

Control-Flow Integrity

20180125

Outline

Vulnerability protection

Stack canaries

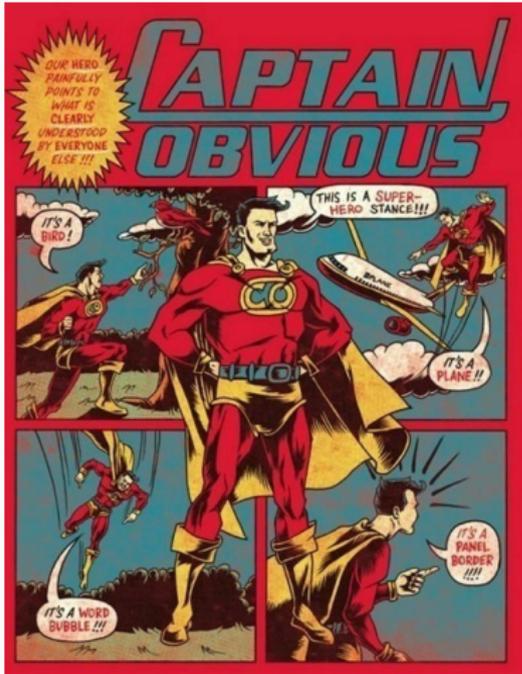
Executable space protection

ASLR

CFI on execution

Vulnerability protection

Write correct code



Some people write fragile code and some people write very structurally sound code, and this is a condition of people.

– K. Thompson

To err is human, but to really foul up requires a computer.

– Anon

Stack canaries

Stack canaries

What it is

A public canary value is placed right above function-local stack buffers in the stack frame.

Its integrity is checked prior to function return.

AKA cookie, stack cookie

What it provides

Ensure the saved base pointer and function return address have not been corrupted

How it looks



Summary

The good

- Pure compiler-based solution (no OS support)
- Most stack-based buffer overflows are countered

The bad

- Protect only variables **above** it in the stack
- Not always active
- Sometimes the cookie can be guessed (see later)

Implementations

VS `/Gs[size]`

If a function requires more than `size` bytes of stack space for local variables, its stack probe is initiated. By default, the compiler generates code that initiates a stack probe when a function requires **more than one page of stack space** (i.e. `/Gs4096`).

GCC `-fstack-protector`

Emit extra code to check for buffer overflows, such as stack smashing attacks. This is done by adding a guard variable to functions with vulnerable objects. This includes functions that call `alloca`, and functions with **buffers larger than 8 bytes**.

Terminator canary

Definition

A **terminator canary** is comprised of common termination symbols, such as `'\0'` (0x00), `'"` (0x0a), `'\r'` (0x0d), EOF (-1)

Example: 0x000a0dff

Effectiveness

The attacker cannot use common C string libraries, since copying functions will terminate on the termination symbols.

- Either the attack is detected (canary does not hold the same value)
- Or it stops it due to termination symbols.

Random canary

Definition

The **loader** chooses a word-sized (32/64 bits) random canary string on program start.

Effectiveness

The randomness makes the value of the canary hard to guess

Behavior

```
1 #include <string.h>
2
3 int main(int argc, char *argv[])
4 {
5     char buf[10];
6     strcpy(buf, argv[1]);
7     return buf[5];
8 }
```

StackGuard effectiveness (Cowan et al., 2000)

Program	without protection	with protection
dip 3.3.7	root shell	program halts
elm 2.4	root shell	program halts
perle 5.003	root shell	program halts
Samba	root shell	program halts
SuperProbe	root shell	program halts
umount / libc 5.3.12	root shell	program halts
wwwcount 2.3	httpd shell	program halts
zgv 2.7	root shell	program halts

Considerations

Efficiency

Canary checks **for every function** causes a performance penalty.

