

2018-10-03

Data Curation Format Profile: netCDF

Sciolla, Sam; Borda, Susan; Hou, Sophie

<https://hdl.handle.net/2027.42/145724>

<http://creativecommons.org/licenses/by-nc/4.0/>

Downloaded from Deep Blue, University of Michigan's institutional repository

Data Curation Format Profile: netCDF
 Research Data Services
 University of Michigan Library

Creator(s): Sam Sciolla (ssciolla@umich.edu), Susan Borda* (sborda@umich.edu)
 *denotes corresponding creator

Core Details about netCDF	
File Extension	.nc
MIME Type	application/netcdf, application/x-netcdf
Structure	Binary, with metadata embedded in a header that can be rendered human-readable by specialized software tools
Versions	netCDF-4/HDF5 (post-4.0.0) 64-bit offset (pre-4.0.0) Classic (pre-3.6.0)
Primary fields or areas of use	Climatology, Meteorology, Oceanography. Earth and Environmental Sciences / Geosciences more broadly. Also used in GIS applications.
Source and affiliation	NetCDF , the software and file format, is developed and managed by Unidata and the University Corporation for Atmospheric Research (UCAR), with funding from the National Science Foundation. The online documentation for the format is open and thorough. The source code for netCDF libraries is available from GitHub (linked in the documentation).
Metadata standards	Climate and Forecast Metadata Conventions
Key questions for curation review	<ul style="list-style-type: none"> ● How are netCDF files organized in the dataset? ● How thorough is the embedded metadata in the headers? ● What steps were involved in producing these netCDF files?
Tools for curation review	<ul style="list-style-type: none"> ● Panoply: Java application available for PC, Mac, and Linux ● NCView: Unix-based application run using a command-line application ● NCAR Command Language: Language that includes the ncdump command for viewing and extraction of file contents ● MATLAB (or GNU Octave): scientific computing software and language with built-in netCDF libraries
Date Created	July 2, 2018
Date updated and summary of changes made	October 3, 2018: Anonymized the researchers from the interview and added the version of MATLAB used. (SB)



Table of Contents

Description of the format	2
Examples of public netCDF data collections	4
Sample data file used in the examples in this profile	4
Key questions to answer when reviewing a netCDF dataset	4
Instructions for resources to use in the curation review of netCDF files	7
Appendix A: Software used by researchers in working with netCDF files	14
Appendix B: File Format Questions for Researchers	15
Bibliography	15

Description of the format

NetCDF is both software and a file format used by researchers in the geosciences to store and analyze data in multi-dimensional arrays. The technology was developed and is maintained by Unidata, part of the University Corporation for Atmospheric Research (UCAR). According to the netCDF documentation, UCAR and Unidata have held the copyright for netCDF since 1993 (Unidata 2018c).

NetCDF is designed to be a self-describing format, meaning that the file itself contains both the data and descriptive information about the data.

A netCDF file is encoded as a binary and contains two parts:

- a header that lists dimensions, variables, and attributes (see example below)
- the data streams

```

Source:
  grn_aer02.clm2.h0.0001-06.nc
Format:
  64bit
Global Attributes:
  title = "CLM History file information"
  comment = "NOTE: None of the variables are weighted by land fraction!"
  Conventions = "CF-1.0"
  history = "Tue Apr 24 15:31:05 2018: ncks -v time,QMELT,H2OSNO,landfrac
data/jamiewa/deep_blue/AOD_009/grn_aer02.clm2.h0.0001-06.nc
created on 06/03/16 04:46:58"
  source = "Community Land Model CLM4.0"
  hostname = "nyx_pgi"
  username = "jamiewa"
  version = "cesml_0_3"
  revision_id = "$id: histFileMod.F90 28831 2011-06-07 06:14:20Z erik $"
  case_title = "UNSET"
  case_id = "grn_aer02"
  Surface_dataset = "surfdata 1.9x2.5 simyr2000_c091005.nc"
  Initial_conditions_dataset = "clm1.BCN.2000-01-01 1.9x2.5_gx1v6_simyr2000_c100309.nc"
  PFT_physiological_constants_dataset = "pft-physiology.c110425.nc"
  RTM_input_dataset = "rdirc 0.5x0.5 simyr2000_c101124.nc"
  Time_constant_3Dvars_filename = "./grn_aer02.clm2.h0.0000-01.nc"
  Time_constant_3Dvars = "ZSO1:DZSO1:WATSAT:SUCSAT:BSW:HKSAT"
  NCO = "4.4.4"
Dimensions:
  time = 1
  lat = 96
  lon = 144
  hist_interval = 2
Variables:
  H2OSNO
  Size: 144x96x1
  Dimensions: lon,lat,time
  Datatype: single
  Attributes:
    long_name = "snow depth (liquid water)"
    units = "mm"
    cell_methods = "time: mean"
    Fillvalue = 999999961690316245365415600208216064
    missing_value = 999999961690316245365415600208216064

