



# **Hierarchical Data Format query language (HDFql)**

## **Reference Manual**

**Version 2.1.0**

July 2019

**Copyright (C) 2016-2019**

This document is part of the Hierarchical Data Format query language (HDFql). For more information about HDFql, please visit the website <http://www.hdfql.com>.

**Disclaimer**

Every effort has been made to ensure that this document is as accurate as possible. The information contained in this document is provided without any express, statutory or implied warranties. The founders of HDFql shall have neither liability nor responsibility to any person or entity with respect to any loss or damage arising from the information in this document or the usage of HDFql.

# TABLE OF CONTENTS

<b>1. INTRODUCTION .....</b>	<b>1</b>
<b>2. INSTALLATION .....</b>	<b>3</b>
2.1 WINDOWS .....	4
2.2 LINUX.....	4
2.3 MACOS .....	5
<b>3. USAGE .....</b>	<b>6</b>
3.1 C .....	6
3.2 C++.....	10
3.3 JAVA .....	14
3.4 PYTHON .....	17
3.5 C# .....	19
3.6 FORTRAN .....	22
3.7 R .....	27
3.8 COMMAND-LINE INTERFACE.....	29
<b>4. CURSOR .....</b>	<b>32</b>
4.1 DESCRIPTION .....	32
4.2 SUBCURSOR.....	36
4.3 EXAMPLES .....	38
<b>5. APPLICATION PROGRAMMING INTERFACE .....</b>	<b>44</b>
5.1 CONSTANTS .....	44
5.2 FUNCTIONS.....	52
5.2.1 HDFQL_EXECUTE .....	58
5.2.2 HDFQL_EXECUTE_GET_STATUS .....	59
5.2.3 HDFQL_ERROR_GET_LINE .....	60
5.2.4 HDFQL_ERROR_GET_POSITION.....	61

5.2.5	HDFQL_ERROR_GET_MESSAGE .....	62
5.2.6	HDFQL_CURSOR_INITIALIZE .....	63
5.2.7	HDFQL_CURSOR_USE .....	64
5.2.8	HDFQL_CURSOR_USE_DEFAULT .....	65
5.2.9	HDFQL_CURSOR_CLEAR .....	66
5.2.10	HDFQL_CURSOR_CLONE .....	67
5.2.11	HDFQL_CURSOR_GET_DATA_TYPE .....	68
5.2.12	HDFQL_CURSOR_GET_COUNT .....	69
5.2.13	HDFQL_SUBCURSOR_GET_COUNT .....	70
5.2.14	HDFQL_CURSOR_GET_POSITION .....	72
5.2.15	HDFQL_SUBCURSOR_GET_POSITION .....	73
5.2.16	HDFQL_CURSOR_FIRST .....	74
5.2.17	HDFQL_SUBCURSOR_FIRST .....	75
5.2.18	HDFQL_CURSOR_LAST .....	76
5.2.19	HDFQL_SUBCURSOR_LAST .....	77
5.2.20	HDFQL_CURSOR_NEXT .....	78
5.2.21	HDFQL_SUBCURSOR_NEXT .....	79
5.2.22	HDFQL_CURSOR_PREVIOUS .....	81
5.2.23	HDFQL_SUBCURSOR_PREVIOUS .....	82
5.2.24	HDFQL_CURSOR_ABSOLUTE .....	83
5.2.25	HDFQL_SUBCURSOR_ABSOLUTE .....	84
5.2.26	HDFQL_CURSOR_RELATIVE .....	86
5.2.27	HDFQL_SUBCURSOR_RELATIVE .....	87
5.2.28	HDFQL_CURSOR_GET_SIZE .....	89
5.2.29	HDFQL_SUBCURSOR_GET_SIZE .....	90
5.2.30	HDFQL_CURSOR_GET .....	91
5.2.31	HDFQL_SUBCURSOR_GET .....	92
5.2.32	HDFQL_CURSOR_GET_TINYINT .....	93

5.2.33	HDFQL_SUBCURSOR_GET_TINYINT .....	94
5.2.34	HDFQL_CURSOR_GET_UNSIGNED_TINYINT.....	96
5.2.35	HDFQL_SUBCURSOR_GET_UNSIGNED_TINYINT .....	97
5.2.36	HDFQL_CURSOR_GET_SMALLINT .....	98
5.2.37	HDFQL_SUBCURSOR_GET_SMALLINT .....	99
5.2.38	HDFQL_CURSOR_GET_UNSIGNED_SMALLINT .....	100
5.2.39	HDFQL_SUBCURSOR_GET_UNSIGNED_SMALLINT.....	101
5.2.40	HDFQL_CURSOR_GET_INT .....	103
5.2.41	HDFQL_SUBCURSOR_GET_INT.....	104
5.2.42	HDFQL_CURSOR_GET_UNSIGNED_INT .....	105
5.2.43	HDFQL_SUBCURSOR_GET_UNSIGNED_INT .....	106
5.2.44	HDFQL_CURSOR_GET_BIGINT.....	107
5.2.45	HDFQL_SUBCURSOR_GET_BIGINT .....	108
5.2.46	HDFQL_CURSOR_GET_UNSIGNED_BIGINT .....	110
5.2.47	HDFQL_SUBCURSOR_GET_UNSIGNED_BIGINT .....	111
5.2.48	HDFQL_CURSOR_GET_FLOAT .....	112
5.2.49	HDFQL_SUBCURSOR_GET_FLOAT .....	113
5.2.50	HDFQL_CURSOR_GET_DOUBLE .....	115
5.2.51	HDFQL_SUBCURSOR_GET_DOUBLE .....	116
5.2.52	HDFQL_CURSOR_GET_CHAR.....	117
5.2.53	HDFQL_SUBCURSOR_GET_CHAR .....	118
5.2.54	HDFQL_VARIABLE_REGISTER .....	119
5.2.55	HDFQL_VARIABLE_TRANSIENT_REGISTER .....	122
5.2.56	HDFQL_VARIABLE_UNREGISTER .....	123
5.2.57	HDFQL_VARIABLE_UNREGISTER_ALL.....	124
5.2.58	HDFQL_VARIABLE_GET_NUMBER.....	125
5.2.59	HDFQL_VARIABLE_GET_DATA_TYPE.....	126
5.2.60	HDFQL_VARIABLE_GET_COUNT.....	127

5.2.61	HDFQL_VARIABLE_GET_SIZE .....	129
5.2.62	HDFQL_VARIABLE_GET_DIMENSION_COUNT .....	130
5.2.63	HDFQL_VARIABLE_GET_DIMENSION .....	131
5.2.64	HDFQL_MPI_GET_SIZE .....	132
5.2.65	HDFQL_MPI_GET_RANK.....	133

## 6. LANGUAGE ..... 135

6.1	DATA TYPES .....	138
6.1.1	TINYINT.....	140
6.1.2	UNSIGNED TINYINT .....	141
6.1.3	SMALLINT .....	142
6.1.4	UNSIGNED SMALLINT .....	142
6.1.5	INT .....	143
6.1.6	UNSIGNED INT .....	143
6.1.7	BIGINT .....	144
6.1.8	UNSIGNED BIGINT .....	145
6.1.9	FLOAT .....	145
6.1.10	DOUBLE .....	146
6.1.11	CHAR.....	147
6.1.12	VARTINYINT .....	147
6.1.13	UNSIGNED VARTINYINT.....	148
6.1.14	VARSMALLINT.....	149
6.1.15	UNSIGNED VARSMALLINT .....	149
6.1.16	VARINT .....	150
6.1.17	UNSIGNED VARINT .....	151
6.1.18	VARBIGINT.....	151
6.1.19	UNSIGNED VARBIGINT .....	152
6.1.20	VARFLOAT.....	153
6.1.21	VARDOUBLE.....	153

6.1.22	VARCHAR .....	154
6.1.23	OPAQUE .....	154
6.2	POST-PROCESSING .....	155
6.2.1	ORDER .....	156
6.2.2	TOP .....	158
6.2.3	BOTTOM .....	160
6.2.4	FROM TO .....	162
6.2.5	STEP .....	164
6.3	REDIRECTING .....	166
6.3.1	FROM .....	166
6.3.2	INTO .....	172
6.4	DATA DEFINITION LANGUAGE (DDL) .....	177
6.4.1	CREATE DIRECTORY .....	177
6.4.2	CREATE FILE .....	179
6.4.3	CREATE GROUP .....	180
6.4.4	CREATE DATASET .....	183
6.4.5	CREATE ATTRIBUTE .....	190
6.4.6	CREATE [SOFT   HARD] LINK .....	193
6.4.7	CREATE EXTERNAL LINK .....	195
6.4.8	ALTER DIMENSION .....	198
6.4.9	RENAME DIRECTORY .....	199
6.4.10	RENAME FILE .....	200
6.4.11	RENAME [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK] .....	202
6.4.12	COPY FILE .....	203
6.4.13	COPY [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK] .....	205
6.4.14	DROP DIRECTORY .....	207
6.4.15	DROP FILE .....	208
6.4.16	DROP [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK] .....	209

6.5	DATA MANIPULATION LANGUAGE (DML) .....	210
6.5.1	INSERT .....	211
6.6	DATA QUERY LANGUAGE (DQL) .....	222
6.6.1	SELECT .....	222
6.7	DATA INTROSPECTION LANGUAGE (DIL) .....	233
6.7.1	SHOW FILE VALIDITY .....	233
6.7.2	SHOW USE DIRECTORY .....	235
6.7.3	SHOW USE FILE .....	236
6.7.4	SHOW ALL USE FILE .....	238
6.7.5	SHOW USE GROUP .....	239
6.7.6	SHOW [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK] .....	241
6.7.7	SHOW TYPE .....	248
6.7.8	SHOW DATA TYPE .....	250
6.7.9	SHOW MASK .....	252
6.7.10	SHOW ENDIANNES .....	253
6.7.11	SHOW CHARSET .....	255
6.7.12	SHOW STORAGE TYPE .....	257
6.7.13	SHOW STORAGE ALLOCATION .....	259
6.7.14	SHOW STORAGE DIMENSION .....	261
6.7.15	SHOW DIMENSION .....	262
6.7.16	SHOW ORDER .....	265
6.7.17	SHOW TAG .....	267
6.7.18	SHOW FILL TYPE .....	269
6.7.19	SHOW FILL VALUE .....	270
6.7.20	SHOW FILE SIZE .....	272
6.7.21	SHOW [DATASET   ATTRIBUTE] SIZE .....	274
6.7.22	SHOW HDFQL VERSION .....	275
6.7.23	SHOW HDF5 VERSION .....	276



6.7.24	SHOW PCRE VERSION .....	277
6.7.25	SHOW ZLIB VERSION .....	278
6.7.26	SHOW MPI VERSION .....	279
6.7.27	SHOW DIRECTORY .....	280
6.7.28	SHOW FILE .....	282
6.7.29	SHOW EXECUTE STATUS .....	283
6.7.30	SHOW LIBRARY BOUNDS .....	284
6.7.31	SHOW CACHE .....	286
6.7.32	SHOW ATOMIC .....	288
6.7.33	SHOW EXTERNAL LINK PREFIX .....	290
6.7.34	SHOW FLUSH .....	291
6.7.35	SHOW DEBUG .....	292
6.8	MISCELLANEOUS .....	293
6.8.1	USE DIRECTORY .....	294
6.8.2	USE FILE .....	295
6.8.3	USE GROUP .....	297
6.8.4	FLUSH .....	300
6.8.5	CLOSE FILE .....	300
6.8.6	CLOSE ALL FILE .....	302
6.8.7	CLOSE GROUP .....	303
6.8.8	SET LIBRARY BOUNDS .....	304
6.8.9	SET CACHE .....	306
6.8.10	SET ATOMIC .....	309
6.8.11	SET EXTERNAL LINK PREFIX .....	310
6.8.12	SET FLUSH .....	311
6.8.13	SET DEBUG .....	312
<b>GLOSSARY .....</b>		<b>314</b>
Application Programming Interface (API) .....		314

Attribute .....	314
Cursor .....	314
Dataset .....	314
Data type .....	315
Endianness.....	315
Group .....	315
Hierarchical Data Format (HDF) .....	315
Hyperslab.....	315
Message Passing Interface (MPI).....	316
Operation .....	316
Parallel HDF5 (PHDF5) .....	316
Post-processing .....	316
Result set .....	316
Result subset .....	317
Subcursor.....	317

# LIST OF TABLES

Table 5.1 – HDFqI constants in C .....	48
Table 5.2 – HDFqI constants in C and their corresponding definitions in C++ .....	49
Table 5.3 – HDFqI constants in C and their corresponding definitions in Java .....	50
Table 5.4 – HDFqI constants in C and their corresponding definitions in Python .....	50
Table 5.5 – HDFqI constants in C and their corresponding definitions in C# .....	50
Table 5.6 – HDFqI constants in C and their corresponding definitions in Fortran .....	51
Table 5.7 – HDFqI constants in C and their corresponding definitions in R .....	51
Table 5.8 – HDFqI functions in C .....	54
Table 5.9 – HDFqI functions in C and their corresponding definitions in C++ .....	55
Table 5.10 – HDFqI functions in C and their corresponding definitions in Java .....	55
Table 5.11 – HDFqI functions in C and their corresponding definitions in Python .....	56
Table 5.12 – HDFqI functions in C and their corresponding definitions in C# .....	56
Table 5.13 – HDFqI functions in C and their corresponding definitions in Fortran .....	57
Table 5.14 – HDFqI functions in C and their corresponding definitions in R .....	57
Table 6.1 – HDFqI operations text formatting conventions .....	135
Table 6.2 – HDFqI operations .....	138
Table 6.3 – HDFqI data types and their corresponding definitions in HDF .....	140
Table 6.4 – HDFqI post-processing options .....	155
Table 6.5 – HDFqI redirecting options .....	166

# LIST OF FIGURES

Figure 3.1 – Illustration of the command-line interface “HDFqI CLI” .....	31
Figure 4.1 – Linearization of a two dimensional dataset into a (one dimensional) cursor .....	36
Figure 4.2 – Cursor populated with data from dataset “my_dataset0” .....	38
Figure 4.3 – Cursor populated with data from dataset “my_dataset1” .....	39
Figure 4.4 – Cursor populated with data from dataset “my_dataset2” .....	40
Figure 4.5 – Cursor and its subcursor populated with data from dataset “my_dataset3” .....	41
Figure 4.6 – Cursor and its subcursors populated with data from dataset “my_dataset4” .....	42
Figure 4.7 – Cursor and its subcursors populated with data from dataset “my_dataset5” .....	43

---

# 1. INTRODUCTION

---

HDFql stands for “Hierarchical Data Format query language” and is the first tool that enables users to manage HDF5<sup>1</sup> files through a high-level language. This language was designed to be simple to use and similar to SQL thus dramatically reducing the learning effort. HDFql can be seen as an alternative to the C API (which contains more than 400 low-level functions that are far from easy to use!) and to existing wrappers for C++, Java, Python, C#, Fortran and R for manipulating HDF5 files. In addition, and whenever possible, it automatically employs parallelism to speed-up operations hiding its inherent complexity from the user.

As an example, imagine that one needs to create an HDF5 file named “myFile.h5” and, inside it, a group named “myGroup” containing a (ZLIB) compressed dataset named “myDataset” of data type float with an initial value of 12.4. In HDFql, this can easily be implemented as follows:

```
create file myFile.h5
use file myFile.h5
create dataset myGroup/myDataset as float enable zlib values(12.4)
```

In contrast, using the C API on the same example is quite cumbersome:

```
hid_t file;
hid_t group;
hid_t dataspace;
hid_t property;
hid_t dataset;
hsize_t dimension;
float value;

file = H5Fcreate("myFile.h5", H5F_ACC_EXCL, H5P_DEFAULT, H5P_DEFAULT);
group = H5Gcreate(file, "myGroup", H5P_DEFAULT, H5P_DEFAULT, H5P_DEFAULT);
dimension = 1;
dataspace = H5Screate_simple(1, &dimension, NULL);
```

---

<sup>1</sup> Hierarchical Data Format is the name of a set of file formats and libraries designed to store and organize large amounts of numerical data. It is supported by The HDF Group, whose mission is to ensure continued development of HDF technologies and the continued accessibility of data currently stored in HDF. Please refer to the website <http://www.hdfgroup.org> for additional information.

```
property = H5Pcreate(H5P_DATASET_CREATE);  
H5Pset_chunk(property, 1, &dimension);  
H5Pset_deflate(property, 9);  
dataset = H5Dcreate(group, "myDataset", H5T_NATIVE_FLOAT, dataspace, H5P_DEFAULT, property,  
H5P_DEFAULT);  
value = 12.4;  
H5Dwrite(dataset, H5T_NATIVE_FLOAT, H5S_ALL, H5S_ALL, H5P_DEFAULT, &value);
```

---

## 2. INSTALLATION

---

The official website of the Hierarchical Data Format query language (HDFqI) is <http://www.hdfqI.com>. Here, the most recent documentation and examples that illustrate how to solve disparate use-cases using HDFqI can be found. In addition, in the download area (<http://www.hdfqI.com/#download>) all versions of HDFqI ever publicly released are available. These versions are packaged as ZIP files, with each one meant for a particular platform (i.e. Windows, Linux or macOS), architecture (i.e. 32 bit or 64 bit), compiler (Microsoft Visual Studio or Gnu Compiler Collection (GCC)) and – optionally – MPI library (i.e. MPICH or Open MPI). When decompressed, such ZIP files typically have the following organization in terms of directories and files contained within:

```
HDFqI-x.y.z
|
+ example (directory that contains C, C++, Java, Python, C#, Fortran and R examples)
|
+ include (directory that contains HDFqI C and C++ header files)
|
+ lib (directory that contains HDFqI C release/debug static and shared libraries)
|
+ bin (directory that contains HDFqI command-line interface and a proper launcher)
|
+ wrapper (directory that contains HDFqI wrappers)
|   |
|   + cpp (directory that contains HDFqI C++ wrapper)
|   |
|   + java (directory that contains HDFqI Java wrapper)
|   |
|   + python (directory that contains HDFqI Python wrapper)
|   |
|   + csharp (directory that contains HDFqI C# wrapper)
|   |
|   + fortran (directory that contains HDFqI Fortran wrapper)
|   |
|   + R (directory that contains HDFqI R wrapper)
|
+ doc (directory that contains HDFqI reference manual)
|
- LICENSE.txt (file that contains information about HDFqI license)
|
```

- *RELEASE.txt* (file that contains information about HDFqL releases)
- |
- *README.txt* (file that contains succinct information about HDFqL)

The following sections provide concise instructions on how to install HDFqL in the different platforms that it currently supports – namely Windows, Linux and macOS.

## 2.1 WINDOWS

- Download the appropriate ZIP file according to the HDFqL version, architecture and compiler of interest from <http://www.hdfql.com/#download>. For instance, if the HDFqL version of interest is 1.0.0 and it is to be used in a machine running Windows 32 bit and, eventually, be linked against C or C++ code using the Microsoft Visual Studio 2010 compiler then the file to download is “HDFqL-1.0.0\_Windows32\_VS-2010.zip”.
- Unzip the downloaded file using Windows Explorer in-build capabilities or a free tool such as 7-Zip (<http://www.7-zip.org>).

## 2.2 LINUX

- Download the appropriate ZIP file according to the HDFqL version, architecture, compiler and (optional) MPI library of interest from <http://www.hdfql.com/#download>. For instance, if the HDFqL version of interest is 1.4.0 and it is to be used in a machine running Linux 64 bit and, eventually, be linked against C, C++, or Fortran code using the GCC 4.9.x compiler with no need to work with HDF5 files in parallel (using an MPI library) then the file to download is “HDFqL-1.4.0\_Linux64\_GCC-4.9.zip”.
- Unzip the downloaded file using the Archive Manager or the KArchive (if in GNOME or KDE respectively), or by opening a terminal and executing “*unzip <downloaded\_zip\_file>*”. If the unzip utility is not installed, it can be done by executing from a terminal:

- In a Red Hat-based distribution:

```
sudo yum install unzip
```



- In a Debian-based distribution:

```
sudo apt-get install unzip
```

## 2.3 MACOS

- Download the appropriate ZIP file according to the HDFqI version, architecture, compiler and (optional) MPI library of interest from <http://www.hdfqI.com/#download>. For instance, if the HDFqI version of interest is 2.1.0 and it is to be used in a machine running macOS 64 bit and, eventually, be linked against C, C++, or Fortran code using the GCC 4.9.x compiler with the need of working with HDF5 files in parallel using MPICH 3.2.x MPI library then the file to download is “HDFqI-2.1.0\_Darwin64\_GCC-4.9\_MPICH-3.2.zip”.
- Unzip the downloaded file using the Archive Utility or by opening a terminal and executing “*unzip <downloaded\_zip\_file>*”. If the unzip utility is not installed, it can be done by executing from a terminal:

```
sudo port install unzip
```

---

## 3. USAGE

---

After following the instructions provided in the chapter [INSTALLATION](#), HDFq1 is ready for usage. It can be used programmatically in C, C++ and Fortran through static and shared libraries; in Java, Python, C# and R through wrappers; and finally, through a command-line interface named “HDFq1CLI”. Moreover, in Linux and macOS, programs written in these programming languages may manipulate HDF5 files both in serial and in parallel<sup>1</sup>, as distributions of HDFq1 built with the serial HDF5 library and the parallel HDF5 (PHDF5) library are available for these platforms. The subsequent sections provide guidance on usage and basic troubleshooting information to solve issues that may arise.

### 3.1 C

HDFq1 can be used in the C programming language through static and shared libraries. These libraries are stored in the directory “lib”. The following short program illustrates how HDFq1 can be used in such language.

```
// include HDFq1 C header file (make sure it can be found by the C compiler)
#include "HDFq1.h"

int main(int argc, char *argv[])
{
    // display HDFq1 version in use
    printf("HDFq1 version: %s\n", HDFQL_VERSION);

    // create an HDF5 file named "my_file.h5"
    hdfql_execute("CREATE FILE my_file.h5");

    // use (i.e. open) HDF5 file "my_file.h5"
    hdfql_execute("USE FILE my_file.h5");

    // create an HDF5 dataset named "my_dataset" of data type int
```

---

<sup>1</sup> Through MPICH (or, alternatively, one of its ABI compatible derivative libraries such as Intel MPI, Cray MPT, MVAPICH2, Parastation MPI) or Open MPI. Both MPICH and Open MPI are freely available, high performance and widely portable implementations of the Message Passing Interface (MPI), a standard for message-passing for distributed memory applications used in parallel computing. Please refer to the website <https://www.mpich.org> and <https://www.open-mpi.org> for additional information.

```
hdfq1_execute("CREATE DATASET my_dataset AS INT VALUES(10)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor with it
hdfq1_execute("SELECT FROM my_dataset");

// move cursor to the first position within the result set
hdfq1_cursor_first(NULL);

// display content of cursor
printf("Dataset value: %d\n", *hdfq1_cursor_get_int(NULL));

return EXIT_SUCCESS;
}
```

Assuming that the program is stored in a file named “example.c”, it must first be compiled before it can be launched from a terminal. To compile the program against the HDFq1 C static library:

- In Windows<sup>2</sup> using Microsoft Visual Studio, by executing from a terminal:

```
cl.exe example.c /I<hdfq1_include_directory> <hdfq1_lib_directory>\HDFq1.lib /link /LTCG
/NODEFAULTLIB:libcmtd.lib
```

- In Linux and macOS using Gnu Compiler Collection (GCC), by executing from a terminal:

- With an HDFq1 non MPI-based distribution:

```
gcc example.c -fopenmp -I<hdfq1_include_directory> <hdfq1_lib_directory>/libHDFq1.a
-lm -ldl
```

- With an HDFq1 MPI-based distribution:

```
gcc example.c -fopenmp -I<hdfq1_include_directory> <hdfq1_lib_directory>/libHDFq1.a
-L<mpi_lib_directory> -lmpi -lm -ldl
```

---

<sup>2</sup> When compiling a program against the HDFq1 C static library in Windows, the functions “hdfq1\_initialize” and “hdfq1\_finalize” must be explicitly called by the program when starting and finishing respectively (otherwise an error may occur such as a segmentation fault). Of note, these functions do not need to be called when compiling the program against the HDFq1 C shared library as this is automatically done by the library itself.

To compile the same program against the HDFqI C shared library:

- In Windows using Microsoft Visual Studio, by executing from a terminal:

```
cl.exe example.c /I<hdfqI_include_directory> <hdfqI_lib_directory>\HDFqI_dll.lib
```

- In Linux and macOS using GCC, by executing from a terminal:

- With an HDFqI non MPI-based distribution:

```
gcc example.c -I<hdfqI_include_directory> -L<hdfqI_lib_directory> -lHDFqI -lm -ldl
```

- With an HDFqI MPI-based distribution:

```
gcc example.c -I<hdfqI_include_directory> -L<hdfqI_lib_directory> -  
L<mpi_lib_directory> -lHDFqI -lmpi -lm -ldl
```

Of note, debug versions of the HDFqI C static and shared libraries are also available. These are stored in the directory “debug” found under the directory “lib”. To compile C programs using debug libraries, the instructions described in the above bullet points should be followed with two modifications: (1) the directory storing the libraries should be updated (“<hdfqI\_lib\_directory>\debug” in Microsoft Visual Studio; “<hdfqI\_lib\_directory>/debug” in GCC); (2) the suffix “D” should be added to the name of the libraries (“HDFqID.lib” and “HDFqI\_dllD.lib” in Microsoft Visual Studio; “libHDFqID.a” and “-lHDFqID” in GCC).

In case the program does not compile, most likely a C compiler is not installed. If a C compiler is missing, the solution is:

- In Windows, download and install a free version of Microsoft Visual Studio from the website <https://www.visualstudio.com/downloads>.
- In Linux, install the GCC C compiler by executing from a terminal:
  - In a Red Hat-based distribution:

```
sudo yum install gcc
```

- In a Debian-based distribution:

```
sudo apt-get install gcc
```

- In macOS, install the GCC C compiler by executing from a terminal (if xcode-select does not support the parameter “--install” (due to being outdated), download and install the Command-Line Tools package from the website <http://developer.apple.com/downloads> which includes GCC instead):

```
xcode-select --install
```

In case the compiled program does not launch, most likely the HDFqI C shared library and/or the MPI shared library was not found (these are needed to launch the program). The solution is:

- In Windows, add the directory where the file “HDFqI\_dll.dll” is located to the environment variable “PATH” by executing from a terminal:

```
set PATH=<hdfqI_lib_directory>;%PATH%
```

- In Linux, add the directories where the files “libHDFqI.so” and (optionally) “libmpi.so” are located to the environment variable “LD\_LIBRARY\_PATH” by executing from a terminal:

- With an HDFqI non MPI-based distribution:

```
export LD_LIBRARY_PATH=<hdfqI_lib_directory>:$LD_LIBRARY_PATH
```

- With an HDFqI MPI-based distribution:

```
export LD_LIBRARY_PATH=<hdfqI_lib_directory>:<mpi_lib_directory>:$LD_LIBRARY_PATH
```

- In macOS, add the directories where the files “libHDFqL.dylib” and (optionally) “libmpi.dylib” are located to the environment variable “DYLD\_LIBRARY\_PATH”<sup>3</sup> by executing from a terminal:
  - With an HDFqL non MPI-based distribution:

```
export DYLD_LIBRARY_PATH=<hdfqL_lib_directory>:$DYLD_LIBRARY_PATH
```

- With an HDFqL MPI-based distribution:

```
export  
DYLD_LIBRARY_PATH=<hdfqL_lib_directory>:<mpi_lib_directory>:$DYLD_LIBRARY_PATH
```

## 3.2 C++

HDFqL can be used in the C++ programming language through static and shared libraries. These libraries are stored in the directory “cpp” found under the directory “wrapper”. The following short program illustrates how HDFqL can be used in such language.

```
// include HDFqL C++ header file (make sure it can be found by the C++ compiler)  
#include <iostream>  
#include "HDFqL.hpp"  
  
int main(int argc, char *argv[])  
{  
    // display HDFqL version in use  
    std::cout << "HDFqL version: " << HDFqL::Version << std::endl;  
  
    // create an HDF5 file named "my_file.h5"  
    HDFqL::execute("CREATE FILE my_file.h5");  
  
    // use (i.e. open) HDF5 file "my_file.h5"  
    HDFqL::execute("USE FILE my_file.h5");  
}
```

---

<sup>3</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information).

```
// create an HDF5 dataset named "my_dataset" of data type int
HDFqL::execute("CREATE DATASET my_dataset AS INT VALUES(10)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor with it
HDFqL::execute("SELECT FROM my_dataset");

// move cursor to the first position within the result set
HDFqL::cursorFirst();

// display content of cursor
std::cout << "Dataset value: " << *HDFqL::cursorGetInt() << std::endl;

return EXIT_SUCCESS;
}
```

Assuming that the program is stored in a file named “example.cpp”, it must first be compiled before it can be launched from a terminal. To compile the program against the HDFqL C++ static library:

- In Windows<sup>4</sup> using Microsoft Visual Studio, by executing from a terminal:

```
cl.exe example.cpp /EHsc /I<hdfqL_include_directory>
<hdfqL_cpp_wrapper_directory>\HDFqL.lib /link /LTCG /NODEFAULTLIB:libcmtd.lib
```

- In Linux and macOS using Gnu Compiler Collection (GCC), by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
g++ example.cpp -fopenmp -I<hdfqL_include_directory>
<hdfqL_cpp_wrapper_directory>/libHDFqL.a -ldl
```

- With an HDFqL MPI-based distribution:

```
g++ example.cpp -fopenmp -I<hdfqL_include_directory>
```

---

<sup>4</sup> When compiling a program against the HDFqL C++ static library in Windows, the functions “HDFqL::initialize” and “HDFqL::finalize” must be explicitly called by the program when starting and finishing respectively (otherwise an error may occur such as a segmentation fault). Of note, these functions do not need to be called when compiling the program against the HDFqL C++ shared library as this is automatically done by the library itself.

```
<hdfql_cpp_wrapper_directory>/libHDFqL.a -L<mpi_lib_directory> -lmpi -ldl
```

To compile the same program against the HDFqL C++ shared library:

- In Windows using Microsoft Visual Studio, by executing from a terminal:

```
cl.exe example.cpp /EHsc /I<hdfql_include_directory>  
<hdfql_cpp_wrapper_directory>\HDFqL_dll.lib
```

- In Linux and macOS using GCC, by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
g++ example.cpp -I<hdfql_include_directory> -L<hdfql_cpp_wrapper_directory> -lHDFqL  
-ldl
```

- With an HDFqL MPI-based distribution:

```
g++ example.cpp -I<hdfql_include_directory> -L<hdfql_cpp_wrapper_directory> -  
L<mpi_lib_directory> -lHDFqL -lmpi -ldl
```

In case the program does not compile, most likely a C++ compiler is not installed. If a C++ compiler is missing, the solution is:

- In Windows, download and install a free version of Microsoft Visual Studio from the website <https://www.visualstudio.com/downloads>.
- In Linux, install the GCC C++ compiler by executing from a terminal:

- In a Red Hat-based distribution:

```
sudo yum install gcc-c++
```

- In a Debian-based distribution:



```
sudo apt-get install g++
```

- In macOS, install the GCC C++ compiler by executing from a terminal (if xcode-select does not support the parameter “--install” (due to being outdated), download and install the Command-Line Tools package from the website <http://developer.apple.com/downloads> which includes GCC instead):

```
xcode-select --install
```

In case the compiled program does not launch, most likely the HDFqI C++ shared library and/or the MPI shared library was not found (these are needed to launch the program). The solution is:

- In Windows, add the directory where the file “HDFqI\_dll.dll” is located to the environment variable “PATH” by executing from a terminal:

```
set PATH=<hdfqI_cpp_wrapper_directory>;%PATH%
```

- In Linux, add the directories where the files “libHDFqI.so” and (optionally) “libmpi.so” are located to the environment variable “LD\_LIBRARY\_PATH” by executing from a terminal:

- With an HDFqI non MPI-based distribution:

```
export LD_LIBRARY_PATH=<hdfqI_cpp_wrapper_directory>:$LD_LIBRARY_PATH
```

- With an HDFqI MPI-based distribution:

```
export  
LD_LIBRARY_PATH=<hdfqI_cpp_wrapper_directory>:<mpi_lib_directory>:$LD_LIBRARY_PATH
```

- In macOS, add the directories where the files “libHDFqI.dylib” and (optionally) “libmpi.dylib” are located to the environment variable “DYLD\_LIBRARY\_PATH”<sup>5</sup> by executing from a terminal:
  - With an HDFqI non MPI-based distribution:

```
export DYLD_LIBRARY_PATH=<hdfqI_cpp_wrapper_directory>:$DYLD_LIBRARY_PATH
```

- With an HDFqI MPI-based distribution:

```
export
DYLD_LIBRARY_PATH=<hdfqI_cpp_wrapper_directory>:<mpi_lib_directory>:$DYLD_LIBRARY_PATH
```

## 3.3 JAVA

HDFqI can be used in the Java programming language through a wrapper named “HDFqI.java”. This wrapper is stored in the directory “java” found under the directory “wrapper”. The following short program illustrates how HDFqI can be used in such language.

```
// import HDFqI package (make sure it can be found by the Java compiler/JVM)
import as.hdfqI.*;

public class Example
{
    public static void main(String args[])
    {
        // display HDFqI version in use
        System.out.println("HDFqI version: " + HDFqI.VERSION);

        // create an HDF5 file named "my_file.h5"
        HDFqI.execute("CREATE FILE my_file.h5");

        // use (i.e. open) HDF5 file "my_file.h5"
```

---

<sup>5</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information).

```
HDFql.execute("USE FILE my_file.h5");

// create an HDF5 dataset named "my_dataset" of data type int
HDFql.execute("CREATE DATASET my_dataset AS INT VALUES(10)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor with it
HDFql.execute("SELECT FROM my_dataset");

// move cursor to the first position within the result set
HDFql.cursorFirst();

// display content of cursor
System.out.println("Dataset value: " + HDFql.cursorGetInt());
}
}
```

Assuming that the program is stored in a file named “Example.java”, it must first be compiled before it can be launched from a terminal. The program can be compiled as follows:

```
javac -classpath <hdfql_java_wrapper_directory> Example.java
```

In case the program does not compile, most likely the Java Development Kit (JDK) is not installed. If the JDK is missing, the solution is to download and install it from the website <http://www.oracle.com/technetwork/java/javase/downloads>.

The compiled program may be launched as follows:

```
java Example
```

In case the compiled program does not launch, most likely the HDFql Java wrapper and/or the MPI shared library was not found (these are needed to launch the program). The solution is:

- In Windows, add the directories where the files “HDFql.java” (i.e. the wrapper) and “HDFql.dll” are located to the environment variables “CLASSPATH” and “PATH” by executing from a terminal:

```
set CLASSPATH=<hdfql_java_wrapper_directory>;.%CLASSPATH%
set PATH=<hdfql_java_wrapper_directory>\as\hdfql;%PATH%
```

- In Linux, add the directories where the files “HDFqL.java”, “libHDFqL.so” and (optionally) “libmpi.so” are located to the environment variables “CLASSPATH” and “LD\_LIBRARY\_PATH” by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
export CLASSPATH=<hdfql_java_wrapper_directory>:.$CLASSPATH
export LD_LIBRARY_PATH=<hdfql_java_wrapper_directory>/as/hdfql:$LD_LIBRARY_PATH
```

- With an HDFqL MPI-based distribution:

```
export CLASSPATH=<hdfql_java_wrapper_directory>:.$CLASSPATH
export
LD_LIBRARY_PATH=<hdfql_java_wrapper_directory>/as/hdfql:<mpi_lib_directory>:$LD_LIBRARY_PATH
```

- In macOS, add the directories where the files “HDFqL.java”, “libHDFqL.dylib” and (optionally) “libmpi.dylib” are located to the environment variables “CLASSPATH” and “DYLD\_LIBRARY\_PATH”<sup>6</sup> by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
export CLASSPATH=<hdfql_java_wrapper_directory>:.$CLASSPATH
export DYLD_LIBRARY_PATH=<hdfql_java_wrapper_directory>/as/hdfql:$DYLD_LIBRARY_PATH
```

- With an HDFqL MPI-based distribution:

```
export CLASSPATH=<hdfql_java_wrapper_directory>:.$CLASSPATH
export
DYLD_LIBRARY_PATH=<hdfql_java_wrapper_directory>/as/hdfql:<mpi_lib_directory>:$DYLD_LIBRARY_PATH
```

---

<sup>6</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information). Alternatively, the Java library path property “java.library.path” should be set with the path where the HDFqL shared library “libHDFqL.dylib” is located when launching the program (e.g. `java -Djava.library.path=<hdfql_java_wrapper_directory>/as/hdfql my_program`).

## 3.4 PYTHON

HDFql can be used in the Python programming language through a wrapper named “HDFql.py”. This wrapper is stored in the directory “python” found under the directory “wrapper”. The following short script illustrates how HDFql can be used in such language.

```
# import HDFql module (make sure it can be found by the Python interpreter)
import HDFql

# display HDFql version in use
print("HDFql version: %s" % HDFql.VERSION)

# create an HDF5 file named "my_file.h5"
HDFql.execute("CREATE FILE my_file.h5")

# use (i.e. open) HDF5 file "my_file.h5"
HDFql.execute("USE FILE my_file.h5")

# create an HDF5 dataset named "my_dataset" of data type int
HDFql.execute("CREATE DATASET my_dataset AS INT VALUES(10)")

# select (i.e. read) data from dataset "my_dataset" and populate cursor with it
HDFql.execute("SELECT FROM my_dataset")

# move cursor to the first position within the result set
HDFql.cursor_first()

# display content of cursor
print("Dataset value: %d" % HDFql.cursor_get_int())
```

Assuming that the script is stored in a file named “example.py” it can be launched by executing the following from a terminal:

```
python example.py
```

In case the script does not launch, most likely (1) the Python interpreter is not installed or (2) the HDFql Python wrapper and/or the MPI shared library was not found (these are needed to launch the script). To fix the former issue, download and install the Python interpreter from the website <http://www.python.org/download>. To fix the latter issue:

- In Windows, add the directory where the file “HDFqI.py” (i.e. the wrapper) is located to the environment variable “PYTHONPATH” by executing from a terminal:

```
set PYTHONPATH=<hdfqI_python_wrapper_directory>;%PYTHONPATH%
```

- In Linux, add the directories where the files “HDFqI.py” and (optionally) “libmpi.so” are located to the environment variables “PYTHONPATH” and “LD\_LIBRARY\_PATH” by executing from a terminal:

- With an HDFqI non MPI-based distribution:

```
export PYTHONPATH=<hdfqI_python_wrapper_directory>:$PYTHONPATH
```

- With an HDFqI MPI-based distribution:

```
export PYTHONPATH=<hdfqI_python_wrapper_directory>:$PYTHONPATH
export LD_LIBRARY_PATH=<mpi_lib_directory>:$LD_LIBRARY_PATH
```

- In macOS, add the directories where the files “HDFqI.py” and (optionally) “libmpi.dylib” are located to the environment variables “PYTHONPATH” and “DYLD\_LIBRARY\_PATH”<sup>7</sup> by executing from a terminal:

- With an HDFqI non MPI-based distribution:

```
export PYTHONPATH=<hdfqI_python_wrapper_directory>:$PYTHONPATH
```

- With an HDFqI MPI-based distribution:

```
export PYTHONPATH=<hdfqI_python_wrapper_directory>:$PYTHONPATH
export DYLD_LIBRARY_PATH=<mpi_lib_directory>:$DYLD_LIBRARY_PATH
```

---

<sup>7</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information).

Besides these steps, a scientific computing package named NumPy for Python must be installed when working with user-defined variables (please refer to the function `hdfq1_variable_register` for additional information). This package can be found at <http://www.scipy.org/scipylib/download.html> along with instructions on how to install and use it.

## 3.5 C#

HDFq1 can be used in the C# programming language through a wrapper named “HDFq1.cs”. This wrapper is stored in the directory “csharp” found under the directory “wrapper”. The following short program illustrates how HDFq1 can be used in such language.

```
// use HDFq1 namespace (make sure it can be found by the C# compiler)
using AS.HDFq1;

public class Example
{
    public static void Main(string []args)
    {
        // display HDFq1 version in use
        System.Console.WriteLine("HDFq1 version: {0}", HDFq1.Version);

        // create an HDF5 file named "my_file.h5"
        HDFq1.Execute("CREATE FILE my_file.h5");

        // use (i.e. open) HDF5 file "my_file.h5"
        HDFq1.Execute("USE FILE my_file.h5");

        // create an HDF5 dataset named "my_dataset" of data type int
        HDFq1.Execute("CREATE DATASET my_dataset AS INT VALUES(10)");

        // select (i.e. read) data from dataset "my_dataset" and populate cursor with it
        HDFq1.Execute("SELECT FROM my_dataset");

        // move cursor to the first position within the result set
        HDFq1.CursorFirst();

        // display content of cursor
        System.Console.WriteLine("Dataset value: {0}", HDFq1.CursorGetInt());
    }
}
```

Assuming that the program is stored in a file named “Example.cs”, it must first be compiled before it can be launched from a terminal. In Windows, the program can be compiled as follows:

- Using Microsoft .NET Framework, by executing from a terminal:

```
csc.exe Example.cs <hdfql_csharp_wrapper_directory>\*.cs
```

- Using Mono, by executing from a terminal:

```
mcs.bat Example.cs <hdfql_csharp_wrapper_directory>\*.cs
```

In Linux and macOS, the program can be compiled using Mono by executing from a terminal (of note, Microsoft .NET Framework does not support these platforms):

```
mcs Example.cs <hdfql_csharp_wrapper_directory>/*.cs
```

In case the program does not compile, most likely a C# compiler is not installed. If a C# compiler is missing, the solution is:

- In Windows, download and install either Microsoft .NET Framework or Mono from the websites <https://www.microsoft.com/net/download/framework> or <http://www.mono-project.com/download>, respectively.
- In Linux and macOS, download and install Mono from the website <http://www.mono-project.com/download>.

Depending on the platform, the compiled program may be launched as follows:

- In Windows, by executing from a terminal:

```
Example.exe
```

- In Linux and macOS, by executing from a terminal:

```
mono Example.exe
```



In case the compiled program does not launch, most likely the HDFq C# wrapper and/or the MPI shared library was not found (these are needed to launch the program). The solution is:

- In Windows, add the directory where the file “HDFq.cs” (i.e. the wrapper) is located to the environment variable “PATH” by executing from a terminal:

```
set PATH=<hdfq_csharp_wrapper_directory>;%PATH%
```

- In Linux, add the directories where the files “HDFq.cs” and (optionally) “libmpi.so” are located to the environment variable “LD\_LIBRARY\_PATH” by executing from a terminal:

- With an HDFq non MPI-based distribution:

```
export LD_LIBRARY_PATH=<hdfq_csharp_wrapper_directory>:$LD_LIBRARY_PATH
```

- With an HDFq MPI-based distribution:

```
export
LD_LIBRARY_PATH=<hdfq_csharp_wrapper_directory>:<mpi_lib_directory>:$LD_LIBRARY_PATH
```

- In macOS, add the directories where the files “HDFq.cs” and (optionally) “libmpi.dylib” are located to the environment variable “DYLD\_LIBRARY\_PATH”<sup>8</sup> by executing from a terminal:

- With an HDFq non MPI-based distribution:

```
export DYLD_LIBRARY_PATH=<hdfq_csharp_wrapper_directory>:$DYLD_LIBRARY_PATH
```

- With an HDFq MPI-based distribution:

---

<sup>8</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information).

```
export
DYLD_LIBRARY_PATH=<hdfq1_csharp_wrapper_directory>:<mpi_lib_directory>:$DYLD_LIBRARY
_PATH
```

## 3.6 FORTRAN

HDFq1 can be used in the Fortran programming language through static and shared libraries. These libraries are stored in the directory “fortran” found under the directory “wrapper”. The following short program illustrates how HDFq1 can be used in such language.

```
PROGRAM Example
! use HDFq1 module (make sure it can be found by the Fortran compiler)
USE HDFq1

! declare variable
INTEGER :: state

! display HDFq1 version in use
WRITE(*, *) "HDFq1 version: ", HDFQL_VERSION

! create an HDF5 file named "my_file.h5"
state = hdfq1_execute("CREATE FILE my_file.h5")

! use (i.e. open) HDF5 file "my_file.h5"
state = hdfq1_execute("USE FILE my_file.h5")

! create an HDF5 dataset named "my_dataset" of data type int
state = hdfq1_execute("CREATE DATASET my_dataset AS INT VALUES(10)")

! select (i.e. read) data from dataset "my_dataset" and populate cursor with it
state = hdfq1_execute("SELECT FROM my_dataset")

! move cursor to the first position within the result set
state = hdfq1_cursor_first()

! display content of cursor
WRITE(*, *) "Dataset value: ", hdfq1_cursor_get_int()
END PROGRAM
```

Assuming that the program is stored in a file named “example.f90”, it must first be compiled before it can be launched from a terminal. To compile the program against the HDFqL Fortran static library:

- In Windows<sup>9</sup> using Intel Fortran Compiler (IFORT), by executing from a terminal:

```
ifort.exe example.f90 /module:<hdfql_fortran_wrapper_directory>\static
<hdfql_fortran_wrapper_directory>\HDFqL.lib /link /LTCG /NODEFAULTLIB:libcmt.lib
```

- In Linux using IFORT, by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
ifort example.f90 -fopenmp -module <hdfql_fortran_wrapper_directory>
<hdfql_fortran_wrapper_directory>/libHDFqL.a
```

- With an HDFqL MPI-based distribution:

```
ifort example.f90 -fopenmp -module <hdfql_fortran_wrapper_directory>
<hdfql_fortran_wrapper_directory>/libHDFqL.a -L<mpi_lib_directory> -lmpi
```

- In Linux and macOS using Gnu Compiler Collection (GCC)<sup>10</sup>, by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
gfortran example.f90 -fopenmp -I<hdfql_fortran_wrapper_directory>
<hdfql_fortran_wrapper_directory>/libHDFqL.a -ldl
```

- With an HDFqL MPI-based distribution:

---

<sup>9</sup> When compiling a program against the HDFqL Fortran static library in Windows, the subroutines “hdfql\_initialize” and “hdfql\_finalize” must be explicitly called by the program when starting and finishing respectively (otherwise an error may occur such as a segmentation fault). Of note, these functions do not need to be called when compiling the program against the HDFqL Fortran shared library as this is automatically done by the library itself.

<sup>10</sup> An incorrect warning is raised by the GCC Fortran compiler when using the HDFqL module (“Warning: Only array FINAL procedures declared for derived type 'hdfql\_cursor' defined at (1), suggest also scalar one”). This warning does not interfere with the final compilation result, though, and it has been solved in the GCC Fortran compiler version 7.0.0 (please refer to [https://gcc.gnu.org/bugzilla/show\\_bug.cgi?id=58175](https://gcc.gnu.org/bugzilla/show_bug.cgi?id=58175) for additional information).

```
gfortran example.f90 -fopenmp -I<hdfql_fortran_wrapper_directory>
<hdfql_fortran_wrapper_directory>/libHDFqL.a -L<mpi_lib_directory> -lmpi -ldl
```

To compile the same program against the HDFqL Fortran shared library:

- In Windows using IFORT, by executing from a terminal:

```
ifort.exe example.f90 /module:<hdfql_fortran_wrapper_directory>
<hdfql_fortran_wrapper_directory>\HDFqL_dll.lib
```

- In Linux using IFORT, by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
ifort example.f90 -module <hdfql_fortran_wrapper_directory> -
L<hdfql_fortran_wrapper_directory> -lHDFqL
```

- With an HDFqL MPI-based distribution:

```
ifort example.f90 -module <hdfql_fortran_wrapper_directory> -
L<hdfql_fortran_wrapper_directory> -L<mpi_lib_directory> -lHDFqL -lmpi
```

- In Linux and macOS using GCC<sup>11</sup>, by executing from a terminal:

- With an HDFqL non MPI-based distribution:

```
gfortran example.f90 -I<hdfql_fortran_wrapper_directory> -
L<hdfql_fortran_wrapper_directory> -lHDFqL -ldl
```

- With an HDFqL MPI-based distribution:

---

<sup>11</sup> An incorrect warning is raised by the GCC Fortran compiler when using the HDFqL module (“Warning: Only array FINAL procedures declared for derived type 'hdfql\_cursor' defined at (1), suggest also scalar one”). This warning does not interfere with the final compilation result, though, and it has been solved in the GCC Fortran compiler version 7.0.0 (please refer to [https://gcc.gnu.org/bugzilla/show\\_bug.cgi?id=58175](https://gcc.gnu.org/bugzilla/show_bug.cgi?id=58175) for additional information).

```
gfortran example.f90 -I<hdfqL_fortran_wrapper_directory> -  
L<hdfqL_fortran_wrapper_directory> -L<mpi_lib_directory> -lHDFqL -lmpi -ldl
```

In case the program does not compile, most likely a Fortran compiler is not installed. If a Fortran compiler is missing, the solution is:

- In Windows, download and install IFORT from the website <https://software.intel.com/en-us/parallel-studio-xe/choose-download/free-trial-cluster-windows-c-fortran>.
- In Linux, download and install IFORT from the website <https://software.intel.com/en-us/parallel-studio-xe/choose-download/free-trial-cluster-linux-fortran>.
- In Linux, install the GCC Fortran compiler by executing from a terminal:

- In a Red Hat-based distribution:

```
sudo yum install gcc-gfortran
```

- In a Debian-based distribution:

```
sudo apt-get install gfortran
```

- In macOS, install the GCC Fortran compiler by executing from a terminal (if xcode-select does not support the parameter “--install” (due to being outdated), download and install the Command-Line Tools package from the website <http://developer.apple.com/downloads> which includes GCC instead):

```
xcode-select --install
```

In case the compiled program does not launch, most likely the HDFqL Fortran shared library and/or the MPI shared library was not found (these are needed to launch the program). The solution is:

- In Windows, add the directory where the file “HDFqL\_dll.dll” is located to the environment variable “PATH” by executing from a terminal:

```
set PATH=<hdfq_fortran_wrapper_directory>;%PATH%
```

- In Linux, add the directories where the files “libHDFq.so” and (optionally) “libmpi.so” are located to the environment variable “LD\_LIBRARY\_PATH” by executing from a terminal:

- With an HDFq non MPI-based distribution:

```
export LD_LIBRARY_PATH=<hdfq_fortran_wrapper_directory>:$LD_LIBRARY_PATH
```

- With an HDFq MPI-based distribution:

```
export
LD_LIBRARY_PATH=<hdfq_fortran_wrapper_directory>:<mpi_lib_directory>:$LD_LIBRARY_PATH
```

- In macOS, add the directories where the files “libHDFq.dylib” and (optionally) “libmpi.dylib” are located to the environment variable “DYLD\_LIBRARY\_PATH”<sup>12</sup> by executing from a terminal:

- With an HDFq non MPI-based distribution:

```
export DYLD_LIBRARY_PATH=<hdfq_fortran_wrapper_directory>:$DYLD_LIBRARY_PATH
```

- With an HDFq MPI-based distribution:

```
export
DYLD_LIBRARY_PATH=<hdfq_fortran_wrapper_directory>:<mpi_lib_directory>:$DYLD_LIBRARY_PATH
```

---

<sup>12</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information).

## 3.7 R

HDFq1 can be used in the R programming language through a wrapper named “HDFq1.R”. This wrapper is stored in the directory “R” found under the directory “wrapper”. The following short script illustrates how HDFq1 can be used in such language.

```
# load HDFq1 R wrapper (make sure it can be found by the R interpreter)
source("HDFq1.R")

# display HDFq1 version in use
print(paste("HDFq1 version:", HDFQ1_VERSION))

# create an HDF5 file named "my_file.h5"
hdfq1_execute("CREATE FILE my_file.h5")

# use (i.e. open) HDF5 file "my_file.h5"
hdfq1_execute("USE FILE my_file.h5")

# create an HDF5 dataset named "my_dataset" of data type int
hdfq1_execute("CREATE DATASET my_dataset AS INT VALUES(10)")

# select (i.e. read) data from dataset "my_dataset" and populate cursor with it
hdfq1_execute("SELECT FROM my_dataset")

# move cursor to the first position within the result set
hdfq1_cursor_first()

# display content of cursor
print(paste("Dataset value:", hdfq1_cursor_get_int()))
```

Assuming that the script is stored in a file named “example.R” it can be launched by executing the following from a terminal:

```
R -f example.R
```

In case the script does not launch, most likely (1) the R interpreter is not installed or (2) the HDFq1 R wrapper and/or the HDFq1 C shared library and/or the MPI shared library was not found (these are needed to launch the script). To fix the former issue, download and install the R interpreter from the website <https://cloud.r-project.org>. To fix the latter issue:

- In Windows, add the directories where the files “HDFq.R” (i.e. the wrapper) and “HDFq\_dll.dll” are located to the environment variable “PATH” by executing from a terminal:

```
set PATH=<hdfq_r_wrapper_directory>;<hdfq_lib_directory>;%PATH%
```

- In Linux, add the directories where the files “HDFq.R”, “libHDFq.so” and (optionally) “libmpi.so” are located to the environment variable “LD\_LIBRARY\_PATH” by executing from a terminal:

- With an HDFq non MPI-based distribution:

```
export
LD_LIBRARY_PATH=<hdfq_r_wrapper_directory>:<hdfq_lib_directory>:$LD_LIBRARY_PATH
```

- With an HDFq MPI-based distribution:

```
export
LD_LIBRARY_PATH=<hdfq_r_wrapper_directory>:<hdfq_lib_directory>:<mpi_lib_directory>:$LD_LIBRARY_PATH
```

- In macOS, add the directories where the files “HDFq.R”, “libHDFq.dylib” and (optionally) “libmpi.dylib” are located to the environment variable “DYLD\_LIBRARY\_PATH”<sup>13</sup> by executing from a terminal:

- With an HDFq non MPI-based distribution:

```
export
DYLD_LIBRARY_PATH=<hdfq_r_wrapper_directory>:<hdfq_lib_directory>:$DYLD_LIBRARY_PATH
```

- With an HDFq MPI-based distribution:

---

<sup>13</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information).



```
export
DYLD_LIBRARY_PATH=<hdfql_r_wrapper_directory>:<hdfql_lib_directory>:<mpi_lib_directory>:$DYLD_LIBRARY_PATH
```

Besides these steps, a package named `bit64` for R must be installed when working with user-defined variables to store 64 bit integers as these are not natively supported by R (please refer to the function `hdfql_variable_register` for additional information). This package can be found at <https://cran.r-project.org/web/packages/bit64> along with instructions on how to install and use it.

## 3.8 COMMAND-LINE INTERFACE

A command-line interface named “HDFqlCLI” is available and can be used for manipulating HDF5 files from a terminal. It is stored in the directory “bin”. To launch the command-line interface, open a terminal (“cmd” if in Windows, “xterm” if in Linux, or “Terminal” if in macOS), go to the directory “bin”, and type:

- In Windows:

```
HDFqlCLI.exe
```

- In Linux and macOS:

```
./HDFqlCLI
```

The list of parameters accepted by the command-line interface can be viewed by launching it with the parameter “--help”. At the time of writing, this list includes the following parameters:

- --help (show the list of parameters accepted by HDFqlCLI and exit)
- --version (show the version of HDFqlCLI and exit)
- --debug (show debug information when executing HDFql operations)
- --no-path (do not show group path currently in use in HDFqlCLI prompt)
- --execute=X (execute HDFql operation(s) “X” and exit)

- `--execute-file=X` (execute HDFqI operation(s) stored in file “X” and exit)
- `--save-file=X` (save executed HDFqI operation(s) to file “X”)

In case the command-line interface does not launch, most likely the HDFqI shared library (which is needed to launch the interface) was not found. Depending on the platform, the solution is:

- In Windows, to either:
  - Add the directory where the file “HDFqI\_dll.dll” is located to the environment variable “PATH” by executing from a terminal:

```
set PATH=<hdfqI_lib_directory>;%PATH%
```

- Execute the batch file named “launch.bat” which properly sets up the environment variable “PATH” and launches the command-line interface from a terminal.
- In Linux, to either:
  - Add the directory where the file “libHDFqI.so” is located to the environment variable “LD\_LIBRARY\_PATH” by executing from a terminal:

```
export LD_LIBRARY_PATH=<hdfqI_lib_directory>:$LD_LIBRARY_PATH
```

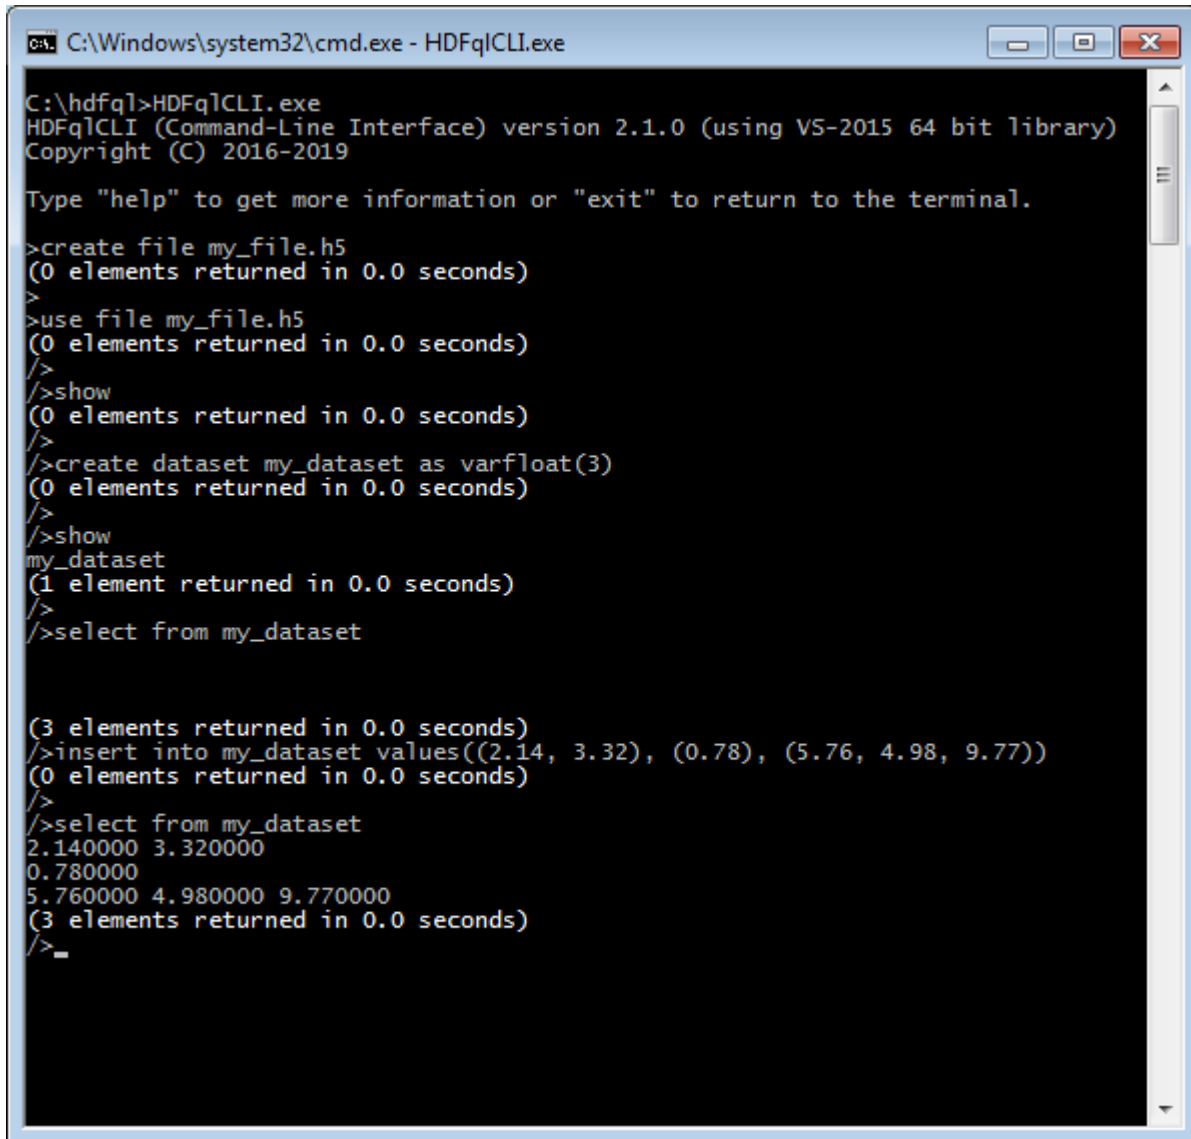
- Execute the bash script file named “launch.sh” which properly sets up the environment variable “LD\_LIBRARY\_PATH” and launches the command-line interface from a terminal.
- In macOS, to either:
  - Add the directory where the file “libHDFqI.dylib” is located to the environment variable “DYLD\_LIBRARY\_PATH”<sup>14</sup> by executing from a terminal:

---

<sup>14</sup> Starting from version 10.11 (i.e. El Capitan), Apple introduced a security feature named System Integrity Protection (SIP) which may prevent setting the environment variable “DYLD\_LIBRARY\_PATH” and, ultimately, launching the program. To overcome this, SIP should be disabled (please refer to [https://developer.apple.com/library/archive/documentation/Security/Conceptual/System\\_Integrity\\_Protection\\_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html](https://developer.apple.com/library/archive/documentation/Security/Conceptual/System_Integrity_Protection_Guide/ConfiguringSystemIntegrityProtection/ConfiguringSystemIntegrityProtection.html) for additional information).

