

TCmalloc (undetected overflows, blocks corruption)		Block offset: 8 MB in 9 MB Area
0	0x1dce008	'AAA800#'
1	0x1dce010	'AAA800#'
2	0x1dce018	'AAA800#'
3	0x1dce020	'AAAAAAAAAAAAAAAAAAAAA800#'
4	0x1dce028	'AAAAAAA800#'
0	0x1dce008	'ZZ800#'
1	0x1dce010	'ZZ800#'
2	0x1dce018	'ZZ800#'
3	0x1dce020	'ZZ800#'
4	0x1dce028	'ZZ800#'

TCmalloc “segregates” its freelist metadata and overflows are undetected. It re-uses freed blocks in their natural order so the second buffer overflow replaces the first overflow. The second pass freeing and re-using overflowed blocks is also undetected. There are no integrity checks. Google is not afraid of memory errors, maybe because that's not its data that are compromised [17] [30].

MIMMalloc (undetected OOB, blocks+freelist corruption)		Block offset: 128 KB in 64 MB Area
0	0x4b040020080	'AAA@'
1	0x4b040020088	'AAA@'
2	0x4b040020090	'AAAAAAAAAAAAAAAAAAAAAAAAAAAAA@'
3	0x4b040020098	'AAAAAAAAAAAAA@'
4	0x4b0400200a0	'AAAAAAA@'
0	0x4b0400200a8	'ZZ@'
1	0x4b0400200b0	'ZZ@'
2	0x4b0400200b8	'ZZ@'
3	0x4b0400200c0	'ZZ@'
4	0x4b0400200c8	'ZZZZZZZZZZ@'

MIMMalloc keeps using in-place freelist metadata (like GLibC). No freed block are re-used because it keeps allocating new space for (much) longer than our loop. Not reallocating blocks is why the two overflows are undetected even after two free/malloc loops – and despite the fact that MIMMalloc uses compiler options designed to catch buffer overflows [35] as recommended by the NSA [28].

Enabling MIMMalloc “secure” options (encoded freelist, random blocks) requires a recompilation:

MIM “secure” (undetected OOB, blocks/free corruption)		Block offset: 0-4 KB in 63 MB Area
0	0x2bbec020580	'AAAZU''
1	0x2bbec020500	''
2	0x2bbec020940	''
3	0x2bbec0207c0	''
4	0x2bbec020400	''
0	0x2bbec020e00	'ZZuU''
1	0x2bbec020dc0	''
2	0x2bbec020a40	''
3	0x2bbec020280	''
4	0x2bbec020a00	''

Overflows and `freelist` corruptions are also undetected by MIMalloc “secure”. Like with Glibc, metadata corruption will cause an `abort()` only after a new allocation meets a corrupted `freelist`. Microsoft just delays the checks made by `free()` and `malloc()` with a pseudo-random numbers generator to pick a new random block when `malloc()` is invoked. Detecting corruption long after damage has been done lets attackers to their job and hide their tracks to stay undetected.

All these allocators could detect corruption with instant integrity checks. They just don't do it. As a result, application/system buffer overflows and metadata corruption are... undetected.

Memory corruption can generate immediate consequences (crash, control-flow hijacking) or long-delayed visible effects (crash, garbage state and output). Primary causes are increasingly more difficult to find as they get older – hidden behind smoke and mirrors (“*garbage in, garbage out*”).

Then developers are blamed. In reality, they might be totally innocent – the OS allocator is guilty:

SLIMalloc (on-the-fly <u>detection, blocking and reporting</u> – no overflow, no corruption)	
<pre>> OOB: memset() accessed:40 block-size:8 OOB:32 caller ./oob.c:9 main() 0 0x41e254c6a700 'AAAAAAAA' 1 0x41e254c6a708 '' 2 0x41e254c6a710 '' 3 0x41e254c6a718 '' 4 0x41e254c6a720 '' > OOB: memset() accessed:40 block-size:8 OOB:32 caller ./oob.c:14 main() 0 0x41e254c6a720 'ZZZZZZZZ' 1 0x41e254c6a718 '' 2 0x41e254c6a710 '' 3 0x41e254c6a708 '' 4 0x41e254c6a700 ''</pre>	<pre>Here heap->opt.rnd_block = false so block-address randomization is not active (all blocks are allocated with their natural order and increment - instead of at random addresses) but the result is the same in both cases. Changing options is done on-the-fly, on a per- heap basis, without recompilation.</pre>

SLIMalloc is the only allocator to (1) traverse the first allocation/free pass unharmd, and to (2) complete the second allocation/free pass without damage, by (3) detecting, blocking, locating and reporting the exact location and precise cause of the OOB memory access violation on-the-fly.

JE/TC/MIMalloc “segregate” metadata in their block-area... exposed to undetected OOBs, defeating the very purpose of segregated metadata, which was initially to serve security. [14]

50 years ago, *C/asm* programmers had the technical background to understand the implications:

“If an error in an operating system program allows a penetration program to work, that program will work every time it is executed – without detection.” [6]
 – Electronics Systems Division, Air Force Systems Command, Computer Security Developments Summary, report MCI-74-1 (1973)

The flawed memory-allocator vendors (claiming to be the root cause of 70-90% of all vulnerabilities for decades) now want to secure the world with unsafe [34] “*memory-safe*” languages. Ahem, really?