

CT437

COMPUTER SECURITY AND FORENSIC COMPUTING

BUFFER OVERFLOW CASE STUDY – THE HEARTBLEED BUG

Dr. Michael Schukat



OÉ Gaillimh
NUI Galway

A Bug with its own Website (heartbleed.com) and Icon

2

The Heartbleed Bug

The Heartbleed Bug is a serious vulnerability in the popular OpenSSL cryptographic software library. This weakness allows stealing the information protected, under normal conditions, by the SSL/TLS encryption used to secure the Internet. SSL/TLS provides communication security and privacy over the Internet for applications such as web, email, instant messaging (IM) and some virtual private networks (VPNs).

The Heartbleed bug allows anyone on the Internet to read the memory of the systems protected by the vulnerable versions of the OpenSSL software. This compromises the secret keys used to identify the service providers and to encrypt the traffic, the names and passwords of the users and the actual content. This allows attackers to eavesdrop on communications, steal data directly from the services and users and to impersonate services and users.



What leaks in practice?

We have tested some of our own services from attacker's perspective. We attacked ourselves from outside, without leaving a trace. Without using any privileged information or credentials we were able to steal from ourselves the secret keys used for our X.509 certificates, user names and passwords, instant messages, emails and business critical documents and communication.

How to stop the leak?

As long as the vulnerable version of OpenSSL is in use it can be abused. [Fixed OpenSSL](#) has been released and now it has to be deployed. Operating system vendors and distribution, appliance vendors, independent software vendors have to adopt the fix and notify their users. Service providers and users have to install the fix as it becomes available for the operating systems, networked appliances and software they use.

X X X

Overview Heartbleed

3

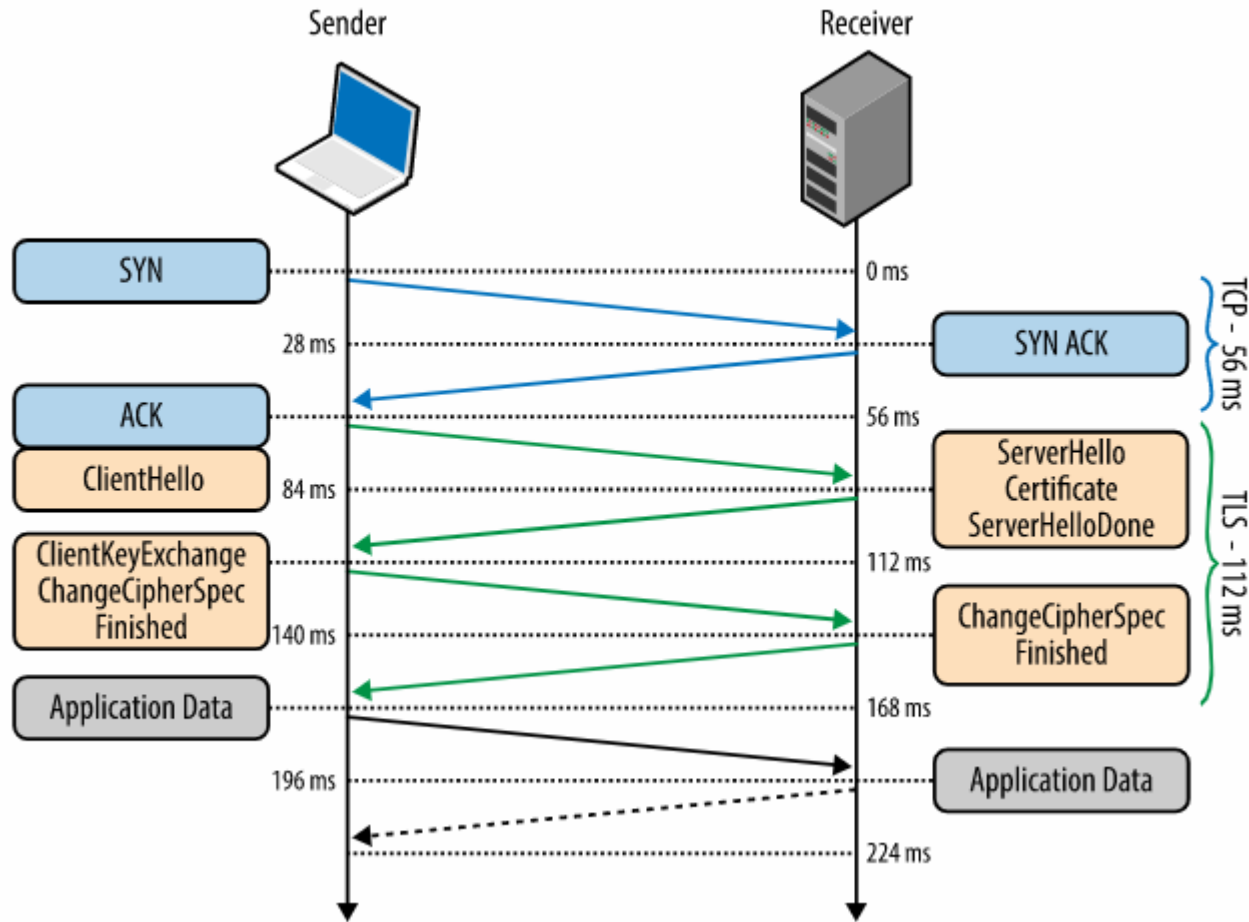
- ❑ Discovered in 2014
- ❑ Exploits a bug in the OpenSSL implementation of the TLS “heartbeat hello” extension
- ❑ Can affect both client and server side