≈ 8% for Apache

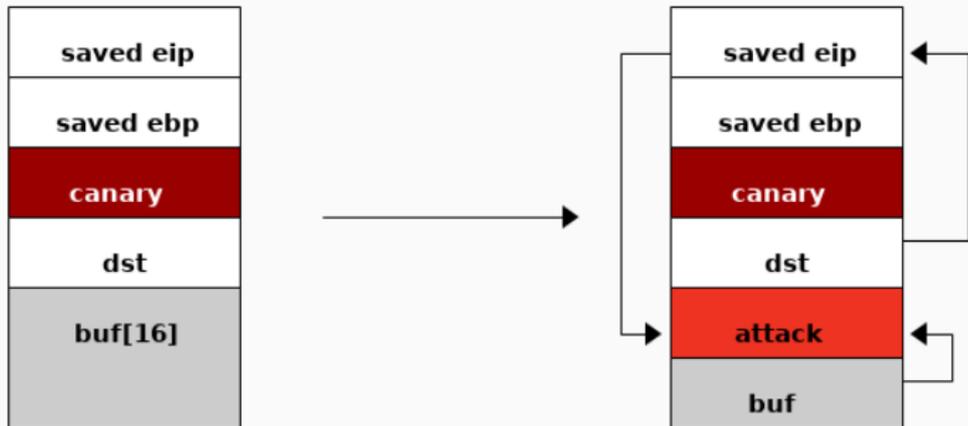
PointGuard

Canaries are also places next to

- function pointers
- setjmp buffers

Greater performance impact

Defeating canaries



Example vulnerable on prior versions

```
1 int f (char ** argv)
2 {
3     int pipa;           // useless variable
4     char *p;
5     char a[30];
6
7     p=a;
8
9     printf ("p=%x\t -- before 1st strcpy\n",p);
10    strcpy(p,argv[1]);    // <= vulnerable strcpy()
11    printf ("p=%x\t -- after 1st  strcpy\n",p);
12    strncpy(p,argv[2],16);
13    printf("After second strcpy ;)\n");
14 }
15
16 int main (int argc, char ** argv) {
17     f(argv);
18     execl("back_to_vul","",0); //<-- The exec that fails
19     printf("End of program\n");
20 }
```

Weakness of canary randomization

Canary is randomized whenever `libc` is loaded.

That is every time, `execve()` is used – but **not** when `fork()` is used

Brute-forcing the canary

Technique :: Byte-per-byte brute-forcing

- On average ≈ 512 attempts
- Brute-force + timing analysis
- Incorrect guesses fail fast, correct guesses fail slow

Limitations

- Need the canary to stay the same (i.e. forking daemons)

Executable space protection

Broad idea

- C does not specify what happens when a data pointer is used as if it were a function pointer
- Self-modifying code is pretty rare (efficient JIT compilers)

Idea

- Mark data memory as non-executable
- Needs OS support

Implementations

OS	Date	Version	Name(s)
OpenBSD	2003	3.3	W^X
Windows	2004	XP	DEP
FreeBSD	2004	5.3	
Linux	2004	2.6	PaX, ExecShield
macOS	2005	10.4	
macOS	2007	> 10.5	

Implementation details

NX/XD/XN bit

Modern AMD/Intel/ARM machines have a dedicated bit which flags memory pages as writable or else executable.

When set, the page is **not** executable

x86's original 32-bits table did not have such a mechanism.

Other implementations

- On x86, the mechanism is sometimes emulated (through CS segment)
- PaX NX also emulates the functionality on 32-bits

Limitations

Warning

Data Execution Prevention does nothing to prevent a buffer overflow to rewrite the saved frame pointer or the saved instruction pointer (aka. return address).

A single call to `SqlExe("drop table ...")` is thus manageable.

Counterattacks

- Indirect code injection
- Jump-to-libc attacks
- Data-only attacks

Return-oriented programming

Definition

Return oriented programming (ROP) is an exploit technique

1. Gains control of the call stack
2. Executes carefully chose machine instruction sequences **already present** called gadgets

Remarks

- There exist Turing-complete sets of gadgets
- This is an extension to `return-into-libc` attacks

Overlapping instructions (J. Kinder)

Other instructions are embedded inside your instructions.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
B8	00	03	c1	BB	B9	00	00	00	05	03	c1	EB	F4	03	c3	c3
mov eax,0xBBC10300				mov eax,0x05000000				add		jmp -10		add		ret		
add		mov ebx, 0xB9				add eax,0xF4EBC103				add		ret				
↑ jump in the middle																

This can be used to find gadgets inside your code, e.g. `jmp esp (0xffe4)`

Gadgets

- Gadgets ending with a ret are typically found in function epilogues
- Tools (ropper, ROPgadget, ...) help in finding gadgets and ROP chains to

Origin

- Intended instructions
- Unaligned bytes

Build

- String gadgets into units of functionality (loads/stores, jumps, arithmetic)
- Goal : execute another shellcode

