



The HDF Group



HDF5 Update

High-Data Rate MX Meeting

NSLS-II at BNL

May 26, 2016

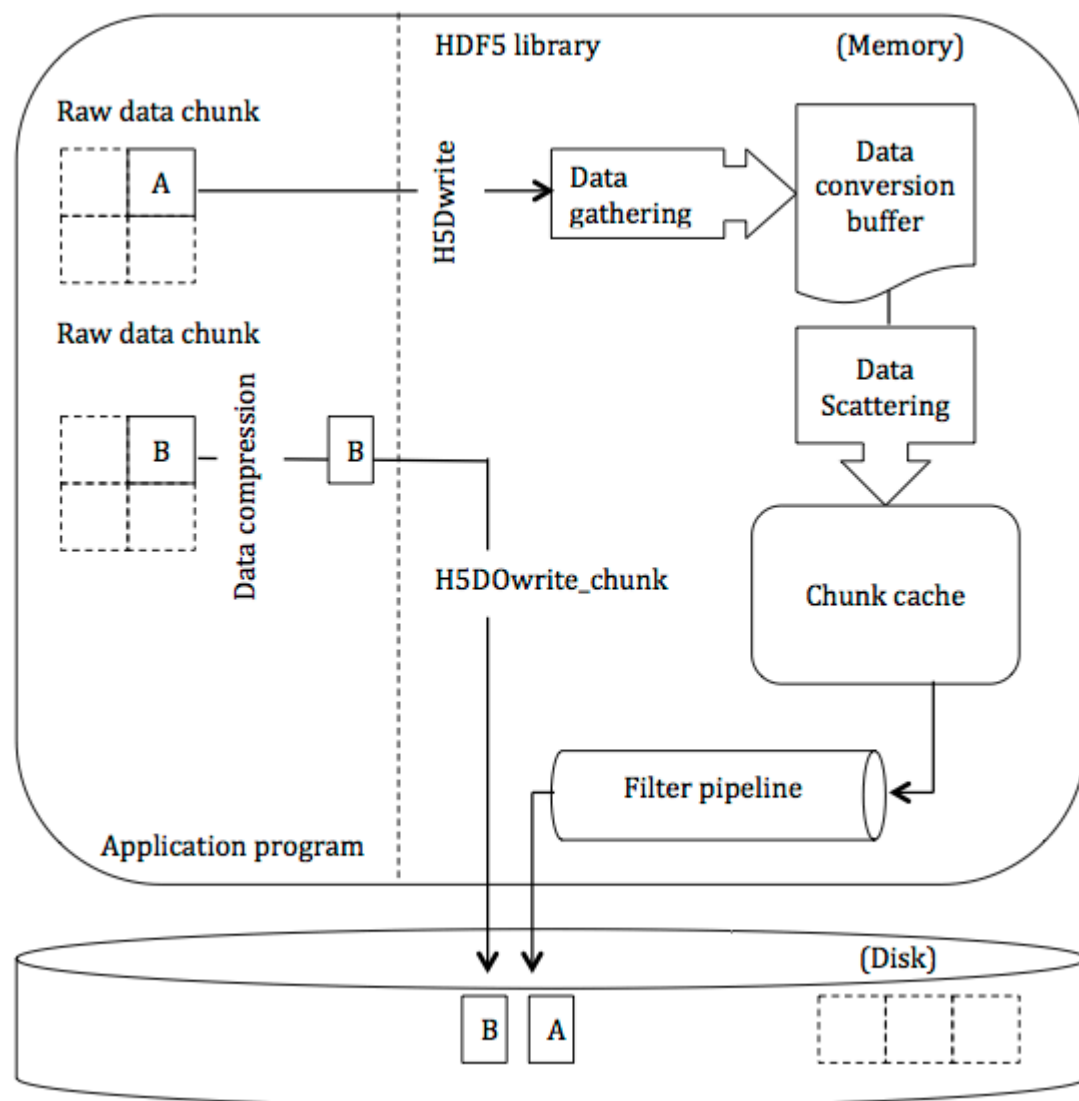
epournal@hdfgroup.org



- HDF5 features requested and sponsored by Synchrotron Community (2012 -2016)
 - Direct chunk I/O (PSI, Dectris)
 - Dynamically loaded filters (DESY)
 - Single Write/ Multiple Reader or SWMR (DLS, ESRF, DESY)
 - Virtual Dataset (VDS) (DLS, DESY, XFEL)
- New requirements and features
- HDF5 roadmap for 2016