```

Figure 1: Sample netCDF (doi:[10.7302/Z24Q7S64](https://doi.org/10.7302/Z24Q7S64)) file header in Octave

Dimensions are often spatial and temporal in nature, and variables describe values specific to a study or a model used by researchers. Together data corresponding to dimensions and variables can be used to investigate and visualize how phenomena change across space and time. For instance, values for surface temperature might be mapped according to values for latitude and longitude, a process which could be repeated again for each step in a time series.

By convention, each dimension usually has a corresponding one-dimensional variable — called a coordinate variable — which shares a name with the dimension. Coordinate variables provide an efficient way for applications (such as NCView, discussed later in this document) to refer to a specific point along dimensions.

Attributes — which can be associated with variables and the entire file, in the case of “global attributes” — describe in more detail the information embedded in headers, specifying things like human-readable names (the “long_name” attribute), units, and data sources. Other attributes can illuminate the research and data manipulation process: “cell_methods” attributes can signify how a variable’s values were computed, and the “history” global attribute provides a place for a log of changes to a netCDF file. “[Appendix A: Attribute Conventions](#)” of The NetCDF User’s Guide lists common attributes and their uses (Unidata 2018b).

For more details on the netCDF structure, refer to the “The Components of a NetCDF Data Set” section of [The NetCDF User’s Guide](#), included in the netCDF documentation (Unidata 2018h).

Examples of public netCDF data collections

- Lawrence Livermore National Laboratory: “[Earth System Grid Federation’s Data Portal](#)”
- NOAA: “[World Ocean Atlas 2009 Data in NetCDF format](#)”

Sample data file used in the examples in this profile

Ward, Jamie L., Flanner, Mark G., Bergin, Mike, Dibb, Jack E., Polashenski, Chris M., Soja, Amber J., Thomas, Jennie L. (2018). Data in support of the study "Modeled Response of Greenland Snowmelt to the Presence of Biomass Burning-Based Absorbing Aerosols in the Atmosphere and Snow". University of Michigan Deep Blue Data Repository. doi:[10.7302/Z24Q7S64](https://doi.org/10.7302/Z24Q7S64)

Key questions to answer when reviewing a netCDF dataset

Datasets containing netCDF files can differ greatly in appearance, as researchers employ varying structures, amounts of data, levels of description, and models and tools in their creation. Thus, rather than creating a list of characteristics that “well-formed” netCDF files need to have, we have created a list of questions that we hope will help data curators determine the clarity and usefulness of a dataset to a discipline and direct conversation with researchers on ways it, and the netCDF files within it, might be improved.

For each overarching question, there will be some background information, gleaned from our own research and conversations with researchers; a statement of what your goal will likely be as a data curator and reviewer of the dataset; and a series of sub-questions to direct your investigation.

It will be easier to answer these questions once you have looked carefully through the header of a few netCDF files from the dataset and explored their data using a netCDF viewer. For an overview of the tools and methods we recommend for accomplishing these tasks, read the next section, [Instructions for tools to use in reviewing netCDF files](#). In addition, [Appendix A](#) provides a more detailed discussion of what tools the researchers we spoke with are using to interact with netCDF files.

1. How are netCDF files organized in the dataset?	Goal	To assess the structure of the dataset as a whole and how well metadata, documentation, and file-naming conventions explain that structure
	Questions to answer	<ul style="list-style-type: none"> • How many total netCDF files are included in the dataset, and are there any zip files or directories used to divide them? • Are there any clues as to the reasons or rationale for the file or directory divisions?

