

Welcome to Biowulf

High performance computing @ NIH

Accelerate biomedical research at the NIH by providing convenient access to large scale **computational resources** and **scientific software** (and the **know-how** to use them).

Biowulf staff - staff@hpc.nih.gov



Steve Bailey



Tim Miller



David Hoover, Ph.D.



Helen Ashdown



Simon Carlson-
Thies



Gennady Denisov,
Ph.D.

**Picture
unavailable**

Afif Elghraoui



Ali Erfani



Andrew Fant, Ph.D.



Beide (Peter) Fu,
Ph.D.



Jonathan Goodson,
Ph.D.



Patsy Jones

**Picture
unavailable**

Charles Lehr



Jean Mao, Ph.D.



Nitish Narula



David O'Brien

**Picture
unavailable**


Ifeanyi Okoye



Wolfgang Resch,
Ph.D.



Antonio Ulloa, Ph.D.



Renbin Yang, Ph.D.



Qi Yu, Ph.D.

Computational resources: The Biowulf system

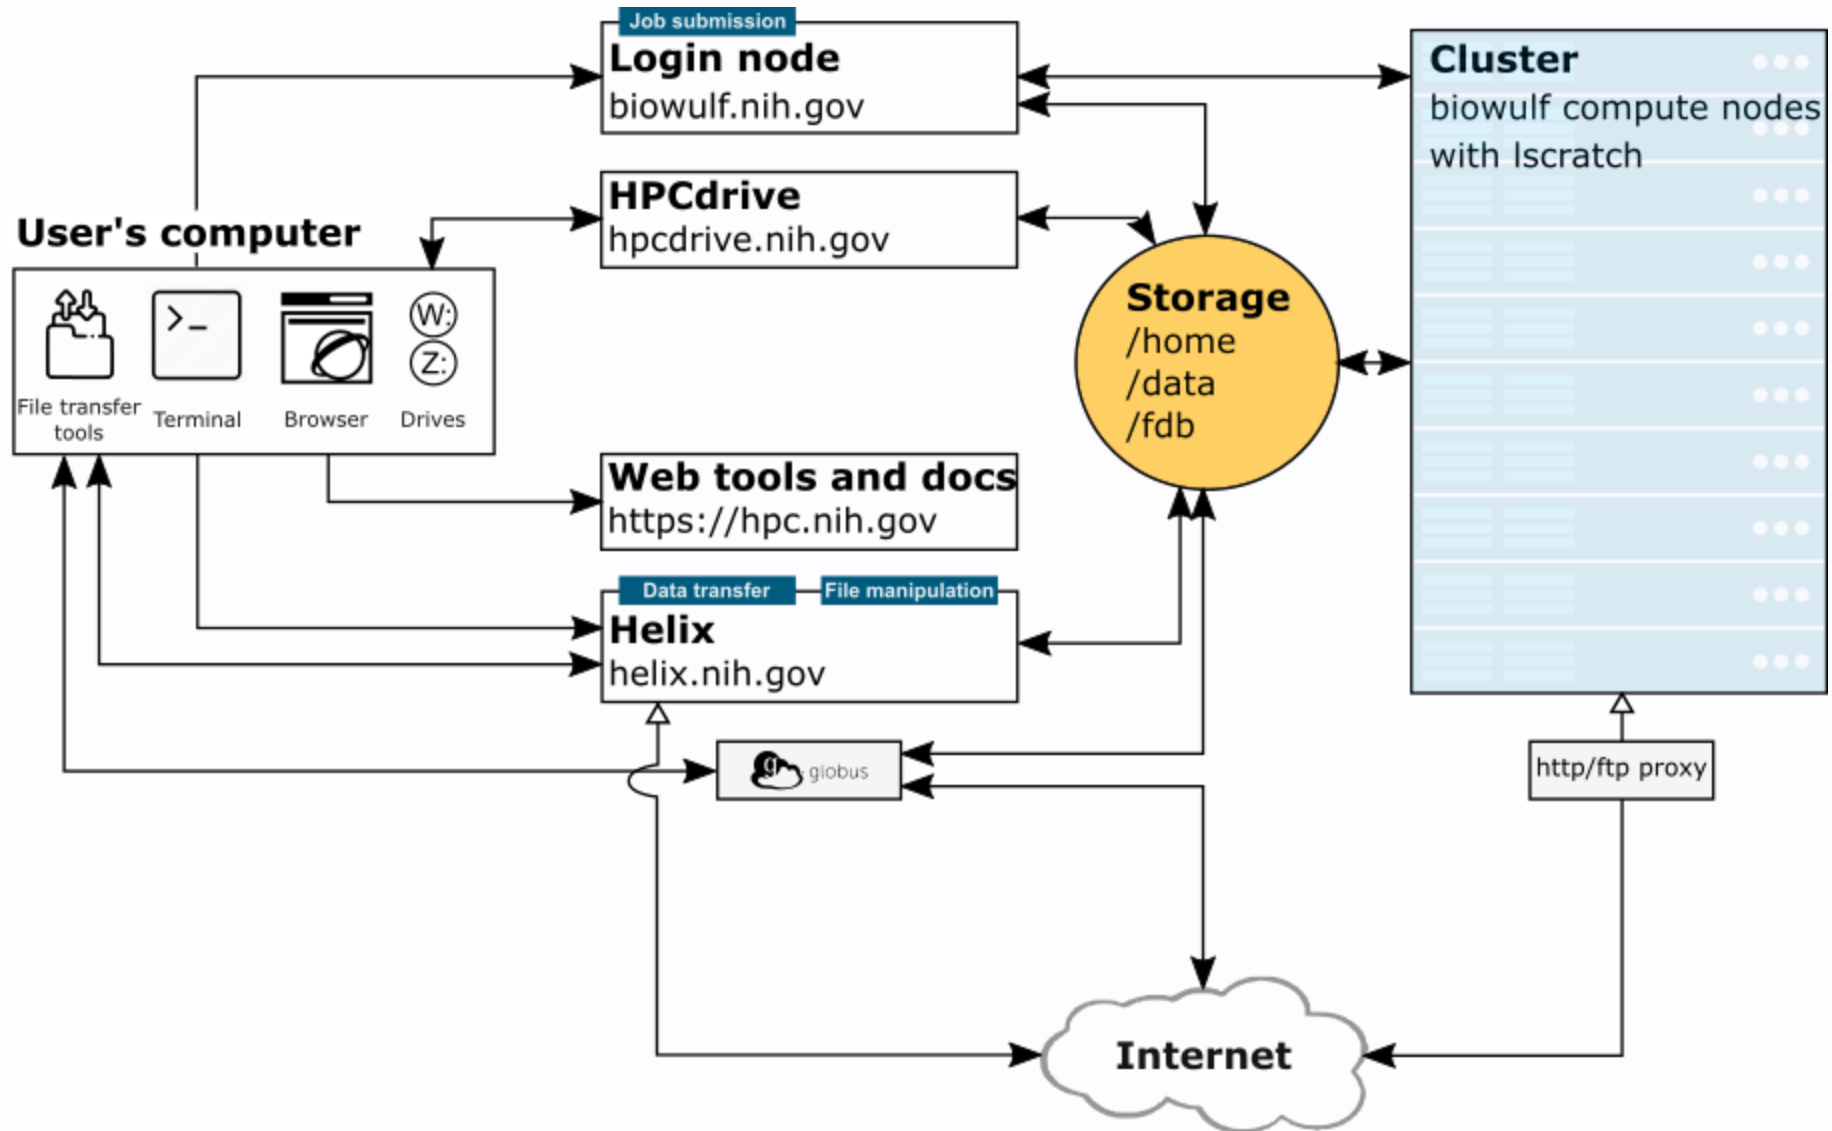
What is Biowulf?



