IceCube Maintenance and Operations
Fiscal Year 2015 Mid-Year Report

October 1, 2014 – March 31, 2015

Submittal Date: March 31, 2015

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 FY2015 Mid-Year Report is submitted as required by the NSF Cooperative Agreement ANT-0937462. This report covers the six-month period beginning October 1, 2014 and concluding March 31, 2015. The status information provided in the report covers actual common fund contributions received through March 18, 2015 and the full 86-string IceCube detector (IC86) performance through March 1, 2015.

Table of Contents

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 $6,900,000 was released to UW–Madison to cover the costs of maintenance and operations in FY2015: $914,550 was directed to the U.S. Common Fund account based on the number of U.S. Ph.D Authors in the last version of the institutional MoU's, and the remaining $5,985,450 was directed to the IceCube M&O Core account (Table 1).

FY2015	Funds Awarded to UW
IceCube M&O Core account	$5,985,450
U.S. Common Fund account	$914,550
TOTAL NSF Funds	**$6,900,000**

Table 1: NSF IceCube M&O Funds - FY2015

Of the IceCube M&O FY2015 Core funds, $925,475 were committed to the U.S. subawardee institutions based on their statement of work and budget plan. The institutions submit invoices to receive reimbursement against their actual IceCube M&O costs. Table 2 summarizes M&O responsibilities and total FY2015 funds for the subawardee institutions.

Institution	Major Responsibilities	Funds
Lawrence Berkeley National Laboratory	DAQ maintenance, computing infrastructure	$81,516
Pennsylvania State University	Computing and data management, simulation production	$41,488
University of California at Berkeley	Detector calibration, monitoring coordination	80,895
University of Delaware, Bartol Institute	IceTop calibration, monitoring and maintenance	$133,797
University of Maryland at College Park	IceTray software framework, online filter, simulation software	$518,906
University of Alabama at Tuscaloosa	Detector calibration, reconstruction and analysis tools	$45,501
Michigan State University	Simulation software, simulation production	$23,372
Total		**$925,475**

Table 2: IceCube M&O Subawardee Institutions – FY2015 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. The following table presents the financial status five months into FY2015, which is Year 5 of the award, and shows an estimated balance at the end of FY2015.

Total awarded funds to the University of Wisconsin (UW) for supporting IceCube M&O from the beginning of FY2011 through the end of FY2015 are $34,694K. Total actual cost as of February 28, 2015 is $29,395K and open commitments are $927K. The current balance as of February 28, 2015 is $4,372K. With a projection of $4,181K for the remaining expenses during the final seven months of FY2015, the estimated unspent funds at the end of FY2015 are $191K, which is 2.8% of the FY2015 budget (Table 3).

(a)	(b)	(c)	(d)= a - b - c	(e)	(f) = d – e
YEARS 1-5 Budget Oct.'10-Sep.'15	**Actual Cost To Date** through Feb. 28, 2015	**Open Commitments** on Feb. 28, 2015	**Current Balance** on Feb. 28, 2015	**Remaining Projected Expenses** through Sept. 2015	**End of FY2015 Forecast Balance** on Sept. 30, 2015
$34,694K	**$29,395K**	**$927K**	**$4,372K**	**$4,181K**	**$191K**

Table 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.

Table 4 summarizes the planned and actual Common Fund contributions for the period of April 1, 2014–March 31, 2015, based on v16.0 of the IceCube Institutional Memorandum of Understanding, from March 2014. Actual Common Fund contributions are $43K less than planned, mostly due to the departure of the University of Bonn from IceCube Collaboration.

	Ph.D. Authors	Planned Contribution	Actual Received
Total Common Funds	**129**	**$1,790,150**	**$1,671,395**
U.S. Contribution	67	$914,550	$914,550
Non-U.S. Contribution	62	$861,950	$818,777

Table 4: Planned and Actual CF Contributions for the period of April 1, 2014–March 31, 2015

Section II – Maintenance and Operations Status and Performance

Detector Operations and Maintenance

Detector Performance – During the period from July 1, 2014, to March 1, 2015, the 86-string detector configuration (IC86) operated in with the full detector configuration for an unprecedented 97.9% of the time. The hardware upgrades of the 2013–14 polar summer season, along with continued data acquisition software improvements, have increased the detector stability and resulted in partial detector runs accounting for only 0.97% of detector time. Figure 1 shows the cumulative detector time usage over the reporting period. Maintenance, commissioning, and verification required 1.05% of detector time. The unexpected downtime due to failures is 0.09%.

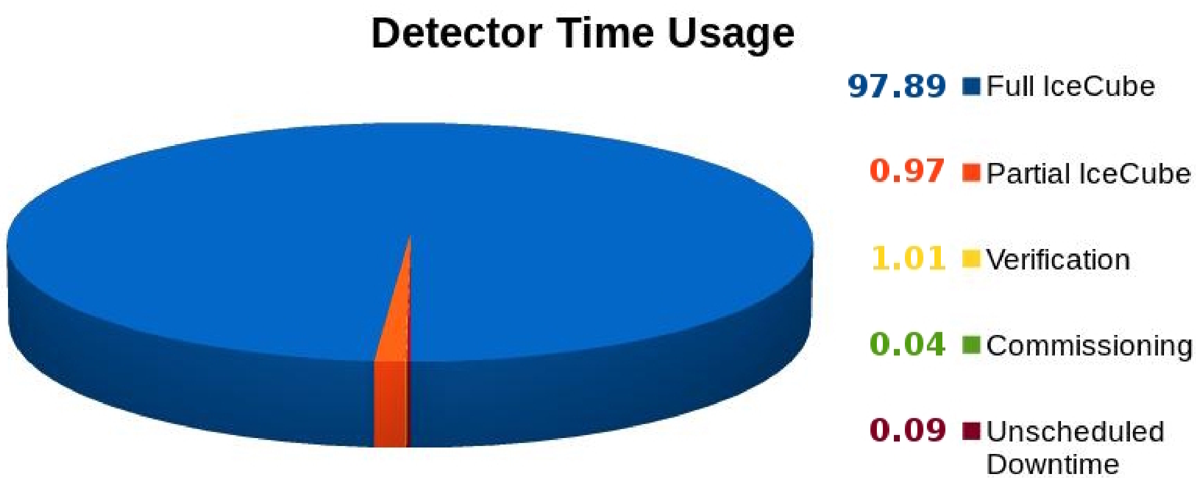

Figure 1: Cumulative IceCube Detector Time Usage, July 1, 2014 – March 1, 2015

Over the reporting period of July 1, 2014, to March 1, 2015, an average detector uptime of 99.5% was achieved, comparable to previous years, as shown in Figure 6. Of the total, 93.6% of the time the runs were successful, meaning the DAQ did not fail within the set 8-hour run duration and used the full 86-string detector. Now implemented is the ability to track portions of failed runs as good, and this allows the recovery of data from all but the last few minutes of runs that fail. This feature brings the average "clean uptime" for this reporting period up to 97.2% for standard analysis-ready, full-detector data, as shown in Figure 2.

