

How to Backdoor Diffie-Hellman

David Wong

NCC Group, June 2016

Abstract

Lately, several backdoors in cryptographic constructions, protocols and implementations have been surfacing in the wild: Dual-EC in RSA's B-Safe product, a modified Dual-EC in Juniper's operating system ScreenOS and a non-prime modulus in the open-source tool socat. Many papers have already discussed the fragility of cryptographic constructions not using nothing-up-my-sleeve numbers, as well as how such numbers can be safely picked. However, the question of how to introduce a backdoor in an already secure, safe and easy to audit implementation has so far rarely been researched (in the public).

We present a new way of building a Nobody-But-Us (NOBUS) Diffie-Hellman backdoor by using a composite modulus with a smooth order. We then explain how we were able to implement a proof of concept with Socat and OpenSSL in order to exploit our backdoor on the TLS protocol.

Keywords: Diffie-Hellman, Ephemeral, DHE, NOBUS, Backdoor, Discrete Logarithm, Small Subgroup Attack, Pohlig-Hellman, Pollard Rho, Factorization, Pollard's p-1, ECM, Dual-EC, Juniper, socat

Update (December 2016): Dorey et al.¹

¹Kristen Dorey, Nicholas Chang-Fong, and Aleksander Essex. *Indiscreet Logs: Persistent Diffie-Hellman Backdoors in TLS*. Cryptology ePrint

have pointed an attack on our first contribution, as well as an improvement for the exploitation of our second contribution. This work has been updated to reflect these advances.

1 Introduction

Around Christmas 2015 *Juniper*, a networking hardware company, released an out-of-cycle security bulletin². Two vulnerabilities were disclosed without much details to help us grasp the seriousness of the situation. Fortunately, at this period of the year many researchers were home with nothing else to do but to solve this puzzle. By quickly comparing both the patched and vulnerable binaries, the two issues were pinpointed. While one of the vulnerabilities was a simple "master"-password implemented at a crucial step of the product's authentication, the other discovery was a bit more subtle: a unique value was modified. More accurately, a number in the source code was replaced. The introduction of the vulnerability was so simple, and due to the fact that the number was stored as a string of hexadecimal digits, the trivial use of the UNIX command line tool `strings` was enough to discover it.

The special value ended up being a constant

Archive, Report 2016/999. <http://eprint.iacr.org/2016/999>. 2016

²<https://kb.juniper.net/InfoCenter/index?page=content&id=JSA10713>

number was at issue. This time it was the public modulus, an integer used to generate the ephemeral Diffie-Hellman keys of both parties during socat’s TLS handshakes. This algorithm had been, contrary to Dual-EC, considered secure from the start. But as it turned out, badly understood as well: as the *Logjam*⁷ paper had demonstrated earlier in the previous year, most servers would use Diffie-Hellman key exchanges to perform *ephemeral handshakes*, and the same servers would generate their ephemeral keys from hardcoded defaults (often the same ones) provided by various TLS libraries. The paper raised a wave of discussion around how developers should use Diffie-Hellman, at the same time scaring people away from 1024 bit DH: “We estimate that even in the 1024-bit case, the computations are plausible given nation-state resources”.

Securely integrating DH in a protocol is unfortunately not well understood. Defensive approaches are discussed in several RFCs^{8,9}, but few papers so far have taken the point of view of the attacker. The combination of the current trend of increasing the bitsize of DH parameters with the now old trend of using open source libraries’ defaults to generate ephemeral Diffie-Hellman keys would give opportunist attackers a valid excuse to submit their bigger (more secure) and backdoored

parameters into open-source or closed-source libraries. This work is about generating such backdoors and implementing them in TLS, showing how easy and subtle the process is. The working code along with explanations on how to reproduce our setup is available on Github¹⁰.

In section 2, we will first briefly talk about the several attacks possible on Diffie-Hellman, from small subgroup attacks to Pohlig Hellman’s algorithm. In section 3 we will introduce our first attempt at a DH backdoor. We will present another DH backdoor in section 4 by using the ideas of the previous section with a composite modulus. In section 5 we will see another method using a composite modulus that allows us to choose a specific generator, allowing us to only modify the modulus value when implementing our backdoor. In section 6 we will explain how we implemented the backdoor in TLS and how we exploited it. We will then see in section 7 how to detect such backdoors and how to prevent them. Eventually we will wrap it all up in section 8.

2 Attacks on Diffie-Hellman and the Discrete Logarithm

To attack a Diffie-Hellman key exchange, one could extract the secret key \mathbf{a} from one of the peer’s public key $y_a = g^a \pmod{p}$. One could then compute the shared key $g^{ab} \pmod{p}$ using the other peer’s public key $y_b = g^b \pmod{p}$.

The naive way to go about this is to compute each power of g (while tracking the exponent) until the public key is found. This is called *trial multiplication* and would need on average $\frac{q}{2}$ operations to find a solution (with

⁷David Adrian et al. “Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice”. In: *22nd ACM Conference on Computer and Communications Security*. <https://weakdh.org/imperfect-forward-secrecy-ccs15.pdf>. Oct. 2015

⁸Eric Rescorla. *RFC 2631: Diffie-Hellman Key Agreement Method*. RFC 2631. <https://rfc-editor.org/rfc/rfc2631.txt>. 2013. DOI: [10.17487/rfc2631](https://doi.org/10.17487/rfc2631)

⁹Robert Zuccherato. *RFC 2785: Methods for Avoiding the Small-Subgroup Attacks on the Diffie-Hellman Key Agreement Method for S/MIME*. RFC 2785. <https://rfc-editor.org/rfc/rfc2785.txt>. 2013. DOI: [10.17487/rfc2785](https://doi.org/10.17487/rfc2785)

¹⁰https://github.com/mimoo/Diffie-Hellman_Backdoor

q the order of the base). More efficiently, algorithms that compute discrete logarithm in expected \sqrt{q} steps like Shank's *baby-step giant-step* (deterministic), *Pollard rho* or *Pollard Kangaroo* (both probabilistic) can be used. Because of the memory required for *baby-step giant-step*, Pollard's algorithms are often preferred. While both are parallelizable, *Pollard Kangaroo* is used when the order is unknown or known to be in a small interval. For larger orders the Index Calculus or Number Field Sieve (NFS) algorithms are the most efficient. But so far, computing a discrete logarithm in polynomial time on a classical computer is still an open problem.

