Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

DIRAC @ PNNL

MALACHI SCHRAM AND VIKAS BANSAL
ON BEHALF OF PNNL BELLE II, PROJECT8, MINICLEAN TEAMS
Pacific Northwest National Laboratory

DIRAC experiments and development

- Current work:
  - Belle II (DOE-HEP)
  - Integrated End-to-End Performance Prediction and Diagnosis for Extreme Scientific Workflows (DOE-ASCR)
  - Project8 (DOE-NP)
  - MiniCLEAN (DOE-HEP)

- Demo and evaluations:
  - SuperCDMS (DOE-HEP)
  - nEXO (DOE-NP)

Belle II @ PNNL

- Core Distributed Data Management System development team
  - (see Miyake-san's talk)
- Belle II DIRAC servers @ PNNL:
  - DDM production slave servers
  - DDM development slave servers
  - BelleDIRAC development master server
  - BelleDIRAC migrate master server
  - BelleDIRAC certification master server
- Current Core DIRAC focus:
  - Evaluating the new FTS3 Agent
  - Investigating/evaluating the LFC client

Belle II Computing on HPC

- NERSC has two clusters with grid access:
  - ***Edison***
    - 133,824 compute cores
    - ~2,000 compute cores for shared queue
  - ***Cori***
    - Phase 1 (Haswell): 2,004 compute nodes
    - Phase II (KNL): 9,688 compute nodes, 658,784 cores in total (68 cores per node).
- **NERSC Allocation for 2017**
  - Used to develop/test/validate the use of HPC for Belle II
- Parasitic DIRAC agent is now submitting job to Cori
  - Additional development is required

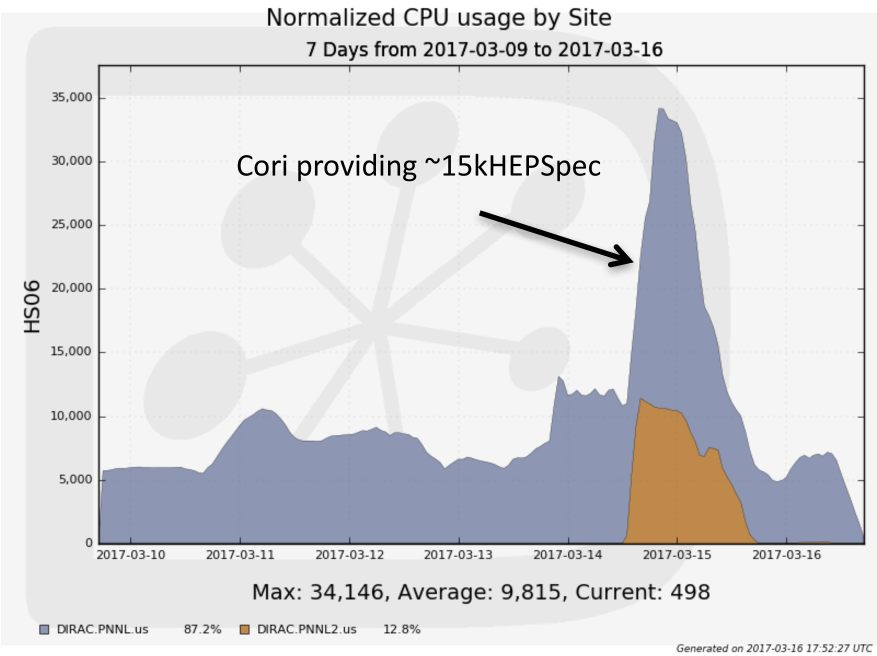

Project Title	US Belle II HPC Workflow	**Funded by DOE Office of Science?**	Y
Organization	Pacific Northwest National Laboratory (PNNL)	**DOE Manager**	Kevin Flood, Glen Crawford
DOE Office & Program	HEP - HEP Other Research	**MPP Hours requested in ERCAP**	10,000,000
Science Category	High Energy Physics	**SRUs requested in ERCAP**	1,000
Project class	DOE Base Funding		

