Scalable HDF5

Quincey Koziol

ATPESC
August 2, 2019
koziol@lbl.gov

Why HDF5?

- **Have you ever asked yourself:**
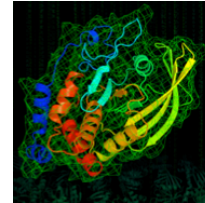
  – How will I deal with one-file-per-process in the exascale era?

  – Do I need to be an "MPI and Lustre pro" to do my research?

  – Where is my checkpoint file?

- **HDF5 *hides I/O complexity* so you can *concentrate on your science***

  – Optimized I/O to single shared file*

* Explorations for "multi-file" HDF5 storage are underway as well.

Goal

- **Introduce you to HDF5**
  - HDF5 Overview
  - HDF5 Programming Overview
  - Intro to Scalable HDF5

WHAT IS HDF5?

What is HDF5?

- **HDF5 == Hierarchical Data Format, v5**
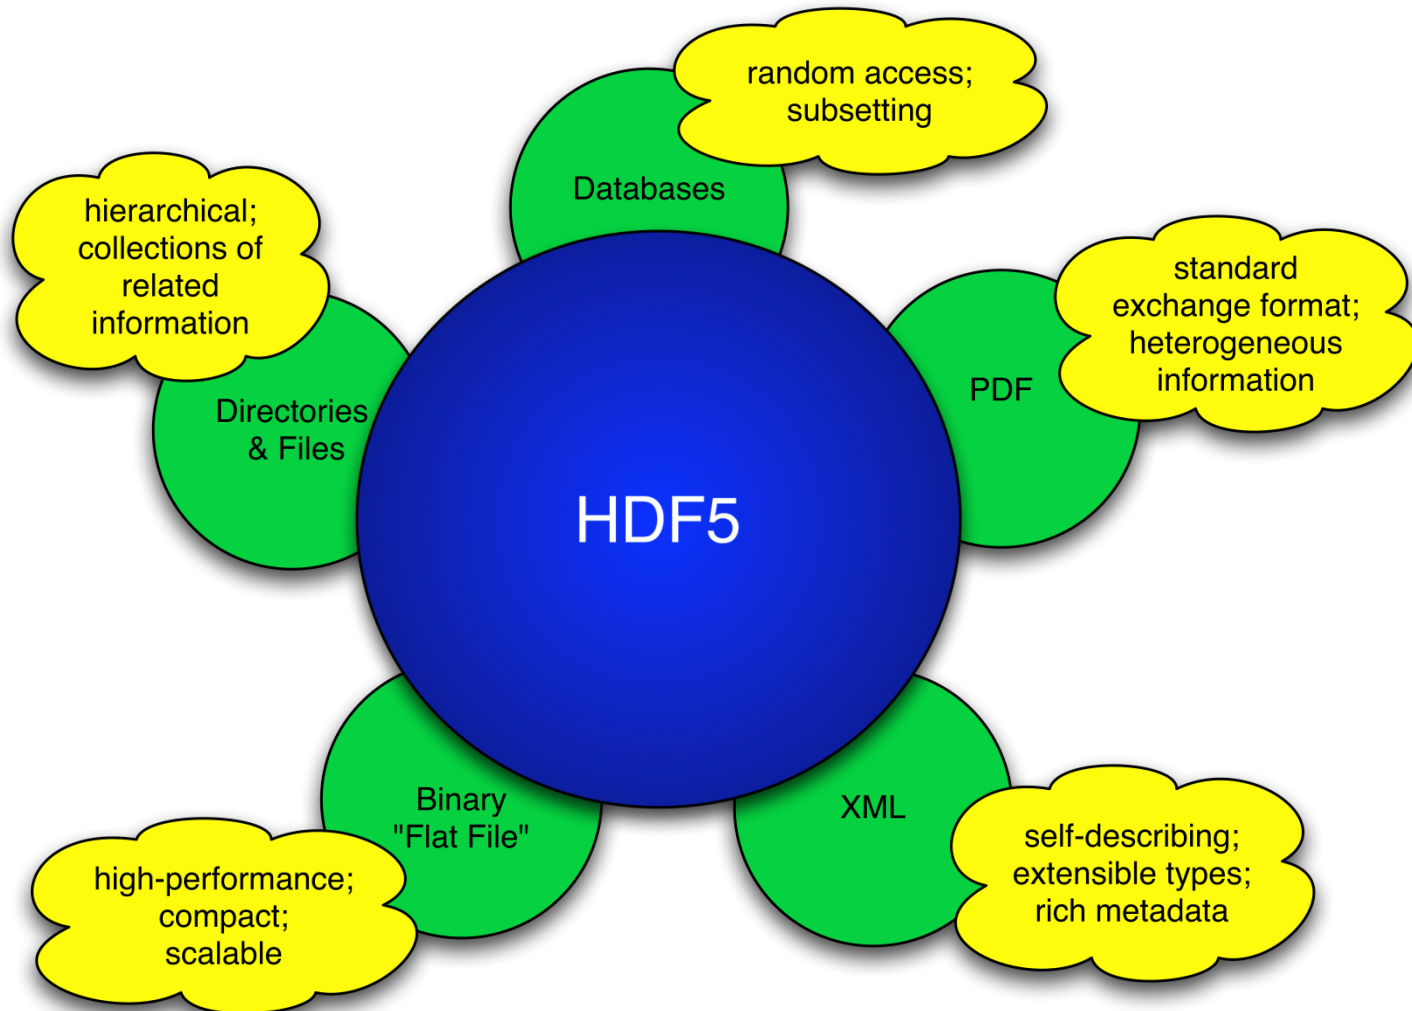
- **Open file format**
  - Designed for high volume and complex data

- **Open source software**
  - Works with data in the format

- **An extensible data model**
  - Structures for data organization and specification

HDF5 is like …

random access; subsetting

Databases

hierarchical; collections of related information

standard exchange format; heterogeneous information

Directories & Files

PDF

HDF5

Binary "Flat File"

XML

high-performance; compact; scalable

self-describing; extensible types; rich metadata

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

HDF5 is designed …

- **for high volume and/or complex data**

- **for every size and type of system (portable)**

- **for flexible, efficient storage and I/O**

- **to enable applications to evolve in their use of HDF5 and to accommodate new models**

- **to support long-term data preservation**

HDF5 Ecosystem
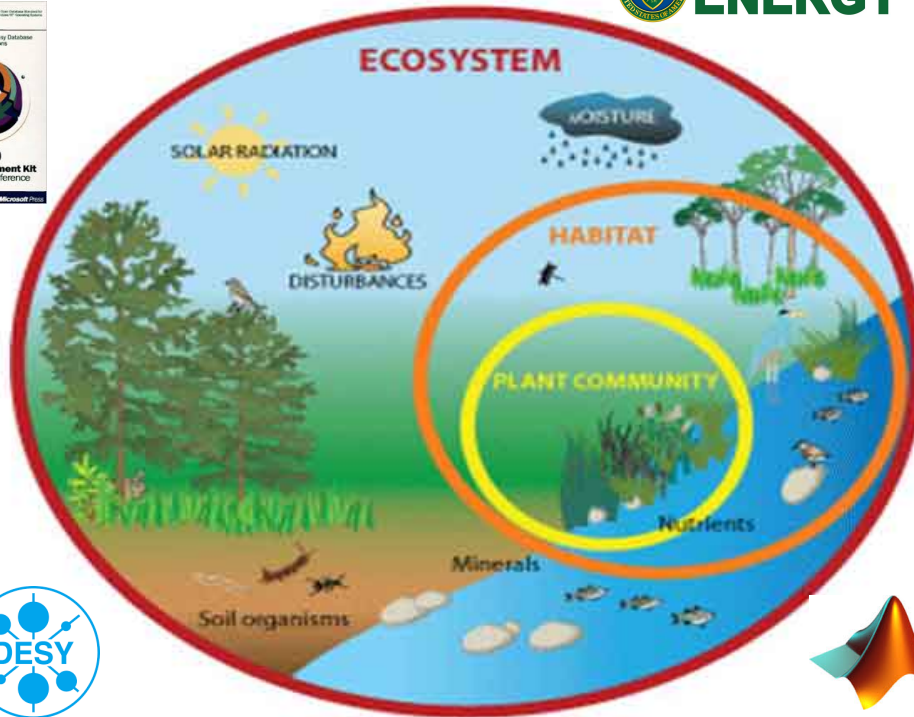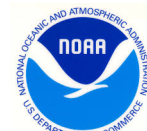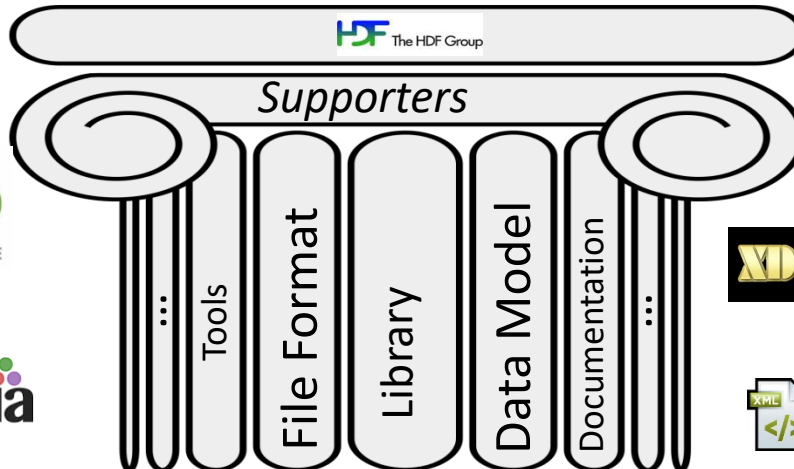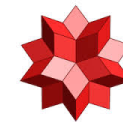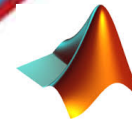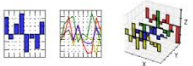
HDF5 DATA MODEL

HDF5 Data Model

Dataset

Group

Attribute

Link

Datatype

Dataspace
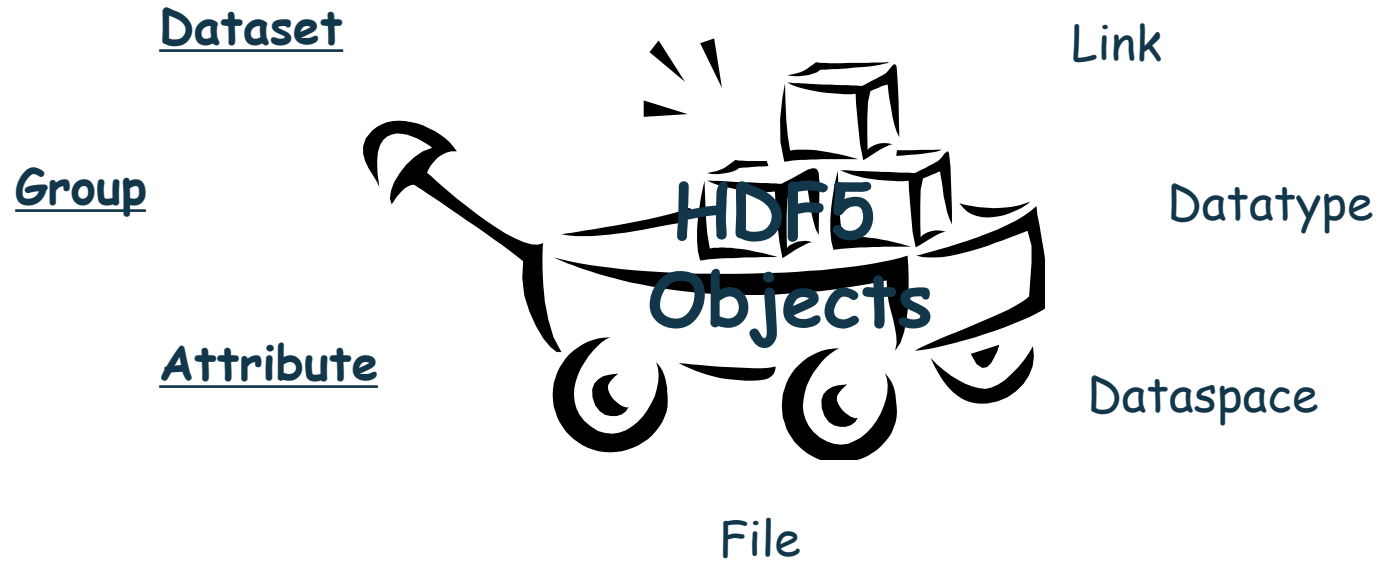
File

HDF5 Objects

HDF5 File

An HDF5 file is a **container** that holds data objects.
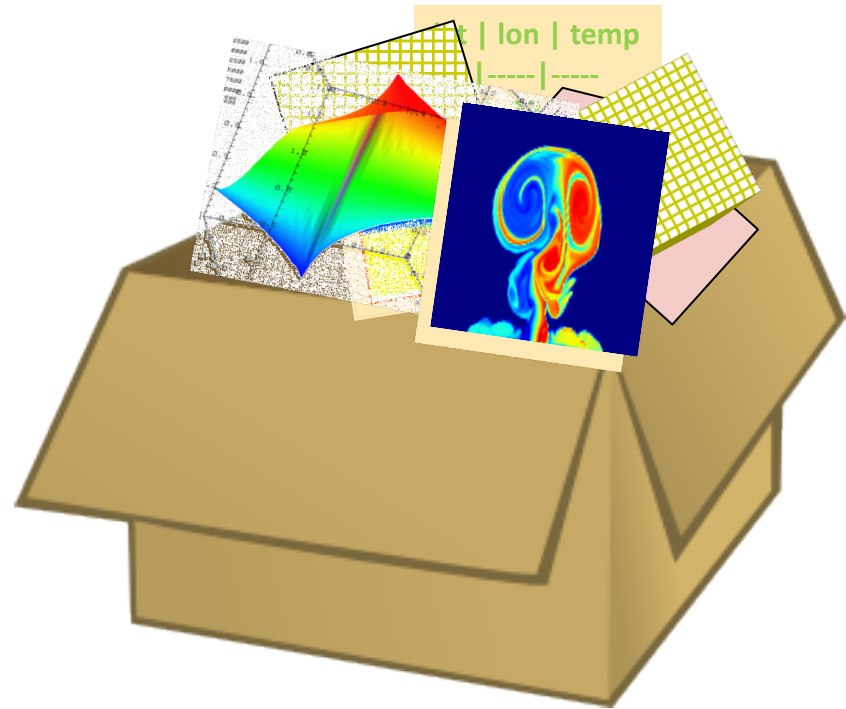
HDF5 Dataset

HDF5 Datatype

Integer: 32-bit, LE

HDF5 Dataspace

Rank	Dimensions
3	Dim[0] = 7
	Dim[1] = 4
	Dim[2] = 5

Specifications for single data element and array dimensions
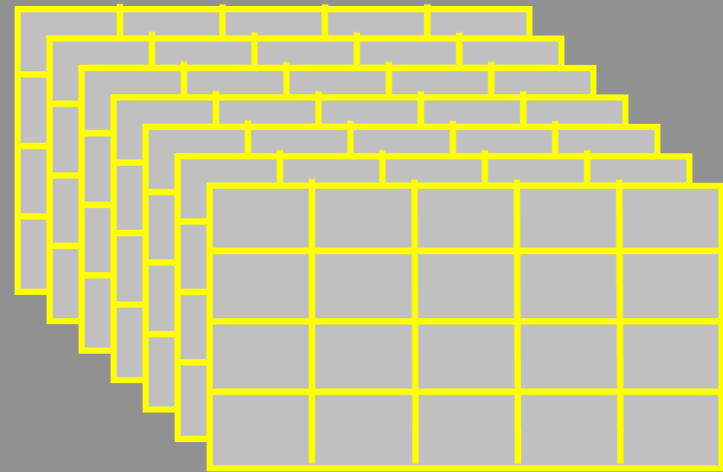
Multi-dimensional array of identically typed data elements

- HDF5 datasets **organize and contain** data elements.
  - HDF5 datatype describes individual data elements.
  - HDF5 dataspace describes the logical layout of the data elements.

HDF5 Dataspace

- **Describes the logical layout of the elements in an HDF5 dataset**
  - NULL
    - No elements (i.e., the empty/null set)
  - Scalar
    - Single element (a "point", without dimensionality)
  - Simple (*most common*)
    - A rectangular array:
      - Rank = number of dimensions
      - Dimension sizes = number of elements in each dimension
      - Maximum number of elements in each dimension
        » may be fixed or unlimited

HDF5 Dataspace – Two Roles

Spatial information for Datasets and Attributes

– Rank and dimension sizes
– Permanent part of object definition
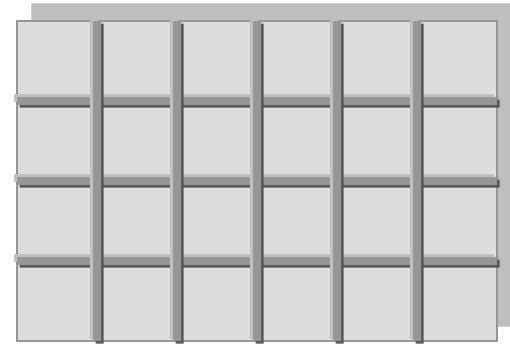
Rank = 2

Dimensions = 4 x 6

Partial I/O: Dataspace and selection describe application's data buffer and data elements participating in I/O

Rank = 1

Dimension = 10

Start = 5

Count = 3

HDF5 Datatypes

- **Describe individual data elements in an HDF5 dataset**

- Wide range of datatypes supported:

  – Integer

  – Float

  – Enum

  – Array (similar to matrix)

  – Variable-length sequence (e.g., strings, C++ vectors)

  – Compound (similar to C structs)

  – User-defined (e.g. 12-bit integer, 16-bit float, etc.)

