

MAY 15th, 2025| OAK RIDGE, TENNESSEE

Experimental Computing Laboratory (ExCL) Monthly Meeting

PRESENTED BY EXCL

Steve Moulton, Aaron Young, Jeff Vetter Advanced Computing Systems Research Section



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

Infrastructure/System Status

- news command implemented as a shell command. Up to the minute information will be available there. Login banner will show whenever there is news you have not read. This will eventually be integrated with the <u>news@excl.groups.io</u> mailing list.
- Mi300a. Faraday, deployed and ready for use.
- Aaron has done considerable work on account reporting and accounting. This will allow us to spend less time on account management as more of the process is automated.
- Continued vulnerability management. Setting gateway machines to auto update (Thursday night 10 PM) and automatically reboot if necessary.
- U280. Deployed to milan3. Updated slurm helper functions and environment setup prolog.
- ThinLinc updated to 4.18.
- Siemens remix and license deployed. No change to selections.
- Tiny Pilot updated and it now has a nice run script menu which can run REISUB to remotely reset a system.



News

Welcome to the Experimental Computing Laboratory

Support:

Support Documentation: https://docs.excl.ornl.gov/ Email support: excl-help@ornl.gov Getting Started Cheat Sheet https://docs.excl.ornl.gov/#excl-cheat-sheet Chat/community: https://ornlccsd.slack.com/#excl Latest ExCL systems news: type news on command line there is new news

Citation Information:

The following acknowledgment should be included in publications and presentations that contain work performed using ExCL resources:

This research used resources of the Experimental Computing Laboratory (ExCL) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-000R22725.

See https://docs.excl.ornl.gov/acknowledgment for details.

zenith x86_64 22.04.5 LTS (Jammy Jellyfish) 64 cores 125Gi



News

7ry@zenith:~\$ news

```
Tuesday May 8, 2025
```

The MI300A system (host name faraday) is available for ExCL users. As usual you have to log in through the login node. Make sure that you "module load rocmmod" to set up your environment.

```
Monday May 7, 2025
```

 Upgraded oswald00,02,03 and radeon to Ubuntu 24.04 as the easiest path to supporting gprofng. Modified kernel configuration to allow users to use these tools accessing all kernel performance counters.
 ExCL maintenance outage has been deferred.

```
Tue April 29, 2025
. U280 deployed to Milan3.
```

Wed April 16, 2025

 equinox gpus are displaying an error on use, and are not available.
 We hope reseating the nvlink cabling will address this (as it has for others). We will be doing this sometime the week of April 21.



MI300A What is it?

- Supermicro AS -4145GH-TNMR
 - No configuration options wrt memory or other addons.
- 4 APU (Accelerated Processing Unit) (combined CPU, GPU and HBM3 memory)
 - 912 CDNA 3 GPU units
 - 96 Zen 4 cores
 - 512 GB unified HBM3 (128 per APU)
- Supermicro designed and built system (we have 4U air cooled, also available as 2U liquid cooled)
 - Rather than the normal PCIe 5.0 slots, riser cards that connect into specialized backplane connectors are used (but they are PCIe 5.0).
 - To add hardware we will need to purchase riser cards, and lots of heads-up time.
- Ubuntu 24.04 LTS; ROCM 6.4.0



What it looks like



4U AS -4145GH-TNMR (Air-cooled)





Why is it interesting

- Higher capability AMD GPUs (interesting for IRIS)
- 256 MB AMD Infinity Cache[™] between CPU and GPU cores
- Lots of memory
- Infinity network interconnect.
- Anticipated uses
 - Memory hierarchy performance studies
 - 4 fast GPUs probably great for AI performance
 - It is fast (LLVM built in 5 minutes)



How to access

- Usual login process through login.excl.ornl.gov.
- Host name is faraday
- module load rocmmod
- A good environment test program is at <u>https://github.com/jungwonkim/amd-toy</u>.
 - If you see any non-zero status values, you may not be in the render group. Contact <u>help@excl.ornl.gov</u> for help.
- Documentation starting on docs.excl.ornl.gov; not much there yet except for a billboard. <u>faraday | ExCL User Docs</u>



Power Monitoring the Mi300a

 Will improve monitoring with CheckMK integration and will add documentation for how to configure power settings.

Every 1.0s: rocm-smishowpower	<pre>faraday.ftpn.ornl.g nqx@faraday:[12:40 PM]~/Faraday/iris\$ rocm-smishowsclkrange</pre>
Big System Management Interface GPU[0] : Current Socket Graphics Package Power (W): 80.0 GPU[1] : Current Socket Graphics Package Power (W): 79.0 GPU[2] : Current Socket Graphics Package Power (W): 102.0 GPU[3] : Current Socket Graphics Package Power (W): 77.0	ROCm System Management Interface ====================================
End of ROCm SMI Log	End of ROCM SMI Log nqx@faraday:[12:40 PM]~/Faraday/iris\$

	======================================
	======================================
GPU[0]	: Max Graphics Package Power (W): 760.0
GPU[1]	: Max Graphics Package Power (W): 760.0
GPU[2]	: Max Graphics Package Power (W): 760.0
GPU[3]	: Max Graphics Package Power (W): 760.0



Online Documentation

- Available models: <u>https://www.supermicro.com/en/accelerators/amd</u>
- Datasheet on Faraday: <u>https://www.supermicro.com/datasheet/datasheet_H13_QuadAPU.pdf</u>
- Hardware documentation: <u>https://www.supermicro.com/manuals/superserver/4U/MNL-2754.pdf</u>



Questions/Discussions/Comments?



