
Secure Programming I



“If our software is buggy, what does
that say about its security?”

—*Robert H. Morris*

It's All About the Software

- Most penetrations are due to buggy software
- Crypto won't help there: bad software defeats good crypto
- Design matters, too

Security Isn't About Cryptography

- Cryptography is necessary (and we'll talk about it)
- But it's possible for cryptographic software to have bugs, too
- In one recent study, 80% of mobile apps had problems with their cryptography

Three Separate Aspects

- Enforcing security
- Avoiding bugs
- Proper components and proper composition

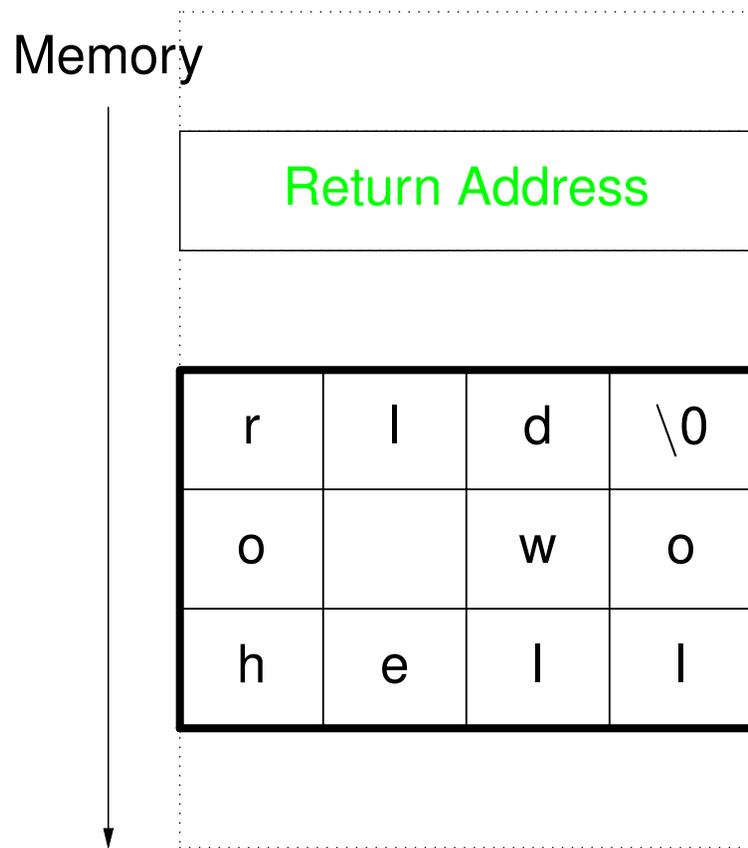
Avoiding Bugs

- Many simple bugs can be exploited to cause security problems
- (The C language is a large part of the problem)
- Much of the trouble is compounded by poor specifications

Buffer Overflows

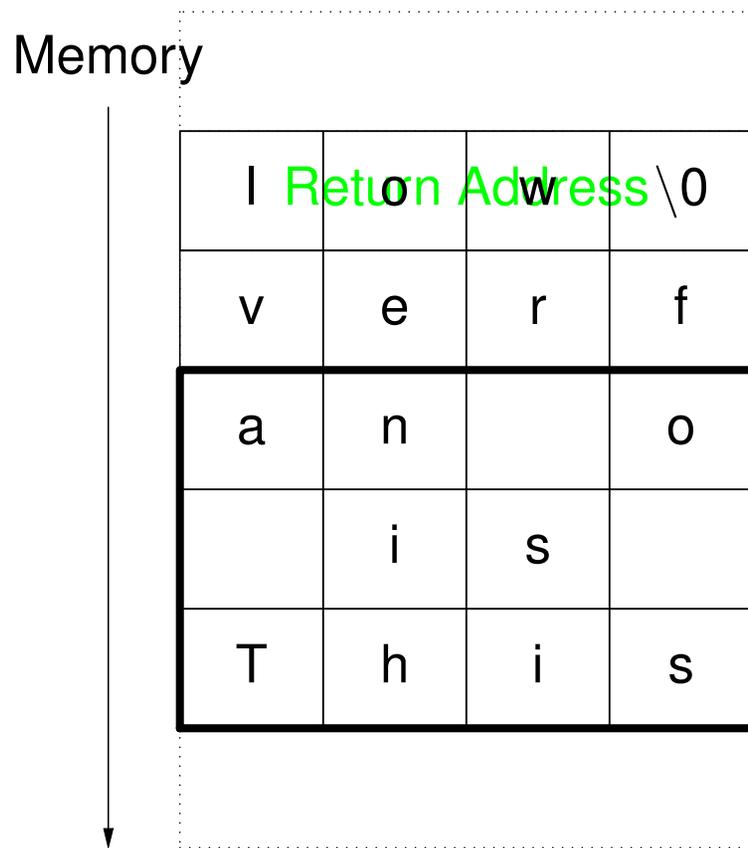
- Once responsible for about half of all security vulnerabilities; still a serious problem
- Fundamental problems:
 - Character strings in C are actually arrays of **chars**
 - There is no array bounds checking done in C
- Attacker's goal: overflow the array in a controlled fashion

