

Lustre 2.2

Lustre 2.2 was released at the end of Q1 2012. This is its support matrix.

Server Support	Client Support
RHEL 6, CentOS 6	RHEL 6, CentOS 6, RHEL 5, CentOS 5, SLES 11

As well as ongoing stability improvements, there are a number of new features in the Lustre 2.2 release.

Feature	Summary	JIRA	Notes
Statahead	Performance improvement for ls -l. This feature accelerates sequential directory traversals by pre-fetching object attributes from the MDS. Fan Yong presented this at LUG 2011.	LU389	Statahead LUG 2011
Asynchronous glimpse lock (AGL)	AGL is an optimization of Statahead. It asynchronously fetches size and block attributes from the OSTs for better ls -l performance. Fan Yong from Whamcloud presented this feature at LUG 2011.	LU925	AGL LUG 2011
Client parallel checksums	Improves speed of client checksums by distributing the workload over all cores in a CPU by using the ptlrpcd feature. Also improves the accuracy by making checksums work properly for memory mapped pages.	LU884	Design Doc from Jira
Imperative Recovery	IR increases speed of recovery for larger deployments by actively notifying clients of server restarts for faster reconnections than the previous timeout mechanism. This feature was presented at LUG 2011 by Jinshan Xiong.	LU580	IR LUG 2011
Large xattrs	AKA Wide Striping. This feature increases the maximum stripe count from 160 to 2000 and was presented at LUG 2011 by Oleg Drokin. This feature effectively increases the maximum file size under ldxfs to 2000 * 16TB = 32PB.	LU80	Wide Striping LUG 2011
Mds-survey	This tool adds Metadata support in echo_client for benchmarking metadata performance. Now the MDS can be performance tested without requiring any clients to create metadata operations. This feature also includes a user friendly script similar to obdfilter_survey for easier testing.	LU593 /LU633	
Multi threaded ptlrpcd	This performance enhancement implements a general ptlrpcd pool. The ptlrpcd threads in this pool are shared by all the asynchronous RPCs on the client, such as BRW requests, data checksums, async glimpse locks, statahead, etc.	LU1019	Jira Design Doc
OSD API	The OSD restructuring project will allow alternate backend file systems to be used by Lustre. This feature is the first step in moving the new code to a released version of Lustre, and adds the OSD API to the OSS stack. This was presented to the EOFS LUG 2011 by Alex Zhuravlev.	LU909	EOFS LUG 2011
Parallel directory operations (pDirOps)	pDirOps introduces more granular locking on ext4 directory structures, thus allowing more concurrent operations simultaneously. This feature was presented at LUG 2011 by Liang Zhen.	LU50	PDirOps LUG 2011
Quota protocol /code improvements	This feature prepares the Lustre quota protocol for future quota changes which allow clients to get quota directly from OSTs instead of the MDS. It also does some reorganization to remove old/dead code, and adds new tests for quota accounting.	LU847	
Layout Lock Part 1	This patch creates a generation number for the file layout. This number will be used in conjunction the layout lock to allow clients to detect when the file layout has changed. This is an important feature for HSM and is the first of many HSM features to be landed in subsequent releases.	LU169	LL Patch