


Using Big Data Technologies with Earth Science Data in HDF5


HDF5 Scalable Solutions



SUMMARY


- HDF5 is Open Source software for managing big heterogeneous data <https://hdfgroup.org>
- HDF5 is used for storing Earth Science data
- HDF5 data can reside on local or parallel file systems (FS), in system memory, in Cloud and/or Object Store
- Wherever HDF5 file resides one can achieve scalable access to data
- HDF5 Virtual Object Layer (VOL) and Virtual File Drivers (VFD) are key components for achieving scalability along with other big data technologies, for example, Apache Drill and Hadoop

HDF5 1.12 Library Architecture



Highly Scalable Data Server (HSDS)

- HSDS is Open Source software
- C, Fortran, C++, Java, Python HDF5 application can use HSDS to achieve scalability in Cloud via REST VOL
- For more information see <https://github.com/HDFGroup/hsds>




- Load balancer – distributes requests to Service nodes
- Service Nodes – processes requests from clients (with help from Data Nodes)
- Data Nodes – responsible for partition of Object Store
- Object Store: Base storage service (e.g. Ceph)
- Client SDKs for Python and C are drop-in replacements for libraries used with local files.
- No significant code change to access local and cloud based data.

HDFS VFD

- HDFS VFD can access HDF5 files in Hadoop Distributed File System (HDFS) by creating access via new HDF5 1.12.0 API `H5Pset_fapl_hdfs`
- HDF5 command line tools can be used to extract metadata and raw data from HDF5 files on HDFS
- Hadoop streaming can collect data from multiple HDF5 files

- HDF5-based software in any language can access HDF5 data on HDFS, Object Store, AWS, AZURE, and Google Cloud, and in regular file systems using VOL or VFD mechanism
- For available HDF5 VOL connectors see <https://bitbucket.hdfgroup.org/projects/HDF5VOL>
- For HSDS and S3 VFDs use HDF5 1.12.0 and HDF5 1.10.6 releases

