

IceCube Management and Operations Plan

PY4: April 1, 2024 – March 31, 2025

IceCube M&O Plan PY4

Revision	Date	Comments
1.0	Jul 2020	Submitted with M&O proposal. Adapted from 2019 M&O Plan.
2.0	Apr 2021	Updated for PY1
3.0	May 2022	Updated for PY2
4.0	Mar 2023	Updated for PY3
5.0	Feb 2024	Updated for PY4

1 Introduction

This document describes the Management and Operations (M&O) Plan for the IceCube Neutrino Observatory (ICNO), focused on the fourth project year, PY4, of the current five-year M&O award. The core M&O team works to develop, maintain, and implement necessary hardware and software to ensure the reliability and capability of the facility to capitalize on its science and discovery potential. These efforts build upon the experience gained from fifteen highly successful years of safe and cost-effective management of the ICNO to continually adapt to its evolving scientific mission.




Figure 1. Schematic view of the IceCube detector. The detector must be calibrated and continuously monitored to ensure collection of high-quality scientific data.

ICNO instruments a cubic kilometer of the Antarctic glacier and square kilometer on the surface above at the geographic South Pole as a neutrino and cosmic-ray detector and is shown in Figure 1. Composed of over 5000 Digital Optical Modules (DOMs), the detector, including the data acquisition, control, data handling and data filtering systems, operates continuously with high duty factor (>99%). In more than a decade of operation, IceCube data have been used in numerous discoveries in astrophysics, from the discovery of astrophysical neutrinos [1, 2], multi-messenger observations of the extremely energetic source TXS0506 [3], the detection of neutrino point source NGC 1068 [4], and the observation of neutrinos from our own Galaxy [5]. The broad scientific reach of ICNO (Figure 3) extends into other areas, notably neutrino oscillation physics [6, 7, 8].

Now more than ever, realization of our scientific vision requires the attention of a highly technically competent and dedicated team to oversee the reliable operation of the ICNO facility. Our approach to the planning and execution of IceCube management and operations is based on nearly a decade of experience, over which time the combined teamwork of the centralized M&O organization WIPAC and the IceCube

Collaboration has streamlined the process. Tasks range from detector hardware and firmware maintenance to characterization of the ice optics; from production of massive, simulated data sets to organizing outreach activities and communicating scientific results to the press. Each aspect of M&O is proven through extensive use to maximize the facility’s scientific and educational potential.

This plan describes the major systems that make up the IceCube M&O effort and outlines the plans for each system for PY4 beginning April 1, 2024 and running through March 31, 2025. These systems directly support the science output of the entire ICNO and IceCube scientific collaboration and are outlined in the ICNO M&O Work Breakdown Structure (WBS) (Table 1). Subsequent sections of this plan describe the core functionality of each WBS area and describe the management and operations effort required and plans for this period.

WBS	Title	Description
2.1	Program Coordination	Program, financial and organizational coordination for the overall M&O effort.
2.2	Core Detector Operations	Maintain and operate core detector hardware and online software systems used to collect and monitor experimental data from the detector at the South Pole.
2.3	Computing and Data Management Services	Cyberinfrastructure, cybersecurity, and administrative support to enable IceCube scientists to fully analyze ICNO collected data.
2.4	Data Processing and Simulation Services	Provide the necessary tools and personnel to perform and verify the production of large data sets for data analysis by the collaboration.
2.5	Software	Provides the software packages used by the collaboration at all levels of data analysis and production, including data filtering, reconstruction and simulation.
2.6	Calibration	Critical optical module and array calibrations needed to accurately reconstruct and extract neutrino signals from ICNO data.

Table 1. IceCube M&O work breakdown structure (WBS).

1.1 IceCube Infrastructure

Several key systems go into making up the ICNO infrastructure that constitutes the observatory [9, 10].




Figure 2. Schematic of the IceCube Digital Optical Module (DOM).

All systems require regular maintenance, support, and improvements to deliver high-quality science:

- *In-Ice and IceTop Digital Optical Module (DOM) arrays*: The DOMs, the combination of a photomultiplier light sensor, signal digitization and communication electronics and firmware, are the heart of ICNO. Over 5000 DOMs (Figure 2), are deployed in the deep ice below the South Pole and in ice-tanks on the surface to make up the In-Ice and IceTop arrays. All sensors are connected via a cable network to the SPS computing system in the ICL. The stable operation of all DOMs requires regular module monitoring and calibration to ensure each is delivering well-characterized, high-quality science data.
- *IceCube Laboratory (ICL)*: The ICL building at the South Pole station houses various IceCube systems, including cable connections to DOMs. Provided and maintained by the United States Antarctic Program, the ICL also provides power and network connections for IceCube systems housed there.
- *IceCube South Pole System (SPS)*: The SPS hardware includes DOMHub computers, commodity server class computers, remote console and power equipment, GPS units, network hardware, and UPSs. Four Iridium RUDICS modems provide low-bandwidth connectivity 24/7 for detector control and monitoring. This computing system is maintained by IceCube system administrators and winterovers to deliver a highly reliable computing platform for IceCube systems and operations. A “mirror” system - South Pole Test System, or SPTS - is maintained in the northern hemisphere and provides a testing environment for validation of new software releases and hardware upgrades without impacting the live South Pole System.
- *Online Data Systems*: The control, readout, monitoring, calibration, online filtering, transport, and storage of the experimental data produced by the array of DOMs is managed by a suite of custom software components running on the SPS hardware and consisting of a data acquisition system, data filtering system, an experiment-wide control system and a data movement and archiving system. These systems require regular software and hardware maintenance to ensure long-term reliability and stable operation of the experiment.
- *Data Warehouse and Storage Infrastructure*: The storage infrastructure at UW–Madison’s data center consists of 14 PB of online disk storage anticipated following the second phase of storage upgrades in PY3, that remain current with modern technologies and deliver the performance and scalability for handling future facility extensions. System administrators experienced in managing disk enclosures, storage networks, servers, and cluster file system software maintain and operate the storage infrastructure. They ensure that data is available for data processing and analysis tasks and that it is delivered with maximum performance.
- *Central Computing Resources*: The current IceCube HTC cluster at UW–Madison consists of nearly 200 servers providing a total of around 7000 CPU job slots and 432 GPU job slots. The HTCondor software, a state-of-the-art workload management system developed at the computer sciences department of UW–Madison, handles job scheduling at the HTC cluster.
- *Core Data Center Infrastructure*: The data center infrastructure is the glue that connects the major computing resources of IceCube. Required core infrastructure systems include distributed authentication, DNS, and email. Also, a large number of servers and services need to be deployed and maintained, such as database services, web services or tailored application servers, to fulfill science needs. System administrators maintain the UW–Madison data center infrastructure services, includes patching, monitoring, troubleshooting core services, and responding to user needs among other tasks.
- *Distributed Computing Infrastructure*: Dedicated computing resources at the level of several thousand CPU cores are needed to perform the required simulation and analysis of IceCube data.

IceCube relies on distributed resources available from collaborating institutions and wider grid computing resources to provide this computing. Support personnel at all sites coordinate and manage the distributed computing effort to produce the needed simulation. In addition, IT professionals at the UW–Madison data center manage the IceCube Grid infrastructure and middleware tools needed to exploit distributed resources and to provide efficient remote access to the data.

- *Physics Software*: IceCube’s physics software codebase covers a wide range of responsibilities and is used directly by the IceCube collaboration, including for online filtering, real-time systems, offline data reprocessing, and offline simulation generation. This software provides a wide range of functionality, from the core framework ("IceTray") to user-defined simulation, reconstruction, and analysis frameworks.

1.2 Five-year Scientific Vision

IceCube neutrino analyses span nine orders of magnitude in neutrino energy, utilizing a detector that operates continuously and is sensitive to signals from the entire sky. High-profile results include the discovery of cosmic neutrinos, some reaching energies beyond 10 PeV [11], and the measurement of the oscillation of atmospheric neutrinos in a previously unexplored energy range from 5 to 55 GeV [12]. The wide variety of physics analyses that are performed using IceCube data is evident in Figure 3. This spectrum of IceCube science is made possible by the expertise and dedication of the ICNO M&O team.

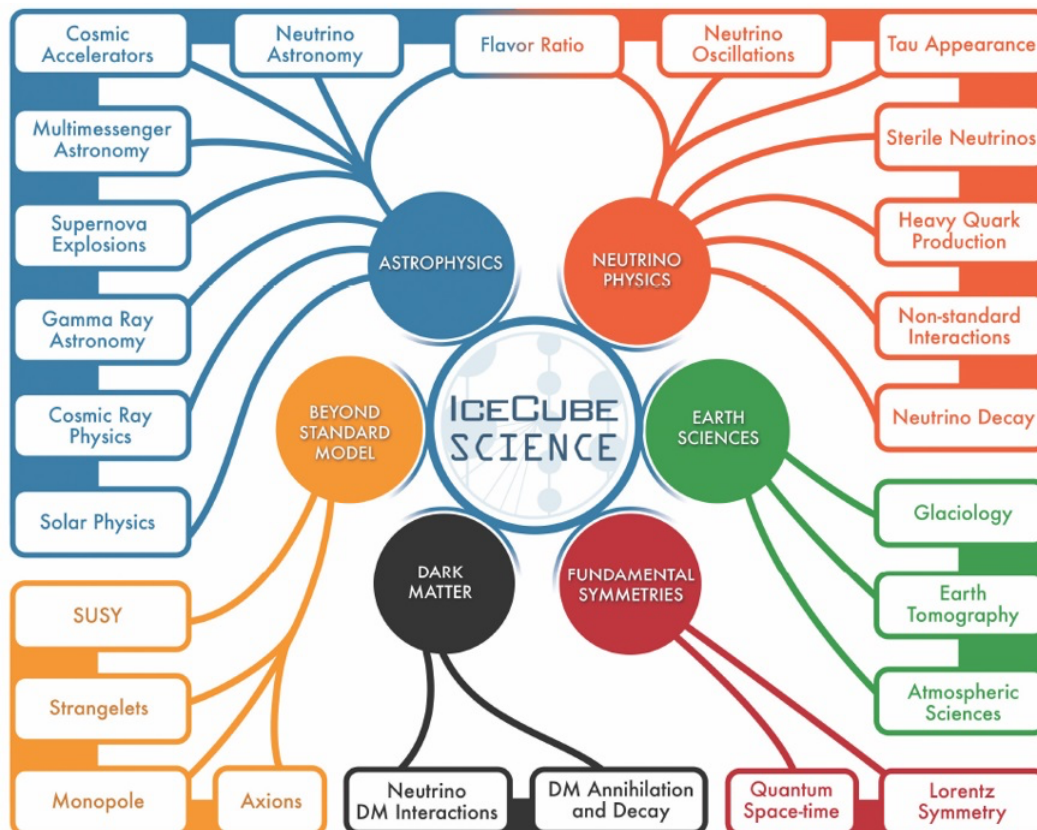


Figure 3. Spectrum of IceCube science. IceCube data are used to deliver results spanning multiple domains, from astrophysics, to neutrino physics, to searches for physics beyond our current Standard Model, to Earth Sciences.

IceCube was primarily designed as a discovery instrument but has kept a steady cadence of ground-breaking results consistently gaining performance enhancements and new functionality through operational improvements to the existing instrumentation. In this 2021–2026 M&O cycle, these improvements include:

- *Integration of the IceCube Upgrade.* The IceCube Upgrade is the first extension of the in-ice instrumentation since the IceCube detector was completed in 2010. These additions will require upgrades and new functionality within existing IceCube operational systems, as these new devices are intended to be fully integrated with existing IceCube DOMs. With the addition of 7 strings containing novel, modern instrumentation and calibration devices, the overall IceCube science program will:
 - reach new levels of sensitivity in measurements of neutrino oscillation parameters using the naturally provided beam of atmospheric neutrinos; and
 - deploy a new suite of calibration devices to further our understanding of the properties of the glacial ice in which IceCube is located. These calibration improvements will yield improved neutrino signal efficiencies, improved angular resolution, and will allow IceCube to reanalyze over a decade of historical data with higher sensitivities.
- *Enhanced surface array.* To combat the snow accumulation on the IceTop tanks that has reduced their sensitivity to air showers, a scintillator panel and radio antenna array is planned to supplement surface detection. These new detectors will restore the low-energy response to the surface array and enable precision studies of conventional and prompt atmospheric cosmic ray signals. This new suite of instrumentation will be fully integrated into existing IceCube detector systems.
- *Continued multi-messenger astrophysics searches.* Planned expansions to the existing IceCube real-time alert system will bring a larger number of public alerts that are richer in astrophysical neutrino signals, as well as enable new searches with the next generation of observatories such as the Rubin Observatory and the next generation of gravitational wave detectors. Our approach to the management of the IceCube M&O effort is to maximize our scientific discovery potential by coordinating the talents and resources from our collaborating institutions with those of the core M&O team. In this section, we describe the organization that performs the M&O functions for IceCube in this distributed model and how we provide traceability and accountability for these tasks. Leveraging the substantial efforts of the collaboration requires investment in the highest quality, experienced M&O operations staff.

