

CSE 484 / 584M

Lab 1: Buffer Overflows

Checkpoint (Splits 1-3): Tuesday, April 17, 2018 at 4:30pm

Final Due: Wednesday, April 25, 2018 at 4:30pm

Goal

- The goal of this assignment is to gain hands-on experience with the effects of buffer overflow bugs. All of the work must be done on the machine `codered.cs.washington.edu` (see instructions below for connecting).
- You are given the source code for seven exploitable programs, whose binaries are stored in the “bin” directory (`/bin/target1, …, /bin/target7`). Each target program `[i]` is installed as `setuid hax0red[i]`. Your goal is to write seven exploit programs (`spl0it1, …, spl0it7`). Program `spl0it[i]` will execute program `/bin/target[i]`, giving it certain input that should result in a shell run with the same permissions as user `hax0red[i]`. If `target[i]` had `setuid root` then `spl0it[i]` would result in a root shell. We don't do that in this case for obvious security reasons, so instead you get the permissions of the `hax0red[i]` user.
- The skeletons for splits 1 through 7 are provided in the `~/splits/` directory. Note that the exploit programs are very short, so there is no need to write a lot of code here.
- Splits 1-7 are required. Split 8 is extra credit.

The Environment

- You will test your exploit programs on a remote machine running Debian Linux hosted at the domain `codered.cs.washington.edu`
- To connect to the machine, each group must first respond to the Lab 1 Google Form that was sent to the course mailing list. Please form groups of up to 3 people and only respond once per group with the following information:
 - 1) A username for your group
 - 2) The UW NetIDs of your group members.
 - 3) RSA **public** key (from a key pair that you control).
- Expect some delay between sending the information and account creation, so please answer the form early and plan accordingly. Please double check that your answers are free of typos to minimize any delays. Also, do **NOT** send the private key (that should remain only on trusted machines).

Here is an example of a public key:

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDAQDTKPi45wxeSezgO5JmG8HiuAQH6R3kqQTe
OeTbntWxliiClrahwlnkv26PAIaQKNdRbVH1fgX9kyUfsdj5JAvvNFuxpfY+GVVZKFI5M3Cuz
AynIymBjqnDn6Auq+tuSl8O4osb/0L9zDeQzOxQ+ed6iVDuPPkBLoX+XyuNUyYKV46xCIH
OS6ao+6CkZXhp4VTz4LUvb1s8DIUcaD8/bbigxxZH3eKRQH2arV9AqP1LoC2T3azLTkHvCrc
ImpjVW/pxf5+nbkRb1SSkkHFvFPdd+0us12yGOp1xBbo2kuKWSdcBgd4eiGHQsO+VWi23R9
2bcOh/DxRZumdMyaDBMGY/ user@localhost
```

Generating Key Pairs

Linux/Mac:

To generate the key-pair run the following (it is strongly suggested to use a passphrase):

```
ssh-keygen -t rsa -f <key_name>
```

To use ssh (after we have created your account):

