LUX-ZEPLIN: data-intensive search for Dark Matter

Maria Elena Monzani

HEP-CCE All-Hands Mtg.
Dec. 19 2023

The LUX-ZEPLIN (LZ) Dark Matter Experiment

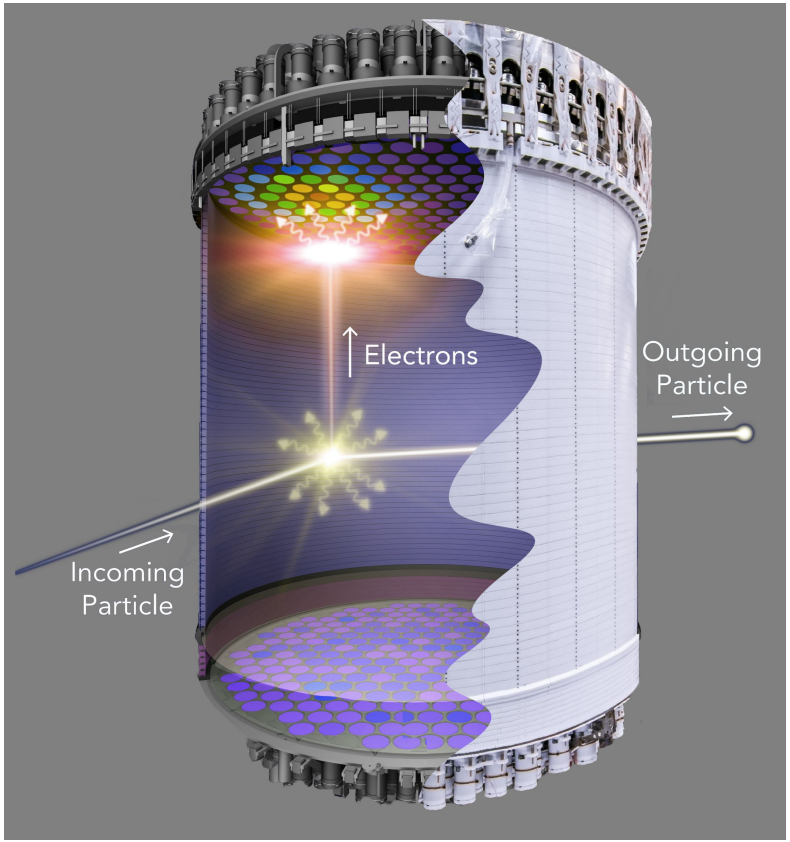

LZ is a 10-ton Liquid Xenon TPC

- Located underground at SURF, South Dakota
- Initial science run data in winter/spring 2022
- Set world-record WIMP sensitivity in July 2022
- (5 weeks turnaround between run and results)!
- LZ data is stored and processed at NERSC

Data Throughput (order of magnitude)

- Fermi-LAT (>2008): 0.3 PB/year
- LZ (2021-2028+): 3.5 PB/year, 7+ years
- ATLAS (>2010): 3.2 PB/year (raw)
- PS: extreme “needle in a haystack” problem!

Construction and Data Taking Timeline

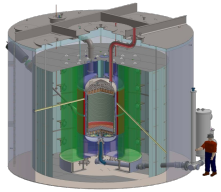

CD3 and TDR

TPC Complete

Underground

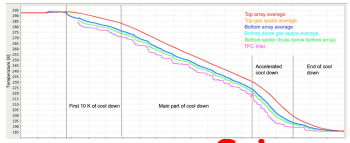

Cold GXe

Science Running!

MDC1

MDC2

MDC3

2017 2018 2019 2020 2021 2022 2023

Cryostat arrives

Sealed Up

LXe fill

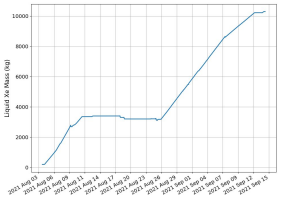

First results!

World-leading WIMP sensitivity (July 10, 2022)