2 Organization and Program Management

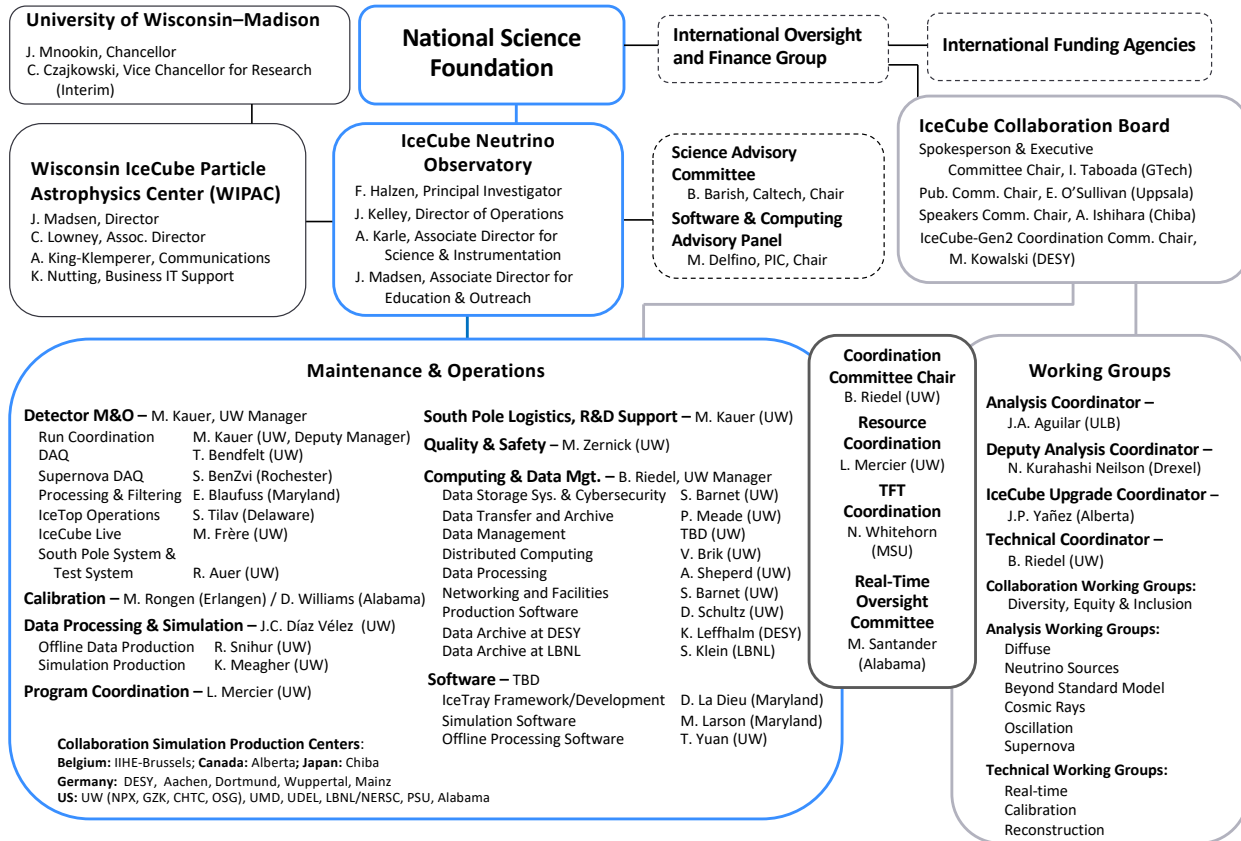
The IceCube M&O management organization integrates the IceCube Collaboration with the host institution, UW–Madison (Figure 4). The principal investigator (PI) is locally responsible to the UW Vice Chancellor for Research (VCR) and to the National Science Foundation for the overall scientific direction of the ICNO. The director of operations appoints technical professionals to serve as managers of the two M&O functions that are predominantly centered at UW–Madison: Detector Operations and Computing and Data Management. The managers in these areas work with their scientific colleagues to ensure the detector operates reliably and the data collected can be analyzed in a timely manner. The collaboration spokesperson appoints collaborating scientists to serve as coordinators for physics analyses. These appointments are subject to the consent of the collaboration.

The entire M&O scope of work is outlined in a work breakdown structure, or WBS (see Fig. 5). WBS tasks are defined in detailed MoUs, proposed by collaborating institutions and reviewed by the M&O

coordinators. The institutional leads at collaborating institutions are responsible for ensuring that the work outlined in their MoUs is completed on schedule. These MoUs are revised twice a year at the collaboration meetings and include a list of the physics group members and a head count of faculty, scientists, postdocs, and graduate students.

The IceCube Collaboration plays a leading role in IceCube, guiding both science and M&O goals. The benefits of this distributed organizational model are 1) the ability to draw highly qualified and specialized personnel from collaborating institutions to perform specific tasks in support of science or M&O, and 2) the education and training opportunities available through hands-on IceCube participation for faculty, postdocs, and students from multiple collaborating institutions.

The IceCube Collaboration Board (ICB) is the policy-making entity that guides and governs the scientific activities of the collaboration. It consists of a representative from each collaborating institution. It establishes and, as necessary, amends governance procedures and has oversight authority over science policy and goals, membership, data access, publications, representation of the IceCube Collaboration at conferences, analysis teams, and education and outreach (E&O).



February 28, 2024

Figure 4 IceCube M&O organizational structure.

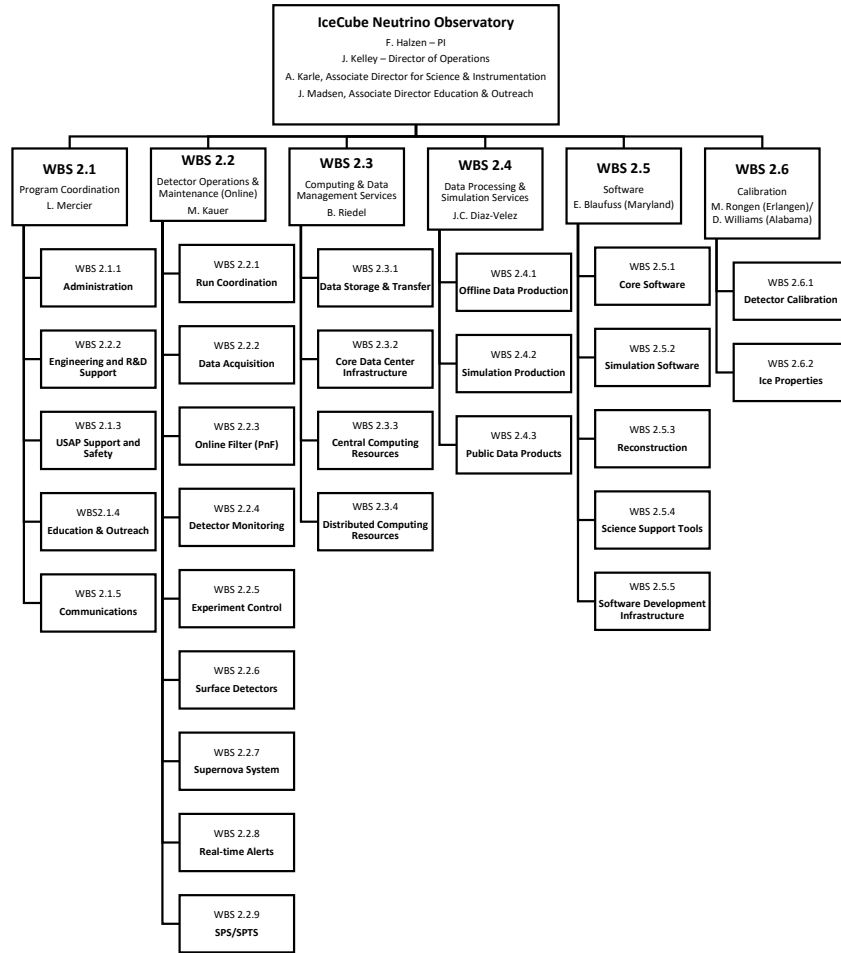


Figure 5: IceCube M&O work breakdown structure (WBS).

Executive Committee. The spokesperson, in consultation with the ICB, the PI, and the director of operations, appoints and chairs an executive committee of the collaboration board. The committee advises the spokesperson in proposing actions to the ICB and in making interim decisions, and its members represent major groups, functions, and competencies within the collaboration.

U.S. National Science Foundation (NSF): The NSF is the executive agent with responsibility for seeing that the IceCube detector meets its objectives, requirements, and technical performance standards. The NSF has a special role in IceCube because of its responsibilities in managing operations of the Amundsen-Scott South Pole Station. The IceCube M&O award is co-funded by the Office of Polar Programs (GEO/OPP) and by the Particle Astrophysics Program within the Division of Physics (MPS/PHY). The respective program directors provide continuous oversight and guidance through direct communication with the IceCube PI and director of operations.

International Oversight and Finance Group (IOFG): The IOFG was created in 2004 to provide oversight and financial support for the ICNO (including construction, M&O, and research phases). The group organizes oversight reviews to discuss detector performance and physics. A representative of the NSF chairs the IOFG, and membership is composed of representatives of the funding agencies from the partner countries supporting IceCube M&O.

University of Wisconsin–Madison: The university provides financial and management oversight and is the home institution for the IceCube M&O activities. The lead executive officer of UW–Madison, the Chancellor, delegates responsibility for research activities to the Vice Chancellor for Research (VCR), who maintains oversight of the ICNO and appoints the IceCube Director of Operations.

The IceCube PI and the Director of Operations report directly to and meet regularly with the VCR and meet as needed with the University’s IceCube leadership team, which includes the Chancellor, Provost, VCR, and Vice Chancellor for Finance and Administration. The meetings provide a forum for the IceCube PI, the Director of Operations, and the Associate Director of Science and Instrumentation to inform the University leadership team of significant issues pertinent to the management of the ICNO.

IceCube’s Associate Director for Science and Instrumentation reports to the Director of Operations and advises on matters related to science, coordination committees, and instrumentation. IceCube’s Associate Director for Education and Outreach reports to the Director of Operations and leads the IceCube E&O program by working with NSF and the IceCube Collaboration to establish E&O priorities and strategies.

Wisconsin IceCube Particle Astrophysics Center (WIPAC): IceCube’s M&O falls under the purview of WIPAC, which is the primary interface to the university administrative and support systems. WIPAC, a center within the Office of the VCRGE, coordinates the multiple roles of the university, such as lead and host institution for the IceCube construction project, for most US IceCube analysis efforts, and for IceCube M&O. WIPAC provides administrative services such as accounting, purchasing, and human resources; coordinates E&O activities; oversees engineering and R&D efforts; and collaborates with the largest participating research group.

Subawards: UW–Madison established subcontracts with some of the key U.S. collaborating institutions. These subawards provide critical support for IceCube M&O through key coordination positions. The IceCube M&O roles and responsibilities of the U.S. institutional subawards are described in Table 2.

Institution	Major Responsibilities
Lawrence Berkeley National Lab	Computing infrastructure, long-term data archival, DOM firmware support
Pennsylvania State University	Simulation production, DAQ firmware support
University of Delaware, Bartol Institute	IceTop calibration, monitoring and maintenance; IceTop simulation production
University of Maryland, College Park	Overall software coordination, IceTray software framework, online filter, simulation software and production
University of Alabama at Tuscaloosa	Detector calibration, reconstruction and analysis tools
Michigan State University	Simulation production, Northern Test System (NTS) maintenance

Table 2 IceCube M&O U.S. sub-award institutions and their major responsibilities.

The committees, panels, and boards responsible for program administration are described below.

2.1 Advisory Committees

Science Advisory Committee: The primary goal of IceCube M&O is to ensure that IceCube meets its high-level science goals and serves the collaboration in a changing environment. In consultation with the collaboration, the PI and the spokesperson appoint a panel of external experts, the Scientific Advisory Committee (SAC). The SAC's role is to meet regularly and review the performance of the M&O organization and make recommendations on scientific goals and other matters that may affect ICNO scientific activities. The current chairperson is Barry Barish from Caltech.

Software and Computing Advisory Panel: The IceCube Software and Computing Advisory Panel (SCAP) is composed of experts in the fields of software development and scientific computing. The SCAP advises the IceCube spokesperson, director of operations, global computing coordinator, and software coordinator on the most efficient and effective computing resources for IceCube, including online computing, online and offline data processing and filtering, offline computing facilities, and simulations and analysis tools support. The spokesperson and the director of operations appoint the SCAP members and chairperson. The SCAP meets semi-annually. The current chairperson is Lothar Bauerdick (FNAL).

2.2 M&O Coordination Structures

IceCube Coordination Committee (ICC): This committee provides high-level coordination of core M&O tasks and service contributions of collaborating institutions necessary for reaching IceCube science goals, ensuring that core and in-kind M&O functions meet the needs of the physics analysis working groups. The committee facilitates communication between the analysis working groups and the operations coordinators to support the collaboration towards its science reach. The coordination committee comprises M&O coordinators, working group leaders, key management personnel, and other technical experts. The chair and committee members work with institutional leads to advertise the list of needed service tasks and negotiate involvement through specific MoUs. The Coordination Committee tracks service task execution and milestones and ensures that students and postdocs working on those tasks get visibility by providing reports at the biweekly IceCube technical phone calls. The committee meets every month by teleconference and at each collaboration meeting.

Trigger Filter Transmit (TFT) Board: The TFT Board's purpose is to ensure that the IceCube detector operates in a configuration that meets the physics needs of the collaboration while respecting the limited computational and bandwidth resources available from the South Pole. Ahead of the yearly physics run transition in May, when the standard data-taking configurations for the year are deployed, the board issues a request for proposals for the upcoming season, coordinates production of Monte Carlo data sets to match the expected detector configuration, sets deadlines for physics working groups to draft proposals, and evaluates proposals for changes and additions to the set of detector triggers and online event filters. Following the transition to the new season's configuration, the TFT board requests brief reports from all physics working groups summarizing the status and quality of selected events.

Detector Operations Coordination: The detector operations working group is responsible for day-to-day operation of the detector, including data acquisition, filtering, transmission, offline processing, calibration, and maintenance. The working group is responsible for maintaining a high detector uptime and ensuring high-quality data are delivered to the collaboration. The Detector Operations Manager is responsible for coordinating group activities via a weekly teleconference.

The Run Coordinator reviews proposals for nonstandard operations of the detector, including commissioning and calibration runs, and tracks detector uptime. Subsystem experts are also involved, including online systems software engineers, calibration group members, IT professionals, and other physicists.

Calibration Coordination: The calibration group designs calibration runs for the detector, analyzes calibration data, and provides tools for utilizing the calibrations to ensure correct and efficient analysis of IceCube data. The calibration working group lead coordinates analysis of calibration data, such as DOM calibrations and in situ LED “flasher” runs, via weekly teleconferences held with students and postdoctoral researchers.

Physics Analysis Coordination: The physics analysis coordinator is responsible for oversight of the scientific output of the IceCube Collaboration. While not part of the organized M&O structure, the analysis coordinator provides key guidance to help align M&O resources with the scientific needs of the collaboration. The analysis coordinator is responsible for oversight of the physics working groups, unblinding and approval of analysis of experimental data for publication, and preparation of publications. These tasks require coordination with the M&O teams to produce simulated data samples, allocate storage and computing resources for analysis work, and help guide DAQ and filtering priorities for the detector.

Software Coordination: The software coordinator facilitates the development and maintenance of IceCube’s various software systems, including core software, online processing, online filtering, simulation, reconstruction, data acquisition, databases, online monitoring, and data transfer from the Pole. The coordinator provides software development training opportunities for collaboration physicists to ensure high-quality software. The software coordinator, along with a core development “strike team,” addresses the most urgent needs of the collaboration by participating in quarterly week-long code sprints. Recently, the IceCube open-source program was formed, where developers can make software developed for IceCube widely available to the larger scientific community.

