Heartbleed Attack Lab

Modified for COMP435, Fall 2017 by Cynthia Sturton, UNC-Chapel Hill.

1 Overview

The Heartbleed bug (CVE-2014-0160) is a severe implementation flaw in the OpenSSL library, which enables attackers to steal data from the memory of the victim server. The contents of the stolen data depend on what is there in the memory of the server. It could potentially contain private keys, TLS session keys, user names, passwords, credit cards, etc. The vulnerability is in the implementation of the Heartbeat protocol, which is used by SSL/TLS to keep the connection alive.

The objective of this lab is for students to understand how serious this vulnerability is, how the attack works, and how to fix the problem. The affected OpenSSL version range is from `1.0.1` to `1.0.1f`. The version in our Ubuntu VM is `1.0.1`.

2 Lab Environment

In this lab there are two machines: the attacker client and the victim server. You can set up two VMs that can communicate, or you can use a single machine that runs both the server process and the client code. In either case you'll use the pre-built `SEEDUbuntu12.04` VM. The instructions for using two separate VMs follow. (There are no special instructions needed for using a single VM, just create the VM and you are ready.)

The website used in this attack can be any HTTPS website that uses SSL/TLS. However, since it is illegal to attack a real website, we have set up a website in our VM, and conduct the attack on our own VM. We use an open-source social network application called `ELGG`, and host it in the following URL: `https://www.heartbleedlabelgg.com`. Warning: Do not under any circumstance, attack real websites.

Important Note: If you have updated the version of OpenSSL installed on the VM (you might have done this while working on the Crypto Hash Lab), this lab won't work; the newest version of OpenSSL has patched the Heartbleed vulnerability. You can reinstall the VM to revert back to the earlier version of OpenSSL.

Using Two VMS You will create two VMs, one called the attacker machine and one called the victim server. The VMs need to use the `NAT-Network` adapter for the network setting. This can be done by going to the VM settings, picking Network, and clicking the Adaptor tag to switch the adapter to `NAT-Network`. Make sure both VMs are on the same NAT-Network.

