

Collective Metadata I/O

The purpose of this page is to briefly describe the new HDF5 Collective Metadata I/O feature and provide a gateway to available documentation. The page includes the following sections:

- [Collective Metadata I/O Overview](#)
- [Collective Metadata I/O User and Resource Documents](#)
- [HDF5 Library APIs](#)

Collective Metadata I/O Overview

Calls for HDF5 metadata can result in many small reads and writes. On metadata reads, collective metadata I/O can improve performance by allowing the library to perform optimizations when reading the metadata, by having one rank read the data and broadcasting it to all other ranks.

Collective metadata I/O improves metadata write performance through the construction of an MPI derived datatype that is then written collectively in a single call.

Collective Metadata I/O User and Resource Documents

HDF5 Collective Metadata I/O User Document (*This document is not yet available.*)

Until an —*HDF5 Collective Metadata I/O User Document* becomes available, users may find the following resources helpful:

[RFC: Collective Metadata Reads](#) (PDF)

[RFC: Collective Metadata Writes](#) (PDF) Taken together, these papers discuss the motivation, design, implementation, and API for HDF5's Collective Metadata I/O feature.

HDF5 Library APIs

New Collective Metadata I/O Functions

H5P_SET_COLL_METADATA_WRITE	Establishes I/O mode property setting, collective or independent, for metadata writes
H5P_GET_COLL_METADATA_WRITE	Retrieves I/O mode property setting for metadata writes
H5P_SET_ALL_COLL_METADATA_OPS	Establishes I/O mode, collective or independent, for metadata read operations
H5P_GET_ALL_COLL_METADATA_OPS	Retrieves I/O mode for metadata read operations

Additional API Reference

Functions with No Access Property List Parameter that May Generate Metadata Reads

This page lists the functions associated with HDF5 operations that can issue metadata reads from the metadata cache, but that currently provide no access property list parameter by which to set a collective requirement individually for those read operations. The need for this document is discussed in [H5P_SET_ALL_COLL_METADATA_OPS](#).