iCER SERVICE REPORT

APR 2017

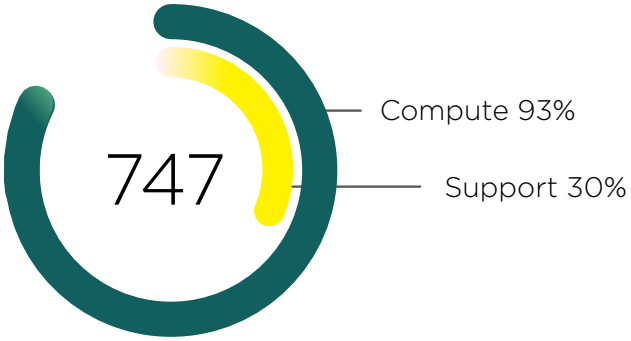

RESEARCHERS USED ICER SERVICES

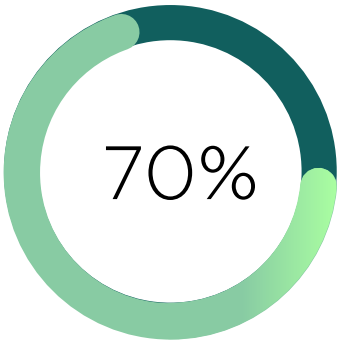

PERCENTAGE OF BUY-IN ON LACONIA

NUMBER OF USERS USING iCER SERVICES IN APRIL

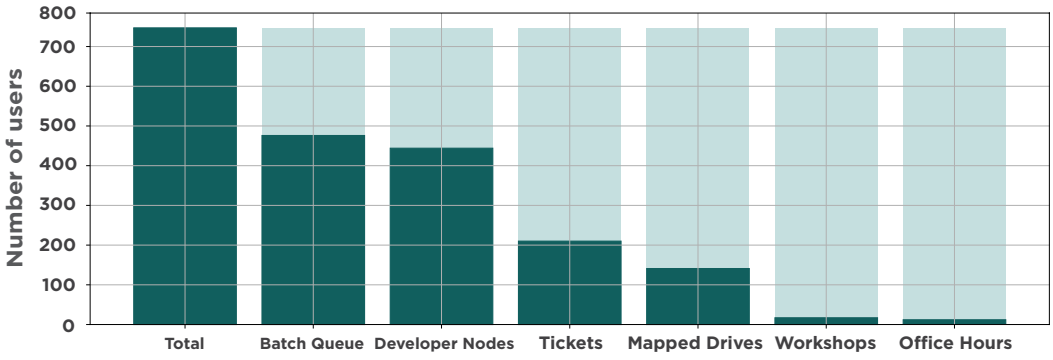

Batch Queue:
Users that run jobs on the main clusters

Developer Nodes:
Users that log into one of our developer nodes

Mapped Drives:
Users that map HPC drives to their local computer (using samba)

Tickets:
Users with active support tickets

Workshops:
Users that attended iCER supported workshops

Office Hours:
Users that attended iCER open office Hours (Mondays and Thursdays 1-2pm)

NUMBER OF USERS USING iCER SUPPORT SERVICES IN APRIL

This figure shows a breakdown of users that use iCER support services. These support services include support tickets, iCER workshops and office hours.

List of iCER workshops in April:

> Webcast: Writing a Successful XSEDE Allocation Proposal
> MSU Science Festival: Tour MSU's Supercomputer
> Monthly Workshop: Introduction to Linux/Unix
> Monthly Workshop: Introduction to HPCC
> XSEDE Webcast: MPI Foundations I & II