```
export DYLD_LIBRARY_PATH=<hdfqI_lib_directory>:$DYLD_LIBRARY_PATH
```

- Execute the bash script file named “launch.sh” which properly sets up the environment variable “DYLD\_LIBRARY\_PATH” and launches the command-line interface from a terminal.



The screenshot shows a Windows command prompt window titled "C:\Windows\system32\cmd.exe - HDFqI CLI.exe". The prompt is at "C:\hdfqI>". The user has entered the command "HDFqI CLI.exe", which has launched the "HDFqI CLI (Command-Line Interface) version 2.1.0 (using VS-2015 64 bit library)". The copyright notice "(C) 2016-2019" is displayed. A message prompts the user to type "help" for more information or "exit" to return to the terminal. The user then enters a series of commands: "create file my\_file.h5", "use file my\_file.h5", "show", "create dataset my\_dataset as varfloat(3)", "show", "insert into my\_dataset values((2.14, 3.32), (0.78), (5.76, 4.98, 9.77))", and "select from my\_dataset". The output shows the creation of the file and dataset, the insertion of three rows of data, and the retrieval of those three rows. The prompt ends with a underscore character.

```
C:\Windows\system32\cmd.exe - HDFqI CLI.exe
C:\hdfqI>HDFqI CLI.exe
HDFqI CLI (Command-Line Interface) version 2.1.0 (using VS-2015 64 bit library)
Copyright (C) 2016-2019

Type "help" to get more information or "exit" to return to the terminal.

>create file my_file.h5
(0 elements returned in 0.0 seconds)
>
>use file my_file.h5
(0 elements returned in 0.0 seconds)
/>
/>show
(0 elements returned in 0.0 seconds)
/>
/>create dataset my_dataset as varfloat(3)
(0 elements returned in 0.0 seconds)
/>
/>show
my_dataset
(1 element returned in 0.0 seconds)
/>
/>select from my_dataset

(3 elements returned in 0.0 seconds)
/>insert into my_dataset values((2.14, 3.32), (0.78), (5.76, 4.98, 9.77))
(0 elements returned in 0.0 seconds)
/>
/>select from my_dataset
2.140000 3.320000
0.780000
5.760000 4.980000 9.770000
(3 elements returned in 0.0 seconds)
/>_
```

Figure 3.1 – Illustration of the command-line interface “HDFqI CLI”

---

## 4. CURSOR

---

Generally speaking, a cursor is a control structure that is used to iterate through the results returned by a query (that was previously executed). It can be seen as an effective means to abstract the programmer from low-level implementation details of accessing data stored in specific structures. This chapter provides a description of cursors and subcursors in HDFqI, as well as examples and illustrations to demonstrate these two concepts in practice.

### 4.1 DESCRIPTION

HDFqI provides cursors which offer several ways to traverse result sets according to specific needs. The following list enumerates these functionalities (please refer to their links for further information):

- First (moves cursor to the first position within the result set – [hdfqI\\_cursor\\_first](#))
- Last (moves cursor to the last position within the result set – [hdfqI\\_cursor\\_last](#))
- Next (moves cursor to the next position within the result set – [hdfqI\\_cursor\\_next](#))
- Previous (moves cursor to the previous position within the result set – [hdfqI\\_cursor\\_previous](#))
- Absolute (moves cursor to an absolute position within the result set – [hdfqI\\_cursor\\_absolute](#))
- Relative (moves cursor to a relative position within the result set – [hdfqI\\_cursor\\_relative](#))

Besides their traversal functionalities, a particular feature of cursors in HDFqI is that they store result sets returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations. To retrieve values from result sets, the functions starting with “hdfqI\_cursor\_get” can be used. These and remaining functions offered by cursors can be found in [Table 5.8](#) (each of these begins with the prefix “hdfqI\_cursor”).

When a certain operation is executed, HDFqI stores the result set returned by this operation in its default cursor. This cursor is available to the programmer and is automatically created and initialized upon loading the HDFqI library by a program. If additional cursors are needed, they can be created like this (in C):

```
// create a cursor named "my_cursor"  
HDFQL_CURSOR my_cursor;
```

As a side note, additional cursors are created in C++, Java, Python, C#, Fortran and R as follows:

```
// create a cursor named "myCursor" in C++  
HDFqL::Cursor myCursor;
```

```
// create a cursor named "myCursor" in Java  
HDFqLCursor myCursor = new HDFqLCursor();
```

```
# create a cursor named "my_cursor" in Python  
my_cursor = HDFqL.Cursor()
```

```
// create a cursor named "myCursor" in C#  
HDFqLCursor myCursor = new HDFqLCursor();
```

```
! create a cursor named "my_cursor" in Fortran  
TYPE (HDFQL_CURSOR) :: my_cursor
```

```
# create a cursor named "my_cursor" in R  
my_cursor <- hdfql_cursor()
```

Before an additional (i.e. user-defined) cursor is used to store and eventually traverse a result set, it must be properly initialized (refer to the function [hdfql\\_cursor\\_initialize](#) for further information). The initialization of a cursor is only required in C and performed once, while in C++, Java, Python, C#, Fortran and R such initialization is redundant as it is done automatically when creating a cursor. Initializing a cursor can be done like this (in C):

```
// initialize a cursor named "my_cursor"  
hdfql_cursor_initialize(&my_cursor);
```

To switch between different cursors (to be used for separate needs), the function `hdfql_cursor_use` may be employed (in C):

```
// use a cursor named "my_cursor"
hdfql_cursor_use(&my_cursor);
```

The following C snippet illustrates usage of the HDFqL default cursor and a user-defined cursor, as well as some typical operations performed on/by these.

```
// create a cursor named "my_cursor"
HDFQL_CURSOR my_cursor;

// create an HDF5 dataset named "my_dataset0" of data type int with an initial value of 8
hdfql_execute("CREATE DATASET my_dataset0 AS INT VALUES(8)");

// create an HDF5 dataset named "my_dataset1" of data type float with initial values of 3.2,
// 5.3, 7.4 and 9.5
hdfql_execute("CREATE DATASET my_dataset1 AS FLOAT(4) VALUES(3.2, 5.3, 7.4, 9.5)");

// select (i.e. read) data from dataset "my_dataset0" and populate HDFqL default cursor with it
hdfql_execute("SELECT FROM my_dataset0");

// initialize cursor "my_cursor"
hdfql_cursor_initialize(&my_cursor);

// use cursor "my_cursor"
hdfql_cursor_use(&my_cursor);

// select (i.e. read) data from dataset "my_dataset1" and populate cursor "my_cursor" with it
hdfql_execute("SELECT FROM my_dataset1");

// use HDFqL default cursor
hdfql_cursor_use(NULL);

// display number of elements in HDFqL default cursor
printf("Number of elements in HDFqL default cursor is %d\n", hdfql_cursor_get_count(NULL));

// move HDFqL default cursor to the next position within the result set
hdfql_cursor_next(NULL);
```

```
// display element of HDFqL default cursor
printf("Current element of HDFqL default cursor is %d\n", *hdfql_cursor_get_int(NULL));

// display number of elements in cursor "my_cursor"
printf("Number of elements in cursor \"my_cursor\" is %d\n",
hdfql_cursor_get_count(&my_cursor));

// use cursor "my_cursor"
hdfql_cursor_use(&my_cursor);

// display elements of cursor "my_cursor"
while(hdfql_cursor_next(NULL) == HDFQL_SUCCESS)
{
    printf("Current element of cursor \"my_cursor\" is %f\n", *hdfql_cursor_get_float(NULL));
}
```

The output of executing the snippet would be similar to this:

```
Number of elements in HDFqL default cursor is 1
Current element of HDFqL default cursor is 8
Number of elements in cursor "my_cursor" is 4
Current element of cursor "my_cursor" is 3.2
Current element of cursor "my_cursor" is 5.3
Current element of cursor "my_cursor" is 7.4
Current element of cursor "my_cursor" is 9.5
```

When populating a cursor with data from a dataset or attribute with two or more dimensions, the data is always linearized into a single dimension. The linearization process is depicted in [Figure 4.1](#). Subsequently, if need be, it is up to the programmer to access the data (stored in the cursor) according to its original dimensions. In this case, the [SHOW DIMENSION](#) operation – which returns the original dimensions of a dataset or attribute – may be useful to help in the task of going from one dimension to the original dimensions.

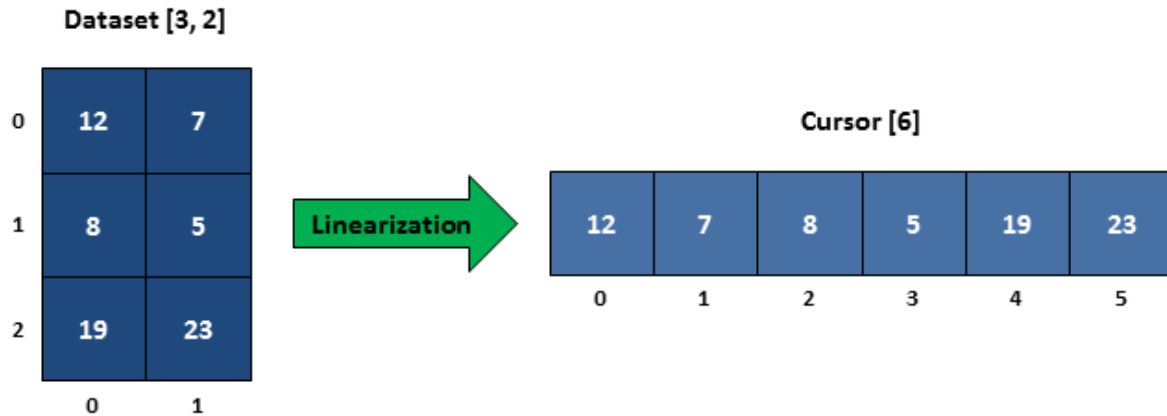


Figure 4.1 – Linearization of a two dimensional dataset into a (one dimensional) cursor

## 4.2 SUBCURSOR

HDFql also provides subcursors – they are meant to complement (i.e. help) cursors in the task of storing data of type [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) and [HDFQL\\_OPAQUE](#). In practice, when a result set is of one of these data types, only the first element of the result set is stored in the cursor (as expected), while all elements of the result set are stored in the subcursor. In other words, each position of the cursor stores the first element of the result set and also points to a subcursor that in turn stores all the elements of the result set. The values stored in a subcursor (which are also known as a result subset) can be accessed with the functions starting with “hdfql\_subcursor\_get” (enumerated in [Table 5.8](#)). Similar to cursors, HDFql subcursors offer several ways to traverse result subsets, namely:

- First (moves subcursor to the first position within the result subset – [hdfql\\_subcursor\\_first](#))
- Last (moves subcursor to the last position within the result subset – [hdfql\\_subcursor\\_last](#))
- Next (moves subcursor to the next position within the result subset – [hdfql\\_subcursor\\_next](#))
- Previous (moves subcursor to the previous position within the result subset – [hdfql\\_subcursor\\_previous](#))
- Absolute (moves subcursor to an absolute position within the result subset – [hdfql\\_subcursor\\_absolute](#))
- Relative (moves subcursor to a relative position within the result subset – [hdfql\\_subcursor\\_relative](#))



The following C snippet illustrates usage of the HDFql subcursors, as well as some typical operations performed on/by these.

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of one dimension
// (size 4)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(4)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((7, 8, 5), (9), (6, 1, 2), (4, 0))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the next position within the result set
while(hdfql_cursor_next(NULL) == HDFQL_SUCCESS)
{
    // display element of the cursor in use
    printf("Current element of cursor is %d\n", *hdfql_cursor_get_int(NULL));

    // move the subcursor in use to the next position within the result subset
    while(hdfql_subcursor_next(NULL) == HDFQL_SUCCESS)
    {
        // display element of the subcursor in use
        printf("    Current element of subcursor is %d\n", *hdfql_subcursor_get_int(NULL));
    }
}
```

The output of executing the snippet would be similar to this:

```
Current element of cursor is 7
    Current element of subcursor is 7
    Current element of subcursor is 8
    Current element of subcursor is 5
Current element of cursor is 9
    Current element of subcursor is 9
Current element of cursor is 6
    Current element of subcursor is 6
    Current element of subcursor is 1
    Current element of subcursor is 2
Current element of cursor is 4
    Current element of subcursor is 4
    Current element of subcursor is 0
```

## 4.3 EXAMPLES

The following C snippets demonstrate how HDFql cursors and subcursors are populated with (variable) data stored in HDF5 datasets or attributes, along with illustrations to facilitate understanding of the populating process and its final result.

```
// create an HDF5 dataset named "my_dataset0" of data type short
hdfql_execute("CREATE DATASET my_dataset0 AS SMALLINT");

// insert (i.e. write) a value into dataset "my_dataset0"
hdfql_execute("INSERT INTO my_dataset0 VALUES(7)");

// select (i.e. read) data from dataset "my_dataset0" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset0");
```

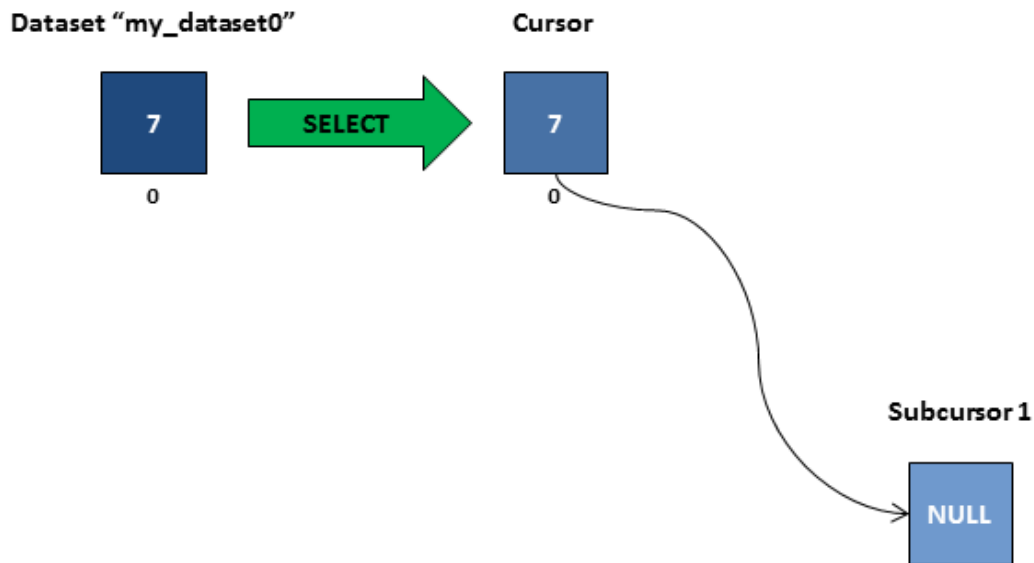


Figure 4.2 – Cursor populated with data from dataset “my\_dataset0”

```
// create an HDF5 dataset named "my_dataset1" of data type float of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset1 AS FLOAT(3)");

// insert (i.e. write) values into dataset "my_dataset1"
hdfql_execute("INSERT INTO my_dataset1 VALUES(5.5, 8.1, 4.9)");

// select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset1");
```

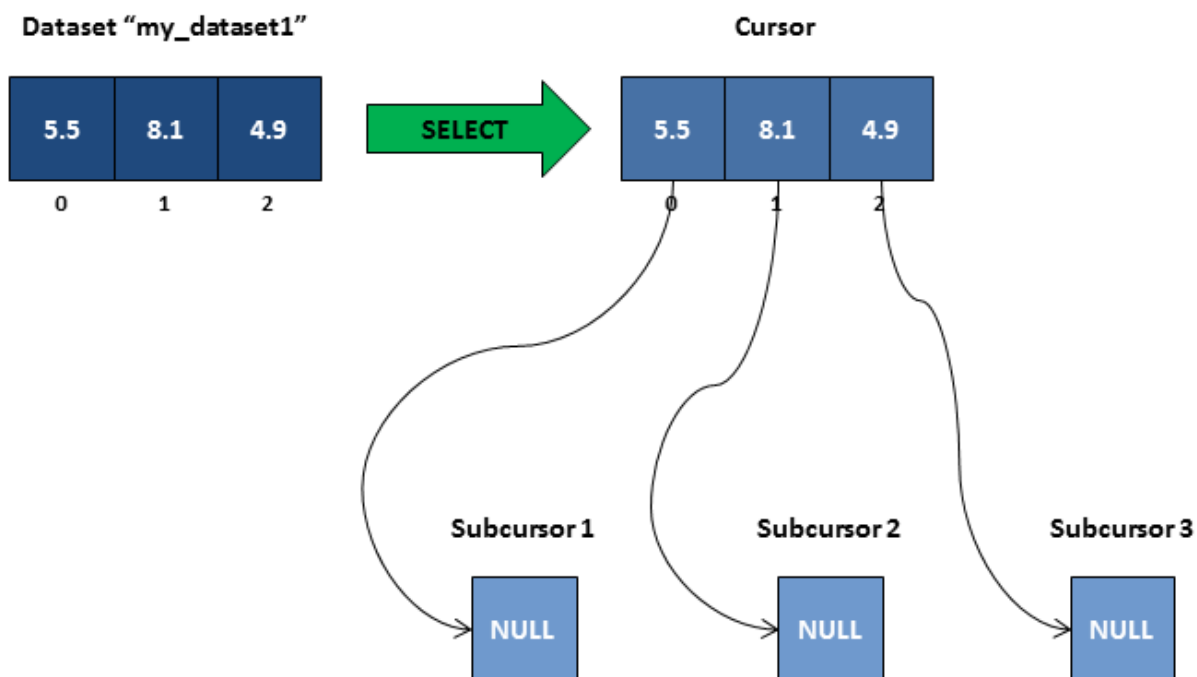


Figure 4.3 – Cursor populated with data from dataset "my\_dataset1"

```
// create an HDF5 dataset named "my_dataset2" of data type double of two dimensions (size 3x2)
hdfql_execute("CREATE DATASET my_dataset2 AS DOUBLE(3, 2)");

// insert (i.e. write) values into dataset "my_dataset2"
hdfql_execute("INSERT INTO my_dataset2 VALUES((3.2, 1.3), (0, 0.2), (9.1, 6.5))");

// select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset2");
```

Dataset "my\_dataset2"

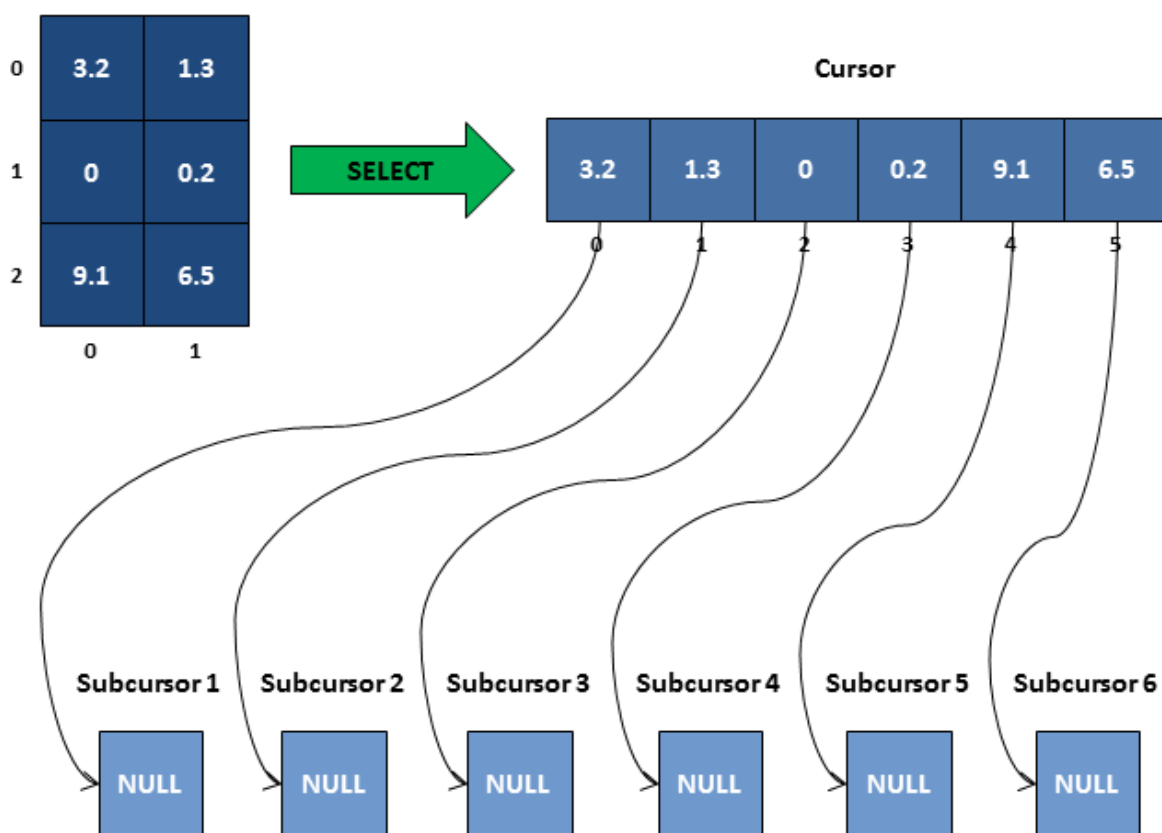


Figure 4.4 – Cursor populated with data from dataset "my\_dataset2"

```
// create an HDF5 dataset named "my_dataset3" of data type variable-length short
hdfql_execute("CREATE DATASET my_dataset3 AS VARSMALLINT");

// insert (i.e. write) values into dataset "my_dataset3"
hdfql_execute("INSERT INTO my_dataset3 VALUES(7, 9, 3)");

// select (i.e. read) data from dataset "my_dataset3" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset3");
```

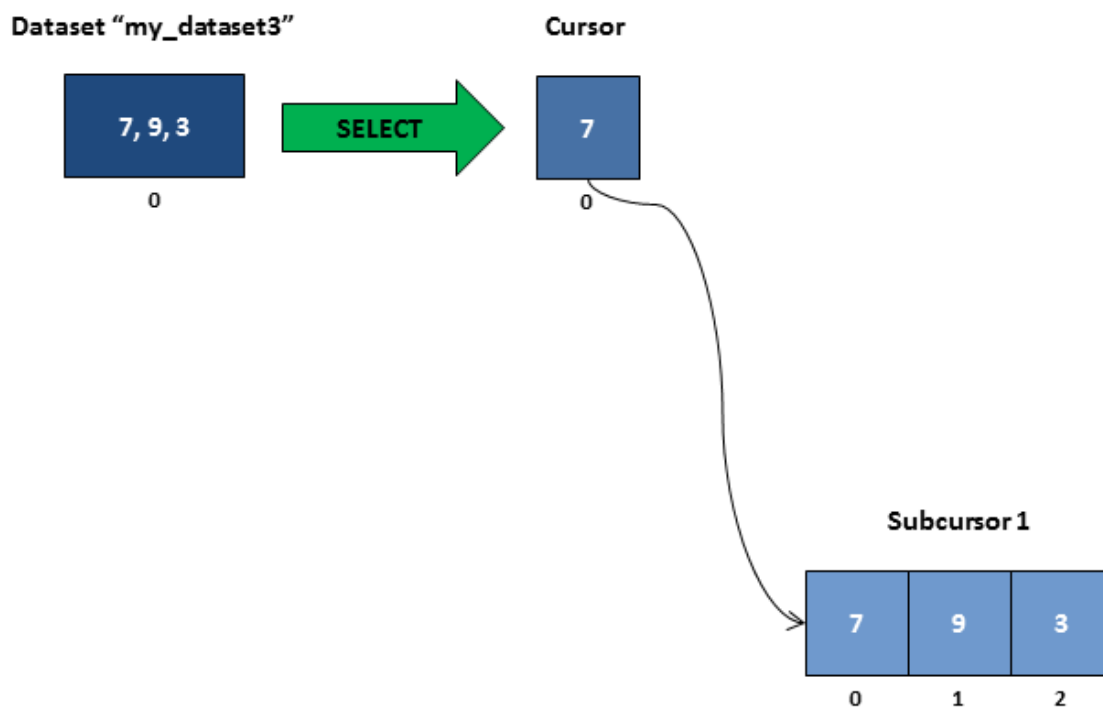


Figure 4.5 – Cursor and its subcursor populated with data from dataset "my\_dataset3"

```
// create an HDF5 dataset named "my_dataset4" of data type variable-length float of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset4 AS VARFLOAT(3)");

// insert (i.e. write) values into dataset "my_dataset4"
hdfql_execute("INSERT INTO my_dataset4 VALUES((5.5), (8.1, 2.2), (4.9, 3.4, 5.6))");

// select (i.e. read) data from dataset "my_dataset4" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset4");
```

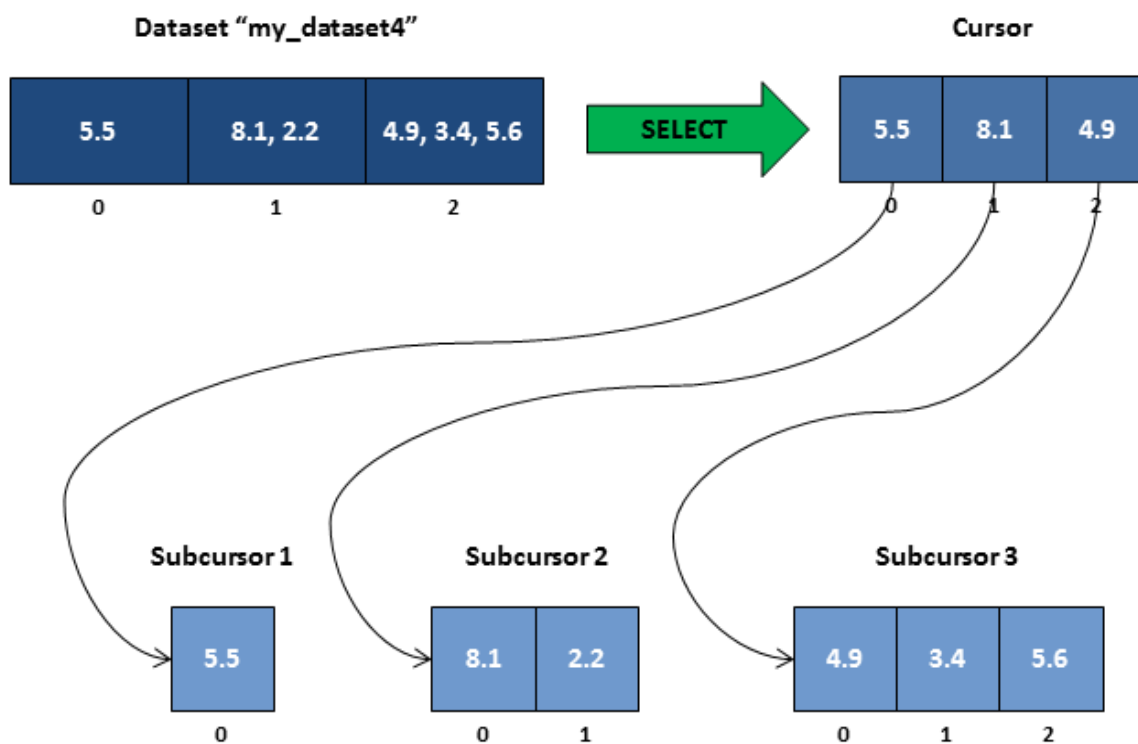


Figure 4.6 – Cursor and its subcursors populated with data from dataset "my\_dataset4"

```
// create an HDF5 dataset named "my_dataset5" of data type variable-length double of two
dimensions (size 3x2)
hdfql_execute("CREATE DATASET my_dataset5 AS VARDOUBLE(3, 2)");

// insert (i.e. write) values into dataset "my_dataset5"
hdfql_execute("INSERT INTO my_dataset5 VALUES(((3.2, 8, 6.7), (1.3, 0.2)), ((0), (0.2, 1.5)),
((9.1, 2, 4, 7), (6.5)))");

// select (i.e. read) data from dataset "my_dataset5" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset5");
```

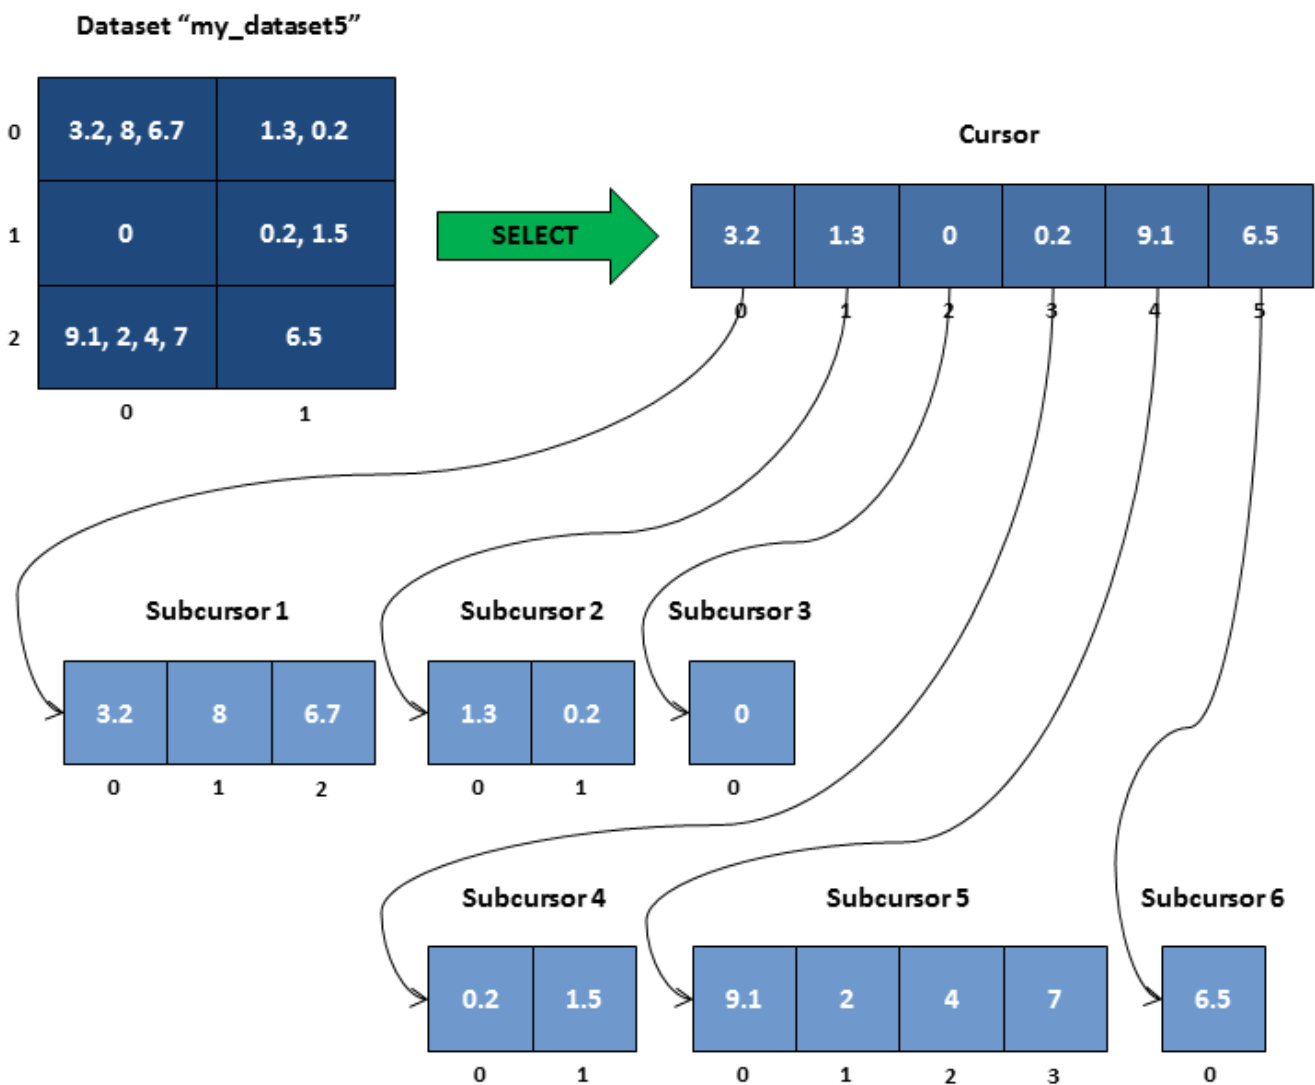


Figure 4.7 – Cursor and its subcursors populated with data from dataset "my\_dataset5"

---

## 5. APPLICATION PROGRAMMING INTERFACE

---

An application programming interface (API) specifies how software components should interact with each other. In practice, an API comes in the form of a library that includes specifications for functions, data structures, object classes, constants and variables. A good API makes it easier to develop a program by providing all the building blocks. This chapter is devoted to describing HDFqI API and how to use it through practical examples in C, C++, Java, Python, C#, Fortran and R.

### 5.1 CONSTANTS

A constant is an identifier whose associated value cannot typically be altered by the program during its execution. Using a constant instead of specifying a value multiple times in the program not only simplifies code maintenance, but can also supply a meaningful name for it. Constants in the C programming languages follow a naming convention of writing all words in uppercase and separating each word with an underscore (\_). The following table summarizes all existing HDFqI constants in C.

HDFqI Constant in C	Description	Data Type	Value
HDFQL_VERSION	Represents the HDFqI version in use	char *	2.1.0
HDFQL_YES	Represents the concept “Yes”	int	0
HDFQL_NO	Represents the concept “No”	int	-1
HDFQL_ENABLED	Represents the concept “Enabled”	int	0
HDFQL_DISABLED	Represents the concept “Disabled”	int	-1
HDFQL_UNLIMITED	Represents the concept “Unlimited”	int	-1
HDFQL_UNDEFINED	Represents the concept “Undefined”	int	-1
HDFQL_GLOBAL	Represents the concept “Global”	int	1
HDFQL_LOCAL	Represents the concept “Local”	int	2
HDFQL_TRACKED	Represents the HDF5 tracked creation order	int	1



	strategy		
HDFQL_INDEXED	Represents the HDF5 indexed creation order strategy	int	2
HDFQL_CONTIGUOUS	Represents the HDF5 contiguous storage type (layout)	int	1
HDFQL_COMPACT	Represents the HDF5 compact storage type (layout)	int	2
HDFQL_CHUNKED	Represents the HDF5 chunked storage type (layout)	int	4
HDFQL_EARLY	Represents the HDF5 early storage allocation	int	1
HDFQL_INCREMENTAL	Represents the HDF5 incremental storage allocation	int	2
HDFQL_LATE	Represents the HDF5 late storage allocation	int	4
HDFQL_DIRECTORY	Represents a directory	int	1
HDFQL_FILE	Represents a file	int	2
HDFQL_GROUP	Represents the HDF5 group object type	int	4
HDFQL_DATASET	Represents the HDF5 dataset object type	int	8
HDFQL_ATTRIBUTE	Represents the HDF5 attribute object type	int	16
HDFQL_SOFT_LINK	Represents the HDF5 soft link object type	int	32
HDFQL_EXTERNAL_LINK	Represents the HDF5 external link object type	int	64
HDFQL_TINYINT	Represents the tiny integer data type ( <b>TINYINT</b> )	int	1
HDFQL_UNSIGNED_TINYINT	Represents the unsigned tiny integer data type ( <b>UNSIGNED TINYINT</b> )	int	2
HDFQL_SMALLINT	Represents the small integer data type ( <b>SMALLINT</b> )	int	4
HDFQL_UNSIGNED_SMALLINT	Represents the unsigned small integer data type ( <b>UNSIGNED SMALLINT</b> )	int	8
HDFQL_INT	Represents the integer data type ( <b>INT</b> )	int	16
HDFQL_UNSIGNED_INT	Represents the unsigned integer data type ( <b>UNSIGNED INT</b> )	int	32

HDFQL_BIGINT	Represents the big integer data type ( <b>BIGINT</b> )	int	64
HDFQL_UNSIGNED_BIGINT	Represents the unsigned big integer data type ( <b>UNSIGNED BIGINT</b> )	int	128
HDFQL_FLOAT	Represents the float data type ( <b>FLOAT</b> )	int	256
HDFQL_DOUBLE	Represents the double data type ( <b>DOUBLE</b> )	int	512
HDFQL_CHAR	Represents the char data type ( <b>CHAR</b> )	int	1024
HDFQL_VARTINYINT	Represents the variable-length tiny integer data type ( <b>VARTINYINT</b> )	int	2048
HDFQL_UNSIGNED_VARTINYINT	Represents the unsigned variable-length tiny integer data type ( <b>UNSIGNED VARTINYINT</b> )	int	4096
HDFQL_VARSMALLINT	Represents the variable-length small integer data type ( <b>VARSMALLINT</b> )	int	8192
HDFQL_UNSIGNED_VARSMALLINT	Represents the unsigned variable-length small integer data type ( <b>UNSIGNED VARSMALLINT</b> )	int	16384
HDFQL_VARINT	Represents the variable-length integer data type ( <b>VARINT</b> )	int	32768
HDFQL_UNSIGNED_VARINT	Represents the unsigned variable-length integer data type ( <b>UNSIGNED VARINT</b> )	int	65536
HDFQL_VARBIGINT	Represents the variable-length big integer data type ( <b>VARBIGINT</b> )	int	131072
HDFQL_UNSIGNED_VARBIGINT	Represents the unsigned variable-length big integer data type ( <b>UNSIGNED VARBIGINT</b> )	int	262144
HDFQL_VARFLOAT	Represents the variable-length float data type ( <b>VARFLOAT</b> )	int	524288
HDFQL_VARDOUBLE	Represents the variable-length double data type ( <b>VARDOUBLE</b> )	int	1048576
HDFQL_VARCHAR	Represents the variable-length char data type ( <b>VARCHAR</b> )	int	2097152
HDFQL_OPAQUE	Represents the opaque data type ( <b>OPAQUE</b> )	int	4194304
HDFQL_BITFIELD	Represents the bitfield data type	int	8388608
HDFQL_ENUMERATION	Represents the enumeration data type	int	16777216
HDFQL_COMPOUND	Represents the compound data type	int	33554432

HDFQL_LITTLE_ENDIAN	Represents the little endian byte ordering	int	1
HDFQL_BIG_ENDIAN	Represents the big endian byte ordering	int	2
HDFQL_ASCII	Represents the ASCII character encoding	int	1
HDFQL_UTF8	Represents the UTF8 character encoding	int	2
HDFQL_FILL_DEFAULT	Represents the default fill type	int	1
HDFQL_FILL_USER_DEFINED	Represents the user defined fill type	int	2
HDFQL_FILL_UNDEFINED	Represents the undefined fill type	int	4
HDFQL_EARLIEST	Represents the HDF5 library bound earliest	int	1
HDFQL_LATEST	Represents the HDF5 library bound latest	int	2
HDFQL_VERSION_18	Represents the HDF5 library bound version 18	int	4
HDFQL_SUCCESS	Represents an operation that succeeded	int	0
HDFQL_ERROR_PARSE	Represents an operation that failed due to a parsing error	int	-1
HDFQL_ERROR_NOT_FOUND	Represents an operation that failed due to an object (e.g. directory, file, group, dataset) not being found	int	-2
HDFQL_ERROR_NO_ACCESS	Represents an operation that failed due to an object (e.g. directory, file, group, dataset) not being accessible	int	-3
HDFQL_ERROR_NOT_OPEN	Represents an operation that failed due to an object (e.g. file) not being opened	int	-4
HDFQL_ERROR_INVALID_FILE	Represents an operation that failed due to a file being invalid (i.e. not a valid HDF5 file)	int	-5
HDFQL_ERROR_NOT_SUPPORTED	Represents an operation that failed due to not being supported	int	-6
HDFQL_ERROR_NOT_ENOUGH_SPACE	Represents an operation that failed due to the machine not having enough (storage) space	int	-7
HDFQL_ERROR_NOT_ENOUGH_MEMORY	Represents an operation that failed due to the machine not having enough (RAM) memory	int	-8
HDFQL_ERROR_ALREADY_EXISTS	Represents an operation that failed due to an object (e.g. directory, file, group, dataset) already	int	-9

	existing		
HDFQL_ERROR_EMPTY	Represents an operation that failed due to its internal structure being empty	int	-10
HDFQL_ERROR_FULL	Represents an operation that failed due to its internal structure being full	int	-11
HDFQL_ERROR_BEFORE_FIRST	Represents an operation that failed due to trying to position/access an element before the first one	int	-12
HDFQL_ERROR_AFTER_LAST	Represents an operation that failed due to trying to position/access an element after the last one	int	-13
HDFQL_ERROR_OUTSIDE_LIMIT	Represents an operation that failed due to being outside the limit	int	-14
HDFQL_ERROR_NO_ADDRESS	Represents an operation that failed due to a user-defined variable having no address (i.e. is NULL)	int	-15
HDFQL_ERROR_UNEXPECTED_TYPE	Represents an operation that failed due to an object (e.g. group, dataset) being of an unexpected type	int	-16
HDFQL_ERROR_UNEXPECTED_DATA_TYPE	Represents an operation that failed due to a user-defined variable being of an unexpected data type	int	-17
HDFQL_ERROR_UNEXPECTED_STORAGE_TYPE	Represents an operation that failed due to a dataset being of an unexpected storage type (layout)	int	-18
HDFQL_ERROR_DANGLING_LINK	Represents an operation that failed due to an object being a dangling (soft or external) link	int	-19
HDFQL_ERROR_NOT_REGISTERED	Represents an operation that failed due to a user-defined variable not being registered	int	-20
HDFQL_ERROR_INVALID_REGULAR_EXPRESSION	Represents an operation that failed due to a regular expression being invalid	int	-21
HDFQL_ERROR_UNKNOWN	Represents an operation that failed due to an unknown/unexpected error	int	-99

Table 5.1 – HDFqI constants in C

HDFq also supports other programming languages namely C++, Java, Python, C#, Fortran and R through wrappers. The below tables provide examples on how HDFq constants are defined in these programming languages.

In C++, the prefix “HDFQL\_” of the name of constants (defined in C) is replaced by the namespace “HDFq” and its underscores ( \_ ) are discarded. The remainder of the name of constants follows the upper camel-case convention. The following table lists a subset of HDFq constants as defined in C and details how these are defined/can be used in C++.

HDFq Constant in C	Corresponding Definition in C++
HDFQL_VERSION	HDFq::Version
HDFQL_SUCCESS	HDFq::Success
HDFQL_ERROR_PARSE	HDFq::ErrorParse
HDFQL_TINYINT	HDFq::TinyInt
HDFQL_UNSIGNED_BIGINT	HDFq::UnsignedBigInt
HDFQL_UTF8	HDFq::Utf8

Table 5.2 – HDFq constants in C and their corresponding definitions in C++

In Java, the prefix “HDFQL\_” of the name of constants (defined in C) is replaced by the class “HDFq”. The remainder of the name of constants remains exactly the same. The following table lists a subset of HDFq constants as defined in C and details how these are defined/can be used in Java.

HDFq Constant in C	Corresponding Definition in Java
HDFQL_VERSION	HDFq.VERSION
HDFQL_SUCCESS	HDFq.SUCCESS
HDFQL_ERROR_PARSE	HDFq.ERROR_PARSE
HDFQL_TINYINT	HDFq.TINYINT
HDFQL_UNSIGNED_BIGINT	HDFq.UNSIGNED_BIGINT
HDFQL_UTF8	HDFq.UTF8

Table 5.3 – HDFqI constants in C and their corresponding definitions in Java

In Python, the prefix “HDFQL\_” of the name of constants (defined in C) is replaced by the class “HDFqI”. The remainder of the name of constants remains exactly the same. The following table lists a subset of HDFqI constants as defined in C and details how these are defined/can be used in Python.

HDFqI Constant in C	Corresponding Definition in Python
HDFQL_VERSION	HDFqI.VERSION
HDFQL_SUCCESS	HDFqI.SUCCESS
HDFQL_ERROR_PARSE	HDFqI.ERROR_PARSE
HDFQL_TINYINT	HDFqI.TINYINT
HDFQL_UNSIGNED_BIGINT	HDFqI.UNSIGNED_BIGINT
HDFQL_UTF8	HDFqI.UTF8

Table 5.4 – HDFqI constants in C and their corresponding definitions in Python

In C#, the prefix “HDFQL\_” of the name of constants (defined in C) is replaced by the class “HDFqI” and its underscores (\_) are discarded. The remainder of the name of constants follows the upper camel-case convention. The following table lists a subset of HDFqI constants as defined in C and details how these are defined/can be used in C#.

HDFqI Constant in C	Corresponding Definition in C#
HDFQL_VERSION	HDFqI.Version
HDFQL_SUCCESS	HDFqI.Success
HDFQL_ERROR_PARSE	HDFqI.ErrorParse
HDFQL_TINYINT	HDFqI.TinyInt
HDFQL_UNSIGNED_BIGINT	HDFqI.UnsignedBigInt
HDFQL_UTF8	HDFqI.Utf8

Table 5.5 – HDFqI constants in C and their corresponding definitions in C#

In Fortran, the name of constants is the same as in C and can be written in any case. The following table lists a subset of HDFq constants as defined in C and details how these are defined/can be used in Fortran.

HDFq Constant in C	Corresponding Definition in Fortran
HDFQL_VERSION	HDFQL_VERSION
HDFQL_SUCCESS	HDFQL_SUCCESS
HDFQL_ERROR_PARSE	HDFQL_ERROR_PARSE
HDFQL_TINYINT	HDFQL_TINYINT
HDFQL_UNSIGNED_BIGINT	HDFQL_UNSIGNED_BIGINT
HDFQL_UTF8	HDFQL_UTF8

Table 5.6 – HDFq constants in C and their corresponding definitions in Fortran

In R, the name of constants is the same as in C. The following table lists a subset of HDFq constants as defined in C and details how these are defined/can be used in R.

HDFq Constant in C	Corresponding Definition in R
HDFQL_VERSION	HDFQL_VERSION
HDFQL_SUCCESS	HDFQL_SUCCESS
HDFQL_ERROR_PARSE	HDFQL_ERROR_PARSE
HDFQL_TINYINT	HDFQL_TINYINT
HDFQL_UNSIGNED_BIGINT	HDFQL_UNSIGNED_BIGINT
HDFQL_UTF8	HDFQL_UTF8

Table 5.7 – HDFq constants in C and their corresponding definitions in R

## 5.2 FUNCTIONS

A function is a group of instructions that together perform a specific task, requiring direction back to the caller on completion of the task. Any given function might be called at any point during a program's execution, including by other functions or itself. It provides better modularity of a program and a high degree of code reusing. The following table summarizes all existing HDFql functions in C.

HDFql Function in C	Description
<code>hdfql_execute</code>	Execute a script (composed of one or more operations)
<code>hdfql_execute_get_status</code>	Get status of the last executed operation
<code>hdfql_error_get_line</code>	Get error line of the last executed operation
<code>hdfql_error_get_position</code>	Get error position of the last executed operation
<code>hdfql_error_get_message</code>	Get error message of the last executed operation
<code>hdfql_cursor_initialize</code>	Initialize a cursor for subsequent use
<code>hdfql_cursor_use</code>	Set the cursor to be used for storing the result of operations
<code>hdfql_cursor_use_default</code>	Set HDFql default cursor as the one to be used for storing the result of operations
<code>hdfql_cursor_clear</code>	Clear (i.e. empty) the cursor in use
<code>hdfql_cursor_clone</code>	Clone (i.e. duplicate) a cursor into another one
<code>hdfql_cursor_get_data_type</code>	Get data type of the cursor in use
<code>hdfql_cursor_get_count</code>	Get number of elements (i.e. result set size) stored in the cursor in use
<code>hdfql_subcursor_get_count</code>	Get number of elements (i.e. result subset size) stored in the subcursor in use
<code>hdfql_cursor_get_position</code>	Get current position of cursor in use within result set
<code>hdfql_subcursor_get_position</code>	Get current position of subcursor in use within result subset
<code>hdfql_cursor_first</code>	Move the cursor in use to the first position within result set
<code>hdfql_subcursor_first</code>	Move the subcursor in use to the first position within result subset
<code>hdfql_cursor_last</code>	Move the cursor in use to the last position within result set
<code>hdfql_subcursor_last</code>	Move the subcursor in use to the last position within result subset



<a href="#">hdfql_cursor_next</a>	Move the cursor in use one position forward from its current position
<a href="#">hdfql_subcursor_next</a>	Move the subcursor in use one position forward from its current position
<a href="#">hdfql_cursor_previous</a>	Move the cursor in use one position backward from its current position
<a href="#">hdfql_subcursor_previous</a>	Move the subcursor in use one position backward from its current position
<a href="#">hdfql_cursor_absolute</a>	Move the cursor in use to an absolute position within the result set
<a href="#">hdfql_subcursor_absolute</a>	Move the subcursor in use to an absolute position within the result subset
<a href="#">hdfql_cursor_relative</a>	Move the cursor in use to a relative position within result set
<a href="#">hdfql_subcursor_relative</a>	Move the subcursor in use to a relative position within result subset
<a href="#">hdfql_cursor_get_size</a>	Get current element size (in bytes) of the cursor in use
<a href="#">hdfql_subcursor_get_size</a>	Get current element size (in bytes) of the subcursor in use
<a href="#">hdfql_cursor_get</a>	Get current element of the cursor in use as a generic (typeless) pointer
<a href="#">hdfql_subcursor_get</a>	Get current element of the subcursor in use as a generic (typeless) pointer
<a href="#">hdfql_cursor_get_tinyint</a>	Get current element of the cursor in use as a <a href="#">TINYINT</a>
<a href="#">hdfql_subcursor_get_tinyint</a>	Get current element of the subcursor in use as a <a href="#">TINYINT</a>
<a href="#">hdfql_cursor_get_unsigned_tinyint</a>	Get current element of the cursor in use as an <a href="#">UNSIGNED TINYINT</a>
<a href="#">hdfql_subcursor_get_unsigned_tinyint</a>	Get current element of the subcursor in use as an <a href="#">UNSIGNED TINYINT</a>
<a href="#">hdfql_cursor_get_smallint</a>	Get current element of the cursor in use as a <a href="#">SMALLINT</a>
<a href="#">hdfql_subcursor_get_smallint</a>	Get current element of the subcursor in use as a <a href="#">SMALLINT</a>
<a href="#">hdfql_cursor_get_unsigned_smallint</a>	Get current element of the cursor in use as an <a href="#">UNSIGNED SMALLINT</a>
<a href="#">hdfql_subcursor_get_unsigned_smallint</a>	Get current element of the subcursor in use as an <a href="#">UNSIGNED SMALLINT</a>
<a href="#">hdfql_cursor_get_int</a>	Get current element of the cursor in use as an <a href="#">INT</a>
<a href="#">hdfql_subcursor_get_int</a>	Get current element of the subcursor in use as an <a href="#">INT</a>
<a href="#">hdfql_cursor_get_unsigned_int</a>	Get current element of the cursor in use as an <a href="#">UNSIGNED INT</a>
<a href="#">hdfql_subcursor_get_unsigned_int</a>	Get current element of the subcursor in use as an <a href="#">UNSIGNED INT</a>
<a href="#">hdfql_cursor_get_bigint</a>	Get current element of the cursor in use as a <a href="#">BIGINT</a>

<code>hdfql_subcursor_get_bigint</code>	Get current element of the subcursor in use as a <b>BIGINT</b>
<code>hdfql_cursor_get_unsigned_bigint</code>	Get current element of the cursor in use as an <b>UNSIGNED BIGINT</b>
<code>hdfql_subcursor_get_unsigned_bigint</code>	Get current element of the subcursor in use as an <b>UNSIGNED BIGINT</b>
<code>hdfql_cursor_get_float</code>	Get current element of the cursor in use as a <b>FLOAT</b>
<code>hdfql_subcursor_get_float</code>	Get current element of the subcursor in use as a <b>FLOAT</b>
<code>hdfql_cursor_get_double</code>	Get current element of the cursor in use as a <b>DOUBLE</b>
<code>hdfql_subcursor_get_double</code>	Get current element of the subcursor in use as a <b>DOUBLE</b>
<code>hdfql_cursor_get_char</code>	Get current element of the cursor in use as a <b>CHAR</b>
<code>hdfql_subcursor_get_char</code>	Get current element of the subcursor in use as a <b>CHAR</b>
<code>hdfql_variable_register</code>	Register a variable for subsequent use
<code>hdfql_variable_transient_register</code>	Register a variable in a transient way for subsequent use
<code>hdfql_variable_unregister</code>	Unregister a variable
<code>hdfql_variable_unregister_all</code>	Unregister all variables
<code>hdfql_variable_get_number</code>	Get number of a variable
<code>hdfql_variable_get_data_type</code>	Get data type of a variable
<code>hdfql_variable_get_count</code>	Get number of elements (i.e. result set size) stored in a variable
<code>hdfql_variable_get_size</code>	Get size (in bytes) of a variable
<code>hdfql_variable_get_dimension_count</code>	Get number of dimensions of a variable
<code>hdfql_variable_get_dimension</code>	Get size of a certain dimension of a variable
<code>hdfql_mpi_get_size</code>	Get number (i.e. size) of processes associated to the MPI communicator
<code>hdfql_mpi_get_rank</code>	Get number (i.e. rank) of the calling process associated to the MPI communicator

Table 5.8 – HDFqI functions in C

In C++, the prefix “hdfql\_” of the name of functions (defined in C) is replaced by the namespace “HDFqL” and its underscores (\_) are discarded. The remainder of the name of functions follows the lower camel-case convention. The following table lists a subset of HDFqL functions as defined in C and details how these are defined/can be used in C++.

HDFqL Function in C	Corresponding Definition in C++
hdfql_execute	HDFqL::execute
hdfql_cursor_next	HDFqL::cursorNext
hdfql_cursor_get_tinyint	HDFqL::cursorGetTinyInt
hdfql_cursor_get_unsigned_int	HDFqL::cursorGetUnsignedInt
hdfql_subcursor_get_big_int	HDFqL::subcursorGetBigInt
hdfql_variable_get_number	HDFqL::variableGetNumber

Table 5.9 – HDFqL functions in C and their corresponding definitions in C++

In Java, the prefix “hdfql\_” of the name of functions (defined in C) is replaced by the class “HDFqL” and its underscores (\_) are discarded. The remainder of the name of functions follows the lower camel-case convention. The following table lists a subset of HDFqL functions as defined in C and details how these are defined/can be used in Java.

HDFqL Function in C	Corresponding Definition in Java
hdfql_execute	HDFqL.execute
hdfql_cursor_next	HDFqL.cursorNext
hdfql_cursor_get_tinyint	HDFqL.cursorGetTinyInt
hdfql_cursor_get_unsigned_int	HDFqL.cursorGetUnsignedInt
hdfql_subcursor_get_big_int	HDFqL.subcursorGetBigInt
hdfql_variable_get_number	HDFqL.variableGetNumber

Table 5.10 – HDFqL functions in C and their corresponding definitions in Java

In Python, the prefix “hdfqI\_” of the name of functions (defined in C) is replaced by the class “HDFqI”. The remainder of the name of functions remains exactly the same. The following table lists a subset of HDFqI functions as defined in C and details how these are defined/can be used in Python.

HDFqI Function in C	Corresponding Definition in Python
hdfqI_execute	HDFqI.execute
hdfqI_cursor_next	HDFqI.cursor_next
hdfqI_cursor_get_tinyint	HDFqI.cursor_get_tinyint
hdfqI_cursor_get_unsigned_int	HDFqI.cursor_get_unsigned_int
hdfqI_subcursor_get_big_int	HDFqI.subcursor_get_big_int
hdfqI_variable_get_number	HDFqI.variable_get_number

Table 5.11 – HDFqI functions in C and their corresponding definitions in Python

In C#, the prefix “hdfqI\_” of the name of functions (defined in C) is replaced by the class “HDFqI” and its underscores (\_) are discarded. The remainder of the name of functions follows the upper camel-case convention. The following table lists a subset of HDFqI functions as defined in C and details how these are defined/can be used in C#.

HDFqI Function in C	Corresponding Definition in C#
hdfqI_execute	HDFqI.Execute
hdfqI_cursor_next	HDFqI.CursorNext
hdfqI_cursor_get_tinyint	HDFqI.CursorGetTinyInt
hdfqI_cursor_get_unsigned_int	HDFqI.CursorGetUnsignedInt
hdfqI_subcursor_get_big_int	HDFqI.SubcursorGetBigInt
hdfqI_variable_get_number	HDFqI.VariableGetNumber

Table 5.12 – HDFqI functions in C and their corresponding definitions in C#

In Fortran, the name of functions is the same as in C and can be written using any case. The following table lists a subset of HDFqI functions as defined in C and details how these are defined/can be used in Fortran.

HDFqI Function in C	Corresponding Definition in Fortran
hdfqI_execute	hdfqI_execute
hdfqI_cursor_next	hdfqI_cursor_next
hdfqI_cursor_get_tinyint	hdfqI_cursor_get_tinyint
hdfqI_cursor_get_unsigned_int	hdfqI_cursor_get_unsigned_int
hdfqI_subcursor_get_big_int	hdfqI_subcursor_get_big_int
hdfqI_variable_get_number	hdfqI_variable_get_number

Table 5.13 – HDFqI functions in C and their corresponding definitions in Fortran

In R, the name of functions is the same as in C. The following table lists a subset of HDFqI functions as defined in C and details how these are defined/can be used in R.

HDFqI Function in C	Corresponding Definition in R
hdfqI_execute	hdfqI_execute
hdfqI_cursor_next	hdfqI_cursor_next
hdfqI_cursor_get_tinyint	hdfqI_cursor_get_tinyint
hdfqI_cursor_get_unsigned_int	hdfqI_cursor_get_unsigned_int
hdfqI_subcursor_get_big_int	hdfqI_subcursor_get_big_int
hdfqI_variable_get_number	hdfqI_variable_get_number

Table 5.14 – HDFqI functions in C and their corresponding definitions in R

## 5.2.1 HDFQL\_EXECUTE

### Syntax

```
int hdfql_execute(const char *script)
```

### Description

Execute a script named *script*. A script can be composed of one or more operations – in case of multiple operations these can either be separated with a semicolon (;) or an end of line (EOL) terminator. In HDFql, operations are case insensitive meaning that, for example, operation “SHOW DATASET” is equivalent to “show dataset” or any other case variation. If a certain operation raises an error, any subsequent operations within *script* are not executed. Please refer to [Table 6.2](#) for a complete enumeration of HDFql operations.

### Parameter(s)

*script* – string containing one or more operations to execute. Multiple operations are either separated with a semicolon (;) or an end of line (EOL) terminator.

### Return

int – depending on the success in executing *script*, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_PARSE](#), [HDFQL\\_ERROR\\_NOT\\_FOUND](#), [HDFQL\\_ERROR\\_NO\\_ACCESS](#), [HDFQL\\_ERROR\\_NOT\\_OPEN](#), [HDFQL\\_ERROR\\_INVALID\\_FILE](#), [HDFQL\\_ERROR\\_NOT\\_SUPPORTED](#), [HDFQL\\_ERROR\\_NOT\\_ENOUGH\\_SPACE](#), [HDFQL\\_ERROR\\_NOT\\_ENOUGH\\_MEMORY](#), [HDFQL\\_ERROR\\_ALREADY\\_EXISTS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_FULL](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#), [HDFQL\\_ERROR\\_AFTER\\_LAST](#), [HDFQL\\_ERROR\\_OUTSIDE\\_LIMIT](#), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#), [HDFQL\\_ERROR\\_UNEXPECTED\\_TYPE](#), [HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#), [HDFQL\\_ERROR\\_UNEXPECTED\\_STORAGE\\_TYPE](#), [HDFQL\\_ERROR\\_DANGLING\\_LINK](#), [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#), [HDFQL\\_ERROR\\_INVALID\\_REGULAR\\_EXPRESSION](#) or [HDFQL\\_ERROR\\_UNKNOWN](#).

