

# CS 489 / 698: Software and Systems Security

## **Module: Defenses against Common Vulnerabilities**

Lecture: entropy / moving-target defense

Meng Xu (*University of Waterloo*)

Fall 2024

# Outline

- 1 Introduction
- 2 Stack canary
- 3 Randomizing memory addresses
- 4 Entropies in heap allocators
- 5 Security through diversity

# Why entropy in security?

Nondeterminism is useful in software security when

- it has no impact on the intended finite state machine BUT
- limits attackers' abilities to program the weird machine.

# Why entropy in security?

Nondeterminism is useful in software security when

- it has no impact on the intended finite state machine BUT
- limits attackers' abilities to program the weird machine.

**In this slide deck:** we will examine some standard / deployed practices of safely introducing nondeterminism to boost system and software security.

# Choosing pills, a lot of pills




**Figure:** Red pill vs Blue pill. Credits / Trademark: The Matrix Movie

# Outline

- 1 Introduction
- 2 Stack canary**
- 3 Randomizing memory addresses
- 4 Entropies in heap allocators
- 5 Security through diversity


# Recap: stack overflow

```
1 int main() {  
2   char buf[16];  
3   scanf("%s", buf);  
4 }
```



# Solution 1: program analysis

```
1 int main() {  
2     char buf[16];  
3     scanf("%s", buf);  
4 }
```






# Solution 1: program analysis

```
1 int main() {  
2     char buf[16];  
3     scanf("%s", buf);  
4 }
```


```
1 int main() {  
2     char buf[16];  
3     - scanf("%s", buf);  
4     + scanf("%15s", buf);  
5 }
```



## Solution 2: exploit mitigation

```
1 int main() {  
2     char buf[16];  
3     scanf("%s", buf);  
4 }
```

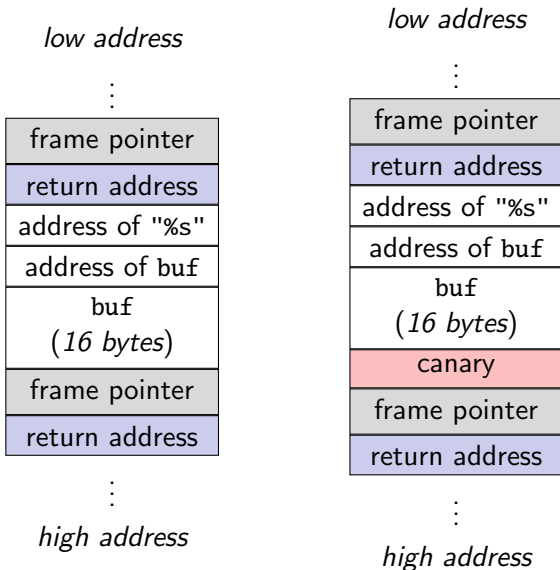
---



## Solution 2: exploit mitigation

```
1 int main() {  
2   char buf[16];  
3   scanf("%s", buf);  
4 }
```

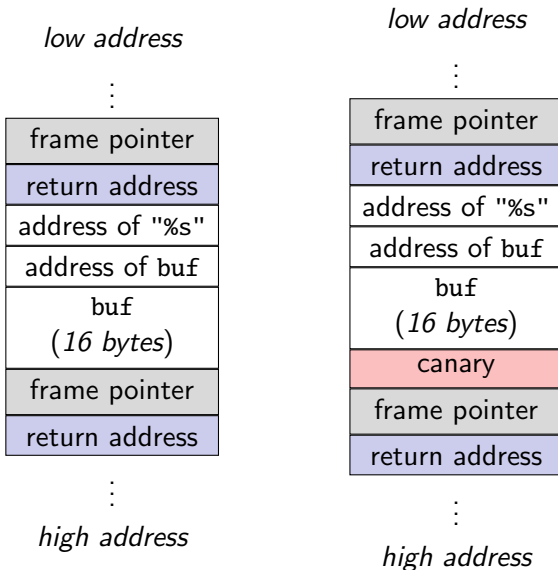
---



## Solution 2: exploit mitigation

```
1 int main() {  
2     char buf[16];  
3     scanf("%s", buf);  
4 }
```

- On function entry, push canary value **X** onto stack.
- On function return, check canary value is still **X**.



