

Overview

Last week

- Spotting memory corruption bugs is hard!
 - Format string attacks are easier to spot
 - Undefined behaviour (eg integer overflow or null dereferencing) allows weird compiler behaviour
- Countermeasures
 - Stack canaries
 - ASLR
 - distinguishing writable & executable memory
 - NX aka $W\oplus X$, W^X , DEP

Today

- more advanced countermeasures

More advanced defences

[See SoK Eternal War in Memory paper]

Types of (building blocks for) attacks

- **Code corruption attack**
Overwrite the original program code in memory
Impossible with $W\oplus X$
- **Control-flow hijack attack**
Overwrite a **code pointer**, eg **return address**, **jump address**, **function pointer**, or **pointer in vtable** of C++ object
- **Data-only attack**
Overwrite some data, eg `bool isAdmin;`
- **Information leak**
Only reading some data; e.g. Heartbleed attack on TLS

Control flow hijack via code pointers

- A compiler translates **function calls** in source code to **call <address>** or **JSR <address>** in machine code where **<address>** is the location of the code for the function.
- For a function call **f (...)** in C a static address (or offset) of the code for **f** may be known **at compile time**.
If compiler can hard-code this static address in the binary, **W \oplus X** can prevent attackers from corrupting this address
- For a **virtual function call o.m (...)** in C++ the address of the code for **m** typically has to be determined **at runtime**, by inspecting the virtual function table (**vtable**)
W \oplus X does not prevent attackers from corrupting code pointers in these tables

Classification of defences [SoK paper Eternal War in Memory]

- **Probabilistic methods**

Basic idea: **add randomness to make attacks harder**

- in location where certain data is located (eg ASLR),
or in the way data is represented in memory (eg pointer encryption)

- **Memory Safety**

Basic idea: **do additional bookkeeping & add runtime checks to prevent some illegal memory access**

- **Control-Flow Hijack Defenses**

Basic idea: **do additional bookkeeping & add runtime check to prevent strange control flow**

More randomness: Pointer Encryption (PointGuard)


- Many buffer overflow attacks involve corrupting pointers, **pointers to data** or **code pointers**
- To complicate this: **store pointers encrypted in main memory, unencrypted in registers**
 - simple & fast encryption scheme: eg. XOR with a fixed value, randomly chosen when a process starts
- Attacker can still corrupt encrypted pointers in memory, but these will not decrypt to predictable values
 - This uses ***encryption to ensure integrity.*** Normally NOT a good idea, but here it works.
- More extreme variant: **Data Space Randomisation (DSR)**
 - store not just pointers encrypted in main memory, but store all data encrypted in memory
 - Some **AMD** chips support this under name **SME (Secure Memory Encryption)** that uses AES

Recent trends on pointer encryption/authentication

- **Pointer Authentication on Qualcomm ARMv8.3**
if not all 64 bits are needed for pointers, remaining bits can be used for a **PAC (Pointer Authentication Code)**
 - 3 – 24 bits PACs using fast **QARMA** cipher
- Joan Daemen's PhD student Yanis Belkheyar in our group works on lightweight ciphers suitable for pointer encryption for **Intel's Cryptographic Capability Computing (C³)**
 - Lightweight can be lightweight in 1) power consumption, 2) surface area of hardware implementation, or 3) time.
For pointer encryption/authentication, time (aka latency) is crucial.

More memory safety

Additional book-keeping of meta-data
& extra runtime checks to prevent illegal memory access




Different possibilities

- add information to **pointer** about size of **memory chunks** it points to (**fat pointers**)
- add information to **memory chunks** about their size (**Spatial safety with object bounds**)
- ...

