

Efficient Padding Oracle Attacks on Cryptographic Hardware

Romain Bardou^{1,*}, Riccardo Focardi^{2,**}, Yusuke Kawamoto^{3,*},
Lorenzo Simonato^{2,***}, Graham Steel^{1,*}, and Joe-Kai Tsay^{4,*}

¹ INRIA Project ProSecCo, Paris, France

² DAIS, Università Ca' Foscari, Venezia, Italy

³ School of Computer Science, University of Birmingham, UK

⁴ Department of Telematics, NTNU, Norway

Abstract. We show how to exploit the encrypted key import functions of a variety of different cryptographic devices to reveal the imported key. The attacks are padding oracle attacks, where error messages resulting from incorrectly padded plaintexts are used as a side channel. In the asymmetric encryption case, we modify and improve Bleichenbacher's attack on RSA PKCS#1v1.5 padding, giving new cryptanalysis that allows us to carry out the 'million message attack' in a mean of 49 000 and median of 14 500 oracle calls in the case of cracking an unknown valid ciphertext under a 1024 bit key (the original algorithm takes a mean of 215 000 and a median of 163 000 in the same case). We show how implementation details of certain devices admit an attack that requires only 9 400 operations on average (3 800 median). For the symmetric case, we adapt Vaudenay's CBC attack, which is already highly efficient. We demonstrate the vulnerabilities on a number of commercially available cryptographic devices, including security tokens, smartcards and the Estonian electronic ID card. The attacks are efficient enough to be practical: we give timing details for all the devices found to be vulnerable, showing how our optimisations make a qualitative difference to the practicality of the attack. We give mathematical analysis of the effectiveness of the attacks, extensive empirical results, and a discussion of countermeasures.

1 Introduction

Tamper-resistant cryptographic security devices such as smartcards, USB keys, and Hardware Security Modules (HSMs) are an increasingly common component of distributed systems deployed in insecure environments. Such a device must offer an API to the outside world that allows the keys stored on the device to be used for cryptographic functions and permits key management operations,

* Work partially carried out while at LSV, CNRS & INRIA & ENS-Cachan, France.

** Work partially supported by the RAS Project *TESLA: Techniques for Enforcing Security in Languages and Applications*.

*** Now at Google Inc.

but without compromising security. The most commonly used standard for designing cryptographic device interfaces, RSA PKCS#11 [25], is known to have vulnerabilities if the attacker is assumed to have access to the full API, and can therefore make attacks by combining commands in unexpected ways [5, 6, 8]. In this paper, we describe a different way to attack keys stored on the device using only decryption queries performed by a single function, usually the `C_UnwrapKey` function for encrypted key import. These attacks are cryptanalytic rather than purely logical, and hence require multiple command calls to the interface, but the attacker only needs access to one seemingly innocuous command, subverting the typical countermeasure of introducing access control policies permitting only limited access to the interface.

We will show how the `C_UnwrapKey` command from the PKCS#11 API is often implemented on commercially available devices in such a way that it offers a ‘padding oracle’, i.e. a side channel allowing him to see whether a decryption has succeeded or not. We give two varieties of the attack: the first for when the imported key is encrypted under a public key using RSA PKCS#1 v1.5 padding, which is still by far the most common and often the only available mechanism on the devices we obtained, and the second for when the key is encrypted under a symmetric key using CBC and PKCS#5 padding. The first attack is based on Bleichenbacher’s well-known attack [3]. Although commonly known as the ‘million message attack’, in practice Bleichenbacher’s attack requires only about 215 000 oracle calls on average against a 1024 bit modulus when the ciphertext under attack is known to be a valid PKCS#1 v1.5 block. This is however not efficient enough to be practical on low power devices such as smartcards which perform RSA operations rather slowly. We give a modified algorithm which results in an attack which is 4 times faster on average than the original, with a median attack time over 10 times faster. We also show how the implementation details of some devices can be exploited to create stronger oracles, where our algorithm requires only 9400 mean (3800 median) calls to the oracle. At the heart of our techniques is a small but significant theorem that allows not just multiplication (as in the original attack) but also division to be used to manipulate a PKCS#1 v1.5 ciphertext and learn about the plaintext. In the second attack we use Vaudenay’s technique [27] which is already highly efficient. Countermeasures to such chosen ciphertext attacks are well known: one should use an encryption scheme proven to be secure against them. We discuss the availability of such modes in current cryptographic hardware and examine what other countermeasures could be used while such modes are still not available.

In summary, our contributions are the following: i) new results on PKCS#1 v1.5 cryptanalysis that, when combined with the ‘parallel threads’ technique of Klima-Pokorny-Rosa [26] (which on its own contributes a 38% improvement on mean and 52% on median) results in an improved version of Bleichenbacher’s algorithm giving a fourfold (respectively tenfold) improvement in mean (respectively median) attack time compared to the original algorithm (measured over 1000 runs with randomly generated 1024 bit RSA keys and randomly generated conforming plaintexts); ii) demonstration of the attacks on a variety of

cryptographic hardware including USB security tokens, smartcards and the Estonian electronic ID card, where we found various implementations of the oracle, and adapted our algorithm to each one, resulting in attacks with as few as 9400 mean (3800 median) oracle calls on the most vulnerable devices; iii) analysis of the complexity of the attacks, empirical data, and countermeasures. In the full version available online [1], we discuss manufacturer reaction.

In the next section, we describe the padding attacks relevant to this work and describe our modifications to Bleichenbacher's algorithm. The results on commercial devices are described in section 3. We discuss countermeasures in section 4. Finally we conclude with a discussion of future work in section 5.

2 Padding Oracle Attacks

A padding oracle attack is a particular type of side channel attack where the attacker is assumed to have access to an oracle which returns true just when a chosen ciphertext corresponds to a correctly padded plaintext under a given scheme.