Basic example

```
1  #include <stdio.h>
2  #include <string.h>
3  #include <stdlib.h>
4
5  void not_called(int pseudo_arg)
6  {
7      printf("Enjoy your shell!\n");
8      system("/bin/bash");
9  }
10
11 void vulnerable_function(char* string)
12 {
13     char buffer[100];
14     strcpy(buffer, string);
15 }
16
17 int main(int argc, char** argv)
18 {
19     vulnerable_function(argv[1]);
20     return 0;
21 }
```

More involved example

```
1  #include <stdio.h>
2  #include <string.h>
3  #include <stdlib.h>
4
5  char* not_used = "/bin/sh";
6
7  void not_called(int pseudo_arg) {
8      printf("Not quite a shell...\n");
9      system("/bin/date");
10 }
11
12 void vulnerable_function(char* string) {
13     char buffer[100];
14     strcpy(buffer, string);
15 }
16
17 int main(int argc, char** argv) {
18     vulnerable_function(argv[1]);
19     return 0;
20 }
```

kBouncer (Pappas et al., 2013)

Observation 1

- ROP attacks issue returns to *non-call-preceded* addresses
- Make all return instructions target call-preceded addresses

Observation 2

- ROP attacks are built of long sequences of short gadgets
- Do not allow long sequences of short gadgets

Based on stack history, decide to abort

State-of-the-art

Lightweight ROP countermeasures are still exploitable

Stronger defenses

- G-Free (K. Onarlioglu et al. 2010) remove unintended return instructions and encrypt return addresses

ASLR

Address-space Layout Randomization

Definition

ASLR is a technique to prevent exploitation of memory corruption vulnerabilities.

It rearranges the address space positions of a process, e.g., the base of the executable, the stack, the heap, and libraries.

Limitations

- Needs OS support
- ASLR + NX needs PIE

How it works

Most everything can be randomized that way :

- code
- global variables
- heap allocations, ...

ASLR basically consists of randomly distributing the fundamental parts of a process (executable base, stack pointers, libraries, ...)

Is it enabled ?

```
1 | ldd /bin/ls
```

```
linux-vdso.so.1 (0x00007ffec37b9000)  
libc.so.6 => /usr/lib/libc.so.6 (0x00007fab19bdc000)  
/lib64/ld-linux-x86-64.so.2 => /usr/lib64/ld-linux-x86-64.so.2 (0x00007fab1a3b9000)
```

```
1 | ldd /bin/ls
```

```
linux-vdso.so.1 (0x00007ffe46556000)  
libc.so.6 => /usr/lib/libc.so.6 (0x00007fced1697000)  
/lib64/ld-linux-x86-64.so.2 => /usr/lib64/ld-linux-x86-64.so.2 (0x00007fced1e74000)
```

What is actually randomized ?

```
1 | cat /proc/self/maps | grep -E 'stack|heap|libc'
```

- Run 1

560c97822000-560c97843000	rw-p	0	00:00	0	[heap]
7f7d15487000-7f7d15635000	r-xp	0	fe:02	2885816	/usr/lib/libc-2.26.so
7f7d15635000-7f7d15834000	—p	001ae000	fe:02	2885816	/usr/lib/libc-2.26.so
7f7d15834000-7f7d15838000	r-p	001ad000	fe:02	2885816	/usr/lib/libc-2.26.so
7f7d15838000-7f7d1583a000	rw-p	001b1000	fe:02	2885816	/usr/lib/libc-2.26.so
7ffc083b9000-7ffc083da000	rw-p	0	00:00	0	[stack]

- Run 2

55b060201000-55b060222000	rw-p	0	00:00	0	[heap]
7f8df42e3000-7f8df4491000	r-xp	0	fe:02	2885816	/usr/lib/libc-2.26.so
7f8df4491000-7f8df4690000	—p	001ae000	fe:02	2885816	/usr/lib/libc-2.26.so
7f8df4690000-7f8df4694000	r-p	001ad000	fe:02	2885816	/usr/lib/libc-2.26.so
7f8df4694000-7f8df4696000	rw-p	001b1000	fe:02	2885816	/usr/lib/libc-2.26.so
7ffe7a775000-7ffe7a796000	rw-p	0	00:00	0	[stack]

Implementations

OS	Date	Version
OpenBSD	2003	3.3
Linux	2005	2.6.12
Windows	2007	Vista
macOS	2007	> 10.5

