

# Improved Versioning, Building, and Distribution of Lustre

LUG 2016

Christopher J. Morrone, Giuseppe Di Natale

April 5, 2016




# Community Release Package

---

lustre-2.8.0-3.10.0\_327.3.1.el7\_lustre.x86\_64.x86\_64.rpm

# Parts of RPM File Name



# Airing of Grievances

---

- RPM Release field abuse
- Tarball version incomplete
- Package tied to single kernel
- Lustre unbuildable without git
- Incompatible with Koji
- Release Candidate abomination
- Multiple version headers
- Complicated code

# More Version Examples

- Development package
  - lustre-2.8.50-3.10.0\_327.10.1.el7\_lustre.g941eb0e.x86\_64\_**g2502560**.x86\_64.rpm
- Third party package (--with-downstream-release=)
  - lustre-2.8.0-**9acme**\_3.10.0\_327.4.5.el7.x86\_64.x86\_64.rpm
- Third party development package
  - lustre-2.8.0-**9acme**\_3.10.0\_327.4.5.el7.x86\_64\_**gdf14e91**.x86\_64.rpm

# Inconsistent Versions

---

- lustre-2.8.0.tar.gz
- lustre-2.8.0-1\_gdf14e91.src.rpm
- lustre-2.8.0-9acme\_3.10.0\_327.4.5.el7.x86\_64\_gdf14e91.x86\_64.rpm

# Almost Consistent Versions

- lustre-2.8.0\_9acme\_1\_gdf14e91.tar.gz
- lustre-2.8.0\_9acme\_1\_gdf14e91-1.src.rpm
- lustre-2.8.0\_9acme\_1\_gdf14e91-1\_**3.10.0\_327.4.5.el7.x86\_64**.x86\_64.rpm

ee813db LU-7699 build: Replace version\_tag.pl with LUSTRE-VERSION-GEN

# Third Party Versions

- After 2.8.0: --with-downstream-release is **OBSOLETE**
- Consult *Documentation/versioning.txt*
- Summary:
  - Edit DEFAULT\_VERSION in *LUSTRE-VERSION-GEN*
  - Use well-formed tag (e.g. 2.8.0\_9acme)



# Release Candidate Abomination


2.8.0-RC1

2.8.0-RC2

2.8.0-RC3

2.8.0-RC4

2.8.0-RC5



lustre-2.8.0-<kernel>.x86\_64.rpm

See rationale in LU-7699

# Consistent Versions

- lustre-2.8.0\_9acme\_1\_gdf14e91.tar.gz
- lustre-2.8.0\_9acme\_1\_gdf14e91-1.src.rpm
- lustre-2.8.0\_9acme\_1\_gdf14e91-1.x86\_64.rpm




No kernel version in the Release!

Requires completion of LU-5614

## More Consistent Versions


- lustre-2.8.0.tar.gz
- lustre-2.8.0-1.src.rpm
- lustre-2.8.0-1.x86\_64.rpm



No kernel version in the Release!

Requires completion of LU-5614

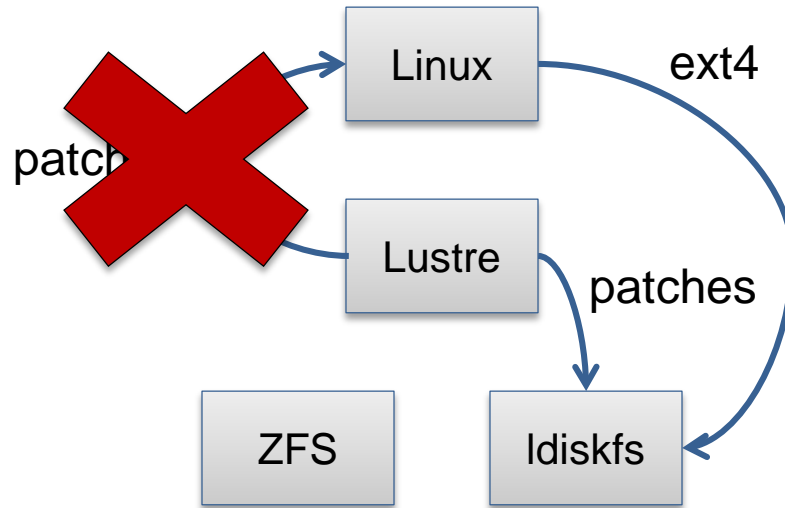
## LU-5614 – use %kernel\_module\_package for weak-updates



LU-5614 under way for over 18 months!

