

Minutes, Spring 2017 Habanero Operating Committee

Wednesday, January 25, 2017

Attendees: Kyle Mandli, Tian Zheng, Danny Siegel, Brent Stockwell, Lorenzo Sironi, Raj Bose, Harmen Bussemaker, Stefano Fusi, Ingrid Richter, Halayn Hescocock, Rob Lane, George Garrett, Marley Bauce

Kyle Mandli opens the Committee meeting by setting its task: To ensure that all 32 members get a fair share of usage time with the machine, and that any concerns are posed to the Research Computing Services team.

Rob begins presenting on the overview of operation statistics, stating eventually he will cover the current and potential business rules of usage.

System Overview

There are four major components of the machine: execute nodes, head nodes, storage, and the network.

There are 176 standard nodes, 32 high memory nodes, and 14 GPU nodes, adding up to 222 nodes total across four racks, with integrated power and cooling, and located in the bottom of the new Jerome L. Greene Building. These were primarily purchased by 32 user groups, with 25 nodes purchased by central university funds. For a further description of the execute nodes, please view the presentation slides.

There are 2 Submit (where you launch your jobs from), 2 Data Transfer (for data getting in and out of the system, for unmanaged data transfer), and 2 Management nodes (separated users from scheduling software, unlike for Yeti, which makes management functions much easier for the RCS team).

The storage system is comprised of 407TB, which is very fast – a DataDirect Networks system.

For Network, Habanero has an EDR Infiniband that runs at 96 gigabytes/second (compared to Yeti's FDR Infiniband of 54 gigabytes/second).

In addition to 222 execute nodes, we have an extra node (so in fact 223), which will be used as a visualization server, which will be launched hopefully by February 2017, with a range of associated email announcements. It will have the same configuration as the GPU node. RCS has not yet decided what softwares will be used here, though likely will include MatLab and Schrodinger. It will have a web interface called Nice Visualization Software. Harmen suggests adding R Studio as well.

Business Rules

This is a service run for the University with investments from groups sitting around the committee table, and it is incumbent to establish rules that work for everyone. The original business rules were copied from those set forth by the Yeti Operating Committee, and can be changed as we see fit.

One major change is that Yeti operates on Torque, while Habanero runs on Slurm, which require that rules are slightly different in order to adjust to the new platform.

Each account has access to three types of execute nodes: Nodes owned by the account, nodes owned by other accounts, and public nodes.

For nodes owned by the account, there are the fewest comparative restrictions, and your group has priority access (since they are your nodes). Tian asks whether there are also physical nodes for Yeti, to which Rob replies affirmatively.

For nodes owned by other accounts, there are the most comparative restrictions, with priority access saved for those node owners. If anyone else is in the queue that owns that particular node, they jump ahead of you; your job cannot last longer than 12 hours. Rob explains that if you own a node, you do not want for it to be tied up for a long amount of time, and sees 12 hours as a reasonable maximum wait time. Lorenzo asks whether it is possible that a larger job would get scheduled ahead of a smaller job that had been in the queue longer because he has seen what looks like this behavior himself. Rob believes this should not happen, though recognizes that scheduling gets confusing easily. If any users experience troubles with scheduling, they should email hpc-support@columbia.edu.

For public nodes, this is a combination of the first two categories: There are few access restrictions, but no users have high priority access, meaning all users have equal chances for access using the standard job priority method.

To further explain the 12 Hour Rule: If your job asks for 12 hours of walltime or less, it can run on any node. However, if your job asks for an excess of 12 hours of walltime, it can only run on nodes owned by its own account, or on the public nodes. However, nodes owned by other researchers will make up the vast majority of the cluster. In other words, jobs under 12 hours have access to a massive computational resource, while jobs longer than 12 hours will have access to significantly less of the machine. Jobs could potentially run up to 5 full days' time. Brent suggests a longer maximum walltime. Rob clarifies that Yeti also had a 5 day maximum walltime. If any users require a longer job, this is a decision that must be made by the Chair of the Habanero Operating Committee, who can authorize in the moment or defer for larger group debate via the Habanero Executive Committee. Historically, most user requests to exceed the 5-day walltime are not considered "a big deal."

Jobs are assigned to one or more "partitions" (or queues), with each account/group having two partitions. One gives users priority to their own nodes, and one that gives lesser access to public nodes. There is a third "shared partition" for all users to run very short jobs. For further details regarding job partitions, please see the corresponding presentation slide.