  – More …

HDF5 Dataset
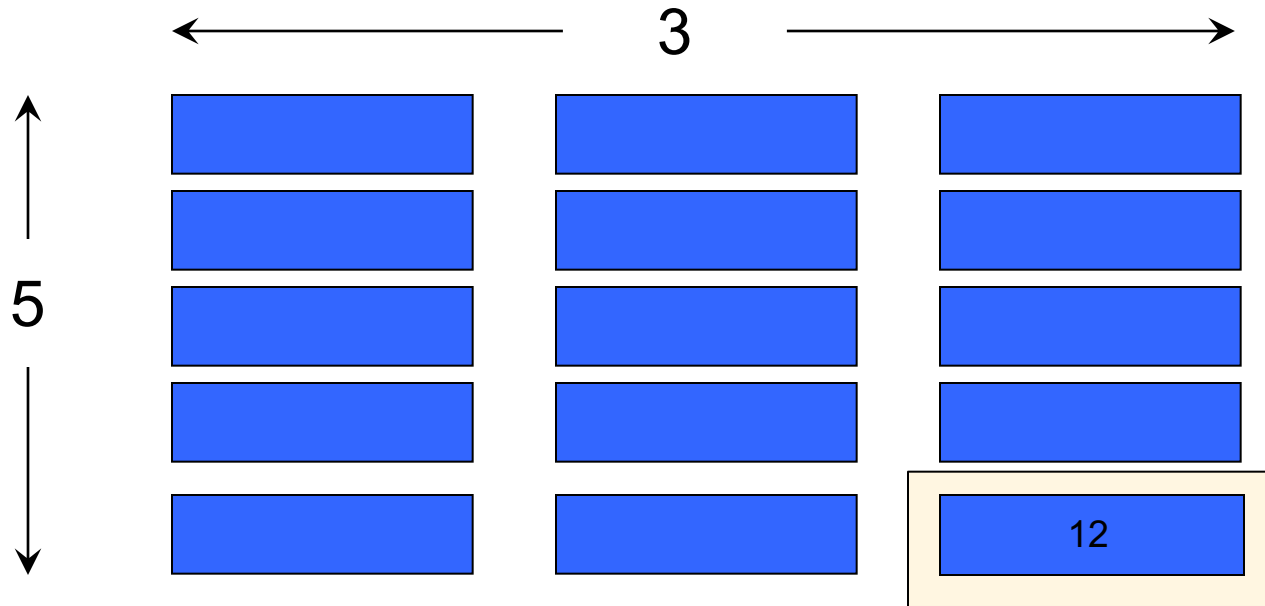
Datatype: 32-bit Integer

Dataspace: Rank = 2
              Dimensions = 3 x 5

Note that this is declared in C as: "array[5][3]" and as "array(3)(5)" in FORTRAN

HDF5 Dataset with Compound Datatype

Compound Datatype:

uint16   char   int32   2x3x2 array of float32
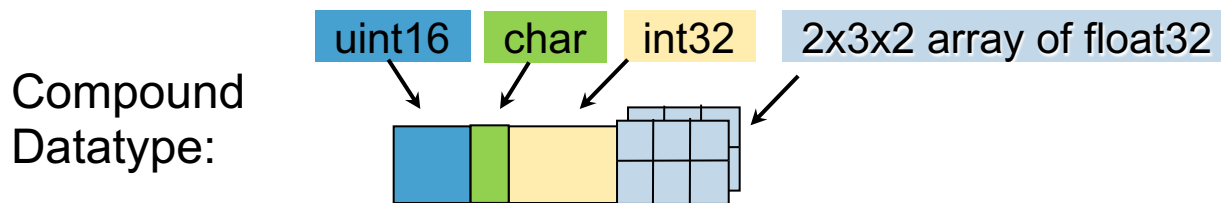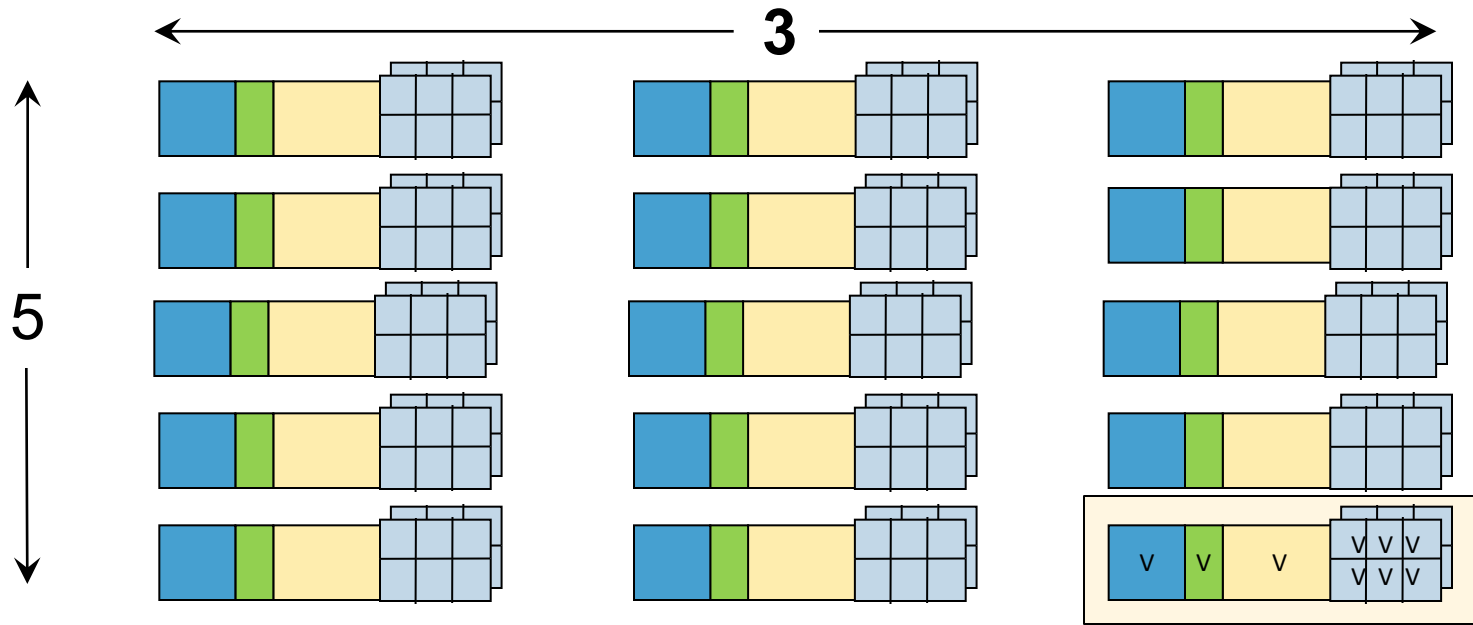
Dataspace: Rank = 2
              Dimensions = 3 x 5

How are data elements stored? (1/2)

Buffer in memory   Data in the file

Contiguous (default)
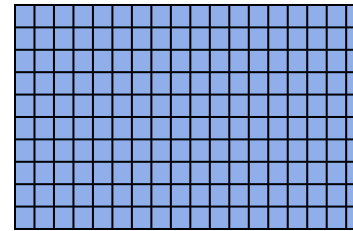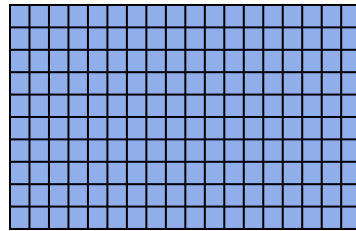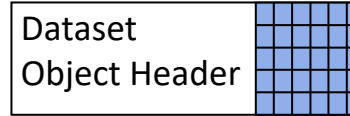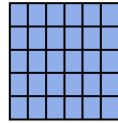
Data elements stored physically adjacent to each other
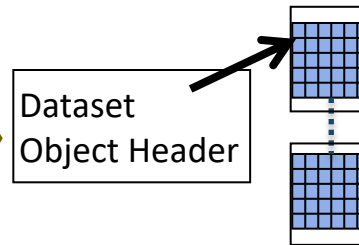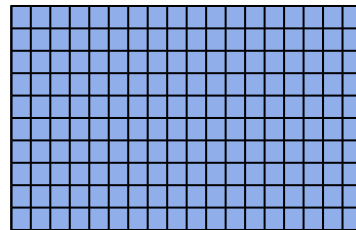
Chunked

Better access time for subsets; extendible; can have filters (e.g. compression)

Chunked & Compressed

Improves storage efficiency, transmission speed

U.S. DEPARTMENT OF **ENERGY** | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory
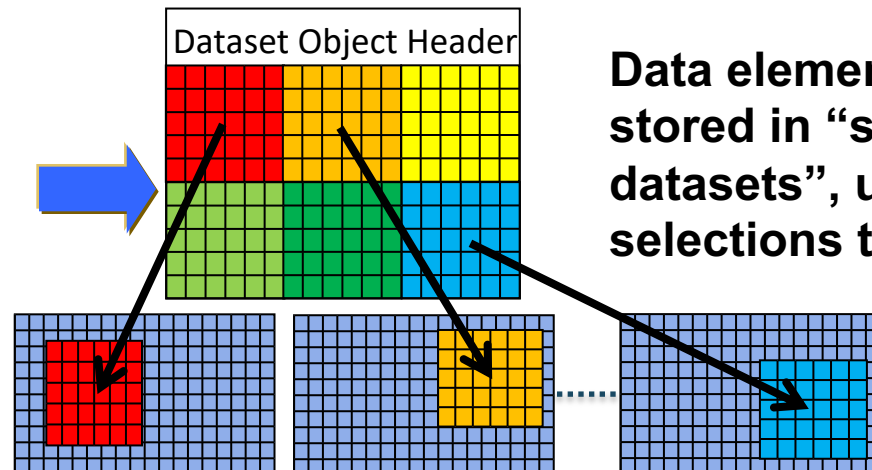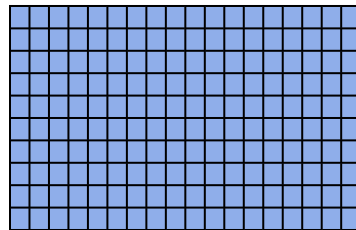
How are data elements stored? (2/2)

HDF5 Attributes

- **Attributes "decorate" HDF5 objects**

- **Typically contain *user* metadata**

- **Similar to Key-Values:**

  – Have a unique <u>name</u> (for that object) and a <u>value</u>

- **Analogous to a dataset**

  – "Value" is an array described by a datatype and a dataspace

  – But do not support partial I/O operations; nor can they be compressed or extended

HDF5 Groups and Links

HDF5 groups and links **organize** data objects.

Experiment Notes:
Serial Number: 99378920
Date: 3/13/09
Configuration: Standard 3

/

Every HDF5 file has a root group

Parameters
10;100;1000

Viz

SimOut

lat	lon	temp
12	23	3.1
15	24	4.2
17	21	3.6

Timestep
36,000
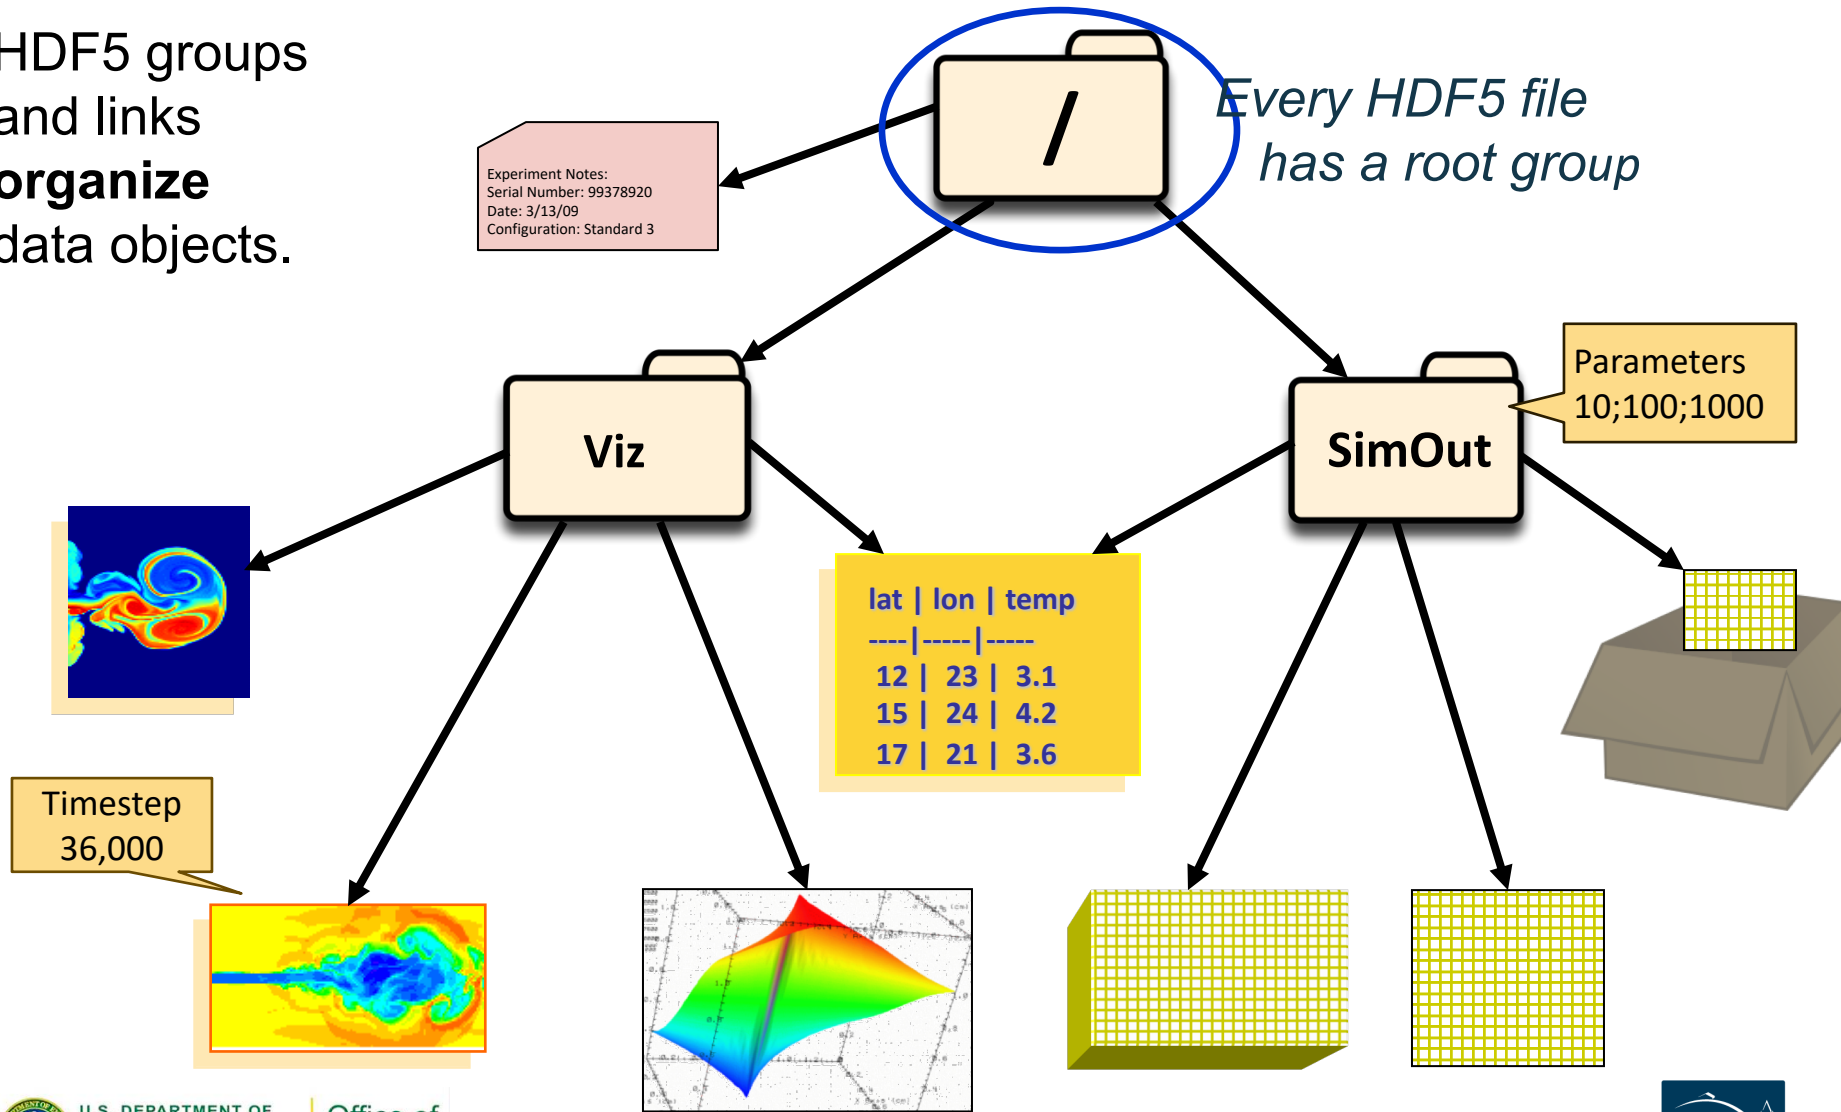
HDF5 SOFTWARE

HDF5 Home Page

HDF5 home page: **http://hdfgroup.org/solutions/HDF5/**
- Latest release: HDF5 1.10.5 (1.10.6 coming Fall 2019)

HDF5 source code:
- Written in C
- Includes optional C++, FORTRAN, and Java APIs, and High Level APIs
- Also command-line utilities (h5dump, h5repack, h5diff, …) and compile scripts

HDF5 pre-built binaries:
- When possible, include C, C++, FORTRAN , and High Level libraries.
  - Check ./lib/libhdf5.settings file for installed options
- Built with and require the SZIP and ZLIB external libraries

Useful HDF5 Tools For New Users

h5dump:
   Tool to "dump" or display contents of HDF5 files

h5cc, h5c++, h5fc:
   Scripts to compile applications

HDFView:
   Java browser to view HDF5 files
   http://www.hdfgroup.org/hdf-java-html/hdfview/

HDF5 Examples in C, C++, FORTRAN, Java, Python, Matlab:
   https://portal.hdfgroup.org/display/HDF5/HDF5+Examples

