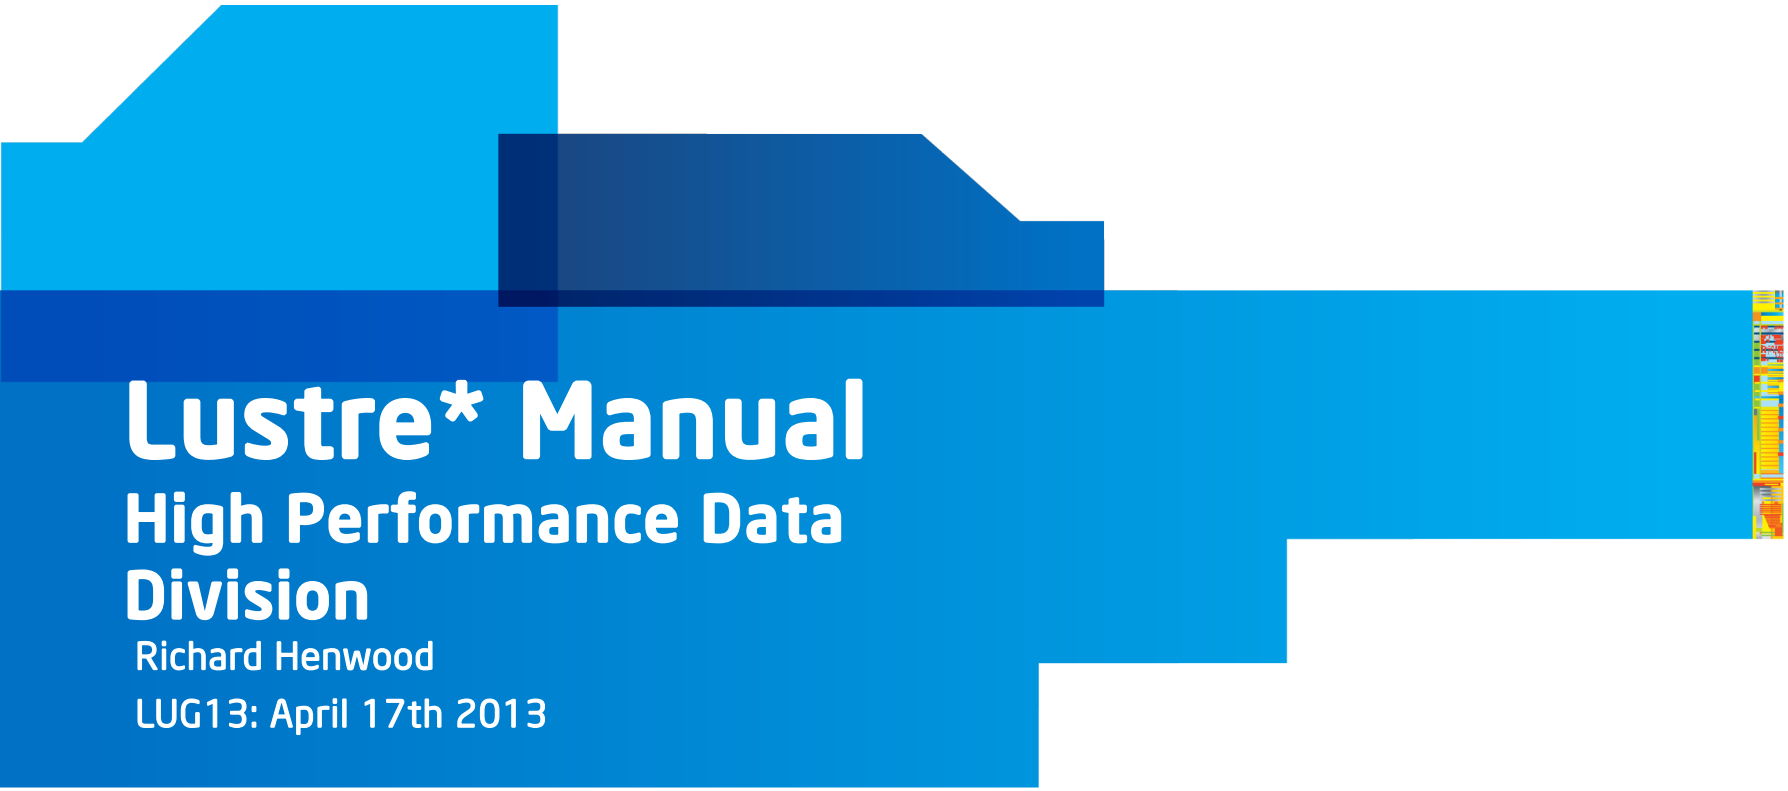

Lustre* Manual

High Performance Data Division

Richard Henwood

LUG13: April 17th 2013

Overview

What is in the Lustre manual this instant?

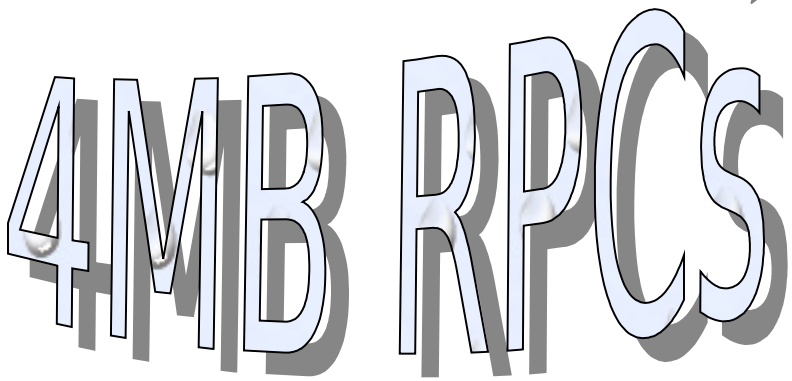

But that's not all...

Multiple MDTs (DNE)	yesterday
Lustre Tuning parameters	yesterday
LFSCK	this morning
Wireshark	this morning
Change-logs	tomorrow

JobStats, mds-survey, tuning, debugging etc, etc, etc

History

For me, it all started with:

```
# lctl get_param llite.*.stats
snapshot_time             1308343279.169704 secs.usecs
dirty_pages_hits          14819716 samples [regs]
dirty_pages_misses        81473472 samples [regs]
read_bytes                36502963 samples [bytes] 1 26843582 55488794
write_bytes               22985001 samples [bytes] 0 125912 3379002
brw_read                  2279 samples [pages] 1 1 2270
ioctl                     186749 samples [regs]
open                      3304805 samples [regs]
close                     3331323 samples [regs]
seek                      48222475 samples [regs]
fsync                     963 samples [regs]
truncate                  9073 samples [regs]
setxattr                  19059 samples [regs]
getxattr                  61169 samples [regs]
```

