

# A Brief History of HDF5

MIKE FOLK

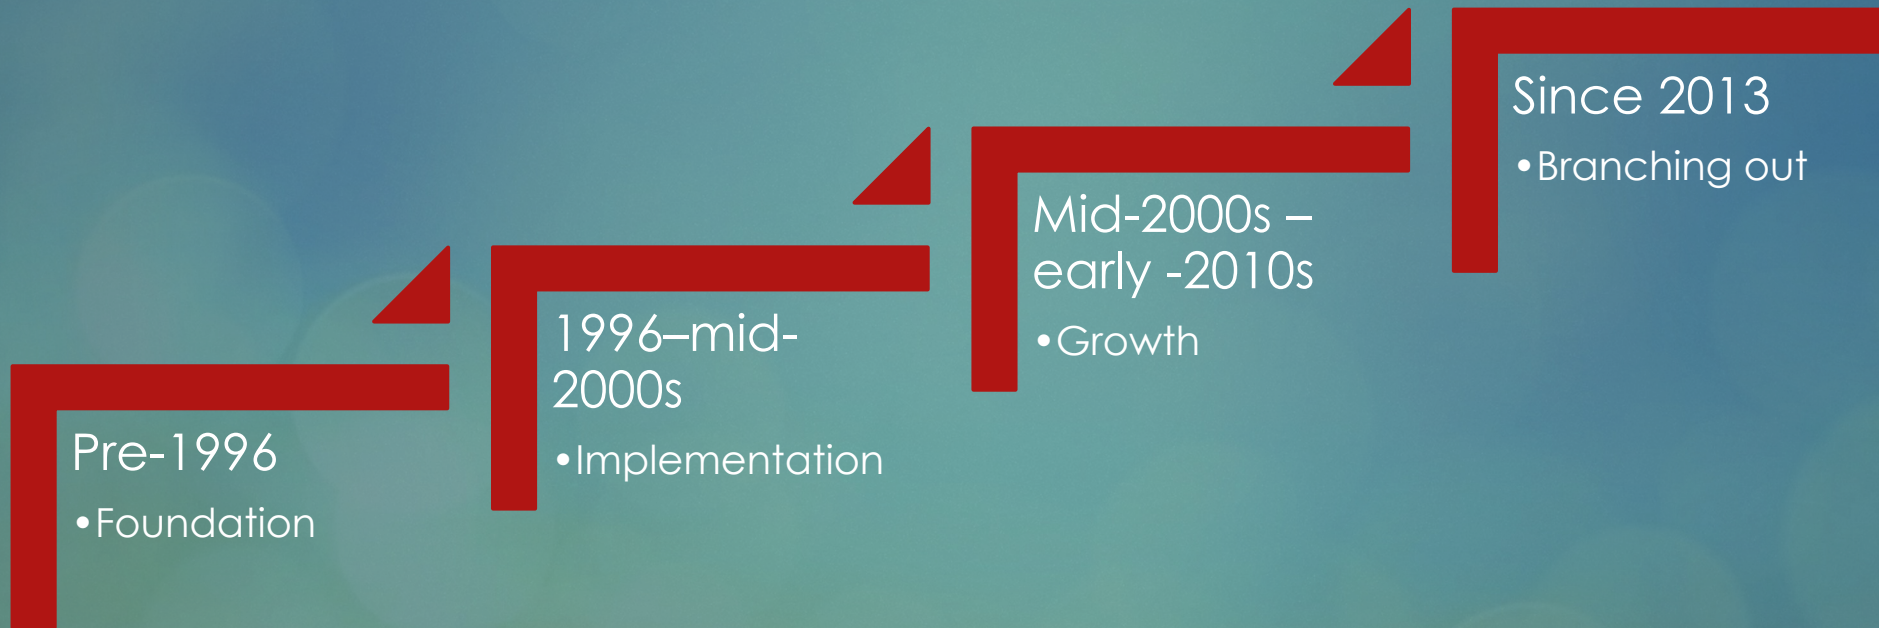
AUGUST 17, 2023

August 16-18, 2023

HUG23



# Overlapping periods






# Foundations – HDF

PRE-1996

# NCSA

4

HUG23  
August 18, 2023



# Data technologies



Network – ftp,  
gopher, telnet



I/O mostly to  
disk and serial



Typical  
workstation  
disk 500 MB.

Local/global attributes



Raster images



Multidimensional arrays  
Groups  
Tables

----

Compression, Chunking



HUG23  
August 16-  
18, 2023

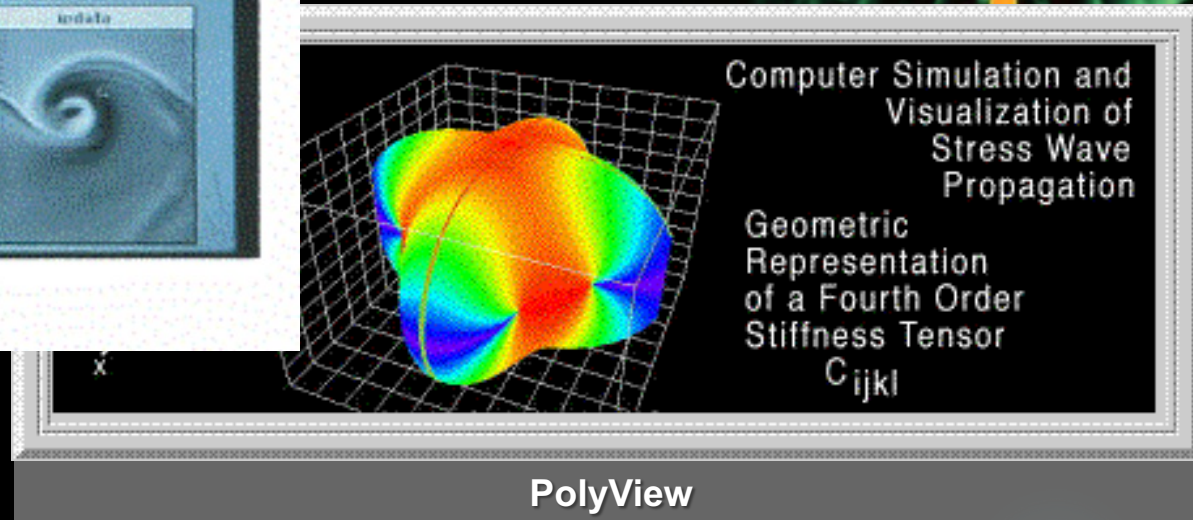

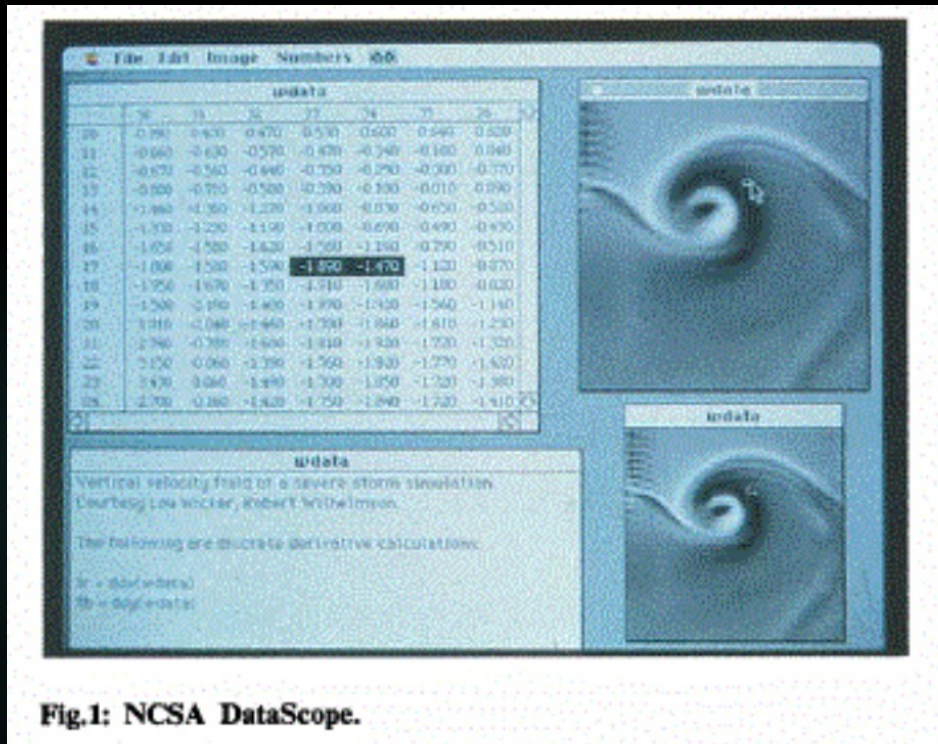
4.0  
a.k.a.  
HDF4

# NCSA desktop visualization software


“NCSA Image”

7

HUG23



# Earth Observing System





# Implementation

1996-MID 2000'S

FIRST CRISIS, AND CREATING HDF5

# Technical shortcomings of HDF

10

HUG23  
August 16-  
18, 2023

- ▶ Limits on object & file size (<2GB)
- ▶ Limited number of objects (<20K)
- ▶ Rigid data models
- ▶ Code complexity
- ▶ I/O performance

# How to maintain a nuclear stockpile in the absence of testing?

11

HUG23  
August 16-  
18, 2023



Accelerated Strategic Computing Initiative (ASCI)

“I’ll be back...”

