


NetCDF-4: A New Data Model, Programming Interface, and Format Using HDF5


Russ Rew, Ed Hartnett, John Caron
UCAR Unidata Program Center
Mike Folk, Robert McGrath, Quincey Koziol
NCSA and The HDF Group, Inc.

Final Project Review, August 9, 2005



THG,
Inc.

Motivation: Why is this area of work important?



Jim Gray,
Distinguished
Engineer at
Microsoft,
1998 Turing Award
winner

While the commercial world has standardized on the relational data model and SQL, no single standard or tool has critical mass in the scientific community. There are many parallel and competing efforts to build these tool suites – at least one per discipline. Data interchange outside each group is problematic. In the next decade, as data interchange among scientific disciplines becomes increasingly important, a common HDF-like format and package for all the sciences will likely emerge.

“Scientific Data Management in the Coming Decade,” Jim Gray, David T. Liu, Maria A. Nieto-Santisteban, Alexander S. Szalay, Gerd Heber, David DeWitt, Cyberinfrastructure Technology Watch Quarterly, Volume 1, Number 2, February 2005

Preservation of scientific data



MacKenzie Smith,
Associate Director
for Technology at
the MIT Libraries,
Project director at
MIT for DSpace, a
groundbreaking
digital repository
system

... the ephemeral nature of both data formats and storage media threatens our very ability to maintain scientific, legal, and cultural continuity, not on the scale of centuries, but considering the unrelenting pace of technological change, from one decade to the next. ... And that's true not just for the obvious items like images, documents, and audio files, but also for scientific images, ... and simulations. In the scientific research community, standards are emerging here and there — HDF (Hierarchical Data Format), NetCDF (network Common Data Form), FITS (Flexible Image Transport System) — but much work remains to be done to define a common cyberinfrastructure.

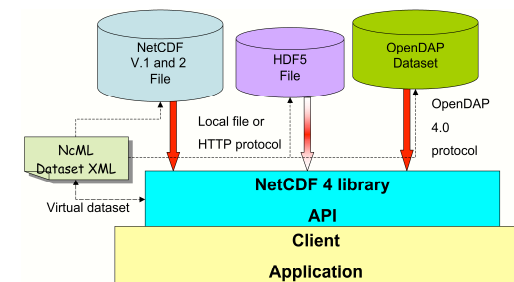


“Eternal Bits: How can we preserve digital files and save our collective memory?,” MacKenzie Smith, IEEE Spectrum, July 2005

Overview

- Background: What are Unidata, netCDF, HDF5, netCDF-4?
- What were project's goals?
- What was accomplished?
- What remains to be done?
- How soon will netCDF-4 reach TRL-7?
- Are the benefits worth the cost?
- What follow-on activities will continue?

Unidata: A Community Endeavor

- Community of educators and researchers at 120 universities, 30 other institutions, international in scope
- Managed by the University Corporation for Atmospheric Research
- Mission: *providing data, tools, support, and community leadership for enhanced earth-system education and research*
- Atmospheric science community, expanding to oceanography, hydrology, other geosciences
- Unidata Program Center: 25 staff, 15 developers



What are netCDF and HDF5?

- **Data Models** for science: useful abstractions for variables, dimensions, attributes, and coordinates
- **Application Programming Interfaces** for storing and accessing scientific data in programs in C, Fortran, Java, C++, Perl, Python, ...
- **File Formats** for self-describing portable binary data

Most users need not know any details about the formats to access netCDF or HDF5 data

Why file formats instead of databases?

- Traditional database systems have lacked
 - support for N-dimensional arrays
 - good tools for scientific analysis and visualization
 - ability to handle large data volumes efficiently using common access patterns in scientific programs
 - simple programming language interfaces for data access
- Unlike database systems, files do not require
 - the expertise of a separate database administrator
 - understanding database features such as query languages, schema declarations, nested transactions, ...
- *“Some scientists use databases for some of their work, but as a general rule, most scientists do not ... databases have to improve a lot before they are worth a second look.”* Jim Gray, et al

Scientific data access requirements

- Preserving backward compatibility, for both APIs and format, is *sacrosanct*.
- Simplicity of the interface and generality for multiple disciplines are also desirable.
- Scientific data is most useful if it is:

<i>self-describing</i>	for independent use
<i>portable</i>	for current and future platforms
<i>directly accessible</i>	for efficient access to subsets
<i>appendable</i>	for incremental creation
<i>sharable</i>	for concurrent access and writing
<i>archivable</i>	for future uses of past archives

NetCDF-3 and HDF5

	NetCDF-3	HDF5
Availability	Free	Free
Development and maintenance	UCAR Unidata	NCSA, HDF Group
Primary funding	NSF	NASA, DOE
Advantages	Popular, simple, lots of tools, multiple implementations	Powerful, high-performance, efficient for storage, extensible
Primary uses	Climate, forecast, ocean models, data archives, remote access	Satellite data, computational fluid dynamics, parallel computing

History of netCDF



netCDF developed
at Unidata

1988

1991

1996

netCDF 2.0
released



netCDF 3.0
released

2004

2005

netCDF 3.6.0
released



netCDF 4.0
alpha released

Goals of netCDF/HDF combination

- Create **netCDF-4**, combining desirable characteristics of netCDF-3 and HDF5, while taking advantage of their separate strengths
 - Widespread use and simplicity of netCDF-3
 - Generality and performance of HDF5
- Make netCDF more suitable for high-performance computing, large datasets
- Provide simple high-level application programming interface (API) for HDF5
- Demonstrate benefits of combination in advanced Earth science modeling efforts

