DUNE analysis training

Heidi Schellman and David Demuth amd Claire David for the collaboration

July 28,2022

Contacts

- Computing Consortium Lead: Heidi Schellman (Heidi.Schellman@oregonstate.edu)
- Team Training & Documentation: Claire David (cdavid@fnal.gov) , Michael Kirby (kirby@fnal.gov), David DeMuth (david.demuth@vcsu.edu) , Tom Junk (trk@fnal.gov), Ken Herner (kherner@fnal.gov) Steve Timm (timm@fnal.gov)

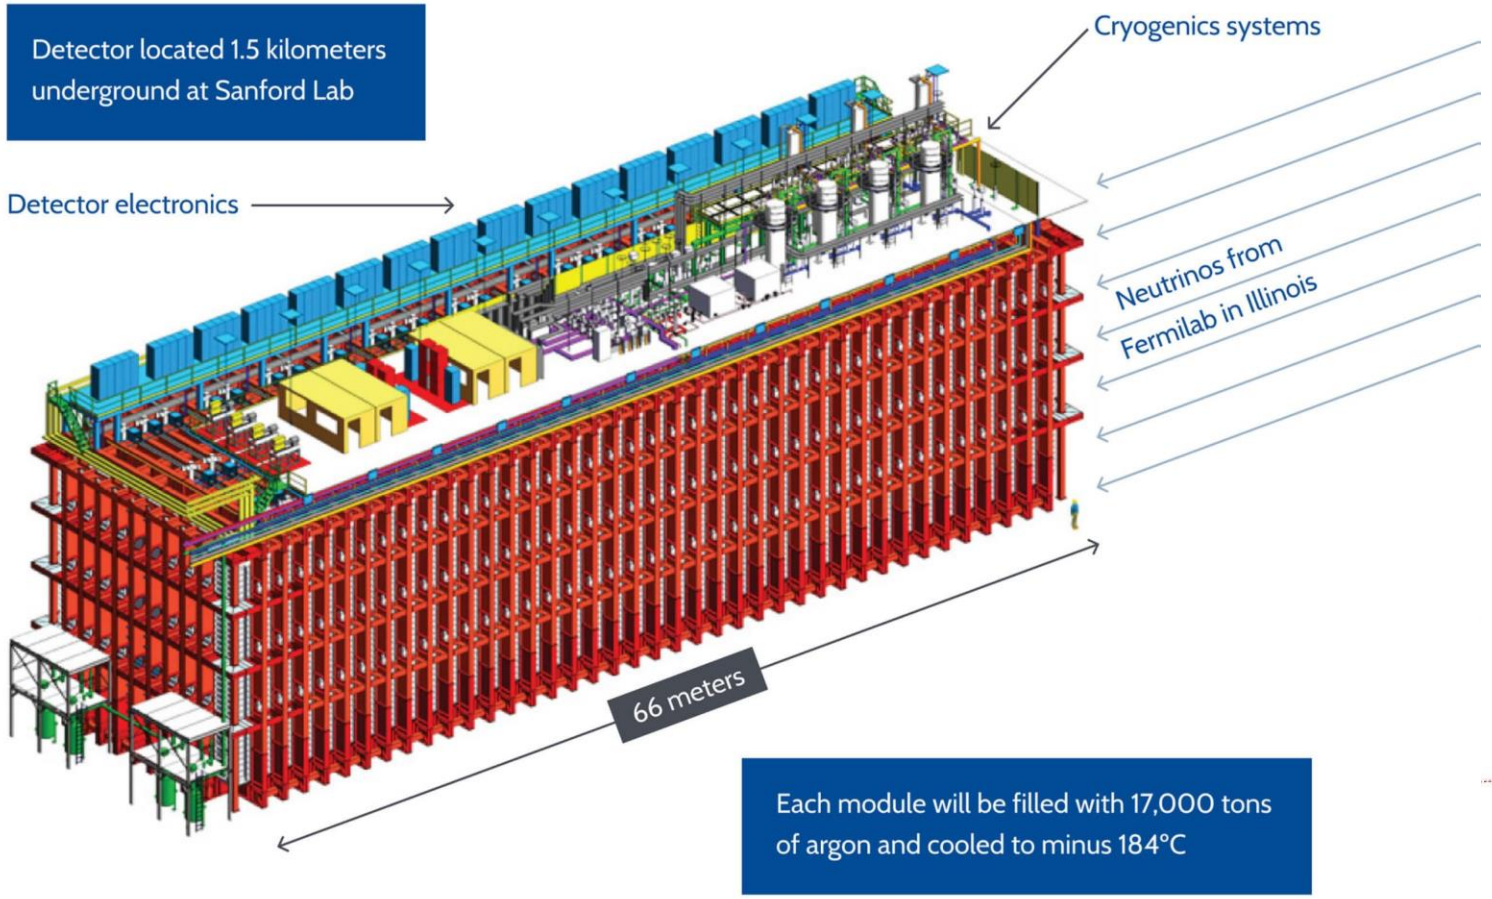

DEEP UNDERGROUND NEUTRINO EXPERIMENT

DEEP UNDERGROUND NEUTRINO EXPERIMENT

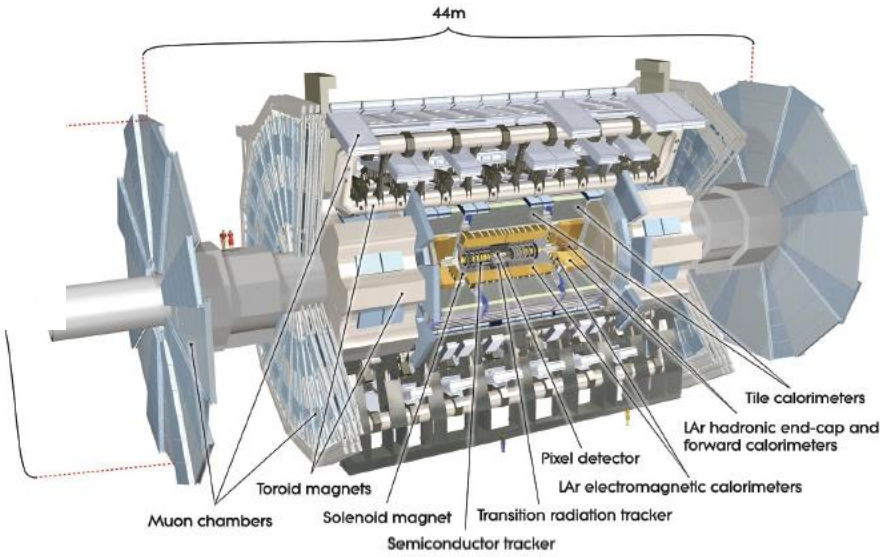

Heidi Schellman
Oregon State University for the DUNE collaboration

7/26/2022

What is DUNE

- A suite of large liquid Argon TPC with with electronic and photon readout + a near detector
- Ran prototypes at CERN in 2018-2020
- Currently planning a second prototype run in late 2022 of 2 700 T prototypes
- A suite of other smaller prototype detectors are running or planned at several institutions
- The full detector will start operating in 2029 with 1 and then 2 modules and a near detector.

DUNE offline computing

- Future: 30 PB of raw data per year
- ~1500 scientists, ~700 have active computing accounts
- A few million very large trigger records/year (4-8 GB per “event”)
  - ~ 5x5x1 mm^3 voxels over a volume of $4x10^4$ m^3
  - TPC data for the first far detector module are 450 (5000x2560x14 bit) 2-D arrays. Second module is slightly larger.
- Currently setting up to run prototypes at CERN - data rates of 2-3 GB/s

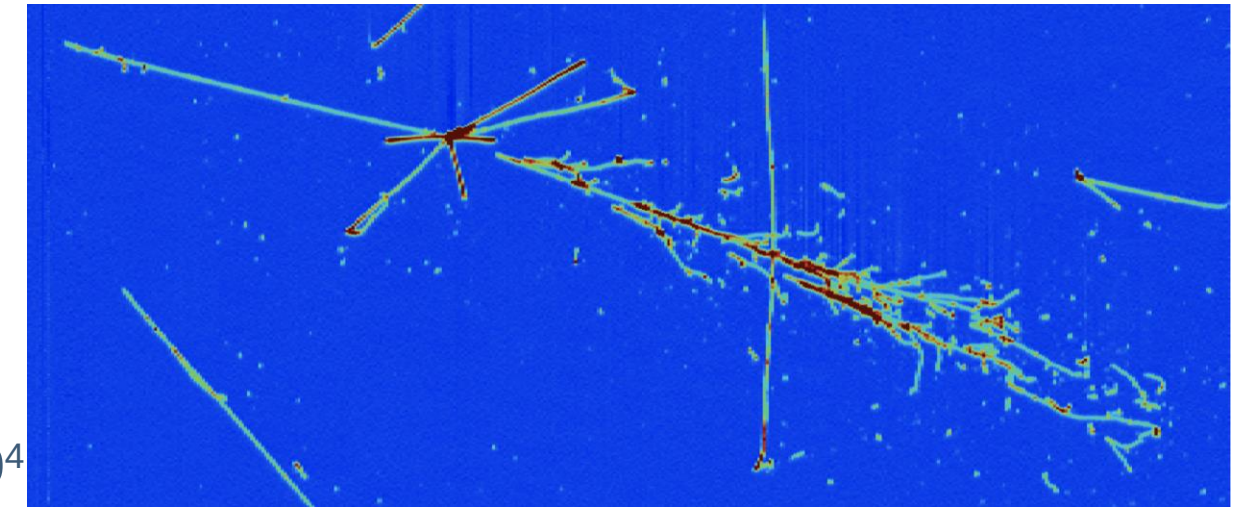

Test beam in a subsection of one of the 450 readout planes

Shared and Unique Tools and Challenges

Shared

- OSG/WLCG
- CVMFS
- ROOT
- Rucio
- GEANT4
- LArSoft (joint with many others)
- Neutrino event generators
- People – lots of them
- Many training materials

Unique

- Data are produced 1300 km away almost a mile down in a mine with limited space and utilities
- Memory management with 4 GB “events”
- Supernova→ ½ PB in 100 seconds
- Calibration of 4x10,000 m^3 of flowing charged liquid

Disk compared to CMS

DUNE is ~10% of an HL-LHC expt.

