

# Parallel-NetCDF: A High Performance API for NetCDF File Access

## Overview

Parallel-NetCDF is a library providing high-performance I/O while still maintaining file-format compatibility with Unidata's NetCDF.

NetCDF gives scientific programmers a space-efficient and portable means for storing data. However, it does so in a serial manner, making it difficult to achieve high I/O performance. By making some small changes to the API specified by NetCDF, we can use MPI-IO and its collective operations.

## Download

Our latest release is 1.0.1. This release has much better support for cross-compilation, builds more easily on BGL, and introduces the nonblocking I/O API.

- bzip2ed tarball: (3.3 MB) <http://www.mcs.anl.gov/parallel-netcdf/parallel-netcdf-1.0.1.tar.bz2>
- gzipped tarball: (13 MB) <http://www.mcs.anl.gov/parallel-netcdf/parallel-netcdf-1.0.1.tar.gz>

**Test Releases:** We have a lot of new work we want to see more widely tested. Please report any issues you find to the parallel-netcdf mailing list.

- bzip2ded tarball: (3.3 MB) [parallel-netcdf-1.0.2pre1.tar.bz2](#)
- gzipped tarball: (13 MB) [parallel-netcdf-1.0.2pre1.tar.gz](#)

## Documentation

- Our Parallel NetCDF API (postscript, 158k) document describes the API we are using. We have tweaked the programming interface to be more friendly to parallel i/o while maintaining file format compatibility with the serial version of NetCDF.
- Our SC2003 Paper about Parallel-NetCDF (PDF, 97k) discusses our library and presents some performance results.
- Here is Jianwei Li's presentation (PDF, 167k) from the SC2003 conference.

## A note about Large File Support

As of parallel-netcdf-0.9.2, we ship with support for "CDF-2" formatted data. With this format, even 32 bit platforms can create netcdf datasets greater than 2GB in size. See the file README.large\_files in the source tree for more information.

The maintainers of the serial NetCDF library added support for the CDF-2 format in netcdf-3.6.0. The support was based largely on this patch from Greg Sjaardema.

## Supporting Software

Parallel-NetCDF makes use of several other technologies.

- ROMIO, an implementation of MPI-IO, provides optimized collective and noncontiguous operations. It also provides an abstract interface for a large number of parallel file systems.
- One of those file systems ROMIO supports is PVFS, a high performance parallel filesystem for linux clusters.

## Mailing List

We discuss the design and use of the Parallel-NetCDF library on the `parallel-netcdf@?` mailing list. Anyone interested in developing or using `parallel-netcdf` is encouraged to join. Send mail to `majordomo@?` with the body `subscribe parallel-netcdf`.

You can browse old mailing list messages at the `parallel-netcdf` mailing list archives

## In the news

- Forrest Hoffman wrote an article about `parallel netcdf` in the July 2004 issue of Linux Magazine.
- The HDF group at NCSA ported a serial NetCDF code to one using Parallel-NetCDF. They posted a writeup a writeup of their efforts . It's a little old but does provide some additional information to supplement `doc/porting_notes.txt`

## Project Members

- Bill Gropp, Rob Latham, Rob Ross, Rajeev Thakur (Argonne National Lab)
- Alok Choudhary, Jianwei Li, Wei-keng Liao (Northwestern University)