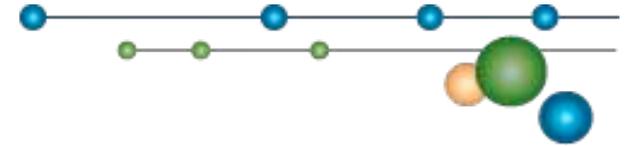




CFS Engineering Update Miami, April 2007

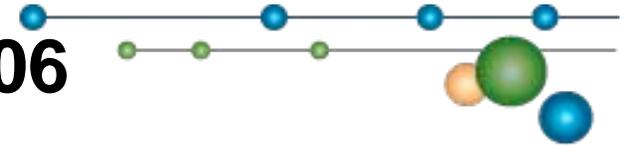
**Peter Bojanic
VP Engineering
Cluster File Systems, Inc.**

Contents



- **Accomplishments since LUG 2006**
- **Lustre Releases**
- **Engineering Practice**
- **Lustre Community**
- **Top Goals for 2007**
- **Challenges for 2007**

Accomplishments since LUG 2006



Release	Date	Major Improvements
1.4.7	06-08-29	nfs exports; MD performance improvements
1.4.8	06-12-22	large ext3 partitions; RAID 5 improvements
1.4.9	07-02-12	SLES 10 support; patchless zero copy sockIn
1.4.10	07-04-04	ldiskfs2 (fast e2fsck; mballo3)
Lustre 1.6.0 beta[34567]	07-01-17	mountconfig; patchless client; optimal stripe allocation

Other Accomplishments



- **Lustre Operations Manual**
 - Introduced with Lustre 1.4.6
 - Major usability improvements for Lustre 1.6.0
- **Hendrix project**
 - CMD, GSS/Kerberos, OFED
- **Ltest improvements**
 - Enabled mixing and matching client and server branches for patchless client testing

