

# **IceCube Maintenance and Operations Final Project Report**

**October 1, 2010 – March 31, 2016**

**Submittal Date: June 21, 2016**

---

University of Wisconsin–Madison

This report is submitted in accordance with the reporting requirements set forth in the IceCube Maintenance and Operations Cooperative Agreement, ANT-0937462.

## **Foreword**

This Final Project Report is submitted as required by the NSF Cooperative Agreement ANT-0937462. This report focuses on the last 6 months of the award covering the period of October 1, 2015 through March 31, 2016 but includes a summary of major highlights of the entire 66-month M&O program covering the period of October 1, 2010 through March 31, 2016. The status information provided in the report covers actual common fund contributions received through March 31, 2016 and the full 86-string IceCube detector (IC86) performance through March 1, 2016.

## Table of Contents

<b>Foreword</b>	<b>2</b>
<b>Section I – Financial/Administrative Performance</b>	<b>4</b>
<b>Section II – Maintenance and Operations Status and Performance</b>	<b>6</b>
<b>Detector Operations and Maintenance</b>	<b>6</b>
<b>Computing and Data Management</b>	<b>14</b>
<b>Data Release</b>	<b>21</b>
<b>Program Management</b>	<b>23</b>
<b>Section III – Project Governance and Upcoming Events</b>	<b>25</b>

## Section I – Financial/Administrative Performance

The University of Wisconsin–Madison is maintaining three separate accounts with supporting charge numbers for collecting IceCube M&O funding and reporting related costs: 1) NSF M&O Core account, 2) U.S. Common Fund account, and 3) Non-U.S. Common Fund account.

A total amount of \$3,450,000 was released to UW–Madison to cover the costs of maintenance and operations during the first half of FY2016: \$498,225 was directed to the U.S. Common Fund account based on the total number of U.S. Ph.D. authors, and the remaining \$2,951,775 was directed to the IceCube M&O Core account (Figure 1).

<b>FY2016 (Oct 1, 2015 – March 31, 2016)</b>	<b>Funds Awarded to UW</b>
IceCube M&O Core account	\$2,951,775
U.S. Common Fund account	\$498,225
<b>TOTAL NSF Funds</b>	<b>\$3,450,000</b>

**Figure 1: NSF IceCube M&O Funds - FY2016**

Of the IceCube M&O FY2016 Core funds, \$567,554 were committed to seven U.S. subawardee institutions. The institutions submit invoices to receive reimbursement against their actual IceCube M&O costs. Figure 2 summarizes M&O responsibilities and total FY2016 funds for the seven subawardee institutions.

<b>Institution</b>	<b>Major Responsibilities</b>	<b>Funds</b>
Lawrence Berkeley National Laboratory	Data Acquisition maintenance, computing infrastructure	\$40,538
Pennsylvania State University	Data acquisition firmware support, simulation production	\$29,357
University of California at Berkeley	Detector calibration, monitoring coordination	\$110,086
University of Delaware, Bartol Institute	IceTop calibration, monitoring and maintenance	\$93,797
University of Maryland at College Park	IceTray software framework, online filter, simulation software	\$259,453
University of Alabama at Tuscaloosa	Detector calibration, reconstruction and analysis tools	\$22,430
Michigan State University	Simulation software, simulation production	\$11,893
<b>Total</b>		<b>\$567,554</b>

**Figure 2: IceCube M&O Subawardee Institutions – FY2016 Major Responsibilities and Funding**

### *IceCube NSF M&O Award Budget, Actual Cost and Forecast*

The current IceCube NSF M&O 5-year award was established at the beginning of Federal Fiscal Year 2011, on October 1, 2010, and was extended by 6 months by NSF, through March 31, 2016. The following table presents the financial status six months into FY2016, and shows an estimated balance at the end of the entire 66-month award.

Total awarded funds to the University of Wisconsin (UW) for supporting IceCube M&O from the beginning of FY2011 through the end of the award are \$38,144K. Total actual cost as of March 31, 2016 is \$37,875K; open commitments against purchase orders and subaward agreements are \$268K. The current balance as of March 31, 2016 is \$268K. With a projection of \$217K for the remaining expenses during the 90-day close-out period, the estimated unspent funds at the end of the award are \$51K or 1.5% of the FY2016 budget (Figure 3), which will be used for part of the recent capital equipment purchase order of CPUs and a disk server.

(a)	(b)	(c)	(d)= a - b - c	(e)	(f) = d - e
<b>YEARS 1-6 Budget</b> Oct.'10-Mar'16	<b>Actual Cost To Date</b> through Mar. 31, 2016	<b>Open Commitments</b> on Mar. 31, 2016	<b>Current Balance</b> on Mar. 31, 2016	<b>Remaining Estimated Expenses</b> through Mar. 31, 2016	<b>End of award Estimated Balance</b> on Mar 31, 2016
\$38,144K	\$37,875K	\$268K	\$268K	\$217K	\$51K

**Figure 3: IceCube NSF M&O Award Budget, Actual Cost and Forecast**

***IceCube M&O Common Fund Contributions***

The IceCube M&O Common Fund was established 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 contributed to the Common Fund, based on the total number of the institution’s Ph.D. authors, at the established rate of \$13,650 per Ph.D. author. The Collaboration updates the Ph.D. author count twice a year before each collaboration meeting in conjunction with the update to the IceCube Memorandum of Understanding for M&O.

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 and are listed in the Maintenance & Operations Plan.

Figure 4 summarizes the planned and actual Common Fund contributions for the period of April 1, 2015–March 31, 2016, based on v18.0 of the IceCube Institutional Memorandum of Understanding, from April 2015. Actual Common Fund 2015-2016 contributions are \$62k less than planned. The final non-U.S. contributions are underway, and it is anticipated that most of the planned contributions will be fulfilled.

	<b>Ph.D. Authors</b>	<b>Planned Contribution</b>	<b>Actual Received</b>
<b>Total Common Funds</b>	<b>137</b>	<b>\$1,870,050</b>	<b>\$1,766,052</b>
U.S. Contribution	73	\$996,450	\$996,450
Non-U.S. Contribution	64	\$873,600	\$810,790

**Figure 4: Planned and Actual CF Contributions for the period of April 1, 2015–March 31, 2016**

## Section II – Maintenance and Operations Status and Performance

### *Detector Operations and Maintenance*

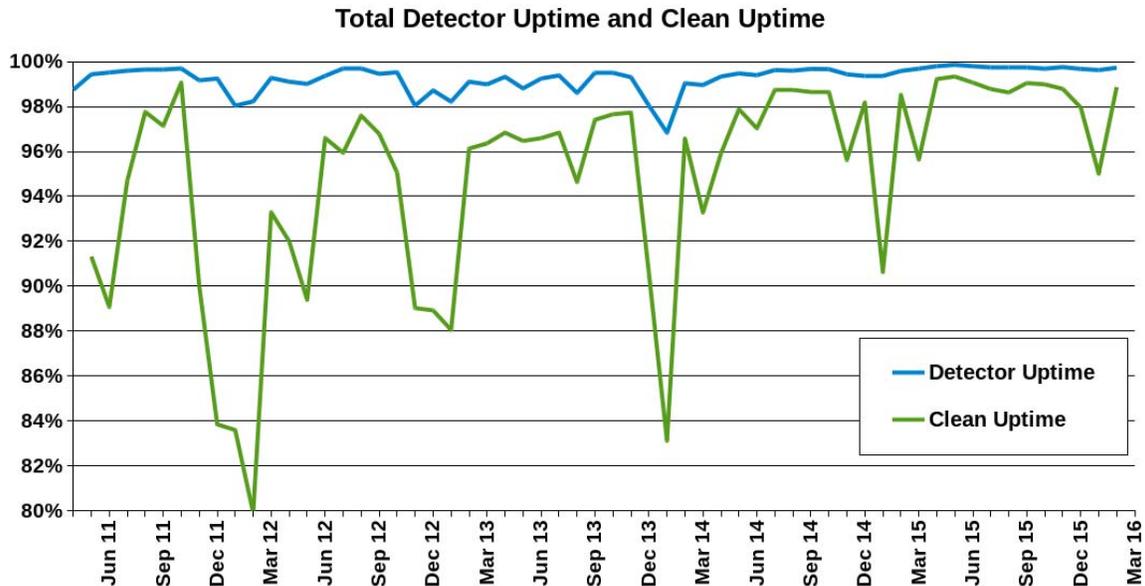


Figure 1: Total IceCube Detector Uptime and Clean Uptime

### *Major Highlights of the 5-Year M&O Program*

The operation of IceCube during the initial M&O period has focused both on improving detector stability and increasing the physics potential of the experiment, via continued development in the data acquisition systems, event filtering, detector calibration, and experiment control and monitoring. The average uptime since the start of full-detector operations has been 99.3% (see Figure 1).

