

Open Geospatial Consortium

Approval Date: 07 September 2012

Publication Date: 02 October 2012

External identifier of this OGC® document: <http://www.opengis.net/doc/IS/netcdf-enhanced/1.0>

Internal reference number of this OGC® document: OGC 11-038r2

Version: 1.0

Category: OGC® Standard

Editor: Ben Domenico

OGC Network Common Data Form (NetCDF) NetCDF Enhanced Data Model Extension Standard

Copyright notice

Copyright © 2012 Open Geospatial Consortium.
To obtain additional rights of use, visit <http://www.opengeospatial.org/legal/>.

Warning

This document is an OGC Member approved international standard. This document is available on a royalty free, non-discriminatory basis. Recipients of this document are invited to submit, with their comments, notification of any relevant patent rights of which they are aware and to provide supporting documentation.

Document type: OGC® Standard
Document subtype:
Document stage: Approved for public release
Document language: English

License Agreement

Permission is hereby granted by the Open Geospatial Consortium, ("Licensor"), free of charge and subject to the terms set forth below, to any person obtaining a copy of this Intellectual Property and any associated documentation, to deal in the Intellectual Property without restriction (except as set forth below), including without limitation the rights to implement, use, copy, modify, merge, publish, distribute, and/or sublicense copies of the Intellectual Property, and to permit persons to whom the Intellectual Property is furnished to do so, provided that all copyright notices on the intellectual property are retained intact and that each person to whom the Intellectual Property is furnished agrees to the terms of this Agreement.

If you modify the Intellectual Property, all copies of the modified Intellectual Property must include, in addition to the above copyright notice, a notice that the Intellectual Property includes modifications that have not been approved or adopted by LICENSOR.

THIS LICENSE IS A COPYRIGHT LICENSE ONLY, AND DOES NOT CONVEY ANY RIGHTS UNDER ANY PATENTS THAT MAY BE IN FORCE ANYWHERE IN THE WORLD.

THE INTELLECTUAL PROPERTY IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF THIRD PARTY RIGHTS. THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE DO NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE INTELLECTUAL PROPERTY WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION OF THE INTELLECTUAL PROPERTY WILL BE UNINTERRUPTED OR ERROR FREE. ANY USE OF THE INTELLECTUAL PROPERTY SHALL BE MADE ENTIRELY AT THE USER'S OWN RISK. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR ANY CONTRIBUTOR OF INTELLECTUAL PROPERTY RIGHTS TO THE INTELLECTUAL PROPERTY BE LIABLE FOR ANY CLAIM, OR ANY DIRECT, SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM ANY ALLEGED INFRINGEMENT OR ANY LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR UNDER ANY OTHER LEGAL THEORY, ARISING OUT OF OR IN CONNECTION WITH THE IMPLEMENTATION, USE, COMMERCIALIZATION OR PERFORMANCE OF THIS INTELLECTUAL PROPERTY.

This license is effective until terminated. You may terminate it at any time by destroying the Intellectual Property together with all copies in any form. The license will also terminate if you fail to comply with any term or condition of this Agreement. Except as provided in the following sentence, no such termination of this license shall require the termination of any third party end-user sublicense to the Intellectual Property which is in force as of the date of notice of such termination. In addition, should the Intellectual Property, or the operation of the Intellectual Property, infringe, or in LICENSOR's sole opinion be likely to infringe, any patent, copyright, trademark or other right of a third party, you agree that LICENSOR, in its sole discretion, may terminate this license without any compensation or liability to you, your licensees or any other party. You agree upon termination of any kind to destroy or cause to be destroyed the Intellectual Property together with all copies in any form, whether held by you or by any third party.

Except as contained in this notice, the name of LICENSOR or of any other holder of a copyright in all or part of the Intellectual Property shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Intellectual Property without prior written authorization of LICENSOR or such copyright holder. LICENSOR is and shall at all times be the sole entity that may authorize you or any third party to use certification marks, trademarks or other special designations to indicate compliance with any LICENSOR standards or specifications. This Agreement is governed by the laws of the Commonwealth of Massachusetts. The application to this Agreement of the United Nations Convention on Contracts for the International Sale of Goods is hereby expressly excluded. In the event any provision of this Agreement shall be deemed unenforceable, void or invalid, such provision shall be modified so as to make it valid and enforceable, and as so modified the entire Agreement shall remain in full force and effect. No decision, action or inaction by LICENSOR shall be construed to be a waiver of any rights or remedies available to it.