2.1 Pollard Rho

The algorithm that interests us here is *Pollard Rho*: it is fast in relatively small orders, it is parallelizable and it takes very little amount of memory to run. The idea comes from the birthday paradox and the following equation (where x is the secret key we are looking for; and a, a', b and b' are known):

$$g^{xa+b} = g^{xa'+b'} \pmod{p}$$

$$\implies x = (a - a')^{-1}(b' - b) \pmod{p - 1}$$

The birthday paradox tells us that by looking for a random collision we can quickly find one in $\mathcal{O}(\sqrt{p})$. A random function is used to efficiently step through various g^{xa+b} until two values repeat themselves, it is then straightforward to calculate x . Cycle-finding algorithms are used to avoid storing every iteration of the algorithm (two different iterations of g^{xa+b} are started and end up in a loop past a certain step) and the technique of distinguished points is used to parallelize the algorithm. (Machines only save and share particular iterations, for example iterations starting with a chosen number of zeros.)

2.2 Pohlig-Hellman

In 1978, Pohlig and Hellman discovered a shortcut to the discrete logarithm problem¹¹: if you know the complete factorization of the order of the group, and all of the factors are relatively small, then the discrete logarithm can be quickly computed.

The idea is to find the value of the secret key x modulo the divisors of the group's order by reducing the public key $y = g^x \pmod{p}$ in subgroups of order dividing the group order. Thanks to the Chinese Remainder Theorem (CRT) stated later on, the secret key can then be reassembled in the group order. Summed up below is the full Pohlig-Hellman algorithm (with φ being Euler's totient function):

1. Determine the prime factorization of the order of the group

$$\varphi(p) = \prod p_i^{k_i}$$

2. Determine the value of x modulo $p_i^{k_i}$ for each i
3. Recompute $x \pmod{\varphi(p)}$ with the CRT

The central idea of Pohlig and Hellman's algorithm is in how they determine the value of the secret key x modulo each factor $p_i^{k_i}$ of the order. One way of doing it is to try to reduce the public key to the subgroup we're looking at by computing:

$$y^{\varphi(p)/p_i^{k_i}} \pmod{p}$$

Computing the discrete logarithm of that value, we get $x \pmod{p_i^{k_i}}$. This works because

¹¹Stephen Pohlig and Martin Hellman. *An Improved Algorithm for Computing Logarithms over GF(p) and its Cryptographic Significance*. <http://www-ee.stanford.edu/~hellman/publications/28.pdf>. 1978

of the following observation (note that x can be written $x_1 + p_i^{k_i} x_2$ for some x_1 and x_2):

$$\begin{aligned} y^{\varphi(p)/p_i^{k_i}} &= (g^x)^{\varphi(p)/p_i^{k_i}} \pmod{p} \\ &= g^{(x_1 + p_i^{k_i} x_2)\varphi(p)/p_i^{k_i}} \pmod{p} \\ &= g^{x_1\varphi(p)/p_i^{k_i}} g^{x_2\varphi(p)} \pmod{p} \\ &= g^{x_1\varphi(p)/p_i^{k_i}} \pmod{p} \\ &= (g^{\varphi(p)/p_i^{k_i}})^{x_1} \pmod{p} \end{aligned}$$

The value we obtain is a generator of the subgroup of order $p_i^{k_i}$ raised to the power x_1 . By computing the discrete logarithm of this value we will obtain x_1 , which is the value of x modulo $p_i^{k_i}$. Generally we will use the *Pollard Rho* algorithm to compute that discrete logarithm.

The Chinese Remainder Theorem, sometimes used for good¹² will be of use here for evil. The following theorem states why it is possible for us to find a solution to our problem once we find a solution modulo each power prime factor of the order.

Theorem 1. Suppose $m = \prod_{i=1}^k m_i$ with m_1, \dots, m_k pairwise co-prime.

For any (a_1, \dots, a_k) there exists an x such that:

$$\begin{cases} x = a_1 \pmod{m_1} \\ \vdots \\ x = a_k \pmod{m_k} \end{cases}$$

There is a simple way to recover $x \pmod{m}$ which is stated in the following theorem:

Theorem 2. Moreover there exists a unique

¹²Shinde and Fadewar. *Faster RSA Algorithm for Decryption Using Chinese Remainder Theorem*. <http://www.techscience.com/doi/10.3970/icces.2008.005.255.pdf>

solution for $x \pmod{m}$:

$$x = \sum_{j=1}^k a_j * \left(\prod_{i \neq j} m_i \bar{m}_i \right) \pmod{m}$$

with $\bar{m}_j = m_j^{-1} \pmod{m_j}$

At first, it might be kind of hard to grasp where that formula is coming from. But let's see where it does by starting with only two equations. Keep in mind that we want to find the value of x modulo $m = m_1 m_2$

$$\left. \begin{aligned} x &= a_1 \pmod{m_1} \\ x &= a_2 \pmod{m_2} \end{aligned} \right\} \implies x = ? \pmod{m}$$

How can we start building the value of x ?

If $x = a_1 m_2 \pmod{m}$,

$$\text{then } \begin{cases} x = a_1 m_2 \pmod{m_1} \\ x = 0 \pmod{m_2} \end{cases}$$

Not quite what we want, but we are getting there. Let's add to it:

If $x = a_1 m_2 \bar{m}_2 \pmod{m}$

\bar{m}_2 the integer congruent to $m_2^{-1} \pmod{m_1}$

$$\text{then } \begin{cases} x = a_1 m_2 \bar{m}_2 = a_1 \pmod{m_1} \\ x = 0 \pmod{m_2} \end{cases}$$

That's almost what we want! Half of what we want actually. We just need to do the same thing for the other side of the equation, and we have:

$$\begin{aligned} &= a_1 m_2 \bar{m}_2 \pmod{m_1} \\ &= a_1 \pmod{m_1} \end{aligned}$$

↑

$$\boxed{x = a_1 m_2 \bar{m}_2 + a_2 m_1 \bar{m}_1 \pmod{m}}$$

↓

$$\begin{aligned} &= a_2 m_1 \bar{m}_1 \pmod{m_2} \\ &= a_2 \pmod{m_2} \end{aligned}$$

with \bar{m}_2 the integer congruent to $m_2^{-1} \pmod{m_1}$ and \bar{m}_1 the integer congruent to $m_1^{-1} \pmod{m_2}$.

Everything works as we wanted! Now you should understand better how we came up with that general formula. There have been improvements to it with the Garner's algorithm¹³ but this method is so fast anyway that it is not the bottleneck of the whole attack.

