



FEBRUARY 19th 2026 | OAK RIDGE TENNESSEE

Experimental Computing Laboratory (ExCL) Monthly Meeting

Steve Moulton, Aaron Young, Jeffrey Vetter

Advanced Computing Systems Research Section



U.S. DEPARTMENT
of **ENERGY**

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY



Maxwell (was hyp-100) being redeployed from ORC to ExCL

- 2x EPYC 7763 64 core processors
 - will probably run hyperthreaded for 256 cores total
- 1 TB Main Memory
- 4 NVIDIA A100 GPUs (40GB, no nvlink)
- 4 AMD MI100 GPUs (32GB)

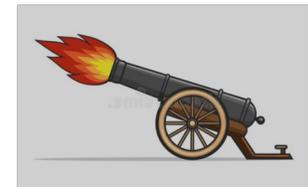
Well suited for general loads, LLM work, and heterogeneous accelerator studies

This system is tricky to configure – will try to leave current OS release in place

Will be integrated (i.e. home directories and other resources) into ExCL

ThinLinc Usage and Updates

- The ThinLinc standard license has reverted from 10 concurrent users to 3 concurrent users.
- We have updated ThinLinc to 4.20.0.
- We have contacted ThinLinc for licensing in our somewhat different environment
 - We now have 90-day licenses for 10 concurrent users for ExCL systems and are working on purchasing licenses.
- Please close your sessions (log off) so licenses are not consumed by idle sessions if not needed.
 - Else we will guess on which sessions to discontinue
 - ThinLinc sessions are persistent (much like tmux)

