

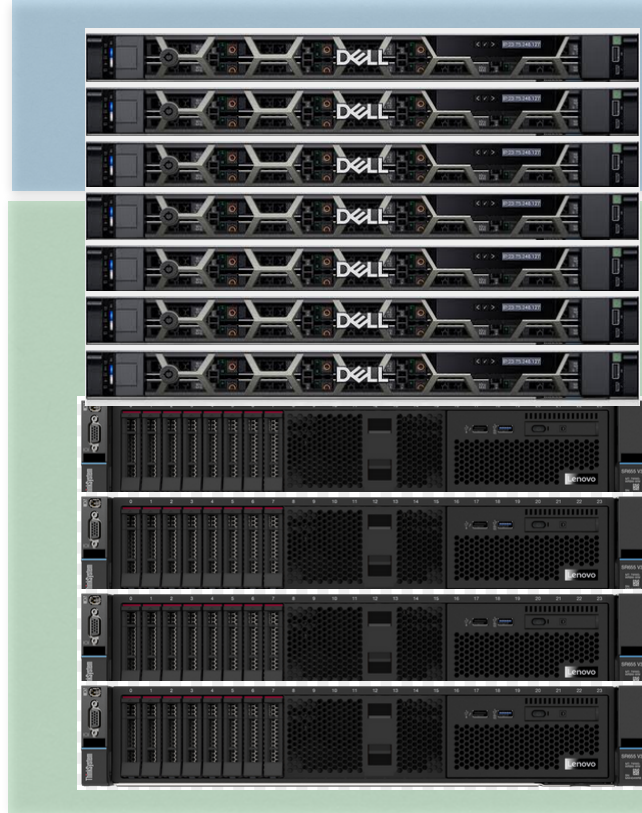
# **New Cluster Introduction**

Replacing Jalapeno

Duncan Mortimer, 24/4/2024

# New Hardware

364 Cores  
12 K class 12GB GPUs




3 Management Hosts  
(Web and software hosting)

2 Remote Desktop Hosts  
2 Compute + 4 GPU Hosts

288 Cores  
12 A class 12GB GPUs  
2 H class up-to 80GB GPUs

# Accessing Ood



## Connecting to and using

- **Open On Demand** - <https://ood.fmrib.ox.ac.uk>  
Web based VNC
- SSH - via [clint.fmrib.ox.ac.uk](https://clint.fmrib.ox.ac.uk)
- SFTP/SCP/rsync - via [sftp.fmrib.ox.ac.uk](https://sftp.fmrib.ox.ac.uk) (or file transfer via Ood)
- No legacy X11/VNC

# Changes

## Software

- Rocky Linux 9 (Red Hat clone) (Jalapeno is 7.3) - very old software may no longer work
- All research software distributed by a much faster mechanism via different locations

No /opt/fmrib - Check your .bashrc to see if you setup Conda from /opt/fmrib

Ask [computing-help@win.ox.ac.uk](mailto:computing-help@win.ox.ac.uk) if you have Conda environments you use regularly on the cluster - central installation could speed up running

MATLAB much faster to launch

- Software configured using 'module' command
  - Only fsl\_sub loaded by default (a few system configuration modules too)
- Some module names may have changed (use `module avail` to list)
- Older software versions may no longer be installed



# Changes

## Cluster

- SLURM (<https://slurm.schedmd.com/>)
- Queues mirror BMRC's cluster:
  - veryshort.q + short.q > short (1 day 6 hours)
  - long.q + verylong.q + bigmem.q > long (10 days)
  - interactive.q > interactive (10 days) - only to be used via OOD
  - gpu.q >
    - gpu\_short (4 hours)
    - gpu\_long (2.5 days)
- Recommend `fsl_sub -T` is used to specify number of minutes required
- Use `fsl_sub -R <GB>` to request more than default 16GB per slot
- /tmp is private to each job
- No K40 GPU class
- No 'infinite' queues

# Changes

## Remote desktop environment

- Runs as a cluster job
- Xfce (jalapeno is lxqt)
- Session is limited in CPU (core equivalents) and memory - programs requesting more than you have requested may be killed
- /tmp is private to the session (cannot access files from other sessions or users)
- GPUs available for CUDA tasks - please use sparingly

# Documentation

## Links

- Overview: [https://www.win.ox.ac.uk/research/it/latest-win-it-news/ood\\_cluster](https://www.win.ox.ac.uk/research/it/latest-win-it-news/ood_cluster)
- OOD: <https://www.win.ox.ac.uk/research/it/i-want-to/wfh/remote-desktops/fmrib-remote-desktops/fmrib-ondemand>
- SSH connections: <https://www.win.ox.ac.uk/research/it/i-want-to/wfh/ssh>
- Cluster: <https://www.win.ox.ac.uk/research/it/i-want-to/analyse-my-data/use-hpc-compute-resources/fmrib-compute-cluster>
- Environment modules: <https://www.win.ox.ac.uk/research/it/i-want-to/analyse-my-data/environment-modules>
- BASH Customisation: <https://www.win.ox.ac.uk/research/it/i-want-to/analyse-my-data/bashprofile>
- Python/Conda: [https://www.win.ox.ac.uk/research/it/i-want-to/analyse-my-data/python\\_conda](https://www.win.ox.ac.uk/research/it/i-want-to/analyse-my-data/python_conda)