Table of Contents

Abstract.....	1
i. Keywords	2
ii. Preface.....	2
iii. Submitting organizations.....	2
iv. Submitters.....	2
1. Scope	3
2. Conformance	3
3. References	4
4. Terms and Definitions	4
5. Conventions.....	7
6. Clauses not Containing Normative Material	8
7. Clauses containing normative material.....	8
7.1 Requirement Class netcdf-enhanced	9
7.1.1 NetCDF classic data model.....	9
7.1.2 Data types.....	10
7.1.3 Top-level, unnamed group	10
7.1.4 Groups.....	10
7.1.5 Variables	10
7.1.6 Dimensions	10
8. Media Types for any data encoding(s)	11
9. References	11
Annex A: Conformance Class Abstract Test Suite (Normative)	11
A.1 Conformance class: netcdf-enhanced	12
A.1.1 Conformance Test 1: http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/data-model-elements	12
A.1.2 Conformance Test 2: http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/classic-data-model	12
A.1.3 Conformance Test 3: http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/data-types	12
A.1.4 Conformance Test 4: http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/primitive-data-types	12
A.1.5 Conformance Test 5: http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/user-defined-data-types	12
A.1.6 Conformance Test 6: http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/top-level-group	13
A.1.7 Conformance Test 7: http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/shared-dimensions	13
Annex B: Revision history.....	14
Figure 1: UML Diagram for netCDF Enhanced Data Model	9

Abstract

When there are multiple valid representations for the same information, humans may understand and ignore gratuitous differences. For example, a human reader recognizes that the strings "Temperature" and "temperature" very likely refer to the same quantity. It is impractical, however, to develop software that has to deal with multiple representations of the same information. A *convention* that declares one representation is favored above others (even though all representations may be technically equivalent) makes interoperable software practical, because it only needs to deal with one way to represent the information.

Standard data models and data storage formats allow for the creation of data conventions that foster this interoperability. A *data model* is an abstract conceptual layer that describes how different types of data should be represented, independent of any particular software application or hardware configuration. A *data format* is a description of how data organized according to a particular data model should be stored on physical storage media by software applications. *Data conventions* can then build upon the infrastructure provided by the data model and data format to imbue specific data collections with meaning that is understood by a wide variety of software applications.

This standard deals with enhancements to the netCDF (Network Common Data Form) data model for array-oriented scientific data..

Two important data models for netCDF are:

- the “classic” netCDF model, used for netCDF-3 and earlier versions
- an enhanced data model, used in netCDF-4 and later versions.

The netCDF classic data model is defined in OGC 10-091r3, “NetCDF Core.”

This document specifies a netCDF extension standard for the enhanced data model. The OGC netCDF encoding supports electronic encoding of geospatial data, specifically digital geospatial information representing space- and time-varying phenomena.

NetCDF (network Common Data Form) is a data model for array-oriented scientific data. The netCDF classic data model is specified in the netCDF core specification. This standard specifies the enhanced data model. A freely distributed collection of access libraries implementing support for that data model in a machine-independent format are available. Together, the interfaces, libraries, and format support the creation, access, and sharing of multi-dimensional scientific data.

i. Keywords

ogcdoc, netcdf, space-time, netcdf-enhanced

ii. Preface

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. The Open Geospatial Consortium shall not be held responsible for identifying any or all such patent rights.

Recipients of this draft document are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of the standard set forth in this document, and to provide supporting documentation.

iii. Submitting organizations

The following organizations submitted this Document to the Open Geospatial Consortium Inc.