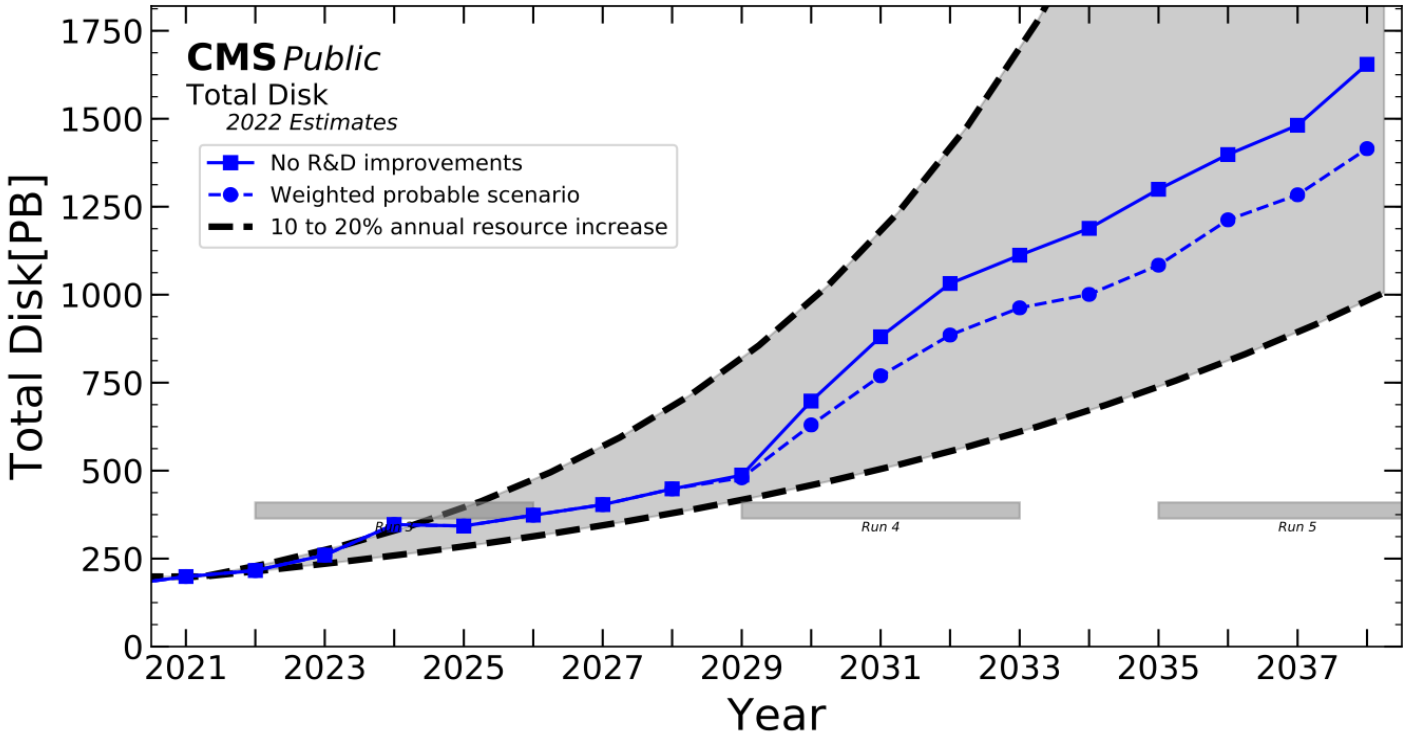

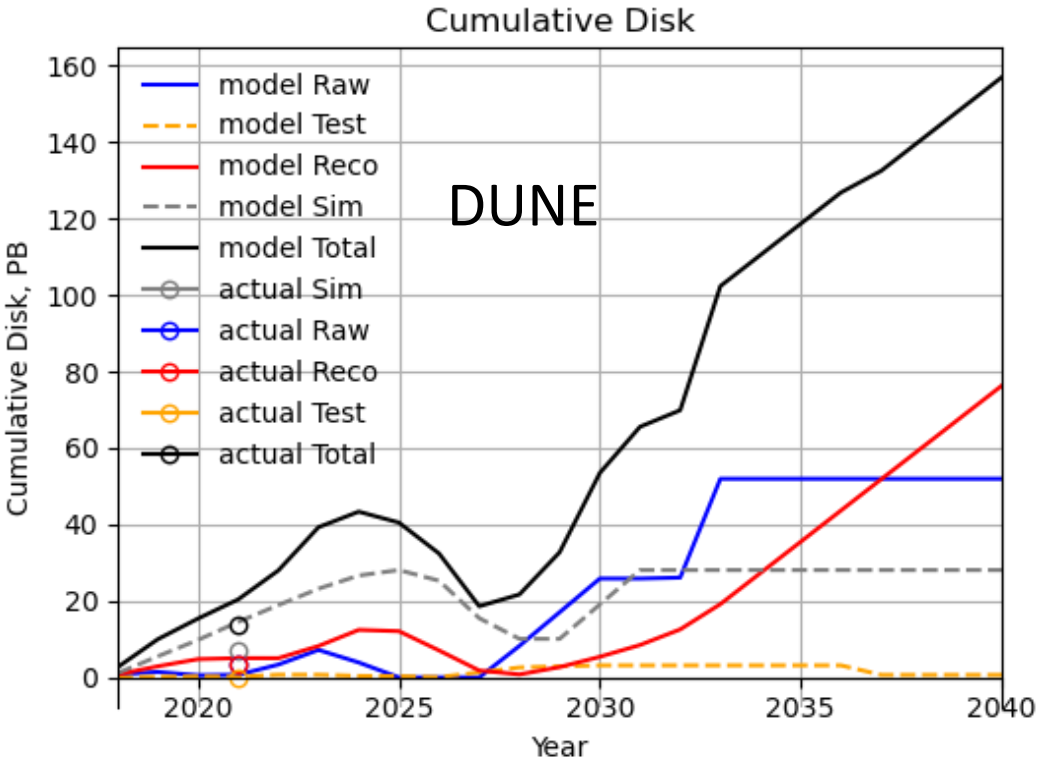

DUNE Users

- As an experiment that is starting up, much of our effort is in **development of simulation and reconstruction** algorithms.
- New users need to know how to **run and modify the reconstruction and simulation algorithms**, not just run final sample analysis.
- At the same time, students and postdocs need to do analysis and write papers. Often on other experiments. **Most DUNE collaborators are on at least one other experiment.**

What HEP analyzers need to know (computing). Items in brown were covered in our tutorial

- ❑C++ (take a real course)
- ❑Python python-novice-inflammation
- ❑unix shell shell-novice
- ❑code standards
- ❑basic graphics python-novice-gapminder
- ❑compilers
- ❑linkers
- ❑build systems such as cmake make-novice
- ❑code repositories such as git git-novice
- ❑component and continuous integration testing
- ❑code release strategies
- ❑debuggers and profilers
- ❑basics of computer security
- ❑basics of operating systems
- ❑basics of storage systems
- ❑basics of data curation
- ❑basics of distributed computing
- ❑basics of workflows
- ❑basics of networking
- ❑basics of databases sql-novice-survey
- ❑basics of algorithms
- ❑basics of data structures
- ❑basics of machine learning
- ❑parallel computing
- ❑GPUs and FPGAs
- ❑Containers

Here basics means knowing what they are and why they exist

And after all of that: Actually doing physics

General

- ❑Basic statistics
- ❑Corrections and calibrations
- ❑Unfolding techniques
- ❑Feature discovery
- ❑Parameter fitting
- ❑Random #’s and Monte Carlo generators
- ❑Numerical methods
- ❑Geometry and alignment

Experiment specific

- ❑frameworks
- ❑algorithms
- ❑data catalogs
- ❑computing resources
- ❑Documentation sources
- ❑Code repositories

Items in brown were covered in our tutorial

DUNE Trainings

- We do 2-3 training sessions in general infrastructure per year associated with collaboration meetings
  - Individual physics groups also run tutorials on specialized topics
- The main tutorials are 2-3 ½ days
- The goal is to cover the very basics needed to log in, find code, build it and submit simple batch jobs
- Recently have been online but mainly synchronous
  - Around 50 people/session
  - Use the HSF/Carpentries format
  - Record and store zoom sessions and post edited version to the course
- People can and do go back and look at the materials if they missed the session or need a refresher.
- https://indico.fnal.gov/event/54191/ and https://dune.github.io/computing-training-basics

Some demographic info from the intake survey May 2022

Where are you career wise?

