

**NAME**

`mremap` – remap a virtual memory address

**SYNOPSIS**

```
#define _GNU_SOURCE /* See feature_test_macros(7) */
#include <sys/mman.h>

void *mremap(void *old_address, size_t old_size,
             size_t new_size, int flags, ... /* void *new_address */);
```

**DESCRIPTION**

`mremap()` expands (or shrinks) an existing memory mapping, potentially moving it at the same time (controlled by the *flags* argument and the available virtual address space).

*old\_address* is the old address of the virtual memory block that you want to expand (or shrink). Note that *old\_address* has to be page aligned. *old\_size* is the old size of the virtual memory block. *new\_size* is the requested size of the virtual memory block after the resize. An optional fifth argument, *new\_address*, may be provided; see the description of **MREMAP\_FIXED** below.

If the value of *old\_size* is zero, and *old\_address* refers to a shareable mapping (see **mmap(2)** **MAP\_SHARED**), then `mremap()` will create a new mapping of the same pages. *new\_size* will be the size of the new mapping and the location of the new mapping may be specified with *new\_address*; see the description of **MREMAP\_FIXED** below. If a new mapping is requested via this method, then the **MREMAP\_MAYMOVE** flag must also be specified.

In Linux the memory is divided into pages. A user process has (one or) several linear virtual memory segments. Each virtual memory segment has one or more mappings to real memory pages (in the page table). Each virtual memory segment has its own protection (access rights), which may cause a segmentation violation if the memory is accessed incorrectly (e.g., writing to a read-only segment). Accessing virtual memory outside of the segments will also cause a segmentation violation.

`mremap()` uses the Linux page table scheme. `mremap()` changes the mapping between virtual addresses and memory pages. This can be used to implement a very efficient **realloc(3)**.

The *flags* bit-mask argument may be 0, or include the following flag:

**MREMAP\_MAYMOVE**

By default, if there is not sufficient space to expand a mapping at its current location, then `mremap()` fails. If this flag is specified, then the kernel is permitted to relocate the mapping to a new virtual address, if necessary. If the mapping is relocated, then absolute pointers into the old mapping location become invalid (offsets relative to the starting address of the mapping should be employed).

**MREMAP\_FIXED** (since Linux 2.3.31)

This flag serves a similar purpose to the **MAP\_FIXED** flag of **mmap(2)**. If this flag is specified, then `mremap()` accepts a fifth argument, *void \*new\_address*, which specifies a page-aligned address to which the mapping must be moved. Any previous mapping at the address range specified by *new\_address* and *new\_size* is unmapped. If **MREMAP\_FIXED** is specified, then **MREMAP\_MAYMOVE** must also be specified.

If the memory segment specified by *old\_address* and *old\_size* is locked (using **mlock(2)** or similar), then this lock is maintained when the segment is resized and/or relocated. As a consequence, the amount of memory locked by the process may change.

**RETURN VALUE**

On success `mremap()` returns a pointer to the new virtual memory area. On error, the value **MAP\_FAILED** (that is, *(void \*) -1*) is returned, and *errno* is set appropriately.

**ERRORS****EAGAIN**

The caller tried to expand a memory segment that is locked, but this was not possible without exceeding the **RLIMIT\_MEMLOCK** resource limit.

**EFAULT**

"Segmentation fault." Some address in the range *old\_address* to *old\_address+old\_size* is an invalid virtual memory address for this process. You can also get **EFAULT** even if there exist mappings that cover the whole address space requested, but those mappings are of different types.

**EINVAL**

An invalid argument was given. Possible causes are:

- \* *old\_address* was not page aligned;
- \* a value other than **MREMAP\_MAYMOVE** or **MREMAP\_FIXED** was specified in *flags*;
- \* *new\_size* was zero;
- \* *new\_size* or *new\_address* was invalid;
- \* the new address range specified by *new\_address* and *new\_size* overlapped the old address range specified by *old\_address* and *old\_size*;
- \* **MREMAP\_FIXED** was specified without also specifying **MREMAP\_MAYMOVE**;
- \* *old\_size* was zero and *old\_address* does not refer to a shareable mapping (but see **BUGS**);
- \* *old\_size* was zero and the **MREMAP\_MAYMOVE** flag was not specified.

**ENOMEM**

The memory area cannot be expanded at the current virtual address, and the **MREMAP\_MAYMOVE** flag is not set in *flags*. Or, there is not enough (virtual) memory available.

**CONFORMING TO**

This call is Linux-specific, and should not be used in programs intended to be portable.

**NOTES**

Prior to version 2.4, glibc did not expose the definition of **MREMAP\_FIXED**, and the prototype for **mremap()** did not allow for the *new\_address* argument.

If **mremap()** is used to move or expand an area locked with **mlock(2)** or equivalent, the **mremap()** call will make a best effort to populate the new area but will not fail with **ENOMEM** if the area cannot be populated.

**BUGS**

Before Linux 4.14, if *old\_size* was zero and the mapping referred to by *old\_address* was a private mapping (**mmap(2)** **MAP\_PRIVATE**), **mremap()** created a new private mapping unrelated to the original mapping. This behavior was unintended and probably unexpected in user-space applications (since the intention of **mremap()** is to create a new mapping based on the original mapping). Since Linux 4.14, **mremap()** fails with the error **EINVAL** in this scenario.

**SEE ALSO**

**brk(2)**, **getpagesize(2)**, **getrlimit(2)**, **mlock(2)**, **mmap(2)**, **sbrk(2)**, **malloc(3)**, **realloc(3)**

Your favorite text book on operating systems for more information on paged memory (e.g., *Modern Operating Systems* by Andrew S. Tanenbaum, *Inside Linux* by Randolph Bentson, *The Design of the UNIX Operating System* by Maurice J. Bach)

**COLOPHON**

This page is part of release 5.05 of the Linux *man-pages* project. A description of the project, information about reporting bugs, and the latest version of this page, can be found at <https://www.kernel.org/doc/man-pages/>.