

# AAMU High-Performance Computing Technical Support Center

**Author:** BenchChem Technical Support Team. **Date:** December 2025

## *Compound of Interest*

Compound Name: AAMU

Cat. No.: B011769

[Get Quote](#)

Welcome to the **AAMU** High-Performance Computing (HPC) Technical Support Center. This guide is designed for researchers, scientists, and drug development professionals to troubleshoot common issues encountered during their computational experiments. The information is presented in a question-and-answer format to directly address specific problems.

## Quick Links

- --INVALID-LINK--
- --INVALID-LINK--
- --INVALID-LINK--
- --INVALID-LINK--

## Job Submission & Scheduling

This section addresses issues related to submitting your jobs to the HPC cluster and managing them within the scheduling system.

**Q1:** My job submission failed immediately. What should I check first?

**A1:** Immediate job submission failures are often due to errors in your submission script or incorrect resource requests. Here's a systematic approach to troubleshoot:

- Check Submission Script Syntax: Ensure that all scheduler directives (e.g., #SBATCH for Slurm or #PBS for PBS) are correctly formatted. Common mistakes include typos and incorrect syntax.
- Verify Resource Requests: Make sure your requests for resources like walltime, memory, and number of nodes/cores are within the limits of the partition or queue you are using.[\[1\]](#)[\[2\]](#)[\[3\]](#)
- Check Account and Permissions: Confirm that you are using the correct project or account allocation and that you have the necessary permissions to submit to the requested queue.[\[1\]](#)[\[2\]](#)[\[3\]](#)
- File Paths: Double-check that the paths to your executable and input/output files are correct.[\[1\]](#)
- Use Verification Tools: Some schedulers offer a dry-run or verification option to check your submission script for errors without actually submitting the job.[\[1\]](#)[\[2\]](#)

Q2: My job is stuck in the queue and not starting. Why is it pending?

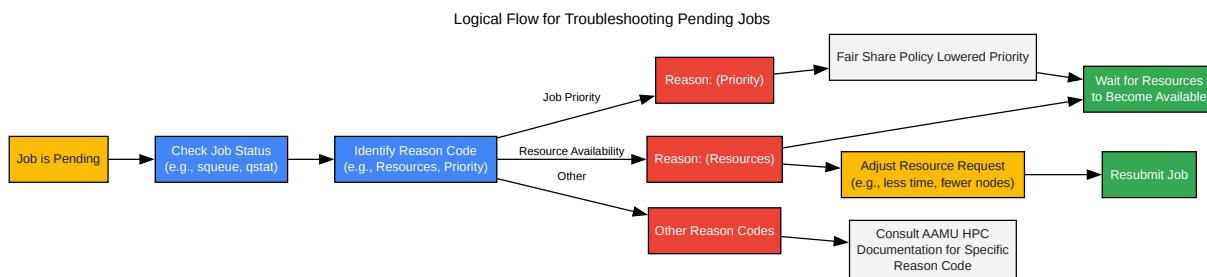
A2: Jobs can remain in a pending or queued state for several reasons. The status is often due to resource availability or scheduling policies.[\[4\]](#)[\[5\]](#)[\[6\]](#)

- High Cluster Load: The cluster may be busy, and there are not enough available resources (nodes, cores, memory, or GPUs) to start your job.[\[5\]](#)[\[6\]](#)
- Resource Request Size: Very large resource requests (e.g., many nodes or a very long walltime) can lead to longer queue times as the scheduler waits for those resources to become free simultaneously.[\[3\]](#)[\[4\]](#)[\[5\]](#)
- Fair Share Policies: HPC systems often use a "fair share" policy to ensure equitable access for all users. If you have recently run many jobs, your priority may be temporarily lowered.[\[4\]](#)
- Partition/Queue Limits: You may have reached a personal or group limit on the number of concurrently running jobs or total cores in use.[\[7\]](#)

- Incorrect Resource Specification: Requesting more resources than are available on any single node (e.g., more memory than any node possesses) will cause the job to pend indefinitely.[4]

To diagnose the reason, use the scheduler's command to check the status and reason for your pending job (e.g., squeue --start in Slurm).

### Troubleshooting Pending Jobs



[Click to download full resolution via product page](#)

Caption: Troubleshooting workflow for a pending HPC job.

## Runtime Errors & Debugging

This section covers common errors that occur while your job is running and provides guidance on debugging.