Stack Frame



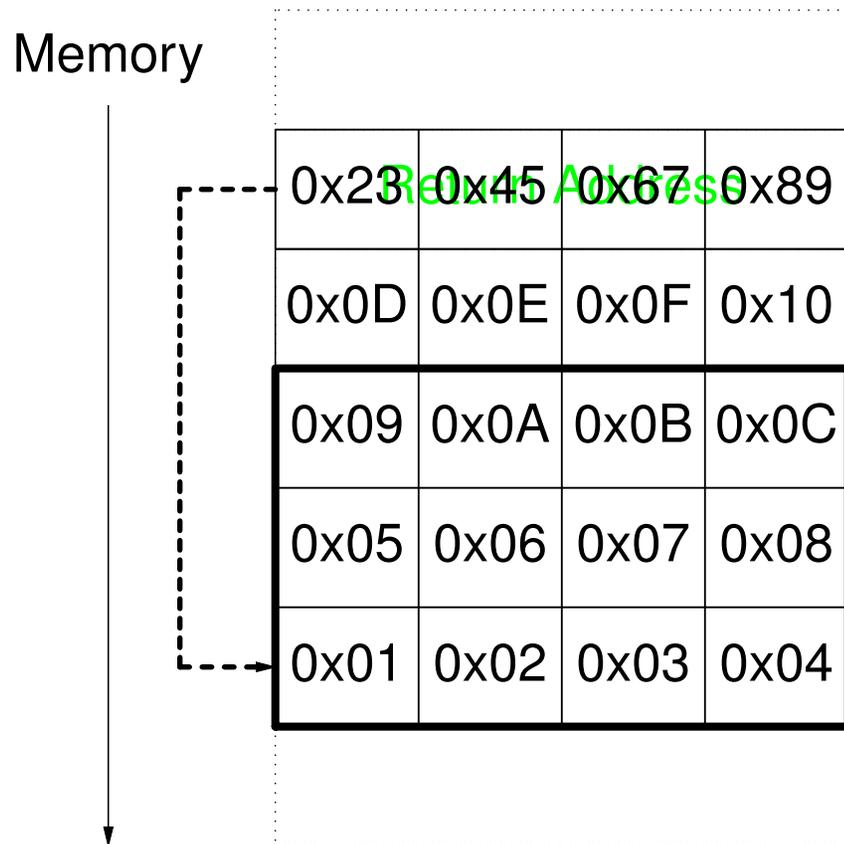
When a function is called, the return address is stored on the stack. Lower in memory, all local variables are stored.

Buffer Overflow



If the array bounds are exceeded, the return address can be overwritten.

Buffer Overflow Attack



Put code in the early part of the buffer, then change the return address to point to it. When the function exits, the injected code is executed.

How Can Such Things Happen?

- C has lots of built-in functions that don't check array bounds
- Programmers frequently don't check, either
- The attacker supplies too-long input

Sample Problematic Code

```
char line[512];
```

```
...
```

```
gets(line);
```

That's from the 4.3BSD `fingerd` command, exploited by the first Internet Worm in 1988...

Bad versus Good

<code>gets()</code>	<code>fgets()</code>
<code>strcpy()</code>	<code>strncpy()</code>
<code>strcat()</code>	<code>strncat()</code>
<code>sprintf()</code>	<code>snprintf()</code>

Java vs. C

- Java checks array bounds
- C# checks array bounds
- Go checks array bounds
- More or less everything *but* C and C++ check...

