

Caveat Implementor! Key Recovery Attacks on MEGA

Conference Paper**Author(s):**

Albrecht, Martin R.; Haller, Miro; Mareková, Lenka; Paterson, Kenneth G.

Publication date:

2023

Permanent link:

<https://doi.org/10.3929/ethz-b-000614276>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

Lecture Notes in Computer Science 14008, https://doi.org/10.1007/978-3-031-30589-4_7

Caveat Implementor! Key Recovery Attacks on MEGA

Martin R. Albrecht¹, Miro Haller²[0000–0001–8796–5064],
Lenka Mareková³[0000–0001–8616–4150], and Kenneth
G. Paterson²[0000–0002–5145–4489]

¹ King’s College London

`martin.albrecht@kcl.ac.uk`

² Applied Cryptography Group, ETH Zurich

`kenny.paterson@inf.ethz.ch`, `miro.haller@ethz.ch`

³ Information Security Group, Royal Holloway, University of London

`lenka.marekova.2018@rhul.ac.uk`

Abstract. MEGA is a large-scale cloud storage and communication platform that aims to provide end-to-end encryption for stored data. A recent analysis by Backendal, Haller and Paterson (IEEE S&P 2023) invalidated these security claims by presenting practical attacks against MEGA that could be mounted by the MEGA service provider. In response, the MEGA developers added lightweight sanity checks on the user RSA private keys used in MEGA, sufficient to prevent the previous attacks.

We analyse these new sanity checks and show how they themselves can be exploited to mount novel attacks on MEGA that recover a target user’s RSA private key with only slightly higher attack complexity than the original attacks. We identify the presence of an ECB encryption oracle under a target user’s master key in the MEGA system; this oracle provides our adversary with the ability to partially overwrite a target user’s RSA private key with chosen data, a powerful capability that we use in our attacks. We then present two distinct types of attack, each type exploiting different error conditions arising in the sanity checks and in subsequent cryptographic processing during MEGA’s user authentication procedure. The first type appears to be novel and exploits the manner in which the MEGA code handles modular inversion when recomputing $u = q^{-1} \bmod p$. The second can be viewed as a small subgroup attack (van Oorschot and Wiener, EUROCRYPT 1996, Lim and Lee, CRYPTO 1998). We prototype the attacks and show that they work in practice.

As a side contribution, we show how to improve the RSA key recovery attack of Backendal-Haller-Paterson against the unpatched version of MEGA to require only 2 logins instead of the original 512.

We conclude by discussing wider lessons about secure implementation of cryptography that our work surfaces.

1 Introduction

MEGA is a cloud storage and communication platform with over 265 million user accounts and more than 10 million daily users [15], advertising itself as secure and private by design. The platform distinguishes itself from other major providers by offering end-to-end encryption for stored data. On MEGA, user files should remain confidential even if the storage provider is malicious or has been compromised through a breach, implying security in a strong threat model. The security of MEGA in this setting was recently analysed in detail by [2], which describes five attacks on the cryptographic protocol used by MEGA to authenticate users and encrypt user data. The first two of these attacks completely broke the confidentiality of user files. Shortly after, [31] significantly improved the first attack in [2], reducing its requirement of 512 user logins to just 6.

At their heart, the attacks in [2] exploit the lack of both key separation and integrity protection for stored keys in the MEGA design: a single user master key is used to encrypt both the user’s RSA private key (used during user authentication) and the user’s file encryption keys themselves; meanwhile AES in ECB mode is used for the encryption. This allowed the authors of [2] to corrupt the RSA private key in certain ways that leaked useful information during the authentication protocol, as well as to “cut and paste” AES-ECB blocks from file encryption keys into the RSA private key.

The authors of [2] proposed an immediate and non-invasive mitigation step in the form of adding a MAC to the existing construction.⁴ In response, MEGA chose to not implement this or any of the other originally suggested countermeasures. Instead, MEGA added extra sanity checks in the client software to do more validation of payloads during or after decryption [18]. These checks were sufficient to prevent the specific attacks of [2, 31].

Shortly after MEGA released their patch addressing the attacks of [2], they made one other change which (as we will show below) further increased the attack surface of their code: they added detailed error reporting during the decryption and sanity checking processes done by the client as part of the authentication protocol [19]. The errors produced during these steps are mostly distinguishable from one another and the error messages are sent to the server in place of the usual authentication response. A malicious storage provider can exploit this verbose behaviour, triggering the errors by supplying specially crafted inputs in an attempt to learn something about the decrypted data.

⁴ This by itself does not suffice for authenticated encryption security, but presents the “immediate” level of countermeasures, i.e. the most easily achievable solution in the short term. [2] outlines further levels of countermeasures termed “minimal” and “recommended”, which provide better guarantees but require more fundamental changes to the MEGA platform.

1.1 Contributions

In the MEGA infrastructure, each user has a master key k_M that is used with AES-ECB mode to encrypt multiple items, including the user’s RSA private key and individual file encryption keys (in a special obfuscated format). In this work we describe two new attacks on the patched MEGA infrastructure in the malicious server setting which achieve an AES-ECB *decryption* capability under k_M . These attacks can be used to recover individual 128-bit blocks of a target user’s RSA private key. Combining this with lattice techniques, we can efficiently recover the entirety of the target user’s RSA private key after recovering four specific blocks. Once this private key is recovered, the adversary can trivially decrypt the RSA ciphertexts appearing in file sharing messages to recover the keys needed to decrypt any files shared with the target user. The attacks can also be used to recover individual file encryption keys directly. As with [2], these attacks exploit the lack of key separation and integrity protection in the MEGA design, showing that the patch and further changes made by MEGA in response to [2] were not only insufficient but actively harmful.

Both attacks make use of an ECB *encryption* oracle that is present in the MEGAdrop feature, a part of the MEGA system that is supposed to be independent of the authentication protocol, yet uses the same master key k_M . This feature enables the receiving of shared files from unregistered users. In short, MEGAdrop encrypts a newly shared file’s encryption key to a user’s public RSA key, but the user’s client then silently re-encrypts that file encryption key under k_M using AES-ECB whenever the user is logged in. Since a malicious server can arbitrarily choose the file encryption key when sharing files with the user and then observe the resulting AES-ECB ciphertext, this provides the ECB encryption oracle that we need. For technical reasons explained later, we obtain two AES-ECB encrypted 128-bit blocks for each use of the oracle. Notably, the ECB encryption oracle can be realised without any user interaction. Details can be found in Section 2.

The attacks also exploit the distinguishable errors arising during user authentication. We describe the individual errors in detail in Section 2. Both attacks can be seen as *key overwriting* attacks, since they rely on manipulating the values that are interpreted as the RSA private key by the client, and on including the target AES-ECB ciphertext block in a particular position in the encoded and encrypted RSA private key. This causes the errors that are triggered during client-side cryptographic processing to depend on the target plaintext block. User interaction is formally required for these attacks, which is why we measure their cost in terms of the number of login attempts they need (they are otherwise computationally inexpensive). As a secondary measure of attack complexity, we account for the number of ECB encryption oracle calls needed.

The first attack, described in Section 3, exploits an implicit error in the computation of modular inverses when sanity checking the RSA private key. It is an (un)fortunate consequence of an otherwise harmless bug in the code (not checking whether an inverse exists) which is caught by the client and reported

to the server. The malicious server can use this oracle repeatedly to learn the value of the target AES-ECB plaintext block modulo a number of small primes, which enables recovery of the full block using the CRT. The attack requires on average $2^{9.29}$ login attempts per recovered AES-ECB plaintext block and 66 ECB encryption oracle queries per attacked user.

The second attack, described in Section 4, relies on how RSA decryption is carried out by the client during user authentication. It exploits a legacy artefact in the code that changes the resulting RSA plaintext length if a certain byte condition on the plaintext does not hold, in combination with an explicit error arising from a plaintext length check that is again reported to the server. The core idea is as follows. Because the user’s RSA private key is encrypted as a sequence of AES-ECB blocks, we can use the ECB encryption oracle to overwrite parts of that key – including p , q , and d – on the granularity of 128-bit blocks. The attack exploits this capability to mount a small order subgroup attack [12, 36] by overwriting the RSA primes p , q with values such that $(p - 1)(q - 1)$ has known small prime factors. The attack also overwrites d with a value that is completely known except in the target plaintext block. By also choosing the RSA ciphertext, the server can force the client’s RSA decryption to take place in any one of the small subgroups corresponding to each of the small prime factors of $(p - 1)(q - 1)$. Then, the malicious server can use the length check oracle repeatedly to learn the overwritten value of d , and hence the target plaintext block, modulo each of the small primes. The final step again combines these values using the CRT to recover the target block.

We present two main versions of the second attack: one that is simpler but which requires a large amount of precomputation and one that is more complex but only requires negligible precomputation. On average, these versions require $2^{11.24}$ and $2^{11.63}$ login attempts per block, respectively. In both versions, this second attack requires a smaller number (up to 15) of ECB encryption oracle queries per attacked user than our first attack does. Further, this second attack exploits different errors from the first one and also relies on behaviours resulting from the “legacy” check on the second byte of plaintext. We include this attack to showcase that the existence of such checks and differentiated error reporting increases the attack surface.

Since the two attacks work on a per-block basis, we discuss how best to recover the entire RSA private key of the target user with the help of lattice techniques in Section 5. This reduces the number of blocks that need to be recovered using either of the two attacks to 4 instead of the 9 that would be required if the attacks were used directly to e.g. recover all of p . The attack complexity of recovering the full RSA private key using our first attack is then $2^{11.29}$ login attempts on average.

As a side contribution, we show in Section 6 how to combine the ECB encryption oracle obtained from the MEGAdrop feature with the second attack in [2] to recover a target user’s RSA private key from an *unpatched* MEGA client using

only 2 logins (compared to the 512 logins needed in [2] and the 6 needed in [31]). This shows that the original, unpatched MEGA system was even weaker than previously thought.

We conclude by briefly discussing attack mitigation in Section 7, noting the problematic nature of relying on easy-to-implement countermeasures that do not properly address the core security vulnerabilities. In that section, we also draw wider lessons from our work.

1.2 Related work

The work of [2] provided a detailed overview of the MEGA infrastructure as well as attacks on confidentiality and integrity of user data stored on the platform. The follow-up work [31] significantly reduced the amount of user interaction required by the first attack of [2] but was already prevented by MEGA’s patches. The attacks in this work draw inspiration from the small-order subgroup attacks on DH [12, 36] and the key overwriting attacks on OpenPGP [5, 10]. The use of a plaintext checking oracle is reminiscent of Bleichenbacher’s attack on RSA with PKCS#1 v1.5 encoding [4] but we target private key recovery rather than plaintext recovery.

1.3 Validation

We have verified the presence of the ECB encryption oracle, implemented the attacks and verified them in practice on test account data, using a TLS-MitM setup with mitmproxy [35] to minimise interaction with the real MEGA servers and a locally-run MEGA web client (version 4.21.4) [16]. We made a single modification to the web client to automatically simulate repeated client login attempts after one initial manual login. The attacks were able to recover arbitrary AES-ECB-encrypted blocks of the test user’s RSA private key with query costs consistent with our analysis (averaging $2^{9.30}$ login attempts for the first attack). We also implemented a proof of concept for recovering the entire RSA private key given four known blocks using lattice techniques. The code is available as supplementary material.⁵

1.4 Disclosure

We contacted MEGA to inform them of the vulnerabilities in their system on 29.09.2022. We suggested a 90-day disclosure period. We also suggested mitigations, stressing the importance of providing proper cryptographic integrity for data stored under users’ master keys. MEGA acknowledged receipt of our disclosure on 30.09.2022. They said they would begin working on fixes and liaise with us before deploying them. On 28.11.2022, MEGA informed us that they were working on hardening their client software, which would include changing

⁵ Available at <https://github.com/MEGA-caveat/mega-caveat-poc>.

how private keys are stored, removing the ECB encryption oracle as well as replacing the `asmcrypto.js` library. We provided high-level feedback on the proposed changes. The upgrade should mitigate against our specific attacks as well as potential future attacks, though we have not reviewed the changes in detail. Given the scale of the changes, we agreed to move the disclosure to 06.03.2023 to coincide with the rolling-out of the upgraded client software and the publication of this paper. The upgrade was implemented in version 4.32.4 of the web client [30]. MEGA awarded a bug bounty.