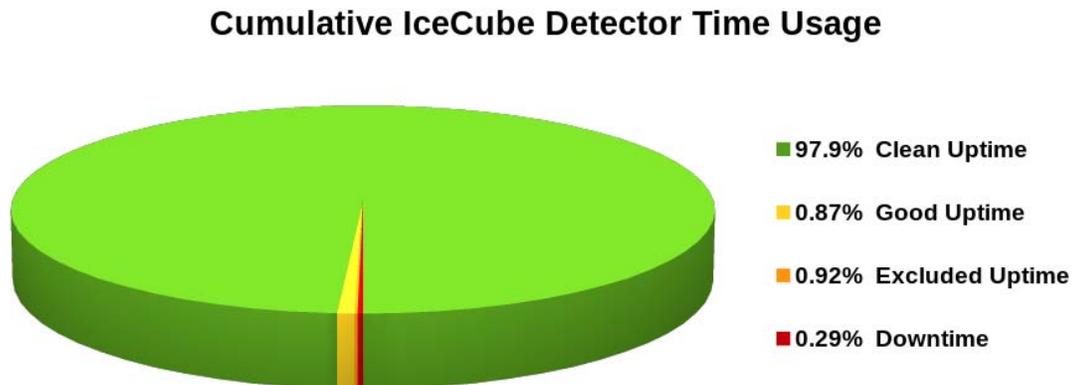
Some of the key improvements and upgrades during this M&O period include:

- implementation of the DAQ “hitspooling” feature that allows the capture of short periods (up to one hour) of *all* photons detected by the DOMs, providing opportunities for a number of new physics analyses as well as improved detector characterization;
- implementation of the supernova DAQ “muon subtraction” feature that improves the sensitivity of supernova alerts to the wider community;
- tracking the good portion of failed runs in IceCube Live, allowing us to achieve our longstanding goal of 95% clean detector uptime;
- upgrade of the single-board computers (SBCs) in all DOMHubs, providing higher performance with lower power consumption;

- rollout of the superDST event compression format, allowing more events to be transferred over satellite and providing an efficient data archive format;
- discovery of an azimuthal anisotropy in the scattering properties of the ice, allowing more accurate directional and energy reconstruction of neutrino events; and
- expansion of the real-time alert program via both satellite connectivity improvements and new online filters to expand IceCube’s multi-messenger search for the sources of cosmic neutrinos.

While the DOMs are frozen into the glacial ice, our ability to continuously improve the online software systems and ICL hardware have allowed us to continue to expand the capabilities of the experiment.

**Detector Performance** – During the period from September 1, 2015, to March 1, 2016, the full 86-string detector configuration (IC86) operated for 97.91% of the time. Continued operational and software improvements have contributed to increased detector stability and resulted in an unprecedented detector uptime of 99.71%. Figure 2 shows the cumulative detector time usage over the reporting period. The good uptime was 0.87% of the time and includes partial detector (not all 86 strings in operation) analysis-ready data. Excluded uptime includes maintenance, commissioning, and verification data and required 0.92% of detector time. The unexpected detector downtime was limited to 0.29%.



**Figure 2: Cumulative IceCube Detector Time Usage, September 1, 2015 – March 1, 2016**

The current feature to track portions of failed runs as good, allows the recovery of data from all but the last few minutes of runs that fail, and the recent implementation of continuous data-taking, have improved detector stability and decreased detector downtime. These features increased the average “clean uptime” for this reporting period to 97.91% of full-detector, analysis-ready data, as shown in Figure 1. We are now regularly exceeding our target clean uptime of 95%.

About 0.2% of the loss in clean uptime is due to the failed portions of runs that are not usable for analysis. There is around 1% of clean uptime loss due to runs not using the full-detector configuration. This occurs when certain components of the detector are excluded from the run configuration during required repairs and maintenance. This data is still good for analyses that have less strict requirements on the active detector volume. There is approximately a 1% loss of clean uptime due to maintenance, commissioning, and verification runs, and short runs that are less than 10 minutes in duration. The experiment control system and DAQ have recently implemented 32-hour periods of continuous data-taking. This new feature has eliminated approximately 90–120 seconds of downtime between each run transition, gaining roughly 0.5% of uptime.

Upcoming improvements include a restructuring of the DAQ functionality so that dropped DOMs can be recovered “in-situ”. This will allow for continuous data-taking to run for longer than the current 32-hour periods. The combined effect will reduce the partial detector configuration time and downtime and increase the clean uptime by at least 0.5%.

The IceCube Run Monitoring system, I3Moni, provides a comprehensive set of tools for assessing and reporting data quality. 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 original monolithic monitoring system processes data from various SPS subsystems, packages them in files for transfer to the Northern Hemisphere, and reprocesses them in the north for display on the monitoring web pages. In a new monitoring system under development (I3Moni 2.0), all detector subsystems report their data directly to IceCube Live. Major advantages of this new approach include: higher quality of the monitoring alerts; simplicity and easier maintenance; flexibility, modularity, and scalability; faster data presentation to the end user; and a significant improvement in the overall longevity of the system implementation over the lifetime of the experiment.

The I3Moni 2.0 infrastructure for collecting the monitoring data is in place at SPS, and monitoring quantities are now being collected from the five major subsystems: data acquisition (DAQ), supernova DAQ, processing and filtering (PnF), calibration and verification (CnV) and the DOM monitoring (HubMoni). The previous I3Moni 2.0 workshop, held at the University of Wisconsin in September 2015, focused on the few remaining elements required for the beta release that include the integration and collection of subsystem quantities as well as the development and integration of sophisticated quality control systems. The workshop also focused on forging out a plan and timeline for the I3Moni 2.0 public release. Since then, 24 issues and requested features have been resolved. The primary improvements include a fully integrated and comprehensive web display showing all quantities and quality control test results, and a subsystem for storing the data representation and quality control test results for quick web page loading. Since December 2015, the I3Moni 2.0 beta release has been active, and the public release candidate is scheduled for mid-2016.

Development of IceCube Live, the experiment control and monitoring system, is still quite active. A few weeks after the previous release's deployment (2.8.3), IceCube Live was using the new Iridium RUDICS link to transfer moderate and high priority monitoring information, with significantly better bandwidth and latency than the older system. This reporting period has seen

two major releases with the following highlighted features:

- Live v2.9.1 (December 2015): 41 separate issues and feature requests have been resolved. Improvements were made to the command-line toolset used to control the data-taking process, and IceCube Live is now able to use the new centralized messaging system I3MS. This release also marks the start of the Moni2.0 beta phase, allowing a wider audience to access and test the new run monitoring tools.
- Live v2.9.2 (March 2016): 39 separate issues and feature requests have been resolved. HitSpool requests can now be initiated and monitored from the Northern IceCube Live website. Many Moni2.0 improvements were implemented in preparation for the upcoming public release. As of this release, IceCube Live instances use dedicated MongoDB servers to store the heavy load of new monitoring quantities, to satisfy our growing performance requirements and enhance the user experience.

Features planned for the next few releases include: continuing development of Moni2.0 into its public release; centralizing DOM problems and developing an interface to access that information; and creating or improving dedicated monitoring pages for the JADE, OFU, GFU, SNDAQ subsystems. The uptime for the I3Live experiment control system during the reporting period was above 99.999%.

The IceCube Data Acquisition System (DAQ) has reached a stable state, and consequently the frequency of software releases has slowed to the rate of 3–4 per year. Nevertheless, the DAQ group continues to develop new features and patch bugs. During the reporting period of October 2015–March 2016, the following accomplishments are noted:

- Delivery of five updates to the DAQ v4.9 release throughout the Pole summer season, which included minor bug fixes and code modifications needed to incorporate the new scintillator and IceACT sources.
- Delivery of the next DAQ release in late March or early April, 2016 which will include improved code to recover DOMs which have stopped producing data; further enhancement of the HitSpool caching system; and more work on separating the hub-based data processing from the run-based DAQ system to pare detector downtime to the absolute minimum.

The supernova data acquisition system (SNDAQ) found that 99.61% of the available data from July 31st, 2015 through March 2nd, 2016 met the minimum analysis criteria for run duration and data quality for sending triggers. An additional 0.06% of the data is available in short physics runs with less than 10-minute duration. While forming a trigger is not possible in these runs, the data are available for reconstructing a supernova signal.

A new SNDAQ release (2015-12-15) improved the leap second handling; SNDAQ alert thresholds were adjusted; and the SNDAQ alert transmission was changed from email to ZeroMQ. In the event of a significant supernova alarm, the HitSpool cached raw data are automatically retrieved and sent to the North for analysis. Efforts towards running an automated

analysis were finalized in December 2015. Atmospheric muon-corrected lightcurves and significances are calculated from HitSpool data within 48 hours after the alarm. The SNDAQ build system was completely revised and simplified and is now in line with other IceCube software. Efforts to include a data-driven trigger that is independent of an assumed signal shape are under way.