Indirect Buffer Overflows

```
void f(char *s)
{
    sprintf(s, "....");
}
```

```
void g()
{
    char buf[128];

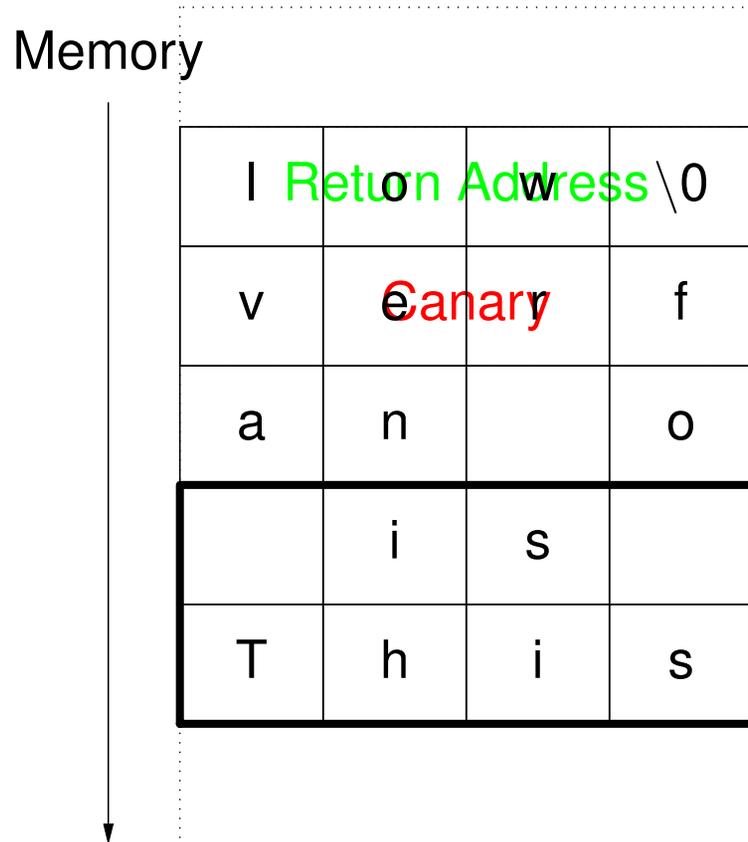
    f(buf);
}
```

Function `f` doesn't even know the size of the array!

Canaries

- Compiler trick—available for gcc and Microsoft compilers
- Generate a random “canary” value at the start of each program execution
- Insert that value between the return address and the rest of the stack frame
- Check if it’s intact before returning
- Any stack-smash attack will have to overwrite the canary to get to the return address
- Remember: the canary’s value is different for each run of the program

Overwriting a Canary



If the random canary is overwritten, the program will abort

ASLR: Address Space Layout Randomization

- Put stack at different random location each time program is executed
- Put heap at different random location as well
- Defeats attempts to address known locations
- But—makes debugging harder

Non-Executable Data Areas

- Modern computer architectures have permission bits for memory pages: can only execute code if the “execute” bit is set
- Defense: on pages with the “write” bit set, don’t set “execute”
- The stack is writable, so code injected by the attacker won’t be executable
- Called “DEP” (Data Execution Prevention) or “ $W \oplus X$ ”

Checking Code

- Look for suspect calls
- Use static checkers
- Use language feature like Perl's "taint" mode

“Taint” Mode

- Optional feature in Perl
- Variables whose value comes from an untrusted source are marked as “tainted”
- Variables whose value comes (in part) from a tainted variable are also tainted
- Tainted values cannot be passed to sensitive routines
- Taints are removed via regular expression match; it is presumed that the expression will sanitize the data, and remove anything dangerous

Stack versus Heap or BSS Storage

- Easiest to exploit if the buffer is on the stack
- Exploits for heap- or BSS-resident buffers are also possible, though they're harder
- Heap and BSS attacks not preventable with canaries (but there are analogous techniques to protect `malloc()`-allocated storage)
- Some operating systems can make such memory pages non-executable, which is a big help—but that breaks some applications

Issues for the Attacker

- Finding vulnerable programs
- NUL bytes
- Uncertainty about addresses

Finding Vulnerable Programs

- Use `nm` and `grep` to spot use of dangerous routines
- Probe via very-long inputs
- Look at source or disassembled/decompiled code

NUL Bytes

- C strings can't have embedded 0 bytes
- Some instructions do have 0 bytes, perhaps as part of an operand
- Solution: use different instruction sequence

Address Uncertainty

- Pad the evil instructions with NOPs
- This is called a *landing zone* or a *NOP sled*
- Set the return address to anywhere in the landing zone

Buffer Overflow: Summary

- You *must* check buffer lengths
- Where you can, use the safer library functions
- Write your own safe string library (there's no commonly-available standard)
- Use C++ and **class String**
- Use Java
- Use *anything* but raw C!