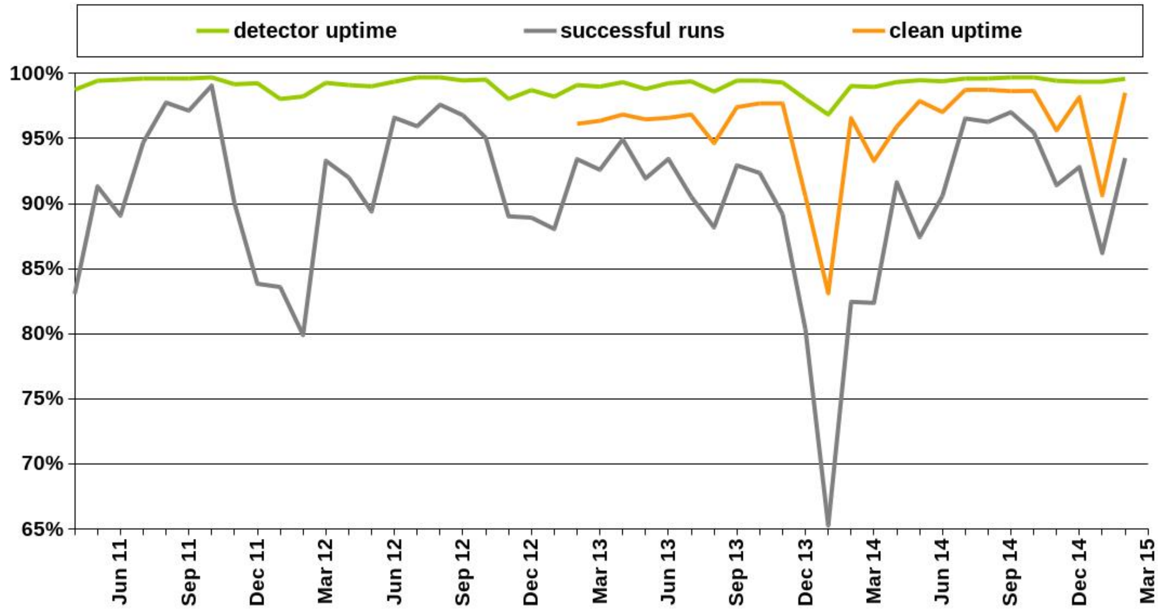

Figure 2: Total IceCube Detector Uptime and Clean Uptime

About 1% of the loss in clean uptime is due to the failed portions of runs that are not usable for analysis and runs less than 10 minutes long. About 1% of the loss in clean uptime is 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. Additionally, the experiment control system and DAQ are in the commissioning phase towards a non-stop (continuous) run configuration. Continuous data-taking will become the standard mode beginning in 2015; this will eliminate the approximately 90–120 seconds of downtime between run transitions, gaining roughly 0.5% of uptime.

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 ROOT 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 via a ZeroMQ messaging system. 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 three major subsystems: data acquisition (DAQ), supernova DAQ, and processing and filtering (PnF). A workshop in June 2014 focused on test integration and collection of quantities as well as the development and integration of sophisticated quality control systems. I3Moni 2.0 is currently an alpha release phase, with the beta release expected by Fall 2015.

Development of IceCube Live, the experiment control and monitoring system, is still quite active, and this reporting period has seen three major releases. The primary features of these releases are:

- Live v2.6 (August 2014): 41 separate issues and requested features resolved. Multiple improvements to IceCube Live's message transport service were made, including the forwarding of neutrino event data to the north via the low-latency ITS Iridium link to support multi-wavelength follow-up observations. In addition, the web site's plotting tools have been entirely refactored, allowing for better analysis and easier maintenance.

- Live v2.7 (October 2014) and v2.8 (February 2015): 63 separate issues and requested features were resolved. New features include integration of the IceCube Live low-latency chat service with the popular Slack communication platform, allowing the IceCube winterover operators to participate in a group chat widely used by scientists in the north. Other features include the ability to restart the experiment control portion of IceCube Live without stopping IceCube data-taking. As usual, user interface improvements were made to facilitate data analysis and support the winterovers' effort. Version 2.8 also includes the alpha release of the I3Moni 2.0 monitoring system.

Features planned for the next few releases include support for the Iridium RUDICS mode, potentially offering lower latency and much higher bandwidth than the Iridium short-burst data mode currently used by IceCube Live as its lowest-latency transport link. A switch to the production phase of the I3Moni 2.0 monitoring system will take place when the preliminary testing phase is complete. The uptime for the I3Live experiment control system during the reporting period was 99.998%.

The IceCube Data Acquisition System (DAQ) has reached a stable state, and consequently the frequency of software releases has slowed to the rate of 1–2 per year. Nevertheless, the DAQ group continues to develop new features and patch bugs. During the reporting period of July 2014–March 2015, the following accomplishments are noted:

- Delivery of the DAQ release on August 21, 2014. This release includes fixes to pDAQ's monitoring and run switching code, the latter of which is integral to the upcoming "continuous run" mode. Cluster configuration files are now dynamically reloaded between runs, enabling operators to modify software operation without relaunching the entire cluster. A second Hinterland release was necessary when the system-wide ZeroMQ library was updated.

- Delivery of the DAQ release on January 27, 2015. This release replaces the mix of semi-related operational commands with a unified 'pdaq' command with subcommands for the various functions ('pdaq launch', 'pdaq kill', etc.); updates all dependent libraries to their latest versions and eliminates some obsolete dependencies; includes internal monitoring code to track down StringHub problems; and fixes some minor monitoring and continual run bugs.

- Delivery of the DAQ release on March 9, 2015. This release includes optimization of the ClusterTrigger (now 5 times faster) along with system-wide speedup of some core functionality; minor improvements to StringHub; major improvements to the "replay" feature augmented by 8 hours of detector-wide data from back from the South Pole enabling better intra- and inter-system testing of all operations software in the North.

- Development toward the next DAQ release. This release will include major improvements to the StringHub component to eliminate the out-of-order and LBM overflow problems which result in discarding the first 20 seconds or so of each run, as well as a new priority-queue-based sorter which should speed up the input side of most components.

The supernova data acquisition system (SNDAQ) found that 98.7% of the available data from July 1st, 2014 through March 1st, 2015, met the minimum analysis criteria for run duration and data quality for sending triggers. An additional 0.1% of the data is available in short physics runs with < 10 minute duration. While forming a trigger is not possible in these runs, the data are available for reconstructing a supernova signal.

New SNDAQ releases (July 2014–March 2015) were installed. The early releases featured several minor fixes to handle the real-time cosmic ray muon rejection and to improve the logging. With release v7 (October 13, 2014), the muon rejection was running in a stable way in its final data-taking configuration. An upgrade to the ZeroMQ messaging system demonstrated that a coordination of the deployment sequence between pDAQ and SNDAQ is very important. When the supernova DAQ server failed due to hardware problems, the procedure to switch to a spare machine was tested and improved; SNDAQ can now be installed and run in a transparent way on any suitable spare machine. Recently, in the v9 release, the latency for sending alarm emails was reduced. The winter-over training manual was reworked with emphasis on improving its clarity.