Simulation Coordination: The simulation production coordinator, working with the technical leads from analysis working groups and the IceCube Coordination Committee, coordinates and guides the computational production of simulated physics data by working groups to balance analysis needs against the availability of computing resources. The coordinator produces quarterly reports on global simulation production status. The coordination committee meets every month to discuss simulation status and evaluate large simulation requests to determine priorities.

Real-time Coordination: The real-time oversight committee (ROC) is responsible for overseeing IceCube’s generation of real-time neutrino alerts and for coordinating the collaboration’s response to alerts from other observatories. The ROC also reviews real-time programs to ensure they are complete, correct, prompt and well-integrated into detector operations. The ROC is also charged with making rapid decisions regarding alert responses, rapid data analyses, and dissemination of alert information, ensuring that needed IceCube alert information is widely available to observatories for follow-up observations.

2.3 Engineering, Technical Support, and R&D Efforts

Ongoing engineering support for the IceCube detector continues with the maintenance and operation of the South Pole System, the South Pole Test System, and the Cable Test System. The latter two systems are at UW–Madison and enable the development of new detector functionality and facilitate investigations into

operational issues, such as communication disruptions and electromagnetic interference. Technical support provides for coordination, communication, and assessment of impacts of activities carried out by external groups engaged in experiments, or potential experiments, at the South Pole.

R&D supports subject matter experts to troubleshoot issues and coordinate internal and external efforts to enhance ICL electronics and computing performance, and interface with externally funded R&D activities that utilize IceCube facilities and infrastructure. The integration of the IceCube Upgrade detector elements into the ongoing ICNO operations will also be supported through this engineering team.

2.4 Environmental Compliance, Health, and Safety

The Environmental Compliance, Health, and Safety standards followed by IceCube are described in Section 9. The Quality and Safety officer is responsible for training, compliance, and reporting.

2.5 USAP Logistical Support

The ICNO has successfully worked with the Antarctic Support Contractor, ASC (formerly RPSC), since 2003 to coordinate logistical resources required to construct and operate the IceCube detector through cost-effective means. Requirements are detailed in the annual submission of the Support Information Package (SIP) to ASC each March / April. The yearly plan is finalized in September through ICNO concurrence with the ASC-generated Research Support Plan (RSP). Throughout the planning process there are weekly teleconferences and close coordination between the project office and ASC.

2.6 Communications, Education, and Public Outreach Coordination

The associate director for E&O leads the ICNO communications and E&O efforts, with guidance from the IceCube PI, the ICB, the ICNO E&O team, and an external advisory panel. The E&O program builds on established successes to address NSF priorities to integrate IceCube research with formal and informal education and to reach audiences that are underrepresented in STEM fields and communities. It seeks out opportunities to develop collaboration-wide E&O programs to leverage resources at 58 institutions in 14 countries and to capitalize on the potential of the ICNO to excite and engage the broader community. More traditional modes including comics featuring the adventures of penguins Rosie and Gibbs are available in a growing number of languages are being utilized along with more cutting-edge approaches. These include virtual reality experiences and an augmented reality app to provide real-time alerts for interesting events that also go out to the multimessenger community.

IceCube communications convey important science results that range from accessible summaries of IceCube publications to national and international press conferences, weekly updates from the South Pole, ICNO highlights, and opportunities for public engagement to the broader community through web, video, print, and social media channels.

2.7 Financial Management

We manage the following four different sources of funds of the IceCube M&O program, providing accountability in separate dedicated accounts through an audit trail:

NSF M&O Core and U.S. Common Fund. Covers M&O core activities, travel, M&S and services for UW–Madison, six U.S. subawards, and the U.S. contribution to the M&O Common Fund. This award mostly supports detector operations and maintenance, computing and data management, and program management.

NSF Base Grants. Supports M&O activities done mostly by graduate students and postdocs, such as detector calibration, monitoring, filtering and triggering, data quality, reconstruction, and simulations.

U.S. Institutional In-Kind. Mostly covers M&O activities done by faculty members, different fellowships, and university-funded activities.

Non-U.S. Institutional In-Kind. Institutional in-kind contributions from non-U.S. collaborators (European, Canadian, and Asia-Pacific).

IceCube M&O Common Funds. These accounts were created in 2007, the start of formal operations, to enable collaborating institutions to contribute to the costs of maintaining the computing hardware and software required to manage experimental data prior to processing for analysis. Each institution contributes based on the total number of its Ph.D. authors, updated twice a year at collaboration meetings. Effective April 1, 2010, the annual established rate per Ph.D. author is \$13,650. U.S. Common Fund contributions are covered by the NSF M&O award, while non-U.S. Common Fund contributions are paid by collaborating non-U.S. institutions to the Host Institution.

2.8 Staffing

2.8.1 Workforce Diversity

The IceCube M&O management organization is a beneficiary of the UW–Madison human resources system, which includes strategies to recruit, develop, and retain a diverse workforce. UW–Madison is committed to hiring the right talent to ensure that the university continues to be a world-class institution of higher education. The university’s goal is to provide opportunities for talented people from all backgrounds to help us maintain a highly productive, welcoming, empowering, and inclusive community. UW–Madison encourages women, minorities, veterans, and people with disabilities to apply for all of our vacancies. ICNO will continue to strive to attract outstanding candidates from underrepresented groups and support them after hiring.

2.8.2 Senior and Critical Personnel

Our senior and critical personnel (Table 3) form the leadership team that ensures the success of IceCube’s M&O and the timely exploitation of its scientific discovery and E&O potential. Our critical personnel form the core team that balances resources from the central M&O cooperative agreement and from collaboration members to maximize value and efficiency to IceCube. This is the core team that designed, built, and successfully installed the IceCube Neutrino Observatory. The full M&O staffing including personnel contributed as in-kind labor resources from collaborating institutions is included as an Appendix to the annual report.

Name, Title	Title / Institution	Responsibilities and Qualifications
<i>IceCube M&O Senior Personnel</i>		
Francis Halzen	Hilldale and Gregory Breit Distinguished Professor UW – Madison	Responsible for overall success of ICNO; IceCube PI from development through construction and transition to operations.
Principal Investigator, ICNO		
John Kelley	Scientist, UW – Madison	Ensures operation program meets performance goals, needs of scientific collaboration; ensures reporting to stakeholders (NSF, IOFG); former Detector Operations lead.
Director of Operations		
Albrecht Karle	Professor, UW – Madison	Oversees technical performance and ensures scientific objectives are met; Associate Director and overall scientific/technical lead during IceCube construction.
Associate Director for Science and Instrumentation		

IceCube M&O Plan PY4

James Madsen	Executive Director and Director (Interim), WIPAC (UW – Madison)	Plans and executes E&O and communications activities; over 20 years of experience with Broader Impacts and communications.
Associate Director for E&O		
<i>IceCube Neutrino Observatory Critical Personnel</i>		
Ignacio Taboada	Professor, GA Tech	Leads collaboration governance, coordinates member resources to support IceCube scientific program, interface to scientific community.
Spokesperson, IceCube Collaboration		
Benedikt Riedel	Scientist, UW – Madison	Coordination of IceCube computing local and distributed hardware, middleware, and software.
IceCube Computing Coordinator & Coordination Committee Chair (Interim)		
Juan-Antonio Aguilar-Sanchez	Professor, Université Libre de Bruxelles	Co-coordinators of IceCube physics analysis tasks; oversee analysis working groups.
IceCube Analysis Coordinator		
Naoko Kurahashi-Neilson	Professor, Drexel	
IceCube Deputy Analysis Coordinator		
Nathan Whitehorn	Professor, Michigan State University	Coordinator of IceCube operational resources associated with detector systems, online systems, data transfer from South Pole to Northern Hemisphere.
TFT Committee Chair		
Marcos Santander	Professor, University of Alabama	Oversee on-going real-time alerts; interface with consumers of IceCube real-time data products within scientific community.
IceCube Real-time Coordinator		
IceCube Management and Operations WBS L2 Leads		
Matt Kauer	Scientist, UW – Madison	Lead and manage the detector operations group; responsible for detector and online systems.
Detector Operations Manager		
Benedikt Riedel	Scientist, UW – Madison	Manage IceCube local and distributed computing clusters.
Computing and Data Systems Manager		
Juan Carlos Diaz-Velez	Scientist, UW – Madison	Oversee simulation production and data processing services.
Data Processing and Simulation Services Manager		
Erik Blaufuss	Scientist, University of Maryland	Oversee and maintain IceCube user software (simulation, reconstruction, analysis); manage software releases.
Software Coordinator		
Martin Rongen	Professor, Friedrich-Alexander-Universität Erlangen-Nürnberg	Manage IceCube calibration system including coordinating calibration runs and managing collaboration labor resources working on calibration.
Dawn Williams	Professor, University of Alabama	
Calibration Coordinator		

Table 3 Key IceCube personnel including M&O and Collaboration leadership.

3 Core Detector Operations

The detector operations group is responsible for the maintenance of the hardware and software systems at the South Pole and for delivery of high-quality, vetted data to the scientific collaboration. Custom surface electronics and computer systems in the ICL require regular replacement and upgrades to keep them running smoothly. The online software systems, such as data acquisition, IceCube Live, and processing and filtering, must all be continuously maintained to ensure stability (Figure 6). The detector systems are also regularly enhanced to support evolving science needs; a recent example is storing untriggered detector data to search for low-energy neutrino bursts coincident with gravitational wave alerts.

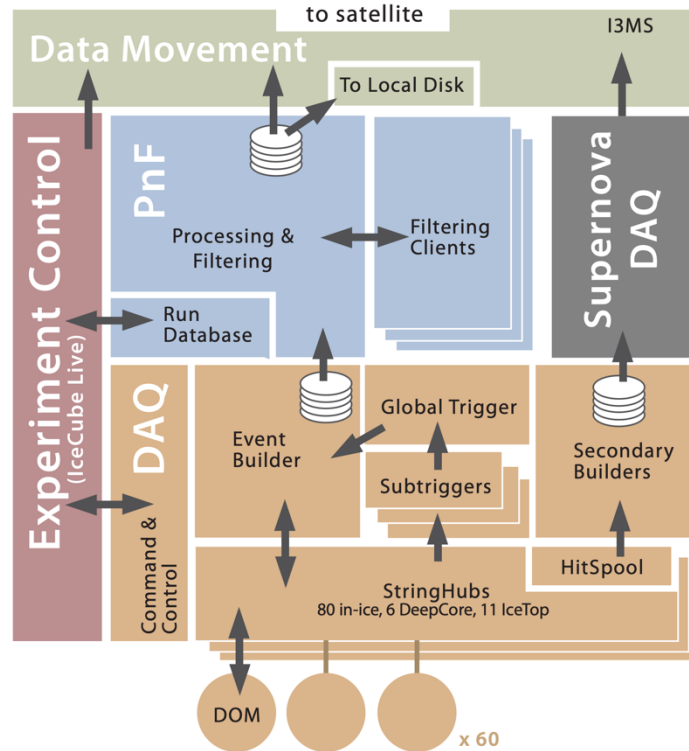


Figure 6: The IceCube online data systems, tracing the data flow from DOMs, through the data acquisition, online filters, and data movement systems, all directed by the experiment control system, IceCube Live.

The team of scientists and software developers in the operations group, with the support of the IceCube Winterovers, monitor detector performance and work to address any issues that arise. In the previous M&O period, we have successfully addressed hardware reliability issues in a number of subsystems, e.g., the DOM power supplies. Furthermore, several software systems that were fragile and/or difficult to maintain have been replaced. Our approach has proven successful, achieving detector uptimes higher than 99%.

The architecture of the online systems will allow us to integrate the additional IceCube Upgrade strings into the existing detector. All subsystems will need to be modified to some extent to support new sensors and calibration devices. However, the current design is flexible enough to support this without a complete redesign, saving significant time and effort. Once the Upgrade strings are deployed, operation of the detector will include both original and Upgrade strings. To streamline ongoing hardware maintenance and enhance reliability, the original, aging DOMHub computers will be replaced with backward compatible Upgrade-style hubs.

Other maintenance activities that fall under the purview of the operations group include enhancing the IceTop surface array to address degraded efficiency from snow accumulation and serving as the caretaker of the Askaryan Radio Array.

3.1 South Pole System and Test Systems

The South Pole System (SPS) is a computing system developed and maintained by the detector operations group that supports the IceCube online data acquisition and filtering tasks at the South Pole. SPS hardware includes commodity server-class computers as well as network hardware such as switches and firewalls. SPS provides infrastructure services such as server name resolution, email, system monitoring, and databases. To maximize availability, SPS incorporates UPSs to handle short power outages, and remotely managed power distribution units are available to enable power cycling of any component in the ICL remotely from the South Pole station. The system is designed to be very robust since it must reliably support data taking with minimal intervention; at the same time, it is simple, since it has to be operated by the WinterOvers, who are newly hired and trained each year.

The commodity computer servers in the SPS are generally replaced every four to five years in order to profit from technological advances that maximize computing power, to minimize the risk of component failure, and to ensure a reliable supply of spare components. Other components with a finite life cycle, such as UPS batteries, must also be regularly replaced. Certain aging infrastructure components that have not been replaced during the previous M&O cycle also have been or will be replaced during this M&O period, including the UPS and network switches, after approximately 10 years of service.

The UPS and network switch life-cycle replacements were scheduled over PY1 to PY3 but were delayed due to the limited pole seasons during COVID-19. The network switch replacement was successfully completed during the 2023–24 pole season, while the UPS replacement is still ongoing due to significant cargo delivery delays (see PY3 annual report) and will be completed by early PY4. The firewall life-cycle replacement is scheduled for PY4 along with an upgrade to the remotely accessible power distribution units (PDUs) and cable management.

The life-cycle replacement of the SPS servers, originally planned for PY3 and PY4 (50% each season), is now deferred until after the IceCube Upgrade installation field season of 2025-26. The hardware is operating stably with a low failure rate, and performance has been verified to scale to the necessary levels for the Upgrade trigger. This ultimately reduces risk to Upgrade integration by reducing cargo needs and ICL disruption during these field seasons.

The South Pole Test System (SPTS) is a test system located on the UW–Madison campus that replicates the basic functionality and performance characteristics of the production SPS and is a reasonable proxy for the detector and DAQ systems. It provides an environment to build and verify software subsystems and perform hardware and software evaluations prior to deployment on the production system. Engineers evaluate software and firmware DAQ updates on the SPTS via one or more of the following: (a) eight real DOMs kept at subfreezing temperatures for evaluation of firmware and software updates; (b) a full string of DOM mainboards connected to a DOMHub for string-level tests; (c) a full-length IceCube cable connected to several DOMs for communications and device driver testing; and (d) a special playback mode of real untriggered data streams coming directly from the deployed DOMs that have been captured at SPS and for use in SPTS to generate realistic load conditions.

