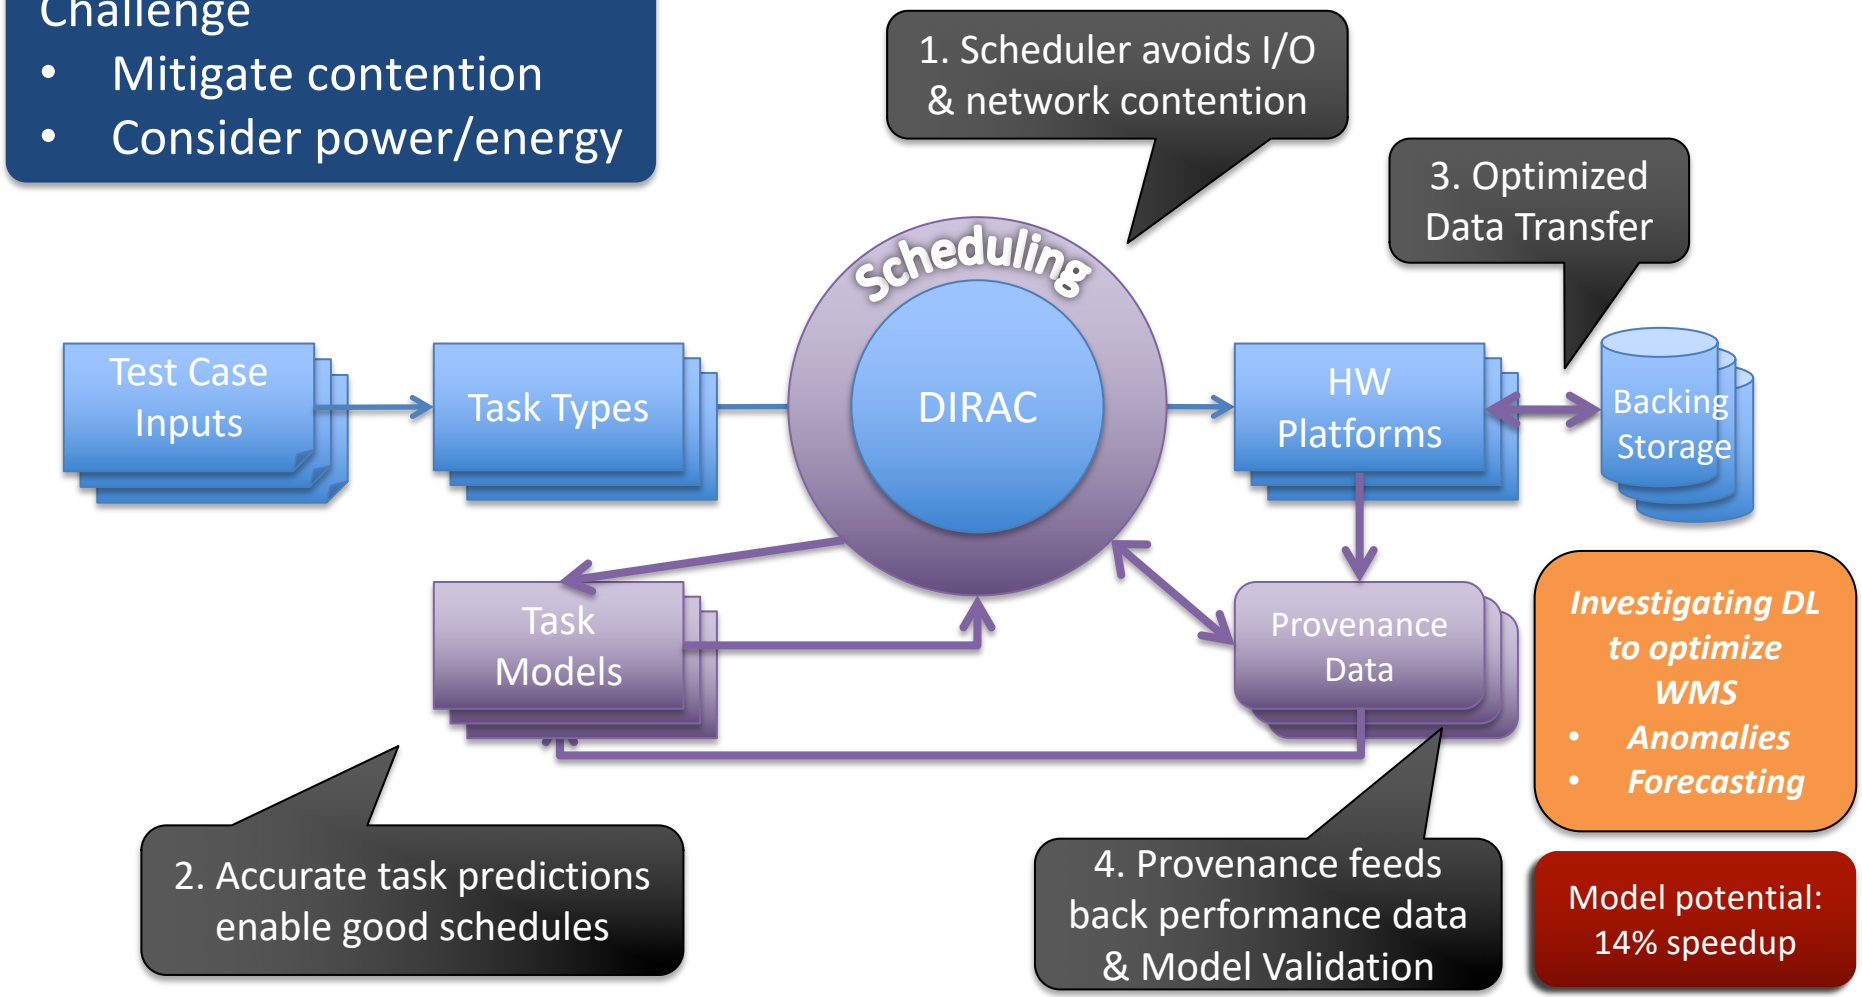

IPPD/2 Activities

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Challenge

- Data movement bottlenecks
- Heterogenous hardware

Observe

Monitor performance costs, task constraints, data movement

Analyze

Task models predict performance, to enable scheduling

Analyze

Detect anomalies during runtime, to proactively guide execution

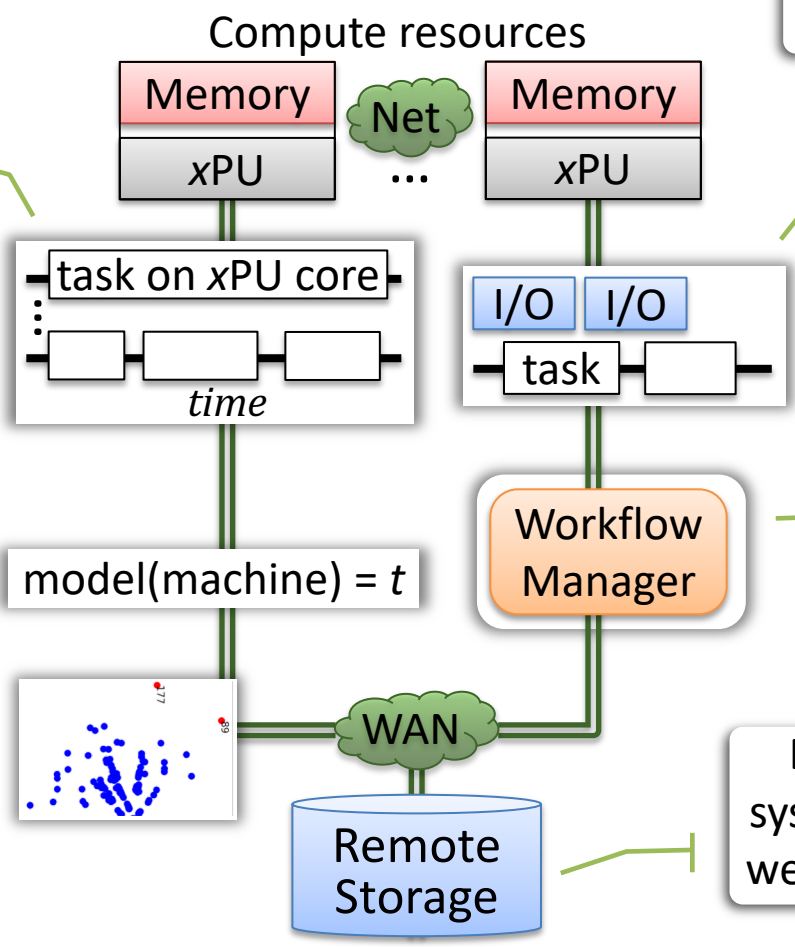

Accelerate

Transparent data prefetching overlaps data movement with computation

Speedup: 15-25%

Accelerate

Scheduling evens I/O load and reduces response time across different hardware

Speedup: +10-15%