- FreeBSD still has no support in `-CURRENT`

Impact on execution

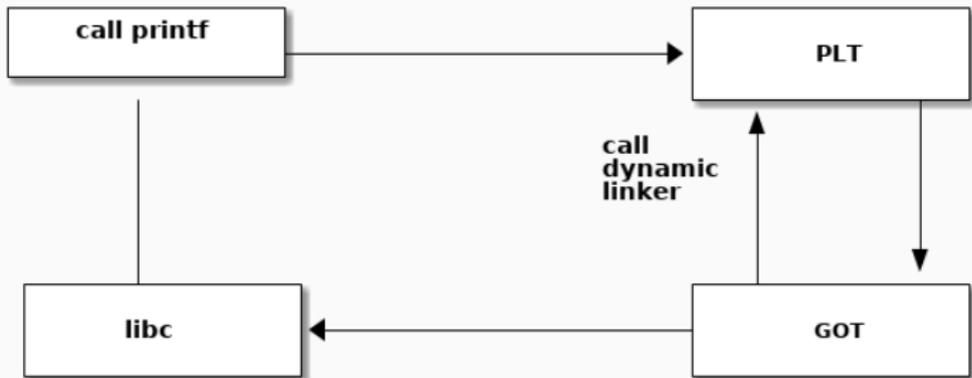
ASLR has a moderate impact ($\approx 3\%$) on performance

Attacking ASLR

- Parts of addresses are not randomized (i.e. GOT)
- Data and BSS segments are mapped to static locations.
Most applications have at least one interesting global
- Any info leak disclosing location can be used to "guess" the where gadgets are.

.got & .plt

- GOT : Global Offset Table
- PLT : Procedure Linking Table



Further protections: RELRO

Definition

RELRO is a generic mitigation technique to harden the data sections of an ELF binary/process.

Partial RELRO

- `gcc -Wl,-z,relro`
- Reorders the binary :
.got, .dtors precede data sections
- non-PLT GOT is RO
- GOT **still writable**

Full RELRO

- `gcc -Wl,-z,relro,-z,now`
- Partial RELRO + GOT is read-only

Definition

KASLR randomizes the kernel code location in memory when the system boots

Problem

The kernel cannot change its distribution in memory throughout its operating time, ie until the next time the system is restarted a new random distribution in memory will not be performed.

Definition (OpenBSD)

Kernel binary files are generated by distributing the kernel's internal files in a random order each time the system is restarted or updated, so each system will work every time it is booted with a unique kernel totally different from other systems at binary level

Our immune systems work better when they are unique. Otherwise one airline passenger from Singapore with a new flu could wipe out Europe (they should fly to Washington instead).

Our computers should be more immune.
– Theo de Raadt

CFI on execution

General idea

Compiler generates a static over-approximation of licit jump sites for **all** dynamic jumps.

At runtime, it is checked that jump targets are authorized.

Example (U. Erlingsson et al.)

```
1 bool lt(int x, int y)
2 {
3     return x < y;
4 }
5
6 bool gt(int x, int y)
7 {
8     return x > y;
9 }
10
11 sort2(int a[], int b[], int len)
12 {
13     sort(a, len, lt);
14     sort(b, len, gt);
15 }
```

Property

The CFI security policy dictates that software execution must follow a CFG path determined ahead of time.

The CFI security policy needs be **conservative**: i.e. all valid executions should be allowed event at the cost of allowing invalid executions.

Overhead and slowdown

Code-size increase

≈ 8%

Execution slowdown

0%–45% (mean: 16%)

Lightweight CFI

Control-flow destinations must be aligned on multi-word boundaries.

- Allow all basic blocks
- Basically only disallows jumping into overlapping instructions

Other measures

Sanitizers are runtime checkers dedicated to specific bugs

Memory sanitization (ASan)

Detect out-of-bound and use-after-free bugs

Undefined behavior sanitization (UBSan)

Detects the used of undefined behaviors at runtime

Impact

- 73% processing time
- 340% memory usage

Pre-summary

Protection	Exploitation
NX	easy
ASLR	feasible
canaries	depends
NX + ASLR	feasible
NX + canaries	depends
ASLR + canaries	hard
All 3	hard

Summary

Memory corruption vulnerabilities are well-addressed by the combination of

- W~X
- Stack canaries
- ASLR

Using only one of these techniques is **not enough**.

Compilers are including more advanced measures (CFI, sanitizers) to further mitigate these issues.

Questions ?



<https://rbonichon.github.io/teaching/2018/asi36/>