The SPS system administrator and the Winterovers are responsible for maintenance and operations of SPS. This includes preventive maintenance, troubleshooting, and upgrades. The SPS system administrator evaluates all new hardware and software infrastructure on SPTS to validate performance and reliability.

The operations group also provides support to the wider collaboration for adding new functionality to SPTS, and ultimately to SPS, in response to new science objectives.

The IceCube Upgrade Northern Test System (NTS) is an extension of SPTS located at Michigan State University for the testing and integration of Upgrade hardware into the IceCube data acquisition system. The NTS installation contains Upgrade DOMs, FieldHub readout computers, and the Upgrade surface communication/power/timing system. NTS is seamlessly integrated into SPTS via a virtual private network (VPN) connection. After the Upgrade construction project is complete, an engineer will continue to maintain NTS as a testing and validation resource for this newly deployed portion of the detector.

3.2 Run Coordination

In order to maintain the high level of analysis-ready “clean” uptime, nonstandard operations on the detector are carefully planned and reviewed. This includes commissioning runs for new software releases, calibration runs, and test runs for the annual “physics run” transition. The IceCube run coordinator is responsible for reviewing nonstandard operation requests and approving them as they see fit. All operations on the detector, including maintenance, are recorded to an email-based logbook by the Winterovers or operations expert.

The Winterover coordinator is responsible for the annual hiring, training, and regular management of the two IceCube Winterovers. The Winterover training, which occurs over a period of two months prior to deployment, makes heavy use of SPTS and involves full-time instruction in system administration, detector electronics, software subsystem operations, calibration, emergency procedures, and safety. In addition to the winterover coordinator, other operations experts regularly lead portions of the training relevant to their subsystem.

3.3 Data Acquisition

The data acquisition (DAQ) system collects raw hits from the individual DOMs, forming triggers with all the relevant data for physics events in the detector. Diagnostic and calibration data are also collected, as are raw counting rates for all DOMs, used for the supernova triggers. Performance of the DAQ hardware and software is a major driver of the quality of data for physics analysis.

The basic data acquisition (DAQ) hardware surface component is the DOMHub, a rack-mounted computer that houses both commercial and custom hardware and is connected to up to 64 DOMs. A GPS master clock system provides accurate timing to the DOMHubs through a series of custom clock fanout units. The DOMHubs, their internal components, and associated cabling must be maintained to prevent malfunctions and repaired quickly as necessary to minimize detector downtime and maintain a stream of high-quality data.

The DAQ software systems, responsible for readout of the DOMs and forming triggered events, consists of DOM software, DOM readout card device drivers, DOMHub software, and the DAQ software components. DAQ software engineers are accountable for the uptime of the DAQ and the integrity, correctness, and completeness of the data it produces. They regularly test and upgrade DAQ software, respond to new feature requests, and provide appropriate documentation for the operators. The software engineers also maintain interfaces to other online systems, including the supernova DAQ and detector monitoring.

To improve reliability in the case of server failure, architectural improvements in the DAQ have moved toward decoupling the DOMHub data-taking component from the higher-level trigger and event builder components. This “hitpooling” system also provides the ability to save up to two weeks of all untriggered sensor data. DAQ software engineers will continue to develop this system, including dedicated components that allow hits from the Upgrade FieldHubs and existing DOMHubs to feed into a unified DAQ trigger

system (Figure 7). The Upgrade FieldHub integration may also incorporate new triggering layers, such as a software-defined local coincidence system that mimics similar functionality built into the existing IceCube cables. Software development milestones to support the integration of the Upgrade have been re-baselined with the planned deployment in 2025–26 and are tracked now as part of the Upgrade project (WBS 1.6).

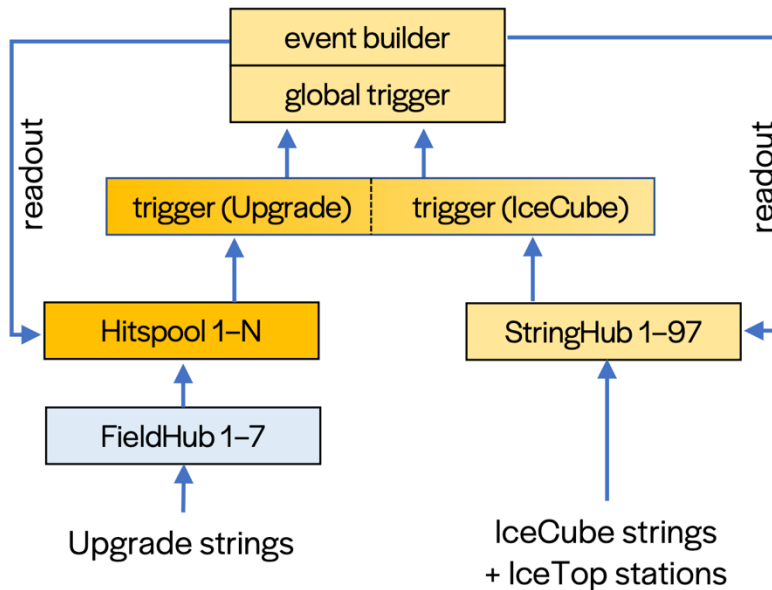


Figure 7: Integration of the IceCube Upgrade into the existing DAQ. Both original and Upgrade hits feed into a single global trigger and are built into combined events.

All custom hardware components in the DOMHubs and clock fanout system contain obsolete parts and cannot easily be remanufactured. While the failure rate of these components is currently low, the design of new DOM readout hardware for the Upgrade provides a path forward to simplify long-term maintenance of the detector. After completion of the Upgrade deployment, we plan to retrofit the existing detector by replacing the nearly 20-year-old DOMHubs with new Upgrade-style “FieldHubs” located in the ICL. The FieldHubs have been designed with hardware that is backward compatible with the current DOMs. A firmware engineer will adapt the communications protocol in the FieldHub FPGAs as needed. The legacy clock fanout system will also be replaced by the Upgrade White Rabbit timing system, unifying the timing and communications system for the entire detector. The DOM power system will also be upgraded to a rack-based system, and as a side benefit, we estimate that this will save up to 3 kW of power. While preparations for this DOMHub replacement continue, since the Upgrade deployment is now in PY5, installation of the replacement hardware is no longer planned for this M&O period.

3.4 Supernova Data Acquisition

The supernova DAQ (SNDAQ) is a parallel data acquisition system that monitors the DOM noise rates for the signature of a Galactic supernova neutrino burst—a coherent rise in the rates in all sensors. A rate excess triggers a supernova candidate alarm to be issued via satellite, and untriggered hitspool data are saved. Highly significant alarms are forwarded to the Supernova Early Warning System (SNEWS) [13]. A recently implemented reverse link also saves IceCube untriggered data in case of a community SNEWS alert.

Collaboration scientists are responsible for core SNDAQ development as in-kind contributions. DAQ and IceCube Live software engineers are responsible for integrating SNDAQ into the experiment control,

monitoring, and DAQ systems. A Python-based rewrite of SNDAQ, PySNDAQ is under development and will simplify ongoing maintenance compared to the legacy system written in C++. We plan for a first online deployment of PySNDAQ in PY4, initially in parallel with the existing SNDAQ and for followup analysis of gravitational wave alerts from LIGO.

3.5 Detector Monitoring and Experiment Control

IceCube Live is the software system that integrates control and monitoring of all of the detector's critical subsystems into a single, virtual command center. It provides an interface for monitoring the detector both via automated alerts and with interactive screens displaying current and historical states of the detector and associated subsystems. Web-based and command-line user interfaces provide maximum accessibility and flexibility to the operators located both locally at the South Pole and remotely in the Northern Hemisphere. IceCube Live is mirrored at SPTS to test upgrades and changes before deployment.

Software engineers are accountable for the uptime of IceCube Live and for maintaining, troubleshooting, supporting, and evolving the interface to subsystems that control the detector. The engineers support physics working groups and operators to add needed functionality; a recent example is the integration of the real-time alert system.

Operation and monitoring of the detector and real-time neutrino alerts to the scientific community both require a 24/7 network connection to the South Pole, but high-bandwidth satellite coverage is limited to approximately 10 hours a day. IceCube has developed a messaging system technology (I3MS) to move critical monitoring and alert traffic off of the station link, using Iridium RUDICS satellite technology. I3MS also provides a real-time communications link between the winterovers and northern operations experts via a chat bridge between IceCube Live and a commercial messaging service used widely by IceCube. A software engineer updates and maintains the operations-critical I3MS software.

IceCube detector monitoring is the system within IceCube Live that provides a comprehensive set of tools for assessing and reporting data quality. The monitoring coordinator oversees development and testing of the monitoring system, which is implemented by IceCube Live software engineers and other operations subsystem experts. IceCube collaborators participate in daily monitoring shift duties by reviewing information presented on the web pages and evaluating and reporting the data quality for each run. The shift takers, frequently graduate students, compile reports on detector performance during their shift. A summary of the monitoring shift is given at weekly teleconferences, where experts determine if the detector is operating as expected and take actions to correct malfunctions as needed.

Software engineers will update both the experiment control and monitoring portions of IceCube Live to support the new calibration devices and sensors of the IceCube Upgrade. Special calibration runs will be controlled the same way as standard data taking, to ensure proper tracking of calibration-generated events and light in the detector. The Upgrade sensors, most of which include multiple photomultiplier tubes, have additional monitoring requirements tailored to their specific hardware. These monitoring quantities will be seamlessly integrated into the existing database system, and new tests and visualizations will be created for monitoring shift takers to validate data quality.

3.6 Online Filtering and Real-time Alert System

Triggered events from the DAQ are immediately analyzed by the online processing and filtering system (PnF). The volume of raw data produced by DAQ exceeds the bandwidth available in IceCube's satellite allowance. The PnF system calibrates the raw waveform data from the DOMs, extracts the time of arrival and amplitude for the light signals observed by the DOMs, collects and reports data quality and monitoring information to IceCube Live, and applies the "Level 1" filters to the all data events. By selecting events of

interest to physics analysis with this Level 1 filter and preserving only the time/amplitude information extracted from DOM waveforms, data volumes are greatly reduced. Additionally, all events are saved using this compressed format at the South Pole, allowing for reduced archival data sample sizes.

The online PnF system also hosts dedicated neutrino signal searches that quickly identify likely neutrino event candidates. These events are immediately transferred to the north via the I3MS messaging system. Dedicated receivers at the UW data center perform further analyses of these events and issue a wide variety of alerts to the astrophysical community. Most notably, following the detection of a high-energy neutrino candidate with a high probability of astrophysical origin, a GCN alert [14] is automatically issued. Roughly 30 alerts are produced per year, including IC-170922A, which launched the multimessenger observation campaign of TXS 0506+056, and are generally publicly available within ~30 seconds of detection, enabling immediate responses by other observatories.

PnF system experts maintain the online system at SPS to ensure filters are being properly applied; test new filters, features, and alerts at SPTS; and respond to and debug unexpected errors. This effort ensures that the online filtering system and real-time alerts produce the highest quality data. Each year, the TFT board revises the Level 1 filter settings to ensure that they meet the evolving physics needs of the collaboration and that the most effective and robust reconstruction and filtering tools are used in online settings. PnF system experts implement these filters, ensure CPU and data volume match requested values, and prepare well-tested releases for deployment to SPS.

PnF developers are also working to implement several system enhancements. With the arrival of the IceCube Upgrade, the PnF system will need to calibrate and analyze data from several new sensor types, integrating both old and new DOMs into a unified Level 1 filter selection. Additionally, the PnF system was designed and built on libraries and standards established more than 10 years ago. The effort to update these underlying libraries will modernize the system to help reduce the maintenance burden over the next decade of operation. Finally, collaboration physicists and software engineers continue to work together to define fast, robust analysis schemes that can alert other observatories for follow-up observation of interesting events, localized in time and/or direction. Supporting these ever-evolving classes of neutrino alerts, including dedicated alert generation and catalog tools at the northern data center, will continue to be a priority for PnF system experts.

3.7 Surface Array Maintenance

The IceTop surface array requires its own tools for calibration, monitoring, reconstruction and simulation. The IceTop DOMs are embedded in ice contained in tanks on the surface, which are subject to environmental changes that must be monitored. In addition, specialized modes of operation are required to maximize IceTop's science potential, which includes study of solar particle activity and high-altitude weather in addition to the cosmic-ray science. The IceTop Operations Coordinator is responsible for monitoring of the physical condition of the IceTop detectors, including coordination of annual surveys of snow accumulation above the tanks and surrounding environmental conditions at the South Pole. The IceTop Operations Coordinator also monitors the quality of IceTop data and coordinates any corrective actions required to address malfunctions or other conditions that degrade IceTop data.

The snow accumulation above the IceTop tanks increases the energy threshold of the detector for cosmic-ray air showers and decreases the trigger rate by approximately 15% per year, negatively impacting the science capabilities of the array. Initial maintenance and operations included snow management plans that involved removal of the snow from the tanks; however, the support cost of this approach has proven burdensome, and it has been discontinued. We have begun to implement a plan to restore the efficiency and

science potential of the surface array by installing low-cost scintillator panels and radio antennas within the IceTop detector footprint. All elements of the array are elevated above the snow surface.

The instrumentation for the surface array upgrade will be provided as in-kind contributions. Scientists and a project engineer will support development of the surface array upgrade hardware and integration into the power, communications, and timing systems of the IceCube Laboratory. A DAQ software engineer will support integration of the new instrumentation into the online data systems. Key elements of both the hardware infrastructure and software architecture are shared between the surface array and the IceCube Upgrade, meaning integration of one supports the other and vice-versa.

A prototype station of the surface array upgrade has been operated successfully since 2018 and upgraded to the production configuration in 2020. All elements of the station were successfully elevated during the 2022–23 pole season. Further deployments of the array have been limited by COVID-19 and other logistics constraints, but we are exploring options to deploy additional stations during PY4.

3.8 Caretaking of the Askaryan Radio Array

The Askaryan Radio Array (ARA) is an ultra-high-energy neutrino detector at the South Pole [16], distinct from the IceCube Neutrino Observatory. Five detector stations search for the broadband radiofrequency Askaryan emission from neutrino interactions in the ice. ARA is no longer under construction and is in a stable data-taking mode in its current configuration. Since the ARA power and optical fiber connections terminate in the ICL, ARA already utilizes the IceCube infrastructure, including networking, data movement, and server support.

