Responses to “Questions to ICNO/M&O”

Questions dated March 17, 2020

The responses below address questions asked by the NSF Integrated Project Team as part of the March 17, 2020 NSF Annual Site Visit of the IceCube Neutrino Observatory (ICNO) Management and Operations (M&O).

Question #1: What does the current M&O computing infrastructure consist of? - e.g., summary of scope of activities, assumptions around labor if included and a diagram of the nodes would be helpful.

The current M&O cyberinfrastructure broadly covers: (a) the ICNO South Pole detector and computing systems for capture, processing, storage, and transmission of the IceCube experimental data; (b) central computing systems, storage and CPU and GPU processing resources, at UW-Madison; (c) distributed processing resources (collaboration, Open Science Grid (OSG)); (d) long-term archival storage at DESY-Zeuthen and NERSC (LBNL); (e) custom firmware and software for all systems; (f) the development systems supporting firmware and software development and maintenance (e.g. code repositories, continuous integration testing systems, and build systems). These elements span the WBS L2 areas 2.2 Detector Operations, 2.3 Computing and Data Management, 2.4 Data Processing, and 2.5 Software. The detailed description at Level 3 of the deliverables and associated tasks is provided in the “ICNO WBS Dictionary” in the Appendix B of this document. Figure A1 (Appendix A) shows a logical map of the global computing resources managed by ICNO M&O.

M&O CyberInfrastructure Hardware Resources and Associated Support

South Pole System (SPS)

SPS consists of the cyberinfrastructure needed to operate the detector, including the DAQ, detector monitoring system (i3Live), first-level data filtering for data transfer via satellite (PnF), data transfer and archival infrastructure, and associated database services. It comprises 45 Dell R740 servers and 40 uninterruptible power supplies (UPSes) with a 5-year lifecycle replacement, and 40 network switches. IceCube Winter-Over experiment operators and R. Auer are the primary maintainers of SPS.

South Pole Test System (SPTS)

SPTS is a scaled-down “mirror” version of SPS hosted at UW-Madison. It is used to verify changes to SPS infrastructure, e.g. master clock replacement, lifecycle burn-in testing, DAQ development, etc. It includes 36 Dell R740 servers, 8 UPSes, and 12 networking switches.

Data Warehouse

The data warehouse is the central hub for all IceCube data. All experimental, simulation, analysis, and user data, including raw data from Pole on disk, will at some point be stored in the warehouse. The majority of data are currently simulation and experimental data. The data warehouse contains 15 PB online disk storage (11 PB usable) and 6 database servers hosting over 50 databases. It is maintained by S. Barnet, V. Brik, and A. Sheperd.

High Throughput Computing Cluster

The high-throughput computing cluster is hosted at UW-Madison and provides centralized computing for the entire collaboration. It has three primary responsibilities:

1. Provide the common base of offline processing of experimental raw data to a level usable by analyzers for higher-level analysis;
2. Production of Monte Carlo simulated datasets;
3. Dedicated computing resources for use by collaborators for higher-level data processing and analysis.

The cluster includes approximately 5000 CPU and 350 GPU cores. It features a direct connection to the data warehouse storage resources to lower the need for data transfer, in particular for I/O-bound analyses. S. Barnet and A. Sheperd maintain this cluster.

Distributed Computing

Distributed computing resources are essential for data processing and simulation production. IceCube's resources are spread across allocated resources, e.g. NSF's XSEDE program, opportunistic resources, e.g. Open Science Grid, and externally dedicated resources, e.g. DESY-Zeuthen. The resources are globally distributed, see Figure X. We have developed a lightweight system--pyglidein--that allows us to tie together all these resources with proven technology coming out of OSG.

