Report of the IceCube Software and Computing Advisory Panel

March 23-24, 2009

Stuart Anderson (chair), Sridhara Dasu, Michael Ernst, Tom Paul, John Pretz (absent Lothar Bauerdick, ex-officio)

Introduction

The first meeting of the IceCube Software and Computing Advisory Panel (SCAP) was held at Madison Wisconsin on March 23-24, 2009. The Panel heard a number of presentations regarding the current status and readiness of the IceCube Project and Collaboration IT systems to fully exploit the scientific content of the initial IceCube experimental data. The panel was asked to evaluate the following 5 areas: Computing and Data Management, UW Data Center Operations & Plans, Core Software Maintenance and Upgrades, Distributed Simulation Production, and Organization & Management.

Executive Summary

The panel was impressed with the open and honest internal assessments of the current situation and the willingness to seek and consider an independent assessment. The primary findings and recommendations of the panel are as follows:

1) The panel commends the IceCube project for accepting criticism of the original software framework and switching to a much more flexible and user friendly system

(IceTray-v2) to support the physics user community. In the experience of the panel the need to do such a full scale re-write is often necessary in large experiments.

2) It is critical that the collaboration implement a formal process for establishing computational requirements necessary to meet its scientific goals. It is clear that the lack of agreed upon requirements has prevented meaningful planning of the necessary infrastructure (CPU cycles, storage, and facilities) which has led to a number of problems, including: unnecessary stress on the IT support staff and overloaded computational resources resulting in severe system outages. This is something that must be addressed immediately to avoid repeating problems for the IC59 analysis and beyond.

3) As the IceCube experiment continues to transition from a centrally managed construction project to a distributed collaboration model it is very important that all critical tasks are clearly identified, distributed, and tracked with clear deliverables and schedules. The panel encourages the establishment of regularly updated white papers that clearly layout the scientific goals of the collaboration and the required infrastructure. This should form the basis of establishing detailed collaboration MOUs. The panel further recommends that an internal self-assessment of previous period MOU performance be made available within the management structure of the collaboration.

4) As the size and criticality of the IceCube Data Center continues to grow, the panel recommends that a search committee be formed to implement an international search for a data center operations manager with significant experience in operating facilities at this scale that support a large scientific community.

5) The panel is very concerned about the current interaction between the IT support staff and the physics user community. While this is something that is already recognized internally we would like to emphasize that it is critical to manage the expectations of users through clear communication channels, both human and technical. For example, additional monitoring facilities should be made available on-line for users to determine what services they should expect to be currently up and running from their perspective.

6) The panel strongly recommends that an open data format be established for reduced data products that will support both the long-term archiving of valuable

scientific data and enable more efficient analysis tasks that do not need to process all of the data.

7) The panel has serious concerns about the cost benefit ratio for the current disk storage solution being proposed and the projected growth needed. For example, IceCube is currently paying 4x the cost for the same disk storage units as FermiLab. This problem is aggravated by the lack of a clear end-to-end planning process that would prevent scrambling to buy individual units to put out immediate fires.

8) Given the apparent consensus that access to at least 10% of the raw data, now estimated at 2TB/day for IC86, will be required at the data center, the panel encourages the IceCube project to look into a more integrated tape management system to solve both its raw data management needs and to provide reliable tape backups.

Computing and Data Management

- *What is the current grid usage (processing and MC production)?*

Findings: Grid usage via grid tools is not significant, although grid resources are being used from the local batch farms for simulation production. There is no grid usage for experiment data processing.

Comments: Lack of grid usage is not surprising given that effort invested in grid tools is small. The local usage at sites that support grid tools, by using local batch systems directly, seems to be working fine for simulation production. The budgeted 0.2 FTE for simulation production per site is adequate, but scaling this to many sites takes too many FTEs. In the long term grid tools usage can potentially save FTEs.

Recommendations: Continue the usage as is and defer proper grid tools deployment to a later time.

- *How can we optimize the usage of computing resources (limited repetition but maximum freedom for analysis)?*

Findings: Current storage resources and planned increases are inadequate to scale up from IC40 to IC59 and IC86. The raw data includes HLC, SLC and overheads at about the same level. The overhead information includes trigger debugging information that is rarely used, and long headers and time stamps. Level-1 data is stored long term although it is fully contained in Level-2 Level-2 data per event is large, and most users need only a small fraction of it.