12

HUG23  
August 16-  
18, 2023



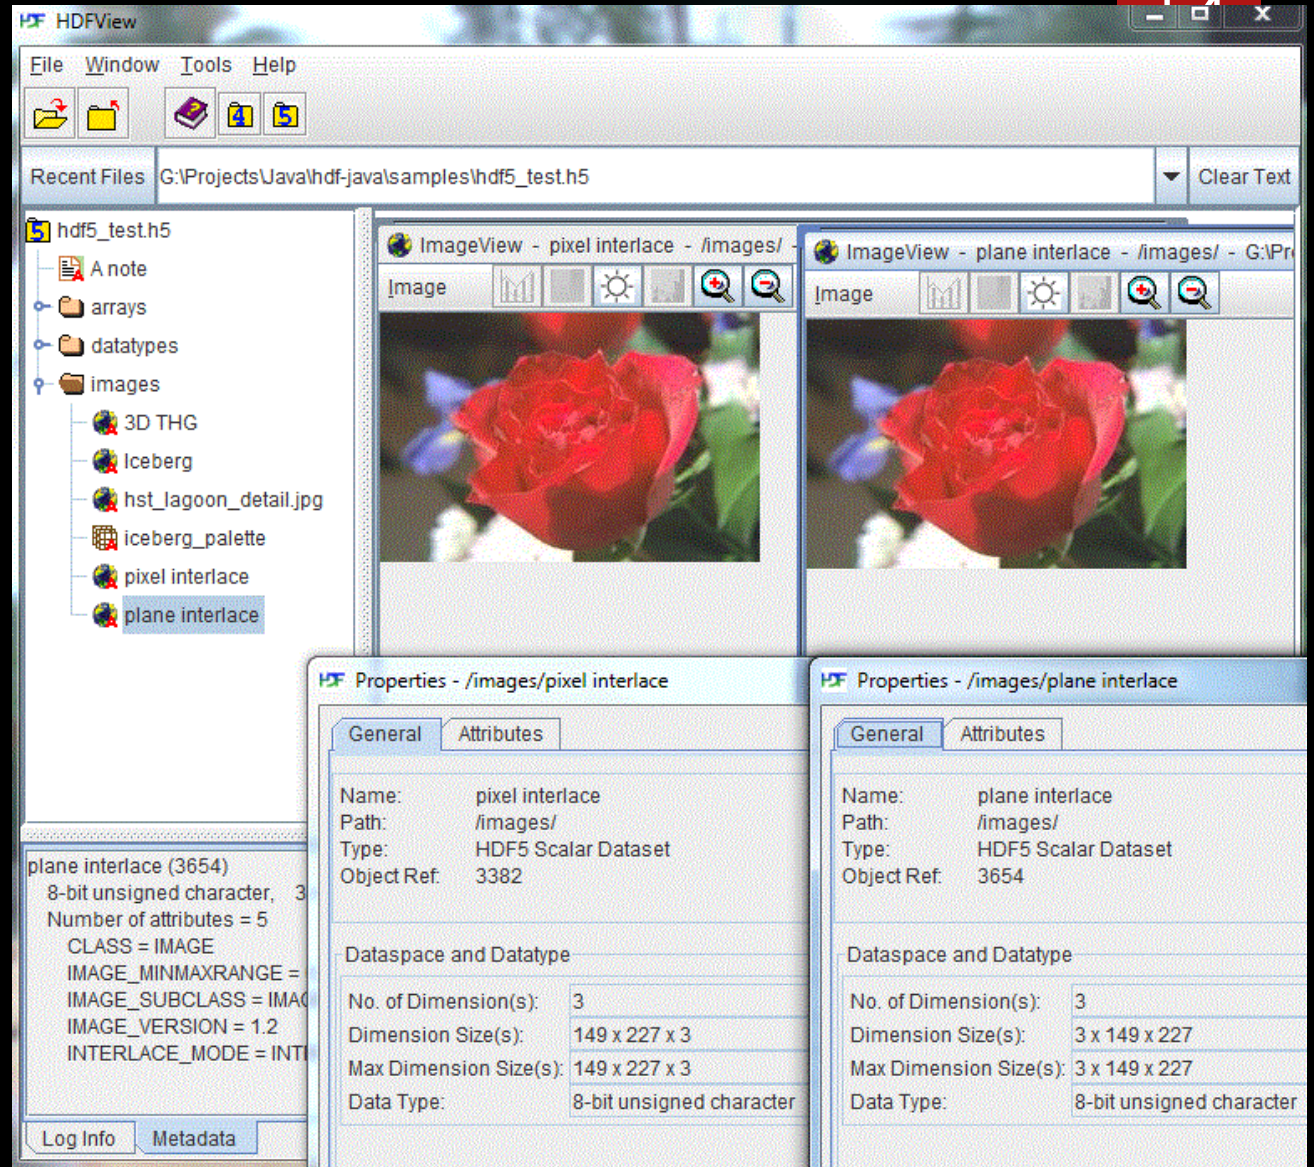
"I'll be back..."™

# First release!

- ▶ Release of HDF5 version 1.0.0 (Nov '98)
  - ▶ **“We are pleased to announce the first official release of HDF5, version 1.0.0...**
  - ▶ **This first release ONLY officially includes support for the serial implementation** of HDF5. The parallel implementation is in place, but we have encountered some problems with it.”
- ▶ Parallel Release of HDF5 version 1.0.1 (Jan '99)
  - ▶ **“The first official release of the PARALLEL IMPLEMENTATION of HDF5,** version 1.0.1. It uses the ROMIO version of the MPIO interface for parallel I/O support.”


Java interface 1.0  
(Nov '99)

HDFView  
1.0 (Jan 2000)



**visit**

**SOFTWARE**  
**that delivers parallel**



# Sponsors




Limit Point Systems





# Sponsors





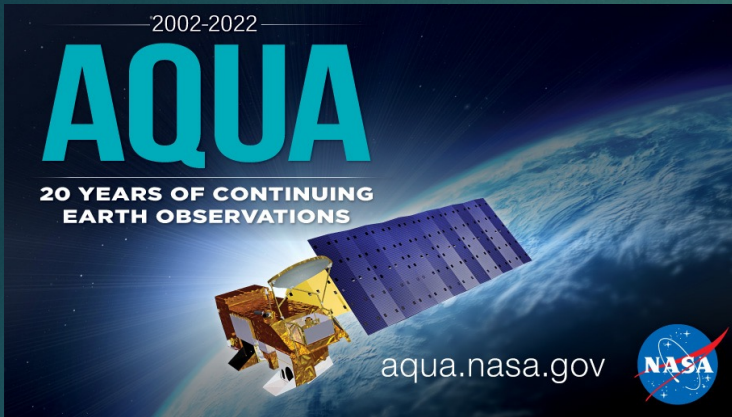
HUG2  
August  
18, 20

## HDF AND HDF-EOS WORKSHOP I

September 8-10, 1997  
GSFC, Maryland

# EOS Aqua and Aura

HUG23  
August 16-  
18, 2023



# R&D 100



Reed Business Information

September 2002 Vol. 44, No. 9

**DEPARTMENTS**

- 10 News & Developments
- 17 Data Management & Analysis
- 70 Regulatory Update
- 71 Emerging Technologies

**R&D**

The R&D Community

DNA Quantifier

Ceramic Composite Rolling Circle Amplifier

Laser Doppler Velocimeter RF Personal Monitor Epoxy Electrocoat

Atomic Force Microscope Thickness Mapper Compact Communicator

Sphigmomanometer

CO<sub>2</sub> Washer Simple Syringe

Ultrasound Doppler Droplet Dispenser

Ophthalmic Analyzer Brain Imager

Dig Gene Test Nerve Stimulator

Quiet Foam Magnetism Analyzer

Landmine Detector

Laser Diode Array

FT-Raman Spectrometer

Super Squeeze

Ultrasonic Spectrometer

Art Restorer

Indoor Air Tester Nanosilumina Fiber Holographic Memory

Self-Cleaning Glass Fast X-Ray Sensor

SNP Genotyper

**The 10th Annual R&D 100 Awards**


Honoring Today's Technological Innovations

# Growth

20

HUG23  
August 16-  
18, 2023

# Motivation: Why is this area of work important?

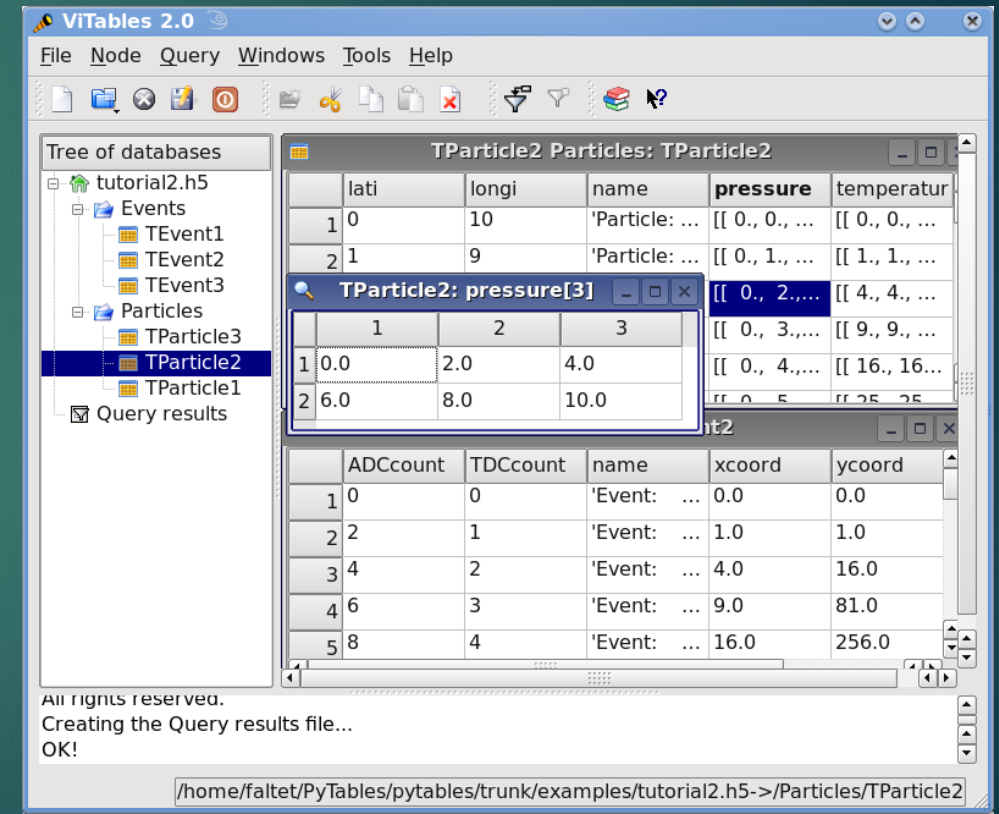



Jim Gray,  
Distinguished  
Engineer at  
Microsoft,  
1998 Turing Award  
winner

*While the commercial world has standardized on the relational data model and SQL, no single standard or tool has critical mass in the scientific community. There are many parallel and competing efforts to build these tool suites – at least one per discipline. Data interchange outside each group is problematic. In the next decade, as data interchange among scientific disciplines becomes increasingly important, a common HDF-like format and package for all the sciences will likely emerge.*

“Scientific Data Management in the Coming Decade,” Jim Gray, David T. Liu, Maria A. Nieto-Santisteban, Alexander S. Szalay, Gerd Heber, David DeWitt, Cyberinfrastructure Technology Watch Quarterly, Volume 1, Number 2, February 2005

# 2002: Pytables 1.0



The screenshot shows the ViTables 2.0 application window. On the left is a 'Tree of databases' showing a hierarchy: tutorial2.h5 > Events > TEvent1, TEvent2, TEvent3; and Particles > TParticle3, TParticle2, TParticle1. The 'Query results' checkbox is checked. The main area displays two tables. The top table is titled 'TParticle2 Particles: TParticle2' and has columns: lati, longi, name, pressure, and temperatur. The bottom table has columns: ADCcount, TDCcount, name, xcoord, and ycoord. A small pop-up window titled 'TParticle2: pressure[3]' is also visible, showing a 2x3 grid of values.