Direct chunk write: H5Dwrite_chunk



Performance results for H5Dwrite_chunk

Test result on Lunux 2.6, x86_64

Each dataset contained 100 chunks, written by chunks

Dataset size (MB)	95.37		762.94		2288.82	
Size after compression (MB)	64.14		512.94		1538.81	
Dataset dimensionality	100x1000x250		100x2000x1000		100x2000x3000	
Chunk dimensionality	1000x250		2000x1000		2000x3000	
Datatype	4-byte integer		4-byte integer		4-byte integer	
	speed ¹	time ²	speed	time	speed	time
H5Dwrite writes without compression filter	77.27	1.23	97.02	7.86	91.77	24.94
H5Dwrite_chunk writes uncompressed data	79	1.21	95.71	7.97	89.17	25.67
H5Dwrite writes with compression filter	2.68	35.59	2.67	285.75	2.67	857.24
H5Dwrite_chunk writes compressed data	77.19	0.83	78.56	6.53	96.28	15.98
Unix writes compressed data to Unix file	76.49	0.84	95	5.4	98.59	15.61

1 Speed in MB/s
2 Time in seconds



Dynamically loaded filters

- Problems with using custom filters
 - “Off the shelf” tools do not work with the third-party filters
 - h5dump, MATLAB and IDL, etc.
 - HDF5 tools cannot read file created by
 - h5py, PyTables , etc.
- Solution
 - Modify HDF5 source with your code and distribute it
 - And what will happen if a user wants filters from the different distributions???? Oh.... No....
 - Use a 1.8.11 and later
 - Provide maintained library of HDF5 compression filters