* Webinars and video conferences not included

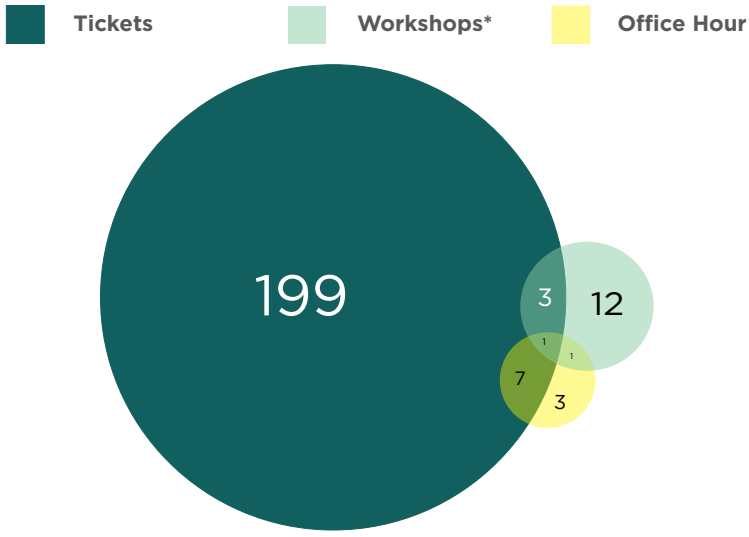

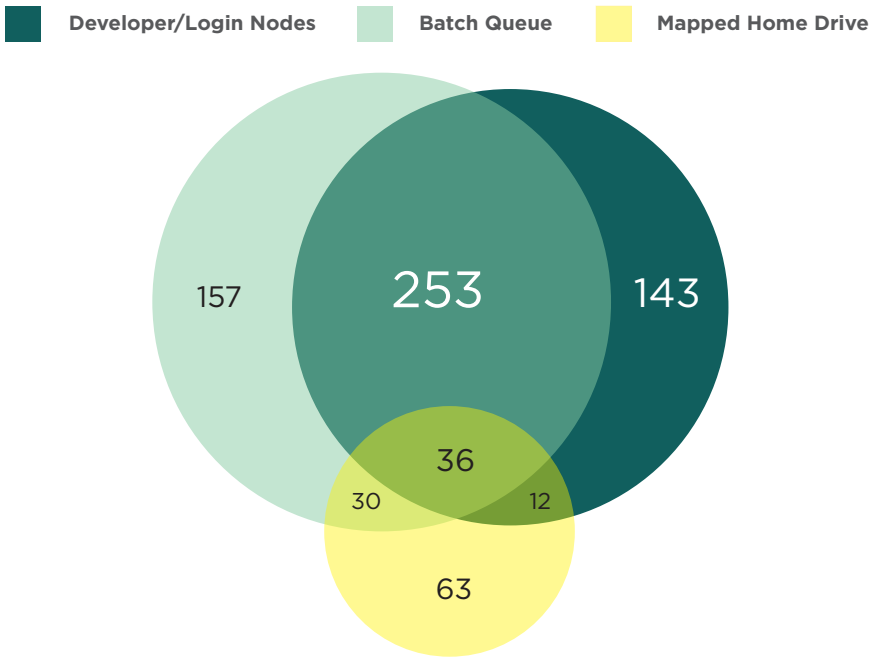

NUMBER OF USERS USING iCER COMPUTE SERVICES IN APRIL

This figure shows a breakdown of users that use iCER compute services:

289 users (253+36) use the developer nodes to submit jobs to the queue.

155 interactive users (143+12) only use iCER developer nodes to do their work. This includes users:
> Only need access to software (ex. Matlab, mathematica)
> Still in software development process and have not submitted a job
> Find development nodes are sufficient for their research.

63 users only used the iCER file systems to store their files.

218 researchers (143+12+63) used iCER hardware outside of the batch queue.