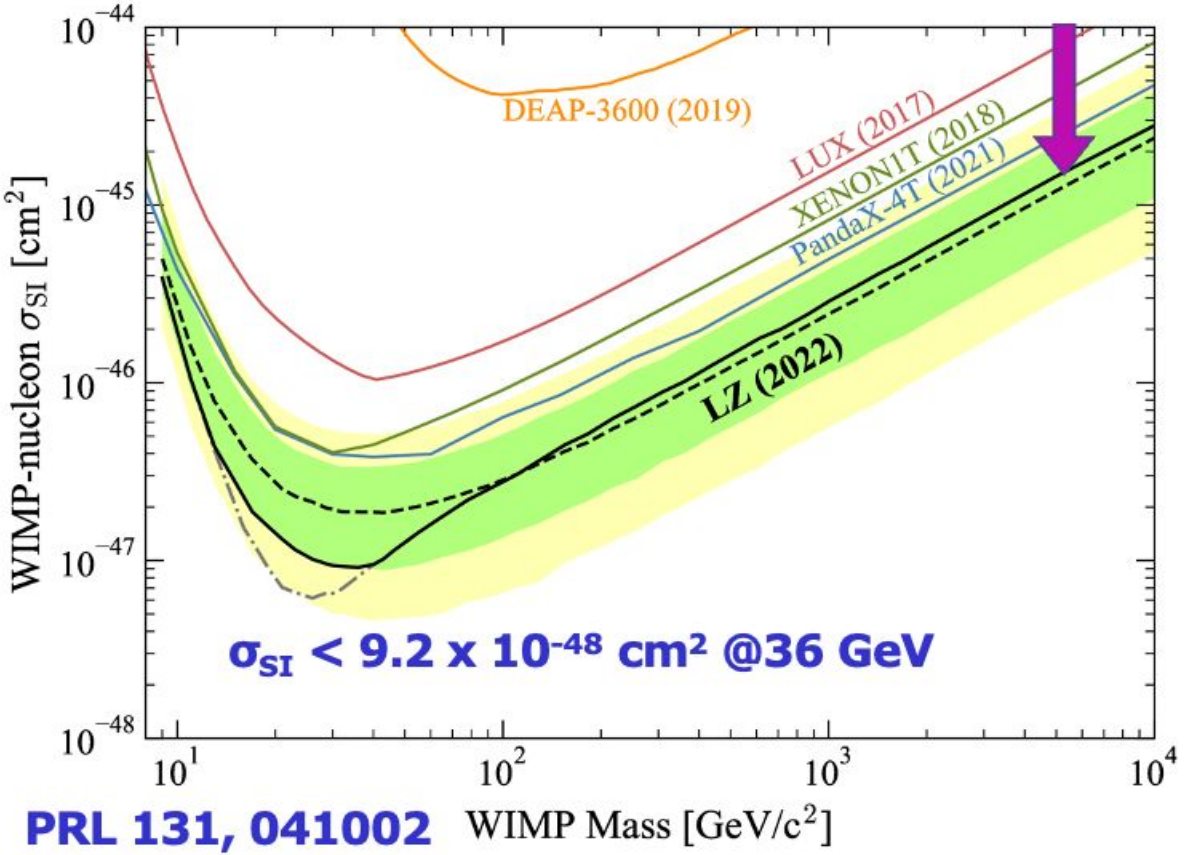

An extension through 2028 was endorsed by P5

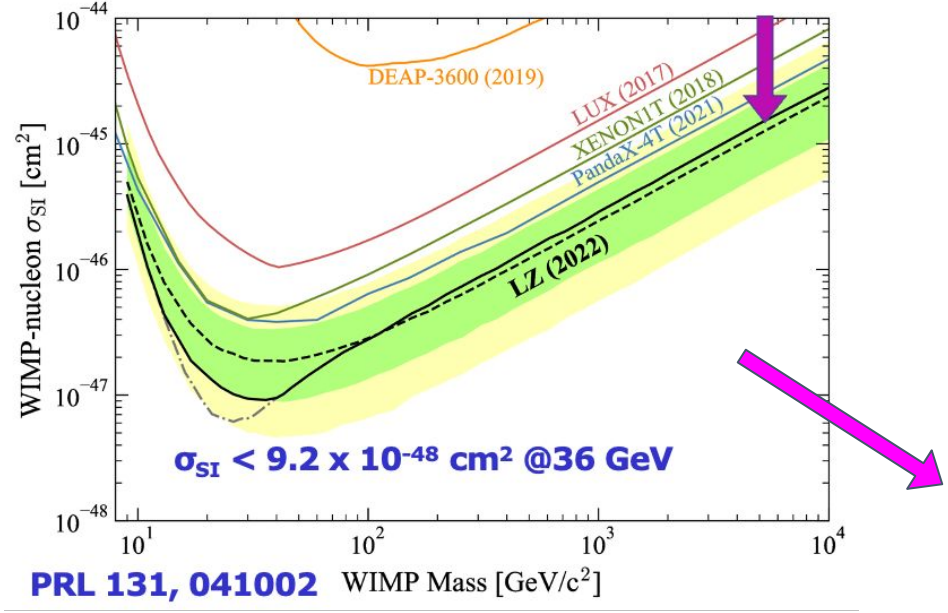

PRL 131, 041002

x15 more data in the full run!

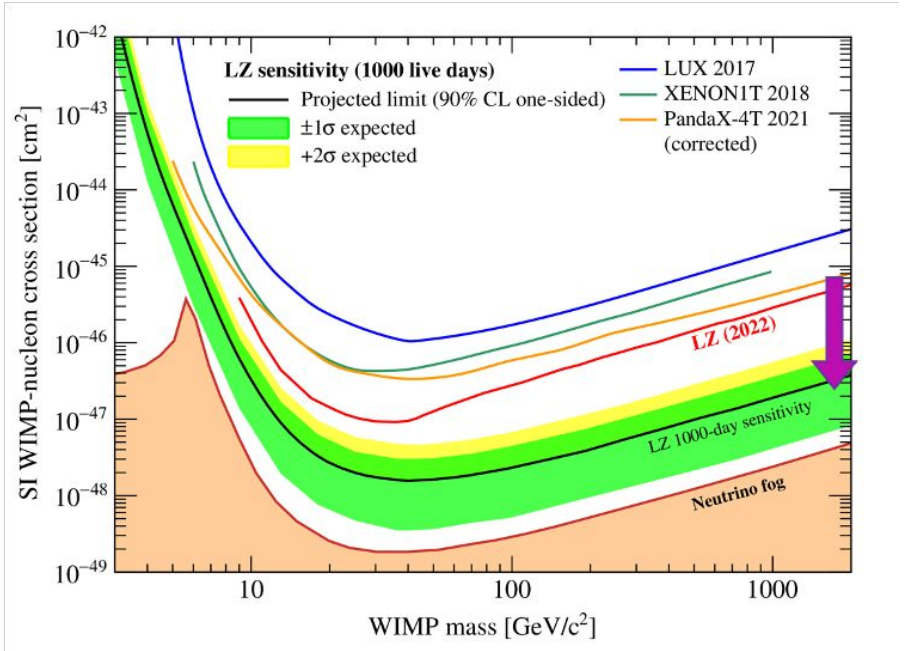

LZ: Offline Computing and Software

Data is staged at SURF and transferred to the remote data centers

- Fully redundant data center design (each site can run data processing and simulation production… and store a complete copy of all the data)!
- Data rate: ~3 PB/year, including raw, reconstructed, calibrations, etc.
- All detector data are processed automatically 24/7 at the USDC.
- Data can be reprocessed on-demand based on calibrations and analysis.
- Reconstructed and simulated data is then made available to all analyzers.

Temporary Storage:

SANFORD LAB HOMESTAKE

US Data Center (USDC):

- Prompt Processing
- Long-term Archiving
- Supercomputers!

NERSC

UK Data Center (UKDC):

- Data Reprocessing
- Sims Production
- Distributed CPUs!

GridPP
UK Computing for Particle Physics

Reconstructed & simulated data can be analyzed at either data center

- NERSC and GridPP have diverging CPU architectures. All LZ software & analysis tools can run seamlessly on either architecture.
- System choice is based on user preference, but several team members have become proficient at both supercomputers and distributed computing.

Offline Requirements and Design Principles

Store all raw & reconstructed data from LZ

- 2 “live” copies of all raw & reconstructed data at NERSC and UKDC
- 1 “tape” archive of all raw data at NERSC before bias mitigation
- At least 1 backup of all versions of reconstructed data at NERSC

Process detector data early and often

- Automatic prompt-processing at USDC upon data reception
- Redundant capabilities to reprocess/simulate multiple times based on calibration/analysis results (rerun 1 year of data in 1 month)

Time is of the essence! Rapid (<1 day) turnaround

- Very limited computing resources are available at SURF (RAID array for storage and “first look” online quality monitoring tool)
- Full-scale detector health assessment happens at NERSC. Quasi-real-time analysis feedback during commissioning

US Data Center (USDC):

- Prompt Processing
- Long-term Archiving
- Supercomputers!

NERSC

UK Data Center (UKDC):

- Data Reprocessing
- Sims Production
- Distributed CPUs!

GridPP
UK Computing for Particle Physics

Resilience, Reliability, Robustness

Superfacility uptime: uptime of LZ x uptime of NERSC services

- Downtime is expensive:
  - Defensive Engineering
  - Reputation with Science partners
- We use so many part of NERSC, downtime or degradation anywhere (DTNs, CFS, Slurm, SPIN, etc.) impacts entire workflow

Impact on commissioning, operations, calibration, detector health

- Example: SURF underground days are Mon-Thu or Tue-Fri
- We performed the leveling of the detector on a Mon-Thu week
- However, there was a scheduled Cori outage that Wed
- We needed to be able to look at/analyze data every night
- (heroic effort from NERSC to keep us running on Gerty that week)

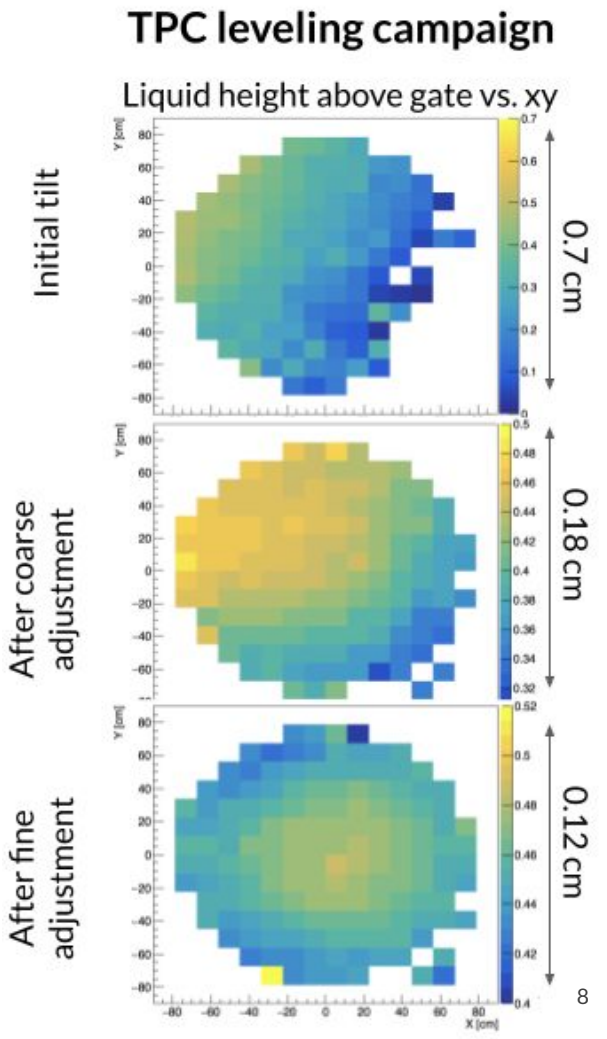

Sustainability of LZ computing to 2028+

Portable workflows

- Goal: maximize uptime, guarantee fast turnaround (<1 day)
- Plan: a “backup” system in the US to mitigate NERSC downtime