X X

Recap TLS 1.2 Handshake (Server Authentication only)

4



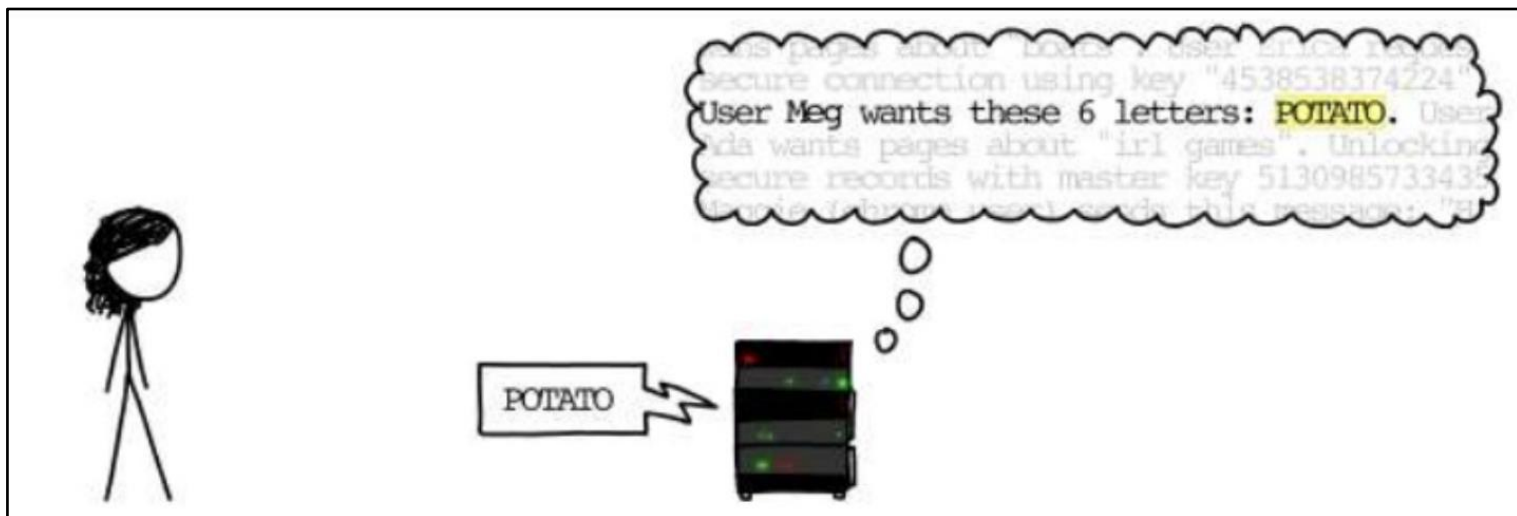
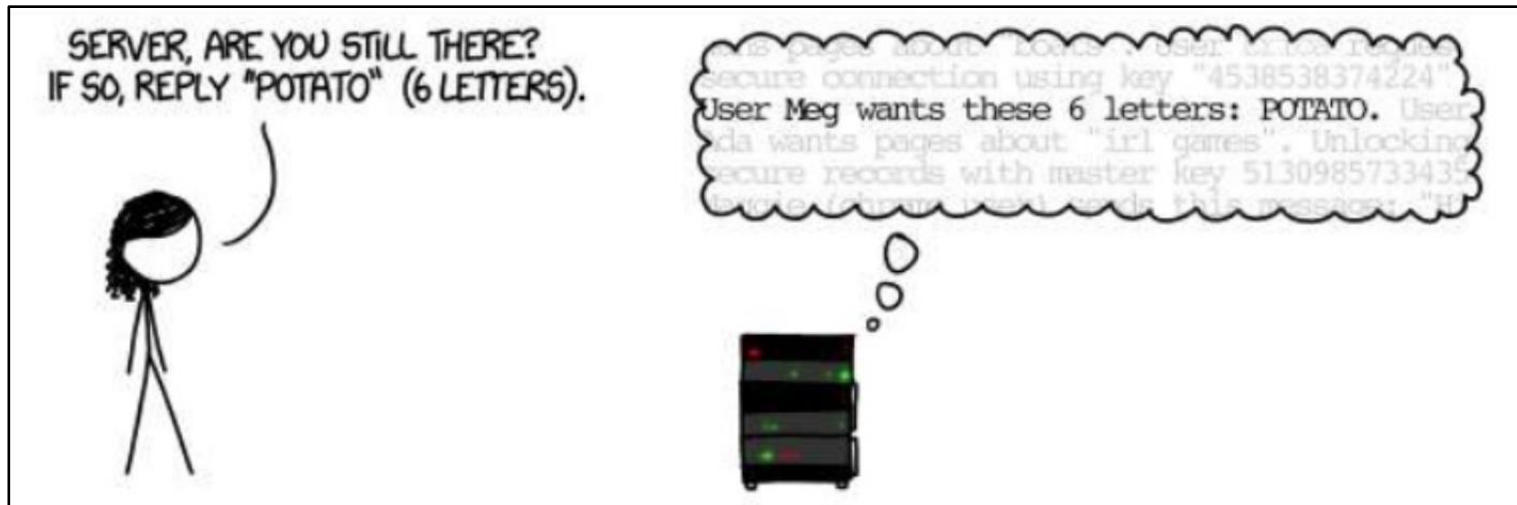
TLS Heartbeat Extension

