SDSC
SAN DIEGO SUPERCOMPUTER CENTER

IceCube

Using Distributed Computing and Storage on PNRP

Mahidhar Tatineni
Director of User Services, SDSC

4th GLOBAL RESEARCH PLATFORM WORKSHOP
Oct 9-10, 2023

Acknowledgements:

Frank Würthwein Director, SDSC
IceCube work to utilize PNRP:
David Schultz, Igor Sfiligoi, Benedikt Riedel, Fabio Andrijauskas, Derek Weitzel, and Frank Würthwein
DOI: https://doi.org/10.48550/arXiv.2308.07999

Outline

- IceCube scientific motivation
- Compute & Storage requirements
- Overview of PNRP resources
- IceCube modifications to use PNRP compute (GPUs) and storage (Origins)
- Summary

Outline

- IceCube scientific motivation
- Compute & Storage requirements
- Overview of PNRP resources
- IceCube modifications to use PNRP compute (GPUs) and storage (Origins)
- Summary

IceCube

A cubic kilometer of ice at the south pole is instrumented with 5160 optical sensors.

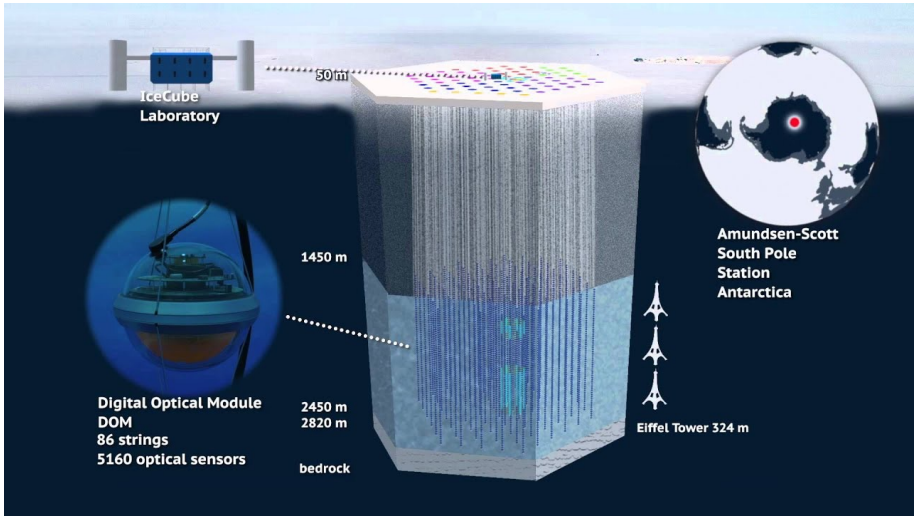

A facility with very diverse science goals

Astrophysics:
- Discovery of astrophysical neutrinos
- First evidence of neutrino point source (TXS)
- Cosmic rays with surface detector

Particle Physics:
- Atmospheric neutrino oscillation
- Neutrino cross sections at TeV scale
- New physics searches at highest energies

Earth Science:
- Glaciology
- Earth tomography

Restrict this talk to high energy Astrophysics

High Energy Astrophysics Science case for IceCube

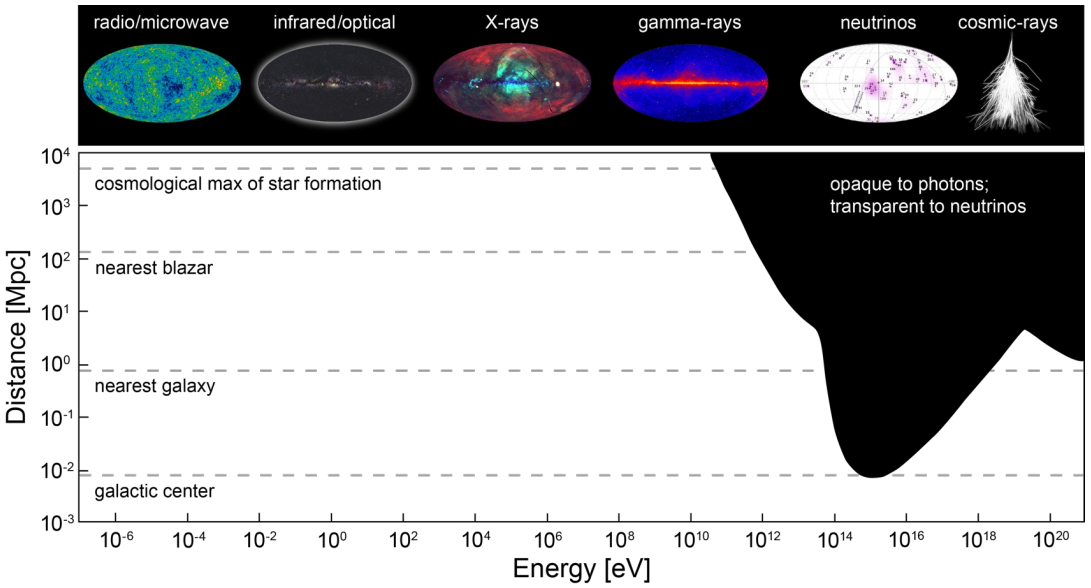

Universe is opaque to light at highest energies and distances.