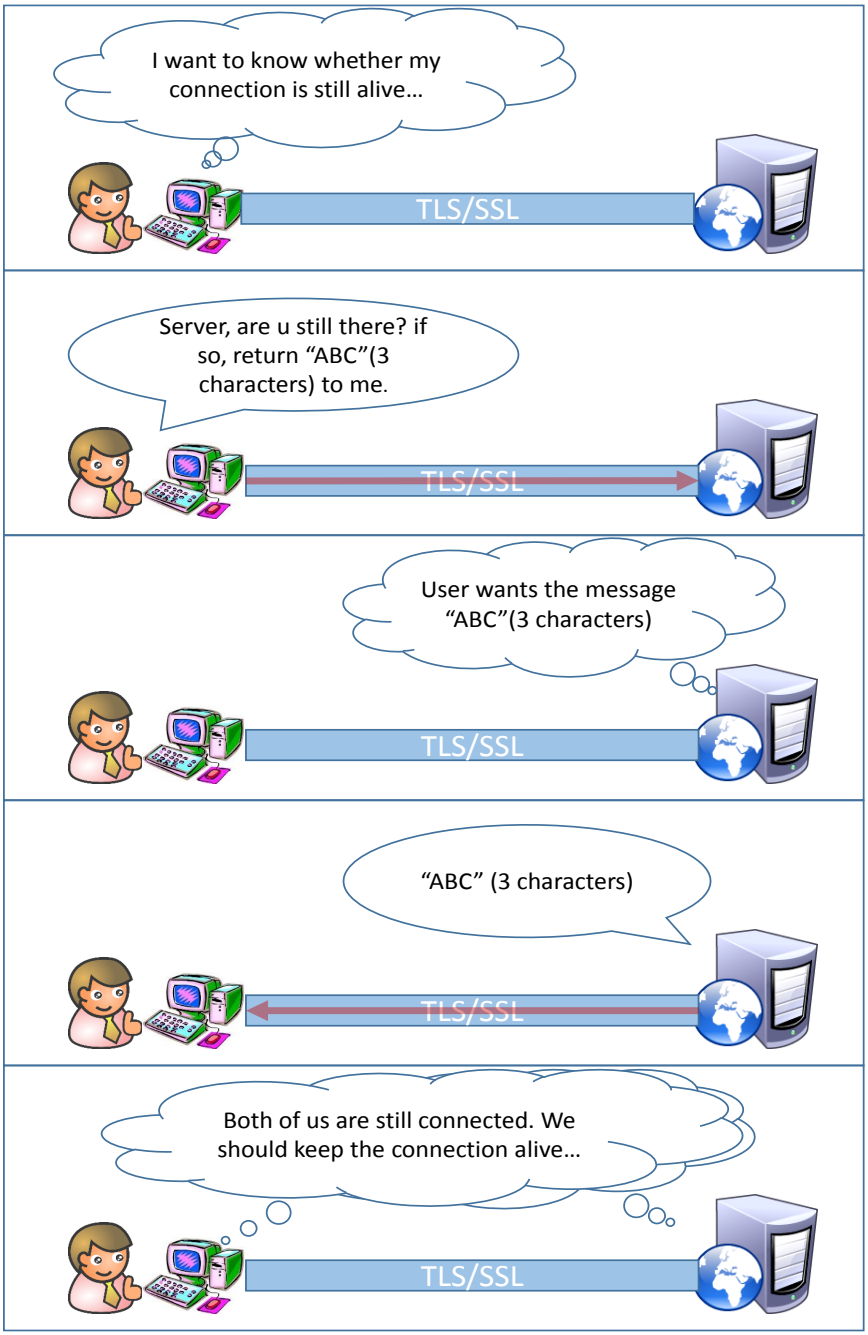

Figure 1: Overview of the Heartbeat Protocol

We need to modify the `/etc/hosts` file on the attacker machine to map the server name to the IP address of the server VM. Search the following line in `/etc/hosts`, and replace the IP address `127.0.0.1` with the actual IP address of the server VM that hosts the `ELGG` application.

```
127.0.0.1 www.heartbleedlabelgg.com
```

3 Lab Tasks

Complete the following tasks and submit your observations and screenshots in a PDF file, which you will submit on Sakai. Your answers should be short. You will be limited in some cases to answers of roughly 2-4 sentences. Your screenshots should be small, less than 1 MB.

Before working on the lab tasks, you need to understand how the heartbeat protocol works. The heartbeat protocol consists of two message types: HeartbeatRequest packet and HeartbeatResponse packet. Client sends a HeartbeatRequest packet to the server. When the server receives it, it sends back a copy of the received message in the HeartbeatResponse packet. The goal is to keep the connection alive. The protocol is illustrated in Figure 1.

3.1 Task 1: Launch the Heartbleed Attack.

In this task, students will launch the Heartbleed attack on our social network site and see what kind of damages can be achieved. The actual damage of the Heartbleed attack depends on what kind of information is stored in the server memory. If there has not been much activity on the server, you will not be able to steal useful data. Therefore, we need to interact with the web server as legitimate users. Let us do it as the administrator, and do the followings:

1. Visit `https://www.heartbleedlabelgg.com` from your browser.
2. Login as the site administrator. (User Name:`admin`; Password:`seedelgg`)
3. Add Boby as friend. (Go to `More -> Members` and click `Boby -> Add Friend`)
4. Send Boby a private message.

Include your CS login in any content you send or post. For example, in a private message to Boby, include your CS login in the subject line, message data, or both. You will be asked to submit screenshots of the data you recover, and as proof that you completed the task (rather than reusing someone else's screenshot), we will look for your login to appear in any message data.