Comments: Before tightening the trigger to reduce potentially useful HLC and SLC data, it is useful to work on reducing the event overhead, including the trigger debugging information. Level-1 data need not be stored long term. The reduced datasets that users will inevitably make are going to repeat core information from level-2 multiple times.

Recommendations: Make tough decisions needed to trim the event overhead.

Plan to remove level-1 datasets in a timely way. Estimate cost for the CPU time incurred to recreate level-1 datasets to the storage needs and optimize the time you can afford to keep the level-1 datasets.

Create a common level-3 dataset, which contains the reconstructed event information without the expanded raw data, for bulk of analysis jobs. Replicated copies of this data can provide high throughput access to "hot data" as needed. The limited number of users who need the raw data will have better access to the level-2 data.

Decide on which level data for simulation need to be saved, by estimating the costs of CPU needed to recreate the data versus storage costs.

- *How should we ensure long-term data availability?*

Findings: Currently *boost* is used for serialization of classes. The format has two deficiencies: 1) it is not a very common format and is tied to the definition of classes as they exist in the current software and 2) the unstructured files need to be read serially even when only a small portion of the data is needed for a particular analysis job. There are plans to change the data format to use NASA HDF5 format or CERN ROOT IO.

Comments: The plans to change the data formats are sound.

Recommendations: We encourage investigation of standard public domain data formats and their early adoption. This will certainly ease long-term data availability.

- *Should we start to consider commercial vendors?*

Findings: No serious investigation of commercial vendors has begun. PDSF has been considered.

Comments: Neither the commercial vendors nor the national labs can provide resources at lower cost than what you would pay by yourself, if they fully pass on their costs.

Recommendation: Considering going to third-party vendors for computing resources only if they are willing to provide very low cost or free resources with altruistic motivation or for getting positive advertising.

- *What are the best opportunities for increasing capacity and performance?*

Findings: The costs incurred by the IceCube group for both storage and compute resources are factors of four higher than being incurred by other large installations. Some of the performance issues are being addressed by deploying new software products, reorganizing the resources on the network and targeted upgrades.

Comments: Capacity improvement requires more careful planning than was presented.

Tape capacity needs a big increase due to the increase in the raw data volume. The plan to put tapes in cabinets for archival purposes do not jive with plans to potentially reprocess up to 10% of that data. Frequent full backup of production data did not make sense to us. It is preferable to backup such data once. Easily reproduced data need not be backed up. User data backups may become a big issue in the long run.

Disk storage cost containment for purchases via more competitive bidding process, pooling together with other groups in purchases, and improved planning for deployment are needed to meet the large increase in capacity needed.

Performance increase for data throughput from storage systems can be met by (a) data format changes addressed earlier and (b) replicated files, more easily than by deploying expensive hardware. Simulation production can be offloaded to collaborating sites with increased use of grid resources.

Recently deployed Lustre storage system seems to be much improved compared to earlier product. However, Lustre needs to be further studied. Questions to be addressed include scalability of metadata server, which seems to be a single point failure as currently deployed. Does the Lustre system offer replication of files? Can it be automatic, e.g., does copy-on-write features exist? Can hot files be automatically replicated by setting a policy?

Recommendations: Explore and exploit opportunities for buying hardware at discount to increase compute systems capacity. Improve data formats, study analysis jobs usage patterns to improve the job efficiency. Dedicate more manpower to studying the newly deployed Lustre system capabilities such as redundant metadata servers, file replica mangers, etc.

UW Data Center Operations & Plans

Findings: The Data Center at UW with computing resources centered around the HPC farm of currently 300 cores (2.4 GHz AMD/Opteron machines, ~4 years old), disk storage offering a usable capacity of ~450TB and a tape library with a gross capacity of ~1.3PB is the primary resource for IC centrally managed production, simulation, analysis and on/off-site data serving. There is a nominal split as to 50% of the resources are allocated for production and 50% is allocated for analysis. Depending on operational requirements the actual allocation is changed according to the needs.

In addition to the HPC farm and storage resources there are central components that are vital to the IceCube (IC) detector and data analysis operation located at the UW data center. Examples include (but are not limited to) a database hosting information about the IC dataset inventory. To reduce the required capacity for data on expensive disk with data that is considered active but with an anticipated access profile that is characterized by low probability access is needed and small quantities (<10%) a Hierarchical Storage Management (HSM) solution is being evaluated.