Only gravitational waves and neutrinos can pinpoint most violent events in universe.

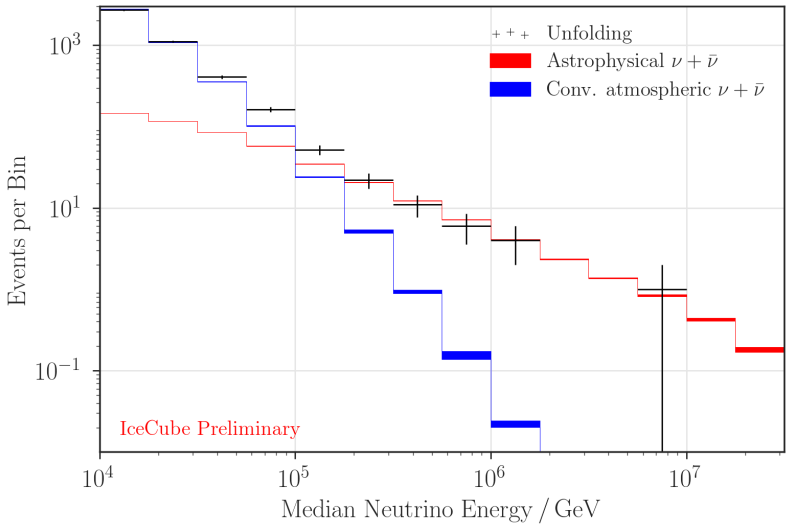

Fortunately, highest energy neutrinos are of cosmic origin.

Effectively "background free" as long as energy is measured correctly.

High energy neutrinos from outside the solar system

First 28 very high energy neutrinos from outside the solar system

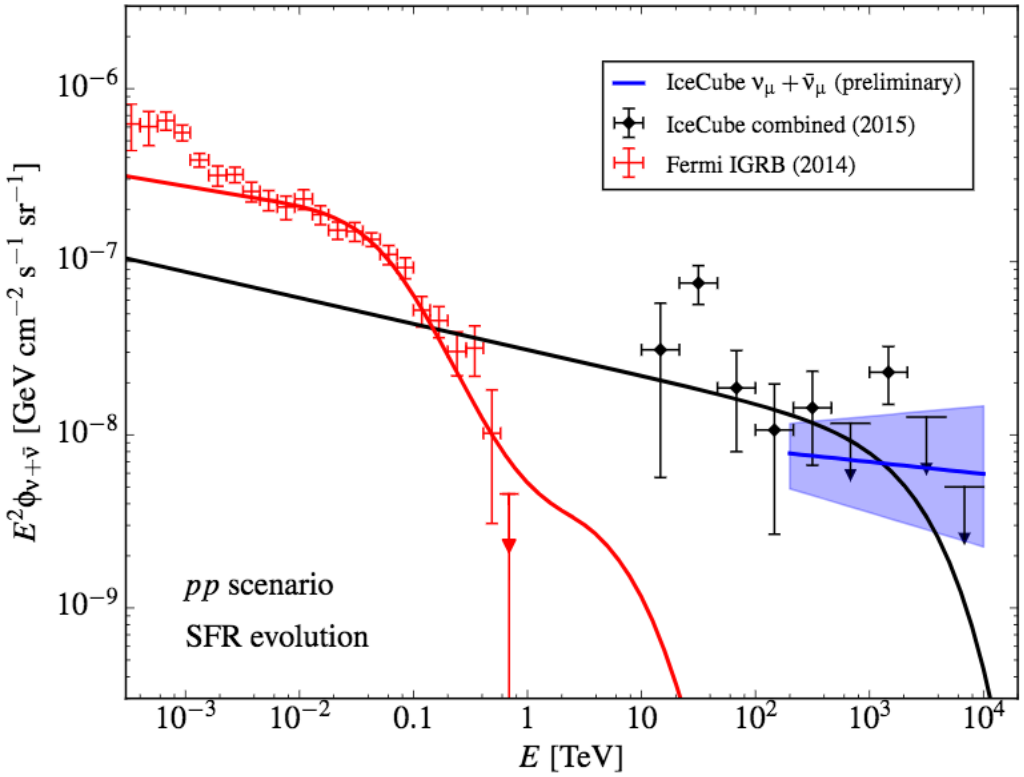

Science 342 (2013). DOI: 10.1126/science.1242856

Red curve is the photon flux spectrum measured with the Fermi satellite.

Black points show the corresponding high energy neutrino flux spectrum measured by IceCube.

This demonstrates both the opaqueness of the universe to high energy photons, and the ability of IceCube to detect neutrinos above the maximum energy we can see light due to this opaqueness.

Understanding the Origin

We now know high energy events happen in the universe. What are they?