# Original use of canary



**Figure:** Canaries in coal-mining. Credits / Trademark: Alamy Stock Photo

# The default implementation in GCC

```
1  int main() {
2    char buf[16];
3    scanf("%s", buf);
4 }

1  extern uintptr_t __stack_chk_guard;
2  noreturn void __stack_chk_fail(void);
3
4  int main() {
5    uintptr_t canary = __stack_chk_guard;
6
7    char buf[16];
8    scanf("%s", buf);
9
10   if ((canary = canary ^ __stack_chk_guard) != 0) {
11     __stack_chk_fail();
12   }
13 }
```

# The default implementation in GCC

```
1 int main() {  
2     char buf[16];  
3     scanf("%s", buf);  
4 }
```

```
1 extern uintptr_t __stack_chk_guard;  
2 noreturn void __stack_chk_fail(void);  
3  
4 int main() {  
5     uintptr_t canary = __stack_chk_guard;  
6  
7     char buf[16];  
8     scanf("%s", buf);  
9  
10    if ((canary = canary ^ __stack_chk_guard) != 0) {  
11        __stack_chk_fail();  
12    }  
13 }
```

- The `__stack_chk_guard` and `__stack_chk_fail` symbols are normally supplied by a GCC library called `libssp`.
- You also have the option of specifying your own value for stack canaries.

# Design choices of stack canaries



## Design choices of stack canaries

- Which value should we use as canary?
  - deterministic? secret? random?

## Design choices of stack canaries

- Which value should we use as canary?
  - deterministic? secret? random?
- What is the granularity of the canary invocation?
  - per function? per execution?

# Design choices of stack canaries

- Which value should we use as canary?
  - deterministic? secret? random?
- What is the granularity of the canary invocation?
  - per function? per execution?
- When to do the integrity check?
  - on function return? is that enough?

# Design choices of stack canaries

- Which value should we use as canary?
  - deterministic? secret? random?
- What is the granularity of the canary invocation?
  - per function? per execution?
- When to do the integrity check?
  - on function return? is that enough?
- How much randomness is needed?
  - 1 byte? 8 bytes? 64 bytes?

## Limitations of stack canary

- Vulnerable to information leak
  - e.g., using a buffer over read to retrieve the canary value

## Limitations of stack canary

- Vulnerable to information leak
  - e.g., using a buffer over read to retrieve the canary value
- Limited protection for frame pointer and return address only
  - other stack variables are not protected

## Limitations of stack canary

- Vulnerable to information leak
  - e.g., using a buffer over read to retrieve the canary value
- Limited protection for frame pointer and return address only
  - other stack variables are not protected
- Unable to defend against arbitrary writes
  - i.e., non-continuous overrides


# Outline

- 1 Introduction
- 2 Stack canary
- 3 Randomizing memory addresses**
- 4 Entropies in heap allocators
- 5 Security through diversity



# Back to the example

```
1 int main() {  
2     char buf[1024];  
3     scanf("%s", buf);  
4 }
```




# Back to the example

```
1 int main() {  
2     char buf[1024];  
3     scanf("%s", buf);  
4 }
```

Meaningful values  
for return address:

- Shellcode (stack)
- system() in libc



# Back to the example



```

1 int main() {
2     char buf[1024];
3     scanf("%s", buf);
4 }

```

Meaningful values  
for return address:

- Shellcode (stack)
- system() in libc

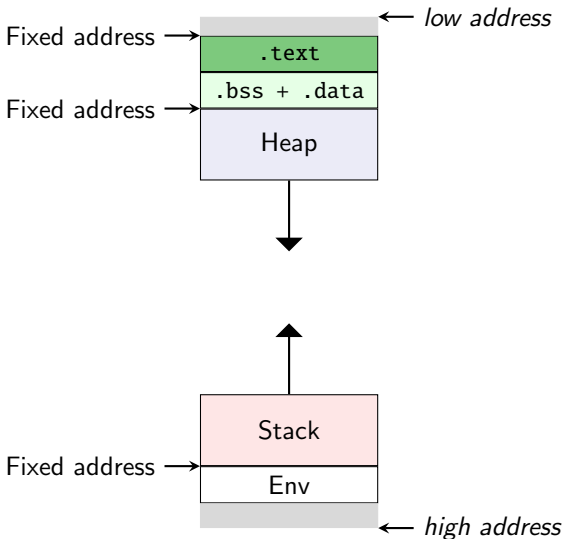


## Randomize the addresses


**ASLR** — Address Space Layout Randomization, is a system-level protection that **randomly** arranges the address space positions of key data areas of a process, including the base of the executable and the positions of the stack, heap and libraries.

**PIE** — Position Independent Executable, is a body of machine code that executes properly **regardless of its absolute address**. This is also known as position-independent code (PIC).

