Homework Responses

Benedikt Riedel
UW-Madison

SCAP 2025
3 April 2025

IceCube
South Pole Neutrino Observatory

Q1

Provide us with a starting point for a resource need model that can be elaborated on over time to formally predict the computing (CPU & GPU) and storage (disk & archive) need vs time.

a. What are the parameters that drive the needs?
b. Which of them have known values? Which require more work to assign a value?
c. Which are the two parameters that have the largest risk to the total cost of computing?
d. Which are the two parameters that have the largest opportunities for R&D to reduce their values and thus the cost?

Q1 - Storage & Archive

- Assuming 30% increase from Upgrade
- Experimental and Simulation Data
  - 1 TB/day archival stream - ~5 PB Today, 475 TB/year IC+Upgrade
  - 110 GB/day over the satellite; Pass2 Offline Processing 2x satellite bandwidth
    - ~1.7 PB + ~0.7 TB archival stream for testing Pass3
  - Pass3 offline processing: aiming for O(1Hz) for each WG filter selection
    - Testing processing is 20GB/day
  - Pass2 data will need to available for another 2-3 years as analyzers shift from Pass2 to Pass3
    - 3 years 1.7 PB
  - Simulation is 6-7x data historically, currently ~1.5x
- User and Analysis Data
  - User: ~1 TB/user, 400ish users, ~870 TB
  - Analysis: ~2.9 PB, final analysis samples, extra working space for on-going analyses
  - Current version of filesystem grew out of Ceph testing, so overall “free space” is not necessarily free given inhomogeneity in hardware

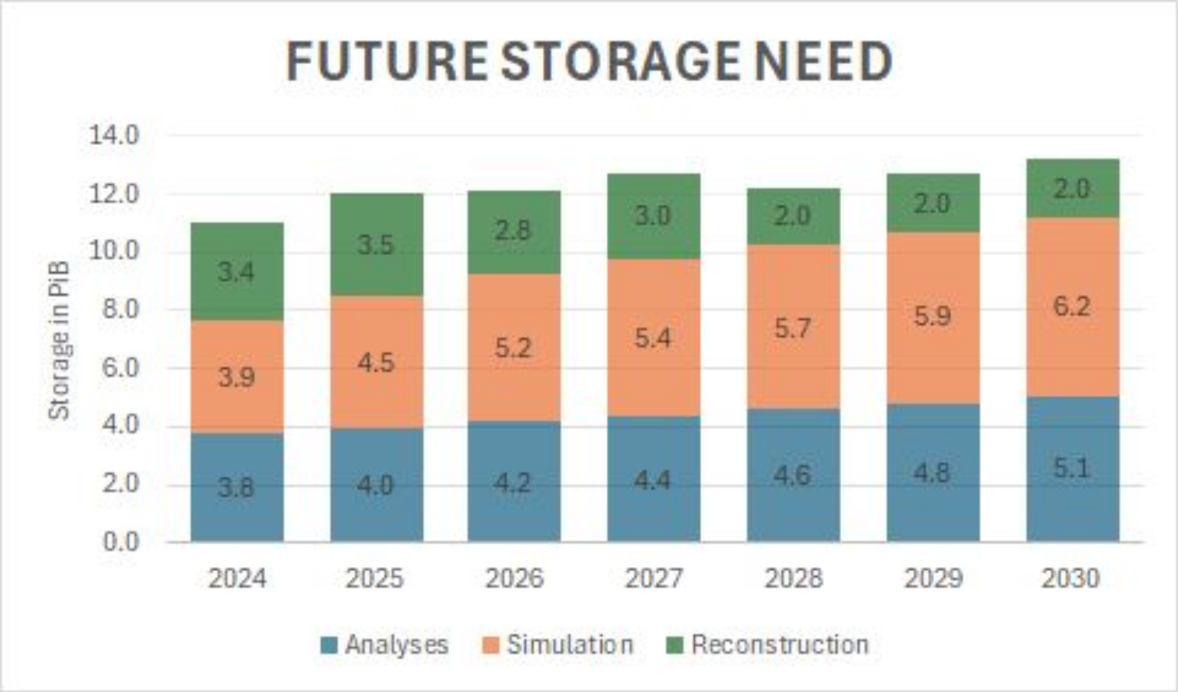

Q1 - CPU & GPU

- Dedicated resources (NPX)
  - 5000 cores, 216 GPUs
  - # of CPU cores and GPU will decline over time as older hardware ages out and is not replaced
  - Per core and GPU performance will increase
- Allocated Resources (ACCESS), In-kind (Campus Clusters, etc.), and Opportunistic
  - 5000 cores, 300 GPUs
  - Overall supply of GPUs will increase, but we are predicting an increased demand as well
    - Question about FP32 performance as more things shift towards ML
    - “Poor” utilization of resources