### Example(s)

```
// declare variable
int status;

// execute script (composed of only one operation - i.e. SHOW USE FILE)
status = hdfql_execute("SHOW USE FILE");

// display message about the status of executed script (i.e. successful or not)
```

```
if (status == HDFQL_SUCCESS)
    printf("Execution was successful\n");
else
    printf("Execution was not successful and returned status is %d\n", status);
```

```
// execute script (composed of two operations - i.e. USE FILE my_file.h5 and SHOW)
hdfql_execute("USE FILE my_file.h5 ; SHOW");
```

## 5.2.2 HDFQL\_EXECUTE\_GET\_STATUS

### Syntax

```
int hdfql_execute_get_status(void)
```

### Description

Get status of the last executed operation. In other words, this function returns the status of the last call of [hdfql\\_execute](#).

### Parameter(s)

None

### Return

int – depending on the success of the last executed operation, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_PARSE](#), [HDFQL\\_ERROR\\_NOT\\_FOUND](#), [HDFQL\\_ERROR\\_NO\\_ACCESS](#), [HDFQL\\_ERROR\\_NOT\\_OPEN](#), [HDFQL\\_ERROR\\_INVALID\\_FILE](#), [HDFQL\\_ERROR\\_NOT\\_SUPPORTED](#), [HDFQL\\_ERROR\\_NOT\\_ENOUGH\\_SPACE](#), [HDFQL\\_ERROR\\_NOT\\_ENOUGH\\_MEMORY](#), [HDFQL\\_ERROR\\_ALREADY\\_EXISTS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_FULL](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#), [HDFQL\\_ERROR\\_AFTER\\_LAST](#), [HDFQL\\_ERROR\\_OUTSIDE\\_LIMIT](#), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#), [HDFQL\\_ERROR\\_UNEXPECTED\\_TYPE](#), [HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#), [HDFQL\\_ERROR\\_UNEXPECTED\\_STORAGE\\_TYPE](#), [HDFQL\\_ERROR\\_DANGLING\\_LINK](#), [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#), [HDFQL\\_ERROR\\_INVALID\\_REGULAR\\_EXPRESSION](#) or [HDFQL\\_ERROR\\_UNKNOWN](#).

### Example(s)

```
// declare variable
int status;
```

```
// execute script (composed of only one operation - i.e. SHOW USE DIRECTORY)
hdfql_execute("SHOW USE DIRECTORY");

// get status of last executed script (i.e. SHOW USE DIRECTORY)
status = hdfql_execute_get_status();

// display message about the status of last executed script (i.e. successful or not)
if (status == HDFQL_SUCCESS)
    printf("Execution was successful\n");
else
    printf("Execution was not successful and returned status is %d\n", status);
```

### 5.2.3 HDFQL\_ERROR\_GET\_LINE

#### Syntax

int hdfql\_error\_get\_line(void)

#### Description

Get error line of the last executed operation. In other words, this function returns the number of the line (in the script) where an error was raised during the last call of [hdfql\\_execute](#). The first line in the script is designated as number one (1).

#### Parameter(s)

None

#### Return

int – number of the line (in the script) where an error has occurred during the last executed operation. If the last executed operation was successful, the number of the line will be [HDFQL\\_UNDEFINED](#).

#### Example(s)

```
// execute script (composed of only one operation - i.e. CREATE FILE my_file.h5 - which is
syntactically correct)
hdfql_execute("CREATE FILE my_file.h5");

// display number of the line where an error occurred during the last executed operation
```



```
(should be -1 - i.e. HDFQL_UNDEFINED)
printf("Error line number is %d\n", hdfql_error_get_line());

// execute script (composed of only one operation - i.e. CREATE FILEX my_file.h5 - which is
syntactically incorrect due to a typo in "FILEX")
hdfql_execute("CREATE FILEX my_file.h5");

// display number of the line where an error occurred during the last executed operation
(should be 1)
printf("Error line number is %d\n", hdfql_error_get_line());
```

## 5.2.4 HDFQL\_ERROR\_GET\_POSITION

### Syntax

```
int hdfql_error_get_position(void)
```

### Description

Get error position of the last executed operation. In other words, this function returns the position in the line where an error was raised during the last call of [hdfql\\_execute](#). The first position in the line is designated as number one (1).

### Parameter(s)

None

### Return

int – position in the line where an error has occurred during the last executed operation. If the last executed operation was successful, the position in the line will be [HDFQL\\_UNDEFINED](#).

### Example(s)

```
// execute script (composed of only one operation - i.e. CREATE FILE my_file.h5 - which is
syntactically correct)
hdfql_execute("CREATE FILE my_file.h5");

// display position in the line where an error occurred during the last executed operation
(should be -1 - i.e. HDFQL_UNDEFINED)
printf("Error position is %d\n", hdfql_error_get_position());
```

```
// execute script (composed of only one operation - i.e. CREATE FILEX my_file.h5 - which is
syntactically incorrect due to a typo in "FILEX")
hdfql_execute("CREATE FILEX my_file.h5");

// display position in the line where an error occurred during the last executed operation
(should be 8)
printf("Error position is %d\n", hdfql_error_get_position());
```

## 5.2.5 HDFQL\_ERROR\_GET\_MESSAGE

### Syntax

```
char *hdfql_error_get_message(void)
```

### Description

Get error message of the last executed operation. In other words, this function returns the message of the error that was raised during the last call of [hdfql\\_execute](#).

### Parameter(s)

None

### Return

char \* – pointer to the message of an error that has occurred during the last executed operation. If the last executed operation was successful, the pointer will be NULL.

### Example(s)

```
// execute script (composed of only one operation - i.e. CREATE FILE my_file.h5 - which is
syntactically correct)
hdfql_execute("CREATE FILE my_file.h5");

// display message of an error that occurred during the last executed operation (should be
"NULL")
printf("%s\n", hdfql_error_get_message());

// execute script (composed of only one operation - i.e. CREATE FILEX my_file.h5 - which is
```

```
syntactically incorrect due to a typo in "FILEX")
hdfql_execute("CREATE FILEX my_file.h5");

// display message of an error that occurred during the last executed operation (should be
"Unknown token "FILEX"")
printf("%s\n", hdfql_error_get_message());
```

## 5.2.6 HDFQL\_CURSOR\_INITIALIZE

### Syntax

```
int hdfql_cursor_initialize(HDFQL_CURSOR *cursor)
```

### Description

Initialize a cursor named *cursor* for subsequent use. Before a new cursor is used for the first time, it should always be initialized (otherwise unexpected errors may arise such as a segmentation fault). The initialization of a cursor sets its data type attribute to undefined ([HDFQL\\_UNDEFINED](#)), its current element to NULL, and resets its count and position attributes to zero and minus one respectively, making it ready for usage. Of note, the process of initializing a cursor is only required in C and performed once, while in other programming languages supported by HDFqI – namely C++, Java, Python, C#, Fortran and R – such initialization is redundant (in other words, it is not needed) as it is done automatically when creating a cursor.

### Parameter(s)

*cursor* – pointer to a cursor (previously declared) to initialize with default values. If the pointer is NULL (in C), the cursor in use is initialized instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqI wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is initialized instead).

### Return

int – depending on the success in initializing *cursor*, it can either be [HDFQL\\_SUCCESS](#) or [HDFQL\\_ERROR\\_UNKNOWN](#).

### Example(s)

```
// create a cursor named "my_cursor"
HDFQL_CURSOR my_cursor;
```

```
// initialize cursor "my_cursor"
hdfql_cursor_initialize (&my_cursor);

// use cursor "my_cursor"
hdfql_cursor_use (&my_cursor);

// display number of elements in cursor "my_cursor" (should be 0)
printf("Number of elements in cursor is %d\n", hdfql_cursor_get_count (NULL));
```

## 5.2.7 HDFQL\_CURSOR\_USE

### Syntax

```
int hdfql_cursor_use(const HDFQL_CURSOR *cursor)
```

### Description

Set the cursor named *cursor* as the one to be used for storing results of operations.

### Parameter(s)

*cursor* – pointer to a cursor to use for storing the result of operations. If the pointer is NULL (in C), the HDFqL default cursor is used instead (i.e. equivalent of calling the function [hdfql\\_cursor\\_use\\_default](#)). The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively.

### Return

int – depending on the success in using *cursor*, it can either be [HDFQL\\_SUCCESS](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

### Example(s)

```
// create a cursor named "my_cursor"
HDFQL_CURSOR my_cursor;

// use cursor "my_cursor"
hdfql_cursor_use (&my_cursor);

// initialize cursor "my_cursor"
hdfql_cursor_initialize (NULL);
```

```
// display data type of cursor "my_cursor" (should be -1 - i.e. HDFQL_UNDEFINED)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type(NULL));

// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// display (again) data type of cursor "my_cursor" (should be 2097152 - i.e. HDFQL_VARCHAR)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type(NULL));

// use HDFql default cursor
hdfql_cursor_use(NULL);

// display data type of HDFql default cursor (should be -1 - i.e. HDFQL_UNDEFINED)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type(NULL));
```

## 5.2.8 HDFQL\_CURSOR\_USE\_DEFAULT

### Syntax

```
int hdfql_cursor_use_default(void)
```

### Description

Set HDFql default cursor as the one to be used for storing results of operations.

### Parameter(s)

None

### Return

int – depending on the success in using HDFql default cursor, it can either be [HDFQL\\_SUCCESS](#) or [HDFQL\\_ERROR\\_UNKNOWN](#).

### Example(s)

```
// create a cursor named "my_cursor"
HDFQL_CURSOR my_cursor;

// initialize cursor "my_cursor"
```

```
hdfql_cursor_initialize (&my_cursor);

// use cursor "my_cursor"
hdfql_cursor_use (&my_cursor);

// display data type of cursor "my_cursor" (should be -1 - i.e. HDFQL_UNDEFINED)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type (NULL));

// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// display (again) data type of cursor "my_cursor" (should be 2097152 - i.e. HDFQL_VARCHAR)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type (NULL));

// use HDFql default cursor
hdfql_cursor_use_default ();

// display data type of HDFql default cursor (should be -1 - i.e. HDFQL_UNDEFINED)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type (NULL));
```

## 5.2.9 HDFQL\_CURSOR\_CLEAR

### Syntax

```
int hdfql_cursor_clear(HDFQL_CURSOR *cursor)
```

### Description

Clear (i.e. empty) a cursor named *cursor*. Specifically, this function removes all elements (i.e. result set) stored in the cursor, specifies its data type attribute to undefined ([HDFQL\\_UNDEFINED](#)), changes its current element to NULL, and resets its count and position attributes to zero and minus one respectively.

### Parameter(s)

*cursor* – pointer to a cursor to clear (i.e. empty). If the pointer is NULL (in C), the cursor in use is cleared instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is cleared instead).

## **Return**

int – depending on the success in clearing *cursor*, it can either be `HDFQL_SUCCESS` or `HDFQL_ERROR_NOT_REGISTERED`.

## **Example(s)**

```
// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// display number of elements in the cursor in use (should be 1)
printf("Number of elements in cursor is %d\n", hdfql_cursor_get_count(NULL));

// clear the cursor in use
hdfql_cursor_clear(NULL);

// display (again) number of elements in the cursor in use (should be 0)
printf("Number of elements in cursor is %d\n", hdfql_cursor_get_count(NULL));
```

## **5.2.10 HDFQL\_CURSOR\_CLONE**

### **Syntax**

```
int hdfql_cursor_clone(const HDFQL_CURSOR *cursor_original, HDFQL_CURSOR *cursor_clone)
```

### **Description**

Clone (i.e. duplicate) a cursor named *cursor\_original* into another one named *cursor\_clone*. In other words, *cursor\_clone* will be an exact copy of *cursor\_original*, meaning that it will have the same data type, count and position values, store the same result set, and have the same current element as the original cursor.

### **Parameter(s)**

*cursor\_original* – pointer to a cursor to clone. If the pointer is NULL (in C), the cursor in use is the one to be cloned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is the one to be cloned instead).

*cursor\_clone* – pointer to the cursor that will be a clone (i.e. duplicate) of the original cursor.

## **Return**

int – depending on the success in cloning *cursor\_original* into *cursor\_clone*, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_NOT\\_ENOUGH\\_MEMORY](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

## **Example(s)**

```
// create a cursor named "my_cursor"
HDFQL_CURSOR my_cursor;

// initialize cursor "my_cursor"
hdfql_cursor_initialize(&my_cursor);

// show (i.e. get) current working directory and populate cursor in use (i.e. HDFql default
// cursor) with it
hdfql_execute("SHOW USE DIRECTORY");

// clone the cursor in use (i.e. HDFql default cursor) into the cursor "my_cursor"
hdfql_cursor_clone(NULL, &my_cursor);

// use cursor "my_cursor"
hdfql_cursor_use(&my_cursor);

// display number of elements in the cursor in use (should be 1)
printf("Number of elements in cursor is %d\n", hdfql_cursor_get_count(NULL));
```

### **5.2.11 HDFQL\_CURSOR\_GET\_DATA\_TYPE**

#### **Syntax**

```
int hdfql_cursor_get_data_type(const HDFQL_CURSOR *cursor)
```

#### **Description**

Get the data type of a cursor named *cursor*. If the cursor has never been populated or has been initialized or cleared, the returned data type is undefined ([HDFQL\\_UNDEFINED](#)). Please refer to [Table 6.3](#) for a complete enumeration of HDFql data types.



## **Parameter(s)**

*cursor* – pointer to a cursor to get its data type. If the pointer is NULL (in C), the data type of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqI wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the data type of the cursor in use is returned instead).

## **Return**

int – depending on the data type of the cursor or its state (i.e. whether it has never been populated or has been initialized or cleared), it can either be `HDFQL_TINYINT`, `HDFQL_UNSIGNED_TINYINT`, `HDFQL_SMALLINT`, `HDFQL_UNSIGNED_SMALLINT`, `HDFQL_INT`, `HDFQL_UNSIGNED_INT`, `HDFQL_BIGINT`, `HDFQL_UNSIGNED_BIGINT`, `HDFQL_FLOAT`, `HDFQL_DOUBLE`, `HDFQL_CHAR`, `HDFQL_VARTINYINT`, `HDFQL_UNSIGNED_VARTINYINT`, `HDFQL_VARSMALLINT`, `HDFQL_UNSIGNED_VARSMALLINT`, `HDFQL_VARINT`, `HDFQL_UNSIGNED_VARINT`, `HDFQL_VARBIGINT`, `HDFQL_UNSIGNED_VARBIGINT`, `HDFQL_VARFLOAT`, `HDFQL_VARDOUBLE`, `HDFQL_VARCHAR`, `HDFQL_OPAQUE` or `HDFQL_UNDEFINED`.

## **Example(s)**

```
// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// display data type of the cursor in use (should be 2097152 - i.e. HDFQL_VARCHAR)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type(NULL));

// clear the cursor in use
hdfql_cursor_clear(NULL);

// display (again) data type of the cursor in use (should be -1 - i.e. HDFQL_UNDEFINED)
printf("Data type of cursor is %d\n", hdfql_cursor_get_data_type(NULL));
```

## **5.2.12 HDFQL\_CURSOR\_GET\_COUNT**

### **Syntax**

```
int hdfql_cursor_get_count(const HDFQL_CURSOR *cursor)
```

## **Description**

Get the number of elements (i.e. result set size) stored in a cursor named *cursor*. If the result set stores data from a dataset or attribute that does not have a dimension (i.e. if it is scalar), the returned number of elements is one. Otherwise, if the result set stores data from a dataset or attribute that has dimensions, the returned number of elements equals the multiplication of all its dimensions' sizes (e.g. if a cursor stores a result set of two dimensions of size 10x3, the number of elements is 30). If the cursor has never been populated or has been initialized or cleared, the returned number of elements is zero.

## **Parameter(s)**

*cursor* – pointer to a cursor to get its number of elements (i.e. result set size). If the pointer is NULL (in C), the number of elements of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the number of elements of the cursor in use is returned instead).

## **Return**

int – number of elements (i.e. result set size) stored in the cursor.

## **Example(s)**

```
// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// display number of elements in the cursor in use (should be 1)
printf("Number of elements in cursor is %d\n", hdfql_cursor_get_count(NULL));

// clear the cursor in use
hdfql_cursor_clear(NULL);

// display (again) number of elements in the cursor in use (should be 0)
printf("Number of elements in cursor is %d\n", hdfql_cursor_get_count(NULL));
```

## **5.2.13 HDFQL\_SUBCURSOR\_GET\_COUNT**

### **Syntax**

```
int hdfql_subcursor_get_count(const HDFQL_CURSOR *cursor)
```

## **Description**

Get the number of elements (i.e. result subset size) stored in the subcursor in use. If the cursor that the subcursor belongs to has never been populated or has been initialized or cleared, the returned number of elements is zero.

## **Parameter(s)**

*cursor* – pointer to a cursor to get the number of elements (i.e. result subset size) stored in the subcursor in use. If the pointer is NULL (in C), the number of elements of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the number of elements of the subcursor of the cursor in use is returned instead).

## **Return**

int – number of elements (i.e. result subset size) stored in the subcursor.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// display number of elements in the cursor in use (should be 4 - i.e. 2x2)
printf("Number of elements in cursor is %d\n", hdfql_cursor_get_count(NULL));

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display number of elements in the subcursor in use (should be 3)
printf("Number of elements in subcursor is %d\n", hdfql_subcursor_get_count(NULL));

// move the cursor in use to next position within the result set (i.e. second position)
hdfql_cursor_next(NULL);

// display number of elements in the subcursor in use (should be 1)
```

```
printf("Number of elements in subcursor is %d\n", hdfql_subcursor_get_count(NULL));
```

## 5.2.14 HDFQL\_CURSOR\_GET\_POSITION

### Syntax

```
int hdfql_cursor_get_position(const HDFQL_CURSOR *cursor)
```

### Description

Get current position of a cursor named *cursor* within the result set. The first element of the result set is at position zero, while the last element is located at the position returned by [hdfql\\_cursor\\_get\\_count](#) - 1. If the cursor has never been populated or has been initialized or cleared, or in case the result set is empty, the returned current position is minus one. If the cursor was moved before the first element or after the last element, the returned current position is minus one or the number of elements in the result set, respectively.

### Parameter(s)

*cursor* – pointer to a cursor to get its current position within the result set. If the pointer is NULL (in C), the current position of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current position of the cursor in use is returned instead).

### Return

int – current position of the cursor within the result set.

### Example(s)

```
// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// display position of the cursor in use within the result set (should be 0)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));

// clear the cursor in use
```

```
hdfql_cursor_clear(NULL);

// display (again) position of the cursor in use within the result set (should be -1)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));
```

## 5.2.15 HDFQL\_SUBCURSOR\_GET\_POSITION

### Syntax

```
int hdfql_subcursor_get_position(const HDFQL_CURSOR *cursor)
```

### Description

Get current position of the subcursor in use within the result subset. The first element of the result subset is at position zero, while the last element is located at the position returned by `hdfql_subcursor_get_count` - 1. If the cursor that the subcursor belongs to has never been populated or has been initialized or cleared, or in case the result subset is empty, the returned current position is minus one. If the subcursor was moved before the first element or after the last element, the returned current position is minus one or the number of elements in the result subset, respectively.

### Parameter(s)

*cursor* – pointer to a cursor to get the current position of the subcursor in use within the result subset. If the pointer is NULL (in C), the current position of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current position of the subcursor of the cursor in use is returned instead).

### Return

int – current position of the subcursor within the result subset.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");
```

```
// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// display position of the subcursor in use within the result subset (should be -1)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));

// move the subcursor in use to the next position within the result subset (two times)
hdfql_subcursor_next(NULL);
hdfql_subcursor_next(NULL);

// display (again) position of the subcursor in use within the result subset (should be 1)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));
```

## 5.2.16 HDFQL\_CURSOR\_FIRST

### Syntax

```
int hdfql_cursor_first(HDFQL_CURSOR *cursor)
```

### Description

Move a cursor named *cursor* to the first position within the result set. In other words, the cursor will point to the first element of the result set and its position is set to zero. If the result set is empty, an error is returned and its position remains unchanged (i.e. remains minus one).

### Parameter(s)

*cursor* – pointer to a cursor to move to the first position within the result set. If the pointer is NULL (in C), the cursor in use is moved to the first position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is moved to the first position instead).

### Return

int – depending on the success in moving the cursor to the first position within the result set, it can either be [HDFQL\\_SUCCESS](#) or [HDFQL\\_ERROR\\_EMPTY](#).

## **Example(s)**

```
// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// display position of the cursor in use within the result set (should be -1)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// display (again) position of the cursor in use within the result set (should be 0)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));
```

## **5.2.17 HDFQL\_SUBCURSOR\_FIRST**

### **Syntax**

```
int hdfql_subcursor_first(HDFQL_CURSOR *cursor)
```

### **Description**

Move the subcursor in use to the first position within the result subset. In other words, the subcursor will point to the first element of the result subset and its position is set to zero. If the result subset is empty, an error is returned and its position remains unchanged (i.e. remains minus one).

### **Parameter(s)**

*cursor* – pointer to a cursor to move the subcursor in use to the first position within the result subset. If the pointer is NULL (in C), the subcursor of the cursor in use is moved to the first position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the subcursor of the cursor in use is moved to the first position instead).

### **Return**

int – depending on the success in moving the subcursor to the first position within the result subset, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// display position of the subcursor in use within the result subset (should be -1)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));

// move the subcursor in use to the first position within the result subset
hdfql_subcursor_first(NULL);

// display (again) position of the subcursor in use within the result subset (should be 0)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));
```

## **5.2.18 HDFQL\_CURSOR\_LAST**

### **Syntax**

```
int hdfql_cursor_last(HDFQL_CURSOR *cursor)
```

### **Description**

Move a cursor named *cursor* to the last position within the result set. In other words, the cursor will point to the last element of the result set and its position is set to the value returned by [hdfql\\_cursor\\_get\\_count](#) - 1. If the result set is empty, an error is returned and its position remains unchanged (i.e. remains minus one).

### **Parameter(s)**

*cursor* – pointer to a cursor to move to the last position within the result set. If the pointer is NULL (in C), the cursor in use is moved to the last position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL



wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is moved to the last position instead).

## **Return**

int – depending on the success in moving the cursor to the last position within the result set, it can either be [HDFQL\\_SUCCESS](#) or [HDFQL\\_ERROR\\_EMPTY](#).

## **Example(s)**

```
// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// display position of the cursor in use within the result set (should be -1)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));

// move the cursor in use to the last position within the result set
hdfql_cursor_last(NULL);

// display position of the cursor in use within the result set (should be 0)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));
```

## **5.2.19 HDFQL\_SUBCURSOR\_LAST**

### **Syntax**

```
int hdfql_subcursor_last(HDFQL_CURSOR *cursor)
```

### **Description**

Move the subcursor in use to the last position within the result subset. In other words, the subcursor will point to the last element of the result subset and its position is set to the value returned by [hdfql\\_subcursor\\_get\\_count](#) - 1. If the result subset is empty, an error is returned and its position remains unchanged (i.e. remains minus one).

### **Parameter(s)**

*cursor* – pointer to a cursor to move the subcursor in use to the last position within the result subset. If the pointer is NULL (in C), the subcursor of the cursor in use is moved to the last position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is

mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the subcursor of the cursor in use is moved to the last position instead).

## **Return**

int – depending on the success in moving the subcursor to the last position within the result subset, it can either be `HDFQL_SUCCESS`, `HDFQL_ERROR_EMPTY`, `HDFQL_ERROR_BEFORE_FIRST` or `HDFQL_ERROR_AFTER_LAST`.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// display position of subcursor in use within the result subset (should be -1)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));

// move the subcursor in use to the last position within the result subset
hdfql_subcursor_last(NULL);

// display (again) position of subcursor in use within the result subset (should be 2)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));
```

## **5.2.20 HDFQL\_CURSOR\_NEXT**

### **Syntax**

```
int hdfql_cursor_next(HDFQL_CURSOR *cursor)
```

## **Description**

Move a cursor named *cursor* one position forward from its current position. In other words, the cursor will point to the next element of the result set and its position is incremented by one. If the result set is empty or the cursor is in the last position, an error is returned and its position remains unchanged (i.e. remains minus one) or is set to the value returned by [hdfql\\_cursor\\_get\\_count](#), respectively.

## **Parameter(s)**

*cursor* – pointer to a cursor to move one position forward from its current position. If the pointer is NULL (in C), the cursor in use is moved one position forward from its current position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is moved one position forward from its current position instead).

## **Return**

int – depending on the success in moving the cursor one position forward from its current position, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

## **Example(s)**

```
// show (i.e. get) current working directory
hdfql_execute("SHOW USE DIRECTORY");

// move the cursor in use to the next position within the result set
hdfql_cursor_next(NULL);

// display position of cursor within the result set (should be 0)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));
```

## **5.2.21 HDFQL\_SUBCURSOR\_NEXT**

### **Syntax**

```
int hdfql_subcursor_next(HDFQL_CURSOR *cursor)
```

## **Description**

Move the subcursor in use one position forward from its current position. In other words, the subcursor will point to the next element of the result subset and its position is incremented by one. If the result subset is empty or the subcursor is in the last position, an error is returned and its position remains unchanged (i.e. remains minus one) or is set to the value returned by [hdfql\\_subcursor\\_get\\_count](#), respectively

## **Parameter(s)**

*cursor* – pointer to a cursor to move the subcursor in use one position forward from its current position. If the pointer is NULL (in C), the subcursor of the cursor in use is moved one position forward from its current position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the subcursor of the cursor in use is moved one position forward from its current position instead).

## **Return**

int – depending on the success in moving the subcursor one position forward from its current position, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// display position of subcursor in use within the result subset (should be -1)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));

// move the subcursor in use to the next position within the result subset (two times)
hdfql_subcursor_next(NULL);
hdfql_subcursor_next(NULL);
```

```
// display (again) position of subcursor in use within the result subset (should be 1)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));
```

## 5.2.22 HDFQL\_CURSOR\_PREVIOUS

### Syntax

```
int hdfql_cursor_previous(HDFQL_CURSOR *cursor)
```

### Description

Move a cursor named *cursor* one position backward from its current position. In other words, the cursor will point to the previous element of the result set and its position is decremented by one. If the result set is empty or the cursor is in the first position, an error is returned and its position remains unchanged (i.e. remains minus one) or is set to minus one, respectively.

### Parameter(s)

*cursor* – pointer to a cursor to move one position backward from its current position. If the pointer is NULL (in C), the cursor in use is moved one position backward from its current position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is moved one position backward from its current position instead).

### Return

int – depending on the success in moving the cursor one position backward from its current position, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#) or [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#).

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type float of two dimensions (size 2x10)
hdfql_execute("CREATE DATASET my_dataset AS FLOAT(2, 10)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the last position within the result set
```

```
hdfql_cursor_last(NULL);

// move the cursor in use to the previous position within the result set
hdfql_cursor_previous(NULL);

// display position of cursor in use within the result set (should be 18 - i.e. 2x10-1-1)
printf("Position of cursor is %d\n", hdfql_cursor_get_position(NULL));
```

## 5.2.23 HDFQL\_SUBCURSOR\_PREVIOUS

### Syntax

```
int hdfql_subcursor_previous(HDFQL_CURSOR *cursor)
```

### Description

Move the subcursor in use one position backward from its current position. In other words, the subcursor will point to the previous element of the result subset and its position is decremented by one. If the result subset is empty or the subcursor is in the first position, an error is returned and its position remains unchanged (i.e. remains minus one) or is set to minus one, respectively.

### Parameter(s)

*cursor* – pointer to a cursor to move the subcursor in use one position backward from its current position. If the pointer is NULL (in C), the subcursor of the cursor in use is moved one position backward from its current position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the subcursor of the cursor in use is moved one position backward from its current position instead).

### Return

int – depending on the success in moving the subcursor one position backward from its current position, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");
```

```
// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// move the subcursor in use to the last position within the result subset
hdfql_subcursor_last(NULL);

// move the subcursor in use to the previous position within the result subset (two times)
hdfql_subcursor_previous(NULL);
hdfql_subcursor_previous(NULL);

// display position of the subcursor within the result subset (should be 0 - i.e. 3-1-1-1)
printf("Position of subcursor is %d\n", hdfql_subcursor_get_position(NULL));
```

## 5.2.24 HDFQL\_CURSOR\_ABSOLUTE

### Syntax

```
int hdfql_cursor_absolute(HDFQL_CURSOR *cursor, int position)
```

### Description

Move a cursor named *cursor* to an absolute position *position* within the result set. The first element of the result set is at position zero, while the last element is located at the position returned by [hdfql\\_cursor\\_get\\_count](#) - 1. An attempt to move the cursor before the first element will return an error and set the position of the cursor to minus one, while an attempt to move the cursor after the last element will return an error and set the position of the cursor to number of elements in the result set.

### Parameter(s)

*cursor* – pointer to a cursor to move to an absolute position within the result set. If the pointer is NULL (in C), the cursor in use is moved to an absolute position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is moved to an absolute position instead).

*position* – absolute position to which to move the cursor. If *position* is positive, the cursor will position itself with reference to the beginning of the result set. If *position* is negative, the cursor will position itself with reference to the end of the result set.

## **Return**

int – depending on the success in moving the cursor to an absolute position within the result set, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

## **Example(s)**

```
// create five HDF5 groups named "g1", "g2", "g3", "g4" and "g5"
hdfql_execute("CREATE GROUP g1, g2, g3, g4, g5");

// show (i.e. get) all existing groups and populate cursor in use with these (should be "g1",
// "g2", "g3", "g4", "g5")
hdfql_execute("SHOW GROUP");

// move the cursor in use to absolute position 2 within the result set
hdfql_cursor_absolute(NULL, 2);

// display current element of the cursor in use within the result set (should be "g3")
printf("Current element of cursor is %s", hdfql_cursor_get_char(NULL));

// move the cursor in use to absolute position -2 within the result set
hdfql_cursor_absolute(NULL, -2);

// display current element of the cursor in use within the result set (should be "g4")
printf("Current element of cursor is %s", hdfql_cursor_get_char(NULL));
```

## **5.2.25 HDFQL\_SUBCURSOR\_ABSOLUTE**

### **Syntax**

```
int hdfql_subcursor_absolute(HDFQL_CURSOR *cursor, int position)
```

### **Description**

Move the subcursor in use to an absolute position *position* within the result subset. The first element of the result subset is at position zero, while the last element is located at the position returned by [hdfql\\_subcursor\\_get\\_count](#) - 1. An attempt



to move the subcursor before the first element will return an error and set the position of the subcursor to minus one, while an attempt to move the subcursor after the last element will return an error and set the position of the subcursor to number of elements in the result subset.

### **Parameter(s)**

*cursor* – pointer to a cursor to move the subcursor in use to an absolute position within the result subset. If the pointer is NULL (in C), the subcursor of the cursor in use is moved to an absolute position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the subcursor of the cursor in use is moved to an absolute position instead).

*position* – absolute position to which to move the subcursor. If *position* is positive, the subcursor will position itself with reference to the beginning of the result subset. If *position* is negative, the subcursor will position itself with reference to the end of the result subset.

### **Return**

int – depending on the success in moving the subcursor to an absolute position within the result subset, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

### **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// move the subcursor in use to absolute position 2 within the result subset
hdfql_subcursor_absolute(NULL, 2);

// display current element of the subcursor in use within the result subset (should be 5)
printf("Current element of subcursor is %d", hdfql_subcursor_get_int(NULL));
```

```
// move the subcursor in use to absolute position -2 within the result subset
hdfql_subcursor_absolute(NULL, -2);

// display current element of the subcursor in use within the result subset (should be 8)
printf("Current element of subcursor is %d", hdfql_subcursor_get_int(NULL));
```

## 5.2.26 HDFQL\_CURSOR\_RELATIVE

### Syntax

```
int hdfql_cursor_relative(HDFQL_CURSOR *cursor, int position)
```

### Description

Move a cursor named *cursor* to a relative position *position* with respect to its current position. The first element of the result set is at position zero, while the last element is located at the position returned by [hdfql\\_cursor\\_get\\_count](#) - 1. An attempt to move the cursor before the first element will return an error and set the position of the cursor to minus one, while an attempt to move the cursor after the last element will return an error and set the position of the cursor to number of elements in the result set.

### Parameter(s)

*cursor* – pointer to a cursor to move to a relative position with respect to its current position. If the pointer is NULL (in C), the cursor in use is moved to a relative position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the cursor in use is moved to a relative position instead).

*position* – relative position to which to move the cursor. If *position* is positive, the cursor will go forward in the result set relative to its current position. If *position* is negative, the cursor will go backward in the result set relative to its current position.

### Return

int – depending on the success in moving the cursor to a relative position with respect to its current position, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

## **Example(s)**

```
// create five HDF5 groups named "g1", "g2", "g3", "g4" and "g5"
hdfql_execute("CREATE GROUP g1, g2, g3, g4, g5");

// show (i.e. get) all existing groups and populate cursor in use with these (should be "g1",
// "g2", "g3", "g4", "g5")
hdfql_execute("SHOW GROUP");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// move the cursor in use to relative position 2 within the result set
hdfql_cursor_relative(NULL, 2);

// display current element of the cursor within the result set (should be "g3")
printf("Current element of cursor is %s", hdfql_cursor_get_char(NULL));

// move the cursor in use to relative position -2 within the result set
hdfql_cursor_relative(NULL, -2);

// display current element of the cursor within the result set (should be "g1")
printf("Current element of cursor is %s", hdfql_cursor_get_char(NULL));
```

## **5.2.27 HDFQL\_SUBCURSOR\_RELATIVE**

### **Syntax**

```
int hdfql_subcursor_relative(HDFQL_CURSOR *cursor, int position)
```

### **Description**

Move the subcursor in use to a relative position *position* with respect to its current position. The first element of the result subset is at position zero, while the last element is located at the position returned by [hdfql\\_subcursor\\_get\\_count](#) - 1. An attempt to move the subcursor before the first element will return an error and set the position of the subcursor to minus one, while an attempt to move the subcursor after the last element will return an error and set the position of the subcursor to number of elements in the result set.

## **Parameter(s)**

*cursor* – pointer to a cursor to move the subcursor in use to a relative position with respect to its current position. If the pointer is NULL (in C), the subcursor of the cursor in use is moved to a relative position instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the subcursor of the cursor in use is moved to a relative position instead).

*position* – relative position to which to move the subcursor. If *position* is positive, the subcursor will go forward in the result set relative to its current position. If *position* is negative, the subcursor will go backward in the result set relative to its current position.

## **Return**

int – depending on the success in moving the subcursor to a relative position with respect to its current position, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of two dimensions
(size 2x2)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(2, 2)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((7, 8, 5), (9)), ((6, 1, 2, 3), (4, 0)))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// move the subcursor in use to the first position within the result subset
hdfql_subcursor_first(NULL);

// move the subcursor in use to relative position 2 within the result subset
hdfql_subcursor_relative(NULL, 2);

// display current element of the subcursor in use within the result subset (should be 5)
printf("Current element of subcursor is %d", hdfql_subcursor_get_int(NULL));

// move the subcursor in use to relative position -1 within the result subset
```

```
hdfql_subcursor_relative(NULL, -1);

// display current element of the subcursor in use within the result subset (should be 8)
printf("Current element of subcursor is %d", hdfql_subcursor_get_int(NULL));
```

## 5.2.28 HDFQL\_CURSOR\_GET\_SIZE

### Syntax

```
int hdfql_cursor_get_size(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element size (in bytes) of a cursor named *cursor*. If the result set is empty or the cursor is located before or after the first or last element of the result set, an error is returned instead.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element size (in bytes). If the pointer is NULL (in C), the current element size of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element size of the cursor in use is returned instead).

### Return

int – depending on the success in getting the current element size (in bytes) of the cursor, it can either be  $\geq 0$  (i.e. the size itself), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

### Example(s)

```
// create an HDF5 group named "my_group"
hdfql_execute("CREATE GROUP my_group");

// show (i.e. get) all existing groups and populate cursor in use with these (should be
"my_group")
hdfql_execute("SHOW GROUP");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);
```

```
// display current element size (in bytes) of the cursor in use within the result set (should be 8 - i.e. 8x1)
printf("Current element size (in bytes) of cursor is %d\n", hdfql_cursor_get_size(NULL));
```

## 5.2.29 HDFQL\_SUBCURSOR\_GET\_SIZE

### Syntax

```
int hdfql_subcursor_get_size(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element size (in bytes) of the subcursor in use. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, an error is returned instead.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element size (in bytes) of the subcursor in use. If the pointer is NULL (in C), the current element size of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element size of the subcursor of the cursor in use is returned instead).

### Return

int – depending on the success in getting the current element size (in bytes) of the subcursor, it can either be  $\geq 0$  (i.e. the size itself), [HDFQL\\_ERROR\\_EMPTY](#), [HDFQL\\_ERROR\\_BEFORE\\_FIRST](#) or [HDFQL\\_ERROR\\_AFTER\\_LAST](#).

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length float of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARFLOAT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES ((5.5, 2.2), (8.1), (4.9, 3.4, 5.6))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
```

```
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to the first position within the result set
hdfql_cursor_first(NULL);

// move the subcursor in use to the first position within the result subset
hdfql_subcursor_first(NULL);

// display current element size (in bytes) of the subcursor within the result subset (should be
4 - i.e. 4x1)
printf("Current element size (in bytes) of subcursor is %d\n", hdfql_subcursor_get_size(NULL));
```

## 5.2.30 HDFQL\_CURSOR\_GET

### Syntax

```
void *hdfql_cursor_get(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element of a cursor named *cursor* as a generic (typeless) pointer. It is up to the programmer to interpret the returned pointer according to their needs. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL. Of note, this function is only available in C and C++ (the remaining HDFqI wrappers – namely Java, Python, C#, Fortran and R – do not support it).

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as a generic (typeless) pointer. If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in the C++ HDFqI wrapper is NULL. While in C *cursor* is mandatory, in C++ it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

void \* – generic (typeless) pointer to the current element of the cursor. If there is no current element, the pointer is NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type float of one dimension (size 3)
```

```
hdfql_execute("CREATE DATASET my_dataset AS FLOAT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES (5.5, 8.1, 4.9)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a float (should be 5.5)
printf("Current element of cursor is %f\n", (float *) hdfql_cursor_get(NULL));
```

## 5.2.31 HDFQL\_SUBCURSOR\_GET

### Syntax

void \*hdfql\_subcursor\_get(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of the subcursor in use as a generic (typeless) pointer. It is up to the programmer to interpret the returned pointer according to their needs. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL. Of note, this function is only available in C and C++ (the remaining HDFqL wrappers – namely Java, Python, C#, Fortran and R – do not support it).

### Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as a generic (typeless) pointer. If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in the C++ HDFqL wrapper is NULL. While in C *cursor* is mandatory, in C++ it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

### Return

void \* – generic (typeless) pointer to the current element of the subcursor. If there is no current element, the pointer is NULL.



## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length float of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARFLOAT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES ((5.5, 2.2), (8.1), (4.9, 3.4, 5.6))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a float (should be 5.5)
printf("Current element of subcursor is %f\n", (float *) hdfql_subcursor_get(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a float (should be 2.2)
printf("Current element of subcursor is %f\n", (float *) hdfql_subcursor_get(NULL));
```

## **5.2.32 HDFQL\_CURSOR\_GET\_TINYINT**

### **Syntax**

```
char *hdfql_cursor_get_tinyint(const HDFQL_CURSOR *cursor)
```

### **Description**

Get the current element of a cursor named *cursor* as a [TINYINT](#). In other words, the current element is interpreted as a “char” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

## **Parameter(s)**

*cursor* – pointer to a cursor to get the current element as a [TINYINT](#). If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

## **Return**

char \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type char of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS TINYINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(12, 34, 23)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a char (should be 12)
printf("Current element of cursor is %d\n", *hdfql_cursor_get_tinyint(NULL));
```

## **5.2.33 HDFQL\_SUBCURSOR\_GET\_TINYINT**

### **Syntax**

char \*hdfql\_subcursor\_get\_tinyint(const HDFQL\_CURSOR \**cursor*)

### **Description**

Get the current element of the subcursor in use as a [TINYINT](#). In other words, the current element is interpreted as a “char” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

## **Parameter(s)**

*cursor* – pointer to a cursor to get the current element of the subcursor in use as a [TINYINT](#). If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

## **Return**

char \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length char of one dimension
(size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARTINYINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a char (should be 5)
printf("Current element of cursor is %d\n", *hdfql_cursor_get_tinyint(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a char (should be 5)
printf("Current element of subcursor is %d\n", *hdfql_subcursor_get_tinyint(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a char (should be 2)
printf("Current element of subcursor is %d\n", *hdfql_subcursor_get_tinyint(NULL));
```

## 5.2.34 HDFQL\_CURSOR\_GET\_UNSIGNED\_TINYINT

### Syntax

unsigned char \*hdfql\_cursor\_get\_unsigned\_tinyint(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of a cursor named *cursor* as an **UNSIGNED TINYINT**. In other words, the current element is interpreted as an “unsigned char” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as a **UNSIGNED TINYINT**. If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

unsigned char \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type unsigned char of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED TINYINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(12, 34, 23)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned char (should be 12)
printf("Current element of cursor is %u\n", *hdfql_cursor_get_unsigned_tinyint(NULL));
```

## 5.2.35 HDFQL\_SUBCURSOR\_GET\_UNSIGNED\_TINYINT

### Syntax

unsigned char \*hdfql\_subcursor\_get\_unsigned\_tinyint(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of the subcursor in use as an **UNSIGNED TINYINT**. In other words, the current element is interpreted as an “unsigned char” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as an **UNSIGNED TINYINT**. If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

### Return

unsigned char \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length unsigned char of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED VARTINYINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned char (should be 5)
printf("Current element of cursor is %u\n", *hdfql_cursor_get_unsigned_tinyint(NULL));
```

```
// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned char (should be 5)
printf("Current element of subcursor is %u\n", *hdfql_subcursor_get_unsigned_tinyint(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned char (should be 2)
printf("Current element of subcursor is %u\n", *hdfql_subcursor_get_unsigned_tinyint(NULL));
```

## 5.2.36 HDFQL\_CURSOR\_GET\_SMALLINT

### Syntax

short \*hdfql\_cursor\_get\_smallint(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of a cursor named *cursor* as a **SMALLINT**. In other words, the current element is interpreted as a “short” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as a **SMALLINT**. If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

short \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type short of one dimension (size 3)
```

```
hdfql_execute("CREATE DATASET my_dataset AS SMALLINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(12, 34, 23)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a short (should be 12)
printf("Current element of cursor is %d\n", *hdfql_cursor_get_smallint(NULL));
```

## 5.2.37 HDFQL\_SUBCURSOR\_GET\_SMALLINT

### Syntax

short \*hdfql\_subcursor\_get\_smallint(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of the subcursor in use as a [SMALLINT](#). In other words, the current element is interpreted as a "short" C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as a [SMALLINT](#). If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

### Return

short \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

## Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length short of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARSMALLINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a short (should be 5)
printf("Current element of cursor is %d\n", *hdfql_cursor_get_smallint(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a short (should be 5)
printf("Current element of subcursor is %d\n", *hdfql_subcursor_get_smallint(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a short (should be 2)
printf("Current element of subcursor is %d\n", *hdfql_subcursor_get_smallint(NULL));
```

## 5.2.38 HDFQL\_CURSOR\_GET\_UNSIGNED\_SMALLINT

### Syntax

unsigned short \*hdfql\_cursor\_get\_unsigned\_smallint(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of a cursor named *cursor* as an [UNSIGNED SMALLINT](#). In other words, the current element is interpreted as an “unsigned short” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.



## **Parameter(s)**

*cursor* – pointer to a cursor to get the current element as an **UNSIGNED SMALLINT**. If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

## **Return**

unsigned short \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type unsigned short of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED SMALLINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(12, 34, 23)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned short (should be 12)
printf("Current element of cursor is %u\n", *hdfql_cursor_get_unsigned_smallint(NULL));
```

## **5.2.39 HDFQL\_SUBCURSOR\_GET\_UNSIGNED\_SMALLINT**

### **Syntax**

unsigned short \*hdfql\_subcursor\_get\_unsigned\_smallint(const HDFQL\_CURSOR \**cursor*)

### **Description**

Get the current element of the subcursor in use as an **UNSIGNED SMALLINT**. In other words, the current element is interpreted as an “unsigned short” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

## **Parameter(s)**

*cursor* – pointer to a cursor to get the current element of the subcursor in use as an **UNSIGNED SMALLINT**. If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

## **Return**

unsigned short \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length unsigned short of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED VARSMALLINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned short (should be 5)
printf("Current element of cursor is %u\n", *hdfql_cursor_get_unsigned_smallint(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned short (should be 5)
printf("Current element of subcursor is %u\n", *hdfql_subcursor_get_unsigned_smallint(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned short (should be 2)
printf("Current element of subcursor is %u\n", *hdfql_subcursor_get_unsigned_smallint(NULL));
```

## 5.2.40 HDFQL\_CURSOR\_GET\_INT

### Syntax

```
int *hdfql_cursor_get_int(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element of a cursor named *cursor* as an [INT](#). In other words, the current element is interpreted as an “int” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as an [INT](#). If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

int \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type int of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS INT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(12, 34, 23)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned short (should be 12)
printf("Current element of cursor is %d\n", *hdfql_cursor_get_int(NULL));
```

## 5.2.41 HDFQL\_SUBCURSOR\_GET\_INT

### Syntax

```
int *hdfql_subcursor_get_int(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element of the subcursor in use as an [INT](#). In other words, the current element is interpreted as an “int” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as an [INT](#). If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

### Return

int \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length int of one dimension
(size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES ((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);
```

```
// display current element of the cursor in use as an int (should be 5)
printf("Current element of cursor is %d\n", *hdfql_cursor_get_int(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an int (should be 5)
printf("Current element of subcursor is %d\n", *hdfql_subcursor_get_int(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an int (should be 2)
printf("Current element of subcursor is %d\n", *hdfql_subcursor_get_int(NULL));
```

## 5.2.42 HDFQL\_CURSOR\_GET\_UNSIGNED\_INT

### Syntax

unsigned int \*hdfql\_cursor\_get\_unsigned\_int(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of a cursor named *cursor* as an **UNSIGNED INT**. In other words, the current element is interpreted as an “unsigned int” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as an **UNSIGNED INT**. If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

unsigned int \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type unsigned int of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED INT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES (12, 34, 23)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned int (should be 12)
printf("Current element of cursor is %u\n", *hdfql_cursor_get_unsigned_int(NULL));
```

## **5.2.43 HDFQL\_SUBCURSOR\_GET\_UNSIGNED\_INT**

### **Syntax**

unsigned int \*hdfql\_subcursor\_get\_unsigned\_int(const HDFQL\_CURSOR \*cursor)

### **Description**

Get the current element of the subcursor in use as an **UNSIGNED INT**. In other words, the current element is interpreted as an “unsigned int” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### **Parameter(s)**

*cursor* – pointer to a cursor to get the current element of the subcursor in use as an **UNSIGNED INT**. If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

## **Return**

unsigned int \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length unsigned int of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED VARINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES ((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned int (should be 5)
printf("Current element of cursor is %u\n", *hdfql_cursor_get_unsigned_int(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned int (should be 5)
printf("Current element of subcursor is %u\n", *hdfql_subcursor_get_unsigned_int(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned int (should be 2)
printf("Current element of subcursor is %u\n", *hdfql_subcursor_get_unsigned_int(NULL));
```

## **5.2.44 HDFQL\_CURSOR\_GET\_BIGINT**

### **Syntax**

long long \*hdfql\_cursor\_get\_bigint(const HDFQL\_CURSOR \*cursor)

## **Description**

Get the current element of a cursor named *cursor* as a **BIGINT**. In other words, the current element is interpreted as a “long long” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

## **Parameter(s)**

*cursor* – pointer to a cursor to get the current element as a **BIGINT**. If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

## **Return**

long long \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type long long of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS BIGINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(12, 34, 23)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a long long (should be 12)
printf("Current element of cursor is %lld\n", *hdfql_cursor_get_bigint(NULL));
```

## **5.2.45 HDFQL\_SUBCURSOR\_GET\_BIGINT**

### **Syntax**

long long \*hdfql\_subcursor\_get\_bigint(const HDFQL\_CURSOR \*cursor)



## Description

Get the current element of the subcursor in use as a **BIGINT**. In other words, the current element is interpreted as a “long long” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

## Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as a **BIGINT**. If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

## Return

long long \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

## Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length long long of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARBIGINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a long long (should be 5)
printf("Current element of cursor is %lld\n", *hdfql_cursor_get_bigint(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a long long (should be 5)
printf("Current element of subcursor is %lld\n", *hdfql_subcursor_get_bigint(NULL));
```

```
// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a long long (should be 2)
printf("Current element of subcursor is %lld\n", *hdfql_subcursor_get_bigint(NULL));
```

## 5.2.46 HDFQL\_CURSOR\_GET\_UNSIGNED\_BIGINT

### Syntax

unsigned long long \*hdfql\_cursor\_get\_unsigned\_bigint(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of a cursor named *cursor* as an **UNSIGNED BIGINT**. In other words, the current element is interpreted as an “unsigned long long” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as an **UNSIGNED BIGINT**. If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

unsigned long long \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type unsigned long long of one dimension
(size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED BIGINT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(12, 34, 23)");
```

```
// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned long long (should be 12)
printf("Current element of cursor is %llu\n", *hdfql_cursor_get_unsigned_bigint(NULL));
```

## 5.2.47 HDFQL\_SUBCURSOR\_GET\_UNSIGNED\_BIGINT

### Syntax

unsigned long long \*hdfql\_subcursor\_get\_unsigned\_bigint(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of the subcursor in use as an **UNSIGNED BIGINT**. In other words, the current element is interpreted as an “unsigned long long” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as an **UNSIGNED BIGINT**. If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

### Return

unsigned long long \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length unsigned long long of
one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS UNSIGNED VARBIGINT(3)");
```

```
// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((5, 2), (8), (4, 3, 9))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as an unsigned long long (should be 5)
printf("Current element of cursor is %llu\n", *hdfql_cursor_get_unsigned_bigint(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned long long (should be 5)
printf("Current element of subcursor is %llu\n", *hdfql_subcursor_get_unsigned_bigint(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as an unsigned long long (should be 2)
printf("Current element of subcursor is %llu\n", *hdfql_subcursor_get_unsigned_bigint(NULL));
```

## 5.2.48 HDFQL\_CURSOR\_GET\_FLOAT

### Syntax

float \*hdfql\_cursor\_get\_float(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of a cursor named *cursor* as a [FLOAT](#). In other words, the current element is interpreted as a “float” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as a [FLOAT](#). If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers

is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

## **Return**

float \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type float of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS FLOAT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(5.5, 8.1, 4.9)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a float (should be 5.5)
printf("Current element of cursor is %f\n", *hdfql_cursor_get_float(NULL));
```

## **5.2.49 HDFQL\_SUBCURSOR\_GET\_FLOAT**

### **Syntax**

float \*hdfql\_subcursor\_get\_float(const HDFQL\_CURSOR \*cursor)

### **Description**

Get the current element of the subcursor in use as a **FLOAT**. In other words, the current element is interpreted as a “float” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### **Parameter(s)**

*cursor* – pointer to a cursor to get the current element of the subcursor in use as a **FLOAT**. If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java,

Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

## **Return**

float \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length float of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARFLOAT(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES ((7.5, 3.1), (4.5), (4.9, 3.2, 9.7, 8.8))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a float (should be 7.5)
printf("Current element of cursor is %f\n", *hdfql_cursor_get_float(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a float (should be 7.5)
printf("Current element of subcursor is %f\n", *hdfql_subcursor_get_float(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a float (should be 3.1)
printf("Current element of subcursor is %f\n", *hdfql_subcursor_get_float(NULL));
```

## 5.2.50 HDFQL\_CURSOR\_GET\_DOUBLE

### Syntax

```
double *hdfql_cursor_get_double(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element of a cursor named *cursor* as a [DOUBLE](#). In other words, the current element is interpreted as a “double” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as a [DOUBLE](#). If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

double \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type double of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS DOUBLE(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES(5.5, 8.1, 4.9)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a double (should be 5.5)
printf("Current element of cursor is %f\n", *hdfql_cursor_get_double(NULL));
```

## 5.2.51 HDFQL\_SUBCURSOR\_GET\_DOUBLE

### Syntax

```
double *hdfql_subcursor_get_double(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element of the subcursor in use as a [DOUBLE](#). In other words, the current element is interpreted as a “double” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as a [DOUBLE](#). If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

### Return

double \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type variable-length double of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARDOUBLE(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES((7.5, 3.1), (4.5), (4.9, 3.2, 9.7, 8.8))");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a double (should be 7.5)
printf("Current element of cursor is %f\n", *hdfql_cursor_get_double(NULL));
```



```
// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a double (should be 7.5)
printf("Current element of subcursor is %f\n", *hdfql_subcursor_get_double(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a double (should be 3.1)
printf("Current element of subcursor is %f\n", *hdfql_subcursor_get_double(NULL));
```

## 5.2.52 HDFQL\_CURSOR\_GET\_CHAR

### Syntax

```
char *hdfql_cursor_get_char(const HDFQL_CURSOR *cursor)
```

### Description

Get the current element of a cursor named *cursor* as a [CHAR](#). In other words, the current element is interpreted as a “char” C data type and returned as a pointer of such type. If the result set is empty or the cursor is located before or after the first or last element of the result set, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element as a [CHAR](#). If the pointer is NULL (in C), the current element of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFqL wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the cursor in use is returned instead).

### Return

char \* – pointer to the current element of the cursor. If there is no current element, the pointer will be NULL.

### Example(s)

```
// create an HDF5 dataset named "my_dataset" of data type char of one dimension (size 3)
```

```
hdfql_execute("CREATE DATASET my_dataset AS CHAR(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES (Red)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a char (should be "Red")
printf("Current element of cursor is %s\n", hdfql_cursor_get_char(NULL));
```

## 5.2.53 HDFQL\_SUBCURSOR\_GET\_CHAR

### Syntax

char \*hdfql\_subcursor\_get\_char(const HDFQL\_CURSOR \*cursor)

### Description

Get the current element of the subcursor in use as a [CHAR](#). In other words, the current element is interpreted as a “char” C data type and returned as a pointer of such type. If the result subset is empty or the subcursor is located before or after the first or last element of the result subset, the returned element is NULL.

### Parameter(s)

*cursor* – pointer to a cursor to get the current element of the subcursor in use as a [CHAR](#). If the pointer is NULL (in C), the current element of the subcursor of the cursor in use is returned instead. The equivalent of a NULL pointer in C++, Java, Python, C#, Fortran and R HDFql wrappers is NULL, null, None, null, 0 and NULL, respectively. While in C *cursor* is mandatory, in C++, Java, Python, C#, Fortran and R it is optional (when not provided, the current element of the subcursor of the cursor in use is returned instead).

### Return

char \* – pointer to the current element of the subcursor. If there is no current element, the pointer will be NULL.

## **Example(s)**

```
// create an HDF5 dataset named "my_dataset" of data type variable-length char of one dimension
(size 3)
hdfql_execute("CREATE DATASET my_dataset AS VARCHAR(3)");

// insert (i.e. write) values into dataset "my_dataset"
hdfql_execute("INSERT INTO my_dataset VALUES (Red, Green, Blue)");

// select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset");

// move the cursor in use to next position within the result set (i.e. first position)
hdfql_cursor_next(NULL);

// display current element of the cursor in use as a char (should be "Red")
printf("Current element of cursor is %s\n", hdfql_cursor_get_char(NULL));

// move the subcursor in use to next position within the result subset (i.e. first position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a char (should be "Red")
printf("Current element of subcursor is %s\n", hdfql_subcursor_get_char(NULL));

// move the subcursor in use to next position within the result subset (i.e. second position)
hdfql_subcursor_next(NULL);

// display current element of the subcursor in use as a char (should be "Green")
printf("Current element of subcursor is %s\n", hdfql_subcursor_get_char(NULL));
```

## **5.2.54 HDFQL\_VARIABLE\_REGISTER**

### **Syntax**

```
int hdfql_variable_register(const void *variable)
```

### **Description**

Register a variable named *variable* for subsequent use. In other words, for HDFql to be able to read or write from/to a user-defined variable it must first be registered. If the operation was successful, *variable* is registered and a number is assigned to it. This number – calculated by HDFql – starts with zero and is incremented by one every time a new variable is

registered. If *variable* is registered more than once, only one number is assigned to it (namely the number assigned upon the first registering). Of note, currently up to eight variables can be registered at any given time (trying to register more than this number will raise an [HDFQL\\_ERROR\\_FULL](#)). In C, C++ and Fortran any variable may be registered as long as HDFqI can properly read and write values from/to it by having direct access to the memory associated with these – otherwise unexpected errors may arise such as a segmentation fault. The following restrictions apply to other programming languages (supported by HDFqI):

- In Java, only a variable that is an array of “byte”, “short”, “int”, “long”, “float”, “double” or “String” data type or its corresponding wrapper class “Byte”, “Short”, “Integer”, “Long”, “Float” or “Double” may be registered. Any attempt to register a variable that is not an array or of the data type/wrapper class previously enumerated will return an error ([HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#)).
- In Python, only a variable that is a NumPy array of “int8”, “uint8”, “int16”, “uint16”, “int32”, “uint32”, “int64”, “uint64”, “float32”, “float64”, “Ssize” or “ubyte” data type may be registered. Any attempt to register a variable that is not a NumPy array or of the data type previously enumerated will return an error ([HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#)). Please refer to <http://www.numpy.org> for additional information.
- In C#, only a variable that is an array of “SByte”, “Byte”, “Int16”, “UInt16”, “Int32”, “UInt32”, “Int64”, “UInt64”, “Single”, “Double” or “String” data type or its alias “sbyte”, “byte”, “short”, “ushort”, “int”, “uint”, “long”, “ulong”, “float”, “double” or “string” may be registered. Any attempt to register a variable that is not an array or of the data type/alias previously enumerated will return an error ([HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#)).
- In R, only a variable that is a vector, matrix or array of “integer”, “integer64” (through package bit64), “numeric”, “double”, “character” or “raw” data type may be registered. Any attempt to register a variable that is not a vector, matrix or array or of the data type previously enumerated will return an error ([HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#)).

An important aspect to remember when working with a variable is that it should not change address from the moment it has been registered until used in the intended operation (e.g. [SELECT](#)) or function (e.g. [HDFQL\\_VARIABLE\\_GET\\_NUMBER](#)), as HDFqI will not be able to identify the variable. In this case, the operation or function will raise an error ([HDFQL\\_ERROR\\_NOT\\_REGISTERED](#)). In case a variable needs to change its address (for whatever the reason), first unregister it via the function [hdfql\\_variable\\_unregister](#), change its address, and register it again. In general, it is advisable to register a variable just before executing the HDFqI operation or function which employs it, and to unregister it as soon as it is no longer used (this is especially relevant in C# where variables are pinned when registered and thus cannot be moved by the Garbage Collector).

## **Parameter(s)**

*variable* – variable to register for subsequent use.

## **Return**

int – depending on the success in registering the variable for subsequent use, it can either be  $\geq 0$  (i.e. the number assigned to the variable when successfully registered), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#), [HDFQL\\_ERROR\\_FULL](#) or [HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#).

## **Example(s)**

```
// declare variables
char script[1024];
short data[3];
int number;

// create an HDF5 dataset named "my_dataset" of data type short of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS SMALLINT(3)");

// populate variable "data" with certain values
data[0] = 21;
data[1] = 18;
data[2] = 75;

// register variable "data" for subsequent use (by HDFq1)
number = hdfql_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset "my_dataset"
sprintf(script, "INSERT INTO my_dataset VALUES FROM MEMORY %d", number);

// execute script
hdfql_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfql_variable_unregister(data);
```

## 5.2.55 HDFQL\_VARIABLE\_TRANSIENT\_REGISTER

### Syntax