IPPD: Integrated Performance Prediction & Diagnosis

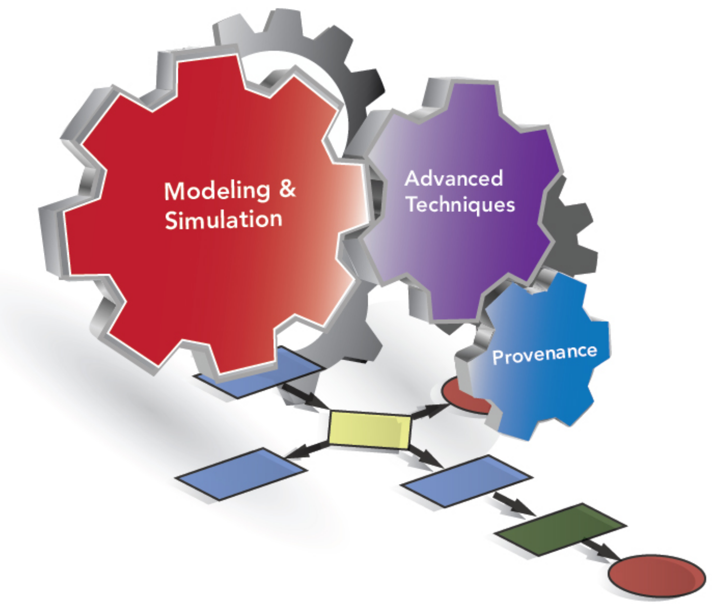

- Aim to provide an integrated approach to the modeling of extreme scale scientific workflows
- Brings together researchers working on modeling / simulation / empirical analysis, workflows and domain scientists
- Builds upon existing research much of which has focused to date on large-scale HPC systems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

Explore in advance **–** Design-space exploration & Sensitivity Analyses

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

Key research areas in IPPD

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- **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**
  - 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.