HDF5 PROGRAMMING MODEL AND API

HDF5 Software Layers & Storage

ATPESC 2019 - http://bit.ly/ATPESC-2019-HDF5
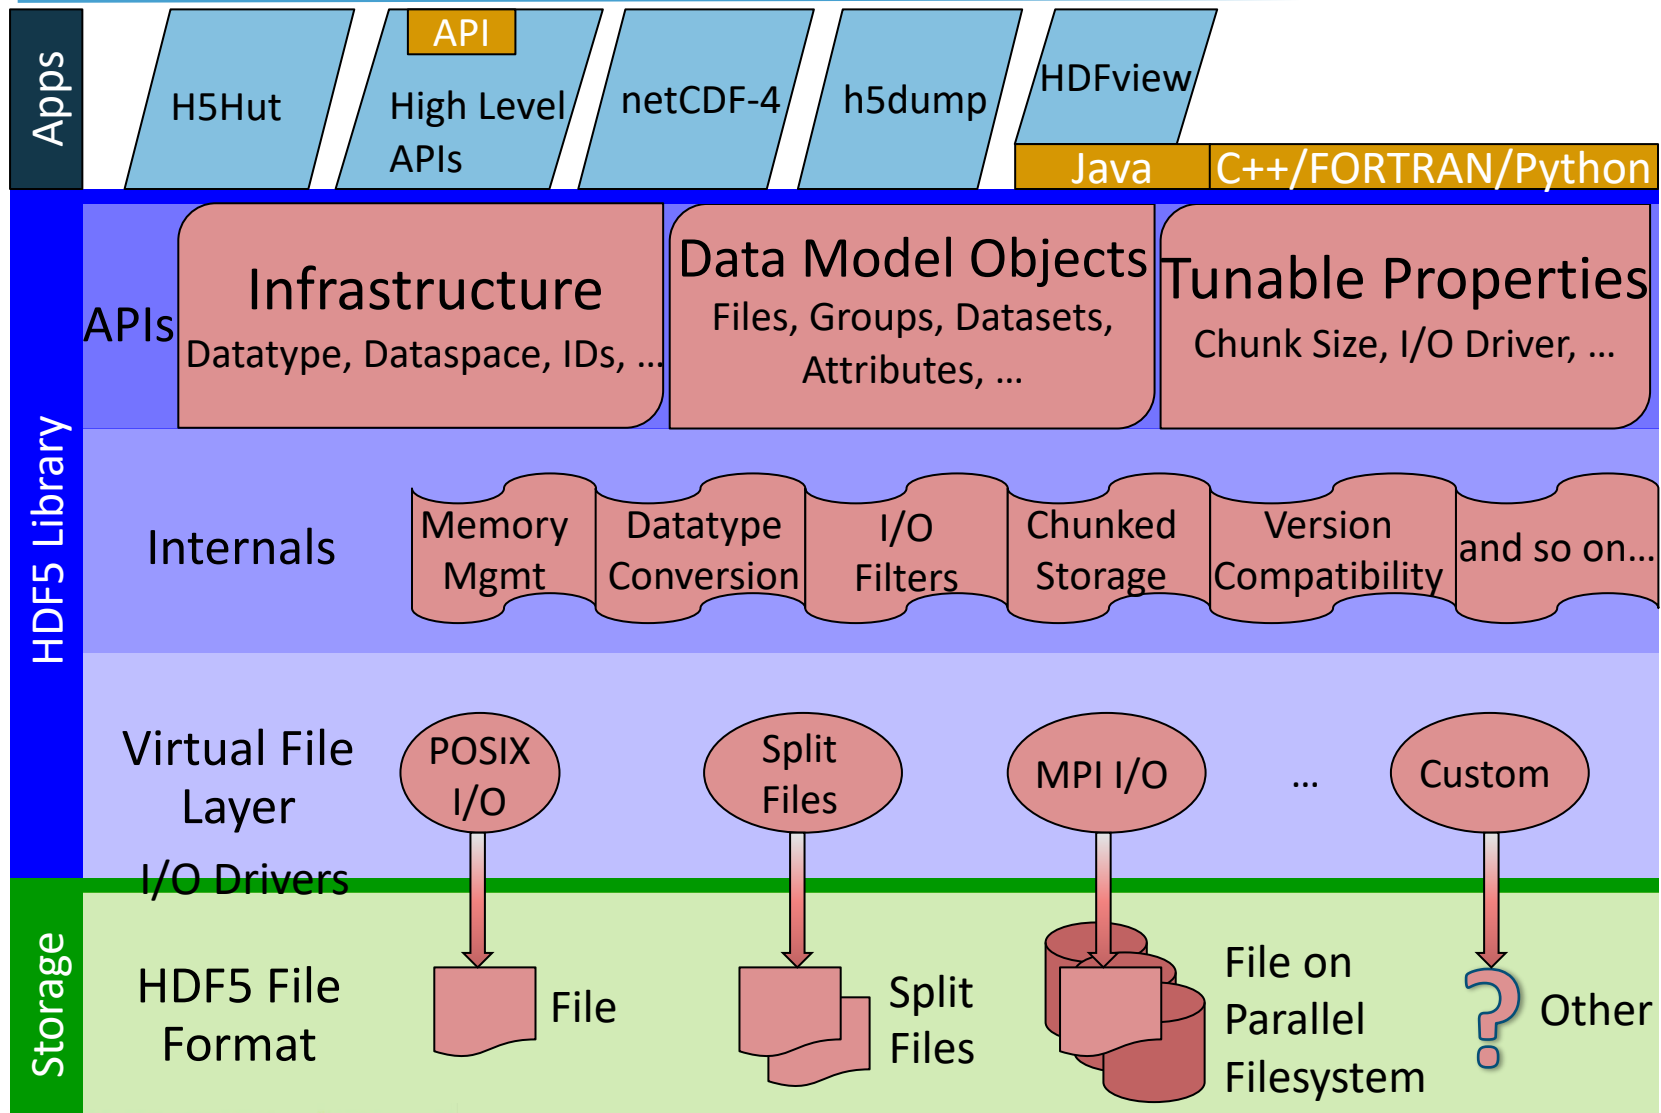
The General HDF5 API

- **C, FORTRAN, Java, C++, and .NET bindings**
- **IDL, MATLAB, Python (H5Py, PyTables)**
- **C routines begin with prefix: H5?**

  ? is a character corresponding to the type of object the function acts on

Example Functions:

H5D: **D**ataset interface   *e.g.,* **H5Dread**

H5F: **F**ile interface   *e.g.,* **H5Fopen**

H5S: data**S**pace interface *e.g.,* **H5Sclose**

The HDF5 API

- **For flexibility, the API is extensive**
  - ✓ 300+ functions

- **This can be daunting… but there is hope**
  - ✓ A few functions can do a lot
  - ✓ Start simple
  - ✓ Build up knowledge as more features are needed

Victorinox
Swiss Army
Cybertool 34

General Programming Paradigm

- **Object is opened or created**
- **Object is accessed, possibly many times**
- **Object is closed**

- **Properties of object are <u>optionally</u> defined**
  - ✓ Creation properties (e.g., use chunking storage)
  - ✓ Access properties

Basic Functions

H5Fcreate (H5Fopen)   *create (open) File*

  H5Screate_simple/H5Screate *create dataSpace*

    H5Dcreate (H5Dopen) *create (open) Dataset*

      H5Dread, H5Dwrite   *access Dataset*

    H5Dclose   *close Dataset*

  H5Sclose   *close dataSpace*

H5Fclose   *close File*

Other Common Functions

DataSpaces: H5Sselect_hyperslab (Partial I/O)
                       H5Sselect_elements (Partial I/O)
                       H5Dget_space

DataTypes: H5Tcreate, H5Tcommit, H5Tclose
                       H5Dget_type, H5Tequal, H5Tget_native_type

Groups: H5Gcreate / H5Gopen, H5Gclose

Attributes: H5Acreate / H5Aopen_name,
                       H5Aread, H5Awrite, H5Aclose

Property lists: H5Pcreate
                       H5Pset_chunk, H5Pset_deflate,
                       H5Pclose

PARALLEL HDF5

Terminology

- **DATA – "problem-size" data, e.g., large arrays**

- **METADATA – is an overloaded term**

- **In this presentation:**

    Metadata "=" HDF5 metadata

    – For each piece of application metadata, there are many associated pieces of HDF5 metadata

    – There are also other sources of HDF5 metadata

        • Chunk indices, heaps to store group links and indices to look them up, object headers, etc.

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Why Parallel HDF5?

- **Take advantage of high-performance parallel I/O while reducing complexity**
  - Use a well-defined high-level I/O layer instead of POSIX or MPI-IO
  - Use only a single or a few shared files
    - "Friends don't let friends use file-per-process!" ☺

- **Code, Performance and Data Portability**
  - Rely on HDF5 to optimize for underlying storage system

What We'll Cover

- Parallel vs. Serial HDF5

- Implementation Layers

- HDF5 files (= composites of data & metadata) in a parallel file system

- PHDF5 I/O modes: collective vs. independent

- Data and Metadata I/O

What We Won't Cover

- Consistency semantics
- Virtual Object Layer (VOL)
- Automatic tuning
- Single Writer / Multiple-Reader (SWMR)
- Virtual Datasets (VDS)
- Asynchronous I/O
- Independent Metadata Modification
- …

Come see me this evening or after the presentation!

(MPI-)Parallel vs. Serial HDF5

- PHDF5 allows multiple processes in an MPI application to perform I/O to a single HDF5 file

- Uses a standard parallel I/O interface (MPI-IO)

- Portable to different platforms

- PHDF5 files <u>ARE</u> HDF5 files conforming to the HDF5 File Format Specification

- The PHDF5 API consists of:
  - The standard HDF5 API
  - A few extra knobs and calls
  - A parallel "etiquette"

Standard HDF5 "Skeleton"

H5Fcreate (H5Fopen)   create (open) File

      H5Screate_simple/H5Screate   create dataSpace

            H5Dcreate (H5Dopen)   create (open) Dataset

               H5Dread, H5Dwrite   access Dataset

            H5Dclose   close Dataset

      H5Sclose   close dataSpace

H5Fclose   close File
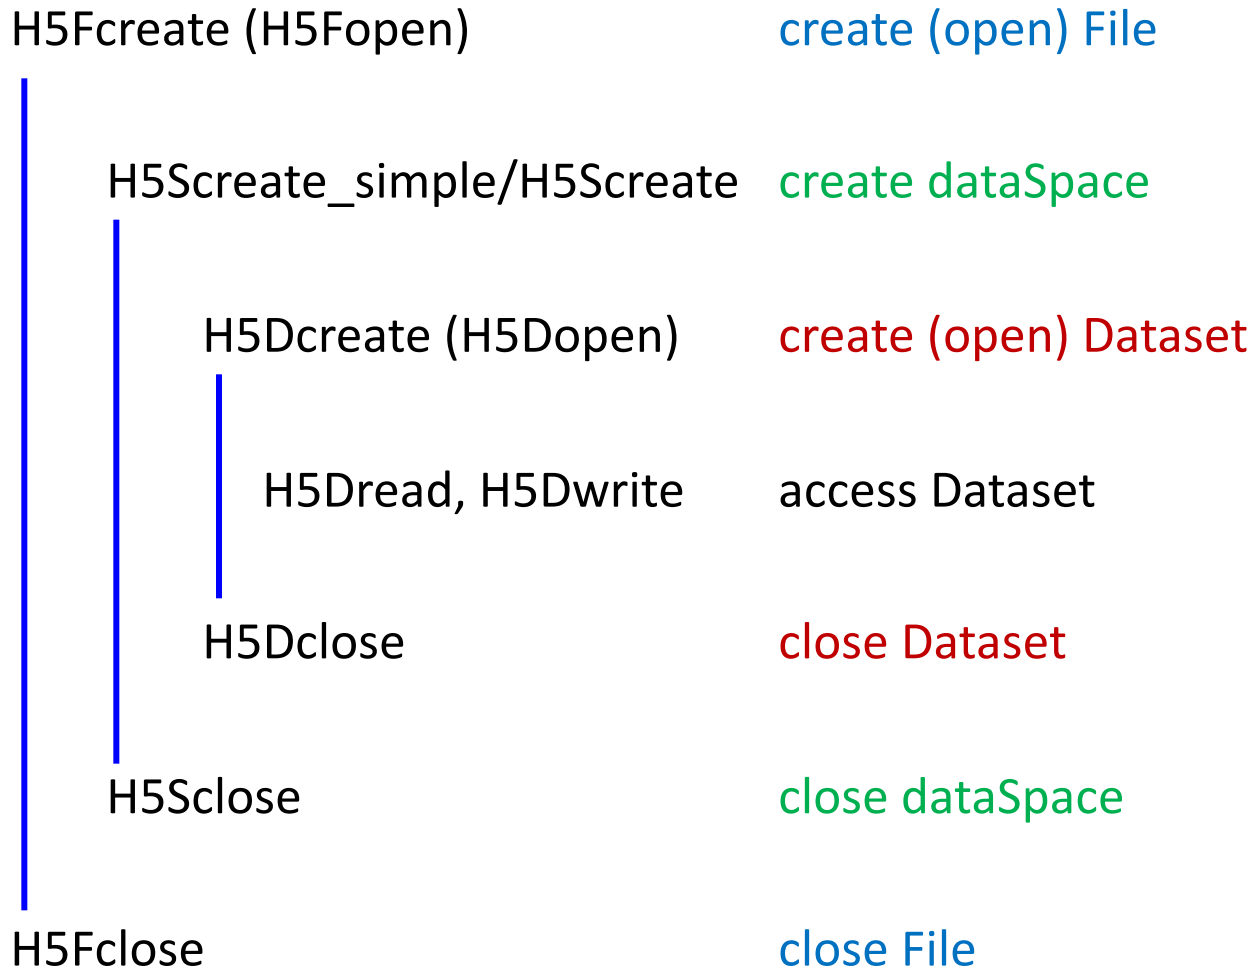
PHDF5 Implementation Layers
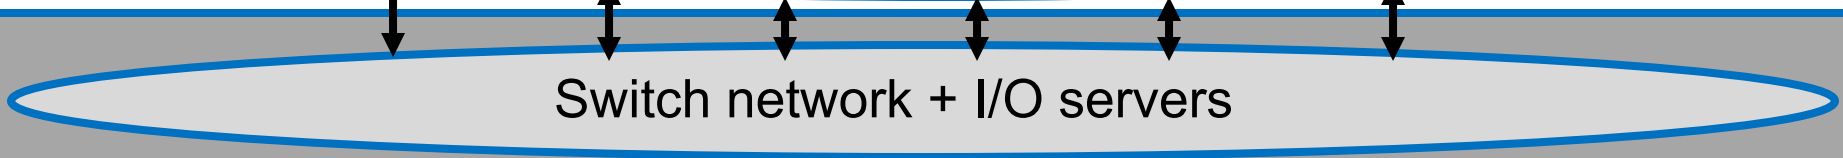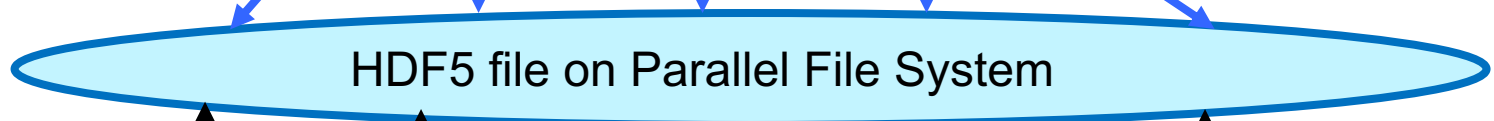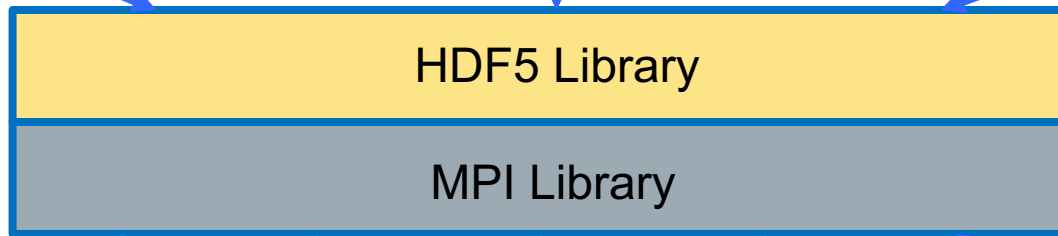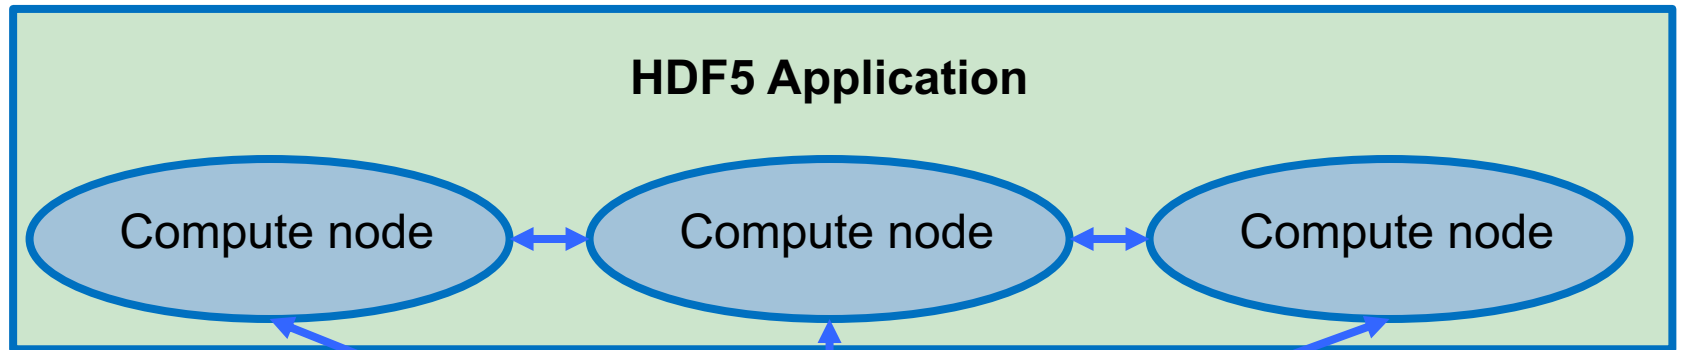
Example of a PHDF5 C Program

Starting with a simple serial HDF5 program:

```
fapl_id = H5Pcreate(H5P_FILE_ACCESS);

file_id = H5Fcreate(FNAME, …, fapl_id);
space_id = H5Screate_simple(…);
dset_id = H5Dcreate(file_id, DNAME, H5T_NATIVE_INT,
                           space_id, …);


status = H5Dwrite(dset_id, H5T_NATIVE_INT, …, xf_id…);
```

Example of a PHDF5 C Program

A parallel HDF5 program has a few extra calls:

```
MPI_Init(&argc, &argv);

fapl_id = H5Pcreate(H5P_FILE_ACCESS);
H5Pset_fapl_mpio(fapl_id, comm, info);
file_id = H5Fcreate(FNAME, …, fapl_id);
space_id = H5Screate_simple(…);
dset_id = H5Dcreate(file_id, DNAME, H5T_NATIVE_INT,
                             space_id, …);
xf_id = H5Pcreate(H5P_DATASET_XFER);
H5Pset_dxpl_mpio(xf_id, H5FD_MPIO_COLLECTIVE);
status = H5Dwrite(dset_id, H5T_NATIVE_INT, …, xf_id…);

MPI_Finalize();
```

PHDF5 Etiquette

- **PHDF5 opens a shared file with an MPI communicator**
  - Returns a file ID (as usual)
  - All future access to the file via that file ID

