

Experimental Computing Laboratory Meeting July 25 , 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.
- Infrastructure moves to new hypervisor continuing
 - Part of Centos7 deprecation. Support ended June 30, 2024. This will continue to consume considerable SE effort.
 - Migration of all infrastructure systems from Centos 7 to Rocky 9.
- login-new.ornl.gov in the testing phase to replace login.ornl.gov after testing is completed.

Plans for next three months

- Continuing documentation improvements
 - <https://docs.excl.ornl.gov>
- Looking into acquiring Grace Hopper system (stalled on vendor)
- Stabilize non-ORNL-managed UID & GID assignments ExCL-wide to simplify service deployment
 - Also has implications for docker groups
- Fully deprecate /noback directories
 - Will be removing from all backups, snaps, and mirrors
 - Contact excl-help@ornl.gov if you need assistance moving files (we know tricks).

Systems in the pipeline (1)

~~• 2 * H100 System~~

- ~~– Order recently changed to H100 NVL (64GB to 96GB GPUs; 64GB no longer available) at no additional cost~~
- ~~– 1TB main memory~~
- ~~– By end of fiscal year~~
- ~~– https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/h100/PB-11773-001_v01.pdf~~

Hudson.ftpn.ornl.gov deployed and is in research use

Some instabilities in gpu0 have been noted (causing failure); cold restarting system has addressed that. Please apprise excl-help@ornl.gov of any challenges that system is presenting – if we see repeated issues on different code bases we will submit a ticket to nvidia.

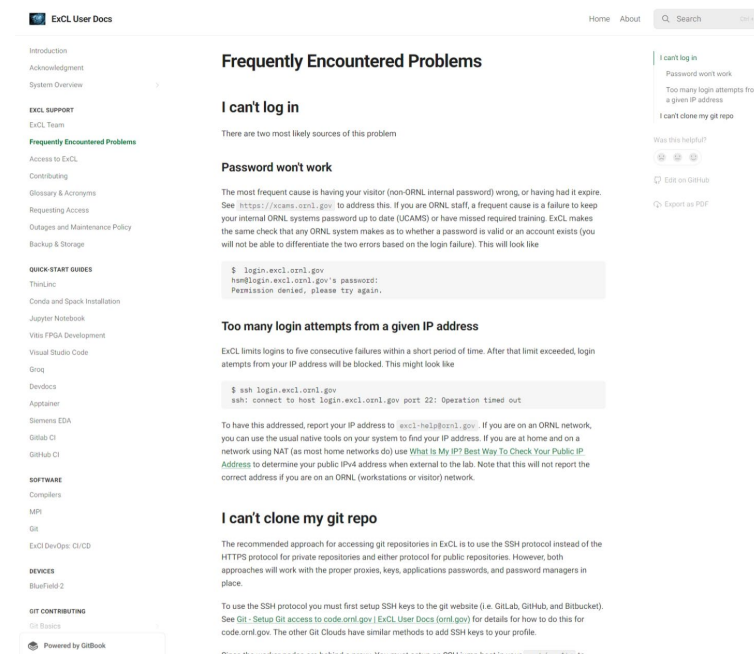
Systems in the pipeline (2)

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”*

Documentation

- A guide for “which hosts to use for what” is in the works, and will be made prominent in docs.excl.ornl.gov when completed.
 - If you are interested in contributing to this list, let me know.
- Frequently encountered issues page added <https://docs.excl.ornl.gov/excl-support/frequentproblems>
 - I can't log in
 - I can't clone my git repo



ZFS/NFS Storage Quotas. See [Backup & Storage](#). Updated!

- We have added quotas to the filesystems in ZFS to avoid runaway storage usage.
- Quotas in ZFS are easy to view and set at any level in the filesystem.
- The quota applies to **your files only, not** including snapshots of the data.
This is changed by using refquotas instead of quotas.
- ZFS is smart and only stores changes in data. It also stores data in a compressed format, so your disk usage is less than your apparent amount.
- Home and project directories start with 512G per directory, and higher quotas can be requested via excl-help@ornl.gov. We can also help by giving a breakdown of file usage and helping clean up large usages.
- Use of /scratch/\$USER for large build artifacts that are local to a node and do not require snapshotting can reduce ZFS storage usage.

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.](#)

SSH-Agent and SSH Forwarding

- SSH-Agents cache SSH keys with passphrases, allowing them to be reused during the session.
 - This is not needed with keys without a passphrase, since they can be used without decrypting.
- SSH Forwarding: SSH agents can forward SSH keys to a remote system, making the keys available there as well.
- How to get started?
 - [Set up an SSH-Agent.](#)
 - Add key to agent
 - ``ssh-add`` or ``ssh-add [file]`` for non-default filenames.
 - Check loaded keys with ``ssh-add -l``.
 - Setup SSH forwarding in ssh config.

Host *
ForwardAgent yes
 - Log in and verify key is still available.
 - Note: Do not launch an SSH-agent on the remote system when using SSH Forwarding, as the new agent will hide the forwarded keys.

Zenith 2 and Embedded FPGA Development



- Embedded FPGA Development Boards
 - Arty-A7 100T
 - Alchitry
 - Polarfire SoC Icicle Kit
- Integrated and reservable via Slurm with helpful fish functions to reserve boards.

```
7ry@zenith2:~$ fish
Welcome to 🐟, the friendly interactive shell
7ry@zenith2:~$ sfpgarun-icebox fish
Welcome to 🐟, the friendly interactive shell
7ry@zenith2:~(fpgarun-icebox)$ echo $$SLURM_FPGA_FILE
/dev/icebox
7ry@zenith2:~(fpgarun-icebox)$ tio $$SLURM_FPGA_FILE
[tio 11:33:12] tio v1.32
[tio 11:33:12] Press ctrl-t q to quit
[tio 11:33:12] Connected
```

```
icicle-kit-es login:
[tio 11:33:16] Disconnected
7ry@zenith2:~(fpgarun-icebox)$ exit
7ry@zenith2:~$
```

```
7ry@zenith2:~$ sinfo -o 'PartitionName,NodeList:20,Gres:50,GresUsed:90' --partition zenith,milan,gpu
```

PARTITION	NODELIST	GRES	GRES_USED
milan	milan1	groq:card:1	gpu:0,fpga:0,groq:card:0
milan	milan0	gpu:A100:2(S:0)	gpu:A100:0(Idx:N/A),fpga:0,groq:0
milan	milan3	(null)	gpu:0,fpga:0,groq:0
gpu	zenith	gpu:RTX3090:1,gpu:RD6900:1	gpu:RTX3090:0(Idx:N/A),gpu:RD6900:0(Idx:N/A),fpga:0,groq:0
zenith	zenith2	fpga:arty-a7:1,fpga:alchitry:1,fpga:icebox:1	gpu:0,fpga:arty-a7:1(Idx:0),fpga:alchitry:0(Idx:N/A),fpga:icebox:1(Idx:2),groq:0
zenith	zenith	gpu:RTX3090:1,gpu:RD6900:1	gpu:RTX3090:0(Idx:N/A),gpu:RD6900:0(Idx:N/A),fpga:0,groq:0



Questions/Discussions/Comments?

- IRIS setup scripts gitlab repository for maintenance of scripts

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.