Current status

- 118000 words, 13000 lines
- Available as pdf, html, epub
- ~90 commits

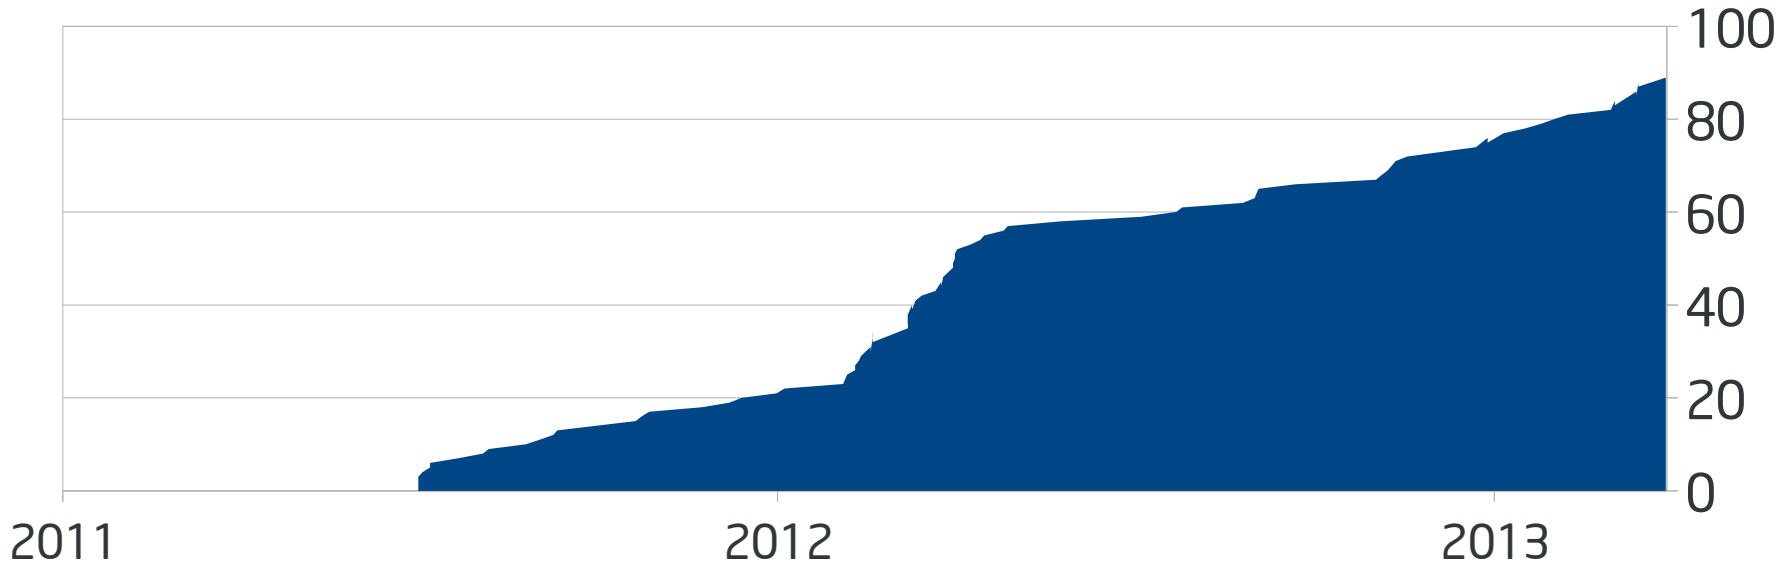

Contributing

Engineering work-flow

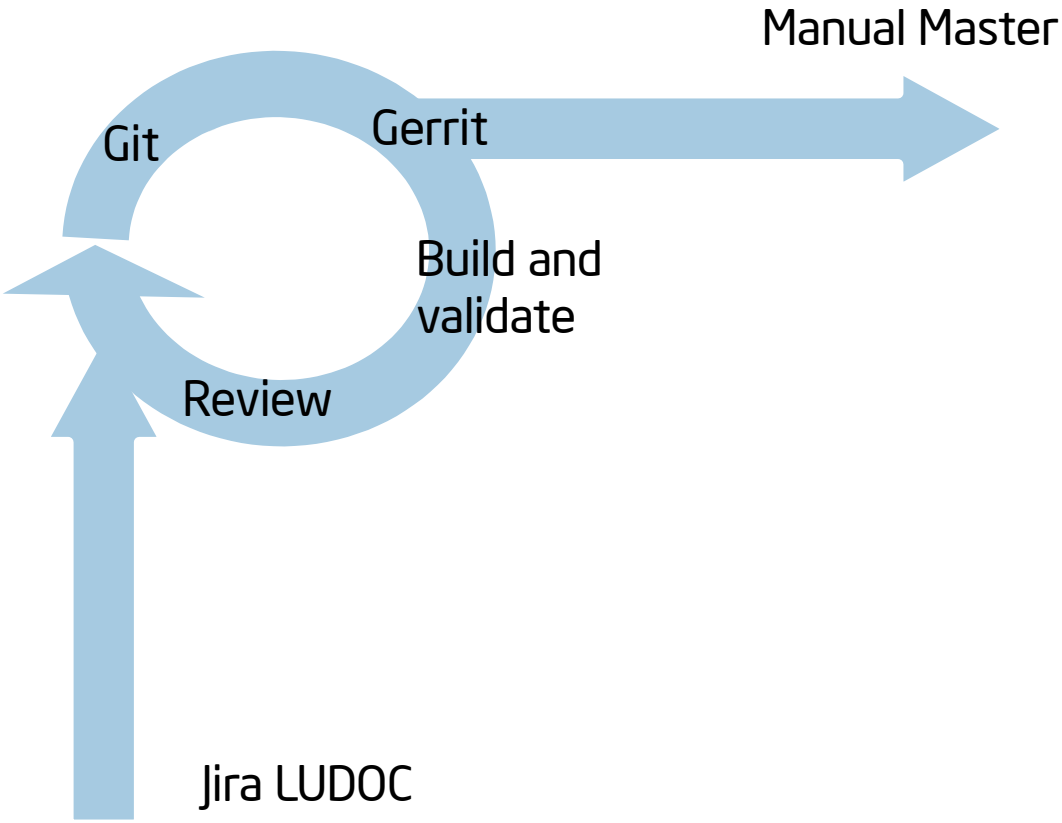

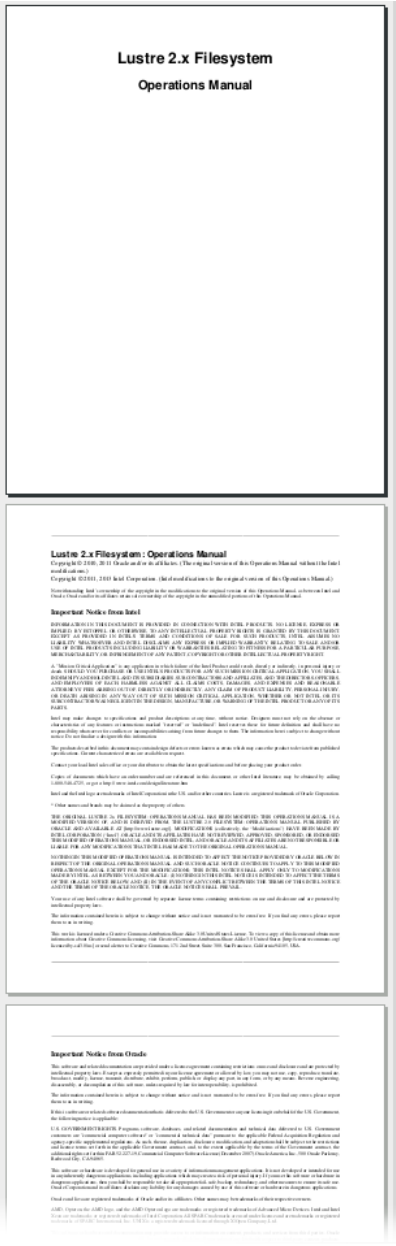

Lustre Manual archaeology

Docbook

- Industry standard
- xml

- 400 tags
- xml editing

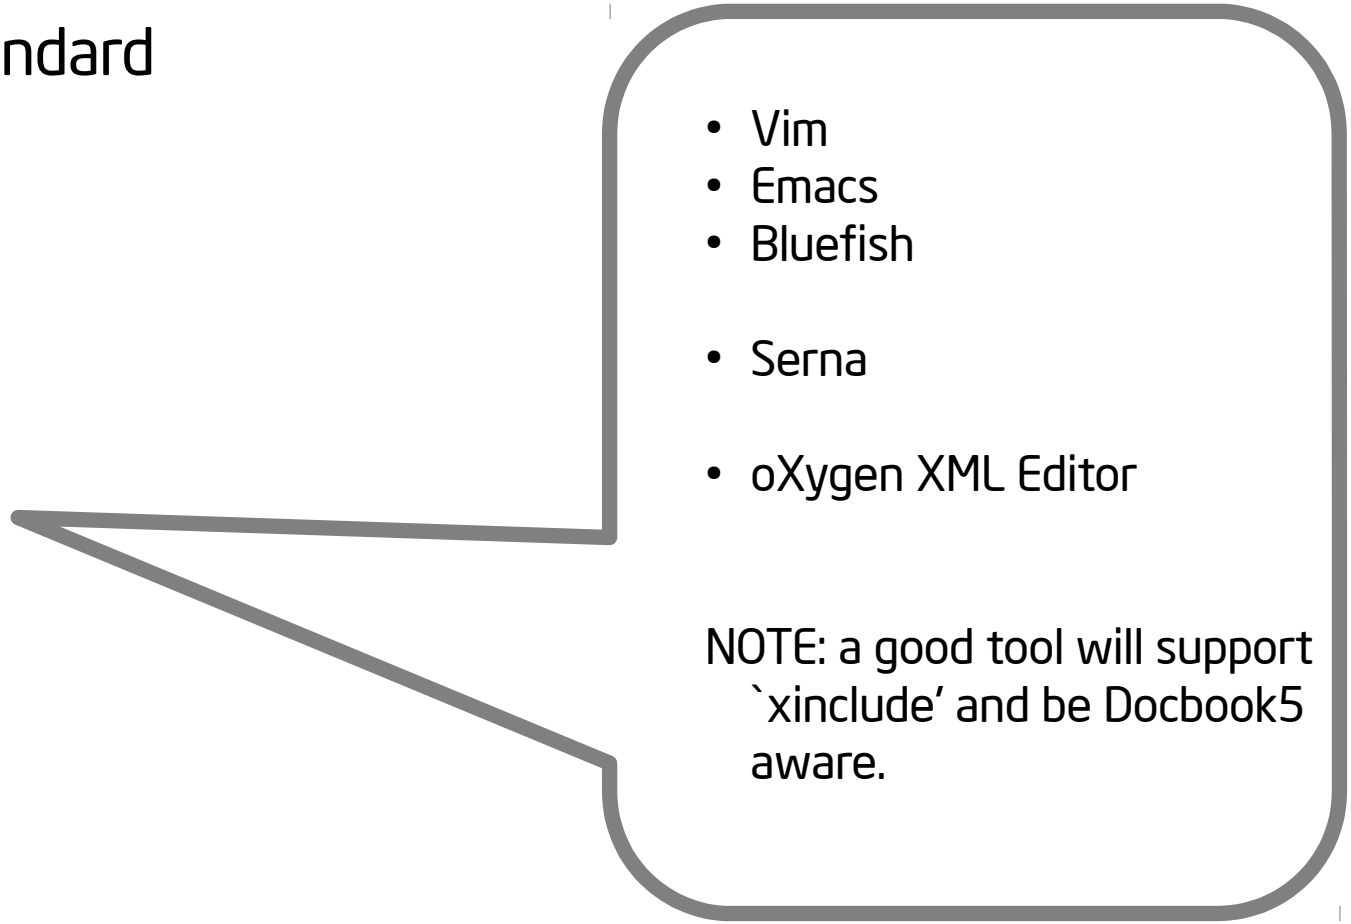

