



HDFq1

the easy way to manage HDF5 data

This presentation was prepared by Rick (Mr. HDFq1) who can't be here today and is presented by Gerd Heber (The HDF Group)

AGENDA



- What it is
- Who is using it
- How it works & examples
- Current features
- What's next!



WHAT IT IS

- HDFq1 stands for “Hierarchical Data Format query language” and is a high-level language to manage HDF5 data: *declarative*
- Designed to be as simple and powerful as SQL: adapted to HDF5
- Dramatically reduces the learning effort and time needed to handle HDF5: more time to focus on the problem at hand
- Reads and writes HDF5 data across multiple programming languages and platforms: a single *guest language* for multiple *host languages*
- Pre/post processes (HDF5) data using all nodes and cores available (and whenever possible): seamless *parallelism*



WHO IS USING IT

- Universities
- Research centers
- Biotech & renewable energy companies
- Automotive industry (electric vehicles)
- ... you name it! 😊



HOW IT WORKS

- HDFq1 allows the execution of many operations to properly manage HDF5 (data). These are grouped in four categories:
 - Data Definition Language (DDL): create files, create groups, rename datasets, alter (i.e. extend) dimensions of datasets, copy attributes, ...
 - Data Manipulation Language (DML): insert (i.e. write) data into datasets or attributes
 - Data Query Language (DQL): select (i.e. read) read data from datasets or attributes
 - Data Introspection Language (DIL): get group names, get dataset names (eventually stored in a certain group), get dimensions of attributes, ...

HOW IT WORKS (DDL)

- *CREATE FILE my_file.h5*

Create an HDF5 file named "my_file.h5"

- *CREATE FILE experiment.h5 IN PARALLEL*

Create an HDF5 file named "experiment.h5"
in parallel (i.e. using MPI)

- *CREATE GROUP countries*

Create a group named "countries"

- *CREATE DATASET values AS FLOAT(20, 40) ENABLE ZLIB*

Create a compressed dataset named "values" of
data type float of two dimensions (size 20x40)



HOW IT WORKS (DML)

- *INSERT INTO my_dataset VALUES(3, 5, 7)*

Insert (i.e. write) values 3, 5 and 7 into dataset "my_dataset"

- *INSERT INTO measurements VALUES FROM EXCEL FILE values.xlsx*

Insert (i.e. write) values from Excel file "values.xlsx" into dataset "measurements"

- *INSERT DIRECTLY INTO raw VALUES(10, 20)*

Insert (i.e. write) values 10 and 20 directly (i.e. bypass several internal processing steps of the HDF5 library itself) into dataset "raw"

- *INSERT INTO dset(0:::1) VALUES FROM MEMORY 0*

Insert (i.e. write) values from a user-defined variable (that was previously registered and assigned to number 0) into the first chunk of dataset "dset" (using an hyperslab selection)



HOW IT WORKS (DQL)

- *SELECT FROM values*

Select (i.e. read) data from dataset "values" and populate cursor in use with it
- *SELECT FROM measurements INTO EXCEL FILE values.xlsx*

Select (i.e. read) data from dataset "measurements" and write it into an Excel file "values.xlsx"
- *SELECT DIRECTLY FROM raw*

Select (i.e. read) data directly (i.e. bypass several internal processing steps of the HDF5 library itself) from dataset "raw" and populate cursor in use with it
- *SELECT FROM dset(3) INTO MEMORY 0*

Select (i.e. read) 4th value of dataset "dset" (using a point selection) and write it into a user-defined variable (that was previously registered and assigned to number 0)

HOW IT WORKS (DIL)

- *SHOW* 
- *SHOW DATASET my_group/* 
- *SHOW LIKE *** 
- *SHOW ATTRIBUTE group2 LIKE **/1|3* 



CURRENT FEATURES

- Portable across programming languages (C, C++, Java, Python, C#, Fortran and R) and platforms (Windows, Linux and macOS)
- Supports direct chunk write and read
- Supports both point and (irregular) hyperslab selections
- Supports parallel HDF5 (i.e. HDF5 + MPI)
- Supports reading data from a text, binary or Excel file and writing it into an HDF5 dataset/attribute (and vice-versa)
- And many more...



WHAT'S NEXT

- Support HDF5 library version 1.10.x
- Support virtual datasets (VDS)
- Support single-write multiple-readers (SWMR)
- Support dynamically loaded filters
- Support HDF5 data stored in Amazon E3 and Hadoop Distributed File System (HDFS)
- Support additional programming languages (e.g. Go, Rust, JavaScript) and tools (e.g. MATLAB) through proper wrappers



WHAT'S NEXT (cont.)

- Support pre/post processing data via pre-defined functions (using all nodes & cores available whenever possible) and user-defined functions.
 - INSERT INTO dset0 **MIN**(VALUES) FROM FILE input.txt ☞ call pre-defined function “MIN” to return the minimum value stored in text file “input.txt” and write the value into object “dset0”
 - SELECT FROM **AVG**(dset1) IN PARALLEL ☞ call pre-defined function “AVG” to return the average value of the data stored in object “dset1” in parallel (i.e. using MPI)
 - SELECT FROM **DUMMY**(dset2) ☞ load shared library “HDFq1DUMMY.so” (by HDFq1) and call user-defined function “DUMMY” to process the data stored in object “dset2” and return the result
- Support reading data from an HDF5 dataset/attribute and writing it into another HDF5 dataset/attribute (eventually using a function to pre/post process the data *en passant*).
 - INSERT INTO upper_names UPPER(VALUES) **FROM names** ☞ call pre-defined function “UPPER” to convert values stored in object “names” in upper-case and write the result into object “upper_names”
 - SELECT FROM AVG(dset) **INTO ATTRIBUTE dset/average** ☞ call pre-defined function “AVG” to return the average value of the data stored in object “dset” and write the value into attribute “average” attached to object “dset” (HDFq1 creates the attribute if it does not exist)



ADDITIONAL INFORMATION

- HDFq1 can be used both in commercial and non-commercial products, and completely free of charge
- Discuss on <https://forum.hdfgroup.org/c/hdf-tools>
- Support inquiries: <https://www.hdfq1.com/#contact>
- All versions of HDFq1 ever publicly released are available at <https://www.hdfq1.com/releases>
- Each released version of HDFq1 contains:
 - Libraries and wrappers for the C, C++, Java, Python, C#, Fortran and R programming languages
 - Examples that illustrate how to use HDFq1 in these programming languages
 - A command-line interface tool named “HDFq1CLI”
 - A complete reference manual describing HDFq1 and its operations

HDFql

the easy way to manage HDF5 data

New users, partnerships or potential sponsors always welcome

www.hdfql.com