5

- ❑ Originally TLS had no provisions to keep a client / server connection alive without continuous data transfer
 - ▣ Idle connections would timeout instead and a (computationally) expensive handshake (224 ms in the previous example) or a reconnect would have to take place
- ❑ The heartbeat extension provides a protocol for “keep-alive” messages that prevent a timeout
 - ▣ One endpoint could send out a *HeartbeatRequest* message, which would be immediately responded with a *HeartbeatResponse* message

Heartbeat with incoming Message (correctly) buffered

6



Heartbeat Request / Response Message

7

- The Heartbeat protocol messages consist of their type and an arbitrary payload and padding.

heartbeat_request or heartbeat_response

- struct {
HeartbeatMessageType type;
uint16 payload_length;
opaque payload[HeartbeatMessage.payload_length];
opaque padding[padding_length];
} HeartbeatMessage;

16+ bytes of random
content, ignored by receiver

- The sender composes a request message containing a payload with a specified length (i.e. *payload_length*)
- The receiver returns a response message containing a copy of the sender's payload (with length *payload_length*)
- “opaque” is a typedef (i.e., unsigned char)

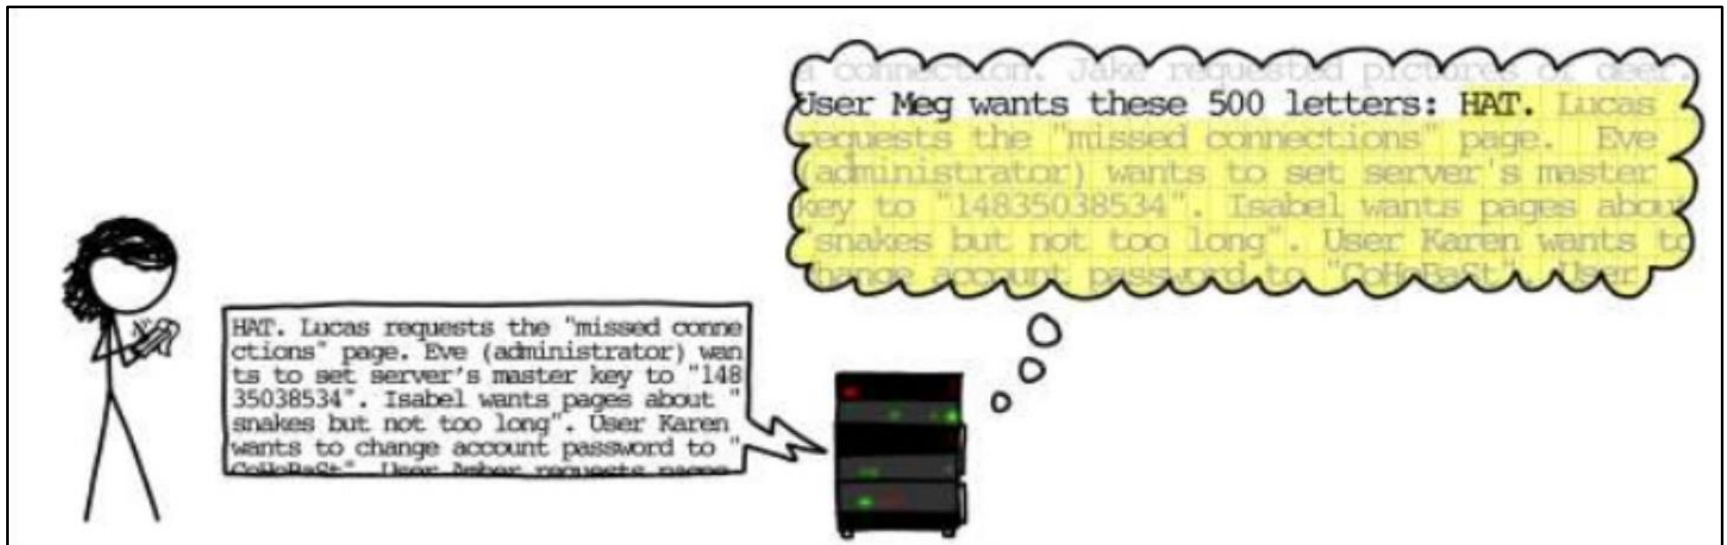
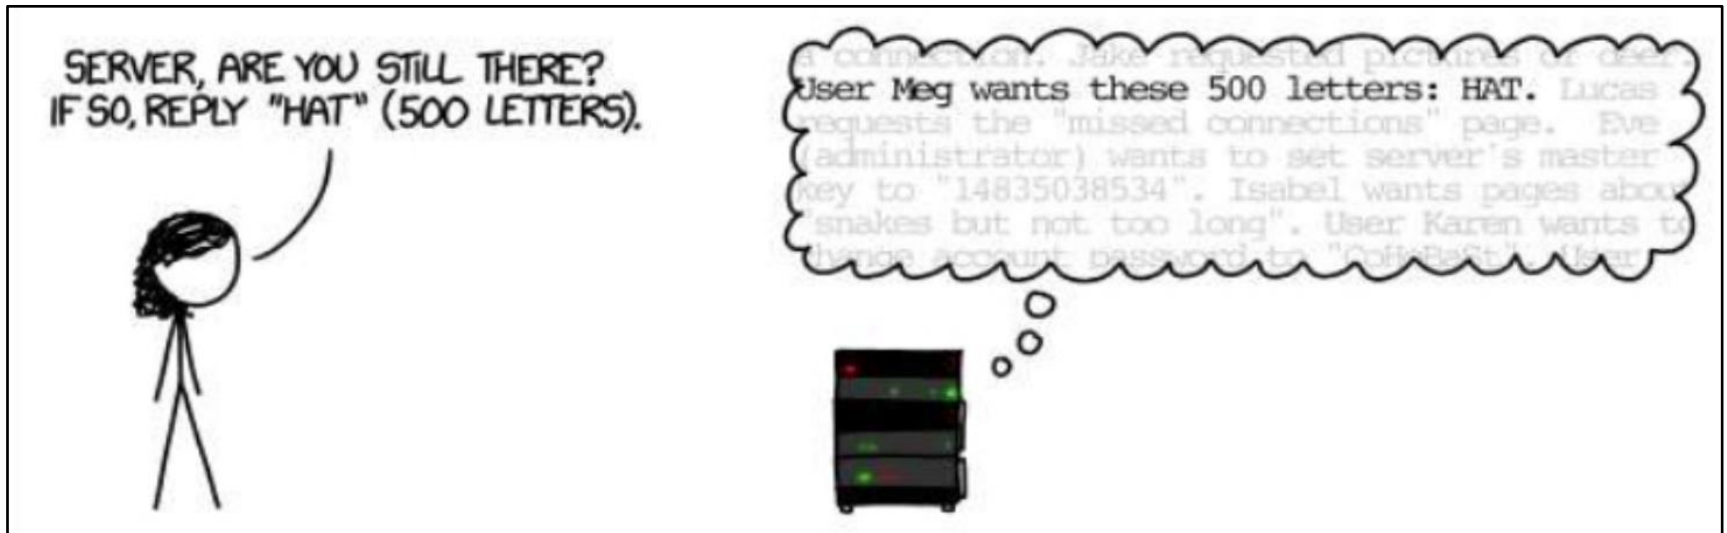
Heartbleed Exploit

8

- The server receives a Heartbeat request message and copies it into memory, for further processing
 - ▣ However, memory also contains information from other sessions including tokens, keys, session IDs etc.
- If *payload_length* is actually larger than the *payload[..]*, the server will copy memory content beyond the payload array into the response message's payload array (let's call it *ret_payload*), which is then sent back to the sender
- *memcpy(ret_payload, payload, payload_length);*
- Remember, this is C (and not Java or Python), so array boundaries are not checked!
- This is a typical **buffer over read attack**

