Abstracts for the HDF5 Tutorials
HDF-EOS Workshop, November 30 – December 2, 2005

Introduction to HDF5 Data and Programming Models (1 hour, Elena Pourmal, Barbara Jones)

This Tutorial is designed for new HDF5 users. We will cover basic HDF5 Data Model objects and their properties; we will give an overview of the HDF5 Libraries and APIs, and discuss the HDF5 programming model. Simple C and Fortran examples will be used to illustrate HDF5 concepts. Participants will work with the Tutorial examples and exercises during the hands-on sessions.

Advanced HDF5 features (2 hours, Elena Pourmal, Barbara Jones)

This Tutorial is designed for the HDF5 users with some HDF5 experience. It will cover advanced features of the HDF5 library for achieving better I/O performance and efficient storage. The following HDF5 features will be discussed: partial I/O, compression and other filters including new n-bit and scale+offset filters, and data storage options. Significant time will be devoted to the discussion of complex HDF5 datatypes such as strings, variable-length, array and compound datatypes. Participants will work with the Tutorial examples and exercises during the hands-on sessions.

Introduction to Parallel HDF5 (1 hour, Albert Cheng, Kent Yang, Elena Pourmal)

This Tutorial is designed for the users who have exposure to MPI I/O and basic concepts of HDF5 and would like to learn about Parallel HDF5 Library. The Tutorial will cover Parallel HDF5 design and programming model. Several C and Fortran examples will be used to illustrate the basic ideas of the Parallel HDF5 programming model. Some performance issues including collective chunked I/O will be discussed. Participants will work with the Tutorial examples and exercises during the hands-on sessions.

Introduction to HDF5 tools (30 minutes, Peter Cao)

This Tutorial is designed for anyone who needs to work with data stored in HDF5 files. The Tutorial will cover functionality and handy features of the HDF5 utilities h5dump, h5diff, h5repack, h5stat, h5repart. It will also cover HDF5 Java browsing and editing tool HDFView.

New features in HDF5 1.8.0 (2 hours, John Mainzer, Raymond Lu, Pedro Nunes, James Laird)

This Tutorial targets HDF5 application developers and anyone who is interested in the new HDF5 Library features. The following new features available in 1.8.0 will be discussed:

HDF5 cache

New metadata cache to improve performance and memory usage in the HDF5 library has been implemented. We will give overview of the implementation and introduce new APIs that can be used by HDF5 application programmers to analyze and tune cache performance.

Error handling

Error handling mechanism was redesigned to give HDF5 application developers better control over the HDF5 Library error stack. New Error APIs that allow integration of application's error reporting with the HDF5 error reporting will be introduced in this part of Tutorial.

Dimension Scales

We will give overview of Dimension Scales data model and library APIs and discuss how this feature can be used by the HDF5 applications.

Packet APIs

We will introduce new Packet APIs that provide efficient way of creating, storing and retrieving fixed length and variable length records in an HDF5 file

Datatype conversions

This part of Tutorial discusses how HDF5 Library implements and handles data conversions between arithmetic data types of the same or different classes.