



Removal of Dead Code Solution Architecture

Added by Richard Henwood, last edited by Richard Henwood on Feb 27, 2014

Introduction

Removal of dead code from the Lustre* file system code base is an important and urgent task for reasons stated in the Problem Statement of the [Statement of Work](#). This document is concerned with describing the Use Cases, Requirements and System Analysis.

A important product of this project is developing tools (as plugins for Clang) that will be of value to all Lustre developers into the future.

Requirements Description

Lustre software is developed in a open and collaborative way. Code contributors have different levels of experience. Code reviewers have apply different style guidelines. Lustre software development community needs to:

- Apply style guidelines evenly to contributed patches.
- Reduce the count of unused or incorrectly lines of code.

Use Cases

A software developer needs to provide a patch for the Lustre software code base.

Checks for code style requirements are completed automatically.

The developer has deviations from the style requirements made available to them in the context of their patch.

A software developer needs to reduce the lines of unused code.

A tool is available to short list code that maybe unused under all circumstances.

The developer uses the output of the tool to create patches removing unused code.

A software developer needs to ensure header files are correctly included.

A tool is available to analyse the Lustre source code and identify redundant header inclusions.

The developer uses the output of the tool to create patches removing unused code.

Solution Proposal

Automated style-checker

A automated style-checker service will be implemented. style-checker will use the perl script (checkpatch.pl) in the Lustre software to check a patch against a set of best practice style guidelines. When an author pushes a patch into Gerrit the patch must be verified for style issues using the checkpatch.pl script. Deviations from the style must be provided in line with the patch in Gerrit. Example output is:

```

maiboo Patch Set 17: Verified tests received by maiboo, run on CentOS release 6.4/x86_64. (https://maiboo.whartoncloud.com/test_sessions/00000400-2)
Andreas Dilger Patch Set 17: Looks good to me, approved (6 inline comments) There are lots of places that the whitespace could be cleaned up more (just a
Mike Pershin Patch Set 17: Looks good to me, but someone else must approve
Andrew Korty Patch Set 17: Looks good to me, but someone else must approve
John L. Hammond Patch Set 17: I would prefer that you didn't submit this Needs rebasing.
HPDD Checkpatch

Patch Set 17:

(19 comments)

19 style warnings. For more details please see https://wiki.hpdd.intel.com/display/PUB/Test+Coding+Style.

lustre/include/dt_object.h

Line 1121: (style) line over 80 characters
Line 1311: (style) line over 80 characters

lustre/include/md_object.h

Line 67: (style) code indent should use tabs where possible

lustre/mdd/mdd_dir.c

Line 500: (style) please, no space before tabs
Line 507: (style) "(foo*)" should be "(foo *)"
Line 532: (style) line over 80 characters

lustre/mdd/mdd_object.c

Line 258: (style) space required after that ',' (ctx:VxV)
Line 269: (style) code indent should use tabs where possible
Line 290: (style) space required after that ',' (ctx:VxV)

lustre/mdt/mdt_handler.c

Line 4682: (style) code indent should use tabs where possible
Line 4683: (style) code indent should use tabs where possible