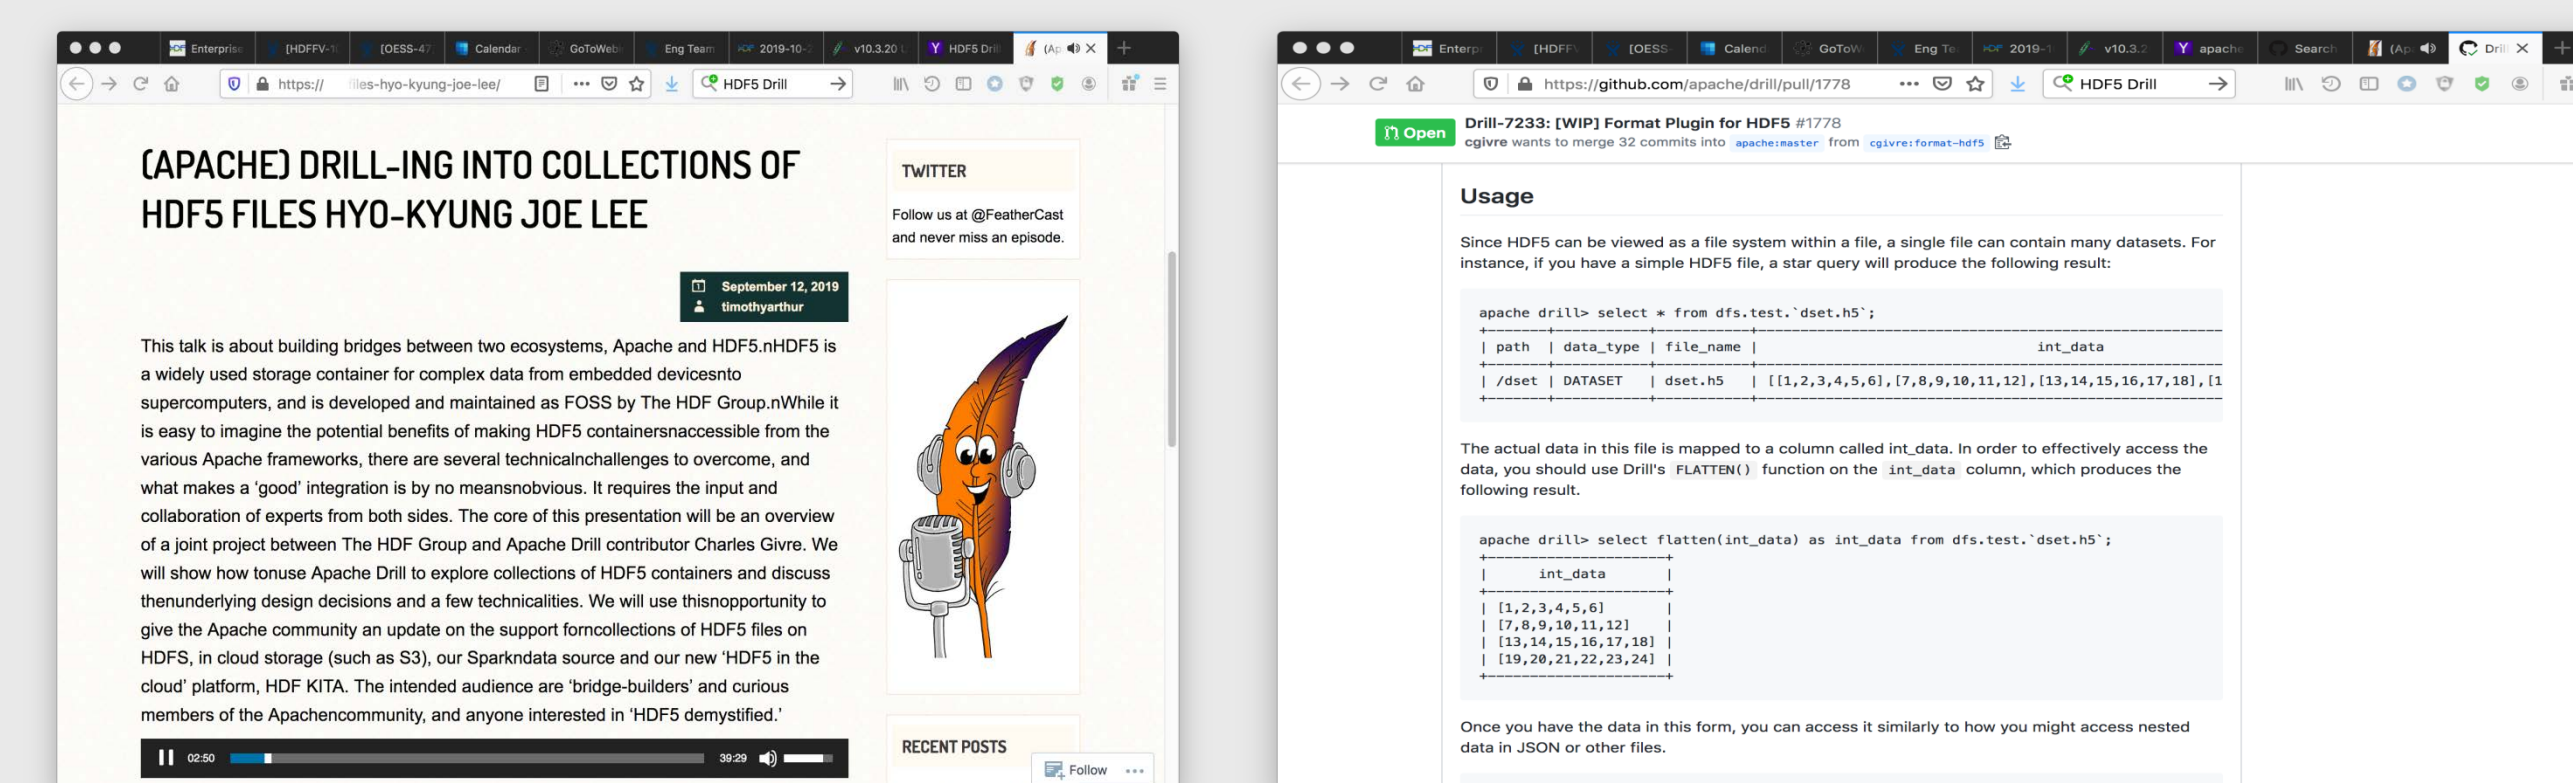
HDF5 and HPC



S3 VFD

- S3 VFD can access HDF5 files in Amazon Web Service (AWS) S3 bucket by creating access via new HDF5 1.12.0 API `H5Pset_fapl_ros3`
- `h5dump` and `h5ls` tools have a flag to access HDF5 file on S3
`h5ls -vfd=ros3 <S3-URL>/file.h5`
- S3 VFD uses "range get" commands to get "bytes" from HDF5 file stored in S3 bucket
- Optimizations are planned to utilize parallel access to S3 buckets

HDF5 Apache Drill Working with HDF5 file collections



Drill HDF5 Format Plugin on Github
<https://github.com/apache/drill/pull/1778>

HDF5 and Deep Memory Hierarchy (DMH)

- HDF5 can use the VOL layer to take advantage of DMH to achieve scalability
- HDF5 VOL Data Elevator connector saves data in burst buffers and then moves it behind the scenes to permanent storage.
- "Data Elevator: Low-contention Data Movement in Hierarchical Storage System", Bin Dong, Suren Byna, Kesheng Wu, Prabhat, Hans Johansen, Jeffrey N. Johnson, and Noel Keen, Lawrence Berkeley National Laboratory, USA

Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author[s] and do not necessarily reflect the views of Raytheon or the National Aeronautics and Space Administration.

More information on HDF5
<https://hdfgroup.org>