		<ul style="list-style-type: none"> To what extent is the organizational structure explained by the metadata, documentation, or file-naming conventions?
	Background	<p>NetCDF datasets often contain several netCDF files, but how they are organized will differ depending on the study and data management practices of researchers. Individuals we spoke with said that multiple files can represent differences in variables; different steps in a time series (though changes over time can also be contained within a single file); or different runs, which can then be averaged together (sometimes called multi-member ensembles) (Michigan Researchers A & B).</p> <p>Ideally, documentation accompanying the dataset would explain the purpose of each file and its relationship with the rest of the dataset. Clear file and directory names can also improve the user's ability to navigate through and comprehend the dataset as a whole.</p>

2. How thorough is the embedded metadata in the headers?	Goal	To assess how detailed and readable (both by humans and computers) the file headers are and whether they are likely to meet the needs of anticipated users of the dataset
	Questions to answer	<ul style="list-style-type: none"> To what extent are long_name, standard_name, units, and other attributes clear and unambiguous? Does there appear to be any empty fields or missing information? Are there fields in which additional detail might help users better understand the data? What conventions (if any) are declared in the global attributes, and does the header conform to those conventions? <div style="border: 1px solid black; padding: 2px; margin: 5px 0;"> Global Attributes: Conventions = "CF-1.0" </div> <ul style="list-style-type: none"> Do you see evidence that a particular disciplinary standard (such as metadata, controlled vocabulary, etc.) is being used? If not, would the data benefit from the use of a particular disciplinary standard?
	Background	NetCDF headers enable netCDF files to be self-describing, as they can contain information about the dimensions, variables, and the file itself. Whether this information is

		<p>manually entered or automatically populated by a simulation (or a combination of the two) can be unclear. Embedded metadata can improve the readability for humans and computers of a netCDF file by specifying human-readable names (long_name attributes), units, methods used to manipulate data, and details about their source and provenance. Much of this information is stored in attributes that apply to variables or the entire dataset. “Appendix A: Attribute Conventions” of The NetCDF User’s Guide lists common attributes and their uses (Unidata 2018b).</p> <p>Header information can also indicate compliance with metadata conventions for netCDF, the most ubiquitous set being the Climate and Forecast (CF) conventions. A major feature of the conventions is the management of “standard_name” attributes, which make possible comparison between datasets by defining unique identifiers for variables. Disciplinary repositories like the one affiliated with the Coupled Model Intercomparison Project (CMIP) require data providers to adhere to these conventions, among other stipulations. CF lists on its website a couple of online tools that check netCDF files’ compliance to current and past iterations of the conventions. While these conventions seem widely used, how important adhering to them is will likely depend on researchers’ needs and/or intentions to share or distribute the files.</p>
--	--	---

3. What steps were involved in producing these netCDF files?	Goal	To assess how much of the research process that produced and made use of these netCDF files is represented and explained by the dataset
	Questions to answer	<ul style="list-style-type: none"> ● Are configuration files and scripts for analysis or plotting included in the dataset? ● Do the global attributes in the headers of the netCDF files clearly indicate who the author is, what data sources and/or models were used, and how the files were changed during post-processing? ● Does the descriptive metadata or documentation describe (at least in brief) how the files were

		produced, modified, analyzed, and used to create visualizations?
	Background	<p>The production of netCDF files represents only one step in a provenance chain that often involves creating configuration files, using supercomputing resources, post-processing, and plotting for visualization and analysis. To facilitate transparency and reproducibility, these steps might be represented by files included in the dataset and spelled out by information contained in netCDF headers, descriptive metadata, and documentation for the dataset.</p> <p>Depending on the complexity of the research and how researchers anticipate their data being reused, it may be valuable for researchers to include in the dataset the parameters or configuration files used to set up a simulation or scripts created to plot data contained within netCDF files. If these files are not initially included, a conversation with the researchers will likely answer the question of whether these files would or would not add to the usefulness of the dataset for others in the field.</p> <p>Within the global attributes portion of a netCDF header, researchers have the ability to embed information about themselves as authors, data sources, and simulation models used. In addition, changes to netCDF files can be tracked using the “history” attribute; such post-processing changes are common, as multiple files are combined, variables are removed to reduce file sizes, or information is added to headers. Appendix A includes information about command-line utilities that make use of the history global attribute.</p>

Instructions for resources to use in the curation review of netCDF files

To conduct a review of a dataset with netCDF files, you will want to open a sample of files in a viewer, either Panoply or NCView, and browse the header metadata using Panoply, MATLAB (or GNU Octave), or NCAR Command Language. How to use the basic functionality of these five tools and a compliance checker is described below.

