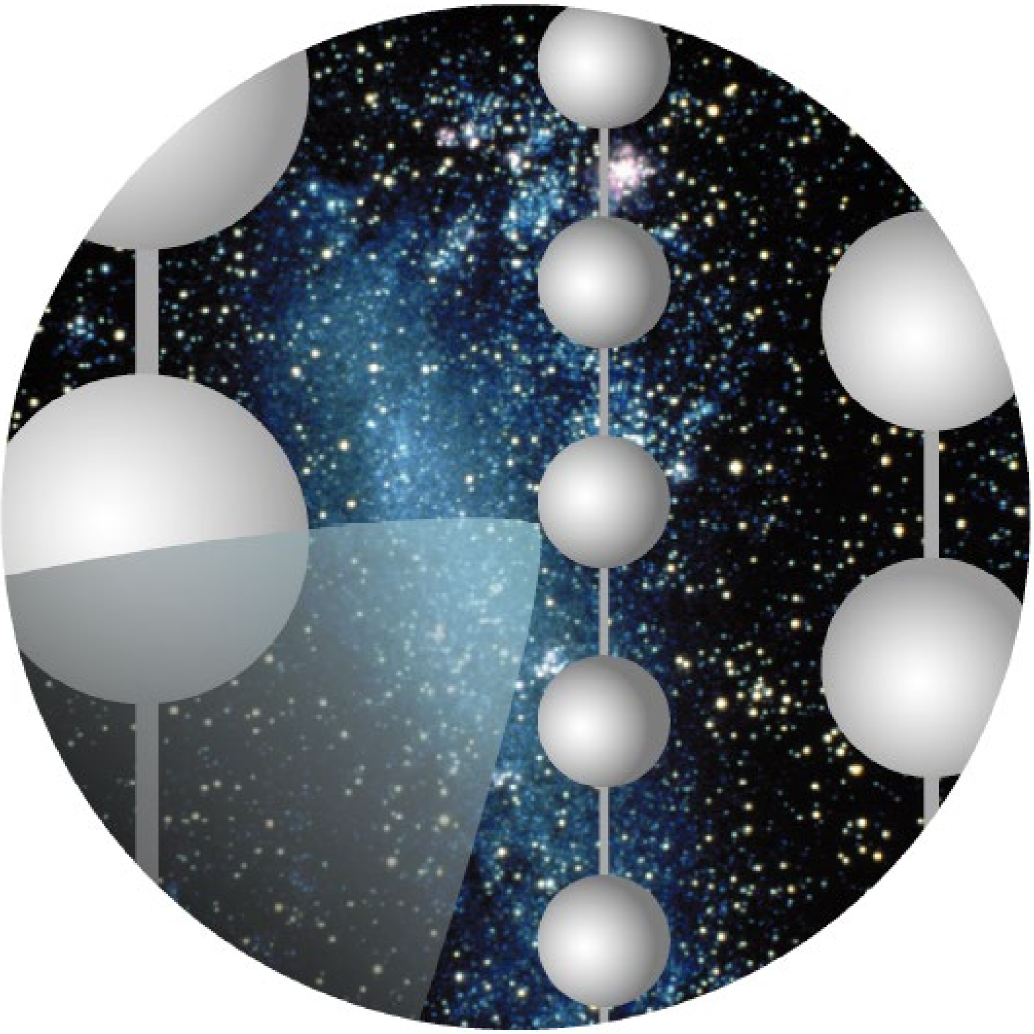

IceCube

Neutrino Observatory

Management & Operations Plan
December 2019
Revision 4.0

IceCube

MANAGEMENT & OPERATIONS PLAN

SUBMITTED BY:

Francis Halzen
IceCube Principal Investigator
University of Wisconsin–Madison

Kael Hanson
Co-PI and IceCube Director of Operations
University of Wisconsin–Madison

Albrecht Karle
Co-PI and Associate Director for Science and Instrumentation
University of Wisconsin–Madison

James Madsen
Associate Director for Education and Outreach
University of Wisconsin–River Falls

Catherine Vakhnina	Resource Coordinator
Paolo Desiati	Coordination Committee Chair
John Kelley	Detector Operations Manager
Benedikt Riedel	Computing and Data Management Services Manager
Juan Carlos Diaz-Velez	Data Processing and Simulation Services Manager
Alex Olivas	Software Coordinator
Summer Blot, Keiichi Mase	Calibration Coordinator

Revision History

Revision	Date Revised	Section Revised	Action
1.0	06/9/2016		First version
2.0	05/31/2017		Updated for M&O PY2
3.0	12/12/2018		Updated for M&O PY3
4.0	12/##/2019		Updated for M&O PY4

Table of Contents

List of Acronyms and Terms

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

ICL	IceCube Laboratory (South Pole)
IOFG	International Oversight and Finance Group
JADE	Java archival and data exchange
LED	Light emitting diode
M&O	Management and operations
M&OP	Management & Operations Plan
mDFL	Mobile/modular dark freezer lab
MoU	Memorandum of Understanding, between UW–Madison and all collaborating institutions
MPS	NSF Directorate for Mathematical & Physical Sciences
MREFC	Major Research Equipment & Facilities Construction
MSPS	Megasamples per second
OPP	NSF Office of Polar Programs
OSG	Open Science Grid
PA	NSF Particle Astrophysics Program
PCTS	Physical Sciences Laboratory cable test system
PHY	NSF Division of Physics
Physics working group	Physics working groups perform high-level analysis and develop specific analysis tools
PLR	NSF Division of Polar Programs
PMT	Photomultiplier tube
PnF	Processing and filtering
SAC	Science Advisory Committee
SCAP	IceCube Software & Computing Advisory Panel
Science DMZ	A secure computer subnetwork designed for high-volume data transfers
SIP	Support Information Package
SNEWS	Supernova Early Warning System network
SNDAQ	Supernova data acquisition
SPS	South Pole System (at the South Pole)
SPTR	IceCube dedicated high-capacity South Pole TDRS relay system.
SPTS	South Pole Test System (at UW Madison)
TDRS	A Tracking and Data Relay Satellite is type of a satellite that forms part of the TDRSS
TDRSS	The Tracking and Data Relay Satellite System is a network of communications satellites
TFT Board	Trigger Filter and Transmission Board
TS	Test statistic
UPS	Uninterruptible power supply

USAP	United States Antarctic Program
UW	University of Wisconsin–Madison, host institution of the IceCube Collaboration
VCRGE	Office of the Vice Chancellor for Research and Graduate Education, at UW–Madison
WBS	Work breakdown structure
WIMPs	Weakly interacting massive dark matter particles
WIPAC	Wisconsin IceCube Particle Astrophysics Center (former IRC)
XSEDE	Extreme Science and Engineering Discovery Environment

1 Preface

In the past few years, the burgeoning field of multimessenger astronomy has moved from theory to observations beginning with the joint observation of gravitational waves and EM emission from a black hole merger in 2015. Two years later, IceCube alerted the community to a single high energy neutrino event which led to multimessenger observations in EM and neutrinos of the flaring blazar TXS0506+056. This single object has revealed a wealth of information on a class of sources that may produce a significant fraction of the UHE cosmic rays. The observations also pose a number of questions. IceCube science continues to expand beyond astrophysics into neutrino particle properties: measurements of neutrino mixing parameters and constraints on sterile neutrinos are competitive with or exceed accelerator measurements for certain regions of parameter space. IceCube is responding to these new demands on both technical and social fronts. New institutions are joining the IceCube Collaboration, either in the capacity of associate members to work on specific analysis topics or as full members interested in contributing resources to the facility to exploit a broader range of science. New MoUs are being produced between IceCube and other astrophysical observatories to exchange archival and real-time data. The M&O organization and the Collaboration together have made enormous advances in better understanding how the detector responds and how to better use the detector for both existing and novel scientific aims.

Realization of our scientific vision now more than ever requires the attention of a highly technically competent and dedicated team to oversee the reliable operation of the IceCube Neutrino Observatory facility. Our approach to the planning and execution of IceCube management and operations is based on nearly a decade of experience, over which time the combined teamwork of the centralized M&O organization WIPAC and the IceCube Collaboration has streamlined the process. Tasks range from detector hardware and firmware maintenance to characterization of the ice optics; from production of massive simulated data sets to organizing outreach activities and communicating scientific results to the press. Each aspect of the M&O is proven through extensive use to maximize the facility's scientific and educational potential.

This Management & Operations Plan (M&OP) describes management, roles and responsibilities, lines of authority and communications, critical or significant project activities, and performance objectives and milestones for the second plan year of the current cooperative agreement (April 2017 – March 2018). The M&OP identifies the budget allocation of the various funding sources, including the direct NSF funding provided through this award and the Common Fund.

Section 2 reviews the scientific vision and objectives that IceCube is designed to achieve and provides a timeline of key milestones. Section 3, Technical Approach, specifies the M&O requirements necessary for IceCube to achieve its design objectives. Section 4, Management Approach, identifies the tasks required to meet the technical requirements and explains how we will perform them. Section 5, Cost Overview, provides a breakdown of costs by funding source.

2 Achievement of Scientific Vision

Recently, and only eight years since IceCube's completion, we have published our most important results to date: the discovery of the first source of high-energy neutrinos. These results, which are further explained below, were made possible by the efforts of several teams within our collaboration, including important contributions of our team, as well as partnerships with several telescopes across the electromagnetic spectrum.

Furthermore, these findings have confirmed cosmic neutrinos as essential astronomical messengers for revealing an unobstructed view of the universe at wavelengths where the universe is opaque to light. The large neutrino flux observed shows an energy density of neutrinos that matches the energy density of photons in the non-thermal universe. The prospects for astronomy are extraordinary, with IceCube observations indicating a more prominent than anticipated role of proton acceleration relative to electrons.

At the same time, the detector, now operating in the 5 GeV to 10 PeV energy range, has already achieved a performance that is significantly superior to what had been expected, with a neutrino collection area that is larger by a factor of 2 to 3, depending on the energy, and an angular resolution of high-energy muon tracks that is less than 0.5 degrees. We have implemented new methods for energy measurement that is on the order of 10% for particle showers.

2.1 Neutrino Emission from the First Identified Cosmic Ray Accelerator

The writing of this report coincides with yet another moment of great excitement in IceCube and within the international multimessenger community. Only six years after the discovery of a flux of astrophysical high-energy neutrinos, and two years after announcing the identification of the first source of high-energy neutrinos, detailed imaging by radio interferometry of the “blazar” TXS 0506+056 revealed a merger of two galaxies instead. Recently, the 300-TeV neutrino IC190830 pointed at a similar source. These new results are made possible by the efforts of several teams within our collaboration, including important contributions by UW–Madison as well as partnerships with several telescopes across the electromagnetic spectrum.

Astrophysical neutrinos unambiguously trace protons and nuclei throughout the universe. The discovery of high-energy extraterrestrial neutrinos with a flux at the very high level of those expected from cosmic-ray accelerators brightened the prospect for identifying the sources. The original observations specialized to neutrinos originating inside the detector, mostly showers initiated by electron and tau neutrinos. After that, the conventional method of selecting upgoing muon tracks delivered a sample of muon neutrinos with similar statistics but superior angular reconstruction. Though limited to muon flavor, this analysis represented a watershed towards the identification of the sources. It revealed secondary muon tracks that deposited several-hundred-TeV of energy inside the detector, indicating parent neutrinos exceeding an energy of 10 PeV. Both analyses for identifying cosmic neutrinos were designed and initially performed at UW–Madison. Increasingly, the data suggests that an excess of events is observed below 100 TeV over the rate obtained by extrapolating a power-law fit to the data above 100 TeV.

After ten years of data, the observed neutrino flux is no longer consistent with an isotropic distribution of the arrival directions: evidence emerges for an asymmetry in the sky map associated with four extragalactic sources, including TXS 05060+056. For the first time, the hottest spot in the map coincides with the most significant source in a list of candidate sources that is fixed prior to unblinding the data. This source is NGC1068, a relatively nearby active galaxy.

We observe equal contributions to the astrophysical flux of the three neutrino flavors, consistent with the observation of extragalactic sources whose flux has equilibrated in the three flavors after propagation over cosmic distances. This oscillation measurement at PeV energy provides a powerful test of the three-flavor scenario, which will be performed with improved precision after the deployment of the IceCube Upgrade. In a variety of analyses, we searched for a possible correlation of the cosmic events with the plane of the galaxy. The UW–Madison group performed a joint analysis of the IceCube and HAWC data and, though we find results consistent with the prediction for the MGRO 1908 and 1853 sources, the neutrino statistics are still too low to reach 3σ evidence.

The production of PeV neutrinos is inevitably associated with the production of PeV gamma rays. Hadronic accelerators produce fluxes of neutral and charged pions that are the parents of gamma rays and neutrinos, respectively. Gamma rays are attenuated over cosmic distances by the extragalactic background light. PeV gamma rays are predominantly absorbed by pair production in interactions with cosmic microwave and infrared background photons. They cascade down in energy to a diffuse gamma-ray flux in the GeV-TeV energy range. It is intriguing that this flux provides an excellent match to the extragalactic gamma-ray flux observed by the Fermi satellite. This implies that most, possibly all, of the energy in the nonthermal universe originates in hadronic accelerators. The matching relative magnitudes of the diffuse extragalactic gamma-ray flux detected by Fermi and the high-energy neutrino flux measured by IceCube may also suggest that they originated in common sources.

A better understanding of the neutrino flux and of the IceCube detector has also allowed us to develop new methods to identify the sources and build on the discovery of cosmic neutrinos to launch a new astronomy. During the last few years, the IceCube multiwavelength program has grown from simple supernova alerts and offline follow-ups with a few observatories to a more elaborated set of high-energy neutrino alerts. Since April 2016, IceCube is able to alert the international community of the observation of a very high energy neutrino within a minute of its detection deep in Antarctica's ice. After the tenth such alert, of a 290-TeV neutrino detected on Sept. 22, 2017, about 20 observatories around the world performed follow-up observations. The detection of high-energy gamma rays as well as radiation from radio waves to X-rays associated with this neutrino presented strong evidence for a known blazar as the origin of the neutrino. And the analysis of archival IceCube data around this location identified a more spectacular neutrino flare that confirmed the discovery.

Subsequently, the IceCube archival data revealed a three-month burst of 13 cosmic neutrinos in 2014-15 that dominates the neutrino flux of the source over the 9.5 years of observations. The original identification of the source as a blazar was puzzling because it required a major accretion event onto the rotating supermassive black hole to provide the target material to produce the 2014 neutrino burst. Subsequent high-resolution radio images of the source with the Very Long Baseline Array (VLBA) brought to light a merger of two galaxies, revealed by the interaction of two jets entangled in the source. Recently, the blazar PKS 1502+106 was found in the direction of a 300-TeV neutrino alert, IC-190730. Owens Valley Radio Array (OVRA) radio observations at 15 GHz indicate that the neutrino also coincides with the highest flux density of a flare that started five years ago. This matches the similar long-term outburst seen from TXS 0506+056 and may indicate merger activity. Also, the dominant hotspot in the 10-year IceCube neutrino sky map, NGC 1068 (Messier 77), is a Seyfert galaxy undergoing a major accretion event onto the black hole observed by the ALMA array. A few-percent fraction of such special sources, now labeled as gamma-ray blazars, is sufficient to accommodate the diffuse cosmic neutrino flux observed by IceCube.

The detection of the first source of high-energy neutrinos and the first evidence for a source in the sky map underscore the case for a next-generation detector, one that will instrument a volume ten times that of IceCube, to advance neutrino astronomy to an era of precision measurements that will provide a totally new view of the extreme universe.

2.2 Cosmic Neutrinos and the Impact of IceCube Results

To date, IceCube has isolated more than 100 high-energy cosmic neutrinos, with energies between 100 TeV and 10 PeV, from more than a million atmospheric neutrinos and hundreds of billions of cosmic-ray muons. In order to filter out this huge atmospheric background, our searches for astrophysical neutrinos focus on high-energy events that start in the detector or that originate in the Northern Hemisphere. The large neutrino flux measured by IceCube also implies that a significant fraction, possibly all, of the energy in the nonthermal universe is generated in powerful hadronic accelerators powered by objects such as black holes or neutron stars.

A study of the neutrino arrival directions reveals that the observed neutrino flux is consistent with an isotropic distribution indicating an extragalactic origin. Several searches for a Galactic component and Galactic sources are ongoing. 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. For cosmic sources, the gamma rays cascade in the microwave and infrared background and reach Earth with reduced energies that are conveniently within the sensitivity range of the Fermi satellite. Figure 2.2-1 shows the gamma-ray flux accompanying an $E^{-2.15}$ neutrino spectrum with an exponential cutoff around a PeV. While it describes the IceCube data only on average, the gamma-ray spectrum, after propagation, matches the extragalactic isotropic diffuse gamma-ray background measured by Fermi. We here assumed equal production of pions of all three charges in the cosmic beam dump. This exercise indicates that the contribution of gamma rays accompanying IceCube neutrinos to Fermi's extragalactic flux is significant, suggesting a common origin of some of the sources at some level. In fact, as mentioned in the previous section, in July 2018, IceCube, gamma-ray telescopes Fermi and MAGIC, and several other experiments announced the detection of neutrinos and photons from blazar TXS 0506+056. These results constitute the first-ever identification of a likely source of extragalactic neutrinos and of high-energy cosmic rays.

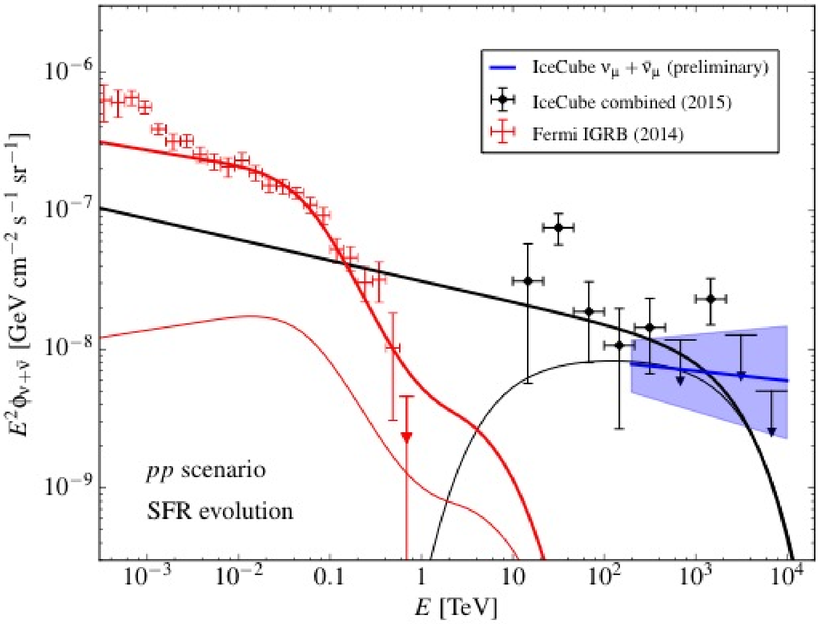

Figure 2.2-1: *The astrophysical neutrino flux (black line) observed by IceCube and the corresponding gamma-ray flux (red line) observed by Fermi after cascading through the extragalactic background light. The calculation assumes that the decay products of neutral and charged pions from hadronic interactions are responsible for the non-thermal emission in the Universe. The black data points are obtained from neutrino events starting inside the detector. Also shown is the best fit to the flux above 220 TeV of high-energy muon neutrinos penetrating the Earth.*

2.3 IceCube Science Beyond Cosmic Neutrinos

When considering the goals of detector operations, it is important to recognize that IceCube data impacts science beyond neutrino astronomy. The science reach of IceCube is illustrated in Figure 2.3-1. Every box in the diagram is covered either by publications or by an ongoing analysis.