2.1 Bleichenbacher's Attack

Bleichenbacher's padding oracle attack, published in 1998, applies to RSA encryption with PKCS#1 v1.5 padding [3]. Let n, e be an RSA public key and d be the corresponding private key, i.e. $n = pq$ and $ed \equiv 1 \pmod{\phi(n)}$. Let k be the byte length of n , so $2^{8(k-1)} \leq n < 2^{8k}$. Suppose we want to encrypt a plaintext block P where P is l bytes long. Under PKCS#1 v1.5 we first generate a pseudorandom non-zero padding string PS which is $k - 3 - l$ bytes long. We allow l to be at most $k - 11$, so there will be at least 8 bytes of padding. The block for encryption is now created as

$$0x00, 0x02, PS, 0x00, P$$

We call a correctly padded plaintext and a ciphertext that encrypts a correctly padded plaintext *PKCS conforming* or just *conforming*. For the attack, imagine, as above, that the attacker has access to an oracle that tells him just when an encrypted block decrypts to give a conforming plaintext, and assume he is trying to obtain the message $m = c^d \pmod n$, where c is an arbitrary integer. He is going to choose integers s , calculate $c' = c \cdot s^e \pmod n$ and then send c' to the padding oracle. If c' is conforming then he learns that the first two bytes of $m \cdot s$ are $0x00, 0x02$. Hence, if we let $B = 2^{8(k-2)}$, $2B \leq m \cdot s \pmod n < 3B$. The idea is to repeat the process for many values of s until only a single plaintext is possible.

2.2 Improving the Bleichenbacher Attack

Let us first review in a little more detail the original attack algorithm. We are trying to obtain message $m = c^d \pmod n$ from ciphertext c . In step 1 (Blinding),

we search for a random integer value s_0 such that $c(s_0)^e \bmod n$ is conforming, by accessing the padding oracle. We let $c_0 = c(s_0)^e \bmod n$ and $m_0 = (c_0)^d \bmod n$. Note that $m_0 = ms_0 \bmod n$. Thus, if we recover m_0 we can compute the target m as $m_0(s_0)^{-1} \bmod n$. If the target ciphertext is already conforming, we can set s_0 to 1 and skip this step.

We let $B = 2^{8(k-2)}$. If c_0 is conforming, $2B \leq m_0 < 3B$. Thus, we set the initial set M_0 of possible intervals for the plaintext as $\{[2B, 3B - 1]\}$. In step 2, we search for s_i such that $c(s_i)^e \bmod n$ is conforming. In step 3, we apply the s_i we found to narrow the set of possible intervals M_i containing the value of the plaintext, and in step 4 we either compute the solution or jump back to step 2.

We are interested in improving step 2, i.e. the search for s_i . We give step 2 of the original algorithm below, and omit the other steps.

Step 2a. If $i = 1$ (i.e. we are searching for s_1), search for the smallest positive integer $s_1 \geq n/(3B)$ such that $c_0(s_1)^e \bmod n$ is conforming. It can be shown that smaller values of s_1 never give a conforming ciphertext.

Step 2b. If $i > 1$ and $|M_{i-1}| > 1$, search for the smallest positive integer $s_i > s_{i-1}$ such that $c_0(s_i)^e \bmod n$ is conforming.

Step 2c. If $i > 1$ and $|M_{i-1}| = 1$, i.e. $M_{i-1} = \{[a, b]\}$, choose small r_i, s_i such that

$$r_i \geq 2 \frac{bs_{i-1} - 2B}{n} \quad \text{and} \quad \frac{2B + r_i n}{b} \leq s_i < \frac{3B + r_i n}{a}$$

until $c_0(s_i)^e \bmod n$ is conforming. Intuitively, the bounds for s_i derive from the fact that we want $c_0(s_i)^e \bmod n$ conforming, i.e. $2B \leq m_0 s_i - r_i n < 3B$, for some r_i , and from the assumption $a \leq m_0 \leq b$. As explained in the original paper, the constraint on r_i aims at dividing the remaining interval in half so to maximize search performance.

Some features of the algorithm’s behaviour were already known from the original paper. For example, step 2a/b will in general be executed only very few times (in roughly 90% of our trials, step 2b was executed a maximum of once, and in 32% of cases not at all). However, a lot of the expected calls are here, since each time we just search naïvely for the next s_i , which takes an expected $1/Pr(P)$ calls where $Pr(P)$ is the probability of a random ciphertext decrypting to give a conforming block. Step 2c, meanwhile, is highly efficient, but is only applicable if there is only one interval left. Furthermore it cannot be directly applied to the original interval $\{[2B, 3B - 1]\}$ (since the bound on r_i, s_i collapses and we end up with the same search as in step 2a). Based on this observation, we devised a new method for narrowing down the initial interval so that ‘step 2c-like’ reasoning could be applied to speed up the search for s_1 .

Trimming M_0 . First observe that as well as multiplying the value of the decrypted plaintext ($\bmod n$) by some integer s , we can also divide it by an integer t by multiplying the original ciphertext by $t^{-e} \bmod n$. Multiplication modulo n is a group operation on $(\mathbb{Z}_n)^*$, so inverses are unique. If the original plaintext

was divisible by t , the result $m_0 \cdot t^{-1} \bmod n$ will just be m_0/t , otherwise it will be some other value in the group that we in general cannot predict without knowing m_0 . The following holds.

Proposition 1. *Let u and t be two coprime positive integers such that $u < \frac{3}{2}t$ and $t < \frac{2n}{9B}$. If m_0 and $m_0 \cdot ut^{-1} \bmod n$ are PKCS conforming, then m_0 is divisible by t .*

Proof. We have $m_0u < m_0\frac{3}{2}t < 3B\frac{3}{2}t < n$. Thus, $m_0u \bmod n = m_0u$. Let $x = m_0 \cdot ut^{-1} \bmod n$. We know $x < 3B$ since it is conforming. Thus $xt < 3Bt < n$ and $xt \bmod n = xt$. Now, $xt = xt \bmod n = m_0u \bmod n = m_0u$ which implies t divides m_0 .