While almost all information is available through Panoply, there are differences in the order and format of the information, and MATLAB provides the version of netCDF while Panoply, NCVIEW, and NCL's ncdump output do not.

Use Panoply to quickly and easily view files. Use MATLAB (or GNU Octave) to view file headers which will indicate whether to use the Compliance Checker or not. NCVIEW and NCL were specifically mentioned by the researchers interviewed as tools they use for viewing and investigating netCDF files.

1) Panoply

Website of the tool: <https://www.giss.nasa.gov/tools/panoply/>.

Version used in this profile: 4.9.0

What does this tool do?

Panoply can be used to view the data stream, plot any geo-referenced and other arrays, and browse the metadata.

Who supports this tool?

The Goddard Institute for Space Studies, a unit of the National Aeronautics and Space Administration (NASA)

How can I get access to this tool?

As a Java application, Panoply can be easily downloaded and run on a PC or a Mac, provided that you have a Java Runtime Environment (JRE) installed. The following link provides instructions on running the program and links to resources on installing a JRE:

<https://www.giss.nasa.gov/tools/panoply/download/>

The **Research Data Services PC** also has Panoply installed and should be accessible on any enabled account. You should be able to find the application file by following this path in File Explorer: This PC → C: → Program Files (x86) → PanoplyWin-4.9.0 → PanoplyWin → Panoply.exe.

How do I use this tool?

Once you have opened Panoply, a prompt will automatically appear asking you to select a file. After you have chosen the desired file, you will see a window like the one shown in **Figure 2**.




Figure 2: Sample netCDF (doi:[10.7302/Z24Q7S64](https://doi.org/10.7302/Z24Q7S64)) file open in Panoply

When the top line in the left panel is selected (and highlighted in blue), the right panel will display information contained within the header of the netCDF document, divided into dimensions, variables, and global attributes. Underneath the top line on the left panel, you will see a list of the variables defined in the header, with each line also including any human-readable name provided (“Long Name”) and an indication of how the variable can be plotted in Panoply (“Type,” with typical values being “—”, “1D”, “2D”, or “Geo2D”). Clicking on the line for a variable will cause the right panel to display information in the header specific to that variable.

When a variable with one or more dimensions is highlighted, the “Create Plot” option will be enabled in the top left of the window. After clicking on the option, you will see a window asking you to select how you would like to plot or render the data. Choose the most appropriate option given your knowledge of the research and header information, and a plot will appear in a new window. **Figure 3** provides an example of the first option, a georeferenced Latitude-Longitude plot.

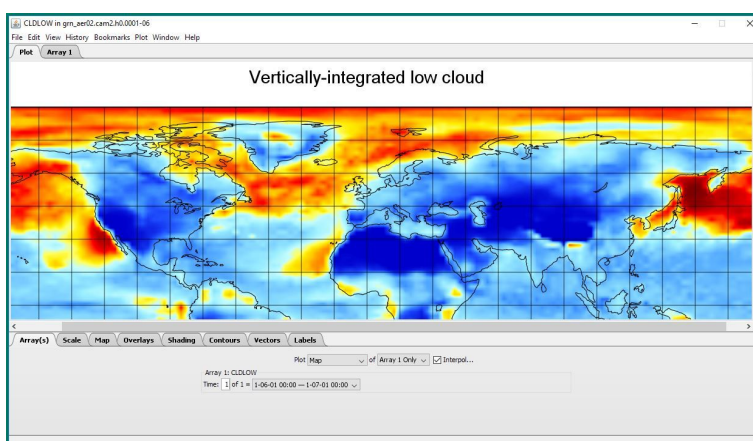


Figure 3: A georeferenced Latitude-Longitude plot from sample netCDF (doi:[10.7302/Z24Q7S64](https://doi.org/10.7302/Z24Q7S64)) in Panoply

At the bottom of the window, there are numerous options that allow you to tweak the rendering and then create exportable images of the plots. If there are multiple values associated with the time dimension in this netCDF file, you can view the plot of the variable at a different time by typing in a number or selecting an option from the drop-down menu on the “Array(s)” tab. You can also click on the “Array 1” tab at the top left-hand corner to view a matrix displaying the individual values at particular points along the dimensional axes.

