

Lustre 2.4

Lustre 2.4.0 was released in May of 2013. This is its test matrix.

Server Support	Client Support
RHEL 6.4, CentOS 6.4	RHEL 6.4, CentOS 6.4, SLES 11 SP2

As well as ongoing stability improvements, there are a number of new features in the Lustre 2.4.0 release.

Feature	Summary	JIRA	Notes
Network Request Scheduler (NRS)	The Network Request Scheduler manages incoming RPC requests on a server to provide improved and consistent performance. It does this primarily by ordering RPC request execution to avoid client starvation and to present a workload to the backend filesystem that can be ordered more easily. It may also change RPC concurrency as active fan-out varies to reduce latency seen by the client and limit request buffering on the server.	LU-398	
4MB RPC	The maximum bulk IO RPC size from a client to the OST was 1MB in previous releases. The 2.4 release increases the maximum bulk IO RPC size to 4MB to allow the back end storage to optimize the IO requests. An additional benefit is the reduction in the round trip time to send the data over high latency links. The default bulk IO RPC size is left at 1MB in this release to avoid potential performance loss in existing systems, but the larger 4MB size is available for testing and use in environments where this may show improvements (WAN, RAID devices that benefit from larger IO sizes).	LU-1431	
DNE Phase I: Remote Directories	This phase introduces a useful minimum of distributed metadata functionality. The purpose primarily to ensure that efforts concentrate on clean code restructuring for DNE. The phase focuses on extensive testing to shake out bugs and oversights not only in the implementation but also in administrative procedures. DNE brings new usage patterns that must necessarily adapt to manage multiple metadata servers.	LU-1187	
FID on OST Infrastructure	Allow objects on OST to be identified by a unique FID, instead of object id that is not unique between OSTs. This is required for DNE functionality and will also be needed in the future for other features.	LU-1445	
Client IO Performance Improvements	A number of performance enhancements have been made to the Lustre client IO code to improve IO performance and reduce CPU usage compared to what is available in the previous Lustre 2.x releases.	LU-744	
LFSCK Phase 1.5	Implement verification and rebuilding of FID-in-dirent and linkEA for the single-MDT case. This is needed for maximum metadata performance for upgraded 1.8 MDTs, or if a 2.x MDT has had a file-level backup and restore.	LU-1866	
OSD Server Stack	Restructuring of the Lustre MDS, OSS, and MGS server code modules to work with the OSD API. This allows the Lustre MDT and OST to use the ZFS backing filesystem, and to simplify the implementation of projects such as DNE. It also prepares the way for future OSD types. The MDT stack has new LOD/OSP modules (analogous to LOV/OSC) to communicate with the OSTs to more transparently handle both local and remote OSD devices.	LU-1711 , LU-1301 , LU-1302 , LU-1303 , LU-1304	
Quota Enforcement	<p>Rewrite of server side quota architecture to deal with OSD and DNE restructuring to work with both ldiskfs and ZFS backing filesystems, and to work with the DNE feature.</p> <ul style="list-style-type: none">• sharing the connection to MDT0 with FID-on-OST• having one single connection between MDTs which can be used by both cross-MDT operations and quota acquire/release• change sanity-quota.sh to support DNE• add more test cases to exercise quota + DNE <p>This work also addresses some enhancements to the original implementation for future features.</p>	LU-1842 , LU-2183	
ZFS OSD Utilities	Lustre user utilities <code>mkfs.lustre</code> and <code>mount.lustre</code> require an update in order to properly function with a ZFS-based OSD.	LU-1581	
HSM Client Side Support Features	<p>Dirty Flag, NDT LVB Support, Layout Lock, Temporary File Support, Flags, Layout Swapping, Activity Monitoring</p> <p>This work represents the majority of the client side support needed for HSM. The remaining server side support and support for HSM to be functional will be completed in the Lustre 2.5 release.</p>	LU-941 , LU-1710 , LU-1876 , LU-2441 , LU-1338 , LU-2016 , LU-2060	
Wireshark Support for LNet and Lustre RPCs	Upgrade the Wireshark support to handle new versions of Wireshark and Lustre. LNet and Lustre plugins updated for Wireshark 1.4 (backwards compatible with Wireshark 1.2.x) are available under the <code>lustre/contrib/wireshark/</code> directory.	LU-1434	
LNNet Networks Hashing	This feature changed the table in LNet which keeps track of all remote networks into a hash table. This speeds up access to these entries when dealing with a very large number of remote networks (over 1000). In addition, a new module parameter, <code>rnet_hhtable_size</code> , was added to allow customers to control the size of this hash table.	LU-2466	

Ability to Disable Pinging	A new ptlrpc module parameter, "suppress_pings", is introduced to provide an option for reducing excessive OBD_PING messages in large clusters. The parameter is a switch and affects all MDTs and OSTs on a node. (MGS pings can not be suppressed.) By default, it is off (zero), giving a behavior identical to previous implementations. If it is on (non-zero), all clients of the affected targets who understand OBD_CONNECT_PINGLESS will know, at connect time, that pings are not required and will suppress keep-alive pings. This feature requires a mechanism external to Lustre to notify the server that clients have failed.	LU-2467	
----------------------------	--	-------------------------	--