



---

## man Pages(9F) : DDI and DKI Kernel Functions

---

Sun Microsystems, Inc.  
901 San Antonio Road  
Palo Alto, CA 94303-4900  
U.S.A.

Part No: 805-3182-10  
October 1998

Copyright 1998 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, SunDocs, Java, the Java Coffee Cup logo, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

**RESTRICTED RIGHTS:** Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 1998 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, Californie 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, SunDocs, Java, le logo Java Coffee Cup, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



# Contents

---

**PREFACE xvii**

Intro(9F) 2

adjmsg(9F) 52

alloca(9F) 53

anocancel(9F) 56

aphysio(9F) 57

ASSERT(9F) 59

backq(9F) 60

bcanput(9F) 61

bcmp(9F) 62

bcopy(9F) 63

bioclone(9F) 65

biodone(9F) 68

bioerror(9F) 70

biofini(9F) 71

bioinit(9F) 72

biomodified(9F) 73

bioreset(9F) 74

biosize(9F) 75

biowait(9F) 76  
bp\_mapin(9F) 77  
bp\_mapout(9F) 78  
btop(9F) 79  
btopr(9F) 80  
bufcall(9F) 81  
bzero(9F) 84  
canput(9F) 85  
canputnext(9F) 86  
clrbuf(9F) 87  
cmn\_err(9F) 88  
condvar(9F) 94  
copyb(9F) 97  
copyin(9F) 99  
copymsg(9F) 101  
copyout(9F) 103  
csx\_AccessConfigurationRegister(9F) 105  
csx\_ConvertSize(9F) 107  
csx\_ConvertSpeed(9F) 109  
csx\_CS\_DDI\_Info(9F) 111  
csx\_DeregisterClient(9F) 113  
csx\_DupHandle(9F) 114  
csx\_Error2Text(9F) 117  
csx\_Event2Text(9F) 119  
csx\_FreeHandle(9F) 121  
csx\_Get8(9F) 122  
csx\_GetFirstClient(9F) 123  
csx\_GetFirstTuple(9F) 125

csx_GetHandleOffset(9F)	127
csx_GetMappedAddr(9F)	128
csx_GetStatus(9F)	129
csx_GetTupleData(9F)	134
csx_MakeDeviceNode(9F)	137
csx_MapLogSocket(9F)	140
csx_MapMemPage(9F)	142
csx_ModifyConfiguration(9F)	144
csx_ModifyWindow(9F)	147
csx_Parse_CISTPL_BATTERY(9F)	150
csx_Parse_CISTPL_BYTEORDER(9F)	152
csx_Parse_CISTPL_CFTABLE_ENTRY(9F)	154
csx_Parse_CISTPL_CONFIG(9F)	162
csx_Parse_CISTPL_DATE(9F)	165
csx_Parse_CISTPL_DEVICE(9F)	167
csx_Parse_CISTPL_DEVICEGEO(9F)	171
csx_Parse_CISTPL_DEVICEGEO_A(9F)	173
csx_Parse_CISTPL_FORMAT(9F)	175
csx_Parse_CISTPL_FUNCE(9F)	178
csx_Parse_CISTPL_FUNCID(9F)	188
csx_Parse_CISTPL_GEOMETRY(9F)	191
csx_Parse_CISTPL_JEDEC_C(9F)	193
csx_Parse_CISTPL_LINKTARGET(9F)	195
csx_Parse_CISTPL_LONGLINK_A(9F)	197
csx_Parse_CISTPL_LONGLINK_MFC(9F)	199
csx_Parse_CISTPL_MANFID(9F)	201
csx_Parse_CISTPL_ORG(9F)	203
csx_Parse_CISTPL_SPCL(9F)	205

csx\_Parse\_CISTPL\_SWIL(9F) 207  
csx\_Parse\_CISTPL\_VERS\_1(9F) 209  
csx\_Parse\_CISTPL\_VERS\_2(9F) 211  
csx\_ParseTuple(9F) 213  
csx\_Put8(9F) 215  
csx\_RegisterClient(9F) 216  
csx\_ReleaseConfiguration(9F) 220  
csx\_RepGet8(9F) 222  
csx\_RepPut8(9F) 224  
csx\_RequestConfiguration(9F) 226  
csx\_RequestIO(9F) 231  
csx\_RequestIRQ(9F) 237  
csx\_RequestSocketMask(9F) 240  
csx\_RequestWindow(9F) 242  
csx\_ResetFunction(9F) 247  
csx\_SetEventMask(9F) 249  
csx\_SetHandleOffset(9F) 251  
csx\_ValidateCIS(9F) 252  
datamsg(9F) 254  
ddi\_add\_intr(9F) 256  
ddi\_add\_softintr(9F) 259  
ddi\_binding\_name(9F) 266  
ddi\_btop(9F) 267  
ddi\_copyin(9F) 268  
ddi\_copyout(9F) 271  
ddi\_create\_minor\_node(9F) 274  
ddi\_device\_copy(9F) 276  
ddi\_device\_zero(9F) 278

ddi_devid_compare(9F)	280
ddi_dev_is_needed(9F)	283
ddi_dev_is_sid(9F)	285
ddi_dev_nintrs(9F)	287
ddi_dev_nregs(9F)	288
ddi_dev_regsize(9F)	289
ddi_dma_addr_bind_handle(9F)	290
ddi_dma_addr_setup(9F)	294
ddi_dma_alloc_handle(9F)	296
ddi_dma_buf_bind_handle(9F)	298
ddi_dma_buf_setup(9F)	302
ddi_dma_burstsizes(9F)	304
ddi_dma_coff(9F)	305
ddi_dma_curwin(9F)	306
ddi_dma_devalign(9F)	307
ddi_dmae(9F)	308
ddi_dmae_alloc(9F)	308
ddi_dma_free(9F)	313
ddi_dma_free_handle(9F)	314
ddi_dma_getwin(9F)	315
ddi_dma_htoc(9F)	317
ddi_dma_mem_alloc(9F)	318
ddi_dma_mem_free(9F)	321
ddi_dma_movwin(9F)	322
ddi_dma_nextcookie(9F)	324
ddi_dma_nextseg(9F)	326
ddi_dma_nextwin(9F)	328
ddi_dma_numwin(9F)	330

ddi\_dma\_segtocookie(9F) 331  
ddi\_dma\_set\_sbus64(9F) 333  
ddi\_dma\_setup(9F) 335  
ddi\_dma\_sync(9F) 337  
ddi\_dma\_unbind\_handle(9F) 339  
ddi\_enter\_critical(9F) 340  
ddi\_ffs(9F) 341  
ddi\_get8(9F) 342  
ddi\_get\_cred(9F) 344  
ddi\_get\_driver\_private(9F) 345  
ddi\_get\_instance(9F) 346  
ddi\_get\_lbolt(9F) 347  
ddi\_get\_parent(9F) 348  
ddi\_get\_pid(9F) 349  
ddi\_get\_time(9F) 350  
ddi\_in\_panic(9F) 351  
ddi\_intr\_hilevel(9F) 352  
ddi\_io\_get8(9F) 353  
ddi\_iomin(9F) 355  
ddi\_iopb\_alloc(9F) 356  
ddi\_io\_put8(9F) 358  
ddi\_io\_rep\_get8(9F) 360  
ddi\_io\_rep\_put8(9F) 362  
ddi\_mapdev(9F) 364  
ddi\_mapdev\_intercept(9F) 367  
ddi\_mapdev\_set\_device\_acc\_attr(9F) 369  
ddi\_map\_regs(9F) 371  
ddi\_mem\_alloc(9F) 373

ddi\_mem\_get8(9F) 375  
ddi\_mem\_put8(9F) 377  
ddi\_mem\_rep\_get8(9F) 379  
ddi\_mem\_rep\_put8(9F) 381  
ddi\_mmap\_get\_model(9F) 383  
ddi\_model\_convert\_from(9F) 385  
ddi\_node\_name(9F) 387  
ddi\_peek(9F) 388  
ddi\_poke(9F) 391  
ddi\_prop\_create(9F) 393  
ddi\_prop\_exists(9F) 398  
ddi\_prop\_get\_int(9F) 400  
ddi\_prop\_lookup(9F) 402  
ddi\_prop\_op(9F) 407  
ddi\_prop\_update(9F) 411  
ddi\_put8(9F) 415  
ddi\_regs\_map\_free(9F) 417  
ddi\_regs\_map\_setup(9F) 418  
ddi\_remove\_minor\_node(9F) 420  
ddi\_rep\_get8(9F) 421  
ddi\_report\_dev(9F) 423  
ddi\_rep\_put8(9F) 424  
ddi\_root\_node(9F) 426  
ddi\_segmap(9F) 427  
ddi\_slaveonly(9F) 430  
ddi\_soft\_state(9F) 431  
ddi\_umem\_alloc(9F) 436  
delay(9F) 438

devmap\_default\_access(9F) 440  
devmap\_devmem\_setup(9F) 443  
devmap\_do\_ctxmgt(9F) 447  
devmap\_set\_ctx\_timeout(9F) 450  
devmap\_setup(9F) 451  
devmap\_unload(9F) 454  
disksort(9F) 457  
drv\_getparm(9F) 459  
drv\_hztousec(9F) 461  
drv\_priv(9F) 462  
drv\_usectohz(9F) 463  
drv\_usecwait(9F) 464  
dupb(9F) 465  
dupmsg(9F) 468  
enableok(9F) 469  
esballoc(9F) 470  
esbcall(9F) 472  
flushband(9F) 473  
flushq(9F) 474  
freeb(9F) 476  
freemsg(9F) 477  
freerbuf(9F) 478  
freezestr(9F) 479  
geterror(9F) 480  
getmajor(9F) 481  
getminor(9F) 482  
get\_pktiopb(9F) 483  
getq(9F) 485

getrbuf(9F) 486  
hat\_getkpfnum(9F) 487  
inb(9F) 488  
insq(9F) 490  
IOC\_CONVERT\_FROM(9F) 492  
kmem\_alloc(9F) 493  
kstat\_create(9F) 495  
kstat\_delete(9F) 498  
kstat\_install(9F) 499  
kstat\_named\_init(9F) 500  
kstat\_queue(9F) 501  
linkb(9F) 503  
makecom(9F) 504  
makedevice(9F) 506  
max(9F) 507  
min(9F) 508  
mkiocb(9F) 509  
mod\_install(9F) 512  
msgdsize(9F) 513  
msgpullup(9F) 514  
mt-streams(9F) 515  
mutex(9F) 517  
nochpoll(9F) 520  
nodev(9F) 521  
noenable(9F) 522  
nulldev(9F) 523  
OTHERQ(9F) 524  
outb(9F) 526

pci\_config\_get8(9F) 528  
pci\_config\_setup(9F) 530  
physio(9F) 531  
pm\_busy\_component(9F) 533  
pm\_create\_components(9F) 535  
pm\_get\_normal\_power(9F) 536  
pollwakeup(9F) 538  
proc\_signal(9F) 539  
ptob(9F) 541  
pullupmsg(9F) 542  
put(9F) 544  
putbq(9F) 545  
putctl1(9F) 546  
putctl(9F) 547  
putnext(9F) 549  
putnextctl1(9F) 550  
putnextctl(9F) 551  
putq(9F) 553  
qbufcall(9F) 554  
qenable(9F) 556  
qprocson(9F) 557  
qreply(9F) 558  
qsize(9F) 560  
qtimeout(9F) 561  
qunbufcall(9F) 563  
quntimeout(9F) 564  
qwait(9F) 565  
qwriter(9F) 567

RD(9F) 569  
rmalloc(9F) 570  
rmallocmap(9F) 573  
rmalloc\_wait(9F) 575  
rmfree(9F) 576  
rmvb(9F) 577  
rmvq(9F) 579  
rwlock(9F) 581  
SAMESTR(9F) 584  
scsi\_abort(9F) 585  
scsi\_alloc\_consistent\_buf(9F) 586  
scsi\_cname(9F) 588  
scsi\_destroy\_pkt(9F) 590  
scsi\_dmaget(9F) 591  
scsi\_errmsg(9F) 593  
scsi\_free\_consistent\_buf(9F) 596  
scsi\_hba\_attach\_setup(9F) 597  
scsi\_hba\_init(9F) 600  
scsi\_hba\_lookup\_capstr(9F) 601  
scsi\_hba\_pkt\_alloc(9F) 603  
scsi\_hba\_probe(9F) 605  
scsi\_hba\_tran\_alloc(9F) 606  
scsi\_ifgetcap(9F) 607  
scsi\_init\_pkt(9F) 611  
scsi\_log(9F) 615  
scsi\_pktalloc(9F) 617  
scsi\_poll(9F) 619  
scsi\_probe(9F) 620

scsi\_reset(9F) 622  
scsi\_reset\_notify(9F) 623  
scsi\_setup\_cdb(9F) 625  
scsi\_slave(9F) 626  
scsi\_sync\_pkt(9F) 628  
scsi\_transport(9F) 629  
scsi\_unprobe(9F) 631  
scsi\_vu\_errmsg(9F) 632  
semaphore(9F) 635  
sprintf(9F) 637  
stoi(9F) 639  
strchr(9F) 640  
strcmp(9F) 641  
strcpy(9F) 642  
strlen(9F) 643  
strlog(9F) 644  
strqget(9F) 646  
strqset(9F) 648  
STRUCT\_DECL(9F) 649  
swab(9F) 655  
testb(9F) 656  
timeout(9F) 658  
uiomove(9F) 660  
unbufcall(9F) 662  
unlinkb(9F) 663  
untimeout(9F) 664  
ureadc(9F) 666  
uwritec(9F) 667

va\_arg(9F) 668

vsprintf(9F) 671

WR(9F) 674

**Index 676**



# PREFACE

---

## Overview

A man page is provided for both the naive user, and sophisticated user who is familiar with the SunOS operating system and is in need of on-line information. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

The following contains a brief description of each section in the man pages and the information it references:

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2 of this volume.
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
- Section 5 contains miscellaneous documentation such as character set tables.
- Section 6 contains available games and demos.

- Section 7 describes various special files that refer to specific hardware peripherals, and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.
- Section 9 provides reference information needed to write device drivers in the kernel operating systems environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver/Kernel Interface (DKI).
- Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer may include in a device driver.
- Section 9F describes the kernel functions available for use by device drivers.
- Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the `intro` pages for more information and detail about each section, and `man(1)` for more information about man pages in general.

**NAME** This section gives the names of the commands or functions documented, followed by a brief description of what they do.

**SYNOPSIS** This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full pathname is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.

The following special characters are used in this section:

- [ ] The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.
- . . . Ellipses. Several values may be provided for the previous argument, or the previous argument can be specified multiple times, for example, `'filename...'`.

| Separator. Only one of the arguments separated by this character can be specified at time.

{ } Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.

**PROTOCOL**

This section occurs only in subsection 3R to indicate the protocol description file.

**DESCRIPTION**

This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES.. Interactive commands, subcommands, requests, macros, functions and such, are described under USAGE.

**IOCTL**

This section appears on pages in Section 7 only. Only the device class which supplies appropriate parameters to the ioctl (2) system call is called `ioctl` and generates its own heading. `ioctl` calls for a specific device are listed alphabetically (on the man page for that specific device). `ioctl` calls are used for a particular class of devices all of which have an `io` ending, such as `mtio(7D)`

**OPTIONS**

This lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.

**OPERANDS**

This section lists the command operands and describes how they affect the actions of the command.

**OUTPUT**

This section describes the output - standard output, standard error, or output files - generated by the command.

**RETURN VALUES**

If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in

tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.

## **ERRORS**

On failure, most functions place an error code in the global variable `errno` indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.

## **USAGE**

This section is provided as a guidance on use. This section lists special rules, features and commands that require in-depth explanations. The subsections listed below are used to explain built-in functionality:

- Commands
- Modifiers
- Variables
- Expressions
- Input Grammar

## **EXAMPLES**

This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command line entry and machine response is shown. Whenever an example is given, the prompt is shown as `example%` or if the user must be superuser, `example#`. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS and USAGE sections.

## **ENVIRONMENT VARIABLES**

This section lists any environment variables that the command or function affects, followed by a brief description of the effect.

## **EXIT STATUS**

This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion and

values other than zero for various error conditions.

**FILES**

This section lists all filenames referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.

**ATTRIBUTES**

This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See `attributes(5)` for more information.

**SEE ALSO**

This section lists references to other man pages, in-house documentation and outside publications.

**DIAGNOSTICS**

This section lists diagnostic messages with a brief explanation of the condition causing the error.

**WARNINGS**

This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.

**NOTES**

This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.

**BUGS**

This section describes known bugs and wherever possible, suggests workarounds.

CHAPTER

---

# Kernel Functions for Drivers

<b>NAME</b>	Intro – introduction to DDI/DKI functions
<b>DESCRIPTION</b>	<p>Section 9F describes the kernel functions available for use by device drivers.</p> <p>In this section, the information for each driver function is organized under the following headings:</p> <ul style="list-style-type: none"> <li>■ <b>NAME</b> summarizes the function’s purpose.</li> <li>■ <b>SYNOPSIS</b> shows the syntax of the function’s entry point in the source code. <code>#include</code> directives are shown for required headers.</li> <li>■ <b>INTERFACE LEVEL</b> describes any architecture dependencies.</li> <li>■ <b>ARGUMENTS</b> describes any arguments required to invoke the function.</li> <li>■ <b>DESCRIPTION</b> describes general information about the function.</li> <li>■ <b>RETURN VALUES</b> describes the return values and messages that can result from invoking the function.</li> <li>■ <b>CONTEXT</b> indicates from which driver context (user, kernel, interrupt, or high-level interrupt) the function can be called.</li> <li>■ A driver function has <i>user context</i> if it was directly invoked because of a user thread. The <code>read(9E)</code> entry point of the driver, invoked by a <code>read(2)</code> system call, has user context.</li> <li>■ A driver function has <i>kernel context</i> if was invoked by some other part of the kernel. In a block device driver, the <code>strategy(9E)</code> entry point may be called by the page daemon to write pages to the device. The page daemon has no relation to the current user thread, so in this case <code>strategy(9E)</code> has kernel context.</li> <li>■ <i>Interrupt context</i> is kernel context, but also has an interrupt level associated with it. Driver interrupt routines have interrupt context.</li> <li>■ <i>High-level interrupt context</i> is a more restricted form of interrupt context. If <code>ddi_intr_hilevel(9F)</code> indicates that an interrupt is high-level, driver interrupt routines added for that interrupt with <code>ddi_add_intr(9F)</code> run in high-level interrupt context. These interrupt routines are only allowed to call <code>ddi_trigger_softintr(9F)</code>, <code>mutex_enter(9F)</code> and <code>mutex_exit(9F)</code>. Furthermore, <code>mutex_enter(9F)</code> and <code>mutex_exit(9F)</code> may only be called on mutexes initialized with the <code>ddi_iblock_cookie</code> returned by <code>ddi_get_iblock_cookie(9F)</code>.</li> <li>■ <b>SEE ALSO</b> indicates functions that are related by usage and sources, and which can be referred to for further information.</li> <li>■ <b>EXAMPLES</b> shows how the function can be used in driver code.</li> </ul>

**STREAMS Kernel  
Function Summary**

Every driver MUST include `<sys/ddi.h>` and `<sys/sunddi.h>`, in that order, and as the last files the driver includes.

The following table summarizes the STREAMS functions described in this section.

Routine	Type
adjmsg	DDI/DKI
allocb	DDI/DKI
backq	DDI/DKI
bcanput	DDI/DKI
bcanputnext	DDI/DKI
bufcall	DDI/DKI
canput	DDI/DKI
canputnext	DDI/DKI
clrbuf	DDI/DKI
copyb	DDI/DKI
copymsg	DDI/DKI
datamsg	DDI/DKI
dupb	DDI/DKI
dupmsg	DDI/DKI
enableok	DDI/DKI
esballoc	DDI/DKI
esbcall	DDI/DKI
flushband	DDI/DKI
flushq	DDI/DKI
freeb	DDI/DKI
freemsg	DDI/DKI
freezestr	DDI/DKI
getq	DDI/DKI
insq	DDI/DKI
linkb	DDI/DKI
msgdsiz	DDI/DKI

Routine	Type
msgpullup	DDI/DKI
mt-streams	Solaris DDI
noenable	DDI/DKI
OTHERQ	DDI/DKI
pullupmsg	DDI/DKI
put	DDI/DKI
putbq	DDI/DKI
putctl	DDI/DKI
putctl1	DDI/DKI
putnext	DDI/DKI
putnextctl	DDI/DKI
putq	DDI/DKI
qbufcall	Solaris DDI
qenable	DDI/DKI
qprocson	DDI/DKI
qprocsoff	DDI/DKI
qreply	DDI/DKI
qsize	DDI/DKI
qtimeout	Solaris DDI
qunbufcall	Solaris DDI
quntimeout	Solaris DDI
qwait	Solaris DDI
qwait_sig	Solaris DDI
qwriter	Solaris DDI
RD	DDI/DKI
rmvb	DDI/DKI
rmvq	DDI/DKI
SAMESTR	DDI/DKI
strlog	DDI/DKI
strqget	DDI/DKI

Routine	Type
strqset	DDI/DKI
testb	DDI/DKI
unbufcall	DDI/DKI
unfreezestr	DDI/DKI
unlinkb	DDI/DKI
WR	DDI/DKI

The following table summarizes the functions not specific to STREAMS.

Routine	Type
ASSERT	DDI/DKI
anocancel	Solaris DDI
aphysio	Solaris DDI
bcmp	DDI/DKI
bcopy	DDI/DKI
biodone	DDI/DKI
bioclone	Solaris DDI
biofini	Solaris DDI
bioinit	Solaris DDI
biomodified	Solaris DDI
biosize	Solaris DDI
bioerror	Solaris DDI
bioreset	Solaris DDI
biowait	DDI/DKI
bp_mapin	DDI/DKI
bp_mapout	DDI/DKI
btop	DDI/DKI
btopr	DDI/DKI
bzero	DDI/DKI
cmn_err	DDI/DKI

Routine	Type
copyin	DDI/DKI
copyout	DDI/DKI
cv_broadcast	Solaris DDI
cv_destroy	Solaris DDI
cv_init	Solaris DDI
cv_signal	Solaris DDI
cv_timedwait	Solaris DDI
cv_wait	Solaris DDI
cv_wait_sig	Solaris DDI
ddi_add_intr	Solaris DDI
ddi_add_softintr	Solaris DDI
ddi_btop	Solaris DDI
ddi_btopr	Solaris DDI
ddi_copyin	Solaris DDI
ddi_copyout	Solaris DDI
ddi_create_minor_node	Solaris DDI
ddi_dev_is_sid	Solaris DDI
ddi_dev_nintrs	Solaris DDI
ddi_dev_nregs	Solaris DDI
ddi_dev_regsize	Solaris DDI
ddi_device_copy	Solaris DDI
ddi_device_zero	Solaris DDI
ddi_devmap_segmap	Solaris DDI
ddi_dma_addr_bind_handle	Solaris DDI
ddi_dma_addr_setup	Solaris DDI
ddi_dma_alloc_handle	Solaris DDI
ddi_dma_buf_bind_handle	Solaris DDI
ddi_dma_buf_setup	Solaris DDI
ddi_dma_burstsizes	Solaris DDI
ddi_dma_coff	Solaris SPARC DDI

Routine	Type
ddi_dma_curwin	Solaris SPARC DDI
ddi_dma_dealign	Solaris DDI
ddi_dma_free	Solaris DDI
ddi_dma_free_handle	Solaris DDI
ddi_dma_getwin	Solaris DDI
ddi_dma_htoc	Solaris SPARC DDI
ddi_dma_mem_alloc	Solaris DDI
ddi_dma_mem_free	Solaris DDI
ddi_dma_movwin	Solaris SPARC DDI
ddi_dma_nextcookie	Solaris DDI
ddi_dma_nextseg	Solaris DDI
ddi_dma_nextwin	Solaris DDI
ddi_dma_numwin	Solaris DDI
ddi_dma_segtocookie	Solaris DDI
ddi_dma_set_sbus64	Solaris DDI
ddi_dma_setup	Solaris DDI
ddi_dma_sync	Solaris DDI
ddi_dma_unbind_handle	Solaris DDI
ddi_dmae	Solaris x86 DDI
ddi_dmae_1stparty	Solaris x86 DDI
ddi_dmae_alloc	Solaris x86 DDI
ddi_dmae_disable	Solaris x86 DDI
ddi_dmae_enable	Solaris x86 DDI
ddi_dmae_getattr	Solaris x86 DDI
ddi_dmae_getcnt	Solaris x86 DDI
ddi_dmae_getlim	Solaris x86 DDI
ddi_dmae_prog	Solaris x86 DDI
ddi_dmae_release	Solaris x86 DDI
ddi_dmae_stop	Solaris x86 DDI
ddi_enter_critical	Solaris DDI

Routine	Type
ddi_exit_critical	Solaris DDI
ddi_ffs	Solaris DDI
ddi_fls	Solaris DDI
ddi_get16	Solaris DDI
ddi_get32	Solaris DDI
ddi_get64	Solaris DDI
ddi_get8	Solaris DDI
ddi_get_cred	Solaris DDI
ddi_get_driver_private	Solaris DDI
ddi_get_iblock_cookie	Solaris DDI
ddi_get_instance	Solaris DDI
ddi_get_name	Solaris DDI
ddi_get_parent	Solaris DDI
ddi_get_soft_iblock_cookie	Solaris DDI
ddi_get_soft_state	Solaris DDI
ddi_getb	Solaris DDI
ddi_getl	Solaris DDI
ddi_getll	Solaris DDI
ddi_getlongprop	Solaris DDI
ddi_getlongprop_buf	Solaris DDI
ddi_getprop	Solaris DDI
ddi_getproplen	Solaris DDI
ddi_getw	Solaris DDI
ddi_intr_hilevel	Solaris DDI
ddi_io_get16	Solaris DDI
ddi_io_get32	Solaris DDI
ddi_io_get8	Solaris DDI
ddi_io_getb	Solaris DDI
ddi_io_getl	Solaris DDI
ddi_io_getw	Solaris DDI

Routine	Type
ddi_io_put16	Solaris DDI
ddi_io_put32	Solaris DDI
ddi_io_put8	Solaris DDI
ddi_io_putb	Solaris DDI
ddi_io_putl	Solaris DDI
ddi_io_putw	Solaris DDI
ddi_io_rep_get16	Solaris DDI
ddi_io_rep_get32	Solaris DDI
ddi_io_rep_get8	Solaris DDI
ddi_io_rep_getb	Solaris DDI
ddi_io_rep_getl	Solaris DDI
ddi_io_rep_getw	Solaris DDI
ddi_io_rep_put16	Solaris DDI
ddi_io_rep_put32	Solaris DDI
ddi_io_rep_put8	Solaris DDI
ddi_io_rep_putb	Solaris DDI
ddi_io_rep_putl	Solaris DDI
ddi_io_rep_putw	Solaris DDI
ddi_iomin	Solaris DDI
ddi_iopb_alloc	Solaris DDI
ddi_iopb_free	Solaris DDI
ddi_map_regs	Solaris DDI
ddi_mapdev	Solaris DDI
ddi_mapdev_intercept	Solaris DDI
ddi_mapdev_nointercept	Solaris DDI
ddi_mapdev_set_device_acc_attr	Solaris DDI
ddi_mem_alloc	Solaris DDI
ddi_mem_free	Solaris DDI
ddi_mem_get16	Solaris DDI
ddi_mem_get32	Solaris DDI

Routine	Type
ddi_mem_get64	Solaris DDI
ddi_mem_get8	Solaris DDI
ddi_mem_getb	Solaris DDI
ddi_mem_getl	Solaris DDI
ddi_mem_getll	Solaris DDI
ddi_mem_getw	Solaris DDI
ddi_mem_put16	Solaris DDI
ddi_mem_put32	Solaris DDI
ddi_mem_put64	Solaris DDI
ddi_mem_put8	Solaris DDI
ddi_mem_putb	Solaris DDI
ddi_mem_putl	Solaris DDI
ddi_mem_putll	Solaris DDI
ddi_mem_putw	Solaris DDI
ddi_mem_rep_get16	Solaris DDI
ddi_mem_rep_get32	Solaris DDI
ddi_mem_rep_get64	Solaris DDI
ddi_mem_rep_get8	Solaris DDI
ddi_mem_rep_getb	Solaris DDI
ddi_mem_rep_getl	Solaris DDI
ddi_mem_rep_getll	Solaris DDI
ddi_mem_rep_getw	Solaris DDI
ddi_mem_rep_put16	Solaris DDI
ddi_mem_rep_put32	Solaris DDI
ddi_mem_rep_put64	Solaris DDI
ddi_mem_rep_put8	Solaris DDI
ddi_mem_rep_putb	Solaris DDI
ddi_mem_rep_putl	Solaris DDI
ddi_mem_rep_putll	Solaris DDI
ddi_mem_rep_putw	Solaris DDI

Routine	Type
ddi_mmap_get_model	Solaris DDI
ddi_model_convert_from	Solaris DDI
ddi_node_name	Solaris DDI
ddi_peek16	Solaris DDI
ddi_peek32	Solaris DDI
ddi_peek64	Solaris DDI
ddi_peek8	Solaris DDI
ddi_peekc	Solaris DDI
ddi_peekd	Solaris DDI
ddi_peekl	Solaris DDI
ddi_peeks	Solaris DDI
ddi_poke16	Solaris DDI
ddi_poke32	Solaris DDI
ddi_poke64	Solaris DDI
ddi_poke8	Solaris DDI
ddi_pokec	Solaris DDI
ddi_poked	Solaris DDI
ddi_pokel	Solaris DDI
ddi_pokes	Solaris DDI
ddi_prop_create	Solaris DDI
ddi_prop_exists	Solaris DDI
ddi_prop_free	Solaris DDI
ddi_prop_get_int	Solaris DDI
ddi_prop_lookup	Solaris DDI
ddi_prop_lookup_byte_array	Solaris DDI
ddi_prop_lookup_int_array	Solaris DDI
ddi_prop_lookup_string	Solaris DDI
ddi_prop_lookup_string_array	Solaris DDI
ddi_prop_modify	Solaris DDI
ddi_prop_op	Solaris DDI

Routine	Type
ddi_prop_remove	Solaris DDI
ddi_prop_remove_all	Solaris DDI
ddi_prop_undefine	Solaris DDI
ddi_prop_update	Solaris DDI
ddi_prop_update_byte_array	Solaris DDI
ddi_prop_update_int	Solaris DDI
ddi_prop_update_int_array	Solaris DDI
ddi_prop_update_string	Solaris DDI
ddi_prop_update_string_array	Solaris DDI
ddi_ptob	Solaris DDI
ddi_put16	Solaris DDI
ddi_put32	Solaris DDI
ddi_put64	Solaris DDI
ddi_put8	Solaris DDI
ddi_putb	Solaris DDI
ddi_putl	Solaris DDI
ddi_putll	Solaris DDI
ddi_putw	Solaris DDI
ddi_regs_map_free	Solaris DDI
ddi_regs_map_setup	Solaris DDI
ddi_remove_intr	Solaris DDI
ddi_remove_minor_node	Solaris DDI
ddi_remove_softintr	Solaris DDI
ddi_rep_get16	Solaris DDI
ddi_rep_get32	Solaris DDI
ddi_rep_get64	Solaris DDI
ddi_rep_get8	Solaris DDI
ddi_rep_getb	Solaris DDI
ddi_rep_getl	Solaris DDI
ddi_rep_getll	Solaris DDI

Routine	Type
ddi_rep_getw	Solaris DDI
ddi_rep_put16	Solaris DDI
ddi_rep_put32	Solaris DDI
ddi_rep_put64	Solaris DDI
ddi_rep_put8	Solaris DDI
ddi_rep_putb	Solaris DDI
ddi_rep_putl	Solaris DDI
ddi_rep_putll	Solaris DDI
ddi_rep_putw	Solaris DDI
ddi_report_dev	Solaris DDI
ddi_root_node	Solaris DDI
ddi_segmap	Solaris DDI
ddi_segmap_setup	Solaris DDI
ddi_set_driver_private	Solaris DDI
ddi_slaveonly	Solaris DDI
ddi_soft_state	Solaris DDI
ddi_soft_state_fini	Solaris DDI
ddi_soft_state_free	Solaris DDI
ddi_soft_state_init	Solaris DDI
ddi_soft_state_zalloc	Solaris DDI
ddi_trigger_softintr	Solaris DDI
ddi_umem_alloc	Solaris DDI
ddi_umem_free	Solaris DDI
ddi_unmap_regs	Solaris DDI
delay	DDI/DKI
devmap_default_access	Solaris DDI
devmap_devmem_setup	Solaris DDI
devmap_do_ctxmgt	Solaris DDI
devmap_load	Solaris DDI
devmap_set_ctx_timeout	Solaris DDI

Routine	Type
devmap_setup	Solaris DDI
devmap_umem_setup	Solaris DDI
devmap_unload	Solaris DDI
disksort	Solaris DDI
drv_getparm	DDI/DKI
drv_hztousec	DDI/DKI
drv_priv	DDI/DKI
drv_usectohz	DDI/DKI
drv_usecwait	DDI/DKI
free_pktiopb	Solaris DDI
freerbuf	DDI/DKI
get_pktiopb	Solaris DDI
geterror	DDI/DKI
getmajor	DDI/DKI
getminor	DDI/DKI
getrbuf	DDI/DKI
hat_getkpfnum	DKI only
inb	Solaris x86 DDI
inl	Solaris x86 DDI
inw	Solaris x86 DDI
kmem_alloc	DDI/DKI
kmem_free	DDI/DKI
kmem_zalloc	DDI/DKI
kstat_create	Solaris DDI
kstat_delete	Solaris DDI
kstat_install	Solaris DDI
kstat_named_init	Solaris DDI
kstat_queue	Solaris DDI
kstat_runq_back_to_waitq	Solaris DDI
kstat_runq_enter	Solaris DDI

Routine	Type
kstat_runq_exit	Solaris DDI
kstat_waitq_enter	Solaris DDI
kstat_waitq_exit	Solaris DDI
kstat_waitq_to_runq	Solaris DDI
makecom_g0	Solaris DDI
makecom_g0_s	Solaris DDI
makecom_g1	Solaris DDI
makecom_g5	Solaris DDI
makedevice	DDI/DKI
max	DDI/DKI
min	DDI/DKI
minphys	Solaris DDI
mod_info	Solaris DDI
mod_install	Solaris DDI
mod_remove	Solaris DDI
mutex_destroy	Solaris DDI
mutex_enter	Solaris DDI
mutex_exit	Solaris DDI
mutex_init	Solaris DDI
mutex_owned	Solaris DDI
mutex_tryenter	Solaris DDI
nochpoll	Solaris DDI
nodev	DDI/DKI
nulldev	DDI/DKI
numtos	Solaris DDI
outb	Solaris x86 DDI
outl	Solaris x86 DDI
outw	Solaris x86 DDI
pci_config_get16	Solaris DDI
pci_config_get32	Solaris DDI

Routine	Type
pci_config_get64	Solaris DDI
pci_config_get8	Solaris DDI
pci_config_getb	Solaris DDI
pci_config_getl	Solaris DDI
pci_config_getw	Solaris DDI
pci_config_put16	Solaris DDI
pci_config_put32	Solaris DDI
pci_config_put64	Solaris DDI
pci_config_put8	Solaris DDI
pci_config_putb	Solaris DDI
pci_config_putl	Solaris DDI
pci_config_putw	Solaris DDI
pci_config_setup	Solaris DDI
pci_config_teardown	Solaris DDI
physio	Solaris DDI
pollwakeup	DDI/DKI
proc_ref	Solaris DDI
proc_signal	Solaris DDI
proc_unref	Solaris DDI
ptob	DDI/DKI
repinsb	Solaris x86 DDI
repinsd	Solaris x86 DDI
repinsw	Solaris x86 DDI
repoutsb	Solaris x86 DDI
repoutsd	Solaris x86 DDI
repoutsw	Solaris x86 DDI
rmalloc	DDI/DKI
rmalloc_wait	DDI/DKI
rmallocmap	DDI/DKI
rmallocmap_wait	DDI/DKI

Routine	Type
rmfree	DDI/DKI
rmfreemap	DDI/DKI
rw_destroy	Solaris DDI
rw_downgrade	Solaris DDI
rw_enter	Solaris DDI
rw_exit	Solaris DDI
rw_init	Solaris DDI
rw_read_locked	Solaris DDI
rw_tryenter	Solaris DDI
rw_tryupgrade	Solaris DDI
scsi_abort	Solaris DDI
scsi_alloc_consistent_buf	Solaris DDI
scsi_cname	Solaris DDI
scsi_destroy_pkt	Solaris DDI
scsi_dmafree	Solaris DDI
scsi_dmaget	Solaris DDI
scsi_dname	Solaris DDI
scsi_errmsg	Solaris DDI
scsi_free_consistent_buf	Solaris DDI
scsi_hba_attach	Solaris DDI
scsi_hba_attach_setup	Solaris DDI
scsi_hba_detach	Solaris DDI
scsi_hba_fini	Solaris DDI
scsi_hba_init	Solaris DDI
scsi_hba_lookup_capstr	Solaris DDI
scsi_hba_pkt_alloc	Solaris DDI
scsi_hba_pkt_free	Solaris DDI
scsi_hba_probe	Solaris DDI
scsi_hba_tran_alloc	Solaris DDI
scsi_hba_tran_free	Solaris DDI

Routine	Type
scsi_ifgetcap	Solaris DDI
scsi_ifsetcap	Solaris DDI
scsi_init_pkt	Solaris DDI
scsi_log	Solaris DDI
scsi_mname	Solaris DDI
scsi_pktalloc	Solaris DDI
scsi_pktfree	Solaris DDI
scsi_poll	Solaris DDI
scsi_probe	Solaris DDI
scsi_realloc	Solaris DDI
scsi_reset	Solaris DDI
scsi_reset_notify	Solaris DDI
scsi_resfree	Solaris DDI
scsi_rname	Solaris DDI
scsi_slave	Solaris DDI
scsi_sname	Solaris DDI
scsi_sync_pkt	Solaris DDI
scsi_transport	Solaris DDI
scsi_unprobe	Solaris DDI
scsi_unslave	Solaris DDI
sema_destroy	Solaris DDI
sema_init	Solaris DDI
sema_p	Solaris DDI
sema_p_sig	Solaris DDI
sema_try	Solaris DDI
sema_v	Solaris DDI
sprintf	Solaris DDI
stoi	Solaris DDI
strchr	Solaris DDI
strcmp	Solaris DDI

Routine	Type
strcpy	Solaris DDI
strlen	Solaris DDI
strncmp	Solaris DDI
strncpy	Solaris DDI
swab	DDI/DKI
timeout	DDI/DKI
uiomove	DDI/DKI
untimeout	DDI/DKI
ureadc	DDI/DKI
uwritec	DDI/DKI
va_arg	Solaris DDI
va_end	Solaris DDI
va_start	Solaris DDI
vcmn_err	DDI/DKI
vsprintf	Solaris DDI

**LIST OF  
FUNCTIONS**

Name	Description
<b>ASSERT(9F)</b>	expression verification
<b>IOC_CONVERT_FROM(9F)</b>	determine if there is a need to translate M_IOCTL contents.
<b>Intro(9F)</b>	introduction to DDI/DKI functions
<b>OTHERQ(9F)</b>	get pointer to queue's partner queue
<b>RD(9F)</b>	get pointer to the read queue
<b>SAMESTR(9F)</b>	test if next queue is in the same stream
<b>SIZEOF_PTR(9F)</b>	See <b>STRUCT_DECL(9F)</b>

<b>SIZEOF_SIZE(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_BUF(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_DECL(9F)</b>	32bit application data access macros
<b>STRUCT_FADDR(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_FGET(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_FGETP(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_FSET(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_FSETP(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_HANDLE(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_INIT(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>STRUCT_SET_HANDLE(9F)</b>	See <b>STRUCT_DECL(9F)</b>
<b>WR(9F)</b>	get pointer to the write queue for this module or driver
<b>adjmsg(9F)</b>	trim bytes from a message
<b>allocb(9F)</b>	allocate a message block
<b>anocancel(9F)</b>	prevent cancellation of asynchronous I/O request
<b>aphysio(9F)</b>	perform asynchronous physical I/O
<b>assert(9F)</b>	See <b>ASSERT(9F)</b>
<b>backq(9F)</b>	get pointer to the queue behind the current queue
<b>bcanput(9F)</b>	test for flow control in specified priority band
<b>bcanputnext(9F)</b>	See <b>canputnext(9F)</b>
<b>bcmp(9F)</b>	compare two byte arrays

<b>bcopy(9F)</b>	copy data between address locations in the kernel
<b>bioclone(9F)</b>	clone another buffer
<b>biodone(9F)</b>	release buffer after buffer I/O transfer and notify blocked threads
<b>bioerror(9F)</b>	indicate error in buffer header
<b>biofini(9F)</b>	uninitialize a buffer structure
<b>bioinit(9F)</b>	initialize a buffer structure
<b>biomodified(9F)</b>	check if a buffer is modified
<b>bioreset(9F)</b>	reuse a private buffer header after I/O is complete
<b>biosize(9F)</b>	returns size of a buffer structure
<b>biowait(9F)</b>	suspend processes pending completion of block I/O
<b>bp_mapin(9F)</b>	allocate virtual address space
<b>bp_mapout(9F)</b>	deallocate virtual address space
<b>btop(9F)</b>	convert size in bytes to size in pages (round down)
<b>btopr(9F)</b>	convert size in bytes to size in pages (round up)
<b>bufcall(9F)</b>	call a function when a buffer becomes available
<b>bzero(9F)</b>	clear memory for a given number of bytes
<b>canput(9F)</b>	test for room in a message queue
<b>canputnext(9F)</b>	test for room in next module's message queue

<b>clrbuf(9F)</b>	erase the contents of a buffer
<b>cmn_err(9F)</b>	display an error message or panic the system
<b>condvar(9F)</b>	condition variable routines
<b>copyb(9F)</b>	copy a message block
<b>copyin(9F)</b>	copy data from a user program to a driver buffer
<b>copymsg(9F)</b>	copy a message
<b>copyout(9F)</b>	copy data from a driver to a user program
<b>csx_AccessConfigurationRegister(9F)</b>	read or write a PC Card Configuration Register
<b>csx_CS_DDI_Info(9F)</b>	obtain DDI information
<b>csx_ConvertSize(9F)</b>	convert device sizes
<b>csx_ConvertSpeed(9F)</b>	convert device speeds
<b>csx_DeregisterClient(9F)</b>	remove client from Card Services list
<b>csx_DupHandle(9F)</b>	duplicate access handle
<b>csx_Error2Text(9F)</b>	convert error return codes to text strings
<b>csx_Event2Text(9F)</b>	convert events to text strings
<b>csx_FreeHandle(9F)</b>	free access handle
<b>csx_Get16(9F)</b>	See <b>csx_Get8(9F)</b>
<b>csx_Get32(9F)</b>	See <b>csx_Get8(9F)</b>
<b>csx_Get64(9F)</b>	See <b>csx_Get8(9F)</b>
<b>csx_Get8(9F)</b>	read data from device address
<b>csx_GetEventMask(9F)</b>	See <b>csx_SetEventMask(9F)</b>

<b>csx_GetFirstClient(9F)</b>	return first or next client
<b>csx_GetFirstTuple(9F)</b>	return Card Information Structure tuple
<b>csx_GetHandleOffset(9F)</b>	return current access handle offset
<b>csx_GetMappedAddr(9F)</b>	return mapped virtual address
<b>csx_GetNextClient(9F)</b>	See <b>csx_GetFirstClient(9F)</b>
<b>csx_GetNextTuple(9F)</b>	See <b>csx_GetFirstTuple(9F)</b>
<b>csx_GetStatus(9F)</b>	return the current status of a PC Card and its socket
<b>csx_GetTupleData(9F)</b>	return the data portion of a tuple
<b>csx_MakeDeviceNode(9F)</b>	create and remove minor nodes on behalf of the client
<b>csx_MapLogSocket(9F)</b>	return the physical socket number associated with the client handle
<b>csx_MapMemPage(9F)</b>	map the memory area on a PC Card
<b>csx_ModifyConfiguration(9F)</b>	modify socket and PC Card Configuration Register
<b>csx_ModifyWindow(9F)</b>	modify window attributes
<b>csx_ParseTuple(9F)</b>	generic tuple parser
<b>csx_Parse_CISTPL_BATTERY(9F)</b>	parse the Battery Replacement Date tuple
<b>csx_Parse_CISTPL_BYTEORDER(9F)</b>	parse the Byte Order tuple
<b>csx_Parse_CISTPL_CFTABLE_ENTRY(9F)</b>	parse 16-bit Card Configuration Table Entry tuple
<b>csx_Parse_CISTPL_CONFIG(9F)</b>	parse Configuration tuple

<b>csx_Parse_CISTPL_DATE(9F)</b>	parse the Card Initialization Date tuple
<b>csx_Parse_CISTPL_DEVICE(9F)</b>	parse Device Information tuples
<b>csx_Parse_CISTPL_DEVICEGEO(9F)</b>	parse the Device Geo tuple
<b>csx_Parse_CISTPL_DEVICEGEO_A(9F)</b>	parse the Device Geo A tuple
<b>csx_Parse_CISTPL_DEVICE_A(9F)</b>	See <b>csx_Parse_CISTPL_DEVICE(9F)</b>
<b>csx_Parse_CISTPL_DEVICE_OA(9F)</b>	See <b>csx_Parse_CISTPL_DEVICE(9F)</b>
<b>csx_Parse_CISTPL_DEVICE_OC(9F)</b>	See <b>csx_Parse_CISTPL_DEVICE(9F)</b>
<b>csx_Parse_CISTPL_FORMAT(9F)</b>	parse the Data Recording Format tuple
<b>csx_Parse_CISTPL_FUNCE(9F)</b>	parse Function Extension tuple
<b>csx_Parse_CISTPL_FUNCID(9F)</b>	parse Function Identification tuple
<b>csx_Parse_CISTPL_GEOMETRY(9F)</b>	parse the Geometry tuple
<b>csx_Parse_CISTPL_JEDEC_A(9F)</b>	See <b>csx_Parse_CISTPL_JEDEC_C(9F)</b>
<b>csx_Parse_CISTPL_JEDEC_C(9F)</b>	parse JEDEC Identifier tuples
<b>csx_Parse_CISTPL_LINKTARGET(9F)</b>	parse the Link Target tuple
<b>csx_Parse_CISTPL_LONGLINK_A(9F)</b>	parse the Long Link A and C tuples
<b>csx_Parse_CISTPL_LONGLINK_C(9F)</b>	See <b>csx_Parse_CISTPL_LONGLINK_A(9F)</b>
<b>csx_Parse_CISTPL_LONGLINK_MFC(9F)</b>	parse the Multi-Function tuple
<b>csx_Parse_CISTPL_MANFID(9F)</b>	parse Manufacturer Identification tuple

<b>csx_Parse_CISTPL_ORG(9F)</b>	parse the Data Organization tuple
<b>csx_Parse_CISTPL_SPCL(9F)</b>	parse the Special Purpose tuple
<b>csx_Parse_CISTPL_SWIL(9F)</b>	parse the Software Interleaving tuple
<b>csx_Parse_CISTPL_VERS_1(9F)</b>	parse Level-1 Version/Product Information tuple
<b>csx_Parse_CISTPL_VERS_2(9F)</b>	parse Level-2 Version and Information tuple
<b>csx_Put16(9F)</b>	See <b>csx_Put8(9F)</b>
<b>csx_Put32(9F)</b>	See <b>csx_Put8(9F)</b>
<b>csx_Put64(9F)</b>	See <b>csx_Put8(9F)</b>
<b>csx_Put8(9F)</b>	write to device register
<b>csx_RegisterClient(9F)</b>	register a client
<b>csx_ReleaseConfiguration(9F)</b>	release PC Card and socket configuration
<b>csx_ReleaseIO(9F)</b>	See <b>csx_RequestIO(9F)</b>
<b>csx_ReleaseIRQ(9F)</b>	See <b>csx_RequestIRQ(9F)</b>
<b>csx_ReleaseSocketMask(9F)</b>	See <b>csx_RequestSocketMask(9F)</b>
<b>csx_ReleaseWindow(9F)</b>	See <b>csx_RequestWindow(9F)</b>
<b>csx_RemoveDeviceNode(9F)</b>	See <b>csx_MakeDeviceNode(9F)</b>
<b>csx_RepGet16(9F)</b>	See <b>csx_RepGet8(9F)</b>
<b>csx_RepGet32(9F)</b>	See <b>csx_RepGet8(9F)</b>
<b>csx_RepGet64(9F)</b>	See <b>csx_RepGet8(9F)</b>
<b>csx_RepGet8(9F)</b>	read repetitively from the device register

<b>csx_RepPut16(9F)</b>	See <b>csx_RepPut8(9F)</b>
<b>csx_RepPut32(9F)</b>	See <b>csx_RepPut8(9F)</b>
<b>csx_RepPut64(9F)</b>	See <b>csx_RepPut8(9F)</b>
<b>csx_RepPut8(9F)</b>	write repetitively to the device register
<b>csx_RequestConfiguration(9F)</b>	configure the PC Card and socket
<b>csx_RequestIO(9F)</b>	request or release I/O resources for the client
<b>csx_RequestIRQ(9F)</b>	request or release IRQ resource
<b>csx_RequestSocketMask(9F)</b>	set or clear the client's client event mask
<b>csx_RequestWindow(9F)</b>	request or release window resources
<b>csx_ResetFunction(9F)</b>	reset a function on a PC card
<b>csx_SetEventMask(9F)</b>	set or return the client event mask for the client
<b>csx_SetHandleOffset(9F)</b>	set current access handle offset
<b>csx_ValidateCIS(9F)</b>	validate the Card Information Structure (CIS)
<b>cv_broadcast(9F)</b>	See <b>condvar(9F)</b>
<b>cv_destroy(9F)</b>	See <b>condvar(9F)</b>
<b>cv_init(9F)</b>	See <b>condvar(9F)</b>
<b>cv_signal(9F)</b>	See <b>condvar(9F)</b>
<b>cv_timedwait(9F)</b>	See <b>condvar(9F)</b>
<b>cv_timedwait_sig(9F)</b>	See <b>condvar(9F)</b>
<b>cv_wait(9F)</b>	See <b>condvar(9F)</b>
<b>cv_wait_sig(9F)</b>	See <b>condvar(9F)</b>

<b>datamsq(9F)</b>	test whether a message is a data message
<b>ddi_add_intr(9F)</b>	hardware interrupt handling routines
<b>ddi_add_softintr(9F)</b>	software interrupt handling routines
<b>ddi_binding_name(9F)</b>	return driver binding name
<b>ddi_btop(9F)</b>	page size conversions
<b>ddi_btopr(9F)</b>	See <b>ddi_btop(9F)</b>
<b>ddi_copyin(9F)</b>	copy data to a driver buffer
<b>ddi_copyout(9F)</b>	copy data from a driver
<b>ddi_create_minor_node(9F)</b>	create a minor node for this device
<b>ddi_dev_is_needed(9F)</b>	inform the system that a device's component is required
<b>ddi_dev_is_sid(9F)</b>	tell whether a device is self-identifying
<b>ddi_dev_nintrs(9F)</b>	return the number of interrupt specifications a device has
<b>ddi_dev_nregs(9F)</b>	return the number of register sets a device has
<b>ddi_dev_regsize(9F)</b>	return the size of a device's register
<b>ddi_device_copy(9F)</b>	copy data from one device register to another device register
<b>ddi_device_zero(9F)</b>	zero fill the device
<b>ddi_devid_compare(9F)</b>	Kernel interfaces for device ids
<b>ddi_devid_free(9F)</b>	See <b>ddi_devid_compare(9F)</b>

<b>ddi_devid_init(9F)</b>	See <b>ddi_devid_compare(9F)</b>
<b>ddi_devid_register(9F)</b>	See <b>ddi_devid_compare(9F)</b>
<b>ddi_devid_sizeof(9F)</b>	See <b>ddi_devid_compare(9F)</b>
<b>ddi_devid_unregister(9F)</b>	See <b>ddi_devid_compare(9F)</b>
<b>ddi_devid_valid(9F)</b>	See <b>ddi_devid_compare(9F)</b>
<b>ddi_devmap_segmap(9F)</b>	See <b>devmap_setup(9F)</b>
<b>ddi_dma_addr_bind_handle(9F)</b>	binds an address to a DMA handle
<b>ddi_dma_addr_setup(9F)</b>	easier DMA setup for use with virtual addresses
<b>ddi_dma_alloc_handle(9F)</b>	allocate DMA handle
<b>ddi_dma_buf_bind_handle(9F)</b>	binds a system buffer to a DMA handle
<b>ddi_dma_buf_setup(9F)</b>	easier DMA setup for use with buffer structures
<b>ddi_dma_burstsizes(9F)</b>	find out the allowed burst sizes for a DMA mapping
<b>ddi_dma_coff(9F)</b>	convert a DMA cookie to an offset within a DMA handle
<b>ddi_dma_curwin(9F)</b>	report current DMA window offset and size
<b>ddi_dma_dealign(9F)</b>	find DMA mapping alignment and minimum transfer size
<b>ddi_dma_free(9F)</b>	release system DMA resources
<b>ddi_dma_free_handle(9F)</b>	free DMA handle
<b>ddi_dma_getwin(9F)</b>	activate a new DMA window
<b>ddi_dma_htoc(9F)</b>	convert a DMA handle to a DMA address cookie

<code>ddi_dma_mem_alloc(9F)</code>	allocate memory for DMA transfer
<code>ddi_dma_mem_free(9F)</code>	free previously allocated memory
<code>ddi_dma_movwin(9F)</code>	shift current DMA window
<code>ddi_dma_nextcookie(9F)</code>	retrieve subsequent DMA cookie
<code>ddi_dma_nextseg(9F)</code>	get next DMA segment
<code>ddi_dma_nextwin(9F)</code>	get next DMA window
<code>ddi_dma_numwin(9F)</code>	retrieve number of DMA windows
<code>ddi_dma_segtocookie(9F)</code>	convert a DMA segment to a DMA address cookie
<code>ddi_dma_set_sbus64(9F)</code>	allow 64bit transfers on SBus
<code>ddi_dma_setup(9F)</code>	setup DMA resources
<code>ddi_dma_sync(9F)</code>	synchronize CPU and I/O views of memory
<code>ddi_dma_unbind_handle(9F)</code>	unbinds the address in a DMA handle
<code>ddi_dmae(9F)</code>	system DMA engine functions
<code>ddi_dmae_1stparty(9F)</code>	See <code>ddi_dmae(9F)</code>
<code>ddi_dmae_alloc(9F)</code>	See <code>ddi_dmae(9F)</code>
<code>ddi_dmae_disable(9F)</code>	See <code>ddi_dmae(9F)</code>
<code>ddi_dmae_enable(9F)</code>	See <code>ddi_dmae(9F)</code>
<code>ddi_dmae_getattr(9F)</code>	See <code>ddi_dmae(9F)</code>
<code>ddi_dmae_getcnt(9F)</code>	See <code>ddi_dmae(9F)</code>
<code>ddi_dmae_getlim(9F)</code>	See <code>ddi_dmae(9F)</code>
<code>ddi_dmae_prog(9F)</code>	See <code>ddi_dmae(9F)</code>

<b>ddi_dmae_release(9F)</b>	See <b>ddi_dmae(9F)</b>
<b>ddi_dmae_stop(9F)</b>	See <b>ddi_dmae(9F)</b>
<b>ddi_enter_critical(9F)</b>	enter and exit a critical region of control
<b>ddi_exit_critical(9F)</b>	See <b>ddi_enter_critical(9F)</b>
<b>ddi_ffs(9F)</b>	find first (last) bit set in a long integer
<b>ddi_fls(9F)</b>	See <b>ddi_ffs(9F)</b>
<b>ddi_get16(9F)</b>	See <b>ddi_get8(9F)</b>
<b>ddi_get32(9F)</b>	See <b>ddi_get8(9F)</b>
<b>ddi_get64(9F)</b>	See <b>ddi_get8(9F)</b>
<b>ddi_get8(9F)</b>	read data from the mapped memory address, device register or allocated DMA memory address
<b>ddi_get_cred(9F)</b>	returns a pointer to the credential structure of the caller
<b>ddi_get_driver_private(9F)</b>	get or set the address of the device's private data area
<b>ddi_get_iblock_cookie(9F)</b>	See <b>ddi_add_intr(9F)</b>
<b>ddi_get_instance(9F)</b>	get device instance number
<b>ddi_get_lbolt(9F)</b>	returns the value of lbolt
<b>ddi_get_name(9F)</b>	See <b>ddi_binding_name(9F)</b>
<b>ddi_get_parent(9F)</b>	find the parent of a device information structure
<b>ddi_get_pid(9F)</b>	returns the process
<b>ddi_get_soft_iblock_cookie(9F)</b>	See <b>ddi_add_softintr(9F)</b>

<code>ddi_get_soft_state(9F)</code>	See <code>ddi_soft_state(9F)</code>
<code>ddi_get_time(9F)</code>	returns the current time in seconds
<code>ddi_getb(9F)</code>	See <code>ddi_get8(9F)</code>
<code>ddi_getl(9F)</code>	See <code>ddi_get8(9F)</code>
<code>ddi_getll(9F)</code>	See <code>ddi_get8(9F)</code>
<code>ddi_getlongprop(9F)</code>	See <code>ddi_prop_op(9F)</code>
<code>ddi_getlongprop_buf(9F)</code>	See <code>ddi_prop_op(9F)</code>
<code>ddi_getprop(9F)</code>	See <code>ddi_prop_op(9F)</code>
<code>ddi_getproplen(9F)</code>	See <code>ddi_prop_op(9F)</code>
<code>ddi_getw(9F)</code>	See <code>ddi_get8(9F)</code>
<code>ddi_in_panic(9F)</code>	determine if system is in panic state
<code>ddi_intr_hilevel(9F)</code>	indicate interrupt handler type
<code>ddi_io_get16(9F)</code>	See <code>ddi_io_get8(9F)</code>
<code>ddi_io_get32(9F)</code>	See <code>ddi_io_get8(9F)</code>
<code>ddi_io_get8(9F)</code>	read data from the mapped device register in I/O space
<code>ddi_io_getb(9F)</code>	See <code>ddi_io_get8(9F)</code>
<code>ddi_io_getl(9F)</code>	See <code>ddi_io_get8(9F)</code>
<code>ddi_io_getw(9F)</code>	See <code>ddi_io_get8(9F)</code>
<code>ddi_io_put16(9F)</code>	See <code>ddi_io_put8(9F)</code>
<code>ddi_io_put32(9F)</code>	See <code>ddi_io_put8(9F)</code>
<code>ddi_io_put8(9F)</code>	write data to the mapped device register in I/O space
<code>ddi_io_putb(9F)</code>	See <code>ddi_io_put8(9F)</code>

<b>ddi_io_putl(9F)</b>	See <b>ddi_io_put8(9F)</b>
<b>ddi_io_putw(9F)</b>	See <b>ddi_io_put8(9F)</b>
<b>ddi_io_rep_get16(9F)</b>	See <b>ddi_io_rep_get8(9F)</b>
<b>ddi_io_rep_get32(9F)</b>	See <b>ddi_io_rep_get8(9F)</b>
<b>ddi_io_rep_get8(9F)</b>	read multiple data from the mapped device register in I/O space
<b>ddi_io_rep_getb(9F)</b>	See <b>ddi_io_rep_get8(9F)</b>
<b>ddi_io_rep_getl(9F)</b>	See <b>ddi_io_rep_get8(9F)</b>
<b>ddi_io_rep_getw(9F)</b>	See <b>ddi_io_rep_get8(9F)</b>
<b>ddi_io_rep_put16(9F)</b>	See <b>ddi_io_rep_put8(9F)</b>
<b>ddi_io_rep_put32(9F)</b>	See <b>ddi_io_rep_put8(9F)</b>
<b>ddi_io_rep_put8(9F)</b>	write multiple data to the mapped device register in I/O space
<b>ddi_io_rep_putb(9F)</b>	See <b>ddi_io_rep_put8(9F)</b>
<b>ddi_io_rep_putl(9F)</b>	See <b>ddi_io_rep_put8(9F)</b>
<b>ddi_io_rep_putw(9F)</b>	See <b>ddi_io_rep_put8(9F)</b>
<b>ddi_iomin(9F)</b>	find minimum alignment and transfer size for DMA
<b>ddi_iopb_alloc(9F)</b>	allocate and free non-sequentially accessed memory
<b>ddi_iopb_free(9F)</b>	See <b>ddi_iopb_alloc(9F)</b>
<b>ddi_map_regs(9F)</b>	map or unmap registers
<b>ddi_mapdev(9F)</b>	create driver-controlled mapping of device

<b>ddi_mapdev_intercept(9F)</b>	control driver notification of user accesses
<b>ddi_mapdev_nointercept(9F)</b>	See <b>ddi_mapdev_intercept(9F)</b>
<b>ddi_mapdev_set_device_acc_attr(9F)</b>	set the device attributes for the mapping
<b>ddi_mem_alloc(9F)</b>	allocate and free sequentially accessed memory
<b>ddi_mem_free(9F)</b>	See <b>ddi_mem_alloc(9F)</b>
<b>ddi_mem_get16(9F)</b>	See <b>ddi_mem_get8(9F)</b>
<b>ddi_mem_get32(9F)</b>	See <b>ddi_mem_get8(9F)</b>
<b>ddi_mem_get64(9F)</b>	See <b>ddi_mem_get8(9F)</b>
<b>ddi_mem_get8(9F)</b>	read data from mapped device in the memory space or allocated DMA memory
<b>ddi_mem_getb(9F)</b>	See <b>ddi_mem_get8(9F)</b>
<b>ddi_mem_getl(9F)</b>	See <b>ddi_mem_get8(9F)</b>
<b>ddi_mem_getll(9F)</b>	See <b>ddi_mem_get8(9F)</b>
<b>ddi_mem_getw(9F)</b>	See <b>ddi_mem_get8(9F)</b>
<b>ddi_mem_put16(9F)</b>	See <b>ddi_mem_put8(9F)</b>
<b>ddi_mem_put32(9F)</b>	See <b>ddi_mem_put8(9F)</b>
<b>ddi_mem_put64(9F)</b>	See <b>ddi_mem_put8(9F)</b>
<b>ddi_mem_put8(9F)</b>	write data to mapped device in the memory space or allocated DMA memory
<b>ddi_mem_putb(9F)</b>	See <b>ddi_mem_put8(9F)</b>
<b>ddi_mem_putl(9F)</b>	See <b>ddi_mem_put8(9F)</b>
<b>ddi_mem_putll(9F)</b>	See <b>ddi_mem_put8(9F)</b>

<code>ddi_mem_putw(9F)</code>	See <code>ddi_mem_put8(9F)</code>
<code>ddi_mem_rep_get16(9F)</code>	See <code>ddi_mem_rep_get8(9F)</code>
<code>ddi_mem_rep_get32(9F)</code>	See <code>ddi_mem_rep_get8(9F)</code>
<code>ddi_mem_rep_get64(9F)</code>	See <code>ddi_mem_rep_get8(9F)</code>
<code>ddi_mem_rep_get8(9F)</code>	read multiple data from mapped device in the memory space or allocated DMA memory
<code>ddi_mem_rep_getb(9F)</code>	See <code>ddi_mem_rep_get8(9F)</code>
<code>ddi_mem_rep_getl(9F)</code>	See <code>ddi_mem_rep_get8(9F)</code>
<code>ddi_mem_rep_getll(9F)</code>	See <code>ddi_mem_rep_get8(9F)</code>
<code>ddi_mem_rep_getw(9F)</code>	See <code>ddi_mem_rep_get8(9F)</code>
<code>ddi_mem_rep_put16(9F)</code>	See <code>ddi_mem_rep_put8(9F)</code>
<code>ddi_mem_rep_put32(9F)</code>	See <code>ddi_mem_rep_put8(9F)</code>
<code>ddi_mem_rep_put64(9F)</code>	See <code>ddi_mem_rep_put8(9F)</code>
<code>ddi_mem_rep_put8(9F)</code>	write multiple data to mapped device in the memory space or allocated DMA memory
<code>ddi_mem_rep_putb(9F)</code>	See <code>ddi_mem_rep_put8(9F)</code>
<code>ddi_mem_rep_putl(9F)</code>	See <code>ddi_mem_rep_put8(9F)</code>
<code>ddi_mem_rep_putll(9F)</code>	See <code>ddi_mem_rep_put8(9F)</code>
<code>ddi_mem_rep_putw(9F)</code>	See <code>ddi_mem_rep_put8(9F)</code>
<code>ddi_mmap_get_model(9F)</code>	return data model type of current thread
<code>ddi_model_convert_from(9F)</code>	determine data model type mismatch
<code>ddi_node_name(9F)</code>	return the devinfo node name
<code>ddi_peek(9F)</code>	read a value from a location

<b>ddi_peek16(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_peek32(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_peek64(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_peek8(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_peekc(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_peekd(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_peekl(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_peeks(9F)</b>	See <b>ddi_peek(9F)</b>
<b>ddi_poke(9F)</b>	write a value to a location
<b>ddi_poke16(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_poke32(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_poke64(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_poke8(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_pokec(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_poked(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_pokel(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_pokes(9F)</b>	See <b>ddi_poke(9F)</b>
<b>ddi_prop_create(9F)</b>	create, remove, or modify properties for leaf device drivers
<b>ddi_prop_exists(9F)</b>	check for the existence of a property
<b>ddi_prop_free(9F)</b>	See <b>ddi_prop_lookup(9F)</b>
<b>ddi_prop_get_int(9F)</b>	lookup integer property
<b>ddi_prop_lookup(9F)</b>	look up property information
<b>ddi_prop_lookup_byte_array(9F)</b>	See <b>ddi_prop_lookup(9F)</b>

<code>ddi_prop_lookup_int_array(9F)</code>	See <code>ddi_prop_lookup(9F)</code>
<code>ddi_prop_lookup_string(9F)</code>	See <code>ddi_prop_lookup(9F)</code>
<code>ddi_prop_lookup_string_array(9F)</code>	See <code>ddi_prop_lookup(9F)</code>
<code>ddi_prop_modify(9F)</code>	See <code>ddi_prop_create(9F)</code>
<code>ddi_prop_op(9F)</code>	get property information for leaf device drivers
<code>ddi_prop_remove(9F)</code>	See <code>ddi_prop_create(9F)</code>
<code>ddi_prop_remove_all(9F)</code>	See <code>ddi_prop_create(9F)</code>
<code>ddi_prop_undefine(9F)</code>	See <code>ddi_prop_create(9F)</code>
<code>ddi_prop_update(9F)</code>	update properties
<code>ddi_prop_update_byte_array(9F)</code>	See <code>ddi_prop_update(9F)</code>
<code>ddi_prop_update_int(9F)</code>	See <code>ddi_prop_update(9F)</code>
<code>ddi_prop_update_int_array(9F)</code>	See <code>ddi_prop_update(9F)</code>
<code>ddi_prop_update_string(9F)</code>	See <code>ddi_prop_update(9F)</code>
<code>ddi_prop_update_string_array(9F)</code>	See <code>ddi_prop_update(9F)</code>
<code>ddi_ptob(9F)</code>	See <code>ddi_btop(9F)</code>
<code>ddi_put16(9F)</code>	See <code>ddi_put8(9F)</code>
<code>ddi_put32(9F)</code>	See <code>ddi_put8(9F)</code>
<code>ddi_put64(9F)</code>	See <code>ddi_put8(9F)</code>
<code>ddi_put8(9F)</code>	write data to the mapped memory address, device register or allocated DMA memory address
<code>ddi_putb(9F)</code>	See <code>ddi_put8(9F)</code>
<code>ddi_putl(9F)</code>	See <code>ddi_put8(9F)</code>
<code>ddi_putll(9F)</code>	See <code>ddi_put8(9F)</code>

<code>ddi_putw(9F)</code>	See <code>ddi_put8(9F)</code>
<code>ddi_regs_map_free(9F)</code>	free a previously mapped register address space
<code>ddi_regs_map_setup(9F)</code>	set up a mapping for a register address space
<code>ddi_remove_intr(9F)</code>	See <code>ddi_add_intr(9F)</code>
<code>ddi_remove_minor_node(9F)</code>	remove a minor node for this <code>dev_info</code>
<code>ddi_remove_softcintr(9F)</code>	See <code>ddi_add_softcintr(9F)</code>
<code>ddi_rep_get16(9F)</code>	See <code>ddi_rep_get8(9F)</code>
<code>ddi_rep_get32(9F)</code>	See <code>ddi_rep_get8(9F)</code>
<code>ddi_rep_get64(9F)</code>	See <code>ddi_rep_get8(9F)</code>
<code>ddi_rep_get8(9F)</code>	read data from the mapped memory address, device register or allocated DMA memory address
<code>ddi_rep_getb(9F)</code>	See <code>ddi_rep_get8(9F)</code>
<code>ddi_rep_getl(9F)</code>	See <code>ddi_rep_get8(9F)</code>
<code>ddi_rep_getll(9F)</code>	See <code>ddi_rep_get8(9F)</code>
<code>ddi_rep_getw(9F)</code>	See <code>ddi_rep_get8(9F)</code>
<code>ddi_rep_put16(9F)</code>	See <code>ddi_rep_put8(9F)</code>
<code>ddi_rep_put32(9F)</code>	See <code>ddi_rep_put8(9F)</code>
<code>ddi_rep_put64(9F)</code>	See <code>ddi_rep_put8(9F)</code>
<code>ddi_rep_put8(9F)</code>	write data to the mapped memory address, device register or allocated DMA memory address
<code>ddi_rep_putb(9F)</code>	See <code>ddi_rep_put8(9F)</code>

<b>ddi_rep_putl(9F)</b>	See <b>ddi_rep_put8(9F)</b>
<b>ddi_rep_putll(9F)</b>	See <b>ddi_rep_put8(9F)</b>
<b>ddi_rep_putw(9F)</b>	See <b>ddi_rep_put8(9F)</b>
<b>ddi_report_dev(9F)</b>	announce a device
<b>ddi_root_node(9F)</b>	get the root of the dev_info tree
<b>ddi_segmap(9F)</b>	set up a user mapping using seg_dev
<b>ddi_segmap_setup(9F)</b>	See <b>ddi_segmap(9F)</b>
<b>ddi_set_driver_private(9F)</b>	See <b>ddi_get_driver_private(9F)</b>
<b>ddi_slaveonly(9F)</b>	tell if a device is installed in a slave access only location
<b>ddi_soft_state(9F)</b>	driver soft state utility routines
<b>ddi_soft_state_fini(9F)</b>	See <b>ddi_soft_state(9F)</b>
<b>ddi_soft_state_free(9F)</b>	See <b>ddi_soft_state(9F)</b>
<b>ddi_soft_state_init(9F)</b>	See <b>ddi_soft_state(9F)</b>
<b>ddi_soft_state_zalloc(9F)</b>	See <b>ddi_soft_state(9F)</b>
<b>ddi_trigger_softintr(9F)</b>	See <b>ddi_add_softintr(9F)</b>
<b>ddi_umem_alloc(9F)</b>	allocate and free page-aligned kernel memory
<b>ddi_umem_free(9F)</b>	See <b>ddi_umem_alloc(9F)</b>
<b>ddi_unmap_regs(9F)</b>	See <b>ddi_map_regs(9F)</b>
<b>delay(9F)</b>	delay execution for a specified number of clock ticks
<b>devmap_default_access(9F)</b>	default driver memory access function

<b>devmap_devmem_setup(9F)</b>	set driver memory mapping parameters
<b>devmap_do_ctxmgt(9F)</b>	perform device context switching on a mapping
<b>devmap_load(9F)</b>	See <b>devmap_unload(9F)</b>
<b>devmap_set_ctx_timeout(9F)</b>	set the timeout value for the context management callback
<b>devmap_setup(9F)</b>	set up a user mapping to device memory using the devmap framework
<b>devmap_umem_setup(9F)</b>	See <b>devmap_devmem_setup(9F)</b>
<b>devmap_unload(9F)</b>	control validation of memory address translations
<b>disksort(9F)</b>	single direction elevator seek sort for buffers
<b>drv_getparm(9F)</b>	retrieve kernel state information
<b>drv_hztousec(9F)</b>	convert clock ticks to microseconds
<b>drv_priv(9F)</b>	determine driver privilege
<b>drv_usectohz(9F)</b>	convert microseconds to clock ticks
<b>drv_usecwait(9F)</b>	busy-wait for specified interval
<b>dupb(9F)</b>	duplicate a message block descriptor
<b>dupmsg(9F)</b>	duplicate a message
<b>enableok(9F)</b>	reschedule a queue for service
<b>esballoc(9F)</b>	allocate a message block using a caller-supplied buffer

<b>esbcall(9F)</b>	call function when buffer is available
<b>flushband(9F)</b>	flush messages for a specified priority band
<b>flushq(9F)</b>	remove messages from a queue
<b>free_pktiopb(9F)</b>	See <b>get_pktiopb(9F)</b>
<b>freeb(9F)</b>	free a message block
<b>freemsg(9F)</b>	free all message blocks in a message
<b>freerbuf(9F)</b>	free a raw buffer header
<b>freezestr(9F)</b>	freeze, thaw the state of a stream
<b>get_pktiopb(9F)</b>	allocate/free a SCSI packet in the iopb map
<b>geterror(9F)</b>	return I/O error
<b>getmajor(9F)</b>	get major device number
<b>getminor(9F)</b>	get minor device number
<b>getq(9F)</b>	get the next message from a queue
<b>getrbuf(9F)</b>	get a raw buffer header
<b>hat_getkpfnum(9F)</b>	get page frame number for kernel address
<b>inb(9F)</b>	read from an I/O port
<b>inl(9F)</b>	See <b>inb(9F)</b>
<b>insq(9F)</b>	insert a message into a queue
<b>intro(9F)</b>	See <b>Intro(9F)</b>
<b>inw(9F)</b>	See <b>inb(9F)</b>

<b>kmem_alloc(9F)</b>	allocate kernel memory
<b>kmem_free(9F)</b>	See <b>kmem_alloc(9F)</b>
<b>kmem_zalloc(9F)</b>	See <b>kmem_alloc(9F)</b>
<b>kstat_create(9F)</b>	create and initialize a new kstat
<b>kstat_delete(9F)</b>	remove a kstat from the system
<b>kstat_install(9F)</b>	add a fully initialized kstat to the system
<b>kstat_named_init(9F)</b>	initialize a named kstat
<b>kstat_queue(9F)</b>	update I/O kstat statistics
<b>kstat_runq_back_to_waitq(9F)</b>	See <b>kstat_queue(9F)</b>
<b>kstat_runq_enter(9F)</b>	See <b>kstat_queue(9F)</b>
<b>kstat_runq_exit(9F)</b>	See <b>kstat_queue(9F)</b>
<b>kstat_waitq_enter(9F)</b>	See <b>kstat_queue(9F)</b>
<b>kstat_waitq_exit(9F)</b>	See <b>kstat_queue(9F)</b>
<b>kstat_waitq_to_runq(9F)</b>	See <b>kstat_queue(9F)</b>
<b>linkb(9F)</b>	concatenate two message blocks
<b>makecom(9F)</b>	make a packet for SCSI commands
<b>makecom_g0(9F)</b>	See <b>makecom(9F)</b>
<b>makecom_g0_s(9F)</b>	See <b>makecom(9F)</b>
<b>makecom_g1(9F)</b>	See <b>makecom(9F)</b>
<b>makecom_g5(9F)</b>	See <b>makecom(9F)</b>
<b>makedevice(9F)</b>	make device number from major and minor numbers
<b>max(9F)</b>	return the larger of two integers
<b>min(9F)</b>	return the lesser of two integers

<b>minphys(9F)</b>	See <b>physio(9F)</b>
<b>mkiocb(9F)</b>	allocates a STREAMS ioctl block for M_IOCTL messages in the kernel.
<b>mod_info(9F)</b>	See <b>mod_install(9F)</b>
<b>mod_install(9F)</b>	add, remove or query a loadable module
<b>mod_remove(9F)</b>	See <b>mod_install(9F)</b>
<b>msgdsize(9F)</b>	return the number of bytes in a message
<b>msgpullup(9F)</b>	concatenate bytes in a message
<b>mt-streams(9F)</b>	STREAMS multithreading
<b>mutex(9F)</b>	mutual exclusion lock routines
<b>mutex_destroy(9F)</b>	See <b>mutex(9F)</b>
<b>mutex_enter(9F)</b>	See <b>mutex(9F)</b>
<b>mutex_exit(9F)</b>	See <b>mutex(9F)</b>
<b>mutex_init(9F)</b>	See <b>mutex(9F)</b>
<b>mutex_owned(9F)</b>	See <b>mutex(9F)</b>
<b>mutex_tryenter(9F)</b>	See <b>mutex(9F)</b>
<b>nochpoll(9F)</b>	error return function for non-pollable devices
<b>nodev(9F)</b>	error return function
<b>noenable(9F)</b>	prevent a queue from being scheduled
<b>nulldev(9F)</b>	zero return function
<b>numtos(9F)</b>	See <b>stoi(9F)</b>
<b>otherq(9F)</b>	See <b>OTHERQ(9F)</b>

<b>outb(9F)</b>	write to an I/O port
<b>outl(9F)</b>	See <b>outb(9F)</b>
<b>outw(9F)</b>	See <b>outb(9F)</b>
<b>pci_config_get16(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_get32(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_get64(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_get8(9F)</b>	read or write single datum of various sizes to the PCI Local Bus Configuration space
<b>pci_config_getb(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_getl(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_getll(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_getw(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_put16(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_put32(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_put64(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_put8(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_putb(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_putl(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_putll(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_putw(9F)</b>	See <b>pci_config_get8(9F)</b>
<b>pci_config_setup(9F)</b>	setup or tear down the resources for enabling accesses to the PCI Local Bus Configuration space
<b>pci_config_teardown(9F)</b>	See <b>pci_config_setup(9F)</b>
<b>physio(9F)</b>	perform physical I/O

<code>pm_busy_component(9F)</code>	control device components' availability for power management
<code>pm_create_components(9F)</code>	create or destroy power-manageable components
<code>pm_destroy_components(9F)</code>	See <code>pm_create_components(9F)</code>
<code>pm_get_normal_power(9F)</code>	get or set a device component's normal power level
<code>pm_idle_component(9F)</code>	See <code>pm_busy_component(9F)</code>
<code>pm_set_normal_power(9F)</code>	See <code>pm_get_normal_power(9F)</code>
<code>pollwakeup(9F)</code>	inform a process that an event has occurred
<code>proc_ref(9F)</code>	See <code>proc_signal(9F)</code>
<code>proc_signal(9F)</code>	send a signal to a process
<code>proc_unref(9F)</code>	See <code>proc_signal(9F)</code>
<code>ptob(9F)</code>	convert size in pages to size in bytes
<code>pullupmsg(9F)</code>	concatenate bytes in a message
<code>put(9F)</code>	call a STREAMS put procedure
<code>putbq(9F)</code>	place a message at the head of a queue
<code>putctl(9F)</code>	send a control message to a queue
<code>putctl1(9F)</code>	send a control message with a one-byte parameter to a queue
<code>putnext(9F)</code>	send a message to the next queue

<b>putnextctl(9F)</b>	send a control message to a queue
<b>putnextctl1(9F)</b>	send a control message with a one-byte parameter to a queue
<b>putq(9F)</b>	put a message on a queue
<b>qbufcall(9F)</b>	call a function when a buffer becomes available
<b>qenable(9F)</b>	enable a queue
<b>qprocsoff(9F)</b>	See <b>qprocson(9F)</b>
<b>qprocson(9F)</b>	enable, disable put and service routines
<b>qreply(9F)</b>	send a message on a stream in the reverse direction
<b>qsize(9F)</b>	find the number of messages on a queue
<b>qtimeout(9F)</b>	execute a function after a specified length of time
<b>qunbufcall(9F)</b>	cancel a pending qbufcall request
<b>quntimeout(9F)</b>	cancel previous qtimeout function call
<b>qwait(9F)</b>	STREAMS wait routines
<b>qwait_sig(9F)</b>	See <b>qwait(9F)</b>
<b>qwriter(9F)</b>	asynchronous STREAMS perimeter upgrade
<b>rd(9F)</b>	See <b>RD(9F)</b>
<b>repinsb(9F)</b>	See <b>inb(9F)</b>
<b>repinsd(9F)</b>	See <b>inb(9F)</b>
<b>repinsw(9F)</b>	See <b>inb(9F)</b>

<b>repoutsb(9F)</b>	See <b>outb(9F)</b>
<b>reputsd(9F)</b>	See <b>outb(9F)</b>
<b>reputsw(9F)</b>	See <b>outb(9F)</b>
<b>rmalloc(9F)</b>	allocate space from a resource map
<b>rmalloc_wait(9F)</b>	allocate space from a resource map, wait if necessary
<b>rmallocmap(9F)</b>	allocate and free resource maps
<b>rmallocmap_wait(9F)</b>	See <b>rmallocmap(9F)</b>
<b>rmfree(9F)</b>	free space back into a resource map
<b>rmfreemap(9F)</b>	See <b>rmallocmap(9F)</b>
<b>rmvb(9F)</b>	remove a message block from a message
<b>rmvq(9F)</b>	remove a message from a queue
<b>rw_destroy(9F)</b>	See <b>rwlock(9F)</b>
<b>rw_downgrade(9F)</b>	See <b>rwlock(9F)</b>
<b>rw_enter(9F)</b>	See <b>rwlock(9F)</b>
<b>rw_exit(9F)</b>	See <b>rwlock(9F)</b>
<b>rw_init(9F)</b>	See <b>rwlock(9F)</b>
<b>rw_read_locked(9F)</b>	See <b>rwlock(9F)</b>
<b>rw_tryenter(9F)</b>	See <b>rwlock(9F)</b>
<b>rw_tryupgrade(9F)</b>	See <b>rwlock(9F)</b>
<b>rwlock(9F)</b>	readers/writer lock functions
<b>samestr(9F)</b>	See <b>SAMESTR(9F)</b>
<b>scsi_abort(9F)</b>	abort a SCSI command

<b>scsi_alloc_consistent_buf(9F)</b>	allocate an I/O buffer for SCSI DMA
<b>scsi_cname(9F)</b>	decode a SCSI name
<b>scsi_destroy_pkt(9F)</b>	free an allocated SCSI packet and its DMA resource
<b>scsi_dmafree(9F)</b>	See <b>scsi_dmaget(9F)</b>
<b>scsi_dmaget(9F)</b>	SCSI dma utility routines
<b>scsi_dname(9F)</b>	See <b>scsi_cname(9F)</b>
<b>scsi_errmsg(9F)</b>	display a SCSI request sense message
<b>scsi_free_consistent_buf(9F)</b>	free a previously allocated SCSI DMA I/O buffer
<b>scsi_hba_attach(9F)</b>	See <b>scsi_hba_attach_setup(9F)</b>
<b>scsi_hba_attach_setup(9F)</b>	SCSI HBA attach and detach routines
<b>scsi_hba_detach(9F)</b>	See <b>scsi_hba_attach_setup(9F)</b>
<b>scsi_hba_fini(9F)</b>	See <b>scsi_hba_init(9F)</b>
<b>scsi_hba_init(9F)</b>	SCSI Host Bus Adapter system initialization and completion routines
<b>scsi_hba_lookup_capstr(9F)</b>	return index matching capability string
<b>scsi_hba_pkt_alloc(9F)</b>	allocate and free a scsi_pkt structure
<b>scsi_hba_pkt_free(9F)</b>	See <b>scsi_hba_pkt_alloc(9F)</b>
<b>scsi_hba_probe(9F)</b>	default SCSI HBA probe function

<b>scsi_hba_tran_alloc(9F)</b>	allocate and free transport structures
<b>scsi_hba_tran_free(9F)</b>	See <b>scsi_hba_tran_alloc(9F)</b>
<b>scsi_ifgetcap(9F)</b>	get/set SCSI transport capability
<b>scsi_ifsetcap(9F)</b>	See <b>scsi_ifgetcap(9F)</b>
<b>scsi_init_pkt(9F)</b>	prepare a complete SCSI packet
<b>scsi_log(9F)</b>	display a SCSI-device-related message
<b>scsi_mname(9F)</b>	See <b>scsi_cname(9F)</b>
<b>scsi_pktalloc(9F)</b>	SCSI packet utility routines
<b>scsi_pktfree(9F)</b>	See <b>scsi_pktalloc(9F)</b>
<b>scsi_poll(9F)</b>	run a polled SCSI command on behalf of a target driver
<b>scsi_probe(9F)</b>	utility for probing a scsi device
<b>scsi_realloc(9F)</b>	See <b>scsi_pktalloc(9F)</b>
<b>scsi_reset(9F)</b>	reset a SCSI bus or target
<b>scsi_reset_notify(9F)</b>	notify target driver of bus resets
<b>scsi_resfree(9F)</b>	See <b>scsi_pktalloc(9F)</b>
<b>scsi_rname(9F)</b>	See <b>scsi_cname(9F)</b>
<b>scsi_setup_cdb(9F)</b>	setup SCSI command descriptor block (CDB)
<b>scsi_slave(9F)</b>	utility for SCSI target drivers to establish the presence of a target
<b>scsi_sname(9F)</b>	See <b>scsi_cname(9F)</b>

<b>scsi_sync_pkt(9F)</b>	synchronize CPU and I/O views of memory
<b>scsi_transport(9F)</b>	request by a SCSI target driver to start a command
<b>scsi_unprobe(9F)</b>	free resources allocated during initial probing
<b>scsi_unslave(9F)</b>	See <b>scsi_unprobe(9F)</b>
<b>scsi_vu_errmsg(9F)</b>	display a SCSI request sense message
<b>sema_destroy(9F)</b>	See <b>semaphore(9F)</b>
<b>sema_init(9F)</b>	See <b>semaphore(9F)</b>
<b>sema_p(9F)</b>	See <b>semaphore(9F)</b>
<b>sema_p_sig(9F)</b>	See <b>semaphore(9F)</b>
<b>sema_tryop(9F)</b>	See <b>semaphore(9F)</b>
<b>sema_v(9F)</b>	See <b>semaphore(9F)</b>
<b>semaphore(9F)</b>	semaphore functions
<b>sprintf(9F)</b>	format characters in memory
<b>stoi(9F)</b>	convert between an integer and a decimal string
<b>strchr(9F)</b>	find a character in a string
<b>strcmp(9F)</b>	compare two null-terminated strings.
<b>strcpy(9F)</b>	copy a string from one location to another.
<b>strlen(9F)</b>	determine the number of non-null bytes in a string
<b>strlog(9F)</b>	submit messages to the log driver

<b>strncmp(9F)</b>	See <b>strcmp(9F)</b>
<b>strncpy(9F)</b>	See <b>strcpy(9F)</b>
<b>strqget(9F)</b>	get information about a queue or band of the queue
<b>strqset(9F)</b>	change information about a queue or band of the queue
<b>swab(9F)</b>	swap bytes in 16-bit halfwords
<b>testb(9F)</b>	check for an available buffer
<b>timeout(9F)</b>	execute a function after a specified length of time
<b>uiomove(9F)</b>	copy kernel data using uio structure
<b>unbufcall(9F)</b>	cancel a pending bufcall request
<b>unfreezestr(9F)</b>	See <b>freezestr(9F)</b>
<b>unlinkb(9F)</b>	remove a message block from the head of a message
<b>untimeout(9F)</b>	cancel previous timeout function call
<b>ureadc(9F)</b>	add character to a uio structure
<b>uwritec(9F)</b>	remove a character from a uio structure
<b>va_arg(9F)</b>	handle variable argument list
<b>va_copy(9F)</b>	See <b>va_arg(9F)</b>
<b>va_end(9F)</b>	See <b>va_arg(9F)</b>
<b>va_start(9F)</b>	See <b>va_arg(9F)</b>
<b>vcmn_err(9F)</b>	See <b>cmn_err(9F)</b>
<b>vsprintf(9F)</b>	format characters in memory

Intro(9F)

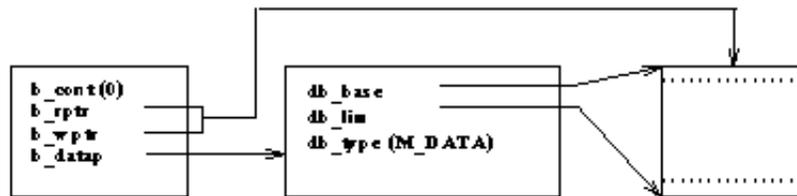
Kernel Functions for Drivers

**wr(9F)**

See **wR(9F)**

<b>NAME</b>	adjmsg – trim bytes from a message
<b>SYNOPSIS</b>	#include <sys/stream.h>  int <b>adjmsg</b> (mblk_t *mp, ssize_t len);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>mp</b> Pointer to the message to be trimmed.  <b>len</b> The number of bytes to be removed.
<b>DESCRIPTION</b>	The <b>adjmsg()</b> function removes bytes from a message. $ len $ (the absolute value of <i>len</i> ) specifies the number of bytes to be removed. The <b>adjmsg()</b> function only trims bytes across message blocks of the same type.  The <b>adjmsg()</b> function finds the maximal leading sequence of message blocks of the same type as that of <i>mp</i> and starts removing bytes either from the head of that sequence or from the tail of that sequence. If <i>len</i> is greater than 0, <b>adjmsg()</b> removes bytes from the start of the first message block in that sequence. If <i>len</i> is less than 0, it removes bytes from the end of the last message block in that sequence.  The <b>adjmsg()</b> function fails if $ len $ is greater than the number of bytes in the maximal leading sequence it finds.  The <b>adjmsg()</b> function may remove any except the first zero-length message block created during adjusting. It may also remove any zero-length message blocks that occur within the scope of $ len $ .
<b>RETURN VALUES</b>	The <b>adjmsg()</b> function returns: 1       Successful completion.  0       An error occurred.
<b>CONTEXT</b>	The <b>adjmsg()</b> function can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>STREAMS Programming Guide</i>

<b>NAME</b>	allocb – allocate a message block														
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  mblk_t *allocb(size_t size, uint_t pri); Architecture independent level 1 (DDI/DKI).</pre>														
<b>PARAMETERS</b>	<p><b>size</b>     The number of bytes in the message block.</p> <p><b>pri</b>       Priority of the request (no longer used).</p>														
<b>DESCRIPTION</b>	<p><b>allocb()</b> tries to allocate a STREAMSmessage block. Buffer allocation fails only when the system is out of memory. If no buffer is available, the <b>bufcall(9F)</b> function can help a module recover from an allocation failure.</p> <p>A STREAMSmessage block is composed of three structures. The first structure is a message block (<b>mblk_t</b>). See <b>msgb(9S)</b>. The <b>mblk_t</b> structure points to a data block structure (<b>dblk_t</b>). See <b>datab(9S)</b>. Together these two structures describe the message type (if applicable) and the size and location of the third structure, the data buffer. The data buffer contains the data for this message block. The allocated data buffer is at least double-word aligned, so it can hold any C data structure.</p> <p>The fields in the <b>mblk_t</b> structure are initialized as follows:</p> <table border="0"> <tr> <td><b>b_cont</b></td> <td>set to NULL</td> </tr> <tr> <td><b>b_rptr</b></td> <td>points to the beginning of the data buffer</td> </tr> <tr> <td><b>b_wptr</b></td> <td>points to the beginning of the data buffer</td> </tr> <tr> <td><b>b_datap</b></td> <td>points to the <b>dblk_t</b> structure</td> </tr> </table> <p>The fields in the <b>dblk_t</b> structure are initialized as follows:</p> <table border="0"> <tr> <td><b>db_base</b></td> <td>points to the first byte of the data buffer</td> </tr> <tr> <td><b>db_lim</b></td> <td>points to the last byte + 1 of the buffer</td> </tr> <tr> <td><b>db_type</b></td> <td>set to M_DATA</td> </tr> </table> <p>The following figure identifies the data structure members that are affected when a message block is allocated.</p>	<b>b_cont</b>	set to NULL	<b>b_rptr</b>	points to the beginning of the data buffer	<b>b_wptr</b>	points to the beginning of the data buffer	<b>b_datap</b>	points to the <b>dblk_t</b> structure	<b>db_base</b>	points to the first byte of the data buffer	<b>db_lim</b>	points to the last byte + 1 of the buffer	<b>db_type</b>	set to M_DATA
<b>b_cont</b>	set to NULL														
<b>b_rptr</b>	points to the beginning of the data buffer														
<b>b_wptr</b>	points to the beginning of the data buffer														
<b>b_datap</b>	points to the <b>dblk_t</b> structure														
<b>db_base</b>	points to the first byte of the data buffer														
<b>db_lim</b>	points to the last byte + 1 of the buffer														
<b>db_type</b>	set to M_DATA														

**RETURN VALUES**

A pointer to the allocated message block of type `M_DATA` on success.