2 Oracles

2.1 Notation

We begin by establishing some notation that we use throughout. Concatenation is denoted by \parallel . $[m]_k$ denotes an encryption of m under the key k , where the algorithm is determined by the context. B denotes bytes, and for x , $|x|_B$ denotes the length of x in bytes and $|x|_b$ denotes the length of x in bits.⁶ For a tuple $X = (x_0, \dots, x_{n-1})$, $|X| = n$ denotes its size. For a byte string $\mathbf{m} = b_0 \parallel b_1 \parallel \dots \parallel b_{n-1}$ of length n and $s, t \in \mathbb{N}$, we define $\mathbf{m}[s] := b_s$, and $\mathbf{m}[s : t] := b_s \parallel \dots \parallel b_{t-1}$ for $s < t$. An empty object is denoted by `null`, and a zero byte by `00`. $\text{ZeroPad}(\mathbf{m}, n) := 00 \parallel 00 \parallel \dots \parallel \mathbf{m}$ such that $|\mathbf{m}|_B = n$, i.e. left-pad \mathbf{m} with zero bytes. If it is necessary to distinguish between a byte representation and other types, \mathbf{m} (as opposed to m) denotes a byte string. Conversion between byte strings and integers remains implicit, so we may write $m \leftarrow \mathbf{m}$ and vice-versa. $(\mathbb{Z}/n\mathbb{Z})^\times$ denotes the multiplicative group of integers modulo n . By $x \leftarrow_\$ S$ we denote x sampled uniformly at random from S . In our attacks, B denotes a target plaintext block, which is a byte string with $|B|_b = 128$. To differentiate it from a value computed while attempting to recover this block (which could be different if the attack is not correct), we denote the computed value by B^* .

2.2 ECB encryption oracle

MEGA's webclient exposes an ECB encryption oracle under a user's master key \mathbf{k}_M . This oracle allows MEGA, or anyone controlling their infrastructure, to encrypt 32 bytes of chosen plaintext in AES-ECB mode under the target user's master key \mathbf{k}_M in a single query. Since AES-ECB without any additional measures does not provide any integrity protection, ciphertexts containing blocks that the adversary queried to the oracle cannot be distinguished without additional tests on the expected structure of the plaintext.

The oracle stems from code related to the MEGAdrop feature as shown in Fig. 1. MEGAdrop enables anyone to upload files to a folder in the cloud storage of

⁶ For $x \in \mathbb{Z}$, the value of $|x|_b$ as understood by the MEGA client implementations is not always exact. In the big integer representation used by the web client, $|x|_b$ is normally rounded up to the closest multiple of 8 or 32.

the recipient without needing an account on MEGA. The recipient activates MEGAdrop for one of their folders and obtains a link that they can share with others. Unlike shared folders, senders do not see any file stored in the MEGAdrop upload folder.

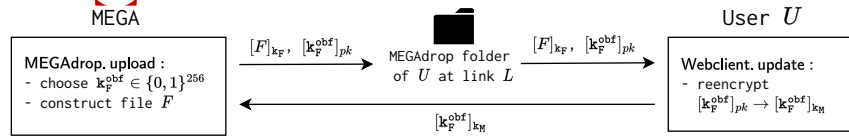


Fig. 1. Overview of the ECB encryption oracle under a user's master key k_M .

The left algorithm of Fig. 2 describes the upload feature of MEGAdrop. The adversary can pick some file key k_F , nonce N_F , and file F during the upload process for the MEGAdrop folder at the link L . The upload feature locally encrypts the file with AES-CCM using k_F and some nonce N_F picked by the client. Backendal et al. describe MEGA's encryption in more detail⁷ on lines 2-11 of Fig. 2 in [2].

MEGAdrop.upload(k_F, N_F, F, L)	Webclient.update()
1 : $[F]_{k_F}, T_{cond} \leftarrow \text{File.enc}(k_F, N_F, F)$	1 : while true do
2 : $k_F^{obf} \leftarrow \text{ObfKey}(k_F, N_F, T_{cond})$	2 : $\tau \leftarrow \text{Server.fetch_update}(k_M, sk)$
3 : $pk \leftarrow \text{Server.lookup}(L)$	3 : if $\tau \neq \perp$ then
4 : $[k_F^{obf}]_{pk} \leftarrow \text{RSA.Enc}(pk, k_F^{obf})$	4 : $[F]_{k_F}, [k_F^{obf}]_{pk} \leftarrow \tau$
5 : $\text{Server.upload}([F]_{k_F}, [k_F^{obf}]_{pk})$	5 : $k_F^{obf} \leftarrow \text{RSA.Dec}(sk, [k_F^{obf}]_{pk})$
	6 : $[k_F^{obf}]_{k_M} \leftarrow \text{AES-ECB.Enc}(k_M, k_F^{obf})$
	7 : $\text{Server.upload}([k_F^{obf}]_{k_M})$
	8 : endif
	9 : endwhile

Fig. 2. MEGAdrop pseudocode. MEGAdrop.upload encrypts a file F with key k_F and nonce N_F , uploaded to the MEGAdrop folder with link L . Webclient.update shows how active clients regularly poll for updates and re-encrypt node keys immediately.

To instantiate an ECB encryption oracle, the adversary sets k_F^{obf} to 32 bytes of its choosing. Since $k_F^{obf} = (k_F \oplus x) \parallel x$ for $x = N_F \parallel T_{cond}$, the obfuscated key defines the values for k_F , N_F , and T_{cond} used in the file encryption (cf. Fig. 4 in [2]). The adversary can use the file reconstruction part of the framing attack described in [2] to obtain a file F that, when encrypted with k_F and N_F , produces the MAC tag

⁷ For instance, we omit the file attributes in our description for simplicity.

value T_{cond} . Consequently, the adversary can run $\text{MEGAdrop.upload}(k_F, N_F, F, L)$ to upload k_F^{obf} , encrypted under the receiver’s public RSA key, to the server.

Section 9.12 of MEGA’s security white paper [17] states that to “conserve CPU cycles, RSA-encrypted keys are transformed into AES-encrypted keys when encountered”. Indeed, the webclient regularly polls for new files in the background and, when encountering an RSA-encrypted key $[k_F^{\text{obf}}]_{pk}$, re-encrypts k_F^{obf} with k_M and AES-ECB to produce an AES-ECB ciphertext that we denote by $[k_F^{\text{obf}}]_{k_M}$. It then uploads this updated key to the server (cf. [25]) as shown in the right half of Fig. 2. Therefore, the malicious server can learn the AES-ECB plaintext-ciphertext pair $(k_F^{\text{obf}}, [k_F^{\text{obf}}]_{k_M})$.

While testing this oracle in mitmproxy [35], we noticed that the server can pretend that a new file was uploaded to a MEGAdrop folder. The webclient re-encrypts the key as described in Fig. 2 even if the recipient does not use MEGAdrop and the file has an invalid path. Thus, we have an efficient ECB encryption oracle that does not require any user interaction and leaves no persistent traces in the user’s cloud storage. It encrypts 32 B per query and can be accessed repeatedly.

2.3 Oracles from decoding and decryption error reports

Consider the authentication and session ID exchange that takes place every time a user logs into their account, summarized in Fig. 3 and described in more detail in [2]. Let k_e be the user’s 128-bit symmetric encryption key derived from their password, k_M the user’s 128-bit symmetric master key and (pk, sk) the user’s 2048-bit RSA keypair.

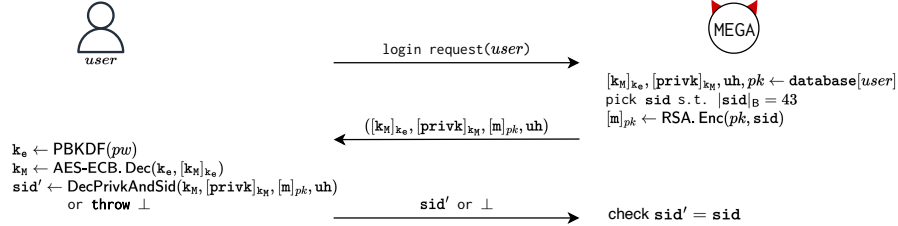


Fig. 3. Simplified overview of the MEGA login procedure.

Here, we focus on one part of this exchange, namely when the server responds to the user’s request with the tuple $([k_M]_{k_e}, [\text{privk}]_{k_M}, [m]_{pk}, \text{uh})$, where $[k_M]_{k_e}$ and $[\text{privk}]_{k_M}$ are AES-ECB-encrypted, $[m]_{pk}$ is RSA-encrypted and uh is in plaintext. Then, privk encodes the secret key sk for RSA-CRT as shown in Fig. 4, m encodes the session ID sid and uh is an 11-byte user handle string. The exact alignments of the fields in privk with respect to the AES-ECB block boundaries will be important in our attacks. The processing done by the client after it decrypts $[k_M]_{k_e}$ is shown in Fig. 5. This is the updated behaviour resulting from the patches

described in Section 1 and converted into pseudocode as faithfully as possible, i.e. in some cases surfacing lower-level processing if it is relevant.

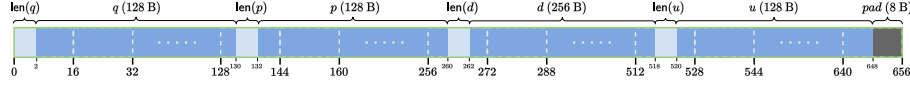


Fig. 4. Encoding of the RSA secret key together with the block boundaries marking the start of different 16-byte AES-ECB blocks. Each length encoding field consists of 2 bytes, meaning that data fields start progressively further into AES-ECB blocks.

In Fig. 5, we adopt the notation “**require condition else error**” to mean that the client checks the *condition* and if it is not satisfied, it aborts and outputs the *error* to the server. Decoding between base64-strings, bytes and integers is left implicit unless relevant to some error. Computation of $a^{-1} \bmod b$ should be understood to return `null` if $\gcd(a, b) \neq 1$.

In `DecodePrivk(privk)`, the function `Parse(privk)` sequentially reads through the bytes of `privk` whose expected form, shown in Fig. 4, is `len(q) || q || len(p) || p || len(d) || d || len(u) || u || pad` where `len(x)` denotes the two-byte big-endian length encoding of the byte-length of `x` and `pad` is padding, and returns the tuple of integers $P = (q, p, d, u)$. If `DecPrivkAndSid(.)` returns successfully, then `sid` is sent to the server in the requests that follow. Notice that in addition to `DecPrivkAndSid(.)` returning a range of different error messages depending on the processing of secret values, it also modifies the resulting plaintext depending on whether the second byte of the RSA-decrypted value is `00` or not (line 9 of `DecryptSid(., .)`), a quirk that is explained in the original code only with the comment “Old bogus padding workaround” [24].

Caught and uncaught exceptions. Some of the errors shown in Fig. 5 are implicit, i.e. they are a result of lower-level exceptions caught at a higher level (the ones corresponding to \perp_5 and \perp_6). In all cases, they are shown at the exact place where the code aborts.

Further, due to some lower-level bugs in `asmcrypto.js` [20], the `bigint` and `crypto` library used by the web client, there are cases where the implementation never terminates:

- In `DecodePrivk(privk)` during `Parse(privk)`, if one of q, p, d, u is 0.
- In `DecodePrivk(privk)` during the computation of $q^{-1} \bmod p$ [22], if $q \bmod p = 0$. We observed that this is because the implementation of $\gcd(0, p)$ never terminates. The same issue arises during the computation of $d^{-1} \bmod (p - 1)(q - 1)$ if $d \bmod (p - 1)(q - 1) = 0$.