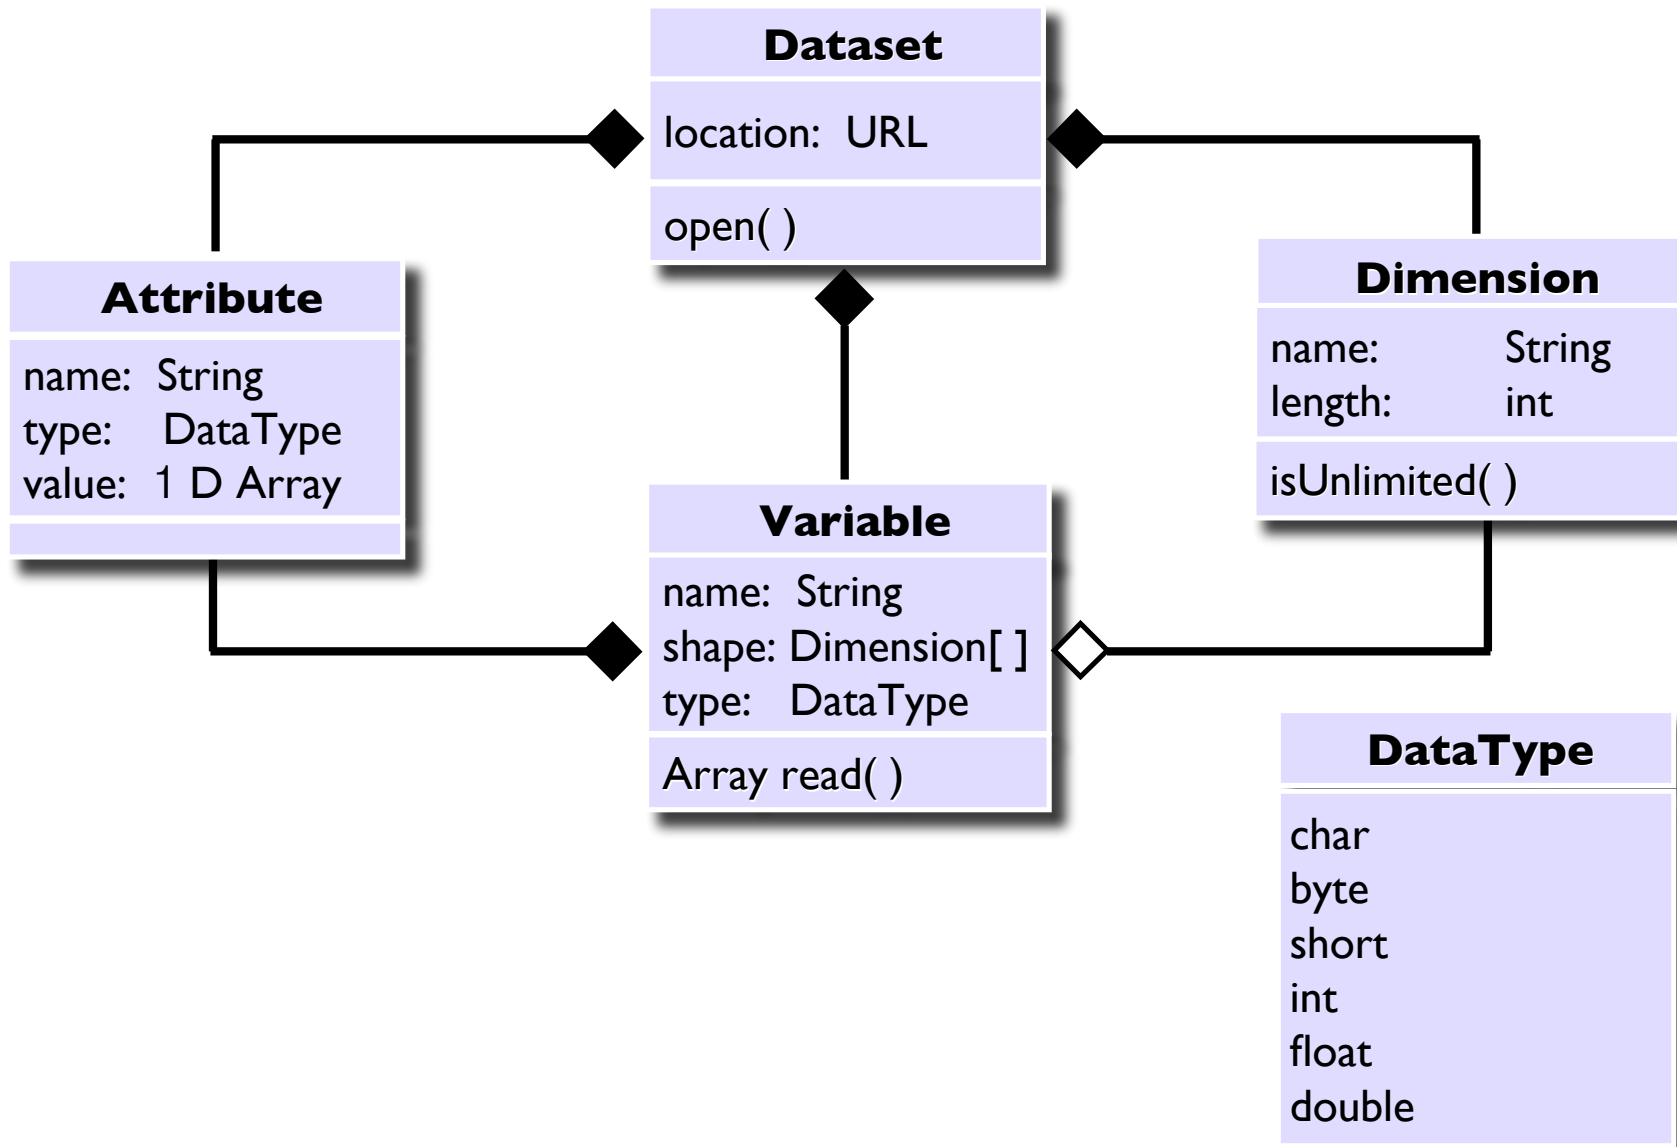
What is netCDF-4?

- A NASA-funded effort to improve
 - Interoperability among scientific data representations
 - Integration of observations and model outputs
 - I/O for high-performance computing
- A new data model for scientific data
- A set of documented programming interfaces (APIs) for using the model
- Freely available software implementing the netCDF-4 APIs, extending netCDF-3, and using HDF5 for storage
- A new format for netCDF data based on HDF5

