

CSCI2467: Systems Programming Concepts

Slideset 12: Stack attacks and defenses

Source: CS:APP section 3.10, Bryant & O'Hallaron

Course Instructors:

Matthew Toups
Caitlin Boyce

Course Assistants:

Saroj Duwal
David McDonald

Spring 2020



THE UNIVERSITY of
NEW ORLEANS

Our last lab assignment

- Lab4 (attacklab) is out
 - Due: Monday, April 20, 11:59pm
 - like before, you will defeat “phases” of an unknown, unique binary program
 - unlike before, no “explosion” penalty

Attacklab handin!

Yes you must hand in to Autolab

- You must hand in a commented, plain-text version of your solutions (eg. `phase1.txt`)
- Use autolab submit button
- After you submit, your score will show up on scoreboard
- Please submit your new solution every time you solve a phase
- Comments must explain how and why your solution works!
- Convince us you know! Otherwise we will not award points


x86-64 Procedure Summary

■ Important Points

- Stack is the right data structure for procedure call / return
 - If P calls Q, then Q returns before P

■ Recursion (& mutual recursion) handled by normal calling conventions

- Can safely store values in local stack frame and in callee-saved registers
 - Put function arguments at top of stack
 - Result return in `%rax`
- ## ■ Pointers are addresses of values
- On stack or global




x86-64 Example Addresses

address range $\sim 2^{47}$

local
p1
p3
p4
p2
big_array
huge_array
main()
useless()

0x00007ffe4d3be87c
0x00007f7262a1e010
0x00007f7162a1d010
0x000000008359d120
0x000000008359d010
0x0000000080601060
0x0000000000601060
0x000000000040060c
0x0000000000400590



Buffer overflows: big security implications

- What we just saw is generally called a *buffer overflow*
- Why a big deal?
 - For decades, has been the #1 technical cause of security vulnerabilities
(#1 overall cause is social engineering)
- Most common form:
 - unchecked lengths on string inputs
 - particularly for bounded character arrays on the stack
(sometimes referred to as “stack smashing”)

Vulnerable buffer code

```
/* Echo Line */  
void echo()  
{  
    char buf[4]; /* Way too small! */  
    gets(buf);  
    puts(buf);  
}
```

```
void call_echo() {  
    echo();  
}
```

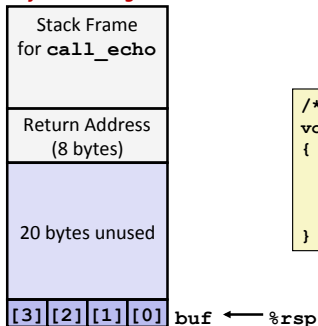
← btw, how big
is big enough?

```
unix> ./bufdemo-nsp  
Type a string:012345678901234567890123  
012345678901234567890123
```

```
unix> ./bufdemo-nsp  
Type a string:0123456789012345678901234  
Segmentation Fault
```


Buffer overflow stack

Before call to gets




```
/* Echo Line */  
void echo()  
{  
    char buf[4]; /* Way too small! */  
    gets(buf);  
    puts(buf);  
}
```

```
echo:  
    subq  $24, %rsp  
    movq  %rsp, %rdi  
    call  gets  
    . . .
```

Buffer overflow stack: example

Before call to gets



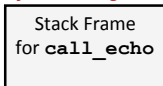
```
void echo()
{
    char buf[4];
    gets(buf);
    . . .
}

echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    . . .
```

call_echo:

```
. . .
4006f1: callq 4006cf <echo>
4006f6: add $0x8,%rsp
. . .
```

After call to gets



```
void echo()
{
    char buf[4];
    gets(buf);
    . . .
}

echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    . . .
```

Buffer overflow stack: example

After call to gets

Stack Frame for <code>call_echo</code>			
00	00	00	00
00	40	06	00
33	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

`buf` ← `%rsp`

register tm clones:

```
. . .  
400600: mov    %rsp,%rbp  
400603: mov    %rax,%rdx  
400606: shr   $0x3f,%rdx  
40060a: add   %rdx,%rax  
40060d: sar   %rax  
400610: jne   400614  
400612: pop   %rbp  
400613: retq
```

- “Returns” to unrelated code
- Lots of things happen, without modifying critical state
- Eventually executes `retq` back to main

Exploits based on buffer overflows

- ***Buffer overflow bugs can allow remote machines to execute arbitrary code on victim machines***
- **Distressingly common in real programs**
 - Programmers keep making the same mistakes ☹
 - Recent measures make these attacks much more difficult
- **Examples across the decades**
 - Original “Internet worm” (1988)
 - “IM wars” (1999)
 - Twilight hack on Wii (2000s)
 - ... and many, many more
- **You will learn some of the tricks in attacklab**
 - Hopefully to convince you to never leave such holes in your programs!!