The hypothesis:

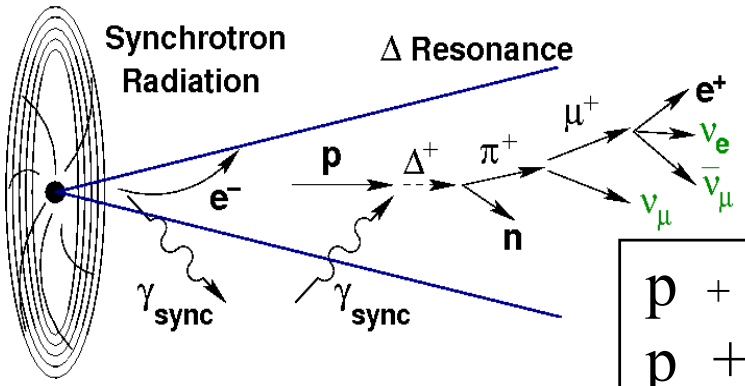

The same cosmic events produce neutrinos and photons

$$p + \gamma \rightarrow \Delta^{+} \rightarrow p + \pi^{0} \rightarrow p + \gamma\gamma$$
$$p + \gamma \rightarrow \Delta^{+} \rightarrow n + \pi^{+} \rightarrow n + \mu + \nu$$

We detect the electrons or muons from neutrino that interact in the ice.
Neutrino interact very weakly => **need a very large array of ice instrumented** to maximize chances that a cosmic neutrino interacts inside the detector.
Need pointing accuracy to point back to origin of neutrino.
Telescopes the world over then try to identify the source in the area IceCube is pointing to given the neutrino.

Multi-messenger Astrophysics

First evidence of an origin

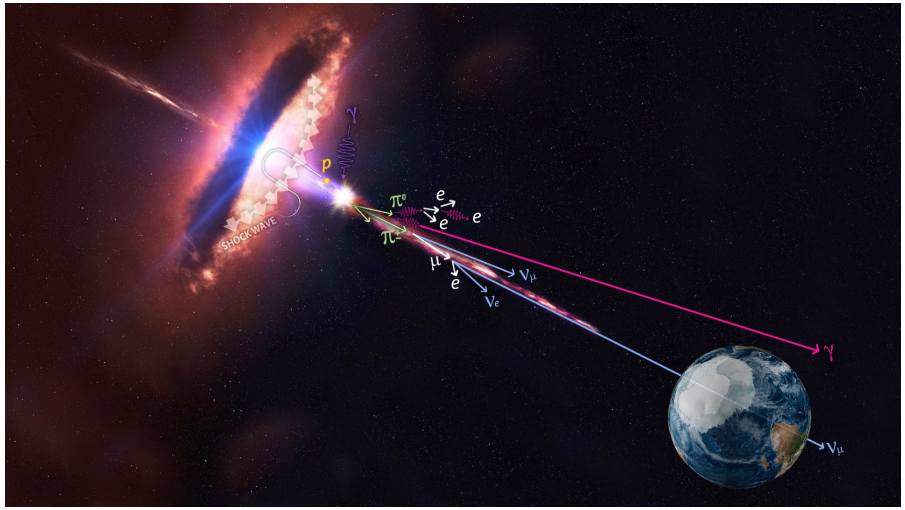

IceCube alerted the astronomy community of the observation of a single high energy neutrino on September 22 2017.

A blazar designated by astronomers as TXS 0506+056 was subsequently identified as most likely source in the direction IceCube was pointing. Multiple telescopes saw light from TXS at the same time IceCube saw the neutrino.

Science 361, 147-151 (2018). DOI:10.1126/science.aat2890

First location of a source of very high energy neutrinos.

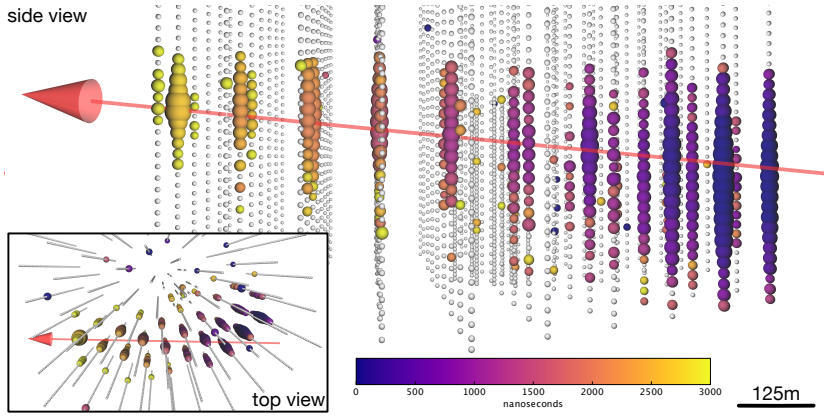

Neutrino produced high energy muon near IceCube. Muon produced light as it traverses IceCube volume. Light is detected by array of phototubes of IceCube.

IceCube's Future Plans

Near term:

add more phototubes to deep core to increase granularity of measurements.

Longer term:

- Extend instrumented volume at smaller granularity.
- Extend even smaller granularity deep core volume.
- Add surface array.

