

THE INTERNET of

~~THINGS~~

SHIT...

OR NET OF A MILLION SPIES

# Overflows & Format Strings

# Announcements...

- Homework 1 due in ~1 week
- Project 1 release Real Soon Now:
  - Practical exploitation of vulnerable SUID programs running in a VM
    - Have you done Project 0?

# Internet of Shit...


- A device produced by the lowest bidder...
  - That you then connect through the network
- This has a very wide **attack surface**
  - Methods where an attacker might access a vulnerability
- And its often incredibly **cost sensitive**
  - Very little support after purchase
    - So things don't get patched
  - No way for the user to tell what is "secure" or "not"
    - But they can tell what is cheaper!
    - And often it is **insanely** insecure:  
Default passwords on telnet of admin/admin...  
Trivial buffer overflows

# Net Of A Million Spies...

- Device only communicates through a central service
  - Greatly reduces the attack surface but...
- Most of the companies running the service are "Data Asset" companies
  - Make their money from advertising, not the product themselves
    - May actually subsidize the product considerably
  - Some you know about: Google, Amazon
  - Some you may not: Salesforce
- Only exception of note is Apple:
  - I may talk about HomeKit later...  
But you still have to trust that the HomeKit product doesn't report to a third party.





 **Traveler Information**

**Traveler 1 - Adults (age 18 to 64)**

To comply with the [TSA Secure Flight program](#), the traveler information listed here must exactly match the information on the government-issued photo ID that the traveler presents at the airport.

Title (optional):	First Name:	Middle Name:	Last Name:
<input type="text" value="Dr."/>	<input type="text" value="Alice"/>	<input type="text"/>	<input type="text" value="Smith"/>

Gender:	Date of Birth:	Travelers are required to enter a middle name/initial if one is listed on their government-issued photo ID.
<input type="text" value="Female"/>	<input type="text" value="01/24/93"/>	


Some younger travelers are not required to present an ID when traveling within the U.S. [Learn more](#)

**Known Traveler Number/Pass ID (optional):** [?](#)

**Redress Number (optional):** [?](#)


Seat Request:

No Preference  Aisle  Window







 **Traveler Information**

**Traveler 1 - Adults (age 18 to 64)**


To comply with the [TSA Secure Flight program](#), the traveler information listed here must exactly match the information on the government-issued photo ID that the traveler presents at the airport.


Title (optional): <input type="text" value="Dr."/>	First Name: <input type="text" value="Alice"/>	Middle Name: <input type="text"/>	Last Name: <input type="text" value="Smithoooooooooooo"/>
-------------------------------------------------------	---------------------------------------------------	--------------------------------------	--------------------------------------------------------------

Travelers are required to enter a middle name/initial if one is listed on their government-issued photo ID.

