

NAME

vm_map_sync - push dirty pages to their pager

SYNOPSIS

```
#include <sys/param.h>
#include <vm/vm.h>
#include <vm/vm_map.h>
```

int

```
vm_map_sync(vm_map_t map, vm_offset_t start, vm_offset_t end, boolean_t syncio,
             boolean_t invalidate);
```

DESCRIPTION

The **vm_map_sync()** function forces any dirty cached pages in the range *start* to *end* within the *map* to be pushed to their underlying pager.

If *syncio* is TRUE, dirty pages are written synchronously.

If *invalidate* is TRUE, any cached pages are also freed.

The range provided must be contiguous, it MUST NOT contain holes. The range provided MUST NOT contain any sub-map entries.

RETURN VALUES

The **vm_map_sync()** function returns KERN_SUCCESS if successful.

Otherwise, KERN_INVALID_ADDRESS will be returned if the function encountered a sub-map entry; KERN_INVALID_ARGUMENT will be returned if the function encountered a hole in the region provided, or if an entry could not be found for the given start address.

SEE ALSO

vm_map(9)

AUTHORS

This manual page was written by Bruce M Simpson <bms@spc.org>.