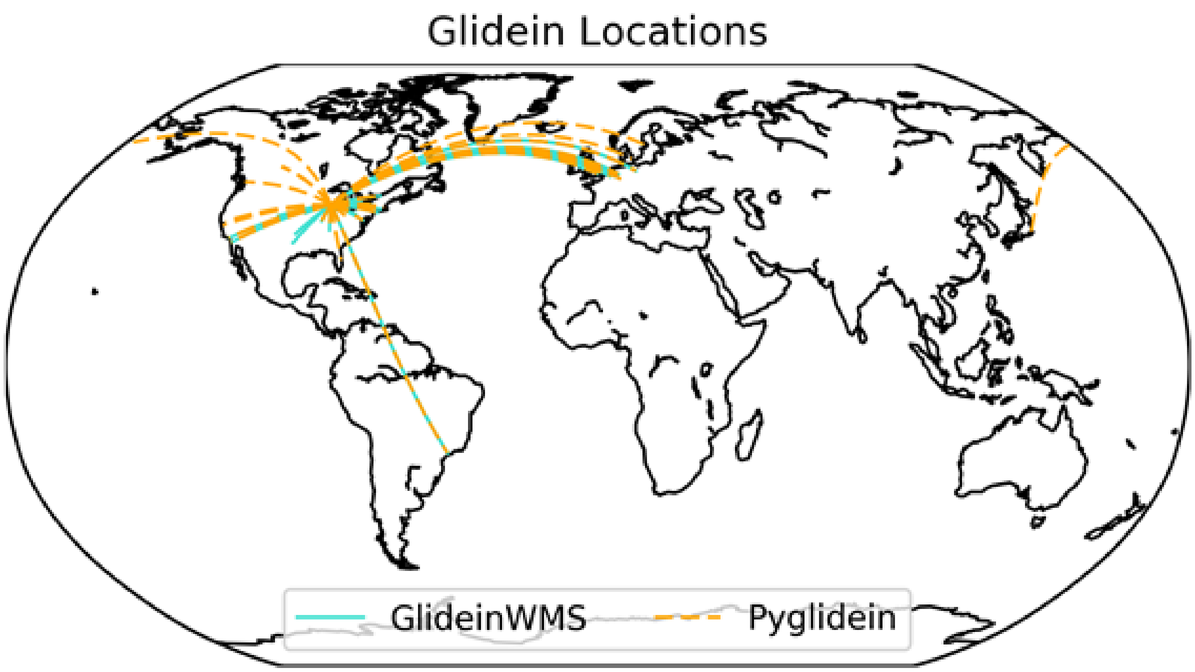

Figure 1 IceCube Distributed Computing Resources

The M&O team has created a workflow management system--IceProd. This system allows us to manage production overflows, establish data provenance, and maintain metadata for the production workflow.

We have shown that our infrastructure can readily scale to enormous scales using cloud resources. These resources are comparatively expensive, so adding them to the overall resource mix is prohibitively expensive for general purpose use. Applications that require a large number of resources for short periods can be off-loaded to the cloud, such as reconstructions used for MMA applications.

M&O CyberInfrastructure Software Resources

IceCube software is classified at the top-level into five major subsets:

> **Data Acquisition (DAQ) Software (300k lines of code):** this software is a mix of firmware and embedded software written in VHDL and C, respectively, that runs inside the DOMs; Linux kernel code running on the readout computers (DOMHubs); Java data processing and movement applications; C++-based processing application for the supernova data stream; and Python distributed control framework. With the notable exception of the supernova subsystem, this codebase is developed and maintained by dedicated software professionals with M&O core funding, 4 FTE.
>
> **Detector Systems Monitoring IceCube, Live, (100k lines of code):** Python infrastructure for data quality monitoring and database of detector and run metadata. This also includes multi-level message priority transport system. This is also maintained by professional software developers supported by M&O core, 2 FTE.
>
> **Data Movement, JADE, (30k lines of code):** Java codebase that supports bulk data transport from South Pole to the Data Warehouse in Madison, transport of data products to long-term-archival storage at NERSC, and associated metadata. A full-time developer is supported by M&O core as well as external funding, 0.5 FTE on M&O core.
>
> **Data Processing, Simulation, and Analysis Software (1 million lines of code):** this includes the core infrastructure *IceTray*, the online filters running at South Pole, reconstruction algorithms, simulation algorithms, and higher-level analysis code written ostly in C++ with Python wrappers, 7 FTE on M&O core.
>
> **Distributed and Grid MiddleWare, IceProd, (30k lines of code):** software that enables massive data processing batches to be run on distributed computing, cloud computing, and supercomputing resources. Mostly Python, 2 FTE.

We comment on the appropriateness of the development team size, roughly 15 FTE. There is no golden rule for required development effort based on high-level software code base attributes such as lines of code, rather it requires a deeper consideration of aspects such as reliability, performance, user base, interaction with specialized computing resources (e.g. GPU shader code development). From personal experience (Kael Hanson) managing the DAQ software development effort for multiple years during the first years of IceCube operations, a team size of 4 programmers was sufficient development effort to maintain the roughly 300,000 lines of DAQ code and still make progress in adding major new features such as hit spooling. The DAQ needed to be very reliable and had many specialized aspects such as interaction with the readout hardware but, compared to the million lines of code in the processing, simulation, and analysis software was extremely uniform. An increase from 15 to 20 FTE is appropriate given the scope of IceCube software.