RCS has attempted generosity to start: If jobs are less than 12 hours, users can have a maximum of 100 nodes. If a job lasts between 12 hours and 5 days, users have access to 50 nodes (so long jobs have half the computational power of short jobs). Halayn clarifies that these nodes are per *person*, not per *group*. One practical consideration is that the largest group in the cluster bought 48 nodes, which may require special rules to be set up eventually depending on how the system is used.

Every job is assigned a priority, with priority determined by target share and recent use. Members of the same group are treated exactly the same. Recent use is measured by the number of core hours used "recently" (this term "recently" can be discussed and redefined); recent use is calculated at both the group and user level. The half-life weight is currently set at two weeks' time. This is a major difference from Yeti, as usage history only goes back for one week. Recent use counts for more than past overall use when determining priority. Lorenzo asks whether if the job is asking for an entire node but only using 2 cores, whether users will be charged for the full node or only 2 cores; Rob responds that users would be charged only for the cores that they request. Rob also explains that users will gain access to the number of cores that they ask for – so users cannot ask for one core and then use 24, as this would slow performance for all users.

If a user's recent use is less than the target share, the job's priority increases; but, conversely, if the recent use is more than the target share, the user's job priority decreases. This is recalculated every scheduling iteration, which occurs every 30 seconds – 1 minute.

Rob concludes that Habanero has been busy, but not yet busy with a backlog. Brent asks whether having another rule could allow users to do longer jobs on only their own node; Rob believes this is doable via a third partition with a longer walltime on the group's own nodes. Harmen believes the ability to go onto the public queue is the only reason to limit jobs to 5 days. Rob believes this should be a shared system, meaning users should not monopolize their own nodes, and that nodes should be made available to their own users. That said, researchers running long jobs risk losing what they have if the job is forced to end early (or if there is a power issue). Rob believes that the Executive Committee should vote on what walltime lengths are permissible under normal conditions.

Lorenzo asks whether it is possible to dedicate one or two public nodes to a development partition or "test queue." RCS has not yet done this, although other University clusters do have this, allowing researchers to do test jobs with high-priority access but very tight walltime restrictions and resources available.

Support Services

If users are not sure where to start, email hpc-support@columbia.edu.

User documentation can be found at <http://hpc.cc.columbia.edu> then "HPC Support."

RCS holds Office Hours on the first Monday of every month from 3-5pm in the Science & Engineering Library (4th floor, Northwest Corner Building).

A Habanero information session will happen next Tuesday, January 31st from 1-3pm in the Science & Engineering Library. This is a repeat of the December presentation, so no need to attend a second time.

RCS will do Group Information Sessions where they visit interested labs, groups, departments, schools, etc. Topics can be general and introductory or tailored to the specific group's needs. Interested parties can email hpc-support@columbia.edu to request visits.

Benchmarks

High Performance LINPACK (HPL) measures compute performance and is used to build the TOP500 list. See the corresponding presentation slide for more information. Testing ran nearly across the entire cluster (across 219 of the 222 nodes). RCS also measured communications operations for a range of message sizes.

To see precise testing specifics, please view the corresponding presentation slide.

George clarifies that the aforementioned tests do not include GPUs. With GPUs, the theoretical maximum is 269 teraflops for the entire system.

Rob shows a visualization of usage activity since the machine launched in November. The theoretical maximum is 127k. Habanero is far more homogeneous compared to Yeti, which means we should be able to get far closer to the theoretical maximum.

Other Issues

Kyle opens the floor for general group discussion.

Lorenzo asks about the plan for opening to new buy-in. Rob says yes, and we need to figure out when to schedule this round, which is dependent upon perceived demand. RCS will need to be fairly confident that there will be interest, noting that two years ago we delayed the process due to insufficient interest. Timing could be anytime from a few months from now to later in 2017; Halayn clarifies this would be an expansion round which would not require a very large buy-in. It would be adding-onto the existing cluster, and would try to keep the cluster as homogenous possible.

Tian suggests it would be beneficial for the Habanero Operating Committee to coordinate across user groups, especially in light of ongoing faculty recruitment efforts and how to time recruitment with the buy-ins. RCS is regularly called to speak with faculty recruits to ensure they understand the resources available to them. Halayn reminds that SRCPAC has a formal policy on this, which ultimately asks current faculty to tell us immediately if they are bringing a new faculty aboard so that we can plan.

Kyle suggests we proceed with the current business rules, and we can review them more carefully at the Fall 2017 Habanero Operating Committee meeting (which should be scheduled in July 2017).

Kyle closes the meeting until we meet next in six months.