



AMD Presentation Roadmap and Introduction to System

Bob Robey
Data Center GPUs, AMD
Oct 21, 2025

Training Workshop Guidelines

– help us create a positive environment

Please help us create a positive collaborative environment for this event. Our top goal is to attract and build the community in high performance computing. Please consider not only the technical aspects of your interactions, but also how it affects other participants.

- Give equal time to all participants
- Please, help others with technical issues and without negative comments
- And most importantly, respect all participants, regardless of background, language, culture

Introduction to Local System and AMD Accelerator Cloud (AAC) for Exercises

Contacts for System Issues -- AAC

- Please contact either Bob Robey Bob.Robey@amd.com or Giacomo Capodaglio Giacomo.Capodaglio@amd.com about any issues with the training system during the workshop.
- Also, can use the chat and send a private message to either of us or to one of the other presenters for questions about specific topics.

AAC notes

Account requests

- We'll be using ssh public keys for access to the AAC system
 - Need ssh public keys, preferred username, email, address (business), organization name
 - Information will only be used to create accounts and for account issues
- To generate a ssh key
 - `ssh-keygen -t ed25519`
 - Hit return and set password if preferred
 - `cat $HOME/.ssh/id_ed25519.pub`
 - Send the string for creating account access
 - You can do this from the terminal in Windows®, from wsl, from Linux® or MacOS
 - There are apps (particularly Windows) that claim to make this easier, but partially obscure the process
- Accounts will stay active for some time after the workshop for participants to try the exercises when it is convenient

Logging into AAC

- Accessing AAC
- `ssh <USERNAME>@aac6.amd.com -i <identityfile>`
 - -i option only needed if not using default ssh key
 - Username only needed if different than local username
- If prompted for password, type "Ctrl-C" and contact us for assistance
 - If using a password, be very careful. Three failures will put you IP address in a blacklist.
- To simplify the login to AAC, you can add the following to your `.ssh/config` file.
- # AAC
 - Host aac
 - User <USERNAME>
 - Hostname aac1.amd.com
 - #IdentityFile <path>/<private key file> # customize location of key file
 - IdentityFile id_ed25519 # file is in .ssh directory
 - ServerAliveInterval 600
 - ServerAliveCountMax 30
- The `ServerAlive*` lines in the config file may be added to avoid timeouts when idle.
- Now you can login with `ssh aac -p <###>`

File copying

You should be able to copy files in or out with the `scp` command.

- In: `scp -i <path>/<keyfile> <file> USER@aac6.amd.com:~/<path>/<file>`
- Out: `scp -i <path>/<keyfile> USER@aac6.amd.com:~/path/to/your/file ./`

You can also use `rsync` command

- `rsync -avz -e "ssh -i <path>/<keyfile>" <file> <USER>@aac6.amd.com:~/path/to/your/files`

AAC Environment

- AMD Accelerator Cloud (AAC) Model Installation Environment
 - Ubuntu® 24.04
 - ROCm version 6.4.1
 - GCC is version 13
 - Clang/LLVM is version 19

- Modules using Lua modules
 - Hierarchical module list
 - Initial set on initial login

```
----- /shared/apps/modules/ubuntu/lmodfiles/base -----  
clang/base      miniconda3/25.3.1  rocm/6.4.1  
gcc/base        miniforge3/24.9.0  turbovnc/3.0.3
```

- rocm/6.4.1 -- ROCm software stack including hip and hip libraries
- Select rocm/6.4.1 for list of software built and dependent on it

Environment -- modules

Modules using Lua modules

- ROCm

amdclang/19.0.0-6.4.1	rocprofiler-compute/6.4.1 (D)
amdflang-new/rocm-afar-7.0.5	rocprofiler-sdk/6.4.1
hipfort/6.4.1	rocprofiler-systems/6.4.1 (D)
opencl/6.4.1	

- amdclang/19.0.0-6.4.1 -- OpenMP® compiler
- amdflang-new/rocm-afar-7.0.5
- hipfort/6.4.1-- Fortran wrappers for hip calls
- opencl/6.4.1 – OpenCL™ language
- rocprofiler-compute/6.4.1 Kernel profiler, built-in version
- rocprofiler-systems/6.4.1 Timeline profiler, built-in version
- rocprofiler-sdk/6.4.1 -- Trace decoder and profiler software development toolkit

Environment – ROCmPlus modules

- ROCmPlus-6.4.1

adios/2.10.1

cupy/13.6.0

fftw/3.3.10

ftorch/dev

hdf5/1.14.6

hip-python/13.6.0

hipfort_from_source/6.4.1

hipifly/dev

hpctoolkit/2024.01.99-next

hypre/2.33.0

jax/0.6.0

kokkos/4.6.01

mpi4py/4.0.3

netcdf-c/4.9.3

netcdf-fortran/4.6.2

openmpi/5.0.7-ucc1.4.4-ucx1.18.1

petsc/3.23.1

pytorch/2.7.1

rocprofiler-compute/develop

rocprofiler-systems/develop

scorep/9.0

tau/dev

tensorflow/merge-250318

Environment -- modules

Modules using Lua modules

- ROCm
- [amdflang-new](#)

AMD is working on a Fortran compiler and soon will include it in ROCm. During the training, **Participants will have hands-on access to the beta**. The beta itself contains early access to open-sourced pre-release features. Going forward, AMD is able to provide continued early access via AFARs (Advanced Feature Access Releases) as improvements and features are added to the beta. As with any beta, we looking to gather feedback on functionality, usability and user experience.