The Heartbleed Attack

9



Heartbleed Exploit Extract (Python Code)

10

□ <https://gist.github.com/eelsivart/10174134>

Heartbleed (CVE-2014-0160) Test & Exploit Python Script

```
heartbleed.py Raw
1  #!/usr/bin/python
2
3  # Modified by Travis Lee
4  # Last Updated: 4/21/14
5  # Version 1.16
6  #
7  # -changed output to display text only instead of hexdump and made it easier to read
8  # -added option to specify number of times to connect to server (to get more data)
9  # -added option to send STARTTLS command for use with SMTP/POP/IMAP/FTP/etc...
10 # -added option to specify an input file of multiple hosts, line delimited, with or without a port specified (host:port)
11 # -added option to have verbose output
12 # -added capability to automatically check if STARTTLS/STLS/AUTH TLS is supported when smtp/pop/imap/ftp ports are entered and automatically
13 # -added option for hex output
14 # -added option to output raw data to a file
15 # -added option to output ascii data to a file
16 # -added option to not display returned data on screen (good if doing many iterations and outputting to a file)
17 # -added tls version auto-detection
18 # -added an extract rsa private key mode (orig code from epixoip. will exit script when found and enables -d (do not display returned data
19 # -requires following modules: gmpy, pyasn1
20
21 # Quick and dirty demonstration of CVE-2014-0160 by Jared Stafford (jspenguin@jspenguin.org)
22 # The author disclaims copyright to this source code.
23
24 import sys
25 import struct
26 import socket
27 import time
28 import select
29 import re
```

What can be leaked?

11

```
38 20 2E 4E 45 R 2.0.50727; .NET CLR 3.5.30729;
30 37 32 39 38 T CLR 3.5.30729;
2E 30 2E 33 30 .NET CLR 3.0.30
43 65 6E 74 65 729; Media Centre
6E 66 6F 50 61 r PC 5.0.0; Infopath
2E 30 43 3E 20 th.2; .NET4.0C;
48 6F 73 74 3A .NET4.0C)..Host:
6E 65 79 77 65 [redacted]
8E 65 63 74 69 11.com; Connecti
69 76 65 0D 0A on: keep-Alive..
20 73 69 64 65 cookie: doc-side
6D 79 77 6F 72 bar=245px; mywor
3D 66 61 6C 73 [redacted]
49 44 3D 43 32 e; SESSIONID=C2
41 34 42 44 31 7E804DA97D5A48D1
46 33 37 3E 20 B9728EDA58F37;
73 69 64 65 67 [redacted]
35 3E 20 2F 70 ar.width=285; /p

63 63 6E 3D (direct) | utmccn-
63 6D 64 3D (direct) | utmced-
64 2E 74 6F [redacted]
42 4E 42 34 [redacted]
67 41 30 30 e3[redacted]
39 7E 03 69 [redacted]
3D 41 37 47 [redacted]
4F 4A 58 46 s-SRK5-O8ID-OJXF
37 62 64 32 12525dfa358e7bd2
38 33 66 38 97249f7b50b883f8
6E 3E 20 4A a44fa928e11in; j

4C 39 [redacted]
5 30 .NET CLR 2.0.50
3 2E 727; .NET CLR 3.
3 4C 5.30729; .NET CL
5 64 R 3.0.30729; Med
E 30 1a Center PC 6.0
E 4E ; Infopath.2; .N
0 45 ET4.0C; .NET4.0E
4 69 ).Accept-Encod
1 74 ng: gzip, deflat
2 61 e..Host: [redacted]
0 0A [redacted]
5 70 Connection: Keep
A 20 -Alive..Cur[redacted]
E 7A [redacted]
```

Server details

```
0 00 00 00 00 00
E 78 73 72 66 2E at [redacted].xsr[redacted]
0 35 52 4B 35 2D token=A7G5-SRK5-
7 65 63 36 63 34 O8ID-OJXF|7ec6a4
7 33 38 35 3E 36 da22594b48738566
9 30 33 31 31 35 c331f7969a903115
3 45 53 53 45 4F 9b|out; 1SESSIO
9 45 36 37 44 47 NID=FC3103986706
3 44 31 34 37 38 A80561B4B43D1478
1 CD 12 93 24 C6 SDF0....1.A...S.
0 34 03 03 03 03 .....GT....
C 2E 63 6F 6D 7C [redacted]
6 65 72 72 61 6C [redacted]
E 6E 6E 73 73 6E [redacted]
```