```
ssh -i <path_to_private_key> <username>@codeder.cs.washington.edu
```

Windows:

To generate the key-pair use PuTTYgen. It comes installed with PuTTY.

- 1) Open PuTTYgen
- 2) Select the type of key as SSH-2 RSA
- 3) Click Generate and move the mouse around to generate entropy
- 4) (Optional but recommended) Enter a Passphrase + Confirmation of Passphrase
- 5) Click save the private key
- 6) Copy the text of the public key to post/email from the box at the top

To ssh (with PuTTY):

On the left side, select Connection->SSH->Auth. In this pane, browse to your private key, and then login as usual. You may want to save the session for a quicker login next time. (Note, if you generated your ssh key pairs using Linux and you want to use it in windows, you will need to use PuTTYgen to convert it from .pem to .ppk before using it)

The Targets

- The targets are stored in /bin and their corresponding sources in ~/sources/. You are free to study the source code of each target. DO NOT recompile the targets!
- Your exploits should assume that the compiled target programs are installed in /bin. Do not move the targets.
- Each target[i] is setuid hax0red[i], which means that they run as hax0red[i] regardless of who runs it. The one exception is when they're run under a debugger. Allowing users to debug a setuid executable is a security flaw, so setuid programs temporarily lose their setuid-ness under a debugger. This means that you can only get a hax0red[i] shell when your exploits are ran outside of gdb. However, if you get a user shell inside gdb, you should get a hax0red[i] shell outside of gdb.

The Exploits

The ~/sploits/ directory contains the source for the exploits which you are to write, along with a Makefile for building them. Also included is shellcode.h, which gives Aleph One's shellcode.

The Assignment

You are to write an exploit for each target #1-7. Each exploit, when run on the remote machine, should yield a hax0red[i] shell (/bin/sh). To confirm this is working, run the command `whoami` in the shell, and you should see the hax0red[i] user.

Extra Credit

Target8 is extra credit! You can see that the source code is exactly the same as target0, except this time, the stack is not executable. You might want to try a return2libc attack. Here's a good tutorial for it [RET2LIBC](#) (starting from page 52).

Hints

- Read Aleph One's "[Smashing the Stack for Fun and Profit.](#)" Carefully! We also recommend reading Chien and Szor's "[Blended Attacks](#)" paper. These readings will help you have a good understanding of what happens to the stack, program counter, and relevant registers before and after a function call, but you may wish to experiment as well. It will be helpful to have a solid understanding of the basic buffer overflow exploits before reading the more advanced exploits.
- Read scut's "[format strings](#)" paper. You may also wish to read <http://seclists.org/bugtraq/2000/Sep/214>.
- gdb (or cgdb, which gives you a view of the source code you're debugging; we also have `peda gdb` installed, if that's your thing, just add it to your `.gdbinit` with `echo "source /home/peda/peda.py" >> ~/.gdbinit`) is your best friend in this assignment,

particularly to understand what's going on. Specifically, note the "disassemble" and "stepi" commands. You may find the 'x' command useful to examine memory (and the different ways you can print the contents such as /a /i after x. If you use peda, telescope is a really useful command. To check the informations stored on stack. In peda, if you lost the original nice display, you can use context to tell it to reprint it.) The 'info register' command is helpful in printing out the contents of registers such as ebp and esp. The 'info frame' command also tells you useful information, such as where the return EIP is saved.

- A useful way to run gdb is to use the -e and -s command line flags; for example, the command `cgdb -e sploit3 -s /bin/target3 -d ~/sources` in the vm tells gdb to execute sploit3, use the symbol file in target3, and the -d shows you the source code of the target as you step through it. These flags let you trace the execution of the target3 after the sploit has forked off the execve process. When running gdb using these command line flags, be sure to first issue 'catch exec' then 'run' the program before you set any breakpoints; the command 'run' naturally breaks the execution at the first execve call before the target is actually exec-ed, so you can set your breakpoints when gdb catches the execve. Note that if you try to set breakpoints before entering the command 'run', you'll get a segmentation fault.
- If you wish, you can instrument your code with arbitrary assembly using the `asm ()` pseudo function.
- Make sure that your exploits work within the remote environment we provided.
- **Start early!!!** Theoretical knowledge of exploits does not readily translate into the ability to write working exploits. Target1 is relatively simple and the other problems are quite a bit more complicated.

Warnings

Aleph One gives code that calculates addresses on the target's stack based on addresses on the exploit's stack. Addresses on the exploit's stack can change based on how the exploit is executed (working directory, arguments, environment, etc.); in our testing, we do not guarantee to execute your exploits as bash does. **You must therefore hard-code target stack locations in your exploits.** You should **not** use a function such as `get_sp()` in the exploits you hand in.

Deliverables

- You may work in groups of up to **three** people. Make sure your submission includes the name of all your group members.
- **In a bid to get you to start early, spoils 1-3 are due by April 17, 2018 at 4:30pm.**
- Since we have access to your remote home directory, you won't need to submit any code. However, to let us know when you're done, please submit a text file named `<netid_member_1>_<netid_member_2>_<netid_member_3>.txt` (make sure to use your uw netid, not student number!) with the result of running `md5sum sploit[i].c` for all spoils in your spoils directory.

- E.g., for `spl0it1.c`, you would type `md5sum spl0it1.c` in your terminal, and copy the entire line into the text file.
- Put one md5sum per line and for check point, you should submit three md5sums and for final project due, you should submit md5sums for all spl0its you did. For example, your final submission for lab 1 should look like:

```
1. 4737d20c261692df595200844a3512b2 spl0it1.c
2. 11d7dce12f121e129eaec3d050f13b06 spl0it2.c
3. cb3ce1e6f810f189f7d51317ad995bc2 spl0it3.c
4. 6e16373a1782f1f4b49c04c1a4b32146 spl0it4.c
5. 85e0b36652ff26d043540037f4b66220 spl0it5.c
6. 675a0b52944aaee683c97379bd08919d spl0it6.c
7. 3635c9af3f463668617dbcff0152e685 spl0it7.c
```

- **DO NOT INCLUDE ANYTHING ELSE** in your text file because we will use script to compare the md5 with what you have in the home directory.
- **AFTER SUBMITTING THIS, DO NOT CHANGE YOUR CODE** in your home directory -- make copies if you'd like to experiment further. Otherwise, the md5 hashes will not match when we check it. If you're concerned you might mess this up, you're welcome to submit your code to the dropbox as well, as a backup. But submitting code is not required.
- Turn in your text file online using the Canvas. One submission per group is enough. (All members can submit and we will only grade the latest one)

Misc

- **Please try to access the remote machine early and let us know if you have any problems! Send us a reminder email if you don't have access within 24 hours of sending your public key. But please do allow 24 hours :)**
- You may wish to backup or write your code elsewhere. We suggest using SCP or SFTP to access your files. For Windows, WinSCP is a great tool. SCP and SFTP run on top of SSH, so use your SSH parameters (port, key, etc.) to connect.
- There's lots of online documentation for GDB. Here's one you might start with: [GDB Notes \(formerly hosted at CMU\)](#)
- The "crash course in x86 and gdb" slides: [section_lecture.pdf](#)

Credits

This project was originally designed for Dan Boneh and John Mitchell's CS155 course at Stanford, and was then also extended by Hovav Shacham at UCSD. Thanks Dan, John, and Hovav!