

**Lustre User Group Meeting - 2006-04-20
Requirements**

<i>Category</i>	<i>Description</i>	<i>ID</i>	<i>Votes 1</i>	<i>Votes 2</i>
Committed	UID mapping	12		999
Committed	Portals over IB to IP router code with OpenIB	14		999
Committed	Improved metadata performance	5		999
Committed	User/Administrator Documentation and Training	6		999
Committed	Cluster metadata servers	7		999
Committed	Online OST expansion	8		999
Committed	wide and n:1 file striping optimization	9		999
Committed	Improved timeout mechanism; allow reconnect instead of eviction after network failure	10		999
Committed	Security (esp. for Lustre over WAN); Kerberos, object protection (capabilities)	11		999
Committed	Patchless clients	13		999
Committed	Automatically manage amount of free space on OSTs (aka QoS)	15		999
Committed	Stable router behaviour (e.g. IB to IP failure use cases)	16		999
Committed	Multi-hop routing	17		999
Management	Understandable and documented error messages; troubleshooting	39	41	59
Management	OST stripe management: 1) Pools; 2) Join files; 3) background migration)	35	24	55
Management	Improved Logging, Debugging and Diagnostic tools; NID logic; per-client stats	38	21	46
Backup	Lustre HPSS copy, using DMAPI	1	14	43
Management	I/O, metadata performance profiling, analysis, and reporting	33	15	40
Management	Cluster monitoring tool	37	16	39
Performance	NUMA awareness	59	21	39
Management	Lustre internals documentation	46	15	34
Stability	Version based recovery	64	11	34
Management	Global health check	43	16	30
Management	Multiple mount protection	36	14	29
Feature	Large and more usable EAs (e.g. DRM information)	21	7	25
Management	Automatic tuning of cluster based on size	32	10	24
Management	Forced global unmount	42	12	24
Stability	Official support for clients and OSTs on the same server	63	9	24
Performance	Improved implementation of SRP (RDMA)	52	11	21
Backup	s-tar support for user striping; metadata backup	2	9	19
Feature	MPI-IO (ADIO); HDF5, NET CDF	31	9	19
Feature	Iron ext3	66	8	19
Stability	Failed router avoidance for all LNDs	62	10	18
Management	Wider range of version/multi-vendor interoperability	40	10	17

Performance	Multi-threaded client metadata operations	56	6	17
Feature	Arbitrary striping patterns	22	2	15
Performance	Server side data cache	53	7	15
Management	Cluster inventory tool (what clients are connected)	34	9	13
Performance	Full 64-bit SCSI addressing	51	8	13
Stability	Network failover (e.g. from Infiniband to TCP0)	61	7	12
Feature	Technology Independent Performance (aka Pools)	24	5	11
Backup	Snapshot capability	3	5	10
Feature	Directory based quotas	26	10	10
Feature	Multi-resident metadata and objects for parallel access (HSM, blunt replication, versioning, persistent caching)	30	10	10
Feature	Compressed file storage	20	3	9
Stability	Complete failover/recovery routers	65	7	9
Feature	Clustered NFS exports (for Linux 2.6)	19	8	8
Feature	Quality of service - use of Lustre in a multi-clustered environment; policy for managing priority of service	23	7	7
Management	Configuration/tuning guidelines and tools for different purposes (e.g. scratch vs. production)	41	7	7
Management	Developers guide for application developers (how to code for distributed file systems)	47	4	7
Performance	30 GB/s per client	57	6	7
Management	Report files associated with stripes on an OST (lfs find)	44	6	6
Feature	Cooperative caching	28	1	5
Performance	Petaflop system	58	5	5
Feature	Content addressability of data; files sharing identical extents	29		3
Management	Maximum object size	45		3
Performance	Defragmentation tool; report and fix fragmentation issues (esp. w/ Linux 2.4 systems)	50	3	3
Feature	Mount Lustre subdirectories	25	2	2
Feature	I/O pattern recognition for applications	27	2	2
Management	Storage management attributes (e.g. creation date, retention period)	48	2	2
Performance	Loosen POSIX behaviour	54	2	2
Performance	Improve fsck capability	55	2	2
Stability	Data integrity protection; checksums	60	2	2
Backup	Commercial backup package support for user striping	4	1	1
Feature	Search engine support (e.g. Beagle, Spotlight)	18		0
Management	Auto-reporting of system configuration and performance metrics to CFS	49		0