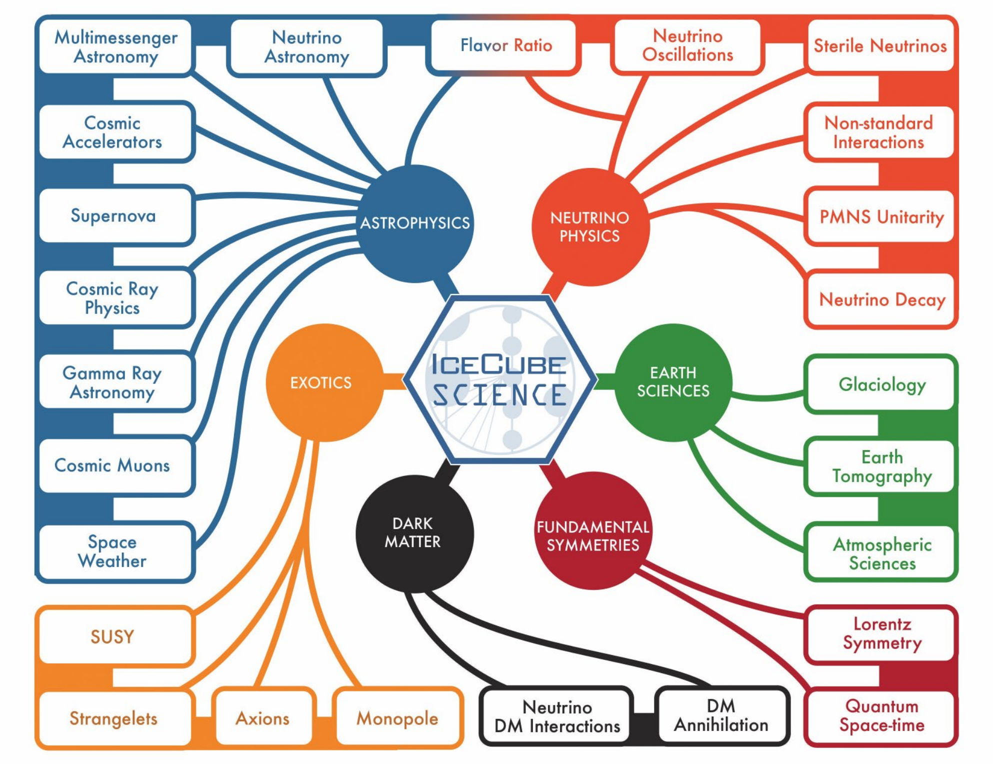

Figure 2.3-1: *Research topics covered by IceCube working groups.*

Most prominently, with only three years of observations of the atmospheric neutrino beam with the DeepCore subsystem, we are measuring neutrino oscillations reaching a precision in the range of accelerator results; see Figure 2.3-2.

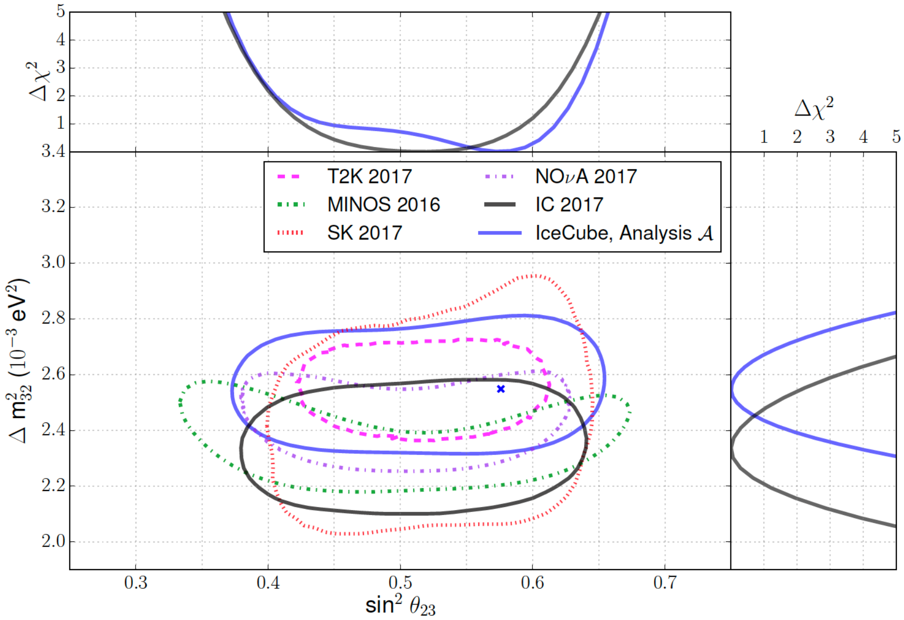

Figure 2.3-2: *Oscillation parameters measured with three years of DeepCore data (2012-2014). Two independent analyses have been performed based on high statistics and high-quality event reconstruction, respecitevely. The 90% allowed region using the data sample from Analysis A in blue compared to other experiments. The best fit point from Analysis A is shown as the blue cross mark. The IceCube 2017 result shown in the previous annual report, represented in black, uses the data sample from Analysis B. The top and right plots are the 1-d $\Delta\chi^2$ profiles of the measured oscillation parameters. The measurement is performed in the energy range of 5-55 GeV, one order of magnitude above the energy range of previous experiments creating an opportunity for detecting new neutrino physics.*

IceCube also has two new searches for light sterile neutrinos using high-energy events. One of these searches focuses on eV-scale sterile neutrinos motivated by the persistent LSND and MiniBooNE anomalies and the other focuses on heavy sterile neutrinos motivated by neutrino mass generation models. Both of these searches yield leading results in the search for sterile neutrino. The result of the eV-scale sterile neutrino is shown in Figure 2.3-3 as a function of the sterile neutrino mixing, on the horizontal axis, and the mass-squared difference, on the vertical axis. The top panel shows the result at the 90% CL as a closed dashed contour and the 99% CL as a solid black line. Points inside the dashed contour are preferred over the no-sterile hypothesis, and points to the right of the solid line are disfavored at the 99% CL. The lower panel compares our result to global analyses of the neutrino data. The parameter space preferred by the MiniBooNE and LSND anomalies is shown in light blue and is excluded by our analysis. Parameter points preferred by disappearance experiments, such as the recent DANSS and NEOS signatures, are shown in light green and are not disfavored by our analysis.

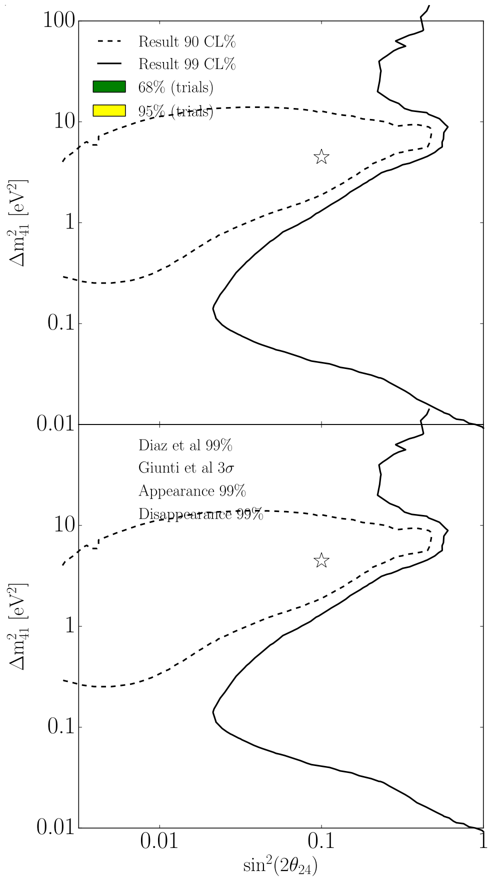

Figure 2.3-3: *Results from the IceCube search for light sterile neutrinos using high-energy atmospheric neutrinos. The 90% (dashed solid line) CL contour is shown with bands containing 68% (green) and 95% (yellow) of the 90% contours in simulated pseudo-experiments, respectively. We also show the 99% CL contour as a solid line. Preferred*

regions for the global neutrino data set are shown as green and red closed contours in the lower panel. These are also shown separated for appearance (light blue contour) and disappearance (light green area).

Our second analysis searches for heavy sterile neutrinos and is complementary to previous searches performed with low-energy neutrinos in IceCube. These constrain the muon, $|U_{\mu4}|^2$, and tau neutrino, $|U_{\tau4}|^2$; flavor content of the heavy neutrino mass state. These mixings have been parameterized in terms of mixing angles. Figure 2.3-4 shows the result of our analysis at the 90% CL as a solid gold line, where the horizontal mixing angle is related to the muon flavor component and the vertical mixing angle to the tau component. Previous results from IceCube-DeepCore are shown as a dotted line, and results from SuperKamiokande are shown as a dashed line; these results were previously the strongest constraints of these parameters. The best-fit point of this analysis is marked as a star and is compatible with the no-sterile neutrino hypothesis. Our result improves on the previous results available in the literature.

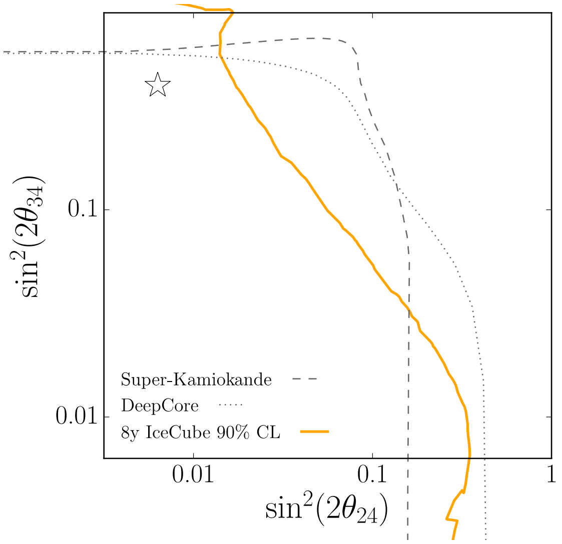

Figure 2.3-4: *Results from the IceCube search for heavy sterile neutrinos using high-energy atmospheric neutrinos. The 90% (golden solid line) CL contour. Other experiments are shown as dashed and dotted lines***.**

2.4 Five-Year Roadmap

The previous five-year roadmap focused on the transition of the ICNO facility from a construction phase to one of stable operations. During this successful transition, not only has IceCube discovered a signal of high-energy neutrinos clearly standing out from the spectrum of atmospherics but it has also found the first evidence for a source of these neutrinos. This enormously successful result nevertheless places new exigencies on the facility operations. While the operations team fully intends to continue to improve on the methodology which helped bring about this success, the current five-year period will be much more than a mere continuation of the previous one. The sharing of data products with the community in the form of bidirectional real-time alerts and access to archived data has expanded and will continue to expand as new MoUs are signed with other astrophysical observatories. New algorithms for identification and reconstruction of the high-energy neutrinos are being developed to improve the number and quality (e.g., angular position, energy, and lepton flavor) of the samples. And then of course there is the multitude of ideas that may be forming in the minds of collaboration members at this very moment, or may have not yet formed. Fully exploiting the observatory's science potential means that the facility operations team must be prepared to respond to as-yet-unknown requests for new detector configurations, ice measurements, upgrades, and data products.

The ICL-based farm of processing servers and data acquisition compute hosts no longer present bottlenecks and are not foreseen to in the near future. As a result, the cycle of ICL computing upgrades has been relaxed, and a single 100% upgrade of the commodity compute elements in ICL will performed in the next five years, split across two polar seasons. On the contrary, access to sufficient computing resources in the north remains a challenge for IceCube. The complicated optics of the ice demands large numbers of graphics processing cores to track photons in the ice for the purposes of event reconstruction and simulation of detector response. Generation of hadronic interactions in the cores of extensive air showers high above the polar plateau needs large farms of CPUs. In the upcoming years, IceCube M&O is extending access by the collaboration to computing resources within the central cluster at Madison, providing better frameworks to access dedicated and opportunistic distributed resources, and making progress towards access to massive supercomputing and cloud computing resources.

The discovery of the first source of high-energy astrophysical neutrinos is a clear demonstration of the validity and success of the NSF large facilities program and the ongoing investments in facility operations. New outreach methods that leverage mass media channels will build upon the existing successful E&O and communications program supported by M&O.

In the next few years, the remaining review in PY4 will re-evaluate the M&O and in particular examine the quality of response to requests given in the previous PY2 review. The outcome of this second review will inform NSF's decision to rebid the M&O at the conclusion of this five-year cycle.

3 Technical Approach

As a discovery instrument with multiple scientific objectives, IceCube requires many varied search strategies. It looks for steady point sources of muon neutrinos in the northern sky—for example, active galactic nuclei or supernova remnants. Other searches target transient point sources such as gamma-ray bursts or supernovae in progress. Yet another objective is to characterize the recently discovered extraterrestrial neutrino flux coming from the entire sky and to follow up detections of high-energy astrophysical neutrinos with multimessenger observations by promptly alerting other telescopes. To achieve these multiple objectives, IceCube must be properly calibrated and continuously monitored to ensure high-quality data. It also requires computing and facilities infrastructure, and the corresponding maintenance and updates necessary to achieve high standards of reliability and quality.

This section sets the technical M&O requirements and specifications ensuring IceCube reliably and continuously provides the capability to achieve its scientific objectives.

3.1 Detector Description and Performance

Required Capabilities. IceCube is designed to detect muons and cascades over a wide energy range. The string spacing was chosen in order to reliably detect and reconstruct muons with energies over 1 TeV and to precisely calibrate the detector using flashing LEDs and atmospheric muons. Because of the attenuation and scattering of light, a certain density of sensors is required to obtain many measurements along each track, which is important for pointing accuracy, background rejection, and energy measurement. The optical properties of the South Pole ice have been measured with various calibration devices and are used for modeling the detector response to charged particles. Muon reconstruction algorithms allow measurement of the direction and energy of tracks that come from all directions.

The depth requirement was driven by two constraints: a) to deploy below the region where air bubbles contribute to light scattering (1400 m), and b) to maximize the use of the remaining depth without risking too close an approach to bedrock (2800 m). Exploratory measurements with the Antarctic Muon And Neutrino Detector Array II (AMANDA-II) verified that the ice is clearer in the region below 2100 m. The greater clarity helps with reconstruction, and the greater depth minimizes background events.

Some of the high-level design goals include:

- Angular resolution for muons (E^{-2} spectrum): < 1° (Actual: 0.5°)

- Angular resolution for muons at 1000 TeV: < 0.7° (Actual: 0.4°)
- Muon Effective area at 10 TeV: 0.9 km^2 (Actual: 0.9 – 1 km^2)
- Livetime: >95% (Actual IC86-2018 physics run: 99.79%)

Infrastructure. The final configuration of the detector (Figure 3.1-1) consists of 86 strings with an instrumented depth range from 1450 m to 2450 m below the surface. There are 60 optical sensors mounted on each string, with equal spacing for standard strings. On the eight strings of the DeepCore subarray, 50 sensors are deployed at a smaller spacing of 7 m between 2100 m and 2450 m, with 10 sensors above 1950 m for additional veto functions. In addition, there are 324 sensors deployed in 162 IceTop detector tanks on the surface of the ice directly above the strings. The sensors are connected to the IceCube Lab (ICL) with a cable containing copper wires, one twisted pair for each pair of sensors. The ICL supports all data processing infrastructure to trigger, build events, and process the data.

M&O Requirements. 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 will undergo changes in time. Calibration constants change over time, data rates change due to the seasonal fluctuations of the atmosphere, and sensors may develop defects and need quick attention to avoid serious system-wide problems. The major effort is required for maintenance and operation of the complex computer systems in the ICL and for data management.

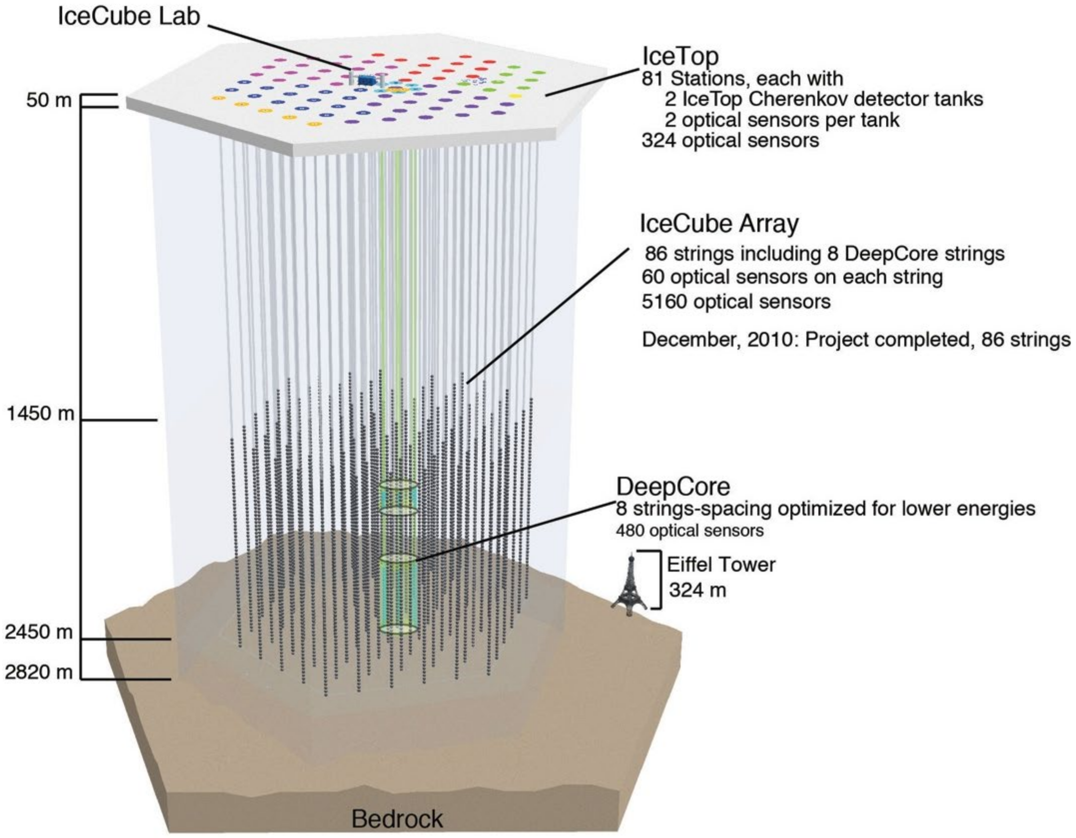

Figure 3.1-1. Schematic View of IceCube Detector. *The detector must be calibrated and continuously monitored to ensure collection of high-quality scientific data.*

3.1.1 Digital Optical Modules (DOMs)

Required Capabilities. Each sensor is required to detect the Cherenkov light emitted by charged particles with high sensitivity and a time resolution of a few nanoseconds (ns) and high dynamic range. Requirements include:

- Time resolution: 5 nsec (Actual: ~3 ns)
- Time synchronization to master clock: < 3 ns (Actual: 1.5 ns)
- Noise rate (with deadtime): 500 Hz (Actual: ~350 Hz)
- Linear dynamic range: 200 PE/15 ns (Actual: ~500 PE/15 ns)
- Failure rate (permanent failures): < 5%/15yr (Forecast: < 2.5%/15yr)
- Deadtime within run: < 1% (Actual: < 0.01%)

For IceCube, timing precision at the level of a few nanoseconds is necessary to maximize the accuracy of angular reconstruction; when looking for point sources of neutrinos in the sky, having two tracks pointing to the same spot within 0.5 degrees is more significant than having them point to the same spot within 1 degree, because random background tracks are four times more likely to occur within 1 degree.

The dynamic range of 200 photoelectrons (PE) per 15 ns is relevant in IceCube DOMs in order to measure light near high-energy tracks, which is directly proportional to their energy (loss). For extremely high energies, the light will saturate nearby DOMs, and the energy must be determined with more distant DOMs, requiring a precise simulation of the photon propagation over large distances.

For IceTop DOMs, the dynamic range is important because cosmic ray air showers are studied across a wide energy spectrum (about four orders of magnitude), and the signals grow with shower energy.

