

Proposal for C2x

WG14 N3813

Title: Memory allocation with size feedback

Author, affiliation: Chris Kennelly, Google; Justin King, Google; Charles Munger, Google

Date: 2026-02-02

Proposal category: New features

Target audience: Developers working on C codebases, allocator developers

Abstract: Provide a function that returns a size along with the allocated memory block, to standardize a number of existing extensions and permit more effective use of memory.

Prior art: Numerous allocator extensions

Memory allocation with size feedback

Summary

Depending on the malloc implementation (referred to as the allocator), the requested allocation size may be less than the actual allocated size. This results in memory being unused which is equal to the difference between the requested allocation size and the actual allocated size. Some allocators, such as the one used by glibc, return mmap'ed memory directly when the requested size is over a certain threshold. In the worst case, this can result in close to a full page size of memory being unused.

To make use of this unused memory, various allocator extensions may be used. These extensions fall into two categories: querying size classes before allocating (`malloc_good_size/nallocx`) and querying the actual allocated size after allocating (`malloc_usable_size/malloc_size/_msize`). Some allocators support both, one, or rarely none. Both approaches have shortcomings for this use case which are explained in detail later. We propose standardizing on a third approach which avoids these shortcomings, by introducing sibling functions for `malloc` and `aligned_alloc` which return the actual size of the allocation. This is also known as size feedback.

Background & Motivation

Querying Size Classes

As mentioned in the summary, some allocators have exposed extensions for querying size classes before allocating. The most prominent extensions are: `nallocx` (jemalloc/tcmalloc) and `malloc_good_size` (Apple). No similar extension is exposed by glibc or Microsoft's CRT. This can be used as a basis for implementing size feedback, however we have avoided standardizing on this approach for a few reasons. This approach has a couple downsides as it requires two function calls to perform an allocation: `nallocx/malloc_good_size` followed by `malloc`. The first downside is that the logic performed by `nallocx/malloc_good_size` to determine the size class will have to be repeated by `malloc` as well, as it is unaware a size class has already been determined. The second downside is the potential mismatch between the size class picked by `nallocx/malloc_good_size` and `malloc`. For example, the allocator could choose to return a slightly larger memory block that is in its freelist which would result in unused memory.

Querying Allocation Size

As mentioned in the summary, some allocators have exposed extensions for querying the actual size after allocating. The most prominent extensions are: `malloc_usable_size` (glibc/jemalloc/tcmalloc/mimalloc), `malloc_size` (Apple), and `_msize` (Microsoft's CRT). This can be used as a basis for implementing size feedback, however we have avoided standardizing on this approach for a few reasons. This approach requires three function calls to perform an allocation: `malloc` followed by `malloc_usable_size/malloc_size/_msize` potentially followed by `realloc`.

1. The logic for determining the size of the allocation by `malloc_usable_size/malloc_size/_msize` is the same logic used by `free` and *avoided by `free_sized/free_aligned_sized`*.
2. The program cannot actually make use of the potential extra memory reported by `malloc_usable_size/malloc_size/_msize` without also calling `realloc` (or an allocator extension equivalent) to expand the allocation. Some allocators do not explicitly forbid it, but glibc's allocator and Address Sanitizer's allocator explicitly do. The call to `realloc` will also perform the same size lookup required by `malloc_usable_size/malloc_size/_msize`. Android's version of the Scudo allocator intentionally reports the requested size rather than the actual size from `malloc_usable_size`, so that it doesn't have to copy the full set of bytes into a new allocation when `realloc` is used.

Use Cases

Dynamic Array

A common data structure is a resizable array that grows on insertions by reallocating. To achieve amortized linear time, implementations allocate larger blocks of memory than they actually need to hold the current set of items - if they could use additional space at low/no cost, fewer resizes would be needed, minimizing copying and allocator churn.

Arena/Bump allocators

Repeated small allocations with the same lifetime can be grouped together by using `malloc` to obtain a large block of memory, and then adding the size of each small allocation to the returned pointer until it is full. If needed, allocate another block and repeat. However, chosen block sizes can waste substantial memory if the requested sizes are a poor fit for underlying size classes used by `malloc`.

I/O Buffers

When streaming data to/from a file, the network, or some other system, we often don't know how much data we'll move in advance. Larger buffers use more memory but often increase throughput and decrease syscall and device overhead; if we can obtain a larger buffer than required at no additional cost, we want to do so.

Workarounds

Various allocators provide facilities for querying the anticipated size of a future allocation so that a request can be made precisely for it, or querying the actual size (so the allocation can be realloc'd in place).

