

What is the best way to write metadata for each process?

For an HDF5 file that is being written to in parallel, I am trying to simulate the writing of attributes/metadata on a per-process basis. These attributes contain metadata about an individual process's data (they are not shared across all processes). I know that HDF5 attributes must be written collectively with Parallel HDF5. Is there another better way to write them?

Most applications do either collective calls, or have one process write the attributes (then close the file). Usually the best performance, when possible, is to have one process create all the data layout in the HDF5 file and then close the file. Then have all the processes open the file for writing.