

# The HDF5-iRODS Module

## A Data Grid System for Object Level Access

Peter Cao, The HDF Group (xcao@hdfgroup.org)  
Michael Wan, San Diego Supercomputer Center (mwan@sdsc.edu)

### DATA CHALLENGES

- Simulations can generate very large and complex datasets.
- Researchers at different sites need fast access to both recent and historical data.
- Storage, networking, and compute platforms vary; some cannot handle full datasets.
- Frequently, only subsets of the data are of interest.


**HDF5, iRODS, and the HDF5-iRODS module address these challenges.**

### HDF5

Hierarchical Data Format Version 5 (HDF5) is a unique technology suite that makes it possible to manage extremely large and complex data collections. More than 600 organizations, over 200 types of applications, and millions of individuals are using HDF5. Terabytes of data are stored in HDF5 every day.

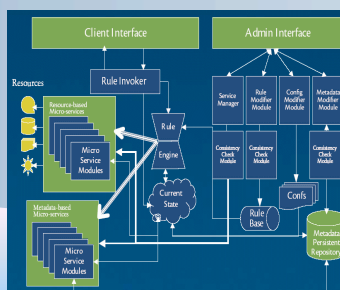
The HDF5 suite includes:

- A versatile data model
- A portable file format
- A library optimized for access time and storage space
- Tools and applications to manage, manipulate, view, and analyze data in HDF5 files



### iRODS

The Integrated Rule-Oriented Data System (iRODS) is a data grid system developed by the Data Intensive Cyber Environments (DICE) group. The most powerful feature of iRODS is the Distributed Rule Engine, which allows users to automate enforcement of management policies by applying iRODS Rules that control the execution of all data access and manipulation operations at distributed sites.



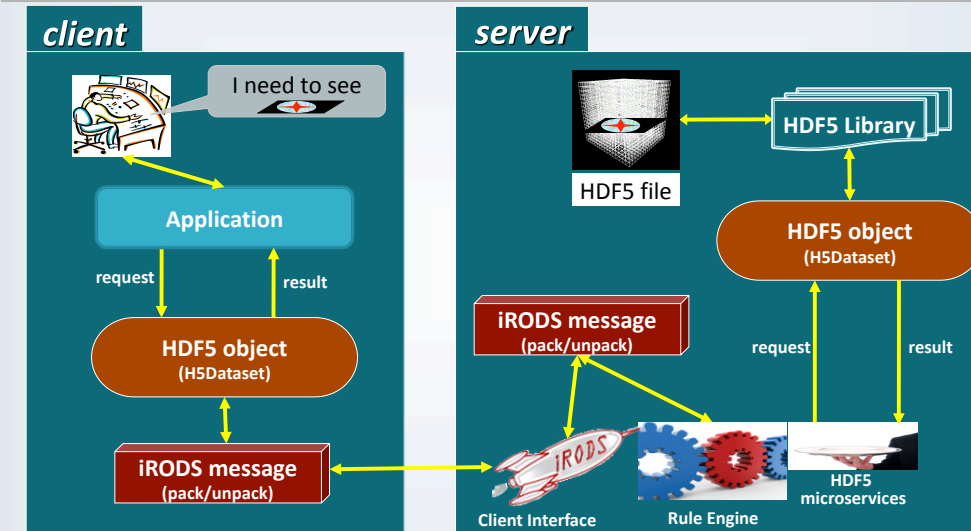
iRODS equips users to handle a full range of tasks:

- Manage distributed data
- Extract metadata
- Move data efficiently
- Share data securely
- Publish data in digital library
- Archive data for long-term preservation

### HDF5-iRODS MODULE

The HDF5-iRODS module components, together with HDF5 and iRODS, implement a client-server system that provides interactive and efficient access to HDF5 files managed by a remote iRODS server. Applications on the local machine use client functions to access specific data and metadata in HDF5 files stored remotely. Only the requested data and metadata are transferred to the local machine, not the entire file.

The HDF5-iRODS module includes two main parts: a set of HDF5 micro-services and a set of HDF5 objects. The HDF5 micro-services perform simple well-defined HDF5 tasks, such as *open a file, read from a dataset, read group attributes, or close a file*. The HDF5 objects, representing objects in HDF5 files, are used to specify requests from the client and to transfer results (data) from the server.

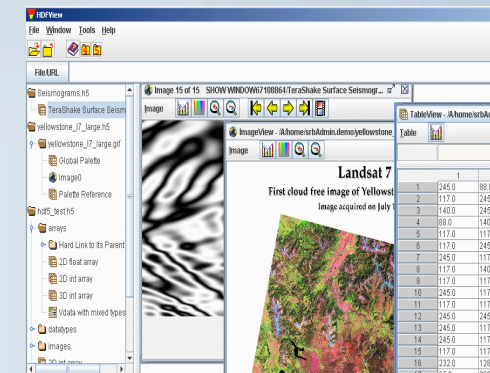


The figure above depicts how the HDF5-iRODS client-server system works. A user on a local (client) machine requests a slice of data from an HDF5 file managed by a remote iRODS server. To serve this request, H5DATASET\_OP\_READ is set in an H5Dataset (HDF5 object). The HDF5 object is then packed into an iRODS message and sent to the iRODS server. The server unpacks the message and checks the rule engine for matches. It finds and executes the associated HDF5 micro-service, msiH5Dataset\_read, which calls H5Dataset.read() to get the requested data from the HDF5 file. The HDF5 object, which contains the requested slice of data, is packed into an iRODS message and returned to the client, where it is unpacked and delivered to the application.

### BENEFITS of the HDF5-iRODS MODULE

- **Reduces storage needed on local machine.**  
Terabytes of data reside remotely; only small subsets are staged locally.
- **Facilitates data sharing.**  
Scientists can easily access updated data after a new simulation run, as well as prior results. Clients do not require the HDF5 library; all HDF5 calls are handled by the iRODS server.
- **Supports fast browsing of data objects.**  
Users can examine the structure of a file without loading the data content.
- **Provides rapid access to selected data content and metadata.**  
By transferring only the selected data content and metadata, access time is reduced.


### HDFView APPLICATION



HDFView, a visual tool for browsing and editing HDF files, was extended to use the HDF5-iRODS module so that users can view HDF5 files stored remotely on the iRODS server.

### islice APPLICATION

The islice tool uses the HDF5-iRODS module to extract a slice of data from a FLASH file stored remotely on the server. The slice of interest is transferred and stored on the local (client) system.



	Data Transfer Time	
	Network Bandwidth	Slice of interest
Full dataset	3.3 min	100 MB/sec
	33 min	10 MB/sec
	5.5 hours	1 MB/sec

### ACKNOWLEDGMENTS

This project was sponsored by CIP/NLADR, an NSF PACI Project in support of NCSA-SDSC collaboration, and managed by the CyberInfrastructure Partnership (CIP), a joint effort led by NCSA and SDSC. The work was carried out by The HDF Group and the SDSC SRB team. The ASC/Alliance Center for Astrophysical Thermonuclear Flashes at the University of Chicago provided the FLASH simulation data (HDF5 files) and other assistance. The islice tool was based on extract\_slice\_from\_chkpt, a slice tool previously developed by Paul Ricker (NCSA/UIUC).

### PROJECT WEBSITE

<http://www.hdfgroup.org/projects/irods>