By Proposition 1, if we find coprime positive integers u and t , $u < \frac{3}{2}t$ and $t < \frac{2n}{9B}$ such that for a PKCS conforming m_0 , $m_0 \cdot ut^{-1} \bmod n$ is also conforming, then we know that m_0 is divisible by t and $m_0 \cdot ut^{-1} \bmod n = m_0\frac{u}{t}$. As a consequence

$$2B \cdot t/u \leq m_0 < 3B \cdot t/u.$$

Note that since we already know $2B \leq m_0 < 3B$ we can restrict our search to t and u such that $2/3 < u/t < 3/2$. We apply this by constructing a list of suitable fractions u/t that we call ‘trimmers’. In practice, we use a few thousand trimmers and take $t \leq 2^{12}$ as the implementations typically satisfy $n \geq 2^{8k-1}$. For each trimmer u/t , we submit $c_0u^et^{-e}$ to the padding oracle. If the oracle succeeds, we can trim the bounds of M_0 .

A large denominator t allows for a more efficient trimming. The trimming process can be thus optimised by taking successful trimming fractions $u_1/t_1, \dots, u_n/t_n$, computing the lowest common multiple t' of t_1, \dots, t_n , using this value as a denominator and then searching for the highest and lowest numerators u_h, u_l that imply a valid padding, giving $2B \cdot t'/u_l \leq m < 3B \cdot t'/u_h$.

Skipping Holes. In the original algorithm step 2a, the search for the first s_1 starts at the value $\lceil n/3B \rceil$. However, note that to be conforming we require in fact that $m \cdot s \geq n + 2B$. Since $3B - 1 \geq m$ we get $(3B - 1)s \geq n + 2B$. So we can start with $s = \lceil (n + 2B)/(3B - 1) \rceil$. On its own this does not save us much: about 8000 queries depending on the exact value of the modulus. However, when we have already applied the trimming rule above to reduce the upper bound on M_0 to some b , this translates immediately into a better start bound for s_1 of $(n + 2B)/b$.

Observe that in general for a successful s we must have $2B \leq ms - jn < 3B$ for some natural number j . Given that we have trimmed the first interval M_0 to the range $[a, b]$, this gives us a series of bounds

$$\frac{2B + jn}{b} \leq s < \frac{3B + jn}{a}$$

Observe further that when

$$\frac{3B + jn}{a} < \frac{2B + (j + 1)n}{b}$$

we have a ‘hole’ of values where a suitable s cannot possibly be. When used in combination with the trimming rule, we found that we frequently obtain a list of many such holes. We use this list to skip out the holes during the search for the s_1 . Note that this is similar to the reasoning used to calculate s values in step 2c, except that here we are concerned with finding the smallest possible s_1 in order to have the fewest possible intervals remaining when searching for s_2 . As we show in the results below, the combination of the trimming and hole skipping techniques is highly effective, in particular against more permissive oracles than a strict PKCS padding oracle.

2.3 Existing Optimisations

In addition to our original modifications, we also implemented changes proposed by Klima, Pokorny and Rosa (KPR) [26]. These are mainly aimed at improving performance in step 2b, because they were concerned with attacking a weaker oracle where most time was spent in step 2b (see below). They are therefore naturally complementary to our optimisation of step 2a.

Parallel Thread Method. The parallel thread method consists of omitting step 2b in the case where there are several intervals in M_{i-1} , and instead forking a separate thread for each interval and using the method of step 2c to search for s_i . As soon as one thread finds a hit, all threads are halted and the new intervals are calculated. If there is still more than one interval remaining, new threads are launched. In practice, since access to the oracle may not be parallelisable, the actions of each thread can be executed stepwise. This heuristic is quite powerful in practice, as we will see below.

Tighter Bounds and Beta Method. KPR were concerned with attacking the weaker ‘bad version’ oracle found in implementations of SSL patched against the original vulnerability. This meant that when the oracle succeeds, they could be sure of the length of the unpadded plaintext, since it must be the right length for the SSL ‘pre-master secret’. This allowed them to tighten the $2B$ and $3B - 1$ bounds. We also implemented this optimisation where possible, since it has no significant cost, but its effects are not significant. We implemented a further proposal of KPR, the so-called ‘Beta Method’ that we do not have space to describe here, but again found that it caused little improvement in practice.

2.4 Stronger and Weaker Oracles

In order to capture behaviour found in real devices (see section 3), we define stronger and weaker Bleichenbacher oracles, i.e. oracles which return true for a greater or smaller proportion of values x such that $2B \leq x < 3B$. We characterise them by three Booleans specifying the tests they apply or skip on the decrypted plaintext. The first Boolean corresponds to the test for a 0 somewhere after the first ten bytes. The second Boolean corresponds to the check for 0s in

the non-zero padding. The third Boolean corresponds to a check of the plaintext length against some specific value (e.g. 16 bytes for an encrypted AES-128 key). More precisely, we say an oracle is FFF if it returns true only on correctly padded plaintexts of a specific fixed length, like the the KPR ‘bad version’ oracle found in some old versions of SSL. An oracle is FFT if it returns true on a correctly padded plaintext of any length. This is the standard PKCS oracle used by Bleichenbacher. An oracle is FTT if it returns true on a correctly padded plaintext of any length and additionally on an otherwise correctly padded plaintext containing a zero in the eight byte padding. An oracle is TFT if it returns true on a correctly padded plaintext of any length and on plaintexts containing no 0s after the first byte. The most permissive oracle, TTT, returns true on any plaintext starting with 0x00, 0x02. We will see in the next section how all these oracles arise in practice.

