

CMD & 1.6.x Client Requests Changes

Lai Siyao <lsy@clusterfs.com>

2006.03.12

1 Introduction

This document describes how 1.6.x client choose different requests format to interoperate with both 1.6 and 1.8 server. Understanding of cmd request layout is needed.

2 Requirements

1.6.x: release that introduces clients that are upward compatible with 1.8 md server. At the connect time 1.6.x client determines whether server supports fids (this is implemented through OBD_CONNECT_FID connect flag). If fids aren't supported (i.e., server is 1.6), 1.6.x client behaves like normal 1.6 client. If fids are supported, cmd3 formatted requests are used to talk with server (This might be a good point to introduce req_capsule interface on client, because it would allow better unification of code between 1.6 and 1.6.x modes.).

3 Functional specification

1.6.x client needs to interoperate with both 1.6 and 1.8 md server, and during this change we'd better remove all explicite use of buffer offset:

- introduce md_ops for 1.6.x client, and there are two categories of md_ops for mdc: mdc_md_ops_16 and mdc_md_ops_18.
- introduce req_capsule interface for 1.8 client and 1.6.x client (only those for 1.8 md server)
- get rid of lustre_shrink_reply(), use req_capsule instead.

4 Use cases

4.1 1.6.x client getattr from 1.6 server

Client calls `md_getattr` to getattr from md server, which finally calls `mdc_getattr_16`.

4.2 1.6.x client getattr from 1.8 server

Client calls `mdc_pack_req` to prepare request and reply buffers according to `MDS_GETATTR` request layout, then calls `req_capsule_client_get(pill, &RMF_MDT_BODY)` to get `mdt_body` and fill it with proper values. After receiving reply, access all reply buffers via `req_capsule_server_get()`.

4.3 1.8 server shrink reply

`mdt_shrink_reply` calls `req_server_shrink()` for all buffers after `mdt_body`.

5 Logic specification

5.1 md_ops for 1.6 client mdc

5.2 req_capsule for 1.6.x and 1.8 client

- add `mdc_pack_req()` to pack request and reply buffers according to the layout.
- use `req_capsule` to access all request and reply buffers in client.

5.3 get rid of lustre_shrink_reply

Introduce `req_capsule_server_shrink()` to shrink all buffers after `mdt_body`.

6 State management

N/A

7 Alternatives

- introduce `req_capsule` for 1.6.x client code to 1.6 md server, however these code will be deserted after 1.8.

8 Focus for inspections

- all places use req_capsule to access buffers now.