After you have interacted with the server as a legitimate user, you can launch the attack and see what information you can get out of the victim server. Writing the program to launch the Heartbleed attack from scratch is not easy, because it requires the low-level knowledge of the Heartbeat protocol. Fortunately, other people have already written the attack code. Therefore, we will use the existing code to gain first-hand experience in the Heartbleed attack. The code that we use is called `attack.py`, which was originally written by Jared Stafford. We made some small changes to the code for educational purposes. You can download the code from the lab's web site (`http://www.cis.syr.edu/~wedu/seed/Labs_12.04/Networking/Heartbleed/attack.py`), change its permission so the file is executable. You can then run the attack code as follows:

```
$ ./attack.py www.heartbleedlabelgg.com
```

You may need to run the attack code multiple times to get useful data. Your task is to retrieve the following information from the target server.

1. Admin user name and password.
2. User's activity (what the user has done).
3. The exact content of a private message.

Submit a screenshot: For each piece of secret data that you steal from the Heartbleed attack, submit a screenshot showing the attack successfully revealing the data. Make sure your CS login is part of the content of the private message revealed by your attack.

3.2 Task 2: Find the Cause of the Heartbleed Vulnerability

In this task, students will compare the outcome of the benign packet and the malicious packet sent by the attacker code to find out the fundamental cause of the Heartbleed vulnerability.

The Heartbleed attack is based on the Heartbeat request. This request just sends some data to the server, and the server will copy the data to its response packet, so all the data are echoed back. In the normal case, suppose that the request includes 3 bytes of data "ABC", so the length field has a value 3. The server will place the data in the memory, and copy 3 bytes from the beginning of the data to its response packet. In the attack scenario, the request may contain 3 bytes of data, but the length field may say 1003. When the server constructs its response packet, it copies from the starting of the data (i.e. "ABC"), but it copies 1003 bytes, instead of 3 bytes. These extra 1000 types obviously do not come from the request packet; they come from the server's private memory, and they may contain other user's information, secret keys, password, etc.

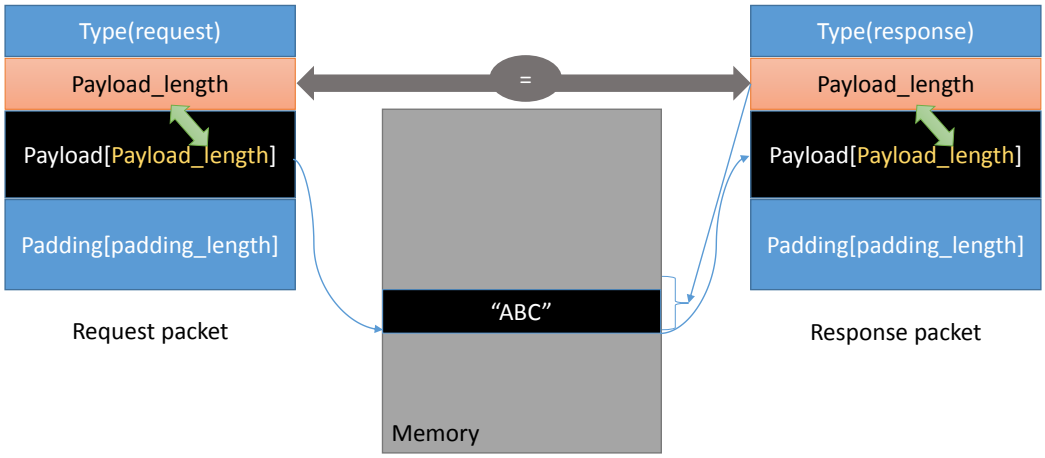

Figure 2: The Benign Heartbeat Communication

In this task, we will play with the length field of the request. First, let's understand how the Heartbeat response packet is built from Figure 2. When the Heartbeat request packet comes, the server will parse the packet to get the payload and the `Payload_length` value (which is highlighted in Figure 2). Here, the payload is only a 3-byte string `"ABC"` and the `Payload_length` value is exactly 3. The server program will blindly take this length value from the request packet. It then builds the response packet by pointing to the memory storing `"ABC"` and copy `Payload_length` bytes to the response payload. In this way, the response packet would contain a 3-byte string `"ABC"`.

