

CSE 484 / CSE M 584: **Software Security, Buffer Overflows**

Fall 2024

Franziska (Franzi) Roesner
franzi@cs

UW Instruction Team: David Kohlbrenner, Yoshi Kohno, Franziska Roesner. Thanks to Dan Boneh, Dieter Gollmann, Dan Halperin, John Manferdelli, John Mitchell, Vitaly Shmatikov, Bennet Yee, and many others for sample slides and materials ...

Announcements

- Things Due (**submit to Gradescope**):
 - Ethics Form: Due today!
 - Homework #1: Due Friday
 - Research Readings (CSE M 584): Due Thursday (and every Thursday thereafter)
- Lab 1
 - Out later this week – **don't miss section!**
- Office hours
 - See course website for planned schedule (**double check before you go!**)

Bugs, Vulnerabilities, and Exploits

- Bug
 - Not working quite right
- Vulnerability
 - A malfunction that can be used for an adversary's goals
- Exploit
 - The mechanical set of operations to make use of a vulnerability

Adversarial Failures

- Software bugs are bad
 - Consequences can be serious
- Even worse when an **intelligent adversary** wishes to **exploit** them!
 - Intelligent adversaries: Force bugs into “**worst possible**” conditions/states
 - Intelligent adversaries: Pick their targets

A Few Words on Ethics

- Ethics questions can be subtle complex!
- *Ex: If you find a vulnerability in a deployed product, what should you?*
- HW1 will ask you to learn + consider different ethical frameworks:
 - **Consequentialist:** considers **consequences** of different decisions
 - **Deontological:** considers **duties and rights** of different stakeholders
- You might come to the same conclusion with different frameworks, or you may come to different conclusions. **We are not looking for a single “right answer”.**

Memory Corruption Bugs

- **Buffer overflows bugs:** Big class of bugs
 - Normal conditions: Can sometimes cause systems to fail
 - Adversarial conditions: Attacker able to violate security of your system (control, obtain private information, ...)
- Stack, Heap both possibilities

BUFFER OVERFLOWS

A Bit of History: Morris Worm

- Worm was released in 1988 by Robert Morris
 - Graduate student at Cornell, son of NSA chief scientist
 - Convicted under Computer Fraud and Abuse Act,
 - 3 years probation and 400 hours of community service
 - Now an EECS professor at MIT
- Worm was intended to propagate slowly and harmlessly measure the size of the Internet
- Due to a coding error, it created new copies as fast as it could and overloaded infected machines
- \$10-100M worth of damage (in 1988)

Morris Worm and Buffer Overflow

- One of the worm's propagation techniques was a **buffer overflow attack** against a vulnerable version of `fingerd` on VAX systems
 - By sending special string to finger daemon, worm caused it to execute code creating a new worm copy

Buffer overflows remain a common source of vulnerabilities and exploits today!
(Especially in embedded systems.)

Attacks on Memory Buffers

- **Buffer** is a pre-defined data storage area inside computer memory (stack or heap)
- Typical situation:
 - A function takes some input that it writes into a **pre-allocated buffer**.
 - The developer **forgets to check** that the size of the input isn't larger than the size of the buffer.
 - **Uh oh.**
 - “Normal” bad input: crash
 - “Adversarial” bad input : take control of execution

Stack Buffers



buf

uh oh!

- Suppose Web server contains this function

```
void func(char *str) {  
    char buf[126];  
    ...  
    strcpy(buf, str);  
    ...  
}
```

- No bounds checking on `strcpy()`
- If `str` is longer than 126 bytes
 - Program may crash
 - Attacker may change program behavior

Example: Changing Flags

buf

! (:-) !)


- Suppose Web server contains this function

```
void func(char *str) {  
    byte authenticated = 0;  
    char buf[126];  
    ...  
    strcpy(buf, str);  
    ...  
}
```

- **Authenticated** variable non-zero when user has extra privileges
- Morris worm also overflowed a buffer to overwrite an authenticated flag in fingerd

Refresher: Memory Layout

- **Text region:** Executable code of the program
- **Heap:** Dynamically allocated data
- **Stack:** Local variables, function return addresses; grows and shrinks as functions are called and return



