

Computer Security

Coursework Exercise CW1

Web Server and Application Security

In this coursework exercise we will guide you through an attack against a vulnerable machine. You will take the role of “Mallet” attacking a web site that was programmed by “Alice”. The successful run through this attack will require some reading about tools and attacks. The questions are intended to encourage you to experiment with the given machines and tools, to give a better understanding of security mechanisms and exploits.

This coursework is worth 50 points in total and accounts for 12.5% of the final mark for this course.

Preparation

The Virtual Machines. For this exercise we are using Virtual Machines (VMs) inside VirtualBox. This provides a convenient way to experiment with systems that might be vulnerable and to test them against known attacks. The VMs we are using come preconfigured and preinstalled with useful tools. The victim is called Alice, and the attacker is called Mallet; they each have a machine. The two virtual machines we are using in this exercise are accessible from all DICE machines in `/group/teaching/cs/`, in a compressed *virtual appliance* format (about 900Mb each). The directory is an AFS mount so you may need to type the full path into a command line or into the folder GUI to see the files. Tab completion will not work and the containing folders will appear empty.

Notations for names. We will use names like “alice” in several ways. In the examples we use we introduce one protagonist, which serves as the victim; her name is Alice, written with a capital A. The victim Alice has a computer; the name of her system is **alice** and is printed bold except in user inputs. Finally on the machine **alice** there is a user account named *alice* which is printed italic except in user inputs. The same notations apply for Mallet.

Setup The machines come as *.ova files, which include the virtual hard drive as well as all necessary settings for VirtualBox. VirtualBox can be downloaded at <https://www.virtualbox.org/> and is already installed on all DICE computers. Once VirtualBox is running, select *File*→*Import Appliance* and select the file `alice.v1.1.ova` or `mallet.v1.1.ova`. The only property in the configuration screen you should change is the location of the virtual hard drive. On DICE computers a home directory account does not have enough quota to store the hard drive, so you will need to store the virtual disk locally on the workstation. The disks take up about 2.3G each once expanded. Go to the last line *Hard Disk Controller*→*Virtual Disk Image*. Change the path from, e.g.,

```
/afs/inf.ed.ac.uk/user/sXX/sXXXXXX/VirtualBox VMs/mallet.1/mallet.v1.1-disk1.vmdk
```

to

```
/tmp/sXXXXXX/VirtualBox VMs/mallet.1/mallet.v1.1-disk1.vmdk
```

(you may need to create the directory `/tmp/sXXXXXX` first). The import operation can take several minutes and need about 3G of disk space per machine.

IMPORTANT: On DICE machines the virtual disk has to be stored on the local disk. That means that you will not be able to access the virtual machines on any other DICE computer than the one you did this setup on. If you want to use your virtual machines from another computer you will have to log in to the machine using `ssh -X <computername>`, or copy the disk files to another machine with `scp`; so make sure you remember the name of the computer you used!

Exercise

We will go through an attack against the computer **alice**. To begin, you will need to run both virtual machines **alice** and **mallet** at the same time. Once **mallet** is started, login with the user *mallet* and password **mallet**. You will find a standard Gnome desktop environment. On **alice** you can login as *alice* with the password **alice**. However, Part A only requires **alice** to be running, a login is not necessary yet.

Part A: Scanning (10 points)

The first thing we want to do is to have a look at the victim's system and its vulnerabilities. To do this attackers can use freely available software tools. One of these tools is **Nmap**, the *Network Mapper*, which is installed on **mallet**. More information for the command can be obtained on the man-page. Some of the functionalities of Nmap require root privilege. In this case it helps to execute the command as `sudo nmap <parameters>` and provide *mallet*'s password **mallet**.

1. From **mallet** run a scan of **alice** with the tool `nmap`. Make sure to include the parameters `-O` and `--version-all -sV` to get as much information about the system **alice** as possible.

```
sudo nmap -O --version-all -sV alice
```

2. List all the open port numbers and the MAC address.
3. Nmap offers a variety of different scanning methods. In particular: TCP SYN scan, TCP ACK scan and the XMAS scan. Explain the differences and the advantages of using each of these three scanning approaches.