35 responses

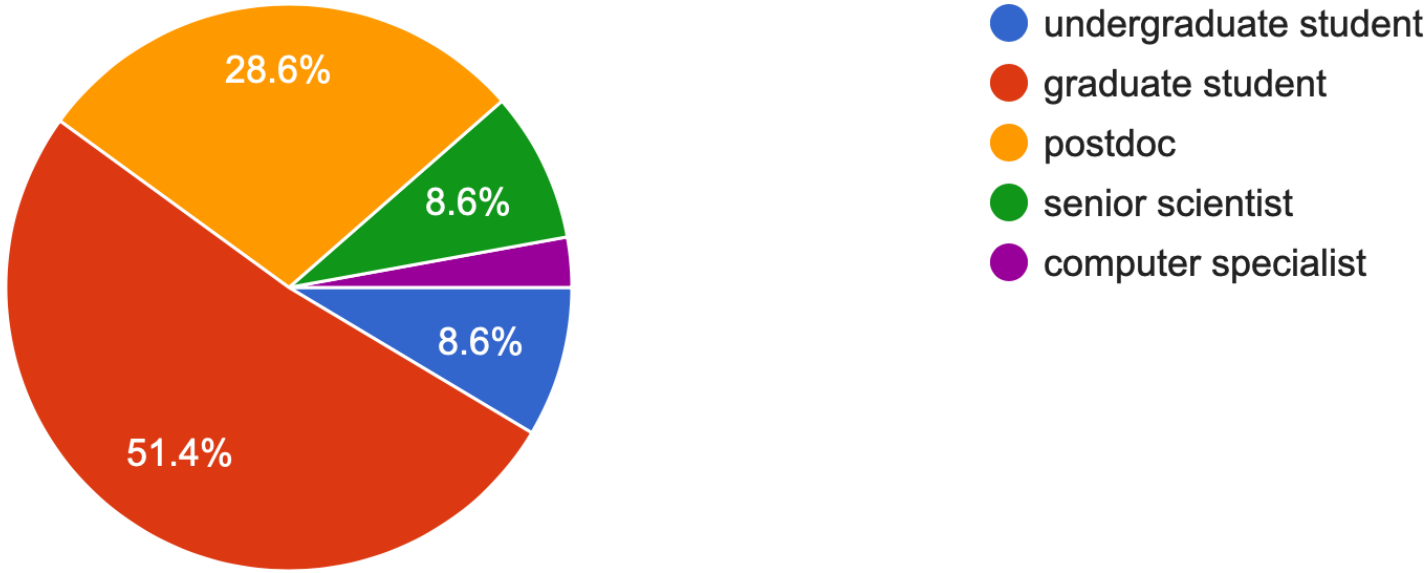

What are you working on

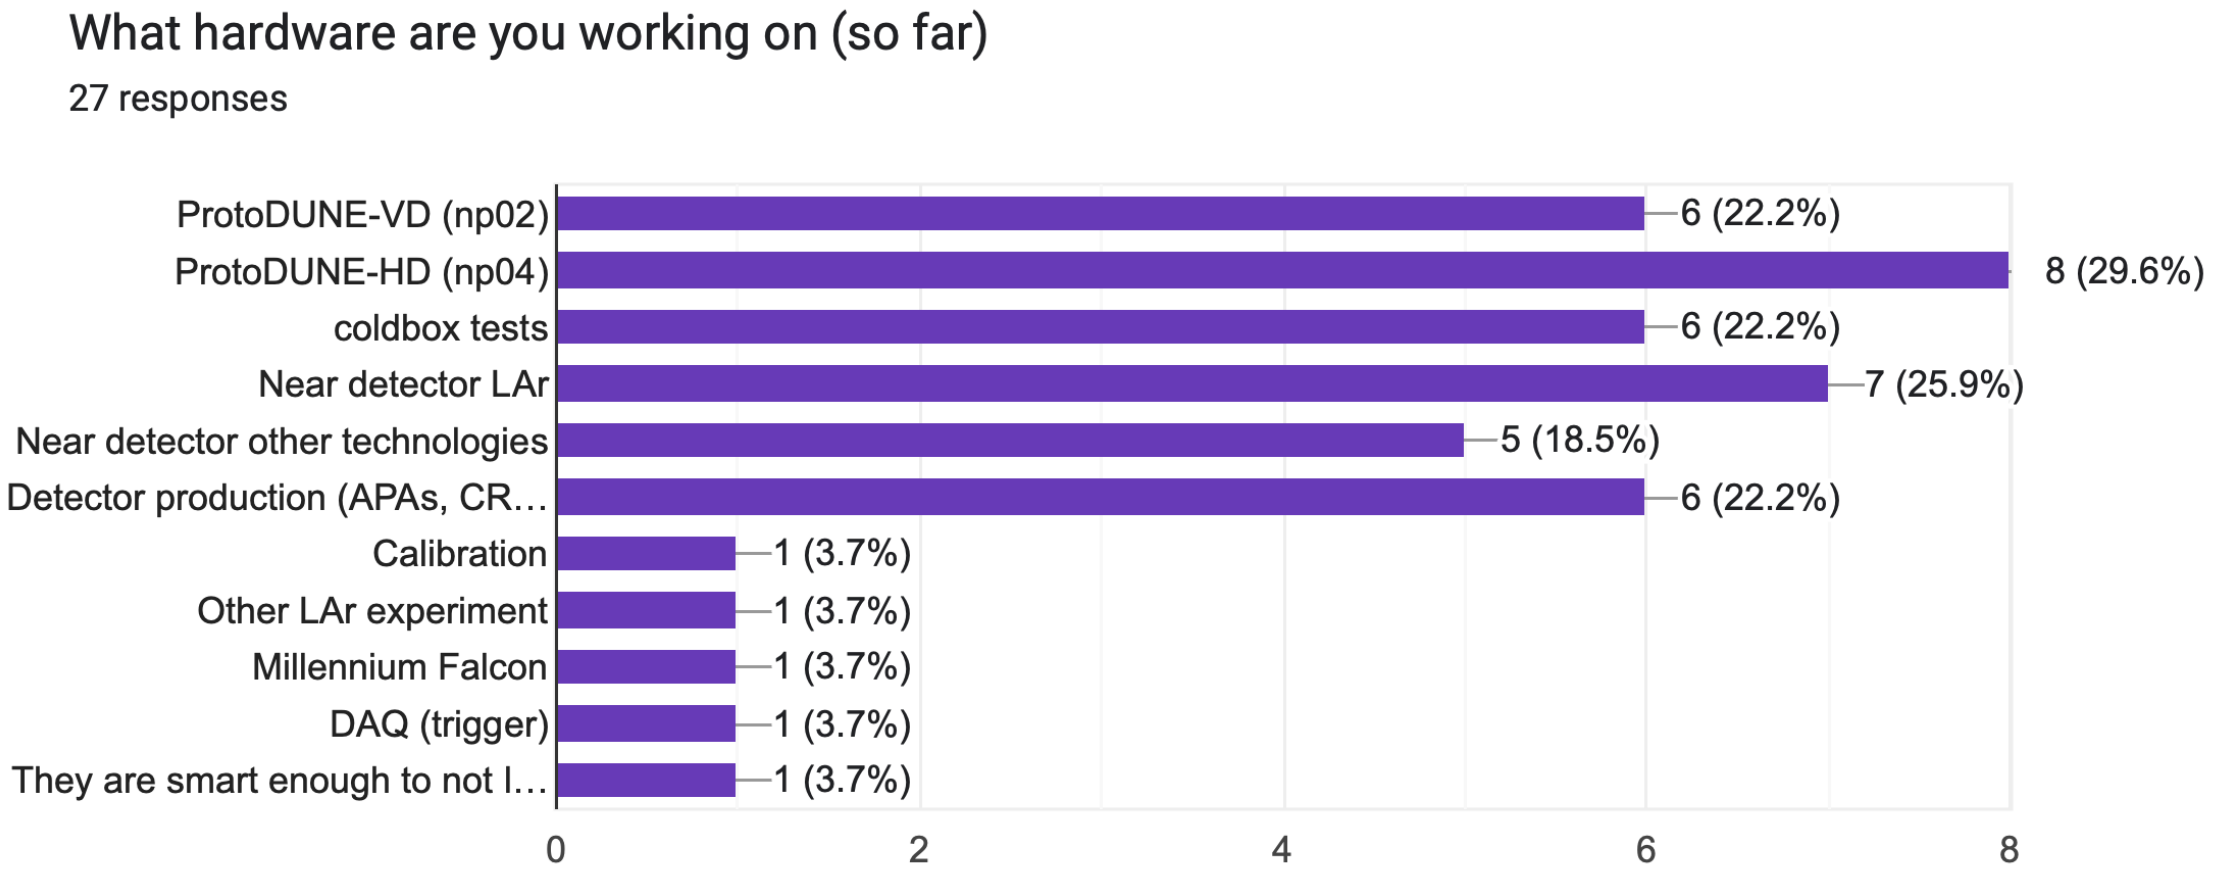

What physics groups are you working in?

