

H5F_SET_MDC_CONFIG

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

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

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

H5F_SET_MDC_CONFIG

Attempts to configure metadata cache of target file

Procedure:

H5F_SET_MDC_CONFIG (file_id, config_ptr)

Signature:

```
herr_t H5Fset_mdc_config(hid_t file_id, H5AC_cache_config_t *config_ptr)
```

Parameters:

<i>hid_t</i> file_id	IN: Identifier of the target file
<i>H5AC_cache_config_t</i> *config_ptr	IN: Pointer to the instance of <i>H5AC_cache_config_t</i> containing the desired configuration. The fields in this structure can be placed in several categories: General configuration Increment configuration Decrement configuration Parallel configuration Specific fields are described below.
General configuration fields:	
<i>int</i> version	IN: Integer field indicating the the version of the <i>H5AC_cache_config_t</i> in use. This field should be set to <i>H5AC__CURR_CACHE_CONFIG_VERSION</i> (defined in <i>H5ACpublic.h</i>).

<p><i>hbool_t</i> rpt_fcn_enabled</p>	<p>IN: Boolean flag indicating whether the adaptive cache resize report function is enabled. This field should almost always be set to disabled (0). Since resize algorithm activity is reported via stdout, it MUST be set to disabled (0) on Windows machines.</p> <p>The report function is not supported code, and can be expected to change between versions of the library. Use it at your own risk.</p>
<p><i>hbool_t</i> open_trace_File</p>	<p>IN: Boolean field indicating whether the <code>trace_file_name</code> field should be used to open a trace file for the cache.</p> <p>The trace file is a debugging feature that allows the capture of top level metadata cache requests for purposes of debugging and/or optimization. This field should normally be set to 0, as trace file collection imposes considerable overhead.</p> <p>This field should only be set to 1 when the <code>trace_file_name</code> contains the full path of the desired trace file, and either there is no open trace file on the cache, or the <code>close_trace_file</code> field is also 1.</p> <p>The trace file feature is unsupported unless used at the direction of The HDF Group. It is intended to allow The HDF Group to collect a trace of cache activity in cases of occult failures and/or poor performance seen in the field, so as to aid in reproduction in the lab. If you use it absent the direction of The HDF Group, you are on your own.</p>
<p><i>hbool_t</i> close_trace_file</p>	<p>IN: Boolean field indicating whether the current trace file (if any) should be closed.</p> <p>See the above comments on the <code>open_trace_file</code> field. This field should be set to 0 unless there is an open trace file on the cache that you wish to close.</p> <p>The trace file feature is unsupported unless used at the direction of The HDF Group. It is intended to allow The HDF Group to collect a trace of cache activity in cases of occult failures and/or poor performance seen in the field, so as to aid in reproduction in the lab. If you use it absent the direction of The HDF Group, you are on your own.</p>
<p><i>char</i> trace_file_name[]</p>	<p>IN: Full path of the trace file to be opened if the <code>open_trace_file</code> field is set to 1.</p> <p>In the parallel case, an ascii representation of the mpi rank of the process will be appended to the file name to yield a unique trace file name for each process.</p> <p>The length of the path must not exceed <code>H5AC__MAX_TRACE_FILE_NAME_LEN</code> characters.</p> <p>The trace file feature is unsupported unless used at the direction of The HDF Group. It is intended to allow The HDF Group to collect a trace of cache activity in cases of occult failures and/or poor performance seen in the field, so as to aid in reproduction in the lab. If you use it absent the direction of The HDF Group, you are on your own.</p>