```
int hdfql_variable_transient_register(const void *variable)
```

### Description

Register a variable named *variable* in a transient way for subsequent use. This function is similar to [hdfql\\_variable\\_register](#), except that after the execution of a script (via the function [hdfql\\_execute](#)) which uses *variable*, *variable* is automatically unregistered (by HDFq1) thus alleviating the programmer from doing it.

### Parameter(s)

*variable* – variable to register in a transient way for subsequent use.

### Return

int – depending on the success in registering the variable in a transient way for subsequent use, it can either be  $\geq 0$  (i.e. the number assigned to the variable when successfully registered), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#), [HDFQL\\_ERROR\\_FULL](#) or [HDFQL\\_ERROR\\_UNEXPECTED\\_DATA\\_TYPE](#).

### Example(s)

```
// declare variables
char script[1024];
short data[3];
int number;

// create an HDF5 dataset named "my_dataset" of data type short of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS SMALLINT(3)");

// populate variable "data" with certain values
data[0] = 21;
data[1] = 18;
data[2] = 75;

// register variable "data" in a transient way for subsequent use (by HDFq1)
number = hdfql_variable_transient_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset "my_dataset"
sprintf(script, "INSERT INTO my_dataset VALUES FROM MEMORY %d", number);
```

```
// execute script (variable "data" is automatically unregistered immediately after the
execution of the script - i.e. there is no need to explicitly unregister the variable)
hdfql_execute(script);
```

## 5.2.56 HDFQL\_VARIABLE\_UNREGISTER

### Syntax

```
int hdfql_variable_unregister(const void *variable)
```

### Description

Unregister a variable named *variable*. In other words, HDFql will free up any memory that may have been allocated to manage the variable as well as the number assigned to it (the number may then be assigned to a new variable registered subsequently). In general, it is advisable to unregister a variable as soon as it is no longer used by HDFql (this is especially relevant in C# as variables are unpinned when unregistered and thus may again be moved by the Garbage Collector). If *variable* has never been registered or has already been unregistered, an error is returned.

### Parameter(s)

*variable* – variable to unregister.

### Return

int – depending on the success in unregistering the variable, it can either be [HDFQL\\_SUCCESS](#), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

### Example(s)

```
// declare variables
char script[1024];
short data[3];
int number;

// create an HDF5 dataset named "my_dataset" of data type short of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset AS SMALLINT(3)");

// populate variable "data" with certain values
```

```
data[0] = 21;
data[1] = 18;
data[2] = 75;

// register variable "data" for subsequent use (by HDFq1)
number = hdfql_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset "my_dataset"
sprintf(script, "INSERT INTO my_dataset VALUES FROM MEMORY %d", number);

// execute script
hdfql_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfql_variable_unregister(data);
```

## 5.2.57 HDFQL\_VARIABLE\_UNREGISTER\_ALL

### Syntax

int hdfql\_variable\_unregister\_all(void)

### Description

Unregister all the variables that may have been registered previously. In other words, HDFq1 will free up any memory that may have been allocated to manage the variables as well as the numbers assigned to them (the numbers may then be assigned to new variables registered subsequently). In general, it is advisable to unregister variables as soon as they are no longer used by HDFq1 (this is especially relevant in C# as variables are unpinned when unregistered and thus may again be moved by the Garbage Collector).

### Parameter(s)

None

### Return

int – depending on the success in unregistering all the variables that may have been registered previously, it can either be [HDFQL\\_SUCCESS](#) or [HDFQL\\_ERROR\\_UNKNOWN](#).



## **Example(s)**

```
// declare variables
short data0[3];
float data1[5];

// register variable "data0" for subsequent use (by HDFq1)
hdfq1_variable_register(data0);

// register variable "data1" for subsequent use (by HDFq1)
hdfq1_variable_register(data1);

// display number of variable "data0" (should be 0)
printf("Number of variable is %d\n", hdfq1_variable_get_number(data0));

// display number of variable "data1" (should be 1)
printf("Number of variable is %d\n", hdfq1_variable_get_number(data1));

// unregister all the variables (i.e. variables "data0" and "data1") as they are no longer
used/needed (by HDFq1)
hdfq1_variable_unregister_all();
```

## **5.2.58 HDFQL\_VARIABLE\_GET\_NUMBER**

### **Syntax**

```
int hdfq1_variable_get_number(const void *variable)
```

### **Description**

Get the number of a variable named *variable*. This refers to the number that was calculated by HDFq1 and assigned to the variable upon registering it with the function [hdfq1\\_variable\\_register](#). If *variable* has never been registered or has been unregistered, an error is returned.

### **Parameter(s)**

*variable* – variable to get the number (calculated by HDFq1) assigned to it.

## **Return**

int – depending on the success in getting the number assigned to the variable, it can either be  $\geq 0$  (i.e. the number assigned to the variable), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

## **Example(s)**

```
// declare variables
short data0[3];
float data1[5];

// register variable "data0" for subsequent use (by HDFq1)
hdfql_variable_register(data0);

// register variable "data1" for subsequent use (by HDFq1)
hdfql_variable_register(data1);

// display number of variable "data0" (should be 0)
printf("Number of variable is %d\n", hdfql_variable_get_number(data0));

// display number of variable "data1" (should be 1)
printf("Number of variable is %d\n", hdfql_variable_get_number(data1));
```

## **5.2.59 HDFQL\_VARIABLE\_GET\_DATA\_TYPE**

### **Syntax**

```
int hdfql_variable_get_data_type(const void *variable)
```

### **Description**

Get the data type of a variable named *variable*. This function should help the programmer to better handle the content stored in *variable*. The data type refers to the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation redirected into memory – and not the data type of *variable* declared in the program. If *variable* has never been registered, populated (through the redirection of the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation into memory), or in case it has been unregistered, the returned data type is undefined ([HDFQL\\_UNDEFINED](#)). Please refer to [Table 6.3](#) for a complete enumeration of HDFq1 data types.

## **Parameter(s)**

*variable* – variable to get its data type.

## **Return**

int – depending on the success in getting the data type of the variable, it can either be [HDFQL\\_TINYINT](#), [HDFQL\\_UNSIGNED\\_TINYINT](#), [HDFQL\\_SMALLINT](#), [HDFQL\\_UNSIGNED\\_SMALLINT](#), [HDFQL\\_INT](#), [HDFQL\\_UNSIGNED\\_INT](#), [HDFQL\\_BIGINT](#), [HDFQL\\_UNSIGNED\\_BIGINT](#), [HDFQL\\_FLOAT](#), [HDFQL\\_DOUBLE](#), [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#), [HDFQL\\_VARCHAR](#), [HDFQL\\_OPAQUE](#), [HDFQL\\_BITFIELD](#), [HDFQL\\_ENUMERATION](#), [HDFQL\\_COMPOUND](#), [HDFQL\\_UNDEFINED](#), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

## **Example(s)**

```
// declare variables
char script[1024];
char data[1024];

// register variable "data" for subsequent use (by HDFql)
hdfql_variable_register(data);

// prepare script to show (i.e. get) current working directory and populate variable "data"
with it
sprintf(script, "SHOW USE DIRECTORY INTO MEMORY %d", hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// display data type of variable "data" (should be 2097152 - i.e. HDFQL_VARCHAR)
printf("Data type of variable is %d\n", hdfql_variable_get_data_type(data));
```

## **5.2.60 HDFQL\_VARIABLE\_GET\_COUNT**

### **Syntax**

```
int hdfql_variable_get_count(const void *variable)
```

## **Description**

Get the number of elements (i.e. result set size) stored in a variable named *variable*. This function should help the programmer to better handle the content stored in *variable*. If the result set stores data from a dataset or attribute that does not have a dimension (i.e. if it is scalar), the returned number of elements is one. Otherwise, if the result set stores data from a dataset or attribute that has dimensions, the returned number of elements equals the multiplication of all its dimensions' sizes (e.g. if a variable stores a result set of two dimensions of size 10x3, the number of elements is 30). Of note, in case a hyperslab or point selection is specified (in a [DATA QUERY LANGUAGE \(DQL\)](#) operation) the number of elements of the selection will be returned instead. If *variable* has never been populated (through the redirection of the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation into memory), the returned number of elements is zero.

## **Parameter(s)**

*variable* – variable to get its number of elements (i.e. result set size).

## **Return**

int – depending on the success in getting the number of elements of the variable, it can either be  $\geq 0$  (i.e. the number of elements), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

## **Example(s)**

```
// declare variables
char script[1024];
int data[5][3];

// create an HDF5 dataset named "my_dataset" of data type int of two dimensions (size 5x3)
hdfql_execute("CREATE DATASET my_dataset AS INT(5, 3)");

// register variable "data" for subsequent use (by HDFqL)
hdfql_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset INTO MEMORY %d", hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// display number of elements in variable "data" (should be 15 - i.e. 5x3)
```

```
printf("Number of elements in variable is %d\n", hdfql_variable_get_count(data));
```

## 5.2.61 HDFQL\_VARIABLE\_GET\_SIZE

### Syntax

```
int hdfql_variable_get_size(const void *variable)
```

### Description

Get the size (in bytes) of a variable named *variable*. This function should help the programmer to better handle the content stored in *variable*. The size (in bytes) refers to the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation redirected into memory – and not the size (in bytes) that *variable* has in the program. If *variable* has never been registered or has been unregistered, an error is returned. If *variable* has never been populated (through the redirection of the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation into memory), the returned size is zero. Please refer to [Table 6.3](#) for a complete enumeration of HDFq1 data types and their corresponding sizes (in bytes).

### Parameter(s)

*variable* – variable to get its size (in bytes).

### Return

int – depending on the success in getting the size (in bytes) of the variable, it can either be  $\geq 0$  (i.e. the size itself), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

### Example(s)

```
// declare variables
char script[1024];
int data[5][3];

// create an HDF5 dataset named "my_dataset" of data type int of two dimensions (size 5x3)
hdfql_execute("CREATE DATASET my_dataset AS INT(5, 3)");

// register variable "data" for subsequent use (by HDFq1)
hdfql_variable_register(data);
```

```
// prepare script to select (i.e. read) data from dataset "my_dataset" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset INTO MEMORY %d", hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// display size (in bytes) of variable "data" (should be 60 - i.e. 5x3x4)
printf("Size (in bytes) of variable is %d\n", hdfql_variable_get_size(data));
```

## 5.2.62 HDFQL\_VARIABLE\_GET\_DIMENSION\_COUNT

### Syntax

```
int hdfql_variable_get_dimension_count(const void *variable)
```

### Description

Get the number of dimensions of a variable named *variable*. This function should help the programmer to better handle the content stored in *variable*. The number of dimensions refers to the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation redirected into memory – and not the number of dimensions that *variable* has in the program. If *variable* has never been registered or has been unregistered, an error is returned. If *variable* has never been populated (through the redirection of the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation into memory), the returned number of dimensions is zero.

### Parameter(s)

*variable* – variable to get its number of dimensions.

### Return

int – depending on the success in getting the number of dimensions of the variable, it can either be  $\geq 0$  (i.e. the number of dimensions), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#) or [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#).

### Example(s)

```
// declare variables
char script[1024];
int data[5][3];
```

```
// create an HDF5 dataset named "my_dataset" of data type int of two dimensions (size 5x3)
hdfql_execute("CREATE DATASET my_dataset AS INT(5, 3)");

// register variable "data" for subsequent use (by HDFqL)
hdfql_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset INTO MEMORY %d", hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// display number of dimensions of variable "data" (should be 2)
printf("Number of dimensions in variable is %d\n", hdfql_variable_get_dimension_count(data));
```

## 5.2.63 HDFQL\_VARIABLE\_GET\_DIMENSION

### Syntax

int hdfql\_variable\_get\_dimension(const void \*variable, int index)

### Description

Get the size of a certain dimension specified in *index* of a variable named *variable*. This function should help the programmer to better handle the content stored in *variable*. The size of a certain dimension refers to the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation redirected into memory – and not the size of a certain dimension that *variable* has in the program. The index of the first dimension is zero (*index* must be between 0 and the value returned by [hdfql\\_variable\\_get\\_dimension\\_count](#) - 1). If *variable* has never been registered, populated (through the redirection of the result of a [DATA QUERY LANGUAGE \(DQL\)](#) or [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operation into memory), or in case it has been unregistered, an error is returned.

### Parameter(s)

*variable* – variable to get the size of one of its dimensions.

*index* – index of the dimension to get its size.

## **Return**

int – depending on the success in getting the size of a certain dimension of the variable, it can either be  $\geq 0$  (i.e. the size of a certain dimension itself), [HDFQL\\_ERROR\\_NO\\_ADDRESS](#), [HDFQL\\_ERROR\\_NOT\\_REGISTERED](#) or [HDFQL\\_ERROR\\_OUTSIDE\\_LIMIT](#).

## **Example(s)**

```
// declare variables
char script[1024];
int data[5][3];

// create an HDF5 dataset named "my_dataset" of data type int of two dimensions (size 5x3)
hdfql_execute("CREATE DATASET my_dataset AS INT(5, 3)");

// register variable "data" for subsequent use (by HDFql)
hdfql_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset INTO MEMORY %d", hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// display size of the first dimension of variable "data" (should be 5)
printf("Size of first dimension of variable is %d\n", hdfql_variable_get_dimension(0));

// display size of the second dimension of variable "data" (should be 3)
printf("Size of second dimension of variable is %d\n", hdfql_variable_get_dimension(1));
```

## **5.2.64 HDFQL\_MPI\_GET\_SIZE**

### **Syntax**

```
int hdfql_mpi_get_size(void)
```

### **Description**

Get the number (i.e. size) of processes associated to the default MPI communicator (MPI\_COMM\_WORLD). In other words, this function returns the number of MPI processes that are specified upon launching a program in parallel using



“mpiexec” (or a similar launcher). Of note, this function is basically a wrapper of the MPI function “MPI\_Comm\_size” (please refer to [https://www.mpich.org/static/docs/v3.2/www3/MPI\\_Comm\\_size.html](https://www.mpich.org/static/docs/v3.2/www3/MPI_Comm_size.html) or [https://www.open-mpi.org/doc/v2.1/man3/MPI\\_Comm\\_size.3.php](https://www.open-mpi.org/doc/v2.1/man3/MPI_Comm_size.3.php) for additional information in case the MPI library used is MPICH (or, alternatively, one of its ABI compatible derivative libraries) or Open MPI).

### **Parameter(s)**

None

### **Return**

int – depending on the success in getting the number of processes associated to the default MPI communicator (MPI\_COMM\_WORLD), it can either be  $\geq 1$  (i.e. the number of processes) or `HDFQL_UNDEFINED` (in case MPI itself was not initialized properly, in case of an HDFqL non MPI-based distribution, or if it was executed in Windows as HDFqL does not support the parallel HDF5 (PHDF5) library in this platform currently).

### **Example(s)**

```
// display number (i.e. size) of MPI processes (if the program is launched as, e.g., "mpiexec -  
n 5 my_program", the message "Number (i.e. size) of MPI processes is 5" will be displayed five  
times)  
printf("Number (i.e. size) of MPI processes is %d\n", hdfql_mpi_get_size());
```

## **5.2.65 HDFQL\_MPI\_GET\_RANK**

### **Syntax**

```
int hdfql_mpi_get_rank(void)
```

### **Description**

Get the number (i.e. rank) of the calling process associated to the default MPI communicator (MPI\_COMM\_WORLD). In other words, this function returns the number of the MPI process assigned to a particular instance of a program that was launched in parallel using “mpiexec” (or a similar launcher). Of note, this function is basically a wrapper of the MPI function “MPI\_Comm\_rank” (please refer to [https://www.mpich.org/static/docs/v3.2/www3/MPI\\_Comm\\_rank.html](https://www.mpich.org/static/docs/v3.2/www3/MPI_Comm_rank.html) or [https://www.open-mpi.org/doc/v2.1/man3/MPI\\_Comm\\_rank.3.php](https://www.open-mpi.org/doc/v2.1/man3/MPI_Comm_rank.3.php) for additional information in case the MPI library used is MPICH (or, alternatively, one of its ABI compatible derivative libraries) or Open MPI).

## **Parameters(s)**

None

## **Return**

int – depending on the success in getting the number (i.e. rank) of the calling process associated to the default MPI communicator (MPI\_COMM\_WORLD), it can either be  $\geq 0$  (i.e. the number of the calling process) or [HDFQL\\_UNDEFINED](#) (in case MPI itself was not initialized properly, in case of an HDFqI non MPI-based distribution, or if in Windows as HDFqI does not support the parallel HDF5 (PHDF5) library in this platform currently).

## **Example(s)**

```
// display number (i.e. rank) of the MPI process (if the program is launched as, e.g., "mpiexec  
-n 3 my_program", the message "Number (i.e. rank) of the MPI process is X" will be displayed  
three times where X is 0, 1 or 2 (not necessarily in this order))  
printf("Number (i.e. rank) of the MPI process is %d\n", hdfql_mpi_get_rank());
```

---

## 6. LANGUAGE

---

HDFql is a high-level language to manage HDF5 files in a simple and natural way. It was designed to be similar to SQL (wherever possible) so that its learning effort is kept at minimum while still providing great power and flexibility to the programmer. This chapter describes data types, post-processing options to further transform result sets, redirecting options to read/write data/result sets from/into disparate input/output sources, and operations (i.e. the language itself) available in HDFql. It also introduces text formatting conventions used throughout this chapter to describe HDFql operations (Table 6.1), and a summary of existing operations (Table 6.2). Before continuing, it is highly recommended to first read the HDF5 User's Guide available at [https://support.hdfgroup.org/HDF5/doc/UG/HDF5\\_Users\\_Guide.pdf](https://support.hdfgroup.org/HDF5/doc/UG/HDF5_Users_Guide.pdf) to facilitate the understanding of the current chapter.

Convention	Description	Example
<b>Bold</b>	Keyword that must be typed exactly as shown	<b>CREATE</b>
<i>Italic</i>	Value that the programmer must supply	<i>dataset_name</i>
Between brackets ([ ])	Optional keyword/value	<b>[DATASET]</b>
Between braces ({ })	Logical grouping of keywords/values (to ease understanding)	{ <b>[TRUNCATE]</b> <b>BINARY FILE</b> <i>file_name</i> }
Separated with a pipe ( )	Set of keywords/values from which one must be chosen	<b>GROUP   DATASET   ATTRIBUTE</b>
Asterisk (*)	Keyword/value that can be supplied several times	<i>group_name</i> [, <i>group_name</i> ]*

Table 6.1 – HDFql operations text formatting conventions

Operation	Description
<b>CREATE DIRECTORY</b>	Create a directory
<b>CREATE FILE</b>	Create an HDF5 file
<b>CREATE GROUP</b>	Create an HDF5 group

CREATE DATASET	Create an HDF5 dataset
CREATE ATTRIBUTE	Create an HDF5 attribute
CREATE [SOFT   HARD] LINK	Create an HDF5 soft or hard link
CREATE EXTERNAL LINK	Create an HDF5 external link
ALTER DIMENSION	Alter (i.e. change) dimensions of an existing HDF5 dataset
RENAME DIRECTORY	Rename (or move) an existing directory
RENAME FILE	Rename (or move) an existing file
RENAME [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK]	Rename (or move) an existing HDF5 group, dataset, attribute, (soft) link or external link
COPY FILE	Copy an existing file
COPY [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK]	Copy an existing HDF5 group, dataset, attribute, (soft) link or external link
DROP DIRECTORY	Drop (i.e. delete) an existing directory
DROP FILE	Drop (i.e. delete) an existing file
DROP [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK]	Drop (i.e. delete) an existing HDF5 group, dataset, attribute, (soft) link or external link
INSERT	Insert (i.e. write) data into an HDF5 dataset or attribute
SELECT	Select (i.e. read) data from an HDF5 dataset or attribute
SHOW FILE VALIDITY	Get validity of a file (i.e. whether it is a valid HDF5 file or not)
SHOW USE DIRECTORY	Get working directory currently in use
SHOW USE FILE	Get HDF5 file currently in use
SHOW ALL USE FILE	Get all HDF5 files in use (i.e. open)
SHOW USE GROUP	Get HDF5 group currently in use
SHOW [GROUP   DATASET   ATTRIBUTE   [SOFT] LINK   EXTERNAL LINK]	Get HDF5 objects (i.e. groups, datasets, attributes, (soft) links or external links) or check existence of an object
SHOW TYPE	Get type of an HDF5 object (i.e. group, dataset or attribute)
SHOW DATA TYPE	Get data type of an HDF5 dataset or attribute

SHOW MASK	Get (filter) mask of an HDF5 dataset
SHOW ENDIANNES	Get endianness of an HDF5 dataset or attribute
SHOW CHARSET	Get charset of an HDF5 dataset or attribute
SHOW STORAGE TYPE	Get storage type (layout) of an HDF5 dataset
SHOW STORAGE ALLOCATION	Get storage allocation of an HDF5 dataset
SHOW STORAGE DIMENSION	Get storage dimensions of an HDF5 dataset
SHOW DIMENSION	Get dimensions of an HDF5 dataset or attribute
SHOW ORDER	Get (creation) order strategy of an HDF5 group or dataset
SHOW TAG	Get tag of an HDF5 dataset or attribute
SHOW FILL TYPE	Get fill type of an HDF5 dataset
SHOW FILL VALUE	Get fill values of an HDF5 dataset
SHOW FILE SIZE	Get size (in bytes) of a file
SHOW [DATASET   ATTRIBUTE] SIZE	Get size (in bytes) of an HDF5 dataset or attribute
SHOW HDFQL VERSION	Get version of HDFq library
SHOW HDF5 VERSION	Get version of HDF5 library used by HDFq
SHOW PCRE VERSION	Get version of PCRE library used by HDFq
SHOW ZLIB VERSION	Get version of ZLIB library used by HDFq
SHOW MPI VERSION	Get version of MPI library used by HDFq
SHOW DIRECTORY	Get directory names within a directory
SHOW FILE	Get file names within a directory or check existence of a file
SHOW EXECUTE STATUS	Get status of the last executed operation
SHOW LIBRARY BOUNDS	Get library bound values for creating or opening HDF5 files
SHOW CACHE	Get cache parameters for accessing HDF5 files or datasets
SHOW ATOMIC	Get atomicity for accessing HDF5 files in an MPI environment
SHOW EXTERNAL LINK PREFIX	Get prefix to prepend to file names stored in HDF5 external links

SHOW FLUSH	Get status of the automatic flushing
SHOW DEBUG	Get status of the debug mechanism
USE DIRECTORY	Use a directory for subsequent operations
USE FILE	Use (i.e. open) an HDF5 file for subsequent operations
USE GROUP	Use (i.e. open) an HDF5 group for subsequent operations
FLUSH	Flush the entire virtual HDF5 file (global) or only the HDF5 file (local) currently in use
CLOSE FILE	Close HDF5 file currently in use
CLOSE ALL FILE	Close all HDF5 files in use
CLOSE GROUP	Close HDF5 group currently in use
SET LIBRARY BOUNDS	Set library bound values for creating and opening HDF5 files
SET CACHE	Set cache parameters for accessing HDF5 files or datasets
SET CACHE	Set atomicity for accessing HDF5 files in an MPI environment to enabled or disabled
SET EXTERNAL LINK PREFIX	Set prefix to prepend to file names stored in HDF5 external links
SET FLUSH	Set automatic flushing of the entire virtual HDF5 file or only the HDF5 file currently in use to enabled or disabled
SET DEBUG	Set debug mechanism to enabled or disabled

Table 6.2 – HDFql operations

## 6.1 DATA TYPES

A data type is a classification identifying one of various types of data such as integer, real or string, which determines the possible values for that type, the operations that can be done on values of that type, the meaning of the data, and the way values of that type can be stored. In other words, a data type is a classification of data that tells HDFql how the user

intends to use it. The following table summarizes all existing HDFq data types and how these map with the HDF5 data types<sup>1</sup>.

HDFq	HDF5	Range of Values
TINYINT	H5T_NATIVE_CHAR	-128 to 127 (1 byte)
UNSIGNED TINYINT	H5T_NATIVE_UCHAR	0 to 255 (1 byte)
SMALLINT	H5T_NATIVE_SHORT	-32,768 to 32,767 (2 bytes)
UNSIGNED SMALLINT	H5T_NATIVE_USHORT	0 to 65,535 (2 bytes)
INT	H5T_NATIVE_INT	-2,147,483,648 to 2,147,483,647 (4 bytes)
UNSIGNED INT	H5T_NATIVE_UINT	0 to 4,294,967,295 (4 bytes)
BIGINT	H5T_NATIVE_LLONG	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 (8 bytes)
UNSIGNED BIGINT	H5T_NATIVE_ULLONG	0 to 18,446,744,073,709,551,615 (8 bytes)
FLOAT	H5T_NATIVE_FLOAT	-3.4E + 38 to 3.4E + 38 (4 bytes)
DOUBLE	H5T_NATIVE_DOUBLE	-1.79E + 308 to 1.79E + 308 (8 bytes)
CHAR	H5T_C_S1	0 to 255 (1 byte)
VARTINYINT	H5T_NATIVE_CHAR	-128 to 127 (size * 1 byte)
UNSIGNED VARTINYINT	H5T_NATIVE_UCHAR	0 to 255 (size * 1 byte)

<sup>1</sup> For a detailed explanation of HDF5 data types please refer to [https://support.hdfgroup.org/HDF5/doc1.8/UG/HDF5\\_Users\\_Guide-ResponsiveHTML5/index.html#t=HDF5\\_Users\\_Guide/Datatypes/HDF5\\_Datatypes.htm](https://support.hdfgroup.org/HDF5/doc1.8/UG/HDF5_Users_Guide-ResponsiveHTML5/index.html#t=HDF5_Users_Guide/Datatypes/HDF5_Datatypes.htm).

<b>VARSMALLINT</b>	H5T_NATIVE_SHORT	-32,768 to 32,767 (size * 2 bytes)
<b>UNSIGNED VARSMALLINT</b>	H5T_NATIVE_USHORT	0 to 65,535 (size * 2 bytes)
<b>VARINT</b>	H5T_NATIVE_INT	-2,147,483,648 to 2,147,483,647 (size * 4 bytes)
<b>UNSIGNED VARINT</b>	H5T_NATIVE_UINT	0 to 4,294,967,295 (size * 4 bytes)
<b>VARBIGINT</b>	H5T_NATIVE_LLONG	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 (size * 8 bytes)
<b>UNSIGNED VARBIGINT</b>	H5T_NATIVE_ULLONG	0 to 18,446,744,073,709,551,615 (size * 8 bytes)
<b>VARFLOAT</b>	H5T_NATIVE_FLOAT	-3.4E + 38 to 3.4E + 38 (size * 4 bytes)
<b>VARDOUBLE</b>	H5T_NATIVE_DOUBLE	-1.79E + 308 to 1.79E + 308 (size * 8 bytes)
<b>VARCHAR</b>	H5T_C_S1	0 to 255 (size * 1 byte)
<b>OPAQUE</b>	H5T_OPAQUE	0 to 255 (1 byte)

Table 6.3 – HDFqI data types and their corresponding definitions in HDF

### 6.1.1 TINYINT

The HDFqI TINYINT data type corresponds to the HDF5 H5T\_NATIVE\_CHAR data type. It may store a value between -128 and 127, and occupies 1 byte in memory. Depending on the programming language (supported by HDFqI), the TINYINT data type is represented by:

- In C, the “char” data type.
- In C++, the “char” data type.
- In Java, the “byte” data type or its corresponding wrapper class “Byte”.
- In Python, the “int8” NumPy data type.



- In C#, the “SByte” data type or its alias “sbyte”.
- In Fortran, the “INTEGER(KIND = 1)” data type.
- In R<sup>2</sup>, the “integer” data type.

### 6.1.2 UNSIGNED TINYINT

The HDFqI UNSIGNED TINYINT data type corresponds to the HDF5 H5T\_NATIVE\_UCHAR data type. It may store a value between 0 and 255, and occupies 1 byte in memory. Depending on the programming language (supported by HDFqI), the UNSIGNED TINYINT data type is represented by:

- In C, the “unsigned char” data type.
- In C++, the “unsigned char” data type.
- In Java<sup>3</sup>, the “byte” data type or its corresponding wrapper class “Byte”.
- In Python, the “uint8” NumPy data type.
- In C#, the “Byte” data type or its alias “byte”.
- In Fortran<sup>4</sup>, the “INTEGER(KIND = 1)” data type.
- In R<sup>5</sup>, the “integer” data type.

---

<sup>2</sup> By design, R does not have a data type that may store a value between -128 and 127 with exactly 1 byte in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

<sup>3</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Java.

<sup>4</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Fortran.

<sup>5</sup> By design, R does not have a data type that may store a value between 0 and 255 with exactly 1 byte in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

### 6.1.3 SMALLINT

The HDFqI SMALLINT data type corresponds to the HDF5 H5T\_NATIVE\_SHORT data type. It may store a value between -32,768 and 32,767, and occupies 2 bytes in memory. Depending on the programming language (supported by HDFqI), the SMALLINT data type is represented by:

- In C, the “short” data type.
- In C++, the “short” data type.
- In Java, the “short” data type or its corresponding wrapper class “Short”.
- In Python, the “int16” NumPy data type.
- In C#, the “Int16” data type or its alias “short”.
- In Fortran, the “INTEGER(KIND = 2)” data type.
- In R<sup>6</sup>, the “integer” data type.

### 6.1.4 UNSIGNED SMALLINT

The HDFqI UNSIGNED SMALLINT data type corresponds to the HDF5 H5T\_NATIVE\_USHORT data type. It may store a value between 0 and 65,535, and occupies 2 bytes in memory. Depending on the programming language (supported by HDFqI), the UNSIGNED SMALLINT data type is represented by:

- In C, the “unsigned short” data type.
- In C++, the “unsigned short” data type.
- In Java<sup>7</sup>, the “short” data type or its corresponding wrapper class “Short”.
- In Python, the “uint16” NumPy data type.

---

<sup>6</sup> By design, R does not have a data type that may store a value between -32,768 and 32,767 with exactly 2 bytes in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

<sup>7</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Java.

- In C#, the “UInt16” data type or its alias “ushort”.
- In Fortran<sup>8</sup>, the “INTEGER(KIND = 2)” data type.
- In R<sup>9</sup>, the “integer” data type.

### 6.1.5 INT

The HDFqI INT data type corresponds to the HDF5 H5T\_NATIVE\_INT data type. It may store a value between -2,147,483,648 and 2,147,483,647, and occupies 4 bytes in memory. Depending on the programming language (supported by HDFqI), the INT data type is represented by:

- In C, the “int” data type.
- In C++, the “int” data type.
- In Java, the “int” data type or its corresponding wrapper class “Integer”.
- In Python, the “int32” NumPy data type.
- In C#, the “Int32” data type or its alias “int”.
- In Fortran, the “INTEGER(KIND = 4)” or “INTEGER” data types.
- In R, the “integer” data type.

### 6.1.6 UNSIGNED INT

The HDFqI UNSIGNED INT data type corresponds to the HDF5 H5T\_NATIVE\_UINT data type. It may store a value between 0 and 4,294,967,295, and occupies 4 bytes in memory. Depending on the programming language (supported by HDFqI), the UNSIGNED INT data type is represented by:

---

<sup>8</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Fortran.

<sup>9</sup> By design, R does not have a data type that may store a value between 0 and 65,535 with exactly 2 bytes in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

- In C, the “unsigned int” data type.
- In C++, the “unsigned int” data type.
- In Java<sup>10</sup>, the “int” data type or its corresponding wrapper class “Integer”.
- In Python, the “uint32” NumPy data type.
- In C#, the “UInt32” data type or its alias “uint”.
- In Fortran<sup>11</sup>, the “INTEGER(KIND = 4)” or “INTEGER” data types.
- In R<sup>12</sup>, the “integer” data type.

### 6.1.7 BIGINT

The HDFqI BIGINT data type corresponds to the HDF5 H5T\_NATIVE\_LLONG data type. It may store a value between -9,223,372,036,854,775,808 and 9,223,372,036,854,775,807, and occupies 8 bytes in memory. Depending on the programming language (supported by HDFqI), the BIGINT data type is represented by:

- In C, the “long long” data type.
- In C++, the “long long” data type.
- In Java, the “long” data type or its corresponding wrapper class “Long”.
- In Python, the “int64” NumPy data type.
- In C#, the “Int64” data type or its alias “long”.
- In Fortran, the “INTEGER(KIND = 8)” data type.
- In R, the “integer64” bit64 data type.

---

<sup>10</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned in Java.

<sup>11</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned in Fortran.

<sup>12</sup> By design, R does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned in R.

### 6.1.8 UNSIGNED BIGINT

The HDFqI UNSIGNED BIGINT data type corresponds to the HDF5 H5T\_NATIVE\_ULLONG data type. It may store a value between 0 and 18,446,744,073,709,551,615, and occupies 8 bytes in memory. Depending on the programming language (supported by HDFqI), the UNSIGNED BIGINT data type is represented by:

- In C, the “unsigned long long” data type.
- In C++, the “unsigned long long” data type.
- In Java<sup>13</sup>, the “long” data type or its corresponding wrapper class “Long”.
- In Python, the “uint64” NumPy data type.
- In C#, the “UInt64” data type or its alias “ulong”.
- In Fortran<sup>14</sup>, the “INTEGER(KIND = 8)” data type.
- In R<sup>15</sup>, the “integer64” bit64 data type.

### 6.1.9 FLOAT

The HDFqI FLOAT data type corresponds to the HDF5 H5T\_NATIVE\_FLOAT data type. It may store a value between  $-3.4E + 38$  and  $3.4E + 38$ , and occupies 4 bytes in memory. Depending on the programming language (supported by HDFqI), the FLOAT data type is represented by:

- In C, the “float” data type.
- In C++, the “float” data type.

---

<sup>13</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned in Java.

<sup>14</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned in Fortran.

<sup>15</sup> By design, R does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned in R.

- In Java, the “float” data type or its corresponding wrapper class “Float”.
- In Python, the “float32” NumPy data type.
- In C#, the “Single” data type or its alias “float”.
- In Fortran, the “REAL(KIND = 4)” or “REAL” data types.
- In R<sup>16</sup>, the “numeric” or “double” data types.

### 6.1.10 DOUBLE

The HDFqI DOUBLE data type corresponds to the HDF5 H5T\_NATIVE\_DOUBLE data type. It may store a value between  $-1.79E + 308$  and  $1.79E + 308$ , and occupies 8 bytes in memory. Depending on the programming language (supported by HDFqI), the DOUBLE data type is represented by:

- In C, the “double” data type.
- In C++, the “double” data type.
- In Java, the “double” data type or its corresponding wrapper class “Double”.
- In Python, the “float64” NumPy data type.
- In C#, the “Double” data type or its alias “double”.
- In Fortran, the “REAL(KIND = 8)” or “DOUBLE PRECISION” data types.
- In R, the “numeric” or “double” data types.

---

<sup>16</sup> By design, R does not have a data type that may store a value between  $-3.4E + 38$  and  $3.4E + 38$  with exactly 4 byte in memory. As a substitute, the R “numeric” or “double” data types may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 8 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

### 6.1.11 CHAR

The HDFqI CHAR data type corresponds to the HDF5 H5T\_C\_S1 data type. It may store a value between 0 and 255, and occupies  $size * 1$  byte in memory ( $size$  being the length of the string). The CHAR data type is useful for storing fixed-length strings. Depending on the programming language (supported by HDFqI), the CHAR data type is represented by:

- In C, the “unsigned char [ $size$ ]” data type.
- In C++, the “unsigned char [ $size$ ]” data type.
- In Java, the “byte [ $size$ ]” data type or its corresponding wrapper class “Byte [ $size$ ]”.
- In Python, the “Ssize” NumPy data type.
- In C#, the “Byte [ $size$ ]” data type or its alias “byte [ $size$ ]”.
- In Fortran, the “CHARACTER(LEN =  $size$ )” data type.
- In R<sup>17</sup>, the “integer” data type.

### 6.1.12 VARTINYINT

The HDFqI VARTINYINT data type corresponds to the HDF5 H5T\_NATIVE\_CHAR data type. It may store a value between -128 and 127, and occupies  $size * 1$  byte in memory ( $size$  being the number of elements composing the VARTINYINT data type). Depending on the programming language (supported by HDFqI), the VARTINYINT data type is represented by:

- In C, the “char” data type.
- In C++, the “char” data type.
- In Java, the “byte” data type or its corresponding wrapper class “Byte”.
- In Python, the “int8” NumPy data type.
- In C#, the “SByte” data type or its alias “sbyte”.

---

<sup>17</sup> By design, R does not have a data type that may store a value between 0 and 255 with exactly 1 byte in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

- In Fortran, the “INTEGER(KIND = 1)” data type.
- In R<sup>18</sup>, the “integer” data type.

### 6.1.13 UNSIGNED VARTINYINT

The HDFqI UNSIGNED VARTINYINT data type corresponds to the HDF5 H5T\_NATIVE\_UCHAR data type. It may store a value between 0 and 255, and occupies *size* \* 1 byte in memory (*size* being the number of elements composing the VARTINYINT data type). Depending on the programming language (supported by HDFqI), the UNSIGNED VARTINYINT data type is represented by:

- In C, the “unsigned char” data type.
- In C++, the “unsigned char” data type.
- In Java<sup>19</sup>, the “byte” data type or its corresponding wrapper class “Byte”.
- In Python, the “uint8” NumPy data type.
- In C#, the “Byte” data type or its alias “byte”.
- In Fortran<sup>20</sup>, the “INTEGER(KIND = 1)” data type.
- In R<sup>21</sup>, the “integer” data type.

---

<sup>18</sup> By design, R does not have a data type that may store a value between -128 and 127 with exactly 1 byte in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

<sup>19</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Java.

<sup>20</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Fortran.

<sup>21</sup> By design, R does not have a data type that may store a value between 0 and 255 with exactly 1 byte in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).



### 6.1.14 VARSMALLINT

The HDFqI VARSMALLINT data type corresponds to the HDF5 H5T\_NATIVE\_SHORT data type. It may store a value between -32,768 and 32,767, and occupies  $size * 2$  bytes in memory ( $size$  being the number of elements composing the VARSMALLINT data type). Depending on the programming language (supported by HDFqI), the VARSMALLINT data type is represented by:

- In C, the “short” data type.
- In C++, the “short” data type.
- In Java, the “short” data type or its corresponding wrapper class “Short”.
- In Python, the “int16” NumPy data type.
- In C#, the “Int16” data type or its alias “short”.
- In Fortran, the “INTEGER(KIND = 2)” data type.
- In R<sup>22</sup>, the “integer” data type.

### 6.1.15 UNSIGNED VARSMALLINT

The HDFqI UNSIGNED VARSMALLINT data type corresponds to the HDF5 H5T\_NATIVE\_USHORT data type. It may store a value between 0 and 65,535, and occupies  $size * 2$  bytes in memory ( $size$  being the number of elements composing the VARSMALLINT data type). Depending on the programming language (supported by HDFqI), the UNSIGNED VARSMALLINT data type is represented by:

- In C, the “unsigned short” data type.
- In C++, the “unsigned short” data type.
- In Java<sup>23</sup>, the “short” data type or its corresponding wrapper class “Short”.

---

<sup>22</sup> By design, R does not have a data type that may store a value between -32,768 and 32,767 with exactly 2 bytes in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

<sup>23</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Java.

- In Python, the “uint16” NumPy data type.
- In C#, the “UInt16” data type or its alias “ushort”.
- In Fortran<sup>24</sup>, the “INTEGER(KIND = 2)” data type.
- In R<sup>25</sup>, the “integer” data type.

## 6.1.16 VARINT

The HDFqI VARINT data type corresponds to the HDF5 H5T\_NATIVE\_INT data type. It may store a value between -2,147,483,648 and 2,147,483,647, and occupies *size* \* 4 bytes in memory (*size* being the number of elements composing the VARINT data type). Depending on the programming language (supported by HDFqI), the VARINT data type is represented by:

- In C, the “int” data type.
- In C++, the “int” data type.
- In Java, the “int” data type or its corresponding wrapper class “Integer”.
- In Python, the “int32” NumPy data type.
- In C#, the “Int32” data type or its alias “int”.
- In Fortran, the “INTEGER(KIND = 4)” data type.
- In R, the “integer” data type.

---

<sup>24</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Fortran.

<sup>25</sup> By design, R does not have a data type that may store a value between 0 and 65,535 with exactly 2 bytes in memory. As a substitute, the R “integer” data type may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 4 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

### 6.1.17 UNSIGNED VARINT

The HDFqI UNSIGNED VARINT data type corresponds to the HDF5 H5T\_NATIVE\_UINT data type. It may store a value between 0 and 4,294,967,295, and occupies *size* \* 4 bytes in memory (*size* being the number of elements composing the UNSIGNED VARINT data type). Depending on the programming language (supported by HDFqI), the UNSIGNED VARINT data type is represented by:

- In C, the “unsigned int” data type.
- In C++, the “unsigned int” data type.
- In Java<sup>26</sup>, the “int” data type or its corresponding wrapper class “Integer”.
- In Python, the “uint32” NumPy data type.
- In C#, the “UInt32” data type or its alias “uint”.
- In Fortran<sup>27</sup>, the “INTEGER(KIND = 4)” data type.
- In R<sup>28</sup>, the “integer” data type.

### 6.1.18 VARBIGINT

The HDFqI VARBIGINT data type corresponds to the HDF5 H5T\_NATIVE\_LLONG data type. It may store a value between -9,223,372,036,854,775,808 and 9,223,372,036,854,775,807, and occupies *size* \* 8 bytes in memory (*size* being the number of elements composing the VARBIGINT data type). Depending on the programming language (supported by HDFqI), the VARBIGINT data type is represented by:

- In C, the “long long” data type.
- In C++, the “long long” data type.

---

<sup>26</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Java.

<sup>27</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Fortran.

<sup>28</sup> By design, R does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in R.

- In Java, the “long” data type or its corresponding wrapper class “Long”.
- In Python, the “int64” NumPy data type.
- In C#, the “Int64” data type or its alias “long”.
- In Fortran, the “INTEGER(KIND = 8)” data type.
- In R, the “integer64” bit64 data type.

### 6.1.19 UNSIGNED VARBIGINT

The HDFqI UNSIGNED VARBIGINT data type corresponds to the HDF5 H5T\_NATIVE\_ULLONG data type. It may store a value between 0 and 18,446,744,073,709,551,615, and occupies *size* \* 8 bytes in memory (*size* being the number of elements composing the UNSIGNED VARBIGINT data type). Depending on the programming language (supported by HDFqI), the UNSIGNED VARBIGINT data type is represented by:

- In C, the “unsigned long long” data type.
- In C++, the “unsigned long long” data type.
- In Java<sup>29</sup>, the “long” data type or its corresponding wrapper class “Long”.
- In Python, the “uint64” NumPy data type.
- In C#, the “UInt64” data type or its alias “ulong”.
- In Fortran<sup>30</sup>, the “INTEGER(KIND = 8)” data type.
- In R<sup>31</sup>, the “integer64” bit64 data type.

---

<sup>29</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Java.

<sup>30</sup> Although there has been some effort to specify unsigned data types in Fortran, nothing concrete is available. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Fortran.

<sup>31</sup> By design, R does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in R.

## 6.1.20 VARFLOAT

The HDFqI VARFLOAT data type corresponds to the HDF5 H5T\_NATIVE\_FLOAT data type. It may store a value between  $-3.4E + 38$  and  $3.4E + 38$ , and occupies  $size * 4$  bytes in memory ( $size$  being the number of elements composing the VARFLOAT data type). Depending on the programming language (supported by HDFqI), the VARFLOAT data type is represented by:

- In C, the “float” data type.
- In C++, the “float” data type.
- In Java, the “float” data type or its corresponding wrapper class “Float”.
- In Python, the “float32” NumPy data type.
- In C#, the “Single” data type or its alias “float”.
- In Fortran, the “REAL(KIND = 4)” data type.
- In R<sup>32</sup>, the “numeric” or “double” data types.

## 6.1.21 VARDOUBLE

The HDFqI VARDOUBLE data type corresponds to the HDF5 H5T\_NATIVE\_DOUBLE data type. It may store a value between  $-1.79E + 308$  and  $1.79E + 308$ , and occupies  $size * 8$  bytes in memory ( $size$  being the number of elements composing the VARDOUBLE data type). Depending on the programming language (supported by HDFqI), the VARDOUBLE data type is represented by:

- In C, the “double” data type.
- In C++, the “double” data type.
- In Java, the “double” data type or its corresponding wrapper class “Double”.
- In Python, the “float64” NumPy data type.

---

<sup>32</sup> By design, R does not have a data type that may store a value between  $-3.4E + 38$  and  $3.4E + 38$  with exactly 4 byte in memory. As a substitute, the R “numeric” or “double” data types may be used with the penalties of more memory being unnecessarily reserved (as this data type occupies 8 bytes in memory) and lower performance (as bytes alignment must be made by HDFqI).

- In C#, the “Double” data type or its alias “double”.
- In Fortran, the “REAL(KIND = 8)” or “DOUBLE PRECISION” data types.
- In R, the “numeric” or “double” data types.

### 6.1.22 VARCHAR

The HDFqI VARCHAR data type corresponds to the HDF5 H5T\_C\_S1 data type. It may store a value between 0 and 255, and occupies  $size * 1$  byte in memory ( $size$  being the length of the string). The VARCHAR data type is useful for storing variable-length strings. Depending on the programming language (supported by HDFqI), the VARCHAR data type is represented by:

- In C, the “unsigned char \*” data type.
- In C++, the “unsigned char \*” data type.
- In Java, the “String” object.
- In Python, the “Ssize” NumPy data type.
- In C#, the “String” data type or its alias “string”.
- In Fortran, the “CHARACTER(LEN = \*)” data type.
- In R, the “character” data type.

### 6.1.23 OPAQUE

The HDFqI OPAQUE data type corresponds to the HDF5 H5T\_OPAQUE data type. It may store a value between 0 and 255, and occupies 1 byte in memory. The OPAQUE data type is useful for representing data that should not be interpreted/rearranged by the HDF5 library when reading/writing it from/into in an HDF5 dataset or attribute. Depending on the programming language (supported by HDFqI), the OPAQUE data type is represented by:

- In C, the “unsigned char” data type.
- In C++, the “unsigned char” data type.

- In Java<sup>33</sup>, the “byte” data type or its corresponding wrapper class “Byte”.
- In Python, the “ubyte” NumPy data type.
- In C#, the “Byte” data type or its alias “byte”.
- In Fortran, the “CHARACTER” data type.
- In R, the “raw” data type.

## 6.2 POST-PROCESSING

Post-processing options enable transforming results of a query according to the programmer’s needs such as ordering or truncating. These options may be used to create a (linear) pipeline to further process result sets returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations. In case a pipeline is composed of two or more options, the order in which they are used affects the final outcome (e.g. usage of ORDER ASC followed by TOP 2 in a result set composed of 4, 2, 3 and 1, returns 1 and 2; usage of these same two options inverted – i.e. TOP 2 followed by ORDER ASC – returns 2 and 4 instead). The next subsections describe the post-processing options provided by HDFql.

Post-processing Option	Description
<a href="#">ORDER</a>	Order (i.e. sort) a result set in an ascending, descending or reverse way
<a href="#">TOP</a>	Truncate a result set after a certain given position in a topmost way
<a href="#">BOTTOM</a>	Truncate a result set after a certain given position in a bottommost way
<a href="#">FROM TO</a>	Retain a result set within a certain given range
<a href="#">STEP</a>	Step (i.e. jump) the result set at every given position

Table 6.4 – HDFql post-processing options

---

<sup>33</sup> By design, Java does not support unsigned data types. Therefore, the programmer is responsible for making the conversion from a signed number to its equivalent unsigned number in Java.

## 6.2.1 ORDER

### Syntax

**ORDER** {{ASC | DESC | REV} | {, {ASC | DESC | REV}}} | {{ASC | DESC | REV}, {ASC | DESC | REV}}

### Description

Order (i.e. sort) a result set in an ascending, descending or reverse way by specifying either the keyword ASC, DESC or REV respectively. When in an ascending or descending order, HDFqL automatically uses all available CPU cores to speed-up the task completion<sup>34</sup>. Additionally, if the result set is of data type [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) or [HDFQL\\_OPAQUE](#), the result subset can be ordered (i.e. sorted) in an ascending, descending or reverse way by specifying a comma (,) and either the keyword ASC, DESC or REV, respectively.

### Parameter(s)

None

### Return

The result set and/or subset is ordered (i.e. sorted) in an ascending, descending or reverse way depending on whether the keyword ASC, DESC or REV is specified respectively.

### Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type float of four dimensions (size
5x8x4x7)
CREATE DATASET my_dataset0 AS FLOAT(5, 8, 4, 7)

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these
(should be 5, 8, 4, 7)
SHOW DIMENSION my_dataset0

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these in
ascending order (should be 4, 5, 7, 8)
```

<sup>34</sup> Through a parallelized Quicksort algorithm.



```
SHOW DIMENSION my_dataset0 ORDER ASC
```

```
# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these in  
descending order (should be 8, 7, 5, 4)
```

```
SHOW DIMENSION my_dataset0 ORDER DESC
```

```
# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these in  
reversed order (should be 7, 4, 8, 5)
```

```
SHOW DIMENSION my_dataset0 ORDER REV
```

```
# create an HDF5 dataset named "my_dataset1" of data type double of two dimensions (size 3x2)
```

```
CREATE DATASET my_dataset1 AS DOUBLE(3, 2)
```

```
# insert (i.e. write) values into dataset "my_dataset1"
```

```
INSERT INTO my_dataset1 VALUES((3.2, 1.3), (0, 0.2), (9.1, 6.5))
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should  
be 3.2, 1.3, 0, 0.2, 9.1, 6.5)
```

```
SELECT FROM my_dataset1
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it in  
ascending order (should be 0, 0.2, 1.3, 3.2, 6.5, 9.1)
```

```
SELECT FROM my_dataset1 ORDER ASC
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it in  
descending order (should be 9.1, 6.5, 3.2, 1.3, 0.2, 0)
```

```
SELECT FROM my_dataset1 ORDER DESC
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it in  
reversed order (should be 6.5, 9.1, 0.2, 0, 1.3, 3.2)
```

```
SELECT FROM my_dataset1 ORDER REV
```

```
# create an HDF5 dataset named "my_dataset2" of data type variable-length double of one  
dimension (size 3)
```

```
CREATE DATASET my_dataset2 AS VARDOUBLE(3)
```

```
# insert (i.e. write) values into dataset "my_dataset2"
```

```
INSERT INTO my_dataset2 VALUES((3.2, 1.3), (0, 0.2), (9.1, 7.4, 6.5))
```

```
# select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it (should  
be 3.2, 1.3, 0, 0.2, 9.1, 7.4, 6.5)
```

```
SELECT FROM my_dataset2
```

```
# select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it in
ascending order on the result subset only (should be 1.3, 3.2, 0, 0.2, 6.5, 7.4, 9.1)
SELECT FROM my_dataset2 ORDER , ASC

# select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it in
descending order on the result subset only (should be 3.2, 1.3, 0.2, 0, 9.1, 7.4, 6.5)
SELECT FROM my_dataset2 ORDER , DESC

# select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it in
reversed order on the result set only (should be 9.1, 7.4, 6.5, 0, 0.2, 3.2, 1.3)
SELECT FROM my_dataset2 ORDER REV

# select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it in
reversed order on the result subset only (should be 1.3, 3.2, 0.2, 0, 6.5, 7.4, 9.1)
SELECT FROM my_dataset2 ORDER , REV

# select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it in
reversed order on both the result set and result subset (should be 6.5, 7.4, 9.1, 0.2, 0, 1.3,
3.2)
SELECT FROM my_dataset2 ORDER REV, REV
```

## 6.2.2 TOP

### Syntax

TOP {*top\_value* | {, *subtop\_value*} | {*top\_value*, *subtop\_value*}}

### Description

Truncate a result set after position *top\_value* in a topmost way. In other words, all elements after position *top\_value* are discarded from the result set. Additionally, if the result set is of data type [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) or [HDFQL\\_OPAQUE](#), the result subset can be truncated in a topmost way by specifying a comma (,) and *subtop\_value*.

### Parameter(s)

*top\_value* – optional integer that specifies the position of the truncation of a result set in a topmost way. If negative, the TOP option will behave as the BOTTOM option with a positive *top\_value*.

*subtop\_value* – optional integer that specifies the position of the truncation of a result set in a topmost way. If negative, the TOP option will behave as the BOTTOM option with a positive *subtop\_value*. Of note, this parameter is only applicable for a result set of one of the aforementioned data types and ignored otherwise.

## **Return**

The result set and/or subset is truncated in a topmost way in function of the position provided.

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type float of four dimensions (size
5x8x4x7)
CREATE DATASET my_dataset0 AS FLOAT(5, 8, 4, 7)

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these
(should be 5, 8, 4, 7)
SHOW DIMENSION my_dataset0

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the
topmost (i.e. first) dimension (should be 5)
SHOW DIMENSION my_dataset0 TOP 1

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the two
topmost dimensions (should be 5, 8)
SHOW DIMENSION my_dataset0 TOP 2

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the two
bottommost dimensions (should be 4, 7)
SHOW DIMENSION my_dataset0 TOP -2

# create an HDF5 dataset named "my_dataset1" of data type variable-length int of one dimension
(size 3) with initial values of 12, 14 and 16 for the first position, 18 for the second
position, and 20, 22, 24 and 26 for the third position
CREATE DATASET my_dataset1 AS VARINT(3) VALUES((12, 14, 16), (18), (20, 22, 24, 26))

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should
be 12, 14, 16, 18, 20, 22, 24, 26)
SELECT FROM my_dataset1

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with values of
the topmost (i.e. first) position (should be 12, 14, 16)
SELECT FROM my_dataset1 TOP 1
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with values of
the two topmost positions (should be 12, 14, 18, 20, 22)
SELECT FROM my_dataset1 TOP , 2

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with the
topmost value of the two bottommost positions (should be 18, 20)
SELECT FROM my_dataset1 TOP -2, 1
```

## 6.2.3 BOTTOM

### Syntax

**BOTTOM** {*bottom\_value* | {, *subbottom\_value*} | {*bottom\_value*, *subbottom\_value*}}

### Description

Truncate a result set after position *bottom\_value* in a bottommost way. In other words, all elements before position *bottom\_value* are discarded from the result set. Additionally, if the result set is of data type [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) or [HDFQL\\_OPAQUE](#), the result subset can be truncated in a bottommost way by specifying a comma (,) and *subbottom\_value*.

### Parameter(s)

*bottom\_value* – optional integer that specifies the position of the truncation of a result set in a bottommost way. If negative, the BOTTOM option will behave as the TOP option with a positive *bottom\_value*.

*subbottom\_value* – optional integer that specifies the position of the truncation of a result set in a bottommost way. If negative, the BOTTOM option will behave as the TOP option with a positive *subbottom\_value*. Of note, this parameter is only applicable for a result set of one of the aforementioned data types and ignored otherwise.

### Return

The result set and/or subset is truncated in a bottommost way in function of the position provided.

## Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type float of four dimensions (size
5x8x4x7)
CREATE DATASET my_dataset0 AS FLOAT(5, 8, 4, 7)

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these
(should be 5, 8, 4, 7)
SHOW DIMENSION my_dataset0

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the
bottommost (i.e. last) dimension (should be 7)
SHOW DIMENSION my_dataset0 BOTTOM 1

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the two
bottommost dimensions (should be 4, 7)
SHOW DIMENSION my_dataset0 BOTTOM 2

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the two
topmost dimensions (should be 5, 8)
SHOW DIMENSION my_dataset0 BOTTOM -2

# create an HDF5 dataset named "my_dataset1" of data type variable-length int of one dimension
(size 3) with initial values of 12, 14 and 16 for the first position, 18 for the second
position, and 20, 22, 24 and 26 for the third position
CREATE DATASET my_dataset1 AS VARINT(3) VALUES((12, 14, 16), (18), (20, 22, 24, 26))

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should
be 12, 14, 16, 18, 20, 22, 24, 26)
SELECT FROM my_dataset1

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with values of
the bottommost (i.e. last) position (should be 20, 22, 24, 26)
SELECT FROM my_dataset1 BOTTOM 1

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with values of
the two bottommost positions (should be 14, 16, 18, 24, 26)
SELECT FROM my_dataset1 BOTTOM , 2

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with the
bottommost value of the two topmost positions (should be 16, 18)
SELECT FROM my_dataset1 BOTTOM -2, 1
```

## 6.2.4 FROM TO

### Syntax

**FROM** {*from\_value* | {, *subfrom\_value*} | {*from\_value*, *subfrom\_value*}} **TO** {*to\_value* | {, *subto\_value*} | {*to\_value*, *subto\_value*}}

### Description

Retain a result set from *from\_value* to *to\_value*. In other words, all elements before position *from\_value* and after position *to\_value* are discarded from the result set. The first element of the result set is at position zero, while the last element is located at the position returned by `hdfql_cursor_get_count` - 1. Additionally, if the result set is of data type `HDFQL_CHAR`, `HDFQL_VARTINYINT`, `HDFQL_UNSIGNED_VARTINYINT`, `HDFQL_VARSMALLINT`, `HDFQL_UNSIGNED_VARSMALLINT`, `HDFQL_VARINT`, `HDFQL_UNSIGNED_VARINT`, `HDFQL_VARBIGINT`, `HDFQL_UNSIGNED_VARBIGINT`, `HDFQL_VARFLOAT`, `HDFQL_VARDOUBLE` or `HDFQL_OPAQUE`, the result subset can be retained by specifying a comma (,), *subfrom\_value* and/or *subto\_value*.

### Parameter(s)

*from\_value* – optional integer that specifies the starting position to retain elements of a result set. If negative, the FROM option will retain elements of a result set starting from its end.

*subfrom\_value* – optional integer that specifies the starting position to retain elements of a result set. If negative, the FROM option will retain elements of a result set starting from its end. Of note, this parameter is only applicable for a result set of one of the aforementioned data types and ignored otherwise.

*to\_value* – optional integer that specifies the ending position to retain elements of a result set. If negative, the TO option will retain elements of a result set starting from its end.

*subto\_value* – optional integer that specifies the ending position to retain elements of a result set. If negative, the TO option will retain elements of a result set starting from its end. Of note, this parameter is only applicable for a result set of one of the aforementioned data types and ignored otherwise.

### Return

The result set and/or subset is retained in function of the position provided.

## Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type float of four dimensions (size
5x8x4x7)
CREATE DATASET my_dataset0 AS FLOAT(5, 8, 4, 7)

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these
(should be 5, 8, 4, 7)
SHOW DIMENSION my_dataset0

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the
first, second and third dimensions (should be 5, 8, 4)
SHOW DIMENSION my_dataset0 FROM 0 TO 2

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the
second and third dimensions (should be 8, 4)
SHOW DIMENSION my_dataset0 FROM 1 TO 2

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the
second, third and fourth dimensions (should be 8, 4, 7)
SHOW DIMENSION my_dataset0 FROM -3 TO -1

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with the
second and third dimensions (should be 8, 4)
SHOW DIMENSION my_dataset0 FROM 1 TO -2

# create an HDF5 dataset named "my_dataset1" of data type variable-length int of one dimension
(size 3) with initial values of 12, 14 and 16 for the first position, 18 for the second
position, and 20, 22, 24 and 26 for the third position
CREATE DATASET my_dataset1 AS VARINT(3) VALUES((12, 14, 16), (18), (20, 22, 24, 26))

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should
be 12, 14, 16, 18, 20, 22, 24, 26)
SELECT FROM my_dataset1

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with values of
the second position (should be 18)
SELECT FROM my_dataset1 FROM 1 TO 1

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with values of
the second and third positions (should be 18, 20, 22, 24, 26)
SELECT FROM my_dataset1 FROM -2 TO -1
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with the second
and third values of all positions (should be 14, 16, 22, 24)
SELECT FROM my_dataset1 FROM , 1 TO , 2

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with second
bottommost and bottommost values of the first position (should be 14, 16)
SELECT FROM my_dataset1 FROM 0, -2 TO 0, -1
```

## 6.2.5 STEP

### Syntax

STEP {step\_value | {, substep\_value} | {step\_value, substep\_value}}

### Description

Step (i.e. jump) the result set at every *step\_value* position. In other words, all elements between steps are discarded from the result set. Additionally, if the result set is of data type [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) or [HDFQL\\_OPAQUE](#), the result subset can be stepped (i.e. jumped) by specifying a comma (,) and *substep\_value*.

### Parameter(s)

*step\_value* – optional integer that specifies the position to step (i.e. jump) a result set. If *step\_value* is negative, the STEP option will step (i.e. jump) the result set starting from its end.

*substep\_value* – optional integer that specifies the position to step (i.e. jump) a result set. If *substep\_value* is negative, the STEP option will step (i.e. jump) the result set starting from its end. Of note, this parameter is only applicable for a result set of one of the aforementioned data types and ignored otherwise.

### Return

The result set and/or subset is stepped (i.e. jumped) in function of the position provided.

### Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type float of four dimensions (size
5x8x4x7)
```



```
CREATE DATASET my_dataset0 AS FLOAT(5, 8, 4, 7)

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these
# (should be 5, 8, 4, 7)
SHOW DIMENSION my_dataset0

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with these
# (should be 5, 8, 4, 7)
SHOW DIMENSION my_dataset0 STEP 1

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with every
# second dimension (should be 5, 4)
SHOW DIMENSION my_dataset0 STEP 2

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with every
# second dimension starting from the end (should be 8, 7)
SHOW DIMENSION my_dataset0 STEP -2

# show (i.e. get) dimensions of dataset "my_dataset0" and populate cursor in use with every
# third dimension (should be 5, 7)
SHOW DIMENSION my_dataset0 STEP 3

# create an HDF5 dataset named "my_dataset1" of data type variable-length int of one dimension
# (size 3) with initial values of 12, 14 and 16 for the first position, 18 for the second
# position, and 20, 22, 24 and 26 for the third position
CREATE DATASET my_dataset1 AS VARINT(3) VALUES((12, 14, 16), (18), (20, 22, 24, 26))

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should
# be 12, 14, 16, 18, 20, 22, 24, 26)
SELECT FROM my_dataset1

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with values of
# every second position (should be 12, 14, 16, 20, 22, 24, 26)
SELECT FROM my_dataset1 STEP 2

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with every
# third value of all positions (should be 12, 18, 20, 26)
SELECT FROM my_dataset1 STEP , 3

# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with every
# second value of every second position (should be 12, 16, 22, 26)
SELECT FROM my_dataset1 STEP 2, -2
```

## 6.3 REDIRECTING

Redirecting options enable reading data from the cursor in use, a (text or binary) file or memory (i.e. user-defined variable) and writing it into an HDF5 dataset or attribute through a [CREATE DATASET](#), [CREATE ATTRIBUTE](#) or [INSERT](#) operation. It also enables writing result sets (i.e. data) returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations into the cursor in use, a (text or binary) file or memory. The next subsections describe the redirecting options provided by HDFqI.

