

NetCDF-4 Update

Ed Hartnett, Unidata/UCAR
NetCDF Workshop, July 25 – 26, 2011



Background of NetCDF-4

- Originally funded by NASA for 2 years.
- Merger (of sorts) of HDF5 and netCDF.
- NetCDF-4.0 released in 2008, with HDF5 1.8.0.
- This presentation gives update since last workshop (October, 2010) for netCDF-4 issues.

4.1.2 Release

- Released in March, 2011
- Added dispatch architecture.
- Upgrade to HDF5-1.8.6-patch1.
- Switch to subversion.
- Build system changes.
- Performance changes.
- Parallel I/O changes.

4.1.2 Release: Build System

- Refactor of build system to simplify and bring more into line with standard practice for free software products using automake/autoconf/libtool.
- No more specifying locations with configure options. Instead using environment variables.
- NetCDF now builds shared libraries by default.
- Result of refactor is more robust and simple build system, which does not meddle with matters best left to the user.

4.1.2 Release: Build Example

Old way:

```
./configure --with-hdf5=/some/place
```

New way:

```
CPPFLAGS=-I/some/place/include
```

```
LD_FLAGS=-L/some/place/lib ./configure
```

4.1.2 Release: Performance

- Huge speedup by changing the way netCDF-4 was scanning the file on open. For large files it is now much faster to open.
- Fixed very slow algorithm for checking for duplicate attribute names. This matters when thousands of attributes (or more) are being used.
- Improved speed of adding variables, dimensions, and other objects.

4.1.2 Release: Parallel I/O

- Problem with previous releases: including sequential build of netCDF within parallel I/O program.
- To fix, a new header file was included for parallel I/O users.
- Parallel I/O users note: you now **MUST** include `netcdf_par.h` instead of (or at least before) `netcdf.h`.

4.1.3 Release

- First in a shorter release cycle (quarterly).
- No more `--with-hdf5` configure options! Now we do things the standard way (CPPFLAGS/LDFLAGS).
- Upgrade to HDF5 1.8.7.
- Build of 32-bit Windows DLLs and tool executables on mingw32/Linux system.
- Bug fixes.
- Switch to Jira site for netCDF project planning:
www.unidata.ucar.edu/jira

Since 4.1.3

- Better default chunking algorithm.
- Switching documentation to doxygen.
- Bug fixes, trying to make windows builds work better.

Poster at the EGU

- Last April, Ed went to the European Geosciences Union in Vienna.
- Poster “Recent Developments in NetCDF Libraries and Formats” received a lot of attention.
- Got a chance to talk to many European users about how they are using netCDF.
- Better than GRIB2 compression was a commonly-expressed need.

4.2 Release (Q3, 2011)

- In Q3 we will release 4.2, with as many features as we can get in.
- We are working on better compression (wavelet) for netCDF-4.
- Diskless files.
- 64-bit DLL builds with mingw.

Now We Are Cool!

- To be extra cool, future releases are to be named, and only assigned a version number when released.
- Although this allows us to create extremely pretentious sounding version names, this is not the only benefit.
- The Jira site has our next few releases listed for the C library.

Challenges and Ideas

- Java compatibility – can we keep up with HDF5? Or should we go with some protocol tunneling so that the Java code calls the C library?



More Challenges and Ideas

- Can we collaborate more with the netCDF community? I welcome suggestions!
- We have relatively few resources for a product of this importance. We must use our time wisely.

Acknowledgments

- NetCDF-4 would not have achieved the success that it has without the hard work and dedication of Russ, John, and others at Unidata.
- HDF5 provides the great features that netCDF-4 exposes. We too rarely take the chance to thank the HDF5 team for their efforts.
- NASA funded netCDF-4 for the first two years; their continued use and support for netCDF-4 is much appreciated.
- Most of the netCDF-4 engineering effort was supported by the NSF funded Unidata Program Center.

Suggestions?

- Suggestions for new netCDF-4 features are always welcome.
- Sending suggestions to support-netcdf@unidata.ucar.edu ensures that the entire team sees them, and that they don't get lost.
- There is also value in posting suggestions to the netCDF mailing list to get other opinions.

Questions?



More Future Plans for NetCDF-4

- Better exposure of dispatch interface, so users can write their own netCDF read/write dispatch layers.
- Refactor pnetcdf as full libdispatch layers.
- Internal documentation with doxygen.
- Upgrade to HDF5 1.10.
- Better interoperability (with NPP, for example).
- Better Fortran support when Fortran APIs are rewritten for Fortran 2003.

