



JANUARY 15TH 2026 | OAK RIDGE TENNESSEE

Experimental Computing Laboratory (ExCL) Monthly Meeting

PRESENTED BY

Steve Moulton, Aaron Young, Jeffrey Vetter

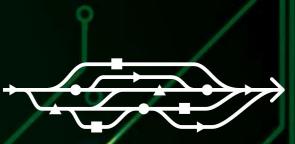
Advanced Computing Systems Research Section



U.S. DEPARTMENT
of ENERGY

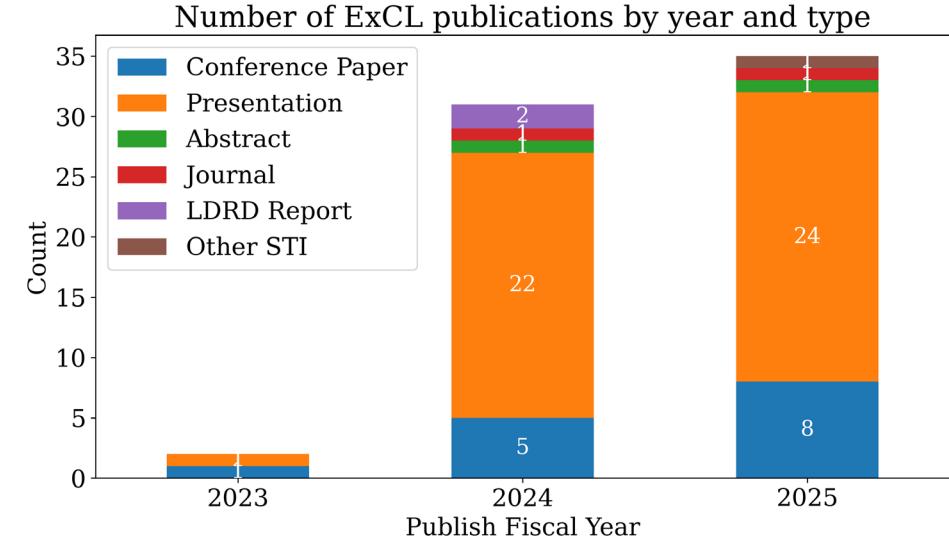
ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE U.S. DEPARTMENT OF ENERGY





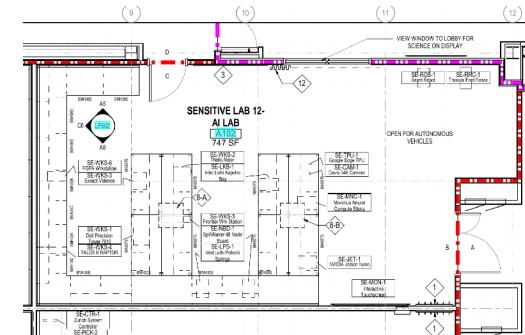
Experimental Computing Laboratory

- Modernized infrastructure services.
- Designed and started implementation of TRC microelectronics and AI lab (3700/A102)
- 98 users across 34 groups, 13 divisions and 7 directorates. 66 external users across 29 organizations. 51 new users (users are periodically archived).
- 11 papers, 2 journals, 27 presentations, 1 abstract, 1 other STI cited ExCL in Resolution in the past year.
- Emphasizes on reliability and storage stability, user self-service wherever possible, with consulting to help enable research goals.



3700 (TRC - Translational Research Capability) Lab A102

- **Purpose:** Lab space dedicated to embedded systems, microelectronics, AI, and neuromorphic computing.
 - **Core Capabilities:** Direct ExCL network access (deployment in progress), support for exposed power (≤ 120 V), and FPGA-PC integration.
 - **Key Equipment:** Zenith systems, DGX Spark, neuromorphic testbeds, SDRs, and photonic-ready infrastructure.
 - **Future Vision:** Expansion into quantum, photonics, and automated AI-driven experimentation



ExCL Newsletter

- Inaugural issue of the ExCL Newsletter was distributed.
- The goal of this newsletter is to highlight research conducted on ExCL, notify of upcoming ExCL outages, and provide reminders for best practices.
- First edition features ExCL at SC 25.

EXPERIMENTAL COMPUTING LABORATORY
MONTHLY NEWSLETTER

VOLUME 1 ISSUE 1

MONDAY, JANUARY 12, 2026

ExCL Update

ExCL at SC 2025

ExCL had a strong presence at SC25. Here are some highlights. See <https://www.excl.ornl.gov/excl/sc25/> for full details.

- **Best Paper at WACCPD**
SULI intern Tatiana Melnichenko and team recognized for work on Mojo's vendor-neutral GPU programming across NVIDIA H100 and AMD MI300A, enabled by ExCL resources.
- **Top-5 ACM Undergraduate Research Competition Finalist**
Tatiana's poster selected as a finalist, showcasing strong internship engagement in ExCL projects.
- **3rd Place ACM SRC Poster Award**
Matthew Chung, mentored by Narasinga Rao, earned 3rd place for work on a heterogeneous eigen solver using IRIS on ExCL systems.
- **SC Best Poster Nomination**
Poster on Julia with an Intelligent Runtime System nominated for Best Poster, highlighting dynamic optimization strategies for heterogeneous systems.
- **AI-Driven HPC Development Papers**
Two papers—**ChatHPC** and **ChatPort**—demonstrated how AI accelerates HPC software development, leveraging ExCL to fine-tune **CodeJulia** with DOE-critical libraries (Kokkos, OpenMP, ADIOS2, IRIS, TAU).
- **Short Talks at Workshops**
Philip W. Fackler: GPU-accelerated phonon calculations at XLOOP 2025.
William Godoy: Mojo's MLIR-based performance-portable GPU programming at Python HPPSS.
- **BoF and Tutorials**
Philip W. Fackler and William Godoy presented at Julia for HPC BoF and tutorial, with ExCL resources powering CI for **JACC_il** performance-portable Julia library.

