1 Introduction

This proposal requests funding for management and operations of the IceCube Neutrino Observatory, to ensure that it continues taking unique, high-quality data, and has the capacity to develop and implement necessary hardware and software upgrades to capitalize on the science and discovery potential of the facility. The management and operations (M&O) core team will work with established partners and build upon the experience gained from over eight highly successful years safely and cost-effectively managing the IceCube Neutrino Observatory (ICNO). Performance of the ICNO already exceeds its original design and operation goals, yet there are still opportunities for improvements: decreasing time from data collection to analysis and publication; enhancements of the real-time alerting system; improvements to the online system; gains in access to additional computing resources; and mitigation strategies to deal with the increasing snow coverage of the IceTop array.

The wide spectrum of IceCube science is made possible by dedicated efforts of the IceCube M&O personnel. The distributed model for analyzing IceCube data originates with investigators in the collaborating institutions and is channeled through designated working groups into results, by consensus, submitted for publication. The working groups cover a broad range of topics, from detector calibration to physics beyond the Standard Model. They closely collaborate with IceCube M&O personnel in order to design the selections (filters) and quality cuts that are applied to the raw data in order for the M&O group to deliver science-ready data of interest. They may also request simulations specialized beyond the global detector simulation that are focused on their specific targeted issues.

The close collaboration between M&O and IceCube science personnel has resulted in steady improvements in almost every aspect of detector performance, including energy and angular resolution, the efficiency of data filters, the accuracy of detector simulations, and the stability of the software for acquiring and analyzing the data. A significant part of this process is driven by our continued efforts to understand the optics of the ice, with steadily improving precision made possible in part by utilizing GPU clusters.

Many of the striking IceCube results reflect the synergistic collaboration between M&O staff and IceCube scientists, of which the largest group in the collaboration is co-located with M&O personnel at the University of Wisconsin–Madison. This synergy has resulted in an excellent performance of the instrument over a wide dynamic range. IceCube neutrino analyses span six orders of magnitude in energy, including the discovery of cosmic neutrinos, some reaching energies close to 10 PeV (see Figure 1), and the measurement of the oscillation of atmospheric neutrinos in a previously unexplored energy range from 10 to 60 GeV.

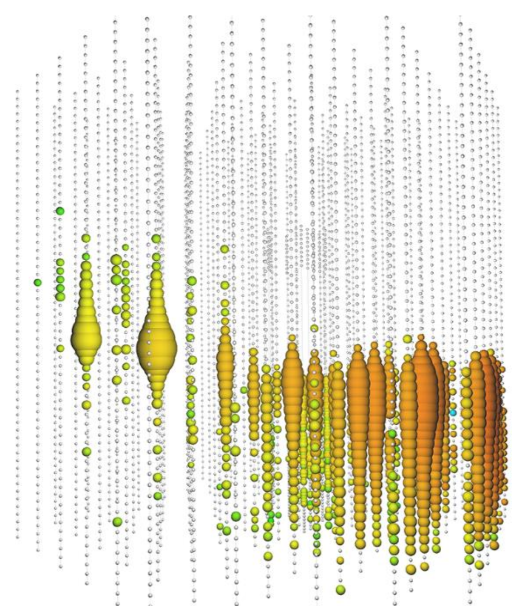

Figure 1: This example event is a secondary muon track depositing 2.6 PeV inside IceCube, detected on June 11, 2014.

It is important to recall that IceCube was primarily designed as a discovery instrument, and from this point of view it has also obtained world-leading results due to ongoing performance improvements. IceCube has set world-best limits on the interaction cross section of dark matter particles with ordinary matter for a number of leading theoretical speculations. It has improved the limits on sterile neutrinos by one order of magnitude over accelerator searches in the parameter region of the LSND [1] evidence, basically ruling out the existence of eV-energy sterile neutrinos. These are only two examples of a long list of searches for physics beyond the Standard Model, but for anything new arising in neutrino physics, it is likely IceCube will have something to say about it—not just by capitalizing on the current detector but also by designing novel ways to use this unusual instrument in close collaboration between the core M&O group and the scientists.

1.1 The IceCube Neutrino Observatory: an Overview

IceCube uses Cherenkov light, emitted by charged particles moving through the ice at super-luminal speed, to realize the enormous detection volume required for observing astrophysical neutrinos. Ice is an abundant resource, naturally occurring in glacial sheets 3-km thick at the South Pole. Cosmic-ray muons created in cosmic-ray interactions high in the atmosphere above Antarctica penetrate the deep ice and, along with cascades of neutrino-induced muons and other high-energy particles, emit Cherenkov light at the level of 100,000 visible photons per GeV of shower energy. Such quantities of photons, coupled with long optical attenuation lengths in South Pole ice and large-area photomultipliers (PMTs), makes it possible to instrument cubic kilometers of ice with a wide spacing of detectors.

The ICNO encompasses the physical detector along with the corresponding data and computing infrastructure, both described in further detail later in this section. Operations are headquartered at the Wisconsin IceCube Particle Astrophysics Center (WIPAC), a research center of UW–Madison that supports dedicated staff responsible primarily for managing the ICNO. IceCube activities are supported by a Cooperative Agreement between the National Science Foundation and UW–Madison that provides core funding for operations at WIPAC and funding for several research groups in the US via subawards. These subawards and additional in-kind and cash contributions from IceCube collaborators support core M&O tasks, as outlined later in Section 3.3.1.

1.1.1 The IceCube Detector

After seven seasons at the South Pole, beginning in 2004 and ending in 2011, 5160 optical sensors were buried in the ice along 86 deep vertical lines and 324 more sensors deployed in 162 tanks at the surface, forming the IceCube deep-ice and IceTop detector arrays. This achievement required 28,000 person-days and the transportation of 4.7 million pounds of cargo to the South Pole during a time of simultaneous intense activities to modernize the South Pole station and construct other large-scale projects at the site. The construction of IceCube, on schedule and within budget, and its science operation are widely recognized as an enormous success, garnering special recognition by the British journal *Physics World* who awarded the project their 2013 Breakthrough of the Year for the discovery of cosmic neutrinos. The digital optical sensors, despite their remote deployment in an inhospitable environment, have proven their sound design with high post-deployment survival rates of over 98% and an overall detector online availability for recording astrophysical events, commonly referred to as detector "uptime," of 99.8%. A schematic representation of the detector array facility is shown in the left side of Figure 2.

1.1.1.1 The Digital Optical Module

The basic detection element used throughout IceCube is the digital optical module, or DOM (see right side of Figure 2). Encased in a ½" thick glass pressure sphere for protection against the extreme pressure at depth, the DOM features a 10" PMT, embedded high-voltage generation, a flasher calibration board, and a mainboard containing the analog and digital processing circuitry for PMT pulses. The mainboard also hosts an FPGA and embedded processor to control the conversion process, buffer the digital data, and transmit the data to computing at the surface. Digital waveform signal capture is exploited by advanced deconvolution techniques developed by the IceCube Collaboration to reconstruct the complicated pattern of photon arrival times of high-energy cascades in the ice. However, most PMT pulses result from single incident photons and do not require transmission of large digital waveform packets. In these cases, salient pulse features, the leading edge times and amplitudes, are extracted from the digital waveform in the DOM acquisition buffer and compressed to a handful of bytes for storage and transmission.

1.1.1.2 Deep-ice Array

DOMs were deployed along a copper cable lowered into each hole. Wire bundles within the cable carry power and digital data between the in-ice DOMs and the surface. There are two further subdivisions of the deep-ice array: 80 strings with mean spacing of 125 meters each carry 60 DOMs spaced 17 meters apart vertically, which span the entire cubic kilometer volume of instrumented ice, and 6 more densely spaced strings, called DeepCore, deployed in the central core of the array, to achieve superior performance at lower energies through increased photodetector density and by using the surrounding strings to veto backgrounds. The DeepCore strings are arrayed with interstring spacing of 75 meters and inter-DOM spacing on each string of 7 meters and are concentrated at depths greater than 2000 meters where the ice is extremely clear.

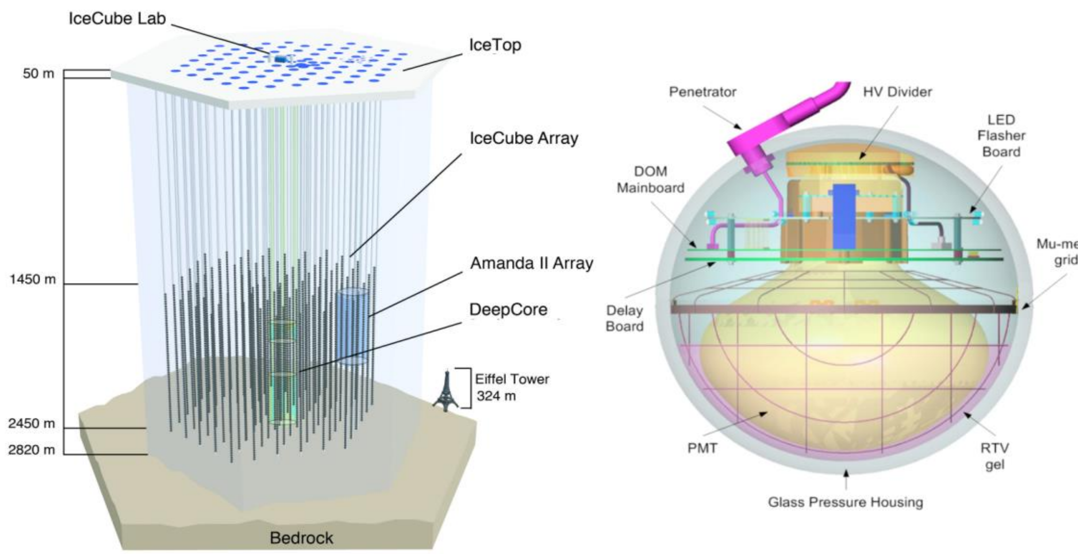

Figure 2: Schematic of the IceCube detector (left) and the digital optical module, or DOM (right).

Particle tracks in the deep ice trigger the array and are reconstructed based on the time-evolved pattern of Cherenkov photon hits in the DOMs. Cosmic-ray muons and neutrino-induced electrons, muons, and tau leptons above several tens of GeV satisfy the trigger requirement, which induces the assembly of array data into event structures. Constant monitoring of the background counting rate of DOMs in the deep-ice array also allows searching for transient bursts of low-energy (10 MeV) neutrinos liberated during stellar collapse events in our galaxy.

1.1.1.3 IceTop

The surface array, IceTop, also utilizes Cherenkov radiation as the fundamental detection strategy. In this case, however, the DOMs are embedded in tanks containing clear ice, deployed at the surface. Together with the deep-ice array, IceTop forms a large 3-D cosmic-ray air shower detector with sensitivity to primary cosmic ray composition in the energy range above 1 PeV. In addition, the surface array is utilized to veto backgrounds from cosmic-ray induced activity in the deep-ice array, which enhances the capabilities of IceCube for all-sky detection of high-energy extraterrestrial neutrinos.

1.1.2 Central Electronics and Data Processing System

Digital sensor array readout is done under the supervision of the IceCube surface data acquisition system (DAQ); see Figure 3. The DOMs are themselves data acquisition units, triggering and acquiring PMT signals autonomously. Aggregation of data at the array level begins in the IceCube Laboratory (ICL). A network of 100 custom readout DOMHubs collect data from the array at a rate of 150 MB/s. The aggregated stream is examined by trigger algorithms and, when patterns of interest appear in the stream, data surrounding the trigger time is extracted, built into an event structure, and written to disk. These triggers occur at a rate of approximately 2.7 kHz and are dominated by cosmic muon-induced events.