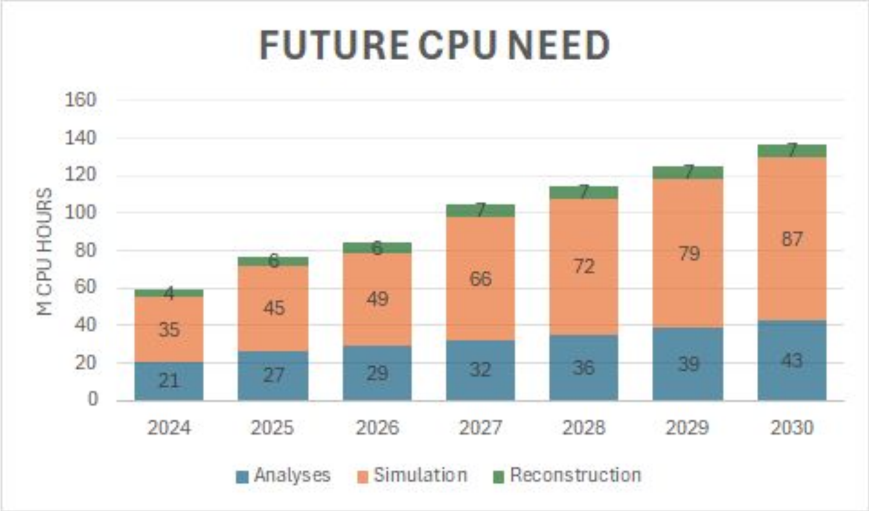

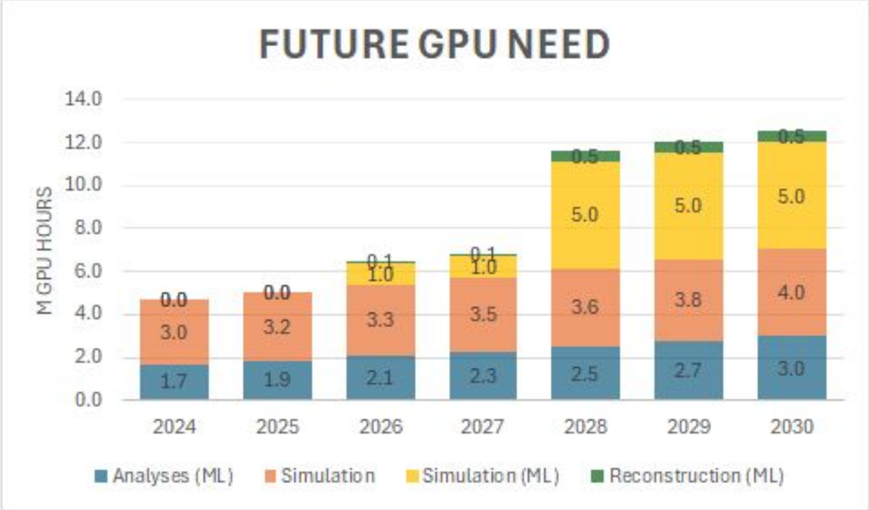

Demand Assumptions:
Start from last year's usage
CPU - 30% bump from upgrade, 10% yearly increase to simulation and analysis
GPU - 5% yearly simulation increase, 10% yearly analysis increase, reconstruction adds some ML in 2026, and additional ML in 2028 (which is 10x in simulation)

Q1 - Key drivers of computing needs (1)

- Detector trigger rates and online filter rates
  - Drive raw data rates
  - Largely fixed for Gen1
  - Expected ~30% increase from Upgrade addition
- Pass3 Offline processing
  - *STORAGE* Number and efficiency of WG proposed filters
    - Charge to WGs was to aim for O(1Hz)
    - *RISK*: Some push back from WG to save significantly more (1500 Hz)
    - These choices have multiplicative effects to “filtered simulation” samples
  - CPU/GPU
    - Choices made in filtering requests when to move to ML/AI algorithms over traditional/simpler selections
  - This requires more work to assign good values given that this is currently in progress