2.3 Small Subgroup Attacks

The attack we just visited is a passive attack: the knowledge of one Diffie-Hellman exchange between two parties is enough to obtain the following shared key. But instead of reducing one party's public key to an element of different subgroups, there is another clever attack called a small subgroup attack that creates the different subgroup generators directly and sends them to one peer successively to obtain its private key. It is an active attack that doesn't work against ephemeral protocols that renew the Diffie-Hellman public key for every new key exchange. This is for example the case with TLS when using ephemeral Diffie-Hellman (DHE) as a key exchange during the handshake.

The attack is straight forward and summed up below:

1. Determine the prime factorization of the order of the group

$$\varphi(p) = \prod p_i^{k_i}$$

2. Find a generator for every subgroup of order $p_i^{k_i}$, this can be done by picking a random element α and computing

$$\alpha^{\varphi(p)/p_i^{k_i}} \pmod{p}$$

¹³<http://www.csee.umbc.edu/~lomonaco/s08/441/handouts/GarnerAlg.pdf>

3. Send generators one by one as your public keys in different Diffie-Hellman key exchanges
4. Determine the value of x modulo $p_i^{k_i}$ for each shared key computed
5. Recompute $x \pmod{\varphi(p)}$ with the CRT

The fourth step can be done by having access to an oracle telling you what the shared key computed by the victim is. In TLS this is done by brute-forcing the possible solutions and seeing which one has been used by the victim in his following encrypted messages (for example the MAC computation in the Finish message during the handshake). With these constraints the attack would be weaker than Pohlig-Hellman since the brute-force is slower than *Pollard Rho*, or even trial multiplication. Because of the previously stated limitations and the fact that this attack only works for rather small subgroups, we won't use it in this work.

3 A First Backdoor Attempt in Prime Groups

The naive approach to creating a backdoor would be to weaken the parameters enough to make the computation of discrete logarithms affordable. Making the modulus a prime of a special form ($r^e + s$ with small r and s) would facilitate the Special Number Field Sieve (SNFS) algorithm. Having a small modulus would also allow for easier pre-computation of the General Number Field Sieve (GNFS) algorithm. It is believed¹⁴ that the NSA has enough power to achieve the

¹⁴David Adrian et al. "Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice". In: *22nd ACM Conference on Computer and Communications Security*. <https://weakdh.org/imperfect-forward-secrecy-ccs15.pdf>. Oct. 2015

first pre-computing phases of GNFS on 1024 bit primes which would then allow them to compute discrete logarithms in such large groups in the matter of seconds. But these ideas are pure computational advantages that involve no secret key to make the use of efficient backdoors possible. Moreover they are downright not practical: the attacker would have to re-do the pre-computing phase entirely for every different modulus, and the next generation of recommended modulus bitsize (2048+) would make these kind of computational advantages fruitless.

Another approach could be to use a generator of a smaller subgroup (without publishing what smaller subgroup we use) so that algorithms like *Pollard Rho* would be cost-effective again.

$$\varphi(p) = p - 1 = \boxed{p_1} \times \cdots \times p_k$$


\uparrow
 order
 $y = g^x \pmod{p}$

But then algorithms like *Pollard Kangaroo* that run in the same amount of time as *Pollard Rho* and that do not require the knowledge of the base's order could be used as well by anyone willing to try. This makes it a poorly hidden backdoor that we cannot qualify as NOBUS.

Section 4 uses both of these ideas to insert a backdoor into Diffie-Hellman by using a composite modulus. GNFS and SNFS can then be used modulo the factors of the composite modulus, or better as we will see, the generator's "small" subgroups can be concealed modulo the factors.

Back to our prime modulus. A second idea would be to set the scene for the Pohlig-Hellman algorithm to work. This can be done

by fixing a prime modulus p such that $p - 1$ is B-smooth with B small enough for discrete logarithms in bases of order B to be possible.



But this design is flawed in the same ways as the previous ones were: anyone can compute the order of the group (by subtracting 1 from p) and try to factor it. Choosing p such that $p - 1$ would include factors small enough to use one of the $\mathcal{O}(\sqrt{p})$ would make it dangerously factorisable. Using the *Elliptic Curve Method* (ECM), a factorization algorithm which complexity only depends on the size of the smallest factor (or for a full factorization, on the size of the second largest factor), the latest records¹⁵ were able to find factors of around 300 bits. This necessary lower bound on the factors makes it unfeasible to use any of the $\mathcal{O}(\sqrt{p})$ algorithms that would take, for example, more than 2^{150} operations to solve the discrete logarithm of 300 bit orders.

Our contribution in section 5 uses a composite modulus to hide the smoothness of the order as long as the modulus cannot be factored. This method is preferred from the other methods as it is not reversible and might only need a change of modulus. For example, in many DH parameters or implementations, $g = 2$ as a generator is often used. While the first method will not allow any easy ways to find a specific generator, our second method will.

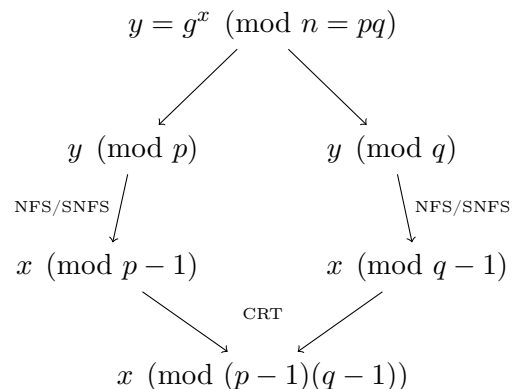
¹⁵<http://www.loria.fr/~zimmerma/records/top50.html>

4 A Second Backdoor Attempt With a Composite Modulus and a Hidden Subgroup

Update: this section was first presented as a way to construct a Nobody-But-Us backdoor, improving on the previous section. Unfortunately, Dorey et al.¹⁶ have pointed an attack on this method that makes the backdoor reversible. This section has been left as is, but it should be noted that there exist a way to reverse it according to Coron et al.¹⁷.


Our first backdoor gets around the previous problems using a composite modulus $n = pq$ with p and q large enough to avoid the factorization of n . This requires the same precautions used to secure RSA instances, with n typically reaching 2048 bits and with two factors p and q nearing the same size.

With the factorization of n known, the discrete logarithm problem can be reduced modulo p and q and solved there, before being reconstructed modulo $\varphi(n)$ with the help of the CRT theorem.



p and q could be hand-picked as SNFS primes, or we could use GNFS to compute the discrete logarithm modulo p and q . But a more efficient way exists to ease the discrete logarithm problem. Choosing a generator g such that both g modulo p and g modulo q generate “small” subgroups, would allow us to compute two discrete logarithms in two small subgroups instead of one discrete logarithm in one large group.