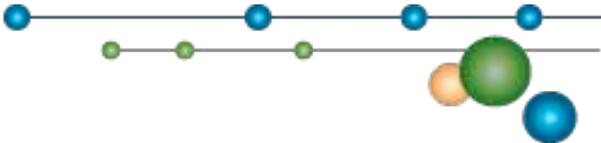
Lustre Releases



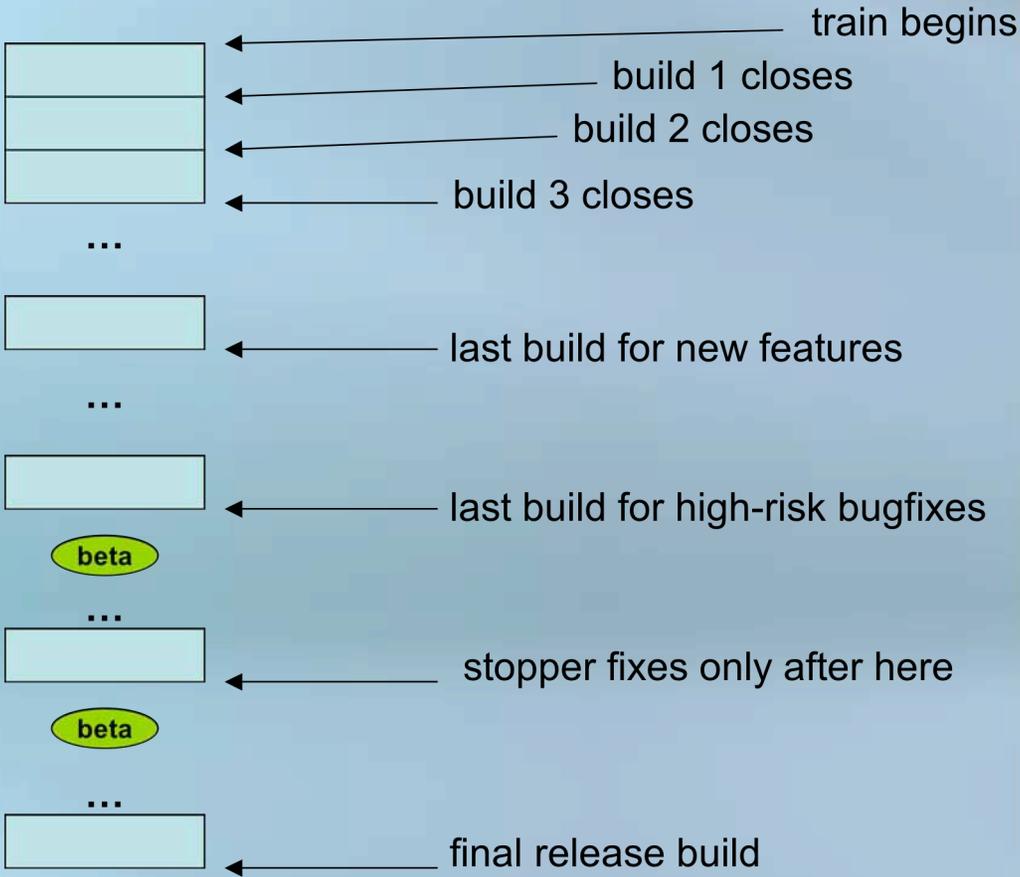
lustre™

- **Focus on more predictable, stable, and robust releases**
- **Product Management Group**
 - Decisions are made by a broader representation of CFS groups
 - Includes executive management, business development, marketing, engineering, support
- **Release Management Group**
 - Maintains control over development and maintenance branches
 - Coordinates landings by all engineers
 - Shepherd releases forward through testing process
- **Lustre releases are planned and managed according to a Train Model**

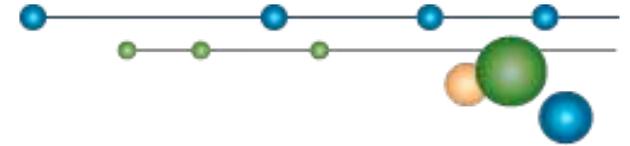
Train Model



- 6-month major releases
- 2-month maintenance releases
- Bi-weekly builds
- Narrowing window of risk
- Train may get held for important landings
- Otherwise, patches need to wait for the next train



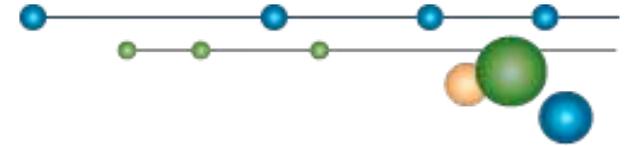
Gold Standard Lustre



lustre™

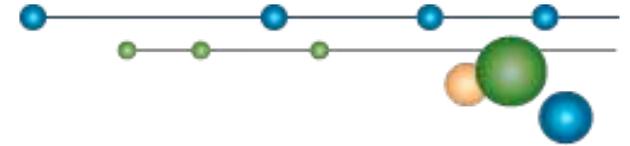
- **Getting major customers and partners much more closely aligned with Lustre releases**
 - There are too many “versions” of Lustre out there
 - They don’t all have known bug fixes
 - They are not tested by CFS
- **Avoiding patch soup**
 - Testing on ALC Test and XT systems prior to release
 - Raising the bar in terms of quality and robustness, especially on very large scale systems
- **Stamp of approval by major customers and partners**

Releases - 1.4.x



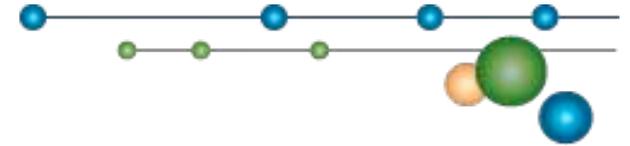
Release	Date	Major Improvements
1.4.10.1	ASAP	Critical update to address bug 12181: dirty pages not being flushed in a timely manner
1.4.11	07-04-23	Maintenance release, especially for Cray Unicors 2.0
1.4.12	07-06-22	Maintenance release
1.4.13	07-08-22	Maintenance release
1.4.?	Not planned	May be considered; maintaining this branch will consume resources and slow us down on other efforts

Releases - 1.6.x



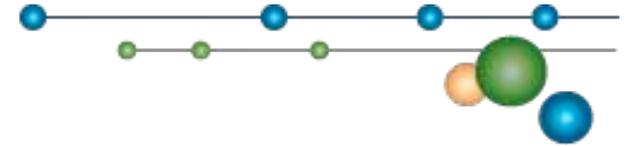
Release	Date	Major Improvements
1.6.0	ASAP	Mountconf; patchless client; optimal stripe allocation
1.6.1	07-05-25	RHEL5 and vanilla 2.6.18 support
1.6.2	07-07-25	Metadata performance improvements; adaptive timeouts; online defrag
1.6.3	07-09-25	Version based recovery
1.6.?	Not planned	Will be considered depending on progress with future Lustre releases

Releases - Future



Release	Date	Major Improvements
1.8.0	07-08-31	GSS/Kerberos; compatibility for Lustre 2.0; RAID 6 improvements
1.10.0	07-12-15	User space OSS; support for user space Solaris OSS and OFED LND
2.0	08-Q3	Clustered Metadata Data