The IceCube WinterOvers support continued ARA data-taking via the following minimal activities: maintenance of the ARA data-taking server in the ICL; support for ARA data archiving at the South Pole and in the northern data warehouse; and as-needed troubleshooting support. IceCube deployers may perform limited ARA station maintenance during the austral summer pole seasons, as support permits.

3.9 Project Year 4 Plan

The following tasks are planned for the 2024-25 South Pole summer field season:

- Annual data archival hard drive exchange
- ICL firewall life-cycle replacement
- ICL rack power distribution unit upgrade
- Askaryan Radio Array DAQ upgrade
- IceACT DAQ upgrade
- Additional surface array station deployments, support permitting

Data acquisition (DAQ) development during PY4 is focused on integration of the IceCube Upgrade DOM data streams using SPTS and NTS, and planning for IceCube calibration device support. While this effort is largely supported by M&O DAQ software engineers, the milestones are tracked as part of the Upgrade project.

Initial components of the Python replacement of the supernova DAQ (PySNDAQ) are successfully running at SPTS and SPS. The full replacement at SPS is planned for PY4.

4 Northern Hemisphere Cyberinfrastructure

4.1 Data Analysis Computing Infrastructure

IceCube data analysis computing utilizes both interactive and batch infrastructures. There are nine servers available in the interactive infrastructure used by researchers to perform data analysis and develop new methods—eight focused on traditional data analysis and one dedicated to machine-learning-based analysis. The batch infrastructure consists of an IceCube-dedicated high-throughput computing (HTC) cluster located at UW–Madison and a globally distributed resource pool. The current IceCube HTC cluster consists of nearly 200 servers providing a total of about 4000 CPU and 300 GPU job slots.

These dedicated resources are not sufficient to ensure timely data processing and simulation production. To increase overall capacity, IceCube relies on distributed resources available at collaboration institutions, opportunistic computing consortia, such as OSG, and allocation-based high-performance computing (HPC) consortia, such as XSEDE/ACCESS. The current capacity plan foresees providing 10% of the CPUs needed for simulation at the UW–Madison data center, 15% at DESY-Zeuthen, and 75% at other collaboration sites and compute consortia, such as PATH. For GPUs, the plan is to provide 50% of the needs through M&O-funded resources and 50% at other collaboration sites and compute consortia. A majority of the M&O-funded resources will remain at UW–Madison, while others will be hosted at MSU and UMD. M&O IT has and will continue to seek collaboration to procure additional resources.

Working with researchers at the University of California San Diego and the San Diego Supercomputing Center (SDSC), and with additional funding from NSF and Internet2's Exploring Clouds for the Acceleration of Science (E-CAS), IceCube generated the largest GPU pool in history using cloud resources from three major vendors: Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure [16, 17]. We determined that IceCube would be able to utilize future exa-scale resources and that the cloud would be too expensive to use for IceCube's production computing needs. This work was awarded the 2020 HPCWire's Editor's Editors' Best Use of HPC in the Cloud (Use Case) award¹.

HTCondor [18], a state-of-the-art workload management software system developed at UW–Madison, handles job scheduling at the HTC cluster and the distributed computing resources. The IceCube system administrators who maintain and operate the cluster collaborate closely with the HTCondor team, providing feedback on specific use cases and ensuring the system fulfills IceCube's needs. They support users by providing guidance on HTC best practices and support the delivery of science-ready data by ensuring that offline processing tasks run as prioritized, and that the end-to-end infrastructure stack (computing/network/storage) is available and delivers optimum performance.

Using infrastructure funds provided by GCP and in collaboration with SDSC, we are currently testing the use of Kubernetes and cloud resources for our data analysis and simulation infrastructure, particularly those that are time-sensitive, such as the Multi-Messenger Astrophysics data stream and analysis. We have used the funds to benchmark different GPU and GPU-sharing paradigms and are working on understanding how ARM-based CPUs could supplement our infrastructure. Our work GPU benchmarking and sharing work was awarded the 2022 HPCWire's Editor's Editors' Best Use of HPC in the Cloud (Use Case) award².

IceCube has been selected as one of 21 codes and 'grand challenge'-class science problems as part of the "Characteristic Science Applications" with Leadership-Class Computing Facility program at the Texas

¹ <https://www.hpcwire.com/20-hpcwire-awards-readers-editors-choice/>

² <https://www.hpcwire.com/2022-readers-editors-choice-cloud-use-case/>

Advanced Computing Center. This will allow us to play a role in shaping the next generation of NSF's MREFC investment into the nation's cyberinfrastructure.

Currently, 1.5 full-time system administrators spread over four people and 1.5 full-time software engineers maintain this infrastructure which includes resource aggregation and usage monitoring as well as workflow management systems. To fulfill the data analysis computing capacity plan, M&O will require a significant investment into the underlying cyberinfrastructure of IceCube's distributed computing grid. We will need an additional software engineer to allow for expansion, proper load balancing between sites, and improvements to the workflow management—in particular, data management for production workloads.

4.2 Data Center Infrastructure

Data management and analysis requires a supportive data center infrastructure. This includes the physical space to locate, power, and cool the hardware and additional cyberinfrastructure for business services, such as document management and e-mail, authentication services, monitoring, network connectivity, and cybersecurity.

IceCube computing facilities are currently hosted in two UW–Madison locations, one off-campus at a co-location facility contracted through UW-Madison and one on-campus at the Physics Department. The facilities provide the total capacity to power and cool IT equipment. UW-Madison has recently started a new program for research cyberinfrastructure that will allow us a bigger footprint in the off-campus leased facility at no cost for the foreseeable future. The space, power, and cooling for these facilities are in-kind contributions from WIPAC and UW–Madison to this project.

The IceCube Collaboration currently spans 13 countries. Many of the data products and services are hosted at the UW–Madison data center, hence excellent global network connectivity is essential. A large part of IceCube's computing requirements are being met using distributed resources; therefore, reliable, high-speed access to the UW–Madison data center storage infrastructure is required to use these resources effectively. With our move from WIPAC headquarters to a co-location facility, we also transitioned from self-administering our network infrastructure to having a large portion of the administration performed by UW–Madison central information technology services.

We will continue to maintain a cybersecurity program at the UW–Madison data center. We are currently taking part of NSF's TrustedCI Cybersecurity Center of Excellence framework cohort to ensure we are in compliance with and use the best cybersecurity practices.

These tasks require a full-time system administrator and half of a full-time software engineer. A large portion of their tasks are to maintain virtualization infrastructure, cybersecurity, and networking.

4.2.1 Data Management Infrastructure

Scientific computing and data management is a constantly evolving ecosystem. IceCube computing efforts (see Fig. 8 for an overview) have been and will continue to be focused on providing a dependable and robust platform for scientists to make discoveries.

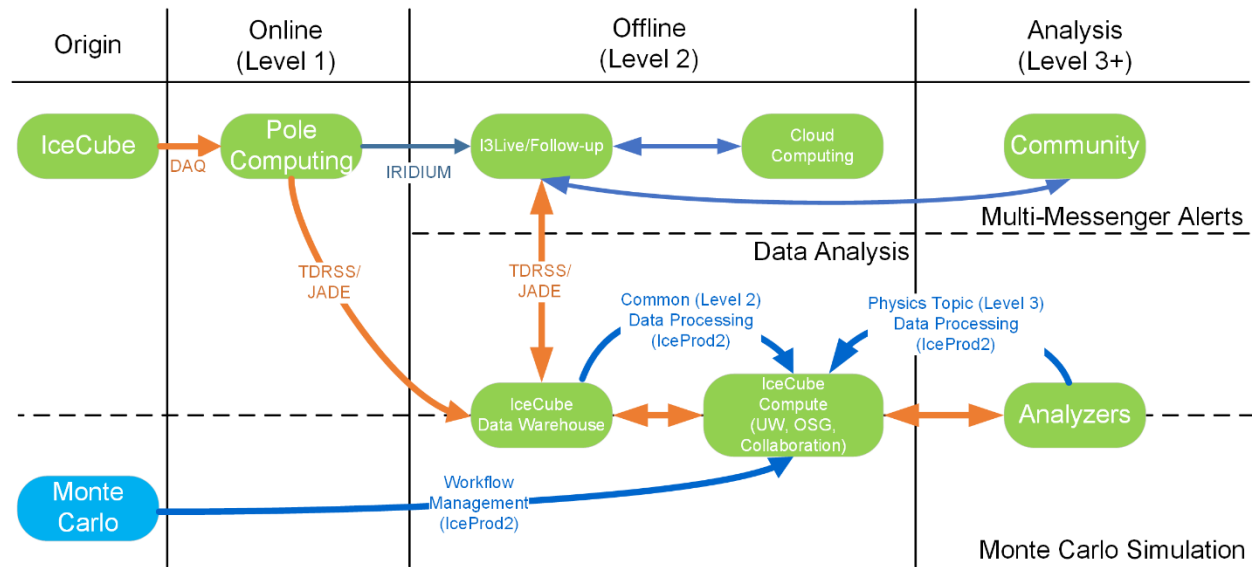


Figure 8 IceCube computing elements and data flow for experimental data and Monte Carlo simulation.

The data generated by the detector is its most precious output. One of IceCube M&O's central missions is to safeguard the data while in transit, especially from the South Pole, and stored, in user-accessible or long-term archival storage. The 11 petabyte (PB) IceCube data set will continue to grow as new data are collected, processed, and analyzed at a rate of about 1 PB per year.

Data Warehouse: The data warehouse consists of two pieces: distributed parallel file systems (DPFS) and database management systems (DBMS). The DPFS consists of a number of disk storage servers organized together using a cluster file system architecture. This provides the required performance and scalability for handling expansions and turnover efficiently and securely. System administrators experienced in managing disk enclosures, storage networks, servers, and cluster file system software maintain and operate the storage infrastructure. They ensure that data are available and that they are accessible by data processing and analysis tasks with maximum performance. The storage system administrators also handle periodic hardware and software upgrades to the storage infrastructure and take care of cluster file system operations such as accounting, quota management, and disk server load balancing. To reduce overall costs as the data set grows and accommodate the decrease unit cost of storage has flattened, we have changed our procurement policy to replace the hardware for the DPFS in five year cycles rather than an incremental approach. This reduces the unit cost of storage significantly, at the cost needing better data management.

IceCube's DBMS is dominated by non-data-analysis services, such as workflow management and detector operations. While these are typically hidden from data scientists, they are essential for operating the detector, processing data, and generating Monte Carlo simulation. Within the last year, we have centralized IceCube's DBMS. Overall, it allows for better oversight and maintenance of the various databases that the DBMS handles.

For experimental data archival, we utilize magnetic tape storage facilities located at NERSC and DESY-Zeuthen. While UW-Madison could operate a magnetic tape archive, the high purchasing and maintenance costs make using existing systems at collaboration institutions significantly more cost-effective.

The most recent DPFS hardware refresh was in mid-2018. We saw marked improvements by going to a common hardware platform and recent improvements to the clustering file system. The current file system was chosen in the early years of IceCube. While the file system will see continual improvements and

support, it is focused on high-performance computing applications and typically requires a significant human effort.

In the last 15 years, several new cluster file systems have emerged. These have been inspired by storage services available from commercial cloud computing providers. They provide significantly more flexibility, better resilience against hardware failure, and better scaling capabilities. They are also software extendable and are built with a distributed data analysis workflow in mind. In PY2 we migrated the user and analysis filesystem from Ceph to Lustre. The process was not as smooth as we had hoped with an extended two-week downtime. The experience since the transition has been trouble-free. In PY3, we migrated the experimental and simulation data filesystems from Lustre to Ceph; this transition went much more smoothly. Final relocation of existing simulation datasets that are still in use will be finished in early PY4.

In total, we dedicate 1.5 full-time system administrators across three employees to maintain the DPFS and DBMS. Besides regular maintenance tasks, e.g. replacing failed hardware, updating software, etc., this includes deploying additional hardware, monitoring, and ensuring overall performance and availability.

4.2.2 Experimental Data Management and Archival

Experimental data from the South Pole are retrieved over bandwidth- and time-limited satellite links or by transferring physical storage media during the summer season. The M&O team has developed the JADE software application to handle the data movement from the South Pole to the central data warehouse at UW–Madison and archival sites. For the satellite data transfers, the service makes use of the Iridium satellite systems for high-priority, low-volume data, e.g. realtime neutrino alerts, and the dedicated high-capacity TDRSS satellite system for the bulk of the filtered data. The unfiltered data stream are stored on two different physical media at the South Pole and shipped to the UW–Madison data center once a year during the austral summer.

To ensure integrity of all data, the software maintains checksums of all files before transfer. If data has not been transferred successfully from the South Pole, it will retransmit the files. The data will not be removed from the South Pole until data integrity has been assured, i.e. the checksums of data arriving in the data warehouse or being stored on disk at the South Pole match the initial checksum.

JADE runs on several servers to achieve higher reliability and scalability. A core software engineer will maintain the transfer and archiving software and manage the daily operations of data transfer.

Recently, we deployed new data archival software—Long-Term Archive (LTA)—that is integrated into the JADE system. LTA automatically replicates the unfiltered and filtered data to their respective archives. The filtered data is archived on magnetic tape at DESY-Zeuthen. This also allows for “local” access in Europe. Once the unfiltered data arrives at UW–Madison, it is read off the physical media, bundled, and transferred to the magnetic tape archival system at NERSC. The development of this system was funded through a separate grant from NSF (OAC-#1841479). In PY2 we used this system for the first time to archive a year of raw data. We were able to complete this task much faster than in the past using LTA.

For PY4, we are planning a secondary copy of the raw data stream stored at NERSC on the Ranch system at Texas Advanced Computing Center (TACC). We made this decision after we encountered issues during the yearly renewal of our NERSC tape allocation and in preparation for a potential reprocessing of the raw data using updated calibration. Without a secondary copy, moving the data from NERSC to other computing resources would take roughly 6 months by itself. This will allow us to cut this time in half at least and allow us to better utilize resources for the deployment of NSF-funded Leadership Class Computing Facility at TACC in the near future.

Two full-time staff members, one software engineer, and one system administrator focus on data management and archival. The software engineer focuses on maintaining JADE and LTA, including handling any error states that might occur at the South Pole that the winterovers cannot handle. The system administrator focuses on the network and storage infrastructure needed to create archives, including archiving other data sources such as the document management system.

4.2.3 Data Access and Exploration Infrastructure

The M&O IT team has been working on a catalog of all simulation and experimental data files to enable a richer metadata catalog for files, e.g. tagging files used in a data analysis, and to allow for better tracking of files, including file integrity over time. This is a first step towards an overall metadata catalog for all IceCube events. This will be needed for several projects, such as an open data service and improved data exploration infrastructure. At the current funding levels, this metadata catalog will be outside the scope of M&O funding.

