Report from the 6th meeting of the IceCube Software and Computing Panel (SCAP) on June 4-5, 2018

L. Bauerdick (remote), P. Couvares, S. Dasu, M. Delfino (chair) and M. Livny (ex-officio)

General Comments

The charge to the committee, the membership with their affiliations and the agenda of the meeting are given in Appendix 1.

The panel thanks the speakers for their thoroughly prepared presentations, which kept well to a very tight schedule. The panel appreciated the mix of presentations and of their content, giving status and general background information as well as information specific to the panel's charge. The panel thanks all participants for the openness of the discussions and the overnight preparation of answers to the panel's questions.

The IceCube collaboration continues to pioneer the field of neutrino astronomy. Efforts in instrumentation, data acquisition and real-time alerts, operations, analysis and simulation, software, and data processing and computing have been consolidated, enabling IceCube to produce superb scientific results. This is illustrated by the latest results, including multi-messenger observations with other telescopes, which the panel greatly appreciated being directly informed about.

As in previous occasions, the panel recognizes the unique culture of the collaboration and how it has been facilitating the ongoing stream of scientific discoveries. However, the panel continues to be concerned that some elements of this culture may hinder or negatively impact the ability of the computing enterprise to support the foreseen evolution of IceCube's science.

The panel acknowledges that its present charge reflects some of these previously expressed concerns, allowing it to make more focused recommendations on these points.

The panel maintains its previous global recommendation that more attention and effort should be devoted to developing and defining goals, predicting impact, and tracing progress and trends in all aspects of the computing enterprise, including effective and efficient coordination with the analysis/science activities, realistic accounting for how human resources are utilized, and tracking of data processing and computing resource consumption together with its impact on the science mission.

Findings, comments and recommendations according to the Charge

Charge 1:
Review on the effectiveness of M&O teams to deliver software and computing services to the collaboration based on the priorities established with science working groups.

Substantial progress has been made on a number of important issues:

- Improvements at the South Pole that enable a prompt alert system.
- Successful migration from 222 W. Washington to a hosting facility.
- Successful ramp-up of a worldwide distributed computing and data transfer environment. This has been helped by a clever and pragmatic approach through *pyglidein*.

- Consolidation of the NERSC and DESY Tape Archives.
- Considerable use of OSG and XSEDE allocation and promising tests on supercomputers.
- Successful massive reprocessing campaign, with indirect benefits and savings in simulation.
- Successful deployment of IceProd2, albeit without multi-user authentication/authorization infrastructure (AAI).

Advances bring benefits, but also opportunities and challenges:

- Users have come to expect the ability to “log-in” to their experiment's “system”. This applies to job-control tools like IceProd, data repositories, software repositories, etc. Such a “log-in” should ideally be based on a “single-sign-on” system.
- Job control becomes more complex (see remarks under charge 3).
- Finding specific data elements, particularly by non-experts, needs to be facilitated. Data management and metadata-related activities will become more and more important.

The natural dynamics in the scientific domain also generate quasi-continuous challenges:

- Software management, quality control and performance improvements continue to be a challenge.
- Refocusing the M&O software team on a few specific items to be solved in a defined period of time is a correct measure, especially considering the limited human resources available. There is evidence of some early successes of this strategy, but it needs to be brought to a steady-state operation, which in turn requires the identification and prioritization of items to be worked on.

The working environment of a large collaborative endeavor such as IceCube requires looking for mechanisms which stimulate and reward the involvement of individuals for the overall benefit of the project:

- The introduction of an “Impact Award” is considered a good idea.
- The panel believes it must be complemented by broader management actions.

In summary, it is satisfactory that a steady course on strategies in the software and computing domain, set forth two or even four years ago, has been maintained in order to achieve results that positively impact science activities. At the same time, the necessary flexibility is being maintained in order to take into account emerging needs, such as requirements on public data releases.

Having said this, the panel maintains the concern, already expressed in the previous two meetings, that some organizational structures and some resource identification and management methods are not yielding the necessary results.

The considerations above bring the focus on:

- the organization of the Computing activities
- the coordination mechanisms between the Science and the Computing activities, in particular the ICC, and
- the resource coordination mechanisms between the WIPAC M&O team and the IceCube collaboration at large.