Refresher: Stack Buffers

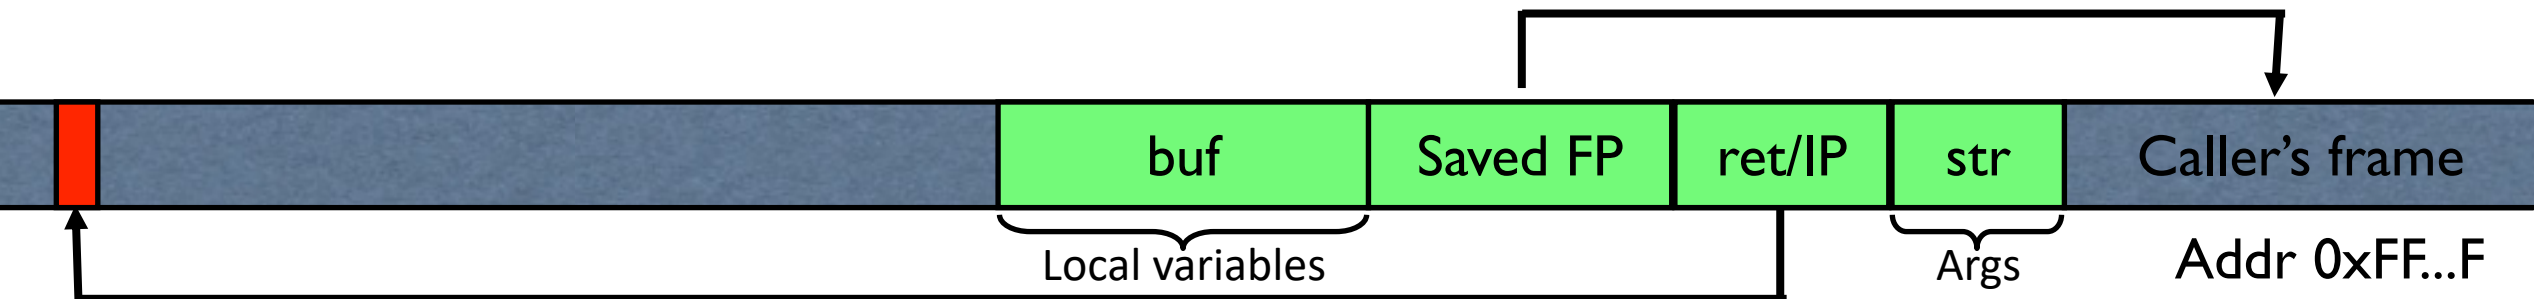
- Suppose Web server contains this function:

```
void func(char *str) {  
    char buf[126];  
    strcpy(buf, str);  
}
```

Allocate local buffer
(126 bytes reserved on stack)

Copy argument into local buffer

- When this function is invoked, a new **frame** (activation record) is pushed onto the stack.



Execute code at this address after func() finishes


What if Buffer is Overstuffed?

- Memory pointed to by str is copied onto stack...

```
void func(char *str) {  
    char buf[126];  
    strcpy(buf, str);  
}
```


strcpy does NOT check whether the string at *str contains fewer than 126 characters

- If a string longer than 126 bytes is copied into buffer, it will overwrite adjacent stack locations.



Executing Attack Code

- Suppose buffer contains attacker-created string
 - For example, str points to a string received from the network as the URL



- When function exits, code in the buffer will be executed, giving attacker a shell (**“shellcode”**)
 - **Root shell** if the victim program is setuid root

Buffer Overflows Can Be Tricky...

- Overflow portion of the buffer must contain **correct address of attack code** in the RET position
 - The value in the RET position must point to the beginning of attack assembly code in the buffer
 - Otherwise application will (probably) crash with segfault
 - **Attacker must correctly guess in which stack position his/her buffer will be when the function is called**

Problem: No Bounds Checking

- strcpy does not check input size
 - strcpy(buf, str) simply copies memory contents into buf starting from *str until “\0” is encountered, ignoring the size of area allocated to buf
- Many C library functions are unsafe
 - strcpy(char *dest, const char *src)
 - strcat(char *dest, const char *src)
 - gets(char *s)
 - scanf(const char *format, ...)
 - printf(const char *format, ...)

Does Bounds Checking Help?

- `strncpy(char *dest, const char *src, size_t n)`
 - For `strncpy` (unlike `strcpy`), no more than `n` characters will be copied from `*src` to `*dest`
- Potential overflow in `htpasswd.c` (Apache 1.3):

```
strcpy(record, user);  
strcat(record, ":");  
strcat(record, cpw);
```

Copies username (“user”) into buffer (“record”), then appends “:” and hashed password (“cpw”)

- Published fix:

```
strncpy(record, user, MAX_STRING_LEN-1);  
strcat(record, ":");  
strncat(record, cpw, MAX_STRING_LEN-1);
```

In-Class Activity

In-Class Activity Gradescope – 9/30 Activity

<https://www.gradescope.com/courses/881751/assignments/5072027>

(This is the first one that will be graded.)

Reminder that you can submit up to half of them a week late.)

Does Bounds Checking Help?