University Corporation for Atmospheric Research (UCAR)

iv. Submitters

All questions regarding this submission should be directed to the editor or the submitters:

Name	Company
Ben Domenico	UCAR

1. Scope

NetCDF (Network Common Data Form) is a data model for array-oriented scientific data. There is a freely distributed collection of access libraries implementing support for that data model in a machine-independent format. Together, the interfaces, libraries, and format support the creation, access, and sharing of scientific data.

Background information regarding the overall landscape of netCDF standards is presented in the CF-netCDF Primer, OGC 10-091r3, “CF-netCDF: Core and Extensions.” This standard is an extension to the core specification for the netCDF Classic data model in OGC 10-090r3, “NetCDF Core.”

This document specifies the netCDF Enhanced Data Model.

2. Conformance

This Standard defines the netCDF enhanced data model.

Standardization targets are netCDF implementations (currently: encodings).

This document establishes a single requirements class, *netcdf-enhanced*, of

<http://www.opengis.net/spec/netcdf/1.0/req/netcdf-enhanced>

with a single pertaining conformance class, *netcdf enhanced*, with URI

<http://www.opengis.net/spec/netcdf/1.0/conf/netcdf-enhanced>.

All requirements-classes and conformance-classes described in this document are owned by the standard identified as <http://www.opengis.net/spec/netcdf-enhanced/1.0> .

Annex A (normative) specifies conformance tests which shall be exercised by any encoded dataset claiming to implement an OGC netCDF encoding.

Conformance with this standard shall be checked using all the relevant tests specified in Annex A (normative) of this document. The framework, concepts, and methodology for testing, and the criteria to be achieved to claim conformance are specified in the OGC Compliance Testing Policies and Procedures and the OGC Compliance Testing web site¹.

¹ www.opengeospatial.org/cite

3. References

The following normative documents contain provisions that, through reference in this text, constitute provisions of this document. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. For undated references, the latest edition of the normative document referred to applies.

The netCDF classic data model is defined in OGC 10-091r3, “NetCDF Core.” The standard defined in this document is an extension to that core specification. Consequently all provisions stated in the NetCDF Core apply to the NetCDF Enhanced Data Model defined herein.

In addition to the normative reference noted above, portions of this document are taken from the Unidata NetCDF User’s Guide

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

and from the NASA Earth Science Data Systems (ESDS) Standards Process Group (SPG) document ESDS-RFC-022, NetCDF-4/HDF5 File Format,

<http://earthdata.nasa.gov/our-community/esdswg/standards-process-spg/rfc/esds-rfc-022>

4. Terms and Definitions

This document uses the terms defined in Sub-clause 5.3 of [OGC 06-121r8], which is based on the ISO/IEC Directives, Part 2, Rules for the structure and drafting of International Standards. In particular, the word “shall” (not “must”) is the verb form used to indicate a requirement to be strictly followed to conform to this standard.

For the purposes of this document, the following additional terms and definitions apply. These definitions are all of definition type “property.” The base URI to be used with the URI tokens is:

http://www.opengis.net/def/property/netcdf/1.0/{URI_token}

Note:

The netCDF data types are based on the IETF XDR standard (<http://tools.ietf.org/html/rfc4506>). Because the XDR standard requires at least 4 bytes for every data value, netCDF instead packs arrays of byte and short data types in order to improve storage efficiency.

4.1

byte*

A byte is an eight-bit integer. The byte type differs from the char type in that it is intended for eight-bit data, and the zero byte has no special significance, as it may for character data. URI token: byte.

4.2

char*

A char is an 8-bit byte that contains uninterpreted character data. Typically a char contains a 7-bit ASCII character, but the character encoding is application-specific. URI token: char.

4.3

compound type

Compound types group together named fields (also called "members"), that may be of different types, into a structure that may be accessed as a unit. It has a fixed total size. It consists of zero or more named members that do not overlap with other members.

- Each member has a name distinct from other members.
- Each member has its own datatype.
- Each member is referenced by an index number between zero and N-1, where N is the number of members in the compound datatype.
- Each member has a fixed byte offset, which is the first byte (smallest byte address) of that member in the compound datatype.
- In addition to other other user-defined data types or atomic datatypes, a member can be a small fixed-size array of any type with up to four fixed-size "dimensions" (not associated with named netCDF dimensions).