The noise rate affects the trigger rate, the bandwidth, and most importantly the reconstruction quality and the sensitivity to neutrino bursts from the core collapse of supernovae. Aside from the goal of a low noise rate, it is equally important that the noise is predictable, stable and free of spikes.

Infrastructure—the As-Built DOM. Each sensor consists of a 25-cm photomultiplier tube (PMT), connected to a waveform recording data acquisition circuit capable of resolving pulses with nanosecond precision and performing within the requirements as listed above.

Each DOM (Figure 3.1-2) triggers autonomously on single photons and sends time-stamped, packetized hit data to the surface. A 33-cm-diameter pressurized glass sphere holds the Hamamatsu R7081-02 photomultiplier tube plus associated electronics. These electronics include a high voltage generator, a resistive divider PMT base, a flasher board (containing 12 light emitting diodes, with programmable drivers), and a main board containing a complete data acquisition (DAQ) system. The DAQ includes two separate waveform digitizer systems. The first is the analog transient waveform digitizer (ATWD), which uses a custom switched-capacitor array chip to collect 128 samples of the PMT output at 300 megasamples per second (MSPS). The ATWD has three independent channels for each PMT providing 16 bits of dynamic range. The second digitizer system uses a commercial 40 MSPS 10-bit ADC chip to record 6.4 µsec of data after each trigger.

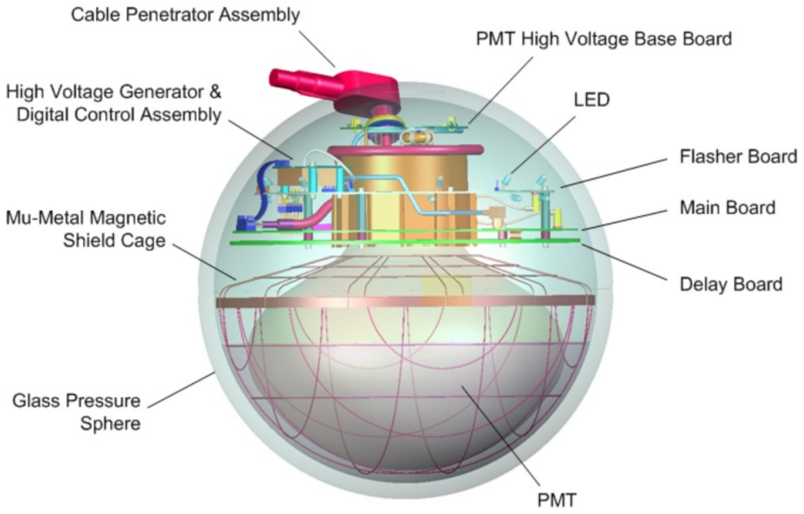

Figure 3.1-2. Digital Optical Module. *As the heart of the detector, DOMs require regular monitoring to detect performance issues that affect the quality of physics data.*

M&O Requirements. The system parameters, such as gains of all amplifiers, noise rates, time resolution, master clock synchronization, photodetection efficiency, and trigger thresholds need to be monitored from run to run, and even in shorter time intervals. Due to the large number of sensors, even occasional perturbations of individual sensors can have detrimental effects on the data quality. While overall a high reliability and stability has been achieved, experience shows that regular monitoring and a rigorous assessment of the observed and often complex issues is required to ensure high data quality. Detailed calibration programs need to be performed on all sensors at regular time intervals. Higher-level tests with LED flashers and downward-going cosmic ray muons are used to verify the system time stability between neighboring DOMs and monitor the DOM charge response.

3.1.2 IceTop

Required Capabilities. The IceTop surface detector array is designed to detect cosmic ray air showers in the energy range from 500 TeV to energies well beyond 1 EeV. Full trigger efficiency is required above 1 PeV for events with the core in the array. Coincidences with the In-Ice detector string array, the main detector of IceCube, allow performance of 3 tasks: a) cosmic ray physics over a wide energy range, b) special cross-calibrations, and c) certain veto functions. The ice in the tanks must be clear and remain clear without cracks over many years. The stations are exposed to and must survive annual temperature cycles down to below -50°C.

Infrastructure—the As-built IceTop Detector. The surface air shower array, IceTop, consists of ice Cherenkov detector tanks each containing two DOMs, which are operated at different gains for increased dynamic range. Two such tanks are associated with each string. The tanks are embedded in the snow just below the surface to minimize drifting of snow. IceTop detects and measures the position and direction of cosmic ray air showers, which also contain muons that penetrate to IceCube depth.

M&O Requirements. The DOMs used in the IceTop tanks must be serviced like all other DOMs. However, the lower gain of every other sensor and the different noise condition from cosmic rays result in different observables and make the IceTop array a complete detector system on its own. Special expertise is needed to service the IceTop array, both at the DOM level as well as at the DAQ level.

The increase of the snow layer on top of the tanks negatively affects the detector efficiency and energy threshold. Annual measurements of the depth of snow on all tanks must be performed, and this information is updated in a database used for reconstruction and simulation. Furthermore, the decision that snow maintenance will be minimized has necessitated augmentation of the IceTop stations with additional detector elements to mitigate this issue and restore IceTop efficiency to prior levels.

Comparing the IceCube (In-Ice) measurement of muons with the IceTop system is one important test of proper calibration and of the reconstruction software. This is an ongoing comparison through the life of IceCube to make sure that everything continues to function as designed, i.e., calibrations or reconstructions or their interfaces have not become corrupted.

3.1.3 Central Electronics and Data Processing System (Counting House)

Required Capabilities. The array of DOMs in the deep ice and in IceTop needs to be supplied with power, communication and control functions. All sensors are connected to the central data acquisition electronics by cables. A pair of DOMs shares one twisted pair of copper wires. The data are collected in the ICL, located at the geometric center of the IceTop array. Data include full waveforms for all hits in time coincidence between two neighboring DOMs, plus summaries of isolated hits. The data streams from the sensors arrive asynchronously via a digital communications protocol. In the ICL, higher multiplicity coincidences are formed to trigger on muons or cascades in the deep ice, or air showers observed in IceTop. The bandwidth allocation depends on the satellite bandwidth availability at the South Pole. It is a system requirement to store data locally in case of an extended failure of the satellite transmission system.

Infrastructure—Data Acquisition and Data Processing System. An overview of the system architecture is given in Figure 3.1-3. Each string (5 cm diameter and typically 3 km long cable) is connected to one DOMHub, a computer with custom PCI cards that perform the three low-level functions listed below. The central data acquisition performs three functions:

- receive data streams from DOMs, perform format changes, form event triggers and build events;
- provide power and slow control functions to DOMs; and
- perform synchronization of all DOM clocks with the system master clock.

M&O Requirements. While the system is designed to perform most functions automatically, the maintenance and operation require professional staff to ensure long-term reliability and stable operation of the experiment.

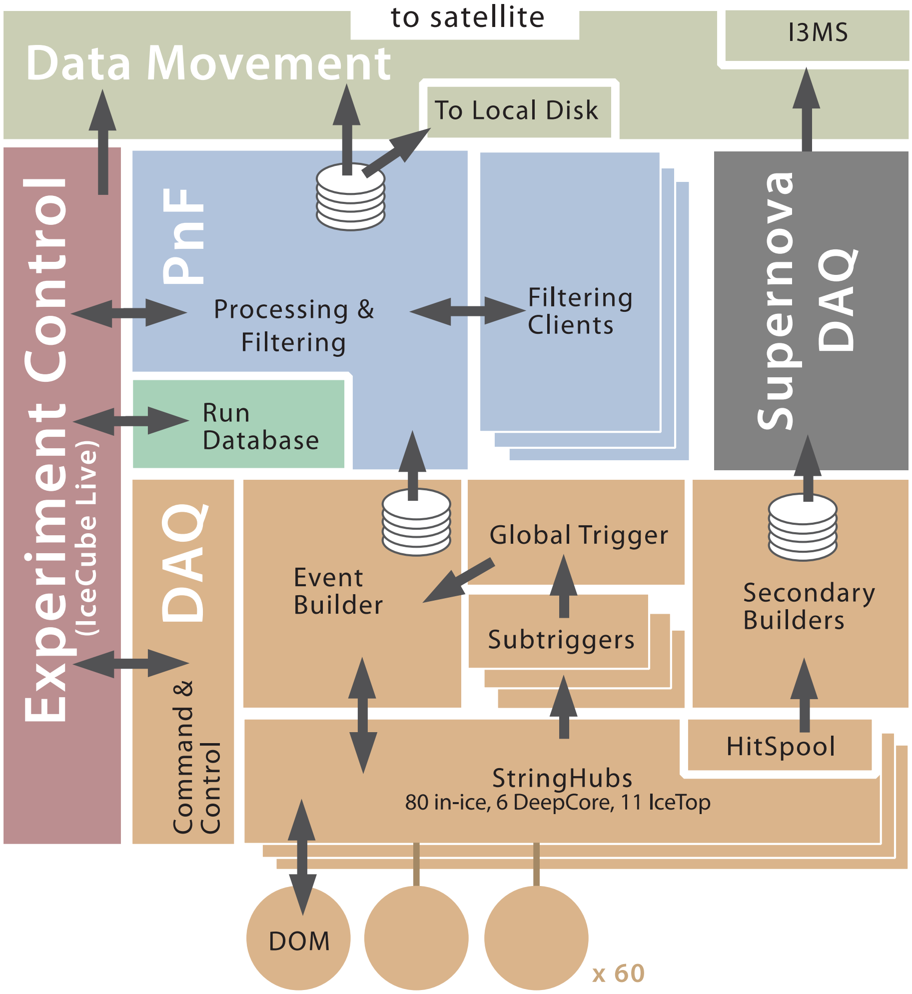

Figure 3.1-3. Detector Data System Architecture. *The data system controls the detector and collects, processes, transmits and stores IceCube and IceTop scientific data.*

3.2 IceCube Infrastructure

3.2.1 United States Antarctic Program (USAP) Infrastructure

Required Capabilities. The IceCube Laboratory (ICL) is one of the core facilities that make up the IceCube Observatory at the South Pole. It fulfills the requirement for a centralized computing facility and physical interface between the IceCube cables and the DOMHubs and associated data processing equipment. Stable electrical power to the IceCube detector is required as a sustained power outage could lead to damage of both surface electronics and in-ice electronics. Additional infrastructure that is required for IceCube management and operations functions are the South Pole Station and the cargo and logistics capability provided by the NSF Antarctic support contractor. IceCube also requires network access to the South Pole, and within the South Pole Station network, for data transfer and communications for network remote access, email, and other basic services. In addition, IceCube needs the capability of transferring data from the South Pole to the IceCube data warehouse in Wisconsin through a number of different pathways depending on the priority of the data.

Infrastructure. The IceCube computing systems located in the ICL (Figure 3.2-1) produce in excess of 30 kW of waste heat that must be removed from the data center. To reduce energy consumption of the data center, the cold external air is used for cooling through an air mixing and handling system. Due to the very high density of equipment in the ICL, a failure of the cooling system can result in critical damaging temperatures within 30 minutes. A high level of reliability and monitoring of the cooling system is therefore required. The NSF support contractor is responsible for the operations, maintenance, monitoring, and response to incidents involving the cooling system. The communications infrastructure, in the form of high-speed satellite data connections and the physical backbone at the South Pole, is also maintained by the NSF Antarctic support contractor.

M&O Requirements. The basic framework of frequent communications (regular conference calls during season planning), one-on-one contacts (NSF support contractor program manager, NSF program officer), Support Information Package (SIP) development, and ad hoc meetings ensure that the USAP program will continue to provide IceCube with needed USAP infrastructure.

Figure 3.2-1. IceCube Laboratory (ICL). *The ICL at the South Pole houses the online computing system that is critical to collecting, storing, and filtering data from IceCube.*

3.2.2 IceCube South Pole System (SPS)

Required Capabilities. IceCube requires a surface computing system that supports the data acquisition and filtering tasks carried out by the real-time systems. Data archive systems must be able to handle the Level 0 data volume generated from the IceCube detector, approximately 300 TB per year.

Infrastructure. The SPS hardware includes DOMHub computers, commodity server class computers, remote console and power equipment, GPS units, network hardware, and UPSs. Four Iridium RUDICS modems provide low-bandwidth connectivity 24/7 for detector control and monitoring.

M&O Requirements. The SPS must maintain very high reliability to support IceCube data taking in a robust manner with minimal intervention. Also, the system must be manageable for the winterover operators, who are different and are re-trained every year.

System administrators, in conjunction with on-site winterovers, are responsible for the maintenance and operations of the South Pole computing system. This includes preventive maintenance, troubleshooting, and upgrades.

3.2.3 IceCube UW Infrastructure

3.2.3.1 South Pole Test System (SPTS)

Required Capabilities. IceCube requires a test system that replicates the basic functionality and performance characteristics of the production SPS computing system and includes a reasonable proxy representation of the detector and DAQ systems.

The SPTS provides an environment to build and verify software subsystems and perform hardware and software evaluations prior to deployment on the production system at the South Pole. The system must adequately represent the live detector systems at a small fraction of the cost in hardware and maintenance resources.

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

Evaluation of software and firmware DAQ updates is accomplished on the SPTS via one or more of the following means:

- 8 real DOMs kept at subfreezing temperatures for evaluation of firmware and software updates,
- a full string of DOM main boards connected to a DOMHub for string-level tests,
- a full-length IceCube cable for communication tests (PCTS) and three walk-in freezers for environmental tests (mDFLs) located at Physical Sciences Laboratory (PSL), and
- a special playback mode of real untriggered data streams coming directly from the deployed DOMs that have been captured at the SPS for use in the SPTS to generate realistic load conditions.

M&O Requirements. System administrators manage the test system responding to software developers and other engineers' requirements. They are responsible for hardware and software maintenance and operations on the SPTS.

3.2.3.2 Data Warehouse and Storage Infrastructure

Required Capabilities. IceCube generates about 1 TB of raw data every day. Only about 10% of that can be transferred out via satellite due to limited bandwidth. A reliable system is needed to store all the generated data to enable long-term archiving and to manage transfers. The data generated by the detector is its most precious output, so the archiving systems have to ensure its integrity.

IceCube requires a data warehouse at UW–Madison consisting of software to facilitate receiving and archiving of data from the South Pole, standards for organizing the data, such as directory structure and metadata, and a high-performance storage system that enables efficient offline data processing and analysis.

Including simulation, IceCube generates close to 1 PB of data every year. Out of that, about 700 TB need to be archived and preserved for the long term due to their uniqueness or their relevance to reproducing published scientific results.

Infrastructure. The storage infrastructure at UW–Madison's data center 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.

IceCube software engineers have developed a software application named JADE to handle the IceCube data end-to-end. JADE manages the data collection and storage at the South Pole, the satellite data transfers, the data ingest at the UW–Madison data warehouse, and the replication to external archive sites for long-term preservation.

In order to provide cost-effective long-term data preservation services, we leverage large data centers at collaborating institutions NERSC and DESY-Zeuthen that already operate large automated tape libraries as part of their services to other experiments.

The NERSC archive is a hierarchical storage management system (HSM) that has been running since 1998. It currently holds 100 PB of data and handles about 100 TB of input/output every day. The HPSS software is used to manage the tape backend. NERSC actively partners with Globus, DSI, and HPSS developers to improve HPSS functionality.

DESY has been providing automated tape archive services for particle physics experiments for more than 20 years. The tape archive system at the DESY-Zeuthen site, where part of the IceCube data is replicated, stores a total of about 15 million files and 2.5 PB of data for several experiments. The OSM software is used to manage the tape backend.

M&O Requirements. System administrators experienced in managing disk enclosures, storage networks, servers, and cluster file system software maintain and operate the storage infrastructure. They ensure that data is available for data processing and analysis tasks and that it is delivered with maximum performance.

As the collected data set grows and new analyses are developed, the load on data access services increases as well. Part of this data processing and analysis demand comes from a more intense usage of distributed computing (Grid) resources. The IceCube data warehouse and storage infrastructure, and in particular the data export services responsible for providing remote access to the data, will need to evolve to cope with the extra load and maintain high performance and reliability.

The JADE data handling software and services will require maintenance and support, to tailor them to the evolving needs of the scientific community. Metadata and cataloging needs will also evolve as the data set grows and new analyses appear.

3.2.3.3 Central Computing Resources

Required Capabilities. IceCube requires a core high-throughput computing (HTC) cluster to perform offline processing and analysis of real data and for the production of simulation data sets. The system must be closely coupled to the storage infrastructure for efficient data processing.

Infrastructure. The current IceCube HTC cluster at UW–Madison consists of nearly 200 servers providing a total of around 7000 CPU job slots and 432 GPU job slots. The HTCondor software, a state-of-the-art workload management system developed at the computer sciences department of UW–Madison, handles job scheduling at the HTC cluster.

M&O Requirements. The IceCube system administrators support users by providing guidance on HTC best practices. They also support the delivery of science-ready data by ensuring that offline processing tasks run with appropriate priority and that the end-to-end infrastructure stack (computing, network, and storage) is available and delivers optimum performance.

3.2.3.4 Core Data Center Infrastructure

Required Capabilities. The data center infrastructure is the glue that connects the major computing resources of IceCube. Required core infrastructure systems include distributed authentication, DNS, and e-mail. Also, a large number of servers and services need to be deployed and maintained, such as database services, web services or tailored application servers, to fulfill science needs.

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

Infrastructure. IceCube computing facilities are currently hosted in various Madison locations. Core services and storage are located at WIPAC headquarters, at 222 W Washington Ave. Compute nodes are mostly located at the Physics department building on campus, and a portion of the GPU compute nodes are located at the Wisconsin Institutes for Discovery building. UW–Madison is deploying a new state-of-the-art data center in a co-location facility in Fitchburg. We have already started to migrate the IceCube storage and core services from WIPAC headquarters into this new facility. This will improve availability and will facilitate future capacity expansions. The plan is to complete this migration before the end of the year.

M&O Requirements. Several systems administrators share duties to maintain the UW–Madison data center infrastructure services. This includes patching, monitoring, troubleshooting core services, and responding to user needs among other tasks.

3.2.4 Distributed Computing Infrastructure

Required Capabilities. The analysis of experimental data requires a suitable amount of simulated data that reproduces the detector response to a well-defined set of physics events. The IceCube observatory's event rate is overwhelmingly dominated by cosmic-ray-induced background events, which must be eliminated through a complex event selection process. A large amount of simulated data needs to be generated in order to perform high-quality physics analyses. Dedicated computing resources at the level of several thousand CPU cores are needed to perform the required simulation and analysis tasks. The available capacity will increase with time along with rising needs.

Infrastructure. In order to reach this capacity, IceCube relies on distributed resources available from collaborating institutions. 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 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 that 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 archiving of IceCube data sets.

M&O Requirements. Support personnel at all sites coordinate and manage the distributed computing effort to produce the needed simulation. In addition, IT professionals at the UW–Madison data center manage the IceCube Grid infrastructure and tools needed to exploit distributed resources and to provide efficient remote access to the data.

3.3 Overview of Events to Publications

Reconstructing neutrino events with energies from 10 GeV to 1000 PeV, the energy range in which IceCube's science portfolio includes neutrino physics, dark matter searches, and the observation of cosmic neutrinos, requires precise recording of everything from single photons up to large pulses lasting several microseconds. Proper maintenance and operation of the detector and its supporting infrastructure (Sections 3.1 and 3.2) allow for capture of the targeted events, analysis of the data, and publication of results that contribute to science and education (Figure 3.3-1).

