


Guidance for Computational Fluid Dynamics for Aircraft Design on AWS

AWS ParallelCluster is an open source cluster management tool and makes it easy to setup and manage a high performance computing (HPC) cluster on AWS. Here, it is used for computational fluid dynamics (CFD) simulations.



- 1** **AWS ParallelCluster** deploys the HeadNode string instance type and Compute queues (visualization, compute, high memory).
- 2** Connect to the HeadNode through Secure Shell (SSH) or **Amazon DCV**. The HeadNode is responsible for running Slurm scheduler and giving you an environment to submit jobs.
- 3** When a job is submitted, Slurm scales up the appropriate number of instances in a queue with the given CPU and Memory.
- 4** You can define a queue for visualization. This is useful for model preparation and meshing. Meshing is a process of making a 2D or 3D grid to analyze it with simulation. Typically these instances have graphics processing units (GPU's) attached such as the G4dn family.
- 5** To run the solve portion of Computational fluid dynamics (CFD), we recommend using an Elastic Fabric Adapter (EFA) enabled instance such as the hpc6a.48xlarge. This queue is setup with EFA enabled and is in a placement group, ensuring low latency inter-node communication.
- 6** For some CFD models, higher memory instances are required such as the hpc6id.32xlarge which has 64 cores and 1024 GB of memory. These can run on a high-memory queue.
- 7** **Amazon FSx for Lustre** provides a shared parallel filesystem (/shared) that allows instances to work in parallel. We recommend **FSx for Lustre** as the scratch filesystem. **Amazon Elastic File System** (Amazon EFS) is used to store CFD applications (/apps) such as Siemens Simcenter STAR-CCM+, Ansys Fluent, or OpenFOAM.
- 8** Long term storage of the model results is accomplished using **Amazon Simple Storage Service** (Amazon S3). The results can be called back at any point for future analysis. Data is synced automatically between **FSx for Lustre** and **Amazon S3** through a Data Repository Association (DRA).