Optimizing data storage

- Portable workflows require exposing/exporting the dataset
- Robustness of data movement to UKDC is a crucial need

Scaling up HEP AI/ML applications

- Extreme “needle in a haystack” problem for DM identification

Accelerating HEP simulations

- Simulation time is dominated by raytracing of optical photons
- Offload raytracing to the GPU (crucial to design G3 experiment)

1. Workflow Portability Pilot

Multiple options for alternate data center (s) on the US

- ANL: hoping similar interfaces and protocols to NERSC
- FermiGrid: simplify the data movement issues with GridPP
- SLAC S3DF: same architecture as Perlmutter (AMD Milan)
    - Additional benefit: synergies with DESC and LCLS-II

Workflow portability pilot

Maria Elena Monzani | October 20, 2023

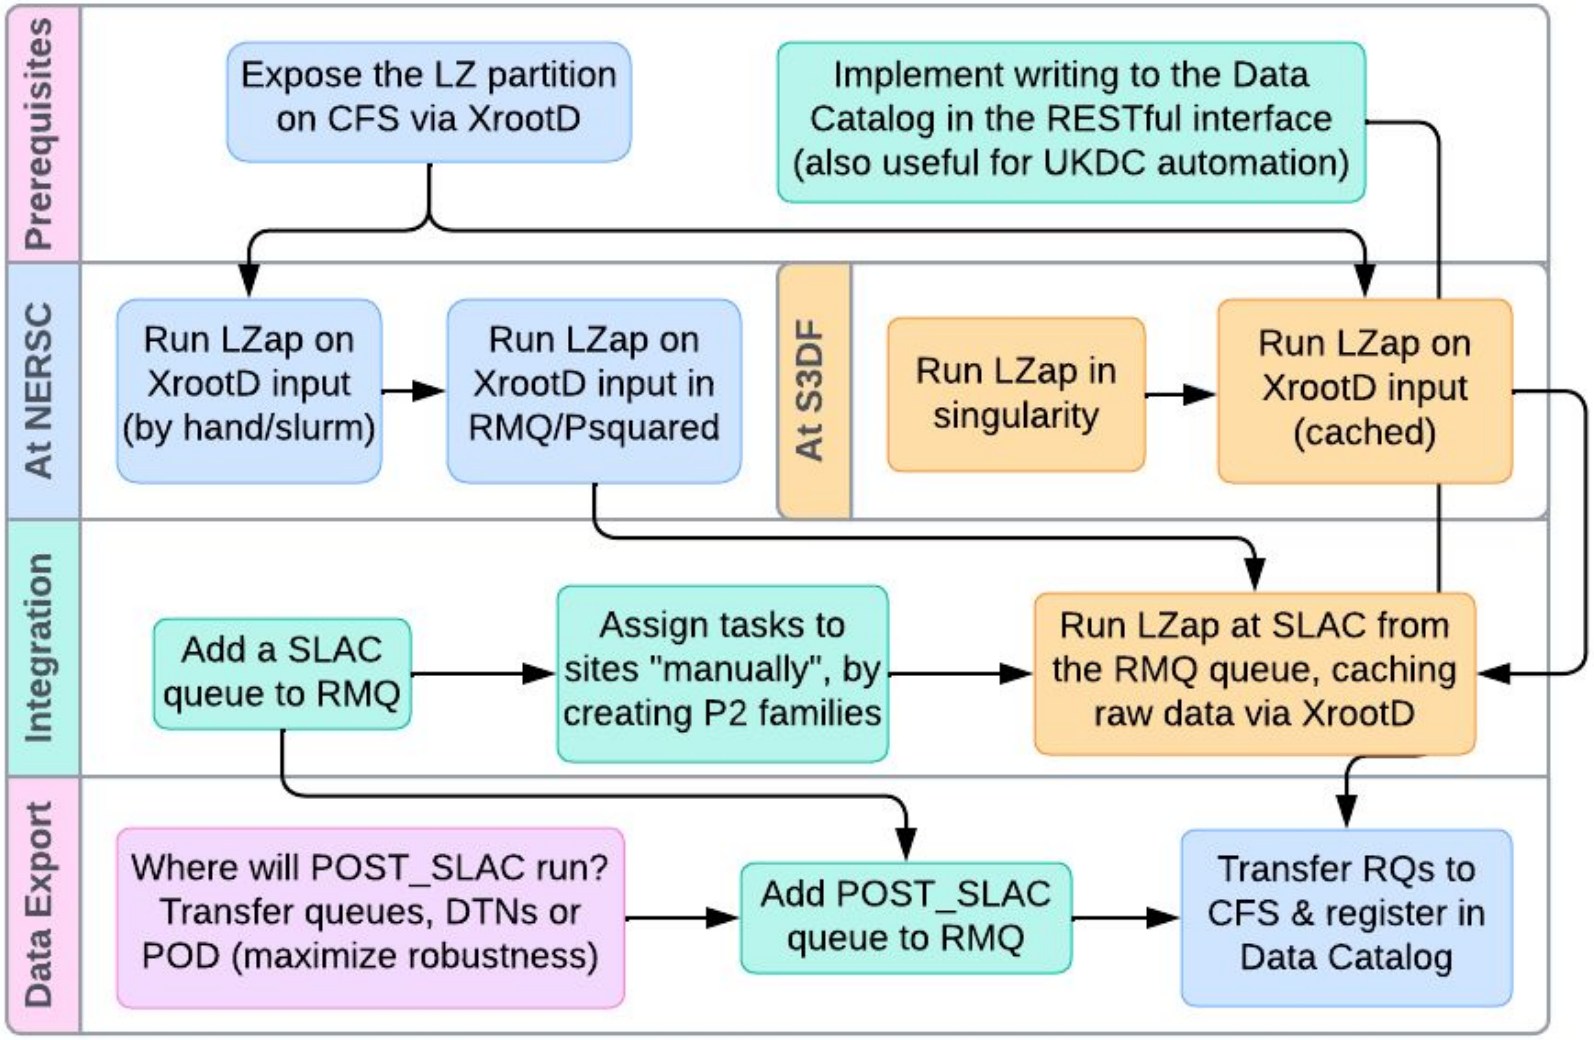

1. Workflow Portability Pilot

Multiple options for alternate data center (s) on the US

- ANL: hoping similar interfaces and protocols to NERSC
- FermiGrid: simplify the data movement issues with GridPP
- SLAC S3DF: same architecture as Perlmutter (AMD Milan)
    - Additional benefit: synergies with DESC and LCLS-II