Distribution Support

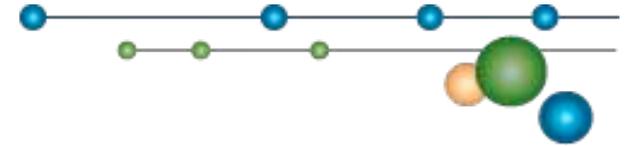


Lustre™

Lustre Release	RHEL3	RHEL4	RHEL5	SLES9	SLES10	Vanilla 2.6.18+
1.4	C	C,S	–	C,S	C,S	–
1.6	C*	C,S	C,S	C,S	C,S	C,S
1.8	–	C,S	C,S	P	C,S	C,S
1.10	–	P	C,S	P	C,S	C,S
2.0	–	P	P,S?	P	P,S?	P,S?

Legend: C=client; P=patchless client only; S=server; *=support caveat

Engineering Practice



lustre™

- **Team**
 - New hire concentrations around Boulder, Beijing, Kiev
 - Some growth since last year, most notably a maturing engineering management team
- **Process**
 - Still fully subscribed to PSP/TSP
 - Considerably more process oriented
- **Quality**
 - LTS development slower than needed, but moving forward
 - Manual testing on XT at Cray; ALC Test at LLNL
 - Continued testing support by HP, Bull, ORNL, and many more
- **Support**
 - Growth in this team; capacity being ramped up in Europe
 - RT introduced; being phased back out