A `NULL` pointer on failure.

**CONTEXT**

`allocb()` can be called from user or interrupt context.

**EXAMPLES****EXAMPLE 1 allocb() Code Sample**

Given a pointer to a queue (*q*) and an error number (*err*), the `send_error()` routine sends an `M_ERROR` type message to the stream head.

If a message cannot be allocated, `NULL` is returned, indicating an allocation failure (line 8). Otherwise, the message type is set to `M_ERROR` (line 10). Line 11 increments the write pointer (`bp->b_wptr`) by the size (one byte) of the data in the message.

A message must be sent up the read side of the stream to arrive at the stream head. To determine whether *q* points to a read queue or to a write queue, the `q->q_flag` member is tested to see if `QREADR` is set (line 13). If it is not set, *q* points to a write queue, and in line 14 the `RD(9F)` function is used to find the corresponding read queue. In line 15, the `putnext(9F)` function is used to send the message upstream, returning 1 if successful.

```

1  send_error(q,err)
2  queue_t *q;
3  unsigned char err;
4  {
5  mblk_t *bp;
6
7  if ((bp = allocb(1, BPRI_HI)) == NULL) /* allocate msg. block */
8      return(0);
9
10 bp->b_datap->db_type = M_ERROR;          /* set msg type to M_ERROR */
11 *bp->b_wptr++ = err;                    /* increment write pointer */
12
13 if (!(q->q_flag & QREADR))              /* if not read queue */
14     q = RD(q);                          /* get read queue */
15 putnext(q, bp);                         /* send message upstream */
16 return(1);
17 }

```

**SEE ALSO**

`RD(9F)`, `bufcall(9F)`, `esballoc(9F)`, `esbcall(9F)`, `putnext(9F)`,  
`testb(9F)`, `datadb(9S)`, `msgb(9S)`

*Writing Device Drivers*

*STREAMS Programming Guide*

**NOTES**

The *pri* argument is no longer used, but is retained for compatibility with existing drivers.

<b>NAME</b>	anocancel – prevent cancellation of asynchronous I/O request
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int anocancel( );</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>anocancel()</b> should be used by drivers that do not support canceling asynchronous I/O requests. <b>anocancel()</b> is passed as the driver cancel routine parameter to <b>aphysio(9F)</b> .
<b>RETURN VALUES</b>	<b>anocancel()</b> returns ENXIO.
<b>SEE ALSO</b>	<b>aread(9E)</b> , <b>awrite(9E)</b> , <b>aphysio(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	aphysio – perform asynchronous physical I/O
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt; #include &lt;sys/uio.h&gt; #include &lt;sys/aio_req.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p><b>int</b> <b>aphysio</b>(<b>int</b> (*<i>strat</i>)(<b>struct</b> buf *), <b>int</b> (*<i>cancel</i>)(<b>struct</b> buf *), <b>dev_t</b> <i>dev</i>, <b>int</b> <i>rw</i>, <b>void</b> (*<i>mincnt</i>)(<b>struct</b> buf *), <b>struct</b> aio_req *<i>aio_reqp</i>);</p>
<b>PARAMETERS</b>	<p><b>strat</b>            Pointer to device strategy routine.</p> <p><b>cancel</b>           Pointer to driver cancel routine. Used to cancel a submitted request. The driver must pass the address of the function <b>anocancel</b>(9F) because cancellation is not supported.</p> <p><b>dev</b>                The device number.</p> <p><b>rw</b>                Read/write flag. This is either <b>B_READ</b> when reading from the device or <b>B_WRITE</b> when writing to the device.</p> <p><b>mincnt</b>            Routine which bounds the maximum transfer unit size.</p> <p><b>aio_reqp</b>          Pointer to the <b>aio_req</b>(9S) structure which describes the user I/O request.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p><b>aphysio()</b> performs asynchronous I/O operations between the device and the address space described by <i>aio_reqp</i>→<i>aio_uio</i>.</p> <p>Prior to the start of the transfer, <b>aphysio()</b> verifies the requested operation is valid. It then locks the pages involved in the I/O transfer so they can not be paged out. The device strategy routine, <i>strat</i>, is then called one or more times to perform the physical I/O operations. <b>aphysio()</b> does not wait for each transfer to complete, but returns as soon as the necessary requests have been made.</p> <p><b>aphysio()</b> calls <i>mincnt</i> to bound the maximum transfer unit size to a sensible default for the device and the system. Drivers which do not provide their own local <i>mincnt</i> routine should call <b>aphysio()</b> with <b>minphys</b>(9F). <b>minphys</b>(9F) is the system <i>mincnt</i> routine. <b>minphys</b>(9F) ensures the transfer size does not exceed any system limits.</p>

If a driver supplies a local *mincnt* routine, this routine should perform the following actions:

- If *bp→b\_bcount* exceeds a device limit, set *bp→b\_bcount* to a value supported by the device.
- Call **minphys**(9F) to ensure that the driver does not circumvent additional system limits.

**RETURN VALUES**

**aphysio()** returns:

- 1       Upon success.
- 2       Upon failure.

**CONTEXT**

**aphysio()** can be called from user context only.

**SEE ALSO**

**aread**(9E), **awrite**(9E), **strategy**(9E), **anocancel**(9F), **biodone**(9F), **biowait**(9F), **minphys**(9F), **physio**(9F), **aio\_req**(9S), **buf**(9S), **uio**(9S)

*Writing Device Drivers*

**WARNINGS**

It is the driver's responsibility to call **biodone**(9F) when the transfer is complete.

**BUGS**

Cancellation is not supported in this release. The address of the function **anocancel**(9F) must be used as the *cancel* argument.

<b>NAME</b>	ASSERT, assert – expression verification
<b>SYNOPSIS</b>	<pre>#include &lt;sys/debug.h&gt;  void ASSERT(<i>EX</i>);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b><i>EX</i></b> boolean expression.
<b>DESCRIPTION</b>	<b>ASSERT()</b> is a macro which checks to see if the expression <i>EX</i> is true. If it is not, then <b>ASSERT()</b> causes an error message to be logged to the console and the system to panic. <b>ASSERT()</b> works only if the preprocessor symbol <b>DEBUG</b> is defined.
<b>CONTEXT</b>	<b>ASSERT()</b> can be used from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	backq – get pointer to the queue behind the current queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  queue_t *backq(queue_t *cq);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>cq</b>     The pointer to the current queue. <code>queue_t</code> is an alias for the <code>queue(9S)</code> structure.</p>
<b>DESCRIPTION</b>	<p><b>backq()</b> returns a pointer to the queue preceding <i>cq</i> (the current queue). If <i>cq</i> is a read queue, <b>backq()</b> returns a pointer to the queue downstream from <i>cq</i>, unless it is the stream end. If <i>cq</i> is a write queue, <b>backq()</b> returns a pointer to the next queue upstream from <i>cq</i>, unless it is the stream head.</p>
<b>RETURN VALUES</b>	<p>If successful, <b>backq()</b> returns a pointer to the queue preceding the current queue. Otherwise, it returns <code>NULL</code>.</p>
<b>CONTEXT</b>	<p><b>backq()</b> can be called from user or interrupt context.</p>
<b>SEE ALSO</b>	<p><code>queue(9S)</code></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	bcanput – test for flow control in specified priority band
<b>SYNOPSIS</b>	#include <sys/stream.h>  int bcanput(queue_t *q, unsigned char pri);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the message queue. <b>pri</b> Message priority.
<b>DESCRIPTION</b>	<b>bcanput()</b> searches through the stream (starting at <i>q</i> ) until it finds a queue containing a service routine where the message can be enqueued, or until it reaches the end of the stream. If found, the queue containing the service routine is tested to see if there is room for a message of priority <i>pri</i> in the queue.  If <i>pri</i> is 0, <b>bcanput()</b> is equivalent to a call with <b>canput(9F)</b> .  <i>canputnext(q)</i> and <i>bcanputnext(q, pri)</i> should always be used in preference to <i>canput(q→q_next)</i> and <i>bcanput(q→q_next, pri)</i> respectively.
<b>RETURN VALUES</b>	1 If a message of priority <i>pri</i> can be placed on the queue. 0 If the priority band is full.
<b>CONTEXT</b>	<b>bcanput()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>bcanputnext(9F)</b> , <b>canput(9F)</b> , <b>canputnext(9F)</b> , <b>putbq(9F)</b> , <b>putnext(9F)</b>  <i>Writing Device Drivers</i>  <i>STREAMS Programming Guide</i>
<b>WARNINGS</b>	Drivers are responsible for both testing a queue with <b>bcanput()</b> and refraining from placing a message on the queue if <b>bcanput()</b> fails.

<b>NAME</b>	bcmp – compare two byte arrays
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt;  int bcmp(const void *s1, const void *s2, size_t len);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>s1</b> Pointer to the first character string.</p> <p><b>s2</b> Pointer to the second character string.</p> <p><b>len</b> Number of bytes to be compared.</p>
<b>DESCRIPTION</b>	<b>bcmp()</b> compares two byte arrays of length <i>len</i> .
<b>RETURN VALUES</b>	<b>bcmp()</b> returns 0 if the arrays are identical, or 1 if they are not.
<b>CONTEXT</b>	<b>bcmp()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>strcmp(9F)</b> <i>Writing Device Drivers</i>
<b>NOTES</b>	Unlike <b>strcmp(9F)</b> , <b>bcmp()</b> does not terminate when it encounters a null byte.

<b>NAME</b>	bcopy – copy data between address locations in the kernel
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt;  void bcopy(const void *from, void *to, size_t bcount);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>from</b> Source address from which the copy is made.</p> <p><b>to</b> Destination address to which copy is made.</p> <p><b>bcount</b> The number of bytes moved.</p>
<b>DESCRIPTION</b>	<p><b>bcopy()</b> copies <i>bcount</i> bytes from one kernel address to another. If the input and output addresses overlap, the command executes, but the results may not be as expected.</p> <p>Note that <b>bcopy()</b> should never be used to move data in or out of a user buffer, because it has no provision for handling page faults. The user address space can be swapped out at any time, and <b>bcopy()</b> always assumes that there will be no paging faults. If <b>bcopy()</b> attempts to access the user buffer when it is swapped out, the system will panic. It is safe to use <b>bcopy()</b> to move data within kernel space, since kernel space is never swapped out.</p>
<b>CONTEXT</b>	<b>bcopy()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Copying data between address locations in the kernel:</p> <p>An I/O request is made for data stored in a RAM disk. If the I/O operation is a read request, the data is copied from the RAM disk to a buffer (line 8). If it is a write request, the data is copied from a buffer to the RAM disk (line 15). <b>bcopy()</b> is used since both the RAM disk and the buffer are part of the kernel address space.</p> <pre> 1 #define RAMDNBLK 1000                /* blocks in the RAM disk */ 2 #define RAMDBSIZ 512                /* bytes per block */ 3 char ramdblks[RAMDNBLK][RAMDBSIZ]; /* blocks forming RAM 4                                     /* disk 5 6 7 8     bcopy(&amp;ramdblks[bp-&gt;b_blkno][0], bp-&gt;b_un.b_addr, 9         bp-&gt;b_bcount); 10</pre>

```
11 else                               /* else write request, */
12                                     /* copy data from a */
13                                     /* system buffer to RAM disk */
14                                     /* data block */
15     bcopy(bp->b_un.b_addr, &ramdblk[bp->b_blkno][0],
16           bp->b_bcount);
```

**SEE ALSO** `copyin(9F)`, `copyout(9F)`

*Writing Device Drivers*

**WARNINGS** The *from* and *to* addresses must be within the kernel space. No range checking is done. If an address outside of the kernel space is selected, the driver may corrupt the system in an unpredictable way.

<b>NAME</b>	bioclone – clone another buffer
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  struct buf *bioclone(struct buf *bp, off_t off, size_t len, dev_t dev, daddr_t blkno, int (*iodone, struct buf *, struct buf *bp_mem, int sleepflag);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>bp</b> Pointer to the <b>buf(9S)</b> structure describing the original I/O request.</p> <p><b>off</b> Offset within original I/O request where new I/O request should start.</p> <p><b>len</b> Length of the I/O request.</p> <p><b>dev</b> Device number.</p> <p><b>blkno</b> Block number on device.</p> <p><b>iodone</b> Specific <b>biodone(9F)</b> routine.</p> <p><b>bp_mem</b> Pointer to a buffer structure to be filled in or NULL.</p> <p><b>sleepflag</b> Determines whether caller can sleep for memory. Possible flags are <b>KM_SLEEP</b> to allow sleeping until memory is available, or <b>KM_NOSLEEP</b> to return NULL immediately if memory is not available.</p>
<b>DESCRIPTION</b>	<p><b>bioclone()</b> returns an initialized buffer to perform I/O to a portion of another buffer. The new buffer will be set up to perform I/O to the range within the original I/O request specified by the parameters <i>off</i> and <i>len</i>. An offset 0 starts the new I/O request at the same address as the original request. <i>off + len</i> must not exceed <i>b_bcount</i>, the length of the original request. The device number <i>dev</i> specifies the device to which the buffer is to perform I/O. <i>blkno</i> is the block number on device. It will be assigned to the <i>b_blkno</i> field of the cloned buffer structure. <i>iodone</i> lets the driver identify a specific <b>biodone(9F)</b> routine to be called by the driver when the I/O is complete. <i>bp_mem</i> determines from where the space for the buffer should be allocated. If <i>bp_mem</i> is NULL, <b>bioclone()</b> will allocate a new buffer using <b>getrbuf(9F)</b>. If <i>sleepflag</i> is set to <b>KM_SLEEP</b>, the driver may sleep until space is freed up. If <i>sleepflag</i> is set to <b>KM_NOSLEEP</b>, the driver will not sleep. In either case, a pointer to the allocated space is returned or NULL to indicate that no space was available. After the transfer is completed, the buffer has to be freed using <b>freerbuf(9F)</b>. If <i>bp_mem</i> is not NULL, it will be used as the space for the</p>

buffer structure. The driver has to ensure that *bp\_mem* is initialized properly either using `getrbuf(9F)` or `bioinit(9F)`.

If the original buffer is mapped into the kernel virtual address space using `bp_mapin(9F)` before calling `bp_clone()`, a clone buffer will share the kernel mapping of the original buffer. An additional `bp_mapin()` to get a kernel mapping for the clone buffer is not necessary.

The driver has to ensure that the original buffer is not freed while any of the clone buffers is still performing I/O. The `biodone()` function has to be called on all clone buffers before it is called on the original buffer.

#### RETURN VALUES

The `bioclone()` function returns a pointer to the initialized buffer header, or `NULL` if no space is available.

#### CONTEXT

`bioclone()` can be called from user or interrupt context. Drivers must not allow `bioclone()` to sleep if called from an interrupt routine.

#### EXAMPLES

**EXAMPLE 1** : Using `bioclone()`

A device driver can use `bioclone()` for disk striping. For each disk in the stripe, a clone buffer is created which performs I/O to a portion of the original buffer.

```
static int
stripe_strategy(struct buf *bp)
{
    ...
    bp_orig = bp;
    bp_1 = bioclone(bp_orig, 0, size_1, dev_1, blkno_1,
                   stripe_done, NULL, KM_SLEEP);
    fragment++;
    ...
    bp_n = bioclone(bp_orig, offset_n, size_n, dev_n,
                   blkno_n, stripe_done, NULL, KM_SLEEP);
    fragment++;
    /* submit bp_1 ... bp_n to device */
    xxstrategy(bp_x);
    return (0);
}

static uint_t
xxintr(caddr_t arg)
{
    ...
    /*
     * get bp of completed subrequest. biodone(9F) will
     * call stripe_done()
     */
    biodone(bp);
    return (0);
}

static int
```

```
stripe_done(struct buf *bp)
{
    ...
    freerbuf(bp);
    fragment--;
    if (fragment == 0) {
        /* get bp_orig */
        biodone(bp_orig);
    }
    return (0);
}
```

**SEE ALSO** [biodone\(9F\)](#), [bp\\_mapin\(9F\)](#), [freerbuf\(9F\)](#), [getrbuf\(9F\)](#), [buf\(9S\)](#)

*Writing Device Drivers*

<b>NAME</b>	biodone – release buffer after buffer I/O transfer and notify blocked threads
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt;  void biodone(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to a <b>buf(9S)</b> structure.
<b>DESCRIPTION</b>	<p><b>biodone()</b> notifies blocked processes waiting for the I/O to complete, sets the <b>B_DONE</b> flag in the <b>b_flags</b> field of the <b>buf(9S)</b> structure, and releases the buffer if the I/O is asynchronous. <b>biodone()</b> is called by either the driver interrupt or <b>strategy(9E)</b> routines when a buffer I/O request is complete.</p> <p><b>biodone()</b> provides the capability to call a completion routine if <b>bp</b> describes a kernel buffer. The address of the routine is specified in the <b>b_iodone</b> field of the <b>buf(9S)</b> structure. If such a routine is specified, <b>biodone()</b> calls it and returns without performing any other actions. Otherwise, it performs the steps above.</p>
<b>CONTEXT</b>	<b>biodone()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b></p> <p>Generally, the first validation test performed by any block device <b>strategy(9E)</b> routine is a check for an end-of-file (EOF) condition. The <b>strategy(9E)</b> routine is responsible for determining an EOF condition when the device is accessed directly. If a <b>read(2)</b> request is made for one block beyond the limits of the device (line 10), it will report an EOF condition. Otherwise, if the request is outside the limits of the device, the routine will report an error condition. In either case, report the I/O operation as complete (line 27).</p> <pre>1  #define RAMDNBLK 1000 /* Number of blocks in RAM disk */ 2  #define RAMDBSIZ 512 /* Number of bytes per block */ 3  char ramdbls[RAMDNBLK][RAMDBSIZ]; /* Array containing RAM disk */ 4 5  static int 6  ramdstrategy(struct buf *bp) 7  { 8      daddr_t blkno = bp-&gt;b_blkno; /* get block number */ 9 10     if ((blkno &lt; 0)    (blkno &gt;= RAMDNBLK)) { 11         /* 12          * If requested block is outside RAM disk</pre>

```

13         * limits, test for EOF which could result
14         * from a direct (physio) request.
15         */
16         if ((blkno == RAMDNBLK) && (bp->b_flags & B_READ)) {
17             /*
18              * If read is for block beyond RAM disk
19              * limits, mark EOF condition.
20              */
21             bp->b_resid = bp->b_bcount; /* compute return value */
22
23             } else { /* I/O attempt is beyond */
24                 bp->b_error = ENXIO; /* limits of RAM disk */
25                 bp->b_flags |= B_ERROR; /* return error */
26             }
27         biodone(bp); /* mark I/O complete (B_DONE) */
28         /*
29          * Wake any processes awaiting this I/O
30          * or release buffer for asynchronous
31          * (B_ASYNC) request.
32          */
33         return (0);
34     }
    ...

```

**SEE ALSO** [read\(2\)](#), [strategy\(9E\)](#), [biowait\(9F\)](#), [ddi\\_add\\_intr\(9F\)](#), [delay\(9F\)](#), [timeout\(9F\)](#), [untimeout\(9F\)](#), [buf\(9S\)](#)

*Writing Device Drivers*

**WARNINGS** After calling **biodone()**, *bp* is no longer available to be referred to by the driver. If the driver makes any reference to *bp* after calling **biodone()**, a panic may result.

**NOTES** Drivers that use the `b_iodone` field of the [buf\(9S\)](#) structure to specify a substitute completion routine should save the value of `b_iodone` before changing it, and then restore the old value before calling **biodone()** to release the buffer.

<b>NAME</b>	bioerror – indicate error in buffer header
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt; #include &lt;sys/ddi.h&gt;  void bioerror(struct buf *bp, int error);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>bp</b> Pointer to the <b>buf(9S)</b> structure describing the transfer.</p> <p><b>error</b> Error number to be set, or zero to clear an error indication.</p>
<b>DESCRIPTION</b>	<p>If <i>error</i> is non-zero, <b>bioerror()</b> indicates an error has occurred in the <b>buf(9S)</b> structure. A subsequent call to <b>geterror(9F)</b> will return <i>error</i>.</p> <p>If <i>error</i> is 0, the error indication is cleared and a subsequent call to <b>geterror(9F)</b> will return 0.</p>
<b>CONTEXT</b>	<b>bioerror()</b> can be called from any context.
<b>SEE ALSO</b>	<b>strategy(9E)</b> , <b>geterror(9F)</b> , <b>getrbuf(9F)</b> , <b>buf(9S)</b>

<b>NAME</b>	biofini – uninitialized a buffer structure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void biofini(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the buffer header structure.
<b>DESCRIPTION</b>	The <b>biofini()</b> function uninitialized a <b>buf(9S)</b> structure. If a buffer structure has been allocated and initialized using <b>kmem_alloc(9F)</b> and <b>bioinit(9F)</b> it needs to be uninitialized using <b>biofini()</b> before calling <b>kmem_free(9F)</b> . It is not necessary to call <b>biofini()</b> before freeing a buffer structure using <b>freerbuf(9F)</b> because <b>freerbuf()</b> will call <b>biofini()</b> directly.
<b>CONTEXT</b>	The <b>biofini()</b> function can be called from any context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b>   Using <b>biofini()</b></p> <pre>struct buf *bp = kmem_alloc(biosize(), KM_SLEEP); bioinit(bp); /* use buffer */ biofini(bp); kmem_free(bp, biosize());</pre>
<b>SEE ALSO</b>	<b>bioinit(9F)</b> , <b>bioreset(9F)</b> , <b>biosize(9F)</b> , <b>freerbuf(9F)</b> , <b>kmem_alloc(9F)</b> , <b>kmem_free(9F)</b> , <b>buf(9S)</b>
	<i>Writing Device Drivers</i>

<b>NAME</b>	bioinit – initialize a buffer structure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void bioinit(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the buffer header structure.
<b>DESCRIPTION</b>	The <b>bioinit()</b> function initializes a <b>buf(9S)</b> structure. A buffer structure contains state information which has to be initialized if the memory for the buffer was allocated using <b>kmem_alloc(9F)</b> . This is not necessary for a buffer allocated using <b>getrbuf(9F)</b> because <b>getrbuf()</b> will call <b>bioinit()</b> directly.
<b>CONTEXT</b>	The <b>bioinit()</b> function can be called from any context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using <b>bioinit()</b></p> <pre>struct buf *bp = kmem_alloc(biosize(), KM_SLEEP); bioinit(bp); /* use buffer */</pre>
<b>SEE ALSO</b>	<p><b>biofini(9F)</b>, <b>bioreset(9F)</b>, <b>biosize(9F)</b>, <b>getrbuf(9F)</b>, <b>kmem_alloc(9F)</b>, <b>buf(9S)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	biomodified – check if a buffer is modified
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  intbiomodified(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the buffer header structure.
<b>DESCRIPTION</b>	<p>The <b>biomodified()</b> function returns status to indicate if the buffer is modified. The <b>biomodified()</b> function is only supported for paged- I/O request, that is the <code>B_PAGEIO</code> flag must be set in the <code>b_flags</code> field of the <b>buf(9S)</b> structure. The <b>biomodified()</b> function will check the memory pages associated with this buffer whether the Virtual Memory system's modification bit is set. If at least one of these pages is modified, the buffer is indicated as modified. A filesystem will mark the pages <code>unmodified</code> when it writes the pages to the backing store. The <b>biomodified()</b> function can be used to detect any modifications to the memory pages while I/O is in progress.</p> <p>A device driver can use <b>biomodified()</b> for disk mirroring. An application is allowed to mmap a file which can reside on a disk which is mirrored by multiple submirrors. If the file system writes the file to the backing store, it is written to all submirrors in parallel. It must be ensured that the copies on all submirrors are identical. The <b>biomodified()</b> function can be used in the device driver to detect any modifications to the buffer by the user program during the time the buffer is written to multiple submirrors.</p>
<b>RETURN VALUES</b>	<p>The <b>biomodified()</b> function returns the following values:</p> <ul style="list-style-type: none"> <li>1       Buffer is modified.</li> <li>0       Buffer is not modified.</li> <li>-1      Buffer is not used for paged I/O request.</li> </ul>
<b>CONTEXT</b>	<b>biomodified()</b> can be called from any context.
<b>SEE ALSO</b>	<b>bp_mapin(9F)</b> , <b>buf(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	bioreset – reuse a private buffer header after I/O is complete
<b>SYNOPSIS</b>	<pre>#include &lt;sys/buf.h&gt; #include &lt;sys/ddi.h&gt;  void bioreset(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI)
<b>PARAMETERS</b>	<b>bp</b> Pointer to the <b>buf(9S)</b> structure.
<b>DESCRIPTION</b>	<b>bioreset()</b> is used by drivers that allocate private buffers with <b>getrbuf(9F)</b> or <b>kmem_alloc(9F)</b> and want to reuse them in multiple transfers before freeing them with <b>freerbuf(9F)</b> or <b>kmem_free(9F)</b> . <b>bioreset()</b> resets the buffer header to the state it had when initially allocated by <b>getrbuf()</b> or initialized by <b>bioinit(9F)</b> .
<b>CONTEXT</b>	<b>bioreset()</b> can be called from any context.
<b>SEE ALSO</b>	<b>strategy(9E)</b> , <b>bioinit(9F)</b> , <b>biofini(9F)</b> , <b>freerbuf(9F)</b> , <b>getrbuf(9F)</b> , <b>kmem_alloc(9F)</b> , <b>kmem_free(9F)</b> , <b>buf(9S)</b>
<b>NOTES</b>	<i>bp</i> must not describe a transfer in progress.

<b>NAME</b>	biosize – returns size of a buffer structure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  size_t biosize(void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	The <b>biosize()</b> function returns the size in bytes of the <b>buf(9S)</b> structure. The <b>biosize()</b> function is used by drivers in combination with <b>kmem_alloc(9F)</b> and <b>bioinit(9F)</b> to allocate buffer structures embedded in other data structures.
<b>CONTEXT</b>	The <b>biosize()</b> function can be called from any context.
<b>SEE ALSO</b>	<b>biofini(9F)</b> , <b>bioinit(9F)</b> , <b>getrbuf(9F)</b> , <b>kmem_alloc(9F)</b> , <b>buf(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	biowait – suspend processes pending completion of block I/O
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt;  int biowait(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the <code>buf</code> structure describing the transfer.
<b>DESCRIPTION</b>	<p>Drivers allocating their own <code>buf</code> structures with <code>getrbuf(9F)</code> can use the <code>biowait()</code> function to suspend the current thread and wait for completion of the transfer.</p> <p>Drivers must call <code>biodone(9F)</code> when the transfer is complete to notify the thread blocked by <code>biowait()</code>. <code>biodone()</code> is usually called in the interrupt routine.</p>
<b>RETURN VALUES</b>	<p>0 Upon success</p> <p><code>non-zero</code> Upon I/O failure. <code>biowait()</code> calls <code>geterror(9F)</code> to retrieve the error number which it returns.</p>
<b>CONTEXT</b>	<code>biowait()</code> can be called from user context only.
<b>SEE ALSO</b>	<code>biodone(9F)</code> , <code>geterror(9F)</code> , <code>getrbuf(9F)</code> , <code>buf(9S)</code> <i>Writing Device Drivers</i>

<b>NAME</b>	bp_mapin - allocate virtual address space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt;  void bp_mapin(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the buffer header structure.
<b>DESCRIPTION</b>	<p><b>bp_mapin()</b> is used to map virtual address space to a page list maintained by the buffer header during a paged- I/O request. <b>bp_mapin()</b> allocates system virtual address space, maps that space to the page list, and returns the starting address of the space in the <code>bp-&gt;b_un.b_addr</code> field of the <b>buf(9S)</b> structure. Virtual address space is then deallocated using the <b>bp_mapout(9F)</b> function.</p> <p>If a null page list is encountered, <b>bp_mapin()</b> returns without allocating space and no mapping is performed.</p>
<b>CONTEXT</b>	<b>bp_mapin()</b> can be called from user and kernel contexts.
<b>SEE ALSO</b>	<b>bp_mapout(9F)</b> , <b>buf(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	bp_mapout – deallocate virtual address space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt;  void bp_mapout(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the buffer header structure.
<b>DESCRIPTION</b>	<b>bp_mapout()</b> deallocates system virtual address space allocated by a previous call to <b>bp_mapin(9F)</b> . <b>bp_mapout()</b> should only be called on buffers which have been allocated and are owned by the device driver. It must not be called on buffers passed to the driver through the <b>strategy(9E)</b> entry point (for example a filesystem). Because <b>bp_mapin(9F)</b> does not keep a reference count, <b>bp_mapout()</b> will wipe out any kernel mapping that a layer above the device driver might rely on.
<b>CONTEXT</b>	<b>bp_mapout()</b> can be called from user context only.
<b>SEE ALSO</b>	<b>strategy(9E)</b> , <b>bp_mapin(9F)</b> , <b>buf(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	btop – convert size in bytes to size in pages (round down)
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  unsigned long btop(unsigned long numbytes);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>numbytes</b> Number of bytes.
<b>DESCRIPTION</b>	<b>btop()</b> returns the number of memory pages that are contained in the specified number of bytes, with downward rounding in the case that the byte count is not a page multiple. For example, if the page size is 2048, then <b>btop(4096)</b> returns 2, and <b>btop(4097)</b> returns 2 as well. <b>btop(0)</b> returns 0.
<b>RETURN VALUES</b>	The return value is always the number of pages. There are no invalid input values, and therefore no error return values.
<b>CONTEXT</b>	<b>btop()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>btopr(9F)</b> , <b>ddi_btop(9F)</b> , <b>ptob(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	btopr – convert size in bytes to size in pages (round up)
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  unsigned long btopr(unsigned long <i>numbytes</i>);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b><i>numbytes</i></b> Number of bytes.
<b>DESCRIPTION</b>	<b>btopr()</b> returns the number of memory pages contained in the specified number of bytes memory, rounded up to the next whole page. For example, if the page size is 2048, then <b>btopr(4096)</b> returns 2, and <b>btopr(4097)</b> returns 3.
<b>RETURN VALUES</b>	The return value is always the number of pages. There are no invalid input values, and therefore no error return values.
<b>CONTEXT</b>	<b>btopr()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>btop(9F)</b> , <b>ddi_btopr(9F)</b> , <b>ptob(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	bufcall – call a function when a buffer becomes available
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/stream.h&gt;</pre> <p>bufcall_id_t bufcall(size_t size, uint_t pri, void (*func)(void *arg), void *arg);</p>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>size</b>    Number of bytes required for the buffer.</p> <p><b>pri</b>     Priority of the <code>allocb(9F)</code> allocation request (not used).</p> <p><b>func</b>    Function or driver routine to be called when a buffer becomes available.</p> <p><b>arg</b>     Argument to the function to be called when a buffer becomes available.</p>
<b>DESCRIPTION</b>	<b>bufcall()</b> serves as a <code>timeout(9F)</code> call of indeterminate length. When a buffer allocation request fails, <b>bufcall()</b> can be used to schedule the routine <i>func</i> , to be called with the argument <i>arg</i> when a buffer becomes available. <i>func</i> may call <code>allocb()</code> or it may do something else.
<b>RETURN VALUES</b>	If successful, <b>bufcall()</b> returns a <code>bufcall</code> ID that can be used in a call to <code>unbufcall()</code> to cancel the request. If the <b>bufcall()</b> scheduling fails, <i>func</i> is never called and 0 is returned.
<b>CONTEXT</b>	<b>bufcall()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b>    Calling a function when a buffer becomes available:</p> <p>The purpose of this <code>srv(9E)</code> service routine is to add a header to all <code>M_DATA</code> messages. Service routines must process all messages on their queues before returning, or arrange to be rescheduled</p> <p>While there are messages to be processed (line 13), check to see if it is a high priority message or a normal priority message that can be sent on (line 14). Normal priority message that cannot be sent are put back on the message queue (line 34). If the message was a high priority one, or if it was normal priority and <code>canputnext(9F)</code> succeeded, then send all but <code>M_DATA</code> messages to the next module with <code>putnext(9F)</code> (line 16).</p> <p>For <code>M_DATA</code> messages, try to allocate a buffer large enough to hold the header (line 18). If no such buffer is available, the service routine must be rescheduled</p>

for a time when a buffer is available. The original message is put back on the queue (line 20) and `bufcall` (line 21) is used to attempt the rescheduling. It will succeed if the rescheduling succeeds, indicating that `qenable` will be called subsequently with the argument `q` once a buffer of the specified size (`sizeof (struct hdr)`) becomes available. If it does, `qenable(9F)` will put `q` on the list of queues to have their service routines called. If `bufcall()` fails, `timeout(9F)` (line 22) is used to try again in about a half second.

If the buffer allocation was successful, initialize the header (lines 25–28), make the message type `M_PROTO` (line 29), link the `M_DATA` message to it (line 30), and pass it on (line 31).

Note that this example ignores the bookkeeping needed to handle `bufcall()` and `timeout(9F)` cancellation for ones that are still outstanding at close time.

```

1  struct hdr {
2      unsigned int h_size;
3      int         h_version;
4  };
5
6  void xxxsrv(q)
7      queue_t *q;
8  {
9      mblk_t *bp;
10     mblk_t *mp;
11     struct hdr *hp;
12
13     while ((mp = getq(q)) != NULL) { /* get next message */
14         if (mp->b_datap->db_type >= QPCTL || /* if high priority */
15             canputnext(q)) { /* normal & can be passed */
16             if (mp->b_datap->db_type != M_DATA)
17                 putnext(q, mp); /* send all but M_DATA */
18             else {
19                 bp = allocb(sizeof(struct hdr), BPRI_LO);
20                 if (bp == NULL) { /* if unsuccessful */
21                     putbq(q, mp); /* put it back */
22                     if (!bufcall(sizeof(struct hdr), BPRI_LO,
23                                 qenable, q)) /* try to reschedule */
24                         timeout(qenable, q, drv_usectohz(500000));
25                     return (0);
26                 }
27                 hp = (struct hdr *)bp->b_wptr;
28                 hp->h_size = msgdsize(mp); /* initialize header */
29                 hp->h_version = 1;
30                 bp->b_wptr += sizeof(struct hdr);
31                 bp->b_datap->db_type = M_PROTO; /* make M_PROTO */
32                 bp->b_cont = mp; /* link it */
33                 putnext(q, bp); /* pass it on */
34             }
35         } else { /* normal priority, canputnext failed */
36             putbq(q, mp); /* put back on the message queue */
37             return (0);
38         }
39     }
40 }

```

```
38 }
```

**SEE ALSO**

`srv(9E)`, `allocb(9F)`, `canputnext(9F)`, `esballoc(9F)`, `esbcall(9F)`,  
`putnext(9F)`, `qenable(9F)`, `testb(9F)`, `timeout(9F)`, `unbufcall(9F)`

*Writing Device Drivers*

*STREAMS Programming Guide*

**WARNINGS**

Even when *func* is called by `bufcall()`, `allocb(9F)` can fail if another module or driver had allocated the memory before *func* was able to call `allocb(9F)`.

<b>NAME</b>	<b>bzero</b> – clear memory for a given number of bytes
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt;  void <b>bzero</b>(void *<i>addr</i>, size_t <i>bytes</i>);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b><i>addr</i></b> Starting virtual address of memory to be cleared.</p> <p><b><i>bytes</i></b> The number of bytes to clear starting at <i>addr</i>.</p>
<b>DESCRIPTION</b>	<b>bzero()</b> clears a contiguous portion of memory by filling it with zeros.
<b>CONTEXT</b>	<b>bzero()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>bcopy(9F)</b> , <b>clrbuf(9F)</b> , <b>kmem_zalloc(9F)</b> <i>Writing Device Drivers</i>
<b>WARNINGS</b>	The address range specified must be within the kernel space. No range checking is done. If an address outside of the kernel space is selected, the driver may corrupt the system in an unpredictable way.

<b>NAME</b>	canput – test for room in a message queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  int canput(queue_t *q);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the message queue.
<b>DESCRIPTION</b>	<b>canput()</b> searches through the stream (starting at <i>q</i> ) until it finds a queue containing a service routine where the message can be enqueued, or until it reaches the end of the stream. If found, the queue containing the service routine is tested to see if there is room for a message in the queue.  canputnext( <i>q</i> ) and bcanputnext( <i>q</i> , <i>pri</i> ) should always be used in preference to canput( <i>q</i> → <i>q_next</i> ) and bcanput( <i>q</i> → <i>q_next</i> , <i>pri</i> ) respectively.
<b>RETURN VALUES</b>	1        If the message queue is not full.  0        If the queue is full.
<b>CONTEXT</b>	<b>canput()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	bcanput(9F), bcanputnext(9F), canputnext(9F), putbq(9F), putnext(9F)  <i>Writing Device Drivers</i>  <i>STREAMS Programming Guide</i>
<b>WARNINGS</b>	Drivers are responsible for both testing a queue with <b>canput()</b> and refraining from placing a message on the queue if <b>canput()</b> fails.

<b>NAME</b>	canputnext, bcanputnext – test for room in next module’s message queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  int canputnext(queue_t * q); int bcanputnext(queue_t * q, unsigned char pri);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to a message queue belonging to the invoking module. <b>pri</b> Minimum priority level.
<b>DESCRIPTION</b>	The invocation canputnext( q ); is an atomic equivalent of the canput( q →q_next ); routine. That is, the STREAMS framework provides whatever mutual exclusion is necessary to insure that dereferencing q through its q_next field and then invoking canput(9F) proceeds without interference from other threads.  bcanputnext( q , pri ); is the equivalent of the bcanput( q →q_next , pri ); routine.  canputnext( q ); and bcanputnext( q , pri ); should always be used in preference to canput( q →q_next ); and bcanput( q → q_next , pri ); respectively.  See canput(9F) and bcanput(9F) for further details.
<b>RETURN VALUES</b>	1        If the message queue is not full. 0        If the queue is full.
<b>CONTEXT</b>	canputnext() and bcanputnext() can be called from user or interrupt context.
<b>WARNINGS</b>	Drivers are responsible for both testing a queue with canputnext() or bcanputnext() and refraining from placing a message on the queue if the queue is full.
<b>SEE ALSO</b>	bcanput(9F) , canput(9F)  <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	clrbuf – erase the contents of a buffer
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt;  void clrbuf(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the <b>buf(9S)</b> structure.
<b>DESCRIPTION</b>	<b>clrbuf()</b> zeros a buffer and sets the <b>b_resid</b> member of the <b>buf(9S)</b> structure to 0. Zeros are placed in the buffer starting at <b>bp→b_un.b_addr</b> for a length of <b>bp→b_bcount</b> bytes. <b>b_un.b_addr</b> and <b>b_bcount</b> are members of the <b>buf(9S)</b> data structure.
<b>CONTEXT</b>	<b>clrbuf()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>getrbuf(9F)</b> , <b>buf(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	cmn_err, vcmn_err – display an error message or panic the system
<b>SYNOPSIS</b>	<pre>#include &lt;sys/cmn_err.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void cmn_err(int level, char * format ...);  #include &lt;sys/varargs.h&gt;  void vcmn_err(int level, char * format, va_list ap);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	
<b>cmn_err()</b>	<p><b>level</b>            A constant indicating the severity of the error condition.</p> <p><b>format</b>            The message to be displayed.</p>
<b>vcmn_err()</b>	<p><b>vcmn_err()</b> takes <i>level</i> and <i>format</i> as described for <b>cmn_err()</b> , but its third argument is different:</p> <p><b>ap</b>                The variable argument list passed to the function.</p>
<b>DESCRIPTION</b>	
<b>cmn_err()</b>	<p><b>cmn_err()</b> displays a specified message on the console. <b>cmn_err()</b> can also panic the system. When the system panics, it attempts to save recent changes to data, display a “panic message” on the console, attempt to write a core file, and halt system processing. See the <i>CE_PANIC level</i> below.</p> <p><i>level</i> is a constant indicating the severity of the error condition. The four severity levels are:</p> <p><b>CE_CONT</b>            Used to continue another message or to display an informative message not associated with an error. Note that multiple <b>CE_CONT</b> messages without a newline may or may not appear on the system console or in the system buffer as a single line message. A single line message may be produced by constructing the message with <b>sprintf(9F)</b> or <b>vsprintf(9F)</b> before calling <b>cmn_err()</b> .</p> <p><b>CE_NOTE</b>            Used to display a message preceded with <b>NOTICE</b> . This message is used to report system events that do not</p>

necessarily require user action, but may interest the system administrator. For example, a message saying that a sector on a disk needs to be accessed repeatedly before it can be accessed correctly might be noteworthy.

**CE\_WARN** Used to display a message preceded with `WARNING`. This message is used to report system events that require immediate attention, such as those where if an action is not taken, the system may panic. For example, when a peripheral device does not initialize correctly, this level should be used.

**CE\_PANIC** Used to display a message preceded with `panic`, and to panic the system. Drivers should specify this level only under the most severe conditions or when debugging a driver. A valid use of this level is when the system cannot continue to function. If the error is recoverable, or not essential to continued system operation, do not panic the system.

*format* is the message to be displayed. It is a character string which may contain plain characters and conversion specifications. By default, the message is sent both to the system console and to the system buffer.

Each conversion specification in *format* is introduced by the `%` character, after which the following appear in sequence:

An optional decimal digit specifying a minimum field width for numeric conversion. The converted value will be right-justified and padded with leading zeroes if it has fewer characters than the minimum.

An optional `l` (`ll`) specifying that a following `d`, `D`, `o`, `O`, `x`, `X`, or `u` conversion character applies to a `long` (`long long`) integer argument. An `l` (`ll`) before any other conversion character is ignored.

A character indicating the type of conversion to be applied:

`d`, `D`, `o`, `O`, `x`, `X` The integer argument is converted to signed decimal (`d`, `D`), unsigned octal (`o`, `O`), unsigned hexadecimal (`x`, `X`), or unsigned decimal (`u`), respectively, and displayed. The letters `abcdef` are used for `x` and `X` conversion.

`c` The character value of the argument is displayed.

`b` The `%b` conversion specification allows bit values to be displayed meaningfully. Each `%b` takes an integer value and a format string from the argument list. The first character of the format string should be the output base encoded as a control character. This base is used to display the integer argument. The remaining groups of characters in the format

string consist of a bit number (between 1 and 32, also encoded as a control character) and the next characters (up to the next control character or '\\0') give the name of the bit field. The string corresponding to the bit fields set in the integer argument is displayed after the numerical value. See EXAMPLE section.

**s** The argument is taken to be a string (character pointer), and characters from the string are displayed until a null character is encountered. If the character pointer is `NULL`, the string `<null string>` is used in its place.

**%** Copy a % ; no argument is converted.

The first character in *format* affects where the message will be written:

**!** the message goes only to the system buffer.

**^** the message goes only to the console.

**?** If *level* is also `CE_CONT`, the message is always sent to the system buffer, but is only written to the console when the system has been booted in verbose mode. See `kernel(1M)`. If neither condition is met, the '?' character has no effect and is simply ignored.

To display the contents of the system buffer, use the `dmesg(1M)` command.

`cmn_err()` appends a \ to each *format*, except when *level* is `CE_CONT`.

**vcmn\_err()** `vcmn_err()` is identical to `cmn_err()` except that its last argument, *ap*, is a pointer to a variable list of arguments. *ap* contains the list of arguments used by the conversion specifications in *format*. *ap* must be initialized by calling `va_start(9F)`. `va_end(9F)` is used to clean up and must be called after each traversal of the list. Multiple traversals of the argument list, each bracketed by `va_start(9F)` and `va_end(9F)`, are possible.

## RETURN VALUES

None. However, if an unknown *level* is passed to `cmn_err()`, the following panic error message is displayed:

```
panic: unknown level in cmn_err (level= level , msg= format )
```

## CONTEXT

`cmn_err()` can be called from user or kernel context.

## EXAMPLES

**EXAMPLE 1** Using `cmn_err()`

This first example shows how `cmn_err()` can record tracing and debugging information only in the system buffer (lines 17); display problems with a device only on the system console (line 23); or display problems with the device on both the system console and in the system buffer (line 28).

```

1 struct reg {
2     uchar_t data;
3     uchar_t csr;
4 };
5
6 struct xxstate {
7     ...
8     dev_info_t *dip;
9     struct reg *regp;
10    ...
11 };
12
13 dev_t dev;
14 struct xxstate *xsp;
15 ...
16 #ifdef DEBUG /* in debugging mode, log function call */
17     cmn_err(CE_CONT, "!!%s%d: xxopen function called.",
18         ddi_binding_name(xsp->dip), getminor(dev));
19 #endif /* end DEBUG */
20 ...
21 /* display device power failure on system console */
22 if ((xsp->regp->csr & POWER) == OFF)
23     cmn_err(CE_NOTE, "^OFF.",
24         ddi_binding_name(xsp->dip), getminor(dev));
25 ...
26 /* display warning if device has bad VTOC */
27 if (xsp->regp->csr & BADVTOC)
28     cmn_err(CE_WARN, "%s%d: xxopen: Bad VTOC.",
29         ddi_binding_name(xsp->dip), getminor(dev));

```

**EXAMPLE 2** Using the %b conversion specification

This example shows how to use the %b conversion specification. Because of the leading '?' character in the format string, this message will always be logged, but it will only be displayed when the kernel is booted in verbose mode.

```

cmn_err(CE_CONT, "?reg=0x%b\
", regval, "\\020\\3Intr\\2Err\\1Enable");

```

**EXAMPLE 3** Using *regval*

When *regval* is set to (decimal) 13, the following message would be displayed:

```

reg=0xd<Intr,,Enable>

```

**EXAMPLE 4** Error Routine

The third example is an error reporting routine which accepts a variable number of arguments and displays a single line error message both in the system buffer and on the system console. Note the use of `vsprintf()` to construct the error message before calling `cmn_err()` .

```
#include <sys/varargs.h>
#include <sys/ddi.h>
#include <sys/sunddi.h>
#define MAX_MSG 256;

void
xxerror(dev_info_t *dip, int level, const char *fmt, ...)
{
    va_list      ap;
    int          instance;
    char         buf[MAX_MSG], *name;

    instance = ddi_get_instance(dip);
    name = ddi_binding_name(dip);

    /* format buf using fmt and arguments contained in ap */

    va_start(ap, fmt);
    vsprintf(buf, fmt, ap);
    va_end(ap);

    /* pass formatted string to cmn_err(9F) */

    cmn_err(level, "%s%d: %s", name, instance, buf);
}

```

**SEE ALSO** `dmesg(1M)` , `kernel(1M)` , `printf(3S)` , `ddi_binding_name(9F)` , `sprintf(9F)` , `va_arg(9F)` , `va_end(9F)` , `va_start(9F)` , `vsprintf(9F)`

*Writing Device Drivers*

**WARNINGS** `cmn_err()` with the `CE_CONT` argument can be used by driver developers as a driver code debugging tool. However, using `cmn_err()` in this capacity can change system timing characteristics.

**NOTES** At times, a driver may encounter error conditions requiring the attention of a primary or secondary system console monitor. These conditions may mean halting multiuser processing; however, this must be done with caution. Except during the debugging stage, a driver should never stop the system.

See the “Debugging” chapter in *Writing Device Drivers*

For severities of `CE_NOTE` and `CE_WARN`, the maximum message length is 256 bytes excluding "Note:" or "Warning:" respectively.

Any message greater than 128 bytes in length is divided into separate 128 byte messages.

**BUGS**

`cmn_err()` does not provide all of the functionality provided by `printf(3S)`

<b>NAME</b>	condvar, cv_init, cv_destroy, cv_wait, cv_signal, cv_broadcast, cv_wait_sig, cv_timedwait, cv_timedwait_sig – condition variable routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ksynch.h&gt;  void cv_init(kcondvar_t * cvp, char * name, kcv_type_t type, void * arg); void cv_destroy(kcondvar_t * cvp); void cv_wait(kcondvar_t * cvp, kmutex_t * mp); void cv_signal(kcondvar_t * cvp); void cv_broadcast(kcondvar_t * cvp); int cv_wait_sig(kcondvar_t * cvp, kmutex_t * mp); clock_t cv_timedwait(kcondvar_t * cvp, kmutex_t * mp, clock_t timeout); clock_t cv_timedwait_sig(kcondvar_t * cvp, kmutex_t * mp, clock_t timeout);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>cvp</b>            A pointer to an abstract data type kcondvar_t .</p> <p><b>mp</b>             A pointer to a mutual exclusion lock ( kmutex_t ), initialized by <b>mutex_init(9F)</b> and held by the caller.</p> <p><b>name</b>           Descriptive string. This is obsolete and should be NULL . (Non- NULL strings are legal, but they're a waste of kernel memory.)</p> <p><b>type</b>           The constant CV_DRIVER .</p> <p><b>arg</b>            A type-specific argument, drivers should pass arg as NULL .</p> <p><b>timeout</b>        A time, in absolute ticks since boot, when <b>cv_timedwait()</b> or <b>cv_timedwait_sig()</b> should return.</p>
<b>DESCRIPTION</b>	<p>Condition variables are a standard form of thread synchronization. They are designed to be used with mutual exclusion locks (mutexes). The associated mutex is used to ensure that a condition can be checked atomically and that the thread can block on the associated condition variable without missing either a change to the condition or a signal that the condition has changed. Condition variables must be initialized by calling <b>cv_init()</b> , and must be deallocated by calling <b>cv_destroy()</b> .</p>

The usual use of condition variables is to check a condition (for example, device state, data structure reference count, etc.) while holding a mutex which keeps other threads from changing the condition. If the condition is such that the thread should block, `cv_wait()` is called with a related condition variable and the mutex. At some later point in time, another thread would acquire the mutex, set the condition such that the previous thread can be unblocked, unblock the previous thread with `cv_signal()` or `cv_broadcast()`, and then release the mutex.

`cv_wait()` suspends the calling thread and exits the mutex atomically so that another thread which holds the mutex cannot signal on the condition variable until the blocking thread is blocked. Before returning, the mutex is reacquired.

`cv_signal()` signals the condition and wakes one blocked thread. All blocked threads can be unblocked by calling `cv_broadcast()`. You must acquire the mutex passed into `cv_wait()` before calling `cv_signal()` or `cv_broadcast()`.

The function `cv_wait_sig()` is similar to `cv_wait()` but returns 0 if a signal (for example, by `kill(2)`) is sent to the thread. In any case, the mutex is reacquired before returning.

The function `cv_timedwait()` is similar to `cv_wait()`, except that it returns -1 without the condition being signaled after the timeout time has been reached.

The function `cv_timedwait_sig()` is similar to `cv_timedwait()`, and `cv_wait_sig()`, except that it returns -1 without the condition being signaled after the timeout time has been reached, or 0 if a signal (for example, by `kill(2)`) is sent to the thread.

For both `cv_timedwait()` and `cv_timedwait_sig()`, time is in absolute clock ticks since the last system reboot. The current time may be found by calling `ddi_get_1bolt(9F)`.

## RETURN VALUES

0	For <code>cv_wait_sig()</code> and <code>cv_timedwait_sig()</code> indicates that the condition was not necessarily signaled and the function returned because a signal (as in <code>kill(2)</code> ) was pending.
-1	For <code>cv_timedwait()</code> and <code>cv_timedwait_sig()</code> indicates that the condition was not necessarily signaled and the function returned because the timeout time was reached.
>0	For <code>cv_wait_sig()</code> , <code>cv_timedwait()</code> or <code>cv_timedwait_sig()</code> indicates that the condition was met and the function returned due to a call to <code>cv_signal()</code> or <code>cv_broadcast()</code> .

## CONTEXT

These functions can be called from user, kernel or interrupt context. In most cases, however, `cv_wait()`, `cv_timedwait()`, `cv_wait_sig()`, and

**cv\_timedwait\_sig()** should not be called from interrupt context, and cannot be called from a high-level interrupt context.

If **cv\_wait()** , **cv\_timedwait()** , **cv\_wait\_sig()** , or **cv\_timedwait\_sig()** are used from interrupt context, lower-priority interrupts will not be serviced during the wait. This means that if the thread that will eventually perform the wakeup becomes blocked on anything that requires the lower-priority interrupt, the system will hang.

For example, the thread that will perform the wakeup may need to first allocate memory. This memory allocation may require waiting for paging I/O to complete, which may require a lower-priority disk or network interrupt to be serviced. In general, situations like this are hard to predict, so it is advisable to avoid waiting on condition variables or semaphores in an interrupt context.

## EXAMPLES

**EXAMPLE 1** : Waiting for a flag value in a driver's unit

Here the condition being waited for is a flag value in a driver's unit structure. The condition variable is also in the unit structure, and the flag word is protected by a mutex in the unit structure.

```
\011mutex_enter(&un->un_lock);
\011while (un->un_flag & UNIT_BUSY)
        cv_wait(&un->un_cv, &un->un_lock);
\011un->un_flag |= UNIT_BUSY;
\011mutex_exit(&un->un_lock);
```

**EXAMPLE 2** : Unblocking threads blocked by the code in Example 1

At some later point in time, another thread would execute the following to unblock any threads blocked by the above code.

```
\011
mutex_enter(&un->un_lock);
un->un_flag &= ~UNIT_BUSY;
cv_broadcast(&un->un_cv);
mutex_exit(&un->un_lock);
```

## SEE ALSO

**kill(2)** , **ddi\_get\_lbolt(9F)** , **mutex(9F)** , **mutex\_init(9F)**

*Writing Device Drivers*

<b>NAME</b>	copyb – copy a message block
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  mblk_t *copyb(mblk_t *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the message block from which data is copied.
<b>DESCRIPTION</b>	<b>copyb()</b> allocates a new message block, and copies into it the data from the block that <i>bp</i> denotes. The new block will be at least as large as the block being copied. <b>copyb()</b> uses the <i>b_rptr</i> and <i>b_wptr</i> members of <i>bp</i> to determine how many bytes to copy.
<b>RETURN VALUES</b>	If successful, <b>copyb()</b> returns a pointer to the newly allocated message block containing the copied data. Otherwise, it returns a <code>NULL</code> pointer.
<b>CONTEXT</b>	<b>copyb()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> : Using copyb</p> <p>For each message in the list, test to see if the downstream queue is full with the <b>canputnext(9F)</b> function (line 21). If it is not full, use <b>copyb</b> to copy a header message block, and <b>dupmsg(9F)</b> to duplicate the data to be retransmitted. If either operation fails, reschedule a timeout at the next valid interval.</p> <p>Update the new header block with the correct destination address (line 34), link the message to it (line 35), and send it downstream (line 36). At the end of the list, reschedule this routine.</p> <pre> 1  struct retrans { 2      mblk_t          *r_mp; 3      int             r_address; 4      queue_t        *r_outq; 5      struct retrans *r_next; 6  }; 7 8  struct protoheader { 9      ... 9      int             h_address; 10     ... 10  }; 11 12  mblk_t *header; 13 14  void</pre>

```

15 retransmit(struct retrans *ret)
16 {
17     mblk_t *bp, *mp;
18     struct protoheader *php;
19
20     while (ret) {
21         if (!canputnext(ret->r_outq)) {           /* no room */
22             ret = ret->r_next;
23             continue;
24         }
25         bp = copyb(header);                       /* copy header msg. block */
26         if (bp == NULL)
27             break;
28         mp = dupmsg(ret->r_mp);                   /* duplicate data */
29         if (mp == NULL) {                       /* if unsuccessful */
30             freeb(bp);                          /* free the block */
31             break;
32         }
33         php = (struct protoheader *)bp->b_rptr;
34         php->h_address = ret->r_address;         /* new header */
35         bp->bp_cont = mp;                       /* link the message */
36         putnext(ret->r_outq, bp);              /* send downstream */
37         ret = ret->r_next;
38     }
39     /* reschedule */
40     (void) timeout(retransmit, (caddr_t)ret, RETRANS_TIME);
41 }

```

**SEE ALSO** [allocb\(9F\)](#), [canputnext\(9F\)](#), [dupmsg\(9F\)](#)

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	copyin – copy data from a user program to a driver buffer
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt;  int copyin(const void *userbuf, void *driverbuf, size_t cn);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b><i>userbuf</i></b>            User program source address from which data is transferred.</p> <p><b><i>driverbuf</i></b>           Driver destination address to which data is transferred.</p> <p><b><i>cn</i></b>                    Number of bytes transferred.</p>
<b>DESCRIPTION</b>	<p><b>copyin()</b> copies data from a user program source address to a driver buffer. The driver developer must ensure that adequate space is allocated for the destination address.</p> <p>Addresses that are word-aligned are moved most efficiently. However, the driver developer is not obligated to ensure alignment. This function automatically finds the most efficient move according to address alignment.</p>
<b>RETURN VALUES</b>	<p>Under normal conditions a 0 is returned indicating a successful copy. Otherwise, a -1 is returned if one of the following occurs:</p> <ul style="list-style-type: none"> <li>■ paging fault; the driver tried to access a page of memory for which it did not have read or write access</li> <li>■ invalid user address, such as a user area or stack area</li> <li>■ invalid address that would have resulted in data being copied into the user block</li> </ul> <p>If a -1 is returned to the caller, driver entry point routines should return EFAULT.</p>
<b>CONTEXT</b>	<b>copyin()</b> can be called from user context only.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b>    An <b>ioctl()</b> Routine</p> <p>A driver <b>ioctl(9E)</b> routine (line 10) can be used to get or set device attributes or registers. In the <b>XX_GETREGS</b> condition (line 17), the driver copies the current device register values to a user data area (line 18). If the specified argument contains an invalid address, an error code is returned.</p>

```

1 struct device {          /* layout of physical device registers */
2     int     control;    /* physical device control word */
3     int     status;    /* physical device status word */
4     short   recv_char; /* receive character from device */
5     short   xmit_char; /* transmit character to device */
6 };
7
8 extern struct device xx_addr[]; /* phys. device regs. location */
9
10 xx_ioctl(dev_t dev, int cmd, int arg, int mode,
11         cred_t *cred_p, int *rval_p)
12     ...
13 {
14     register struct device *rp = &xx_addr[getminor(dev) >> 4];
15     switch (cmd) {
16
17     case XX_GETREGS: /* copy device regs. to user program */
18         if (copyin(arg, rp, sizeof(struct device)))
19             return(EFAULT);
20         break;
21         ...
22     }
23     ...
24 }

```

**SEE ALSO** [ioctl\(9E\)](#), [bcopy\(9F\)](#), [copyout\(9F\)](#), [ddi\\_copyin\(9F\)](#), [ddi\\_copyout\(9F\)](#), [uiomove\(9F\)](#).

*Writing Device Drivers*

**NOTES** Driver writers who intend to support layered ioctls in their [ioctl\(9E\)](#) routines should use [ddi\\_copyin\(9F\)](#) instead.

Driver defined locks should not be held across calls to this function.

This should not be used from a streams driver. See [M\\_COPYIN](#) and [M\\_COPYOUT](#) in *STREAMS Programming Guide*.

<b>NAME</b>	copymsg – copy a message
<b>SYNOPSIS</b>	#include <sys/stream.h>  mblk_t *copymsg(mblk_t *mp);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>mp</b> Pointer to the message to be copied.
<b>DESCRIPTION</b>	<b>copymsg()</b> forms a new message by allocating new message blocks, and copying the contents of the message referred to by <i>mp</i> (using the <b>copyb(9F)</b> function). It returns a pointer to the new message.
<b>RETURN VALUES</b>	If the copy is successful, <b>copymsg()</b> returns a pointer to the new message. Otherwise, it returns a NULL pointer.
<b>CONTEXT</b>	<b>copymsg()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<b>EXAMPLE 1</b> : Using copymsg  The routine <b>lctouc()</b> converts all the lowercase ASCII characters in the message to uppercase. If the reference count is greater than one (line 8), then the message is shared, and must be copied before changing the contents of the data buffer. If the call to the <b>copymsg()</b> function fails (line 9), return NULL (line 10), otherwise, free the original message (line 11). If the reference count was equal to 1, the message can be modified. For each character (line 16) in each message block (line 15), if it is a lowercase letter, convert it to an uppercase letter (line 18). A pointer to the converted message is returned (line 21).

```

1 mblk_t *lctouc(mp)
2   mblk_t *mp;
3   {
4     mblk_t *cmp;
5     mblk_t *tmp;
6     unsigned char *cp;
7
8     if (mp->b_datap->db_ref > 1) {
9       if ((cmp = copymsg(mp)) == NULL)
10        return (NULL);
11      freemsg(mp);
12    } else {
13      cmp = mp;
14    }
15    for (tmp = cmp; tmp; tmp = tmp->b_cont) {
16      for (cp = tmp->b_rptr; cp < tmp->b_wptr; cp++) {
17        if ((*cp <= 'z') && (*cp >= 'a'))

```

```
18                                     *cp -= 0x20;
19                                     }
20                                     }
21     return(cmp);
22 }
```

**SEE ALSO** `allocb(9F)`, `copyb(9F)`, `msgb(9S)`

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	copyout – copy data from a driver to a user program
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt;  int copyout(const void *driverbuf, void *userbuf, size_t cn);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>driverbuf</b>      Source address in the driver from which the data is transferred.</p> <p><b>userbuf</b>        Destination address in the user program to which the data is transferred.</p> <p><b>cn</b>              Number of bytes moved.</p>
<b>DESCRIPTION</b>	<p><b>copyout()</b> copies data from driver buffers to user data space.</p> <p>Addresses that are word-aligned are moved most efficiently. However, the driver developer is not obligated to ensure alignment. This function automatically finds the most efficient move algorithm according to address alignment.</p>
<b>RETURN VALUES</b>	<p>Under normal conditions a 0 is returned to indicate a successful copy. Otherwise, a -1 is returned if one of the following occurs:</p> <ul style="list-style-type: none"> <li>■ paging fault; the driver tried to access a page of memory for which it did not have read or write access</li> <li>■ invalid user address, such as a user area or stack area</li> <li>■ invalid address that would have resulted in data being copied into the user block</li> </ul> <p>If a -1 is returned to the caller, driver entry point routines should return EFAULT.</p>
<b>CONTEXT</b>	<b>copyout()</b> can be called from user context only.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> An <b>ioctl()</b> Routine</p> <p>A driver <b>ioctl(9E)</b> routine (line 10) can be used to get or set device attributes or registers. In the <b>XX_GETREGS</b> condition (line 17), the driver copies the current device register values to a user data area (line 18). If the specified argument contains an invalid address, an error code is returned.</p>

```

1 struct device {          /* layout of physical device registers */
2     int     control;     /* physical device control word */
3     int     status;      /* physical device status word */
4     short   recv_char;   /* receive character from device */
5     short   xmit_char;   /* transmit character to device */
6 };
7
8 extern struct device xx_addr[]; /* phys. device regs. location */
9     . . .
10 xx_ioctl(dev_t dev, int cmd, int arg, int mode,
11     cred_t *cred_p, int *rval_p)
12     . . .
13 {
14     register struct device *rp = &xx_addr[getminor(dev) >> 4];
15     switch (cmd) {
16
17     case XX_GETREGS:      /* copy device regs. to user program */
18         if (copyout(rp, arg, sizeof(struct device)))
19             return(EFAULT);
20         break;
21         . . .
22     }
23     . . .
24 }

```

**SEE ALSO** [ioctl\(9E\)](#), [bcopy\(9F\)](#), [copyin\(9F\)](#), [ddi\\_copyin\(9F\)](#), [ddi\\_copyout\(9F\)](#), [uiomove\(9F\)](#)

*Writing Device Drivers*

**NOTES** Driver writers who intend to support layered ioctls in their [ioctl\(9E\)](#) routines should use [ddi\\_copyout\(9F\)](#) instead.

Driver defined locks should not be held across calls to this function.

This should not be used from a streams driver. See [M\\_COPYIN](#) and [M\\_COPYOUT](#) in *STREAMS Programming Guide*.

<b>NAME</b>	csx_AccessConfigurationRegister – read or write a PC Card Configuration Register
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_AccessConfigurationRegister(client_handle_t ch, access_config_reg_t *acr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>acr</b> Pointer to an <b>access_config_reg_t</b> structure.</p>
<b>DESCRIPTION</b>	This function allows a client to read or write a PC Card Configuration Register.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>access_config_reg_t</b> are:</p> <pre>uint32_t Socket; /* socket number*/ uint32_t Action; /* register access operation*/ uint32_t Offset; /* config register offset*/ uint32_t Value; /* value read or written*/</pre> <p>The fields are defined as follows:</p> <p><b>Socket</b> Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p> <p><b>Action</b> May be set to <b>CONFIG_REG_READ</b> or <b>CONFIG_REG_WRITE</b>. All other values in the <b>Action</b> field are reserved for future use. If the <b>Action</b> field is set to <b>CONFIG_REG_WRITE</b>, the <b>Value</b> field is written to the specified configuration register. Card Services does not read the configuration register after a write operation. For that reason, the <b>Value</b> field is only updated by a <b>CONFIG_REG_READ</b> request.</p> <p><b>Offset</b> Specifies the byte offset for the desired configuration register from the PC Card configuration register base specified in <b>csx_RequestConfiguration(9F)</b>.</p> <p><b>Value</b> Contains the value read from the PC Card Configuration Register for a read operation. For a write operation, the</p>

`Value` field contains the value to write to the configuration register. As noted above, on return from a write request, the `Value` field is the value written to the PC Card and not any changed value that may have resulted from the write request (that is, no read after write is performed).

A client must be very careful when writing to the COR (Configuration Option Register) at offset 0. This has the potential to change the type of interrupt request generated by the PC Card or place the card in the reset state. Either request may have undefined results. The client should read the register to determine the appropriate setting for the interrupt mode (Bit 6) before writing to the register.

If a client wants to reset a PC Card, the `csx_ResetFunction(9F)` function should be used. Unlike `csx_AccessConfigurationRegister()`, the `csx_ResetFunction(9F)` function generates a series of event notifications to all clients using the PC Card, so they can re-establish the appropriate card state after the reset operation is complete.

**RETURN VALUES**

<code>CS_SUCCESS</code>	Successful operation.
<code>CS_BAD_ARGS</code>	Specified arguments are invalid. Client specifies an <code>Offset</code> that is out of range or neither <code>CONFIG_REG_READ</code> or <code>CONFIG_REG_WRITE</code> is set.
<code>CS_UNSUPPORTED_MODE</code>	Client has not called <code>csx_RequestConfiguration(9F)</code> before calling this function.
<code>CS_BAD_HANDLE</code>	Client handle is invalid.
<code>CS_NO_CARD</code>	No PC card in socket.
<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

`csx_ParseTuple(9F)`, `csx_RegisterClient(9F)`,  
`csx_RequestConfiguration(9F)`, `csx_ResetFunction(9F)`

*PCCard 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_ConvertSize – convert device sizes
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_ConvertSize(convert_size_t *cs);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b>cs</b> Pointer to a <code>convert_size_t</code> structure.
<b>DESCRIPTION</b>	<b>csx_ConvertSize()</b> is a Solaris-specific extension that provides a method for clients to convert from one type of device size representation to another, that is, from <i>devsize</i> format to <i>bytes</i> and vice versa.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>convert_size_t</code> are:</p> <pre>uint32_t    Attributes; uint32_t    bytes; uint32_t    devsize;</pre> <p>The fields are defined as follows:</p> <p><b>Attributes</b>      This is a bit-mapped field that identifies the type of size conversion to be performed. The field is defined as follows:</p> <pre>                  CONVERT_BYTES_TO_DEVSIZE                   Converts <i>bytes</i> to <i>devsize</i> format.</pre> <pre>                  CONVERT_DEVSIZE_TO_BYTES                   Converts <i>devsize</i> format to <i>bytes</i>.</pre> <p><b>bytes</b>            If <code>CONVERT_BYTES_TO_DEVSIZE</code> is set, the value in the <code>bytes</code> field is converted to a <i>devsize</i> format and returned in the <code>devsize</code> field.</p> <p><b>devsize</b>          If <code>CONVERT_DEVSIZE_TO_BYTES</code> is set, the value in the <code>devsize</code> field is converted to a <i>bytes</i> value and returned in the <code>bytes</code> field.</p>
<b>RETURN VALUES</b>	<pre>CS_SUCCESS                      Successful operation. CS_BAD_SIZE                      Invalid <i>bytes</i> or <i>devsize</i>.</pre>

CS\_UNSUPPORTED\_FUNCTION            No PCMCIA hardware installed.

**CONTEXT**            This function may be called from user or kernel context.

**SEE ALSO**            *csx\_ModifyWindow(9F)*, *csx\_RequestWindow(9F)*  
*PCCard 95 Standard*, *PCMCIA/JEIDA*

<b>NAME</b>	csx_ConvertSpeed – convert device speeds
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_ConvertSpeed(convert_speed_t *cs);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b>cs</b> Pointer to a <code>convert_speed_t</code> structure.
<b>DESCRIPTION</b>	This function is a Solaris-specific extension that provides a method for clients to convert from one type of device speed representation to another, that is, from <i>devspeed</i> format to <i>nS</i> and vice versa.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>convert_speed_t</code> are:</p> <pre>uint32_t    Attributes; uint32_t    nS; uint32_t    devspeed;</pre> <p>The fields are defined as follows:</p> <p><b>Attributes</b>      This is a bit-mapped field that identifies the type of speed conversion to be performed. The field is defined as follows:</p> <pre>                  CONVERT_NS_TO_DEVSPEED                   Converts <i>nS</i> to <i>devspeed</i> format</pre> <pre>                  CONVERT_DEVSPEED_TO_NS                   Converts <i>devspeed</i> format to <i>nS</i></pre> <p><b>nS</b>                If <code>CONVERT_NS_TO_DEVSPEED</code> is set, the value in the <code>nS</code> field is converted to a <i>devspeed</i> format and returned in the <code>devspeed</code> field.</p> <p><b>devspeed</b>        If <code>CONVERT_DEVSPEED_TO_NS</code> is set, the value in the <code>devspeed</code> field is converted to an <i>nS</i> value and returned in the <code>nS</code> field.</p>
<b>RETURN VALUES</b>	<pre>CS_SUCCESS                      Successful operation. CS_BAD_SPEED                    Invalid <i>nS</i> or <i>devspeed</i>.</pre>

CS_BAD_ATTRIBUTE	Bad Attributes value.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_ModifyWindow(9F)`, `csx_RequestWindow(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_CS_DDI_Info – obtain DDI information						
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_CS_DDI_Info(cs_ddi_info_t *cdi);</pre>						
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)						
<b>PARAMETERS</b>	<b>cdi</b> Pointer to a cs_ddi_info_t structure.						
<b>DESCRIPTION</b>	This function is a Solaris-specific extension that is used by clients that need to provide the <i>xx_getinfo</i> driver entry point (see <i>getinfo</i> (9E)). It provides a method for clients to obtain DDI information based on their socket number and client driver name.						
<b>STRUCTURE MEMBERS</b>	<p>The structure members of cs_ddi_info_t are:</p> <pre>uint32_t  Socket;          /* socket number */ char*     driver_name;    /* unique driver name */ dev_info_t *dip;         /* dip */ int32_t   instance;      /* instance */</pre> <p>The fields are defined as follows:</p> <p><b>Socket</b> This field must be set to the physical socket number that the client is interested in getting information about.</p> <p><b>driver_name</b> This field must be set to a string containing the name of the client driver to get information about.</p> <p>If <b>csx_CS_DDI_Info()</b> is used in a client's <i>xx_getinfo</i> function, then the client will typically extract the <b>Socket</b> value from the <i>*arg</i> argument and it <i>must</i> set the <b>driver_name</b> field to the same string used with <b>csx_RegisterClient</b>(9F).</p> <p>If the <b>driver_name</b> is found on the <b>Socket</b>, the <b>csx_CS_DDI_Info()</b> function returns both the <b>dev_info</b> pointer and the <b>instance</b> fields for the requested driver instance.</p>						
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>CS_SUCCESS</td> <td>Successful operation.</td> </tr> <tr> <td>CS_BAD_SOCKET</td> <td>Client not found on Socket.</td> </tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td> <td>No PCMCIA hardware installed.</td> </tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_SOCKET	Client not found on Socket.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.						
CS_BAD_SOCKET	Client not found on Socket.						
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.						

