Storage benchmarking for HPC

Mike Mesnier, James Hendricks, Raja R. Sambasivan, Matthew Wachs,
Gregory Ganger, Garth Gibson
Contact email: mmesnier@ece.cmu.edu

The storage accesses of an HPC application are difficult to simulate with existing benchmarks, particularly when the application contains coordination among its compute nodes (e.g., synchronization and collective operations). This work proposes a new environment for storage benchmarking that explicitly models such coordination using a workflow specification language. Specifications are then input into a distributed workload simulator in order to approximate the I/O of an HPC application on a given storage system.

This work will extend an existing workflow specification language with a description of the storage accesses of each compute node. The model may include constructs for data sources (e.g., disk drives and file systems), transformations on the data at compute nodes, and data sinks. Then, the I/O phases can be described in terms data "flows" between sources and sinks. However, to be accurate, the model must also consider the performance affecting attributes of a workload (e.g., read/write ratio, request size, and randomness) and the simulators must be modified to replay these flows accordingly.

Selecting a workflow modeling environment and providing the appropriate extensions for storage is one contribution of this work. In addition, this work will foster industry involvement. Specifically, a repository of workload specifications for representative HPC applications (e.g., life sciences, earth sciences, etc.) can be maintained and used by others to test their storage system. In addition, domain experts (e.g., oil/gas) will be encouraged to contribute specifications of their workloads.

In summary, we believe the HPC community will benefit from more representative storage benchmarks based on workflow specification. We are building such a benchmark and will provide reference specifications for applications in various fields of computational science.