2) MATLAB

Website of the tool: <https://www.mathworks.com/products/matlab.html>

Version used in this profile: R2017a

What does this tool do?

MATLAB has built-in libraries for handling data in netCDF files. While plotting netCDF data would require deeper knowledge of MATLAB syntax, extracting the netCDF header information is simple, requiring only a single command. MATLAB is installed and configured on all enabled accounts on the **Research Data Services PC**.

How do I get access to this tool?

MATLAB is a commercial product, but may be available to you as a part of an institutional subscription. Ask the unit that handles software purchases at your institution for information.

How do I use this tool?

To begin, open MATLAB and use the file toolbar, located just above the “Current Folder” and “Command Window” panes, to navigate to the directory containing the netCDF file from which you want to extract header information. Once you are in the correct directory, you should see the file in the Current Folder pane on the left.

Then, in the Command Window, type the following, replacing “file_name” with the name of the file surrounded by quotation marks:

```
>> ncdisp("file_name")
```

After hitting enter, the header information will appear in the Command Window. From there, you can easily copy and paste it into a text document for later inspection. The MATLAB interface should look something like **Figure 8** below after entering the ncdisp command.

```

>> ncdisp('grn_aer02.clm2.h0.0001-06.nc')
Source: C:\Users\asciolla\Desktop\netCDF_experimentation\grn_aer02.clm2.h0.0001-06.nc
Format: 48bit
Global Attributes:
Conventions = 'CF-1.0'
source = 'CAM'
case = 'grn_aer02'
title = 'UNSET'
loopname = 'jamaieva'
host = 'nyx4412.aero-ia.u'
version = '2.0.0'
revision_id = '1418'
initial_file = '/scratch/Climate_flux/jamaieva/cmdata/atm/cam/mio/Cv/cam_0000-01-01_1-9kz_0_12k_007900.nc'
topography_file = '/scratch/Climate_flux/jamaieva/cmdata/atm/cam/topo/2000-01-01-1-9kz_0_12k_007900.nc'
history = 'Tue Apr 24 15:31:44 2016: ncks --v time,NetFluxTot,SHEX,SHEFX,FMS,THEFIT,IS,T,O,CLOW,LANDFAC,FMSO grn_a
Fri Jan 6 16:01:20 2017: ncsp2 -O -s NetFluxTot-FMS-FMS-SHEX-SHEFX grn_aer02.clm2.h0.0001-06.nc /data/jamaieva/
Thu Dec 1 15:47:16 2016: ncsp2 -O -s NetFluxTot-FMS-FMS-SHEX-SHEFX grn_aer02.clm2.h0.0001-06.nc /data/jamaieva/
Mon Jun 6 17:16:10 2016: ncsp2 -O -s FMS-FMS grn_aer02.clm2.h0.0001-06.nc /data/jamaieva/good_box/grn_aer02/

ncp_openmp_thread_number = 1
BCD = '4.4.4'
Dimensions:
time = 1 (UNLIMITED)
lat = 96
lon = 144
lev = 26
sband = 2
Variables:
cldcov
Size: 144x96x1
Dimensions: lon,lat,time
Datatype: single
Attributes:
Sampling_Sequence = 'rad_levs'
units = 'fraction'
long_name = 'Vertically-integrated low cloud'
cell_methods = 'time: mean'
FMS
Size: 144x96x1
Dimensions: lon,lat,time
Datatype: single
Attributes:
Sampling_Sequence = 'rad_levs'
units = 'M/m^3'

```

Figure 8: Result of the ncdisp command in MATLAB

3) GNU Octave

Website of the tool: <https://www.gnu.org/software/octave/>

Version used in this profile: Octave-4.4.0

What does this tool do?

GNU Octave is an open source software similar to MATLAB.

How do I get access to this tool?

GNU Octave is freely available online for download.

To work with netCDF files you will need to download “netcdf package”

<https://octave.sourceforge.io/netcdf/index.html>

How do I use this tool?