Keys?

```
09 00 79 09 0E 07 43 trw[suspect]ryingw
25 37 44 25 32 30 6F 20and%2Fpe%70%20o
64 61 79 2E 25 32 30 r%20thursday.%20
6E 75 6D 62 65 72 3A .co-e103-number:
30 20 65 31 20 34 30 7701705.co-e104
25 32 30 74 58 65 25 string:if%20the%
73 25 32 30 67 69 76 20day%20is%20gily
30 74 6F 64 61 79 27 en%20as%20today
30 6F 72 25 62 30 74 s%20day%20or%20t
25 32 30 64 61 79 25 omorrow's%20day%
74 61 74 25 62 30 72 20thermostat%20r
30 77 69 6C 6C 25 32 esponse%20will%2
69 6C 67 25 62 30 77 0be%20end%20ing%20w
25 32 30 6E 65 78 74 1th%20%22%20next
61 79 25 32 60 25 32 %20Thursday%20%2
6D 6F 72 72 6F 77 25 0%2F%20tomorrow%
70 74 25 32 60 6E 61 22%20except%20ona
30 74 68 65 65 72 30 me%20of%20th%20
30 20 65 31 30 35 3D day.%20.co-e105-
53 53 45 44 0A 63 30 string:PASSED.CO
62 65 72 3A 31 32 0A -e106-number:12.
```

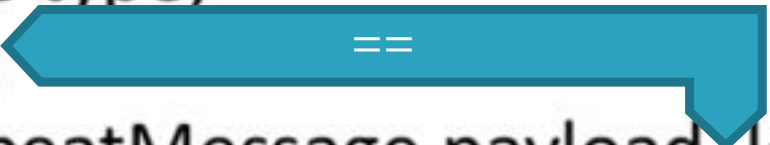
CSRF tokens

```
8 65 [redacted]
3 61 [redacted]
1 65 [redacted]
2 73 [redacted]
2 61 [redacted]
1 4C [redacted]
0 5A [redacted]
4E 26 [redacted]
69 70 [redacted]
26 6E [redacted]
28 48 [redacted]
A 41 [redacted]
2 41 [redacted]
66 4D [redacted]
48 4F [redacted]
69 5A [redacted]
1 76 [redacted]
7 55 [redacted]
69 68 [redacted]
64 67 [redacted]
6 72 [redacted]
1 46 [redacted]
0 53 [redacted]
11 70 [redacted]
13 65 [redacted]
13 74 [redacted]
15 72 [redacted]
11 20 [redacted]
```

What happened next?

12

- ❑ The Heartbleed bug was fixed (of course)
- ❑ Further checks and balances were added to validate that payload length was correct

```
struct {  
    HeartbeatMessageType type;  
    uint16 payload_length;  ==  
    opaque payload[HeartbeatMessage.payload_length];  
    opaque padding[padding_length];  
} HeartbeatMessage;
```

Heartbleed Impact

13

- The Heartbleed vulnerability was in all versions of OpenSSL released between March 2012 and April 2014
 - ▣ It was a zero-day (i.e., a vulnerability unknown to its owners, developers or anyone capable of mitigating it) **for almost 2 years**
- According to CVE-2014-0160, the following operating system distributions were potentially affected:
 - ▣ Debian Wheezy (stable)
 - ▣ Ubuntu 12.04.4 LTS
 - ▣ CentOS 6.5
 - ▣ Fedora 18
 - ▣ OpenBSD 5.3
 - ▣ FreeBSD 10.0
 - ▣ NetBSD 5.0.2
 - ▣ OpenSUSE 12.2

Lessons learnt

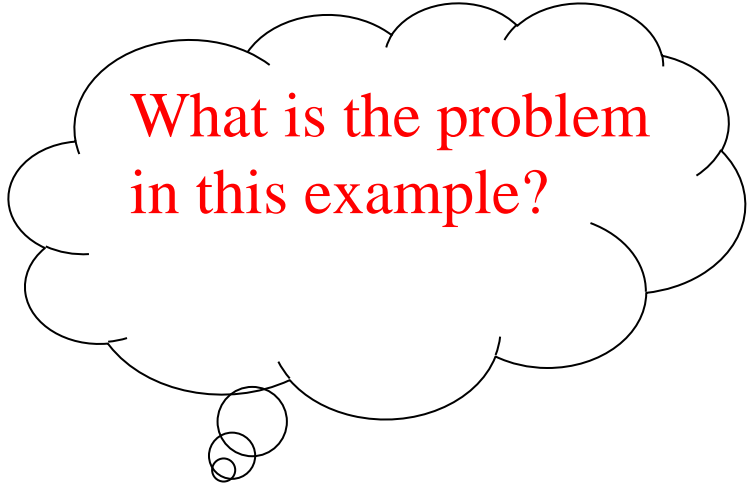
14

- *OpenSSL core developer Ben Laurie claimed that a security audit of OpenSSL would have caught Heartbleed*
- *Some other quotes from the security community:*
 - ▣ *“Think about it, OpenSSL only has two fulltime people to write, maintain, test, and review 500,000 lines of business-critical code”*
 - ▣ *“The mystery is not that a few overworked volunteers missed this bug; the mystery is why it hasn't happened more often”*
 - ▣ *“There should be a continuous effort to simplify the code, because otherwise just adding capabilities will slowly increase the software complexity. The code should be refactored over time to make it simple and clear, not just constantly add new features. The goal should be code that is “obviously right”, as opposed to code that is so complicated that “I can't see any problems”*