URI token: compoundType.

4.4

double*

A double is an IEEE double-precision floating point (64 bits.). The double type can hold values between about -1.7×10^{308} and 1.7×10^{308} , with external representation as 64-bit IEEE standard normalized double-precision, floating-point numbers. URI token: double.

4.5

enumerated (enum) type

Enumerated types are used to represent a small number of named values more concisely than strings. URI token: enumeratedType.

4.6

float*

A float is an IEEE single-precision floating point (32 bits). The float type can hold values between about -3.4×10^{38} and 3.4×10^{38} , with external representation as 32-bit IEEE normalized single-precision floating-point numbers. URI token: float.

4.7

group

A group provides a nested scope for names. URI token: group.

4.8

int*

An int is a 32-bit signed integer. An int can hold values between -2147483648 and 2147483647. URI token: int

4.9

int64

An int64 is an eight-byte signed integer. An int64 can hold values between -9223372036854775808 and 9223372036854775807. URI token: int64.

4.10

long*

The long type is deprecated. It is synonymous with int. URI token: long.

4.11

opaque type

Opaque is a collection of objects of a known size. Each object is the same size. Nothing is known to netCDF about the contents of these blobs of data, except their size in bytes, and the name of the type. URI token: opaqueType.

4.12

real*

The real type is synonymous with float. URI token: real.

4.13

short*

A short is a 16-bit signed integer. A short holds values between -32768 and 32767. URI token: short.

4.14

string

A string is a variable-length array of Unicode characters. A string is stored in a netCDF file as UTF-8 encoded Unicode. URI token: string.

4.15

unsigned int

An unsigned int is a 32-bit unsigned integer. An unsigned int holds values between 0 and 4294967296. URI token: unsignedInt

4.16

unsigned int64

An unsigned int64 is an unsigned 8-byte integer type. An unsigned int64 can hold values between 0 and 18446744073709551616. URI token: unsignedInte64.

4.17

unsigned short

An unsigned short is a 16-bit unsigned integer. An unsigned short holds values between 0 and 65536. URI token: unsignedShort.

4.18

user defined type

User defined types are: compound, enumerated, variable-length, and opaque. URI token: userDefinedType.

4.19

variable length type

The variable length type represents a one-dimensional array of variable length. URI token: variableLengthType.

*Note that, for purposes of completeness and readability, the definitions of all the primitive data types are included here. However, *char*, *byte*, *short*, *double*, *float*, and *real* are actually part of the netCDF classic data model and hence are part of the core netCDF specification.

5. Conventions

Some abbreviated terms frequently used in conjunction with enhanced netCDF:

API	Application Program Interface
BNF	Backus-Naur Form
CF	Climate and Forecast Conventions
ESDSWG	NASA Earth Standards Data Systems Working Groups
ES	Earth Sciences

GIS	Geographic Information System
HDF5	Hierarchical Data Format version 5
NcML	NetCDF Markup Language
NcML-GML	NetCDF Markup Language – Geography Markup Language
NetCDF	Network Common Data Form
NetCDF-4	NetCDF Release 4
ISO	International Organization for Standardization
OGC	Open Geospatial Consortium
SPG	NASA Standards Process Group
UML	Unified Modeling Language
XML	eXtended Markup Language
WFS	Web Feature Service
WCS	Web Coverage Service

6. Clauses not Containing Normative Material

There are no clauses with non normative material.

7. Clauses containing normative material

The netCDF enhanced data model adds Groups and User-Defined Types to the classic netCDF data model, but backward compatibility is preserved. The netCDF enhanced data model is also known as the netCDF-4 data model.

The UML representation of the netCDF enhanced data model below highlights in red the additions to the netCDF classic data model, showing netCDF enhanced data model is conceptually a set of simple extensions to the classic netCDF data model.

A file has one top-level unnamed group. Each group may contain zero or more named subgroups, user-defined types, variables, dimensions, and attributes. Variables may also have attributes. Variables may share dimensions, indicating a common grid. One or more dimensions may be of unlimited length.

