Architect of an Open World™

Lustre Static Code Analysis with Coverity

September 25, 2012

Sebastien Buisson

Parallel File Systems
Extreme Computing R&D

Static Code Analysis with Coverity

- Why static code analysis is useful?
- Tool for analysis: Coverity
- Coverity applied to Lustre
- Defects found by Coverity
- Benefits for the whole Lustre Community

Why static code analysis is useful?

Architect of an Open World™

Why static code analysis is useful?

From http://wiki.whamcloud.com/display/PUB/Project+Ideas

“Run Lustre code through static analysis tools to identify potential latent bugs in the Lustre code. These are often hard to find through testing, and easily fixed once found.”

From W. S. Humphrey, "Using a Defined and Measured Personal Software Process," IEEE Software, May, 1996

“Even experienced programmers typically make a mistake for every seven to ten lines of code they develop.”

Tool for analysis: Coverity

Architect of an Open World™

Tool for analysis: Coverity

- How it works

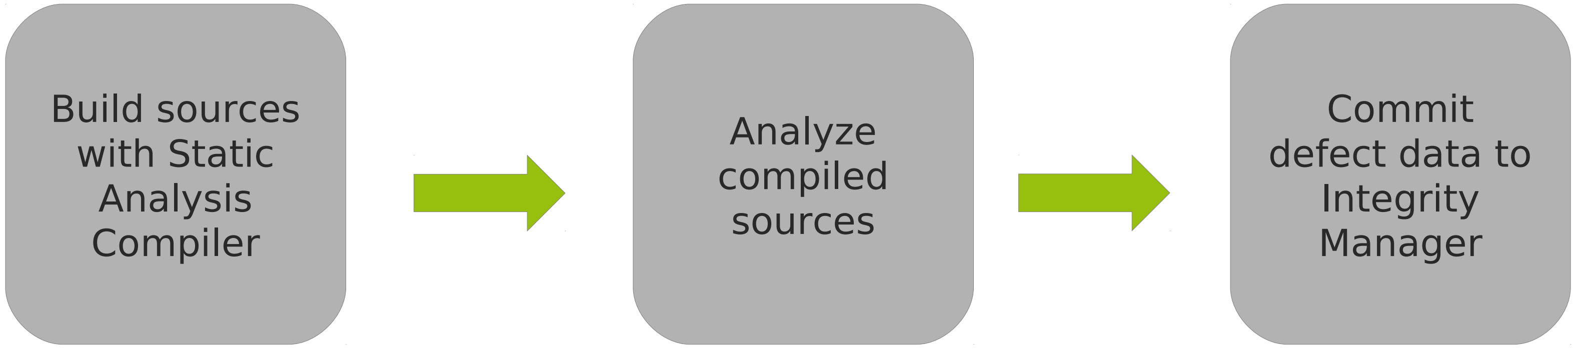
Build sources with Static Analysis Compiler
Analyze compiled sources
Commit defect data to Integrity Manager

- Various checkers involved:
  - STRING_SIZE
  - RESOURCE_LEAK
  - UNINIT
  - ...

