

ORACLE®

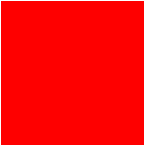


ORACLE®



Lustre Community

Dan Ferber
Lustre Community Manager
April 16, 2010



The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



What is Oracle's Open Source Strategy for Lustre?

- The Core Lustre file system technology, including both Lustre 1.8 and Lustre 2, will remain open source, licensed under GPL 2.0
- Oracle intends for Sun Lustre Storage Systems built with Lustre 2 to include both the core file system and other components that may or may not be open source



Is Oracle going to support a Lustre community?

- Yes, Oracle intends to continue hosting the Lustre community
- Oracle is hosting the annual Lustre User Group meeting, scheduled for April 14-16 2010; see <http://lug2010.org/> for more information
- Oracle intends to continue hosting and improving the community web site, public git source code repository with canonical Lustre releases, mailing lists, and Bugzilla issue tracking system
- The Lustre Group will continue to accept software defect reports from the community and address them to improve the quality of Lustre for all
- The Lustre Group intends to accept community patches that are submitted based on established contributions guidelines for inclusion in the canonical release branches



Lustre Community Program

- Promote growth of Lustre features, performance, quality, and stability through community collaboration
 - Feature and fix contributions
 - Testing, bug finding, and use cases
 - Enhanced documentation
 - Best practices
 - Increased Lustre knowledge in community
 - Technical discussions, workshops, user groups
 - Feedback



Lustre Community

Sample Program Activities

- Technical information collaborations (partial list only)
 - Oak Ridge Lustre Development and Workshops
 - CEA HSM
 - DARPA HPCS
 - NRL changelogs, replication, and WAN
 - LLNL and Sandia - Multiple Collaborations
 - Indiana University
 - Annual Lustre User Group each Spring in United States
 - Autumn Workshop in Europe
 - ISC and Supercomputing
 - lustre-discuss, lustre-devel, lustre.org, and git repo



Lustre Community Development

- Lustre Knowledge
 - Lustre operations manual
 - ORNL Lustre internals manual
 - Lustre architecture documents and presentations
 - Lustre Internals Documentation (LID)
- Other Lustre Community Resources
 - ORNL (and other) Lustre workshop slides and papers
 - Lustre User Group Slides and Videos
 - Bugzilla, and Lustre downloads
 - Searchable lustre-discuss archive
 - Lustre Quick Start
 - Lustre technical papers

lustre.org Lustre Community Site



[Log In / Create Account](#) [Help](#)

Search Lustre.org

High Performance and Scalability

For the world's largest and most complex computing environments, the Lustre™ file system redefines high performance, scaling to tens of thousands of nodes and petabytes of storage with groundbreaking I/O and metadata throughput.

More on Lustre performance, service, and support at the [Lustre product page](#).

Lustre 2.0 Beta-1 Available

We are pleased to announce that Lustre 2.0 Beta-1 is available for download. This is the first Beta step in a series of milestone based pre-releases as we move towards Lustre 2.0 GA. New milestone releases will be planned for every 4-6 weeks.

Lustre User Group 2010

Plan to join us for LUG 2010 April 14-16 at the beautiful [Seascape](#) resort and conference center on Monterey Bay, California. This year's event will feature a one day Lustre User Advanced seminar, followed by two days of informative presentations by Lustre developers and users. [Registration](#) information and the LUG 2010 [agenda](#) are now available.

Lustre 1.8.2

Lustre 1.8.2 is now GA and available for download. Lustre 1.8.2 introduces support for RHEL 5.4, offers several minor improvements (including 16TB LUN support), and provides a number of bug fixes, including the short read-ahead bug for 32-bit clients (see [Bug 21506](#)). Learn about the 1.8 family of features - [Adaptive Timeouts](#), [OSS Read Cache](#), [OST Pools](#) and [Version-based Recovery](#) and why you should upgrade.



Send feedback on the lustre wiki to: lustre-wiki-feedback@sun.com

DOWNLOAD

Find out about released and pre-release versions of Lustre and download free open source software.

[Get Lustre](#) [Lustre Support Matrix](#)

LEARN

Find out about Lustre current and upcoming features, publications and training.

[Lustre 1.8](#) [Lustre 2.0](#)

USE

Access information about installing, configuring, administering and troubleshooting Lustre.

[Lustre Quick Start](#) [Install & Use Lustre](#)

CONTRIBUTE

Access guidelines for contributing code and testing Lustre along with developer resources and tools.

