Contribution ID: **2** Type: **Submitted talk**

Expanding HDF5 capabilities to support multi-threading access and new types of storage

Tuesday 19 September 2023 10:00 (30 minutes)

Enabling multi-threaded access to data stored in HDF5 and efficient storage of sparse and variable-length data are long-standing requests from the HDF5 user community. Lifeboat, LLC has been working closely with The HDF Group on design and implementation of the new capabilities.

In our talk we will report on the progress we made toward multi-threaded concurrency in HDF5 since the last European HUG at ITER in May 2022. We will also present proposed extensions to the HDF5 File format and public APIs to support sparse and variable-length data storage in HDF5. The proposed sparse storage is agnostic to memory structures used to represent sparse data in RAM (e.g., sparse matrix), and provides storage savings and portability between different memory formats. New implementation of variable-length data in HDF5 will allow significant improvements in I/O performance and will finally enable compression of the variable-length data in HDF5.

Website

Primary authors: Ms POURMAL, Elena (Lifeboat, LLC); Mr MAINZER, John (Lifeboat, LLC)

Presenter: Ms POURMAL, Elena (Lifeboat, LLC)

Session Classification: Day 1