In Table 1, we show performance of the standard Bleichenbacher algorithm on these oracles, apart from FFF for which it is far too slow to obtain meaningful statistics. Attacking the strongest oracles TTT and TFT is substantially easier than the standard oracle. We can explain this by observing that for the original oracle, on a 1024 bit block, the probability $Pr(P)$ of a random ciphertext decrypting to give a conforming block is equal to the probability that the first two blocks are 0x00, 0x02, the next 8 bytes are non-zero, and there is a zero somewhere after that. We let $Pr(A)$ be the probability that the first two bytes are 0x00, 0x02, i.e $Pr(A) \approx 2^{-16}$. We identify $Pr(P|A)$, the probability of a ciphertext giving a valid plaintext provided the first two bytes are 0x00, 0x02, i.e

$$\left(\frac{255}{256}\right)^8 \cdot \left(1 - \left(\frac{255}{256}\right)^{118}\right) \approx 0.358$$

$Pr(P)$ is therefore $0.358 \cdot 2^{-16}$. Bleichenbacher estimates that, if no blinding phase is required, the attack on a 128 byte plaintext will take

$$2/Pr(P) + 16 \cdot 128/Pr(P|A)$$

oracle calls. So we have

$$(2 \cdot 2^{16} + 16 \cdot 128)/Pr(P|A) = 371843$$

In the case of, say, the TTT oracle, $Pr(P|A)$ is 1, since any block starting 0x00, 0x02 will be accepted. Hence we have

$$2^{17} + 16 \cdot 128 = 133120$$

oracle queries. This is higher than what we were able to achieve in practice in both cases, but the discrepancy is not surprising since the analysis Bleichenbacher uses is a heuristic approximation of the upper bound rather than the mean. However, it gives an explanation of why the powerful oracle gives such a big improvement in run times: improvements in the oracle to $Pr(P|A)$ make a

multiplicative difference to the run time. Additionally, the expected number of intervals at the end of step 2a is $\lceil s_1 \cdot B/n \rceil$ [3, p. 7], so if s_1 is less than 2^{16} , the expected number of intervals is one. For the FFT oracle, the expected value of s_1 (calculated as $1/2 \cdot 1/Pr(P)$) is about 91 500, between 2^{16} and 2^{17} , whereas for TTT it is 2^{15} . That means that in the TTT case we can often jump step 2b and go straight to step 2c, giving a total of

$$2^{16} + 16 \cdot 128 = 34816$$

i.e. the TTT oracle is about 10 times more powerful than the FFT oracle, which is fairly close to what we see in practice (our mean for FFT is about 5.5 times that for TTT).

In comparison, if the modulus is 2048 bit long, then $Pr(P|A) \approx 0.599$. Because the modulus is longer, the probability that 0x00 appears after the 8 non-zero bytes is higher than in the 1024 bit case. Furthermore, following the same argument as above, we obtain that the attack on a 2048 bit plaintext will take about 335 065 calls to the FFT oracle, fewer than in the 1024 bit case. Note however that RSA private key operations slow down by roughly a factor of four when key length is doubled.

Table 1. Performance of the original and modified algorithms

Oracle	Original algorithm		Modified algorithm			
	Mean	Median	Mean	Median	Trimmers	Mean skipped
FFF	-	-	18 040 221	12 525 835	50 000	7 321
FFT	215 982	163 183	49 001	14 501	1 500	65 944
FTT	159 334	111 984	39 649	11 276	2 000	61 552
TFT	39 536	24 926	10 295	4 014	600	20 192
TTT	38 625	22 641	9 374	3 768	500	18 467

2.5 Performance of the Modified Algorithm

Referring again to Table 1, we give a summary of our experiments with our modified algorithm. As well as mean and median, we give the number of trimming fractions tried and the average number of oracle calls saved by the hole skipping modification we presented in section 2.2. Observe that as the oracles become stronger, the contribution of the KPR ‘parallel threads’ method becomes less significant and our hole skipping technique more significant. This is to be expected, since as discussed above, for the stronger oracles, fewer runs need to use step 2b. Similarly, when trimming the first interval M_0 , we find that more fractions can be used because of the more permissive oracle, hence we find more holes to skip. For the most restrictive oracle, FFF, the addition of our trimming method slightly improves on the results of KPR (which were 20 835 297 mean and 13 331 256 median). Note also that the trimming technique contributes more

than just the oracle calls saved by the hole skipping, it also slightly improves performance on all subsequent stages of the algorithm. We know this because we can compare performance using only the parallel threads optimisation, where we obtain a mean of 113 667 and a median of 78 674 (on the FFT oracle). In Figure 1, we give the density distribution for 1000 runs of the original algorithm and our optimised algorithm on the classical FFT oracle, with medians marked. Notice the change in shape: we have a much thinner tail.




Fig. 1. Graph comparing distribution of oracle calls for original (lower peak, thinner line) and optimised version of the algorithm on the FFT oracle. Median is marked for each.

2.6 Vaudenay's Attack

Vaudenay's attack on CBC mode symmetric-key encryption [27] is somewhat simpler and highly efficient. Recall first the operation of CBC mode [9]: given some block cipher with encryption, decryption functions $E(\cdot), D(\cdot)$ and a fixed block size of b bytes, suppose we want to encrypt a message P of length $l = j \cdot b$ for some integer j , i.e. $P = P_1, \dots, P_j$. In CBC mode, we first choose a fresh *initialisation vector* IV . The first encrypted block is defined as $C_1 = E(IV \oplus P_1)$, and subsequent blocks as $C_i = E(C_{i-1} \oplus P_i)$. The need for padding arises because l is not always a multiple of b . Suppose $l = j \cdot b + r$. Then we need to encrypt the last r bytes of the message in a b bytes block in such a way that on decryption, we can recognise that only the first r bytes are to be considered part of the plaintext. One way to do this is the so-called RC5 padding, also known as PKCS padding and described in RFC 5652 [12]. The r bytes are encoded into the leftmost bytes of the final block, and then the final $b - r$ bytes are filled with the value $b - r$. Under this padding scheme, if the plaintext length should happen to be an exact multiple of the block size, then we add a whole block of padding bytes b .