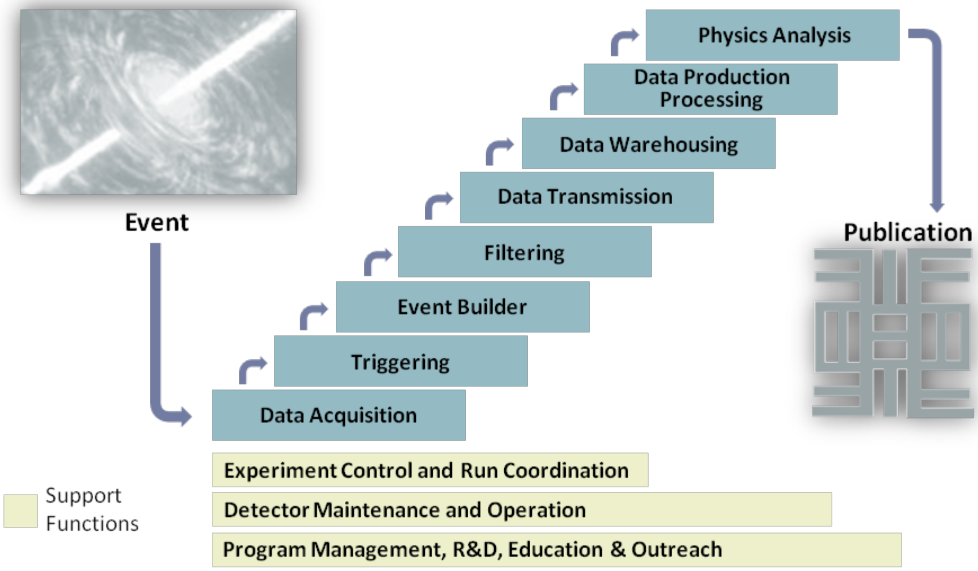

Figure 3.3-1. IceCube Path to Science. *Our approach to IceCube M&O is structured to support all tasks required to produce science—from event to publication.*

Detector maintenance and operations provides the online framework for the capture of astrophysical events of interest. This process begins at the DOMs in the ice, which continuously time-stamp and digitize photomultiplier tube pulses originating either from passage of charged particles or from intrinsic light background in the DOMs themselves. Data is digitally transferred to the surface where computers order the hits in time and trigger on resulting patterns of scientific interest to separate them from the intrinsic noise. Once a trigger is issued, hits close to the trigger times are assembled into file-based event data structures by event builder processes and sent to the PnF filtering system, which may be regarded as a second level trigger. Filtering further selects out potentially interesting events such as upgoing muons, high-energy neutrino candidates, and several others from triggered events for daily transmission via satellite. All triggered events are recorded to disk storage, however, which are physically shipped north at the end of the austral summer season. A separate process (JADE) takes care of managing the data streams, buffering data, sending the PnF stream to the satellite and writing the bulk of the data locally on disk.

In 2016, PnF has been equipped with the ability to generate real-time alerts to be sent to external observatories when extremely interesting events are found in the data stream. Optical observatories, for example, may elect to examine the area of sky where these targets of opportunity occur and make follow-up searches for associated optical signals. It is known that many astrophysical objects such as blazars emit energetic gamma rays often in flares of time scales of weeks, days and even less than an hour. These multimessenger collaborations are intended to aid in the eventual identification of objects that produce high-energy neutrinos. One such alert on September 22, 2017 resulted in a successful and broadbased follow up program leading to the coincident observation of gamma rays from a distant galaxy with the results published in Science magazine. The alert was issued a mere 43 seconds after the event was recorded in the detector and the observations made a prime example of the importance of real time technology data processing and multimessenger astronomy.

A run coordinator oversees and configures the online systems through a global experiment control system called *IceCube Live* to focus data collection on areas of scientific interest prioritized by the IceCube Collaboration. Data filters are annually proposed by working groups in the Collaboration and are reviewed by the Trigger Filter and Transmission (TFT) Board that allocates resources such as computing and data transfer bandwidth.

Each data stream is reprocessed after transmission to the Northern Hemisphere data center, where more computing power is available and more sophisticated reconstruction algorithms can be applied. The reprocessing takes place within only a few weeks after the data are taken at the South Pole. At this point, the science-ready data are available to the IceCube scientific collaboration. The refined data streams are first evaluated by the channel working groups for initial analysis and for possible recommendation for further filtering. The physics working groups typically only access the processed data for high-level analysis and development of specific tools needed to execute the analyses. The analysis coordinator manages the analysis process, which typically includes formal analysis and unblinding proposals and an approval process. The Publication Committee manages the publication review processes.

4 Management Approach

Our approach to IceCube M&O—from science events to publication—is to maximize the scientific discovery potential by drawing on talent and resources from Collaboration institutions to support both M&O and science tasks. The first part of this section (Section 4.1) 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. The second part (Section 4.2) identifies the tasks required to meet the technical requirements and specifications discussed in Section 3, and explains how we perform each task.

4.1 Organization

The IceCube M&O management organization integrates the IceCube Collaboration and the host institution, the University of Wisconsin–Madison (Figure 4.1-1). The principal investigator is responsible to the UW vice chancellor for research and the National Science Foundation for the overall scientific direction of the IceCube Neutrino Observatory. The Collaboration spokesperson appoints collaborating scientists to serve as the coordinators of science working groups as well as the overall analysis coordinator. These appointments are subject to the concurrence of the Collaboration. The director of operations appoints technical professionals to serve as M&O coordinators. The managers in these areas work with their scientific counterparts to ensure the detector, data processing, and software systems operate reliably and the data taken by the detector can be analyzed in a timely, consistent manner.

The IceCube spokesperson and the director of operations are jointly responsible for the success of the IceCube M&O program, with the spokesperson directly accountable to the Collaboration and the director of operations accountable to the National Science Foundation through the University of Wisconsin–Madison as the host institution for the M&O program.

The spokesperson-appointed coordinators and the director of operations-appointed managers are successful through the efforts of collaborating scientists, technical professionals, and managerial and administrative support staff. The entire M&O scope of work is sorted in a Work Breakdown Structure - WBS (included as Appendix 1 of this plan), and the WBS tasks are defined in a detailed Memorandum of Understanding (MoU) approved by the IceCube collaborating institutions.

Every task in the MoU is assigned to an institution. The principal investigators (PIs) at the institutions are responsible for ensuring that the work is completed on schedule. If an institution is not able to fulfill an agreed upon commitment, the institutional PI is responsible for ensuring that the work is assigned to another institution before there are adverse impacts to the M&O program. The institutional MoUs also include a list of the physics group members and a head count of faculty, scientists, postdocs, and graduate students. The institutional MoUs are revised twice a year at the IceCube Collaboration meetings. (A summary of the most current MoU head count, level of committed contribution, and a summary of the collaborating institutions evolvement over time are included as Appendix 2 of this plan).

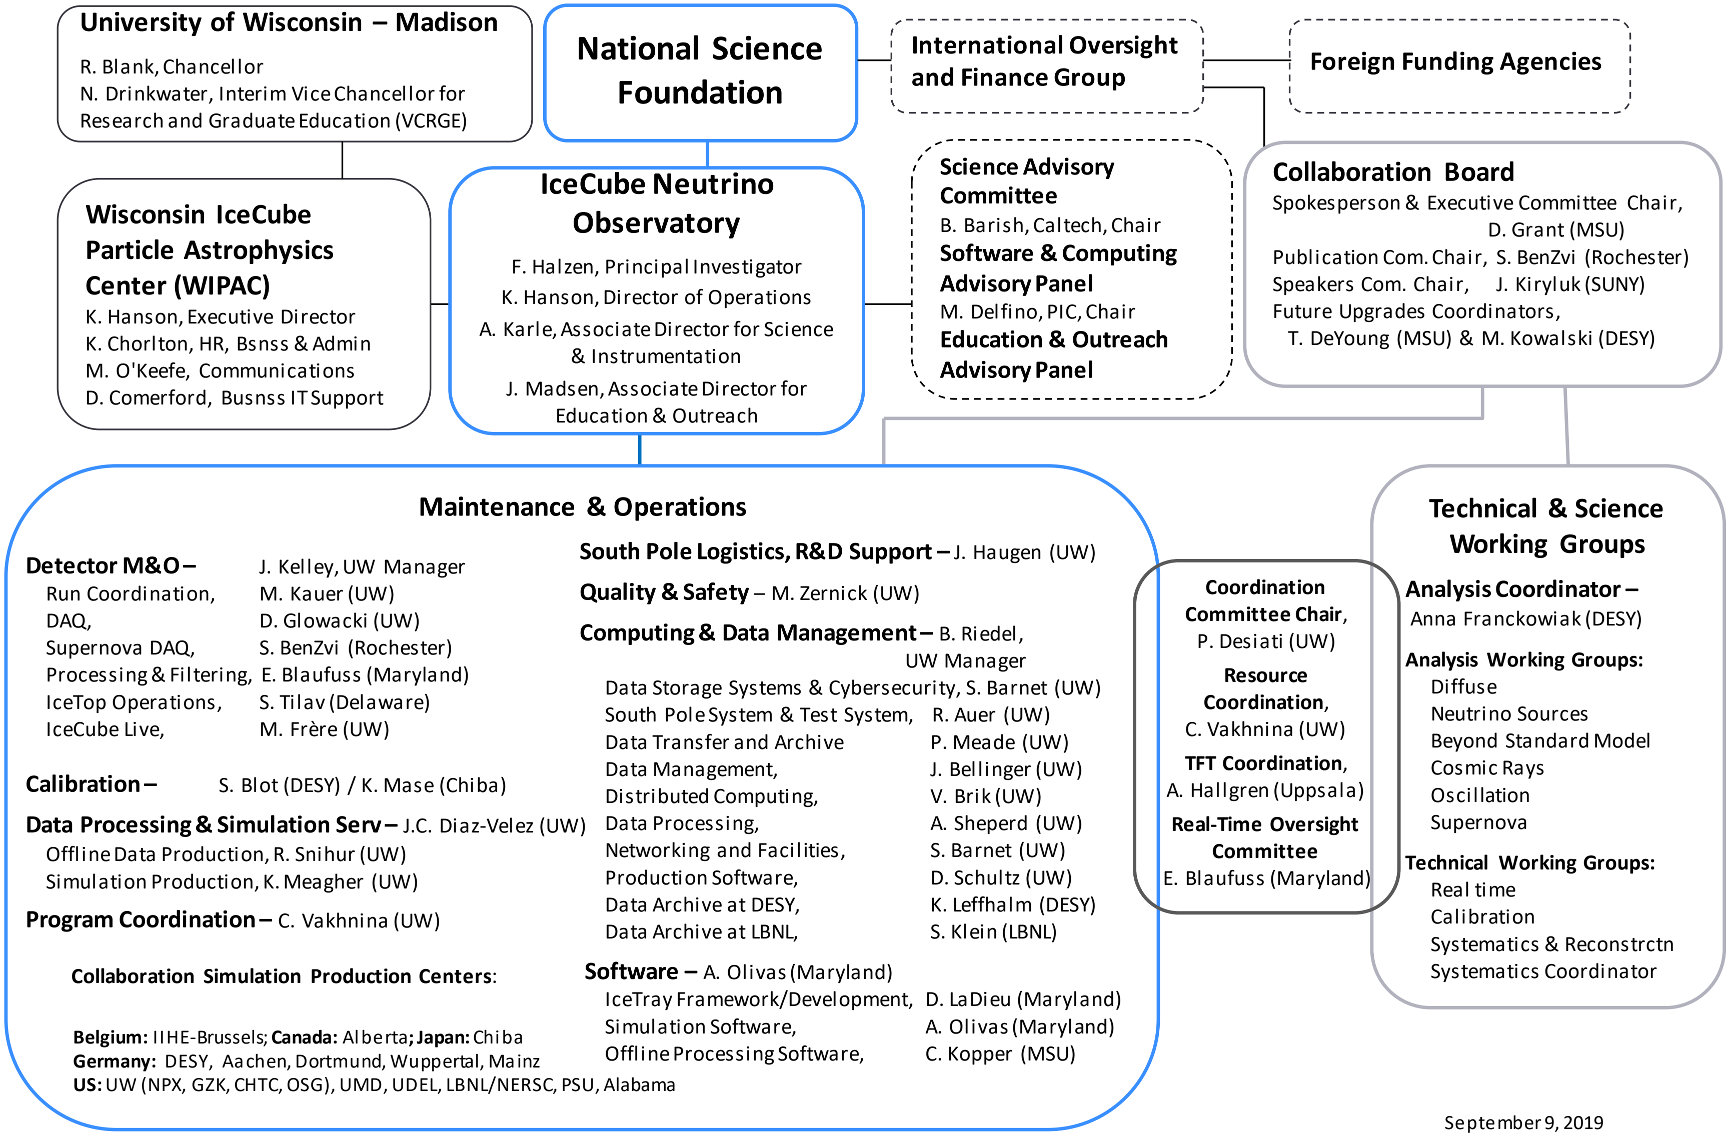

Figure 4.1-1. IceCube Organization. *Our organization maximizes the use of both Collaboration resources and core resources managed by UW while maintaining clear lines of accountability to the NSF.*

4.1.1 The 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. The NSF has a special role in IceCube because of its Host Laboratory responsibilities in managing operation of the Amundson-Scott South Pole Station. These responsibilities include safety; physical qualification; transport of personnel, fuel and equipment; and the provision of housing, food service, support personnel, logistical support, IT support, and general infrastructure support. 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.

4.1.2 International Oversight and Finance Group (IOFG)

The International Oversight and Finance Group (IOFG) was created in 2004 to provide oversight and financial support for the IceCube Neutrino Observatory (including construction phase, management & operations and research phases). The group organizes annual oversight reviews of the operations and meets annually to discuss detector performance and physics. The group also sets policies for receiving periodic progress reports on all aspects of the detector operation and by all the performers in the collaboration, and for conducting external reviews when appropriate.

Membership. A representative of the National Science Foundation chairs the IOFG. Membership comprises representatives of the funding agencies in the partner countries supporting the construction and operation of IceCube Neutrino Observatory, currently the funding agencies from Belgium, Germany, Sweden, and the United States. The IOFG is informed by the spokesperson of the Collaboration, the director of operations, the principal investigator and others as appropriate.

Decisions. The IOFG is committed to operate through discussion and consensus. The executive agent (the NSF) will make final decisions on matters before the group related to the operation of IceCube.

Issues that may come before the group include:

- Approval of a formal charter for the group.
- Review of Memoranda of Understanding (MoU) between the various institutions.
- Concurrence on the Management & Operations Plan.
- Funding issues.
- Concurrence on the Collaboration’s plans for new membership in the collaboration.
- Data sharing and data management policies.
- Coordination regarding press releases and education and outreach activities.
- Input on seasonal flight and personnel logistics planning.
- Other matters related to successful operation of the IceCube Neutrino Observatory for science.

4.1.3 University of Wisconsin–Madison

IceCube Oversight. The lead executive officer of the University of Wisconsin–Madison is the Chancellor. The Chancellor delegates responsibility for research activities to the vice chancellor for research and graduate education (VCRGE). The VCRGE maintains oversight of the IceCube Neutrino Observatory and appoints the IceCube director of operations. The IceCube principal investigator and the director of operations report directly to VCRGE. The director of operations contacts the vice chancellor for research when significant developments occur or important issues arise.

The IceCube associate director for science and instrumentation reports to the director of operations and advises primarily on matters related to science, coordination committee and instrumentation.

The IceCube associate director for education and outreach (E&O) reports to the director of operations and leads the IceCube E&O program. The associate director for E&O works with the NSF and the IceCube Collaboration to establish E&O priorities and strategies and to provide support for ongoing activities and to respond to outside requests.

Wisconsin IceCube Particle Astrophysics Center. The IceCube operations organization is located within the Wisconsin IceCube Particle Astrophysics Center (WIPAC). WIPAC is IceCube's primary interface to the UW administrative and support systems, established within the Office of the VCRGE to coordinate the multiple roles of the university:

- Lead institution for the IceCube construction project;
- Host institution for initiating and continuing IceCube management and operations;
- Administrator of services such as accounting, purchasing and human resources;
- Coordinating institution for IceCube education and outreach activities; and
- Collaborating institution with the largest participating research group.

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

4.1.4 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 Collaboration 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 Collaboration institutions. The institutions collaborating in the IceCube Neutrino Observatory are listed in the IceCube Governance Document (included as Appendix 4 of this plan).

IceCube Collaboration Board. 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 as described in detail at the IceCube Governance Document (included as Appendix 4 of this plan). 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 topical and general conferences, analysis teams, and education and outreach. The principal investigator is an ex-officio member of the ICB.

Executive Committee. The spokesperson, in consultation with the ICB, the PI and the director of operations, appoints and chairs the Executive Committee of the ICB (Figure 4.1-2). The term of the members is two years. The job of this committee is to advise the spokesperson in proposing actions to the ICB and in making interim decisions. The members of the Executive Committee represent major groups, functions and competencies within the Collaboration.

	Name and Institution	**Area of Expertise/Responsibility**
Spokesperson	Darren Grant, University of Alberta	Overall direction of IceCube Collaboration
Member	Olga Botner, Uppsala University	Former Spokesperson
	Greg Sullivan, University of Maryland	Former Spokesperson
	Albrecht Karle, University of Wisconsin–Madison	All aspects of detector operation, Associate Director for Science & Instrumentation, liaison with R&D
	Elisa Resconi, Technische Universität München	Multimessenger astrophysics
	Marek Kowalski, Humboldt-Universität zu Berlin	High-energy neutrino astrophysics
	Segev BenZvi, Rochester	PubComm chair
	Tyce DeYoung, Michigan State University	Future upgrades: neutrino oscillations
	Anna Franckowiak, DESY	Analysis Coordinator
Ex-Officio Member	Francis Halzen, Principal Investigator, University of Wisconsin	Neutrino astronomy & high-energy physics, overall scientific direction
	Kael Hanson, Director of Operations University of Wisconsin–Madison	Project and Operations Management, NSF Primary Contact for IceCube Operations

Figure 4.1-2. Executive Committee of Collaboration Board

IceCube Collaboration Meetings. IceCube Collaboration meetings are held at least twice a year with one meeting in Europe or Asia Pacific and one in the United States. These meetings serve as a forum for the presentation of scientific results, and for communicating project progress and status to the entire collaboration. Official Collaboration Board meetings are conducted during these meetings.

Collaboration Institution Tasks. Tasks are rotated in a fair and equitable manner, taking account of the special interests and capabilities of each institution. Tracking and transparency is provided as part of the MoU Scope of Work Summary (included as Appendix 2 of this plan). This summary matrix provides a breakdown of tasks by WBS Level 2 and by collaborating institution that provides the foundations of the MoU with each institution.

4.1.5 Key Personnel

Our key personnel form the leadership team that ensures the success of the IceCube M&O and the timely exploitation of its scientific discovery and education and outreach potential. This section discusses the roles and responsibilities of these personnel. Key personnel (Figure 4.1-3) are employees of the Host Institution, University of Wisconsin–Madison. UW–Madison will seek concurrence from the NSF prior to any changes in the appointments.

Name	**Position**	**Responsibilities**
Francis Halzen	Principal Investigator	Responsible for the overall success of the IceCube Neutrino Observatory
Kael Hanson	Director of Operations, Co-Principal Investigator	Ensures operations meet established performance goals and the needs of NSF and the IceCube Collaboration.
Albrecht Karle	Co-Principal Investigator, Associate Director for Science and Instrumentation	Supervises technical performance of the IceCube detector infrastructure, ensuring that it meets IceCube science objectives
James Madsen	Associate Director for Education and Outreach	Responsible for planning and executing of education and outreach activities

Figure 4.1-3. IceCube Key Personnel

4.1.6 Advisory Committees

4.1.6.1 Science Advisory Committee

In consultation with the collaboration, the principal investigator and the spokesperson appoint a Scientific Advisory Committee (SAC) of external experts. The role of the SAC is to make recommendations on the IceCube scientific goals and on any other matters that may affect the scientific activities of the IceCube Neutrino Observatory. The SAC typically meets annually. The current chairperson is Barry Barish from Caltech.