**CONTEXT**

This function may be called from user or kernel context.

**EXAMPLES**

**EXAMPLE 1** : Using `csx_CS_DDI_Info`

The following example shows how a client might call the `csx_CS_DDI_Info()` in the client's `xx_getinfo` function to return the dip or the instance number:

```
static int
pcepp_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg,
              void **result)
{
    int                error = DDI_SUCCESS;
    pcepp_state_t     *pps;
    cs_ddi_info_t     cs_ddi_info;

    switch (cmd) {

    case DDI_INFO_DEVT2DEVINFO:
        cs_ddi_info.Socket = getminor((dev_t)arg) & 0x3f;
        cs_ddi_info.driver_name = pcepp_name;
        if (csx_CS_DDI_Info(&cs_ddi_info) != CS_SUCCESS)
            return (DDI_FAILURE);
        if (!(pps = ddi_get_soft_state(pcepp_soft_state_p,
                                       cs_ddi_info.instance))) {
            *result = NULL;
        } else {
            *result = pps->dip;
        }
        break;

    case DDI_INFO_DEVT2INSTANCE:
        cs_ddi_info.Socket = getminor((dev_t)arg) & 0x3f;
        cs_ddi_info.driver_name = pcepp_name;
        if (csx_CS_DDI_Info(&cs_ddi_info) != CS_SUCCESS)
            return (DDI_FAILURE);
        *result = (void *)cs_ddi_info.instance;
        break;

    default:
        error = DDI_FAILURE;
        break;

    }

    return (error);
}
```

**SEE ALSO**

`getinfo(9E)`, `csx_RegisterClient(9F)`, `ddi_get_instance(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_DeregisterClient – remove client from Card Services list								
<b>SYNOPSIS</b>	#include <sys/pccard.h>  int32_t csx_DeregisterClient(client_handle_t ch);								
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)								
<b>PARAMETERS</b>	<b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .								
<b>DESCRIPTION</b>	This function removes a client from the list of registered clients maintained by Card Services. The Client Handle returned by <b>csx_RegisterClient(9F)</b> is passed in the <code>client_handle_t</code> argument.  The client must have returned all requested resources before this function is called. If any resources have not been released, <code>CS_IN_USE</code> is returned.								
<b>RETURN VALUES</b>	<table border="0"> <tr> <td><code>CS_SUCCESS</code></td> <td>Successful operation.</td> </tr> <tr> <td><code>CS_BAD_HANDLE</code></td> <td>Client handle is invalid.</td> </tr> <tr> <td><code>CS_IN_USE</code></td> <td>Resources not released by this client.</td> </tr> <tr> <td><code>CS_UNSUPPORTED_FUNCTION</code></td> <td>No PCMCIA hardware installed.</td> </tr> </table>	<code>CS_SUCCESS</code>	Successful operation.	<code>CS_BAD_HANDLE</code>	Client handle is invalid.	<code>CS_IN_USE</code>	Resources not released by this client.	<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.
<code>CS_SUCCESS</code>	Successful operation.								
<code>CS_BAD_HANDLE</code>	Client handle is invalid.								
<code>CS_IN_USE</code>	Resources not released by this client.								
<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.								
<b>CONTEXT</b>	This function may be called from user or kernel context.								
<b>SEE ALSO</b>	<b>csx_RegisterClient(9F)</b>  <i>PC Card 95 Standard, PCMCIA/JEIDA</i>								
<b>WARNINGS</b>	Clients should be prepared to receive callbacks until Card Services returns from this request successfully.								

<b>NAME</b>	csx_DupHandle – duplicate access handle																
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_DupHandle(acc_handle_t handle1, acc_handle_t *handle2, uint32_t flags);</pre>																
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)																
<b>PARAMETERS</b>	<p><b>handle1</b>            The access handle returned from <b>csx_RequestIO(9F)</b> or <b>csx_RequestWindow(9F)</b> that is to be duplicated.</p> <p><b>handle2</b>            A pointer to the newly-created duplicated data access handle.</p> <p><b>flags</b>                The access attributes that will be applied to the new handle.</p>																
<b>DESCRIPTION</b>	<p>This function duplicates the handle, <i>handle1</i>, into a new handle, <i>handle2</i>, that has the access attributes specified in the <i>flags</i> argument. Both the original handle and the new handle are active and can be used with the common access functions.</p> <p>Both handles must be explicitly freed when they are no longer necessary.</p> <p>The <i>flags</i> argument is bit-mapped. The following bits are defined:</p> <table border="0"> <tr> <td>WIN_ACC_NEVER_SWAP</td> <td>Host endian byte ordering</td> </tr> <tr> <td>WIN_ACC_BIG_ENDIAN</td> <td>Big endian byte ordering</td> </tr> <tr> <td>WIN_ACC_LITTLE_ENDIAN</td> <td>Little endian byte ordering</td> </tr> <tr> <td>WIN_ACC_STRICT_ORDER</td> <td>Program ordering references</td> </tr> <tr> <td>WIN_ACC_UNORDERED_OK</td> <td>May re-order references</td> </tr> <tr> <td>WIN_ACC_MERGING_OK</td> <td>Merge stores to consecutive locations</td> </tr> <tr> <td>WIN_ACC_LOADCACHING_OK</td> <td>May cache load operations</td> </tr> <tr> <td>WIN_ACC_STORECACHING_OK</td> <td>May cache store operations</td> </tr> </table> <p>WIN_ACC_BIG_ENDIAN and WIN_ACC_LITTLE_ENDIAN describe the endian characteristics of the device as big endian or little endian, respectively. Even though most of the devices will have the same endian characteristics as their busses, there are examples of devices with an I/O processor that has opposite endian characteristics of the busses. When WIN_ACC_BIG_ENDIAN or WIN_ACC_LITTLE_ENDIAN is set, byte swapping will automatically be performed by the system if the host machine and the device data formats have opposite endian characteristics. The implementation may take advantage of hardware platform byte swapping capabilities. When WIN_ACC_NEVER_SWAP is specified, byte swapping will not be invoked in the data access functions. The ability to specify the order in which the CPU will reference data is</p>	WIN_ACC_NEVER_SWAP	Host endian byte ordering	WIN_ACC_BIG_ENDIAN	Big endian byte ordering	WIN_ACC_LITTLE_ENDIAN	Little endian byte ordering	WIN_ACC_STRICT_ORDER	Program ordering references	WIN_ACC_UNORDERED_OK	May re-order references	WIN_ACC_MERGING_OK	Merge stores to consecutive locations	WIN_ACC_LOADCACHING_OK	May cache load operations	WIN_ACC_STORECACHING_OK	May cache store operations
WIN_ACC_NEVER_SWAP	Host endian byte ordering																
WIN_ACC_BIG_ENDIAN	Big endian byte ordering																
WIN_ACC_LITTLE_ENDIAN	Little endian byte ordering																
WIN_ACC_STRICT_ORDER	Program ordering references																
WIN_ACC_UNORDERED_OK	May re-order references																
WIN_ACC_MERGING_OK	Merge stores to consecutive locations																
WIN_ACC_LOADCACHING_OK	May cache load operations																
WIN_ACC_STORECACHING_OK	May cache store operations																

provided by the following *flags* bits. Only one of the following bits may be specified:

WIN_ACC_STRICT_ORDER	The data references must be issued by a CPU in program order. Strict ordering is the default behavior.
WIN_ACC_UNORDERED_OK	The CPU may re-order the data references. This includes all kinds of re-ordering (that is, a load followed by a store may be replaced by a store followed by a load).
WIN_ACC_MERGING_OK	The CPU may merge individual stores to consecutive locations. For example, the CPU may turn two consecutive byte stores into one halfword store. It may also batch individual loads. For example, the CPU may turn two consecutive byte loads into one halfword load. Setting this bit also implies re-ordering.
WIN_ACC_LOADCACHING_OK	The CPU may cache the data it fetches and reuse it until another store occurs. The default behavior is to fetch new data on every load. Setting this bit also implies merging and re-ordering.
WIN_ACC_STORECACHING_OK	The CPU may keep the data in the cache and push it to the device (perhaps with other data) at a later time. The default behavior is to push the data right away. Setting this bit also implies load caching, merging, and re-ordering.

These values are advisory, not mandatory. For example, data can be ordered without being merged or cached, even though a driver requests unordered, merged and cached together.

## RETURN VALUES

CS_SUCCESS	Successful operation.
CS_FAILURE	Error in <i>flags</i> argument or handle could not be duplicated for some reason.

CS\_UNSUPPORTED\_FUNCTION

No PCMCIA hardware  
installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

`csx_Get8(9F)`, `csx_GetMappedAddr(9F)`, `csx_Put8(9F)`,  
`csx_RepGet8(9F)`, `csx_RepPut8(9F)`, `csx_RequestIO(9F)`,  
`csx_RequestWindow(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Error2Text – convert error return codes to text strings
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Error2Text(error2text_t *er);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b>er</b> Pointer to an error2text_t structure.
<b>DESCRIPTION</b>	This function is a Solaris-specific extension that provides a method for clients to convert Card Services error return codes to text strings.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of error2text_t are:</p> <pre>uint32_t    item;                                    /*the error code*/ char        test[CS_ERROR_MAX_BUFSIZE];         /*the error code*/</pre> <p>A pointer to the text for the Card Services error return code in the <code>item</code> field is returned in the <code>text</code> field if the error return code is found. The client is not responsible for allocating a buffer to hold the text. If the Card Services error return code specified in the <code>item</code> field is not found, the <code>text</code> field will be set to a string of the form:</p> <pre>"{unknown Card Services return code}"</pre>
<b>RETURN VALUES</b>	<pre>CS_SUCCESS                                         Successful operation.  CS_UNSUPPORTED_FUNCTION                         No PCMCIA hardware installed.</pre>
<b>CONTEXT</b>	This function may be called from user or kernel context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b>    : Using the csxError2Text function</p> <pre>if ((ret = csx_RegisterClient(&amp;client_handle, &amp;client_reg)) != CS_SUCCESS) {     error2text_t error2text;     error2text.item = ret;     csx_Error2Text(&amp;error2text);     cmn_err(CE_CONT, "RegisterClient failed %s (0x%x)", error2text.text, ret); }</pre>
<b>SEE ALSO</b>	<b>csx_Event2Text(9F)</b>

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Event2Text – convert events to text strings
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Event2Text(event2text_t *ev);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b>ev</b> Pointer to an event2text_t structure.
<b>DESCRIPTION</b>	This function is a Solaris-specific extension that provides a method for clients to convert Card Services events to text strings.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of event2text_t are:</p> <pre>event_t      event;                                /*the event code*/ char         text[CS_EVENT_MAX_BUFSIZE]        /*the event code*/</pre> <p>The fields are defined as follows:</p> <pre>event                The text for the event code in the event field is returned in                       the text field.  text                 The text string describing the name of the event.</pre>
<b>RETURN VALUES</b>	<pre>CS_SUCCESS                                        Successful operation.  CS_UNSUPPORTED_FUNCTION                         No PCMCIA hardware installed.</pre>
<b>CONTEXT</b>	This function may be called from user or kernel context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> : Using <b>csx_Event2Text()</b></p> <pre>xx_event(event_t event, int priority, event_callback_args_t *eca) {     event2text_t    event2text;      event2text.event = event;     csx_Event2Text(&amp;event2text);     cmn_err(CE_CONT, "event %s (0x%x)", event2text.text, (int)event); }</pre>
<b>SEE ALSO</b>	<b>csx_event_handler(9E)</b> , <b>csx_Error2Text(9F)</b>

*PC Card 95 Standard, PCMCIA/JEIDA*



<b>NAME</b>	csx_Get8, csx_Get16, csx_Get32, csx_Get64 – read data from device address
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  uint8_t csx_Get8(acc_handle_t handle, uint32_t offset); uint16_t csx_Get16(acc_handle_t handle, uint32_t offset); uint32_t csx_Get32(acc_handle_t handle, uint32_t offset); uint64_t csx_Get64(acc_handle_t handle, uint64_t offset);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>handle</b> The access handle returned from <code>csx_RequestIO(9F)</code> , <code>csx_RequestWindow(9F)</code> , or <code>csx_DupHandle(9F)</code> .</p> <p><b>offset</b> The offset in bytes from the base of the mapped resource.</p>
<b>DESCRIPTION</b>	<p>These functions generate a read of various sizes from the mapped memory or device register.</p> <p>The <code>csx_Get8()</code> , <code>csx_Get16()</code> , <code>csx_Get32()</code> , and <code>csx_Get64()</code> functions read 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively, from the device address represented by the handle, <i>handle</i> , at an offset in bytes represented by the offset, <i>offset</i> .</p> <p>Data that consists of more than one byte will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte swapping if the host and the device have incompatible endian characteristics.</p>
<b>RETURN VALUES</b>	These functions return the value read from the mapped address.
<b>CONTEXT</b>	These functions may be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<p><code>csx_DupHandle(9F)</code> , <code>csx_GetMappedAddr(9F)</code> , <code>csx_Put8(9F)</code> , <code>csx_RepGet8(9F)</code> , <code>csx_RepPut8(9F)</code> , <code>csx_RequestIO(9F)</code> , <code>csx_RequestWindow(9F)</code></p> <p><i>PC Card 95 Standard</i>, PCMCIA/JEIDA</p>

<b>NAME</b>	csx_GetFirstClient, csx_GetNextClient – return first or next client
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_GetFirstClient(get_firstnext_client_t * fnc); int32_t csx_GetNextClient(get_firstnext_client_t * fnc);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b>fnc</b> Pointer to a <code>get_firstnext_client_t</code> structure.
<b>DESCRIPTION</b>	The functions <code>csx_GetFirstClient()</code> and <code>csx_GetNextClient()</code> return information about the first or subsequent PC cards, respectively, that are installed in the system.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>get_firstnext_client_t</code> are:</p> <pre>uint32_t      Socket;          /* socket number */ uint32_t      Attributes;     /* attributes */ client_handle_t client_handle; /* client handle */ uint32_t      num_clients;    /* number of clients */</pre> <p>The fields are defined as follows:</p> <p><b>Socket</b> If the <code>CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY</code> attribute is set, return information only on the PC card installed in this socket.</p> <p><b>Attributes</b> This field indicates the type of client. The field is bit-mapped; the following bits are defined:</p> <p><b>CS_GET_FIRSTNEXT_CLIENT_ALL_CLIENTS</b> Return information on all clients.</p> <p><b>CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY</b> Return client information for the specified socket only.</p> <p><b>client_handle</b> The client handle of the PC card driver is returned in this field.</p> <p><b>num_clients</b> The number of clients is returned in this field.</p>

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_BAD_SOCKET	Socket number is invalid.
CS_NO_CARD	No PC Card in socket.
CS_NO_MORE_ITEMS	PC Card driver does not handle the CS_EVENT_CLIENT_INFO event.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

**csx\_event\_handler(9E)**

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_GetFirstTuple, csx_GetNextTuple – return Card Information Structure tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_GetFirstTuple(client_handle_t ch, tuple_t * tu); int32_t csx_GetNextTuple(client_handle_t ch, tuple_t * tu);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code> .</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure.</p>
<b>DESCRIPTION</b>	The functions <code>csx_GetFirstTuple()</code> and <code>csx_GetNextTuple()</code> return the first and next tuple, respectively, of the specified type in the Card Information Structure (CIS) for the specified socket.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>tuple_t</code> are:</p> <pre>uint32_t Socket; /* socket number */ uint32_t Attributes; /* Attributes */ cisdata_t DesiredTuple; /* tuple to search for or flags */ cisdata_t TupleCode; /* tuple type code */ cisdata_t TupleLink; /* tuple data body size */</pre> <p>The fields are defined as follows:</p> <p>Socket</p> <p>Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p> <p>Attributes</p> <p>This field is bit-mapped. The following bits are defined:</p> <p><b>TUPLE_RETURN_LINK</b></p> <p>Return link tuples if set. The following are link tuples and will only be returned by this function if the <code>TUPLE_RETURN_LINK</code> bit in the <code>Attributes</code> field is set:</p> <pre>\011CISTPL_NULL\011CISTPL_LOGLINK_MFC \011CISTPL_LOGLINK_A\011CISTPL_LINKTARGET \011CISTPL_LOGLINK_C\011CISTPL_NO_LINK \011CISTPL_LOGLINK_CB\011CISTPL_END</pre>

**TUPLE\_RETURN\_IGNORED\_TUPLES**

Return ignored tuples if set. Ignored tuples will be returned by this function if the `TUPLE_RETURN_IGNORED_TUPLES` bit in the `Attributes` field is set, see `tuple(9S)` for more information. The `CIS` is parsed from the location setup by the previous `csx_GetFirstTuple()` or `csx_GetNextTuple()` request.

`DesiredTuple`

This field is the tuple value desired. If it is `RETURN_FIRST_TUPLE`, the very first tuple of the `CIS` is returned (if it exists). If this field is set to `RETURN_NEXT_TUPLE`, the very next tuple of the `CIS` is returned (if it exists). If the `DesiredTuple` field is any other value on entry, the `CIS` is searched in an attempt to locate a tuple which matches.

`TupleCode`, `TupleLink`

These fields are the values returned from the tuple found. If there are no tuples on the card, `CS_NO_MORE_ITEMS` is returned.

Since the `csx_GetFirstTuple()`, `csx_GetNextTuple()`, and `csx_GetTupleData(9F)` functions all share the same `tuple_t` structure, some fields in the `tuple_t` structure are unused or reserved when calling this function and these fields must not be initialized by the client.

**RETURN VALUES**

<code>CS_SUCCESS</code>	Successful operation.
<code>CS_BAD_HANDLE</code>	Client handle is invalid.
<code>CS_NO_CARD</code>	No PC Card in socket.
<code>CS_NO_CIS</code>	No Card Information Structure (CIS) on PC card.
<code>CS_NO_MORE_ITEMS</code>	Desired tuple not found.
<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.

**CONTEXT**

These functions may be called from user or kernel context.

**SEE ALSO**

`csx_GetTupleData(9F)`, `csx_ParseTuple(9F)`, `csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95Standard*, *PCMCIA/JEIDA*



<b>NAME</b>	csx_GetMappedAddr – return mapped virtual address	
<b>SYNOPSIS</b>	#include <sys/pccard.h>	
	int32_t csx_GetMappedAddr(acc_handle_t handle, void **addr);	
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)	
<b>PARAMETERS</b>	<p><b>handle</b> The access handle returned from <code>csx_RequestIO(9F)</code>, <code>csx_RequestWindow(9F)</code>, or <code>csx_DupHandle(9F)</code>.</p> <p><b>addr</b> The virtual or I/O port number represented by the handle.</p>	
<b>DESCRIPTION</b>	This function returns the mapped virtual address or the mapped I/O port number represented by the handle, <i>handle</i> .	
<b>RETURN VALUES</b>	CS_SUCCESS	The resulting address or I/O port number can be directly accessed by the caller.
	CS_FAILURE	The resulting address or I/O port number can not be directly accessed by the caller; the caller must make all accesses to the mapped area via the common access functions.
	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
<b>CONTEXT</b>	This function may be called from user, kernel, or interrupt context.	
<b>SEE ALSO</b>	<code>csx_DupHandle(9F)</code> , <code>csx_Get8(9F)</code> , <code>csx_Put8(9F)</code> , <code>csx_RepGet8(9F)</code> , <code>csx_RepPut8(9F)</code> , <code>csx_RequestIO(9F)</code> , <code>csx_RequestWindow(9F)</code> <i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

<b>NAME</b>	csx_GetStatus – return the current status of a PC Card and its socket
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_GetStatus(client_handle_t ch, get_status_t *gs);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>gs</b> Pointer to a <b>get_status_t</b> structure.</p>
<b>DESCRIPTION</b>	This function returns the current status of a PC Card and its socket.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>get_status_t</b> are:</p> <pre>uint32_t Socket;          /* socket number*/ uint32_t CardState;      /* "live" card status for this client*/ uint32_t SocketState;   /* latched socket values */ uint32_t raw_CardState; /* raw live card status */</pre> <p>The fields are defined as follows:</p> <p><b>Socket</b> Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p>

CardState      The `CardState` field is the bit-mapped output data returned from Card Services. The bits identify what Card Services thinks the current state of the installed PC Card is. The bits are:

- CS\_STATUS\_WRITE\_PROTECTED**  
Card is write protected
- CS\_STATUS\_CARD\_LOCKED**  
Card is locked
- CS\_STATUS\_EJECTION\_REQUEST**  
Ejection request in progress
- CS\_STATUS\_INSERTION\_REQUEST**  
Insertion request in progress
- CS\_STATUS\_BATTERY\_DEAD**  
Card battery is dead
- CS\_STATUS\_BATTERY\_DEAD**  
Card battery is dead (BVD1)
- CS\_STATUS\_BATTERY\_LOW**  
Card battery is low (BVD2)
- CS\_STATUS\_CARD\_READY**  
Card is READY
- CS\_STATUS\_CARD\_INSERTED**  
Card is inserted
- CS\_STATUS\_REQ\_ATTN**  
Extended status attention request
- CS\_STATUS\_RES\_EVT1**  
Extended status reserved event status
- CS\_STATUS\_RES\_EVT2**  
Extended status reserved event status
- CS\_STATUS\_RES\_EVT3**  
Extended status reserved event status
- CS\_STATUS\_VCC\_50**  
5.0 Volts Vcc Indicated

**CS\_STATUS\_VCC\_33**

3.3 Volts Vcc Indicated

**CS\_STATUS\_VCC\_XX**

X.X Volts Vcc Indicated

The state of the `CS_STATUS_CARD_INSERTED` bit indicates whether the PC Card associated with this driver instance, not just any card, is inserted in the socket. If an I/O card is installed in the specified socket, card state is returned from the PRR (Pin Replacement Register) and the ESR (Extended Status Register) (if present). If certain state bits are not present in the PRR or ESR, a simulated state bit value is returned as defined below:

**CS\_STATUS\_WRITE\_PROTECTED**

Not write protected

**CS\_STATUS\_BATTERY\_DEAD**

Power good

**PCS\_STATUS\_BATTERY\_LOW**

Power good

**CS\_STATUS\_CARD\_READY**

Ready

**CS\_STATUS\_REQ\_ATTN**

Not set

**CS\_STATUS\_RES\_EVT1**

Not set

**CS\_STATUS\_RES\_EVT2**

Not set

**CS\_STATUS\_RES\_EVT3**

Not set

`SocketState` The `SocketState` field is a bit-map of the current card and socket state. The bits are:

- CS SOCK STATUS WRITE PROTECT CHANGE**  
Write Protect
- ECS SOCK STATUS CARD LOCK CHANGE**  
Card Lock Change
- CS SOCK STATUS EJECTION PENDING**  
Ejection Request
- CS SOCK STATUS INSERTION PENDING**  
Insertion Request
- CS SOCK STATUS BATTERY DEAD CHANGE**  
Battery Dead
- CS SOCK STATUS BATTERY LOW CHANGE**  
Battery Low
- CS SOCK STATUS CARD READY CHANGE**  
Ready Change
- CS SOCK STATUS CARD INSERTION CHANGE**  
Card is inserted

The state reported in the `SocketState` field may be different from the state reported in the `CardState` field. Clients should normally depend only on the state reported in the `CardState` field.

The state reported in the `SocketState` field may be different from the state reported in the `CardState` field. Clients should normally depend only on the state reported in the `CardState` field.

`raw_CardState` The `raw_CardState` field is a Solaris-specific extension that allows the client to determine if any card is inserted in the socket. The bit definitions in the `raw_CardState` field are identical to those in the `CardState` field with the exception that the `CS STATUS CARD INSERTED` bit in the `raw_CardState` field is set whenever any card is inserted into the socket.

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_BAD_SOCKET	Error getting socket state.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_NO_CARD	will not be returned if there is no PC Card present in the socket.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

**csx\_RegisterClient(9F)**

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_GetTupleData – return the data portion of a tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_GetTupleData(client_handle_t ch, tuple_t *tu);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code>.</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure.</p>
<b>DESCRIPTION</b>	This function returns the data portion of a tuple, as returned by the <code>csx_GetFirstTuple(9F)</code> and <code>csx_GetNextTuple(9F)</code> functions.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>tuple_t</code> are:</p> <p>The fields are defined as follows:</p> <pre>uint32_t    Socket;                /* socket number */ uint32_t    Attributes;           /* tuple attributes*/ cisdata_t   DesiredTuple;        /* tuple to search for*/ cisdata_t   TupleOffset;        /* tuple data offset*/ cisdata_t   TupleDataMax;       /* max tuple data size*/ cisdata_t   TupleDataLen;       /* actual tuple data length*/ cisdata_t   TupleData[CIS_MAX_TUPLE_DATA_LEN]; /* tuple body data buffer*/ cisdata_t   TupleCode;          /* tuple type code*/ cisdata_t   TupleLink;          /* tuple link */</pre> <p><b>Socket</b> Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p> <p><b>Attributes</b> Initialized by <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>; the client must not modify the value in this field.</p> <p><b>DesiredTuple</b> Initialized by <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>; the client must not modify the value in this field.</p> <p><b>TupleOffset</b> This field allows partial tuple information to be retrieved, starting anywhere within the tuple.</p>

TupleDataMax	This field is the size of the tuple data buffer that Card Services uses to return raw tuple data from <code>csx_GetTupleData(9F)</code> . It can be larger than the number of bytes in the tuple data body. Card Services ignores any value placed here by the client.
TupleDataLen	This field is the actual size of the tuple data body. It represents the number of tuple data body bytes returned.
TupleData	This field is an array of bytes containing the raw tuple data body contents.
TupleCode	Initialized by <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code> ; the client must not modify the value in this field.
TupleLink	Initialized by <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code> ; the client must not modify the value in this field.

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_BAD_ARGS	Data from prior <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code> is corrupt.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_NO_MORE_ITEMS	Card Services was not able to read the tuple from the PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

`csx_GetFirstTuple(9F)`, `csx_ParseTuple(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_MakeDeviceNode, csx_RemoveDeviceNode – create and remove minor nodes on behalf of the client
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_MakeDeviceNode(client_handle_t ch, make_device_node_t * dn);  int32_t csx_RemoveDeviceNode(client_handle_t ch, remove_device_node_t * dn);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>dn</b> Pointer to a <code>make_device_node_t</code> or <code>remove_device_node_t</code> structure.</p>
<b>DESCRIPTION</b>	<b>csx_MakeDeviceNode()</b> and <b>csx_RemoveDeviceNode()</b> are Solaris-specific extensions to allow the client to request that device nodes in the filesystem are created or removed, respectively, on its behalf.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>make_device_node_t</code> are:</p> <pre>uint32_t      Action;          /* device operation */ uint32_t      NumDevNodes;    /* number of nodes to create */ devnode_desc_t *devnode_desc; /* description of device nodes */</pre> <p>The structure members of <code>remove_device_node_t</code> are:</p> <pre>uint32_t      Action;          /* device operation */ uint32_t      NumDevNodes;    /* number of nodes to remove */ devnode_desc_t *devnode_desc; /* description of device nodes */</pre> <p>The structure members of <code>devnode_desc_t</code> are:</p> <pre>char          *name;          /* device node path and name */ int32_t       spec_type;     /* device special type (block or char) */ int32_t       minor_num;    /* device node minor number */ char          *node_type;    /* device node type */</pre> <p>The <code>Action</code> field is used to specify the operation that <b>csx_MakeDeviceNode()</b> and <b>csx_RemoveDeviceNode()</b> should perform.</p>

The following `Action` values are defined for `csx_MakeDeviceNode()` :

**CREATE\_DEVICE\_NODE**

Create `NumDevNodes` minor nodes

The following `Action` values are defined for `csx_RemoveDeviceNode()` :

**REMOVE\_DEVICE\_NODE**

Remove `NumDevNodes` minor nodes

**REMOVE\_ALL\_DEVICE\_NODES**

Remove all minor nodes for this client

For `csx_MakeDeviceNode()` , if the `Action` field is:

`CREATE_DEVICE_NODE`

The `NumDevNodes` field must be set to the number of minor devices to create, and the client must allocate the quantity of `devnode_desc_t` structures specified by `NumDevNodes` and fill out the fields in the `devnode_desc_t` structure with the appropriate minor node information. The meanings of the fields in the `devnode_desc_t` structure are identical to the parameters of the same name to the `ddi_create_minor_node(9F)` DDI function.

For `csx_RemoveDeviceNode()` , if the `Action` field is:

`REMOVE_DEVICE_NODE`

The `NumDevNodes` field must be set to the number of minor devices to remove, and the client must allocate the quantity of `devnode_desc_t` structures specified by `NumDevNodes` and fill out the fields in the `devnode_desc_t` structure with the appropriate minor node information. The meanings of the fields in the `devnode_desc_t` structure are identical to the parameters of the same name to the `ddi_remove_minor_node(9F)` DDI function.

`REMOVE_ALL_DEVICE_NODES`

The `NumDevNodes` field must be set to 0 and the `devnode_desc_t` structure pointer must be set to `NULL` . All device nodes for this client will be removed from the filesystem.

**RETURN VALUES**

<code>CS_SUCCESS</code>	Successful operation.
<code>CS_BAD_HANDLE</code>	Client handle is invalid.

CS_BAD_ATTRIBUTE	The value of one or more arguments is invalid.
CS_BAD_ARGS	Action is invalid.
CS_OUT_OF_RESOURCE	Unable to create or remove device node.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** These functions may be called from user or kernel context.

**SEE ALSO** `csx_RegisterClient(9F)` , `ddi_create_minor_node(9F)` ,  
`ddi_remove_minor_node(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_MapLogSocket – return the physical socket number associated with the client handle
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_MapLogSocket(client_handle_t ch, map_log_socket_t *ls);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>ls</b> Pointer to a <b>map_log_socket_t</b> structure.</p>
<b>DESCRIPTION</b>	This function returns the physical socket number associated with the client handle.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>map_log_socket_t</b> are:</p> <pre>uint32_t    LogSocket;    /* logical socket number */ uint32_t    PhyAdapter;   /* physical adapter number */ uint32_t    PhySocket;    /* physical socket number */</pre> <p>The fields are defined as follows:</p> <p><b>LogSocket</b> Not used by this implementation of Card Services and can be set to any arbitrary value.</p> <p><b>PhyAdapter</b> Returns the physical adapter number, which is always 0 in the Solaris implementation of Card Services.</p> <p><b>PhySocket</b> Returns the physical socket number associated with the client handle. The physical socket number is typically used as part of an error or message string or if the client creates minor nodes based on the physical socket number.</p>
<b>RETURN VALUES</b>	<p><b>CS_SUCCESS</b> Successful operation.</p> <p><b>CS_BAD_HANDLE</b> Client handle is invalid.</p> <p><b>CS_UNSUPPORTED_FUNCTION</b> No PCMCIA hardware installed.</p>
<b>CONTEXT</b>	This function may be called from user or kernel context.

**SEE ALSO**

`csx_RegisterClient(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_MapMemPage – map the memory area on a PC Card												
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_MapMemPage(window_handle_t wh, map_mem_page_t *mp);</pre>												
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)												
<b>PARAMETERS</b>	<p><b>wh</b> Window handle returned from <b>csx_RequestWindow(9F)</b>.</p> <p><b>mp</b> Pointer to a <b>map_mem_page_t</b> structure.</p>												
<b>DESCRIPTION</b>	This function maps the memory area on a PC Card into a page of a window allocated with the <b>csx_RequestWindow(9F)</b> function.												
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>map_mem_page_t</b> are:</p> <pre>uint32_t    CardOffset;    /* card offset */ uint32_t    Page;         /* page number */</pre> <p>The fields are defined as follows:</p> <p><b>CardOffset</b> The absolute offset in bytes from the beginning of the PC Card to map into system memory.</p> <p><b>Page</b> Used internally by Card Services; clients must set this field to 0 before calling this function.</p>												
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>CS_SUCCESS</td> <td>Successful operation.</td> </tr> <tr> <td>CS_BAD_HANDLE</td> <td>Client handle is invalid.</td> </tr> <tr> <td>CS_BAD_OFFSET</td> <td>Offset is invalid.</td> </tr> <tr> <td>CS_BAD_PAGE</td> <td>Page is not zero.</td> </tr> <tr> <td>CS_NO_CARD</td> <td>No PC Card in socket.</td> </tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td> <td>No PCMCIA hardware installed.</td> </tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_HANDLE	Client handle is invalid.	CS_BAD_OFFSET	Offset is invalid.	CS_BAD_PAGE	Page is not zero.	CS_NO_CARD	No PC Card in socket.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.												
CS_BAD_HANDLE	Client handle is invalid.												
CS_BAD_OFFSET	Offset is invalid.												
CS_BAD_PAGE	Page is not zero.												
CS_NO_CARD	No PC Card in socket.												
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.												
<b>CONTEXT</b>	This function may be called from user or kernel context.												

**SEE ALSO**

`csx_ModifyWindow(9F)`, `csx_ReleaseWindow(9F)`,  
`csx_RequestWindow(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_ModifyConfiguration – modify socket and PC Card Configuration Register
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_ModifyConfiguration(client_handle_t ch, modify_config_t *mc);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>mc</b> Pointer to a <b>modify_config_t</b> structure.</p>
<b>DESCRIPTION</b>	This function allows a socket and PC Card configuration to be modified. This function can only modify a configuration requested via <b>csx_RequestConfiguration(9F)</b> .
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>modify_config_t</b> are:</p> <pre>uint32_t Socket; /* socket number */ uint32_t Attributes; /* attributes to modify */ uint32_t Vpp1; /* Vpp1 value */ uint32_t Vpp2; /* Vpp2 value */</pre> <p>The fields are defined as follows:</p> <p><b>Socket</b> Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p> <p><b>Attributes</b> This field is bit-mapped. The following bits are defined:</p> <p><b>CONF_ENABLE_IRQ_STEERING</b> Enable IRQ steering. Set to connect the PC Card IREQ line to a previously selected system interrupt.</p> <p><b>CONF_IRQ_CHANGE_VALID</b> IRQ change valid. Set to request the IRQ steering enable to be changed.</p> <p><b>CONF_VPP1_CHANGE_VALID</b> Vpp1 change valid. These bits are set to request a change to the corresponding voltage level for the PC Card.</p>

**CONF\_VPP2\_CHANGE\_VALID**

Vpp2 change valid. These bits are set to request a change to the corresponding voltage level for the PC Card.

**CONF\_VSOVERRIDE**

Override VS pins. For Low Voltage keyed cards, must be set if a client desires to apply a voltage inappropriate for this card to any pin. After card insertion and prior to the first `csx_RequestConfiguration(9F)` call for this client, the voltage levels applied to the card will be those specified by the Card Interface Specification. (See WARNINGS.)

Vpp1, Vpp2

Represent voltages expressed in tenths of a volt. Values from 0 to 25.5 volts may be set. To be valid, the exact voltage must be available from the system. To be compliant with the *PC Card 95 Standard*, *PCMCIA/JEIDA*, systems must always support 5.0 volts for both  $V_{CC}$  and  $V_{PP}$ . (See WARNINGS.)

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid or <code>csx_RequestConfiguration(9F)</code> not done.
CS_BAD_SOCKET	Error getting/setting socket hardware parameters.
CS_BAD_VPP	Requested $V_{PP}$ is not available on socket.
CS_NO_CARD	No PC Card in socket.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

`csx_RegisterClient(9F)`, `csx_ReleaseConfiguration(9F)`,  
`csx_ReleaseIO(9F)`, `csx_ReleaseIRQ(9F)`,  
`csx_RequestConfiguration(9F)`, `csx_RequestIO(9F)`,  
`csx_RequestIRQ(9F)`

*PC Card 95 Standard*, *PCMCIA/JEIDA*

**WARNINGS**

1. `CONF_VSOVERRIDE` is provided for clients that have a need to override the information provided in the CIS. The client must exercise caution when setting this as it overrides any voltage level protection provided by Card Services.
2. Using `csx_ModifyConfiguration()` to set  $V_{pp}$  to 0 volts may result in the loss of a PC Card's state. Any client setting  $V_{pp}$  to 0 volts is responsible for insuring that the PC Card's state is restored when power is re-applied to the card.

**NOTES**

Mapped IO addresses can only be changed by first releasing the current configuration and IO resources with `csx_ReleaseConfiguration(9F)` and `csx_ReleaseIO(9F)`, requesting new IO resources and a new configuration with `csx_RequestIO(9F)`, followed by `csx_RequestConfiguration(9F)`.

IRQ priority can only be changed by first releasing the current configuration and IRQ resources with `csx_ReleaseConfiguration(9F)` and `csx_ReleaseIRQ(9F)`, requesting new IRQ resources and a new configuration with `csx_RequestIRQ(9F)`, followed by `csx_RequestConfiguration(9F)`.

$V_{cc}$  can not be changed using `csx_ModifyConfiguration()`.  $V_{cc}$  may be changed by first invoking `csx_ReleaseConfiguration(9F)`, followed by `csx_RequestConfiguration(9F)` with a new  $V_{cc}$  value.

<b>NAME</b>	csx_ModifyWindow – modify window attributes
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_ModifyWindow(window_handle_t wh, modify_win_t *mw);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>wh</b> Window handle returned from <code>csx_RequestWindow(9F)</code>.</p> <p><b>mw</b> Pointer to a <code>modify_win_t</code> structure.</p>
<b>DESCRIPTION</b>	<p>This function modifies the attributes of a window allocated by the <code>csx_RequestWindow(9F)</code> function.</p> <p>Only some of the window attributes or the access speed field may be modified by this request. The <code>csx_MapMemPage(9F)</code> function is also used to set the offset into PC Card memory to be mapped into system memory for paged windows. The <code>csx_RequestWindow(9F)</code> and <code>csx_ReleaseWindow(9F)</code> functions must be used to change the window base or size.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>modify_win_t</code> are:</p> <pre>uint32_t      Attributes;          /* window flags */ uint32_t      AccessSpeed;        /* window access speed */</pre> <p>The fields are defined as follows:</p> <p><b>Attributes</b> This field is bit-mapped and defined as follows:</p> <p><b>WIN_MEMORY_TYPE_CM</b> Window points to Common Memory area. Set this to map the window to Common Memory.</p> <p><b>WIN_MEMORY_TYPE_AM</b> Window points to Attribute Memory area. Set this to map the window to Attribute Memory.</p> <p><b>WIN_ENABLE</b> Enable Window. The client must set this to enable the window.</p>

**WIN\_ACCESS\_SPEED\_VALID**

AccessSpeed valid. The client must set this when the AccessSpeed field has a value that the client wants set for the window.

AccessSpeed The bit definitions for this field use the format of the extended speed byte of the Device ID tuple. If the mantissa is 0 (noted as reserved in the *PC Card 95 Standard*), the lower bits are a binary code representing a speed from the following list:

Code	Speed
<b>0</b>	Reserved: do not use
<b>1</b>	250 nsec
<b>2</b>	200 nsec
<b>3</b>	150 nsec
<b>4</b>	100 nsec
<b>5 - 7</b>	Reserved: do not use

It is recommended that clients use the **csx\_ConvertSpeed(9F)** function to generate the appropriate AccessSpeed values rather than manually perturbing the AccessSpeed field.

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Window handle is invalid.
CS_NO_CARD	No PC Card in socket.
CS_BAD_OFFSET	Error getting/setting window hardware parameters.
CS_BAD_WINDOW	Error getting/setting window hardware parameters.
CS_BAD_SPEED	AccessSpeed is invalid.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** | This function may be called from user or kernel context.

**SEE ALSO** | `csx_ConvertSpeed(9F)`, `csx_MapMemPage(9F)`, `csx_ReleaseWindow(9F)`,  
`csx_RequestWindow(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_BATTERY – parse the Battery Replacement Date tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_BATTERY(client_handle_t ch, tuple_t *tu, cistpl_date_t *cb);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code>.</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure (see <code>tuple(9S)</code>) returned by a call to <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>.</p> <p><b>cb</b> Pointer to a <code>cistpl_battery_t</code> structure which contains the parsed CISTPL_BATTERY tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Battery Replacement Date tuple, CISTPL_BATTERY, into a form usable by PC Card drivers.</p> <p>The CISTPL_BATTERY tuple is an optional tuple which shall be present only in PC Cards with battery-backed storage. It indicates the date on which the battery was replaced, and the date on which the battery is expected to need replacement. Only one CISTPL_BATTERY tuple is allowed per PC Card.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>cistpl_date_t</code> are:</p> <pre>uint32_t    rday;    /* date battery last replaced */ uint32_t    xday;    /* date battery due for replacement */</pre> <p>The fields are defined as follows:</p> <p><code>rday</code> This field indicates the date on which the battery was last replaced.</p> <p><code>xday</code> This field indicates the date on which the battery should be replaced.</p>
<b>RETURN VALUES</b>	<p>CS_SUCCESS Successful operation.</p> <p>CS_BAD_HANDLE Client handle is invalid.</p> <p>CS_UNKNOWN_TUPLE Parser does not know how to parse tuple.</p>

CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_BYTEORDER – parse the Byte Order tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_BYTEORDER(client_handle_t ch, tuple_t *tu, cistpl_byteorder_t *cbo);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code>.</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure (see <code>tuple(9S)</code>) returned by a call to <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>.</p> <p><b>cbo</b> Pointer to a <code>cistpl_byteorder_t</code> structure which contains the parsed CISTPL_BYTEORDER tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Byte Order tuple, CISTPL_BYTEORDER, into a form usable by PC Card drivers.</p> <p>The CISTPL_BYTEORDER tuple shall only appear in a partition tuple set for a memory-like partition. It specifies two parameters: the order for multi-byte data, and the order in which bytes map into words for 16-bit cards.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>cistpl_byteorder_t</code> are:</p> <pre>uint32_t    order;    /* byte order code */ uint32_t    map;     /* byte mapping code */</pre> <p>The fields are defined as follows:</p> <p><code>order</code> This field specifies the byte order for multi-byte numeric data.</p> <p><b>TPLBYTEORD_LOW</b> Little endian order</p> <p><b>TPLBYTEORD_VS</b> Vendor specific</p> <p><code>map</code> This field specifies the byte mapping for 16-bit or wider cards.</p>

	<b>TPLBYTEMAP_LOW</b>	
	Byte zero is least significant byte	
	<b>TPLBYTEMAP_HIGH</b>	
	Byte zero is most significant byte	
	<b>TPLBYTEMAP_VS</b>	
	Vendor specific mapping	
<b>RETURN VALUES</b>	<b>CS_SUCCESS</b>	Successful operation.
	<b>CS_BAD_HANDLE</b>	Client handle is invalid.
	<b>CS_UNKNOWN_TUPLE</b>	Parser does not know how to parse tuple.
	<b>CS_NO_CARD</b>	No PC Card in socket.
	<b>CS_NO_CIS</b>	No Card Information Structure (CIS) PC Card.
	<b>CS_UNSUPPORTED_FUNCTION</b>	No PCMCIA hardware installed.
<b>CONTEXT</b>	This function may be called from user or kernel context.	
<b>SEE ALSO</b>	<b>csx_GetFirstTuple(9F)</b> , <b>csx_GetTupleData(9F)</b> , <b>csx_RegisterClient(9F)</b> , <b>csx_ValidateCIS(9F)</b> , <b>tuple(9S)</b>	
	<i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

<b>NAME</b>	csx_Parse_CISTPL_CFTABLE_ENTRY – parse 16-bit Card Configuration Table Entry tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_CFTABLE_ENTRY(client_handle_t ch, tuple_t *tu, cistpl_cftable_entry_t *cft);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>cft</b> Pointer to a <b>cistpl_cftable_entry_t</b> structure which contains the parsed CISTPL_CFTABLE_ENTRY tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the 16 bit Card Configuration Table Entry tuple, CISTPL_CFTABLE_ENTRY, into a form usable by PC Card drivers.</p> <p>The CISTPL_CFTABLE_ENTRY tuple is used to describe each possible configuration of a PC Card and to distinguish among the permitted configurations. The CISTPL_CONFIG tuple must precede all CISTPL_CFTABLE_ENTRY tuples.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_cftable_entry_t</b> are:</p> <pre>uint32_t          flags;      /* valid descriptions */ uint32_t          ifc;       /* interface description */                                /* information */ uint32_t          pin;       /* values for PRR */ uint32_t          index;    /* configuration index number */ cistpl_cftable_entry_pd_t  pd;  /* power requirements */                                /* description */ cistpl_cftable_entry_speed_t  speed; /* device speed description */ cistpl_cftable_entry_io_t   io;   /* device I/O map */ cistpl_cftable_entry_irq_t  irq;  /* device IRQ utilization */ cistpl_cftable_entry_mem_t  mem;  /* device memory space */ cistpl_cftable_entry_misc_t  misc; /* miscellaneous                                /* device features */</pre> <p>The <b>flags</b> field is defined and bit-mapped as follows:</p> <p><b>CISTPL_CFTABLE_TPCE_DEFAULT</b></p> <p>This is a default configuration</p>

**CISTPL\_CFTABLE\_TPCE\_IF**

If configuration byte exists

**CISTPL\_CFTABLE\_TPCE\_FS\_PWR**

Power information exists

**CISTPL\_CFTABLE\_TPCE\_FS\_TD**

Timing information exists

**CISTPL\_CFTABLE\_TPCE\_FS\_IO**

I/O information exists

**CISTPL\_CFTABLE\_TPCE\_FS\_IRQ**

IRQ information exists

**CISTPL\_CFTABLE\_TPCE\_FS\_MEM**

MEM space information exists

**CISTPL\_CFTABLE\_TPCE\_FS\_MISC**

MISC information exists

**CISTPL\_CFTABLE\_TPCE\_FS\_STCE\_EV**

STCE\_EV exists

**CISTPL\_CFTABLE\_TPCE\_FS\_STCE\_PD**

STCE\_PD exists

If the `CISTPL_CFTABLE_TPCE_IF` flag is set, the `ifc` field is bit-mapped and defined as follows:

**CISTPL\_CFTABLE\_TPCE\_IF\_MEMORY**

Memory interface

**CISTPL\_CFTABLE\_TPCE\_IF\_IO\_MEM**

IO and memory

**CISTPL\_CFTABLE\_TPCE\_IF\_CUSTOM\_0**

Custom interface 0

**CISTPL\_CFTABLE\_TPCE\_IF\_CUSTOM\_1**

Custom interface 1

**CISTPL\_CFTABLE\_TPCE\_IF\_CUSTOM\_2**

Custom interface 2

**CISTPL\_CFTABLE\_TPCE\_IF\_CUSTOM\_3**

Custom interface 3

**CISTPL\_CFTABLE\_TPCE\_IF\_MASK**

Interface type mask

**CISTPL\_CFTABLE\_TPCE\_IF\_BVD**

BVD active in PRR

**CISTPL\_CFTABLE\_TPCE\_IF\_WP**

WP active in PRR

**CISTPL\_CFTABLE\_TPCE\_IF\_RDY**

RDY active in PRR

**CISTPL\_CFTABLE\_TPCE\_IF\_MWAIT**

WAIT - mem cycles

pin is a value for the Pin Replacement Register.

index is a configuration index number.

The structure members of `cistpl_cftable_entry_pd_t` are:

```
uint32_t          flags;          /* which descriptions are valid */
cistpl_cftable_entry_pwr_t pd_vcc; /* VCC power description */
cistpl_cftable_entry_pwr_t pd_vpp1; /* Vpp1 power description */
cistpl_cftable_entry_pwr_t pd_vpp2; /* Vpp2 power description */
```

This `flags` field is bit-mapped and defined as follows:

**CISTPL\_CFTABLE\_TPCE\_FS\_PWR\_VCC**

Vcc description valid

**CISTPL\_CFTABLE\_TPCE\_FS\_PWR\_VPP1**

Vpp1 description valid

**CISTPL\_CFTABLE\_TPCE\_FS\_PWR\_VPP2**

Vpp2 description valid

The structure members of `cistpl_cftable_entry_pwr_t` are:

```

uint32_t    nomV;           /* nominal supply voltage */
uint32_t    nomV_flags;
uint32_t    minV;          /* minimum supply voltage */
uint32_t    minV_flags;
uint32_t    maxV;          /* maximum supply voltage */
uint32_t    maxV_flags;
uint32_t    staticI;       /* continuous supply current */
uint32_t    staticI_flags;
uint32_t    avgI;          /* max current required averaged over 1 sec. */
uint32_t    avgI_flags;
uint32_t    peakI;         /* max current required averaged over 10ms */
uint32_t    peakI_flags;
uint32_t    pdownI;        /* power down supply current required */
uint32_t    pdownI_flags;

```

`nomV`, `minV`, `maxV`, `staticI`, `avgI`, `peakI`, `peakI_flag`, and `pdownI` are defined and bit-mapped as follows:

**CISTPL\_CFTABLE\_PD\_NOMV**

Nominal supply voltage

**CISTPL\_CFTABLE\_PD\_MINV**

Minimum supply voltage

**CISTPL\_CFTABLE\_PD\_MAXV**

Maximum supply voltage

**CISTPL\_CFTABLE\_PD\_STATICI**

Continuous supply current

**CISTPL\_CFTABLE\_PD\_AVGI**

Maximum current required averaged over 1 second

**CISTPL\_CFTABLE\_PD\_PEAKEI**

Maximum current required averaged over 10mS

#### **CISTPL\_CFTABLE\_PD\_PDOWNI**

Power down supply current required

nomV\_flags, minV\_flags, maxV\_flags, staticI\_flags, avgI\_flags, peakI\_flags, and pdownI\_flags are defined and bit-mapped as follows:

#### **CISTPL\_CFTABLE\_PD\_EXISTS**

This parameter exists

#### **CISTPL\_CFTABLE\_PD\_MUL10**

Multiply return value by 10

#### **CISTPL\_CFTABLE\_PD\_NC\_SLEEP**

No connection on sleep/power down

#### **CISTPL\_CFTABLE\_PD\_ZERO**

Zero value required

#### **CISTPL\_CFTABLE\_PD\_NC**

No connection ever

The structure members of `cistpl_cftable_entry_speed_t` are:

```
uint32_t    flags;           /* which timing information is present */
uint32_t    wait;           /* max WAIT time in device speed format */
uint32_t    nS_wait;        /* max WAIT time in nS */
uint32_t    rdybsy;         /* max RDY/BSY time in device speed format */
uint32_t    nS_rdybsy;      /* max RDY/BSY time in nS */
uint32_t    rsvd;           /* max RSVD time in device speed format */
uint32_t    nS_rsvd;        /* max RSVD time in nS */
```

The `flags` field is bit-mapped and defined as follows:

#### **CISTPL\_CFTABLE\_TPCE\_FS\_TD\_WAIT**

WAIT timing exists

#### **CISTPL\_CFTABLE\_TPCE\_FS\_TD\_RDY**

RDY/BSY timing exists

#### **CISTPL\_CFTABLE\_TPCE\_FS\_TD\_RSVD**

RSVD timing exists

The structure members of `cistpl_cftable_entry_io_t` are:

```
uint32_t    flags;          /* direct copy of TPCE_IO byte in tuple */
uint32_t    addr_lines;    /* number of decoded I/O address lines */
uint32_t    ranges;        /* number of I/O ranges */
cistpl_cftable_entry_io_range_t
            range[CISTPL_CFTABLE_ENTRY_MAX_IO_RANGES];
```

The `flags` field is defined and bit-mapped as follows:

**CISTPL\_CFTABLE\_TPCE\_FS\_IO\_BUS**

Bus width mask

**CISTPL\_CFTABLE\_TPCE\_FS\_IO\_BUS8**

8-bit flag

**CISTPL\_CFTABLE\_TPCE\_FS\_IO\_BUS16**

16-bit flag

**CISTPL\_CFTABLE\_TPCE\_FS\_IO\_RANGE**

IO address ranges exist

The structure members of `cistpl_cftable_entry_io_range_t` are:

```
uint32_t    addr;          /* I/O start address */
uint32_t    length;        /* I/O register length */
```

The structure members of `cistpl_cftable_entry_irq_t` are:

```
uint32_t    flags;          /* direct copy of TPCE_IR byte in tuple */
uint32_t    irqs;          /* bit mask for each allowed IRQ */
```

The structure members of `cistpl_cftable_entry_mem_t` are:

```
uint32_t    flags;          /* memory descriptor type and host addr info */
uint32_t    windows;        /* number of memory space descriptors */
cistpl_cftable_entry_mem_window_t
            window[CISTPL_CFTABLE_ENTRY_MAX_MEM_WINDOWS];
```

The `flags` field is defined and bit-mapped as follows:

**CISTPL\_CFTABLE\_TPCE\_FS\_MEM3**

Space descriptors

**CISTPL\_CFTABLE\_TPCE\_FS\_MEM2**

host\_addr=card\_addr

**CISTPL\_CFTABLE\_TPCE\_FS\_MEM1**

Card address=0 any host address

**CISTPL\_CFTABLE\_TPCE\_FS\_MEM\_HOST**

If host address is present in MEM3

The structure members of `cistpl_cftable_entry_mem_window_t` are:

```
uint32_t    length;    /* length of this window */
uint32_t    card_addr; /* card address */
uint32_t    host_addr; /* host address */
```

The structure members of `cistpl_cftable_entry_misc_t` are:

```
uint32_t    flags;    /* miscellaneous features flags */
```

The `flags` field is defined and bit-mapped as follows:

**CISTPL\_CFTABLE\_TPCE\_MI\_MTC\_MASK**

Max twin cards mask

**CISTPL\_CFTABLE\_TPCE\_MI\_AUDIO**

Audio on BVD2

**CISTPL\_CFTABLE\_TPCE\_MI\_READONLY**

R/O storage

**CISTPL\_CFTABLE\_TPCE\_MI\_PWRDOWN**

Powerdown capable

**CISTPL\_CFTABLE\_TPCE\_MI\_DRQ\_MASK**

DMAREQ mask

**CISTPL\_CFTABLE\_TPCE\_MI\_DRQ\_SPK**

DMAREQ on SPKR  
**CISTPL\_CFTABLE\_TPCE\_MI\_DRQ\_IOIS**  
 DMAREQ on IOIS16  
**CISTPL\_CFTABLE\_TPCE\_MI\_DRQ\_INP**  
 DMAREQ on INPACK  
**CISTPL\_CFTABLE\_TPCE\_MI\_DMA\_8**  
 DMA width 8 bits  
**CISTPL\_CFTABLE\_TPCE\_MI\_DMA\_16**  
 DMA width 16 bits

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

**csx\_GetFirstTuple(9F)**, **csx\_GetTupleData(9F)**,  
**csx\_Parse\_CISTPL\_CONFIG(9F)**, **csx\_RegisterClient(9F)**,  
**csx\_ValidateCIS(9F)**, **tuple(9S)**

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_CONFIG – parse Configuration tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_CONFIG(client_handle_t ch, tuple_t *tu, cistpl_config_t *cc);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>cc</b> Pointer to a <b>cistpl_config_t</b> structure which contains the parsed CISTPL_CONFIG tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Configuration tuple, CISTPL_CONFIG, into a form usable by PC Card drivers. The CISTPL_CONFIG tuple is used to describe the general characteristics of 16-bit PC Cards containing I/O devices or using custom interfaces. It may also describe PC Cards, including Memory Only cards, which exceed nominal power supply specifications, or which need descriptions of their power requirements or other information.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_config_t</b> are:</p> <pre>uint32_t    present;    /* register present flags */ uint32_t    nr;        /* number of config registers found */ uint32_t    hr;        /* highest config register index found */ uint32_t    regs[CISTPL_CONFIG_MAX_CONFIG_REGS]; /* reg offsets */ uint32_t    base;      /* base offset of config registers */ uint32_t    last;      /* last config index */</pre> <p>The fields are defined as follows:</p> <p><b>present</b> This field indicates which configuration registers are present on the PC Card.</p> <p><b>CONFIG_OPTION_REG_PRESENT</b> Configuration Option Register present</p> <p><b>CONFIG_STATUS_REG_PRESENT</b> Configuration Status Register present</p>

**CONFIG\_PINREPL\_REG\_PRESENT**

Pin Replacement Register present

**CONFIG\_COPY\_REG\_PRESENT**

Copy Register present

**CONFIG\_EXSTAT\_REG\_PRESENT**

Extended Status Register present

**CONFIG\_IOBASE0\_REG\_PRESENT**

IO Base 0 Register present

**CONFIG\_IOBASE1\_REG\_PRESENT**

IO Base 1 Register present

**CONFIG\_IOBASE2\_REG\_PRESENT**

IO Base2 Register present

**CONFIG\_IOBASE3\_REG\_PRESENT**

IO Base3 Register present

**CONFIG\_IOLIMIT\_REG\_PRESENT**

IO Limit Register present

**nr** This field specifies the number of configuration registers that are present on the PC Card.

**hr** This field specifies the highest configuration register number that is present on the PC Card.

**regs** This array contains the offset from the start of Attribute Memory space for each configuration register that is present on the PC Card. If a configuration register is not present on the PC Card, the value in the corresponding entry in the **regs** array is undefined.

**base** This field contains the offset from the start of Attribute Memory space to the base of the PC Card configuration register space.

**last** This field contains the value of the last valid configuration index for this PC Card.

**RETURN VALUES**

**CS\_SUCCESS** Successful operation.

CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_Parse_CISTPL_CFTABLE_ENTRY(9F)`, `csx_RegisterClient(9F)`,  
`csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

**NOTES** PC Card drivers should not attempt to use configurations beyond the "last" member in the `cistpl_config_t` structure.

<b>NAME</b>	csx_Parse_CISTPL_DATE – parse the Card Initialization Date tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_DATE(client_handle_t ch, tuple_t *tu, cistpl_date_t *cd);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>cd</b> Pointer to a <b>cistpl_date_t</b> structure which contains the parsed CISTPL_DATE tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Card Initialization Date tuple, CISTPL_DATE, into a form usable by PC Card drivers.</p> <p>The CISTPL_DATE tuple is an optional tuple. It indicates the date and time at which the card was formatted. Only one CISTPL_DATE tuple is allowed per PC Card.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_date_t</b> are:</p> <pre>uint32_t    time; uint32_t    day;</pre> <p>The fields are defined as follows:</p> <p><b>time</b> This field indicates the time at which the PC Card was initialized.</p> <p><b>day</b> This field indicates the date the PC Card was initialized.</p>
<b>RETURN VALUES</b>	<p>CS_SUCCESS Successful operation.</p> <p>CS_BAD_HANDLE Client handle is invalid.</p> <p>CS_UNKNOWN_TUPLE Parser does not know how to parse tuple.</p> <p>CS_NO_CARD No PC Card in socket.</p>

CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_DEVICE, csx_Parse_CISTPL_DEVICE_A, csx_Parse_CISTPL_DEVICE_OC, csx_Parse_CISTPL_DEVICE_OA – parse Device Information tuples
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_DEVICE(client_handle_t ch, tuple_t * tu, cistpl_device_t * cd);  int32_t csx_Parse_CISTPL_DEVICE_A(client_handle_t ch, tuple_t * tu, cistpl_device_t * cd);  int32_t csx_Parse_CISTPL_DEVICE_OC(client_handle_t ch, tuple_t * tu, cistpl_device_t * cd);  int32_t csx_Parse_CISTPL_DEVICE_OA(client_handle_t ch, tuple_t * tu, cistpl_device_t * cd);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b> )returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b> .</p> <p><b>cd</b> Pointer to a <b>cistpl_device_t</b> structure which contains the parsed CISTPL_DEVICE, CISTPL_DEVICE_A, CISTPL_DEVICE_OC, or CISTPL_DEVICE_OA tuple information upon return from these functions, respectively.</p>
<b>DESCRIPTION</b>	<p><b>csx_Parse_CISTPL_DEVICE()</b> and <b>csx_Parse_CISTPL_DEVICE_A()</b> parse the 5 volt Device Information tuples, CISTPL_DEVICE and CISTPL_DEVICE_A, respectively, into a form usable by PC Card drivers.</p> <p><b>csx_Parse_CISTPL_DEVICE_OC()</b> and <b>csx_Parse_CISTPL_DEVICE_OA()</b> parse the Other Condition Device Information tuples, CISTPL_DEVICE_OC and CISTPL_DEVICE_OA, respectively, into a form usable by PC Card drivers.</p> <p>The CISTPL_DEVICE and CISTPL_DEVICE_A tuples are used to describe the card's device information, such as device speed, device size, device type, and address space layout information for Common Memory or Attribute Memory space, respectively.</p>

**STRUCTURE MEMBERS**

The CISTPL\_DEVICE\_OC and CISTPL\_DEVICE\_OA tuples are used to describe the information about the card's device under a set of operating conditions for Common Memory or Attribute Memory space, respectively.

The structure members of `cistpl_device_t` are:

```
uint32_t      num_devices;          /* number of devices found */
cistpl_device_node_t devnode[CISTPL_DEVICE_MAX_DEVICES];
```

The structure members of `cistpl_device_node_t` are:

```
uint32_t      flags;                /* flags specific to this device */
uint32_t      speed;                /* device speed in device speed code format */
uint32_t      nS_speed;            /* device speed in nS */
uint32_t      type;                /* device type */
uint32_t      size;                /* device size */
uint32_t      size_in_bytes;       /* device size in bytes */
```

The fields are defined as follows:

`flags` This field indicates whether or not the device is writable, and describes a Vcc voltage at which the PC Card can be operated.

**CISTPL\_DEVICE\_WPS**

Write Protect Switch bit is set

Bits which are applicable only for CISTPL\_DEVICE\_OC and CISTPL\_DEVICE\_OA are:

**CISTPL\_DEVICE\_OC\_MWAIT**

Use MWAIT

**CISTPL\_DEVICE\_OC\_Vcc\_MASK**

Mask for Vcc value

**CISTPL\_DEVICE\_OC\_Vcc5**

5.0 volt operation

**CISTPL\_DEVICE\_OC\_Vcc33**

3.3 volt operation

**CISTPL\_DEVICE\_OC\_VccXX**

X.X volt operation

**CISTPL\_DEVICE\_OC\_VccYY**

Y.Y volt operation

speed	The device speed value described in the device speed code unit. If this field is set to <code>CISTPL_DEVICE_SPEED_SIZE_IGNORE</code> , then the speed information will be ignored.
nS_speed	The device speed value described in nanosecond units.
size	The device size value described in the device size code unit. If this field is set to <code>CISTPL_DEVICE_SPEED_SIZE_IGNORE</code> , then the size information will be ignored.
size_in_bytes	The device size value described in byte units.
type	This is the device type code field which is defined as follows:  <b>CISTPL_DEVICE_DTYPE_NULL</b> No device  <b>CISTPL_DEVICE_DTYPE_ROM</b> Masked ROM  <b>CISTPL_DEVICE_DTYPE_OTPROM</b> One Time Programmable ROM  <b>CISTPL_DEVICE_DTYPE_EPROM</b> UV EPROM  <b>CISTPL_DEVICE_DTYPE_EEPROM</b> EEPROM  <b>CISTPL_DEVICE_DTYPE_FLASH</b> FLASH  <b>CISTPL_DEVICE_DTYPE_SRAM</b> Static RAM  <b>CISTPL_DEVICE_DTYPE_DRAM</b> Dynamic RAM

**CISTPL\_DEVICE\_DTYPE\_FUNCSPEC**

Function-specific memory address range

**CISTPL\_DEVICE\_DTYPE\_EXTEND**

Extended type follows

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

These functions may be called from user or kernel context.

**SEE ALSO**

**csx\_GetFirstTuple(9F)** , **csx\_GetTupleData(9F)** ,  
**csx\_Parse\_CISTPL\_JEDEC\_C(9F)** , **csx\_RegisterClient(9F)** ,  
**csx\_ValidateCIS(9F)** , **tuple(9S)**

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_DEVICEGEO – parse the Device Geo tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_DEVICEGEO(client_handle_t ch, tuple_t *tp, cistpl_devicegeo_t *pt);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tp</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>pt</b> Pointer to a <b>cistpl_devicegeo_t</b> structure which contains the parsed Device Geo tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Device Geo tuple, <b>CISTPL_DEVICEGEO</b>, into a form usable by PC Card drivers.</p> <p>The <b>CISTPL_DEVICEGEO</b> tuple describes the device geometry of common memory partitions.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_devicegeo_t</b> are:</p> <pre>uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].bus; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].ebs; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].rbs; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].wbs; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].part; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].hwil;</pre> <p>The fields are defined as follows:</p> <pre>info[CISTPL_DEVICEGEO_MAX_PARTITIONS].bus</pre> <p>This field indicates the card interface width in bytes for the given partition.</p> <pre>info[CISTPL_DEVICEGEO_MAX_PARTITIONS].ebs</pre> <p>This field indicates the minimum erase block size for the given partition.</p> <pre>info[CISTPL_DEVICEGEO_MAX_PARTITIONS].rbs</pre> <p>This field indicates the minimum read block size for the given partition.</p>

```
info[CISTPL_DEVICEGEO_MAX_PARTITIONS].wbs
```

This field indicates the minimum write block size for the given partition.

```
info[CISTPL_DEVICEGEO_MAX_PARTITIONS].part
```

This field indicates the segment partition subdivisions for the given partition.

```
info[CISTPL_DEVICEGEO_MAX_PARTITIONS].hwil
```

This field indicates the hardware interleave

## RETURN VALUES

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

## CONTEXT

This function may be called from user or kernel context.

## SEE ALSO

**csx\_GetFirstTuple(9F)**, **csx\_GetNextTuple(9F)**,  
**csx\_GetTupleData(9F)**, **csx\_Parse\_CISTPL\_DEVICEGEO\_A(9F)**,  
**csx\_RegisterClient(9F)**, **tuple(9S)**

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_DEVICEGEO_A – parse the Device Geo A tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_DEVICEGEO_A(client_handle_t ch, tuple_t *tp, cistpl_devicegeo_t *pt);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tp</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>pt</b> Pointer to a <b>cistpl_devicegeo_t</b> structure which contains the parsed Device Geo A tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Device Geo A tuple, <b>CISTPL_DEVICEGEO_A</b>, into a form usable by PC Card drivers.</p> <p>The <b>CISTPL_DEVICEGEO_A</b> tuple describes the device geometry of attribute memory partitions.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_devicegeo_t</b> are:</p> <pre>uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].bus; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].ebs; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].rbs; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].wbs; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].part; uint32_t    info[CISTPL_DEVICEGEO_MAX_PARTITIONS].hwil;</pre> <p>The fields are defined as follows:</p> <pre>info[CISTPL_DEVICEGEO_MAX_PARTITIONS].bus</pre> <p>This field indicates the card interface width in bytes for the given partition.</p> <pre>info[CISTPL_DEVICEGEO_MAX_PARTITIONS].ebs</pre> <p>This field indicates the minimum erase block size for the given partition.</p> <pre>info[CISTPL_DEVICEGEO_MAX_PARTITIONS].rbs</pre> <p>This field indicates the minimum read block size for the given partition.</p>

`info[CISTPL_DEVICEGEO_MAX_PARTITIONS].wbs`

This field indicates the minimum write block size for the given partition.

`info[CISTPL_DEVICEGEO_MAX_PARTITIONS].part`

This field indicates the segment partition subdivisions for the given partition.

`info[CISTPL_DEVICEGEO_MAX_PARTITIONS].hwil`

This field indicates the hardware interleave for the given partition.

## RETURN VALUES

<code>CS_SUCCESS</code>	Successful operation.
<code>CS_BAD_HANDLE</code>	Client handle is invalid.
<code>CS_UNKNOWN_TUPLE</code>	Parser does not know how to parse tuple.
<code>CS_NO_CARD</code>	No PC Card in socket.
<code>CS_NO_CIS</code>	No Card Information Structure (CIS) on PC Card.
<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.

## CONTEXT

This function may be called from user or kernel context.

## SEE ALSO

`csx_GetFirstTuple(9F)`, `csx_GetNextTuple(9F)`,  
`csx_GetTupleData(9F)`, `csx_Parse_CISTPL_DEVICEGEO(9F)`,  
`csx_RegisterClient(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_FORMAT – parse the Data Recording Format tuple
<b>SYNOPSIS</b>	#include <sys/pccard.h>  int32_t csx_Parse_CISTPL_FORMAT(client_handle_t ch, tuple_t *tu, cistpl_format_t *pt);
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>pt</b> Pointer to a <b>cistpl_format_t</b> structure which contains the parsed CISTPL_FORMAT tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Data Recording Format tuple, CISTPL_FORMAT, into a form usable by PC Card drivers.</p> <p>The CISTPL_FORMAT tuple indicates the data recording format for a device partition.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_format_t</b> are:</p> <pre>uint32_t    type; uint32_t    edc_length; uint32_t    edc_type; uint32_t    offset; uint32_t    nbytes; uint32_t    dev.disk.bksize; uint32_t    dev.disk.nblocks; uint32_t    dev.disk.edcloc; uint32_t    dev.mem.flags; uint32_t    dev.mem.reserved; caddr_t     dev.mem.address; uint32_t    dev.mem.edcloc;</pre> <p>The fields are defined as follows:</p> <p><b>type</b> This field indicates the type of device:</p> <p style="text-align: center;"><b>TPLFMTTYPE_DISK</b> disk-like device</p>

	<b>TPLFMTTYPE_MEM</b> memory-like device
	<b>TPLFMTTYPE_VS</b> vendor-specific device
edc_length	This field indicates the error detection code length.
edc_type	This field indicates the error detection code type.
offset	This field indicates the offset of the first byte of data in this partition.
nbytes	This field indicates the number of bytes of data in this partition
dev.disk.bksize	This field indicates the block size, for disk devices.
dev.disk.nblocks	This field indicates the number of blocks, for disk devices.
dev.disk.edcloc	This field indicates the location of the error detection code, for disk devices.
dev.mem.flags	This field provides flags, for memory devices. Valid flags are:
	<b>TPLFMTFLAGS_ADDR</b> address is valid
	<b>TPLFMTFLAGS_AUTO</b> automatically map memory region
dev.mem.reserved	This field is reserved.
dev.mem.address	This field indicates the physical address, for memory devices.
dev.mem.edcloc	This field indicates the location of the error detection code, for memory devices.
<b>RETURN VALUES</b>	
CS_SUCCESS	Successful operation.

CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_FUNCE – parse Function Extension tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_FUNCE(client_handle_t ch, tuple_t *tu, cistpl_funce_t *cf, uint32_t fid);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code>.</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure (see <code>tuple(9S)</code>) returned by a call to <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>.</p> <p><b>cf</b> Pointer to a <code>cistpl_funce_t</code> structure which contains the parsed CISTPL_FUNCE tuple information upon return from this function.</p> <p><b>fid</b> The function ID code to which this CISTPL_FUNCE tuple refers. See <code>csx_Parse_CISTPL_FUNCID(9F)</code>.</p>
<b>DESCRIPTION</b>	<p>This function parses the Function Extension tuple, CISTPL_FUNCE, into a form usable by PC Card drivers.</p> <p>The CISTPL_FUNCE tuple is used to describe information about a specific PC Card function. The information provided is determined by the Function Identification tuple, CISTPL_FUNCID, that is being extended. Each function has a defined set of extension tuples.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>cistpl_funce_t</code> are:</p> <pre>uint32_t    function;           /* type of extended data */ uint32_t    subfunction; union {     struct serial {         uint32_t  ua;           /* UART in use */         uint32_t  uc;           /* UART capabilities */     } serial;     struct modem {         uint32_t  fc;           /* supported flow control methods */         uint32_t  cb;           /* size of DCE command buffer */         uint32_t  eb;           /* size of DCE to DCE buffer */         uint32_t  tb;           /* size of DTE to DCE buffer */     } modem;     struct data_modem {         uint32_t  ud;           /* highest data rate */         uint32_t  ms;           /* modulation standards */         uint32_t  em;           /* err correct proto and non-CCITT modulation */         uint32_t  dc;           /* data compression protocols */     } data_modem; }</pre>

```

        uint32_t cm;      /* command protocols */
        uint32_t ex;      /* escape mechanisms */
        uint32_t dy;      /* standardized data encryption */
        uint32_t ef;      /* miscellaneous end user features */
        uint32_t ncd;     /* number of country codes */
        uchar_t  cd[16];  /* CCITT country code */
    } data_modem;
    struct fax {
        uint32_t uf;      /* highest data rate in DTE/UART */
        uint32_t fm;      /* CCITT modulation standards */
        uint32_t fy;      /* standardized data encryption */
        uint32_t fs;      /* feature selection */
        uint32_t ncf;     /* number of country codes */
        uchar_t  cf[16];  /* CCITT country codes */
    } fax;
    struct voice {
        uint32_t uv;      /* highest data rate */
        uint32_t nsr;     /* voice sampling rates (*100) */
        uint32_t sr[16];  /* voice sampling rates (*100) */
        uint32_t nss;     /* voice sample sizes (*10) */
        uint32_t ss[16];  /* voice sample sizes (*10) */
        uint32_t nsc;     /* voice compression methods */
        uint32_t sc[16];  /* voice compression methods */
    } voice;
    struct lan {
        uint32_t tech;    /* network technology */
        uint32_t speed;   /* media bit or baud rate */
        uint32_t media;   /* network media supported */
        uint32_t con;     /* open/closed connector standard */
        uint32_t id_sz;   /* length of lan station id */
        uchar_t  id[16]; /* station ID */
    } lan;
} data;

```

The fields are defined as follows:

**function** This field identifies the type of extended information provided about a function by the CISTPL\_FUNCE tuple. This field is defined as follows:

**TPLFE\_SUB\_SERIAL**

Serial port interface

**TPLFE\_SUB\_MODEM\_COMMON**

Common modem interface

**TPLFE\_SUB\_MODEM\_DATA**

Data modem services

**TPLFE\_SUB\_MODEM\_FAX**

Fax modem services

**TPLFE\_SUB\_VOICE**

Voice services

**TPLFE\_CAP\_MODEM\_DATA**

Capabilities of the data modem interface

**TPLFE\_CAP\_MODEM\_FAX**

Capabilities of the fax modem interface

**TPLFE\_CAP\_MODEM\_VOICE**

Capabilities of the voice modem interface

**TPLFE\_CAP\_SERIAL\_DATA**

Serial port interface for data modem services

**TPLFE\_CAP\_SERIAL\_FAX**

Serial port interface for fax modem services

**TPLFE\_CAP\_SERIAL\_VOICE**

Serial port interface for voice modem services

subfunction

This is for identifying a sub-category of services provided by a function in the CISTPL\_FUNCE tuple. The numeric value of the code is in the range of 1 to 15.

ua

This is the serial port UART identification and is defined as follows:

**TPLFE\_UA\_8250**

Intel 8250

**TPLFE\_UA\_16450**

NS 16450

**TPLFE\_UA\_16550**

NS 16550

uc

This identifies the serial port UART capabilities and is defined as follows:

**TPLFE\_UC\_PARITY\_SPACE**

Space parity supported

**TPLFE\_UC\_PARITY\_MARK**

Mark parity supported

**TPLFE\_UC\_PARITY\_ODD**

Odd parity supported

**TPLFE\_UC\_PARITY\_EVEN**

Even parity supported

**TPLFE\_UC\_CS5**

5 bit characters supported

**TPLFE\_UC\_CS6**

6 bit characters supported

**TPLFE\_UC\_CS7**

7 bit characters supported

**TPLFE\_UC\_CS8**

8 bit characters supported

**TPLFE\_UC\_STOP\_1**

1 stop bit supported

**TPLFE\_UC\_STOP\_15**

1.5 stop bits supported

fc

This identifies the modem flow control methods and is defined as follows:

**TPLFE\_FC\_TX\_XONOFF**

Transmit XON/XOFF

**TPLFE\_FC\_RX\_XONOFF**

Receiver XON/XOFF

**TPLFE\_FC\_TX\_HW**

Transmit hardware flow control (CTS)

**TPLFE\_FC\_RX\_HW**

Receiver hardware flow control (RTS)

**TPLFE\_FC\_TRANS**

Tranparent flow control

**ms**

This identifies the modem modulation standards and is defined as follows:

**TPLFE\_MS\_BELL103**

300bps

**TPLFE\_MS\_V21**

300bps (V.21)

**TPLFE\_MS\_V23**

600/1200bps (V.23)

**TPLFE\_MS\_V22AB**

1200bps (V.22A V.22B)

**TPLFE\_MS\_BELL212**

2400bsp (US Bell 212)

**TPLFE\_MS\_V22BIS**

2400bps (V.22bis)

**TPLFE\_MS\_V26**

2400bps leased line (V.26)

**TPLFE\_MS\_V26BIS**

2400bps (V.26bis)

**TPLFE\_MS\_V27BIS**

4800/2400bps leased line (V.27bis)

**TPLFE\_MS\_V29**

9600/7200/4800 leased line (V.29)

**TPLFE\_MS\_V32**

Up to 9600bps (V.32)

em	<p><b>TPLFE_MS_V32BIS</b> Up to 14400bps (V.32bis)</p> <p><b>TPLFE_MS_VFAST</b> Up to 28800 V.FAST</p> <p>This identifies modem error correction/detection protocols and is defined as follows:</p> <p><b>TPLFE_EM_MNP</b> MNP levels 2-4</p> <p><b>TPLFE_EM_V42</b> CCITT LAPM (V.42)</p>
dc	<p>This identifies modem data compression protocols and is defined as follows:</p> <p><b>TPLFE_DC_V42BI</b> CCITT compression V.42</p> <p><b>TPLFE_DC_MNP5</b> MNP compression (uses MNP 2, 3 or 4)</p>
cm	<p>This identifies modem command protocols and is defined as follows:</p> <p><b>TPLFE_CM_AT1</b> ANSI/EIA/TIA 602 "Action" commands</p> <p><b>TPLFE_CM_AT2</b> ANSI/EIA/TIA 602 "ACE/DCE IF Params"</p> <p><b>TPLFE_CM_AT3</b> ANSI/EIA/TIA 602 "Ace Parameters"</p> <p><b>TPLFE_CM_MNP_AT</b> MNP specification AT commands</p> <p><b>TPLFE_CM_V25BIS</b> V.25bis calling commands</p> <p><b>TPLFE_CM_V25A</b> V.25bis test procedures</p>

ex	<p><b>TPLFE_CM_DMCL</b> DMCL command mode</p> <p>This identifies the modem escape mechanism and is defined as follows:</p> <p><b>TPLFE_EX_BREAK</b> BREAK support standardized</p> <p><b>TPLFE_EX_PLUS</b> +++ returns to command mode</p> <p><b>TPLFE_EX_UD</b> User defined escape character</p>
dy	<p>This identifies modem standardized data encryption and is a reserved field for future use and must be set to 0.</p>
ef	<p>This identifies modem miscellaneous features and is defined as follows:</p> <p><b>TPLFE_EF_CALLERID</b> Caller ID is supported</p>
fm	<p>This identifies fax modulation standards and is defined as follows:</p> <p><b>TPLFE_FM_V21C2</b> 300bps (V.21-C2)</p> <p><b>TPLFE_FM_V27TER</b> 4800/2400bps (V.27ter)</p> <p><b>TPLFE_FM_V29</b> 9600/7200/4800 leased line (V.29)</p> <p><b>TPLFE_FM_V17</b> 14.4K/12K/9600/7200bps (V.17)</p> <p><b>TPLFE_FM_V33</b> 4.4K/12K/9600/7200 leased line (V.33)</p>
fs	<p>This identifies the fax feature selection and is defined as follows:</p>

**TPLFE\_FS\_T3**

Group 2 (T.3) service class

**TPLFE\_FS\_T4**

Group 3 (T.4) service class

**TPLFE\_FS\_T6**

Group 4 (T.6) service class

**TPLFE\_FS\_ECM**

Error Correction Mode

**TPLFE\_FS\_VOICEREQ**

Voice requests allowed

**TPLFE\_FS\_POLLING**

Polling support

**TPLFE\_FS\_FTP**

File transfer support

**TPLFE\_FS\_PASSWORD**

tech

This file is a password technology type and is defined as follows:

password support

**TPLFE\_LAN\_TECH\_ARCNET**

Arcnet

**TPLFE\_LAN\_TECH\_ETHERNET**

Ethernet

**TPLFE\_LAN\_TECH\_TOKENRING**

Token Ring

**TPLFE\_LAN\_TECH\_LOCALTALK**

Local Talk

**TPLFE\_LAN\_TECH\_FDDI**

FDDI/CDDI

**TPLFE\_LAN\_TECH\_ATM**

ATM

media

**TPLFE\_LAN\_TECH\_WIRELESS**

Wireless

This identifies the LAN media type and is defined as follows:

**TPLFE\_LAN\_MEDIA\_INHERENT**

Generic interface

**TPLFE\_LAN\_MEDIA\_UTP**

Unshielded twisted pair

**TPLFE\_LAN\_MEDIA\_STP**

Shielded twisted pair

**TPLFE\_LAN\_MEDIA\_THIN\_COAX**

Thin coax

**TPLFE\_LAN\_MEDIA\_THICK\_COAX**

Thick coax

**TPLFE\_LAN\_MEDIA\_FIBER**

Fiber

**TPLFE\_LAN\_MEDIA\_SSR\_902**

Spread spectrum radio 902-928 MHz

**TPLFE\_LAN\_MEDIA\_SSR\_2\_4**

Spread spectrum radio 2.4 GHz

**TPLFE\_LAN\_MEDIA\_SSR\_5\_4**

Spread spectrum radio 5.4 GHz

**TPLFE\_LAN\_MEDIA\_DIFFUSE\_IR**

Diffuse infra red

**TPLFE\_LAN\_MEDIA\_PTP\_IR**

Point to point infra red

**RETURN VALUES**

CS\_SUCCESS

Successful operation.