To effect Vaudenay's attack, suppose that the attacker has some ciphertext C_1, \dots, C_n and access to an oracle that returns true just when a ciphertext decrypts with valid padding. To attack a given block C_i , we first prepend a random block $R = r_1, \dots, r_b$. We then ask the padding oracle to decrypt $R \parallel C_i$.

If the padding is valid most probably the final byte is 1, hence the final byte p_m of the plaintext P_i satisfies $p_b = r_b \oplus 1$. If the padding is not accepted, we iterate over i setting $r'_b = r_b \oplus i$ and retrying the oracle until eventually it is accepted. There is a small chance that the final byte of an accepted block is not 1, but this is easily detected. Having discovered the last byte, it is easy to extend the attack to obtain p_{b-1} by tweaking r_{b-1} , and so on for the whole block. Given this ‘block decryption oracle’ we can then apply it to all the blocks of the message. Overall, the attack requires $O(nb)$ steps, and hence is highly efficient.

Since the original attack appeared, many variations have been found on other padding schemes and block cipher modes [2, 7, 14, 17, 20, 22]. Bond and French recently showed that the attack could be applied to the `C.UnwrapKey` command as implemented on a hardware security module (HSM) [4]. We will show in the next section that many cryptographic devices are indeed vulnerable to variants of the attack.

Table 2. Attack Results on Tokens

Device	PKCS#11 version	PKCS#1 v1.5 Attack		CBC-PAD Attack	
		Token	Session	Token	Session
Aladdin eTokenPro	2.01	✓	✓	✓	✓
Feitian ePass 2000	2.11	×	×	N/A	N/A
Feitian ePass 3003	2.20	×	×	N/A	N/A
Gemalto Cyberflex	2.01	✓	N/A	N/A	N/A
RSA Securid 800	2.20	✓	N/A	N/A	N/A
Safenet Ikey 2032	2.01	✓	✓	N/A	N/A
SATA DKey	2.11	×	×	×	×
Siemens CardOS	2.11	✓	✓	N/A	N/A

3 Attacking Real Devices

We applied the optimised versions of the attacks of Bleichenbacher and Vaudenay presented in section 2 to the unwrap functionality of PKCS#11 devices. RSA PKCS#11, which describes the ‘Cryptoki’ API for cryptographic hardware, was first published in 1995 (v1.0). The latest official version is v2.20 (2004) which runs to just under 400 pages [25]. Adoption of the standard is almost ubiquitous in commercial cryptographic tokens and smartcards, even if other additional interfaces are frequently offered. In a PKCS#11-based API, applications initiate a *session* with the cryptographic token, by supplying a PIN. Once a session is initiated, the application may access the *objects* stored on the token, such as keys and certificates. Objects are referenced in the API via *handles*, which can be thought of as pointers to or names for the objects. In general, the value of the handle, e.g. for a secret key, does not reveal any information about the actual value of the key. Objects have *attributes*, which may be bitstrings e.g. the value

of a key, or Boolean flags signalling properties of the object, e.g. whether the key may be used for encryption (`CKA_ENCRYPT`¹), or for encrypting other keys, for signing, verification, and other uses. New objects can be created by calling a key generation command, or by *unwrapping* an encrypted key packet using the `C.UnwrapKey` command, which takes a handle, a ciphertext and a *template* as input. A template is a partial description of the key to be imported, giving notably its length. The device attempts to decrypt the ciphertext using the key referred to by the handle. If it succeeds, it creates a new key on the device using the extracted plaintext and the template, and returns a new handle.

Observe that a padding check immediately following the decryption could give rise to an oracle that may be used to determine the value of the newly stored key. To test for such an oracle on a device, we create a key with the `CKA_UNWRAP` attribute set to allow the `C.UnwrapKey` operation, create encrypted key packets with deliberately placed padding errors, call the function on these ciphertexts and observe the return codes. For the case of asymmetric key unwrapping, constructing test ciphertexts is easy since the public key of the pair is always obtainable via a query to the PKCS#11 interface. For symmetric key unwrapping, it is not quite so trivial since the device may create unwrapping keys marked with the Boolean key attribute `CKA_SENSITIVE` which prevents them from being read via the PKCS#11 interface. In this case there are various tricks we can use: we can try to set the attribute `CKA_ENCRYPT` and then use the PKCS#11 function `C.Encrypt` to construct the test packets if a suitable mode is available, or if the device does not allow this, we can explicitly try to create a key with `CKA_SENSITIVE` set to false, assuming the same unwrap algorithm will be used as for sensitive keys. In the event, we were always able to find some way to do this with the devices under test.

3.1 Smartcards and Security Tokens

In Table 2 we give results from implementing the attacks on all the commercially available smartcards and USB tokens we were able to obtain that offer a PKCS#11 interface and support the unwrap operation. A tick means not only that we were able to construct a padding oracle, but that we were actually able to execute the attack and extract the correct encrypted key. A cross notes that the attack fails. We explain these failures below. Not applicable (N/A) means that the token did not support the cryptographic mechanisms and/or unwrap modes required for this attack. Note that relatively few devices support unwrap under symmetric key algorithms. We tested the attacks using both token keys and session keys for the unwrapping. The exact semantics of the difference between these key types is not completely clear from the standard: there is an attribute `CKA_TOKEN` which when set to true indicates a token key and when false

¹ Throughout the paper we will refer to commands, attributes, return codes and mechanisms by their names as defined in the PKCS#11 standard, so `C_` prefixes a (cryptoki) command, `CKA_` prefixes a cryptoki attribute, `CKR_` prefixes a cryptoki return code and `CKM_` prefixes a cryptoki mechanism.

indicates a session key. Session keys are destroyed when the session is ended, whereas token keys persist. However, we have noticed that devices often enforce very different policies for token keys and session keys, so it seemed pertinent to test both types.