Tool for analysis: Coverity

What do we get

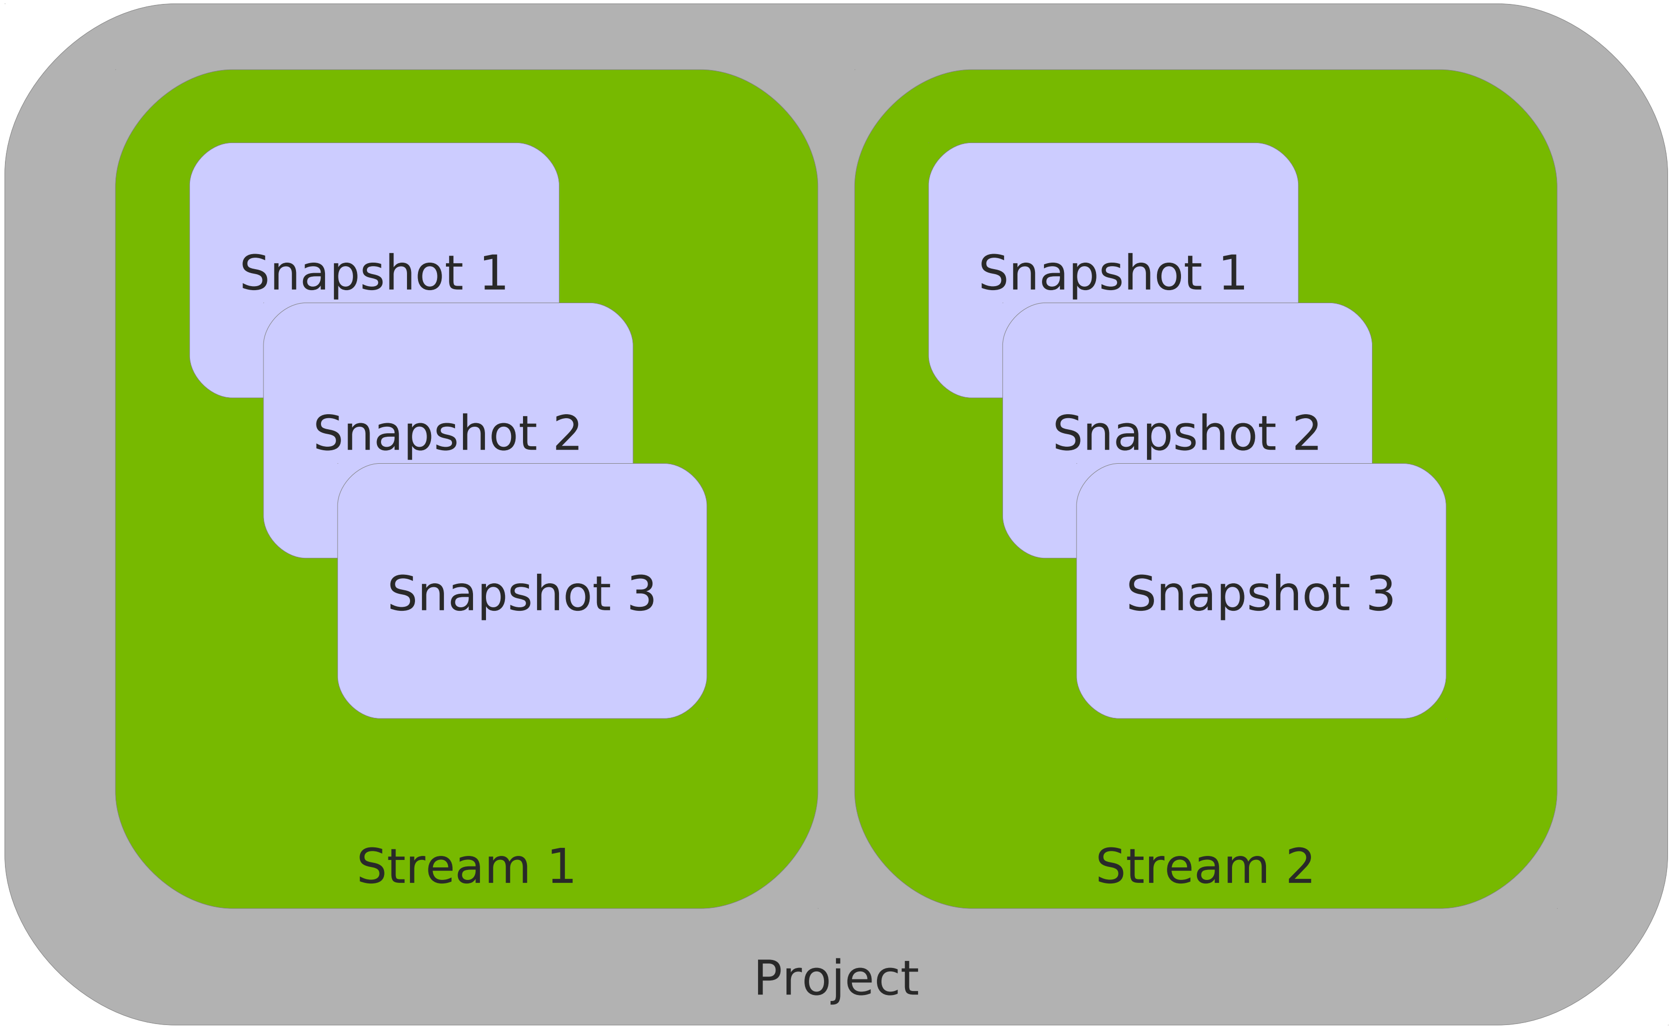

Tool for analysis: Coverity

- What do we get
  - GUI

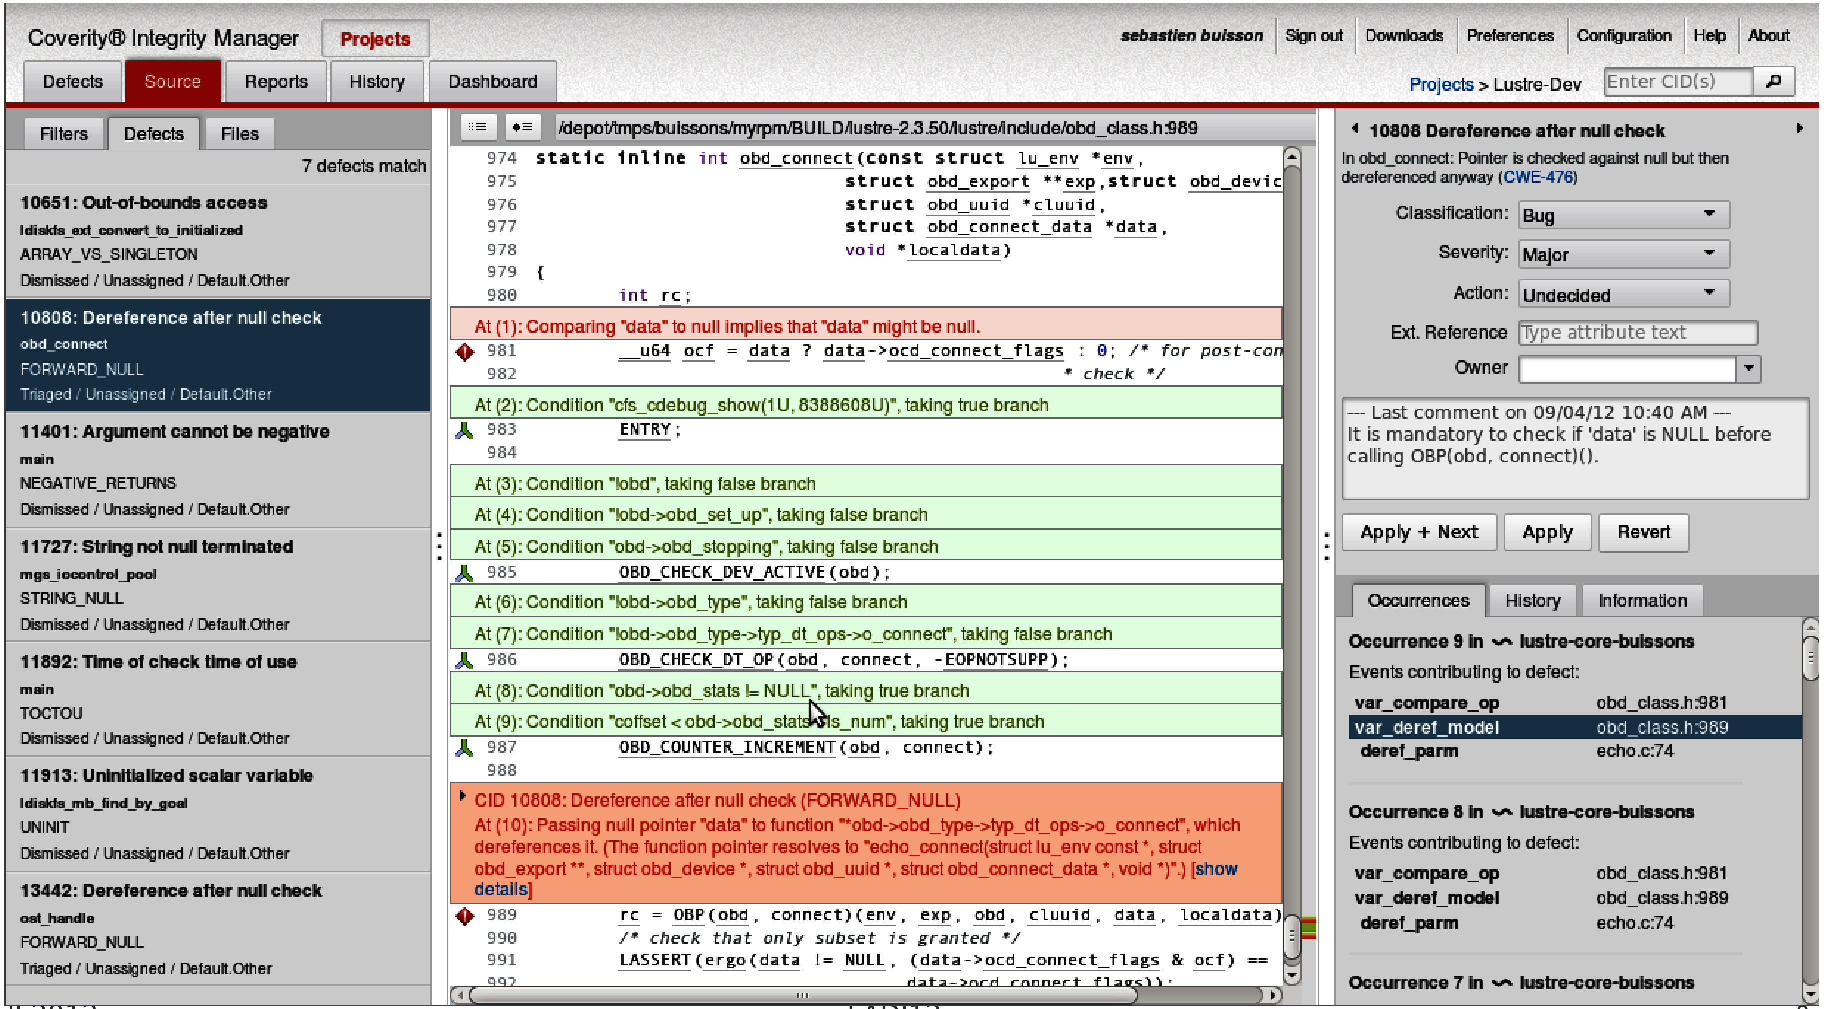

Coverity® Integrity Manager | **Projects**

Defects | Source | Reports | History | Dashboard

sebastien buisson | Sign out | Downloads | Preferences | Configuration | Help | About

Projects > Lustre-Dev | Enter CID(s)

Filters | Defects | Files

7 defects match

10651: Out-of-bounds access
ldiskfs_ext_convert_to_initialized
ARRAY_VS_SINGLETON
Dismissed / Unassigned / Default.Other

10808: Dereference after null check
obd_connect
FORWARD_NULL
Triaged / Unassigned / Default.Other

11401: Argument cannot be negative
main
NEGATIVE_RETURNS
Dismissed / Unassigned / Default.Other

11727: String not null terminated
mgs_iocontrol_pool
STRING_NULL
Dismissed / Unassigned / Default.Other

11892: Time of check time of use
main
TOCTOU
Dismissed / Unassigned / Default.Other

11913: Uninitialized scalar variable
ldiskfs_mb_find_by_goal
UNINIT
Dismissed / Unassigned / Default.Other