Example: “Morris worm” (1988)

- First Internet-based malware
(spread via network using buffer overflow exploits)

- Used several vulnerabilities to spread
 - gets() called in some UNIX services (fingerd)
these services take input from users
... and with certain crafted input, would execute root shell
- Once the worm takes over a system, it scans the network for other computers to attack
 - Morris worm took over an estimated 6000 computers
(10% of the internet at the time!)
 - Described in June 1989 article in *Communications of the ACM*

Computer Intruder Is Put on Probation And Fined \$10,000

By JOHN MARKOFF, Special to The New York Times
Published: May 5, 1990

SYRACUSE, May 4— Saying the punishment of prison did not fit the crime, a Federal judge today placed a 24-year-old computer science student on three years' probation for intentionally disrupting a nationwide computer network. The student, Robert Tappan Morris, was also fined \$10,000 and ordered to perform 400 hours of community service.

The sentencing of Mr. Morris had been awaited with great interest by computer security experts and those who try to evade them.

The case, which began when Mr. Morris wrote a program that copied itself wildly in thousands of separate machines in November 1988, has become a symbol of the vulnerabilities of the computer networks that serve as the nation's highways in the age of instant information.

Legal experts said the Government's decision to prosecute Mr. Morris, after an eight-month debate in the Justice Department, sent a strong message that tampering with computers, even when not intentionally destructive, was not acceptable. When Mr. Morris was found guilty last January, he became the first person convicted by a jury under the Federal Computer Fraud and Abuse Act of 1986.

- Worm – a program that:
 - can run by itself
 - can propagate a full working version of itself to other computers
- Virus – code that:
 - adds itself to other programs
 - does not run independently
- Both are designed to spread on their own

1990s: the word gets out

- 1995: “How to write buffer overflows” published
- 1996: First *complete*, public explanation of buffer overflows:
“Smashing The Stack For Fun And Profit”

.oO Phrack 49 Oo.

Volume Seven, Issue Forty-Nine

File 14 of 16

BugTraq, r00t, and Underground.Org
bring you

XX
Smashing The Stack For Fun And Profit
XX

by Aleph One
aleph1@underground.org

`smash the stack` [C programming] n. On many C implementations it is possible to corrupt the execution stack by writing past the end of an array declared auto in a routine. Code that does this is said to smash the stack, and can cause return from the routine to jump to a random address. This can produce some of

Later worms


Buffer overflow a common tactic for malware

- 2001: Code Red worm
 - buffer overflow in Microsoft IIS
- 2003: SQL Slammer
 - buffer overflow MS SQL server
 - hit 75,000 victims within 10 minutes
 - 376 bytes

Worms are not the only (in)famous uses of buffer overflows...

■ July, 1999

- Microsoft launches MSN Messenger (instant messaging system).
- Messenger clients can access popular AOL Instant Messaging Service (AIM) servers



In Cyberspace, Rivals Skirmish Over Messaging

By SAUL HANSELL

America Online closed its on-line service yesterday to new software from two of its fiercest rivals, Microsoft and Yahoo, that had been designed to tap into one of America Online's most popular features: instant messages.

On Thursday, both the Microsoft Corporation and Yahoo introduced

■ August 1999

- Mysteriously, Messenger clients can no longer access AIM servers
- Microsoft and AOL begin the IM war:
 - AOL changes server to disallow Messenger clients
 - Microsoft makes changes to clients to defeat AOL changes
 - At least 13 such skirmishes
- What was really happening?
 - AOL had discovered a buffer overflow bug in their own AIM clients
 - They exploited it to detect and block Microsoft: the exploit code returned a 4-byte signature (the bytes at some location in the AIM client) to server
 - When Microsoft changed code to match signature, AOL changed signature location

IM wars

From: Phil Bucking <philbucking@yahoo.com>
Subject: AOL exploiting buffer overrun bug in their own software!
To: rms@pharlap.com

Mr. Smith,

I am writing you because I have discovered something that I think you might find interesting because you are an Internet security expert with experience in this area. I have also tried to contact AOL but received no response.

