

Secure Software Development

Countermeasures: Exploitation Prevention & Privilege Minimization

Daniel Gruss, Vedad Hadzic, Lukas Maar, Stefan Gast, Marcel Nageler

02.12.2022

Winter 2023/24, www.iaik.tugraz.at



1. Stack Buffer Overflows
2. Code Injection Attacks
3. Code Reuse Attacks
4. Memory Safety
5. Privilege Minimization
6. In-process Sandboxing
7. Process Sandboxing

Attacker's perspective

- 🔍 Vulnerability discovery
- 👤 Exploitation
- 🔑 Privilege elevation

Defender's perspective

- 🔍 Vulnerability prevention
- 👤 **Exploit prevention** (today)
- 🔑 **Privilege minimization** (today)



Attacker's perspective

Vulnerability discovery

- buffer/integer overflow, use-after-free, format strings, type confusion

Exploitation

- Data corruption, shellcode, code reuse, ROP, return-to-libc

Privilege elevation

- exploit suid binaries, kernel exploits, crack root PW hash ;)

Attacker's perspective

Vulnerability discovery

- buffer/integer overflow, use-after-free, format strings, type confusion

Exploitation

- Data corruption, shellcode, code reuse, ROP, return-to-libc

Privilege elevation

- exploit suid binaries, kernel exploits, crack root PW hash ;)

Defender's perspective

Vulnerability prevention

- Code quality, memory safety, type safety, error handling ...

Exploit prevention

- Compiler/runtime defenses, hardware defenses

Privilege minimization

- System call filtering, sandboxing, virtualization



 Attacker triggered a vulnerability



- 👤 Attacker triggered a vulnerability
 - Part 1: Can we prevent exploitation?



 Attacker triggered a vulnerability

- Part 1: Can we prevent exploitation? → **Exploit Prevention**



- 👤 Attacker triggered a vulnerability
 - Part 1: Can we prevent exploitation? → **Exploit Prevention**
- 🔑 Attacker gained arbitrary code execution



- 👤 Attacker triggered a vulnerability
 - Part 1: Can we prevent exploitation? → **Exploit Prevention**
- 🔑 Attacker gained arbitrary code execution
 - Part 2: Can we prevent further damage?



- 👤 Attacker triggered a vulnerability
 - Part 1: Can we prevent exploitation? → **Exploit Prevention**
- 🔑 Attacker gained arbitrary code execution
 - Part 2: Can we prevent further damage? → **Privilege Minimization**



- 👤 Attacker triggered a vulnerability
 - Part 1: Can we prevent exploitation? → **Exploit Prevention**
- 🔑 Attacker gained arbitrary code execution
 - Part 2: Can we prevent further damage? → **Privilege Minimization**
- 👍 Defenses must be cheap!

Stack Buffer Overflows



```
void printName(char* buffer) {
    char name[16];
    strcpy(name, buffer);
    printf("Hello %s\n", name);
}

int main(int argc, char* argv[]) {
    if(argc > 1) printName(argv[1]);
    return 0;
}
```



```
void printName(char* buffer) {  
    char name[16];  
    strcpy(name, buffer);  
    printf("Hello %s\n", name);  
}
```

```
int main(int argc, char* argv[]) {  
    if(argc > 1) printName(argv[1]);  
    return 0;  
}
```

👁 Observation 1: Buffer overflows are mostly linear



```
void printName(char* buffer) {
    char name[16];
    strcpy(name, buffer);
    printf("Hello %s\n", name);
}

int main(int argc, char* argv[]) {
    if(argc > 1) printName(argv[1]);
    return 0;
}
```

- 👁 Observation 1: Buffer overflows are mostly linear
- Cannot hit arbitrary memory, unlike format string vulnerabilities



```
void printName(char* buffer) {
    char name[16];
    strcpy(name, buffer);
    printf("Hello %s\n", name);
}

int main(int argc, char* argv[]) {
    if(argc > 1) printName(argv[1]);
    return 0;
}
```

- 👁 Observation 1: Buffer overflows are mostly linear
 - Cannot hit arbitrary memory, unlike format string vulnerabilities
- 👁 Observation 2: Attackers typically overwrite code pointer (return address)



```
void printName(char* buffer) {
    char name[16];
    strcpy(name, buffer);
    printf("Hello %s\n", name);
}

int main(int argc, char* argv[]) {
    if(argc > 1) printName(argv[1]);
    return 0;
}
```

- 👁 Observation 1: Buffer overflows are mostly linear
 - Cannot hit arbitrary memory, unlike format string vulnerabilities
- 👁 Observation 2: Attackers typically overwrite code pointer (return address)
- ❓ How can we detect linear buffer overflows?





- If the mine canary is dead, get out immediately



💡 Idea: Introduce a canary on the stack that signals a hazard





- 💡 Idea: Introduce a canary on the stack that signals a hazard
 - Hazard = corrupted return address



- 💡 Idea: Introduce a canary on the stack that signals a hazard
 - Hazard = corrupted return address
- ⚙️ Implementation



- 💡 Idea: Introduce a canary on the stack that signals a hazard
 - Hazard = corrupted return address
- ⚙️ Implementation
 - Simple compiler extension



💡 Idea: Introduce a canary on the stack that signals a hazard

- Hazard = corrupted return address

⚙️ Implementation

- Simple compiler extension
- Function prologue: push a random value (the canary), after the return address



💡 Idea: Introduce a canary on the stack that signals a hazard

- Hazard = corrupted return address

⚙️ Implementation

- Simple compiler extension
- Function prologue: push a random value (the canary), after the return address
- Linear buffer overflow can only overwrite return address when also overwriting canary

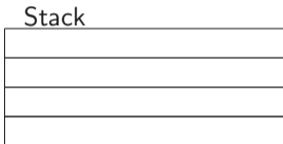


💡 Idea: Introduce a canary on the stack that signals a hazard

- Hazard = corrupted return address

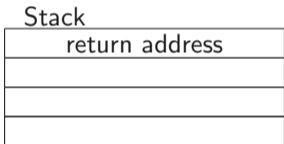
⚙️ Implementation

- Simple compiler extension
- Function prologue: push a random value (the canary), after the return address
- Linear buffer overflow can only overwrite return address when also overwriting canary
- Function epilogue: check if canary is valid (unmodified) before doing `retq`



```
<func >:  
  Setup stack frame  
  PUSH canary  
  [...]   
  ; check canary  
  RET
```

```
<main >:  
  [...]   
  CALL func  
  [...]
```



```
<func >:
```

```
  Setup stack frame
```

```
  PUSH canary
```

```
  [...]
```

```
  ; check canary
```

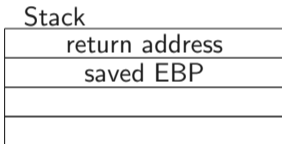
```
  RET
```

```
<main >:
```

```
  [...]
```

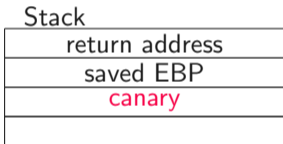
```
  CALL func
```

```
  [...]
```

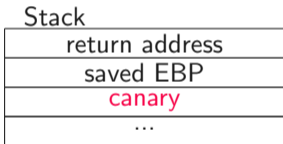


```
<func >:  
  Setup stack frame  
  PUSH canary  
  [...]  
  ; check canary  
  RET
```

```
<main >:  
  [...]  
  CALL func  
  [...]
```

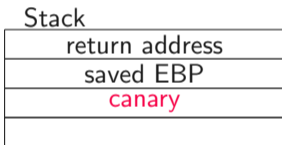


```
<func >:  
  Setup stack frame  
  PUSH canary  
  [...]   
  ; check canary  
  RET  
  
<main >:  
  [...]   
  CALL func  
  [...]
```

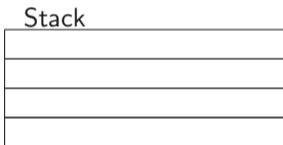
```
<func >:  
  Setup stack frame  
  PUSH canary  
  [...]  
  ; check canary  
  RET
```

```
<main >:  
  [...]  
  CALL func  
  [...]
```



```
<func >:  
  Setup stack frame  
  PUSH canary  
  [...]  
  ; check canary  
  RET
```

```
<main >:  
  [...]  
  CALL func  
  [...]
```



```
<func >:  
  Setup stack frame  
  PUSH canary  
  [...]   
  ; check canary  
  RET
```

```
<main >:  
  [...]   
  CALL func  
  [...] 
```



```
#include <stdio.h>
#include <string.h>

void printName(char* buffer) {
    char name[16];
    strcpy(name, buffer);
    printf("Hello %s\n", name);
}

int main(int argc, char* argv[]) {
    if(argc > 1) printName(argv[1]);
    return 0;
}
```



```
% gcc -o stack -fno-stack-protector stack.c
```



```
% gcc -o stack -fno-stack-protector stack.c
% ./stack AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Hello AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```



```
% gcc -o stack -fno-stack-protector stack.c
% ./stack AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Hello AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[1] 12345 segmentation fault (core dumped) ./stack
```



```
% gcc -o stack -fno-stack-protector stack.c
% ./stack AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Hello AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[1] 12345 segmentation fault (core dumped) ./stack
```

```
% gcc -o stack -fstack-protector stack.c
```




```
% gcc -o stack -fno-stack-protector stack.c
% ./stack AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Hello AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[1] 12345 segmentation fault (core dumped) ./stack
```

```
% gcc -o stack -fstack-protector stack.c
% ./stack AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Hello AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```



```
% gcc -o stack -fno-stack-protector stack.c
% ./stack AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Hello AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[1] 12345 segmentation fault (core dumped) ./stack
```

```
% gcc -o stack -fstack-protector stack.c
% ./stack AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Hello AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
*** stack smashing detected ***: ./stack terminated
```



```
% objdump -d stack
```



```
% objdump -d stack
0000000004005d6 <printName>:
  // function prologue
  ...
4005e2: mov     %fs:0x28,%rax    // load canary value
4005eb: mov     %rax,-0x8(%rbp) // store canary on stack
4005ef: xor     %eax,%eax
  ...
  // function epilogue
40061b: mov     -0x8(%rbp),%rax // load canary from stack
40061f: xor     %fs:0x28,%rax   // compare
400628: je     40062f <printName+0x59>
40062a: callq  4004a0 <__stack_chk_fail@plt>
40062f: leaveq
400630: retq
```



★ Properties





★ Properties

- Detects linear stack buffer overflows corrupting return address





★ Properties

- Detects linear stack buffer overflows corrupting return address
- Does not detect





★ Properties

- Detects linear stack buffer overflows corrupting return address
- Does not detect
 - overflowing one buffer into the other





★ Properties

- Detects linear stack buffer overflows corrupting return address
- Does not detect
 - overflowing one buffer into the other
 - overflowing co-located variables





★ Properties

- Detects linear stack buffer overflows corrupting return address
- Does not detect
 - overflowing one buffer into the other
 - overflowing co-located variables
 - arbitrary write access, e.g. `buffer[input] = input2;`





★ Properties

- Detects linear stack buffer overflows corrupting return address
- Does not detect
 - overflowing one buffer into the other
 - overflowing co-located variables
 - arbitrary write access, e.g. `buffer[input] = input2;`
 - format string vulnerabilities ...





★ Properties

- Detects linear stack buffer overflows corrupting return address
- Does not detect
 - overflowing one buffer into the other
 - overflowing co-located variables
 - arbitrary write access, e.g. `buffer[input] = input2;`
 - format string vulnerabilities ...
- Probabilistic defense





★ Properties



- Detects linear stack buffer overflows corrupting return address
- Does not detect
 - overflowing one buffer into the other
 - overflowing co-located variables
 - arbitrary write access, e.g. `buffer[input] = input2;`
 - format string vulnerabilities ...
- Probabilistic defense
 - Ineffective if attacker can guess or leak the canary value



★ Properties

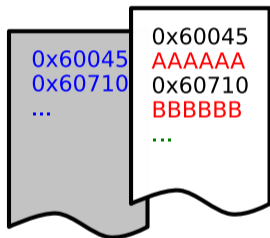


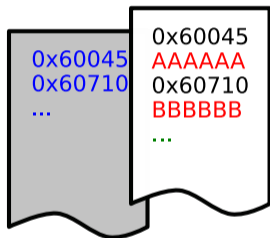
- Detects linear stack buffer overflows corrupting return address
- Does not detect
 - overflowing one buffer into the other
 - overflowing co-located variables
 - arbitrary write access, e.g. `buffer[input] = input2;`
 - format string vulnerabilities ...
- Probabilistic defense
 - Ineffective if attacker can guess or leak the canary value
- Animal welfare compatible (no bird has to die ;)



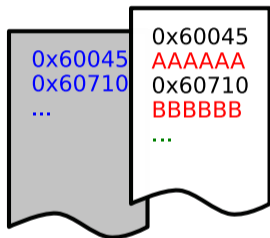
★ Properties

- Detects linear stack buffer overflows corrupting return address
 - Does not detect
 - overflowing one buffer into the other
 - overflowing co-located variables
 - arbitrary write access, e.g. `buffer[input] = input2;`
 - format string vulnerabilities ...
 - Probabilistic defense
 - Ineffective if attacker can guess or leak the canary value
 - Animal welfare compatible (no bird has to die ;)
- ❓ If we push/pop a canary for each return address, why not just duplicate return addresses instead?