2. Data Storage/Data Movement

Data Movement is a vulnerability in our infrastructure

- Expose all datasets via xrootd, to support portable workflow
- If pilot is successful, automate “exposure” of datasets
- Upgrade data movement framework from SPADE to Rucio
- Improve integration with GridPP (diverging IM, certificates)

3. Scaling up HEP AI/ML applications

Extreme needle in a haystack problem:

- Identify a handful of DM events (if nature cooperates)
- Expected background is of order ~5-10 billion events
- Background rejection problem with a rarity of order 10^{-9}
- Ideal playground for the development of novel ML algorithms
- Rare/unmodeled backgrounds can spoil bias mitigation schema

Approach: anomaly detection at the 10^{-9} sensitivity

- Collaboration with Stanford ICME (School of Engineering)
- Tools: event clustering and resilient-VAEs (in recursive mode)
- Challenge: train ML models on the waveform (multi-PB dataset)
- There are currently no machines with a multi-PB scale RAM

UMAP + DBSCAN (credit: Maris Arthurs)

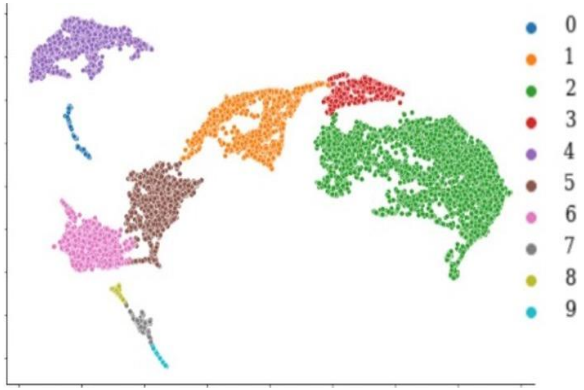

VAE on full WF (credit: Tyler Anderson)

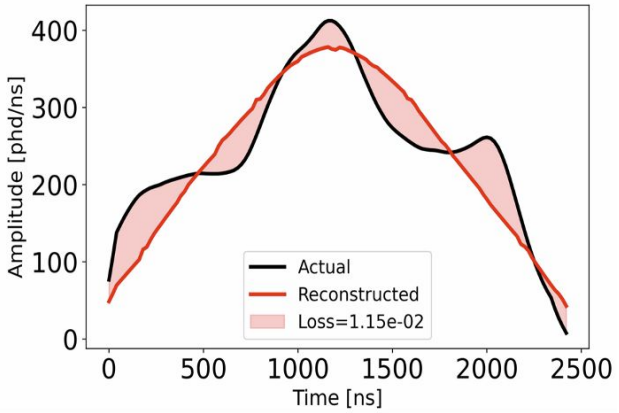

What will happen after LZ?

LZ taking data through 2028. Analysis through 2030?

P5 endorsed an “ultimate” Dark Matter experiment

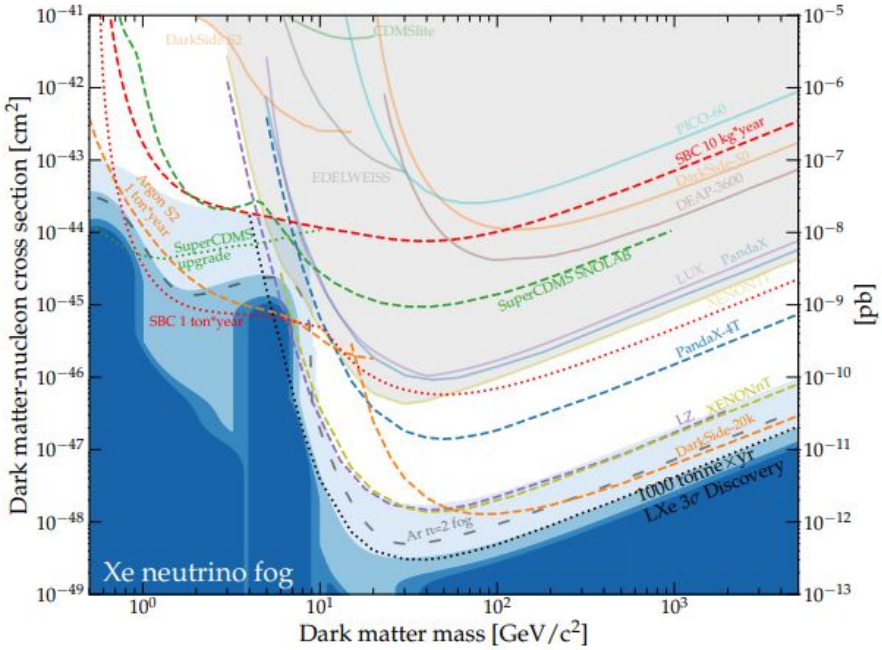

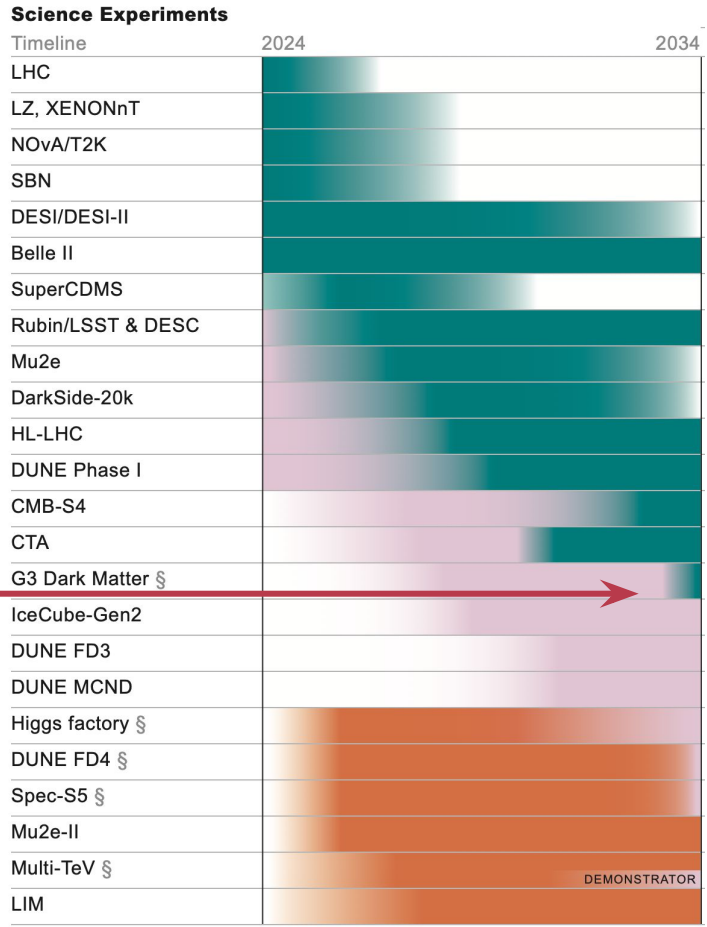

4. Accelerating HEP Simulations

LZ taking data through 2028. Analysis through 2030?

