May 7, 2013

Al Kellie, Director
Computational and Information Systems Laboratory (CISL)
National Center for Atmospheric Research (NCAR)
Boulder, CO 80307-3000

Dear Al:

This is the customary summary letter to you after the latest meeting of the CISL High performance computing Advisory Panel (CHAP). Our meeting took place in the VisLab of the NCAR Mesa Laboratory on May 2, 2013. As usual, we heard presentations from Anke Kamrath, Dave Hart, and Rich Loft on their respective areas of managing operations, computer usage, and technology development. We also heard a presentation by John Dennis on one major Accelerated Scientific Development (ASD) project that used one third of the 72,288 cores on the Yellowstone computer.

Overall, the Yellowstone computing system at the NCAR-Wyoming Supercomputing Center (NWSC) has been delivering a mammoth amount of successful computing to the NWSC users, beginning with ASD projects and now transitioning more generally to all the user communities of the NWSC. Inevitable glitches in both hardware and software implementation of such a powerful new system are being dealt with by CISL personnel in concert with vendors such as IBM and Mellanox. Most of these glitches are expected to be solved over periods of days to weeks.

One quite serious problem is in regard to cables in the Mellanox interconnect fabric of Yellowstone. Some of the original cables in this first to be delivered FDR14 InfiniBand Switch were found to be defective. About 150 of these cables were identified and replaced recently, but continuing failures of cables at a rate of about 10 per week are indicative of a more general problem. The effect on system performance is quite deleterious, particularly for jobs requiring more than 1000 cores. This problem threatens the large core count experiments on Yellowstone as future applications are readied to advance to higher core counts.

Effectively dealing with the continuing cable failures of the Mellanox fabric may require a wholesale replacement of over 2000 optical cables. CISL is currently evaluating a determination of root causes as well as drafting plans for an expeditious as possible replacement process, including round the clock teams of installers. The panel felt this was an important consideration, as the replacement would likely entail an unprecedented downtime. Some panel members expressed a preferred time window around the beginning of the CESM Breckenridge workshop.

The CHAP feels that the total replacement of the relevant Mellanox cables should be undertaken as a resolution of last resort, if it is determined to be an absolute necessity, and should be accomplished with a minimum of downtime. Any significant long-term impact on Yellowstone's performance and throughput has a serious impact on the investment in Yellowstone and the science it would produce, and CHAP supports whatever steps are necessary on the part of UCAR

and NCAR in dealing with the vendors of the Yellowstone system regarding the cable issue.

The CHAP noted that, with such a huge increase in computing and data manipulation and storage capacities, users are on a learning curve with respect to understanding how best to use the NWSC facilities. The CHAP encourages CISL to continue its proactive outreach efforts to support users in this regard, and to continue its evaluation of and experimentation with more effective ways to help users manage their workflow and data flow to optimize scientifically productive use of the resources.

The CHAP began its executive session with CISL upper management and Sarah Ruth of the National Science Foundation by getting a briefing from Al Kellie on NCAR and CISL budget prospects for the rest of FY 2013 and beyond. Nothing had been decided, but a number of high-level discussions were in progress on how to deal with the impacts of sequestration among other things. Further information would be forthcoming and made public in the weeks to come.

The CHAP then began its usual task of reviewing University computing requests totaling some 47 million core-hours. This amount was somewhat less than the total available core-hours, as most ASD users and some initial users having 200,000 core-hours without CHAP review were not submitting requests for this allocation cycle. A significant number of requests lacked adequate detail or had confusing information, such that only partial allocations were made to them. A total of 35 million core-hours were recommended for approval to the CISL Director. The CHAP expects CISL to reach out to requesting scientists and provide advice and feedback in order to help improve the overall quality of their future requests. This improvement will be important as the available resources once again become oversubscribed, and it will provide the panel with appropriate justifications for their desired allocations.

The next meeting of the CHAP will be on October 17, 2013. The CHAP looks forward to hearing about an optimally performing Yellowstone system at that time; and they also hope to review a significantly better formulated set of computing requests for some outstanding scientific projects to be accomplished on Yellowstone.

Respectfully submitted,

Bert Semtner
CHAP Chair

cc: Anke Kamrath, Rich Loft, Dave Hart, Maura Hagan, Sarah Ruth, CHAP members