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

1. Scheduler avoids I/O & network contention

3. Optimized Data Transfer

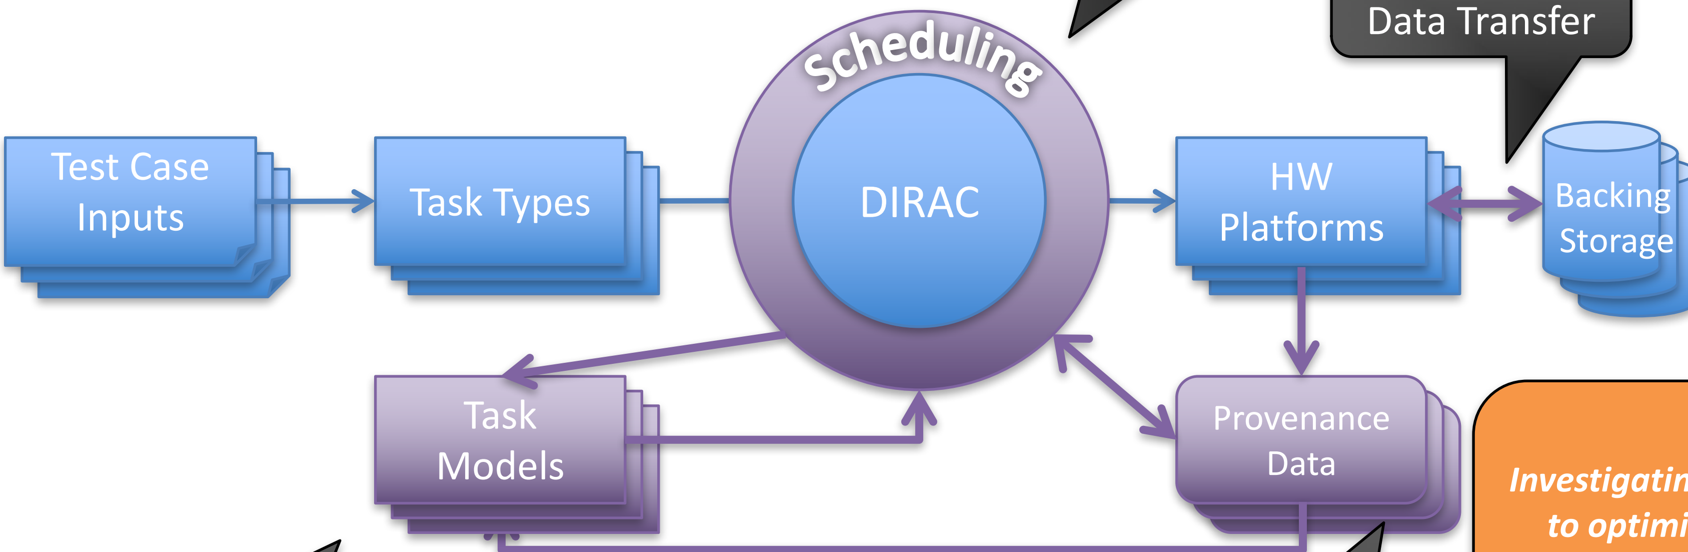

2. Accurate task predictions enable good schedules

4. Provenance feeds back performance data & Model Validation

Investigating DL to optimize DDMS
- ***Anomalies***
- ***Forecasting***

The Tritium Endpoint Method w/ Project 8

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

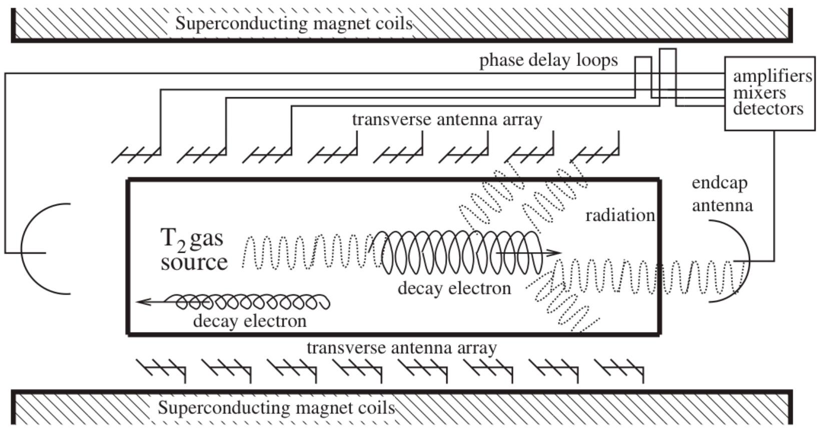

Invert *f(E)*

Tritium β-decay electron endpoint

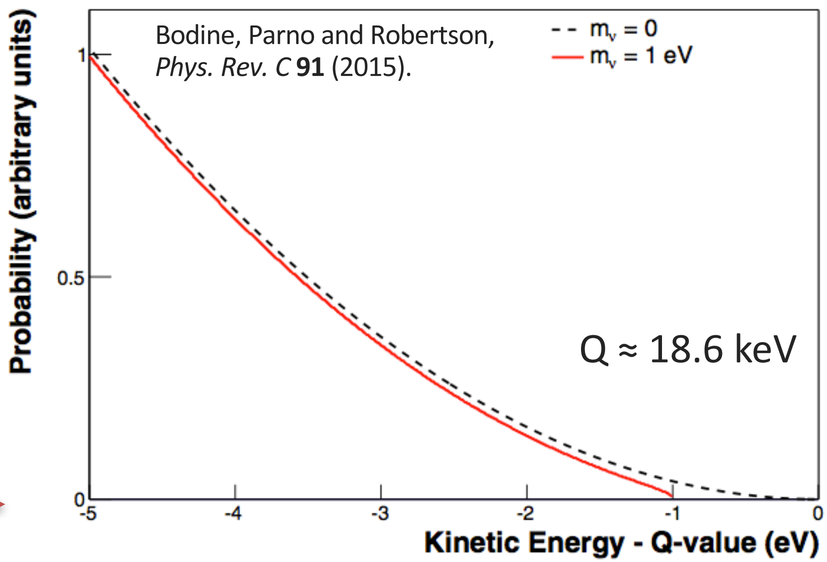

- Tritium Beta Decay: $^3H \rightarrow {}^3He^+ + e^- + \nu_e + Q$.
- High-precision spectroscopy on the e^-.
- Neutrino mass manifests as a deviation at the energy endpoint.
- Fit the spectral shape with $m_{\nu e}^2$ as a free parameter:

Project 8 and DIRAC

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- Project 8 currently uses DIRAC for all aspects of data and work management:
  - Data transfers from experiment DAQ at University of Washington (UW) to storage element at PNNL
  - 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

- Works in progress with DIRAC data management:
  - Integration of simulation metadata, etc.
  - Automation of fast DAQ data analysis
    - How to propagate metadata and ancestry in a transformation?