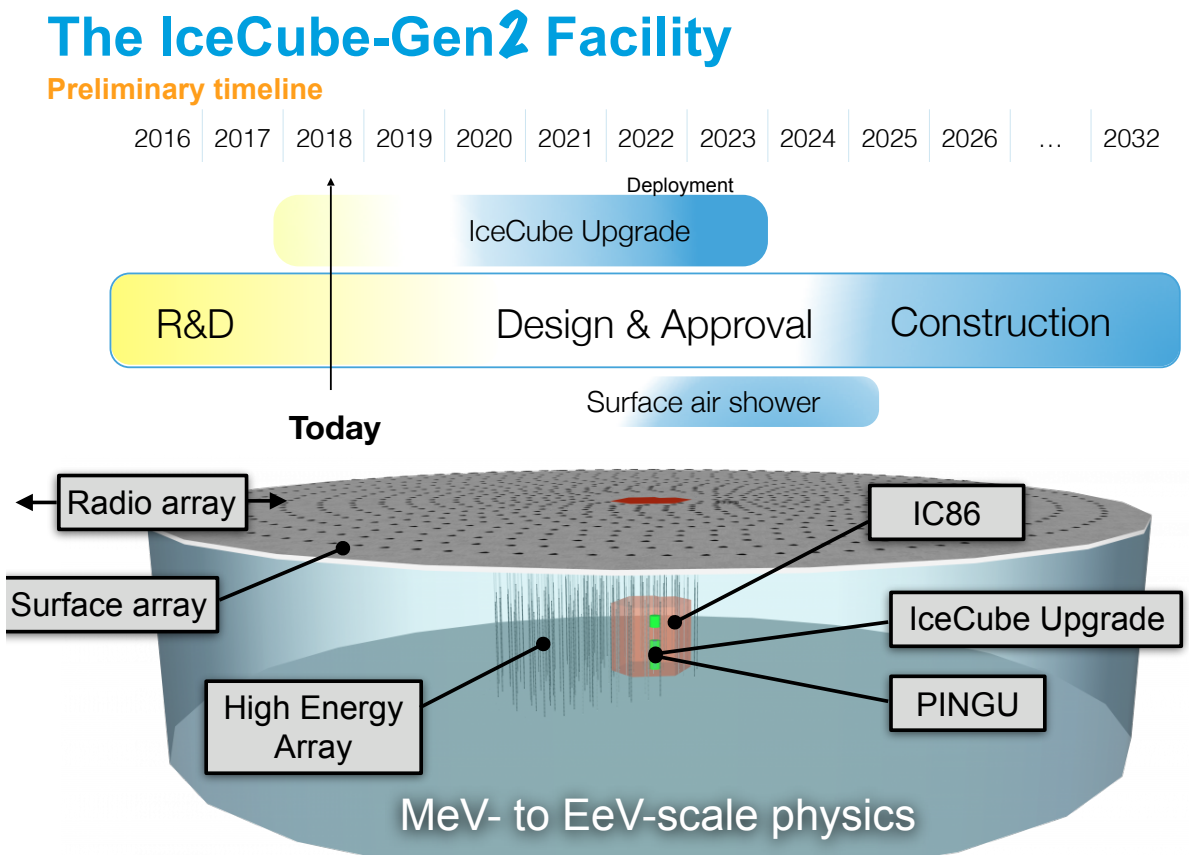

Improve detector for low & high energy neutrinos

Outline

- IceCube scientific motivation
- Compute & Storage requirements
- Overview of PNRP resources
- IceCube modifications to use PNRP compute (GPUs) and storage (Origins)
- Summary

Compute & Storage Requirements

- Understanding detector systematic effects is a continuous process. IceCube's largest systematic effect comes from the optical properties of the ice (where detector is embedded)
- Requires Monte Carlo simulations to quantify potential changes
- Most computationally intensive part of the IceCube simulation workflow is a photon propagation code that performs well on GPUs.
- IceCube utilizes OSG resources extensively, and the PNRP cluster is a big source of the GPU time.
- Results are copied back to central storage servers located at the University of Wisconsin–Madison (UW).
- PNRP storage Origins have also been used at scale (more in upcoming slides)
- Cloud bursting has also been explored. Details in:

  https://doi.org/10.48550/arXiv.2107.03963

IceCube: Global Computing & Data Infrastructure

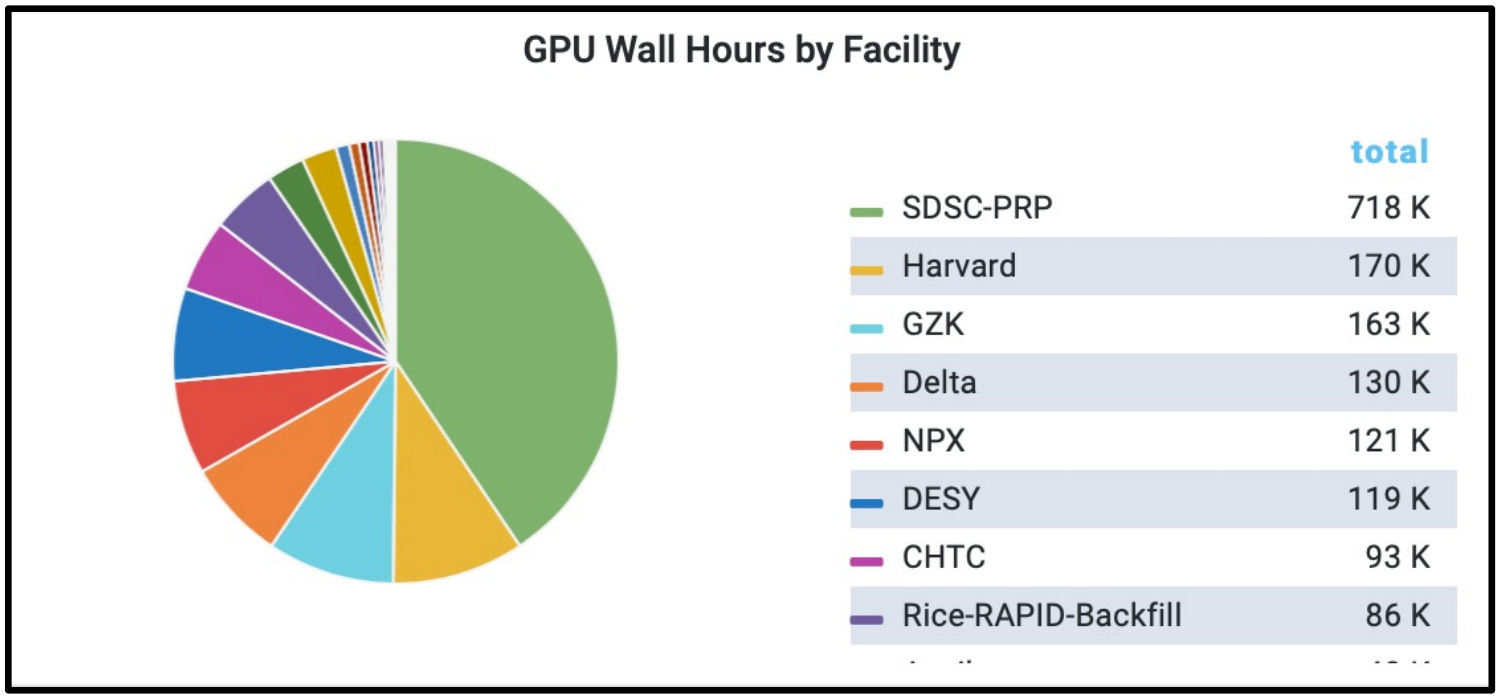

IceCube usage for a year

Many funding agencies contribute hardware in many places.

Distribution of computing matches distribution of researchers that collaborate on the science.

OSG Compute Federation

SDSC
SAN DIEGO SUPERCOMPUTER CENTER

OSG federates 100's of clusters worldwide

Owners determine policy of use.

Many allow opportunistic use of spare capacity.

OSG Data Federation

18 Caches in 4 continents … 6 in R&E network backbone

9 Data Origins in 2 continents … both public and private

Data owners integrate their data on their storage or transfer their data to public storage

Outline

- IceCube scientific motivation
- Compute & Storage requirements
- Overview of PNRP resources
- IceCube modifications to use PNRP compute (GPUs) and storage (Origins)
- Summary