CS\_BAD\_HANDLE

Client handle is invalid.

CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_Parse_CISTPL_FUNCID(9F)`, `csx_RegisterClient(9F)`,  
`csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_FUNCID – parse Function Identification tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_FUNCID(client_handle_t ch, tuple_t *tu, cistpl_funcid_t *cf);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code>.</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure (see <code>tuple(9S)</code>) returned by a call to <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>.</p> <p><b>cf</b> Pointer to a <code>cistpl_funcid_t</code> structure which contains the parsed CISTPL_FUNCID tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Function Identification tuple, CISTPL_FUNCID, into a form usable by PC Card drivers.</p> <p>The CISTPL_FUNCID tuple is used to describe information about the functionality provided by a PC Card. Information is also provided to enable system utilities to decide if the PC Card should be configured during system initialization. If additional function specific information is available, one or more function extension tuples of type CISTPL_FUNCE follow this tuple (see <code>csx_Parse_CISTPL_FUNCE(9F)</code>).</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>cistpl_funcid_t</code> are:</p> <pre>uint32_t    function;    /* PC Card function code */ uint32_t    sysinit;    /* system initialization mask */</pre> <p>The fields are defined as follows:</p> <p><code>function</code> This is the function type for CISTPL_FUNCID:</p> <p><b>TPLFUNC_MULTI</b> Vendor-specific multifunction card</p> <p><b>TPLFUNC_MEMORY</b> Memory card</p> <p><b>TPLFUNC_SERIAL</b> Serial I/O port</p>

**TPLFUNC\_PARALLEL**

Parallel printer port

**TPLFUNC\_FIXED**

Fixed disk, silicon or removable

**TPLFUNC\_VIDEO**

Video interface

**TPLFUNC\_LAN**

Local Area Network adapter

**TPLFUNC\_AIMS**

Auto Incrementing Mass Storage

**TPLFUNC\_SCSI**

SCSI bridge

**TPLFUNC\_SECURITY**

Security cards

**TPLFUNC\_VENDOR\_SPECIFIC**

Vendor specific

**TPLFUNC\_UNKNOWN**

Unknown function(s)

sysinit This field is bit-mapped and defined as follows:

**TPLINIT\_POST**

POST should attempt configure

**TPLINIT\_ROM**

Map ROM during sys init

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.

CS\_NO\_CIS

No Card Information Structure (CIS) on PC Card.

CS\_UNSUPPORTED\_FUNCTION

No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

`csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_Parse_CISTPL_FUNCID(9F)`, `csx_RegisterClient(9F)`,  
`csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_GEOMETRY – parse the Geometry tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_GEOMETRY(client_handle_t ch, tuple_t *tu, cistpl_geometry_t *pt);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>pt</b> Pointer to a <b>cistpl_geometry_t</b> structure which contains the parsed CISTPL_GEOMETRY tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Geometry tuple, CISTPL_GEOMETRY, into a form usable by PC Card drivers.</p> <p>The CISTPL_GEOMETRY tuple indicates the geometry of a disk-like device.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_geometry_t</b> are:</p> <pre>uint32_t    spt; uint32_t    tpc; uint32_t    ncyl;</pre> <p>The fields are defined as follows:</p> <p><b>spt</b> This field indicates the number of sectors per track.</p> <p><b>tpc</b> This field indicates the number of tracks per cylinder.</p> <p><b>ncyl</b> This field indicates the number of cylinders.</p>
<b>RETURN VALUES</b>	<p><b>CS_SUCCESS</b> Successful operation.</p> <p><b>CS_BAD_HANDLE</b> Client handle is invalid.</p> <p><b>CS_UNKNOWN_TUPLE</b> Parser does not know how to parse tuple.</p> <p><b>CS_NO_CARD</b> No PC Card in socket.</p>

CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_JEDEC_C, csx_Parse_CISTPL_JEDEC_A – parse JEDEC Identifier tuples
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_JEDEC_C(client_handle_t ch, tuple_t * tu, cistpl_jedec_t * cj);  int32_t csx_Parse_CISTPL_JEDEC_A(client_handle_t ch, tuple_t * tu, cistpl_jedec_t * cj);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI )
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b> )returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b> .</p> <p><b>cj</b> Pointer to a <b>cistpl_jedec_t</b> structure which contains the parsed CISTPL_JEDEC_C or CISTPL_JEDEC_A tuple information upon return from these functions, respectively.</p>
<b>DESCRIPTION</b>	<p><b>csx_Parse_CISTPL_JEDEC_C()</b> and <b>csx_Parse_CISTPL_JEDEC_A()</b> parse the JEDEC Identifier tuples, CISTPL_JEDEC_C and CISTPL_JEDEC_A, respectively, into a form usable by PC Card drivers.</p> <p>The CISTPL_JEDEC_C and CISTPL_JEDEC_A tuples are optional tuples provided for cards containing programmable devices. They describe information for Common Memory or Attribute Memory space, respectively.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_jedec_t</b> are:</p> <pre>uint32_t      nid; /* # of JEDEC identifiers present */ jedec_ident_t jid[CISTPL_JEDEC_MAX_IDENTIFIERS];</pre> <p>The structure members of <b>jedec_ident_t</b> are:</p> <pre>uint32_t      id; /* manufacturer id */ uint32_t      info; /* manufacturer specific info */</pre>
<b>RETURN VALUES</b>	<p>CS_SUCCESS Successful operation.</p> <p>CS_BAD_HANDLE Client handle is invalid.</p> <p>CS_UNKNOWN_TUPLE Parser does not know how to parse tuple.</p>

CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure ( CIS )on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

These functions may be called from user or kernel context.

**SEE ALSO**

`csx_GetFirstTuple(9F)` , `csx_GetTupleData(9F)` ,  
`csx_Parse_CISTPL_DEVICE(9F)` , `csx_RegisterClient(9F)` ,  
`csx_ValidateCIS(9F)` , `tuple(9S)`

*PC Card 95 Standard* , PCMCIA/JEIDA

<b>NAME</b>	csx_Parse_CISTPL_LINKTARGET – parse the Link Target tuple								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_LINKTARGET(client_handle_t ch, tuple_t *tu, cistpl_linktarget_t *pt);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)								
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>pt</b> Pointer to a <b>cistpl_linktarget_t</b> structure which contains the parsed CISTPL_LINKTARGET tuple information upon return from this function.</p>								
<b>DESCRIPTION</b>	<p>This function parses the Link Target tuple, CISTPL_LINKTARGET, into a form usable by PCCard drivers.</p> <p>The CISTPL_LINKTARGET tuple is used to verify that tuple chains other than the primary chain are valid. All secondary tuple chains are required to contain this tuple as the first tuple of the chain.</p>								
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_linktarget_t</b> are:</p> <pre>uint32_t    length; char        tpltg_tag[CIS_MAX_TUPLE_DATA_LEN];</pre> <p>The fields are defined as follows:</p> <p><b>length</b> This field indicates the number of bytes in <b>tpltg_tag</b>.</p> <p><b>tpltg_tag</b> This field provides the Link Target tuple information.</p>								
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>CS_SUCCESS</td> <td>Successful operation.</td> </tr> <tr> <td>CS_BAD_HANDLE</td> <td>Client handle is invalid.</td> </tr> <tr> <td>CS_UNKNOWN_TUPLE</td> <td>Parser does not know how to parse tuple.</td> </tr> <tr> <td>CS_NO_CARD</td> <td>No PC Card in socket.</td> </tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_HANDLE	Client handle is invalid.	CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.	CS_NO_CARD	No PC Card in socket.
CS_SUCCESS	Successful operation.								
CS_BAD_HANDLE	Client handle is invalid.								
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.								
CS_NO_CARD	No PC Card in socket.								

CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`  
*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_LONGLINK_A, csx_Parse_CISTPL_LONGLINK_C – parse the Long Link A and C tuples
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_LONGLINK_A(client_handle_t ch, tuple_t * tu, cistpl_longlink_ac_t * pt);  int32_t csx_Parse_CISTPL_LONGLINK_C(client_handle_t ch, tuple_t * tu, cistpl_longlink_ac_t * pt);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b> )returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b> .</p> <p><b>pt</b> Pointer to a <b>cistpl_longlink_ac_t</b> structure which contains the parsed CISTPL_LONGLINK_A or CISTPL_LONGLINK_C tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Long Link A and C tuples, CISTPL_LONGLINK_A and CISTPL_LONGLINK_A, into a form usable by PC Card drivers.</p> <p>The CISTPL_LONGLINK_A and CISTPL_LONGLINK_C tuples provide links to Attribute and Common Memory.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_longlink_ac_t</b> are:</p> <pre>uint32_t flags; uint32_t tp11_addr;</pre> <p>The fields are defined as follows:</p> <p><b>flags</b> This field indicates the type of memory:</p> <p style="padding-left: 40px;">CISTPL_LONGLINK_AC_AM long link to Attribute Memory</p> <p style="padding-left: 40px;">CISTPL_LONGLINK_AC_CM long link to Common Memory</p> <p><b>tp11_addr</b> This field provides the offset from the beginning of the specified address space.</p>

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

**csx\_GetFirstTuple(9F)** , **csx\_GetTupleData(9F)** ,  
**csx\_RegisterClient(9F)** , **csx\_ValidateCIS(9F)** , **tuple(9S)**

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_LONGLINK_MFC - parse the Multi-Function tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_LONGLINK_MFC(client_handle_t ch, tuple_t *tu, cistpl_longlink_mfc_t *pt);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>pt</b> Pointer to a <b>cistpl_longlink_mfc_t</b> structure which contains the parsed CISTPL_LONGLINK_MFC tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Multi-Function tuple, CISTPL_LONGLINK_MFC, into a form usable by PC Card drivers.</p> <p>The CISTPL_LONGLINK_MFC tuple describes the start of the function-specific CIS for each function on a multi-function card.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_longlink_mfc_t</b> are:</p> <pre>uint32_t    nfuncs; uint32_t    nregs; uint32_t    function[CIS_MAX_FUNCTIONS].tas uint32_t    function[CIS_MAX_FUNCTIONS].addr</pre> <p>The fields are defined as follows:</p> <p><b>nfuncs</b></p> <p>This field indicates the number of functions on the PC card.</p> <p><b>nregs</b></p> <p>This field indicates the number of configuration register sets.</p> <pre>function[CIS_MAX_FUNCTIONS].tas</pre> <p>This field provides the target address space for each function on the PC card. This field can be one of:</p>

**RETURN VALUES**

CISTPL\_LONGLINK\_MFC\_TAS\_AM

CIS in attribute memory

CISTPL\_LONGLINK\_MFC\_TAS\_CM

CIS in common memory

function[CIS\_MAX\_FUNCTIONS].addr

This field provides the target address offset for each function on the PC card.

CS\_SUCCESS

Successful operation.

CS\_BAD\_HANDLE

Client handle is invalid.

CS\_UNKNOWN\_TUPLE

Parser does not know how to parse tuple.

CS\_NO\_CARD

No PC Card in socket.

CS\_NO\_CIS

No Card Information Structure (CIS) on PC Card.

CS\_UNSUPPORTED\_FUNCTION

No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO****csx\_GetFirstTuple(9F)**, **csx\_GetTupleData(9F)**,  
**csx\_RegisterClient(9F)**, **csx\_ValidateCIS(9F)**, **tuple(9S)***PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_Parse_CISTPL_MANFID – parse Manufacturer Identification tuple												
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_MANFID(client_handle_t ch, tuple_t *tu, cistpl_manfid_t *cm);</pre>												
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)												
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>cm</b> Pointer to a <b>cistpl_manfid_t</b> structure which contains the parsed CISTPL_MANFID tuple information upon return from this function.</p>												
<b>DESCRIPTION</b>	<p>This function parses the Manufacturer Identification tuple, CISTPL_MANFID, into a form usable by PC Card drivers.</p> <p>The CISTPL_MANFID tuple is used to describe the information about the manufacturer of a PC Card. There are two types of information, the PC Card's manufacturer and a manufacturer card number.</p>												
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_manfid_t</b> are:</p> <pre>uint32_t  manf; /* PCMCIA assigned manufacturer code */ uint32_t  card; /* manufacturer information                 (part number and/or revision) */</pre>												
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>CS_SUCCESS</td> <td>Successful operation.</td> </tr> <tr> <td>CS_BAD_HANDLE</td> <td>Client handle is invalid.</td> </tr> <tr> <td>CS_UNKNOWN_TUPLE</td> <td>Parser does not know how to parse tuple.</td> </tr> <tr> <td>CS_NO_CARD</td> <td>No PC Card in socket.</td> </tr> <tr> <td>CS_NO_CIS</td> <td>No Card Information Structure (CIS) on PC card.</td> </tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td> <td>No PCMCIA hardware installed.</td> </tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_HANDLE	Client handle is invalid.	CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.	CS_NO_CARD	No PC Card in socket.	CS_NO_CIS	No Card Information Structure (CIS) on PC card.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.												
CS_BAD_HANDLE	Client handle is invalid.												
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.												
CS_NO_CARD	No PC Card in socket.												
CS_NO_CIS	No Card Information Structure (CIS) on PC card.												
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.												

**CONTEXT** | This function may be called from user or kernel context.

**SEE ALSO** | `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_ORG – parse the Data Organization tuple								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_ORG(client_handle_t ch, tuple_t *tu, cistpl_org_t *pt);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)								
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>pt</b> Pointer to a <b>cistpl_org_t</b> structure which contains the parsed CISTPL_ORG tuple information upon return from this function.</p>								
<b>DESCRIPTION</b>	<p>This function parses the Data Organization tuple, CISTPL_ORG, into a form usable by PC Card drivers.</p> <p>The CISTPL_ORG tuple provides a text description of the organization.</p>								
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_org_t</b> are:</p> <pre>uint32_t  type; char      desc[CIS_MAX_TUPLE_DATA_LEN];</pre> <p>The fields are defined as follows:</p> <p><b>type</b></p> <p>This field indicates type of data organization.</p> <p><b>desc[CIS_MAX_TUPLE_DATA_LEN]</b></p> <p>This field provides the text description of this organization.</p>								
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>CS_SUCCESS</td> <td>Successful operation.</td> </tr> <tr> <td>CS_BAD_HANDLE</td> <td>Client handle is invalid.</td> </tr> <tr> <td>CS_UNKNOWN_TUPLE</td> <td>Parser does not know how to parse tuple.</td> </tr> <tr> <td>CS_NO_CARD</td> <td>No PC Card in socket.</td> </tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_HANDLE	Client handle is invalid.	CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.	CS_NO_CARD	No PC Card in socket.
CS_SUCCESS	Successful operation.								
CS_BAD_HANDLE	Client handle is invalid.								
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.								
CS_NO_CARD	No PC Card in socket.								

CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_SPCL – parse the Special Purpose tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_SPCL(client_handle_t ch, tuple_t *tu, cistpl_spcl_t *csp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code>.</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure (see <code>tuple(9S)</code>) returned by a call to <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>.</p> <p><b>csp</b> Pointer to a <code>cistpl_spcl_t</code> structure which contains the parsed CISTPL_SPCL tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Special Purpose tuple, CISTPL_SPCL, into a form usable by PC Card drivers.</p> <p>The CISTPL_SPCL tuple is identified by an identification field that is assigned by PCMCIA or JEIDA. A sequence field allows a series of CISTPL_SPCL tuples to be used when the data exceeds the size that can be stored in a single tuple; the maximum data area of a series of CISTPL_SPCL tuples is unlimited. Another field gives the number of bytes in the data field in this tuple.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>cistpl_data_t</code> are:</p> <pre>uint32_t id; /* tuple contents identification */ uint32_t seq; /* data sequence number */ uint32_t bytes; /* number of bytes following */ uchar_t data[CIS_MAX_TUPLE_DATA_LEN];</pre> <p>The fields are defined as follows:</p> <p><b>id</b> This field contains a PCMCIA or JEIDA assigned value that identifies this series of one or more CISTPL_SPCL tuples. These field values are assigned by contacting either PCMCIA or JEIDA.</p> <p><b>seq</b> This field contains a data sequence number. CISTPL_SPCL_SEQ_END is the last tuple in sequence.</p> <p><b>bytes</b> This field contains the number of data bytes in the <code>data[CIS_MAX_TUPLE_DATA_LEN]</code>.</p> <p><b>data</b> The data component of this tuple.</p>

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

**csx\_GetFirstTuple(9F)**, **csx\_GetTupleData(9F)**,  
**csx\_RegisterClient(9F)**, **csx\_ValidateCIS(9F)**, **tuple(9S)**

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_SWIL – parse the Software Interleaving tuple												
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_SWIL(client_handle_t ch, tuple_t *tu, cistpl_swil_t *pt);</pre>												
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)												
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <code>csx_RegisterClient(9F)</code>.</p> <p><b>tu</b> Pointer to a <code>tuple_t</code> structure (see <code>tuple(9S)</code>) returned by a call to <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code>.</p> <p><b>pt</b> Pointer to a <code>cistpl_swil_t</code> structure which contains the parsed CISTPL_SWIL tuple information upon return from this function.</p>												
<b>DESCRIPTION</b>	<p>This function parses the Software Interleaving tuple, CISTPL_SWIL, into a form usable by PC Card drivers.</p> <p>The CISTPL_SWIL tuple provides the software interleaving of data within a partition on the card.</p>												
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <code>cistpl_swil_t</code> are:</p> <pre>uint32_t    intrlv;</pre> <p>The fields are defined as follows:</p> <pre>intrlv      This field provides the software interleaving for a partition.</pre>												
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>CS_SUCCESS</td> <td>Successful operation.</td> </tr> <tr> <td>CS_BAD_HANDLE</td> <td>Client handle is invalid.</td> </tr> <tr> <td>CS_UNKNOWN_TUPLE</td> <td>Parser does not know how to parse tuple.</td> </tr> <tr> <td>CS_NO_CARD</td> <td>No PC Card in socket.</td> </tr> <tr> <td>CS_NO_CIS</td> <td>No Card Information Structure (CIS) on PC Card.</td> </tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td> <td>No PCMCIA hardware installed.</td> </tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_HANDLE	Client handle is invalid.	CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.	CS_NO_CARD	No PC Card in socket.	CS_NO_CIS	No Card Information Structure (CIS) on PC Card.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.												
CS_BAD_HANDLE	Client handle is invalid.												
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.												
CS_NO_CARD	No PC Card in socket.												
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.												
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.												

**CONTEXT** | This function may be called from user or kernel context.

**SEE ALSO** | `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_VERS_1 – parse Level-1 Version/Product Information tuple	
<b>SYNOPSIS</b>	#include <sys/pccard.h>	
	<pre>int32_t csx_Parse_CISTPL_VERS_1(client_handle_t ch, tuple_t *tu, cistpl_vers_1_t *cv1);</pre>	
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)	
<b>PARAMETERS</b>	<b>ch</b>	Client handle returned from <b>csx_RegisterClient(9F)</b> .
	<b>tu</b>	Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b> ) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b> .
	<b>cv1</b>	Pointer to a <b>cistpl_vers_1_t</b> structure which contains the parsed CISTPL_VERS_1 tuple information upon return from this function.
<b>DESCRIPTION</b>	<p>This function parses the Level-1 Version/Product Information tuple, CISTPL_VERS_1, into a form usable by PC Card drivers.</p> <p>The CISTPL_VERS_1 tuple is used to describe the card Level-1 version compliance and card manufacturer information.</p>	
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_vers_1_t</b> are:</p> <pre>uint32_t major; /* major version number */ uint32_t minor; /* minor version number */ uint32_t ns; /* number of information strings */ char pi[CISTPL_VERS_1_MAX_PROD_STRINGS] [CIS_MAX_TUPLE_DATA_LEN]; /* pointers to product information strings */</pre>	
<b>RETURN VALUES</b>	<b>CS_SUCCESS</b>	Successful operation.
	<b>CS_BAD_HANDLE</b>	Client handle is invalid.
	<b>CS_UNKNOWN_TUPLE</b>	Parser does not know how to parse tuple.
	<b>CS_NO_CARD</b>	No PC Card in socket.
	<b>CS_NO_CIS</b>	No Card Information Structure (CIS) on PC Card.
	<b>CS_UNSUPPORTED_FUNCTION</b>	No PCMCIA hardware installed.

**CONTEXT** | This function may be called from user or kernel context.

**SEE ALSO** | `csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Parse_CISTPL_VERS_2 – parse Level-2 Version and Information tuple
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_Parse_CISTPL_VERS_2(client_handle_t ch, tuple_t *tu, cistpl_vers_2_t *cv2);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>cv2</b> Pointer to a <b>cistpl_vers_2_t</b> structure which contains the parsed CISTPL_VERS_2 tuple information upon return from this function.</p>
<b>DESCRIPTION</b>	<p>This function parses the Level-2 Version and Information tuple, CISTPL_VERS_2, into a form usable by PC Card drivers.</p> <p>The CISTPL_VERS_2 tuple is used to describe the card Level-2 information which has the logical organization of the card's data.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cistpl_vers_2_t</b> are:</p> <pre>uint32_t vers; /* version number */ uint32_t comply; /* level of compliance */ uint32_t dindex; /* byte address of first data byte in card */ uint32_t vspec8; /* vendor specific (byte 8) */ uint32_t vspec9; /* vendor specific (byte 9) */ uint32_t nhdr; /* number of copies of CIS present on device */ char oem[CIS_MAX_TUPLE_DATA_LEN]; /* Vendor of software that formatted card */ char info[CIS_MAX_TUPLE_DATA_LEN]; /* Informational message about card */</pre>
<b>RETURN VALUES</b>	<p><b>CS_SUCCESS</b> Successful operation.</p> <p><b>CS_BAD_HANDLE</b> Client handle is invalid.</p> <p><b>CS_UNKNOWN_TUPLE</b> Parser does not know how to parse tuple.</p> <p><b>CS_NO_CARD</b> No PC Card in socket.</p> <p><b>CS_NO_CIS</b> No Card Information Structure (CIS) on PC Card.</p>

CS\_UNSUPPORTED\_FUNCTION            No PCMCIA hardware installed.

**CONTEXT**            This function may be called from user or kernel context.

**SEE ALSO**            *csx\_GetFirstTuple*(9F), *csx\_GetTupleData*(9F),  
*csx\_RegisterClient*(9F), *csx\_ValidateCIS*(9F), *tuple*(9S)

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_ParseTuple – generic tuple parser
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_ParseTuple(client_handle_t ch, tuple_t *tu, cisparsed_t *cp, cisdata_t cd);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>tu</b> Pointer to a <b>tuple_t</b> structure (see <b>tuple(9S)</b>) returned by a call to <b>csx_GetFirstTuple(9F)</b> or <b>csx_GetNextTuple(9F)</b>.</p> <p><b>cp</b> Pointer to a <b>cisparsed_t</b> structure that unifies all tuple parsing structures.</p> <p><b>cd</b> Extended tuple data for some tuples.</p>
<b>DESCRIPTION</b>	This function is the generic tuple parser entry point.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cisparsed_t</b> are:</p> <pre>typedef union cisparsed_t {     cistpl_config_t      cistpl_config;     cistpl_device_t     cistpl_device;     cistpl_vers_1_t     cistpl_vers_1;     cistpl_vers_2_t     cistpl_vers_2;     cistpl_jedec_t      cistpl_jedec;     cistpl_format_t     cistpl_format;     cistpl_geometry_t   cistpl_geometry;     cistpl_byteorder_t  cistpl_byteorder;     cistpl_date_t       cistpl_date;     cistpl_battery_t    cistpl_battery;     cistpl_org_t        cistpl_org;     cistpl_manfid_t     cistpl_manfid;     cistpl_funcid_t     cistpl_funcid;     cistpl_funcc_t      cistpl_funcc;     cistpl_cftable_entry_t cistpl_cftable_entry;     cistpl_linktarget_t cistpl_linktarget;     cistpl_longlink_ac_t cistpl_longlink_ac;     cistpl_longlink_mfc_t cistpl_longlink_mfc;     cistpl_spcl_t       cistpl_spcl;     cistpl_swil_t       cistpl_swil;     cistpl_bar_t        cistpl_bar;     cistpl_devicegeo_t  cistpl_devicegeo;     cistpl_longlink_cb_t cistpl_longlink_cb;     cistpl_get_tuple_name_t cistpl_get_tuple_name; } cisparsed_t;</pre>

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_UNKNOWN_TUPLE	Parser does not know how to parse tuple.
CS_NO_CARD	No PC Card in socket.
CS_BAD_CIS	Generic parser error.
CS_NO_CIS	No Card Information Structure (CIS) on PC Card.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

`csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`,  
`csx_Parse_CISTPL_BATTERY(9F)`, `csx_Parse_CISTPL_BYTEORDER(9F)`,  
`csx_Parse_CISTPL_CFTABLE_ENTRY(9F)`,  
`csx_Parse_CISTPL_CONFIG(9F)`, `csx_Parse_CISTPL_DATE(9F)`,  
`csx_Parse_CISTPL_DEVICE(9F)`, `csx_Parse_CISTPL_FUNCE(9F)`,  
`csx_Parse_CISTPL_FUNCID(9F)`, `csx_Parse_CISTPL_JEDEC_C(9F)`,  
`csx_Parse_CISTPL_MANFID(9F)`, `csx_Parse_CISTPL_SPCL(9F)`,  
`csx_Parse_CISTPL_VERS_1(9F)`, `csx_Parse_CISTPL_VERS_2(9F)`,  
`csx_RegisterClient(9F)`, `csx_ValidateCIS(9F)`, `tuple(9S)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_Put8, csx_Put16, csx_Put32, csx_Put64 – write to device register
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  void csx_Put8(acc_handle_t handle, uint32_t offset, uint8_t value); void csx_Put16(acc_handle_t handle, uint32_t offset, uint16_t value); void csx_Put32(acc_handle_t handle, uint32_t offset, uint32_t value); void csx_Put64(acc_handle_t handle, uint32_t offset, uint64_t value);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>handle</b>            The access handle returned from <b>csx_RequestIO(9F)</b> , <b>csx_RequestWindow(9F)</b> , or <b>csx_DupHandle(9F)</b> .</p> <p><b>offset</b>            The offset in bytes from the base of the mapped resource.</p> <p><b>value</b>             The data to be written to the device.</p>
<b>DESCRIPTION</b>	<p>These functions generate a write of various sizes to the mapped memory or device register.</p> <p>The <b>csx_Put8()</b> , <b>csx_Put16()</b> , <b>csx_Put32()</b> , and <b>csx_Put64()</b> functions write 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively, to the device address represented by the handle, <i>handle</i> , at an offset in bytes represented by the offset, <i>offset</i> .</p> <p>Data that consists of more than one byte will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions may be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<p><b>csx_DupHandle(9F)</b> , <b>csx_Get8(9F)</b> , <b>csx_GetMappedAddr(9F)</b> , <b>csx_RepGet8(9F)</b> , <b>csx_RepPut8(9F)</b> , <b>csx_RequestIO(9F)</b> , <b>csx_RequestWindow(9F)</b></p> <p><i>PC Card 95 Standard</i> , PCMCIA/JEIDA</p>

<b>NAME</b>	csx_RegisterClient – register a client
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_RegisterClient(client_handle_t *ch, client_reg_t *cr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b>     Pointer to a client_handle_t structure.</p> <p><b>mc</b>     Pointer to a client_reg_t structure.</p>
<b>DESCRIPTION</b>	This function registers a client with Card Services and returns a unique client handle for the client. The client handle must be passed to <b>csx_DeregisterClient(9F)</b> when the client terminates.
<b>STRUCTURE MEMBERS</b>	<p>The structure members of client_reg_t are:</p> <pre>uint32_t                    Attributes; uint32_t                    EventMask; event_callback_args_t      event_callback_args; uint32_t                    Version;                    /* CS version to expect */ csfunction_t                *event_handler; ddi_iblock_cookie_t        *iblk_cookie;               /* event iblk cookie */ ddi_idevice_cookie_t       *idev_cookie;               /* event idev cookie */ dev_info_t                  *dip;                       /* client's dip */ char                         driver_name[MODMAXNAMELEN];</pre> <p>The fields are defined as follows:</p> <p>Attributes</p> <p>This field is bit-mapped and defined as follows:</p> <pre>INFO_MEM_CLIENT Memory client device driver.</pre> <pre>INFO_MTD_CLIENT Memory Technology Driver client.</pre> <pre>INFO_IO_CLIENT IO client device driver.</pre> <pre>INFO_CARD_SHARE Generate artificial CS_EVENT_CARD_INSERTION and CS_EVENT_REGISTRATION_COMPLETE events.</pre>

INFO\_CARD\_EXCL

Generate artificial CS\_EVENT\_CARD\_INSERTION and CS\_EVENT\_REGISTRATION\_COMPLETE events.

INFO\_MEM\_CLIENT

INFO\_MTD\_CLIENT

INFO\_IO\_CLIENT

These bits are mutually exclusive (that is, only one bit may be set), but one of the bits must be set.

INFO\_CARD\_SHARE

INFO\_CARD\_EXCL

If either of these bits is set, the client will receive a CS\_EVENT\_REGISTRATION\_COMPLETE event when Card Services has completed its internal client registration processing and after a successful call to `csx_RequestSocketMask(9F)`.

Also, if either of these bits is set, and if a card of the type that the client can control is currently inserted in the socket (and after a successful call to `csx_RequestSocketMask(9F)`), the client will receive an artificial CS\_EVENT\_CARD\_INSERTION event. See `csx_event_handler(9E)` for valid event definitions and for additional information about handling events.

event\_callback\_args

The event\_callback\_args\_t structure members are:

```
void *client_data;
```

The client\_data field may be used to provide data available to the event handler (see `csx_event_handler(9E)`). Typically, this is the client driver's soft state pointer.

Version

This field contains the specific Card Services version number that the client expects to use. Typically, the client will use the CS\_VERSION macro to specify to Card Services which version of Card Services the client expects.

event\_handler

The client event callback handler entry point is passed in the `event_handler` field.

`iblk_cookie`

`idev_cookie`

These fields must be used by the client to set up mutexes that are used in the client's event callback handler when handling high priority events.

`dip`

The client must set this field with a pointer to the client's dip.

`driver_name`

The client must copy a driver-unique name into this member. This name must be identical across all instances of the driver.

## RETURN VALUES

`CS_SUCCESS`

Successful operation.

`CS_BAD_ATTRIBUTE`

No client type or more than one client type specified.

`CS_OUT_OF_RESOURCE`

Card Services is unable to register client.

`CS_BAD_VERSION`

Card Services version is incompatible with client.

`CS_BAD_HANDLE`

Client has already registered for this socket.

`CS_UNSUPPORTED_FUNCTION`

No PCMCIA hardware installed.

## CONTEXT

This function may be called from user or kernel context.

**SEE ALSO**

`csx_DeregisterClient(9F)`, `csx_RequestSocketMask(9F)`

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_ReleaseConfiguration – release PC Card and socket configuration
<b>SYNOPSIS</b>	#include <sys/pccard.h>  int32_t csx_ReleaseConfiguration(client_handle_t ch, release_config_t *rc);
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b>ch</b> Client handle returned from csx_RegisterClient(9F). <b>rc</b> Pointer to a release_config_t structure.
<b>DESCRIPTION</b>	This function returns a PC Card and socket to a simple memory only interface and sets the card to configuration zero by writing a 0 to the PC card's COR (Configuration Option Register).  Card Services may remove power from the socket if no clients have indicated their usage of the socket by an active csx_RequestConfiguration(9F) or csx_RequestWindow(9F).  Card Services is prohibited from resetting the PC Card and is not required to cycle power through zero (0) volts.  After calling csx_ReleaseConfiguration() any resources requested via the request functions csx_RequestIO(9F), csx_RequestIRQ(9F), or csx_RequestWindow(9F) that are no longer needed should be returned to Card Services via the corresponding csx_ReleaseIO(9F), csx_ReleaseIRQ(9F), or csx_ReleaseWindow(9F) functions. csx_ReleaseConfiguration() must be called to release the current card and socket configuration before releasing any resources requested by the driver via the request functions named above.
<b>STRUCTURE MEMBERS</b>	The structure members of release_config_t are: <pre>uint32_t Socket; /* socket number */</pre> The Socket field is not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.
<b>RETURN VALUES</b>	CS_SUCCESS  Successful operation.  CS_BAD_HANDLE  Client handle is invalid or csx_RequestConfiguration(9F) not done.

CS\_BAD\_SOCKET

Error getting or setting socket hardware parameters.

CS\_NO\_CARD

No PC card in socket.

CS\_UNSUPPORTED\_FUNCTION

No PCMCIA hardware installed.

**CONTEXT**

This function may be called from user or kernel context.

**SEE ALSO**

`csx_RegisterClient(9F)`, `csx_RequestConfiguration(9F)`,  
`csx_RequestIO(9F)`, `csx_RequestIRQ(9F)`, `csx_RequestWindow(9F)`

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_RepGet8, csx_RepGet16, csx_RepGet32, csx_RepGet64 – read repetitively from the device register
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  void csx_RepGet8(acc_handle_t handle, uint8_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);  void csx_RepGet16(acc_handle_t handle, uint16_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);  void csx_RepGet32(acc_handle_t handle, uint32_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);  void csx_RepGet64(acc_handle_t handle, uint64_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI )
<b>PARAMETERS</b>	<p><b>handle</b>            The access handle returned from <code>csx_RequestIO(9F)</code> , <code>csx_RequestWindow(9F)</code> , or <code>csx_DupHandle(9F)</code> .</p> <p><b>hostaddr</b>         Source host address.</p> <p><b>offset</b>            The offset in bytes from the base of the mapped resource.</p> <p><b>repcount</b>         Number of data accesses to perform.</p> <p><b>flags</b>             Device address flags.</p>
<b>DESCRIPTION</b>	<p>These functions generate multiple reads of various sizes from the mapped memory or device register.</p> <p>The <code>csx_RepGet8()</code> , <code>csx_RepGet16()</code> , <code>csx_RepGet32()</code> , and <code>csx_RepGet64()</code> functions generate <i>repcount</i> reads of 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively, from the device address represented by the handle, <i>handle</i> , at an offset in bytes represented by the offset, <i>offset</i> . The data read is stored consecutively into the buffer pointed to by the host address pointer, <i>hostaddr</i> .</p> <p>Data that consists of more than one byte will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte swapping if the host and the device have incompatible endian characteristics.</p> <p>When the <i>flags</i> argument is set to <code>CS_DEV_AUTOINCR</code> , these functions increment the device offset, <i>offset</i> , after each datum read operation. However,</p>

when the *flags* argument is set to `CS_DEV_NO_AUTOINCR`, the same device offset will be used for every datum access. For example, this flag may be useful when reading from a data register.

**CONTEXT**

These functions may be called from user, kernel, or interrupt context.

**SEE ALSO**

`csx_DupHandle(9F)`, `csx_Get8(9F)`, `csx_GetMappedAddr(9F)`,  
`csx_Put8(9F)`, `csx_RepPut8(9F)`, `csx_RequestIO(9F)`,  
`csx_RequestWindow(9F)`

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_RepPut8, csx_RepPut16, csx_RepPut32, csx_RepPut64 – write repetitively to the device register
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  void csx_RepPut8(acc_handle_t handle, uint8_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);  void csx_RepPut16(acc_handle_t handle, uint16_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);  void csx_RepPut32(acc_handle_t handle, uint32_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);  void csx_RepPut64(acc_handle_t handle, uint64_t * hostaddr, uint32_t offset, uint32_t repcount, uint32_t flags);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI )
<b>PARAMETERS</b>	<p><b>handle</b>            The access handle returned from <code>csx_RequestIO(9F)</code> , <code>csx_RequestWindow(9F)</code> , or <code>csx_DupHandle(9F)</code> .</p> <p><b>hostaddr</b>        Source host address.</p> <p><b>offset</b>            The offset in bytes from the base of the mapped resource.</p> <p><b>repcount</b>        Number of data accesses to perform.</p> <p><b>flags</b>            Device address flags.</p>
<b>DESCRIPTION</b>	<p>These functions generate multiple writes of various sizes to the mapped memory or device register.</p> <p>The <code>csx_RepPut8()</code> , <code>csx_RepPut16()</code> , <code>csx_RepPut32()</code> , and <code>csx_RepPut64()</code> functions generate <i>repcount</i> writes of 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively, to the device address represented by the handle, <i>handle</i> , at an offset in bytes represented by the offset, <i>offset</i> . The data written is read consecutively from the buffer pointed to by the host address pointer, <i>hostaddr</i> .</p> <p>Data that consists of more than one byte will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte swapping if the host and the device have incompatible endian characteristics.</p> <p>When the <i>flags</i> argument is set to <code>CS_DEV_AUTOINCR</code> , these functions increment the device offset, <i>offset</i> , after each datum write operation. However,</p>

when the *flags* argument is set to `CS_DEV_NO_AUTOINCR`, the same device offset will be used for every datum access. For example, this flag may be useful when writing to a data register.

**CONTEXT**

These functions may be called from user, kernel, or interrupt context.

**SEE ALSO**

`csx_DupHandle(9F)`, `csx_Get8(9F)`, `csx_GetMappedAddr(9F)`,  
`csx_Put8(9F)`, `csx_RepGet8(9F)`, `csx_RequestIO(9F)`,  
`csx_RequestWindow(9F)`

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_RequestConfiguration – configure the PC Card and socket
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_RequestConfiguration(client_handle_t ch, config_req_t *cr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>cr</b> Pointer to a <b>config_req_t</b> structure.</p>
<b>DESCRIPTION</b>	<p>This function configures the PC Card and socket. It must be used by clients that require I/O or IRQ resources for their PC Card.</p> <p><b>csx_RequestIO(9F)</b> and <b>csx_RequestIRQ(9F)</b> must be used before calling this function to specify the I/O and IRQ requirements for the PC Card and socket if necessary. <b>csx_RequestConfiguration()</b> establishes the configuration in the socket adapter and PC Card, and it programs the Base and Limit registers of multi-function PC Cards if these registers exist. The values programmed into these registers depend on the IO requirements of this configuration.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>config_req_t</b> are:</p> <pre>uint32_t Socket;          /* socket number */ uint32_t Attributes;     /* configuration attributes */ uint32_t Vcc;            /* Vcc value */ uint32_t Vpp1;           /* Vpp1 value */ uint32_t Vpp2;           /* Vpp2 value */ uint32_t IntType;        /* socket interface type - mem or IO */ uint32_t ConfigBase;     /* offset from start of AM space */ uint32_t Status;         /* value to write to STATUS register */ uint32_t Pin;            /* value to write to PRR */ uint32_t Copy;           /* value to write to COPY register */ uint32_t ConfigIndex;    /* value to write to COR */ uint32_t Present;        /* which config registers present */ uint32_t ExtendedStatus; /* value to write to EXSTAT register */</pre> <p>The fields are defined as follows:</p> <p>Socket</p> <p>Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p> <p>Attributes</p>

This field is bit-mapped. It indicates whether the client wishes the IRQ resources to be enabled and whether Card Services should ignore the VS bits on the socket interface. The following bits are defined:

CONF\_ENABLE\_IRQ\_STEERING

Enable IRQ Steering. Set to connect the PC Card IREQ line to a system interrupt previously selected by a call to `csx_RequestIRQ(9F)`. If `CONF_ENABLE_IRQ_STEERING` is set, once `csx_RequestConfiguration()` has successfully returned, the client may start receiving IRQ callbacks at the IRQ callback handler established in the call to `csx_RequestIRQ(9F)`.

CONF\_VSOVERRIDE

Override VS pins. After card insertion and prior to the first successful `csx_RequestConfiguration()`, the voltage levels applied to the card shall be those indicated by the card's physical key and/or the VS[2:1] voltage sense pins. For Low Voltage capable host systems (hosts which are capable of VS pin decoding), if a client desires to apply a voltage not indicated by the VS pin decoding, then `CONF_VSOVERRIDE` must be set in the `Attributes` field; otherwise, `CS_BAD_VCC` shall be returned.

Vcc, Vpp1, Vpp2

These fields all represent voltages expressed in tenths of a volt. Values from zero (0) to 25.5 volts may be set. To be valid, the exact voltage must be available from the system. PC Cards indicate multiple `VCC` voltage capability in their CIS via the `CISTPL_CFTABLE_ENTRY` tuple. After card insertion, Card Services processes the CIS, and when multiple `VCC` voltage capability is indicated, Card Services will allow the client to apply `VCC` voltage levels which are contrary to the VS pin decoding without requiring the client to set `CONF_VSOVERRIDE`.

IntType

This field is bit-mapped. It indicates how the socket should be configured. The following bits are defined:

SOCKET\_INTERFACE\_MEMORY

Memory only interface.

SOCKET\_INTERFACE\_MEMORY\_AND\_IO

Memory and I/O interface.

ConfigBase

This field is the offset in bytes from the beginning of attribute memory of the configuration registers.

Present

This field identifies which of the configuration registers are present. If present, the corresponding bit is set. This field is bit-mapped as follows:

CONFIG\_OPTION\_REG\_PRESENT

Configuration Option Register (COR) present

CONFIG\_STATUS\_REG\_PRESENT

Configuration Status Register (CCSR) present

CONFIG\_PINREPL\_REG\_PRESENT

Pin Replacement Register (PRR) present

CONFIG\_COPY\_REG\_PRESENT

Socket and Copy Register (SCR) present

CONFIG\_ESR\_REG\_PRESENT

Extended Status Register (ESR) present

Status, Pin, Copy, ExtendedStatus

These fields represent the initial values that should be written to those registers if they are present, as indicated by the `Present` field.

The `Pin` field is also used to inform Card Services which pins in the PC Card's PRR (Pin Replacement Register) are valid. Only those bits which are set are considered valid. This affects how status is returned by the `csx_GetStatus(9F)` function. If a particular signal is valid in the PRR, both the *mask* (STATUS) bit and the *change* (EVENT) bit must be set in the `Pin` field. The following PRR bit definitions are provided for client use:

PRR_WP_STATUS	WRITE PROTECT mask
PRR_READY_STATUS	READY mask
PRR_BVD2_STATUS	BVD2 mask
PRR_BVD1_STATUS	BVD1 mask
PRR_WP_EVENT	WRITE PROTECT changed
PRR_READY_EVENT	READY changed

PRR\_BVD2\_EVENT            BVD2 changed

PRR\_BVD1\_EVENT            BVD1 changed

ConfigIndex

This field is the value written to the COR (Configuration Option Register) for the configuration index required by the PC Card. Only the least significant six bits of the ConfigIndex field are significant; the upper two (2) bits are ignored. The interrupt type in the COR is always set to *level* mode by Card Services.

## RETURN VALUES

CS\_SUCCESS

Successful operation.

CS\_BAD\_HANDLE

Client handle is invalid or **csx\_RequestConfiguration()** not done.

CS\_BAD\_SOCKET

Error in getting or setting socket hardware parameters.

CS\_BAD\_VCC

Requested VCC is not available on socket.

CS\_BAD\_VPP

Requested VPP is not available on socket.

CS\_NO\_CARD

No PC Card in socket.

CS\_BAD\_TYPE

I/O and memory interface not supported on socket.

CS\_CONFIGURATION\_LOCKED

**csx\_RequestConfiguration()** already done.

CS\_UNSUPPORTED\_FUNCTION

No PCMCIA hardware installed.

**CONTEXT** This function may be called from user or kernel context.

**SEE ALSO** `csx_AccessConfigurationRegister(9F)`, `csx_GetStatus(9F)`,  
`csx_RegisterClient(9F)`, `csx_ReleaseConfiguration(9F)`,  
`csx_RequestIO(9F)`, `csx_RequestIRQ(9F)`

*PC Card 95 Standard, PCMCIA/JEIDA*

<b>NAME</b>	csx_RequestIO, csx_ReleaseIO – request or release I/O resources for the client
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_RequestIO(client_handle_t ch, io_req_t * ir);  int32_t csx_ReleaseIO(client_handle_t ch, io_req_t * ir);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>ir</b> Pointer to an <b>io_req_t</b> structure.</p>
<b>DESCRIPTION</b>	<p>The functions <b>csx_RequestIO()</b> and <b>csx_ReleaseIO()</b> request or release, respectively, I/O resources for the client.</p> <p>If a client requires I/O resources, <b>csx_RequestIO()</b> must be called to request I/O resources from Card Services; then <b>csx_RequestConfiguration(9F)</b> must be used to establish the configuration. <b>csx_RequestIO()</b> can be called multiple times until a successful set of I/O resources is found. <b>csx_RequestConfiguration(9F)</b> only uses the last configuration specified.</p> <p><b>csx_RequestIO()</b> fails if it has already been called without a corresponding <b>csx_ReleaseIO()</b> .</p> <p><b>csx_ReleaseIO()</b> releases previously requested I/O resources. The Card Services window resource list is adjusted by this function. Depending on the adapter hardware, the I/O window might also be disabled.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>io_req_t</b> are:</p> <pre>uint32_t      Socket;          /* socket number*/  uint32_t      Baseport1.base;  /* IO range base port address */ acc_handle_t  Baseport1.handle; /* IO range base address or port num */ uint32_t      NumPorts1;      /* first IO range number contiguous ports */ uint32_t      Attributes1;    /* first IO range attributes */  uint32_t      Baseport2.base;  /* IO range base port address */ acc_handle_t  Baseport2.handle; /* IO range base address or port num */ uint32_t      NumPorts2;      /* second IO range number contiguous ports */ uint32_t      Attributes2;    /* second IO range attributes */  uint32_t      IOAddrLines;    /* number of IO address lines decoded */</pre> <p>The fields are defined as follows:</p> <p>Socket</p>

Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.

BasePort1.base

BasePort1.handle

BasePort2.base

BasePort2.handle

Two I/O address ranges can be requested by **csx\_RequestIO()**. Each I/O address range is specified by the `BasePort`, `NumPorts`, and `Attributes` fields. If only a single I/O range is being requested, the `NumPorts2` field must be reset to 0.

When calling **csx\_RequestIO()**, the `BasePort.base` field specifies the first port address requested. Upon successful return from **csx\_RequestIO()**, the `BasePort.handle` field contains an access handle, corresponding to the first byte of the allocated I/O window, which the client must use when accessing the PC Card's I/O space via the common access functions. A client *must not* make any assumptions as to the format of the returned `BasePort.handle` field value.

If the `BasePort.base` field is set to 0, Card Services returns an I/O resource based on the available I/O resources and the number of contiguous ports requested. When `BasePort.base` is 0, Card Services aligns the returned resource in the host system's I/O address space on a boundary that is a multiple of the number of contiguous ports requested, rounded up to the nearest power of two. For example, if a client requests two I/O ports, the resource returned will be a multiple of two. If a client requests five contiguous I/O ports, the resource returned will be a multiple of eight.

If multiple ranges are being requested, at least one of the `BasePort.base` fields must be non-zero.

NumPorts

This field is the number of contiguous ports being requested.

Attributes

This field is bit-mapped. The following bits are defined:

`IO_DATA_WIDTH_8`

I/O resource uses 8-bit data path.

`IO_DATA_WIDTH_16`

I/O resource uses 16-bit data path.

`WIN_ACC_NEVER_SWAP`

Host endian byte ordering.

`WIN_ACC_BIG_ENDIAN`

Big endian byte ordering

`WIN_ACC_LITTLE_ENDIAN`

Little endian byte ordering.

`WIN_ACC_STRICT_ORDER`

Program ordering references.

`WIN_ACC_UNORDERED_OK`

May re-order references.

`WIN_ACC_MERGING_OK`

Merge stores to consecutive locations.

`WIN_ACC_LOADCACHING_OK`

May cache load operations.

`WIN_ACC_STORECACHING_OK`

May cache store operations.

For some combinations of host system busses and adapter hardware, the width of an I/O resource can not be set via **RequestIO()** ; on those systems, the host bus cycle access type determines the I/O resource data path width on a per-cycle basis.

`WIN_ACC_BIG_ENDIAN` and `WIN_ACC_LITTLE_ENDIAN` describe the endian characteristics of the device as big endian or little endian, respectively. Even though most of the devices will have the same endian characteristics as their busses, there are examples of devices with an I/O processor that has opposite endian characteristics of the busses. When `WIN_ACC_BIG_ENDIAN` or `WIN_ACC_LITTLE_ENDIAN` is set, byte swapping will automatically be performed by the system if the host machine and the device data formats have opposite endian characteristics. The implementation may take advantage of hardware platform byte swapping capabilities.

When `WIN_ACC_NEVER_SWAP` is specified, byte swapping will not be invoked in the data access functions. The ability to specify the order in which the CPU will reference data is provided by the following `Attributes` bits. Only one of the following bits may be specified:

`WIN_ACC_STRICT_ORDER`

The data references must be issued by a CPU in program order. Strict ordering is the default behavior.

`WIN_ACC_UNORDERED_OK`

The CPU may re-order the data references. This includes all kinds of re-ordering (that is, a load followed by a store may be replaced by a store followed by a load).

`WIN_ACC_MERGING_OK`

The CPU may merge individual stores to consecutive locations. For example, the CPU may turn two consecutive byte stores into one halfword store. It may also batch individual loads. For example, the CPU may turn two consecutive byte loads into one halfword load. `IO_MERGING_OK_ACC` also implies re-ordering.

`WIN_ACC_LOADCACHING_OK`

The CPU may cache the data it fetches and reuse it until another store occurs. The default behavior is to fetch new data on every load.

`WIN_ACC_LOADCACHING_OK` also implies merging and re-ordering.

`WIN_ACC_STORECACHING_OK`

The CPU may keep the data in the cache and push it to the device (perhaps with other data) at a later time. The default behavior is to push the data right away. `WIN_ACC_STORECACHING_OK` also implies load caching, merging, and re-ordering.

**These values are advisory, not mandatory. For example, data can be ordered without being merged or cached, even though a driver requests unordered, merged and cached together. All other bits in the `Attributes` field must be set to 0 .**

`IOAddrLines`

This field is the number of I/O address lines decoded by the PC Card in the specified socket.

On some systems, multiple calls to `csx_RequestIO()` with different `BasePort` , `NumPorts` , and/or `IOAddrLines` values will have to be made to find an acceptable combination of parameters that can be used by Card Services to allocate I/O resources for the client. (See `NOTES` ).

**RETURN VALUES**

CS\_SUCCESS

**Successful operation.**

CS\_BAD\_ATTRIBUTE

**Invalid Attributes specified.**

CS\_BAD\_BASE

**BasePort value is invalid.**

CS\_BAD\_HANDLE

**Client handle is invalid.**

CS\_CONFIGURATION\_LOCKED

**csx\_RequestConfiguration(9F) has already been done.**

CS\_IN\_USE

**csx\_RequestIO() has already been done without a corresponding csx\_ReleaseIO() .**

CS\_NO\_CARD

**No PC Card in socket.**

CS\_BAD\_WINDOW

**Unable to allocate I/O resources.**

CS\_OUT\_OF\_RESOURCE

**Unable to allocate I/O resources.**

CS\_UNSUPPORTED\_FUNCTION

**No PCMCIA hardware installed.****CONTEXT**

These functions may be called from user or kernel context.

**SEE ALSO****csx\_RegisterClient(9F) , csx\_RequestConfiguration(9F)**

*PC Card 95 Standard*, PCMCIA/JEIDA

**NOTES**

It is important for clients to try to use the minimum amount of I/O resources necessary. One way to do this is for the client to parse the CIS of the PC Card and call **csx\_RequestIO()** first with any `IOAddrLines` values that are 0 or that specify a minimum number of address lines necessary to decode the I/O space on the PC Card. Also, if no convenient minimum number of address lines can be used to decode the I/O space on the PC Card, it is important to try to avoid system conflicts with well-known architectural hardware features.

<b>NAME</b>	csx_RequestIRQ, csx_ReleaseIRQ – request or release IRQ resource
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_RequestIRQ(client_handle_t ch, irq_req_t * ir);  int32_t csx_ReleaseIRQ(client_handle_t ch, irq_req_t * ir);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI )
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>ir</b> Pointer to an <b>irq_req_t</b> structure.</p>
<b>DESCRIPTION</b>	<p>The function <b>csx_RequestIRQ()</b> requests an IRQ resource and registers the client's IRQ handler with Card Services.</p> <p>If a client requires an IRQ , <b>csx_RequestIRQ()</b> must be called to request an IRQ resource as well as to register the client's IRQ handler with Card Services. The client will not receive callbacks at the IRQ callback handler until <b>csx_RequestConfiguration(9F)</b> or <b>csx_ModifyConfiguration(9F)</b> has successfully returned when either of these functions are called with the <b>CONF_ENABLE_IRQ_STEERING</b> bit set.</p> <p>The function <b>csx_ReleaseIRQ()</b> releases a previously requested IRQ resource.</p> <p>The Card Services IRQ resource list is adjusted by <b>csx_ReleaseIRQ()</b> . Depending on the adapter hardware, the host bus IRQ connection might also be disabled. Client IRQ handlers always run above lock level and so should take care to perform only Solaris operations that are appropriate for an above-lock-level IRQ handler.</p> <p><b>csx_RequestIRQ()</b> fails if it has already been called without a corresponding <b>csx_ReleaseIRQ()</b> .</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>irq_req_t</b> are:</p> <pre>uint32_t          Socket;          /* socket number */ uint32_t          Attributes;      /* IRQ attribute flags */ csfunction_t      *irq_handler;    /* IRQ handler */ caddr_t           irq_handler_arg; /* IRQ handler argument */ ddi_iblock_cookie_t *iblk_cookie; /* IRQ interrupt block cookie */ ddi_idevice_cookie_t *idev_cookie; /* IRQ interrupt device cookie */</pre> <p>The fields are defined as follows:</p> <p>Socket</p>

Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.

#### Attributes

This field is bit-mapped. It specifies details about the type of IRQ desired by the client. The following bits are defined:

IRQ\_TYPE\_EXCLUSIVE

IRQ is exclusive to this socket.

IRQ\_ISR\_ADDRESS\_PROVIDED

IRQ handler address provided.

IRQ\_TYPE\_EXCLUSIVE

This bit *must* be set. It indicates that the system IRQ is dedicated to this PC Card.

IRQ\_ISR\_ADDRESS\_PROVIDED

This bit *must* be set. It indicates that the `irq_handler` field contains the address of the client's IRQ handler.

The client IRQ callback handler entry point is passed in the `irq_handler` field.

`irq_handler_arg`

The client can use the `irq_handler_arg` field to pass client-specific data to the client IRQ callback handler.

`iblk_cookie`

`idev_cookie`

These fields must be used by the client to set up mutexes that are used in the client's IRQ callback handler.

For a specific **csx\_ReleaseIRQ()** call, the values in the `irq_req_t` structure must be the same as those returned from the previous **csx\_RequestIRQ()** call; otherwise, `CS_BAD_ARGS` is returned and no changes are made to Card Services resources or the socket and adapter hardware.

#### RETURN VALUES

`CS_SUCCESS`

Successful operation.

CS\_BAD\_ARGS

IRQ description does not match allocation.

CS\_BAD\_ATTRIBUTE

IRQ\_TYPE\_EXCLUSIVE and IRQ\_ISR\_ADDRESS\_PROVIDED not set.

CS\_BAD\_HANDLE

Client handle is invalid or **csx\_RequestConfiguration(9F)** not done.

CS\_BAD\_IRQ

Unable to allocate IRQ resources.

CS\_IN\_USE

**csx\_RequestIRQ()** already done or a previous **csx\_RequestIRQ()** has not been done for a corresponding **csx\_ReleaseIRQ()** .

CS\_CONFIGURATION\_LOCKED

**csx\_RequestConfiguration(9F)** already done or **csx\_ReleaseConfiguration(9F)** has not been done.

CS\_NO\_CARD

No PC Card in socket.

CS\_UNSUPPORTED\_FUNCTION

No PCMCIA hardware installed.

#### CONTEXT

These functions may be called from user or kernel context.

#### SEE ALSO

**csx\_ReleaseConfiguration(9F)** , **csx\_RequestConfiguration(9F)**

*PC Card Card 95 Standard* , PCMCIA/JEIDA

<b>NAME</b>	csx_RequestSocketMask, csx_ReleaseSocketMask – set or clear the client's client event mask
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_RequestSocketMask(client_handle_t ch, request_socket_mask_t * sm); int32_t csx_ReleaseSocketMask(client_handle_t ch, release_socket_mask_t * rm);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI )
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>sm</b> Pointer to a <b>request_socket_mask_t</b> structure.</p> <p><b>rm</b> Pointer to a <b>release_socket_mask_t</b> structure.</p>
<b>DESCRIPTION</b>	<p>The function <b>csx_RequestSocketMask()</b> sets the client's client event mask and enables the client to start receiving events at its event callback handler. Once this function returns successfully, the client can start receiving events at its event callback handler. Any pending events generated from the call to <b>csx_RegisterClient(9F)</b> will be delivered to the client after this call as well. This allows the client to set up the event handler mutexes before the event handler gets called.</p> <p><b>csx_RequestSocketMask()</b> must be used before calling <b>csx_GetEventMask(9F)</b> or <b>csx_SetEventMask(9F)</b> for the client event mask for this socket.</p> <p>The function <b>csx_ReleaseSocketMask()</b> clears the client's client event mask.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>request_socket_mask_t</b> are:</p> <pre>uint32_t Socket; /* socket number */ uint32_t EventMask; /* event mask to set or return */</pre> <p>The structure members of <b>release_socket_mask_t</b> are:</p> <pre>uint32_t Socket; /* socket number */</pre> <p>The fields are defined as follows:</p> <p><b>Socket</b> Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p> <p><b>EventMask</b> This field is bit-mapped. Card Services performs event notification based on this field. See</p>

**csx\_event\_handler(9E)** for valid event definitions and for additional information about handling events.

**RETURN VALUES**

CS_SUCCESS	Successful operation.
CS_BAD_HANDLE	Client handle is invalid.
CS_IN_USE	<b>csx_ReleaseSocketMask()</b> has not been done.
CS_BAD_SOCKET	<b>csx_RequestSocketMask()</b> has not been done.
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.

**CONTEXT**

These functions may be called from user or kernel context.

**SEE ALSO**

**csx\_event\_handler(9E)** , **csx\_GetEventMask(9F)** ,  
**csx\_RegisterClient(9F)** , **csx\_SetEventMask(9F)**

*PC Card 95 Standard* , PCMCIA/JEIDA

<b>NAME</b>	csx_RequestWindow, csx_ReleaseWindow – request or release window resources
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_RequestWindow(client_handle_t ch, window_handle_t * wh, win_req_t * wr);  int32_t csx_ReleaseWindow(window_handle_t wh);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>wh</b> Pointer to a <b>window_handle_t</b> structure.</p> <p><b>wr</b> Pointer to a <b>win_req_t</b> structure.</p>
<b>DESCRIPTION</b>	<p>The function <b>csx_RequestWindow()</b> requests a block of system address space be assigned to a PC Card in a socket.</p> <p>The function <b>csx_ReleaseWindow()</b> releases window resources which were obtained by a call to <b>csx_RequestWindow()</b> . No adapter or socket hardware is modified by this function.</p> <p>The <b>csx_MapMemPage(9F)</b> and <b>csx_ModifyWindow(9F)</b> functions use the window handle returned by <b>csx_RequestWindow()</b> . This window handle must be freed by calling <b>csx_ReleaseWindow()</b> when the client is done using this window.</p> <p>The PC Card Attribute or Common Memory offset for this window is set by <b>csx_MapMemPage(9F)</b> .</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>win_req_t</b> are:</p> <pre>uint32_t      Socket;                /* socket number */ uint32_t      Attributes;            /* window flags */ uint32_t      Base.base;             /* requested window base address */ acc_handle_t  Base.handle;           /* returned handle for base of window */ uint32_t      Size;                  /* window size requested/granted */ uint32_t      win_params.AccessSpeed; /* window access speed */ uint32_t      win_params.IOAddrLines; /* IO address lines decoded */ uint32_t      ReqOffset;             /* required window offset */</pre> <p>The fields are defined as follows:</p> <p>Socket</p>

Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.

#### Attributes

**This field is bit-mapped. It is defined as follows:**

WIN_MEMORY_TYPE_IO	Window points to I/O space
WIN_MEMORY_TYPE_CM	Window points to Common Memory space
WIN_MEMORY_TYPE_AM	Window points to Attribute Memory space
WIN_ENABLE	Enable window
WIN_DATA_WIDTH_8	Set window to 8-bit data path
WIN_DATA_WIDTH_16	Set window to 16-bit data path
WIN_ACC_NEVER_SWAP	Host endian byte ordering
WIN_ACC_BIG_ENDIAN	Big endian byte ordering
WIN_ACC_LITTLE_ENDIAN	Little endian byte ordering
WIN_ACC_STRICT_ORDER	Program ordering references
WIN_ACC_UNORDERED_OK	May re-order references
WIN_ACC_MERGING_OK	Merge stores to consecutive locations
WIN_ACC_LOADCACHING_OK	May cache load operations
WIN_ACC_STORECACHING_OK	May cache store operations
WIN_MEMORY_TYPE_IO	
WIN_MEMORY_TYPE_CM	
WIN_MEMORY_TYPE_AM	These bits select which type of window is being requested. One of these bits must be set.
WIN_ENABLE	The client must set this bit to enable the window.
WIN_ACC_BIG_ENDIAN	
WIN_ACC_LITTLE_ENDIAN	These bits describe the endian characteristics of the device as big endian or little endian, respectively. Even though most of the devices will have the same endian characteristics as their busses, there are examples of devices with an I/O processor that has opposite endian characteristics of the busses. When either of these bits are set, byte swapping will automatically be performed by the system if the host machine and the device data formats have opposite endian characteristics. The implementation may take advantage of hardware platform byte swapping capabilities.
WIN_ACC_NEVER_SWAP	When this is specified, byte swapping will not be invoked in the data access functions.

The ability to specify the order in which the CPU will reference data is provided by the following `Attributes` bits, only one of which may be specified:

<code>WIN_ACC_STRICT_ORDER</code>	The data references must be issued by a CPU in program order. Strict ordering is the default behavior.
<code>WIN_ACC_UNORDERED_OK</code>	The CPU may re-order the data references. This includes all kinds of re-ordering (that is, a load followed by a store may be replaced by a store followed by a load).
<code>WIN_ACC_MERGING_OK</code>	The CPU may merge individual stores to consecutive locations. For example, the CPU may turn two consecutive byte stores into one halfword store. It may also batch individual loads. For example, the CPU may turn two consecutive byte loads into one halfword load. This bit also implies re-ordering.
<code>WIN_ACC_LOADCACHING_OK</code>	The CPU may cache the data it fetches and reuse it until another store occurs. The default behavior is to fetch new data on every load. This bit also implies merging and re-ordering.
<code>WIN_ACC_STORECACHING_OK</code>	The CPU may keep the data in the cache and push it to the device (perhaps with other data) at a later time. The default behavior is to push the data right away. This bit also implies load caching, merging, and re-ordering.

These values are advisory, not mandatory. For example, data can be ordered without being merged or cached, even though a driver requests unordered, merged and cached together.

All other bits in the `Attributes` field must be set to 0.

On successful return from `csx_RequestWindow()`, `WIN_OFFSET_SIZE` is set in the `Attributes` field when the client must specify card offsets to `csx_MapMemPage(9F)` that are a multiple of the window size.

`Base.base`

This field must be set to 0 on calling `csx_RequestWindow()`.

`Base.handle`

On successful return from **csx\_RequestWindow()**, the `Base.handle` field contains an access handle corresponding to the first byte of the allocated memory window which the client must use when accessing the PC Card's memory space via the common access functions. A client must *not* make any assumptions as to the format of the returned `Base.handle` field value.

#### Size

On calling **csx\_RequestWindow()**, the `Size` field is the size in bytes of the memory window requested. `Size` may be zero to indicate that Card Services should provide the smallest sized window available. On successful return from **csx\_RequestWindow()**, the `Size` field contains the actual size of the window allocated.

#### `win_params.AccessSpeed`

This field specifies the access speed of the window if the client is requesting a memory window. The `AccessSpeed` field bit definitions use the format of the extended speed byte of the Device ID tuple. If the mantissa is 0 (noted as reserved in the *PC Card 95 Standard*), the lower bits are a binary code representing a speed from the following table:

Code	Speed
0	(Reserved - do not use).
1	250 nsec
2	200 nsec
3	150 nsec
4	100 nse
5-7	(Reserved—do not use.)

To request a window that supports the WAIT signal, OR-in the `WIN_USE_WAIT` bit to the `AccessSpeed` value before calling this function.

It is recommended that clients use the **csx\_ConvertSpeed(9F)** function to generate the appropriate `AccessSpeed` values rather than manually perturbing the `AccessSpeed` field.

#### `win_params.IOAddrLines`

If the client is requesting an I/O window, the `IOAddrLines` field is the number of I/O address lines decoded by the PC Card in the specified

socket. Access to the I/O window is not enabled until `csx_RequestConfiguration(9F)` has been invoked successfully.

ReqOffset

This field is a Solaris-specific extension that can be used by clients to generate optimum window offsets passed to `csx_MapMemPage(9F)`.

## RETURN VALUES

<code>CS_SUCCESS</code>	Successful operation.
<code>CS_BAD_ATTRIBUTE</code>	Attributes are invalid.
<code>CS_BAD_SPEED</code>	Speed is invalid.
<code>CS_BAD_HANDLE</code>	Client handle is invalid.
<code>CS_BAD_SIZE</code>	Window size is invalid.
<code>CS_NO_CARD</code>	No PC Card in socket.
<code>CS_OUT_OF_RESOURCE</code>	Unable to allocate window.
<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.

## CONTEXT

These functions may be called from user or kernel context.

## SEE ALSO

`csx_ConvertSpeed(9F)`, `csx_MapMemPage(9F)`, `csx_ModifyWindow(9F)`, `csx_RegisterClient(9F)`, `csx_RequestConfiguration(9F)`

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_ResetFunction – reset a function on a PC card	
<b>SYNOPSIS</b>	#include <sys/pccard.h>	
	int32_t csx_ResetFunction(client_handle_t ch, reset_function_t *rf);	
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)	
<b>PARAMETERS</b>	<b>ch</b>	Client handle returned from <b>csx_RegisterClient(9F)</b> .
	<b>rf</b>	Pointer to a <code>reset_function_t</code> structure.
<b>DESCRIPTION</b>	<b>csx_ResetFunction()</b> requests that the specified function on the PC card initiate a reset operation.	
<b>STRUCTURE MEMBERS</b>	The structure members of <code>reset_function_t</code> are:	
	uint32_t	Socket; /* socket number */
	uint32_t	Attributes; /* reset attributes */
	The fields are defined as follows:	
	Socket	Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.
	Attributes	Must be 0.
<b>RETURN VALUES</b>	CS_SUCCESS	Card Services has noted the reset request.
	CS_IN_USE	This Card Services implementation does not permit configured cards to be reset.
	CS_BAD_HANDLE	Client handle is invalid.
	CS_NO_CARD	No PC card in socket.
	CS_BAD_SOCKET	Specified socket or function number is invalid.
	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
<b>CONTEXT</b>	This function may be called from user or kernel context.	

**SEE ALSO** | `csx_event_handler(9E)`, `csx_RegisterClient(9F)`

| *PC Card 95 Standard*, PCMCIA/JEIDA

**NOTES** | `csx_ResetFunction()` has not been implemented in this release and always returns CS\_IN\_USE.

<b>NAME</b>	csx_SetEventMask, csx_GetEventMask – set or return the client event mask for the client
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_SetEventMask(client_handle_t ch, sockevent_t * se); int32_t csx_GetEventMask(client_handle_t ch, sockevent_t * se);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b> .</p> <p><b>se</b> Pointer to a <b>sockevent_t</b> structure</p>
<b>DESCRIPTION</b>	<p>The function <b>csx_SetEventMask()</b> sets the client or global event mask for the client.</p> <p>The function <b>csx_GetEventMask()</b> returns the client or global event mask for the client.</p> <p><b>csx_RequestSocketMask(9F)</b> must be called before calling <b>csx_SetEventMask()</b> for the client event mask for this socket.</p>
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>sockevent_t</b> are:</p> <pre>uint32_t   uint32_t   /* attribute flags for call */ uint32_t   EventMask; /* event mask to set or return */ uint32_t   Socket;    /* socket number if necessary */</pre> <p>The fields are defined as follows:</p> <p>Attributes</p> <p>This is a bit-mapped field that identifies the type of event mask to be returned. The field is defined as follows:</p> <p><b>CONF_EVENT_MASK_GLOBAL</b> Client's global event mask. If set, the client's global event mask is returned.</p> <p><b>CONF_EVENT_MASK_CLIENT</b> Client's local event mask. If set, the client's local event mask is returned.</p> <p>EventMask</p>

This field is bit-mapped. Card Services performs event notification based on this field. See `csx_event_handler(9E)` for valid event definitions and for additional information about handling events.

Socket

Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.

#### RETURN VALUES

<code>CS_SUCCESS</code>	Successful operation.
<code>CS_BAD_HANDLE</code>	Client handle is invalid.
<code>CS_BAD_SOCKET</code>	<code>csx_RequestSocketMask(9F)</code> not called for <code>CONF_EVENT_MASK_CLIENT</code> .
<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.

#### CONTEXT

These functions may be called from user or kernel context.

#### SEE ALSO

`csx_event_handler(9E)`, `csx_RegisterClient(9F)`, `csx_ReleaseSocketMask(9F)`, `csx_RequestSocketMask(9F)`

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	csx_SetHandleOffset – set current access handle offset
<b>SYNOPSIS</b>	#include <sys/pccard.h>  int32_t csx_SetHandleOffset(acc_handle_t handle, uint32_t offset);
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b>handle</b> Access handle returned by <b>csx_RequestIRQ(9F)</b> or <b>csx_RequestIO(9F)</b> .  <b>offset</b> New access handle offset.
<b>DESCRIPTION</b>	This function sets the current offset for the access handle, <i>handle</i> , to <i>offset</i> .
<b>RETURN VALUES</b>	CS_SUCCESS        Successful operation.
<b>CONTEXT</b>	This function may be called from user or kernel context.
<b>SEE ALSO</b>	<b>csx_GetHandleOffset(9F)</b> , <b>csx_RequestIO(9F)</b> , <b>csx_RequestIRQ(9F)</b>  <i>PC Card 95 Standard</i> , PCMCIA/JEIDA

<b>NAME</b>	csx_ValidateCIS – validate the Card Information Structure (CIS)								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/pccard.h&gt;  int32_t csx_ValidateCIS(client_handle_t ch, cisinfo_t *ci);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)								
<b>PARAMETERS</b>	<p><b>ch</b> Client handle returned from <b>csx_RegisterClient(9F)</b>.</p> <p><b>ci</b> Pointer to a <b>cisinfo_t</b> structure.</p>								
<b>DESCRIPTION</b>	This function validates the Card Information Structure (CIS) on the PC Card in the specified socket.								
<b>STRUCTURE MEMBERS</b>	<p>The structure members of <b>cisinfo_t</b> are:</p> <pre>uint32_t Socket; /* socket number to validate CIS on */ uint32_t Chains; /* number of tuple chains in CIS */ uint32_t Tuples; /* total number of tuples in CIS */</pre> <p>The fields are defined as follows:</p> <p><b>Socket</b> Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</p> <p><b>Chains</b> This field returns the number of valid tuple chains located in the CIS. If 0 is returned, the CIS is not valid.</p> <p><b>Tuples</b> This field is a Solaris-specific extension and it returns the total number of tuples on all the chains in the PC Card's CIS.</p>								
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>CS_SUCCESS</td> <td>Successful operation.</td> </tr> <tr> <td>CS_NO_CIS</td> <td>No CIS on PC Card or CIS is invalid.</td> </tr> <tr> <td>CS_NO_CARD</td> <td>No PC Card in socket.</td> </tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td> <td>No PCMCIA hardware installed.</td> </tr> </table>	CS_SUCCESS	Successful operation.	CS_NO_CIS	No CIS on PC Card or CIS is invalid.	CS_NO_CARD	No PC Card in socket.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.								
CS_NO_CIS	No CIS on PC Card or CIS is invalid.								
CS_NO_CARD	No PC Card in socket.								
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.								
<b>CONTEXT</b>	This function may be called from user or kernel context.								

**SEE ALSO**

`csx_GetFirstTuple(9F)`, `csx_GetTupleData(9F)`, `csx_ParseTuple(9F)`,  
`csx_RegisterClient(9F)`

*PC Card 95 Standard*, PCMCIA/JEIDA

<b>NAME</b>	datamsg – test whether a message is a data message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;</pre> <p>int <b>datamsg</b>(unsigned char <i>type</i>);</p>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>type</b> The type of message to be tested. The <code>db_type</code> field of the <b>datab</b>(9S) structure contains the message type. This field may be accessed through the message block using <code>mp-&gt;b_datap-&gt;db_type</code>.</p>
<b>DESCRIPTION</b>	<b>datamsg</b> () tests the type of message to determine if it is a data message type ( <code>M_DATA</code> , <code>M_DELAY</code> , <code>M_PROTO</code> , or <code>M_PCPROTO</code> ).
<b>RETURN VALUES</b>	<p><b>datamsg</b> returns</p> <p>1 if the message is a data message</p> <p>0 otherwise.</p>
<b>CONTEXT</b>	<b>datamsg</b> () can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> The <b>put</b>(9E) routine enqueues all data messages for handling by the <b>srv</b>(9E) (service) routine. All non-data messages are handled in the <b>put</b>(9E) routine.</p> <pre> 1 xxxput(q, mp) 2     queue_t *q; 3     mblk_t *mp; 4 { 5     if (datamsg(mp-&gt;b_datap-&gt;db_type)) { 6         putq(q, mp); 7         return; 8     } 9     switch (mp-&gt;b_datap-&gt;db_type) { 10    case M_FLUSH: 11        ... 12    } </pre>
<b>SEE ALSO</b>	<p><b>put</b>(9E), <b>srv</b>(9E), <b>allocb</b>(9F), <b>datab</b>(9S), <b>msgb</b>(9S)</p> <p><i>Writing Device Drivers</i></p>

*STREAMS Programming Guide*

<b>NAME</b>	ddi_add_intr, ddi_get_iblock_cookie, ddi_remove_intr – hardware interrupt handling routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_get_iblock_cookie(dev_info_t * dip, uint_t inumber, ddi_iblock_cookie_t * iblock_cookiep);  int ddi_add_intr(dev_info_t * dip, uint_t inumber, ddi_iblock_cookie_t * iblock_cookiep, ddi_idevice_cookie_t * idevice_cookiep, uint_t (* int_handler)(caddr_t), caddr_t int_handler_arg);  void ddi_remove_intr(dev_info_t * dip, uint_t inumber, ddi_iblock_cookie_t iblock_cookie);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
ddi_get_iblock_cookie()	<p><b>dip</b> Pointer to dev_info structure.</p> <p><b>inumber</b> Interrupt number.</p> <p><b>iblock_cookiep</b> Pointer to an interrupt block cookie.</p>
ddi_add_intr()	<p><b>dip</b> Pointer to dev_info structure.</p> <p><b>inumber</b> Interrupt number.</p> <p><b>iblock_cookiep</b> Optional pointer to an interrupt block cookie where a returned interrupt block cookie is stored.</p> <p><b>idevice_cookiep</b> Optional pointer to an interrupt device cookie where a returned interrupt device cookie is stored.</p> <p><b>int_handler</b> Pointer to interrupt handler.</p> <p><b>int_handler_arg</b> Argument for interrupt handler.</p>
ddi_remove_intr()	<p><b>dip</b> Pointer to dev_info structure.</p> <p><b>inumber</b> Interrupt number.</p>

***iblock\_cookie*** Block cookie which identifies the interrupt handler to be removed.

## DESCRIPTION

**ddi\_get\_iblock\_cookie()** retrieves the interrupt block cookie associated with a particular interrupt specification. This routine should be called before **ddi\_add\_intr()** to retrieve the interrupt block cookie needed to initialize locks (**mutex(9F)**, **rwlock(9F)**) used by the interrupt routine. The interrupt number *inumber* determines which interrupt specification to retrieve the cookie for. *inumber* is associated with information provided either by the device (see **sbus(4)**) or the hardware configuration file (see **vme(4)**, **sysbus(4)**, **isa(4)**, **eisa(4)**, and **driver.conf(4)**). If only one interrupt is associated with the device, *inumber* should be 0.

On a successful return, *\*iblock\_cookiep* contains information needed for initializing locks associated with the interrupt specification corresponding to *inumber* (see **mutex\_init(9F)** and **rw\_init(9F)**). The driver can then initialize locks acquired by the interrupt routine before calling **ddi\_add\_intr()** which prevents a possible race condition where the driver's interrupt handler is called immediately *after* the driver has called **ddi\_add\_intr()** but *before* the driver has initialized the locks. This may happen when an interrupt for a different device occurs on the same interrupt level. If the interrupt routine acquires the lock before the lock has been initialized, undefined behavior may result.

**ddi\_add\_intr()** adds an interrupt handler to the system. The interrupt number *inumber* determines which interrupt the handler will be associated with. (Refer to **ddi\_get\_iblock\_cookie()** above.)

On a successful return, *iblock\_cookiep* contains information used for initializing locks associated with this interrupt specification (see **mutex\_init(9F)** and **rw\_init(9F)**). Note that the interrupt block cookie is usually obtained using **ddi\_get\_iblock\_cookie()** to avoid the race conditions described above (refer to **ddi\_get\_iblock\_cookie()** above). For this reason, *iblock\_cookiep* is no longer useful and should be set to **NULL**.

On a successful return, *idevice\_cookiep* contains a pointer to a **ddi\_idevice\_cookie\_t** structure (see **ddi\_idevice\_cookie(9S)**) containing information useful for some devices that have programmable interrupts. If *idevice\_cookiep* is set to **NULL**, no value is returned.

The routine *intr\_handler*, with its argument *int\_handler\_arg*, is called upon receipt of the appropriate interrupt. The interrupt handler should return **DDI\_INTR\_CLAIMED** if the interrupt was claimed, **DDI\_INTR\_UNCLAIMED** otherwise.

	If successful, <b>ddi_add_intr()</b> will return <code>DDI_SUCCESS</code> ; if the interrupt information cannot be found, it will return <code>DDI_INTR_NOTFOUND</code> .
<b>ddi_remove_intr()</b>	<b>ddi_remove_intr()</b> removes an interrupt handler from the system. Unloadable drivers should call this routine during their <b>detach(9E)</b> routine to remove their interrupt handler from the system.  The device interrupt routine for this instance of the device will not execute after <b>ddi_remove_intr()</b> returns. <b>ddi_remove_intr()</b> may need to wait for the device interrupt routine to complete before returning. Therefore, locks acquired by the interrupt handler should not be held across the call to <b>ddi_remove_intr()</b> or deadlock may result.
<b>RETURN VALUES</b>	<b>ddi_add_intr()</b> and <b>ddi_get_iblock_cookie()</b> return: <code>DDI_SUCCESS</code> On success.  <code>DDI_INTR_NOTFOUND</code> On failure to find the interrupt.
<b>CONTEXT</b>	<b>ddi_add_intr()</b> , <b>ddi_remove_intr()</b> , and <b>ddi_get_iblock_cookie()</b> can be called from user or kernel context.
<b>SEE ALSO</b>	<b>driver.conf(4)</b> , <b>eisa(4)</b> , <b>isa(4)</b> , <b>sbus(4)</b> , <b>sysbus(4)</b> , <b>vme(4)</b> , <b>attach(9E)</b> , <b>detach(9E)</b> , <b>ddi_intr_hilevel(9F)</b> , <b>mutex(9F)</b> , <b>mutex_init(9F)</b> , <b>rw_init(9F)</b> , <b>rwlock(9F)</b> , <b>ddi_idevice_cookie(9S)</b>  <i>Writing Device Drivers</i>
<b>NOTES</b>	<b>ddi_get_iblock_cookie()</b> must not be called <i>after</i> the driver adds an interrupt handler for the interrupt specification corresponding to <i>inumber</i> .
<b>BUGS</b>	The <i>idevice_cookiep</i> should really point to a data structure that is specific to the bus architecture that the device operates on. Currently only VMEbus and SBus are supported and a single data structure is used to describe both.