4.1.6.2 Software & Computing Advisory Panel

The IceCube Software & 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 on-line computing; on-line and off-line data processing and filtering; off-line computing facilities; and simulations and analysis tools support. The spokesperson and the director of operations appoint the SCAP members and the chairperson. Meetings are held once each year. The current chairperson is Manuel Delfino from Port d'Informació Científica (PIC).

4.1.7 M&O Coordination Boards and Organizations

The purpose of coordinating structures within the ICNO organization is to ensure that M&O tasks from raw data to publications are properly planned and executed. These organizations make certain that the resources committed in their areas of activity are realized and used efficiently and effectively. Examples include the following.

4.1.7.1 Coordination Committee

A close relationship between the operational activities and the scientific investigations is a prerequisite for achieving the science goals of the IceCube Collaboration and the proper operation of the observatory. Establishing appropriate science requirements and in-kind manpower to operate the experiment's infrastructure and to develop the tools to achieve the identified science goals is the main task of the IceCube Coordination Committee.

The charge of the Coordination Committee is to provide high-level coordination of tasks related to M&O activities and to provide technical support for the IceCube physics data analyses. To achieve this, the committee manages the dependencies between operational areas and physics working groups, in order to plan for the necessary operational support to pursue IceCube science goals and keep track of the essential in-kind contribution service tasks. The tasks identified as having high-priority deliverables are tracked to make sure that sufficient support is provided for their completion. Coordination between working group technical leaders and the coordinators of the operational areas is the principal means for achieving an optimized use of the experiment infrastructure in achieving science goals and prompt responses to critical situations that require immediate solutions.

The committee is composed of the M&O coordinators, the spokesperson-appointed analysis coordinator and working group technical leaders (shown in Figure 4.1-4), key personnel, and others as needed. The chair of the Coordination Committee works with the M&O coordinators to provide a list of service tasks needed in the specific operational areas, with priority level and estimated labor contribution (FTE); with the working group technical leaders to establish the science requirements, plan for simulation needs and computing resource allocation, and a list of service tasks needed to accomplish the requirements (see organizational chart in Figure 4.1-1); and with the members of the committee to establish specific communication channels with institutional leaders to advertise the list of needed service tasks and negotiate involvement through specific MoUs. The MoUs are typically updated and renewed twice a year, at each collaboration meeting.

The Coordination Committee makes sure that the "negotiated" deliverables and the pledges by the institutional leaders match and that service tasks are executed. The chair of the Coordination Committee sets up milestones for each of the assigned tasks to benchmark the status of service tasks. Students and postdocs working on those tasks get visibility and credit by providing reports at the IceCube weekly phone calls.

The committee typically meets on a monthly basis to address technical and resource issues, and to advance strategic goals. The committee is the primary point for determining priorities and resolving resource conflicts that arise at lower levels in the organization but also for making sure that operational activities are in correct synchronization with the changing science goals. Issues that cannot be resolved by the Coordination Committee are resolved by the spokesperson and director of operations.

Role/Area	**Representative**	**Affiliation**
ICC Chair	Paolo Desiati	UW–Madison
M&O Coordinators:		
Detector Operations	John Kelley	UW–Madison
IceTop Operations	Serap Tilav	U of Delaware
Computing	Benedikt Riedel	UW–Madison
Calibration Coordination	Summer Blot	DESY
	Allan Hallgren	Chiba
Software Coordination	Alex Olivas	UMD
Simulation Production	Juan Carlos Díaz Vélez	UW–Madison
Data Processing	Rob Snihur	UW–Madison
Icetop Simulation Production	Dennis Soldin	U of Delaware
Resource Coordinator	Catherine Vakhnina	UW–Madison
WIPAC Director	Kael Hanson	UW–Madison
Assoc. Director for Science and Instr.	Albrecht Karle	UW–Madison
Working Group Technical Leaders:		
Neutrino Sources WG	Martin Wolf	TUM
	Michael Larson	UMD
CR WG	Katherine Rawlins	Alaska Anchorage
Oscillations WG	Philipp Eller	PSU
Diffuse WG	Lu Lu	Chiba U.
Beyond Standard Model WG	Carlos Argüelles	MIT
Systematics/Reconstruction WG	Jakob van Santen	DESY
Science Collaboration:		
IceCube PI	Francis Halzen	UW-Madison
Spokesperson	Darren Grant	U of Alberta
Analysis Coordinator/Online Filtering	Anna Franckowiak	DESY
Real-time Alerts	Erik Blaufuss	UMD
TFT Board Chair	Allan Hallgren (interim)	U of Uppsala
IceCube Upgrade:		
	Ty DeYoung	MSU
	Dawn Williams	U of Alabama
	Tom Stuttard	NBI
Other Key Members:		
	Claudio Kopper	MSU
	Juan Pablo Yañez Garza	U of Alberta
	Ty DeYoung	MSU
	James Bellinger	UW-Madison
	Steve Barnet	UW-Madison

Figure 4.1-4. Coordination Committee

4.1.7.2 Trigger Filter and Transmission (TFT) Board

The role of the TFT Board is to maximize transmission of scientifically valuable data within the constrained resources of the South Pole system in support of IceCube's scientific objectives. It coordinates proposals for revisions or introduction of new DAQ trigger and software settings, online filter streams, and level 2 offline data processing. It acts as an interface between the IceCube collaboration, specifically the analysis groups and the core operations groups of IceCube. The TFT will also recommend settings for initial offline data processing in the north. The annual review cycle continuously enhances the science output of the detector. The TFT Board works with the Coordination Committee for identifying service tasks needed to accomplish its goals.

4.1.7.3 Detector Operations Working Group

The detector operations working group is responsible for day-to-day operation of the detector, including data acquisition, filtering, transmission, offline processing, calibration, and maintenance. The working group is responsible for maintaining a high detector uptime and ensuring high-quality data are delivered to the Collaboration. The detector operations manager is responsible for coordinating group activities via a weekly teleconference. The run coordinator reviews proposals for nonstandard operations 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.

4.1.7.4 Analysis Coordination Working Groups

The responsibility of the working groups is twofold: a) coordinate analysis activities amongst more than 50 institutions in the collaboration and b) provide a framework for coordinating analysis with operations and technology development for an integrated focus on IceCube science and technology issues and needs. The working groups provide specialized expertise and general support to M&O tasks that include maintaining the data warehouse; developing data preparation scripts; and supporting detector calibration and verification of its performance. Tasks for each collaboration member are described in general in their MoUs. The Collaboration assigns a leader responsible for each functional area to coordinate Collaboration institution resources in that area.

4.1.7.5 Real-Time Oversight Committee

Real-time coordination of alerts to and from IceCube provides one of the best opportunities to identify potential sources of IceCube's astrophysical neutrino sample by identification of transient phenomena in electromagnetic signals or gravitational waves. These alerts take two forms: alerts generated by IceCube and those generated by other observatories. In both cases, careful oversight is required to ensure that IceCube's public or private response under MOUs is both prompt and scientifically correct. In order to address this need for oversight, the Real-Time Oversight Committee (ROC) has been formed. This group is charged with:

- Overview of our real-time programs to ensure they are complete, correct, prompt, and well integrated with each other;
- Rapid decision on significant real-time alerts and responses within IceCube.

These actions are to be handled as quickly as possible by the available members of the ROC (not requiring all be present for a quick response) and without outside approval from working groups or the ICB. Reporting on activities to the physics working groups, analysis call, collaboration email, and/or ICB occurs promptly following significant actions.

4.1.8 Physics Analysis Coordination

Physics analysis includes tasks that are not included in the M&O core and in-kind budgets but are essential to complete the process from science event to publication. These tasks are supported through research grants to the collaborating groups.

4.1.8.1 Analysis Coordinator

IceCube reaches its greatest potential both in achieving its scientific objectives and in education and outreach by balancing centralized M&O resources with resources distributed among Collaboration members and maximizing the benefits of the specialized expertise of each collaborating institution, both in M&O and in analysis.

The distributed model is illustrated in Figure 4.1-5. The work is performed in working groups, which are divided into technical and physics analysis working groups. The technical working groups develop tools and maintain filters that can be used by multiple analyses, whereas physics analysis working groups develop the high-level analysis strategies and debate the statistical interpretation of results and updates on physics scenarios. Systematic uncertainties are an increasingly dominant component of IceCube analyses. Sophisticated tools to analyze the impact of systematics on analyses are necessary to fully exploit the physics potential of the IceCube detector as we move beyond the initial discovery phase and are overseen by the systematics coordinator. After the discovery of the first neutrino source candidate, the blazar TXS 0506+056, a large effort coordinated by the neutrino sources science group has gone into developing new multimessenger analyses. The physics analysis working groups have introduced formal technical coordinator roles to track data release and simulation needs. IceCube data analysis is coordinated by the IceCube Collaboration under the leadership of the analysis coordinator, a position appointed by the spokesperson with concurrence of the ICB. Analysis funding is provided directly to the IceCube collaborating groups by their respective funding agencies.

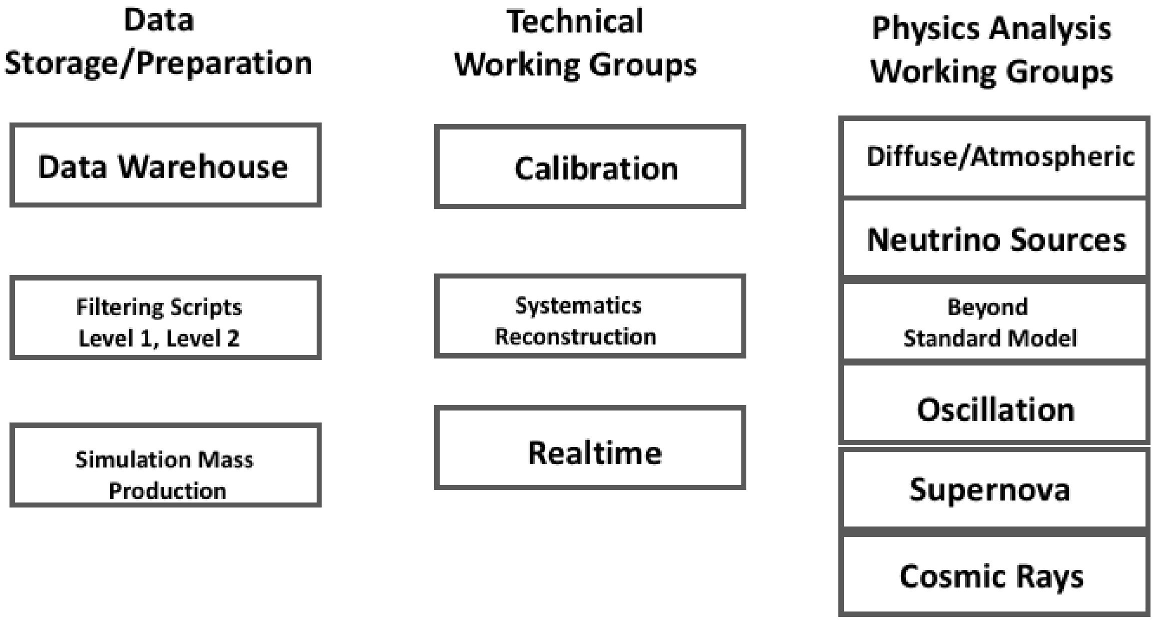

Figure 4.1-5. Data Analysis. *Shown is a schematic view of the distributed data analysis organization.*

The analysis coordinator uses several communication mechanisms to coordinate analysis activities and ensure high-quality data analysis using the best resources available to the Collaboration. The physics working groups hold weekly or biweekly teleconferences, supplemented by weekly plenary teleconferences on topics of more general interest, such as the impact of changes in simulation, software, and calibration. The weekly data analysis teleconference discusses analyses in the final stage, when they are ready for review by the full collaboration. The analysis teleconference also reviews the outline for publications, approves plots for use in conferences, and reviews the release of public data. To improve communication between the working groups, regular updates of the working group activities are presented by the working group convenors at the weekly analysis call and by the technical coordinators at the weekly technical call.

Conscious and unconscious biases can impact the robustness of physics analyses and results. IceCube has adopted the Collaboration policy to perform analyses in a “blind” manner to prevent the analyzer from biasing the result toward their own preconceptions while their analysis is under development. Application of blindness can be performed in different ways, including time and direction scrambling, and restricting analysis development to smaller subsets of the data sample. The blinding policy for IceCube does not prevent full exploration of the data, especially for calibration, verification, and reconstruction. Moreover, in the event of multiple analyses of the same data sample, the unblinding of one analysis does not bias the status of any other analysis. Application of this policy is neither centralized nor controlled by a specific authority; rather, the physics working group assigned to perform the analysis is responsible for blinding the final answer while analysis procedures are being set. Once the analysis is approved by the Collaboration, the permission to unblind is granted, and the final results are produced.

As multiple analyses are performed on the same data set and new analyses often develop out of existing work, reproducibility is crucial. The analysis coordinator is working with the convenors of the working groups to improve documentation of analyses to ensure reproducibility: a new requirement for unblinding in the short term will be to archive analysis software and data used in a centralized location. In the longer term, analyzers will also be required to demonstrate that their results can be reproduced by someone else in their working group. The analysis coordinator is also encouraging the development of technical papers, to document methods, and intragroup and intergroup software and simulation coordination efforts, in conjunction with the ICC chair. One major technical effort planned for 2020 is to publicly release a new set of high-quality astrophysical neutrino track data from the first ten years of IceCube data.

4.1.8.2 Publication Committee

After discussion and positive reception by the Collaboration of the results of an analysis, a working group produces a draft paper with supporting web pages. To be acceptable for publication, physics papers must have significantly better sensitivity than previous IceCube published results and/or demonstrate a substantially improved method. The Publication Committee regulates and manages the review process for IceCube papers. It consists of senior physicists, the analysis coordinator, and the Collaboration spokesperson. The Publication Committee sets standards and procedures for publication of papers and conference proceedings to ensure a high standard of quality and integrity for IceCube scientific papers. Moreover, through organized review panels, the Committee participates actively in the refereeing process of each paper and contribution to conference proceedings.

4.1.9 Milestones

On an annual basis, the management and operation of the ICNO involves the following periodic activities (Figure 4.1-6):

Month	Activity
January – December	Assess annual computing needs, augment central cluster as required
March	Submission of Support Information Package (SIP) for following polar season
March	Recruitment of winterover experiment operators for following polar season
April	Submission of interim 6-month project report to NSF program officers
April – May	Spring meeting of the IceCube Collaboration
May	Start of IceCube annual run
June	Annual submission of M&O Plan
June – September	Purchase and testing of equipment for following polar season
August – September	Winterover training
Early September	Season safety and readiness review
Late September	Shipment of equipment to Pole
September – October	Fall meeting of the IceCube Collaboration
October	Submission of annual project report to NSF program officers
November – February	Polar season – maintenance/upgrades

Figure 4.1-6. Management & Operations Periodic Milestones

In addition, several high-level project milestones have been identified (Figure 4.1-7). Detail on longer-term project technical goals will be further developed in subsequent M&O Plan documents.

Date	Milestone
Nov 2016	Completion of XSEDE computing allocation
Sept 2017	Completion of long-term archive baseline data import to NERSC
Jan 2018	Deployment of scintillator panels based on SiPM photodetectors
Jan 2018	SPS computing upgrade phase I complete
Mar 2018	PY2 M&O review
TBD	Integration of IceCube HTC into OSG
Jan 2019	SPS computing upgrade phase II complete
Jan 2019	Deployment of scintillator panels with new data acquisition system
Mar 2020	PY4 M&O review
May 2021	Final project report M&O 2016-2021

Figure 4.1-7. Management & Operations High-Level Milestones

4.1.10 Reports and Reviews

The IceCube Neutrino Observatory reports are distributed within the IceCube organization, including the Collaboration, host institution, and various advisory and oversight committees, and are submitted to the National Science Foundation.

Annual Report. The annual report will describe progress made based on objectives in the annual M&O Plan. Significant differences between planned and actual accomplishments will be discussed. The report will consist of a summary of work accomplished during the reporting period, including major technical accomplishments, an assessment of current or anticipated problem areas and corrective actions, and progress in the area of project governance.

Interim Report. The midyear interim report will include a brief summary of the status of all M&O activities, including a section on the overall status and performance of the data handling and detector systems. It will also include highlights and accomplishments, specific comments on detector performance such as uptime and scheduled maintenance, failures, software releases and test results, major procurements planned or placed, an assessment of the overall labor effort, and any other performance data that is needed to characterize the overall data system performance.

Final Report. The final report will include a summary of all 60 months of the IceCube M&O award.

Common Fund Annual Report. The Common Fund (CF) report is prepared by the IceCube resource coordinator on an annual basis. The IceCube M&O Common Fund was created 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. The Common Fund report summarizes the status of past CF contributions and expenditures. In addition, the report includes a list of the major annual upgrades to the South Pole System (SPS), South Pole Test System (SPTS), UW data warehouse and UW data center.

Annual Reviews. NSF will conduct reviews of the IceCube Management & Operations activities through annual site visits of cognizant program officers that will address management issues, cost and performance objectives, and scientific and technical performance. The NSF may also conduct site visits and reviews on special topics. An external panel review covering, at a minimum, project management, cost and performance objectives, and scientific and technical performance will be organized after the second and fourth project years to inform NSF's decision on potential pathways for the support of IceCube M&O activities beyond 2021. NSF will invite IOFG members to participate in this review.

4.2 Management & Operations Plan

Building on our past experience, we have developed a plan to maintain and operate the detector and manage our collaboration resources to go from raw data to physics publications in a timely and efficient manner. Our plan maximizes IceCube's scientific potential and educational value by distributing both analysis and M&O tasks among collaborators. This structure draws the best expertise from collaborating institutions while also offering opportunities to educate scientists and engineers through hands-on experience with IceCube.

We provide accountability mechanisms in MoUs and strong leadership to coordinate distributed resources. In this section we present our plan by explaining how we will perform each task required to meet the technical requirements and specifications described at a top level in Section 3 and listed in detail in this section.

The operations organization has six primary elements: program coordination, detector maintenance and operations, computing and data management services, data processing and simulation services, software, and calibration:

1) Program Coordination: management and administration, engineering and R&D support, USAP support and safety, coordination of education and outreach, communications and other services typically provided by a scientific host laboratory.

2) Detector Maintenance and Operations: run coordination and winterover personnel, data acquisition (DAQ), online filters (PnF), detector monitoring, experiment control, surface detector operations, and supernova system.

3) Computing and Data Management: filtering data at South Pole for satellite transmission, incorporating data into the data warehouse; maintenance of data warehouse and UW data center and support the distributed computing infrastructure. Maintenance of data archiving system, networking and security infrastructure, core online/offline software code repository and build system.

4) Data Processing and Simulation Services: manage and execute production of offline data processing to generate general-purpose Level 2 data (with refined event reconstruction) and science working group specific Level 3 data (with event reduction defined by channel working groups). Manage and execute production of simulation data to provide sufficient cosmic-ray-induced and neutrino-induced events in the IceCube observatory for all physics working groups. Produce simulation data at the general-purpose Level 2 and provide the tools for working groups to produce Level 3 data. Coordination of the regular release of science-level experimental data from refereed-journal published results and the release of real-time alerts authorized by the Real-Time Oversight Committee.

5) Software: managing simulation software tools and maintain detector simulation software (IceSim), maintain and verify simulation of event generation, photon propagation and geometry calibration. Develop core common reconstruction tools in order to process raw waveform data to ultimately reconstruct muon tracks, shower events, direction, energy, and background probability of in-ice events, as well as to reconstruct cosmic-ray air showers. Develop and maintain high-level analysis tools to maximize the efficiency of turning reconstructed data into physics results. Perform data quality checks to support final selection of science-ready data and coordinate, develop and monitor common reconstruction for offline data processing.

6) Calibration: translation of detected quantities such as time and charge into physical quantities such as particle type, energy, and direction requires an understanding of the detector response to photons travelling through the glacial ice medium. Moreover, it demands an understanding of the ice itself. This area manages the activities and data needed to achieve both of these goals.