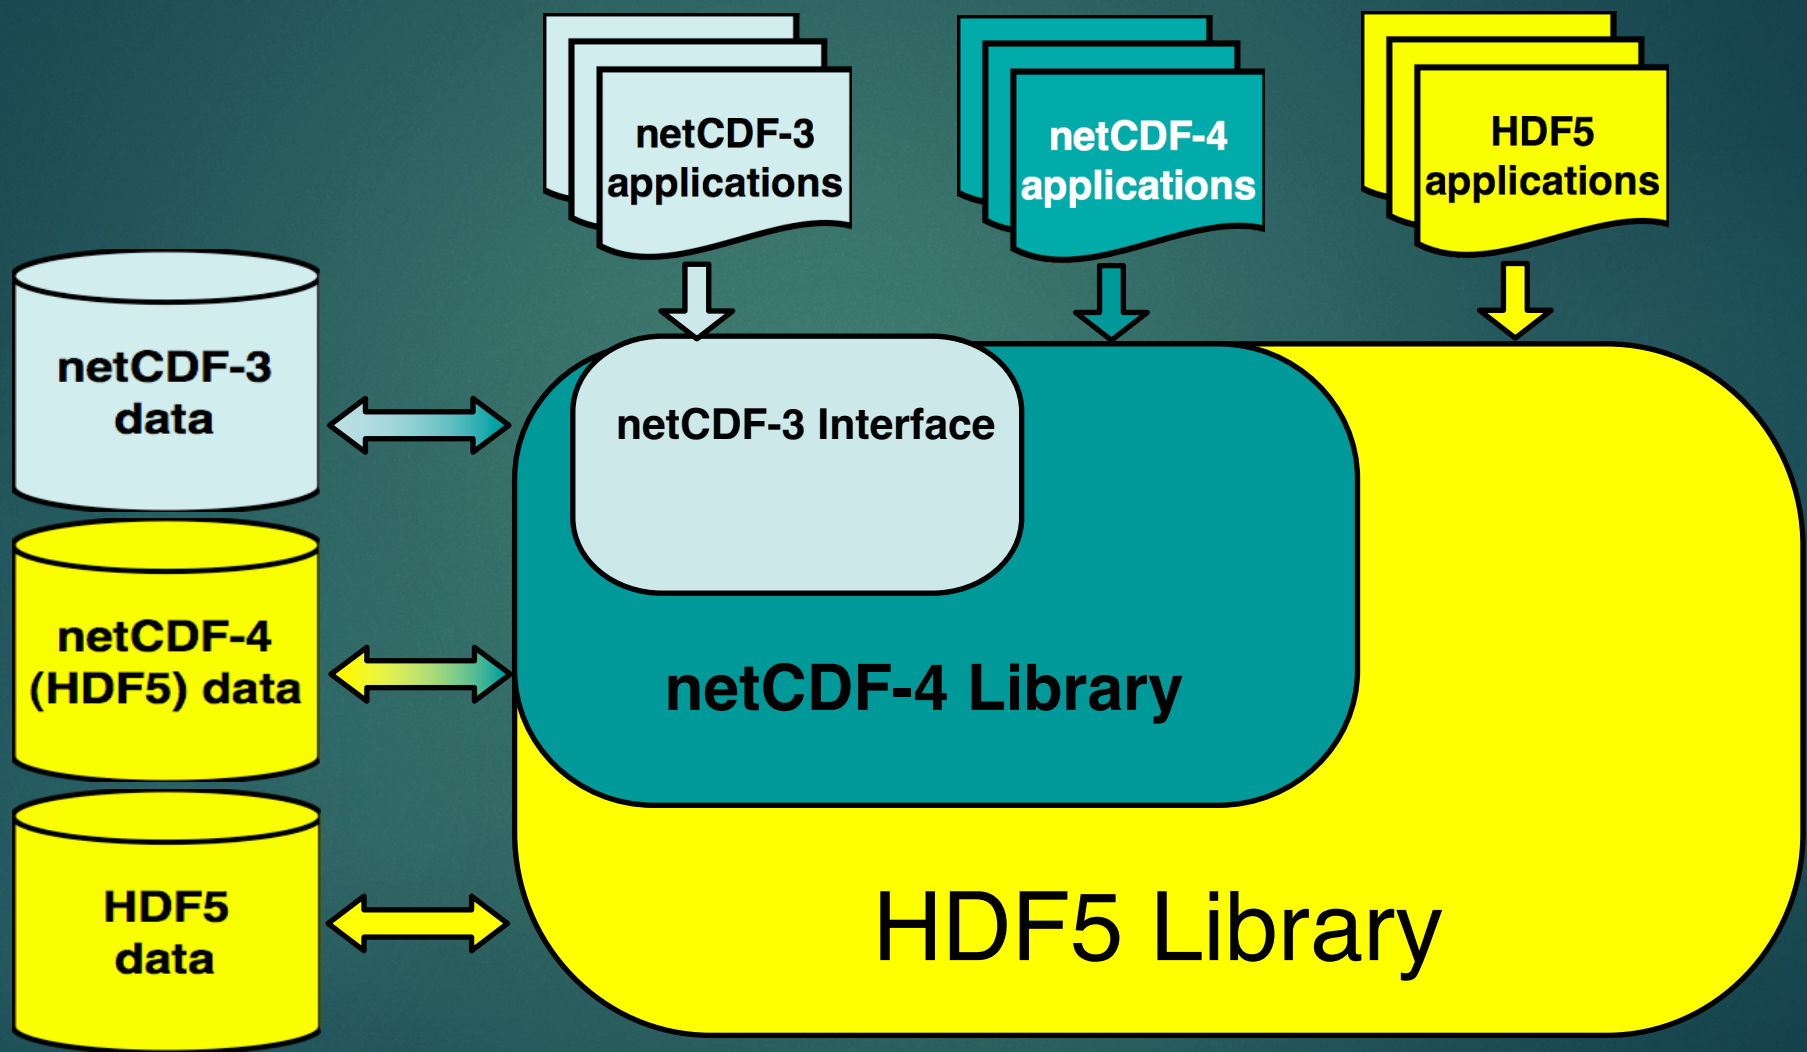
lati	longi	name	pressure	temperatur
1 0	10	'Particle: ...	[[ 0., 0., ...	[[ 0., 0., ...
2 1	9	'Particle: ...	[[ 0., 1., ...	[[ 1., 1., ...

ADCcount	TDCcount	name	xcoord	ycoord
1 0	0	'Event: ...	0.0	0.0
2 2	1	'Event: ...	1.0	1.0
3 4	2	'Event: ...	4.0	16.0
4 6	3	'Event: ...	9.0	81.0
5 8	4	'Event: ...	16.0	256.0

All rights reserved.  
Creating the Query results file...  
OK!

/home/faltet/PyTables/pytables/trunk/examples/tutorial2.h5->/Particles/TParticle2





**NCSA-NARA investigations of  
HDF5 for long term retention of  
engineering data**

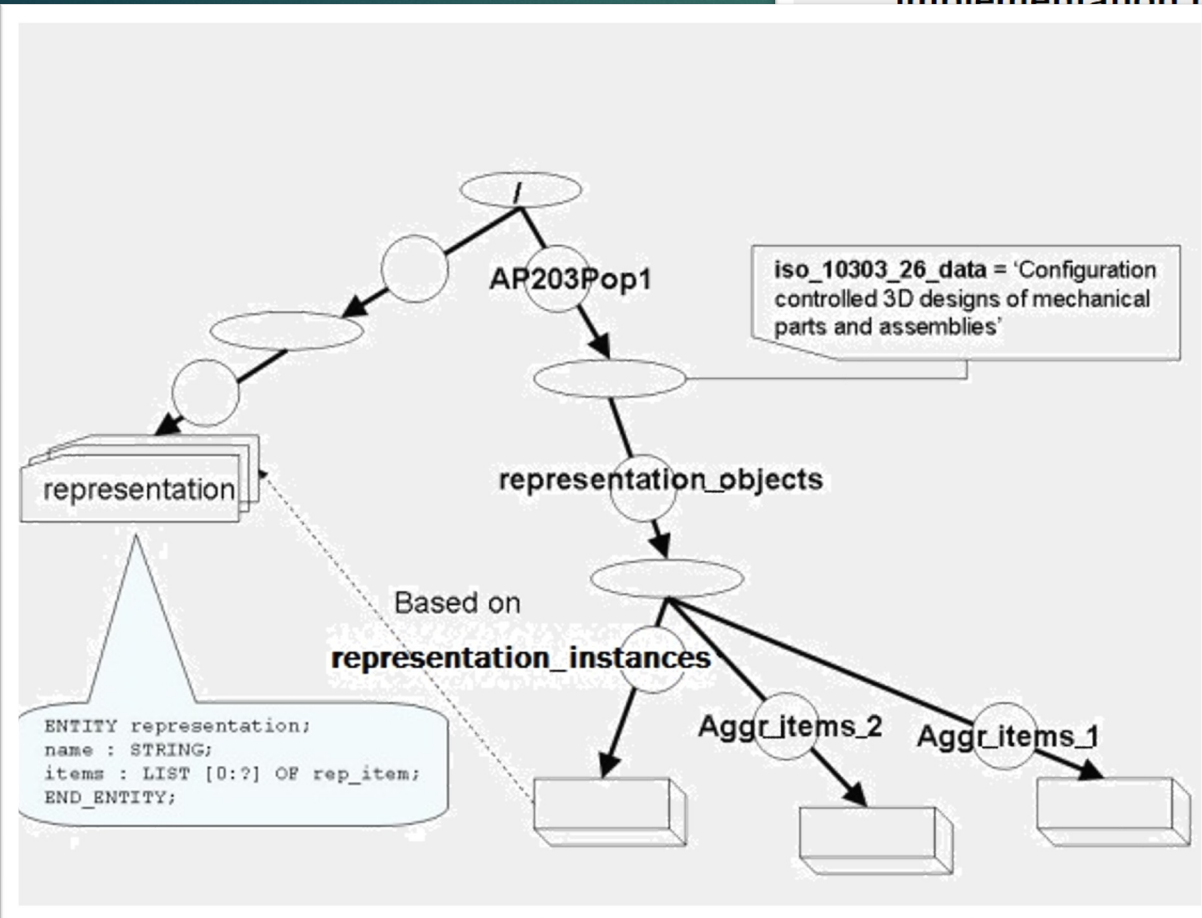


First edition 2011-12-15

## Industrial automation systems — Product data representation and exchange — Part 26: Implementation methods: Binary representation of EXPRESS-driven data

Industrielle et intégration — Représentation et échange de données de produits —  
Mise en oeuvre: Représentation binaire de données menées par EXPRESS

[ISO/TS 10303-26:2011\(E\)](#)



# Nexus & CGNS




**NeXus: a common data format for  
neutron, x-ray, and muon science**






# HDF Time History

HUG23  
August 16-  
18, 2023



## HDF Packet





# BioHDF - Open Binary File Standards for Bioinformatics

*By Todd Smith, Chairman and CEO*

October 29, 2004

Entertainment

29

# THE POLAR EXPRESS

August 16-18, 2023

U.S. \$23



# Branching out

MID-2000S – LATE 2010S

# Spinoff?

31

- ▶ The HDF Group prepares to leave U of I to form an independent non-profit.

HUG23  
August 16-  
18, 2023




# A hedge fund





# Broad lab support




# Lab instruments and software

34


- ▶ Keysight, Mathworks, Labview, need features
- ▶ IVI\* File Format Specification

\*Interchangeable Virtual Instruments





- ▶ Data exchange standard for the reservoir life cycle.