The team that is running the data center comprises 8 IT professionals with an aggregate workforce of 7-8 FTEs with different levels of expertise in fields relevant to the IT needs of the collaboration. Computing in IC is based on a distributed system that, at least for simulation, is heavily reliant on resources provided by collaborating institutes. Though some of the larger institutions have made their resources "grid-enabled" IC is running centrally managed production in local mode with well identified administrators at the institutions available to support the production effort.

As to the organization and with a construction completeness level of ~70% IC is moving from the construction phase into M&O. A detailed proposal has been prepared and is ready for submission to the funding agency. While completing detector construction the collaboration is taking and processing data. Regarding the latter the goal for offline processing is to have this completed within 3 months during the construction phase and gradually reduce the time to 2 weeks towards the model of a pipeline process during subsequent detector operations.

Comments: Various presentations made to the panel clearly indicate excellent progress with detector construction at the South Pole with well defined preparation and execution phases, and with complex logistics involved when carrying out a project at this scale in such a challenging environment. While the activities at the pole are moving ahead as planned the panel felt that the collaboration struggles to clearly identify and articulate their computing needs, and manage the data center operations in a way that satisfies the needs. Regarding the latter the panel agrees with a statement made in the self assessment as to "the growth in requirements and the experience with initial operations indicates that the current technical approach is not consistent with the initial funding levels". It became apparent during several presentations and in particular during discussions that there is a lack of detailed information needed for a solid planning process in terms of resource provisioning at the required performance level. Planning on the IT side is heavily based on assumptions the IT professionals come up with and do not necessarily match the existing and upcoming requirements. The consequence is that too many changes have to be made on short notice, without proper planning and proper communication with the user base.

Referring to the presented cost profile for equipment acquired in the past and cost projections for future upgrades indicate that the IT group is not getting hardware components at a competitive price level. Widespread experience in computing in HENP has proven that compute servers and high-density storage servers can be procured for up to 1/3 of the amount that was presented. In view of disk being the major cost driver the integration of a HSM solution was presented for data that is accessed at low probability and small quantities. However, apart from the operational potential of a disk/tape based storage hierarchy and a description of some product studies that were focusing on technical integration work no details were presented as to how the HSM would fit into the IC production and analysis workflow. It is conceivable that occasional access to data on tape can be accommodated with such a hierarchy but it will only deliver the anticipated performance if the requirements such as access profile and quantities are well understood and provided the characteristics of the underlying hardware is matching the performance expectations. This needs thorough investigations by experts from relevant application fields and the IT group using real-world use cases.

The panel noted that the communication process between the physics community and the IT service providers is far from optimal. This holds in both directions. As was said above, while the group running the data center is lacking detailed requirements they are not making essential information appropriately available to individuals and groups relying on the IT infrastructure. This appears to be the case for information that relates to the status of services as they are deployed and their plans as to short and medium term modifications and upgrades. Another important aspect is that there is no clear and agreed upon definition as to what the service level is users can expect from the IT infrastructure and the group managing the services.

It appears modifications and upgrades to infrastructure components in production is done without much information and coordination with the user base, and without having the typical multi-tier deployment hierarchy in place that is used in industry and other professional data centers. Development and integration instances of vital services should be used to ensure that the quality and the performance of what was planned is achievable and to limit the risk of unexpected extended disruptions of the production instance of these services that are likely to occur if there is no integration phase preceding the migration of new components and services into production.

As it was mentioned in the findings the data center is operated by 8 individuals. Given the size and the level of complexity this appears to be about the right number. However, the panel feels that due to the lack of detailed information as to the skills of these people the panel cannot judge as to whether the positions are appropriately filled.

In summary, the IT group is largely in a re-active mode and constantly in danger of falling behind regarding all important aspects of IT provisioning. Based on their own statements they are barely able to provide the basic level of services and are unable to improve important and forward looking aspects such as participation in physics oriented and coordination meetings that allow them to better understand the physics driven requirements and to work on essential infrastructure components like monitoring and a presentation layer targeting the needs of production coordinators and the general user community as to the status of the IT infrastructure and the services they rely on.

To the panel a point of concern is the fact that there was no convincing mechanism presented that guarantees the timely delivery of services and capacities.