Q3: My job failed with an "out of memory" error. What should I do?

A3: "Out of memory" errors are common in HPC and occur when your job attempts to use more memory than you allocated in your submission script.[7][8]

Experimental Protocol for Resolving Memory Issues:

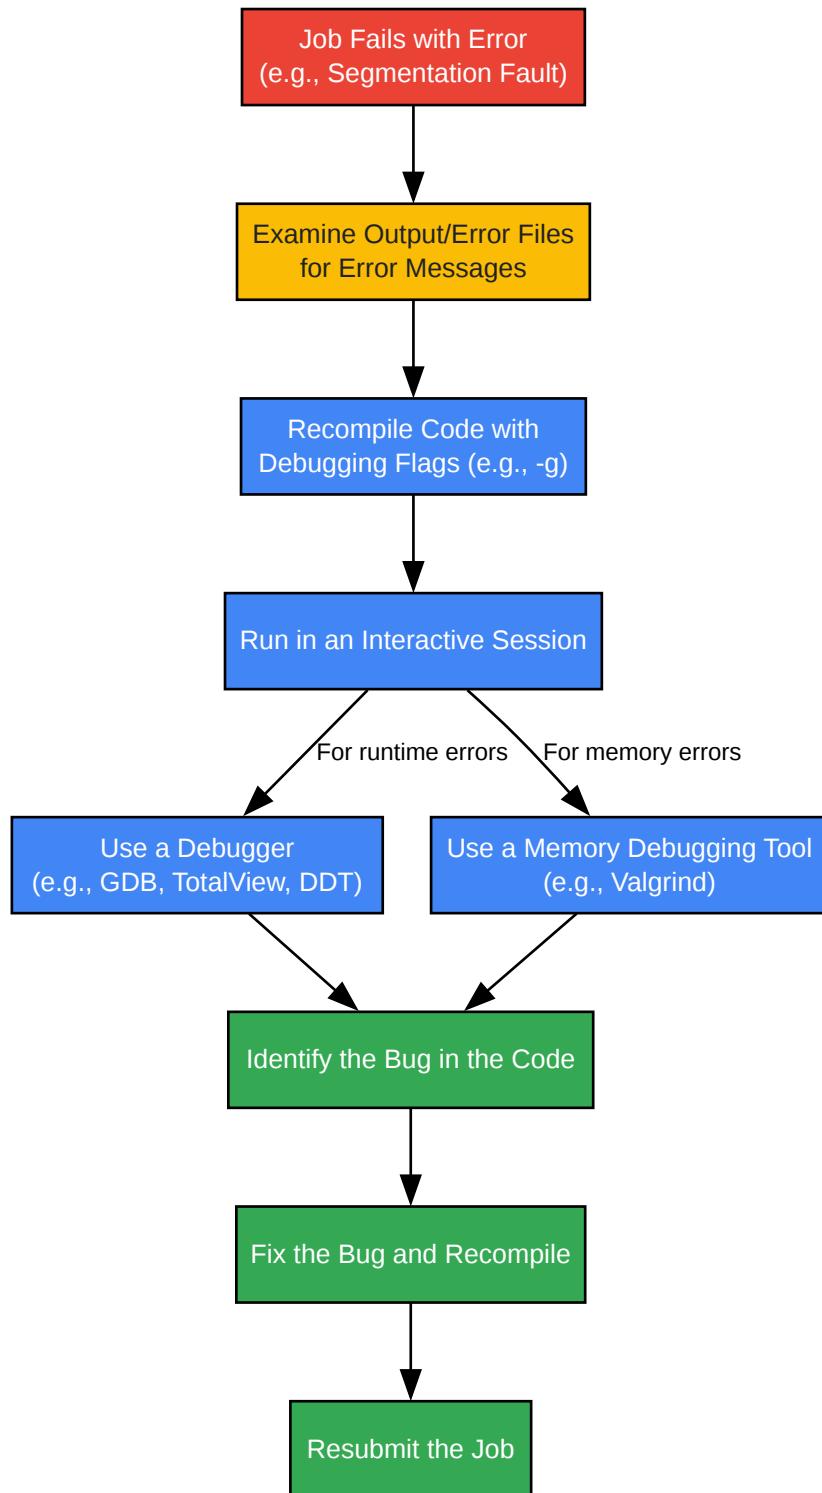
- Analyze Memory Usage: Use the scheduler's accounting tools (e.g., sacct or qacct) to check the actual memory used by your failed job.[\[1\]](#) Many HPC centers also provide a command like seff to analyze the efficiency of a completed job.[\[7\]](#)
- Increase Memory Request: Re-submit your job with a higher memory request. It is good practice to request slightly more memory than the previous run to provide a buffer.
- Optimize Your Code: If your memory usage is unexpectedly high, profile your code to identify memory leaks or inefficient data structures.[\[9\]](#)[\[10\]](#)[\[11\]](#) Tools like Valgrind can help detect memory errors and leaks in C/C++ applications.[\[9\]](#)[\[11\]](#)
- Parallelism Strategy: For distributed memory (MPI) jobs, you can often solve memory issues by increasing the number of nodes while keeping the problem size per node the same. This distributes the memory load across more machines.