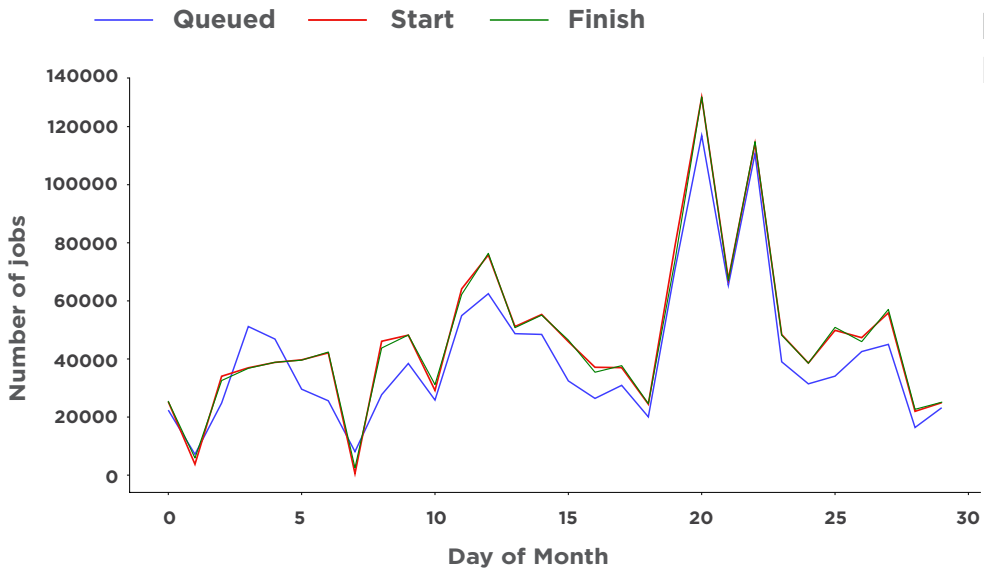

DAILY SCHEDULER ACTIVITY IN APRIL

This figure shows the activity of the batch scheduling system by day. There were two scheduler issues early in this month, and the HPCC team restared the scheduler to address the problem. On a typical day, the scheduler processes approximately 134552 jobs. This includes jobs that are queued, jobs that start and jobs that end. Put in another way, the scheduler manages approximately 93 jobs per minute.