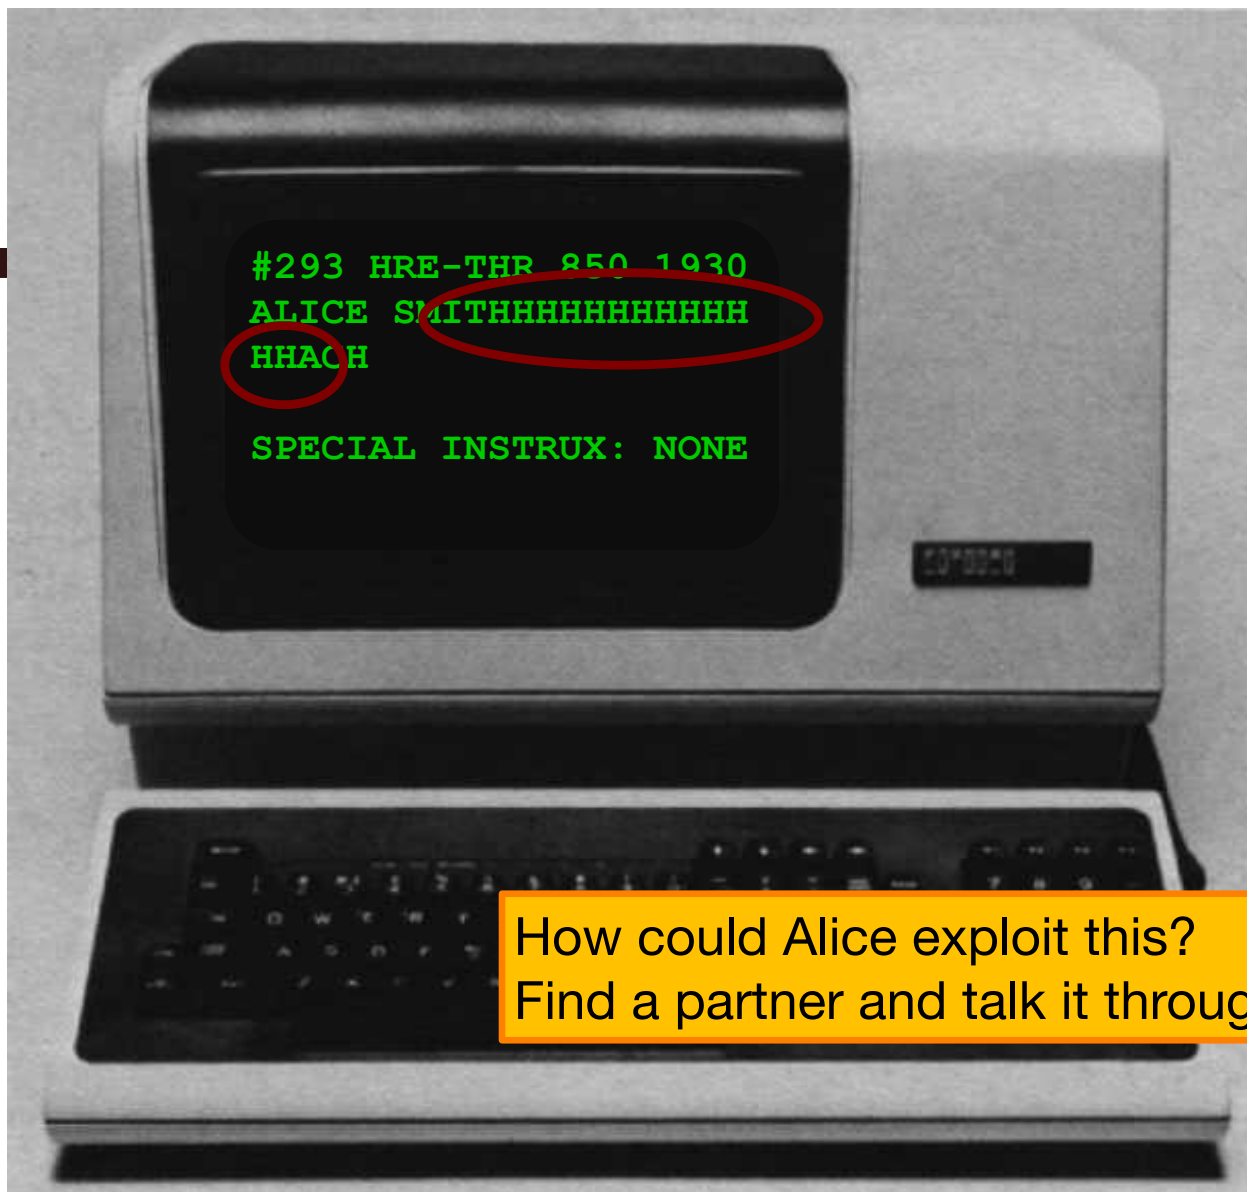
Gender: <input type="text" value="Female"/>	Date of Birth: <input type="text" value="01/24/93"/>
------------------------------------------------	---------------------------------------------------------

Some younger travelers are not required to present an ID when traveling within the U.S. [Learn more](#)


**+ Known Traveler Number/Pass ID (optional):** 

**+ Redress Number (optional):** 

Seat Request:  
 No Preference  Aisle  Window



How could Alice exploit this?  
Find a partner and talk it through.