Hardware on NRP is Global

NRP integrates hardware in USA, EU, and Asia

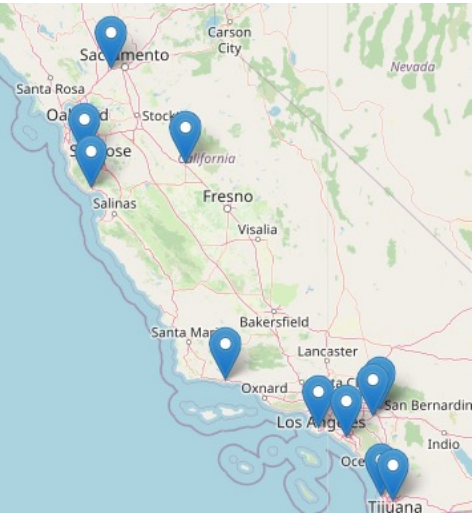

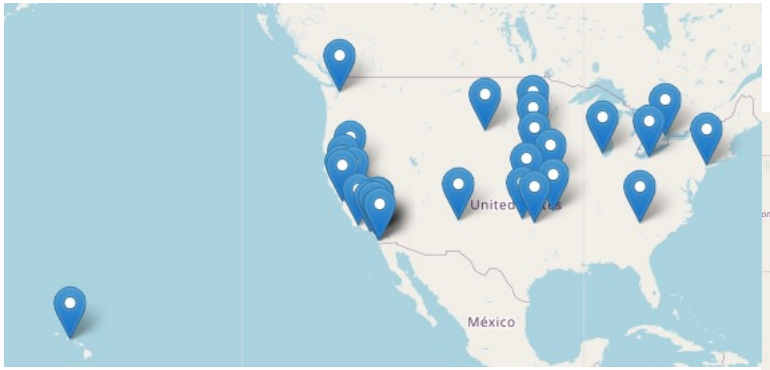

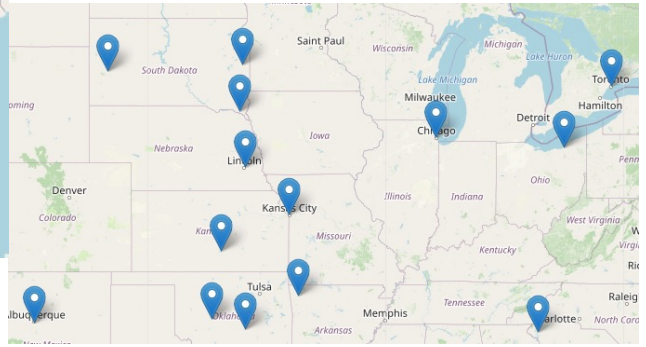

Cat-II: Prototype National Research Platform (PNRP)