Similarly, there are cases when the implementation returns incorrect output:

DecPrivkAndSid($k_M, [\text{privk}]_{k_M}, [m]_{pk}, \text{uh}$):	DecryptSid($sk, [m]_{pk}$):
<pre> 1 : require $\text{uh} _{\mathbb{B}} = 11$ else \perp_1 2 : $\text{privk} \leftarrow \text{AES-ECB.Dec}(k_M, [\text{privk}]_{k_M})$ 3 : $sk \leftarrow \text{DecodePrivk}(\text{privk})$ 4 : $m \leftarrow \text{DecryptSid}(sk, [m]_{pk})$ 5 : require $m _{\mathbb{B}} = 255$ else $(\perp_2, m _{\mathbb{B}})$ 6 : require $m[16 : 27] = \text{uh}$ else \perp_3 7 : $\text{sid} \leftarrow m[0 : 43]$ 8 : return sid </pre>	<pre> 1 : $N, e, d, p, q, d_p, d_q, u \leftarrow sk$ 2 : $c \leftarrow [m]_{pk}$ 3 : require $c < N$ else \perp_7 4 : $x \leftarrow c^{d_p} \bmod p; y \leftarrow c^{d_q} \bmod q$ 5 : $t \leftarrow x - y \bmod p$ 6 : $h \leftarrow u \cdot t \bmod p$ 7 : $m \leftarrow h \cdot q + y \bmod 2^{ N _{\mathbb{B}}}$ 8 : $m \leftarrow \text{ZeroPad}(m, N _{\mathbb{B}})$ 9 : if $m[1] \neq 00$ then 10 : $m' \leftarrow 00 \parallel m$ 11 : else 12 : $m' \leftarrow m$ 13 : return $m'[2 : m' _{\mathbb{B}}]$ </pre>
DecodePrivk(privk):	
<pre> 1 : $P, \text{pad} \leftarrow \text{Parse}(\text{privk})$ 2 : require $P = 4 \wedge \text{pad} _{\mathbb{B}} < 16$ else \perp_4 3 : $q, p, d, u \leftarrow P$ 4 : $N \leftarrow p \cdot q$ 5 : $e \leftarrow d^{-1} \bmod (p-1)(q-1)$ 6 : $d_p \leftarrow d \bmod p; d_q \leftarrow d \bmod q$ 7 : $u' \leftarrow q^{-1} \bmod p$ 8 : require $u' \neq \text{null}$ else \perp_5 9 : $\text{cond} \leftarrow p _{\mathbb{B}}, q _{\mathbb{B}}, u _{\mathbb{B}} > 1000 \wedge d _{\mathbb{B}} > 2000$ 10 : require $\text{cond} \wedge (u' = u)$ else \perp_4 11 : $sk \leftarrow N, e, d, p, q, d_p, d_q, u$ 12 : require $e \neq \text{null}$ else \perp_6 13 : return sk </pre>	

Fig. 5. Client decoding and decryption to process the session ID, derived from [24,26–28].

- In $\text{DecryptSid}(sk, c)$ during the computation of $x \leftarrow c^{d_p} \bmod p$ (and likewise $y \leftarrow c^{d_q} \bmod q$), there are several issues.
 - If p is even, the code computes $x = 0$ regardless of the other input values, because modular power computations were not implemented for even moduli [21].
 - If $|p|_{\mathbb{B}} > 1024$, the implementation of Montgomery reduction [23] does not return correct values, and so the output x is also incorrect.

We were forced to work around some of these implementation errors in our attacks.

3 Attack based on modular inverse computation

Our first attack enables block-by-block plaintext recovery of AES-ECB blocks encrypted under k_M . In particular, this enables RSA private key recovery, i.e. the recovery of `privk`. Let $[B]_{k_M}$ be such a target ciphertext block with unknown target plaintext block B , for example corresponding to an unknown block of q from `privk`.

This attack is in the malicious server setting, or equivalently the TLS-MitM setting, and makes use of the ECB encryption oracle described in Section 2.2. It exploits the error type \perp_5 , which arises on line 7 and line 8 of `DecodePrivk(privk)` in Fig. 5 when $\gcd(p, q) \neq 1$. To get to this point, the server must submit inputs such that none of the previous error types are triggered. The server will only replace the $[\text{privk}]_{k_M}$ value and expect to abort before executing `DecryptSid(·, ·)`, so the only condition that must be satisfied is the one on line 2, which requires that the decrypted `privk` parses into 4 values without too much extra padding. Then, error \perp_5 can be distinguished from any of the errors that could follow, though with overwhelming probability this will be error \perp_4 from line 10 due to the server overwriting parts of `privk`.

The main idea behind this attack rests in the observation that if the server can construct $[\text{privk}^*]_{k_M}$ such that the decrypted and decoded p is divisible by a small prime r , and the decrypted and decoded q contains the target block B in its least-significant position, then the outputting of error \perp_5 leaks that $\gcd(p, q) \neq 1$ and thus (if some further conditions are satisfied), that $q \bmod r = 0$. From this, the server can learn the value of $B \bmod r$. Repeating this for a sufficient number of different primes r_i and combining the values using the Chinese Remainder Theorem (CRT), the server can learn the value of $B \bmod r_0 \cdot \dots \cdot r_{n-1}$. If $|r_0 \cdot \dots \cdot r_{n-1}|_b \geq 128$, the server recovers B .

In the following subsections, we describe two versions of the attack in more detail, starting with the simple, block-aligned version and then describing an attack that is more general and resistant to simple fixes. Both versions have been implemented and verified using our TLS-MitM setup described in Section 1.3.

3.1 Block-aligned, small-length version

The attack proceeds in two distinct phases. The first phase calls the ECB encryption oracle to obtain a set of chosen-plaintext blocks, which are then combined with a target block to form the ciphertexts submitted to the client as part of the second phase. The second phase relies on the client making a number of online login attempts. The ECB encryption oracle calls are shown as $[x]_{k_M} \leftarrow \text{OECB}_{k_M}(x)$ (if x consists of 2ℓ blocks, this call will involve ℓ uses of the actual oracle described in Section 2.2). The content of the modified ciphertexts

that will be submitted to the client is shown in Fig. 6. Note that we aim to reduce the number of $\text{OECB}_{\text{KM}}()$ calls by ensuring most of the content consists of all-zero blocks (or blocks containing the value 1), which only need to be queried once.

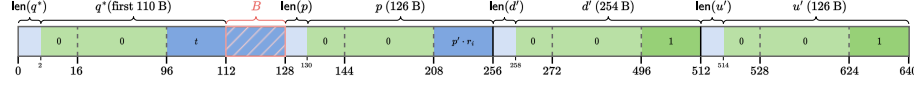


Fig. 6. The plaintext content of $\text{ct}_{i,t}$, where the all-zero blocks are light green, the blocks containing 1 are dark green and the placement of the target block B is in red.

Precomputation using the ECB encryption oracle. Take $\{r_0, \dots, r_{n-1}\} = \{7, 11, \dots, 103\}$, $n = 24$ small odd primes such that their product $R = \prod_{i=0}^{n-1} r_i$ has $|R|_{\text{B}} \geq 128$. Let $[B]_{\text{KM}}$ be the target ciphertext block and denote by B^* the plaintext block computed as part of this attack.

1. Generate a random prime p' such that $|p'|_{\text{B}} = 256$.
2. Let $d' \leftarrow 1$, $u' \leftarrow 1$ and encode them as byte strings d', u' such that $|d'|_{\text{B}} = 254$, $|u'|_{\text{B}} = 126$.
3. Let $\text{rest} \leftarrow \text{len}(d') \parallel d' \parallel \text{len}(u') \parallel u'$ and obtain $[\text{rest}]_{\text{KM}} \leftarrow \text{OECB}_{\text{KM}}(\text{rest})$.
4. For $i \in \{0, \dots, n-1\}$, do the following:
 - (a) Compute $p \leftarrow p' \cdot r_i$ and encode it as a byte string p such that $|p|_{\text{B}} = 126$.⁸
 - (b) Let $\text{ptp}_i \leftarrow \text{len}(p) \parallel p$ and obtain $[\text{ptp}_i]_{\text{KM}} \leftarrow \text{OECB}_{\text{KM}}(\text{ptp}_i)$.
5. For $t \in \{0, \dots, r_{n-1} - 1\}$, do the following:
 - (a) Compute $q^* \leftarrow 2^{128} \cdot t$ and encode it as a byte string q^* such that $|q^*|_{\text{B}} = 126$.
 - (b) Let $\text{ptq}_t \leftarrow \text{len}(q^*) \parallel \text{q}^*[0 : 110]$, which skips the last block of q^* to make space for the target. Obtain $[\text{ptq}_t]_{\text{KM}} \leftarrow \text{OECB}_{\text{KM}}(\text{ptq}_t)$.
6. For $i \in \{0, \dots, n-1\}$, do the following:
 - (a) For $t \in \{0, \dots, r_i - 1\}$, do the following:

$$\text{Store } \text{ct}_{i,t} \leftarrow [\text{ptq}_t]_{\text{KM}} \parallel [B]_{\text{KM}} \parallel [\text{ptp}_i]_{\text{KM}} \parallel [\text{rest}]_{\text{KM}}.$$

⁸ We include the prime p' for several reasons. First, because of one of the uncaught errors, we must make sure that $q \bmod p \neq 0$. Further, to avoid false positives from error \perp_5 , we need the $\gcd(p, q) \neq 1$ signal to be equivalent to $\gcd(p, q) = r_i$.

Online attack. Suppose we have a set of $\text{ct}_{i,t}$ as described above.

1. For $i \in \{0, \dots, n-1\}$, do the following:
 - (a) For $t \in \{0, \dots, r_i-1\}$, do the following:
 - i. When the client initiates a login, respond to the client's request with $([\mathbf{k}_M]_{\mathbf{k}_e}, \text{ct}_{i,t}, [\mathbf{m}]_{pk}, \mathbf{uh})$, where everything but $\text{ct}_{i,t}$ is as it would be in an honest response.
 - ii. If the client returns \perp_5 , save the value of t and break out of this loop.
 - (b) Save $B_i^* \leftarrow -2^{128} \cdot t \bmod r_i$.
2. Then, compute $B^* \bmod R$ by solving the system $B^* \equiv B_i^* \pmod{r_i}$ for $i \in \{0, \dots, n-1\}$ using CRT.

Correctness. Notice that for each decrypted $\text{ct}_{i,t}$, $\text{DecodePrivk}(\cdot)$ results in $p \leftarrow p' \cdot r_i$ and $q \leftarrow 2^{128} \cdot t + B$. The error \perp_5 will be triggered if and only if $\gcd(p, q) \neq 1$, which is equivalent to $\gcd(p, q) = r_i$, since p' is a prime larger than q . Hence \perp_5 is triggered if and only if $q \bmod r_i = 0$, and so if and only if $B \equiv -2^{128} \cdot t \pmod{r_i}$. This means that for the computed value B_i^* we have $B_i^* \equiv B \pmod{r_i}$. It follows that $B^* \equiv B \pmod{R}$. Since R is such that $|R|_{\text{b}} \geq 128$ and $|B|_{\text{b}} = 128$, we deduce that $B^* = B$ (over the integers).

