
Enabling a Strict Consistency Semantics Model in Parallel HDF5

Introduced with
HDF5 Release 1.8.9
in
May 2012



Copyright Notice and License Terms for HDF5 (Hierarchical Data Format 5) Software Library and Utilities

HDF5 (Hierarchical Data Format 5) Software Library and Utilities
Copyright 2006-2012 by The HDF Group.

NCSA HDF5 (Hierarchical Data Format 5) Software Library and Utilities
Copyright 1998-2006 by the Board of Trustees of the University of Illinois.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted for any purpose (including commercial purposes) provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions, and the following disclaimer in the documentation and/or materials provided with the distribution.
3. In addition, redistributions of modified forms of the source or binary code must carry prominent notices stating that the original code was changed and the date of the change.
4. All publications or advertising materials mentioning features or use of this software are asked, but not required, to acknowledge that it was developed by The HDF Group and by the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign and credit the contributors.
5. Neither the name of The HDF Group, the name of the University, nor the name of any Contributor may be used to endorse or promote products derived from this software without specific prior written permission from The HDF Group, the University, or the Contributor, respectively.

DISCLAIMER: THIS SOFTWARE IS PROVIDED BY THE HDF GROUP AND THE CONTRIBUTORS "AS IS" WITH NO WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED. In no event shall The HDF Group or the Contributors be liable for any damages suffered by the users arising out of the use of this software, even if advised of the possibility of such damage.

Contributors: National Center for Supercomputing Applications (NCSA) at the University of Illinois, Fortner Software, Unidata Program Center (netCDF), The Independent JPEG Group (JPEG), Jean-loup Gailly and Mark Adler (gzip), and Digital Equipment Corporation (DEC).

Portions of HDF5 were developed with support from the Lawrence Berkeley National Laboratory (LBNL) and the United States Department of Energy under Prime Contract No. DE-AC02-05CH11231.

Portions of HDF5 were developed with support from the University of California, Lawrence Livermore National Laboratory (UC LLNL). The following statement applies to those portions of the product and must be retained in any redistribution of source code, binaries, documentation, and/or accompanying materials:

This work was partially produced at the University of California, Lawrence Livermore National Laboratory (UC LLNL) under contract no. W-7405-ENG-48 (Contract 48) between the U.S. Department of Energy (DOE) and The Regents of the University of California (University) for the operation of UC LLNL.

DISCLAIMER: This work was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately-owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Contents

1. Introduction	4
2. Motivation.....	6
3. PHDF5 Metadata Consistency Semantics.....	7
4. PHDF5 Raw Data Consistency Semantics	9
5. Adding New Routines to the Library	10

1. Introduction

Consistency semantics are rules that define the outcome of multiple, possibly concurrent, accesses to an object or data structure by one or more processes in a computer system. An I/O library defines a set of consistency semantics to let users know what to expect when processes access data managed by that library. Parallel HDF5 (PHDF5) provides users with an option to have multiple processes perform I/O to an HDF5 file at the same time. In most cases the file resides on a parallel file system, which allows striping of data over multiple servers to achieve high performance and provides mechanisms for concurrent data access by many processes. Therefore, an application that utilizes parallel HDF5 needs to be aware of the consistency semantics provided by the HDF5 Library. These consistency semantics define when the changes a process makes are actually visible to itself (if it tries to read back that data) or to other processes that access the same file with independent or collective I/O operations. This document describes the following:

- The consistency semantics model for the Parallel HDF5 Library when it comes to metadata and raw data access
- Some issues with the model
- New features that have been added to the library which can be used to overcome the issues with the model

The HDF5 library provides a set of virtual file drivers that perform the lower level I/O accesses to an HDF5 file, including an MPI-I/O file driver for performing parallel file access. Each of those drivers has certain consistency rules that define the outcome of accessing an HDF5 file with it. There are several scenarios that arise from parallel access to a parallel file system using the MPI-based file driver, but the most common cases include:

1. Each process accesses its own separate file: there are no conflicts in this case, which means data written by a process is visible immediately to that process.
2. All processes open the file(s) in *read-only* mode: all processes will read exactly what is in the file and no possibility for conflicts exists.
3. Multiple processes access the same file with read-write operations: this is the scenario discussed in this document.