After this first gathering of information we proceed with a more automatic method. The program **OpenVAS** is a collection of tools to automatically scan for vulnerabilities. It consists of a daemon and a user interface. To start the daemon type `sudo openvasd` into the command-line. Once it is loaded start the user interface by typing `openvas-client` into another command-line-window. All options are already adjusted properly. Note: Do not connect to the openVAS-server without creating the task as described below or the client will freeze.

4. Perform a full scan of the system **alice**: Hit the *Scan assistant* button “?” at the upper left corner, provide reasonable descriptions for the task and the scope and hit *Execute*. In the next window you will be asked to provide the password for the server, which is again user *mallet* and password **mallet**. This test may take a while and cause virtual box to temporarily stop responding. Report back the 4 most dangerous security risks found and explain why they are the most dangerous.

Turn In: Answers for 2, 3, and 4.

Part B: Session stealing (10 points)

To find and exploit a vulnerability on a website it is helpful to know the structure of the site. That is, to know what the different pages are and how they communicate with one another. To do this, Mallet has already created a user account for herself on the web application. The username is *mallet* and the password is `mallet123`.

Clarification (10 Feb): Use the web page at `http://alice/`

1. Explore the entire website by clicking on all the links and buttons to see what they do.
2. For each page transition (link or button) on the website, provide the information below. Some page transitions have different behavior based on the state the page is in. So you may need to provide more than one record for each link to fully describe the behavior of the site. It may help to draw the behavior of the page as a state transition diagram.
 - Page State: Does this link or button's behaviors depend on earlier events or actions. (For example, buttons behave differently if the user is logged out vs. logged in.)
 - Link Text: The text on the link or button. (For example: "Home".)
 - Data: What data is being transferred between the pages and how it is being transferred (GET or POST).
 - Effect: Describe any effect clicking the link has. (For example: causing a message to appear on the list of messages.)

Example record:

- Page State: Initial start page with no content and not logged in.
- Link Text: "Home"
- Data: No data is sent
- Effect: Page reloads with no change.

You should turn in either: a) a state transition diagram which contains all the information described above, or b) a set of text records containing all the information described above with at least one blank line between the records.

Having collected some information data, we are now going to try *black-box* attacks against **alice**. That means we are not assuming any knowledge about the code itself. Since this application is programmed in PHP this is a reasonable assumption; PHP code is server-side code and thus never delivered to the user's machine.

Our first attack point is the transmission of a new message. This is done using the URL and the GET method of HTML. We can see that the information included into the URL after pressing the submit button on the message board includes the message in plain-text, the screen that should be displayed and a parameter that is called 'sid'. Not included are for example the name of the user, which is probably stored in a database at the server. It is not hard to guess that the 'sid' parameter is the session id, which links the request to the user in the database on the server. Note that it is not verifiable whether the URL sent to the server (including the data) was generated by the HTML form, or simply typed in directly by the user.

3. By logging in and out several times find out how the session id is computed. Describe the algorithm used to compute sid when a user logs in (do not over think this).

Sending data to a server using the session id of another user is called "session stealing".

4. Steal another session: On the machine alice log into the web page `http://alice` as *alice* with the password `alice123`. Now get back to **mallet** and forge a URL with data to publish a message as *alice*. Describe all details of your method. What are the preconditions and limitations of this attack?

Turn In: Answers to 2, 3, and 4 above.

Part C: SQL Injection (15 points)

Another category of attack is the so called *white-box attack*. For this type of attack we assume we know at least a part of the source code. For this exercise we assume that Alice carelessly gave away the SQL query containing the lookup for the password. It is:

```
SELECT id FROM users WHERE username='$uname' AND password = '$pwd'
```

The variables \$uname and \$pwd contain the user input from the two input fields in the login form. Alice does not check them further but just passes them to the SQL command. This gives us the chance to insert code into this command. SQL code entered into the input fields will be inserted into this SQL command and executed. This attack is called SQL injection. A usual injection string consists of the following parts:

- characters to close the opened environments (e.g. ' or closing brackets)
- the malicious code
- characters to prevent syntax error resulting from the injection: Either reopen all necessary environments or hide following characters as comment using the sequence " -- " (note the blanks).

As an example entering the username `dont know' OR '1'='1' --` will result in the query being

```
SELECT id FROM users WHERE username='dont know' OR '1'='1' -- ' AND password = '$pwd'
```

which will be true for every user in the users-table.