Accelerate

Dynamically auto-tune file system parameters to perform well under different workloads

Speedup: 15-35%

The Tritium Endpoint Method w/ Project8

Cyclotron Radiation Emission Spectroscopy
(measure frequency of radiation from magnetically trapped electrons)

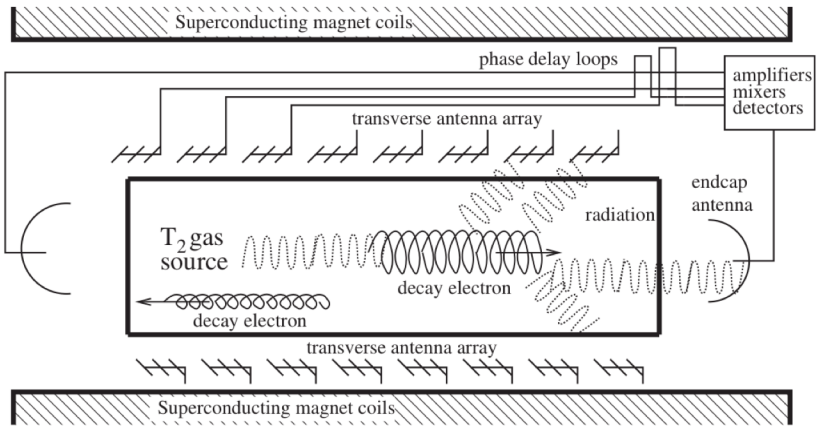

Invert *f(E)*

Tritium β-decay electron endpoint

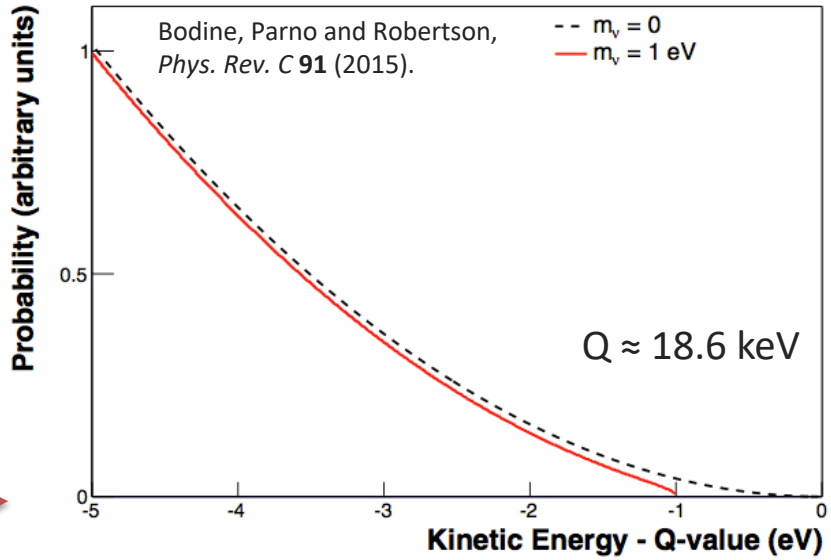

- Tritium Beta Decay: $^3H \rightarrow {}^3He^+ + e^- + \nu_e + Q$.
- High-precision spectroscopy on the e^-.
- Neutrino mass manifests as a deviation at the energy endpoint.
- Fit the spectral shape with m_{ve}^2 as a free parameter:

Project8 Data Taking and Computing

- Expected data rates to increase at each experimental phase:
  - Phase 1 & 2 (now): ~0.5PB
  - Phase 3 (FY2021): 10-20PB
  - Phase 4 (FY2025-2030): ExaBytes
- Processed data samples will be distributed within the Continental U.S.
- Project8 leverages metadata information in conjunction with the transformation system to automate the production efforts.
- PNNL leads the design, deployment, and operations of the computing effort
- Leveraging virtualization expertize to deploy a robust and scalable solution using containers, Kubernetes, and HELM. Used for core Project8 computing system and PNNL and Yale.