Upcoming Outages

- **Tuesday, January 13th** full system outage with rebooting to update kernels and libraries. The planned duration is from 9:00 am to 11:00 am.
- **Tuesday, January 27th** shutdown of all systems in 5100 J116 for a planned EPO test. Outage will start at 6 AM and last until around 10 AM.

Reminders

- Login excl.ornl.gov is not for computation or compiling. It is a VM configured for ssh and [thinlinc](https://excl.ornl.gov) access only. Running large scale work on this host affects all other ExCL users
- Please don't run [VSCode](https://excl.ornl.gov) on the login node either via Remote-SSH or [thinlinc](https://excl.ornl.gov). You can connect to a worker node via login as a jump host. Instructions at <https://docs.excl.ornl.gov/> – [Visual Studio Code | ExCL User Docs](https://docs.excl.ornl.gov/)
- Remember that snapshots going back some [time](https://excl.ornl.gov) are available in [~zfs/snapshot](https://excl.ornl.gov) (which will not show up on a normal directory walk). This is a read-only file system from which you can recover older versions of files. Snapshot names include date and time that snapshot was made
 - Hourly snapshots are kept for forty-eight hours
 - Daily (midnight or so) snapshots are kept for thirty days
 - Monthly snapshots are kept for six months
- There are now system specific notes in the boxed message of the day (MOTD) shown on every system. Suggestions about what other information might be presented per system (specialized hardware, etc) happily received.
- Remember to use the [news](https://excl.ornl.gov) command when you see new notices (which will show up in [yellow](https://excl.ornl.gov) in the MOTD). Outages, hardware changes and other issues will be found there. We keep this small and up to date.
- If you need solo access to a resource for performance measurement, or need additional storage space, do not hesitate to ask.
- Resource sharing is ad hoc so please reach out to coordinate usage.

ExCL Team: Steve Moulton and Aaron Young

OAK RIDGE
National Laboratory

PYX Discussion—Why pyx for ExCL?

Accelerating AI/HPC Research on ExCL

- **Current Pain Points:**
 - Manual package builds across heterogeneous nodes (GPUs, CUDA versions, OS).
 - Senior staff diverted to debugging; interns require training to onboard.
 - Fragile scaling risks publication timelines.
- **Impact:** New projects, increased ExCL usage, and AI/LLM focus will amplify these challenges.
- **Goal:** Reduce environment friction to speed up “time-to-science.”

PYX Discussion—pyx Introduction

Accelerating AI/HPC Research on *ExCL*

A Python-native package registry
from the creators of uv.

⚡ Fast

Speed up installs from PyPI, PyTorch, and your own private sources with optimized artifacts and uv-native metadata APIs. An order of magnitude faster than other private registries.

🔒 Secure

Create dedicated index URLs to filter packages by popularity, age, vulnerabilities, and more. Encode your own compliance rules and ensure reproducible builds on the server.

⭐ Modern

A singular focus on Python means best-in-class support for cutting-edge standards. Direct integration with uv means zero configuration and seamless authentication.

GPU-aware

Get the right, pre-built versions of PyTorch, vLLM, FlashAttention, DeepSpeed, and more — all with consistent metadata and optimal configuration — based on your hardware.

PYX Discussion—Proposed Solution & Outcomes

pyx: Private GPU-Aware Package Registry

- **Key Features:**
 - Auto-detect hardware; serve pre-built wheels (no manual compile) for specific hardware configuration.
 - Shared build cache: build once, reuse everywhere.
 - Resilient PyPI mirror: faster installs, full visibility, avoids download throttling.
 - Private PyPI repository for research packages. Speeds up sharing with collaborators and scaling-up experiments.
- **Expected Gains:**
 - **Time-to-Science:** 15+ min → <10 sec python environment startup time.
 - **Support Volume:** Reduced the number of support tickets as hardware is auto-detected by uv and the correct environment is served automatically.
 - **Focus:** Shift from maintenance to innovation.
- **Next Steps:** Ask around ORNL to gauge interest, and if there is interest, deploy a pilot in ExCL.
 - General pyx docs: [technical documentation](#)
 - GPU Index: [documentation around the GPU index specifically.](#)
 - Learn More: [Episode #520 - pyx - the other side of the uv coin \(announcing pyx\) | Talk Python To Me Podcast](#)

Preemptable Slurm Queues

What's new

- Added **nvidia-long** as a **preemptable GPU queue**
- Standard GPU work continues in **nvidia** (non-preemptable)

How it works

- nvidia-long runs when GPUs are idle
- Jobs **may be preempted** by higher-priority work in the nvidia queue
- Preempted jobs can **requeued and restarted**

User action required

- Jobs are **NOT requeue-able by default**
- **Must add --requeue** to enable restart after preemption
- Recommended: handle SIGTERM + checkpoint



Best use cases

- Long-running, checkpointed jobs
- Training, parameter sweeps, Monte Carlo

Why this matters

- Improves overall GPU utilization
- Enables opportunistic long jobs without impacting priority work

Documentation

- [Slurm | ExCL User Docs](#)

Other ongoing efforts

- ExCL microelectronics and neuromorphic laboratory (3700/A102)
 - Networking to be completed by 1/20. This has been the major blocker to moving resources there.
 - Networking will be on the ExCL VLAN.
 - Moderate/workstations wired and Wi-Fi available now.
 - More work chairs (stools) on order.
- GPUs to deploy (consumer grade but high power requirements)
 - Nvidia GeForce RTX 5090
 - AMD Radeon RX 9070 XT
- Working with vendor to procure additional AI capability



Questions/Projects/Comments/Discussions