Q1 - Key drivers of computing needs (2)

- CORSIKA/BG simulation
  - Inefficiency of background production - avoid wasting time on background events that have no chance of passing any filters
  - ML could narrow down the parameter space that needs to be biased in CORSIKA generation and identify what events are likely to pass a filter prior to photon propagation.
  - ML could provide a means to predict the expected photon yield and timing distribution to allow for faster simulation for development of analyses - Full simulation would still be needed for detailed studies
- GPU photon tracking/ice modeling efficiency
  - Current tools have limited ability to change parameters mid-run (Snowstorm model)

Q1 c) - Highest Risk

Top risks

- CORSIKA/BG simulation production
- Inability to get the collaboration to reduce data rates at offline processing

Also worrying

- GPU demand between AI and simulation
- GPU inefficient usage - risk losing/reduced allocations at computing centers
- Age of dedicated resources

Q1 d) - R&D Opportunities

- Improvements to CORSIKA/BG
  - Bias cuts in background production
  - Can be improvement of existing tools or ML based biasing functions
  - Aim to target events that will pass event selections (not easily rejected bkg events)
- Improving GPU efficiencies
  - GPU photon simulation
  - ML/AI workflows
  - Keeping work queues full

Q2

Give us a presentation of what has been done for each of the recommendations from last time. (something like a few bullet points for each item)

Q2 - 2021-1

Recommendation:

The increase in centrally-managed full-time professional Research Software Engineer effort at the level of 5 FTE is considered essential by the committee to fully and effectively exploit the enormous scientific potential of IceCube. We encourage IceCube management to continue making a strong case in NSF to secure the required funding.

Response:

Requested 5 FTE, received 1.5 additional FTE funding from NSF for 2021-26 M&O cycle.

Q2 - 2021-2

Recommendation:

Create an S&C strategic plan that allows making a crisp case for IceCube's S&C support needs, flowing down from its science goals and capabilities, and that pro-actively plans for the significant investments needed in the areas of S&C developments, that formulates the science vision and the S&C technology goals to achieve them. The plan should help secure the required additional personnel resources and expertise.

• As part of this planning process, develop a set of quantitative science-reach goals that flow down to needed efforts and resources, thus justified by a detailed estimation of improvement goals. This should include Gen2 instrument, software, calibrations, etc.

• With this vision and strategy, determine how to create an organizational structure that will best support the mission

Response:

IceCube does not formally prioritize analyses or science topics. There are high level science goals (cosmic neutrino astrophysics, neutrino property measurements, etc.) and resource limitations (computing at the pole, satellite bandwidth, etc.). These are difficult to translate into research needs/resource requests given the diverse pool of science goals and methods and science goals evolve. So far, the supply for resources, in particular simulation, has been deemed acceptable. With a shifting science priorities (southern sky), this may not be the case over the long-run.

Leadership has been proactive in creating a longer term view of resources and goals. As part of this the Director of M&O and the Analysis coordinator have regular meetings regarding the needs of analyzers from M&O and to determine resource needs. At the same time, IceCube analysis is bottoms up in that limited information about an analysis is known outside the science working group. So far, there has been no incentive or need to determine the computational cost of a given analysis.

We created goals and metrics for the individual WBSs during the 2021 M&O renewal. Overall these goals and metrics have been met or exceeded.

Q2 - 2021-3

Recommendation:

Keep the momentum and use the effort invested in developing and documenting the ten-year data release to publish yearly updates to this comprehensive dataset.

Response:

The 10 year data set has been released on dataverse and the icecube website. We are continuing to release new datasets (14 year data set, etc.) and working on improving existing datasets with new reconstructions

Q2 - 2021-4

Recommendation:

The IceCube collaboration should be proactive with embedding themselves within the broader community of software development and ML. A proactive engagement with the broader community in S&C is important to position IceCube to affect the software tools that your science critically depends upon, benefit broadly from knowledge and technologies developed by other experiments, and to help in the professional development of your early career researchers (e.g. working collaboratively on modern open source tools from computer science and data science).

Response:

IceCube has been working on extending their collaborations with computer scientists, data scientists, other physics collaborations (namely LHC, and LIGO), etc. through involvement in proposals, e.g. EWMS with CHTC/HTCondor, A3D3 with LHC, LIGO, and LCCF's CAS program. So far, this has led to joint developed on deployment mechanisms for ML-Inference-as-a-Service infrastructure, testing of EWMS for biology applications, and planning for future funding opportunities (CSSI Framework).