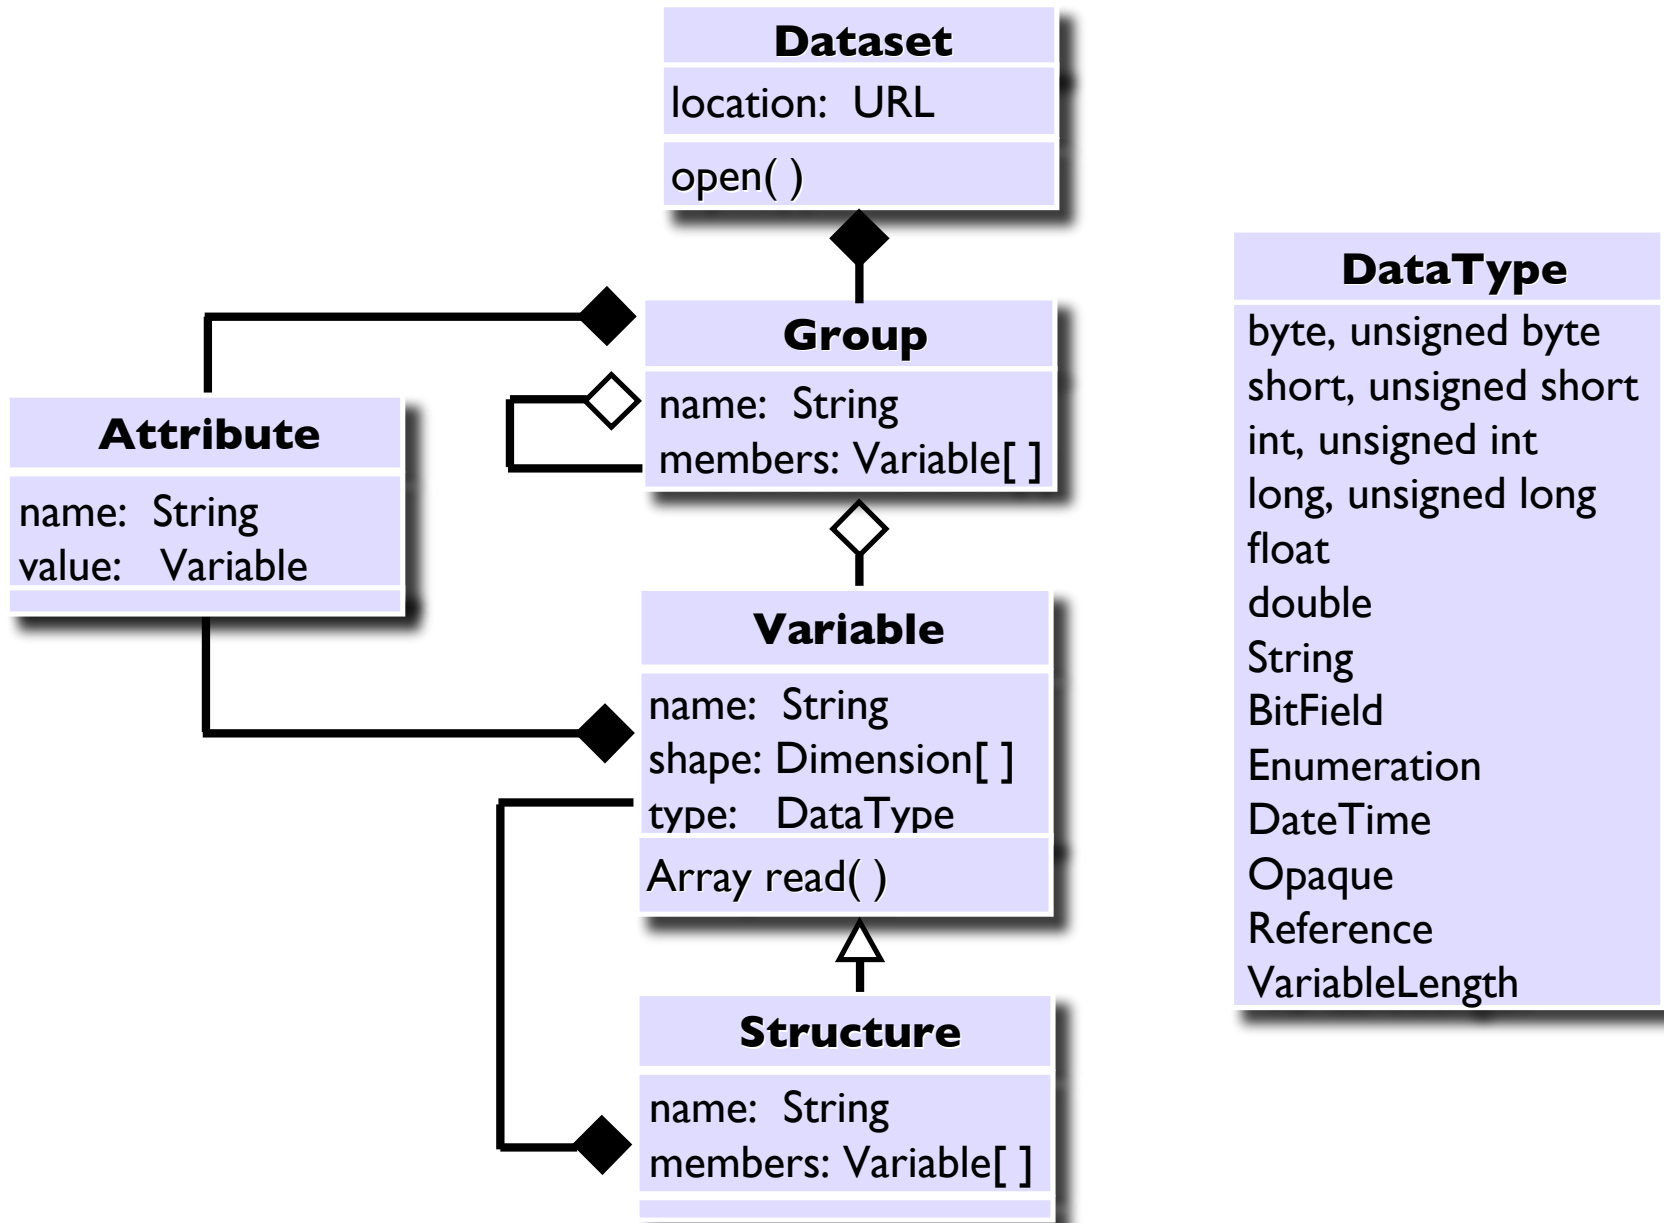
NetCDF-3 and NetCDF-4 Data Models

- NetCDF-3 models multidimensional arrays of primitive types with Variables, Dimensions, and Attributes, with one unlimited dimension
- NetCDF-4 implements an extended data model with enhancements made possible with HDF5:
 - Structure types: like C structures, except portable
 - Multiple unlimited dimensions
 - Groups: containers providing hierarchical scopes for variables, dimensions, attributes, and other Groups
 - Variable-length objects: for soundings, ragged arrays, ...
 - New primitive types: Strings, unsigned types, opaque