Cost. First, we count the cost of recovering the target in terms of ECB encryption oracle calls, assuming that each repeated value (such as an all-zero block) is only queried once. As can be seen in Fig. 6, the encoding of q^* , p , d' and u' is block-aligned. The value **rest** consists of four non-zero blocks: two blocks that include a length encoding, and two identical blocks containing the value 1. Next, ptp_i also has four non-zero blocks: one length-encoding block and three blocks for $p' \cdot r_i$ since $|p' \cdot r_i|_{\text{b}} < 263 < 3 \cdot 128$; similarly, ptq_t has two non-zero blocks: one length-encoding block and one block for t since $|t|_{\text{b}} \leq |r_{n-1}|_{\text{b}} < 128$. Finally, notice that $\text{ptp}_i[0 : 16]$ is the same for all i , and similarly $\text{ptq}_t[0 : 16]$ is the same for all t , so the length-encoding blocks can be reused. Recalling that each use of the oracle returns two blocks of ciphertext, together the attack requires $\lceil \frac{1}{2} \cdot (1 + 3 + 2 + n \cdot 3 + r_{n-1}) \rceil = 91 \approx 2^{6.5}$ queries. Further, the result of these queries can be reused when recovering multiple blocks for a given target user.

Second, we count the number of online login attempts. On average, the attack requires $\frac{1}{2} \cdot \sum_{i=0}^{n-1} r_i = 627 \approx 2^{9.29}$ logins ($2^{10.29}$ in the worst case).⁹

3.2 Full-length version

The attack in Section 3.1 could technically be prevented by a number of simple checks, e.g. by moving the check on bit lengths before the client computes

⁹ Note that the attack can be easily modified to use one less login for each r_i . This is because, in the online phase, if the server does not get a positive answer from the oracle for any of the values $t \in \{0, \dots, r_i-2\}$, it means that the value r_i-1 is the correct one and so does not need to be submitted explicitly.

$q^{-1} \bmod p$ (and so possibly triggers \perp_5), by ensuring that $|p|_b, |q|_b = 1024$ or that $d, u \neq 1$. However, none of these changes would prevent this type of attack: here we provide a more general version that would still work if these changes were made. The content of the modified ciphertexts that will be submitted to the client is shown in Fig. 7.

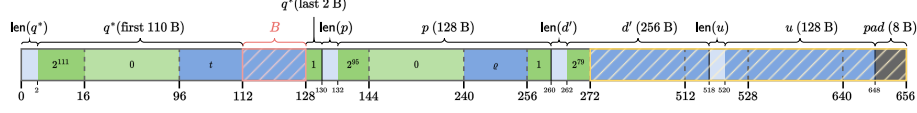


Fig. 7. The plaintext content of $\text{ct}_{i,t}$, where the all-zero blocks are light green, the parts containing fixed values are dark green, the placement of the target block B is in red and the placement of the unmodified values from ct is in yellow.

Precomputation using the ECB encryption oracle. As before, take $\{r_0, \dots, r_{n-1}\} = \{7, 11, \dots, 103\}$, $n = 24$ small odd primes such that their product $R = \prod_{i=0}^{n-1} r_i$ has $|R|_b \geq 128$. Let $[B]_{k_M}$ be the target ciphertext block and denote by B^* the plaintext block computed as part of this attack. Let $\text{ct} \leftarrow [\text{privk}]_{k_M}$ be the original ciphertext encrypting the user's private RSA key.

1. Let $d' \leftarrow 2^{2047}$ and encode it as a byte string d' such that $|d'|_B = 256$.
2. Let $\text{ptd} \leftarrow 00\ 00\ 00\ 01 \parallel \text{len}(d') \parallel d'[0 : 10]$ and get $[\text{ptd}]_{k_M} \leftarrow \text{OECB}_{k_M}(\text{ptd})$.
3. Let $[\text{rest}]_{k_M} \leftarrow \text{ct}[272 : |\text{ct}|_B]$. The slice begins with the ciphertext block that encrypts the most-significant full block of the original d .
4. For $i \in \{0, \dots, n-1\}$, do the following:
 - (a) Compute $p \leftarrow 2^{1023} + 2^{32} \cdot \varrho + 1$ for ϱ such that $p \equiv 0 \pmod{r_i}$ and p/r_i is prime. Encode it as a byte string p such that $|p|_B = 128$.
 - (b) Let $\text{ptp}_i \leftarrow 00\ 01 \parallel \text{len}(p) \parallel p[0 : 124]$ and get $[\text{ptp}_i]_{k_M} \leftarrow \text{OECB}_{k_M}(\text{ptp}_i)$.
5. For $t \in \{0, \dots, r_{n-1} - 1\}$, do the following:
 - (a) Compute $q^* \leftarrow 2^{1023} + 2^{128+16} \cdot t + 1$ and encode it as a byte string q^* such that $|q^*|_B = 128$.
 - (b) Let $\text{ptq}_t \leftarrow \text{len}(q^*) \parallel q^*[0 : 110]$ and obtain $[\text{ptq}_t]_{k_M} \leftarrow \text{OECB}_{k_M}(\text{ptq}_t)$.
6. For $i \in \{0, \dots, n-1\}$, do the following:
 - (a) For $t \in \{0, \dots, r_i - 1\}$, do the following:

$$\text{Store } \text{ct}_{i,t} \leftarrow [\text{ptq}_t]_{k_M} \parallel [B]_{k_M} \parallel [\text{ptp}_i]_{k_M} \parallel [\text{ptd}]_{k_M} \parallel [\text{rest}]_{k_M}.$$

Correctness. In this version, the precomputation must construct a modified ciphertext such that all values q, p, d, u are of the expected bit length. Recall that the plaintext encoding has the form: $\text{len}(q) \parallel q \parallel \text{len}(p) \parallel p \parallel \text{len}(d) \parallel d \parallel \text{len}(u) \parallel u \parallel \text{pad}$. Since each value is encoded by prefixing a two-byte length field and the original lengths are either 1024 bits or 2048, the values in the resulting plaintext are not block-aligned. This is why we construct the “partial” block ptd in Step 2 separately: it is composed of the final 4 bytes of p , $\text{len}(d')$ and the first 10 bytes of d' . Similarly, the block-aligned plaintext ptp in Step 4b begins with another partial block which consists of the final 2 bytes of q^* , $\text{len}(p)$ and the first 12 bytes of p .¹⁰ Finally, the modified blocks are “stitched” together in Step 6a as in the simple version of the attack, ensuring that the target B is interpreted as the last “full” block of q .

Cost. Finding p of the correct form for each i in Step 4a is easy and takes $326 \approx 2^{8.35}$ trials on average for the given primes r_i . This step is independent of user data and so can be reused to attack multiple users. With reference to Fig. 7, note that both ptp_i and ptq_t will likely have two non-zero blocks each. We assume the reuse of the length-encoding blocks as in Section 3.1. Thus the attack requires $\lceil \frac{1}{2} \cdot (1 + 1 + 2 + n + r_{n-1}) \rceil = 66 \approx 2^{6.04}$ queries.

Online attack. Suppose we have a set of $\text{ct}_{i,t}$ as described above.

1. For $i \in \{0, \dots, n-1\}$, do the following:
 - (a) For $t \in \{0, \dots, r_i-1\}$, do the following:
 - i. When the client initiates a login, respond to the client’s request with $([\mathbf{k}_M]_{\mathbf{k}_a}, \text{ct}_{i,t}, [\mathbf{m}]_{pk}, \mathbf{uh})$, where everything but $\text{ct}_{i,t}$ is as it would be in an honest response.
 - ii. If the client returns \perp_5 , save the value of t and break out of this loop.
 - (b) Save $B_i^* \leftarrow (2^{16})^{-1} \cdot (-2^{1023} - 2^{128+16} \cdot t - 1) \bmod r_i$.
2. Then, compute $B^* \bmod R$ by solving the system $B^* \equiv B_i^* \pmod{r_i}$ for $i \in \{0, \dots, n-1\}$ using CRT.

Correctness. Recall that for each decrypted $\text{ct}_{i,t}$, $\text{DecodePrivk}(\cdot)$ gets $p \leftarrow 2^{1023} + 2^{32} \cdot \varrho + 1$ and $q \leftarrow 2^{1023} + 2^{128+16} \cdot t + 2^{16} \cdot B + 1$. The overwritten values are encoded so that the parsing succeeds, and there are no other explicit errors that could be triggered before the error we are using for the attack.¹¹ The error \perp_5 will be triggered if and only if $\gcd(p, q) \neq 1$, which is equivalent to $\gcd(p, q) = r_i$ with high probability, since p/r_i is a large prime and the probability that

¹⁰ That is, $\text{ptd}[0 : 4] = \text{p}[124 : 128]$ for all p , and $\text{ptp}_i[0 : 2] = \mathbf{q}^*[126 : 128]$ for all \mathbf{q}^* .

¹¹ There is a possibility that $d^* \bmod (p-1)(q-1) = 0$ where $d^* \leftarrow d' + (d \bmod 2^{1968})$ and d is the original value encrypted in ct . Because of the uncaught non-termination bug arising during the computation of $(d^*)^{-1} \bmod (p-1)(q-1)$, in this case the attack would fail, but this is highly unlikely to happen in practice.

$q \equiv 0 \pmod{(p/r_i)}$ is $\approx 1/(p/r_i) \leq 2^{-1016}$. Hence \perp_5 is triggered if and only if $q \bmod r_i = 0$, and hence if and only if $B = (2^{16})^{-1} \cdot (-2^{1023} - 2^{128+16} \cdot t - 1) \bmod r_i$. Thus we have $B_i^* \equiv B \pmod{r_i}$. The rest of the analysis follows as for the simpler version of the attack.

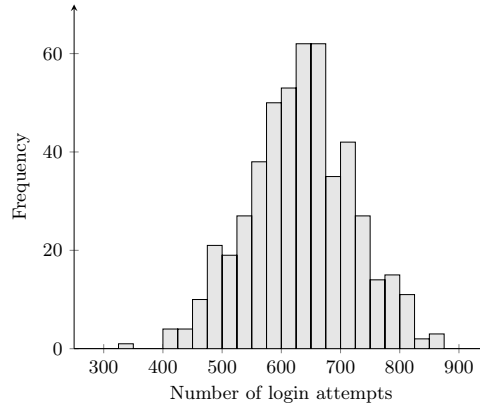


Fig. 8. Number of login attempts used by the attack over 500 runs.

Cost. The attack requires the same number of online login attempts as the simpler version in Section 3.1. We confirmed this in our implementation: in 500 runs of the attack recovering random ECB-encrypted blocks, the average number of login attempts required by the full version of the attack was $632 \approx 2^{9.30}$. The histogram is shown in Fig. 8.

4 Attack based on small subgroups

Here, we present our second AES-ECB decryption attack. In terms of login attempts it is less efficient than the attack in Section 3. However it requires fewer uses of the ECB encryption oracle. Further, it exploits a number of additional errors and also behaviours resulting from the “legacy” check on the second byte of the RSA plaintext.

The attack is also in the malicious server/TLS-MitM setting and uses the ECB encryption oracle from Section 2.2 with the aim of recovering blocks of d from the original `privk` (or any other AES-ECB-encrypted blocks that can be placed in their position). It exploits the errors \perp_2 and \perp_3 arising on line 5 and line 6 of `DecPrivkAndSid($k_M, [\text{privk}^*]_{k_M}, c^*, \text{uh}^*$)` in Fig. 5 for an adversarially supplied `privk*` (created with the help of the ECB encryption oracle), `c*` and `uh*`. It also requires working around some of the uncaught exceptions described in Section 2.3. To reach the needed error, the checks that trigger the earlier errors $\perp_1, \perp_4, \perp_5, \perp_6$ and \perp_7 must all be satisfied: `uh*` must be a UTF-8 string of size 11,

privk^* must encode q^*, p^*, d^*, u^* of sufficient length such that $\gcd(q^*, p^*) = 1$ and $\gcd(d^*, (p^* - 1)(q^* - 1)) = 1$ so that the corresponding inverses exist, $u^* = (q^*)^{-1} \bmod p^*$ and $c^* < N^*$ where $N^* = p^* \cdot q^*$.

