

A Roadmap for Implementing NetCDF4 Configure/build

March 30, 2005

Overview

The general question is how we want the netCDF-4 source code to be built by users. The particular technical challenge is how to link with HDF5 in the right way. The solution should meet the general requirements of ease of use and robustness.

This problem is related to the question of how the final product will be distributed. The ultimate goal is for users to be able to obtain netcdf-4 and HDF5 in a single download. This will not be the only option.

This note discusses what we plan to do.

Roadmap

The configuration and build can be divided into two phases, the initial development through initial release, and the final product. The former will be completed summer and fall of 2005. The latter will be done later.

For the initial phase, the configuration will be:

- HDF5 and netcdf-4 are two separate source trees.
- netcdf-4 should be built using 'h5cc' or 'h5pcc'

For the final phase, netcdf-4 will be configured as a part of the HDF5 source, i.e., under the directory 'hdf5' in the source. Details of this will be determined later.

Initial Phase

There are two approaches that have been proposed:

- 1) set up netcdf4 configure/make to use the 'h5cc' script
- 2) set up netcdf4 configure/make to use regular compilers, with '--with-hdf5' to point to HDF5

We currently have a version of #1 implemented. It is working, but there are at least 2 irritating problems that have yet to be addressed. We do not know how difficult it will be to fix the problems. However, most users will not have trouble with them.

#2 is thought to be fairly simple, but we don't have a detailed technical assessment.

Known Problems with h5cc

1. 'h5cc' has all the warnings enabled

This bug has been reported several times, and really should be fixed. The fix requires modification to how the h5cc script is generated--probably some sed or other shell hacking.

2. Make in sub-directories fails in some cases

Automake doesn't understand 'h5cc', and is not doing dependency checking correctly. In some cases, make fails because of this. This has been a pain for developers, but is not an issue for most users.

There are at least possible fixes:

- a) the autoconf mail list suggests fixing 'depcomp'

- b) disabling dependency checking

These have not been explored in detail, and we don't know how difficult it will be to fix this.

We should realize that there may be other problems we don't know about.

Not using 'h5cc'

The alternative to using 'h5cc' is to rewrite the netcdf-4 autoconf and possibly other files to accept the '--with-hdf5' argument, and do whatever is necessary to find the right compilers, libraries, and settings.

Ideally, this would be able to use information from, say, libhdf5.settings, to create the right make and link steps.

It is not clear how much work this rewrite would be. Without trying to do it, we can't really know if this is easier or harder to use than 'h5cc'.

Conclusion for Short Term

We will distribute netcdf-4 designed to be built with 'h5cc'. Essentially, keep it the way it is implemented now.

We will address the problems with h5cc at some time in the future, but not right now.

Final Phase

After the netcdf-4 product becomes stable, we will work to integrate the two libraries, in order to provide a single download. The technical and non-technical details will be worked out in the intervening time.

Proposed Roadmap

About six months after initial release, some combined installations will be available. Users will be able to download one tarball from the Unidata site, and install both netCDF and HDF-5 libraries. They will still have to link to both libraries.

Within a year of initial release, a cvs merge of code bases will allow users to create a single library, with a single download and install procedure. NetCDF-4 will be a optionally enabled part of the regular HDF5 distributions. Users will only need to link to one library, the HDF5 one.

Additionally, a netCDF-only release will be available from Unidata. It will allow us to distribute betas and other updates, and allow users to install independently from HDF5 installation. Users may build HDF5 without netCDF, then build netCDF, then link to both libraries. Only enthusiastic beta testers and users in dire need of the latest release will attempt this.

As netCDF beta code becomes release code, it will be merged into HDF5 releases.