13442: Dereference after null check
ost_handle
FORWARD_NULL
Triaged / Unassigned / Default.Other

/depot/tmps/buissons/myrpm/BUILD/lustre-2.3.50/lustre/include/obd_class.h:989

```
974 static inline int obd_connect(const struct lu_env *env,
975                               struct obd_export **exp,struct obd_devic
976                               struct obd_uuid *cluuid,
977                               struct obd_connect_data *data,
978                               void *localdata)
979 {
980         int rc;
At (1): Comparing "data" to null implies that "data" might be null.
981         __u64 ocf = data ? data->ocd_connect_flags : 0; /* for post-con
982                                                  * check */
At (2): Condition "cfs_cdebug_show(1U, 8388608U)", taking true branch
983         ENTRY;
984
At (3): Condition "!obd", taking false branch
At (4): Condition "!obd->obd_set_up", taking false branch
At (5): Condition "obd->obd_stopping", taking false branch
985         OBD_CHECK_DEV_ACTIVE(obd);
At (6): Condition "!obd->obd_type", taking false branch
At (7): Condition "!obd->obd_type->typ_dt_ops->o_connect", taking false branch
986         OBD_CHECK_DT_OP(obd, connect, -EOPNOTSUPP);
At (8): Condition "obd->obd_stats != NULL", taking true branch
At (9): Condition "coffset < obd->obd_stats->ls_num", taking true branch
987         OBD_COUNTER_INCREMENT(obd, connect);
988
CID 10808: Dereference after null check (FORWARD_NULL)
At (10): Passing null pointer "data" to function "*obd->obd_type->typ_dt_ops->o_connect", which dereferences it. (The function pointer resolves to "echo_connect(struct lu_env const *, struct obd_export **, struct obd_device *, struct obd_uuid *, struct obd_connect_data *, void *)".) [show details]
989         rc = OBP(obd, connect)(env, exp, obd, cluuid, data, localdata)
990         /* check that only subset is granted */
991         LASSERT(ergo(data != NULL, (data->ocd_connect_flags & ocf) ==
992                      data->ocd_connect_flags));
```

10808 Dereference after null check

In obd_connect: Pointer is checked against null but then dereferenced anyway (CWE-476)

Classification:	Bug
Severity:	Major
Action:	Undecided
Ext. Reference	Type attribute text
Owner	

--- Last comment on 09/04/12 10:40 AM ---
It is mandatory to check if 'data' is NULL before calling OBP(obd, connect)().

Apply + Next | Apply | Revert

Occurrences | History | Information

Occurrence 9 in lustre-core-buissons
Events contributing to defect:

var_compare_op	obd_class.h:981
var_deref_model	obd_class.h:989
deref_parm	echo.c:74

Occurrence 8 in lustre-core-buissons
Events contributing to defect:

var_compare_op	obd_class.h:981
var_deref_model	obd_class.h:989
deref_parm	echo.c:74

Occurrence 7 in lustre-core-buissons

Tool for analysis: Coverity

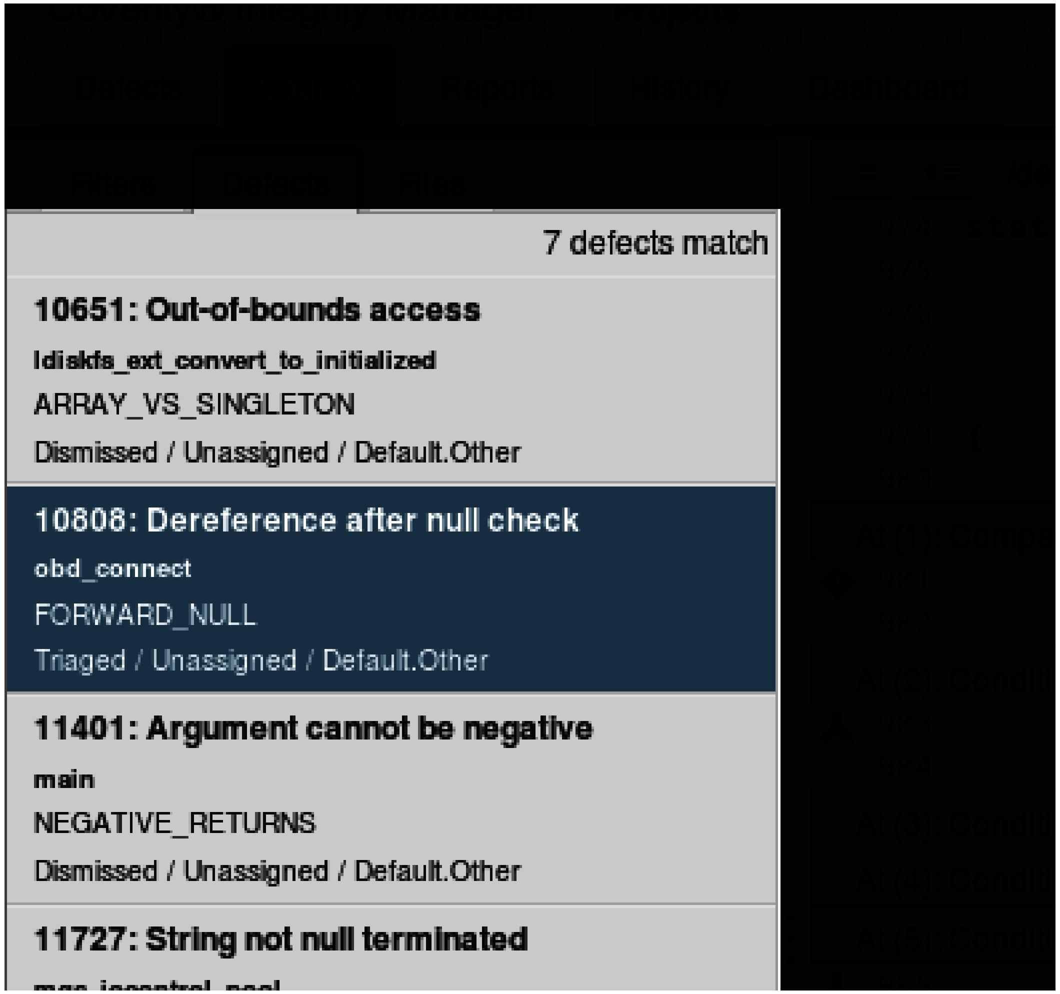
7 defects match

10651: Out-of-bounds access
ldiskfs_ext_convert_to_initialized
ARRAY_VS_SINGLETON
Dismissed / Unassigned / Default.Other

10808: Dereference after null check
obd_connect
FORWARD_NULL
Triaged / Unassigned / Default.Other

11401: Argument cannot be negative
main
NEGATIVE_RETURNS
Dismissed / Unassigned / Default.Other