Under these constraints, observe that $\text{DecryptSid}(sk, c^*)$ behaves differently depending on whether the second byte of the decrypted value $m^* \leftarrow (c^*)^{d^*} \bmod N^*$ is 00, where m^* is first zero-padded to the length of N^* to form \mathbf{m}^* . Suppose the server supplied p^*, q^* such that $|N^*|_{\mathbb{B}} = 256$. Let $\mathbf{m} \leftarrow \text{DecryptSid}(sk, c^*)$ and \mathbf{m}' be the intermediate value such that $\mathbf{m} = \mathbf{m}'[2 : |\mathbf{m}'|_{\mathbb{B}}]$. Then, based on the error returned by the client, the server can distinguish the following two cases:

- Case $\perp_2, 254$): This means that $|\mathbf{m}|_{\mathbb{B}} = 254$, so $|\mathbf{m}'|_{\mathbb{B}} = 256 = |N^*|_{\mathbb{B}} = |\mathbf{m}^*|_{\mathbb{B}}$, so the condition on line 9 was not satisfied, i.e. $\mathbf{m}^*[1] = 00$.¹²
- Case \perp_3 : This means that $|\mathbf{m}|_{\mathbb{B}} = 255$, so $|\mathbf{m}'|_{\mathbb{B}} = 257 = |N^*|_{\mathbb{B}} + 1 = |\mathbf{m}^*|_{\mathbb{B}} + 1$, which can only arise if $\mathbf{m}' = 00 \parallel \mathbf{m}^*$ and so $\mathbf{m}^*[1] \neq 00$.

A similar case analysis can be done for arbitrary values of $|N^*|_{\mathbb{B}}$; then the errors may be swapped. However due to the bugs in the modular power implementation in MEGA code, the attack actually only works for $|N^*|_{\mathbb{B}} \leq 256$.

We explain next how to exploit this behavioural difference to leak information about a target user's RSA private key.

The server constructs $[\text{privk}^*]_{\mathbf{k}_M}$ using the ECB encryption oracle such that in the “ d ” field it knows the plaintext for all blocks except the least-significant full block. That block will be the target of the attack; it can be an arbitrary AES-ECB-encrypted block $[B]_{\mathbf{k}_M}$. Let d^* denote the “ d ” component constructed in this way. The server must also precompute p^*, q^* of a special form and a number of values \mathbf{m}^* with $\mathbf{m}^*[1] = 00$ such that it can interpret one of the errors arising on decryption of a corresponding ciphertext as confirmation of a correct “guess”.

At a high level, the primes p^* and q^* are constructed so that $(p^* - 1)(q^* - 1)$ contains small prime factors r_i of a given bit length such that their product is at least 128 bits.¹³ Let $\mathbb{G} = (\mathbb{Z}/N^*\mathbb{Z})^\times$ so that $|\mathbb{G}| = (p^* - 1)(q^* - 1)$. For each factor r_i , the server computes $g_i \in \mathbb{G}$ such that g_i has order r_i and such that a value $t_i \in \{1, \dots, r_i - 1\}$ (or a set of such values \mathcal{T}_i) exists with the property that $g_i^{t_i} \bmod N^*$ has second byte 00 after zero-padding to the length of N^* . The value of u^* is then set to $(q^*)^{-1} \bmod p^*$.

Then, in the online phase of the attack, the server submits privk^* constructed using the ECB encryption oracle to contain q^*, p^*, d^*, u^* . For each r_i , it sets $x_i = 1, 2, \dots, r_i - 1$ and submits $c_{i,t}^* \leftarrow g_i^{x_i} \bmod N^*$ until the client returns the

¹² Note that the server does not know whether this is because prior to zero-padding, we have $|m^*|_{\mathbb{B}} \leq |N^*|_{\mathbb{B}} - 2$ and therefore trivially $\mathbf{m}^*[1] = 00$ or because $|m^*|_{\mathbb{B}} = |N^*|_{\mathbb{B}}$ and $\mathbf{m}^*[1] = 00$. However, the root cause is immaterial to our attack.

¹³ The factors do not need to be common between $(p^* - 1)$ and $(q^* - 1)$, and can be freely distributed between the two.

error that confirms the second byte of the decrypted value was 00 (which is $(\perp_2, 254)$ in the case that $|N^*|_B = 256$ which we will use in the attack). Then, based on the precomputed values it learns that, for the specific value x_i triggering the error, $x_i \cdot d^* \equiv t_i \pmod{r_i}$. Here d^* is a value that is known except for its least significant full block, where it contains B . From this equation, the value of $B \pmod{r_i}$ can be recovered. Finally, using CRT and taking some care with non-block-aligned inputs, allows recovery of the block B .

The attack is described in more detail in the following subsections, first a simpler but less-efficient version and then the full version. The ECB encryption oracle calls are shown as $[x]_{k_M} \leftarrow \text{OECB}_{k_M}(x)$ as before. Since both versions of the attack must “stitch” AES-ECB blocks together to create the final ciphertext, we provide the algorithm in Fig. 9 to avoid repetition. This algorithm combines the chosen values q^*, p^*, d', u^* so that they parse as expected, with the target block B being placed in the position of the least-significant full block of d^* and overwriting the corresponding block of d' . This is visualised in Fig. 10.

Stitch($q^*, p^*, d', u^*, [B]_{k_M}$)

```

1 :  $pt_0 \leftarrow \text{len}(q^*) \parallel q^* \parallel \text{len}(p^*) \parallel p^* \parallel \text{len}(d') \parallel d'[0 : 234]$ 
2 :  $pad \leftarrow \{0, 1\}^8$  // random padding, could also be 00s
3 :  $pt_1 \leftarrow d'[250 : 256] \parallel \text{len}(u^*) \parallel u^* \parallel pad$ 
4 :  $[pt_0]_{k_M} \leftarrow \text{OECB}_{k_M}(pt_0)$ 
5 :  $[pt_1]_{k_M} \leftarrow \text{OECB}_{k_M}(pt_1)$ 
6 :  $ct^* \leftarrow [pt_0]_{k_M} \parallel [B]_{k_M} \parallel [pt_1]_{k_M}$ 
7 : return  $ct^*$ 
```

Fig. 9. Combining modified values produced using the ECB encryption oracle with the target ciphertext block in the correct format, reusing known AES-ECB blocks where possible. This assumes that $|q^*|_B = |p^*|_B = |u^*|_B = 128$ and $|d'|_B = 256$, as is the case for legitimate MEGA keys.

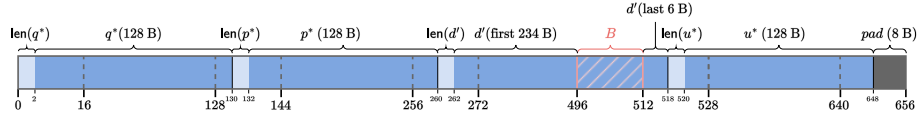


Fig. 10. The plaintext content of ct^* , with the placement of the target block B in red.

4.1 Simplified version

This version of the attack assumes a single t_i value per factor, which simplifies the presentation but imposes a high cost at the precomputation stage. Further, there is a non-negligible probability of the attack aborting and thus failing to complete. We will remove this restriction in the full version of the attack below.

Precomputation Take $\{r_0, \dots, r_{n-1}\}$ where each r_i is a prime such that $|r_i|_b = 8$, and n is such that $\left| \prod_{i=0}^{n-1} r_i \right|_b \geq 128$. This imposes the constraint $16 \leq n \leq 19$. Let $[B]_{\mathbf{k}_M}$ be the target ciphertext block.

1. Find primes p^*, q^* such that $|p^*|_b = |q^*|_b = 1024$ and

$$p^* = 2 \cdot \left(\prod_{i=0}^{\lceil n/2 \rceil - 1} r_i \right) \cdot p' + 1, \quad q^* = 2 \cdot \left(\prod_{i=\lceil n/2 \rceil}^{n-1} r_i \right) \cdot q' + 1$$

where p', q' is each a product of 2-4 large primes.¹⁴ Encode p^*, q^* as byte strings $\mathbf{p}^*, \mathbf{q}^*$.

2. Set $N^* \leftarrow p^* \cdot q^*$ and $\mathbb{G} \leftarrow (\mathbb{Z}/N^*\mathbb{Z})^\times$.
3. For $i \in \{0, \dots, n-1\}$:
 - (a) Find $g_i \in \mathbb{G}$ of order r_i , e.g. by sampling $h \leftarrow \mathbb{G}$ and computing $g_i \leftarrow h^{(p^*-1)(q^*-1)/r_i} \bmod N^*$ until $g_i \neq 1$.
 - (b) Find a value $t_i \in \{1, \dots, r_i - 1\}$ such that for $m \leftarrow g_i^{t_i} \bmod N^*$; $\mathbf{m} \leftarrow \text{ZeroPad}(m, N^*)$, we have $\mathbf{m}[1] = 00$. If no such t_i is found or there are multiple possible values, restart the precomputation.
4. Compute $u^* \leftarrow (q^*)^{-1} \bmod p^*$ and encode it as a byte string \mathbf{u}^* with $|\mathbf{u}^*|_b = 1024$.
5. Let $d' \leftarrow 2^{2047} + 1$ and encode it as a byte string \mathbf{d}' with $|\mathbf{d}'|_b = 2048$.
6. Obtain $\mathbf{ct}^* \leftarrow \text{Stitch}(\mathbf{q}^*, \mathbf{p}^*, \mathbf{d}', \mathbf{u}^*, [B]_{\mathbf{k}_M})$. Let $d^* \leftarrow d' + 2^{48} \cdot B$ (where B is the unknown target block) denote the unknown value in the “ d ” field that will arise on decrypting \mathbf{ct}^* .¹⁵

Success probability. For random $m \in \mathbb{G}$ we have $\Pr[\mathbf{m}[1] = 00] = 2^{-8}$. For each factor r_i the probability that Step 3b finds exactly one suitable t_i is $(r_i - 1) \cdot 2^{-8} \cdot (1 - 2^{-8})^{r_i - 2}$, which is greater than 0.18 for $2^7 < r_i < 2^8$. However, this needs to occur for all n factors where $n \geq 16$ to get a product of sufficient length to recover B using CRT, so the overall success probability is of the order

¹⁴ These primes could repeat, the goal here is to avoid $(p^* - 1)(q^* - 1)$ having any other small factors except for r_0, \dots, r_{n-1} .

¹⁵ Note that by the choice of d' , overwriting the least significant full block of d' with B is equivalent to adding $2^{48} \cdot B$ to d' .

$\approx 2^{-39}$ or less. To reduce the required amount of precomputation, in Section 4.2 we increase the bit length of each factor to ensure that there is at least one suitable t_i for each r_i and provide a strategy to disambiguate between multiple fitting t_i values.

Online attack Let $R = \prod_{i=0}^{n-1} r_i$ and $\text{ct}^*, \{g_i\}_{i \in \mathcal{I}}, \{t_i\}_{i \in \mathcal{I}}$ be as computed before, for $\mathcal{I} = \{0, \dots, n-1\}$.

1. When the client initiates a login, respond to the client's request with $([\mathbf{k}_M]_{\mathbf{k}_e}, \text{ct}^*, [\mathbf{m}]_{pk}, \mathbf{uh})$, where everything but ct^* is as it would be in an honest response. If the client returns \perp_6 , abort.
2. For $i \in \{0, \dots, n-1\}$, do the following:
 - (a) For $x \in \{1, \dots, r_i-1\}$, do the following:
 - i. Compute $c_{i,x}^* \leftarrow (g_i)^x \bmod N^*$.
 - ii. When the client initiates a login, respond to the client's request with $([\mathbf{k}_M]_{\mathbf{k}_e}, \text{ct}^*, c_{i,x}^*, \mathbf{uh})$, where everything but ct^* and $c_{i,x}^*$ is as it would be in an honest response.¹⁶
 - iii. If the client returns $(\perp_2, 254)$, save the value of x and break out of this loop.
 - (b) If there is a saved value x , then we have $d^* \equiv x^{-1} \cdot t_i \pmod{r_i}$ for unknown d^* .
3. Then, use CRT to compute $d^* \bmod R$ from the values collected in Step 2b. Recall that by construction $d^* = d' + 2^{48} \cdot B$, so $d^* = 2^{2047} + 2^{48} \cdot B + 1$. Hence compute

$$B \equiv (2^{48})^{-1} \cdot (d^* - 2^{2047} - 1) \pmod{R},$$

to recover the target plaintext block since $|R|_{\mathbf{b}} \geq 128$.