P5 endorsed an "ultimate" Dark Matter experiment

- Multi-purpose observatory for a multitude of dark matter models, neutrinoless double beta decay, and astrophysical neutrinos
- Fully probe WIMP parameter space into the neutrino fog (50-100 tonne experiment)
- A x10 scale-up from LZ: will need accurate simulations to design the "ultimate" experiment
- This level of accuracy requires raytracing on the GPU, which is needed in the next ~few years

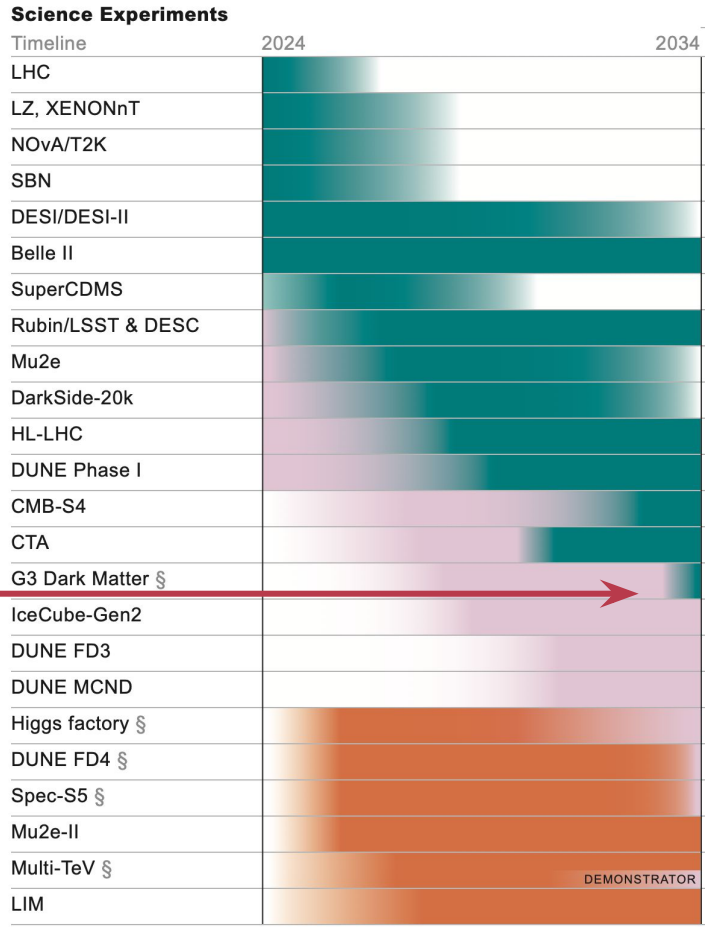

Let's work together! (pretty please?)

We are a "small experiment" with a "large dataset"

- Long-range sustainability plan is not itself sustainable…
- …unless we collaborate with other teams and experiments
- …and potentially with computing centers beyond NERSC

Excellent alignment of LZ needs with HEP-CCE themes

- Workflow portability is crucial to LZ's realtime needs
- Data storage/data movement upgrades are also needed
- Ideal testbed for ML at scale (full multi-PB raw dataset)
- Raytracing on the GPU needed for G3 detector design

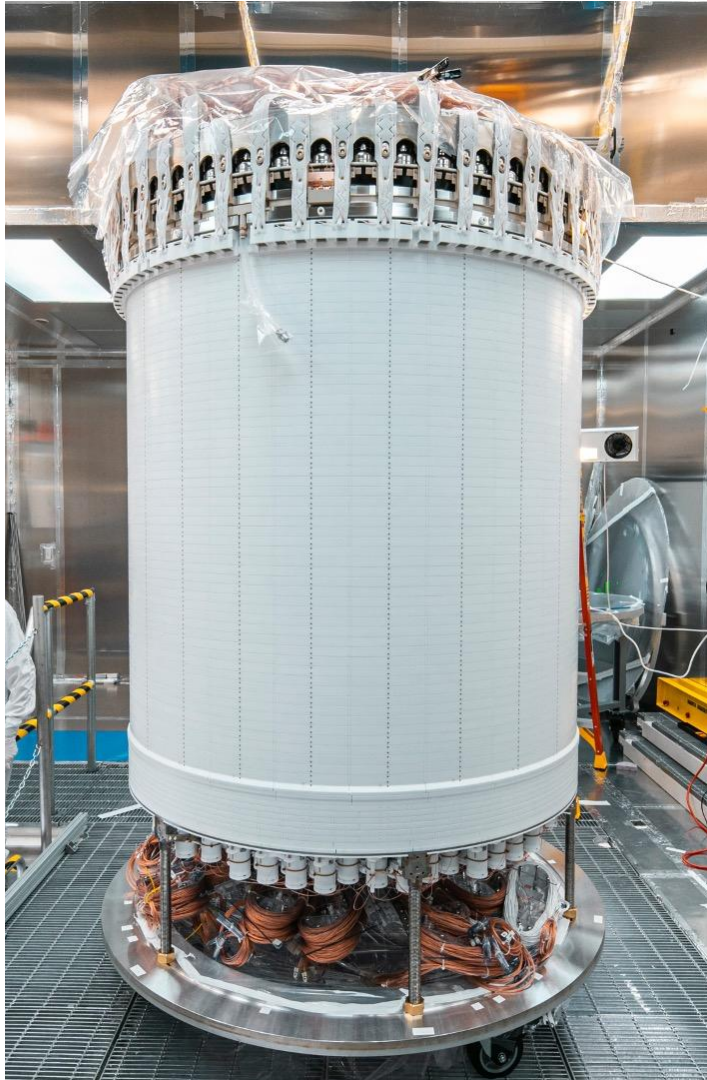

The LZ Collaboration at SURF in 2019 (220 NERSC users)

Thanks to our sponsors and participating institutions!

U.S. Department of Energy
Office of Science

Science and Technology Facilities Council

Sanford Underground Research Facility
South Dakota Science and Technology Authority

FCT Fundação para a Ciência e a Tecnologia

Institute for Basic Science

Backup slides: SPIN & detector pictures

Extensive use of SPIN-based services

Supporting both production tools and user access!

- Data transfer (SPADE)
- Job submission engine (PSQUARED)
- Monitoring data movement and processing (SPADE/PSQUARED)
- Offline event viewer
- PREM (Offline Data Quality Monitor)
- Databases, database mirrors, and associated web service interfaces
- Data Catalog and its interfaces
- Code Quality and Software Release validation
- Web Services use SAML/NGINX authentication tools