Recommendations: Data Center Operations with particular focus on how can reliability and performance be improved, data archiving strategies and cost:

1) Review the workforce budget and the assignment of tasks to individuals on a regular basis. Prepare realistic estimates as to the schedule of modifications to existing services and the deployment of new hardware and software components. The IT group should present their plans to the IC service coordinators and agree upon (and document) the deliverables, milestones and schedules. There are clear indications as to the need to intensify work with the collaboration, in particular with S/W, DAQ and production coordinators to ensure timely delivery of services. Apply proven industrial Quality Assurance mechanisms to improve stability and performance of IT services to provide predictable and high-quality results.

The panel noted that the position of the IT group manager needs to be filled. IC management should consider forming an international search committee to maximize the potential of finding the most qualified person for this critical position. Another important aspect IC management should take a look into is the organizational integration of the IT group, in particular with respect to chain of authority, reporting structure and user interfaces. Though this aspect was not addressed in detail during the review, the panel is under the impression that the current reporting structure is suboptimal in terms of ensuring proper integration of the IT group into the rest of the experiment's operations. The proposed matrix organization looks promising regarding the project oriented component however, as line management is, in particular during the M&O phase, equally important, IC management should also consider putting a workable solution in place for this component. The proposed organization as outlined in the diagram does not necessarily indicate that this is the case.

2) Negotiate service level agreements with the user base and deploy service monitoring and implement actions according to Service Level Agreements (SLA) to correct SLA violations, by that manage service level expectations. Present the actual status of the services to users in a comprehensive manner at a prominent place of the IC information pages. Make information as to modifications to production systems and services transparent to at least service coordinators, and eventually to the entire user community.

3) Take precautions to minimize the risk by conducting risk assessments with the results to be shared with management and user representatives. For modifications

and upgrades take a staged approach when introducing new services and new hardware components (utilize a multi-tier hierarchy w/ dev, integration and production platform).

4) Plan for and execute disaster recovery mechanisms, i.e. in the area of storage hardware and storage management solutions. The recent introduction of the Lustre based storage management solution as one of the most critical resources needs further validation to ensure proper stability and performance.

5) The installation of computing and storage hardware is close to exhausting the physical and electrical infrastructure available at the main data center (UW, 222 W Washington, 5th floor). The panel understands explorations currently underway indicate that there are at least two possible suitable locations allowing the ramp-up of capacities as they will be needed. While this satisfies a mandatory prerequisite for expansion the IT group should prepare a detailed plan including multi-year cost projections for each of the alternatives. Again, this needs reliable capacity projections for the next ~5 years, information that seems currently not to be available at the required level of accuracy.

6) As to the archiving strategy and in view of disk being the major cost driver the integration of a HSM solution was presented for data that is accessed at low probability and small quantities. However, apart from the operational potential of a disk/tape based storage hierarchy and a description of some product studies that were focusing on technical integration work no details were presented as to how the HSM would fit into the IC production and analysis workflow. It is conceivable that occasional access to data on tape can be accommodated with such a hierarchy but it will only deliver the anticipated performance if the requirements such as access profile and quantities are well understood and provided the characteristics of the underlying hardware is matching the performance expectations. The panel recommends thorough investigations to be carried out by experts from relevant application fields and the IT group using real-world use cases.

7) In view of unnecessarily high cost of compute and storage server hardware the IT group should make use of free market principals to lower the capital investment volume. In some cases the same products are used in quantities by other DOE and NSF sponsored programs/institutions, while the latter pay a lot less for the same or equivalent component. The IT group should plan on having frequent interactions with other sizable data centers (e.g. Fermilab, DESY) to exchange information on procurement, and resource and service provisioning.

Core Software Maintenance and Upgrades

Question 1: How do we define core software?

Findings: The IceCube core software encompasses a diverse collection of applications, including DAQ, calibration software, online processing and filtering (PnF), database systems, simulation production tools, SPADE, the IceCube Live monitoring and control system, the IceTray analysis framework, and the simulation and reconstruction software.

Comments: Though the extent of the core software systems is substantial, it appears to the panel that, overall, the systems designs have benefited from collaboration between physicists and software professionals. At this point, most of the core software appears to have achieved a reasonably mature state.

