Upstreaming Lustre to Linux

Three topics for discussion.

Neil Brown
Senior Kernel Engineer
SUSE

Today I would like input on three topics.

1. What is the general level of commitment, or buy-in, to this?
2. What, in general terms, in the sequence of steps?
3. What are some possible pain points that need special attentions?

What is the general level of commitment?

1. Are you in favor? Keen? Excited?
2. Is your management supportive?
3. What concerns hold you back, which I can try to alleviate?

What, in general terms, in the sequence of steps?

1. Finish code cleanup
2. Update upstreaming tree to match “master”
3. Land client without RDMA upstream: Lnet and Lustre
4. Add RDMA klnd
5. Add server support over native EXT4
6. Add required enhancements to VFS and EXT4

A) osd-zfs remain out of mainline (while zfs does).

B) Development patches must continue to flow into mainline

C) At some point, development work must go to mainline first, then into lustre-release/master

D) Eventually, the “lustre” package has no kernel code, except backports.

What are some possible pain points that need special attentions?

1. Job-id from environment
2. Private logging infrastructure
3. RDMA?
4. lu_ref reference tracking
5. Removal of unused interfaces (Dead code)
6. “mount -t lustre” for client only, server needs something else.
7. Replace percpt locking with RCU
8. Gerrit ….

Thank You !!


SUSE
We adapt. You succeed.