

Spring 2018 Habanero Operating Committee Minutes  
March 6, 2018

**Attendees:** Kyle Mandli (Chair), George Garrett, Sander Antoniadis, CJ Wright, Marley Bauce, Michael Weisner, Arie Zask, Tomas Rube, Ryan Abernathy, Jochen Weber, Zack Fuller, Marc Spiegelman, Tian Zheng.

Kyle Mandli, Chair of the Habanero Operating Committee, opens the meeting by welcoming members to the Spring 2018 meeting, and turns the floor over to George Garrett, Manager of Research Computing Services within CUIT.

### **Naming New Cluster**

SRCPAC's RFP and Design Subcommittee is underway to develop specifications and select vendors for the high-performance computing cluster that will succeed Habanero and will be set for installation come Fall 2018. More information about the corresponding buy-in process will be announced via the SRCPAC ListServ in April.

The Habanero Operating Committee is requesting suggestions for new names for the HPC cluster. Names should be one-word verbs (i.e. people, places, or things). With naming suggestions, please email [srcpac@columbia.edu](mailto:srcpac@columbia.edu).

### **Habanero Expansion Update**

There are four ways to participate in the University's high-performance computing initiative: Purchasing nodes, Renting nodes, the Free Tier (with limited access), and the Education Tier (supported by the Engineering School and Arts and Sciences).

The first Habanero round launched in 2016 with 222 nodes. In 2017, we expanded to add 80 new nodes (including 58 standard servers, 9 high memory servers, and 12 GPUs). The current cluster has 2 submit nodes, 2 data transfer nodes, and 2 management nodes.

The Visualization server is underutilized. It has remote GUI access to Habanero storage, with the same configuration as GPU node. NICE Desktop Cloud Visualization Software.

In Spring 2018, Habanero will expand its storage because researchers have begun running out of space. This will bring storage available to 740TB. Any researchers who purchased storage in this round and need a quota increase before the new drives arrive, please email [hpc-support@columbia.edu](mailto:hpc-support@columbia.edu).

During the last downtime, Slurm was updated to Version 17, which is more efficient and has multiple bug fixes. Some of the intermittent errors have appeared less frequently.

A new Test Queue has been added – a high-priority queue with a 4-hour maximum wall time, with a maximum of two jobs per user. The settings are currently at up to 24 cores, although this can be reduced per the Operating Committee's decision.

Jupyterhub and Docker (for complicated software stacks, such as LIGO, which have different library dependencies not compatible with our standard set-up) are being tested. Users interested in participating in this testing should email [hpc-support@columbia.edu](mailto:hpc-support@columbia.edu). Ryan suggests having Jupyterlab as another option on the web interface, and otherwise seems very happy with the current set-up. Once the service is finished with testing, information about accessing will be posted to the website.

Yeti Round 1 was retired in November 2017, with Yeti Round 2 scheduled to retire in March 2019. The RFP process has initiated for the next cluster, which will live in the Morningside Data Center; this will be announced via the SRCPAC ListServ within the next month, with the purchase round in mid-May to mid-June.

### **Overview of Business Rules**

Any special usage requests should be sent to [hpc-support@columbia.edu](mailto:hpc-support@columbia.edu) and RCS will contact Kyle Mandli as necessary for approval.

Nodes owned by individual users have the fewest restrictions, with priority access for the dedicated node owners. Vice versa, nodes owned by other accounts carry the most restrictions.

For Public nodes (approximately 25), there are few restrictions, but nobody has priority.

The maximum wall time is 5 days on nodes owned by your group or on public nodes.

Every job is assigned a priority: “Fair Share” is determined by target share (determined by the number of nodes owned by each account) and recent use (number of core hours used recently; calculated at the individual and group level). If recent use is less than the target share, the job priority increases; users using more of their target share see their priority decrease.

The Operating Committee requests an open data approach to the scheduler, whereby all users can view aggregate data regarding other jobs in the queue. George understands this request, and will endeavor to provide more visibility regarding jobs.

### **Support**

The RCS team offers a number of support methods, including open office hours and group information sessions (more information can be found on the corresponding presentation slides). Senior Statistics Department students have been teaching the junior students on using HPC, which is why RCS has not visited their Department recently.

Ryan notes that his group members have reported intermittent problems with nodes dying and jobs being terminated because of it. Ryan asks whether this issue is on RCS’ radar, and what can be done to alleviate this issue in the future; his students have expressed frustration. Sander responds that some of the issues may have been alleviated after upgrades applied during the previous downtime approximately three weeks ago. Performance failures can be difficult to isolate and RCS is now benchmarking performance once nodes restart to ensure they are operating to consistent standards; therefore, these incidents should have only been temporary, and Ryan’s students should not experience these same performance issues in the future.

Ryan then asks whether there are persistent long-term challenges discussed by RCS that the larger Operating Committee would be interested in hearing. Sander responds that rebooted nodes may not always perform up to historic or consistent standards, and RCS is not yet certain why, and cannot always isolate why this happens. Ryan suggests more communication with users about *why* nodes are failing so they understand the problems, as more information is always better than less.

Marc then requests more improved diagnostics to get better visibility into what is happening on nodes so that users can better understand node health; Sander clarifies this diagnostic information would allow users to ascertain general server health, not individual node health.

Ryan asks for quantifications during future meetings of support tickets (to “peer behind the curtain” to understand the process) regarding how many tickets were opened in the semester, the average time to have them closed, the number currently outstanding, etc. another attendee suggests a chat room for users to chat with one another if they are using the same node. Ryan clarifies that this would help with efficiency in case RCS answers one question, and the same question is asked by another user, we can avoid redundancy by directing users to previously-answered questions. George suggests developing a FAQ on the RCS website.

### **Cluster Usage** (As of March 1, 2018)

44 research groups, including 1,080 users (this has grown from 740 at the Fall 2017 Operating Committee meeting: 60 new users per month). There are 7 Renters, 63 Free Tier Users, 9 courses using the Education Tier since its launch, and 2.097M jobs completed so far. Over 2M of these jobs are 1-49 cores, and 555 required 1,000+ cores, meaning most jobs are relatively small.

When the new Habanero cluster launched, usage vastly increased, which coincides with Phase 1 Yeti retirement. 174,000 core hours are theoretically possible per day, and we have come close to this maximum a few times so far. Astronomy is, once again, the heaviest users of the Habanero machine.

Tian expresses concern about open access to other group’s historical use of the job scheduler. RCS will evaluate how and whether this can be addressed and whether groups could perhaps opt-in or opt-out of access to this data.

### **HPC Publications Reporting**

Every year, Chris Marianetti, Chair of SRCPAC, presents an annual report to the Research Computing Executive Committee; this presentation always includes quantifications of publications emerging out of HPC-hosted research, and the Committee greatly appreciates this tally, which motivates research leadership to continue their commitments to research computing at Columbia. Therefore, all users of the Habanero and Yeti clusters are strongly requested to provide citation information for any published manuscripts. To report this information, please email [srcpac@columbia.edu](mailto:srcpac@columbia.edu).

### **General Discussion**

Kyle requests a “cook book” of different scripts or tools to help users easily and quickly determine disk quotas, and making them public so that those responsible for different users can monitor disk usage of different groups and of the cluster overall, so that the PIs can take appropriate action.