

# OGC Hierarchical Data Format Version 5 (HDF5®) Core Standard

# Table of Contents

1. Scope .....	5
2. Conformance .....	6
3. References .....	7
4. Terms and Definitions .....	8
4.1. address .....	8
4.2. array .....	8
4.3. array type .....	8
4.4. array variable .....	8
4.5. atomic type .....	8
4.6. bit .....	8
4.7. byte .....	8
4.8. composite type .....	8
4.9. data .....	9
4.10. data element .....	9
4.11. data model .....	9
4.12. data processing system .....	9
4.13. data type .....	9
4.14. directed graph .....	9
4.15. directed path .....	9
4.16. entity .....	9
4.17. encoding .....	9
4.18. file .....	10
4.19. graph .....	10
4.20. HDF5 attribute .....	10
4.21. HDF5 committed datatype .....	10
4.22. HDF5 container .....	10
4.23. HDF5 dataset .....	10
4.24. HDF5 dataset region reference .....	10
4.25. HDF5 dataspace .....	10
4.26. HDF5 datatype .....	10
4.27. HDF5 entity .....	11
4.28. HDF5 group .....	11
4.29. HDF5 link .....	11
4.30. HDF5 object .....	11
4.31. HDF5 object graph .....	11
4.32. HDF5 object reference .....	11
4.33. HDF5 path name .....	11
4.34. instance .....	11

4.35. metadata	11
4.36. multigraph	11
4.37. process	12
4.38. program	12
4.39. programming language	12
4.40. rank	12
4.41. record	12
4.42. rooted directed graph	12
5. Conventions	13
5.1. Identifiers	13
5.2. UML Notation	13
5.3. Abbreviations	13
6. HDF5 Data Model Overview (Informative)	14
6.1. HDF5 Link	14
6.2. HDF5 Group	15
6.3. HDF5 Datatype	15
6.4. HDF5 Dataspace	16
6.5. HDF5 Dataset	16
6.6. HDF5 Attribute	16
7. HDF5 Data Model Requirements (Normative)	17
7.1. Requirements Class: Core	17
7.1.1. Requirement 1	18
7.1.2. Requirement 2	18
7.1.3. Requirement 3	18
7.1.4. Requirement 4	18
7.1.5. Requirement 5	19
7.1.6. Requirement 6	19
7.1.7. Requirement 7	19
7.1.8. Requirement 8	19
7.1.9. Requirement 9	19
7.1.10. Requirement 10	20
7.1.11. Requirement 11	20
7.1.12. Requirement 12	20
7.1.13. Requirement 13	20
7.1.14. Requirement 14	20
8. Media Types for any data encoding(s)	22
Annex A: Conformance Class Abstract Test Suite (Normative)	23
A.1. Conformance Class: Core	23
A.1.1. Requirement 1	23
A.1.2. Requirement 2	23
A.1.3. Requirement 3	23

A.1.4. Requirement 4 .....	23
A.1.5. Requirement 5 .....	23
A.1.6. Requirement 6 .....	24
A.1.7. Requirement 7 .....	24
A.1.8. Requirement 8 .....	24
A.1.9. Requirement 9 .....	24
A.1.10. Requirement 10 .....	24
A.1.11. Requirement 11 .....	25
A.1.12. Requirement 12 .....	25
A.1.13. Requirement 13 .....	25
A.1.14. Requirement 14 .....	25
Annex B: Mapping Between NetCDF Enhanced and HDF5 Data Models (Informative) .....	26
Annex C: Revision History .....	28
Annex D: Bibliography .....	29

## Open Geospatial Consortium

Submission Date: 2019-03-27

Approval Date: 2019-05-28

Publication Date: 2019-10-28

External identifier of this OGC® document: <http://www.opengis.net/doc/IS/HDF5/data-model/1.0>

Internal reference number of this OGC® document: 18-043r3

Version: 1.0

Category: OGC® Implementation Specification

Editors: Aleksandar Jelenak, Ted Habermann, Gerd Heber

### OGC Hierarchical Data Format Version 5 (HDF5®) Core Standard

#### Copyright notice

Copyright © 2019 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:	Encoding
Document stage:	Approved for public release
Document language:	English

## The HDF Group

The company listed above has granted the Open Geospatial Consortium (OGC) a nonexclusive, royalty-free, paid up, worldwide license to copy and distribute this document and to modify this document and distribute copies of the modified version under an Apache 2.0 license (see below).

## License Agreement

The standard is licensed under the Apache License, Version 2.0 (the "License"). You are directed to the License for specific details at: <http://www.apache.org/licenses/LICENSE-2.0>.

