

Question 1 *Buffer Overflow Mitigations*

(10 min)

Buffer overflow mitigations generally fall into two categories: (1) eliminating the cause and (2) alleviating the damage. This question is about techniques in the second category.

Several requirements must be met for a buffer overflow to succeed. Each requirement listed below can be combated with a different countermeasure. With each mitigation you discuss, think about *where* it can be implemented—common targets include the compiler and the operating system (OS). Also discuss limitations, pitfalls, and costs of each mitigation.

- (a) The attacker needs to overwrite the return address on the stack to change the control flow. Is it possible to prevent this from happening or at least detect when it occurs?
- (b) The overwritten return address must point to a valid instruction sequence. The attacker often places the malicious code to execute in the vulnerable buffer. However, the buffer address must be known to set up the jump target correctly. One way to find out this address is to observe the program in a debugger. This works because the address tends to same across multiple program runs. What could be done to make it harder to accurately find out the address of the start of the malicious code?
- (c) Attackers often store their malicious code inside the same buffer that they overflow. What mechanism could prevent the execution of the malicious code? What type of code would break with this defense in place?

Solution:

- (a) **Stack Canaries.** A *canary* or *canary word* is a known value placed between the local variables and control data on the stack. Before reading the return address, code inserted by the compiler checks the canary against the known value. Since a successful buffer overflows needs to overwrite the canary before reaching the return address, and the attacker cannot predict the canary value, the canary validation will fail and stop execution prior to the jump.

As an example, consider the following function.

```
void vuln()
{
    char buf[32];
    gets(buf);
}
```

The compiler will take this function and generate:

```
/* This number is randomly set before each run. */
int MAGIC = rand();

void vuln()
{
    int canary = MAGIC;
    char buf[32];
    gets(buf);
    if (canary != MAGIC)
        HALT();
}
```

Limitations.

- Canaries only protect against stack smashing attacks, not against heap overflows or format string vulnerabilities.
- Local variables, such as function pointers and authentication flags, can still be overwritten.
- No protection against buffer *underflows*. This can be problematic in combination with the previous point.
- If the attack occurs before the end of the function, the canary validation does not even take place. This happens for example when an exception handler on the stack gets invoked before the function returns.
- A canary generated from a low-entropy pool can be predictable. In 2011 research showed that the Windows canary implementation only relied on 1 bit of entropy.

Cost. The canary has to be validated on each function return. The performance overhead is only a few percent since a canary is only needed in functions with local arrays. To determine whether to use the canary, Windows additionally applies heuristics (which unfortunately can also be subverted.)

- (b) **Address Randomization.** When the OS loader puts an executable into memory, it maps the different sections (text, data/BSS, heap, stack) to fixed memory

locations. In the mitigation technique called *address space layout randomization* (ASLR), rather than deterministically allocating the process layout, the OS randomizes the starting base of each section. This randomization makes it more difficult for an attacker to predict the addresses of jump targets. For instance, the OS might decide to start stack frames from somewhere other than the highest memory address.

Limitations.

- Entropy reduction attacks can significantly lower the efficacy of ASLR. For example, reducing factors are page alignment requirements (stack: 16 bytes, heap: 4096 bytes).
- Address space information disclosure techniques can force applications to leak known addresses (e.g., DLL addresses).
- Revealing addresses via brute-forcing can also be an effective technique when an application does not terminate, e.g., when a block that catches exceptions exists.
- Techniques known as *heap spraying* and *JIT spraying* allow an attacker to inject code at predictable locations.
- Like the canary defense, ASLR also does not defend against local data manipulation.
- Not all applications work properly with ASLR. In Windows, some opt out via the `/DYNAMICBAS` linker flag.

Cost. The overhead incurred by ASLR is negligible.

- (c) **Executable Space Protection.** Modern CPUs include a feature to mark certain memory regions non-executable. AMD calls this feature the NX (**n**o **e**xecute) bit and Intel the XD (**e**xecute **d**isable) bit. The idea is to combat buffer overflows where the attacker injects their own code.

PaX pioneered this technique in 2000 with per-page non-executable page support, protecting binary data, the heap, and the stack. OpenBSD implemented a form of this called $W\oplus X$ (write x-or execute) in 2003. Since Service Pack 2 in 2004, Windows features *data execution prevention* (DEP), an executable space protection mechanism that uses the NX or XD bit to mark pages, which are intended to only contain data, as non-executable.

Limitations.

- An attacker does not have to inject their own code. It is also possible to leverage existing instruction sequences in memory and jump to them. See [part 2](#) for details.

- The defense mechanism disallows execution of code generated at runtime, such as during JIT compilation or self-modifying code (SMC).
- If code is loaded at predictable addresses, it is possible to turn non-executable into executable code, e.g., via system functions like `VirtualAlloc` or `VirtualProtect` on Windows.

Cost. There is no measurable overhead due to the hardware support of modern CPUs.


Question 2 *Arc Injection* (10 min)

Imagine that you are trying to exploit a buffer overflow, but you notice that none of the code you are injecting will execute for some reason. How frustrating! You still really want to run some malicious code, so what could you try instead?

Hint: In a stack smashing attack, you can overwrite the return address with any address of your choosing.

Solution: Rather than injection code, the main idea of *arc injection* is to inject data. It is a powerful technique that bypasses numerous protection mechanisms, in particular executable space protection (1a:exec). By injecting malicious data that existing instructions later operates on, an attacker can still manipulate the execution path.

For example, an attacker can overwrite the return address with a function in `libc`, such as `system(const char* cmd)` whose single argument `cmd` is the new program to spawn. The attacker also has to setup the arguments (i.e., the data) appropriately. Recall that function arguments are pushed in reverse order on the stack before pushing the return address. Consider the example below, where an attacker overwrites the return address with the address of `system` (denoted by `&system`) to spawn a shell.



The first figure on the left is the stack layout before the attack. The second figure in the middle represents the state after having overflowed the buffer. Here, the return address is overwritten with `&system`. The value above is the location of the return address, from the perspective of `system`'s stack frame. But since the attacker plans

on spawning a shell that blocks to take evil commands (e.g., `rm -rf /`), this value will never be used — hence any dummy value will suffice. The argument to `system` is the address of attacker-supplied data, in this case a pointer to the string `/bin/sh`. Finally, the third figure displays the stack state after transferring control to `system`, which happened by popping `&system` into the program counter (and decrementing the stack pointer). At this point, the attacker can execute commands using the shell.

A more sophisticated version of arc injection is called *return-oriented programming* (ROP). It is based on the observations that the virtual memory space (which has the C library) offers many little code snippets, *gadgets*, that can be parsed as a valid sequence of instructions and end with a `ret` instruction. Recall that the `ret` instruction is equivalent to `popl %eip`, i.e., it writes the top of the stack into the program counter. The attacker does not even have to jump to the start of a function, any arbitrary location in the middle works as long as it terminates with a `ret`.

Shacham et al. showed that these small gadgets can be combined to perform arbitrary computation. In our above example, a basic combination of two gadgets would involve writing the starting address of the next gadget at the value of `dummy`. When the first gadget finishes, the next one is loaded by executing `ret`.

Setting up the stack is very tricky to get right manually, but the paper referenced above actually wrote a compiler to transform code from a language as expressive as C-like into mixture of gadgets to be pushed on the stack!

A final note: do not hesitate to ask for help! Our office hours exist to help you. Please visit us if you have any questions or doubts about the material.