- **Different files can be opened via different communicators**

- **<u>All</u> processes must participate in <u>collective</u> PHDF5 APIs**

- **<u>All</u> HDF5 APIs that modify the HDF5 namespace and structural metadata are collective!**
  - File ops., group structure, dataset dimensions, object life-cycle, etc.

https://support.hdfgroup.org/HDF5/doc/RM/CollectiveCalls.html

Parallel HDF5 Tutorial Examples

- **For more examples how to write different data patterns see:**

  https://portal.hdfgroup.org/display/HDF5/Introduction+to+Parallel+HDF5

In a Parallel File System
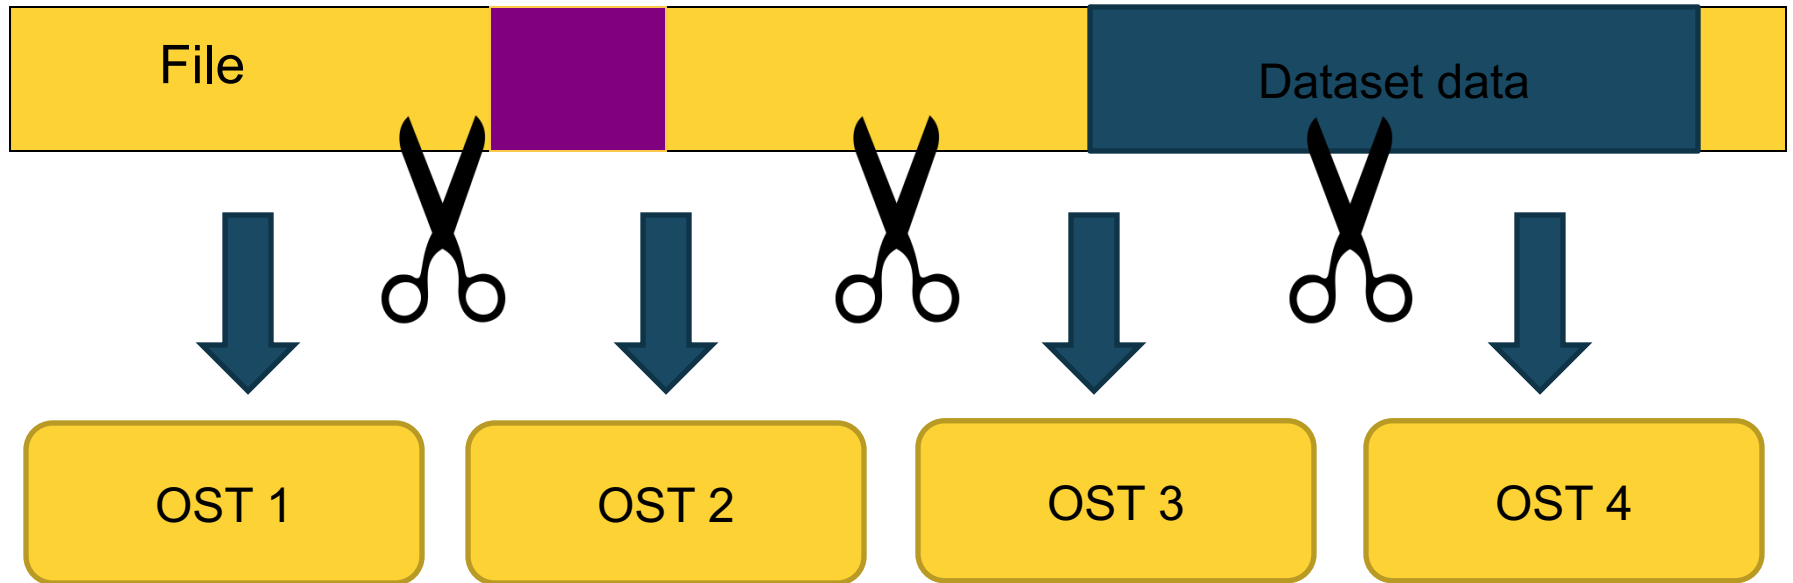
The file is striped over multiple "disks" (e.g. Lustre OSTs) depending on the stripe size and stripe count with which the file was created.

And it gets worse before it gets better…

HDF5 - Contiguous Storage

- **HDF5 Object header separate from dataset data**
- **Data stored in one contiguous block in HDF5 file**

HDF5 - Chunked Storage

- **Dataset data is divided into equal-sized blocks (chunks).**

- **Each chunk is stored separately in the HDF5 file, located by a "chunk index".**
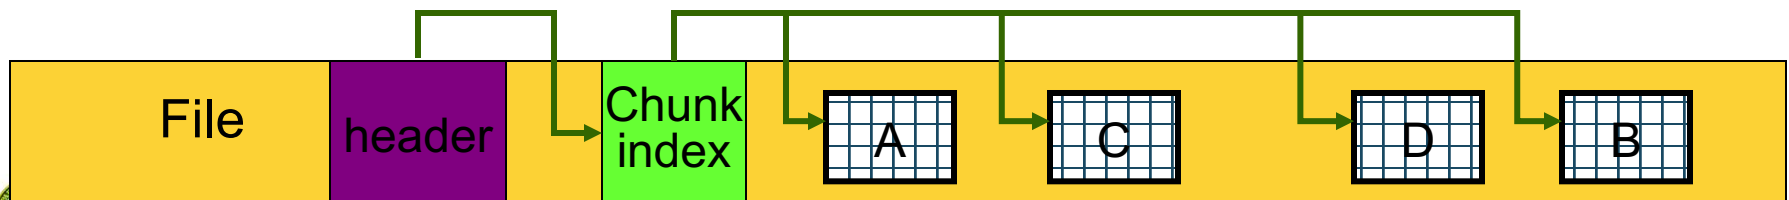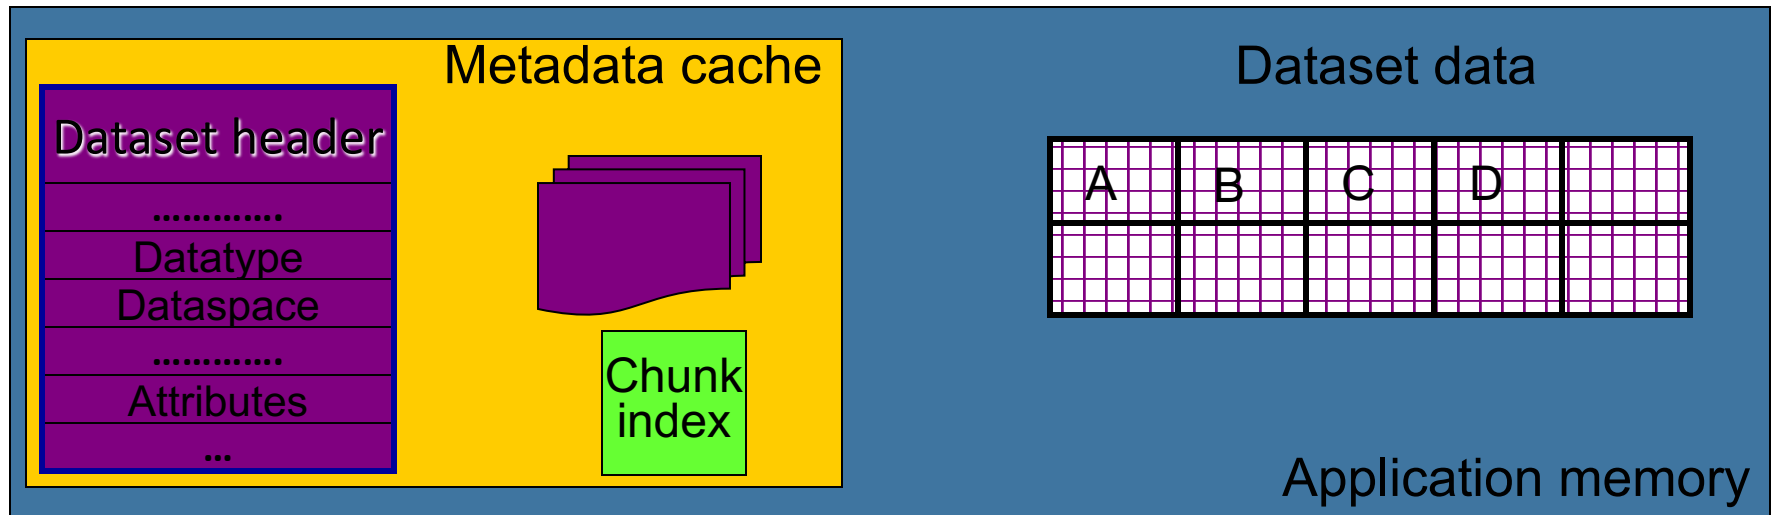
HDF5 In a Parallel File System
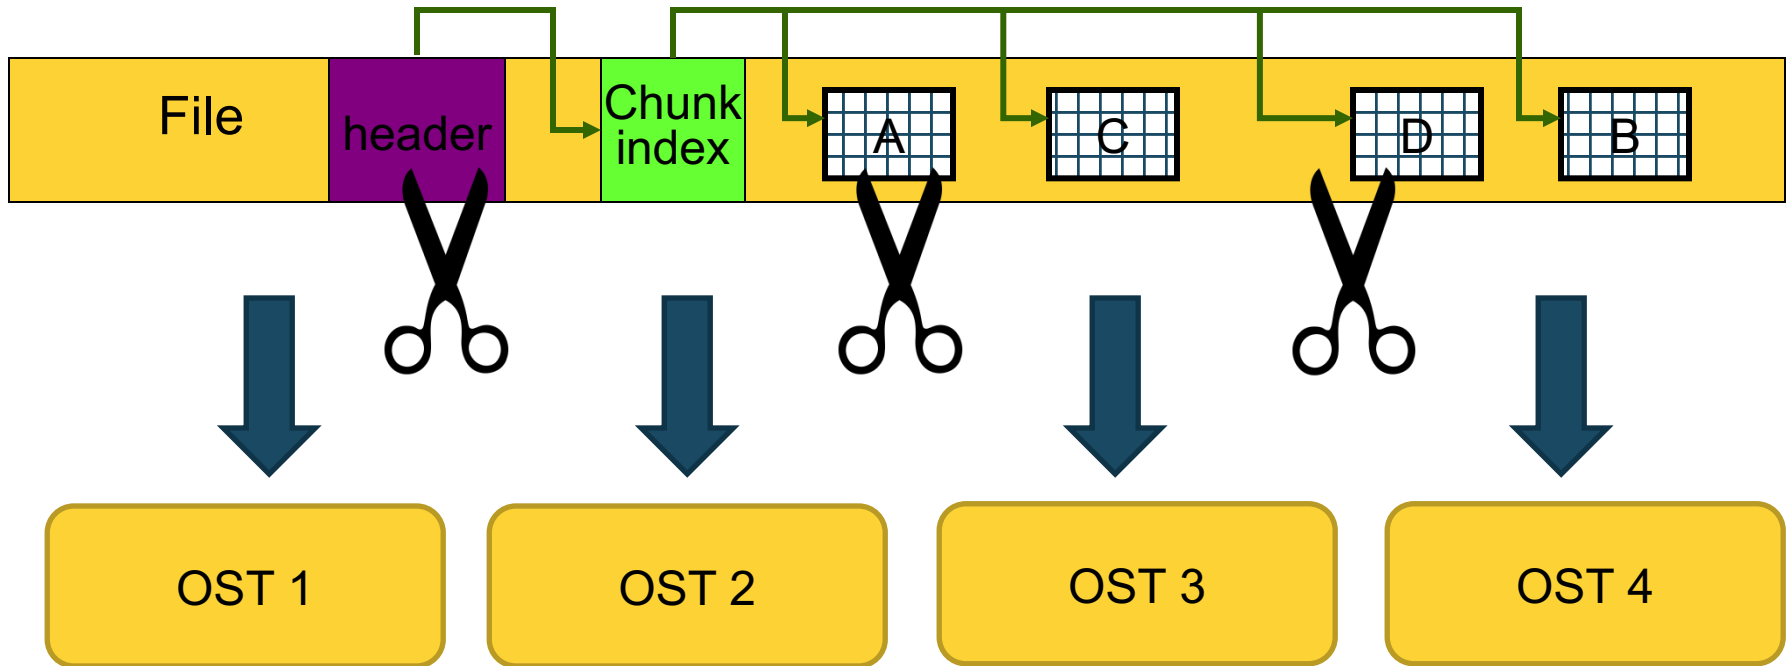
The file is striped over multiple OSTs depending on the stripe size and stripe count with which the file was created.

Collective vs. Independent I/O
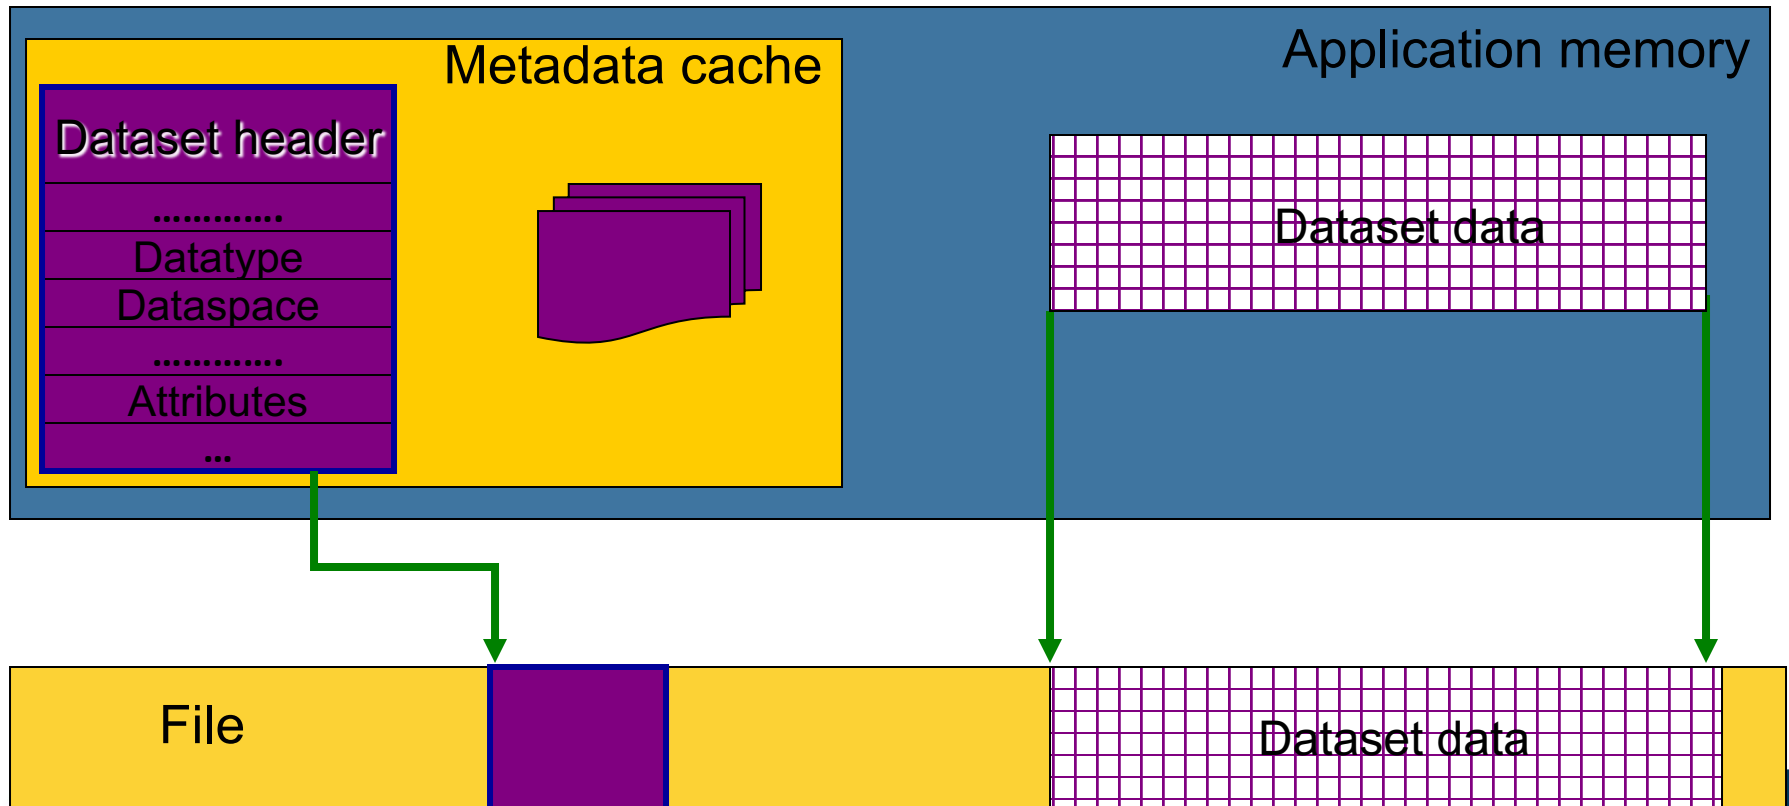
- **Collective I/O attempts to combine multiple smaller independent I/O ops into fewer larger ops.**
  - Neither mode is preferable *a priori*
- **MPI definition of collective calls:**
  - All processes of the communicator must participate in calls in the same order:

    <u>Process 1</u>   <u>Process 2</u>

    call A(); ➔ call B();   call A(); ➔ call B(); **right**

    call A(); ➔ call B();   call B(); ➔ call A(); **wrong**

  - Independent calls are not collective ☺
  - Collective calls are not necessarily synchronous, nor must they require communication
    - It could be that only internal state for the communicator changes

U.S. DEPARTMENT OF **ENERGY** | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Data and Metadata I/O

Data

- **"Problem-sized"**
- **I/O can be independent or collective**

- **Improvement targets:**
  - Avoid unnecessary I/O
  - I/O frequency
  - Layout on disk
    - Different I/O strategies for chunked layout
  - Aggregation and balancing
  - Alignment

Metadata

- **Small**
- **Reads can be independent or collective**
- **<u>All</u> modifications must be collective**

- **Improvement targets:**
  - Metadata design
  - Use the latest library version, if possible
  - Metadata cache
    - In desperate cases, take control of evictions

Don't Forget! It's a Multi-Layer Problem

Application

HDF5
(Disable truncate in H5Fclose)

MPI-IO
(Number of collective buffer nodes, Collective buffer size)

Lustre Parallel File System
(Stripe factor and Stripe size)