Due to a December power shutdown at the University of Mainz and a subsequent change of operating system, the server that manages the alarm handling for SNEWS was moved, and the related software was updated. This serves as an interim solution before the alarm handling and physics related monitoring pages are moved to a central server at UW–Madison. This move is well on its way, *e.g.* new monitoring webpages summarizing higher level HitSpool analyses are now available. In addition, the working group is beginning to develop criteria to integrate significant SNDAQ alerts with other IceCube data using tools such as the Astrophysical Multimessenger Observatory Network (AMON). The SNDAQ monitoring is now working in I3Live Moni2.0 and will be expanded to provide the history of SNEWS alerts as well as supernova data acquisition runs in the future; efforts to include the HitSpool monitoring in I3Live are under way.

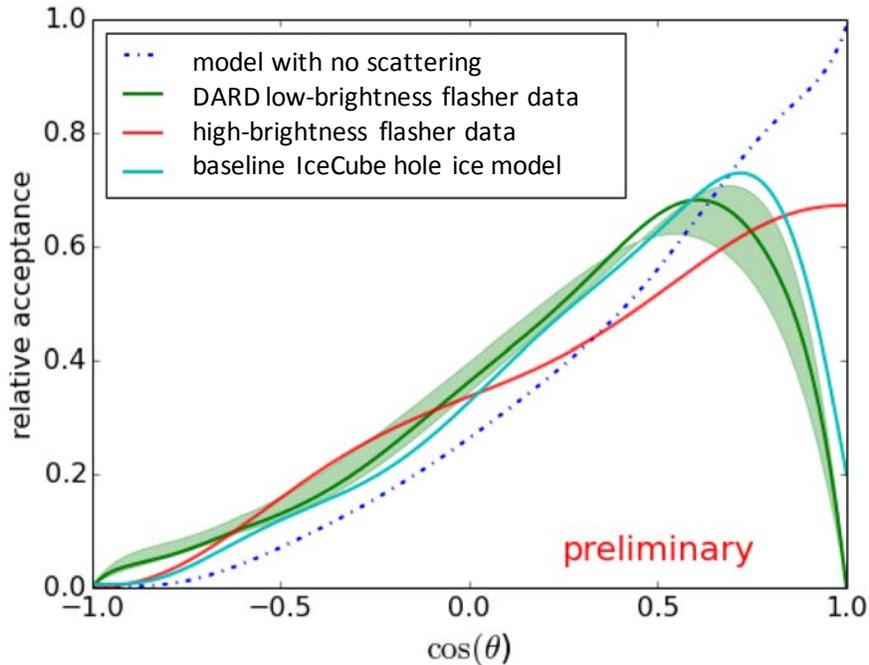
The online filtering system (“PnF”) performs real-time reconstruction and selection of events collected by the data acquisition system and sends them for transmission north via the data movement system. Primary activities for this system during this report period focused on continued expansion of data quality checks in the Moni2.0 system, training of new winterovers, and support for special operations during the austral summer season (flashers and scintillator deployment.) New software release versions were tested and released to support realtime alert system testing and development, additional Moni2.0 data quality checks, additional alerts for winterover monitoring of the system and retirement of the CnV system (V15-11-00 and V16-01-00). Current development work is targeted to address system load issues seen during tests of the full Moni2.0 system and development work for the upcoming 2016 run configuration transition.

A weekly calibration call keeps collaborators abreast of issues in both in-ice and offline DOM calibration. Keiichi Mase (Chiba University) joined the calibration group as co-convener with Dawn Williams (Univ. of Alabama) and organizes a second call for the calibration work in the Asia/Pacific IceCube institutions.

Long-term detector stability has been examined with the in-ice calibration laser called the standard candle, LED flashers and muons. The detector response has been consistent within 1% since detector completion. The DOM calibration procedure has also been updated to include the new surface detectors installed at the South Pole in the 2015–2016 pole season, and the software for analyzing calibration data has been updated to be completely Python-based and compatible with the updated monitoring system.

The effect of refrozen “hole ice” on the angular response of the DOM continues to be studied with LED flashers. A large set of flasher data with the tilted LEDs was collected at the South Pole in January 2016, which will assist in these studies, and serve as a cross-verification of data from the horizontal LEDs which was used to develop IceCube’s current ice model. A new run

with the camera system at the bottom of string 80 was collected in order to look at the bubbles in the hole ice. PMT collection efficiency continues to be studied in the laboratory as well.



**Figure 2: Hole ice effect on DOM relative acceptance as a function of incident photon zenith angle from various studies including: low-brightness flashers in ice (DARD); a multi-parameter fit to bright flasher data; and the simulation-based default hole ice model.**

The IceTop group has continued work on the characterization of average IceTop DOM collected charge for short- and long-term monitoring. However, characterization of the charge distributions is complex not only due to the gradual snow buildup on tanks, but also due to the seasonal modulations. Periodic variations due to temperature effects are smoothed out when divided by the bi-weekly Vertical-Equivalent Muon (VEM) calibrations, but a +/-7% fluctuation for high-gain DOMs and +/-25% fluctuations for low-gain DOMs remain. An upward trend in high-gain DOM mean charge is being investigated in detail. Since the VEM calibration does not show such a trend, we suspect this is an effect of the snow accumulation (see Fig. 4). The average snow depth across the array has now reached 1.67m.

A recent decision by the NSF and the support contractor has been taken to stop any further snow management efforts. With this in mind, we are investigating experimental techniques to restore the full operational efficiency of the IceCube surface component using plastic scintillator panels on the snow surface above the buried IceTop tanks. Four of these panels were built, calibrated, and cold tested at UW–Madison.

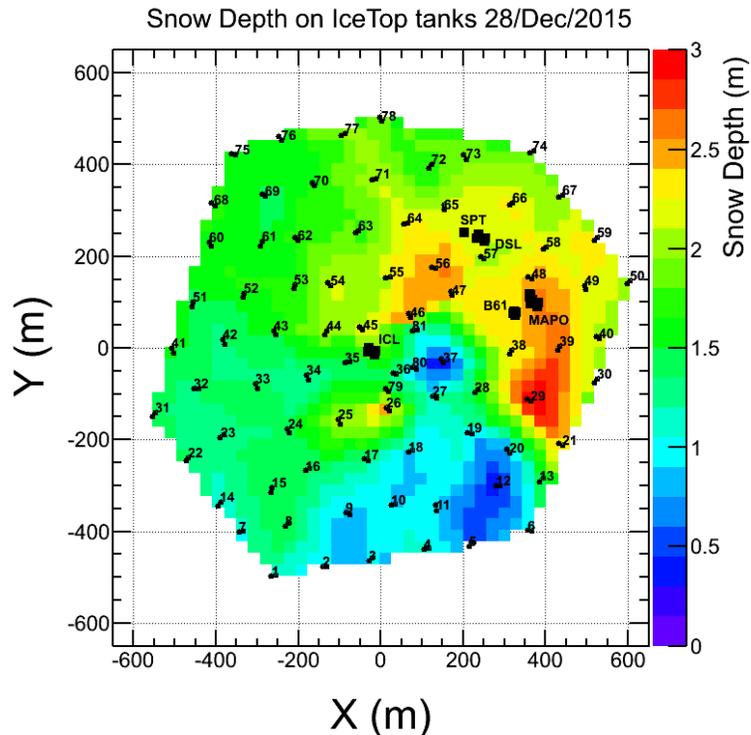
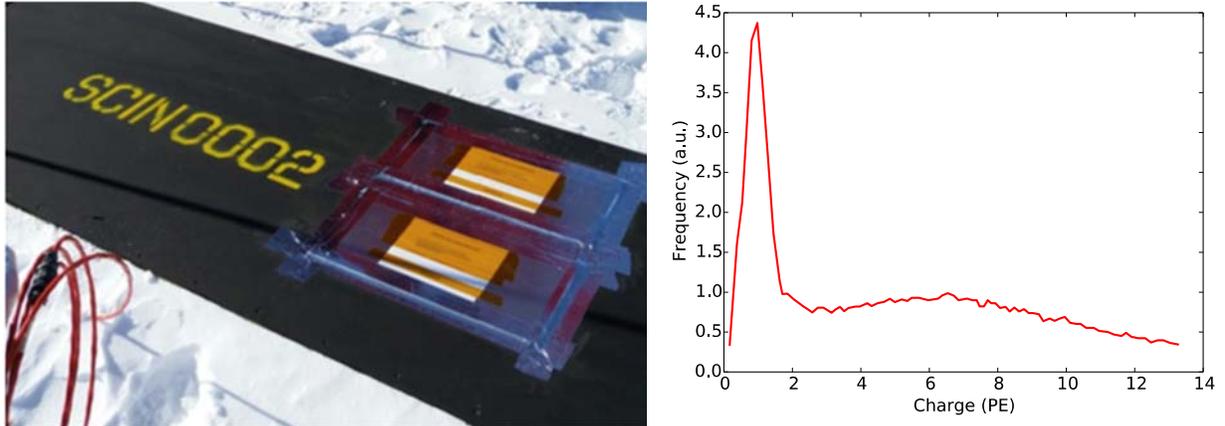


Figure 4. Snow accumulation over IceTop tanks as of December 2015.