Depending on the file driver being used, the consistency semantics vary. For example, the MPI-POSIX driver guarantees sequential consistency¹ because POSIX I/O provides that guarantee. Consider a scenario with two processes accessing data at the same location in the same file:

Process 0	Process 1
1) <code>write ()</code>	1)
2) <code>MPI_Barrier()</code>	2) <code>MPI_Barrier()</code>
3)	3) <code>read ()</code>

POSIX I/O guarantees that Process 1 will read what Process 0 has written: the atomicity of the read and write calls and the synchronization using the barrier ensures that Process 1 will call the read function after Process 0 is finished with the write function.

¹ The term sequential consistency means that all operations to access data on disk appear to be atomic and consistent with the order of operations in a program.

On the other hand, the MPI-I/O driver does not give that guarantee. If we take the example above and substitute `MPI_File_write_at()` and `MPI_File_read_at()` for the `write()` and `read()` calls respectively, MPI-I/O with its default settings does not guarantee sequential consistency. In the example, Process 1 is not guaranteed that it will read what Process 0 writes. In fact, the read might occur before, after, or even during the write operation. MPI-I/O does, however, provide different options to ensure sequential consistency or user-imposed consistency, which we will discuss further in the next sections. For a full discussion on MPI-I/O consistency semantics, please refer to the MPI-2.2 standard (<http://www.mpi-forum.org/docs/> section 13.6, page 437, line 44).

2. Motivation

MPI provides the user with options to impose different levels of consistency. If the user enables file atomicity, all MPI access functions are guaranteed to execute atomically. For example:

Process 0	Process 1
1) MPI_File_set_atomicsity (fh, TRUE)	1) MPI_File_set_atomicsity (fh, TRUE)
2) MPI_File_write_at ()	2) MPI_File_read_at ()

The read by process 1 will either read the entire data before the write by process 0 occurs or after (but not during).

If the user wishes to impose an ordering between the read and write operations, a barrier (or any other method to ensure ordering) can be used. For example:

Process 0	Process 1
1) MPI_File_set_atomicsity (fh, TRUE)	1) MPI_File_set_atomicsity (fh, TRUE)
2) MPI_File_write_at ()	2)
3) MPI_Barrier ()	3) MPI_Barrier ()
4)	4) MPI_File_read_at ()

This will ensure that Process 1 reads what Process 0 has written (the read happens logically after the write).

There are other ways for the user to impose this level of consistency on file access operations without enabling file atomicity with `MPI_File_set_atomicsity`:

1. Close the file after the write operation (with all processes) and re-open it, and then issue the read operation.
2. Ensure that no write sequence on any process is concurrent with a read or write sequence on any other process, as shown in the next table. This is a general characterization of the sync-barrier-sync construct and is a more complicated way to impose a strict consistency semantics model. Please refer to the MPI standard for a more detailed discussion on this method.

Process 0	Process 1
1) MPI_File_write_at ()	1)
2) MPI_File_sync ()	2) MPI_File_sync ()
3) MPI_Barrier ()	3) MPI_Barrier ()
4) MPI_File_sync ()	4) MPI_File_sync ()
5)	5) MPI_File_read_at ()

The previous sequence will ensure that process 1 reads what process 0 has written.

3. PHDF5 Metadata Consistency Semantics

HDF5 files can contain several types of metadata. For a discussion of these different types of metadata, please refer to the “*HDF5 Metadata*” document in the Advanced Topics section of the documentation (http://www.hdfgroup.org/HDF5/doc/Advanced/HDF5_Metadata/index.html). In the parallel version of the HDF5 Library, there are several constraints that limit the ability to update HDF5 metadata on disk. To handle synchronization issues, all HDF5 operations that could potentially modify the metadata in the HDF5 file are required to be collective. A list of those routines is available in the *HDF5 Reference Manual* (<http://www.hdfgroup.org/HDF5/doc/RM/CollectiveCalls.html>).