Storage Hardware
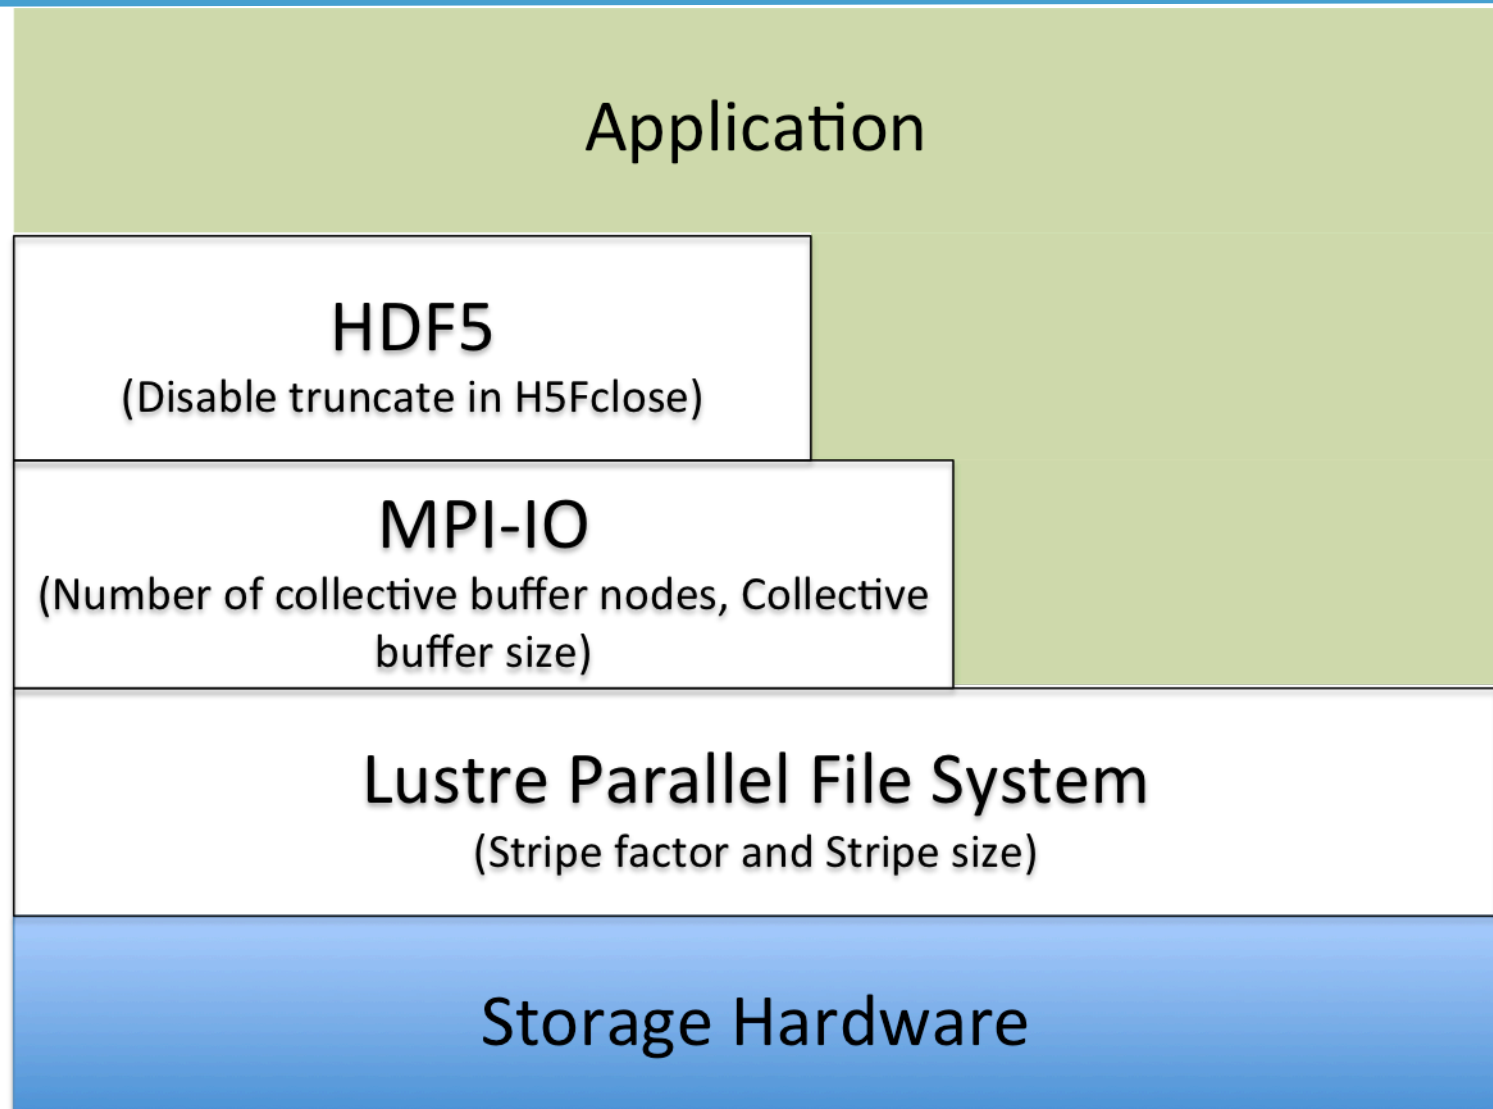
DIAGNOSTICS AND INSTRUMENTATION

A Textbook Problem

User reported:

- **Independent data transfer mode is much slower than the collective data transfer mode**

- **Data array is tall and thin: 230,000 rows by 4 columns**
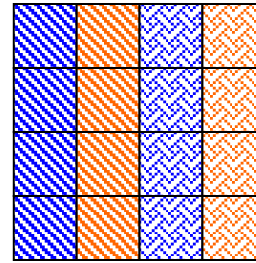
230,000 rows
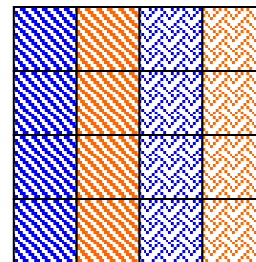
Symptoms

Writing to one dataset

- 4 MPI processes ➔ 4 columns
- Datatype is 8-byte floats (doubles)
- 4 processes x 1000 rows x 8 bytes = 32,000 bytes

```
% mpirun -np 4 ./a.out 1000
```
  ➢ Execution time: 1.783798 s.

```
% mpirun -np 4 ./a.out 2000
```
  ➢ Execution time: 3.838858 s. (linear scaling)

- **2 sec. extra for 1000 more rows = 32,000 bytes.**

    16KB/sec ➔ Way too slow!!!

"Poor Man's Debugging"

- **Build a version of PHDF5 with**
- **`./configure --enable-debug --enable-parallel …`**
- **This allows the tracing of MPIO I/O calls in the HDF5 library such as `MPI_File_read_xx` and `MPI_File_write_xx`**
- **Don't forget to `% setenv H5FD_mpio_Debug "rw"`**

- **You'll get something like this…**

Independent and Contiguous

```
% setenv H5FD_mpio_Debug 'rw'
% mpirun -np 4 ./a.out 1000      # Indep.; contiguous.
in H5FD_mpio_write   mpi_off=0      size_i=96
in H5FD_mpio_write   mpi_off=0      size_i=96
in H5FD_mpio_write   mpi_off=0      size_i=96
in H5FD_mpio_write   mpi_off=0      size_i=96
in H5FD_mpio_write   mpi_off=2056   size_i=8
in H5FD_mpio_write   mpi_off=2048   size_i=8
in H5FD_mpio_write   mpi_off=2072   size_i=8
in H5FD_mpio_write   mpi_off=2064   size_i=8
in H5FD_mpio_write   mpi_off=2088   size_i=8
in H5FD_mpio_write   mpi_off=2080   size_i=8
…
```
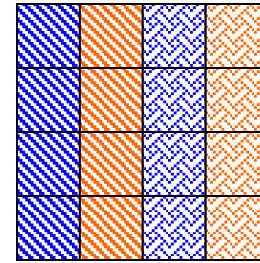
- **A total of 4000 of these 8 bytes writes == 32,000 bytes.**

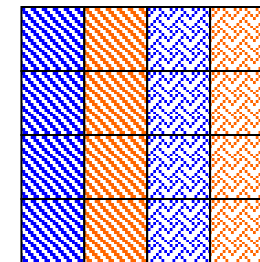# Plenty of Independent and Small I/Os

Diagnosis:

- **Each process writes one element of one row, skips to next row, writes one element, and so on.**

- **Each process issues 230,000 writes of 8 bytes each.**

230,000 rows

Chunked by Column

```
% setenv H5FD_mpio_Debug 'rw'
% mpirun -np 4 ./a.out 1000      # Indep., Chunked by column.
in H5FD_mpio_write   mpi_off=0          size_i=96
in H5FD_mpio_write   mpi_off=0          size_i=96
in H5FD_mpio_write   mpi_off=0          size_i=96     Metadata
in H5FD_mpio_write   mpi_off=0          size_i=96
in H5FD_mpio_write   mpi_off=3688        size_i=8000
in H5FD_mpio_write   mpi_off=11688       size_i=8000
in H5FD_mpio_write   mpi_off=27688       size_i=8000   Dataset elements
in H5FD_mpio_write   mpi_off=19688       size_i=8000
in H5FD_mpio_write   mpi_off=96          size_i=40
in H5FD_mpio_write   mpi_off=136         size_i=544
in H5FD_mpio_write   mpi_off=680         size_i=120    Metadata
in H5FD_mpio_write   mpi_off=800         size_i=272
…
```
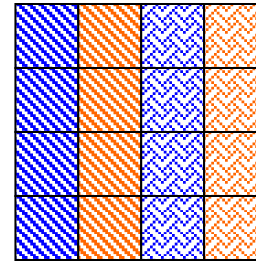
- **Execution time: 0.011599 s.**

Use Collective Mode or Chunked Storage

Remedy:

- **Collective I/O will combine many small independent calls into fewer but bigger calls**

- **Chunks of columns speeds up too**

230,000 rows
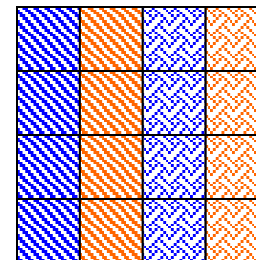
Collective vs. Independent Write
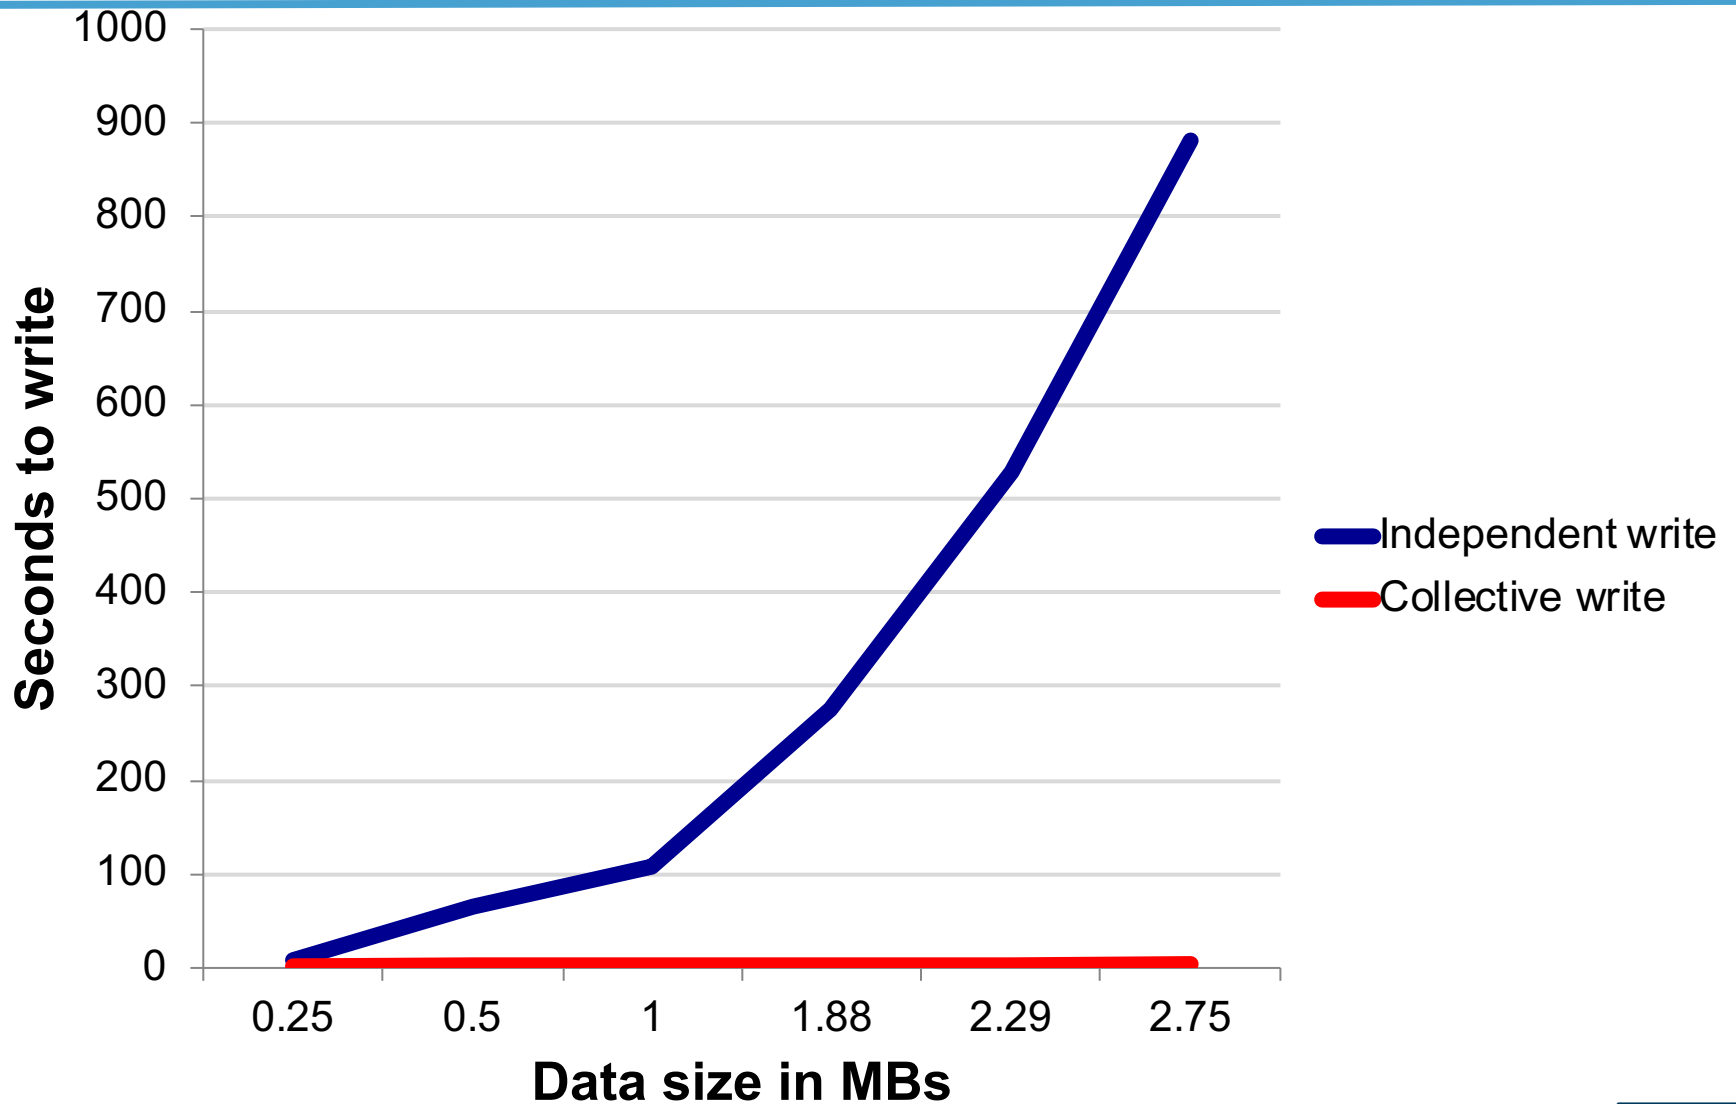
Hands-On Examples – Dataset I/O

- **h5par_ex1a.c**
  - Independent I/O to contiguous datasets

- **h5par_ex1b.c**
  - Collective I/O to contiguous datasets

- **h5par_ex1c.c**
  - Independent I/O to chunked datasets

- **Try modifying h5par_ex1c.c to use collective I/O**

Hands-On Examples – Metadat I/O

- **h5par_ex2a.c**
  - Independent group creation – Hangs!