In Table 3 we give the class of padding oracle found in each device in the PKCS#1 v1.5 case. To obtain this table we construct padded plaintexts with a single padding error and observed the return code from the token. Note that we give separate entries for token and session keys in this table only when there is a difference in the device's behaviour in the two cases. We report median attack time, computed from the results of table 1 and from a measure of the unwrap rate of the hardware. Notice how the tenfold improvement in median attack time of our modified algorithm makes attacks even against FFT oracles on slow devices quite practical. Unwrap calls using session keys are often many times faster than token keys though it is not clear why, unless perhaps these devices are carrying out session key operations in the driver software rather than on the card.

We will briefly discuss each line of Table 2 in turn. The **Aladdin eToken Pro** supports both unwrapping modes required, though the `CBC_PAD` unwrap mode does not conform to the standard: a block containing a final byte of `0x00` is accepted. According to the standard, if the final byte of the plaintext is zero and it falls at the end of a block, then an entire block of padding should be added (see section 2). This causes a small problem for the attack since it gives us an extra possibility for the last byte, but we easily adapted the attack to take account of this. The PKCS#1 v1.5 padding implementation ignores zeros in the first 8 bytes of the padding and gives a separate error when the length of the extracted key does not match the requested one (`CKR_TEMPLATE_INCONSISTENT`). Based on this we can build an FTT oracle. The **Feitian** tokens do not support `CBC_PAD` modes. They also do not implement PKCS#1 v1.5 padding correctly: in our tests, any block with `0x02` in the second byte was accepted, except for very large values (e.g. for one key, anything between `0x00` and `0xE2` in the first byte was accepted). The result is that the attack does not succeed. The **Gemalto Cyberflex** smartcard does not allow unwrapping under symmetric keys. However, it seems to implement standard PKCS#1 v1.5 padding correctly, and the Bleichenbacher attack succeeds (FFT oracle, since the length is ignored). The **RSA SecurID** device does not support unwrapping using symmetric keys, hence the Vaudenay attack is not possible. However, the Bleichenbacher attack works perfectly. In fact, the RSA token implements a perfect TTT oracle. The device also supports OAEP, but not in a way that prevents the attack (see next paragraph). The **Safenet ikey2032** implements an asymmetric key unwrapping. The padding oracle derived is more accepting than the Bleichenbacher oracle since the 0s in the first 8 bytes of the padding string are ignored (FTT oracle). The **SATA DKey** does not implement standard padding checks. In `CBC_PAD` mode, only the last byte is checked: it seems that as long as the last byte n is less than the number of bytes in a block, the padding is accepted and the final n bytes discarded. This means we cannot use the attack to recover the whole

key, just the final byte. In PKCS#1 v1.5 mode, many incorrectly padded blocks were accepted, and we were unable to deduce the rationale. For example, any block with the *first* byte equal to 0x02 is accepted. The wide range of accepted blocks prevents the attack. The **Siemens CardOS** supports only unwrapping under asymmetric keys. The Bleichenbacher attack works perfectly: with token keys the oracle is TTT, while with session keys it is FFT.

Table 3. Oracle Details and Median Attack Times

Device	Token		Session	
	Oracle	Time	Oracle	Time
Aladdin eTokenPro	FFT	21m	FFT	17m
Gemalto Cyberflex	FFT	92m	N/A	N/A
RSA Securid 800	TTT	13m	N/A	N/A
Safenet Ikey 2032	FFT	88m	FFT	17m
Siemens CardOS	TTT	21m	FFT	89s

Attacking OAEP Mode Unwrapping. A solution to the Bleichenbacher attack is to use OAEP mode encryption, which was first added to PKCS#1 in v2.0 (1998) and is recommended for all new applications since v2.1 (2002). RSA OAEP was included as a mechanism in PKCS#11 in version 2.10 (1999). However, out of the tokens tested (all of which are currently available products), only one, the RSA SecureID, supports OAEP encryption. The standard PKCS#1 v2.1 notes that it is dangerous to allow two mechanisms to be enabled on the same key [24, p. 14], since “an opponent might be able to exploit a weakness in the implementation of RSAES-PKCS1-v1.5 to recover messages encrypted with either scheme.”. An examination of the developer’s manual for the RSA SecurID reveals that for private keys generated by the token, the relevant attribute “CKA_ALLOWED_MECHANISMS is always set to the following mechanism list : CKM_RSA_PKCS, CKM_RSA_PKCS_OAEP, and CKM_RSA_X_509.”. We created a key wrapped under OAEP and then performed Bleichenbacher’s attack on it using a PKCS#1 v1.5 unwrap oracle. The attack is only slightly complicated by the fact that the initial encrypted block does not yield a valid block when decrypted, requiring us to use the ‘blinding phase’ where many ciphertexts are derived from the original to obtain one that passes the padding oracle. In our tests this added only a few hundred seconds to the attack.

3.2 HSMs

Hardware Security Modules are widely used in banking and similar sectors where a large amount of cryptographic processing has to be done securely at high speed (verifying PIN numbers, signing transactions, etc.). A typical HSM retails for around 20 000 Euros hence is unfortunately too expensive for our laboratory

budget. HSMs process RSA operations at considerable speed: over 1000 decryptions per second for 1024 bit keys. Even in the case of the FFF oracle, which requires 12 000 000 queries, this would result in a median attack time of 12 000 seconds, or just over three hours.

We hope to be able to give details of HSM testing soon in the online version of the paper [1].

3.3 Estonian ID Card

Estonia's Citizenship and Migration Board completed the issuing of more than 1 million national electronic ID (eID) cards in 2006 [16]. The eID is the primary national identification document in Estonia and it is mandatory for all Estonian citizens and alien residents 15 years and older to have one [10]. The card contains two RSA key pairs [13]. One key pair is intended to be mainly used for authentication (e.g., for mutual authentication with TLS/SSL) but can also be used for encrypting and signing email (e.g., with S/MIME). The other key pair is attributed only to be used for digital signatures. Only this latter key pair can be used for legally binding digital signatures [16]. Since January 1, 2011, the eID cards contain 2048 bit RSA keys, therefore these cards comply with NIST's recommendation [18]. However, cards issued before January 1, 2011 continue to use 1024 bit keys.