Lustre Manual archaeology

Docbook tags: an example with <replaceable>

```
<title>Synopsis</title>
<screen>
    lctl lfsck_stop -M |
        --device <replaceable>MDT_device</replaceable> \
                    [-h | --help]
</screen>
</section>
```

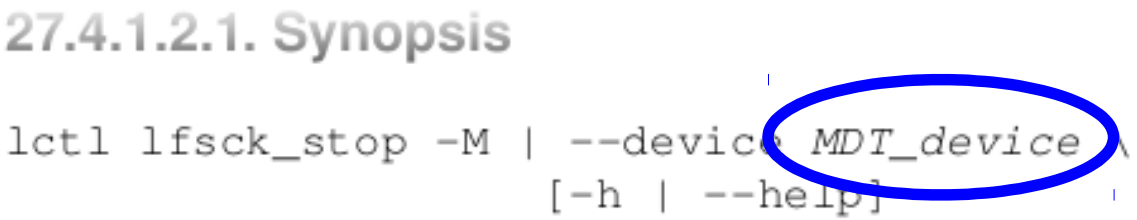

Lustre Manual specifics

```
<* conditional='l23'>Lustre 2.3| specific text.</*>
```

25.2. Binding MDS Service Thread to CPU Partitions

introduced in Lustre 2.3

With the introduction of Node Affinity (Node Affinity) in Lustre 2.3, MDS threads can be bound to particular CPU Partitions (CPTs). Default values for bindings are selected automatically to provide good overall performance for a given CPU count. However, an administrator can deviate from these setting if they choose.