- `strncpy(char *dest, const char *src, size_t n)`
 - For `strncpy` (unlike `strcpy`), no more than `n` characters will be copied from `*src` to `*dest`
 - Programmer has to supply the right value of `n`
- Potential overflow in `htpasswd.c` (Apache 1.3):

```
strcpy(record, user);  
strcat(record, ":");  
strcat(record, cpw);
```

Copies username (“user”) into buffer (“record”), then appends “:” and hashed password (“cpw”)

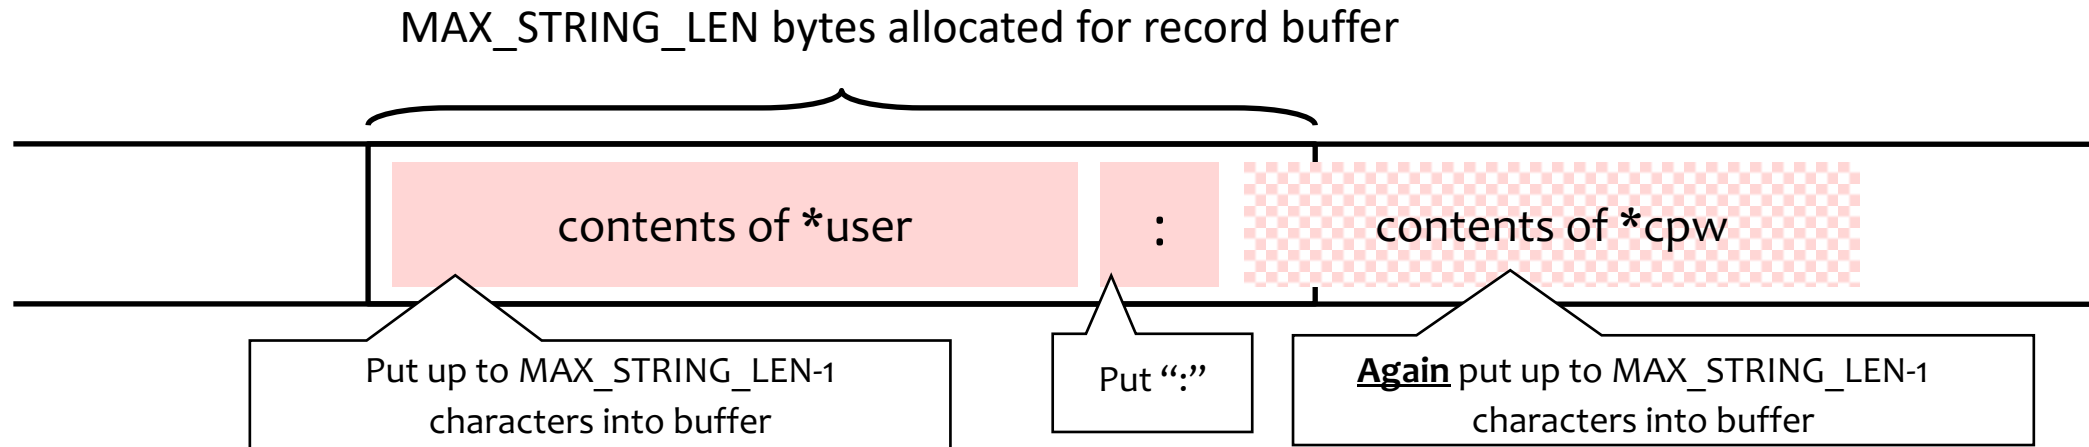
- Published fix:

```
strncpy(record, user, MAX_STRING_LEN-1);  
strcat(record, ":");  
strncat(record, cpw, MAX_STRING_LEN-1);
```

Misuse of strncpy in httpasswd “Fix”

- Published “fix” for Apache httpasswd overflow:

```
strncpy(record,user,MAX_STRING_LEN-1);  
strcat(record,":")  
strncat(record,cpw,MAX_STRING_LEN-1);
```



What About This?

- Home-brewed range-checking string copy

```
void mycopy(char *input) {
    char buffer[512]; int i;

    for (i=0; i<=512; i++)
        buffer[i] = input[i];
}

void main(int argc, char *argv[]) {
    if (argc==2)
        mycopy(argv[1]);
}
```

In-Class Activity

In-Class Activity Gradescope → 9/30

Off-by-One Overflow

- Home-brewed range-checking string copy

```
void mycopy(char *input) {
    char buffer[512]; int i;

    for (i=0; i<=512; i++)
        buffer[i] = input[i];
}

void main(int argc, char *argv[]) {
    if (argc==2)
        mycopy(argv[1]);
}
```

This will copy **513** characters into buffer. Oops!

- 1-byte overflow: can't change RET, but can change pointer to previous stack frame...