<b>NAME</b>	ddi_add_softintr, ddi_get_soft_iblock_cookie, ddi_remove_softintr, ddi_trigger_softintr – software interrupt handling routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_get_soft_iblock_cookie(dev_info_t * dip, int preference, ddi_iblock_cookie_t * iblock_cookiep);  int ddi_add_softintr(dev_info_t * dip, int preference, ddi_softintr_t * idp, ddi_iblock_cookie_t * iblock_cookiep, ddi_idevice_cookie_t * idevice_cookiep, uint_t(* int_handler)(caddr_t int_handler_arg), caddr_t int_handler_arg);  void ddi_remove_softintr(ddi_softintr_t id);  void ddi_trigger_softintr(ddi_softintr_t id);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
ddi_get_soft_iblock_cookie()	<p><b>dip</b> Pointer to a dev_info structure.</p> <p><b>preference</b> The type of soft interrupt to retrieve the cookie for.</p> <p><b>iblock_cookiep</b> Pointer to a location to store the interrupt block cookie.</p>
ddi_add_softintr()	<p><b>dip</b> Pointer to dev_info structure.</p> <p><b>preference</b> A hint value describing the type of soft interrupt to generate.</p> <p><b>idp</b> Pointer to a soft interrupt identifier where a returned soft interrupt identifier is stored.</p> <p><b>iblock_cookiep</b> Optional pointer to an interrupt block cookie where a returned interrupt block cookie is stored.</p> <p><b>idevice_cookiep</b> Optional pointer to an interrupt device cookie where a returned interrupt device cookie is stored (not used).</p> <p><b>int_handler</b> Pointer to interrupt handler.</p> <p><b>int_handler_arg</b> Argument for interrupt handler.</p>

<b>ddi_remove_softintr()</b>	<b>id</b>	The identifier specifying which soft interrupt handler to remove.
<b>ddi_trigger_softintr()</b>	<b>id</b>	The identifier specifying which soft interrupt to trigger and which soft interrupt handler will be called.
<b>DESCRIPTION</b>		
<b>ddi_get_soft_iblock_cookie()</b>	<p><b>ddi_get_soft_iblock_cookie()</b> retrieves the interrupt block cookie associated with a particular soft interrupt preference level. This routine should be called before <b>ddi_add_softintr()</b> to retrieve the interrupt block cookie needed to initialize locks ( <b>mutex(9F)</b> , <b>rwlock(9F)</b> ) used by the software interrupt routine. <i>preference</i> determines which type of soft interrupt to retrieve the cookie for. The possible values for <i>preference</i> are:</p> <p>DDI_SOFTINT_LOW            Low priority soft interrupt.</p> <p>DDI_SOFTINT_MED            Medium priority soft interrupt.</p> <p>DDI_SOFTINT_HIGH           High priority soft interrupt.</p> <p>On a successful return, <i>iblock_cookiep</i> contains information needed for initializing locks associated with this soft interrupt (see <b>mutex_init(9F)</b> and <b>rw_init(9F)</b> ). The driver can then initialize mutexes acquired by the interrupt routine before calling <b>ddi_add_softintr()</b> which prevents a possible race condition where the driver's soft interrupt handler is called immediately <i>after</i> the driver has called <b>ddi_add_softintr()</b> but <i>before</i> the driver has initialized the mutexes. This can happen when a soft interrupt for a different device occurs on the same soft interrupt priority level. If the soft interrupt routine acquires the mutex before it has been initialized, undefined behavior may result.</p>	
<b>ddi_add_softintr()</b>	<p><b>ddi_add_softintr()</b> adds a soft interrupt to the system. The user specified hint <i>preference</i> identifies three suggested levels for the system to attempt to allocate the soft interrupt priority at. The value for <i>preference</i> should be the same as that used in the corresponding call to <b>ddi_get_soft_iblock_cookie()</b> . Refer to the description of <b>ddi_get_soft_iblock_cookie()</b> above.</p> <p>The value returned in the location pointed at by <i>idp</i> is the soft interrupt identifier. This value is used in later calls to <b>ddi_remove_softintr()</b> and <b>ddi_trigger_softintr()</b> to identify the soft interrupt and the soft interrupt handler.</p> <p>The value returned in the location pointed at by <i>iblock_cookiep</i> is an interrupt block cookie which contains information used for initializing mutexes associated with this soft interrupt (see <b>mutex_init(9F)</b> and <b>rw_init(9F)</b> ).</p>	

Note that the interrupt block cookie is normally obtained using **ddi\_get\_soft\_iblock\_cookie()** to avoid the race conditions described above (refer to the description of **ddi\_get\_soft\_iblock\_cookie()** above). For this reason, *iblock\_cookiep* is no longer useful and should be set to `NULL` .

*idevice\_cookiep* is not used and should be set to `NULL` .

The routine *int\_handler* , with its argument *int\_handler\_arg* , is called upon receipt of a software interrupt. Software interrupt handlers must not assume that they have work to do when they run, since (like hardware interrupt handlers) they may run because a soft interrupt occurred for some other reason. For example, another driver may have triggered a soft interrupt at the same level. For this reason, before triggering the soft interrupt, the driver must indicate to its soft interrupt handler that it should do work. This is usually done by setting a flag in the state structure. The routine *int\_handler* checks this flag, reachable through *int\_handler\_arg* , to determine if it should claim the interrupt and do its work.

The interrupt handler must return `DDI_INTR_CLAIMED` if the interrupt was claimed, `DDI_INTR_UNCLAIMED` otherwise.

If successful, **ddi\_add\_softintr()** will return `DDI_SUCCESS` ; if the interrupt information cannot be found, it will return `DDI_FAILURE` .

**ddi\_remove\_softintr()**

**ddi\_remove\_softintr()** removes a soft interrupt from the system. The soft interrupt identifier *id* , which was returned from a call to **ddi\_add\_softintr()** , is used to determine which soft interrupt and which soft interrupt handler to remove. Drivers must remove any soft interrupt handlers before allowing the system to unload the driver.

**ddi\_trigger\_softintr()**

**ddi\_trigger\_softintr()** triggers a soft interrupt. The soft interrupt identifier *id* is used to determine which soft interrupt to trigger. This function is used by device drivers when they wish to trigger a soft interrupt which has been set up using **ddi\_add\_softintr()** .

**RETURN VALUES**

**ddi\_add\_softintr()** and **ddi\_get\_soft\_iblock\_cookie()** return:

`DDI_SUCCESS`                    on success

`DDI_FAILURE`                    on failure

**CONTEXT**

These functions can be called from user or kernel context.

**ddi\_trigger\_softintr()** may be called from high-level interrupt context as well.

**EXAMPLES**

**EXAMPLE 1**    device using high-level interrupts

In the following example, the device uses high-level interrupts. High-level interrupts are those that interrupt at the level of the scheduler and above. High

level interrupts must be handled without using system services that manipulate thread or process states, because these interrupts are not blocked by the scheduler. In addition, high level interrupt handlers must take care to do a minimum of work because they are not preemptable. See `ddi_intr_hilevel(9F)`.

In the example, the high-level interrupt routine minimally services the device, and enqueues the data for later processing by the soft interrupt handler. If the soft interrupt handler is not currently running, the high-level interrupt routine triggers a soft interrupt so the soft interrupt handler can process the data. Once running, the soft interrupt handler processes all the enqueued data before returning.

The state structure contains two mutexes. The high-level mutex is used to protect data shared between the high-level interrupt handler and the soft interrupt handler. The low-level mutex is used to protect the rest of the driver from the soft interrupt handler.

```

struct xxstate {
    ...
    ddi_softintr_t\011\011    id;
\011    ddi_iblock_cookie_t\011 high_iblock_cookie;
\011    kmutex_t\011\011\011    high_mutex;
\011    ddi_iblock_cookie_t\011 low_iblock_cookie;
\011    kmutex_t\011\011\011    low_mutex;
\011    int\011\011\011\011    softint_running;
    ...
};
struct xxstate *xsp;
static uint_t xxsoftintr(caddr_t);
static uint_t xxhighintr(caddr_t);
...

```

#### EXAMPLE 2 sample attach() routine

The following code fragment would usually appear in the driver's `attach(9E)` routine. `ddi_add_intr(9F)` is used to add the high-level interrupt handler and `ddi_add_softintr()` is used to add the low-level interrupt routine.

```

static uint_t
xxattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    \011struct xxstate *xsp;
    \011...
    /* get high-level iblock cookie */
    \011if (ddi_get_iblock_cookie(dip,
inumber
,
\011\011    &xsp->high_iblock_cookie) != DDI_SUCCESS) {
\011\011\011    /* clean up */
\011\011\011    return (DDI_FAILURE); /* fail attach */

```

```

\011}
\011
\011/* initialize high-level mutex */
\011mutex_init(&xsp->high_mutex, "xx high mutex", MUTEX_DRIVER,
\011\011      (void *)xsp->high_iblock_cookie);
\011
\011/* add high-level routine - xxhighintr() */
\011if (ddi_add_intr(dip,
inumber
, NULL, NULL,
\011\011      xxhighintr, (caddr_t) xsp) != DDI_SUCCESS) {
\011\011\011      /* cleanup */
\011\011\011      return (DDI_FAILURE); /* fail attach */
\011}

\011/* get soft iblock cookie */
\011if (ddi_get_soft_iblock_cookie(dip, DDI_SOFTINT_MED,
\011\011      &xsp->low_iblock_cookie) != DDI_SUCCESS) {
\011\011\011      /* clean up */
\011\011\011      return (DDI_FAILURE); /* fail attach */
\011}

\011/* initialize low-level mutex */
\011mutex_init(&xsp->low_mutex, "xx low mutex", MUTEX_DRIVER,
\011\011      (void *)xsp->low_iblock_cookie);

\011      /* add low level routine - xxsoftintr() */
\011      if ( ddi_add_softintr(dip, DDI_SOFTINT_MED, &xsp->id,
\011\011      NULL, NULL, xxsoftintr, (caddr_t) xsp) != DDI_SUCCESS) {
\011\011\011      /* cleanup */
\011\011\011      return (DDI_FAILURE); /* fail attach */
\011}
\011
\011      ...
}

```

**EXAMPLE 3** High-level interrupt routine

The next code fragment represents the high-level interrupt routine. The high-level interrupt routine minimally services the device, and enqueues the data for later processing by the soft interrupt routine. If the soft interrupt routine is not already running, **ddi\_trigger\_softintr()** is called to start the routine. The soft interrupt routine will run until there is no more data on the queue.

```

static uint_t
xxhighintr(caddr_t arg)
{
    struct xxstate *xsp = (struct xxstate *) arg;
\011    int need_softint;
\011    ...
\011    mutex_enter(&xsp->high_mutex);
\011/*

```

```

\011     * Verify this device generated the interrupt
\011     * and disable the device interrupt.
\011     * Enqueue data for xxsoftintr() processing.
\011     */

\011     /* is xxsoftintr() already running ? */
\011\011     \011if (xsp->softint_running)
\011\011         need_softint = 0;
\011     else
\011\011         need_softint = 1;
\011     mutex_exit(&xsp->high_mutex);

\011     /* read-only access to xsp->id, no mutex needed */
\011     if (need_softint)
\011\011         ddi_trigger_softintr(xsp->id);
\011     ...
\011     return (DDI_INTR_CLAIMED);
}

static uint_t
xxsoftintr(caddr_t arg)
{
    struct xxstate *xsp = (struct xxstate *) arg;
    ...
\011     mutex_enter(&xsp->low_mutex);
    mutex_enter(&xsp->high_mutex);

    /* verify there is work to do */
    if (
work queue empty
|| xsp->softint_running ) {
\011\011         mutex_exit(&xsp->high_mutex);
\011\011         mutex_exit(&xsp->low_mutex);
\011\011         return (DDI_INTR_UNCLAIMED);
    }

    xsp->softint_running = 1;

\011     while (
data on queue
) {
\011\011         ASSERT(mutex_owned(&xsp->high_mutex));

\011\011         /* de-queue data */

\011\011         mutex_exit(&xsp->high_mutex);
\011\011
\011\011         /* Process data on queue */

\011\011         mutex_enter(&xsp->high_mutex);
\011     }

\011     xsp->softint_running = 0;
\011     mutex_exit(&xsp->high_mutex);
\011     mutex_exit(&xsp->low_mutex);

\011     return (DDI_INTR_CLAIMED);
}

```

```
}
```

**SEE ALSO**

**ddi\_add\_intr(9F)** , **ddi\_intr\_hilevel(9F)** , **ddi\_remove\_intr(9F)** ,  
**mutex\_init(9F)**

*Writing Device Drivers*

**NOTES**

**ddi\_add\_softintr()** may not be used to add the same software interrupt handler more than once. This is true even if a different value is used for *int\_handler\_arg* in each of the calls to **ddi\_add\_softintr()** . Instead, the argument passed to the interrupt handler should indicate what service(s) the interrupt handler should perform. For example, the argument could be a pointer to the device's soft state structure, which could contain a 'which\_service' field that the handler examines. The driver must set this field to the appropriate value before calling **ddi\_trigger\_softintr()** .

<b>NAME</b>	ddi_binding_name, ddi_get_name – return driver binding name
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  char * ddi_binding_name(dev_info_t * dip); char * ddi_get_name(dev_info_t * dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>dip</b> A pointer to the device's dev_info structure.
<b>DESCRIPTION</b>	<b>ddi_binding_name()</b> and <b>ddi_get_name()</b> return the driver binding name. This is the name used to select a driver for the device. This name is typically derived from the device name property or the device compatible property. The name returned may be a driver alias or the driver name.
<b>RETURN VALUES</b>	<b>ddi_binding_name()</b> and <b>ddi_get_name()</b> return the name used to bind a driver to a device.
<b>CONTEXT</b>	<b>ddi_binding_name()</b> and <b>ddi_get_name()</b> can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	ddi_node_name(9F) <i>Writing Device Drivers</i>
<b>WARNINGS</b>	The name returned by <b>ddi_binding_name()</b> and <b>ddi_get_name()</b> is read-only.

<b>NAME</b>	ddi_btop, ddi_btopr, ddi_ptob – page size conversions
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>unsigned long ddi_btop(dev_info_t *dip, unsigned long bytes);</p> <p>unsigned long ddi_btopr(dev_info_t *dip, unsigned long bytes);</p> <p>unsigned long ddi_ptob(dev_info_t *dip, unsigned long pages);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p>This set of routines use the parent nexus driver to perform conversions in page size units.</p> <p><b>ddi_btop()</b> converts the given number of bytes to the number of memory pages that it corresponds to, rounding down in the case that the byte count is not a page multiple.</p> <p><b>ddi_btopr()</b> converts the given number of bytes to the number of memory pages that it corresponds to, rounding up in the case that the byte count is not a page multiple.</p> <p><b>ddi_ptob()</b> converts the given number of pages to the number of bytes that it corresponds to.</p> <p>Because bus nexus may possess their own hardware address translation facilities, these routines should be used in preference to the corresponding DDI/DKI routines <b>btop(9F)</b> , <b>btopr(9F)</b> , and <b>ptob(9F)</b> , which only deal in terms of the pagesize of the main system MMU.</p>
<b>RETURN VALUES</b>	<p><b>ddi_btop()</b> and <b>ddi_btopr()</b> return the number of corresponding pages. <b>ddi_ptob()</b> returns the corresponding number of bytes. There are no error return values.</p>
<b>CONTEXT</b>	This function can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Find the size (in bytes) of one page</p> <pre>pagesize = ddi_ptob(dip, 1L);</pre>
<b>SEE ALSO</b>	<p><b>btop(9F)</b> , <b>btopr(9F)</b> , <b>ptob(9F)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_copyin – copy data to a driver buffer
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int <b>ddi_copyin</b>(const void *buf, void *driverbuf, size_t cn, int flags);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>buf</b>                    Source address from which data is transferred.</p> <p><b>driverbuf</b>            Driver destination address to which data is transferred.</p> <p><b>cn</b>                     Number of bytes transferred.</p> <p><b>flags</b>                 Set of flag bits that provide address space information about <i>buf</i>.</p>
<b>DESCRIPTION</b>	<p>This routine is designed for use in driver <b>ioctl</b>(9E) routines for drivers that support layered ioctls. <b>ddi_copyin()</b> copies data from a source address to a driver buffer. The driver developer must ensure that adequate space is allocated for the destination address.</p> <p>The <i>flags</i> argument is used to determine the address space information about <i>buf</i>. If the <b>FKIOCTL</b> flag is set, this indicates that <i>buf</i> is a kernel address, and <b>ddi_copyin()</b> behaves like <b>bcopy</b>(9F). Otherwise <i>buf</i> is interpreted as a user buffer address, and <b>ddi_copyin()</b> behaves like <b>copyin</b>(9F).</p> <p>Addresses that are word-aligned are moved most efficiently. However, the driver developer is not obliged to ensure alignment. This function automatically finds the most efficient move according to address alignment.</p>
<b>RETURN VALUES</b>	<p><b>ddi_copyin()</b> returns 0, indicating a successful copy. It returns -1 if one of the following occurs:</p> <ul style="list-style-type: none"> <li>■ paging fault; the driver tried to access a page of memory for which it did not have read or write access</li> <li>■ invalid user address, such as a user area or stack area</li> <li>■ invalid address that would have resulted in data being copied into the user block</li> </ul> <p>If -1 is returned to the caller, driver entry point routines should return <b>EFAULT</b>.</p>

**CONTEXT**

**ddi\_copyin()** can be called from user or kernel context only.

**EXAMPLES****EXAMPLE 1 ddi\_copyin() example**

A driver `ioctl(9E)` routine (line 12) can be used to get or set device attributes or registers. For the `XX_SETREGS` condition (line 25), the driver copies the user data in `arg` to the device registers. If the specified argument contains an invalid address, an error code is returned.

```

1  struct device { /* layout of physical device registers */
2      int     control; /* physical device control word */
3      int     status; /* physical device status word */
4      short   recv_char; /* receive character from device */
5      short   xmit_char; /* transmit character to device */
6  };
7  struct device_state {
8      volatile struct device *regsp; /* pointer to device registers */
9      kmutex_t reg_mutex; /* protect device registers */
10     . . .
11 };
12 static void *statep; /* for soft state routines */
13
14 xxioctl(dev_t dev, int cmd, int arg, int mode,
15         cred_t *cred_p, int *rval_p)
16 {
17     struct device_state *sp;
18     volatile struct device *rp;
19     struct device reg_buf; /* temporary buffer for registers */
20     int instance;
21
22     instance = getminor(dev);
23     sp = ddi_get_soft_state(statep, instance);
24     if (sp == NULL)
25         return (ENXIO);
26     rp = sp->regsp;
27     . . .
28     switch (cmd) {
29
30     case XX_GETREGS: /* copy data to temp. regs. buf */
31         if (ddi_copyin(arg, &reg_buf,
32             sizeof (struct device), mode) != 0) {
33             return (EFAULT);
34         }
35
36         mutex_enter(&sp->reg_mutex);
37         /*
38          * Copy data from temporary device register
39          * buffer to device registers.
40          * e.g. rp->control = reg_buf.control;
41          */
42         mutex_exit(&sp->reg_mutex);
43
44         break;
45     }
46 }

```

**SEE ALSO** `ioctl(9E)`, `bcopy(9F)`, `copyin(9F)`, `copyout(9F)`, `ddi_copyout(9F)`, `uiomove(9F)`

*Writing Device Drivers*

**NOTES** The value of the *flags* argument to `ddi_copyin()` should be passed through directly from the *mode* argument of `ioctl()` untranslated.

Driver defined locks should not be held across calls to this function.

This should not be used from a streams driver. See `M_COPYIN` and `M_COPYOUT` in *STREAMS Programming Guide*.

<b>NAME</b>	ddi_copyout – copy data from a driver								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_copyout(const void *driverbuf, void *buf, size_t cn, int flags);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).								
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="padding-right: 1em;"><b>driverbuf</b></td> <td>Source address in the driver from which the data is transferred.</td> </tr> <tr> <td style="padding-right: 1em;"><b>buf</b></td> <td>Destination address to which the data is transferred.</td> </tr> <tr> <td style="padding-right: 1em;"><b>cn</b></td> <td>Number of bytes to copy.</td> </tr> <tr> <td style="padding-right: 1em;"><b>flags</b></td> <td>Set of flag bits that provide address space information about <i>buf</i>.</td> </tr> </table>	<b>driverbuf</b>	Source address in the driver from which the data is transferred.	<b>buf</b>	Destination address to which the data is transferred.	<b>cn</b>	Number of bytes to copy.	<b>flags</b>	Set of flag bits that provide address space information about <i>buf</i> .
<b>driverbuf</b>	Source address in the driver from which the data is transferred.								
<b>buf</b>	Destination address to which the data is transferred.								
<b>cn</b>	Number of bytes to copy.								
<b>flags</b>	Set of flag bits that provide address space information about <i>buf</i> .								
<b>DESCRIPTION</b>	<p>This routine is designed for use in driver <code>ioctl1(9E)</code> routines for drivers that support layered <code>ioctls</code>. <code>ddi_copyout()</code> copies data from a driver buffer to a destination address, <i>buf</i>.</p> <p>The <i>flags</i> argument is used to determine the address space information about <i>buf</i>. If the <code>FKIOCTL</code> flag is set, this indicates that <i>buf</i> is a kernel address, and <code>ddi_copyout()</code> behaves like <code>bcopy(9F)</code>. Otherwise <i>buf</i> is interpreted as a user buffer address, and <code>ddi_copyout()</code> behaves like <code>copyout(9F)</code>.</p> <p>Addresses that are word-aligned are moved most efficiently. However, the driver developer is not obliged to ensure alignment. This function automatically finds the most efficient move algorithm according to address alignment.</p>								
<b>RETURN VALUES</b>	<p>Under normal conditions, 0 is returned to indicate a successful copy. Otherwise, -1 is returned if one of the following occurs:</p> <ul style="list-style-type: none"> <li>■ paging fault; the driver tried to access a page of memory for which it did not have read or write access</li> <li>■ invalid user address, such as a user area or stack area</li> <li>■ invalid address that would have resulted in data being copied into the user block</li> </ul>								

If -1 is returned to the caller, driver entry point routines should return EFAULT.

**CONTEXT**

**ddi\_copyout()** can be called from user or kernel context only.

**EXAMPLES****EXAMPLE 1 ddi\_copyout() example**

A driver `ioctl(9E)` routine (line 12) can be used to get or set device attributes or registers. In the `XX_GETREGS` condition (line 25), the driver copies the current device register values to another data area. If the specified argument contains an invalid address, an error code is returned.

```

1  struct device {          /* layout of physical device registers */
2  int     control;        /* physical device control word */
3  int     status;        /* physical device status word */
4  short   recv_char;     /* receive character from device */
5  short   xmit_char;     /* transmit character to device */
6  };

7  struct device_state {
8  volatile struct device *regsp; /* pointer to device registers */
9  kmutex_t reg_mutex;         /* protect device registers */
10 };

11 static void *statep; /* for soft state routines */

12 xxioctl(dev_t dev, int cmd, int arg, int mode,
13         cred_t *cred_p, int *rval_p)
14 {
15     struct device_state *sp;
16     volatile struct device *rp;
17     struct device reg_buf; /* temporary buffer for registers */
18     int instance;

19     instance = getminor(dev);
20     sp = ddi_get_soft_state(statep, instance);
21     if (sp == NULL)
22         return (ENXIO);
23     rp = sp->regsp;
24     . . .
25     switch (cmd) {
26     case XX_GETREGS: /* copy registers to arg */
27         mutex_enter(&sp->reg_mutex);
28         /*
29          * Copy data from device registers to
30          * temporary device register buffer
31          * e.g. reg_buf.control = rp->control;
32          */
33         mutex_exit(&sp->reg_mutex);
34         if (ddi_copyout(&reg_buf, arg,
35             sizeof (struct device), mode) != 0) {
36             return (EFAULT);
37         }
38     }

```

```
37         break;
38     }
39 }
```

**SEE ALSO** `ioctl(9E)`, `bcopy(9F)`, `copyin(9F)`, `copyout(9F)`, `ddi_copyin(9F)`, `uiomove(9F)`

*Writing Device Drivers*

**NOTES** The value of the *flags* argument to `ddi_copyout()` should be passed through directly from the *mode* argument of `ioctl()` untranslated.

Driver defined locks should not be held across calls to this function.

This should not be used from a streams driver. See `M_COPYIN` and `M_COPYOUT` in *STREAMS Programming Guide*.

<b>NAME</b>	ddi_create_minor_node – create a minor node for this device																																		
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stat.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_create_minor_node(dev_info_t *dip, char *name, int spec_type, minor_t minor_num, char *node_type, int is_clone);</pre>																																		
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).																																		
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><b>dip</b></td> <td>A pointer to the device's dev_info structure.</td> </tr> <tr> <td style="padding-right: 20px;"><b>name</b></td> <td>The name of this particular minor device.</td> </tr> <tr> <td style="padding-right: 20px;"><b>spec_type</b></td> <td>S_IFCHR or S_IFBLK for character or block minor devices respectively.</td> </tr> <tr> <td style="padding-right: 20px;"><b>minor_num</b></td> <td>The minor number for this particular minor device.</td> </tr> <tr> <td style="padding-right: 20px;"><b>node_type</b></td> <td>Any string that uniquely identifies the type of node. The following predefined node types are provided with this release:</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_SERIAL</td> <td>For serial ports</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_SERIAL_MB</td> <td>For on board serial ports</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_SERIAL_DO</td> <td>For dial out ports</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_SERIAL_MB_DO</td> <td>For on board dial out ports</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_BLOCK</td> <td>For hard disks</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_BLOCK_CHAN</td> <td>For hard disks with channel or target numbers</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_CD</td> <td>For CDROM drives</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_CD_CHAN</td> <td>For CDROM drives with channel or target numbers</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_FD</td> <td>For floppy disks</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_TAPE</td> <td>For tape drives</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_NET</td> <td>For network devices</td> </tr> <tr> <td style="padding-left: 40px;">DDI_NT_DISPLAY</td> <td>For display devices</td> </tr> </table>	<b>dip</b>	A pointer to the device's dev_info structure.	<b>name</b>	The name of this particular minor device.	<b>spec_type</b>	S_IFCHR or S_IFBLK for character or block minor devices respectively.	<b>minor_num</b>	The minor number for this particular minor device.	<b>node_type</b>	Any string that uniquely identifies the type of node. The following predefined node types are provided with this release:	DDI_NT_SERIAL	For serial ports	DDI_NT_SERIAL_MB	For on board serial ports	DDI_NT_SERIAL_DO	For dial out ports	DDI_NT_SERIAL_MB_DO	For on board dial out ports	DDI_NT_BLOCK	For hard disks	DDI_NT_BLOCK_CHAN	For hard disks with channel or target numbers	DDI_NT_CD	For CDROM drives	DDI_NT_CD_CHAN	For CDROM drives with channel or target numbers	DDI_NT_FD	For floppy disks	DDI_NT_TAPE	For tape drives	DDI_NT_NET	For network devices	DDI_NT_DISPLAY	For display devices
<b>dip</b>	A pointer to the device's dev_info structure.																																		
<b>name</b>	The name of this particular minor device.																																		
<b>spec_type</b>	S_IFCHR or S_IFBLK for character or block minor devices respectively.																																		
<b>minor_num</b>	The minor number for this particular minor device.																																		
<b>node_type</b>	Any string that uniquely identifies the type of node. The following predefined node types are provided with this release:																																		
DDI_NT_SERIAL	For serial ports																																		
DDI_NT_SERIAL_MB	For on board serial ports																																		
DDI_NT_SERIAL_DO	For dial out ports																																		
DDI_NT_SERIAL_MB_DO	For on board dial out ports																																		
DDI_NT_BLOCK	For hard disks																																		
DDI_NT_BLOCK_CHAN	For hard disks with channel or target numbers																																		
DDI_NT_CD	For CDROM drives																																		
DDI_NT_CD_CHAN	For CDROM drives with channel or target numbers																																		
DDI_NT_FD	For floppy disks																																		
DDI_NT_TAPE	For tape drives																																		
DDI_NT_NET	For network devices																																		
DDI_NT_DISPLAY	For display devices																																		

	DDI_PSEUDO	For pseudo devices
	<b><i>is_clone</i></b>	If the device is a clone device then this flag is set to CLONE_DEV else it is set to 0.
<b>DESCRIPTION</b>	<b>ddi_create_minor_node()</b> provides the necessary information to enable the system to create the /dev and /devices hierarchies. The <i>name</i> is used to create the minor name of the block or character special file under the /devices hierarchy. At-sign (@), slash (/), and space are not allowed. The <i>spec_type</i> specifies whether this is a block or character device. The <i>minor_num</i> is the minor number for the device. The <i>node_type</i> is used to create the names in the /dev hierarchy that refers to the names in the /devices hierarchy. See <b>disks(1M)</b> , <b>ports(1M)</b> , <b>tapes(1M)</b> , <b>devlinks(1M)</b> . Finally <i>is_clone</i> determines if this is a clone device or not.	
<b>RETURN VALUES</b>	<b>ddi_create_minor_node()</b> returns:	
	DDI_SUCCESS	if it was able to allocate memory, create the minor data structure, and place it into the linked list of minor devices for this driver.
	DDI_FAILURE	if minor node creation failed.
<b>EXAMPLES</b>	<b>EXAMPLE 1</b> create a data structure describing a minor device with minor number of 0	
	The following example creates a data structure describing a minor device called <i>foo</i> which has a minor number of 0. It is of type DDI_NT_BLOCK (a block device) and it is not a clone device.	
	<pre>ddi_create_minor_node(dip, "foo", S_IFBLK, 0, DDI_NT_BLOCK, 0);</pre>	
<b>SEE ALSO</b>	<b>add_drv(1M)</b> , <b>devlinks(1M)</b> , <b>disks(1M)</b> , <b>drvconfig(1M)</b> , <b>ports(1M)</b> , <b>tapes(1M)</b> , <b>attach(9E)</b> , <b>ddi_remove_minor_node(9F)</b>	
	<i>Writing Device Drivers</i>	

<b>NAME</b>	ddi_device_copy – copy data from one device register to another device register								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int ddi_device_copy(ddi_acc_handle_t src_handle, caddr_t src_addr, ssize_t src_advcnt, ddi_acc_handle_t dest_handle, caddr_t dest_addr, ssize_t dest_advcnt, size_t bytecount, uint_t dev_datsz);</p>								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).								
<b>PARAMETERS</b>	<p><b>src_handle</b>      The data access handle of the source device.</p> <p><b>src_addr</b>        Base data source address.</p> <p><b>src_advcnt</b>      Number of <i>dev_datsz</i> units to advance on every access.</p> <p><b>dest_handle</b>     The data access handle of the destination device.</p> <p><b>dest_addr</b>      Base data destination address.</p> <p><b>dest_advcnt</b>    Number of <i>dev_datsz</i> units to advance on every access.</p> <p><b>bytecount</b>      Number of bytes to transfer.</p> <p><b>dev_datsz</b>      The size of each data word. Possible values are defined as:</p> <table border="0" style="margin-left: 40px;"> <tr> <td>DDI_DATA_SZ01_ACC</td> <td>1 byte data size</td> </tr> <tr> <td>DDI_DATA_SZ02_ACC</td> <td>2 bytes data size</td> </tr> <tr> <td>DDI_DATA_SZ04_ACC</td> <td>4 bytes data size</td> </tr> <tr> <td>DDI_DATA_SZ08_ACC</td> <td>8 bytes data size</td> </tr> </table>	DDI_DATA_SZ01_ACC	1 byte data size	DDI_DATA_SZ02_ACC	2 bytes data size	DDI_DATA_SZ04_ACC	4 bytes data size	DDI_DATA_SZ08_ACC	8 bytes data size
DDI_DATA_SZ01_ACC	1 byte data size								
DDI_DATA_SZ02_ACC	2 bytes data size								
DDI_DATA_SZ04_ACC	4 bytes data size								
DDI_DATA_SZ08_ACC	8 bytes data size								
<b>DESCRIPTION</b>	<p><b>ddi_device_copy()</b> copies <i>bytecount</i> bytes from the source address, <i>src_addr</i>, to the destination address, <i>dest_addr</i>. The attributes encoded in the access handles, <i>src_handle</i> and <i>dest_handle</i>, govern how data is actually copied from the source to the destination. Only matching data sizes between the source and destination are supported.</p> <p>Data will automatically be translated to maintain a consistent view between the source and the destination. The translation may involve byte-swapping if the source and the destination devices have incompatible endian characteristics.</p> <p>The <i>src_advcnt</i> and <i>dest_advcnt</i> arguments specifies the number of <i>dev_datsz</i> units to advance with each access to the device addresses. A value of 0 will</p>								

use the same source and destination device address on every access. A positive value increments the corresponding device address by certain number of data size units in the next access. On the other hand, a negative value decrements the device address.

The *dev\_datsz* argument determines the size of the data word on each access. The data size must be the same between the source and destination.

**RETURN VALUES**

**ddi\_device\_copy()** returns:

DDI_SUCCESS	Successfully transferred the data.
DDI_FAILURE	The byte count is not a multiple <i>dev_datsz</i> .

**CONTEXT**

**ddi\_device\_copy()** can be called from user, kernel, or interrupt context.

**SEE ALSO**

**ddi\_regs\_map\_free(9F)**, **ddi\_regs\_map\_setup(9F)**

*Writing Device Drivers*

<b>NAME</b>	ddi_device_zero – zero fill the device								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_device_zero(ddi_acc_handle_t handle, caddr_t dev_addr, size_t bytecount,     ssize_t dev_advcnt, uint_t dev_datsz);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).								
<b>PARAMETERS</b>	<p><b>handle</b>           The data access handle returned from setup calls, such as <code>ddi_regs_map_setup(9F)</code>.</p> <p><b>dev_addr</b>         Beginning of the device address.</p> <p><b>bytecount</b>       Number of bytes to zero.</p> <p><b>dev_advcnt</b>       Number of <code>dev_datsz</code> units to advance on every access.</p> <p><b>dev_datsz</b>        The size of each data word. Possible values are defined as:</p> <table border="0" style="margin-left: 40px;"> <tr> <td>DDI_DATA_SZ01_ACC</td> <td>1 byte data size</td> </tr> <tr> <td>DDI_DATA_SZ02_ACC</td> <td>2 bytes data size</td> </tr> <tr> <td>DDI_DATA_SZ04_ACC</td> <td>4 bytes data size</td> </tr> <tr> <td>DDI_DATA_SZ08_ACC</td> <td>8 bytes data size</td> </tr> </table>	DDI_DATA_SZ01_ACC	1 byte data size	DDI_DATA_SZ02_ACC	2 bytes data size	DDI_DATA_SZ04_ACC	4 bytes data size	DDI_DATA_SZ08_ACC	8 bytes data size
DDI_DATA_SZ01_ACC	1 byte data size								
DDI_DATA_SZ02_ACC	2 bytes data size								
DDI_DATA_SZ04_ACC	4 bytes data size								
DDI_DATA_SZ08_ACC	8 bytes data size								
<b>DESCRIPTION</b>	<p><b>ddi_device_zero()</b> function fills the given, <code>bytecount</code>, number of byte of zeroes to the device register or memory.</p> <p>The <code>dev_advcnt</code> argument determines the value of the device address, <code>dev_addr</code>, on each access. A value of 0 will use the same device address, <code>dev_addr</code>, on every access. A positive value increments the device address in the next access while a negative value decrements the address. The device address is incremented and decremented in <code>dev_datsz</code> units.</p> <p>The <code>dev_datsz</code> argument determines the size of data word on each access.</p>								
<b>RETURN VALUES</b>	<p><b>ddi_device_zero()</b> returns:</p> <p>DDI_SUCCESS    Successfully zeroed the data.</p> <p>DDI_FAILURE    The byte count is not a multiple of <code>dev_datsz</code>.</p>								

**CONTEXT** | **ddi\_device\_zero()** can be called from user, kernel, or interrupt context.

**SEE ALSO** | **ddi\_regs\_map\_free(9F)**, **ddi\_regs\_map\_setup(9F)**

*Writing Device Drivers*

<b>NAME</b>	ddi_devid_compare, ddi_devid_free, ddi_devid_init, ddi_devid_register, ddi_devid_sizeof, ddi_devid_unregister, ddi_devid_valid – Kernel interfaces for device ids	
<b>SYNOPSIS</b>	<pre>int ddi_devid_compare(ddi_devid_t devid1, ddi_devid_t devid2);  size_t ddi_devid_sizeof(ddi_devid_t devid);  int ddi_devid_init(dev_info_t * dip, ushort_t devid_type, ushort_t nbytes, void * id, ddi_devid_t * retdevid);  void ddi_devid_free(ddi_devid_t devid);  int ddi_devid_register(dev_info_t * dip, ddi_devid_t devid);  void ddi_devid_unregister(dev_info_t * dip);  int ddi_devid_valid(ddi_devid_t devid);</pre>	
<b>PARAMETERS</b>	<b>devid</b>	The device id address.
	<b>devid1</b>	The first of two device id addresses to be compared calling <b>ddi_devid_compare()</b> .
	<b>devid2</b>	The second of two device id addresses to be compared calling <b>ddi_devid_compare()</b> .
	<b>dip</b>	A dev_info pointer, which identifies the device.
	<b>devid_type</b>	The following device id types may be accepted by the <b>ddi_devid_init()</b> function:
	DEVID_SCSI3_WWN	World Wide Name associated with SCSI-3 devices.
	DEVID_SCSI_SERIAL	Vendor ID and serial number associated with a SCSI device. Note: This may only be used if known to be unique; otherwise a fabricated device id must be used.
	DEVID_ENCAP	Device ID of another device. This is for layered device driver usage.
	DEVID_FAB	Fabricated device ID .
	<b>nbytes</b>	The length in bytes of device ID .

**DESCRIPTION**

**retdevid** The return address of the device ID created by **ddi\_devid\_init()** .

The following routines are used to provide unique identifiers, device ID s, for devices. Specifically, kernel modules use these interfaces to identify and locate devices, independent of the device's physical connection or its logical device name or number.

**ddi\_devid\_compare()** compares two device ID s byte-by-byte and determines both equality and sort order.

**ddi\_devid\_sizeof()** returns the number of bytes allocated for the passed in device ID ( *devid* ).

**ddi\_devid\_init()** allocates memory and initializes the opaque device ID structure. This function does not store the *devid* . If the device id is not derived from the device's firmware, it is the driver's responsibility to store the *devid* on some reliable store. When a *devid\_type* of either `DEVID_SCSI3_WWN` , `DEVID_SCSI_SERIAL` , or `DEVID_ENCAP` is accepted, an array of bytes ( *id* ) must be passed in ( *nbytes* ).

When the *devid\_type* `DEVID_FAB` is used, the array of bytes ( *id* ) must be NULL and the length ( *nbytes* ) must be zero. The fabricated device ids, `DEVID_FAB` will be initialized with the machine's host id and a timestamp.

Drivers must free the memory allocated by this function, using the **ddi\_devid\_free()** function.

**ddi\_devid\_free()** frees the memory allocated by the **ddi\_devid\_init()** function.

**ddi\_devid\_register()** registers the device ID address ( *devid* ) with the DDI framework, associating it with the `dev_info` passed in ( *dip* ). The drivers must register device ID s at attach time. See **attach(9E)** .

**ddi\_devid\_unregister()** removes the device ID address from the `dev_info` passed in ( *dip* ). Drivers must use this function to unregister the device ID when devices are being detached. This function does not free the space allocated for the device ID . The driver must free the space allocated for the device ID , using the **ddi\_devid\_free()** function. See **detach(9E)** .

**ddi\_devid\_valid()** validates the device ID ( *devid* ) passed in. The driver must use this function to validate any fabricated device ID that has been stored on a device.

**RETURN VALUES**

**ddi\_devid\_init()** returns the following values:  
`DDI_SUCCESS` Success.

DDI\_FAILURE Out of memory. An invalid *devid\_type* was passed in.

**ddi\_devid\_valid()** returns the following values:

DDI\_SUCCESS Valid device ID .

DDI\_FAILURE Invalid device ID .

**ddi\_devid\_register()** returns the following values:

DDI\_SUCCESS Success.

DDI\_FAILURE Failure. The device ID is already registered or the device ID is invalid.

**ddi\_devid\_valid()** returns the following values:

DDI\_SUCCESS Valid device ID .

DDI\_FAILURE Invalid device ID .

**ddi\_devid\_compare()** returns the following values:

-1 The device ID pointed to by *devid1* is less than the device ID pointed to by *devid2* .

0 The device ID pointed to by *devid1* is equal to the device ID pointed to by *devid2* .

1 The device ID pointed to by *devid1* is greater than the device ID pointed to by *devid2* .

**ddi\_devid\_sizeof()** returns the size of the *devid* in numbers of bytes.

#### CONTEXT

These functions can be called from a user context only.

#### ATTRIBUTES

See **attributes(5)** for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

#### SEE ALSO

**devid\_compare(3)**, **devid\_deviceid\_to\_nmlist(3)**, **devid\_free(3)**, **devid\_free\_nmlist(3)**, **devid\_get(3)**, **devid\_get\_minor\_name(3)**, **devid\_sizeof(3)**, **libdevid(4)**, **attributes(5)**, **attach(9E)**, **detach(9E)**

*Writing Device Drivers*

<b>NAME</b>	ddi_dev_is_needed – inform the system that a device’s component is required
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dev_is_needed(dev_info_t *dip, int component, int level);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                    A pointer to the device’s dev_info structure.</p> <p><b>component</b>            The component of the driver which is needed.</p> <p><b>level</b>                    The power level at which the component is needed.</p>
<b>DESCRIPTION</b>	<p>The <b>ddi_dev_is_needed()</b> function informs the system that a device component is needed at the specified power level. The <i>level</i> argument must be non-zero.</p> <p>This function sets a component to the required level and sets all of the devices on which it depends (see <b>pm(7D)</b>) to their normal power levels. If component 0 of the device is at power level 0, the <b>ddi_dev_is_needed()</b> call will result in component 0 being returned to normal power and the device being resumed via <b>attach(9E)</b> before <b>di_dev_is_needed()</b> returns.</p> <p>The state of the device should be examined before each physical access. The <b>ddi_dev_is_needed()</b> function should be called to set a component to the required power level if the operation to be performed requires the component to be at a power level other than its current level.</p> <p>The <b>ddi_dev_is_needed()</b> may cause re-entry of the driver. Deadlock may result if driver locks are held across the call to <b>ddi_dev_is_needed()</b>.</p>
<b>RETURN VALUES</b>	<p>The <b>ddi_dev_is_needed()</b> function returns:</p> <p>DDI_SUCCESS    Power successfully set to the requested level.</p> <p>DDI_FAILURE    An error occurred.</p>
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b>    disk driver code</p> <p>A hypothetical disk driver might include this code:</p> <pre>static int xxdisk_spun_down(struct xxstate *xsp) {     return (xsp-&gt;power_level[DISK_COMPONENT] &lt; POWER_SPUN_UP); }</pre>

```

static int
xxdisk_strategy(struct buf *bp)
{
    ...

    mutex_enter(&xxstate_lock);
    /*
     * Since we have to drop the mutex, we have to do this in a loop
     * in case we get preempted and the device gets taken away from
     * us again
     */
    while (device_spun_down(sp)) {
        mutex_exit(&xxstate_lock);
        if (ddi_dev_is_needed(xsp->mydip,
            XXDISK_COMPONENT, XXPOWER_SPUN_UP) != DDI_SUCCESS) {
            bioerror(bp, EIO);
            biodone(bp);
            return (0);
        }
        mutex_enter(&xxstate_lock);
    }
    xsp->device_busy++;
    mutex_exit(&xxstate_lock);

    ...
}

```

**CONTEXT** This function can be called from user or kernel context.

**SEE ALSO** [pm\(7D\)](#), [attach\(9E\)](#), [detach\(9E\)](#), [power\(9E\)](#), [pm\\_busy\\_component\(9F\)](#), [pm\\_create\\_components\(9F\)](#), [pm\\_destroy\\_components\(9F\)](#), [pm\\_idle\\_component\(9F\)](#)

*Writing Device Drivers*

<b>NAME</b>	ddi_dev_is_sid – tell whether a device is self-identifying
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dev_is_sid(dev_info_t *dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>dip</b> A pointer to the device's dev_info structure.
<b>DESCRIPTION</b>	<b>ddi_dev_is_sid()</b> tells the caller whether the device described by <i>dip</i> is self-identifying, that is, a device that can unequivocally tell the system that it exists. This is useful for drivers that support both a self-identifying as well as a non-self-identifying variants of a device (and therefore must be probed).
<b>RETURN VALUES</b>	<p>DDI_SUCCESS    Device is self-identifying.</p> <p>DDI_FAILURE    Device is not self-identifying.</p>
<b>CONTEXT</b>	<b>ddi_dev_is_sid()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b></p> <pre>1  ... 2  int 3  bz_probe(dev_info_t *dip) 4  { 5  ... 6  if (ddi_dev_is_sid(dip) == DDI_SUCCESS) { 7  /* 8   * This is the self-identifying version (OpenBoot). 9   * No need to probe for it because we know it is there. 10 * The existence of dip &amp;&amp; ddi_dev_is_sid() proves this. 11 */ 12     return (DDI_PROBE_DONTCARE); 13 } 14 /* 15 * Not a self-identifying variant of the device. Now we have to 16 * do some work to see whether it is really attached to the 17 * system. 18 */ 19 ...</pre>

**SEE ALSO** | `probe(9E)` *Writing Device Drivers*

<b>NAME</b>	ddi_dev_nintrs – return the number of interrupt specifications a device has
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dev_nintrs(dev_info_t *dip, int *resultp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_dev_nintrs()</b> returns the number of interrupt specifications a device has in <i>*resultp</i> .
<b>RETURN VALUES</b>	<b>ddi_dev_nintrs()</b> returns: DDI_SUCCESS    A successful return. The number of interrupt specifications that the device has is set in <i>resultp</i> . DDI_FAILURE    The device has no interrupt specifications.
<b>CONTEXT</b>	<b>ddi_dev_nintrs()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>isa(4)</i> , <i>sbus(4)</i> , <i>vme(4)</i> , <i>ddi_add_intr(9F)</i> , <i>ddi_dev_nregs(9F)</i> , <i>ddi_dev_regsizes(9F)</i> <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_dev_nregs – return the number of register sets a device has
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_dev_nregs(dev_info_t *dip, int *resultp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                    A pointer to the device's dev_info structure.</p> <p><b>resultp</b>                Pointer to an integer that holds the number of register sets on return.</p>
<b>DESCRIPTION</b>	The function <b>ddi_dev_nregs()</b> returns the number of sets of registers the device has.
<b>RETURN VALUES</b>	<p><b>ddi_dev_nregs()</b> returns:</p> <p><b>DDI_SUCCESS</b>    A successful return. The number of register sets is returned in <i>resultp</i>.</p> <p><b>DDI_FAILURE</b>    The device has no registers.</p>
<b>CONTEXT</b>	<b>ddi_dev_nregs()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<p>ddi_dev_nintrs(9F), ddi_dev_regsize(9F)</p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_dev_regsize – return the size of a device’s register
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_dev_regsize(dev_info_t *dip, uint_t rnumber, off_t *resultp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                   A pointer to the device’s dev_info structure.</p> <p><b>rnumber</b>               The ordinal register number. Device registers are associated with a dev_info and are enumerated in arbitrary sets from 0 on up. The number of registers a device has can be determined from a call to ddi_dev_nregs(9F).</p> <p><b>resultp</b>               Pointer to an integer that holds the size, in bytes, of the described register (if it exists).</p>
<b>DESCRIPTION</b>	<b>ddi_dev_regsize()</b> returns the size, in bytes, of the device register specified by <i>dip</i> and <i>rnumber</i> . This is useful when, for example, one of the registers is a frame buffer with a varying size known only to its proms.
<b>RETURN VALUES</b>	<p><b>ddi_dev_regsize()</b> returns:</p> <p>DDI_SUCCESS    A successful return. The size, in bytes, of the specified register, is set in <i>resultp</i>.</p> <p>DDI_FAILURE    An invalid (nonexistent) register number was specified.</p>
<b>CONTEXT</b>	<b>ddi_dev_regsize()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	ddi_dev_nintrs(9F), ddi_dev_nregs(9F) <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_dma_addr_bind_handle – binds an address to a DMA handle	
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_addr_bind_handle(ddi_dma_handle_t handle, struct as *as, caddr_t addr, size_t len, uint_t flags, int (*callback) (caddr_t) , caddr_t arg, ddi_dma_cookie_t *cookiep, uint_t *ccountp);</pre>	
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).	
<b>PARAMETERS</b>	<b>handle</b>	The DMA handle previously allocated by a call to <b>ddi_dma_alloc_handle(9F)</b> .
	<b>as</b>	A pointer to an address space structure. This parameter should be set to <b>NULL</b> , which implies kernel address space.
	<b>addr</b>	Virtual address of the memory object.
	<b>len</b>	Length of the memory object in bytes.
	<b>flags</b>	Valid flags include:
	DDI_DMA_WRITE	Transfer direction is from memory to I/O.
	DDI_DMA_READ	Transfer direction is from I/O to memory.
	DDI_DMA_RDWR	Both read and write.
	DDI_DMA_REDZONE	Establish an MMU redzone at end of the object.
	DDI_DMA_PARTIAL	Partial resource allocation.
	DDI_DMA_CONSISTENT	Nonsequential, random, and small block transfers.
	<b>callback</b>	The address of a function to call back later if resources are not currently available. The following special function addresses may also be used.

	DDI_DMA_SLEEP	Wait until resources are available.
	DDI_DMA_DONTWAIT	Do not wait until resources are available and do not schedule a callback.
<b>arg</b>	Argument to be passed to the callback function, <i>callback</i> , if such a function is specified.	
<b>cookiep</b>	A pointer to the first <code>ddi_dma_cookie(9S)</code> structure.	
<b>ccountp</b>	Upon a successful return, <i>ccountp</i> points to a value representing the number of cookies for this DMA object.	

**DESCRIPTION**

**ddi\_dma\_addr\_bind\_handle()** allocates DMA resources for a memory object such that a device can perform DMA to or from the object. DMA resources are allocated considering the device's DMA attributes as expressed by `ddi_dma_attr(9S)` (see `ddi_dma_alloc_handle(9F)`).

**ddi\_dma\_addr\_bind\_handle()** fills in the first DMA cookie pointed to by *cookiep* with the appropriate address, length, and bus type. *ccountp* is set to the number of DMA cookies representing this DMA object. Subsequent DMA cookies must be retrieved by calling `ddi_dma_nextcookie(9F)` the number of times specified by *\*countp-1*.

When a DMA transfer completes, the driver frees up system DMA resources by calling `ddi_dma_unbind_handle(9F)`.

The *flags* argument contains information for mapping routines.

DDI\_DMA\_WRITE, DDI\_DMA\_READ, DDI\_DMA\_RDWR

These flags describe the intended direction of the DMA transfer.

DDI\_DMA\_STREAMING

This flag should be set if the device is doing sequential, unidirectional, block-sized, and block-aligned transfers to or from memory. The alignment and padding constraints specified by the *minxfer* and *burstsizes* fields in the DMA attribute structure, `ddi_dma_attr(9S)` (see `ddi_dma_alloc_handle(9F)`) is used to allocate the most effective hardware support for large transfers.

DDI\_DMA\_CONSISTENT

This flag should be set if the device accesses memory randomly, or if synchronization steps using `ddi_dma_sync(9F)` need to be as efficient as possible. I/O parameter blocks used for communication between a device and a driver should be allocated using `DDI_DMA_CONSISTENT`.

`DDI_DMA_REDZONE`

If this flag is set, the system attempts to establish a protected red zone after the object. The DMA resource allocation functions do not guarantee the success of this request as some implementations may not have the hardware ability to support a red zone.

`DDI_DMA_PARTIAL`

Setting this flag indicates the caller can accept resources for part of the object. That is, if the size of the object exceeds the resources available, only resources for a portion of the object are allocated. The system indicates this condition by returning status `DDI_DMA_PARTIAL_MAP`. At a later point, the caller can use `ddi_dma_getwin(9F)` to change the valid portion of the object for which resources are allocated. If resources were allocated for only part of the object, `ddi_dma_addr_bind_handle()` returns resources for the first DMA window. Even when `DDI_DMA_PARTIAL` is set, the system may decide to allocate resources for the entire object (less overhead) in which case `DDI_DMA_MAPPED` is returned.

The callback function *callback* indicates how a caller wants to handle the possibility of resources not being available. If *callback* is set to `DDI_DMA_DONTWAIT`, the caller does not care if the allocation fails, and can handle an allocation failure appropriately. If *callback* is set to `DDI_DMA_SLEEP`, the caller wishes to have the allocation routines wait for resources to become available. If any other value is set and a DMA resource allocation fails, this value is assumed to be the address of a function to be called when resources become available. When the specified function is called, *arg* is passed to it as an argument. The specified callback function must return either `DDI_DMA_CALLBACK_RUNOUT` or `DDI_DMA_CALLBACK_DONE`. `DDI_DMA_CALLBACK_RUNOUT` indicates that the callback function attempted to allocate DMA resources but failed. In this case, the callback function is put back on a list to be called again later. `DDI_DMA_CALLBACK_DONE` indicates that either the allocation of DMA resources was successful or the driver no longer wishes to retry.

The callback function is called in interrupt context. Therefore, only system functions accessible from interrupt context are available. The callback function must take whatever steps are necessary to protect its critical resources, data structures, queues, and so on.

**RETURN VALUES****ddi\_dma\_addr\_bind\_handle()** returns:

DDI_DMA_MAPPED	Successfully allocated resources for the entire object.
DDI_DMA_PARTIAL_MAP	Successfully allocated resources for a part of the object. This is acceptable when partial transfers are permitted by setting the DDI_DMA_PARTIAL flag in <i>flags</i> .
DDI_DMA_INUSE	Another I/O transaction is using the DMA handle.
DDI_DMA_NORESOURCES	No resources are available at the present time.
DDI_DMA_NOMAPPING	The object cannot be reached by the device requesting the resources.
DDI_DMA_TOOBIG	The object is too big. A request of this size can never be satisfied on this particular system. The maximum size varies depending on machine and configuration.

**CONTEXT**

**ddi\_dma\_addr\_bind\_handle()** can be called from user, kernel, or interrupt context, except when *callback* is set to DDI\_DMA\_SLEEP, in which case it can only be called from user or kernel context.

**SEE ALSO**

**ddi\_dma\_alloc\_handle(9F)**, **ddi\_dma\_free\_handle(9F)**,  
**ddi\_dma\_getwin(9F)**, **ddi\_dma\_mem\_alloc(9F)**, **ddi\_dma\_mem\_free(9F)**,  
**ddi\_dma\_nextcookie(9F)**, **ddi\_dma\_sync(9F)**,  
**ddi\_dma\_unbind\_handle(9F)**, **ddi\_dma\_attr(9S)**, **ddi\_dma\_cookie(9S)**

*Writing Device Drivers*

**NOTES**

If the driver permits partial mapping with the DDI\_DMA\_PARTIAL flag, the number of cookies in each window may exceed the size of the device's scatter/gather list as specified in the *dma\_attr\_sgllen* field in the **ddi\_dma\_attr(9S)** structure. In this case, each set of cookies comprising a DMA window will satisfy the DMA attributes as described in the **ddi\_dma\_attr(9S)** structure in all aspects. The driver should set up its DMA engine and perform one transfer for each set of cookies sufficient for its scatter/gather list, up to the number of cookies for this window, before advancing to the next window using **ddi\_dma\_getwin(9F)**.

<b>NAME</b>	ddi_dma_addr_setup – easier DMA setup for use with virtual addresses
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_addr_setup(dev_info_t *dip, struct as *as, caddr_t addr, size_t len, uint_t flags, int (*waitfp) (caddr_t), caddr_t arg, ddi_dma_lim_t *lim, ddi_dma_handle_t *handlep);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                   A pointer to the device's dev_info structure.</p> <p><b>as</b>                     A pointer to an address space structure. Should be set to NULL, which implies kernel address space.</p> <p><b>addr</b>                  Virtual address of the memory object.</p> <p><b>len</b>                   Length of the memory object in bytes.</p> <p><b>flags</b>                 Flags that would go into the ddi_dma_req structure (see ddi_dma_req(9S)).</p> <p><b>waitfp</b>                The address of a function to call back later if resources aren't available now. The special function addresses DDI_DMA_SLEEP and DDI_DMA_DONTWAIT (see ddi_dma_req(9S)) are taken to mean, respectively, wait until resources are available or, do not wait at all and do not schedule a callback.</p> <p><b>arg</b>                   Argument to be passed to a callback function, if such a function is specified.</p> <p><b>lim</b>                   A pointer to a DMA limits structure for this device (see ddi_dma_lim_sparc(9S) or ddi_dma_lim_x86(9S)). If this pointer is NULL, a default set of DMA limits is assumed.</p> <p><b>handlep</b>               Pointer to a DMA handle. See ddi_dma_setup(9F) for a discussion of handle.</p>
<b>DESCRIPTION</b>	ddi_dma_addr_setup() is an interface to ddi_dma_setup(9F). It uses its arguments to construct an appropriate ddi_dma_req structure and calls ddi_dma_setup(9F) with it.
<b>RETURN VALUES</b>	See ddi_dma_setup(9F) for the possible return values for this function.

**CONTEXT** | **ddi\_dma\_addr\_setup()** can be called from user or interrupt context, except when *waitfp* is set to `DDI_DMA_SLEEP`, in which case it can be called from user context only.

**SEE ALSO** | `ddi_dma_buf_setup(9F)`, `ddi_dma_free(9F)`, `ddi_dma_htoc(9F)`,  
`ddi_dma_setup(9F)`, `ddi_dma_sync(9F)`, `ddi_iopb_alloc(9F)`,  
`ddi_dma_lim_sparc(9S)`, `ddi_dma_lim_x86(9S)`, `ddi_dma_req(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_alloc_handle – allocate DMA handle				
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_dma_alloc_handle(dev_info_t *dip, ddi_dma_attr_t *attr, int (*callback) (caddr_t), caddr_t arg, ddi_dma_handle_t *handlep);</pre>				
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).				
<b>PARAMETERS</b>	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>attr</b> Pointer to a DMA attribute structure for this device (see ddi_dma_attr(9S)).</p> <p><b>callback</b> The address of a function to call back later if resources aren't available now. The following special function addresses may also be used.</p> <table border="0" style="margin-left: 40px;"> <tr> <td>DDI_DMA_SLEEP</td> <td>Wait until resources are available.</td> </tr> <tr> <td>DDI_DMA_DONTWAIT</td> <td>Do not wait until resources are available and do not schedule a callback.</td> </tr> </table> <p><b>arg</b> Argument to be passed to a callback function, if such a function is specified.</p> <p><b>handlep</b> Pointer to the DMA handle to be initialized.</p>	DDI_DMA_SLEEP	Wait until resources are available.	DDI_DMA_DONTWAIT	Do not wait until resources are available and do not schedule a callback.
DDI_DMA_SLEEP	Wait until resources are available.				
DDI_DMA_DONTWAIT	Do not wait until resources are available and do not schedule a callback.				
<b>DESCRIPTION</b>	<p><b>ddi_dma_alloc_handle()</b> allocates a new DMA handle. A DMA handle is an opaque object used as a reference to subsequently allocated DMA resources. <b>ddi_dma_alloc_handle()</b> accepts as parameters the device information referred to by <i>dip</i> and the device's DMA attributes described by a <b>ddi_dma_attr(9S)</b> structure. A successful call to <b>ddi_dma_alloc_handle()</b> fills in the value pointed to by <i>handlep</i>. A DMA handle must only be used by the device for which it was allocated and is only valid for one I/O transaction at a time.</p> <p>The callback function, <i>callback</i>, indicates how a caller wants to handle the possibility of resources not being available. If <i>callback</i> is set to <b>DDI_DMA_DONTWAIT</b>, then the caller does not care if the allocation fails, and can handle an allocation failure appropriately. If <i>callback</i> is set to <b>DDI_DMA_SLEEP</b>, then the caller wishes to have the the allocation routines wait for resources to become available. If any other value is set, and a DMA</p>				

resource allocation fails, this value is assumed to be a function to call at a later time when resources may become available. When the specified function is called, it is passed *arg* as an argument. The specified callback function must return either `DDI_DMA_CALLBACK_RUNOUT` or `DDI_DMA_CALLBACK_DONE`. `DDI_DMA_CALLBACK_RUNOUT` indicates that the callback routine attempted to allocate DMA resources but failed to do so, in which case the callback function is put back on a list to be called again later. `DDI_DMA_CALLBACK_DONE` indicates either success at allocating DMA resources or the driver no longer wishes to retry.

The callback function is called in interrupt context. Therefore, only system functions that are accessible from interrupt context is available. The callback function must take whatever steps necessary to protect its critical resources, data structures, queues, and so forth.

When a DMA handle is no longer needed, `ddi_dma_free_handle(9F)` must be called to free the handle.

**RETURN VALUES**

**ddi\_dma\_alloc\_handle()** returns:

<code>DDI_SUCCESS</code>	Successfully allocated a new DMA handle.
<code>DDI_DMA_BADATTR</code>	The attributes specified in the <code>ddi_dma_attr(9S)</code> structure make it impossible for the system to allocate potential DMA resources.
<code>DDI_DMA_NORESOURCES</code>	No resources are available.

**CONTEXT**

**ddi\_dma\_alloc\_handle()** can be called from user, kernel, or interrupt context, except when *callback* is set to `DDI_DMA_SLEEP`, in which case it can be called from user or kernel context only.

**SEE ALSO**

`ddi_dma_addr_bind_handle(9F)`, `ddi_dma_buf_bind_handle(9F)`, `ddi_dma_burstsizes(9F)`, `ddi_dma_free_handle(9F)`, `ddi_dma_unbind_handle(9F)`, `ddi_dma_attr(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_buf_bind_handle – binds a system buffer to a DMA handle	
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int ddi_dma_buf_bind_handle(ddi_dma_handle_t handle, struct buf *bp, uint_t flags, int (*callback)(caddr_t), caddr_t arg, ddi_dma_cookie_t *cookiep, uint_t *ccountp);</p>	
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).	
<b>PARAMETERS</b>	<b>handle</b>	The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).
	<b>bp</b>	A pointer to a system buffer structure (see buf(9S)).
	<b>flags</b>	Valid flags include:
	DDI_DMA_WRITE	Transfer direction is from memory to I/O
	DDI_DMA_READ	Transfer direction is from I/O to memory
	DDI_DMA_RDWR	Both read and write
	DDI_DMA_REDZONE	Establish an MMU redzone at end of the object.
	DDI_DMA_PARTIAL	Partial resource allocation
	DDI_DMA_CONSISTENT	Nonsequential, random, and small block transfers.
	DDI_DMA_STREAMING	Sequential, unidirectional, block-sized, and block-aligned transfers.
	<b>callback</b>	The address of a function to call back later if resources are not available now. The following special function addresses may also be used.
	DDI_DMA_SLEEP	Wait until resources are available.
	DDI_DMA_DONTWAIT	Do not wait until resources are available and do not schedule a callback.

**arg** Argument to be passed to the callback function, *callback*, if such a function is specified.

**cookiep** A pointer to the first `ddi_dma_cookie(9S)` structure.

**ccountp** Upon a successful return, *ccountp* points to a value representing the number of cookies for this DMA object.

**DESCRIPTION**

`ddi_dma_buf_bind_handle()` allocates DMA resources for a system buffer such that a device can perform DMA to or from the buffer. DMA resources are allocated considering the device's DMA to attributes as expressed by `ddi_dma_attr(9S)` (see `ddi_dma_alloc_handle(9F)`).

`ddi_dma_buf_bind_handle()` fills in the first DMA cookie pointed to by *cookiep* with the appropriate address, length, and bus type. *\*ccountp* is set to the number of DMA cookies representing this DMA object. Subsequent DMA cookies must be retrieved by calling `ddi_dma_nextcookie(9F)` *\*ccountp*-1 times.

When a DMA transfer completes, the driver should free up system DMA to resources by calling `ddi_dma_unbind_handle(9F)`.

The *flags* argument contains information for mapping routines.

DDI\_DMA\_WRITE, DDI\_DMA\_READ, DDI\_DMA\_RDWR

These flags describe the intended direction of the DMA transfer.

DDI\_DMA\_STREAMING

This flag should be set if the device is doing sequential, unidirectional, block-sized, and block-aligned transfers to or from memory. The alignment and padding constraints specified by the *minxfer* and *burstsizes* fields in the DMA attribute structure, `ddi_dma_attr(9S)` (see `ddi_dma_alloc_handle(9F)`) is used to allocate the most effective hardware support for large transfers.

DDI\_DMA\_CONSISTENT

This flag should be set if the device accesses memory randomly, or if synchronization steps using `ddi_dma_sync(9F)` need to be as efficient as possible. I/O parameter blocks used for communication between a device and a driver should be allocated using `DDI_DMA_CONSISTENT`.

DDI\_DMA\_REDZONE

If this flag is set, the system attempts to establish a protected red zone after the object. The DMA resource allocation functions do not guarantee the success of this request as some implementations may not have the hardware ability to support a red zone.

#### DDI\_DMA\_PARTIAL

Setting this flag indicates the caller can accept resources for part of the object. That is, if the size of the object exceeds the resources available, only resources for a portion of the object are allocated. The system indicates this condition returning status `DDI_DMA_PARTIAL_MAP`. At a later point, the caller can use `ddi_dma_getwin(9F)` to change the valid portion of the object for which resources are allocated. If resources were allocated for only part of the object, `ddi_dma_addr_bind_handle()` returns resources for the first DMA window. Even when `DDI_DMA_PARTIAL` is set, the system may decide to allocate resources for the entire object (less overhead) in which case `DDI_DMA_MAPPED` is returned.

The callback function, *callback*, indicates how a caller wants to handle the possibility of resources not being available. If *callback* is set to `DDI_DMA_DONTWAIT`, the caller does not care if the allocation fails, and can handle an allocation failure appropriately. If *callback* is set to `DDI_DMA_SLEEP`, the caller wishes to have the allocation routines wait for resources to become available. If any other value is set, and a DMA resource allocation fails, this value is assumed to be the address of a function to call at a later time when resources may become available. When the specified function is called, it is passed *arg* as an argument. The specified callback function must return either `DDI_DMA_CALLBACK_RUNOUT` or `DDI_DMA_CALLBACK_DONE`. `DDI_DMA_CALLBACK_RUNOUT` indicates that the callback function attempted to allocate DMA resources but failed to do so. In this case the callback function is put back on a list to be called again later. `DDI_DMA_CALLBACK_DONE` indicates either a successful allocation of DMA resources or that the driver no longer wishes to retry.

The callback function is called in interrupt context. Therefore, only system functions accessible from interrupt context are available. The callback function must take whatever steps necessary to protect its critical resources, data structures, queues, etc.

#### RETURN VALUES

`ddi_dma_buf_bind_handle()` returns:

<code>DDI_DMA_MAPPED</code>	Successfully allocated resources for the entire object.
<code>DDI_DMA_PARTIAL_MAP</code>	Successfully allocated resources for a part of the object. This is acceptable when partial transfers are permitted

		by setting the <code>DDI_DMA_PARTIAL</code> flag in <i>flags</i> .
	<code>DDI_DMA_INUSE</code>	Another I/O transaction is using the DMA handle.
	<code>DDI_DMA_NORESOURCES</code>	No resources are available at the present time.
	<code>DDI_DMA_NOMAPPING</code>	The object cannot be reached by the device requesting the resources.
	<code>DDI_DMA_TOOBIG</code>	The object is too big. A request of this size can never be satisfied on this particular system. The maximum size varies depending on machine and configuration.
<b>CONTEXT</b>	<b>ddi_dma_buf_bind_handle()</b> can be called from user, kernel, or interrupt context, except when <i>callback</i> is set to <code>DDI_DMA_SLEEP</code> , in which case it can be called from user or kernel context only.	
<b>SEE ALSO</b>	<code>ddi_dma_addr_bind_handle(9F)</code> , <code>ddi_dma_alloc_handle(9F)</code> , <code>ddi_dma_free_handle(9F)</code> , <code>ddi_dma_getwin(9F)</code> , <code>ddi_dma_nextcookie(9F)</code> , <code>ddi_dma_sync(9F)</code> , <code>ddi_dma_unbind_handle(9F)</code> , <code>buf(9S)</code> , <code>ddi_dma_attr(9S)</code> , <code>ddi_dma_cookie(9S)</code>	
	<i>Writing Device Drivers</i>	
<b>NOTES</b>	If the driver permits partial mapping with the <code>DDI_DMA_PARTIAL</code> flag, the number of cookies in each window may exceed the size of the device's scatter/gather list as specified in the <code>dma_attr_sgllen</code> field in the <code>ddi_dma_attr(9S)</code> structure. In this case, each set of cookies comprising a DMA window will satisfy the DMA attributes as described in the <code>ddi_dma_attr(9S)</code> structure in all aspects. The driver should set up its DMA engine and perform one transfer for each set of cookies sufficient for its scatter/gather list, up to the number of cookies for this window, before advancing to the next window using <code>ddi_dma_getwin(9F)</code> .	

<b>NAME</b>	ddi_dma_buf_setup – easier DMA setup for use with buffer structures
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_buf_setup(dev_info_t *dip, struct buf *bp, uint_t flags, int (*waitfp) (caddr_t), caddr_t arg, ddi_dma_lim_t *lim, ddi_dma_handle_t *handlep);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>            A pointer to the device's dev_info structure.</p> <p><b>bp</b>             A pointer to a system buffer structure (see buf(9S)).</p> <p><b>flags</b>          Flags that go into a ddi_dma_req structure (see ddi_dma_req(9S)).</p> <p><b>waitfp</b>        The address of a function to call back later if resources aren't available now. The special function addresses DDI_DMA_SLEEP and DDI_DMA_DONTWAIT (see ddi_dma_req(9S)) are taken to mean, respectively, wait until resources are available, or do not wait at all and do not schedule a callback.</p> <p><b>arg</b>            Argument to be passed to a callback function, if such a function is specified.</p> <p><b>lim</b>            A pointer to a DMA limits structure for this device (see ddi_dma_lim_sparc(9S) or ddi_dma_lim_x86(9S)). If this pointer is NULL, a default set of DMA limits is assumed.</p> <p><b>handlep</b>       Pointer to a DMA handle. See ddi_dma_setup(9F) for a discussion of handle.</p>
<b>DESCRIPTION</b>	ddi_dma_buf_setup() is an interface to ddi_dma_setup(9F). It uses its arguments to construct an appropriate ddi_dma_req structure and calls ddi_dma_setup() with it.
<b>RETURN VALUES</b>	See ddi_dma_setup(9F) for the possible return values for this function.
<b>CONTEXT</b>	ddi_dma_buf_setup() can be called from user or interrupt context, except when waitfp is set to DDI_DMA_SLEEP, in which case it can be called from user context only.

**SEE ALSO**

`ddi_dma_addr_setup(9F)`, `ddi_dma_free(9F)`, `ddi_dma_htoc(9F)`,  
`ddi_dma_setup(9F)`, `ddi_dma_sync(9F)`, `physio(9F)`, `buf(9S)`,  
`ddi_dma_lim_sparc(9S)`, `ddi_dma_lim_x86(9S)`, `ddi_dma_req(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_burstsizes – find out the allowed burst sizes for a DMA mapping
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_burstsizes(ddi_dma_handle_t handle);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>handle</b> A DMA handle that was filled in by a successful call to <b>ddi_dma_setup(9F)</b> .
<b>DESCRIPTION</b>	<b>ddi_dma_burstsizes()</b> returns the allowed burst sizes for a DMA mapping. This value is derived from the <code>dlim_burstsizes</code> member of the <b>ddi_dma_lim_sparc(9S)</b> structure, but it shows the allowable burstsizes <i>after</i> imposing on it the limitations of other device layers in addition to device's own limitations.
<b>RETURN VALUES</b>	<b>ddi_dma_burstsizes()</b> returns a binary encoded value of the allowable DMA burst sizes. See <b>ddi_dma_lim_sparc(9S)</b> for a discussion of DMA burst sizes.
<b>CONTEXT</b>	This function can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>ddi_dma_devalign(9F)</b> , <b>ddi_dma_setup(9F)</b> , <b>ddi_dma_lim_sparc(9S)</b> , <b>ddi_dma_req(9S)</b>  <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_dma_coff – convert a DMA cookie to an offset within a DMA handle
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_dma_coff(ddi_dma_handle_t handle, ddi_dma_cookie_t *cookiep, off_t *offp);</pre>
<b>INTERFACE LEVEL</b>	Solaris SPARC DDI (Solaris SPARC DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The <i>handle</i> filled in by a call to ddi_dma_setup(9F).</p> <p><b>cookiep</b>           A pointer to a DMA cookie (see ddi_dma_cookie(9S)) that contains the appropriate address, length and bus type to be used in programming the DMA engine.</p> <p><b>offp</b>                A pointer to an offset to be filled in.</p>
<b>DESCRIPTION</b>	<p><b>ddi_dma_coff()</b> converts the values in DMA cookie pointed to by <i>cookiep</i> to an offset (in bytes) from the beginning of the object that the DMA <i>handle</i> has mapped.</p> <p><b>ddi_dma_coff()</b> allows a driver to update a DMA cookie with values it reads from its device's DMA engine after a transfer completes and convert that value into an offset into the object that is mapped for DMA.</p>
<b>RETURN VALUES</b>	<p><b>ddi_dma_coff()</b> returns:</p> <p>DDI_SUCCESS        Successfully filled in <i>offp</i>.</p> <p>DDI_FAILURE        Failed to successfully fill in <i>offp</i>.</p>
<b>CONTEXT</b>	<b>ddi_dma_coff()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	ddi_dma_setup(9F), ddi_dma_sync(9F), ddi_dma_cookie(9S) <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_dma_curwin – report current DMA window offset and size
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int ddi_dma_curwin(ddi_dma_handle_t handle, off_t *offp, uint_t *lenp);</p>
<b>INTERFACE LEVEL</b>	Solaris SPARC DDI specific (Solaris SPARC DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The DMA handle filled in by a call to ddi_dma_setup(9F).</p> <p><b>offp</b>              A pointer to a value which will be filled in with the current offset from the beginning of the object that is mapped for DMA.</p> <p><b>lenp</b>              A pointer to a value which will be filled in with the size, in bytes, of the current window onto the object that is mapped for DMA.</p>
<b>DESCRIPTION</b>	<b>ddi_dma_curwin()</b> reports the current DMA window offset and size. If a DMA mapping allows partial mapping, that is if the DDI_DMA_PARTIAL flag in the ddi_dma_req(9S) structure is set, its current (effective) DMA window offset and size can be obtained by a call to <b>ddi_dma_curwin()</b> .
<b>RETURN VALUES</b>	<p><b>ddi_dma_curwin()</b> returns:</p> <p>DDI_SUCCESS    The current length and offset can be established.</p> <p>DDI_FAILURE    Otherwise.</p>
<b>CONTEXT</b>	<b>ddi_dma_curwin()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	ddi_dma_movwin(9F), ddi_dma_setup(9F), ddi_dma_req(9S) <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_dma_devalign – find DMA mapping alignment and minimum transfer size
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_dma_devalign(ddi_dma_handle_t handle, uint_t *alignment, uint_t *minxfr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The DMA handle filled in by a successful call to <b>ddi_dma_setup(9F)</b>.</p> <p><b>alignment</b>        A pointer to an unsigned integer to be filled in with the minimum required alignment for DMA. The alignment is guaranteed to be a power of two.</p> <p><b>minxfr</b>            A pointer to an unsigned integer to be filled in with the minimum effective transfer size (see <b>ddi_iomin(9F)</b>, <b>ddi_dma_lim_sparc(9S)</b> and <b>ddi_dma_lim_x86(9S)</b>). This also is guaranteed to be a power of two.</p>
<b>DESCRIPTION</b>	<b>ddi_dma_devalign()</b> determines after a successful DMA mapping (see <b>ddi_dma_setup(9F)</b> ) the minimum required data alignment and minimum DMA transfer size.
<b>RETURN VALUES</b>	<p><b>ddi_dma_devalign()</b> returns:</p> <p>DDI_SUCCESS    The <i>alignment</i> and <i>minxfr</i> values have been filled.</p> <p>DDI_FAILURE    The handle was illegal.</p>
<b>CONTEXT</b>	<b>ddi_dma_devalign()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><b>ddi_dma_setup(9F)</b>, <b>ddi_iomin(9F)</b>, <b>ddi_dma_lim_sparc(9S)</b>, <b>ddi_dma_lim_x86(9S)</b>, <b>ddi_dma_req(9S)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_dmae, ddi_dmae_alloc, ddi_dmae_release, ddi_dmae_prog, ddi_dmae_disable, ddi_dmae_enable, ddi_dmae_stop, ddi_dmae_getcnt, ddi_dmae_1stparty, ddi_dmae_getlim, ddi_dmae_getattr – system DMA engine functions				
<b>SYNOPSIS</b>	<pre>int ddi_dmae_alloc(dev_info_t * dip, int chnl, int (* callback)(caddr_t), caddr_t arg);  int ddi_dmae_release(dev_info_t * dip, int chnl);  int ddi_dmae_prog(dev_info_t * dip, struct ddi_dmae_req * dmaereq, ddi_dma_cookie_t * cookiep, int chnl);  int ddi_dmae_disable(dev_info_t * dip, int chnl);  int ddi_dmae_enable(dev_info_t * dip, int chnl);  int ddi_dmae_stop(dev_info_t * dip, int chnl);  int ddi_dmae_getcnt(dev_info_t * dip, int chnl, int * countp);  int ddi_dmae_1stparty(dev_info_t * dip, int chnl);  int ddi_dmae_getlim(dev_info_t * dip, ddi_dma_lim_t * limitsp);  int ddi_dmae_getattr(dev_info_t * dip, ddi_dma_attr_t * attrp);</pre>				
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).				
<b>PARAMETERS</b>	<p><b>dip</b> A dev_info pointer that identifies the device.</p> <p><b>chnl</b> A DMA channel number. On ISA or EISA buses this number must be 0, 1, 2, 3, 5, 6, or 7.</p> <p><b>callback</b> The address of a function to call back later if resources are not currently available. The following special function addresses may also be used:</p> <table border="0" style="margin-left: 2em;"> <tr> <td>DDI_DMA_SLEEP</td> <td>Wait until resources are available.</td> </tr> <tr> <td>DDI_DMA_DONTWAIT</td> <td>Do not wait until resources are available and do not schedule a callback.</td> </tr> </table> <p><b>arg</b> Argument to be passed to the callback function, if specified.</p> <p><b>dmaereq</b> A pointer to a DMA engine request structure. See ddi_dmae_req(9S).</p>	DDI_DMA_SLEEP	Wait until resources are available.	DDI_DMA_DONTWAIT	Do not wait until resources are available and do not schedule a callback.
DDI_DMA_SLEEP	Wait until resources are available.				
DDI_DMA_DONTWAIT	Do not wait until resources are available and do not schedule a callback.				

<b><i>cookiep</i></b>	A pointer to a <code>ddi_dma_cookie(9S)</code> object, obtained from <code>ddi_dma_segtocookie(9F)</code> , which contains the address and count.
<b><i>countp</i></b>	A pointer to an integer that will receive the count of the number of bytes not yet transferred upon completion of a DMA operation.
<b><i>limitsp</i></b>	A pointer to a DMA limit structure. See <code>ddi_dma_lim_x86(9S)</code> .
<b><i>attrp</i></b>	A pointer to a DMA attribute structure. See <code>ddi_dma_attr(9S)</code> .

**DESCRIPTION**

There are three possible ways that a device can perform DMA engine functions:

**Bus master DMA**

If the device is capable of acting as a true bus master, then the driver should program the device's DMA registers directly and not make use of the DMA engine functions described here. The driver should obtain the DMA address and count from `ddi_dma_segtocookie(9F)`. See `ddi_dma_cookie(9S)` for a description of a DMA cookie.

**Third-party DMA**

This method uses the system DMA engine that is resident on the main system board. In this model, the device cooperates with the system's DMA engine to effect the data transfers between the device and memory. The driver uses the functions documented here, except `ddi_dmae_1stparty()`, to initialize and program the DMA engine. For each DMA data transfer, the driver programs the DMA engine and then gives the device a command to initiate the transfer in cooperation with that engine.

**First-party DMA**

Using this method, the device uses its own DMA bus cycles, but requires a channel from the system's DMA engine. After allocating the DMA channel, the `ddi_dmae_1stparty()` function may be used to perform whatever configuration is necessary to enable this mode.

**ddi\_dmae\_alloc()**

The `ddi_dmae_alloc()` function is used to acquire a DMA channel of the system DMA engine. `ddi_dmae_alloc()` allows only one device at a time to have a particular DMA channel allocated. It must be called prior to any other

system DMA engine function on a channel. If the device allows the channel to be shared with other devices, it must be freed using **ddi\_dmae\_release()** after completion of the DMA operation. In any case, the channel must be released before the driver successfully detaches. See **detach(9E)**. No other driver may acquire the DMA channel until it is released.

If the requested channel is not immediately available, the value of *callback* determines what action will be taken. If the value of *callback* is `DDI_DMA_DONTWAIT`, **ddi\_dmae\_alloc()** will return immediately. The value `DDI_DMA_SLEEP` will cause the thread to sleep and not return until the channel has been acquired. Any other value is assumed to be a callback function address. In that case, **ddi\_dmae\_alloc()** returns immediately, and when resources might have become available, the callback function is called (with the argument *arg*) from interrupt context. When the callback function is called, it should attempt to allocate the DMA channel again. If it succeeds or no longer needs the channel, it must return the value `DDI_DMA_CALLBACK_DONE`. If it tries to allocate the channel but fails to do so, it must return the value `DDI_DMA_CALLBACK_RUNOUT`. In this case, the callback function is put back on a list to be called again later.

**ddi\_dmae\_prog()**

The **ddi\_dmae\_prog()** function programs the DMA channel for a DMA transfer. The `ddi_dmae_req` structure contains all the information necessary to set up the channel, except for the memory address and count. Once the channel has been programmed, subsequent calls to **ddi\_dmae\_prog()** may specify a value of `NULL` for *dmaareqp* if no changes to the programming are required other than the address and count values. It disables the channel prior to setup, and enables the channel before returning. The DMA address and count are specified by passing **ddi\_dmae\_prog()** a cookie obtained from **ddi\_dma\_segtocookie(9F)**. Other DMA engine parameters are specified by the DMA engine request structure passed in through *dmaareqp*. The fields of that structure are documented in **ddi\_dmae\_req(9S)**.