Here the data is split into two independent paths. First, the disk files produced are written to archival disk storage, awaiting physical transfer north once per year. In addition, the data is sent to an online compute farm of 22 nodes based in the ICL for near-real-time processing, event reconstruction, and filtering. Neutrino candidates and other event signatures of interest such as background count rate increases associated with potential supernova explosions, are identified within minutes, and notifications are dispatched to other astrophysical observatories worldwide via the Iridium satellite system with 24/7 link availability to the South Pole. Data acquisition, processing, and storage are supported by the commodity computing and network infrastructure referred to as SPS, for South Pole System. A mirror test system, SPTS for South Pole Test System, is maintained at UW–Madison to evaluate test deployments of software and hardware before they go "live" on the real detector system.

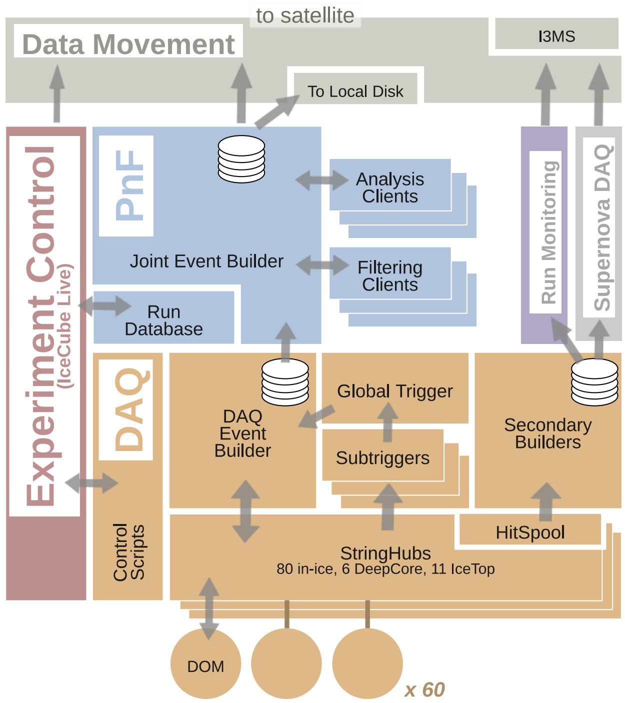

Figure 3: Diagram of the IceCube online data systems.

Despite the presence of two on-site, year-round experiment operators, the IceCube winterovers, the isolated nature of the IceCube online system presents unique challenges to operations. Collaborators who share the distributed duties of detector monitoring as well as the various subsystem experts are usually only able to access the experiment via satellite internet connection, often with very limited bandwidth. An experiment control system abstracts many aspects of the remote communication and provides worldwide monitoring and control interfaces to the operators and monitors.

1.1.3 Data Transfer, Processing, and Storage in the Northern Hemisphere

Filtered events that do not trigger real-time alerts may still be events of interest, and approximately 100 GB per day are queued for daily transmission to a storage cluster operated at UW–Madison via the high-bandwidth TDRSS bulk-data channel. Once in the Madison data warehouse, filtered data is further processed to a level suitable for scientific analysis, L2 or "Level 2." At this point, the data is still generic and common to all science analyses. Analysis-specific processing begins at L3. Long-term management of the data in online storage systems is discussed in the separate data management plan.

1.1.4 Simulation

Conversion of event rates into physical fluxes ultimately relies on knowledge of detector characteristics numerically evaluated by running Monte Carlo simulations that model fundamental particle physics, the interaction of particles with matter, transport of optical photons through glacial ice, and detector response and electronics. Vast datasets containing simulations of background and signal events must be produced and cataloged for use by the data analysts. The computationally expensive numerical models necessitate a distributed computing model comprising multiple clusters of thousands of CPU cores and specialized GPUs (graphical processing units).

1.1.5 Software Assets

From the firmware that controls the low-level acquisition of PMT waveforms to the algorithms that model detector response, software spans many levels and programming environments and is a significant asset of the ICNO. Over 100 distinct software projects have been contributed by the collaboration. However, core software frameworks have been developed and are maintained by engineering effort supported by the M&O program, which additionally oversees the management of interdependencies, scheduling of periodic code reviews, and support for development and release environments, including version control repositories, automated build systems, and unit testing frameworks.

1.2 The Era of Neutrino Astronomy—a Strategic Vision for ICNO Science

Because neutrinos originate in the decay of pions, astrophysical neutrinos unambiguously trace the still enigmatic cosmic sources where protons and nuclei are accelerated throughout the Universe; some reach energies exceeding the reach of the LHC by more than a factor of ten million. The discovery of high-energy cosmic neutrinos by IceCube [2] with a flux at the level anticipated for those associated with cosmic-ray accelerators [3] brightens the prospect for identifying the sources of the highest energy particles.

The original observation specialized to neutrinos interacting inside the detector, mostly showers initiated by electron and tau neutrinos. Since then, the conventional method of selecting upgoing muon tracks, created below the detector by muon neutrinos that have traversed the Earth, has delivered a sample of cosmic neutrinos with similar statistics but with superior angular reconstruction accuracy of better than 0.5 degrees [4]. Though limited to muon flavor, this high-resolution sample represents a watershed towards the identification of the sources. It has revealed secondary muon tracks that deposit energies of several hundred TeV inside the detector, indicating PeV-energy parent neutrinos.

While the original data suggested the possibility of clustering near the Galactic center and the Galactic plane, the statistical evidence for this was not compelling and has not increased after doubling the number of events [5, 6, 7]. Instead, the observed cosmic neutrino events are rather consistent with an isotropic neutrino flux with equal contributions from each neutrino flavor [8], suggesting the observation of extra-galactic sources whose flux has equilibrated in the three flavors after propagation over cosmic distances [9, 10]. Increasingly, a variety of analyses [8, 11] indicate that the cosmic neutrino flux dominates the atmospheric background above an energy that may be as low as 30 TeV with a spectrum that cannot be described as a single power. This is reinforced by the fact that fitting the excess flux in different ranges of energy yields different values for the spectral index.

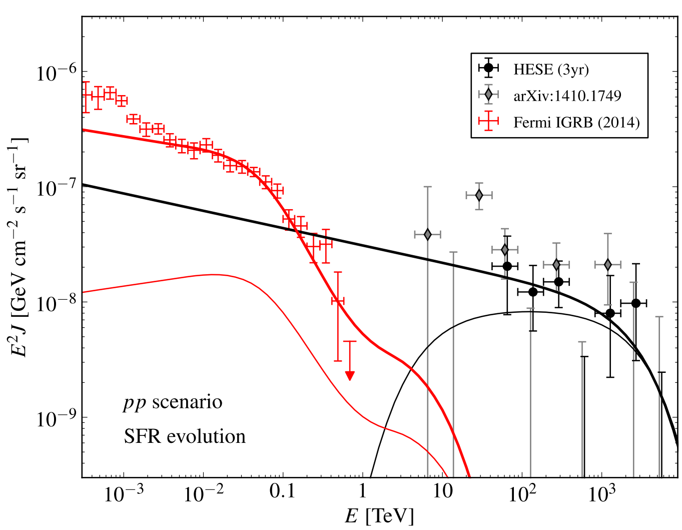

Figure 4: Two models of the astrophysical neutrino flux (black lines) observed by IceCube and the corresponding cascaded gamma-ray flux (red lines) observed by Fermi. The models assume that the decay products of neutral and charged pions from pp interactions are responsible for the non-thermal emission in the Universe [13]. Further details in [11].

The production of PeV neutrinos is inevitably associated with the production of PeV gamma rays: hadronic accelerators produce fluxes of both neutral and charged pions that are the parents of gamma rays and neutrinos, respectively. Gamma rays are attenuated over cosmic distances by interaction with the diffuse radiation backgrounds of the Universe. They suffer absorption through pair production processes in interactions with cosmic microwave and infrared background photons and cascade down in energy to a diffuse gamma-ray flux in the GeV–TeV energy range. It is intriguing that the gamma-ray flux anticipated on the basis of IceCube's neutrino measurements provides an excellent match to the extragalactic gamma-ray flux observed by the Fermi satellite; see Figure 4 [12]. The matching relative magnitudes of the diffuse extragalactic gamma-ray flux observed by Fermi and the high-energy neutrino flux measured by IceCube suggest that they originate in common sources at some level. The observed extragalactic flux exceeds IceCube's limit on the emission of neutrinos from gamma-ray bursts by one order of magnitude. As a result, the active galactic nuclei and starburst galaxies that dominate the extragalactic Fermi sky have emerged as best-bet candidates for accommodating the neutrino observations.

IceCube's DeepCore infill has enhanced its scientific potential as an atmospheric neutrino detector. DeepCore uses the surrounding IceCube strings as a veto in order to observe the muon tracks produced in contained neutrino events. With an energy threshold below 10 GeV, DeepCore is sensitive to the disappearance of atmospheric muon neutrinos and covers with unprecedented statistics the first oscillation minimum near 20 GeV. Its instrumented volume is on the order of 10 Mton. Higher energy atmospheric neutrinos detected by IceCube serve as a control sample of neutrinos that do not oscillate.

The results of a measurement of neutrino oscillations with three years of data of the completed IceCube Neutrino Observatory have been published [14]. In a livetime of 953 days, 5174 events were found compared to an expectation of 6830 without oscillations. The resulting oscillation parameters achieve a precision comparable to that of the Super-Kamiokande experiment, which performs a similar analysis but at energies lower by one order of magnitude. Therefore, the measurement represents a novel test of the three-neutrino oscillation formalism.

The study of neutrino oscillations and the ongoing correlation studies between the arrival directions of cosmic neutrinos and high-energy photons are only two of a large range of analyses that, on average, have produced on the order of 20 publications per year over the last five years. They also illustrate that, at best, these discoveries are in their initial stages. For instance, we have developed analysis techniques that specialize to DeepCore data that will result in improved measurements of the atmospheric oscillation parameters. The accumulating cosmic neutrino events indicate an enhancement of the flux below 100 TeV, which may hint at another component in the cosmic neutrino spectrum. The 10,000 TeV event indicates that the spectrum continues uninterrupted through the Glashow resonance, where 6,400 TeV anti-electron neutrinos produce real weak bosons in interactions with atomic electrons in the ice. Continued operation of the ICNO through this proposal will create the opportunity to observe the next Galactic supernova explosion with unprecedented statistics and precision. In addition, improved data and analysis tools will continue to result in rapid progress since no analyses are statistics limited at this point.

2 Results from Prior Work

The current proposal covers the continued maintenance and operation of the ICNO previously performed with funding under NSF grant ANT-093764, titled "IceCube Maintenance and Operations 2010–2015," in the amount of $34,693,749.00.

Intellectual Merit. In the last five years, the team submitting the present proposal has operated the IceCube detector and delivered science-ready data that have resulted in about one hundred scientific publications in refereed journals over the same period (see References section). During this period, the detector has been stable, and all aspects of its performance have improved: uptime, telescope effective area, pointing and energy resolution, and time to delivery of science-ready data as well as simulation. Continued progress is made possible by a close collaboration of the core M&O team with the IceCube Collaboration on improved calibration of the detector, including the characterization of the optical properties of the ice and the development of improved triggering, filtering, and data analysis tools. The development of highly efficient data summary methods has resulted in significantly improved quality and turn-around time from data-taking to scientific analysis. Precision of scientific results is still limited in certain areas by computing. However, rapid progress has resulted from well-timed upgrades of IceCube computing capabilities and access to increased resources elsewhere. Success has been greatly enhanced by conversion to GPU-based production for simulation data.

In summary, the detector as it was completed in December 2010 would not have yielded its major successes—discovery of cosmic neutrinos and competitive results covering neutrino physics from 10 GeV to 10 PeV—had it not been for the steady improvement of the detector's performance and all aspects of its data acquisition and analysis that was realized by the current M&O team.

