

CS 489/698 Assignment 4

TA: Liyi Zhang (l392zhang@uwaterloo.ca)
Office hours: to be announced on Piazza posts

Due date: April 5, 2024

Escaping seccomp-based Sandboxes

Environment preparation

We provide a similar VM setup as you experienced with Assignment 1 and Assignment 3. To be specific, you can use the [portal.sh](#) script to create a new VM, destroy an existing VM, or ssh into a running VM. To learn more about portal.sh and how to use it, please visit <http://ugster72d.student.cs.uwaterloo.ca:8000>.

Unfortunately, in this assignment, we cannot re-use the VMs you have in A1 and A3 due to different system configurations needed. This means that you will need to **register again for the new platform with a new pair of keys**. Again, **take care of your private key!**

- Your initial status on the system is unregistered. Please run `./portal.sh register` first.
- If for whatever reasons you decide to reset your VM, the way to do it is through `./portal.sh destroy`. Please do not use `./portal.sh register` for this purpose.
- NOTE: if you lose your private key after successful registration, you will be locked out of the system completely. There is no way back. You can request a key reset on Piazza, and every reset means a **1 pt** deduction on the whole assignment.

After you register to the `ugster` platform and have your VM launched, please SSH into your VM and run the following command from the directory `/home/vagrant` to setup the VM for A4:

```
wget http://ugster72d.student.cs.uwaterloo.ca/a4/a4_setup.sh && \  
  chmod +x a4_setup.sh && ./a4_setup.sh
```

This will create four binary executable files `sandbox[1-4]` and a text file `flag` in `/home/vagrant`.

Tasks

We have provided four binary executable files: `sandbox1`, `sandbox2`, `sandbox3`, and `sandbox4`. Each of these programs implements a `seccomp`-based sandbox and employs some distinct filtering rules. The programs are designed to execute external programs `spl0it1`, `spl0it2`, `spl0it3`, and `spl0it4` respectively. Additionally, we have provided a `flag` file that stores a flag you required to get.

Your task is to write four exploit programs in C language, namely `spl0it1.c`, `spl0it2.c`, `spl0it3.c`, and `spl0it4.c`, and compile them into four binary executable files `spl0it1`, `spl0it2`, `spl0it3`, and `spl0it4`, respectively, in such a way that when running the sandbox programs, the contents of the `flag` file are written into standard output. You are expected to submit the four source code files of your exploit programs with a `Makefile`, such that we can get your compiled binary executable files by simply typing `make` in the terminal.

Below is an example demonstrating how to test your code and run the sandbox programs:

```
// compile your exploit programs
// (should produces spl0it[1-4])
$ make

// run the sandbox programs
// (your exploit programs should be in the same folder as the sandbox programs)
$ ./sandbox1
the flag should be written to standard output and displayed here

$ ./sandbox2
the flag should be written to standard output and displayed here

$ ./sandbox3
the flag should be written to standard output and displayed here

$ ./sandbox4
the flag should be written to standard output and displayed here
```

NOTE: When grading your submissions, we will replace the contents of the `flag` file. Therefore, **you should NOT hardcode the contents of the flag file in your program.**

Each `spl0it` carries an equal weight of **10 pts**, making a total of 40 pts.

Deliverables

All assignment submissions take place on LEARN dropbox.

You are required to hand in a single compressed `.zip` file that contains the following:

- **sploit1.c, sploit2.c, sploit3.c, sploit4.c:** The exploit programs.
- **Makefile:** A Makefile that compiles all your exploit programs.
- **write-up.pdf (optional):** An optional write-up in PDF format (see below).

Write-up. We use an auto-grader to check the code submitted for this assignment. While efficient, the auto-grader can only provide a binary pass/fail result, which rules out the possibility of awarding partial marks for each task. As a result, we also solicit a write-up submissions.

The write-up can include information about any of the exploits you attempted as long as you believe the information is relevant for the grading. Typical things to be put in the write-up include:

- How the code you submitted is expected to work
- How you manage to get certain hardcoded information in the code
- Explanation on critical steps / algorithms in the code
- Any special situations the TAs need to be aware when running the code
- If you do not complete the full task, how far you have explored

On the other hand, if you are confident that all your code will work out of the box and can tolerate a zero score for any tasks on which the auto-grader fails to execute your code, you do not need to submit a write-up (or you can omit certain tasks in the write-up).