I am a developer who has been working on a revolutionary new instant messaging client that should be released later this year.

...

It appears that the AIM client has a buffer overrun bug. By itself this might not be the end of the world, as MS surely has had its share. But AOL is now *exploiting their own buffer overrun bug* to help in its efforts to block MS Instant Messenger.

....

Since you have significant credibility with the press I hope that you can use this information to help inform people that behind AOL's friendly exterior they are nefariously compromising peoples' security.

It appears that the AIM client has a buffer overrun bug. By itself this might not be the end of the world, as MS surely has had its share. But AOL is now *exploiting their own buffer overrun bug* to help in its efforts to block MS Instant Messenger.

Twilight hack (2008)²

- First method of running “homebrew” apps on Nintendo Wii
- Using a special saved-game file with a custom name for Epona (Link’s horse) containing exploit code (code injection)



²http://www.wiibrew.org/wiki/Twilight_Hack

iPhone jailbreaks (2007-present)

- Apple locks down iPods / iPhones / etc
 - Why? Restrict apps to iTunes store, lock devices to carriers, prevent malware
- Device owners try to circumvent “jail”
 - Why? Run other apps, evade censorship, customize OS, unlock device, etc



iPhone jailbreaks³

- How? Many techniques
 - Buffer overruns and integer overflows are common
- Apple's response? Issue OS update to prevent jailbreak (patch vulnerabilities)
 - So new vulnerabilities are found, new jailbreak released

Apple plugs critical iPhone jailbreak holes

Reuters Staff

4 MIN READ



Apple today patched the two vulnerabilities used to jailbreak Apple's newest iOS 4 operating system, bugs that security researchers warned could be used to hijack iPhones , iPod Touches or iPads .

The patches came just 10 days after a group published a site that automatically exploited and then jailbroke any iOS 4 device that used the mobile Safari browser to surf to jailbreakme.com

Also last week, other researchers confirmed that the first exploit of the pair leveraged a flaw in Safari's parsing of fonts in PDF documents to compromise the browser. A second vulnerability was exploited to break out of the isolating "sandbox" and gain full, or "root,"

● Class updates

① Memory layout

② Buffer Overflow

- Vulnerability
- History
- Protection
 - Bug-free code?
 - ASLR and NX
 - Stack canaries
- Arms race!
 - Return-oriented programming (ROP)

③ Attacklab (lab4)

Defense 1: avoid overflow vulnerabilities when writing programs

```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    fgets(buf, 4, stdin);
    puts(buf);
}
```

- For example, use library routines that limit string lengths
 - `fgets` instead of `gets`
 - `strncpy` instead of `strcpy`
 - Don't use `scanf` with `%s` conversion specification
 - Use `fgets` to read the string
 - Or use `%ns` where `n` is a suitable integer

Defense 2: System-level protections


Randomized stack offsets

■ Randomized stack offsets

- At start of program, allocate random amount of space on stack
- Shifts stack addresses for entire program
- Makes it difficult for hacker to predict beginning of inserted code
- E.g.: 5 executions of memory allocation code

local 0x7ffe4d3be87c 0x7fff75a4f9fc 0x7ffeadb7c80c 0x7ffeaea2fdac 0x7ffcd452017c

- Stack repositioned each time program executes




Defense 2: System-level protections

Non-executable stack

■ Nonexecutable code segments

- In traditional x86, can mark region of memory as either “read-only” or “writeable”
 - Can execute anything readable
- X86-64 added explicit “execute” permission
- Stack marked as non-executable



Any attempt to execute this code will fail

Canary in a coal mine



Source: U.S. State Department/Doug Thompson

Defense 3: Canary example

Disassembly shows additional instructions in every function

echo:

```
40072f:  sub    $0x18,%rsp
400733:  mov    %fs:0x28,%rax
40073c:  mov    %rax,0x8(%rsp)
400741:  xor    %eax,%eax
400743:  mov    %rsp,%rdi
400746:  callq 4006e0 <gets>
40074b:  mov    %rsp,%rdi
40074e:  callq 400570 <puts@plt>
400753:  mov    0x8(%rsp),%rax
400758:  xor    %fs:0x28,%rax
400761:  je     400768 <echo+0x39>
400763:  callq 400580 <__stack_chk_fail@plt>
400768:  add    $0x18,%rsp
40076c:  retq
```