11727: String not null terminated

Tool for analysis: Coverity

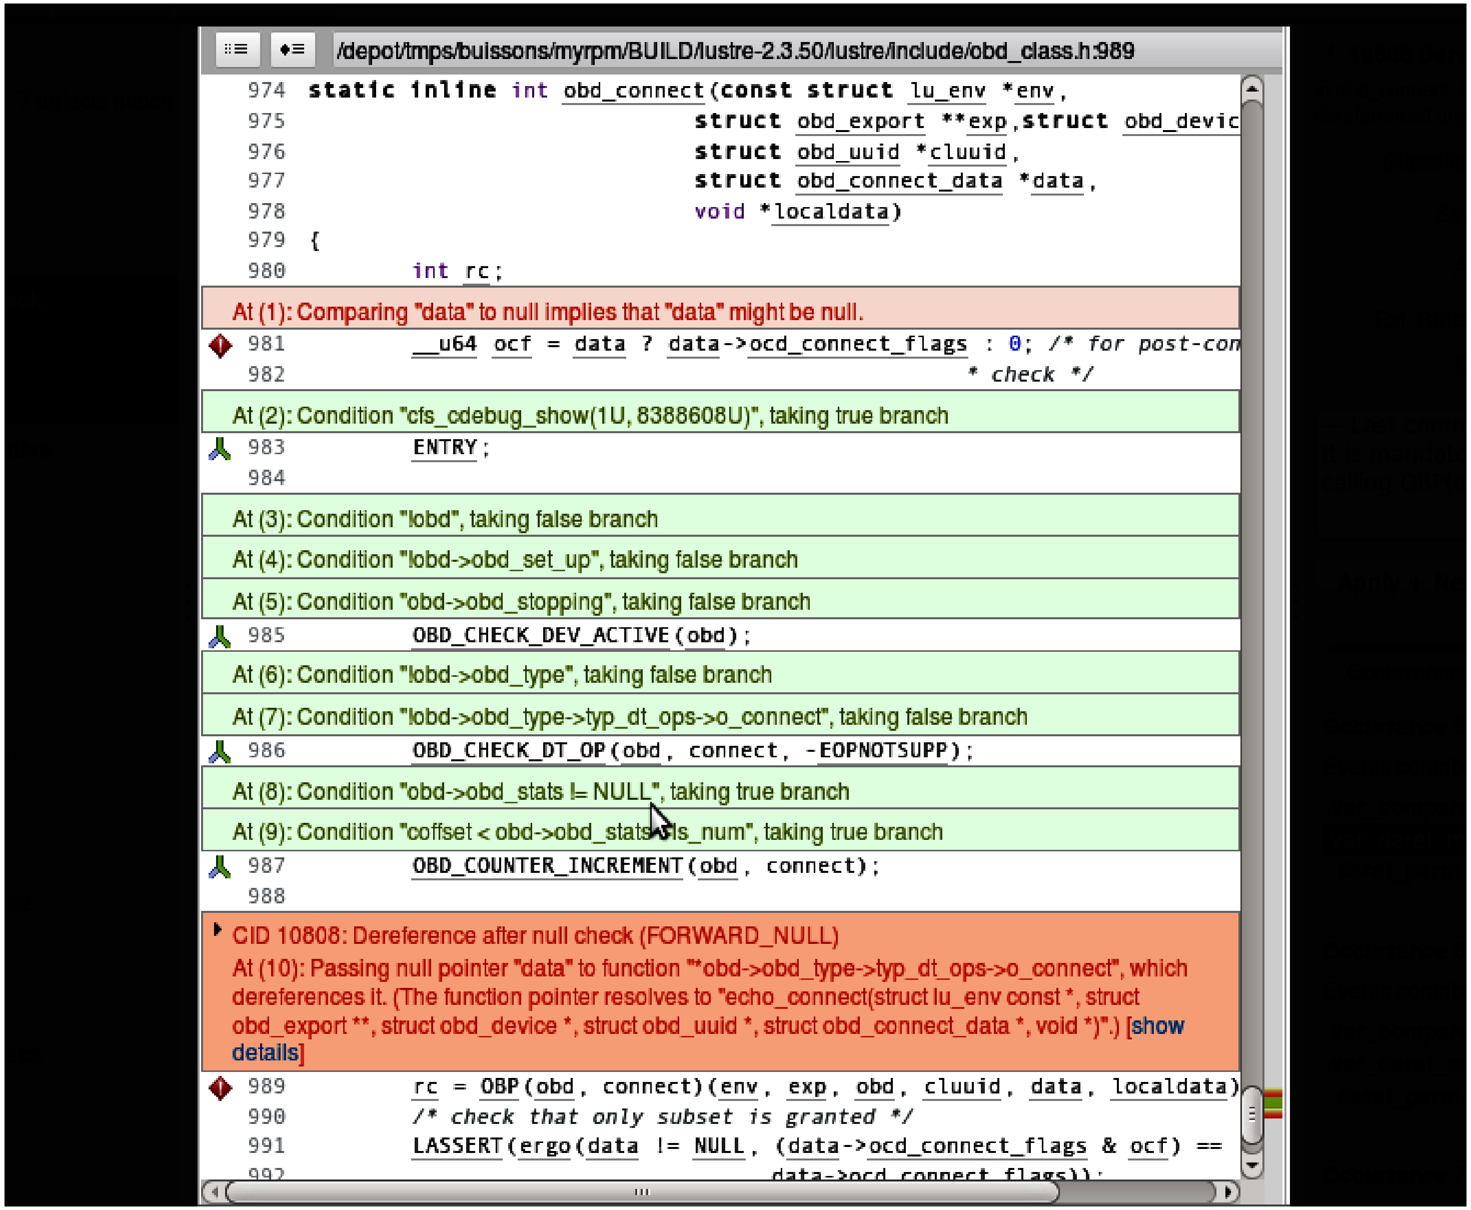

/depot/tmps/buissons/myrpm/BUILD/lustre-2.3.50/lustre/include/obd_class.h:989

```
974 static inline int obd_connect(const struct lu_env *env,
975                               struct obd_export **exp,struct obd_devic
976                               struct obd_uuid *cluuid,
977                               struct obd_connect_data *data,
978                               void *localdata)
979 {
980         int rc;
```

At (1): Comparing "data" to null implies that "data" might be null.

```
981         __u64 ocf = data ? data->ocd_connect_flags : 0; /* for post-con
982                                                          * check */
```

At (2): Condition "cfs_cdebug_show(1U, 8388608U)", taking true branch

```
983         ENTRY;
984
```

At (3): Condition "!obd", taking false branch

At (4): Condition "!obd->obd_set_up", taking false branch

At (5): Condition "obd->obd_stopping", taking false branch

```
985         OBD_CHECK_DEV_ACTIVE(obd);
```

At (6): Condition "!obd->obd_type", taking false branch

At (7): Condition "!obd->obd_type->typ_dt_ops->o_connect", taking false branch

```
986         OBD_CHECK_DT_OP(obd, connect, -EOPNOTSUPP);
```

At (8): Condition "obd->obd_stats != NULL", taking true branch

At (9): Condition "coffset < obd->obd_stats->ls_num", taking true branch

```
987         OBD_COUNTER_INCREMENT(obd, connect);
988
```

CID 10808: Dereference after null check (FORWARD_NULL)

At (10): Passing null pointer "data" to function "*obd->obd_type->typ_dt_ops->o_connect", which dereferences it. (The function pointer resolves to "echo_connect(struct lu_env const *, struct obd_export **, struct obd_device *, struct obd_uuid *, struct obd_connect_data *, void *)".) [show details]

```
989         rc = OBP(obd, connect)(env, exp, obd, cluuid, data, localdata)
990         /* check that only subset is granted */
991         LASSERT(ergo(data != NULL, (data->ocd_connect_flags & ocf) ==
992                      data->ocd_connect_flags));
```