Project 8 and DIRAC

- Project 8 currently uses DIRAC for all aspects of data and work management:
  - RAW data transfers from experiment DAQ at University of Washington (UW) to storage element at PNNL and Yale
  - Automated analysis of calibration and slow control data, with results returned to operations web server at UW
  - Analysis of fast DAQ time series data
  - User access to data and CPUs at PNNL

- Work in progress:
  - Updating the metadata for the different data products to automatically job processing
  - How to propagate metadata and ancestry in a transformation?

Current Project8 DIRAC Computing Setup

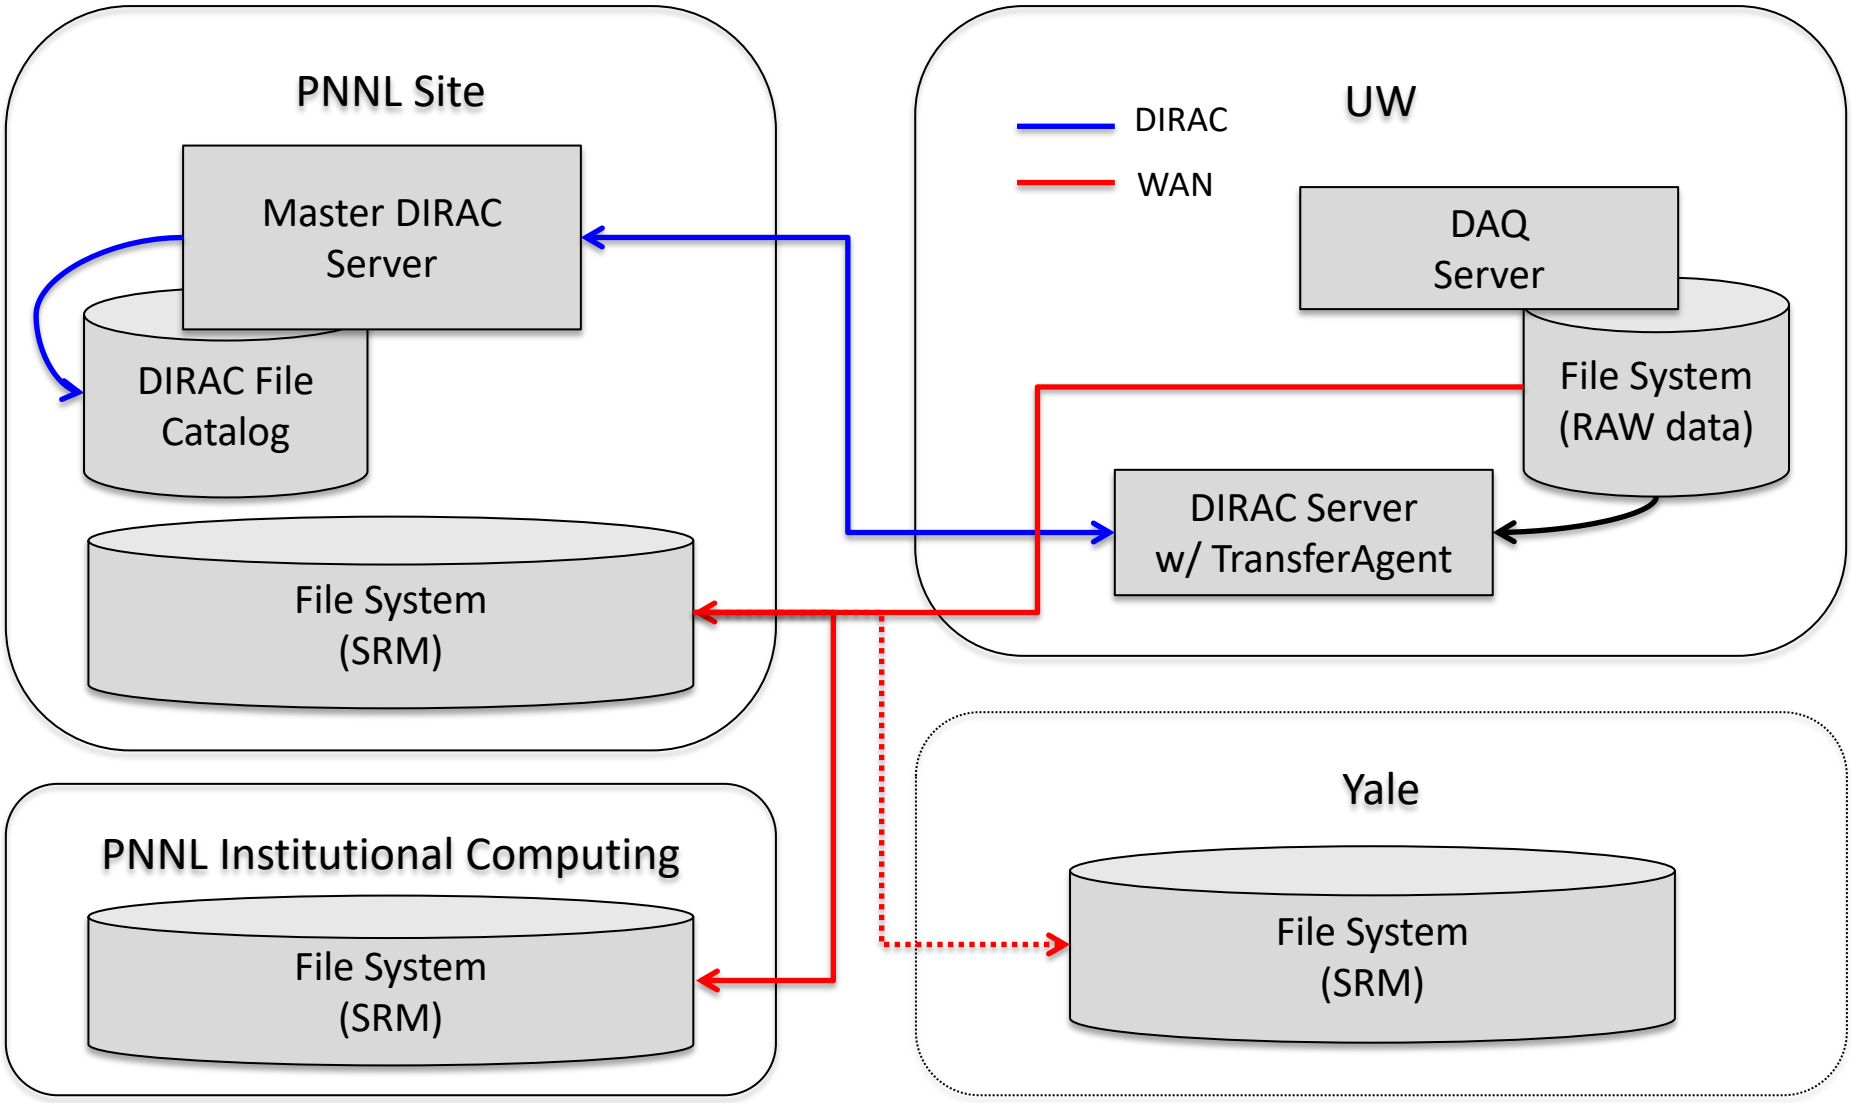

DIRAC servers for Project8

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- Master DIRAC Server @ PNNL:
  - Web Portal: dirac-project8-prod.hep.pnnl.gov:8443/DIRAC
  - DIRAC File Catalog (DFC): File location and metadata
- DIRAC slave server at UW:
  - Dedicated agent to automate data transfers from UW to PNNL
- DIRAC slave server at PNNL Institutional Computing
  - Provides access to 100TB of disk and 100TB of tape storage
- DIRAC slave servers will provide access to PNNL:
  - Opportunistic HEP Cloud (>4000 cores)

Migrating to new Hardware

- Developed tools to deploy DIRAC system using HELM, Kubernetes, and containers
- Automate workflow using the TS by carefully assigning metadata

