NAME

contigmalloc, **contigfree** - manage contiguous kernel physical memory

SYNOPSIS

#include <sys/types.h>
#include <sys/malloc.h>

*void **
contigmalloc(*unsigned long size, struct malloc_type *type, int flags, vm_paddr_t low, vm_paddr_t high, unsigned long alignment, vm_paddr_t boundary*);

void
contigfree(*void *addr, unsigned long size, struct malloc_type *type*);

#include <sys/param.h>
#include <sys/domainset.h>

*void **
contigmalloc_domainset(*unsigned long size, struct malloc_type *type, struct domainset *ds, int flags, vm_paddr_t low, vm_paddr_t high, unsigned long alignment, vm_paddr_t boundary*);

DESCRIPTION

The **contigmalloc**() function allocates *size* bytes of contiguous physical memory that is aligned to *alignment* bytes, and which does not cross a boundary of *boundary* bytes. If successful, the allocation will reside between physical addresses *low* and *high*. The returned pointer points to a wired kernel virtual address range of *size* bytes allocated from the kernel virtual address (KVA) map.

The **contigmalloc_domainset**() variant allows the caller to additionally specify a numa(4) domain selection policy. See domainset(9) for some example policies.

The *flags* parameter modifies **contigmalloc**()'s behaviour as follows:

M_ZERO
: Causes the allocated physical memory to be zero filled.

M_NOWAIT
: Causes **contigmalloc**() to return NULL if the request cannot be immediately fulfilled due to resource shortage.

Other flags (if present) are ignored.

The **contigfree**() function deallocates memory allocated by a previous call to **contigmalloc**() or **contigmalloc_domainset**().

IMPLEMENTATION NOTES

The **contigmalloc**() function does not sleep waiting for memory resources to be freed up, but instead actively reclaims pages before giving up. However, unless M_NOWAIT is specified, it may select a page for reclamation that must first be written to backing storage, causing it to sleep.

The **contigfree**() function does not accept NULL as an address input, unlike free(9).

RETURN VALUES

The **contigmalloc**() function returns a kernel virtual address if allocation succeeds, or NULL otherwise.

EXAMPLES

```
void *p;
p = contigmalloc(8192, M_DEVBUF, M_ZERO, 0, (1L << 22),
   32 * 1024, 1024 * 1024);
```

Ask for 8192 bytes of zero-filled memory residing between physical address 0 and 4194303 inclusive, aligned to a 32K boundary and not crossing a 1M address boundary.

DIAGNOSTICS

The **contigmalloc**() function will panic if *size* is zero, or if *alignment* or *boundary* is not a power of two.

SEE ALSO

malloc(9), memguard(9)