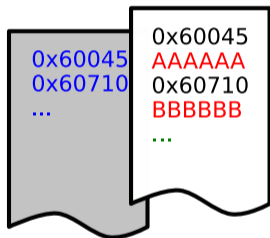




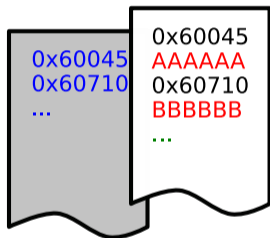
💡 Idea: duplicate return address on separate stack



- 💡 Idea: duplicate return address on separate stack
- ⚙️ Implementation: compiler extension similar to canaries



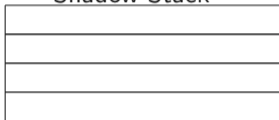
- 💡 Idea: duplicate return address on separate stack
- ⚙️ Implementation: compiler extension similar to canaries
 - Prologue: push return address also on shadow stack



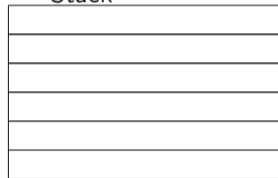
- 💡 Idea: duplicate return address on separate stack
- ⚙️ Implementation: compiler extension similar to canaries
 - Prologue: push return address also on shadow stack
 - Epilogue: verify return address before doing `retq`

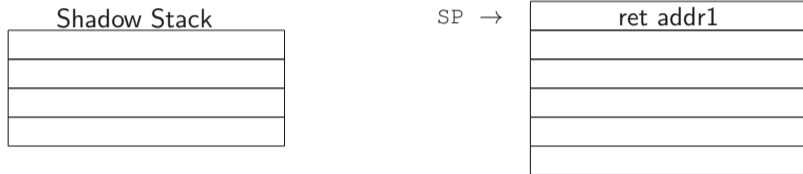


Shadow Stack

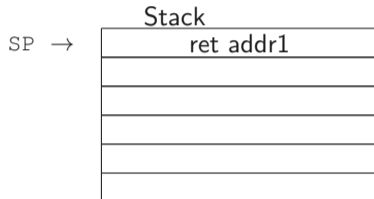
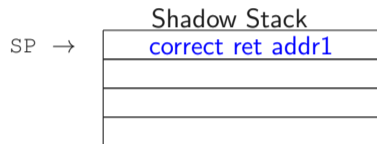


Stack

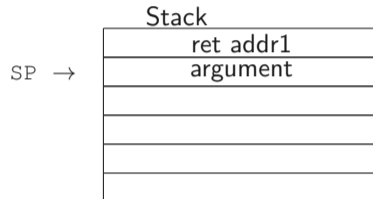
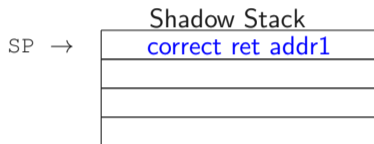




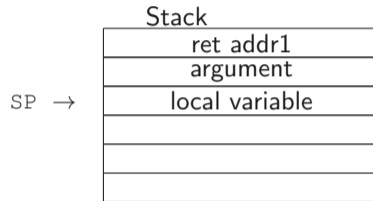
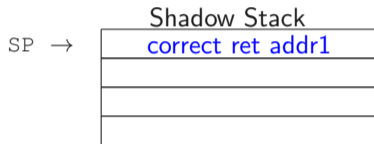
- Shadow stack duplicates all return addresses



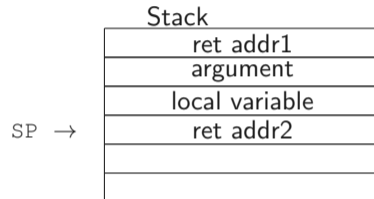
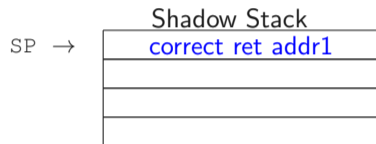
- Shadow stack duplicates all return addresses



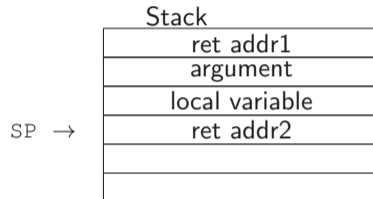
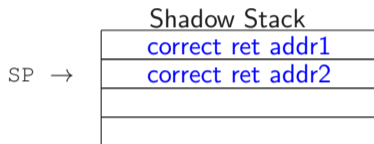
- Shadow stack duplicates all return addresses



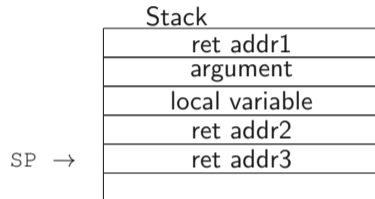
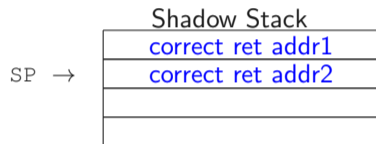
- Shadow stack duplicates all return addresses



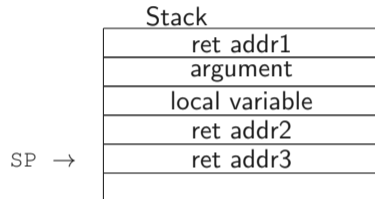
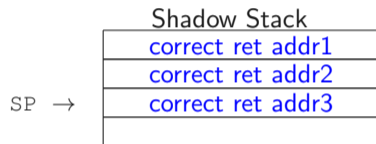
- Shadow stack duplicates all return addresses



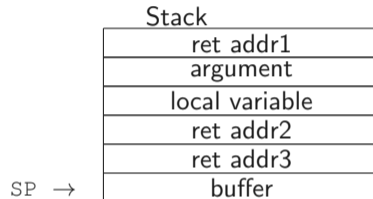
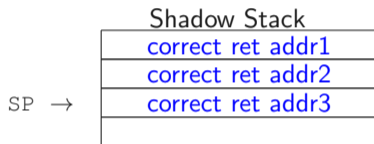
- Shadow stack duplicates all return addresses



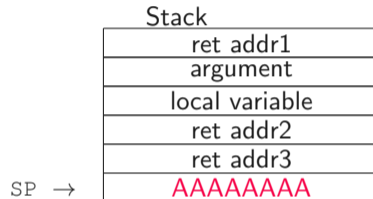
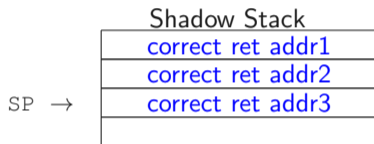
- Shadow stack duplicates all return addresses



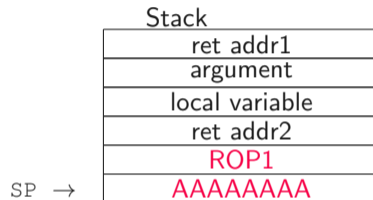
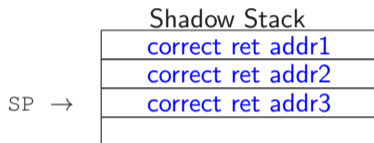
- Shadow stack duplicates all return addresses



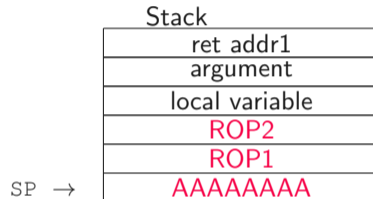
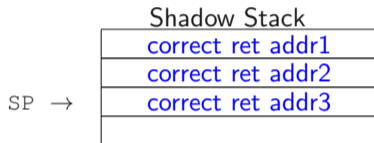
- Shadow stack duplicates all return addresses
- Attacker injects ROP chain



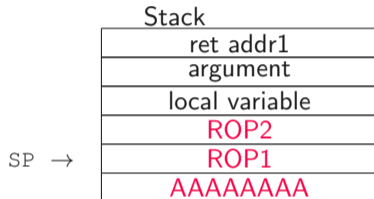
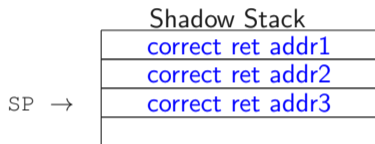
- Shadow stack duplicates all return addresses
- Attacker injects ROP chain



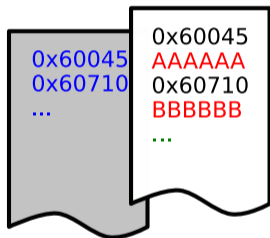
- Shadow stack duplicates all return addresses
- Attacker injects ROP chain



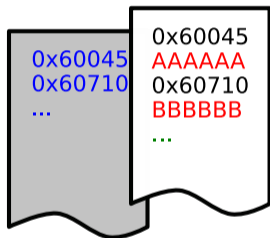
- Shadow stack duplicates all return addresses
- Attacker injects ROP chain



- Shadow stack duplicates all return addresses
- Attacker injects ROP chain
- Program crashes because of shadow stack mismatch

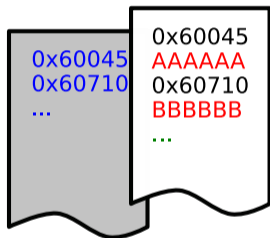


★ Properties



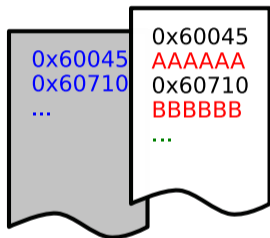
★ Properties

- Detect if buffer overflow corrupts return address



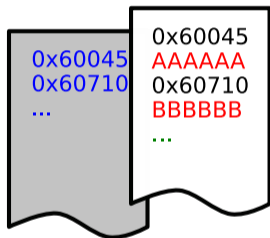
★ Properties

- Detect if buffer overflow corrupts return address
- Does not prevent



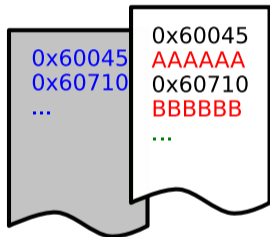
★ Properties

- Detect if buffer overflow corrupts return address
- Does not prevent
 - attacks not affecting return addresses



★ Properties

- Detect if buffer overflow corrupts return address
- Does not prevent
 - attacks not affecting return addresses
 - attacks on shadow stack



★ Properties

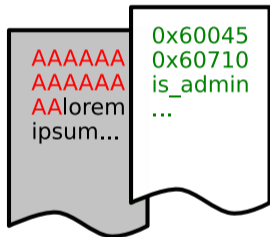
- Detect if buffer overflow corrupts return address
- Does not prevent
 - attacks not affecting return addresses
 - attacks on shadow stack

❓ Why duplicate at all if we assume shadow stack is secure?

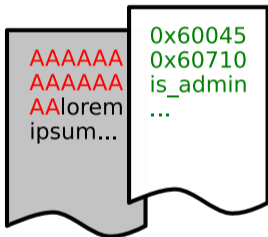




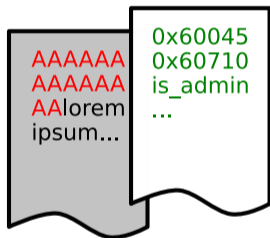
💡 Inverse idea: store unsafe buffers on separate stack



- 💡 Inverse idea: store unsafe buffers on separate stack
 - Safe stack only contains **return addresses** and **sensible variables** that cannot overflow



- 💡 Inverse idea: store unsafe buffers on separate stack
 - Safe stack only contains **return addresses** and **sensible variables** that cannot overflow
- ⚙️ Implementation: compiler extension similar to shadow stack

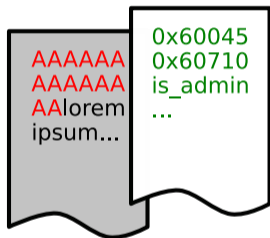


★ Properties



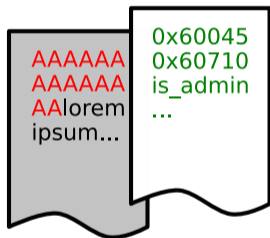
★ Properties

- Buffer overflow cannot corrupt return addresses / safe variables



★ Properties

- Buffer overflow cannot corrupt return addresses / safe variables
- Does not prevent



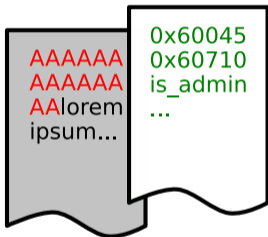
★ Properties

- Buffer overflow cannot corrupt return addresses / safe variables
- Does not prevent
 - overflowing one unsafe buffer into the other



★ Properties

- Buffer overflow cannot corrupt return addresses / safe variables
- Does not prevent
 - overflowing one unsafe buffer into the other
 - format string vulnerabilities ...



★ Properties

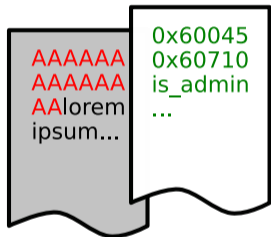
- Buffer overflow cannot corrupt return addresses / safe variables
- Does not prevent
 - overflowing one unsafe buffer into the other
 - format string vulnerabilities ...
 - attacks on safe stack



★ Properties

- Buffer overflow cannot corrupt return addresses / safe variables
- Does not prevent
 - overflowing one unsafe buffer into the other
 - format string vulnerabilities ...
 - attacks on safe stack

❓ Why not protect shadow/safe stack in hardware?



★ Properties

- Buffer overflow cannot corrupt return addresses / safe variables
- Does not prevent
 - overflowing one unsafe buffer into the other
 - format string vulnerabilities ...
 - attacks on safe stack

❓ Why not protect shadow/safe stack in hardware?

- Control-Flow Enforcement Technology (CET) for Intel (and AMD)

Code Injection Attacks



- Exploit buffer overflow



- Exploit buffer overflow
- Inject custom code that spawns a shell → Shellcode



- Exploit buffer overflow
- Inject custom code that spawns a shell → Shellcode
- Corrupt code pointer to execute shellcode



Data Execution Prevention



Data Execution Prevention (DEP)





Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)





Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

- 👁 Observation: Von-Neumann CPUs mix code and data memory
 - This allows code injection into data memory





Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

- 👁 Observation: Von-Neumann CPUs mix code and data memory
 - This allows code injection into data memory
- 💡 Idea 1: make data memory non-executable



Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

- 👁 Observation: Von-Neumann CPUs mix code and data memory
 - This allows code injection into data memory
- 💡 Idea 1: make data memory non-executable
- 💡 Idea 2: make code memory non-writable



Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

- 👁 Observation: Von-Neumann CPUs mix code and data memory
 - This allows code injection into data memory
- 💡 Idea 1: make data memory non-executable
- 💡 Idea 2: make code memory non-writable
- ⚙ Implementation



Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

👁 Observation: Von-Neumann CPUs mix code and data memory

- This allows code injection into data memory

💡 Idea 1: make data memory non-executable

💡 Idea 2: make code memory non-writable

⚙ Implementation

- Set writable memory to non-executable, e.g.: stack, heap, data, ...



Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

👁 Observation: Von-Neumann CPUs mix code and data memory

- This allows code injection into data memory

💡 Idea 1: make data memory non-executable

💡 Idea 2: make code memory non-writable

⚙ Implementation

- Set writable memory to non-executable, e.g.: stack, heap, data, ...
- Usually done by the program loader (using `mmap`, `mprotect`)



Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

👁 Observation: Von-Neumann CPUs mix code and data memory

- This allows code injection into data memory

💡 Idea 1: make data memory non-executable

💡 Idea 2: make code memory non-writable

⚙ Implementation

- Set writable memory to non-executable, e.g.: stack, heap, data, ...
- Usually done by the program loader (using `mmap`, `mprotect`)
- Hardware support in the page tables



Data Execution Prevention (DEP) \approx Write-Xor-Execute ($W\oplus X$)

👁 Observation: Von-Neumann CPUs mix code and data memory

- This allows code injection into data memory

💡 Idea 1: make data memory non-executable

💡 Idea 2: make code memory non-writable

⚙ Implementation

- Set writable memory to non-executable, e.g.: stack, heap, data, ...
- Usually done by the program loader (using `mmap`, `mprotect`)
- Hardware support in the page tables
 - Intel: XD-bit, AMD: NX-bit, ARM: XN-bit



```
#include <stdio.h>
#include <string.h>

char code[] = "\x31\xc0\x48\xbb\xd1\x9d\x96\x91\xd0\x8c\x97\xff\x48\xf7\xdb\x53\x54\x5f\x99\x52\x57\x54\x5e\xb0\x3b\x0f\x05";

int main()
{
    printf("len:%d bytes\n", strlen(code));
    (*(void(*)()) code)();
    return 0;
}
```



```
% gdb ./shellcode
```



```
% gdb ./shellcode  
(gdb) run
```



```
% gdb ./shellcode
(gdb) run
Starting program: /home/shellcode
len:27 bytes

Program received signal SIGSEGV, Segmentation fault.
0x0000000000601040 in code ()
```



```
% gdb ./shellcode
(gdb) run
Starting program: /home/shellcode
len:27 bytes

Program received signal SIGSEGV, Segmentation fault.
0x000000000601040 in code ()
```

```
% execstack -s ./shellcode
```




```
% gdb ./shellcode
(gdb) run
Starting program: /home/shellcode
len:27 bytes

Program received signal SIGSEGV, Segmentation fault.
0x000000000601040 in code ()
```

```
% execstack -s ./shellcode
% gdb ./shellcode
```



```
% gdb ./shellcode
(gdb) run
Starting program: /home/shellcode
len:27 bytes

Program received signal SIGSEGV, Segmentation fault.
0x000000000601040 in code ()
```

```
% execstack -s ./shellcode
% gdb ./shellcode
(gdb) run
```



```
% gdb ./shellcode
(gdb) run
Starting program: /home/shellcode
len:27 bytes

Program received signal SIGSEGV, Segmentation fault.
0x000000000601040 in code ()
```

```
% execstack -s ./shellcode
% gdb ./shellcode
(gdb) run
Starting program: /home/shellcode
len:27 bytes
process 9494 is executing new program: /bin/dash
$
```



★ Properties of DEP/W⊕X



★ Properties of DEP/W⊕X

- Prevents code injection attacks



★ Properties of DEP/W⊕X

- Prevents code injection attacks
- Hardly any runtime overhead



★ Properties of DEP/W⊕X

- Prevents code injection attacks
- Hardly any runtime overhead
- Requires hardware support



★ Properties of DEP/W⊕X

- Prevents code injection attacks
- Hardly any runtime overhead
- Requires hardware support
- Does not prevent code reuse attacks (since no code is injected)



★ Properties of DEP/W⊕X

- Prevents code injection attacks
 - Hardly any runtime overhead
 - Requires hardware support
 - Does not prevent code reuse attacks (since no code is injected)
- ❓ How to protect just-in-time (JIT) compiled code? JIT compiler needs to modify code at runtime ...

Code Reuse Attacks





- 👁 Observation: Many exploits need knowledge of addresses (ROP, ret2libc ...)



- 👁 Observation: Many exploits need knowledge of addresses (ROP, ret2libc ...)
- 💡 Idea: randomize program to make exploit development (much) more difficult



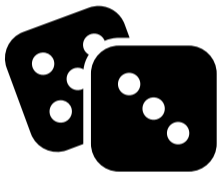
- 👁 Observation: Many exploits need knowledge of addresses (ROP, ret2libc ...)
- 💡 Idea: randomize program to make exploit development (much) more difficult
- ★ General properties



- 👁 Observation: Many exploits need knowledge of addresses (ROP, ret2libc ...)
- 💡 Idea: randomize program to make exploit development (much) more difficult
- ★ General properties
 - Probabilistic defense



- 👁 Observation: Many exploits need knowledge of addresses (ROP, ret2libc ...)
- 💡 Idea: randomize program to make exploit development (much) more difficult
- ★ General properties
 - Probabilistic defense
 - Can be broken by information leakage (e.g., via side channels)



- 👁 Observation: Many exploits need knowledge of addresses (ROP, ret2libc ...)
- 💡 Idea: randomize program to make exploit development (much) more difficult
- ★ General properties
 - Probabilistic defense
 - Can be broken by information leakage (e.g., via side channels)
 - ❓ How big is the entropy?





💡 Randomize the memory layout



- 💡 Randomize the memory layout
 - Attacker cannot guess location of libc, stack, heap ...



- 💡 Randomize the memory layout
 - Attacker cannot guess location of libc, stack, heap ...
- ⚙️ Implementation



- 💡 Randomize the memory layout
 - Attacker cannot guess location of libc, stack, heap ...
- ⚙️ Implementation
 - At program startup move various segments to a random position



💡 Randomize the memory layout

- Attacker cannot guess location of libc, stack, heap ...

⚙️ Implementation

- At program startup move various segments to a random position
 - Stack, heap, shared memory



💡 Randomize the memory layout

- Attacker cannot guess location of libc, stack, heap ...

⚙️ Implementation

- At program startup move various segments to a random position
 - Stack, heap, shared memory
 - Shared libraries



💡 Randomize the memory layout

- Attacker cannot guess location of libc, stack, heap ...

⚙️ Implementation

- At program startup move various segments to a random position
 - Stack, heap, shared memory
 - Shared libraries
 - Main executable (optional)



💡 Randomize the memory layout

- Attacker cannot guess location of libc, stack, heap ...

⚙️ Implementation

- At program startup move various segments to a random position
 - Stack, heap, shared memory
 - Shared libraries
 - Main executable (optional)
- Randomization done by operating system (e.g., on `mmap`)



💡 Randomize the memory layout

- Attacker cannot guess location of libc, stack, heap ...

⚙️ Implementation

- At program startup move various segments to a random position
 - Stack, heap, shared memory
 - Shared libraries
 - Main executable (optional)
- Randomization done by operating system (e.g., on `mmap`)
 - Linux `/proc/sys/kernel/randomize_va_space`



```
#include <stdio.h>
#include <stdlib.h>

int main() {
    int x;
    printf("Stack: %p\n", &x);
    printf("Heap: %p\n", malloc(10));
    return 0;
}
```



```
#include <stdio.h>
#include <stdlib.h>

int main() {
    int x;
    printf("Stack: %p\n", &x);
    printf("Heap: %p\n", malloc(10));
    return 0;
}
```

```
% ./aslr
Stack: 0x7ffcc2666e74
Heap: 0x1dd9420
```



```
#include <stdio.h>
#include <stdlib.h>

int main() {
    int x;
    printf("Stack: %p\n", &x);
    printf("Heap: %p\n", malloc(10));
    return 0;
}
```

```
% ./aslr
Stack: 0x7ffcc2666e74
Heap: 0x1dd9420
% ./aslr
Stack: 0x7ffcbf0c1ae4
Heap: 0x124b420
```



```
% cat /proc/self/maps
```



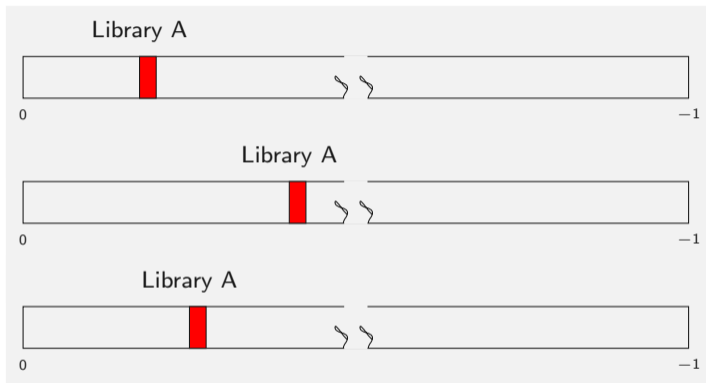
```
% cat /proc/self/maps
00400000-0040c000 r-xp 00000000 fd:00 395191      /bin/cat
0060b000-0060c000 r--p 0000b000 fd:00 395191      /bin/cat
0060c000-0060d000 rw-p 0000c000 fd:00 395191      /bin/cat
00b0c000-00b2d000 rw-p 00000000 00:00 0        [heap]
7efcbb558000-7efcbb87e000 r--p 00000000 fd:00 11534857 /usr/lib/locale/locale-archive
7efcbb87e000-7efcbb87e000 r-xp 00000000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7efcbb87e000-7efcbb87e000 r--p 001c0000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7efcbb87e000-7efcbb87e000 r--p 001c0000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7efcbb87e000-7efcbb87e000 rw-p 001c4000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7efcbb87e000-7efcbb87e000 rw-p 00000000 00:00 0
7efcbb87e000-7efcbb87e000 r-xp 00000000 fd:00 4588089 /lib/x86_64-linux-gnu/ld-2.23.so
7efcbb87e000-7efcbb87e000 rw-p 00000000 00:00 0
7efcbb87e000-7efcbb87e000 rw-p 00000000 00:00 0
7efcbb87e000-7efcbb87e000 r--p 00025000 fd:00 4588089 /lib/x86_64-linux-gnu/ld-2.23.so
7efcbb87e000-7efcbb87e000 rw-p 00026000 fd:00 4588089 /lib/x86_64-linux-gnu/ld-2.23.so
7efcbb87e000-7efcbb87e000 rw-p 00000000 00:00 0
7ffff84c6000-7ffff84e7000 rw-p 00000000 00:00 0        [stack]
7ffff8536000-7ffff8538000 r--p 00000000 00:00 0        [vvar]
7ffff8538000-7ffff853a000 r-xp 00000000 00:00 0        [vdso]
fffffffff60000-fffffffff601000 r-xp 00000000 00:00 0 [vsyscall]
```




```
% cat /proc/self/maps
00400000-0040c000 r-xp 00000000 fd:00 395191      /bin/cat
0060b000-0060c000 r--p 0000b000 fd:00 395191      /bin/cat
0060c000-0060d000 rw-p 0000c000 fd:00 395191      /bin/cat
00799000-007ba000 rw-p 00000000 00:00 0        [heap]
7fec1f08d000-7fec1f3b3000 r--p 00000000 fd:00 11534857 /usr/lib/locale/locale-archive
7fec1f3b3000-7fec1f573000 r-xp 00000000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7fec1f573000-7fec1f773000 --p 001c0000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7fec1f773000-7fec1f777000 r--p 001c0000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7fec1f777000-7fec1f779000 rw-p 001c4000 fd:00 4587769 /lib/x86_64-linux-gnu/libc-2.23.so
7fec1f779000-7fec1f77d000 rw-p 00000000 00:00 0
7fec1f77d000-7fec1f7a3000 r-xp 00000000 fd:00 4588089 /lib/x86_64-linux-gnu/ld-2.23.so
7fec1f96d000-7fec1f970000 rw-p 00000000 00:00 0
7fec1f980000-7fec1f9a2000 rw-p 00000000 00:00 0
7fec1f9a2000-7fec1f9a3000 r--p 00025000 fd:00 4588089 /lib/x86_64-linux-gnu/ld-2.23.so
7fec1f9a3000-7fec1f9a4000 rw-p 00026000 fd:00 4588089 /lib/x86_64-linux-gnu/ld-2.23.so
7fec1f9a4000-7fec1f9a5000 rw-p 00000000 00:00 0
7ffeffa30000-7ffeffa51000 rw-p 00000000 00:00 0        [stack]
7ffeffa7f000-7ffeffa81000 r--p 00000000 00:00 0        [vvar]
7ffeffa81000-7ffeffa83000 r-xp 00000000 00:00 0        [vdso]
fffffffff60000-fffffffff601000 r-xp 00000000 00:00 0 [vsyscall]
```



- Same library code is randomized to **different addresses** at each program start



- Same library code is randomized to **different addresses** at each program start
- ❓ How does randomized code remain functional?