- **h5par_ex2b.c**
  - Collective group and dataset creation

- **Look at the difference for the two example codes**

Back Into the Real World…

- **Two kinds of tools:**
  - I/O benchmarks for measuring a system's I/O capabilities
  - I/O profilers for characterizing applications' I/O behavior

- **Two examples:**
  - h5perf (in the HDF5 source code distro)
  - [Darshan](#) (from Argonne National Laboratory)

- **Profilers have to compromise between**
  - A lot of detail ➔ large trace files and overhead
  - Aggregation ➔ loss of detail, but low overhead

I/O Patterns
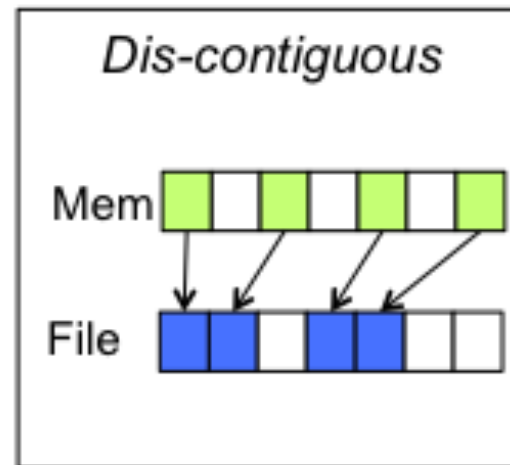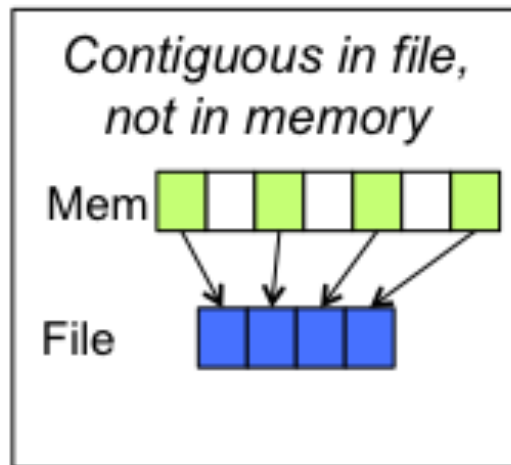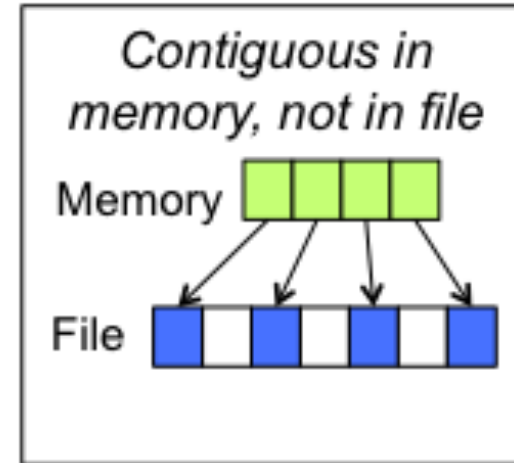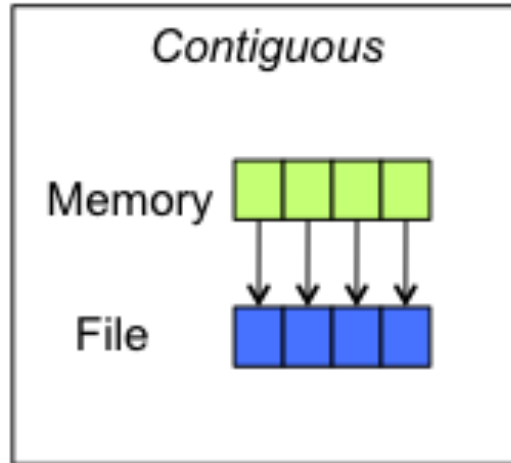
h5perf(_serial)

- **Measures performance of a filesystem for different I/O patterns and APIs**
- **Three File I/O APIs for the price of one!**
  - POSIX I/O (open/write/read/close…)
  - MPI-I/O (MPI_File_{open,write,read,close})
  - HDF5 (H5Fopen/H5Dwrite/H5Dread/H5Fclose)
- **An indication of I/O speed ranges and HDF5 overheads**
- **Expectation management…**

A Serial Run

h5perf_serial, 3 iterations, 1 GB dataset, 1 MB transfer buffer, HDF5 dataset contiguous storage, HDF5 SVN trunk, NCSA BW

A Parallel Run

**h5perf, 3 MPI processes, 3 iterations, 3 GB dataset (total),
1 GB per process, 1 GB transfer buffer,
HDF5 dataset contiguous storage, HDF5 SVN trunk, NCSA BW**
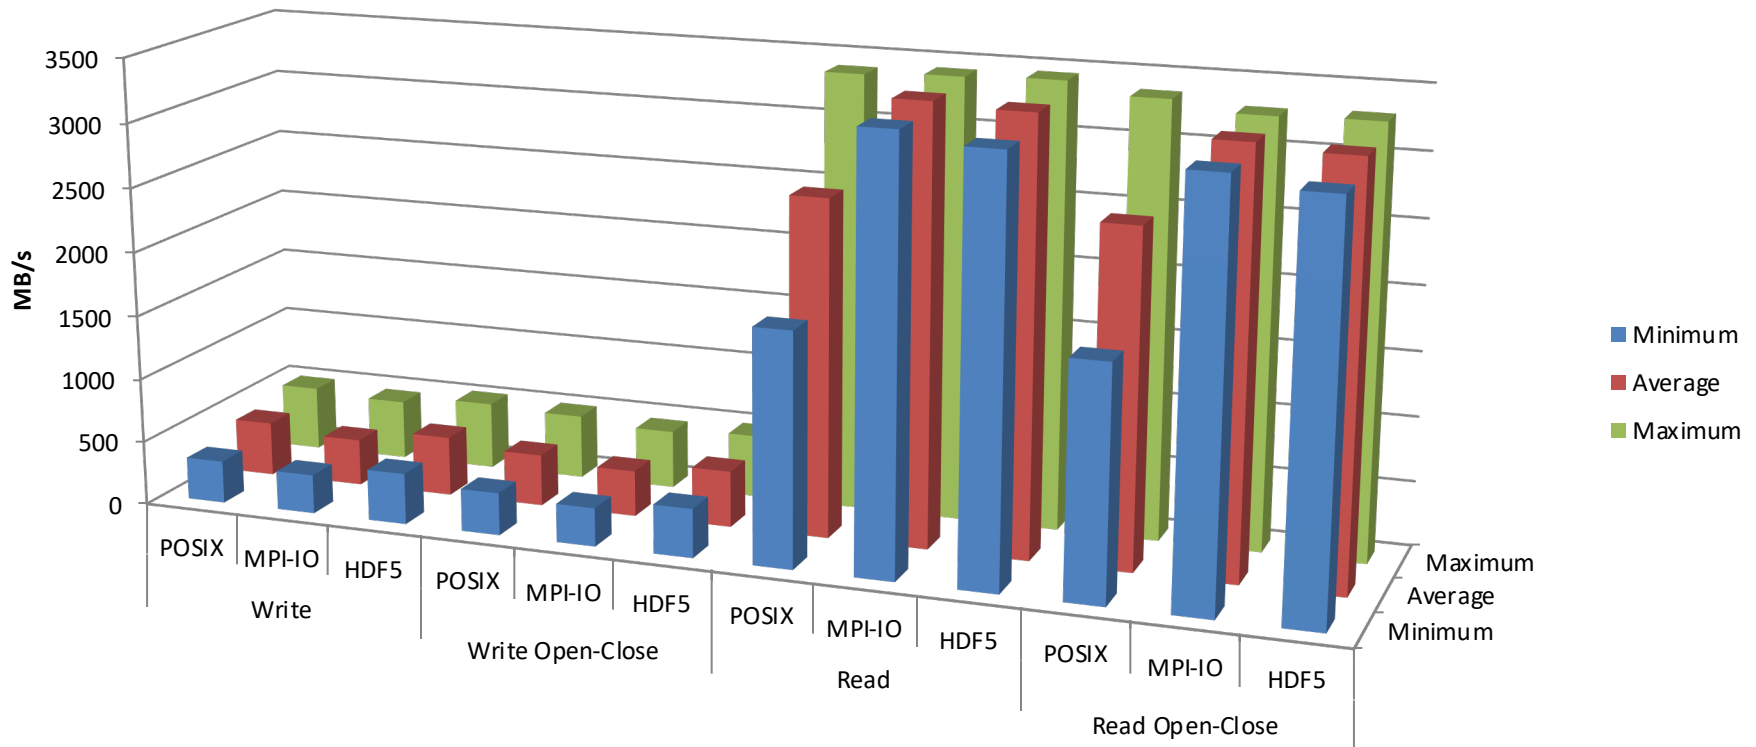
Darshan (ANL)

- **Design goals:**
  - Transparent integration with user environment
  - Negligible impact on application performance
- **Provides aggregate figures for:**
  - Operation counts (POSIX, MPI-IO, HDF5, PnetCDF)
  - Datatypes and hint usage
  - Access patterns: alignments, sequentiality, access size
  - Cumulative I/O time, intervals of I/O activity
- **Does not provide I/O behavior over time**
- **Excellent starting point, maybe not your final stop**

Darshan Sample Output
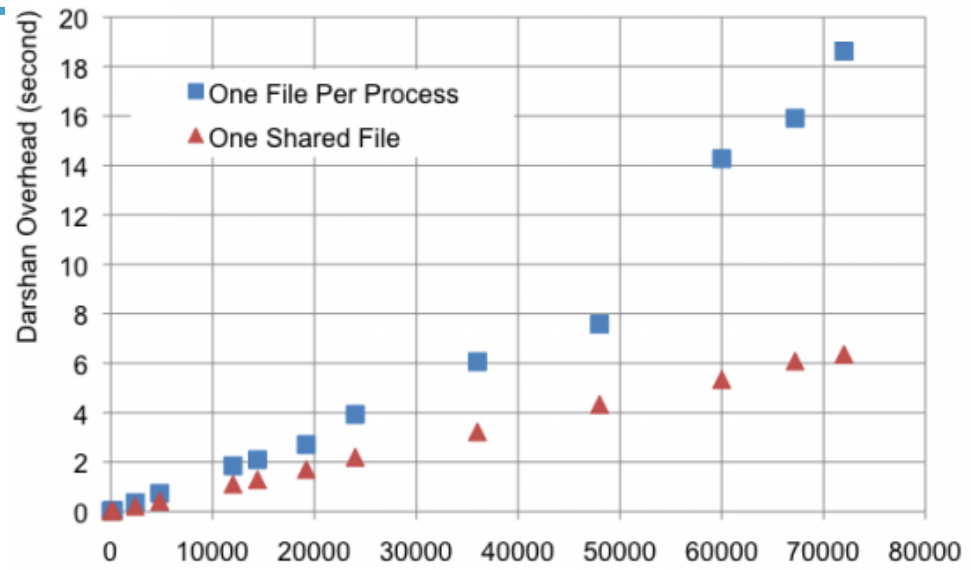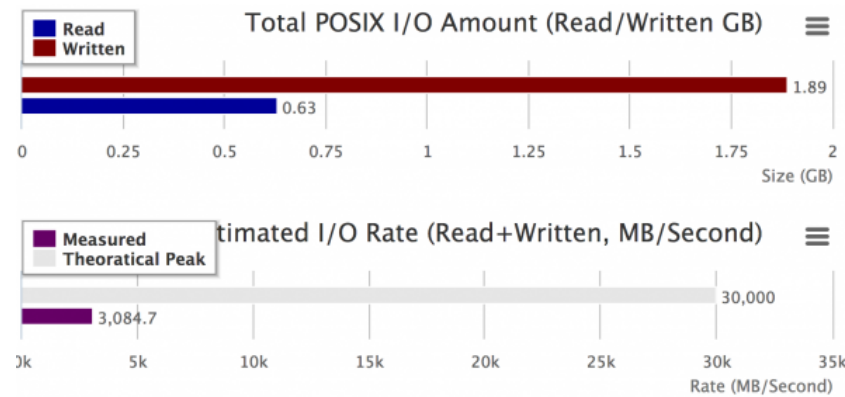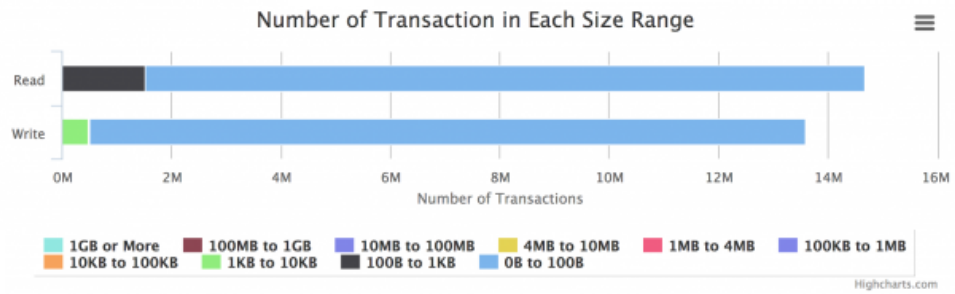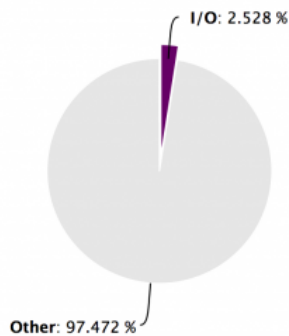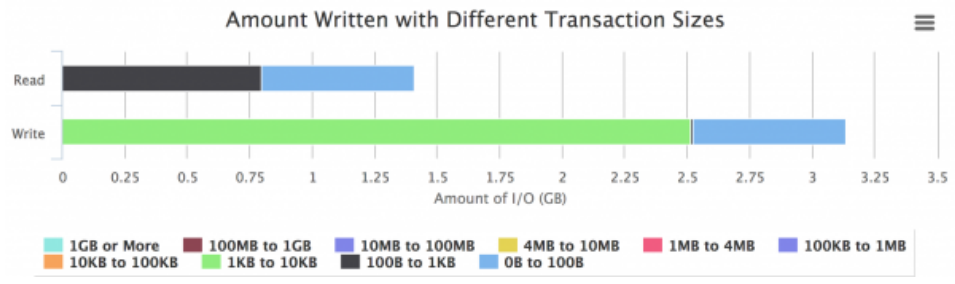
Source: NERSC

EXAMPLES

Standard Questions

- **What I/O layers are involved and how much control do I have over them?**

- **Which ones do I tackle in which order?**
  - Are there any low-hanging fruit?

- **What's my baseline (for each layer) and what are my metrics?**

- **Which tool(s) will give me the information I need?**

- **When do I stop?**

- **New information ➔ New answers (maybe) : Need to keep an open mind!**

Reference:

[Trillion Particles, 120,000 cores, and 350 TBs:](#)
[Lessons Learned from a Hero I/O Run on Hopper](#),
By Suren Byna (LBNL) et al., 2015.

Examples

VPIC

Layers

Application

HDF5
(Disable truncate in H5Fclose)

MPI-IO
(Number of collective buffer nodes, Collective buffer size)

Lustre Parallel File System
(Stripe factor and Stripe size)

Storage Hardware

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory
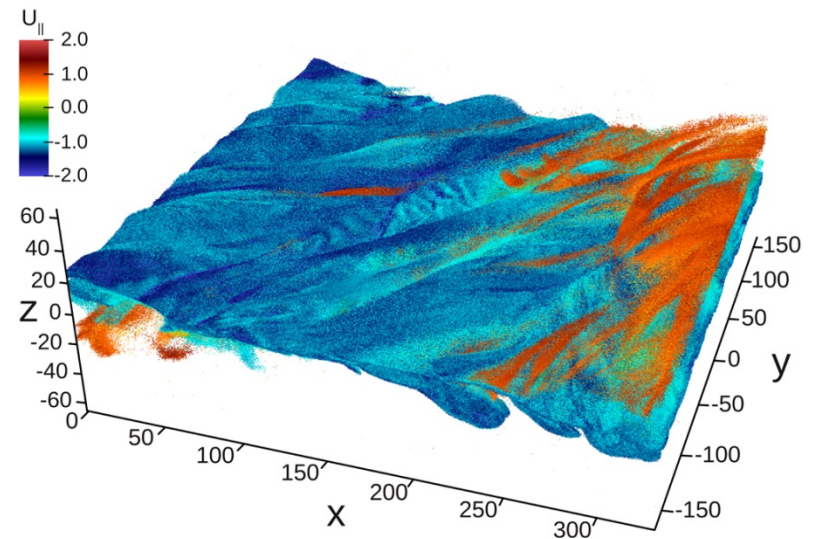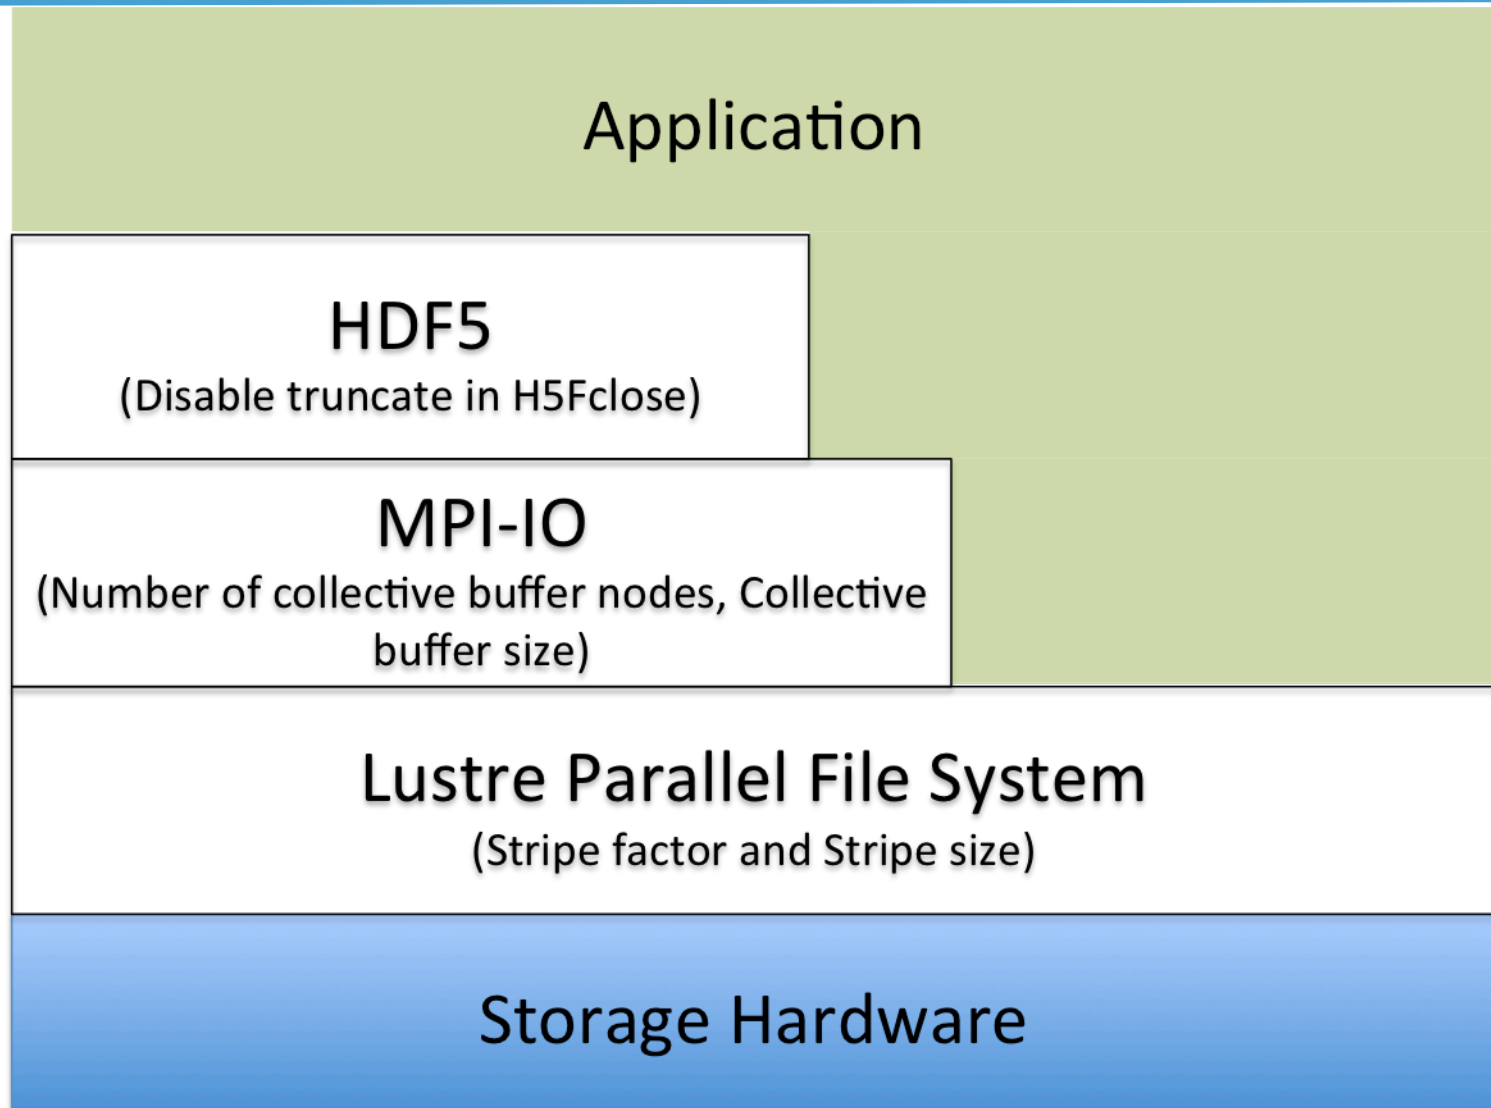
"Application I/O Structure"