Q4: My job terminated with a "segmentation fault" or other unexpected error. How can I debug this?

A4: A segmentation fault occurs when your program tries to access a memory location that it's not allowed to access.[\[9\]](#) This is a common bug in C/C++ and Fortran codes.

#### Debugging Workflow for Application Crashes

## Debugging Workflow for Application Crashes

[Click to download full resolution via product page](#)

Caption: A systematic workflow for debugging crashing HPC jobs.

## Debugging Tools Comparison

Tool	Primary Use Case	Key Features
GDB	Serial code debugging	Command-line, step-through execution, inspect variables.
Valgrind (Memcheck)	Memory error detection	Detects memory leaks, out-of-bounds access, use of uninitialized memory. <a href="#">[9]</a> <a href="#">[11]</a>
TotalView/DDT	Parallel (MPI/OpenMP) debugging	GUI-based, scalable to many processes, advanced features like reverse debugging. <a href="#">[9]</a> <a href="#">[12]</a>
STAT/ATP	Scalable stack trace analysis	Gathers stack traces from all processes of a parallel application to identify hangs or crashes. <a href="#">[9]</a>

## Software & Environment Management

Properly managing your software environment is crucial for reproducible research. Most HPC systems use environment modules to manage different software packages and versions.[\[13\]](#)[\[14\]](#)[\[15\]](#)

Q5: How do I find and load the software I need?

A5: Use the module command to manage your software environment.

- Find Available Software: To see all available software modules, use the command:

To search for a specific package (e.g., python), use:

- Load a Module: To load the default version of a software package, use:[\[16\]](#)

For a specific version, which is highly recommended for reproducibility, use:[\[16\]](#)

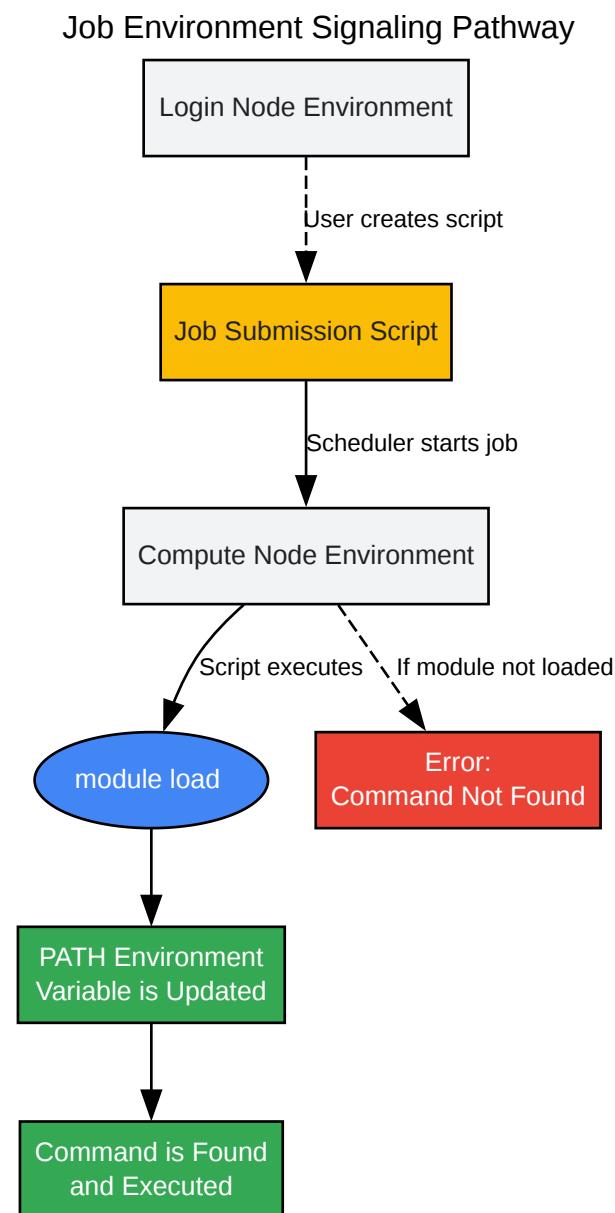
- List Loaded Modules: To see which modules are currently loaded in your environment, use:

- Unload All Modules: To unload all currently loaded modules, use:

Q6: My job script fails with a "command not found" error, but the command works when I'm logged in.

A6: This typically happens because the necessary software module is not loaded within the job script's environment. [17] The environment on a compute node is separate from your login shell environment.

Solution: Always include the necessary module load commands at the beginning of your job submission script, before you call any executables from that software package. [16] Signaling Pathway for a Job's Software Environment



[Click to download full resolution via product page](#)

Caption: How a job script sets up the software environment on a compute node.

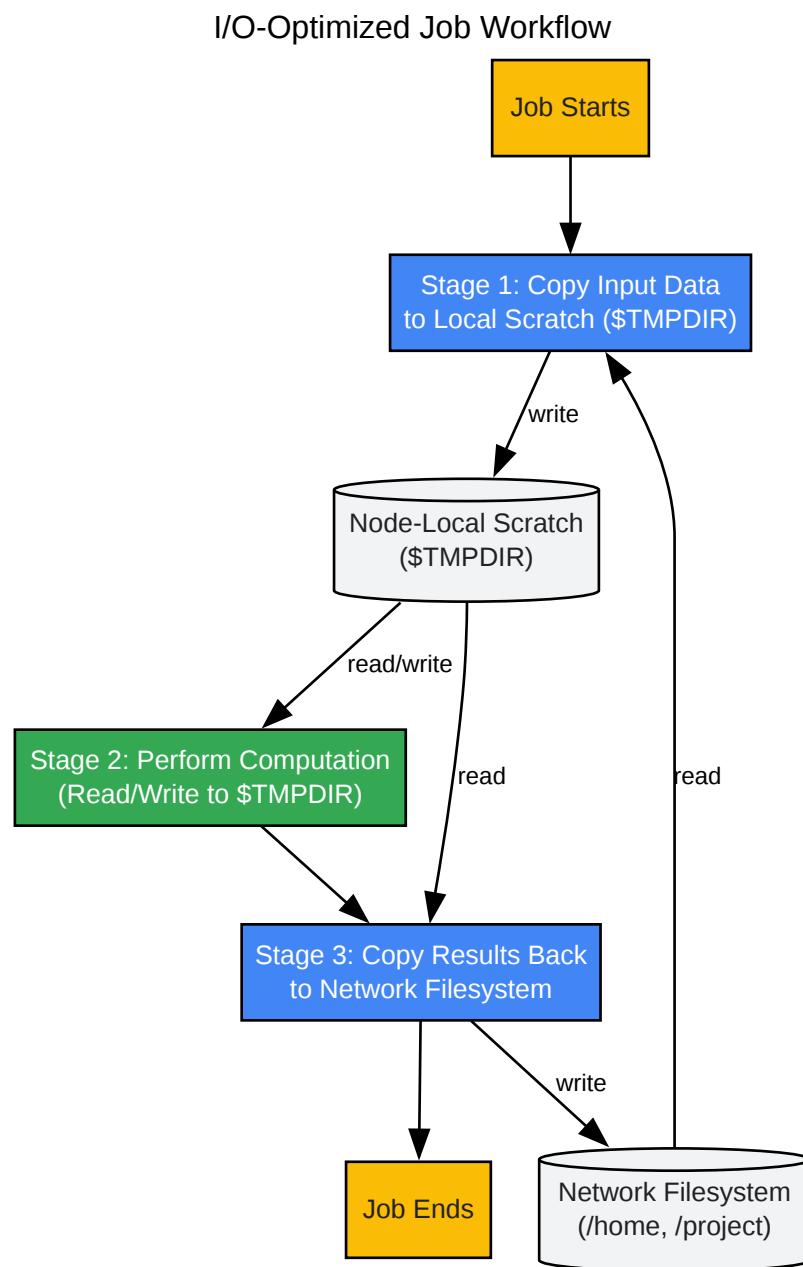
## File & Data Management

Efficiently managing your files and data I/O (Input/Output) is critical for performance on an HPC cluster.