Cost. In the worst case, the main cost of the online attack is $\sum_{i \in \mathcal{I}} (r_i - 1)$ login attempts. This is bounded from above by $n \cdot (2^8 - 1) \approx 2^{12.24}$ for $n \leq 19$. In the average case, for each i we expect Step 2a to conclude after approximately $\frac{1}{2} \cdot 2^8$ trials, so the overall bound becomes $n \cdot 2^7 \approx 2^{11.24}$ for $n \leq 19$.

Probability of abort. Note that the attack aborts if it receives error \perp_6 . This error is returned whenever the decrypted $d^* = d' + 2^{48} \cdot B$ is such that $\gcd(d^*, (p^* -$

¹⁶ An honest response refers to the data that an honest server would have sent. Note that in this case, the “honest” \mathbf{uh} will not match the value recovered from $c_{i,x}^*$, but this check only comes after the errors triggered by the attack. The attacker could equally replace the \mathbf{uh} value with an arbitrary 11-byte UTF-8 string.

$1)(q^* - 1)) \neq 1$. Since d^* is odd by construction,¹⁷ the error can only be caused if at least one of the following is true:

- $d^* \equiv 0 \pmod{r_i}$ for at least one r_i ,
- $d^* \equiv 0 \pmod{p'_j}$ for at least one $p'_j \mid p'$, or
- $d^* \equiv 0 \pmod{q'_k}$ for at least one $q'_k \mid q'$.

The values p'_j, q'_k are large primes by construction, so the probability of an abort being caused by those cases is negligible. However, each factor r_i is only 8 bits in size, which means that assuming a random B the probability that the attack aborts because $d^* \equiv 0 \pmod{r_i}$ for at least one r_i is bounded by $n \cdot 2^{-7} \approx 0.15$ with $n \leq 19$. In Section 4.2, we discuss strategies for avoiding the abort.

Correctness. Now, assume the attack does not abort. By construction, the values of q^*, p^*, d', u^* pass the check on bit length, we have $\gcd(q^*, p^*) = 1$, $u^* = (q^*)^{-1} \pmod{p^*}$ and all $c_{i,x}^* < N^*$. During $\text{DecryptSid}(sk, c_{i,x}^*)$, the client will compute $m = (c_{i,x}^*)^{d^*} \pmod{N^*} = (g_i)^{x \cdot d^*} \pmod{N^*}$. If it is the case that $m = (g_i)^{t_i} \pmod{N^*}$ and therefore $x \cdot d^* \equiv t_i \pmod{r_i}$, the second byte of zero-padded m will be 00 and so the client will return $(\perp_2, 254)$ to the server. Otherwise, it will proceed with the computation and with very high probability return \perp_3 , since the uh value will not match the relevant substring of m . Hence the attack recovers the target plaintext block.

4.2 Full version

Here, we provide strategies to improve the running time and the success probability of our second attack. First, we discuss the use of multiple t_i values per factor r_i , incorporate this into the attack and show the effect of this strategy. For practical purposes, this strategy is already sufficient to reduce the precomputation cost and the likelihood of aborts.

In this version of the attack, we increase the bit length of the factors r_i . As a result, the probability of finding a suitable t_i value during precomputation is increased. However this also implies that there will be more than one such value. We therefore have to also amend the online part of the attack to provide a way of determining which $t \in \mathcal{T}_i$ value has caused the expected error for a given x . There are multiple ways in which this could be achieved, and here we describe one option.

Take r_i, \mathcal{T}_i and assume that we got the $(\perp_2, 254)$ error for some $x \in \{1, \dots, r_i - 1\}$. We can test each potential value $t_j \in \mathcal{T}_i$ by submitting another query $c_{i,x_j}^* \leftarrow (g_i)^{x_j} \pmod{N^*}$ where $x_j \leftarrow x \cdot t_j^{-1} \pmod{r_i}$. If the guess for t_j is correct, we have

¹⁷ This is also why we cannot make the block-aligned simplification for this attack, because if we aligned it such that the least-significant block of d^* is full and therefore placed our target block B there, then if $B \equiv 0 \pmod{2}$ the client would output error \perp_6 on all queries.

$x \cdot d^* \equiv t_j \pmod{r_i}$, and so decryption of c_{i,x_j}^* will produce $(g_i)^{x_j \cdot d^*} \bmod N^* = (g_i)^{x \cdot t_j^{-1} \cdot d^*} \bmod N^* = g_i$ as the plaintext. Then, as long as g_i is such that its second byte is not 00, which we can ensure in the precomputation phase, the check that produces \perp_2 will pass. Since the server knows g_i and is able to set `uh` to arbitrary 11-byte values, it can also make sure to pass the check that produces \perp_3 , and therefore get 43 bytes of g_i from the client via the returned `sid` value when the guess is correct. However, if the guess is not correct, it is very unlikely that the server-modified `uh` would match the resulting plaintext, leading to \perp_3 . So the server can distinguish between the two cases.

Precomputation Take $\{r_0, \dots, r_{n-1}\}$ where each r_i is a prime such that $|r_i|_b = 12$, and n is such that for $R \leftarrow \prod_{i=0}^{n-1} r_i$ we have $|R|_b \geq 128$, so $11 \leq n \leq 12$. Let $[B]_{\mathbb{K}_M}$ be the target ciphertext block.

1. Find primes p^*, q^* such that $|p^*|_b = |q^*|_b = 1024$ and

$$p^* = 2 \cdot \left(\prod_{i=0}^{\lceil n/2 \rceil - 1} r_i \right) \cdot p' + 1, \quad q^* = 2 \cdot \left(\prod_{i=\lceil n/2 \rceil}^{n-1} r_i \right) \cdot q' + 1$$

where p', q' is each a product of 2-4 large primes. Encode p^*, q^* as byte strings $\mathbf{p}^*, \mathbf{q}^*$.

2. Set $N^* \leftarrow p^* \cdot q^*$ and $\mathbb{G} \leftarrow (\mathbb{Z}/N^*\mathbb{Z})^\times$.
3. For $i \in \{0, \dots, n-1\}$:
 - (a) Find $g \in \mathbb{G}$ of order r_i , e.g. by sampling $h \leftarrow \$ \mathbb{G}$ and computing $g \leftarrow h^{(p^*-1)(q^*-1)/r_i} \bmod N^*$ until $g \neq 1$.
 - (b) Initialise $\mathcal{T}_i = \emptyset$.
 - (c) For $t \in \{1, \dots, r_i - 1\}$, do the following:
 - i. Let $\mathbf{g}' \leftarrow \text{ZeroPad}(g^t, N^*)$ for $g^t \leftarrow g^t \bmod N^*$.
 - ii. If $\mathbf{g}'[1] = 00$, add t to \mathcal{T}_i . Else if $\mathbf{g}'[17 : \alpha]$ for some $\alpha \geq 28$ is a valid UTF-8 string of size 11¹⁸, save $g_i \leftarrow g^t$, $a \leftarrow t$ and $\mathbf{uh}_i^* \leftarrow \mathbf{g}'[17 : \alpha]$.
 - (d) If $\mathcal{T}_i = \emptyset$ or a is undefined, restart the precomputation.
 - (e) Shift \mathcal{T}_i by replacing each $t \in \mathcal{T}_i$ by $t \cdot a^{-1} \bmod r_i$. This ensures that the values in \mathcal{T}_i are with respect to the new generator g_i instead of g .
4. Compute $u^* \leftarrow (q^*)^{-1} \bmod p^*$ and encode it as a byte string \mathbf{u}^* with $|\mathbf{u}^*|_b = 1024$.

¹⁸ Note that an 11 B byte string interpreted as a valid UTF-8 string will likely not be a string of size 11, i.e. a string consisting of 11 characters, since not all byte values are interpreted as text and non-ASCII characters require multiple bytes to encode [38].

5. Compute $d' \leftarrow 2^{2047} + 2^{48+128} \cdot \delta + 1$ for $\delta < R$ such that $d' \equiv 0 \pmod{R}$. Encode it as a byte string \mathbf{d}' with $|\mathbf{d}'|_{\text{b}} = 2048$.
6. Obtain $\mathbf{ct}^* \leftarrow \text{Stitch}(\mathbf{q}^*, \mathbf{p}^*, \mathbf{d}', \mathbf{u}^*, [B]_{\mathbf{k}_M})$.

Success probability. Increasing the bit length of the factors means that now for each factor r_i the probability that Step 3(c)ii finds at least one suitable t is $1 - (1 - 2^{-8})^{r_i-1}$, which is greater than 0.9996 for $2^{11} < r_i < 2^{12}$. Across all n factors for $n \leq 12$, it is still greater than 0.99. Next, the probability that a random 11-byte string is a valid UTF-8 string is ≈ 0.001634 . Hence for each factor r_i the probability that at least one such string will be found is $1 - (1 - 0.001634)^{r_i-1} > 0.9648$, and across all factors it is at least 0.65. In practice, if the precomputation fails at this point, it can simply be re-run again with different r_i values.

Cost. This version tests all possible values of t for every r_i , so overall it must check at most $n \cdot 2^{12} \approx 2^{15}$ values of g^t (these can however be cycled through for each r_i). The prime generation is a one-time cost in the sense that the values can be reused in attacks on multiple users. Finally, since d' will be composed mostly of zero-blocks, building the ciphertext \mathbf{ct}^* requires up to 15 uses of the ECB encryption oracle (which, recall, produces 2 blocks at a time).

Online attack Let $\mathbf{ct}^*, \{g_i\}_{i \in \mathcal{I}}, \{\mathcal{T}_i\}_{i \in \mathcal{I}}$ be the values computed before where $\mathcal{I} = \{0, \dots, n-1\}$.

1. When the client initiates a login, respond to the client's request with $([\mathbf{k}_M]_{\mathbf{k}_0}, \mathbf{ct}^*, [\mathbf{m}]_{pk}, \mathbf{uh})$, where everything but \mathbf{ct}^* is as it would be in an honest response. If the client returns \perp_6 , abort.
2. For $i \in \mathcal{I}$, do the following:
 - (a) For $x \in \{2, \dots, r_i - 1\}$, do the following:
 - i. Compute $c_{i,x}^* \leftarrow (g_i)^x \bmod N^*$.
 - ii. When the client initiates a login, respond to the client's request with $([\mathbf{k}_M]_{\mathbf{k}_0}, \mathbf{ct}^*, c_{i,x}^*, \mathbf{uh})$, where everything but \mathbf{ct}^* and $c_{i,x}^*$ is as it would be in an honest response.
 - iii. If the client returns $(\perp_2, 254)$, save the value of x and break out of this loop.
 - (b) If $\mathcal{T}_i = \{t_i\}$ has a single element, skip this step. Otherwise, for $t \in \mathcal{T}_i$, do the following:
 - i. Let $x' \leftarrow x \cdot t^{-1} \bmod r_i$.
 - ii. Compute $c_{i,x'}^* \leftarrow (g_i)^{x'} \bmod N^*$.