In the case of the offline software framework, IceTray, we note that a final design was achieved only after a first discarded attempt, after which more completely understood requirements were employed to arrive at a superior implementation.

Question 2: How do we justify our scripting and programming languages?

Findings: The core software applications employ quite a wide variety of technologies, including C++, java, python, MySQL, Django and others, as well a number auxiliary libraries such as boost and ROOT.

Comments: It appears that technology choices were made to some extent on an ad hoc basis, taking account not only the essential needs of the projects at hand but also the preferences and expertise of the developers who happened to be involved.

Though quite diverse, the technologies employed do appear to be sound at least in terms of maturity and prospects for longevity.

Question 3: How do we maintain heterogeneous software systems?

Findings: As noted in Questions 1 and 2, the core software encompasses a wide range of application areas, cuts across several computing environments from data acquisition to mass production, and further employs a number of different technologies.

Comments: Maintenance of such a heterogeneous software system will likely require retention of software experts, either dedicated software professional or software-savvy physicists with broad skill-sets. The M&O proposal should take this into account. The panel was pleased to see that several of the core software components, including the PnF, monitoring and verification, simulation and reconstruction, are all based upon the IceTray framework, providing some underlying commonality. IceTray thus constitutes a crucial centerpiece of the software effort, and long-term maintenance of it needs to be ensured.

The panel was also heartened to see that efforts are underway to render the user-side of the offline software palatable to non-C++ experts.

Question 4: What are the opportunities for improvement?

Findings: The panel was made aware of problems in maintaining code orphaned by transient personnel, and of a general failure to produce high quality and up-to-date documentation, both at the detailed technical level at a level appropriate for end-users.

The panel was, however, pleased to learn that yearly "bootcamps" are held to bring new members up to speed in using the analysis tools.

Overall recommendations: It appears that the diverse applications and technologies comprising the core software systems will require ongoing experienced manpower to provide support during steady state operations. The panel encourages the Collaboration to take account of this in preparing its M&O proposal, to retain an appropriate level of software professionals, and to negotiate detailed MoU's which can ensure software commitments from groups with manpower possessing the appropriate skill sets.

The identified problems of documentation and orphaned code may, in particular, be able to benefit from an MoU mechanism which strictly holds institutes to account. For instance, beginner-level documentation and software "bootcamps" need not pose an undue burden on the software experts; experienced postdocs and advanced students can take on the role of organizing documentation and providing instruction in order to fulfill MoU commitments.

Similarly, the panel suggests employing an enforceable MoU mechanism to require institutes whose members contribute code to the software base to arrange for long-term support and maintenance of their software contributions.

Distributed Simulation Production

What is the scope of the simulation production task?

Findings: There are two interrelated tasks here. The first is *simulation development*, the generation of a reliable and accurate simulation software chain. This task involves a cycle of the verification of simulated datasets, identification of discrepancies, and resolution of these discrepancies. The second task, the *simulation generation*, is the actual usage of that chain to produce datasets for analysis. The simulation generation task is well in hand and the focus of the collaboration's work should be on developing a function simulation development cycle that works well with the distributed-developer approach the collaboration has taken.

Comments: In the simulation generation task, the IceCube collaboration has been successful, building a distributed system where individual institutions contribute processing nodes, which is controllable from a central location. This system has been exercised successfully and the issues are understood. In the M&O phase of the experiment, the simulation generation will largely be supported by M&O funds, and the simulation development will be distributed across participating institutions as part of their base grants. The close association of simulation development with scientific analysis drives the assignment of simulation development to the base grants. The collaboration identified the distributed development as a cause for concern because the collaboration has had some difficulty in getting delivery of simulation components.

How is the distributed simulation production organized?

Findings: The organization of the simulation development effort is not clear. The collaboration has admitted problems with getting institutions to deliver what they have promised. This lack of organization was cited by the collaboration as a cause for delay in simulation production (a summer of no simulation production, for instance). The software development to address discrepancies appears slowed due to manpower limitations and insufficient coordination.

On the other hand, the simulation generation effort appears to be organized sufficiently. Institutions have made their CPU contributions and the whole distributed system has been utilized successfully. Issues appear isolated and under control.

Comments: One cause for concern is that the decision of what datasets to generate is not formal. The production coordinator makes these decisions in informal consultation with the working groups, but this mechanism appears unsustainable as the competition for resources increases. Furthermore, simulation requests are reconciled to hardware realities in an informal way. This will likely become a problem as the demand for computing resources ramps up.