These four scintillators have now been deployed during the 2015–16 pole season (see Figure 5) at IceTop stations 12 and 62. The initial deployment used spare tank freeze control cables left in place after deployment in order to read out the scintillator data and connect into the IceCube DAQ, with minimal changes necessary. The scintillators do not participate in forming IceCube triggers, but are read out and included into the data stream when other triggers are formed either by IceTop or the in-ice detector. Additionally, a prototype air Cherenkov telescope (IceACT) was installed on the IceCube Lab. During the polar night, IceACT can be used to cross-calibrate IceTop by detecting the Cherenkov emission from cosmic ray air showers; it may also prove useful as a supplementary veto technique. While engineering studies have begun with this prototype instrument, full operation requires dark sky conditions. Results of 2016 winter data taking will be presented in the next M&O report.

Current surface array work is focused on analysis of the data being collected by the scintillator prototypes. Development of a new version of the scintillators is underway, which will include a new digitization and readout system, a possibly different photodetection technology (SiPMs), and a streamlined, lighter housing to ease deployment.



**Figure 5: Left: Scintillator deployed over IceTop station 62. Right: Charge spectrum of scintillator, showing single photoelectron peak at 1PE and broad cosmic-ray muon peak at ~6PE.**

**IC86 Physics Runs** – The fifth season of the 86-string physics run, IC86–2015, began on May 18, 2015. Detector settings were updated using the latest yearly DOM calibrations from March 2015, and new precision DOM gain calibrations were integrated into the online processing system. DAQ trigger settings did not change from IC86–2014. Filter changes include new starting-event filters as well as a real-time neutrino stream, both of which will be used as follow-up multi-messenger alerts to other observatories. The removal of a minimum-bias component to the full-sky-starting-track filter led to a reduction in average satellite bandwidth usage. Preparations for the IC86–2016 physics run are now underway, with the yearly full-detector calibration completed in March 2016.

The last DOM failures (2 DOMs) occurred during a power outage on May 22, 2013. No DOMs have failed during this reporting period. The total number of active DOMs remains 5404 (98.5% of deployed DOMs).

**TFT Board** – The TFT board is in charge of adjudicating SPS resources according to scientific need, as well as assigning CPU and storage resources at UW for mass offline data processing (a.k.a. Level 2). TFT management of the offline processing has resulted in a latency of only 2–4 weeks after data-taking. Working groups within IceCube submit proposals requesting data processing, satellite bandwidth and data storage, and the use of various IceCube triggers for IC86–2015. Sophisticated online filtering data selection techniques are used on SPS to preserve bandwidth for other science objectives. Over the past three years, a new data compression algorithm (SuperDST) has allowed IceCube to send a larger fraction of the triggered events over TDRSS than in previous seasons. The additional data enhances the science of IceCube in the search for neutrino sources from the Southern sky, including the Galactic Center. Furthermore, this compressed stream serves as the archival format for IceCube raw data.

Starting with IC86–2015, we implemented changes to the methodology for producing online quasi-real-time alerts. Neutrino candidate events at a rate of 3 mHz are now sent via Iridium satellite, so that neutrino coincident multiplets (and thus candidates for astrophysical transient sources) can be rapidly calculated and distributed in the northern hemisphere. This change will enable significant flexibility in the type of fast alerts produced by IceCube.

Since the IC86–2015 run start in May, the average TDRSS daily transfer rate is approximately 70 GB/day, reduced from 90 GB/day in IC86-2014 by optimization of online filter settings. IceCube is a heavy user of the available bandwidth, and we will continue to moderate our usage without compromising the physics data.

Several new triggers and filters are under development for the IC86–2016 physics run. A new IceTop trigger using the infill tanks in the core of the array will target low-energy cosmic ray air showers, and a modified calibration trigger will collect minimum bias data from the newly deployed scintillator panels. A new monopole filter will collect an event sample to

***Operational Communications*** – Communication with the IceCube winterovers, timely delivery of detector monitoring information, and login access to SPS are critical to IceCube’s high-uptime operations. Several technologies are used for this purpose, including interactive chat, ssh/scp, the IMCS e-mail system, and IceCube’s own Iridium modem(s).

With the TDRS-F5 satellite retirement, the total high-bandwidth satellite coverage was reduced by approximately 6 hours per day, because TDRS-F6 now overlaps with the GOES pass. This significantly reduced the amount of e-mail that could be transmitted over the satellites, and more e-mail traffic had to be moved over the IMCS (Iridium) link. This, combined with issues with a Microsoft Exchange upgrade, led the contractor to eliminate 24/7 e-mail delivery service except for operational e-mail and monitoring data. Latency of the system still proved unpredictable, and communication with personnel at the pole this summer season was challenging due to these restrictions.

We have now developed our own Iridium RUDICS-based transport software (IceCube Messaging System, or I3MS) and have moved IceCube’s monitoring data to our own Iridium modems as of the 2015–16 austral summer season. The technical details of this plan have been reviewed and approved by NSF and the contractor. The system is scalable, in that more modems can be added the future to increase capacity.

***Personnel*** – No changes.

## ***Computing and Data Management***

### ***Major Highlights of the 5-Year M&O Program***

- **New system for handling IceCube data archiving and transfer.** The specialized software that manages the data lifecycle of the experiment, including archive at the South Pole, transfer to the North and storage at the UW-Madison data warehouse, has been rewritten in order to make it more efficient, easy to maintain and operate. One of the main features of the new software is that it manages the data lifecycle end-to-end within one single software system. This allows us to do a more consistent cataloguing of the IceCube data products and will ease the future development of a central data catalog that will streamline data discovery and metadata queries. Together with this major software rewrite, we have upgraded the hardware systems that are in charge of data archiving at the South Pole. These systems have moved from using tapes into using disks as archive

media. These software and hardware improvements have significantly increased the stability and reliability of the data archive system at the South Pole.

- **MoU with NERSC/LBNL for IceCube long-term data archive services.** Curating a multi-petabyte tape archive for many years is an activity that entails a number of very specialized operational procedures. One example of these would be the periodic tape media migration that must be done to avoid losing data when tapes age or when technologies become obsolete. Due to this, long term data archiving and curation is a service that can have a very high cost if not done at the proper scale. In 2014 we started exploring the possibility of partnering with a large storage facility to implement the IceCube long-term data archive. The main motivation was aiming for a better service at lower cost, since large facilities that routinely manage data at the level of 100 petabytes benefit of economies of scale that ultimately make the process more efficient and economical. A Memorandum of Understanding between UW-Madison and NERSC/LBNL was signed in December 2015 for NERSC to provide long-term data archive services for IceCube in the next five years.
- **Major version upgrade of Lustre software in the UW-Madison data warehouse.** Lustre is a distributed filesystem software used for large-scale cluster computing. The IceCube data warehouse relies entirely on this software for providing scalable and high performance data services. Until the start of 2014, all of the production filesystems were using Lustre version 1.8, released in May 2009. In 2014 we started a massive migration of the data from the old filesystems to brand new filesystems configured with the last stable Lustre version 2.5. This migration implied the transfer of close to 3 PB of data, and was performed transparently for users as data was moved in the background. The migration was successfully completed on June 2015. Since then, all the filesystems in the data warehouse are running an up to date Lustre version that provides richer functionality and improved resiliency.
- **GPU cluster for simulation production.** In October 2012, the IceCube Collaboration approved the use of direct photon propagation for the mass production of simulation data. The direct propagation method requires the use of GPU (graphical processing units) hardware to deliver adequate performance. The first IceCube GPU cluster was installed in the data center at UW-Madison in 2012, containing 48 NVidia Tesla M2070 GPU cards. Since then, the cluster has been successively expanded following the increasing demand. It currently consists of 400 GPU cards that provide a total compute capacity more than 20 times larger than the original 2012 cluster and substantially larger than the GPU capacity in many of the top-level research supercomputer systems in the US.
- **Expansion of the IceCube distributed computing infrastructure.** The amount of simulated data that IceCube analyses require demands of very large amounts of computing resources. The strategy to build such large computing infrastructure relies strongly on being able to use distributed resources in an efficient manner. Many of the IceCube collaboration sites are able to provide access to local computing resources. The available infrastructure is therefore potentially very large, provided we can handle heterogeneous distributed resources effectively. In the last years we have strongly

focused in adopting the standards and methodologies that allow us to leverage the power of large distributed computing systems. By actively participating in the Open Science Grid (OSG) project and interfacing our simulation production systems to the Grid services provided by OSG, we have been able to steadily expand our distributed infrastructure in the last years, including many sites in the US, Europe and Canada.