## **i. Abstract**

Hierarchical Data Format Version 5 (HDF5®) is a data model, a programming interface, and a storage model for keeping and managing data. It supports an unlimited variety of data types, and is designed to be flexible and efficient for large and complex data. HDF5 is extensible via customizing data types, allowing communities and their applications to evolve in the use of HDF5.

This document describes the HDF5 data model as an encoding standard particularly suitable to scientific and engineering geospatial applications that employ multidimensional numeric arrays to describe temporally and spatially varying phenomena. The data model is simple yet versatile, capable of supporting complex data relationships and dependencies through its grouping and linking mechanisms. It is also self-describing by accommodating user-defined metadata.

## **ii. Keywords**

The following are keywords to be used by search engines and document catalogues.

ogcdoc, OGC document, hdf, hdf5, hierarchical, data model, multidimensional array, big data, data cube

## **iii. Preface**

This document describes the HDF5 data model as an OGC Implementation Specification. It contains the required terms and definitions of the HDF5 data model and its conformance classes and tests.

HDF5 is already in use by the OGC community. The NetCDF Enhanced Data Model (OGC 11-038r2) can be fully encoded with the HDF5 data model because its storage format is HDF5. However, the HDF5 data model has several features exceeding those of the NetCDF Enhanced model (see [Annex B](#)). Furthermore, OGC members are starting to use HDF5 as a standard format for more complex types of data. In particular, HDF5 has been proposed as at least two standards for point cloud data in the OGC Point Cloud Domain Working Group [1]; has been demonstrated as a GDAL compliant format for images (KEA) [2]; and used as an International Hydrographic Organization's open standard for the Bathymetric Attributed Grid (BAG) format [3]. Given this situation, it is appropriate to consider the standardization of other elements of the HDF5 ecosystem which will all be based on this document.

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 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.

## **iv. Submitting organizations**

The following organizations submitted this Implementation Specification to the Open Geospatial Consortium (OGC):

- The HDF Group, Inc.

- U.S. National Aeronautics and Space Administration (NASA)

**v. Submitters**

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

<b>Name</b>	<b>Affiliation</b>
Aleksandar Jelenak	The HDF Group, Inc.
Ted Habermann	The HDF Group, Inc.



# Chapter 1. Scope

This document defines the HDF5 data model and the conformance requirements every implementation must comply with.

# Chapter 2. Conformance

This standard defines a single requirements class, [Core](#), with URI:

<http://www.opengis.net/spec/HDF5/data-model/1.0/req/core>

with a single accompanying conformance class, [Core](#) with URI:

<http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core>

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 site.

# Chapter 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.

[ISO2382] ISO/IEC: ISO 2382:2015 Information technology — Vocabulary (2015)

[ISO19118] ISO/TC211: ISO 19118:2011 Geographic information — Encoding (2011)

[ISO11179] ISO/IEC: ISO 11179-1:2004(E) Information technology — Metadata registries (MDR) — Part 1: Framework (2004)

# Chapter 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.

## 4.1. address

Value that identifies a location.

## 4.2. array

An instance of an array type where each element or appropriate group of elements may be referenced randomly and independently of the others [ISO2382]. A non-degenerate array has a rectilinear shape described by its rank (dimensionality) and extent. There are two degenerate cases: empty arrays (so-called null arrays), and arrays that consist of only a single element (scalars).

## 4.3. array type

Composite type whose components are of the same data type [ISO2382].

## 4.4. array variable

A variable capable of assuming values of an array type.

## 4.5. atomic type

Data type, each data object of which consists of a single non-decomposable data value. [ISO2382]

## 4.6. bit

Either of the digits 0 or 1 when used in the binary system [ISO2382].

## 4.7. byte

String that consists of a number of bits, treated as a unit, and usually representing a character or a part of a character [ISO2382]. The number of bits in a byte is fixed for a given data processing system. The number of bits in a byte is usually eight (8).

## 4.8. composite type

Data type that has a structure composed of the data structures of one or more data types and that

has its own set of permissible operations [\[ISO2382\]](#).

## 4.9. data

Reinterpretable representation of information in a formalized manner suitable for communication, interpretation, or processing [\[ISO2382\]](#).

## 4.10. data element

Unit of data that is considered in context to be indivisible [\[ISO2382\]](#).

## 4.11. data model

Graphical and/or lexical representation of data, specifying their properties, structure and inter-relationships [\[ISO11179\]](#).

## 4.12. data processing system

One or more computers, peripheral equipment, and software that perform data processing [\[ISO2382\]](#).

## 4.13. data type

Defined set of data objects of a specified data structure and a set of permissible operations, such that these data objects act as operands in the execution of any one of these operations [\[ISO2382\]](#).

## 4.14. directed graph

A graph where the edges are represented by ordered pairs of vertices and an implied edge direction (or orientation) from the first to the second vertex.

