

# Notes on OpenSSL remote memory corruption

*Posted on [June 27, 2022](#) by [guidovranken](#)*

OpenSSL version 3.0.4, released on June 21st 2022, is susceptible to remote memory corruption which can be triggered trivially by an attacker. BoringSSL, LibreSSL and the OpenSSL 1.1.1 branch are not affected. Furthermore, only x64 systems with AVX512 support are affected. The bug is fixed (<https://github.com/openssl/openssl/pull/18626>) in the repository but a new release is still pending.

Somewhat peculiarly, almost nobody is talking about this. If RCE exploitation is possible this makes it worse than Heartbleed in an isolated severity assessment, though the potential blast radius is limited by the fact that many people are still using the 1.1.1 tree rather than 3, libssl has forked into LibreSSL and BoringSSL, the vulnerability has only existed for a week (HB existed for years) and an AVX512-capable CPU is required.

This post gives some background on how this bug came to be and some notes on its exploitability.

On May 31st 2022 I found and reported an issue with constant-time Montgomery modular exponentiation in OpenSSL and BoringSSL (but not LibreSSL). For some values of  $B$ ,  $E$ ,  $M$  for which  $B \wedge E \% M \neq 0$ , some functions would return  $M$  instead of 0; the result was not fully reduced.

It was found that there are four distinct code paths affected by this:

- RSAZ 1024  
([https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn\\_exp.c#L668-L669](https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn_exp.c#L668-L669)).
- RSAZ 512  
([https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn\\_exp.c#L678](https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn_exp.c#L678)).
- Dual 1024 RSAZ  
([https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn\\_exp.c#L1448-L1452](https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn_exp.c#L1448-L1452)).
- Default constant-time Montgomery modular exponentiation  
([https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn\\_exp.c#L687-L1124](https://github.com/openssl/openssl/blob/b2feb9f0e394da6570346598837f1b01eb58c028/crypto/bn/bn_exp.c#L687-L1124)).

David Benjamin of Google analyzed the issue extensively.

(<https://boringssl.googleusercontent.com/boringssl/+13c9d5c69d04485a7a8840c12185c832026c8315>).

and found that the bug does not constitute a security risk (at least not internally; external callers might end up in incorrect states depending on what they are trying to compute). Interestingly, he also found an apparent bug in the paper

(<https://boringssl.googleusercontent.com/boringssl/+801a801024febe1a33add5ddaa719e257d97aba5>) by Shay Gueron upon which the RSAZ code is based.

This took my [fuzzer \(https://github.com/guidovranken/cryptofuzz\)](https://github.com/guidovranken/cryptofuzz) a long time to find, because the odds of finding  $B^M \pmod E = 0$  for large, N-bit values of B, E, M where those values are semi-random are small, so I proceeded to add a [modular exponentiation solver, \(https://github.com/guidovranken/cryptofuzz/blob/dd20f87e98ca92e9e31f8dd9d0a80efb5c249df7/expmod.cpp\)](https://github.com/guidovranken/cryptofuzz/blob/dd20f87e98ca92e9e31f8dd9d0a80efb5c249df7/expmod.cpp) so the bug can now be found quite quickly, and hopefully it helps finding more, similar bugs in the future (in OpenSSL or elsewhere). (Z3's performance in solving equations involving modular exponentiation is quite poor so I had to roll my own).

The [fix](https://github.com/openssl/openssl/commit/10d8a109be0fe50315e4eeb0676f6571914cd47a#diff-854689acd04b8f1b65120880bebcd98d519e89b601328820f276ec0e5c164c4f)

[\(https://github.com/openssl/openssl/commit/10d8a109be0fe50315e4eeb0676f6571914cd47a#diff-854689acd04b8f1b65120880bebcd98d519e89b601328820f276ec0e5c164c4f\)](https://github.com/openssl/openssl/commit/10d8a109be0fe50315e4eeb0676f6571914cd47a#diff-854689acd04b8f1b65120880bebcd98d519e89b601328820f276ec0e5c164c4f) that was applied to the dual 1024 RSAZ code is wrong because the [reduction function is called \(https://github.com/openssl/openssl/blob/eea820f3e239a4c11d618741fd5d00a6bc877347/crypto/bn/rsaz\\_exp\\_x2.c#L260-L261\)](https://github.com/openssl/openssl/blob/eea820f3e239a4c11d618741fd5d00a6bc877347/crypto/bn/rsaz_exp_x2.c#L260-L261) with num set to the bit size, where it should be number of BN\_ULONG elements (which are always 8 bytes large, because that is the size of an unsigned long on x64 systems, which is the only architecture which can have AVX512 support). So with the input sizes being 1024 bits, 8192 bytes are accessed (read from or written to) instead of 128.

On to the bug internals. There are 5 distinct arrays involved. 3 arrays are overwritten.