To view netCDF headers you will need to first install and load the netcdf package. Enter the following commands at the Octave prompt:

```
>> pkg install -forge netcdf
```

```
>> pkg load netcdf
```

Once netcdf package is loaded to view the header of a netCDF file (Figure 1):

```
>> ncdisp("file_name")
```

For example:

```
>> ncdisp("grn_aer02.clm2.h0.0001-06.nc")
```

4) NCVIEW

Website of the tool: http://meteora.ucsd.edu/~pierce/ncview_home_page.html

Version used in this profile: 2.1.7 Released March 29, 2016

What does this tool do?

Ncview is a visual browser for netCDF format files that can be used for a quick review of their contents. Researchers pointed to NcView as a way to quickly explore trends in the data of a netCDF file.

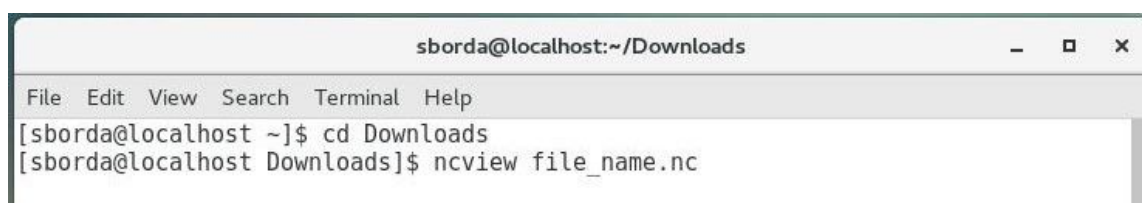
How can I get access to this tool?

NcView can be downloaded from: http://meteora.ucsd.edu/~pierce/ncview_home_page.html. As a Unix-based application, the viewer is simpler to set up on a Mac computer, but it can be run on a PC using tools such as Cygwin or a virtual machine.

The **Research Data Services PC** has a Linux virtual machine set up with NcView installed for convenient use. For assistance in configuring and logging in to the virtual machine on your account, speak with Data Workflows Specialist Susan Borda.

How do I use this tool?

NcView must be run from a command line. First, log on to the virtual machine, and, using the Firefox Web Browser application, download the netCDF file you wish to view, ensuring it is saved in the Downloads directory. Then click on Terminal (the virtual machine's command-line application), and use the Unix commands shown in **Figure 4**, where "file_name.nc" is the name of the file you want to view. Each command should be typed after the "\$" symbol and executed by hitting the Enter button.



```

sborda@localhost:~/Downloads
File Edit View Search Terminal Help
[sborda@localhost ~]$ cd Downloads
[sborda@localhost Downloads]$ ncview file_name.nc

```

Figure 4: Unix commands to start NcView in Terminal

Once NcView is open, you will see a window like the one on the left in **Figure 5**. The middle section labeled "Var: " lists all the variables defined within the netCDF file. Clicking one of the rectangles will cause the rest of the panel to populate, as shown in **Figure 5**'s right image. The panel will include the full name of the variable at the top, a scale for the variable's values in the middle, and the variable's dimensions (and their minimums, maximums, and units) at the bottom.




Figure 5: Before and after selecting a variable from sample netCDF (doi:[10.7302/Z24Q7S64](https://doi.org/10.7302/Z24Q7S64)) in NcView's main panel

In addition, a plot will appear in a separate window like the one shown in **Figure 6**. The initial area covered by the plot will vary based on information provided in the netCDF header. Additional modifications can be made using the options listed just above the color-coded scale in the right image of **Figure 5**. By hovering your cursor over the plot, you can also view the values at particular coordinates in the top section of NcView's main panel.




Figure 6: The plot window in NcView

If there are multiple values associated with the time dimension in this netCDF file, you can use the toolbar directly above the plotting options to either step through the different time values or display changes over time as a movie.

5) NCAR Command Language

Website of the tool: <https://www.ncl.ucar.edu/>

Version used in this profile: 6.4.0 Released on February 28, 2017

What does this tool do?