<code>hbool_t evictions_enabled</code>	<p>IN: A boolean flag indicating whether evictions from the metadata cache are enabled. This flag is initially set to enabled (1).</p> <p>In rare circumstances, the raw data throughput requirements may be so high that the user wishes to postpone metadata writes so as to reserve I/O throughput for raw data. The <code>evictions_enabled</code> field exists to allow this. However, this is an extreme step, and you have no business doing it unless you have read the User Guide section on metadata caching, and have considered all other options carefully.</p> <p>The <code>evictions_enabled</code> field may not be set to disabled (0) unless all adaptive cache resizing code is disabled via the <code>incr_mode</code>, <code>flag_h_incr_mode</code>, and <code>decr_mode</code> fields.</p> <p>When this flag is set to disabled (0), the metadata cache will not attempt to evict entries to make space for new entries, and thus will grow without bound.</p> <p>Evictions will be re-enabled when this field is set back to 1. This should be done as soon as possible.</p>
<code>hbool_t set_initial_size</code>	IN: Boolean flag indicating whether the cache should be forced to the user specified initial size.
<code>size_t initial_size</code>	IN: If <code>set_initial_size</code> is set to 1, then <code>initial_size</code> must contain the desired initial size in bytes. This value must lie in the closed interval <code>[min_size, max_size]</code> . (see below)
<code>double min_clean_fraction</code>	<p>IN: This field specifies the minimum fraction of the cache that must be kept either clean or empty.</p> <p>The value must lie in the interval <code>[0.0, 1.0]</code>. 0.01 is a good place to start in the serial case. In the parallel case, a larger value is needed -- see Metadata Caching in HDF5 in the collection "Advanced Topics in HDF5."</p>
<code>size_t max_size</code>	IN: Upper bound (in bytes) on the range of values that the adaptive cache resize code can select as the maximum cache size.
<code>size_t min_size</code>	IN: Lower bound (in bytes) on the range of values that the adaptive cache resize code can select as the maximum cache size.
<code>long int epoch_length</code>	IN: Number of cache accesses between runs of the adaptive cache resize code. 50,000 is a good starting number.
Increment configuration fields:	
<code>enum H5C_cache_incr_mode incr_mode</code>	<p>IN: Enumerated value indicating the operational mode of the automatic cache size increase code. At present, only two values are legal:</p> <p><code>H5C_incr_off</code>: Automatic cache size increase is disabled, and the remaining increment fields are ignored.</p> <p><code>H5C_incr_threshold</code>: Automatic cache size increase is enabled using the hit rate threshold algorithm.</p>
<code>double lower_hr_threshold</code>	<p>IN: Hit rate threshold used by the hit rate threshold cache size increment algorithm.</p> <p>When the hit rate over an epoch is below this threshold and the cache is full, the maximum size of the cache is multiplied by increment (below), and then clipped as necessary to stay within <code>max_size</code>, and possibly <code>max_increment</code>.</p> <p>This field must lie in the interval <code>[0.0, 1.0]</code>. 0.8 or 0.9 is a good starting point.</p>

<code>double increment</code>	<p>IN: Factor by which the hit rate threshold cache size increment algorithm multiplies the current maximum cache size to obtain a tentative new cache size.</p> <p>The actual cache size increase will be clipped to satisfy the <code>max_size</code> specified in the general configuration, and possibly <code>max_increment</code> below.</p> <p>The parameter must be greater than or equal to 1.0 -- 2.0 is a reasonable value.</p> <p>If you set it to 1.0, you will effectively disable cache size increases.</p>
<code>hbool_t apply_max_increment</code>	<p>IN: Boolean flag indicating whether an upper limit should be applied to the size of cache size increases.</p>
<code>size_t max_increment</code>	<p>IN: Maximum number of bytes by which cache size can be increased in a single step -- if applicable.</p>
<code>enum H5C_cache_flash_incr_mode flash_incr_mode</code>	<p>IN: Enumerated value indicating the operational mode of the flash cache size increase code. At present, only the following values are legal:</p> <p><code>H5C_flash_incr__off</code>: Flash cache size increase is disabled.</p> <p><code>H5C_flash_incr__add_space</code>: Flash cache size increase is enabled using the add space algorithm.</p>
<code>double flash_threshold</code>	<p>IN: The factor by which the current maximum cache size is multiplied to obtain the minimum size entry / entry size increase which may trigger a flash cache size increase.</p> <p>At present, this value must lie in the range [0.1, 1.0].</p>
<code>double flash_multiple</code>	<p>IN: The factor by which the size of the triggering entry / entry size increase is multiplied to obtain the initial cache size increment. This increment may be reduced to reflect existing free space in the cache and the <code>max_size</code> field above.</p> <p>At present, this field must lie in the range [0.1, 10.0].</p>
Decrement configuration fields:	
<code>enum H5C_cache_decr_mode decr_mode</code>	<p>IN: Enumerated value indicating the operational mode of the automatic cache size decrease code. At present, the following values are legal:</p> <p><code>H5C_decr__off</code>: Automatic cache size decrease is disabled.</p> <p><code>H5C_decr__threshold</code>: Automatic cache size decrease is enabled using the hit rate threshold algorithm.</p> <p><code>H5C_decr__age_out</code>: Automatic cache size decrease is enabled using the ageout algorithm.</p> <p><code>H5C_decr__age_out_with_threshold</code>: Automatic cache size decrease is enabled using the ageout with hit rate threshold algorithm</p>
<code>double upper_hr_threshold</code>	<p>IN: Hit rate threshold for the hit rate threshold and ageout with hit rate threshold cache size decrement algorithms.</p> <p>When <code>decr_mode</code> is <code>H5C_decr__threshold</code>, and the hit rate over a given epoch exceeds the supplied threshold, the current maximum cache size is multiplied by decrement to obtain a tentative new (and smaller) maximum cache size.</p> <p>When <code>decr_mode</code> is <code>H5C_decr__age_out_with_threshold</code>, there is no attempt to find and evict aged out entries unless the hit rate in the previous epoch exceeded the supplied threshold.</p> <p>This field must lie in the interval [0.0, 1.0].</p> <p>For <code>H5C_incr__threshold</code>, .9995 or .99995 is a good place to start.</p> <p>For <code>H5C_decr__age_out_with_threshold</code>, .999 might be more useful.</p>