The panel learned with interest about the re-organization of the analysis Working Groups and the categorization into Science and Technical types. This re-organization is, however, too recent for the panel to evaluate its impact, although there are indications that it is positive. In particular, the panel notes that the Analysis Coordinator can now maintain a global vision and efficiently and effectively synthesize key information about needs and priorities and represent the domain in coordination bodies.

The lack of improvement in the ability of the ICC to generate appropriate coordination results may indicate, however, that there are additional organizational deficiencies.

The panel is quite concerned about the organizational structure in the M&O Software and Computing domain. The panel's concerns have one common thread: the lack of a manager with enough available time to maintain and execute an overall, global vision of this area. This lack of global vision manifests itself in various manners, for example:

- Lack of an efficient and effective representation in coordination bodies
- Inability to express the work being done in architectural terms, rather than technological terms.
- Difficulties in generating important common threads across subdomains or activities.

(Example: Generating a coherent metadata architecture spanning from DAQ to Data Releases.)

Recommendations:

2018-1 IceCube high level management is strongly urged to review the organizational structure of the M&O software and computing domain. In particular, management is invited to name a manager dedicated to maintain and execute an overall, global vision of this area (a “Global Computing Coordinator” reporting to the IceCube Neutrino Observatory Director of Operations.)

2018-2 IceCube high level management is strongly urged to review the functions of the ICC committee pertaining to the coordination between computing and analysis. In particular, management is invited to define a small, dedicated coordination group, co-chaired by the Global Computing Coordinator and the Analysis Coordinator, which would be responsible for preparing issues and recommend actions and priorities, to be brought for review and approval by the ICC.

2018-3 IceCube high level management is strongly urged to empower the Global Computing Coordinator with the ability to dialog with responsibles of collaborating institutions in order to harvest additional resources for the software and computing domain. This should be done through the definition of specific work items to be accomplished within well defined periods of time, in a spirit similar – but not limited to – the way the Software area has been recently refocused.

2018-4 The IceProd2 team is urged to complete the implementation of multi-user features, which should enable the inclusion of collaboration-wide activities. A wider use of an IceCube AAI framework should be considered, preferably in the direction of a single-sign-on for as many IceCube resources and services as possible.

2018-5 Architectural and Technical frameworks should be defined to consolidate data management and metadata related activities. Evaluation of software products to support the architectural framework should be performed as focused, time limited activities.

2018-6 The focused action scheme of the Software area should be brought into steady-state operation.

Charge 2:
Evaluate our new procedures to establish an internal policy of science result reproducibility and on public data release.

The panel heard with interest the steps being taken towards ensuring science result reproducibility and the related issue of public data releases. The panel recognizes that there are complex issues involved, ranging from technical to policies. Nevertheless, it is essential that IceCube develops these techniques. There is a lot of activity in this domain in the scientific community, and IceCube is invited to get informed, participate and profit from it.

Recommendations:

2018-7 Science reproducibility and public data releases should be considered different aspects of a more global Data Management and Preservation framework. An end-to-end architecture, from DAQ to public data releases, should be arrived at, possibly in incremental steps which are coherently orchestrated. This should be coordinated with metadata-related activities.

2018-8 IceCube high level management should take note that fully establishing and maintaining these policies requires non-negligible human resources which are currently not identified.

Charge 3:
Evaluate our expansion of computing capabilities to reach IceCube science goals.

As mentioned previously, substantial progress has been made in deploying a working distributed computing environment. A variety of resources are now contributing to IceCube, including OSG, XSEDE and the TITAN supercomputer. This success is due in part to the pragmatic use of *pyglidein*. GPU computing continues to be very important, and has grown by a factor of 2 since 2014.

The organic expansion in the resources available has also revealed some difficulties. One clear example is the I/O overload on the centralized disk server used to store intermediate results from the components of workflows. Another is the need for more detailed monitoring information for identifying reasons for job failure and as an enabler for more sophisticated job scheduling.

Opening up the environment to more self-service use is a positive step. But again, it requires improvements in monitoring and scheduling (per task, per Working Group, per RAM requirements, etc.)