Attack Vector. Unlike the cryptographic devices discussed above, the Estonian eID card does not allow the import of keys, so our attack here does not rely on the unwrap operation. Instead we consider attacks using the padding oracle provided by the decryption function of the DigiDoc software, part of the official ID software package developed by the Estonian Certification Center, Estonia's only CA [11]. We note that the attack succeeds with any application that returns whether decryption with the eID card succeeds. Our experiments were conducted using the Java library of DigiDoc, called *JDigiDoc*. DigiDoc encrypts data using a hybrid encryption scheme, where a 128-bit AES key is encrypted under a public key. First we tested the Estonian ID card's decryption function using raw PKCS#11 calls and confirmed that it checks padding correctly. We then observed that with the default configuration, when attempting to decrypt, e.g., an encrypted email, JDigiDoc writes a log file of debug information that includes the padding errors for the 128-bit AES key that is encrypted under the public key. Any application built on JDigiDoc, that reveals whether decryption succeeds, e.g., by leaking the contents of the log file, provides an attacker with a suitable padding oracle. The information in JDigiDoc's log file gives an attacker access to essentially an FFT oracle but with additional length information. The length information allows us to adjust the $2B$ and $3B - 1$ bounds used in the attack, though in our experiments this made little difference.

In tests, the Estonian ID card, using 2048 bit keys, was able to perform 100 decryptions in 340 seconds. This means that for our optimised attack, where 28 300 decryptions are required, we would need about 96 200 seconds, or about 27 hours to decrypt an arbitrary valid ciphertext. For ID cards using 1024 bit

keys, each decryption should be four times faster, while 49 000 decryptions are required; therefore we estimate a time of about 41 700 seconds, or about 11 hours and 30 minutes to decrypt an arbitrary valid ciphertext. To forge a signature, we require, due to the extra blinding step, a mean of 109 000 oracle calls and a median of 69 000 oracle calls to get a valid signature on an arbitrary message, giving an expected time of 103 hours on a 2048 bit Estonian eID. On a card using 1024 bit keys, we require a mean of 203 000 calls and a median of 126 000 calls; therefore expect to sign an arbitrary message in around 48 hours.

4 Countermeasures

A general countermeasure to the Bleichenbacher and Vaudenay attacks has been well known for years: use authenticated encryption. There are no such modes for symmetric key encryption in the current version of PKCS#11, but version 2.30, which is still at the draft stage, includes GCM and CCM (mechanisms `CKM_AES_GCM` and `CKM_AES_CCM`). While these modes have their critics [23], they do in theory provide secure authenticated encryption and hence could form the basis of secure symmetric key unwrap mechanisms. Unfortunately, in the current draft (v7), they are given only as modes for `C_Encrypt`. Adoption of these modes for `C_UnwrapKey` would provide a great opportunity to give the option of specifying authenticated data along with the encrypted key to allow secure transfer of attributes between devices. This would greatly enhance the flexibility of secure configurations of PKCS#11. To prevent the Bleichenbacher attack one must simply switch to OAEP, which is already in the standard. PKCS#11 should follow PKCS#1's long-held position of recommending OAEP exclusively for all new applications. Care must also be taken to remind developers not to allow the two modes to be used on the same key, as is the case in RSA's own SecureID device. In fact, the minutes of the 2003 PKCS workshop suggest that there was a consensus to include the single mechanism recommendation in version 2.20 [21], but it does not appear in the final draft. Note that care must be taken when implementing OAEP as otherwise there may also be a padding oracle attack which is even more efficient than our modified Bleichenbacher attack [15], though we are yet to find such an oracle on a PKCS#11 device.

If unauthenticated unwrap modes need to be maintained for backwards compatibility reasons, there are various options available. For the CBC case, Black and Urtubia note that the 10^* padding, where the plaintext is followed by a single 1 bit and then only 0 bits until the end of the block, leaks no information from failed padding checks while still allowing length of the plaintext to be determined unambiguously [2]. Paterson and Watson suggest a refinement that additionally preserves a notion of indistinguishability, by ensuring that no padded blocks are invalid [19]. They also give appropriate security proofs for the two schemes. If PKCS#1 v1.5 needs to be maintained, we have seen that an implementation of the padding check that rejects anything other than a conforming plaintext containing a key of the correct length with a single error code gives the weakest possible (FFF) oracle. This may be enough for some applications, but

one is well advised to remember the maxim that attacks only get better, never worse. An alternative approach would be to adopt ‘SSL style’ countermeasures, proceeding to import a randomly generated key in the case where a block contains invalid padding. However, this may not fix the hole: if an attacker is able to replay the same block and detect that two different keys have been imported, he knows there is a padding error. One could also decide to ignore padding errors completely and always import just the number of bytes corresponding to the size of the key required, but this looks dangerous: if the same block can be passed off as several different kinds of key, this might open the possibility of attacking weaker algorithms to obtain keys for stronger ones. Thus it seems clear that authenticated encryption is by far the superior solution.

We give manufacturer response in the full version of the paper [1]. There is a broad spectrum: while some manufacturers offer mitigations and state a clear need to get authenticated encryption into the standard and adopted as soon as possible, others see their responsibility as ending as soon as they conform to the PKCS#11 standard, however vulnerable it might be.

5 Conclusions

We have demonstrated a modified version of the Bleichenbacher RSA PKCS#1 v1.5 attack that allows the ‘million message attack’ to be carried out in a few tens of thousands of messages in many cases. We have implemented and tested this and the Vaudenay CBC attack on a variety of contemporary cryptographic hardware, enabling us to determine the value of encrypted keys under import. We have shown that the way the `C_UnwrapKey` command from the PKCS#11 standard is implemented on many devices gives rise to an especially powerful error oracle that further reduces the complexity of the Bleichenbacher attack. In the worst case, we found devices for which our algorithm requires a median of only 3 800 oracle calls to determine the value of the imported key. Vulnerable devices include eID cards, smartcards and USB tokens.