Before using **ddi\_dmae\_prog()**, you must allocate system DMA resources using DMA setup functions such as **ddi\_dma\_buf\_setup(9F)**. **ddi\_dma\_segtocookie(9F)** can then be used to retrieve a cookie which contains the address and count. Then this cookie is passed to **ddi\_dmae\_prog()**.

**ddi\_dmae\_disable()**

The **ddi\_dmae\_disable()** function disables the DMA channel so that it no longer responds to a device's DMA service requests.

**ddi\_dmae\_enable()**

The **ddi\_dmae\_enable()** function enables the DMA channel for operation. This may be used to re-enable the channel after a call to **ddi\_dmae\_disable()**. The channel is automatically enabled after successful programming by **ddi\_dmae\_prog()**.

<b>ddi_dmae_stop()</b>	The <b>ddi_dmae_stop()</b> function disables the channel and terminates any active operation.
<b>ddi_dmae_getcnt()</b>	The <b>ddi_dmae_getcnt()</b> function examines the count register of the DMA channel and sets <i>*countp</i> to the number of bytes remaining to be transferred. The channel is assumed to be stopped.
<b>ddi_dmae_1stparty()</b>	In the case of ISA and EISA buses, <b>ddi_dmae_1stparty()</b> configures a channel in the system's DMA engine to operate in a "slave" ("cascade") mode.  When operating in <b>ddi_dmae_1stparty()</b> mode, the DMA channel must first be allocated using <b>ddi_dmae_alloc()</b> and then configured using <b>ddi_dmae_1stparty()</b> . The driver then programs the device to perform the I/O, including the necessary DMA address and count values obtained from <b>ddi_dma_segtocookie(9F)</b> .
<b>ddi_dmae_getlim()</b>	The <b>ddi_dmae_getlim()</b> function fills in the DMA limit structure, pointed to by <i>limitsp</i> , with the DMA limits of the system DMA engine. Drivers for devices that perform their own bus mastering or use first-party DMA must create and initialize their own DMA limit structures; they should not use <b>ddi_dmae_getlim()</b> . The DMA limit structure must be passed to the DMA setup routines so that they will know how to break the DMA request into windows and segments (see <b>ddi_dma_nextseg(9F)</b> and <b>ddi_dma_nextwin(9F)</b> ). If the device has any particular restrictions on transfer size or granularity (such as the size of disk sector), the driver should further restrict the values in the structure members before passing them to the DMA setup routines. The driver must not relax any of the restrictions embodied in the structure after it is filled in by <b>ddi_dmae_getlim()</b> . After calling <b>ddi_dmae_getlim()</b> , a driver must examine, and possibly set, the size of the DMA engine's scatter/gather list to determine whether DMA chaining will be used. See <b>ddi_dma_lim_x86(9S)</b> and <b>ddi_dmae_req(9S)</b> for additional information on scatter/gather DMA.
<b>ddi_dmae_getattr</b>	The <b>ddi_dmae_getattr()</b> function fills in the DMA attribute structure, pointed to by <i>attrp</i> , with the DMA attributes of the system DMA engine. Drivers for devices that perform their own bus mastering or use first-party DMA must create and initialize their own DMA attribute structures; they should not use <b>ddi_dmae_getattr()</b> . The DMA attribute structure must be passed to the DMA resource allocation functions to provide the information necessary to break the DMA request into DMA windows and DMA cookies. See <b>ddi_dma_nextcookie(9F)</b> and <b>ddi_dma_getwin(9F)</b> .

**RETURN VALUES**

DDI\_SUCCESS

Upon success, for all of these routines.

DDI\_FAILURE                    May be returned due to invalid arguments.

DDI\_DMA\_NORESOURCES        May be returned by **ddi\_dmae\_alloc()** if the requested resources are not available and the value of *dmae\_waitfp* is not DDI\_DMA\_SLEEP .

**CONTEXT**

If **ddi\_dmae\_alloc()** is called from interrupt context, then its *dmae\_waitfp* argument and the callback function must not have the value DDI\_DMA\_SLEEP . Otherwise, all these routines may be called from user or interrupt context.

**ATTRIBUTES**

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

**SEE ALSO**

**eisa(4)** , **isa(4)** , **attributes(5)** , **ddi\_dma\_buf\_setup(9F)** , **ddi\_dma\_getwin(9F)** , **ddi\_dma\_nextcookie(9F)** , **ddi\_dma\_nextseg(9F)** , **ddi\_dma\_nextwin(9F)** , **ddi\_dma\_segtocookie(9F)** , **ddi\_dma\_setup(9F)** , **ddi\_dma\_attr(9S)** , **ddi\_dma\_cookie(9S)** , **ddi\_dma\_lim\_x86(9S)** , **ddi\_dma\_req(9S)** , **ddi\_dmae\_req(9S)**

<b>NAME</b>	ddi_dma_free – release system DMA resources
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_free(ddi_dma_handle_t handle);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>handle</b> The handle filled in by a call to ddi_dma_setup(9F).
<b>DESCRIPTION</b>	<b>ddi_dma_free()</b> releases system DMA resources set up by <b>ddi_dma_setup(9F)</b> . When a DMA transfer completes, the driver should free up system DMA resources established by a call to <b>ddi_dma_setup(9F)</b> . This is done by a call to <b>ddi_dma_free()</b> . <b>ddi_dma_free()</b> does an implicit <b>ddi_dma_sync(9F)</b> for you so any further synchronization steps are not necessary.
<b>RETURN VALUES</b>	<p><b>ddi_dma_free()</b> returns:</p> <p>DDI_SUCCESS    Successfully released resources</p> <p>DDI_FAILURE    Failed to free resources</p>
<b>CONTEXT</b>	<b>ddi_dma_free()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><b>ddi_dma_addr_setup(9F)</b>, <b>ddi_dma_buf_setup(9F)</b>, <b>ddi_dma_htoc(9F)</b>, <b>ddi_dma_sync(9F)</b>, <b>ddi_dma_req(9S)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_dma_free_handle – free DMA handle
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_dma_free_handle(ddi_dma_handle_t *handle);</pre>
<b>PARAMETERS</b>	<p><b>handle</b>            A pointer to the DMA handle previously allocated by a call to <code>ddi_dma_alloc_handle(9F)</code>.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p><code>ddi_dma_free_handle()</code> destroys the DMA handle pointed to by <i>handle</i>. Any further references to the DMA handle will have undefined results. Note that <code>ddi_dma_unbind_handle(9F)</code> must be called prior to <code>ddi_dma_free_handle()</code> to free any resources the system may be caching on the handle.</p>
<b>CONTEXT</b>	<code>ddi_dma_free_handle()</code> can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<code>ddi_dma_alloc_handle(9F)</code> , <code>ddi_dma_unbind_handle(9F)</code> <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_dma_getwin – activate a new DMA window
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_getwin(ddi_dma_handle_t handle, uint_t win, off_t *offp, size_t *lenp, ddi_dma_cookie_t *cookiep, uint_t *ccountp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The DMA handle previously allocated by a call to <code>ddi_dma_alloc_handle(9F)</code>.</p> <p><b>win</b>                Number of the window to activate.</p> <p><b>offp</b>             Pointer to an offset. Upon a successful return, <i>offp</i> will contain the new offset indicating the beginning of the window within the object.</p> <p><b>lenp</b>             Upon a successful return, <i>lenp</i> will contain the size, in bytes, of the current window.</p> <p><b>cookiep</b>          A pointer to the first <code>ddi_dma_cookie(9S)</code> structure.</p> <p><b>ccountp</b>          Upon a successful return, <i>ccountp</i> will contain the number of cookies for this DMA window.</p>
<b>DESCRIPTION</b>	<p><b>ddi_dma_getwin()</b> activates a new DMA window. If a DMA resource allocation request returns <code>DDI_DMA_PARTIAL_MAP</code> indicating that resources for less than the entire object were allocated, the current DMA window can be changed by a call to <b>ddi_dma_getwin()</b>.</p> <p>The caller must first determine the number of DMA windows, <i>N</i>, using <code>ddi_dma_numwin(9F)</code>. <b>ddi_dma_getwin()</b> takes a DMA window number from the range <i>[0..N-1]</i> as the parameter <i>win</i> and makes it the current DMA window.</p> <p><b>ddi_dma_getwin()</b> fills in the first DMA cookie pointed to by <i>cookiep</i> with the appropriate address, length, and bus type. <i>*ccountp</i> is set to the number of DMA cookies representing this DMA object. Subsequent DMA cookies must be retrieved using <code>ddi_dma_nextcookie(9F)</code>.</p> <p><b>ddi_dma_getwin()</b> takes care of underlying resource synchronizations required to shift the window. However accessing the data prior to or after moving the window requires further synchronization steps using <code>ddi_dma_sync(9F)</code>.</p>

**ddi\_dma\_getwin()** is normally called from an interrupt routine. The first invocation of the DMA engine is done from the driver. All subsequent invocations of the DMA engine are done from the interrupt routine. The interrupt routine checks to see if the request has been completed. If it has, the interrupt routine returns without invoking another DMA transfer. Otherwise, it calls **ddi\_dma\_getwin()** to shift the current window and start another DMA transfer.

**RETURN VALUES**

**ddi\_dma\_getwin()** returns:

DDI\_SUCCESS Resources for the specified DMA window are allocated.

DDI\_FAILURE *win* is not a valid window index.

**CONTEXT**

**ddi\_dma\_getwin()** can be called from user, kernel, or interrupt context.

**SEE ALSO**

**ddi\_dma\_addr\_bind\_handle(9F)**, **ddi\_dma\_alloc\_handle(9F)**,  
**ddi\_dma\_buf\_bind\_handle(9F)**, **ddi\_dma\_nextcookie(9F)**,  
**ddi\_dma\_numwin(9F)**, **ddi\_dma\_sync(9F)**, **ddi\_dma\_unbind\_handle(9F)**,  
**ddi\_dma\_cookie(9S)**

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_htoc – convert a DMA handle to a DMA address cookie
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_dma_htoc(ddi_dma_handle_t handle, off_t off, ddi_dma_cookie_t *cookiep);</pre>
<b>INTERFACE LEVEL</b>	Solaris SPARC DDI specific (Solaris SPARC DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The handle filled in by a call to <code>ddi_dma_setup(9F)</code>.</p> <p><b>off</b>                An offset into the object that <i>handle</i> maps.</p> <p><b>cookiep</b>           A pointer to a <code>ddi_dma_cookie(9S)</code> structure.</p>
<b>DESCRIPTION</b>	<code>ddi_dma_htoc()</code> takes a DMA handle (established by <code>ddi_dma_setup(9F)</code> ), and fills in the cookie pointed to by <i>cookiep</i> with the appropriate address, length, and bus type to be used to program the DMA engine.
<b>RETURN VALUES</b>	<p><code>ddi_dma_htoc()</code> returns:</p> <p>DDI_SUCCESS        Successfully filled in the cookie pointed to by <i>cookiep</i>.</p> <p>DDI_FAILURE        Failed to successfully fill in the cookie.</p>
<b>CONTEXT</b>	<code>ddi_dma_htoc()</code> can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><code>ddi_dma_addr_setup(9F)</code>, <code>ddi_dma_buf_setup(9F)</code>,  <code>ddi_dma_setup(9F)</code>, <code>ddi_dma_sync(9F)</code>, <code>ddi_dma_cookie(9S)</code></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_dma_mem_alloc – allocate memory for DMA transfer				
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_mem_alloc(ddi_dma_handle_t handle, size_t length, ddi_device_acc_attr_t *accattrp, uint_t flags, int (*waitfp) (caddr_t), caddr_t arg, caddr_t *kaddrp, size_t *real_length, ddi_acc_handle_t *handlep);</pre>				
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).				
<b>PARAMETERS</b>	<p><b>handle</b>            The DMA handle previously allocated by a call to <b>ddi_dma_alloc_handle(9F)</b>.</p> <p><b>length</b>            The length in bytes of the desired allocation.</p> <p><b>accattrp</b>          Pointer to a device access attribute structure of this device (see <b>ddi_device_acc_attr(9S)</b>).</p> <p><b>flags</b>             Data transfer mode flags. Possible values are:</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="padding-right: 2em;">DDI_DMA_STREAMING</td> <td>Sequential, unidirectional, block-sized, and block-aligned transfers.</td> </tr> <tr> <td style="padding-right: 2em;">DDI_DMA_CONSISTENT</td> <td>Nonsequential transfers of small objects.</td> </tr> </table> <p><b>waitfp</b>            The address of a function to call back later if resources are not available now. The callback function indicates how a caller wants to handle the possibility of resources not being available. If callback is set to <b>DDI_DMA_DONTWAIT</b>, the caller does not care if the allocation fails, and can handle an allocation failure appropriately. If callback is set to <b>DDI_DMA_SLEEP</b>, the caller wishes to have the allocation routines wait for resources to become available. If any other value is set and a DMA resource allocation fails, this value is assumed to be the address of a function to be called when resources become available. When the specified function is called, <i>arg</i> is passed to it as an argument. The specified callback function must return either <b>DDI_DMA_CALLBACK_RUNOUT</b> or <b>DDI_DMA_CALLBACK_DONE</b>. <b>DDI_DMA_CALLBACK_RUNOUT</b> indicates that the callback function attempted to allocate DMA resources but failed. In this case, the callback function</p>	DDI_DMA_STREAMING	Sequential, unidirectional, block-sized, and block-aligned transfers.	DDI_DMA_CONSISTENT	Nonsequential transfers of small objects.
DDI_DMA_STREAMING	Sequential, unidirectional, block-sized, and block-aligned transfers.				
DDI_DMA_CONSISTENT	Nonsequential transfers of small objects.				

is put back on a list to be called again later. `DDI_DMA_CALLBACK_DONE` indicates that either the allocation of DMA resources was successful or the driver no longer wishes to retry. The callback function is called in interrupt context. Therefore, only system functions accessible from interrupt context are available.

The callback function must take whatever steps are necessary to protect its critical resources, data structures, queues, and so on.

<b><i>arg</i></b>	Argument to be passed to the callback function, if such a function is specified.
<b><i>kaddrp</i></b>	On successful return, <i>kaddrp</i> points to the allocated memory.
<b><i>real_length</i></b>	The amount of memory, in bytes, allocated. Alignment and padding requirements may require <code>ddi_dma_mem_alloc()</code> to allocate more memory than requested in <i>length</i> .
<b><i>handlep</i></b>	Pointer to a data access handle.

## DESCRIPTION

`ddi_dma_mem_alloc()` allocates memory for DMA transfers to or from a device. The allocation will obey the alignment, padding constraints and device granularity as specified by the DMA attributes (see `ddi_dma_attr(9S)`) passed to `ddi_dma_alloc_handle(9F)` and the more restrictive attributes imposed by the system.

*flags* should be set to `DDI_DMA_STREAMING` if the device is doing sequential, unidirectional, block-sized, and block-aligned transfers to or from memory. The alignment and padding constraints specified by the *minxfer* and *burstsizes* fields in the DMA attribute structure, `ddi_dma_attr(9S)` (see `ddi_dma_alloc_handle(9F)`) will be used to allocate the most effective hardware support for large transfers. For example, if an I/O transfer can be sped up by using an I/O cache, which has a minimum transfer of one cache line, `ddi_dma_mem_alloc()` will align the memory at a cache line boundary and it will round up *real\_length* to a multiple of the cache line size.

*flags* should be set to `DDI_DMA_CONSISTENT` if the device accesses memory randomly, or if synchronization steps using `ddi_dma_sync(9F)` need to be as efficient as possible. I/O parameter blocks used for communication between a device and a driver should be allocated using `DDI_DMA_CONSISTENT`.

The device access attributes are specified in the location pointed by the *accattrp* argument (see `ddi_device_acc_attr(9S)`).

The data access handle is returned in *handlep*. *handlep* is opaque – drivers may not attempt to interpret its value. To access the data content, the driver must invoke `ddi_get8(9F)` or `ddi_put8(9F)` (depending on the data transfer direction) with the data access handle.

DMA resources must be established before performing a DMA transfer by passing *kaddrp* and *real\_length* as returned from `ddi_dma_mem_alloc()` and the flag `DDI_DMA_STREAMING` or `DDI_DMA_CONSISTENT` to `ddi_dma_addr_bind_handle(9F)`. In addition, to ensure the consistency of a memory object shared between the CPU and the device after a DMA transfer, explicit synchronization steps using `ddi_dma_sync(9F)` or `ddi_dma_unbind_handle(9F)` are required.

**RETURN VALUES**

`ddi_dma_mem_alloc()` returns:

`DDI_SUCCESS` Memory successfully allocated.

`DDI_FAILURE` Memory allocation failed.

**CONTEXT**

`ddi_dma_mem_alloc()` can be called from user or interrupt context, except when *waitfp* is set to `DDI_DMA_SLEEP`, in which case it can be called from user context only.

**SEE ALSO**

`ddi_dma_addr_bind_handle(9F)`, `ddi_dma_alloc_handle(9F)`,  
`ddi_dma_mem_free(9F)`, `ddi_dma_sync(9F)`,  
`ddi_dma_unbind_handle(9F)`, `ddi_get8(9F)`, `ddi_put8(9F)`,  
`ddi_device_acc_attr(9S)`, `ddi_dma_attr(9S)`

*Writing Device Drivers*

**WARNINGS**

If `DDI_NEVERSWAP_ACC` is specified, memory can be used for any purpose; but if either endian mode is specified, you must use `ddi_get/put*` and never anything else.

<b>NAME</b>	ddi_dma_mem_free – free previously allocated memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_dma_mem_free(ddi_acc_handle_t *handlep);</pre>
<b>PARAMETERS</b>	<p><b>handlep</b>            Pointer to the data access handle previously allocated by a call to ddi_dma_mem_alloc(9F).</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_dma_mem_free()</b> deallocates the memory acquired by <b>ddi_dma_mem_alloc(9F)</b> . In addition, it destroys the data access handle <i>handlep</i> associated with the memory.
<b>CONTEXT</b>	<b>ddi_dma_mem_free()</b> can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<b>ddi_dma_mem_alloc(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_dma_movwin – shift current DMA window
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_movwin(ddi_dma_handle_t handle, off_t *offp, uint_t *lenp, ddi_dma_cookie_t *cookiep);</pre>
<b>INTERFACE LEVEL</b>	Solaris SPARC DDI specific (Solaris SPARC DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The DMA handle filled in by a call to <b>ddi_dma_setup(9F)</b>.</p> <p><b>offp</b>              A pointer to an offset to set the DMA window to. Upon a successful return, it will be filled in with the new offset from the beginning of the object resources are allocated for.</p> <p><b>lenp</b>               A pointer to a value which must either be the current size of the DMA window (as known from a call to <b>ddi_dma_curwin(9F)</b> or from a previous call to <b>ddi_dma_movwin()</b>). Upon a successful return, it will be filled in with the size, in bytes, of the current window.</p> <p><b>cookiep</b>           A pointer to a DMA cookie (see <b>ddi_dma_cookie(9S)</b>). Upon a successful return, <i>cookiep</i> is filled in just as if an implicit <b>ddi_dma_htoc(9F)</b> had been made.</p>
<b>DESCRIPTION</b>	<p><b>ddi_dma_movwin()</b> shifts the current DMA window. If a DMA request allows the system to allocate resources for less than the entire object by setting the <b>DDI_DMA_PARTIAL</b> flag in the <b>ddi_dma_req(9S)</b> structure, the current DMA window can be shifted by a call to <b>ddi_dma_movwin()</b>.</p> <p>The caller must first determine the current DMA window size by a call to <b>ddi_dma_curwin(9F)</b>. Using the current offset and size of the window thus retrieved, the caller of <b>ddi_dma_movwin()</b> may change the window onto the object by changing the offset by a value which is some multiple of the size of the DMA window.</p> <p><b>ddi_dma_movwin()</b> takes care of underlying resource synchronizations required to <i>shift</i> the window. However, if you want to <i>access</i> the data prior to or after moving the window, further synchronizations using <b>ddi_dma_sync(9F)</b> are required.</p> <p>This function is normally called from an interrupt routine. The first invocation of the DMA engine is done from the driver. All subsequent invocations of the DMA engine are done from the interrupt routine. The interrupt routine checks</p>

to see if the request has been completed. If it has, it returns without invoking another DMA transfer. Otherwise it calls **ddi\_dma\_movwin()** to shift the current window and starts another DMA transfer.

**RETURN VALUES**

**ddi\_dma\_movwin()** returns:

DDI\_SUCCESS    The current length and offset are legal and have been set.

DDI\_FAILURE    Otherwise.

**CONTEXT**

**ddi\_dma\_movwin()** can be called from user or interrupt context.

**SEE ALSO**

**ddi\_dma\_curwin(9F)**, **ddi\_dma\_htoc(9F)**, **ddi\_dma\_setup(9F)**,  
**ddi\_dma\_sync(9F)**, **ddi\_dma\_cookie(9S)**, **ddi\_dma\_req(9S)**

*Writing Device Drivers*

**WARNINGS**

The caller must guarantee that the resources used by the object are inactive prior to calling this function.

<b>NAME</b>	ddi_dma_nextcookie – retrieve subsequent DMA cookie
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_dma_nextcookie(ddi_dma_handle_t handle, ddi_dma_cookie_t *cookiep);</pre>
<b>PARAMETERS</b>	<p><b>handle</b>            The handle previously allocated by a call to <code>ddi_dma_alloc_handle(9F)</code>.</p> <p><b>cookiep</b>           A pointer to a <code>ddi_dma_cookie(9S)</code> structure.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p><code>ddi_dma_nextcookie()</code> retrieves subsequent DMA cookies for a DMA object. <code>ddi_dma_nextcookie()</code> fills in the <code>ddi_dma_cookie(9S)</code> structure pointed to by <code>cookiep</code>. The <code>ddi_dma_cookie(9S)</code> structure must be allocated prior to calling <code>ddi_dma_nextcookie()</code>.</p> <p>The DMA cookie count returned by <code>ddi_dma_buf_bind_handle(9F)</code>, <code>ddi_dma_addr_bind_handle(9F)</code>, or <code>ddi_dma_getwin(9F)</code> indicates the number of DMA cookies a DMA object consists of. If the resulting cookie count, <i>N</i>, is larger than 1, <code>ddi_dma_nextcookie()</code> must be called <i>N</i>-1 times to retrieve all DMA cookies.</p>
<b>CONTEXT</b>	<code>ddi_dma_nextcookie()</code> can be called from user, kernel, or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b>    process a scatter-gather list of I/O requests</p> <p>This example demonstrates the use of <code>ddi_dma_nextcookie()</code> to process a scatter-gather list of I/O requests.</p> <pre>/* setup scatter-gather list with multiple DMA cookies */ ddi_dma_cookie_t dmacookie; uint_t    ccount; ...  status = ddi_dma_buf_bind_handle(handle, bp, DDI_DMA_READ,     NULL, NULL, &amp;dmacookie, &amp;ccount);  if (status == DDI_DMA_MAPPED) {      /* program DMA engine with first cookie */      while (--ccount &gt; 0) {         ddi_dma_nextcookie(handle, &amp;dmacookie);         /* program DMA engine with next cookie */     } }</pre>

```
}  
...
```

**SEE ALSO**

`ddi_dma_addr_bind_handle(9F)`, `ddi_dma_alloc_handle(9F)`,  
`ddi_dma_buf_bind_handle(9F)`, `ddi_dma_unbind_handle(9F)`,  
`ddi_dma_cookie(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_nextseg – get next DMA segment
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int <b>ddi_dma_nextseg</b>(ddi_dma_win_t win, ddi_dma_seg_t seg, ddi_dma_seg_t *nseg);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>win</b>                    A DMA window.</p> <p><b>seg</b>                     The current DMA segment or NULL.</p> <p><b>nseg</b>                    A pointer to the next DMA segment to be filled in. If <b>seg</b> is NULL, a pointer to the first segment within the specified window is returned.</p>
<b>DESCRIPTION</b>	<p><b>ddi_dma_nextseg()</b> gets the next DMA segment within the specified window <i>win</i>. If the current segment is NULL, the first DMA segment within the window is returned.</p> <p>A DMA segment is always required for a DMA window. A DMA segment is a contiguous portion of a DMA window (see <b>ddi_dma_nextwin</b>(9F)) which is entirely addressable by the device for a data transfer operation.</p> <p>An example where multiple DMA segments are allocated is where the system does not contain DVMA capabilities and the object may be non-contiguous. In this example the object will be broken into smaller contiguous DMA segments. Another example is where the device has an upper limit on its transfer size (for example an 8-bit address register) and has expressed this in the DMA limit structure (see <b>ddi_dma_lim_sparc</b>(9S) or <b>ddi_dma_lim_x86</b>(9S)). In this example the object will be broken into smaller addressable DMA segments.</p>
<b>RETURN VALUES</b>	<p><b>ddi_dma_nextseg()</b> returns:</p> <p>DDI_SUCCESS                Successfully filled in the next segment pointer.</p> <p>DDI_DMA_DONE                There is no next segment. The current segment is the final segment within the specified window.</p> <p>DDI_DMA_STALE                <i>win</i> does not refer to the currently active window.</p>
<b>CONTEXT</b>	<b>ddi_dma_nextseg()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	For an example, see <b>ddi_dma_segtocookie</b> (9F).

**SEE ALSO**

`ddi_dma_addr_setup(9F)`, `ddi_dma_buf_setup(9F)`,  
`ddi_dma_nextwin(9F)`, `ddi_dma_segtocookie(9F)`, `ddi_dma_sync(9F)`,  
`ddi_dma_lim_sparc(9S)`, `ddi_dma_lim_x86(9S)`, `ddi_dma_req(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_nextwin – get next DMA window
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_dma_nextwin(ddi_dma_handle_t handle, ddi_dma_win_t win, ddi_dma_win_t *<i>nwin</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>handle</i></b>            A DMA handle.</p> <p><b><i>win</i></b>                The current DMA window or NULL.</p> <p><b><i>nwin</i></b>                A pointer to the next DMA window to be filled in. If <i>win</i> is NULL, a pointer to the first window within the object is returned.</p>
<b>DESCRIPTION</b>	<p><b>ddi_dma_nextwin()</b> shifts the current DMA window <i>win</i> within the object referred to by <i>handle</i> to the next DMA window <i>nwin</i>. If the current window is NULL, the first window within the object is returned. A DMA window is a portion of a DMA object or might be the entire object. A DMA window has system resources allocated to it and is prepared to accept data transfers. Examples of system resources are DVMA mapping resources and intermediate transfer buffer resources.</p> <p>All DMA objects require a window. If the DMA window represents the whole DMA object it has system resources allocated for the entire data transfer. However, if the system is unable to setup the entire DMA object due to system resource limitations, the driver writer may allow the system to allocate system resources for less than the entire DMA object. This can be accomplished by specifying the DDI_DMA_PARTIAL flag as a parameter to <b>ddi_dma_buf_setup(9F)</b> or <b>ddi_dma_addr_setup(9F)</b> or as part of a <b>ddi_dma_req(9S)</b> structure in a call to <b>ddi_dma_setup(9F)</b>.</p> <p>Only the window that has resources allocated is valid per object at any one time. The currently valid window is the one that was most recently returned from <b>ddi_dma_nextwin()</b>. Furthermore, because a call to <b>ddi_dma_nextwin()</b> will reallocate system resources to the new window, the previous window will become invalid. It is a <i>severe</i> error to call <b>ddi_dma_nextwin()</b> before any transfers into the current window are complete.</p> <p><b>ddi_dma_nextwin()</b> takes care of underlying memory synchronizations required to shift the window. However, if you want to access the data before</p>

or after moving the window, further synchronizations using `ddi_dma_sync(9F)` are required.

**RETURN VALUES**

**ddi\_dma\_nextwin()** returns:

DDI\_SUCCESS Successfully filled in the next window pointer.

DDI\_DMA\_DONE There is no next window. The current window is the final window within the specified object.

DDI\_DMA\_STALE *win* does not refer to the currently active window.

**CONTEXT**

**ddi\_dma\_nextwin()** can be called from user or interrupt context.

**EXAMPLES**

For an example see `ddi_dma_segtocookie(9F)`.

**SEE ALSO**

`ddi_dma_addr_setup(9F)`, `ddi_dma_buf_setup(9F)`,  
`ddi_dma_nextseg(9F)`, `ddi_dma_segtocookie(9F)`, `ddi_dma_sync(9F)`,  
`ddi_dma_req(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_numwin – retrieve number of DMA windows
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_numwin(ddi_dma_handle_t handle, uint_t *nwinp);</pre>
<b>PARAMETERS</b>	<p><b>handle</b>            The DMA handle previously allocated by a call to <code>ddi_dma_alloc_handle(9F)</code>.</p> <p><b>nwinp</b>             Upon a successful return, <i>nwinp</i> will contain the number of DMA windows for this object.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_dma_numwin()</b> returns the number of DMA windows for a DMA object if partial resource allocation was permitted.
<b>RETURN VALUES</b>	<p><b>ddi_dma_numwin()</b> returns:</p> <p>DDI_SUCCESS            Successfully filled in the number of DMA windows.</p> <p>DDI_FAILURE            DMA windows are not activated.</p>
<b>CONTEXT</b>	<b>ddi_dma_numwin()</b> can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<p><code>ddi_dma_addr_bind_handle(9F)</code>, <code>ddi_dma_alloc_handle(9F)</code>,  <code>ddi_dma_buf_bind_handle(9F)</code>, <code>ddi_dma_unbind_handle(9F)</code></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_dma_segtocookie – convert a DMA segment to a DMA address cookie
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_segtocookie(ddi_dma_seg_t seg, off_t *offp, off_t *lenp,     ddi_dma_cookie_t *cookiep);</pre>
<b>PARAMETERS</b>	<p><b>seg</b>                   A DMA segment.</p> <p><b>offp</b>                  A pointer to an <code>off_t</code>. Upon a successful return, it is filled in with the offset. This segment is addressing within the object.</p> <p><b>lenp</b>                  The byte length. This segment is addressing within the object.</p> <p><b>cookiep</b>               A pointer to a DMA cookie (see <code>ddi_dma_cookie(9S)</code>).</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_dma_segtocookie()</b> takes a DMA segment and fills in the cookie pointed to by <i>cookiep</i> with the appropriate address, length, and bus type to be used to program the DMA engine. <b>ddi_dma_segtocookie()</b> also fills in <i>*offp</i> and <i>*lenp</i> , which specify the range within the object.
<b>RETURN VALUES</b>	<p><b>ddi_dma_segtocookie()</b> returns:</p> <p>DDI_SUCCESS            Successfully filled in all values.</p> <p>DDI_FAILURE            Failed to successfully fill in all values.</p>
<b>CONTEXT</b>	<b>ddi_dma_segtocookie()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>CODE EXAMPLE 1 ddi_dma_segtocookie() example</b></p> <pre>for (win = NULL; (retw = ddi_dma_nextwin(handle, win, &amp;nwin)) !=     DDI_DMA_DONE; win = nwin) {     if (retw != DDI_SUCCESS) {         /* do error handling */     } else {         for (seg = NULL; (rets = ddi_dma_nextseg(nwin, seg, &amp;nseg)) !=             DDI_DMA_DONE; seg = nseg) {             if (rets != DDI_SUCCESS) {                  /* do error handling */             } else {</pre>

```
    ddi_dma_segtocookie(nseg, &off, &len, &cookie);  
    /* program DMA engine */  
    }  
  }  
}
```

**SEE ALSO** [ddi\\_dma\\_nextseg\(9F\)](#), [ddi\\_dma\\_nextwin\(9F\)](#), [ddi\\_dma\\_sync\(9F\)](#),  
[ddi\\_dma\\_cookie\(9S\)](#)

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_set_sbus64 – allow 64-bit transfers on SBus
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_set_sbus64(ddi_dma_handle_t handle, uint_t burstsizes);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The handle filled in by a call to <code>ddi_dma_alloc_handle(9F)</code>.</p> <p><b>burstsizes</b>       The possible burst sizes the device's DMA engine can accept in 64-bit mode.</p>
<b>DESCRIPTION</b>	<p><b>ddi_dma_set_sbus64()</b> informs the system that the device wishes to perform 64-bit data transfers on the SBus. The driver must first allocate a DMA handle using <code>ddi_dma_alloc_handle(9F)</code> with a <code>ddi_dma_attr(9S)</code> structure describing the DMA attributes for a 32-bit transfer mode.</p> <p><i>burstsizes</i> describes the possible burst sizes the device's DMA engine can accept in 64-bit mode. It may be distinct from the burst sizes for 32-bit mode set in the <code>ddi_dma_attr(9S)</code> structure. The system will activate 64-bit SBus transfers if the SBus supports them. Otherwise, the SBus will operate in 32-bit mode.</p> <p>After DMA resources have been allocated (see <code>ddi_dma_addr_bind_handle(9F)</code> or <code>ddi_dma_buf_bind_handle(9F)</code>), the driver should retrieve the available burst sizes by calling <code>ddi_dma_burstsizes(9F)</code>. This function will return the burst sizes in 64-bit mode if the system was able to activate 64-bit transfers. Otherwise burst sizes will be returned in 32-bit mode.</p>
<b>RETURN VALUES</b>	<p><b>ddi_dma_set_sbus64()</b> returns:</p> <p>DDI_SUCCESS   Successfully set the SBus to 64-bit mode.</p> <p>DDI_FAILURE   64-bit mode could not be set.</p>
<b>CONTEXT</b>	<b>ddi_dma_set_sbus64()</b> can be called from user, kernel, or interrupt context.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SBus

**SEE ALSO**

`attributes(5)`, `ddi_dma_addr_bind_handle(9F)`,  
`ddi_dma_alloc_handle(9F)`, `ddi_dma_buf_bind_handle(9F)`,  
`ddi_dma_burstsizes(9F)`, `ddi_dma_attr(9S)`

**NOTES**

64-bit SBus mode is activated on a per SBus slot basis. If there are multiple SBus cards in one slot, they all must operate in 64-bit mode or they all must operate in 32-bit mode.

<b>NAME</b>	ddi_dma_setup – setup DMA resources
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_setup(dev_info_t *dip, ddi_dma_req_t *dmareqp, ddi_dma_handle_t *handlep);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                   A pointer to the device's dev_info structure.</p> <p><b>dmareqp</b>               A pointer to a DMA request structure (see <code>ddi_dma_req(9S)</code>).</p> <p><b>handlep</b>               A pointer to a DMA handle to be filled in. See below for a discussion of a handle. If <i>handlep</i> is NULL, the call to <b>ddi_dma_setup()</b> is considered an advisory call, in which case no resources are allocated, but a value indicating the legality and the feasibility of the request is returned.</p>
<b>DESCRIPTION</b>	<p><b>ddi_dma_setup()</b> allocates resources for a memory object such that a device can perform DMA to or from that object.</p> <p>A call to <b>ddi_dma_setup()</b> informs the system that device referred to by <i>dip</i> wishes to perform DMA to or from a memory object. The memory object, the device's DMA capabilities, the device driver's policy on whether to wait for resources, are all specified in the <code>ddi_dma_req</code> structure pointed to by <i>dmareqp</i>.</p> <p>A successful call to <b>ddi_dma_setup()</b> fills in the value pointed to by <i>handlep</i>. This is an opaque object called a DMA handle. This handle is then used in subsequent DMA calls, until <b>ddi_dma_free(9F)</b> is called.</p> <p>Again a DMA handle is opaque—drivers may <i>not</i> attempt to interpret its value. When a driver wants to enable its DMA engine, it must retrieve the appropriate address to supply to its DMA engine using a call to <b>ddi_dma_htoc(9F)</b>, which takes a pointer to a DMA handle and returns the appropriate DMA address.</p> <p>When DMA transfer completes, the driver should free up the the allocated DMA resources by calling <b>ddi_dma_free()</b>.</p>
<b>RETURN VALUES</b>	<b>ddi_dma_setup()</b> returns:

DDI_DMA_MAPPED	Successfully allocated resources for the object. In the case of an <i>advisory</i> call, this indicates that the request is legal.
DDI_DMA_PARTIAL_MAP	Successfully allocated resources for a <i>part</i> of the object. This is acceptable when partial transfers are allowed using a flag setting in the <code>ddi_dma_req</code> structure (see <code>ddi_dma_req(9S)</code> and <code>ddi_dma_movwin(9F)</code> ).
DDI_DMA_NORESOURCES	When no resources are available.
DDI_DMA_NOMAPPING	The object cannot be reached by the device requesting the resources.
DDI_DMA_TOOBIG	The object is too big and exceeds the available resources. The maximum size varies depending on machine and configuration.

**CONTEXT** `ddi_dma_setup()` can be called from user or interrupt context, except when the `dmr_fp` member of the `ddi_dma_req` structure pointed to by *dmareqp* is set to `DDI_DMA_SLEEP`, in which case it can be called from user context only.

**SEE ALSO** `ddi_dma_addr_setup(9F)`, `ddi_dma_buf_setup(9F)`, `ddi_dma_free(9F)`, `ddi_dma_htoc(9F)`, `ddi_dma_movwin(9F)`, `ddi_dma_sync(9F)`, `ddi_dma_req(9S)`

*Writing Device Drivers*

**NOTES** The construction of the `ddi_dma_req` structure is complicated. Use of the provided interface functions such as `ddi_dma_buf_setup(9F)` simplifies this task.

<b>NAME</b>	ddi_dma_sync – synchronize CPU and I/O views of memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_sync(ddi_dma_handle_t handle, off_t offset, size_t length, uint_t type);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The <i>handle</i> filled in by a call to <code>ddi_dma_alloc_handle(9F)</code>.</p> <p><b>offset</b>           The offset into the object described by the <i>handle</i>.</p> <p><b>length</b>           The length, in bytes, of the area to synchronize. When <i>length</i> is zero, the entire range starting from <i>offset</i> to the end of the object has the requested operation applied to it.</p> <p><b>type</b>             Indicates the caller's desire about what view of the memory object to synchronize. The possible values are <code>DDI_DMA_SYNC_FORDEV</code>, <code>DDI_DMA_SYNC_FORCPU</code> and <code>DDI_DMA_SYNC_FORKERNEL</code>.</p>
<b>DESCRIPTION</b>	<p><code>ddi_dma_sync()</code> is used to selectively synchronize either a DMA device's or a CPU's view of a memory object that has DMA resources allocated for I/O . This may involve operations such as flushes of CPU or I/O caches, as well as other more complex operations such as stalling until hardware write buffers have drained.</p> <p>This function need only be called under certain circumstances. When resources are allocated for DMA using <code>ddi_dma_addr_bind_handle()</code> or <code>ddi_dma_buf_bind_handle()</code>, an implicit <code>ddi_dma_sync()</code> is done. When DMA resources are deallocated using <code>ddi_dma_unbind_handle(9F)</code>, an implicit <code>ddi_dma_sync()</code> is done. However, at any time between DMA resource allocation and deallocation, if the memory object has been modified by either the DMA device or a CPU and you wish to ensure that the change is noticed by the party that <i>didnot</i> do the modifying, a call to <code>ddi_dma_sync()</code> is required. This is true independent of any attributes of the memory object including, but not limited to, whether or not the memory was allocated for consistent mode I/O (see <code>ddi_dma_mem_alloc(9F)</code>) or whether or not DMA resources have been allocated for consistent mode I/O (see <code>ddi_dma_addr_bind_handle(9F)</code> or <code>ddi_dma_buf_bind_handle(9F)</code>).</p> <p>This cannot be stated too strongly. If a consistent view of the memory object must be ensured between the time DMA resources are allocated for the object</p>

and the time they are deallocated, you *must* call **ddi\_dma\_sync()** to ensure that either a CPU or a DMA device has such a consistent view.

What to set `type` to depends on the view you are trying to ensure consistency for. If the memory object is modified by a CPU, and the object is going to be read by the DMA engine of the device, use `DDI_DMA_SYNC_FORDEV`. This ensures that the device's DMA engine sees any changes that a CPU has made to the memory object. If the DMA engine for the device has *written* to the memory object, and you are going to *read* (with a CPU) the object (using an extant virtual address mapping that you have to the memory object), use `DDI_DMA_SYNC_FORCPU`. This ensures that a CPU's view of the memory object includes any changes made to the object by the device's DMA engine. If you are only interested in the kernel's view (kernel-space part of the CPU's view) you may use `DDI_DMA_SYNC_FORKERNEL`. This gives a hint to the system—that is, if it is more economical to synchronize the kernel's view only, then do so; otherwise, synchronize for CPU.

**RETURN VALUES**

**ddi\_dma\_sync()** returns:

`DDI_SUCCESS` Caches are successfully flushed.

`DDI_FAILURE` The address range to be flushed is out of the address range established by `ddi_dma_addr_bind_handle(9F)` or `ddi_dma_buf_bind_handle(9F)`.

**CONTEXT**

**ddi\_dma\_sync()** can be called from user or interrupt context.

**SEE ALSO**

`ddi_dma_addr_bind_handle(9F)`, `ddi_dma_alloc_handle(9F)`,  
`ddi_dma_buf_bind_handle(9F)`, `ddi_dma_mem_alloc(9F)`,  
`ddi_dma_unbind_handle(9F)`

*Writing Device Drivers*

<b>NAME</b>	ddi_dma_unbind_handle – unbinds the address in a DMA handle
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_dma_unbind_handle(ddi_dma_handle_t handle);</pre>
<b>PARAMETERS</b>	<p><b>handle</b>            The DMA handle previously allocated by a call to <code>ddi_dma_alloc_handle(9F)</code>.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p><b>ddi_dma_unbind_handle()</b> frees all DMA resources associated with an existing DMA handle. When a DMA transfer completes, the driver should call <b>ddi_dma_unbind_handle()</b> to free system DMA resources established by a call to <code>ddi_dma_buf_bind_handle(9F)</code> or <code>ddi_dma_addr_bind_handle(9F)</code>. <b>ddi_dma_unbind_handle()</b> does an implicit <code>ddi_dma_sync(9F)</code> making further synchronization steps unnecessary.</p>
<b>RETURN VALUES</b>	<p>DDI_SUCCESS    on success</p> <p>DDI_FAILURE    on failure</p>
<b>CONTEXT</b>	<b>ddi_dma_unbind_handle()</b> can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<p><code>ddi_dma_addr_bind_handle(9F)</code>, <code>ddi_dma_alloc_handle(9F)</code>,  <code>ddi_dma_buf_bind_handle(9F)</code>, <code>ddi_dma_free_handle(9F)</code>,  <code>ddi_dma_sync(9F)</code></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_enter_critical, ddi_exit_critical – enter and exit a critical region of control
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  unsigned int ddi_enter_critical(void);  void ddi_exit_critical(unsignedint ddi);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>ddi</b>                    The returned value from the call to <b>ddi_enter_critical()</b> must be passed to <b>ddi_exit_critical()</b> .</p>
<b>DESCRIPTION</b>	<p>Nearly all driver operations can be done without any special synchronization and protection mechanisms beyond those provided by, for example, mutexes (see <b>mutex(9F)</b> ). However, for certain devices there can exist a very short critical region of code which <i>must</i> be allowed to run uninterrupted. The function <b>ddi_enter_critical()</b> provides a mechanism by which a driver can ask the system to guarantee to the best of its ability that the current thread of execution will neither be preempted nor interrupted. This stays in effect until a bracketing call to <b>ddi_exit_critical()</b> is made (with an argument which was the returned value from <b>ddi_enter_critical()</b> ).</p> <p>The driver may not call any functions external to itself in between the time it calls <b>ddi_enter_critical()</b> and the time it calls <b>ddi_exit_critical()</b> .</p>
<b>RETURN VALUES</b>	<b>ddi_enter_critical()</b> returns an opaque unsigned integer which must be used in the subsequent call to <b>ddi_exit_critical()</b> .
<b>CONTEXT</b>	This function can be called from user or interrupt context.
<b>WARNINGS</b>	<p>Driver writers should note that in a multiple processor system this function does not temporarily suspend other processors from executing. This function also cannot guarantee to actually block the hardware from doing such things as interrupt acknowledge cycles. What it <i>can</i> do is guarantee that the currently executing thread will not be preempted.</p> <p>Do not write code bracketed by <b>ddi_enter_critical()</b> and <b>ddi_exit_critical()</b> that can get caught in an infinite loop, as the machine may crash if you do.</p>
<b>SEE ALSO</b>	<p><b>mutex(9F)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_ffs, ddi_fls – find first (last) bit set in a long integer
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_ffs(long mask); int ddi_fls(long mask);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>mask</b> A 32-bit argument value to search through.
<b>DESCRIPTION</b>	The function <b>ddi_ffs()</b> takes its argument and returns the shift count that the first (least significant) bit set in the argument corresponds to. The function <b>ddi_fls()</b> does the same, only it returns the shift count for the last (most significant) bit set in the argument.
<b>RETURN VALUES</b>	<p>0            No bits are set in mask.</p> <p><i>N</i>          Bit <i>N</i> is the least significant ( <code>ddi_ffs</code> ) or most significant ( <code>ddi_fls</code> ) bit set in mask. Bits are numbered from 1 to 32 , with bit 1 being the least significant bit position and bit 32 the most significant position.</p>
<b>CONTEXT</b>	This function can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_get8, ddi_get16, ddi_get32, ddi_get64, ddi_getb, ddi_getw, ddi_getl, ddi_getll – read data from the mapped memory address, device register or allocated DMA memory address
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  uint8_t ddi_get8(ddi_acc_handle_t handle, uint8_t * dev_addr); uint16_t ddi_get16(ddi_acc_handle_t handle, uint16_t * dev_addr); uint32_t ddi_get32(ddi_acc_handle_t handle, uint32_t * dev_addr); uint64_t ddi_get64(ddi_acc_handle_t handle, uint64_t * dev_addr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The data access handle returned from setup calls, such as ddi_regs_map_setup(9F) .</p> <p><b>dev_addr</b>        Base device address.</p>
<b>DESCRIPTION</b>	<p>The ddi_get8() , ddi_get16() , ddi_get32() , and ddi_get64() functions read 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, from the device address, dev_addr .</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>RETURN VALUES</b>	These functions return the value read from the mapped address.
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	ddi_put8(9F) , ddi_regs_map_free(9F) , ddi_regs_map_setup(9F) , ddi_rep_get8(9F) , ddi_rep_put8(9F)
<b>NOTES</b>	The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
ddi_getb	ddi_get8
ddi_getw	ddi_get16
ddi_getl	ddi_get32
ddi_getll	ddi_get64

<b>NAME</b>	ddi_get_cred – returns a pointer to the credential structure of the caller
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  cred_t *ddi_get_cred(void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_get_cred()</b> returns a pointer to the user credential structure of the caller.
<b>RETURN VALUES</b>	<b>ddi_get_cred()</b> returns a pointer to the caller's credential structure.
<b>CONTEXT</b>	<b>ddi_get_cred()</b> can be called from user context only.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_get_driver_private, ddi_set_driver_private – get or set the address of the device's private data area
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>void ddi_set_driver_private(dev_info_t * dip, caddr_t data); caddr_t ddi_get_driver_private(dev_info_t * dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
<b>ddi_get_driver_private()</b>	<b>dip</b> Pointer to device information structure to get from.
<b>ddi_set_driver_private()</b>	<b>dip</b> Pointer to device information structure to set. <b>data</b> Data area address to set.
<b>DESCRIPTION</b>	<p><b>ddi_get_driver_private()</b> returns the address of the device's private data area from the device information structure pointed to by <i>dip</i> .</p> <p><b>ddi_set_driver_private()</b> sets the address of the device's private data area in the device information structure pointed to by <i>dip</i> with the value of <i>data</i> .</p>
<b>RETURN VALUES</b>	<b>ddi_get_driver_private()</b> returns the contents of <code>devi_driver_data</code> . If <b>ddi_set_driver_private()</b> has not been previously called with <i>dip</i> , an unpredictable value is returned.
<b>CONTEXT</b>	These functions can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_get_instance – get device instance number
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_get_instance(dev_info_t *dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>dip</b> Pointer to dev_info structure.
<b>DESCRIPTION</b>	<p><b>ddi_get_instance()</b> returns the instance number of the device corresponding to <i>dip</i>.</p> <p>The system assigns an instance number to every device. Instance numbers for devices attached to the same driver are unique. This provides a way for the system and the driver to uniquely identify one or more devices of the same type. The instance number is derived by the system from different properties for different device types in an implementation specific manner.</p> <p>Once an instance number has been assigned to a device, it will remain the same even across reconfigurations and reboots. Therefore, instance numbers seen by a driver may not appear to be in consecutive order. For example, if device <code>f000</code> has been assigned an instance number of 0 and device <code>f001</code> has been assigned an instance number of 1, if <code>f000</code> is removed, <code>f001</code> will continue to be associated with instance number 1 (even though <code>f001</code> is now the only device of its type on the system).</p>
<b>RETURN VALUES</b>	<b>ddi_get_instance()</b> returns the instance number of the device corresponding to <i>dip</i> .
<b>CONTEXT</b>	<b>ddi_get_instance()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><code>path_to_inst(4)</code></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_get_lbolt - returns the value of lbolt
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  clock_t ddi_get_lbolt (void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_get_lbolt()</b> returns the value of <code>lbolt</code> where <code>lbolt</code> is an integer that represents the number of clock ticks since the last system reboot. This value is used as a counter or timer inside the system kernel. The tick frequency can be determined by using <b>drv_usec2ohz(9F)</b> which converts microseconds into clock ticks.
<b>RETURN VALUES</b>	<b>ddi_get_lbolt()</b> returns the value of <code>lbolt</code> .
<b>CONTEXT</b>	This routine can to be called from any context.
<b>SEE ALSO</b>	<b>ddi_get_time(9F)</b> , <b>drv_getparm(9F)</b> , <b>drv_usec2ohz(9F)</b> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	ddi_get_parent – find the parent of a device information structure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  dev_info_t *ddi_get_parent(dev_info_t *dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>dip</b> Pointer to a device information structure.
<b>DESCRIPTION</b>	<b>ddi_get_parent()</b> returns a pointer to the device information structure which is the parent of the one pointed to by <i>dip</i> .
<b>RETURN VALUES</b>	<b>ddi_get_parent()</b> returns a pointer to a device information structure.
<b>CONTEXT</b>	<b>ddi_get_parent()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_get_pid - returns the process ID
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  pid_t ddi_get_pid (void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_get_pid()</b> the process ID of the current process. This value can be used to allow only a select process to perform a certain operation. It can also be used to determine if a device context belongs to the current process.
<b>RETURN VALUES</b>	<b>ddi_get_pid()</b> returns process ID.
<b>CONTEXT</b>	This routine can to be called from user context only.
<b>SEE ALSO</b>	<b>drv_getparm(9F)</b> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	ddi_get_time – returns the current time in seconds
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  time_t ddi_get_time (void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_get_time()</b> returns the current time in seconds since 00:00:00 UTC, January 1, 1970. This value can be used to set of wait or expiration intervals.
<b>RETURN VALUES</b>	<b>ddi_get_time()</b> returns the time in seconds.
<b>CONTEXT</b>	This routine can to be called from any context.
<b>SEE ALSO</b>	<b>ddi_get_lbolt(9F)</b> , <b>drv_getparm(9F)</b> , <b>drv_usectohz(9F)</b> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	ddi_in_panic – determine if system is in panic state
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_in_panic(void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p>Drivers controlling devices on which the system may dump a kernel core image in the event of a panic may determine if the system is panicing by calling <b>ddi_in_panic()</b>.</p> <p>When the system is panicing, the calls of functions scheduled by <b>timeout(9F)</b> and <b>ddi_trigger_softintr(9F)</b> will never occur. Neither can <b>delay(9F)</b> be relied upon, since it is implemented via <b>timeout(9F)</b>.</p> <p>Drivers that need to enforce a time delay such as SCSI bus reset delay time must busy-wait when the system is panicing.</p>
<b>RETURN VALUES</b>	<b>ddi_in_panic()</b> returns 1 if the system is in panic, or 0 otherwise.
<b>CONTEXT</b>	<b>ddi_in_panic()</b> may be called from any context.
<b>SEE ALSO</b>	<b>dump(9E)</b> , <b>delay(9F)</b> , <b>ddi_trigger_softintr(9F)</b> , <b>timeout(9F)</b>
	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_intr_hilevel – indicate interrupt handler type
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_intr_hilevel(dev_info_t *dip, uint_t inumber);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                    Pointer to dev_info structure.</p> <p><b>inumber</b>                Interrupt number.</p>
<b>DESCRIPTION</b>	<p><b>ddi_intr_hilevel()</b> returns non-zero if the specified interrupt is a "high level" interrupt.</p> <p>High level interrupts must be handled without using system services that manipulate thread or process states, because these interrupts are not blocked by the scheduler.</p> <p>In addition, high level interrupt handlers must take care to do a minimum of work because they are not preemptable.</p> <p>A typical high level interrupt handler would put data into a circular buffer and schedule a soft interrupt by calling <b>ddi_trigger_softintr()</b>. The circular buffer could be protected by using a mutex that was properly initialized for the interrupt handler.</p> <p><b>ddi_intr_hilevel()</b> can be used before calling <b>ddi_add_intr()</b> to decide which type of interrupt handler should be used. Most device drivers are designed with the knowledge that the devices they support will always generate low level interrupts, however some devices, for example those using SBus or VME bus level 6 or 7 interrupts must use this test because on some machines those interrupts are high level (above the scheduler level) and on other machines they are not.</p>
<b>RETURN VALUES</b>	<p><b>non-zero</b>                indicates a high-level interrupt.</p>
<b>CONTEXT</b>	These functions can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><b>ddi_add_intr(9F)</b>, <b>mutex(9F)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	ddi_io_get8, ddi_io_get16, ddi_io_get32, ddi_io_getb, ddi_io_getw, ddi_io_getl – read data from the mapped device register in I/O space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  uint8_t ddi_io_get8(ddi_acc_handle_t handle, int dev_port); uint16_t ddi_io_get16(ddi_acc_handle_t handle, int dev_port); uint32_t ddi_io_get32(ddi_acc_handle_t handle, int dev_port);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The data access handle returned from setup calls, such as <code>ddi_regs_map_setup(9F)</code> .</p> <p><b>dev_port</b>        The device port.</p>
<b>DESCRIPTION</b>	<p>These routines generate a read of various sizes from the device port, <code>dev_port</code> , in I/O space. The <code>ddi_io_get8()</code> , <code>ddi_io_get16()</code> , and <code>ddi_io_get32()</code> functions read 8 bits, 16 bits, and 32 bits of data, respectively, from the device port, <code>dev_port</code> .</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<code>isa(4)</code> , <code>ddi_io_put8(9F)</code> , <code>ddi_io_rep_get8(9F)</code> , <code>ddi_io_rep_put8(9F)</code> , <code>ddi_regs_map_free(9F)</code> , <code>ddi_regs_map_setup(9F)</code> , <code>ddi_device_acc_attr(9S)</code>
<b>NOTES</b>	<p>For drivers using these functions, it may not be easy to maintain a single source to support devices with multiple bus versions. For example, devices may offer I/O space in ISA bus (see <code>isa(4)</code> )but memory space only in PCI local bus. This is especially true in instruction set architectures such as x86 where accesses to the memory and I/O space are different.</p> <p>The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:</p>

Previous Name	New Name
ddi_io_getb	ddi_io_get8
ddi_io_getw	ddi_io_get16
ddi_io_getl	ddi_io_get32

<b>NAME</b>	ddi_iomin – find minimum alignment and transfer size for DMA
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>int ddi_iomin(dev_info_t *dip, int initial, int streaming);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                   A pointer to the device's dev_info structure.</p> <p><b>initial</b>                The initial minimum DMA transfer size in bytes. This may be zero or an appropriate dlim_minxfer value for device's ddi_dma_lim structure (see ddi_dma_lim_sparc(9S) or ddi_dma_lim_x86(9S)). This value must be a power of two.</p> <p><b>streaming</b>             This argument, if non-zero, indicates that the returned value should be modified to account for streaming mode accesses (see ddi_dma_req(9S) for a discussion of streaming versus non-streaming access mode).</p>
<b>DESCRIPTION</b>	<b>ddi_iomin()</b> , finds out the minimum DMA transfer size for the device pointed to by <i>dip</i> . This provides a mechanism by which a driver can determine the effects of underlying caches as well as intervening bus adapters on the granularity of a DMA transfer.
<b>RETURN VALUES</b>	<b>ddi_iomin()</b> returns the minimum DMA transfer size for the calling device, or it returns zero, which means that you cannot get there from here.
<b>CONTEXT</b>	This function can be called from user or interrupt context.
<b>SEE ALSO</b>	ddi_dma_devalign(9F), ddi_dma_setup(9F), ddi_dma_sync(9F), ddi_dma_lim_sparc(9S), ddi_dma_lim_x86(9S), ddi_dma_req(9S)
	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_iopb_alloc, ddi_iopb_free – allocate and free non-sequentially accessed memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_iopb_alloc(dev_info_t * dip, ddi_dma_lim_t * limits, uint_t length, caddr_t * iopbp);  void ddi_iopb_free(caddr_t iopbp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
<b>ddi_iopb_alloc()</b>	<p><b>dip</b>            A pointer to the device's dev_info structure.</p> <p><b>limits</b>        A pointer to a DMA limits structure for this device (see ddi_dma_lim_sparc(9S) or ddi_dma_lim_x86(9S) ). If this pointer is NULL , a default set of DMA limits is assumed.</p> <p><b>length</b>        The length in bytes of the desired allocation.</p> <p><b>iopbp</b>         A pointer to a caddr_t . On a successful return, *iopbp points to the allocated storage.</p>
<b>ddi_iopb_free()</b>	<p><b>iopbp</b>         The iopbp returned from a successful call to ddi_iopb_alloc() .</p>
<b>DESCRIPTION</b>	<p><b>ddi_iopb_alloc()</b> allocates memory for DMA transfers and should be used if the device accesses memory in a non-sequential fashion, or if synchronization steps using <b>ddi_dma_sync(9F)</b> should be as lightweight as possible, due to frequent use on small objects. This type of access is commonly known as <i>consistent</i> access. The allocation will obey the alignment and padding constraints as specified in the <i>limits</i> argument and other limits imposed by the system.</p> <p>Note that you still must use DMA resource allocation functions (see <b>ddi_dma_setup(9F)</b> ) to establish DMA resources for the memory allocated using <b>ddi_iopb_alloc()</b> .</p> <p>In order to make the view of a memory object shared between a CPU and a DMA device consistent, explicit synchronization steps using</p>

**ddi\_dma\_sync(9F)** or **ddi\_dma\_free(9F)** are still required. The DMA resources will be allocated so that these synchronization steps are as efficient as possible.

**ddi\_iopb\_free()** frees up memory allocated by **ddi\_iopb\_alloc()** .

**RETURN VALUES**

**ddi\_iopb\_alloc()** returns:

DDI\_SUCCESS    Memory successfully allocated.

DDI\_FAILURE    Allocation failed.

**CONTEXT**

These functions can be called from user or interrupt context.

**SEE ALSO**

**ddi\_dma\_free(9F)** , **ddi\_dma\_setup(9F)** , **ddi\_dma\_sync(9F)** ,  
**ddi\_mem\_alloc(9F)** , **ddi\_dma\_lim\_sparc(9S)** , **ddi\_dma\_lim\_x86(9S)** ,  
**ddi\_dma\_req(9S)**

*Writing Device Drivers*

**NOTES**

This function uses scarce system resources. Use it selectively.

<b>NAME</b>	ddi_io_put8, ddi_io_put16, ddi_io_put32, ddi_io_putw, ddi_io_putl, ddi_io_putb – write data to the mapped device register in I/O space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_io_put8(ddi_acc_handle_t handle, int dev_port, uint8_t value); void ddi_io_put16(ddi_acc_handle_t handle, int dev_port, uint16_t value); void ddi_io_put32(ddi_acc_handle_t handle, int dev_port, uint32_t value);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The data access handle returned from setup calls, such as <b>ddi_regs_map_setup(9F)</b> .</p> <p><b>dev_port</b>         The device port.</p> <p><b>value</b>             The data to be written to the device.</p>
<b>DESCRIPTION</b>	<p>These routines generate a write of various sizes to the device port, <i>dev_port</i> , in I/O space. The <b>ddi_io_put8()</b> , <b>ddi_io_put16()</b> , and <b>ddi_io_put32()</b> functions write 8 bits, 16 bits, and 32 bits of data, respectively, to the device port, <i>dev_port</i> .</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<b>isa(4)</b> , <b>ddi_io_get8(9F)</b> , <b>ddi_io_rep_get8(9F)</b> , <b>ddi_io_rep_put8(9F)</b> , <b>ddi_regs_map_setup(9F)</b> , <b>ddi_device_acc_attr(9S)</b>
<b>NOTES</b>	<p>For drivers using these functions, it may not be easy to maintain a single source to support devices with multiple bus versions. For example, devices may offer I/O space in ISA bus (see <b>isa(4)</b> )but memory space only in PCI local bus. This is especially true in instruction set architectures such as x86 where accesses to the memory and I/O space are different.</p> <p>The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed</p>

so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
ddi_io_putb	ddi_io_put8
ddi_io_putw	ddi_io_put16
ddi_io_putl	ddi_io_put32

<b>NAME</b>	ddi_io_rep_get8, ddi_io_rep_get16, ddi_io_rep_get32, ddi_io_rep_getw, ddi_io_rep_getb, ddi_io_rep_getl – read multiple data from the mapped device register in I/O space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_io_rep_get8(ddi_acc_handle_t handle, uint8_t * host_addr, int dev_port, size_t repcount);  void ddi_io_rep_get16(ddi_acc_handle_t handle, uint16_t * host_addr, int dev_port, size_t repcount);  void ddi_io_rep_get32(ddi_acc_handle_t handle, uint32_t * host_addr, int dev_port, size_t repcount);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The data access handle returned from setup calls, such as <b>ddi_regs_map_setup(9F)</b> .</p> <p><b>host_addr</b>        Base host address.</p> <p><b>dev_port</b>         The device port.</p> <p><b>repcount</b>         Number of data accesses to perform.</p>
<b>DESCRIPTION</b>	<p>These routines generate multiple reads from the device port, <i>dev_port</i> , in I/O space. <i>repcount</i> data is copied from the device port, <i>dev_port</i> , to the host address, <i>host_addr</i> . For each input datum, the <b>ddi_io_rep_get8()</b> , <b>ddi_io_rep_get16()</b> , and <b>ddi_io_rep_get32()</b> functions read 8 bits, 16 bits, and 32 bits of data, respectively, from the device port. <i>host_addr</i> must be aligned to the datum boundary described by the function.</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<b>isa(4)</b> , <b>ddi_io_get8(9F)</b> , <b>ddi_io_put8(9F)</b> , <b>ddi_io_rep_put8(9F)</b> , <b>ddi_regs_map_free(9F)</b> , <b>ddi_regs_map_setup(9F)</b> , <b>ddi_device_acc_attr(9S)</b>

**NOTES**

For drivers using these functions, it may not be easy to maintain a single source to support devices with multiple bus versions. For example, devices may offer I/O space in ISA bus (see `isa(4)`) but memory space only in PCI local bus. This is especially true in instruction set architectures such as x86 where accesses to the memory and I/O space are different.

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
<code>ddi_io_rep_getb</code>	<code>ddi_io_rep_get8</code>
<code>ddi_io_rep_getw</code>	<code>ddi_io_rep_get16</code>
<code>ddi_io_rep_getl</code>	<code>ddi_io_rep_get32</code>

<b>NAME</b>	ddi_io_rep_put8, ddi_io_rep_put16, ddi_io_rep_put32, ddi_io_rep_putw, ddi_io_rep_putl, ddi_io_rep_putb - write multiple data to the mapped device register in I/O space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_io_rep_put8(ddi_acc_handle_t handle, uint8_t * host_addr, int dev_port, size_t repcount);  void ddi_io_rep_put16(ddi_acc_handle_t handle, uint16_t * host_addr, int dev_port, size_t repcount);  void ddi_io_rep_put32(ddi_acc_handle_t handle, uint32_t * host_addr, int dev_port, size_t repcount);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The data access handle returned from setup calls, such as ddi_regs_map_setup(9F) .</p> <p><b>host_addr</b>        Base host address.</p> <p><b>dev_port</b>         The device port.</p> <p><b>repcount</b>         Number of data accesses to perform.</p>
<b>DESCRIPTION</b>	<p>These routines generate multiple writes to the device port, <i>dev_port</i> , in I/O space. <i>repcount</i> data is copied from the host address, <i>host_addr</i> , to the device port, <i>dev_port</i> . For each input datum, the <b>ddi_io_rep_put8()</b> , <b>ddi_io_rep_put16()</b> , and <b>ddi_io_rep_put32()</b> functions write 8 bits, 16 bits, and 32 bits of data, respectively, to the device port. <i>host_addr</i> must be aligned to the datum boundary described by the function.</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	isa(4) , ddi_io_get8(9F) , ddi_io_put8(9F) , ddi_io_rep_get8(9F) , ddi_regs_map_setup(9F) , ddi_device_acc_attr(9S)

**NOTES**

For drivers using these functions, it may not be easy to maintain a single source to support devices with multiple bus versions. For example, devices may offer I/O space in ISA bus (see `isa(4)`) but memory space only in PCI local bus. This is especially true in instruction set architectures such as x86 where accesses to the memory and I/O space are different.

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
<code>ddi_io_rep_putb</code>	<code>ddi_io_rep_put8</code>
<code>ddi_io_rep_putw</code>	<code>ddi_io_rep_put16</code>
<code>ddi_io_rep_putl</code>	<code>ddi_io_rep_put32</code>

<b>NAME</b>	ddi_mapdev – create driver-controlled mapping of device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int <b>ddi_mapdev</b>(dev_t <i>dev</i>, off_t <i>offset</i>, struct as *<i>asp</i>, caddr_t *<i>addrp</i>, off_t <i>len</i>, uint_t <i>prot</i>, uint_t <i>maxprot</i>, uint_t <i>flags</i>, cred_t *<i>cred</i>, struct ddi_mapdev_ctl *<i>ctl</i>, ddi_mapdev_handle_t *<i>handlep</i>, void *<i>devprivate</i>);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>dev</i></b>                   The device whose memory is to be mapped.</p> <p><b><i>offset</i></b>                 The offset within device memory at which the mapping begins.</p> <p><b><i>as</i></b>                     An opaque pointer to the user address space into which the device memory should be mapped.</p> <p><b><i>addrp</i></b>                 Pointer to the starting address within the user address space to which the device memory should be mapped.</p> <p><b><i>len</i></b>                    Length (in bytes) of the memory to be mapped.</p> <p><b><i>prot</i></b>                   A bit field that specifies the protections.</p> <p><b><i>maxprot</i></b>               Maximum protection flag possible for attempted mapping.</p> <p><b><i>flags</i></b>                 Flags indicating type of mapping.</p> <p><b><i>cred</i></b>                  Pointer to the user credentials structure.</p> <p><b><i>ctl</i></b>                    A pointer to a <b>ddi_mapdev_ctl(9S)</b> structure. The structure contains pointers to device driver-supplied functions that manage events on the device mapping.</p> <p><b><i>handlep</i></b>               An opaque pointer to a device mapping handle. A handle to the new device mapping is generated and placed into the location pointed to by <i>handlep</i>. If the call fails, the value of <i>handlep</i> is undefined.</p> <p><b><i>devprivate</i></b>           Driver private mapping data. This value is passed into each mapping call back routine.</p>

**DESCRIPTION**

Future releases of Solaris will provide this function for binary and source compatibility. However, for increased functionality, use `devmap_setup(9F)` instead. See `devmap_setup(9F)` for details.

`ddi_mapdev()` sets up user mappings to device space. The driver is notified of user events on the mappings via the entry points defined by *ctl*.

The user events that the driver is notified of are:

- access**            User has accessed an address in the mapping that has no translations.
- duplication**     User has duplicated the mapping. Mappings are duplicated when the process calls `fork(2)`.
- unmapping**        User has called `munmap(2)` on the mapping or is exiting. See `mapdev_access(9E)`, `mapdev_dup(9E)`, and `mapdev_free(9E)` for details on these entry points.

The range to be mapped, defined by *offset* and *len* must be valid.

The arguments *dev*, *asp*, *addrp*, *len*, *prot*, *maxprot*, *flags*, and *cred* are provided by the `segmap(9E)` entry point and should not be modified. See `segmap(9E)` for a description of these arguments. Unlike `ddi_segmap(9F)`, the drivers `mmap(9E)` entry point is not called to verify the range to be mapped.

With the handle, device drivers can use `ddi_mapdev_intercept(9F)` and `ddi_mapdev_nointercept(9F)` to inform the system of whether or not they are interested in being notified when the user process accesses the mapping. By default, user accesses to newly created mappings will generate a call to the `mapdev_access()` entry point. The driver is always notified of duplications and unmaps.

The device may also use the handle to assign certain characteristics to the mapping. See `ddi_mapdev_set_device_acc_attr(9F)` for details.

The device driver can use these interfaces to implement a device context and control user accesses to the device space. `ddi_mapdev()` is typically called from the `segmap(9E)` entry point.

**RETURN VALUES**

`ddi_mapdev()` returns zero on success and non-zero on failure. The return value from `ddi_mapdev()` should be used as the return value for the drivers `segmap()` entry point.

**CONTEXT**

This routine can be called from user or kernel context only.

**SEE ALSO**

`fork(2)`, `mmap(2)`, `munmap(2)`, `mapdev_access(9E)`, `mapdev_dup(9E)`, `mapdev_free(9E)`, `mmap(9E)`, `segmap(9E)`, `ddi_mapdev_intercept(9F)`, `ddi_mapdev_nointercept(9F)`,

`ddi_mapdev_set_device_acc_attr(9F)`, `ddi_segmap(9F)`,  
`ddi_mapdev_ctl(9S)`

*Writing Device Drivers*

**NOTES**

Only mappings of type `MAP_PRIVATE` should be used with `ddi_mapdev()`.

<b>NAME</b>	ddi_mapdev_intercept, ddi_mapdev_nointercept – control driver notification of user accesses
<b>SYNOPSIS</b>	<pre>#include &lt;sys/sunddi.h&gt;  int ddi_mapdev_intercept(ddi_mapdev_handle_t handle, off_t offset, off_t len); int ddi_mapdev_nointercept(ddi_mapdev_handle_t handle, off_t offset, off_t len);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            An opaque pointer to a device mapping handle.</p> <p><b>offset</b>            An offset in bytes within device memory.</p> <p><b>len</b>                Length in bytes.</p>
<b>DESCRIPTION</b>	<p>Future releases of Solaris will provide these functions for binary and source compatibility. However, for increased functionality, use <code>devmap_load(9F)</code> or <code>devmap_unload(9F)</code> instead. See <code>devmap_load(9F)</code> and <code>devmap_unload(9F)</code> for details.</p> <p>The <code>ddi_mapdev_intercept()</code> and <code>ddi_mapdev_nointercept()</code> functions control whether or not user accesses to device mappings created by <code>ddi_mapdev(9F)</code> in the specified range will generate calls to the <code>mapdev_access(9E)</code> entry point. <code>ddi_mapdev_intercept()</code> tells the system to intercept the user access and notify the driver to invalidate the mapping translations. <code>ddi_mapdev_nointercept()</code> tells the system to not intercept the user access and allow it to proceed by validating the mapping translations.</p> <p>For both routines, the range to be affected is defined by the <code>offset</code> and <code>len</code> arguments. Requests affect the entire page containing the <code>offset</code> and all pages up to and including the page containing the last byte as indicated by <code>offset + len</code>.</p> <p>Supplying a value of 0 for the <code>len</code> argument affects all addresses from the <code>offset</code> to the end of the mapping. Supplying a value of 0 for the <code>offset</code> argument and a value of 0 for <code>len</code> argument affect all addresses in the mapping.</p> <p>To manage a device context, a device driver would call <code>ddi_mapdev_intercept()</code> on the context about to be switched out, switch contexts, and then call <code>ddi_mapdev_nointercept()</code> on the context switched in.</p>
<b>RETURN VALUES</b>	<p><code>ddi_mapdev_intercept()</code> and <code>ddi_mapdev_nointercept()</code> return the following values:</p> <p>0                    Successful completion.</p>

**Non-zero**            An error occurred.

**EXAMPLES**

**EXAMPLE 1**    managing a device context that is one page in length

The following shows an example of managing a device context that is one page in length.

```

ddi_mapdev_handle_t cur_hdl;
static int
xxmapdev_access(ddi_mapdev_handle_t handle, void *devprivate,
                off_t offset)
{
    \011int err;
    \011/* enable access callbacks for the current mapping */
    \011if (cur_hdl != NULL) {
    \011\011if ((err = ddi_mapdev_intercept(cur_hdl, offset, 0)) != 0)
    \011\011\011return (err);
    \011}
    \011/* Switch device context - device dependent*/
    \011...
    \011/* Make handle the new current mapping */
    \011cur_hdl = handle;
    \011/*
    \011 * Disable callbacks and complete the access for the
    \011 * mapping that generated this callback.
    \011 */
    \011return (ddi_mapdev_nointercept(handle, offset, 0));
}

```

**CONTEXT**

These routines can be called from user or kernel context only.

**SEE ALSO**

`mapdev_access(9E)` , `ddi_mapdev(9F)`

*Writing Device Drivers*

<b>NAME</b>	ddi_mapdev_set_device_acc_attr – set the device attributes for the mapping
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_mapdev_set_device_acc_attr(ddi_mapdev_handle_t mapping_handle, off_t offset, off_t len, ddi_device_acc_attr_t *accattrp, uint_t rnumber);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>mapping_handle</b> A pointer to a device mapping handle.</p> <p><b>offset</b> The offset within device memory to which the device access attributes structure applies.</p> <p><b>len</b> Length (in bytes) of the memory to which the device access attributes structure applies.</p> <p><b>*accattrp</b> Pointer to a <b>ddi_device_acc_attr(9S)</b> structure. Contains the device access attributes to be applied to this range of memory.</p> <p><b>rnumber</b> Index number to the register address space set.</p>
<b>DESCRIPTION</b>	<p>Future releases of Solaris will provide this function for binary and source compatibility. However, for increased functionality, use <b>devmap(9E)</b> instead. See <b>devmap(9E)</b> for details.</p> <p>The <b>ddi_mapdev_set_device_acc_attr()</b> function assigns device access attributes to a range of device memory in the register set given by <b>rnumber</b>.</p> <p><b>*accattrp</b> defines the device access attributes. See <b>ddi_device_acc_attr(9S)</b> for more details.</p> <p><b>mapping_handle</b> is a mapping handle returned from a call to <b>ddi_mapdev(9F)</b>.</p> <p>The range to be affected is defined by the <b>offset</b> and <b>len</b> arguments. Requests affect the entire page containing the <b>offset</b> and all pages up to and including the page containing the last byte as indicated by <b>offset+len</b>. Supplying a value of 0 for the <b>len</b> argument affects all addresses from the <b>offset</b> to the end of the mapping. Supplying a value of 0 for the <b>offset</b> argument and a value of 0 for the <b>len</b> argument affect all addresses in the mapping.</p>
<b>RETURN VALUES</b>	<p>The <b>ddi_mapdev_set_device_acc_attr()</b> function returns the following values:</p> <p>DDI_SUCCESS                      The attributes were successfully set.</p>

DDI\_FAILURE

It is not possible to set these attributes for this mapping handle.

**CONTEXT**

This routine can be called from user or kernel context only.

**SEE ALSO**

`segmap(9E)`, `ddi_mapdev(9F)`, `ddi_segmap_setup(9F)`,  
`ddi_device_acc_attr(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_map_regs, ddi_unmap_regs – map or unmap registers
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_map_regs(dev_info_t * dip, uint_t rnumber, caddr_t * kaddrp, off_t offset, off_t len);  void ddi_unmap_regs(dev_info_t * dip, uint_t rnumber, caddr_t * kaddrp, off_t offset, off_t len);</pre>
<b>PARAMETERS</b>	
<b>ddi_map_regs()</b>	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>rnumber</b> Register set number.</p> <p><b>kaddrp</b> Pointer to the base kernel address of the mapped region (set on return).</p> <p><b>offset</b> Offset into register space.</p> <p><b>len</b> Length to be mapped.</p>
<b>ddi_unmap_regs()</b>	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>rnumber</b> Register set number.</p> <p><b>kaddrp</b> Pointer to the base kernel address of the region to be unmapped.</p> <p><b>offset</b> Offset into register space.</p> <p><b>len</b> Length to be unmapped.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p><b>ddi_map_regs()</b> maps in the register set given by <i>rnumber</i> . The register number determines which register set will be mapped if more than one exists. The base kernel virtual address of the mapped register set is returned in <i>kaddrp</i> . <i>offset</i> specifies an offset into the register space to start from and <i>len</i> indicates the size of the area to be mapped. If <i>len</i> is non-zero, it overrides the length given in the register set description. See the discussion of the <code>reg</code> property in</p>

**sbus(4)** and **vme(4)** for more information on register set descriptions. If *len* and *offset* are 0, the entire space is mapped.

**ddi\_unmap\_regs()** undoes mappings set up by **ddi\_map\_regs()**. This is provided for drivers preparing to detach themselves from the system, allowing them to release allocated mappings. Mappings must be released in the same way they were mapped (a call to **ddi\_unmap\_regs()** must correspond to a previous call to **ddi\_map\_regs()**). Releasing portions of previous mappings is not allowed. *rnumber* determines which register set will be unmapped if more than one exists. The *kaddrp*, *offset* and *len* specify the area to be unmapped. *kaddrp* is a pointer to the address returned from **ddi\_map\_regs()**; *offset* and *len* should match what **ddi\_map\_regs()** was called with.

**RETURN VALUES**

**ddi\_map\_regs()** returns:  
DDI\_SUCCESS on success.

**CONTEXT**

These functions can be called from user or interrupt context.

**SEE ALSO**

**sbus(4)**, **vme(4)**  
*Writing Device Drivers*

<b>NAME</b>	ddi_mem_alloc, ddi_mem_free – allocate and free sequentially accessed memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_mem_alloc(dev_info_t * dip, ddi_dma_lim_t * limits, uint_t length, uint_t flags, caddr_t * kaddrp, uint_t * real_length);  void ddi_mem_free(caddr_t kaddr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
<b>ddi_mem_alloc()</b>	<p><b>dip</b> A pointer to the device's dev_info structure.</p> <p><b>limits</b> A pointer to a DMA limits structure for this device (see ddi_dma_lim_sparc(9S) or ddi_dma_lim_x86(9S) ). If this pointer is NULL, a default set of DMA limits is assumed.</p> <p><b>length</b> The length in bytes of the desired allocation.</p> <p><b>flags</b> The possible flags 1 and 0 are taken to mean, respectively, wait until memory is available, or do not wait.</p> <p><b>kaddrp</b> On a successful return, *kaddrp points to the allocated memory.</p> <p><b>real_length</b> The length in bytes that was allocated. Alignment and padding requirements may cause ddi_mem_alloc() to allocate more memory than requested in length .</p>
<b>ddi_mem_free()</b>	<p><b>kaddr</b> The memory returned from a successful call to ddi_mem_alloc() .</p>
<b>DESCRIPTION</b>	<p>ddi_mem_alloc() allocates memory for DMA transfers and should be used if the device is performing sequential, unidirectional, block-sized and block-aligned transfers to or from memory. This type of access is commonly known as <i>streaming</i> access. The allocation will obey the alignment and padding constraints as specified by the limits argument and other limits imposed by the system.</p>

Note that you must still use DMA resource allocation functions (see `ddi_dma_setup(9F)`) to establish DMA resources for the memory allocated using `ddi_mem_alloc()`. `ddi_mem_alloc()` returns the actual size of the allocated memory object. Because of padding and alignment requirements, the actual size might be larger than the requested size. `ddi_dma_setup(9F)` requires the actual length.

In order to make the view of a memory object shared between a CPU and a DMA device consistent, explicit synchronization steps using `ddi_dma_sync(9F)` or `ddi_dma_free(9F)` are required.

`ddi_mem_free()` frees up memory allocated by `ddi_mem_alloc()`.

**RETURN VALUES**

`ddi_mem_alloc()` returns:

`DDI_SUCCESS` Memory successfully allocated.

`DDI_FAILURE` Allocation failed.

**CONTEXT**

`ddi_mem_alloc()` can be called from user or interrupt context, except when *flags* is set to 1, in which case it can be called from user context only.