For example, we could pick p and q such that $p - 1 = 2p_1p_2$ and $q - 1 = 2q_1q_2$ with p_1 and q_1 two small prime factors and p_2, q_2 two large prime factors. Lagrange’s theorem tells us that the possible orders of the subgroups are divisors of the group order. This means we can probably find an element g of order p_1q_1 to be our Diffie-Hellman generator.



By reducing the discrete logarithm problem $y = g^x$ modulo p and q with our new backdoored generator, we can compute x modulo $p - 1$ and $q - 1$ more easily and then recompute an equivalent secret key modulo $(p - 1)(q - 1)$.

¹⁶Kristen Dorey, Nicholas Chang-Fong, and Aleksander Essex. *Indiscreet Logs: Persistent Diffie-Hellman Backdoors in TLS*. Cryptology ePrint Archive, Report 2016/999. <http://eprint.iacr.org/2016/999>. 2016

¹⁷Jean-Sebastien Coron et al. *Cryptanalysis of the RSA Subgroup Assumption from TCC 2005*. Cryptology ePrint Archive, Report 2010/650. <http://eprint.iacr.org/2010/650>. 2010

This will find the exact original secret key with a probability of $\frac{1}{4p_2q_2}$, which is tiny, but this doesn't matter since the shared key we will compute with that solution and the other peer's public key will be a valid shared key.

Proof. Let $a + k_ap_1q_1$ be Alice's public key for $k_a \in \mathbb{Z}$ and let $b + k_bp_1q_1$ be Bob's public key for $k_b \in \mathbb{Z}$,

then Bob's shared key will be $(g^{a+k_ap_1q_1})^{b+k_bp_1q_1} = g^{ab} \pmod{n}$.

Let $a + k_cp_1q_1$ be the solution we found for $k_c \in \mathbb{Z}$,

then the shared key we will compute will be $(g^{b+k_bp_1q_1})^{a+k_cp_1q_1} = g^{ab} \pmod{n}$, which is the same as Bob's shared key. \square

Update: here, the research of Dorey et al.¹⁸ pointed that p_1 and p_2 can be found in $\mathcal{O}(\max(p_1, p_2))$ operations with Coron et al's attack¹⁹, which makes this backdoor as easy to reverse as it is to exploit. Thus this can no longer be called a Nobody-But-Us backdoor. Fortunately, our contribution in section 5 still holds.

We used the *Pollard Rho* function in Sage 6.10 on a Macbook Pro with an i7 Intel Core @ 3.1GHz to compute discrete logarithms modulo safe primes of diverse bitsizes. The results are summed up in the table below.

order size	expected complexity	time
40 bits	2^{20}	01s
45 bits	2^{22}	04s
50 bits	2^{25}	34s

¹⁸Kristen Dorey, Nicholas Chang-Fong, and Aleksander Essex. *Indiscreet Logs: Persistent Diffie-Hellman Backdoors in TLS*. Cryptology ePrint Archive, Report 2016/999. <http://eprint.iacr.org/2016/999>. 2016

¹⁹Jean-Sebastien Coron et al. *Cryptanalysis of the RSA Subgroup Assumption from TCC 2005*. Cryptology ePrint Archive, Report 2010/650. <http://eprint.iacr.org/2010/650>. 2010

A stronger and more clever attacker would parallelize this algorithm on more powerful machines to obtain better numbers. To be able to exploit the backdoor "live" we want a running-time close to zero. Using a 80 bit integer as our generator's order, someone with no knowledge of the factorization of the modulus would take around 2^{40} operations to compute a discrete logarithm while this would take us on average 2^{21} thanks to the trapdoor. A more serious adversary with a higher computation power and a care for security might want to choose a 200 bit integer as the generator's order. For that he would need to be able to perform 2^{50} operations instantaneously if he would want to tamper with the encrypted communications following the key exchange, while an outsider would have to perform an "impossible" number of 2^{100} operations. The size of the two primes p and q , and of the resulting $n = pq$, should be chosen large enough to resist against the same attacks as RSA. That is a n of 2048 bits with p and q both being 1024 bit long would suffice.

To use such a backdoor, one must not only generate two primes p and q to satisfy the previous shape, but also find a specific generator g . This is not a hard task, unless you want to use a specific generator g . For example many libraries use $g = 2$ by default, implementing this backdoor would mean changing both the modulus and the generator. This is because the probability that an element in a group of order q is the generator of a subgroup of order d is $\frac{d}{q}$. This means that with our example $g = 2$, we would need to generate many modulus hoping that $g = 2$ as a generator would work. The probability that it would work for each try would be :

$$\frac{p_1p_2}{(p-1)(q-1)} \sim \frac{1}{pq} = \frac{1}{n}$$

This is obviously too small of a probability for us to try to generate many parameters until one admits our targeted g as a generator of our “small” subgroup. This is a problem if we want to only replace the modulus of an implementation to activate our backdoor. Since changing only one value would be more subtle than changing two values, our next contribution revise the way we generate the backdoor parameters to solve this problem.

5 A Composite Modulus for a NOBUS Backdoor with a B-Smooth Order


Let’s start again with a composite modulus $n = pq$, but this time let’s choose p and q such that $p - 1$ and $q - 1$ are both B-smooth with B small enough for the discrete logarithm to be doable in subgroups of order B. We’ll see later how to choose B.

Let $p - 1 = p_1 \times \dots \times p_k \times 2$ and $q - 1 = q_1 \times \dots \times q_l \times 2$ such that $\gcd(p - 1, q - 1) = 2$ and such that $p_i \leq B$ and $q_j \leq B$ for all $i \in \llbracket 1, k \rrbracket$ and $j \in \llbracket 1, l \rrbracket$ respectively. This makes the order of the group $\varphi(n) = (p - 1)(q - 1)$ B-smooth.

Constructing the Diffie-Hellman modulus this way permits anyone with both the knowledge of the order factorization and the ability of computing the discrete logarithm in subgroups of order B, to compute the discrete logarithm modulo n by using the Pohlig-Hellman method.

Since $p - 1$ and $q - 1$ are both B-smooth, they are susceptible to be factored with the *Pollard’s $p-1$* factorization algorithm, a factorization algorithm that can find a factor p of n if $p - 1$ is partially-smooth. RSA counters this problem using safe primes of the form $p = q + 1$ with q prime as well, but this would break our backdoor. Instead, as a way of countering *Pol-*

lard’s $p-1$ we can add a large factor to both $p - 1$ and $q - 1$ that we will call p_{big} and q_{big} respectively.



