

Data On MDT Scope Statement

Introduction

The following scope statement applies to the Data on MDT Design project within the Technical Proposal by High Performance Data Division of Intel for OpenSFS Contract SFS-DEV-003 signed Friday 23rd August, 2013.

Project Goals

Lustre performance is currently optimized for large files. This results in additional RPC round-trips to the OSTs, which hurt small file performance. This project aims to correct this deficiency by allowing the data for small files to be placed on the MDT so that these additional RPCs can be eliminated and performance correspondingly improved. Used in conjunction with DNE, this will preserve efficiency without sacrificing horizontal scale.

Users or system administrators will set a layout policy that locates files to be stored entirely on an MDT. Files that grow beyond this size will be migrated onto OSTs.

In Scope

- Implement a new layout for files with data on MDT.
- Client IO stack must allow IO requests to an MDT device.
- Design of a tunable for specifying size limit for files on a MDT.
- A mechanism to perform migration from MDT to OST. If you start writing to a file that was layed out on the MDT and it becomes too large, it will be migrated.
- Explicitly allocating files on the MDT by default directory striping.
- A discussion of interoperability of old clients and servers with Data on MDT.

Out of Scope

- Automatically locating small files on the MDT after the file has started to be written.

Project Constraints

- Mike Pershin is the engineer with the expertise best suited to this task.

Key Deliverables

- Signed-off Milestone document for the project phases:
 - Solution Architecture.
 - High-level Design.
 - Implementation assessment