 **Traveler Information**

**Traveler 1 - Adults (age 18 to 64)**

To comply with the [TSA Secure Flight program](#), the traveler information listed here must exactly match the information on the government-issued photo ID that the traveler presents at the airport.

Title (optional):	First Name:	Middle Name:	Last Name:
Dr. <input type="text"/>	Alice <input type="text"/>	<input type="text"/>	Smith <input type="text"/> First <input type="text"/>

Gender:  Date of Birth:

Travelers are required to enter a middle name/initial if one is listed on their government-issued photo ID.

Some younger travelers are not required to present an ID when traveling within the U.S. [Learn more](#)

+ **Known Traveler Number/Pass ID (optional):** [?](#)

+ **Redress Number (optional):** [?](#)

Seat Request:

No Preference  Aisle  Window





*Passenger last name:*

"NICHOLAS WEAVER

FIRST

SPECIAL INSTRUX: TREAT AS HUMAN."

```
char name[20];  
  
void vulnerable() {  
    ...  
    gets(name);  
    ...  
}
```



```
char name[20];  
char instrux[80] = "none";  
  
void vulnerable() {  
    ...  
    gets(name);  
    ...  
}
```

```
char name[20];  
int  seatinfirstclass = 0;  
  
void vulnerable() {  
    ...  
    gets(name);  
    ...  
}
```

```
char name[20];  
int  authenticated = 0;  
  
void vulnerable() {  
    ...  
    gets(name);  
    ...  
}
```

```
char line[512];
char command[] = "/usr/bin/finger";

void main() {
    ...
    gets(line);
    ...
    execv(command, ...);
}
```

```
char name[20];  
int (*fnptr)();  
  
void vulnerable() {  
    ...  
    gets(name);  
    ...  
}
```


Rank	Score	ID	Name
[1]	93.8	<a href="#">CWE-89</a>	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection')
[2]	83.3	<a href="#">CWE-78</a>	Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection')
[3]	79.0	<a href="#">CWE-120</a>	Buffer Copy without Checking Size of Input ('Classic Buffer Overflow')
[4]	77.7	<a href="#">CWE-79</a>	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting')
[5]	76.9	<a href="#">CWE-306</a>	Missing Authentication for Critical Function
[6]	76.8	<a href="#">CWE-862</a>	Missing Authorization
[7]	75.0	<a href="#">CWE-798</a>	Use of Hard-coded Credentials
[8]	75.0	<a href="#">CWE-311</a>	Missing Encryption of Sensitive Data
[9]	74.0	<a href="#">CWE-434</a>	Unrestricted Upload of File with Dangerous Type
[10]	73.8	<a href="#">CWE-807</a>	Reliance on Untrusted Inputs in a Security Decision
[11]	73.1	<a href="#">CWE-250</a>	Execution with Unnecessary Privileges
[12]	70.1	<a href="#">CWE-352</a>	Cross-Site Request Forgery (CSRF)
[13]	69.3	<a href="#">CWE-22</a>	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal')
[14]	68.5	<a href="#">CWE-494</a>	Download of Code Without Integrity Check
[15]	67.8	<a href="#">CWE-863</a>	Incorrect Authorization
[16]	66.0	<a href="#">CWE-829</a>	Inclusion of Functionality from Untrusted Control Sphere

```
void vulnerable() {  
    char buf[64];  
    ...  
    gets(buf);  
    ...  
}
```



```
void still_vulnerable() {  
    char *buf = malloc(64);  
    ...  
    gets(buf);  
    ...  
}
```

## IE's Role in the Google-China War



By Richard Adhikari  
TechNewsWorld  
01/15/10 12:25 PM PT

**The hack attack on Google that set off the company's ongoing standoff with China appears to have come through a zero-day flaw in Microsoft's Internet Explorer browser. Microsoft has released a security advisory, and researchers are hard at work studying the exploit. The attack appears to consist of several files, each a different piece of malware.**

Computer security companies are scurrying to cope with the fallout from the Internet Explorer (IE) flaw that led to cyberattacks on [Google](#) (Nasdaq: GOOG) and its corporate and individual customers.