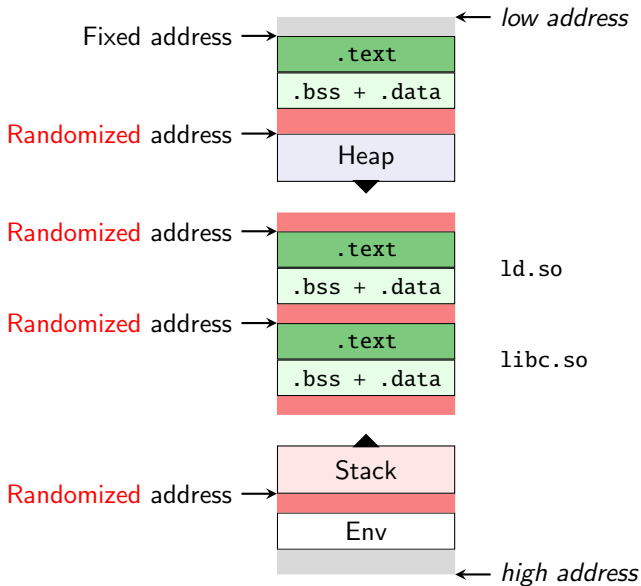
# Base case: static program



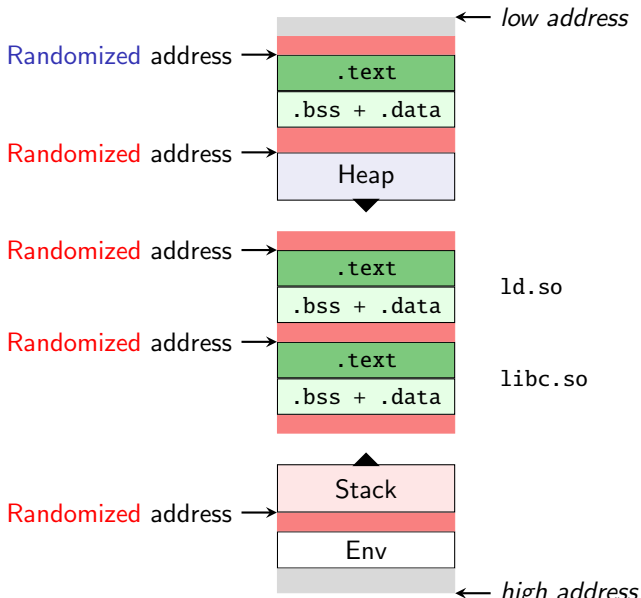
# Static program + shared libraries



# Static program + shared libraries + ASLR

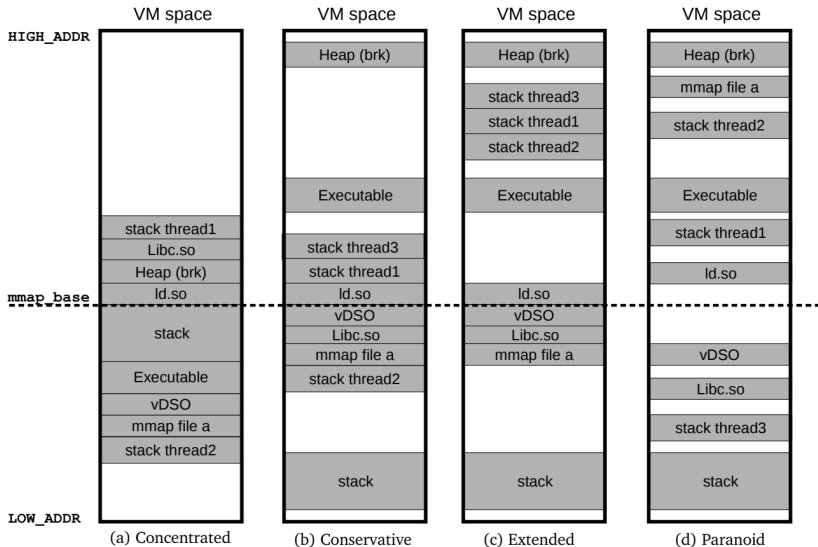


# Static program + shared libraries + ASLR + PIE





# Paranoid randomization



**Figure:** Different level of randomization proposed by the [ASLR-NG project](#)

# Limitations of ASLR + PIE

- Limited entropy
  - visualized by the [ASLR-NG project](#)

# Limitations of ASLR + PIE

- Limited entropy
  - visualized by the [ASLR-NG project](#)
- Memory layout inheritance
  - Child processes inherit/share the memory layout of the parent.