Since all operations that modify metadata in the HDF5 file are collective, all processes will have the same dirty metadata entries in their cache (in other words, metadata that is inconsistent with what is on disk). There is no communication necessary to mark metadata as dirty in a process's cache, as all the calls are collective, and all processes see the same stream of changes to metadata. Operations that only read metadata (for example, `H5Dopen`) are not required to be collective, and so processes are allowed to have different clean metadata entries (in other words, metadata that is in sync with what is on disk).

Within the HDF5 library, the metadata cache running on process 0 of the file communicator is responsible for managing changes to the metadata in the HDF5 file. All the other caches must retain dirty metadata until the process 0 cache tells them that the metadata is clean (in other words, synchronized with data on disk). The actual flush of metadata entries to disk is currently implemented in two ways:

A – Single Process (Process 0) writes:

1. perform a barrier
2. process 0 writes each piece of metadata individually
3. process 0 broadcasts the list of flushed metadata to other processes
4. all processes mark the flushed metadata entries as clean in their cache

B – Distributed:

1. perform a barrier
2. process 0 broadcasts the list of metadata to flush to other processes
3. each process determines which pieces of metadata it should write
4. each process writes each piece of metadata independently
5. all processes participate in an exchange of whether any errors occurred
6. all processes mark all flushed metadata entries as clean in their cache

Please refer to the “*Metadata Caching in HDF5*” entry in the Advanced Topics section of the documentation (<http://www.hdfgroup.org/HDF5/doc/Advanced/MetadataCache/index.html>) for more details.

At first glance, both of these implementations appear to guarantee the sequential consistency of access to the file's metadata, since there can be no conflicting access by different processes to the same metadata at once. However, considering the default consistency semantics of MPI-I/O (which is the interface used to read/write metadata on disk for parallel HDF5), sequential consistency can be broken with one particular access pattern:

1. All processes have a dirty metadata entry M in their cache.
2. Using either approach described above, metadata entry M is flushed to disk by process P and marked as clean on all processes.

3. Process Q evicts metadata entry M because it needs more space in its cache.
4. Process Q needs metadata entry M again, and since it was evicted it needs to read it from disk

Process Q was not the process that wrote metadata entry M to disk, so there is no guarantee that metadata entry M written out by process P actually reached disk before the read call on process Q is issued.

To put this in an HDF5 example, consider the following scenario:

Process 0	Process 1
1) <code>H5Acreate()</code>	1) <code>H5Acreate()</code>
2) <code>H5Awrite()</code>	2) <code>H5Awrite()</code>
3) ...	3) ...
4)	4) <code>H5Aread()</code>

All processes create an attribute and write data to that attribute collectively. Process 1 comes in later and reads the same data that was written in the collective write call. Suppose that before process 1 issues the `H5Aread()` call, a metadata flush is triggered and using either approach A or B (described earlier), process 0 ends up writing the metadata of the attribute to disk, and then informs process 1 that this piece of metadata is clean, so process 1 evicts it from its cache. When process 1 calls `H5Aread()`, it goes back to disk to read the metadata for the attribute, specifically the data written in step 2. However, since MPI, in its default mode, is not atomic and does not guarantee sequential consistency, process 1 might actually read the old, the new one, or even an undefined value. This scenario can be translated to the following simple access to the same region in the file:

Process 0	Process 1
1) <code>MPI_File_write_at()</code>	1)
2) <code>MPI_Barrier()</code>	2) <code>MPI_Barrier()</code>
3)	3) <code>MPI_File_read_at()</code>

In this case, as explained in the previous section, MPI-I/O does not guarantee that process 1 will actually read what process 0 has written, and therefore the consistency semantics for metadata operations are not guaranteed to be sequential with the current implementation.

One solution to the above issue is to enable MPI atomic mode. With atomic mode enabled, all file operations are guaranteed to execute with sequential consistency semantics. The main drawback to enabling atomic mode is that the performance drops significantly when reading/writing to the file.

The race condition described above is rare and has not yet been reported by any application that has used parallel HDF5. While it's important to handle the issue, it might be valuable to switch the metadata writes from the cache to be collective first, and then look at the consistency issue later, when more efficient solutions may be available in later releases of MPI. Additionally, this issue will disappear when the metadata server implementation we are currently working on is deployed.