## 4.15. directed path

A finite sequence of directed graph edges (arcs, lines) where for all edges but the last the edge's second vertex is identical to the next edge's first vertex.

## 4.16. entity

Any concrete or abstract thing that exists, did exist, or might exist, including associations among these things [\[ISO2382\]](#).

## 4.17. encoding

Conversion of data into a series of codes [\[ISO11179\]](#).

## 4.18. file

Named set of records stored or processed as a unit [\[ISO2382\]](#).

## 4.19. graph

A non-empty, finite set of vertices (nodes) and a set of edges (arcs, lines) that connect pairs of vertices. An edge can also connect one vertex to itself.

## 4.20. HDF5 attribute

A named array variable ascribed to an HDF5 object, its owner or assignee.

## 4.21. HDF5 committed datatype

An immutable kind of HDF5 object that is used to store an HDF5 datatype definition, which can be referenced by multiple HDF5 datasets and HDF5 attributes.

## 4.22. HDF5 container

A mapping of an HDF5 object graph into storage such as one or more files in a file system, a buffer in memory, or an object store.

## 4.23. HDF5 dataset

An array variable linked to one or more HDF5 groups.

## 4.24. HDF5 dataset region reference

An HDF5 datatype whose values represent references to a subset of elements of an HDF5 dataset. Logically, such a reference can be thought of as a pair which consists of a reference to an HDF5 dataset and a boolean array variable of the same shape as the dataset where boolean *true/false* might indicate the inclusion/exclusion in/from the subset.

## 4.25. HDF5 dataspace

1. A rectilinear multidimensional lattice with finite extent (Simple).
2. A single point (Scalar).
3. The empty set (Null).

## 4.26. HDF5 datatype

A description of the data type of a data element including its storage layout as a pattern of bits.

## 4.27. HDF5 entity

A member of the HDF5 data model.

## 4.28. HDF5 group

An HDF5 object that is a collection of zero or more HDF5 links.

## 4.29. HDF5 link

A named, unidirectional association between an HDF5 group and an HDF5 object.

## 4.30. HDF5 object

An HDF5 group, dataset or committed datatype; an HDF5 entity that can be linked to zero or more HDF5 groups and assigned zero or more HDF5 attributes.

## 4.31. HDF5 object graph

A rooted, directed multigraph whose edges consist of HDF5 links, and whose vertices consist of all HDF5 objects which can be reached by traversing the edges starting from the HDF5 root group.

## 4.32. HDF5 object reference

An HDF5 datatype whose values represent references to HDF5 objects in an HDF5 container.

## 4.33. HDF5 path name

A finite-length sequence of HDF5 link names. HDF5 path names are used to locate and refer to HDF5 objects.

## 4.34. instance

Concrete module created from a generic module by generic instantiation [\[ISO2382\]](#).

## 4.35. metadata

Data about data or data elements, possibly including their data descriptions, and data about data ownership, access paths, access rights and data volatility [\[ISO2382\]](#).

## 4.36. multigraph

A graph with multiple edges connecting the same two vertices.

## 4.37. process

Predetermined course of events that occur during the execution of all or part of a program [\[ISO2382\]](#).

## 4.38. program

Syntactic unit that conforms to the rules of a particular programming language and that is composed of declarations and statements or instructions needed to solve a certain function, task, or problem [\[ISO2382\]](#).

## 4.39. programming language

Artificial language for expressing programs [\[ISO2382\]](#).

## 4.40. rank

A positive integer denoting the number of dimensions (dimensionality) of an array variable.

## 4.41. record

Set of data elements treated as a unit [\[ISO2382\]](#).

## 4.42. rooted directed graph

A graph with a single distinguished vertex, referred to as the root, and where all its other vertices can be reached by at least one directed path from the root.



# Chapter 5. Conventions

These sections provide details and examples for any conventions used in the document. Examples of conventions are symbols, abbreviations, use of XML schema, or special notes regarding how to read the document.

## 5.1. Identifiers

The normative provisions in this specification are denoted by the URI: <http://www.opengis.net/spec/HDF5/data-model/1.0>.

Any partial URI that may appear in this document's requirements and conformance tests shall be considered relative to this base.

## 5.2. UML Notation

All the diagrams included in this document use the Unified Modeling Language (UML) static structure diagrams, as described in Subclause 5.2 of OGC Web Services Common [OGC 06-121r8].