To exploit this backdoor we can reduce our public key y modulo p and q , as we did in our first method, and proceed with Pohlig-Hellman’s algorithm there. This is not a necessary step but this will reduce the size of the numbers in our calculations, speeding up the attack. We then carry on with CRT to recompute the private key modulo its order, which can be picked at a secure maximum of $\frac{(p-1)(q-1)}{2}$, which brings around the same security promises of a safe-prime modulus. This is because of the following isomorphy we have: $(\mathbb{Z}_n)^* \simeq (\mathbb{Z}_p)^* \times (\mathbb{Z}_q)^*$, with the product’s orders $s = |(\mathbb{Z}_p)^*|$ and $t = |(\mathbb{Z}_q)^*|$ not being coprimes ($\gcd(p - 1, q - 1) = 2$). This results in a non-cyclic group with an upper-bound on possible subgroup orders of $\text{lcm}(s, t) = \frac{(p-1)(q-1)}{2}$.

To decide how big p_{big} and q_{big} should be, we can look at the world’s records for the *Pollard’s $p-1$* factorization algorithm²⁰, the largest B2 parameter (the large factor) used in a factorization is $10^{15} \sim 50\text{bits}$ in 2015. As with our previous method, we could use much larger factors of around 100 bits to avoid any powerful adversaries and have a agreeable 2^{51} computations on average to solve the discrete logarithm problem in these large subgroups.

²⁰<http://www.loria.fr/~zimmerma/records/Pminus1.html>

This method brings the security of our overall scheme to the one of a perfectly secure Diffie-Hellman. Its security also relies on the RSA’s assumption that factoring n is difficult if n is large enough. More than that, the probability of having a targeted element be a valid generator can be as large as $\frac{1}{2}$ in our example of a secure subgroup of order $\frac{(p-1)(q-1)}{2}$. This will allow us to easily generate a backdoored modulus that will fit a specific generator, thus increasing the stealthiness of the implementation phase of our scheme. In the case where the large two subgroups of order p_{big} and q_{big} need to be avoided, one could think about trying to generate many modulus hopping that the targeted g would fit. This can be done with probability $\frac{1}{2^{p_{big}q_{big}}}$ for each try, which is way too low. But worse, this would give a free start to someone trying to factor the modulus using *Pollard’s $p-1$* . Another way would be to pass a fake order to the program, forcing it to generate small ephemeral private keys upper-bounded by $\frac{\varphi(n)}{p_{big}q_{big}}$. After this, proceeding to use Pohlig-Hellman over the small factors, ignoring p_{big} and q_{big} , would be enough to find the private key. This can be done in OpenSSL or libraries making use of it by passing the fake order to OpenSSL via `dh->q`. Of course doing this would bring back our first method’s issues by having to add extra lines of code to our malicious patch.

6 Implementing and Exploiting the Backdoor in TLS

Theoretically, any application including Diffie-Hellman might be backdoored using our method. As TLS is one of the most well known protocols using Diffie-Hellman it is particularly interesting to abuse for a field test of our work. Most TLS applications making use of the Diffie-Hellman algorithm for the handshake – although this is an algorithm rarely used in

TLS – would have their DH public key and parameters baked into user’s or server’s generated certificates. Interestingly, the parameters of the – much more commonly used – *ephemeral* version of Diffie-Hellman used to add the properties of *Perfect Forward Secrecy*, are rarely chosen by end users and thus never engraved into user’s or server’s certificates. Furthermore, most libraries implementing the TLS protocol (socat, Apache, NGINX, ...) have predefined or hardcoded ephemeral DH parameters. Developers using those libraries will rarely generate their own parameters and will use the default ones. This was the source of many discussions after being pointed out by *Logjam*²¹ last year, creating a movement of awareness, pushing people to migrate to bigger parameters and increase the bitsizes of application’s Diffie-Hellman modulus from 1024 or lower to 2048+ bits. This trend seems like the perfect excuse to submit a backdoored patch claiming to improve the security of a library. We’ll first see how TLS works with ephemeral Diffie-Hellman in the next section. Followed will be a demonstration on how the backdoor was implemented in real open source libraries. Finally we’ll explain how our setup worked to make use of the backdoor.

6.1 Background

An ephemeral handshake allows two parties to negotiate a “fresh” set of keys for every new TLS connection. This has become the preferred way of using TLS as it increases its security, providing the property that we call *Forward Secrecy* or *Perfect Forward Secrecy*, that is: if the long term key is compromised, recorded past communications won’t be af-

²¹David Adrian et al. “Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice”. In: *22nd ACM Conference on Computer and Communications Security*. <https://weakdh.org/imperfect-forward-secrecy-ccs15.pdf>. Oct. 2015

ected and future communications will still resist passive attacks. This is done by using one of the two Diffie-Hellman algorithms provided by TLS: “normal” Diffie-Hellman present in the ciphersuites containing DHE in their names and Elliptic Curve Diffie-Hellman (ECDH) present in the ciphersuites containing ECDHE in their names. Note that the concept of “*ephemeral*” is not defined the same by everyone: the default behavior of OpenSSL, up until recent versions, has been to generate the ephemeral DH key at boot time and cache it until reboot, unless specified not to do so. Such behavior would greatly speed up our attack.

At the start of a new *ephemeral handshake*, both the server and the client will send each other their ephemeral DH (DHE) public keys via a *ServerKeyExchange* and a *ClientKeyExchange* message respectively. The server will dictate as well what the DHE parameters are via the same *ServerKeyExchange* message.

```

▼ TLSv1.2 Record Layer: Handshake Protocol: Server Key Exchange
  Content Type: Handshake (22)
  Version: TLS 1.2 (0x0303)
  Length: 654
  ▼ Handshake Protocol: Server Key Exchange
    Handshake Type: Server Key Exchange (12)
    Length: 650
    ▼ Diffie-Hellman Server Params
      p Length: 128
      p: 9bd3030031d1db2287ef9e74c9ab7b646e38bc5196901b1d...
      g Length: 128
      g: 2e422f97728cc9b301014c6ee5624b37e49ceed3eedaf052...
      Pubkey Length: 128
      Pubkey: 239e9299b93ec58ab009b01b2e348529fea0f35b4ce6084e...
      ▶ Signature Hash Algorithm: 0x0601
      Signature Length: 256

```

Figure 3: The serverKeyExchange message

```

▼ TLSv1.2 Record Layer: Handshake Protocol: Client Key Exchange
  Content Type: Handshake (22)
  Version: TLS 1.2 (0x0303)
  Length: 134
  ▼ Handshake Protocol: Client Key Exchange
    Handshake Type: Client Key Exchange (16)
    Length: 130
    ▼ Diffie-Hellman Client Params
      Pubkey Length: 128
      Pubkey: 16194547c991725cadfd623c1293459b83ed526e7a65f09c...