**SEE ALSO**

`ddi_dma_free(9F)`, `ddi_dma_setup(9F)`, `ddi_dma_sync(9F)`,  
`ddi_iopb_alloc(9F)`, `ddi_dma_lim_sparc(9S)`, `ddi_dma_lim_x86(9S)`,  
`ddi_dma_req(9S)`

*Writing Device Drivers*

<b>NAME</b>	ddi_mem_get8, ddi_mem_get16, ddi_mem_get32, ddi_mem_get64, ddi_mem_getw, ddi_mem_getl, ddi_mem_getll, ddi_mem_getb – read data from mapped device in the memory space or allocated DMA memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  uint8_t ddi_mem_get8(ddi_acc_handle_t handle, uint8_t * dev_addr); uint16_t ddi_mem_get16(ddi_acc_handle_t handle, uint16_t * dev_addr); uint32_t ddi_mem_get32(ddi_acc_handle_t handle, uint32_t * dev_addr); uint64_t ddi_mem_get64(ddi_acc_handle_t handle, uint64_t * dev_addr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The data access handle returned from setup calls, such as ddi_regs_map_setup(9F) .</p> <p><b>dev_addr</b>        Base device address.</p>
<b>DESCRIPTION</b>	<p>These routines generate a read of various sizes from memory space or allocated DMA memory. The <b>ddi_mem_get8()</b> , <b>ddi_mem_get16()</b> , <b>ddi_mem_get32()</b> , and <b>ddi_mem_get64()</b> functions read 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, from the device address, <i>dev_addr</i> , in memory space.</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	ddi_mem_put8(9F) , ddi_mem_rep_get8(9F) , ddi_mem_rep_put8(9F) , ddi_regs_map_setup(9F) , ddi_device_acc_attr(9S)
<b>NOTES</b>	The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
ddi_mem_getb	ddi_mem_get8
ddi_mem_getw	ddi_mem_get16
ddi_mem_getl	ddi_mem_get32
ddi_mem_getll	ddi_mem_get64

<b>NAME</b>	ddi_mem_put8, ddi_mem_put16, ddi_mem_put32, ddi_mem_put64, ddi_mem_putb, ddi_mem_putw, ddi_mem_putl, ddi_mem_putll – write data to mapped device in the memory space or allocated DMA memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_mem_put8(ddi_acc_handle_t handle, uint8_t * dev_addr, uint8_t value);  void ddi_mem_put16(ddi_acc_handle_t handle, uint16_t * dev_addr, uint16_t value);  void ddi_mem_put32(ddi_acc_handle_t handle, uint32_t * dev_addr, uint32_t value);  void ddi_mem_put64(ddi_acc_handle_t handle, uint64_t * dev_addr, uint64_t value);</pre>
<b>PARAMETERS</b>	<p><b>handle</b>            The data access handle returned from setup calls, such as ddi_regs_map_setup(9F) .</p> <p><b>dev_addr</b>        Base device address.</p> <p><b>value</b>            The data to be written to the device.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p>These routines generate a write of various sizes to memory space or allocated DMA memory. The ddi_mem_put8() , ddi_mem_put16() , ddi_mem_put32() , and ddi_mem_put64() functions write 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, to the device address, dev_addr , in memory space.</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	ddi_mem_get8(9F) , ddi_mem_rep_get8(9F) , ddi_regs_map_setup(9F) , ddi_device_acc_attr(9S)
<b>NOTES</b>	The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
ddi_mem_putb	ddi_mem_put8
ddi_mem_putw	ddi_mem_put16
ddi_mem_putl	ddi_mem_put32
ddi_mem_putll	ddi_mem_put64

<b>NAME</b>	ddi_mem_rep_get8, ddi_mem_rep_get16, ddi_mem_rep_get32, ddi_mem_rep_get64, ddi_mem_rep_getw, ddi_mem_rep_getl, ddi_mem_rep_getll, ddi_mem_rep_getb – read multiple data from mapped device in the memory space or allocated DMA memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_mem_rep_get8(ddi_acc_handle_t handle, uint8_t * host_addr, uint8_t * dev_addr, size_t recount, uint_t flags);  void ddi_mem_rep_get16(ddi_acc_handle_t handle, uint16_t * host_addr, uint16_t * dev_addr, size_t recount, uint_t flags);  void ddi_mem_rep_get32(ddi_acc_handle_t handle, uint32_t * host_addr, uint32_t * dev_addr, size_t recount, uint_t flags);  void ddi_mem_rep_get64(ddi_acc_handle_t handle, uint64_t * host_addr, uint64_t * dev_addr, size_t recount, uint_t flags);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The data access handle returned from setup calls, such as <b>ddi_regs_map_setup(9F)</b> .</p> <p><b>host_addr</b>       Base host address.</p> <p><b>dev_addr</b>       Base device address.</p> <p><b>recount</b>        Number of data accesses to perform.</p> <p><b>flags</b>           Device address flags:</p> <p>                  DDI_DEV_ADDR_INCREMENT   Automatically increment the device address, <i>dev_addr</i> , during data accesses.</p> <p>                  DDI_DEV_ADDR_INCREMENT   Do not advance the device address, <i>dev_addr</i> , during data accesses.</p>
<b>DESCRIPTION</b>	These routines generate multiple reads from memory space or allocated DMA memory. <i>recount</i> data is copied from the device address, <i>dev_addr</i> , in memory space to the host address, <i>host_addr</i> . For each input datum, the <b>ddi_mem_rep_get8()</b> , <b>ddi_mem_rep_get16()</b> , <b>ddi_mem_rep_get32()</b> , and <b>ddi_mem_rep_get64()</b> functions read 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, from the device address, <i>dev_addr</i> . <i>dev_addr</i> and <i>host_addr</i> must be aligned to the datum boundary described by the function.

Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.

When the *flags* argument is set to `DDI_DEV_AUTOINCR`, these functions will treat the device address, *dev\_addr*, as a memory buffer location on the device and increments its address on the next input datum. However, when the *flags* argument is set to `DDI_DEV_NO_AUTOINCR`, the same device address will be used for every datum access. For example, this flag may be useful when reading from a data register.

**CONTEXT**

These functions can be called from user, kernel, or interrupt context.

**SEE ALSO**

`ddi_mem_get8(9F)`, `ddi_mem_put8(9F)`, `ddi_mem_rep_put8(9F)`, `ddi_regs_map_setup(9F)`, `ddi_device_acc_attr(9S)`

**NOTES**

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
<code>ddi_mem_rep_getb</code>	<code>ddi_mem_rep_get8</code>
<code>ddi_mem_rep_getw</code>	<code>ddi_mem_rep_get16</code>
<code>ddi_mem_rep_getl</code>	<code>ddi_mem_rep_get32</code>
<code>ddi_mem_rep_getll</code>	<code>ddi_mem_rep_get64</code>

<b>NAME</b>	ddi_mem_rep_put8, ddi_mem_rep_put16, ddi_mem_rep_put32, ddi_mem_rep_put64, ddi_mem_rep_putw, ddi_mem_rep_putl, ddi_mem_rep_putll, ddi_mem_rep_putb - write multiple data to mapped device in the memory space or allocated DMA memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_mem_rep_put8(ddi_acc_handle_t handle, uint8_t * host_addr, uint8_t * dev_addr, size_t recount, uint_t flags);  void ddi_mem_rep_put16(ddi_acc_handle_t handle, uint16_t * host_addr, uint16_t * dev_addr, size_t recount, uint_t flags);  void ddi_mem_rep_put32(ddi_acc_handle_t handle, uint32_t * host_addr, uint32_t * dev_addr, size_t recount, uint_t flags);  void ddi_mem_rep_put64(ddi_acc_handle_t handle, uint64_t * host_addr, uint64_t * dev_addr, size_t recount, uint_t flags);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The data access handle returned from setup calls, such as <code>ddi_regs_map_setup(9F)</code> .</p> <p><b>host_addr</b>       Base host address.</p> <p><b>dev_addr</b>        Base device address.</p> <p><b>recount</b>        Number of data accesses to perform.</p> <p><b>flags</b>           Device address flags:</p> <p>                  DDI_DEV_AUTOINCR                   Automatically increment the device address, <i>dev_addr</i> , during data accesses.</p> <p>                  DDI_DEV_NO_AUTOINCR                   Do not advance the device address, <i>dev_addr</i> , during data accesses.</p>
<b>DESCRIPTION</b>	These routines generate multiple writes to memory space or allocated DMA memory. <i>recount</i> data is copied from the host address, <i>host_addr</i> , to the device address, <i>dev_addr</i> , in memory space. For each input datum, the <code>ddi_mem_rep_put8()</code> , <code>ddi_mem_rep_put16()</code> , <code>ddi_mem_rep_put32()</code> , and

**ddi\_mem\_rep\_put64()** functions write 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, to the device address. *dev\_addr* and *host\_addr* must be aligned to the datum boundary described by the function.

Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.

When the *flags* argument is set to `DDI_DEV_AUTOINCR`, these functions will treat the device address, *dev\_addr*, as a memory buffer location on the device and increments its address on the next input datum. However, when the *flags* argument is set to `DDI_DEV_NO_AUTOINCR`, the same device address will be used for every datum access. For example, this flag may be useful when writing from a data register.

**CONTEXT**

These functions can be called from user, kernel, or interrupt context.

**SEE ALSO**

`ddi_mem_get8(9F)`, `ddi_mem_put8(9F)`, `ddi_mem_rep_get8(9F)`, `ddi_regs_map_setup(9F)`, `ddi_device_acc_attr(9S)`

**NOTES**

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
<code>ddi_mem_rep_putb</code>	<code>ddi_mem_rep_put8</code>
<code>ddi_mem_rep_putw</code>	<code>ddi_mem_rep_put16</code>
<code>ddi_mem_rep_putl</code>	<code>ddi_mem_rep_put32</code>
<code>ddi_mem_rep_putll</code>	<code>ddi_mem_rep_put64</code>

<b>NAME</b>	ddi_mmap_get_model – return data model type of current thread						
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  uint_t ddi_mmap_get_model(void);</pre>						
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).						
<b>DESCRIPTION</b>	<p><b>ddi_mmap_get_model()</b> returns the C Language Type Model which the current thread expects. <b>ddi_mmap_get_model()</b> is used in combination with <b>ddi_model_convert_from(9F)</b> in the <b>mmap(9E)</b> driver entry point to determine whether there is a data model mismatch between the current thread and the device driver. The device driver might have to adjust the shape of data structures before exporting them to a user thread which supports a different data model.</p>						
<b>RETURN VALUES</b>	<table border="0"> <tr> <td>DDI_MODEL_ILP32</td> <td>Current thread expects 32-bit (<i>ILP32</i>) semantics.</td> </tr> <tr> <td>DDI_MODEL_LP64</td> <td>Current thread expects 64-bit (<i>LP64</i>) semantics.</td> </tr> <tr> <td>DDI_FAILURE</td> <td>The <b>ddi_mmap_get_model()</b> function was not called from the <b>mmap(9E)</b> entry point.</td> </tr> </table>	DDI_MODEL_ILP32	Current thread expects 32-bit ( <i>ILP32</i> ) semantics.	DDI_MODEL_LP64	Current thread expects 64-bit ( <i>LP64</i> ) semantics.	DDI_FAILURE	The <b>ddi_mmap_get_model()</b> function was not called from the <b>mmap(9E)</b> entry point.
DDI_MODEL_ILP32	Current thread expects 32-bit ( <i>ILP32</i> ) semantics.						
DDI_MODEL_LP64	Current thread expects 64-bit ( <i>LP64</i> ) semantics.						
DDI_FAILURE	The <b>ddi_mmap_get_model()</b> function was not called from the <b>mmap(9E)</b> entry point.						
<b>CONTEXT</b>	The <b>ddi_mmap_get_model()</b> function can only be called from the <b>mmap(9E)</b> driver entry point.						
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> : Using <b>ddi_mmap_get_model()</b></p> <p>The following is an example of the <b>mmap(9E)</b> entry point and how to support 32-bit and 64-bit applications with the same device driver.</p> <pre>struct data32 {     int len;     caddr32_t addr; };  struct data {     int len;     caddr_t addr; };  xxmmap(dev_t dev, off_t off, int prot) {     struct data dtc; /* a local copy for clash resolution */     struct data *dp = (struct data *)shared_area;  #ifdef _MULTI_DATAMODEL     switch (ddi_model_convert_from(ddi_mmap_get_model())) {</pre>						

```
    case DDI_MODEL_ILP32:
    {
        struct data32 *da32p;

        da32p = (struct data32 *)shared_area;
        dp = &dtc;
        dp->len = da32p->len;
        dp->address = da32p->address;
        break;
    }
    case DDI_MODEL_NONE:
        break;
}
#endif /* _MULTI_DATAMODEL */
/* continues along using dp */
...
}
```

**SEE ALSO** [mmap\(9E\)](#), [ddi\\_model\\_convert\\_from\(9F\)](#)  
*Writing Device Drivers*

<b>NAME</b>	ddi_model_convert_from – determine data model type mismatch				
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  uint_t ddi_model_convert_from(uint_t model);</pre>				
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).				
<b>PARAMETERS</b>	<b>model</b> The data model type of the current thread.				
<b>DESCRIPTION</b>	<p><b>ddi_model_convert_from()</b> is used to determine if the current thread uses a different C Language Type Model than the device driver. The 64-bit version of Solaris will require a 64-bit kernel to support both 64-bit and 32-bit user mode programs. The difference between a 32-bit program and a 64-bit program is in its C Language Type Model: a 32-bit program is ILP32 (integer, longs, and pointers are 32-bit) and a 64-bit program is LP64 (longs and pointers are 64-bit). There are a number of driver entry points such as <b>ioctl(9E)</b> and <b>mmap(9E)</b> where it is necessary to identify the C Language Type Model of the user-mode originator of a kernel event. For example any data which flows between programs and the device driver or vice versa need to be identical in format. A 64-bit device driver may need to modify the format of the data before sending it to a 32-bit application. <b>ddi_model_convert_from()</b> is used to determine if data that is passed between the device driver and the application requires reformatting to any non-native data model.</p>				
<b>RETURN VALUES</b>	<table border="0"> <tr> <td style="vertical-align: top;">DDI_MODEL_ILP32</td> <td style="vertical-align: top;">A conversion to/from ILP32 is necessary.</td> </tr> <tr> <td style="vertical-align: top;">DDI_MODEL_NONE</td> <td style="vertical-align: top;">No conversion is necessary. Current thread and driver use the same data model.</td> </tr> </table>	DDI_MODEL_ILP32	A conversion to/from ILP32 is necessary.	DDI_MODEL_NONE	No conversion is necessary. Current thread and driver use the same data model.
DDI_MODEL_ILP32	A conversion to/from ILP32 is necessary.				
DDI_MODEL_NONE	No conversion is necessary. Current thread and driver use the same data model.				
<b>CONTEXT</b>	<b>ddi_model_convert_from()</b> can be called from any context.				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> : Using <b>ddi_model_convert_from()</b> in the <b>ioctl()</b> entry point to support both 32-bit and 64-bit applications.</p> <p>The following is an example how to use <b>ddi_model_convert_from()</b> in the <b>ioctl()</b> entry point to support both 32-bit and 64-bit applications.</p> <pre>struct passargs32 {     int len;     caddr32_t addr;</pre>				

```

};

struct passargs {
    int len;
    caddr_t addr;
};

xxioctl(dev_t dev, int cmd, intptr_t arg, int mode,
        cred_t *credp, int *rvalp) {
    struct passargs pa;

#ifdef _MULTI_DATAMODEL
    switch (ddi_model_convert_from(mode & FMODELS)) {
        case DDI_MODEL_ILP32:
            {
                struct passargs32 pa32;

                ddi_copyin(arg, &pa32, sizeof (struct passargs32), mode);
                pa.len = pa32.len;
                pa.address = pa32.address;
                break;
            }
        case DDI_MODEL_NONE:
            ddi_copyin(arg, &pa, sizeof (struct passargs), mode);
            break;
    }
#else /* _MULTI_DATAMODEL */
    ddi_copyin(arg, &pa, sizeof (struct passargs), mode);
#endif /* _MULTI_DATAMODEL */

    do_ioctl(&pa);
    ....
}

```

**SEE ALSO** [ioctl\(9E\)](#), [mmap\(9E\)](#), [ddi\\_mmap\\_get\\_model\(9F\)](#)

*Writing Device Drivers*

<b>NAME</b>	ddi_node_name – return the devinfo node name
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  char *ddi_node_name(dev_info_t *dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>dip</b> A pointer the device's dev_info structure.
<b>DESCRIPTION</b>	<b>ddi_node_name()</b> returns the device node name contained in the dev_info node pointed to by <i>dip</i> .
<b>RETURN VALUES</b>	<b>ddi_node_name()</b> returns the device node name contained in the dev_info structure.
<b>CONTEXT</b>	<b>ddi_node_name()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>ddi_binding_name(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_peek, ddi_peek8, ddi_peek16, ddi_peek32, ddi_peek64, ddi_peekc, ddi_peeks, ddi_peekl, ddi_peekd – read a value from a location
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  #include &lt;sys/sunddi.h&gt;  int ddi_peek8(dev_info_t *dip, int8_t *addr, int8_t *valuep); int ddi_peek16(dev_info_t *dip, int16_t *addr, int16_t *valuep); int ddi_peek32(dev_info_t *dip, int32_t *addr, int32_t *valuep); int ddi_peek64(dev_info_t *dip, int64_t *addr, int64_t *valuep);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b> A pointer to the device's <code>dev_info</code> structure.</p> <p><b>addr</b> Virtual address of the location to be examined.</p> <p><b>valuep</b> Pointer to a location to hold the result. If a null pointer is specified, then the value read from the location will simply be discarded.</p>
<b>DESCRIPTION</b>	<p>These routines cautiously attempt to read a value from a specified virtual address, and return the value to the caller, using the parent nexus driver to assist in the process where necessary.</p> <p>If the address is not valid, or the value cannot be read without an error occurring, an error code is returned.</p> <p>The routines are most useful when first trying to establish the presence of a device on the system in a driver's <code>probe(9E)</code> or <code>attach(9E)</code> routines.</p>
<b>RETURN VALUES</b>	<p><b>DDI_SUCCESS</b> The value at the given virtual address was successfully read, and if <code>valuep</code> is non-null, <code>*valuep</code> will have been updated.</p> <p><b>DDI_FAILURE</b> An error occurred while trying to read the location. <code>*valuep</code> is unchanged.</p>
<b>CONTEXT</b>	These functions can be called from user or interrupt context.

**EXAMPLES**

**EXAMPLE 1** Checking to see that the status register of a device is mapped into the kernel address space:

```
\011
if (ddi_peek8(dip, csr, (int8_t *)0) != DDI_SUCCESS) {
    cmn_err(CE_WARN, "Status register not mapped");
    return (DDI_FAILURE);
\011}
```

**EXAMPLE 2** Reading and logging the device type of a particular device:

```
int
xx_attach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    ...
    /* map device registers */
    ...
    if (ddi_peek32(dip, id_addr, &id_value) != DDI_SUCCESS) {
        cmn_err(CE_WARN, "%s%d: cannot read device identifier",
            ddi_get_name(dip), ddi_get_instance(dip));
        goto failure;
    } else
        cmn_err(CE_CONT, "!!s%d: device type 0x%x\
",
            ddi_get_name(dip), ddi_get_instance(dip), id_value);
\011    ...
\011    ...

    ddi_report_dev(dip);
    return (DDI_SUCCESS);

failure:
    /* free any resources allocated */
    ...
    return (DDI_FAILURE);
}
```

**SEE ALSO**

**attach(9E)**, **probe(9E)**, **ddi\_poke(9F)**

*Writing Device Drivers*

**NOTES**

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
ddi_peekc	ddi_peek8
ddi_peeks	ddi_peek16
ddi_peekl	ddi_peek32
ddi_peekd	ddi_peek64

<b>NAME</b>	ddi_poke, ddi_poke8, ddi_poke16, ddi_poke32, ddi_poke64, ddi_pokec, ddi_pokes, ddi_pokel, ddi_poked – write a value to a location
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_poke8(dev_info_t *dip, int8_t *addr, int8_t value); int ddi_poke16(dev_info_t *dip, int16_t *addr, int16_t value); int ddi_poke32(dev_info_t *dip, int32_t *addr, int32_t value); int ddi_poke64(dev_info_t *dip, int64_t *addr, int64_t value);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>            A pointer to the device's dev_info structure.</p> <p><b>addr</b>            Virtual address of the location to be written to.</p> <p><b>value</b>           Value to be written to the location.</p>
<b>DESCRIPTION</b>	<p>These routines cautiously attempt to write a value to a specified virtual address, using the parent nexus driver to assist in the process where necessary.</p> <p>If the address is not valid, or the value cannot be written without an error occurring, an error code is returned.</p> <p>These routines are most useful when first trying to establish the presence of a given device on the system in a driver's <b>probe(9E)</b> or <b>attach(9E)</b> routines.</p> <p>On multiprocessing machines these routines can be extremely heavy-weight, so use the <b>ddi_peek(9F)</b> routines instead if possible.</p>
<b>RETURN VALUES</b>	<p>DDI_SUCCESS    The value was successfully written to the given virtual address.</p> <p>DDI_FAILURE    An error occurred while trying to write to the location.</p>
<b>CONTEXT</b>	These functions can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>attach(9E)</b> , <b>probe(9E)</b> , <b>ddi_peek(9F)</b> <i>Writing Device Drivers</i>

**NOTES**

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
ddi_pokec	ddi_poke8
ddi_pokes	ddi_poke16
ddi_pokel	ddi_poke32
ddi_poked	ddi_poke64

<b>NAME</b>	ddi_prop_create, ddi_prop_modify, ddi_prop_remove, ddi_prop_remove_all, ddi_prop_undefine – create, remove, or modify properties for leaf device drivers
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_prop_create(dev_t dev, dev_info_t * dip, int flags, char * name, caddr_t valuep, int length);  int ddi_prop_undefine(dev_t dev, dev_info_t * dip, int flags, char * name);  int ddi_prop_modify(dev_t dev, dev_info_t * dip, int flags, char * name, caddr_t valuep, int length);  int ddi_prop_remove(dev_t dev, dev_info_t * dip, char * name);  void ddi_prop_remove_all(dev_info_t * dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
<b>ddi_prop_create()</b>	<p><b>dev</b> dev_t of the device.</p> <p><b>dip</b> dev_info_t pointer of the device.</p> <p><b>flags</b> flag modifiers. The only possible flag value is DDI_PROP_CANSLEEP: Memory allocation may sleep.</p> <p><b>name</b> name of property.</p> <p><b>valuep</b> pointer to property value.</p> <p><b>length</b> property length.</p>
<b>ddi_prop_undefine()</b>	<p><b>dev</b> dev_t of the device.</p> <p><b>dip</b> dev_info_t pointer of the device.</p> <p><b>flags</b> flag modifiers. The only possible flag value is DDI_PROP_CANSLEEP: Memory allocation may sleep.</p> <p><b>name</b> name of property.</p>

<b>ddi_prop_modify()</b>	<p><b>dev</b>            dev_t of the device.</p> <p><b>dip</b>            dev_info_t pointer of the device.</p> <p><b>flags</b>          flag modifiers. The only possible flag value is DDI_PROP_CANSLEEP: Memory allocation may sleep.</p> <p><b>name</b>          name of property.</p> <p><b>valuep</b>        pointer to property value.</p> <p><b>length</b>        property length.</p>
<b>ddi_prop_remove()</b>	<p><b>dev</b>            dev_t of the device.</p> <p><b>dip</b>            dev_info_t pointer of the device.</p> <p><b>name</b>          name of property.</p>
<b>ddi_prop_remove_all()</b>	<p><b>dip</b>            dev_info_t pointer of the device.</p>
<b>DESCRIPTION</b>	<p>Device drivers have the ability to create and manage their own properties as well as gain access to properties that the system creates on behalf of the driver. A driver uses <b>ddi_getprop(9F)</b> to query whether or not a specific property exists.</p> <p>Property creation is done by creating a new property definition in the driver's property list associated with <i>dip</i> .</p> <p>Property definitions are stacked; they are added to the beginning of the driver's property list when created. Thus, when searched for, the most recent matching property definition will be found and its value will be return to the caller.</p>
<b>ddi_prop_create()</b>	<p><b>ddi_prop_create()</b> adds a property to the device's property list. If the property is not associated with any particular <i>dev</i> but is associated with the physical device itself, then the argument <i>dev</i> should be the special device DDI_DEV_T_NONE . If you do not have a <i>dev</i> for your device (for example during <b>attach(9E)</b> time), you can create one using <b>makedevice(9F)</b> with a major number of DDI_MAJOR_T_UNKNOWN . <b>ddi_prop_create()</b> will then make the correct <i>dev</i> for your device.</p> <p>For boolean properties, you must set <i>length</i> to 0 . For all other properties, the <i>length</i> argument must be set to the number of bytes used by the data structure representing the property being created.</p>

	<p>Note that creating a property involves allocating memory for the property list, the property name and the property value. If <i>flags</i> does not contain <code>DDI_PROP_CANSLEEP</code>, <b>ddi_prop_create()</b> returns <code>DDI_PROP_NO_MEMORY</code> on memory allocation failure or <code>DDI_SUCCESS</code> if the allocation succeeded. If <code>DDI_PROP_CANSLEEP</code> was set, the caller may sleep until memory becomes available.</p>	
<b>ddi_prop_undefine()</b>	<p><b>ddi_prop_undefine()</b> is a special case of property creation where the value of the property is set to undefined. This property has the effect of terminating a property search at the current devinfo node, rather than allowing the search to proceed up to ancestor devinfo nodes. See <b>ddi_prop_op(9F)</b> .</p> <p>Note that undefining properties does involve memory allocation, and therefore, is subject to the same memory allocation constraints as <b>ddi_prop_create()</b> .</p>	
<b>ddi_prop_modify()</b>	<p><b>ddi_prop_modify()</b> modifies the length and the value of a property. If <b>ddi_prop_modify()</b> finds the property in the driver's property list, allocates memory for the property value and returns <code>DDI_PROP_SUCCESS</code> . If the property was not found, the function returns <code>DDI_PROP_NOT_FOUND</code> .</p> <p>Note that modifying properties does involve memory allocation, and therefore, is subject to the same memory allocation constraints as <b>ddi_prop_create()</b> .</p>	
<b>ddi_prop_remove()</b>	<p><b>ddi_prop_remove()</b> unlinks a property from the device's property list. If <b>ddi_prop_remove()</b> finds the property (an exact match of both <i>name</i> and <i>dev</i> ), it unlinks the property, frees its memory, and returns <code>DDI_PROP_SUCCESS</code> , otherwise, it returns <code>DDI_PROP_NOT_FOUND</code> .</p>	
<b>ddi_prop_remove_all()</b>	<p><b>ddi_prop_remove_all()</b> removes the properties of all the <code>dev_t</code> 's associated with the <i>dip</i> . It is called before unloading a driver.</p>	
<b>RETURN VALUES</b>		
<b>ddi_prop_create()</b>	<code>DDI_PROP_SUCCESS</code>	on success.
	<code>DDI_PROP_NO_MEMORY</code>	on memory allocation failure.
	<code>DDI_PROP_INVALID_ARG</code>	if an attempt is made to create a property with <i>dev</i> equal to <code>DDI_DEV_T_ANY</code> or if <i>name</i> is <code>NULL</code> or <i>name</i> is the <code>NULL</code> string.
<b>ddi_prop_undefine()</b>	<code>DDI_PROP_SUCCESS</code>	on success.
	<code>DDI_PROP_NO_MEMORY</code>	on memory allocation failure.

	DDI_PROP_INVALID_ARG	if an attempt is made to create a property with <i>dev</i> DDI_DEV_T_ANY or if <i>name</i> is NULL or <i>name</i> is the NULL string.
<b>ddi_prop_modify()</b>	DDI_PROP_SUCCESS	on success.
	DDI_PROP_NO_MEMORY	on memory allocation failure.
	DDI_PROP_INVALID_ARG	if an attempt is made to create a property with <i>dev</i> equal to DDI_DEV_T_ANY or if <i>name</i> is NULL or <i>name</i> is the NULL string.
	DDI_PROP_NOT_FOUND	on property search failure.
<b>ddi_prop_remove()</b>	DDI_PROP_SUCCESS	on success.
	DDI_PROP_INVALID_ARG	if an attempt is made to create a property with <i>dev</i> equal to DDI_DEV_T_ANY or if <i>name</i> is NULL or <i>name</i> is the NULL string.
	DDI_PROP_NOT_FOUND	on property search failure.
<b>CONTEXT</b>	If DDI_PROP_CANSLEEP is set, these functions can only be called from user context; otherwise, they can be called from interrupt or user context.	
<b>EXAMPLES</b>	<b>EXAMPLE 1</b> : Creating a property	
	The following example creates a property called <i>nblocks</i> for each partition on a disk.	
	<pre> for (minor = 0; minor &lt; 8; minor++) { \011    (void) ddi_prop_create(makedevice(DDI_MAJOR_T_UNKNOWN, minor), \011        dev, DDI_PROP_CANSLEEP, "nblocks", 8192, sizeof(int)); \011    ... } </pre>	
<b>SEE ALSO</b>	<b>driver.conf(4)</b> , <b>attach(9E)</b> , <b>ddi_getproplen(9F)</b> , <b>ddi_prop_op(9F)</b> , <b>makedevice(9F)</b>	

ddi\_prop\_create(9F)

Kernel Functions for Drivers

*Writing Device Drivers*

<b>NAME</b>	ddi_prop_exists – check for the existence of a property
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int ddi_prop_exists(dev_t match_dev, dev_info_t *dip, uint_t flags, char *name);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>match_dev</b> Device number associated with property or DDI_DEV_T_ANY.</p> <p><b>dip</b> Pointer to the device info node of device whose property list should be searched.</p> <p><b>flags</b> Possible flag values are some combination of:</p> <p>DDI_PROP_DONTPASS Do not pass request to parent device information node if the property is not found.</p> <p>DDI_PROP_NOTPROM Do not look at PROM properties (ignored on platforms that do not support PROM properties).</p> <p><b>name</b> String containing the name of the property.</p>
<b>DESCRIPTION</b>	<p><b>ddi_prop_exists()</b> checks for the existence of a property regardless of the property value data type.</p> <p>Properties are searched for based on the <i>dip</i>, <i>name</i>, and <i>match_dev</i>. The property search order is as follows:</p> <ol style="list-style-type: none"> <li>1. Search software properties created by the driver.</li> <li>2. Search the software properties created by the system (or nexus nodes in the device info tree).</li> <li>3. Search the driver global properties list.</li> <li>4. If DDI_PROP_NOTPROM is not set, search the PROM properties (if they exist).</li> <li>5. If DDI_PROP_DONTPASS is not set, pass this request to the parent device information node.</li> <li>6. Return 0 if not found and 1 if found.</li> </ol>

Usually, the *match\_dev* argument should be set to the actual device number that this property is associated with. However, if the *match\_dev* argument is `DDI_DEV_T_ANY`, then **ddi\_prop\_exists()** will match the request regardless of the *match\_dev* the property was created with. That is the first property whose name matches *name* will be returned. If a property was created with *match\_dev* set to `DDI_DEV_T_NONE` then the only way to look up this property is with a *match\_dev* set to `DDI_DEV_T_ANY`. PROM properties are always created with *match\_dev* set to `DDI_DEV_T_NONE`.

*name* must always be set to the name of the property being looked up.

**RETURN VALUES**

**ddi\_prop\_exists()** returns 1 if the property exists and 0 otherwise.

**CONTEXT**

These functions can be called from user or kernel context.

**EXAMPLES**

**EXAMPLE 1** : Using **ddi\_prop\_exists()**

The following example demonstrates the use of **ddi\_prop\_exists()**.

```
/*
 * Enable "whizzy" mode if the "whizzy-mode" property exists
 */
if (ddi_prop_exists(xx_dev, xx_dip, DDI_PROP_NOTPROM,
    "whizzy-mode") == 1) {
    xx_enable_whizzy_mode(xx_dip);
} else {
    xx_disable_whizzy_mode(xx_dip);
}
```

**SEE ALSO**

**ddi\_prop\_get\_int(9F)**, **ddi\_prop\_lookup(9F)**, **ddi\_prop\_remove(9F)**,  
**ddi\_prop\_update(9F)**

*Writing Device Drivers*

<b>NAME</b>	ddi_prop_get_int – lookup integer property
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int ddi_prop_get_int(dev_t <i>match_dev</i>, dev_info_t *<i>dip</i>, uint_t <i>flags</i>, char *<i>name</i>, int <i>defvalue</i>);</p>
<b>PARAMETERS</b>	<p><b><i>match_dev</i></b> Device number associated with property or DDI_DEV_T_ANY.</p> <p><b><i>dip</i></b> Pointer to the device info node of device whose property list should be searched.</p> <p><b><i>flags</i></b> Possible flag values are some combination of:</p> <p>DDI_PROP_DONTPASS Do not pass request to parent device information node if property not found.</p> <p>DDI_PROP_NOTPROM Do not look at PROM properties (ignored on platforms that do not support PROM properties).</p> <p><b><i>name</i></b> String containing the name of the property.</p> <p><b><i>defvalue</i></b> An integer value that is returned if the property cannot be found.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p><b>ddi_prop_get_int()</b> searches for an integer property and, if found, returns the value of the property.</p> <p>Properties are searched for based on the <i>dip</i>, <i>name</i>, <i>match_dev</i>, and the type of the data (integer). The property search order is as follows:</p> <ol style="list-style-type: none"> <li>1. Search software properties created by the driver.</li> <li>2. Search the software properties created by the system (or nexus nodes in the device info tree).</li> <li>3. Search the driver global properties list.</li> <li>4. If DDI_PROP_NOTPROM is not set, search the PROM properties (if they exist).</li> </ol>

5. If DDI\_PROP\_DONTPASS is not set, pass this request to the parent device information node.
6. Return DDI\_PROP\_NOT\_FOUND.

Usually, the *match\_dev* argument should be set to the actual device number that this property is associated with. However, if the *match\_dev* argument is DDI\_DEV\_T\_ANY, then **ddi\_prop\_get\_int()** will match the request regardless of the *match\_dev* the property was created with. If a property was created with *match\_dev* set to DDI\_DEV\_T\_NONE, then the only way to look up this property is with a *match\_dev* set to DDI\_DEV\_T\_ANY. PROM properties are always created with *match\_dev* set to DDI\_DEV\_T\_NONE.

*name* must always be set to the name of the property being looked up.

The return value of the routine is the value of the property. If the property is not found, the argument *defvalue* is returned as the value of the property.

**RETURN VALUES**

**ddi\_prop\_get\_int()** returns the value of the property. If the property is not found, the argument *defvalue* is returned.

**CONTEXT**

**ddi\_prop\_get\_int()** can be called from user or kernel context.

**EXAMPLES**

**EXAMPLE 1** : Using **ddi\_prop\_get\_int()**

The following example demonstrates the use of **ddi\_prop\_get\_int()**.

```
/*
 * Get the value of the integer "width" property, using
 * our own default if no such property exists
 */
width = ddi_prop_get_int(xx_dev, xx_dip, 0, "width",
    XX_DEFAULT_WIDTH);
```

**SEE ALSO**

**ddi\_prop\_exists(9F)**, **ddi\_prop\_lookup(9F)**, **ddi\_prop\_remove(9F)**,  
**ddi\_prop\_update(9F)**

*Writing Device Drivers*

<b>NAME</b>	ddi_prop_lookup, ddi_prop_lookup_int_array, ddi_prop_lookup_string_array, ddi_prop_lookup_string, ddi_prop_lookup_byte_array, ddi_prop_free – look up property information
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_prop_lookup_int_array(dev_t match_dev, dev_info_t * dip, uint_t flags, char * name, int ** datap, uint_t * nelements);  int ddi_prop_lookup_string_array(dev_t match_dev, dev_info_t * dip, uint_t flags, char * name, char *** datap, uint_t * nelements);  int ddi_prop_lookup_string(dev_t match_dev, dev_info_t * dip, uint_t flags, char * name, char ** datap);  int ddi_prop_lookup_byte_array(dev_t match_dev, dev_info_t * dip, uint_t flags, char * name, uchar_t ** datap, uint_t * nelements);  void ddi_prop_free(void * data);</pre>
<b>PARAMETERS</b>	<p><b>match_dev</b> Device number associated with property or DDI_DEV_T_ANY.</p> <p><b>dip</b> Pointer to the device info node of device whose property list should be searched.</p> <p><b>flags</b> Possible flag values are some combination of:</p> <p style="padding-left: 20px;">DDI_PROP_DONTPASS Do not pass request to parent device information node if the property is not found.</p> <p style="padding-left: 20px;">DDI_PROP_NOTPROM Do not look at PROM properties (ignored on platforms that do not support PROM properties).</p> <p><b>name</b> String containing the name of the property.</p> <p><b>nelements</b> The address of an unsigned integer which, upon successful return, will contain the number of elements accounted for in the memory pointed at by <i>datap</i>. The elements are either integers, strings or bytes depending on the interface used.</p>

	<p><b>datap</b></p> <p><b>ddi_prop_lookup_int_array()</b> The address of a pointer to an array of integers which, upon successful return, will point to memory containing the integer array property value.</p> <p><b>ddi_prop_lookup_string_array()</b> The address of a pointer to an array of strings which, upon successful return, will point to memory containing the array of strings. The array of strings is formatted as an array of pointers to NULL terminated strings, much like the <i>argv</i> argument to <code>execve(2)</code>.</p> <p><b>ddi_prop_lookup_string()</b> The address of a pointer to a string which, upon successful return, will point to memory containing the NULL terminated string value of the property.</p> <p><b>ddi_prop_lookup_byte_array()</b> The address of pointer to an array of bytes which, upon successful return, will point to memory containing the byte array value of the property.</p>
<p><b>INTERFACE LEVEL DESCRIPTION</b></p>	<p>Solaris DDI specific (Solaris DDI).</p> <p>The property look up routines search for and, if found, return the value of a given property. Properties are searched for based on the <i>dip</i>, <i>name</i>, <i>match_dev</i>, and the type of the data (integer, string or byte). The property search order is as follows:</p> <ol style="list-style-type: none"> <li>1. Search software properties created by the driver.</li> <li>2. Search the software properties created by the system (or nexus nodes in the device info tree).</li> <li>3. Search the driver global properties list.</li> <li>4. If <code>DDI_PROP_NOTPROM</code> is not set, search the PROM properties (if they exist).</li> <li>5. If <code>DDI_PROP_DONTPASS</code> is not set, pass this request to the parent device information node.</li> <li>6. Return <code>DDI_PROP_NOT_FOUND</code>.</li> </ol> <p>Usually, the <i>match_dev</i> argument should be set to the actual device number that this property is associated with. However, if the <i>match_dev</i> argument is <code>DDI_DEV_T_ANY</code>, the property look up routines will match the request regardless of the actual <i>match_dev</i> the property was created with. If a property</p>

was created with *match\_dev* set to `DDI_DEV_T_NONE`, then the only way to look up this property is with a *match\_dev* set to `DDI_DEV_T_ANY`. PROM properties are always created with *match\_dev* set to `DDI_DEV_T_NONE`.

*name* must always be set to the name of the property being looked up.

For the routines `ddi_prop_lookup_int_array()`, `ddi_prop_lookup_string_array()`, `ddi_prop_lookup_string()`, and `ddi_prop_lookup_byte_array()`, *datap* is the address of a pointer which, upon successful return, will point to memory containing the value of the property. In each case *datap* points to a different type of property value. See the individual descriptions of the routines below for details on the different return values. *nelementsp* is the address of an unsigned integer which, upon successful return, will contain the number of integer, string or byte elements accounted for in the memory pointed at by *datap*.

All of the property look up routines may block to allocate memory needed to hold the value of the property.

When a driver has obtained a property with any look up routine and is finished with that property, it must be freed by calling `ddi_prop_free()`. `ddi_prop_free()` must be called with the address of the allocated property. For instance, if one called `ddi_prop_lookup_int_array()` with *datap* set to the address of a pointer to an integer, `&my_int_ptr`, then the companion free call would be `ddi_prop_free( my_int_ptr )`.

#### **ddi\_prop\_lookup\_int\_array()**

This routine searches for and returns an array of integer property values. An array of integers is defined to *nelementsp* number of 4 byte long integer elements. *datap* should be set to the address of a pointer to an array of integers which, upon successful return, will point to memory containing the integer array value of the property.

#### **ddi\_prop\_lookup\_string\_array()**

This routine searches for and returns a property that is an array of strings. *datap* should be set to address of a pointer to an array of strings which, upon successful return, will point to memory containing the array of strings. The array of strings is formatted as an array of pointers to null-terminated strings, much like the *argv* argument to `execve(2)`.

#### **ddi\_prop\_lookup\_string()**

This routine searches for and returns a property that is a null-terminated string. *datap* should be set to the address of a pointer to string which, upon

successful return, will point to memory containing the string value of the property.

### **ddi\_prop\_lookup\_byte\_array()**

This routine searches for and returns a property that is an array of bytes. *datap* should be set to the address of a pointer to an array of bytes which, upon successful return, will point to memory containing the byte array value of the property.

### **ddi\_prop\_free()**

Frees the resources associated with a property previously allocated using **ddi\_prop\_lookup\_int\_array()**, **ddi\_prop\_lookup\_string\_array()**, **ddi\_prop\_lookup\_string()**, or **ddi\_prop\_lookup\_byte\_array()**.

## RETURN VALUES

The functions **ddi\_prop\_lookup\_int\_array()**, **ddi\_prop\_lookup\_string\_array()**, **ddi\_prop\_lookup\_string()**, and **ddi\_prop\_lookup\_byte\_array()** return the following values:

DDI_PROP_SUCCESS	Upon success.
DDI_PROP_INVALID_ARG	If an attempt is made to look up a property with <i>match_dev</i> equal to <code>DDI_DEV_T_NONE</code> , <i>name</i> is NULL or <i>name</i> is the null string.
DDI_PROP_NOT_FOUND	Property not found.
DDI_PROP_UNDEFINED	Property explicitly not defined (see <b>ddi_prop_undefine(9F)</b> ).
DDI_PROP_CANNOT_DECODE	The value of the property cannot be decoded.

## CONTEXT

These functions can be called from user or kernel context.

## EXAMPLES

**EXAMPLE 1** Using **ddi\_prop\_lookup()** :

The following example demonstrates the use of **ddi\_prop\_lookup()**.

```
\011
int\011*options;
int\011noptions;

/*
 * Get the data associated with the integer "options" property
```

```
    * array, along with the number of option integers
    */
\011if (ddi_prop_lookup_int_array(DDI_DEV_T_ANY, xx_dip, 0,
    "options", &options, &noptions) == DDI_PROP_SUCCESS) {
    /*
    * Do "our thing" with the options data from the property
    */
    xx_process_options(options, noptions);

    /*
    * Free the memory allocated for the property data
    */
    ddi_prop_free(options);
}
```

**SEE ALSO**

**execve(2)**, **ddi\_prop\_exists(9F)**, **ddi\_prop\_get\_int(9F)**,  
**ddi\_prop\_remove(9F)**, **ddi\_prop\_undefine(9F)**,  
**ddi\_prop\_update(9F)**

*Writing Device Drivers*

<b>NAME</b>	ddi_prop_op, ddi_getprop, ddi_getlongprop, ddi_getlongprop_buf, ddi_getproplen – get property information for leaf device drivers
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_prop_op(dev_t dev, dev_info_t * dip, ddi_prop_op_t prop_op, int flags, char * name, caddr_t valuep, int * lengthp);  int ddi_getprop(dev_t dev, dev_info_t * dip, int flags, char * name, int defvalue);  int ddi_getlongprop(dev_t dev, dev_info_t * dip, int flags, char * name, caddr_t valuep, int * lengthp);  int ddi_getlongprop_buf(dev_t dev, dev_info_t * dip, int flags, char * name, caddr_t valuep, int * lengthp);  int ddi_getproplen(dev_t dev, dev_info_t * dip, int flags, char * name, int * lengthp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dev</b> Device number associated with property or DDI_DEV_T_ANY as the <i>wildcard</i> device number.</p> <p><b>dip</b> Pointer to a device info node.</p> <p><b>prop_op</b> Property operator.</p> <p><b>flags</b> Possible flag values are some combination of:</p> <ul style="list-style-type: none"> <li>DDI_PROP_DONTPASS do not pass request to parent device information node if property not found</li> <li>DDI_PROP_CANSLEEP the routine may sleep while allocating memory</li> <li>DDI_PROP_NOTPROM do not look at PROM properties (ignored on architectures that do not support PROM properties)</li> </ul> <p><b>name</b> String containing the name of the property.</p>

**valuep** If *prop\_op* is `PROP_LEN_AND_VAL_BUF`, this should be a pointer to the users buffer. If *prop\_op* is `PROP_LEN_AND_VAL_ALLOC`, this should be the *address* of a pointer.

**lengthp** On exit, *\*lengthp* will contain the property length. If *prop\_op* is `PROP_LEN_AND_VAL_BUF` then before calling `ddi_prop_op()`, *lengthp* should point to an `int` that contains the length of callers buffer.

**defvalue**The value that `ddi_getprop()` returns if the property is not found.

## DESCRIPTION

`ddi_prop_op()` gets arbitrary-size properties for leaf devices. The routine searches the device's property list. If it does not find the property at the device level, it examines the *flags* argument, and if `DDI_PROP_DONTPASS` is set, then `ddi_prop_op()` returns `DDI_PROP_NOT_FOUND`. Otherwise, it passes the request to the next level of the device info tree. If it does find the property, but the property has been explicitly undefined, it returns `DDI_PROP_UNDEFINED`. Otherwise it returns either the property length, or both the length and value of the property to the caller via the *valuep* and *lengthp* pointers, depending on the value of *prop\_op*, as described below, and returns `DDI_PROP_SUCCESS`. If a property cannot be found at all, `DDI_PROP_NOT_FOUND` is returned.

Usually, the *dev* argument should be set to the actual device number that this property applies to. However, if the *dev* argument is `DDI_DEV_T_ANY`, the *wildcard dev*, then `ddi_prop_op()` will match the request based on *name* only (regardless of the actual *dev* the property was created with). This property/*dev* match is done according to the property search order which is to first search software properties created by the driver in *last-in, first-out* (LIFO) order, next search software properties created by the *system* in LIFO order, then search PROM properties if they exist in the system architecture.

Property operations are specified by the *prop\_op* argument. If *prop\_op* is `PROP_LEN`, then `ddi_prop_op()` just sets the callers length, *\*lengthp*, to the property length and returns the value `DDI_PROP_SUCCESS` to the caller. The *valuep* argument is not used in this case. Property lengths are 0 for boolean properties, `sizeof(int)` for integer properties, and size in bytes for long (variable size) properties.

If *prop\_op* is `PROP_LEN_AND_VAL_BUF`, then *valuep* should be a pointer to a user-supplied buffer whose length should be given in *\*lengthp* by the caller. If the requested property exists, `ddi_prop_op()` first sets *\*lengthp* to the property length. It then examines the size of the buffer supplied by the caller, and if it is large enough, copies the property value into that buffer, and returns `DDI_PROP_SUCCESS`. If the named property exists but the buffer supplied is too small to hold it, it returns `DDI_PROP_BUF_TOO_SMALL`.

If *prop\_op* is `PROP_LEN_AND_VAL_ALLOC`, and the property is found, **ddi\_prop\_op()** sets *lengthp* to the property length. It then attempts to allocate a buffer to return to the caller using the `kmem_alloc(9F)` routine, so that memory can be later recycled using `kmem_free(9F)`. The driver is expected to call `kmem_free()` with the returned address and size when it is done using the allocated buffer. If the allocation is successful, it sets *valuep* to point to the allocated buffer, copies the property value into the buffer and returns `DDI_PROP_SUCCESS`. Otherwise, it returns `DDI_PROP_NO_MEMORY`. Note that the *flags* argument may affect the behavior of memory allocation in **ddi\_prop\_op()**. In particular, if `DDI_PROP_CANSLEEP` is set, then the routine will wait until memory is available to copy the requested property.

**ddi\_getprop()** returns boolean and integer-size properties. It is a convenience wrapper for **ddi\_prop\_op()** with *prop\_op* set to `PROP_LEN_AND_VAL_BUF`, and the buffer is provided by the wrapper. By convention, this function returns a 1 for boolean (zero-length) properties.

**ddi\_getlongprop()** returns arbitrary-size properties. It is a convenience wrapper for **ddi\_prop\_op()** with *prop\_op* set to `PROP_LEN_AND_VAL_ALLOC`, so that the routine will allocate space to hold the buffer that will be returned to the caller via *valuep*.

**ddi\_getlongprop\_buf()** returns arbitrary-size properties. It is a convenience wrapper for **ddi\_prop\_op()** with *prop\_op* set to `PROP_LEN_AND_VAL_BUF` so the user must supply a buffer.

**ddi\_getproplen()** returns the length of a given property. It is a convenience wrapper for **ddi\_prop\_op()** with *prop\_op* set to `PROP_LEN`.

## RETURN VALUES

**ddi\_prop\_op()** **ddi\_getlongprop()** **ddi\_getlongprop\_buf()** **ddi\_getproplen()** return:

<code>DDI_PROP_SUCCESS</code>	Property found and returned.
<code>DDI_PROP_NOT_FOUND</code>	Property not found.
<code>DDI_PROP_UNDEFINED</code>	Property already explicitly undefined.
<code>DDI_PROP_NO_MEMORY</code>	Property found, but unable to allocate memory. <i>lengthp</i> points to the correct property length.
<code>DDI_PROP_BUF_TOO_SMALL</code>	Property found, but the supplied buffer is too small. <i>lengthp</i> points to the correct property length.

**ddi\_getprop()** returns:

The value of the property or the value passed into the routine as *defvalue* if the property is not found. By convention, the value of zero length properties (boolean properties) are returned as the integer value 1.

**CONTEXT** These functions can be called from user or interrupt context, provided `DDI_PROP_CANSLEEP` is not set; if it is set, they can be called from user context only.

**SEE ALSO** `ddi_prop_create(9F)` , `kmem_alloc(9F)` , `kmem_free(9F)`

*Writing Device Drivers*

<b>NAME</b>	ddi_prop_update, ddi_prop_update_int_array, ddi_prop_update_int, ddi_prop_update_string_array, ddi_prop_update_string, ddi_prop_update_byte_array - update properties
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_prop_update_int_array(dev_t dev, dev_info_t * dip, char * name, int * data, uint_t nelements);  int ddi_prop_update_int(dev_t dev, dev_info_t * dip, char * name, int data);  int ddi_prop_update_string_array(dev_t dev, dev_info_t * dip, char * name, char ** data, uint_t nelements);  int ddi_prop_update_string(dev_t dev, dev_info_t * dip, char * name, char * data);  int ddi_prop_update_byte_array(dev_t dev, dev_info_t * dip, char * name, uchar_t * data, uint_t nelements);</pre>
<b>PARAMETERS</b>	<p><b>dev</b> Device number associated with the device.</p> <p><b>dip</b> Pointer to the device info node of device whose property list should be updated.</p> <p><b>name</b> String containing the name of the property to be updated.</p> <p><b>nelements</b> The number of elements contained in the memory pointed at by <i>data</i> .</p>
<b>ddi_prop_update_int_array()</b>	<b>data</b> A pointer an integer array with which to update the property.
<b>ddi_prop_update_int()</b>	<b>data</b> An integer value with which to update the property.
<b>ddi_prop_update_string_array()</b>	<b>data</b> A pointer to a string array with which to update the property. The array of strings is formatted as an array of pointers to NULL terminated strings, much like the <i>argv</i> argument to <b>execve(2)</b> .
<b>ddi_prop_update_string()</b>	<b>data</b> A pointer to a string value with which to update the property.

<b>ddi_prop_update_byte_array()</b>	A pointer to a byte array with which to update the property.
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<p>The property update routines search for and, if found, modify the value of a given property. Properties are searched for based on the <i>dip</i> , <i>name</i> , <i>dev</i> , and the type of the data (integer, string or byte). The driver software properties list is searched. If the property is found, it is updated with the supplied value. If the property is not found on this list, a new property is created with the value supplied. For example, if a driver attempts to update the "foo" property, a property named "foo" is searched for on the driver's software property list. If "foo" is found, the value is updated. If "foo" is not found, a new property named "foo" is created on the driver's software property list with the supplied value even if a "foo" property exists on another property list (such as a PROM property list).</p> <p>Every property value has a data type associated with it: byte, integer, or string. A property should be updated using a function with the same corresponding data type as the property value. For example, an integer property must be updated using either <b>ddi_prop_update_int_array()</b> or <b>ddi_prop_update_int()</b> . Attempts to update a property with a function that does correspond to the property value data type will result in the creation of another property with the same name. However, the data type of the new property value will correspond to the data type called out in the function name.</p> <p>Usually, the <i>dev</i> argument should be set to the actual device number that this property is associated with. If the property is not associated with any particular <i>dev</i> , then the argument <i>dev</i> should be set to <code>DDI_DEV_T_NONE</code> . This property will then match a look up request (see <b>ddi_prop_lookup(9F)</b> )with the <i>match_dev</i> argument set to <code>DDI_DEV_T_ANY</code> . If no <i>dev</i> is available for the device (for example during <b>attach(9E)</b> time), one can be created using <b>makedevice(9F)</b> with a major number of <code>DDI_MAJOR_T_UNKNOWN</code> . The update routines will then generate the correct <i>dev</i> when creating or updating the property.</p> <p><i>name</i> must always be set to the name of the property being updated.</p> <p>For the routines <code>ddi_prop_update_int_array()</code> , <code>ddi_prop_update_string_array()</code> , <code>ddi_prop_update_string()</code> , and <code>ddi_prop_update_byte_array()</code> <i>data</i> is a pointer which points to memory containing the value of the property. In each case <i>*data</i> points to a different type of property value. See the individual descriptions of the routines below for details concerning the different values. <i>nelements</i> is an unsigned integer which contains the number of integer, string, or byte elements accounted for in the memory pointed at by <i>*data</i> .</p>

- For the routine `ddi_prop_update_int()`, *data* is the new value of the property.
- ddi\_prop\_update\_int\_array()** Updates or creates an array of integer property values. An array of integers is defined to be *n* elements of 4 byte long integer elements. *data* must be a pointer to an integer array with which to update the property.
- ddi\_prop\_update\_int()** Update or creates a single integer value of a property. *data* must be an integer value with which to update the property.
- ddi\_prop\_update\_string\_array()** Updates or creates a property that is an array of strings. *data* must be a pointer to a string array with which to update the property. The array of strings is formatted as an array of pointers to NULL terminated strings, much like the *argv* argument to `execve(2)`.
- ddi\_prop\_update\_string()** Updates or creates a property that is a single string value. *data* must be a pointer to a string with which to update the property.
- ddi\_prop\_update\_byte\_array()** Updates or creates a property that is an array of bytes. *data* should be a pointer to a byte array with which to update the property.

The property update routines may block to allocate memory needed to hold the value of the property.

#### RETURN VALUES

All of the property update routines return:

DDI_PROP_SUCCESS	On success.
DDI_PROP_INVALID_ARG	If an attempt is made to update a property with <i>name</i> set to NULL or <i>name</i> set to the null string.
DDI_PROP_CANNOT_ENCODE	If the bytes of the property cannot be encoded.

#### CONTEXT

These functions can only be called from user or kernel context.

#### EXAMPLES

##### EXAMPLE 1 Updating Properties

The following example demonstrates the use of `ddi_prop_update()`.

```
int\011options[4];

\011/*
\011 * Create the "options" integer array with
\011 * our default values for these parameters
\011 */
\011options[0] = XX_OPTIONS0;
\011options[1] = XX_OPTIONS1;
```

```
\011options[2] = XX_OPTIONS2;  
\011options[3] = XX_OPTIONS3;  
\011i = ddi_p  
rop_update_int_array(xx_dev, xx_dip, "options",  
\011\011&options, sizeof (options) / sizeof (int));
```

**SEE ALSO**

**execve(2)**, **attach(9E)**, **ddi\_prop\_lookup(9F)**, **ddi\_prop\_remove(9F)**, **makedevice(9F)**

*Writing Device Drivers*

<b>NAME</b>	ddi_put8, ddi_put16, ddi_put32, ddi_put64, ddi_putb, ddi_putl, ddi_putll, ddi_putw – write data to the mapped memory address, device register or allocated DMA memory address
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_put8(ddi_acc_handle_t handle, uint8_t *dev_addr, uint8_t value); void ddi_put16(ddi_acc_handle_t handle, uint16_t *dev_addr, uint16_t value); void ddi_put32(ddi_acc_handle_t handle, uint32_t *dev_addr, uint32_t value); void ddi_put64(ddi_acc_handle_t handle, uint64_t *dev_addr, uint64_t value);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The data access handle returned from setup calls, such as <code>ddi_regs_map_setup(9F)</code> .</p> <p><b>value</b>            The data to be written to the device.</p> <p><b>dev_addr</b>         Base device address.</p>
<b>DESCRIPTION</b>	<p>These routines generate a write of various sizes to the mapped memory or device register. The <code>ddi_put8()</code> , <code>ddi_put16()</code> , <code>ddi_put32()</code> , and <code>ddi_put64()</code> functions write 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, to the device address, <code>dev_addr</code> .</p> <p>Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.</p>
<b>CONTEXT</b>	These functions can be called from user, kernel, or interrupt context.
<b>SEE ALSO</b>	<code>ddi_get8(9F)</code> , <code>ddi_regs_map_free(9F)</code> , <code>ddi_regs_map_setup(9F)</code> , <code>ddi_rep_get8(9F)</code> , <code>ddi_rep_put8(9F)</code> , <code>ddi_device_acc_attr(9S)</code>
<b>NOTES</b>	The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
ddi_putb	ddi_put8
ddi_putw	ddi_put16
ddi_putl	ddi_put32
ddi_putll	ddi_put64

**NAME** ddi\_regs\_map\_free – free a previously mapped register address space

**SYNOPSIS** #include <sys/ddi.h>  
#include <sys/sunddi.h>

```
void ddi_regs_map_free(ddi_acc_handle_t *handle);
```

**INTERFACE LEVEL** Solaris DDI specific (Solaris DDI).

**PARAMETERS** *handle* Pointer to a data access handle previously allocated by a call to a setup routine such as ddi\_regs\_map\_setup(9F).

**DESCRIPTION** ddi\_regs\_map\_free() frees the mapping represented by the data access handle *handle*. This function is provided for drivers preparing to detach themselves from the system, allowing them to release allocated system resources represented in the handle.

**CONTEXT** ddi\_regs\_map\_free() must be called from user or kernel context.

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	PCI Local Bus, SBus, ISA, EISA

**SEE ALSO** attributes(5), ddi\_regs\_map\_setup(9F)  
*Writing Device Drivers*

<b>NAME</b>	ddi_regs_map_setup – set up a mapping for a register address space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_regs_map_setup(dev_info_t *dip, uint_t rnumber, caddr_t *addrp, offset_t offset, offset_t len, ddi_device_acc_attr_t *accattrp, ddi_acc_handle_t *handlep);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>rnumber</b> Index number to the register address space set.</p> <p><b>addrp</b> Pointer to the mapping address base.</p> <p><b>offset</b> Offset into the register address space.</p> <p><b>len</b> Length to be mapped.</p> <p><b>accattrp</b> Pointer to a device access attribute structure of this device (see ddi_device_acc_attr(9S)).</p> <p><b>handlep</b> Pointer to a data access handle.</p>
<b>DESCRIPTION</b>	<p><b>ddi_regs_map_setup()</b> maps in the register set given by <i>rnumber</i>. The register number determines which register set is mapped if more than one exists.</p> <p><i>offset</i> specifies the starting location within the register space and <i>len</i> indicates the size of the area to be mapped. If <i>len</i> is non-zero, it overrides the length given in the register set description. If both <i>len</i> and <i>offset</i> are 0, the entire space is mapped. The base of the mapped register space is returned in <i>addrp</i>.</p> <p>The device access attributes are specified in the location pointed by the <i>accattrp</i> argument (see ddi_device_acc_attr(9S) for details).</p> <p>The data access handle is returned in <i>handlep</i>. <i>handlep</i> is opaque; drivers should not attempt to interpret its value. The handle is used by the system to encode information for subsequent data access function calls to maintain a consistent view between the host and the device.</p>
<b>RETURN VALUES</b>	<p><b>ddi_regs_map_setup()</b> returns:</p> <p>DDI_SUCCESS Successfully set up the mapping for data access.</p>

DDI\_FAILURE

Invalid register number *rnumber*,  
offset *offset*, or length *len*.

DDI\_REGS\_ACC\_CONFLICT

Cannot enable the register mapping  
due to access conflicts with other  
enabled mappings.**CONTEXT****ddi\_regs\_map\_setup()** must be called from user or kernel context.**ATTRIBUTES**See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	PCI Local Bus, SBus, ISA, EISA

**SEE ALSO****attributes(5)**, **ddi\_regs\_map\_free(9F)**, **ddi\_device\_acc\_attr(9S)***Writing Device Drivers*

<b>NAME</b>	<code>ddi_remove_minor_node</code> – remove a minor node for this <code>dev_info</code>
<b>SYNOPSIS</b>	<code>void ddi_remove_minor_node(dev_info_t *dip, char *name);</code>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b> A pointer to the device's <code>dev_info</code> structure.</p> <p><b>name</b> The name of this minor device. If <i>name</i> is <code>NULL</code>, then remove all minor data structures from this <code>dev_info</code>.</p>
<b>DESCRIPTION</b>	<b>ddi_remove_minor_node()</b> removes a data structure from the linked list of minor data structures that is pointed to by the <code>dev_info</code> structure for this driver.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Removing a minor node</p> <p>This will remove a data structure describing a minor device called <code>dev1</code> which is linked into the <code>dev_info</code> structure pointed to by <code>dip</code>:</p> <pre>ddi_remove_minor_node(dip, "dev1");</pre>
<b>SEE ALSO</b>	<code>attach(9E)</code> , <code>detach(9E)</code> , <code>ddi_create_minor_node(9F)</code> <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_rep_get8, ddi_rep_get16, ddi_rep_get32, ddi_rep_get64, ddi_rep_getw, ddi_rep_getl, ddi_rep_getll, ddi_rep_getb – read data from the mapped memory address, device register or allocated DMA memory address
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_rep_get8(ddi_acc_handle_t handle, uint8_t *host_addr, uint8_t *dev_addr, size_t recount, uint_t flags);  void ddi_rep_get16(ddi_acc_handle_t handle, uint16_t *host_addr, uint16_t *dev_addr, size_t recount, uint_t flags);  void ddi_rep_get32(ddi_acc_handle_t handle, uint32_t *host_addr, uint32_t *dev_addr, size_t recount, uint_t flags);  void ddi_rep_get64(ddi_acc_handle_t handle, uint64_t *host_addr, uint64_t *dev_addr, size_t recount, uint_t flags);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI ).
<b>PARAMETERS</b>	<p><b>handle</b>           The data access handle returned from setup calls, such as <code>ddi_regs_map_setup(9F)</code> .</p> <p><b>host_addr</b>       Base host address.</p> <p><b>dev_addr</b>        Base device address.</p> <p><b>recount</b>        Number of data accesses to perform.</p> <p><b>flags</b>           Device address flags:</p> <p style="padding-left: 40px;">DDI_DEV_AUTOINCR Automatically increment the device address, <i>dev_addr</i> , during data accesses.</p> <p style="padding-left: 40px;">DDI_DEV_NO_AUTOINCR Do not advance the device address, <i>dev_addr</i> , during data accesses.</p>
<b>DESCRIPTION</b>	These routines generate multiple reads from the mapped memory or device register. <i>recount</i> data is copied from the device address, <i>dev_addr</i> , to the host address, <i>host_addr</i> . For each input datum, the <code>ddi_rep_get8()</code> , <code>ddi_rep_get16()</code> , <code>ddi_rep_get32()</code> , and <code>ddi_rep_get64()</code> functions read 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively, from the device address,

*dev\_addr* . *dev\_addr* and *host\_addr* must be aligned to the datum boundary described by the function.

Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.

When the *flags* argument is set to `DDI_DEV_AUTOINCR` , these functions treat the device address, *dev\_addr* , as a memory buffer location on the device and increment its address on the next input datum. However, when the *flags* argument is to `DDI_DEV_NO_AUTOINCR` , the same device address will be used for every datum access. For example, this flag may be useful when reading from a data register.

**RETURN VALUES**

These functions return the value read from the mapped address.

**CONTEXT**

These functions can be called from user, kernel, or interrupt context.

**SEE ALSO**

`ddi_get8(9F)` , `ddi_put8(9F)` , `ddi_regs_map_free(9F)` ,  
`ddi_regs_map_setup(9F)` , `ddi_rep_put8(9F)`

**NOTES**

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
<code>ddi_rep_getb</code>	<code>ddi_rep_get8</code>
<code>ddi_rep_getw</code>	<code>ddi_rep_get16</code>
<code>ddi_rep_getl</code>	<code>ddi_rep_get32</code>
<code>ddi_rep_getll</code>	<code>ddi_rep_get64</code>

<b>NAME</b>	ddi_report_dev - announce a device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_report_dev(dev_info_t *dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>dip</b> a pointer the device's dev_info structure.
<b>DESCRIPTION</b>	<b>ddi_report_dev()</b> prints a banner at boot time, announcing the device pointed to by <i>dip</i> . The banner is always placed in the system logfile (displayed by <b>dmesg(1M)</b> ), but is only displayed on the console if the system was booted with the verbose (-v) argument.
<b>CONTEXT</b>	<b>ddi_report_dev()</b> can be called from user context.
<b>SEE ALSO</b>	<b>dmesg(1M)</b> , <b>kernel(1M)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	ddi_rep_put8, ddi_rep_put16, ddi_rep_put32, ddi_rep_put64, ddi_rep_putb, ddi_rep_putw, ddi_rep_putl, ddi_rep_putll – write data to the mapped memory address, device register or allocated DMA memory address
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void ddi_rep_put8(ddi_acc_handle_t handle, uint8_t *host_addr, uint8_t *dev_addr, size_t repcount, uint_t flags);  void ddi_rep_put16(ddi_acc_handle_t handle, uint16_t *host_addr, uint16_t *dev_addr, size_t repcount, uint_t flags);  void ddi_rep_put32(ddi_acc_handle_t handle, uint32_t *host_addr, uint32_t *dev_addr, size_t repcount, uint_t flags);  void ddi_rep_put64(ddi_acc_handle_t handle, uint64_t *host_addr, uint64_t *dev_addr, size_t repcount, uint_t flags);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>           The data access handle returned from setup calls, such as <b>ddi_regs_map_setup(9F)</b>.</p> <p><b>host_addr</b>        Base host address.</p> <p><b>dev_addr</b>         Base device address.</p> <p><b>repcount</b>        Number of data accesses to perform.</p> <p><b>flags</b>            Device address flags:</p> <p style="padding-left: 40px;">DDI_DEV_AUTOINCR Automatically increment the device address, <i>dev_addr</i>, during data accesses.</p> <p style="padding-left: 40px;">DDI_DEV_NO_AUTOINCR Do not advance the device address, <i>dev_addr</i>, during data accesses.</p>
<b>DESCRIPTION</b>	These routines generate multiple writes to the mapped memory or device register. <i>repcount</i> data is copied from the host address, <i>host_addr</i> , to the device address, <i>dev_addr</i> . For each input datum, the <b>ddi_rep_put8()</b> , <b>ddi_rep_put16()</b> , <b>ddi_rep_put32()</b> , and <b>ddi_rep_put64()</b> functions write 8 bits, 16 bits, 32 bits,

and 64 bits of data, respectively, to the device address, *dev\_addr*. *dev\_addr* and *host\_addr* must be aligned to the datum boundary described by the function.

Each individual datum will automatically be translated to maintain a consistent view between the host and the device based on the encoded information in the data access handle. The translation may involve byte-swapping if the host and the device have incompatible endian characteristics.

When the *flags* argument is set to `DDI_DEV_AUTOINCR`, these functions treat the device address, *dev\_addr*, as a memory buffer location on the device and increment its address on the next input datum. However, when the *flags* argument is set to `DDI_DEV_NO_AUTOINCR`, the same device address will be used for every datum access. For example, this flag may be useful when writing to a data register.

**CONTEXT**

These functions can be called from user, kernel, or interrupt context.

**SEE ALSO**

`ddi_get8(9F)`, `ddi_put8(9F)`, `ddi_regs_map_free(9F)`,  
`ddi_regs_map_setup(9F)`, `ddi_rep_get8(9F)`,  
`ddi_device_acc_attr(9S)`

**NOTES**

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
<code>ddi_rep_putb</code>	<code>ddi_rep_put8</code>
<code>ddi_rep_putw</code>	<code>ddi_rep_put16</code>
<code>ddi_rep_putl</code>	<code>ddi_rep_put32</code>
<code>ddi_rep_putll</code>	<code>ddi_rep_put64</code>

<b>NAME</b>	ddi_root_node – get the root of the dev_info tree
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <pre>dev_info_t *ddi_root_node(void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>DESCRIPTION</b>	<b>ddi_root_node()</b> returns a pointer to the root node of the device information tree.
<b>RETURN VALUES</b>	<b>ddi_root_node()</b> returns a pointer to a device information structure.
<b>CONTEXT</b>	<b>ddi_root_node()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_segmap, ddi_segmap_setup – set up a user mapping using seg_dev
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt;  #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_segmap(dev_t dev, off_t offset, struct as * asp, caddr_t * addrp, off_t len, uint_t prot, uint_t maxprot, uint_t flags, cred_t *credp);  int ddi_segmap_setup(dev_t dev, off_t offset, struct as * asp, caddr_t * addrp, off_t len, uint_t prot, uint_t maxprot, uint_t flags, cred_t *credp, ddi_device_acc_attr_t *accattrp, uint_t rnumber);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dev</b> The device whose memory is to be mapped.</p> <p><b>offset</b> The offset within device memory at which the mapping begins.</p> <p><b>asp</b> An opaque pointer to the user address space into which the device memory should be mapped.</p> <p><b>addrp</b> Pointer to the starting address within the user address space to which the device memory should be mapped.</p> <p><b>len</b> Length (in bytes) of the memory to be mapped.</p> <p><b>prot</b> A bit field that specifies the protections. Some combinations of possible settings are:</p> <pre>PROT_READ     Read access is desired.  PROT_WRITE     Write access is desired.  PROT_EXEC     Execute access is desired.  PROT_USER     User-level access is desired (the mapping is being done as a result of a mmap(2) system call).  PROT_ALL     All access is desired.</pre>

**maxprot** Maximum protection flag possible for attempted mapping (the `PROT_WRITE` bit may be masked out if the user opened the special file read-only). If `(maxprot & prot) != prot` then there is an access violation.

**flags** Flags indicating type of mapping. Possible values are (other bits may be set):

`MAP_PRIVATE`

Changes are private.

`MAP_SHARED`

Changes should be shared.

`MAP_FIXED`

**credp** The user specified an address in `*addrp` rather than letting the system pick an address. Pointer to user credential structure.

#### `ddi_segmap_setup()`

**dev\_acc\_attr** Pointer to a `ddi_device_acc_attr(9S)` structure which contains the device access attributes to apply to this mapping.

**rnumber** Index number to the register address space set.

#### DESCRIPTION

Future releases of Solaris will provide this function for binary and source compatibility. However, for increased functionality, use `ddi_devmap_segmap(9F)` instead. See `ddi_devmap_segmap(9F)` for details.

`ddi_segmap()` and `ddi_segmap_setup()` set up user mappings to device space. When setting up the mapping, the `ddi_segmap()` and `ddi_segmap_setup()` routines call the `mmap(9E)` entry point to validate the range to be mapped. When a user process accesses the mapping, the drivers `mmap(9E)` entry point is again called to retrieve the page frame number that needs to be loaded. The mapping translations for that page are then loaded on behalf of the driver by the DDI framework.

`ddi_segmap()` is typically used as the `segmap(9E)` entry in the `cb_ops(9S)` structure for those devices that do not choose to provide their own `segmap(9E)` entry point. However, some drivers may have their own `segmap(9E)` entry point to do some initial processing on the parameters and then call `ddi_segmap()` to establish the default memory mapping.

**ddi\_segmap\_setup()** is used in the drivers `segmap(9E)` entry point to set up the mapping and assign device access attributes to that mapping. *rnumber* specifies the register set representing the range of device memory being mapped. See `ddi_device_acc_attr(9S)` for details regarding what device access attributes are available.

**ddi\_segmap\_setup()** cannot be used directly in the `cb_ops(9S)` structure and requires a driver to have a `segmap(9E)` entry point.

**RETURN VALUES**

**ddi\_segmap()** and **ddi\_segmap\_setup()** return the following values:

0 Successful completion.

Non-zero An error occurred. In particular, they return ENXIO if the range to be mapped is invalid.

**CONTEXT**

**ddi\_segmap()** and **ddi\_segmap\_setup()** can be called from user or kernel context only.

**SEE ALSO**

`mmap(2)` , `mmap(9E)` , `segmap(9E)` , `ddi_mapdev(9F)` , `cb_ops(9S)` , `ddi_device_acc_attr(9S)`

*Writing Device Drivers*

**NOTES**

If driver notification of user accesses to the mappings is required, the driver should use `ddi_mapdev(9F)` instead.

<b>NAME</b>	<b>ddi_slaveonly</b> – tell if a device is installed in a slave access only location
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int ddi_slaveonly(dev_info_t *dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>dip</b> A pointer to the device's dev_info structure.
<b>DESCRIPTION</b>	<b>ddi_slaveonly()</b> tells the caller if the bus, or part of the bus that the device is installed on, does not permit the device to become a DMA master, that is, whether the device has been installed in a slave access only slot.
<b>RETURN VALUES</b>	<b>DDI_SUCCESS</b> The device has been installed in a slave access only location. <b>DDI_FAILURE</b> The device has not been installed in a slave access only location.
<b>CONTEXT</b>	<b>ddi_slaveonly()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	ddi_soft_state, ddi_get_soft_state, ddi_soft_state_fini, ddi_soft_state_free, ddi_soft_state_init, ddi_soft_state_zalloc – driver soft state utility routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void * ddi_get_soft_state(void *state, int item);  void ddi_soft_state_fini(void **state_p);  void ddi_soft_state_free(void *state, int item);  int ddi_soft_state_init(void **state_p, size_t size, size_t n_items);  int ddi_soft_state_zalloc(void *state, int item);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>state_p</b> Address of the opaque state pointer which will be initialized by <b>ddi_soft_state_init()</b> to point to implementation dependent data.</p> <p><b>size</b> Size of the item which will be allocated by subsequent calls to <b>ddi_soft_state_zalloc()</b> .</p> <p><b>n_items</b> A hint of the number of items which will be preallocated; zero is allowed.</p> <p><b>state</b> An opaque pointer to implementation-dependent data that describes the soft state.</p> <p><b>item</b> The item number for the state structure; usually the instance number of the associated devinfo node.</p>
<b>DESCRIPTION</b>	<p>Most device drivers maintain state information with each instance of the device they control; for example, a soft copy of a device control register, a mutex that must be held while accessing a piece of hardware, a partition table, or a unit structure. These utility routines are intended to help device drivers manage the space used by the driver to hold such state information.</p> <p>For example, if the driver holds the state of each instance in a single state structure, these routines can be used to dynamically allocate and deallocate a separate structure for each instance of the driver as the instance is attached and detached.</p> <p>To use the routines, the driver writer needs to declare a state pointer, <i>state_p</i> , which the implementation uses as a place to hang a set of per-driver structures; everything else is managed by these routines.</p>

The routine **ddi\_soft\_state\_init()** is usually called in the drivers **\_init(9E)** routine to initialize the state pointer, set the size of the soft state structure, and to allow the driver to pre-allocate a given number of such structures if required.

The routine **ddi\_soft\_state\_zalloc()** is usually called in the drivers **attach(9E)** routine. The routine is passed an item number which is used to refer to the structure in subsequent calls to **ddi\_get\_soft\_state()** and **ddi\_soft\_state\_free()**. The item number is usually just the instance number of the devinfo node, obtained with **ddi\_get\_instance(9F)**. The routine attempts to allocate space for the new structure, and if the space allocation was successful, **DDI\_SUCCESS** is returned to the caller.

A pointer to the space previously allocated for a soft state structure can be obtained by calling **ddi\_get\_soft\_state()** with the appropriate item number.

The space used by a given soft state structure can be returned to the system using **ddi\_soft\_state\_free()**. This routine is usually called from the drivers **detach(9E)** entry point.

The space used by all the soft state structures allocated on a given state pointer, together with the housekeeping information used by the implementation can be returned to the system using **ddi\_soft\_state\_fini()**. This routine can be called from the drivers **\_fini(9E)** routine.

The **ddi\_soft\_state\_zalloc()**, **ddi\_soft\_state\_free()** and **ddi\_get\_soft\_state()** routines coordinate access to the underlying data structures in an MT-safe fashion, thus no additional locks should be necessary.

## RETURN VALUES

<b>ddi_get_soft_state()</b>	<b>NULL</b>	The requested state structure was not allocated at the time of the call.
	<b>pointer</b>	The pointer to the state structure.
<b>ddi_soft_state_init()</b>	<b>0</b>	The allocation was successful.
	<b>EINVAL</b>	Either the <i>size</i> parameter was zero, or the <i>state_p</i> parameter was invalid.
<b>ddi_soft_state_zalloc()</b>	<b>DDI_SUCCESS</b>	The allocation was successful.
	<b>DDI_FAILURE</b>	The routine failed to allocate the storage required; either the <i>state</i> parameter was invalid, the item number was negative,

or an attempt was made to allocate an item number that was already allocated.

**CONTEXT**

**ddi\_soft\_state\_init()** , and **ddi\_soft\_state\_alloc()** can be called from user context only, since they may internally call **kmem\_zalloc(9F)** with the **KM\_SLEEP** flag.

The **ddi\_soft\_state\_fini()** , **ddi\_soft\_state\_free()** and **ddi\_get\_soft\_state()** routines can be called from any driver context.

**EXAMPLES****CODE EXAMPLE 1** Creating and Removing Data Structures

The following example shows how the routines described above can be used in terms of the driver entry points of a character-only driver. The example concentrates on the portions of the code that deal with creating and removing the driver's data structures.

```
typedef struct {
    volatile caddr_t *csr;      /* device registers */
    kmutex_t      csr_mutex;   /* protects 'csr' field */
    unsigned int  state;
    dev_info_t    *dip;        /* back pointer to devinfo */
} devstate_t;
static void *statep;

int
_init(void)
{
    int error;

    error = ddi_soft_state_init(&statep, sizeof (devstate_t), 0);
    if (error != 0)
\011\011    return (error);
    if ((error = mod_install(&modlinkage)) != 0)
\011\011    ddi_soft_state_fini(&statep);
    return (error);
}

int
_fini(void)
{
    int error;

    if ((error = mod_remove(&modlinkage)) != 0)
\011\011    return (error);
    ddi_soft_state_fini(&statep);
    return (0);
}

static int
xxattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    int instance;
    devstate_t *softc;
```

```

        switch (cmd) {
        case DDI_ATTACH:
\011     instance = ddi_get_instance(dip);
        if (ddi_soft_state_zalloc(statep, instance) != DDI_SUCCESS)
\011         return (DDI_FAILURE);
\011     softc = ddi_get_soft_state(statep, instance);
\011     softc->dip = dip;\011
\011     ...
\011     return (DDI_SUCCESS);
        default:
\011     return (DDI_FAILURE);
        }
    }

static int
xxdetach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    int instance;

    switch (cmd) {

        case DDI_DETACH:
\011     instance = ddi_get_instance(dip);
\011     ...
        ddi_soft_state_free(statep, instance);
        return (DDI_SUCCESS);

        default:
            return (DDI_FAILURE);
        }
    }

static int
xxopen(dev_t *devp, int flag, int otyp, cred_t *cred_p)
{
    devstate_t *softc;
    int instance;

    instance = getminor(*devp);
    if ((softc = ddi_get_soft_state(statep, instance)) == NULL)
\011     return (ENXIO);
    ...
    softc->state |= XX_IN_USE;
    ...
    return (0);
}

```

**SEE ALSO**

***\_fini(9E)***, ***\_init(9E)***, ***attach(9E)***, ***detach(9E)***,  
***ddi\_get\_instance(9F)***, ***getminor(9F)***, ***kmem\_zalloc(9F)***

*Writing Device Drivers*

**WARNINGS**

There is no attempt to validate the `item` parameter given to **ddi\_soft\_state\_zalloc()** other than it must be a positive signed integer. Therefore very large item numbers may cause the driver to hang forever waiting for virtual memory resources that can never be satisfied.

**NOTES**

If necessary, a hierarchy of state structures can be constructed by embedding state pointers in higher order state structures.

**DIAGNOSTICS**

All of the messages described below usually indicate bugs in the driver and should not appear in normal operation of the system.

```
WARNING: ddi_soft_state_zalloc: bad handle
WARNING: ddi_soft_state_free: bad handle
WARNING: ddi_soft_state_fini: bad handle
```

The implementation-dependent information kept in the state variable is corrupt.

```
WARNING: ddi_soft_state_free: null handle
WARNING: ddi_soft_state_fini: null handle
```

The routine has been passed a null or corrupt state pointer. Check that **ddi\_soft\_state\_init()** has been called.

```
WARNING: ddi_soft_state_free: item %d not in range [0..%d]
```

The routine has been asked to free an item which was never allocated. The message prints out the invalid item number and the acceptable range.