## 5.3. Abbreviations

The following abbreviations are being used throughout this document:

### **GDAL**

Geospatial Data Abstraction Library

### **HDF5**

Hierarchical Data Format Version 5

### **IEC**

International Electrotechnical Commission

### **ISO**

International Organization for Standardization

### **NetCDF**

Network Common Data Form

### **OGC**

Open Geospatial Consortium

### **UML**

Unified Modeling Language

# Chapter 6. HDF5 Data Model Overview (Informative)

This clause describes the HDF5 data model in more detail. Most of the presented material is taken from [4] and [5] which are recommended as the definitive sources of information.

The HDF5 data model consists of six entities: Group, Dataset, Link, Datatype, Dataspace, and Attribute, as illustrated in Figure 1.




Figure 1. Entities of the HDF5 data model.

Each of these entities are explained below in separate subclauses.

## 6.1. HDF5 Link

An HDF5 link connects an HDF5 group, its *source*, with another HDF5 object (group, dataset, or committed datatype) and is the means by which group membership is established. The object that a link connects to is referred to as the link's *target*.

An HDF5 link has a name associated with it. The name is a Unicode string that identifies the link within its source HDF5 group. A link's name is used to resolve the link, i.e., to locate the link's target. Because of this link names within one HDF5 group must be unique although there is no limit on how many links can have the same source HDF5 group and/or the same target. It is possible for HDF5 links to "dangle", i.e., their targets can be undefined or unavailable.

HDF5 links are the mean of traversing an HDF5 object graph, as depicted in Figure 2. Each vertex of the graph represents either an HDF5 group, dataset, or committed datatype. The identifiers associated with each HDF5 object in the graph are only to aid explanation in this text and are not part of the data model. However, an implementation would need to adopt some form of identification in order to distinguish different HDF5 objects.

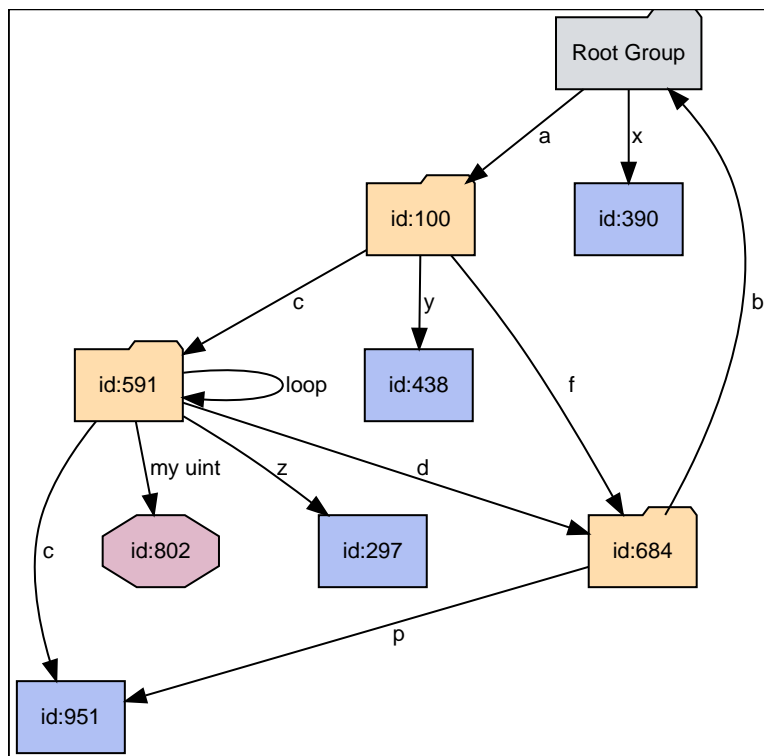


Figure 2. An example of HDF5 object graph as a rooted directed multigraph. Folder-shaped objects represent HDF5 groups; rectangles are HDF5 datasets; and the octagon is an HDF5 committed datatype.

One of the HDF5 groups in the graph is distinguished from the other groups because all the other objects can be reached by following one or more HDF5 links only starting from it. This HDF5 group has a special name: root group, and the data model allows only one such group to exist.

The names of the sequence of HDF5 links which lead to a specific HDF5 object form that object's HDF5 *path name*. One HDF5 object can have more than one path name. HDF5 path names starting from the root group are called *absolute*, while those starting from any other HDF5 group are called *relative*. For example, these are all HDF5 absolute path names to reach the HDF5 dataset **id:951**: (a, c, c), (a, c, d, p), and (a, f, p). Few examples of HDF5 relative path names to the same dataset are: (f, p) from the **id:100** group, and (d, p) or (c) from the **id:591** group.

Cyclical path names (the start and end group are the same) are possible. Two examples in Figure 2: the HDF5 absolute path (a, f, b) from the root group, and the relative path (loop) of the **id:591** group.

## 6.2. HDF5 Group

The HDF5 Group is a container of zero or more HDF5 links and has a functionality akin to a file system directory. With the exception of the root group, any other group of an HDF5 object graph is the link target of one HDF5 group.