Tool for analysis: Coverity

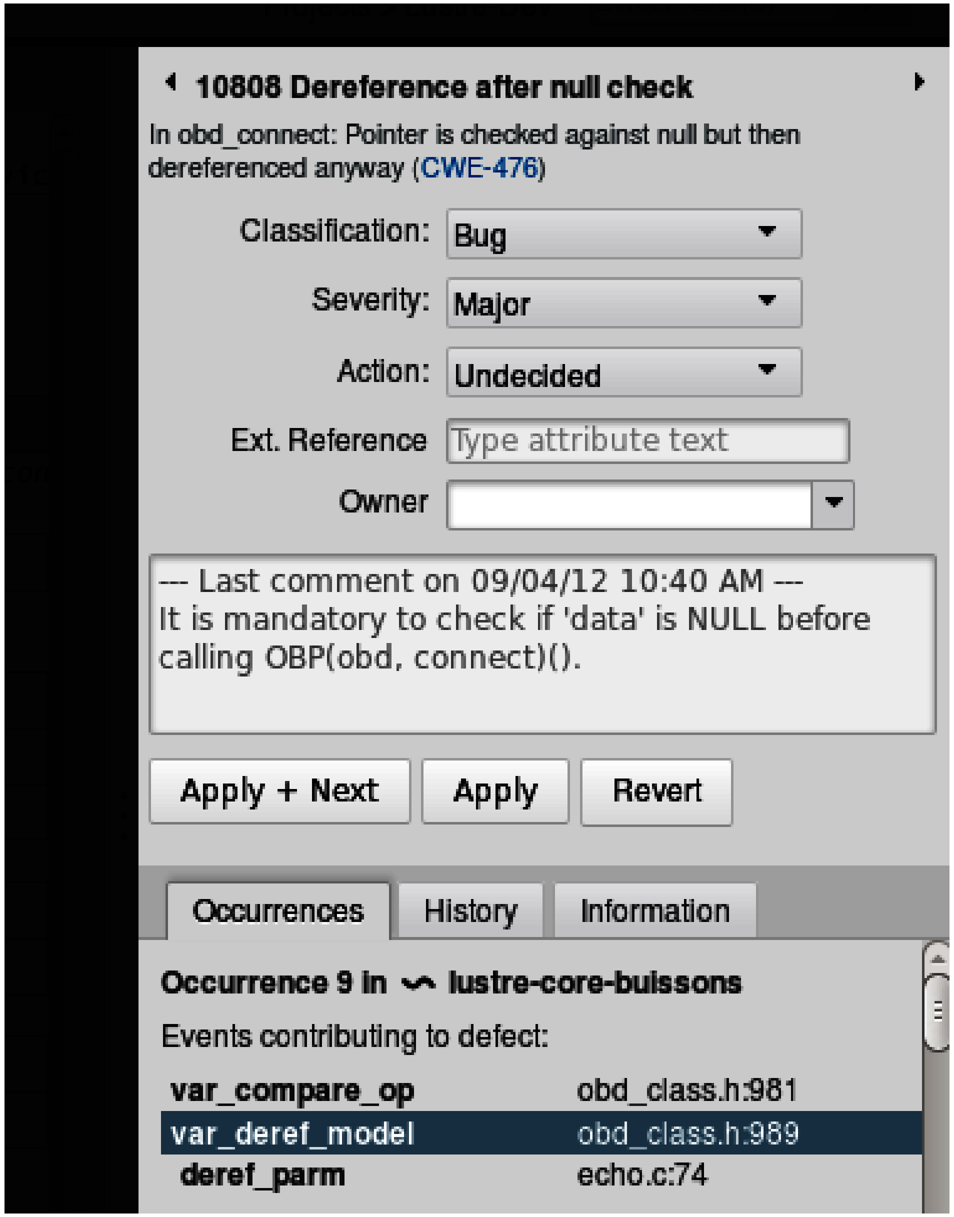

10808 Dereference after null check

In obd_connect: Pointer is checked against null but then dereferenced anyway (CWE-476)

Classification:	Bug
Severity:	Major
Action:	Undecided
Ext. Reference	Type attribute text
Owner	

--- Last comment on 09/04/12 10:40 AM ---
It is mandatory to check if 'data' is NULL before calling OBP(obd, connect)().

Occurrence 9 in lustre-core-buissons

Events contributing to defect:

var_compare_op	obd_class.h:981
var_deref_model	obd_class.h:989
deref_parm	echo.c:74

Tool for analysis: Coverity

- Defect categories
  - API usage errors
  - Code maintainability issues
  - Concurrent data access violations
  - Control flow issues
  - Error handling issues
  - Incorrect expression
  - Integer handling issues
  - Memory - corruptions
  - Memory - illegal accesses
  - Null pointer dereferences
  - Program hangs
  - Resource leaks
  - Security best practices violations
  - Uninitialized variables

Coverity applied to Lustre

Architect of an Open World™

Coverity applied to Lustre

Necessary setup for Lustre

- Models

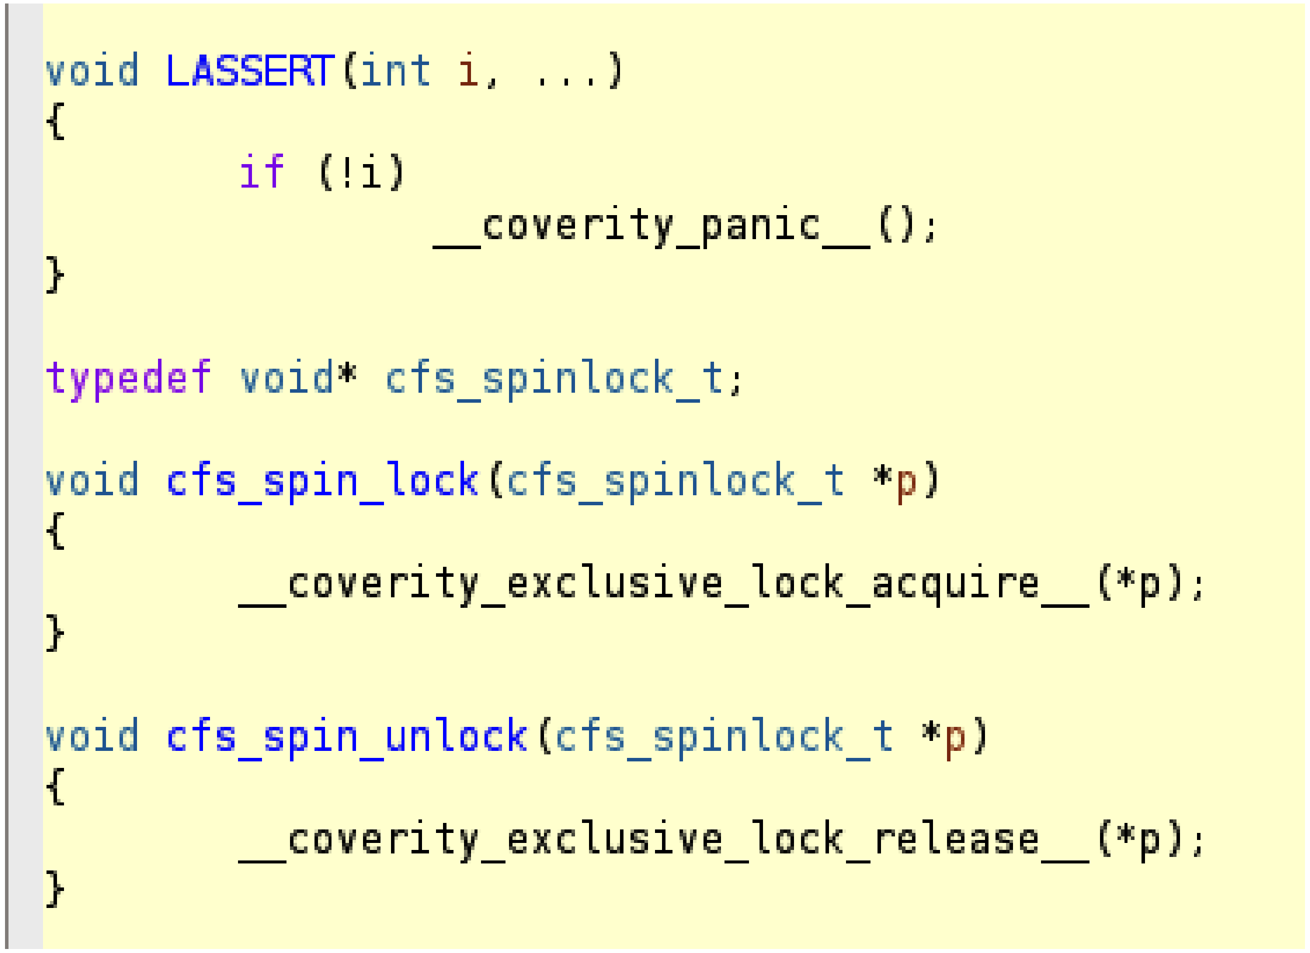

```
void LASSERT(int i, ...)
{
        if (!i)
                __coverity_panic__();
}

typedef void* cfs_spinlock_t;

void cfs_spin_lock(cfs_spinlock_t *p)
{
        __coverity_exclusive_lock_acquire__(*p);
}

void cfs_spin_unlock(cfs_spinlock_t *p)
{
        __coverity_exclusive_lock_release__(*p);
}
```

