

H5P_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](#)

H5P_SET_MDC_CONFIG

Set the initial metadata cache configuration in the indicated File Access Property List to the supplied value

Procedure:

H5P_SET_MDC_CONFIG (plist_id, config_ptr)

Signature:

```
herr_t H5Pset_mdc_config(hid_t  
    plist_id, H5AC_cache_config_t *config_ptr)
```

Parameters:

<i>hid_t</i> plist_id	IN: Identifier of the file access property list
<i>H5AC_cache_config_t</i> *config_ptr	IN: Pointer to the instance of H5AC_cache_config_t containing the desired configuration

The fields of the H5AC_cache_config_t structure are discussed below:

General configuration section:

<i>int</i> version	IN: Integer field indicating the the version of the H5AC_cache_config_t in use. This field should be set to H5AC__CURR_CACHE_CONFIG_VERSION (defined in H5ACpublic.h).
<i>hbool_t</i> rpt_fcn_enabled	IN: Boolean flag indicating whether the adaptive cache resize report function is enabled. This field should almost always be set to FALSE. Since resize algorithm activity is reported via stdout, it MUST be set to FALSE on Windows machines. The report function is not supported code, and can be expected to change between versions of the library. Use it at your own risk.
<i>hbool_t</i> open_trace_file	IN: Boolean field indicating whether the trace_file_name field should be used to open a trace file for the cache. 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 FALSE, as trace file collection imposes considerable overhead. This field should only be set to TRUE when the trace_file_name contains the full path of the desired trace file, and either there is no open trace file on the cache, or the close_trace_file field is also TRUE. 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.
<i>hbool_t</i> close_trace_file	IN: Boolean field indicating whether the current trace file (if any) should be closed See the above comments on the open_trace_file field. This field should be set to FALSE unless there is an open trace file on the cache that you wish to close. 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.
<i>char</i> trace_file_name[]	IN: Full path of the trace file to be opened if the open_trace_file field is TRUE 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. The length of the path must not exceed H5AC__MAX_TRACE_FILE_NAME_LEN characters. 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.

<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 <code>TRUE</code>.</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 <code>FALSE</code> unless all adaptive cache resizing code is disabled via the <code>incr_mode</code>, <code>flash_incr_mode</code>, and <code>decr_mode</code> fields.</p> <p>When this flag is set to <code>FALSE</code>, 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 <code>TRUE</code>. 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 created with a user specified initial size
<code>size_t initial_size</code>	IN: If <code>set_initial_size</code> is <code>TRUE</code> , <code>initial_size</code> must contain the desired initial size in bytes. This value must lie in the closed interval [<code>min_size</code> , <code>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 [0.0, 1.0]. 0.01 is a good place to start in the serial case. In the parallel case, a larger value is needed -- see the overview of the metadata cache in the "Metadata Caching in HDF5" section of the <i>HDF5 User's Guide</i> for details.</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 section:

<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 <code>increment</code> (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 [0.0, 1.0]. 0.8 or 0.9 is a good place to start.</p>

<i>double</i> increment	<p>IN: Factor by which the hit rate threshold cache size increment algorithm multiplies the current cache max 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>
<i>hbool_t</i> apply_max_increment	<p>IN: Boolean flag indicating whether an upper limit should be applied to the size of cache size increases.</p>
<i>size_t</i> max_increment	<p>IN: Maximum number of bytes by which cache size can be increased in a single step -- if applicable.</p>
<i>enum H5C_cache_flash_incr_mode</i> flash_incr_mode	<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>
<i>double</i> flash_threshold	<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>
<i>double</i> flash_multiple	<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 section:

<i>enum H5C_cache_decr_mode</i> decr_mode	<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>
<i>double</i> upper_hr_threshold	<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 <code>min_size</code> specified in the general configuration, and possibly <code>max_decrement</code> 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 section:

<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. 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 max cache size times the min clean fraction.</p>
<i>int</i> metadata_write_strategy	<p>IN: Desired metadata write strategy. The valid values for this field are:</p> <p><code>H5AC_METADATA_WRITE_STRATEGY__PROCESS_0_ONLY</code>: Specifies that only process zero is allowed to write dirty metadata to disk.</p> <p><code>H5AC_METADATA_WRITE_STRATEGY__DISTRIBUTED</code>: Specifies that process zero still makes the decisions as to what entries should be flushed, but the actual flushes are distributed across the processes in the computation to the extent possible.</p> <p>The <code>src/H5ACpublic.h</code> include file in the HDF5 library has detailed information on each strategy.</p>

Description:

`H5P_SET_MDC_CONFIG` attempts to set the initial metadata cache configuration to the supplied value. It will fail if an invalid configuration is detected. This configuration is used when the file is opened.

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 have not read and understood that documentation, you really should not be using this API call.

Returns:

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

Example:

Coming Soon!

--- Last Modified: July 23, 2019 | 08:38 AM