- `mds_num_cpts=[EXPRESSION]` binds the default MDS service threads to CPTs defined by `EXPRESSION`. For example `mdt_num_cpts=[0-3]` will bind the MDS service threads to `CPT[0,1,2,3]`.

Contributing continued...

- Low barriers to entry with submissions to

  http://jira.hpdd.intel.com/browse/LUDOC

- Lots of opportunities for improvements, large and small

What to expect in future:

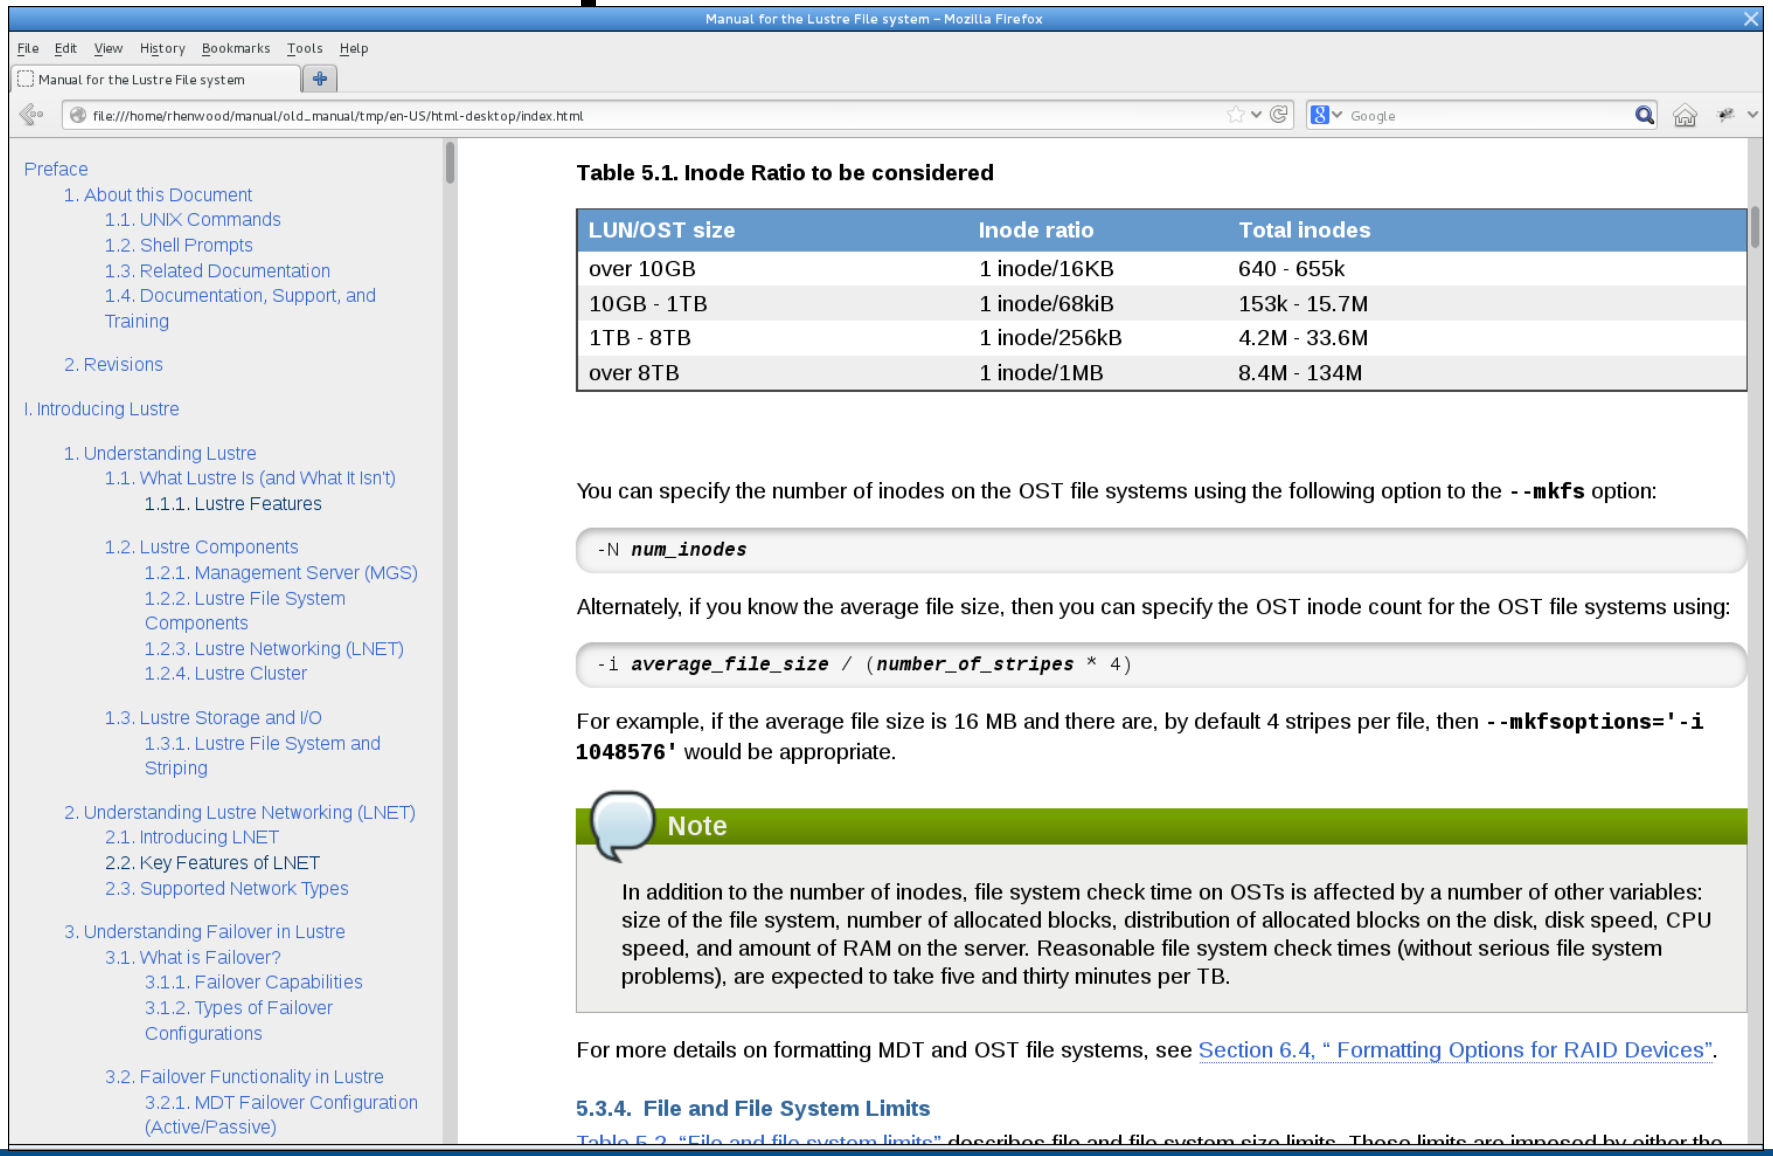

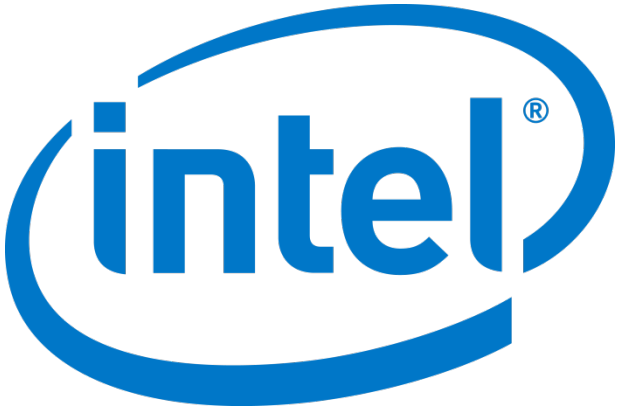
intel