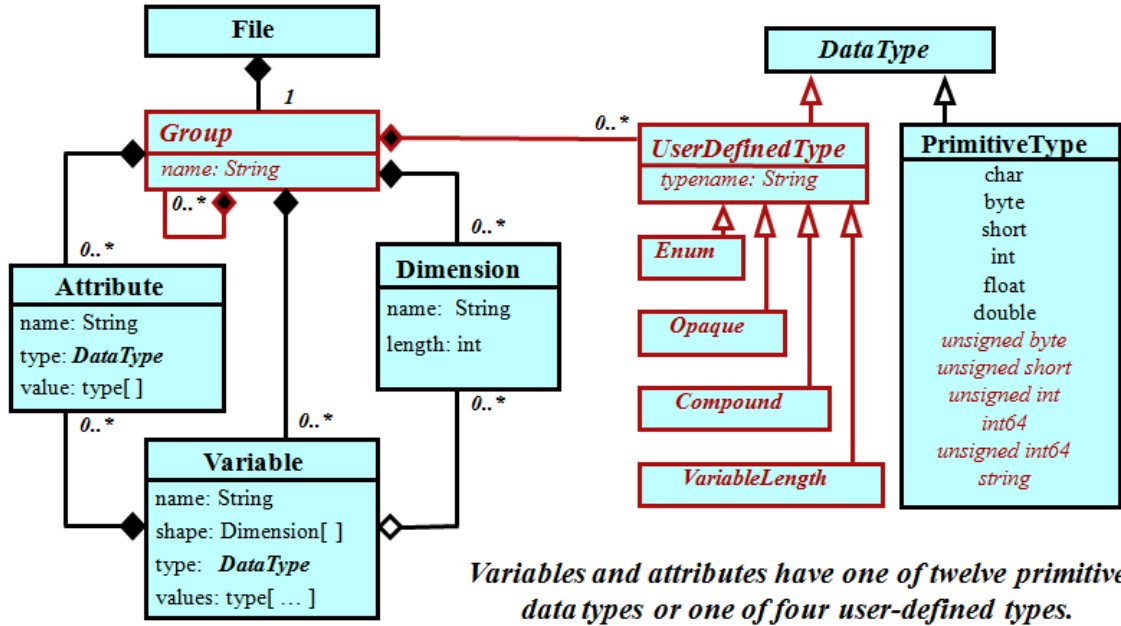


Figure 1: UML Diagram for netCDF Enhanced Data Model

7.1 Requirement Class netcdf-enhanced

There is only one requirements class for the Enhanced NetCDF Data Model. It is netcdf-enhanced.

<http://www.opengis.net/spec/netcdf-enhanced/1.0/req/req-class/netcdf-enhanced>

Req 1 <http://www.opengis.net/spec/netcdf-enhanced/1.0/req/netcdf-enhanced/data-model-elements>

NetCDF enhanced datasets shall represent information using groups, dimensions, variables, and attributes using the data types as shown in the UML diagram in Figure 1.

7.1.1 NetCDF classic data model

Req 2 <http://www.opengis.net/spec/netcdf-enhanced/1.0/req/netcdf-enhanced/classic-data-model>

The data shall conform to the netCDF classic abstract model as specified in the document OGC 10-090r3, “NetCDF Core” with the extensions specified in the requirements that

follow. Related conformance test cases are defined in Annex A of OGC 10-090r3, “NetCDF Core.”

7.1.2 Data types

Req 3 <http://www.opengis.net/spec/netcdf-enhanced/1.0/req/netcdf-enhanced/data-types>

Variables and attributes shall have one of 12 primitive data types or one of four user-defined data types

Req 4 <http://www.opengis.net/spec/netcdf-enhanced/1.0/req/netcdf-enhanced/primitive-data-types>

A primitive data type shall be one of char, byte, short, int, int64, float, double, unsigned byte, unsigned short, unsigned int, unsigned int64, string.

Note: This extends the corresponding netCDF core requirement which only allows the char, byte, short, int, float, and double primitive data types.