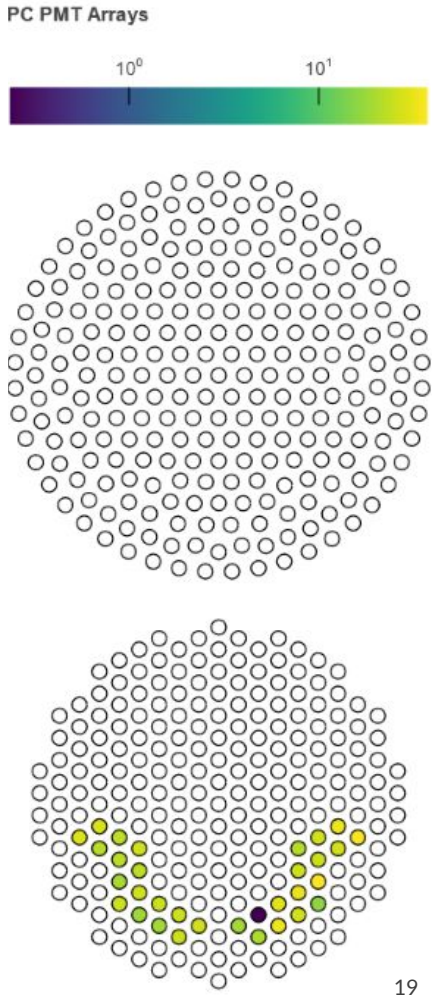

Monitoring Data Movement

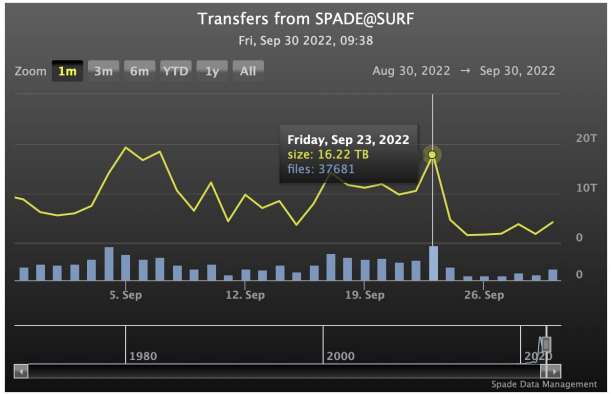

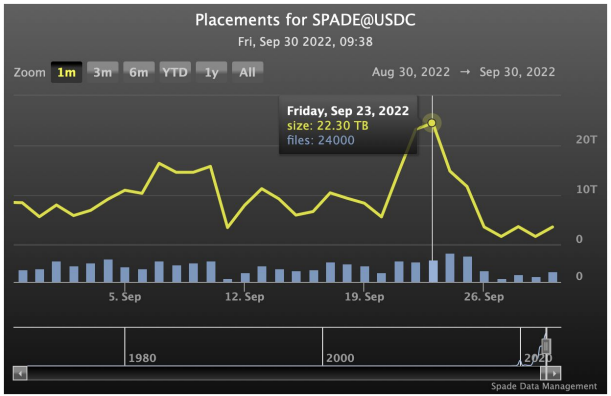

Heartbeat Monitoring

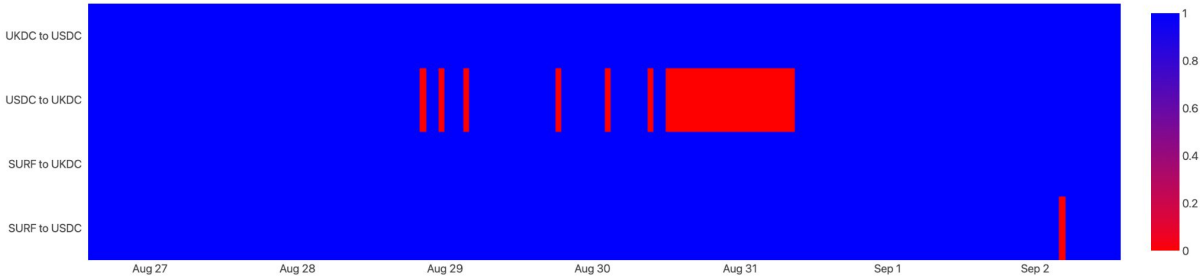

Data Transfer Latency

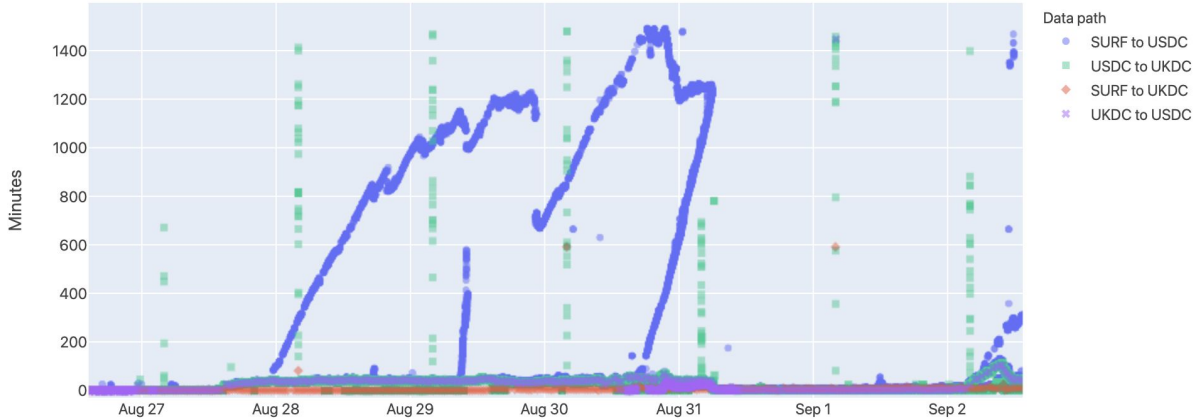

Monitoring Prompt Processing

Family	Samples	Status	Progress
preSR2_autodata_Y2022_W33_PROD-1	41		622 / 2111 / 1 + expand
preSR2_autodata_Y2022_W32_PROD-1	122		2010 / 482 / 5475 / 71 + expand
preSR2_autodata_Y2022_W31_PROD-1	101		90141 / 4411 / 19 + expand
preSR2_autodata_Y2022_W30_PROD-1	105		706 / 3494 / 145 + expand