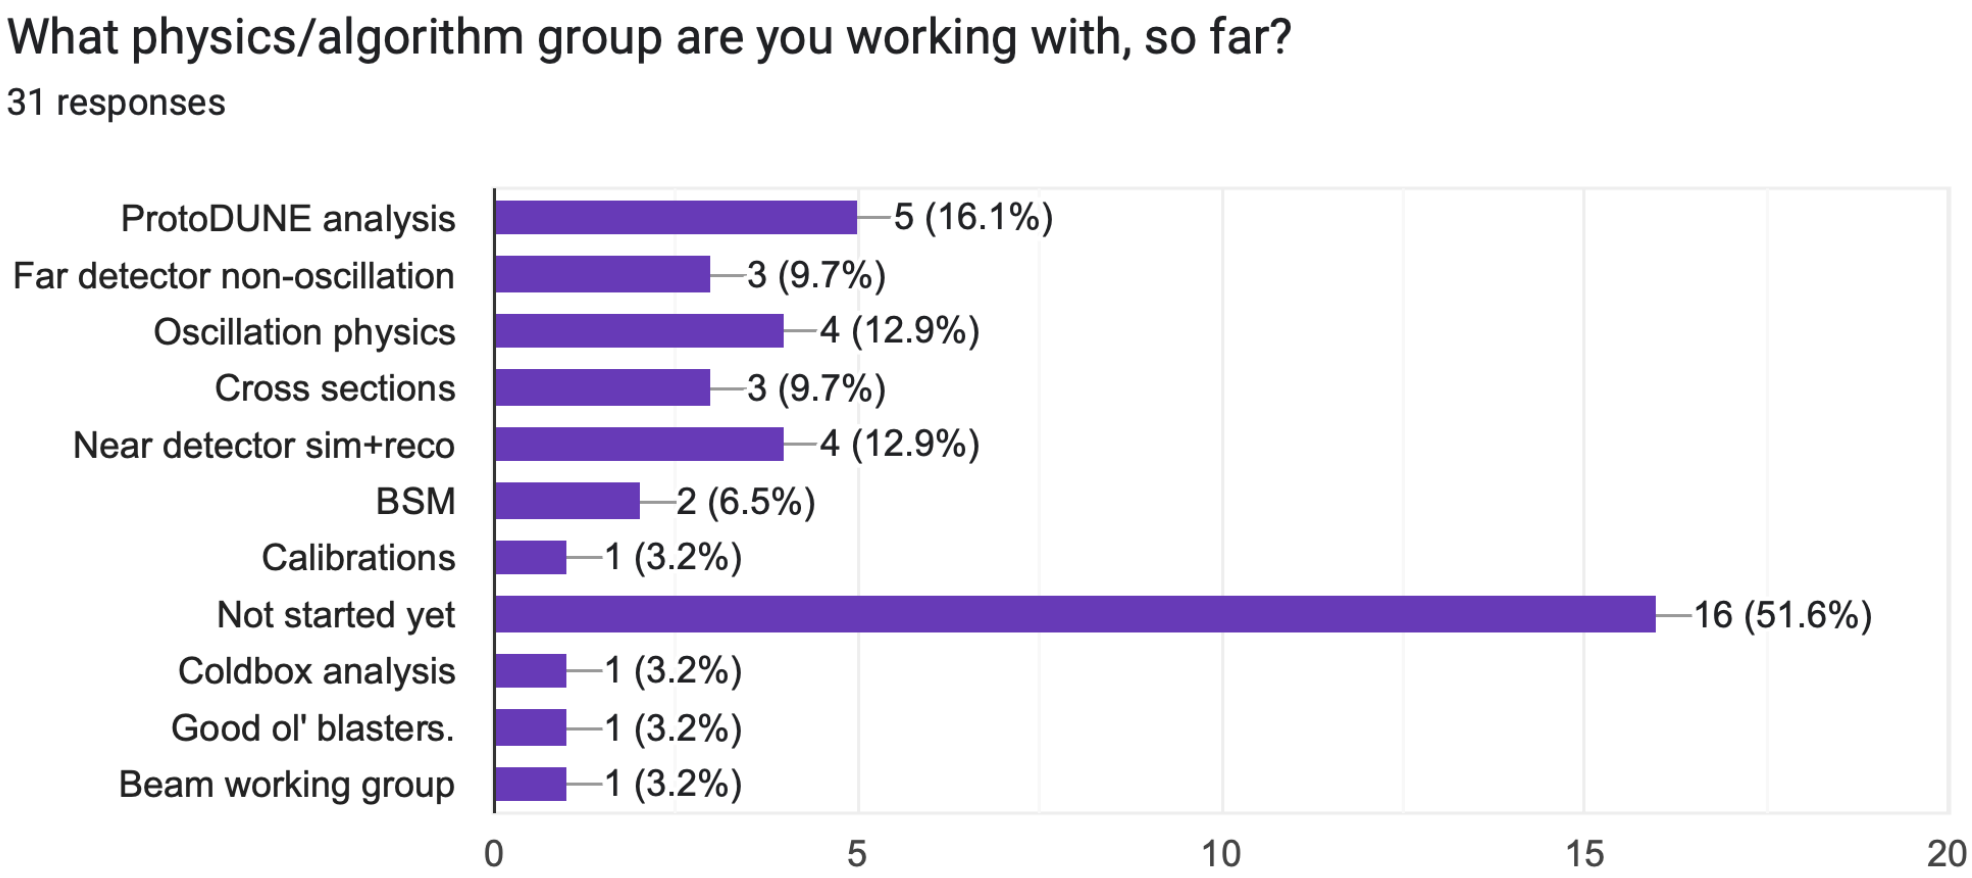

What are you already familiar with?

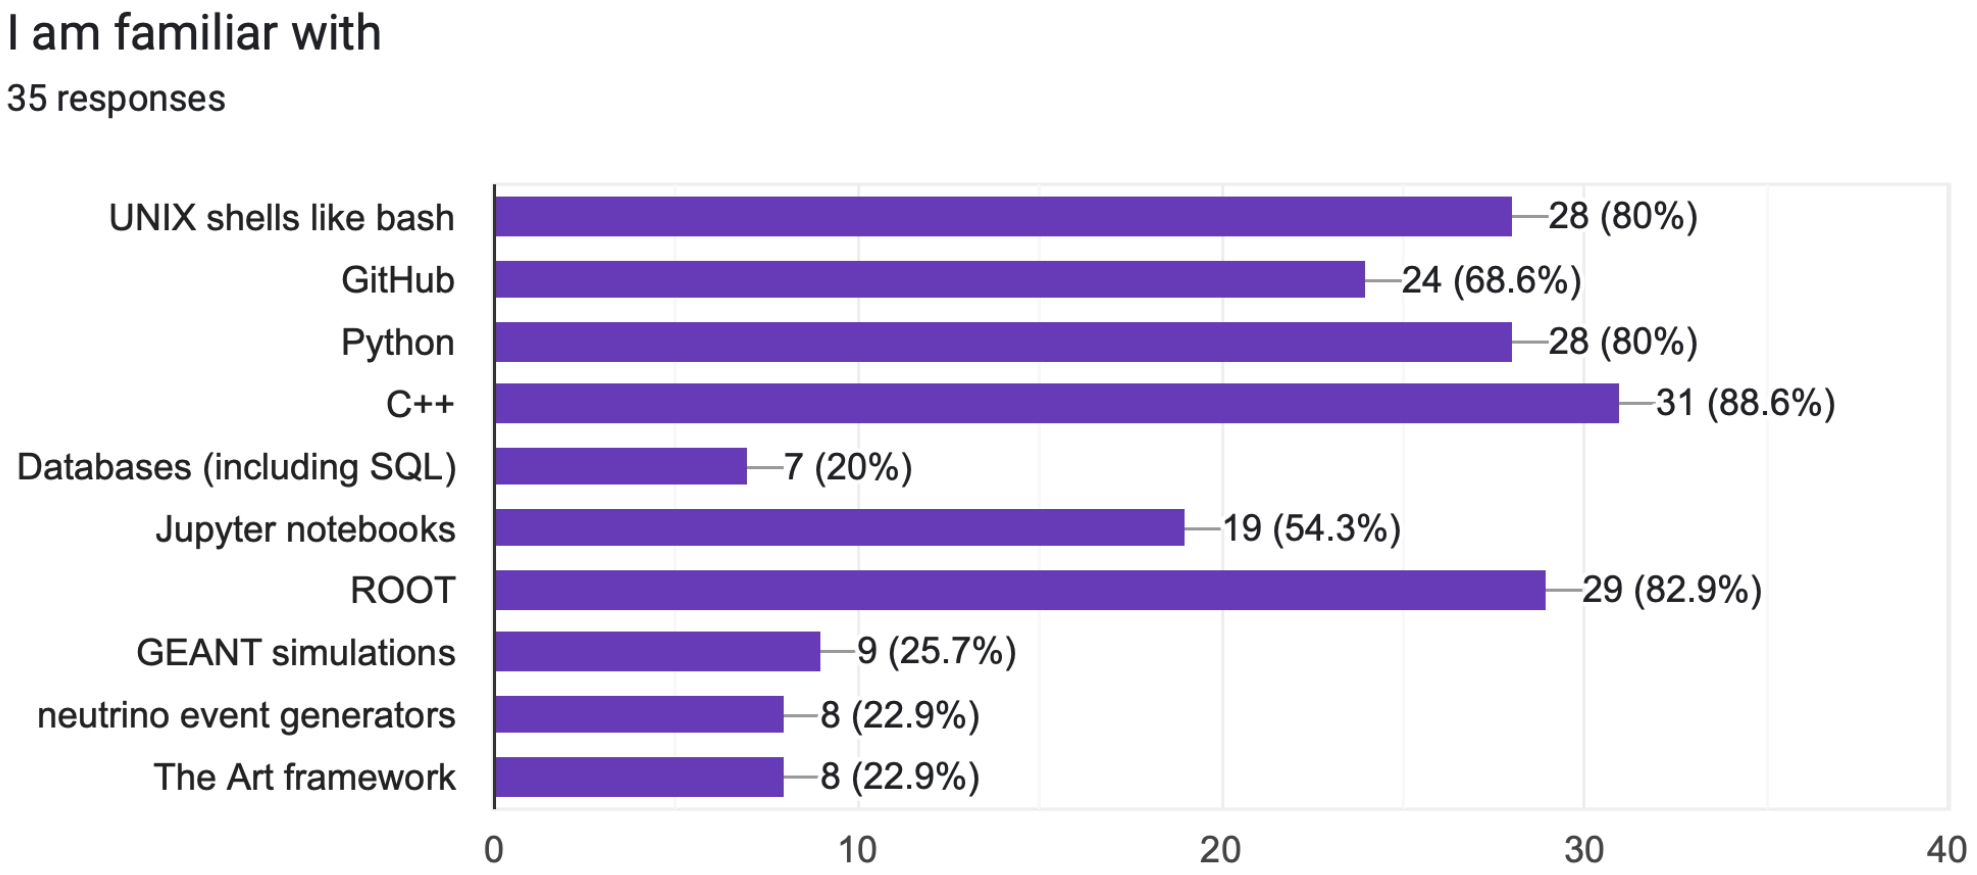

Any experience with large scale batch jobs?