Related Problem: Buffer Overflow / Stack Overflow

```
#include <string.h>
void foo (char *bar)
{
    char c[12];
    strcpy(c, bar);
}

int main (int argc, char **argv)
{
    foo(argv[1]);
    return(1);
}
```



What is the problem
in this example?

Example for a Stack Overflow

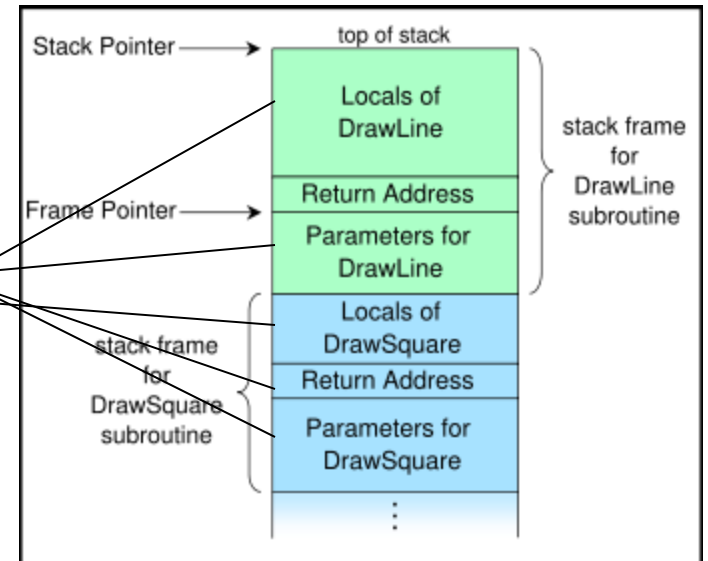
```
#include <string.h>
void foo (char *bar)
{
    char c[12];
    strcpy(c, bar);
}

int main (int argc, char **argv)
{
    foo(argv[1]);
    return(1);
}
```

- Lets assume the compiled program is called test
- Test is invoked from command line (next slide):
 - “> test hello” will work fine
 - “> test AAAAAAAAAAAAAA AAAAAAAAAA” (> 11 characters) may crash the program