Question #2: A picture of the current GPU demands, and the factors contributing to outpacing the supply of existing GPUs while performing strictly the M&O-related computing tasks, and estimated timeframe of when this will be evident?

Eight years ago IceCube simulation switched from using precomputed "photon tables" to simulate the detection of Cherenkov light by the optical modules to direct photon propagation made possible using the massively parallel architectures available in commercial GPU hardware. It is only by using direct photon propagation that simulation *and reconstruction* algorithms can model the known complexities of the ice optics and provide the high fidelity simulated datasets and reconstructions that maximize the overall performance of the facility for both astrophysics and neutrino physics.

Simulation is constrained by the limited availability of GPU computing infrastructure due to the requirement, for a number of scientific analyses, of enough simulated background, i.e. cosmic ray events of which there are nearly 100 billion each year. To keep pace with the data coming in, IceCube M&O computing requires approximately 2000 GPU (NVIDIA GTX1080 normalized), a factor 2.5 to 3 times what is currently available. Increasing the simulation livetime would allow us to improve the overall sensitivity of analyses.

The most recent optical models of the ice have been so complex that they can no longer be parametrized for use in reconstruction algorithms. This required the development of the first reconstruction algorithms that utilize direct photon propagation. Further development of these algorithms, in particular speeding them up, will hopefully allow us to further decrease the angular error to 0.1 degrees.

Additionally, new workflows have begun to emerge now with increased use of machine learning/artificial intelligence (ML/AI) techniques. These new workflows either require different interaction with the resources, i.e. interactive rather than batch, or have significantly higher resource requirements than traditional workflows. One particular issue with a more interactive workflow in the current scheme is data management and access. Users will need to access a large portion of the data to be able to "train" ML/AI algorithms. This will require either direct access to the data warehouse, i.e. co-location of data and compute, or a data management and access layer that users can easily interact with. These methods have the potential to revolutionize the entire analysis software stack. In particular, they have shown potential to produce better analysis samples and allow for better separation between signal and background. Having better and cleaner analysis could help IceCube produce significantly improved science results.

The core software that exploits these GPU resources, CLSim and PPC, and the ML/AI codes all require maintenance by specialized developers with knowledge of GPU programming.

Question #3: What is your plan for the computing infrastructure upgrade in the $$$ amount(s) and additional nodes to the above-mentioned diagram?

In the short-run (1-2 years), the planning is to limit updates to the compute infrastructure, i.e. CPU and GPU, as much as possible. There will be a limited investment in new architecture, such as the rumored NVIDIA "Ampere" architecture. This will be limited to an investment on the order of $25-50k over that time period. This limited investment in compute infrastructure is needed to allow for a replacement of the data warehouse, roughly $1.2M, once the current data warehouse hardware will no longer be under warranty (2023). The long-run may require more significant limits on infrastructure investments, i.e. reduced to the bare minimum replacement.

This reduced investment will have a negative effect on the rate of science results. The lack of replacement of dedicated resources will increase resource contention. This will require students to invest more time in learning how to use distributed resources rather than analyzing data.

We are investigating additional funding streams through UW and NSF. For example, we are in discussions with UW central IT to submit an MRI proposal for GPU resources. We also submitted an unsuccessful Mid-Scale Research Infrastructrure R-1 proposal for a GPU cluster with SDSC in 2019. We are planning to resubmit said proposal once the Mid-Scale program solicitations are published again. None of the additional resources are guaranteed, for example in 2019 IceCube’s XSEDE allocation was reduced. Adding resources through M&O would allow us to better exploit the science potential of IceCube data.

The software and hardware infrastructure goes hand in hand. We will need additional effort on the software side to fully utilize resources that we are allocated through programs like XSEDE. There are significant changes emerging out of international exa-scale computing efforts in terms of software interfaces that will need to be investigated and adopted to be able to use these resources.

Question #4: What is needed for fully satisfying the M&O needs in computing? - e.g., estimate all costs labor, hardware, software, and include escalation?

Software: A fully-staffed computing portfolio across WBS 2.3, 2.4, and 2.5 requires the following software engineering resources, averaged over the next five years, in addition to the current M&O profile: (a) 0.7 FTE to develop and maintain middleware for cloud computing and supercomputing resources; (b) 1.8 FTE to maintain online and offline filtering and processing; (c) 0.85 FTE to maintain simulation algorithms, not photon propagation related; (d) 0.2 FTE to support visualization tools; (e) 0.85 FTE to develop and maintain GPU codes. The costs are broken down with a 2% escalation in the Table 1 below.

