

Parallel Directory Operations Scope Statement

 This document was presented for review on 2011-09-16. This document was signed off on 2011-09-28.

Introduction

The following scope statement applies to the Parallel Directory Operations Scope Statement project within the SFS-DEV-001 contract/SOW dates 08/01/2011.

Problem Statement

Many HPC workloads create numerous files in a single directory and contention of the single directory lock degrades performance. The Lustre filesystem provides parallel locking at the directory level for name entry operations (add, remove and lookup). However, the underlying Idiskfs only provides a single lock for a directory. The Parallel Directory Operations project will replace a single directory lock with a parallel directory lock in Idiskfs. This project will deliver a significant and measurable performance improvement to name entry operations.

Project Goals

1. Minimize new code in Idiskfs to support Parallel Directory Operations.
2. Additional Idiskfs code will be simple to maintain.
3. WC-Lustre build that includes file-level locking for add, remove and lookup name entry operations.
4. Benchmark add, remove and lookup name entry operations for WC-Lustre 2.x baseline.
5. Show performance improvement for add, remove and lookup name entry operations with Parallel Directory Operations WC-Lustre 2.x.
6. Code landed on WC-Lustre master branch.

NOTE: Parallel Directory Operations work takes place independent of Project Orion and will land independently.

In-Scope

- Parallel Directory Operations code development will take place against WC-Lustre 2.x baseline.
- Modifications to Idiskfs to provide Parallel Directory Operations.
- Modifications to OSD to exploit Parallel Directory Idiskfs.
- New locking behavior documented in the WC-Lustre 2.x manual.
- Parallel Directory Operations supported on all Linux distributions that WC-Lustre supports.

Out of Scope

- Only EXT4 back-end filesystem will be supported.
- As a rule, patches to EXT4 will not be prepared for up-stream inclusion.

Project Constraints

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

Key Deliverables

- Signed Milestone documents for project phases:
 - Solution Architecture.
 - High-Level Design.

- 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.
- Benchmark test results.

Key Milestones

- Scope Statement Delivery 2011-09-14
- Solution Architecture 2011-09-28
- High-Level Design 2011-10-19
- Implementation 2011-11-23
- Acceptance Testing 2011-12-07
- Demonstration and close 2011-12-21
- Land on WC-Lustre Master 2011-12-31

Glossary

WC-Lustre – Whamcloud community release of Lustre.