An HDF5 object linked into a group is said to be a *member* of that group. The same HDF5 object can be a member of one or more groups. An example of this in Figure 2 is the **id:951** dataset which is a member of both **id:591** and **id:684** groups.

## 6.3. HDF5 Datatype

An HDF5 datatype describes an individual data element of an HDF5 dataset or attribute. It specifies

the set of possible values a data element can have and the set of operations that can be performed on it. The data model recognizes 10 different classes of HDF5 datatypes that fall into two major categories: atomic and composite, as depicted in [Figure 3](#). Expected set of data values for each datatype class is listed in [Table 3](#).




Figure 3. Classification of HDF5 datatypes.

An HDF5 dataset or attribute has an HDF5 datatype associated with it. The datatype may either be private to that dataset or attribute, or be linked into an HDF5 object graph and be available for reuse by other HDF5 datasets and attributes. A linked HDF5 datatype is referred to as *committed datatype*. An example of committed datatype in [Figure 2](#) is the `id:802` object with the HDF5 absolute path (`a, c, my uint`).

## 6.4. HDF5 Dataspace

An HDF5 dataspace describes the dimensionality (rank) and extent (the size of each dimension) of the data array associated with an HDF5 dataset or attribute object. The data model recognizes three different types of HDF5 dataspace which properties are listed in [Table 2](#).

## 6.5. HDF5 Dataset

The HDF5 Dataset is the main entity of the data model for storing data. An HDF5 dataset is comprised of a multidimensional array, an HDF5 datatype describing the array's data elements, and an HDF5 dataspace specifying the array's rank and extent.

## 6.6. HDF5 Attribute

HDF5 groups, datasets, and committed datatypes may be assigned zero or more HDF5 attributes. Each HDF5 attribute is private to the assignee HDF5 object.

An HDF5 attribute is a multidimensional array with a name. The array requires an HDF5 datatype and an HDF5 dataspace, the same as an HDF5 dataset.

An HDF5 attribute is functionally similar to an HDF5 dataset with the following differences:

- An attribute can only be accessed via the HDF5 object it was assigned to.
- Attribute names are significant only within the HDF5 object.
- Attributes cannot have attributes.

HDF5 attributes support the self-descriptiveness of the data encoded with the HDF5 data model by holding user-created metadata.

# Chapter 7. HDF5 Data Model Requirements (Normative)

The requirements defining compliance with the HDF5 data model are stated in the following subclauses. This is not an encoding but a set of concepts that could be implemented in numerous ways or technology.

## 7.1. Requirements Class: Core

The core requirements every implementation of the HDF5 data model must fulfill are listed in [Table 1](#). Their normative definitions are given in the subsequent subclauses.

Table 1. List of the requirements for the requirement class Core.

Requirements Class	
<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core</a>	
Target type	HDF5 container
Dependency	None
<b>Requirement 1</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/file-content">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/file-content</a> The content is based on the HDF5 data model.
<b>Requirement 2</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/root-group">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/root-group</a> Exactly one root group exists.
<b>Requirement 3</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/reachability">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/reachability</a> Reachability of HDF5 groups, datasets, and committed datatypes.
<b>Requirement 4</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatypes</a> HDF5 datasets and attributes have one of the supported HDF5 datatypes.
<b>Requirement 5</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/committed-datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/committed-datatypes</a> Every HDF5 committed datatype belong to one of the HDF5 Datatype classes.
<b>Requirement 6</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspaces">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspaces</a> HDF5 datasets and attributes have valid HDF5 dataspace.
<b>Requirement 7</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/attributes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/attributes</a> HDF5 attributes are attached only to HDF5 groups, datasets, and committed datatypes.
<b>Requirement 8</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/unique-attribute-name">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/unique-attribute-name</a> HDF5 attributes of one HDF5 object have unique names.
<b>Requirement 9</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/name-charset">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/name-charset</a> Character set of HDF5 link and HDF5 attribute names.

<b>Requirement 10</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/group-link-names">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/group-link-names</a> HDF5 link names in an HDF5 group.
<b>Requirement 11</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-source">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-source</a> Which HDF5 objects can be HDF5 link sources.
<b>Requirement 12</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-target">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-target</a> HDF5 link targets can be undefined.
<b>Requirement 13</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspace-dimsizes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspace-dimsizes</a> Simple HDF5 dataspace must report sizes of its dimensions.
<b>Requirement 14</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatype-classes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatype-classes</a> HDF5 datatype classes.