The zero-day attack that exploited IE is part of a lethal cocktail of malware that is keeping researchers very busy.

"We're discovering things on an up-to-the-minute basis, and we've seen about a dozen files dropped on infected PCs so far," Dmitri Alperovitch, vice president of research at [McAfee Labs](#), told TechNewsWorld.

The attacks on Google, which appeared to originate in China, have sparked a feud between the Internet giant and the nation's government over censorship, and it could result in Google pulling away from its business dealings in the country.

### Pointing to the Flaw

The vulnerability in IE is an invalid pointer reference, [Microsoft](#) (Nasdaq: MSFT) said in [security advisory 979352](#), which it issued on Thursday. Under certain conditions, the invalid pointer can be accessed after an object is deleted, the advisory states. In specially crafted attacks, like the ones launched against Google and its customers, IE can allow remote execution of code when the flaw is exploited.

# Disclaimer: x86-32

- For this class, we are going to use 32b x86...
  - Why?
- It is both common and weak...
  - Almost everyone in this class has access to an x86 system: Mac, Linux, Windows...  
And can run a 32b x86 virtual machine
  - 64b x86 systems generally include a lot better "mitigations":  
System defenses designed to limit exploitation in this manner
- But these attacks do apply to other microarchitectures
  - Phones are 64b ARM: Can still be exploited in this manner
  - The Internet of Things is mostly 32b or 64b ARM...  
and often neglects to include the mitigations!

