

# PIO Library for HPC Performance with NetCDF-based Codes

Ed Hartnett  
CIRES/NOAA  
9/7/17



Cooperative Institute for Research in Environmental Sciences  
UNIVERSITY OF COLORADO BOULDER and NOAA



# Problems with IO on HPC

- Distributing data across many processors results in IO delay.
- Expensive processors wait for IO to complete.
- Conceptually simple fix: subset of tasks to do IO.
- Expensive to code.

# Introducing The PIO Library

- Free C/Fortran software library from NCAR.
- Now includes async writes from multiple computational components.
- Sits on top of pnetcdf/netCDF.
- API mirrors netCDF API for easy use.
- Includes distributed array support.

# Future Plans

- Upcoming major release with full testing and documentation of new features.
- Performance testing.
- Packaging within netCDF as optional additionally supported create/open mode.