- Within a module



- Within a module
 - Compiler replaces absolute addresses with (rip-)relative addresses
 - Code can be executed from virtually any offset
 - Shared libraries: compile flags `-fpic`, `-fPIC`
 - Executable: compile flags `-fpie`, `-fPIE`



- Within a module
 - Compiler replaces absolute addresses with (rip-)relative addresses
 - Code can be executed from virtually any offset
 - Shared libraries: compile flags `-fpic`, `-fPIC`
 - Executable: compile flags `-fpie`, `-fPIE`
- Across modules



- Within a module
 - Compiler replaces absolute addresses with (rip-)relative addresses
 - Code can be executed from virtually any offset
 - Shared libraries: compile flags `-fpic`, `-fPIC`
 - Executable: compile flags `-fpie`, `-fPIE`
- Across modules
 - Runtime linker resolves addresses via:
 - Global Offset Table (GOT) for arbitrary addresses
 - Procedure Linkage Table (PLT) for function calls



```
#include <stdio.h>
```

```
int main() {  
    printf("Hi\n");  
}
```




```
#include <stdio.h>
```

```
int main() {  
    printf("Hi\n");  
}
```

```
% gcc -o main -static main.c
```



```
#include <stdio.h>
```

```
int main() {  
    printf("Hi\n");  
}
```

```
% gcc -o main -static main.c
```

```
% objdump -d main
```



```
#include <stdio.h>
```

```
int main() {  
    printf("Hi\n");  
}
```

```
% gcc -o main -static main.c
```

```
% objdump -d main
```

```
00000000004009ae <main>:
```

```
4009ae: 55                push   %rbp  
4009af: 48 89 e5          mov    %rsp,%rbp  
4009b2: bf a4 11 4a 00    mov    $0x4a11a4,%edi    # Address of Hi  
4009b7: e8 24 f2 00 00    callq 40fbe0 <_IO_puts> # Direct call  
4009bc: b8 00 00 00 00    mov    $0x0,%eax  
4009c1: 5d                pop    %rbp  
4009c2: c3                retq
```



```
#include <stdio.h>
```

```
int main() {  
    printf("Hi\n");  
}
```

```
% gcc -o main -fPIE main.c
```



```
#include <stdio.h>
```

```
int main() {  
    printf("Hi\n");  
}
```

```
% gcc -o main -fPIE main.c
```

```
% objdump -d main
```



```
#include <stdio.h>
```

```
int main() {  
    printf("Hi\n");  
}
```

```
% gcc -o main -fPIE main.c
```

```
% objdump -d main
```

```
00000000000000526 <main>:
```

```
   526: 55                push   %rbp  
   527: 48 89 e5          mov    %rsp,%rbp  
   52a: 48 8d 3d 93 00 00 00 lea   0x93(%rip),%rdi  # Address of Hi  
   531: e8 ca fe ff ff   callq 400 <puts@plt> # PLT  
   536: b8 00 00 00 00   mov    $0x0,%eax  
   53b: 5d                pop    %rbp  
   53c: c3                retq
```



★ Properties





★ Properties

- Cheap



★ Properties

- Cheap
- Make exploit development harder



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module
- Requires operating system support



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module
- Requires operating system support
- ASLR on code requires position independence



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module
- Requires operating system support
- ASLR on code requires position independence
- Quite limited entropy on 32-bit systems



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module
- Requires operating system support
- ASLR on code requires position independence
- Quite limited entropy on 32-bit systems → brute-force



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module
- Requires operating system support
- ASLR on code requires position independence
- Quite limited entropy on 32-bit systems → brute-force
- Information leak breaks ASLR



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module
- Requires operating system support
- ASLR on code requires position independence
- Quite limited entropy on 32-bit systems → brute-force
- Information leak breaks ASLR

❓ How to protect ASLR against information leakage?



★ Properties

- Cheap
- Make exploit development harder
- Does not prevent exploitation within a module
- Requires operating system support
- ASLR on code requires position independence
- Quite limited entropy on 32-bit systems → brute-force
- Information leak breaks ASLR

❓ How to protect ASLR against information leakage?

- Execute-only memory (non-readable)



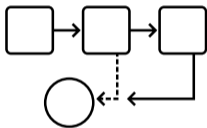
Change the randomization of the code segment

- You should generate **two binaries** with **ASLR enabled**
- One binary should have **randomization** for stack, heap, and code
- The other binary should only have **randomization** for stack and heap, but **not for code**
- Both binaries must run for at **least 5 seconds** (e.g., `sleep(5)`; before return) but **not** longer than **10 seconds**
- Upload your binaries at `https://challenges.sasectf.student.iaik.tugraz.at/aslr/index.php`
- If it is correct, you will get the flag
- Test system is Debian 11.8, kernel 5.10.0-26



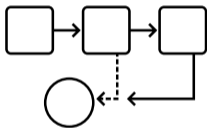


👁 Observation: most attacks corrupt code pointers



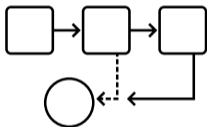


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers



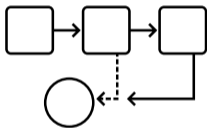


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)



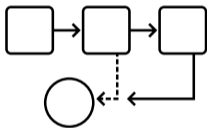


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow



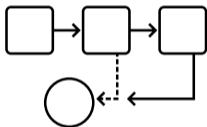


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...



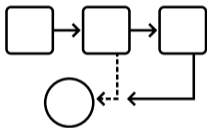


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...
- Code pointers everywhere



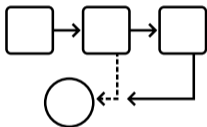


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...
- Code pointers everywhere
 - Return addresses



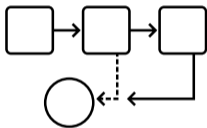


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...
- Code pointers everywhere
 - Return addresses
 - Function pointers



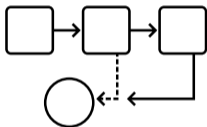


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...
- Code pointers everywhere
 - Return addresses
 - Function pointers
 - C++ vtables



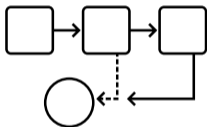


- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...
- Code pointers everywhere
 - Return addresses
 - Function pointers
 - C++ vtables
 - GOT entries



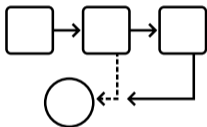


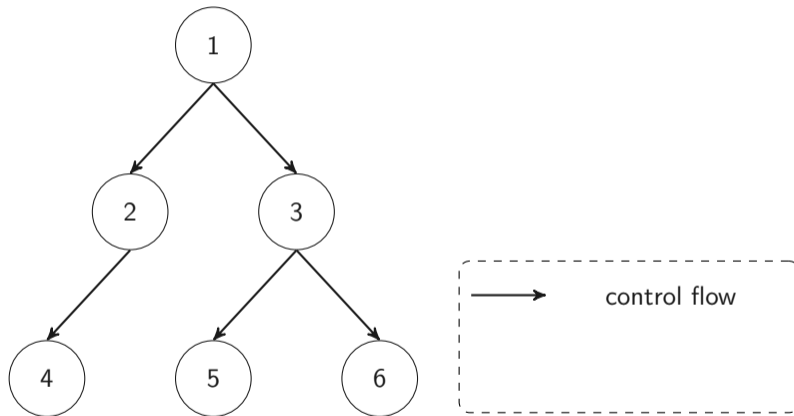
- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...
- Code pointers everywhere
 - Return addresses
 - Function pointers
 - C++ vtables
 - GOT entries
 - Signal handlers

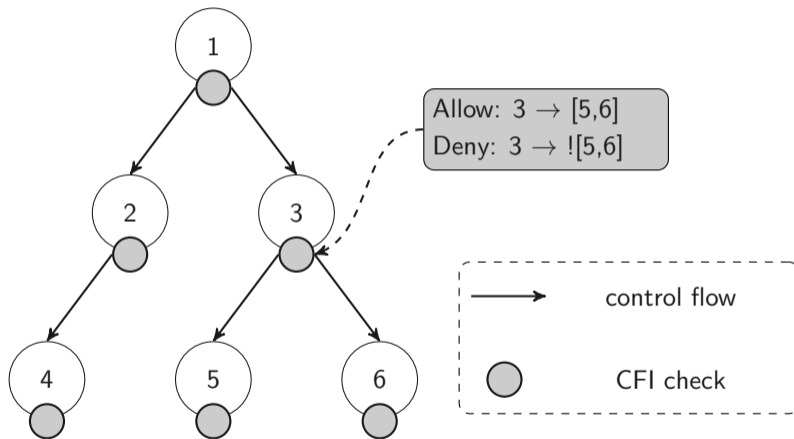


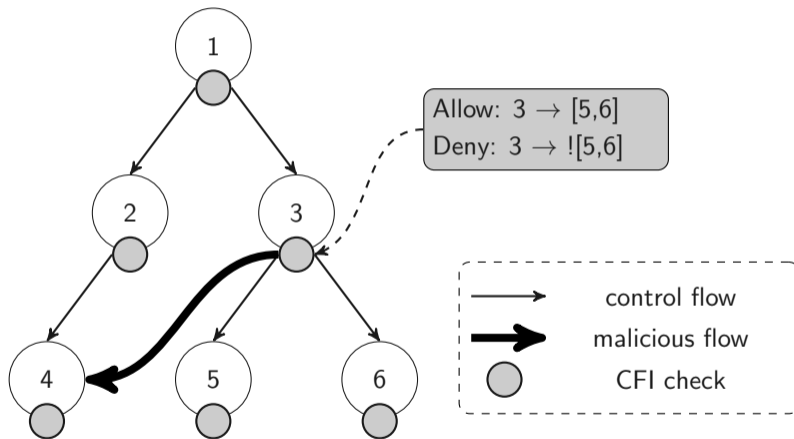


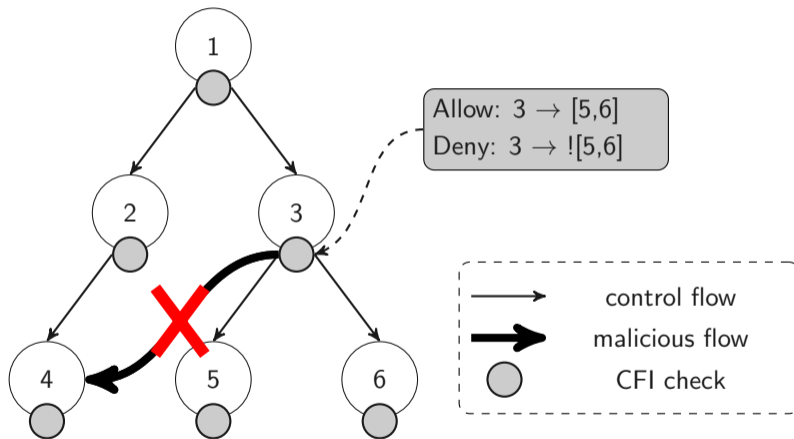
- 👁 Observation: most attacks corrupt code pointers
- 💡 Idea: specifically protect code pointers
 - CFI: Program must stay inside its approximated Control Flow Graph (CFG)
 - Attacker cannot (arbitrarily) change control flow
 - Prevent ret2libc, ROP, JOP, ...
- Code pointers everywhere
 - Return addresses
 - Function pointers
 - C++ vtables
 - GOT entries
 - Signal handlers
- We need to protect all of them!













- CFI on backward edges



- CFI on backward edges
 - Return addresses shall only point to the real caller



- CFI on backward edges
 - Return addresses shall only point to the real caller
 - Solved via shadow stack / safe stack



- CFI on backward edges
 - Return addresses shall only point to the real caller
 - Solved via shadow stack / safe stack
- CFI on forward edges



- CFI on backward edges
 - Return addresses shall only point to the real caller
 - Solved via shadow stack / safe stack
- CFI on forward edges
 - GOT entries shall only contain legitimate pointers



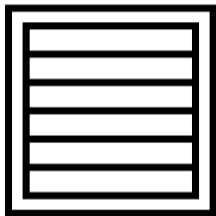
- CFI on backward edges
 - Return addresses shall only point to the real caller
 - Solved via shadow stack / safe stack
- CFI on forward edges
 - GOT entries shall only contain legitimate pointers
 - Function pointers shall only call legitimate functions



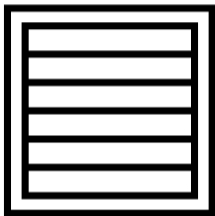
- CFI on backward edges
 - Return addresses shall only point to the real caller
 - Solved via shadow stack / safe stack
- CFI on forward edges
 - GOT entries shall only contain legitimate pointers
 - Function pointers shall only call legitimate functions
 - C++ vtable shall only call legitimate members



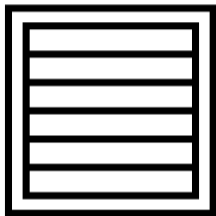
- CFI on backward edges
 - Return addresses shall only point to the real caller
 - Solved via shadow stack / safe stack
- CFI on forward edges
 - GOT entries shall only contain legitimate pointers
 - Function pointers shall only call legitimate functions
 - C++ vtable shall only call legitimate members
 - ❓ How to determine *legitimate* targets?



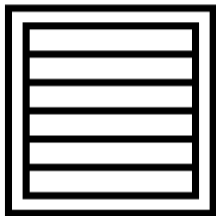
👁 Observation: only **runtime linker** populates GOT



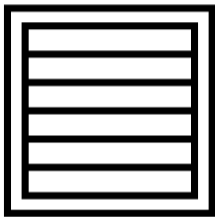
- 👁 Observation: only **runtime linker** populates GOT
 - Any other GOT manipulation is either an accident or an attack



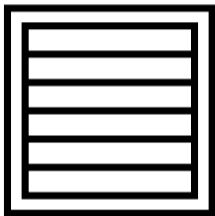
- 👁 Observation: only **runtime linker** populates GOT
 - Any other GOT manipulation is either an accident or an attack
- 💡 Idea: Make GOT **read-only**



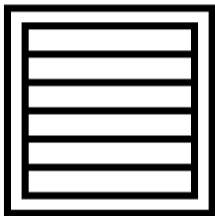
- 👁 Observation: only **runtime linker** populates GOT
 - Any other GOT manipulation is either an accident or an attack
- 💡 Idea: Make GOT **read-only**
- ⚙ Implementation



- 👁 Observation: only **runtime linker** populates GOT
 - Any other GOT manipulation is either an accident or an attack
- 💡 Idea: Make GOT **read-only**
- ⚙ Implementation
 - Linker populates *all* GOT entries at program start



- 👁 Observation: only **runtime linker** populates GOT
 - Any other GOT manipulation is either an accident or an attack
- 💡 Idea: Make GOT **read-only**
- ⚙ Implementation
 - Linker populates *all* GOT entries at program start → slowdown



- 👁 Observation: only **runtime linker** populates GOT
 - Any other GOT manipulation is either an accident or an attack
- 💡 Idea: Make GOT **read-only**
- ⚙ Implementation
 - Linker populates *all* GOT entries at program start → slowdown
 - Compiler flag `-Wl,-z,relro` "relocations read-only"



What are **legitimate targets** for an indirect function call?