Approach

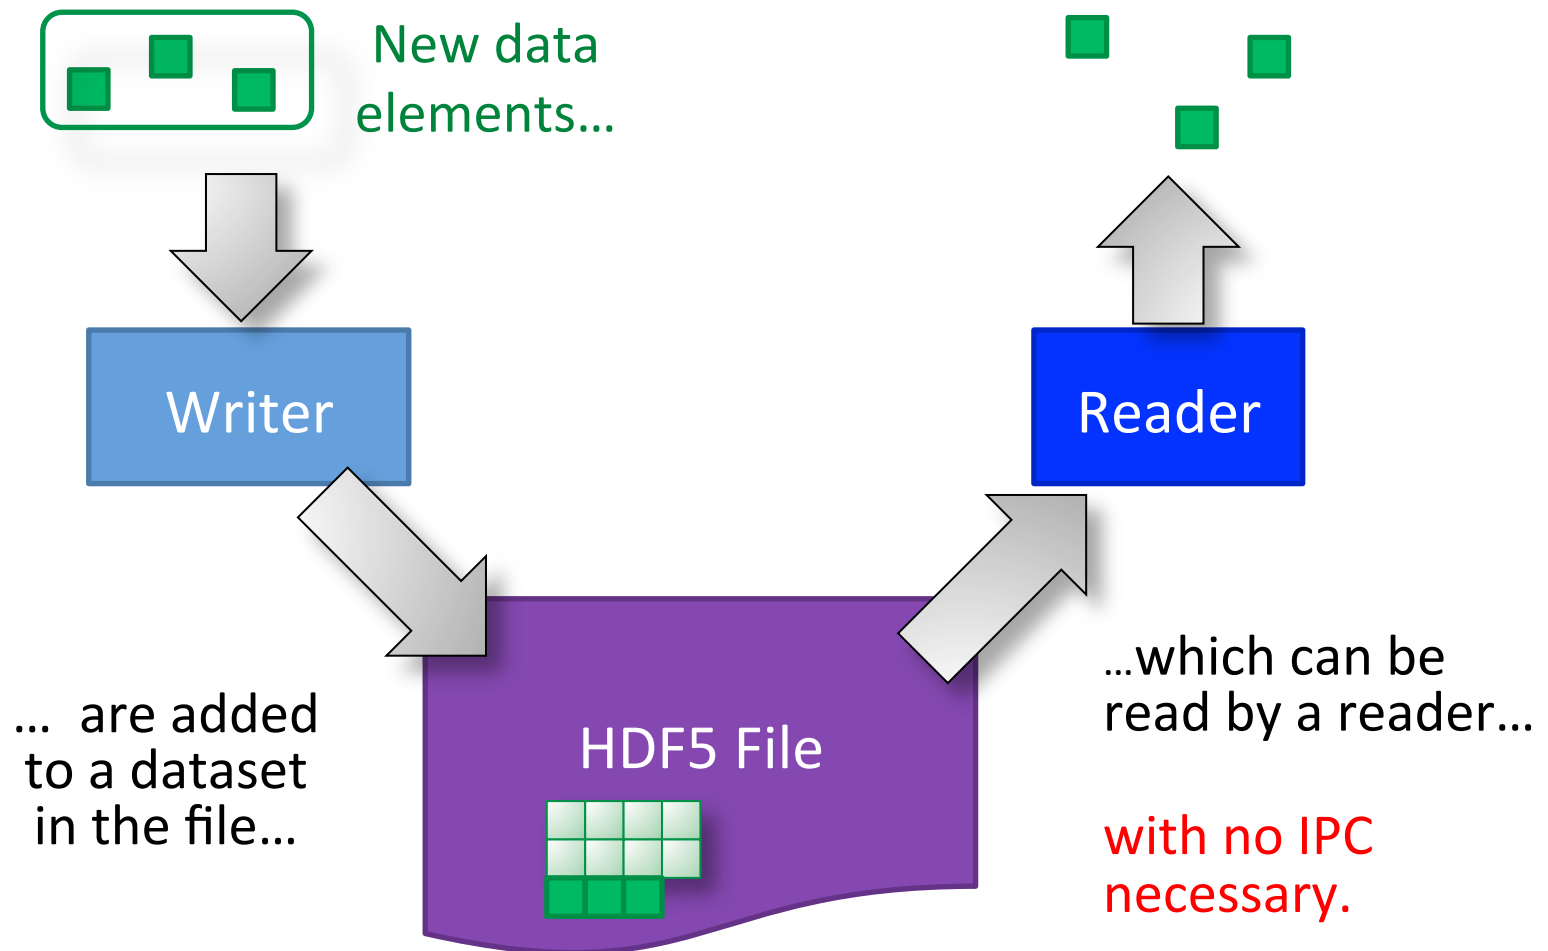
- There are predefined default locations where the HDF5 library searches the shared libraries or DLLs with the HDF5 filter functions.
- The default location may be overwritten by an environment variable.
- Encoding: Standard programming model to register compression filter on write operation
- Decoding: Transparent on read; no user action required.
- Works great, but....

User's headache and possible solutions

- Where to find HDF5 plugins?
- No official distribution site
 - Example of HDF5 filters distributions on Github
 - <https://github.com/dectris/HDF5Plugin>
 - <https://github.com/Blosc/hdf5-blosc>
 - <https://github.com/nexusformat/HDF5-External-Filter-Plugins>
 - https://svn.hdfgroup.org/hdf5_plugins/
- The HDF Group can help with
 - Maintaining common repository
 - Testing plugin with the releases
 - Distributing source and binaries for Linux, Mac and Windows
 - Free and licensed (for fee) binary distributions