Background Info: The Call Stack

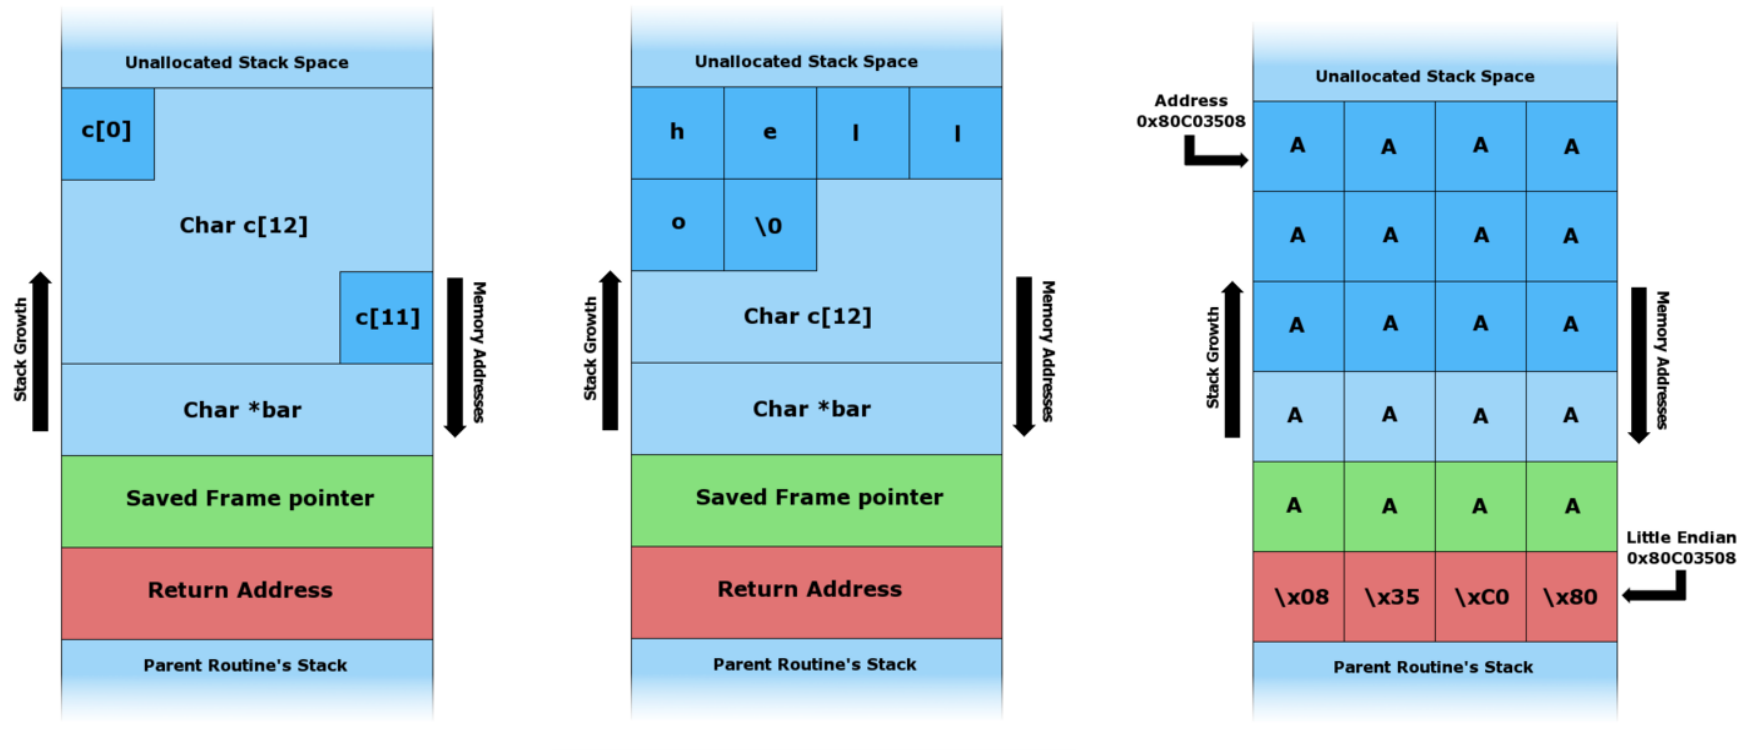
```
main()
{
    DrawSquare(1,1,4,4);
    ...
}
...
void DrawSquare(int lux, int luy, int rbx, int rby)
{
    int t1, t2;
    DrawLine(lux, luy, rbx, luy);
    DrawLine(lux, luy, lux, rby);
    ...
}
...
void DrawLine(int p1x, int p1y, int p2x, int p2y)
{
    int temp1;
    ...
}
```



Background Info: The Call Stack

- Each stack frame contains a stack pointer to the top of the frame immediately below
 - ▣ The stack pointer is a mutable register
- The stack frame is the collection of all data on the stack associated with one subprogram call. The stack frame generally includes:
 - ▣ The return address
 - ▣ Argument variables passed on the stack
 - ▣ Local variables
- A frame pointer of a given invocation of a function is a copy of the stack pointer as it was before the function was invoked
- If a stack frame is corrupted, i.e. overwritten, arguments, variables and / or return address do change
- If the return address is manipulated, the program can crash, or malware can be executed (with the return address being the start address of the malware in memory)

Example for a Stack Overflow



Buffer Overflow Countermeasures

- ❑ Use a programming language that supports automatic bounds checking of buffers
 - ❑ Java or Python, but NOT C
- ❑ Use a language specific library module that implements info validation in the form of safe buffer handling
- ❑ Compilers can produce a warning when an unsafe function call is made, or can add code for buffer overflow detection
- ❑ An Operating System can enforce more stringent memory access control so that buffer overflows cannot infringe into the protected areas of the main memory

Buffer Overflow Mitigation using Electric Fence / Boundary Checks

21

- Here each data object (i.e., array) is guarded by a boundary signature that is checked for its integrity every time that object is accessed
- If the signature has changed as shown below, the data object is deemed to be corrupted, and an alarm will be raised

Before	0xDA	0xEF	Array[0]	Array[1]	...	Array[n]	0xFF	0xED
Attack	0xAA	0xAA	0xAA	0xAA	...	0xAA	0xAA	0xAA

Example Code

22

```
...
char boundary0 = 0xDA;
char boundary1 = 0xEF
char array[n];
char boundary2 = 0xFF;
char boundary3 = 0xED;
...
// Access array[] only if boundary is intact.
if ((boundary0 == 0xDA) && (boundary1 == 0xEF) && (boundary2 == 0xFF) && (boundary3 == 0xED))
{
    // Access array
    ...
}
else
{
    // Error handling
    ...
}
```