History of Buffer Overflows

- Long-recognized as a security issue
- First very visible exploit: Robert T. Morris' Internet Worm, November 1988.
- Popularized by Aleph One in November 1996; serious threat since then
 - 👉 The attack is theoretically difficult, but there are canned exploit kits available

Hoare's Turing Award Lecture, 1980

The first principle was security: . . . A consequence of this principle is that every occurrence of every subscript of every subscripted variable was on every occasion checked at run time against both the upper and the lower declared bounds of the array. . . . I note with fear and horror that even in 1980, language designers and users have not learned this lesson. In any respectable branch of engineering, failure to observe such elementary precautions would have long been against the law.

Can We Afford Array-Bounds Checking?

- Of course—spend the Moore's Law benefit on something besides better video games
- Compiler optimizations often make the expense a lot less than you'd think
- It's hard to do in C, though, because of array vs. pointer semantics
- Things like `*p++ = *q++` are hard to check efficiently
- A bounds-checking C compiler has been written, but it's largely unused

The Role of Specifications

- Contrast this:
 - “File names may be up to 1024 bytes long”with
 - “File names may be up to 1024 bytes long; longer file names *must* be rejected”
- The second form alerts the programmer to the real requirement
- Just as important, the second form alerts the *tester* to the requirement
- Testing is done against requirements!

Format String Errors

- Suppose `str` is input to the program

- Wrong:

```
printf(str);
```

- Right:

```
printf("%s", str);
```

- Format strings can be dangerous...
- Note: other functions (i.e., `syslog`) also take format strings

What's the Problem?

- Minor problem: metacharacters can confuse log files
- Here's an embedded newline in a username

👉 12:34:56 Permission denied: user
12:34:xx Watch this crash!

- Bigger problem: %n

The %n Problem

- Rather complex; I won't try to explain the details here
- Fundamental issue: %n writes to a variable the number of bytes printed thus far

- The statement

```
printf("Hello\n%n", &cnt)
```

stores a 6 in integer variable **cnt**

- This can be used to overwrite memory locations
- Use tricks involving other references to (non-existent!) other arguments to let you write to someplace “useful”

Return-Oriented Programming

- A more complex way to exploit buffer overflows
- Find a sequence of “gadgets”—sequences of instructions that do something useful—that end in a “return” instruction
- Push a sequence of gadget addresses onto the stack
- Each gadget will be executed in sequence

The Underlying Issues

- Problem 1: C has strange semantics
- The *only* defense is to know the language thoroughly
- You also have to know possible exploits
- There are integer overflow attacks, too
- Problem 2: programs don't always validate their inputs

Input Validation

- Trust *nothing* supplied by the user
- Must define inputs before they can be checked
- “A program whose behavior has not been specified cannot be buggy, only surprising.”
- Example: is a newline a valid character in a username?

Defenses

- Rigorously check all inputs against the specification
- Before that, of course, you need a spec
- (Specs can be buggy, too)
- Alternatively, use an earlier filter or check against a known-good list

Filtering

- Example: `fgets ()` stops at a newline; you can't find any embedded
- 👉 But watch for unterminated buffer—what if the input line is too long?
- Note that `argv` has no such guarantee
- Email: check recipient name against valid user list—no funny characters there

Being Careful Near the Shell

- If user input is being passed to the shell, be especially careful
- Watch for `popen()` and `system()`

- Dangerous characters include:

``~$^&()={ } [] | ; : ' " ? < > \`

- That's most of the special characters!
- You're always much better off with a “good” list than a “bad” list
- Example: on some Unix systems, `^` was treated the same as `|`.
Why? Because on some models of Teletype—the *ancient* hard-copy terminal!—`^` printed as `↑`, which looked similar to `|`

Knowing the Semantics

- Sometimes check that there are no / characters in a program name
- Why? To ensure that the reference is to a given directory
- Do you need to check \ as well?
- 👉 Will the program ever run on Windows? Note that URLs on Windows use /, but the file system uses \

Summary

- Trust nothing
- Specify acceptable inputs
- Check everything
- Understand the semantics of anything you invoke
- Try to use a better language than C