Broader Impacts. Our data access plan has resulted in a wide variety of collaborative efforts with an increasingly broad user community, which have contributed to significant cross-disciplinary research. Additionally, we have MoUs in place for coordinated observations with the leading cosmic-ray, gamma-ray, and gravitational wave observatories (see Attachment 6). These have already resulted in six publications. The M&O core group and the IceCube Collaboration have successfully created opportunities to introduce undergraduate students to research and have trained a large number of successful scientists, technicians, and academics. Documentaries, articles in popular science publications, public talks, activities at science festivals, scaled detector models, and innovative partnerships are just a few of our undertakings in regularly engaging the public. The diversity and impact of this public engagement, which will be discussed further in later sections, have been boosted by the discovery of cosmic neutrinos.

3 Management and Operation of ICNO

The primary effort of this proposal is to manage a workforce of scientists, engineers, technicians, and administrators to ensure that M&O tasks are properly defined and that necessary resources are available. This section highlights requirements, previous successes, and plans for the future within five main areas of responsibility: management, operations, budgeting, planning, and education and outreach.

3.1 Management

Our approach to IceCube M&O is to maximize the scientific discovery potential by coordinating talent and resources from the collaborating institutions with that of the core M&O team to support both M&O and science tasks. This section describes how we are organized to perform the M&O functions for IceCube in this distributed model and how we provide accountability for task execution.

3.1.1 Organization

The IceCube M&O management organization integrates the IceCube Collaboration and the host institution, UW–Madison (see Figure 5). The principal investigator (PI) is responsible to the UW vice chancellor for research and to the National Science Foundation for the overall scientific direction of the IceCube Neutrino Observatory. The director of operations appoints technical professionals to serve as managers of the two M&O functions that are predominately centered at UW–Madison: detector M&O 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 way. The collaboration spokesperson appoints collaborating scientists to serve as coordinators for physics analysis. 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 (included as Attachment 1). WBS tasks are defined in detailed MoUs, proposed by collaborating institutions and reviewed by 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 collaboration meetings and include a list of the physics group members and a head count of faculty, scientists, postdocs, and graduate students. (An MoU summary is included as Attachment 2, and further details about institutional representation and committee structure can be found in the IceCube Governance Document included as Attachment 5).

3.1.1.1 The IceCube Collaboration

The collaboration plays a leading role in IceCube, guiding both science and M&O. 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 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 and authority over science policy and goals, membership, data access, publications, representation of IceCube at conferences, analysis teams, and education and outreach (E&O).

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.

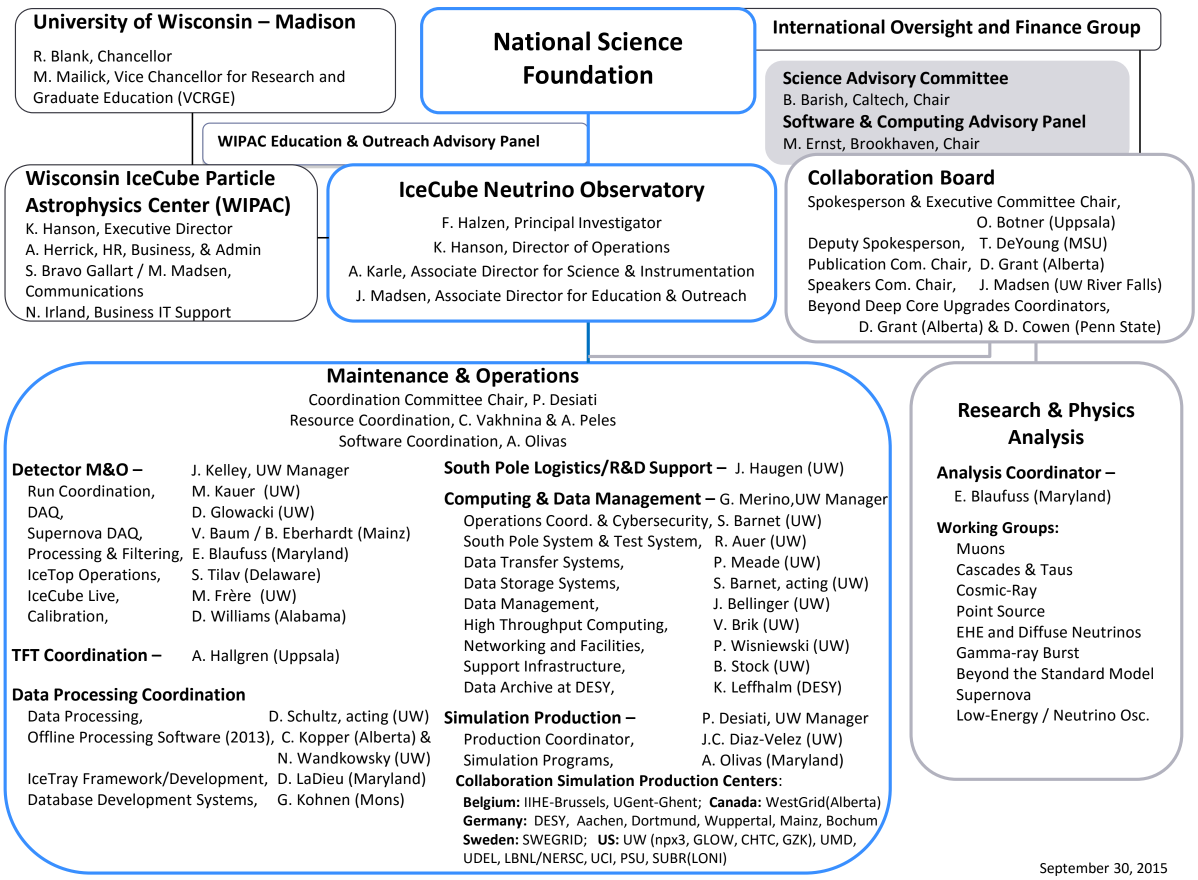

Figure 5: The IceCube organization maximizes the use of both collaboration resources and core resources managed by UW while maintaining clear lines of accountability to the NSF.

3.1.1.2 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 operation of the Amundsen-Scott South Pole Station. The IceCube M&O award is co-funded by the Division of Polar Programs and the Particle Astrophysics Program within the Division of Physics. The respective program directors provide continuous oversight and guidance through direct communication with the IceCube PI and director of operations.

3.1.1.3 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 comprised of representatives of the funding agencies in the partner countries supporting IceCube M&O. (For more details, see the IceCube Collaboration Governance Document in Attachment 5.)

3.1.1.4 University of Wisconsin–Madison

IceCube Oversight. The lead executive officer of UW–Madison, the chancellor, delegates responsibility for research activities to the vice chancellor for research and graduate education (VCRGE), who maintains oversight of the ICNO and appoints the IceCube director of operations.

The IceCube PI and the director of operations report directly to the VCRGE and report regularly to the university's IceCube leadership team, which includes the chancellor, provost, VCRGE, and vice chancellor for administration/budget, planning and analysis. 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. IceCube's M&O falls under the purview of WIPAC, which is the primary interface to the university administrative and support systems, established within the Office of the VCRGE to coordinate the multiple roles of the university, such as lead and host institution for the IceCube construction project and for IceCube M&O. WIPAC provides administrative services such as accounting, purchasing, and human resources, coordinates E&O activities, and collaborates with the largest participating research group.

3.1.1.5 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 1.

Table 1: IceCube M&O U.S. Subawards – Major Responsibilities

Institution	Major Responsibilities
Lawrence Berkeley National Lab.	Data acquisition maintenance, computing infrastructure, long-term data archive
Pennsylvania State Univ.	Simulation production, DAQ firmware support
Univ. of Delaware, Bartol Institute	IceTop calibration, monitoring and maintenance
Univ. of Maryland at College Park	Overall software coordination, IceTray software framework, online filter, simulation software
Univ. of Alabama at Tuscaloosa	Detector calibration, reconstruction and analysis tools
Michigan State Univ.	Simulation software, simulation production
Univ. of Wisconsin–River Falls	Education and outreach coordination

3.1.2 Planning, Coordination, and Performance Assessment

3.1.2.1 Advisory Committees

3.1.2.1.1 Science Advisory Committee.

The primary goal of IceCube M&O is to ensure that IceCube meets its 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 review the performance of the M&O organization and make recommendations on scientific goals and other matters that may affect ICNO scientific activities. The SAC meets annually. The current chairperson is Barry Barish from Caltech.

3.1.2.1.2 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 and director of operations 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 the chairperson. The SCAP meets annually. The current chairperson is Michael Ernst from Brookhaven National Laboratory.

3.1.2.2 M&O Coordination Structure

3.1.2.2.1 Coordination Committee

The Coordination Committee provides high-level coordination of core M&O tasks and service contributions of collaborating institutions, ensuring that core and in-kind M&O functions meet the needs of the physics analysis working groups. The 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 ensures that students and postdocs working on those tasks get visibility by providing reports at the IceCube weekly phone calls. The committee meets every two months by teleconference and at each collaboration meeting.

3.1.2.2.2 Trigger Filter Transmit 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 SPS. Ahead of the yearly physics run transition in May, at which time 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 or additions to the set of detector triggers and online event filters. Following the transition to the new season's configuration, the TFT board tasks all physics working groups to provide brief reports summarizing the status and quality of selected events.

3.1.2.2.3 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 non-standard operations on 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 experts, and other physicists.

3.1.2.2.4 Calibration Coordination

The calibration group designs calibration runs for the detector, analyzes calibration data, and provides tools for utilizing the calibrations in order 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 a weekly teleconference held with students and postdoctoral researchers.

3.1.2.2.5 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.

3.1.2.2.6 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 a high quality of software. The software coordinator, along with a core development "strike team," addresses the most urgent needs of the collaboration by participating in week-long monthly "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.

3.1.2.2.7 Simulation Coordination

The simulation production coordinator, working with the simulation coordination committee, guides the computational production of simulated physics data and works with the collaboration to balance analysis needs versus the availability of computing resources. Production site managers assist with simulation production on their local institutional computing systems. The coordinator monitors data benchmarks and verifies simulation releases before full production. The coordinator produces quarterly reports on global simulation production status and a weekly status report on simulation production. The coordination committee meets every two weeks to discuss simulation status and to establish milestones.

3.2 ICNO Maintenance and Operations

The operation of the IceCube Neutrino Observatory involves the maintenance and development of the detector hardware and online software running at the South Pole; the computing infrastructure both at the IceCube Lab and in the Northern Hemisphere, including distributed computing; and the core data processing, simulation, and analysis software framework and tools. Operations experts must respond dynamically to support a wide variety of science goals, continuously improving in order to maximize the potential of the experiment. This section describes the tasks involved in each of these areas, detailing the strategies we will employ to expand upon our past operational and scientific successes.

3.2.1 Detector Maintenance and Operations

All subsystems in the IceCube infrastructure require effort to maintain and operate. Even though some hardware systems are frozen into the ice, the overall system changes in time. Calibration constants drift over time, data rates change due to the seasonal fluctuations of the atmosphere, and sensors may display defects and need quick attention to avoid serious system-wide problems. Furthermore, detector operations must respond to new discoveries and analysis techniques to ensure IceCube continues to achieve its maximal science potential.

Several performance metrics are used to track the operation of the detector (see Table 2), in particular, the uptimes of the DAQ, the IceCube Live monitoring and experiment control system, and the supernova data acquisition system. Continued improvements in online software and operational methodology have led to significant increases in the most important of these performance metrics (see Figure 6), and we are now regularly exceeding our aggressive detector uptime objectives.

