



COLUMBIA UNIVERSITY
Information Technology

Habanero Operating Committee

Spring 2018 Meeting

March 6, 2018

Meeting Called By: **Kyle Mandli**, Chair



COLUMBIA|RESEARCH

Shared Research Computing
Policy Advisory Committee

Introduction

George Garrett

Manager, Research Computing Services

shinobu@columbia.edu

The HPC Support Team

Research Computing Services

hpc-support@columbia.edu

Agenda

1. Habanero Expansion Update
2. Storage Expansion
3. Additional Updates
4. Business Rules
5. Support Services
6. Current Usage
7. HPC Publications Reporting
8. Feedback

Habanero



Habanero - Ways to Participate

Four Ways to Participate

1. Purchase
2. Rent
3. Free Tier
4. Education Tier



Habanero Expansion Update

Habanero HPC Cluster

- 1st Round Launched in 2016 with 222 nodes (5328 cores)
- Expansion nodes went live on December 1st, 2017
 - Added 80 more nodes (1920 cores)
 - 12 New Research Groups onboarded
- Total: 302 nodes (7248 cores) after expansion

Habanero Expansion Equipment

- 80 nodes (1920 cores)
 - Same CPUs (24 cores per server)
 - 58 Standard servers (128 GB)
 - 9 High Memory servers (512 GB)
 - 13 GPU servers each with 2 x Nvidia P100 modules
- 240 TB additional storage purchased



Compute Nodes - Types (Post-Expansion)

Type	Quantity
Standard	234
High Memory	41
GPU Servers	27
Total	302

Head Nodes

2 Submit nodes

- Submit jobs to compute nodes

2 Data Transfer nodes (10 Gb)


- scp, rdist, Globus

2 Management nodes

- Bright Cluster Manager, Slurm

HPC - Visualization Server

- Remote GUI access to Habanero storage
- Reduce need to download data
- Same configuration as GPU node (2 x K80)
- NICE Desktop Cloud Visualization software



Habanero Storage Expansion (Spring 2018)

- Researchers purchased around **100 TB** additional storage
- Placing order with vendor (DDN) in March
- Install new drives after purchasing process completes
- Total Habanero storage after expansion: 740 TB

Contact us if you need quota increase prior to equipment delivery.

Additional Updates

- **Scheduler upgrade**
 - Slurm 16.05 to 17.2
 - More efficient
 - Bug fixes
- **New [test queue](#) added**
 - High priority queue dedicated to interactive testing
 - 4 hour max walltime
 - Max 2 jobs per user
- **Jupyterhub and Docker being piloted**
 - Contact us if interested in testing

Additional Updates (Continued)

- **Yeti cluster updates**
 - Yeti round 1 was retired in November 2017
 - Yeti round 2 slated for retirement in March 2019

- **New HPC cluster**
 - RFP process
 - Purchase round to commence in late Spring 2018

Business Rules

- Business rules set by Habanero Operating Committee
- Any rules that require revision can be adjusted
- If you have special requests, i.e. longer walltime or temporary bump in priority or resources, contact us and we will raise with the Habanero OC chair as needed

Nodes

For each account there are three types of execute nodes

1. Nodes owned by the account
2. Nodes owned by other accounts
3. Public nodes

Nodes

1. Nodes owned by the account
 - Fewest restrictions
 - Priority access for node owners

Nodes

2. Nodes owned by other accounts
 - Most restrictions
 - Priority access for node owners

Nodes

3. Public nodes

- Few restrictions
- No priority access

Public nodes: 25 total (3 GPU, 3 High Mem, 19 Standard)

Job wall time limits

- Your maximum wall time is **5 days** on nodes your group owns and on public nodes
- Your maximum wall time on other group's nodes is **12 hours**

12 Hour Rule

- If your job asks for 12 hours of walltime or less, it can run on any node
- If your job asks for more than 12 hours of walltime, it can only run on nodes owned by its own account or public nodes

Fair share

- Every job is assigned a priority
- Two most important factors in priority
 1. Target share
 2. Recent use

Target Share

- Determined by number of nodes owned by account
- All members of account have same target share

Recent Use

- Number of cores*hours used "recently"
- Calculated at group and user level
- Recent use counts for more than past use
- Half-life weight currently set to two weeks

Job Priority

- If recent use is less than target share, job priority goes up
- If recent use is more than target share, job priority goes down
- Recalculated every scheduling iteration

Business Rules

Questions regarding business rules?

Support Services

Email support

hpc-support@columbia.edu

User Documentation

- hpc.cc.columbia.edu
- Click on "Habanero Documentation"
- <https://confluence.columbia.edu/confluence/display/rcs/Habanero+HPC+Cluster+User+Documentation>

Office Hours

HPC support staff are available to answer your Habanero questions in person on the first Monday of every month.