4.2.1 Program Management

4.2.1.1 Program Administration

The primary program administration task is to ensure that the resources needed to perform each task, regardless of source, are available when needed and used efficiently to accomplish the task requirements and achieve IceCube's scientific objectives.

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.

Financial Management. We manage IceCube finances, including NSF funding, a Common Fund supported by cash from the European and Asian Pacific collaborating institutions, and in-kind contributions from collaborating institutions, providing accountability through an audit trail for all funds regardless of source. The complete description of the funding sources can be found in section 5.1 "Funding Sources."

Performance Management and Reporting. In cooperation with NSF, we establish performance measures that are meaningful to evaluating our performance against M&O objectives. Performance measures are shown in Figure 4.2-1. We also establish with NSF a set of reporting deliverables that fulfill NSF internal and external requirements for oversight.

Key Performance Indicator	Annual Objective	Rationale
Detector Uptime	99%	Key performance measure of time that the detector was sensitive to transient astrophysical events or signals
Detector Clean Uptime	95%	Key indicator of production of pristine data for physics analysis with no contamination and no serious alerts
Supernova System Uptime	95%	Performance measure of time that the detector was sensitive to Galactic core-collapse supernova events
IceCube Live Uptime	99.9%	Critical to ability to resolve detector performance issues
South Pole System Uptime	99%	Critical to collection and storage of data
Latency Raw to L2 Processed Data	2 weeks	Demonstrates maturity in the processing pipeline, reduces time to release of highly significant astrophysical phenomena to the community

Figure 4.2-1. Proposed Performance Measures

4.2.1.2 Engineering and R&D Support

The engineering and R&D tasks are limited to the minimum tasks required to support day-to-day operations of the detector. R&D supports efforts to address electromagnetic interference (EMI) in the Dark Sector, snow depth mitigation for the IceTop array, enhancements to the performance of the IceCube Laboratory electronics and computing, and the ability to interface with externally funded R&D activities, especially those that intend to use the IceCube facilities, infrastructure, or data flow.

4.2.1.3 USAP Support and Safety

IceCube personnel prepare detailed support requirements and identify the most cost effective approach to meeting the requirements, through the annual planning cycle, direct communication with the NSF Antarctic Support Contractor, and the submission of the Support Information Package (SIP). Safety of personnel and equipment is vital to the smooth operation of a facility. This area is also responsible for ensuring that IceCube M&O personnel and collaborators participating in the USAP comply with USAP and follow good practice when deploying to the South Pole or otherwise working in the context of the program. Pre-season planning reviews will streamline logistics and address safety concerns.

4.2.1.4 Education and Outreach (E&O) Coordination

As a part of the IceCube Collaboration MoUs, each member contributes support to E&O. The associate director for education and outreach, working with NSF and IceCube leadership, establishes E&O priorities, provides support to ongoing activities, and responds to outside requests that support priorities by identifying appropriate resources within the Collaboration, assigning tasks, and providing oversight. Figure 4.2-2 describes examples of ongoing and high-impact IceCube E&O activities.

E&O Activity Title	Description	Benefit
Enhance STEM interest and understanding of K-12 students	Reach high school students and teachers through IceCube Masterclasses, internships, and the Upward Bound program, and K-12 classrooms through webcasts from the South Pole.	Provide connections to ongoing science and to working scientists and staff, including targeted programs for underrepresented groups.
Undergraduate research experiences and high school teacher South Pole deployments	Increase STEM awareness through undergraduate research experiences and South Pole deployments for high school teachers who are integrated into the UWRF Upward Bound (UB) program.	Provide firsthand science experiences for undergraduates and high school teachers, provide role models to inspire next generation STEM professionals.
Developing captivating media to communicate science	Craft accessible multimedia resources that describe IceCube science and technology. 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 , and an IceCube virtual reality game for the Oculus system.	Contribute to the NSF Broader Impact goal to "build the capacity of the Nation to address societal challenges using a suite of formal, informal, and broadly available STEM educational mechanisms."
Communication workshops	Build internal E&O capacity by developing and implementing semiannual workshops on communication skills and DEI in conjunction with IceCube Collaboration meetings and develop and lead the Multimessenger Diversity Network (MDN).	Strengthen ability of STEM practitioners to communicate science and technology in accessible language appropriate for an intended audience of STEM professionals, increase awareness and support, and broaden STEM participation at all levels through MDN training and activities.

Figure 4.2-2. Examples of E&O Activities

4.2.1.5 Communications

IceCube communication efforts, in support of the collaboration, include disseminating information about IceCube and our activities through press releases, news articles, and other means to reach audiences that are growing both in number and in diversity. We will also increase audience reach through our digital channels (website and social networks), produce more multimedia content, and develop additional multilingual materials, especially in Spanish and Portuguese. Finally, we will continue working with organizations representing underrepresented communities. We will increase the visibility of IceCube in their communities through articles in their media and participation in their training and outreach programs, and by providing opportunities for field deployments at the South Pole .

During the fourth year of this grant, we expect to produce 2-3 press releases and several dozen news articles for our website. In addition, we will continue production of digital media to explain the IceCube detector and its science, from neutrino astronomy to neutrino oscillations and other topics, to lay audiences. One goal of these products is also to show the diversity of the IceCube team, especially highlighting the contribution of young researchers. Social media will receive greater emphasis with regular posts designed to appeal to more diverse audiences.

4.2.2 Detector Maintenance and Operations

The IceCube detector maintenance and operations manager is accountable for the overall performance of the people, hardware and processes required to execute the operational plan of the detector at the South Pole in order to acquire high-data quality, meet necessary data throughput rates, provide appropriate technical documentation, maintain a problem reporting system, maintain a software library and revision history, and demonstrate overall system sustainability.

The detector operations manager holds weekly phone calls with the detector operations group on run coordination and detector operations matters, prepares periodic reports to NSF, prepares budgets, manages expenses, serves as a member of the Coordination Committee, resolves personnel matters, organizes planning for the austral summer, supports the SPTS, and is generally responsible for the overall coordination and performance of the detector through management of subsystem leads.

4.2.2.1 Run Coordination

During normal operations, the run coordinator ensures that data is being taken with high uptime and that the data is of the highest quality, with emphasis on data stability. The austral summer brings increased activity to the detector through planned maintenance of the computing, networking, and detector subsystems as well as dedicated calibration campaigns.

The run coordinator oversees the detector-related activities of subsystem experts and operators both at the South Pole and in the Northern Hemisphere, carefully documenting the run operation and auditing its effects on the data. All special operations requests are reviewed by the run coordinator to ensure the stability of the detector. Documentation and communication of run coordination activities include weekly monitoring reports from monitoring shifters, automated e-mail alerts on error conditions, weekly winterover reports, and other communications with stakeholders using a variety of media.

4.2.2.2 Winterover Operations

Maintaining the IceCube detector system is a complex challenge and requires our on-site operators to have a complete understanding of each individual detector subsystem. Their daily routine includes monitoring detector hardware, data acquisition software, and physics data quality assurance. On-site operators need to respond immediately to any problems with the data center in order to ensure continuous data taking. This includes replacing failed equipment as well as monitoring overall system load and identifying potential bottlenecks in the data processing. They also maintain a large disk array at the South Pole to redundantly archive and index the generated data.

In addition to their day-to-day duties in keeping the detector systems running, on-site operators regularly participate in outreach activities organized by the NSF or IceCube's Education & Outreach department at WIPAC.

4.2.2.3 Data Acquisition (DAQ)

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 the SPTS and/or PCTS prior to deployment at the SPS.

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 investigate alternative solutions using modern electronics as a contingency plan in case an upgrade is necessary to maintain the operation of the detector.

DAQ software collects raw hits from the individual DOMs, rejecting noise hits and forming triggers with all the relevant data for physics events in the detector. Additionally, all raw hits from the detector are buffered for a limited time and can be saved for analysis in special cases (HitSpooling). Diagnostic and calibration data are also collected, as are raw counting rates for all DOMs, used for the supernova triggers. Performance and reliability 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. The software engineers also maintain interfaces to other online systems, including the supernova DAQ and detector monitoring. In addition to adding features to the DAQ to support evolving science needs, DAQ software engineers also refactor and optimize the DAQ to facilitate seamless integration of future extensions of IceCube, such as the IceCube Upgrade.

Collaboration physicists from physics working groups, using Monte Carlo simulation of signals, develop new triggering algorithms for use in the DAQ. Physics working groups propose new trigger algorithms to the Trigger Filter and Transmission (TFT) Board. Once approved by the board, the triggers are adapted, tested and deployed within the DAQ triggering system.

The DOM firmware consists of a low-level FPGA design that controls the DOM hardware. A DOM firmware engineer supplies required FPGA modifications, maintains the code base, and updates documentation as needed. New physics requirements and hardware/software upgrades during the experimental program may require additional features in either the DOM or DOR FPGA designs.

A small fraction of DOMs (0.4%) have malfunctioned and must be operated as part of normal data taking in a nonstandard configuration. A typical solution is to bypass the failed or malfunctioning component within the DOM or to bypass the DOM completely. The detector operations group, working with the winterovers, excludes problem DOMs from the array and creates new standard run configurations as needed, tracking problem DOMs and performing studies on problem DOMs to develop solutions or workarounds that minimize impact of malfunctions on data quality.

4.2.2.4 Online Filters (Processing and Filtering—PnF)

The volume of data produced by the data acquisition system exceeds the limited bandwidth available in IceCube’s TDRSS satellite allowance. An online processing and filtering (PnF) system is used to apply a set of first-level event selections to the collected data, transmitting only those selected events. PnF system expertise is required to maintain the online system, ensure filters are being properly applied, and respond to and debug unexpected errors.

PnF system experts 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 that allow for more efficient use of IceCube’s satellite bandwidth and have 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 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. The online follow-up system sends a continuous stream of neutrino events to servers in the Northern Hemisphere, using the I3MS Iridium system, and a real-time alert system is in place that notifies other experiments of astrophysical neutrino event candidates within minutes, enabling multimessenger follow-up observations.

4.2.2.5 Detector Monitoring

The IceCube Live detector monitoring system provides a comprehensive set of tools for assessing and reporting of data quality. It collects raw subsystem data on the SPS during and on completion of a run. It then sends these data to the Northern Hemisphere via satellite where they are processed and presented through a web-based user interface. The system is critical to the ability to perform short-term and long-term analyses of detector performance.

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 or if action needs to be taken to correct malfunctions.

Problems can occur with individual DOMs, groups of DOMs, DOMHubs (entire strings), or racks of DOMHubs (groups of strings). Detector operators and winterovers must be alerted immediately when a problem occurs since the loss of a single DOM affects the overall quality of the data. Automatic alerting and automatic diagnosis of the problem help to limit the amount of time of a detector outage or degradation in data quality. The detector operations group and IceCube Live software engineers work with the SPS system administrator to maintain and develop the automatic alert paging and e-mail system.

4.2.2.6 Experiment Control

The IceCube Live experiment control system integrates control 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 for displaying the current and historical state 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 on SPTS to test upgrades and changes before deployment. Data quality designations for each run period are collected and indicate to the collaboration which data can be included for further processing and analysis.

The IceCube Live software engineers are accountable for uptime of IceCube Live and for maintaining, troubleshooting, supporting and evolving the interface to subsystems that control and monitor the detector. The software engineers continue to develop IceCube Live to integrate all subsystems, and add features as the behavior of the detector changes. During stable operations, the software engineers support physics working groups and operators to add needed functionality and respond to evolving science needs.

4.2.2.7 Surface Detector 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 and 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 the study of solar particle activity and high-altitude weather in addition to 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.

A plan has been developed for restoring the efficiency of the surface array and enhancing its scientific potential by installing low-cost scintillator modules and radio antennas above the IceTop stations. During the 2017/18 austral season, two co-located surface array stations have been deployed and connected to the ICL using a new fiber and power network. During the 2019/20 season, one station will be upgraded with production electronics (the other will be removed). The full surface array plan is under review by the NSF.

IceACT consists of two prototype air Cherenkov telescopes (ACT) deployed on the roof of the ICL and in the field at the prototype surface array. The instrument uses a 50.7 cm Fresnel lens with a 12º field of view and a SiPM camera to detect the atmospheric Cherenkov light from cosmic ray air showers during dark, cloud-free periods of the polar night. Air showers are measured in coincidence with IceTop. Coincident air shower events will be used to extend IceTop's energy reach to below 1 PeV, to validate IceCube's absolute pointing, and to calibrate the veto efficiency of IceTop for atmospheric muons and neutrinos. The electronics of the telescopes will be upgraded during the 2019/20 season. An internal review of IceACT scientific goals and technical plans is underway in order to define future efforts.

4.2.2.8 Supernova Operations

The supernova data acquisition (SNDAQ) receives the single photoelectron trigger scalar data produced by IceCube DAQ software and looks for a rate excess over the entire detector. For runs with no rate excess, the data are compressed to monitor the entire detector. In the event that an excess is found, an alarm is issued and sent via the IceCube Messaging System (I3MS), and more detailed data are saved, including all untriggered DOM readouts (HitSpooling). High-significance alarms are sent automatically to the Supernova Early Warning System (SNEWS), and external alerts from SNEWS trigger automatic archiving of IceCube HitSpool data.

Collaboration scientists are responsible for SNDAQ development. Core software engineers are responsible for integrating SNDAQ into the experiment control, monitoring, and DAQ systems. The supernova working group maintains a shift system to ensure that at least one monitor is checking alerts at all times.

4.2.2.9 Real-time Alerts

The real-time alert system is a collection of event filtering, communication, and analysis components responsible for detecting candidate astrophysical neutrino events and alerting the wider scientific community for potential multiwavelength follow-up observations. Events are identified online with the PnF system and transferred to follow-up analysis clients in the Northern Hemisphere within approximately 30 seconds, using the low-latency Iridium RUDICS system (I3MS).

Event times, reconstructed directions, and energy estimates are sent directly to other astronomical observatories, released to the AMON multimessenger network, and/or published online via the GCN transient alert system. System oversight and prompt decision-making on individual alert disposition is handled by the Real-time Oversight Committee. Core software engineers are responsible for the communications infrastructure facilitating the alert system (I3MS), integrating the real-time alert system into IceCube Live, and supporting the HitSpool archival system.

The real-time alert system is also responsible for automatic follow-up of alerts from other observatories, e.g., gravitational-wave alerts from the LIGO/VIRGO Collaboration. Rapid follow-up neutrino searches are run automatically, and requests are sent via IceCube Live and I3MS to archive HitSpool data around the time of the alert.

4.2.3 Computing and Data Management

The computing and data manager is accountable for the overall performance of the personnel, hardware, software, and processes required to support IceCube computing and data management from detector event to publication. The manager holds weekly teleconferences on operations issues, provides input to status reports to NSF, prepares and manages budgets, serves as a member of the Coordination Committee, and develops long-term strategies to maximize the benefit to IceCube science from evolving computing and data management technologies.

4.2.3.1 Data Storage and Transfer

Data is transferred from the South Pole using two mechanisms: 1) short messages and monitoring information over a system using Iridium satellites and 2) over the dedicated high-capacity SPTR (South Pole TDRS Relay) system for the bulk of IceCube data. About 100 GB of filtered data is transferred daily via satellite. The total amount of data generated by IceCube in one day is about ten times larger, at 1TB. Two copies of this data are archived on disk at the South Pole and are shipped to UW–Madison once a year, during the austral summer, for long-term archive.

IT specialists monitor and archive the data transfer from the South Pole. They address specific high-priority transfer needs that are requested by the Detector Operations team during data taking or any data related issues that might arise for the near real-time data processing.

A software engineer maintains the data handling software JADE. JADE manages the data collection and storage at the South Pole, the satellite data transfers, the data ingest at the UW–Madison data warehouse and the replication to external archive sites for long-term preservation. One of the goals of this project is to improve the functionality of the data handling software by adding missing key features such as the capability to efficiently index the metadata. A new File 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. A Final Analysis Sample application will also be developed that will make use of the File Catalog to uniquely identify and track the metadata associated with the final data samples produced by researchers in the analysis phase. These tools are part of our effort to improve traceability and reproducibility of the IceCube scientific results.

Data from the detector is processed, analyzed, and stored in intermediate and final stages at the UW–Madison disk storage infrastructure. System administrators operate the data storage infrastructure and ensure that active data is available and that the system provides the required performance and capacity. The storage system administrators also handle periodic hardware and software upgrades to the storage infrastructure and take care of cluster file system operations such as accounting, quota management, disk server load balance, etc. In order to benefit from technological improvements in storage density and energy efficiency, the baseline plan is to replace the disk systems every five years.

Of the 1 PB of data generated every year, about 700 TB will need to be archived and preserved in the 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 operating them. In order to provide cost-effective long-term data preservation services, we leverage large data centers at collaborating institutions NERSC and DESY-Zeuthen that already operate on a large scale and can provide long-term data archiving and curation as a service.

The first version of the new IceCube long-term archive (LTA) service part of the JADE application software started operating in September 2016, transferring data products to the NERSC and DESY tape storage facilities. In the first year of operation, about 2 PB of data have been archived to NERSC and 600 TB to DESY. The plan is to keep this archive stream constantly active while developing further JADE functionality that will allow us to steadily increase the performance and automation of archive data flow. As part of this activity we are collaborating with researchers at the Open Science Grid project to evaluate a 3rd party tool, Rucio, to manage data replication to multiple sites and consider it for a possible integration with JADE. Rucio is being used to manage data replication for the ATLAS and CMS experiments at CERN and the XENON experiment at LNGS. The goal is to build more robust software by leveraging existing tools with proven reliability and focus our in-house development effort on delivering IceCube specific functionality more effectively.

Having multiple copies of valuable data ensures its integrity and preservation. In addition, archived data is catalogued such that it is readily discoverable and accessible in the future. The consolidation of the operation of the LTA service represents an important step towards one of the overall goals for the IceCube computing: rolling out improved data management services that ensure long-term preservation and usability of the data as well as ease of discoverability.

4.2.3.2 Core Data Center Infrastructure

IceCube requires a flexible and highly available set of computer systems to support operations. Some are highly visible and transversal, such as e-mail or authentication services. Others are more focused and less visible but play equally vital roles in science outcomes. Examples of these are application servers to host real-time alert services or experiment monitoring services for remote shifters.

Given the size and distributed nature of the IceCube computing infrastructure, the network is a core component. One of the key issues in designing and operating the IceCube data center network is to ensure that it is optimized for high-volume bulk data transfers while keeping critical infrastructure secure. For bulk data movement, a Science DMZ has been implemented to provide adequate performance and appropriate compensating controls to ensure the security of those services.

UW–Madison networking provides an excellent service connecting the IceCube data centers on campus with dedicated links up to 100 Gbps in capacity and providing high bandwidth Wide Area Networking (WAN). The new Fitchburg data center, where the IceCube data and core services are being migrated, will be equipped with the highest capacity network connection on campus, at 200 Gbps.

IceCube network administrators work in close coordination with UW–Madison networking experts to ensure the IceCube network performs optimally and with maximum uptime. This includes maintenance, support, configuration, and customization of the interface network elements when necessary. They also monitor the health of the devices and configurations to identify system bottlenecks and potential hardware problems and analyze security logs for suspicious behavior and traffic signatures. During operations, network administrators respond to the needs of scientists, software developers, project engineers, and detector operators to maximize network reliability and provide customized solutions to optimize performance.

IT specialists maintain a cybersecurity program at the UW–Madison data center. These activities 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 that is 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 remain a priority. In addition, we 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.

4.2.3.3 Central Computing Resources

In order to efficiently process and analyze the large amounts of real and simulated data involved in the IceCube analysis process, a large HTC cluster is available for the Collaboration at the UW–Madison data center. 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 key part of the IceCube simulation chain. 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 5500 CPU and 432 GPU job slots. The cluster uses the HTCondor software for job scheduling and management.

IceCube system administrators maintain and operate the cluster. They collaborate closely with the HTCondor development team at the computer sciences department of UW–Madison, providing feedback on specific use cases and ensuring the system fulfills IceCube's evolving needs.