What are **legitimate targets** for an indirect function call? Possible answers:





What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function





What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise





What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI





What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI
- A2: functions with the **same signature** as the function pointer





What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI
- A2: functions with the **same signature** as the function pointer
 - Better, prevents some type confusion attacks





What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI
- A2: functions with the **same signature** as the function pointer
 - Better, prevents some type confusion attacks
- A3: only functions which **could be assigned** to the function pointer



What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI
- A2: functions with the **same signature** as the function pointer
 - Better, prevents some type confusion attacks
- A3: only functions which **could be assigned** to the function pointer
 - Even better, still a bit imprecise



What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI
- A2: functions with the **same signature** as the function pointer
 - Better, prevents some type confusion attacks
- A3: only functions which **could be assigned** to the function pointer
 - Even better, still a bit imprecise
- A4: only the function which is **actually assigned** to the function pointer



What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI
- A2: functions with the **same signature** as the function pointer
 - Better, prevents some type confusion attacks
- A3: only functions which **could be assigned** to the function pointer
 - Even better, still a bit imprecise
- A4: only the function which is **actually assigned** to the function pointer
 - Fully precise, more expensive



What are **legitimate targets** for an indirect function call? Possible answers:

- A1: **any** function
 - Simple but imprecise
 - Attacker can invoke arbitrary functions without violating CFI
- A2: functions with the **same signature** as the function pointer
 - Better, prevents some type confusion attacks
- A3: only functions which **could be assigned** to the function pointer
 - Even better, still a bit imprecise
- A4: only the function which is **actually assigned** to the function pointer
 - Fully precise, more expensive
 - Code pointer integrity



Examples for Hardware/Software CFI mechanisms



Examples for Hardware/Software CFI mechanisms

- A1: **any** function is valid



Examples for Hardware/Software CFI mechanisms

- A1: **any** function is valid
 - Hardware CFI: Control-Flow Enforcement Technology (CET)



Examples for Hardware/Software CFI mechanisms

- A1: **any** function is valid
 - Hardware CFI: Control-Flow Enforcement Technology (CET)
- A2: functions with the **same signature** are valid



Examples for Hardware/Software CFI mechanisms

- A1: **any** function is valid
 - Hardware CFI: Control-Flow Enforcement Technology (CET)
- A2: functions with the **same signature** are valid
 - Software CFI: `clang -fsanitize=cfi`



Examples for Hardware/Software CFI mechanisms

- A1: **any** function is valid
 - Hardware CFI: Control-Flow Enforcement Technology (CET)
- A2: functions with the **same signature** are valid
 - Software CFI: `clang -fsanitize=cfi`
 - Hardware CFI: Arm Pointer Authentication, FinelBT



- Every function entry marked with `endbr64` instruction





- Every function entry marked with `endbr64` instruction
- Call instructions only succeed towards `endbr64` instructions





- Every function entry marked with `endbr64` instruction
- Call instructions only succeed towards `endbr64` instructions
- Supported by recent Intel and AMD CPUs





- Every function entry marked with `endbr64` instruction
- Call instructions only succeed towards `endbr64` instructions
- Supported by recent Intel and AMD CPUs



```
% objdump -d main
0000000000001149 <test>:
   1149:  f3 0f 1e fa  endbr64
   114d:  55           push   %rbp
   114e:  48 89 e5     mov    %rsp,%rbp
   ...

0000000000001160 <test2>:
   1160:  f3 0f 1e fa  endbr64
   1164:  55           push   %rbp
   1165:  48 89 e5     mov    %rsp,%rbp
   ...
```




```
#include <iostream>

class A {
public: virtual const char* name() { return "A"; };
};

class B {
public: const char* name() { return "B"; };
private: virtual const char* secret() { return "secret"; };
};

int main() {
    A* a = new A();
    std::cout << a->name() << std::endl;
    B* b = new B();
    std::cout << b->name() << std::endl;

    a = (A*)b; // type confusion vulnerability
    std::cout << a->name() << std::endl;
}
```



```
% ./test  
A
```



```
% ./test
```

```
A
```

```
B
```



```
% ./test  
A  
B  
secret
```



```
% ./test
```

```
A
```

```
B
```

```
secret
```

```
% clang++ -flto -fsanitize=cfi -fvisibility=hidden \  
-fno-sanitize-trap=all tc.cpp -o tc
```



```
% ./test
```

```
A
```

```
B
```

```
secret
```

```
% clang++ -flto -fsanitize=cfi -fvisibility=hidden \  
-fno-sanitize-trap=all tc.cpp -o tc
```

```
% ./tc
```

```
A
```



```
% ./test
```

```
A
```

```
B
```

```
secret
```

```
% clang++ -flto -fsanitize=cfi -fvisibility=hidden \  
-fno-sanitize-trap=all tc.cpp -o tc
```

```
% ./tc
```

```
A
```

```
B
```

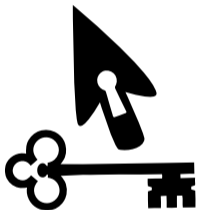


```
% ./test
A
B
secret
```

```
% clang++ -flto -fsanitize=cfi -fvisibility=hidden \
  -fno-sanitize-trap=all tc.cpp -o tc
% ./tc
A
B
tc.cpp:21:9: runtime error: control flow integrity check for
type 'A' failed during cast to unrelated type
(vtable address 0x00000042bd40)
0x00000042bd40: note: vtable is of type 'B'
00 00 00 00  d0 4b 42 00 00 00 00 00  01 1b 03 3b cc 11 00
                ^
```



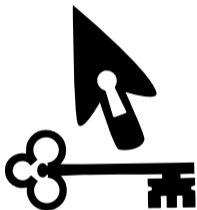

ARMv8.3 hardware-based pointer authentication





ARMv8.3 hardware-based pointer authentication

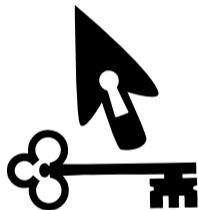
👁 Observation: 64-bit architectures do not use all bits





ARMv8.3 hardware-based pointer authentication

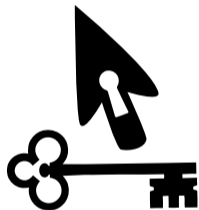
- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused





ARMv8.3 hardware-based pointer authentication

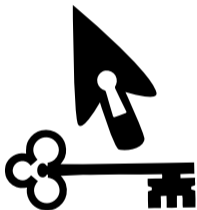
- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused
- 💡 Idea: repurpose bits for cryptographically authenticating pointers





ARMv8.3 hardware-based pointer authentication

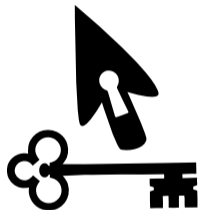
- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused
- 💡 Idea: repurpose bits for cryptographically authenticating pointers
- ⚙ Implementation





ARMv8.3 hardware-based pointer authentication

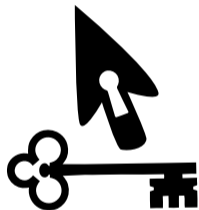
- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused
- 💡 Idea: repurpose bits for cryptographically authenticating pointers
- ⚙ Implementation
 - Compiler secures a pointer with special PAC instructions

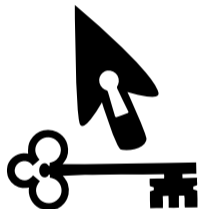




ARMv8.3 hardware-based pointer authentication

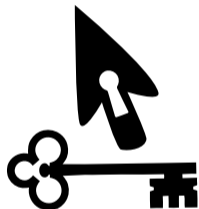
- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused
- 💡 Idea: repurpose bits for cryptographically authenticating pointers
- ⚙ Implementation
 - Compiler secures a pointer with special `PAC` instructions
 - Hardware computes cryptographic **MAC**





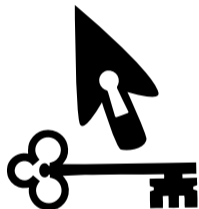
ARMv8.3 hardware-based pointer authentication

- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused
- 💡 Idea: repurpose bits for cryptographically authenticating pointers
- ⚙ Implementation
 - Compiler secures a pointer with special `PAC` instructions
 - Hardware computes cryptographic **MAC**
 - Hardware stores **parts** of the MAC in **unused** pointer bits



ARMv8.3 hardware-based pointer authentication

- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused
- 💡 Idea: repurpose bits for cryptographically authenticating pointers
- ⚙ Implementation
 - Compiler secures a pointer with special `PAC` instructions
 - Hardware computes cryptographic `MAC`
 - Hardware stores `parts` of the `MAC` in `unused` pointer bits
 - Before dereferencing (calling) a pointer, `AUT` instruction authenticates

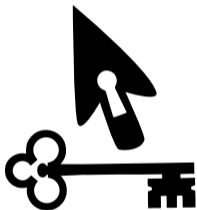


ARMv8.3 hardware-based pointer authentication

- 👁 Observation: 64-bit architectures do not use all bits
 - E.g. 48-bit virtual address space → 16 bits unused
- 💡 Idea: repurpose bits for cryptographically authenticating pointers
- ⚙ Implementation
 - Compiler secures a pointer with special `PAC` instructions
 - Hardware computes cryptographic **MAC**
 - Hardware stores **parts** of the MAC in **unused** pointer bits
 - Before dereferencing (calling) a pointer, **AUT** instruction authenticates
 - Hardware **invalidates** pointer or **faults** on authentication failure



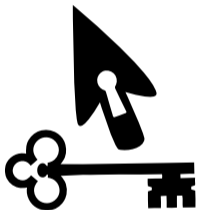
MAC algorithm has three inputs





MAC algorithm has three inputs

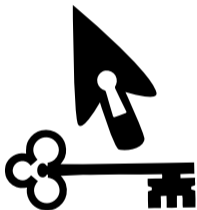
- Pointer value

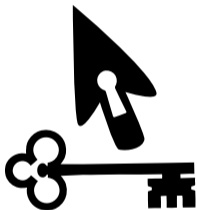




MAC algorithm has three inputs

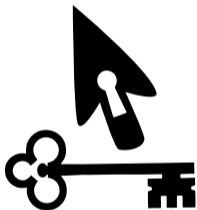
- Pointer value
- Secret (process-dependent) key





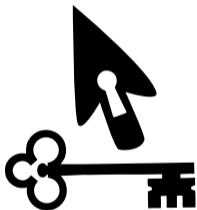
MAC algorithm has three inputs

- Pointer value
- Secret (process-dependent) key
- User-defined *context*



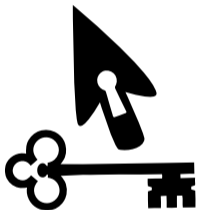
MAC algorithm has three inputs

- Pointer value
- Secret (process-dependent) key
- User-defined *context*
- *Context* can be used to distinguish



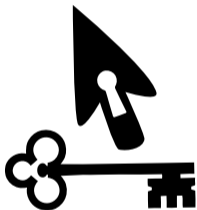
MAC algorithm has three inputs

- Pointer value
- Secret (process-dependent) key
- User-defined *context*
- *Context* can be used to distinguish
 - Function signatures



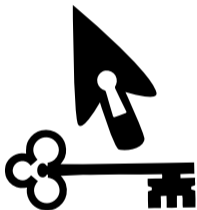
MAC algorithm has three inputs

- Pointer value
- Secret (process-dependent) key
- User-defined *context*
- *Context* can be used to distinguish
 - Function signatures
 - Types of data pointers



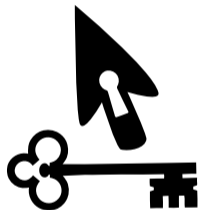
MAC algorithm has three inputs

- Pointer value
- Secret (process-dependent) key
- User-defined *context*
- *Context* can be used to distinguish
 - Function signatures
 - Types of data pointers
 - Stack frames

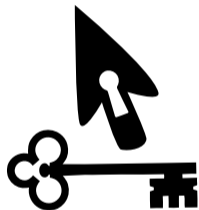


MAC algorithm has three inputs

- Pointer value
- Secret (process-dependent) key
- User-defined *context*
- *Context* can be used to distinguish
 - Function signatures
 - Types of data pointers
 - Stack frames
 - ...

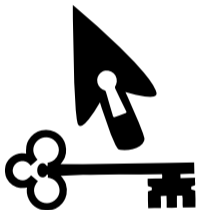


★ Properties



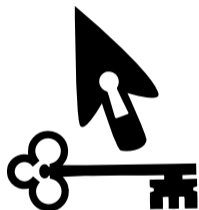
★ Properties

- Cryptographically-enforced CFI



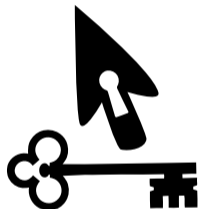
★ Properties

- Cryptographically-enforced CFI
- Effectiveness depends on additional *context* input



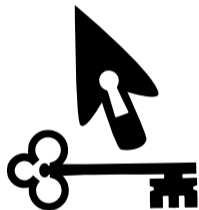
★ Properties

- Cryptographically-enforced CFI
- Effectiveness depends on additional *context* input
 - Precision: if *context* is zero, attacker can exchange all authenticated pointers (Similar to Intel CET)



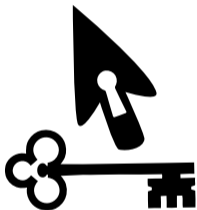
★ Properties

- Cryptographically-enforced CFI
- Effectiveness depends on additional *context* input
 - Precision: if *context* is zero, attacker can exchange all authenticated pointers (Similar to Intel CET)
 - Security: If attacker can control *context* → forge arbitrary pointers



★ Properties

- Cryptographically-enforced CFI
- Effectiveness depends on additional *context* input
 - Precision: if *context* is zero, attacker can exchange all authenticated pointers (Similar to Intel CET)
 - Security: If attacker can control *context* → forge arbitrary pointers
- Authentication code (MAC) is truncated to upper pointer bits



★ Properties

- Cryptographically-enforced CFI
- Effectiveness depends on additional *context* input
 - Precision: if *context* is zero, attacker can exchange all authenticated pointers (Similar to Intel CET)
 - Security: If attacker can control *context* → forge arbitrary pointers
- Authentication code (MAC) is truncated to upper pointer bits
 - Imprecision, allows brute-force/collision attacks



★ Properties



★ Properties

- CFI: attacker cannot escape control flow graph (CFG)



★ Properties

- CFI: attacker cannot escape control flow graph (CFG)
 - Defeats ROP, JOP, ...