- **Total control over all layers**
- **Challenge: large output files**
- **Metric: write speed (throughput)**
- **Computationally intensive ➜ Need an I/O kernel**
- **H5Part multiple dataset writes**

```
h5pf = H5PartOpenFileParallel (fname, H5PART_WRITE |
                H5PART_FS_LUSTRE, MPI_COMM_WORLD);
H5PartSetStep (h5pf, step);
H5PartSetNumParticlesStrided (h5pf, np_local, 8);

H5PartWriteDataFloat32 (h5pf, "dX", Pf);
H5PartWriteDataFloat32 (h5pf, "dY", Pf+1);
H5PartWriteDataFloat32 (h5pf, "dZ", Pf+2);
H5PartWriteDataInt32   (h5pf, "i",  Pi+3);
H5PartWriteDataFloat32 (h5pf, "Ux", Pf+4);
H5PartWriteDataFloat32 (h5pf, "Uy", Pf+5);
H5PartWriteDataFloat32 (h5pf, "Uz", Pf+6);
H5PartWriteDataFloat32 (h5pf, "q", Pf+7);

H5PartCloseFile (h5pf);
```

- **"Game plan":**
  - MPI-IO / Lustre tuning
    - Low hanging fruit (relatively)
    - Pair MPI aggregators with Lustre OSTs
    - Match MPI-IO buffer sizes and Lustre stripe size
  - Worry about HDF5 (H5Part)

I/O Aggregation

Application MPI Domains 20,000 6 threads per domain

MPI Aggregators

Lustre OSTs
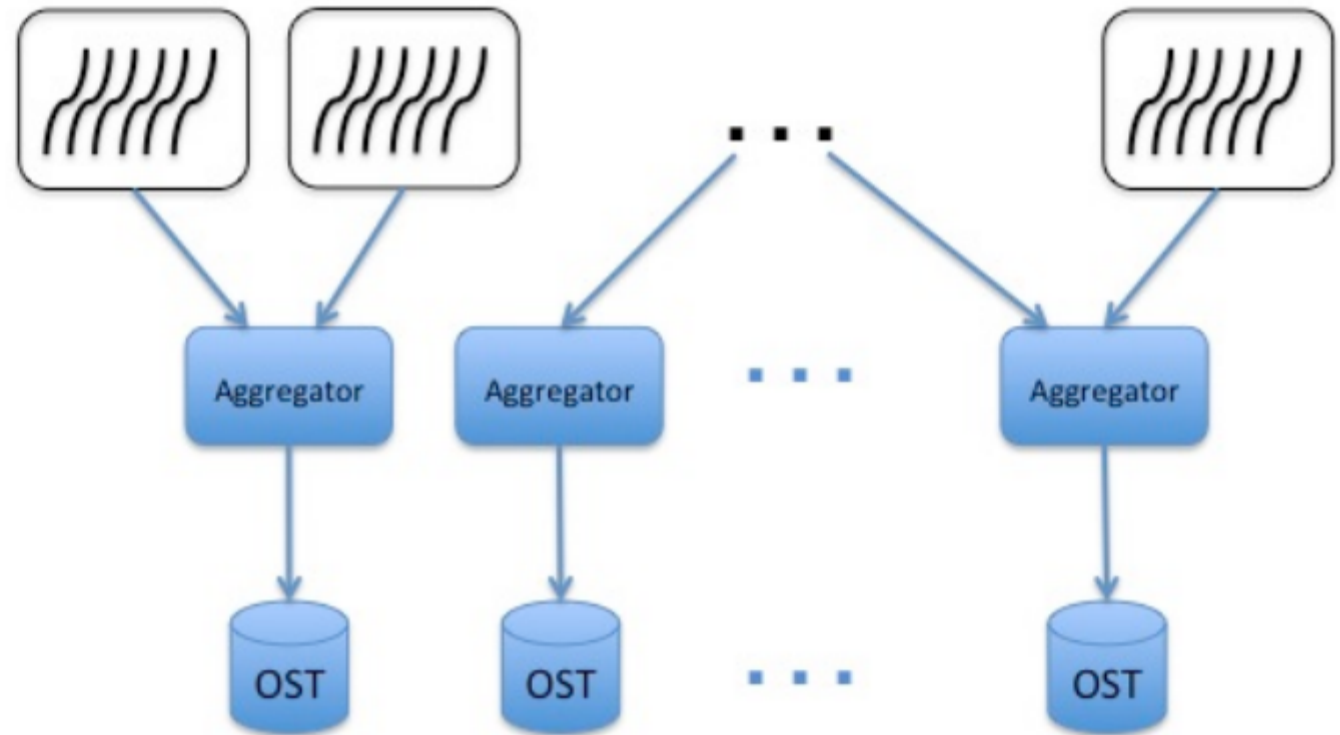
Closing HDF5 File…

Q: How long does it take to close/flush an HDF5 file?

A: A lot longer than you might expect!
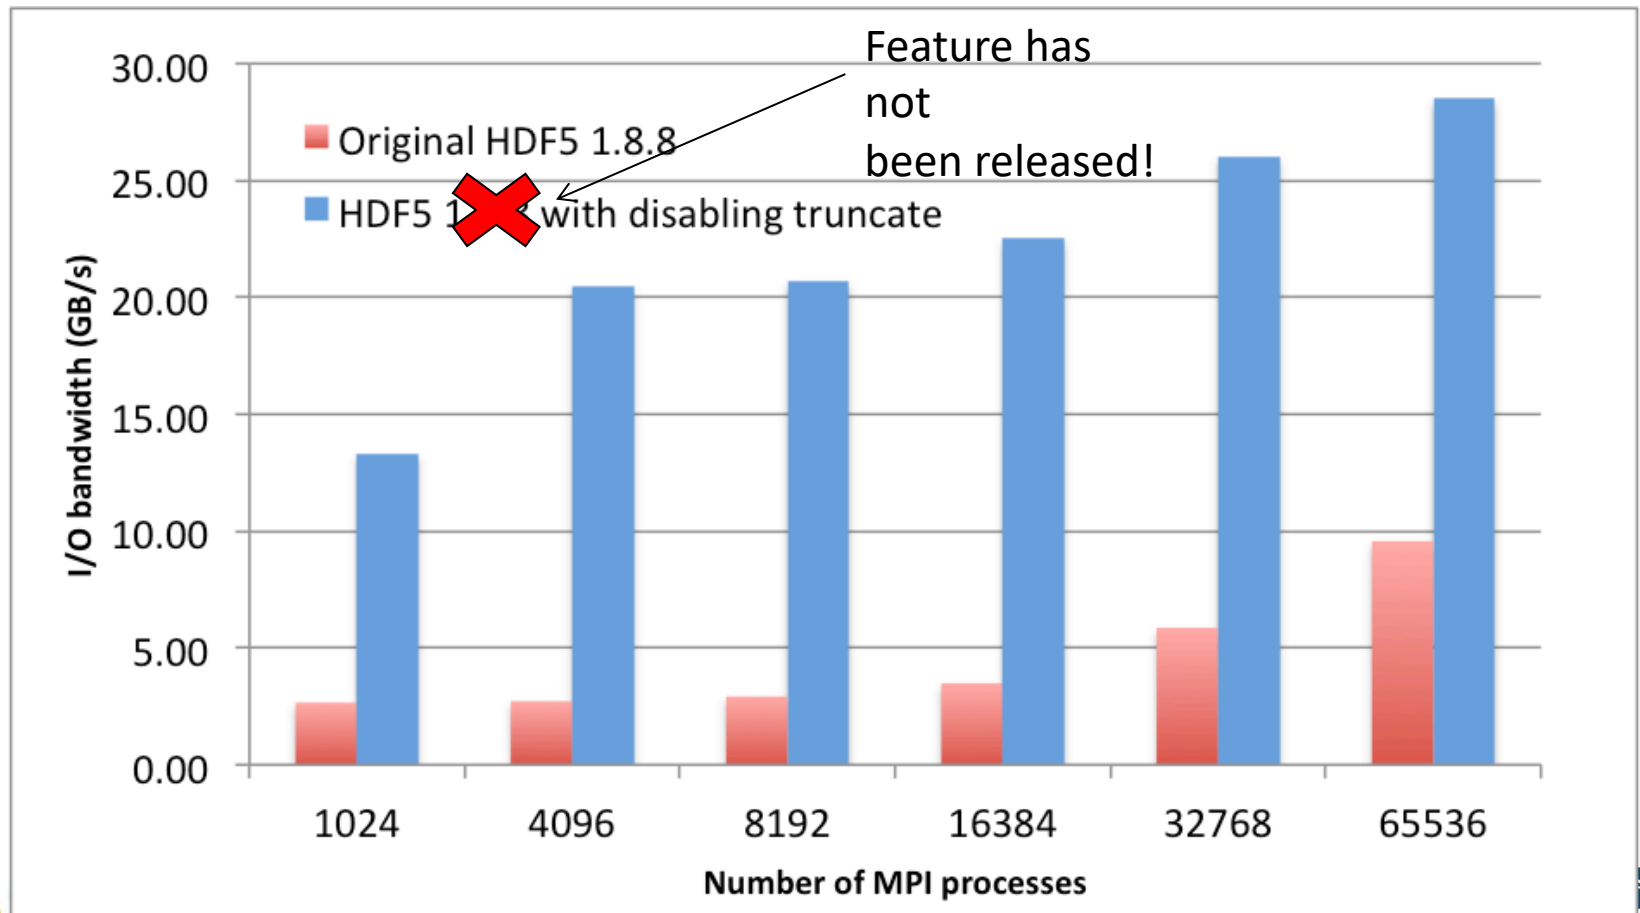
File Truncation (Today)
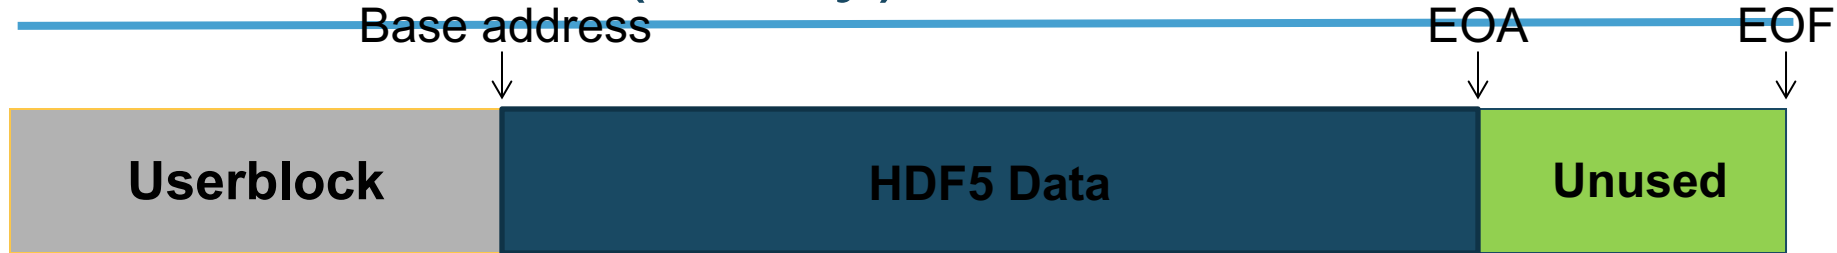
A call to `H5Fflush` or `H5Fclose` triggers a call to `ftruncate` (serial) or `MPI_File_set_size` (parallel), which can be fairly expensive.

*Currently, only **one** number is stored in the file and used for error detection.*

File Truncation (Tomorrow)

Base address   EOA   EOF

| Userblock | HDF5 Data | Unused |
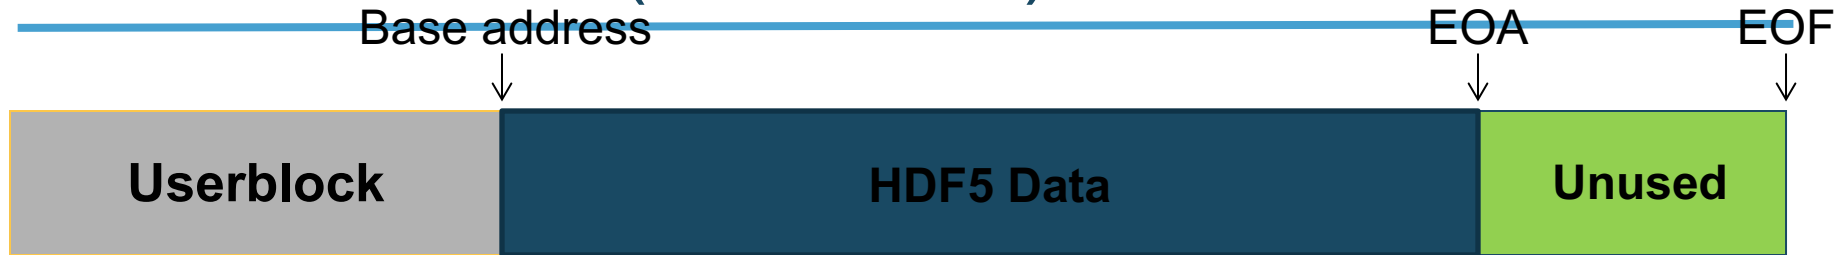
A call to `H5Fflush` or `H5Fclose` triggers both values (EOA, EOF) to be saved in the file and **no** truncation takes place, IF the file was created with the "avoid truncation" property set.

Continue   allocation from here:

Base address   EOA   EOF

| Userblock | HDF5 Data | Unused |
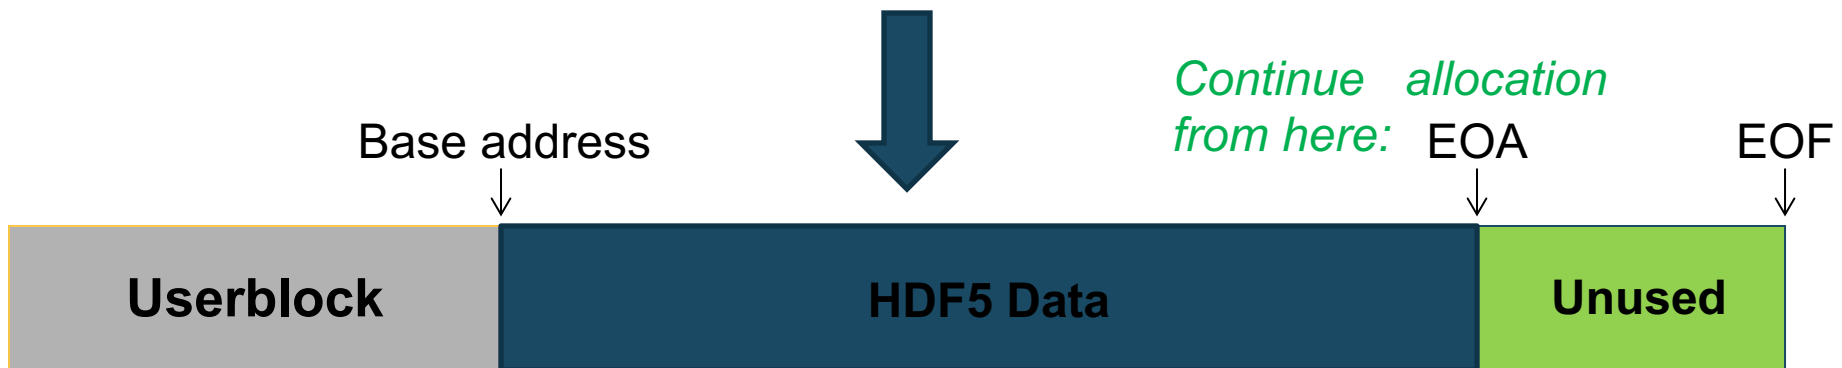
Caveat: Incompatible with older versions of the library. Requires HDF5 library version 1.12 or later.

U.S. DEPARTMENT OF **ENERGY** | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Multi-Dataset I/O - Motivation

- **HDF5 accesses elements in one dataset at a time**

- **Many HPC applications access data in multiple datasets in every time step**

- **Frequent small-size dataset access ➔ Big Trouble (≠Big Data)**

- **Parallel file systems tend not to like that.**

- **Idea: Let users to do more I/O per HDF5 call!**

- **Two New API routines:**
  - `H5Dread_multi()`
  - `H5Dwrite_multi()`
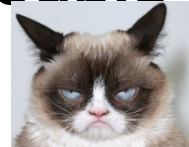
Not a new idea: PnetCDF has supported this for some time…

Sample Results

The plot shows the performance difference between using a single `H5Dwrite()` multiple times and using `H5Dwrite_multi()` once on 30 chunked datasets.