COMPARISON BETWEEN NUMBER OF USERS USING iCER SUPPORT AND COMPUTE SERVICE IN APRIL

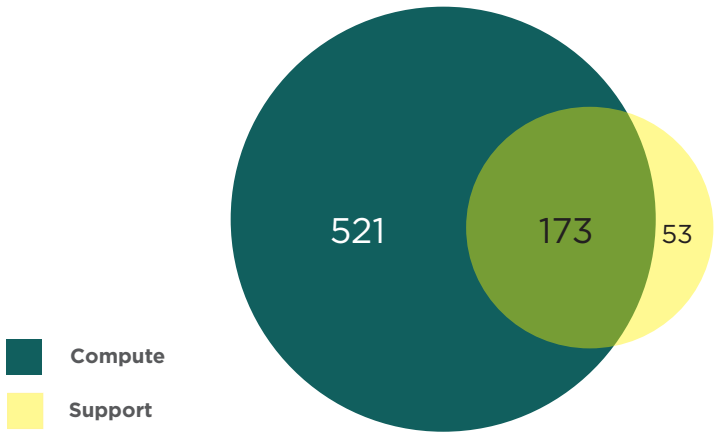

NUMBER OF MAPPED HOME DIRECTORIES PER SERVER IN APRIL

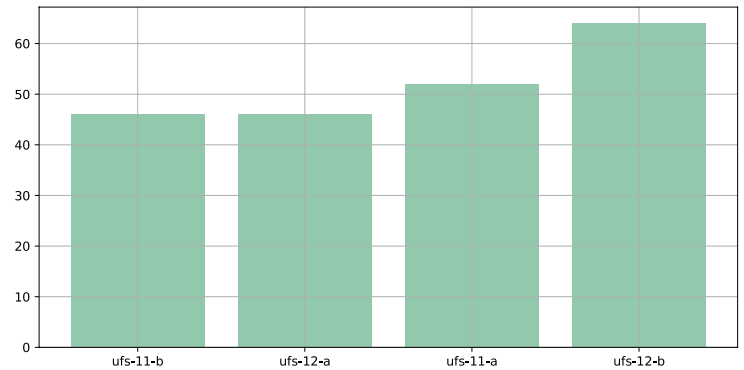

DAILY TICKET STATUS IN APRIL

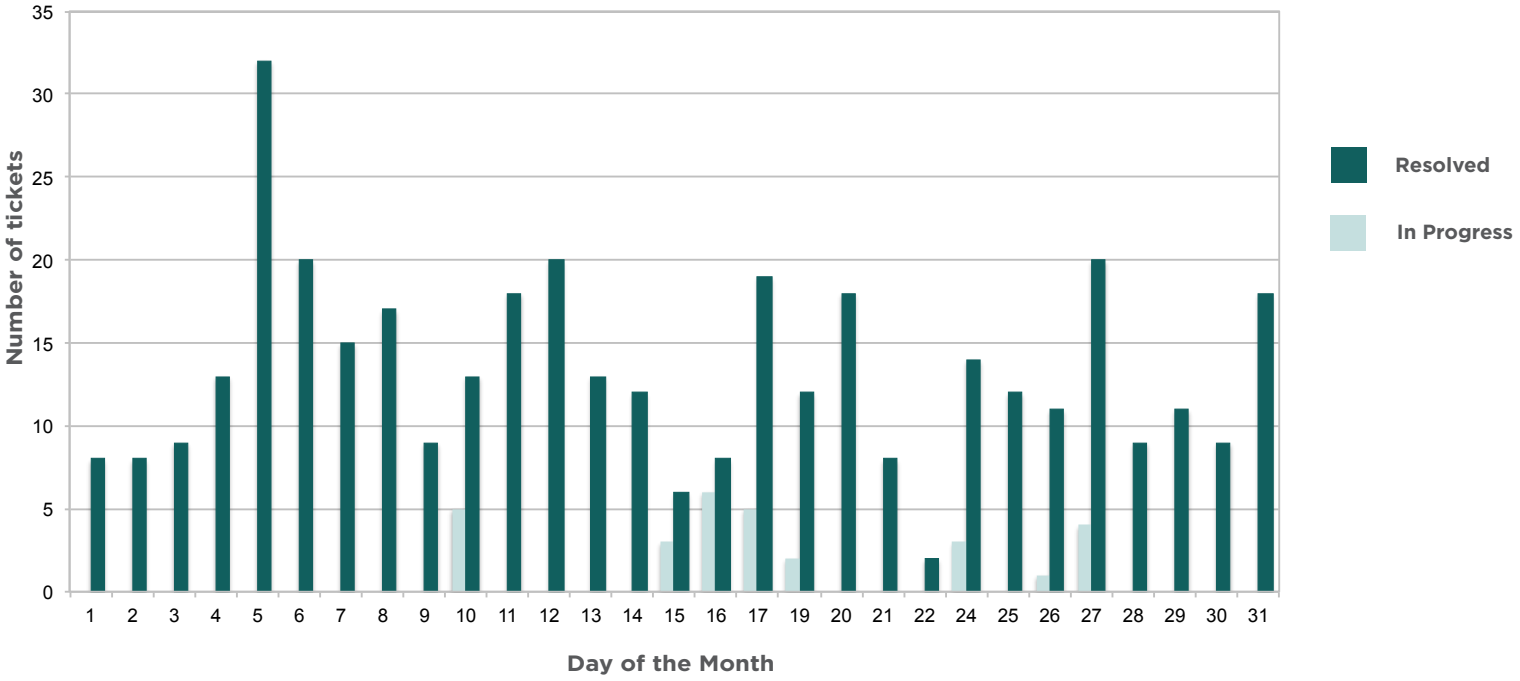

TICKET STATUS SUMMARY IN APRIL

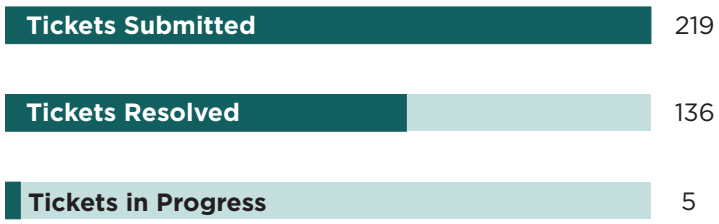

73 tickets are stalled because of no user's responses.