Redirecting Option	Description
<a href="#">FROM</a>	Read data from the cursor in use, a file or memory and write it into an HDF5 dataset or attribute
<a href="#">INTO</a>	Write result sets into the cursor in use, a file or memory

Table 6.5 – HDFqI redirecting options

### 6.3.1 FROM

#### Syntax

**FROM** {**CURSOR** | {[**DOS** | **UNIX**] [**TEXT**] **FILE** *file\_name* [**SEPARATOR** {*separator\_value* | {, *subseparator\_value*} | {*separator\_value*, *subseparator\_value*}}]} | **BINARY FILE** *file\_name*} | **MEMORY** *variable\_number* [**SIZE** *variable\_size*]}}

#### Description

Read data from the cursor in use (default behavior when no redirecting option is specified), a (text or binary) file or memory (i.e. user-defined variable) and write it into an HDF5 dataset or attribute through a [CREATE DATASET](#), [CREATE ATTRIBUTE](#) or [INSERT](#) operation. This procedure (which is known as input redirecting option) can be performed from:

- The cursor in use. Example: “*CREATE DATASET my\_dataset AS FLOAT VALUES FROM CURSOR*” or “*INSERT INTO my\_dataset VALUES FROM CURSOR*”.
- A text file using optional parameters such as which terminator to use – DOS (CR+LF) or UNIX (LF) – for the end of line (EOL), or the separator to use between elements (of the data). Example: “*CREATE DATASET my\_dataset AS FLOAT VALUES FROM TEXT FILE my\_file.txt*” or “*INSERT INTO my\_dataset VALUES FROM TEXT FILE my\_file.txt*”.

- A binary file. Example: *“CREATE DATASET my\_dataset AS FLOAT VALUES FROM BINARY FILE my\_file.bin”* or *“INSERT INTO my\_dataset VALUES FROM BINARY FILE my\_file.bin”*.
- A user-defined variable that was previously registered through the function [hdfql\\_variable\\_register](#). Example: *“CREATE DATASET my\_dataset AS FLOAT VALUES FROM MEMORY 0”* or *“INSERT INTO my\_dataset VALUES FROM MEMORY 2”*. Of note, when working in Java, HDFqL has to copy each element of the Java variable into the HDF5 dataset or attribute (managed by the underlying HDFqL C library) as the JVM does not provide a direct access to the memory associated to the variable, which induces a performance penalty. This penalty is not present when working in other programming languages supported by HDFqL – namely C, C++, Python, C#, Fortran and R – as these provide a way for the underlying HDFqL C library to access the variable directly.

### **Parameter(s)**

*file\_name* – optional string that specifies the name of a text or binary file to read data from.

*separator\_value* – optional string that specifies the separator to use between elements (of the data) when reading these from a text file. If not specified, its default value is a comma (,).

*subseparator\_value* – optional string that specifies the subseparator to use between elements (of the data) when reading these from a text file. The subseparator is only applicable when the data type of the HDF5 dataset or attribute is either [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) or [HDFQL\\_OPAQUE](#), and ignored otherwise. If not specified, its default value is a space.

*variable\_number* – optional integer that specifies the number of the variable whose data will be written into the HDF5 dataset or attribute. The number is returned by the function [hdfql\\_variable\\_register](#) upon registering the variable or, subsequently, returned by the function [hdfql\\_variable\\_get\\_number](#).

*variable\_size* – optional integer that specifies the maximum size (in bytes) of the data stored in the variable to write into the HDF5 dataset or attribute. In other words, only the first *variable\_size* bytes of the data stored in the variable will be written into the dataset or attribute. Of note, *variable\_size* may be smaller than the size (in bytes) of the dataset or attribute (in this case, the remainder of the dataset or attribute will be zeroed if a number or emptied if a string).

### **Example(s)**

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5

# show (i.e. get) HDF5 file currently in use and populate cursor in use with it
```

**SHOW USE FILE**

```
# create an HDF5 dataset named "my_dataset0" of data type variable-length char with initial values from the cursor in use
```

```
CREATE DATASET my_dataset0 AS VARCHAR VALUES FROM CURSOR
```

```
# select (i.e. read) data from dataset "my_dataset0" and populate cursor in use with it (should be "my_file.h5")
```

```
SELECT FROM my_dataset0
```

```
# create an HDF5 dataset named "my_dataset1" of data type char of one dimension (size 3)
```

```
CREATE DATASET my_dataset1 AS TINYINT(3)
```

```
# insert (i.e. write) values from a text file named "my_file0.txt" into dataset "my_dataset1" (assume that the file "my_file0.txt" exists and contains "65,66,67")
```

```
INSERT INTO my_dataset1 VALUES FROM FILE my_file0.txt
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should be 65, 66, 67)
```

```
SELECT FROM my_dataset1
```

```
# insert (i.e. write) values from a text file named "my_file1.txt" into dataset "my_dataset1" (assume that the file "my_file1.txt" exists and contains "90**92**94")
```

```
INSERT INTO my_dataset1 VALUES FROM TEXT FILE my_file1.txt SEPARATOR **
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should be 90, 92, 94)
```

```
SELECT FROM my_dataset1
```

```
# insert (i.e. write) values from binary file "my_file.bin" into dataset "my_dataset1" (assume that the file "my_file.bin" exists and contains "ABC")
```

```
INSERT INTO my_dataset1 VALUES FROM BINARY FILE my_file.bin
```

```
# select (i.e. read) data from dataset "my_dataset1" and populate cursor in use with it (should be 65, 66, 67)
```

```
SELECT FROM my_dataset1
```

```
// declare variables
```

```
char script[1024];
```

```
double data[3][2];
```

```
int x;
int y;

// create an HDF5 dataset named "my_dataset2" of data type double of two dimensions (size 3x2)
hdfql_execute("CREATE DATASET my_dataset2 AS DOUBLE(3, 2)");

// populate variable "data" with certain values
data[0][0] = 3.2;
data[0][1] = 1.3;
data[1][0] = 0;
data[1][1] = 0.2;
data[2][0] = 9.1;
data[2][1] = 6.5;

// register variable "data" for subsequent use (by HDFq1)
hdfql_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset "my_dataset2"
sprintf(script, "INSERT INTO my_dataset2 VALUES FROM MEMORY %d",
hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfql_variable_unregister(data);

// select (i.e. read) data from dataset "my_dataset2" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset2");

// display content of cursor in use (should be 3.2, 1.3, 0, 0.2, 9.1, 6.5)
while(hdfql_cursor_next(NULL) == HDFQL_SUCCESS)
{
    printf("%f\n", *hdfql_cursor_get_double(NULL));
}
```

```
// declare variables
char script[1024];
HDFQL_VARIABLE_LENGTH data[3];

// create an HDF5 dataset named "my_dataset3" of data type variable-length double of one
dimension (size 3)
```

```
hdfq1_execute("CREATE DATASET my_dataset3 AS VARDOUBLE(3)");

// allocate memory in variable "data"
data[0].address = malloc(2 * sizeof(double));
data[0].count = 2;
data[1].address = malloc(3 * sizeof(double));
data[1].count = 3;
data[2].address = malloc(1 * sizeof(double));
data[2].count = 1;

// populate variable "data" with certain values
*((double *) data[0].address + 0) = 3.2;
*((double *) data[0].address + 1) = 1.3;
*((double *) data[1].address + 0) = 0;
*((double *) data[1].address + 1) = 0.2;
*((double *) data[1].address + 2) = 9.1;
*((double *) data[2].address + 0) = 6.5;

// register variable "data" for subsequent use (by HDFq1)
hdfq1_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset "my_dataset3"
sprintf(script, "INSERT INTO my_dataset3 VALUES FROM MEMORY %d",
hdfq1_variable_get_number(data));

// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);

// select (i.e. read) data from dataset "my_dataset3" and populate cursor in use with it
hdfq1_execute("SELECT FROM my_dataset3");

// display content of cursor in use (should be 3.2, 1.3, 0, 0.2, 9.1, 6.5)
while(hdfq1_cursor_next(NULL) == HDFQL_SUCCESS)
{
    while(hdfq1_subcursor_next(NULL) == HDFQL_SUCCESS)
    {
        printf("%f\n", *hdfq1_subcursor_get_double(NULL));
    }
}

// release memory allocated in variable "data"
```

```
free(data[0].address);
free(data[1].address);
free(data[2].address);
```

```
// declare variables
char script[1024];
char *data[3];

// create an HDF5 dataset named "my_dataset4" of data type variable-length char of one
dimension (size 3)
hdfq1_execute("CREATE DATASET my_dataset4 AS VARCHAR(3)");

// allocate memory in variable "data"
data[0] = malloc(13 * sizeof(char));
data[1] = malloc(5 * sizeof(char));
data[2] = malloc(7 * sizeof(char));

// populate variable "data" with certain values
strcpy(data[0], "Hierarchical");
strcpy(data[1], "Data");
strcpy(data[2], "Format");

// register variable "data" for subsequent use (by HDFq1)
hdfq1_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset "my_dataset4"
sprintf(script, "INSERT INTO my_dataset4 VALUES FROM MEMORY %d",
hdfq1_variable_get_number(data));

// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);

// select (i.e. read) data from dataset "my_dataset4" and populate cursor in use with it
hdfq1_execute("SELECT FROM my_dataset4");

// display content of cursor in use (should be "Hierarchical", "Data", "Format")
while(hdfq1_cursor_next(NULL) == HDFQL_SUCCESS)
{
    printf("%s\n", hdfq1_cursor_get_char(NULL));
}
```

```

}

// release memory allocated in variable "data"
free(data[0]);
free(data[1]);
free(data[2]);

```

## 6.3.2 INTO

### Syntax

**INTO** {**CURSOR** | {[**TRUNCATE**] [**DOS** | **UNIX**] [**TEXT**] **FILE** *file\_name* [**SEPARATOR** {*separator\_value* | {*subseparator\_value*} | {*separator\_value*, *subseparator\_value*}}] [**SPLIT** *split\_value*]} | {[**TRUNCATE**] **BINARY FILE** *file\_name*} | {**MEMORY** *variable\_number* [**SIZE** *variable\_size*]}}

### Description

Write result sets (i.e. data) returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations into the cursor in use (default behavior when no redirecting option is specified), a (text or binary) file or memory (i.e. user-defined variable). This procedure (which is known as output redirecting option) can be performed into:

- The cursor in use. Example: “*SELECT FROM my\_dataset INTO CURSOR*” or “*SHOW USE DIRECTORY INTO CURSOR*”.
- A text file using optional parameters such as which terminator to use – DOS (CR+LF) or UNIX (LF) – for the end of line (EOL), which separator to use between elements (of the result set), or the number of elements to write per line before starting writing remaining elements in a new line. Example: “*SELECT FROM my\_dataset INTO TEXT FILE my\_file.txt*” or “*SHOW USE DIRECTORY INTO TEXT FILE output.txt*”.
- A binary file. Example: “*SELECT FROM my\_dataset INTO BINARY FILE my\_file.bin*” or “*SHOW USE DIRECTORY INTO BINARY FILE output.bin*”. When redirecting data of type [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#) or [HDFQL\\_VARDOUBLE](#) into a binary file, each result subset to be written is preceded by its number of elements (as a C “unsigned int” data type with a 4 bytes size). This is to enable a correct interpretation/parsing of the binary file when reading it afterwards.
- A user-defined variable that was previously registered through the function [hdfql\\_variable\\_register](#). Example: “*SELECT FROM my\_dataset INTO MEMORY 0*” or “*SHOW USE DIRECTORY INTO MEMORY 2*”. When redirecting data of type [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#),



[HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) or [HDFQL\\_VARCHAR](#) into a user-defined variable, the programmer is responsible for releasing the memory (allocated by HDFqL) afterwards. Of note, when working in Java, HDFqL has to copy each element of the result set (managed by the underlying HDFqL C library) into the Java variable as the JVM does not provide a direct access to the memory associated to the variable, which induces a performance penalty. This penalty is not present when working in other programming languages supported by HDFqL – namely C, C++, Python, C#, Fortran and R – as these provide a way for the underlying HDFqL C library to access the memory of the variable directly.

When redirecting a result set into a file that already exists, the result set is appended to it. To overwrite an existing file, specify the keyword TRUNCATE (all data stored in the file will be permanently lost).

### **Parameter(s)**

*file\_name* – optional string that specifies the name of a text or binary file to redirect (i.e. write) a result set into.

*separator\_value* – optional string that specifies the separator to use between elements (of the result set) when redirecting (i.e. writing) these in a text file. If not specified, its default value is a comma (,).

*subseparator\_value* – optional string that specifies the subseparator to use between elements (of the result subset) when redirecting (i.e. writing) these in a text file. The subseparator is only applicable when the data type of the result set is either [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) or [HDFQL\\_OPAQUE](#), and ignored otherwise. If not specified, its default value is a space.

*split\_value* – optional integer that specifies the number of elements (of the result set) to redirect (i.e. write) per line before starting writing remaining elements in a new line in a text file. If *split\_value* is specified it must be equal to or greater than zero (otherwise an error will be raised). Otherwise, if it is not specified, no splitting is done which means that all elements (of the result set) are redirected (i.e. written) in the same line.

*variable\_number* – optional integer that specifies the number of the variable that will store the result set (i.e. data) returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations. The number is returned by the function [hdfql\\_variable\\_register](#) upon registering the variable or, subsequently, returned by the function [hdfql\\_variable\\_get\\_number](#).

*variable\_size* – optional integer that specifies the maximum size (in bytes) to use from the variable to store the result set (i.e. data). In other words, only the first *variable\_size* bytes of the variable will be used to store the result set. Of note, *variable\_size* may be smaller than the size (in bytes) of the result set (in this case, the remainder of the result set is discarded). If *variable\_size* is specified it must be equal to or greater than zero (otherwise an error will be raised).

Otherwise, if it is not specified, the variable must have enough space to store the entire result set (otherwise an error may occur such as a segmentation fault).

### **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type char of one dimension (size 3)
CREATE DATASET my_dataset0 AS TINYINT(3)

# insert (i.e. write) values into dataset "my_dataset0"
INSERT INTO my_dataset0 VALUES(65, 66, 67)

# select (i.e. read) data from dataset "my_dataset0" and populate cursor in use with it (should
be 65, 66, 67)
SELECT FROM my_dataset0

# select (i.e. read) data from dataset "my_dataset0" and populate cursor in use with it (should
be 65, 66, 67)
SELECT FROM my_dataset0 INTO CURSOR

# select (i.e. read) data from dataset "my_dataset0" and write it into a text file named
"my_file0.txt" using default separator "," (should be "65,66,67," in one single line)
SELECT FROM my_dataset0 INTO FILE my_file0.txt

# select (i.e. read) data from dataset "my_dataset0" and write it into a text file named
"my_file1.txt" using separator "***" (should be "65**66**67**" in one single line)
SELECT FROM my_dataset0 INTO TEXT FILE my_file1.txt SEPARATOR **

# select (i.e. read) data from dataset "my_dataset0" and write it into a text file named
"my_file2.txt" splitting every two values in a new line using a UNIX-based EOL terminator
(should be "65,65" in the first line and "67" in the second line)
SELECT FROM my_dataset0 INTO UNIX TEXT FILE my_file2.txt SPLIT 2

# select (i.e. read) data from dataset "my_dataset0" and write it into a binary file (truncate
it if it already exists) named "my_file.bin" (should be "ABC")
SELECT FROM my_dataset0 INTO TRUNCATE BINARY FILE my_file.bin
```

```
// declare variables
char script[1024];
double data[3][2];
int x;
int y;
```

```
// create an HDF5 dataset named "my_dataset1" of data type double of two dimensions (size 3x2)
hdfq1_execute("CREATE DATASET my_dataset1 AS DOUBLE(3, 2)");

// insert (i.e. write) values into dataset "my_dataset1"
hdfq1_execute("INSERT INTO my_dataset1 VALUES((3.2, 1.3), (0, 0.2), (9.1, 6.5))");

// register variable "data" for subsequent use (by HDFq1)
hdfq1_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset1" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset1 INTO MEMORY %d", hdfq1_variable_get_number(data));

// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);

// display content of variable "data" (should be 3.2, 1.3, 0, 0.2, 9.1, 6.5)
for(x = 0; x < 3; x++)
{
    for(y = 0; y < 2; y++)
    {
        printf("%d\n", data[x][y]);
    }
}
```

```
// declare variables
char script[1024];
HDFQL_VARIABLE_LENGTH data[3];
int x;
int y;
int count;

// create an HDF5 dataset named "my_dataset2" of data type variable-length double of one
dimension (size 3)
hdfq1_execute("CREATE DATASET my_dataset2 AS VARDOUBLE(3)");

// insert (i.e. write) values into dataset "my_dataset2"
hdfq1_execute("INSERT INTO my_dataset2 VALUES((3.2, 1.3), (0, 0.2), (9.1, 7.4, 6.5))");
```

```
// register variable "data" for subsequent use (by HDFql)
hdfql_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset2" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset2 INTO MEMORY %d", hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFql)
hdfql_variable_unregister(data);

// display content of cursor in use (should be 3.2, 1.3, 0, 0.2, 9.1, 7.4, 6.5)
for(x = 0; x < 3; x++)
{
    count = data[x].count;
    for(y = 0; y < count; y++)
    {
        printf("%f\n", *((double *) data[x].address + y));
    }
}

// release memory allocated (by HDFql) in variable "data"
for(x = 0; x < 3; x++)
{
    free(data[x].address);
}
```

```
// declare variables
char script[1024];
char *data[3];
int x;

// create an HDF5 dataset named "my_dataset3" of data type variable-length char of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset3 AS VARCHAR(3)");

// insert (i.e. write) values into dataset "my_dataset3"
hdfql_execute("INSERT INTO my_dataset3 VALUES(\"Hierarchical\", \"Data\", \"Format\")");
```

```
// register variable "data" for subsequent use (by HDFq1)
hdfq1_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset3" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset3 INTO MEMORY %d", hdfq1_variable_get_number(data));

// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);

// display content of cursor in use (should be "Hierarchical", "Data", "Format")
for(x = 0; x < 3; x++)
{
    printf("%s\n", data[x]);
}

// release memory allocated (by HDFq1) in variable "data"
for(x = 0; x < 3; x++)
{
    free(data[x]);
}
```

## 6.4 DATA DEFINITION LANGUAGE (DDL)

Data Definition Language (DDL) is, generally speaking, syntax for defining and modifying structures that store data. In HDFq1, the DDL assembles the operations that enable the creation, alteration, renaming, copying and deletion of HDF5 files, groups, datasets, attributes and links. These operations begin either with the keyword CREATE, ALTER, RENAME, COPY or DROP.

### 6.4.1 CREATE DIRECTORY

#### Syntax

**CREATE DIRECTORY** *directory\_name* [, *directory\_name*]\*

## **Description**

Create a directory named *directory\_name*. Multiple directories can be created at once by separating these with a comma (.). If *directory\_name* already exists, it will not be overwritten, no subsequent directories are created, and an error is raised. In case *directory\_name* has intermediate directories that do not exist, besides *directory\_name* being created, all these intermediate directories will be created on the fly (e.g. when creating the directory "my\_directory/my\_subdirectory/my\_subsubdirectory", besides "my\_subsubdirectory" being created, "my\_directory" and "my\_subdirectory" will be created in case they do not exist).

## **Parameter(s)**

*directory\_name* – mandatory string that specifies the name of the directory to create. Multiple directories are separated with a comma (.). As a general rule, in case *directory\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes ("). Otherwise, if it is not surrounded by double-quotes, the directory will not be created and an error is raised. This rule also applies to any other HDFql operation that works with directory names (e.g. RENAME DIRECTORY).

## **Return**

Nothing

## **Example(s)**

```
# create a directory named "my_directory0" (the directory will not be overwritten if it already exists)
CREATE DIRECTORY my_directory0

# create a directory named "my_directory1" in a root directory named "data" (neither directory will be overwritten if they already exist; directory "data" will be created on the fly if it does not exist)
CREATE DIRECTORY /data/my_directory1

# create two directories named "my_directory2" and "my_directory3" (neither directory will be overwritten if they already exist)
CREATE DIRECTORY my_directory2, my_directory3

# create a directory named "this is a long directory name" (the directory will not be overwritten if it already exists)
CREATE DIRECTORY "this is a long directory name"
```

## 6.4.2 CREATE FILE

### Syntax

**CREATE** [**TRUNCATE**] [**PARALLEL**] **FILE** *file\_name* [, *file\_name*]\*

[**LIBRARY BOUNDS** [**FROM** {**EARLIEST** | **LATEST** | **V18** | **DEFAULT**}] [**TO** {**LATEST** | **V18** | **DEFAULT**}]

### Description

Create an HDF5 file named *file\_name*. Multiple files can be created at once by separating these with a comma (,). If *file\_name* already exists, it will not be overwritten, no subsequent files are created, and an error is raised. To overwrite an existing file, specify the keyword **TRUNCATE** (all data stored in the file will be permanently lost). In case the keyword **PARALLEL**<sup>35</sup> is specified, HDFql creates the file using all the MPI processes specified upon launching the program (that employs HDFql). In case the keyword **LIBRARY BOUNDS** is specified, HDFql creates the file using these bounds (instead of the library bounds that may have been set through the operation [SET LIBRARY BOUNDS](#)).

### Parameter(s)

*file\_name* – mandatory string that specifies the name of the HDF5 file to create. Multiple files are separated with a comma (,). As a general rule, in case *file\_name* is composed of spaces, special characters or reserved keywords (e.g. **SELECT**), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the file will not be created and an error is raised. This rule also applies to any other HDFql operation that works with file names (e.g. **RENAME FILE**).

### Return

Nothing

### Example(s)

```
# create an HDF5 file named "my_file0.h5" (the file will not be overwritten if it already exists)
CREATE FILE my_file0.h5

# create an HDF5 file named "my_file1.h5" in a root directory named "data" (the file will not be overwritten if it already exists)
```

<sup>35</sup> This option is not allowed in Windows as HDFql does not support the parallel HDF5 (PHDF5) library in this platform currently.

```

CREATE FILE /data/my_file1.h5

# create two HDF5 files named "my_file2.h5" and "my_file3.h5" (both files will be overwritten
if they already exist)
CREATE TRUNCATE FILE my_file2.h5, my_file3.h5

# create an HDF5 file named "my_file4.h5" (the file will not be overwritten if it already
exists) with the latest version of the HDF5 library
CREATE FILE my_file4.h5 LIBRARY BOUNDS FROM LATEST TO LATEST

# create an HDF5 file named "this is a long file name.h5" (the file will not be overwritten if
it already exists)
CREATE FILE "this is a long file name.h5"

# create an HDF5 file named "my_file5.h5" (the file will not be overwritten if it already
exists) in parallel (i.e. all the MPI processes specified upon launching the program (that
employs HDFql) will collectively create the file - e.g. if the program is launched as "mpiexec
-n 3 my_program", all three MPI processes will participate in the creation of the file)
CREATE PARALLEL FILE my_file5.h5

```

### 6.4.3 CREATE GROUP

#### Syntax

**CREATE** [**TRUNCATE**] **GROUP** [*file\_name*] *group\_name* [, [*file\_name*] *group\_name*]\*

[**ORDER** {**TRACKED** | **INDEXED**}]

[**STORAGE COMPACT** *object\_max\_compact* **DENSE** *object\_min\_dense*]

[**ATTRIBUTE** [**ORDER** {**TRACKED** | **INDEXED**}] [**STORAGE COMPACT** *attribute\_max\_compact* **DENSE** *attribute\_min\_dense*]]

#### Description

Create an HDF5 group named *group\_name*. Multiple groups can be created at once by separating these with a comma (,). If *group\_name* already exists, it will not be overwritten, no subsequent groups are created, and an error is raised. To overwrite an existing group, specify the keyword **TRUNCATE** (all data stored in the group will be permanently lost). In case *group\_name* has intermediate groups that do not exist, besides *group\_name* being created, all these intermediate groups will be created on the fly (e.g. when creating the group "my\_group/my\_subgroup/my\_subsubgroup", besides



“my\_subsubgroup” being created, “my\_group” and “my\_subgroup” will be created in case they do not exist). By default, *group\_name* does not track objects (i.e. groups, datasets, (soft) links or external links) stored within it by their creation order. To track the creation order of objects stored in *group\_name*, the keyword ORDER TRACKED must be specified. In case the keyword ORDER INDEXED is specified, objects stored within *group\_name* are also tracked by their creation order and using an index (to speed-up retrieval of object names). By default, *group\_name* does not track attributes stored within it by their creation order. To track the creation order of attributes stored in *group\_name*, the keyword ATTRIBUTE ORDER TRACKED must be specified. In case the keyword ATTRIBUTE ORDER INDEXED is specified, attributes stored within *group\_name* are also tracked by their creation order and using an index (to speed-up retrieval of attribute names)

### **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file in which the group is created. If *file\_name* is specified, the file is opened on the fly, the group is created within it and, afterwards, the file is closed. Otherwise, if it is not specified, the group is created in the file currently in use. As a general rule, in case *file\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the group will not be created and an error is raised. This rule also applies to any other HDFql operation that works with group names (e.g. RENAME GROUP).

*group\_name* – mandatory string that specifies the name of the HDF5 group to create. Multiple groups are separated with a comma (,). As a general rule, in case *group\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the group will not be created and an error is raised. This rule also applies to any other HDFql operation that works with group names (e.g. RENAME GROUP).

*object\_max\_compact* – optional integer that specifies the maximum number of links (i.e. objects) to store in the compact format. In case the number of links (stored in *group\_name*) exceeds *object\_max\_compact*, the storage of links switches to the dense format. If not specified, its default value is 8 (defined by the HDF5 library).

*object\_min\_dense* – optional integer that specifies the minimum number of links (i.e. objects) to store in the dense format. In case the number of links (stored in *group\_name*) falls below *object\_min\_dense*, the storage of links switches to the compact format. If not specified, its default value is 6 (defined by the HDF5 library).

*attribute\_max\_compact* – optional integer that specifies the maximum number of attributes to store in the compact format. In case the number of attributes (stored in *group\_name*) exceeds *attribute\_max\_compact*, the storage of attributes switches to the dense format. If not specified, its default value is 8 (defined by the HDF5 library).

*attribute\_min\_dense* – optional integer that specifies the minimum number of attributes to store in the dense format. In case the number of attributes (stored in *group\_name*) falls below *attribute\_min\_dense*, the storage of attributes switches to the compact format. If not specified, its default value is 6 (defined by the HDF5 library).

## **Return**

Nothing

## **Example(s)**

```
# create an HDF5 group named "my_group0" (the group will not be overwritten if it already exists)
CREATE GROUP my_group0

# create an HDF5 group named "my_subgroup0" in a root group named "my_group1" (neither group will be overwritten if they already exist; group "my_group1" will be created on the fly if it does not exist)
CREATE GROUP /my_group1/my_subgroup0

# create two HDF5 groups named "my_group2" and "my_group3" (both groups will be overwritten if they already exist)
CREATE TRUNCATE GROUP my_group2, my_group3

# create an HDF5 group named "this is a long group name" (the group will not be overwritten if it already exists)
CREATE GROUP "this is a long group name"
```

```
# create an HDF5 group named "my_group4" that tracks the objects' (i.e. groups and datasets) creation order within the group and using compact storage
CREATE GROUP my_group4 ORDER TRACKED STORAGE COMPACT 10 DENSE 7

# create an HDF5 group named "my_group5" that indexes the attributes' creation order
CREATE GROUP my_group5 ATTRIBUTE ORDER INDEXED
```

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5

# create an HDF5 group named "my_group6" in the HDF5 file currently in use (i.e. file "my_file.h5")
```

```

CREATE GROUP my_group6

# close HDF5 file currently in use (i.e. file "my_file.h5")
CLOSE FILE

# create an HDF5 group named "my_group7" in file "my_file.h5"
CREATE GROUP my_file.h5 my_group7

```

## 6.4.4 CREATE DATASET

### Syntax

```

CREATE [TRUNCATE] [EARLY | INCREMENTAL | LATE] [CONTIGUOUS | COMPACT | {CHUNKED [(chunk_dim [,
chunk_dim]*)]}] DATASET [file_name] dataset_name [, [file_name] dataset_name]* AS [NATIVE | LITTLE ENDIAN |
BIG ENDIAN | ASCII | UTF8] {TINYINT | UNSIGNED TINYINT | SMALLINT | UNSIGNED SMALLINT | INT |
UNSIGNED INT | BIGINT | UNSIGNED BIGINT | FLOAT | DOUBLE | CHAR | VARTINYINT | UNSIGNED
VARTINYINT | VARSMALLINT | UNSIGNED VARSMALLINT | VARINT | UNSIGNED VARINT | VARBIGINT |
UNSIGNED VARBIGINT | VARFLOAT | VARDOUBLE | VARCHAR | OPAQUE} [(UNLIMITED | {dataset_dim [TO
dataset_max_dim | UNLIMITED]})] [, UNLIMITED | {dataset_dim [TO {dataset_max_dim | UNLIMITED}]})]*]

[TAG tag_value]

[FILL {(fill_value [, fill_value]* ) | UNDEFINED}]

[ATTRIBUTE [ORDER {TRACKED | INDEXED}] [STORAGE COMPACT attribute_max_compact DENSE
attribute_min_dense]]

[ENABLE [NBIT PRECISION nbit_precision_value OFFSET nbit_offset_value] [SCALEOFFSET scaleoffset_value]
[SHUFFLE] [ZLIB [LEVEL zlib_level]] [FLETCHER32]]

[VALUES {(initial_value [, initial_value]* ) | input_redirecting_option}]

```

### Description

Create an HDF5 dataset named *dataset\_name*. Multiple datasets can be created at once by separating these with a comma (,). If *dataset\_name* already exists, it will not be overwritten, no subsequent datasets are created, and an error is raised. To overwrite an existing dataset, specify the keyword **TRUNCATE** (all data stored in the dataset will be permanently lost). In case *dataset\_name* has intermediate groups that do not exist, besides *dataset\_name* being created, all these intermediate groups will be created on the fly (e.g. when creating the dataset “my\_group/my\_subgroup/my\_dataset”,

besides “my\_dataset” being created as a dataset, “my\_group” and “my\_subgroup” will be created as groups in case they do not exist). By default, *dataset\_name* does not track attributes stored within it by their creation order. To track the creation order of attributes stored in *dataset\_name*, the keyword `ATTRIBUTE ORDER TRACKED` must be specified. In case the keyword `ATTRIBUTE ORDER INDEXED` is specified, attributes stored within *dataset\_name* are also tracked by their creation order and using an index (to speed-up retrieval of attribute names).

By default, if no storage type (layout) is specified and (1) the dataset is not extendible and (2) no HDF5 pre-defined filter is used, the dataset will be created as contiguous. To specify a certain storage type (layout), one of the following keywords may be employed:

- `CONTIGUOUS` – the data is stored in the HDF5 file in one contiguous block.
- `COMPACT` – the data is stored in the object header of the dataset. This storage type (layout) should only be used for data with a size limit of 65520 bytes (otherwise an error is raised).
- `CHUNKED` – the data is stored in equal-sized blocks or chunks of a pre-defined size. This storage type (layout) should be used when the dataset is extendible and/or HDF5 pre-defined filters are specified (otherwise an error is raised).

By default, if no storage allocation is specified, the dataset will have an early, incremental or late storage allocation depending on whether its storage type (layout) is compact, chunked or contiguous, respectively. To specify a certain storage allocation, one of the following keywords may be employed:

- `EARLY` – the space necessary to store the entire dataset is immediately allocated (i.e. reserved) in the HDF5 file.
- `INCREMENTAL` – the space necessary to store the dataset is incrementally allocated (i.e. reserved) according to the ongoing needs in the HDF5 file.
- `LATE` – the space necessary to store the entire dataset is only allocated (i.e. reserved) in the HDF5 file when data is written into the dataset for the first time.

To create an extendible dataset<sup>36</sup>, the keyword `TO` may be employed when specifying the dimensions that are extendible (i.e. that can grow) along with the initial size of the dimension (*dataset\_dim*) and the maximum size (*dataset\_max\_dim*) that it may grow to. If a dimension is expected to grow infinitely, the keyword `UNLIMITED` should be specified. Of note, when a dimension has an initial size of one and is expected to grow infinitely, the keyword `TO` along with *dataset\_dim* and *dataset\_max\_dim* may simply be replaced by the keyword `UNLIMITED`.

---

<sup>36</sup> An extendible HDF5 dataset is one whose one or more dimensions can grow. These dimensions start with an initial size and may be increased in a later stage. To be able to create an extendible dataset, the storage type (layout) of the dataset must be chunked (otherwise an error is raised). In case the storage type (layout) is not specified, HDFqI will automatically set it to chunked and calculate an appropriate value for the chunk size.

In case the keyword ENABLE is specified, one or more HDF5 pre-defined filters<sup>37</sup> may be used to create a (linear) pipeline by additionally specifying one or more of the following keywords:

- **NBIT** – Compresses the data of an n-bit data type (including arrays and the n-bit fields of compound data types) by packing n-bit data on output (i.e. stripping off all unused bits) and unpacking on input (i.e. restoring the extra bits required by the computation). This filter may only be used for integer and floating-point data types (otherwise an error is raised).
- **SCALEOFFSET** – Compresses the data by performing a scale and/or offset operation on each element and truncates the result to a minimum number of bits. This filter may only be used for integer and floating-point data types (otherwise an error is raised).
- **SHUFFLE** – Rearranges the bytes in the chunk by de-interlacing a block of data, which may lead to a better compression ratio. This filter is usually used in conjunction with the ZLIB filter.
- **ZLIB** – Compresses the data using the ZLIB library which is based on the Deflate lossless data compression algorithm.
- **FLETCHER32** – Adds a checksum to each chunk to detect data corruption. In case a chunk gets corrupted, any attempt to read it afterwards will raise an error.

### **Parameter(s)**

*chunk\_dim* – optional integer that specifies the chunk size of the dimension in question. Multiple chunk sizes are separated with a comma (,). If *chunk\_dim* is specified it must be equal to or greater than one (otherwise an error will be raised). Otherwise, if it is not specified and in case the keyword CHUNKED is specified, HDFqI will automatically calculate an appropriate value<sup>38</sup> and assign it to *chunk\_dim*.

*file\_name* – optional string that specifies the name of the HDF5 file in which the dataset is created. If *file\_name* is specified, the file is opened on the fly, the dataset is created within it and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset is created in the file currently in use. As a general rule, in case *file\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the dataset will not be created and an error is raised. This rule also applies to any other HDFqI operation that works with dataset names (e.g. RENAME DATASET).

---

<sup>37</sup> To be able use HDF5 pre-defined filters the storage type (layout) of the HDF5 dataset must be chunked (otherwise an error is raised). In case the storage type (layout) is not specified, HDFqI will automatically set it to chunked and calculate an appropriate value for the chunk size.

<sup>38</sup> This calculated value may not be optimal as it is based on a best guess approach with the main purpose of alleviating the programmer from specifying it. In case performance is critical, the chunk size of the dimension in question should be explicitly specified taking into account how the data (stored in the HDF5 dataset) is accessed as it greatly influences performance (HDFqI does not have enough information on how this access is ultimately done).

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset to create. Multiple datasets are separated with a comma (,). As a general rule, in case *dataset\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“). Otherwise, if it is not surrounded by double-quotes, the dataset will not be created and an error is raised. This rule also applies to any other HDFql operation that works with dataset names (e.g. RENAME DATASET).

*dataset\_dim* – optional integer that specifies the size of the dimension. Multiple dimensions are separated with a comma (,). If not specified, the size of the dimension is zero.

*dataset\_max\_dim* – optional integer that specifies the maximum size of the dimension. Multiple dimensions are separated with a comma (,). To specify an unlimited size, the keyword UNLIMITED should be specified for this purpose. If *dataset\_max\_dim* is specified it must be equal to or greater than *dataset\_dim* and the keyword CHUNKED should be specified (otherwise an error will be raised).

*tag\_value* – optional string that specifies the value of a tag attached to the HDF5 dataset. Of note, the specification of a tag is only available for a dataset of data type [HDFQL\\_OPAQUE](#) (any other data type will raise an error).

*fill\_value* – optional integer, float or string that specifies the (default) value to return in case of reading the HDF5 dataset when no data has ever been written into it. Multiple fill values are separated with a comma (,). If not specified, the dataset will be zeroed or emptied depending on whether the dataset is a number or a string, respectively.

*attribute\_max\_compact* – optional integer that specifies the maximum number of attributes to store in the compact format. In case the number of attributes (stored in *dataset\_name*) exceeds *attribute\_max\_compact*, the storage of attributes switches to the dense format. If not specified, its default value is 8 (defined by the HDF5 library).

*attribute\_min\_dense* – optional integer that specifies the minimum number of attributes to store in the dense format. In case the number of attributes (stored in *dataset\_name*) falls below *attribute\_min\_dense*, the storage of attributes switches to the compact format. If not specified, its default value is 6 (defined by the HDF5 library).

*nbit\_precision\_value* – optional integer that specifies the precision of the N-bit filter.

*nbit\_offset\_value* – optional integer that specifies the offset of the N-bit filter.

*scaleoffset\_value* – optional integer that specifies the offset of the scale-offset filter. The *scaleoffset\_value* must be equal to or greater than zero (otherwise an error is raised). In case the HDF5 dataset is of integer data type, *scaleoffset\_value* specifies the number of bits to retain (of note, if *scaleoffset\_value* is zero, the HDF5 library automatically calculates the number of bits required for lossless compression). In case the dataset is of floating-point data type, *scaleoffset\_value* specifies the number of digits after the decimal point to retain.

*zlib\_level* – optional integer that specifies the compression level of the ZLIB filter. The *zlib\_level* must be between 0 (no compression) and 9 (best compression) (otherwise an error is raised). If not specified and in case the keyword ZLIB is specified, its default value is 9.

*initial\_value* – optional integer, float or string to write into the created HDF5 dataset. Multiple values are separated with a comma (,).

*input\_redirecting\_option* – optional option that specifies a file or memory to read data from in order to write it into the created HDF5 dataset (please refer to the subsection [FROM](#) for additional information).

## **Return**

Nothing

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type int (the dataset will not be
overwritten if it already exists)
CREATE DATASET my_dataset0 AS INT

# create an HDF5 dataset named "my_dataset1" of data type char in a root group named "my_group"
(the dataset will not be overwritten if it already exists)
CREATE DATASET /my_group/my_dataset1 AS CHAR

# create two HDF5 datasets named "my_dataset2" and "my_dataset3" of data type short (both
datasets will be overwritten if they already exist)
CREATE TRUNCATE DATASET my_dataset2, my_dataset3 AS SMALLINT

# create an HDF5 dataset named "this is a long dataset name" of data type float (the dataset
will not be overwritten if it already exists)
CREATE DATASET "this is a long dataset name" AS FLOAT
```

```
# create an HDF5 dataset named "my_dataset4" of data type unsigned long long using the big
endian representation
CREATE DATASET my_dataset4 AS BIG ENDIAN UNSIGNED BIGINT

# create an HDF5 dataset named "my_dataset5" of data type int using the little endian
representation with an initial value of 80178
CREATE DATASET my_dataset5 AS LITTLE ENDIAN INT VALUES(80178)
```

```
# create an HDF5 dataset named "my_dataset6" of data type char using an ASCII representation
CREATE DATASET my_dataset6 AS ASCII CHAR
```

```
# create an HDF5 dataset named "my_dataset7" of data type float of one dimension (size 1024)
with a fill value (i.e. default) of 85.2
```

```
CREATE DATASET my_dataset7 AS FLOAT(1024) FILL(85.2)
```

```
# create a compact HDF5 dataset named "my_dataset8" of data type double of three dimensions
(size 2x5x10)
```

```
CREATE COMPACT DATASET my_dataset8 AS DOUBLE(2, 5, 10)
```

```
# create a chunked (size 20x100) HDF5 dataset named "my_dataset9" of data type unsigned char of
two dimensions (size 500x1000)
```

```
CREATE CHUNKED(20, 100) DATASET my_dataset9 AS UNSIGNED TINYINT(500, 1000)
```

```
# create an HDF5 dataset named "my_dataset10" of data type int of two dimensions (size 20x400)
using the N-bit data compression filter
```

```
CREATE DATASET my_dataset10 AS INT(20, 400) ENABLE NBIT PRECISION 16 OFFSET 4
```

```
# create an HDF5 dataset named "my_dataset11" of data type float of one dimension (size 500000)
using both the ZLIB data compression and Fletcher32 checksum error detection filters
```

```
CREATE DATASET my_dataset11 AS FLOAT(500000) ENABLE ZLIB LEVEL 5 FLETCHER32
```

```
# create an HDF5 dataset named "my_dataset12" of data type variable-length float
```

```
CREATE DATASET my_dataset12 AS VARFLOAT
```

```
# create an HDF5 dataset named "my_dataset13" of data type variable-length short of one
dimension (size 5) with initial values from a text file named "my_file.txt"
```

```
CREATE DATASET my_dataset13 AS VARSMALLINT(5) VALUES FROM FILE my_file.txt
```

```
# create an HDF5 dataset named "my_dataset14" of data type variable-length char with an initial
value of "Hierarchical Data Format"
```

```
CREATE DATASET my_dataset14 AS VARCHAR VALUES("Hierarchical Data Format")
```

```
# create an HDF5 dataset named "my_dataset15" of data type opaque
```

```
CREATE DATASET my_dataset15 AS OPAQUE
```



```
# create an HDF5 dataset named "my_dataset16" of data type opaque of one dimension (size 6)
with initial (ASCII) values of 72, 68, 70, 0, 113 and 108 (i.e. "HDF\0ql")
```

```
CREATE DATASET my_dataset16 AS OPAQUE(6) VALUES(72, 68, 70, 0, 113, 108)
```

```
# create an HDF5 dataset named "my_dataset17" of data type opaque of two dimensions (size
10x1024) with a tag value "Raw data"
```

```
CREATE DATASET my_dataset17 AS OPAQUE(10, 1024) TAG "Raw data"
```

```
# create a chunked (size 2) HDF5 dataset named "my_dataset18" of data type float of one
dimension (size 5 and extendible up to 10)
```

```
CREATE CHUNKED(2) DATASET my_dataset18 AS FLOAT(5 TO 10)
```

```
# create a chunked (with an automatically calculated size) HDF5 dataset named "my_dataset19" of
data type variable-length int of one dimension (size 1 and extendible to an unlimited size)
```

```
CREATE CHUNKED DATASET my_dataset19 AS VARINT(UNLIMITED)
```

```
# create a chunked (with an automatically calculated size) HDF5 dataset named "my_dataset20" of
data type double of three dimensions (first dimension with size 3 and extendible up to 5;
second dimension with size 7; third dimension with size 20 and extendible to an unlimited size)
```

```
CREATE CHUNKED DATASET my_dataset20 AS DOUBLE(3 TO 5, 7, 20 TO UNLIMITED)
```

```
# use (i.e. open) an HDF5 file named "my_file.h5"
```

```
USE FILE my_file.h5
```

```
# create an HDF5 dataset named "my_dataset21" of data type double in the HDF5 file currently in
use (i.e. file "my_file.h5")
```

```
CREATE DATASET my_dataset21 AS DOUBLE
```

```
# close HDF5 file currently in use (i.e. file "my_file.h5")
```

```
CLOSE FILE
```

```
# create an HDF5 dataset named "my_dataset22" of data type int in file "my_file.h5"
```

```
CREATE DATASET my_file.h5 my_dataset22 AS INT
```

## 6.4.5 CREATE ATTRIBUTE

### Syntax

```
CREATE [TRUNCATE] ATTRIBUTE [file_name] attribute_name [, [file_name] attribute_name]* AS [NATIVE | LITTLE
ENDIAN | BIG ENDIAN | ASCII | UTF8] {TINYINT | UNSIGNED TINYINT | SMALLINT | UNSIGNED SMALLINT | INT
| UNSIGNED INT | BIGINT | UNSIGNED BIGINT | FLOAT | DOUBLE | CHAR | VARTINYINT | UNSIGNED
VARTINYINT | VARSMALLINT | UNSIGNED VARSMALLINT | VARINT | UNSIGNED VARINT | VARBIGINT |
UNSIGNED VARBIGINT | VARFLOAT | VARDOUBLE | VARCHAR | OPAQUE} [(attribute_dim [, attribute_dim]*)]

[TAG tag_value]

[VALUES {(initial_value [, initial_value]*)} | input_redirecting_option]
```

### Description

Create an HDF5 attribute named *attribute\_name*. Multiple attributes can be created at once by separating these with a comma (,). If *attribute\_name* already exists, it will not be overwritten, no subsequent attributes are created, and an error is raised. To overwrite an existing attribute, specify the keyword TRUNCATE (all data stored in the attribute will be permanently lost).

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file in which the attribute is created. If *file\_name* is specified, the file is opened on the fly, the attribute is created within it and, afterwards, the file is closed. Otherwise, if it is not specified, the attribute is created in the file currently in use. As a general rule, in case *file\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the attribute will not be created and an error is raised. This rule also applies to any other HDFql operation that works with attribute names (e.g. RENAME ATTRIBUTE).

*attribute\_name* – mandatory string that specifies the name of the HDF5 attribute to create. Multiple attributes are separated with a comma (,). As a general rule, in case *attribute\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the attribute will not be created and an error is raised. This rule also applies to any other HDFql operation that works with attribute names (e.g. RENAME ATTRIBUTE).

*attribute\_dim* – optional integer that specifies the size of the dimension. Multiple dimensions are separated with a comma (,).

*tag\_value* – optional string that specifies the value of a tag attached to the HDF5 attribute. Of note, the specification of a tag is only available for an attribute of data type [HDFQL\\_OPAQUE](#) (any other data type will raise an error).

*initial\_value* – optional integer, float or string to write into the created HDF5 attribute. Multiple values are separated with a comma (,). In case *initial\_value* is not specified, the element in question will be zeroed or emptied depending on whether the attribute is a number or a string, respectively.

*input\_redirecting\_option* – optional option that specifies a file or memory to read data from in order to write it into the created HDF5 attribute (please refer to the subsection [FROM](#) for additional information).

## **Return**

Nothing

## **Example(s)**

```
# create an HDF5 attribute named "my_attribute0" of data type int (the attribute will not be
overwritten if it already exists)
CREATE ATTRIBUTE my_attribute0 AS INT

# create an HDF5 attribute named "my_attribute1" of data type char in a root object (either a
group or dataset) named "my_object" (the attribute will not be overwritten if it already
exists)
CREATE ATTRIBUTE /my_object/my_attribute1 AS CHAR

# create two HDF5 attributes named "my_attribute2" and "my_attribute3" of data type short (both
attributes will be overwritten if they already exist)
CREATE TRUNCATE ATTRIBUTE my_attribute2, my_attribute3 AS SMALLINT

# create an HDF5 attribute named "this is a long attribute name" of data type float (the
attribute will not be overwritten if it already exists)
CREATE ATTRIBUTE "this is a long attribute name" AS FLOAT
```

```
# create an HDF5 attribute named "my_attribute4" of data type unsigned long long using the big
endian representation
CREATE ATTRIBUTE my_attribute4 AS BIG ENDIAN UNSIGNED BIGINT

# create an HDF5 attribute named "my_attribute5" of data type int using the little endian
representation with an initial value of 80178
CREATE ATTRIBUTE my_attribute5 AS LITTLE ENDIAN INT VALUES(80178)
```

```
# create an HDF5 attribute named "my_attribute6" of data type char using an UTF8 representation
CREATE ATTRIBUTE my_attribute6 AS UTF8 CHAR
```

```
# create an HDF5 attribute named "my_attribute7" of data type float of one dimension (size 512)
CREATE ATTRIBUTE my_attribute7 AS FLOAT(512)

# create an HDF5 attribute named "my_attribute8" of data type unsigned char of two dimensions
(size 500x1000)
CREATE ATTRIBUTE my_attribute8 AS UNSIGNED TINYINT(500, 1000)
```

```
# create an HDF5 attribute named "my_attribute9" of data type variable-length float
CREATE ATTRIBUTE my_attribute9 AS VARFLOAT

# create an HDF5 attribute named "my_attribute10" of data type variable-length short of one
dimension (size 5) with initial values from a text file named "my_file.txt"
CREATE ATTRIBUTE my_attribute10 AS VARSMALLINT(5) VALUES FROM FILE my_file.txt

# create an HDF5 attribute named "my_attribute11" of data type variable-length char with an
initial value of "Hierarchical Data Format"
CREATE ATTRIBUTE my_attribute11 AS VARCHAR VALUES("Hierarchical Data Format")
```

```
# create an HDF5 attribute named "my_attribute12" of data type opaque
CREATE ATTRIBUTE my_attribute12 AS OPAQUE

# create an HDF5 attribute named "my_attribute13" of data type opaque of one dimension (size 6)
with initial (ASCII) values 72, 68, 70, 0, 113 and 108 (i.e. "HDF\0q1")
CREATE ATTRIBUTE my_attribute13 AS OPAQUE(6) VALUES(72, 68, 70, 0, 113, 108)

# create an HDF5 attribute named "my_attribute14" of data type opaque of two dimensions (size
10x1024) with a tag value "Raw data"
CREATE ATTRIBUTE my_attribute14 AS OPAQUE(10, 1024) TAG "Raw data"
```

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5
```

```
# create an HDF5 attribute named "my_attribute15" of data type double in the HDF5 file
currently in use (i.e. file "my_file.h5")
CREATE ATTRIBUTE my_attribute15 AS DOUBLE

# close HDF5 file currently in use (i.e. file "my_file.h5")
CLOSE FILE

# create an HDF5 attribute named "my_attribute16" of data type int in file "my_file.h5"
CREATE ATTRIBUTE my_file.h5 my_attribute16 AS INT
```

## 6.4.6 CREATE [SOFT | HARD] LINK

### Syntax

CREATE [TRUNCATE] [SOFT | HARD] LINK [*file\_name*] *link\_name* [, [*file\_name*] *link\_name*]\* TO *object\_name* [, *object\_name*]\*

### Description

Create an HDF5 soft or hard link named *link\_name* to a group or dataset named *object\_name*. Multiple links can be created at once by separating these with a comma (,). If *link\_name* already exists, it will not be overwritten, no subsequent links are created, and an error is raised. To overwrite an existing link, specify the keyword TRUNCATE. If neither the keyword SOFT nor HARD is specified, a soft link is created by default. To create a hard link, the keyword HARD must be specified.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file in which the soft or hard link is created. If *file\_name* is specified, the file is opened on the fly, the soft or hard link is created within it and, afterwards, the file is closed. Otherwise, if it is not specified, the soft or hard link is created in the file currently in use. As a general rule, in case *file\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the link will not be created and an error is raised. This rule also applies to any other HDFqL operation that works with link names (e.g. RENAME LINK).

*link\_name* – mandatory string that specifies the name of the HDF5 soft or hard link to create. Multiple links are separated with a comma (,). As a general rule, in case *link\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“”). Otherwise, if it is not surrounded by double-quotes, the link will

not be created and an error is raised. This rule also applies to any other HDFqL operation that works with link names (e.g. RENAME LINK).

*object\_name* – mandatory string that specifies the name of the HDF5 group or dataset that *link\_name* points to. Multiple objects are separated with a comma (.). As a general rule, in case *object\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes (“). Otherwise, if it is not surrounded by double-quotes, the link will not be created and an error is raised. This rule also applies to any other HDFqL operation that works with link names (e.g. RENAME LINK).

## **Return**

Nothing

## **Example(s)**

```
# create an HDF5 group named "my_group0"
CREATE GROUP my_group0

# create an HDF5 dataset named "my_dataset0" of data type variable-length unsigned int
CREATE DATASET my_dataset0 AS UNSIGNED VARINT

# create an HDF5 soft link named "my_link0" to group "my_group0" (the soft link will not be
overwritten if it already exists)
CREATE LINK my_link0 TO my_group0

# create an HDF5 soft link named "my_link1" to dataset "my_dataset0" (the soft link will not be
overwritten if it already exists)
CREATE SOFT LINK my_link1 TO my_dataset0

# create two HDF5 soft links named "my_link2" and "my_link3" to dataset "my_dataset0" and group
"my_group0" respectively (both soft links will be overwritten if they already exist)
CREATE TRUNCATE SOFT LINK my_link2, my_link3 TO my_dataset0, my_group0

# create an HDF5 soft link named "this is a long link name" to dataset "my_dataset0" (the soft
link will not be overwritten if it already exists)
CREATE LINK "this is a long link name" TO my_dataset0
```

```
# create an HDF5 group named "my_group1"
CREATE GROUP my_group1
```

```
# create an HDF5 dataset named "my_dataset1" of data type variable-length unsigned int
CREATE DATASET my_dataset1 AS UNSIGNED VARINT

# create an HDF5 hard link named "my_link4" to group "my_group1" (the hard link will not be
overwritten if it already exists)
CREATE HARD LINK my_link4 TO my_group1

# create an HDF5 hard link named "my_link5" to dataset "my_dataset1" (the hard link will not be
overwritten if it already exists)
CREATE HARD LINK my_link5 TO my_dataset1

# create two HDF5 hard links named "my_link6" and "my_link7" to dataset "my_dataset1" and group
"my_group1" respectively (both hard links will be overwritten if they already exist)
CREATE TRUNCATE HARD LINK my_link6, my_link7 TO my_dataset1, my_group1
```

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5

# create an HDF5 soft link named "my_link8" to object "my_object0" in the HDF5 file currently
in use (i.e. file "my_file.h5")
CREATE LINK my_link8 TO my_object0

# close HDF5 file currently in use (i.e. file "my_file.h5")
CLOSE FILE

# create an HDF5 soft link named "my_link9" to object "my_object1" in file "my_file.h5"
CREATE LINK my_file.h5 my_link9 TO my_object1
```

## 6.4.7 CREATE EXTERNAL LINK

### Syntax

**CREATE** [TRUNCATE] **EXTERNAL LINK** [*file\_name*] *link\_name* [, [*file\_name*] *link\_name*]\* **TO** *target\_file\_name*  
*object\_name* [, *target\_file\_name* *object\_name*]\*

## **Description**

Create an HDF5 external link named *link\_name* to a group or dataset named *object\_name* belonging to another HDF5 file named *target\_file\_name*. Multiple external links can be created at once by separating these with a comma (,). If *link\_name* already exists, it will not be overwritten, no subsequent external links are created, and an error is raised. To overwrite an existing external link, specify the keyword TRUNCATE.

## **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file in which the external link is created. If *file\_name* is specified, the file is opened on the fly, the external link is created within it and, afterwards, the file is closed. Otherwise, if it is not specified, the external link is created in the file currently in use. As a general rule, in case *file\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes ("). Otherwise, if it is not surrounded by double-quotes, the link will not be created and an error is raised. This rule also applies to any other HDFqL operation that works with external link names (e.g. RENAME EXTERNAL LINK).

*link\_name* – mandatory string that specifies the name of the HDF5 external link to create. Multiple external links are separated with a comma (,). As a general rule, in case *link\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes ("). Otherwise, if it is not surrounded by double-quotes, the link will not be created and an error is raised. This rule also applies to any other HDFqL operation that works with external link names (e.g. RENAME EXTERNAL LINK).

*target\_file\_name* – mandatory string that specifies the name of the HDF5 file where *object\_name* is stored and *link\_name* points to. Multiple files are separated with a comma (,). As a general rule, in case *target\_file\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes ("). Otherwise, if it is not surrounded by double-quotes, the link will not be created and an error is raised. This rule also applies to any other HDFqL operation that works with link external names (e.g. RENAME EXTERNAL LINK).

*object\_name* – mandatory string that specifies the name of the HDF5 group or dataset (stored in *file\_name*) that *link\_name* points to. As a general rule, in case *object\_name* is composed of spaces, special characters or reserved keywords (e.g. SELECT), it should be surrounded by double-quotes ("). Otherwise, if it is not surrounded by double-quotes, the link will not be created and an error is raised. This rule also applies to any other HDFqL operation that works with external link names (e.g. RENAME EXTERNAL LINK).

## **Return**

Nothing



## Example(s)

```
# use (i.e. open) an HDF5 file named "my_file0.h5"
USE FILE my_file0.h5

# create an HDF5 group named "my_group0"
CREATE GROUP my_group0

# create an HDF5 dataset named "my_dataset0" of data type variable-length unsigned int
CREATE DATASET my_dataset0 AS UNSIGNED VARINT

# create an HDF5 external link named "my_link0" to object "my_group0" in file "my_file0.h5"
(the external link will not be overwritten if it already exists)
CREATE EXTERNAL LINK my_link0 TO my_file0.h5 my_group0

# create an HDF5 external link named "my_link1" to object "my_object0" in file "my_file1.h5"
(the external link will be overwritten if it already exists)
CREATE TRUNCATE EXTERNAL LINK my_link1 TO my_file1.h5 my_object0

# create two HDF5 external links named "my_link2" and "my_link3" to object "my_object1" in file
"my_file1.h5" and object "my_object2" in file "my_file2.h5" respectively (neither external
links will be overwritten if they already exist)
CREATE EXTERNAL LINK my_link2, my_link3 TO my_file1.h5 my_object1, my_file2.h5 my_object2

# create an HDF5 external link named "this is a long external link name" to object "my_object3"
in file "my_file3.h5" (the external link will not be overwritten if it already exists)
CREATE EXTERNAL LINK "this is a long external link name" TO my_file3.h5 my_object3
```

```
# use (i.e. open) an HDF5 file named "my_file4.h5"
USE FILE my_file4.h5

# create an HDF5 external link named "my_link4" in the HDF5 file currently in use (i.e. file
"my_file4.h5") to object "my_object4" in file "my_file5.h5"
CREATE EXTERNAL LINK my_link4 TO my_file5.h5 my_object4

# close HDF5 file currently in use (i.e. file "my_file4.h5")
CLOSE FILE

# create an HDF5 external link named "my_link5" in file "my_file4.h5" to object "my_object5" in
file "my_file6.h5"
CREATE EXTERNAL LINK my_file4.h5 my_link5 TO my_file6.h5 my_object5
```

## 6.4.8 ALTER DIMENSION

### Syntax

**ALTER DIMENSION** [*file\_name*] *dataset\_name* [, [*file\_name*] *dataset\_name*]\* **TO** *dataset\_dim* [, *dataset\_dim*]\*

### Description

Alter (i.e. change) the dimensions of an existing HDF5 dataset named *dataset\_name*. Multiple datasets can have their dimensions altered at once by separating these with a comma (,). If *dataset\_name* was not found or its dimensions could not be altered (due to its storage type not being [HDFQL\\_CHUNKED](#) or for unknown/unexpected reasons), no subsequent datasets are altered, and an error is raised.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset to alter (i.e. change) dimensions. If *file\_name* is specified, the file is opened on the fly, the dimensions of the dataset are altered and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset to alter the dimensions is stored in the file currently in use.