★ Properties

- CFI: attacker cannot escape control flow graph (CFG)
 - Defeats ROP, JOP, ...
 - Still allows more or less code reuse within the CFG



★ Properties

- CFI: attacker cannot escape control flow graph (CFG)
 - Defeats ROP, JOP, ...
 - Still allows more or less code reuse within the CFG
 - Depending on precision of forward edges, attacker can substitute valid pointers



★ Properties

- CFI: attacker cannot escape control flow graph (CFG)
 - Defeats ROP, JOP, ...
 - Still allows more or less code reuse within the CFG
 - Depending on precision of forward edges, attacker can substitute valid pointers
- CFI does not prevent data-only attacks



★ Properties

- CFI: attacker cannot escape control flow graph (CFG)
 - Defeats ROP, JOP, ...
 - Still allows more or less code reuse within the CFG
 - Depending on precision of forward edges, attacker can substitute valid pointers
- CFI does not prevent data-only attacks
 - E.g., is_admin flag, loop counters, syscall arguments?

Memory Safety



👁 Observation: Most attacks due to a memory safety vulnerability



- 👁 Observation: Most attacks due to a memory safety vulnerability
- 💡 Idea: Prevent exploitation of memory safety vulnerability



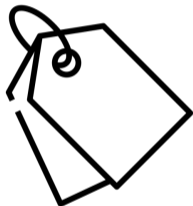
- 👁 Observation: Most attacks due to a memory safety vulnerability
- 💡 Idea: Prevent exploitation of memory safety vulnerability
 - Preventing invalid memory access



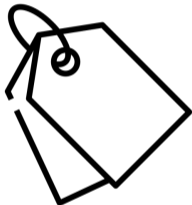
- 👁 Observation: Most attacks due to a memory safety vulnerability
- 💡 Idea: Prevent exploitation of memory safety vulnerability
 - Preventing invalid memory access
 - ARM Memory Tagging Extension (MTE)



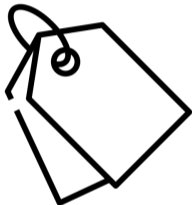
Memory Tagging Extension



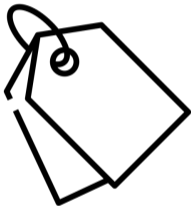
- On allocation



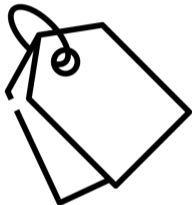
- On allocation
 - e.g., `char *ptr = malloc(8);`



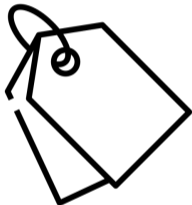
- On allocation
 - e.g., `char *ptr = malloc(8);`
 - Tag memory object



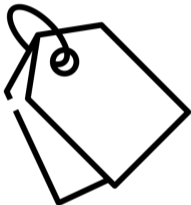
- On allocation
 - e.g., `char *ptr = malloc(8);`
 - Tag memory object
 - Store tag in pointer



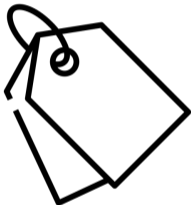
- On allocation
 - e.g., `char *ptr = malloc(8);`
 - Tag memory object
 - Store tag in pointer
- On access



- On allocation
 - e.g., `char *ptr = malloc(8);`
 - Tag memory object
 - Store tag in pointer
- On access
 - Check whether memory object is tagged with the stored tag of the pointer



- On allocation
 - e.g., `char *ptr = malloc(8);`
 - Tag memory object
 - Store tag in pointer
- On access
 - Check whether memory object is tagged with the stored tag of the pointer
 - If correct → Access



- On allocation
 - e.g., `char *ptr = malloc(8);`
 - Tag memory object
 - Store tag in pointer
- On access
 - Check whether memory object is tagged with the stored tag of the pointer
 - If correct → Access
 - If not correct → Fault



```
char *ptr = new char [16]; // memory colored
```



```
ptr[17] = 42; // color mismatch -> overflow
```



```
delete [] ptr; // memory re-coloured on free
```



```
ptr[10] = 10; // color mismatch -> use-after-free
```



★ Security claims



```
char *ptr = new char [16]; // memory colored
```



```
ptr[17] = 42; // color mismatch -> overflow
```



```
delete [] ptr; // memory re-coloured on free
```



```
ptr[10] = 10; // color mismatch -> use-after-free
```



★ Security claims

- Prevents spatial memory violations



```
char *ptr = new char [16]; // memory colored
```



```
ptr[17] = 42; // color mismatch -> overflow
```



```
delete [] ptr; // memory re-coloured on free
```



```
ptr[10] = 10; // color mismatch -> use-after-free
```

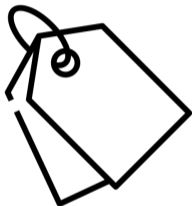


★ Security claims

- Prevents spatial memory violations
- Prevents temporal memory violations



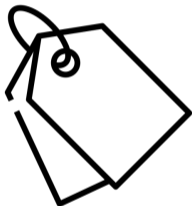
★ Tag size of MTE

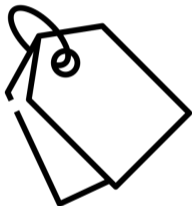




★ Tag size of MTE

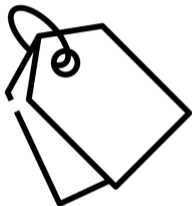
- 4 bits



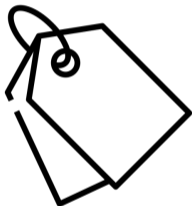


★ Tag size of MTE

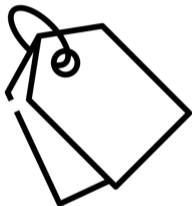
- 4 bits
- 16 distinct tags



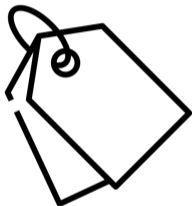
- ★ Tag size of MTE
 - 4 bits
 - 16 distinct tags
- ★ Tag storage



- ★ Tag size of MTE
 - 4 bits
 - 16 distinct tags
- ★ Tag storage
 - Pointer



- ★ Tag size of MTE
 - 4 bits
 - 16 distinct tags
- ★ Tag storage
 - Pointer
 - Only 48 bits used of virtual address

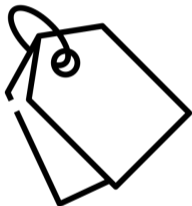


★ Tag size of MTE

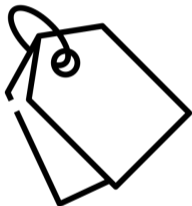
- 4 bits
- 16 distinct tags

★ Tag storage

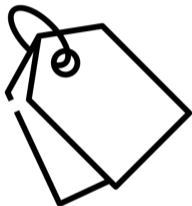
- Pointer
 - Only 48 bits used of virtual address
 - Store 4 bit tag in unused bits



- ★ Tag size of MTE
 - 4 bits
 - 16 distinct tags
- ★ Tag storage
 - Pointer
 - Only 48 bits used of virtual address
 - Store 4 bit tag in unused bits
 - Memory object



- ★ Tag size of MTE
 - 4 bits
 - 16 distinct tags
- ★ Tag storage
 - Pointer
 - Only 48 bits used of virtual address
 - Store 4 bit tag in unused bits
 - Memory object
 - In memory

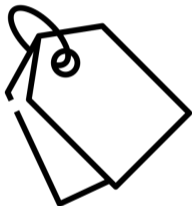


★ Tag size of MTE

- 4 bits
- 16 distinct tags

★ Tag storage

- Pointer
 - Only 48 bits used of virtual address
 - Store 4 bit tag in unused bits
- Memory object
 - In memory
 - 3.125 % memory overhead with 4 bit tag size

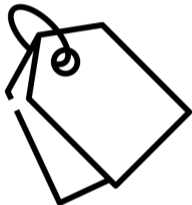


★ Tag size of MTE

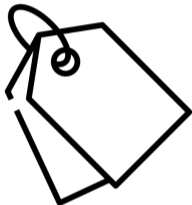
- 4 bits
- 16 distinct tags

★ Tag storage

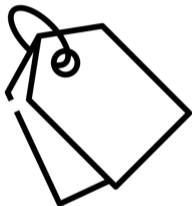
- Pointer
 - Only 48 bits used of virtual address
 - Store 4 bit tag in unused bits
- Memory object
 - In memory
 - 3.125 % memory overhead with 4 bit tag size
 - Two memory lookups for dereferencing a pointer



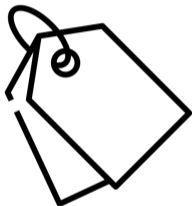
- ❓ Given 16 tags, how likely is it to detect an arbitrary memory access or a use-after-free?



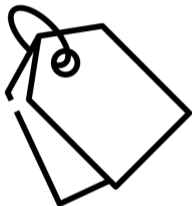
- ❓ Given 16 tags, how likely is it to detect an arbitrary memory access or a use-after-free?
- On random re-coloring



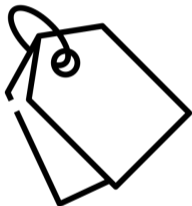
- ❓ Given 16 tags, how likely is it to detect an arbitrary memory access or a use-after-free?
- On random re-coloring
 - $(1 - \frac{1}{2^4}) \cdot 100 = 93.75\%$



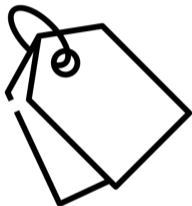
- ❓ Given 16 tags, how likely is it to detect an arbitrary memory access or a use-after-free?
 - On random re-coloring
 - $(1 - \frac{1}{2^4}) \cdot 100 = 93.75\%$
- ❓ Increase tag size to 8 bits?



- ❓ Given 16 tags, how likely is it to detect an arbitrary memory access or a use-after-free?
 - On random re-coloring
 - $(1 - \frac{1}{2^4}) \cdot 100 = 93.75\%$
- ❓ Increase tag size to 8 bits?
 - Detection probability increases to $(1 - \frac{1}{2^8}) \cdot 100 \approx 99.6\%$



- ❓ Given 16 tags, how likely is it to detect an arbitrary memory access or a use-after-free?
 - On random re-coloring
 - $(1 - \frac{1}{2^4}) \cdot 100 = 93.75 \%$
- ❓ Increase tag size to 8 bits?
 - Detection probability increases to $(1 - \frac{1}{2^8}) \cdot 100 \approx 99.6 \%$
 - Memory overhead increases to 6.25 %



- ❓ Given 16 tags, how likely is it to detect an arbitrary memory access or a use-after-free?
 - On random re-coloring
 - $(1 - \frac{1}{2^4}) \cdot 100 = 93.75\%$
- ❓ Increase tag size to 8 bits?
 - Detection probability increases to $(1 - \frac{1}{2^8}) \cdot 100 \approx 99.6\%$
 - Memory overhead increases to 6.25%
 - Currently only 4 bit supported

Attacker's perspective

- 🔍 Vulnerability discovery ✓
- 👤 Exploitation ✓
- 🔑 Privilege elevation ✓

Defender's perspective

- 🔍 Vulnerability prevention ✓
- 👤 Exploit prevention ✓
- 🔑 **Privilege minimization** (today)

Privilege Minimization

How to minimize the impact of **Arbitrary Code Execution**

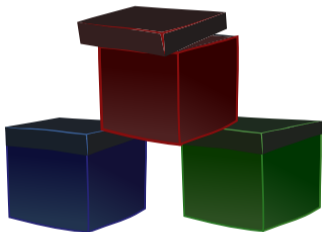
Think inside boxes ...

Think inside boxes ...

MINECRAFT

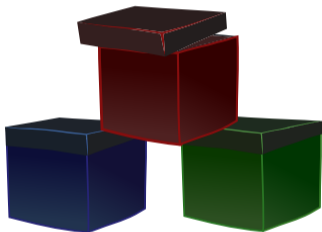


**PROOF THAT HIGH QUALITY COME
IN LOW RESOLUTION**



💡 Everything is a box

- "Principle of least privileges"

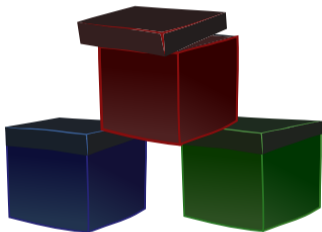


💡 Everything is a box

- "Principle of least privileges"

☐☐ Compartmentalization

- Break large boxes into smaller boxes
- Virtual machines, processes, libraries, functions ...
- Mostly manual effort



💡 Everything is a box

- "Principle of least privileges"

▣ Compartmentalization

- Break large boxes into smaller boxes
- Virtual machines, processes, libraries, functions ...
- Mostly manual effort

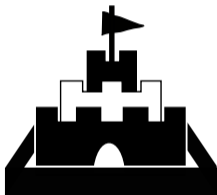
🛡 Isolation

- Isolate boxes from each other
- Safeguard all interfaces
 - File permissions, network firewall ... **system calls**

In-process Sandboxing



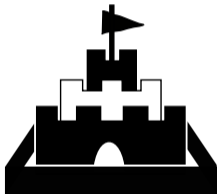
🚩 Goal: confine parts of an application

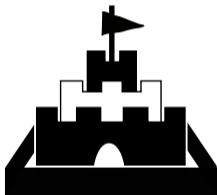




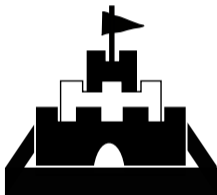
🚩 Goal: confine parts of an application

- E.g., dangerous plugins, libraries, user-provided code ...