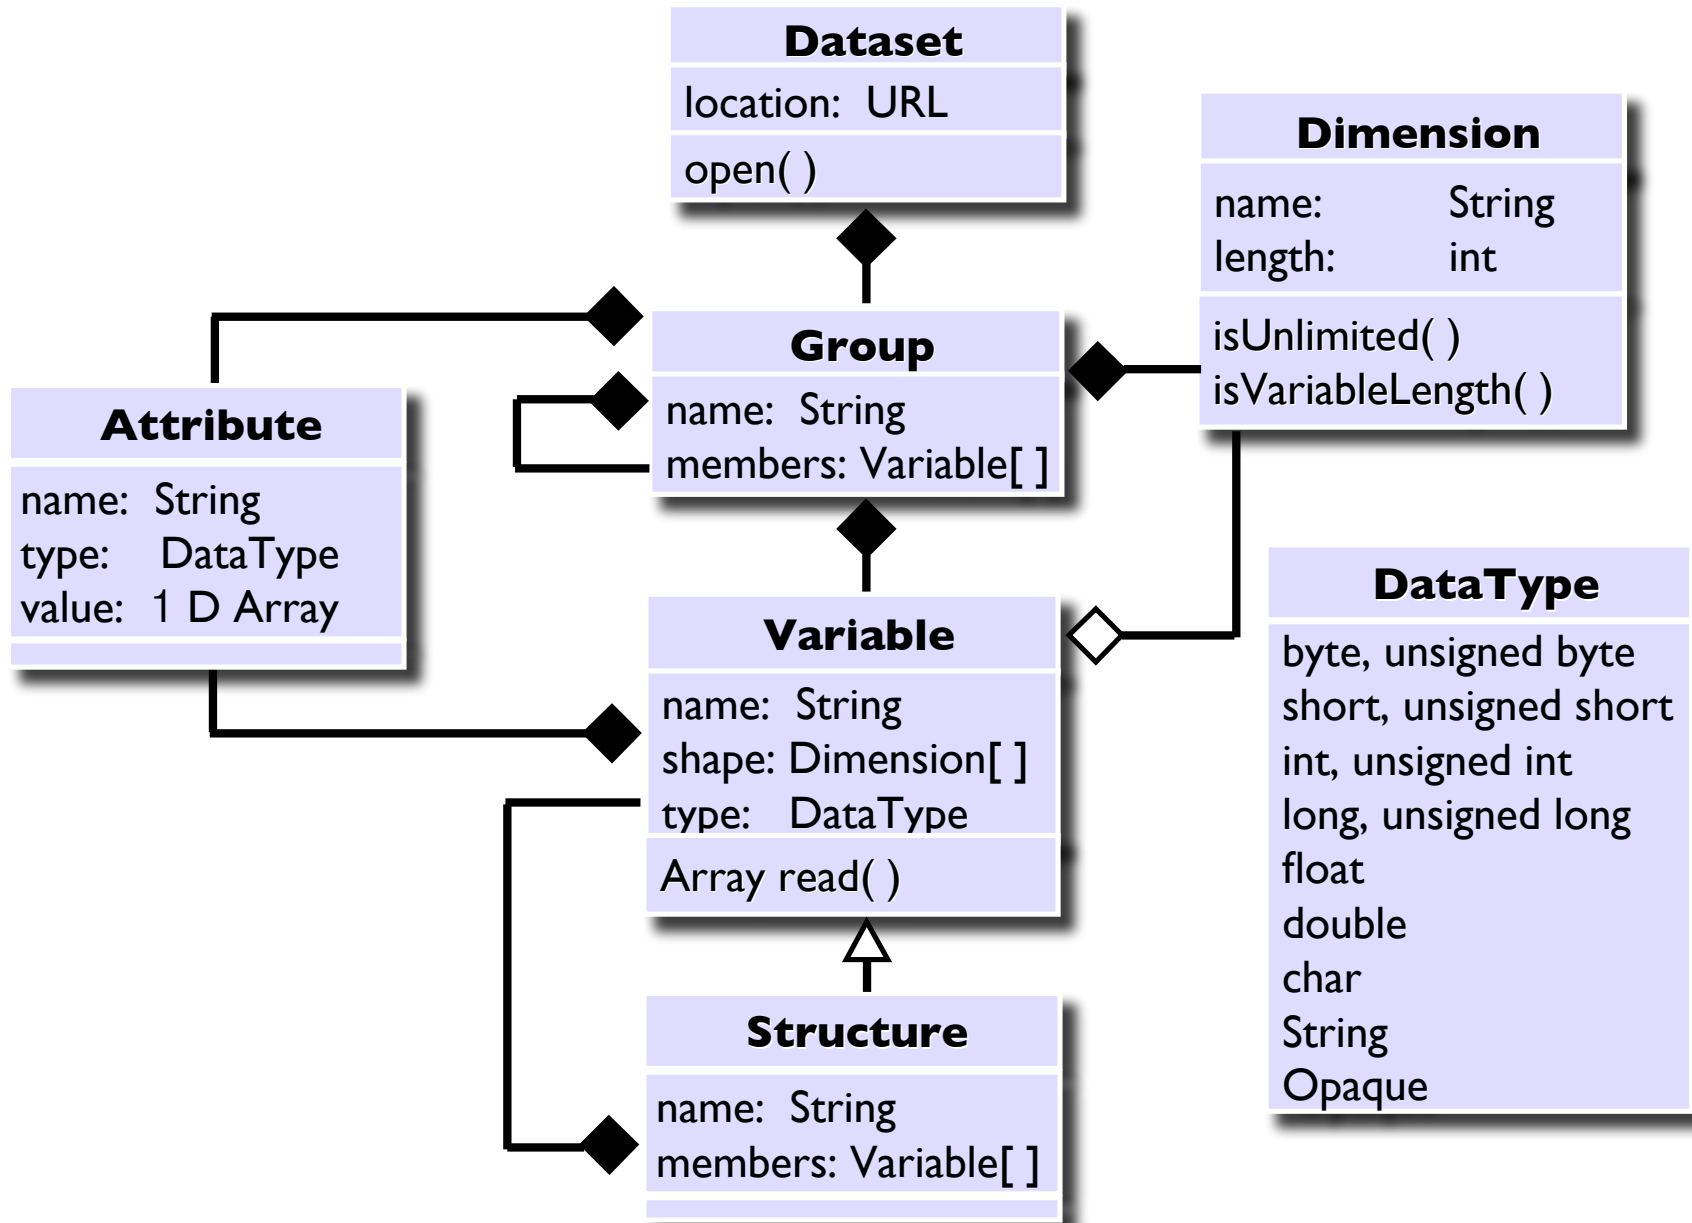
NetCDF-3 Data Model



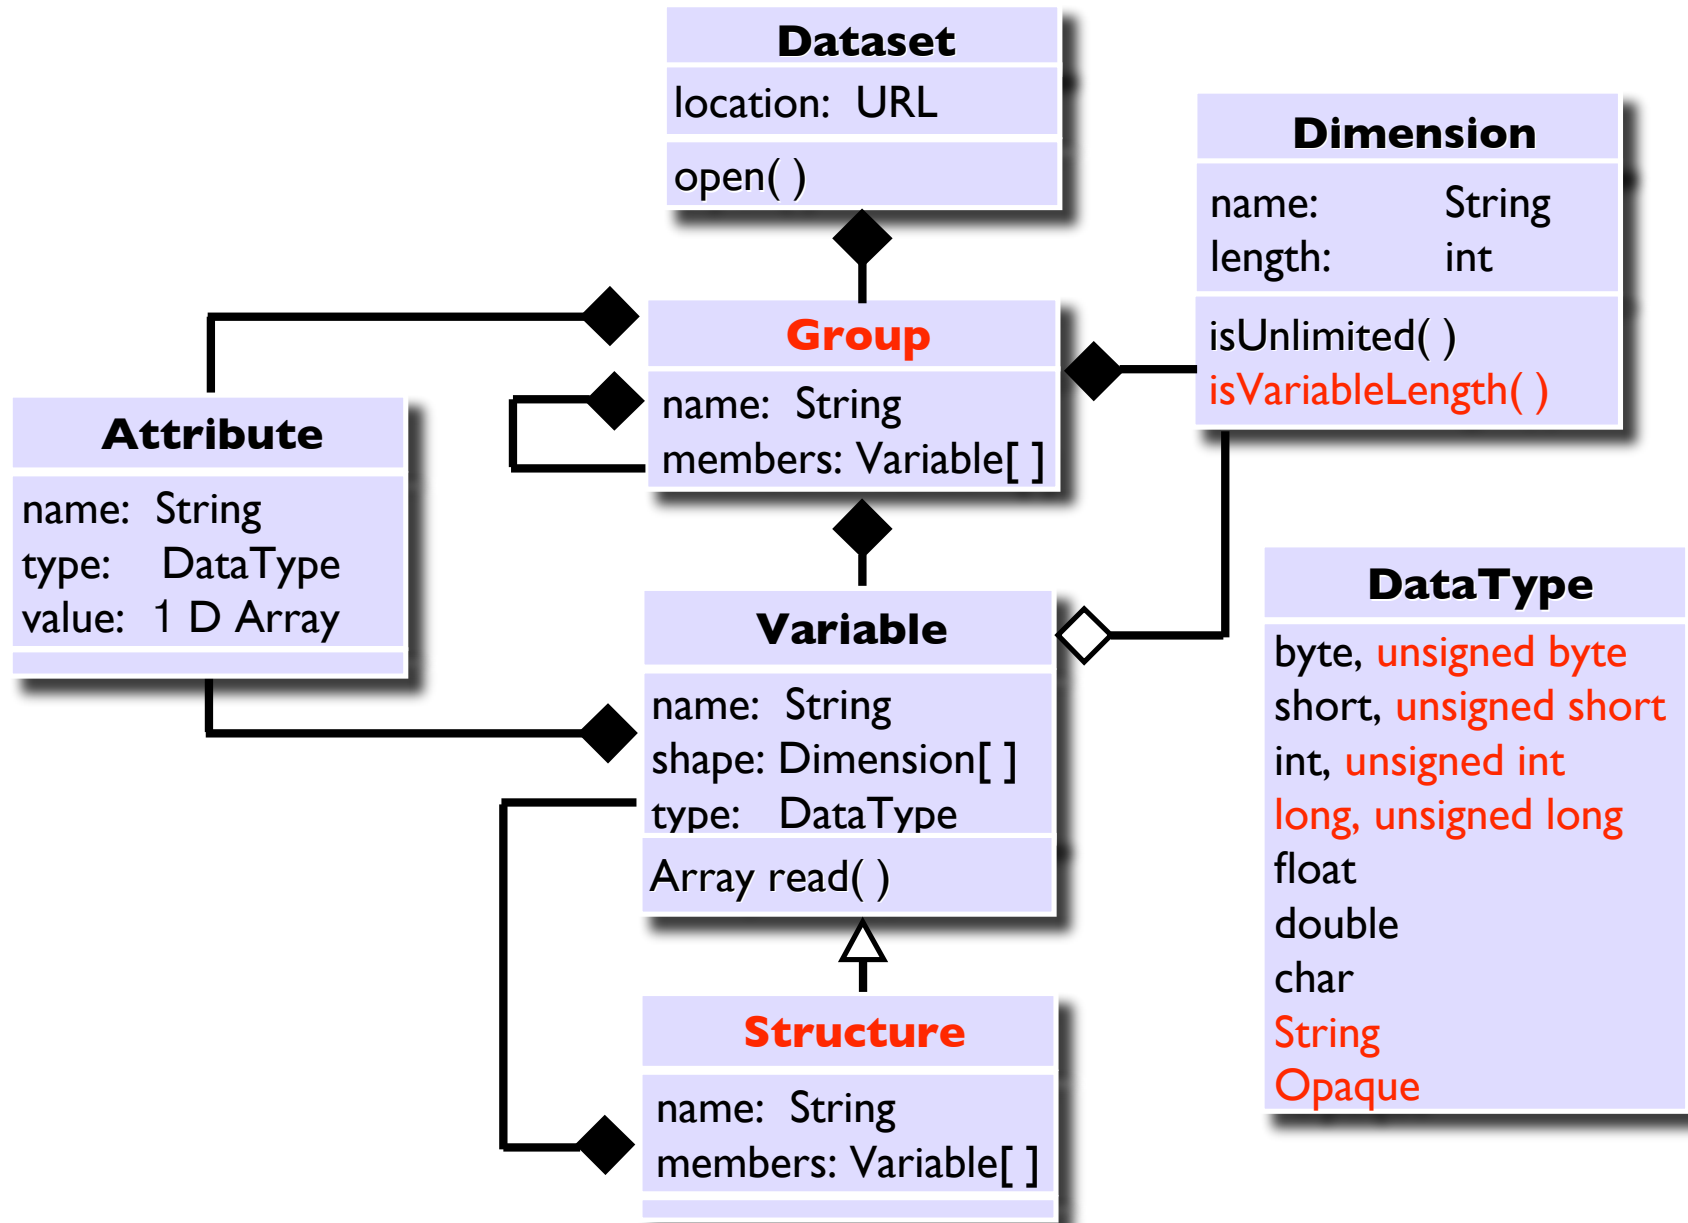
HDF5 Data Model



A Common Data Model



NetCDF-4 Data Model



The Common Data Model

- NetCDF, HDF5, and OPeNDAP developers have begun to discuss moving towards this Common Data Model, providing
 - useful mappings among the three data models
 - opportunities to tweak the data models to mitigate differences
 - a plan to make OPeNDAP the remote access protocol for netCDF-4 and netCDF-4 the persistence format for OPeNDAP
- This is an important long-term effort.

Accomplishments

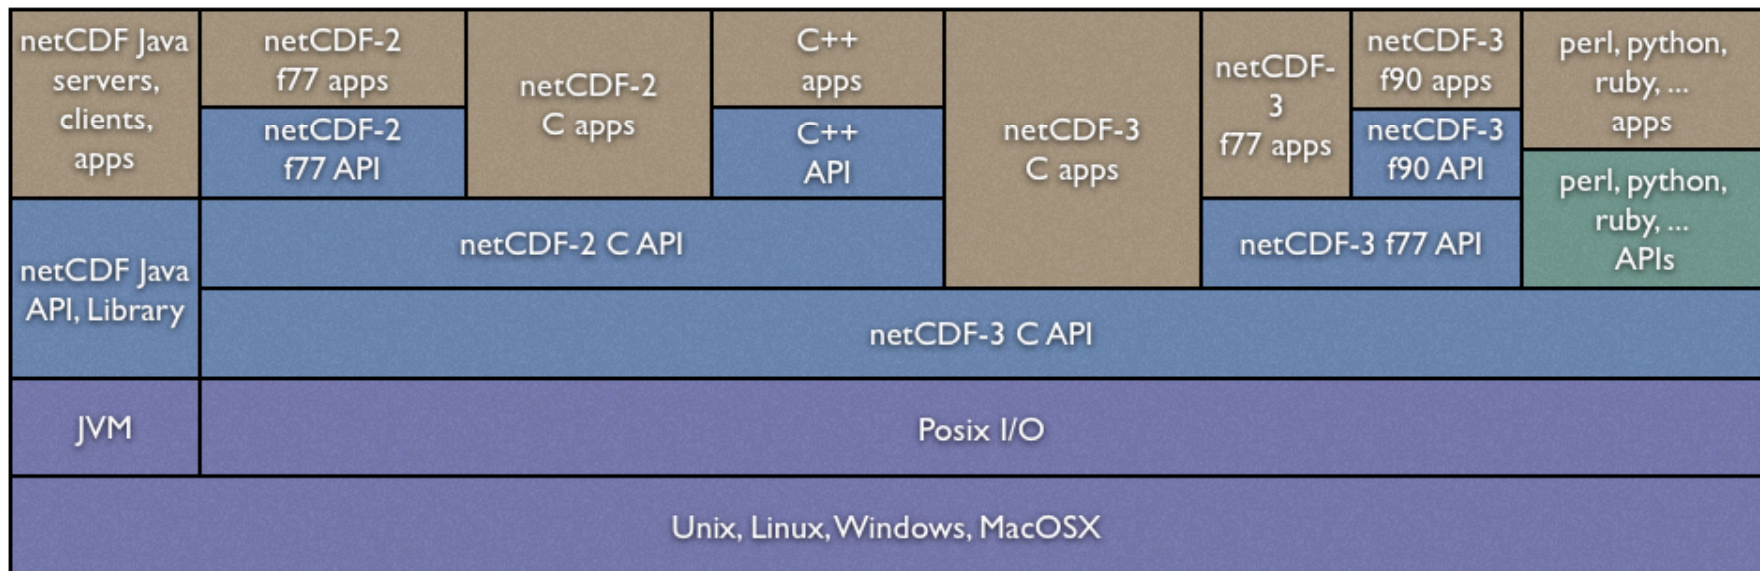
- Design and documentation of netCDF-4 data model
- Implementation of complete support for netCDF-3 API over HDF5 storage layer
- Prototyped netCDF-4 features in netCDF Java
- Implemented netCDF-4 data model over HDF5, including following additions:
 - Parallel I/O interfaces
 - Multiple dynamic dimensions
 - New unsigned integer data types
 - Use of chunking (multidimensional tiling)
 - Dynamic schema modification
 - Groups
 - User-defined compound types (portable C structures)

More accomplishments

- Re-engineered software architecture:
 - Use of autoconf, automake, libtool consistent with HDF5
 - Designed and wrote many new unit tests
- Refactored, converted, and rewrote documentation:
 - Changed from FrameMaker to texinfo and automatically generated HTML, PDF, and info documents
 - Provided new language-independent NetCDF Users Guide
- Determined needed HDF5 enhancements and implemented most of them
 - Dimension scales, for coordinate variables
 - Integer to float conversions during I/O
- Large File Support added to netCDF 3.6 release (users just couldn't wait)
- Better interoperability with HDF5 than planned: can access HDF5 data that uses HDF5 1.8 "Dimension Scales" feature
- Talks with ESRI resulted in netCDF support in ArcGIS 9.2 (a million new netCDF users)