How familiar are you with running large batch jobs

35 responses

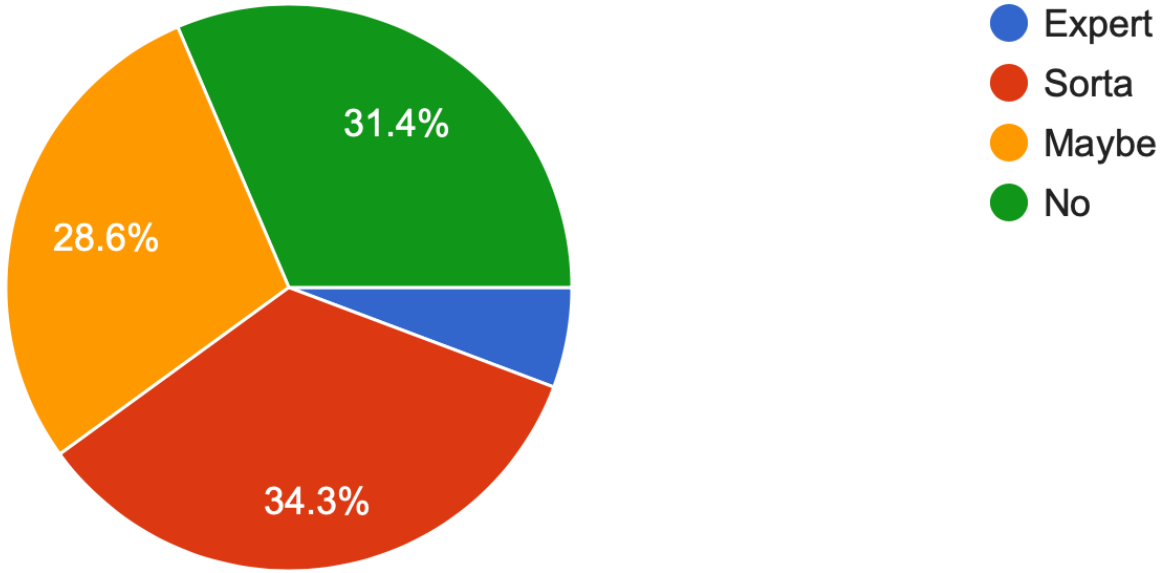

What software packages do you use?

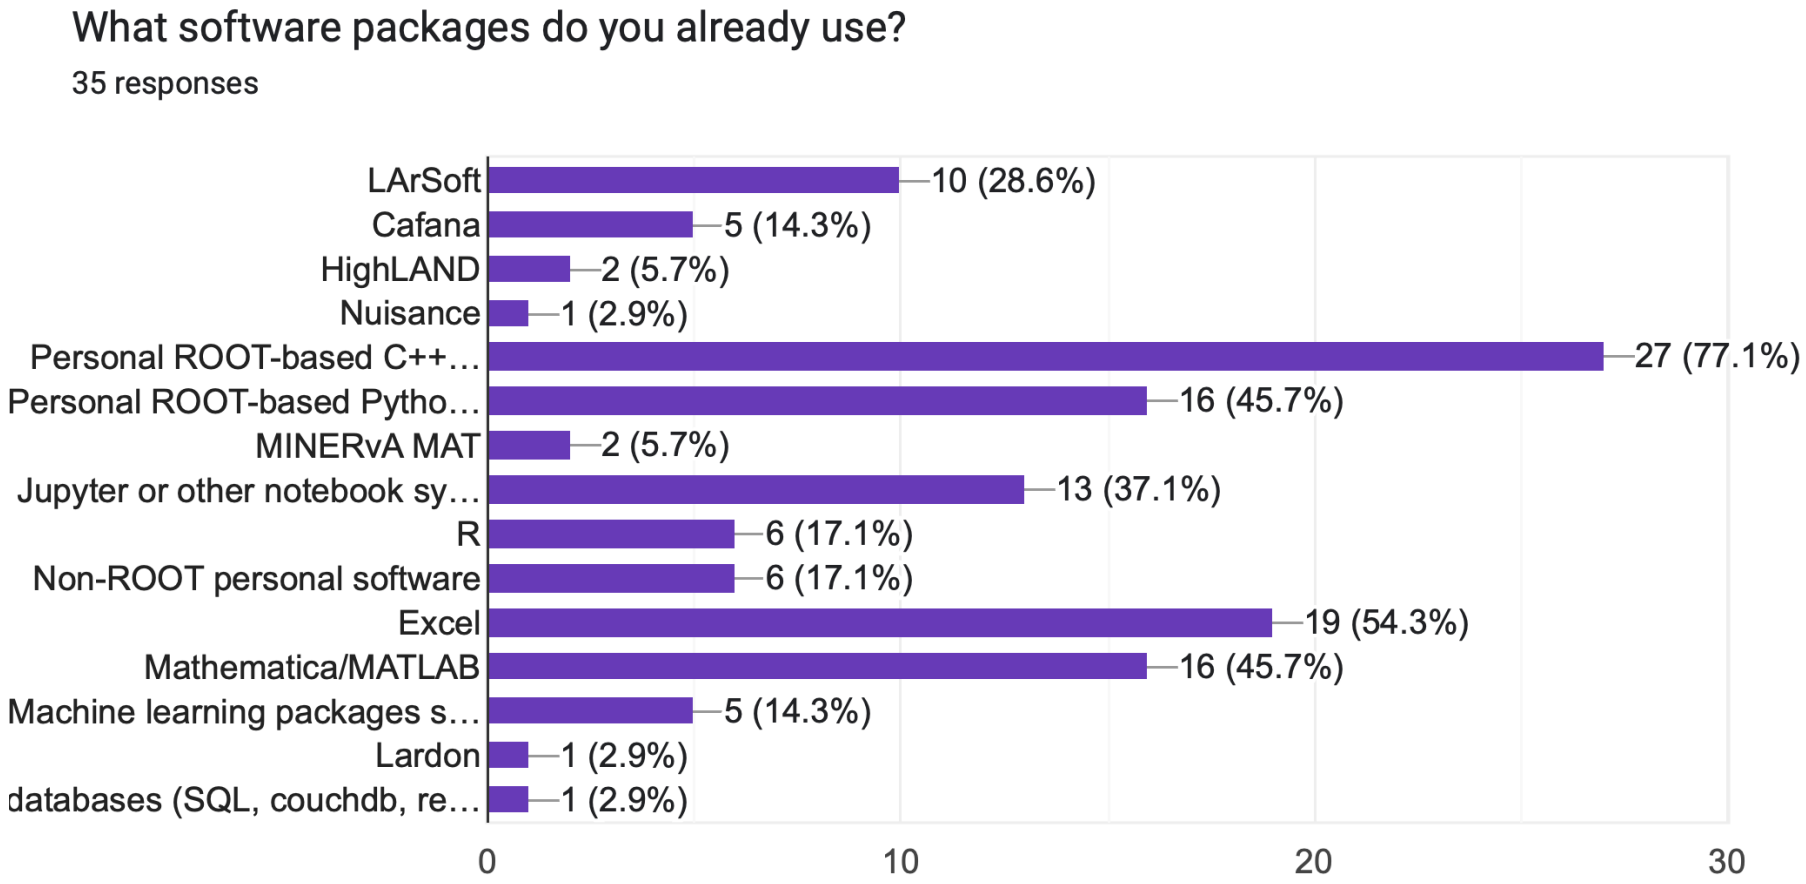

Where do you do your computing?