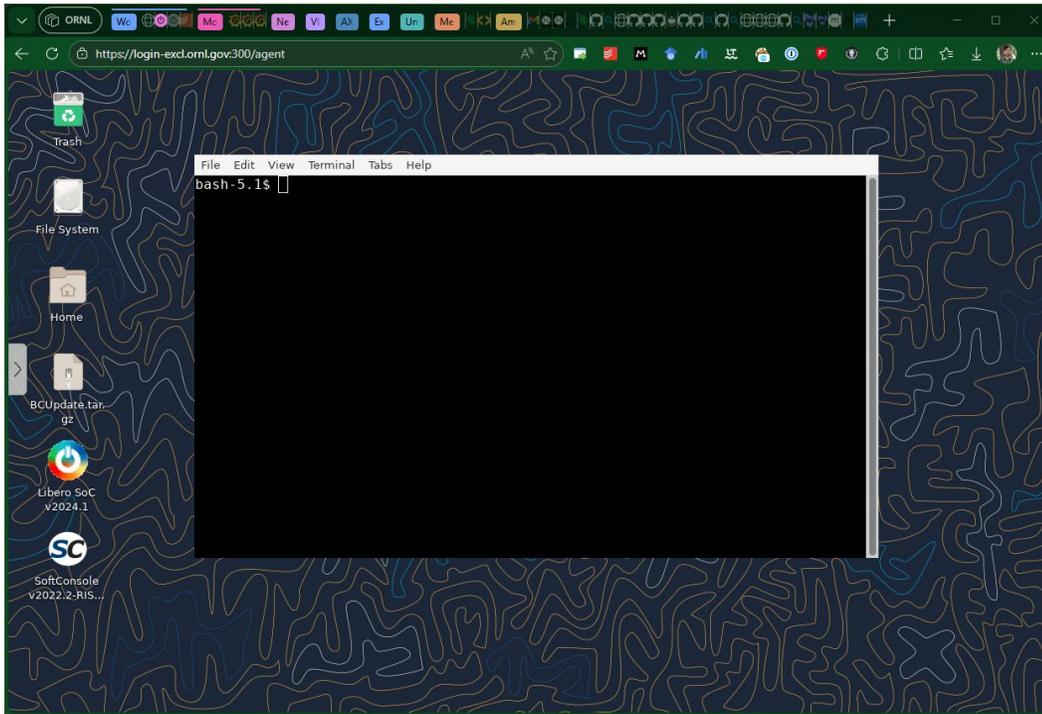


ThinLinc: Broken XFCE Config Warning

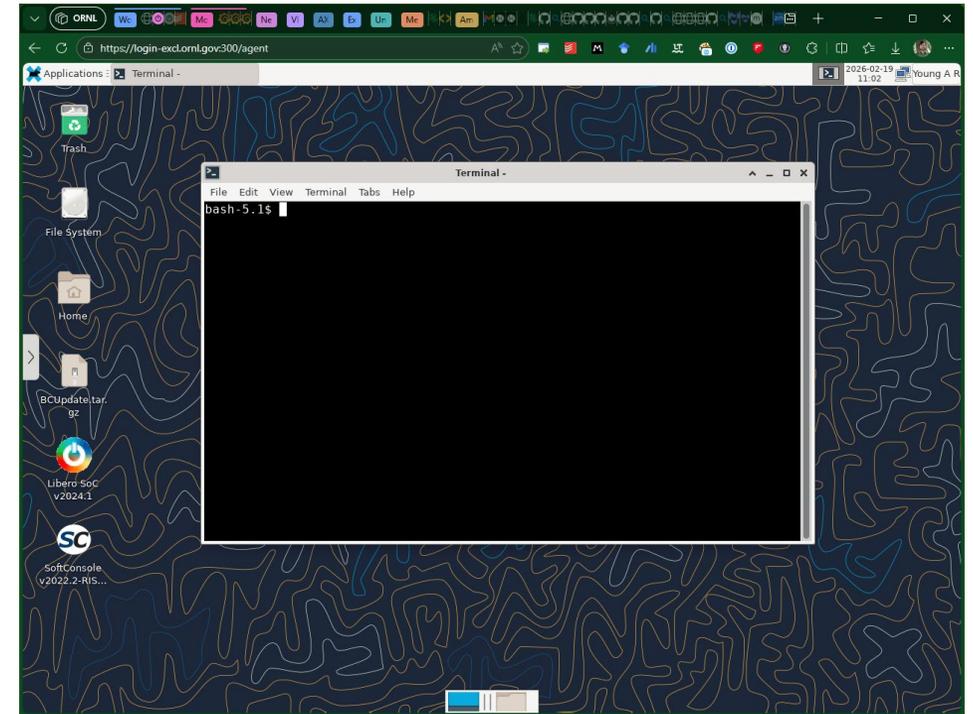
```
rm -rf ~/.config/xfce4
rm -rf ~/.cache/sessions
rm -rf ~/.local/share/xfce4
rm -rf ~/.config/xfce4/panel
```

Warning for broken xfce4 config:

- If your environment is messed up in xfce4 remove config to reset state



Bad



Good

Apptainer Deployment

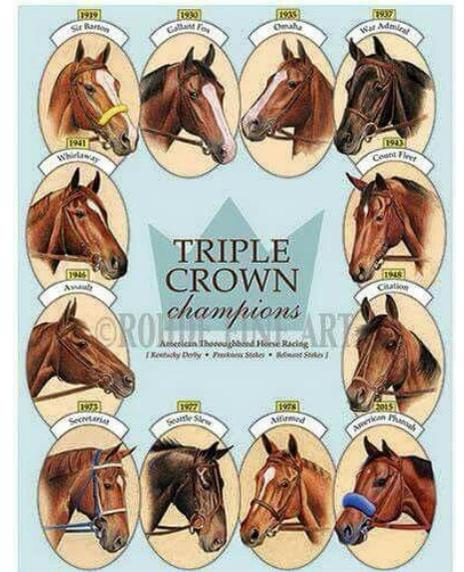
[Apptainer | ExCL User Docs](#)

Apptainer has been deployed and tested on the Triple Crown servers.

- Affirmed, Justify, Pharaoh, Secretariat

Apptainer has also been deployed on Faraday.

As other systems are redeployed, or on request, Apptainer will be deployed on all systems where it makes sense.



Documentation Updates

- [Python | ExCL User Docs # Python Virtual Environments Using uv](#)
- [Marimo | ExCL User Docs # Quick Start](#)
- [Jupyter Notebook | ExCL User Docs # Using FoxyProxy \(Recommended\)](#)

Demonstration:

How to use ThinLinc, FoxyProxy, uv, and port forwarding

Other ongoing efforts

- ExCL microelectronics and neuromorphic laboratory (3700/A102)
 - Networking complete (ExCL VLAN).
 - Stools delivered to random location. Will track it down.
 - Move of devices from 5100/227 (and offices) to be scheduled
- GPUs to deploy (consumer grade but high-power requirements)
 - Nvidia GeForce RTX 5090
 - AMD Radeon RX 9070 XT
- In light of increasing AI inference workload
 - We are restarting discussion with Grok as to whether their accelerator cards are still supported, and what steps needed to bring them back online.
 - Discussions with Tenstorrent

Ideally these devices would shift workload from GPUs.



Questions/Projects/Comments/Discussions

- Julia in slurm question: Make sure to request GPU resources with Slurm if have jobs which require GPU resources.
For example: `--gpus=2`