Further, IceCube collaborators in Europe are part of the GraphNet effort to provide a modular and common platform for ML training for large volume neutrino detectors.

We will be encouraging students and postdocs who are working on AI/ML applications to attend conferences like ACAT and CHEP to present their work.

Q2 - 2021-5

Recommendation:

We recommend having the next follow-up SCAP meeting in 12 to 18 months.

Response:

:(

Q2 - 2021-6

Recommendation:

IceCube should watch and assess quantitatively to what extent collaborators deliver on agreements, the level and efficacy of efforts that are actually being made available, and if this approach covers required roles and expertise sufficiently and in a sustainable way. The SCAP renews the previous recommendation for a general Computing Coordinator.

• The Computing Coordinator should immediately start and lead the efforts to develop a long-term vision and strategy for IceCube S&C that addresses future IceCube scientific program's computing and data requirements and addresses the changing technology landscape. The support of IceCube management and agencies will be vital.

• This plan should include a strategy for addressing future hardware and technology choices, a mechanism to prioritize and fund advanced support and service developments, and an outline of necessary R&D.

• To implement this vision and strategy, determine how to create an S&C organizational structure within IceCube that augments the successful bottom-up communication structures and the ICC, which will best support the science mission. An organization with clearer boundaries may be beneficial.

Response:

With the MoU dashboard, we have a clearer picture of the agreed work. We are still struggling with assessing the impact of the agreed work and whether work has been completed. With the proposed "audit" of the MoUs we hope to improve our understanding of what work has been completed, or needs additional time or support. We are also considering an automatic notification of L2 managers of new or updated tasks, so they can reach out to the designated individual and their supervisor to better understand the motivation and potential impact of the contribution.

There was discussion to create a position of Computing Coordinator. This position was never officially created. It is currently rolled into the position of ICC Chair.

Overall the computing group has been trying to establish a long-term vision. Given the dynamic nature of IceCube's science and external requirements, it has been difficult to execute on a vision without delay due to competing priorities. AI/ML, in particular, has created an open question regarding future GPU demand and purchases. After comments from the panel, we will be focusing on FP32 GPUs for our internal resources to allow for more efficient simulation production.

Q2 - 2021-7

Recommendation:

- We recommend that the core M&O offline software effort, including the additionally requested FTE, be organized as a team effort, avoiding siloed expertise as much as possible.
- The core M&O software team should continue to seek further opening their activities to the broader community. Interact with experts from other facilities dealing with similar issues and actively participate in joint software projects where appropriate. Examples of areas where collaboration could be beneficial are the long-term evolution of the software framework (serialization libraries in particular, etc.) and the workflow management system (richer metadata, data popularity, storage tiering, etc.).
- The committee encourages the IceCube team to readily implement their plan to report on critical software tickets regularly and to use these reports as a tool to implement a transparent prioritization process. We think this will ultimately shorten the time-to-production for important issues.

Response:

The core software team has evolved significantly over the past 4 years. The software team now has weekly meetings and much better communication and organization regarding releases, etc. This was particularly helped by establishing goals and metrics for the current M&O cycle.

See response to 2021-4 on how we tackling community outreach.

Q2 - 2021-8

Recommendation:

Perform a quantitative estimate of resource needs for Gen2. Resource needs should start from quantifying the possible improvement in science; from there, they should flow down to the planned instrumentation upgrades, as well as to any software, reconstruction, and calibration improvements; and then flow down further to needed increases in effort and computing resources to fully exploit the IceCube Gen2 science potential.

Response:

This has been done as part of the TDR process, see https://icecube-gen2.wisc.edu/science/publications/tdr/

Q3

For each WBS area, give an effort table for both the M&O funded FTEs and the in-kind contributed FTE. (i.e. your baseline plan for the next M&O proposal, incl. the +4FTE you presented)

- How would you change the table if you had +2 FTE (where would you put them?)
- How would you change the table if you had -2 FTE (where would you cut?)