Given the projections showing that the computational needs of IceCube science will outstrip projected supply, the panel sees an opportunity for IceCube to

- better measure and understand its current resource supply and usage
- ensure more consistent and efficient use of all existing resources
- aggressively seek significantly more computing resources to enable additional science (e.g., 10 times the presently used resources)
- prepare for additional resources by systematically understanding and fixing any scalability issues today (e.g., i/o, memory) that would prevent efficient use of such additional resources, keeping in mind that they may become available without much prior notice.

Efforts on improving code efficiency are important, but IceCube's involvement, and particularly that of the M&O team, must be carefully gauged against the available human resources.

Recommendations:

2018-9	Efforts should continue in a highly focused manner in order to maintain workflows which can run efficiently on systems where IceCube can request resources. This requires work on the workflows themselves, but also on monitoring and job scheduling and on the handling of intermediate results.
2018-10	IceCube is encouraged to request time from research computing providers, such as OSG, XSEDE, supercomputers, etc. with a target of achieving within 18 months a computing power about 10 times higher than currently available. This should be done through WIPAC and through IceCube collaborating institutions around the globe, and the applications should receive maximum support from their principal investigators. The M&O team should be the catalyst that puts these resources to the best use for IceCube. This requires an immediate corresponding effort to scale the computing and data analysis infrastructure to be able to efficiently and robustly handle such an order of magnitude increase in resources.

Appendix 1

Charge to Committee

1. Review on the effectiveness of M&O teams to deliver software and computing services to the collaboration based on the priorities established with science working groups.

   IceCube science is proceeding at a fast pace and data analyses have shifted towards some degree of consolidation. Science working groups have been restructured to reflect the current analysis operation scheme. All operation activities that rely on in-kind contributions (software, calibration, simulation production) are being revised to properly address the communication between science working groups and operation teams. The aim is to gather adequate computing resource requirements that reflect the true needs of the analyses.

2. Evaluate our new procedures to establish an internal policy of science result reproducibility and on public data release.

   IceCube is about to start implementing an internal procedure aimed to fully document the details of all data analyses that are published in refereed journals. Such internal data repository is designed to make science results reproducible and to be the source of information for the IceCube public data releases.

3. Evaluate our expansion of computing capabilities to reach IceCube science goals.

   IceCube is expanding the distributed computing infrastructure with resources pledged by the collaborating institutions, supercomputers and commercial cloud providers. The tests are promising and are going to lead us towards new opportunities, addressing issues of adaptive and fast response to the IceCube science goals.

Current Membership

- Lothar Bauerdick, FNAL, CMS;
- Peter Couvares, Caltech, LIGO Laboratory;
- Manuel Delfino, Port d'Informació Científica (PIC), (chair);
- Sridhara Dasu, University of Wisconsin, CMS;
- Miron Livny (ex-officio), University of Wisconsin, OSG.

Agenda

Monday 4 June - Supernova Conference Room

Time	Title	Speaker
09:00	Introduction and Project Status	K. Hanson
09:15	Collaboration Report	D. Grant
09:30	IceCube Science - the New Era of MMA	F. Halzen
10:00	Computing Services and Infrastructure	G. Merino
10:50	South Pole Computing / DAQ	J. Kelley
11:05	The Production Framework: IceProd	D. Schultz
11:20	Data Management Tools	P. Meade
11:35	Pass II Processing	J. C. Diaz-Velez
11:50	IceCube Realtime Alerts: Status and Plans	E. Blaufuss
13:00	Internal Data Releases	P. Desiati
13:25	Public Data Releases	E. Blaufuss
13:50	IceCube Analysis Needs	D. Williams
15:00	Simulation Software	A. Olivas
15:40	Simulation Production: Resource Estimation	J. C. Diaz-Velez
16:20	Coordination Between IceCube M&O and Science	P. Desiati
17:00	Closed session panel	
17:45	Open session - questions / requests for info	

Tuesday 5 June, 2018 - Supernova Conference Room

8:30	Responses to questions for SCAP and discussion.
10:30	Panel report writing and preparation for closeout.
13:00	Close out