- Apple: `malloc_good_size+malloc, malloc_size`
- jemalloc/tcmalloc: `nallocx+malloc`
- dlmalloc: `realloc_in_place`
- mimalloc: `mi_expand`
- Microsoft: `_msize, _expand`

Suggested Wording

7.24.3 Memory management functions

The order and contiguity of storage allocated by successive calls to the `aligned_alloc`, `aligned_alloc_at_least`, `alloc_at_least`, `calloc`, `malloc`, and `realloc` functions is unspecified. The pointer returned if the allocation succeeds is suitably aligned so that it may be assigned to a pointer to any type of object with a fundamental alignment requirement and size less than or equal to the size requested, or in the case of `aligned_alloc_at_least` and `alloc_at_least`, the size returned. It may then be used to access such an object or an array of such objects in the space allocated (until the space is explicitly deallocated). [...]

The `alloc_result_t` type

Synopsis

```
C/C++  
#include <stdlib.h>  
typedef struct {
```

```
void* ptr;
size_t size;
} alloc_result_t
```

Description

`ptr` is `NULL` or the address of the resulting allocation.

`size` is the number of bytes available for use at `ptr`. If `ptr` is `NULL`, `size` must be 0.

The `alloc_at_least` function

Synopsis

C/C++

```
#include <stdlib.h>
alloc_result_t alloc_at_least(size_t min_size);
```

Description

NOTE: A conforming implementation may simply call `malloc`, place the resulting address in `alloc_result_t.ptr`, and update `alloc_result_t.size` to be 0 or `min_size` depending on whether the resulting address was `NULL` or not.

Recommended practice

An allocator should return an actual size larger than the requested minimum size if the additional memory could not be otherwise put to use until the returned pointer was freed.

NOTE: `alloc_at_least(n).ptr` is usable everywhere that the result of `malloc(n)` is, however using `alloc_at_least(n)` and ignoring the resulting size should be avoided. If the resulting size is not used, prefer `malloc`.

Returns

The `alloc_at_least` function returns the type `alloc_result_t` which holds `NULL` or a pointer to the allocated space. If the pointer is not `NULL`, the resulting size is guaranteed to be greater than or equal to the requested size.

The `aligned_alloc_at_least` function

Synopsis

C/C++

```
#include <stdlib.h>
alloc_result_t aligned_alloc_at_least(size_t alignment,
                                      size_t min_size);
```

Description

NOTE: A conforming implementation may simply call `aligned_alloc`, place the resulting address in `alloc_result_t.ptr`, and update `alloc_result_t.size` to be 0 or `min_size` depending on whether the resulting address was `NULL` or not.

Recommended practice

An allocator should return an actual size larger than the requested minimum size if the additional memory could not be otherwise put to use until the returned pointer was freed.

NOTE: `aligned_alloc_at_least(a, n).ptr` is usable everywhere that the result of `aligned_alloc(a, n)` is, however using `aligned_alloc_at_least(a, n)` and ignoring the resulting size should be avoided. If the resulting size is not used, prefer `aligned_alloc`.

Returns

The `aligned_alloc_at_least` function returns the type `alloc_result_t` which holds `NULL` or a pointer to the allocated space. If the pointer is not `NULL`, the pointer is aligned to at least the requested alignment and the resulting size is guaranteed to be greater than or equal to the requested size.

The `free_sized` function

Description

If `ptr` is a null pointer or the result obtained from a call to `malloc`, `realloc`, or `calloc`, where `size` is equal to the requested allocation size, this function is equivalent to `free(ptr)`. **If `ptr` was allocated by `alloc_at_least(n)` and `bytes >= size && size >= n`, where `bytes` is the size returned in `alloc_result_t`, this function is equivalent to `free(ptr)`.** Otherwise, the behavior is undefined. The result of an `aligned_alloc` or `aligned_alloc_at_least` call may not be passed to `free_sized`.

The `free_aligned_sized` function

Description

If `ptr` is a null pointer or the result obtained from a call to `aligned_alloc`, where `alignment` is equal to the requested allocation alignment and `size` is equal to the requested allocation size, this function is equivalent to `free(ptr)`. **If `ptr` was allocated by `aligned_alloc_at_least(alignment, n)` and `bytes >= size >= n`, where `bytes` is the size returned in `alloc_result_t`, this function is equivalent to `free(ptr)`.** Otherwise, the behavior is undefined. The result of an `malloc`, `calloc`, **or** `realloc`, or `alloc_at_least` call may not be passed to `free_aligned_sized`.