- Future wish list for DIRAC data management:
  - Expanded use of “dataset” feature that appears to be at least partially implemented in the CLI
  - More command-line tools to facilitate ancestry/descendent data selections, to isolate “ordinary” users from the DIRAC API
  - Similar command-line tools to facilitate dataset operations

Current Project8 DIRAC Data Workflow:

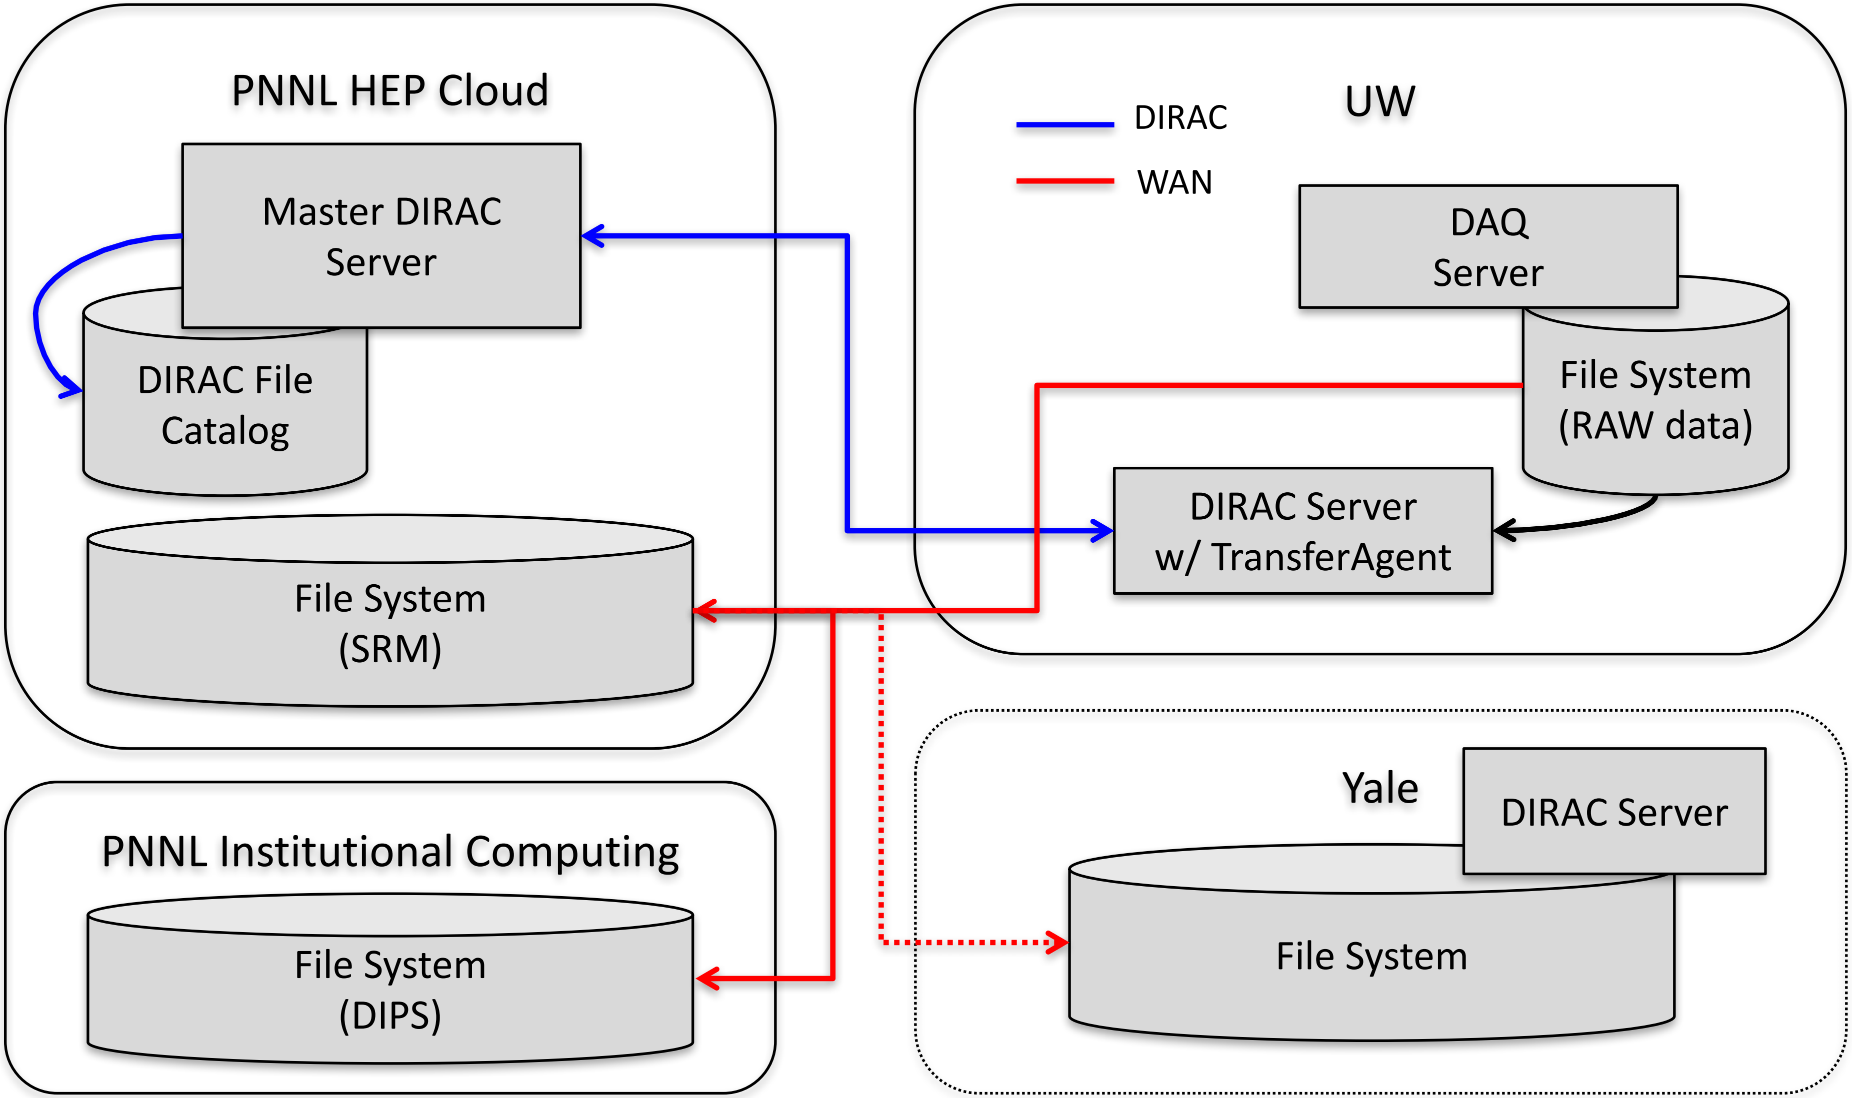

DIRAC servers for Project8

- 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)

Project 8 Dataset Needs

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- Project 8 is still in research and development mode. Our hardware configuration changes frequently, and analysis that is meaningful for one data set may not be meaningful for another, even if all the data is technically in the same format. We therefore have a need for a robust dataset feature to organize the analysis of sub-experiments. The current "dataset" implementation seems close, but we still need:
  - Ability to edit datasets after they are created. Currently have to remove and re-add.
  - Ability to add and remove LFNs "by hand" and not just by metadata cut, *e.g.*, to randomly add background or remove signal to implement a blindness scheme.
  - Custom metadata fields to describe the dataset (*e.g.*, Operator and analyst annotations, links to appropriate e-log entries, data quality flags, list of analyses performed…). This is distinct from metadata that describes the individual data files.

