Presentation 3: Learning from the behavior of UFS

Stefan Gary (Parallel Works)

- There are several levels of cloud cluster configurations that can have a significant impact on the performance of large, multi-node HPC jobs:
    - Cloud instance types,
    - Networking type,
    - Lustre configuration settings,
    - Slurm configuration settings, and
    - Compiler and MPI implementation.

- Working through all these parameters is laborious; we plan to automate the benchmarking process so that a small UFS configuration can be used at the core of an automated workflow to explore the parameter space of cluster configurations.

- There are trade-offs between highest performance and general use configurations.

- Furthermore, instance provisioning times, availability, and spot pricing can vary; we are working toward applying machine learning (ML) to guide users for when and where and how to launch their jobs to match time and cost constraints.

- We already use ML to predict compute costs for certain aspects of live display of on-demand pricing.

Example UFS benchmarking results in the cloud: HPC FV3GFS Benchmark Test Performance

CSP	Processor			Count	Cores	Mins / Model Day
AWS	AMD: CPU	hpc6a.48xlarge	96	70	96	5.60
AWS	Intel: vCPU	c5n.18xlarge	72	102	36	5.95
Azure	Intel: CPU	Standard_HC44rs	44	96	44	6.10
GCP	Intel: vCPU	c2-standard-60	60	144	30	6.86

Credit: Shaxted, M. and Kirandumkara, U. (2022) Multi-Cloud HPC Clusters with Hybrid Workflow, Supercomputing 2022, Dallas TX.