[Contribute to Lustre](#) [Developer Resources](#)

GET INVOLVED

Find out about the Lustre User Group, community projects, Lustre Centers of Excellence and more.

[LUG](#) [LCEs](#) [Mailing Lists](#)



Lustre.org Wiki Enhancements

- Navigation updated and topics refreshed to provide easier access and more complete information
 - DOWNLOAD – Official Lustre software downloads and pre-release versions (Lustre 2.0), Lustre interop and support matrix
 - LEARN – Lustre features (current and upcoming), publication and presentation materials, Lustre training
 - USE – Lustre installation and configuration, administration, troubleshooting, user resources
 - CONTRIBUTE – Find a project, develop, debug, test and submit code. Also contains developer resources
 - GET INVOLVED – Lustre community events and development projects, LCEs and third-party contributions

Accessing Lustre Code



[Log In / Create Account](#) [Help](#)

Search Lustre.org

[Article](#) [Discussion](#) [View Source](#) [History](#)

[Main Page](#)

[Download](#)

[Learn](#)

[Use](#)

[Contribute](#)

[Get Involved](#)

Accessing Lustre Code



(Updated: Jan 2010)

NOTICE: The transition from CVS to Git took place on Monday, December 14. For more information about the transition, see the [Git Transition Notice](#). For details about how to migrate to Git, see [Migrating to Git](#).

We welcome and encourage contributions to the development and testing of a more robust, feature-rich Lustre™. You can obtain the latest bleeding-edge Lustre source code by anonymous Git access.

```
git clone git://git.lustre.org/prime/lustre.git
```

Note: For more information about using Git, including tutorials and guides to help you get started, see the [Git documentation](#) page. For descriptions of the commands you are most likely to need, see the [Commands section at the bottom of this page](#).

See [Contribute](#) for more information about developing, testing, and submitting a patch to the Lustre code.

Note: If you have questions or experience problems, send email to the [Admins](#).

For more information about Git, see the [Git home](#) page.

Naming conventions

Stable development branches are named `b{major}_{minor}` (for example, `b1_6` and `b1_8`). Even-numbered minor releases are considered stable releases. Odd-numbered minor releases correspond to alpha and beta releases and will sometimes be given `v{major}_{minor}_{patch}` tags to provide a point of reference for internal and external testing.

A release branch is created an official release to isolate it from further development and named `b_release_{major}_{minor}_{patch}` (for example, `b_release_1_8_0`). A final release gets a tag in the form `v{major}_{minor}_{patch}` (for example, `v1_8_0` or `v1_6_7_1`).

Work for the next upcoming version is done on the `master` branch.

Lustre [Subsystem Map](#) describes each of the subsystems in the Lustre code.

[Back to Top](#)

Wiki Toolbox

- [What Links Here](#)
- [Related Changes](#)
- [Special Pages](#)
- [Printable Version](#)
- [Permanent Link](#)

Contributing Code to Lustre

lustre® Docadmin My Talk My Preferences My Watchlist My Contributions Log Out Help

Search Lustre.org

Article Discussion Edit History Delete Move Unprotect Watch

Contribute

Access guidelines for contributing code and testing Lustre™ along with developer resources and tools.

Contributing to Lustre

Getting started...

- See [Finding a Project](#) for information about how to select a project, find a bug to fix, help with Lustre testing, or contribute to the Lustre user documentation.
- Read the Lustre [Contribution Policy](#) and sign and return a Contributor Agreement.
- Join, post to, and search [Lustre mailing lists](#) for developers, administrators, and users.

Developing your code...

- See [Accessing Lustre Code](#) for how to download Lustre code.
- See [Building Lustre Code](#) for how to install and build a working version of Lustre.
- See [Applying Lustre Patches to a Kernel](#) for how to apply Lustre patches to an unpatched kernel and how to modify an existing kernel patch.
- Follow Lustre [Coding Guidelines](#) to avoid problems when merging your code.
- Use the Lustre [Documenting Code](#) guidelines to add reference documentation to your Lustre code contribution.

Debugging and testing your code...

- See [Testing Lustre Code](#) for procedures to verify the code works before you submit it.

Submitting your code...

- See [Submitting Patches](#) to find out how to submit your changes to be reviewed for acceptance into a mainline Lustre branch.