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset whose dimensions are to be altered (i.e. changed). Multiple datasets are separated with a comma (,).

*dataset\_dim* – mandatory integer that specifies the new size for the dimension in question. Multiple dimensions are separated with a comma (,). Depending on the prefix of the value specified in *dataset\_dim*, one of the following behaviors applies:

- If its prefix is "+", the dimension will have its size increased by this value.
- If its prefix is "-", the dimension will have its size decreased by this value.
- In case its prefix is neither "+" nor "-", the dimension will have exactly the size of this value.

To preserve the value of a certain dimension (i.e. for its size not to be altered), it should be skipped with a comma (,).

### Return

Nothing

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset" of data type double of three dimensions (first
dimension with size 2 and extendible up to 10; second dimension with size 7; third dimension
with size 20 and extendible to an unlimited size)
CREATE CHUNKED DATASET my_dataset AS DOUBLE(2 TO 10, 7, 20 TO UNLIMITED)

# show (i.e. get) current dimensions of dataset "my_dataset" (should be 2, 7, 20)
SHOW DIMENSION my_dataset

# alter (i.e. change) dimensions of dataset "my_dataset" to set its first dimension size to 6,
and increase the third dimension size by 10 (the second dimension size remains intact)
ALTER DIMENSION my_dataset TO 6, , +10

# show (i.e. get) current dimensions of dataset "my_dataset" (should be 6, 7, 30)
SHOW DIMENSION my_dataset

# alter (i.e. change) dimensions of dataset "my_dataset" to increase its first dimension size
by 2, to set the second dimension size to 3, and to decrease the third dimension size by 5
ALTER DIMENSION my_dataset TO +2, 3, -5

# show (i.e. get) current dimensions of dataset "my_dataset" (should be 8, 3, 25)
SHOW DIMENSION my_dataset
```

## **6.4.9 RENAME DIRECTORY**

### **Syntax**

**RENAME DIRECTORY** *directory\_name* [, *directory\_name*]\* **AS** *new\_directory\_name* [, *new\_directory\_name*]\*

### **Description**

Rename (or move) an existing directory named *directory\_name* as *new\_directory\_name*. Multiple directories can be renamed (or moved) at once by separating these with a comma (,). If *new\_directory\_name* already exists, it will not be overwritten, no subsequent directories are renamed (or moved), and an error is raised.

### **Parameter(s)**

*directory\_name* – mandatory string that specifies the name of the directory to rename (or move). Multiple directories are separated with a comma (,).

*new\_directory\_name* – mandatory string that specifies the new name and/or location (in the file system) to use for renaming and/or moving *directory\_name*. Multiple directories are separated with a comma (,).

## **Return**

Nothing

## **Example(s)**

```
# rename a directory named "my_directory0" as "my_directory1" (the directory "my_directory1"
will not be overwritten if it already exists)
RENAME DIRECTORY my_directory0 AS my_directory1

# rename two directories named "my_directory2" and "my_directory3" as "my_directory4" and
"my_directory5" respectively (neither directory will be overwritten if it already exists)
RENAME DIRECTORY my_directory2, my_directory3 AS my_directory4, my_directory5

# move a directory named "my_directory6" into a root directory named "data" and rename it as
"my_directory7" (the directory "my_directory7" will not be overwritten if it already exists)
RENAME DIRECTORY my_directory6 AS /data/my_directory7

# move a directory named "my_directory8" into a relative directory named "backup" (the
directory "my_directory8" will not be overwritten if it already exists)
RENAME DIRECTORY my_directory8 AS backup/
```

## **6.4.10 RENAME FILE**

### **Syntax**

**RENAME** [**TRUNCATE**] **FILE** *file\_name* [, *file\_name*]\* **AS** *new\_file\_name* [, *new\_file\_name*]\*

### **Description**

Rename (or move) an existing file named *file\_name* as *new\_file\_name*. Multiple files can be renamed (or moved) at once by separating these with a comma (,). If *new\_file\_name* already exists, it will not be overwritten, no subsequent files are renamed (or moved), and an error is raised. To overwrite an existing file, specify the keyword **TRUNCATE** (all data stored in the file will be permanently lost).

## **Parameter(s)**

*file\_name* – mandatory string that specifies the name of the file to rename (or move). Multiple files are separated with a comma (,).

*new\_file\_name* – mandatory string that specifies the new name and/or location (in the file system) to use for renaming and/or moving *file\_name*. Multiple files are separated with a comma (,).

## **Return**

Nothing

## **Example(s)**

```
# rename a file named "my_file0.h5" as "my_file1.h5" (the file "my_file1.h5" will not be
overwritten if it already exists)
RENAME FILE my_file0.h5 AS my_file1.h5

# rename a file named "my_file2.h5" as "my_file3.h5" (the file "my_file3.h5" will be
overwritten if it already exists)
RENAME TRUNCATE FILE my_file2.h5 AS my_file3.h5

# rename two files named "my_file4.h5" and "my_file5.h5" as "my_file6.h5" and "my_file7.h5"
respectively (both files "my_file6.h5" and "my_file7.h5" will be overwritten if they already
exist)
RENAME TRUNCATE FILE my_file4.h5, my_file5.h5 AS my_file6.h5, my_file7.h5

# move a file named "my_file8.h5" into a root directory named "data" and rename it as
"my_file9.h5" (the file "my_file9.h5" will not be overwritten if it already exists in this
directory)
RENAME FILE my_file8.h5 AS /data/my_file9.h5

# move a file named "my_file10.h5" into a relative directory named "backup" (the file
"my_file10.h5" will not be overwritten if it already exists in this directory)
RENAME FILE my_file10.h5 AS backup/
```

## 6.4.11 RENAME [GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK]

### Syntax

**RENAME** [TRUNCATE] [GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK] [*file\_name*] *object\_name* [, [*file\_name*] *object\_name*]\* **AS** *new\_object\_name* [, *new\_object\_name*]\*

### Description

Rename (or move) an existing HDF5 group, dataset, attribute, (soft) link or external link named *object\_name* as *new\_object\_name*. Multiple groups, datasets, attributes, (soft) links or external links can be renamed (or moved) at once by separating these with a comma (,). If *new\_object\_name* already exists, it will not be overwritten, no subsequent objects are renamed (or moved), and an error is raised. To overwrite an existing object, specify the keyword TRUNCATE (all data stored in the object will be permanently lost). In case (1) a group and an attribute or (2) a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword GROUP, DATASET nor ATTRIBUTE is specified, the object to be renamed is the group or dataset, respectively. To explicitly rename an object according to its type, the keyword GROUP, DATASET, ATTRIBUTE, [SOFT] LINK or EXTERNAL LINK must be specified. While the renaming (or moving) of groups and datasets to a different location is supported by the HDF5 library, this is not the case for attributes; HDFq overcomes this limitation by (1) creating a new attribute with the same characteristics as the existing one (e.g. data type, number of dimensions) using the new specified location and name, (2) writing the data from the existing attribute to the newly created attribute, and (3) deleting the existing attribute.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the object to rename (or move). If *file\_name* is specified, the file is opened on the fly, the object is renamed (or moved) and, afterwards, the file is closed. Otherwise, if it is not specified, the object to rename (or move) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the object to rename (or move). Multiple objects are separated with a comma (,).

*new\_object\_name* – mandatory string that specifies the new name and/or location (within the HDF5 file) to use for renaming and/or moving *object\_name*. Multiple objects are separated with a comma (,).

### Return

Nothing

## Example(s)

```
# create two HDF5 groups named "my_group0" and "my_group1"
CREATE GROUP my_group0, my_group1

# create two HDF5 datasets named "my_dataset" and "my_common" of data type short
CREATE DATASET my_dataset, my_common AS SMALLINT

# create two HDF5 attributes named "my_attribute" and "my_common" of data type float
CREATE ATTRIBUTE my_attribute, my_common AS FLOAT

# rename an object named "my_group0" as "my_group" (the object "my_group" will not be
overwritten if it already exists)
RENAME my_group0 AS my_group

# move an object named "my_group1" into object "my_group" and rename it as "my_subgroup" (the
object "my_subgroup" will be overwritten if it already exists in object "my_group")
RENAME TRUNCATE my_group1 AS my_group/my_subgroup

# move two objects named "my_dataset" and "my_attribute" into objects "my_group" and
"my_group/my_subgroup" respectively (both objects "my_dataset" and "my_attribute" will not be
overwritten if they already exist in objects "my_group" and "my_group/my_subgroup")
RENAME my_dataset, my_attribute AS my_group/, my_group/my_subgroup/

# rename attribute "my_common" as "my_attribute" (the attribute "my_attribute" will not be
overwritten if it already exists)
RENAME ATTRIBUTE my_common AS my_attribute

# rename dataset "my_common" as "my_dataset" (the dataset "my_dataset" will not be overwritten
if it already exists)
RENAME DATASET my_common AS my_dataset
```

## 6.4.12 COPY FILE

### Syntax

**COPY** [TRUNCATE] FILE *file\_name* [, *file\_name*]\* **TO** *new\_file\_name* [, *new\_file\_name*]\*

## **Description**

Copy an existing file named *file\_name* to *new\_file\_name*. Multiple files can be copied at once by separating these with a comma (,). If *new\_file\_name* already exists, it will not be overwritten, no subsequent files are copied, and an error is raised. To overwrite an existing file, specify the keyword TRUNCATE (all data stored in the file will be permanently lost).

## **Parameter(s)**

*file\_name* – mandatory string that specifies the name of the file to copy. Multiple files are separated with a comma (,).

*new\_file\_name* – mandatory string that specifies the new name and/or location (in the file system) to use for copying *file\_name*. Multiple files are separated with a comma (,).

## **Return**

Nothing

## **Example(s)**

```
# copy a file named "my_file0.h5" to "my_file1.h5" (the file "my_file1.h5" will not be
overwritten if it already exists)
COPY FILE my_file0.h5 TO my_file1.h5

# copy a file named "my_file2.h5" to "my_file3.h5" (the file "my_file3.h5" will be overwritten
if it already exists)
COPY TRUNCATE FILE my_file2.h5 TO my_file3.h5

# copy two files named "my_file4.h5" and "my_file5.h5" to "my_file6.h5" and "my_file7.h5"
respectively (both files "my_file6.h5" and "my_file7.h5" will be overwritten if they already
exist)
COPY TRUNCATE FILE my_file4.h5, my_file5.h5 TO my_file6.h5, my_file7.h5

# copy a file named "my_file8.h5" into a root directory named "data" and rename it as
"my_file9.h5" (the file "my_file9.h5" will not be overwritten if it already exists in this
directory)
COPY FILE my_file8.h5 TO /data/my_file9.h5

# copy a file named "my_file10.h5" into a relative directory named "backup" (the file
"my_file10.h5" will not be overwritten if it already exists in this directory)
COPY FILE my_file10.h5 TO backup/
```



### 6.4.13 COPY [GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK]

#### Syntax

**COPY** [TRUNCATE] [GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK] [*file\_name*] *object\_name* [, [*file\_name*] *object\_name*]\* **TO** [*target\_file\_name*] *new\_object\_name* [, [*target\_file\_name*] *new\_object\_name*]\*

#### Description

Copy an existing HDF5 group, dataset, attribute, (soft) link or external link named *object\_name* to *new\_object\_name*. Multiple groups, datasets, attributes, (soft) links or external links can be copied at once by separating these with a comma (,). If *new\_object\_name* already exists, it will not be overwritten, no subsequent objects are copied, and an error is raised. To overwrite an existing object, specify the keyword TRUNCATE (all data stored in the object will be permanently lost). In case (1) a group and an attribute or (2) a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword GROUP, DATASET nor ATTRIBUTE is specified, the object to be copied is the group or dataset, respectively. To explicitly copy an object according to its type, the keyword GROUP, DATASET, ATTRIBUTE, [SOFT] LINK or EXTERNAL LINK must be specified.

#### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the object to copy. If *file\_name* is specified, the file is opened on the fly, the object is copied and, afterwards, the file is closed. Otherwise, if it is not specified, the object to copy is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the object to copy. Multiple objects are separated with a comma (,).

*target\_file\_name* – optional string that specifies the name of the HDF5 file in which to copy the object. Multiple files are separated with a comma (,).

*new\_object\_name* – mandatory string that specifies the new name and/or location (within the HDF5 file or in another HDF5 file specified by *target\_file\_name*) to use for copying *object\_name*. Multiple objects are separated with a comma (,).

#### Return

Nothing

## Example(s)

```
# create two HDF5 groups named "my_group0" and "my_group1"
CREATE GROUP my_group0, my_group1

# create two HDF5 datasets named "my_dataset0" and "my_common" of data type short
CREATE DATASET my_dataset0, my_common AS SMALLINT

# create two HDF5 attributes named "my_attribute0" and "my_common" of data type float
CREATE ATTRIBUTE my_attribute0, my_common AS FLOAT

# copy an object named "my_group0" to "my_group2" (the object "my_group2" will not be
overwritten if it already exists)
COPY my_group0 TO my_group2

# copy an object named "my_group1" into object "my_group0" and rename it as "my_subgroup0" (the
object "my_subgroup0" will be overwritten if it already exists in object "my_group0")
COPY TRUNCATE my_group1 TO my_group0/my_subgroup0

# copy two objects named "my_dataset0" and "my_attribute0" into objects "my_group0" and
"my_group0/my_subgroup0" respectively (both objects "my_dataset0" and "my_attribute0" will not
be overwritten if they already exist in objects "my_group0" and "my_group0/my_subgroup0")
COPY my_dataset0, my_attribute0 TO my_group0/, my_group0/my_subgroup0/

# copy attribute "my_common" to "my_attribute1" (the attribute "my_attribute1" will not be
overwritten if it already exists)
COPY ATTRIBUTE my_common TO my_attribute1

# copy dataset "my_common" to "my_dataset1" (the dataset "my_dataset1" will not be overwritten
if it already exists)
COPY DATASET my_common TO my_dataset1
```

```
# copy an object named "my_group3" from the file currently in use to "my_group4" in an HDF5
file named "my.file0.h5" (the object "my_group4" will not be overwritten if it already exists
in the file)
COPY my_group3 TO my_file0.h5 my_group4

# copy an object named "my_group5" from an HDF5 file named "my_file1.h5" to "my_group6" in the
file currently in use (the object "my_group6" will not be overwritten if it already exists in
the file)
COPY my_file1.h5 my_group5 TO my_group6
```

```
# copy an object named "my_group7" from an HDF5 file named "my_file2.h5" to "my_group8" in an
HDF5 file named "my_file3.h5" (the object "my_group8" will not be overwritten if it already
exists in the file)
```

```
COPY my_file2.h5 my_group7 TO my_file3.h5 my_group8
```

## 6.4.14 DROP DIRECTORY

### Syntax

**DROP DIRECTORY** *directory\_name* [, *directory\_name*]\*

### Description

Drop (i.e. delete) an existing directory named *directory\_name*. Multiple directories can be dropped at once by separating these with a comma (,). If *directory\_name* contains directories or files (i.e. if it is not empty), it will not be dropped, no subsequent directories are dropped, and an error is raised.

### Parameter(s)

*directory\_name* – mandatory string that specifies the name of the directory to drop (i.e. delete). Multiple directories are separated with a comma (,).

### Return

Nothing

### Example(s)

```
# create two directories named "my_directory0" and "my_directory1" within the current working
directory
```

```
CREATE DIRECTORY my_directory0, my_directory1
```

```
# create two directories named "my_subdirectory0" and "my_subdirectory1" within the directory
"my_directory0"
```

```
CREATE DIRECTORY my_directory0/my_subdirectory0, my_directory0/my_subdirectory1
```

```
# drop (i.e. delete) directory "my_directory1" within the current working directory
```

```
DROP DIRECTORY my_directory1
```

```
# drop (i.e. delete) directory "my_subdirectory0" within directory "my_directory0"
```

```
DROP DIRECTORY my_directory0/my_subdirectory0
```

## 6.4.15 DROP FILE

### Syntax

**DROP FILE** *file\_name* [, *file\_name*]\*

### Description

Drop (i.e. delete) an existing file named *file\_name*. Multiple files can be dropped at once by separating these with a comma (,). If *file\_name* was not found or could not be dropped (due to unknown/unexpected reasons), no subsequent files are dropped, and an error is raised.

### Parameter(s)

*file\_name* – mandatory string that specifies the name of the file to drop (i.e. delete). Multiple files are separated with a comma (,).

### Return

Nothing

### Example(s)

```
# create two HDF5 files named "my_file0.h5" and "my_file1.h5" within the current working
directory
CREATE FILE my_file0.h5, my_file1.h5

# create two HDF5 files named "my_file2.h5" and "my_file3.h5" within a directory named
"my_directory"
CREATE FILE my_directory/my_file2.h5, my_directory/my_file3.h5

# drop (i.e. delete) file "my_file1.h5" within the current working directory
DROP FILE my_file1.h5

# drop (i.e. delete) file "my_file2.h5" within directory "my_directory"
DROP FILE my_directory/my_file2.h5
```

## 6.4.16 DROP [GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK]

### Syntax

```
DROP {GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK} | {[GROUP | DATASET | ATTRIBUTE |  
[SOFT] LINK | EXTERNAL LINK] [{[file_name] object_name [, [file_name] object_name]*} | {[[file_name]  
object_name] LIKE regular_expression [DEEP deep_value [, deep_value]*]}}
```

### Description

Drop (i.e. delete) an existing HDF5 group, dataset, attribute, (soft) link or external link named *object\_name*. Multiple groups, datasets, attributes, (soft) links or external links can be dropped at once by separating these with a comma (.). If *object\_name* was not found or could not be dropped (due to unknown/unexpected reasons), no subsequent objects are dropped, and an error is raised. In case (1) a group and an attribute or (2) a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword GROUP, DATASET nor ATTRIBUTE is specified, the object to be dropped is the group or dataset, respectively. To explicitly drop an object according to its type, the keyword GROUP, DATASET, ATTRIBUTE, [SOFT] LINK or EXTERNAL LINK must be specified.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the object to drop (i.e. delete). If *file\_name* is specified, the file is opened on the fly, the object is dropped and, afterwards, the file is closed. Otherwise, if it is not specified, the object to drop is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the object to drop (i.e. delete). Multiple objects are separated with a comma (.).

*regular\_expression* – optional string that specifies the regular expression which only names of objects that comply with it are dropped. If *regular\_expression* includes “\*”, recursive search is performed.

*deep\_value* – optional integer that specifies the maximum recursiveness limit.

### Return

Nothing

### Example(s)

```
# create three HDF5 groups named "my_group0", "my_group1" and "my_group2"
```

```
CREATE GROUP my_group0, my_group1, my_group2

# create two HDF5 datasets named "my_dataset0" and "my_dataset1" of data type short in group
"my_group2"
CREATE DATASET my_group2/my_dataset0, my_group2/my_dataset1 AS SMALLINT

# create two HDF5 datasets named "my_dataset2" and "my_common" of data type short
CREATE DATASET my_dataset2, my_common AS SMALLINT

# create two HDF5 attributes named "my_attribute0" and "my_common" of data type float
CREATE ATTRIBUTE my_attribute0, my_common AS FLOAT

# drop (i.e. delete) an object named "my_group0" (and all objects that may eventually be stored
in it)
DROP my_group0

# drop (i.e. delete) attribute "my_common"
DROP ATTRIBUTE my_common

# drop (i.e. delete) all existing datasets in group "my_group2" (should be "my_dataset2" and
"my_dataset3")
DROP DATASET my_group2/

# drop (i.e. delete) all existing groups (should be "my_group1" and "my_group2")
DROP GROUP

# drop (i.e. delete) all existing objects (should be "my_dataset2", "my_common" and
"my_attribute0")
DROP /
```

## 6.5 DATA MANIPULATION LANGUAGE (DML)

Data Manipulation Language (DML) is, generally speaking, syntax for defining and modifying data stored in structures. In HDFq, the DML is composed of only one operation (INSERT), which enables the insertion (i.e. writing) of data into HDF5 datasets or attributes. Moreover, it supports **REDIRECTING** options to redirect the input source according to the programmer's needs.

## 6.5.1 INSERT

### Syntax

```

INSERT [DIRECTLY [MASK mask_value] [SIZE data_size]] INTO [PARALLEL] [DATASET | ATTRIBUTE] [file_name]
object_name [{[start]:[stride]:[count]:[block] [, [start]:[stride]:[count]:[block]]* [, {OR | AND | XOR | NOTA |
NOTB} [start]:[stride]:[count]:[block] [, [start]:[stride]:[count]:[block]]*}] | {coord [, coord]* [, coord [, coord]]*}] |
{chunk_number [, chunk_number]]*}] [, [file_name] object_name [{[start]:[stride]:[count]:[block] [,
[start]:[stride]:[count]:[block]]* [, {OR | AND | XOR | NOTA | NOTB} [start]:[stride]:[count]:[block] [,
[start]:[stride]:[count]:[block]]*}] | {coord [, coord]* [, coord [, coord]]*}] | {chunk_number [, chunk_number]]*}]]*

[VALUES {{value [, value]]*} | input_redirecting_option}]

```

### Description

Insert (i.e. write) data into an HDF5 dataset or attribute named *object\_name*. Multiple datasets or attributes can be written at once by separating these with a comma (,). If *object\_name* was not found or could not be written (due to unknown/unexpected reasons), no subsequent objects are written, and an error is raised. In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword DATASET nor ATTRIBUTE is specified, the object that will have data inserted into it is the dataset. To explicitly insert data into an object according to its type, the keyword DATASET or ATTRIBUTE must be specified. In case the keyword DIRECTLY<sup>39</sup> is specified, HDFq inserts data chunks directly into the dataset bypassing several internal processing steps of the HDF5 library itself (e.g. data conversion, filter pipeline), which can lead to a much faster writing. In case the keyword PARALLEL<sup>40</sup> is specified, HDFq inserts data into a dataset in parallel using all the MPI processes specified upon launching the program (that employs HDFq).

By default, the entire *object\_name* is written when performing an insert operation. To write only a subset (i.e. portion) of *object\_name*, hyperslab and point selections can be used<sup>41</sup>. To enable a (regular) hyperslab selection, the *start*, *stride*, *count* and *block* parameters may be specified and separated with a colon (:). For each dimension of *object\_name*, a set of such parameters may be specified and each set should be separated with a comma (,). Multiple hyperslab selections can

---

<sup>39</sup> Only available for HDF5 datasets as, by design, direct insert (i.e. write) for HDF5 attributes is not supported by the HDF5 library. Moreover, the library does not support inserting (i.e. writing) data directly into a dataset of data type variable-length or in parallel (by specifying the keyword PARALLEL).

<sup>40</sup> This option is not allowed in Windows as HDFq does not support the parallel HDF5 (PHDF5) library in this platform currently. Moreover, due to a limitation of the PHDF5 library, inserting (i.e. writing) data into an HDF5 dataset of data type variable-length in parallel is not supported.

<sup>41</sup> Only available for HDF5 datasets as, by design, both hyperslab and point selections for HDF5 attributes are not supported by the HDF5 library.

be enabled at once (in this case, the hyperslab will be considered irregular). This is enabled by using the following boolean operators:

- OR – adds the new selection to the existing selection.
- AND – retains only the overlapping portions of the new selection and the existing selection.
- XOR – retains only the elements that are members of the new selection or the existing selection, excluding elements that are members of both selections.
- NOT A – retains only elements of the new selection that are not in the existing selection.
- NOT B – retains only elements of the existing selection that are not in the new selection.

To enable a point selection, a set of coordinates may be specified. Each coordinate is separated with a comma (,). More than one set of coordinates (i.e. points) may be specified and each set should be separated with a semicolon (;). Of note, hyperslab and point selections cannot be used both at the same time (i.e. be mixed) in an insert operation. Since hyperslab and point selections can be complicated to set up, it is highly recommended to first read [https://support.hdfgroup.org/HDF5/doc1.8/UG/HDF5\\_Users\\_Guide-Responsive%20HTML5/index.html#t=HDF5\\_Users\\_Guide%2FDataspaces%2FHDF5\\_Dataspaces\\_and\\_Partial\\_I\\_O.htm%23TOC\\_7\\_4\\_1\\_Data\\_Selectionbc-7](https://support.hdfgroup.org/HDF5/doc1.8/UG/HDF5_Users_Guide-Responsive%20HTML5/index.html#t=HDF5_Users_Guide%2FDataspaces%2FHDF5_Dataspaces_and_Partial_I_O.htm%23TOC_7_4_1_Data_Selectionbc-7) and eventually enable the debug mechanism (through the operation `SET DEBUG`) when working with these to obtain debug information in case of errors.

HDFqI provides several ways of inserting data into a dataset or attribute, namely either from a cursor (e.g. `INSERT INTO my_dataset`), direct values (e.g. `INSERT INTO my_dataset VALUES(0, 2, 4, 6, 8)`), or an input redirecting option (e.g. `INSERT INTO my_dataset VALUES FROM FILE my_file.txt`).

### **Parameter(s)**

*mask\_value* – optional integer that specifies which filters have been applied to the data chunk. A filter is skipped if the bit corresponding to the position of the filter in the pipeline is turned on. If *mask\_value* is specified it must be equal to or greater than zero (otherwise an error will be raised). Otherwise, if it is not specified and in case the keyword `DIRECTLY` is specified, its default value is 0 (meaning that all filters have been applied to the data chunk).

*data\_size* – optional integer that specifies the size (in bytes) of the data to insert (i.e. write) into the HDF5 dataset. If *data\_size* is specified it must be greater than zero (otherwise an error will be raised). Otherwise, if it is not specified and in case the keyword `DIRECTLY` is specified, HDFqI automatically calculates the size by multiplying all storage dimensions of the dataset with its data type size.



*file\_name* – optional string that specifies the name of the HDF5 file in which the HDF5 dataset or attribute to insert (i.e. write) data into is stored. If *file\_name* is specified, the file is opened on the fly, the dataset or attribute is inserted and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose data is to be inserted) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute to insert (i.e. write) data into. Multiple datasets or attributes are separated with a comma (,).

*start* – optional integer that specifies the starting location of the hyperslab selection. If not specified, its default value is 0 (i.e. the first position of the dimension in question). If negative, its value will be the last position of the dimension in question minus the value of *start*.

*stride* – optional integer that specifies the number of elements to separate each block to be selected. If not specified, its default value is equal to the value of *block*.

*count* – optional integer that specifies the number of blocks to select along each dimension. If not specified, its default value is 1.

*block* – optional integer that specifies the size of the block selected (i.e. number of elements) from the HDF5 dataset. If not specified, its default value is the size of the dimension in question minus the value of *start* divided by the value of *count*.

*coord* – optional integer that specifies the point of interest (i.e. to insert) for the point selection. If negative, its value will be the last position of the dimension in question minus the value of *coord*.

*chunk\_number* – optional integer that specifies the number of the chunk to insert (i.e. write) data into. Multiple chunk numbers are separated with a comma (,). If *chunk\_number* is specified it must either be between 0 and the storage dimension in question - 1 (otherwise an error will be raised) or negative (in this case its value will be the last position of the storage dimension in question minus the value of *chunk\_number*). Otherwise, if it is not specified and in case the keyword **DIRECTLY** is specified, its default value is 0 (i.e. first chunk of the storage dimension in question).

*value* – optional integer, float or string to insert (i.e. write) into the HDF5 dataset or attribute. Multiple values are separated with a comma (,). In case *value* is not specified, the element in question will be zeroed or emptied depending on whether the dataset/attribute is a number or a string, respectively.

*input\_redirecting\_option* – optional option that specifies a file or memory to read data from in order to write it into an HDF5 dataset or attribute (please refer to the subsection [FROM](#) for additional information).

## Return

Nothing

## Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type short of one dimension (size 3)
CREATE DATASET my_dataset0 AS SMALLINT(3)

# create an HDF5 dataset named "my_dataset1" of data type int of one dimension (size 5)
CREATE DATASET my_dataset1 AS INT(5)

# insert (i.e. write) values 65, 66 and 67 into dataset "my_dataset0"
INSERT INTO my_dataset0 VALUES(65, 66, 67)

# select (i.e. read) data from dataset "my_dataset0" and populate cursor in use with it (should
be 65, 66, 67)
SELECT FROM my_dataset0

# insert (i.e. write) values into dataset "my_dataset1" from cursor in use (should be 65, 66,
67, 0, 0)
INSERT INTO my_dataset1

# create an HDF5 attribute named "my_attribute0" of data type short
CREATE ATTRIBUTE my_attribute0 AS SMALLINT

# insert (i.e. write) value 95 into attribute "my_attribute0"
INSERT INTO my_attribute0 VALUES(95)

# create an HDF5 attribute named "my_attribute1" of data type unsigned short of one dimension
(size 2)
CREATE ATTRIBUTE my_attribute1 AS UNSIGNED SMALLINT(2)

# insert (i.e. write) values 95 and 97 into attribute "my_attribute1"
INSERT INTO my_attribute1 VALUES(95, 97)
```

```
# create an HDF5 dataset named "my_dataset2" of data type float of one dimension (size 512)
CREATE DATASET my_dataset2 AS FLOAT(512)

# insert (i.e. write) values into dataset "my_dataset2" from a text file named "my_file0.txt"
that has values separated with "," (i.e. default separator)
```

```
INSERT INTO my_dataset2 VALUES FROM FILE my_file0.txt
```

```
# insert (i.e. write) values into dataset "my_dataset2" from a text file named "my_file1.txt"
that has a DOS-based end of line (EOL) terminator and values separated with "***"
```

```
INSERT INTO my_dataset2 VALUES FROM DOS TEXT FILE my_file1.txt SEPARATOR **
```

```
# insert (i.e. write) values into dataset "my_dataset2" from a binary file named "my_file.bin"
```

```
INSERT INTO my_dataset2 VALUES FROM BINARY FILE my_file.bin
```

```
# create an HDF5 dataset named "my_dataset3" of data type short of one dimension (size 5)
```

```
CREATE DATASET my_dataset3 AS SMALLINT(5)
```

```
# insert (i.e. write) value 9 into position #3 of dataset "my_dataset3" using a hyperslab
selection
```

```
INSERT INTO my_dataset3(3:::) VALUES(9)
```

```
# select (i.e. read) data from dataset "my_dataset3" and populate cursor in use with it (should
be 0, 0, 0, 9, 0)
```

```
SELECT FROM my_dataset3
```

```
# insert (i.e. write) value 9 into position #4 of dataset "my_dataset3" using a hyperslab
selection
```

```
INSERT INTO my_dataset3(-1:::) VALUES(7)
```

```
# select (i.e. read) data from dataset "my_dataset3" and populate cursor in use with it (should
be 0, 0, 0, 9, 7)
```

```
SELECT FROM my_dataset3
```

```
# insert (i.e. write) values 5 and 3 into positions #1 and #2 of dataset "my_dataset3" using a
hyperslab selection
```

```
INSERT INTO my_dataset3(1:::2) VALUES(5, 3)
```

```
# select (i.e. read) data from dataset "my_dataset3" and populate cursor in use with it (should
be 0, 5, 3, 9, 7)
```

```
SELECT FROM my_dataset3
```

```
# create an HDF5 dataset named "my_dataset4" of data type int of two dimensions (size 3x3)
```

```
CREATE DATASET my_dataset4 AS INT(3, 3)
```

```
# insert (i.e. write) value 8 into position #2 of the first dimension and position #1 of the
second dimension of dataset "my_dataset4" using a hyperslab selection
```

```
INSERT INTO my_dataset4(2:::, 1:::) VALUES(8)
```

```
# select (i.e. read) data from dataset "my_dataset4" and populate cursor in use with it (should
be 0, 0, 0, 0, 0, 0, 0, 0, 8, 0)
SELECT FROM my_dataset4

# insert (i.e. write) value 4 into position #2 of the first dimension and position #0 of the
second dimension, and value 6 into position #2 of the first dimension and position #2 of the
second dimension of dataset "my_dataset4" using a hyperslab selection
INSERT INTO my_dataset4(2::, 0:2:2:1) VALUES(4, 6)

# select (i.e. read) data from dataset "my_dataset4" and populate cursor in use with it (should
be 0, 0, 0, 0, 0, 0, 0, 4, 8, 6)
SELECT FROM my_dataset4

# create an HDF5 dataset named "my_dataset5" of data type short of one dimension (size 10)
CREATE DATASET my_dataset5 AS SMALLINT(10)

# insert (i.e. write) values 90, 91 and 92 into positions #2, #3 and #4, value 93 into
position#5, and values 94 and 95 into positions #7 and #8 of dataset "my_dataset5" using an
irregular hyperslab selection
INSERT INTO my_dataset5(2::3:1 OR 4::2:1 OR 7::2:1) VALUES(90, 91, 92, 93, 94, 95)

# select (i.e. read) data from dataset "my_dataset5" and populate cursor in use with it (should
be 0, 0, 90, 91, 92, 93, 0, 94, 95, 0)
SELECT FROM my_dataset5

# create an HDF5 dataset named "my_dataset6" of data type long long of one dimension (size 15)
CREATE DATASET my_dataset6 AS BIGINT(15)

# insert (i.e. write) values 75 and 77 into positions #5 and #6 of dataset "my_dataset6" using
an irregular hyperslab selection
INSERT INTO my_dataset6(3::4:1 AND 5::3:1) VALUES(75, 77, 79, 81, 83, 85, 87)

# select (i.e. read) data from dataset "my_dataset6" and populate cursor in use with it (should
be 0, 0, 0, 0, 0, 0, 75, 77, 0, 0, 0, 0, 0, 0, 0)
SELECT FROM my_dataset6

# create an HDF5 dataset named "my_dataset7" of data type float of one dimension (size 8)
CREATE DATASET my_dataset7 AS FLOAT(8)

# insert (i.e. write) values 7.5, 7.7 and 7.9 into positions #2, #4 and #7 of dataset
"my_dataset7" using a point selection
INSERT INTO my_dataset7(2; 4; 7) VALUES(7.5, 7.7, 7.9)
```

```
# select (i.e. read) data from dataset "my_dataset7" and populate cursor in use with it (should
be 0, 0, 7.5, 0, 7.7, 0, 0, 7.9)
SELECT FROM my_dataset7

# create an HDF5 dataset named "my_dataset8" of data type double of two dimensions (size 4x3)
CREATE DATASET my_dataset8 AS DOUBLE(4, 3)

# insert (i.e. write) value 15.2 into position #1 of the first dimension and position #2 of the
second dimension, and value 18.5 into position #3 of the first dimension and position #0 of the
second dimension of dataset "my_dataset8" using a point selection
INSERT INTO my_dataset8(1, 2; 3, 0) VALUES(15.2, 18.5)

# select (i.e. read) data from dataset "my_dataset8" and populate cursor in use with it (should
be 0, 0, 0, 0, 0, 15.2, 0, 0, 0, 18.5, 0, 0)
SELECT FROM my_dataset8
```

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5

# create an HDF5 dataset named "my_dataset9" of data type double in the HDF5 file currently in
use (i.e. file "my_file.h5")
CREATE DATASET my_dataset9 AS DOUBLE

# insert (i.e. write) value 6.5 into dataset "my_dataset9"
INSERT INTO my_dataset9 VALUES(6.5)

# select (i.e. read) data from dataset "my_dataset9" and populate cursor in use with it (should
be 6.5)
SELECT FROM my_dataset9

# close HDF5 file currently in use (i.e. file "my_file.h5")
CLOSE FILE

# insert (i.e. write) value 3.2 into dataset "my_dataset9" in file "my_file.h5"
INSERT INTO my_file.h5 my_dataset9 VALUES(3.2)

# select (i.e. read) data from dataset "my_dataset9" in file "my_file.h5" and populate cursor
in use with it (should be 3.2)
SELECT FROM my_file.h5 my_dataset9
```

```
# create a chunked (size 2) HDF5 dataset named "my_dataset10" of data type int of one dimension
(size 6)
CREATE CHUNKED(2) DATASET my_dataset10 AS INT(6)

# insert (i.e. write) values 60 and 61 directly into chunk #0 of dataset "my_dataset10" using a
(filter) mask equal to 8
INSERT DIRECTLY MASK 8 INTO my_dataset10 VALUES(60, 61)

# insert (i.e. write) values 62 and 63 directly into chunk #1 of dataset "my_dataset10" using a
(filter) mask equal to 255 (i.e. 0xFF)
INSERT DIRECTLY MASK 0xFF INTO my_dataset10(1) VALUES(62, 63)

# insert (i.e. write) values 64 and 65 directly into chunk #2 of dataset "my_dataset10" using a
(filter) mask equal to 0 (i.e. default value)
INSERT DIRECTLY INTO my_dataset10(2) VALUES(64, 65)

# select (i.e. read) data from dataset "my_dataset10" and populate cursor in use with it
(should be 60, 61, 62, 63, 64, 65)
SELECT FROM my_dataset10
```

```
// declare variables
char script[1024];
double data[3];

// create an HDF5 dataset named "my_dataset11" of data type double of one dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset11 AS DOUBLE(3)");

// populate variable "data" with certain values
data[0] = 21.1;
data[1] = 18.8;
data[2] = 75.6;

// register variable "data" for subsequent use (by HDFql)
hdfql_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset
"my_dataset11"
sprintf(script, "INSERT INTO my_dataset11 VALUES FROM MEMORY %d",
hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);
```

```
// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);
```

```
// declare variables
char script[1024];
HDFQL_VARIABLE_LENGTH data[3];

// create an HDF5 dataset named "my_dataset12" of data type variable-length double of one
dimension (size 3)
hdfq1_execute("CREATE DATASET my_dataset12 AS VARDOUBLE(3)");

// allocate memory in variable "data"
data[0].address = malloc(2 * sizeof(double));
data[0].count = 2;
data[1].address = malloc(3 * sizeof(double));
data[1].count = 3;
data[2].address = malloc(1 * sizeof(double));
data[2].count = 1;

// populate variable "data" with certain values
*((double *) data[0].address + 0) = 3.2;
*((double *) data[0].address + 1) = 1.3;
*((double *) data[1].address + 0) = 0;
*((double *) data[1].address + 1) = 0.2;
*((double *) data[1].address + 2) = 9.1;
*((double *) data[2].address + 0) = 6.5;

// register variable "data" for subsequent use (by HDFq1)
hdfq1_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset
"my_dataset12"
sprintf(script, "INSERT INTO my_dataset12 VALUES FROM MEMORY %d",
hdfq1_variable_get_number(data));

// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);
```

```
// select (i.e. read) data from dataset "my_dataset12" and populate cursor in use with it
hdfql_execute("SELECT FROM my_dataset12");

// display content of cursor in use (should be 3.2, 1.3, 0, 0.2, 9.1, 6.5)
while(hdfql_cursor_next(NULL) == HDFQL_SUCCESS)
{
    while(hdfql_subcursor_next(NULL) == HDFQL_SUCCESS)
    {
        printf("%f\n", *hdfql_subcursor_get_double(NULL));
    }
}

// release memory allocated in variable "data"
free(data[0].address);
free(data[1].address);
free(data[2].address);
```

```
// declare variables
char script[1024];
char *data[3];

// create an HDF5 dataset named "my_dataset13" of data type variable-length char of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset13 AS VARCHAR(3)");

// allocate memory in variable "data"
data[0] = malloc(13 * sizeof(char));
data[1] = malloc(5 * sizeof(char));
data[2] = malloc(7 * sizeof(char));

// populate variable "data" with certain values
strcpy(data[0], "Hierarchical");
strcpy(data[1], "Data");
strcpy(data[2], "Format");

// register variable "data" for subsequent use (by HDFql)
hdfql_variable_register(data);

// prepare script to insert (i.e. write) values from variable "data" into dataset
"my_dataset13"
sprintf(script, "INSERT INTO my_dataset13 VALUES FROM MEMORY %d",
hdfql_variable_get_number(data));
```



```
// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);

// select (i.e. read) data from dataset "my_dataset13" and populate cursor in use with it
hdfq1_execute("SELECT FROM my_dataset13");

// display content of cursor in use (should be "Hierarchical", "Data", "Format")
while(hdfq1_cursor_next(NULL) == HDFQL_SUCCESS)
{
    printf("%s\n", hdfq1_cursor_get_char(NULL));
}

// release memory allocated in variable "data"
free(data[0]);
free(data[1]);
free(data[2]);
```

```
// assume that the following program is launched in parallel using four MPI processes (e.g.
"mpiexec -n 4 my_program")

// declare variables
char script[1024];
int rank;

// create an HDF5 file named "my_file.h5" in parallel
hdfq1_execute("CREATE PARALLEL FILE my_file.h5");

// use (i.e. open) HDF5 file "my_file.h5" in parallel
hdfq1_execute("USE PARALLEL FILE my_file.h5");

// create an HDF5 dataset named "my_dataset14" of data type int of one dimension (size 4)
hdfq1_execute("CREATE DATASET my_dataset14 AS INT(4)");

// get number (i.e. rank) of the MPI process (should be between 0 and 3)
rank = hdfq1_mpi_get_rank();

// prepare script to insert (i.e. write) values 0, 10, 20 and 30 in parallel into positions #0
// (by MPI process rank 0), #1 (by MPI process rank 1), #2 (by MPI process rank 2) and #3 (by MPI
```

```

process rank 3) of dataset "my_dataset14" using a point selection
sprintf(script, "INSERT INTO PARALLEL my_dataset14(%d) VALUES(%d)", rank, rank * 10);

// execute script
hdfql_execute(script);

```

## 6.6 DATA QUERY LANGUAGE (DQL)

Data Query Language (DQL) is, generally speaking, syntax for retrieving data stored in structures. In HDFqL, the DQL is composed of only one operation (SELECT). It enables retrieval (i.e. reading) of data stored in HDF5 datasets or attributes according to certain criteria. Moreover, it supports both [POST-PROCESSING](#) and [REDIRECTING](#) options to further transform and redirect the result of the operation according to the programmer's needs.

### 6.6.1 SELECT

#### Syntax

```

SELECT [DIRECTLY] FROM [PARALLEL] [DATASET | ATTRIBUTE] [file_name] object_name
[[{[start]:[stride]:[count]:[block] [, [start]:[stride]:[count]:[block]]* [, {OR | AND | XOR | NOTA | NOTB}
[start]:[stride]:[count]:[block] [, [start]:[stride]:[count]:[block]]*}] | {coord [, coord]* [, coord [, coord]]*}] |
{chunk_number [, chunk_number]*}]]

[CACHE [SLOTS {slots_value | FILE | DEFAULT}] [SIZE {size_value | FILE | DEFAULT}] [PREEMPTION
{preemption_value | FILE | DEFAULT}]]

[post_processing_option [post_processing_option]*]

[output_redirecting_option]

```

#### Description

Select (i.e. read) data from an HDF5 dataset or attribute named *object\_name*. In case the keyword **CACHE** is specified, the dataset is read using cache parametrized with the values *slots\_value*, *size\_value* and *preemption\_value* (instead of the dataset cache parameters that may have been set through the operation [SET CACHE](#)). In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword **DATASET** nor **ATTRIBUTE** is specified, the object for which data will be read is the dataset. To explicitly read data from an object

according to its type, the keyword DATASET or ATTRIBUTE must be specified. In case the keyword DIRECTLY<sup>42</sup> is specified, HDFqI reads data chunks directly from the dataset bypassing several internal processing steps of the HDF5 library itself (e.g. data conversion, filter pipeline), which can lead to a much faster reading. In case the keyword PARALLEL<sup>43</sup> is specified, HDFqI reads data from a dataset in parallel using all the MPI processes specified upon launching the program (that employs HDFqI).

By default, the entire *object\_name* is read when performing a select operation. To read only a subset (i.e. portion) of *object\_name*, hyperslab and point selections can be used<sup>44</sup>. To enable a (regular) hyperslab selection, the *start*, *stride*, *count* and *block* parameters may be specified and separated with a colon (:). For each dimension of *object\_name*, a set of such parameters may be specified and each set should be separated with a comma (,). Multiple hyperslab selections can be enabled at once (in this case, the hyperslab will be considered irregular). This is enabled by using the following boolean operators:

- OR – adds the new selection to the existing selection.
- AND – retains only the overlapping portions of the new selection and the existing selection.
- XOR – retains only the elements that are members of the new selection or the existing selection, excluding elements that are members of both selections.
- NOTA – retains only elements of the new selection that are not in the existing selection.
- NOTB – retains only elements of the existing selection that are not in the new selection.

To enable a point selection, a set of coordinates may be specified. Each coordinate is separated with a comma (,). More than one set of coordinates (i.e. points) may be specified and each set should be separated with a semicolon (;). Of note, hyperslab and point selections cannot be used both at the same time (i.e. be mixed) in a select operation. Since hyperslab and point selections can be complicated to set up, it is highly recommended to first read [https://support.hdfgroup.org/HDF5/doc1.8/UG/HDF5\\_Users\\_Guide-Responsive%20HTML5/index.html#t=HDF5\\_Users\\_Guide%2FDataspaces%2FHDF5\\_Dataspaces\\_and\\_Partial\\_I\\_O.htm%23T\\_OC\\_7\\_4\\_1\\_Data\\_Selectionbc-7](https://support.hdfgroup.org/HDF5/doc1.8/UG/HDF5_Users_Guide-Responsive%20HTML5/index.html#t=HDF5_Users_Guide%2FDataspaces%2FHDF5_Dataspaces_and_Partial_I_O.htm%23T_OC_7_4_1_Data_Selectionbc-7) and eventually enable the debug mechanism (through the operation `SET DEBUG`) when working with these to obtain debug information in case of errors.

---

<sup>42</sup> Only available for HDF5 datasets as, by design, direct selection (i.e. read) for HDF5 attributes is not supported by the HDF5 library. Moreover, the library does not support selecting (i.e. reading) data directly from a dataset of data type variable-length or in parallel (by specifying the keyword PARALLEL).

<sup>43</sup> This option is not allowed in Windows as HDFqI does not support the parallel HDF5 (PHDF5) library in this platform currently.

<sup>44</sup> Only available for HDF5 datasets as, by design, both hyperslab and point selections for HDF5 attributes are not supported by the HDF5 library.

HDFqL provides several ways of writing result sets that was read from a dataset or attribute, namely either to a cursor (e.g. “*SELECT FROM my\_dataset*”) or an output redirecting option (e.g. “*SELECT FROM my\_dataset INTO FILE my\_file.txt*”).

### **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file in which the HDF5 dataset or attribute to select (i.e. read) data from is stored. If *file\_name* is specified, the file is opened on the fly, the dataset or attribute is selected and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose data is to be selected is stored) in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute to select (i.e. read) data from.

*start* – optional integer that specifies the starting location of the hyperslab selection. If not specified, its default value is 0 (i.e. the first position of the dimension in question). If negative, its value will be the last position of the dimension in question minus the value of *start*.

*stride* – optional integer that specifies the number of elements to separate each block to be selected. If not specified, its default value is equal to the value of *block*.

*count* – optional integer that specifies the number of blocks to select along each dimension. If not specified, its default value is 1.

*block* – optional integer that specifies the size of the block (i.e. number of elements) selected from the HDF5 dataset. If not specified, its default value is the size of the dimension in question minus the value of *start* divided by the value of *count*.

*coord* – optional integer that specifies the point of interest (i.e. to select) for the point selection. If negative, its value will be the last position of the dimension in question minus the value of *coord*.

*chunk\_number* – optional integer that specifies the number of the chunk to select (i.e. read) data from. Multiple chunk numbers are separated with a comma (,). If *chunk\_number* is specified it must either be between 0 and the storage dimension in question - 1 (otherwise an error will be raised) or negative (in this case its value will be the last position of the storage dimension in question minus the value of *chunk\_number*). Otherwise, if it is not specified and in case the keyword **DIRECTLY** is specified, its default value is 0 (i.e. first chunk of the storage dimension in question).

*slots\_value* – optional integer that specifies the number of chunk slots in the raw data chunk cache for accessing the HDF5 dataset. Due to the hashing strategy, its value should ideally be a prime number. In case the keyword **DEFAULT** is specified, its value is 521 (i.e. default value defined by the HDF5 library). In case the keyword **FILE** is specified, its value will be as defined in the cache slots parameter upon using (i.e. opening) the file. In case the keyword **SLOTS** is not specified, its

current value remains intact. Of note, if *object\_name* is an HDF5 attribute then the cache is ignored (i.e. it is not applicable for accessing attributes).

*size\_value* – optional integer that specifies the total size of the raw data chunk cache in bytes for accessing the HDF5 dataset. In case the keyword DEFAULT is specified, its value is 1048576 (i.e. 1 MB – default value defined by the HDF5 library). In case the keyword FILE is specified, its value will be as defined in the cache size parameter upon using (i.e. opening) the file. In case the keyword SIZE is not specified, its current value remains intact. Of note, if *object\_name* is an HDF5 attribute then the cache is ignored (i.e. it is not applicable for accessing attributes).

*preemption\_value* – optional float that specifies the chunk preemption policy for accessing the HDF5 dataset. Its value must be between 0 and 1. It indicates the weighting according to which chunks which have been fully read or written are penalized when determining which chunks to flush from cache. In case the keyword DEFAULT is specified, its value is 0.75 (i.e. default value defined by the HDF5 library). In case the keyword FILE is specified, its value will be as defined in the cache preemption parameter upon using (i.e. opening) the file. In case the keyword PREEMPTION is not specified, its current value remains intact. Of note, if *object\_name* is an HDF5 attribute then the cache is ignored (i.e. it is not applicable for accessing attributes).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The data selected (i.e. read) from an HDF5 dataset or attribute as an [HDFQL\\_TINYINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_TINYINT](#)), [HDFQL\\_UNSIGNED\\_TINYINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_UNSIGNED\\_TINYINT](#)), [HDFQL\\_SMALLINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_SMALLINT](#)), [HDFQL\\_UNSIGNED\\_SMALLINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_UNSIGNED\\_SMALLINT](#)), [HDFQL\\_INT](#) (in case the data type of the dataset or attribute is [HDFQL\\_INT](#)), [HDFQL\\_UNSIGNED\\_INT](#) (in case the data type of the dataset or attribute is [HDFQL\\_UNSIGNED\\_INT](#)), [HDFQL\\_BIGINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_BIGINT](#)), [HDFQL\\_UNSIGNED\\_BIGINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_UNSIGNED\\_BIGINT](#)), [HDFQL\\_FLOAT](#) (in case the data type of the dataset or attribute is [HDFQL\\_FLOAT](#)), [HDFQL\\_DOUBLE](#) (in case the data type of the dataset or attribute is [HDFQL\\_DOUBLE](#)), [HDFQL\\_CHAR](#) (in case the data type of the dataset or attribute is [HDFQL\\_CHAR](#)), [HDFQL\\_VARTINYINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_VARTINYINT](#)), [HDFQL\\_UNSIGNED\\_VARTINYINT](#) (in case the data type of the dataset or attribute is [HDFQL\\_UNSIGNED\\_VARTINYINT](#)).

`HDFQL_UNSIGNED_VARTINYINT`), `HDFQL_VARSMALLINT` (in case the data type of the dataset or attribute is `HDFQL_VARSMALLINT`), `HDFQL_UNSIGNED_VARSMALLINT` (in case the data type of the dataset or attribute is `HDFQL_UNSIGNED_VARSMALLINT`), `HDFQL_VARINT` (in case the data type of the dataset or attribute is `HDFQL_VARINT`), `HDFQL_UNSIGNED_VARINT` (in case the data type of the dataset or attribute is `HDFQL_UNSIGNED_VARINT`), `HDFQL_VARBIGINT` (in case the data type of the dataset or attribute is `HDFQL_VARBIGINT`), `HDFQL_UNSIGNED_VARBIGINT` (in case the data type of the dataset or attribute is `HDFQL_UNSIGNED_VARBIGINT`), `HDFQL_VARFLOAT` (in case the data type of the dataset or attribute is `HDFQL_VARFLOAT`), `HDFQL_VARDOUBLE` (in case the data type of the dataset or attribute is `HDFQL_VARDOUBLE`), `HDFQL_VARCHAR` (in case the data type of the dataset or attribute is `HDFQL_VARCHAR`) or `HDFQL_OPAQUE` (in case the data type of the dataset or attribute is `HDFQL_OPAQUE`).

### Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type short of one dimension (size 3) with
initial values of 65, 66 and 77
CREATE DATASET my_dataset0 AS SMALLINT(3) VALUES(65, 66, 67)

# select (i.e. read) data from dataset "my_dataset0" and populate cursor in use with it (should
be 65, 66, 67)
SELECT FROM my_dataset0

# create an HDF5 attribute named "my_attribute0" of data type short
CREATE ATTRIBUTE my_attribute0 AS SMALLINT

# select (i.e. read) data from attribute "my_attribute0" and populate cursor in use with it
(should be 0)
SELECT FROM my_attribute0

# create an HDF5 attribute named "my_attribute1" of data type unsigned short of one dimension
(size 2) with initial values of 95 and 97
CREATE ATTRIBUTE my_attribute1 AS UNSIGNED SMALLINT(2) VALUES(95, 97)

# select (i.e. read) data from attribute "my_attribute1" and populate cursor in use with it
(should be 95, 97)
SELECT FROM my_attribute1
```

```
# create an HDF5 dataset named "my_dataset1" of data type float of one dimension (size 512)
CREATE DATASET my_dataset1 AS FLOAT(512)

# select (i.e. read) data from dataset "my_dataset1" and write it into a text file named
```

```
"my_file0.txt" using default separator ","
SELECT FROM my_dataset1 INTO FILE my_file0.txt

# select (i.e. read) data from dataset "my_dataset1" and write it into a text file named
"my_file1.txt" using a DOS-based end of line (EOL) terminator and separator "***"
SELECT FROM my_dataset1 INTO DOS TEXT FILE my_file1.txt SEPARATOR **

# select (i.e. read) data from dataset "my_dataset1" and write it into a binary file named
"my_file.bin"
SELECT FROM my_dataset1 INTO BINARY FILE my_file.bin
```

```
# create an HDF5 dataset named "my_dataset2" of data type short of one dimension (size 5)
CREATE DATASET my_dataset2 AS SMALLINT(5)

# insert (i.e. write) values 0, 5, 3, 9 and 7 into dataset "my_dataset2"
INSERT INTO my_dataset2 VALUES(0, 5, 3, 9, 7)

# select (i.e. read) data from dataset "my_dataset2" using a hyperslab selection (starting from
position #3) and populate cursor in use with it (should be 9, 7)
SELECT FROM my_dataset2(3:::)

# select (i.e. read) data from dataset "my_dataset2" using a hyperslab selection (starting from
position #4) and populate cursor in use with it (should be 7)
SELECT FROM my_dataset2(-1:::)

# select (i.e. read) data from dataset "my_dataset2" using a hyperslab selection (starting from
position #1 with a block of 2) and populate cursor in use with it (should be 5, 3)
SELECT FROM my_dataset2(1:::2)

# create an HDF5 dataset named "my_dataset3" of data type int of two dimensions (size 3x3)
CREATE DATASET my_dataset3 AS INT(3, 3)

# insert (i.e. write) values 0, 0, 0, 0, 0, 0, 4, 8 and 6 into dataset "my_dataset3"
INSERT INTO my_dataset3 VALUES(0, 0, 0, 0, 0, 0, 4, 8, 6)

# select (i.e. read) data from dataset "my_dataset3" using a hyperslab selection (starting from
position #2 of the first dimension and position #1 of the second dimension) and populate cursor
in use with it (should be 8, 6)
SELECT FROM my_dataset3(2:::, 1:::)

# select (i.e. read) data from dataset "my_dataset3" using a hyperslab selection (starting from
position #2 of the first dimension and position #0 of the second dimension with a stride of 2,
```

```
count of 2 and block of 1) and populate cursor in use with it (should be 4, 6)
SELECT FROM my_dataset3(2::, 0:2:2:1)

# create an HDF5 dataset named "my_dataset4" of data type short of one dimension (size 10)
CREATE DATASET my_dataset4 AS SMALLINT(10)

# insert (i.e. write) values 0, 0, 90, 91, 92, 93, 0, 94, 95 and 0 into dataset "my_dataset4"
INSERT INTO my_dataset4 VALUES(0, 0, 90, 91, 92, 93, 0, 94, 95, 0)

# select (i.e. read) data from dataset "my_dataset4" using an irregular hyperslab selection
(starting from position #2 with a count of 3 and block of 1; starting from position #4 with a
count of 2 and block of 1; starting from position #7 with a count of 2 and block of 1) and
populate cursor in use with it (should be 90, 91, 92, 93, 94, 95)
SELECT FROM my_dataset4(2::3:1 OR 4::2:1 OR 7::2:1)

# create an HDF5 dataset named "my_dataset5" of data type long long of one dimension (size 15)
CREATE DATASET my_dataset5 AS BIGINT(15)

# insert (i.e. write) values 0, 0, 0, 0, 0, 75, 77, 0, 0, 0, 0, 0, 0, 0 and 0 into dataset
"my_dataset5"
INSERT INTO my_dataset5 VALUES(0, 0, 0, 0, 0, 75, 77, 0, 0, 0, 0, 0, 0, 0, 0)

# select (i.e. read) data from dataset "my_dataset5" using an irregular hyperslab selection
(starting from position #3 with a count of 4 and block of 1; starting from position #5 with a
count of 3 and block of 1) and populate cursor in use with it (should be 75, 77)
SELECT FROM my_dataset5(3::4:1 AND 5::3:1)

# create an HDF5 dataset named "my_dataset6" of data type float of one dimension (size 8)
CREATE DATASET my_dataset6 AS FLOAT(8)

# insert (i.e. write) values 0, 0, 7.5, 0, 7.7, 0, 0 and 7.9 into dataset "my_dataset6"
INSERT INTO my_dataset6 VALUES(0, 0, 7.5, 0, 7.7, 0, 0, 7.9)

# select (i.e. read) data from dataset "my_dataset6" using a point selection (positions #2, #4
and #7) and populate cursor in use with it (should be 7.5, 7.7, 7.9)
SELECT FROM my_dataset6(2; 4; 7)

# create an HDF5 dataset named "my_dataset7" of data type double of two dimensions (size 4x3)
CREATE DATASET my_dataset7 AS DOUBLE(4, 3)

# insert (i.e. write) values 0, 0, 0, 0, 0, 15.2, 0, 0, 0, 18.5, 0 and 0 into dataset
"my_dataset7"
INSERT INTO my_dataset7 VALUES(0, 0, 0, 0, 0, 15.2, 0, 0, 0, 18.5, 0, 0)
```



```
# select (i.e. read) data from dataset "my_dataset7" using a point selection (position #1 of
the first dimension and position #2 of the second dimension, position #3 of the first dimension
and position #0 of the second dimension) and populate cursor in use with it (should be 15.2,
18.5)
SELECT FROM my_dataset7(1, 2; 3, 0)
```

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5

# create an HDF5 dataset named "my_dataset8" of data type double in the HDF5 file currently in
use (i.e. file "my_file.h5")
CREATE DATASET my_dataset8 AS DOUBLE

# insert (i.e. write) value 6.5 into dataset "my_dataset8"
INSERT INTO my_dataset8 VALUES(6.5)

# select (i.e. read) data from dataset "my_dataset8" and populate cursor in use with it (should
be 6.5)
SELECT FROM my_dataset8

# close HDF5 file currently in use (i.e. file "my_file.h5")
CLOSE FILE

# insert (i.e. write) value 3.2 into dataset "my_dataset8" in file "my_file.h5"
INSERT INTO my_file.h5 my_dataset8 VALUES(3.2)

# select (i.e. read) data from dataset "my_dataset8" in file "my_file.h5" and populate cursor
in use with it (should be 3.2)
SELECT FROM my_file.h5 my_dataset8
```

```
# create a chunked (size 2) HDF5 dataset named "my_dataset9" of data type int of one dimension
(size 6)
CREATE CHUNKED(2) DATASET my_dataset9 AS INT(6)

# insert (i.e. write) values 60, 61, 62, 63, 64 and 65 into dataset "my_dataset9"
INSERT INTO my_dataset9 VALUES(60, 61, 62, 63, 64, 65)

# select (i.e. read) data directly from chunk #0 of dataset "my_dataset9" (should be 60, 61)
SELECT DIRECTLY FROM my_dataset9
```

```
# select (i.e. read) data directly from chunk #1 of dataset "my_dataset9" (should be 62, 63)
SELECT DIRECTLY FROM my_dataset9(1)

# select (i.e. read) data directly from chunk #2 of dataset "my_dataset9" (should be 64, 65)
SELECT DIRECTLY FROM my_dataset9(2)
```