Future supernova DAQ development efforts will include the switch to triggers that are based on cosmic-ray-muon-subtracted data. By studying data taken in previous years, the settings for such triggers were optimized, leading to a proposal that incorporates substantially lowered trigger thresholds. For a supernova in the Magellanic Cloud, e.g., the probability to pass the SNEWS trigger can be improved from 12% to 85% with the new trigger settings. We still need, however, to thoroughly test whether the new settings allow for a similarly stable running as the present configuration.

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. Version V15-02-00 was deployed in February 2015 and supports the IC86–2014 physics run. This release also includes software that produces data-quality monitoring values for the I3Moni 2.0 system which are reported in real-time to the IceCube Live experiment control system. The release also brings an upgrade to the online neutrino alert system and now sends summary information for all high-quality neutrino-alert candidate events north via the ITS system for use in alerts to other astronomical instruments. Final testing and verification of DAQ continuous run operation has also been completed. Current work is now focused on preparation of the filtering system for the transition to the IC86–2015 physics runs.

A weekly calibration call keeps collaborators abreast of issues in both in-ice and offline DOM calibration. Implementation of online DOM gain corrections using the single photoelectron (SPE) peak will be implemented as of the IC86–2015 run transition in May 2015. Online monitoring of the in-ice SPE charge peaks and all baselines using the new I3Moni 2.0 framework has begun (see Figure 3). Previously, these quantities were monitored with a separate system, which is planned to be decommissioned by the end of 2015. This summer, we plan to introduce IceTop charge monitoring into I3Moni 2.0.

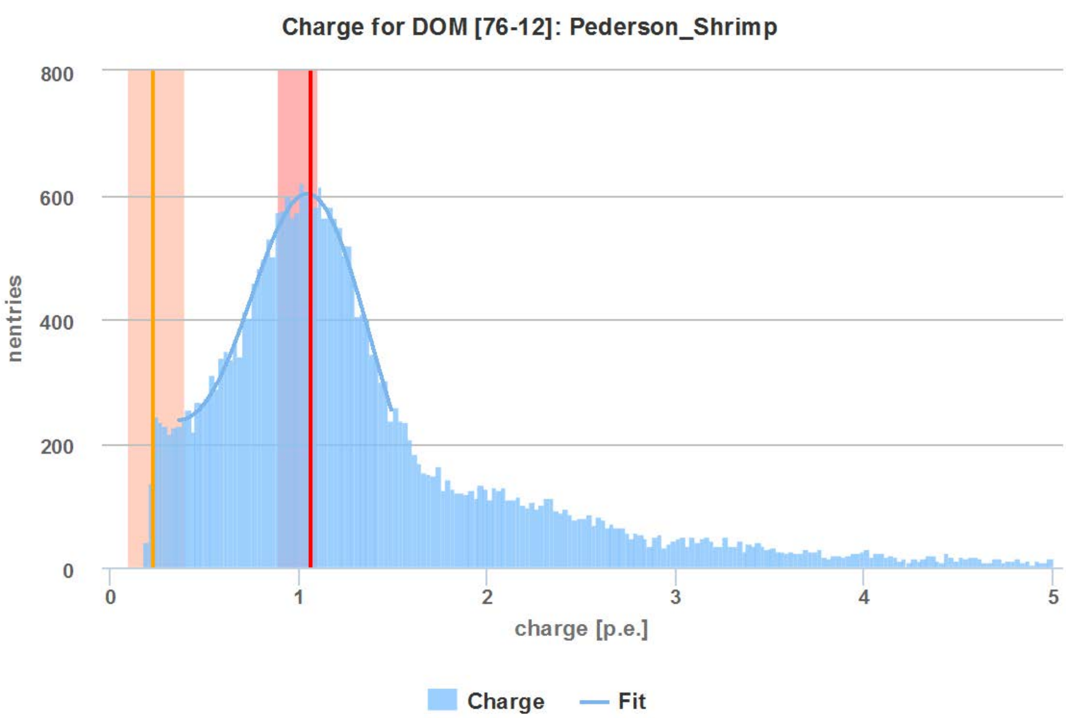

Figure 3: Measured single photoelectron (SPE) charge spectrum in I3Moni 2.0, for DOM 76-12. The vertical lines indicate the fitted threshold (left) and charge peak (right). The peak is used in DOM gain correction.

Flasher calibration data taken during the austral summer seasons continues to be analyzed to reduce ice model systematics, with a major focus on the refrozen hole ice. A new DOM angular sensitivity curve has been fitted to the flasher data, showing the modification in sensitivity as a function of photon angle due to the hole ice, shown in Figure 4. Efforts continue to verify the mean DOM sensitivity and individual DOM sensitivity fluctuations using both flashers and low energy muons.

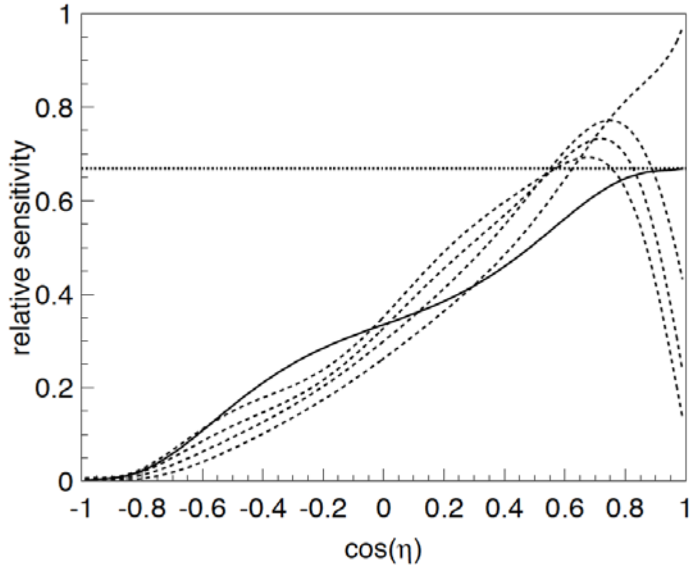

Figure 4: DOM angular sensitivity as a function of photon angle Left is upgoing photons, right is downgoing photons. The solid line represents the new sensitivity curve fitted to flasher data, assuming 50 cm scattering length in the hole. The dashed lines which bend downward at the right represent the previously-used models for various hole ice scattering lengths. The dashed line which trends upward on the right is for no hole ice. Overall the effect of hole ice is to increase the sensitivity of the DOM to downgoing photons, which are upscattered into the DOM.