Coverity applied to Lustre

How we proceed

- Main work on Lustre Master branch + 2.1 in parallel

- Current status for master:
  - Initial analysis done (v2_3_50)
  - Next steps: diff with new tags on master

- Hard to analyze specific components or features

Defects found by Coverity

Architect of an Open World™

Defects found by Coverity

- Statistics on master

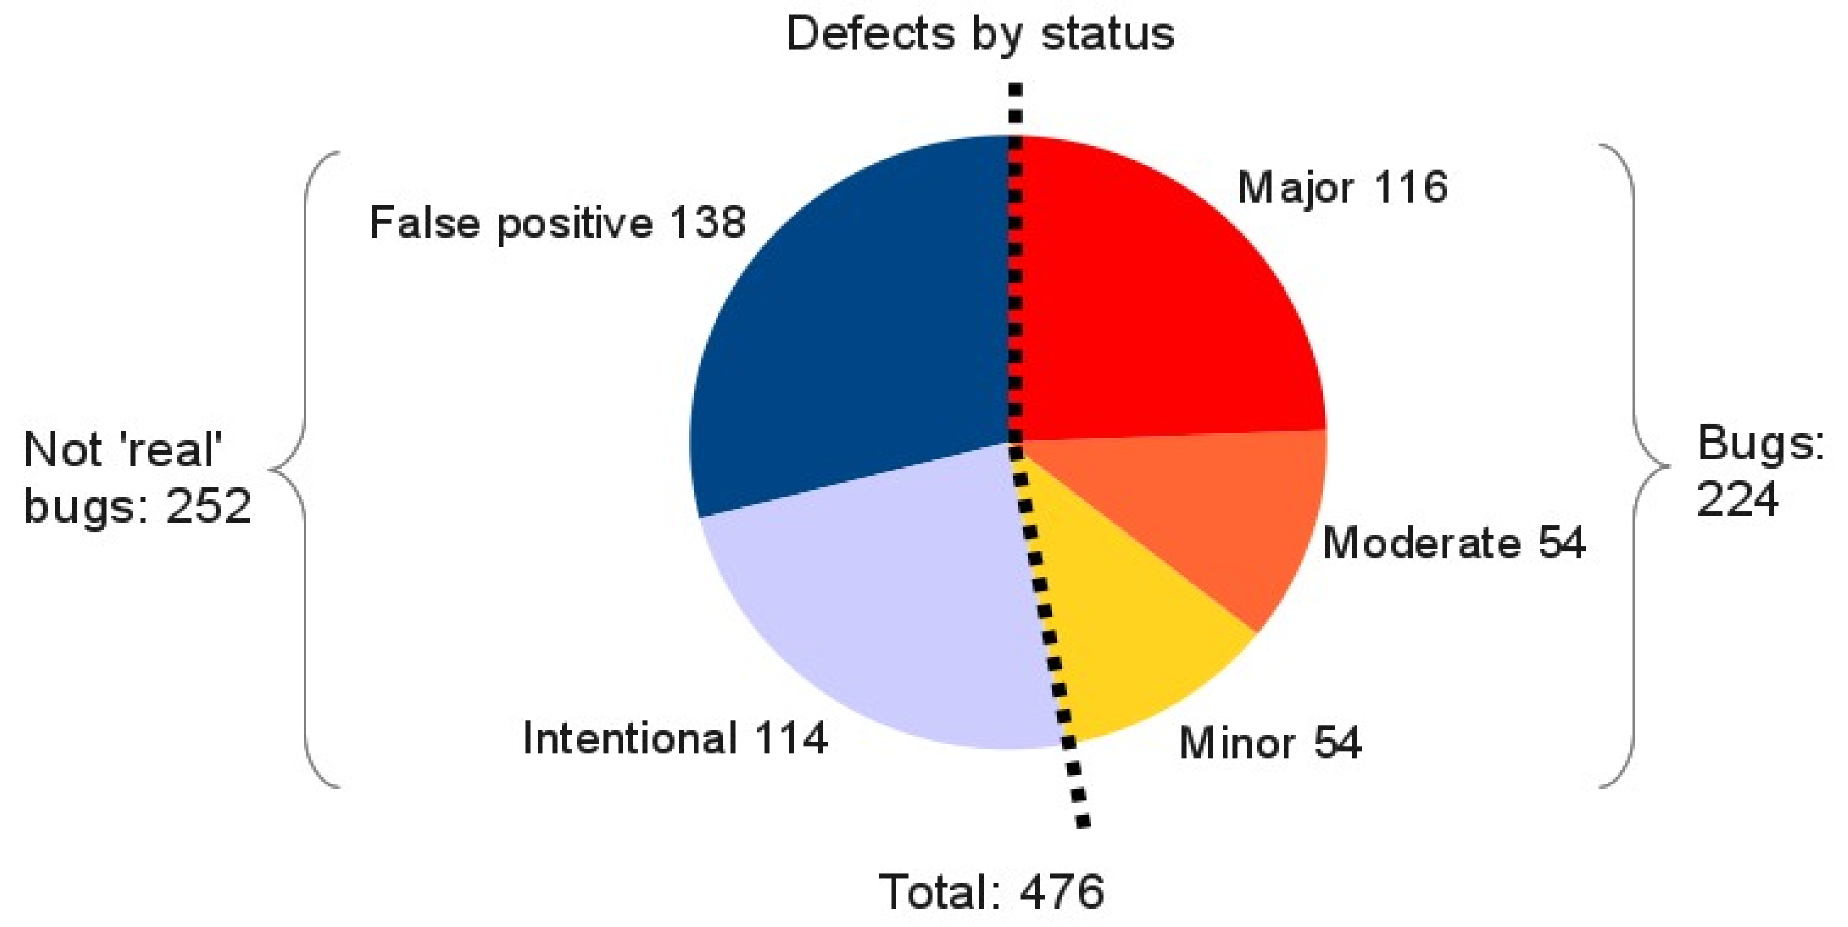

Defects found by Coverity

- Statistics on master

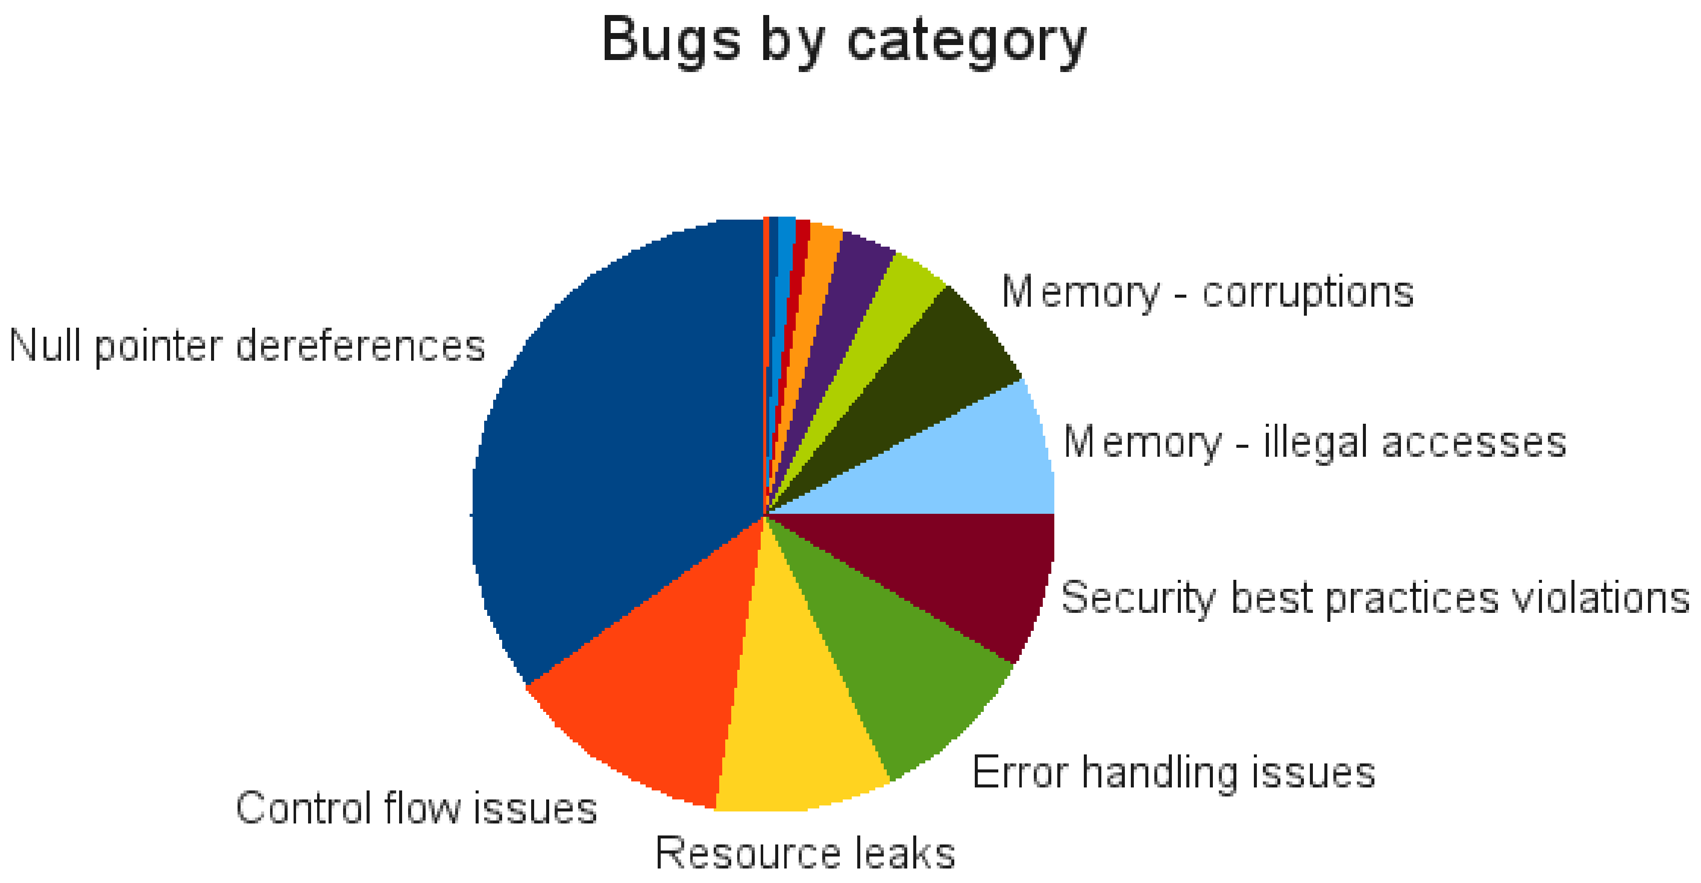

Defects found by Coverity

More on false positives and intentionals

- How can they be avoided?
  - 'fall through' in switch cases: please comment
  - function pointers, like the ones set in cfs_hash_create()
    - Too complex path to follow for Coverity
    - Redesign code?
    - Specific Coverity comments for future analysis

Benefits for the whole Lustre Community

Architect of an Open World™

Benefits for the whole Lustre Community

- Jira tickets opened, patches proposed
  - 8 tickets opened so far
  - Identified with 'coverity' label
  - 3 already merged:
    - LU-1856
    - LU-1884
    - LU-1907
  - Thanks for Intel's responsiveness