```

Figure 4: The clientKeyExchange message

Let **c** and **s** be the client and the server public keys respectively. The following computation is done on each side, right after the key ex-

change, to derive the session keys that will encrypt further communications (including final handshake messages):

1. $premaster_secret = g^{cs} \pmod n$
2. $master_secret = PRF(premaster_secret, \text{“master secret”}, ClientHello.random + ServerHello.random)$
3. $keys = PRF(master_secret, \text{“key expansion”}, ServerHello.random + ClientHello.random)$

Right after trading their ephemeral DH public keys, the TLS peers compute the Diffie-Hellman algorithm by exponentiating the other’s public key with their own private key. The output is stored in a *premaster_secret* variable that is sent into a pseudo-random function (PRF) with the string “master secret” and the public values *random* of both parties taken from their Hello message as parameters. This is because the DH output can be of fluctuating lengths: TLS offers several parameters and algorithms to perform this part of the handshake, passing it through a PRF first aims to normalize its size before deriving the keys from it. The output of that first PRF is then sent into the same PRF repeatedly along with different arguments: the string “key expansion” and the reversed order of the *random* values we just used, until enough bits are produced for the many keys used to encrypt and authenticate the post-handshake communications.

Two authentication keys are first derived, *client_write_MAC_key* and *server_write_MAC_key*, one for each direction. Then two encryption keys are derived as well, *client_write_key* and *server_write_key*. For AEAD ciphers, MAC keys are ignored and two more values after the encryption keys are derived: *client_write_IV* and *server_write_IV*.

6.2 Implementation

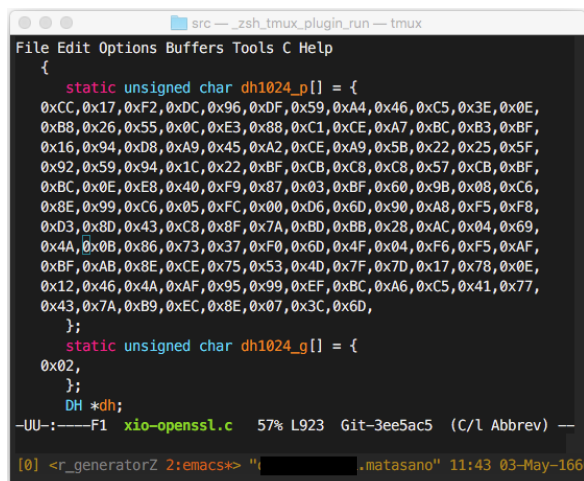


Figure 5: socat’s xio-openssl.c file

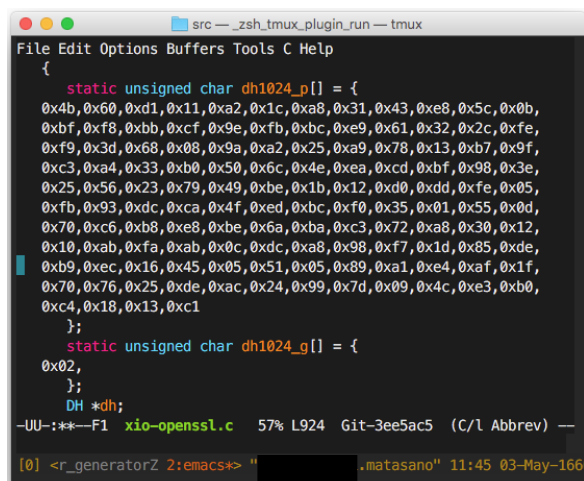


Figure 6: The same backdoored socat file after our changes

socat makes use of OpenSSL underneath the surface, and being a useful nifty command line tool, it is what we first used to quickly test our backdoor parameters. With the help of the `dhparam` argument, we can use an ASN.1 file encoded in the DER format to specify our backdoored DH parameters to the program. But as this exercise is one of implementation, we will show how we modified the

source code of socat to introduce the backdoor to the masses. We started by generating small parameters with the method in section 5 for our tests. Since socat’s generator is 2, we only had to modify the modulus in the `dh1024_p` variable of the `xio-openssl.c` file. Above are the differences between a valid version and a backdoored version of socat.

If we want to submit such a patch as an excuse to increase the size of the DH parameters, we could mimick socat’s case²²:

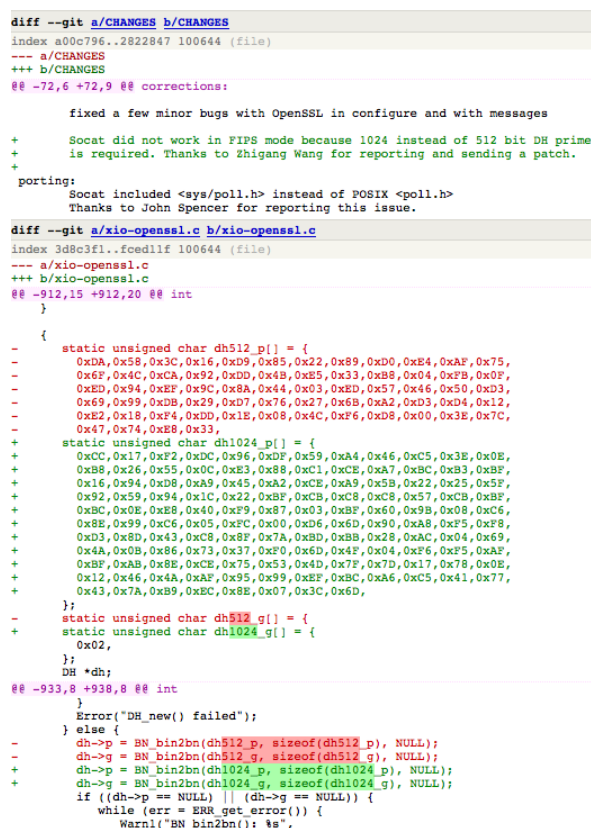


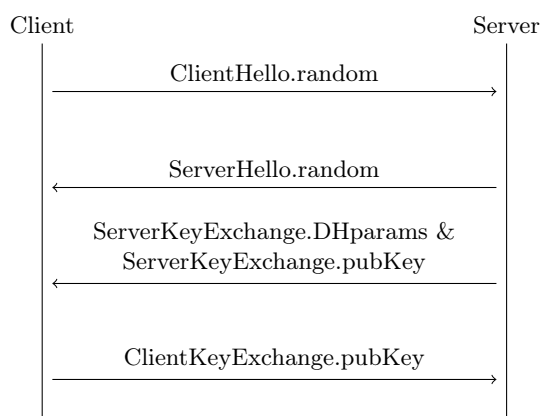
Figure 7: The commit diff page of the socat “presumed” backdoor