# pandas


$$y_{it} = \beta' x_{it} + \mu_i + \epsilon_{it}$$



# H5py

```
import h5py
import numpy as np

# Create an HDF5 file
with h5py.File("mytestfile.hdf5", "w") as f:
    # Create a dataset within the file
    dset = f.create_dataset("mydataset", (100,), dtype='i')
```



38

HUG23  
August 16-  
18, 2023

# Other libraries, language bindings

- ▶ Java
  - ▶ JHDF5
  - ▶ Java HDF5 (Unidata)
  - ▶ jHDF
- ▶ C++
  - ▶ H5cpp (Varga)
  - ▶ H5cpp (European spallation source)
  - ▶ Cclib
  - ▶ HighFive
- ▶ Python
  - ▶ H5py
  - ▶ Pyfive
  - ▶ PyTables
- ▶ Others
  - ▶ C# - HDF5.Pinvoke
  - ▶ D - d\_hdf5
  - ▶ Delphi - Delphi HDF5
  - ▶ Fortran - H5fortran
  - ▶ Golang (Google) - Go-hdf5
  - ▶ Haskell – Haskell HDF5
  - ▶ Javascript - Jsfive
  - ▶ Julia - HDF5.jl
  - ▶ Lisp - Hdf5-cffi
  - ▶ R - HDF5r, rhdf5, H5package
  - ▶ S - SLh5
  - ▶ Amazon S3 - H5Coro
  - ▶ Yorick - Yorick HDF5


# Science



# Trillion Particle VPIC Simulation + FastQuery

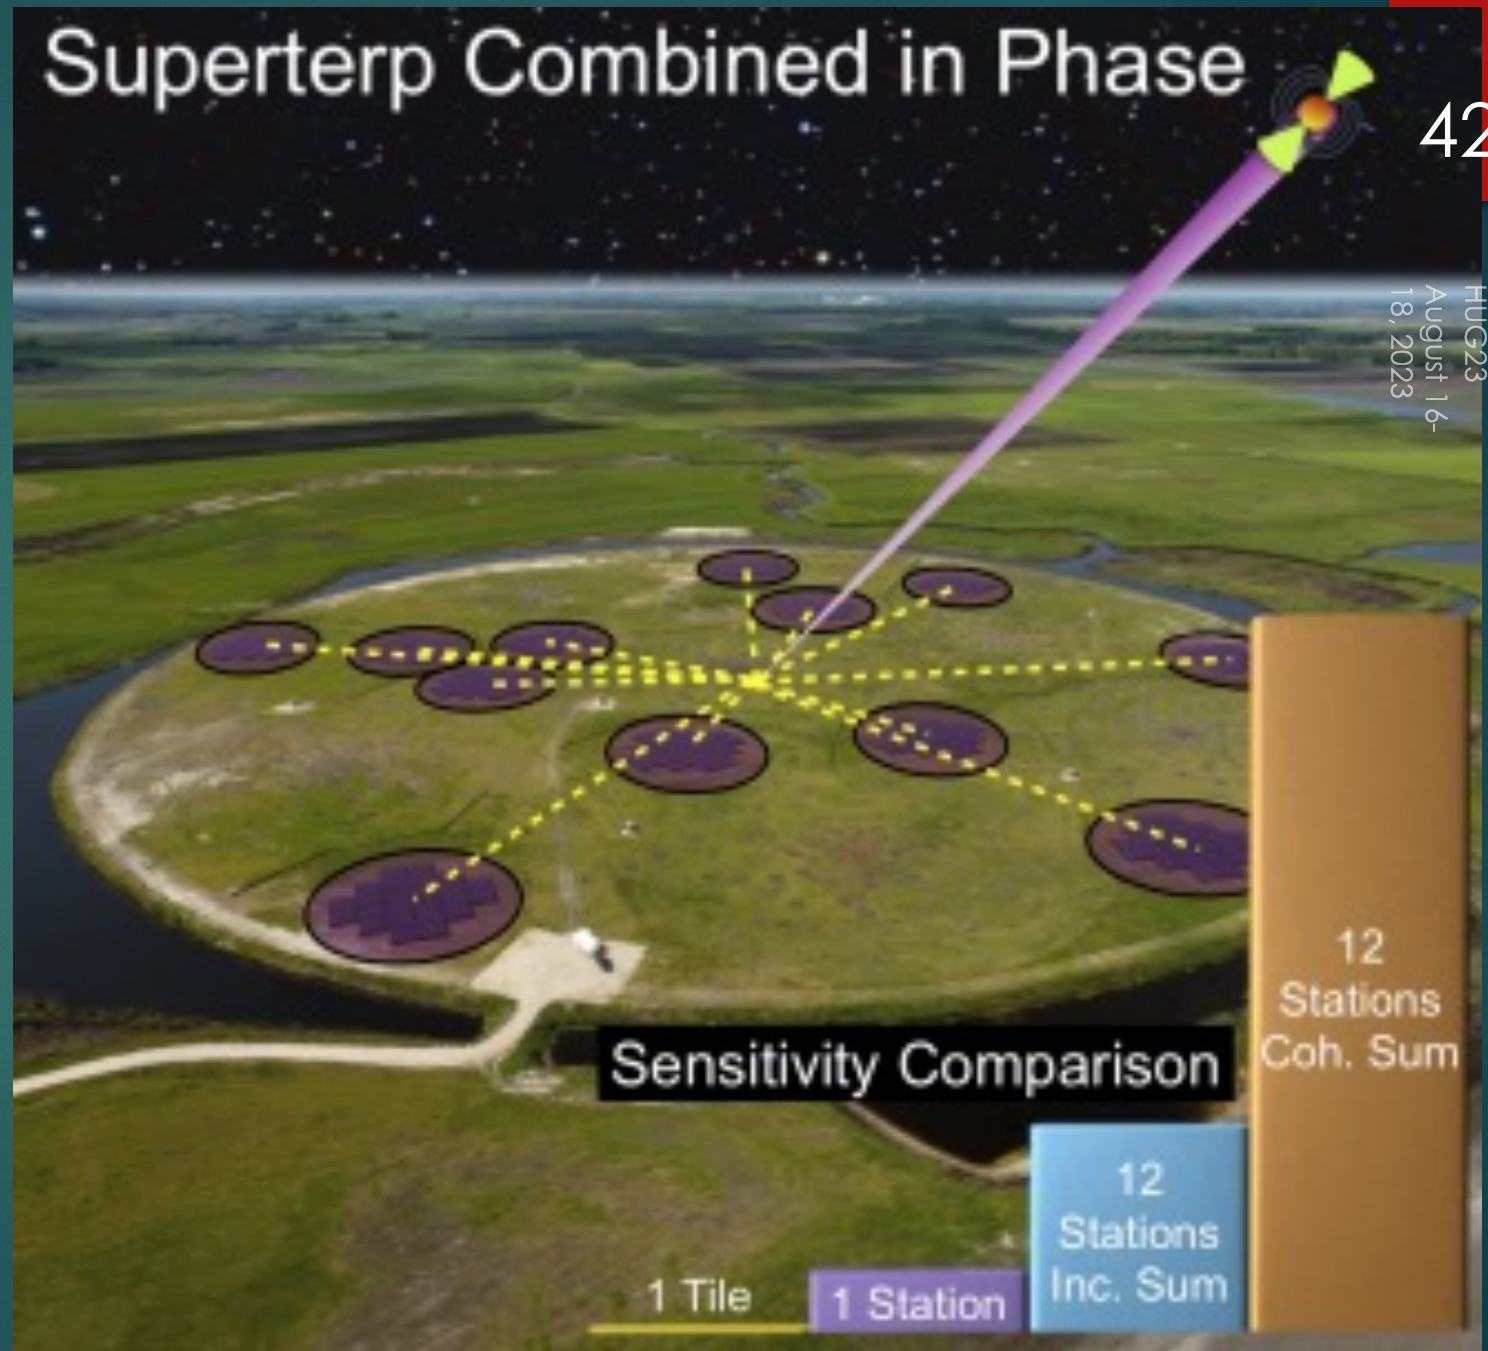
41

- ▶ Trillion particle simulation on 120,000 cores produces 350 TB dataset
- ▶ Parallel HDF5 obtained peak 35GB/s I/O rate and 80% sustained bandwidth
- ▶ Developed FastQuery using FastBit to use multicore hardware
- ▶ FastQuery took 10 min to index and 3 secs to query energetic particles
- ▶ Enabled novel discoveries in plasma physics



# LOFAR

- ▶ “5 seconds of data ... adds up to approximately 200 GB.”



# New challenges in mid-2010's

43

- ▶ Technology change
- ▶ Funding challenges

HUG23  
August 16-  
18, 2023

# 2012 - 2023

## Data technologies don't stand still

- ▶ Amazon, Google, Microsoft clouds mature
- ▶ Storage hierarchies, object storage, etc., gain in HPC and cloud
- ▶ The Exascale Computing Project (ECP) established, “to prepare for the world's first capable exascale ecosystem.”

# NASA and DOE continue major support



**But hedge  
fund  
support  
ends**

Light source  
community funds  
SWMR, VDS,  
direct chunk I/O,  
etc.

## DEVELOPING HDF5 FOR THE SYNCHROTRON COMMUNITY

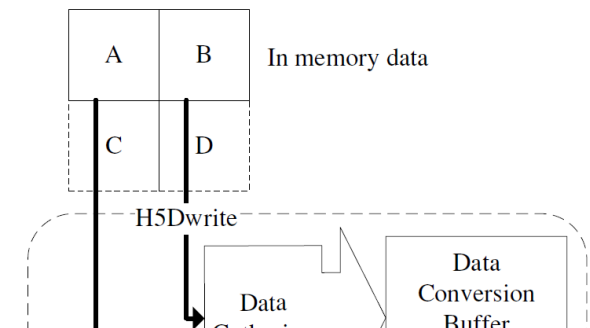
N. Rees, Diamond Light Source, Oxfordshire, UK  
H. Billich, PSI, Villigen, Switzerland  
A. Götz, ESRF, Grenoble, France  
Q. Koziol & E. Pourmal, The HDF Group, Champaign, IL, USA  
M. Rissi, Dectris AG, Baden, Switzerland  
E. Wintersberger DESY, Hamburg, Germany

HDF5  
August 13  
19, 2015

### Abstract

HDF5[1] and NeXus[2] (which normally uses HDF5 as its underlying format) have been widely touted as a standard for storing Photon and Neutron data. They offer many advantages to other common formats and are widely used at many facilities. However, it has been found that the existing implementations of these standards have limited the performance of some recent detector systems. This paper describes how the synchrotron light source community has worked closely with The HDF Group to drive changes to the HDF5 software to make it more suitable for their environment. This includes developments managed by a detector manufacturer (Dectris - for direct chunk writes) as well as synchrotrons (DESY, ESRF and Diamond - for pluggable filters, Single Writer/Multiple Reader and Virtual Data Sets).

gathering and scattering, data conversion, filter pipeline, and chunk cache and writes the data chunk to the file directly (see Fig. 1). This allows the user program to compress the data outside the library – potentially using parallel algorithms or hardware accelerators. It also avoids a number of data copies, which limits any dataflow through the filter pipeline to ~500 MB/sec on typical processors.