- Ongoing effort
  - Our goal is to open Jiras for all defects found by Coverity
  - So far:
    - 51 out of 224 bugs covered
    - 20 out of 252 “false issues” covered

Benefits for the whole Lustre Community

- LU-1855: memory corruption, out-of-bounds access
  - in lustre/ptlrpc/sec.c, sptlrpc_secflags2str():

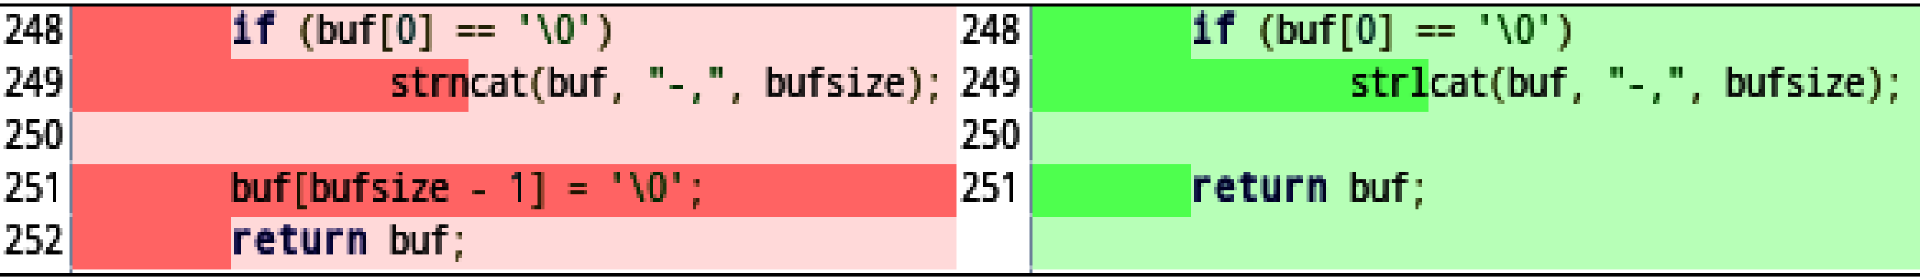

```
        if (buf[0] == '\0')
                strncat(buf, "-,", bufsize);

        buf[bufsize - 1] = '\0';
        return buf;
```

```
        if (buf[0] == '\0')
                strlcat(buf, "-,", bufsize);

        return buf;
```

Benefits for the whole Lustre Community

- LU-1857: memory corruption, unbounded source buffer

- Flaw in liblustreapi: llapi_file_get_stripe()
  - No 'size' parameter
  - Replace with llapi_file_get_layout()

Benefits for the whole Lustre Community

- LU-1884: resource leak
  - in lustre/lmv/lmv_intent.c, lmv_revalidate_slaves()

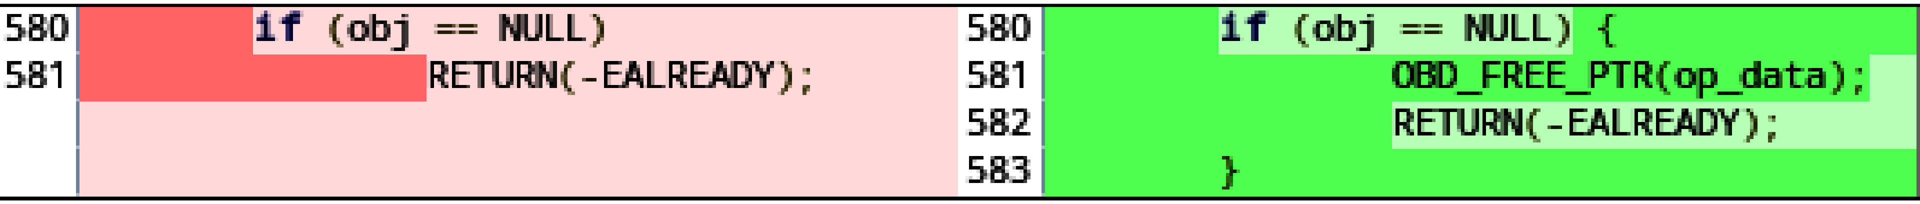

```
        if (obj == NULL)
                RETURN(-EALREADY);
```

```
        if (obj == NULL) {
                OBD_FREE_PTR(op_data);
                RETURN(-EALREADY);
        }
```