Software Task	**PY1**	**PY2**	**PY3**	**PY4**	**PY5**	**TOTAL**
Cloud Programmer	0.0	44.6	178.4	178.4	178.4	579.9
Core software support	178.4	356.9	356.9	356.9	356.9	1,605.9
Simulation programmer (upgrade support)	44.6	178.4	178.4	178.4	178.4	758.4
Visualization programmer (steamshovel)	0.0	44.6	44.6	44.6	44.6	178.4
GPU programmer	89.2	178.4	178.4	178.4	178.4	803.0
					TOTAL	**3,925.6**

Table 1 Additional software development labor required for next M&O cycle, thousands of dollars.

Without these resources, the further development of data analysis and simulation software will be significantly hampered. This would result in the delay in publication of new or improved results. This is especially alarming in light of the Upgrade. Two critical pieces of software that will require additional attention are photon propagation and particle reconstruction software. The photon propagation code will require significant effort to be ready for the Upgrade. We are currently expecting delays in Monte Carlo simulation and thereby science results. Similarly, implementing new reconstruction algorithms based on the improved detector understanding gained from the Upgrade will take significantly longer to implement and thoroughly test.

Hardware: In addition to the current M&O support for computing (CPU and GPU) and storage lifecycle replacements, a fully-equipped system to satisfy M&O needs for simulation and reconstruction would contain 2000 GPUs (NVIDIA 2080Ti normalized).

Given the trend within the HPC community to add significant GPU resources to recently deployed and upcoming systems, we believe that an increase to 500 dedicated GPUs (NVIDIA 2080Ti normalized) to IceCube M&O will suffice. Overall we will be aiming to have 2000 GPUs available for IceCube across the collaboration.

Question #5: If the M&O Project couldn't hire staff (software engineers) or obtain contributed labor, what is done about that? As the facility manager, UW should be trying to address problems like that. If he didn't have budget, why has program management (UW) not shifted funds from elsewhere?

WIPAC as facility manager has been proactive in bringing additional external resources into M&O computing: we have obtained about $200k in funding to maintain software development resources and about $350k in supercomputing and cloud computing hardware resources in addition to exploiting more of the distributed computing resources available within the collaboration.

In the case where no labor was available, the issues were typically addressed in order of importance and resolution was delayed as much as possible. One result was that maintaining existing software was prioritized over working on new projects or refactoring code that may have greater long term benefit.

In asking to shift funds from other areas within M&O one must prioritize which functions of M&O are the most critical to facility operations. Clearly, the stable operation of the detector and online systems that leads to the delivery of high-quality experimental data is the most critical aspect of M&O. Loss of data in this area is unrecoverable. Detector operations has therefore received the top priority, given limited resources. M&O supported effort in calibration is less than 8 FTE. Taking even 2 FTE from this area would significantly hinder progress here and, consequently, have a large negative impact on future science. Internally software resources could be shifted from, say, middleware development to better access more computing hardware to improving the efficiency of the algorithms to better exploit the capabilities of existing hardware resources. The M&O team continually re-evaluates priorities in all these areas and within the software particularly and balances resources accordingly. Asking further support from science base grants is not feasible as none of these base grants contain any support for computing infrastructure. That is assumed to be provided by M&O.

Question #6: We suspect that what you call "a computing infrastructure" is a much larger thing than the only M&O-related needs for the ICNO data collection and processing. Are you including the needs of developing simulations (needed only for the M&O regular operation) in the "GPU demand"?

All of the computing hardware infrastructure, and the vast majority of the software infrastructure, *viz* common simulation production and processing, is considered to be within the scope of IceCube M&O. This is stipulated in the Cooperative Agreement for NSF Award PLR-1600823 paragraphs 2.5.b.2 through 2.5.b.5 which encompasses the support of development and maintenance of software, firmware, and hardware systems from the online systems at South Pole up through “reconstruction tools” and “high-level analysis tools”. As reiterated in the answer to question #7 below, the US science analysis program does not support computing individually for the groups. The non-US groups’ computing is supported through the non-US common fund cash contributions which are used to purchase computing hardware at the South Pole, in the Data Warehouse, and in the IceCube central computing cluster maintained by WIPAC. Were computing infrastructure support (i.e. batch computing and storage) not provided to individual researchers within the collaboration by IceCube M&O there would be no dedicated computing available for analysis of IceCube data.