### 7.1.1. Requirement 1

<b>Requirement 1</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/file-content">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/file-content</a>  Information stored shall be encoded using the entities of the HDF5 data model as shown in <a href="#">Figure 1</a> .
----------------------	---

### 7.1.2. Requirement 2

<b>Requirement 2</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/root-group">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/root-group</a>  Exactly one root group shall exist.
----------------------	--

### 7.1.3. Requirement 3

<b>Requirement 3</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/reachability">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/reachability</a>  Every HDF5 group, dataset, or committed datatype shall be reachable via HDF5 links starting from the root group or by dereferencing a data element of the type HDF5 object reference belonging to an HDF5 dataset or attribute.
----------------------	--

### 7.1.4. Requirement 4

<b>Requirement 4</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatypes</a>  Every HDF5 dataset and attribute shall have an HDF5 datatype belonging to one of the classes shown in <a href="#">Figure 3</a> .
----------------------	---

### 7.1.5. Requirement 5

<b>Requirement 5</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/committed-datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/committed-datatypes</a>  Every HDF5 committed datatype shall belong to one of the classes shown in <a href="#">Figure 3</a> .
----------------------	---

### 7.1.6. Requirement 6

<b>Requirement 6</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspaces">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspaces</a>  Every HDF5 dataset and attribute shall have an HDF5 dataspace of type: Null, Scalar, or Simple.
----------------------	--

HDF5 Dataspace types are explained in [Table 2](#).

*Table 2. HDF5 Dataspace types.*

Type	Definition	Dimensionality (Rank)
Null	Empty set	Not applicable
Scalar	Singleton	Not applicable
Simple	Multidimensional array of data elements	Less than 33

### 7.1.7. Requirement 7

<b>Requirement 7</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/attributes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/attributes</a>  Every HDF5 attribute shall belong to only one of the HDF5 groups, datasets, or committed datatypes.
----------------------	--

### 7.1.8. Requirement 8

<b>Requirement 8</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/unique-attribute-name">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/unique-attribute-name</a>  The names of all HDF5 attributes belonging to the same HDF5 object shall be unique.
----------------------	--

### 7.1.9. Requirement 9

<b>Requirement 9</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/name-charset">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/name-charset</a>  The character set of HDF5 link or HDF5 attribute names shall be Unicode.
----------------------	---

### 7.1.10. Requirement 10

<b>Requirement 10</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/group-link-names">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/group-link-names</a>  The names of all HDF5 links within each HDF5 group shall be unique.
-----------------------	--

### 7.1.11. Requirement 11

<b>Requirement 11</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-source">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-source</a>  Only HDF5 groups shall be sources of HDF5 links.
-----------------------	---

### 7.1.12. Requirement 12

<b>Requirement 12</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-target">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-target</a>  The target of an HDF5 link might be undefined or unavailable.
-----------------------	--

### 7.1.13. Requirement 13

<b>Requirement 13</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspace-dimsizes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspace-dimsizes</a>  A Simple HDF5 dataspace shall report the extent for each of its dimensions.
-----------------------	--

### 7.1.14. Requirement 14

<b>Requirement 14</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatype-classes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatype-classes</a>  HDF5 datatypes shall provide implementation-conforming representations of the data classes described in <a href="#">Table 3</a> .
-----------------------	--

Table 3. HDF5 Datatype classes.

Atomic datatype classes	
Class	Description of value set
Integer	Range of integer values from $-9,223,372,036,854,775,808$ ( $-2^{63}$ ) to $18,446,744,073,709,551,615$ ( $2^{64}-1$ ).
Floating-point	Range of values as specified by the IEEE 754 double-precision binary floating-point.
String	Unicode strings.
Bitfield	Fixed-length sequences of bits.



Enumeration	A set of Unicode string-named integer constants.
Opaque	Unicode string-tagged set of fixed-length sequences of uninterpreted bytes.
Reference	Either the set of HDF5 object references or the set of HDF5 region references.
<b>Composite datatype classes</b>	
<b>Class</b>	<b>Description of value set</b>
Array	Arrays of fixed rank and extent with elements of the same HDF5 datatype class.
Variable-Length Sequence	Variable-length sequences of elements of the same HDF5 datatype class.
Compound	Tuples of values of the same heading, where the heading consists of a fixed-length sequence of (Unicode string field name, field HDF5 datatype) pairs.

# Chapter 8. Media Types for any data encoding(s)

The MIME type `application/x-hdf5` is unregistered and it doesn't appear in the list at <http://www.iana.org/assignments/media-types/index.html>. However, it is in common use as evidenced by appearing in [6], [7], and [8].

