

Inode Iterator and OI Scrub Scope Statement

Review

This document was presented for review on 2011-09-21. The document was accepted on 2011-10-05.

Introduction

The following scope statement applies to the Inode Iterator and OI Scrub Scope Statement project within the SFS-DEV-001 contract/SOW dates 08/01/2011.

Problem Statement

Over the lifetime of an active filesystem, inconsistencies may develop as a result of a restore from file-level backup or non-graceful shutdown. Currently, the Lustre 2.x series is unable to do a file-level MDT backup and restore, because restoring a file-level backup will give files new inode numbers and cause the Object Index (OI) to contain an incorrect FID-to-inode mapping table. The Inode Iterator and IO Scrub project will create a tool to traverse all inodes in an MDT and verify their corresponding entries in the Object Index (OI) are correct. Incorrect entries will be resolved automatically. The Inode Iterator and OI Scrub tool will increase available backup options and provide the foundation for additional distributed filesystem checks for later phases of this project.

Project Goals

1. Demonstrable file level backup and restore.
2. Characterize the performance of fsck without filesystem load.
3. Characterize the performance of fsck with filesystem load.
4. Code landed in WC-Lustre master branch.

In-Scope

- Kernel-space Inode Iterator and OI Scrub tool.
- User-space control code for Inode Iterator and OI Scrub tool.
- Administrative documentation in the form of a man page and update to WC-Lustre 2.x manual.

Out of Scope

- While the design and implementation will strive to be agnostic to the back-end OSD filesystem, only ldiskfs OSDs will be tested and landed for this phase.
- Detection and resolution of internal filesystem inconsistencies is not within scope of this phase.

Project Constraints

- Fan Yong is the only engineer with the correct expertise available for this work.
- Fan Yong does not have root access to the Hyperion test platform.

Key Deliverables

- Signed Milestone documents for project phases:
 - Solution Architecture.
 - Implementation & Test.
 - Acceptance Testing (OpenSFS executed).
- Test Plan.

- Source code that meets feature requirements and runs with WC-Lustre 2.x on customer's site.
- Source code for new test cases.
- inode iterator and IO scrub code landed in the Master WC-Lustre 2.x.

Key Milestones

Fan Yong becomes available from 7th Sept 2011.

- Scope Statement Delivery 2011-09-21
- Solution Architecture 2011-10-05
- Implementation 2011-11-09
- Acceptance Testing 2011-11-30
- Demonstration and close 2011-12-14

Glossary

WC-Lustre – Whamcloud community release of Lustre.