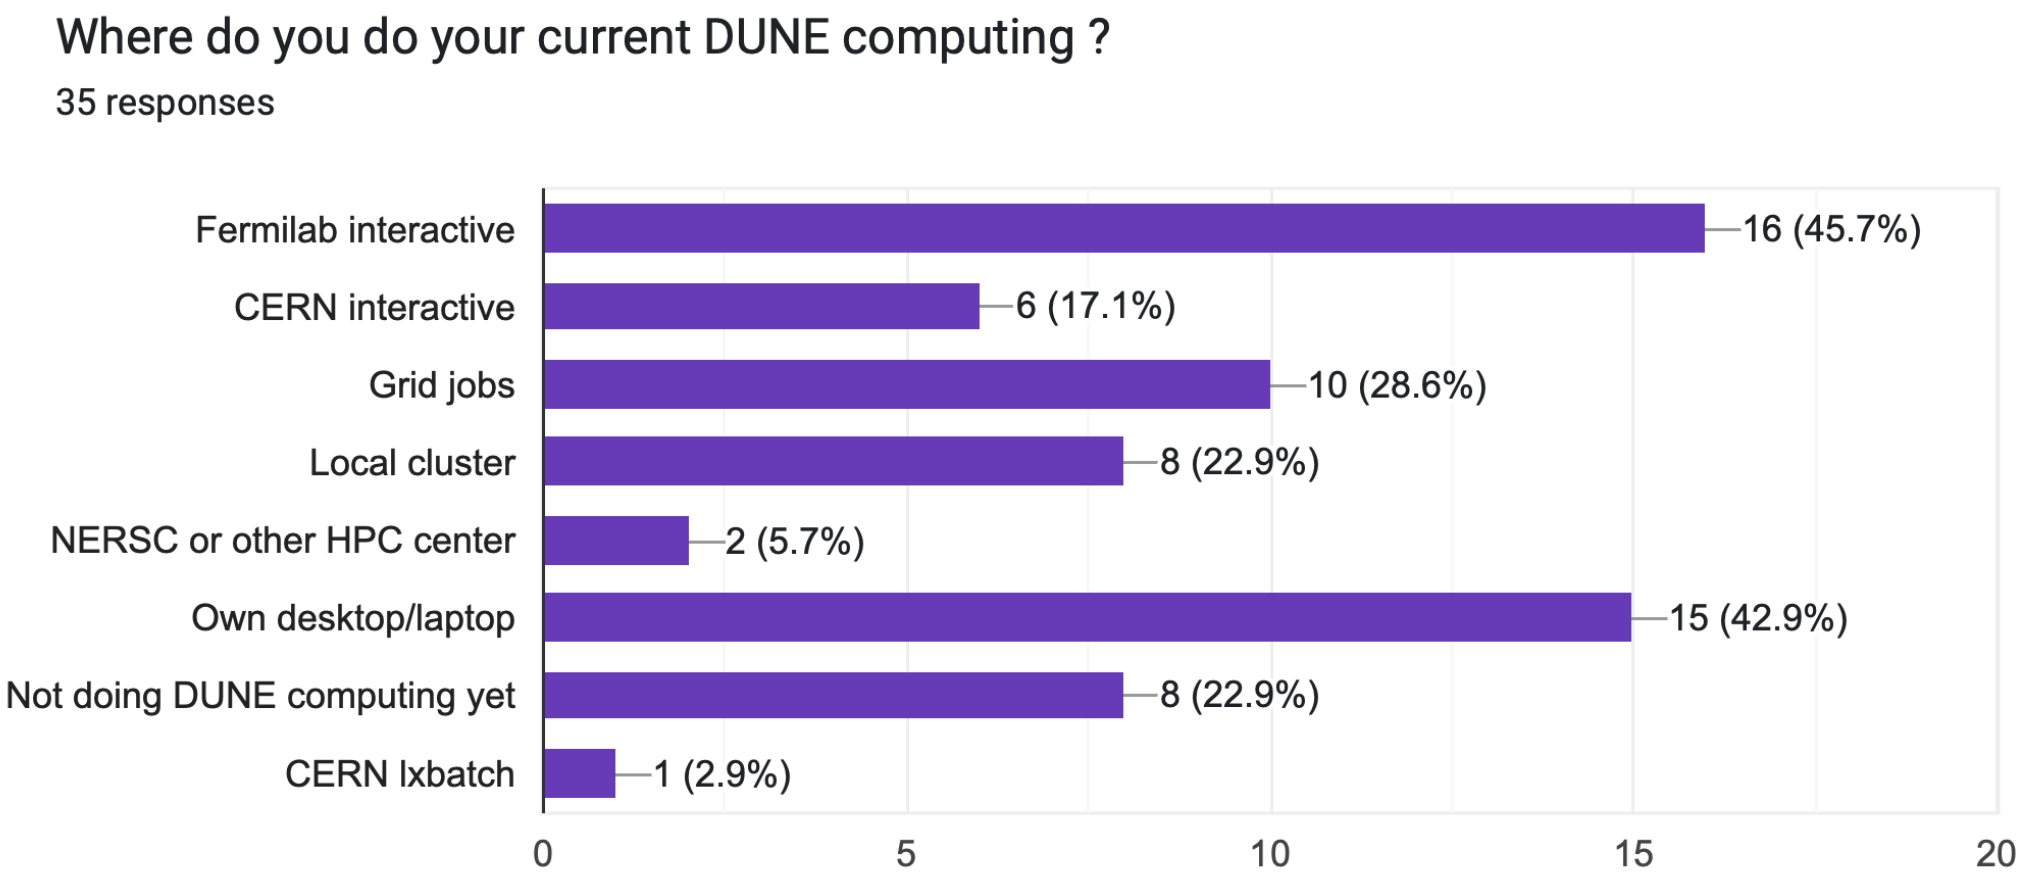

How do you back up your code?

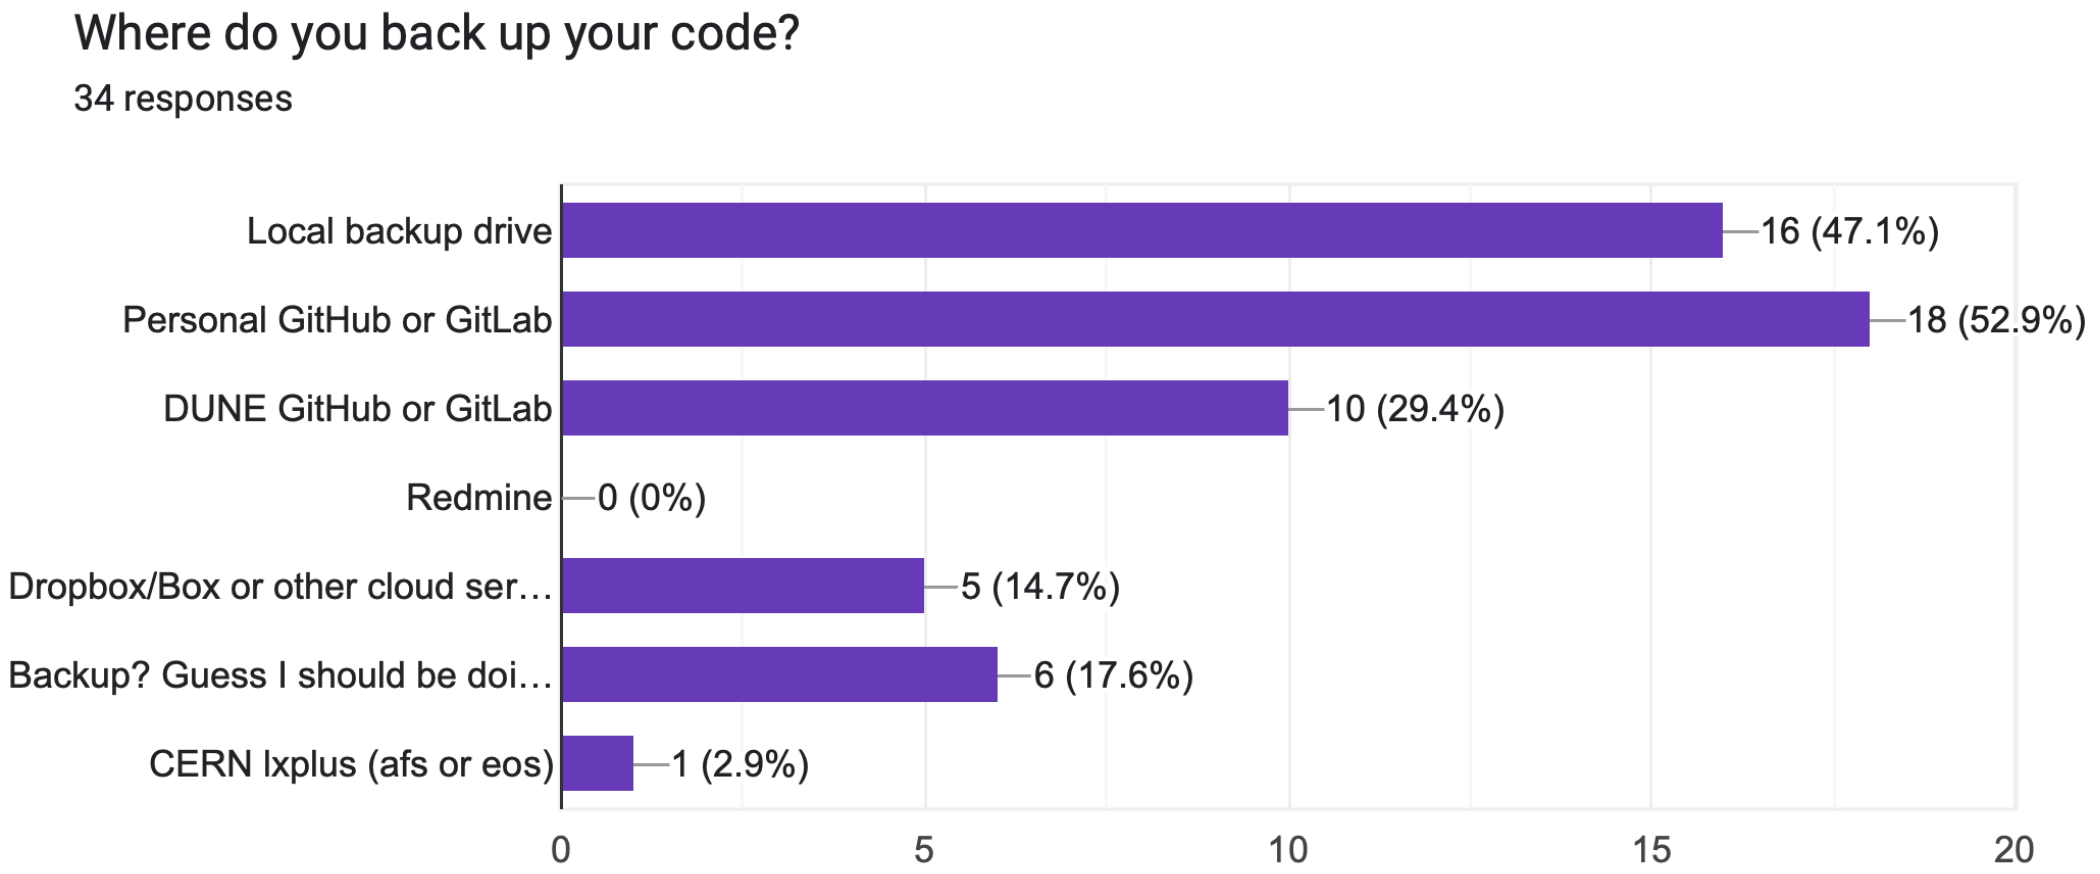

May 22 training schedule

Schedule by Day

Day 1

08:00	Welcome
08:10	DUNE Basics: Storage space and data management
09:45	Coffee
10:15	DUNE Basics: Intro to LArSoft
12:00	Lunch
12:30	DUNE Basics: Grid job submissions + common errors
13:40	Code-makover
14:50	Closing

Day 2

08:00	Welcome
08:10	Code-makover
09:30	Coffee
09:45	Expert in the room: LARSoft: How to Modify a Module
12:00	Lunch
12:30	Code-makeover
13:30	Expert in the room: Grid job submission
14:50	Closing

This lesson is being piloted (Beta version)

Storage Spaces

Overview

Teaching: 45 min
Exercises: 0 min

Questions

- What are the types and roles of DUNE's data volumes?
- What are the commands and tools to handle data?