With the deprecation of the Globus toolkit, IceCube will transition from using GridFTP to HTTPS-based external data access protocol in PY4. In PY2 we deployed the required new authentication and authorization services.

In total, we dedicate one full-time employee, one-half of a software engineer, and one-half of system administrator to these tasks. This includes ingesting new metadata into the file catalog as well as monitoring and maintaining the outward-facing data access infrastructure.

4.3 IceCube Open Data Services

IceCube public data will continue to be in high demand during the period covered by this proposal. The field of multi-messenger astronomy (MMA) will grow rapidly when the Vera Rubin Observatory (VRO) comes online in 2025, the O4 observing run of LIGO, Virgo, and KARGA begins in early 2023, and the beginning of operations at the Cherenkov Telescope Array (CTA). In addition to the alert data stream, IceCube will be delivering new and updated archival events. To meet this demand, the IceCube Collaboration provides public access to reconstructed neutrino event data on several levels. The broadest data set, as described in the IceCube Collaboration data-sharing policy and the data management plan, is primary event data on all events transferred from the detector over the satellite. The data sets will be provided in an open format to be usable by researchers outside of the collaboration. The open data set requires support beyond the scope of IceCube M&O at current funding levels; therefore, we will be applying for external funding for this project. Additionally, we will be investigating the possibility of hosting the open data set through commercial cloud providers.

A more selective general-purpose data set consists of high-quality reconstructed muon neutrino and muon tracks. This set enables scientists outside of the collaboration to test a broad variety of models against IceCube data. Currently, we release the general-purpose data as a tarball of text files on the IceCube website. We are investigating more feature-rich options that connect us more closely to the MMA community, such as NASA's High Energy Astrophysics Science Archive Research Center (HEASARC) [19]. Over time, future versions of these data sets will require continued support as we update reconstructions and incorporate new calibration data from the IceCube Upgrade.

Finally, there are targeted data releases associated with specific publications such as the observation of TXS 0506+056 and the high-energy starting event selection. These event selections are highly specialized to a single object or analysis, and these targeted data releases allow external researchers to reproduce and extend IceCube results. Additionally, IceCube real-time alert events are released publicly in real time as GCN

notices. We will develop an open data portal for live alert events that will incorporate follow-up data from MMA partners, similar to LIGO's GraceDB.

4.4 Cybersecurity

The information security program was developed, implemented, and is maintained to provide an organizational environment to ensure appropriate information security and levels of information-related risk. This program entails ongoing activities to address relevant policies and procedures, technology and mitigation, and training and awareness.

A risk-based approach is used to secure ICNO systems. Information systems are evaluated in terms of sensitivity of information and availability requirements of the asset. Security controls are selected and implemented to reduce risk to acceptable levels. In addition, we inherit security controls from ASC for information systems at the South Pole station, and UW security controls for information systems operated at UW-Madison.

The IceCube detector is the single most valuable asset for ICNO. As such, the primary concern is securing and maintaining the operational capability of the detector as well as day to day operations and data collection. This is followed closely by the data collected by the detector.

Access to the detector and its subsystems is restricted to IceCube personnel with a need to work on the detector itself. Remote access is limited to a small set of machines in the northern hemisphere. These machines are protected by ICNO-operated network-based firewalls in the north and south. In addition, any access to the systems on station must also pass through network firewalls and other security systems operated by ASC. Changes to station security controls are coordinated with ASC via the annual Support Information Package process.

The data collected by the detector are the foundation of all science output. It is critical to collect and preserve the observational data as they are created to avoid missing unique or rare events. To reduce the likelihood of data loss, two copies of the raw data are written to disks at Pole. These disks are shipped to UW-Madison during the austral summer. A filtered copy is written to disk at Pole, and a reduced data set, about 10%, is transferred north via satellite daily. The reduced data set is replicated daily to DESY in Germany when it reaches the north. The raw data are read from disk when they arrive at UW-Madison, where they are read and replicated to NERSC. One copy is also physically stored offline in Madison.

The science data collected and maintained are not sensitive or regulated, and indeed are eventually published. In the course of operating the center, other information is generated and stored. This information is intended for internal use only. We only generate and retain the data necessary for executing administrative processes. This information is stored separately from all computing and research systems, and uses normal IT controls to ensure the confidentiality of the data.

Where widely accepted security practices and standards are not workable, compensating controls are adopted to maintain an appropriate security level. For example, stateful, network-based firewalls have unacceptable performance impacts on large research data flows and therefore data moving machines are frequently placed outside of such protections. To mitigate the risk, a ScienceDMZ architecture [20] is applied as a compensating control to apply equivalent protections.

4.4.1 Cybersecurity Standards and Adherence

We follow standards, practices, and guidance from TrustedCI [21] that are consistent with operations of NSF Major Facilities as well as UW-Madison campus policies, and ASC policies at the station. ICNO

participates in and contributes to NSF security communities via TrustedCI and the Large Facilities security team.

4.4.2 Cybersecurity Breach Reporting Policy

ICNO maintains an incident response plan which includes escalation and notification procedures. To summarize, breaches will be reported to the appropriate parties in a timely manner in accordance with the severity. For incidents with a scope beyond the home institution, external incident response staff will be engaged immediately. For breaches which may impact resources at the South Pole station, ASC and NSF Program Officers will be notified immediately. Significant breaches will be reported to NSF Program Officers within 24 hours. UW security personnel will be notified in a manner consistent with UW reporting policy. ICNO maintains a list of security contacts for all collaborating institutions to facilitate notifications within the collaboration.

4.5 Project Year 4 Plan

Over the course of PY4 we are planning the following:

- Deploy new hardware to IceCube-owned HTC cluster
- Transition business services to Google Workspace
- Transition from GridFTP- to HTTP-based data transfer
- Copying raw data to TACC tape storage system (Ranch)

The hardware for filesystem that hosts the experimental data and simulation is reaching the end of service life, i.e. end of warranty, and needs to be replaced. We are also transitioning the underlying software that controls the distributed filesystem from Lustre to Ceph. We combined the hardware and software transition because the hardware best-suited for Ceph is different than the hardware best-suited for Lustre. The transition to Ceph will provide improved resilience and flexibility for the future.

We will be adding new hardware to the IceCube HTC cluster as most of the hardware was installed in 2015. In the last eight years there have been significant improvements in power efficiency and computational speed. The collaboration is also moving to machine learning based data analysis methods, which require different hardware than traditional data analysis methods.

Google Workspace will allow us to reduce the number of services that we run, monitor, and maintain, particularly those with significant user impact. This will also allow us to automate and generate better business rules regarding access to collaboration resources and documents.

GridFTP as a software package has been deprecated for nearly two years, which is a large cybersecurity concern. Transitioning to HTTP(S) will allow us to modernize our cybersecurity practices and improve the overall user experience.

5 Data Processing and Simulation

5.1 Data Processing

There are several levels of processing that yield increasingly specific data for different types of analyses, starting from the processing done at the South Pole (Level 1). Once the data arrives via satellite, a common processing step is applied to all data (Level 2), which includes a number of energy and directional reconstructions. The M&O team, additionally processes and filters data specific to different analyses (L3). Level 1 and Level 2 processing, including data validation, are managed centrally by M&O personnel. Individual working groups are responsible for delivering tested and validated L3 scripts and for validating the L3 output data.

A software engineer adapts data processing based on the detector configuration and required reconstruction algorithms developed by the collaboration. The software engineer is responsible for executing and monitoring tasks to make the best use of the available computing resources. Close coordination with the run coordinator ensures data-quality issues are well understood and that any poor-quality data are excluded from processing. The software engineer performs additional data validations to detect potential issues with data value and file corruption. Data processing is performed using IceProd [22, 23], an in-house workflow management and data provenance system. This system is operated as a service for the Collaboration. Replication of all the data at the DESY–Zeuthen collaborating institution is subsequently done in a timely manner.

Recent improvements to data processing and detector calibration required that we reprocess nearly 10 years of data. This reprocessing involves retrieving data from long-term archive and developing a new workflow that includes both online and offline processing.

As an added benefit of the reprocessing of data (Pass2), was the opportunity to unify the multiyear data set and to profit from improvements in our understanding of low-level DOM calibration. The reprocessing campaign started on June 1, 2017, was completed in August of 2018, and included a total of seven years of data (2010–2016).

Additional data reprocessing campaigns will be needed as further improvements in calibration and feature extraction are developed. The current refinements do not yield sufficient improvements to warrant the expense of roughly \$1,000,000 in computing time and person-hours. As a result, we are delaying such a reprocessing until deemed necessary. Improvements in calibration resulting from the IceCube Upgrade will certainly require additional reprocessing campaigns. We anticipate at least two additional reprocessing campaigns.

5.2 Monte Carlo Simulation Production

The production of Monte Carlo simulations is coordinated by M&O personnel as a service to the IceCube Collaboration. Such simulations are required for developing analysis methods to identify signal from background, for testing the performance of reconstruction algorithms, and for determining the background contamination of data analysis samples. Ideally, one would generate an order of magnitude more statistics in Monte Carlo as data. IceCube is currently far short of this goal. One of our main goals is to accelerate the workflow in order to catch up and generate at least a comparable amount of Monte Carlo statistics as data.

The Monte Carlo chain starts at generating background and signal events and ends with the common processing steps (L2). This is done to reduce the storage space needed for simulation. Some simulation is stored at intermediate steps in the chain, e.g., triggered data, in order to perform more detailed studies on effects of changes to the processing chain at various stages. Direct photon propagation is currently done on dedicated GPU hardware at several IceCube Collaboration sites and through opportunistic grid computing. The number of such resources continues to grow along with further software optimizations for GPU utilization.

In order to study the systematic effects of uncertainties in our modeling of the detector, we have implemented an approach of generating continuous variation of nuisance parameters such as, optical properties of the bulk ice, DOM efficiency and angular acceptance, etc. instead of generating multiple simulation sets for some specific choices/combinations. The details of this approach are described in [24].

The simulation production model has transitioned from a centrally managed production to a coordinated model. Each physics working group is responsible for determining their own simulation requirements and requesting allocation of computing resources. The simulation is then centrally generated.

A data scientist collects this information, including rough estimates of resource utilization (i.e., storage, CPU/GPU-hours), and discusses data set priorities with the technical leads for each analysis working group within the ICC. The data scientist provides an up-to-date status summary web page with all pending and complete requests where users can monitor the progress of their requests. A software engineer provides technical assistance to working groups and can, at the request of each working group, directly manage a given production.

The resource aggregation and workflow management of Monte Carlo simulations has transitioned from a distributed model to a centrally managed one, significantly reducing the effort needed compared to managing individual sites. Throughput has continually increased due to incorporation of an increasing number of dedicated and opportunistic resources and several code optimizations. New monitoring tools are currently being developed in order to keep track of efficiency and further optimizations. New procedures are also being implemented for allocating resources and priorities to individual simulations.

There is an expected need to periodically re-simulate most signal datasets as new improvements in software and modeling of the optical properties of the ice are developed. As a result, old-obsolete simulations are scheduled to be removed from disk in order to accommodate for storage of these new datasets.

The simulation production team regularly organizes workshops to explore better and more efficient ways to meet the simulation needs of analyzers. This includes both software improvements and new strategies as well as providing the tools to generate targeted simulations optimized for individual analyses. New strategies are being developed for dynamic simulation of systematic uncertainties in our understanding of ice properties, hole-ice, and DOM sensitivity and for determining the impacts of these on physics analyses.

5.3 Computing Resource Needs

The current utilization of resources required for the offline production is approximately 480,000 CPU-hours on the IceCube cluster at the UW–Madison data center. An effort was made in 2019 to clean up filters, reconstructions, and libraries no longer needed in offline reconstruction, resulting in a 36% reduction of CPU utilization and a comparable reduction in memory requirements.

We add roughly 100 TB of storage per year for both the Pole-filtered input data and the output data resulting from the offline production. Additional savings in storage resulted from switching to a more efficient compression in the last couple of years. L2 data are typically available one and a half weeks after data taking. Resource utilization and data volume is expected to increase in the near future with the increase in complexity of the detector due to the Upgrade and reconstruction algorithms and with the additional number of sensors deployed.

The reprocessing of Pass2 utilized 10,905,951 CPU-hours and 520 TB storage for sDST and L2 data. An additional 2,000,000 CPU-hours and 30 TB storage were required to process the Pass2 L2 data to L3. We anticipate a need for future “Pass3” reprocessing of data that will require on the order of 20 M CPU-hours of additional processing in the next couple of years as shown in the table below.

Class	CORSIKA	MuonGun	Diffuse	OscNext	IceTop	Total
CPU (years)	30000	23	1500	2112	1157	34792
GPU (years)	2400	95	680	184	0	3359

Storage (TB)	3330	10	10	60	200	3610
--------------	------	----	----	----	-----	------




Figure 9 (a) CPU (GPU) time per cosmic-ray shower as a function of energy. (b) The ratio of walltime to detector livetime indicates the number of CPU (GPU) units continuously running for 1 year needed to simulate 1 year (10 years) of detector livetime.

Simulation production requirements are primarily dominated by background simulations with CORSIKA [25] given that there is roughly a factor of 10^6 cosmic ray induced muons triggering the detector for each neutrino event (Figure 9). Background simulations for the in-ice array require roughly 30k years of CPU time and about 2.4k years of GPU time to produce and filter. This is in addition to IceTop surface array simulations and signal simulation (including systematics). As an alternative to this amount of background simulation, we can also simulate final-state muons that can be weighted according to a parametrized flux calculated from CORSIKA simulations using the same approach of MUPAGE [26] which was developed by the ANTARES Collaboration [27]. These MuonGun simulations are significantly more efficient to produce, requiring about 6M CPU-hours and comparable GPU time to simulate in order to meet our goals. These simulations need to be validated against CORSIKA, but this requires a significantly smaller data set.

Expanding access to computing resources has been a continuous effort. IceCube has been at the forefront of using novel cyberinfrastructure, including the large-scale use of GPUs. Technical debt has been accumulating over the past decade and is hindering further expansion. The newest and largest GPU-based resources across NSF, the DOE, and PRACE that feature the Power9 platform no longer fully support IceCube's software. The heterogeneity of resources will only increase in the coming years. The upcoming pre-exascale, such as NERSC's Perlmutter system, and exascale class HPC resources, such as Argonne National Lab's Aurora, will introduce new types of GPUs. Given IceCube's increasing demand for resources from Monte Carlo simulation and individual analyzers, the technical debt may affect the scientific output of IceCube substantially. Addressing this will require a concerted effort across IceCube's cyberinfrastructure teams and additional professional resources to ensure that long-term solutions can be implemented, including addressing the technical debt as well improvements needed to support the Upgrade and IceCube-Gen2.