In order to benefit from technological advances and improvements in energy efficiency, the baseline plan is to replace the servers in the HTC cluster every five years. If hardware is viable past this initial 5-year period, we will continue to utilize it either in the HTC cluster or for testing new technologies, such as Kubernetes. A wholesale replacement of hardware has not occurred in several years; rather, there have been incremental changes or additions. In the past four years, only GPU compute capability has been added.

4.2.3.4 Distributed Computing Resources

To obtain the computing resources required to process vast amounts of data, IceCube relies on distributed resources available from Collaboration institutions.

In addition to the collaboration resources, IceCube will continue to tap into opportunistic resources (mostly from Grid projects in the Open Science Grid in the US and European Grid Infrastructure in Europe) as much as possible. This is an effective and efficient way to produce larger statistics of simulated 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) have GPU nodes. This makes them especially attractive for their potential in increasing IceCube simulation capabilities. IceCube started requesting computing time allocations from XSEDE in 2016, which were successfully renewed in 2017, 2018, and 2019. With this activity, we have successfully demonstrated sustained CPU and GPU usage in several XSEDE resources throughout the year. We will continue renewing XSEDE allocations in order to consolidate this contribution to the overall IceCube computing capacity and also to continue collaborating with the HPC community to ensure these specialized resources can be effectively utilized by IceCube.

The efficient use of distributed resources requires coordination among the different sites as well as the use of Grid software such as job meta-scheduling and data access tools. IceCube system administrators maintain and operate the core services of this IceCube Grid infrastructure. Wherever possible, standard tools are used to manage the Grid resources, which engineers then interface with IceCube specific software. In order to manage this process efficiently, it is essential to maintain close contact with the distributed scientific computing community. We ensure this by participating in the Open Science Grid (OSG) project and the National Data Service initiative. One of the goals of this project is to integrate the IceCube HTC cluster into OSG, contributing back to this large research infrastructure that is so important for our mission.

Besides OSG and XSEDE, we are also working on closer relationships with the Pacific Research Platform, pilot projects for the National Research Platform, and San Diego Supercomputing Center (SDSC). These collaborations have been fruitful in 2019, such as a proposal for a GPU cluster and an experiment with cloud resources shown at Supercomputing 2019. The data generated in this large distributed infrastructure is transferred back to the central data warehouse at UW–Madison using high-throughput links and the GridFTP protocol. IceCube system administrators manage the middleware services needed for providing high-performance remote access to the data. The deprecation of the GridFTP protocol has led to searching for alternatives for both mass storage and data transfer. We are currently exploring a number of technologies for the future replacement of the current distributed filesystem.

4.2.3.5 South Pole System (SPS)

The SPS is a computing system developed and maintained by IceCube system administrators that supports the data acquisition and filtering tasks carried out by the detector 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 operating system and configuration management tools used on the SPS are the same as the ones used for all services in the UW–Madison data center. This allows system administrators to apply consistent procedures across systems and efficiently manage version control, patching, software updates, monitoring, and maintenance.

The computer servers in the SPS will be replaced on average every four years in order to profit from technological advances that maximize computing power per kilowatt and minimize the risk of component failures.

The IceCube network connects the detector systems in the ICL and the South Pole station with the USAP network and, through that, with the data center at UW–Madison. It must comply with policies and regulations of NSF and the University of Wisconsin. The SPS system administrators are responsible for uptime and performance optimization of the IceCube network, including maintenance, support, configuration, and customization of the system when necessary. Network support tasks also include monitoring the health of the devices and configurations to identify system bottlenecks and potential hardware problems. The SPS systems are isolated from the USAP and other external networks by means of a firewall. Security logs are monitored for suspicious behavior and traffic signatures.

4.2.3.6 South Pole Test System (SPTS)

The primary purpose of the SPTS is to build and test software and hardware in advance of operational deployment in the South Pole System (SPS). Software developers use the SPTS to debug system changes safely in a non-production environment. The close physical and logical match to the SPS allows system maintainers to verify hardware and identify potential system side effects introduced by software upgrades or configuration changes.

IceCube system administrators are responsible for hardware maintenance and operations of the SPTS. During testing, system administrators support software developers and engineers to maximize hardware reliability and provide customized solutions to increase testing time.

Prior to any hardware upgrade in the SPS, thorough evaluation takes place in the SPTS in order to validate and select the best platform to be deployed in production.

DOM hardware engineers maintain and upgrade the system to ensure maximum uptime when the system is required for testing. They provide support to those wishing to add features as required in response to new science needs and to evolve the functionality of the SPTS as appropriate.

The SPTS operating systems software is based on the Scientific Linux distribution and it is kept in synchronization with the SPS system. System administrators are responsible for system maintenance, troubleshooting, and upgrades for the SPTS operating systems. The same mechanisms as in SPS are used for system configuration control, monitoring, and patch management.

4.2.4 Data Processing and Simulation Services

The data processing and simulations services manager is accountable for the overall processing of calibrated data products, including generation of Monte Carlo simulations and offline filtering of data collected by the IceCube detector to support physics analyses by the IceCube Collaboration and final data products to be made publicly available to the scientific community. The manager holds regular teleconferences on production issues and serves as a member of the Coordination Committee in order to work with managers responsible for computing and data management, calibration, and detector operations to address critical issues related to production and to develop long-term strategies to maximize the benefit to IceCube science.

4.2.4.1 Offline Data Production

Data arriving in the north are compressed and stripped of all unnecessary information to conserve transfer bandwidth. In a first processing step, the data must be unpacked and uncompressed, and calibrations must be applied to these data to convert raw DAQ measurements into physical quantities. The reconstructions used at the South Pole to form the filter decisions must then be reapplied to the calibrated data and all intermediate results stored together with the data to allow studies of the filter performance. A software engineer monitors the execution of the processing scripts and verifies regularly the quality of the data.

The complex reconstructions required allowing the suppression of the high muon background from cosmic-ray-initiated air showers from the neutrino signal are computationally intensive. To make the best use of the computing resources in the IceCube Collaboration, these reconstructions must be run centrally and results made available in the data warehouse for consumption by the different physics analysis working groups. Execution of the processing scripts is actively monitored and the quality of data is regularly verified by members of the production team. Plots of various reconstruction parameters are provided to the Collaboration through a web interface for quality assurance.

4.2.4.2 Simulation Production

Coordination of simulation production and resources involves management of multiple dependencies across M&O and the Collaboration. These include, for example, detector geometry calibration, charge and time calibration, and detector configuration uploaded into the database; maintenance of simulation software; and physics demand and dataset priority agreed with the Collaboration and matched with current computing infrastructure capacity. The simulation production coordinator is responsible for coordinating with other groups in the Collaboration to assess the impact of these tasks on physics analyses and understand issues involving computing infrastructure. The coordinator ensures proper production of data to verify simulation releases before full production is enacted.

In order to detect physics events caused by high-energy neutrinos, the large background of cosmic muons events must be rejected while retaining the highest signal efficiency. Simulation data are essential in this analysis procedure and a large number of cosmic muon events must be produced. The coordinator determines the computing and storage requirements for generation of Monte Carlo simulations and communicates with the data management manager to insure that these needs are met.

The simulation production coordinator along with the software coordinator organize periodic workshops to explore better and more efficient ways to meet the simulation needs of the analyzers. This includes both software improvements and new strategies as well as providing the tools to generate targeted simulations optimized for individual analyses. Simulation strategies help physicists study the impact that systematic uncertainties in our understanding of ice properties, hole-ice, and DOM sensitivity have on physics analyses.

Simulation of Cherenkov light in the ice is done by directly propagating individual photons accounting for the state-of-the-art properties of glacial ice. Such component of simulation needs graphics processing units (GPU). GPUs are an essential component of the distributed computing infrastructure. The addition of data filtering and processing adds further complexity. This requires a dedicated framework to coordinate data set management and result tracking.

The coordinator works with working group technical leads to determine the computing and storage requirements for generating Monte Carlo simulations and communicates with the data management manager to ensure that these needs are met.

IceProd is a database-driven scheduling and management software package that catalogs simulation data sets and optimizes the usage of computing resources. As recommended by the SCAP, the system now incorporates third-party Grid middleware products to reduce long-term maintenance associated with an entirely in-house framework. The new IceProd2 model is based on a single central scheduler that relies on GlideIns running at satellite sites. Production throughput on IceProd has continually increased as a growing number of dedicated and opportunistic resources has become available and a number of code optimizations have been implemented. A new set of monitoring tools is currently being developed in order to keep track of efficiency and further optimizations.

A physicist supports this task by performing runtime basic data checks to verify evident configuration errors; low- and high-level data verification by comparing simulation data from different production sites and different historical simulation releases to experimental data; and analysis-level data checks expected by working groups for the very early stages of physics analyses. The physicist maintains the simulation production web portal to keep all stakeholders informed of simulation production status and issues. This group is currently managed by a staff scientist at UW Madison.

4.2.4.3 Public Data Products

The IceCube Collaboration already provides public access to event reconstruction information for events selected as neutrinos for specific published results as well as real-time alerts to other observatories around the world in order to provide early warning of interesting cosmic events, share data, and combine observations through collaborations defined through Memoranda of Understanding (MoUs). 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.

4.2.5 Software

The management of the software effort is divided into five groups, each with an individual manager overseeing the effort. The five groups are: core, simulation, reconstruction, science support tools, and infrastructure. The IceCube software coordinator is responsible for coordinating maintenance of all five software groups. The software coordinator also conducts regular software training sessions for IceCube collaborators.

Future efforts will focus on improving the speed, efficiency, and robustness of production software through workshops and code sprints. 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. This has recently been made the focus of the IceCube Software Strike Team, a team consisting of a core group of 10 collaborators, concentrating on three high-priority projects aimed at improving production resource utilization.

4.2.5.1 Core Software

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. This group is currently managed by the software coordinator.

A new database system is now in production, which leverages detector status information in the IceCube Live monitoring system. This precludes the need for separate databases previously maintained by the Mons group.

4.2.5.2 Simulation Software

IceCube's simulation software has to cover a wide dynamic range, supporting low energy at the GeV scale as well as ultrahigh energy at the EeV scale. The IceCube simulation software consists of a set of software modules designed to run within the IceTray framework. Each module is responsible for different aspects of the Monte Carlo simulation chain, starting with particle injection and propagation through photon propagation on GPUs as well as detector response and DAQ trigger. Development and maintenance of each of these software components is handled by developers across the IceCube Collaboration. The modular design, allows for individual developers to write each part of the code independently but provide a common interface to other modules in the simulation chain. Recent additions include a project that will vastly improve the generation of systematics neutrino datasets and a project that significantly improves the generation of background datasets. This group is currently managed by a staff scientist at UW Madison.

4.2.5.3 Reconstruction Software

IceCube's reconstruction software runs online for filtering at the Pole and offline in the north, for higher filter levels and as a starting point for analysis. The reconstruction software consists of more than 60 extra projects on top of core software. Recent additions include a project that will better utilize production resources, allowing access to more compute nodes. The reconstruction software is managed by a postdoc from UW–Madison.

4.2.5.4 Science Support Tools

This group is a spin-off from the reconstruction group, focusing on common tools used beyond filtering at L2, including IceCube's open source effort. There are several large projects slated for inclusion into IceCube's open source repository, including SkyLab (point source library) and PISA (oscillation tools). This group is managed by a postdoc at UW–Madison. IceCube's open source effort is managed by a computer scientist at UW–Madison.

4.2.5.5 Software Development Infrastructure

Software development in IceCube is a worldwide, distributed effort with more than 100 contributors and running on several different platforms to maximize grid resources. IceCube relies on several software development tools, such as a central repository (Subversion), a ticketing system (Trac), and a continuous build test system (Buildbot). This infrastructure is managed by an engineer at UMD.

1.1.1. Calibration

IceCube calibration provides a common set of detector calibration constants that translate IceCube DOM signals into recorded charge and time, which are then used to reconstruct neutrino properties such as energy, position, time, direction, and flavor. As photons typically propagate tens to hundreds of meters in ice before being recorded by a DOM, the measurement of the optical properties of the ice is a critical part of the calibration process. Calibration methods and devices include the onboard calibration electronics, LED flashers, which are co-located on each DOM, cosmic ray muons, calibration lasers, and several recoverable borehole loggers. Calibration constants are stored in the IceCube database or in the software repository and documented on the internal wiki and Docushare systems.

1.1.1.1. Detector Calibration

The time synchronization of the DOM internal clocks with the master clock on the surface is continuously monitored with the Reciprocal Active Pulsing procedure, the results of which are available in the IceCube monitoring system. Every year, the in-ice DOM discriminator thresholds, high-voltage settings, and other DOM constants are calibrated using the onboard electronics system (DOMCal) and adjusted using data from cosmic ray muons in the ice. IceTop DOMs are calibrated once per month. The run coordinator organizes the calibration runs, which are performed by the winterover personnel, in order to minimize detector down time. IceCube postdocs and students vet this data under the supervision of the calibration coordinator, report the data at weekly teleconferences, and archive the results on the wiki. The calibration constants are stored in the database. The DOM digitizer baseline and PMT and electronics gain are the most important calibration constants for converting raw DOM waveforms to measured voltage and are set to their final values using data from physics runs. The digitizer baselines are monitored continuously with forced-trigger "beacon" launches. The beacon waveforms are collected in the IceCube Live monitoring system and are stored in the IceCube Live database. The PMT gains are continuously monitored with collected charge from cosmic ray muons interacting in the ice, also collected in IceCube Live. Baselines are stable to within 4 microvolts and the gains for most DOMs are stable to within 1%. Monitoring shifts taken in turn by all collaboration members report any deviations, which are then followed up by the calibration coordinator and the operations coordinator as needed.

4.2.5.6 Ice Properties

The ice consists of two components: the untouched “bulk ice” between strings and the refrozen “hole ice,” which was melted during IceCube construction and drilling. The optical absorption and scattering lengths of the bulk ice are measured using the LED flashers, which are located on each DOM. Flasher data is collected in dedicated calibration runs during the Antarctic summer, organized by the calibration coordinator, with the data stored in the IceCube data warehouse and documented on the wiki. In addition, several borehole logging devices have been developed for use in SPICEcore, which is located about 1 km from IceCube and extends to about 1700 m depth. These devices are used to measure additional bulk ice properties, e.g., ice luminescence and UV transparency, that are not possible with the currently deployed calibration devices in IceCube and that will benefit several IceCube analyses.

It has been demonstrated with LED flasher data that the scattering in the bulk ice is anisotropic, having a preferred direction aligned with the ice flow. More recent studies have shown that there is also a depth dependence to the strength of the anisotropy. Close communication and collaboration with the glaciology community have led to new understanding of the underlying processes that likely cause the observed anisotropy. Developing new models and parameterizations of these processes will be a strong focus of the calibration group moving forward.

The refrozen hole ice is an active area of study, as it modifies the angular acceptance of the DOM due to bubbles trapped in the refrozen column. The calibration of this bubble column is extremely challenging with the current devices deployed in IceCube. Many of the planned IceCube Upgrade calibration devices will have improved sensitivity to measure the properties of this bubble column. The calibration group oversees some of the design efforts of these devices to ensure these local ice properties can be measured with the required precision. Progress of all these analyses is reported in weekly teleconferences, with the results archived on Docushare.

5 Cost Overview

IceCube M&O finance management includes NSF funding, a Common Fund supported by cash payments by European, Canadian and Asian Pacific collaborating institutions, and in-kind contributions from collaborating institutions, providing accountability through an audit trail for all funds regardless of source.

The M&O budgets are 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 Appendix 1 of this plan). These costs are very well understood and are based on actual experience during past years of M&O. There is no explicit budgeting for contingency as was done for the IceCube MREFC project.

5.1 Funding Sources

The NSF IceCube five-year M&O award covers federal fiscal years 2016–2021 (April 1, 2016–March 31, 2021). NSF intends to provide a total of $35,000,000 over the term of five years ($7,000,000 per year), with the support split equally between the Polar programs and physics divisions. The expectation is that annual increases typically expected due to escalation will be offset by efficiencies in the program.

In addition to the NSF M&O award, which also covers the U.S. annual contributions to the Common Fund (CF), other sources of funds for the M&O core activities are the European, Canadian, and Asia/Pacific annual contributions to the CF, NSF analysis base grants and institutional in-kind contributions.

Sources of Funds: there are four different sources of funds for the IceCube M&O program:

NSF M&O Core	This NSF award mostly covers labor, travel and partially capital, M&S & services for: • UW–Madison, six U.S. subawardees, and one UW shared grant • Core activities mostly under detector ops. & maint., computing & data mgmt., program mgmt. • U.S. cash contribution to the M&O Common Fund
NSF Base Grants	The NSF IceCube analysis base grants support labor and travel for: • M&O activities mostly done by graduate students and postdoctoral researchers • Data quality, reconstruction & simulations, calibration, monitoring, filtering & triggering
U.S. Institutional In-Kind	U.S. institutional in-kind contributions mostly cover labor, travel and M&S for: • M&O activities mostly done by faculty and administration members • Fellowships and university funded activities, computing power & cooling
Europe & Asia Pacific In-Kind	Europe, Canada & Asia Pacific in-kind contributions cover labor, travel and HW/SW for: • M&O activities done by non-U.S. scientists, engineers and other team members • Data quality, reconstruction & simulations, calibration, monitoring, filtering & triggering • Non-U.S. cash contribution to the M&O Common Fund that covers most of the capital equipment and service agreements for the computing upgrades

5.1.1 NSF IceCube M&O Award

The following figure describes the NSF M&O award budget by cost categories (Figure 5.1-1) for UW and all US subaward institutions:

Cost Category (including indirect)	**FY2016-FY2021**
Labor	$26,276
Materials & Supplies	$603
Travel	$1,464
Services and Service Agreements	$1,184
Subawards with U.S. collaborating institutions	$4,936
Capital Equipment	$538
Total	**$35,000**

Figure 5.1-1. NSF IceCube M&O Award – *Cost by Category (in $k)*

Labor: The primary basis of estimate for effort level is experience from executing identical or similar tasks in past years. Management judgments applied to estimates include whether past allocations were correct and the extent to which task over time will require the same, more, or fewer resources.

Materials and Supplies (M&S): Expenses related to computing infrastructure are the major cost driver in this category. Cost estimates support several different operational tasks. For example, planned operations require sufficient disk media at the South Pole to store two copies of the raw data, and sufficient resources for the northern data center to perform backups of various storage systems. The plan also includes personal computers, UPS batteries for the South Pole System, physical qualification examinations, calibration and engineering supplies, and spare items such as memory, disk or network expansions. Other M&S expenses include shipping and packing.

Travel: The budget is based on an estimated number of domestic and foreign trips, multiplied by total FTE for each labor category. The travel direct rates take into consideration airfare and transportation, lodging and per diem expenses. Travel expenditures include travel to domestic and foreign IceCube collaboration meetings, training, reviews, IceCube meetings, and travel expenses in Christchurch on the way to/from Antarctica.

Capital Equipment: These include expenditures for computing infrastructure upgrades.

Services and Service Agreements: Computing infrastructure and software maintenance services for the South Pole System, UW data center, data warehouse and networking are the major cost drivers for services and service agreements. This category includes maintenance contracts, licenses, operating systems, warranties, and technical support. Because of the need for high availability and reliability of computing infrastructure, we reduce risk through having service agreements with vendors of major commercial off-the-shelf equipment.

Subawards with U.S. collaborating institutions: The IceCube M&O roles and responsibilities of six U.S. institutional subawards and one UW shared grant are described in Figure 5.1-2.

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 University	Simulation software, simulation production
Univ. of Wisconsin–River Falls	Education and outreach coordination

Figure 5.1-2: IceCube M&O U.S. Subawards and Shared Grant – *PY4 (FY2019-FY2020) Major Responsibilities*

5.1.2 IceCube M&O Common Fund

The IceCube M&O Common Fund (CF) was created in April 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 to the CF based on the total number of the institution's Ph.D. authors.