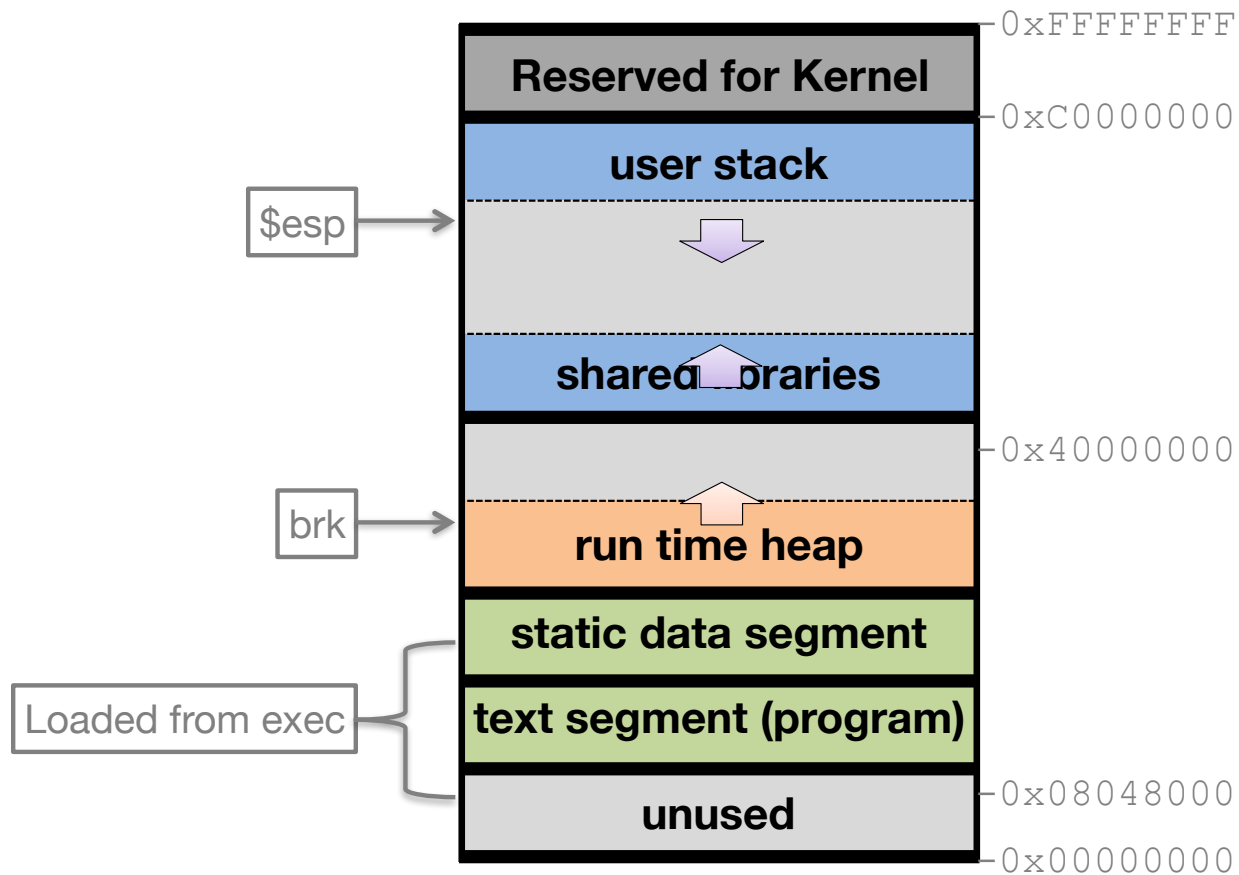
# x86 vs RISC-V

- All RISC architectures are the same except for one or two ‘seems like a good idea at the time’ design decisions  
... But x86 is a very different beast from a programming viewpoint
- RISC-V: 32 general purpose registers (well, 31 + x0...)
  - All operations are on data in registers apart from loads & stores
- x86: only a few registers
  - Operations can be directly on data in memory, including a large number relative to the stack
  - EG, add takes two operands, adds them together, and stores the result in the first
    - The first can be a register or memory location
    - The second can be a register, a memory location, or an immediate...
      - But the first and second can't both be a memory location?!?

# The main x86 registers...

- General purpose: EAX-EDX
  - What you use for computing and other stuff, sorta...
- Indexes & Pointers
  - EBP: “Frame pointer”: points to the start of the current call frame on the stack
  - ESP: “Stack pointer”: points to the current stack
    - PUSH and POP
      - Decrement the stack pointer and store something there
      - Load something and increment the stack pointer
    - Most operations are done with data on the stack...