This section outlines the technical M&O requirements and specifications ensuring that IceCube performs reliably and continuously provides the capability to achieve its scientific objectives. Here and in Section 3.4, we propose improvements to the data acquisition software, detector monitoring, calibration, and IceTop hardware that will enhance the reliability and science potential of the ICNO.

Table 2: Detector Operations Performance Metrics

Key Performance Metric (Uptime)	**Objective**	**Achieved (2015)**	**Description/rationale**
Total Detector	99%	99.7%	Measure of total time that detector was operating
Clean Detector	95%	97.0%	Indicates production of pristine, full-detector data for physics analysis with no contamination or serious alerts
IceCube Live	99.9%	99.9997%	Critical to monitor and resolve detector issues
Supernova System	99%	99.4%	Measure of time that the detector was operating and sensitive to supernovae

3.2.1.1 Data Acquisition

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

The winterovers maintain and repair the DAQ hardware at the South Pole; the monitoring and paging system alerts them to failures of any DOMHub components, at which point they can exclude the faulty hardware from the detector while it is repaired. Upgrades to the hardware are tested at SPTS prior to deployment at SPS. In 2012 and 2013, we upgraded all DOMHub single board computers, providing a significant performance boost and power savings.

A number of custom surface hardware components contain obsolete parts and cannot easily be remanufactured. While the failure rate of these components is currently low relative to the inventory of spares, support engineers will investigate alternative solutions using modern electronics, as a contingency plan in case an upgrade is necessary to maintain the operation of the detector.

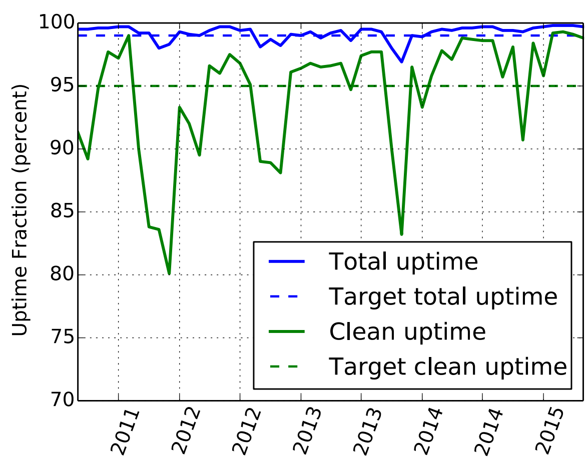

Figure 6: Detector uptime and clean uptime statistics are shown from the start of full-detector operations in May 2011. Clean uptime continues to improve due to improvements in operations software and methodology.

DAQ software 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 software is a major driver of the quality of data for physics analysis.

DAQ software engineers are accountable for the uptime of the DAQ and the integrity, correctness, and completeness of the data it produces. They also provide appropriate documentation for the operators. They regularly test and upgrade DAQ software systems—including DOM software, DOM readout card device drivers, DOMHub software, triggers, event builder, secondary builders, and control scripts—while responding to evolving science needs. 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 will further decouple the DOMHub data-taking component from the higher-level trigger and event builder components. This "hitspooling" system also provides the ability to save short periods of all untriggered sensor data in case of a supernova alarm; we plan to extend this to other interesting transient events.

The DOM firmware consists of a low-level FPGA design that controls the DOM hardware, collecting hit data and buffering it in memory, and supporting a large variety of configurations and controls. A DOM firmware engineer will supply any required FPGA modifications, maintain the code base, and update documentation as needed. Both DOM and DOM readout firmware are relatively stable, but maintenance of these systems is critical to the health of the detector.

A small fraction of DOMs (0.4%) have malfunctioned and must be operated in a non-standard configuration as part of normal data-taking. A fraction of DOMs (1.5%) have also failed, which results in the need for a detector reconfiguration. Since we have instituted procedures that minimize DOM power-cycles, the rate of DOM failures has slowed, the last two failures having occurred in May 2013 during a power outage; however, we will continue to monitor the health of the DOMs closely.

3.2.1.2 Online Filtering and Multimessenger Follow-up Alerts

The volume of data produced by DAQ far exceeds the limited bandwidth available in IceCube's satellite allowance. An online processing and filtering system (PnF) is used to apply a set of initial "Level 1" event selections to the collected data, transmitting only those selected events (approximately 10% of the raw data). Additionally, the online PnF system generates alerts when astrophysical neutrino event candidates are detected, enabling immediate response to these events by other telescopes.

PnF system experts will maintain the online system, ensure filters are being properly applied, and respond to and debug unexpected errors. This effort ensures that the online filtering system produces the highest quality data. Maintenance is performed at the start of each new physics run and on an as-needed basis at other times. This will include requests from the TFT Board to support new analysis priorities and alert categories.

The online PnF system supports several event formats, including highly compressed formats based on information extracted from DOM waveforms that allow for more efficient use of IceCube's satellite bandwidth and has enabled a new set of filters to be added. Additionally, all events are saved using this compressed format at the South Pole, allowing for reduced archival data sample sizes.

Collaboration physicists and software engineers will continue to work together to define fast, robust analysis schemes that can alert other telescopes for follow-up observation of interesting events, localized in time and/or direction. A recent improvement to the online follow-up system sends a continuous stream of neutrino events to servers in the Northern Hemisphere, using IceCube's Iridium satellite connection. A new real-time alert system will be deployed for high-energy starting events (HESE), IceCube's most signal-rich sample of astrophysical neutrinos.

Each year, we evaluate the filters that select events for immediate transmission north for further analysis to ensure that they meet the evolving physics needs of the collaboration and that the most effective reconstruction and filtering tools are in use online. Collaboration physics working group members with expertise in the available analysis tools and goals of the physics program provide filters to the TFT Board for evaluation. The working group members research and write initial proposals, participate in internal working group discussions, make presentations to the TFT Board, and report on the filtered data quality once the season has begun. Filter development is based on data samples generated for the season's physics run in the IceTray software framework, and using reconstruction tools. The TFT Board must approve the filtering system and for deployment at the start of each new year's physics run in May.

3.2.1.3 Experiment Control and Detector Monitoring

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 the 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 will support physics working groups and operators to add needed functionality and to respond to evolving science needs. Improvements to IceCube Live are responsible for the large increase in detector clean uptime shown in Figure 6 in 2014 and 2015. Prior to this time, DAQ runs that ended in failure were classified as bad and removed from the detector clean uptime. However, the system can now track and keep the large fraction of the run before the failure that is good. These software and procedural improvements have allowed us to achieve our target performance metric of 95% clean uptime.

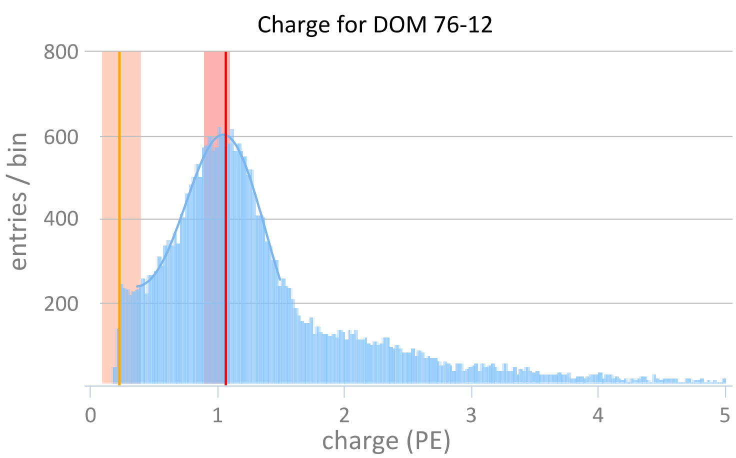

Figure 7: Measured single photoelectron (SPE) charge spectrum in I3Moni 2.0, for DOM 76-12. The vertical lines indicate the fitted threshold (left) and SPE charge (right), compared with acceptable bounds (shaded regions). The peak is used in DOM gain correction.

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. The station's 24/7 satellite link is overcommitted given the demands of station e-mail traffic. IceCube will develop a new messaging system technology (I3MS) to move critical monitoring and alert traffic off of the station link. A software engineer will design, update, and maintain this software, which will provide a higher-bandwidth, scalable solution that will alleviate bandwidth pressure on the station's infrastructure.

IceCube detector monitoring is the system that provides a comprehensive set of tools for assessing and reporting the data quality. Our original monolithic monitoring system processes data from various SPS subsystems, packages them in files for transfer to the Northern Hemisphere, and reprocesses them in the north for display on the monitoring web pages. In a new monitoring system we are developing (I3Moni 2.0), all detector subsystems will report their data directly to IceCube Live. Advantages of this new approach include: higher quality of the monitoring alerts; simplicity and easier maintenance; flexibility, modularity, and scalability; faster data presentation to the end user (see Figure 7); and improvement in the overall longevity of the system over the lifetime of the experiment.

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.

3.2.1.4 Calibration

The correct and efficient analysis of IceCube data relies on the use of a common set of calibrations and calibration tools. The IceCube run coordinator and the calibration working group lead orchestrate many of these tasks since they either require inactivation of detector segments or illumination of the fiducial volume. Often, specialized data sets are produced and analyzed offline, either on computers at the South Pole or in the Northern Hemisphere using collaboration-maintained algorithms.

IceTop DOM calibration runs are taken monthly, and in-ice DOM calibration runs are taken once a year. The results are validated, loaded into a database, and then used as part of online reconstruction and filtering. The winterovers are responsible for running DOMCal, with calibration working group support to validate and interpret the results from DOMCal runs. Software engineering support is required to maintain and upgrade the DOMCal software on the DOMs.

Collaboration graduate students and postdocs perform additional calibration tasks under the supervision of the calibration working group lead. For example, while on site at the South Pole, they take data with flashing in situ light sources to track the overall detector response to Cherenkov-like light and to measure the optical properties of the ice. A recent improvement to the detector calibration refines the gain calibration of each DOM individually with a few-percent adjustment based on single photons collected during normal data-taking (see Figure 7), increasing the accuracy of event energy reconstructions.

As a fundamental component of the IceCube detector, the glacial ice must be understood in detail in order to correctly reconstruct neutrino event direction and energy. The ice is layered, with scattering and absorption features varying due to historical geological dust deposition. The flasher runs provide the underlying calibration data needed to extract the optical properties of the ice. Scientists iteratively fit the flasher data with complex, many-parameter models to determine the scattering and absorption at various depths while correcting for differences in DOM efficiency. Recently, we discovered that the scattering is not isotropic but has a preferred direction aligned with the flow of the ice. Modeling this anisotropy has been crucial to the correct classification of neutrino event topologies, and continued refinements to the ice model will reduce the uncertainty in neutrino event directional reconstruction.

While the DOM's relative efficiency can be determined in situ with flasher data and low-energy muon measurements, determining the absolute efficiency requires a calibrated light source. The efficiency varies as a function of light wavelength, incident direction, and whether the DOM is in air or frozen in ice. Via ongoing laboratory calibration measurements, we will determine the DOM efficiency as a function of incident angle and wavelength to the few-percent level. NIST-calibrated photodetectors will be used as a reference to ensure a known number of photons are incident upon the DOM. The DOM can be rotated in a water tank to vary the angle of incidence. These measurements will complement in situ measurements of the DOM efficiency.

3.2.1.5 IceTop Operations

IceTop by itself includes all aspects of a major experiment, requiring 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. Data rates in individual DOMs are significantly higher, and typical signals are much larger than in the deep detector. 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.

An IceTop data specialist coordinates monitoring of the physical condition of the IceTop detectors, including annual surveys of snow accumulation above the tanks and surrounding environmental conditions at the South Pole. The data specialist also coordinates monitoring the quality of IceTop data and any corrective actions required to address malfunctions or other conditions that degrade IceTop data.

