

How would you create separate files for each compute node in a cluster using HDF5?

1. If you use the Parallel HDF5 library, open a communicator for each node and add just one process to that communicator.
2. You may also try to use the sequential library, but then you have to make sure that only a particular process creates/modifies/closes the file that is written to. This approach has not been tested.

Pros: should be faster than 1)

Cons: multiple files to handle, not tested

Also there can be other approaches if you decide to go with multiple files. For example, each process writes a flat binary file and one can use the HDF5 external datasets storage feature to create a wrapper HDF5 file to combine all data. See the section on Dataset Creation Property Lists under [Property Lists](#) in the HDF5 Tutorial.