# Eliminate the Custom Lustre Kernel Requirement

- LU-20 - patchless server kernel
- LU-684 – replace dev\_rdonly kernel patch with dm-flakey



# Community Build Farm for Lustre



# Goals

---

- For each commit submitted to Gerrit
  - Build on various distributions and architectures
  - Publish build results and products in a timely fashion
  - Test functionality and publish results
- Have a community driven build and test platform that can be rapidly adapted to changing needs

# OpenSFS Build Farm for Lustre

---

- An automated build system backed by Buildbot
  - <http://buildbot.net>
- Configuration open to community collaboration on github
- Hosted on Amazon Web Service (AWS)

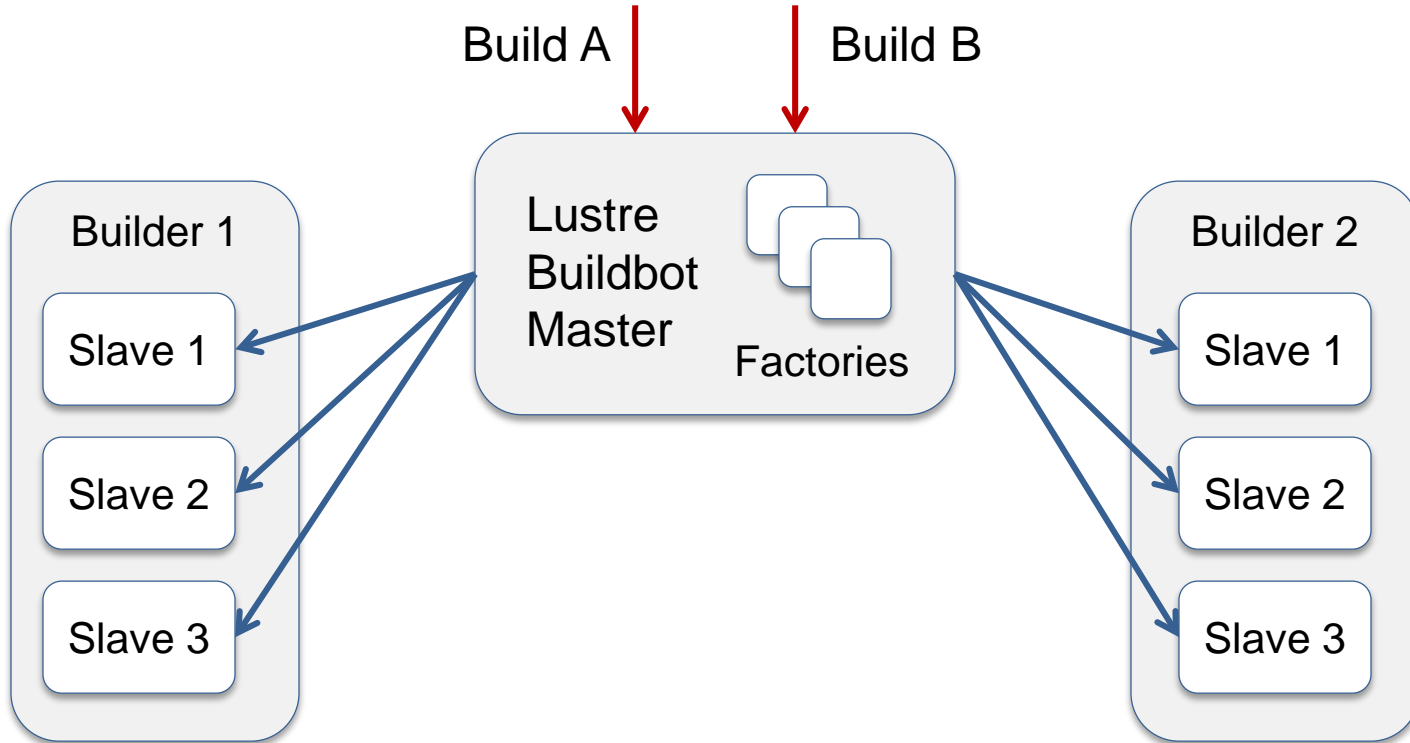


# Why Buildbot?

---

- Flexible and Extendable
  - Open Source
  - Good interoperability with other tools
  - Developer defined workflow
  - If you can script it, Buildbot can run it
- Out of the box support for Gerrit and Amazon EC2
- Hosts a web page to view builds out of the box
- Learned from ZFSonLinux experience with Buildbot

# Buildbot Architecture



# Waterfall

|                  |               | CentOS 6.7 x86_64 (BUILD)<br>build successful shell_1  | CentOS 7.2 x86_64 (BUILD)<br>failed shell_1 shell_2   | SLES 12.1 x86_64 (BUILD)<br>build successful   | Ubuntu 12.04 x86_64 (BUILD)<br>build successful  | Ubuntu 14.04 x86_64 (BUILD)<br>build successful   |
|------------------|---------------|--|---|--|--|---|
| last build       |               | building<br>ETA in<br>~ 53 mins<br>at 20:08<br>ETA in<br>~ 50 mins<br>at 20:05<br>ETA in<br>~ 53 mins<br>at 20:08<br>4 pending | building<br>ETA in<br>~ 4 mins<br>at 19:20<br>ETA in<br>~ 4 mins<br>at 19:20<br>ETA in<br>~ 7 mins<br>at 19:22<br>4 pending | building<br>ETA in<br>~ 7 mins<br>at 19:22<br>ETA in<br>~ 29 mins<br>at 19:44<br>3 pending | building<br>ETA in<br>~ 3 mins<br>at 19:18<br>ETA in<br>~ 3 mins<br>at 19:18<br>ETA in<br>~ 16 mins<br>at 19:31<br>2 pending | building<br>ETA in<br>~ 5 mins<br>at 19:20<br>ETA in<br>~ 5 mins<br>at 19:20<br>ETA in<br>~ 5 mins<br>at 19:20<br>1 pending |
| current activity |               |  |   |  |  |   |