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 of the array by installing low-cost scintillator modules above the IceTop stations. Four prototype modules will be installed in the 2015–16 austral summer season and will use existing IceCube cabling to connect back to the ICL. An engineer will refine the design of the modules for potential deployment in future years. Already, an effort, supported by the UW analysis grant, to separate the electron and muon components in the IceTop detector data has been successful and is under further development.

3.2.1.6 Supernova Operations

Supernova data acquisition (SNDAQ) reads the single photoelectron scaler data produced by the DOMs and searches for a rate excess over the entire detector. A rate excess triggers a supernova candidate alarm is issued via satellite, and detailed data are saved. If monitors conclude that the alarm is significant, an additional alarm is sent to the Supernova Early Warning System (SNEWS) [15]. We also plan a reverse link that will save IceCube untriggered data in case of a community SNEWS alert.

Collaboration scientists are responsible for core SNDAQ development. Core software engineers are responsible for integrating SNDAQ into the experiment control, monitoring, and DAQ systems. A recent improvement in SNDAQ includes real-time integration with the DAQ trigger system. Using information on the muon trigger rate in the in-ice detector, SNDAQ corrects the significance of any fluctuation in the noise rates caused by background muons, increasing the probability that IceCube will alert other observatories via SNEWS to a Galactic core-collapse supernova.

3.2.2 Computing and Data Management

The computing infrastructure supporting IceCube operations is a natural extension of the detector systems. A powerful and complex computing infrastructure is essential to extract science results from IceCube data. The main goals of this proposal in terms of computing and data management are twofold. First, we will increase the capacity of the data processing and analysis services by steadily upgrading the computing facilities at the UW–Madison data center and increasing the amount of distributed resources. Second, we will roll out improved data management services that ensure long-term preservation and usability of the data as well as ease its discoverability.

3.2.2.1 South Pole System (SPS)

The SPS is a computing system developed and maintained by system administrators that supports the data acquisition and filtering tasks carried out by the IceCube real-time systems. The SPS hardware includes DOMHub computers, commodity server class computers and network hardware. It also contains infrastructure services such as DNS, mail, monitoring, and databases. The system is designed to be very robust since it has to 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 IceCube network integrates the detector systems in the ICL and the South Pole station with the USAP network and ultimately the data center at UW–Madison. It complies with policies and regulations of the NSF and UW. The systems are isolated from the USAP and other external networks by a firewall. Security logs are monitored for suspicious behavior and traffic signatures.

System administrators and winterovers are responsible for maintenance and operations of the SPS. This includes preventive maintenance, troubleshooting, and upgrades. The operating system and configuration management tools used in the SPS are the same as those used for all services in the UW–Madison datacenter. This allows applying consistent procedures across systems and efficiently managing version control, patching, software updates, monitoring, and maintenance.

In order to maximize availability, the SPS incorporates uninterruptible power supplies are installed 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 commodity computer servers in the SPS will be replaced every four years in order to profit from technological advances that maximize computing power and minimize the risk of component failures.

3.2.2.2 South Pole Test System (SPTS)

The SPTS is a test system that replicates the basic functionality and performance characteristics of the production SPS and represents a reasonable proxy representation of 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.

The SPTS is a scaled-down version of the SPS and is located on the UW–Madison campus. System infrastructure is essentially identical to that deployed on the production system including power and network devices. All major subsystems are represented although some with a reduced number of nodes. Nonetheless, infrastructure is in place to expand the number of nodes for specific large-scale tests by temporarily integrating servers from the offline cluster.

Engineers evaluate software and firmware DAQ updates on the SPTS via one or more of the following: (a) 8 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, and (c) 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.

Core system administrators are responsible for hardware and software maintenance and operations on the SPTS as well as for adapting the infrastructure to the evolving requirements. Engineers maintain and upgrade the system to ensure maximum uptime when it is required for testing. They provide support for adding new functionality in the SPTS in response to new science needs.

3.2.2.3 Offline Computing Resources

3.2.2.3.1 High-Throughput Computing (HTC) System

System administrators will continue to operate an HTC cluster dedicated to IceCube at UW–Madison to perform offline processing and analysis of real data and production of simulation data sets. The system is closely coupled to the storage infrastructure for efficient data processing.

Graphics processing units (GPUs) have been found to be a very effective resource for simulating photon propagation in the ice. A GPU-based cluster is available at the UW–Madison data center as part of the HTC cluster infrastructure. The current IceCube HTC cluster at UW–Madison consists of nearly 200 servers providing a total of about 4000 CPU and 350 GPU job slots.

The HTCondor [16] 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. 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 they also 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.

In order to benefit from technological advances and improvements in energy efficiency, we will ramp up the cluster capacity each year, following demand, and will keep servers in production for five years.

3.2.2.3.2 Distributed Computing

Dedicated computing resources at the level of several thousand CPU cores are needed to perform required simulation and analysis tasks. In order to reach this capacity, IceCube relies on distributed resources available from collaborating institutions. Figure 8 shows the CPU resources used by IceCube in the last two years, illustrating that distributed resources represent a sizeable contribution of the overall capacity. Also, it shows how the distributed resources provide the important capability of handling peak demand loads. The current capacity plan foresees providing 10% of the CPUs needed for simulation at the UW–Madison data center, 25% at DESY-Zeuthen, and 65% at other collaboration sites. For GPUs, the plan is to provide 50% of the needs at UW–Madison and 50% at other collaboration sites. The current capacity allows generating one year of simulated data in one year of compute time. Our goal is to triple this capacity by the end of the next five-year cycle. This is important in order to have the capability of regenerating large amounts of simulation when major changes in the software occur.

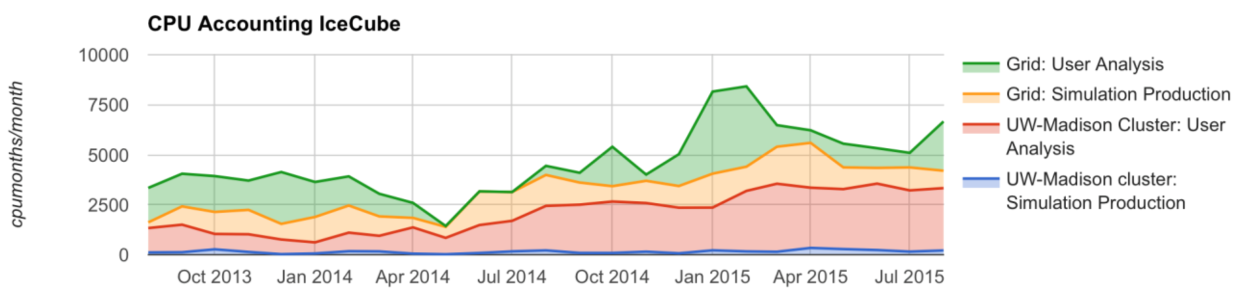

Figure 8: CPU time used by IceCube simulation and analysis processes. The blue and red graphs at the bottom represent the resources used at the dedicated HTC cluster in UW-Madison. The yellow and green graphs correspond to jobs that ran elsewhere in distributed resources (Grid).

In addition to the collaboration resources, IceCube will continue to tap into other opportunistic resources (mostly from Grid projects in the US and Europe) as much as possible. This is an effective and efficient way to produce larger statistics of simulation data, enabling the higher precision required for some analyses. Several of the nation's most advanced computing systems that are part of the Extreme Science and Engineering Discovery Environment (XSEDE) [17] currently have GPU nodes. Preliminary tests exploring the feasibility of using some of these resources for IceCube simulation show promising results. Our goal is to steadily increase the contribution to simulation from supercomputers in the next five years by regularly requesting compute time allocations and developing interfaces for integrating these resources in the IceCube distributed workload management system.

The efficient use of distributed resources requires coordination among the different sites as well as the use of Grid software such as job scheduling and data access tools. System administrators will continue managing the core services of this IceCube Grid infrastructure. Standard tools are used to manage the Grid resources wherever possible, which engineers then interface with IceCube specific software. In order to manage this process efficiently, it is essential to maintain close contact with the computing community. We will ensure this by participating in the Open Science Grid (OSG) project or the National Data Service initiative. Our goal is also to integrate the IceCube HTC cluster into OSG, contributing to this large research infrastructure from which we benefit so much.

Collaborating institution DESY-Zeuthen provides a Tier-1 data center, which assumes a number of core computing activities to complement UW's. The goal is to leverage additional locations to ensure core data processing and analysis services are highly resilient. The DESY-Zeuthen data center supplies significant computing and storage infrastructure for simulation and analysis and also acts as a replication site for the long-term archive of IceCube data sets (see Attachment 6 for further information).

3.2.2.3.3 Data Center Networking and Security

The IceCube Collaboration bridges 12 countries. Many of the data products and services are hosted at the UW–Madison data center, hence excellent global network connectivity is essential. Also, a large part of IceCube's computing requirements will be met using distributed resources, therefore reliable, high-speed access to the UW–Madison data center storage infrastructure is required to use these resources efficiently.

One of the key aspects in designing and operating the IceCube data center network is ensuring optimization for high-volume bulk data transfers while keeping critical infrastructure secure. The bulk data movement servers are deployed in a subnet that is structured to be secure, but without the performance limits that would result from passing data through a firewall. Figure 9 illustrates the volume of outgoing data (left) and incoming (right) of the data center through the high-performance access nodes for the past year.

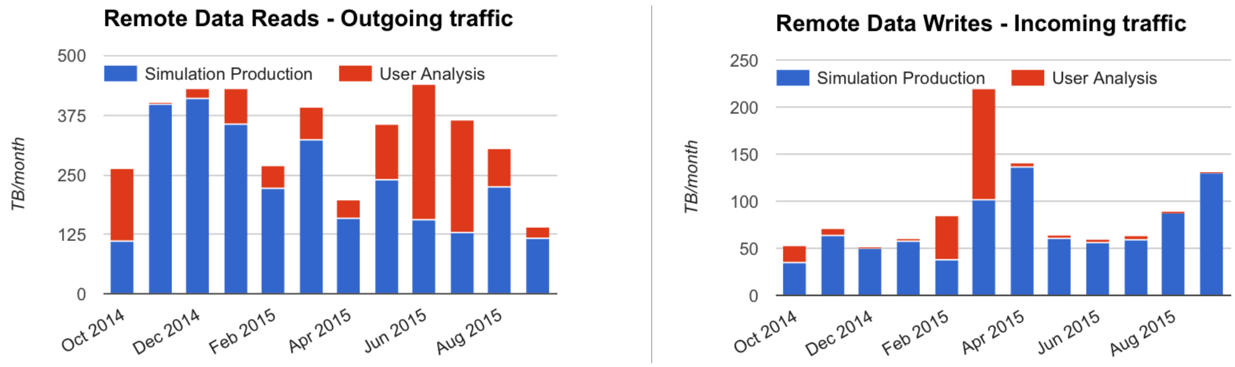

Figure 9: Monthly data volume handled by remote data access service at the UW–Madison data center.

We will continue to maintain a cybersecurity program at the UW-Madison data center. These activities will ensure that the policies and security controls continue to provide access for legitimate users within the IceCube Collaboration but maintain a highly secure network environment robust against hostile attacks. In particular, ensuring secure and stable operation of the detector computing systems at the South Pole and the data handling facilities at UW–Madison will remain a priority. In addition, we will maintain contact with other NSF large facilities and cybersecurity programs such as the Center for Trustworthy Scientific Cyberinfrastructure (CTSC) to share knowledge and ensure that WIPAC's practices are consistent with those accepted in the larger community.

3.2.2.3.4 Data Center Infrastructure

Core computing infrastructure systems include distributed authentication, DNS, and email. Also required are a large number of servers and services, such as database or web services and custom application servers to host services such as real-time alerts and data quality monitoring.