Objectives

- Understanding the data volumes and their properties
- Displaying volume information (total size, available size, mount point, device location)
- Differentiating the commands to handle data between grid accessible and interactive volumes

This lesson is being piloted (Beta version)

Session Video

The session was captured for your asynchronous review.

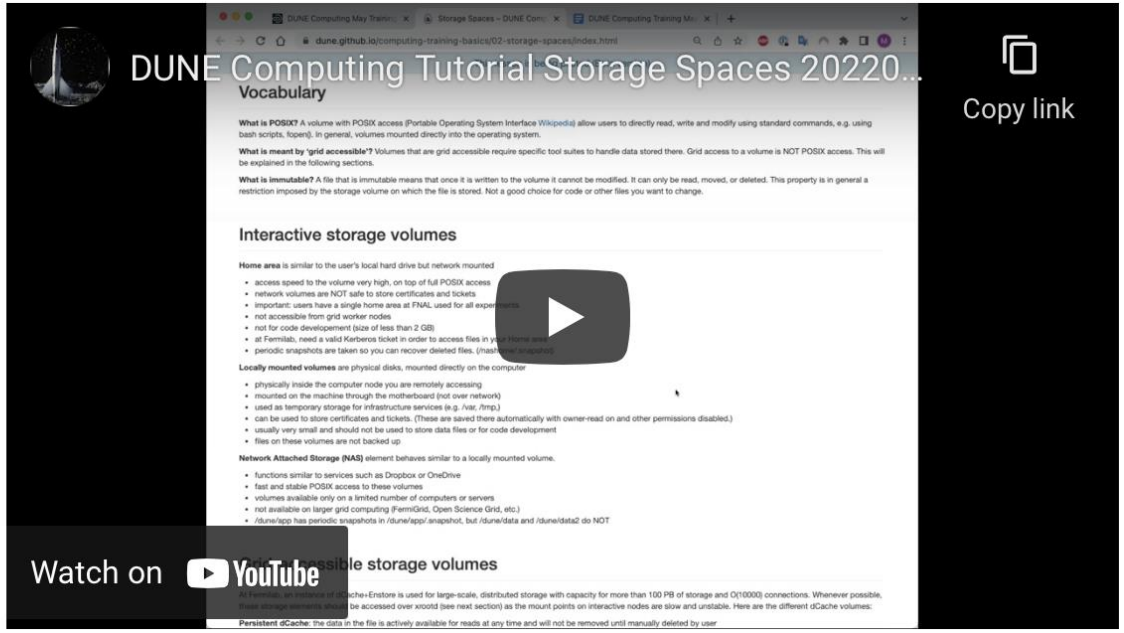

Live Notes

live Notes

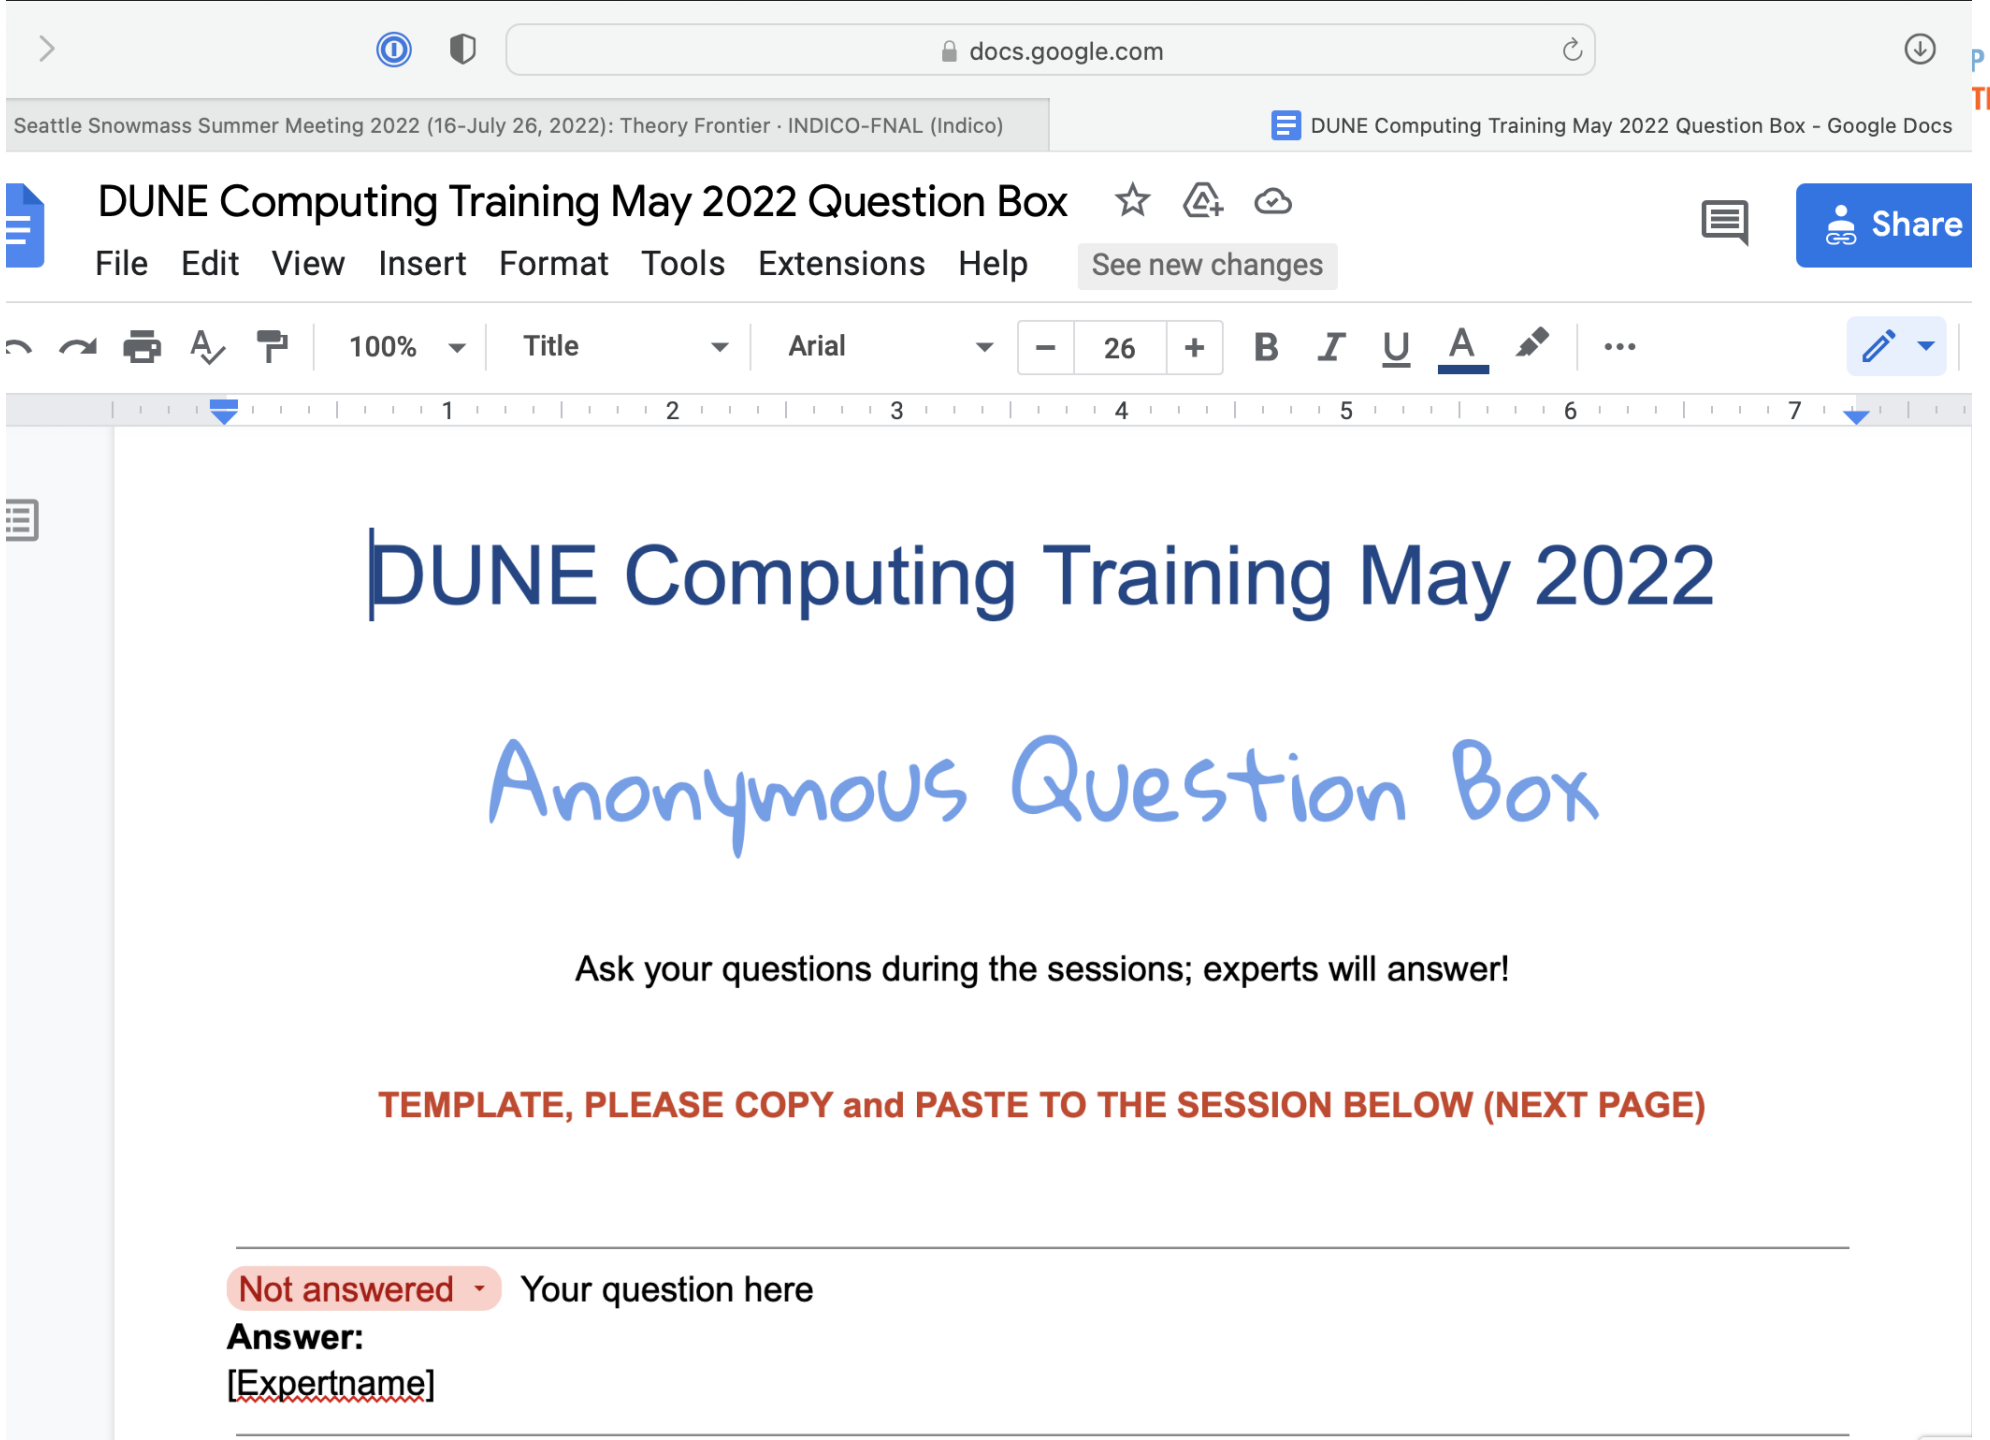

docs.google.com

Seattle Snowmass Summer Meeting 2022 (16-July 26, 2022): Theory Frontier · INDICO-FNAL (Indico)

DUNE Computing Training May 2022 Question Box - Google Docs

DUNE Computing Training May 2022 Question Box

File Edit View Insert Format Tools Extensions Help See new changes

Share

100% Title Arial 26

DUNE Computing Training May 2022

Anonymous Question Box

Ask your questions during the sessions; experts will answer!

TEMPLATE, PLEASE COPY and PASTE TO THE SESSION BELOW (NEXT PAGE)

Not answered Your question here
Answer:
[Expertname]

Vocabulary

What is POSIX? A volume with POSIX access (Portable Operating System Interface Wikipedia) allow users to directly read, write and modify using standard commands, e.g. using bash scripts, fopen(). In general, volumes mounted directly into the operating system.

What is meant by 'grid accessible'? Volumes that are grid accessible require specific tool suites to handle data stored there. Grid access to a volume is NOT POSIX access. This will be explained in the following sections.

What is immutable? A file that is immutable means that once it is written to the volume it cannot be modified. It can only be read, moved, or deleted. This property is in general a restriction imposed by the storage volume on which the file is stored. Not a good choice for code or other files you want to change.

Interactive storage volumes

Home area is similar to the user's local hard drive but network mounted

- access speed to the volume very high, on top of full POSIX access
- network volumes are NOT safe to store certificates and tickets
- important: users have a single home area at FNAL used for all experiments
- not accessible from grid worker nodes
- not for code developement (size of less than 2 GB)
- at Fermilab, need a valid Kerberos ticket in order to access files in your Home area
- periodic snapshots are taken so you can recover deleted files. (/nashome/.snapshot)

Locally mounted volumes are physical disks, mounted directly on the computer

- physically inside the computer node you are remotely accessing
- mounted on the machine through the motherboard (not over network)
- used as temporary storage for infrastructure services (e.g. /var, /tmp,)
- can be used to store certificates and tickets. (These are saved there automatically with owner-read on and other permissions disabled.)
- usually very small and should not be used to store data files or for code development

Full documentation: Understanding Storage Volumes

In the following table, <exp> stands for the experiment (uboone, nova, dune, etc…)

	Quota/Space	Retention Policy	Tape Backed?	Retention Lifetime on disk	Use for	Path	Grid Accessible
Persistent dCache	No/~100 TB/exp	Managed by Experiment	No	Until manually deleted	immutable files w/ long lifetime	/pnfs/<exp>/persistent	Yes
Scratch dCache	No/no limit	LRU eviction - least recently used file deleted	No	Varies, ~30 days (*NOT* guaranteed)	immutable files w/ short lifetime	/pnfs/<exp>/scratch	Yes
Resilient dCache	No/no limit	Periodic eviction if file not accessed	No	Approx 30 days (your experiment may have an active clean up policy)	input tarballs with custom code for grid jobs (do NOT use for grid job outputs)	/pnfs/<exp>/resilient	Yes
Tape backed	dCache No/O(10) PB	LRU eviction (from disk)	Yes	Approx 30 days	Long-term archive	/pnfs/dune/…	Yes
NAS Data	Yes (~1 TB)/ 32+30 TB total	Managed by Experiment	No	Until manually deleted	Storing final analysis samples	/dune/data	No
NAS App	Yes (~100 GB)/ ~15 TB total	Managed by Experiment	No	Until manually deleted	Storing and compiling software	/dune/app	No