Overview of Processing Status, by Week/Year

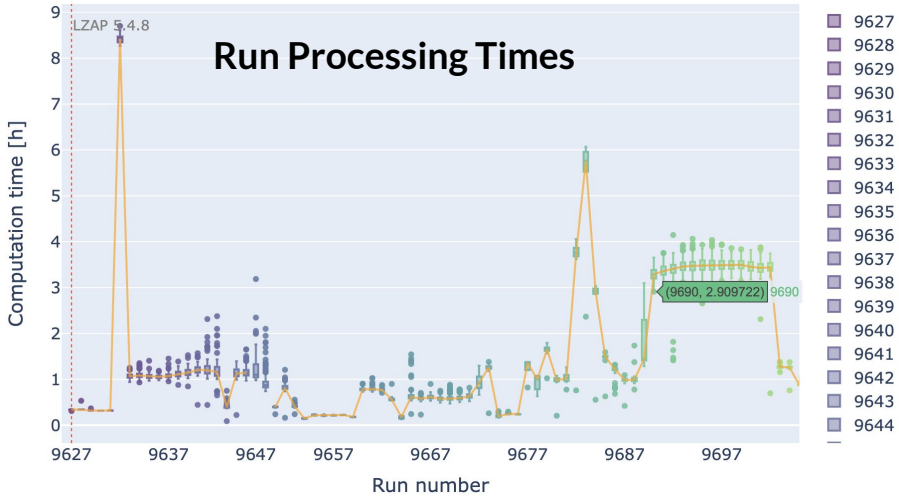

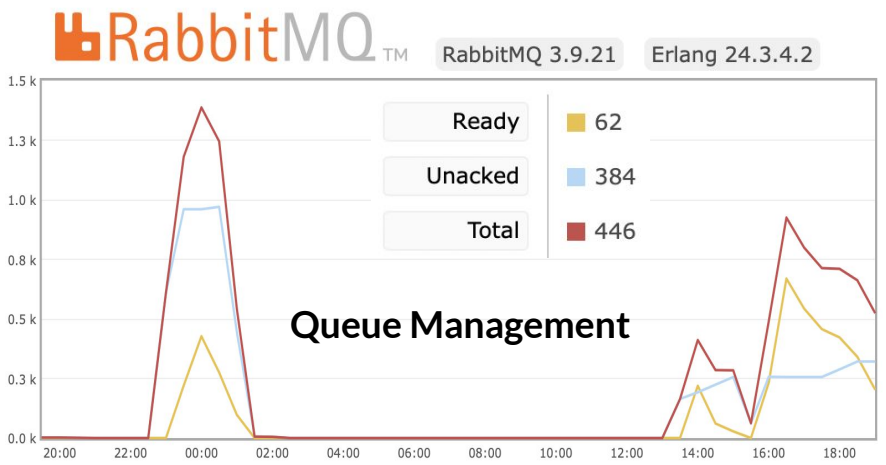

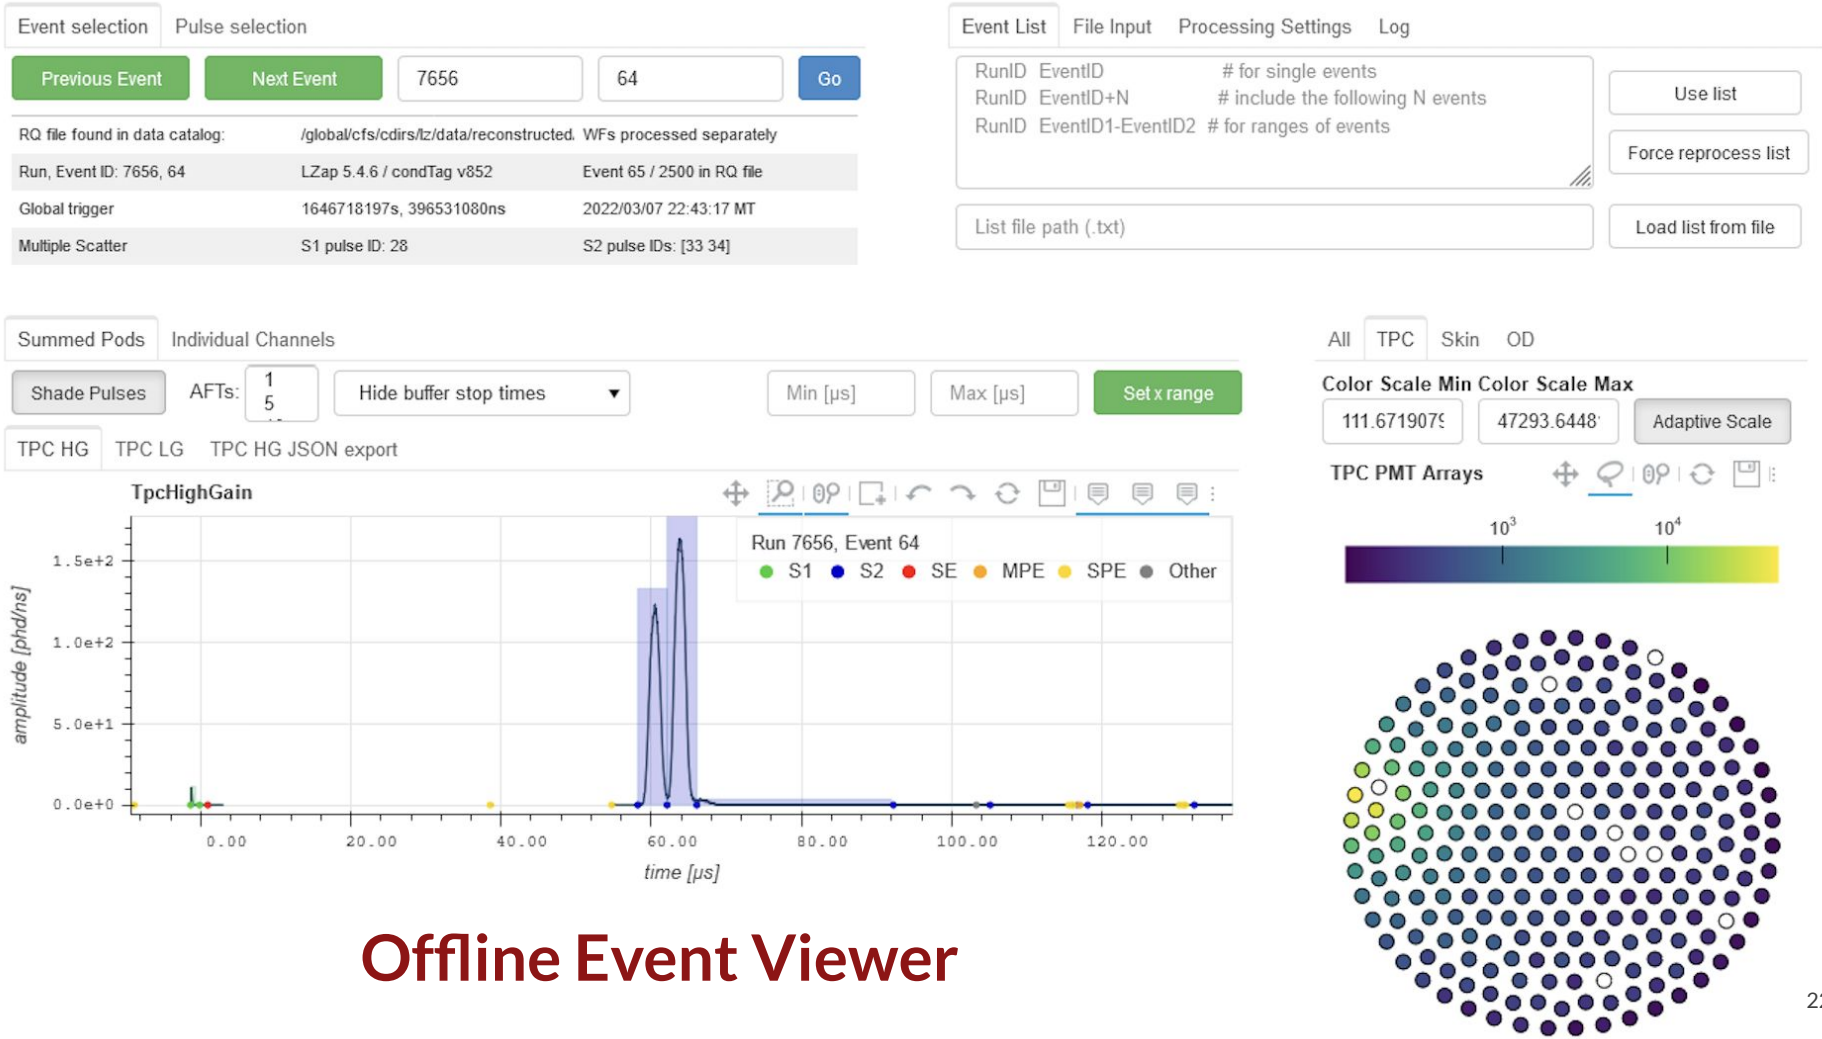

Offline Event Viewer

The Offline Event Viewer is Extremely Popular in LZ!