Design Discussion

Naming Bikeshed

The wording for this proposal uses `alloc_at_least`, but we have selected several alternatives for consideration:

- `alloc_at_least/aligned_alloc_at_least`
- `malloc_at_least/aligned_malloc_at_least`
- `sized_alloc/aligned_sized_alloc`
- `sized_malloc/aligned_sized_malloc`
- `size_returning_alloc/size_returning_aligned_alloc`
- `size_returning_malloc/sized_returning_aligned_malloc`
- `min_size_alloc/aligned_min_size_alloc`

Interaction with Sized Free

For allocations made with `alloc_at_least` and `aligned_alloc_at_least`, we need to relax `free_sized`'s and `free_aligned_sized`'s size argument (7.25.4.5). For allocations of `T`, the size quanta used by the allocator may not be a multiple of `sizeof(T)`, leading to both the original and returned sizes being unrecoverable at the time of deallocation.

Consider the memory allocated by:

```
C/C++  
typedef struct {  
    uint64_t data[2];  
} T;  
  
alloc_result_t r = alloc_at_least(sizeof(T) * 4);  
T* p = r.ptr;  
size_t s = r.ptr / sizeof(T);
```

```
C/C++  
typedef struct {  
    size_t count;  
    uint64_t data[];  
} T;  
  
alloc_result_t r = alloc_at_least(offsetof(T, data[4]));
```

- The memory allocator may return a 72 byte object. Since there is no `k` such that `sizeof(T) * k == 72`, we can't provide that value to `free_sized`. The only option would be storing 72 explicitly, which would be wasteful.
- The memory allocator may instead return an 80 byte object (5 `T`'s): We now cannot represent the original request when deallocating without additional storage.

For allocations made with

```
C/C++  
alloc_result_t r = alloc_at_least(n);
```

we permit `free_sized(r.ptr, s)` where `n <= s <= r.size`.

Where `s` must fall between the requested size `n` and the actual allocated size. This behavior is consistent with jemalloc's `sdallocx` and tcmalloc's size returning extensions. `free_sized` (N2801) recommends implementations accept sizes up to the actual size provided by extensions (i.e., `r.size` in this proposal) already.

Interaction with realloc

If a pointer obtained from `alloc_at_least` is passed to `realloc`, and `realloc` allocates a new block, the new block contains the contents of the previous block up to the returned size from `alloc_at_least`.

Why not realloc alone?

`realloc` must determine from the allocator's metadata the true size of the block. Even if paired with extensions like `malloc_usable_size` to resize to the precise, actual size, these pointer-to-size lookups are costly. Avoiding this lookup was a motivation behind C++14's sized delete and C23's `free_sized` features. When a program can make use of the added space, the best time to determine it is at allocation time when the allocator has all of the relevant metadata available.

While an allocator may obtain memory via facilities like `mmap` under the hood and leave an opportunity to grow an allocation arbitrarily, this is not always possible in practice.

- Some allocators (like TCMalloc) cache deallocated objects rather than having a 1-to-1 correspondence between them and VMAs.
- VMA limitations (for example Linux's `/proc/sys/vm/max_map_count`) may preclude having too many independent regions for allocations, each of which is arbitrarily growable. This motivates coalescing allocations onto fewer address regions, so another allocation may be directly "after" the allocation we wish to grow.

Where is `realloc_at_least`?

Allocators have two options for implementing `realloc` when programs attempt to grow an allocation: extending in-place or performing a new allocation, copying the memory from the old allocation to the new allocation, deallocating the old allocation, and returning a pointer to the new allocation. When the allocator hands back OS memory pages directly for an allocation, the allocator may be able to use platform-specific optimizations to expand or shrink an allocation such as `mremap(2)` on Linux. This is employed by glibc on Linux.

With `alloc_at_least`, the available memory was already reported to the program so it is unlikely the allocator would be able to extend the allocation in-place, leaving the allocator to copy which the program could do itself with `alloc_at_least+memcpy+free_sized`. Thus we have chosen to not introduce `realloc_at_least` as it provides no clear benefit.

Alignment of `alloc_at_least`

The standard requires that allocated memory be “suitably aligned so that it may be assigned to a pointer to any type of object with a fundamental alignment requirement and size less than or equal to the size requested.” To preserve the behavior that the result is usable anywhere the result of the same-sized malloc call is, we have to make the guaranteed alignment match the returned size, not the requested one.

Implementation Experience

TCMalloc has had a C++-style implementation of this since 2018. Since mid-2025, it now has an implementation of `alloc_at_least`.

Interaction with Bounds Checks

This proposal does not impede “hardened” allocators or sanitizers that wish to perform bounds checks, and can replace the use of `malloc_usable_size` and `_msize` which make bounds checking more difficult.

Avoiding Ossification (Hyrum’s Law)

Exposing the underlying size of allocations may result in code depending on the sizes returned in practice. GWP-ASan returns the requested size rather than the actual size in tcmalloc, to detect this possibly buggy code on a sampled basis, and other allocators wishing to be defensive against dependence on implementation details can apply the same strategy.