

# Experimental Computing Laboratory Meeting August 22, 2024, 12:30 PM

Advanced Computing Systems Research Section

Jeffrey Vetter, Section Head

Steve Moulton, Systems Engineer

Aaron Young, Software Engineer

ORNL is managed by UT-Battelle LLC for the US Department of Energy



U.S. DEPARTMENT OF  
**ENERGY**

# Infrastructure/System Status

- Ubuntu 24.04 announced, but not approved for ORNL use. Will not start installations or upgrades until this happens.
  - Not considered stable; rolling upgrades generally broken.
  - Will be installed on request if needed
- login-new.ornl.gov in the testing phase to replace login.ornl.gov on Tuesday.
- Remaining Centos 7 systems awaiting migration:
  - Slurm – awaiting current workload completion
    - Will involve full ExCL reboot, as more recent versions use cgroups2
  - Checkmk – in progress; new agent deployments, odd problems

# Plans for next three months

- Continuing documentation improvements
  - <https://docs.excl.ornl.gov>
- Stabilize non-ORNL-managed UID & GID assignments ExCL-wide to simplify service deployment
  - Also has implications for docker groups
- Fully deprecate /noback directories
  - Removed from all backups, snaps, and mirrors
  - Contact [excl-help@ornl.gov](mailto:excl-help@ornl.gov) if you need assistance moving files (we know tricks).
  - Mount points will be removed October 1, 2024 (no deletions until November).
- Partial day outage – discussed later

# Systems in the pipeline

Order still pending shipping; vendor changed some part specifications. Integrator has been made aware of 9/24 EOFY deadline.

- 4 \* *Mi300A system*
  - *APUs (Integrated CPUs and GPUs)*
  - *128 GB HBM3 memory on each*
  - *Infinity Fabric Links (8)*
  - *PCIe 5.0x16 for peripherals (and who knows what else).*
  - *<https://www.amd.com/en/products/accelerators/instinct/mi300/mi300a.html>*
  - *“in the next few months”*

# Reminder: SSH Keys for Authentication

- Using SSH keys is the preferred way to authenticate your user and to authenticate with private Git repositories.
  - For security, it is recommended to use an SSH keys encrypted with a passphrase.
- Why not passwords?
  - ExCL will block your account after 3 failed attempts. Automatic login tools, e.g. VS Code, can easily exceed this limit using a cached password and auto-reconnect.
  - For git repos with two-factor authentication, an application token/password must be created, and this password must be stored externally and is more cumbersome to use.
- How to get started?
  - Set up a key pair:
    - [Visual Studio Code Remote Development Troubleshooting Tips and Tricks](#)
    - [Generating a new SSH key and adding it to the ssh-agent - GitHub Docs](#)
  - [Add key to Git Hosting Websites.](#)
  - [Setup ExCL worker node proxy via login node.](#)

# Proposed Host layout Change: Milan0

- Currently has two Nvidia A100 GPUs
- Proposal
  - Move 2 AMD M100 GPUs from Cousteau to Milan0
    - Cousteau would then become general compute server
  - Deploy Xilinx U280 in Milan0
  - Requested Intel ARC GPU
- This would provide a nice heterogeneous environment, as well as moving M100s to Pcie5 (which they may or may not be able to take full advantage of).
- Comments?

# Proposed Outage Tuesday August 27 2024 9 AM to 1 PM

- Reconfigure all kernels to do cgroup v2 only
- Deploy updated Slurm configurations, move Slurm controller to Rocky9
- Deploy new login server with new slurm client
  - There will be things missing that you expect; let us know
- Update and reboot file servers and hypervisors
  - Nvhpc, nvidia drivers and rocm updates
- Outage planned for four hours starting at 9 AM; may take longer if problems or shorter if things are easier than expected.
- *Speak now ...*

# Updated Vitis Deployment and Documentation

Vitis FPGA Development | ExCL User Docs (ornl.gov)

ExCL User Docs

Home About

Search

Introduction

Acknowledgment

System Overview

EXCL SUPPORT

ExCL Team

Frequently Encountered Problems

Access to ExCL

Contributing

Glossary & Acronyms

Requesting Access

Outages and Maintenance Policy

Backup & Storage

QUICK-START GUIDES

ThinLinc

Conda and Spack Installation

Jupyter Notebook

Vitis FPGA Development

Visual Studio Code

Groq

Devdocs

Apptainer

Siemens EDA

Gitlab CI

GitHub CI

SOFTWARE

Compilers

MPI

Git

ExCL DevOps: CI/CD

Vitis FPGA Development

Getting started with Vitis FPGA development.

ExCL → User Documentation → Vitis FPGA Development

FPGA Current State

FPGA	State
U250	Attached to <code>spike</code> in Alveo mode.
u55C	Attached to <code>spike</code> in Alveo mode.
u280	Aaron's office

Vitis Development Tools

This page covers how to access the Vitis development tools available in ExCL. The available FPGAs are listed in the [FPGAs](#) section. All Ubuntu 22.04 systems can load the Vitis/Vivado development tools as a module. See [Quickstart](#) to get started. The [virtual systems](#) have [ThinLinc](#) installed, which makes it easier to run graphical applications. See section [Accessing ThinLinc](#) to get started.

Vitis is now primarily deployed as a module for Ubuntu 22.04 systems. You can view available modules and versions with `module avail` and load the most recent version with `module load Vitis`. These modules should be able to work on any Ubuntu 22.04 system in ExCL.

FPGAs	Host System	Slurm GRES Name
<a href="#">Alveo U250</a>	Spike	U250
<a href="#">Alveo U55C</a>	Spike	U55C
<a href="#">Alveo U280</a>		U280

FPGA Current State

Vitis Development Tools

FPGAs

Vitis and FPGA Allocation with Slurm (Recommended Method to Use Tools)

Quickstart

First Steps

Getting specific FPGA information from the Platform.

Accessing systems graphically using ThinLinc

Using Vitis with the Fish Shell (Recommended Approach)

Manually Setting up License

Building and Running FPGA Applications

Setting up the Vitis Environment

Build Targets

Building the Host Program

Building the Device Binary

Analyzing the Build Results

Running Emulation

Running the Application

Hardware Build

Example Makefile


Performance Considerations

Useful References

Useful Commands

Was this helpful?

Edit on GitHub





# IRIS Setup scripts

- Setup scripts for ExCL systems that setup the environment and tools for IRIS accelerator use are available at `/auto/software/iris` and used via:

```
source /auto/software/iris/setup_system.source
```

- To request changes to these scripts, fork the repo [Brisbane / ExCL IRIS Setup Scripts](#), make the modifications, and open a merge request. Once we review the merge request, we will merge the changes and deploy the scripts to ExCL.

# Questions/Discussions/Comments?

- Aaron has methodology to pass SSH credentials into Apptainer containers without copying private keys. Let [excl-help@ornl.gov](mailto:excl-help@ornl.gov) know if you have needs or questions.
- Discussion of github-runner linger; Brett W has solution documented. Aaron and Steve to look at deployment in specific ExCL systems.