

# Streamlining HDF5's AI Workloads Benchmarking

*Wednesday 28 May 2025 10:35 (25 minutes)*

Rapid adoption of artificial intelligence (AI) in scientific computing requires new tools to evaluate I/O performance effectively. HDF5 is one of the data formats frequently used not only in HPC but also in modern AI applications. However, existing benchmarks are insufficient to address the current challenges posed by AI workloads. This talk introduces an extension to the existing HDF5 benchmark, called h5bench, by incorporating the workload characteristics from the MLPerf Storage - DLIO benchmark. This extension allows users to test AI workloads together with traditional HPC benchmarks in the same context without the complexities of installing various machine learning libraries. Our experimental analysis demonstrates that the extension can replicate the existing I/O patterns with easily customizable configurations to perform various scaling tests.

**Presenter:** DJEBAROV, Dlyaver