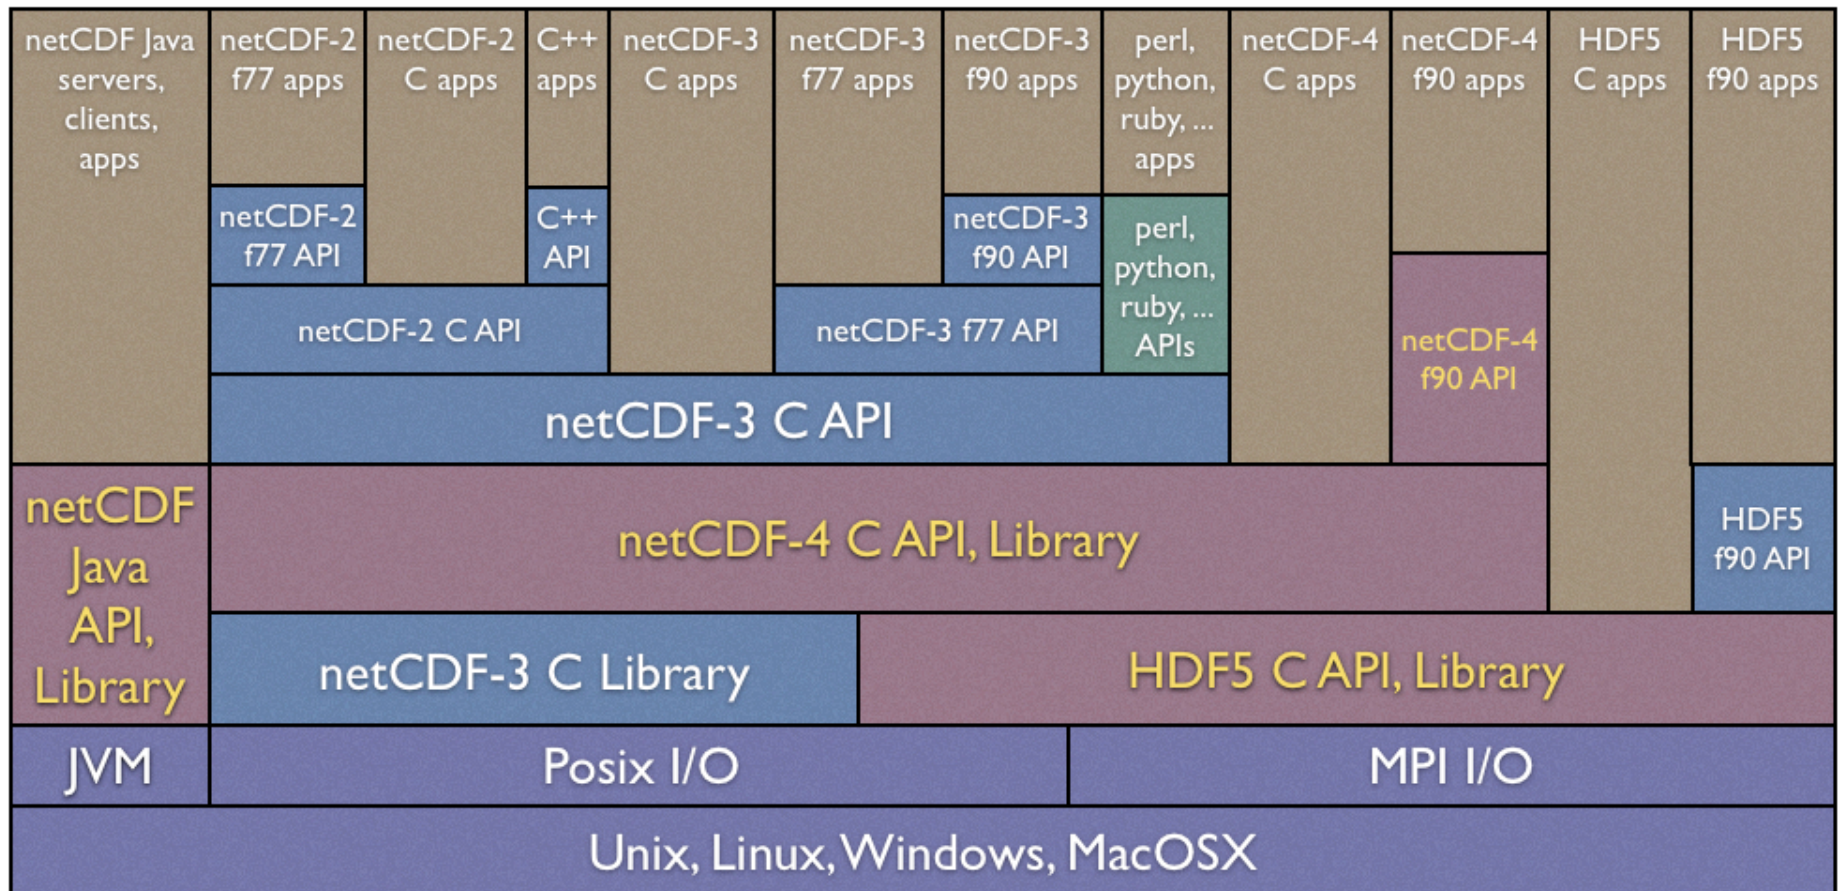
NetCDF-3 Software Architecture

- Core of netCDF-3 is C library, supporting f77, C++, f90, and most other language interfaces
- Java netCDF library is an independent implementation that uses same format



NetCDF-4 Software Architecture

- The netCDF-4 project proposed new C, f90 layers and HDF5 enhancements



- Java netCDF developments have tested usefulness, practicality of Common Data Model for netCDF-4

