

ORACLE®



ORACLE[®]




Lustre 2.0

Robert Read
Senior Technical Lead, Lustre Group

Program **Agenda**

- State of 2.0
- What is and isn't in 2.0?
- What to expect after the GA





Finally time to say...

BETA!



What's new in 2.0

- New architecture on client and MD server
 - Focus has been on stabilizing new code
- Changelogs
 - Enables HSM
 - Example app - lustre-rsync
- SOM preview
- Security preview
 - Kerberos integration
- No lustre_config, try Shine



Since LUG 2009

- 7 Alpha releases
- 10,000 tests
- 491 2.0 bugs fixed
- <30 open release blockers
- 1 new SCM
- 1 Corporate Merger

Daily testing results

	September	October	November	December	January	February	March	Total Runs
ior	80%	100%	98%	96%	98%	100%	100%	324
simul	95%	72%	94%	96%	98%	100%	98%	315
racer	14%	45%	35%	67%	75%	96%	93%	274
sanity	4%	2%	16%	66%	71%	90%	87%	317
sanity_benchmark	45%	75%	88%	80%	82%	94%	90%	315
lustre-rsync_test	88%	85%	94%	0%	5%	85%	87%	264
sanityn	33%	87%	57%	86%	87%	98%	96%	312
liblustre	94%	98%	80%	98%	88%	98%	95%	320
replay_single	8%	15%	14%	86%	76%	98%	93%	316
metabench	NT	100%	NT	98%	100%	100%	100%	168
recovery_small	50%	82%	98%	90%	83%	90%	93%	312
replay_dual	98%	79%	78%	94%	71%	92%	95%	317
insanity	100%	95%	100%	96%	94%	98%	100%	318
sanity_quota	33%	82%	92%	88%	85%	96%	91%	324
sanity_sec	100%	100%	100%	96%	98%	100%	100%	323
performance_sanity	8%	75%	67%	85%	75%	80%	82%	323
replay_vbr	98%	71%	67%	94%	83%	96%	93%	322
replay_ost_single	10%	91%	92%	96%	88%	100%	91%	322
conf_sanity	79%	84%	73%	67%	60%	90%	95%	319
pios	40%	88%	100%	82%	86%	100%	100%	57
parallel_scale	0%	18%	55%	27%	57%	38%	63%	61
lnet_selftest	100%	100%	100%	100%	100%	88%	100%	51
recovery-mds-scale	NT	70%	91%	100%	100%	63%	100%	51
recovery-dbl-scale	NT	44%	91%	91%	71%	50%	63%	51
recovery-rand-scale	NT	78%	100%	100%	71%	75%	50%	51
ost_pools	NT	20%	33%	45%	29%	25%	50%	45
parallel_scale_nfsv4	NT	0%	0%	0%	NT	50%	0%	37
parallel_scale_nfsv3	NT	0%	0%	0%	NT	13%	88%	39
metadata_updates	NT	0%	0%	0%	0%	0%	0%	43