# Annex A: Conformance Class Abstract Test Suite (Normative)

## A.1. Conformance Class: Core

### A.1.1. Requirement 1

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/file-content">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/file-content</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/file-content">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/file-content</a>
<b>Test purpose:</b>	Verify that the stored information is composed of the HDF5 data model entities.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.2. Requirement 2

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/root-group">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/root-group</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/root-group">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/root-group</a>
<b>Test purpose:</b>	Verify that there is only one root group.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.3. Requirement 3

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/reachability">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/reachability</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/reachability">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/reachability</a>
<b>Test purpose:</b>	Verify that all HDF5 groups, datasets, and committed datatypes can be accessed either by using HDF5 links starting from the root group, or via HDF5 object references stored in HDF5 datasets and attributes.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.4. Requirement 4

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/datatypes</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatypes</a>
<b>Test purpose:</b>	Verify that the datatype of every HDF5 dataset and HDF5 attribute belongs to one of the HDF5 datatype classes in <a href="#">Table 3</a> .
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.5. Requirement 5

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/committed-datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/committed-datatypes</a>
-----------------	---

<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/committed-datatypes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/committed-datatypes</a>
<b>Test purpose:</b>	Verify that every HDF5 committed datatype belongs to one of the HDF5 datatype classes in <a href="#">Table 3</a> .
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.6. Requirement 6

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/dataspaces">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/dataspaces</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspaces">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspaces</a>
<b>Test purpose:</b>	Verify that the dataspace of every HDF5 dataset and HDF5 attribute belongs to one of the HDF5 dataspace types in <a href="#">Table 2</a> .
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.7. Requirement 7

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/attributes">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/attributes</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/attributes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/attributes</a>
<b>Test purpose:</b>	Verify that every HDF5 attribute belongs to only one of the HDF5 groups, datasets, or committed datatypes.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.8. Requirement 8

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/unique-attribute-name">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/unique-attribute-name</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/unique-attribute-name">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/unique-attribute-name</a>
<b>Test purpose:</b>	Verify that the names of all HDF5 attributes ascribed to the same HDF5 object are unique.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.9. Requirement 9

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/name-charset">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/name-charset</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/name-charset">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/name-charset</a>
<b>Test purpose:</b>	Verify that the character set of every HDF5 link and HDF5 attribute name is Unicode.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.10. Requirement 10

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/group-link-names">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/group-link-names</a>
-----------------	---

<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/group-link-names">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/group-link-names</a>
<b>Test purpose:</b>	Verify that for every HDF5 group the names of all HDF5 links sourced from it are unique.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.11. Requirement 11

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/link-source">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/link-source</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-source">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-source</a>
<b>Test purpose:</b>	Verify that only HDF5 groups are the sources of HDF5 links.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.12. Requirement 12

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/link-target">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/link-target</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-target">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/link-target</a>
<b>Test purpose:</b>	Verify that the target of an HDF5 link can be undefined or not available.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.13. Requirement 13

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/dataspace-dimsize">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/dataspace-dimsize</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspace-dimsize">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/dataspace-dimsize</a>
<b>Test purpose:</b>	Verify that the HDF5 Simple Dataspace can report the extent (size) of every dimension.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

### A.1.14. Requirement 14

<b>Test id:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/datatype-classes">http://www.opengis.net/spec/HDF5/data-model/1.0/conf/core/datatype-classes</a>
<b>Requirement:</b>	<a href="http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatype-classes">http://www.opengis.net/spec/HDF5/data-model/1.0/req/core/datatype-classes</a>
<b>Test purpose:</b>	Verify that the HDF5 Datatype classes in <a href="#">Table 3</a> are supported.
<b>Test method:</b>	Inspect the implementation to confirm the above requirement.

# Annex B: Mapping Between NetCDF Enhanced and HDF5 Data Models (Informative)

The netCDF Enhanced data model is already an OGC Implementation Standard (OGC 11-038r2). Its storage format, known as netCDF-4, is HDF5 which means the netCDF data model can be represented using the HDF5 data model. The following table provides mapping between the main capabilities of these two data models.

Table 4. Mapping between features of the netCDF Enhanced and HDF5 data models.

netCDF Enhanced	HDF5	Note
Dataset	Container	—
Group	Group	—
Top-level, unnamed group	Root group	The netCDF Enhanced model refers to the HDF5 root group as <i>unnamed</i> due to the backward compatibility requirement with the older netCDF Classic data model (OGC 10-090r3) which lacks grouping capability.
Dimension	Dataset	There is no equivalent in the HDF5 data model. A netCDF dimension is encoded as an HDF5 dataset with several HDF5 attributes that have specific names and values. The attributes' values declare that the dataset represents a netCDF dimension, specify dimension's length, and associate it as a dimension of other HDF5 datasets using HDF5 object references.
Variable	Dataset	—
Attribute	Attribute	The netCDF Enhanced attribute can only store a scalar or one-dimensional array.
Global attribute	Root group attribute	—
Variable attribute	Dataset attribute	—
Variable or attribute data type	Datatype	—
Variable shape	Dataspace	—
Primitive data type	Atomic datatype	netCDF Enhanced primitive datatypes do not include the following HDF5 atomic datatypes: Opaque, Bitfield, and Reference.
User-defined data type	Composite datatype	The netCDF Enhanced user-defined data types do not include the composite HDF5 Array datatype but include the atomic HDF5 Opaque datatype.
Group or variable name	Link name	Only one HDF5 link per netCDF group or variable is allowed, i.e. netCDF groups and variables cannot have multiple HDF5 path names.