4. PHDF5 Raw Data Consistency Semantics

As mentioned previously, the consistency semantics of data access is dictated by the virtual file driver (VFD) used to access the file. In the Parallel HDF5 Library, when the MPI-I/O driver is used, the MPI-I/O consistency semantics model is enforced. HDF5 uses the default MPI semantics which do not guarantee sequential consistency. This means that access to a file happens in a non-atomic and arbitrary order. Although MPI-I/O provides an option for atomic (by enabling atomicity) and ordered access to the file (using sync-barrier-sync or atomicity-barrier), the Parallel HDF5 Library does not provide this functionality through its API. A consistent read after write operation on different processes (a process reading the data that another process wrote) can be achieved with the current API only by closing the file after the write and reopening it. Otherwise, a dataset read operation following a write (to the same dataset elements by a different process) is considered undefined.

5. Adding New Routines to the Library

To match the consistency semantic functionality that MPI offers, new functions have been or will be added to the library's API. The functions that have been added are `H5Fset_mpio_atomicity` and `H5Fget_mpio_atomicity`. The `H5Fsync` function will be added in a future version of the library.

`H5Fset_mpio_atomicity (hid_t file_id, hbool_t flag)`: enables (flag = 1) or disables (flag = 0) atomicity on a file opened with the MPI-IO VFD.

`H5Fget_mpio_atomicity (hid_t file_id, hbool_t *flag)`: retrieves the atomicity setting for a file opened with the MPI-IO VFD.

`H5Fsync (hid_t file_id)`: this call causes all previous writes to the file to be transferred to the storage device.

Setting atomicity to true will call the MPI atomicity function (`MPI_File_set_atomicity`) which sets atomicity to true on the MPI file handle. The sync function will also call its MPI counterpart (`MPI_File_sync`). Using the two routines described above, the user will be able to enforce stricter consistency semantics in a way similar to what MPI provides. To be able to guarantee sequential consistency among different processes, the use of barriers (`MPI_Barrier`) will be required by either:

- setting atomicity to true, then calling `MPI_Barrier()` to synchronize processes after writes/before reads
- using the sync-barrier-sync construct

In certain scenarios, however, just setting atomicity to true in the MPI-I/O file driver does not suffice, because an `H5Dwrite` (for example) may translate into multiple `MPI_File_write_at` functions underneath. This happens in all cases where the high level file access routine translates to multiple lower level file access routines. The following scenarios will encounter this issue:

- Non-contiguous file access using independent I/O
- Partial collective I/O using chunked access
- Collective I/O using filters or when data conversion needs to happen

In any of these scenarios, the following case:

Process 0	Process 1
1) <code>H5Fset_mpio_atomicity (fid, true)</code>	1) <code>H5Fset_mpio_atomicity (fid, true)</code>
2) <code>H5Dwrite()</code>	2) <code>H5Dread()</code>

might not yield the expected result (the read happens as if the write was executed before or after it) if `H5Fset_mpio_atomicity` is implemented by just setting MPI-I/O atomicity to true, because it could translate into the following:

Process 0	Process 1
MPI_File_set_atomicsity (fh, 1)	MPI_File_set_atomicsity (fh, 1)
MPI_File_write_at()	MPI_File_read_at()
MPI_File_write_at()	MPI_File_read_at()
MPI_File_write_at()	MPI_File_read_at()
...	...

This means that atomicity is done per MPI file access operation and not per HDF5 operation, which is probably not what the user wants. The use of barriers after the `H5Dwrite()` and before the `H5Dread()` in addition to setting atomicity to true in the user's application would ensure that the `H5Dread()` and `H5Dwrite()` operations are atomic:

Process 0	Process 1
1) H5Fset_mpio_atomicsity (fid, true)	1) H5Fset_mpio_atomicsity (fid, true)
2) H5Dwrite()	2)
3) MPI_Barrier()	3) MPI_Barrier()
4)	4) H5Dread()

This does, however, also impose additional ordering semantics, where the operations are not only atomic, but also that the `H5Dread()` operation will occur after the `H5Dwrite()` operation. The barrier will guarantee that all underlying write operations will execute atomically before the read operations starts. Future development on the HDF5 library can overcome this limitation, but is deferred currently.