Where is Biowulf?



Architecture



Software

Applications

- multiple versions for ~1000 applications available
- Multiple versions of python with ~500 packages each
- Multiple versions of R with about 1600 packages each
- singularity for containerization
- jupyter, rstudio (desktop, community edition), rstudio-server
- <https://hpc.nih.gov/apps/>

**Know how - training and
outreach**

Training

- Online, self-paced classes (intro to biowulf, bash)
- Videos: https://hpc.nih.gov/training/intro_biowulf/
- Live action classes/recorded (e.g. Deep learning by example, R, matlab, snakemake)
- Documentation (<https://hpc.nih.gov>) and tutorials (e.g. https://hpc.nih.gov/training/gatk_tutorial)
- BTEP Intro to Biowulf at:
<https://bioinformatics.ccr.cancer.gov/btep/courses/introduction-to-unix-on-biowulf-january-2024>

Outreach

- Monthly *Zoom-in* consult
- Meeting with individual groups or 1:1
- Cranky email from staff if/when you break something

How do I ...

Connect to Biowulf

- Commandline sessions
- <https://hpcondemand.nih.gov>

Transfer files to Biowulf

- Large data volumes: **globus**
- Medium/large: **scp/rsync/sftp** (use helix; graphical clients available)
- For convenient small transfers **mount hpcdrive**
- Transfer directly from biowulf to/from NIH
Box/OneDrive: **rclone** or **globus**
- Transfer to/from cloud storage: **globus**

Do interactive analysis on Biowulf

- `sinteractive`
- <https://hpcondemand.nih.gov>

Run a batch job on Biowulf

- 'batch' command for single jobs
- 'swarm' command for parallel jobs
- <https://hpc.nih.gov/docs/userguide.html>

Find commonly used cli tools

```
$ bwulf
usage: bwulf command [options]

bwulf - Unified interface to custom utilities by the NIH HPC staff.

  help|-h|--help      Show this help message.
  version|-v|--version Show this program's version.

Batch System Submission Commands

  sinteractive|sint    Start an interactive session on a compute node.
                      <https://hpc.nih.gov/docs/userguide.html#int>
  svis                 Submit a job for hardware-accelerated visualization.
                      <https://hpc.nih.gov/docs/svis.html>
  spersist            Start an unlimited interactive session (NIMH-only).
                      <https://hpc.nih.gov/docs/nimh.html>
  swarm               Submit an array of similar jobs. <https://hpc.nih.gov/apps/swarm.html>

Other Batch System Commands

  batchlim            Show the current per-user limits for each partition.
  checker             Basic check of SSH permission, quota, conda init, and out of memory jobs.
  dashboard|dsh      Command-line interface to data in the HPC User dashboard <https://hpc.nih.gov/dashboard>.
  freen              Show the free nodes and resources currently available.
  ...
```

Best practices

If you have questions please

email

`staff@hpc.nih.gov`

If you are running jobs on
Biowulf please **read your**
NIH email in case there
are problems.

Be **detailed** and **specific**
when you file an issue (text,
not screenshots, include
jobids, commands, logs, ...)


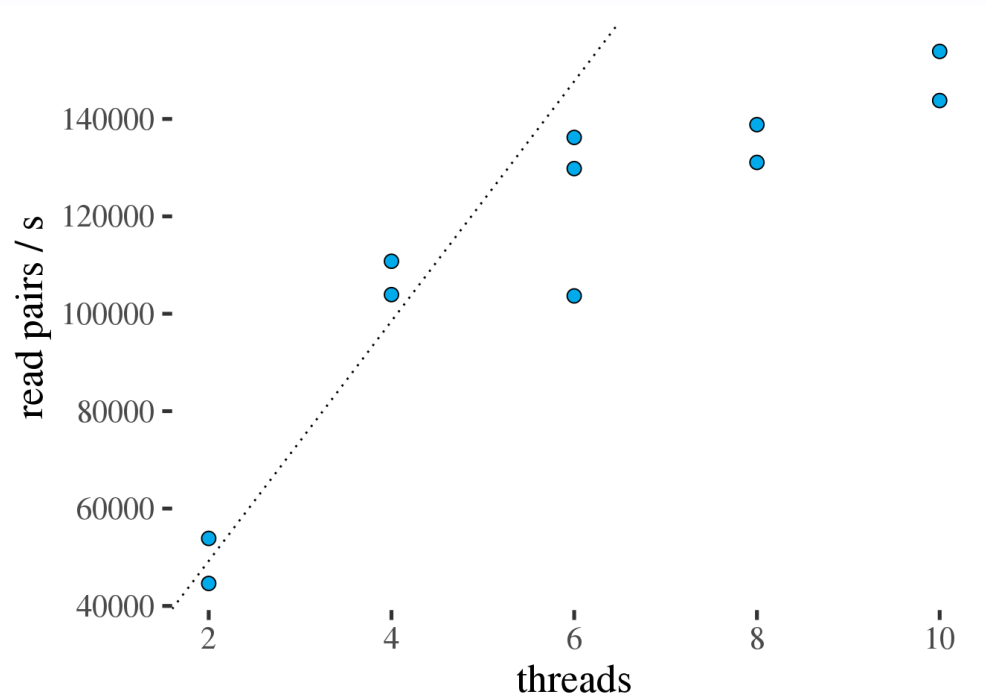
Test small before running
large production jobs

Get familiar with limits on
jobs/CPU/GPU/Memory.
See output of `batchlim` for
current limits