  - in lustre/utils/lfs.c, lfs_flushctx()

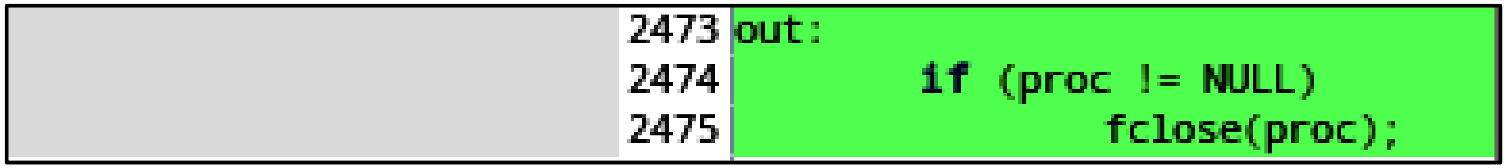

```
out:
        if (proc != NULL)
                fclose(proc);
```

Benefits for the whole Lustre Community

- LU-1889: false 'uninitialized variable' errors
  - in lustre/obdfilter/filter_log.c, filter_recov_log_mds_ost_cb()

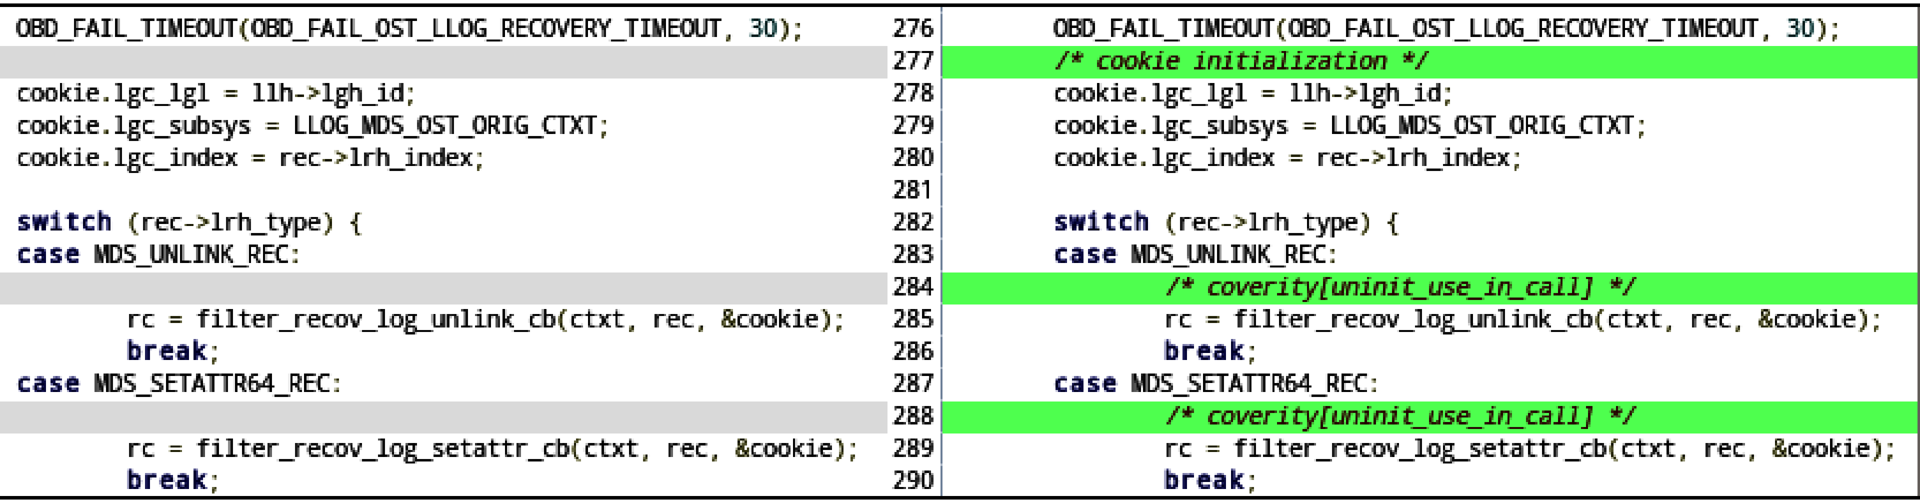

```
OBD_FAIL_TIMEOUT(OBD_FAIL_OST_LLOG_RECOVERY_TIMEOUT, 30);

cookie.lgc_lgl = llh->lgh_id;
cookie.lgc_subsys = LLOG_MDS_OST_ORIG_CTXT;
cookie.lgc_index = rec->lrh_index;

switch (rec->lrh_type) {
case MDS_UNLINK_REC:

        rc = filter_recov_log_unlink_cb(ctxt, rec, &cookie);
        break;
case MDS_SETATTR64_REC:

        rc = filter_recov_log_setattr_cb(ctxt, rec, &cookie);
        break;
```

```
OBD_FAIL_TIMEOUT(OBD_FAIL_OST_LLOG_RECOVERY_TIMEOUT, 30);
/* cookie initialization */
cookie.lgc_lgl = llh->lgh_id;
cookie.lgc_subsys = LLOG_MDS_OST_ORIG_CTXT;
cookie.lgc_index = rec->lrh_index;

switch (rec->lrh_type) {
case MDS_UNLINK_REC:
        /* coverity[uninit_use_in_call] */
        rc = filter_recov_log_unlink_cb(ctxt, rec, &cookie);
        break;
case MDS_SETATTR64_REC:
        /* coverity[uninit_use_in_call] */
        rc = filter_recov_log_setattr_cb(ctxt, rec, &cookie);
        break;
```

  - in lustre/utils/liblustreapi.c, llapi_ping()

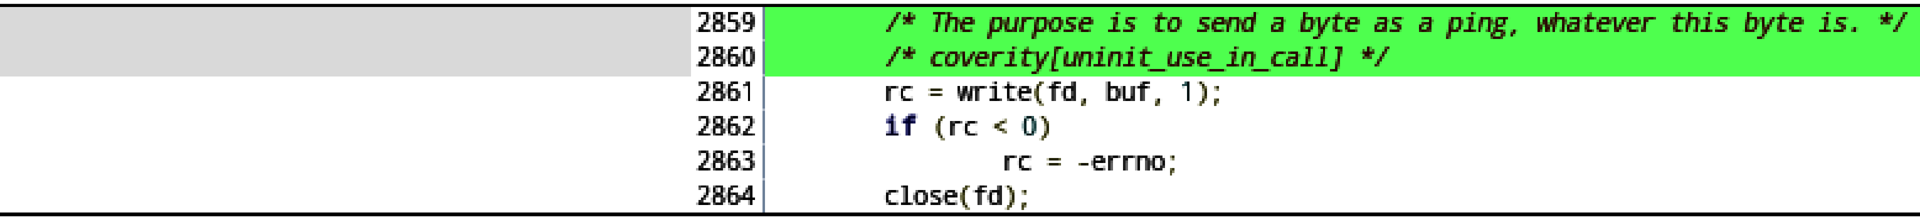

```
        /* The purpose is to send a byte as a ping, whatever this byte is. */
        /* coverity[uninit_use_in_call] */
        rc = write(fd, buf, 1);
        if (rc < 0)
                rc = -errno;
        close(fd);
```

Architect of an Open World™