# Outline

- 1 Introduction
- 2 Stack canary
- 3 Randomizing memory addresses
- 4 Entropies in heap allocators**
- 5 Security through diversity

# Motivation for secure heap allocators

Memory errors are equally (if not more) likely to happen on heap objects which can cause all sorts of unexpected behaviors.

# A heap buffer overflow case

```
1 struct dispatcher {
2     uint64_t counter;
3     int (*action)(uint64_t counter, char *data);
4 }
5
6 int main() {
7     char *p1 = malloc(16);
8     char *p2 = malloc(sizeof(struct dispatcher));
9     p2->counter = 0;
10    p2->action = /* some valid function */;
11
12    scanf("%s", p1);
13    int result = p2->action(p2->counter, p1);
14
15    free(p1);
16    free(p2);
17    return result;
18 }
```

# A heap use-after-free case

```
1 struct dispatcher {
2     uint64_t counter;
3     int (*action)(uint64_t counter, char *data);
4 }
5
6 char *p1;
7
8 void main() {
9     p1 = malloc(16);
10    pthread_create(/* ... */, thread_1);
11    pthread_create(/* ... */, thread_2);
12    /* wait for thread termination */
13 }
```

---

```
1 void thread_1() {
2     scanf("%15s", p1);
3     /* ... compromised here ... */
4     /* use-after-free */
5     free(p1);
6     ((struct dispatcher *)p1)
7     ->action = /* bad function */;
8 }
```

```
1 void thread_2() {
2     char *p2 = malloc(
3         sizeof(struct dispatcher));
4     p2->counter = 0;
5     p2->action = /* good function */;
6     p2->action(p2->counter, p1);
7     free(p2);
8 }
```


# Secure heap allocators

These exploits have **implicit assumptions** on the **layout** of the heap, which can be invalidated by a secure heap allocator.




# Basic allocator example

Initial state:



`p1 = malloc(16);`



`p2 = malloc(sizeof(..));`



`free(p1);`



`p3 = malloc(sizeof(..));`




---


<sup>0</sup>Each square is a 4-byte box

# Allocator + random placement

Initial state:



`p1 = malloc(16);`




`p2 = malloc(sizeof(..));`



`free(p1);`



`p3 = malloc(sizeof(..));`




---


<sup>0</sup>Each square is a 4-byte box

# Allocator + random placement + canary

Initial state:



`p1 = malloc(16);`




`p2 = malloc(sizeof(..));`



`free(p1);`



`p3 = malloc(sizeof(..));`



---

<sup>0</sup>Each square is a 4-byte box

# Outline

- 1 Introduction
- 2 Stack canary
- 3 Randomizing memory addresses
- 4 Entropies in heap allocators
- 5 Security through diversity

## Intuition: gene/DNA diversity

**In biology**, maintaining high **genetic diversity** allows species to adapt to future environmental changes, survive from deadly diseases, and avoid inbreeding.

---


## Intuition: gene/DNA diversity

**In biology**, maintaining high **genetic diversity** allows species to adapt to future environmental changes, survive from deadly diseases, and avoid inbreeding.


---

Similarly, we expect **software diversity** to protect software systems (especially critical systems) from deadly viruses and attacks while also serving as an early signal of being attacked.

# Core architecture



# Core architecture (under attack)





# Challenges of applying diversity-based defenses

- Source of diversity
- Synchronization of diversified instances

# Source of diversity

- Compiler/loader-assisted diversity
  - e.g., direction of stack growth
  - e.g., different canary values
  - e.g., different sanitizer instrumentation

# Source of diversity

- Compiler/loader-assisted diversity
  - e.g., direction of stack growth
  - e.g., different canary values
  - e.g., different sanitizer instrumentation
- N-version programming
  - e.g., different language VM (V8 vs SpiderMonkey)
  - e.g., different applications (nginx vs apache web server)
  - e.g., similar applications from independent vendors/teams

# Source of diversity

- Compiler/loader-assisted diversity
  - e.g., direction of stack growth
  - e.g., different canary values
  - e.g., different sanitizer instrumentation
- N-version programming
  - e.g., different language VM (V8 vs SpiderMonkey)
  - e.g., different applications (nginx vs apache web server)
  - e.g., similar applications from independent vendors/teams
- Platform diversity
  - e.g., different libc implementations (glibc vs musl libc)
  - e.g., Adobe Reader on MacOS and Windows
  - e.g., Server programs on Intel and ARM CPUs

# Mode of synchronization

- Online mode (via rendezvous points)
- Offline mode (via record-and-replay)

The key is to synchronize **all sources of nondeterminism**.

〈 End 〉