While some theoreticians find the lack of a security proof sufficient grounds for rejecting a scheme, some practitioners find the absence of practical attacks sufficient grounds for continuing to use it. We hope that the new results with our modified algorithm will prompt editors to reconsider the inclusion of PKCS#1 v1.5 in contemporary standards such as PKCS#11.

References

1. Bardou, R., Focardi, R., Kawamoto, Y., Simionato, L., Steel, G., Joe-Kai-Tsay: The million message attack in 15 000 messages, or efficient padding oracle attacks on cryptographic hardware. Cryptology ePrint Archive (to appear, 2012), <http://eprint.iacr.org/>
2. Black, J., Urtubia, H.: Side-channel attacks on symmetric encryption schemes: The case for authenticated encryption. In: Boneh, D. (ed.) USENIX Security Symposium, pp. 327–338. USENIX (2002)

3. Bleichenbacher, D.: Chosen Ciphertext Attacks against Protocols Based on the RSA Encryption Standard PKCS #1. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 1–12. Springer, Heidelberg (1998)
4. Bond, M., French, G.: Hidden semantics: why? how? and what to do? Presentation at Fourth Analysis of Security APIs Workshop, ASA-4 (July 2010)
5. Bortolozzo, M., Centenaro, M., Focardi, R., Steel, G.: Attacking and fixing PKCS#11 security tokens. In: Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS 2010), Chicago, Illinois, USA. ACM Press (October 2010)
6. Clulow, J.: On the Security of PKCS #11. In: Walter, C.D., Koç, Ç.K., Paar, C. (eds.) CHES 2003. LNCS, vol. 2779, pp. 411–425. Springer, Heidelberg (2003)
7. Degabriele, J.P., Paterson, K.G.: On the (in)security of ipsec in mac-then-encrypt configurations. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM Conference on Computer and Communications Security, pp. 493–504. ACM (2010)
8. Delaune, S., Kremer, S., Steel, G.: Formal analysis of PKCS#11. In: Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF 2008), Pittsburgh, PA, USA, pp. 331–344. IEEE Computer Society Press (June 2008)
9. Dworkin, M.: Recommendation for block cipher modes of operation: Modes and techniques. NIST Special Publication 800-38A (December 2001)
10. Estonian Certification Center. The estonian ID card and digital signature concept, principles and solutions (March 2003), http://www.id.ee/public/The_Estonian_ID_Card_and_Digital_Signature_Concept.pdf
11. Estonian Informatics Center. Estonian ID-software, <https://installer.id.ee/?lang=eng>
12. Housley, R.: Cryptographic Message Syntax (CMS). RFC 5652 (Standard) (September 2009)
13. ID Süsteemide AS. EstEID specification v2.01, http://www.id.ee/public/EstEID_Spetsifikatsioon_v2.01.pdf
14. Jager, T., Somorovsky, J.: How to break XML encryption. In: Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS), pp. 413–422 (2011)
15. Manger, J.: A Chosen Ciphertext Attack on RSA Optimal Asymmetric Encryption Padding (OAEP) as Standardized in PKCS #1 v2.0. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 230–238. Springer, Heidelberg (2001)
16. Martens, T.: eID interoperability for PEGS, national profile estonia, European Commission’s IDABC programme (November 2007), <http://ec.europa.eu/idabc/en/document/6485/5938>
17. Mitchell, C.J.: Error Oracle Attacks on CBC Mode: Is There a Future for CBC Mode Encryption? In: Zhou, J., López, J., Deng, R.H., Bao, F. (eds.) ISC 2005. LNCS, vol. 3650, pp. 244–258. Springer, Heidelberg (2005)
18. National Institute of Standards and Technology. NIST special publication 800-57, recommendation for key management (March 2007), <http://csrc.nist.gov/publications/PubsSPs.html>
19. Paterson, K.G., Watson, G.J.: Immunising CBC Mode Against Padding Oracle Attacks: A Formal Security Treatment. In: Ostrovsky, R., De Prisco, R., Visconti, I. (eds.) SCN 2008. LNCS, vol. 5229, pp. 340–357. Springer, Heidelberg (2008)
20. Paterson, K.G., Yau, A.: Padding Oracle Attacks on the ISO CBC Mode Encryption Standard. In: Okamoto, T. (ed.) CT-RSA 2004. LNCS, vol. 2964, pp. 305–323. Springer, Heidelberg (2004)
21. Minutes for the April, 2003 PKCS workshop (2003), <ftp://ftp.rsa.com/pub/pkcs/03workshop/minutes.txt>

22. Rizzo, J., Duong, T.: Practical padding oracle attacks. In: Proceedings of the 4th USENIX Conference on Offensive Technologies, WOOT 2010, pp. 1–8. USENIX Association, Berkeley (2010)
23. Rogaway, P.: Evaluation of some blockcipher modes of operation (February 2011), <http://www.cs.ucdavis.edu/~rogaway>; Evaluation carried out for the Cryptography Research and Evaluation Committees (CRYPTREC) for the Government of Japan
24. RSA Security Inc., v2.1. PKCS #1: RSA Cryptography Standard (June 2002)
25. RSA Security Inc., v2.20. PKCS #11: Cryptographic Token Interface Standard (June 2004)
26. Klíma, V., Pokorný, O., Rosa, T.: Attacking RSA-Based Sessions in SSL/TLS. In: Walter, C.D., Koç, Ç.K., Paar, C. (eds.) CHES 2003. LNCS, vol. 2779, pp. 426–440. Springer, Heidelberg (2003)
27. Vaudenay, S.: Security Flaws Induced by CBC Padding - Applications to SSL, IPSEC, WTLS... In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 534–545. Springer, Heidelberg (2002)