Home Area (NFS mount)	Yes (~10 GB)	Centrally Managed by CCD	No	Until manually deleted	Storing global environment scripts (All FNAL Exp)	/nashome/<letter>/<uid>	No

DUNE Computing Training May 2022

Quiz time!

Storage spaces and data management

1. Which volumes are directly accessible (POSIX) from grid worker nodes?
   A. /dune/data
   B. DUNE CVMFS repository
   C. /pnfs/dune/scratch
   D. /pnfs/dune/persistent
   E. None of the above

2. Which data volume is the best location for the output of an analysis-user grid job?
   A. dCache scratch (/pnfs/dune/scratch/users/${USER}/)
   B. dCache persistent (/pnfs/dune/persistent/users/${USER}/)
   C. Enstore tape (/pnfs/dune/tape_backed/users/${USER}/)
   D. user's home area (`~${USER}`)
   E. NFS data volume (/dune/data or /dune/app)

Exit survey info (from 2021)

Your thoughts and suggestions regarding the communication

- I really enjoyed the way it was organised. I never felt off even for a small time even though I was quite busy with some other stuff.
- I think the setup part should have been released at least a week earlier. As someone who is completely new to programming it was extremely overwhelming.
- Appreciated the remainder! Slack would be another (additional ?) great option
- Vocabulary was difficult to follow and sometimes not explained
- Communication was perfect. Updates were given before time so that we could practice beforehand.
- It was fantastic
- I think it worked very well.
- Announcements, reminders, emails, posts on various channels, etc are good.
- I thought that the open document was perfect for posting questions and interacting with other attendees

More exit survey

Your thoughts on the quizzes

- Really Really enjoying!!!
- I really liked them! it was nice when i got to a breakout room where people were actually coworking
- On one quiz I was wrong on almost every question. So I think they give me a baseline of where to get started on my own time.
- Questions were fine, but the wrong answer selections were too obvious so even if you didn't attend the lecture, you could still answer the questions correctly.
- Good idea. It was also very good, interesting and helpful to discuss with some participants in the breakout room.
- More quizzes in the future!
- The quizzes help to learn the presented material in a different and memorable way.

General takeaways

- Getting new people started is hard
  - Accounts take time to set up and are a complex multi-step process
  - People sometimes put off the requirement to get their accounts set up ahead of time and suffered. We've gone to zero-tolerance to avoid holding everything up.
- New people are often not well connected to the collaboration by their home institutions.
  - New people are often asked to do X by their advisors/bosses, when X already has a large group working on it.
  - This course may be their first chance to form connections with others.
- As seen in the survey, most people do come in with a decent grounding in basic computing, it's the complicated structures needed for a large collaboration that cause them woe.
- Participant feedback is that having explanations of why things are the way they are is very useful.
- Participant feedback is also very positive on having the pages and videos to return to afterwards.
- Participants did not always like the fast pace and torrent of information. We should probably cut back a bit.
- Some wanted us to cover machine learning. Probably out of scope for the first introduction.
- Our computing training currently only covers the basics. We need a stronger infrastructure for guiding people to advanced training in experiment of physics specific topics.

Thanks

- Thanks to the HSF Training group
- Thanks to the Fermilab FIFE group and the CERN IT people who provide the tools we are teaching
- Thanks to the coordinators: Claire David and Dave DeMuth,
- and the instructors: Steve Timm, Mike Kirby, Ken Herner and Tom Junk
- And the helpers who answered questions on the fly