- iii. When the client initiates a login, respond to the client's request with $([\mathbf{k}_M]_{\mathbf{k}_e}, \mathbf{ct}^*, \mathbf{c}_{i,x'}^*, \mathbf{uh}_i^*)$, where only $[\mathbf{k}_M]_{\mathbf{k}_e}$ is as it would be in an honest response.
 - iv. If the client returns $\mathbf{sid} = \mathbf{g}_i[1 : 44]$, save the value $t_i \leftarrow t$ and break out of this loop.
- (c) We have that $d^* = d' + 2^{48} \cdot B \equiv x^{-1} \cdot t_i \pmod{r_i}$, and so $B \equiv (2^{48})^{-1} \cdot x^{-1} \cdot t_i \pmod{r_i}$.

3. Then, use CRT to compute $B \bmod R$ from the values collected in Step 2c, which in turn recovers the target plaintext block since $|R|_b \geq 128$.

Success probability. As in the attack in Section 4.1, this attack aborts if it receives error \perp_6 . However, the probability that this happens becomes smaller with the increased bit length of the factors r_i . Assuming a random B , for 12-bit factors the probability of an abort is bounded by $n \cdot 2^{-11} \approx 0.006$ with $n \leq 12$. In the full version of this work we give a more complex attack strategy that avoids the abort altogether.

In practice, the attack's success probability may be impacted by another factor, namely differing implementations of UTF-8 validation. Suppose that the values \mathbf{g} produced in Step 3(c)ii of the precomputation in Section 4.2 have valid UTF-8 substrings of size 11 in Python: this does not guarantee that they will be interpreted as such by the Javascript webclient. This requires implementing additional strategies for disambiguation in case the UTF-8-based one never yields the expected \mathbf{sid} request.¹⁹

Cost. In the worst case, the main cost of the online phase of the attack is the $\sum_{i \in \mathcal{I}} (r_i - 1)$ login attempts needed. This is bounded by $n \cdot (2^{12} - 1) \approx 2^{15.58}$ for $n \leq 12$. In the average case, for each i we expect Step 2a to conclude after at most 2^8 trials and Step 2b to finish after around $\frac{1}{2} \cdot |\mathcal{T}_i| \approx \frac{1}{2} \cdot r_i \cdot 2^{-8}$ trials. Added together, the number of login attempts needed in the average case is bounded by $n \cdot (2^8 + \frac{1}{2} \cdot 2^{12} \cdot 2^{-8}) \approx 2^{11.63}$ for $n \leq 12$. Performing the experimental analysis over a large number of runs as in Section 3.2 would be more difficult due to the interaction between the disambiguation strategies and the web client with automated logins, which causes the web client to freeze or begin sending requests in large batches. This can impact the success rate (in particular, the attack may produce one x or t value that is slightly off) and hinders automating the attack. We stress that this is purely an artefact of our proof-of-concept implementation.

¹⁹ One alternative is to instead for all $t \in \mathcal{T}_i$ submit $x' \leftarrow x \cdot t^{-1} \cdot t_j \bmod r_i$ for some $t_j \in \mathcal{T}_i, t_j \neq t$, and use the original error $(\perp_2, 254)$ as the confirmation signal. This still has a potential for false positives and false negatives, however. A final, and most expensive, failover strategy is then to cycle through all values of x , saving the ones for which the client returns $(\perp_2, 254)$ and then running an offline computation to determine which x values are matched to which t values.

Note that to keep the presentation of the attacks simpler, we have assumed specific values of $|r_i|_b$ and thus constrained the value of n . In reality, using different values would allow making a different tradeoff between the precomputation cost and the number of login attempts needed in the online phase. For instance, using 10-bit primes would lower the (online) worst-case bound to $n \cdot (2^{10} - 1) \approx 2^{13.91}$ for $n \leq 15$, but slightly increase the (online) average-case bound to $n \cdot (2^8 + \frac{1}{2} \cdot 2^{10} \cdot 2^{-8}) \approx 2^{11.92}$ login attempts. It would also make the precomputation phase much less likely to succeed in a single run: the probability of finding suitable t values for all r_i would fall to around 0.11, while the probability of finding generators with suitable UTF-8 substrings for all r_i would only be around 0.0002.

5 Recovering the RSA private key

Our attacks in Sections 3 and 4 can be seen as building generic AES-ECB decryption oracles. In this section, we turn this capability into an RSA private key recovery attack. Naively we would expect to call our costly AES-ECB decryption oracle up to nine times: each factor p, q of N has 1024 bits, but these are not perfectly aligned with AES block boundaries, necessitating to cover (partial) plaintexts from nine different 128-bit blocks. However, using a post-processing stage, we can reduce this number to four.

In particular, as illustrated in Fig. 4, the block alignments of p and q differ. For reasons that will become apparent below we will need to recover at least 512 bits. Based on the specific alignments, we will aim to recover the 512+16 least significant bits of q : 512 bits (i.e. four 128-bit blocks) are recovered using the attacks from Sections 3 and 4 and the least significant 16 bits are “recovered” using exhaustive search (which avoids the query cost of recovering a fifth block). If instead we targeted p , we would need to recover 32 bits using exhaustive search, which would have prohibitive cost. Thus, next, we discuss how to recover the remaining bits of q given the $\ell = 512 + 16$ least significant bits of q . In particular, we will solve the following computational problem.

Definition 1. *Let $N = p \cdot q$ be a 2048-bit RSA modulus with p, q having 1024 bits each. Given ℓ consecutive least significant bits of q , recover q .*

Our approach is a simple combination of exhaustive search, lattice reduction and root finding over \mathbb{Z} following Coppersmith’s method [6]. In particular, we use the Howgrave-Graham variant [8, 9, 14, 29] of this algorithm. Let $\lceil \log_2 q \rceil - \ell < 1024$, $q = 2^{\lceil \log_2 q \rceil - \ell} \cdot r + q'_0$, where r are the bits we are trying to recover and $|q'_0| \leq 2^\ell$ are the known bits of q . Then r satisfies $f'(x) \equiv 0 \pmod q$ for $f'(x) := q'_0 + 2^{\lceil \log_2 q \rceil - \ell} \cdot x \pmod q$. Given this we can consider

$$q_0 := 2^{-\lceil \log_2 q \rceil + \ell} \cdot q'_0 \quad \text{and} \quad f(x) := q_0 + x \pmod q$$

and note that r still satisfies $f(x) \equiv 0 \pmod q$. That is, we translate our problem into one where the most significant bits are known rather than the least significant ones, cf. [29].

From this, the algorithm proceeds by constructing several polynomials that evaluate to zero modulo q or a multiple thereof, such as (powers of) N . In more detail, Let $h \geq 2 \in \mathbb{N}$ and $u < h \in \mathbb{N}$, for $0 \leq i < h$ we let

$$f_i(x) := \begin{cases} N^{u-i} \cdot (q_0 + x)^i & \text{for } 0 \leq i < u, \\ x^{i-u} \cdot (q_0 + x)^u & \text{for } u \leq i < h. \end{cases}$$

For example, picking $h = 4$ and $u = 2$ we get

$$N^2, \quad N \cdot q_0 + N \cdot x, \quad q_0^2 + 2q_0 \cdot x + x^2 \quad \text{and} \quad q_0^2 \cdot x + 2q_0 \cdot x^2 + x^3.$$

First, note that all $f_i(x)$ evaluate to zero modulo q^u at the correct r . Second, note the maximal degree of the $f_i(x)$ is $h - 1$, i.e. $\max_{0 \leq i < h} (\deg(f_i(x))) = h - 1$ and thus each polynomial has at most h coefficients.

Now, letting $X = 2^{\lceil \log_2 q \rceil - \ell}$ and $f_i^{(j)}$ denote the coefficient of x^j in $f_i(x)$, we construct a matrix \mathbf{A} where the entry $A_{i,j} := f_i^{(j)} \cdot X^j$. Continuing with our example, we would have

$$\mathbf{A} := \begin{pmatrix} N^2 & 0 & 0 & 0 \\ N \cdot q_0 & N \cdot X & 0 & 0 \\ q_0^2 & 2q_0 \cdot X & X^2 & 0 \\ 0 & q_0^2 \cdot X & 2q_0 \cdot X^2 & X^3 \end{pmatrix}.$$

Since the matrix is triangular we can read off the determinant $\det(\mathbf{A}) = N^{u \cdot (u+1)/2} \cdot X^{h \cdot (h-1)/2}$. The rows of this matrix \mathbf{A} span a lattice which contains a vector \mathbf{v} of Euclidean norm $\|\mathbf{v}\| \leq \sqrt{h} \cdot (N^{u \cdot (u+1)/2} \cdot X^{h \cdot (h-1)/2})^{1/h}$ by Minkowski's theorem. In other words, there exists an integer-linear combination of the rows of \mathbf{A} that produces a vector with at most this Euclidean norm. Using lattice reduction we can find this shortest vector.²⁰ Now, given a vector of Euclidean norm $\|\mathbf{v}\|$ we know that its ℓ_1 norm, i.e. the sum of the absolute values of its entries, is bounded by $|\mathbf{v}|_1 \leq \sqrt{h} \cdot \|\mathbf{v}\|$. Finally, if $\mathbf{v} \neq 0$ and $|\mathbf{v}|_1 \leq q^u$, we can extract a polynomial that evaluates to zero modulo q^u on r but which evaluated at r is strictly smaller than q^u .²¹ In other words, this polynomial evaluates to zero at r over \mathbb{Z} . The algorithm concludes by finding the roots of this polynomial, which can be accomplished in polynomial time (and efficiently in practice).

To select h and u , by abuse of notation let h also be a formal variable and set $u := 1/2 \cdot h - 1$. As in [8, p.102], we then find a root > 0 of

$$\frac{1024 - \ell}{2048} \cdot h \cdot (h - 1) - u \cdot h + u \cdot (u + 1).$$

²⁰ The traditional presentation of this algorithm invokes the LLL algorithm which gives a short vector that is at most an exponential factor away from the shortest vector. However, the lattice dimensions involved here are in the range where the shortest vector problem (SVP) can be solved efficiently in practice – say, up to dimension 150 [7] – and we may thus simply assume we solve SVP. In any case, the exponential factor is $\approx 1.0219^h$ which is < 3 for $h \leq 50$.

²¹ We extract $g(x)$ as $g^{(j)} := v_j / X^j \in \mathbb{Z}$.

This succeeds for $\ell > 512$ and the solution grows as ℓ approaches 512 from above.

As mentioned above, in our setting we consider $\ell = 512 + 16$: We run the attack from Section 3 on four blocks to recover 512 bits and run an exhaustive search over the remaining 16 bits (which are contained in a non-aligned block). In this setting, we picked $h = 36$ and $u = 18$. In our experiments, using LLL, finding a sufficiently short vector takes about 26 seconds on a Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz using SageMath/FPLLL [33, 34]. In 1024 experiments, we obtained a success rate of 100%. Thus, we expect to be able to recover q in time $2^{16} \cdot 26$ seconds, or about 20 core days.²²

The overall cost of the RSA private key recovery attack is $4 \cdot 2^{9.29} = 2^{11.29} \approx 2500$ login attempts, 66 ECB encryption oracle calls, and about 20 core days of computation (using the attack of Section 3 in combination with the attack in this section).

6 Attacking unpatched clients

We briefly revisit the attacks of [2] against unpatched MEGA clients in the light of our discovery of the ECB encryption oracle described in Section 2.2.

Attack 1 in [2] uses an estimated 512 logins to recover a target user’s RSA private key. The number of logins required was subsequently reduced to 6 in [31] by using more sophisticated lattice techniques.

Attack 2 in [2] then exploits knowledge of that private key to recover two blocks of AES-ECB plaintext per login. This is done by overwriting two blocks of the encrypted version of u with the target AES-ECB ciphertext blocks and selecting a carefully crafted RSA ciphertext in the authentication protocol; the session ID returned by the client in that protocol then leaks the two AES-ECB plaintext blocks. This approach is used to build an efficient procedure for recovering file encryption keys in [2].