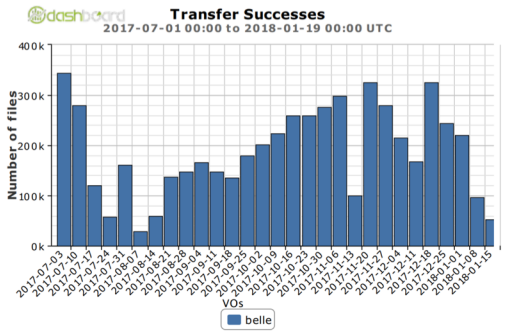

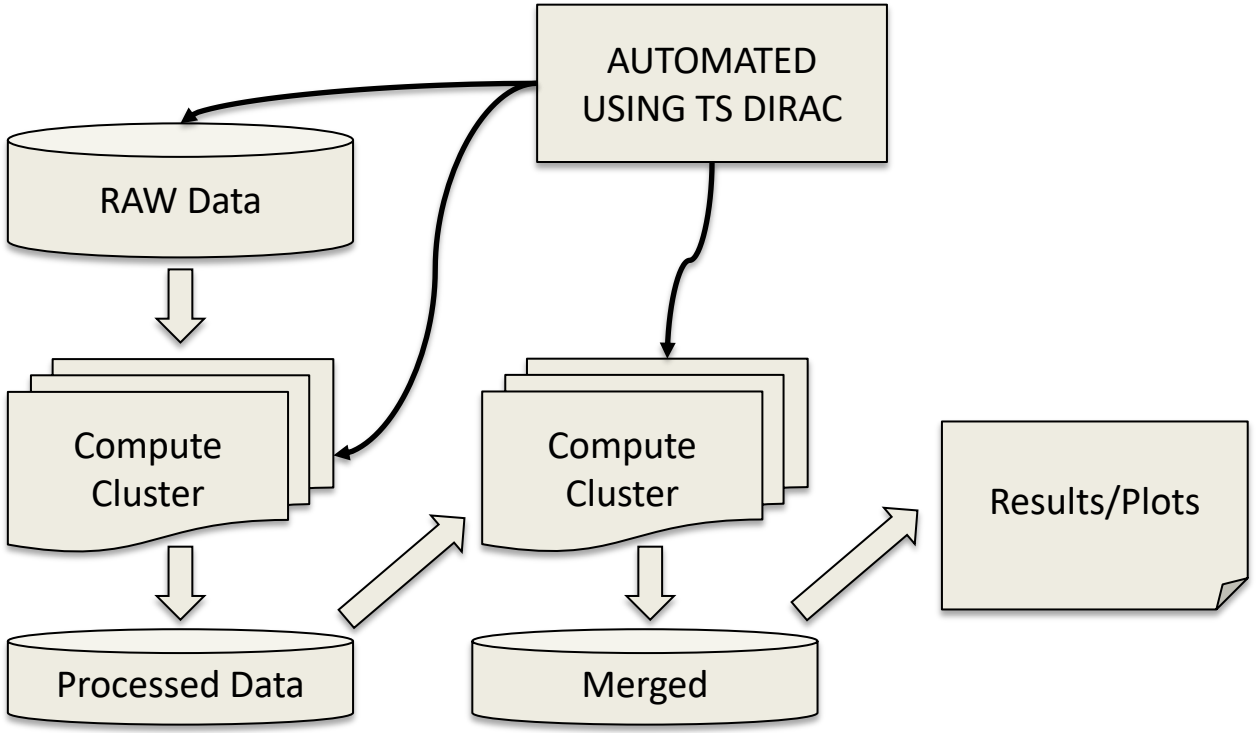

Complex Chemical Signatures

- Asynchronous workflow for experimental and model results
- Automate workflow using the TS by carefully assigning metadata
- Need to create a TS plugin to trigger new comparisons (multiple input files)