5.4 Project Year 4 Plan

For data processing, a major PY4 goal is to implement the transition to a streamlined offline processing. This follows the new online processing implemented in PY3 that sends all events over the satellite after a software retriggering in highly compressed format. The new offline processing will combine elements of the previous L1, L2, and L3 filtering with modern event reconstructions and selections.

After the successful transition of offline processing with the IC86-2024 run start, we will start preparations for a “Pass3” reprocessing. As with the “Pass2” reprocessing described above, this will apply the latest calibrations and processing to historical data to provide a new, unified dataset for analyses. The initial phase of Pass3, which will begin in PY4, will implement the new calibration and online reprocessing.

Simulation efforts in PY4 will focus on continuing to produce specific requested datasets, rollout of the reparametrized MuonGun simulation tool, support for IceCube Upgrade simulation, and any new simulation needed to evaluate the offline processing changes.

We will continue to expand computing resources as support permits. We are evaluating the tradeoffs between adding CPU, general-purpose GPU, or ML-accelerated GPU hardware in terms of benefit to the collaboration given a certain investment.

6 Physics Software

IceCube’s physics software codebase (IceTray) covers a wide range of responsibilities and is used directly by a large fraction of the collaboration analyses, including for online filtering, real-time systems, offline data reprocessing, and offline simulation sample generation. It currently consists of over 100 projects and 1M lines of code, written in both C++ and Python. Their functions range from the core IceTray framework to user-defined simulation and reconstruction modules. Nearly all of IceCube’s more than 15 PBs of data (both archived and active) is stored in an IceTray-custom serialization format. The ongoing effort over the five years of this award will focus largely on adapting current projects to support new detector modules planned for the IceCube Upgrade and envisioned for IceCube-Gen2, ensuring the entire chain, from simulation generation through filtering, operates reliably and efficiently in a wide range of distributed environments.

The IceTray software development is organized in a full and robust IceCube organization in GitHub, with over 100 additional software repositories contributed from the collaboration. Development effort and release planning are coordinated using GitHub enterprise tools for issue tracking, pull requests, code reviews and release coordination in Kanban-style project boards. Additionally, all CI systems are implemented with a collection of self-hosted GitHub Action runners (CI) that provide rapid, automated unit testing of newly submitted software. The software group also supports several open-source software projects that are shared with the larger particle astrophysics community. Many software efforts are supported almost entirely by in-kind contributions, though some key projects will transfer to central data production after development and testing and ultimately become the responsibility of M&O. Our plan is to continue to encourage collaborators to continue to develop and maintain code used for production and analyses within our software organization.

The physics software codebase currently has roughly 4 professional FTEs of effort supporting 8 experts dedicated in the M&O software effort to its development and maintenance. Physics software projects currently fall into one of the following nine categories, with only the first group of four covered by M&O maintenance efforts, while the rest are supported by in-kind contributions from the wider collaboration:

- Core Software, Simulation, Reconstruction, Filtering

- Oscillations, Neutrino Sources, Beyond the Standard Model, Cosmic Ray, Diffuse Neutrinos

The roles and responsibilities of individual, professional software engineers on IceCube can be classified into four broad categories in order of importance: maintenance, feature development, researching future technologies, and educating colleagues. The fraction of time spent in each of the roles will vary, but work-plans for members of the software group are prioritized around these key four responsibilities. The software group has adopted a seasonal release cycle, producing feature releases four times per year, with smaller point-level releases as needed by data and simulation production teams.

6.1 Core Software

The bulk of the projects are written in C++ and Python, which—over the last decade—have been rapidly evolving, where new standards are released every three years. The focus of the software group is to incorporate modern C++ and Python features that improves its scaling performance. The core software projects generally fall into one of the following categories:

- Framework - C++ framework IceTray including python bindings.
- DevOps - Responsible for third-party tool detection, compilation, documentation builds, and CI/CD.
- Serialization - Responsible for the bulk of I/O on on-disk format. (adapted from [28])
- Data Structures – Specialized in-memory and on-disk structures used by downstream filters and analysis modules.
- Data Visualization - Visualization tool that generates high-quality graphics for publications [29].
- DAQ Payload Tools - Tools to read and handle custom IceCube DAQ payloads.
- Waveform Tools - Projects to calibrate and unfold digitized waveforms.
- Analysis Formats - Generic framework to support the conversion of native IceTray classes to popular analysis formats.

Currently, there are 2 FTEs, split between 4 people on core M&O efforts, focused on the maintenance and development of the above projects. The software build system is continually evolving to support tools needed for physics analyses. Deep Learning techniques are proving valuable in many areas of IceCube analysis. It will be necessary to provide support for popular ML analysis formats, such as Apache Arrow, PyTORCH, or Tensorflow based workflows. Work planned for PY4 and PY5 will continue to modernize this software, expand support to modern OS releases and workflows, and integrate the IceCube Upgrade into core software systems.

6.2 Simulation Software

The simulation projects contain various IceTray modules and services necessary for generating both neutrino signal and cosmic-ray background Monte Carlo simulation data samples. IceCube’s simulation chain has to perform equally well for neutrino oscillations at $O(10 \text{ GeV})$ energy as for cosmogenic neutrinos at $O(10 \text{ PeV})$ energy. Efficient resource usage on all available hardware (e.g. x86, Power9, ARM, NVIDIA, AMD, etc...) will continue be a development focus of the core simulation group. Core support, which all collaborators rely on, is needed in at least two main areas: GPU/heterogeneous programming and detector simulation.

- Cosmic Ray Signal - Full hadronic shower generation with CORSIKA for IceTop only.
- Hadronic Shower Background - In-ice background muon and atmospheric neutrino generation with CORSIKA.

- Fast Background - Parameterized single-muon, in-ice background, sacrificing systematics for statistics.
- Neutrino Signal - Injection of neutrino signal, including standard model and BSM oscillations.
- Exotic Signal - Monopole and WIMP generation for the BSM physics working group.
- Photon Propagation - Propagation of photons through the complex ice structures utilizing GPUs.
- Lepton Propagation - Propagation of leptons, which includes energy losses and stochastic generation, the main contributors to Cherenkov radiation and therefore detector signal.
- Detector Simulation - Simulation of the PMT, DOM Mainboard, and DAQ trigger.
- Data Structures - Specialized structures to store Monte Carlo truth information used by downstream analysis modules. New structures and services will need to be developed and maintained to support Upgrade modules.

Of the above projects only the hadronic shower background, neutrino signal, photon propagation, detector simulation, and data structures are supported by M&O, at a level of 1.6 FTE fractionally supporting 5 people. The rest are supported by in-kind contributions from collaborating institutions. Producing sufficient background to meet the analysis needs of all working groups remains one of the most significant challenges for the simulation group. The larger software group also explores the use of GPU-accelerated programming in areas, such as detector simulation, waveform deconvolution, reconstruction, and deep learning algorithms. Work planned for PY4 and PY5 will continue to improve this software, addressing known issues, and expand the simulation to integrate the IceCube Upgrade into simulation software programs.

6.3 Reconstruction Software

The reconstruction projects consist of IceTray modules used from low-level online filtering at the pole to high-level analysis. The reconstruction software effort is a largely in-kin support one, with only 0.5 FTE of dedicated M&O core effort available to support the core reconstruction frameworks. Below are the broad categories that most reconstruction projects fall into:

- First-guess Seeds - Simple first-guess algorithms that typically serve as seeds to more resource intensive reconstructions.
- Framework - Reconstruction framework within IceTray that manages the interaction between first guess hypotheses, minimizers, and resource intensive algorithms.
- Reconstructions - Track and cascade algorithms used both in low-level filtering and analysis. Ensuring all reconstructions are adapted to Upgrade modules is going to be critical for future analyses.
- Tools - Tools that estimate photon arrival times and amplitudes will need to be adapted to support upgrade modules.
- Data Structures - Specialized structures that store reconstruction results used by downstream filters and analysis modules.

The reconstruction software effort is a largely in-kin support one, with only 0.5 FTE of dedicated M&O core effort available to support the core reconstruction frameworks. Work planned for PY4 and PY5 will focus on support of these software packages, addressing known issues, and inclusion of the IceCube Upgrade into reconstruction algorithms.

6.4 Integration of the IceCube Upgrade

The level of effort from core developers and in-kind contributed efforts has been sufficient over the first 3 years of the M&O grant to meet most key operational milestones and performance metrics for the IceCube software effort. However, over this award, first-guess algorithms, core reconstructions, and data structures will need to be updated to support data from the planned Upgrade modules. Support efforts by M&O experts will ensure efficient utilization of distributed resources and manage the increased complexity of joint reconstructions. Core software must be adapted to support the IceCube Upgrade development of new data structures, DAQ payloads, waveform tools, and database structures. Additionally, simulation of the PMT response and electronics is currently the second largest consumer of CPU resources in our simulation chain. The ability to predict resource usage, to allow for efficient scheduling in a distributed environment, and remains a critical task as the number, type, complexity, and density of optical modules increase during the Upgrade. And finally, experience has shown that difficulties from performing joint reconstructions between in-ice/IceTop and in-ice/DeepCore have complicated analyses. Avoiding this in the future after the deployment of the Upgrade modules will ensure timely delivery of analysis results.

Limited staffing has complicated addressing the large increase in software issues and tasks from the integration of the IceCube Upgrade into the IceCube M&O program, both for core software tasks, and support for simulation and analysis of the upcoming IceCube Upgrade has lagged our planned efforts. For in-kind contributions, the delays in the Upgrade deployment have caused planned efforts for Upgrade tasks to be reassigned to meet graduation and career goals of the student and post-doc workforce. Core software efforts have been slowly reducing in level due to pressures from inflation. Addressing these lags in software effort remains a critical task for PY4 and beyond for the M&O management team.

7 Detector Calibration

Calibration of the IceCube sensors and measurement of the properties of the South Pole ice enables the production of high-quality science products from the raw charge and time data collected by the IceCube DOMs. Routine calibration tasks are handled through the detector operations group. A dedicated calibration working group is charged with developing and disseminating models of the bulk and hole ice, measuring the efficiency of the DOMs in the lab and in situ, measuring the SPE behavior in DOMs and organizing LED flasher runs. There are 2.38 FTEs dedicated to calibration in the project, with all additional effort contributed in kind.

7.1 Sensor Calibration

The primary calibration routine, called DOMCal, is performed every month in IceTop and every year in the in-ice array, in order to measure the PMT gain as a function of high voltage, discriminator threshold settings, and other calibration constants. Additionally, the DOM digitizer baselines and SPE distributions are measured every year. These constants are used by the PnF system to translate the raw waveform data into physical units. The SPE distributions (a proxy for PMT gain) and the digitizer baselines are monitored continuously by the detector monitoring system to ensure the stability of these constants. The collaboration recently published a new method of extracting the SPE charge distribution using a deconvolution of the multiple-photoelectron charge distribution [31]. Calibration constants based on this method were originally derived from data from seasons 2011–2017. Currently, the method is being applied to data from 2016 to 2021 with preliminary results indicating that changes are small enough that no update to the calibration constants will be required. The SPE charge distribution fit code is also being updated with improved pulse selection and fitting for the FADC. During PY4 this work will be finalized with additional plans to automate the process in the future and investigate an alternative description of the very low charge region. We will also continue our effort to uniformly apply these calibration constants in data taking and simulation, as part

of the plans to update filtering and processing. There are 0.08 FTEs dedicated to supporting DOMCal tasks and baseline and SPE calibration.

7.2 Measurement of Ice Properties

Measuring the optical properties of the ice is critical to accurate reconstructions of the energy and direction of neutrino events. The ice remains a major source of systematic uncertainty in IceCube science analyses. The collaboration continuously improves the ice model with data from the LED flasher calibration runs. After the development, release for simulation and publication of a major new ice model [32] describing the ice anisotropy effect through the cumulative light deflection caused by the birefringent microstructure of the glacial ice, attention is now shifting to improvements in modeling of the ice layer undulations (called tilt) and an update to the detector geometry. For PY4 we will continue to investigate the sources of the remaining disagreement between data and simulation. For the geometry several approaches using both LED flashers as well as muon events are being pursued. While the potential gains are uncertain, we plan to at least release preliminary models for testing by the simulation and analysis groups. There are 0.65 FTEs dedicated to supporting the implementation of the ice model and supporting use of new ice models in event reconstruction.

7.3 IceCube Upgrade

Significantly improving the detector calibration is one of the primary goals of the IceCube Upgrade. The Upgrade will include an array of new calibration devices including flasher LEDs, cameras, beamed and isotropic LED light sources, and acoustic sensors. Closer spacing of the Upgrade strings will facilitate measurements on a baseline comparable to or less than the optical scattering length in ice. For the improvements in the ice model and detector characterization, we expect to achieve DOM optical efficiency determination in situ to better than 3%; and we expect to reduce uncertainties of angular acceptance of IceCube DOMs due to refrozen hole ice by a factor of at least 2. Additionally, we expect to resolve remaining uncertainties around the new anisotropy model. The updated calibration constants will improve angular resolution of existing IceCube data by a factor of 4 in the cascade channel and by a factor of 2 in the muon channel for commensurate improvement in point source sensitivities and reduction of false alarm rates by a factor of 4. We expect to double the astrophysical cascade event rate to 20 per year and achieve 3σ observation of cosmic tau neutrinos after 12 years. The Upgrade project supports the deployment of Upgrade hardware, collection of calibration data and development of a database to contain new calibration constants. ICNO operations, in turn, will support the implementation of updated calibrations into the simulation and reconstruction software. PY4 will in particular be focused on operations planning for Upgrade calibration modules both during and after deployment, mDOM FAT, and developing module software at the Northern Test System. There are 1.65 FTEs dedicated to supporting the data from LEDs in the Upgrade, including the beamed LED module ("PencilBeam") which will be able to directly test the light deflection intrinsic to the newly proposed birefringence-based mechanism for the ice optical anisotropy. All remaining calibration effort in the Upgrade following the end of Upgrade construction is contributed in-kind.

8 Risk Management Plan and Risk Register

This plan describes how the IceCube Maintenance and Operations team manages and mitigates risk to the existing detector and its data. This plan governs all IceCube M&O Collaborators, including UW-Madison personnel who control budgets and schedules.

