



ASLR randomizes VMAs



$2^N - 1$ Blue Balls and 1 Red Ball in a Jar



Let N be the number of randomized bits; so there are a total of 2^N possible base addresses. There is only one red ball in a jar, which corresponds to the expected base address. Then, what is the probability of selecting the red ball?

- Case 1: Select balls without replacement (**Windows**).
- Case 2: Select balls with replacement (**Linux**).

Case 1: Selecting Balls w/o Replacement

$$Pr[\text{success on 1st trial}] = \frac{1}{2^N}$$

Case 1: Selecting Balls w/o Replacement

$$\begin{aligned} Pr[\text{success only on 3rd trial}] &= \left(\frac{2^N - 1}{2^N}\right) \left(\frac{2^N - 2}{2^N - 1}\right) \left(\frac{1}{2^N - 2}\right) \\ &= \frac{1}{2^N} \end{aligned}$$

$$\begin{aligned} Pr[\text{success only on } k\text{th trial}] &= \left(\frac{2^N - 1}{2^N}\right) \times \dots \times \left(\frac{2^N - k + 1}{2^N - k}\right) \left(\frac{1}{2^N - k + 1}\right) \\ &= \frac{1}{2^N} \end{aligned}$$

Case 1: Selecting Balls w/o Replacement

Expected number of trials before success:

$$\begin{aligned}\sum_{k=1}^{2^N} k \cdot Pr[\text{success only on } k\text{th trial}] &= \sum_{k=1}^{2^N} \frac{k}{2^N} \\ &= \frac{1}{2^N} \sum_{k=1}^{2^N} k \\ &= \frac{1}{2^N} \frac{2^N(2^N + 1)}{2} \\ &= \frac{2^N + 1}{2}\end{aligned}$$

Case 2: Selecting Balls w/ Replacement

$$Pr[\text{success on 1st trial}] = \frac{1}{2^N}$$

$$Pr[\text{success on 2nd trial}] = \left(1 - \frac{1}{2^N}\right) \frac{1}{2^N}$$

$$Pr[\text{success on } k\text{th trial}] = \left(1 - \frac{1}{2^N}\right)^{k-1} \frac{1}{2^N}$$

The classic geometric distribution where $p = \frac{1}{2^N}$.

Case 2: Selecting Balls w/ Replacement

Expected number of trials before success:

$$\begin{aligned} E[x] &= \frac{1}{p} \text{ (for geometric distribution)} \\ &= 2^N \end{aligned}$$

Comparison: Windows vs. Linux

Brute-force attack will succeed in

- $\frac{2^N+1}{2} \approx 2^{N-1}$ trials on Windows.
- 2^N trials on Linux.

Hence, Linux is $\approx 2\times$ safer than Windows against a brute-force attack.

Security vs. Performance?

Windows's ASLR is faster, but less secure than it of Linux.

Attack #2: Exploiting Fixed Addresses

Most binaries (before 2016) had non-randomized segments (VMAs).

Before 2016, compilers created *non-PIE*⁴ executables by default.

⁴Non Position-Independent Executable.

Position-Independent Executable (PIE)

Position-Independent Code (PIC) or PIE is code that runs regardless of its location (e.g., shellcode).

- “gcc” will produce a PIE by default.
- “gcc -fno-pic -no-pie” will produce a non-PIE.

PIE vs. non-PIE

```
...  
lea RAX, qword ptr [RIP - 0x25]  
mov qword ptr [RBP-0x8], RAX  
mov RAX, qword ptr [RBP-0x8]  
mov EDI, 0x2a  
call RAX  
...
```

```
...  
mov qword ptr [RBP-0x8], 0x401106  
mov RAX, qword ptr [RBP-0x8]  
mov EDI, 0x2a  
call RAX  
...
```

Any Libraries Must Be Position-Independent

Any shared objects (.so), such as LIBC, are position-independent.

Why?

Legacy Binaries Are Not a PIE

- 93% of Linux binaries were not a PIE (in 2009).
- Thus, the code sections were not randomized.
- Thus, code reuse attacks (e.g., ROP) are still possible on legacy binaries.

But, why?

Security vs. Performance

- Relative-addressing instructions are slower than absolute-addressing instructions.
- Performance overhead of PIE on x86 is 10% on average⁵.
- Most applications on current x86 are still not PIEs.

⁵Too much PIE is bad for performance, ETH Techreport, 2012.

ROP-based Attack on Legacy Binaries

- Code sections are not randomized, hence we can use ROP.
- But, LIBC address is randomized! Cannot directly return to LIBC functions.

But, still, relative offsets between LIBC functions are the same regardless of ASLR.

Exploitation Idea

- If a LIBC function has been invoked at least once, GOT should contain a concrete address of the function in LIBC.
- Therefore, we will read the GOT entry using ROP and compute the address of `system` by using the relative offset between the LIBC function and `system`.

$$\begin{aligned}(\text{addr of } \text{system}) &= (\text{addr of } \text{open}) \\ &+ (\text{offset from } \text{open} \text{ to } \text{system} \text{ in LIBC})\end{aligned}\tag{1}$$

Possible Defenses?

- Use PIEs.
- Use 64-bit CPU: lots of entropy.
- Detect brute-forcing attacks (as there should be many crashes in a short amount of time).
- Use non-forking servers.

Better ASLR?

Single pointer leakage can reveal the entire memory layout of a VMA.

Can we make it harder?

Fine-grained ASLR

Randomize code within a VMA boundary.

- Function-level randomization.
- Block-level randomization.
- Instruction-level randomization.

Fine-grained ASLR: Design Challenges

- Can we apply fine-grained ASLR without debugging information?
- How often should we apply fine-grained ASLR? Performance impact?

Example: In-Place Code Randomization

Smashing the Gadgets: Hindering Return-Oriented Programming Using In-Place Code Randomization, *Oakland 2012*.

- Instruction reordering.

```
mov ebx, 42
mov eax, [ecx]
```

```
mov eax, [ecx]
mov ebx, 42
```

- Instruction substitution.

```
mov rax, 0
```

```
xor rax, rax
```

- Register re-allocation.

```
mov rax, [rcx]
call [rax]
```

```
mov rbx, [rcx]
call [rbx]
```


Example: Instruction Location Randomization

ILR: Where'd My Gadgets Go?, *Oakland 2012*.

```
mov rbx, 42
mov rax, [rcx]
add rax, rcx
ret
```



A `add rax, rcx`

B `mov rbx, 42`

C `ret`

D `mov rax, [rcx]`

Fall-through map:

A \mapsto C

B \mapsto D

D \mapsto A

Question?