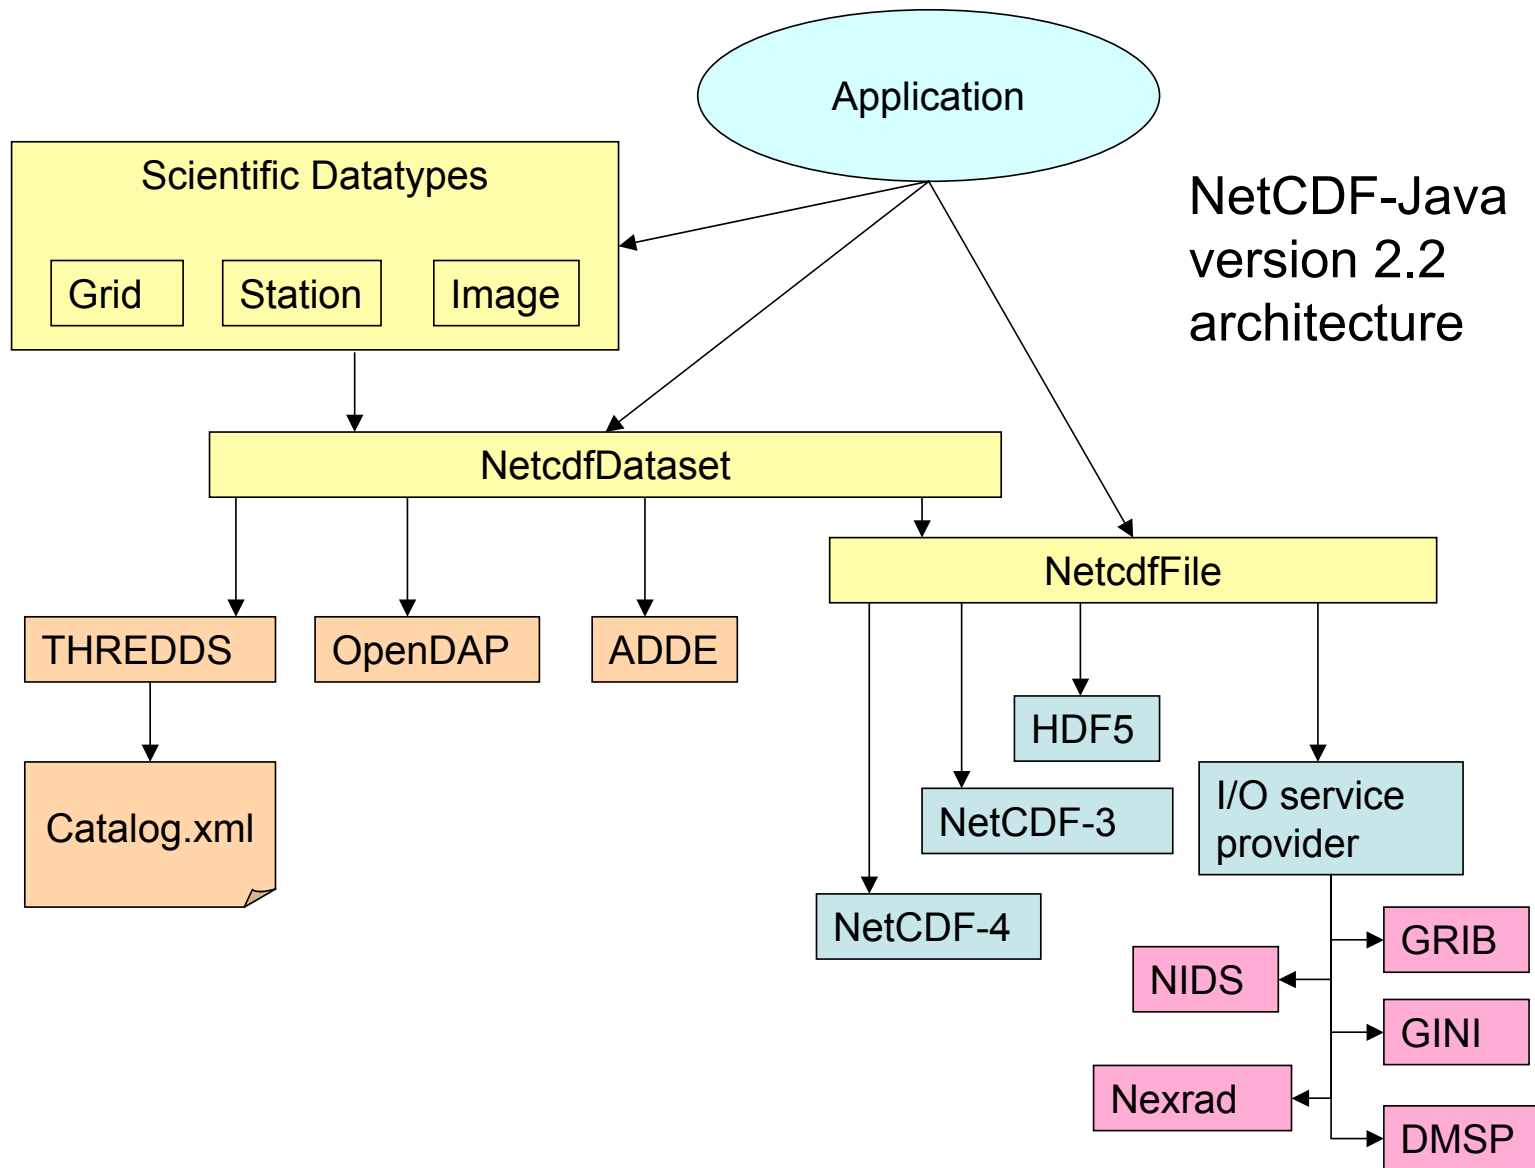
How Are the APIs Changing?

- Current APIs for C, Fortran, Java, and C++ will continue to be supported
- NetCDF-4 features will initially be available only for C and Java interfaces, followed by Fortran-90 and eventually C++
- Access from Fortran-77 to most netCDF-4 features is limited (Structures, for example)
- Advanced Java features are being moved to C-based interfaces during the next year

Advanced Features of Java Interface

- Client access to data servers:
 - HTTPD
 - OPeNDAP
- Java netCDF version 2.2 (in beta release) implements
 - NetCDF-4 Data Model
 - Coordinate system support for general and georeferenced coordinates
 - I/O Framework providing netCDF interface to data in other formats: GRIB, HDF5, GINI, NEXRAD, ...
- Access through NcML virtual datasets to add metadata, aggregate data, subset

NetCDF Java




NetCDF-4 Formats

- Still supports *classic* XDR-based format (1988) and *64-bit offset* format variant (2004)
- New *netCDF-4* format uses HDF5 representation to support
 - Appending along multiple unlimited dimensions
 - Dynamic schema modification
 - Per-variable chunking (tiled storage)
 - Per-variable compression
 - Unicode names
 - “Reader makes right” conversions
- For maximum interoperability with existing operational systems, classic format should still be used, but software transparently supports all three format variants

What remains to be done?

- Release of HDF5 1.8.0, originally expected in July 2005:
 - Access of HDF5 objects in a Group by creation order
 - Bug fixes related to parallel I/O
 - HDF 1.8 enhancements are required for netCDF-4
- Completion of netCDF-4 f90 interface
- Demonstration of netCDF-4 benefits in advanced modeling efforts by enticing WRF and CCSM model developers to test beta release with parallel I/O. Obstacles include:
 - Adequacy of new Argonne/Northwestern pnetcdf 1.0
 - Other priorities higher than improving I/O performance
 - Desire of developers to wait for real release, complete f90 interface
- Provide packed data type as originally envisioned
 - Lack is result of misunderstanding about HDF5 packed bit type

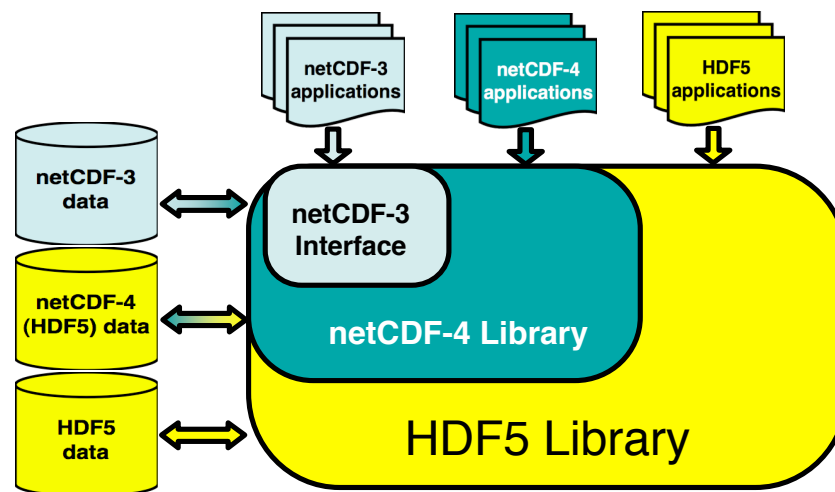


Merging the NetCDF and HDF5 Libraries to Achieve Gains in Performance and Interoperability

