
Newsletter #148

March 31, 2016

Release of HDF5-1.10.0

We are very pleased to announce the release of HDF5-1.10.0.

Links to the HDF5 1.10.0 source code, documentation, and additional materials can be found on the [HDF5 web page](#).

The HDF5 1.10.0 release can be obtained directly from:

<https://www.hdfgroup.org/HDF5/release/obtain5110.html>

User documentation for 1.10.0 can be accessed from:

<https://www.hdfgroup.org/HDF5/doc/>

This release contains many important new features and changes, including those listed below. The features marked with * use new extensions to the HDF5 file format.

- The Single-Writer / Multiple-Reader or SWMR feature enables users to read data while concurrently writing it. *
- The virtual dataset (VDS) feature enables users to access data in a collection of HDF5 files as a single HDF5 dataset and to use the HDF5 APIs to work with that dataset. *

NOTE: There is a known issue with the h5repack utility when using it to modify the layout of a VDS. We understand the issue and are working on a patch for it.

- New indexing structures for chunked datasets were added to support SWMR and to optimize performance. *
- Persistent free file space can now be managed and tracked for better performance. *
- The HDF5 Collective Metadata I/O feature has been added to improve performance when reading and writing data collectively with Parallel HDF5.
- The Java HDF5 JNI has been integrated into HDF5.
- Changes were made in how autotools handles large file support (LFS).

- New options for the storage and filtering of partial edge chunks have been added for performance tuning. *

Files created with the new extensions will not be readable by applications based on the HDF5-1.8 library.

Please refer to the [New Features in HDF5 Release 1.10.0](#) for what is new in this release.

All new and modified APIs are listed in detail in the [HDF5 Software Change from Release to Release](#) document.

For detailed information regarding this release see the [release notes](#).