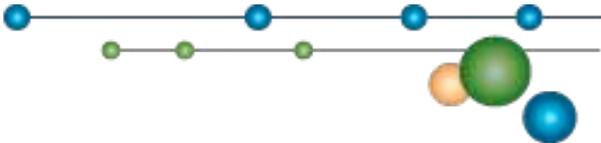
Lustre Community



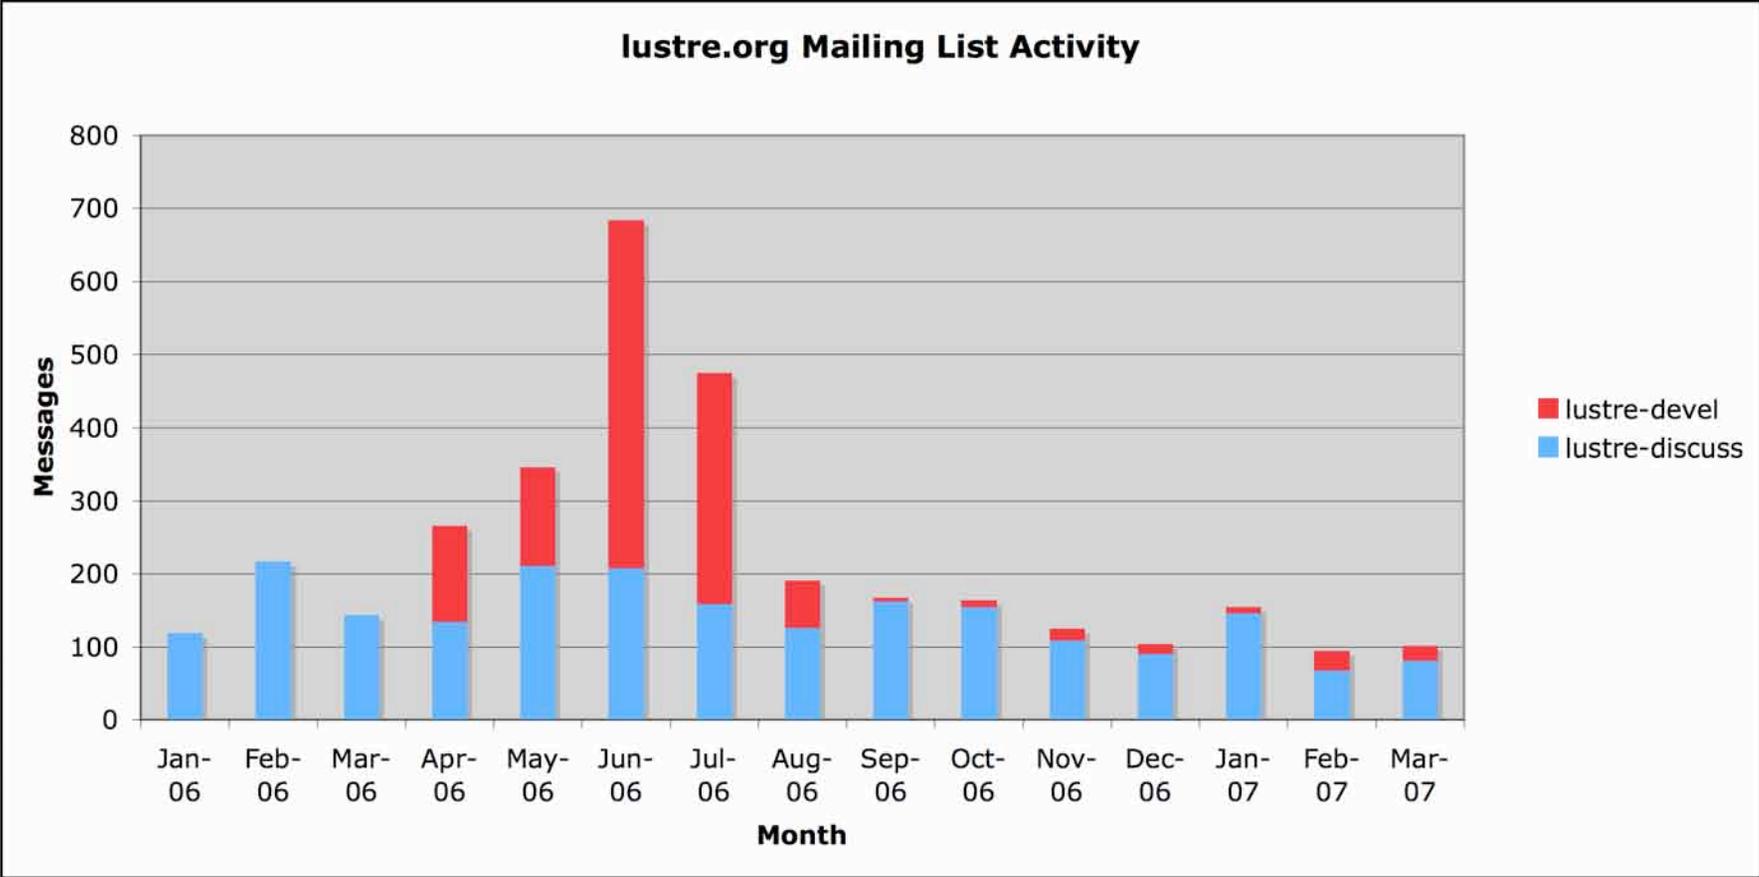
lustre™

- **Openness**
 - Lots of interaction on lustre-devel and other mailing lists
 - Opt-in to publish Bugzilla comments
 - Patches available publicly
 - Considerably more sharing of designs and architectural information
- **Contributions**
 - Ext4, e2fsprogs
 - New allocator
 - OFED
- **Community development**
 - We haven't had the resources to invest that we would like
 - Considering engaging Collab.net as a catalyst
 - Top priority is a public SCM

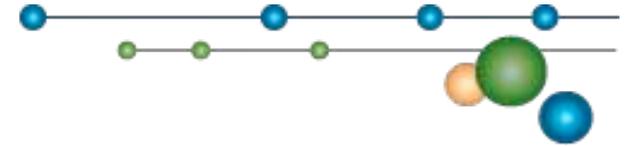
lustre.org Mailing Lists



lustre™



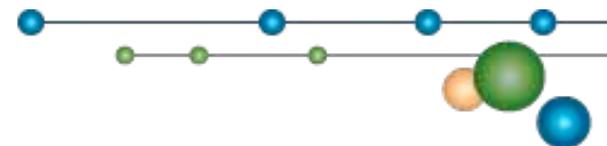
Top Goals for 2007



lustre™

- **Major bug squashing**
 - User level servers
 - Client restructuring
- **Testing**
 - LTS evolutionary development and deployment
 - Relationships with customers and partners (Cray, LLNL, others)
- **Predictable releases**
 - Better estimating
 - Dedicating uninterruptible resources to projects
 - Reliable quality through bug obliteration initiatives and testing

Challenges for 2007



lustre™

- **Whole storage stack**
 - The file system has considerable dependencies from kernel to network to disk; any misbehaviour affects Lustre
 - We need to verify and tune all the layers for screaming Lustre performance
 - Efforts underway include LNET smoke test, DDN 95xx tuning effort
- **Balancing product maintenance and evolution**
 - Fixing bugs while rewriting components
 - We need to do both
- **Growing the engineering team**
 - We have awesome people; we need more
 - Ramping up new engineers faster

Summary



lustre™

- **CFS has accomplished much of what it planned for 2006**
 - But there were some important things that slipped, like Lustre 1.6.0
- **Quality is a top priority**
 - Strategic improvements as with user space OSS
 - Quality processes, testing
- **Better coordination with Customers and Partners**
 - Predictable release schedule
 - Keeping closer in line with official, tested releases
- **Even closer collaboration with community**
 - Development and testing
- **Great engineers are a key ingredient!**