The latest AFARs containing the next generation Fortran compiler also contains later versions of the C and C++ compiler, much of a later ROCm version and profilers.

PRE-PRODUCTION SOFTWARE: The software accessible on this page may be a pre-production version, intended to provide advance access to features that may or may not eventually be included into production version of the software. Accordingly, pre-production software may not be fully functional, may contain errors, and may have reduced or different security, privacy, accessibility, availability, and reliability standards relative to production versions of the software. Use of pre-production software may result in unexpected results, loss of data, project delays or other unpredictable damage or loss. Pre-production software is not intended for use in production, and your use of pre-production software is at your own risk.

Environment -- modules

Modules using Lua modules

Compiler modules set the CC, CXX, FC flags. Only one compiler module can be loaded at a time. hipcc is in the path when the rocm module is loaded. These modules also set the OMPI_CC/CXX/FC environment variables so that the compiler wrappers use the currently set compiler.

Module commands

- module avail
- module load <package>
- module unload <package>

Environment -- Slurm

Slurm configuration with single queue LocalQ

sinfo

PARTITION	AVAIL	TIMELIMIT	NODES	STATE	NODELIST
1CN192C4G1H_MI300A_Ubuntu22*	up	8-00:00:00	5	idle	ppac-pl1-s24-[16,26,30,35],ppac-pl1-s25-40
1CN48C1G1H_MI300A_Ubuntu22	up	8-00:00:00	5	idle	sh5-pl1-s12-[09,12,15,33,36]

25 total GPUs

1CN192C4G1H_MI300A_Ubuntu22 - Node with four MI300As

1CN48C1G1H_MI300A_Ubuntu22 - Node with one MI300A

Slurm commands

- salloc can be used to schedule a long-term interactive session
 - sbatch is used to submit a job to the batch queue
 - squeue will show the jobs running in the batch queue
-
- Please allocate only the GPUs that you need and release them after you are done

AAC6 – Total 25 GPUs – your local system may be different

AAC6 is all pre-production parts

Login or Control Node

AAC6:
pl1vm1mi300ctl01

Slurm

salloc
srun
sbatch

Compute Nodes

Queue 1CN192C4G1H_MI300A_Ubuntu22

ppac-pl1-s24-16

ppac-pl1-s24-26

ppac-pl1-s24-30

ppac-pl1-s24-35

Queue 1CN48C1G1H_MI300A_Ubuntu22

sh5-pl1-s12-09

sh5-pl1-s12-12

sh5-pl1-s12-15

sh5-pl1-s12-33

4 MI300A
192 CPUs
512 GB memory

ppac-pl1-s24-16

1 MI300A
48 CPUs
128 GB memory

sh5-pl1-s12-36

Sample slurm commands

- `sinfo`
- `squeue -u <username>`
- `scancel <jobid>`

Most work should be done on the compute nodes. Examples.

- `srun -p 1CN48C1G1H_MI300A_Ubuntu22 --exclusive my_batch_script.batch`
- `srun -p 1CN48C1G1H_MI300A_Ubuntu22 --ntasks=16 my_batch_script.batch`
- `srun -p 1CN48C1G1H_MI300A_Ubuntu22 --ngpus=1 --ntasks=4 my_batch_script.batch`
- `salloc --ntasks 16 --gpus=1 -t`
- `salloc --nodelist=ppac-pl1-s24-30 --ntasks 1 --gpus=1 --time=01:00:00ime=01:00:00`
- `salloc -N 1 --gpus=4 --cpus-per-gpu=12 --time=02:00:00`
- `sbatch my_batch_script.batch`

Sample slurm batch file

- At the top of the batch file
- Bash script commands follow the #SBATCH section

```
#!/bin/bash
#SBATCH --partition=1CN192C4G1H_MI300A_Ubuntu22
#SBATCH --gpus=1
#SBATCH --ntasks=1
#SBATCH --cpus-per-task=1
#SBATCH --time=01:00:00
#SBATCH -- output=my_job.out
#SBATCH -- error=my_job.out
```

Installed software

- emacs
- vim
- autotools
- cmake
- tmux
- boost
- eigen
- fftw
- gmp
- gsl
- hdf5-openmpi
- lapack
- magma
- matplotlib
- parmetis
- mpfr
- mpi4py
- openblas
- openssl
- swig
- numpy
- scipy
- h5sparse

Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2025 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ROCm, Radeon and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

The OpenMP name and the OpenMP logo are registered trademarks of the OpenMP Architecture Review Board

Ubuntu and the Ubuntu logo are registered trademarks of Canonical® Ltd.

Canonical and the Canonical logo are registered trademarks of Canonical Ltd.

OpenCL is a trademark of Apple Inc. used by permission by Khronos Group, Inc.

Windows is a registered trademark of Microsoft Corporation in the US and/or other countries.

Microsoft is a registered trademark of Microsoft Corporation in the US and/or other countries.