(On Hopper @ NERSC, a Cray XE-6 with Lustre file system)
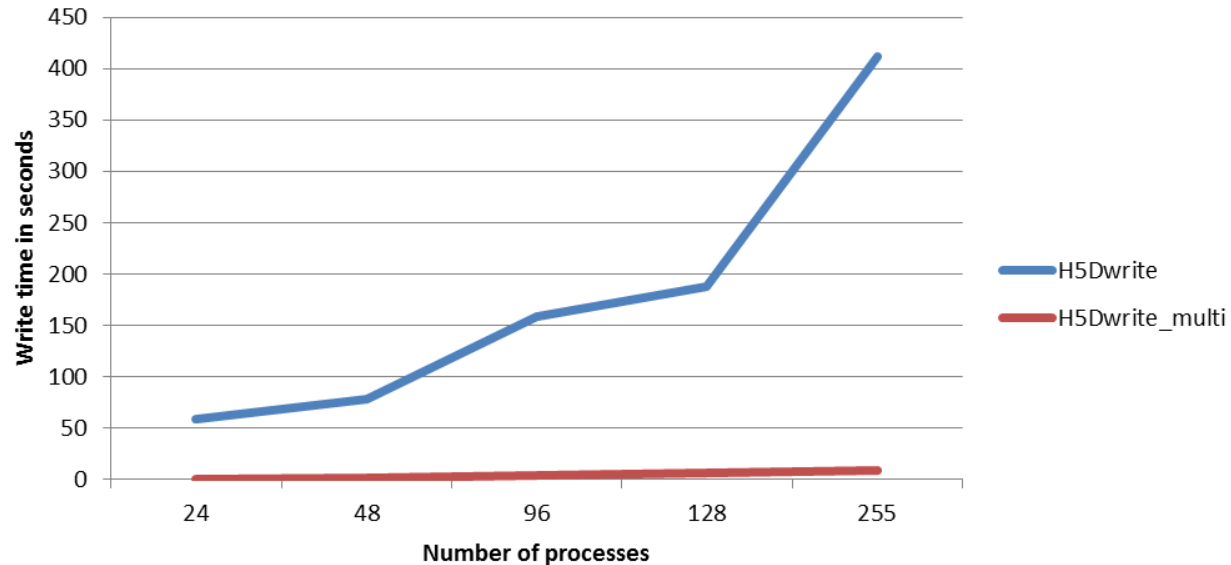
Run	Code	Nodes	Cores	File Size	Stripe Size	Write Time	Total HDF Time	Throughput
Max Throughput	Multiple Dataset	320	5,120	5 TB	1 GB	91.08 s	167.78 s	56.21 GB/s
Code Comparison	Single Dataset	320	5,120	5 TB	128 MB	116.94 s	117.37 s	43.78 GB/s
Hero Run	Single Dataset	9,314	298,048	291 TB	1 GB	5,763.14 s	5,779.89 s	51.81 GB/s

TABLE I
COMPARISON OF VPIC-IO KERNEL PARAMETERS AND OBSERVED IO THROUGHPUT.

Reference:

[Parallel and Large-scale Simulation Enhancements to CGNS](#), By Scot Breitenfeld, The HDF Group, 2015.

Examples

CGNS

CFD Standard

- **CGNS = Computational Fluid Dynamics (CFD) General Notation System**

- **An effort to standardize CFD input and output data including:**
  - Grid (both structured and unstructured), flow solution
  - Connectivity, boundary conditions, auxiliary information.

- **Two parts:**
  - A standard format for recording the data
  - Software that reads, writes, and modifies data in that format.

- **An American Institute of Aeronautics and Astronautics Recommended Practice**

CGNS Storage Evolution

- **CGNS data was originally stored in ADF ('Advanced Data Format')**

- **ADF lacks parallel I/O or data compression capabilities**

- **Doesn't have HDF5's support base and tools**

- **HDF5 superseded ADF as the official storage mechanism**

- **CGNS introduced parallel I/O APIs w/ parallel HDF5 in 2013**

- **Poor performance of the new parallel APIs in most circumstances**

- **In 2014, NASA provided funding for The HDF Group with the goal to improve the under-performing parallel capabilities of the CGNS library.**

U.S. DEPARTMENT OF **ENERGY** | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

CGNS Performance Problems

- **Opening an existing file**
  - CGNS reads the entire HDF5 file structure, loading a lot of (HDF5) metadata
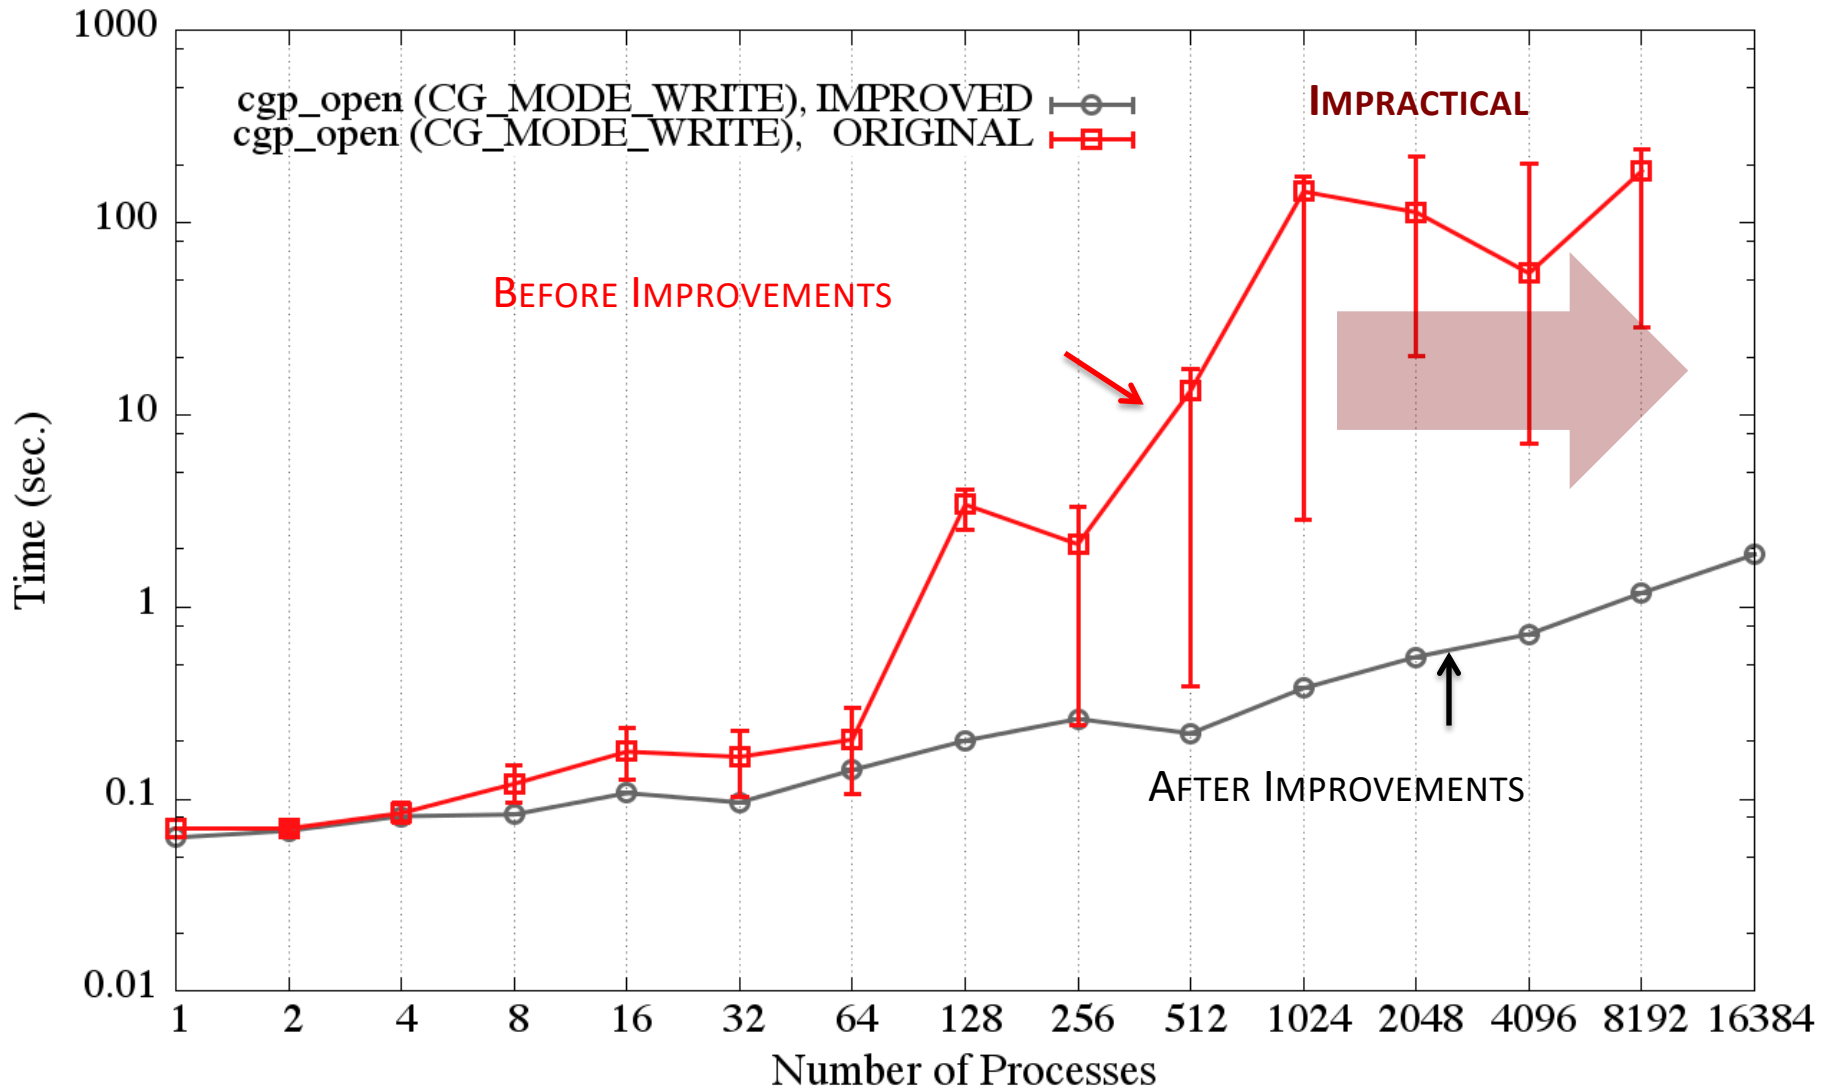 - Reads occur independently on ALL ranks competing for the same metadata
    - ➔"Read Storm"

- **Closing a CGNS file**
  - Triggers HDF5 flush of a large amount of small metadata entries
  - Implemented as iterative, independent writes, an unsuitable workload for parallel file systems

Opening CGNS File…

ATPESC 2019 - http://bit.ly/ATPESC-2019-HDF5

Metadata Read Storm Problem (I)

- **All metadata "write" operations are required to be collective:** ✗   ✓

```
if(0 == rank)
    H5Dcreate("dataset1");
else if(1 == rank)
    H5Dcreate("dataset2");
```

```
/* All ranks have to call */
H5Dcreate("dataset1");
H5Dcreate("dataset2");
```

- Metadata read operations are not required to be collective ✓   ✓

```
if(0 == rank)
    H5Dopen("dataset1");
else if(1 == rank)
    H5Dopen("dataset2");
```

```
/* All ranks have to call */
H5Dopen("dataset1");
H5Dopen("dataset2");
```

Metadata Read Storm Problem (II)

- **Metadata read operations are treated by the library as independent read operations.**

- **Consider a very large MPI job size where all processes want to open a dataset that already exists in the file.**

- **All processes**
  - Call H5Dopen("/G1/G2/D1");
  - Read the same metadata to get to the dataset and the metadata of the dataset itself
    - IF metadata not in cache, THEN read it from disk.
  - Might issue read requests to the file system for the same small metadata.

- **➔ READ STORM**

Avoiding a Read Storm

- **Hint that metadata access is done collectively**

- **A property on an access property list**

- **If set on the file access property list, then all metadata read operations will be required to be collective**

- **Can be set on individual object property list**

- **If set, MPI rank 0 will issue the read for a metadata entry to the file system and broadcast to all other ranks**

Closing a CGNS File…

CGNS Close
Hopper - Lustre

CGNS Close
Cetus - GPFS
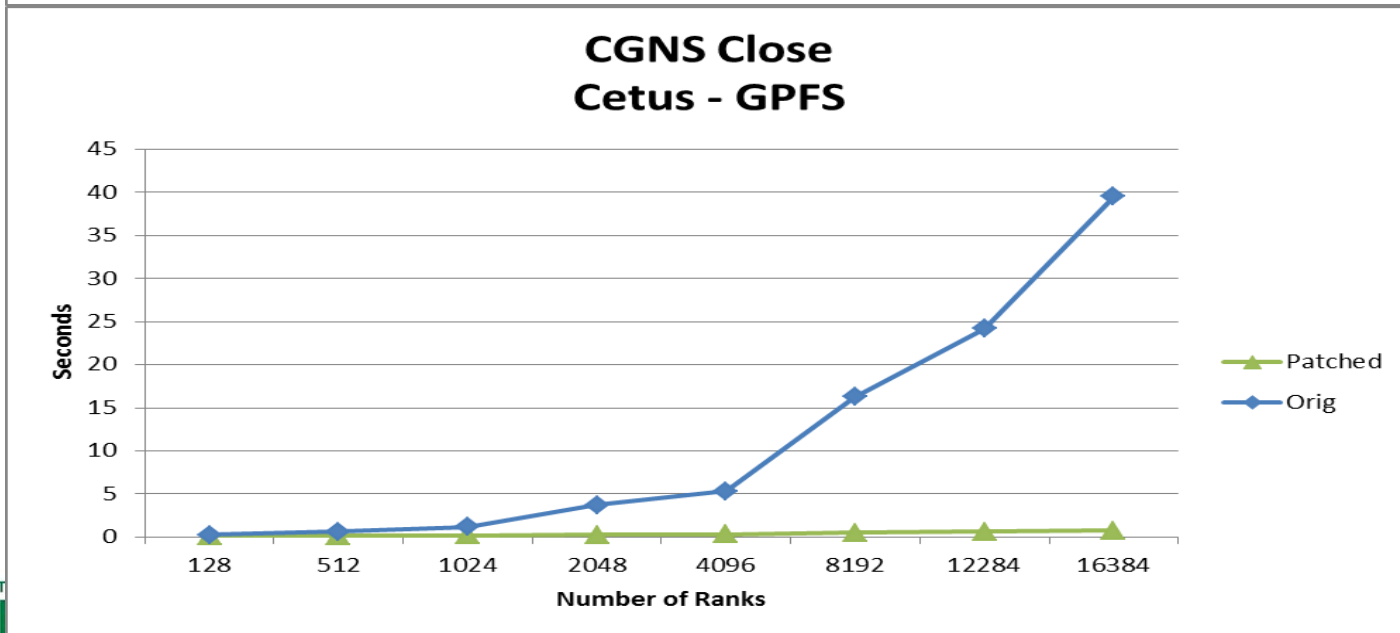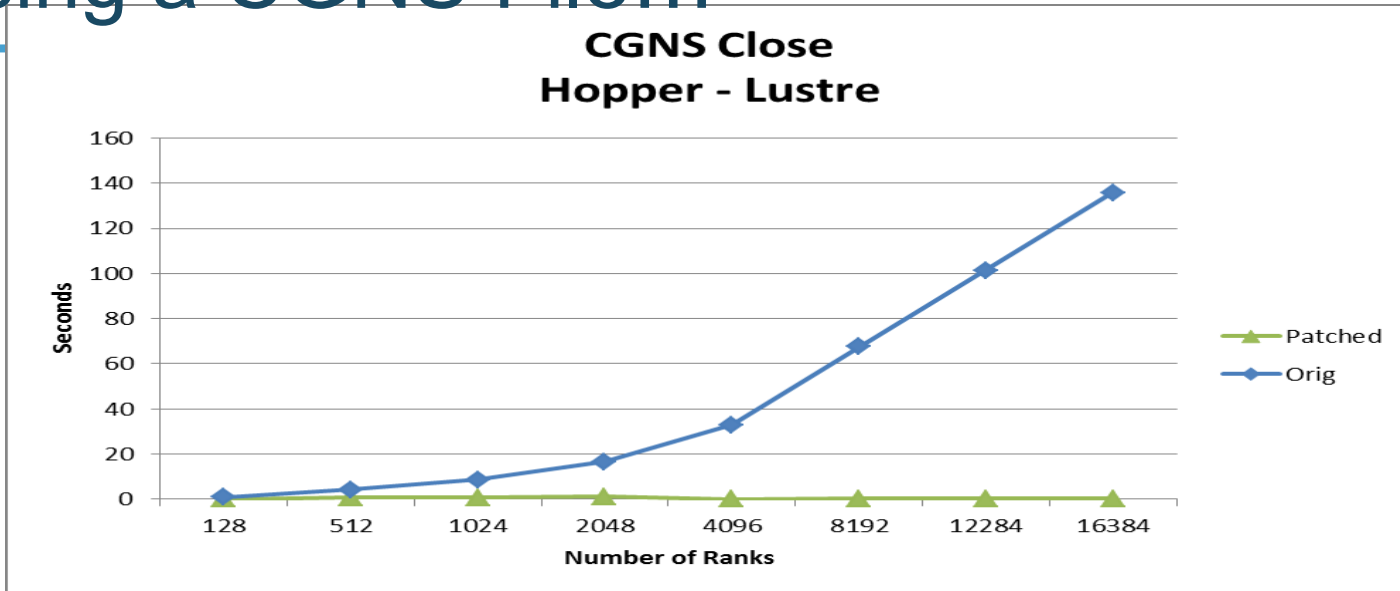
Write Metadata Collectively!

- **Symptoms: Many users reported that `H5Fclose()` is very slow and doesn't scale well on parallel file systems.**

- **Diagnosis: HDF5 metadata cache issues very small accesses (one write per entry). We know that parallel file systems don't do well with small I/O accesses.**

- **Solution: Gather up all the entries of an epoch, create an MPI derived datatype, and issue a single collective MPI write.**

HDF5 Roadmap

- **Concurrency**
  - Single-Writer / Multiple-Reader (SWMR)
  - Asynchronous I/O
  - Internal threading
- **Virtual Object Layer**
- ~~**Virtual Datasets**~~
- **Query & Indexing**
- **Native HDF5 client/server**

- Performance
  - ~~Scalable chunk indices~~
  - ~~Metadata aggregation and Page buffering~~
  - Variable-length records
- Fault tolerance
- Parallel I/O
- I/O Autotuning

Questions, Comments, Feedback?

Thank You!

ExaHDF5 - Features

- **Virtual Object Layer (VOL) integration into HDF5**

- **Caching and prefetching – Data Elevator VOL**

- **Topology-aware I/O**

- **Asynchronous I/O**

- **Independent metadata updates**

- **Workflow supporting features – SWMR**

- **Querying HDF5 data and metadata**

- **Interoperability with other file formats**
  - **PnetCDF/netCDF, ADIOS**

- **Maintenance and release support**

- **ECP engagement w/ AD, ST, and HT**
  - **Consulting and performance tuning for applications**

ExaHDF5 – Development Timeline

ATPESC 2019 - http://bit.ly/ATPESC-2019-HDF5