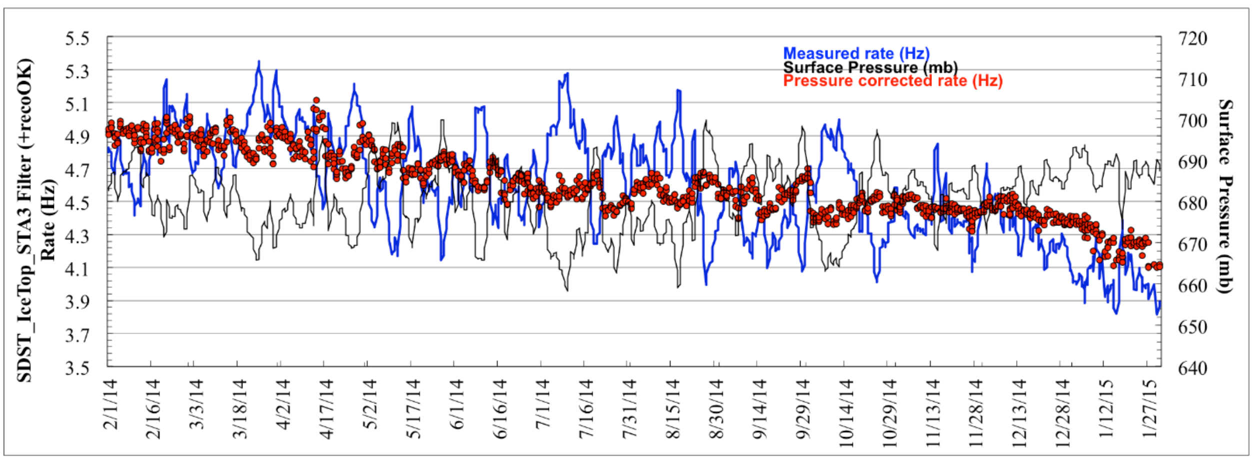

Figure 5: Decrease in IceTop filter rate in 2014 due to snow accumulation. Blue shows the uncorrected rate; red, the rate corrected for barometric pressure.

A procedure for measuring the depth of snow above IceTop surface array tanks is in place and continues to work well. However, despite previous efforts at managing the snow accumulation on the tanks, the increasing snow depth continues to increase the energy threshold for cosmic ray air shower detection each year; a continued steady decrease in filter-level event rate of ~15% per year is shown in Figure 5. Moreover, the snow overburden complicates analyses due to uncertainty in the response to different particle types. During the 2014–15 austral summer season, we conducted a series of measurements with portable muon taggers over IceTop tanks to calibrate the difference in response to muons and electromagnetic particles, and analysis of this data is now underway. Figure 6 shows the current snow depth and the locations of high accumulation.

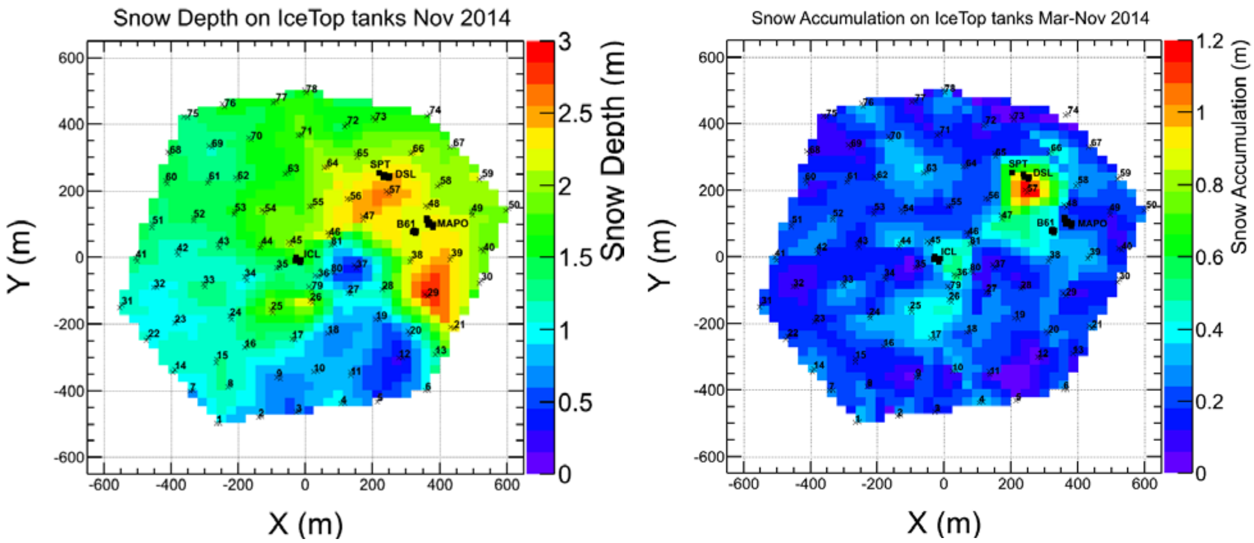

Figure 6: Left, snow depth above IceTop tanks as of November 2014. Right: snow accumulation from the period Mar.–Nov. 2014. Management of the snow accumulation by the contractor has been unsuccessful and has resulted in a 15% decrease per year in event rate, necessitating the development of other methods to restore the detector efficiency.

Furthermore, a 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. We plan to deploy plastic scintillator panels on the snow surface above the buried IceTop tanks; two to four of these panels will be ready for deployment in the 2015–16 pole season. The initial deployment will use 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. These prototypes will use existing DOM mainboards for readout, but modernization of the electronics will be required for a larger deployment over more IceTop tanks.

IC86 Physics Runs – The fourth season of the 86-string physics run, IC86–2014, began on May 6, 2014. Detector settings were updated using the latest yearly DOM calibrations from March 2014, but minimal changes to filter settings were implemented. DAQ trigger settings did not change from IC86–2013. Preparations for the IC86–2015 physics run, expected to begin in May 2015, are underway.

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). The IC86–2013 season was the first instance in which L2 processing was handled by the TFT from beginning to end. The result is a dramatic reduction in the latency for processing of L2. For the IC86–2011 season, L2 latency was several months after the end of the data season. For IC86–2013 and the current physics run IC86–2014, the latency is 2–4 weeks after data-taking.

Working groups within IceCube will submit approximately 20 proposals requesting data processing, satellite bandwidth and data storage, and the use of various IceCube triggers for IC86–2014. Sophisticated online filtering data selection techniques are used on SPS to preserve bandwidth for other science objectives. Over the past three years, new data compression algorithms (SuperDST) have 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.

Starting with IC86–2014, we have begun to implement 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 IC86–2014 run start in May, the average TDRSS daily transfer rate is approximately 90 GB/day. IceCube is a heavy user of the available bandwidth, and we will continue to moderate our usage without compromising the physics data.

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 the IceCube Teleport System (ITS) using a dedicated Iridium modem in short-burst data mode.

Low-bandwidth text chat is used to communicate with the IceCube winterovers to coordinate operations and diagnose problems. Skype has been successfully used for this purpose for many years in IceCube, but during 2014, we were notified of the necessity to stop the use of Skype at the South Pole. Support for Jabber as an alternative has been intermittent and unreliable. We have now implemented our own chat system via I3Live that uses IceCube's own Iridium modem.

With the TDRSS F5 satellite retirement, the total high-bandwidth satellite coverage was reduced by approximately 6 hours per day, because TDRSS 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, has led the contractor to eliminate 24/7 e-mail delivery service except for operational needs. Because IceCube relies on the IMCS link to deliver small quantities of critical monitoring information, a dedicated e-mail queue for this purpose was retained. Nevertheless, we are now developing our own Iridium RUDICS-based transport software and will move IceCube's monitoring traffic to our own modem(s) in the 2015–16 austral summer season.

A common theme is a degradation of the connectivity to pole that is driving IceCube to develop parallel solutions to those formerly provided to the entire South Pole community. This is a discouraging trend that we hope improves in future seasons.

Personnel **–** An open DAQ software developer position was filled in October 2015 by Tim Bendfelt.