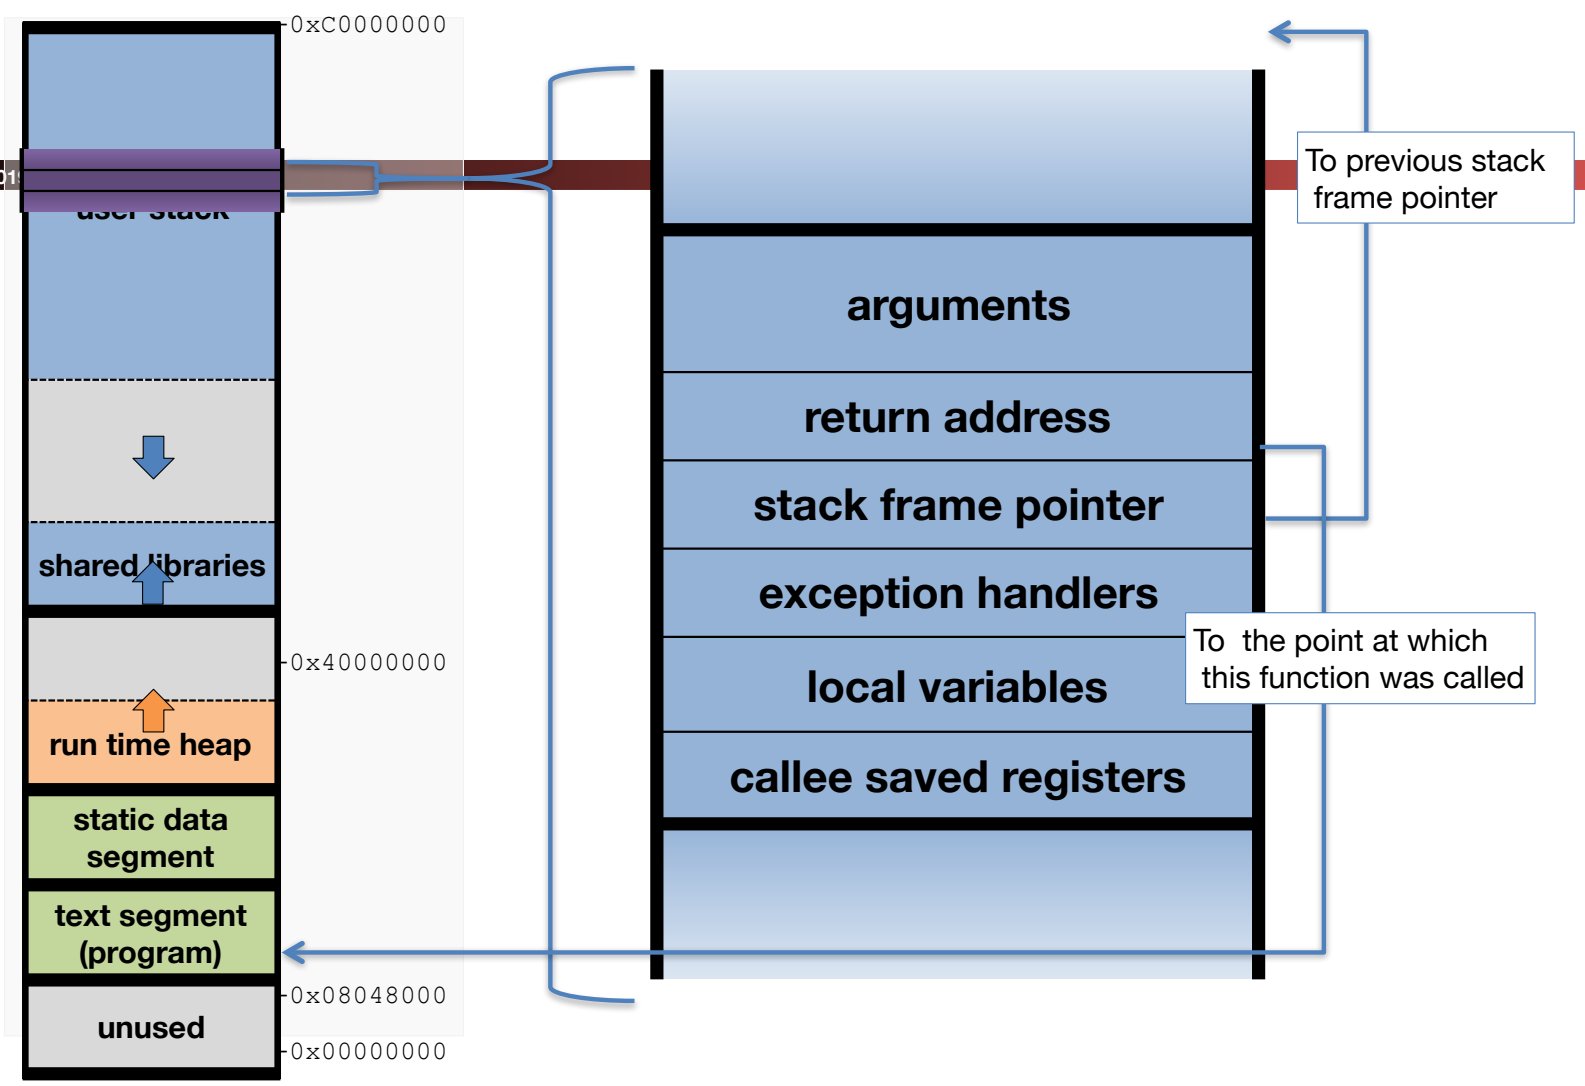
# Linux (32-bit) process memory layout



# x86 function calling

- Place the arguments on the stack
  - Compare with RISC-V where the first arguments are in registers
- CALL the function
  - Which pushes the return address onto the stack (RIP == Return Instruction Pointer)
- do your stuff...
  - Start by saving the old EP on the stack (SFP == Saved Frame Pointer)
- Restore everything
  - Reload EBP, pop ESP as necessary
- RET
  - Which jumps to the return address that is currently pointed to by ESP
  - And can optionally pop the stack a lot further...





```
void safe() {  
    char buf[64];  
    ...  
    fgets(buf, 64, stdin);  
    ...  
}
```

```
void safer() {  
    char buf[64];  
    ...  
    fgets(buf, sizeof(buf), stdin);  
    ...  
}
```

Assume these are both under the control of an attacker.

```
void vulnerable(int len, char *data) {  
    char buf[64];  
    if (len > 64)  
        return;  
    memcpy(buf, data, len);  
}
```

```
memcpy(void *s1, const void *s2, size_t n);
```

size\_t is *unsigned*:  
What happens if len == -1?

```
void safe(size_t len, char *data) {  
    char buf[64];  
    if (len > 64)  
        return;  
    memcpy(buf, data, len);  
}
```

```
void f(size_t len, char *data) {  
    char *buf = malloc(len+2);  
    if (buf == NULL) return;  
    memcpy(buf, data, len);  
    buf[len] = '\n';  
    buf[len+1] = '\0';  
}
```

Is it safe? Talk to your partner.

Vulnerable!