-Oct-2015 11:00




20

**iter**

**china eu india japan korea russia usa**

47





Laser Interferometer  
Gravitational-Wave Observatory  
Supported by the National Science Foundation  
Operated by Caltech and MIT

[About](#) [Learn More](#) [News](#) [Gallery](#) [Educational Resources](#) [For Scientists](#) [Study & Careers](#)

# 2017 Nobel Prize in Physics



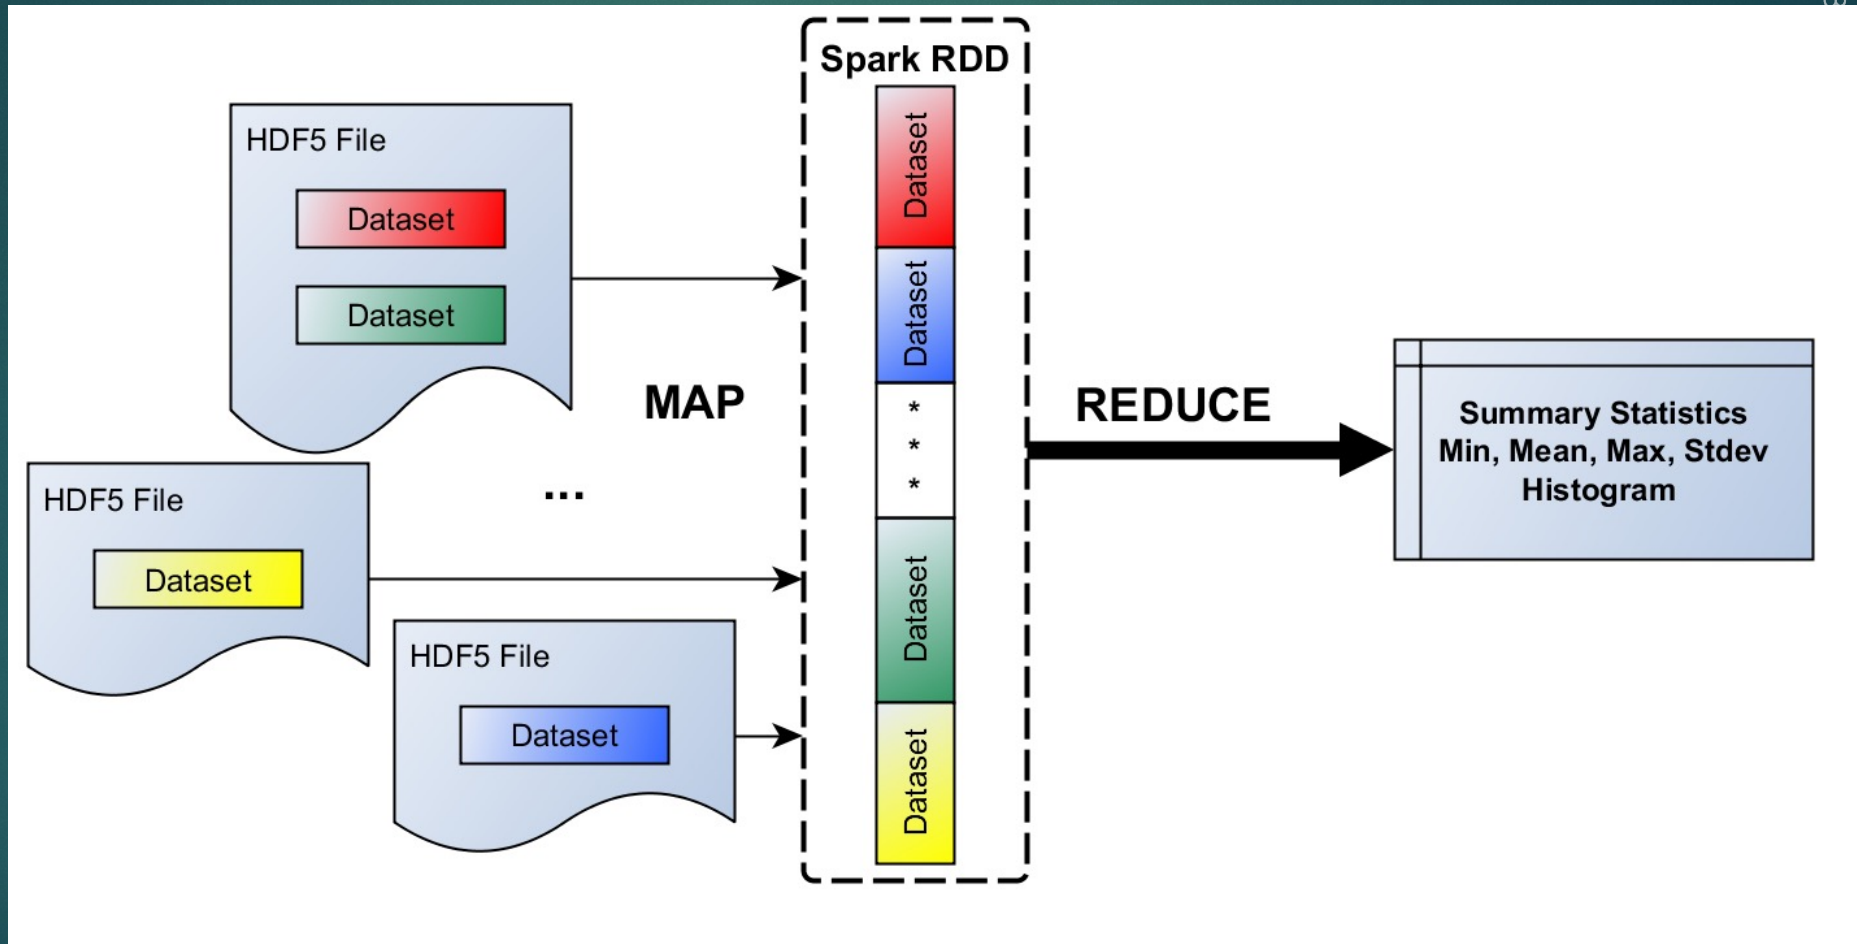
*The 2017 Nobel Prize in Physics has been awarded to LIGO co-founders. (Medal image: Wikipedia. Collage: LIGO Lab)*