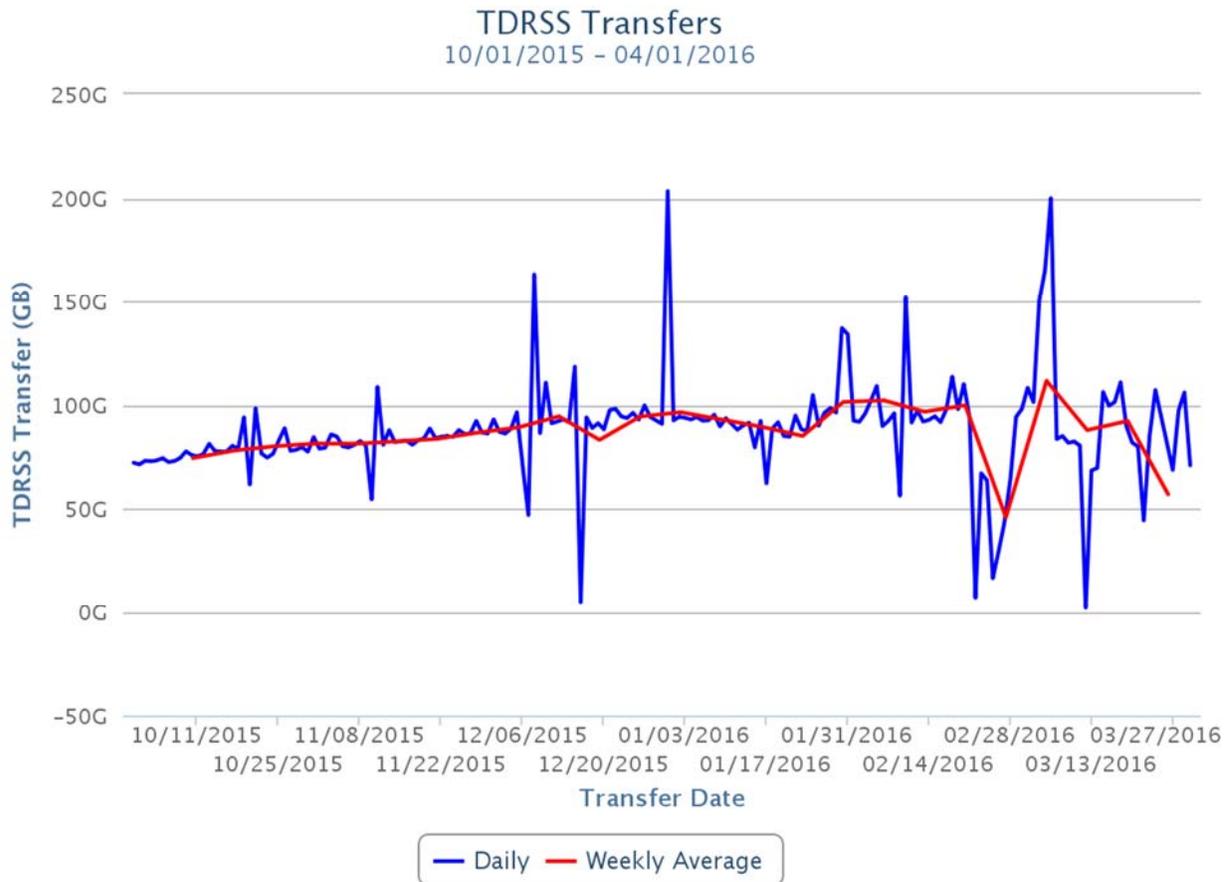
**South Pole System** – The main activities carried out during the 2015-16 South Pole season were the upgrade of the network firewall system, the deployment in production of the new Iridium based RUDICS infrastructure and the migration of the SPS servers configuration into a configuration management system.

The CISCO ASA firewalls operating at SPS until 2015 were deployed close to 10 years ago, so they needed an upgrade. During the 2015-16 season, a pair of Dell SonicWall NSA 3600 firewall units were deployed to replace them. One of the main improvements with the new firewalls is the enhanced failover implementation, which allows operators to apply firmware upgrades with no connection drops.

After the extensive development and testing for the new RUDICS system that took place during 2015, the final system was deployed in the ICL during the 2015-16 season. Two modems were deployed in the server room and connected to the I3MS application server, located in a dedicated network zone.

Another major activity that took place during the last Pole season was the deployment of a new configuration management system to handle more efficiently the provision and configuration changes of SPS servers. The chosen software is Puppet (see <https://puppetlabs.com>) an open source tool that is widely used in commercial and academic IT environments to manage configuration and automate provisioning of large IT installations. One of the processes that puppet facilitates is keeping a strict version control of configuration changes without the need of reinstalling servers.

**Data Movement** – Data movement has performed nominally over the past six months. Figure 6 shows the daily satellite transfer rate and weekly average satellite transfer rate in GB/day through March 2016. The IC86 filtered physics data are responsible for 95% of the bandwidth usage. One can also notice in the Figure that the daily transfer rates have been quite unstable during February and March.



Highcharts.com

**Figure 6: TDRSS Data Transfer Rates, October 1, 2015–March 31, 2016.** The daily transferred volumes are shown in blue and, superimposed in red, the weekly average daily rates are also displayed.

**Data Archive** – The IceCube raw data are archived to two copies on independent hard disks. During the reporting period (October 2015 to March 2016) a total of 200 TB of data were archived to disk, averaging 1.1 TB/day. A total of 14.4 TB of data were sent over TDRSS, averaging 81.2 GB/day.

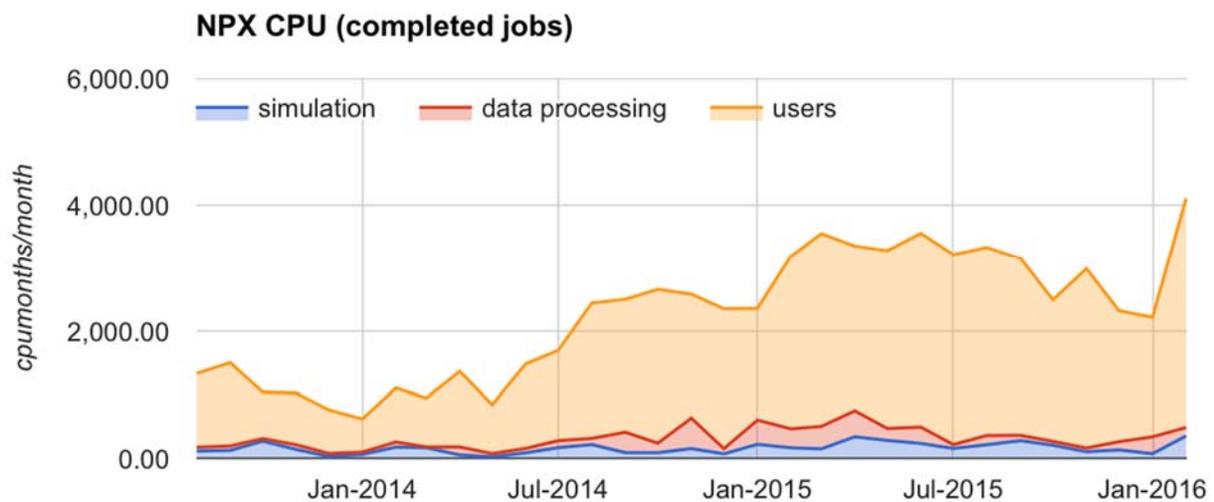
**Computing Infrastructure at UW-Madison** – IceCube computing and storage systems, both at the Pole and in the north, have performed well over the reporting period. The total amount of data stored on disk in the data warehouse is 3938 Terabytes<sup>1</sup> (TB): 1000 TB for experimental data, 2727 TB for simulation and analysis and 211 TB for user data.

The IceCube computing cluster at UW-Madison has continued to deliver reliable data processing services. An expansion of the CPU cluster consisting of two SuperMicro TwinBlade chassis was purchased and deployed in February 2016. Each chassis contains 20 blade servers and each server has two E5-2680v3 processors, 128 GB RAM, 2TB disk and two 10Gbps network interfaces. After this upgrade, the cluster has a total of 7000 CPU cores (Hyperthreading enabled) with an average of 2 GB of memory per core.

<sup>1</sup> 1 Terabyte = 10<sup>12</sup> bytes

Figure 7 shows the CPU time delivered per month by the UW-Madison cluster since August 2013. The different colors represent the three main workloads that run in the cluster: data processing, user analysis and simulation production. Centralized data processing is the highest priority workload, to ensure data products are available as soon as possible. However, it amounts only to about 10% of the cluster capacity. Most of the cluster capacity is used for user analysis, and this is the main facility for this activity within the collaboration. An additional 10% of the capacity is used for simulation production tasks, but the bulk of this activity relies on resources that are external to the main IceCube facility at UW-Madison.

The capacity increases that can be seen around July 2014 and January 2015 come mainly from the fine-tune of default per job memory requirements, which allowed us to increase the overall utilization factor by doing a more efficient resource allocation per slot. The sharp increase around February 2016 corresponds to the deployment of the last hardware purchase.