Q7: My job is running very slowly, and I suspect it's a file I/O issue. How can I confirm and fix this?

A7: Poor I/O performance is a common bottleneck, especially for jobs that read or write many small files or large amounts of data. [\[18\]](#)[\[19\]](#)[\[20\]](#) Protocol for Diagnosing and Optimizing I/O:

- Use Temporary Storage: Whenever possible, perform intensive I/O operations on the local scratch or temporary storage of the compute node (\$TMPDIR or /tmp). This is much faster than writing to a shared network file system. Copy your input files to the temporary directory at the beginning of your job and copy the results back to your home or project directory at the end. [\[21\]](#)2. Avoid Many Small Files: Accessing thousands of small files creates significant metadata overhead on the file system. If possible, aggregate your data into fewer, larger files. [\[21\]](#)Libraries like HDF5 or NetCDF are designed for efficient parallel I/O to structured data files.
- Profile I/O Operations: Use tools like strace to trace the file system calls your application is making. [\[21\]](#)This can reveal if your application is opening and closing files unnecessarily or performing many small read/write operations.
- Check Filesystem Quotas: Ensure you have not exceeded your disk space or file count (inode) quotas in your home or project directories. A full quota can cause jobs to fail or behave unexpectedly. [\[1\]](#)[\[3\]](#) Data Flow for an I/O-Optimized Job



[Click to download full resolution via product page](#)

Caption: Optimal data staging workflow to improve job I/O performance.

#### Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: [info@benchchem.com](mailto:info@benchchem.com) or [Request Quote Online](#).

## References

- 1. Troubleshooting jobs on HPC — ARC Documentation [arcdocs.leeds.ac.uk]
- 2. Debugging jobs - HPC @ QMUL [docs.hpc.qmul.ac.uk]
- 3. Troubleshooting Jobs - NERSC Documentation [docs.nersc.gov]
- 4. Job Troubleshooting & FAQs - FSU Research Computing Center Documentation [docs.rcc.fsu.edu]
- 5. Problems Running Jobs [hpcc.umd.edu]
- 6. 3. FAQ: Troubleshooting compute jobs — VUB-HPC [hpc.vub.be]
- 7. Slurm Errors [hpcc.umd.edu]
- 8. Troubleshooting Batch Problems | Ohio Supercomputer Center [osc.edu]
- 9. Debugging Tools Overview - NERSC Documentation [docs.nersc.gov]
- 10. perforce.com [perforce.com]
- 11. Valgrind and Valgrind4hpc - NERSC Documentation [docs.nersc.gov]
- 12. perforce.com [perforce.com]
- 13. Loading software packages on HPC — Home Page [coding-for-reproducible-research.github.io]
- 14. Introduction to High-Performance Computing: Accessing software [epcced.github.io]
- 15. Accessing software – Introduction to High-Performance Computing [ofisette.github.io]
- 16. Software Environment Setup | User Guides | High Performance Computing Center | Texas Tech [depts.ttu.edu]
- 17. Troubleshooting on Sapelo2 - Research Computing Center Wiki [wiki.gacrc.uga.edu]
- 18. youtube.com [youtube.com]
- 19. src.acm.org [src.acm.org]
- 20. datacore.com [datacore.com]
- 21. Optimizing IO Performance :: Documentation for HPC [docs.hpc.gwdg.de]
- To cite this document: BenchChem. [AAMU High-Performance Computing Technical Support Center]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b011769#troubleshooting-high-performance-computing-jobs-at-aamu]

**Disclaimer & Data Validity:**

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

**Technical Support:** The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [\[Contact our Ph.D. Support Team for a compatibility check\]](#)

**Need Industrial/Bulk Grade?** [Request Custom Synthesis Quote](#)

## BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

### Contact

Address: 3281 E Guasti Rd  
Ontario, CA 91761, United States  
Phone: (601) 213-4426  
Email: [info@benchchem.com](mailto:info@benchchem.com)