High Performance Storage for HPC: Proven HDF5 Scalability and Consistency on Pure Storage FlashBlade

Monday 26 May 2025 11:00 (25 minutes)

The integration of the HDF5 high-level IO library, encompassing Single Writer Multiple Reader (SWMR) and parallel mode (pHDF5), with Pure Storage FlashBlade using NFS, modern Linux Kernels, and Networks, offers a robust and high-performance solution. This combination significantly reduces Total Cost of Ownership (TCO) and technical debt compared to traditional parallel file systems.

HDF5, particularly in SWMR configuration, ensures data consistency during concurrent read and write operations. Pure Storage FlashBlade's NFS provides a high-performance storage solution, ideal for High-Performance Computing (HPC) environments, facilitating efficient access and modification of large datasets. Enhancements within HDF5, such as superblocks, further bolster data consistency. FlashBlade's POSIX-compliant data access guarantees atomicity during read and write operations to direct flash memory modules, a crucial requirement for SWMR operations to prevent inconsistencies.

Pure Storage FlashBlade delivers exceptional throughput, low latency, and supports GPU Direct Storage with NFS over RDMA, essential for HPC. Multi-iteration write and read tests have demonstrated robust throughput and consistency, validating FlashBlade's capability to meet HDF5 demands in both SWMR and parallel configurations. Parallel HDF5 leverages MPI-IO, which supports NFS backends, treating it as a unified file system. This integration ensures parallel processes operate efficiently on shared datasets, with MPI layer compatibility enabling effective data handling across multiple nodes.

With optimized HDF5 settings, SWMR mode operates without exclusive or shared locks, allowing a single writer to handle write operations while multiple readers access data concurrently without conflicts. Comprehensive tests have validated FlashBlade's ability to run HDF5 in both SWMR and parallel modes, ensuring atomic writes and reads. Performance benchmarking using HDF5-specific tools has confirmed high throughput and data consistency across multiple nodes via NFS.

In conclusion, the synergy of HDF5's robust data handling with the high performance and consistency of Pure Storage FlashBlade using NFS creates a powerful solution for managing large, complex datasets in HPC environments.

May we record your session?

Yes

Primary author: KASPARS, Michael (Pure Storage)Co-author: CHOUDHURY, Bikash Roy (Pure Storage)Presenter: KASPARS, Michael (Pure Storage)