IceCube computing facilities are currently hosted in two UW–Madison locations, one off campus at the WIPAC offices and the other on campus at the Physics Department. Together they provide the total capacity to power and cool about 170 kW of IT equipment. This is expected to be sufficient for IceCube's needs in the next five years, as improvements in energy efficiency in new generations of hardware will enable a net capacity growth within a constant power envelope. The space, power, and cooling for these facilities are in-kind contributions from WIPAC and UW–Madison to this project.

3.2.2.4 Data Storage and Transfer

3.2.2.4.1 Data Archive and Transfer Software

A reliable system is needed for managing data transfers and long-term storage. The data generated by the detector is its most precious output, so the archival system is an important safeguard. Also, the network at the South Pole is a very scarce resource and is only accessible via very specific interfaces, such as the satellite channel. The data transfer system must comply with these interfaces to manage this resource with high efficiency.

Core software engineers have developed a software application named JADE has been developed to handle the data at the SPS. JADE manages both the archiving of generated data and the near-real-time transfer of filtered data via satellite. Two copies of the generated data are stored on different disk drives at the South Pole and shipped to the UW–Madison data center once a year during the austral summer for long-term archive. For the daily data transfers, the service makes use of the Iridium satellite systems for high-priority low-volume data or the dedicated high-capacity TDRSS satellite system for the bulk of the data.

JADE runs on several servers to achieve higher reliability and scalability. Because data integrity cannot be guaranteed over satellite transfers, the software maintains checksums of all files. It periodically contacts the data ingest services at UW–Madison's data center to get information on successfully transferred files and retransmits missing or corrupted files as needed. A core software engineer will maintain the transfer and archiving software and manage the daily operations of data transfer.

3.2.2.4.2 Data Ingest and Catalog Software

In the data center at UW–Madison, data needs to be retrieved from the satellite and, after integrity has been verified, entered into the data warehouse and cataloged.

In the initial stages of the project, core software engineers developed an application named INGEST to take care of entering new data files into the data warehouse. A core engineer will operate the INGEST software, monitor the data transfer quality daily, and address any issues that the system might have in interfacing with TDRSS or with the data center storage infrastructure.

As a near-term goal, we will integrate the INGEST functionality into JADE to improve its functionality and operability. Having a unified software system handling data archiving and end-to-end transfer will also simplify the setup and reduce maintenance costs in the long term. A software engineer will integrate the data ingest functionality into JADE and will improve it by adding missing key features such as the capability to efficiently index the metadata. A new catalog service will be developed that will provide users access to the file metadata database and will enable efficient searching of any file produced by IceCube. This tool will also allow managing storage allocations or data retention policies in a more efficient way.

3.2.2.4.3 Data Storage Infrastructure

The complete IceCube data set will continue to grow as data is collected, processed, and analyzed at a rate of about 1 Petabyte (PB) of new data every year, which will increase the load on data access services. The storage infrastructure, and in particular the export services responsible for providing remote access to the data, will need to evolve in order to cope with the extra load and maintain high performance and reliability.

The main facility for processing and analyzing IceCube data is the UW–Madison data center, whose infrastructure consists of disk storage servers organized in 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 will ensure that data is available and that it is delivered to data processing and analysis tasks with maximum performance. The storage system administrators will 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 balance. In order to benefit from technological improvements in storage density and energy efficiency, we will ramp up the storage infrastructure capacity each year following demand and will keep disk systems in production for five years.

3.2.2.4.4 Long Term Preservation and Archive Services

Of the 1 PB of data generated every year, about 700 TB will need to be archived and preserved long term due to their uniqueness or their relevance to reproducing published scientific results. The most cost-effective technology for archiving this data is magnetic tape. Automated tape libraries at the PB scale are not commodity infrastructure in terms of hardware or software. There are very high maintenance and fixed costs associated with them. In order to provide cost-effective long-term data preservation services, we will leverage large data centers at collaborating institutions, such as NERSC and DESY-Zeuthen, that already operate on a large scale and can provide long-term data archive and curation as a service (see Attachment 6 for further information).

A software engineer will develop a software layer to interface the IceCube data warehouse and the remote long-term archive systems. This service will manage the data transfers and will ensure that the data products stored in the long-term archive are registered in the data catalog so that they can be found and accessed transparently by IceCube applications.

3.2.2.4.5 IceCube Open Data Services and Tools

The IceCube Collaboration already provides public access to event reconstruction information for events selected as neutrinos for specific published results. The plan, as described in the IceCube Collaboration data sharing policy, is to also release 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. Documentation and tools will be maintained as well.

As usage increases, the public data sets themselves will need to be maintained so that improvements are implemented and issues corrected. Also, user support will be needed so that queries from external users are addressed. A team of two IT specialists will initially take care of developing and rolling out the IceCube open data service. One of them will have a physics-oriented focus, working on the formats and tool specification as well as providing user support. The other will have a technical focus, working on the implementation of web services and other tools needed to make the data accessible to other communities.

3.2.2.5 Offline Data Processing and Data Quality

The software for processing data for physics analysis comprises submission scripts for processing jobs to the compute elements of the central HTC cluster, processing scripts, database software to monitor job execution, and web pages to display processing progress and quality parameters.

A software engineer adapts data processing based on the detector configuration and required reconstruction algorithms developed by the collaboration. The software engineer also adapts submission and monitors execution to make the best use of the available computing resources. Close coordination with the run coordinator ensures data-quality issues are well understood and any bad data are removed from processing. Improvements to the data processing flow have led to a dramatic reduction in the time frame in which data are available to the collaboration for scientific analysis. By defining the offline data processing with the TFT Board before each physics run start in May, and by reviewing data quality issues weekly at operations calls, the processing delay time has been reduced from 13–15 months to less than two weeks.

3.2.3 Software

3.2.3.1 Core Software Systems

IceCube’s core software library consists of the IceTray framework, a set of basic modules and data containers, and a wide range of open source tools that are used in the development of calibration, simulation, reconstruction, and analysis modules. A robust set of Python bindings is also included, which facilitates the use of advanced analysis environments and 3-D graphical event displays.

The IceCube software coordinator is responsible for coordinating maintenance of all core software as well as the software repository system, continuous-build testing system, and external libraries as newer operating system versions emerge. The software coordinator also conducts regular training sessions for new IceCube collaborators.

Central databases, with mirrors in key locations to enhance efficiency of data access, store key IceCube information such as detector geometry and calibration, configuration, and run summaries. Database locations include the South Pole, Belgium, and Madison. Keeping the contents of these databases well organized, synchronized, operating, and available is key to ensuring that all parts of IceCube data analysis are well understood and repeatable. A software engineer will maintain and extend the database tables and will maintain all code to update and query the database. Continuous support for data insertion at the South Pole and in the Northern Hemisphere provides all necessary information for data processing. In order to simplify long-term maintenance and reduce redundant information, we will develop a new database system that will leverage detector status information already in the IceCube Live monitoring system.

Software development in IceCube is a worldwide, distributed effort with more than 100 contributors and running on several different platforms, thereby maximizing grid resources. Critical software development tools, such as a central repository, ticketing system, and continuous build test system, will be maintained by a computer scientist, using industry standards such as Subversion, Trac, and Buildbot.

Future efforts will focus on improving the speed, efficiency, and robustness of production software through training, workshops, and the integration of modern tools, such as clang's Static Analysis, into IceCube's development workflow. The goal is to maximize efficient use of all of IceCube's computing resources, such as disk space, CPU, and GPU power, while increasing background simulation samples, which is critical for several analyses.

3.2.3.2 Reconstruction and Simulation Software

IceCube's reconstruction software runs online for filtering at the pole, offline in the north, for higher filter levels, and as a starting point for analysis. The reconstruction software is managed directly by the software coordinator and a software engineer. IceCube's simulation software has to cover a wide dynamic range, supporting low energy at the GeV scale as well as ultra-high-energy at the EeV scale, while running on large-scale, heterogeneous computing systems that include batch processing clusters and grids. This requires a dedicated framework (IceProd) to coordinate data set management and result tracking. IceProd is a database-driven scheduling and management software package that catalogs simulation data sets and optimizes the usage of computing resources. Local support by the simulation production coordinators to resolve problems and incompatibilities of different systems is a critical task to maximize resource usage. Computer scientists will maintain configurations of the available resources and adapt to individual policies and restrictions of distributed production sites. As recommended by the SCAP, the system will incorporate third-party Grid middleware products to reduce long-term maintenance associated with an entirely in-house framework.

3.2.4 Environmental Compliance, Health and Safety

The ICNO follows all guidance as stated in the Comprehensive Environmental Evaluation (CEE) that was prepared by the Director of the Office of Polar Programs in 2004. ICNO management is mindful of the unique environment in the Antarctic and advises all USAP participants of USAP policy and proper and safe conduct. ICNO health and safety training is provided by the member institutions of the collaboration. ICNO has an excellent safety record and has operated at the South Pole without a reportable safety incident since November 2009.

3.2.5 USAP Logistical Support

The ICNO has successfully worked with the Antarctic Support Contractor (ASC, formerly RPSC) since 2003 to coordinate logistics resources required to construct and operate the IceCube detector. Requirements are detailed in the annual submission of the Support Information Package (SIP) each March. The yearly plan is finalized in September by ICNO concurrence with the ASC-generated Research Support

Plan (RSP). Further details can be found in the supplementary document "IceCube M&O: Logistical Requirements and Field Plan."

3.3 Budgeting, Staffing, and Workforce Management

This NSF IceCube M&O renewal proposal covers the five-year period of April 1, 2016, through March 31, 2021, at a total funds request of $35,000,000 (for more details, see the Budget Justification and Cost Overview in Attachment 3).

This proposal is for renewing the current NSF IceCube M&O award for UW–Madison, which began on October 1, 2010, and will end on March 31, 2016 (five and a half years, which includes a six-month extension), at a total amount of $38,143,749.

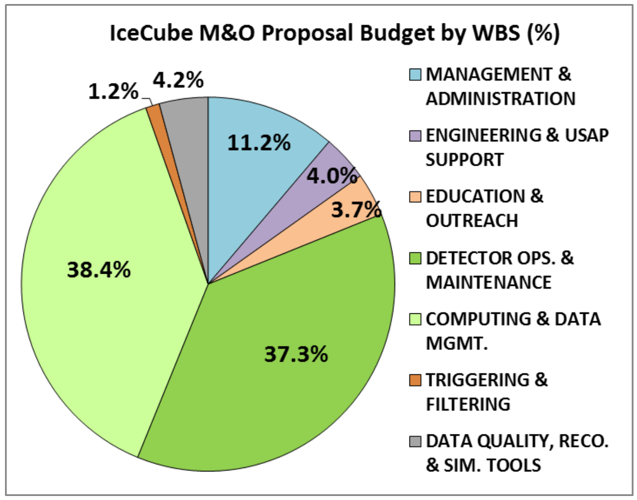

Figure 10: IceCube M&O proposal budget, WBS (%).

The budget in this proposal is based on a detailed, bottom-up analysis of the costs required to complete each task in the M&O Work Breakdown Structure (WBS) (included as Attachment 1); see Figure 10. These costs are very well understood and are based on actual experience during the past eight years of M&O. The two major WBS areas supported by this NSF M&O proposal are "Detector Operations and Maintenance" and "Computing and Data Management." This section identifies the tasks required to meet the technical requirements and specifications discussed in the previous sections.