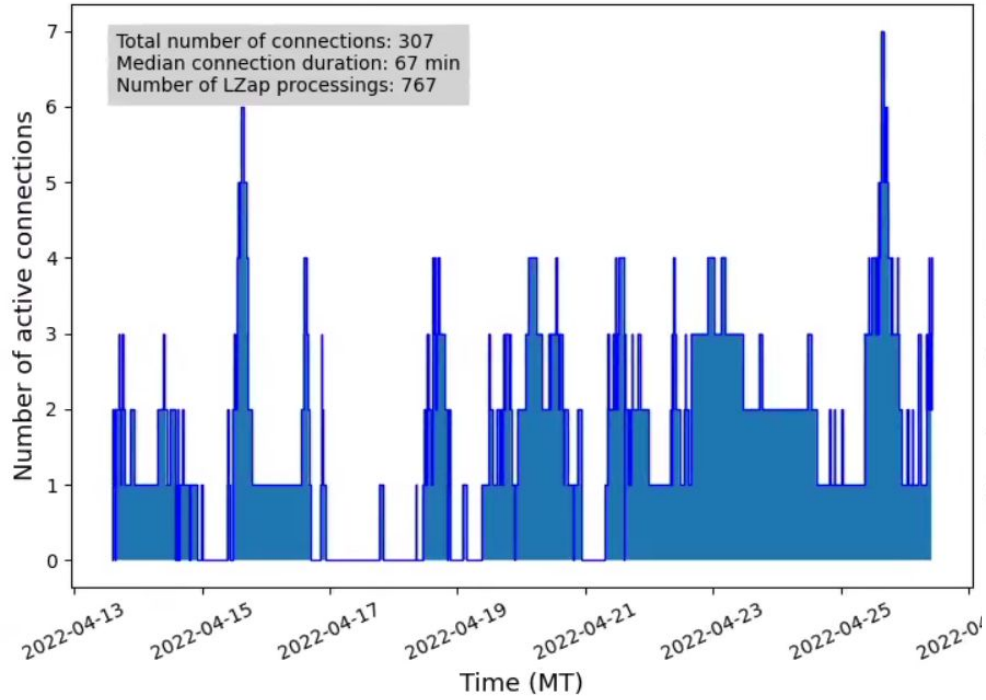

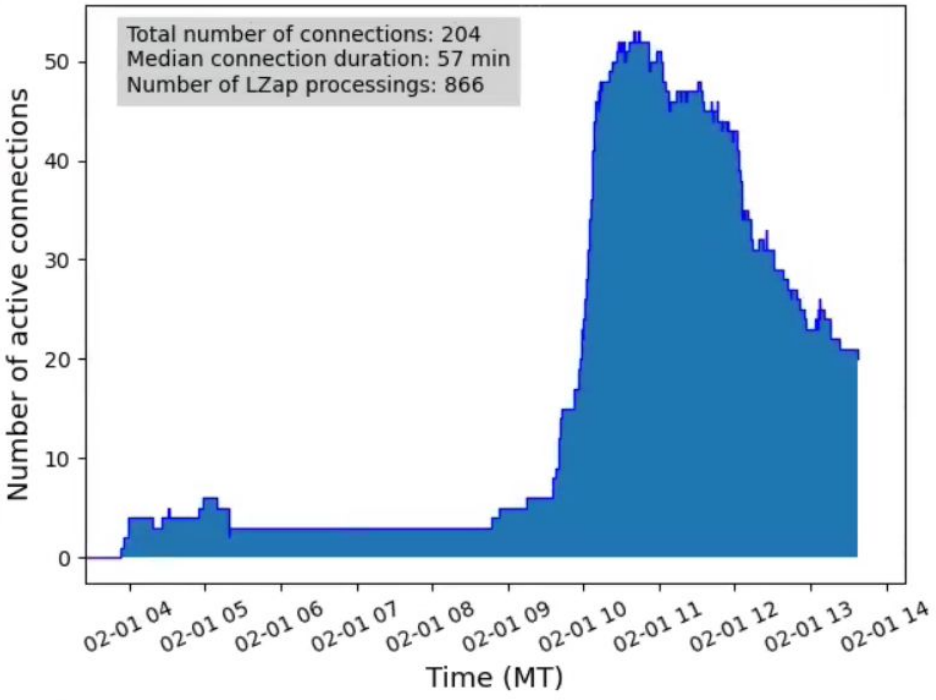

LET'S LOOK AT SOME WAVEFORMS

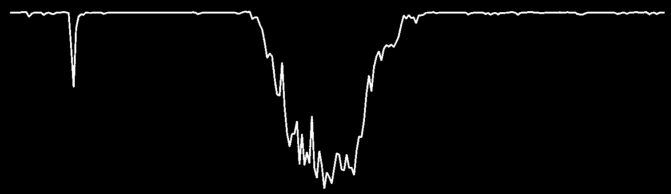

Please enjoy some pretty pictures from LZ detector construction


HOMESTAKE

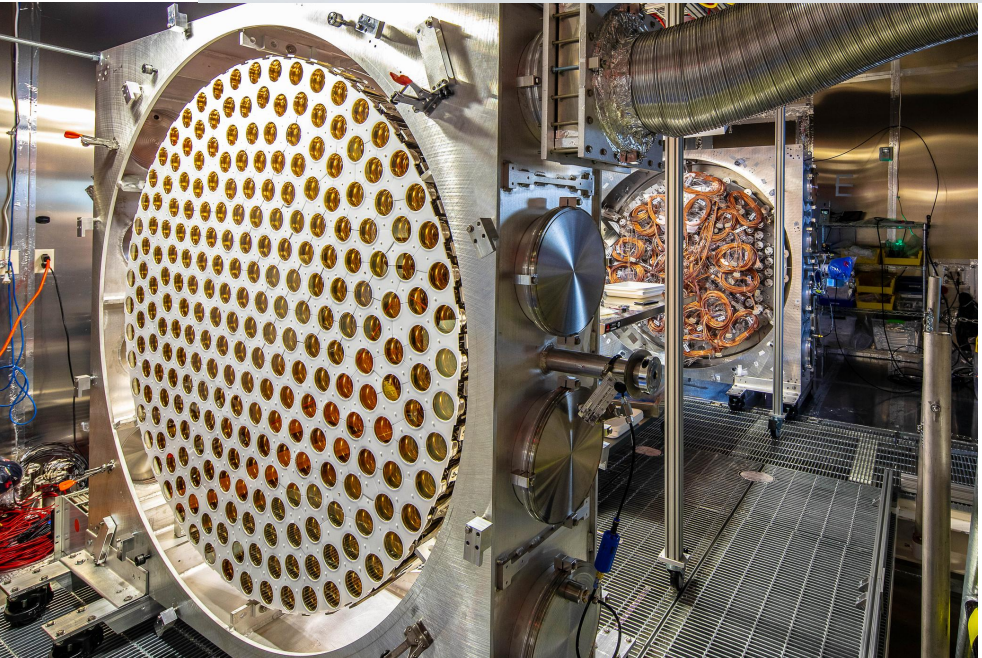

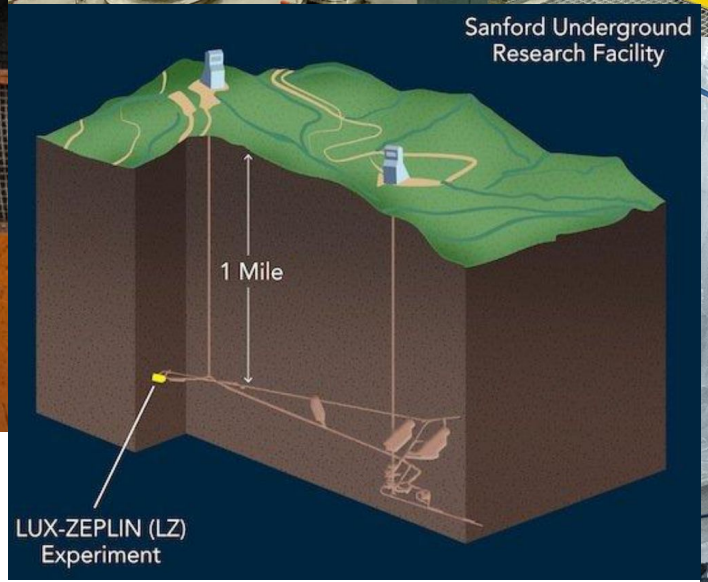
Sanford Underground
Research Facility
1 Mile
LUX-ZEPLIN (LZ)
Experiment