We can launch the HeartBleed attack like what is shown in Figure 3. We keep the same payload (3 bytes), but set the `Payload_length` field to 1003. The server will again blindly take this `Payload_length` value when building the response packet. This time, the server program will point to the string `"ABC"` and copy 1003 bytes from the memory to the response packet as a payload. Besides the string "ABC", the extra 1000 bytes are copied into the response packet, which could be anything from the memory, such as secret

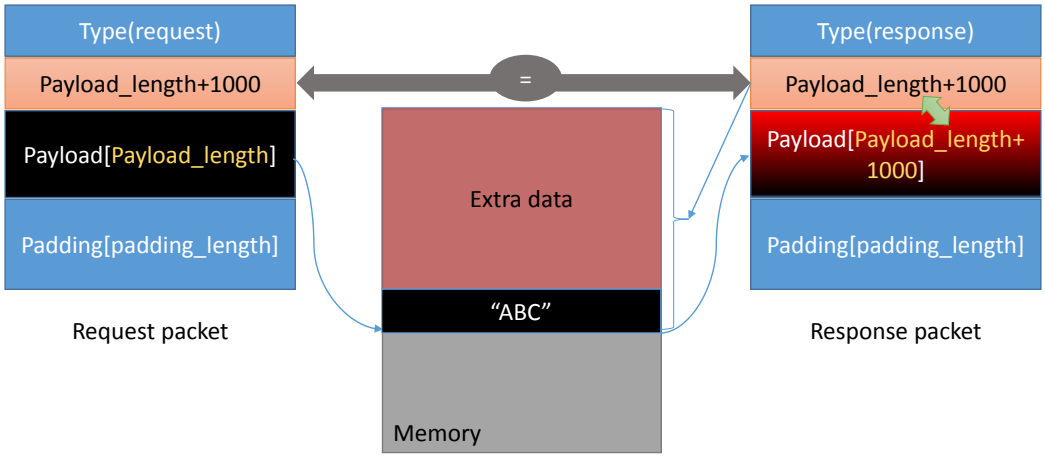

Figure 3: The Heartbleed Attack Communication

activity, logging information, password and so on.

Our attack code allows you to play with different `Payload_length` values. By default, the value is set to a quite large one (`0x4000`), but you can reduce the size using the command option `"-l"` (letter ell) or `"--length"` as shown in the following examples:

```
$./attack.py www.heartbleedlabelgg.com -l 0x015B
$./attack.py www.heartbleedlabelgg.com --length 83
```

Your task is to play with the attack program with different payload length values and answer the following questions:

Question 2.1: As the length variable decreases, what kind of difference can you observe?

Question 2.2: As the length variable decreases, there is a boundary value for the input length variable. At or below that boundary, the Heartbeat query will receive a response packet without attaching any extra data (which means the request is benign). Please find that boundary length. You may need to try many different length values until the web server sends back the reply without extra data. To help you with this, when the number of returned bytes is smaller than the expected length, the program will print `"Server processed malformed Heartbeat, but did not return any extra data."`

3.3 Task 3: Countermeasure and Bug Fix

In this task you will implement the best-practice countermeasure (patching the bug) and describe how the patch works.

3.3.1 Task 3.1

To fix the Heartbleed vulnerability, the best way is to update the OpenSSL library to the newest version. This can be achieved using the following commands. It should be noted that once it is updated, it is hard to

go back to the vulnerable version. Therefore, make sure you have finished the previous tasks before doing the update. You can also take a snapshot of your VM before the update.

```
#sudo apt-get update
#sudo apt-get upgrade
```

Submit a screenshot: Try your attack again after you have updated the OpenSSL library. Take a screenshot of what you observe trying the attack after upgrading OpenSSL.