3.3.1 Program Management and Financial Resource Coordination

Our approach focuses resources on achievement of IceCube's scientific objectives and provides accountability to NSF for taxpayer funds. The approach has four primary elements:

Operations Management and Science Support. We provide leadership to manage the effectiveness and efficiency of all services and ensure communication among the collaboration, NSF, partner funding agencies, and the M&O functions. We prepare strategic plans and conduct formal risk management to achieve objectives.

Computing Infrastructure Management. We manage computing resources to maximize uptime of all computing services and availability of required distributed services, including storage, processing, database, grid, networking, interactive user access, user support, and quota management.

Performance Management and Reporting. In cooperation with NSF, we establish performance measures that are meaningful to evaluating our performance against M&O objectives. We also establish with NSF a set of reporting deliverables that fulfill NSF internal and external requirements for oversight.

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

1. **NSF M&O Core (this proposal):** covers M&O core activities, travel, M&S and services for UW–Madison, 6 U.S. subawards, 1 UW shared grant, 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.
2. **NSF Base Grants:** support M&O activities done mostly by graduate students and postdocs, such as detector calibration, monitoring, filtering and triggering, data quality, reconstruction, and simulations.

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

4. **Europe and Asia Pacific In-Kind**: institutional contributions from non-U.S. collaborators, including labor, travel, and cash contributions to the M&O Common Fund.

IceCube M&O Common Fund (CF) was 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 PhD authors, which is updated twice a year at the collaboration meetings. Effective April 1, 2010, the annual established rate per PhD author is $13,650.

The M&O core activities identified as appropriate for support from the CF and agreed to be of common necessity for reliable operation of the detector and computing infrastructure include: winterover support at the South Pole; hardware and software systems for acquiring, filtering and transmitting data via satellite and disk to the UW data center; data archiving in the central data warehouse; and UW data center operations (as listed in the IceCube M&O Cooperative Agreement with NSF).

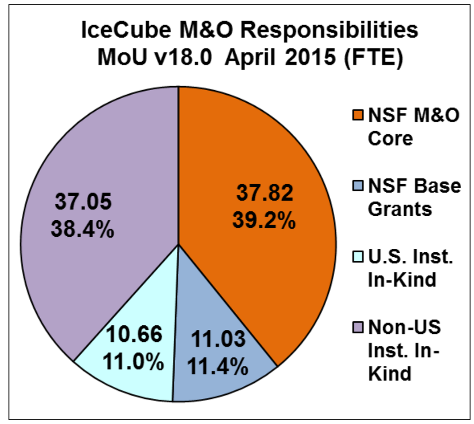

Figure 11: IceCube M&O distributed model, based on the institutional MoUs v18.0 (April, 2015).

Institutional M&O Memoranda of Understanding (MoUs) define the distributed M&O labor contribution from collaborating institutions and represent the transition from a centralized management and funding approach during IceCube's construction phase to a more distributed model for M&O (see Figure 11). The NSF M&O core award and NSF base grants provide about half of the human resources required for IceCube M&O, while the other half is provided by the IceCube Collaboration from U.S. and non-U.S. in-kind contributions.

The distributed model results in increased financial contributions to the Common Fund and in-kind labor contributions to M&O tasks from European and Asia Pacific collaborators. It also results in a greater emphasis on direct NSF funding to U.S. collaborating institutions. (Summary of the MoU Scope of Work is included as Attachment 2). The complete list of M&O activities with names and FTE level can be found in the Cost Overview in Attachment 3.

IceCube M&O Reports describe progress made based on objectives in the annual M&O Plan and differences between planned and actual accomplishments. These reports consist of a summary of work accomplished and include a financial section comparing annual budgets with actual cost and projected costs. Mid-year interim reports include a brief summary of the status of all M&O activities, including performance of data handling and detector systems. They also include highlights and accomplishments, detector uptime, software releases, and test results. A final report is required that includes a detailed summary of the entire period of the IceCube M&O award. In addition, a Common Fund (CF) annual report summarizes the status of past CF contributions and expenditures and lists the major annual upgrades to the SPS, SPTS, UW data warehouse, and UW data center.

3.3.2 Workforce Diversity

The IceCube M&O management organization is a beneficiary of the robust UW–Madison human resource 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 our vacancies. ICNO will continue to strive to attract outstanding candidates from underrepresented groups.

3.3.3 Key and Critical Personnel

Our key and critical personnel (tables below) 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.

Table 3: IceCube Key Personnel

Name, Position & Responsibilities	Qualifications
Francis Halzen, *PI* • Responsible for the overall success of the IceCube Neutrino Observatory	• UW Hilldale and Gregory Breit Distinguished Professor • Director of UW Institute for Elementary Particle Physics • IceCube Principal Investigator, from development, through construction and transition to operations
Kael Hanson*, Co-PI, Dir. of Operations* • Ensures operations meet established performance goals and the needs of NSF and the IceCube Collaboration	• Professor of physics, UW–Madison • ULB (Brussels, Belgium) Professor (Physics) 2009-2015 • IceCube In-Ice Devices Manager during construction • Former Detector Operations Coordinator
Albrecht Karle, *Co-PI, Assoc. Dir. for Science and Instrumentation* • Oversees technical performance and ensuring science objectives are met	• Professor and Chair of Physics Dept., UW–Madison • Associate Director for the IceCube construction project • Technical lead in AMANDA construction/operation • Scientific and technical lead for IceCube construction
James Madsen, *Assoc. Dir. for E&O* • Responsible for planning and executing of E&O activities	• Professor and Chair of Physics Dept., UW-River Falls • Chair of IceCube Speakers Committee, 2011-2015

Table 4: IceCube Critical Personnel

Name, Position & Responsibilities	Qualifications
Olga Botner, *IceCube Spokesperson* • Provides effective governance of the collaboration and coordinates member resources to support IceCube M&O	• Professor of experimental physics, Uppsala University • Experience in development of detectors for astroparticle physics (CERN, AMANDA, IceCube since 1998) • Contributor to IceCube Collaboration Governance Document
John Kelley, *Detector M&O Mgr.* • Manages detector M&O to provide consistently high detector availability and data quality	• Assistant Scientist, UW–Madison • 12 years of astroparticle detector development with IceCube, ARA, and Auger • Both scientific and industry management experience
Gonzalo Merino*, Comp. & Data Mgr.* • Manages computing services for M&O and data management policies	• Particle physicist experienced in scientific computing • Manager of a LHC Tier1 computing center for over 10 years • Participated in the development and roll out of major Grid projects in Europe since 2001
Paolo Desiati, *Sim. Prod. Mgr., Chair of IceCube Coordination Committee* • Oversees management of simulation production for the collaboration	• Senior Scientist, UW–Madison • Managed simulation production during construction • Coordination of M&O tasks in institutional MoUs with collaboration science requirements and operational needs
Allan Hallgren, *Trigger, Filter, Transmission (TFT) Board Chair* • Coordinates triggering and filtering and chairs TFT Board	• Professor of experimental physics, Uppsala University • Early contributions to IceCube, including leading DOM tests • Chair of Particle Physics Section, Swedish Physical Society • Speakers committee member (8 years)

Name, Position & Responsibilities	Qualifications
James Haugen, *South Pole Logistics / R&D Support* • Coordinates ASC support required for M&O activities at the South Pole	• Twelve years with IceCube and 7 deployments to South Pole • Over 30 years of successful technical and management contributions in academia and private industry • Instrumentation Manager for IceCube during construction
Catherine Vakhnina, *Resource Coord.* • Coordinates financial resources and interinstitutional contracts for the collaboration	• Five years of collaboration with IceCube institutional leads • Over 10 years of financial and project management experience in both science and industry • Holds an MBA and Project Mgt. Prof. (PMP) credentials
Alex Olivas, *Software Coord.* • Oversees software coordination for the collaboration	• Research Scientist, University of Maryland • Core IceCube software developer for over 10 years • Previous IceCube simulation lead • Previous IceCube Detector Simulation
Erik Blaufuss, *Analysis Coord.* • Responsible for coordinating physics analysis tasks and working groups within the collaboration	• University of Maryland Research Scientist • Technical lead for online filtering and offline software systems during IceCube construction • Former TFT Board chair and WG leader of GRBs/Transients

3.4 Science and Facility Planning

IceCube is a living instrument with diverse scientific objectives that change as scientific discoveries are made, both inside the IceCube Collaboration and in the larger scientific community. A nimble operations strategy ensures IceCube can respond to new discoveries and analysis techniques to maximize the science potential of the detector. This section presents a strategic vision, highlighting several of the recent and planned improvements to the ICNO operational structure that seek to improve the response time of the IceCube detector to interesting astrophysical events, increase scientific output of the collaboration, and maximize overall operational efficiency.

The core scientific objectives include neutrino point sources searches and measurements of the diffuse extraterrestrial neutrino flux, which require a stable, well-calibrated detector that continuously delivers high-quality data samples. Searches for transient sources such as supernovae or flaring active galactic nuclei require a high detector uptime to ensure IceCube is sensitive during these events and systems that ensure high data quality are available to initiate multimessenger campaigns with other observatories.

To more quickly identify astrophysical events of interest to both the collaboration and others, advanced analysis techniques are being refined and optimized. These improvements will provide real-time alerts to partner observatories and the wider physics community upon detection of astrophysical neutrino candidates, such as those from the high-energy starting event (HESE) selection that enabled IceCube's discovery of astrophysical neutrinos [2]. By implementing these selections in the online filtering system, and upgrading our satellite messaging and alert systems, we can coordinate with optical, gamma-ray, and radio observatories for follow-up observations with delays of about one minute. Such multimessenger searches may allow the first identification of the sources of the observed high-energy neutrino flux.

Ensuring that IceCube is continuously producing high-quality data has driven the effort behind improvements to the IceCube DAQ and IceCube Live systems. The DAQ has focused on maximizing system uptime by supporting "stopless" operation modes, resulting in detector uptime typically exceeding 99%. With planned DAQ system improvements to the hitspooling system, allowing DOM data to be collected independently of DAQ triggers, detector uptime will continue to improve. The IceCube Live monitoring effort has unified all data-quality monitoring information within a single system, allowing for efficient identification of concerns and rapid resolution of issues. Work planned for IceCube Live will incorporate additional measures and alert mechanisms to further streamline data-quality monitoring and ensure that alerts sent to external observers are based on high-quality neutrino events.

All software and operational changes to the IceCube detector systems are verified using the SPTS test system, located at the UW–Madison data center. To increase the reliability of these tests and correspondingly the quality of data from the SPS, we will improve system testing strategies with end-to-end integrated system tests at SPTS using real, prerecorded data, played back on the test system, and use simulated data with known, programmed physics inputs as test data for evaluation of DAQ and online software.

In order to quickly produce high-quality scientific results, efficient organizational approaches, backed by ongoing software development, are needed to streamline the flow of data from raw events collected at the South Pole to publication of analysis results. As an example, the coordination of offline data processing by the TFT Board, combined with real-time monitoring by IceCube Live, has sped up the availability of reconstructed Level 2 data from 13–15 months to 2 weeks. This work will continue to streamline advanced Level 3 reconstructions and production of required simulated data samples.

As the data collected keeps growing and the analyses mature, a deeper understanding of systematic effects increases in importance. One necessary ingredient to address these uncertainties will be larger amounts of simulated events. This requires continued work to improve software efficiency and improvements to the simulation production system to maximize the simulation generation capacity within our budget. IceCube will exploit the usage of external resources such as some of the GPU-capable supercomputers in XSEDE or similar HPC networks.