Req 5 <http://www.opengis.net/spec/netcdf-enhanced/1.0/req/netcdf-enhanced/user-defined-data-types>

A user-defined data type shall be one of enumerated, opaque, compound, variableLength.

7.1.3 Top-level, unnamed group

Req 6 <http://www.opengis.net/spec/netcdf-enhanced/1.0/req/netcdf-enhanced/top-level-group>

A netCDF dataset shall have one top level, unnamed group. In this context an unnamed group is equivalent to a group whose name is the null string.

Note: The top-level unnamed group provides backward compatibility with netCDF classic files, which have no groups. Thus, a netCDF classic file can be thought of as a netCDF enhanced model file with only the top-level unnamed group.

7.1.4 Groups

Each group may contain zero or more named subgroups, user-defined types, variables, dimensions, and attributes.

7.1.5 Variables

Variables may have attributes.

7.1.6 Dimensions

One or more dimensions may be of unlimited length.

7.1.6.1 Shared dimensions

Variables may share dimensions.

Req 7 <http://www.opengis.net/spec/netcdf-enhanced/1.0/req/netcdf-enhanced/shared-dimensions>

Shared dimensions shall indicate a common grid.

8. Media Types for any data encoding(s)

The MIME type `application/x-netcdf` is unregistered, since it doesn't appear in the list at

<http://www.iana.org/assignments/media-types/index.html>

It is, however, in common use; you can find it listed among the

"MIME Types by Content Type" at

<http://webdesign.about.com/od/multimedia/a/mime-types-by-content-type.htm>

and in the list of various MIME types at

<http://www.htaccess-guide.com/adding-mime-types/>

and many other places.

9. References

OGC 10-090r3, *netCDF Classic Core Interface Standard*

OGC 10-091r3, *CF-netCDF Core and Extensions Primer*

IEEE 754-2008, *Standard for Floating-Point Arithmetic*

ISO/IEC 10646, *Universal multiple-octet coded character set (UCS)*

Annex A: Conformance Class Abstract Test Suite (Normative)

A.1 Conformance class: netcdf-enhanced

There is only one conformance class for the Enhanced NetCDF Data Model. It is netcdf-enhanced.

<http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/conf-class/netcdf-enhanced>

A.1.1 Conformance Test 1: <http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/data-model-elements>

Open the dataset and verify that it represents information using groups, dimensions, variables, and attributes using the data types as shown in the UML diagram in Figure 1.

A.1.2 Conformance Test 2:

<http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/classic-data-model>

[Dependencies: netCDF classic abstract model as specified in the document OGC 10-090, “NetCDF Core.” Annex A of OGC 10-090, “NetCDF Core.”]

Open the dataset and verify that data conform to the netCDF classic abstract model as specified in the document OGC 10-090, “NetCDF Core” with the extensions specified in the following conformance tests. Related conformance test cases are defined in Annex A of OGC 10-090, “NetCDF Core.”

A.1.3 Conformance Test 3: <http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/data-types>

Open the dataset and verify that variables and attributes have one of 12 primitive data types or one of four user-defined data types

A.1.4 Conformance Test 4: <http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/primitive-data-types>

Open the dataset and verify that primitive data type are one of char, byte, short, int, int64, float, double, unsigned byte, unsigned short, unsigned int, unsigned int64, string.

Note: This extends the corresponding netCDF core conformance test which only allows the char, byte, short, int, float, and double primitive data types.

A.1.5 Conformance Test 5: <http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/user-defined-data-types>

Open the dataset and verify that each user-defined data types is one of enumerated, opaque, compound, variableLength.

A.1.6 Conformance Test 6: <http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/top-level-group>

Open the dataset and determine that it has one top level, unnamed group. In this context an unnamed group is equivalent to a group whose name is the null string.

A.1.7 Conformance Test 7: <http://www.opengis.net/spec/netcdf-enhanced/1.0/conf/netcdf-enhanced/shared-dimensions>

Open the dataset and verify that each shared dimension indicates a common grid.

Annex B: Revision history

Date	Release	Author	Paragraph modified	Description