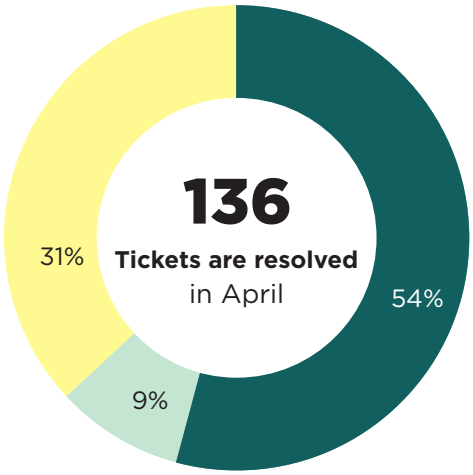

XIAOGE WANG
Research Consultant

APRIL
TICKET HIGHLIGHTS

3 WAYS TO AVOID RESOURCE OVERUTILIZATION ON HPCC

- Tickets resolved within 24 hours
- Tickets resolved within 24-48 hours
- Tickets resolved in more than 48 hours

24

New User Accounts are created
in April

In an effort to better serve our users, we have been analyzing the software that is being used on the HPC by recording which software modules are being loaded using the "module load" command. Clearly this is not a complete view; many users install their own software in their home directories, some modules are automatically loaded as part of a user profile and there will be a bias toward pleasantly parallel codes which will load their required modules every time a job runs (as compared to bigger jobs which would only load the modules once). However, we find this data interesting and wanted to share it with you.