PI: Russell K. Rew, UCAR/Unidata

Description and Objectives

- Extend and merge the Network Common Data Form (netCDF) library and the Hierarchical Data Format-5 (HDF5) library to facilitate access to scientific data and the integration of observations with model representations in multiple disciplines
- Benefit science community by making available packed and larger data sets, providing parallel I/O and greater data management, analysis, and visualization capabilities, and a simpler high-level interface for scientific data



Approach

- Implement netCDF-3 using the public HDF5 API
- Design netCDF-4 API, determining any needed HDF5 additions
- Implement needed HDF5 enhancements
- Implement netCDF-4 using HDF5 as its storage layer, exploiting HDF5 parallel I/O, compound types, chunking
- Test and tune netCDF-4 to achieve efficient I/O performance
- Demonstrate effectiveness of merged software in models

Co-Is/Partners

Mike Folk, NCSA

Science Themes

Atmospheric Composition Carbon cycle
 Climate Solid Earth
 Water & Energy Cycle Weather

Schedule and Deliverables

- Detailed design of netCDF4 (RFC document) (12/03)
- Initial prototype of core library (3/04)
- Parallel I/O support, additional types (10/04)
- Beta release of netCDF-4 as soon as HDF5 allows
- Release of netCDF-4 following HDF5 1.8.0 release

Application/Mission

- Supports scientific data storage, exchange, access, analysis, discovery and visualization using free and open technologies
- Cross-disciplinary research

TRL=5

ESTO

Earth Science Technology Office

How soon will netCDF-4 reach TRL-7?

- Requires release of HDF 1.8 (currently estimated for January 2006)
- A netCDF-4 beta release will be available as soon as HDF5 permits (estimated after October 2005)
- Delay will provide opportunity to
 - finish full f90 API
 - add more Common Data Model tests
 - implement ncdump and ncgen utilities that understand netCDF-4 enhancements
- When integrated into WRF or CCSM models, will be promoted to TRL-7

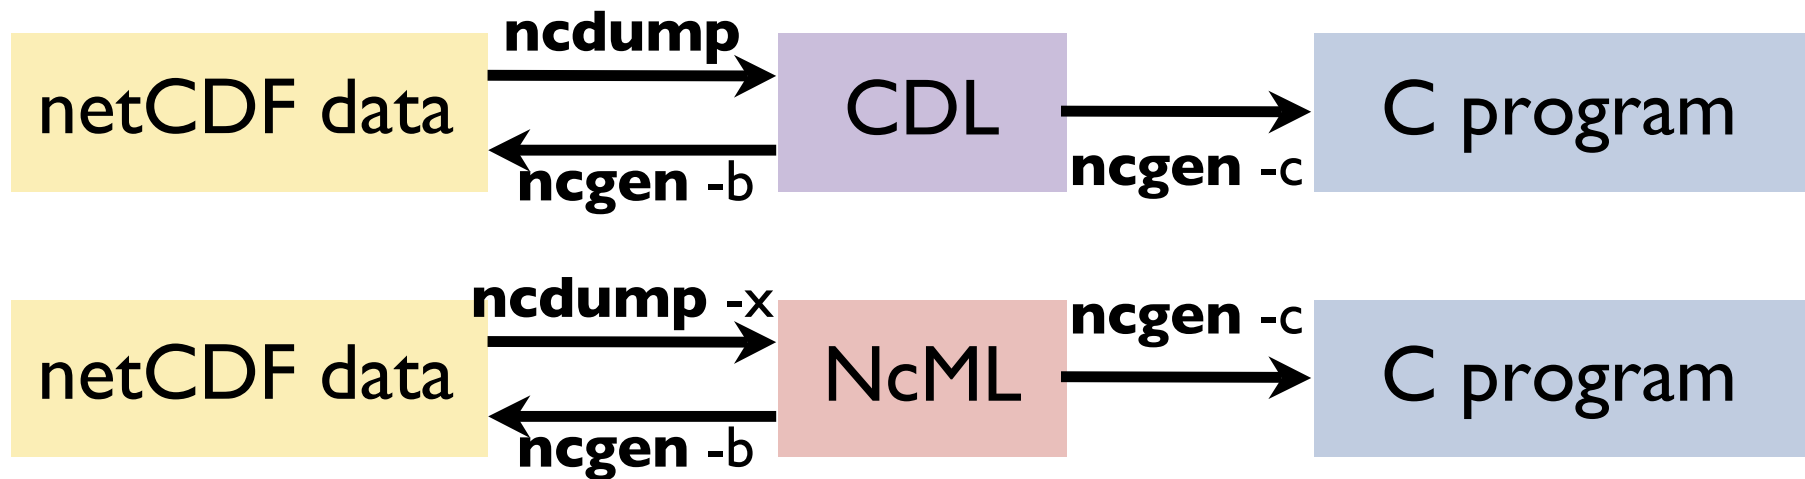
Why not release netCDF-4 beta now?

- Current alpha release must use artifacts to emulate HDF5 enhancements, like access by creation order.
- The artifacts define yet another format, netCDF-4-alpha, that we would rather not continue to support.
- Testers of the alpha release are warned that the beta release and subsequent releases will not correctly read files created with the alpha release that contain development artifacts.

ncdump, ncgen, CDL, and NcML

As resources permit:

- ncdump and ncgen utilities will handle netCDF-4 groups, structs, and new data types
- ncdump and ncgen will support optional use of NcML dialect of XML instead of CDL



What follow-on activities will continue?

- Development and support of HDF5 is the mission of The HDF Group:
to sustain the HDF technologies and to support worldwide HDF user communities with production-level software and services
- Further development and support of netCDF is in Unidata's core mission:
providing data, tools, and community leadership for enhanced Earth-system education and research
- Plans beyond the initial release of netCDF-4 include:
 - Moving Java advanced features to C interface, including access through NcML
 - Providing an extensive set of examples in various language interfaces
 - Designing and implementing a new C++ interface

Papers, Posters, Presentations

2 papers, 5 posters, and 6 presentations:

- E. Hartnett: Introduction to NetCDF Classic and to NetCDF-4, *Extreme I/O Workshop*, San Diego Supercomputing Center, July 2005, presentation.
- R. Rew: The Future of netCDF. *GO-ESSP Workshop 4*, British Atmospheric Data Centre, England, June 2005, presentation.
- J. Caron: NetCDF-Java prototype for a Common Data Model. *HDF/HDF-EOS Workshop VIII*, Aurora, Colorado, October 2004. Poster and presentation.
- E. Hartnett: Merging the NetCDF and HDF5 Libraries to Achieve Gains in Performance and Interoperability. *HDF/HDF-EOS Workshop VIII*, Aurora, Colorado, October 2004. Poster and presentation.
- R. Rew, M. Folk, E. Hartnett, and R. McGrath: Plans for an Enhanced NetCDF-4 Interface to HDF5 Data. *HDF/HDF-EOS Workshop VII*, Silver Springs, September 2003. Poster and presentation.
- R. Rew and E. Hartnett: Merging NetCDF and HDF5. *20th International Conference on Interactive Information Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, Seattle, January 2004. Paper and poster.
- E. Hartnett: Merging the NetCDF and HDF5 Libraries to Achieve Gains in Performance and Interoperability. *2004 Earth Science Technology Conference, Palo Alto, June 2004*. Paper and presentation.
- M. Folk, R. Rew, K. Yang, R. McGrath: NetCDF-4: Combining netCDF and HDF5 Data. *AGU Fall Meeting*, San Francisco, December 2003. Poster.