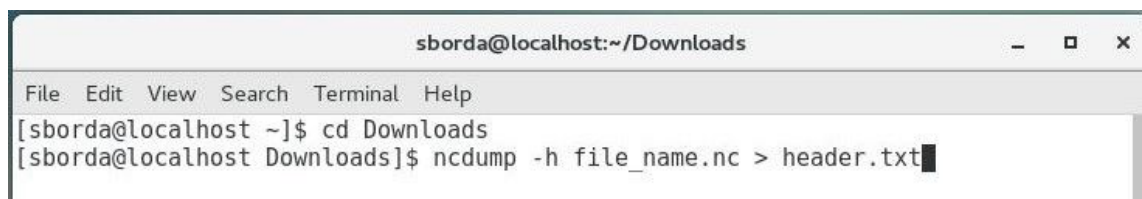
NCAR Command Language (NCL) is a programming language designed to work with netCDF files and other scientific data formats. While researchers commonly write NCL scripts to analyze and plot netCDF data, there are also simple tools available from a command-line application using NCL that allow for quick access to the file's contents, most notably “ncdump.” This command can be used to quickly extract header information for later inspection.

How do I get access to this tool?

NCL can be downloaded here: https://www.ncl.ucar.edu/current_release.shtml. Like NCView, NCL is simpler to run in a Unix-like environment. For convenience, NCL has been built on the **Research Data Services PC** using the virtual machine. Speak with Data Workflows Specialist Susan Borda for assistance configuring the virtual machine on your account.

How do I use this tool?

To use NCL's ncdump command, log on to the virtual machine and, using the Firefox Web Browser, download the netCDF file you would like to work with, ensuring that it is saved in the Downloads directory. Next, open the Terminal command-line application, and enter the commands in **Figure 7**, replacing “file_name.nc” with the name of the desired file.



```

sborda@localhost:~/Downloads
File Edit View Search Terminal Help
[sborda@localhost ~]$ cd Downloads
[sborda@localhost Downloads]$ ncdump -h file_name.nc > header.txt

```

Figure 7: Unix commands for ncdump command in Terminal

The second command will write the output of the ncdump command to a file called “header.txt”, which you can find in the Downloads directory. The output will have a neatly formatted, indented structure that is very similar to what is displayed in Panoply.

6) Compliance Checker

Website of the tool: <http://cfconventions.org/compliance-checker.html>

What does this tool do?

This utility checks that a netCDF file which you supply complies with the CF conformance requirements and recommendations.

How do I get access to this tool?

This tool is freely available online: <http://puma.nerc.ac.uk/cgi-bin/cf-checker.pl>

How do I use this tool?

Just upload your file where indicated and click “Check File.”

Appendix A: Software used by researchers in working with netCDF files

Interacting with netCDF files requires a configured computing environment, to satisfy a number of dependencies, or a specialized viewer. This is because the methods that exist for creating and interacting with these binary files use or rely on lower-level programming languages including C, C++, Fortran 77, and Fortran 90. While this makes it trickier to work with the contents of netCDF files, it allows for quicker and more powerful data manipulation and analysis by using “direct access rather than sequential access,” according to the netCDF documentation (Unidata 2018a).

Researchers at the University of Michigan who work with netCDF files reported primarily using MATLAB and the NCAR Command Language (NCL) to read and analyze data in netCDF files (Michigan Researchers A & B). They also noted that they have seen researchers and students using other languages, including Python and R (Michigan Researchers A & B). The netCDF documentation page also mentions a Java interface. Multiple researchers noted that they have been particularly happy using NCL to make plots for publications, in part because of functionality related to projections (Michigan Researchers A & B).

In addition, researchers discussed other utilities to modify or view netCDF files. Multiple individuals reported using netCDF Operators (NCO), an open-source suite of programs run from a command line that can perform common operations on netCDF files (Michigan Researchers A & B). When using NCO, changes to the file by these programs are tracked by appending entries to a global history attribute in the netCDF header, providing a form of version tracking. Researchers also referred to a similar suite called Climate Data Operators — developed by the Max Planck Institute for Meteorology — that some in the atmospheric sciences employ (Michigan Researchers A & B).

For a brief snapshot of a netCDF file, some researchers make use of NCView, a Unix-based application that presents metadata about the file’s structure (dimensions and variables) and allows for two-dimensional plotting (plus the creation of movies using a time dimension) and the reading of individual data points (Michigan Researcher A). Though the scientists Research Data Services spoke with did not report using them, other viewers do exist, including Panoply, which enables quick access to header information, simple data plotting, and browsing of data in arrays.

Unidata has compiled an extensive list of interfaces, languages, and tools for working with netCDF files, which can be found on its website:

<https://www.unidata.ucar.edu/software/netcdf/usage.html> (Unidata n.d.b).

