



Contribution ID: 55

Type: **not specified**

Can cloud computing accelerate transition of research to operations for workflows with cycling data assimilation?

Thursday, 23 September 2021 15:30 (20 minutes)

In numerical weather prediction, transition of research to operations often requires extensive computational experiments with cycling data assimilation, where short (6-12 hours) model forecasts are cycled in sequence with data assimilation and data handling tasks. While computationally expensive (~100M CPU hours), these cycling experiments scale poorly (~10,000 CPUs) and have to be performed sequentially. This requirement is further exacerbated at NOAA where the need to calibrate sub-seasonal forecasts with the new coupled model (UFS) requires development of the new coupled reanalysis product that spans a 30-year period. Existing on-premises NOAA HPC resources are insufficient to accommodate the additional workload of generating a new reanalysis in a timely manner. To prove feasibility of using the commercial cloud to surge computational resources for reanalysis production, NOAA-PSL has ported a complete cycling system to the AWS and MS Azure cloud providers.

The initial prototype was by-design deployed at a low resolution (1 degree configuration of the coupled ocean, atmosphere, and ice UFS model) and in deterministic-only configuration (GSI 3DVAR for atmosphere and JEDI SOCA 3DVAR for ocean and ice). We used the Cylc workflow engine to orchestrate the cycling tasks and Slurm to manage the job scheduling and resource allocation. We tried to replicate on-premise configuration of the compiler and architecture suite by using Intel's 2018v4 compiler, NOAA-EMC's hpc-stack, and using Intel-based computing instances. On AWS, we configured the cluster using a single master node (m5.large or c5n.2xlarge), six compute nodes (c5n.18xlarge), and file storage based on a shared Elastic Block Storage (EBS) volume or FSx Lustre. On Azure, we configured the cluster using a single master node (Standard_H8 or Standard_HC44rs), four compute nodes (Standard_HC44rs), and file storage based on either BeeOND or Lustre. Observations and initial conditions were staged on an AWS S3 bucket and this bucket was used by both cloud providers. Our initial assumptions about the ability to use singularity containers to easily port the workflow between platforms proved overly optimistic as the container configuration and compile had to match the specifics of the MPI drivers installed on each provider. Our hopes to efficiently utilize spot pricing were also overly optimistic where we experienced poor availability of the preferred instances. Performance and cost were similar between the two providers and comparable to prototype runs on the on-premise HPC (when the cost of waiting in the queue was excluded) demonstrating feasibility of using cloud computing as a surge platform for NWP computations.

Primary author: LAWRENCE, Steve (CIRES/NOAA/PSL)

Co-authors: FROLOV, Sergey (Naval Research Laboratory); WINTERBOTTOM, Henry (CIRES/NOAA/PSL); WHITAKER, Jeff (NOAA/PSL)

Presenter: FROLOV, Sergey (Naval Research Laboratory)

Session Classification: Session 8

Track Classification: 19th Workshop on high performance computing in meteorology