Work Breakdown Structure (WBS) to L3

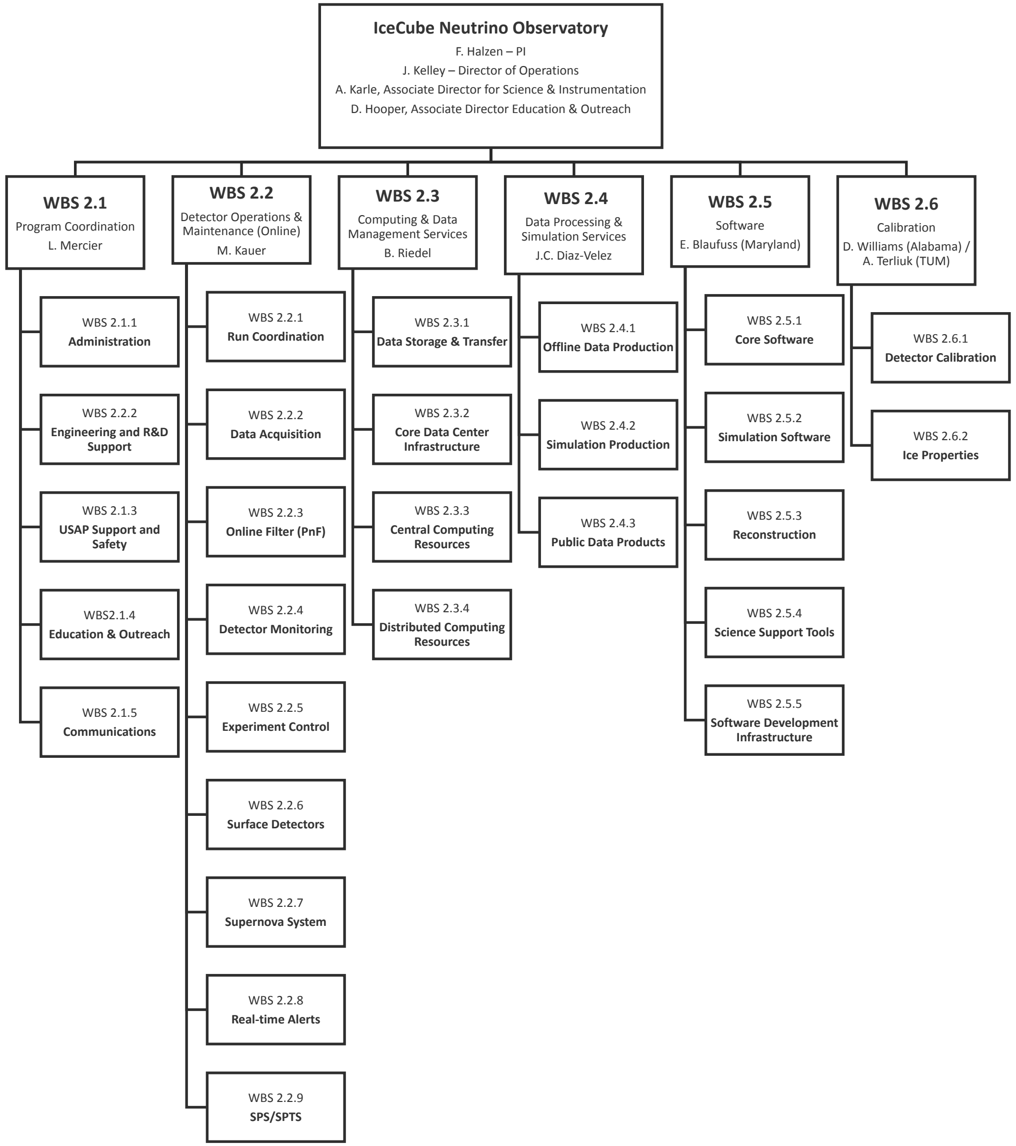

M&O Labor Baseline by WBS and Alternate Scenarios

WBS L2 Area FTE	2.1 Program Coordination	2.2 Detector Operations	2.3 Computing and Data Mgmt.	2.4 Data Processing & Simulation	2.5 Software	2.6 Calibration
M&O proposal baseline (v0)	**6.1**	**10.6** + 1 fw/systems engr.	**8.7** + 1 cybersecurity	**3.2**	**5.7** + 1.5 core sw + 1.0 ML/GPU	**2.4** + 0.4 postdoc
Current in-kind	—	2.0 DAQ 0.2 monitoring	0.2 data mgmt	0.2 simprod	2.0 core O(1)FTE WG filters	2.0 special devices
+2 FTE scenario	—	—	+1 web dev	+0.5 simprod	+0.5 reco framework	—
-2 FTE scenario	—	-0.5 fw/systems	-1 cybersecurity	—	-0.5 core sw	—