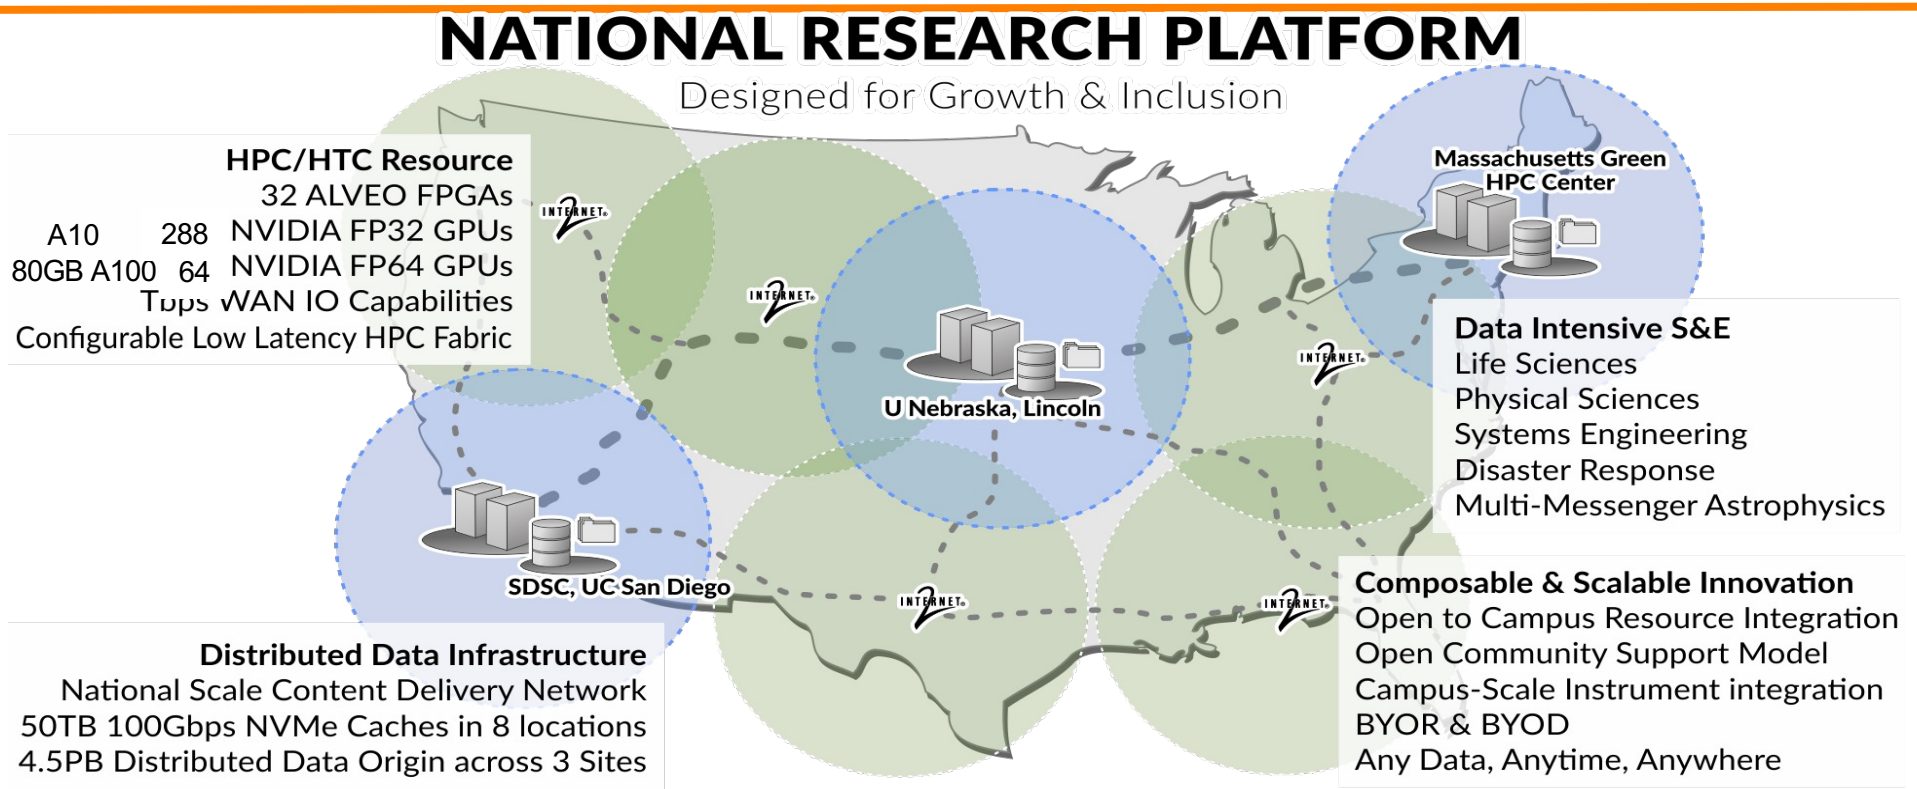

5 year project with $5M hardware & $6.45M people

Supports Nautilus, and thus the core NRP infrastructure

Promises to build on “PRP” functionality, and go beyond

NSF Acceptance Review completed, System in Testbed Phase

PI = Wuerthwein; Co-PIs: DeFanti, Rosing, Tatineni, Weitzel

Funded as NSF 2112167

Outline

- IceCube scientific motivation
- Compute & Storage requirements
- Overview of PNRP resources
- IceCube modifications to use PNRP compute (GPUs) and storage (Origins)
- Summary

IceCube and PNRP

- GPUs significantly improve Photon propagation code (compute intensive part of IceCube workflow) performance.
- PNRP is a significant provider of GPU cycles for IceCube
- XRootD managed storage, configured as OSDF origins add significant storage resources for the IceCube workflow
- Use of OSDF storage was explored in: https://doi.org/10.48550/arXiv.2308.07999

Changes to IceCube Workflow to Use PNRP compute & Storage

- Original implementation used HTCondor for compute and wrapper scripts for data movement/management
- Transitioned to use HTCondor for both aspects, leveraging native HTCondor support for OSDF Origin endpoints
- Authentication based on short lived SciTokens credentials which are automatically copied to each running job and renewed as-needed by HTCondor.

Managing data across Origins

SDSC
SAN DIEGO SUPERCOMPUTER CENTER

- Need to explicitly partition the jobs between the various OSDF Origins operated by PNRP.
- PNRP has large GPU clusters at UNL and MGHPCC. Outputs sent to co-located Origins to minimize latency and increase throughput
- Rest of resources directed output to Origin at SDSC
- Partitioning adds overhead to IceCube workflow
- Work required both on IceCube side and on OSDF to make it work. Once stabilized the error rates were low.

IceCube Data Usage on PNRP

- 160TB of data produced between Dec 2022 and May 2023, distributed among 3 sites.

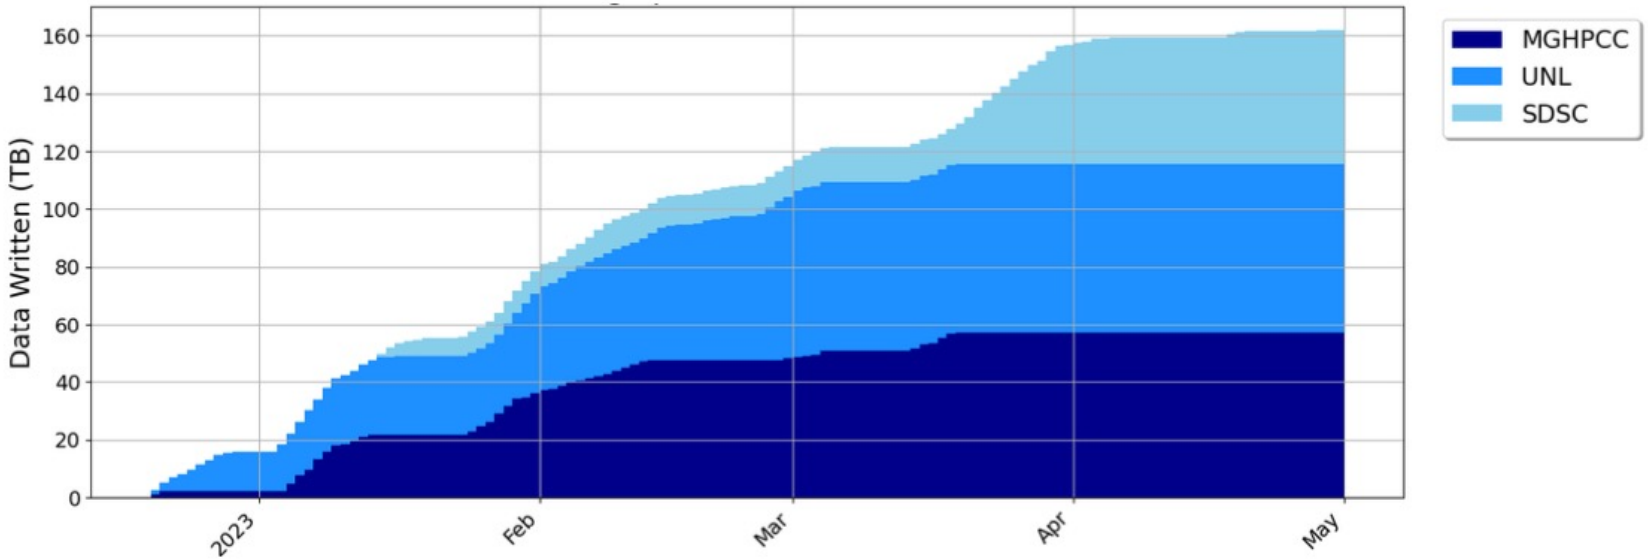

Storage per site, cumulative.