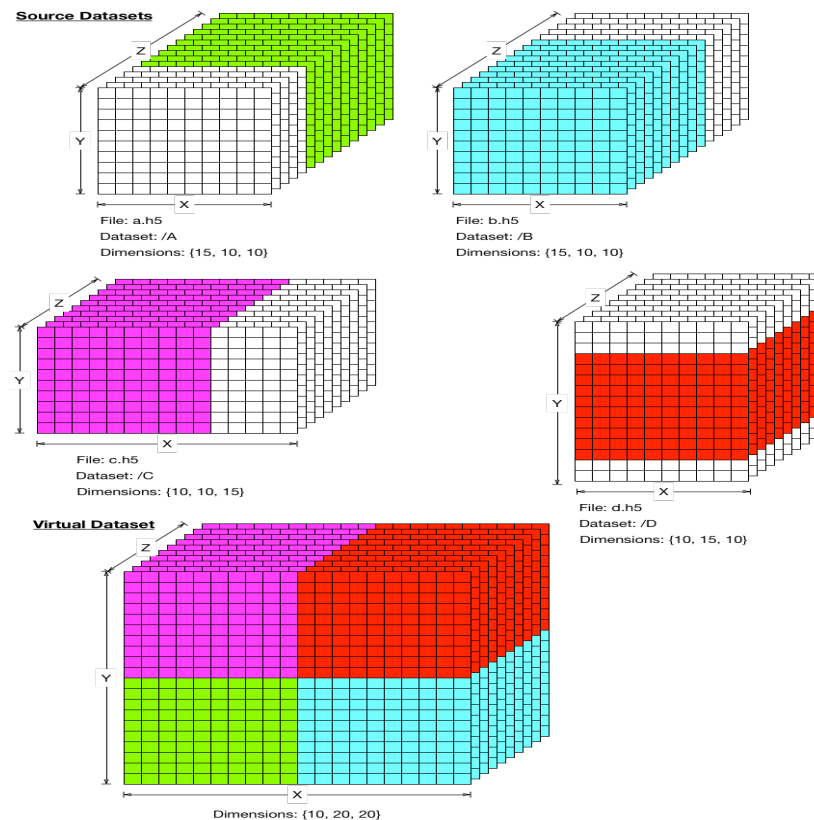


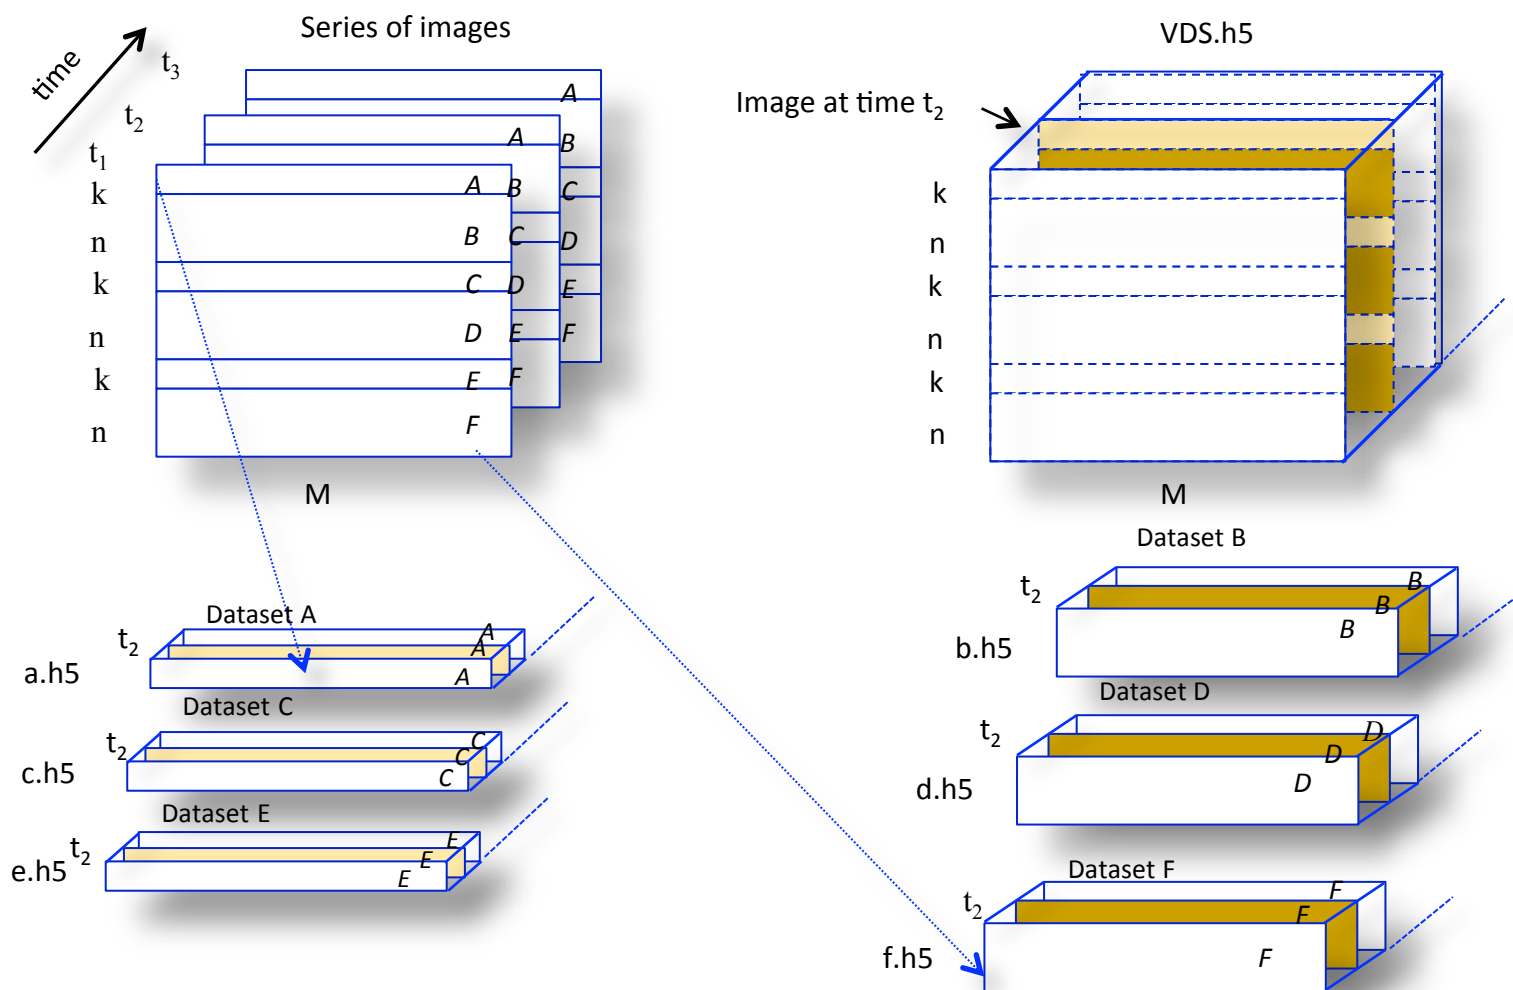
SWMR: Data access to file being written



- Released in HDF5 1.10.0
- Restricted to append-data only scenario
- SWMR doesn't work on NFS
- Files are not compatible with HDF5 1.8.* libraries
- Use h5format_convert tool
 - Converts HDF5 metadata in place
 - No raw data is rewritten

- Data stored in multiple files and datasets can be accessed via one dataset (VDS) using standard HDF5 read/write





- VDS works with SWMR
- File with VDS cannot be accessed by HDF5 1.8.* libraries
- Use h5repack tool to rewrite data (1.10.0-patch1)



New requirements and features?

- Tell us your needs:
 - Multi-threaded compression filters
 - H5DRead_chunk function
 - Full SWMR implementation
 - Performance
 - Backward/forward compatibility
- Other requests?



HDF5 Roadmap for 2016

- May 31 -HDF5 1.10.0-patch1
 - H5repack, Windows builds, Fortran issues on HPC systems
- Mid-summer HDF5 1.10.1
 - Some internal issues found for 1.10.0
- December
 - HPC features that didn't make it into 1.10.0 release



Questions?

Thank you!