6.3 Exploitation

To exploit this kind of backdoor, we first need to obtain a Man-In-The-Middle position be-

²²<http://repo.or.cz/socat.git/commitdiff/281d1bd6515c2f0f8984fc168fb3d3b91c20bdc0>

tween the client and the server. This can be done passively by obtaining posterior access to logs of TLS records, but this was done actively in our proof of concept²³ by using a machine as a proxy to the server and making the client connect to the proxy directly instead of the server. The proxy unintelligently forwards the packets back and forth until a TLS connection is initiated, it then observes the handshake, storing the *random* values at first, until the server decides to send its public Diffie-Hellman parameters to be used in an *ephemeral* key exchange. If the proxy recognizes the backdoor parameters in the server’s *ServerKeyExchange* message, it runs the attack, recovering one party’s private key and computing the session keys out of that information. With the session keys in hand, the proxy can then observe the traffic in clear and even tamper with the messages being exchanged.



Depending on the security margins chosen during the generation of the backdoor, and on the computing power of the attacker, it may be the case that the attacker would not be able to derive the session keys until the first few messages have been exchanged, exempting them from tampering. For better results, the work could be parallelized and the two

²³https://github.com/mimoo/Diffie-Hellman_Backdoor

public keys could be attacked simultaneously as one might be recovered more quickly than the other one. As soon as the private key of one party is recovered, the Diffie-Hellman and the session keys computations are done in a negligible time, and the proxy can start live decrypting and live tampering with the packets. If the attacker really wants to be able to tamper with the first messages, it can delay the end of the handshake by sending *TLS warning alerts* that can keep a handshake alive indefinitely or for a period of time depending on the TLS implementation used by both parties.

Update: exploitation of this backdoor has been improved by Dorey et al.²⁴ who pointed that an active man-in-the-middle attacker can also force both a client and a server to use a DHE cipher suite (to force them to use the backdoor) if both peers already support a DHE cipher suite.

7 Detecting a Backdoor and Defending Against One

In the course of this work several open source libraries were tested for composite modulus with no positive results. TLS handshakes of the full range of IPv4 addresses obtained from scans.io were inspected on March 3rd, 2016. A total of 50,222,805 handshakes were analyzed from which 4,522,263 were augmented with the use of ephemeral Diffie-Hellman. From these numbers, only 30 handshakes used a composite modulus, all of them had a small factor but none of them could be factored in less than 5 hours using the *ECM* or *Pollard’s p-1* fac-

²⁴Kristen Dorey, Nicholas Chang-Fong, and Alexander Essex. *Indiscreet Logs: Persistent Diffie-Hellman Backdoors in TLS*. Cryptology ePrint Archive, Report 2016/999. <http://eprint.iacr.org/2016/999>. 2016

torization algorithms. Most IPs were hosting webpages, in some cases the same one. All administrators were contacted about the issue. Our contributions should withstand any kind of reversing and thus we should not be able to detect any backdoor produced by people who would have reached the same conclusions as ours. The addition of easy to find small factors could have been intentionally done to provide plausible deniability. Interestingly, it is also hard to differentiate a mistake in the modulus generation from a backdoor. From the Handbook of Applied Cryptography²⁵ fact 3.7:

Definition 1. *Let n be chosen uniformly at random from the interval $[1, x]$.*

1. *if $1/2 \leq \alpha \leq 1$, then the probability that the largest prime factor of n is $\leq x^\alpha$ is approximately $1 + \ln(\alpha)$. Thus, for example, the probability that n has a prime factor $> \sqrt{x}$ is $\ln(2) \approx 0.69$*
2. *The probability that the second-largest prime factor of n is $\leq x^{0.2117}$ is about $1/2$.*
3. *The expected total number of prime factors of n is $\ln \ln x + \mathcal{O}(1)$. (If $n = \prod p_i^{e_i}$, the total number of prime factors of n is $\sum e_i$.)*

Since it might be easier to visualize this with numbers:

1. a 1024 bit composite modulus n probability to have a prime factor greater than 512 bits is ≈ 0.69 .
2. the probability that the second-largest prime factor of n is smaller than 217 bits is $1/2$.
3. The total number of prime factor of n is expected to be 7.

²⁵<http://cacr.uwaterloo.ca/hac/about/chap3.pdf>

Considering that a full factorization with ECM runs in a complexity tied to the size of the second largest factor, it might be hard or impossible to do it half of the time. The rest of the time it might take a bit of work, but since the largest factor found using ECM²⁶ is 274 bits, it is possible.

The question of how to avoid these kinds of backdoors or weaknesses is also interesting and well understood, but rarely done correctly. First, it is known that by using safe primes – primes of the form $2q + 1$ with q prime – the generator’s subgroup will have an order close to the modulus’ size. Since it is easy to check if a number is a safe prime, the client should also only accept such moduli. The current state is that most programs don’t even check for prime modulus. As an example, no browser currently warns the user if a composite modulus is detected.

Another way to prevent this is to have a pre-defined list of public parameters²⁷, this would make Diffie-Hellman look similar to Elliptic Curve Diffie-Hellman in the sense that only a few curves are pre-defined and accepted in most exchanges.

Both mitigations can be hard to integrate if the two endpoints of a key exchange are not controlled. For example this is the case between

²⁶<http://www.loria.fr/~zimmerma/records/top50.html>

²⁷Matt Lepinski and Dr. Stephen T. Kent. *RFC 5114: Additional Diffie-Hellman Groups for Use with IETF Standards*. RFC 5114. <https://rfc-editor.org/rfc/rfc5114.txt>. 2015. DOI: [10.17487/rfc5114](https://doi.org/10.17487/rfc5114), Tero Kivinen. *RFC 3526: More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)*. RFC 3526. <https://rfc-editor.org/rfc/rfc3526.txt>. 2015. DOI: [10.17487/rfc3526](https://doi.org/10.17487/rfc3526), Daniel Kahn Gillmor. *IETF Draft: Negotiated Finite Field Diffie-Hellman Ephemeral Parameters for TLS*. Internet-Draft draft-ietf-tls-negotiated-ffdhe-10. <https://tools.ietf.org/html/draft-ietf-tls-negotiated-ff-dhe-10>. Internet Engineering Task Force, 2015

browsers and websites TLS connections where the browser is a different program from what is running on the server. Asserting for these special primes might just break the connection, which would be worse from the user’s perspective. This is why Google Chrome is currently removing DHE from its list of supported cipher suites²⁸, and recommending server administrators to migrate from DHE to ECDHE. This is also one of the recommendations from Logjam²⁹. These security measures might very well prevent this work’s efforts: backdooring ECDHE in a stealthy way as we did with DHE remains an open problem.

