



# NOvA HEPCloud overview

---

v hewes  
HEPCloud meeting  
25th September 2024

# CPU node allocation

---

- NOvA would like to request **75k** CPU node hours on Perlmutter.
- Primary use case is **statistical corrections** for NOvA's 3F and 3+1F **oscillation analyses**.
  - NOvA's frequentist oscillation results all require computationally intensive Feldman-Cousins corrections.
  - Highly beneficial to use HPC resources for these runs, as they are very resource-intensive and often time-sensitive.
  - Most recent 3+1F analysis utilized ~22M CPU core hours across Cori-KNL and Perlmutter nodes.
- NOvA is also set up to run MC production at NERSC, which we can use to leverage our allocation more flexibly around FC corrections.

# GPU node allocation

---

- NOvA would like to request **5k** GPU node hours on Perlmutter.
- This allocation would primarily support **ML training and inference**.
  - NOvA switched to EAF and institutional GPU clusters immediately following Wilson's retirement, but would like to utilize NERSC more actively.
  - As effort ramps up on NOvA's **MC Production 6**, neural network training needs will increase accordingly.
- Switching to Spack allows for **native unification of sim/reco and ML software stacks**.
- Development of **GPU-optimized multi-sample fitting** techniques.
- Will need a system in place to set up NOvA reconstruction experts with NERSC accounts and project access.