Fat pointers

The compiler

- records size information for all pointers
- adds runtime checks for pointer arithmetic & array indexing



Downsides

- Considerable execution time overhead
- Not binary compatible – ie all code needs to be compiled to add this book-keeping for all pointers

More memory safety

Additional book keeping of meta-data
& extra runtime checks to prevent illegal memory access



Different possibilities

- add information to **pointer** about size of **memory chunks** it points to (**fat pointers**)
- add information to **memory chunks** about their size (**Spatial safety with object bounds**)
- keep a shadow administration of this meta-data, separate from the pointers & the existing memory (**SoftBounds**)
- keep a shadow administration of which memory cells have been allocated (**Valgrind, Memcheck, AddressSanitizer or ASan**)
 - to also spot **temporal** bugs, ie. malloc/free bugs

Object-based temporal safety (Valgrind, Memcheck, ASan)

Shadow admin

1	1	1	1	1	1	1	1
0	0	0	0	0	0	0	0
0	0	1	1	1	1	1	1

of allocated memory


s	o	m	e	d	a	t	a
o	l	d	j	u	n	k	x
y	z	h	e	l	l	o	\0

to keep track of which memory is **allocated**, to generate runtime error when code tries to read/write **unallocated** memory

- Can also catch spatial bugs, ie. small buffer overruns, by keeping empty space between allocated chunks (unless overrun is huge)
 - small overrun will end up in this unallocated space
- Cannot spot illegal access via a stale pointer if the data chunk it points to has been re-allocated
 - Eg the last bug, line 3004, on slide 20 from last week

Guard pages to improve memory safety

Allocate chunks with the end at a **page boundary** with a non-readable, non-writable page  between them



Buffer overwrite or overread will cause a memory fault.

Small execution overhead, but **big** memory overhead

Control Flow Integrity (CFI)

Extra bookkeeping & checks to spot unexpected control flow

- **Dynamic return integrity**

Stack canaries, or **shadow stack** that keeps copies of all return addresses, providing extra check against corruption of return addresses

- **Static control flow integrity**

Idea: **determine the control flow graph (cfg) and monitor jumps in the control flow to spot deviant behavior**


If $f()$ never calls $g()$,
because $g()$ does not even occur in the code of $f()$,
then call from $f()$ to $g()$ is suspicious,
as is a return from $g()$ to $f()$

Interrupting execution when this happens prevents (some) attacks.

This can detect some **Return-to-libc** and **ROP** attacks

Static control flow integrity: example code & CFG

```
void f() {  
    ... ; g();  
    ... ; g();  
    ... ; h();  
    ...  
}  
void g() { ..h(); }  
void h() { ... }
```



Before and/or after every control transfer (**function call** or **return**) we could check if it is legal – ie. allowed by the CFG

Some weird returns would still be allowed

- eg if we call `h()` from `g()`, and the return is to `f()`, this would be allowed by the static cfg
- Additional *dynamic* return integrity check can narrow this down to actual call site – using recorded call site on shadow stack

Downsides of static control flow integrity checks

- Requires a **whole program analysis**
- Use of function pointers in C or virtual functions in C++ (that both result in so-called **indirect control transfers**) complicate compile-time analysis of the cfg: we would need
 - a **points-to analysis** to determine where such code pointers can point to
 - eg in C++, if `Animal.eat()` can resolve to `Cat.eat()` or `Dog.eat()`, so both these addresses are valid targets for transferring control
 - or: simply allow transfer to any function entry point

New(er) features of modern OS

Stack canaries, ASLR, and NX are standard, except on very cheap devices (eg in IoT).

Some fancier features are slowly becoming used:

- **Pointer encryption in iOS (2018)**
- **Hardware-enforced Stack Protection in Windows 10 (2020)**
 - with a **shadow stack**,
using Intel **Control-flow Enforcement Technology (CET)**

<https://techcommunity.microsoft.com/t5/windows-kernel-internals/understanding-hardware-enforced-stack-protection/ba-p/1247815>

For more info: Evolution of CFI at Microsoft discussed by Joe Bialek

<https://www.youtube.com/watch?v=oOqpl-2rMTw>

The Evolution of CFI Attacks and Defenses @ OffensiveCON 18

Exam questions: you should be able to

- Explain how simple buffer overflows work & what root causes are
- Spot a *simple* buffer overflow, memory-allocation problem, format string attack, or integer overflow in some C code
- Explain how countermeasures - such as stack canaries, ASLR, non-executable memory, CFI, bounds checkers, pointer encryption - work
- Explain why they might not always work