Thank you

... ask us anything at
`staff@hpc.nih.gov`

Scaling matters

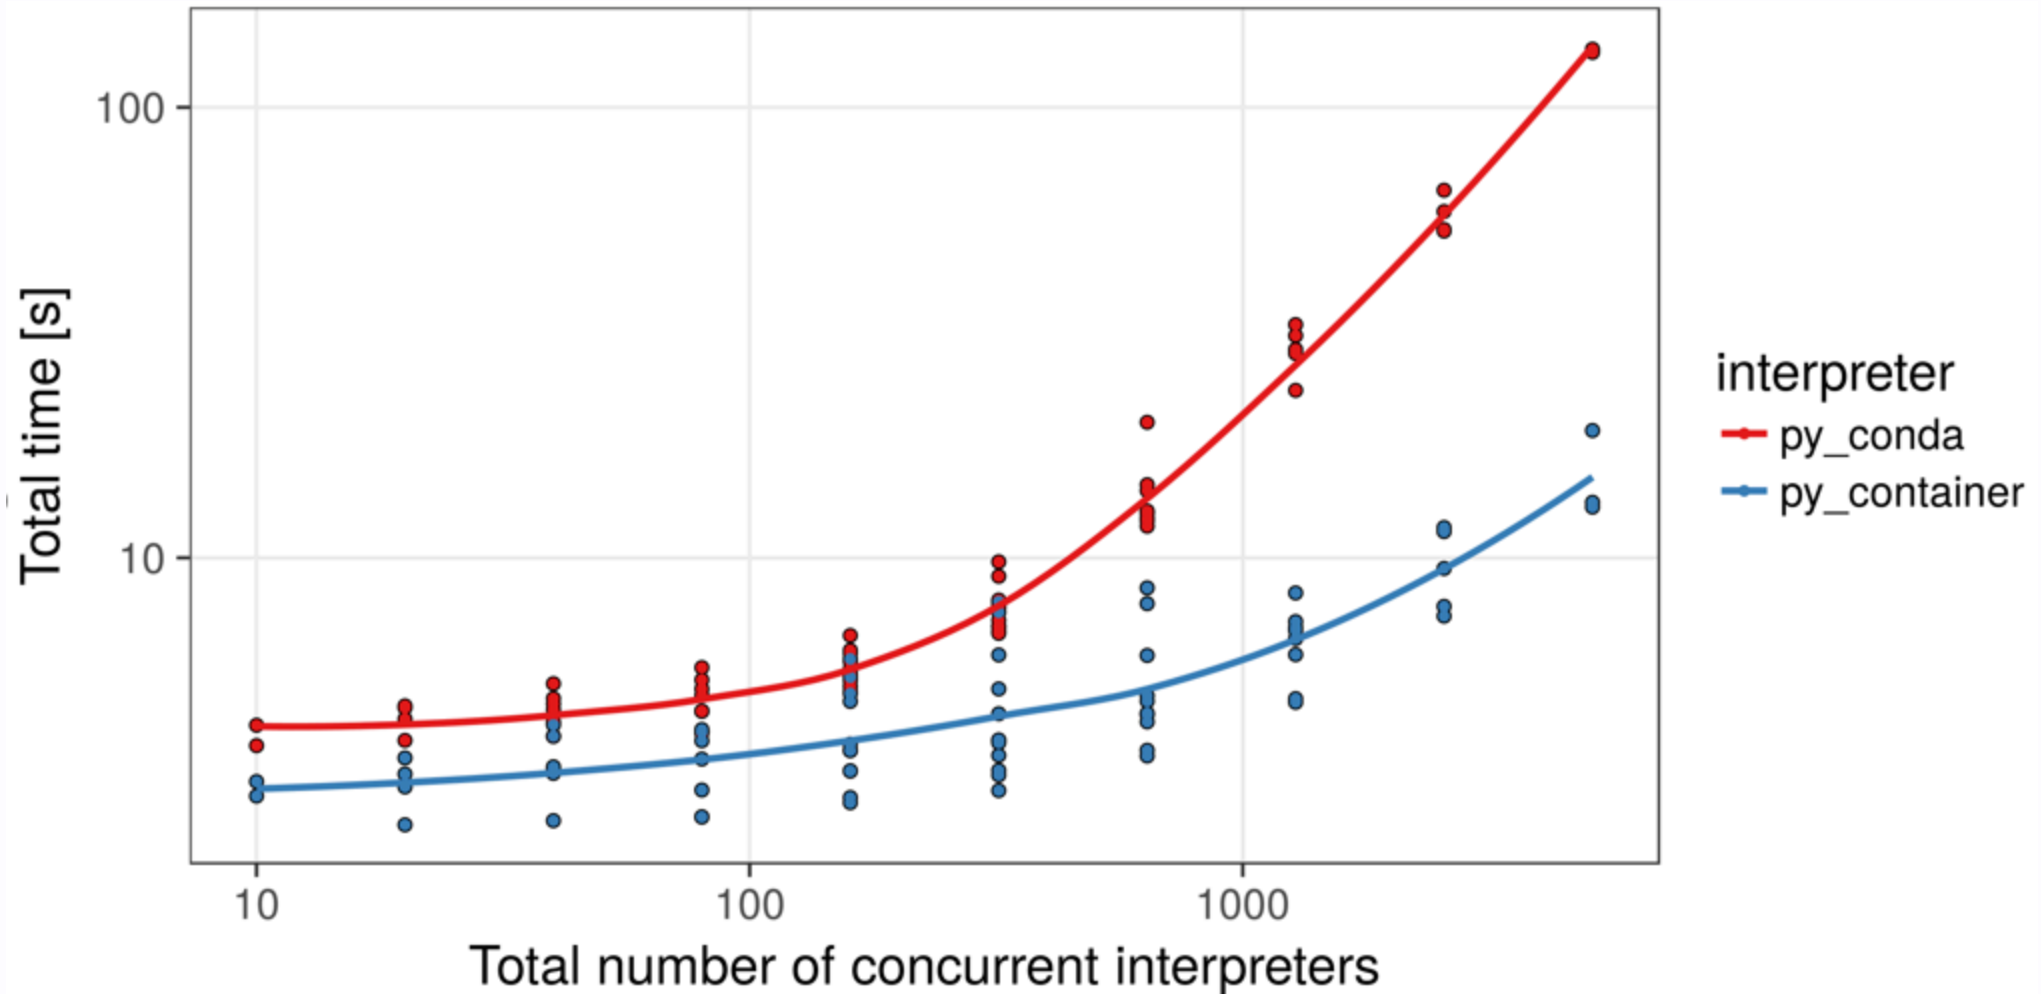


Left: fastp Right: bwa

Be nice to the **file systems**
and use **/lscratch**
and **pipes**

Interpreter startup is
expensive when you run
lots of short processes
concurrently.

For example: concurrent python processes




Make sure you have **space**
in /data for your jobs' output

Make sure your jobs' do not
create more files than
allowed by your **quota** (see
output of `bwu1f`
`checkquota`).