<i>double</i> decrement	<p>IN: In the hit rate threshold cache size decrease algorithm, this parameter contains the factor by which the current max cache size is multiplied to produce a tentative new cache size. The actual cache size decrease will be clipped to satisfy the min_size specified in the general configuration, and possibly max_decrement below.</p> <p>The parameter must be in the interval [0.0, 1.0].</p> <p>If you set it to 1.0, you will effectively disable cache size decreases. 0.9 is a reasonable starting point.</p>
<i>hbool_t</i> apply_max_decrement	<p>IN: Boolean flag indicating whether an upper limit should be applied to the size of cache size decreases.</p>
<i>size_t</i> max_decrement	<p>IN: Maximum number of bytes by which the maximum cache size can be decreased in any single step -- if applicable.</p>
<i>int</i> epochs_before_eviction	<p>IN: In the ageout based cache size reduction algorithms, this field contains the minimum number of epochs an entry must remain unaccessed in cache before the cache size reduction algorithm tries to evict it. 3 is a reasonable value.</p>
<i>hbool_t</i> apply_empty_reserve	<p>IN: Boolean flag indicating whether the ageout based decrement algorithms will maintain an empty reserve when decreasing cache size.</p>
<i>double</i> empty_reserve	<p>IN: Empty reserve as a fraction of maximum cache size if applicable. When so directed, the ageout based algorithms will not decrease the maximum cache size unless the empty reserve can be met.</p> <p>The parameter must lie in the interval [0.0, 1.0]. 0.1 or 0.05 is a good place to start.</p>
Parallel configuration field:	
<i>int</i> dirty_bytes_threshold	<p>IN: Threshold number of bytes of dirty metadata generation for triggering synchronizations of the metadata caches serving the target file in the parallel case.</p> <p>Synchronization occurs whenever the number of bytes of dirty metadata created since the last synchronization exceeds this limit.</p> <p>This field only applies to the parallel case. While it is ignored elsewhere, it can still draw a value out of bounds error.</p> <p>It must be consistent across all caches on any given file.</p> <p>By default, this field is set to 256 KB. It shouldn't be more than half the current maximum cache size times the minimum clean fraction.</p>

Description:

H5F_SET_MDC_CONFIG attempts to configure the file's metadata cache according configuration supplied in **config_ptr*.

See the overview of the metadata cache in the special topics section of the user manual for details on what is being configured. If you haven't read and understood that documentation, you really shouldn't be using this API call.

Returns:

Returns a non-negative value if successful; otherwise returns a negative value.

Example:

History:

Release	Change
---------	--------

1.8.0

Function introduced in this release.

--- Last Modified: July 22, 2020 | 03:04 PM