

h5perf_serial

[Expand all](#) [Collapse all](#)

- [Jump to ...](#)
- [Summary](#)
- [Description](#)
- [Example](#)
- [Switch language ...](#)
- [C](#)
- [C++](#)
- [FORTRAN](#)
- [JAVA](#)

[Summary](#)
[Description](#)
[Example](#)
[JAVA](#)
[FORTRAN](#)
[C++](#)
[C](#)

h5perf_serial

Tests HDF5 serial performance

Syntax:

```
h5perf_serial [-h | --help]  
h5perf_serial [options]
```

Description:

`h5perf_serial` provides tools for testing the performance of the HDF5 Library in serial mode.

See [h5perf_serial, a Serial File System Benchmarking Tool](#) for a complete description of this tool.

The following environment variable can be set to control the specified aspect of `h5perf_serial` behavior:

HDF5_NOCLEANUP	If set, <code>h5perf_serial</code> does not remove data files. (Default: Data files are removed.)
HDF5_PREFIX	Sets the prefix for output data files.

Options and Parameters:

The term *size specifier* is used as follows in this section:

A size specifier is an integer greater than or equal to 0 (zero) followed by a size indicator:

K	for kilobytes (1024 bytes)
M	for megabytes (1024 kilobytes)
G	for gigabytes (1073741824 bytes)

Example:

37M specifies 37 megabytes or 38797312 bytes.

<p>-A <i>api_list</i></p>	<p>Specifies which APIs to test. <i>api_list</i> is a comma-separated list with the following valid values:</p> <table border="1" data-bbox="820 325 1487 424"><tr><td>hdf5</td><td>HDF5 library APIs</td></tr><tr><td>posix</td><td>POSIX APIs</td></tr></table> <p>(Default: All APIs are monitored.)</p> <p>Example: -A hdf5, posix specifies that the HDF5 and POSIX APIs are to be monitored.</p>	hdf5	HDF5 library APIs	posix	POSIX APIs
hdf5	HDF5 library APIs				
posix	POSIX APIs				
<p>-c <i>chunk_size_list</i></p>	<p>Specifies chunked storage and defines chunks dimensions and sizes. (Default: Chunking is off.)</p> <p><i>chunk_size_list</i> is a comma-separated list of size specifiers. For example, a <i>chunk_size_list</i> value of</p> <p style="padding-left: 40px;">2K, 4K, 6M</p> <p>specifies that chunking is turned on and that chunk size is 2 kilobytes by 4 kilobytes by 6 megabytes.</p>				
<p>-e <i>dataset_size_list</i></p>	<p>Specifies dataset dimensionality and dataset dimension sizes. (Default dataset size is 100x200, or 100, 200.)</p> <p><i>dataset_size_list</i> is a comma-separated list of size specifiers, which are defined above.</p> <p>For example, a <i>dataset_size_list</i> value of</p> <p style="padding-left: 40px;">2K, 4K, 6M</p> <p>specifies a 2 kilobytes by 4 kilobytes by 6 megabytes dataset.</p>				
<p>-i <i>iterations</i></p>	<p>Specifies the number of iterations to perform. (Default: A single iteration, 1, is performed.)</p> <p><i>iterations</i> is an integer specifying the number of iterations.</p>				
<p>-r <i>access_order</i></p>	<p>Specifies dimension access order. (Default: 1, 2)</p> <p><i>access_order</i> is a comma-separated list of integers specifying the order of access. For example,</p> <p style="padding-left: 40px;">-r 1, 3, 2</p> <p>specifies the traversal of dimension 1 first, then dimension 3, and finally dimension 2.</p>				
<p>-t</p>	<p>Selects extendable HDF5 dataset dimensions. (Default: Datasets are fixed size.)</p>				
<p>-v <i>file_driver</i></p>	<p>Selects HDF5 driver to be used for HDF5 file access. (Default: sec2)</p> <p>Valid values are as follows:</p> <ul style="list-style-type: none">sec2stdiocoresplitmultifamilydirect				
<p>-w</p>	<p>Specifies the performance of write tests only, read performance will not be tested. (Default: Both write and read tests are performed.)</p>				
<p>-x <i>buffer_size_list</i></p>	<p>Specifies transfer buffer dimensions and sizes. (Default: 10, 20)</p>				

Exit Status:

0	Succeeded.
> 0	An error occurred.

History:

Release	Change
1.8.1	Tool introduced in this release.

--- Last Modified: August 28, 2019 | 10:43 AM