Developer Resources

- [Lustre Design Documents](#) provides access to presentations and design documents describing Lustre features currently under development by the Lustre engineering team.
- The [Lustre Design Document Archive](#) contains older architecture and design documents including feature and architectural descriptions, high-level design documents and detailed-level design documents.
- Lustre [Subsystem Map](#) describes each of the subsystems in the Lustre code.

Questions or Comments?

If you have questions or comments about how you can contribute or about any of the procedures on this page, please let us know [✉](#).

Wiki Toolbox

- What Links Here
- Related Changes
- Upload File
- Special Pages
- Printable Version
- Permanent Link

Main Page

Download

Learn

Use

Contribute

Get Involved

Getting started ...
(including *Finding a Project*)

Developing your code ...

Debugging and testing your code ...

Submitting your code ...

Developer Resources
(including *Lustre Design Document Archive*)



Contributing to Lustre

- Updated information on contributing to Lustre code development
 - How to find a project and connect with other contributors
 - Resources to develop, debug, test and submit code
- Developer resources now includes the Design Document archive:
 - Architecture and design documents
 - HLDs
 - DLDs



Lustre Internals Documentation (LID)

- LID is a new resource that provides detailed descriptions of the Lustre codebase
 - Located on the Lustre wiki – lustre.org/lid
- Key features:
 - Subsystem map
 - Documented subsystems
 - Doxygen-generated API documentation for various Lustre modules
- A few subsystems documented now (CLIO, LNET and ptlrpc), but more will be added in the future
- Lustre team welcomes contributions to the LID

Lustre Internals Documentation (LID)

l-u-s-t-r-e®

LUSTRE INTERNALS
DOCUMENTATION

ORACLE®

- » [LID Home](#)
- » [LID History](#)
- » [Global Index](#)
- » [Page Source](#)

Lustre Internals Documentation

Welcome to the Lustre Internals Documentation (LID) web pages.

The goal is to provide detailed descriptions of the Lustre codebase in an easily accessible format.

The following resources are available:

Glossary

Brief descriptions of [Lustre concepts](#), [objects](#) and [major components](#) indexed in various ways.

Lustre Internals: A Gentle Introduction

[Here](#) is an easy to read overview of the Lustre Internals.

Subsystem Map

The [subsystem map](#) provides links to the doxygen generated API documentation and other documentation for the current Lustre release.

Old Subsystem Map

The [old subsystem map](#) provides brief descriptions of most of the subsystems in an earlier Lustre release.