8.1 Risk Registry

We maintain a risk registry of approximately the top ten risks to the ongoing ICNO M&O. Personnel risks during South Pole deployment are handled separately, with required general and specialized safety training prior to travel to Antarctica. This minimizes the possibility of a major injury that would impact the ICNO operations. Safety is the number one priority for those deploying, and the excellent track record to date is evidence of effectiveness of the safety training.

IceCube operations risk management attempts to leverage the technical experience from the IceCube construction: everyone is responsible for identifying risks, and once a risk is identified someone is assigned formal responsibility for managing it. This risk registry is an evolved (and simplified) version of the project construction risk registries of the IceCube Gen1 construction effort and the IceCube Upgrade construction project. Risks are categorized as per the Research Infrastructure Guide:

- Technical
- External
- Organizational
- Project Management

8.2 Risk and Opportunity Review and Management Process

The risk register is reviewed and updated annually, or what specific circumstances require a new look at the risks. (For example, with the current pandemic, the risks have been examined with an eye towards the difficulty of South Pole access during the 2020-21 Austral Summer.) In addition, individual risks are actively monitored, and the risk status is reported as part of the quarterly reporting process.

The project Quality Manager is responsible for the risk register and works with subject matter experts and the project managers to adjust the cost and risk scoring to an equivalent basis. Risks are manually reviewed in this manner to ensure that entries developed by many personnel across the project are captured and assessed with similar leveling, allowing for project-wide integrated assessment of risk and exposure.

Risk management also include opportunities and the project is actively pursuing identification of additional funding from partners as well as cost-saving opportunities such as identifying alternate, more cost-effective sources of major capital subsystems, e.g. computer CPU and GPU hours. These opportunities are captured in the risk register alongside risks and are used in an equivalent manner to help the project management assess scope as the project evolves. Opportunities are identified actively and will be utilized in order to reduce overall project risk.

Management of risks and opportunities forms an integral part of the Project Execution Plan and is governed by the guidelines of NSF Major Facilities Guide. As such, risk and opportunity review is a necessary component of yearly detail planning and contingency planning. In general, risks might have a horizon date after which they are no longer risks, though the primary risks to IceCube data-taking are not of this form, and remain constant (risk retirement date is shown as Never) throughout the course of the M&O cycle.

8.3 Risk Classification

Risks are classified by their probability of occurrence and by their impact to the project as tabulated below. The risk score is then determined from the risk scoring table.

The risk probability tables, impact assessment levels, and the matrix of probability and impact are all taken directly from the RIG (figures 6.2.7-3, 6.2.7-4, and 6.2.7-5). They are not reproduced here.

Risk mitigation strategies are then defined, and a post-mitigated risk is derived with new impacts and probabilities assigned. This gives a qualitative estimate of the associated degree of risk so that major risks can be identified easily. To quantitatively assess an equivalent cost of the risk, the post-mitigated risk is assigned a total cost exposure, i.e. what amount of money would need to be spent in the event of risk occurrence. The resulting risk exposures do not have calculated costs associated with them since in a project construction those costs would hook directly to the project contingency, and no contingency is planned into the M&O budgets. Some of the risks, if realized, would likely require direct cooperation with the NSF and/or the host institution to manage.

9 Safety Standards

The Environmental, Health, and Safety (EH&S) specific standards followed by the M&O program are listed in Tables 6 and 7. They form the foundation for the IceCube M&O EH&S program which adheres to the procedures, policies, and training in the listed standards. The IceCube M&O Project has an internal audit program that assesses compliance to the requirements of the respective standards. We also work closely with ASC counterparts to refine and assess the IceCube M&O EH&S program, drawing on their expertise and experience particularly with the safety programs adapted to working in the extreme Antarctic conditions.

Title and Part	Section(s)	Description
29 CFR 1910	.38, .120	Emergency response
	.95	Hearing conservation
	.122, .241-244	Tool safety
	.132-133, .138	Personal protective equipment
	.146	Confined space entry
	.147	Control of hazardous energy (lockout/tagout)
	.151	First aid
	.178	Forklift operations
	.179	Overhead and gantry cranes
	.180	Mobile cranes
	.184 .301-308 .1200	Slings Electrical safety Hazardous chemicals
29 CFR 1910 Subpart L		Fire Protection
29 CFR 1926	.96, .101-103	Personal protective equipment
	.50	First aid
	.500-503	Fall protection
	.550	Mobile cranes
49 CFR Subchapter C		Hazardous Materials Regulations

Code	Description
University of Wisconsin Hazard Communication Program	
University of Wisconsin Laboratory Safety Guide	
ANSI Z117.1	Safety requirements for confined spaces
Executive Order 13043 of Apr. 16, 1997	Increasing seat belt use in the United States
Executive Order 12196 of Feb. 26, 1980	Occupational safety and health programs for federal employees and contractors
NASA Standard 8719.13A	Software safety
NEC, MIL-STD 454	Standard General Requirements for Electronic Equipment
NFPA fire standards, codes, and appendices	
NFPA 101 Life Safety Code Handbook	
NFPA 70 NEC Handbook	
NIOSH Publication No. 87-113	A Guide to Safety in Confined Spaces
OSHA Technical Manual, Sec. 10, Chap 3	Pressure vessel guidelines
International Boiler and Pressure Vessel Code	2004 Edition
NPD 8710.5	NASA policy for pressure vessels and pressurized systems
SafeStart Training	

10 Acronyms and Definitions

ADC	Analog-to-digital converter chip
AMANDA	Antarctic Muon and Neutrino Detection Array
AMON	Astrophysical Multimessenger Observatory Network
ATWD	Analog transient wave digitizer
Condor	UW–Madison workload management system for compute-intensive jobs
CF	Common Fund
CTSC	Center for Trustworthy Scientific Cyberinfrastructure
DACS	NSF Division of Acquisition and Cooperative Support
DAQ	Data acquisition system
DOM	Digital optical module
DOMCal	DOM in situ self-calibration system
DOMHub	Surface cable terminus with readout electronics and low-level data acquisition function
DOR	DOM readout electronics PCI card
DSI	Data Systems International
E&O	Education and outreach
EMI	Electromagnetic interference
GCN	Gamma-ray coordinates network
GPU	Graphical processing units
GridFTP	An extension of the standard file transfer protocol (FTP) for use with Grid computing
HPC	High-performance computing
HPSS	High performance storage system
HSM	Hierarchical storage management
HTC	High-throughput computing
I3MS	IceCube messaging system
ICB	IceCube Collaboration Board, the entity that guides and governs the scientific activities
IceCube Live	The system that integrates control of all of the detector’s critical subsystems
IceProd	IceCube simulation production custom-made software

IceSim	IceCube simulation software package tools
IceTray	IceCube core analysis software framework is part of the IceCube core software library
ICL	IceCube Laboratory (South Pole)
IOFG	International Oversight and Finance Group
JADE	Java archival and data exchange
LED	Light emitting diode
M&O	Management and operations
M&OP	Management & Operations Plan
mDFL	Mobile/modular dark freezer lab
MoU	Memorandum of Understanding, between UW–Madison and all collaborating institutions
MPS	NSF Directorate for Mathematical & Physical Sciences
MREFC	Major Research Equipment & Facilities Construction
MSPS	Mega-samples per second
NTS	Northern Test System, IceCube Upgrade extension of SPTS
OPP	NSF Office of Polar Programs
OSG	Open Science Grid
PA	NSF Particle Astrophysics Program
PCTS	Physical Sciences Laboratory cable test system
PHY	NSF Division of Physics
Physics working group (WG)	Physics working groups perform high-level analysis and develop specific analysis tools
PLR	NSF Division of Polar Programs
PMT	Photomultiplier tube
PnF	Processing and filtering
SAC	Science Advisory Committee
SCAP	IceCube Software & Computing Advisory Panel
Science DMZ	A secure computer subnetwork designed for high-volume data transfers
SIP	Support Information Package
SNEWS	Supernova Early Warning System network
SNDAQ	Supernova data acquisition
SPS	South Pole System (at the South Pole)

SPTR	IceCube dedicated high-capacity South Pole TDRS relay system.
SPTS	South Pole Test System (at UW Madison)
TDRS	A Tracking and Data Relay Satellite is type of a satellite that forms part of the TDRSS
TDRSS	The Tracking and Data Relay Satellite System is a network of communications satellites
TFT Board	Trigger Filter and Transmission Board
TS	Test statistic
UPS	Uninterruptible power supply
USAP	United States Antarctic Program
UW	University of Wisconsin–Madison, host institution of the IceCube Collaboration
VCR	Vice Chancellor for Research at UW–Madison
WBS	Work breakdown structure
WIMPs	Weakly interacting massive dark matter particles
WIPAC	Wisconsin IceCube Particle Astrophysics Center (former IRC)
XSEDE	Extreme Science and Engineering

References Cited

- [1] M.G. Aartsen et al. Time-integrated Neutrino Source Searches with 10 years of IceCube Data. *Phys. Rev. Lett.*, 124(5):051103, 2020, 1910.08488.
- [2] M.G. Aartsen et al. Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A. *Science*, 361(6398):eaat1378, 2018.
- [3] M.G. Aartsen et al. (IceCube, Fermi, MAGIC, Agile, ASAS-SN, HAWC, HESS, INTEGRAL, KANATA, KISO, KAPTEYN, Liverpool Telescope, Subaru, SWIFT, VERITAS, VLA collaborations). Multimessenger Observations of a Flaring Blazar Coincident with the High-Energy IceCube Neutrino IceCube-170922A. *Science* 361(6398), 2018, 147.
- [4] M.G. Aartsen et al (IceCube Collaboration). Evidence for Neutrino Emission from the Nearby Active Galaxy NGC 1068. *Science* 378(6619), 2022, 538.
- [5] M.G. Aartsen et al (IceCube Collaboration). Observation of High-Energy Neutrinos from the Galactic Plane. *Science* 380(6652), 2023, 1338.
- [6] M.G. Aartsen et al. Measurement of Atmospheric Neutrino Oscillations at 6-56 GeV with IceCube DeepCore, *Phys. Rev. Lett.* 120, 2018, 071801.
- [7] M.G. Aartsen et al. Measurement of Atmospheric Tau Neutrino Appearance with IceCube DeepCore, *Phys. Rev. D*, 99, 2019, 032007.
- [8] M. G. Aartsen et al. Combined sensitivity to the neutrino mass ordering with JUNO, the IceCube Upgrade, and PINGU. *Phys. Rev. D* 101, 2020, 032006.
- [9] R. Abbasi et al. The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping. *Nucl. Instrum. Meth. A*, 601:294–316, 2009, 0810.4930.
- [10] M.G. Aartsen et al. The IceCube Neutrino Observatory: Instrumentation and Online Systems. *JINST*, 12(03):P03012, 2017, 1612.05093.
- [11] M. G. Aartsen et al. The IceCube Neutrino Observatory - Contributions to ICRC 2017 Part II: Properties of the Atmospheric and Astrophysical Neutrino Flux. 2017, 1710.01191.
- [12] M.G. Aartsen et al. Measurement of Atmospheric Tau Neutrino Appearance with IceCube DeepCore. *Phys. Rev. D*, 99(3):032007, 2019, 1901.05366.
- [13] Pietro Antonioli, Richard Tresch Fienberg, Fabrice Fleuret, Yoshiyuki Fukuda, Walter Fulgione, Alec Habig, Jaret Heise, Arthur B McDonald, Corrinne Mills, Toshio Namba, et al. SNEWS: The supernova early warning system. *New Journal of Physics*, 6(1):114, 2004.
- [14] GCN. *GCN*, 2020 (accessed April 21, 2020).
- [15] P. Allison, et al., and ARA Collaboration. Performance of two Askaryan Radio Array stations and first results in the search for ultrahigh energy neutrinos. *Physical Review D*, 93(8):082003, April 2016, 1507.08991.
- [16] I. Sfiligoi, D. Schultz, B. Riedel, F. Wuerthwein, S. Barnet, and V. Brik. Demonstrating a Pre-Exascale, Cost-Effective Multi-Cloud Environment for Scientific Computing. *to be published in proceedings of PEARC'20*, page arXiv:2004.09492, April 2020, 2004.09492.

- [17] Igor Sfiligoi, Frank Wuerthwein, Benedikt Riedel, and David Schultz. Running a Pre-Exascale, Geographically Distributed, Multi-Cloud Scientific Simulation. *to be published in Proceedings of ISC High Performance 2020*, page arXiv:2002.06667, February 2020, 2002.06667.
- [18] HTCondor. *HTCondor - High Throughput Computing*, 2020 (accessed April 21, 2020).
- [19] HEASARC. *High Energy Astrophysics Science Archive Research Center*, 2020 (accessed April 21, 2020).
- [20] Eli Dart, Lauren Rotman, Brian Tierney, Mary Hester, and Jason Zurawski. The science dmz: A network design pattern for data-intensive science. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC '13, New York, NY, USA, 2013. Association for Computing Machinery.
- [21] TrustedCI. *TrustedCI*, 2020 (accessed April 21, 2020).
- [22] D Schultz. IceProd 2: A next generation data analysis framework for the IceCube neutrino observatory. *Journal of Physics: Conference Series*, 664(6):062056, dec 2015.
- [23] M.G. Aartsen et al. The IceProd framework: Distributed data processing for the IceCube neutrino observatory. *Journal of Parallel and Distributed Computing*, 75:198 – 211, 2015.
- [24] M.G. Aartsen *et al.* Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube. *JCAP10(2019)048* DOI 10.1088/1475-7516/2019/10/048
- [25] D. Heck, G. Schatz, T. Thouw, J. Knapp, and J. N. Capdevielle. CORSIKA: A Monte Carlo code to simulate extensive air showers. 1998. FZKA-6019.
- [26] G. Carminati, M. Bazzotti, S. Biagi, S. Cecchini, T. Chiarusi, A. Margiotta, M. Sioli, and M. Spurio. MUPAGE: a fast atmospheric MUon GEnerator for neutrino telescopes based on Parametric formulas. 7 2009, 0907.5563.
- [27] M. Ageron et al. ANTARES: The first undersea neutrino telescope. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 656(1):11 – 38, 2011.
- [28] Robert Ramey. *Boost.Serialization*, 2020.
- [29] The IceCube Collaboration. *Science Cover*, 2013.
- [30] Rapids.ai. *Rapids Open GPU Data Science*, 2020 (accessed April 21, 2020).
- [31] M.G. Aartsen et al. In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes. 2 2020, 2002.00997.
- [32] R. Abbasi et al. In-situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory. *The Cryosphere* , <https://doi.org/10.5194/tc-2022-174>