1. Perform the described SQL injection. Describe the result and give a possible explanation.
2. Break confidentiality: Describe a sequence of login credentials, that will give you the list of all user names on this forum. Write down your input strings and give a short explanation what happens inside the system.
3. Break availability: Now that we can inject arbitrary SQL commands we can tamper with data that should not be accessible to any user, for example with the user database. Change the password for the user account *alice*. Write down the inputs you used and explain in detail why this works. (Hint: Look at http://www.w3schools.com/sql/sql_update.asp if you need help with the SQL commands.)

Turn In: Answers to 1, 2, and 3 above.

Part D: Firewalls (15 points)

Warning do this part AFTER completing parts A, B, and C. You will be configuring a firewall on **alice**. If you create the wrong set of rules it will block the web server on **alice** and it will become challenging to complete the earlier parts.

Alice uses her computer for two things:

- Using SSH to connect to her computer remotely (ssh).
- As a server for her website (http).

The earlier port scan showed that **alice** has more ports open and responding than are strictly necessary to do the two things above. Help Alice setup her iptables firewall so that it only allows http and ssh connections. When you are done, the website `http://alice/` should be accessible from Firefox on both **mallet** and **alice**. It should be possible to ssh into **alice** from **mallet** using *alice*'s username and password. Finally, an nmap scan of **alice** from **mallet** should show only two open ports (http and ssh).

Similar to Part A, the iptables command requires root privileges, so you will need to execute the command as `sudo iptables`.

You are welcome to use any online resource to complete this part of the exercise. The tutorial at the link below is recommended as is reading the man page for iptables.

```
https://help.ubuntu.com/community/IptablesHowTo
```

Turn In: The list of iptables commands you used to correctly setup the firewall and the output of the command: `sudo iptables-save`.

The following commands are also useful.

To see the current iptables rules use:

```
sudo iptables -L
```

To reset the rules to their default empty state use:

```
sudo iptables -f INPUT
```

To ssh into **alice** from **mallet** use:

```
ssh alice@alice
```

Submission Instructions

You should turn in one plain text file with sections for each part of the assignment. This is so that it is possible to see which part of the file answers which question. Do not submit a Word document or a PDF file - only submit a plain text file.

IMPORTANT: You must create your text file on your local computer, NOT the VM. The VMs we are using are isolated from the host system for everybody's safety. There is no easy way to copy files between the VM and the local computer. The submit command will not work from within the VM. You must use the local command line on a DICE computer to use submit.

Your text file should be named `cw1.txt`. Organize your file into sections for each part of the exercise and clearly label each answer with the number associated with it. An example submission file should look like this:

```
####  
# Part A  
####
```

```
Question 1  
Answer
```

```
Question 2  
Answer
```

```
...
```

```
####  
# Part B  
####  
...
```

Use the submit program to submit the cw1.txt text file. You may submit the assignment as many times as you like before the deadline. To submit from the command line on a DICE machine use the following command:

```
submit cs 1 cw1.txt
```

If you choose to submit an image file for Part B question 2, then name the file cw1-b2.jpg. This is the only question for which we will accept an image. Separately submit the image file using the following command:

```
submit cs 1 cw1-b2.jpg
```

You need to submit by the deadline of **February 26th at 16:00**.

If for any reason you are unable to use the submit command you may also submit to the ITO in person on hardcopy, either a printed document or a solution in clearly legible handwriting marked with your matriculation number at the top of each page.

You're reminded that *late coursework* is not allowed without "good reason", see

<http://www.inf.ed.ac.uk/teaching/years/ug3/CourseGuide/coursework.html>

for more details about this, and the procedure to follow if you must submit late. In particular, if you have a good reason to submit late, please use the ITO support form <http://www.inf.ed.ac.uk/admin/ITO/support/index.html> to make a request.

Good Scholarly Practice:

Please remember the University requirement as regards all assessed work for credit. Details about this can be found at:

<http://www.ed.ac.uk/academic-services/students/undergraduate/discipline/academic-misconduct>
and at:

<http://web.inf.ed.ac.uk/infweb/admin/policies/academic-misconduct>

Furthermore, you are required to take reasonable measures to protect your assessed work from unauthorised access. For example, if you put any such work on a public repository then you must set access

permissions appropriately (generally permitting access only to yourself, or your group in the case of group practicals).

1

¹Exercise originally created by Daniel Franzen and David Aspinall. Modified by Kami Vaniea.