This is not to say, however, that what we call “computing infrastructure” extends to the point where core M&O pays directly for collaboration researchers to develop analysis software. Core M&O provides the software infrastructure, IceTray, as well as providing experimental data and Monte Carlo simulation processed and filtered in a common processing for further use by all scientific collaborators, the so-called “Level 2” data stream. Often, the algorithms for this common processing are developed outside of M&O core by graduate students and postdocs who then leave the collaboration. Because the entire collaboration still relies on the continued, high-quality operation of this common processing, support must be guaranteed by M&O.

Question #7: This question is triggered by Benedikt's statement in Slide 9: "Machine learning, new reconstructions, etc. require GPUs, e.g. student estimated needing half of GPU resources of the entire collaboration for a single analysis." Do you mean here actual science analyses, or this is only for developing the M&O-related simulations? On the student's estimated need for half of the GPU resources - was that estimate validated? Are there other examples of verified need exceeding available resources? Cost estimates? Also, what does student research project have to do with M&O?

The computing model that has been followed throughout the 10+ year operations of IceCube, the model requested in the 2016 M&O proposal document (pp. 20, 22, and 24), and the model affirmed by paragraphs 2.5.b.2 through 2.5.b.5 of the Cooperative Agreement for the current operations cycle, relies on a centralized computing infrastructure supplemented by distributed resources which themselves are supported through middleware guaranteed by IceCube M&O resources. The science grants do not support computing infrastructure. Were there no M&O support for collaboration computing, there would be no coherent infrastructure to support the major investment made by the NSF and its partners to construct the facility.

The particular example referenced was meant to convey, anecdotally, the coming wave of next-generation analyses using ML/AI that are already demonstrating revolutionary improvements in performance. The existing IceCube M&O resource arbitration mechanisms, viz The IceCube Coordination Committee (ICC), has to date not approved such a significant request due to the impact on core operations. However, the anecdote serves to illustrate the underlying message that *IceCube science is currently constrained by limitations in computing resources.*

Question #8: BTW, is SCAP funded by the award, and if so, why aren't they listed in the award under governance? How much budget goes to supporting this advisory body? Where are their reports posted? NSF should be able to review their recommendations. This activity must be transparent.

The role of the SCAP (Software and Computing Advisory Panel) is described in the Cooperative Agreement under section 2.5.c, SCAP reports are posted at:

https://docushare.wipac.wisc.edu/dsweb/View/Collection-12012

and a direct link to the most recent 2018 report was provided as supplementary material to this site visit.

SCAP is not funded under the M&O award but their travel for reviews at UW-Madison is. The domestic ($1800/trip) and foreign ($3,500/trip) direct rates take into consideration airfare, transportation, lodging and per diem plus an applied escalation rate of 1.5% per year.

Question #9 (Slide 9 in C. Vakhnina's Presentation): What is the $254K overrun due to - inflation on labor or anything else?

See chart Figure 2 which shows the planned M&O core-supported labor in the final proposed budget and actual labor to-date. Labor was reduced as the project years progressed in order to hold the requested flat budget in the face of inflation.

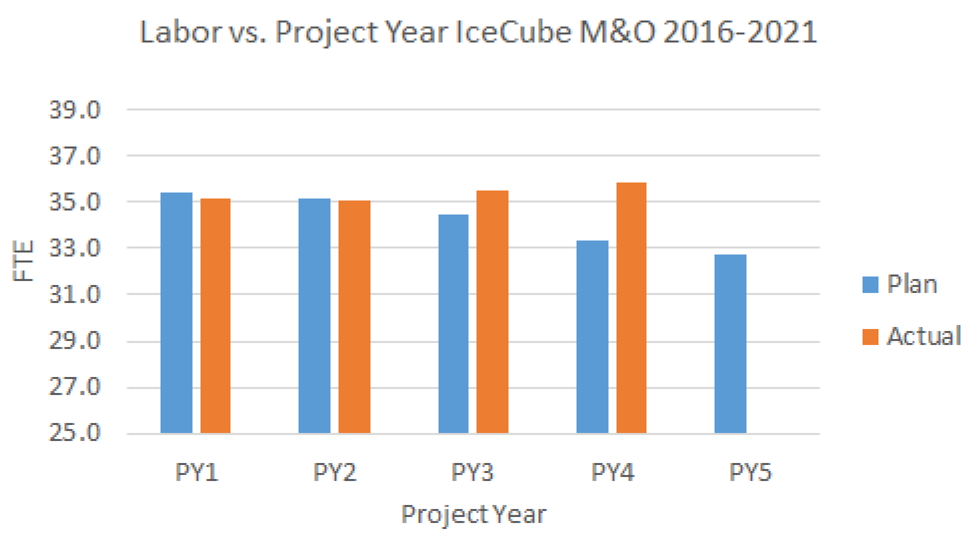

Figure 2 M&O 2016-2021 labor plan vs actual.