| EDT              | changes       | CentOS 6.7 x86_64 (BUILD)  | CentOS 7.2 x86_64 (BUILD)   | SLES 12.1 x86_64 (BUILD)   | Ubuntu 12.04 x86_64 (BUILD)  | Ubuntu 14.04 x86_64 (BUILD)   |
| 19:14:58         |               | installing dependencies<br><a href="#">studio</a>  | building lustre<br><a href="#">studio</a>   | installing dependencies<br><a href="#">studio</a>  | building spl and zfs<br><a href="#">studio</a>   | building lustre<br><a href="#">studio</a>   |
| 19:14:43         |               |  | cloned<br><a href="#">studio</a>  | Build 538  |  |   |
| 19:14:05         |               |  |   | building lustre<br><a href="#">studio</a>  |  |   |
| 19:13:08         |               |  |   | cloned<br><a href="#">studio</a>   |  | cloned<br><a href="#">studio</a>  |
| 19:12:11         |               |  | built spl and zfs<br><a href="#">studio</a>   |  | installed dependencies<br><a href="#">studio</a>   | cloned<br><a href="#">studio</a>  |
|                  |               |  |   |  |  | installed dependencies<br><a href="#">studio</a>  |
|                  |               |  |   |  |  | Build 569   |
|                  |               |  |   |  |  | building lustre<br><a href="#">studio</a>   |
|                  |               |  |   |  |  | cloned<br><a href="#">studio</a>  |
|                  |               |  |   |  |  | installed dependencies<br><a href="#">studio</a>  |
|                  |               |  |   |  |  | Build 568   |
|                  |               |  |   |  |  | building lustre<br><a href="#">studio</a>   |
|                  |               |  |   |  |  | cloned<br><a href="#">studio</a>  |
|                  |               |  |   | installed dependencies<br><a href="#">studio</a>   | Build 534  | installed dependencies<br><a href="#">studio</a>  |
| 19:11:08         |               |  |   | Build 537  | building lustre<br><a href="#">studio</a>  | Build 567   |
| 19:10:03         | James Simmons |  |   |  | cloned<br><a href="#">studio</a>   |   |
|                  |               |  |   |  | installed dependencies<br><a href="#">studio</a>   |   |
|                  |               |  |   |  | Build 533  |   |