# Meanwhile ...

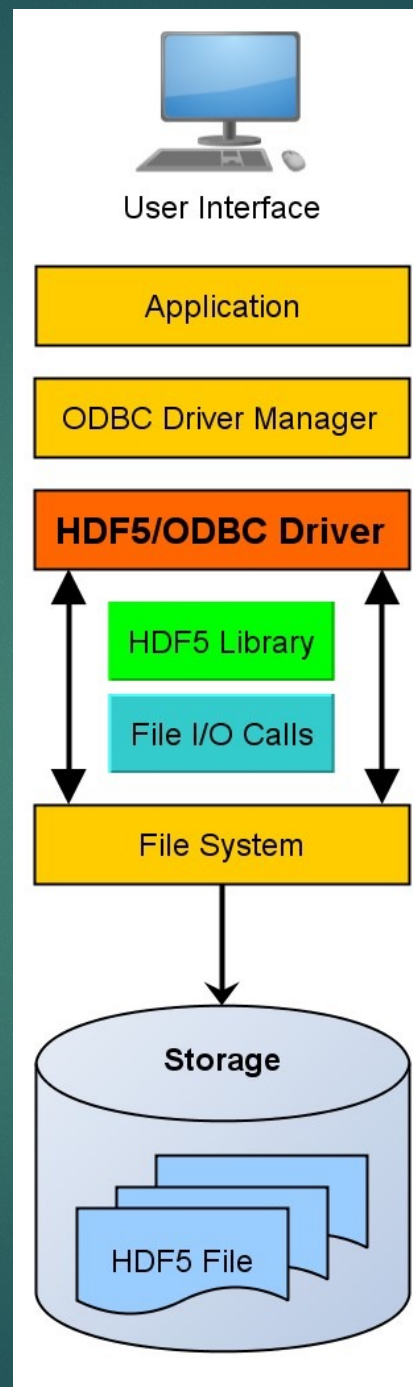
HDF5 CAPABILITIES – NOT THE LIBRARY

# 2015: Spark connector



HUG23  
August 16  
18

# ODBC driver PyHexad

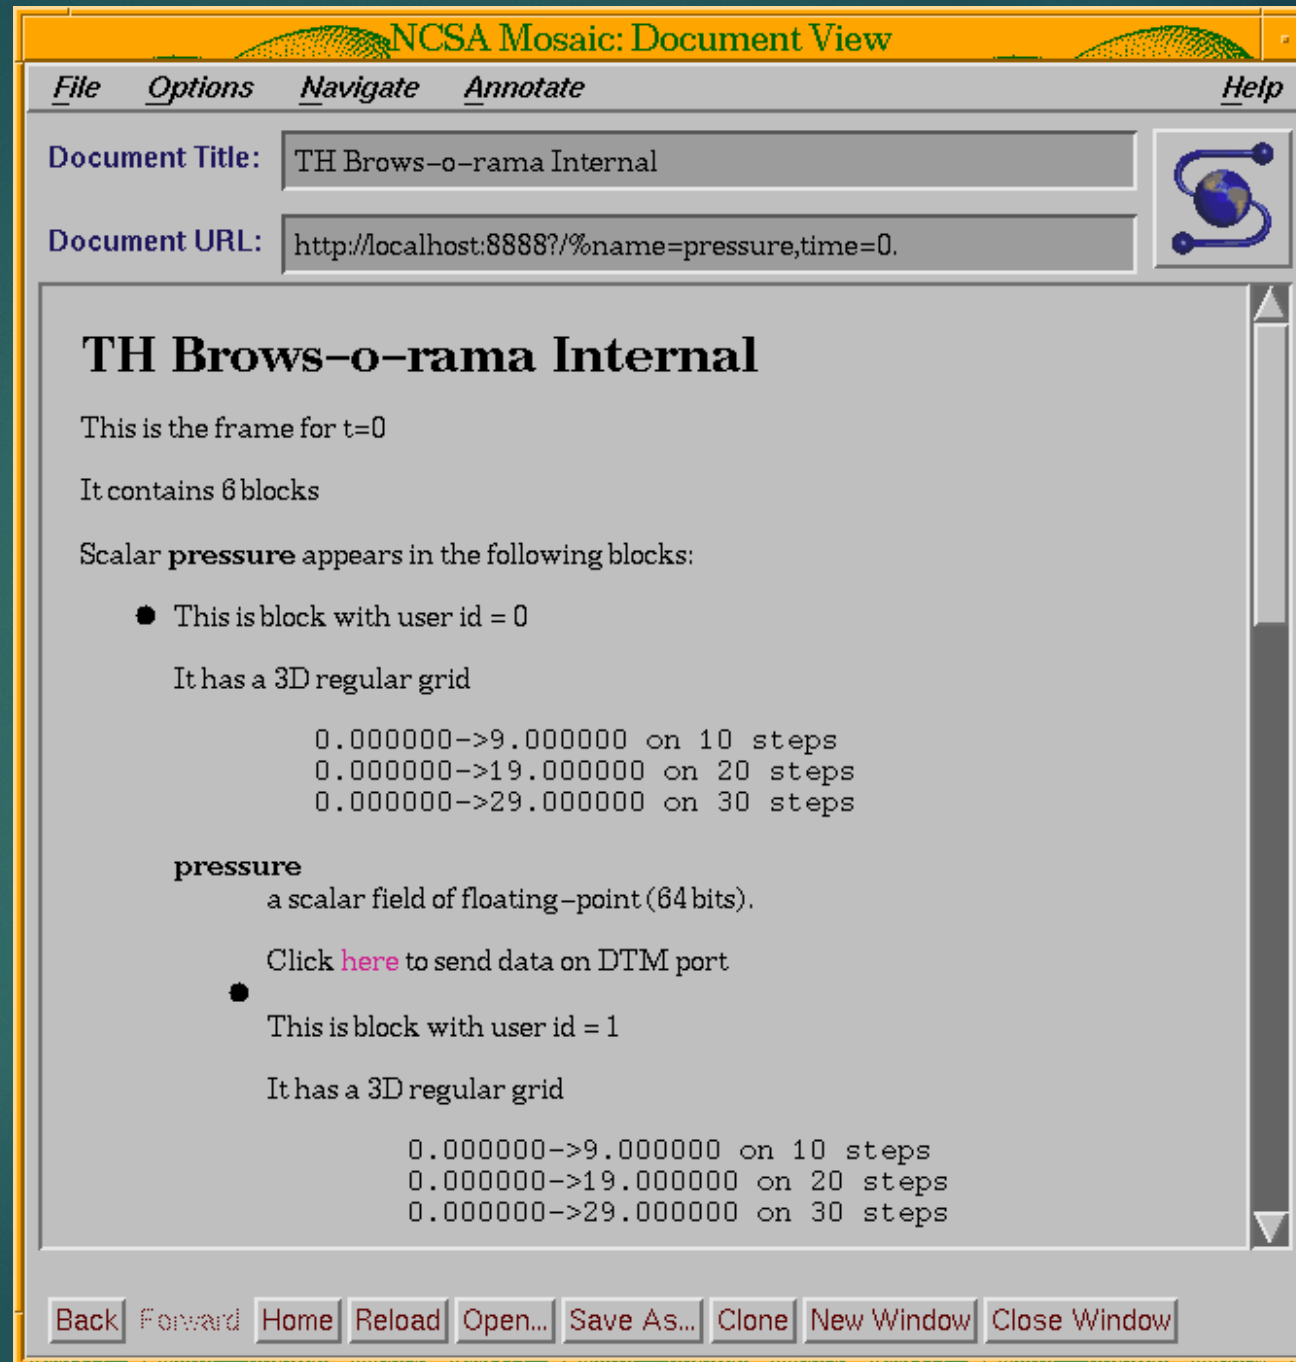


51

HUG23  
August 16-  
18, 2023

**PyHexad**  
Python-based **HDF5**  
**Excel add-in** can  
read or write data  
from Excel

# Web and cloud access over the years




# Scientific Data Browser (Horizon)

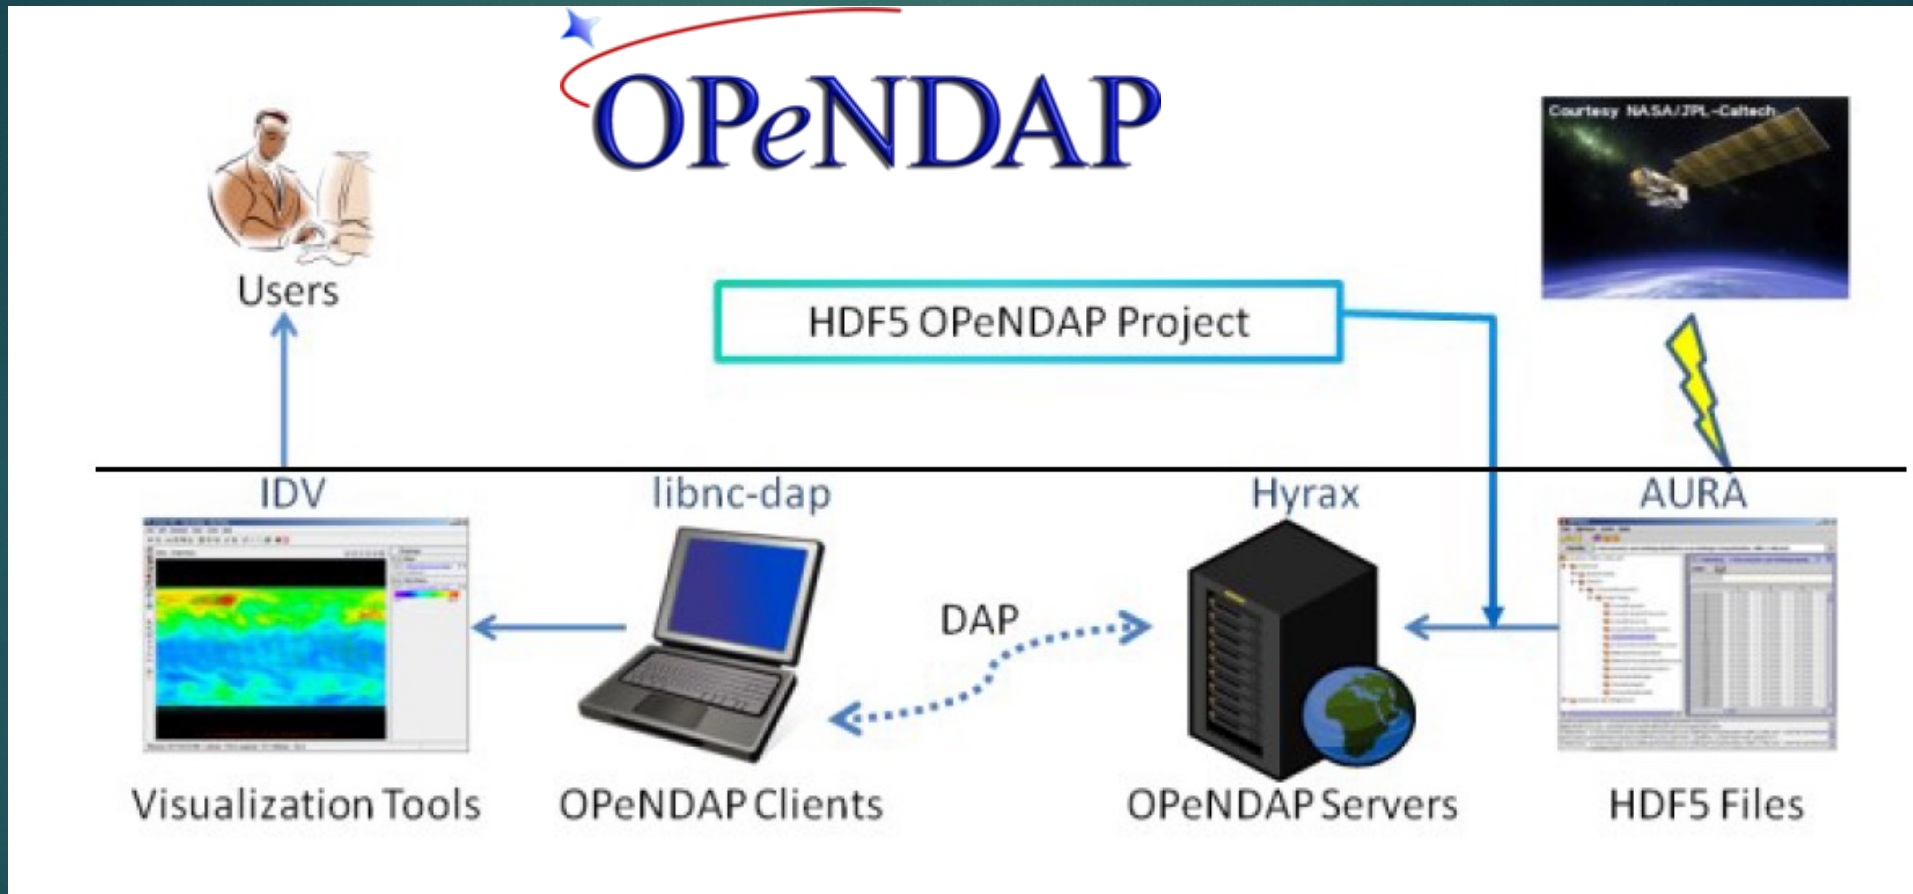
54

HUG23  
August 16-  
18, 2023

- ▶ **The scientific data browser (SDB)** is a web-based technology that enables scientists to stage and access scientific data in a conversational mode.
- ▶ Became DIAL.

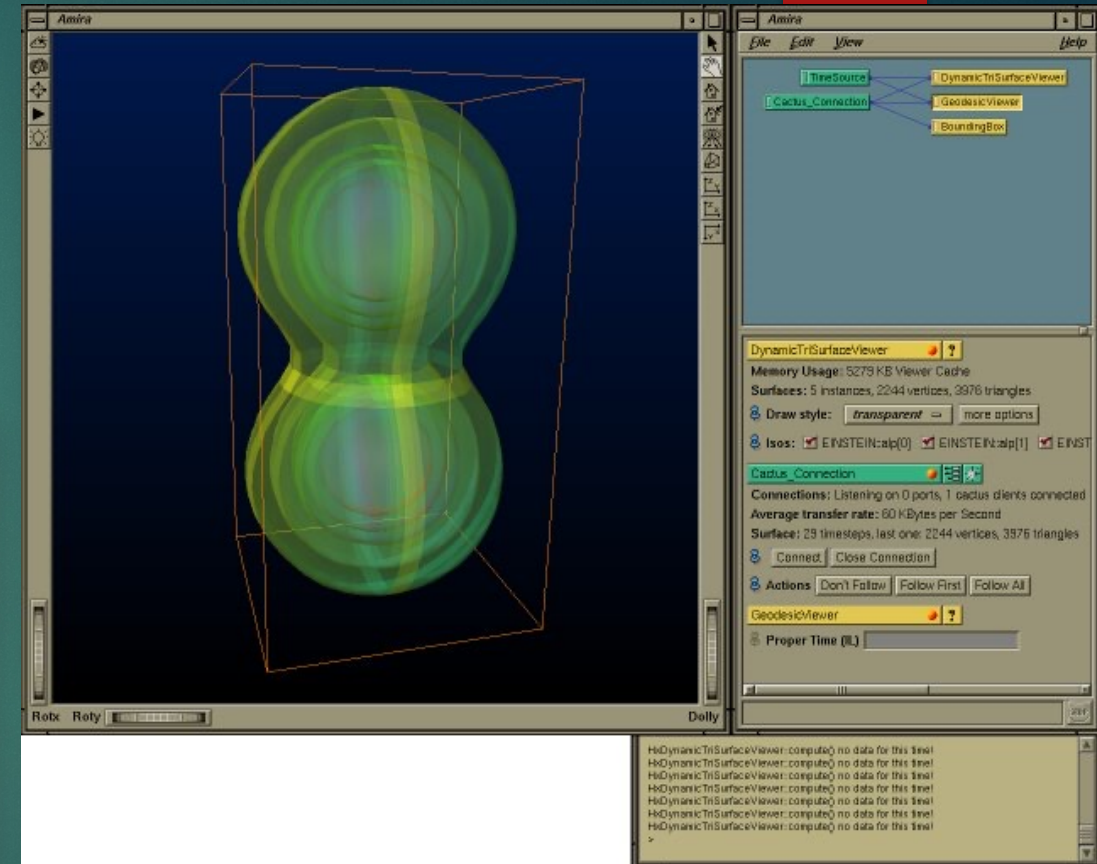
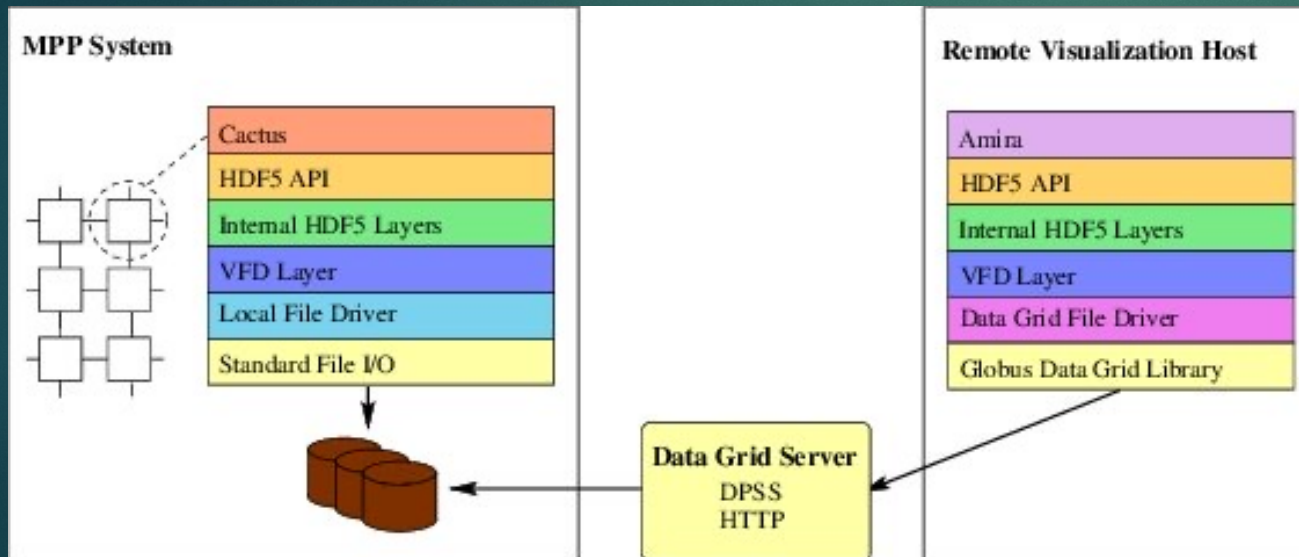