If `len = 0xffffffff`, allocates only 1 byte

## Broward Vote-Counting Blunder Changes Amendment Result

POSTED: 1:34 pm EST November 4, 2004

**BROWARD COUNTY, Fla.** -- The Broward County Elections Department has egg on its face today after a computer glitch misreported a key amendment race, according to WPLG-TV in Miami.

Amendment 4, which would allow Miami-Dade and Broward counties to hold a future election to decide if slot machines should be allowed at racetracks, was thought to be tied. But now that a computer glitch for machines counting absentee ballots has been exposed, it turns out the amendment passed.

"The software is not geared to count more than 32,000 votes in a precinct. So what happens when it gets to 32,000 is the software starts counting backward," said Broward County Mayor Ilene Lieberman.

That means that Amendment 4 passed in Broward County by more than 240,000 votes rather than the 166,000-vote margin reported Wednesday night. That increase changes the overall statewide results in what had been a neck-and-neck race, one for which recounts had been going on today. But with news of Broward's error, it's clear amendment 4 passed.



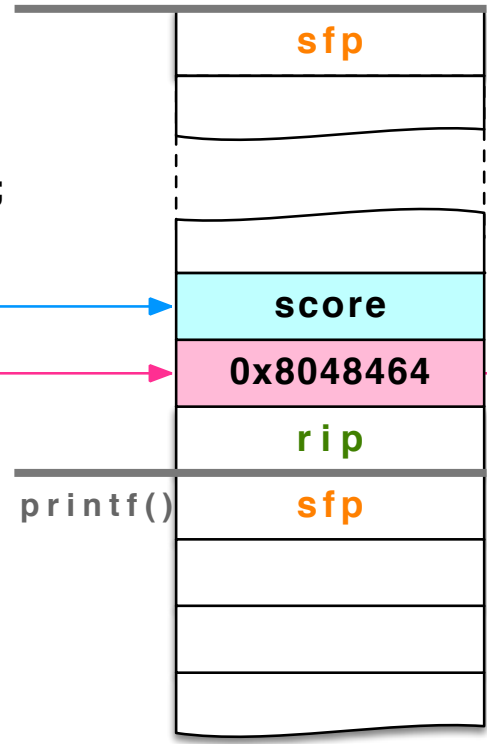
Broward County Mayor Ilene Lieberman says voting counting error is an "embarrassing mistake."

```
void vulnerable() {  
    char buf[64];  
    if (fgets(buf, 64, stdin) == NULL)  
        return;  
    printf(buf);  
}
```



```
printf("you scored %d\n", score);
```

```
printf("you scored %d\n", score);
```