<b>netCDF Enhanced</b>	<b>HDF5</b>	<b>Note</b>
Group or variable hierarchy	Path name	The netCDF Enhanced data model implements strict tree-structured hierarchy, i.e. circular paths through the netCDF dataset object graph — as in the HDF5 data model — are not possible.

The main HDF5 data model's features that provide additional capabilities for describing data are:

- Multiple links (path names) to the same HDF5 object.
- Dangling HDF5 links.
- HDF5 attributes able to store multidimensional arrays of the same maximal rank as HDF5 datasets.
- HDF5 object reference datatype for locating HDF5 objects independent of their path names.
- HDF5 region reference datatype for recording subsets of elements of other HDF5 datasets.
- HDF5 object and region references can be stored in HDF5 datasets and attributes.

# Annex C: Revision History

Date	Release	Editor	Primary clauses modified	Description
2018-05-03	18-043	Aleksandar Jelenak	All	Initial version.
2018-11-04	18-043r1	Aleksandar Jelenak	Title page	<ul style="list-style-type: none"> <li>• New license agreement.</li> <li>• Add a registered trademark sign to <i>HDF5</i> in the title.</li> </ul>
2019-01-04	18-043r2	Aleksandar Jelenak	Front matter & Annex B	Changes in response to received public comments.
2019-01-18	18-043r3	Aleksandar Jelenak	Front matter & Annex B	Clarify the relationship between the NetCDF Enhanced and HDF5 data models.



# Annex D: Bibliography

1. M. Ingram, “Advanced Point Cloud Format Standards.” Presentation in the OGC Point Cloud DWG Session, Washington, D.C., Mar-2016 [Online]. Available: [https://portal.opengeospatial.org/files/?artifact\\_id=67558](https://portal.opengeospatial.org/files/?artifact_id=67558)
2. P. Bunting and S. Gillingham, “The KEA image file format,” *Computers & Geosciences*, vol. 57, no. Supplement C, pp. 54–58, 2013 [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0098300413001015>
3. “Bathymetric Surface Product Specification,” no. S-102. International Hydrographic Organization, 2012 [Online]. Available: [https://www.iho.int/iho\\_pubs/standard/S-102/S-102\\_Ed1.0.0\\_Apr12.pdf](https://www.iho.int/iho_pubs/standard/S-102/S-102_Ed1.0.0_Apr12.pdf)
4. The HDF Group, “HDF5 User’s Guide, Abstract Data Model.” 2017 [Online]. Available: [https://support.hdfgroup.org/HDF5/doc/UG/HDF5\\_Users\\_Guide-Responsive%20HTML5/index.html#t=HDF5\\_Users\\_Guide%2FDataModelAndFileStructure%2FThe\\_HDF5\\_Data\\_Model\\_and\\_File\\_Structure.htm%3Frhtocid%3Dtoc1.1%23TOC\\_1\\_2\\_The\\_Abstract\\_Databc-2](https://support.hdfgroup.org/HDF5/doc/UG/HDF5_Users_Guide-Responsive%20HTML5/index.html#t=HDF5_Users_Guide%2FDataModelAndFileStructure%2FThe_HDF5_Data_Model_and_File_Structure.htm%3Frhtocid%3Dtoc1.1%23TOC_1_2_The_Abstract_Databc-2)
5. M. Folk and E. Pourmal, “HDF5 Data Model, File Format and Library–HDF5 1.6,” NASA ESDIS Standards Office, ESDS-RFC-007v1, 2007 [Online]. Available: <https://earthdata.nasa.gov/standards/hdf5>
6. Scilab GitHub repository, “scilab.xml.” [Online]. Available: <https://github.com/opencollab/scilab/blob/bd38cc4979d86a2240aeef4a277373da97afe516/scilab/desktop/scilab.xml>
7. Wolfram Language & System Documentation Center, “HDF5.” [Online]. Available: <https://reference.wolfram.com/language/ref/format/HDF5.html>
8. Research Data Netherlands, “Essentials 4 Data Support, Data Formats.” [Online]. Available: <https://datasupport.researchdata.nl/en/start-the-course/iii-the-research-phase/data-formats/mime-types/>