

## Grid 2.5 Improvements

Improvement Description	Old Version	New Version
<b>Upgraded operating system to support the use of newer versions of software packages</b>	Red Hat Enterprise Linux 6.5	Red Hat Enterprise Linux 7.4
<b>NoMachine graphical environment upgraded and simplified</b>	NoMachine 5.1.54	NoMachine 6.0.78
<b>Updated and new software titles</b>	Stata14, R 3.1.2, RStudio 0.99, Python 2.7, Python 3.4, SAS, MATLAB 2017a, Mathematica 11.1	Stata15, R 3.4.3 <sup>1</sup> , RStudio 1.1, Python 2.7, Python 3.6.5, Spyder (Python IDE), SAS, MATLAB 2018a, Mathematica 11.2, GitKraken
<b>System optimized to support Intel's Math Kernel Library for faster code execution</b>	not available	Up to 10x speedup for certain math, algebraic, and statistical functions & algorithms. <sup>2</sup>
<b>Increased interactive user quotas in terms of both compute cores and memory</b>	Users restricted to 12 cores and 30 GB of RAM	No restriction on RAM (other than physical machine constraints), and users allowed up to 24 cores at one time, spread across a max of 3 sessions <sup>3</sup>
<b>Increased batch user quotas</b>	Users restricted to 12 cores and 30 GB of RAM, including any use of interactive queue (see above)	No restriction on RAM (other than physical machine constraints) or number of jobs <sup>4</sup> , with max of 16 cores for "short" batch jobs, 8 cores for "long" batch jobs, and 8 cores for "infinite" batch jobs <sup>3</sup>

<sup>1</sup> Our release of R is slightly behind CRAN as we are using Microsoft Open R as packaged with Anaconda. This incorporates the Math Kernel Library as a part of the standard build.

<sup>2</sup> Some of these include: Statistical Functions, Vector Mathematical Functions, Partial Differential Equations support, Non-Linear Optimization Problem Solvers, Data Fitting Functions, BLAS and Sparse BLAS, LAPACK, Intel MKL functions for Deep Neural Networks, ScaLAPACK, Sparse Solver Extended Eigensolver, Conventional DFTs and Cluster DFTs, PBLAS, and BLACS. Our R and Python environments should use this automatically.

<sup>3</sup> Note: in order to be able to offer more cores and RAM to users, we have had to organize both interactive and batch jobs into "short" and "long" queues: for interactive queues, short sessions can last up to 1 day (24 hours), while long sessions can last up to 3 days (72 hours); for batch queues, short queues can last up to 3 days (72 hours), while long sessions can last up to 7 days (168 hours). Infinite batch jobs can last until the next scheduled maintenance reboot, but will be of lower priority than both short and long batch jobs.

<sup>4</sup> There are no limits to the number of batch jobs that can be submitted to the short, long, or infinite queues, but a priority score ("Fairshare") determines your dispatch order based on your past usage/other factors relative to others who also have submitted jobs. This allows fair access to compute by all persons.