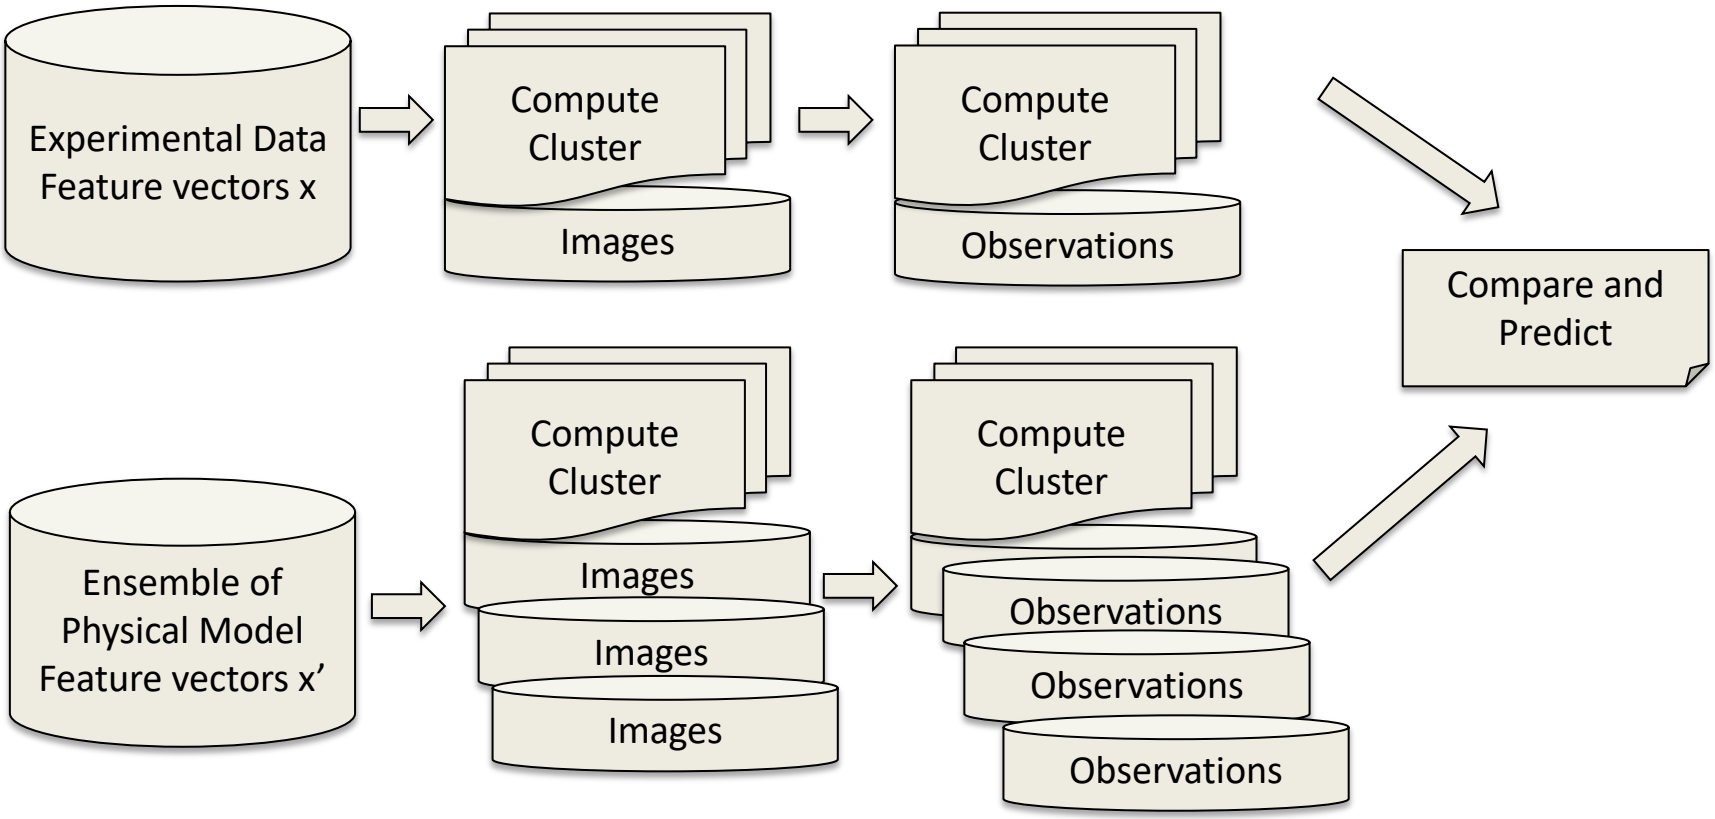

Automating analysis chain using TS

Examples of questions we are asking:

