

# MPI-IO hints for PnetCDF

PnetCDF sits on top of MPI-IO, allowing us to use MPI Info objects to tune behavior. This page gives some more information about the way pnetcdf uses MPI Info objects, and also some situations where hints have proven useful.

## MPI Info and pnetcdf API

The pnetcdf routines `ncmpi_create`, `ncmpi_open`, `nfmpi_create`, and `nfmpi_open` take a hint parameter. That parameter is usually `MPI_INFO_NULL`, which in most cases is just fine.

If you create an MPI-IO Info, and set it up with some "hints", then you can direct the MPI implementation to perhaps take a different action. Implementations can ignore MPI-IO hints, so if you are running on an implementation that does not understand a particular hint, that's ok.

## Using MPI Info objects

At a high-level, here's how PnetCDF uses MPI Info objects Here's how you might be able to use hints in your code.

1. create an info object
2. set hints
3. pass the hints to create/open

## Useful Hints

On the Lustre file system, especially on Cray systems, we've seen quite high file locking costs. For some workloads, the optimizations for noncontiguous access might actually prove to be slower.

- set "romio\_ds\_read" "disable"
- set "romio\_ds\_write" "disable"

or

- try setting "ind\_wr\_buffer\_size" to "16777216"

## Examples

```
INTEGER INFO, IERROR
CHARACTER*(*) KEY, VALUE

MPI_INFO_CREATE(INFO, IERROR)
MPI_INFO_SET(INFO, KEY, VALUE, IERROR)
NFMPI_CREATE(MPI_COMM_WORLD, "myfile.nc", flags, INFO, NCID)
MPI_INFO_FREE(INFO, IERROR)
```

## More Information

[The MPI-2 standard's chapter on MPI Info objects](#)