```

Figure 1: An example of style-checker commenting against a patch in Gerrit

```

1342 LASSERT(dt);
1343 LASSERT(dt->do_ops);
1344 LASSERT(dt->do_ops->do_xattr_del);
1345 return dt->do_ops->do_xattr_del(env, dt, name, th, capa);
1346 }
1347
1348 static inline int dt_declare_xattr_set(const struct lu_env *env,
1349 struct dt_object *dt,
1350 const struct lu_buf *buf,
1351 const char *name, int fl,
1352 struct thandle *th)
1353 {
1354 LASSERT(dt);
1355 LASSERT(dt->do_ops);
1356 LASSERT(dt->do_ops->do_declare_xattr_set);
1357 return dt->do_ops->do_declare_xattr_set(env, dt, buf, name, fl, th);
1358 }
1359
1360 static inline int dt_xattr_set(const struct lu_env *env,
1361 struct dt_object *dt, const struct lu_buf *buf,
1362 const char *name, int fl, struct thandle *th,
1363 struct lustre_capa *capa)
1364 {
1365 LASSERT(dt);
1366 LASSERT(dt->do_ops);
1367 LASSERT(dt->do_ops->do_xattr_set);
1368 return dt->do_ops->do_xattr_set(env, dt, buf, name, fl, th, capa);
1369 }
1370
1371 static inline int dt_xattr_get(const struct lu_env *env,
1372 struct dt_object *dt, struct lu_buf *buf,
1373 const char *name, struct lustre_capa *capa)
1374 {
1375 LASSERT(dt);
1376 LASSERT(dt->do_ops);
1377 LASSERT(dt->do_ops->do_xattr_get);
1378 return dt->do_ops->do_xattr_get(env, dt, buf, name, capa);
1379 }
1380
1292 LASSERT(dt);
1293 LASSERT(dt->do_ops);
1294 LASSERT(dt->do_ops->do_xattr_del);
1295 return dt->do_ops->do_xattr_del(env, dt, name, th,
1296 );
1297
1298 static inline int dt_declare_xattr_set(const struct lu_env
1299 *env, struct dt_object *dt,
1300 const struct lu_buf *buf,
1301 const char *name,
1302 struct thandle *th)
1303 {
1304 LASSERT(dt);
1305 LASSERT(dt->do_ops);
1306 LASSERT(dt->do_ops->do_declare_xattr_set);
1307 return dt->do_ops->do_declare_xattr_set(env, dt,
1308 buf, name, fl, th);
1309 }
1310 static inline int dt_xattr_set(const struct lu_env *env, s
1311 struct dt_object *dt, const struct lu_buf *buf,
1312 struct thandle *th)
1313 {
1314 LASSERT(dt);
1315 LASSERT(dt->do_ops);
1316 LASSERT(dt->do_ops->do_xattr_set);
1317 return dt->do_ops->do_xattr_set(env, dt, buf, nam
1318 );
1319 }
1320 static inline int dt_xattr_get(const struct lu_env *env,
1321 struct dt_object *dt, struct lu_buf *buf, st
1322 struct thandle *th, const char *name)
1323 {
1324 LASSERT(dt);
1325 LASSERT(dt->do_ops);
1326 LASSERT(dt->do_ops->do_xattr_get);
1327 return dt->do_ops->do_xattr_get(env, dt, buf, nam
1328 );
1329 }
1330 static inline int dt_xattr_list(const struct lu_env *env,

```

Figure 2: An example of in-line style comment from style-checker against a specific part of the patch.

Static Analysis to Identify Unused Code

We will develop a Clang plugin to identify lines of code that are never run (i.e. unused and redundant). By pairing this tool with a extensive Lustre workload, lines that are never run under the workload are revealed. This tool does not guarantee that code will never run under any circumstances. The tool provides a short-list of candidates for investigation and potential deletion by a expert developer.

Static Analysis to Identify Omitted use of 'const' in Function Prototypes

Function prototypes should have 'const' for unmodified pointers. Function prototypes that do not have 'const' for unmodified pointers need to be corrected. The CLANG and LLVM compiler tool chain will be extended with a plugin to resolve instances of incorrect const usage.

Static Analysis to Identify Redundant Header Inclusions

Files that have redundant header inclusions need to be corrected. The CLANG and LLVM compiler tool chain will be extended with a plugin to resolve instances of redundant header includes.

Sparse for Linux Specific Checks

Sparse is an open source tool designed to find coding faults in the Linux kernel source. The Sparse tool will be run against the Lustre source code. A large part of the Lustre source code is Linux modules which should follow Linux Sparse guidelines. If any issues are discovered they will be short-listed for correction.

Requirements Analysis

Automated operation for style-checker.

- The style-checker must be run immediately a new patch is pushed into Gerrit.
- The style-checker will be loosely coupled to Gerrit running on it's own machine and communicating with Gerrit over ssh.
- The style-checker will not score patches with -1 if there are issues with the style.
- If the style-checker does not execute on a newly pushed patch (because the style-checker host is down for example), the style-checker will not process any patches it missed when it returns to working order.

Static Analysis

- Static analysis tools will be run ad-hoc.
- Static analysis tools will be available to a developer to run locally.

Static Analysis tools include:

- Identifying unused code.
- Identifying incorrect use of const in function prototypes.
- Identifying incorrect header inclusions.
- Identifying Linux kernel specific coding faults with Sparse.

Acceptance Criteria

Automated check-patch operation will be demonstrated.

Ad-hoc un-used code analysis will be demonstrated.

Ad-hoc incorrect use of const in function prototypes will be demonstrated.

Ad-hoc incorrect header inclusion analysis will be demonstrated.

Ad-hoc Linux specific coding faults analysis will be demonstrated.

[Like](#) Be the first to like this

2 Comments



Andreas Dilger

It would be desirable for the style-checking tool to also review patches that were submitted while it is disconnected from Gerrit. Is it possible to have the list of patches to inspect based on a Gerrit query that can be run periodically instead of watching for new patches to be reviewed?



John Hammond

Yes. That was always the plan.