Finally, the simulation production has suffered particularly during the switch from IBRIX to LUSTRE. The collaboration anticipates that this is an isolated incident, but it underscores the need for close communication between the working groups, the simulation production manager, and the data warehouse personnel. Some of the tension between scientists and data warehouse personnel may be alleviated by better communication.

Recommendations:

1.) The collaboration should consider forming a ‘resources’ committee made up of working group heads, simulation production personnel, offline filtering personnel, analysis personnel, and datacenter professionals. This committee should have the charge of allocating computing and disk space resources. This committee should share responsibility with non-simulation needs (level 2 processing, analysis processing) in order to allow the competition for resources to take place in a formal way.

How is quality ensured across the distributed effort?

Findings: The collaboration is working on techniques to streamline the validation of simulation by applying the verification software to the simulated data sets. This is in addition to existing tracking that is part of the simulation production. The final evaluation is performed by the working groups.

What are the opportunities for improvement?

Recommendations:

1.) The collaboration should formalize the simulation development process.

a. The collaboration should follow its plan to identify software responsibilities and guaranteed CPU allocations specifically in an MOU.

b. The collaboration management, with input from the working groups, should conduct yearly assessment of the MOU with the participating institutions to establish that commitments have been met. These assessments should be saved and the collaboration should consider using these assessments when institution base grants are submitted.

c. The collaboration should adopt a formal sign-off procedure by working groups before significant computing resources are dedicated to simulation. The working groups should conduct initial validation of small datasets. This sign-off may be an appropriate charge for the 'resources' committee.

2.) The simulation production group should work to define their computing and storage requirements to facilitate planning by the data warehouse.

3.) The collaboration has employed software professionals successfully in the framework design and maintenance. The collaboration should consider software experts to help with the large memory footprint issue in simulation.

Organization & Management

Are the requirements and the scope adequately defined?

Findings: Based on the presentations made the computational requirements are currently uncertain by factors of a few. It is clear that the lack of agreed upon requirements has prevented long-term planning of the necessary infrastructure (CPU cycles, storage, and facilities) which has led to a number of problems, including: unnecessary stress on the IT support staff and overloaded computational resources resulting in severe system outages. For example, the recent major instability of the data center appears to have been due to running the storage system at 100% capacity due to inadequate planning.

Comments: Long-term planning of facilities, storage, and CPU-cycles are all strongly dependent on developing a comprehensive requirements document that is agreed to by the full collaboration.

Recommendations: A formal computational requirements document should be created and regularly updated. This document should be changed based on previous experience, but there should be a formal process to agree to and accept changes. This document may also be useful for the collaboration to sign up for computational tasks in their MOU's.

Is the proposed organization appropriate?

Findings: From the presentations and many of the issues discussed above it is apparent that filling the open position for the Data Center IT Manager is a critical role that should be given a high priority.

Comments: In general the IceCube organization appears to be well positioned to continue its transition from a construction project to an ongoing operations phase.

Recommendations: As mentioned in the Distributed Simulation Production section the collaboration could benefit from having a resource allocation committee.

The panel encourages the establishment of regularly updated white papers that clearly layout the scientific goals of the collaboration and the required IT infrastructure. This should form the basis of establishing detailed collaboration MOUs. The panel further recommends that an internal self-assessment of previous period MOU performance be made available within the management structure of the collaboration.

Are the resources planned consistent with the requirements?

Findings: The resource acquisition and allocation plans presented where largely disconnected from the requirements and where to a large extent reactive. For example, the disk storage system had serious stability problems when it was overloaded, backup recoveries where problematic when it was found they where needed, and the cost estimate for additional CPU cycles was solely based on what a typical cluster costs in the past without reference to the current requirements for IC-59 and beyond.

Comments: The lack of planning based on realistic and agreed to requirements is resulting in system instability and unnecessary stress between the resource providers and consumers.

Recommendations: As discussed in several specific facets above IceCube needs to establish a systematic way to obtain and manage its computational requirements. From a management perspective this is an important criteria in hiring the UW Data Center Manager, i.e., selecting someone who recognizes the need, and will insist on the information needed to plan ahead to provide the necessary IT resources to maximize the science output of IceCube.