8 Conclusion

Many cryptographic constructions are not subject to change, unless a breakthrough comes along and the whole construction has to be replaced. Very rarely the excuse of updating a reviewed and considered strong cryptographic implementation to change a single number comes along, and very few people understand such subtle changes. According to the grading system of a whitepaper by Schneier et al³⁰, here is how such a backdoor scores:

- *medium undetectability*: to discover the backdoor one would have to test for the primality of the modulus. A pretty easy task, although not typically performed as seen with the socat’s case where it took

²⁸<https://groups.google.com/a/chromium.org/forum/m/#!topic/blink-dev/AAdv838-koo/discussion>

²⁹David Adrian et al. “Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice”. In: *22nd ACM Conference on Computer and Communications Security*. <https://weakdh.org/imperfect-forward-secrecy-ccs15.pdf>. Oct. 2015

³⁰Bruce Schneier et al. *Surreptitiously Weakening Cryptographic Systems*. Cryptology ePrint Archive, Report 2015/097. <http://eprint.iacr.org/>. 2015

them more than a year to realize the composite modulus.

- *high lack of conspiracy*: in the case of socat only the person who had submitted the vulnerability would be the target of investigation. It turns out he is a regular employee at Oracle.
- *high plausible deniability*: three things help us in the creation of a good story in the socat’s case: reversing bytes of the fake prime gives us a prime, some small factors were found, anyone with weak knowledge of cryptography could have submitted a composite number.
- *medium ease of use*: man-in-the-middleing the attack and observing the first handshake would allow the attacker to take advantage of the backdoor.
- *high severity*: having access to that backdoor lets us observe, or if exploited in real time let us tamper, with any communications made over TLS.
- *medium durability*: system admins would have to update to newer versions to remove the backdoor.
- *high monitorability*: the saboteur cannot detect if other attackers are taking advantage of the backdoor, which is OK since the backdoor in this work are NOBUS ones.
- *high scale*: backdooring an open-source library would allow access to many systems’ and users’ communications.
- *high precision*: the saboteur doesn’t weaken any system, only the saboteur himself can access the backdoor.
- *high control*: like Dual-EC, only the saboteur can exploit the backdoor.

While this work is mostly a fictive exercise, we hope to raise awareness in the need for better toolings and deeper reviews of open source – as well as closed source – implementations of cryptographic algorithms.

Acknowledgements

Many thanks to both Scott Fluhrer and Scott Contini who have been of precious help in the core ideas of this paper.

Thanks as well to Pete L. Clark, Tom Ritter, Mike Brown, Roman Zabicki, Vincent Lynch, Drew Suarez, Ryan Koppenhaver, Divya Natesan, Andy Grant and Jake Massimo for feedback and discussions.

References

- [1] David Adrian et al. “Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice”. In: *22nd ACM Conference on Computer and Communications Security*. <https://weakdh.org/imperfect-forward-secrecy-ccs15.pdf>. Oct. 2015.
- [2] Daniel J. Bernstein, Tanja Lange, and Ruben Niederhagen. *Dual EC: A Standardized Back Door*. Cryptology ePrint Archive, Report 2015/767. <http://eprint.iacr.org/2015/767>. 2015.
- [3] Stephen Checkoway et al. *A Systematic Analysis of the Juniper Dual EC Incident*. Cryptology ePrint Archive, Report 2016/376. <http://eprint.iacr.org/2016/376>. 2016.
- [4] Jean-Sebastien Coron et al. *Cryptanalysis of the RSA Subgroup Assumption from TCC 2005*. Cryptology ePrint Archive, Report 2010/650. <http://eprint.iacr.org/2010/650>. 2010.
- [5] Kristen Dorey, Nicholas Chang-Fong, and Aleksander Essex. *Indiscreet Logs: Persistent Diffie-Hellman Backdoors in TLS*. Cryptology ePrint Archive, Report 2016/999. <http://eprint.iacr.org/2016/999>. 2016.
- [6] Daniel Kahn Gillmor. *IETF Draft: Negotiated Finite Field Diffie-Hellman Ephemeral Parameters for TLS*. Internet-Draft draft-ietf-tls-negotiated-ff-dhe-10. <https://tools.ietf.org/html/draft-ietf-tls-negotiated-ff-dhe-10>. Internet Engineering Task Force, 2015.
- [7] Tero Kivinen. *RFC 3526: More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)*. RFC 3526. <https://rfc-editor.org/rfc/rfc3526.txt>. 2015. DOI: 10.17487/rfc3526.
- [8] Matt Lepinski and Dr. Stephen T. Kent. *RFC 5114: Additional Diffie-Hellman Groups for Use with IETF Standards*. RFC 5114. <https://rfc-editor.org/rfc/rfc5114.txt>. 2015. DOI: 10.17487/rfc5114.
- [9] Stephen Pohlig and Martin Hellman. *An Improved Algorithm for Computing Logarithms over $GF(p)$ and its Cryptographic Significance*. <http://www-ee.stanford.edu/~hellman/publications/28.pdf>. 1978.
- [10] Eric Rescorla. *RFC 2631: Diffie-Hellman Key Agreement Method*. RFC 2631. <https://rfc-editor.org/rfc/rfc2631.txt>. 2013. DOI: 10.17487/rfc2631.

- [11] Bruce Schneier et al. *Surreptitiously Weakening Cryptographic Systems*. Cryptology ePrint Archive, Report 2015/097. <http://eprint.iacr.org/>. 2015.
- [12] Shinde and Fadewar. *Faster RSA Algorithm for Decryption Using Chinese Remainder Theorem*. <http://www.techscience.com/doi/10.3970/icces.2008.005.255.pdf>.
- [13] Shumow and Ferguson. *On the Possibility of a Back Door in the NIST SP800-90 Dual Ec Prng*. Crypto 2007. <http://rump2007.cr.yp.to/15-shumow.pdf>. 2007.
- [14] Robert Zuccherato. *RFC 2785: Methods for Avoiding the Small-Subgroup Attacks on the Diffie-Hellman Key Agreement Method for S/MIME*. RFC 2785. <https://rfc-editor.org/rfc/rfc2785.txt>. 2013. DOI: [10.17487/rfc2785](https://doi.org/10.17487/rfc2785).