MiniCLEAN

- **High mass WIMP dark matter** detector, 150 kg fiducial liquid argon **technology demonstrator**
- **Simple and Scalable Design**
  - Single phase
  - 4π coverage and cold PMTs within liquid cryogen to maximize light collection
- **Low Background**
  - Pulse shape discrimination using fraction of prompt light in signal
  - Self shielding with fiducialization
- **Target Exchange**
  - Liquid argon and liquid neon
  - Allows intrinsic background removal, dark matter nuclei scattering dependence, neutrinos…

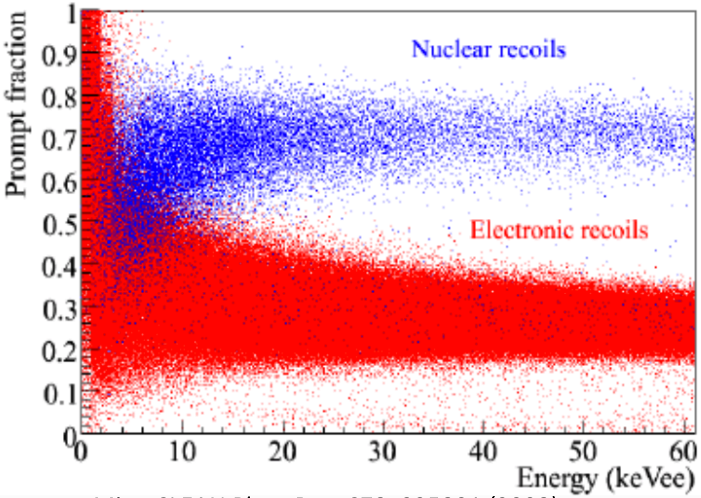

MicroCLEAN Phys. Rev. C78, 035801 (2008)

DIRAC Setup for MiniCLEAN

- Master DIRAC Server @ PNNL:
  - Web Portal: dirac-project8-prod.hep.pnnl.gov:8443/DIRAC
  - DIRAC File Catalog (DFC): File location and metadata
  - Currently using private root CA
- DIRAC slave server at SNOLAB:
  - Dedicated agent to automate data transfers from SNOLAB to PNNL

