

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 1.0.2. This is primarily a bugfix release, collecting all the fixes and improvements since 1.0.1. See the release announcement for more information

- bzip2ed tarball: (3.3 MB) [parallel-netcdf-1.0.2.tar.bz2](#)
- gzipped tarball: (13 MB) [parallel-netcdf-1.0.2.tar.gz](#)

Test Releases: We currently have no test releases, but you can always find the latest code in our subversion repository.

Subversion Access

The Parallel-NetCDF project is now using [Subversion](#) for source-code management. With the change we can also provide read-only access to anyone interested.

```
svn co https://svn.mcs.anl.gov/repos/parallel-netcdf/trunk parallel-netcdf
```

The SSL fingerprint should be df:f5:37:b1:69:11:e0:63:d3:99:a8:e4:de:50:11:01:f5:73:dc:0a

After you've checked out the source, run 'aclocal && autoconf && autoheader' to generate the configure script.

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.
- [Jianwei Li's presentation](#) (PDF, 167k) from the SC2003 conference.

Tuning

- [HintsForPnetcdf](#) describes some low-level turning parameters

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 work from Greg Sjaardema.

File and Variable Limits

Both Parallel-netCDF and NetCDF share limitations on file and variable sizes. More information can be found on the [FileLimits](#) page.

Required Software

Parallel-NetCDF requires an MPI implementation with MPI-IO support. Most MPI libraries have this nowadays. A parallel file system would also go a long way towards achieving highest performance.

Related Projects

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@mcs.anl.gov` mailing list. Anyone interested in developing or using parallel-netcdf is encouraged to join. Send mail to `majordomo@mcs.anl.gov` 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 of their efforts . It's a little old but does provide some additional information to supplement `doc/porting_notes.txt`

Project Members

- Rob Latham, Rob Ross, Rajeev Thakur (Argonne National Lab)
- Alok Choudhary, Wei-keng Liao (Northwestern University)
- Jianwei Li (NWU, since graduated)
- Bill Gropp (formerly ANL, now UIUC)