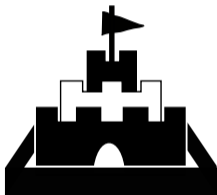




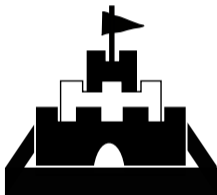
- 🚩 Goal: confine parts of an application
 - E.g., dangerous plugins, libraries, user-provided code ...
- 💡 Dangerous code is not executed natively but interpreted



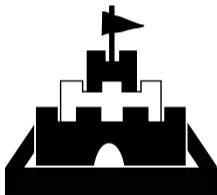
- 🚩 Goal: confine parts of an application
 - E.g., dangerous plugins, libraries, user-provided code ...
- 💡 Dangerous code is not executed natively but interpreted
 - E.g., JavaScript, WebAssembly, ...



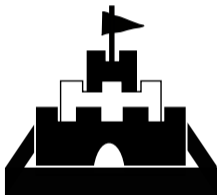
- 🚩 Goal: confine parts of an application
 - E.g., dangerous plugins, libraries, user-provided code ...
- 💡 Dangerous code is not executed natively but interpreted
 - E.g., JavaScript, WebAssembly, ...
- Software-enforced:



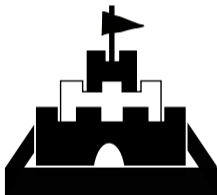
- 🚩 Goal: confine parts of an application
 - E.g., dangerous plugins, libraries, user-provided code ...
- 💡 Dangerous code is not executed natively but interpreted
 - E.g., JavaScript, WebAssembly, ...
- Software-enforced:
 - Software-Fault Isolation (SFI)



- 🚩 Goal: confine parts of an application
 - E.g., dangerous plugins, libraries, user-provided code ...
- 💡 Dangerous code is not executed natively but interpreted
 - E.g., JavaScript, WebAssembly, ...
- Software-enforced:
 - Software-Fault Isolation (SFI)
 - Masks access to restrict memory accesses



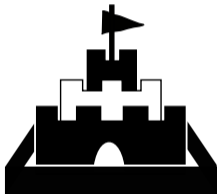
- 🚩 Goal: confine parts of an application
 - E.g., dangerous plugins, libraries, user-provided code ...
- 💡 Dangerous code is not executed natively but interpreted
 - E.g., JavaScript, WebAssembly, ...
- Software-enforced:
 - Software-Fault Isolation (SFI)
 - Masks access to restrict memory accesses
- Hardware-enforced:



- 🚩 Goal: confine parts of an application
 - E.g., dangerous plugins, libraries, user-provided code ...
- 💡 Dangerous code is not executed natively but interpreted
 - E.g., JavaScript, WebAssembly, ...
- Software-enforced:
 - Software-Fault Isolation (SFI)
 - Masks access to restrict memory accesses
- Hardware-enforced:
 - Memory Protection Keys (MPK)

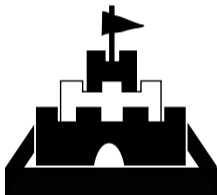


- Tag pages with a key and change permission of the key



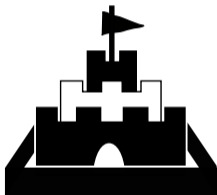


- Tag pages with a key and change permission of the key
- Change access permissions of a key → change permission the tagged pages



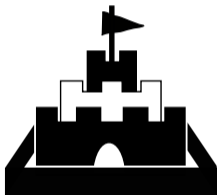


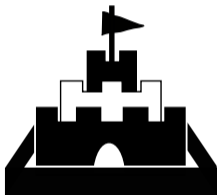
- Tag pages with a key and change permission of the key
- Change access permissions of a key → change permission the tagged pages
- Tagged key stored in page-table



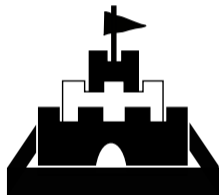


- Tag pages with a key and change permission of the key
- Change access permissions of a key → change permission the tagged pages
- Tagged key stored in page-table
- Keys permissions are stored in dedicated register

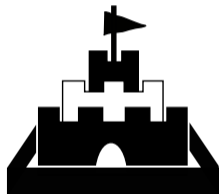




- Tag pages with a key and change permission of the key
- Change access permissions of a key → change permission the tagged pages
- Tagged key stored in page-table
- Keys permissions are stored in dedicated register
- Allow fine-grained permission setting, *i.e.*, read/write-access, of tagged page



- Tag pages with a key and change permission of the key
- Change access permissions of a key → change permission the tagged pages
- Tagged key stored in page-table
- Keys permissions are stored in dedicated register
- Allow fine-grained permission setting, *i.e.*, read/write-access, of tagged page
- Hardware implementations:
 - Intel Protection Keys for Userspace (PKU) and Protection Keys for Supervisor (PKS)
 - ARM Domain Access Control (DAC)



- Tag pages with a key and change permission of the key
- Change access permissions of a key → change permission the tagged pages
- Tagged key stored in page-table
- Keys permissions are stored in dedicated register
- Allow fine-grained permission setting, *i.e.*, read/write-access, of tagged page
- Hardware implementations:
 - Intel Protection Keys for Userspace (PKU) and Protection Keys for Supervisor (PKS)
 - ARM Domain Access Control (DAC)
- Only 4 bit key → 16 distinct access domains

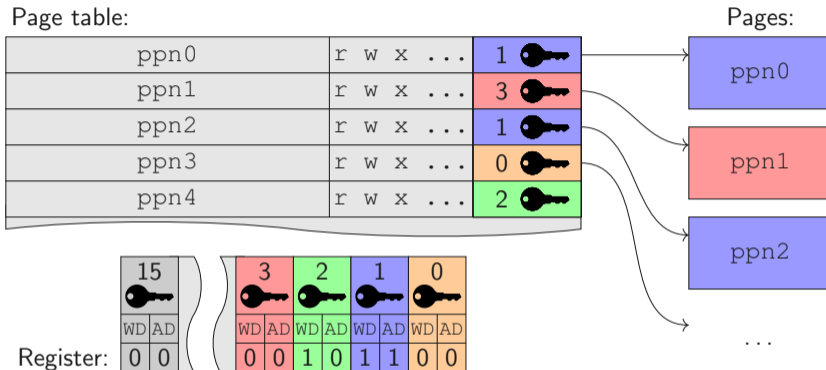


Figure 1: Working principle of MPK, where *AD* and *WD* stands for access and write disable, respectively. Pages tagged with key 0 are write- and access-permitted, while pages tagged with key 1 are write- and access-prohibited, and pages tagged with key 2 are only write-prohibited.

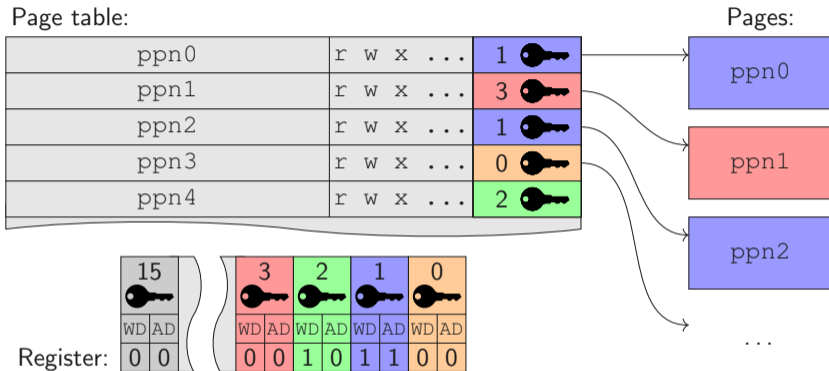


Figure 1: Working principle of MPK, where *AD* and *WD* stands for access and write disable, respectively. Pages tagged with key 0 are write- and access-permitted, while pages tagged with key 1 are write- and access-prohibited, and pages tagged with key 2 are only write-prohibited. **Advantages** over `mprotect`: No TLB flush and page-table walk on permission change → faster.

Process Sandboxing





👁 Observation: Most programs do not need most system calls





- 👁 Observation: Most programs do not need most system calls
- E.g., fork, exec, prctl ...



- 👁 Observation: Most programs do not need most system calls
 - E.g., fork, exec, prctl ...
- 💡 Idea: block unnecessary system calls



- 👁 Observation: Most programs do not need most system calls
 - E.g., fork, exec, prctl ...
- 💡 Idea: block unnecessary system calls
- ⚙ Implementation



- 👁 Observation: Most programs do not need most system calls
 - E.g., fork, exec, prctl ...
- 💡 Idea: block unnecessary system calls
- ⚙ Implementation
 - Program installs seccomp filters on startup



- 👁 Observation: Most programs do not need most system calls
 - E.g., fork, exec, prctl ...
- 💡 Idea: block unnecessary system calls
- ⚙ Implementation
 - Program installs seccomp filters on startup
 - Seccomp supports small *Berkeley Packet Filter (BPF) programs*



- 👁 Observation: Most programs do not need most system calls
 - E.g., fork, exec, prctl ...
- 💡 Idea: block unnecessary system calls
- ⚙ Implementation
 - Program installs seccomp filters on startup
 - Seccomp supports small *Berkeley Packet Filter (BPF) programs*
 - Kernel does the filtering (e.g., executes the BPF program) on every system call



👁 Observation: Most programs do not need most system calls

- E.g., fork, exec, prctl ...

💡 Idea: block unnecessary system calls

⚙ Implementation

- Program installs seccomp filters on startup
- Seccomp supports small *Berkeley Packet Filter (BPF) programs*
- Kernel does the filtering (e.g., executes the BPF program) on every system call
- On a filter violation: deny syscall, send signal, kill program ...



```
#include <stdio.h>          /* printf */
#include <sys/prctl.h>       /* prctl */
#include <linux/seccomp.h>   /* seccomp's constants */
#include <unistd.h>         /* dup2: just for test */

int main() {
    printf("step 1: unrestricted\n");
    prctl(PR_SET_SECCOMP, SECCOMP_MODE_STRICT); // Enable filtering
    printf("step 2: only 'read', 'write', '_exit' and 'sigreturn' syscalls\n");
    dup2(1, 2); // redirect stderr to stdout
    printf("step 3: !! YOU SHOULD NOT SEE ME !!\n");
    return 0;
}
```

<https://blog.yadutaf.fr/2014/05/29/introduction-to-seccomp-bpf-linux-syscall-filter/>



```
dgruss@t460sdg ~ % gcc seccomp.c
dgruss@t460sdg ~ % ./a.out
step 1: unrestricted
step 2: only 'read', 'write', '_exit' and 'sigreturn' syscalls
[1] 19622 killed ./a.out
137 dgruss@t460sdg ~ %
```



```
int main() {
    printf("step 1: init\n");
    prctl(PR_SET_NO_NEW_PRIVS, 1);
    prctl(PR_SET_DUMPABLE, 0);      // ptrace on this process / childs is not allowed
    scmp_filter_ctx ctx;
    ctx = seccomp_init(SCMP_ACT_KILL);           // Denylist everything
    seccomp_rule_add(ctx, SCMP_ACT_ALLOW, SCMP_SYS(rt_sigreturn), 0); // Allowlist
    seccomp_rule_add(ctx, SCMP_ACT_ALLOW, SCMP_SYS(exit), 0);       // Allowlist
    seccomp_rule_add(ctx, SCMP_ACT_ALLOW, SCMP_SYS(exit_group), 0); // Allowlist
    seccomp_rule_add(ctx, SCMP_ACT_ALLOW, SCMP_SYS(read), 0);      // Allowlist
    seccomp_rule_add(ctx, SCMP_ACT_ALLOW, SCMP_SYS(write), 0);     // Allowlist
    seccomp_rule_add(ctx, SCMP_ACT_ALLOW, SCMP_SYS(dup2), 2,      // Allowlist
                    SCMP_A0(SCMP_CMP_EQ, 1), SCMP_A1(SCMP_CMP_EQ, 2)); // Allowlist
    seccomp_load(ctx);
    printf("step 2: only 'write' and dup2(1, 2) syscalls\n");
    dup2(1, 2); // redirect stderr to stdout
    printf("step 3: stderr redirected to stdout\n");
    dup2(2, 42); // redirect stderr to stdout
}
```



```
dgruss@t460sdg ~ % gcc seccomp.c -lseccomp && ./a.out
step 1: init
step 2: only 'write' and dup2(1, 2) syscalls
step 3: stderr redirected to stdout
[1] 23312 invalid system call ./a.out
159 dgruss@t460sdg ~ % █
```



Write a secure wrapper binary

- Usage: `./secwrap <command>`
 - The wrapper shall start the program specified by `<command>`
 - Anything `<command>` does may not be allowed to create new processes!
- Very convenient to use :)
- Upload your wrapper binary at <https://challenges.sasectf.student.iaik.tugraz.at/secwrap/index.php>
 - If it is correct, you will get the flag
 - Test system is Debian 11.8, kernel 5.10.0-26



- Sandbox process runs dangerous code



- Sandbox process runs dangerous code
- Monitor process interacts with sandbox via IPC



- Sandbox process runs dangerous code
- Monitor process interacts with sandbox via IPC
 - Minimal filter: only allow required IPC system calls



- Sandbox process runs dangerous code
- Monitor process interacts with sandbox via IPC
 - Minimal filter: only allow required IPC system calls
- Example: Google *sandbox2*
<https://developers.google.com/sandboxed-api/>



★ Properties





★ Properties

- Protect **system call** interface



★ Properties

- Protect **system call** interface
- Filters can only be **specialized** but not tightened



★ Properties

- Protect **system call** interface
- Filters can only be **specialized** but not tightened
 - Attacker cannot manipulate/unload existing filters



★ Properties

- Protect **system call** interface
- Filters can only be **specialized** but not tightened
 - Attacker cannot manipulate/unload existing filters
- Filter: simple arithmetic operations on system call arguments



★ Properties

- Protect **system call** interface
- Filters can only be **specialized** but not tightened
 - Attacker cannot manipulate/unload existing filters
- Filter: simple arithmetic operations on system call arguments
 - Enhanced filtering is impossible



★ Properties

- Protect **system call** interface
- Filters can only be **specialized** but not tightened
 - Attacker cannot manipulate/unload existing filters
- Filter: simple arithmetic operations on system call arguments
 - Enhanced filtering is impossible
 - E.g., checking for strings, sanitizing paths, dereferencing pointers



★ Properties

- Protect **system call** interface
 - Filters can only be **specialized** but not tightened
 - Attacker cannot manipulate/unload existing filters
 - Filter: simple arithmetic operations on system call arguments
 - Enhanced filtering is impossible
 - E.g., checking for strings, sanitizing paths, dereferencing pointers
- ❓ How do we know which system calls are needed by libc functions such as `pthread_create` ? Implementation defined!