The Collaboration updates the Ph.D. author count twice a year at the collaboration meetings in conjunction with the update to the IceCube M&O responsibilities in the institutional Memorandum of Understanding. Effective April 1, 2010, the annual established rate per Ph.D. author is $13,650.

Common Fund Expenditures. The M&O activities identified as appropriate for support from the Common Fund are those core activities that are agreed to be of common necessity for reliable operation of the IceCube detector and computing infrastructure. The activities directly support the functions of winterover technical support at the South Pole, hardware and software systems for acquiring and filtering data at the South Pole, hardware and software systems for transmitting data via satellite and disks to the UW data center, systems for archiving the data in the central data warehouse at UW, and UW data center operations as listed in the Cooperative Agreement with NSF.

The Common Fund expenditures are divided into two categories: U.S. Common Fund and non-U.S. Common Fund.

Common Fund Contributions. The planned contributions to the IceCube M&O Common Fund during the 12th year of IceCube operations (April 2018 – March 2019), is based on the Ph.D. authors head count in the institutional MoUs v24 from May 2018. The actual contributions were about $91K less than planned due to delayed payments by some of the collaborating institutions (Figure 5.1-3).

IceCube M&O	Ph.D. Authors, May 2018	Planned ($k)	Actual ($k)
Total CF Planned	**137**	**$1,870**	**$1,812**
U.S. Contribution	71	$969	$969
Non-U.S. Contribution	66	$901	$810

Figure 5.1-3. Planned vs. Actual CF Contributions – *Year 12 of M&O, April 1st, 2018 – March 31st, 2019*

The following table provides the most recent detailed breakdown of the Ph.D. authors headcount based on MoUs v.27.0, September 2019 (Figure 5.1.-4).

	Total Ph.D. Authors	Faculty	Scientists / Post Docs	Ph.D. Students
U.S. Institutions Subtotal	79	48	31	50
Non-U.S. Institutions Subtotal	73	44	29	83
Total U.S. & Non-U.S.	**152**	**92**	**60**	**133**

Figure 5.1-4. IceCube Collaboration – *Authors head count based on the institutional Memorandum of Understanding v27.0 (September 2019)*

5.1.3 Institutional In-Kind Contribution

In addition to the U.S. M&O core funds and U.S. base grants support, IceCube MoUs define in-kind contributions of distributed M&O labor and computing resources from collaborating institutions.

This represents a transition from a centralized management and funding approach during IceCube's construction phase to a more distributed model of management and funding for M&O. (Figure 5.1-5).

The distributed model results in increased financial contributions to the Common Fund and in-kind labor contributions to M&O tasks from European, Canadian, and Asia Pacific collaborators. It also results in a greater emphasis on direct NSF funding to U.S. collaborating institutions. In-kind contributions by each collaborating institution are included in the Memorandum of Understanding (MoU). (Summary of the MoU Scope of Work is included as Appendix 2 of this plan).

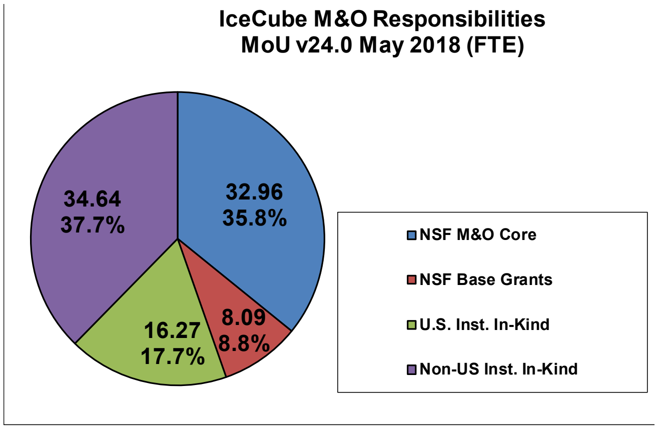

Figure 5.1-5. IceCube M&O Distributed Management and Funding Model (FY2018). *Based on the institutional Memorandum of Understanding v24.0 (May 2018)*

5.2 Computing Infrastructure Upgrade Plan

Computing infrastructure is the major cost driver in IceCube M&O expenses for capital equipment, computing services, and materials and supplies. The annual upgrade plan assumes consolidation of computing and storage infrastructure, with an expectation to upgrade the existing systems both at the South Pole and in the north every 4 years on average. The computing requirements for data analysis and the corresponding multiyear computing capacity planning are presented and reviewed periodically at the Software and Computing Advisory Panel (SCAP) meeting. The materials and supplies upgrade plan supports several different operational items, such as disk media at the South Pole to store raw and filtered data and sufficient spare parts for addressing operational issues in the South Pole System during the winter months. Other expenses include software purchases. Computing infrastructure and software both at the South Pole and at UW are also the major cost drivers for service agreements, which include warranties, technical support, licenses, and work by software programming consultants.

Appendix 3 includes a list of FY2016/2021 major IceCube purchases for the South Pole System (SPS) upgrade, for the South Pole Test System (SPTS) upgrade, and for the UW data warehouse and UW data center upgrades.

Appendix 1: IceCube M&O Work Breakdown Structure

May 2017

Appendix 2: IceCube M&O Memorandum of Understanding

Effort and Authors Head Count Summary

v 27.0 September 15, 2019		Authors Head Count				IceCube Authors: M&O Responsibilities (FTE)						
Institution	**Institutional Lead**	**Ph.D. Authors**	Faculty	Scientists / Post Docs	Ph.D. Students	WBS 2.1 Program Management	WBS 2.2 Detector Operations & Maintenanc	WBS 2.3 Computing & Data Management	WBS 2.4 Data Processing & Simulation	WBS 2.5 Software	WBS 2.6 Calibration	**Total**
University of Alabama*	Dawn Williams	3	(2	1	2)	0.25	0.20	0.00	0.05	0.05	0.30	0.85
University of Alaska	Katherine Rawlins	1	(1	0	0)		0.25			0.20		0.45
Clark Atlanta	George Japaridze	1	(1	0	0)		0.02					0.02
Drexel University	Naoko Kurahashi Neilson	2	(1	1	1)	0.25	0.10			0.55		0.90
Georgia Tech	Ignacio Taboada	1	(1	0	3)	0.05	0.75			0.10		0.90
LBNL*	Spencer Klein	4	(2	2	1)	0.13	0.24	1.08				1.45
Loyola University Chicago	Rasha Abbasi	1	(1	0	0)							0.00
Marquette University	Karen Andeen	2	(1	1	0)	0.20	0.55					0.75
Massachusetts Institute of Technology	Janet Conrad	2	(1	1	3)	0.60	0.00	0.3	0.0	0.60	0.00	1.20
Mercer University	Frank McNally	1	(1	0	0)							0.00
Michigan State University*	Tyce DeYoung	8	(5	3	9)	1.05	0.28	0.35	0.15	0.40	0.40	2.63
Ohio State University	James Beatty	2	(1	1	0)				0.05	0.25		0.30
Pennsylvania State University*	Doug Cowen	6	(2	4	3)	0.70	0.62		0.05	0.45		1.82
South Dakota School of Mines & Technology*	Xinhua Bai	1	(1	0	2)	0.28	0.02		0.20	2.00		2.50
Southern University	Ali Fazely	3	(2	1	0)		0.02	0.30		0.60		0.92
Stony Brook University	Joanna Kiryluk	1	(1	0	1)	0.15	0.10		0.10	0.10		0.45
University of California, Berkeley	Buford Price	1	(1	0	0)	0.10						0.10
University of California, Irvine	Steve Barwick	1	(1	0	1)		0.02					0.02
University of California, Los Angeles	Nathan Whitehorn	1	(1	0	0)	0.15				0.10		0.25
University of Delaware*	Tom Gaisser	9	(6	3	5)	0.55	1.15	0.10	1.05	0.55	0.25	3.65
University of Kansas	Dave Besson	1	(1	0	0)		0.02					0.02
University of Maryland*	Greg Sullivan	6	(3	3	2)	1.15	1.01	0.35	0.50	0.30	0.00	3.31
University of Rochester	Segev BenZvi	1	(1	0	2)	0.25	0.45			0.10		0.80
University of Texas at Arlington	Benjamin Jones	1	(1	0	2)	0.05				1.15	0.60	1.80
University of Wisconsin, River Falls*	Jim Madsen	3	(3	0	0)	0.60				0.10	0.20	0.90
University of Wisconsin, Madison	Albrecht Karle	15	(5	10	13)	2.23	2.56	0.00	1.60	0.65	2.15	9.19
Yale University	Reina Maruyama	1	(1	0	0)	0.05	0.05			0.05		0.15
U.S. Institutions Subtotal		**79**	**(48**	**31**	**50)**	**8.79**	**8.40**	**2.18**	**3.75**	**8.30**	**3.90**	**35.32**
DESY-Zeuthen	Marek Kowalski	10	(5	5	10)	1.20	1.32	1.10	0.00	0.40	0.25	4.27
Karlsruhe Institute of Technology	Andreas Haungs	5	(2	3	5)	0.20	1.30	0.15	0.45	0.35	0.20	2.65
RWTH Aachen	Christopher Wiebusch	2	(1	1	8)	0.75	0.32	0.30	0.55	0.75	0.75	3.42
Universität Dortmund	Wolfgang Rhode	2	(1	1	5)		0.73		0.30	1.00		2.03
Universität Mainz	Lutz Köpke	3	(2	1	5)	0.25	1.45			0.40		2.10
University of Münster	Alexander Kappes	2	(1	1	5)	0.40						0.40
Universität Wuppertal	Klaus Helbing	2	(1	1	5)		0.93			0.40	0.20	1.53
Humboldt Universität Berlin	Marek Kowalski	1	(1	0	1)						0.20	0.20
Universität Bochum	Julia Tjus	2	(1	1	2)	0.10	0.38			0.10		0.58
Technische Universität München	Elisa Resconi	4	(1	3	8)		0.05		1.00	1.10	0.20	2.35
Universite Libre de Bruxelles	J. A. Aguilar Sanchez	5	(3	2	3)	0.85	0.52					1.37
University of Gent	Dirk Ryckbosch	2	(1	1	4)	0.05	0.03		0.55	0.05		0.68
Vrije Universiteit Brussel	Nick van Eijndhoven	3	(3	0	3)	1.45	0.06		0.10	0.45		2.06
Stockholm University	Klas Hultqvist	3	(3	0	2)	0.40	0.11		0.25	0.20		0.96
Uppsala University	Olga Botner	3	(3	0	2)	0.35	0.41				0.15	0.91
University of Alberta	Darren Grant	3	(2	1	1)		0.60	0.40	0.30	0.30	0.60	2.20
University of Oxford	Subir Sarkar	1	(1	0	0)	0.20						0.20
University of Canterbury	Jenni Adams	2	(1	1	2)	0.20	0.15			0.15	0.15	0.65
University of Adelaide	Gary Hill	1	(1	0	1)					0.90		0.90
Chiba University	Shigeru Yoshida	5	(3	2	0)	0.10	0.20	0.80	0.00	0.70	0.10	1.90
Université de Genève	Teresa Montaruli	2	(1	1	2)	0.35	0.10			2.30		2.75
Universität Erlangen-Nürnberg	Gisela Anton	3	(2	1	3)	0.20				0.55	0.05	0.80
Niels Bohr Institute	Jason Koskinen	4	(2	2	1)	0.25	0.55		0.70	0.10		1.60
Sungkyunkwan University	Carsten Rott	2	(1	1	5)	0.15	0.69				0.20	1.04
Queen's University	Ken Clark	1	(1	0	0)					0.10		0.10
Non-U.S. Institutions Subtotal		**73**	**(44**	**29**	**83)**	**7.45**	**9.90**	**2.75**	**4.20**	**10.30**	**3.05**	**37.65**
Total U.S. & Non-U.S.		**152**	**(92**	**60**	**133)**	**16.24**	**18.30**	**4.93**	**7.95**	**18.60**	**6.95**	**72.97**
Changes since last official version are colored red												
*** IceCube M&O Subawardee Institutions**												

IceCube Collaborating Institutions

The following chart summarizes evolvement of the U.S. and non-U.S collaborating institutions over time.

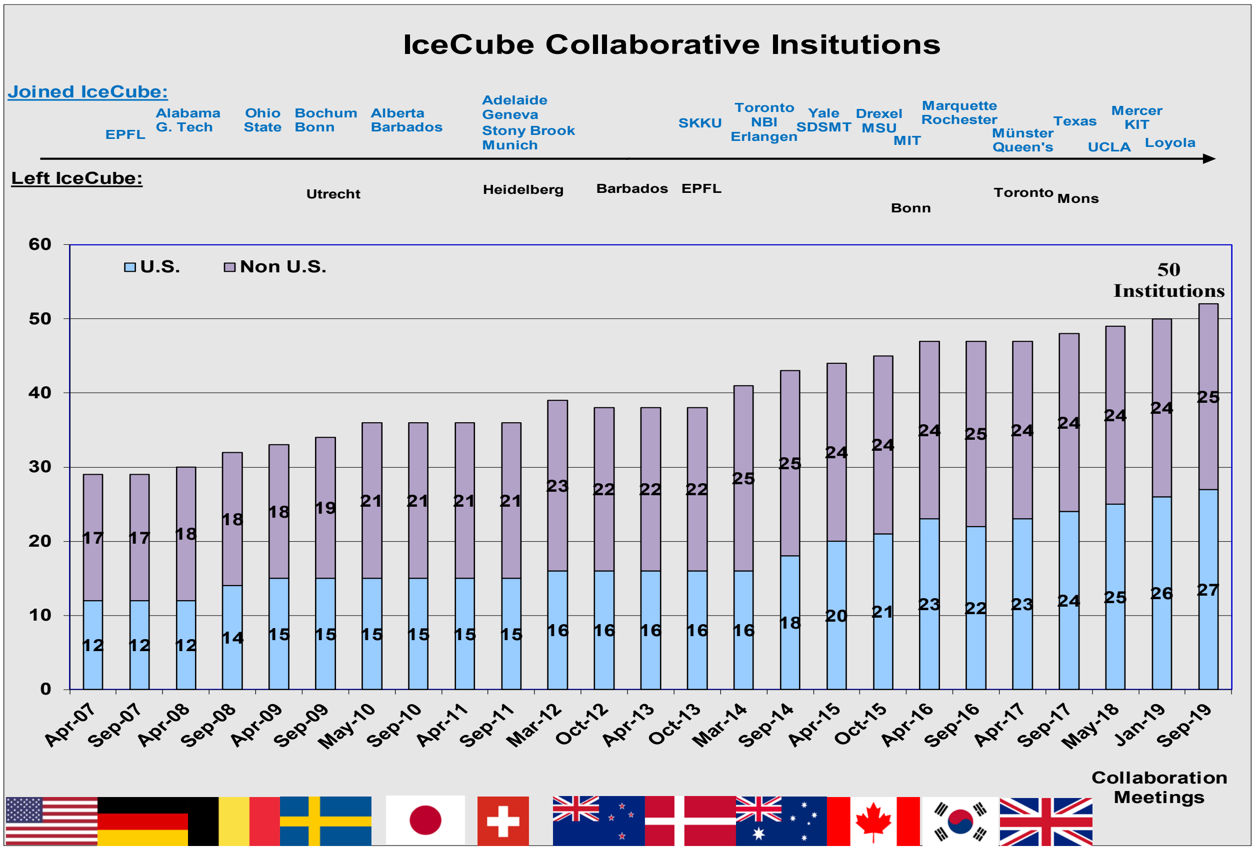

Changes to the IceCube Collaborating institutions in FY2019:

Following the September 2018 Fall collaboration meeting, the Mercer University with Dr. Frank McNally as the institutional lead, and Karlsruhe Institute of Technology (KIT) with Dr. Ralph Engel as the institutional lead were approved as full members of the IceCube Collaboration.

As of April 2019, the IceCube Collaboration consists of 50 institutions in 12 countries (26 U.S. and Canada, 20 Europe and 4 Asia Pacific). The IceCube collaborating institutions are listed in the IceCube Governance Document (included as Appendix 4 of this plan).

Appendix 3: IceCube Computing Infrastructure FY2016/2021 Upgrade Plan

The computing infrastructure systems include the South Pole System (SPS), the South Pole Test System (SPTS), the data warehouse, the UW data center, and networking. The following table summarizes the M&O computing infrastructure upgrade budget by subsystem.

Type of Funds	Sub System	YEAR1 (Direct $k)	YEAR2 (Direct $k)	YEAR3 (Direct $k)	YEAR4 (Direct $k)	YEAR5 (Direct $k)	YEARS 1-5 (Direct $k)
⊟ **US CF**	UW Data Center	$78k	$37k	$83k	$99k	$93k	**$390k**
	South Pole System	$8k	$14k	$21k	$43k	$38k	**$123k**
	South Pole Test System	$5k	$5k	$35k	$32k	$21k	**$98k**
	Data Warehouse	$16k	$64k	$0k	$0k	$80k	**$160k**
	DOMHubs	$5k	$25k	$65k	$36k	$35k	**$166k**
	Networking	$11k	$25k	$9k	$22k	$0k	**$67k**
US CF Total		**$123k**	**$170k**	**$213k**	**$232k**	**$266k**	**$1,004k**
⊟ **Non US CF**	UW Data Center	$457k	$390k	$218k	$282k	$275k	**$1,624k**
	South Pole System	$28k	$187k	$338k	$173k	$15k	**$742k**
	South Pole Test System	$162k	$0k	$0k	$149k	$0k	**$311k**
	Data Warehouse	$37k	$224k	$195k	$87k	$94k	**$638k**
	Networking	$109k	$0k	$0k	$0k	$0k	**$109k**
Non US CF Total		**$794k**	**$801k**	**$751k**	**$692k**	**$385k**	**$3,423k**
Grand Total		**$917k**	**$971k**	**$964k**	**$924k**	**$651k**	**$4,426k**

IceCube M&O Computing Infrastructure Upgrades Budget by Subsystem ($k)

The non-U.S. contribution to the Common Fund covers most of the capital equipment expenditures and their associated service agreements fees. The following table summarizes the M&O computing infrastructure upgrades budget by subsystem and cost category.

Type of Funds	Sub System	Capital Equipment	Materials & Supplies	Services	YEARS 1-5 (Direct$)
⊟ **US CF**	UW Data Center	$175k	$0k	$215k	**$390k**
	South Pole System	$25k	$98k	$0k	**$123k**
	South Pole Test System	$53k	$9k	$36k	**$98k**
	Data Warehouse	$160k	$0k	$0k	**$160k**
	DOMHubs	$125k	$25k	$16k	**$166k**
	Networking	$0k	$0k	$67k	**$67k**
US CF Total		**$538k**	**$132k**	**$334k**	**$1,004k**
⊟ **Non US CF**	UW Data Center	$1,454k	$0k	$170k	**$1,624k**
	South Pole System	$742k	$0k	$0k	**$742k**
	South Pole Test System	$311k	$0k	$0k	**$311k**
	Data Warehouse	$519k	$0k	$119k	**$638k**
	Networking	$109k	$0k	$0k	**$109k**
Non US CF Total		**$3,134k**	**$0k**	**$288k**	**$3,423k**
Grand Total		**$3,672k**	**$132k**	**$623k**	**$4,426k**

IceCube M&O Computing Infrastructure Upgrades Budget by Subsystem and Cost Category ($k)

Following is a summary of the proposed computing infrastructure upgrades per each major subsystem, including annual planned quantities and total U.S. and non-U.S. direct cost.

Unit definition:

- **CPU:** HEP-SPEC06, a CPU benchmark used extensively in other high-energy physics experiments
- **GPU:** “normalized gpu units” which is the computing power of an Nvidia GeForce GTX680 GPU
- **Disk/Tape:** Terabytes, defined as 10^{12} bytes

Appendix 4: IceCube Collaboration Governance Document

Revision 8.8, January 2019