

NetCDF-4 and Parallel-netCDF

A brief history: In the fall of 2008, UCAR released NetCDF4, which took the HDF5 file format and features, and put the serial netcdf interface on top of that. UCAR modified the serial netcdf API to expose some of the new features HDF5 introduced, including a means for parallel I/O.

Some software engineering gave NetCDF-4 the concept of a separation between interface and file format (netcdf API, legacy netcdf or new HDF5 file format). Soon, UCAR extended this interface and gave NetCDF-4 the ability to use parallel-netCDF. Now, with some caveats, the classic netcdf API can read and write the classic netcdf file format either using the classic netcdf library or our parallel-netcdf library.

We think this is terrific, and in no way has UCAR usurped or challenged parallel-netcdf with this move. Instead, UCAR has provided a valuable option to applications developers. In order to transition to parallel I/O, the typical "file per process" or "all I/O from rank 0" approaches need to be re-thought, no matter which underlying library will carry out the I/O. Applications accustomed to serial netcdf can try out pnetcdf while still using an API familiar to them. If they want to use some of the more sophisticated parallel-netcdf features, or try out pnetcdf down the road, only minor edits to the names of the functions will be needed -- the application has already done the hard work of domain decomposition.