```
// declare variables
char script[1024];
double data[3];
int i;

// create an HDF5 dataset named "my_dataset10" of data type double of one dimension (size 3)
with initial values of 21.1, 18.8 and 75.6
hdfq1_execute("CREATE DATASET my_dataset10 AS DOUBLE(3) VALUES(21.1, 18.8, 75.6)");

// register variable "data" for subsequent use (by HDFq1)
hdfq1_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset10" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset10 INTO MEMORY %d", hdfq1_variable_get_number(data));

// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);

// display content of variable "data" (should be 21.1, 18.8, 75.6)
for(i = 0; i < 3; i++)
{
    printf("%f\n", data[i]);
}
```

```
// declare variables
char script[1024];
HDFQL_VARIABLE_LENGTH data[3];
int x;
int y;
int count;
```

```
// create an HDF5 dataset named "my_dataset11" of data type variable-length double of one
dimension (size 3)
hdfq1_execute("CREATE DATASET my_dataset11 AS VARDOUBLE(3)");

// insert (i.e. write) values into dataset "my_dataset11"
hdfq1_execute("INSERT INTO my_dataset11 VALUES((3.2, 1.3), (0, 0.2), (9.1, 7.4, 6.5))");

// register variable "data" for subsequent use (by HDFq1)
hdfq1_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset11" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset11 INTO MEMORY %d", hdfq1_variable_get_number(data));

// execute script
hdfq1_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFq1)
hdfq1_variable_unregister(data);

// display content of cursor in use (should be 3.2, 1.3, 0, 0.2, 9.1, 7.4, 6.5)
for(x = 0; x < 3; x++)
{
    count = data[x].count;
    for(y = 0; y < count; y++)
    {
        printf("%f\n", *((double *) data[x].address + y));
    }
}

// release memory allocated (by HDFq1) in variable "data"
for(x = 0; x < 3; x++)
{
    free(data[x].address);
}
```

```
// declare variables
char script[1024];
char *data[3];
int x;
```

```
// create an HDF5 dataset named "my_dataset12" of data type variable-length char of one
dimension (size 3)
hdfql_execute("CREATE DATASET my_dataset12 AS VARCHAR(3)");

// insert (i.e. write) values into dataset "my_dataset12"
hdfql_execute("INSERT INTO my_dataset12 VALUES(\"Hierarchical\", \"Data\", \"Format\")");

// register variable "data" for subsequent use (by HDFql)
hdfql_variable_register(data);

// prepare script to select (i.e. read) data from dataset "my_dataset12" and populate variable
"data" with it
sprintf(script, "SELECT FROM my_dataset12 INTO MEMORY %d", hdfql_variable_get_number(data));

// execute script
hdfql_execute(script);

// unregister variable "data" as it is no longer used/needed (by HDFql)
hdfql_variable_unregister(data);

// display content of cursor in use (should be "Hierarchical", "Data", "Format")
for(x = 0; x < 3; x++)
{
    printf("%s\n", data[x]);
}

// release memory allocated (by HDFql) in variable "data"
for(x = 0; x < 3; x++)
{
    free(data[x]);
}
```

```
// assume that 1) the following program is launched in parallel using four MPI processes (e.g.
"mpirun -n 4 my_program"), 2) an HDF5 file named "my_file.h5" containing a dataset named
"my_dataset13" of data type int of one dimension (size 4) already exists, and 3) the dataset
stores the values 0, 10, 20 and 30 in positions #0, #1, #2 and #3 respectively

// declare variables
char script[1024];
int rank;

// use (i.e. open) an HDF5 file named "my_file.h5" in parallel
```

```

hdfq1_execute("USE PARALLEL FILE my_file.h5");

// get number (i.e. rank) of the MPI process (should be between 0 and 3)
rank = hdfq1_mpi_get_rank();

// prepare script to select (i.e. read) in parallel positions #0 (by MPI process rank 0), #1
// (by MPI process rank 1), #2 (by MPI process rank 2) and #3 (by MPI process rank 3) from dataset
// "my_dataset13" using a point selection
sprintf(script, "SELECT FROM PARALLEL my_dataset13(%d)", rank);

// execute script
hdfq1_execute(script);

// move the cursor in use to the first position within the result set
hdfq1_cursor_first(NULL);

// display value selected (i.e. read) by each MPI process (should display message "Value read
// by MPI process rank X is Y" four times, where X is 0 and Y is 0, X is 1 and Y is 10, X is 2 and
// Y is 20, or X is 3 and Y is 30 (not necessarily in this order))
printf("Value read by MPI process rank %d is %d\n", rank, *hdfq1_cursor_get_int(NULL));

```

## 6.7 DATA INTROSPECTION LANGUAGE (DIL)

HDFq1 has certain operations that retrieve information about the internals of HDF5 files but also about HDFq1 itself and the runtime environment. These operations are part of the Data Introspection Language (DIL) and they all begin with the keyword **SHOW**. Moreover, these operations support both **POST-PROCESSING** and **REDIRECTING** options to further transform and redirect the result of operations according to the programmer's needs. Typically, a DIL operation has the following syntactical form:

**SHOW** *operation\_name* [*post\_processing\_option* [*post\_processing\_option*]\*] [*output\_redirecting\_option*]

### 6.7.1 SHOW FILE VALIDITY

#### Syntax

**SHOW FILE VALIDITY** *file\_name* [, *file\_name*]\*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

## **Description**

Show (i.e. get) validity of a file named *file\_name*. Multiple files' validities can be checked at once by separating these with a comma (,). If *file\_name* was not found or its validity could not be checked (due to unknown/unexpected reasons), no subsequent files are checked, and an error is raised.

## **Parameter(s)**

*file\_name* – mandatory string that specifies the name of the file whose validity is to be obtained. Multiple files are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The validity of a file as an [HDFQL\\_INT](#), which can either be [HDFQL\\_YES](#) or [HDFQL\\_NO](#) depending on whether the file is a valid HDF5 file or not.

## **Example(s)**

```
# create an HDF5 file named "my_file.h5"
CREATE FILE my_file.h5

# show (i.e. get) validity of file "my_file.h5" (should be 0 - i.e. HDFQL_YES)
SHOW FILE VALIDITY my_file.h5

# show (i.e. get) validity of a file named "not_an_hdf_file.xml" (should be -1 - i.e. HDFQL_NO)
# (assume that the file "not_an_hdf_file.xml" exists and contains XML text)
SHOW FILE VALIDITY not_an_hdf_file.xml

# show (i.e. get) validity of both files "my_file.h5" and "not_an_hdf_file.xml" at once (should
# be 0, -1)
SHOW FILE VALIDITY my_file.h5, not_an_hdf_file.xml
```

## 6.7.2 SHOW USE DIRECTORY

### Syntax

#### SHOW USE DIRECTORY

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### Description

Show (i.e. get) working directory currently in use.

### Parameter(s)

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

The working directory currently in use as an [HDFQL\\_VARCHAR](#).

### Example(s)

```
# set working directory currently in use to "/"
USE DIRECTORY /

# show (i.e. get) current working directory (should be "/")
SHOW USE DIRECTORY

# create a directory named "my_directory"
CREATE DIRECTORY my_directory
```

```
# set working directory currently in use to "my_directory" (more precisely "/my_directory")
USE DIRECTORY my_directory

# show (i.e. get) current working directory (should be "/my_directory")
SHOW USE DIRECTORY

# create two directories named "my_subdirectory0" and "my_subdirectory1" (both directories will
be created in directory "/my_directory")
CREATE DIRECTORY my_subdirectory0, my_subdirectory1

# set directory currently in use to "my_subdirectory0" (more precisely
"/my_directory/my_subdirectory0")
USE DIRECTORY my_subdirectory0

# show (i.e. get) current working directory (should be "/my_directory/my_subdirectory0")
SHOW USE DIRECTORY

# set directory currently in use to "my_subdirectory1" located one level up (more precisely
"/my_directory/my_subdirectory1")
USE DIRECTORY ../my_subdirectory1

# show (i.e. get) current working directory (should be "/my_directory/my_subdirectory1")
SHOW USE DIRECTORY

# set directory currently in use two levels up (should be "/")
USE DIRECTORY ../../

# show (i.e. get) current working directory (should be "/")
SHOW USE DIRECTORY
```

## 6.7.3 SHOW USE FILE

### Syntax

#### SHOW USE FILE

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]



## **Description**

Show (i.e. get) HDF5 file currently in use.

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The HDF5 file currently in use as an [HDFQL\\_VARCHAR](#) or nothing (in case no file is in use).

## **Example(s)**

```
# show (i.e. get) HDF5 file currently in use (i.e. open) (should be empty)
SHOW USE FILE

# use (i.e. open) four HDF5 files named "my_file0.h5", "my_file1.h5", "my_file2.h5" and
"my_file3.h5"
USE FILE my_file0.h5, my_file1.h5, my_file2.h5, my_file3.h5

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be "my_file3.h5")
SHOW USE FILE

# close HDF5 file currently in use (i.e. file "my_file3.h5")
CLOSE FILE

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be my_"file2.h5")
SHOW USE FILE

# close HDF5 file "my_file1.h5"
CLOSE FILE my_file1.h5

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be "my_file2.h5")
SHOW USE FILE
```

```
# close HDF5 file currently in use (i.e. file "my_file2.h5")
CLOSE FILE

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be "my_file0.h5")
SHOW USE FILE

# close HDF5 file currently in use (i.e. file "my_file0.h5")
CLOSE FILE

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be empty)
SHOW USE FILE
```

## 6.7.4 SHOW ALL USE FILE

### Syntax

#### SHOW ALL USE FILE

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### Description

Show (i.e. get) all HDF5 files in use (i.e. open).

### Parameter(s)

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

All HDF5 files in use (i.e. open) as an [HDFQL\\_VARCHAR](#) or nothing (in case no files are in use).

## **Example(s)**

```
# show (i.e. get) all HDF5 files in use (i.e. open) (should be empty)
SHOW ALL USE FILE

# use (i.e. open) three HDF5 files named "my_file0.h5", "my_file1.h5" and "my_file2.h5"
USE FILE my_file0.h5, my_file1.h5, my_file2.h5

# show (i.e. get) all HDF5 files in use (i.e. open) (should be "my_file2.h5", "my_file1.h5",
"my_file0.h5")
SHOW ALL USE FILE

# close all HDF5 files in use (i.e. open)
CLOSE ALL FILE

# show (i.e. get) all HDF5 files in use (i.e. open) (should be empty)
SHOW ALL USE FILE
```

## **6.7.5 SHOW USE GROUP**

### **Syntax**

#### **SHOW USE GROUP**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### **Description**

Show (i.e. get) HDF5 group currently in use.

### **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The HDF5 group currently in use as an [HDFQL\\_VARCHAR](#) or nothing (in case no file is in use).

## **Example(s)**

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5

# show (i.e. get) current working group (should be "/")
SHOW USE GROUP

# create an HDF5 group named "my_group"
CREATE GROUP my_group

# set group currently in use to "my_group" (more precisely "/my_group")
USE GROUP my_group

# show (i.e. get) current working group (should be "/my_group")
SHOW USE GROUP

# create two HDF5 groups named "my_subgroup0" and "my_subgroup1" (both groups will be created
in group "/my_group")
CREATE GROUP my_subgroup0, my_subgroup1

# set group currently in use to "my_subgroup0" (more precisely "/my_group/my_subgroup0")
USE GROUP my_subgroup0

# show (i.e. get) current working group (should be "/my_group/my_subgroup0")
SHOW USE GROUP

# set group currently in use to "." (the group currently in use will not change as "." refers
to the current working group itself)
USE GROUP .

# show (i.e. get) current working group (should be "/my_group/my_subgroup0")
SHOW USE GROUP

# set group currently in use to "my_subgroup1" located one level up (more precisely
```

```

"/my_group/my_subgroup1")
USE GROUP ../my_subgroup1

# set group currently in use two levels up (should be "/")
USE GROUP ../../

```

## 6.7.6 SHOW [GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK]

### Syntax

**SHOW** [GROUP | DATASET | ATTRIBUTE | [SOFT] LINK | EXTERNAL LINK] *[[file\_name] object\_name]*

*[LIKE regular\_expression [DEEP deep\_value [, deep\_value]\*]]*

*[WHERE condition]*

*[ORDER CREATION]*

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### Description

Show (i.e. get) HDF5 objects (i.e. groups, datasets, attributes, (soft) links or external links) within an HDF5 group or dataset named *object\_name* or check the existence of an object named *object\_name*. If *object\_name* is not specified, all objects are returned – to return only objects of type group, dataset, attribute, (soft) link or external link, specify the keyword GROUP, DATASET, ATTRIBUTE, [SOFT] LINK or EXTERNAL LINK respectively. Otherwise, if *object\_name* is specified and the keyword LIKE is not specified, one of the following behaviors applies:

- If it ends with *"/*, *object\_name* will be treated as a group or dataset, and all groups, datasets or attributes stored in *object\_name* are returned.
- If it does not end with *"/*, *object\_name* will be checked for its existence. If it does exist, *object\_name* is returned; otherwise, if it does not exist, an error is raised.

If the keyword LIKE is specified, only objects with names complying with a regular expression named *regular\_expression* will be returned (in HDFql, regular expressions are the ones specified by PCRE which closely follow PERL5 syntax – please refer to <http://www.pcre.org> and <http://perldoc.perl.org/perlre.html> for additional information). If *regular\_expression*

includes `***`, recursive search is performed (i.e. HDFql will search in all existing groups and subgroups for objects). To limit the recursiveness, the keyword `DEEP` may be specified along with a value *deep\_value* representing the maximum recursiveness limit.

A special type of ordering can be performed using the keyword `ORDER CREATION` allowing HDF5 objects (i.e. groups, datasets, attributes, (soft) links or external links) to be returned according to their time of creation<sup>45</sup> – in contrast to the default behavior which returns objects in an ascending order.

### **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file which stores the objects (i.e. groups, datasets, attributes, (soft) links or external links) to show (i.e. `get`) or check for their existence. If *file\_name* is specified, the file is opened on the fly, the objects are obtained or checked for their existence and, afterwards, the file is closed. Otherwise, if it is not specified, the objects to obtain or check for their existence are stored in the file currently in use.

*object\_name* – optional string that specifies the name of the HDF5 group or dataset to show (i.e. `get`) existing objects (i.e. groups, datasets, attributes, (soft) links or external links) within *object\_name* or check the existence of an object named *object\_name*.

*regular\_expression* – optional string that specifies the regular expression which only names of objects that comply with it are returned. If *regular\_expression* includes `***`, recursive search is performed.

*deep\_value* – optional integer that specifies the maximum recursiveness limit.

*condition* – to be defined.

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

---

<sup>45</sup> This assumes that the HDF5 group or dataset storing the objects was created with the option of tracking objects by their time of creation. Otherwise, if the group or dataset was not created with the option of tracking objects by their time of creation, the keyword `ORDER CREATION` is ignored. Please refer to the [CREATE GROUP](#) and [CREATE DATASET](#) operations for additional information.

## Return

The HDF5 objects (i.e. groups, datasets, attributes, (soft) links or external links) within an HDF5 group or dataset or the existence of an object as an [HDFQL\\_VARCHAR](#).

## Example(s)

```
# set group currently in use to "/" (i.e. the root of the HDF5 file)
USE GROUP /

# create two HDF5 groups named "my_group0" and "my_group1" (both groups will be created in
group "/")
CREATE GROUP my_group0, my_group1

# create one HDF5 dataset named "my_dataset0" of data type unsigned short (it will be created
in group "/")
CREATE DATASET my_dataset0 AS UNSIGNED SMALLINT

# create one HDF5 dataset named "my_dataset1" of data type short (it will be created in group
"/my_group0")
CREATE DATASET my_group0/my_dataset1 AS SMALLINT

# create two HDF5 attributes named "my_attribute0" and "my_attribute1" of data type long long
(both attributes will be created in group "/")
CREATE ATTRIBUTE my_attribute0, my_attribute1 AS BIGINT

# create one HDF5 attribute named "my_attribute2" of data type char (it will be created in
group "/my_group0")
CREATE ATTRIBUTE my_group0/my_attribute2 AS TINYINT

# create one HDF5 attribute named "my_attribute3" of data type unsigned char (it will be
created in dataset "/my_dataset0")
CREATE ATTRIBUTE my_dataset0/my_attribute3 AS UNSIGNED TINYINT

# show (i.e. get) all HDF5 objects existing in group "/" (should be "my_group0", "my_group1",
"my_dataset0", "my_attribute0", "my_attribute1")
SHOW

# show (i.e. get) all HDF5 groups existing in group "/" (should be "my_group0", "my_group1")
SHOW GROUP

# show (i.e. get) all HDF5 datasets existing in group "/" (should be "my_dataset0")
SHOW DATASET
```

```
# check if HDF5 object "my_groupX" exists (should raise an error)
SHOW my_groupX

# check if HDF5 object "my_group0" exists (should be "my_group0")
SHOW my_group0

# show (i.e. get) all HDF5 objects existing within group "my_group0" (should be "my_dataset1",
"my_attribute2")
SHOW my_group0/

# show (i.e. get) all HDF5 attributes existing within group "my_group0" (should be
"my_attribute2")
SHOW ATTRIBUTE my_group0/

# show (i.e. get) all HDF5 objects existing within dataset "my_dataset0" (should be
"my_attribute3")
SHOW my_dataset0/
```

```
# create an HDF5 group named "my_group1" that tracks the objects' (i.e. groups and datasets)
creation order within the group
CREATE GROUP my_group1 ORDER TRACKED

# create two HDF5 groups named "my_subgroup1" and "my_subgroup0" (both groups will be created
in group "/my_group1")
CREATE GROUP my_group1/my_subgroup1, my_group1/my_subgroup0

# create two HDF5 datasets named "my_dataset1" and "my_dataset0" of data type float (both
datasets will be created in group "/my_group1")
CREATE DATASET my_group1/my_dataset1, my_group1/my_dataset0 AS FLOAT

# show (i.e. get) all HDF5 objects existing within group "my_group1" (should be "my_dataset0",
"my_dataset1", "my_subgroup0", "my_subgroup1")
SHOW my_group1/

# show (i.e. get) all HDF5 objects existing within group "my_group1" ordered by their time of
creation (should be "my_subgroup1", "my_subgroup0", "my_dataset1", "my_dataset0")
SHOW my_group1/ ORDER CREATION

# create an HDF5 dataset named "my_dataset1" of data type double that tracks the attributes'
creation order within the dataset
CREATE DATASET my_dataset1 AS DOUBLE ATTRIBUTE ORDER TRACKED
```



```
# create two HDF5 attributes named "my_attribute2" and "my_attribute0" of data type int (both
attributes will be created in dataset "/my_dataset1")
CREATE ATTRIBUTE my_dataset1/my_attribute2, my_dataset1/my_attribute0 AS INT

# create an HDF5 attribute named "my_attribute1" of data type short (it will be created in
dataset "/my_dataset1")
CREATE ATTRIBUTE my_dataset1/my_attribute1 AS SMALLINT

# show (i.e. get) all HDF5 objects existing within dataset "my_dataset1" (should be
"my_attribute0", "my_attribute1", "my_attribute2")
SHOW my_dataset1/

# show (i.e. get) all HDF5 objects existing within dataset "my_dataset1" ordered by their time
of creation (should be "my_attribute2", "my_attribute0", "my_attribute1")
SHOW my_dataset1/ ORDER CREATION
```

```
# create an HDF5 group named "my_group2"
CREATE GROUP my_group2

# create two HDF5 groups named "my_subgroup0" and "my_subgroup1" (both groups will be created
in group "/my_group2")
CREATE GROUP my_group2/my_subgroup0, my_group2/my_subgroup1

# create three HDF5 groups in one go named "my_group3" (in root group "/"), "my_subgroup0" (in
group "my_group3") and "my_subsubgroup0" (in group "my_group3/my_subgroup0")
CREATE GROUP my_group3/my_subgroup0/my_subsubgroup0

# create an HDF5 dataset named "my_dataset2" (in root group "/") of data type double
CREATE DATASET my_dataset2 AS DOUBLE

# create an HDF5 dataset named "my_dataset0" (in group "my_group2") of data type int
CREATE DATASET my_group2/my_dataset0 AS INT

# create an HDF5 dataset named "my_dataset1" (in group "my_group2") of data type short
CREATE DATASET my_group2/my_dataset1 AS SMALLINT

# create an HDF5 dataset named "my_dataset0" (in group "my_group3") of data type float
CREATE DATASET my_group3/my_dataset0 AS FLOAT

# create an HDF5 dataset named "my_dataset0" (in group "my_group3/my_subgroup0") of data type
char
```

```
CREATE DATASET my_group3/my_subgroup0/my_dataset0 AS TINYINT

# create an HDF5 attribute named "my_attribute3" (in group "/") of data type long long
CREATE ATTRIBUTE my_attribute3 AS BIGINT

# create an HDF5 attribute named "my_attribute4" (in group "/") of data type unsigned int
CREATE ATTRIBUTE my_attribute4 AS UNSIGNED INT

# create two HDF5 attributes in one go that are both named "my_attribute0" (one in group
"my_group2" and the other in "my_group3") of data type variable float
CREATE ATTRIBUTE my_group2/my_attribute0, my_group3/my_attribute0 AS VARFLOAT

# create an HDF5 attribute named "my_attribute0" (in dataset "my_dataset2") of data type
variable char
CREATE ATTRIBUTE my_dataset2/my_attribute0 AS VARCHAR

# show (i.e. get) all HDF5 objects from group "/" that has "3" in their names (should be
"my_attribute3", "my_group3")
SHOW LIKE 3

# show (i.e. get) all HDF5 attributes from group "/" that has "3" in their names (should be
"my_attribute3")
SHOW ATTRIBUTE LIKE 3

# show (i.e. get) all HDF5 objects recursively starting from group "/" (should be
"my_attribute3", "my_attribute4", "my_dataset2", "my_dataset2/my_attribute0", "my_group2",
"my_group2/my_attribute0", "my_group2/my_dataset0", "my_group2/my_dataset1",
"my_group2/my_subgroup0", "my_group2/my_subgroup1", "my_group3", "my_group3/my_attribute0",
"my_group3/my_dataset0", "my_group3/my_subgroup0", "my_group3/my_subgroup0/my_dataset0",
"my_group3/my_subgroup0/my_subsubgroup0")
SHOW LIKE **

# show (i.e. get) all HDF5 datasets recursively starting from group "/" (should be
"my_dataset2", "my_group2/my_dataset0", "my_group2/my_dataset1", "my_group3/my_dataset0",
"my_group3/my_subgroup0/my_dataset0")
SHOW DATASET LIKE **

# show (i.e. get) all HDF5 datasets recursively starting from group "/" and one level deep at
most (should be "my_dataset2", "my_group2/my_dataset0", "my_group2/my_dataset1",
"my_group3/my_dataset0")
SHOW DATASET LIKE ** DEEP 1

# show (i.e. get) all HDF5 objects recursively starting from group "my_group3" (should be
"my_attribute0", "my_dataset0", "my_subgroup0", "my_subgroup0/my_dataset0",
```

```
"my_subgroup0/my_subsubgroup0")
SHOW my_group3 LIKE **

# show (i.e. get) all HDF5 groups recursively starting from group "my_group3" (should be
"my_subgroup0", "my_subgroup0/my_subsubgroup0")
SHOW GROUP my_group3 LIKE **

# show (i.e. get) all HDF5 objects recursively starting from group "/" that has "2" in their
names (should be "my_dataset2", "my_group2")
SHOW LIKE **/2

# show (i.e. get) all HDF5 groups recursively starting from group "/" that has "1" or "2" in
their names (should be "my_group2", "my_group2/my_subgroup1")
SHOW GROUP LIKE **/1|2

# show (i.e. get) all HDF5 objects recursively starting from group "/" that starts with "sub"
in their names (should be "my_group2/my_subgroup0", "my_group2/my_subgroup1",
"my_group3/my_subgroup0", "my_group3/my_subgroup0/my_subsubgroup0")
SHOW LIKE **/^my_sub
```

```
# use (i.e. open) an HDF5 file named "my_file.h5"
USE FILE my_file.h5

# create an HDF5 group named "my_group3" in the HDF5 file currently in use (i.e. file
"my_file.h5")
CREATE GROUP my_group3

# create two HDF5 datasets named "my_dataset3" and "my_dataset4" of data type double both in
the HDF5 file currently in use (i.e. file "my_file.h5")
CREATE DATASET my_dataset3, my_dataset4 AS DOUBLE

# show (i.e. get) all HDF5 objects existing in group "/" of the HDF5 file currently in use
(i.e. file "my_file.h5") (should be "my_group3", "my_dataset3", "my_dataset4")
SHOW /

# close HDF5 file currently in use (i.e. file "my_file.h5")
CLOSE FILE

# show (i.e. get) all HDF5 objects existing in group "/" of file "my_file.h5" (should be
"my_group3", "my_dataset3", "my_dataset4")
SHOW my_file.h5 /
```

## 6.7.7 SHOW TYPE

### Syntax

**SHOW TYPE** [*file\_name*] *object\_name* [, [*file\_name*] *object\_name*]\*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) type of an object named *object\_name*. Multiple objects' types can be obtained at once by separating these with a comma (,). If *object\_name* was not found or its type could not be checked (due to unknown/unexpected reasons), no subsequent objects are checked, and an error is raised.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the object to show (i.e. get) the type. If *file\_name* is specified, the file is opened on the fly, the type of the object is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the object (whose type is to be obtained) is stored in the file currently in use.

*object\_name* – name of the object whose type is to be obtained. Multiple objects are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

The type of an object as an [HDFQL\\_INT](#), which can either be [HDFQL\\_GROUP](#), [HDFQL\\_DATASET](#), [HDFQL\\_ATTRIBUTE](#), [HDFQL\\_GROUP | HDFQL\\_SOFT\\_LINK](#), [HDFQL\\_DATASET | HDFQL\\_SOFT\\_LINK](#), [HDFQL\\_GROUP | HDFQL\\_EXTERNAL\\_LINK](#), or [HDFQL\\_DATASET | HDFQL\\_EXTERNAL\\_LINK](#) depending on whether the object is a group, dataset, attribute, group and (soft) link at the same time, dataset and (soft) link at the same time, group and external link at the same time, or dataset and external link at the same time, respectively.

## **Example(s)**

```
# create an HDF5 group named "my_object0"
CREATE GROUP my_object0

# create an HDF5 dataset named "my_object1" of data type double
CREATE DATASET my_object1 AS DOUBLE

# create an HDF5 attribute named "my_object2" of data type float
CREATE ATTRIBUTE my_object2 AS FLOAT

# create an HDF5 soft link named "my_object3" to object "my_object0"
CREATE SOFT LINK my_object3 TO my_object0

# create an HDF5 external link named "my_object4" to object "my_object" (assumed to be a
dataset) in file "my_file.h5"
CREATE EXTERNAL LINK my_object4 TO my_file.h5 my_object

# show (i.e. get) type of object "my_object0" (should be 4 - i.e. HDFQL_GROUP)
SHOW TYPE my_object0

# show (i.e. get) type of object "my_object1" (should be 8 - i.e. HDFQL_DATASET)
SHOW TYPE my_object1

# show (i.e. get) type of object "my_object2" (should be 16 - i.e. HDFQL_ATTRIBUTE)
SHOW TYPE my_object2

# show (i.e. get) type of both objects "my_object0" and "my_object2" at once (should be 4, 16)
SHOW TYPE my_object0, my_object2

# show (i.e. get) type of object "my_object3" (should be 36 - i.e. HDFQL_GROUP |
HDFQL_SOFT_LINK)
SHOW TYPE my_object3

# show (i.e. get) type of object "my_object4" (should be 136 - i.e. HDFQL_DATASET |
HDFQL_EXTERNAL_LINK)
SHOW TYPE my_object4
```

## 6.7.8 SHOW DATA TYPE

### Syntax

```
SHOW [DATASET | ATTRIBUTE] DATA TYPE [file_name] object_name [, [file_name] object_name]*  
  
    [post_processing_option [post_processing_option]*]  
  
    [output_redirecting_option]
```

### Description

Show (i.e. get) data type of an HDF5 dataset or attribute named *object\_name*. Multiple objects' data types can be obtained at once by separating these with a comma (,). If *object\_name* was not found or its data type could not be checked (due to unknown/unexpected reasons), no subsequent objects are checked, and an error is raised. In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword DATASET nor ATTRIBUTE is specified, the data type returned belongs to the dataset. To explicitly get the data type of *object\_name* according to its type, the keyword DATASET or ATTRIBUTE must be specified.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset or attribute to show (i.e. get) the data type. If *file\_name* is specified, the file is opened on the fly, the data type of the dataset or attribute is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose data type is to be obtained) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute whose data type is to be obtained. Multiple datasets or attributes are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## Return

The data type of an HDF5 dataset or attribute as an [HDFQL\\_INT](#), which can either be [HDFQL\\_TINYINT](#), [HDFQL\\_UNSIGNED\\_TINYINT](#), [HDFQL\\_SMALLINT](#), [HDFQL\\_UNSIGNED\\_SMALLINT](#), [HDFQL\\_INT](#), [HDFQL\\_UNSIGNED\\_INT](#), [HDFQL\\_BIGINT](#), [HDFQL\\_UNSIGNED\\_BIGINT](#), [HDFQL\\_FLOAT](#), [HDFQL\\_DOUBLE](#), [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#), [HDFQL\\_VARCHAR](#), [HDFQL\\_OPAQUE](#), [HDFQL\\_BITFIELD](#), [HDFQL\\_ENUMERATION](#), [HDFQL\\_COMPOUND](#) or [HDFQL\\_UNDEFINED](#) (please refer to [Table 6.3](#) for additional information about data types).

## Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type double
CREATE DATASET my_dataset0 AS DOUBLE

# show (i.e. get) data type of dataset "my_dataset0" (should be 512 - i.e. HDFQL_DOUBLE)
SHOW DATA TYPE my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type float
CREATE DATASET my_dataset1 AS FLOAT

# show (i.e. get) data type of dataset "my_dataset1" (should be 256 - i.e. HDFQL_FLOAT)
SHOW DATA TYPE my_dataset1

# create an HDF5 dataset named "my_common" of data type short
CREATE DATASET my_common AS SMALLINT

# create an HDF5 attribute named "my_common" of data type int
CREATE ATTRIBUTE my_common AS INT

# show (i.e. get) data type of dataset "my_common" (should be 4 - i.e. HDFQL_SMALLINT)
SHOW DATA TYPE my_common

# show (i.e. get) data type of dataset "my_common" (should be 4 - i.e. HDFQL_SMALLINT)
SHOW DATASET DATA TYPE my_common

# show (i.e. get) data type of attribute "my_common" (should be 16 - i.e. HDFQL_INT)
SHOW ATTRIBUTE DATA TYPE my_common
```

## 6.7.9 SHOW MASK

### Syntax

```
SHOW MASK [file_name] dataset_name[(chunk_number [, chunk_number]*)] [, [file_name]  
dataset_name[(chunk_number [, chunk_number]*)]]*  
  
[post_processing_option [post_processing_option]*]  
  
[output_redirecting_option]
```

### Description

Show (i.e. get) (filter) mask of an HDF5 dataset named *dataset\_name*. Multiple datasets' masks can be obtained at once by separating these with a comma (,). If *dataset\_name* was not found or its mask could not be checked (due to unknown/unexpected reasons), no subsequent datasets are checked, and an error is raised.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset to show (i.e. get) the (filter) mask. If *file\_name* is specified, the file is opened on the fly, the mask of the dataset is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset (whose mask is to be obtained) is stored in the file currently in use.

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset whose (filter) mask is to be obtained. Multiple datasets are separated with a comma (,).

*chunk\_number* – optional integer that specifies the number of the chunk to show (i.e. get) its (filter) mask. Multiple chunk numbers are separated with a comma (,). If *chunk\_number* is specified it must either be between 0 and the storage dimension in question - 1 (otherwise an error will be raised) or negative (in this case its value will be the last position of the storage dimension in question minus the value of *chunk\_number*). Otherwise, if it is not specified, its default value is 0 (i.e. first chunk of the storage dimension in question).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).



## Return

The (filter) mask of an HDF5 dataset as an [HDFQL\\_UNSIGNED\\_INT](#).

## Example(s)

```
# create a chunked (size 2) HDF5 dataset named "my_dataset" of data type int of one dimension
(size 6)
CREATE CHUNKED(2) DATASET my_dataset AS INT(6)

# insert (i.e. write) values 60 and 61 directly into chunk #0 of dataset "my_dataset" using a
(filter) mask equal to 8
INSERT DIRECTLY MASK 8 INTO my_dataset VALUES(60, 61)

# insert (i.e. write) values 62 and 63 directly into chunk #1 of dataset "my_dataset" using a
(filter) mask equal to 255 (i.e. 0xFF)
INSERT DIRECTLY MASK 0xFF INTO my_dataset(1) VALUES(62, 63)

# insert (i.e. write) values 64 and 65 directly into chunk #2 of dataset "my_dataset" using a
(filter) mask equal to 0 (i.e. default value)
INSERT DIRECTLY INTO my_dataset(2) VALUES(64, 65)

# select (i.e. read) data from dataset "my_dataset" and populate cursor in use with it (should
be 60, 61, 62, 63, 64, 65)
SELECT FROM my_dataset

# show (i.e. get) (filter) mask of chunks #0, #1 and #2 of dataset "my_dataset" and populate
cursor in use with it (should be 8, 255, 0)
SHOW MASK my_dataset(0), my_dataset(1), my_dataset(2)
```

## 6.7.10 SHOW ENDIANNESS

### Syntax

**SHOW** [DATASET | ATTRIBUTE] **ENDIANNESS** [*file\_name*] *object\_name* [, [*file\_name*] *object\_name*]\*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

## **Description**

Show (i.e. get) endianness of an HDF5 dataset or attribute named *object\_name*. Multiple objects' endiannesses can be obtained at once by separating these with a comma (,). If *object\_name* was not found or its endianness could not be checked (due to unknown/unexpected reasons), no subsequent objects are checked, and an error is raised. In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword DATASET nor ATTRIBUTE is specified, the endianness returned belongs to the dataset. To explicitly get the endianness of *object\_name* according to its type, the keyword DATASET or ATTRIBUTE must be specified.

## **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset or attribute to show (i.e. get) the endianness. If *file\_name* is specified, the file is opened on the fly, the endianness of the dataset or attribute is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose endianness is to be obtained) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute whose endianness is to be obtained. Multiple datasets or attributes are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The endianness of an HDF5 dataset or attribute as an [HDFQL\\_INT](#), which can either be [HDFQL\\_LITTLE\\_ENDIAN](#), [HDFQL\\_BIG\\_ENDIAN](#) or [HDFQL\\_UNDEFINED](#) depending on whether the endianness is little, big or undefined (i.e. endianness is not applicable to *object\_name*) respectively.

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type int using the native endian
representation (of the machine)
CREATE DATASET my_dataset0 AS INT
```

```

# show (i.e. get) endianness of dataset "my_dataset0" (should be 1 or 2 - i.e.
HDFQL_LITTLE_ENDIAN or HDFQL_BIG_ENDIAN - depending on whether the dataset was created in a
little or big endian machine respectively)
SHOW ENDIANNESS my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type long long using the little endian
representation
CREATE DATASET my_dataset1 AS LITTLE_ENDIAN BIGINT

# show (i.e. get) endianness of dataset "my_dataset1" (should be 1 - i.e. HDFQL_LITTLE_ENDIAN)
SHOW ENDIANNESS my_dataset1

# create an HDF5 dataset named "my_common" of data type short using the big endian
representation
CREATE DATASET my_common AS BIG_ENDIAN SMALLINT

# create an HDF5 attribute named "my_common" of data type int using the little endian
representation
CREATE ATTRIBUTE my_common AS LITTLE_ENDIAN INT

# show (i.e. get) endianness of dataset "my_common" (should be 2 - i.e. HDFQL_BIG_ENDIAN)
SHOW ENDIANNESS my_common

# show (i.e. get) endianness of dataset "my_common" (should be 2 - i.e. HDFQL_BIG_ENDIAN)
SHOW DATASET ENDIANNESS my_common

# show (i.e. get) endianness of attribute "my_common" (should be 1 - i.e. HDFQL_LITTLE_ENDIAN)
SHOW ATTRIBUTE ENDIANNESS my_common

```

## 6.7.11 SHOW CHARSET

### Syntax

**SHOW [DATASET | ATTRIBUTE] CHARSET** [*file\_name*] *object\_name* [, [*file\_name*] *object\_name*]\*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

## **Description**

Show (i.e. get) charset of an HDF5 dataset or attribute named *object\_name*. Multiple objects' charsets can be obtained at once by separating these with a comma (,). If *object\_name* was not found or its charset could not be checked (due to unknown/unexpected reasons), no subsequent objects are checked, and an error is raised. In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword DATASET nor ATTRIBUTE is specified, the charset returned belongs to the dataset. To explicitly get the charset of *object\_name* according to its type, the keyword DATASET or ATTRIBUTE must be specified.

## **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset or attribute to show (i.e. get) the charset. If *file\_name* is specified, the file is opened on the fly, the charset of the dataset or attribute is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose charset is to be obtained) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute whose charset is to be obtained. Multiple datasets or attributes are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The charset of an HDF5 dataset or attribute as an [HDFQL\\_INT](#), which can either be [HDFQL\\_ASCII](#), [HDFQL\\_UTF8](#) or [HDFQL\\_UNDEFINED](#) depending on whether the charset is ASCII, UTF8 or undefined (i.e. *object\_name* is neither of data type [HDFQL\\_CHAR](#) nor [HDFQL\\_VARCHAR](#)) respectively.

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type char
CREATE DATASET my_dataset0 AS CHAR

# show (i.e. get) charset of dataset "my_dataset0" (should be 1 - i.e. HDFQL_ASCII)
```

```
SHOW CHARSET my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type char of one dimension (size 20)
CREATE DATASET my_dataset1 AS UTF8 CHAR(20)

# show (i.e. get) charset of dataset "my_dataset1" (should be 2 - i.e. HDFQL_UTF8)
SHOW CHARSET my_dataset1

# create an HDF5 dataset named "my_common" of data type short
CREATE DATASET my_common AS UTF8 CHAR

# create an HDF5 attribute named "my_common" of data type variable-length char
CREATE ATTRIBUTE my_common AS ASCII VARCHAR

# show (i.e. get) charset of dataset "my_common" (should be 2 - i.e. HDFQL_UTF8)
SHOW CHARSET my_common

# show (i.e. get) data type of dataset "my_common" (should be 2 - i.e. HDFQL_UTF8)
SHOW DATASET CHARSET my_common

# show (i.e. get) charset of attribute "my_common" (should be 1 - i.e. HDFQL_ASCII)
SHOW ATTRIBUTE CHARSET my_common
```

## 6.7.12 SHOW STORAGE TYPE

### Syntax

**SHOW STORAGE TYPE** [*file\_name*] *dataset\_name* [, [*file\_name*] *dataset\_name*]\*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) storage type (layout) of an HDF5 dataset named *dataset\_name*. Multiple datasets' storage types can be obtained at once by separating these with a comma (,). If *dataset\_name* was not found or its storage type could not be checked (due to unknown/unexpected reasons), no subsequent datasets are checked, and an error is raised.

## **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset to show (i.e. get) the storage type (layout). If *file\_name* is specified, the file is opened on the fly, the storage type of the dataset is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset (whose storage type is to be obtained) is stored in the file currently in use.

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset whose storage type (layout) is to be obtained. Multiple datasets are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The storage type (layout) of an HDF5 dataset as an [HDFQL\\_INT](#), which can either be [HDFQL\\_CONTIGUOUS](#), [HDFQL\\_COMPACT](#) or [HDFQL\\_CHUNKED](#) depending on whether the storage is contiguous, compact or chunked respectively.

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type unsigned int
CREATE DATASET my_dataset0 AS UNSIGNED INT

# show (i.e. get) storage type (layout) of dataset "my_dataset0" (should be 1 - i.e.
HDFQL_CONTIGUOUS)
SHOW STORAGE TYPE my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type int of two dimensions (size 5x7)
CREATE CONTIGUOUS DATASET my_dataset1 AS INT(5, 7)

# show (i.e. get) storage type (layout) of dataset "my_dataset1" (should be 1 - i.e.
HDFQL_CONTIGUOUS)
SHOW STORAGE TYPE my_dataset1
```

```
# create an HDF5 dataset named "my_dataset2" of data type double of one dimension (size 8)
CREATE COMPACT DATASET my_dataset2 AS DOUBLE(8)

# show (i.e. get) storage type (layout) of dataset "my_dataset2" (should be 2 - i.e.
HDFQL_COMPACT)
SHOW STORAGE TYPE my_dataset2

# create an HDF5 dataset named "my_dataset3" of data type float of three dimensions (size
3x5x20)
CREATE CHUNKED DATASET my_dataset3 AS FLOAT(3, 5, 20)

# show (i.e. get) storage type (layout) of dataset "my_dataset3" (should be 4 - i.e.
HDFQL_CHUNKED)
SHOW STORAGE TYPE my_dataset3
```

## 6.7.13 SHOW STORAGE ALLOCATION

### Syntax

**SHOW STORAGE ALLOCATION** [*file\_name*] *dataset\_name* [, [*file\_name*] *dataset\_name*]\*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) storage allocation of an HDF5 dataset named *dataset\_name*. Multiple datasets' storage allocation can be obtained at once by separating these with a comma (,). If *dataset\_name* was not found or its storage allocation could not be checked (due to unknown/unexpected reasons), no subsequent datasets are checked, and an error is raised.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset to show (i.e. get) the storage allocation. If *file\_name* is specified, the file is opened on the fly, the storage allocation of the dataset is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset (whose storage allocation is to be obtained) is stored in the file currently in use.

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset whose storage allocation is to be obtained. Multiple datasets are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The storage allocation of an HDF5 dataset as an [HDFQL\\_INT](#), which can either be [HDFQL\\_EARLY](#), [HDFQL\\_INCREMENTAL](#) or [HDFQL\\_LATE](#) depending on whether the storage allocation is early, incremental or late respectively.

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type unsigned int
CREATE DATASET my_dataset0 AS UNSIGNED INT

# show (i.e. get) storage allocation of dataset "my_dataset0" (should be 4 - i.e. HDFQL_LATE)
SHOW STORAGE ALLOCATION my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type int of two dimensions (size 5x7)
CREATE CONTIGUOUS DATASET my_dataset1 AS INT(5, 7)

# show (i.e. get) storage allocation of dataset "my_dataset1" (should be 4 - i.e. HDFQL_LATE)
SHOW STORAGE ALLOCATION my_dataset1

# create an HDF5 dataset named "my_dataset2" of data type double of one dimension (size 8)
CREATE COMPACT DATASET my_dataset2 AS DOUBLE(8)

# show (i.e. get) storage allocation of dataset "my_dataset2" (should be 1 - i.e. HDFQL_EARLY)
SHOW STORAGE ALLOCATION my_dataset2

# create an HDF5 dataset named "my_dataset3" of data type float of three dimensions (size
3x5x20)
CREATE CHUNKED DATASET my_dataset3 AS FLOAT(3, 5, 20)

# show (i.e. get) storage allocation of dataset "my_dataset3" (should be 2 - i.e.
HDFQL_INCREMENTAL)
SHOW STORAGE ALLOCATION my_dataset3
```



## 6.7.14 SHOW STORAGE DIMENSION

### Syntax

**SHOW STORAGE DIMENSION** [*file\_name*] *dataset\_name*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) storage dimensions of an HDF5 dataset named *dataset\_name*.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset to show (i.e. get) the storage dimensions. If *file\_name* is specified, the file is opened on the fly, the storage dimensions of the dataset is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset (whose storage dimensions is to be obtained) is stored in the file currently in use.

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset whose storage dimensions are to be obtained.

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

The storage dimensions of an HDF5 dataset as an [HDFQL\\_BIGINT](#) or nothing (in case the dataset is not chunked – i.e. its storage type is not [HDFQL\\_CHUNKED](#)).

## Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type unsigned int
CREATE DATASET my_dataset0 AS UNSIGNED INT

# show (i.e. get) storage dimensions of dataset "my_dataset0" (should be empty)
SHOW STORAGE DIMENSION my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type int of two dimensions (size 5x7)
CREATE DATASET my_dataset1 AS INT(5, 7)

# show (i.e. get) storage dimensions of dataset "my_dataset1" (should be empty)
SHOW STORAGE DIMENSION my_dataset1

# create an HDF5 dataset named "my_dataset2" of data type double of one dimension (size 8)
CREATE CHUNKED DATASET my_dataset2 AS DOUBLE(8)

# show (i.e. get) storage dimensions of dataset "my_dataset2" (should be 8)
SHOW STORAGE DIMENSION my_dataset2

# create an HDF5 dataset named "my_dataset3" of data type float of three dimensions (size 3x5x20)
CREATE CHUNKED(1, 2, 10) DATASET my_dataset3 AS FLOAT(3, 5, 20)

# show (i.e. get) storage dimensions of dataset "my_dataset3" (should be 1, 2, 10)
SHOW STORAGE DIMENSION my_dataset3
```

## 6.7.15 SHOW DIMENSION

### Syntax

SHOW [DATASET | ATTRIBUTE] [MAX] DIMENSION [*file\_name*] *object\_name*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) dimensions of an HDF5 dataset or attribute named *object\_name*. In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword DATASET nor

ATTRIBUTE is specified, the dimensions returned belong to the dataset. To explicitly get the dimensions of *object\_name* according to its type, the keyword DATASET or ATTRIBUTE must be specified. If *object\_name* does not have a dimension (i.e. if it is scalar), the returned value is one. By default, the returned dimensions refer to the ones that a dataset or an attribute currently has; to return the maximum dimensions that a dataset or an attribute may grow to, the keyword MAX must be specified. If the maximum dimension is unlimited, the returned value is [HDFQL\\_UNLIMITED](#).

### **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset or attribute to show (i.e. get) the dimensions. If *file\_name* is specified, the file is opened on the fly, the dimensions of the dataset or attribute is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose dimensions is to be obtained) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute whose dimensions are to be obtained.

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### **Return**

The dimensions of an HDF5 dataset or attribute as an [HDFQL\\_BIGINT](#).

### **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type unsigned int
CREATE DATASET my_dataset0 AS UNSIGNED INT

# show (i.e. get) dimensions of dataset "my_dataset0" (should be 1)
SHOW DIMENSION my_dataset0

# show (i.e. get) maximum dimensions of dataset "my_dataset0" (should be 1)
SHOW MAX DIMENSION my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type double of one dimension (size 15)
```

```
CREATE DATASET my_dataset1 AS DOUBLE(15)

# show (i.e. get) dimensions of dataset "my_dataset1" (should be 15)
SHOW DIMENSION my_dataset1

# show (i.e. get) maximum dimensions of dataset "my_dataset1" (should be 15)
SHOW MAX DIMENSION my_dataset1

# create an HDF5 attribute named "my_attribute0" of data type int of one dimension (size 1)
CREATE ATTRIBUTE my_attribute0 AS INT(1)

# show (i.e. get) dimensions of attribute "my_attribute0" (should be 1)
SHOW DIMENSION my_attribute0

# show (i.e. get) maximum dimensions of attribute "my_attribute0" (should be 1)
SHOW MAX DIMENSION my_attribute0

# create an HDF5 attribute named "my_attribute1" of data type short of two dimensions (size 2x3)
CREATE ATTRIBUTE my_attribute1 AS SMALLINT(2, 3)

# show (i.e. get) dimensions of attribute "my_attribute1" (should be 2, 3)
SHOW DIMENSION my_attribute1

# show (i.e. get) maximum dimensions of attribute "my_attribute1" (should be 2, 3)
SHOW MAX DIMENSION my_attribute1

# create an HDF5 dataset named "my_dataset2" of data type float of three dimensions (first
dimension with size 2 and extendible up to 10; second dimension with size 5; third dimension
with size 20 and extendible to an unlimited size)
CREATE CHUNKED DATASET my_dataset2 AS FLOAT(3 TO 10, 5, 20 TO UNLIMITED, UNLIMITED)

# show (i.e. get) dimensions of dataset "my_dataset2" (should be 3, 5, 20, 1)
SHOW DIMENSION my_dataset2

# show (i.e. get) maximum dimensions of dataset "my_dataset2" (should be 10, 5, -1, -1)
SHOW MAX DIMENSION my_dataset2
```

## 6.7.16 SHOW ORDER

### Syntax

```
SHOW [ATTRIBUTE] ORDER [file_name] object_name [, [file_name] object_name]*  
  
    [post_processing_option [post_processing_option]*]  
  
    [output_redirecting_option]
```

### Description

Show (i.e. get) (creation) order strategy of an HDF5 group or dataset named *object\_name*. Multiple objects' order strategies can be obtained at once by separating these with a comma (,). If *object\_name* was not found or its order strategy could not be checked (due to unknown/unexpected reasons), no subsequent objects are checked, and an error is raised. By default, the returned order strategy refers to objects (i.e. groups, datasets, (soft) links or external links) within a group; to return the order strategy of attributes within a group or dataset, the keyword **ATTRIBUTE** must be specified.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the group or dataset to show (i.e. get) the (creation) order strategy. If *file\_name* is specified, the file is opened on the fly, the order strategy of the group or dataset is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the group or dataset (whose order strategy is to be obtained) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 group or dataset whose (creation) order strategy is to be obtained. Multiple groups or datasets are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The (creation) order strategy of an HDF5 group or dataset as an [HDFQL\\_INT](#), which can either be [HDFQL\\_TRACKED](#), [HDFQL\\_INDEXED](#) or [HDFQL\\_UNDEFINED](#) depending on whether the order is tracked, indexed or undefined (i.e. *object\_name* was created without any order strategy) respectively.

## **Example(s)**

```
# create an HDF5 group named "my_group0"
CREATE GROUP my_group0

# show (i.e. get) (creation) order strategy of objects within group "my_group0" (should be -1 -
i.e. HDFQL_UNDEFINED)
SHOW ORDER my_group0

# show (i.e. get) (creation) order strategy of attributes within group "my_group0" (should be -
1 - i.e. HDFQL_UNDEFINED)
SHOW ATTRIBUTE ORDER my_group0

# create an HDF5 group named "my_group1" that tracks both the objects' (i.e. groups and
datasets) and the attributes' creation order within the group
CREATE GROUP my_group1 ORDER TRACKED ATTRIBUTE ORDER INDEXED

# show (i.e. get) (creation) order strategy of objects within group "my_group1" (should be 1 -
i.e. HDFQL_TRACKED)
SHOW ORDER my_group1

# show (i.e. get) (creation) order strategy of attributes within group "my_group1" (should be 2
- i.e. HDFQL_INDEXED)
SHOW ATTRIBUTE ORDER my_group1

# create an HDF5 dataset named "my_dataset0" of data type int that tracks the attributes'
creation order within the dataset
CREATE DATASET my_dataset0 AS INT ATTRIBUTE ORDER TRACKED

# show (i.e. get) (creation) order strategy of attributes within dataset "my_dataset0" (should
be 1 - i.e. HDFQL_TRACKED)
SHOW ATTRIBUTE ORDER my_dataset0

# show (i.e. get) (creation) order strategy of attributes within both group "my_group1" and
dataset "my_dataset0" at once (should be 2, 1)
SHOW ATTRIBUTE ORDER my_group1, my_dataset0
```

## 6.7.17 SHOW TAG

### Syntax

```
SHOW [DATASET | ATTRIBUTE] TAG [file_name] object_name [, [file_name] object_name]*  
  
    [post_processing_option [post_processing_option]*]  
  
    [output_redirecting_option]
```

### Description

Show (i.e. get) tag of an HDF5 dataset or attribute named *object\_name*. Multiple objects' tags can be obtained at once by separating these with a comma (,). If *object\_name* was not found or its tag could not be checked (due to its data type not being [HDFQL\\_OPAQUE](#) or for unknown/unexpected reasons), no subsequent objects are checked, and an error is raised. In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword DATASET nor ATTRIBUTE is specified, the tag returned belongs to the dataset. To explicitly get the tag of *object\_name* according to its type, the keyword DATASET or ATTRIBUTE must be specified.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset or attribute to show (i.e. get) the tag. If *file\_name* is specified, the file is opened on the fly, the tag of the dataset or attribute is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose tag is to be obtained) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute whose tag is to be obtained. Multiple datasets or attributes are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

The tag of an HDF5 dataset or attribute as an [HDFQL\\_VARCHAR](#).

## Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type opaque
CREATE DATASET my_dataset0 AS OPAQUE

# show (i.e. get) tag of dataset "my_dataset0" (should be empty)
SHOW TAG my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type opaque of one dimension (size 15)
with a tag value "my_tag1"
CREATE DATASET my_dataset1 AS OPAQUE(15) TAG my_tag1

# show (i.e. get) tag of dataset "my_dataset1" (should be "my_tag1")
SHOW TAG my_dataset1

# create an HDF5 attribute named "my_attribute0" of data type opaque of two dimensions (size
3x5) with a tag value "Hierarchical Data Format"
CREATE ATTRIBUTE my_attribute0 AS OPAQUE(3, 5) TAG "Hierarchical Data Format"

# show (i.e. get) tag of attribute "my_attribute0" (should be "Hierarchical Data Format")
SHOW TAG my_attribute0

# create an HDF5 dataset named "my_common" of data type opaque with a tag value "Dataset tag"
CREATE DATASET my_common AS OPAQUE TAG "Dataset tag"

# create an HDF5 attribute named "my_common" of data type opaque of one dimension (size 10)
with a tag value "Attribute tag"
CREATE ATTRIBUTE my_common AS OPAQUE(10) TAG "Attribute tag"

# show (i.e. get) tag of dataset "my_common" (should be "Dataset tag")
SHOW TAG my_common

# show (i.e. get) tag of dataset "my_common" (should be "Dataset tag")
SHOW DATASET TAG my_common

# show (i.e. get) tag of attribute "my_common" (should be "Attribute tag")
SHOW ATTRIBUTE TAG my_common
```



## 6.7.18 SHOW FILL TYPE

### Syntax

```
SHOW FILL TYPE [file_name] dataset_name [, [file_name] dataset_name]*  
  
    [post_processing_option [post_processing_option]*]  
  
    [output_redirecting_option]
```

### Description

Show (i.e. get) fill type of an HDF5 dataset named *dataset\_name*. Multiple datasets' fill types can be obtained at once by separating these with a comma (,). If *dataset\_name* was not found or its fill type could not be checked (due to unknown/unexpected reasons), no subsequent datasets are checked, and an error is raised.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset to show (i.e. get) the fill type. If *file\_name* is specified, the file is opened on the fly, the fill type of the dataset is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset (whose fill type is to be obtained) is stored in the file currently in use.

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset whose fill type is to be obtained. Multiple datasets are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

The fill type of an HDF5 dataset as an [HDFQL\\_INT](#), which can either be [HDFQL\\_FILL\\_DEFAULT](#), [HDFQL\\_FILL\\_USER\\_DEFINED](#) or [HDFQL\\_FILL\\_UNDEFINED](#) depending on whether the fill is default, user defined or undefined respectively.

## Example(s)

```
# create an HDF5 dataset named "my_dataset0" of data type int
CREATE DATASET my_dataset0 AS INT

# show (i.e. get) fill type of dataset "my_dataset0" (should be 1 - i.e. HDFQL_FILL_DEFAULT)
SHOW FILL TYPE my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type int with a fill value "20"
CREATE DATASET my_dataset1 AS INT FILL(20)

# show (i.e. get) fill type of dataset "my_dataset1" (should be 2 - i.e.
HDFQL_FILL_USER_DEFINED)
SHOW FILL TYPE my_dataset1

# create an HDF5 dataset named "my_dataset2" of data type variable-length char of one dimension
(size 5) with an undefined fill value
CREATE DATASET my_dataset2 AS VARCHAR(5) FILL UNDEFINED

# show (i.e. get) fill type of dataset "my_dataset2" (should be 4 - i.e. HDFQL_FILL_UNDEFINED)
SHOW FILL TYPE my_dataset2
```

## 6.7.19 SHOW FILL VALUE

### Syntax

SHOW FILL VALUE [*file\_name*] *dataset\_name*

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) fill values of an HDF5 dataset named *dataset\_name*.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset to show (i.e. get) the fill values. If *file\_name* is specified, the file is opened on the fly, the fill values of the dataset are obtained and, afterwards, the

file is closed. Otherwise, if it is not specified, the dataset (whose fill values is to be obtained) is stored in the file currently in use.

*dataset\_name* – mandatory string that specifies the name of the HDF5 dataset whose fill values are to be obtained.

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The fill values of an HDF5 dataset as an [HDFQL\\_TINYINT](#) (in case the data type of the dataset is [HDFQL\\_TINYINT](#)), [HDFQL\\_UNSIGNED\\_TINYINT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_TINYINT](#)), [HDFQL\\_SMALLINT](#) (in case the data type of the dataset is [HDFQL\\_SMALLINT](#)), [HDFQL\\_UNSIGNED\\_SMALLINT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_SMALLINT](#)), [HDFQL\\_INT](#) (in case the data type of the dataset is [HDFQL\\_INT](#)), [HDFQL\\_UNSIGNED\\_INT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_INT](#)), [HDFQL\\_BIGINT](#) (in case the data type of the dataset is [HDFQL\\_BIGINT](#)), [HDFQL\\_UNSIGNED\\_BIGINT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_BIGINT](#)), [HDFQL\\_FLOAT](#) (in case the data type of the dataset is [HDFQL\\_FLOAT](#)), [HDFQL\\_DOUBLE](#) (in case the data type of the dataset is [HDFQL\\_DOUBLE](#)), [HDFQL\\_CHAR](#) (in case the data type of the dataset is [HDFQL\\_CHAR](#)), [HDFQL\\_VARTINYINT](#) (in case the data type of the dataset is [HDFQL\\_VARTINYINT](#)), [HDFQL\\_UNSIGNED\\_VARTINYINT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_VARTINYINT](#)), [HDFQL\\_VARSMALLINT](#) (in case the data type of the dataset is [HDFQL\\_VARSMALLINT](#)), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_VARSMALLINT](#)), [HDFQL\\_VARINT](#) (in case the data type of the dataset is [HDFQL\\_VARINT](#)), [HDFQL\\_UNSIGNED\\_VARINT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_VARINT](#)), [HDFQL\\_VARBIGINT](#) (in case the data type of the dataset is [HDFQL\\_VARBIGINT](#)), [HDFQL\\_UNSIGNED\\_VARBIGINT](#) (in case the data type of the dataset is [HDFQL\\_UNSIGNED\\_VARBIGINT](#)), [HDFQL\\_VARFLOAT](#) (in case the data type of the dataset is [HDFQL\\_VARFLOAT](#)), [HDFQL\\_VARDOUBLE](#) (in case the data type of the dataset is [HDFQL\\_VARDOUBLE](#)), [HDFQL\\_VARCHAR](#) (in case the data type of the dataset is [HDFQL\\_VARCHAR](#)) or [HDFQL\\_OPAQUE](#) (in case the data type of the dataset is [HDFQL\\_OPAQUE](#)).