Question #10 (Slide 9 in C. Vakhnina’s presentation): How this overrun will be mitigated? Other than inflation, what are the reasons (and justifications) for the increase of the M&O costs in the current operation?

At a critical time when the IceCube science program is expanding frontiers in fields as diverse as multimessenger astrophysics and neutrino physics, a ramp-down of labor would retard scientific momentum. In addition, NSF’s investment in the IceCube Upgrade must be protected: to support integration of the Upgrade, to exploit the opportunity to achieve precision characterization of the ice, and to lead the field of precision oscillation physics, M&O management opted to maintain labor at PY1 and PY2 levels as evidenced by the chart in Figure 2. Any overruns will be mitigated by use of institutional funds.

Appendix A: IceCube Global Computing Map

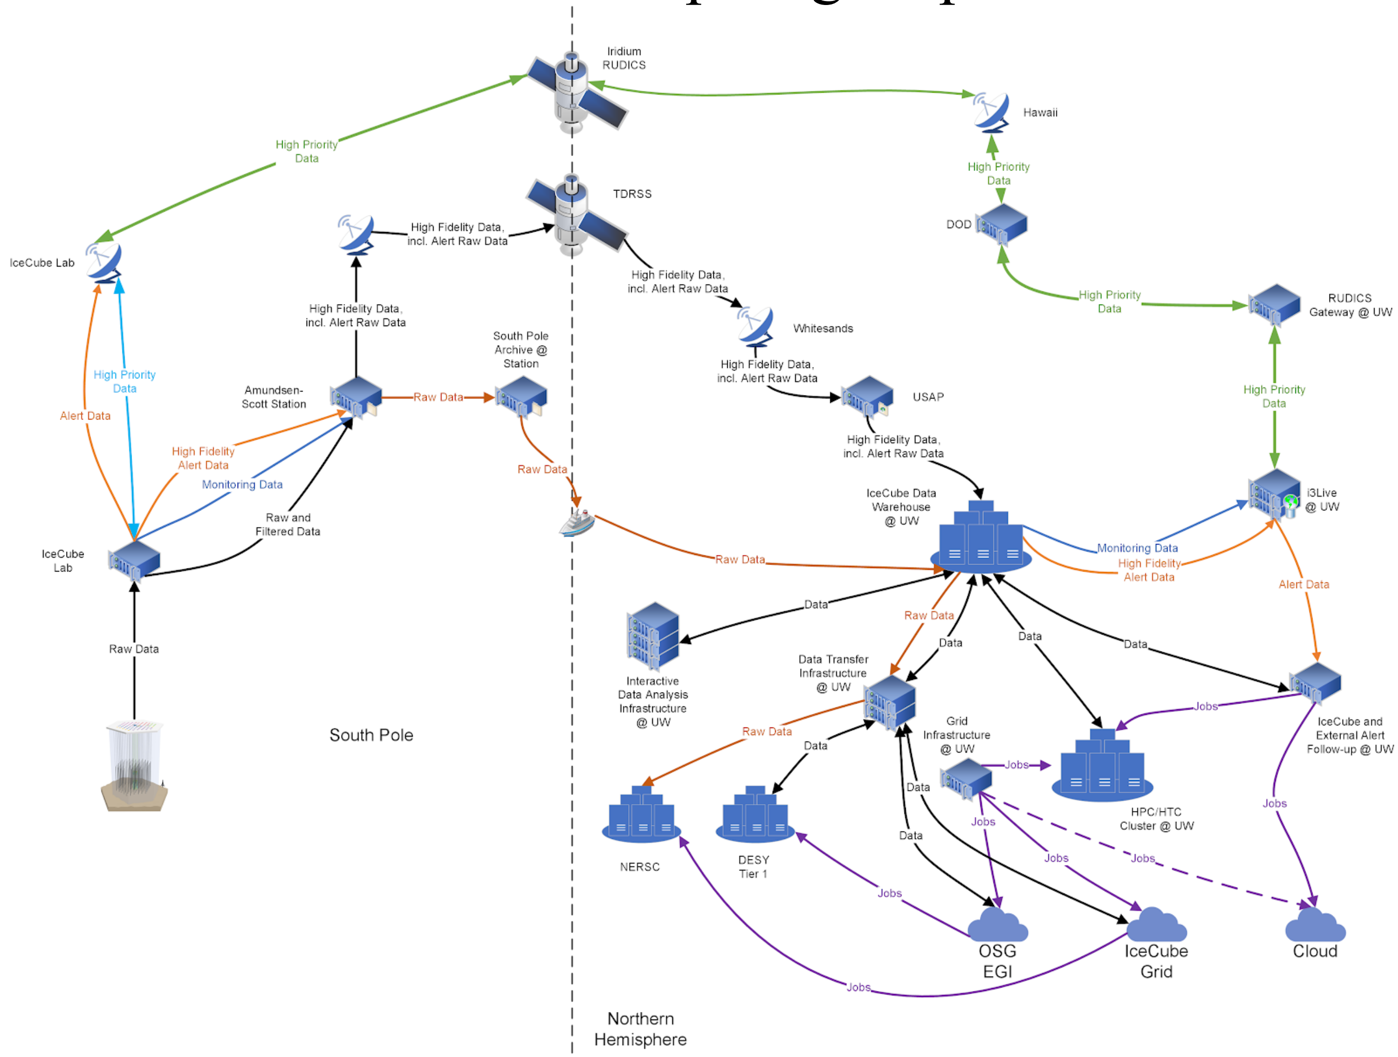

Figure A1: Map of IceCube global computing.

Appendix B: IceCube M&O WBS Dictionary

2.1 Program Coordination

Activities related to the management, administration, and systems engineering.

2.1.1 Administration

Provide oversight and direction in managing and facilitating the IceCube Management & Operations program as well as interfacing to the IceCube collaboration structure. Act as focal point for ensuring continued support for stakeholder related issues and decision processes. Provide ongoing support for daily operations and review processes to ensure accuracy of reporting data while providing leadership. Coordinate and facilitate activities related to program involving subcontract management, financial reporting, and stakeholder related issues.

2.1.2 Engineering and R&D Support

The engineering and R&D tasks required to support day-to-day operations of the detector and ensure its long-term viability. R&D supports efforts to address electromagnetic interference (EMI) in the Dark Sector, enhancements to the performance of the IceCube Laboratory electronics and computing, and the ability to interface with externally funded R&D and science activities, especially those that intend to utilize the IceCube facilities, infrastructure, or data flow.

2.1.3 USAP Support and Safety

South Pole planning which includes grantee deployments, cargo, and detailed on-ice support plans coordinated with the Antarctic Support Contractor (ASC). Gather needed information to submit the Support Information Package (SIP) to ASC / NSF each April. Safety planning to ensure IceCube personnel and equipment operate in a safe and healthful workplace. Conduct yearly training prior to Pole deployments.

