


NetCDF State of the Union

Ward Fisher - NetCDF Team Lead
NSF Unidata
RMAcc HPC Symposium 2026



This material is based upon work supported by NSF Unidata under award #1901712 from the U. S. National Science Foundation. Any opinions, findings and conclusions or recommendations expressed in this material do not necessarily reflect the views of the NSF.

Introduction

- In this talk, I'll be:




Introduction

- In this talk, I'll be:
 - Introducing myself and NSF Unidata



Introduction

- In this talk, I'll be:
 - Introducing myself and NSF Unidata
 - Discussing the current state of the netCDF project.



Introduction

- In this talk, I'll be:
 - Introducing myself and NSF Unidata
 - Discussing the current state of the netCDF project.
 - Provide an overview of some of our challenges, and how we're handling them.

netCDF

Introduction

- In this talk, I'll be:
 - Introducing myself and NSF Unidata
 - Discussing the current state of the netCDF project.
 - Provide an overview of some of our challenges, and how we're handling them.
 - Talk a bit about 'what's next'.

Who am I?

- Ward Fisher
- Lead NetCDF Developer at NSF Unidata
 - My background is Computer Science
 - I've been at Unidata for 14 years.



What is Unidata

One of several UCAR Community Programs. We are funded **primarily** by NSF through a core-funding proposal every five years, originally founded in 1983.

Things we do:

- Facilitate data access
- Develop and provide tools for data management, analysis, and visualization
- Provide comprehensive support, including training
- Engage in community building and advocacy, and bring the community together to address data-related issues.

Things we don't:

- Operate a data center
- Do science
- Serve K-12 education directly
- Provide operational support
- Collect data from sensors
- Develop curricula
- Compete with universities

What is NetCDF?

Released in 1989, **netCDF (Network Common Data Form)** is a set of software libraries and self-describing, machine-independent data formats designed for creating, sharing, and accessing array-oriented scientific data. It is widely used in climate, weather, oceanography, geoscience, and other fields that generate and analyze large, multi-dimensional datasets.

netCDF

What is NetCDF?

Released in 1989, **netCDF (Network Common Data Form)** is a set of software libraries and self-describing, machine-independent data formats designed for creating, sharing, and accessing array-oriented scientific data. It is widely used in climate, weather, oceanography, geoscience, and other fields that generate and analyze large, multi-dimensional datasets.

NetCDF offers an **iron-clad guarantee** that any data written will always be readable.

What is NetCDF?

- NetCDF may refer to one of three things



What is NetCDF?

Data Model

- Primarily of interest to the scientific community.



What is NetCDF?

Data Model

Data Files

- Primarily of interest to the scientific community.
- Of general interest to scientists for sharing/accessing the data model, and developers for writing netCDF data.



What is NetCDF?

Data Model

Data Files

Software Library/API

- Primarily of interest to the scientific community.
- Of general interest to scientists for sharing/accessing the data model, and developers for writing netCDF data.
- Of interest to software developers.

netCDF

What is NetCDF?

Data Model

Data Files

Software Library/API

Why does this matter?

- It's a good reminder that when working on the project roadmap, we have to consider more than just the software/API.

netCDF

What have we been doing?

- Our efforts are prioritized by community needs.



NetCDF: The Past 3-5 Years

Over the last 3–5 years, netCDF has:

- Added more extensive compression options through filter and plugin support
- Added cloud-native data support via `ncZarr`, a netCDF convention using the Zarr file format.
- Too many Quality of Life changes to enumerate. They largely boil down to improved build and test systems, migrating from configure-time configuration to run-time configuration (where possible), and leveraging functionality in downstream libraries (e.g. libhdf5).

NetCDF: The Past 3-5 Years

- Plugins and Filter Support
 - Added support for modern compression algorithms (e.g., zstandard, lossy quantization)
 - Standardized plugin discovery and configuration (build-time and runtime)
 - Easier integration and management of third-party filters and plugins
 - See <https://github.com/ccr> - Community Codec Repository for compression codecs, maintained by Ed Hartnett and Charlie Zender.

NetCDF: The Past 3-5 Years

- Support and Performance in HPC Environments
 - Optimized cache management for high-throughput workflows
 - Enhanced scalability for unlimited dimensions and complex data types
 - Better compatibility with HPC platforms and job schedulers
 - Streamlined error reporting and diagnostics in parallel environments



NetCDF: The Past 3-5 Years

- Introduced and expanded ncZarr support, enabling cloud-native netCDF data access.