Where: Science & Engineering Library, NWC Building

When: 3-5 pm first Monday of the month

RSVP is required: <https://goo.gl/forms/v2EViPPUEXxTRMTX2>

Group Information Sessions

HPC support staff can come and talk to your group

Topics can be general and introductory or tailored to your group.

Contact [hpc-support](#) to discuss setting up a session.

Support Services

Questions regarding support services?

Cluster Usage (As of 03/01/2018)


- 44 Groups
- 1080 Users
- 7 Renters
- 63 Free tier users
- Education tier
 - 9 courses since launch
 - 5 courses in Spring 2018
- 2,097,172 Jobs Completed

Job Size

Cores	1 - 49 cores	50 - 249 cores	250 - 499 cores	500 - 999 cores	1000+ cores
Jobs	2,088,654	5,894	1,590	479	555

Cluster Usage in Core Hours

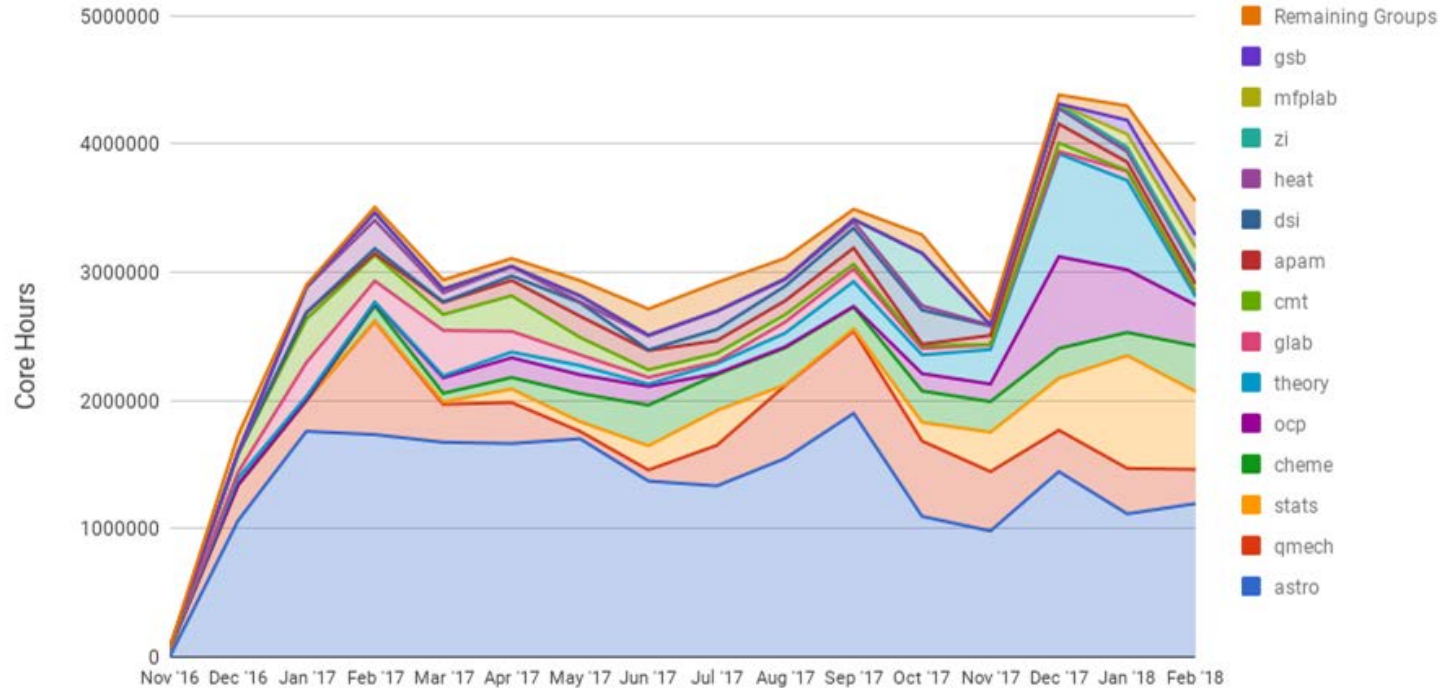
Habanero Usage



Group Utilization

Monthly Group Utilization

Core Hours



HPC Publications Reporting

- Research conducted on the Habanero, Yeti, and/or Hotfoot machines has led to over 100 peer-reviewed publications in top-tier research journals.
- To report new publications utilizing one or more of these machines, please email **srcpac@columbia.edu**

Feedback?

Any feedback about your experience with Habanero?

End of Slides

Questions?

User support: hpc-support@columbia.edu