Variable	Description	Allocated size	Over-read/write	Total	Read/write?
res1	modexp result 1	128	896	8192	Read, then write
m1	Modulus 1	128	896	8192	Read
res2	modexp result 1	128	896	8192	Read, then write
m2	Modulus 2	128	896	8192	Read
storage	Scratch space	1184	7296	8192	Write, then read

- 8192 bytes are read from res1, res2, m1, m2 and storage
- 8192 bytes are written to res1, res2 and storage (this is where the memory corruption takes place)
- If we consider res1\_bn to be a bignum comprising res1[0..8192] (where the last byte is the most significant byte) and m1\_bn to be m1[0..8192], then if res1\_bn < m1\_bn, the contents of res1[0..8192] will be left unchanged after being overwritten. The same applies for res2 and m2.
- This implies that if you can set the most significant bit of m1[8191] to 1 and the most significant bit of res1[8191] to 0, then res1[0..8192] will retain its original state (and no actual corruption takes place). This circumstance may occur by chance. The same applies for res2 and m2.
- Conversely, if res1\_bn >= m1\_bn, then after the write, res1[N] will be one of {res1[N], res1[N] - m1[N], res1[N] - m1[N] - 1}. The same applies for res2 and m2.
- After the overwrites have occurred, storage[N] will be ~(m2[N] - res2[N]) or ~(m2[N] - res2[N]) + 1.
- From this it follows that if you control m2[N] and res2[N], you mostly control storage[N].
- The original contents of storage is never read, so it does not influence the end state in any way.

Summarized:

Variable	Post-write value
res1[N]	$\underline{res1[N].} \text{ or } \underline{res1[N] - m1[n].} \text{ or } \underline{res1[N] - m1[N] - 1}$
res2[N]	$\underline{res2[N].} \text{ or } \underline{res2[N] - m2[n].} \text{ or } \underline{res2[N] - m2[N] - 1}$
storage[N]	$\sim(\underline{m2[N] - res2[N].}) \text{ or } \sim(\underline{m2[N] - res2[N].}) + 1$

OpenSSL vulnerability post-write states

Each of these lemma's are true independent of what the inputs to the modexp function are, and of any other variable or state.

Here is a fuzzer (<https://gist.github.com/guidovranken/b1c46bd9e42e959519009681b261a896>), which demonstrates these invariants.

The (wrapping) subtraction mechanics at play here are interesting because you can use them to apply pointer delta's to a function pointer to make it point to a different function and this can help circumvent ASLR ([https://en.wikipedia.org/wiki/Address\\_space\\_layout\\_randomization](https://en.wikipedia.org/wiki/Address_space_layout_randomization)), because while the function addresses are randomized, their spacing is constant for every given binary.

For example, if:

- a single, particular BN\_ULONG in the res1 array is known to contain an active pointer to a specific function
- and you can control the BN\_ULONG in m1 at the same index (for example via heap spraying)
- and the rest of the state is completely unknown and uncontrolled

then by setting  $m1[N]$  to  $oldfunc - newfunc$ , the original function pointer will be exactly  $newfunc$  after the overflow with a 25% probability:

```

1 void goodfunc(void) { printf("good\n"); }
2 void badfunc(void) { printf("bad\n"); }
3 int main(void)
4 {
5     /* Indices 0..127 are within allocated bounds */
6     /* Beyond that is either unallocated or allocated by different p
7     * of the program.
8     */
9     struct {
10         BN_ULONG storage[1024], res1[1024], m1[1024], res2[1024], m2[
11     } vars;
12
13     /* Randomize everything; not known to the attacker */
14     FILE* fp = fopen("/dev/urandom", "rb");
15     assert(fread(&vars, 1, sizeof(vars), fp));
16     fclose(fp);
17
18     /* Assume res1[345] contains a function pointer used internally k
19     vars.res1[345] = (BN_ULONG)&goodfunc;
20     /* Assume we control m1[345] */
21     vars.m1[345] = (BN_ULONG)&goodfunc - (BN_ULONG)&badfunc;
22
23     bn_reduce_once_in_place(vars.res1, /*carry=*/0, vars.m1, vars.stc
24     bn_reduce_once_in_place(vars.res2, /*carry=*/0, vars.m2, vars.stc
25
26     void (*fnptr)(void) = (void*)vars.res1[345];
27     fnptr(); /* good or bad? */
28     return 0;
29 }

```

OpenSSL makes heavy use of function pointers. Running `find -name '.c' -exec grep 'METH. = {' {} \; | grep -v test` from the repository root shows over 130 data structures that encapsulate a set of function pointers. Delta subtraction may be useful in exploiting this circumstance to make OpenSSL misbehave to varying degrees of severity.

Apart from code execution, there can also be scenario's where private data is leaked to the attacker.

Assume:

- $R^1 = \text{res}[I..J]$
- $M^1 = \text{m}[I..J]$
- $I, K \geq 128$
- $J, L \leq 1023$

Let  $R^1$  be an allocated space which the attacker can read and write (for example an internal TLS state whose value is, to an extent, determined by how the attacker conducts the handshake, and which can be read, to an extent, by how the TLS code subsequently behaves based on its state).

Let  $M^1$  contain some kind of secret information.

Recall that the combination of  $\text{res}[N]$  (before the overwrite) and  $\text{m}[N]$  leaks into  $\text{res}[N]$ . It follows that if the attacker has read-write access to  $R^1$ ,  $M^1$  can be partially or completely inferred.

This is what I've deduced so far. It's possible there are errors in this post; please e-mail [guido@guidovranken.com](mailto:guido@guidovranken.com) (<mailto:guido@guidovranken.com>) and I'll correct them and credit you. You can [follow me on Twitter \(https://twitter.com/GuidoVranken\)](https://twitter.com/GuidoVranken).

This site uses Akismet to reduce spam. [Learn how your comment data is processed.](#)

*[Blog at WordPress.com.](#)*