★ Properties

- Protect **system call** interface
 - Filters can only be **specialized** but not tightened
 - Attacker cannot manipulate/unload existing filters
 - Filter: simple arithmetic operations on system call arguments
 - Enhanced filtering is impossible
 - E.g., checking for strings, sanitizing paths, dereferencing pointers
- ❓ How do we know which system calls are needed by libc functions such as `pthread_create` ? Implementation defined!
- ❓ How can we virtualize resources?





- 💡 Idea: Manage resource usage of a group of processes (and all its children)



- 💡 Idea: Manage resource usage of a group of processes (and all its children)
 - Memory, CPU time, networking, disk I/O ...



- 💡 Idea: Manage resource usage of a group of processes (and all its children)
- Memory, CPU time, networking, disk I/O ...
 - Set limits / priorities



💡 Idea: Manage resource usage of a group of processes (and all its children)

- Memory, CPU time, networking, disk I/O ...
- Set limits / priorities

★ Properties



- 💡 Idea: Manage resource usage of a group of processes (and all its children)
 - Memory, CPU time, networking, disk I/O ...
 - Set limits / priorities
- ★ Properties
 - Can prevent some Denial-of-Service (DoS) attacks



💡 Idea: Manage resource usage of a group of processes (and all its children)

- Memory, CPU time, networking, disk I/O ...
- Set limits / priorities

★ Properties

- Can prevent some Denial-of-Service (DoS) attacks
- Cannot prevent privilege escalation





💡 Idea: Namespace hides (virtualizes) resources from processes



- 💡 Idea: Namespace hides (virtualizes) resources from processes
 - Various namespaces: `mnt`, `pid`, `net`, `ipc`, `uts (hostname)`, `user`



- 💡 Idea: Namespace hides (virtualizes) resources from processes
 - Various namespaces: `mnt`, `pid`, `net`, `ipc`, `uts (hostname)`, `user`
 - How? Namespace translates resource identifiers



- 💡 Idea: Namespace hides (virtualizes) resources from processes
 - Various namespaces: `mnt`, `pid`, `net`, `ipc`, `uts (hostname)`, `user`
 - How? Namespace translates resource identifiers
- Examples:



- 💡 Idea: Namespace hides (virtualizes) resources from processes
 - Various namespaces: `mnt`, `pid`, `net`, `ipc`, `uts (hostname)`, `user`
 - How? Namespace translates resource identifiers
- Examples:
 - Inside namespace: `uid=0 (root)`, `path=/f.txt`



- 💡 Idea: Namespace hides (virtualizes) resources from processes
 - Various namespaces: `mnt`, `pid`, `net`, `ipc`, `uts (hostname)`, `user`
 - How? Namespace translates resource identifiers
- Examples:
 - Inside namespace: `uid=0 (root)`, `path=/f.txt`
 - Outside namespace: `uid=1000 (ssd)`, `path=/home/ssd/f.txt`



💡 Idea: combine 'em all: Docker containers





- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)





- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
 - Docker automatically





- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups



- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups
 - configures seccomp



- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups
 - configures seccomp
- ★ Properties



💡 Idea: combine 'em all: Docker containers

- See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups
 - configures seccomp

★ Properties

- Fast



- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups
 - configures seccomp
- ★ Properties
 - Fast
 - Security depends on proper configuration



- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups
 - configures seccomp
- ★ Properties
 - Fast
 - Security depends on proper configuration
 - Kernel is shared between containers and host



- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups
 - configures seccomp
- ★ Properties
 - Fast
 - Security depends on proper configuration
 - Kernel is shared between containers and host
- ❓ What if one container compromises the host kernel?



- 💡 Idea: combine 'em all: Docker containers
 - See also Linux Containers (LXC)
- Docker automatically
 - creates namespaces and cgroups
 - configures seccomp
- ★ Properties
 - Fast
 - Security depends on proper configuration
 - Kernel is shared between containers and host
- ❓ What if one container compromises the host kernel?
GVisor/Kata to harden Docker containers

Virtualization





💡 Idea: fake the entire system





💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)





💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines



💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines

⚙️ Implementation



💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines

⚙️ Implementation

- Managed by **hypervisor**



💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines

⚙️ Implementation

- Managed by **hypervisor**
 - Xen, VMware, VirtualBox, Hyper-V, Qemu ...



💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines

⚙️ Implementation

- Managed by **hypervisor**
 - Xen, VMware, VirtualBox, Hyper-V, Qemu ...
- Typically **hardware-accelerated**



💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines

⚙️ Implementation

- Managed by **hypervisor**
 - Xen, VMware, VirtualBox, Hyper-V, Qemu ...
- Typically **hardware-accelerated**
- **CloudOS** in the summer term



💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines

⚙️ Implementation

- Managed by **hypervisor**
 - Xen, VMware, VirtualBox, Hyper-V, Qemu ...
- Typically **hardware-accelerated**
- **CloudOS** in the summer term

❓ What if VM compromises hypervisor?



💡 Idea: fake the entire system

- Virtualization of hardware resources (memory, CPU, peripherals ...)
- System runs many isolated virtual machines

⚙️ Implementation

- Managed by **hypervisor**
 - Xen, VMware, VirtualBox, Hyper-V, Qemu ...
- Typically **hardware-accelerated**
- **CloudOS** in the summer term

❓ What if VM compromises hypervisor?

❓ Is there an end to this recursive problem?

Enclaves





👁 Observation: Sandboxes follow **hierarchical** ring model



- 👁 Observation: Sandboxes follow **hierarchical** ring model
 - Higher rings (kernel space) have strictly higher privileges



- 👁 Observation: Sandboxes follow **hierarchical** ring model
- Higher rings (kernel space) have strictly higher privileges
 - Lower rings (user space) need to fully trust higher rings



- 👁 Observation: Sandboxes follow **hierarchical** ring model
- Higher rings (kernel space) have strictly higher privileges
 - Lower rings (user space) need to fully trust higher rings
 - Vulnerability in higher ring is **fatal**



- 👁 Observation: Sandboxes follow **hierarchical** ring model
 - Higher rings (kernel space) have strictly higher privileges
 - Lower rings (user space) need to fully trust higher rings
 - Vulnerability in higher ring is **fatal**
- 💡 Idea: build a reverse sandbox: Enclaves



- 👁 Observation: Sandboxes follow **hierarchical** ring model
 - Higher rings (kernel space) have strictly higher privileges
 - Lower rings (user space) need to fully trust higher rings
 - Vulnerability in higher ring is **fatal**
- 💡 Idea: build a reverse sandbox: Enclaves
 - Only trust **enclave code** (and **hardware**)



- 👁 Observation: Sandboxes follow **hierarchical** ring model
 - Higher rings (kernel space) have strictly higher privileges
 - Lower rings (user space) need to fully trust higher rings
 - Vulnerability in higher ring is **fatal**
- 💡 Idea: build a reverse sandbox: Enclaves
 - Only trust **enclave code** (and **hardware**)
 - Distrust all non-enclave code



- 👁 Observation: Sandboxes follow **hierarchical** ring model
 - Higher rings (kernel space) have strictly higher privileges
 - Lower rings (user space) need to fully trust higher rings
 - Vulnerability in higher ring is **fatal**
- 💡 Idea: build a reverse sandbox: Enclaves
 - Only trust **enclave code** (and **hardware**)
 - Distrust all non-enclave code
 - Host application, kernel, hypervisor



- 👁 Observation: Sandboxes follow **hierarchical** ring model
 - Higher rings (kernel space) have strictly higher privileges
 - Lower rings (user space) need to fully trust higher rings
 - Vulnerability in higher ring is **fatal**
- 💡 Idea: build a reverse sandbox: Enclaves
 - Only trust **enclave code** (and **hardware**)
 - Distrust all non-enclave code
 - Host application, kernel, hypervisor
 - Example: Intel SGX/TDX, AMD-SEV, ARM Trustzone



★ Properties



★ Properties

- Enclaves protect a piece of secure code / data



★ Properties

- Enclaves protect a piece of secure code / data
- Enclaves cannot sandbox untrusted code



★ Properties

- Enclaves protect a piece of secure code / data
- Enclaves cannot sandbox untrusted code
- Can be (mis)used for Digital Rights Management (DRM), hiding malware

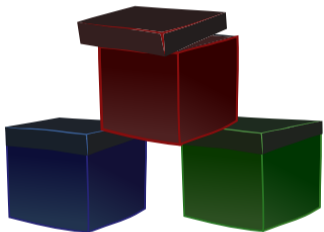


★ Properties

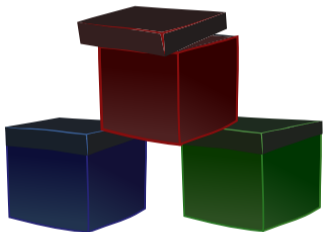
- Enclaves protect a piece of secure code / data
- Enclaves cannot sandbox untrusted code
- Can be (mis)used for Digital Rights Management (DRM), hiding malware

❓ Are we (too) secure now?

Summary & Outlook

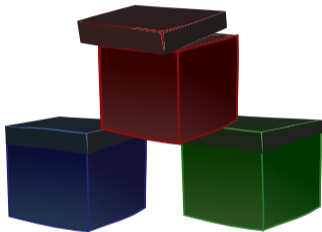


💡 Everything is a box



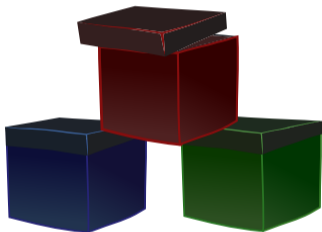
💡 Everything is a box

- **Compartmentalization:** Make boxes as small as possible



💡 Everything is a box

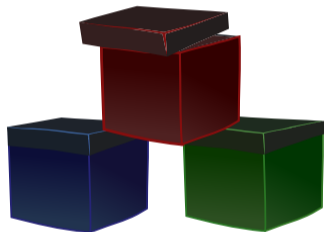
- **Compartmentalization:** Make boxes as small as possible
- **Isolation:** A box shall have minimal permission



💡 Everything is a box

- **Compartmentalization:** Make boxes as small as possible
- **Isolation:** A box shall have minimal permission
- "Principle of least privileges"

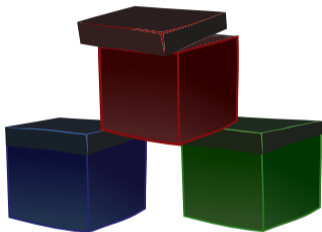
Isolation techniques

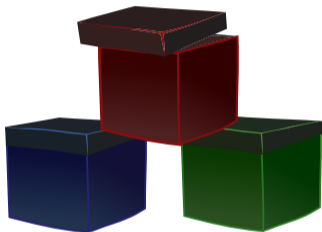


Isolation techniques

🛡 In-process Sandboxing

- Isolate domains within process





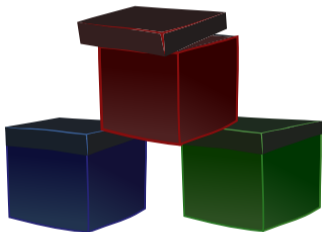
Isolation techniques

🛡 In-process Sandboxing

- Isolate domains within process

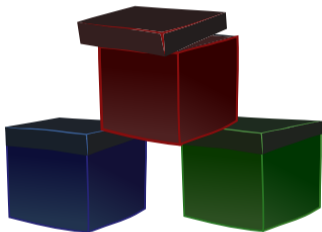
🛡 Process Sandboxing

- Seccomp
- Cgroups
- Namespaces
- Docker container = seccomp + cgroups + namespaces



Isolation techniques

- 🛡 In-process Sandboxing
 - Isolate domains within process
- 🛡 Process Sandboxing
 - Seccomp
 - Cgroups
 - Namespaces
 - Docker container = seccomp + cgroups + namespaces
- 🛡 Virtualization
 - Full system virtualization






Isolation techniques




- 🛡 In-process Sandboxing
 - Isolate domains within process
- 🛡 Process Sandboxing
 - Seccomp
 - Cgroups
 - Namespaces
 - Docker container = seccomp + cgroups + namespaces
- 🛡 Virtualization
 - Full system virtualization
- 🛡 Enclaves
 - Reverse sandbox

Summary

Attacker's perspective

-  Vulnerability discovery
-  Exploitation
-  Privilege elevation

Defender's perspective

-  Vulnerability prevention
-  Exploit prevention
-  Privilege minimization

Questions?

