

What limits are there in HDF5?

We are aware that some of the *limits* questions below do not yet have answers. We will update this page as the information is obtained.

How large can an HDF5 file be?

There are no limitations on the file size.

What is the maximum size of an HDF5 dataset?

The library currently allows up to 32 dimension dataspace, and each dimension can have up to an unsigned 64-bit value. Each datatype can be somewhat arbitrarily large, since you can have array datatypes or compound datatypes that are very large. Multiplying those two factors together gives a theoretical upper limit probably in the thousands of bits of for a dataset's size. However, the library currently only supports 64-bit offsets (although that is easily adjustable when files over 16 exabytes are needed!), so that is the practical upper limit.

How many dimensions can there be in a dataspace and what is the maximum value for a dimension?

The library currently allows up to 32 dimension dataspace, and each dimension can have up to an unsigned 64-bit value.

How large can an attribute be?

Attributes have the same limitations as Datasets.

By default attributes are stored in the Object Header if they are small enough to fit. If they become too large to fit in the Object Header, they will be stored as dense attributes outside of the Object Header.

How many attributes can be attached to an object?

What is the maximum number of identifiers?

How many objects can be contained in a hdf5 file?

No specific limit.

How many links can be in a group?

No specific limit.

As an example, [grplots.c](#) and [grplots.f90](#) each create a file with 100,000 groups in the root group.

How long can a link name be?

How many members can be contained in a compound type?

The size of the compound datatype must not exceed 64KB (the limit of the object header, which is where the datatype is stored).

What is the limit on the chunk size?

The current maximum number of bytes in a chunk is $2^{32}-1$ (4 GB). As a result of this restriction, the number of elements in a chunk cannot be greater than 4 GB. You must also account for the datatype size of each element.

The chunk size can be larger than a dataset's dimension if the dataset's maximum dimension sizes are declared as unlimited or if the chunk size is less than the maximum dimension size. The chunk size for fixed-size datasets cannot exceed the size of a fixed-size dataset. For example, a dataset consisting of a 5 x 4 fixed-size array cannot be defined with 10 x 10 chunks, as the following error will occur:

```
#010: H5Dchunk.c line 550 in H5D__chunk_construct(): chunk size must be <= maximum dimension size for
      fixed-sized dimensions
      major: Dataset
      minor: Unable to initialize object
```

What is the limit on the number of external files that can be added to an EFL?

There is no internal limit.

Is there an object header limit and how does that affect HDF5?

There is a limit (in HDF5-1.8) of the object header, which is 64 KB. The datatype for a dataset is stored in the object header, so there is therefore a limit on the size of the datatype that you can have. (See H5FFV-1089)

Is there a limit on the number of processes you can use with Parallel HDF5?

No, not from HDF5's perspective.

What is the limit on the amount of data you can read in Parallel HDF5?

The limit is 2 GB because the MPI standard specifies that the 'count' parameter passed to `MPI_File_read()` be a 32-bit integer.