**Figure 7: CPU time consumed by the jobs completing monthly since August 2013 in the IceCube CPU cluster at UW-Madison. The main contribution comes from user analysis jobs (in yellow). Data processing (in red) and simulation production (in blue) take about 10% of the resources each.**

The IceCube computing cluster also contains 60 GPU-enabled nodes, providing a total of 400 slots to run GPU jobs. Boosting the internal GPU computing capacity has been a high priority of the project since the Collaboration decided to use GPUs for the photon propagation part of the simulation chain back in 2012. Direct photon propagation was found to provide the precision required, and it happens to be very well suited for the GPU hardware, running about 100 times faster than in CPUs.

Figure 8 shows the GPU time delivered per month by the UW-Madison GPU cluster since January 2013. A steady increase in the overall GPU usage can be observed. The big jump in total usage during the first quarter of 2015 corresponds to the deployment of the last GPU expansion, consisting of 256 Nvidia GTX 980 GPU cards.

The GPU nodes are deployed as part of the main IceCube cluster at UW-Madison. This way,

users can access both resource types, CPUs and GPUs, through the same interface. This simplified access mechanism has had the positive outcome that more collaboration members are now making use of this valuable resource.

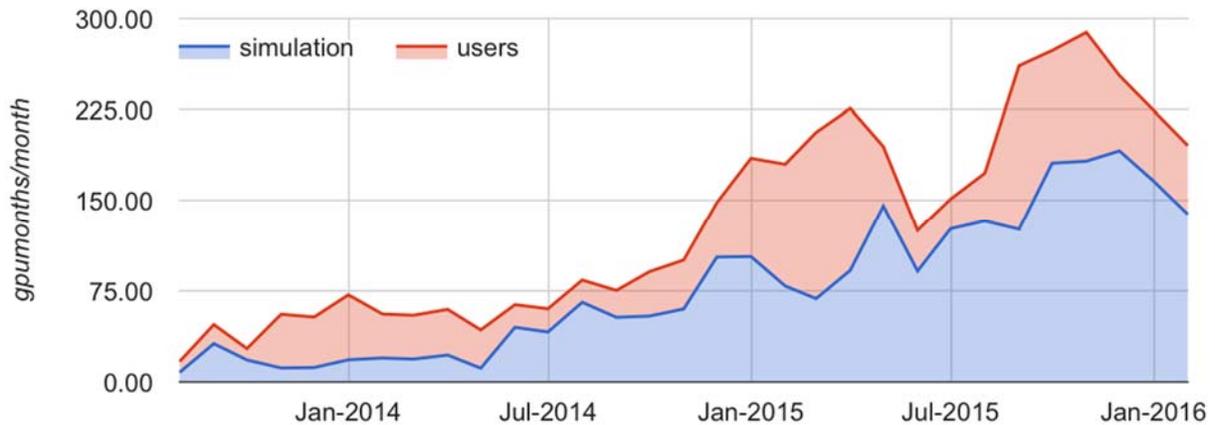


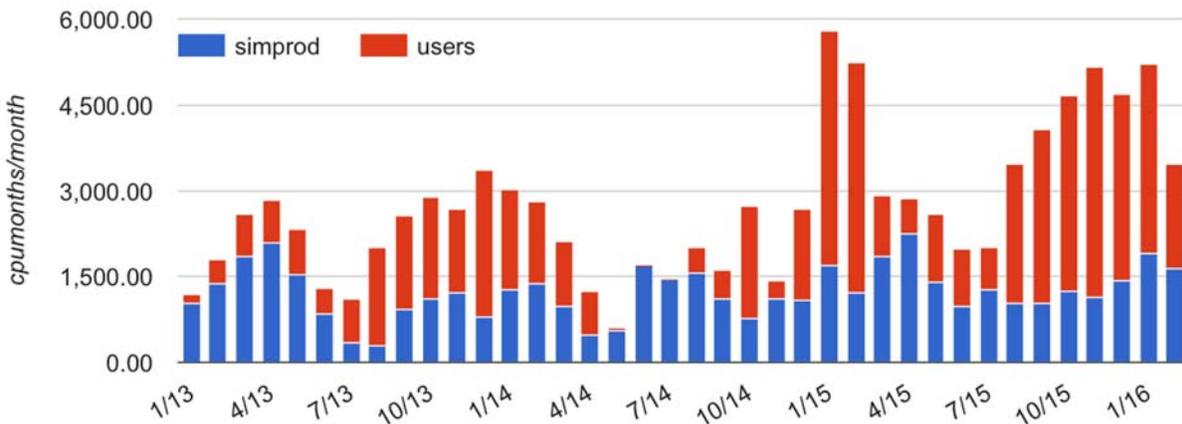
Figure 8: GPU hours consumed by the jobs completing monthly since August 2013 in the IceCube GPU cluster at UW-Madison. The contribution from Simulation Production jobs is shown in blue, and in red the contribution from other IceCube users.

**Distributed Computing** – One of the high level goals within the IceCube computing services is increasing the use of distributed computing (Grid) resources. One important aspect for simplifying the user access to Grid resources is to try and present the several heterogeneous clusters in an as unified as possible way. To pursue this, IceCube makes use of some of the federation technologies within HTCondor<sup>2</sup> such as “flocking” or “glidein” pilot jobs. These allow presenting many heterogeneous clusters as if they were one big cluster with a single interface.

IceCube benefits a lot from the opportunistic access to computing resources in other research institutions that share them. About of half of the simulation production today takes place in opportunistic resources. This is done mainly via the Open Science Grid (OSG) infrastructure. We make extensive use of the GlideinWMS infrastructure (pilot factories, VO frontend nodes) that is operated by the OSG project in order to get IceCube jobs to run at tens of US sites in an opportunistic manner.

Figure 10 shows the CPU time consumed by IceCube jobs in opportunistic distributed resources. A relatively constant use by the Simulation Production activities (in blue) is observed and, on top of that, large peaks in users activity (in red) that show the potential of tapping very large amounts of CPU from this type of resources. During the second half of 2015 for instance, IceCube was using the equivalent of more than 3000 dedicated CPUs for several months in an opportunistic manner.

<sup>2</sup> <http://research.cs.wisc.edu/htcondor/>



**Figure 10: CPU time used by IceCube jobs in opportunistic resources accessed via the Grid. These resources are mostly clusters from other departments at UW-Madison and OSG sites elsewhere in the US. The color code indicates the amount of resources that were used by Simulation Production (blue) as compared to normal users jobs (red).**

Many of the IceCube collaborating institutions provide access to local computing resources. In most cases, these resources do not have a Grid interface and access is only possible by means of a local account. During 2015 we have developed a lightweight version of a glidein pilot job factory that can be deployed as a cron job in the user’s account. This software allows us to seamlessly integrate these “local cluster” resources within the IceCube global workload system so that jobs can run anywhere in a way which is completely transparent for users.

Beyond the computing capacity provided by IceCube institutions, and the opportunistic access to Grid sites that are open to share their idle capacity, IceCube started exploring the possibility of getting additional computing resources from targeted allocation requests submitted to Supercomputing facilities such as the NSF Extreme Science and Engineering Discovery Environment (XSEDE). In October 2015 a research allocation was submitted to XSEDE that obtained positive reviews and was finally awarded a sizeable amount of resources (allocation number TG-PHY150040). The allocation included compute time in two GPU-capable systems: Comet, with 5.543.895 Service Units (SU) granted, and Bridges, with 512.665 SU. The allocation runs until the end of 2016, and our goal is to use this year to demonstrate that we can successfully integrate the IceCube simulation production workload management system with specialized resources such as XSEDE Supercomputers, and that we can make efficient use of those for extended periods of time. With the support of XSEDE ECSS resources, the integration with Comet happened during the month of February, and we started running simulation GPU jobs in that system on March 4<sup>th</sup> 2016. By the time of writing this report, IceCube has used 81.954 SU in Comet.

One of the most critical services in a distributed computing system is the data access. In order to efficiently benefit from potentially very large opportunistic computing capacity peaks, the service at UW-Madison that feeds the input data to those processes and ingests the outputs has to scale accordingly. IceCube data at UW-Madison can be remotely accessed with high performance by means of two main protocols: gridftp and http. The network connectivity of the IceCube data center at UW-Madison has been greatly improved in the last year with the help of

UW-Madison Networking teams. In April 2015 the Internet connectivity was raised to 20Gbps, after being capped to 4Gbps for many years. Today, the data export import rates out of the UW-Madison site routinely exceed the 10Gbps levels. In February 2016 the network backbone connecting the two main data centers at UW-Madison was upgraded from 40Gbps to 80Gbps. This greatly increases the data processing capabilities of our facility, since it boosts the capacity of the network connecting the cluster nodes to the data warehouse.

**Data Reprocessing** - At the end of 2012, the IceCube Collaboration agreed to store the compressed SuperDST as part of the long-term archive of IceCube data. The decision taken was that this change would be implemented from the IC86-2011 run onwards. A server and a partition of the main tape library for input were dedicated to this data reprocessing task. Raw tapes are read to disk and the raw data files processed into SuperDST. A copy is saved in the data warehouse. It is planned to save an additional copy to NERSC. During the reporting period, the system to translate raw tapes into SuperDST files was modified to use the Iceprod framework, to streamline operations and allow access to a greater resource pool. The ideal capacity of the tape system is about 4000 files/day, while the actual rate of the full system is nearer 2400/day when space is available.