Computing and Data Management

Computing Infrastructure **–** 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 2996 TB (Terabytes[1]): 1037 TB for experimental data, 1771 TB for simulation and analysis and 188 TB for user data.

Two large disk expansions were procured and brought online during 2014. The first one is a Dell Compellent system including two SC8000 RAID Controllers and five SC280 disk enclosures, providing a total usable space of 1267 TB. The second expansion includes 6 Dell PowerEdge R515 servers and 18 Dell PowerVault MD1200 disk enclosures, providing a total usable space of 960 TB. This disk expansion targets a net increase in disk space for hosting the new data generated by the experiment as well as the upgrade of those disk servers that are reaching the end of their operative life.

The deployment of these new large disk systems enables the execution of a very important milestone for the maintenance and operations of the Data Warehouse: an upgrade of the Lustre software version. Lustre is a parallel 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. The new disk servers have been configured with brand new filesystems that use the last stable version of Lustre (2.5). This has made possible a safe upgrade by means of a massive data migration from the old filesystems to the new ones. At the same time that this migration is taking place, some of the disk systems that are at the end of their operative lifetime are retired from production. Once the migration process is finalized, the production filesystems in the Data Warehouse will all be Lustre 2.5 and will provide a total of 4000 TB of usable space.

In 2012, the IceCube Collaboration decided to use direct photon propagation for the mass production of simulation data. This requires the use of GPU (graphical processing units) hardware to deliver adequate performance. The first IceCube GPU-based cluster at UW-Madison, containing 48 NVidia Tesla M2070 GPUs, was installed in early 2012. An expansion of the GPU cluster containing 32 NVidia GeForce GTX-690 and 32 AMD ATI Radeon 7970 GPUs was deployed during the second quarter of 2013. Even if the performance of these clusters was very good, they were still far from providing the total computing capacity required for IceCube simulation. The urgent need for gaining access to a larger amount of GPU resources was highlighted during the IceCube Software and Computing Advisory Panel in April 2014 as one of the higher priority resources needs to focus on.

---

[1] 1 Terabyte = 10^{12} bytes

Early in 2014 NVidia released their next generation GPU cards based on the new Maxwell architecture. The first cards released were the low-end range (GTX 750) with a relatively low core count and memory (512 cores, 1GB RAM). A series of tests and benchmarking measurements were performed using IceCube code on these GTX 750 cards and results were very promising, showing up to ~70% improvement in performance per Watt with respect to the previous generations. In September 2014 Nvidia announced the release of the high-end Maxwell GPU card GTX 980, with 2048 cores and 4 GB memory. After confirming its good performance with the IceCube benchmarks, an expansion of the GPU cluster was purchased consisting of 16 SuperMicro 4027GR-TR servers containing a total of 128 GTX 980 GPU cards. The new servers were deployed by December 2014, providing a big boost to the overall IceCube GPU compute capacity. An identical purchase of GPU cards was initiated in February 2015 with an expectation to integrate this hardware into the GPU cluster in mid to late April.

Figure 7 shows the GPU time delivered per month by the UW Madison NPX GPU cluster since January 2013. A steady increase in the overall GPU usage can be observed. The big jump in total usage at the end of the plot corresponds to the deployment of the new 128 Nvidia GTX 980 GPU cards, each of which performs twice as fast as each of the previous generation Nvidia cards.
The last two GPU expansions in 2013 and 2014 have been deployed as part of the main IceCube cluster NPX. This way, users can access both resource types (4000 cpu and 224 GPU job slots) 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.

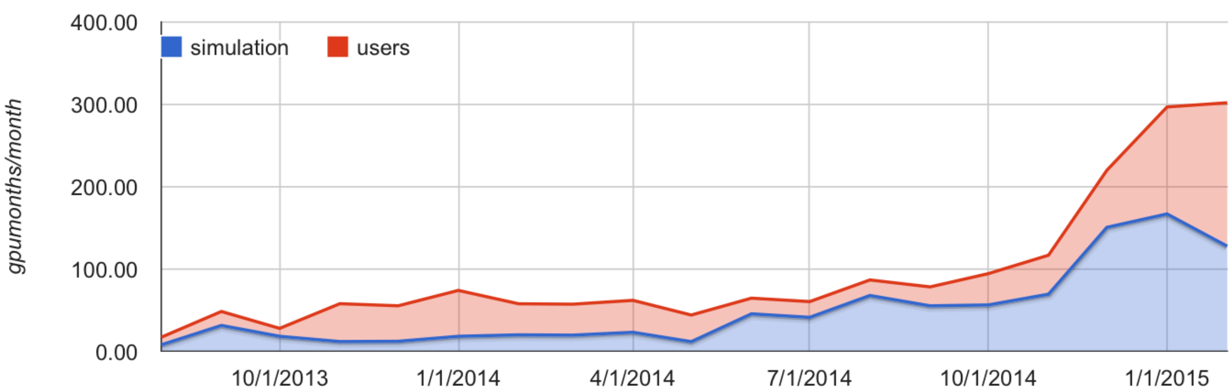

Figure 7: GPU hours consumed by the jobs completing daily in the new NPX GPU cluster. The contribution from Simulation Production jobs is shown in blue, and in red the contribution from other IceCube users.

Other sites in the collaboration are also contributing with GPU cluster time to Simulation Production (see Figure 8). More IceCube sites are currently working to deploy GPU clusters and contribute those to the central Simulation Production.

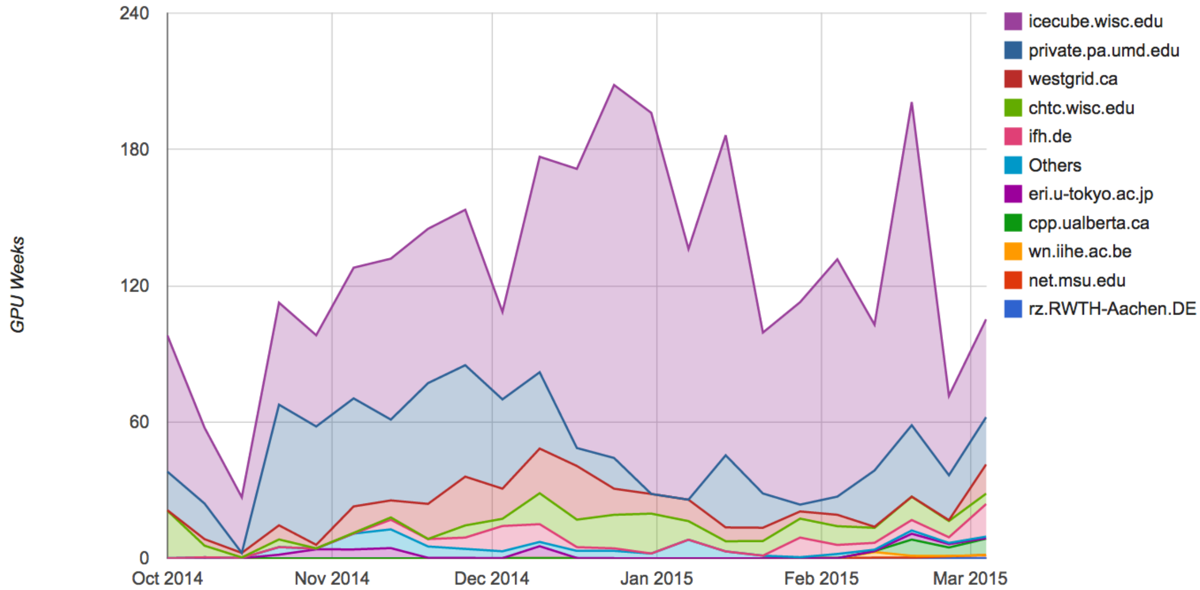

Figure 8: GPU time consumed by IceCube Simulation Production jobs in different clusters. The period displayed spans from 30-Sep-2014 until 3-Mar-2015. The plotted values correspond to GPU weeks consumed per week. The four sites with contributions above 1% are: UW-Madison (68%), UMD (19%), Westgrid Canada (13%) and DESY Zeuthen (5%).