3.3.2 Task 3.2

The objective of this task is to figure out how to fix the Heartbleed bug in the source code. Below we present the sourcecode that was introduced in December 2011. You may view the exact commit which introduced the bug here: https://github.com/openssl/openssl/commit/4817504d069b4c5082161b02a22116ad75f822b1

```
struct ssl3_record_st
{
  unsigned int length;    /* How many bytes available */
  [...]
  unsigned char *data;    /* pointer to the record data */
  [...]
} SSL3_RECORD;

struct
{
  HeartbeatMessageType type; /* 1 byte: request or the response */
  uint16 payload_length;     /* 2 bytes: the length of the payload */
  opaque payload[HeartbeatMessage.payload_length];
  opaque padding[padding_length];
} HeartbeatMessage;
```

The `SSL3_RECORD` is a basic building block of SSL communications. It has a `length` field giving the length of the received message (in this case, a Heartbeat message) and a pointer (`data`) to the actual message. The message is formatted as a `HeartbeatMessage`. This struct has a `payload_length` field which gives the (client-provided) length of the payload. The `payload` field is the client-provided data that should be sent back in the Heartbeat response.

The following code snippet shows how the server copies the data from the request packet to the response packet.

Listing 1: Process the Heartbeat request packet and generate the response packet

```
1 /* Allocate memory for the response, size is 1 byte
2  * message type, plus 2 bytes payload length, plus
3  * payload, plus padding
4  */
5
6 unsigned int payload;
7 unsigned int padding = 16; /* Use minimum padding */
8
```

```
/* Read from type field of HeartbeatMessage first.
 * After this instruction, the pointer
 * p will point to the payload_length field of HeartbeatMessage. */
hbtype = *p++;


/* Read from the payload_length field of HeartbeatMessage.
 * This function reads 16 bits from pointer p and stores
 * the value in the local INT variable "payload". */

n2s(p, payload);

pl=p; // pl points to the beginning of the payload content

if (hbtype == TLS1_HB_REQUEST)
{
    unsigned char *buffer, *bp;
    int r;

    /* Allocate memory for the response, size is 1 byte
     * message type, plus 2 bytes payload length, plus
     * payload, plus padding
     */

    buffer = OPENSSL_malloc(1 + 2 + payload + padding);
    bp = buffer;

    // Enter response type and length
    *bp++ = TLS1_HB_RESPONSE;
    s2n(payload, bp);

    /* Copy payload
     * pl is the pointer that points to the beginning
     * of the payload content */

    memcpy(bp, pl, payload);
    bp += payload;

    // Random padding
    RAND_pseudo_bytes(bp, padding);

    // this function will copy the 3+payload+padding bytes
    // from the buffer and put them into the heartbeat response
    // packet to send back to the request client side.
    OPENSSL_free(buffer);
    r = ssl3_write_bytes(s, TLS1_RT_HEARTBEAT, buffer,
         3 + payload + padding);
}
```

Describe a solution: Please point out the problem from the code in Listing 1 and describe a solution to fix the bug (i.e., what modification is needed to fix the bug). You do not need to recompile the code; just describe how you can fix the problem in your lab report. Your answer should include a snippet of C code (pseudo code will suffice), and a description of where the new code should be placed.

4 Submission

On Sakai in a pdf document titled: **a4_youronyen.pdf** submit the following items:

1. Three screenshots of the attack revealing the admin's user name and password, some user activity, and the exact content of a private message. The private message must contain your CS login.
2. Task 2 Question 2.1: Describe your observation. What difference do you observe as the variable length decreases?
3. Task 2 Question 2.2: What is the boundary at or below which no extra data is returned? Provide your answer as a decimal (not hex) value
4. Task 3.1: Screenshot of attack output after updating the OpenSSL library
5. Task 3.2: Describe how to fix the code. Your answer should include a snippet of code (pseudo code will suffice), a description of how the new code fixes the vulnerability, and instructions for where the new code should be placed.

References

[1] Heartbleed attack - Implementation: https://alexandreborgesbrazil.files.wordpress.com/2014/04/hearbleed_attack_version_a_1.pdf

[2] Heartbleed attack - Interesting explanation: http://xkcd.com/1354/