During 2015 the possibility of partnering with the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL) for archiving data was explored. In December 2015, a Memorandum of Understanding was signed between UW-Madison and NERSC/LBNL by which NERSC agreed to provide long-term archive services for the IceCube data until 2019. By implementing the long-term archive functionality using a storage facility external to the UW-Madison data center, we can aim for a better service at lower cost, since large facilities that routinely manage data at the level of 100 petabytes benefit of economies of scale that ultimately make the process more efficient and economical.

**Offline Data Filtering** – The data collection for the IC86-2015 season started in May 18 2015. A new compilation of data processing scripts had been previously validated and benchmarked with the data taken during the 24-hour test run using the new configuration. The differences with respect to the IC86-2014 season scripts are minimal, therefore we estimate that the resources required for the offline production will be of about 750,000 CPU hours on the IceCube cluster at UW-Madison datacenter. 120TB of storage is required to store both the Pole-filtered input data and the output data resulting from the offline production. Data processing proceeds with no issues and Level2 data products are typically available two weeks after data taking. Replication of all the data at the DESY-Zeuthen collaborating institution is being done in a timely manner comparable to previous seasons. Further checks of the data integrity during replication have added an extra layer of validation to the offline production process.

Much progress has been made on the proposed effort to adopt the tools currently employed by the centralized offline production in the post-offline data processing for the various physics analysis groups. This transition enables the efficient automation and monitoring of production. It also ensures better coordination and improvement in resource management and planning. After multiple successful tests, the three main physics analysis groups in the collaboration have now adopted this framework that facilitates producing Level3 data products in a shorter time scale and using consistent procedures with Level2 production in terms of cataloging and bookkeeping.

**Simulation** – The production of IC86 Monte Carlo began in mid-2012 with simulations of the IC86-2011 detector configuration. A transition to a combined production of IC86-2012, and IC86-2013 simulation started at the beginning of 2014 in tandem with a new release of the IceCube simulation software, IceSim4, including Level 2 filtering of simulation to reduce storage space requirements. IceSim 4 contains improvements to low-level DOM simulation, correlated noise generation, Earth modeling, and lepton propagation. At the beginning of 2016 we are again transitioning to a new major release of IceSim 5 that includes major performance optimizations and a transition to IC86-2014 through IC86-2016. We have progressed toward having 100% of all simulations based on direct photon propagation using GPUs or a hybrid of CPU and spline-photonics for high-energy events. Producing simulations of direct photon propagation using GPUs began with a dedicated pool of computers built for this purpose in addition to the standard CPU-based production. Benchmark performance studies of consumer-class GPU cards have been completed and provided to the collaboration as we scale up the available GPUs for simulation. Starting in 2015, we began incorporating more opportunistic computing including GPU clusters through Open Science Grid and XSEDE. We are currently testing a newly redesigned production management framework IceProd v2 that will scale with the growing availability of computing resources.

The simulation production sites are: CHTC – UW campus (including GZK9000 GPU cluster); Dortmund; DESY-Zeuthen; University of Mainz; EGI – German grid; WestGrid – U. Alberta; SWEGRID – Swedish grid; PSU – Pennsylvania State University; GLOW – Grid Laboratory of Wisconsin; UMD – University of Maryland; RWTH Aachen; IIHE – Brussels; UGent – Ghent; Ruhr-Uni – Bochum; UC Irvine; Michigan State University - ICER; The Extreme Science and Engineering Discovery Environment (XSEDE); Niels Bohr Institute, Copenhagen Denmark; Cori – LBNL; and NPX – UW IceCube.

**Personnel** – Two open positions were filled in January 2016:

- Linux system administrator position – filled by Ian Saunders.
- Offline data processing programmer position – filled by Jan Oertlin.

## ***Data Release***

***Data Use Policy*** – IceCube is committed to the goal of releasing data to the scientific community. The following links contain data sets produced by AMANDA/IceCube researchers along with a basic description. Due to challenging demands on event reconstruction, background rejection and systematic effects, data will be released after the main analyses are completed and results are published by the international IceCube Collaboration.

***Datasets (last release on 20 Aug 2015):*** <http://icecube.wisc.edu/science/data>

The pages below contain information about the data that were collected and links to the data files.

The 79-string IceCube search for dark matter:

<http://icecube.wisc.edu/science/data/ic79-solar-wimp>

Observation of Astrophysical Neutrinos in Four Years of IceCube Data:

<http://icecube.wisc.edu/science/data/HE-nu-2010-2014>

Astrophysical muon neutrino flux in the northern sky with 2 years of IceCube data:

[https://icecube.wisc.edu/science/data/HE\\_NuMu\\_diffuse](https://icecube.wisc.edu/science/data/HE_NuMu_diffuse)

IceCube-59: Search for point sources using muon events:

<https://icecube.wisc.edu/science/data/IC59-point-source>

Search for contained neutrino events at energies greater than 1 TeV in 2 years of data:

[http://icecube.wisc.edu/science/data/HEnu\\_above1tev](http://icecube.wisc.edu/science/data/HEnu_above1tev)

IceCube Oscillations: 3 years muon neutrino disappearance data:

[http://icecube.wisc.edu/science/data/nu\\_osc](http://icecube.wisc.edu/science/data/nu_osc)

Search for contained neutrino events at energies above 30 TeV in 2 years of data:

<http://icecube.wisc.edu/science/data/HE-nu-2010-2012>

IceCube String 40 Data:

<http://icecube.wisc.edu/science/data/ic40>

IceCube String 22–Solar WIMP Data:

<http://icecube.wisc.edu/science/data/ic22-solar-wimp>

AMANDA 7 Year Data:

<http://icecube.wisc.edu/science/data/amanda>

## ***Program Management***

***Management & Administration*** – The primary management and administration effort is to ensure that tasks are properly defined and assigned and that the resources needed to perform each

task are available when needed. Efforts include monitoring that resources are used efficiently to accomplish the task requirements and achieve IceCube’s scientific objectives.

- The FY2015 M&O Plan was submitted in January 2015.
- The detailed M&O Memorandum of Understanding (MoU) addressing responsibilities of each collaborating institution was revised for the collaboration meeting in Stony Brook, NY, April 16-22, 2016.

IceCube M&O – FY2015-FY2016 Milestones Status:

Milestone	Month
Submit for NSF approval, a revised IceCube Maintenance and Operations Plan (M&OP) and send the approved plan to non-U.S. IOFG members.	January 2015
Annual South Pole System hardware and software upgrade is complete.	January 2015
Submit to NSF a mid-year interim report with a summary of the status and performance of overall M&O activities, including data handling and detector systems.	March 2015
Revise the Institutional Memorandum of Understanding (MOU v18.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration Meeting	April 2015
Post the revised institutional MoU’s and Annual Common Fund Report and notify IOFG.	April 2015
Report on Scientific Results at the Spring Collaboration Meeting	Apr 28 - May 2, 2015
IceCube IOFG Meeting	May 3, 2015
Submit for NSF approval an annual report which will describe progress made and work accomplished based on objectives and milestones in the approved annual M&O Plan.	September 2015
Revise the Institutional Memorandum of Understanding (MOU v19.0) - Statement of Work and Ph.D. Authors head count for the fall collaboration meeting	October 2015
Annual South Pole System hardware and software upgrade is complete.	January 2016
Software & Computing Advisory Panel (SCAP) Review at UW-Madison	March 21-22, 2016
Revise the Institutional Memorandum of Understanding (MOU v20.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration Meeting	April 2016
Submit to NSF a final report with a summary of the status and performance of overall M&O activities, including data handling and detector systems.	June 2016

**Engineering, Science & Technical Support** – Ongoing support for the IceCube detector continues with the maintenance and operation of the South Pole Systems, the South Pole Test System, and the Cable Test System. The latter two systems are located at the University of Wisconsin–Madison and enable the development of new detector functionality as well as investigations into various operational issues, such as communication disruptions and electromagnetic interference. Technical support provides for coordination, communication, and assessment of impacts of activities carried out by external groups engaged in experiments or potential experiments at the South Pole.