## Console View

Legend: Passed Failed Warnings Failed Again Running Exception Offline No data

Go

|   |               |  |  |  |  |
|---|---------------|--|--|--|--|
| LU-4423 Inet: Drop unnecessary wrapper function srpc_post_active_rqtbuf |               |  |  |  |  |
| c3cc07426f5b...   | James Simmons | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> | <span style="background-color: #dc3545; width: 100px; height: 15px;"></span> | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> |
| ce72a6a25bf8...   | James Simmons | <input type="text"/>   | <input type="text"/>   | <input type="text"/>   | <input type="text"/>   |
| LU-4423 Inet: remove memset(0) after LIBCFS_ALLOC                       |               |  |  |  |  |
| 6af76bcc1787...   | Li Dongyang   | <input type="text"/>   | <input type="text"/>   | <input type="text"/>   | <input type="text"/>   |
| LU-6215 o2ibld: rdma_create_id() takes extra parameter                  |               |  |  |  |  |
| 5fbdfa85912e...   | Li Dongyang   | <input type="text"/>   | <span style="background-color: #ffc107; width: 100px; height: 15px;"></span> | <input type="text"/>   | <span style="background-color: #ffc107; width: 100px; height: 15px;"></span> |
| LU-6215 gss: key->payload.data is an array in 4.4+ kernels              |               |  |  |  |  |
| 24364e4b3af2...   | James Simmons | <span style="background-color: #ffc107; width: 100px; height: 15px;"></span> | <span style="background-color: #ffc107; width: 100px; height: 15px;"></span> | <input type="text"/>   | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> |
| LU-6215 lprocs: handle seq_printf api change                            |               |  |  |  |  |
| a23af8834e5b...   | James Simmons | <span style="background-color: #ffc107; width: 100px; height: 15px;"></span> | <span style="background-color: #ffc107; width: 100px; height: 15px;"></span> | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> |
| LU-6163 kernel: use do_div64() for 64 bit divides                       |               |  |  |  |  |
| 7c1dff608f0...  | James Nunez   | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> | <span style="background-color: #dc3545; width: 100px; height: 15px;"></span> | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> | <span style="background-color: #28a745; width: 100px; height: 15px;"></span> |
| LU-0000 tests: Resume parallel-grouplock testing                        |               |  |  |  |  |

# Current Builders

- Currently building patch sets
  - CentOS 6.7 x86\_64 (zfs, ldiskfs)
  - CentOS 7.2 x86\_64 (zfs, ldiskfs)
  - SLES 12.1 x86\_64 (zfs)
  - Ubuntu 12.04 x86\_64 (zfs)
  - Ubuntu 14.04 x86\_64 (zfs)
- Community contributed build slaves!
  - PowerPC 64-bit (coming)

# Build Turnaround Time

---

- Timely response for single build
  - Worse case, 40 minutes
  - Best case, 15 minutes
  
- Outstanding Issues
  - Multiple Gerrit commits
  - Slave startup can be improved
  - Buildbot sequentializes slave startup

# Future Plans

---

- Publish build products
  - Promote independent testing
- Supply build status in Gerrit review
- Automated testing
- Build official packages for all tags
- Community development and maintenance!!

# Summary

---

- Versions and Packages
  - Consistent versions
  - Lustre builds through standard tools (Koji)
  - Weak-updates
  - Change to third-party versions
- Buildbot
  - Email: [buildbot-admin@lustre.org](mailto:buildbot-admin@lustre.org)
  - Web: <http://build.lustre.org>
  - Repository and Issue Tracker: <http://github.com/opensfs/lustre-buildbot-config>
  - Community driven
  - Hosted by OpenSFS in AWS
- Questions? Comments? Ideas? Get involved!



