

Improving the performance of H5Gget_info_by_idx and H5Lget_info_by_idx

17 November 2008

If you find that the function `H5Gget_info_by_idx` or `H5Lget_info_by_idx` is slow for the new format file in the release 1.8 of the library, you may want to adjust the metadata cache size to improve the performance. The degree of the performance improvement depends on the machine you are using. On one machine, it can be 50% faster than the old format file in the release 1.6 of the library. On another machine, it can be 10% faster. In the attached C programs, you can find out how to enable the new format file and change the size of metadata cache. Below are our recommended sizes of the metadata cache for the numbers of groups at one location.

Number of groups	recommended metadata cache size
10,000	1MB
100,000	4MB
1,000,000	32MB

For example, if there are 100,000 groups under the root group, try to set the size of the metadata cache to 4MB. The original size of the metadata cache in the library is 1MB. You do not need to adjust this size if there are 10,000 groups.

Attachment 1: To create a number of groups under the root group in the new file format.

```
#include<hdf5.h>

int main(void)
{
    hid_t file, grp, fapl;
    char *base = "group";
    char group_name[16];
    int i;

    /* Indicates using the new file format */
    fapl = H5Pcreate(H5P_FILE_ACCESS);
    H5Pset_libver_bounds(fapl, H5F_LIBVER_LATEST, H5F_LIBVER_LATEST);
```

```

        file = H5Fcreate("/tmp/slu/file.hdf5", H5F_ACC_TRUNC, H5P_DEFAULT,
fapl);

        /* Create some groups under the root group */
        for(i=0; i<123456; i++) {
            sprintf(group_name, "%s_%d", base, i);
            grp = H5Gcreate2(file, group_name, H5P_DEFAULT, H5P_DEFAULT,
H5P_DEFAULT);
            H5Gclose(grp);
        }

        H5Fclose(file);
        H5Pclose(fapl);
        return 0;
    }

```

Attachment 2: To adjust the size of the metadata cache and call `H5Gget_info_by_idx` using the file created in Attachment 1.

```

#include<hdf5.h>

int main(void)
{
    hid_t file;
    H5G_info_t grp_info;
    H5AC_cache_config_t config;
    int i;

    file = H5Fopen("/tmp/slu/file.hdf5", H5F_ACC_RDONLY, H5P_DEFAULT);

    /* Adjust the size of metadata cache */
    config.version = H5AC__CURR_CACHE_CONFIG_VERSION;
    H5Fget_mdc_config(file, &config);
    config.set_initial_size = 1;

```

```
config.initial_size = 4*1024*1024;

config.max_size = 16*1024*1024;

H5Fset_mdc_config(file, &config);

    H5Gget_info_by_idx(file, ".", H5_INDEX_NAME, H5_ITER_INC,
(hsize_t)123450, &grp_info, H5P_DEFAULT);

    H5Fclose(file);

    return 0;

}
```