Appendix B: File Format Interview Questions for Researchers

As a part of developing this document we interviewed several researchers who use the NetCDF format for their data. Our interview questions are available in a [Google Document](#).

Bibliography

Interviews

Michigan Researcher A. 2018. Interview. University of Michigan, Climate and Space Sciences and Engineering. Held on June 1, 2018.

Michigan Researcher B. 2018. Interview. University of Michigan, Climate and Space Sciences and Engineering. Held on June 6, 2018.

Web Resources

Library of Congress. 2017. “NetCDF-3 (Network Common Data Form, version 3).” Sustainability of Digital Formats: Planning for Library of Congress Collections. Last modified on July 27, 2017.
<https://www.loc.gov/preservation/digital/formats/fdd/fdd000330.shtml>

Max-Planck-Institut für Meteorologie. 2018. “Overview: Climate Data Operators.” Last modified on May 9, 2018. <https://code.mpimet.mpg.de/projects/cdo>

National Aeronautics and Space Administration. 2018. “Panoply netCDF, HDF and GRIB Data Viewer.” Goddard Institute for Space Studies. Last modified on June 1, 2018. <https://www.giss.nasa.gov/tools/panoply>

NCO Project. 2018. “Bienvenue sur le netCDF Operator (NCO) site.” Sourceforge. Last modified on June 10, 2018. <http://nco.sourceforge.net>

NetCDF Climate and Forecast Metadata Conventions. n.d. “Compliance Checker.” Accessed on June 25, 2018. <http://cfconventions.org/compliance-checker.html>

NetCDF Climate and Forecast Metadata Conventions. n.d. “NetCDF CF Metadata Conventions.” Accessed on June 25, 2018. <http://cfconventions.org>

Pierce, David W. 2016. “Ncview: a netCDF visual browser.” Last modified on March 29, 2016. http://meteora.ucsd.edu/~pierce/ncview_home_page.html

Program For Climate Model Diagnosis and Intercomparison. n.d. “CMIP5 - Modeling Info - Producing Model Output.” CMIP5 Coupled Model Intercomparison Project. Accessed on June 25, 2018. https://cmip.llnl.gov/cmip5/output_req.html

The National Archives. 2010. “Details for: netCDF-3 64 bit.” PRONOM. Last modified on Nov. 1, 2010. <http://www.nationalarchives.gov.uk/pronom/fmt/283>

Unidata. 2018a. “An Introduction to netCDF.” NetCDF. Last modified on March 15, 2018.

https://www.unidata.ucar.edu/software/netcdf/docs/netcdf_introduction.html

Unidata. 2018b. “Appendix A: Attribute Conventions.” NetCDF.

Last modified on March 15, 2018.

https://www.unidata.ucar.edu/software/netcdf/docs/attribute_conventions.html

Unidata. 2018c. “COPYRIGHT.” NetCDF. Last modified on March 15, 2018.

<https://www.unidata.ucar.edu/software/netcdf/docs/copyright.html>

Unidata. 2018d. “FAQ.” NetCDF. Last modified on March 15, 2018.

<https://www.unidata.ucar.edu/software/netcdf/docs/faq.html>

Unidata. 2018e. “File Structure and Performance.” NetCDF. Last modified on March 15, 2018.

https://www.unidata.ucar.edu/software/netcdf/docs/file_structure_and_performance.html

Unidata. 2018f. “Introduction and Overview.” NetCDF. Last modified on March 15, 2018.

<https://www.unidata.ucar.edu/software/netcdf/docs/index.html>

Unidata. 2018g. “netcdf-c.” GitHub. Last modified on June 22, 2018.

<https://github.com/Unidata/netcdf-c>

Unidata. 2018h. “The Components of a NetCDF Data Set.” NetCDF.

Last modified on March 15, 2018.

https://www.unidata.ucar.edu/software/netcdf/docs/netcdf_data_set_components.html

Unidata. n.d.b “Where is NetCDF Used?” Unidata. Accessed on June 25, 2018.

<https://www.unidata.ucar.edu/software/netcdf/usage.html>

Unidata. n.d.a “Software for Manipulating or Displaying NetCDF Data.” Unidata.

Accessed on June 25, 2018. <https://www.unidata.ucar.edu/software/netcdf/software.html>