2.1.4 Education and Outreach

Provide for the education and public outreach responsibilities of IceCube's missions, projects, and programs. Includes management and coordinated activities, formal education, informal education, public outreach, and website maintenance.

2.1.5 Communications

Media support and dissemination of information about IceCube and its activities through press releases, news articles, and digital channels (website and social networks) to reach audiences that are growing both in number and in diversity.

2.2 Detector Operations and Maintenance

Daily operation, control, and monitoring of the physical systems composing the detector instrument.

2.2.1 Run Coordination

Coordinate and prioritize of special detector operations, calibration runs, and online software upgrades. Ensure data quality delivered to Offline Data Production, including determination of good/bad runs and good/bad DOMs. Track detector performance measures.

2.2.2 Data Acquisition

Maintain and update the data acquisition software including DOM software, DOR card device driver, StringHub, trigger, event builder, and secondary stream builder. Maintain the HitSpool untriggered data buffering and capture system. Interface with the TFT and working groups to define and implement new triggers.

2.2.3 Online Filtering

Maintain and update the processing and filtering system software, including incorporating new filters each physics run start, and interfacing with the TFT and working groups to determine filter CPU and bandwidth constraints. Identify candidate astrophysical events for the real-time alert system.

2.2.4 Detector Monitoring

Maintain and update the IceCube Live detector monitoring system, including the real-time web page status, backend databases, and data quality monitoring subsystem. Maintain system for tracking good/bad runs and good/bad DOMs. Support real-time alert program by providing detector status and alert interfaces.

2.2.5 Experiment Control

Maintain and update the online IceCube LiveControl experiment control system. Provide low-latency communications and messaging to/from the South Pole. Maintain and update the databases and other systems providing detector metadata (e.g. calibration constants).

