

Hardening Your Code

Marshall Clow
Qualcomm Technologies, Inc.
mclow.lists@gmail.com

Twitter: [@mclow](https://twitter.com/mclow)

Blog: cplusplusmusings.wordpress.com (occasional)

CppCon

September 2014

It's a dangerous world

- Active attackers
- Environments that you did not envision
- Good methodology is not enough
- Mistakes can happen

What can you do to
increase your
confidence in your
code?

- Remember your history
- Test, test, test
- Look at compiler warnings
- Static analysis
- Dynamic analysis
- Fuzzing

Version Control

- Remember how you got where you are
- You can go back in time
- Remembers releases

A War Story

- <http://esr.ibiblio.org/?p=6205>
- Summary
 - Code stopped working (tests failed)
 - Reverted last change; still not working
 - Used ‘git bisect’ to find breaking change

Automated Tests

- Have a test suite
- Run it often
- Add to it whenever you can

But .. I don't have
automated tests
(and that sounds like a lot of work)

- Write some tests (even if it's just one)
 - Run it/them! (often)
- A small test suite is better than no test suite
- Add to it whenever you can.

What kind of things should I test?

- Normal operations
- Edge cases
- Error conditions

Someone reports a bug. What do you do?

- Write a test that reproduces the bug.
- Add this to your automated tests
- Run the test; verify that it fails
- Implement a fix
- Verify that the test passes (and all the other tests, too!)

Why are tests important?

- They give you confidence in your code
- They give you the ability to change your code w/o fear.
- “What if I break something?”
- Then the tests should catch it
- Enables refactoring

Compiler Warnings

Who cares about compiler warnings?

- You should.
- The compiler is telling you “there’s something here in your code that isn’t what I expect”
- If you have lots of warnings, it’s really hard to tell when you get a new one.

Toy Example #1

```
unsigned test (unsigned foo) {  
    if (foo >= 0)  
        return foo;  
    return 123;  
}
```

Test with different compilers

- Different compilers have different sets of warnings
- Staying warning-free on multiple compilers can be hard.

When I started working on [Product], there were tens of thousands of warnings in the source code in whatever version of gcc was in Xcode 2.1.

We are now (and have been for almost five years) warning free. With only a dozen or so uses of manually shutting off a warning around a block of code using the pragma for clang, and with maybe a hundred similar pragmas on Windows (because the Windows system header files produce warnings), and we're now running with the default warnings that are enabled in Xcode 5.1, which is significantly stricter than gcc ever was.

How did we do that? By turning on warnings (and treat warnings as errors) and going through and clearing up every single one of them. Took a couple years of me and one other engineer working on it in our "spare time" to get to zero, but staying at zero is significantly easier than getting there was.

Static Analysis

- Multiple commercial products
- Some open-source checkers as well

Toy Example #2

```
char *get_string (int foo) {
    if (foo == 0)
        return NULL;
    return "123";
}

int len (const char *s) {
    return strlen(s);
}

int bar = len(get_string(x));
```

More about static analyzers

- They find “interesting” bugs.
- They are heavyweight tools
 - Expensive to run
 - Expensive to purchase
- They tend to have high false positive rates

Dynamic Analysis

- Build an instrumented version of your program
 - You don't ship this version
- Test it
- Report when things go wrong

Examples

- Assertions
- “Debug mode”
- “Sanitizers”

Sanitizers

- Usually implemented by compiler vendor
 - Address
 - Undefined Behavior
 - Memory
 - Thread

Sanitizers (2)

- Very few (goal: 0) false positives
- Report misbehavior when it happens
- Can report stack crawl, etc.

Using ASAN

- <http://blog.llvm.org/2013/03/testing-libc-with-address-sanitizer.html>
- Ran ~4350 tests; found three bugs
 - One bug in the library being tested
 - Two bugs in the test code
- <https://www.mozilla.org/security/announce/2014/mfsa2014-49.html>

Undefined Behavior

- What is Undefined Behavior?
 - Integer overflow
 - indirection through NULL
 - indexing off the end of an array
 - Other things.

UBSAN

- Catches undefined behavior
- Examples
 - “load of value 123, which is not a valid value for type ‘bool’”
 - “runtime error: index 40 out of bounds for type 'char_type [10]’”
 - <http://blog.llvm.org/2013/04/testing-libc-with-fsanitizeundefined.html>

Fuzzing

- Take a set of valid inputs, and “modify” them.
- Feed the modified inputs into the program, and look for misbehavior
- <http://blog.chromium.org/2012/04/fuzzing-for-security.html>

Fuzzing and Dynamic Analysis

- These two techniques work *very well* together.

Wrapping up

- All of these things can be started small
 - Except maybe the static analysis
 - Improve them over time
 - They work well together (tests and source control, Sanitizers and fuzzing, etc)

Questions?

```
template <typename T>
unsigned char hex_char_to_int ( T val ) {
    char c = static_cast<char> ( val );
    if (c >= '0' && c <= '9') return c - '0';
    else if (c >= 'A' && c <= 'F') return c - 'A' +10;
    else if (c >= 'a' && c <= 'f') return c - 'a' +10;
    else throw non_hex_input();
}
```

```
template <typename T>
unsigned char hex_char_to_int ( T val ) {
    char c = static_cast<char> ( val );
    if (c >= '0' && c <= '9') return c - '0';
    else if (c >= 'A' && c <= 'F') return c - 'A' +10;
    else if (c >= 'a' && c <= 'f') return c - 'a' +10;
    else throw non_hex_input();
    return 0; // added to placate compiler
}
```

```
template <typename T>
unsigned char hex_char_to_int ( T val ) {
    char c = static_cast<char> ( val );
    unsigned retval = 0;
    if (c >= '0' && c <= '9') retval = c - '0';
    else if (c >= 'A' && c <= 'F') retval = c - 'A' +10;
    else if (c >= 'a' && c <= 'f') retval = c - 'a' +10;
    else throw non_hex_input();
    return retval;
}
```