DOI: https://doi.org/10.48550/arXiv.2308.07999

IceCube Data Usage on PNRP

- Each OSDF Origin typically received less than 2 TB of data per day, but there were occasional peaks exceeding 3 TB per day.

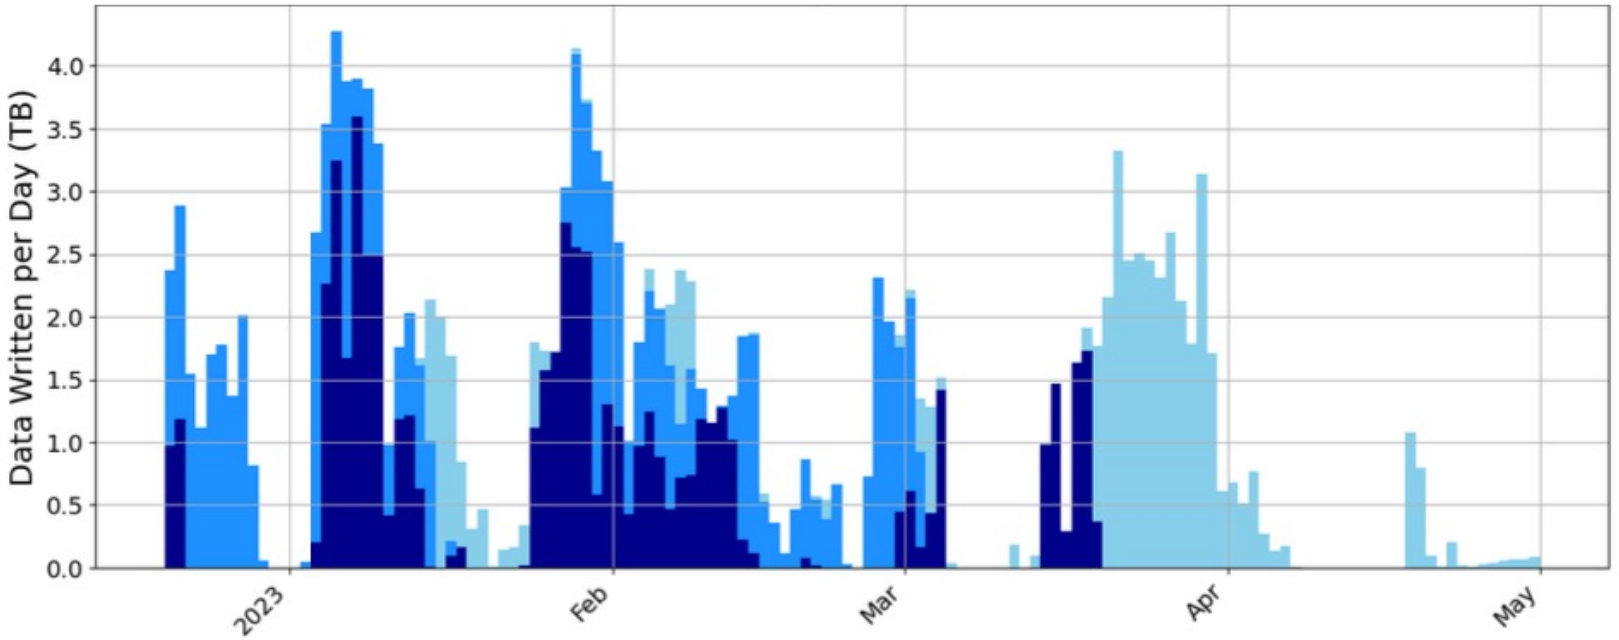

Storage per site, binned by day.

DOI: https://doi.org/10.48550/arXiv.2308.07999

Summary & Conclusion

- The IceCube Neutrino Observatory is the world's premier facility to detect high energy neutrinos and an essential part of multi-messenger astrophysics.
- The computing for IceCube is as globally distributed as its researchers.
- PNRP which itself is a distributed cluster is a significant provider of GPU cycles for IceCube
- Results from PNRP hosted computations are typically copied back to central storage at UW.
- IceCube use of PNRP storage configured as OSDF origins was also demonstrated over a long period.

Acknowledgements

- This work was partially funded by the U.S. National Science Foundation (NSF) under grants OAC-1826967, OAC-1541349, OPP-2042807, OAC-2030508, OAC-2112167 and CNS-1730158.