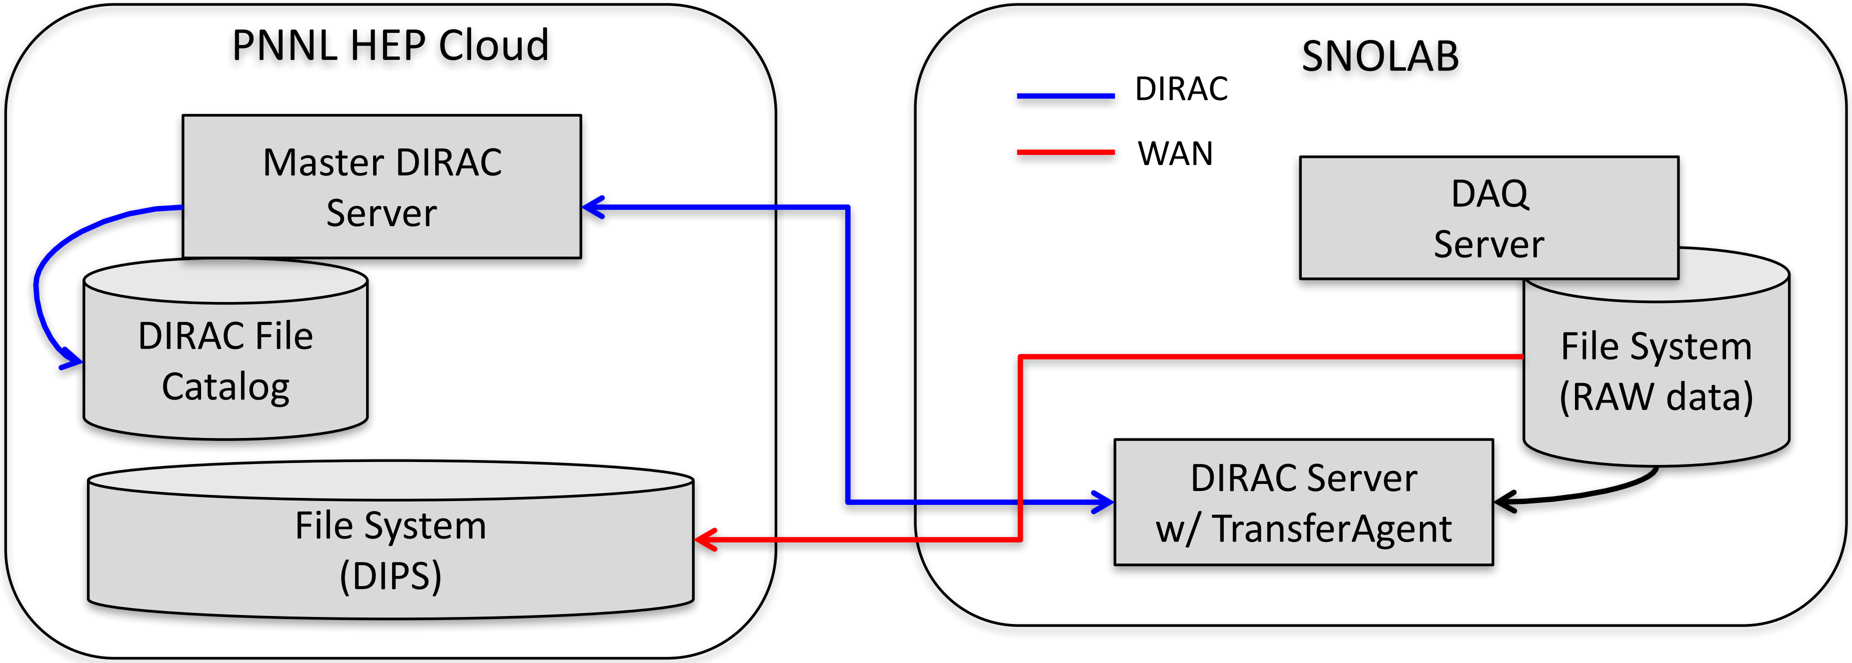

SuperCDMS SNOLAB

- Science goal: Direct detection of ***low-mass*** WIMP-like dark matter

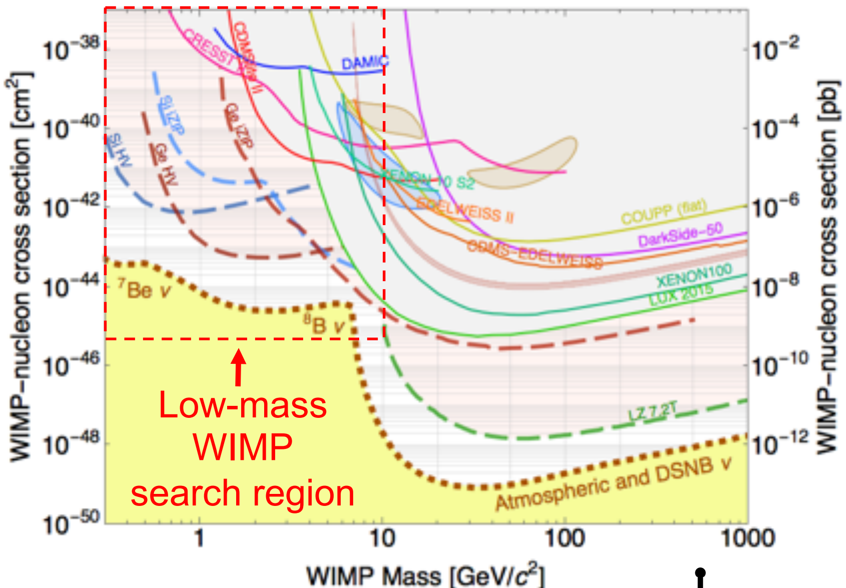

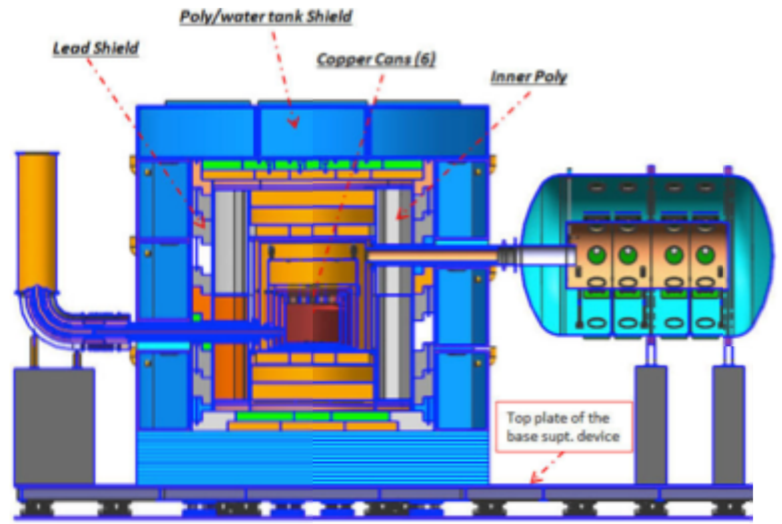

- Facilitated by high voltage (HV) phonon detectors
- Complemented by iZIP recoil discrimination
- Employs Ge & Si targets

HV: Neganov-Luke gain effect

$$P_t = P_r + (E_q/\epsilon)eV = P_r\,(1 + Y \times eV_b/\epsilon)$$

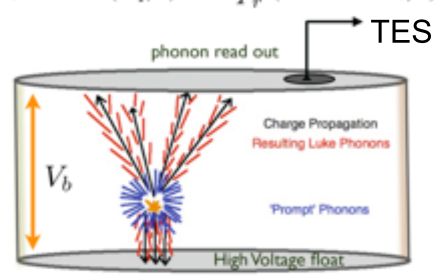

Very-low energy threshold mode proposed by Jeter Hall

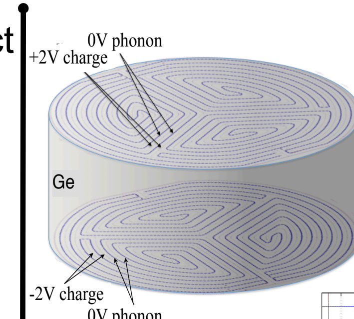

Electron vs. nuclear recoil discrimination

iZIP:
interleaved
Z-sensitive
Ionization
Phonon

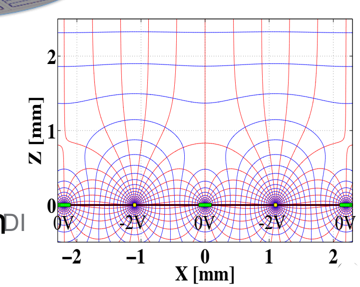

Demo for SuperCDMS @ PNNL

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- DIRAC master server at PNNL for demo
  - 2 CE are configured:
    - PNNL-CE: Access to opportunistic cloud (~3200 cores)
    - FNAL-CE: Access to fermi-grid
  - 3 SE are configured:
    - PNNL-DIP0-SE: local DIRAC storage
    - PNNL-PIC-SE: 100TB disk and 100TB tape allocation
    - FERMILAB-SE: Using volatile storage for now
  - DIRAC File Catalog (DFC)
- CVMFS Stratum-0 and Stratum-1:
  - Initially setup to allow PNNL to run jobs on several internal clusters
    - 1M core hours on PNNL Institutional Computing (PIC)
    - Opportunistic HEP Cloud
- Simulation workflows was successfully tested

Summary of PNNL DIRAC work

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- We are using and developing DIRAC components for 4 active project
- We performed a DIRAC demo and evaluations for SuperCDMS experiment
- We are now working with the PNNL nEXO group to determine if DIRAC is a good solution for them