NetCDF: The Past 3-5 Years

- Introduced and expanded ncZarr support, enabling cloud-native netCDF data access.
 - Uses Zarr on the back-end.



NetCDF: The Past 3-5 Years

- Introduced and expanded ncZarr support, enabling cloud-native netCDF data access.
 - Uses Zarr on the back-end.
 - Currently Zarr v2 spec compliant, work is ongoing towards the Zarr v3 spec.



NetCDF: The Past 3-5 Years

- Introduced and expanded ncZarr support, enabling cloud-native netCDF data access.
 - Uses Zarr on the back-end.
 - Currently Zarr v2 spec compliant, work is ongoing towards the Zarr v3 spec.
 - Unidata has been part of the Zarr governance committee for over 5 years, at this point.

NetCDF: The Past 3-5 Years

- I know for a fact that I have left things out, but these are the broad strokes that bring us to where we are today.



NetCDF: The Past 3-5 Years

- I know for a fact that I have left things out, but these are the broad strokes that bring us to where we are today.



Where do we go from here?

Challenges we're facing

- The biggest challenge: lack of internal resources for development efforts.



Challenges we're facing

- The biggest challenge: lack of internal resources for development efforts.

Implementing New
Features



netCDF

Challenges we're facing

- The biggest challenge: lack of internal resources for development efforts.

Implementing New Features

Expanding CI coverage with new test cases



Challenges we're facing

- The biggest challenge: lack of internal resources for development efforts.

Implementing New Features

Expanding CI coverage with new test cases

Writing really good, really cool documentation

netCDF



Challenges we're facing

- The biggest challenge: lack of internal resources for development efforts.

Implementing New Features

Writing really good, really cool documentation

Expanding CI coverage with new test cases

Exploring new technologies that might be useful

netCDF



Challenges we're facing

- The biggest challenge: lack of internal resources for development efforts.



The Solution?

Our developer Community!



The Solution?

- We have a strong community of contributors active on our Github page

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

(It can always be stronger, hint hint)

There **IS** a bottleneck, however!



The Bottleneck?

Me! (Sorry)



Introductions

- NetCDF Team Lead
 - Software Development
 - C/Fortran/C++ libraries
 - Bug Fixes
 - New Features
 - Build and Test Infrastructure
 - Documentation
 - Also:
 - User Support
 - Project Management
 - Community Relations
 - Presenting at Conferences and giving talks
- Additionally, I do non-netCDF stuff at Unidata



This material is based upon work supported by NSF Unidata under award #2403649 from the U. S. National Science Foundation. Any opinions, findings and conclusions or recommendations expressed in this material do not necessarily reflect the views of the NSF.

The Bottleneck?

Me! (Sorry)



We have a strong community of contributors active on our Github page

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

Issues, conversations, pull requests are wonderful! Send them along! They do, however, require deliberate review and consideration. **Why does it take so long?**



The Bottleneck?

Me! (Sorry)



- What is the chance that my decision will **RUIN NETCDF FOREVER?**



Future Roadmap




Future Roadmap

- Active discussion around this right now on our Github discussions page, and our community mailing lists.



Future Roadmap

- Active discussion around this right now on our Github discussions page, and our community mailing lists.
- Exists in the context of the broader NSF Unidata roadmap, which is evolving.



Future Roadmap

- Short Term: Modernization and Improvements
 - Leveraging emergent ML/AI tools
 - Leverage CMake functionality to make installing netCDF-Fortran, netCDF-C++ interface libraries easier
 - Finalizing Zarr V3 spec functionality.

netCDF

Future Roadmap

- Medium Term:
 - Examining features in HDF5 2.0 (particularly those related to HPC performance).
 - Implementation of the Zarr v3 spec.



Future Roadmap


- Long Term:
 - Unidata is ending the 2nd year of our 5 year funding cycle.
 - We have lots to do as part of our current proposal, largely focused on cloud-based technologies and data access.

netCDF



Future Roadmap

- Long Term:
 - Evolving cloud and HPC technologies, emergent machine-learning based tech, etc.
 - A lot of the short term modernization work is being done to make the long-term work, whatever it ends up being, more viable.



In Summary

NetCDF, and NSF Unidata, remain healthy.



<https://linktr.ee/netcdf>

In Summary

NetCDF, and NSF Unidata, remain healthy.

Relevance doesn't require cutting edge features, it requires useful features that will benefit the community at large.



<https://linktr.ee/netcdf>