South Pole System - All IceCube in-ice instrumentation is connected to electronics read-out systems in the IceCube Lab (ICL), which is located about 1km from the main station. These custom built systems (DOMHubs) provide power to the detector strings and handle raw data read-out. Failure of a single DOMHub results in data loss for an entire string, equaling 60 photo sensors. Remotely controlled, fully redundant power distribution units (PDU) were installed in IceCube Lab during this year's Pole season. These PDUs allow operators to power cycle stuck components from the main station without having to go out to the ICL, reducing partial detector downtime from about one hour to only a few minutes.

After successfully operating the new disk based archival test setup at the South Pole for one year, the entire data archival system was upgraded during the 2014/15 South Pole season. The unreliable tape drives were decommissioned and were replaced by a hard disk based archival system. Three new disk enclosures, each providing 12 hot-swappable 4TB hard disks, provide the capability to archive IceCube data and could also be used as an extended buffer of up to 144TB in case of unexpected satellite outages. Winter Overs on site take care of replacing filled hard disks as part of their daily routine.

The yearly system maintenance of the IceCube compute farm in the ICL was performed, including necessary firmware updates for the Dell PowerEdge compute nodes, applying OS bug fixes and installing all NSF mandated security patches.

A new infrared sky camera was installed on the roof of the ICL to study long-term cloud coverage as part of an evaluation for future experiments. Muon taggers were temporarily installed and test data recorded to study the feasibility of modifying the IceTop detector system during the 2015/16 season. A full 8h worth of unfiltered hitspool data was extracted from the readout systems. This data was transported north, where it allows software developers to simulate full 86-string real-time detector response at the test facility in Madison.

A new RUDICS based iridium system was installed in the ICL to evaluate the possibilities of transporting high-priority detector monitoring data 24/7 in near real-time. Software developments and testing are ongoing efforts, and a fully operational RUDICS system is scheduled to be commissioned during the 2015/16 season.

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, a small tape library for output, 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 and another is written to LTO5 tapes for long term archive. 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. Reprocessing was stopped for several weeks to avoid stepping on the data migration into the new Lustre 2.5 filesystems. The process was restarted early March 2015 under the new framework, which is currently processing about 1600 files/day. We expect to ramp this up to 3000 files/day in the next few days. The ideal capacity of the tape system is about 4000 files/day. Without further problems this should be able to wrap up the project before the end of the year, ideally by Fall.

Distributed Computing – One of the high level goals within the IceCube computing services is promoting the use of distributed computing (Grid) resources. Facilitating the access to these resources to all the collaboration members is one of the needed actions to get there. With this objective, a CernVM-Filesystem[2] (CVMFS) repository has been deployed at UW–Madison hosting the IceCube offline software stack, photon tables and detector conditions data needed for data simulation, reconstruction and analysis. CVMFS enables seamless access to the IceCube software from any computer anywhere in the world by means of HTTP. High scalability and performance can be accomplished by deploying a network of standard web caches. The system became the default software distribution mechanism at UW-Madison facilities in January 2014 and during the last year it was deployed in most of the other IceCube sites.

During the last quarter of 2014, the IceCube repository was consolidated into the Open Science Grid (OSG) CVMFS distribution. This allowed IceCube to make use of the content distribution network that the OSG project was rolling out in coordination with the European Grid Infrastructure (EGI) project. Through this mechanism, the IceCube software is currently being mirrored from the BNL and FNAL sites in the US, from NIKHEF and RAL in Europe and from ASGC in Taiwan. This software distribution infrastructure provides an enormous scalability and resilience for IceCube to ramp up its usage of distributed resources.

[2] http://cernvm.cern.ch/portal/filesystem

Another important aspect of 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, we make use of some of the federation technologies within HTCondor[3] such as Flocking or GlideinWMS. These allow to present users many heterogeneous clusters as if they were one big cluster with a single interface. Opportunistic resources at other UW-Madison departments and in many OSG sites across the US are accessed this way. During 2014, some IceCube sites such as DESY-ZN in Germany or SCINET in Canada were also integrated into this system so any IceCube user job can run in those resources in a completely transparent manner.

Figure 9 shows the CPU time delivered from the previously described distributed systems to IceCube jobs. A relatively constant use by the Simulation Production activities (in blue) is observed and, on top of that, the large peaks in users activity (in red) shows the potential of tapping very large amounts of resources in the Grid. In January for instance, the CPU time used by IceCube jobs in these Grid systems was equivalent to more than 5600 simultaneous jobs continuously running for one month.

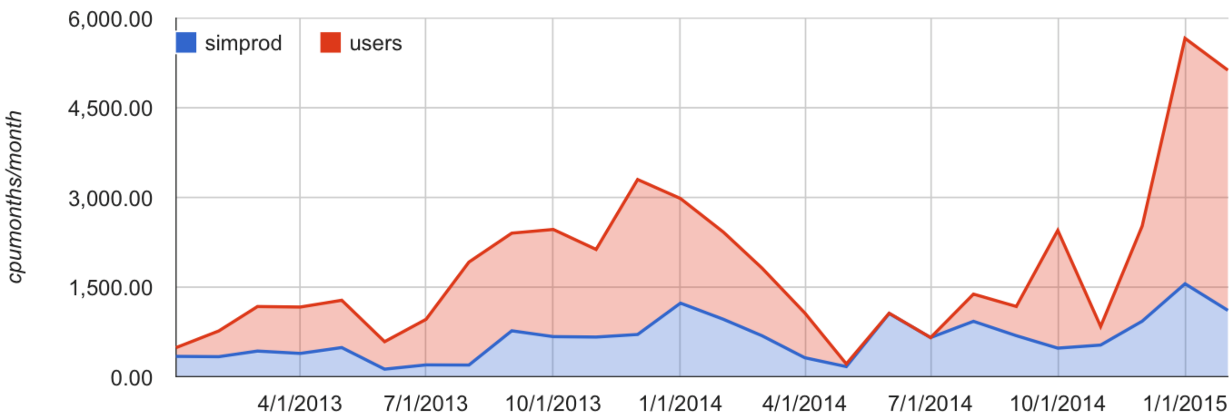

Figure 9: CPU time used by IceCube jobs in opportunistic resources accessed via the Grid. These resources are mostly clusters from other departments at UW-Madison and OSG sites elsewhere in the US. The color code indicates the amount of resources that went to Simulation Production (blue) as compared to normal users jobs (red).