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type int
CREATE DATASET my_dataset0 AS INT
```

```
# show (i.e. get) fill values of dataset "my_dataset0" (should be 0)
SHOW FILL VALUE my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type int with a fill value "20"
CREATE DATASET my_dataset1 AS INT FILL(20)

# show (i.e. get) fill values of dataset "my_dataset1" (should be 20)
SHOW FILL VALUE my_dataset1

# create an HDF5 dataset named "my_dataset2" of data type variable-length char of one dimension
(size 5) with a fill value "Hierarchical Data Format"
CREATE DATASET my_dataset2 AS VARCHAR(5) FILL("Hierarchical Data Format")

# show (i.e. get) fill values of dataset "my_dataset2" (should be "Hierarchical Data Format")
SHOW FILL VALUE my_dataset2

# create an HDF5 dataset named "my_dataset3" of data type variable-length int with fill values
"100" and "200"
CREATE DATASET my_dataset3 AS VARINT FILL(100, 200)

# show (i.e. get) fill values of dataset "my_dataset3" (should be 100, 200)
SHOW FILL VALUE my_dataset3
```

## 6.7.20 SHOW FILE SIZE

### Syntax

**SHOW FILE SIZE** [*file\_name* [, *file\_name*]\*]

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) size (in bytes) of a file named *file\_name*. Multiple files' sizes can be obtained at once by separating several file names with a comma (.). If *file\_name* is not specified, the size of the (HDF5) file currently in use is returned. Otherwise, if *file\_name* is specified, its size is returned. If *file\_name* was not found or its size could not be checked (due to unknown/unexpected reasons), no subsequent files are checked, and an error is raised.

## **Parameter(s)**

*file\_name* – optional string that specifies the name of the file whose size (in bytes) is to be obtained. Multiple files are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The size (in bytes) of a file as an [HDFQL\\_UNSIGNED\\_BIGINT](#).

## **Example(s)**

```
# create an HDF5 file named "my_file.h5"
CREATE FILE my_file.h5

# show (i.e. get) size of file "my_file.h5" (should be 800)
SHOW FILE SIZE my_file.h5

# use (i.e. open) HDF5 file "my_file.h5"
USE FILE my_file.h5

# create an HDF5 group named "my_group"
CREATE GROUP my_group

# flush the entire virtual HDF5 file (global) currently in use
FLUSH

# show (i.e. get) size of the file currently in use (should be greater than 800)
SHOW FILE SIZE
```

## 6.7.21 SHOW [DATASET | ATTRIBUTE] SIZE

### Syntax

```
SHOW [DATASET | ATTRIBUTE] SIZE [file_name] object_name [, [file_name] object_name]*  
  
    [post_processing_option [post_processing_option]*]  
  
    [output_redirecting_option]
```

### Description

Show (i.e. get) size (in bytes) of an HDF5 dataset or attribute named *object\_name*. Multiple objects' sizes can be obtained at once by separating these with a comma (,). If *object\_name* was not found or its size could not be checked (due to unknown/unexpected reasons), no subsequent objects are checked, and an error is raised. In case a dataset and an attribute with identical names (*object\_name*) are stored in the same location (i.e. group) and neither the keyword DATASET nor ATTRIBUTE is specified, the size returned belongs to the dataset. To explicitly get the size of *object\_name* according to its type, the keyword DATASET or ATTRIBUTE must be specified.

### Parameter(s)

*file\_name* – optional string that specifies the name of the HDF5 file which stores the dataset or attribute to show (i.e. get) the size (in bytes). If *file\_name* is specified, the file is opened on the fly, the size of the dataset or attribute is obtained and, afterwards, the file is closed. Otherwise, if it is not specified, the dataset or attribute (whose size is to be obtained) is stored in the file currently in use.

*object\_name* – mandatory string that specifies the name of the HDF5 dataset or attribute whose size (in bytes) is to be obtained. Multiple datasets or attributes are separated with a comma (,).

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

The size (in bytes) of an HDF5 dataset or attribute as an [HDFQL\\_UNSIGNED\\_BIGINT](#).

## **Example(s)**

```
# create an HDF5 dataset named "my_dataset0" of data type float
CREATE DATASET my_dataset0 AS FLOAT

# show (i.e. get) size (in bytes) of dataset "my_dataset0" (should be 4)
SHOW SIZE my_dataset0

# create an HDF5 dataset named "my_dataset1" of data type long long of one dimension (size 3)
CREATE DATASET my_dataset1 AS BIGINT(3)

# show (i.e. get) size (in bytes) of dataset "my_dataset1" (should be 24 - i.e. 8x3)
SHOW SIZE my_dataset1

# create an HDF5 dataset named "my_common" of data type variable-length short with initial
values of 10, 20, 30, 40, 50 and 60
CREATE DATASET my_common AS VARSMALLINT VALUES(10, 20, 30, 40, 50, 60)

# create an HDF5 attribute named "my_common" of data type double of two dimensions (size 2x3)
CREATE ATTRIBUTE my_common AS DOUBLE(2, 3)

# show (i.e. get) size (in bytes) of dataset "my_common" (should be 12 - i.e. 2x6)
SHOW SIZE my_common

# show (i.e. get) size (in bytes) of dataset "my_common" (should be 12 - i.e. 2x6)
SHOW DATASET SIZE my_common

# show (i.e. get) size (in bytes) of attribute "my_common" (should be 48 - i.e. 8x2x3)
SHOW ATTRIBUTE SIZE my_common
```

## **6.7.22 SHOW HDFQL VERSION**

### **Syntax**

#### **SHOW HDFQL VERSION**

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

## **Description**

Show (i.e. get) version of HDFqI library. The format of the version returned is MAJOR.MINOR.REVISION.

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The version of HDFqI library as an [HDFQL\\_VARCHAR](#).

## **Example(s)**

```
# show (i.e. get) version of HDFqI library (should be something similar to "2.1.0")  
SHOW HDFQL VERSION
```

## **6.7.23 SHOW HDF5 VERSION**

### **Syntax**

**SHOW HDF5 VERSION**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### **Description**

Show (i.e. get) version of the HDF5 library used by HDFqI. The format of the version returned is MAJOR.MINOR.REVISION. The HDF5 library refers to the library used to compile HDFqI, and which is shipped with HDFqI (and not the HDF5 library that may be installed in the machine).



## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The version of the HDF5 library used by HDFql as an [HDFQL\\_VARCHAR](#).

## **Example(s)**

```
# show (i.e. get) version of the HDF5 library used by HDFql (should be something similar to  
"1.8.21")  
SHOW HDF5 VERSION
```

## **6.7.24 SHOW PCRE VERSION**

### **Syntax**

**SHOW PCRE VERSION**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### **Description**

Show (i.e. get) version of the PCRE library used by HDFql. The format of the version returned is MAJOR.MINOR. The PCRE library refers to the library used to compile HDFql, and which is shipped with HDFql (and not the PCRE library that may be installed in the machine).

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The version of the PCRE library used by HDFql as an [HDFQL\\_VARCHAR](#).

## **Example(s)**

```
# show (i.e. get) version of the PCRE library used by HDFql (should be something similar to  
"8.41")  
SHOW PCRE VERSION
```

## **6.7.25 SHOW ZLIB VERSION**

### **Syntax**

**SHOW ZLIB VERSION**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### **Description**

Show (i.e. get) version of the ZLIB library used by HDFql. The format of the version returned is MAJOR.MINOR.REVISION. The ZLIB library refers to the library used to compile HDFql, and which is shipped with HDFql (and not the ZLIB library that may be installed in the machine).

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The version of the ZLIB library used by HDFqL as an [HDFQL\\_VARCHAR](#).

## **Example(s)**

```
# show (i.e. get) version of the ZLIB library used by HDFqL (should be something similar to
"1.2.11")
SHOW ZLIB VERSION
```

## **6.7.26 SHOW MPI VERSION**

### **Syntax**

#### **SHOW MPI VERSION**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### **Description**

Show (i.e. get) version of the MPI library used by HDFqL. The information returned depends on the MPI library loaded by HDFqL at runtime (which must be previously installed in the machine). Please refer to [https://www.mpich.org/static/docs/v3.2/www3/MPI\\_Get\\_library\\_version.html](https://www.mpich.org/static/docs/v3.2/www3/MPI_Get_library_version.html) or [https://www.openmpi.org/doc/v2.1/man3/MPI\\_Get\\_library\\_version.3.php](https://www.openmpi.org/doc/v2.1/man3/MPI_Get_library_version.3.php) for additional information in case the MPI library used is MPICH (or, alternatively, one of its ABI compatible derivative libraries) or Open MPI.

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The version of the MPI library used by HDFqI as an [HDFQL\\_VARCHAR](#) or nothing (in case of using an HDFqI non MPI-based distribution or if in Windows as HDFqI does not support the parallel HDF5 (PHDF5) library in this platform currently).

## **Example(s)**

```
# show (i.e. get) version of the MPI library used by HDFqI (e.g. if the MPI library used is
Open MPI, it should be something similar to "Open MPI v2.1.3, package: Open MPI dummy@machine
Distribution, ident: 2.1.3, repo rev: v2.1.2-129-gcfd8f3f, Mar 13, 2018")
SHOW MPI VERSION
```

## **6.7.27 SHOW DIRECTORY**

### **Syntax**

**SHOW DIRECTORY** [*directory\_name*]

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### **Description**

Show (i.e. get) directory names within a directory named *directory\_name* or check the existence of a directory named *directory\_name*.

## **Parameter(s)**

*directory\_name* – optional string that specifies the name of the directory whose directory names are to be obtained or the name of the directory to check for its existence. If *directory\_name* is not specified, all directory names within the current working directory are returned. Otherwise, if *directory\_name* is specified, one of the following behaviors applies:

- If it ends with “/”, all directory names within *directory\_name* are returned.
- If it does not end with “/”, *directory\_name* will be checked for its existence as a directory. If it exists, *directory\_name* is returned; otherwise, if it does not exist, an error is raised.

*post\_processing\_option* – optional option that transforms the result set according to the programmer’s needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The directory names within a directory or the existence of a directory as an [HDFQL\\_VARCHAR](#).

## **Example(s)**

```
# create three directories named "my_directory0", "my_directory1" and "my_directory2" within
the current working directory
CREATE DIRECTORY my_directory0, my_directory1, my_directory2

# create two directories named "my_subdirectory0" and "my_subdirectory1" within the directory
"my_directory0"
CREATE DIRECTORY my_directory0/my_subdirectory0, my_directory0/my_subdirectory1

# show (i.e. get) directory names within the current working directory (should be
"my_directory0", "my_directory1" and "my_directory2")
SHOW DIRECTORY

# show (i.e. get) directory names within directory "my_directory0" (should be
"my_subdirectory0" and "my_subdirectory1")
SHOW DIRECTORY my_directory0/
```

```
# check the existence of a directory named "my_directory0" within the current working directory
(should be "my_directory0" - i.e. it exists)
SHOW DIRECTORY my_directory0
```

## 6.7.28 SHOW FILE

### Syntax

**SHOW FILE** [*object\_name*]

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) file names within a directory named *object\_name* or check the existence of a file named *object\_name*.

### Parameter(s)

*object\_name* – optional string that specifies the name of the directory whose file names are to be obtained or the name of the file to check for its existence. If *object\_name* is not specified, all file names within the current working directory are returned. Otherwise, if *object\_name* is specified, one of the following behaviors applies:

- If it ends with "/", *object\_name* will be treated as a directory and all file names within this directory are returned.
- If it does not end with "/", *object\_name* will be checked for its existence as a file. If it exists, *object\_name* is returned; otherwise, if it does not exist, an error is raised.

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

### Return

The file names within a directory or the existence of a file as an [HDFQL\\_VARCHAR](#).

## **Example(s)**

```
# create three HDF5 files named "my_file0.h5", "my_file1.h5" and "my_file2.h5" within the
current working directory
CREATE FILE my_file0.h5, my_file1.h5, my_file2.h5

# create two HDF5 files named "my_file3.h5" and "my_file4.h5" within a directory named
"my_directory"
CREATE FILE my_directory/my_file3.h5, my_directory/my_file4.h5

# show (i.e. get) file names within the current working directory (should be "my_file0.h5",
"my_file1.h5" and "my_file2.h5")
SHOW FILE

# show (i.e. get) file names within directory "my_directory" (should be "my_file3.h5" and
"my_file4.h5")
SHOW FILE my_directory/

# check the existence of a file named "my_file0.h5" within the current working directory
(should be "my_file0.h5" - i.e. it exists)
SHOW FILE my_file0.h5
```

## **6.7.29 SHOW EXECUTE STATUS**

### **Syntax**

#### **SHOW EXECUTE STATUS**

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### **Description**

Show (i.e. get) status of the last executed operation.

### **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The status of the last executed operation as an [HDFQL\\_INT](#).

## **Example(s)**

```
# show (i.e. get) current working directory (this operation will succeed since it is
syntactically correct)
SHOW USE DIRECTORY

# show (i.e. get) status of the last executed operation (should be 0 - i.e. HDFQL_SUCCESS)
SHOW EXECUTE STATUS

# show (i.e. get) current working directory (this operation will fail since it is syntactically
incorrect due to a typo in "SHOWX")
SHOWX USE DIRECTORY

# show (i.e. get) status of the last executed operation (should be -1 - i.e. HDFQL_ERROR_PARSE)
SHOW EXECUTE STATUS
```

## **6.7.30 SHOW LIBRARY BOUNDS**

### **Syntax**

**SHOW [USE FILE] LIBRARY BOUNDS [FROM | TO]**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### **Description**

Show (i.e. get) library bound values for creating or opening HDF5 files. If neither the keyword FROM nor TO is specified, all library bound values (i.e. from and to) are returned. To return a specific library bound value, the keyword FROM or TO must be specified. In case the keyword USE FILE is not specified, the library bound values returned refers to creating or opening files by default; otherwise, if it is specified, the library bound values of the file currently in use are returned.



## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer’s needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The library bound values for creating or opening HDF5 files as an [HDFQL\\_INT](#).

## **Example(s)**

```
# show (i.e. get) library bound values from (i.e. lower bound) and to (i.e. upper bound)
# (should be EARLIEST and LATEST - i.e. default values defined by the HDF5 library)
SHOW LIBRARY BOUNDS

# show (i.e. get) library bound value from (i.e. lower bound) (should be EARLIEST)
SHOW LIBRARY BOUNDS FROM

# show (i.e. get) library bound value to (i.e. upper bound) (should be LATEST)
SHOW LIBRARY BOUNDS TO

# set library bound from (i.e. lower bound) to LATEST (the library bound to - i.e. upper bound
- remains intact) for subsequent usage (i.e. creating or opening HDF5 files)
SET LIBRARY BOUNDS FROM LATEST

# show (i.e. get) library bound values from (i.e. lower bound) and to (i.e. upper bound)
# (should be LATEST and LATEST)
SHOW LIBRARY BOUNDS

# set both library bounds from (i.e. lower bound) and to (i.e. upper bound) to DEFAULT for
subsequent usage (i.e. creating or opening HDF5 files)
SET LIBRARY BOUNDS FROM DEFAULT TO DEFAULT

# show (i.e. get) library bound values from (i.e. lower bound) and to (i.e. upper bound)
# (should be EARLIEST and LATEST - i.e. default values defined by the HDF5 library)
SHOW LIBRARY BOUNDS
```

```
# use (i.e. open) an HDF5 file named "my_file0.h5" with library bounds from (i.e. lower bound)
and to (i.e. upper bound) set to EARLIEST and LATEST respectively
USE FILE my_file0.h5

# show (i.e. get) library bound value from (i.e. lower bound) (should be EARLIEST)
SHOW USE FILE LIBRARY BOUNDS FROM

# show (i.e. get) library bound value to (i.e. upper bound) (should be LATEST)
SHOW USE FILE LIBRARY BOUNDS TO

# set library bound from (i.e. lower bound) to LATEST (the library bound to - i.e. upper bound
- remains intact) for subsequent usage (i.e. creating or opening HDF5 files)
SET LIBRARY BOUNDS FROM LATEST

# use (i.e. open) an HDF5 file named "my_file1.h5" with both library bounds from (i.e. lower
bound) and to (i.e. upper bound) set to LATEST respectively
USE FILE my_file1.h5

# show (i.e. get) library bound values from (i.e. lower bound) and to (i.e. upper bound) of the
file currently in use (should be LATEST and LATEST)
SHOW USE FILE LIBRARY BOUNDS

# use (i.e. open) an HDF5 file named "my_file2.h5" with library bounds from (i.e. lower bound)
and to (i.e. upper bound) set to EARLIEST (i.e. default value defined by the HDF5 library) and
LATEST respectively
USE FILE my_file2.h5 LIBRARY BOUNDS FROM DEFAULT

# show (i.e. get) library bound values from (i.e. lower bound) and to (i.e. upper bound) of the
file currently in use (should be EARLIEST and LATEST)
SHOW USE FILE LIBRARY BOUNDS
```

## 6.7.31 SHOW CACHE

### Syntax

**SHOW** [[**USE**] **FILE** | **DATASET**] **CACHE** [**SLOTS** | **SIZE** | **PREEMPTION**]

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

## **Description**

Show (i.e. get) cache parameter values for accessing HDF5 files or datasets. If neither the keyword SLOTS, SIZE nor PREEMPTION is specified, all cache parameter values (i.e. for slots, size and preemption) are returned. To return a specific cache parameter value, the keyword SLOTS, SIZE or PREEMPTION must be specified. If neither the keyword FILE, USE FILE nor DATASET is specified, the cache parameters returned refers to files by default (optionally, the keyword FILE may be specified to make the purpose of this operation clearer). To explicitly return cache parameters of datasets or the file currently in use, the keyword DATASET or USE FILE must be specified.

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The cache parameter values for accessing HDF5 files or datasets as an [HDFQL\\_DOUBLE](#) (when returning all cache parameter values or the preemption cache parameter value only) or [HDFQL\\_INT](#) (when returning the slots or size cache parameter value).

## **Example(s)**

```
# show (i.e. get) cache parameter values for accessing HDF5 files (should be 521, 1048576, 0.75)
```

```
SHOW CACHE
```

```
# show (i.e. get) cache preemption value for accessing HDF5 files (should be 0.75)
```

```
SHOW CACHE PREEMPTION
```

```
# show (i.e. get) cache parameter values for accessing HDF5 files (should be 521, 1048576, 0.75)
```

```
SHOW FILE CACHE
```

```
# show (i.e. get) cache slots value for accessing HDF5 datasets (should be 521)
```

```
SHOW DATASET CACHE SLOTS
```

```
# use (i.e. open) an HDF5 file named "my_file0.h5" with cache parameters values previously set
(i.e. with slots, size and preemption values of 521, 1048576 and 0,75 respectively)
USE FILE my_file0.h5

# show (i.e. get) cache parameter values for accessing the HDF5 file currently in use (should
be 521, 1048576, 0,75)
SHOW USE FILE CACHE

# use (i.e. open) an HDF5 file named "my_file1.h5" with cache slots, size and preemption values
of 1523, 262144 and 0.6 respectively
USE FILE my_file1.h5 CACHE SLOTS 1523 SIZE 262144 PREEMPTION 0.6

# show (i.e. get) cache parameter values for accessing the HDF5 file currently in use (should
be 1523, 262144, 0.6)
SHOW USE FILE CACHE

# use (i.e. open) an HDF5 file named "my_file2.h5" with a cache preemption value of 0.9
USE FILE my_file2.h5 CACHE PREEMPTION 0.9

# show (i.e. get) cache parameter values for accessing the HDF5 file currently in use (should
be 521, 1048576, 0.9)
SHOW USE FILE CACHE
```

## 6.7.32 SHOW ATOMIC

### Syntax

**SHOW [USE FILE] ATOMIC**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### Description

Show (i.e. get) atomicity for accessing HDF5 files in an MPI environment. In case the keyword USE FILE is not specified, the atomicity returned refers to files that are subsequently opened; otherwise, if it is specified, the atomicity of the file currently in use is returned.

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The status of the atomicity for accessing HDF5 files as an [HDFQL\\_INT](#), which can either be [HDFQL\\_ENABLED](#) or [HDFQL\\_DISABLED](#) depending on whether the atomicity for accessing files is enabled or disabled respectively.

## **Example(s)**

```
# enable atomicity for accessing HDF5 files in an MPI environment
SET ATOMIC ENABLE

# show (i.e. get) atomicity for accessing HDF5 files in an MPI environment (should be 0 - i.e.
HDFQL_ENABLED)
SHOW ATOMIC

# use (i.e. open) an HDF5 file named "my_file0.h5" with atomicity for accessing it in an MPI
environment
USE PARALLEL FILE my_file0.h5

# show (i.e. get) atomicity of the HDF5 file currently in use in an MPI environment (should be
0 - i.e. HDFQL_ENABLED)
SHOW USE FILE ATOMIC

# disable atomicity for accessing HDF5 files in an MPI environment
SET ATOMIC DISABLE

# show (i.e. get) atomicity for accessing HDF5 files in an MPI environment (should be -1 - i.e.
HDFQL_DISABLED)
SHOW ATOMIC

# use (i.e. open) an HDF5 file named "my_file1.h5" without atomicity for accessing it in an MPI
environment
USE PARALLEL FILE my_file1.h5
```

```
# show (i.e. get) atomicity of the HDF5 file currently in use in an MPI environment (should be
-1 - i.e. HDFQL_DISABLED)
SHOW USE FILE ATOMIC

# use (i.e. open) an HDF5 file named "my_file2.h5" with atomicity for accessing it in an MPI
environment
USE ATOMIC PARALLEL FILE my_file2.h5

# show (i.e. get) atomicity of the HDF5 file currently in use in an MPI environment (should be
0 - i.e. HDFQL_ENABLED)
SHOW USE FILE ATOMIC
```

## 6.7.33 SHOW EXTERNAL LINK PREFIX

### Syntax

#### SHOW EXTERNAL LINK PREFIX

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]

### Description

Show (i.e. get) prefix to prepend to file names stored in HDF5 external links.

### Parameter(s)

*post\_processing\_option* – optional option that transforms the result set according to the programmer’s needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The prefix to prepend to file names stored in HDF5 external links as an [HDFQL\\_VARCHAR](#) or nothing (in case no prefix is specified).

## **Example(s)**

```
# set external link prefix to "/target"
SET EXTERNAL LINK PREFIX /target

# show (i.e. get) external link prefix (should be "/target")
SHOW EXTERNAL LINK PREFIX

# set external link prefix to default (i.e. empty)
SET EXTERNAL LINK PREFIX DEFAULT

# show (i.e. get) external link prefix (should be empty)
SHOW EXTERNAL LINK PREFIX
```

## **6.7.34 SHOW FLUSH**

### **Syntax**

**SHOW FLUSH**

*[post\_processing\_option [post\_processing\_option]\*]*

*[output\_redirecting\_option]*

### **Description**

Show (i.e. get) status of the automatic flushing.

### **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer's needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The status of the automatic flushing as an [HDFQL\\_INT](#), which can either be [HDFQL\\_GLOBAL](#), [HDFQL\\_LOCAL](#) or [HDFQL\\_DISABLED](#) depending on whether the automatic flushing of the entire virtual HDF5 file (global) or only the HDF5 file (local) currently in use is enabled or disabled respectively.

## **Example(s)**

```
# enable automatic flushing of the entire virtual HDF5 file (global) currently in use
SET FLUSH ENABLE

# show (i.e. get) status of the automatic flushing (should be 1 - i.e. HDFQL_GLOBAL)
SHOW FLUSH

# enable automatic flushing of only the HDF5 file (local) currently in use
SET FLUSH LOCAL ENABLE

# show (i.e. get) status of the automatic flushing (should be 2 - i.e. HDFQL_LOCAL)
SHOW FLUSH

# disable automatic flushing of the entire virtual HDF5 file (global) or only the HDF5 file
(local) currently in use
SET FLUSH DISABLE

# show (i.e. get) status of the automatic flushing (should be -1 - i.e. HDFQL_DISABLED)
SHOW FLUSH
```

## **6.7.35 SHOW DEBUG**

### **Syntax**

#### **SHOW DEBUG**

[*post\_processing\_option* [*post\_processing\_option*]\*]

[*output\_redirecting\_option*]



## **Description**

Show (i.e. get) status of the debug mechanism.

## **Parameter(s)**

*post\_processing\_option* – optional option that transforms the result set according to the programmer’s needs such as ordering or truncating (please refer to the section [POST-PROCESSING](#) for additional information). Multiple post processing options are separated with a space.

*output\_redirecting\_option* – optional option that specifies a (text or binary) file or memory (i.e. user-defined variable) to write the result set into. If not specified, the cursor in use is populated with the result set instead (please refer to the chapter [CURSOR](#) and subsection [INTO](#) for additional information).

## **Return**

The status of the debug mechanism as an [HDFQL\\_INT](#), which can either be [HDFQL\\_ENABLED](#) or [HDFQL\\_DISABLED](#) depending on whether the debug mechanism is enabled or disabled respectively.

## **Example(s)**

```
# enable debug mechanism (i.e. debug messages will be displayed when executing operations)
SET DEBUG ENABLE

# show (i.e. get) status of the debug mechanism (should be 0 - i.e. HDFQL_ENABLED)
SHOW DEBUG

# disable debug mechanism (i.e. debug messages will not be displayed when executing operations)
SET DEBUG DISABLE

# show (i.e. get) status of the debug mechanism (should be -1 - i.e. HDFQL_DISABLED)
SHOW DEBUG
```

# 6.8 MISCELLANEOUS

This section assembles all remaining HDFql operations that – due to their heterogeneous nature and functionality – do not fit in the previous sections about the language for data definition, manipulation, querying and introspection.

## 6.8.1 USE DIRECTORY

### Syntax

**USE DIRECTORY** *directory\_name*

### Description

Use a directory named *directory\_name* for subsequent operations. This will change the current working directory to *directory\_name* thus avoiding the need to explicitly provide the full path of this directory when working within it (i.e. subsequent operations are done relatively to this directory, unless otherwise specified). If *directory\_name* was not found or could not be opened (due to unknown/unexpected reasons), an error is raised.

### Parameter(s)

*directory\_name* – mandatory string that specifies the name of the directory to use for subsequent operations.

### Return

Nothing

### Example(s)

```
# set working directory currently in use to "/"
USE DIRECTORY /

# show (i.e. get) current working directory (should be "/")
SHOW USE DIRECTORY

# create a directory named "my_directory"
CREATE DIRECTORY my_directory

# set working directory currently in use to "my_directory" (more precisely "/my_directory")
USE DIRECTORY my_directory

# show (i.e. get) current working directory (should be "/my_directory")
SHOW USE DIRECTORY

# create two directories named "my_subdirectory0" and "my_subdirectory1" (both directories will
be created in directory "/my_directory")
CREATE DIRECTORY my_subdirectory0, my_subdirectory1
```

```

# set directory currently in use to "my_subdirectory0" (more precisely
"/my_directory/my_subdirectory0")
USE DIRECTORY my_subdirectory0

# show (i.e. get) current working directory (should be "/my_directory/my_subdirectory0")
SHOW USE DIRECTORY

# set directory currently in use to "my_subdirectory1" located one level up (more precisely
"/my_directory/my_subdirectory1")
USE DIRECTORY ../my_subdirectory1

# show (i.e. get) current working directory (should be "/my_directory/my_subdirectory1")
SHOW USE DIRECTORY

# set directory currently in use two levels up (should be "/")
USE DIRECTORY ../../

# show (i.e. get) current working directory (should be "/")
SHOW USE DIRECTORY

```

## 6.8.2 USE FILE

### Syntax

**USE** [**READONLY**] [[**ATOMIC**] **PARALLEL**] **FILE** *file\_name* [, *file\_name*]\*

[**LIBRARY BOUNDS** [**FROM** {**EARLIEST** | **LATEST** | **V18** | **DEFAULT**}] [**TO** {**LATEST** | **V18** | **DEFAULT**}}]

[**CACHE** [**SLOTS** {*slots\_value* | **DEFAULT**}] [**SIZE** {*size\_value* | **DEFAULT**}] [**PREEMPTION** {*preemption\_value* | **DEFAULT**}}]

### Description

Use (i.e. open) an HDF5 file named *file\_name* for subsequent operations. Multiple files can be opened at once by separating these with a comma (,). If *file\_name* was not found or could not be opened (due to unknown/unexpected reasons), no subsequent files are opened, and an error is raised. By default, the file is opened with read/write permissions. To open a file with read only permission, the keyword **READONLY** should be specified (any subsequent attempt to write into this file will raise an error). HDFqI tracks opened files in a stack fashion (i.e. LIFO) meaning that the most recently

opened file is the one currently in use. In case the keyword `PARALLEL`<sup>46</sup> is specified, HDFqI opens the file using all the MPI processes specified upon launching the program (that employs HDFqI). In case the keyword `ATOMIC` is specified, all file access operations will appear atomic, guaranteeing sequential consistency in an MPI environment (i.e. the operations will behave as though they were performed in a serial order consistent with the program order). In case the keyword `LIBRARY BOUNDS` is specified, HDFqI opens the file using these bounds (instead of the library bounds that may have been set through the operation `SET LIBRARY BOUNDS`). In case the keyword `CACHE` is specified, HDFqI opens the file using cache parametrized with the *slots\_value*, *size\_value* and *preemption\_value* values (instead of the file cache parameters that may have been set through the operation `SET CACHE`).

### **Parameter(s)**

*file\_name* – mandatory string that specifies the name of the HDF5 file to use (i.e. open) for subsequent operations. Multiple files are separated with a comma (,).

*slots\_value* – optional integer that specifies the number of chunk slots in the raw data chunk cache for accessing the HDF5 file. Due to the hashing strategy, its value should ideally be a prime number. In case the keyword `DEFAULT` is specified, its value is 521 (i.e. default value defined by the HDF5 library). In case the keyword `SLOTS` is not specified, its current value remains intact.

*size\_value* – optional integer that specifies the total size of the raw data chunk cache in bytes for accessing the HDF5 file. In case the keyword `DEFAULT` is specified, its value is 1048576 (i.e. 1 MB – default value defined by the HDF5 library). In case the keyword `SIZE` is not specified, its current value remains intact.

*preemption\_value* – optional float that specifies the chunk preemption policy for accessing the HDF5 file. Its value must be between 0 and 1. It indicates the weighting according to which chunks which have been fully read or written are penalized when determining which chunks to flush from cache. In case the keyword `DEFAULT` is specified, its value is 0.75 (i.e. default value defined by the HDF5 library). In case the keyword `PREEMPTION` is not specified, its current value remains intact.

### **Return**

Nothing

---

<sup>46</sup> This option is not allowed in Windows as HDFqI does not support the parallel HDF5 (PHDF5) library in this platform currently.

## **Example(s)**

```
# use (i.e. open) an HDF5 file named "my_file0.h5" located in the current working directory
USE FILE my_file0.h5

# use (i.e. open) an HDF5 file named "my_file1.h5" located in a root directory named "data"
USE FILE /data/my_file1.h5

# use (i.e. open) two HDF5 files named "my_file2.h5" and "my_file3.h5" with read only
permissions (both files are located in the current working directory)
USE READONLY FILE my_file2.h5, my_file3.h5

# use (i.e. open) an HDF5 file named "my_file4.h5" located in the parent directory with the
latest version of the HDF5 library
USE FILE ../my_file4.h5 LIBRARY BOUNDS FROM LATEST TO LATEST

# use (i.e. open) an HDF5 file named "my_file5.h5" located in the current working directory
with cache slots, size and preemption values of 1523, 262144 and 0.6 respectively
USE FILE my_file5.h5 CACHE SLOTS 1523 SIZE 262144 PREEMPTION 0.6

# use (i.e. open) an HDF5 file named "my_file6.h5" located in the current working directory
with the earliest version of the HDF5 library and a cache preemption value of 0.9
USE FILE my_file6.h5 LIBRARY BOUNDS FROM EARLIEST CACHE PREEMPTION 0.9

# use (i.e. open) an HDF5 file named "my_file7.h5" located in the current working directory in
parallel (i.e. all the MPI processes specified upon launching the program (that employs HDFql)
will collectively open the file - e.g. if the program is launched as "mpiexec -n 4 my_program",
all the four MPI processes will participate in the opening of the file)
USE PARALLEL FILE my_file7.h5

# use (i.e. open) an HDF5 file named "my_file8.h5" located in the current working directory in
parallel with atomicity for accessing it
USE ATOMIC PARALLEL FILE my_file8.h5
```

## **6.8.3 USE GROUP**

### **Syntax**

**USE GROUP** *group\_name*

## **Description**

Use (i.e. open) an HDF5 group named *group\_name* for subsequent operations. This will change the current working group to *group\_name* thus avoiding the need to explicitly provide the full path of this group when working within it (i.e. subsequent operations are done relatively to this group, unless otherwise specified). If *group\_name* was not found or could not be opened (due to unknown/unexpected reasons), an error is raised. Upon using (i.e. opening) an HDF5 file, the group currently in use is `"/` (i.e. the root of the HDF5 file).

## **Parameter(s)**

*group\_name* – mandatory string that specifies the name of the HDF5 group to use (i.e. open) for subsequent operations. Besides the name of the group to be used for subsequent operations, *group\_name* may be composed of special tokens (that are not part of the name of the group itself). These are:

- `"/` to separate multiple groups. Example: `"USE GROUP my_group/my_subgroup/my_subsubgroup"`.
- `."` to refer to the group currently in use. Example: `"USE GROUP ."`.
- `.."` to go up one level from the group currently in use. Example: `"USE GROUP .."`.

## **Return**

Nothing

## **Example(s)**

```
# set group currently in use to "/" (i.e. the root of the HDF5 file)
USE GROUP /

# create two HDF5 groups named "my_group0" and "my_group1" (both groups will be created in
group "/")
CREATE GROUP my_group0, my_group1

# create an HDF5 dataset named "my_dataset0" of data type double (it will be created in group
"/")
CREATE DATASET my_dataset0 AS DOUBLE

# set group currently in use to "my_group0" (more precisely "/"my_group0")
USE GROUP my_group0

# create an HDF5 dataset named "my_dataset1" of data type double (it will be created in group
```

```
"/my_group0")
CREATE DATASET my_dataset1 AS DOUBLE

# create an HDF5 group named "my_subgroup0" (it will be created in group "/my_group0")
CREATE GROUP my_subgroup0

# create an HDF5 dataset named "my_dataset2" of data type variable-length double (it will be
created in group "/my_group0/my_subgroup0")
CREATE DATASET my_subgroup0/my_dataset2 AS VARDOUBLE

# create an HDF5 attribute named "my_attribute0" of data type float (it will be created in
group "/")
CREATE ATTRIBUTE ../my_attribute0 AS FLOAT

# set group currently in use to "my_subgroup0" (more precisely "/my_group0/my_subgroup0")
USE GROUP my_subgroup0

# create an HDF5 attribute named "my_attribute1" of data type char (it will be created in group
"/my_group1")
CREATE ATTRIBUTE ../../my_group1/my_attribute1 AS CHAR

# create an HDF5 attribute named "my_attribute2" of data type variable-length char (it will be
created in group "/")
CREATE ATTRIBUTE /my_attribute2 AS VARCHAR

# set group currently in use to "." (the group currently in use will not change - i.e. it
remains "/my_group0/my_subgroup0" - as "." refers to the current working group itself)
USE GROUP .

# create an HDF5 attribute named "my_attribute3" of data type int (it will be created in group
"/my_group0/my_subgroup0")
CREATE ATTRIBUTE my_attribute3 AS INT

# set group currently in use one level up (should be "/my_group0")
USE GROUP ..

# create an HDF5 attribute named "my_attribute4" of data type short (it will be created in
group "/my_group0")
CREATE ATTRIBUTE my_attribute4 AS SMALLINT
```

## 6.8.4 FLUSH

### Syntax

**FLUSH** [GLOBAL | LOCAL]

### Description

Flush the entire virtual HDF5 file (global) or the specific HDF5 file (local) currently in use. All buffered data will be written into the disk. If neither the keyword GLOBAL nor LOCAL is specified, a global flush is performed by default (optionally, the keyword GLOBAL may be specified to make the purpose of this operation clearer). To perform a local flush, the keyword LOCAL must be specified. If no file is currently used, no flush is performed.

### Parameter(s)

None

### Return

Nothing

### Example(s)

```
# flush the entire virtual HDF5 file (global) currently in use
FLUSH

# flush the entire virtual HDF5 file (global) currently in use
FLUSH GLOBAL

# flush only the HDF5 file (local) currently in use
FLUSH LOCAL
```

## 6.8.5 CLOSE FILE

### Syntax

**CLOSE FILE** [*file\_name* [, *file\_name*]\*]



## **Description**

Close a certain HDF5 file used (i.e. opened) or the HDF5 file currently in use. Multiple files can be closed at once by separating these with a comma (.). If *file\_name* is not in use (i.e. open) or it is not possible to close it (due to unknown/unexpected reasons), no subsequent files are closed, and an error is raised. Before closing a file, all buffered data will be written into it. After closing a file, the file in use will be the one most recently used before the closed file. If *file\_name* is specified, it will be closed regardless of whether it is the file currently in use or not. Otherwise, if *file\_name* is not specified, the file currently in use will be closed.

## **Parameter(s)**

*file\_name* – optional string that specifies the name of the HDF5 file to close. Multiple files are separated with a comma (.). If *file\_name* is specified it must match exactly the name of the file when it was opened (otherwise no file will be closed and an error is raised).

## **Return**

Nothing

## **Example(s)**

```
# use (i.e. open) four HDF5 files named "my_file0.h5", "my_file1.h5", "my_file2.h5" and
"my_file3.h5"
USE FILE my_file0.h5, my_file1.h5, my_file2.h5, my_file3.h5

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be "my_file3.h5")
SHOW USE FILE

# close HDF5 file currently in use (i.e. file "my_file3.h5")
CLOSE FILE

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be "my_file2.h5")
SHOW USE FILE

# close HDF5 file "my_file1.h5"
CLOSE FILE my_file1.h5

# show (i.e. get) HDF5 file currently in use (i.e. open) (should be "my_file2.h5")
SHOW USE FILE

# close HDF5 file currently in use (i.e. file "my_file2.h5")
```

***CLOSE FILE***

```
# show (i.e. get) HDF5 file currently in use (i.e. open) (should be "my_file0.h5")
```

***SHOW USE FILE***

```
# close HDF5 file currently in use (i.e. file "my_file0.h5")
```

***CLOSE FILE***

```
# show (i.e. get) HDF5 file currently in use (i.e. open) (should be empty)
```

***SHOW USE FILE***

## 6.8.6 CLOSE ALL FILE

### Syntax

CLOSE ALL FILE

### Description

Close all HDF5 files in use. All buffered data will be written into the respective files before closing them. If it is not possible to close a file (due to unknown/unexpected reasons), no subsequent files are closed, and an error is raised.

### Parameter(s)

None

### Return

Nothing

### Example(s)

```
# use (i.e. open) three HDF5 files named "my_file0.h5", "my_file1.h5" and "my_file2.h5"
```

```
USE FILE my_file0.h5, my_file1.h5, my_file2.h5
```

```
# show (i.e. get) all HDF5 files in use (i.e. open) (should be "my_file2.h5", "my_file1.h5",  
"my_file0.h5")
```

```
SHOW ALL USE FILE
```

```
# close all HDF5 files in use (i.e. open)
```

```
CLOSE ALL FILE
```

```
# show (i.e. get) all HDF5 files in use (i.e. open) (should be empty)
```

```
SHOW ALL USE FILE
```

## 6.8.7 CLOSE GROUP

### Syntax

**CLOSE GROUP**

### Description

Close the HDF5 group currently in use. After closing it, the group currently in use will be "/" (i.e. the root of the HDF5 file). If no file is currently used, no group is closed.

### Parameter(s)

None

### Return

Nothing

### Example(s)

```
# use (i.e. open) an HDF5 file named "my_file.h5"
```

```
USE FILE my_file.h5
```

```
# show (i.e. get) current working group (should be "/")
```

```
SHOW USE GROUP
```

```
# create an HDF5 group named "my_group"
```

```
CREATE GROUP my_group
```

```
# set group currently in use to "my_group" (more precisely "/my_group")
```

```
USE GROUP my_group
```

```
# show (i.e. get) current working group (should be "/my_group")
```

```
SHOW USE GROUP
```

```
# create an HDF5 dataset named "my_dataset" of data type double (more precisely
"/my_group/my_dataset")
CREATE DATASET my_dataset AS DOUBLE

# set group currently in use to "/" (i.e. the root of the HDF5 file)
CLOSE GROUP

# show (i.e. get) current working group (should be "/")
SHOW USE GROUP

# create an HDF5 dataset named "my_dataset" of data type int (more precisely "/" my_dataset")
CREATE DATASET my_dataset AS INT
```

## 6.8.8 SET LIBRARY BOUNDS

### Syntax

**SET LIBRARY BOUNDS [FROM {EARLIEST | LATEST | V18 | DEFAULT}] [TO {LATEST | V18 | DEFAULT}]**

### Description

Set library bounds<sup>47</sup> for creating and opening HDF5 files. In other words, it sets bounds on library versions to be used when creating objects (the object format versions are determined indirectly from the HDF5 library versions specified in the call). All files that are subsequently created or opened (through the operations [CREATE FILE](#) or [USE FILE](#)) will use the default bound values defined by the HDF5 library or user-defined bound values. These bounds are:

- From – lower bound on the range of possible library versions used to create the object. The library version indirectly specifies the earliest object format version that can be used when creating objects in an HDF5 file. In case the keyword DEFAULT is specified, its value is EARLIEST (i.e. default value defined by the HDF5 library). In case the keyword FROM is not specified (i.e. the lower bound), its current value remains intact.
- To – upper bound on the range of possible library versions used to create the object. The library version indirectly specifies the latest object format version that can be used when creating objects in an HDF5 file. In case the keyword

---

<sup>47</sup> By default (i.e. upon initialization of the HDFql library), the library bounds from (i.e. lower bound) and to (i.e. upper bound) are set to EARLIEST and LATEST respectively.

DEFAULT is specified, its value is LATEST (i.e. default value defined by the HDF5 library). In case the keyword TO is not specified (i.e. the upper bound), its current value remains intact.

### **Parameter(s)**

None

### **Return**

Nothing

### **Example(s)**

```
# use (i.e. open) an HDF5 file named "my_file0.h5" with library bounds from (i.e. lower bound)
and to (i.e. upper bound) set to EARLIEST and LATEST respectively (default values defined by
the HDF5 library)
```

```
USE FILE my_file0.h5
```

```
# set library bound from (i.e. lower bound) to LATEST (the library bound to - i.e. upper bound
- remains intact) for subsequent usage (i.e. creating or opening HDF5 files)
```

```
SET LIBRARY BOUNDS FROM LATEST
```

```
# use (i.e. open) an HDF5 file named "my_file1.h5" with both library bounds from (i.e. lower
bound) and to (i.e. upper bound) set to LATEST
```

```
USE FILE my_file1.h5
```

```
# set library bound to (i.e. upper bound) to V18 (the library bound from - i.e. lower bound -
remains intact) for subsequent usage (i.e. creating or opening HDF5 files)
```

```
SET LIBRARY BOUNDS TO V18
```

```
# use (i.e. open) an HDF5 file named "my_file2.h5" with library bounds from (i.e. lower bound)
and to (i.e. upper bound) set to LATEST and V18 respectively
```

```
USE FILE my_file2.h5
```

```
# set both library bounds from (i.e. lower bound) and to (i.e. upper bound) to DEFAULT for
subsequent usage (i.e. creating or opening HDF5 files)
```

```
SET LIBRARY BOUNDS FROM DEFAULT TO DEFAULT
```

```
# use (i.e. open) an HDF5 file named "my_file3.h5" with library bounds from (i.e. lower bound)
and to (i.e. upper bound) set to EARLIEST and LATEST respectively (default values defined by
the HDF5 library)
```

```
USE FILE my_file3.h5
```

## 6.8.9 SET CACHE

### Syntax

```
SET [FILE | DATASET] CACHE [SLOTS {slots_value | FILE | DEFAULT}] [SIZE {size_value | FILE | DEFAULT}]
[PREEMPTION {preemption_value | FILE | DEFAULT}]
```

### Description

Set cache parameters<sup>48</sup> for accessing HDF5 files or datasets. All files or datasets that are subsequently opened or read (through the operations [USE FILE](#) or [SELECT](#) respectively) will use the default cache parameter values defined by the HDF5 library or user-defined cache parameter values. These cache parameters are:

- Slots – number of chunk slots in the raw data chunk cache.
- Size – total size of the raw data chunk cache in bytes.
- Preemption – chunk preemption policy.

If neither the keyword FILE nor DATASET is specified, the setting of the cache parameters refers to files by default (optionally, the keyword FILE may be specified to make the purpose of this operation clearer). To explicitly set the cache parameters to datasets, the keyword DATASET must be specified.

### Parameter(s)

*slots\_value* – number of chunk slots in the raw data chunk cache for accessing HDF5 files or datasets. Due to the hashing strategy, its value should ideally be a prime number. In case the keyword FILE is specified, its value will be as defined in the cache slots parameter upon using (i.e. opening) the file. In case the keyword DEFAULT is specified, its value is 521 (i.e. default value defined by the HDF5 library). In case the keyword SLOTS is not specified, its current value remains intact.

*size\_value* – total size of the raw data chunk cache in bytes for accessing HDF5 files or datasets. In case the keyword FILE is specified, its value will be as defined in the cache size parameter upon using (i.e. opening) the file. In case the keyword DEFAULT is specified, its value is 1048576 (i.e. 1 MB – default value defined by the HDF5 library). In case the keyword SIZE is not specified, its current value remains intact.

*preemption\_value* – chunk preemption policy for accessing HDF5 files or datasets. Its value must be between 0 and 1. It indicates the weighting according to which chunks which have been fully read or written are penalized when determining

---

<sup>48</sup> By default (i.e. upon initialization of the HDFqI library), the cache parameters slots, size and preemption are set to 521, 1048576 and 0.75 respectively.

which chunks to flush from cache. In case the keyword FILE is specified, its value will be as defined in the cache preemption parameter upon using (i.e. opening) the file. In case the keyword DEFAULT is specified, its value is 0.75 (i.e. default value defined by the HDF5 library). In case the keyword PREEMPTION is not specified, its current value remains intact.

## **Return**

Nothing

## **Example(s)**

```
# use (i.e. open) an HDF5 file named "my_file0.h5" with cache slots, size and preemption values
of 521, 1048576 and 0.75 respectively (default values defined by the HDF5 library)
USE FILE my_file0.h5

# set cache slots and preemption values to 2297 and 0.9 respectively (the cache size value
remains intact) for subsequent usage (i.e. opening HDF5 files)
SET CACHE SLOTS 2297 PREEMPTION 0.9

# use (i.e. open) an HDF5 file named "my_file1.h5" with cache slots, size and preemption values
of 2297, 1048576 and 0.9 respectively
USE FILE my_file1.h5

# set cache slots, size and preemption values to 1523, 262144 and 0.6 respectively for
subsequent usage (i.e. opening HDF5 files)
SET FILE CACHE SLOTS 1523 SIZE 262144 PREEMPTION 0.6

# use (i.e. open) an HDF5 file named "my_file2.h5" with cache slots, size and preemption values
of 1523, 262144 and 0.6 respectively
USE FILE my_file2.h5

# set cache size value to 1048576 (default value defined by the HDF5 library) and preemption
value to 0.4 (the cache slots value remains intact) for subsequent usage (i.e. opening HDF5
files)
SET FILE CACHE SIZE DEFAULT PREEMPTION 0.4

# use (i.e. open) an HDF5 file named "my_file3.h5" with cache slots, size and preemption values
of 1523, 1048576 and 0.4 respectively
USE FILE my_file3.h5
```

```
# select (i.e. read) an HDF5 dataset named "my_dataset0" with cache slots, size and preemption
values of 521, 1048576 and 0.75 respectively (default values defined by the HDF5 library)
SELECT FROM my_dataset0

# set cache slots and preemption values to 2297 and 0.9 respectively (the cache size value
remains intact) for subsequent selection (i.e. reading HDF5 datasets)
SET DATASET CACHE SLOTS 2297 PREEMPTION 0.9

# select (i.e. read) an HDF5 dataset named "my_dataset1" with cache slots, size and preemption
values of 2297, 1048576 and 0.9 respectively
SELECT FROM my_dataset1

# set cache slots, size and preemption values to 1523, 262144 and 0.6 respectively for
subsequent selection (i.e. reading HDF5 datasets)
SET DATASET CACHE SLOTS 1523 SIZE 262144 PREEMPTION 0.6

# select (i.e. read) an HDF5 dataset named "my_dataset2" with cache slots, size and preemption
values of 1523, 262144 and 0.6 respectively
SELECT FROM my_dataset2

# set cache size value to 1048576 (default value defined by the HDF5 library) and preemption
value to 0.4 (the cache slots value remains intact) for subsequent selection (i.e. reading HDF5
datasets)
SET DATASET CACHE SIZE DEFAULT PREEMPTION 0.4

# select (i.e. read) an HDF5 dataset named "my_dataset3" with cache slots, size and preemption
values of 1523, 1048576 and 0.4 respectively
SELECT FROM my_dataset3

# set cache slots, size and preemption values to 3089, 2048 and 0.85 respectively for
subsequent usage (i.e. opening HDF5 files)
SET FILE CACHE SLOTS 3089 SIZE 2048 PREEMPTION 0.85

# set cache slots value to 521 (default value defined by the HDF5 library), size value to 1024,
and preemption value to 0.85 (defined by the cache preemption value for HDF5 files) for
subsequent selection (i.e. reading HDF5 datasets)
SET DATASET CACHE SLOTS DEFAULT SIZE 1024 PREEMPTION FILE

# select (i.e. read) an HDF5 dataset named "my_dataset4" with cache slots, size and preemption
values of 521, 1024 and 0.85 respectively
SELECT FROM my_dataset4
```



## 6.8.10 SET ATOMIC

### Syntax

**SET [USE FILE] ATOMIC {ENABLE | DISABLE | DEFAULT}**

### Description

Set atomicity<sup>49</sup> for accessing HDF5 files in an MPI environment to enabled or disabled. All files that are subsequently opened (through the operation [USE FILE](#)) will have access operations performed in an atomic fashion or not accordingly. If enabled, all file access operations will appear atomic, guaranteeing sequential consistency in an MPI environment (i.e. the operations will behave as though they were performed in a serial order consistent with the program order). If disabled, no enforcement of atomic file access will be done. In case the keyword DEFAULT is specified, the atomicity for accessing files is set to disabled (i.e. equivalent to specifying the keyword DISABLE). In case the keyword USE FILE is specified, subsequent access operations of the file currently in use will be performed in an atomic fashion or not accordingly.

### Parameter(s)

None

### Return

Nothing

### Example(s)

```
# enable atomicity for accessing HDF5 files in an MPI environment
SET ATOMIC ENABLE

# use (i.e. open) an HDF5 file named "my_file0.h5" with atomicity for accessing it in an MPI
environment
USE PARALLEL FILE my_file0.h5

# disable atomicity for accessing HDF5 files in an MPI environment
SET ATOMIC DISABLE

# use (i.e. open) an HDF5 file named "my_file1.h5" without atomicity for accessing it in an MPI
environment
```

---

<sup>49</sup> By default (i.e. upon initialization of the HDFqI library), the atomicity for accessing HDF5 files is set to disabled.

```
USE PARALLEL FILE my_file1.h5

# enable atomicity for accessing the HDF5 file currently in use (i.e. file "my_file1.h5") in an
MPI environment
SET USE FILE ATOMIC

# use (i.e. open) an HDF5 file named "my_file2.h5" with atomicity for accessing it in an MPI
environment
USE ATOMIC PARALLEL FILE my_file2.h5
```

## 6.8.11 SET EXTERNAL LINK PREFIX

### Syntax

SET EXTERNAL LINK PREFIX {*prefix\_value* | DEFAULT}

### Description

Set prefix<sup>50</sup> to prepend to file names stored in HDF5 external links. In other words, before resolving a file name stored in an external link, the prefix *prefix\_value* is prepended to the name. In case the keyword DEFAULT is specified, the prefix to resolve file names is set to empty (i.e. nothing is prepended).

### Parameter(s)

*prefix\_value* – optional string that specifies the prefix to prepend to file names stored in HDF5 external links.

### Return

Nothing

### Example(s)

```
# set directory currently in use to "/data"
USE DIRECTORY /data

# create two HDF5 files named "my_file0.h5" and "my_file1.h5" in the directory currently in use
(i.e. directory "/data") and in a root directory named "target" respectively
```

---

<sup>50</sup> By default (i.e. upon initialization of the HDFqI library), the prefix is set to empty (i.e. nothing is prepended).

```

CREATE FILE my_file0.h5, /target/my_file1.h5

# create an HDF5 dataset named "my_dataset" (in file "my_file1.h5" located in root directory
"target") of data type float with an initial value of 7.8
CREATE DATASET /target/my_file1.h5 my_dataset AS FLOAT VALUES (7.8)

# use (i.e. open) HDF5 file "my_file0.h5" located in the directory currently in use (i.e.
directory "/data")
USE FILE my_file0.h5

# create an HDF5 external link named "my_link" in file "my_file0.h5" to object "my_dataset" in
file "my_file1.h5"
CREATE EXTERNAL LINK my_link TO my_file1.h5 my_dataset

# select (i.e. read) data from object "my_link" and populate cursor in use with it (should
raise an error since "my_link" is a dangling link due to "my_file1.h5" being located in
directory "/target" and not in the directory currently in use)
SELECT FROM my_link

# set external link prefix to "/target"
SET EXTERNAL LINK PREFIX /target

# select (i.e. read) data from object "my_link" and populate cursor in use with it (should be
7.8)
SELECT FROM my_link

```

## 6.8.12 SET FLUSH

### Syntax

**SET FLUSH** {[GLOBAL | LOCAL] **ENABLE**} | **DISABLE** | **DEFAULT**}

### Description

Set automatic flushing<sup>51</sup> of the entire virtual HDF5 file (global) or only the HDF5 file (local) currently in use to enabled or disabled. If enabled, automatic flushing (i.e. all buffered data is written into the disk) will subsequently occur whenever an operation modifying the file is executed. If neither the keyword GLOBAL nor LOCAL is specified, automatic global flushing is set by default (optionally, the keyword GLOBAL may be specified to make the purpose of this operation clearer). To set

---

<sup>51</sup> By default (i.e. upon initialization of the HDFqI library), the automatic flushing is set to disabled.

automatic local flushing, the keyword LOCAL must be specified. In case the keyword DEFAULT is specified, the automatic flushing is set to disabled (i.e. equivalent to specifying the keyword DISABLE).

### **Parameter(s)**

None

### **Return**

Nothing

### **Example(s)**

```
# enable automatic flushing of the entire virtual HDF5 file (global) currently in use
SET FLUSH ENABLE

# enable automatic flushing of the entire virtual HDF5 file (global) currently in use
SET FLUSH GLOBAL ENABLE

# enable automatic flushing of only the HDF5 file (local) currently in use
SET FLUSH LOCAL ENABLE

# disable automatic flushing of the entire virtual HDF5 file (global) or only the HDF5 file
(local) currently in use
SET FLUSH DISABLE
```

## **6.8.13 SET DEBUG**

### **Syntax**

**SET DEBUG {ENABLE | DISABLE | DEFAULT}**

### **Description**

Set debug mechanism<sup>52</sup> to enabled or disabled. If enabled, debug messages will be displayed when executing operations, which should help the programmer to have a better understanding of the parameters HDFqI is receiving, the operation performed, and the return value of this operation. Additionally, debug messages of the HDF5 library itself are displayed in

---

<sup>52</sup> By default (i.e. upon initialization of the HDFqI library), the debug mechanism is set to disabled.

case of an error. In case the keyword DEFAULT is specified, the debug mechanism is set to disabled (i.e. equivalent to specifying the keyword DISABLE).

### **Parameter(s)**

None

### **Return**

Nothing

### **Example(s)**

```
# enable debug mechanism (i.e. debug messages will be displayed when executing operations)
```

```
SET DEBUG ENABLE
```

```
# disable debug mechanism (i.e. debug messages will not be displayed when executing operations)
```

```
SET DEBUG DISABLE
```

---

# GLOSSARY

---

## Application Programming Interface (API)

An Application Programming Interface (API) specifies how software components should interact with each other. In practice, an API comes in the form of a library that includes specifications for functions, data structures, object classes, constants and variables. A good API makes it easier to develop a program by providing all the building blocks.

## Attribute

An (HDF5) attribute is a metadata object describing the nature and/or intended usage of a primary data object. A primary data object may be a group, dataset or committed data type. Attributes are assumed to be very small as data objects go, so storing them as standard (HDF5) datasets would be inefficient.

## Cursor

A cursor is a control structure that is used to iterate through the results returned by a query (that was previously executed). It can be seen as an effective means to abstract the programmer from low-level implementation details of accessing data stored in specific structures. Besides offering several ways to traverse result sets according to specific needs, cursors in HDFqI also store result sets returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations.

## Dataset

A (HDF5) dataset is an object composed of a collection of data elements and metadata that stores a description of the data elements, data layout and all other information necessary to write and read the data. A dataset may be a multidimensional array of data elements and it may have zero or more attributes.

## Data type

A data type is a classification identifying one of various types of data such as integer, real or string, which determines the possible values for that type, the operations that can be done on values of that type, the meaning of the data, and the way values of that type can be stored. In other words, a data type is a classification of data that tells HDFqI how the user intends to use it.

## Endianness

Endianness refers to the ordering of packing bytes into words when stored in memory. In big endian format, whenever addressing memory or storing words byte-wise, the most significant byte – i.e. the byte containing the most significant bit – is stored first (has the lowest address); subsequently, the following bytes are stored in order of decreasing significance, with the least significant byte – i.e. the one containing the least significant bit – stored last (having the highest address). The little endian format reverses this order: the sequence addresses/stores the least significant byte first (lowest address) and the most significant byte last (highest address).

## Group

A (HDF5) group is a container structure which can hold zero or more objects (i.e. datasets, (soft) links, external links and other groups) and have zero or more attributes (attached to it). Every object must be a member of at least one group, except the root group, which (as the sole exception) may not belong to any group.

## Hierarchical Data Format (HDF)

The Hierarchical Data Format (HDF) is the name of a set of file formats and libraries designed to store and organize large amounts of numerical data. It is a versatile data model that can represent complex data objects and a wide variety of metadata. HDF is supported by The HDF Group, whose mission is to ensure continued development of HDF technologies and the continued accessibility of data currently stored in this file format.

## Hyperslab

A hyperslab allows reading or writing a portion (subset) of a dataset (as opposed to its entirety). It can be a selection of logically contiguous collection of points in a dataspace, or it can be a regular pattern of points or blocks in a dataspace.

## Message Passing Interface (MPI)

The Message Passing Interface (MPI) is a standardized means of exchanging messages between multiple computers running a parallel program across distributed memory. It was designed by a group of researchers from academia and industry to work on a wide variety of parallel computing architectures. MPI fosters the development of a parallel software industry, and encourages development of portable and scalable large-scale parallel applications.

## Operation

An operation is an action that can be performed in HDFql such as to create an HDF5 file or read data from a dataset. Operations can be seen as the HDFql language itself. In HDFql, many operations exist and these are categorized into [DATA DEFINITION LANGUAGE \(DDL\)](#), [DATA MANIPULATION LANGUAGE \(DML\)](#), [DATA QUERY LANGUAGE \(DQL\)](#), [DATA INTROSPECTION LANGUAGE \(DIL\)](#) and [MISCELLANEOUS](#).

## Parallel HDF5 (PHDF5)

The Parallel HDF5 (PHDF5) is a parallel version of HDF5 which is the name of a set of file formats and libraries designed to store and organize large amounts of numerical data. It leverages MPI to effectively manipulate HDF5 files in parallel (e.g. write a dataset) across multiple computers. In HDFql, PHDF5 can be explicitly used through the [CREATE FILE](#), [USE FILE](#), [INSERT](#) and [SELECT](#) operations.

## Post-processing

Post-processing options enable transforming results of a query according to the programmer's needs such as ordering or truncating. These options are optional and may be used to create a (linear) pipeline to further process result sets returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations.

## Result set

A result set stores the results (of data type [HDFQL\\_TINYINT](#), [HDFQL\\_UNSIGNED\\_TINYINT](#), [HDFQL\\_SMALLINT](#), [HDFQL\\_UNSIGNED\\_SMALLINT](#), [HDFQL\\_INT](#), [HDFQL\\_UNSIGNED\\_INT](#), [HDFQL\\_BIGINT](#), [HDFQL\\_UNSIGNED\\_BIGINT](#),



[HDFQL\\_FLOAT](#), [HDFQL\\_DOUBLE](#) and [HDFQL\\_VARCHAR](#)) returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations.

## Result subset

A result subset stores the results (of data type [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) and [HDFQL\\_OPAQUE](#)) returned by [DATA QUERY LANGUAGE \(DQL\)](#) and [DATA INTROSPECTION LANGUAGE \(DIL\)](#) operations.

## Subcursor

A subcursor is meant to complement (i.e. help) cursors in the task of storing data of type [HDFQL\\_CHAR](#), [HDFQL\\_VARTINYINT](#), [HDFQL\\_UNSIGNED\\_VARTINYINT](#), [HDFQL\\_VARSMALLINT](#), [HDFQL\\_UNSIGNED\\_VARSMALLINT](#), [HDFQL\\_VARINT](#), [HDFQL\\_UNSIGNED\\_VARINT](#), [HDFQL\\_VARBIGINT](#), [HDFQL\\_UNSIGNED\\_VARBIGINT](#), [HDFQL\\_VARFLOAT](#), [HDFQL\\_VARDOUBLE](#) and [HDFQL\\_OPAQUE](#). In practice, when a result set is of one of these data types, only the first element of the result set is stored in the cursor (as expected), while all elements of the result set are stored in the subcursor. In other words, each position of the cursor stores the first element of the result set and also points to a subcursor that in turn stores all the elements of the result set. Similar to cursors, HDFql subcursors offer several ways to traverse result subsets.