- How do you determine what measurement to perform next?
- Are we missing any key data feature vector variable in our models?
- How can we reduced computational time for expensive physical models (ensemble builder, etc.)
- How can we reduce the number of experiments

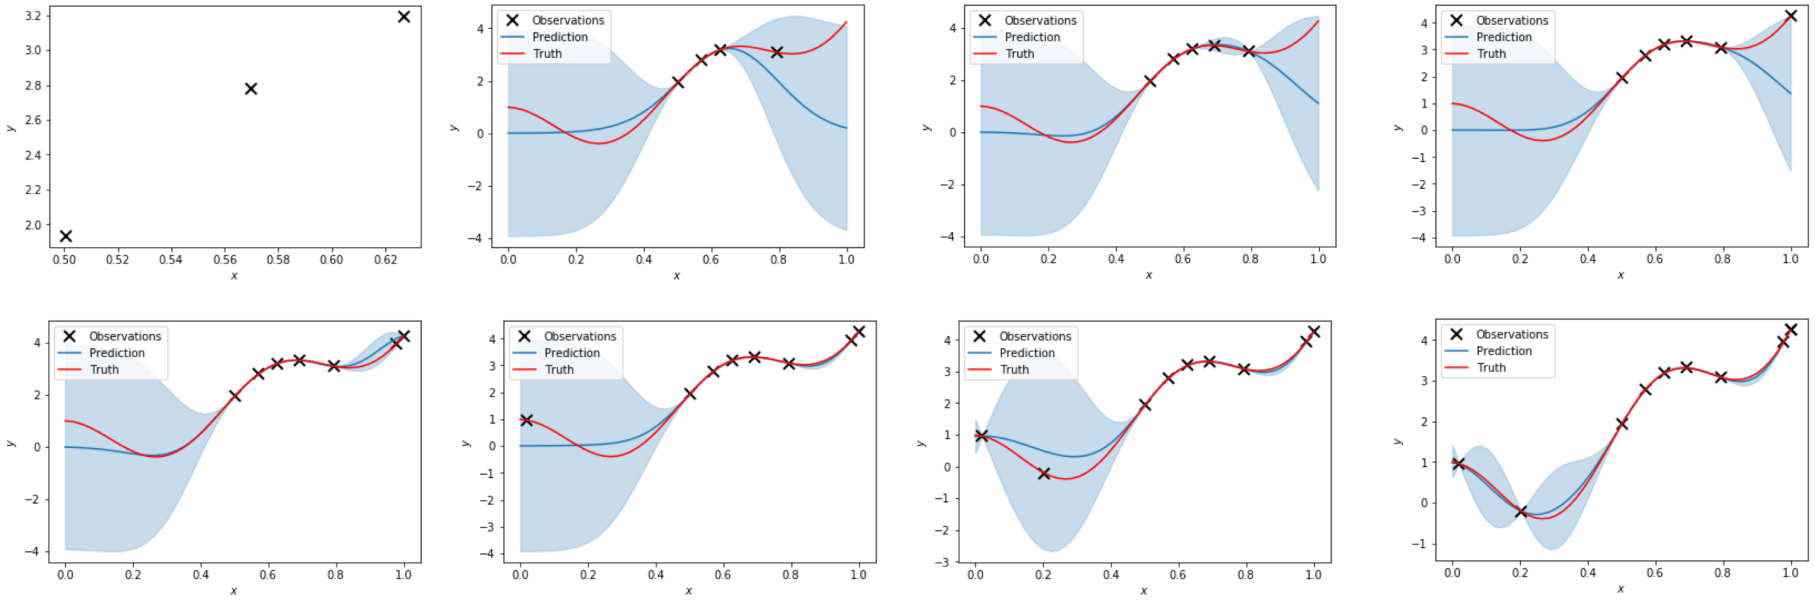

Summary of DIRAC related work at PNNL

- We are using and developing DIRAC components for 4 active projects
- We are investigating ML techniques to improve overall efficiency
- We heavily use the Transformation System to automate workflows
- General interested in using DIRAC for other projects
  - Evaluations are ongoing or planned