Interestingly, however, the RSA private key used in Attack 2 in [2] does not need to be the target user’s true key – it only needs to be a key known to the adversary and any valid RSA private key (in the appropriate format) will do. Hence, an adversary can use the ECB encryption oracle to create a suitably encrypted, known RSA private key. By carefully reusing all-zero blocks for most of q , p and d , the number of ECB encryption oracle calls needed can be made as small as 7. The adversary then applies Attack 2 from [2] with the target AES-ECB ciphertext blocks being selected from those encrypting the least significant bits

²² We note that this computation is “proudly parallel” or “embarrassingly parallel”. This is because for each of our 2^{16} guesses we can run an independent lattice reduction. We also note that the running time is independent of whether the input instance corresponds to a correct or incorrect guess. Moreover, incorrect solutions resulting from incorrect guesses can be filtered out using the known public key.

of q (from the actual private key). With two applications of the attack, the adversary recovers 4 plaintext blocks, or 512 bits of q . Applying the lattice attack from Section 5, the adversary recovers the full RSA private key.

The cost of the attack is 2 login attempts and a small number of ECB encryption oracle calls.

Note that Attack 2 of [2] is prevented by patched MEGA client code because of the requirement that the client-selected 11-byte string `uh` appear in `m` at a specific location and because overwriting u with a target ciphertext would make the check at line 10 in `DecodePrivk(privk)` fail.

7 Discussion and future work

On the one hand, the conclusion to be drawn from this work for practitioners and designers is no different from the one derived from [2]. The root causes at play here were already identified in [2], whose suggestion of protecting the integrity of encrypted keys using a MAC would have prevented the attacks in this work as well. Further, the existence of the ECB encryption oracle in a feature completely separate from the attacked protocol highlights the continued fragility of the MEGA infrastructure, made possible also by the lack of key separation.

However, our attacks also highlight issues going beyond the ones exposed in previous works. First, some of the errors that our attacks exploit as oracles are not explicit, but derive from bugs in the big integer arithmetic provided by `asmcrypto.js`. This presents a challenge already mentioned in [3] which called for a verified big integer library that could serve as a common core for different projects. In the case considered here, such a library would need to be cross-compilable to JavaScript or WebAssembly. We consider this a pressing area for future work.

There are also further lessons to be drawn for a cryptanalytic audience. First, our attacks serve as an additional example of key overwriting attacks [2, 5, 10], a class of attacks that appears to deserve more exploration in terms of targets (deployed in practice) and attack refinement. Moreover, our attacks make use of the detailed and verbose error reporting by MEGA clients. This enables powerful side-channel attacks that can be observed remotely,²³ highlighting the practical significance of these classes of attacks. Finally, our work, along with other recent works attacking widely deployed protocols such as [1, 5, 11, 32, 37], underlines that while it might seem that the “golden age” of cryptographic attacks against deployed protocols is over – given the level of academic involvement and formal rigour that went into the design of TLS 1.3 – the target has simply moved up the stack. As cryptographic applications move beyond “simple” protection of data in transit or at rest, more complex cryptographic solutions are deployed at scale, often without significant input from the cryptographic community. This

²³ In contrast to timing-based side-channel attacks, generally considered less practical in the remote, as opposed to local, setting.

suggests a broad and impactful field for cryptanalysis of targets “in the wild”. It is well known that attacks are typically required to convince practitioners to adopt cryptographic recommendations. This in turn suggests that to achieve the adoption of more secure and formally analysed cryptographic solutions in practice, further cryptanalytical work on the “current generation” of deployed solutions is needed.

Finally, the two attacks presented in this work require a large number of login attempts. This was also the case for the first attack of [2] and used as an argument by MEGA that the attack was not practical. However, later work by [31] reduced the number of login attempts to six, and we have further reduced it to just two. Beyond reinforcing the truism that attacks only get better, this poses the open problem to improve the attacks presented in this work in terms of login attempt complexity.

Acknowledgements

The research of Mareková was carried out in part during a visit to the Applied Cryptography Group at ETH Zürich. She was also supported by the EPSRC and the UK Government as part of the Centre for Doctoral Training in Cyber Security at Royal Holloway, University of London (EP/P009301/1). The work of Paterson was supported in part by a gift from VMware. The work of Albrecht was done while Albrecht was at Royal Holloway.

References

1. Albrecht, M.R., Mareková, L., Paterson, K.G., Stepanovs, I.: Four attacks and a proof for Telegram. In: 43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022. pp. 87–106. IEEE (2022), <https://doi.org/10.1109/SP46214.2022.9833666>
2. Backendal, M., Haller, M., Paterson, K.G.: MEGA: Malleable Encryption Goes Awry. In: 44th IEEE Symposium on Security and Privacy, to appear (2023), <https://eprint.iacr.org/2022/959>
3. Barbosa, M., Barthe, G., Bhargavan, K., Blanchet, B., Cremers, C., Liao, K., Parno, B.: SoK: Computer-aided cryptography. In: 2021 IEEE Symposium on Security and Privacy. pp. 777–795. IEEE Computer Society Press (May 2021). <https://doi.org/10.1109/SP40001.2021.00008>
4. Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In: Krawczyk, H. (ed.) CRYPTO’98. LNCS, vol. 1462, pp. 1–12. Springer, Heidelberg (Aug 1998). <https://doi.org/10.1007/BFb0055716>
5. Bruseghini, L., Paterson, K.G., Huigens, D.: Victory by KO: Attacking OpenPGP using key overwriting. In: ACM Conference on Computer and Communications Security (ACM CCS), to appear (2022), <https://doi.org/10.3929/ethz-b-000545839>
6. Coppersmith, D.: Finding a small root of a bivariate integer equation; factoring with high bits known. In: Maurer [13], pp. 178–189. https://doi.org/10.1007/3-540-68339-9_16

7. Ducas, L., Stevens, M., van Woerden, W.P.J.: Advanced lattice sieving on GPUs, with tensor cores. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part II. LNCS, vol. 12697, pp. 249–279. Springer, Heidelberg (Oct 2021). https://doi.org/10.1007/978-3-030-77886-6_9
8. Howgrave-Graham, N.A.: Computational Mathematics Inspired by RSA. Ph.D. thesis, University of Bath (1998), available at <https://researchportal.bath.ac.uk/en/studentTheses/computational-mathematics-inspired-by-rsa>
9. Howgrave-Graham, N.: Finding small roots of univariate modular equations revisited. In: Darnell, M. (ed.) 6th IMA International Conference on Cryptography and Coding. LNCS, vol. 1355, pp. 131–142. Springer, Heidelberg (Dec 1997)
10. Klima, V., Rosa, T.: Attack on private signature keys of the OpenPGP format, PGP(TM) programs and other applications compatible with OpenPGP. Cryptology ePrint Archive, Report 2002/076 (2002), <https://eprint.iacr.org/2002/076>
11. Len, J., Grubbs, P., Ristenpart, T.: Partitioning oracle attacks. In: Bailey, M., Greenstadt, R. (eds.) USENIX Security 2021. pp. 195–212. USENIX Association (Aug 2021)
12. Lim, C.H., Lee, P.J.: A key recovery attack on discrete log-based schemes using a prime order subgroup. In: Kaliski Jr., B.S. (ed.) CRYPTO’97. LNCS, vol. 1294, pp. 249–263. Springer, Heidelberg (Aug 1997). <https://doi.org/10.1007/BFb0052240>
13. Maurer, U.M. (ed.): EUROCRYPT’96, LNCS, vol. 1070. Springer, Heidelberg (May 1996)
14. May, A.: Using LLL-reduction for solving RSA and factorization problems. pp. 315–348. ISC, Springer, Heidelberg (2010). <https://doi.org/10.1007/978-3-642-02295-1>
15. MEGA: About Us. <https://mega.io/about> (September 2022)
16. MEGA: Mega.nz web client (2022), <https://github.com/meganz/webclient>
17. MEGA: Security White Paper. <https://mega.nz/SecurityWhitepaper.pdf> (June 2022)
18. MEGA: webclient – #15273: Patch for ETH Zurich exploit. <https://github.com/meganz/webclient/commit/d2a0d054d4dbb90f035b3b4b421f780adafaa78e> (June 2022)
19. MEGA: webclient – #15295: Output detailed information about RSA decoding failures. <https://github.com/meganz/webclient/commit/cd4ab89b2cd0e388b0ea55753b86c8808f810138> (June 2022)
20. MEGA: webclient – asmcrypto.js. <https://github.com/meganz/webclient/blob/v4.21.4/js/vendor/asmcrypto.js> (August 2022)
21. MEGA: webclient – asmcrypto.js: Modulus. <https://github.com/meganz/webclient/blob/v4.21.4/js/vendor/asmcrypto.js#L10325> (August 2022)
22. MEGA: webclient – asmcrypto.js: Modulus_inverse. <https://github.com/meganz/webclient/blob/v4.21.4/js/vendor/asmcrypto.js#L10382> (August 2022)
23. MEGA: webclient – asmcrypto.js: mredc. <https://github.com/meganz/webclient/blob/v4.21.4/js/vendor/asmcrypto.js#L9706> (August 2022)
24. MEGA: webclient – asmcrypto.js: RSA_decrypt. <https://github.com/meganz/webclient/blob/v4.21.4/js/vendor/asmcrypto.js#L10746> (August 2022)
25. MEGA: webclient – crypto.js: api_updfkeysync. <https://github.com/meganz/webclient/blob/v4.21.4/js/crypto.js#L3050> (September 2022)
26. MEGA: webclient – crypto.js: crypto_decodeprivkey. <https://github.com/meganz/webclient/blob/v4.21.4/js/crypto.js#L2047> (August 2022)
27. MEGA: webclient – nodedec.js: crypto_rsadecrypt. <https://github.com/meganz/webclient/blob/v4.21.4/nodedec.js#L550> (August 2022)

28. MEGA: webclient – security.js: decryptRsaKeyAndSessionId. <https://github.com/meganz/webclient/blob/v4.21.4/js/security.js#L1231> (August 2022)
29. Micheli, G.D., Heninger, N.: Recovering cryptographic keys from partial information, by example. Cryptology ePrint Archive, Report 2020/1506 (2020), <https://eprint.iacr.org/2020/1506>
30. Ortmann, M.: MEGA security update. <https://web.archive.org/web/20230306125626/https://blog.mega.io/e2ee-security-update/> (Mar 2023)
31. Ryan, K., Heninger, N.: Cryptanalyzing MEGA in six queries. Cryptology ePrint Archive, Report 2022/914 (2022), <https://eprint.iacr.org/2022/914>
32. Shakevsky, A., Ronen, E., Wool, A.: Trust dies in darkness: Shedding light on samsung’s TrustZone keymaster design. Cryptology ePrint Archive, Report 2022/208 (2022), <https://eprint.iacr.org/2022/208>
33. Stein, W., et al.: Sage Mathematics Software Version 9.5. The Sage Development Team (2022), <http://www.sagemath.org>
34. The FPLLL development team: FPLLL, a lattice reduction library (2021), <https://github.com/fplll/fplll>
35. The mitmproxy development team: mitmproxy – an interactive HTTPS proxy (2022), <https://mitmproxy.org/>
36. van Oorschot, P.C., Wiener, M.J.: On Diffie-Hellman key agreement with short exponents. In: Maurer [13], pp. 332–343. https://doi.org/10.1007/3-540-68339-9_29
37. Vanhoef, M., Ronen, E.: Dragonblood: Analyzing the dragonfly handshake of WPA3 and EAP-pwd. In: 2020 IEEE Symposium on Security and Privacy. pp. 517–533. IEEE Computer Society Press (May 2020). <https://doi.org/10.1109/SP40000.2020.00031>
38. Wikipedia: UTF-8. <https://en.wikipedia.org/wiki/UTF-8> (2022)