Understanding Lustre Filesystem Internals

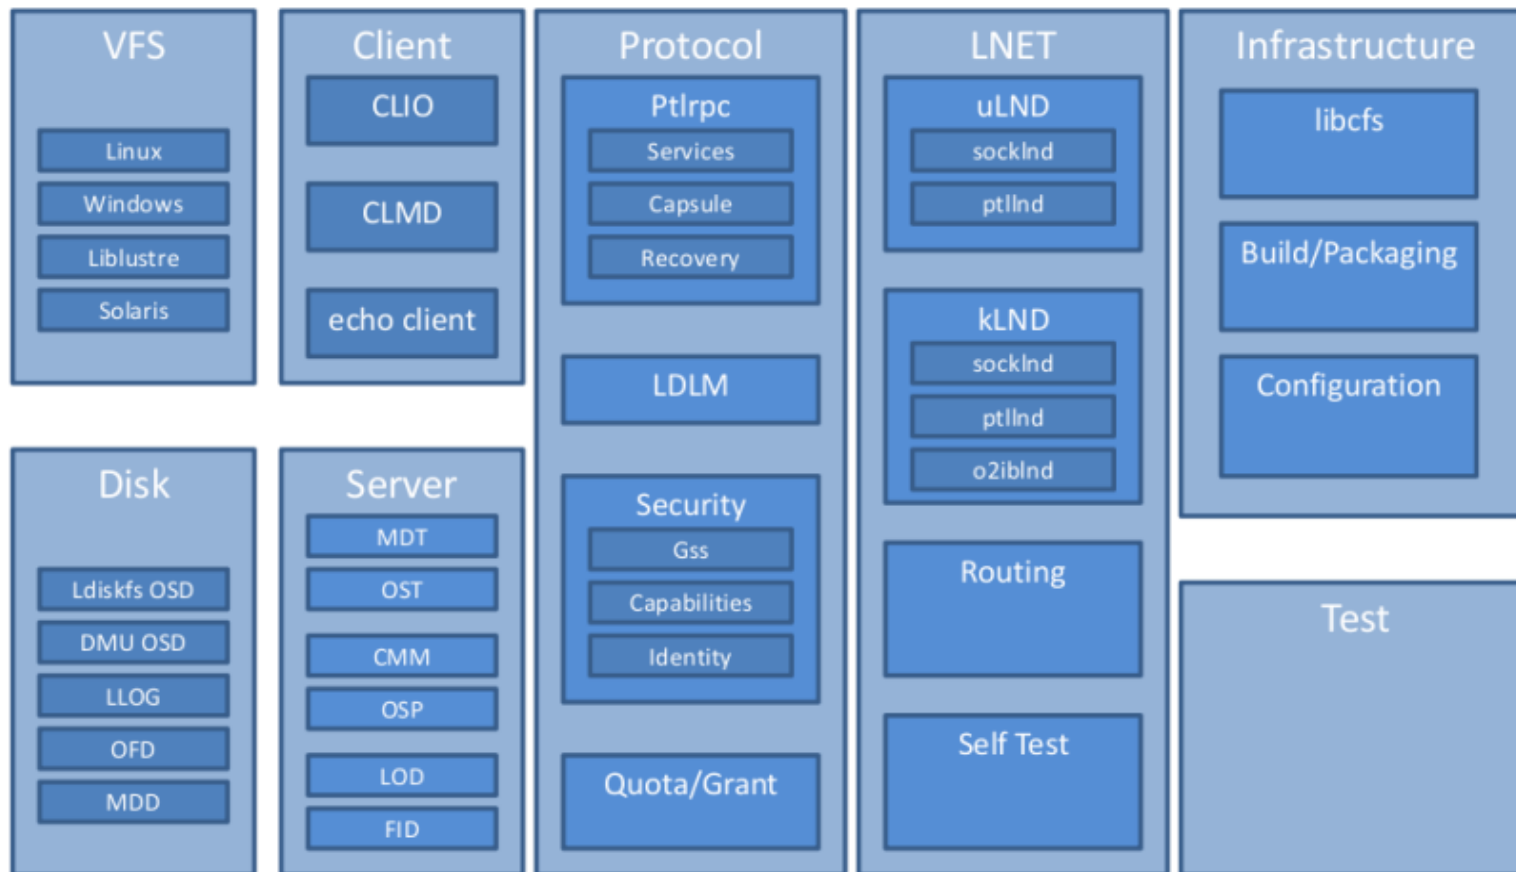
This is document ORNL/TM-2009/117 that was written by a team from the NCCS and Sun.

It is formatted as a single page [here](#).

The same document formatted as one section per page is [here](#).

Lustre Internals Documentation (LID)

Lustre subsystem map



Lustre Internals Documentation (LID)

Doxygen-generated API documentation



[Main Page](#) [Modules](#) [Data Structures](#) [Files](#) [Related Pages](#)

cli

Client objects implement io operations and cache pages. [More...](#)

Modules

cl_object
cl_page
cl_lock
cl_io
cl_page_list <i>Page list used to perform collective operations on a group of pages.</i>
cl_req
cl_env <i>lu_env handling for a client.</i>

Data Structures

struct cl_device_operations <i>Operations for each data device in the client stack. More...</i>
struct cl_device <i>Device in the client stack. More...</i>
struct cache_stats <i>Stats for a generic cache (similar to inode, lu_object, etc. More...</i>
struct cl_site <i>Client-side site. More...</i>

helpers

Type conversion and accessory functions.

void cl_page_slice_add (struct cl_page *page, struct cl_page_slice *slice, struct cl_object *obj, const struct cl_page_operations *ops) <i>Adds page slice to the compound page.</i>
void cl_lock_slice_add (struct cl_lock *lock, struct cl_lock_slice *slice, struct cl_object *obj, const struct cl_lock_operations *ops) <i>Adds lock slice to the compound lock.</i>
void cl_io_slice_add (struct cl_io *io, struct cl_io_slice *slice, struct cl_object *obj, const struct cl_io_operations *ops) <i>Adds io slice to the cl_io.</i>



Some of the Next Steps

- Keep improving the Lustre internals manual and LID pages
- Get more people contributing code
- Reference related community related tools and projects
- Continue dialogues on lustre-discuss and lustre-devel
- Community testing of Lustre 2.0 code drops
- ISC, Oracle Open World, SC10, and LUG
- Community workshops sponsored by Lustre sites
- User run user groups
- What else?



Thank You

Dan Ferber
Oracle Corporation

ORACLE®