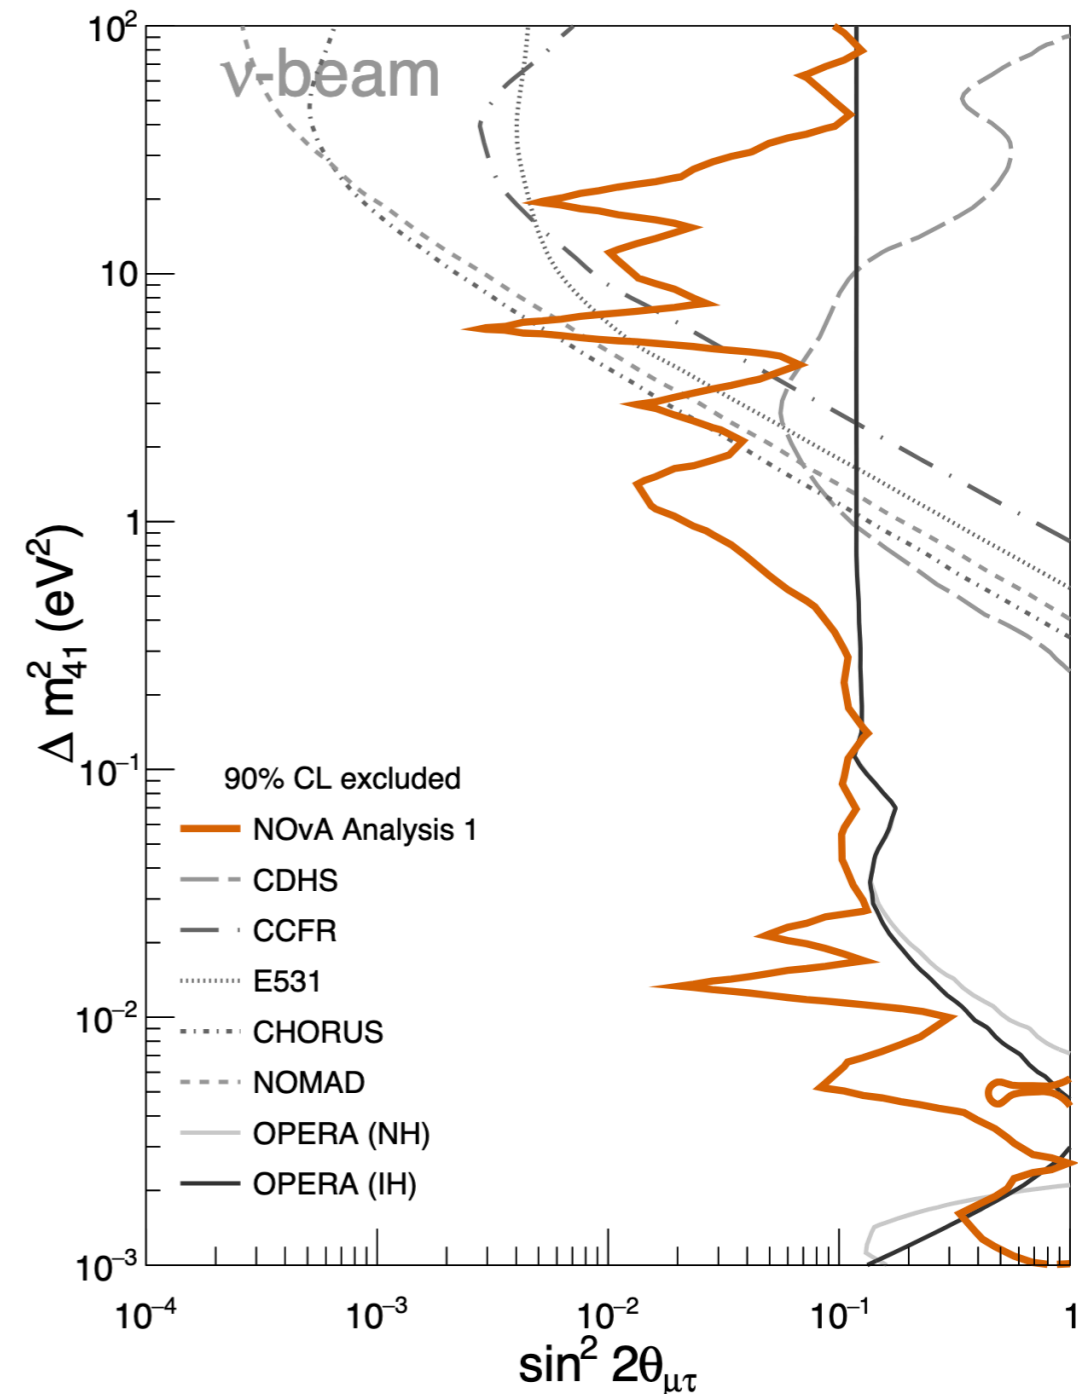
# Disk allocation

---

- NOvA would ideally like to request **15TB** disk space at NERSC, to support HDF5 workflows and input files for production.
- SciDAC-developed workflow to generate HDF5-based file format.
- Each HDF5 MC dataset requires ~1TB disk space.
- File concatenation procedure developed to merge files into one large HDF5 file for highly parallel processing.
- Preferred workflow:
  - Generate dataset of small HDF5 files at FNAL.
  - Transfer to Perlmutter.
  - Utilize HPC CPU nodes to merge HDF5 datasets into single file.
  - Leverage merged HDF5 dataset files for ML training.

# Publication list

- **Dual-Baseline Search for Active-to-Sterile Neutrino Oscillations in NOvA**
  - [arXiv:2409.04553](https://arxiv.org/abs/2409.04553), submitted to PRL (2024)
- **Measurement of the double-differential muon-neutrino charged-current inclusive cross section in the NOvA near detector**
  - *Phys. Rev. D* **107**, 052011 (2023)
- **Improved measurement of neutrino oscillation parameters by the NOvA experiment**
  - *Phys. Rev. D* **106**, 032004 (2022)
- **Search for Active-Sterile Antineutrino Mixing Using Neutral-Current Interactions with the NOvA Experiment**
  - *Phys. Rev. Lett.* **127**, 2018091 (2021)
- **First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA**
  - *Phys. Rev. Lett.* **123**, 151803 (2019)



# NOvA HEPCloud summary

---

<b>Allocation</b>	<b>Amount</b>
<b>CPU node hours</b>	<b>75k</b>
<b>GPU node hours</b>	<b>5k</b>
<b>Disk space</b>	<b>15TB</b>