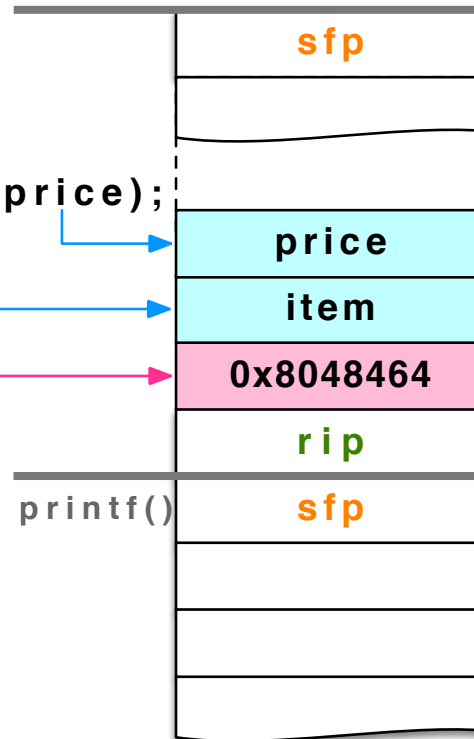


	\0	\n	d
%		d	e
r	o	c	s
	u	o	y

0x8048464

```
printf("a %s costs $%d\n", item, price);
```

```
printf(" a %s costs $%d\n", item, price);
```



\0	\n	d	%
\$		s	t
s	o	c	
s	%		a


0x8048464

# Fun With `printf` format strings...

```
printf("100% dude: ",
```

Format argument is missing!

```
printf("100% dude!");
```



	\0	!	e
d	u	d	
%	0	0	1

0x8048464

# More Fun With `printf` format strings...

```
printf("100% dude!");
```

⇒ *prints value 4 bytes above `retaddr` as integer*

```
printf("100% sir!");
```

⇒ *prints bytes pointed to by that stack entry  
up through first NUL*

```
printf("%d %d %d %d ...");
```

⇒ *prints series of stack entries as integers*

```
printf("%d %s");
```

⇒ *prints value 4 bytes above `retaddr` plus bytes  
pointed to by preceding stack entry*

```
printf("100% nuke'm!");
```

What does the `%n` format do??

`%n` writes the number of characters printed so far into the corresponding format argument.

```
int report_cost(int item_num, int price) {
    int colon_offset;
    printf("item %d:%n $%d\n", item_num,
          &colon_offset, price);
    return colon_offset;
}
```

`report_cost(3, 22)` prints "item 3: \$22"  
and returns the value 7

`report_cost(987, 5)` prints "item 987: \$5"  
and returns the value 9



# Fun With `printf` format strings...

```
printf("100% dude!");
```

⇒ *prints value 4 bytes above `retaddr` as integer*

```
printf("100% sir!");
```

⇒ *prints bytes pointed to by that stack entry  
up through first NUL*

```
printf("%d %d %d %d ...");
```

⇒ *prints series of stack entries as integers*

```
printf("%d %s");
```

⇒ *prints value 4 bytes above `retaddr` plus bytes  
pointed to by preceding stack entry*

```
printf("100% nuke'm!");
```

⇒ ***writes the value 3 to the address pointed to by stack entry***

```
void safe() {  
    char buf[64];  
    if (fgets(buf, 64, stdin) == NULL)  
        return;  
    printf("%s", buf);  
}
```