ALL MODULE LOAD COUNTS <10000000

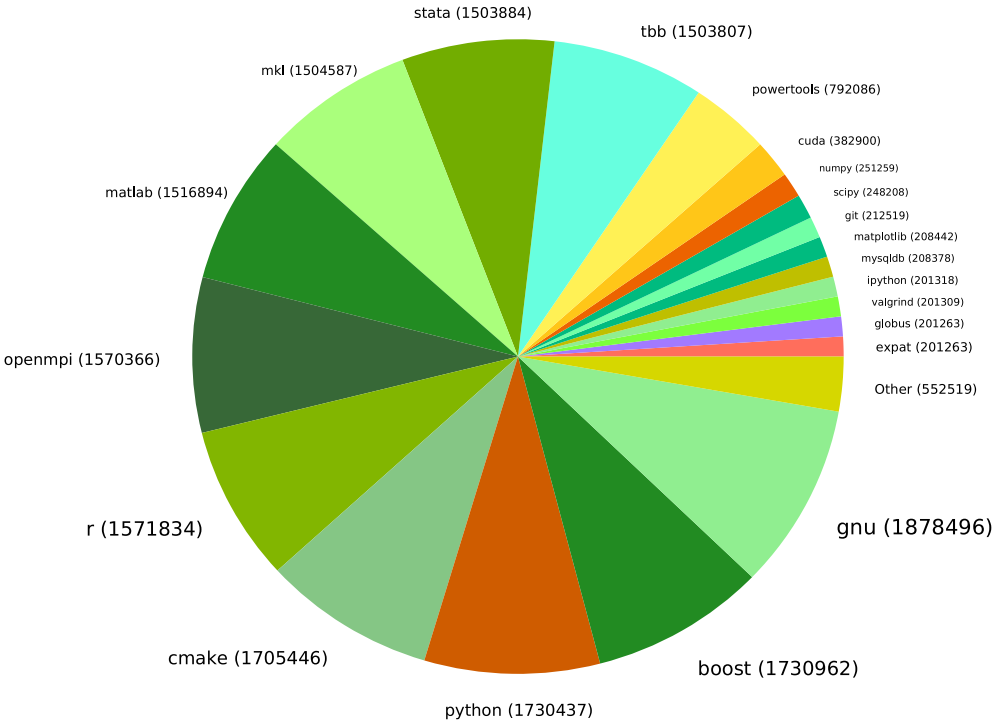

The pie chart shows the most commonly loaded modules. Note again that the biggest ones are the ones included in a user's default profile such as MATLAB, Python, and R. These modules get loaded every time they log in or run a job. As can be seen clearly, the default modules get loaded in an order of magnitude more than many of the other modules.

ALL MODULE LOAD COUNTS <100000

After taking out the default modules, the pie chart on the right shows more modules that users are choosing to include in their .bashrc files and being submitted on a lot of jobs. This group also includes the gateway module which gets loaded every time someone logs onto gateway. This by itself is interesting and shows that we had 22406 gateway connections in April. From our service report we know that 444 unique individuals used a developer node in April. This means that on average each person is logging into gateway approximately 50 times in the month or about 2 times a day (on average).

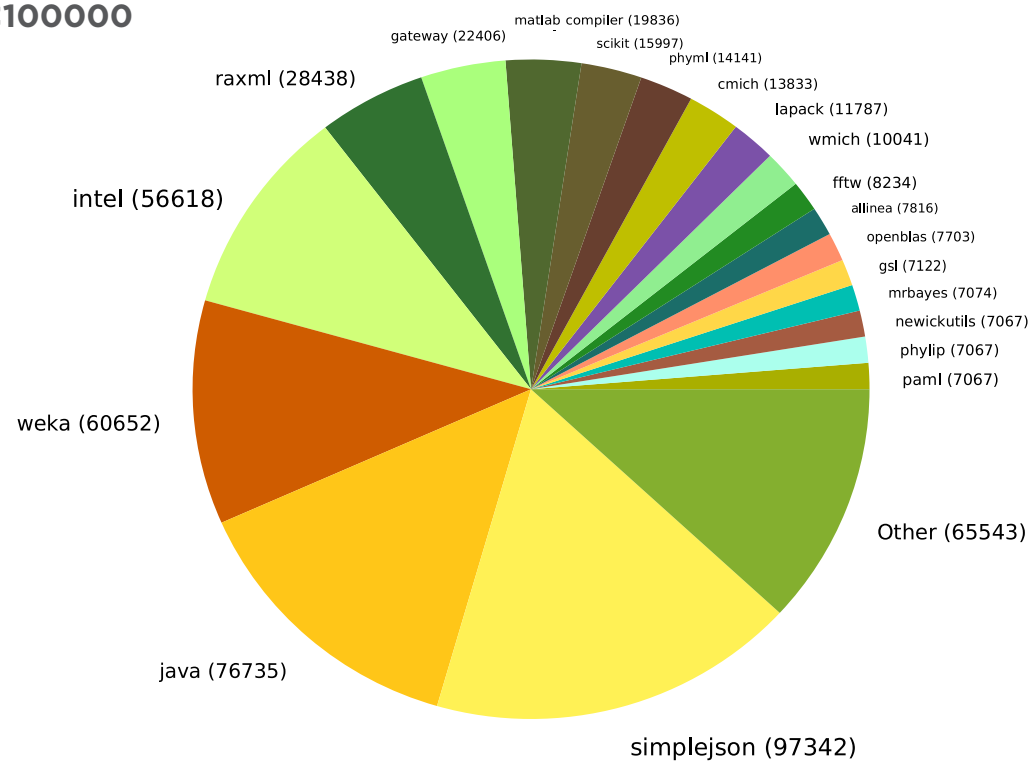

Report Contributors:
Camille Archer
Chun-Min Chang
Jim Leikert
Kelly Osborn
Michael Rometty
Nicol Springer
Xiaoxing (Adele) Han