BOSCO is a software project developed by OSG that is being explored in IceCube as an alternative federation technology for integrating remote clusters that do not provide any access interface other than SSH. In July 2014 a BOSCO server prototype was deployed at UW-Madison that allowed Simulation Production jobs to be run in the OSGConnect[4] infrastructure. The system has been running stably since then and has delivered more than 800.000 CPU hours so far. During the last quarter of 2014 two additional resources were integrated using BOSCO: the Jasper and Guillimin clusters from the Canadian national compute infrastructure for research. The integration of more clusters via BOSCO and testing the software full capabilities is an ongoing activity.

[3] http://research.cs.wisc.edu/htcondor/
[4] http://osgconnect.net/

One of the most critical services in a distributed computing system is the data access. In order to efficiently profit from potentially very large computing power 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 of all of the servers providing these services was upgraded to 10Gpbs during the 3rd quarter of 2014. This way, it was ensured that the total WAN data moving capacity was only limited by the WIPAC global Internet connection. This connection had been capped to 4Gbps for a long time due to a specific network module configuration. As WAN data rates kept growing the cap became more visible and, eventually, an actual bottleneck. In January 2015, this limitation was removed and WIPAC external connectivity bandwidth bumped up to 10Gbps, the current physical layer limit.

Figure 10 shows the amount of data served out of the Data Warehouse via gridftp. The lift of the maximum capacity from 4Gbps to 10Gbps becomes visible after January, when there were periods with 40 to 65 TB served in a single day.

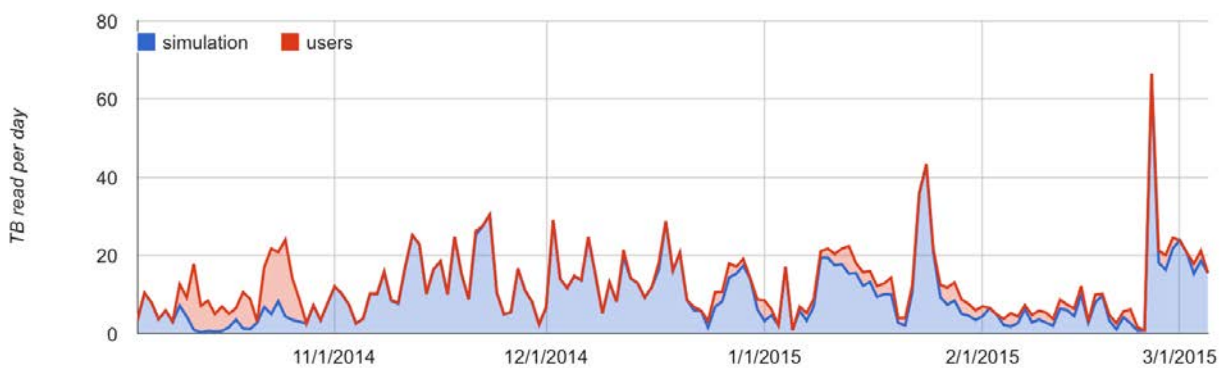

Figure 10: Amount of data in the UW-Madison Data Warehouse accessed by remote processes via gridftp.

Data Movement – Data movement has performed nominally over the past year. Figure 11 shows the daily satellite transfer rate and weekly average satellite transfer rate in GB/day through March 2015. The IC86 filtered physics data are responsible for 95% of the bandwidth usage.

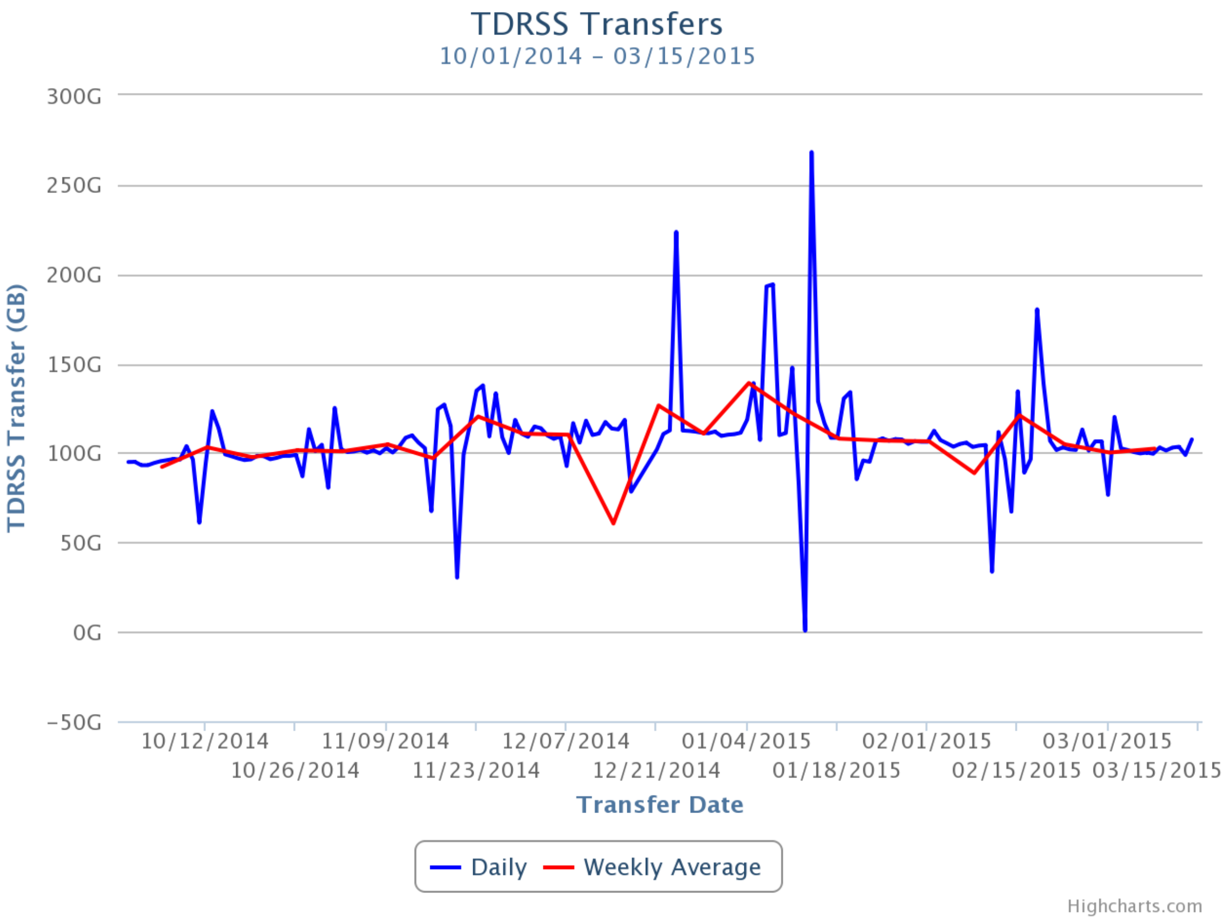

Figure 11: TDRSS Data Transfer Rates, October 1, 2014–March 15, 2015. The daily transferred volumes are shown in blue and, superimposed in red, the weekly average daily rates are also displayed.