**Education & Outreach (E&O)** – The collaboration-wide IceCube E&O efforts have coalesced around four main themes:

- 1) Reaching motivated high school students and teachers through [IceCube Masterclasses](#)

- 2) Providing intensive research experiences for teachers (in collaboration with [PolarTREC](#)) and for undergraduate students (NSF science grants, International Research Experience for Students (IRES), and Research Experiences for Undergraduates (REU) funding)
- 3) Engaging the public through web and print resources, graphic design, webcasts with IceCube staff at the Pole, and displays
- 4) Developing and implementing semiannual communication skills workshops held in conjunction with IceCube Collaboration meetings

The third IceCube Masterclasses were held in March, 2016, with a total of ten institutions participating from the United States, Belgium, and Germany and over 200 students. Accompanying web resources were updated and provided in three languages: English, German, and Spanish. For the first time, the WIPAC team in Madison offered a masterclass in Spanish, attended by 28 students. Pre- and post-masterclass surveys indicate that the students were challenged and motivated by the experience, and especially appreciated the chance to work with real data and active scientists. The new cosmic ray analysis activity developed by Hans Dembinski at the University of Delaware was described in a poster and proceedings at the 2015 International Cosmic Ray Conference at The Hague, Netherlands. We continue to promote the masterclass, encouraging collaborators to host and/or develop new analysis.

[Kate Miller](#), a physics teacher from Washington-Lee High School in Arlington, VA, has been selected to work with the IceCube project and the PolarTREC program, deploying to the South Pole in the 2016-17 season. Kate will contribute to the UW–River Falls Upward Bound program (July 11-22, 2016), helping teach math and science along with former IceCube PolarTREC teachers Liz Ratliff and longtime high school teacher Eric Muhs, who deployed with the AMANDA project in the 2011-12 season. Kate will be the fifth teacher to work and deploy with the IceCube project since the 2009-10 season.

Five undergraduate students (two women, one of whom is African American) have been selected to travel and work at Stockholm University for ten weeks in the summer of 2016, supported by NSF International Research Experience for Students (IRES) funding. The students are from UW–River Falls, UW–Madison, and the University of Minnesota. The UW–River Falls astrophysics NSF Research Experience for Undergraduates program selected six students (two men, one of whom is African American, four students total from two-year colleges) from over 60 applicants for ten-week summer 2016 research experiences. These students will participate in the WIPAC IceCube software and science boot camp. Most IceCube institutions have funding to support undergraduates, who make important contributions to the collaboration. Greater effort is being put into ensuring that they have opportunities to interact with each other. Two of the IRES students from the summer of 2015 will attend the IceCube Collaboration meeting in April, 2016, at Stony Brook.

The IceCube scale model, with one colored LED for each of the 5,160 DOMs, continues to be exhibited, and a second model is under construction at Ghent University in Belgium. A more compact 1m x 1m x 1m model is being developed at York University with art/science professor and three-year collaborator Mark-David Hosale. Five South Pole webcasts, including one in Spanish, were held this past season, bringing stories of life and research at the Pole to hundreds of participants from 22 schools in six countries. Over the past four years since inaugurating the

open webcasts, we have had thousands of people from 65 schools, in 10 countries, participate and interact with researchers at the Pole in English, Spanish and German.

The communication training program targeting PhD students and postdocs, launched at the spring 2015 IceCube Collaboration in Madison, is helping participants develop skills for reaching all types of audiences, from journalists to science-skeptical and lay audiences. The first workshop focused on understanding one's audience and was attended by 20 collaborators. The workshop at the fall 2015 IceCube Collaboration meeting in Denmark developed experience-based learning activities. For the spring 2016 IceCube collaboration meeting at Stony Brook, we have partnered with the [Story Collider](#) group. They will give an interactive presentation on using narratives to explain research, connecting to audiences in meaningful ways, and engaging in cultural conversations about science.

Finally, we have increased the dissemination of IceCube E&O efforts. We presented talks on the E&O program at the Scientific Conference on Antarctic Research meeting in New Zealand (August 2014), the American Geophysical Union (AGU) Fall Meeting in San Francisco (December 2014), the International Teacher-Scientist Partnership Conference in San Francisco (February 2015), the Broader Impact Summit in Madison (April 2015), and the Wisconsin Science Festival teacher workshops (October 2015). Posters were presented on the masterclass at the AGU Fall Meeting in San Francisco (December 2014) and on the new cosmic ray masterclass at the 34th International Cosmic Ray Conference in The Hague (August 2015).

The E&O team works closely with the communication team. Science news summaries of IceCube publications, written at a level accessible to science-literate but non-expert audiences, continue to be produced and posted regularly on the IceCube website and highlighted on social media. Media mentions that have been tracked since January 2015 include over 170 news pieces appearing in 20 different countries, with over 50% of those in the US. Besides national and international coverage in the media, we have documented local news mentions in 9 different U.S. states.

**Ongoing local E&O efforts** - Two examples of the extensive WIPAC E&O efforts beyond the four core areas of collaboration-wide E&O include:

i) High School Internship Program—A dozen postdocs and graduate students mentor about 25 area high school students for ten-week internships during the school year and seven-week internships during the summer. The high school interns use data from real physics experiments to learn about topics in astrophysics, computer programming, and data processing. Post-internship surveys indicate that participants gained new skills, were exposed to the scientific research process, and appreciated working with active scientists. Three high school interns have subsequently been hired as undergraduate researchers at WIPAC.

ii) Science Alliance—WIPAC is an active member in the Science Alliance, a group of researchers, outreach professionals, and volunteers from UW–Madison and the community. Each year, this group organizes and hosts several science festivals and activities in Madison and around Wisconsin that allow us to talk about IceCube to thousands of people every year. Through this partnership, we have also shown *Chasing the Ghost Particle* to over 500 people in

three years. Chasing the Ghost particle has been purchased by 36 planetariums across the country.

### **Section III – Project Governance and Upcoming Events**

The detailed M&O institutional responsibilities and Ph.D. author head count is revised twice a year at the time of the IceCube Collaboration meetings. This is formally approved as part of the institutional Memorandum of Understanding (MoU) documentation. The MoU was last revised in April 2016 for the Spring collaboration meeting in Madison, WI (v20.0), and the next revision (v21.0) will be posted in October 2015 at the Fall collaboration meeting in Mainz, Germany.

#### ***IceCube Collaborating Institutions***

Following the October 2015 fall collaboration meeting, the University of Rochester, with Segev BenZvi as the institutional lead, and Marquette University, with Karen Andeen as the institutional lead, were approved as full members of the IceCube Collaboration.

As of March 2016, the IceCube Collaboration consists of 47 institutions in 12 countries (25 U.S. and Canada, 18 Europe, and 4 Asia Pacific).

The list of current IceCube collaborating institutions can be found on:  
<http://icecube.wisc.edu/collaboration/institutions>

#### ***IceCube Major Meetings and Events***

Software & Computing Advisory Panel (SCAP) Review, Madison, WI	March 21-22, 2016
IceCube Science Advisory Committee (SAC) Meeting, Madison, WI	October 19-20, 2015
IceCube Fall Collaboration Meeting – Copenhagen, Denmark	October 12-16, 2015
IceCube Particle Astrophysics Symposium – UW–Madison	May 4-6, 2015
IceCube IOFG Meeting – Madison, WI	May 3, 2015
IceCube Spring Collaboration Meeting – Madison, WI	April 28 – May 2, 2015
National Academies of Sciences Study on Strategic Vision FOR USAP	October 21, 2014
IceCube Fall Collaboration Meeting – CERN (Geneva), Switzerland	September 15–19, 2014
IceCube Spring Collaboration Meeting – Banff, Canada	March 3-8, 2014
Software & Computing Advisory Panel Meeting (SCAP) – Madison, WI	April 1-2, 2014
IceCube Fall Collaboration Meeting – Munich, Germany	October 8–12, 2013
IceCube Spring Collaboration Meeting – UW–Madison	May 7-11, 2013
IceCube Particle Astrophysics Symposium – UW–Madison	May 13-15, 2013
NSF M&O Midterm Progress and Operations Review	May 15-17, 2013
IceCube Fall Collaboration Meeting – Aachen, Germany	October 1-5, 2012
IceCube Spring Collaboration Meeting – University of California, Berkeley	March 19-23, 2012
NSF M&O Reverse Site Visit	May 24, 2012
IceCube IOFG Meeting – Uppsala, Sweden	September 21, 2011
IceCube Fall Collaboration Meeting at Uppsala University, Sweden	September 19-23, 2011
IceCube Collaboration Meeting, Madison	April 25 – May 2, 2011

### *Acronym List*

CnV	Calibration and Verification
CPU	Central Processing Unit
CVMFS	CernVM-Filesystem
DAQ	Data Acquisition System
DOM	Digital Optical Module
E&O	Education and Outreach
I3Moni	IceCube Run Monitoring system
IceCube Live	The system that integrates control of all of the detector's critical subsystems; also "I3Live"
IceTray	IceCube core analysis software framework, part of the IceCube core software library
MoU	Memorandum of Understanding between UW–Madison and all collaborating institutions
PMT	Photomultiplier Tube
PnF	Processing and Filtering
SNDAQ	Supernova Data Acquisition System
SPE	Single photoelectron
SPS	South Pole System
SuperDST	Super Data Storage and Transfer, a highly compressed IceCube data format
TDRSS	Tracking and Data Relay Satellite System, a network of communications satellites
TFT Board	Trigger Filter and Transmit Board
WIPAC	Wisconsin IceCube Particle Astrophysics Center