# OPeNDAP



# CACTUS VFD stream I/O driver

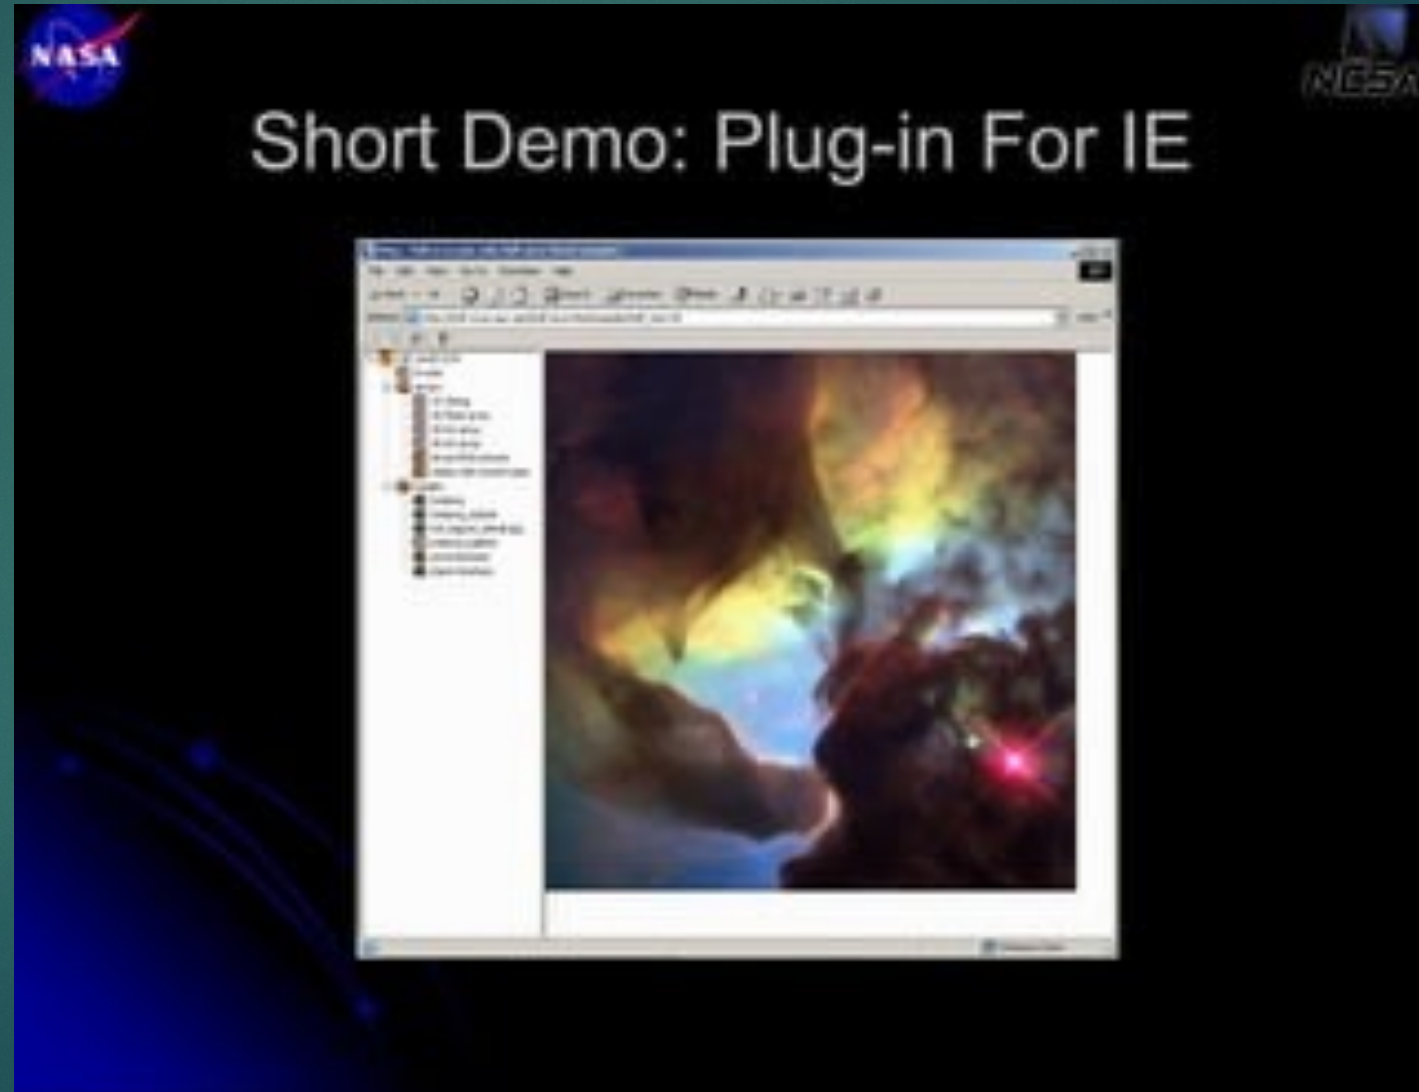
56





# 2006: Windows Browser plugin

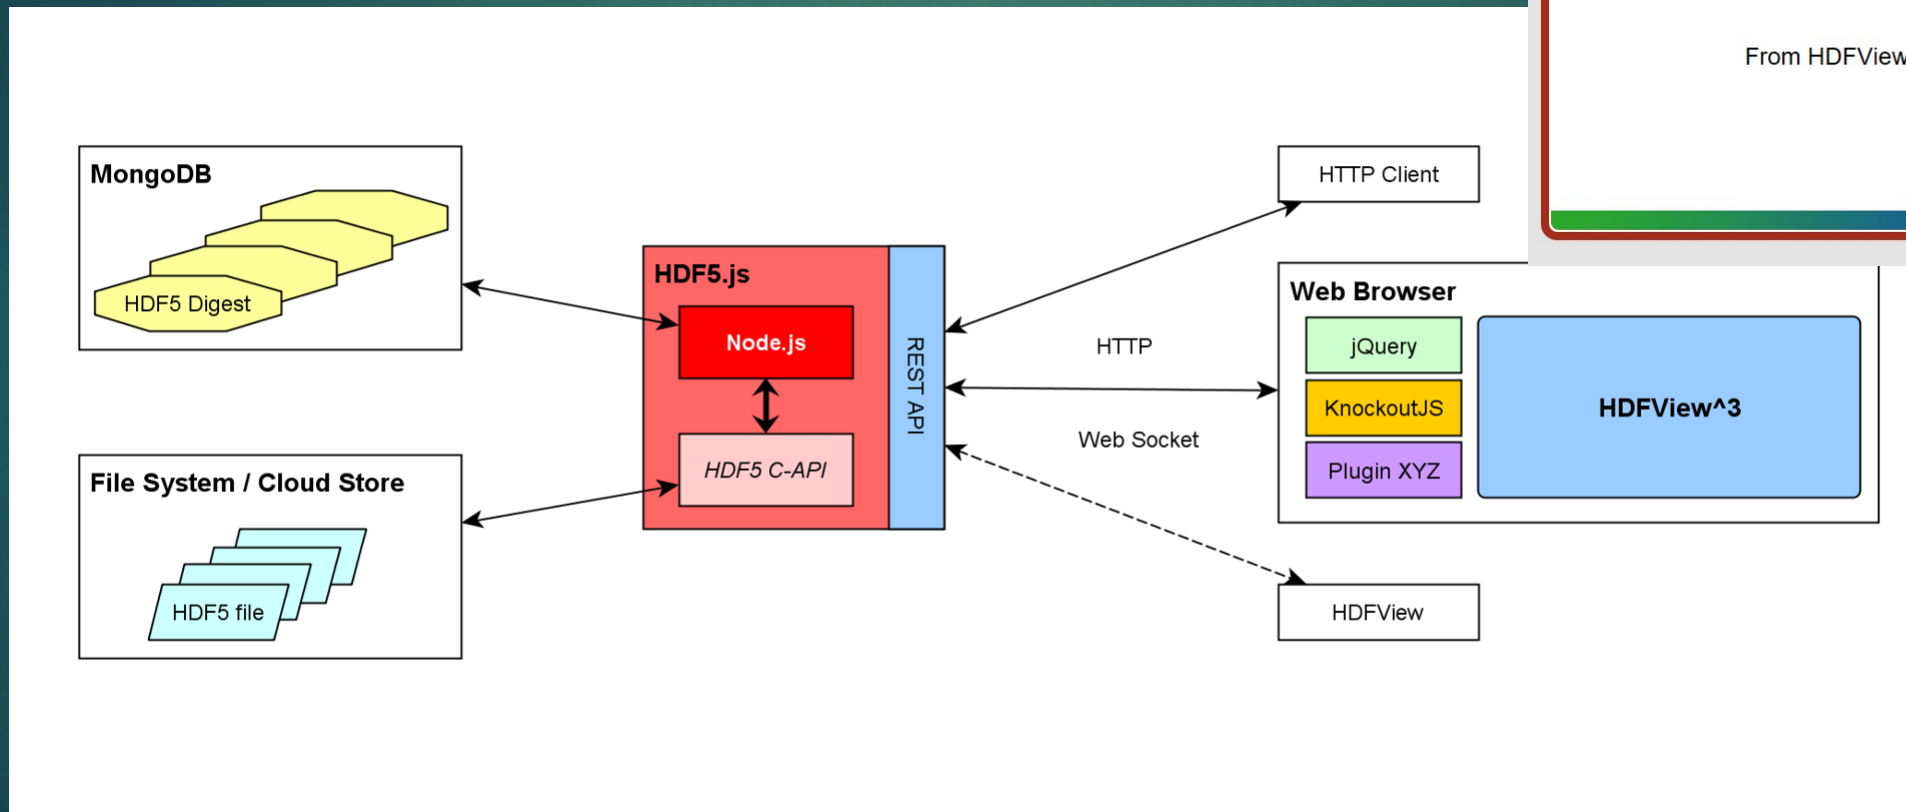
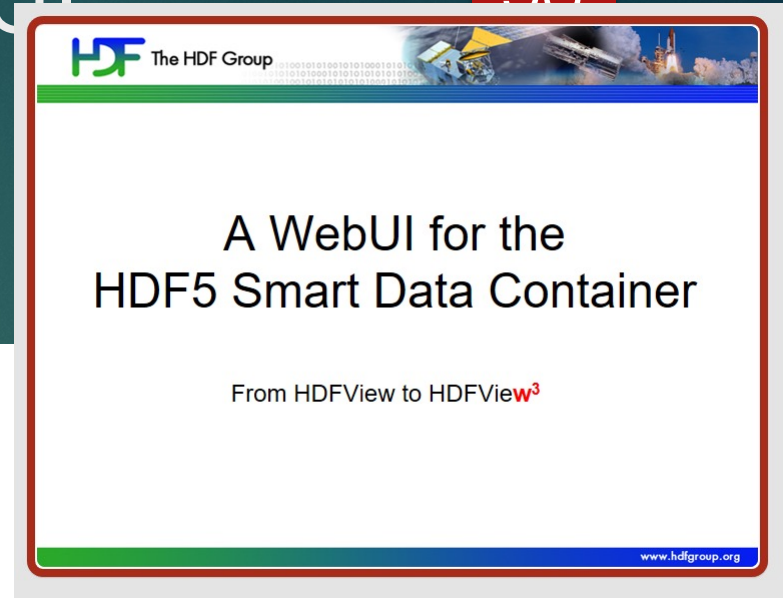
57