***Data Archive* –** The IceCube raw data were archived on LTO4 data tapes until October 2014. In October 2014, archival was switched to a new disk archival system. A total of 8.8 TB of data were written to LTO4 tapes during the reporting period (October 2014 to March 2015). A total of 163.6 TB of data were archived to disk, averaging 1.0 TB/day. A total of 15.5 TB of data were sent over TDRSS, averaging 96.4 GB/day.

***Offline Data Filtering* –** Data collection for the IC86-2014 season is still ongoing, and the resources required for the offline production has been in line with the initial estimates. The complete offline production will require about 750,000 CPU hours on the NPX cluster at UW Madison. 120TB of storage is required to store both the Pole-filtered input data and the output data resulting from the offline production. Replication of all the data at the Deutsches Elektronen-Synchrotron (DESY) 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 3 main physics analysis groups in the collaboration have now adopted this framework. Processing of one complete season of data has been completed for these groups, and production of a second year of data has started. Based on these initial results, it is envisaged that all subsequent post-offline production will be done using the new approach.

The IC86-2015 season is scheduled to start at end of April 2015. In preparation for this, a compilation of processing scripts for the data is in progress. It is expected that there will be minimal difference between these and the scripts from the IC86-2014 season. Collection of 24hrs of data using the new configuration is planned for mid-April. This data will be used for script validation and benchmarking purposes. Results from the processing of the test data will inform final updates to processing scripts and form valuable inputs for production resources planning.

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, and direct generation of Level 2 simulation data is now used to reduce storage space requirements. This transition also included a move to a new release of the IceCube simulation software release, IceSim 4. IceSim 4 contains improvements to low-level DOM simulation, correlated noise generation, Earth modeling, and lepton propagation. We have progressed toward having 100% of all simulations based on direct photon propagation using GPUs or a hybrid of GPU 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 will begin transitioning to generating IC86-2014 datasets and halt production of IC86-2012, and IC86-2013 given that the detector configuration and level2 reconstructions are very similar and representative of all three years.
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; LONI – Louisiana Optical Network Infrastructure; 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; PDSF/Carver/Dirac – LBNL; and NPX – UW IceCube.

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. The following two links give more information about IceCube data formats and policies.

IceCube Open Data: http://icecube.umd.edu/PublicData/I3OpenDataFormat.html

IceCube Policy on Data Sharing:
http://icecube.umd.edu/PublicData/policy/IceCubeDataPolicy.pdf

Datasets (last release on 19 Feb 2015): http://icecube.wisc.edu/science/data
The pages below contain information about the data that were collected and links to the data files.

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

http://icecube.wisc.edu/science/data/HEnu_above1tev

IceCube Oscillations: 3 years muon neutrino disappearance data:

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

IceCube Software Coordination – The software systems spanning the IceCube Neutrino Observatory, from embedded data acquisition codes to high-level scientific data processing, benefit from concerted efforts to manage their complexity. Alex Olivas has accepted the role of IceCube Software Coordinator. In addition to providing comprehensive guidance for the development and maintenance of the software, the Software Coordinator will work in conjunction with the IceCube Coordination Committee, the IceCube Maintenance and Operations Leads, the Analysis Coordinator, and the Working Group Leads to respond to current operational and analysis needs and to plan for anticipated evolution of the IceCube software systems.

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 Geneva, Switzerland, September 15-19, 2014.

IceCube M&O - FY2015 Milestones Status:

Milestone	**Month**
Revise the Institutional Memorandum of Understanding (MOU v17.0) - SOW and Ph.D. Authors head count for the fall collaboration meeting	**September 2014**
Submit for NSF approval, a revised IceCube Maintenance and Operations Plan (M&OP).	**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
Report on Scientific Results at the Spring Collaboration Meeting	April 28-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	September 2015

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. The IceCube detector performance continues to improve as we restore individual DOMs to the array at a faster rate than problem DOMs are removed during normal operations.

Education & Outreach (E&O) - The IceCube Neutrino Observatory's education and outreach efforts continue to build on past successes while developing new opportunities to reach diverse audiences and provide intensive research experiences for undergraduates, high school students, and teachers. Ongoing collaboration-wide education and outreach efforts include:

- Reaching motivated high school students and teachers through IceCube Masterclasses
- 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)
- Supporting the IceCube Collaboration's communications needs through social media, science news, web resources, webcasts, print materials, and displays (icecube.wisc.edu)

Ten IceCube institutions will participate in the 2015 masterclass, double the number from the inaugural 2014 effort that involved 100 students.

PolarTREC teacher Armando Caussade, who deployed to the South Pole with IceCube in January 2015, kept journals and did two webcasts, one in English and one in Spanish, during his visit. His trip received international media attention, and he continues to get requests for presentations about his experience.

The four South Pole webcasts in total reached about 1,600 people, involved more than a dozen USA schools, and had audiences from Germany, India, and Italy.

Four undergraduates have been selected for NSF-IRES-supported 10-week research experiences with European Icecube collaborators at Ruhr-Universität Bochum and the Universität Mainz. An additional NSF REU grant will provide support for 18 more students to do astrophysics research over the next three summers. At least one-third of the participants for both programs will be from two-year colleges and/or underrepresented groups.

The science section of the IceCube website was updated in December 2014, and we have posted almost 40 news stories during FY15. We have added an additional 1500 followers on Facebook and Twitter, and routinely receive thousand visitors per week.

To help disseminate our programs, identify potential partners, and learn from other education and outreach programs, we have attended and presented at national meetings. One poster and one talk was given at the 2014 American Geophysical Union meeting, and classroom resources were presented during the 2015 International Teacher Scientist Partnership meeting. IceCube will also participate the workshop "*Convening on Broader Impacts and Informal Science Education*", organized by the NSF-funded Center for Advancement of Informal Science Education (CAISE), which will take place in April 2015 in Arlington, VA. The IceCube collaboration will host a joint event during the National Association for Broader Impacts meeting in April 2015 to highlight the importance of broader impact work.

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 September 2014 for the Fall collaboration meeting in Geneva, Switzerland (v17.0), and the next revision (v18.0) will be posted in April 2015 at the Spring collaboration meeting in Madison, WI.

IceCube Collaborating Institutions

Following the September 2014 Fall collaboration meeting, the Drexel University with Dr. Naoko Kurahashi Neilson as the institutional lead, and Michigan State University with Dr. Tyce DeYoung as the institutional lead, were approved as full members of the IceCube Collaboration.

As of March 2015, the IceCube Collaboration consists of 44 institutions in 12 countries (22 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/collaborators

IceCube Major Meetings and Events

IceCube Spring Collaboration Meeting – Madison, WI	April 28-May 2, 2015
IceCube IOFG Meeting – Madison, WI	May 3, 2015
IceCube Particle Astrophysics Symposium – UW–Madison	May 4-6, 2015
IceCube Fall Collaboration Meeting – Copenhagen, Denmark	October 12-16, 2015

Acronym List

CVMFS	CernVM-Filesystem
DAQ	Data Acquisition System
DOM	Digital Optical Module
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
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