Keep < **5000 files** in a
single directory level

listing files in large directories



Biowulf storage is not for

archiving

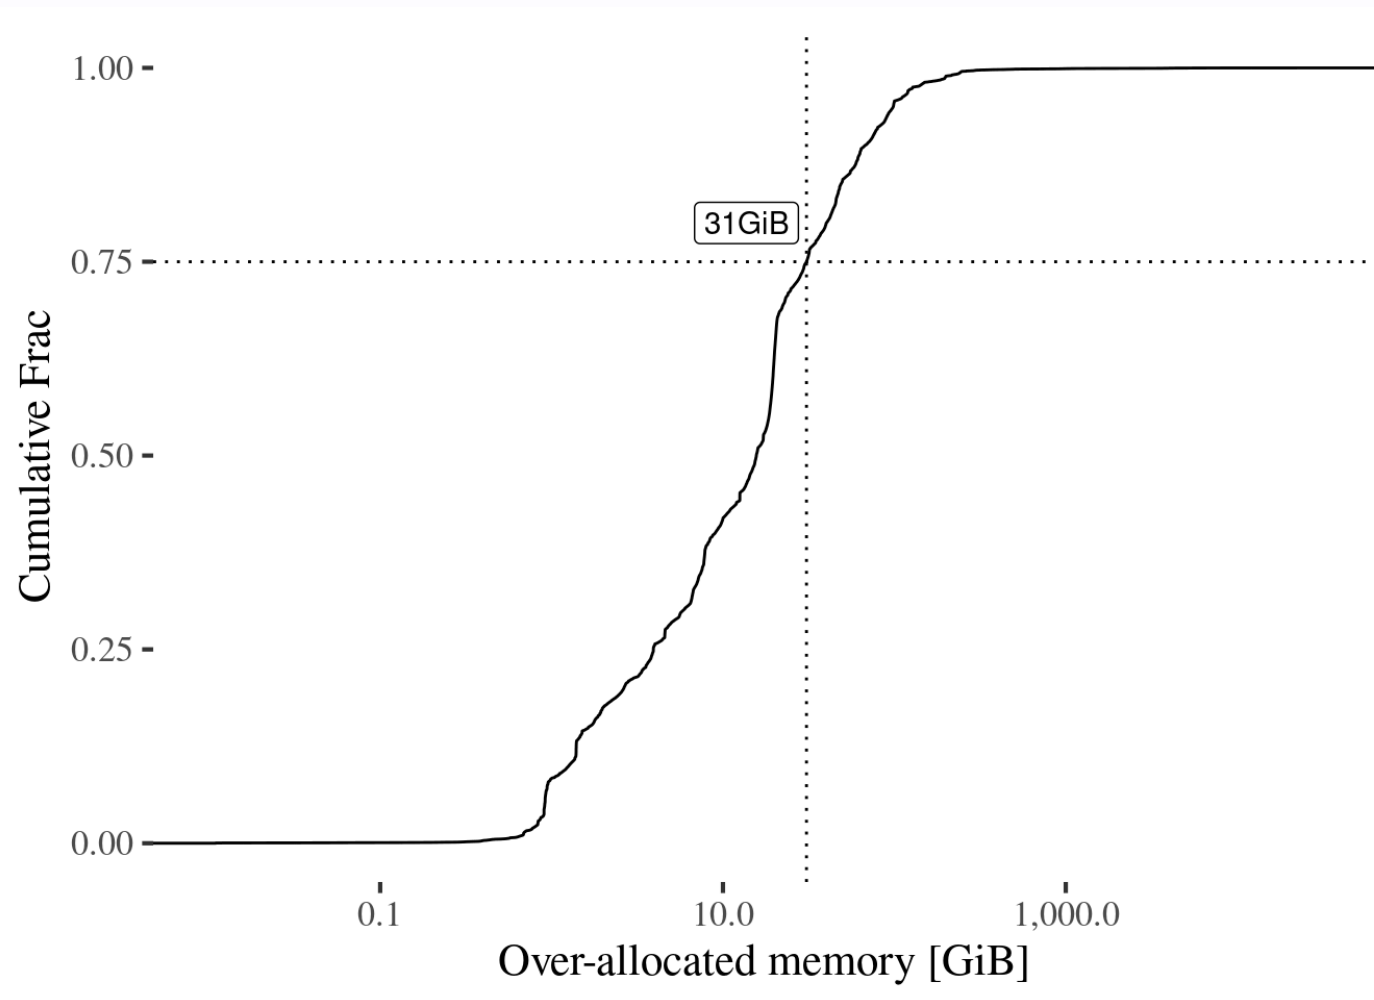
except the object store, kind
of

Backup your data (we
don't)

If you accidentally
delete/modify a file you *may*
be able to
recover it from **snapshots**

Have **good resource estimates** for your jobs
(memory, cpu, time)

Completed jobs in 2023



Just adding more resources
will not always make it run
faster.

Don't **abuse** `bwu1f`
`newwall` to extend walltime
(but it's there if you need it)

Use **job arrays / swarms**
when you can

Try to run jobs with
walltime > 15 min.

Many short jobs overload the
scheduler

Do not run tools that query the scheduler (`queue`, `sacct`, ...) in a loop. It overloads the scheduler. `bwulf dashboard` is ok.

Keep **startup files**
simple (no conda, no
module loads, ...).

Don't `sudo` , `apt-get` ,
`dnf` , or `yum` . It won't work.

Your workstation is probably
better for graphical
applications.

Please **acknowledge**
biowulf in your
publications.

Talk to us.

`staff@hpc.nih.gov`

File systems

103,500 cores / **207,000 CPUs** / 3558 nodes

916 TB memory

3.4 PB local scratch (lscratch)

50.7 PB shared high performance storage

5 PB of object storage (*)

947 GPUs (K80 - A100)

Connected by **FDR/HDR infiniband** @ 56-100 Gb/s.

Connected to NIH backbone @ 2x100Gb/s

Node classes

- Single node, CPU (56-128 CPUs, 240-500GB)
- Multinode, CPU (56 CPUs, 240GB)
- GPU (K80, P100, V100, A100)
- large memory (1.5-3TB)
- visualization, dragen