2.2.6 Surface Detectors

Maintain the specialized calibration and operational systems required for the IceTop surface array. Develop and deploy new scintillator-based devices to mitigate the effects of snow accumulation on the surface array efficiency.

2.2.7 Supernova System

Maintain and update the parallel data acquisition system responsible for detecting Galactic core-collapse supernova. Generate timely alerts to the northern hemisphere including follow-up systems (e.g. SNEWS).

2.2.8 Real-time Alerts

Maintain and update software systems related to generation and reception of real-time alerts other than those generated by the online supernova system.

2.3 Computing and Data Management Services

2.3.1 Data storage and transfer

Definition: Hardware and software infrastructure that enables handling the IceCube data throughout its entire lifecycle: archiving the data at the South Pole, transferring it to the North via satellite, storing it in the main Data Warehouse at UW-Madison and archiving it for long-term preservation. It includes data discovery, access and preservation.
Tasks:

- Maintain data handling software (JADE): Archive at the S. Pole, transfer, ingest to the Data Warehouse and long-term archive. - P. Meade
- Maintain data catalog. Data discovery and metadata web interface - P. Meade
- Operate data handling services - J. Bellinger
- Long term preservation and archive services. Data curation. - J. Bellinger
- Maintain and operate storage infrastructure at UW-Madison – A. Sheperd
- Maintain and operate remote data access services at UW-Madison - V. Brik

2.3.2 Core data center infrastructure

Definition: Operation of the core IT infrastructure at UW-Madison. Includes Networking, core services (email, web, etc) and custom application services.
Tasks:

- Maintain and operate core computing infrastructure at UW-Madison - S. Barnet
- Manage computing facilities at UW-Madison – A. Sheperd
- Manage networking infrastructure at UW-Madison – UW-Madison DoIT
- Cybersecurity - S. Barnet

2.3.3 Core computing resources

Definition: High throughput computing cluster at UW-Madison that provides data processing, analysis and simulation services for the IceCube collaboration.
Tasks: Maintain and operate data processing and analysis cluster - V. Brik

2.3.4 Distributed computing resources

Definition: Infrastructure that enables efficient use of heterogeneous distributed resources available for simulation, data processing and analysis. Includes software and operations of Grid services that provide IceCube users seamless access to distributed resources.
Tasks:

- Distributed resources coordination
- Maintain and operate distributed workload management infrastructure (Grid)
- Maintain and operate remote data access services at UW-Madison

2.3.5 SPS

Definition: Computing infrastructure at the South Pole, including commodity servers to host infrastructure services and online applications (DAQ, filtering, monitoring), OS and configuration for DOMhubs and networking.
Tasks:
- Maintain SPS computing infrastructure
- SPS networking and security

2.3.6 SPTS

Definition: Computing infrastructure at UW-Madison that provides a reduced size replica of the SPS.
Tasks:
- Maintain SPTS computing infrastructure
- SPTS networking and security

2.4 Data Processing and Simulation Services

2.4.1 Offline Data Production

Manage and execute production of offline data processing to generate general purpose Level 2 data (with refined event reconstruction) and science working group specific Level 3 data (with event reduction defined by channel working groups).

2.4.2 Simulation Production

Manage and execute production of simulation data to provide sufficient cosmic ray induced and neutrino induced events in the IceCube Observatory for all physics working groups. Produce simulation data at the general purpose Level 2 and provide the tools for working groups to produce Level 3 data.

2.4.3 Public Data Products

Coordinate the regular release of science level experimental data from refereed-journal published results, and the release of real-time alerts authorized by the Real Time Oversight Committee.

2.5 Software

2.5.1 Core Software

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

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

2.5.2 Simulation Software

Simulation code not related to production infrastructure: physics generators, particle propagators, detector response simulation, photon propagators.

2.5.3 Reconstruction

Event reconstruction and filtering tools of general utility.

2.5.4 Science Support Tools

This new group is a spin-off from the reconstruction group, but is still tightly coupled with the reconstruction effort, since the software will still be bundled together in releases. The purpose of this group is to manage common software tools used at filter and analysis levels beyond L2, including IceCube's open source effort.

2.5.5 Software Development Infrastructure

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

2.6 Calibration

2.6.1 Detector Calibration

Tasks related to understanding of the response of the IceCube detector: coordination of flasher runs; detector geometry; DOM response / linearity; DOM noise studies.

2.6.2 Ice Properties

Bulk and hole ice studies including fitting programs.