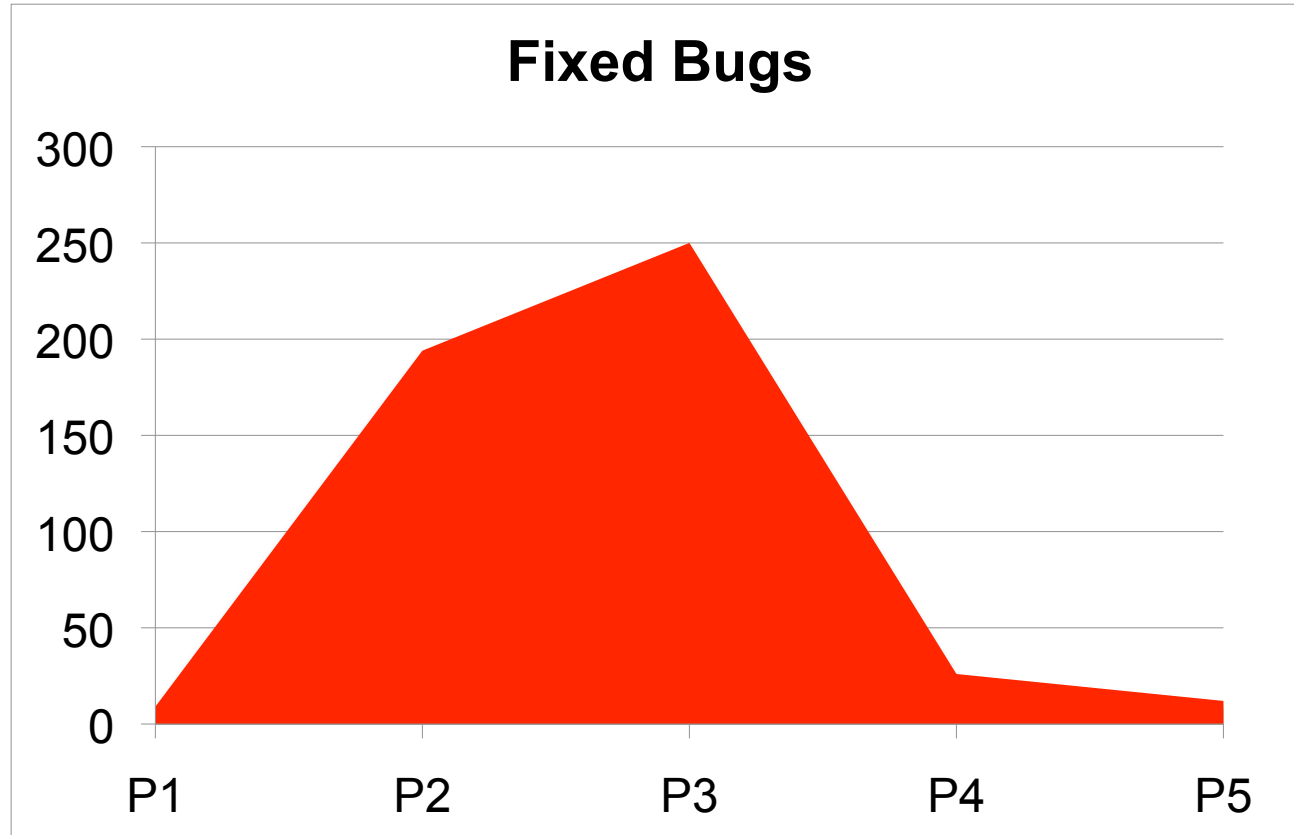


Bug Priorities

- Used bug priorities to manage 2.0 bugs
 - P1 – Show Stopper
 - P2 – LBUG, LASSERT, consistent errors
 - P3 – Non-fatal test errors
 - P4 – Enhancements
 - P5 – Not in scope for this release



491 Bugs fixed and closed






Why only Beta?

- No large scale (more than hundreds) testing yet
 - Can we borrow your top 10 cluster?
- Still working on interoperability issues
 - 1.8 client needs to work as well as 2.0 client in GA
- Recent issue with ext4 vs. ext3
- Recovery issues


2.0 Performance

- Stability is the priority for this release
- Accepting some performance regressions
 - IO seems most affected
 - Some MD issues
 - Still doing analysis
- SMP scaling project already underway
 - Goal is to set new records Lustre 2.



Lessons Learned

- Main branch “master” is now tested daily
- Gatekeepers actively track main branches
 - Single gatekeeper per branch
 - Approves and controls commits
 - Reverts changes when a regression is identified
- Main branches should be kept in releasable state
 - All regression tests should pass
 - Regular stress/scale testing
 - Keep blocker bugs under control






What's After 2.0 GA?

- 2.0.x maintenance releases
 - Bug fixes only
 - Kerberos fixes
 - Possibly some minor performance enhancements
- 2.x release
 - Planned performance enhancements
 - Other features as ready



ORACLE IS THE INFORMATION COMPANY



The preceding 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.