HUG23  
August 16-  
18, 2023

# 2013: RESTful HDF5 and Web UI

58



# 2014: HDF Cloud proposed

59



## What is HDF Cloud?

- Provide HDF functionality as as a Web Service (SaaS)
- A data repository provided by THG that clients can access via:
  - REST HTTP API
  - C/Fortran applications (using modified HDF5Lib)
  - Python scripts (using modified H5py)
  - Cmd Line tools (h5ls, h5import, h5diff, etc.)
  - Web Page
  - HDF Studio

# 2015: H5serv

60

HUG2  
August  
18, 2015

## h5serv

### About the Project

H5serv is a web service that enables HDF5 structured data to be created, updated, and read over HyperText Transfer Protocol (HTTP). H5serv is a reference implementation for the REST-based api as described in this paper: [http://www.hdfgroup.org/pubs/papers/RESTful\\_HDF5.pdf](http://www.hdfgroup.org/pubs/papers/RESTful_HDF5.pdf).

*From H5serv one-pager, June,  
2015*

# 2016: Highly Scalable Data Service proposed

61


HUG23  
August 16-  
18, 2023

The proposed project will support scalable concurrent access as well as eliminating the need to have data persisted on host drivers (i.e. a “share nothing” architecture).

# 2017: HSDS 0.1 released

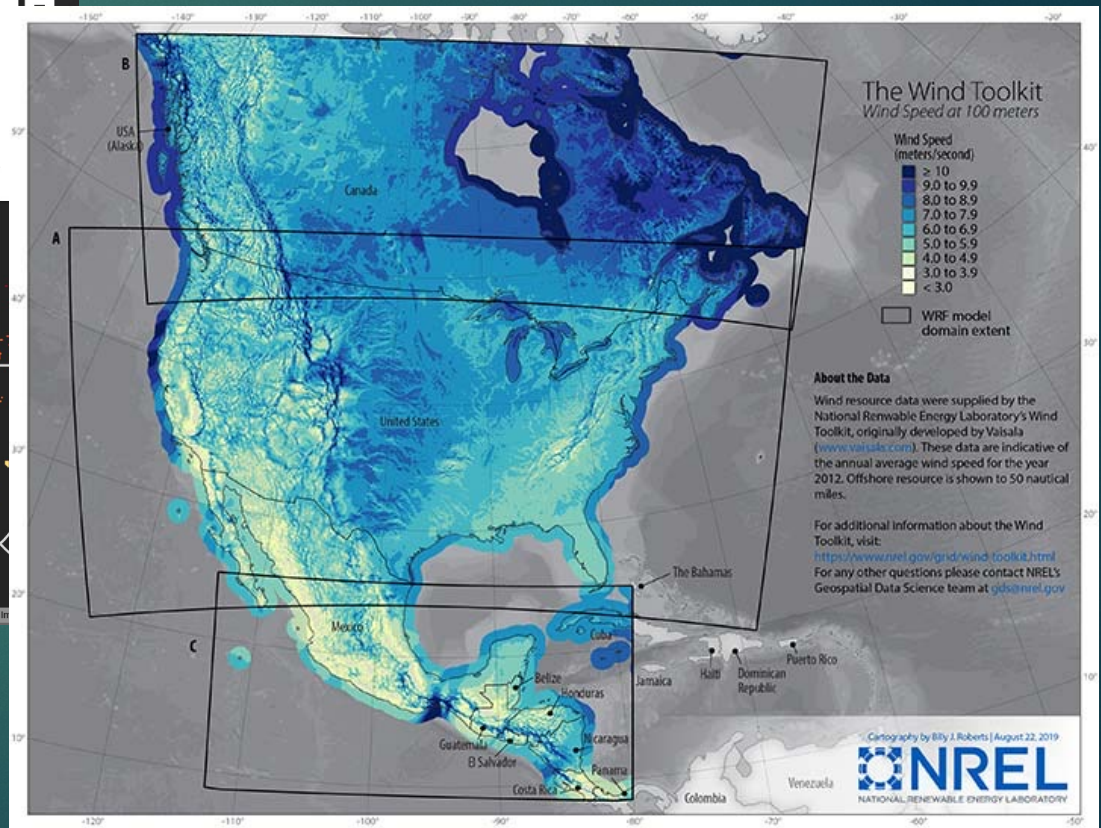
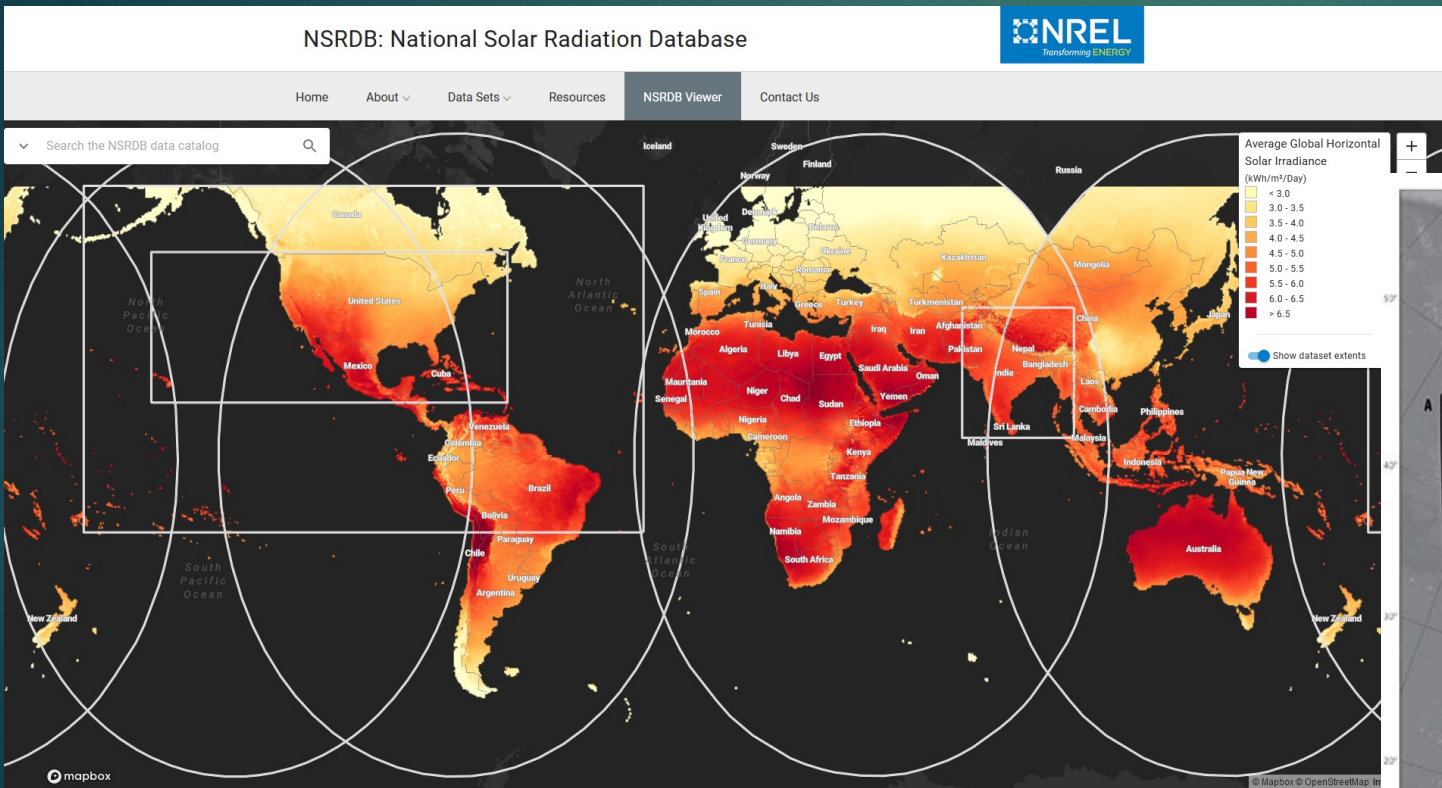
62

HUG23  
August 16-  
18, 2023



# Wind Integration National Dataset (WIND) 63 National Solar Radiation Database (NSRDB)

HUG23  
August 16-18, 2023



Thank you!