

# **New Features in HDF5 1.8.0 - Group Revisions**

HDF and HDF-EOS Workshop IX

November 30, 2005

# Outline of Topics

- Overview of HDF5 groups
- Feature requests
- Problems with current implementation
- Changes and new features
- Backward/forward compatibility

## Overview: Application Level

- Group is a container for links to HDF5 objects
- Identified by a local or global path in the file (ASCII string)
- Iteration over links is in alphanumeric order, by link name
- Groups with a few/no links create a substantial overhead in the file (e.g. empty group occupies 832 bytes)

# Overview: File Format Level

- Group object header information:
  - B-tree address
    - Quickly locates links in group
  - Local Heap address
    - Stores variable-sized information that doesn't fit in B-tree records
      - Link names
      - Soft link “destination” values (i.e. path names in file)
- Note:
  - ASCII character set for link name is assumed
  - No information on number of links in group
  - No information on links' creation order
  - Links must reference objects in same file

# Feature Requests

- Access links in a group according to their creation order
- Non-ASCII link names
  - UTF-8, etc.
- Links to objects in other HDF5 files
  - “External Links”
- Checksum HDF5 file metadata
  - (Not planned to be implemented in 1.8.0 release)

# Deficiencies in Current File Format

- Have to traverse entire B-tree to find number of links and to find  $n^{\text{th}}$  link
- Inefficient local heap storage method
  - Has to be recopied in file when grows
  - May “abuse” metadata cache when large
- Inefficient storage representation for groups with few/no links
  - Each link needs ~64 bytes of storage, but empty group requires 832 bytes and group with 1 link requires 1160 bytes
  - Groups with more than ~8 links storage very efficiently
- If no link ordering required, extra overhead of providing alphanumeric ordering

## New Features: Overview

- Allow multiple indices on links: iterate over links in creation, alphanumeric, or “no” order
- New B-tree implementation to efficiently access  $n^{\text{th}}$  link in alphanumeric or creation order
- Support non-ASCII character sets for link names
- New non-contiguous heaps for storing links
- Compact representation for groups with few links
- Links to objects in other HDF5 files

# New Features: API Changes

- Add new group API routines, with additional parameters:
  - **Creation property**
    - Parameters to control group size
    - Indices to create on links
  - **Access property**
    - Index to use for operation
- “Object copy” API routine
- Add new “Link” object and corresponding APIs
  - **Supports non-ASCII character set encodings**
  - **Allows “anonymous”, temporary objects to be created**
  - **Helps HDF5 Programming Model follow HDF5 Abstract Data Model**

# New Features: File Format Changes

- Completely revised group object header, storing:
  - Number of links in group
  - Multiple B-trees for different indices on links
  - Creation properties
    - Storage transition points, etc.
    - Type of indices available on links
  - When group in “compact” form, links are stored in object header
  - Uses revised B-tree and heap storage for better indexing and storing of links
- Updated information about each link:
  - Link type
    - Hard, soft, external
  - Link name character-set encoding
  - Link creation time

# Backward/Forward Compatibility

- File Format
  - New library will read old files
  - New features of 1.8.\* (e.g. different indexing or non-ASCII character set names) cannot be used on old files
  - Old library will not read new files
- APIs
  - Have not decided yet
    - Add new APIs and leave old APIs intact
    - Modify current ones and add more if necessary