The IceCube detector produces more the 300 TB of data per year. Storage, transportation, and cataloging of this data are critical for maximizing the scientific output of the detector in a cost-efficient manner. In 2014, the IceCube data archive system at the South Pole switched from using tape technology to hard drives. This change solved several operational issues associated with tape media in the low-humidity environment of the South Pole and enabled efficient ingest of all data into the data warehouse in the north. Additionally, with the new disk-based archive system, the data generated at the Pole will be copied every year to two different locations that will store it in state-of-the-art automated tape libraries. This will ensure the preservation and reusability of this data for the entire duration of the project and beyond. Additionally, IceCube will continue efforts to improve public access to its data products, including addressing discoverability, reuse, and preservation. A catalog service will be developed to enable efficient metadata searches, and final data products from publications will be made publicly available with clear, persistent identifiers to ease accessibility and reusability over time.

Accumulation of snow on the IceTop surface detector reduces the detection efficiency for lower-energy cosmic-ray air showers. Managing this snow accumulation has proven burdensome in effort and cost for the South Pole contractor and has been discontinued. We are investigating installation of low-cost scintillator panels over buried IceTop stations to restore the efficiency of IceTop. The scintillators can be connected to the ICL and integrated into the DAQ, and prototypes can use existing IceCube cabling.

The power, network, and satellite communications at the ICL have supported other experiments in the Dark Sector, such as DM-Ice and the Askaryan Radio Array (ARA). Operation of these experiments does not fall under the purview of IceCube. However, careful maintenance and upgrades of our hardware and software ensure that the ICL facility will continue to benefit the wider South Pole science community.

This strategic vision for IceCube will enable efficient operations that supply increased scientific output and robust real-time response to interesting astrophysical events with targeted improvements to the operational systems. These will position IceCube to be a world-leading neutrino observatory for the next five years and beyond.

3.5 Education and Public Outreach

The associate director for E&O leads the ICNO E&O effort, with guidance from the IceCube PI, the IceCube Collaboration Board, the ICNO E&O team, and an external advisory board. The proposed E&O program will build on its established successes (as mentioned in NAS's 2015 strategic vision report [18]) to address NSF priorities to integrate research and formal and informal education and to reach audiences that are underrepresented in STEM fields and communities. We will also increase the number of collaboration-wide E&O programs to leverage resources at 45 institution in 12 countries in the US, Europe, and Asia to capitalize on the potential of the ICNO to excite and engage the broader community.

While the E&O team will continue to seek new opportunities and develop new partnerships to best use available resources and personnel, the following will be areas of focus for the 2016-2020 period:

1) *Reaching high school students and teachers through IceCube Masterclasses and K-12 classrooms through webcasts from the South Pole.* Masterclasses are one-day events held at IceCube Collaboration campuses that give high school students and accompanying teachers a chance to experience real research using IceCube data. The second annual IceCube masterclass, held in March, 2015, doubled the number of participating institutions from five to ten, and reached 175 students. For the 2014–15 season, five South Pole webcasts were held in English, German, and Spanish, reaching 1,600 people in seven countries. **Goals:** Increase masterclass participation to 250 students annually and increase diversity by having one taught in Spanish at a U.S. site; annually review and revamp masterclass resources based on user evaluations; and hold 10 webcasts per year, with at least one in Spanish, averaging 100 participants and attracting diverse audiences.

2) *Increasing STEM awareness through undergraduate research experiences and South Pole deployments for high school teachers who are integrated into the UWRF Upward Bound (UB) program.* Undergraduate research programs increase academic success and build transferrable skills, especially for underrepresented students. Most collaborators provide undergraduate research opportunities, such as working with large data sets, in addition to UWRF's NSF-funded Research Experience for Undergraduates and International Research Experiences for Students programs at UWRF. The ICNO has hosted four high school teachers at the South Pole jointly with the NSF-funded PolarTREC program, the last a native Spanish speaker from Puerto Rico. Following their deployment, these teachers brought their experience to the eight-day UWRF UB summer math and science enrichment program that has served over 400 low-income, predominately first-generation high school students over the last 10 years. **Goals:** Host at least one teacher at the South Pole annually (at least one from an underrepresented group during the 2016–2020 period); continue collaborating with past IceCube PolarTREC teachers to bring ICNO science to the classroom; hold dedicated virtual meetings for undergraduate researchers and increase their visibility through poster prizes at collaboration meetings; and enhance diversity by committing to having a majority of UWRF REU and IRES students from underrepresented groups and/or two year colleges.

3) Producing captivating web and print resources, graphic designs, and displays. To support NSF's strategic objective to "build the capacity of the Nation to address societal challenges using a suite of formal, informal, and broadly available STEM educational mechanisms," we will work with creative professionals to craft multimedia resources that describe IceCube science and technology in an accessible way. Ongoing efforts include creating inviting web content, designing and producing graphics for E&O programs and events, and developing portable scale models of the detector, with LEDs to display data. **Goals:** Annually assess effectiveness of formal and informal learning web resources through user experience groups of different ages and lay audiences attending IceCube outreach events; and add data-driven sound to detector models and assess effectiveness of models in explaining ICNO data.

4) *Building internal E&O capacity by developing and implementing semiannual communication skills workshops in conjunction with IceCube Collaboration meetings.* The ICNO, with its extreme location, and exotic science, is a captivating topic that attracts wide audiences. Still, it takes training to effectively communicate science and technology in accessible language appropriate for an intended audience. The first communication skills workshop was held for early career IceCube researchers at the spring 2015 collaboration meeting in Madison and the second is scheduled for fall 2015 at the collaboration meeting in Copenhagen. The semiannual collaboration meetings, one in North America and one in Europe, are ideal opportunities to engage with international colleagues and build communication skills in a supportive environment. **Goals:** Increase the capacity of communication workshops to engage 50% of young/early

career IceCube researchers; offer a biennial communication workshop for mid- and late-career IceCube scientists; assess the effectiveness of workshops through participant surveys and the outcomes of the participants' broader impacts efforts and make improvements as needed.

4 Broader Impacts of the Proposed Work

The 2016–2020 ICNO broader impacts (BI) program will build on the success of our previous BI program and the expertise developed during our 15-year history. We will continue promoting an extensive E&O program and facilitating access to IceCube data for other communities and disciplines. We will also further develop our multifaceted communication plan, and we will strengthen our efforts in the development of a diverse and skilled STEM workforce. Finally, new efforts will allow integrating ICNO computing resources to the Open Science Grid (OSG), contributing to the largest distributed high-throughput computing infrastructure for research in the US.

Details about specific BI activities and goals are given below. In addition, Section 3.5 describes efforts targeting teachers and learners of all ages, the postdoctoral mentoring plan outlines our program to help postdoctoral researchers excel in their careers, and the data management plan includes an open data access policy and accompanying activities.

Specific metrics and outcomes are detailed for each activity. We will design and implement an evaluation program to increase the validity of assessments with help from an evaluation expert, who will work closely with the BI lead team. A mixed-methods design will be used, including data triangulation and qualitative and quantitative data sources.

- *Contributing to Astrophysics Beyond IceCube*

Strategies to enhance the discovery potential of IceCube include neutrino alerts to the international astrophysics community—Astronomer's Telegram (ATEL) service, the SuperNova Early Warning System (SNEWS), and the Astrophysical Multimessenger Observatory Network (AMON)—and targeted collaborations with other experiments.

The unique blend of technology and physics backgrounds of the ICNO team has allowed us to reduce the delay between detection at the South Pole and the transmitted alert to about two minutes. By 2017, ICNO will reach near-real-time alerts, with delays under a minute. We will also add extremely high energy and high-energy starting neutrino data streams to the current time and angular coincidences alerts.

IceCube has signed agreements with 10 neutrino, cosmic-ray, gamma-ray, and gravitational wave detectors to develop a multimessenger research program (see Attachment 6). Over the last two years, these collaborations resulted in six joint studies, including the coincidental discovery of a supernova. We will expand these relationships to extract the full physics potential of the data, and increase the number of joint studies in search of new discoveries to improve of our understanding of the extreme universe.

- *Interdisciplinary Research with IceCube Data*

Event reconstruction requires a thorough characterization of the properties of South Pole ice, which resulted in a series of papers submitted to glaciology journals. In 2013, a paper published in the *Journal of Glaciology* showed that impurities in the ice, a complication for the detection of neutrinos, could reveal vital stories about past climate changes. ICNO will strengthen its collaboration with glaciologists, especially with researchers from the SPICE core and West Antarctic Ice Sheet (WAIS) Divide projects, both funded by NSF, to improve ice uncertainties in IceCube and develop joint glaciology studies.

- *Sharing Computing Resources and Knowledge with Researchers in the US*

Since 2012, the ICNO has increasingly been using the OSG infrastructure. By 2020, the integration of IceCube computing and storage facilities (excluding those at the Pole) will leverage ICNO resources while contributing to a multidisciplinary infrastructure for research. Accounting metrics for computing resources as well as the number of IceCube analyses benefiting from OSG and the number of external scientists using the ICNO resources will be used to assess this effort.

IceCube also provides pioneering feedback on the use of GPU workload to the HTCondor software program, which is at the core of the OSG. We will continue working with the HTCondor team of developers at UW–Madison to fully develop the potential of GPUs in the OSG, integrate all ICNO GPU resources in the OSG and facilitate a growing number of OSG GPU sites.

- *Fostering a Strong, Diverse, and Skilled STEM Workforce in the US*

A diverse community of hundreds of skilled scientists, technicians, engineers, drillers, IT technicians, managers, and communication and outreach specialists enabled the ICNO to become the largest particle detector and a pioneering tool for astronomy. This community, headquartered at UW–Madison, is spread throughout the US, Europe, and Asia.

The ICNO continues to be fully committed to the development of a strong, diverse, and skilled STEM workforce in the US. Some efforts in this regard are described in Section 3.5 and are accompanied by a) UW–Madison's Affirmative Action Plan, a positive effort to assure that women and minorities are not underrepresented in the university's workforce, b) ICNO research opportunities for undergraduate students in physics, engineering and computing programs at UW–Madison, including members of underrepresented groups, c) WIPAC's support, through institutional funds, for faculty from collaborating HBCU institutions, and d) increased collaboration with associations serving underrepresented audiences (see section below). A steady increase in the number of staff and collaborators as well as of participants and audience members from underrepresented communities will mark success of these efforts.

- *Engaging All Audiences*

ICNO communication efforts, in support of the collaboration, have been continually growing and currently reach audiences across the US and around the world. We will continue disseminating information about IceCube and our activities through press releases, news articles, and other means to reach this growing and increasingly diverse audience. In the next five years, we will double our website and social network audience, currently reaching about 5,000 people per week, while increasing the range of their geographical origins and backgrounds. We will increase the production of content related to non-research activities, such as detector operation, to show the multidisciplinary and multicultural team and expertise behind IceCube. We will also create engaging multimedia resources and develop multilingual materials, especially in Spanish. At least two new multimedia resources will be produced each year, and we will seek funds outside the budget of this proposal to create others. A native Spanish speaker in the M&O core team will create content in Spanish with extra support from other UW–Madison personnel. Other IceCube collaborators will also contribute content in German and possibly other languages.

Finally, we will further strengthen our collaboration with organizations representing underrepresented communities, such as the National Society of Black Physicists, the Association for Women in Science, and the Society for Advancement of Chicanos/Hispanics and Native Americans in Science. We will work with them to increase the visibility of IceCube in their communities through articles in their media and our participation in their training and outreach programs. Our final goal is the recruitment of at least two staff, faculty, or students every year from underrepresented communities.

- *A Growing Team to Lead ICNO BI Efforts*

The ICNO director, PI and IceCube spokesperson lead these BI efforts. A core team coordinates communication and education and outreach efforts, but over 100 staff and IceCube collaborators have contributed to or led ICNO-based BI projects, either in research or operational activities. These do not include local BI efforts at the 45 IceCube institutions. The goal for the next five years is to increase this participation to about 175 people, while increasing the number of programs that coordinate collaboration-wide efforts from 8 to 12.