<b>NAME</b>	ddi_umem_alloc, ddi_umem_free – allocate and free page-aligned kernel memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/sunddi.h&gt;  void * ddi_umem_alloc(size_t size, int flag, ddi_umem_cookie_t * cookiep, ) void ddi_umem_free(ddi_umem_cookie_t cookie);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
<b>ddi_umem_alloc()</b>	<p><b>size</b>     Number of bytes to allocate.</p> <p><b>flag</b>     Used to determine the sleep and pageable conditions. Possible sleep flags are DDI_UMEM_SLEEP which allows sleeping until memory is available, and DDI_UMEM_NOSLEEP which returns NULL immediately if memory is not available. The default condition is to allocate locked memory; this can be changed to allocate pageable memory using the DDI_UMEM_PAGEABLE flag.</p> <p><b>cookiep</b> Pointer to a kernel memory cookie.</p>
<b>ddi_umem_free()</b>	<p><b>cookie</b> A kernel memory cookie allocated in <b>ddi_umem_alloc()</b> .</p>
<b>DESCRIPTION</b>	<p><b>ddi_umem_alloc()</b> allocates page-aligned kernel memory and returns a pointer to the allocated memory. The number of bytes allocated is a multiple of the system page size (roundup of <i>size</i> ). The allocated memory can be used in the kernel and can be exported to user space. See <b>devmap(9E)</b> and <b>devmap_umem_setup(9F)</b> for further information.</p> <p><i>flag</i> determines whether the caller can sleep for memory and whether the allocated memory is locked or not. DDI_UMEM_SLEEP allocations may sleep but are guaranteed to succeed. DDI_UMEM_NOSLEEP allocations do not sleep but may fail (return NULL )if memory is currently unavailable. If DDI_UMEM_PAGEABLE is set, pageable memory will be allocated. These pages can be swapped out to secondary memory devices. The initial contents of memory allocated using <b>ddi_umem_alloc()</b> is zero-filled.</p> <p><i>*cookiep</i> is a pointer to the kernel memory cookie that describes the kernel memory being allocated. A typical use of <i>cookiep</i> is in</p>

<b>RETURN VALUES</b>	<p><b>devmap_umem_setup(9F)</b> when the drivers want to export the kernel memory to a user application.</p> <p>To free the allocated memory, a driver calls <b>ddi_umem_free()</b> with the cookie obtained from <b>ddi_umem_alloc()</b> . <b>ddi_umem_free()</b> releases the entire buffer.</p> <p>Non-null      Successful completion. <b>ddi_umem_alloc()</b> returns a pointer to the allocated memory.</p> <p>NULL            Memory cannot be allocated by <b>ddi_umem_alloc()</b> because <b>DDI_UMEM_NOSLEEP</b> is set and the system is out of resources.</p>
<b>CONTEXT</b>	<p><b>ddi_umem_alloc()</b> can be called from any context if <i>flag</i> is set to <b>DDI_UMEM_NOSLEEP</b> . If <b>DDI_UMEM_SLEEP</b> is set, <b>ddi_umem_alloc()</b> can be called from user and kernel context only. <b>ddi_umem_free()</b> can be called from any context.</p>
<b>SEE ALSO</b>	<p><b>devmap(9E)</b> , <b>condvar(9F)</b> , <b>devmap_umem_setup(9F)</b> , <b>kmem_alloc(9F)</b> , <b>mutex(9F)</b> , <b>rwlock(9F)</b> , <b>semaphore(9F)</b></p> <p><i>Writing Device Drivers</i></p>
<b>WARNINGS</b>	<p>Setting the <b>DDI_UMEM_PAGEABLE</b> flag in <b>ddi_umem_alloc()</b> will result in an allocation of pageable memory. Because these pages can be swapped out to secondary memory devices, drivers should use this flag with care. This memory should not be used for synchronization objects such as locks and condition variables. See <b>mutex(9F)</b> , <b>semaphore(9F)</b> , <b>rwlock(9F)</b> , and <b>condvar(9F)</b> . This memory also should not be accessed in the driver interrupt routines.</p> <p>Memory allocated using <b>ddi_umem_alloc()</b> without setting <b>DDI_UMEM_PAGEABLE</b> flag cannot be paged. Available memory is therefore limited by the total physical memory on the system. It is also limited by the available kernel virtual address space, which is often the more restrictive constraint on large-memory configurations.</p> <p>Excessive use of kernel memory is likely to effect overall system performance. Over-commitment of kernel memory may cause unpredictable consequences.</p> <p>Misuse of the kernel memory allocator, such as writing past the end of a buffer, using a buffer after freeing it, freeing a buffer twice, or freeing an invalid pointer, will cause the system to corrupt data or panic.</p>
<b>NOTES</b>	<p><b>ddi_umem_alloc(0, flag, cookiep)</b> always returns <b>NULL</b> .</p> <p><b>ddi_umem_free(NULL)</b> has no effects on system.</p>

<b>NAME</b>	<b>delay</b> – delay execution for a specified number of clock ticks
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  void <b>delay</b>(clock_t ticks);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>ticks</b> The number of clock cycles to delay.
<b>DESCRIPTION</b>	<p><b>delay()</b> provides a mechanism for a driver to delay its execution for a given period of time. Since the speed of the clock varies among systems, drivers should base their time values on microseconds and use <b>drv_usec2ohz(9F)</b> to convert microseconds into clock ticks.</p> <p><b>delay()</b> uses <b>timeout(9F)</b> to schedule an internal function to be called after the specified amount of time has elapsed. <b>delay()</b> then waits until the function is called.</p> <p><b>delay()</b> does not busy-wait. If busy-waiting is required, use <b>drv_usecwait(9F)</b>.</p>
<b>CONTEXT</b>	<b>delay()</b> can be called from user and kernel contexts.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1    delay() Example</b></p> <p>Before a driver I/O routine allocates buffers and stores any user data in them, it checks the status of the device (line 12). If the device needs manual intervention (such as, needing to be refilled with paper), a message is displayed on the system console (line 14). The driver waits an allotted time (line 17) before repeating the procedure.</p> <pre> 1  struct  device  { /* layout of physical device registers */ 2          int     control; /* physical device control word */ 3          int     status; /* physical device status word */ 4          short   xmit_char; /* transmit character to device */ 5  }; 6 7 8          . . . 9  /* get device registers */ 10     register struct device *rp = ... 11 12     while (rp-&gt;status &amp; NOPAPER) { /* while printer is out of paper */ 13     /* display message and ring bell */ 14     /* on system console */ 15         cmn_err(CE_WARN, "\007", 16                 (getminor(dev) &amp; 0xf)); 17     /* wait one minute and try again */</pre>

```
17         delay(60 * drv_usectohz(1000000));  
18     }
```

**SEE ALSO**

**biodone(9F), biowait(9F), drv\_hztousec(9F), drv\_usectohz(9F),  
drv\_usecwait(9F), timeout(9F), untimeout(9F)**

*Writing Device Drivers*

<b>NAME</b>	devmap_default_access – default driver memory access function
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int devmap_default_access(devmap_cookie_t dhp, void *pvtp, offset_t off, size_t len, uint_t type, uint_t rw);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dhp</b> An opaque mapping handle that the system uses to describe the mapping.</p> <p><b>pvtp</b> Driver private mapping data.</p> <p><b>off</b> User offset within the logical device memory at which the access begins.</p> <p><b>len</b> Length (in bytes) of the memory being accessed.</p> <p><b>type</b> Type of access operation.</p> <p><b>rw</b> Type of access.</p>
<b>DESCRIPTION</b>	<p><b>devmap_default_access()</b> is a function providing the semantics of <b>devmap_access(9E)</b>. The drivers call <b>devmap_default_access()</b> to handle the mappings that do not support context switching. The drivers should call <b>devmap_do_ctxmgt(9F)</b> for the mappings that support context management.</p> <p><b>devmap_default_access()</b> can either be called from <b>devmap_access(9E)</b> or be used as the <b>devmap_access(9E)</b> entry point. The arguments <i>dhp</i>, <i>pvtp</i>, <i>off</i>, <i>len</i>, <i>type</i>, and <i>rw</i> are provided by the <b>devmap_access(9E)</b> entry point and must not be modified.</p>
<b>RETURN VALUES</b>	<p>0 Successful completion.</p> <p>Non-zero An error occurred.</p>
<b>CONTEXT</b>	<b>devmap_default_access()</b> must be called from the driver's <b>devmap_access(9E)</b> entry point.

**EXAMPLES****EXAMPLE 1** Using devmap\_default\_access in devmap\_access.

The following shows an example of using **devmap\_default\_access()** in the **devmap\_access(9E)** entry point.

```

...
#define OFF_DO_CTXMGT  0x40000000
#define OFF_NORMAL    0x40100000
#define CTXMGT_SIZE   0x100000
#define NORMAL_SIZE   0x100000

/*
 * Driver devmap_contextmgt(9E) callback function.
 */
static int
xx_context_mgt(devmap_cookie_t dhp, void *pvtp, offset_t offset,
              size_t length, uint_t type, uint_t rw)
{
    .....
    /*
     * see devmap_contextmgt(9E) for an example
     */
}

/*
 * Driver devmap_access(9E) entry point
 */
static int
xxdevmap_access(devmap_cookie_t dhp, void *pvtp, offset_t off,
               size_t len, uint_t type, uint_t rw)
{
    offset_t diff;
    int err;

    /*
     * check if off is within the range that supports
     * context management.
     */
    if ((diff = off - OFF_DO_CTXMGT) >= 0 && diff < CTXMGT_SIZE) {
        /*
         * calculates the length for context switching
         */
        if ((len + off) > (OFF_DO_CTXMGT + CTXMGT_SIZE))
            return (-1);
        /*
         * perform context switching
         */
        err = devmap_do_ctxmgt(dhp, pvtp, off, len, type,
                              rw, xx_context_mgt);
    }
    /*
     * check if off is within the range that does normal
     * memory mapping.
     */
} else if ((diff = off - OFF_NORMAL) >= 0 && diff < NORMAL_SIZE) {
    if ((len + off) > (OFF_NORMAL + NORMAL_SIZE))
        return (-1);
    err = devmap_default_access(dhp, pvtp, off, len, type, rw);
}

```

```
    } else  
        return (-1);  
    return (err);  
}
```

**SEE ALSO**

**devmap\_access(9E), devmap\_do\_ctxmgt(9F),  
devmap\_callback\_ctl(9S)**

*Writing Device Drivers*

<b>NAME</b>	devmap_devmem_setup, devmap_umem_setup – set driver memory mapping parameters								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int devmap_devmem_setup(devmap_cookie_t dhp, dev_info_t * dip, struct devmap_callback_ctl * callbackops, uint_t rnumber, offset_t roff, size_t len, uint_t maxprot, uint_t flags, ddi_device_acc_attr_t * accattrp);  int devmap_umem_setup(devmap_cookie_t dhp, dev_info_t * dip, struct devmap_callback_ctl * callbackops, ddi_umem_cookie_t cookie, offset_t koff, size_t len, uint_t maxprot, uint_t flags, ddi_device_acc_attr_t * accattrp);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).								
<b>PARAMETERS</b>									
devmap_devmem_setup(0)	<p><b>dhp</b> An opaque mapping handle that the system uses to describe the mapping.</p> <p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>callbackops</b> Pointer to a devmap_callback_ctl(9S) structure. The structure contains pointers to device driver-supplied functions that manage events on the device mapping. The framework will copy the structure to the system private memory.</p> <p><b>rnumber</b> Index number to the register address space set.</p> <p><b>roff</b> Offset into the register address space.</p> <p><b>len</b> Length (in bytes) of the mapping to be mapped.</p> <p><b>maxprot</b> Maximum protection flag possible for attempted mapping. Some combinations of possible settings are:</p> <table border="0" style="margin-left: 40px;"> <tr> <td>PROT_READ</td> <td>Read access is allowed.</td> </tr> <tr> <td>PROT_WRITE</td> <td>Write access is allowed.</td> </tr> <tr> <td>PROT_EXEC</td> <td>Execute access is allowed.</td> </tr> <tr> <td>PROT_USER</td> <td>User-level access is allowed (the mapping is being done as a result of a mmap(2) system call).</td> </tr> </table>	PROT_READ	Read access is allowed.	PROT_WRITE	Write access is allowed.	PROT_EXEC	Execute access is allowed.	PROT_USER	User-level access is allowed (the mapping is being done as a result of a mmap(2) system call).
PROT_READ	Read access is allowed.								
PROT_WRITE	Write access is allowed.								
PROT_EXEC	Execute access is allowed.								
PROT_USER	User-level access is allowed (the mapping is being done as a result of a mmap(2) system call).								

		PROT_READ	Read access is allowed.
		PROT_WRITE	Write access is allowed.
		PROT_EXEC	Execute access is allowed.
		PROT_USER	User-level access is allowed (the mapping is being done as a result of a <code>mmap(2)</code> system call).
		PROT_ALL	All access is allowed.
	<b>flags</b>		Must be set to 0 .
	<b>accattrp</b>		Pointer to a <code>ddi_device_acc_attr(9S)</code> structure. The structure contains the device access attributes to be applied to this range of memory.
<b>devmap_uem_setup()</b>	<b>dhp</b>		An opaque data structure that the system uses to describe the mapping.
	<b>dip</b>		Pointer to the device's <code>dev_info</code> structure.
	<b>callbackops</b>		Pointer to a <code>devmap_callback_ctl(9S)</code> structure. The structure contains pointers to device driver-supplied functions that manage events on the device mapping.
	<b>cookie</b>		A kernel memory <i>cookie</i> (see <code>ddi_uem_alloc(9F)</code> ).
	<b>koff</b>		Offset into the kernel memory defined by <i>cookie</i> .
	<b>len</b>		Length (in bytes) of the mapping to be mapped.
	<b>maxprot</b>		Maximum protection flag possible for attempted mapping. Some combinations of possible settings are:
		PROT_READ	Read access is allowed.
		PROT_WRITE	Write access is allowed.
		PROT_EXEC	Execute access is allowed.
		PROT_USER	User-level access is allowed (the mapping is being done as a result of a <code>mmap(2)</code> system call).
		PROT_ALL	All access is allowed.
	<b>flags</b>		Must be set to 0 .
	<b>accattrp</b>		Pointer to a <code>ddi_device_acc_attr(9S)</code> structure. The structure contains the device access attributes to be applied to this range of memory.
<b>DESCRIPTION</b>	<b>devmap_devmem_setup()</b> and <b>devmap_uem_setup()</b> are used in the <b>devmap(9E)</b> entry point to pass mapping parameters from the driver to the system.		

*dhp* is a device mapping handle that the system uses to store all mapping parameters of a physical contiguous memory. The system copies the data pointed to by *callbacks* to a system private memory. This allows the driver to free the data after returning from either **devmap\_devmem\_setup()** or **devmap\_umem\_setup()**. The driver is notified of user events on the mappings via the entry points defined by **devmap\_callback\_ctl(9S)**. The driver is notified of the following user events:

**Mapping Setup** User has called **mmap(2)** to create a mapping to the device memory.

**Access** User has accessed an address in the mapping that has no translations.

**Duplication** User has duplicated the mapping. Mappings are duplicated when the process calls **fork(2)**.

**Unmapping** User has called **munmap(2)** on the mapping or is exiting, **exit(2)**.

See **devmap\_map(9E)**, **devmap\_access(9E)**, **devmap\_dup(9E)**, and **devmap\_unmap(9E)** for details on these entry points.

By specifying a valid *callbacks* to the system, device drivers can manage events on a device mapping. For example, the **devmap\_access(9E)** entry point allows the drivers to perform context switching by unloading the mappings of other processes and to load the mapping of the calling process. Device drivers may specify **NULL** to *callbacks* which means the drivers do not want to be notified by the system.

The maximum protection allowed for the mapping is specified in *maxprot*. *accattrp* defines the device access attributes. See **ddi\_device\_acc\_attr(9S)** for more details.

**devmap\_devmem\_setup()** is used for device memory to map in the register set given by *rnumber* and the offset into the register address space given by *roff*. The system uses *rnumber* and *roff* to go up the device tree to get the physical address that corresponds to *roff*. The range to be affected is defined by *len* and *roff*. The range from *roff* to *roff + len* must be a physical contiguous memory and page aligned.

Drivers use **devmap\_umem\_setup()** for kernel memory to map in the kernel memory described by *cookie* and the offset into the kernel memory space given by *koff*. *cookie* is a kernel memory pointer obtained from **ddi\_umem\_alloc(9F)**. If *cookie* is **NULL**, **devmap\_umem\_setup()** returns **-1**. The range to be affected is defined by *len* and *koff*. The range from *koff* to *koff + len* must be within the limits of the kernel memory described by *koff + len* and must be page aligned.

Drivers use **devmap\_umem\_setup()** to export the kernel memory allocated by **ddi\_umem\_alloc(9F)** to user space. The system selects a user virtual address that is aligned with the kernel virtual address being mapped to avoid cache incoherence if the mapping is not `MAP_FIXED`.

**RETURN VALUES**

0        Successful completion.

-1        An error occurred.

**CONTEXT**

**devmap\_devmem\_setup()** and **devmap\_umem\_setup()** can be called from user, kernel, and interrupt context.

**SEE ALSO**

**exit(2)**, **fork(2)**, **mmap(2)**, **munmap(2)**, **devmap(9E)**,  
**ddi\_umem\_alloc(9F)**, **ddi\_device\_acc\_attr(9S)**,  
**devmap\_callback\_ctl(9S)**

*Writing Device Drivers*

<b>NAME</b>	devmap_do_ctxmgt – perform device context switching on a mapping
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int devmap_do_ctxmgt(devmap_cookie_t,dhp, void *pvtp, offset_t off, size_t len, uint_t type, uint_t rw, int (*devmap_contextmgt), (devmap_cookie_t,void *,offset_t,size_t,uint_t,uint_t));</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dhp</b> An opaque mapping handle that the system uses to describe the mapping.</p> <p><b>pvtp</b> Driver private mapping data.</p> <p><b>off</b> User offset within the logical device memory at which the access begins.</p> <p><b>len</b> Length (in bytes) of the memory being accessed.</p> <p><b>devmap_contextmgt</b> The address of driver function that the system will call to perform context switching on a mapping. See <a href="#">devmap_contextmgt(9E)</a> for details.</p> <p><b>type</b> Type of access operation. Provided by <a href="#">devmap_access(9E)</a>. Should not be modified.</p> <p><b>rw</b> Direction of access. Provided by <a href="#">devmap_access(9E)</a>. Should not be modified.</p>
<b>DESCRIPTION</b>	<p>Device drivers call <a href="#">devmap_do_ctxmgt()</a> in the <a href="#">devmap_access(9E)</a> entry point to perform device context switching on a mapping. <a href="#">devmap_do_ctxmgt()</a> passes a pointer to a driver supplied callback function, <a href="#">devmap_contextmgt(9E)</a>, to the system that will perform the actual device context switching. If <a href="#">devmap_contextmgt(9E)</a> is not a valid driver callback function, the system will fail the memory access operation which will result in a SIGSEGV or SIGBUS signal being delivered to the process.</p> <p><a href="#">devmap_do_ctxmgt()</a> performs context switching on the mapping object identified by <i>dhp</i> and <i>pvtp</i> in the range specified by <i>off</i> and <i>len</i>. The arguments <i>dhp</i>, <i>pvtp</i>, <i>type</i>, and <i>rw</i> are provided by the <a href="#">devmap_access(9E)</a> entry point and must not be modified. The range from <i>off</i> to <i>off+len</i> must support context switching.</p>

The system will pass through *dhp*, *pvtpt*, *off*, *len*, *type*, and *rw* to **devmap\_contextmgt(9E)** in order to perform the actual device context switching. The return value from **devmap\_contextmgt(9E)** will be returned directly to **devmap\_do\_ctxmgt()**.

**RETURN VALUES**

0 Successful completion.

Non-zero An error occurred.

**CONTEXT**

**devmap\_do\_ctxmgt()** must be called from the driver's **devmap\_access(9E)** entry point.

**EXAMPLES**

**EXAMPLE 1** Using **devmap\_do\_ctxmgt** in the **devmap\_access** entry point.

The following shows an example of using **devmap\_do\_ctxmgt()** in the **devmap\_access(9E)** entry point.

```

...
#define OFF_DO_CTXMGT 0x40000000
#define OFF_NORMAL 0x40100000
#define CTXMGT_SIZE 0x100000
#define NORMAL_SIZE 0x100000

/*
 * Driver devmap_contextmgt(9E) callback function.
 */
static int
xx_context_mgt(devmap_cookie_t dhp, void *pvtpt, offset_t offset,
              size_t length, uint_t type, uint_t rw)
{
    .....
    /*
     * see devmap_contextmgt(9E) for an example
     */
}

/*
 * Driver devmap_access(9E) entry point
 */
static int
xxdevmap_access(devmap_cookie_t dhp, void *pvtpt, offset_t off,
               size_t len, uint_t type, uint_t rw)
{
    offset_t diff;
    int err;

    /*
     * check if off is within the range that supports
     * context management.
     */
    if ((diff = off - OFF_DO_CTXMGT) >= 0 && diff < CTXMGT_SIZE) {
        /*
         * calculates the length for context switching

```

```
    /*
    if ((len + off) > (OFF_DO_CTXMGT + CTXMGT_SIZE))
        return (-1);
    /*
    * perform context switching
    */
    err = devmap_do_ctxmgt(dhp, pvtp, off, len, type,
                          rw, xx_context_mgt);
/*
* check if off is within the range that does normal
* memory mapping.
*/
} else if ((diff = off - OFF_NORMAL) >= 0 && diff < NORMAL_SIZE) {
    if ((len + off) > (OFF_NORMAL + NORMAL_SIZE))
        return (-1);
    err = devmap_default_access(dhp, pvtp, off, len, type, rw);
} else
    return (-1);

return (err);
}
```

**SEE ALSO**

**devmap\_access(9E), devmap\_contextmgt(9E),  
devmap\_default\_access(9F)**

*Writing Device Drivers*

<b>NAME</b>	devmap_set_ctx_timeout – set the timeout value for the context management callback
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void devmap_set_ctx_timeout(devmap_cookie_t <i>dhp</i>, clock_t <i>ticks</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>dhp</i></b> An opaque mapping handle that the system uses to describe the mapping.</p> <p><b><i>ticks</i></b> Number of clock ticks to wait between successive calls to the context management callback function.</p>
<b>DESCRIPTION</b>	<p><b>devmap_set_ctx_timeout()</b> specifies the time interval for the system to wait between successive calls to the driver's context management callback function, <b>devmap_contextmgt(9E)</b>.</p> <p>Device drivers typically call <b>devmap_set_ctx_timeout()</b> in the <b>devmap_map(9E)</b> routine. If the drivers do not call <b>devmap_set_ctx_timeout()</b> to set the timeout value, the default timeout value of 0 will result in no delay between successive calls to the driver's <b>devmap_contextmgt(9E)</b> callback function.</p>
<b>CONTEXT</b>	<b>devmap_set_ctx_timeout()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>devmap_contextmgt(9E)</b> , <b>devmap_map(9E)</b> , <b>timeout(9F)</b>

<b>NAME</b>	devmap_setup, ddi_devmap_segmap – set up a user mapping to device memory using the devmap framework								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int devmap_setup(dev_t dev, offset_t off, ddi_as_handle_t as, caddr_t * addrp, size_t len, uint_t prot, uint_t maxprot, uint_t flags, cred_t * cred);  int ddi_devmap_segmap(dev_t dev, off_t off, ddi_as_handle_t as, caddr_t * addrp, off_t len, uint_t prot, uint_t maxprot, uint_t flags, cred_t * cred);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).								
<b>PARAMETERS</b>	<p><b>dev</b> Device whose memory is to be mapped.</p> <p><b>off</b> User offset within the logical device memory at which the mapping begins.</p> <p><b>as</b> An opaque data structure that describes the address space into which the device memory should be mapped.</p> <p><b>addrp</b> Pointer to the starting address in the address space into which the device memory should be mapped.</p> <p><b>len</b> Length (in bytes) of the memory to be mapped.</p> <p><b>prot</b> A bit field that specifies the protections. Some possible settings combinations are:</p> <table border="0" style="margin-left: 2em;"> <tr> <td>PROT_READ</td> <td>Read access is desired.</td> </tr> <tr> <td>PROT_WRITE</td> <td>Write access is desired.</td> </tr> <tr> <td>PROT_EXEC</td> <td>Execute access is desired.</td> </tr> <tr> <td>PROT_USER</td> <td>User-level access is desired (the mapping is being done as a result of a <code>mmap(2)</code> system call).</td> </tr> </table> <p><b>maxprot</b> Maximum protection flag possible for attempted mapping; the PROT_READ, PROT_WRITE, PROT_EXEC, and PROT_USER bits may be masked out if the user opened the special file read-only.</p> <p><b>flags</b> Flags indicating type of mapping. The following flags can be specified:</p>	PROT_READ	Read access is desired.	PROT_WRITE	Write access is desired.	PROT_EXEC	Execute access is desired.	PROT_USER	User-level access is desired (the mapping is being done as a result of a <code>mmap(2)</code> system call).
PROT_READ	Read access is desired.								
PROT_WRITE	Write access is desired.								
PROT_EXEC	Execute access is desired.								
PROT_USER	User-level access is desired (the mapping is being done as a result of a <code>mmap(2)</code> system call).								

MAP\_PRIVATE Changes are private.

MAP\_SHARED Changes should be shared.

MAP\_FIXED The user specified an address in *\*addrp* rather than letting the system choose an address.

**cred** Pointer to the user credential structure.

**DESCRIPTION**

**devmap\_setup()** and **ddi\_devmap\_segmap()** allow device drivers to use the devmap framework to set up user mappings to device memory. The devmap framework provides several advantages over the default device mapping framework that is used by **ddi\_segmap(9F)** or **ddi\_segmap\_setup(9F)**. Device drivers should use the devmap framework, if the driver wants to:

- use an optimal MMU pagesize to minimize address translations,
- conserve kernel resources,
- receive callbacks to manage events on the mapping,
- export kernel memory to applications,
- set up device contexts for the user mapping if the device requires context switching,
- assign device access attributes to the user mapping, or
- change the maximum protection for the mapping.

**devmap\_setup()** must be called in the **segmap(9E)** entry point to establish the mapping for the application. **ddi\_devmap\_segmap()** can be called in, or be used as, the **segmap(9E)** entry point. The differences between **devmap\_setup()** and **ddi\_devmap\_segmap()** are in the data type used for *off* and *len*.

When setting up the mapping, **devmap\_setup()** and **ddi\_devmap\_segmap()** call the **devmap(9E)** entry point to validate the range to be mapped. The **devmap(9E)** entry point also translates the logical offset (as seen by the application) to the corresponding physical offset within the device address space. If the driver does not provide its own **devmap(9E)** entry point, EINVAL will be returned to the **mmap(2)** system call.

**RETURN VALUES**

0 Successful completion.

Non-zero An error occurred. The return value of **devmap\_setup()** and **ddi\_devmap\_segmap()** should be used directly in the **segmap(9E)** entry point.

**CONTEXT** `devmap_setup()` and `ddi_devmap_segmap()` can be called from user or kernel context only.

**SEE ALSO** `mmap(2)` , `devmap(9E)` , `segmap(9E)` , `ddi_segmap(9F)` , `ddi_segmap_setup(9F)` , `cb_ops(9S)`

*Writing Device Drivers*

<b>NAME</b>	devmap_unload, devmap_load – control validation of memory address translations
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int devmap_load(devmap_cookie_t dhp, offset_t off, size_t len, uint_t type, uint_t rw, int devmap_unload(devmap_cookie_t dhp, offset_t off, size_t len);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dhp</b> An opaque mapping handle that the system uses to describe the mapping.</p> <p><b>off</b> User offset within the logical device memory at which the loading or unloading of the address translations begins.</p> <p><b>len</b> Length (in bytes) of the range being affected.</p>
devmap_load() only	<p><b>type</b> Type of access operation.</p> <p><b>rw</b> Direction of access.</p>
<b>DESCRIPTION</b>	<p><b>devmap_unload()</b> and <b>devmap_load()</b> are used to control the validation of the memory mapping described by <i>dhp</i> in the specified range. <b>devmap_unload()</b> invalidates the mapping translations and will generate calls to the <b>devmap_access(9E)</b> entry point next time the mapping is accessed. The drivers use <b>devmap_load()</b> to validate the mapping translations during memory access.</p> <p>A typical use of <b>devmap_unload()</b> and <b>devmap_load()</b> is in the driver's context management callback function, <b>devmap_contextmgt(9E)</b>. To manage a device context, a device driver calls <b>devmap_unload()</b> on the context about to be switched out. It switches contexts, and then calls <b>devmap_load()</b> on the context switched in. <b>devmap_unload()</b> can be used to unload the mappings of other processes as well as the mappings of the calling process, but <b>devmap_load()</b> can only be used to load the mappings of the calling process. Attempting to load another process's mappings with <b>devmap_load()</b> will result in a system panic.</p> <p>For both routines, the range to be affected is defined by the <i>off</i> and <i>len</i> arguments. Requests affect the entire page containing the <i>off</i> and all pages up to and including the page containing the last byte as indicated by <i>off + len</i>. The arguments <i>type</i> and <i>rw</i> are provided by the system to the calling function (for example, <b>devmap_contextmgt(9E)</b>) and should not be modified.</p>

Supplying a value of 0 for the *len* argument affects all addresses from the *off* to the end of the mapping. Supplying a value of 0 for the *off* argument and a value of 0 for *len* argument affect all addresses in the mapping.

A non-zero return value from either **devmap\_unload()** or **devmap\_load()** will cause the corresponding operation to fail. The failure may result in a SIGSEGV or SIGBUS signal being delivered to the process.

**RETURN VALUES**

0                      Successful completion.  
 Non-zero              An error occurred.

**CONTEXT**

These routines can be called from user or kernel context only.

**EXAMPLES****EXAMPLE 1** Managing a One-Page Device Context

The following shows an example of managing a device context that is one page in length.

```

struct xx_context cur_ctx;

static int
xxdevmap_contextmgt(devmap_cookie_t dhp, void *pvtp, offset_t off,
    size_t len, uint_t type, uint_t rw)
{
    int err;
    devmap_cookie_t cur_dhp;
    struct xx_pvt *p;
    struct xx_pvt *pvp = (struct xx_pvt *)pvtp;
    /* enable access callbacks for the current mapping */
    if (cur_ctx != NULL && cur_ctx != pvp->ctx) {
        p = cur_ctx->pvt;
        /*
         * unload the region from off to the end of the mapping.
         */
        cur_dhp = p->dhp;
        if ((err = devmap_unload(cur_dhp, off, len)) != 0)
            return (err);
    }
    /* Switch device context - device dependent*/
    ...
    /* Make handle the new current mapping */
    cur_ctx = pvp->ctx;
    /*
     * Disable callbacks and complete the access for the
     * mapping that generated this callback.
     */
    return (devmap_load(pvp->dhp, off, len, type, rw));
}

```

**SEE ALSO**

**devmap\_access(9E)**, **devmap\_contextmgt(9E)**

*Writing Device Drivers*

<b>NAME</b>	disksort – single direction elevator seek sort for buffers
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt; void  disksort(struct diskhd *dp, struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dp</b> A pointer to a <code>diskhd</code> structure. A <code>diskhd</code> structure is essentially identical to head of a buffer structure (see <code>buf(9S)</code>). The only defined items of interest for this structure are the <code>av_forw</code> and <code>av_back</code> structure elements which are used to maintain the front and tail pointers of the forward linked I/O request queue.</p> <p><b>bp</b> A pointer to a buffer structure. Typically this is the I/O request that the driver receives in its strategy routine (see <code>strategy(9E)</code>). The driver is responsible for initializing the <code>b_resid</code> structure element to a meaningful sort key value prior to calling <code>disksort()</code>.</p>
<b>DESCRIPTION</b>	<p>The function <code>disksort()</code> sorts a pointer to a buffer into a single forward linked list headed by the <code>av_forw</code> element of the argument <code>*dp</code>.</p> <p>It uses a one-way elevator algorithm that sorts buffers into the queue in ascending order based upon a key value held in the argument buffer structure element <code>b_resid</code>.</p> <p>This value can either be the driver calculated cylinder number for the I/O request described by the buffer argument, or simply the absolute logical block for the I/O request, depending on how fine grained the sort is desired to be or how applicable either quantity is to the device in question.</p> <p>The head of the linked list is found by use of the <code>av_forw</code> structure element of the argument <code>*dp</code>. The tail of the linked list is found by use of the <code>av_back</code> structure element of the argument <code>*dp</code>. The <code>av_forw</code> element of the <code>*bp</code> argument is used by <code>disksort()</code> to maintain the forward linkage. The value at the head of the list presumably indicates the currently active disk area.</p>
<b>CONTEXT</b>	This function can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><code>strategy(9E)</code>, <code>buf(9S)</code></p> <p><i>Writing Device Drivers</i></p>

**WARNINGS**

**disksort()** does no locking. Therefore, any locking is completely the responsibility of the caller.

<b>NAME</b>	drv_getparm - retrieve kernel state information
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  int drv_getparm(unsigned int parm, void *value_p);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>parm</b>            The kernel parameter to be obtained. Possible values are:</p> <p><b>LBOLT</b>            Read the value of <code>lbolt</code>. <code>lbolt</code> is an integer that is unconditionally incremented by one at each clock tick. No special treatment is applied when this value overflows the maximum value of a signed int. When this occurs, its value will be negative, and its magnitude will be decreasing until it again passes zero. It can therefore not be relied upon to provide an indication of the amount of time that passes since the last system reboot, nor should it be used to mark an absolute time in the system. Only the difference between two measurements of <code>lbolt</code> is significant. It is used in this way inside the system kernel for timing purposes.</p> <p><b>PPGRP</b>            Read the process group identification number. This number determines which processes should receive a HANGUP or BREAK signal when detected by a driver.</p> <p><b>UPROCP</b>           Read the process table token value.</p> <p><b>PPID</b>            Read process identification number.</p> <p><b>PSID</b>            Read process session identification number.</p> <p><b>TIME</b>            Read time in seconds.</p> <p><b>UCRED</b>           Return a pointer to the caller's credential structure.</p> <p><b>value_p</b>          A pointer to the data space in which the value of the parameter is to be copied.</p>
<b>DESCRIPTION</b>	<b>drv_getparm()</b> function verifies that <i>parm</i> corresponds to a kernel parameter that may be read. If the value of <i>parm</i> does not correspond to a parameter or corresponds to a parameter that may not be read, -1 is returned. Otherwise, the value of the parameter is stored in the data space pointed to by <i>value_p</i> .

**drv\_getparm()** does not explicitly check to see whether the device has the appropriate context when the function is called and the function does not check for correct alignment in the data space pointed to by *value\_p*. It is the responsibility of the driver writer to use this function only when it is appropriate to do so and to correctly declare the data space needed by the driver.

**RETURN VALUES**

**drv\_getparm()** returns 0 to indicate success, -1 to indicate failure. The value stored in the space pointed to by *value\_p* is the value of the parameter if 0 is returned, or undefined if -1 is returned. -1 is returned if you specify a value other than LBOLT, PPGRP, PPID, PSID, TIME, UCRED, or UPROCP. Always check the return code when using this function.

**CONTEXT**

**drv\_getparm()** can be called from user context only when using PPGRP, PPID, PSID, UCRED, or UPROCP. It can be called from user or interrupt context when using the LBOLT or TIME argument.

**SEE ALSO**

**buf(9S)**

*Writing Device Drivers*

<b>NAME</b>	drv_hztousec – convert clock ticks to microseconds
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt;  clock_t drv_hztousec(clock_t hertz);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>hertz</b> The number of clock ticks to convert.
<b>DESCRIPTION</b>	<p><b>drv_hztousec()</b> converts into microseconds the time expressed by <i>hertz</i>, which is in system clock ticks.</p> <p>The kernel variable <code>lbolt</code>, whose value should be retrieved by calling <code>ddi_get_lbolt(9F)</code>, is the length of time the system has been up since boot and is expressed in clock ticks. Drivers often use the value of <code>lbolt</code> before and after an I/O request to measure the amount of time it took the device to process the request. <b>drv_hztousec()</b> can be used by the driver to convert the reading from clock ticks to a known unit of time.</p>
<b>RETURN VALUES</b>	The number of microseconds equivalent to the <i>hertz</i> parameter. No error value is returned. If the microsecond equivalent to <i>hertz</i> is too large to be represented as a <code>clock_t</code> , then the maximum <code>clock_t</code> value will be returned.
<b>CONTEXT</b>	<b>drv_hztousec()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<code>ddi_get_lbolt(9F)</code> , <code>drv_usectohz(9F)</code> , <code>drv_usecwait(9F)</code> <i>Writing Device Drivers</i>

<b>NAME</b>	drv_priv – determine driver privilege
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt;  int drv_priv(cred_t *cr);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>cr</b> Pointer to the user credential structure.
<b>DESCRIPTION</b>	<b>drv_priv()</b> provides a general interface to the system privilege policy. It determines whether the credentials supplied by the user credential structure pointed to by <i>cr</i> identify a privileged process. This function should only be used when file access modes and special minor device numbers are insufficient to provide protection for the requested driver function. It is intended to replace all calls to <b>suser()</b> and any explicit checks for effective <code>user ID = 0</code> in driver code.
<b>RETURN VALUES</b>	This routine returns 0 if it succeeds, EPERM if it fails.
<b>CONTEXT</b>	<b>drv_priv()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	drv_usecsthz – convert microseconds to clock ticks
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt;  clock_t drv_usecsthz(clock_t <i>microsecs</i>);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b><i>microsecs</i></b> The number of microseconds to convert.
<b>DESCRIPTION</b>	<p><b>drv_usecsthz()</b> converts a length of time expressed in microseconds to a number of system clock ticks. The time arguments to <b>timeout(9F)</b> and <b>delay(9F)</b> are expressed in clock ticks.</p> <p><b>drv_usecsthz()</b> is a portable interface for drivers to make calls to <b>timeout(9F)</b> and <b>delay(9F)</b> and remain binary compatible should the driver object file be used on a system with a different clock speed (a different number of ticks in a second).</p>
<b>RETURN VALUES</b>	The value returned is the number of system clock ticks equivalent to the <i>microsecs</i> argument. No error value is returned. If the clock tick equivalent to <i>microsecs</i> is too large to be represented as a <code>clock_t</code> , then the maximum <code>clock_t</code> value will be returned.
<b>CONTEXT</b>	<b>drv_usecsthz()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>delay(9F)</b> , <b>drv_hztousec(9F)</b> , <b>timeout(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	drv_usecwait – busy-wait for specified interval
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt;  void drv_usecwait(clock_t <i>microsecs</i>);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b><i>microsecs</i></b> The number of microseconds to busy-wait.
<b>DESCRIPTION</b>	<p><b>drv_usecwait()</b> gives drivers a means of busy-waiting for a specified microsecond count. The amount of time spent busy-waiting may be greater than the microsecond count but will minimally be the number of microseconds specified.</p> <p><b>delay(9F)</b> can be used by a driver to delay for a specified number of system ticks, but it has two limitations. First, the granularity of the wait time is limited to one clock tick, which may be more time than is needed for the delay. Second, <b>delay(9F)</b> may only be invoked from user context and hence cannot be used at interrupt time or system initialization.</p> <p>Often, drivers need to delay for only a few microseconds, waiting for a write to a device register to be picked up by the device. In this case, even in user context, <b>delay(9F)</b> produces too long a wait period.</p>
<b>CONTEXT</b>	<b>drv_usecwait()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>delay(9F)</b> , <b>timeout(9F)</b> , <b>untimeout(9F)</b> <i>Writing Device Drivers</i>
<b>NOTES</b>	The driver wastes processor time by making this call since <b>drv_usecwait()</b> does not block but simply busy-waits. The driver should only make calls to <b>drv_usecwait()</b> as needed, and only for as much time as needed. <b>drv_usecwait()</b> does not mask out interrupts.

<b>NAME</b>	dupb – duplicate a message block descriptor
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  mblk_t *dupb(mblk_t *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>.bp</b> Pointer to the message block to be duplicated. mblk_t is an instance of the <b>msgb(9S)</b> structure.
<b>DESCRIPTION</b>	<p><b>dupb()</b> creates a new mblk_t structure (see <b>msgb(9S)</b>) to reference the message block pointed to by <i>bp</i>.</p> <p>Unlike <b>copyb(9F)</b>, dupb() does not copy the information in the dblk_t structure (see <b>datab(9S)</b>), but creates a new mblk_t structure to point to it. The reference count in the dblk_t structure (db_ref) is incremented. The new mblk_t structure contains the same information as the original. Note that b_rptr and b_wptr are copied from the <i>bp</i>.</p>
<b>RETURN VALUES</b>	If successful, <b>dupb()</b> returns a pointer to the new message block. A NULL pointer is returned if <b>dupb()</b> cannot allocate a new message block descriptor or if the db_ref field of the data block structure (see <b>datab(9S)</b> ) has reached a maximum value (255).
<b>CONTEXT</b>	<b>dupb()</b> can be called from user, kernel, or interrupt context.

**EXAMPLES****EXAMPLE 1** Using **dupb()**

This **srv(9E)** (service) routine adds a header to all **M\_DATA** messages before passing them along. **dupb** is used instead of **copyb(9F)** because the contents of the header block are not changed.

For each message on the queue, if it is a priority message, pass it along immediately (lines 10–11). Otherwise, if it is anything other than an **M\_DATA** message (line 12), and if it can be sent along (line 13), then do so (line 14). Otherwise, put the message back on the queue and return (lines 16–17). For all **M\_DATA** messages, first check to see if the stream is flow-controlled (line 20). If it is, put the message back on the queue and return (lines 37–38). If it is not, the header block is duplicated (line 21).

**dupb()** can fail either due to lack of resources or because the message block has already been duplicated 255 times. In order to handle the latter case, the example calls **copyb(9F)** (line 22). If **copyb(9F)** fails, it is due to buffer allocation failure. In this case, **qbufcall(9F)** is used to initiate a callback (lines 30–31) if one is not already pending (lines 26–27).

The callback function, **xxxcallback()**, clears the recorded **qbufcall(9F)** callback id and schedules the service procedure (lines 49–50). Note that the close routine, **xxxclose()**, must cancel any outstanding **qbufcall(9F)** callback requests (lines 58–59).

If **dupb()** or **copyb(9F)** succeed, link the **M\_DATA** message to the new message block (line 34) and pass it along (line 35).

```

1  xxxsrv(q)$
2      queue_t *q;$
3  {$
4      struct xx *xx = (struct xx *)q->q_ptr;$
5      mblk_t *mp;$
6      mblk_t *bp;$
7      extern mblk_t *hdr;$
8  $
9      while ((mp = getq(q)) != NULL) {$
10         if (mp->b_datap->db_type >= QPCTL) {$
11             putnext(q, mp);$
12         } else if (mp->b_datap->db_type != M_DATA) {$
13             if (canputnext(q))$
14                 putnext(q, mp);$
15             else {$
16                 putbq(q, mp);$
17                 return;$
18             }$
19         } else { /* M_DATA */$
20             if (canputnext(q)) {$
21                 if ((bp = dupb(hdr)) == NULL)$
22                     bp = copyb(hdr);$
23                 if (bp == NULL) {$
24                     size_t size = msgdsize(mp);$
25                     putbq(q, mp);$

```

```

26             if (xx->xx_qbufcall_id) {$
27                 /* qbufcall pending */$
28                 return;$
29             }$
30             xx->xx_qbufcall_id = qbufcall(q, size,$
31                 BPRI_MED, xxxcallback, (intptr_t)q);$
32             return;$
33         }$
34         linkb(bp, mp);$
35         putnext(q, bp);$
36     } else {$
37         putbq(q, mp);$
38         return;$
39     }$
40 }$
41 }$
42 }$
43 void$
44 xxxcallback(q)$
45     queue_t *q;$
46     {$
47         struct xx *xx = (struct xx *)q->q_ptr;$
48$
49         xx->xx_qbufcall_id = 0;$
50         qenable(q);$
51     }$
$
52 xxxclose(q, cflag, crp)$
53     queue_t *q;$
54     int cflag;$
55     cred_t *crp;$
56     {$
57         struct xx *xx = (struct xx *)q->q_ptr;$
58         ...$
59         if (xx->xx_qbufcall_id)$
60             qunbufcall(q, xx->xx_qbufcall_id);$
61         ...$
62     }$

```

**SEE ALSO** [srv\(9E\)](#), [copyb\(9F\)](#), [qbufcall\(9F\)](#), [datab\(9S\)](#), [msgb\(9S\)](#)

*Writing Device Drivers STREAMS Programming Guide*

<b>NAME</b>	dupmsg – duplicate a message
<b>SYNOPSIS</b>	#include <sys/stream.h>  mblk_t *dupmsg(mblk_t *mp);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<i>mp</i> Pointer to the message.
<b>DESCRIPTION</b>	<b>dupmsg()</b> forms a new message by copying the message block descriptors pointed to by <i>mp</i> and linking them. <b>dupb(9F)</b> is called for each message block. The data blocks themselves are not duplicated.
<b>RETURN VALUES</b>	If successful, <b>dupmsg()</b> returns a pointer to the new message block. Otherwise, it returns a NULL pointer. A return value of NULL indicates either memory depletion or the data block reference count, <i>db_ref</i> (see <b>datab(9S)</b> ), has reached a limit (255). See <b>dupb(9F)</b> .
<b>CONTEXT</b>	<b>dupmsg()</b> can be called from user, kernel, or interrupt context.
<b>EXAMPLES</b>	<b>EXAMPLE 1</b> Using <b>dupmsg()</b>  See <b>copyb(9F)</b> for an example using <b>dupmsg()</b> .
<b>SEE ALSO</b>	<b>copyb(9F)</b> , <b>copymsg(9F)</b> , <b>dupb(9F)</b> , <b>datab(9S)</b>  <i>Writing Device Drivers</i>  <i>STREAMS Programming Guide</i>

<b>NAME</b>	enableok – reschedule a queue for service
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void enableok(queue_t *q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> A pointer to the queue to be rescheduled.
<b>DESCRIPTION</b>	<b>enableok()</b> enables queue <i>q</i> to be rescheduled for service. It reverses the effect of a previous call to <b>noenable(9F)</b> on <i>q</i> by turning off the <b>QNOENB</b> flag in the queue.
<b>CONTEXT</b>	<b>enableok()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using <b>enableok()</b></p> <p>The <b>qrestart()</b> routine uses two STREAMS functions to restart a queue that has been disabled. The <b>enableok()</b> function turns off the <b>QNOENB</b> flag, allowing the <b>qenable(9F)</b> to schedule the queue for immediate processing.</p> <pre>1 void 2 qrestart(rdwr_q) 3     register queue_t *rdwr_q; 4 { 5     enableok(rdwr_q); 6     /* re-enable a queue that has been disabled */ 7     (void) qenable(rdwr_q); 8 }</pre>
<b>SEE ALSO</b>	<b>noenable(9F)</b> , <b>qenable(9F)</b> <i>Writing Device Drivers STREAMS Programming Guide</i>

<b>NAME</b>	esballoc – allocate a message block using a caller-supplied buffer
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  mblk_t *esballoc(uchar *base, size_t size, uint_t pri, frtn_t *fr_rtnp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>base</b>            Address of user supplied data buffer.</p> <p><b>size</b>            Number of bytes in data buffer.</p> <p><b>pri</b>             Priority of allocation request (to be used by <b>allocb(9F)</b> function, called by <b>esballoc( )</b>).</p> <p><b>fr_rtnp</b>         Free routine data structure.</p>
<b>DESCRIPTION</b>	<p><b>esballoc()</b> creates a STREAMS message and attaches a user-supplied data buffer in place of a STREAMS data buffer. It calls <b>allocb(9F)</b> to get a message and data block header only. The newly allocated message will have both the <b>b_wptr</b> and <b>b_rptr</b> set to the base of the buffer. As when using <b>allocb(9F)</b>, the newly allocated message will have both <b>b_wptr</b> and <b>b_rptr</b> set to the base of the data buffer. The user-supplied data buffer, pointed to by <b>base</b>, is used as the data buffer for the message.</p> <p>When <b>freeb(9F)</b> is called to free the message, the driver's message freeing routine (referenced through the <b>free_rtn</b> structure) is called, with appropriate arguments, to free the data buffer.</p> <p>The <b>free_rtn</b> structure includes the following members:</p> <pre>void (*free_func)();        /* user's freeing routine */ char *free_arg;            /* arguments to free_func() */</pre> <p>Instead of requiring a specific number of arguments, the <b>free_arg</b> field is defined of type <b>char *</b>. This way, the driver can pass a pointer to a structure if more than one argument is needed.</p> <p>The method by which <b>free_func</b> is called is implementation-specific. The module writer must not assume that <b>free_func</b> will or will not be called directly from STREAMS utility routines like <b>freeb(9F)</b> which free a message block.</p>

	<p><code>free_func</code> must not call another modules put procedure nor attempt to acquire a private module lock which may be held by another thread across a call to a STREAMS utility routine which could free a message block. Otherwise, the possibility for lock recursion and/or deadlock exists.</p> <p><code>free_func</code> must not access any dynamically allocated data structure that might no longer exist when it runs.</p>
<b>RETURN VALUES</b>	On success, a pointer to the newly allocated message block is returned. On failure, <code>NULL</code> is returned.
<b>CONTEXT</b>	<b>esballoc()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>allocb(9F)</b> , <b>freeb(9F)</b> , <b>datadb(9S)</b> , <b>free_rtn(9S)</b> <i>Writing Device Drivers STREAMS Programming Guide</i>
<b>WARNINGS</b>	The <code>free_func</code> must be defined in kernel space, should be declared <code>void</code> and accept one argument. It has no user context and must not sleep.

<b>NAME</b>	esbbscall – call function when buffer is available
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  bufcall_id_t esbbscall(uint_t pri, void (*func)(void *arg), void (arg));</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>pri</b> Priority of allocation request (to be used by <b>allocb(9F)</b> function, called by <b>esbbscall()</b>)</p> <p><b>func</b> Function to be called when buffer becomes available.</p> <p><b>arg</b> Argument to <i>func</i>.</p>
<b>DESCRIPTION</b>	<b>esbbscall()</b> , like <b>bufcall(9F)</b> , serves as a <b>timeout(9F)</b> call of indeterminate length. If <b>esbbsalloc(9F)</b> is unable to allocate a message and data block header to go with its externally supplied data buffer, <b>esbbscall()</b> can be used to schedule the routine <i>func</i> , to be called with the argument <i>arg</i> when a buffer becomes available. <i>func</i> may be a routine that calls <b>esbbsalloc(9F)</b> or it may be another kernel function.
<b>RETURN VALUES</b>	On success, a <b>bufcall</b> ID is returned. On failure, 0 is returned. The value returned from a successful call should be saved for possible future use with <b>unbufcall()</b> should it become necessary to cancel the <b>esbbscall()</b> request (as at driver close time).
<b>CONTEXT</b>	<b>esbbscall()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>allocb(9F)</b> , <b>bufcall(9F)</b> , <b>esbbsalloc(9F)</b> , <b>timeout(9F)</b> , <b>datadb(9S)</b> , <b>unbufcall(9F)</b>
	<i>Writing Device Drivers STREAMS Programming Guide</i>

<b>NAME</b>	flushband – flush messages for a specified priority band
<b>SYNOPSIS</b>	#include <sys/stream.h>  void <b>flushband</b> (queue_t *q, unsigned char pri, int flag);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>q</b> Pointer to the queue.</p> <p><b>pri</b> Priority of messages to be flushed.</p> <p><b>flag</b> Valid <i>flag</i> values are:</p> <p>FLUSHDATA Flush only data messages (types M_DATA, M_DELAY, M_PROTO, and M_PCPROTO).</p> <p>FLUSHALL Flush all messages.</p>
<b>DESCRIPTION</b>	<b>flushband()</b> flushes messages associated with the priority band specified by <i>pri</i> . If <i>pri</i> is 0, only normal and high priority messages are flushed. Otherwise, messages are flushed from the band <i>pri</i> according to the value of <i>flag</i> .
<b>CONTEXT</b>	<b>flushband()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>flushq(9F)</b>  <i>Writing Device Drivers STREAMS Programming Guide</i>

<b>NAME</b>	flushq – remove messages from a queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  void flushq(queue_t *q, int flag);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue to be flushed.  <b>flag</b> Valid <i>flag</i> values are:  FLUSHDATA Flush only data messages (types M_DATA M_DELAY M_PROTO and M_PCPROTO).  FLUSHALL Flush all messages.
<b>DESCRIPTION</b>	<b>flushq()</b> frees messages and their associated data structures by calling <b>freemsg(9F)</b> . If the queue's count falls below the low water mark and the queue was blocking an upstream service procedure, the nearest upstream service procedure is enabled.
<b>CONTEXT</b>	<b>flushq()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<b>EXAMPLE 1</b> Using <b>flushq()</b>  This example depicts the canonical flushing code for STREAMS modules. The module has a write service procedure and potentially has messages on the queue. If it receives an M_FLUSH message, and if the FLUSHR bit is on in the first byte of the message (line 10), then the read queue is flushed (line 11). If the FLUSHW bit is on (line 12), then the write queue is flushed (line 13). Then the message is passed along to the next entity in the stream (line 14). See the example for <b>qreply(9F)</b> for the canonical flushing code for drivers.  <pre> 1  /* 2   * Module write-side put procedure. 3   */ 4  xxxwput(q, mp) 5      queue_t *q; 6      mblk_t *mp; 7  { 8      switch(mp-&gt;b_datap-&gt;db_type) { 9          case M_FLUSH: 10             if (*mp-&gt;b_rptr &amp; FLUSHR) 11                 flushq(RD(q), FLUSHALL); 12             if (*mp-&gt;b_rptr &amp; FLUSHW) 13                 flushq(q, FLUSHALL); </pre>

```
14     putnext(q, mp);
15     break;
      . . .
16   }
17 }
```

**SEE ALSO** flushband(9F), freemsg(9F), putq(9F), qreply(9F)  
*Writing Device Drivers STREAMS Programming Guide*

<b>NAME</b>	freeb – free a message block
<b>SYNOPSIS</b>	#include <sys/stream.h>  void <b>freeb</b> (mblk_t *bp);
<b>PARAMETERS</b>	<b>bp</b> Pointer to the message block to be deallocated. mblk_t is an instance of the <b>msgb</b> (9S) structure.
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>DESCRIPTION</b>	<b>freeb</b> () deallocates a message block. If the reference count of the <code>db_ref</code> member of the <b>data</b> (9S) structure is greater than 1, <b>freeb</b> () decrements the count. If <code>db_ref</code> equals 1, it deallocates the message block and the corresponding data block and buffer.  If the data buffer to be freed was allocated with the <b>esballoc</b> (9F), the buffer may be a non-STREAMS resource. In that case, the driver must be notified that the attached data buffer needs to be freed, and run its own freeing routine. To make this process independent of the driver used in the stream, <b>freeb</b> () finds the <b>free_rtn</b> (9S) structure associated with the buffer. The <code>free_rtn</code> structure contains a pointer to the driver-dependent routine, which releases the buffer. Once this is accomplished, <b>freeb</b> () releases the STREAMS resources associated with the buffer.
<b>CONTEXT</b>	<b>freeb</b> () can be called from user or interrupt context.
<b>EXAMPLES</b>	<b>CODE EXAMPLE 1</b> Using <b>freeb</b> ()  See <b>copyb</b> (9F) for an example of using <b>freeb</b> ()().
<b>SEE ALSO</b>	<b>allocb</b> (9F), <b>copyb</b> (9F), <b>dupb</b> (9F), <b>esballoc</b> (9F), <b>free_rtn</b> (9S)  <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	freemsg – free all message blocks in a message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  void freemsg(mblk_t *mp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>mp</b> Pointer to the message blocks to be deallocated. <code>mblk_t</code> is an instance of the <code>msgb(9S)</code> structure.
<b>DESCRIPTION</b>	<code>freemsg()</code> calls <code>freeb(9F)</code> to free all message and data blocks associated with the message pointed to by <code>mp</code> .
<b>CONTEXT</b>	<code>freemsg()</code> can be called from user or interrupt context.
<b>EXAMPLES</b>	<b>CODE EXAMPLE 1</b> Using <code>freemsg()</code> See <code>copymsg(9F)</code> .
<b>SEE ALSO</b>	<code>copymsg(9F)</code> , <code>freeb(9F)</code> , <code>msgb(9S)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	freerbuf – free a raw buffer header
<b>SYNOPSIS</b>	<pre>#include &lt;sys/buf.h&gt; #include &lt;sys/ddi.h&gt;  void freerbuf(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to a previously allocated buffer header structure.
<b>DESCRIPTION</b>	<b>freerbuf()</b> frees a raw buffer header previously allocated by <b>getrbuf(9F)</b> . This function does not sleep and so may be called from an interrupt routine.
<b>CONTEXT</b>	<b>freerbuf()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>getrbuf(9F)</b> , <b>kmem_alloc(9F)</b> , <b>kmem_free(9F)</b> , <b>kmem_zalloc(9F)</b>

<b>NAME</b>	freeze, unfreeze – freeze, thaw the state of a stream
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void freeze(queue_t * q); void unfreeze(queue_t * q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the message queue to freeze/unfreeze.
<b>DESCRIPTION</b>	<p><b>freeze()</b> freezes the state of the entire stream containing the queue pair <i>q</i>. A frozen stream blocks any thread attempting to enter any open, close, put or service routine belonging to any queue instance in the stream, and blocks any thread currently within the stream if it attempts to put messages onto or take messages off of any queue within the stream (with the sole exception of the caller). Threads blocked by this mechanism remain so until the stream is thawed by a call to <b>unfreeze()</b>.</p> <p>Drivers and modules must freeze the stream before manipulating the queues directly (as opposed to manipulating them through programmatic interfaces such as <b>getq(9F)</b>, <b>putq(9F)</b>, <b>putbq(9F)</b>, etc.)</p>
<b>CONTEXT</b>	These routines may be called from any stream open, close, put or service routine as well as interrupt handlers, callouts and call-backs.
<b>SEE ALSO</b>	<p><b>getq(9F)</b>, <b>insq(9F)</b>, <b>putbq(9F)</b>, <b>putq(9F)</b>, <b>rmvq(9F)</b>, <b>strqget(9F)</b>, <b>strqset(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>
<b>NOTES</b>	<p>Calling <b>freeze()</b> to freeze a stream that is already frozen by the caller will result in a single-party deadlock.</p> <p>The caller of <b>unfreeze()</b> must be the thread who called <b>freeze()</b>.</p> <p>There are usually better ways to accomplish things than by freezing the stream. STREAMS utility functions such as <b>getq(9F)</b>, <b>putq(9F)</b>, <b>putbq(9F)</b>, etc. may not be called by the caller of <b>freeze()</b> while the stream is still frozen, as they indirectly freeze the stream to ensure atomicity of queue manipulation.</p>

<b>NAME</b>	geterror – return I/O error
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt; #include &lt;sys/ddi.h&gt;  int geterror(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to a <b>buf(9S)</b> structure.
<b>DESCRIPTION</b>	<b>geterror()</b> returns the error number from the error field of the buffer header structure.
<b>RETURN VALUES</b>	An error number indicating the error condition of the I/O request is returned. If the I/O request completes successfully, 0 is returned.
<b>CONTEXT</b>	<b>geterror()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>buf(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	getmajor – get major device number
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/mkdev.h&gt; #include &lt;sys/ddi.h&gt;</pre> <p>major_t getmajor(dev_t dev);</p>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>dev</b> Device number.
<b>DESCRIPTION</b>	<b>getmajor()</b> extracts the major number from a device number.
<b>RETURN VALUES</b>	The major number.
<b>CONTEXT</b>	<b>getmajor()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>CODE EXAMPLE 1</b> Using <b>getmajor()</b></p> <p>The following example shows both the <b>getmajor()</b> and <b>getminor(9F)</b> functions used in a debug <b>cmn_err(9F)</b> statement to return the major and minor numbers for the device supported by the driver.</p> <pre>dev_t dev;  #ifdef DEBUG cmn_err(CE_NOTE, "Driver Started. Major# = %d, Minor# = %d", getmajor(dev), getminor(dev)); #endif</pre>
<b>SEE ALSO</b>	<b>cmn_err(9F)</b> , <b>getminor(9F)</b> , <b>makedevice(9F)</b>
<b>WARNINGS</b>	No validity checking is performed. If <i>dev</i> is invalid, an invalid number is returned.

<b>NAME</b>	getminor – get minor device number
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/mkdev.h&gt; #include &lt;sys/ddi.h&gt;  minor_t getminor(dev_t dev);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>dev</b> Device number.
<b>DESCRIPTION</b>	<b>getminor()</b> extracts the minor number from a device number.
<b>RETURN VALUES</b>	The minor number.
<b>CONTEXT</b>	<b>getminor()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See the <b>getmajor(9F)</b> manual page for an example of how to use <b>getminor()</b> .
<b>SEE ALSO</b>	<b>getmajor(9F)</b> , <b>makedevice(9F)</b> <i>Writing Device Drivers</i>
<b>WARNINGS</b>	No validity checking is performed. If <i>dev</i> is invalid, an invalid number is returned.

<b>NAME</b>	get_pktiopb, free_pktiopb – allocate/free a SCSI packet in the iopb map																
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  struct scsi_pkt *get_pktiopb(struct scsi_address * ap, caddr_t * datap, int cdblen, int statuslen, int datalen, int readflag, int (* callback);  void free_pktiopb(struct scsi_pkt * pkt, caddr_t datap, int datalen);</pre>																
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).																
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><b><i>ap</i></b></td> <td>Pointer to the target's <code>scsi_address</code> structure.</td> </tr> <tr> <td style="padding-right: 20px;"><b><i>datap</i></b></td> <td>Pointer to the address of the packet, set by this function.</td> </tr> <tr> <td style="padding-right: 20px;"><b><i>cdblen</i></b></td> <td>Number of bytes required for the SCSI command descriptor block (CDB).</td> </tr> <tr> <td style="padding-right: 20px;"><b><i>statuslen</i></b></td> <td>Number of bytes required for the SCSI status area.</td> </tr> <tr> <td style="padding-right: 20px;"><b><i>datalen</i></b></td> <td>Number of bytes required for the data area of the SCSI command.</td> </tr> <tr> <td style="padding-right: 20px;"><b><i>readflag</i></b></td> <td>If non-zero, data will be transferred from the SCSI target.</td> </tr> <tr> <td style="padding-right: 20px;"><b><i>callback</i></b></td> <td>Pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code></td> </tr> <tr> <td style="padding-right: 20px;"><b><i>pkt</i></b></td> <td>Pointer to a <code>scsi_pkt(9S)</code> structure.</td> </tr> </table>	<b><i>ap</i></b>	Pointer to the target's <code>scsi_address</code> structure.	<b><i>datap</i></b>	Pointer to the address of the packet, set by this function.	<b><i>cdblen</i></b>	Number of bytes required for the SCSI command descriptor block (CDB).	<b><i>statuslen</i></b>	Number of bytes required for the SCSI status area.	<b><i>datalen</i></b>	Number of bytes required for the data area of the SCSI command.	<b><i>readflag</i></b>	If non-zero, data will be transferred from the SCSI target.	<b><i>callback</i></b>	Pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code>	<b><i>pkt</i></b>	Pointer to a <code>scsi_pkt(9S)</code> structure.
<b><i>ap</i></b>	Pointer to the target's <code>scsi_address</code> structure.																
<b><i>datap</i></b>	Pointer to the address of the packet, set by this function.																
<b><i>cdblen</i></b>	Number of bytes required for the SCSI command descriptor block (CDB).																
<b><i>statuslen</i></b>	Number of bytes required for the SCSI status area.																
<b><i>datalen</i></b>	Number of bytes required for the data area of the SCSI command.																
<b><i>readflag</i></b>	If non-zero, data will be transferred from the SCSI target.																
<b><i>callback</i></b>	Pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code>																
<b><i>pkt</i></b>	Pointer to a <code>scsi_pkt(9S)</code> structure.																
<b>DESCRIPTION</b>	<p><b>get_pktiopb()</b> allocates a <code>scsi_pkt</code> structure that has a small data area allocated. It is used by some SCSI commands such as <code>REQUEST_SENSE</code>, which involve a small amount of data and require cache-consistent memory for proper operation. It uses <code>ddi_iopb_alloc(9F)</code> for allocating the data area and <code>scsi_realloc(9F)</code> to allocate the packet and DMA resources.</p> <p><i>callback</i> indicates what <b>get_pktiopb()</b> should do when resources are not available:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>NULL_FUNC</code></td> <td>Do not wait for resources. Return a <code>NULL</code> pointer.</td> </tr> <tr> <td style="padding-right: 20px;"><code>SLEEP_FUNC</code></td> <td>Wait indefinitely for resources.</td> </tr> </table> <p><b>Other Values</b> <i>callback</i> points to a function which is called when resources may have become available. <i>callback</i> must return either 0 (indicating that it attempted to allocate resources but failed to do so again), in which case it is put back on a list to be</p>	<code>NULL_FUNC</code>	Do not wait for resources. Return a <code>NULL</code> pointer.	<code>SLEEP_FUNC</code>	Wait indefinitely for resources.												
<code>NULL_FUNC</code>	Do not wait for resources. Return a <code>NULL</code> pointer.																
<code>SLEEP_FUNC</code>	Wait indefinitely for resources.																

	called again later, or 1 indicating either success in allocating resources or indicating that it no longer cares for a retry.
	<b>free_pktiopb()</b> is used for freeing the packet and its associated resources.
<b>RETURN VALUES</b>	<b>get_pktiopb()</b> returns a pointer to the newly allocated <code>scsi_pkt</code> or a NULL pointer.
<b>CONTEXT</b>	If <i>callback</i> is <code>SLEEP_FUNC</code> , then this routine may only be called from user-level code. Otherwise, it may be called from either user or interrupt level. The <i>callback</i> function may not block or call routines that block.  <b>free_pktiopb()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<code>ddi_iopb_alloc(9F)</code> , <code>scsi_alloc_consistent_buf(9F)</code> , <code>scsi_free_consistent_buf(9F)</code> , <code>scsi_pktalloc(9F)</code> , <code>scsi_realloc(9F)</code> , <code>scsi_pkt(9S)</code>  <i>Writing Device Drivers</i>
<b>NOTES</b>	<b>get_pktiopb()</b> and <b>free_pktiopb()</b> are old functions and should be replaced with <code>scsi_alloc_consistent_buf(9F)</code> and <code>scsi_free_consistent_buf(9F)</code> . <b>get_pktiopb()</b> uses scarce resources. Use it selectively.

<b>NAME</b>	getq – get the next message from a queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  mblk_t *getq(queue_t *q);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue from which the message is to be retrieved.
<b>DESCRIPTION</b>	<b>getq()</b> is used by a service ( <b>srv(9E)</b> ) routine to retrieve its enqueued messages.  A module or driver may include a service routine to process enqueued messages. Once the STREAMS scheduler calls <b>srv()</b> it must process all enqueued messages, unless prevented by flow control. <b>getq()</b> obtains the next available message from the top of the queue pointed to by <i>q</i> . It should be called in a <b>while</b> loop that is exited only when there are no more messages or flow control prevents further processing.  If an attempt was made to write to the queue while it was blocked by flow control, <b>getq()</b> back-enables (restarts) the service routine once it falls below the low water mark.
<b>RETURN VALUES</b>	If there is a message to retrieve, <b>getq()</b> returns a pointer to it. If no message is queued, <b>getq()</b> returns a <b>NULL</b> pointer.
<b>CONTEXT</b>	<b>getq()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See <b>dupb(9F)</b> .
<b>SEE ALSO</b>	<b>srv(9E)</b> , <b>bcanput(9F)</b> , <b>canput(9F)</b> , <b>dupb(9F)</b> , <b>putbq(9F)</b> , <b>putq(9F)</b> , <b>qenable(9F)</b>  <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	getrbuf – get a raw buffer header
<b>SYNOPSIS</b>	<pre>#include &lt;sys/buf.h&gt; #include &lt;sys/kmem.h&gt; #include &lt;sys/ddi.h&gt;</pre> <pre>struct buf *getrbuf(int <i>sleepflag</i>);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b><i>sleepflag</i></b> Indicates whether driver should sleep for free space.
<b>DESCRIPTION</b>	<p><b>getrbuf()</b> allocates the space for a buffer header to the caller. It is used in cases where a block driver is performing raw (character interface) I/O and needs to set up a buffer header that is not associated with the buffer cache.</p> <p><b>getrbuf()</b> calls <b>kmem_alloc(9F)</b> to perform the memory allocation. <b>kmem_alloc()</b> requires the information included in the <i>sleepflag</i> argument. If <i>sleepflag</i> is set to <b>KM_SLEEP</b>, the driver may sleep until the space is freed up. If <i>sleepflag</i> is set to <b>KM_NOSLEEP</b>, the driver will not sleep. In either case, a pointer to the allocated space is returned or <b>NULL</b> to indicate that no space was available.</p>
<b>RETURN VALUES</b>	<b>getrbuf()</b> returns a pointer to the allocated buffer header, or <b>NULL</b> if no space is available.
<b>CONTEXT</b>	<b>getrbuf()</b> can be called from user or interrupt context. (Drivers must not allow <b>getrbuf()</b> to sleep if called from an interrupt routine.)
<b>SEE ALSO</b>	<b>bioinit(9F)</b> , <b>freerbuf(9F)</b> , <b>kmem_alloc(9F)</b> , <b>kmem_free(9F)</b>  <i>Writing Device Drivers</i>

<b>NAME</b>	hat_getkpfnum – get page frame number for kernel address
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  uint_t hat_getkpfnum(caddr_t addr);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 2 (DKI only).
<b>PARAMETERS</b>	<b>addr</b> The kernel virtual address for which the page frame number is to be returned.
<b>DESCRIPTION</b>	<p>hat_getkpfnum( ) returns the page frame number corresponding to the kernel virtual address, <i>addr</i>.</p> <p><i>addr</i> must be a kernel virtual address which maps to device memory. <b>ddi_map_regs(9F)</b> can be used to obtain this address. For example, <b>ddi_map_regs(9F)</b> can be called in the driver's <b>attach(9E)</b> routine. The resulting kernel virtual address can be saved by the driver (see <b>ddi_soft_state(9F)</b>) and used in <b>mmap(9E)</b>. The corresponding <b>ddi_unmap_regs(9F)</b> call can be made in the driver's <b>detach(9E)</b> routine. Refer to <b>mmap(9E)</b> for more information.</p>
<b>RETURN VALUES</b>	The page frame number corresponding to the valid virtual address <i>addr</i> . Otherwise the return value is undefined.
<b>CONTEXT</b>	<b>hat_getkpfnum()</b> can be called only from user or kernel context.
<b>SEE ALSO</b>	<b>attach(9E)</b> , <b>detach(9E)</b> , <b>mmap(9E)</b> , <b>ddi_map_regs(9F)</b> , <b>ddi_soft_state(9F)</b> , <b>ddi_unmap_regs(9F)</b>
<b>NOTES</b>	<p><i>Writing Device Drivers</i></p> <p>For some devices, mapping device memory in the driver's <b>attach(9E)</b> routine and unmapping device memory in the driver's <b>detach(9E)</b> routine is a sizeable drain on system resources. This is especially true for devices with a large amount of physical address space. Refer to <b>mmap(9E)</b> for alternative methods.</p>

<b>NAME</b>	inb, inw, inl, repinsb, repinsw, repinsd – read from an I/O port				
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  unsigned char inb(int port); unsigned short inw(int port); unsigned long inl(int port); void repinsb(int port, unsignedchar *addr, int count); void repinsw(int port, unsignedshort *addr, int count); void repinsd(int port, unsignedlong *addr, int count);</pre>				
<b>INTERFACE LEVEL</b>	Solaris x86 DDI specific (Solaris x86 DDI).				
<b>PARAMETERS</b>	<p><b>port</b>            A valid I/O port address.</p> <p><b>addr</b>            The address of a buffer where the values will be stored.</p> <p><b>count</b>           The number of values to be read from the I/O port.</p>				
<b>DESCRIPTION</b>	<p>These routines read data of various sizes from the I/O port with the address specified by <i>port</i> .</p> <p>The <b>inb()</b> , <b>inw()</b> , and <b>inl()</b> functions read 8 bits, 16 bits, and 32 bits of data respectively, returning the resulting values.</p> <p>The <b>repinsb()</b> , <b>repinsw()</b> , and <b>repinsd()</b> functions read multiple 8-bit, 16-bit, and 32-bit values, respectively. <i>count</i> specifies the number of values to be read. A pointer to a buffer will receive the input data; the buffer must be long enough to hold count values of the requested size.</p>				
<b>RETURN VALUES</b>	<b>inb()</b> , <b>inw()</b> , and <b>inl()</b> return the value that was read from the I/O port.				
<b>CONTEXT</b>	These functions may be called from user or interrupt context.				
<b>ATTRIBUTES</b>	See <b>attributes(5)</b> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>x86</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	x86
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Architecture	x86				

**SEE ALSO** | `eisa(4)`, `isa(4)`, `attributes(5)`, `outb(9F)`

| *Writing Device Drivers*

<b>NAME</b>	insq – insert a message into a queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  int insq(queue_t *q, mblk_t *emp, mblk_t *nmp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>q</b> Pointer to the queue containing message <i>emp</i>.</p> <p><b>emp</b> Enqueued message before which the new message is to be inserted. <i>mblk_t</i> is an instance of the <b>msgb(9S)</b> structure.</p> <p><b>nmp</b> Message to be inserted.</p>
<b>DESCRIPTION</b>	<b>insq()</b> inserts a message into a queue. The message to be inserted, <i>nmp</i> , is placed in <i>q</i> immediately before the message <i>emp</i> . If <i>emp</i> is NULL, the new message is placed at the end of the queue. The queue class of the new message is ignored. All flow control parameters are updated. The service procedure is enabled unless QNOENB is set.
<b>RETURN VALUES</b>	<b>insq()</b> returns 1 on success, and 0 on failure.
<b>CONTEXT</b>	<b>insq()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p>This routine illustrates the steps a transport provider may take to place expedited data ahead of normal data on a queue (assume all M_DATA messages are converted into M_PROTO T_DATA_REQ messages). Normal T_DATA_REQ messages are just placed on the end of the queue (line 16). However, expedited T_EXDATA_REQ messages are inserted before any normal messages already on the queue (line 25). If there are no normal messages on the queue, <i>bp</i> will be NULL and we fall out of the <code>for</code> loop (line 21). <b>insq</b> acts like <b>putq(9F)</b> in this case.</p> <pre> 1  #include 2  #include 3 4  static int 5  xxxwput(queue_t *q, mblk_t *mp) 6  { 7      union T_primitives *tp; 8      mblk_t *bp; 9      union T_primitives *ntp; 10 11     switch (mp-&gt;b_datap-&gt;db_type) {</pre>

```

12 case M_PROTO:
13     tp = (union T_primitives *)mp->b_rptr;
14     switch (tp->type) {
15     case T_DATA_REQ:
16         putq(q, mp);
17         break;
18
19     case T_EXDATA_REQ:
20         freezestr(q);
21         for (bp = q->q_first; bp; bp = bp->b_next) {
22             if (bp->b_datap->db_type == M_PROTO) {
23                 ntp = (union T_primitives *)bp->b_rptr;
24                 if (ntp->type != T_EXDATA_REQ)
25                     break;
26             }
27         }
28         (void)insq(q, bp, mp);
29         unfreezestr(q);
30         break;
31     . . . }
32 }
33 }

```

**SEE ALSO** [freezestr\(9F\)](#), [putq\(9F\)](#), [rmvq\(9F\)](#), [unfreezestr\(9F\)](#), [msgb\(9S\)](#)

*Writing Device Drivers*

*STREAMS Programming Guide*

**WARNINGS** If *emp* is non-NULL, it must point to a message on *q* or a system panic could result.

**NOTES** The stream must be frozen using [freezestr\(9F\)](#) before calling [insq\(\)](#).

<b>NAME</b>	IOC_CONVERT_FROM – determine if there is a need to translate M_IOCTL contents.
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  uint_t IOC_CONVERT_FROM(struct iocblk *iobp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI Specific (Solaris DDI)
<b>PARAMETERS</b>	<b><i>iobp</i></b> A pointer to the M_IOCTL control structure.
<b>DESCRIPTION</b>	The IOC_CONVERT_FROM macro is used to see if the contents of the current M_IOCTL message had its origin in a different C Language Type Model.
<b>RETURN VALUES</b>	<p><b>IOC_CONVERT_FROM()</b> returns the following values:</p> <p>IOC_ILP32 This is an LP64 kernel and the M_IOCTL originated in an ILP32 user process.</p> <p>IOC_NONE The M_IOCTL message uses the same C Language Type Model as this calling module or driver.</p>
<b>CONTEXT</b>	IOC_CONVERT_FROM() can be called from user or interrupt context.
<b>SEE ALSO</b>	<p>ddi_model_convert_from(9F)</p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	kmem_alloc, kmem_zalloc, kmem_free – allocate kernel memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/kmem.h&gt;  void *kmem_alloc(size_t size, int flag); void *kmem_zalloc(size_t size, int flag); void kmem_free(void* buf, size_t size);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>size</b>            Number of bytes to allocate.</p> <p><b>flag</b>            Determines whether caller can sleep for memory. Possible flags are KM_SLEEP to allow sleeping until memory is available, or KM_NOSLEEP to return NULL immediately if memory is not available.</p> <p><b>buf</b>             Pointer to allocated memory.</p>
<b>DESCRIPTION</b>	<p><b>kmem_alloc()</b> allocates <i>size</i> bytes of kernel memory and returns a pointer to the allocated memory. The allocated memory is at least double-word aligned, so it can hold any C data structure. No greater alignment can be assumed. <i>flag</i> determines whether the caller can sleep for memory. KM_SLEEP allocations may sleep but are guaranteed to succeed. KM_NOSLEEP allocations are guaranteed not to sleep but may fail (return NULL) if no memory is currently available. The initial contents of memory allocated using <b>kmem_alloc()</b> are random garbage.</p> <p><b>kmem_zalloc()</b> is like <b>kmem_alloc()</b> but returns zero-filled memory.</p> <p><b>kmem_free()</b> frees previously allocated kernel memory. The buffer address and size must exactly match the original allocation. Memory cannot be returned piecemeal.</p>
<b>RETURN VALUES</b>	If successful, <b>kmem_alloc()</b> and <b>kmem_zalloc()</b> return a pointer to the allocated memory. If KM_NOSLEEP is set and memory cannot be allocated without sleeping, <b>kmem_alloc()</b> and <b>kmem_zalloc()</b> return NULL .
<b>CONTEXT</b>	<b>kmem_alloc()</b> and <b>kmem_zalloc()</b> can be called from interrupt context only if the KM_NOSLEEP flag is set. They can be called from user context with any valid <i>flag</i> . <b>kmem_free()</b> can be called from user or interrupt context.

**SEE ALSO** `copyout(9F)` , `freerbuf(9F)` , `getrbuf(9F)`

*Writing Device Drivers*

**WARNINGS**

Memory allocated using `kmem_alloc()` is not paged. Available memory is therefore limited by the total physical memory on the system. It is also limited by the available kernel virtual address space, which is often the more restrictive constraint on large-memory configurations.

Excessive use of kernel memory is likely to affect overall system performance. Overcommitment of kernel memory will cause the system to hang or panic.

Misuse of the kernel memory allocator, such as writing past the end of a buffer, using a buffer after freeing it, freeing a buffer twice, or freeing a null or invalid pointer, will corrupt the kernel heap and may cause the system to corrupt data or panic.

The initial contents of memory allocated using `kmem_alloc()` are random garbage. This random garbage may include secure kernel data. Therefore, uninitialized kernel memory should be handled carefully. For example, never `copyout(9F)` a potentially uninitialized buffer.

**NOTES**

`kmem_alloc(0, flag)` always returns `NULL` . `kmem_free(NULL, 0)` is legal.

<b>NAME</b>	kstat_create – create and initialize a new kstat
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/kstat.h&gt;</pre> <p>kstat_t *kstat_create(char *module, int instance, char *name, char *class, uchar_t type, ulong_t ndata, uchar_t ks_flag);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>module</b>           The name of the provider's module (such as "sd", "esp", ...). The "core" kernel uses the name "unix".</p> <p><b>instance</b>           The provider's instance number, as from <b>ddi_get_instance(9F)</b>. Modules which do not have a meaningful instance number should use 0.</p> <p><b>name</b>                A pointer to a string that uniquely identifies this structure. Only KSTAT_STRLEN – 1 characters are significant.</p> <p><b>class</b>               The general class that this kstat belongs to. The following classes are currently in use: disk, tape, net, controller, vm, kvm, hat, streams, kstat, and misc.</p> <p><b>type</b>                The type of kstat to allocate. Valid types are:</p> <p style="margin-left: 2em;">KSTAT_TYPE_NAMED Allows more than one data record per kstat.</p> <p style="margin-left: 2em;">KSTAT_TYPE_INTR Interrupt; only one data record per kstat.</p> <p style="margin-left: 2em;">KSTAT_TYPE_IO I/O; only one data record per kstat</p> <p><b>ndata</b>               The number of type-specific data records to allocate.</p> <p><b>flag</b>                A bit-field of various flags for this kstat. <i>flag</i> is some combination of:</p> <p style="margin-left: 2em;">KSTAT_FLAG_VIRTUAL Tells <b>kstat_create()</b> not to allocate memory for the kstat data section; instead, the driver will set the <i>ks_data</i> field to point to the data it wishes to export. This provides a convenient way to export existing data structures.</p>

<b>DESCRIPTION</b>	<p><code>KSTAT_FLAG_WRITABLE</code> Makes the <code>kstat</code> data section writable by root.</p> <p><code>KSTAT_FLAG_PERSISTENT</code> Indicates that this <code>kstat</code> is to be persistent over time. For persistent <code>kstats</code>, <code>kstat_delete(9F)</code> simply marks the <code>kstat</code> as dormant; a subsequent <code>kstat_create()</code> reactivates the <code>kstat</code>. This feature is provided so that statistics are not lost across driver close/open (such as raw disk I/O on a disk with no mounted partitions.) Note: Persistent <code>kstats</code> cannot be virtual, since <code>ks_data</code> points to garbage as soon as the driver goes away.</p> <p><code>kstat_create()</code> is used in conjunction with <code>kstat_install(9F)</code> to allocate and initialize a <code>kstat(9S)</code> structure. The method is generally as follows:</p> <p><code>kstat_create()</code> allocates and performs necessary system initialization of a <code>kstat(9S)</code> structure. <code>kstat_create()</code> allocates memory for the entire <code>kstat</code> (header plus data), initializes all header fields, initializes the data section to all zeroes, assigns a unique <code>kstat</code> ID (KID), and puts the <code>kstat</code> onto the system's <code>kstat</code> chain. The returned <code>kstat</code> is marked invalid because the provider (caller) has not yet had a chance to initialize the data section.</p> <p>After a successful call to <code>kstat_create()</code> the driver must perform any necessary initialization of the data section (such as setting the name fields in a <code>kstat</code> of type <code>KSTAT_TYPE_NAMED</code>). Virtual <code>kstats</code> must have the <code>ks_data</code> field set at this time. The provider may also set the <code>ks_update</code>, <code>ks_private</code>, and <code>ks_lock</code> fields if necessary.</p> <p>Once the <code>kstat</code> is completely initialized, <code>kstat_install(9F)</code> is used to make the <code>kstat</code> accessible to the outside world.</p>
<b>RETURN VALUES</b>	If successful, <code>kstat_create()</code> returns a pointer to the allocated <code>kstat</code> . <code>NULL</code> is returned upon failure.
<b>CONTEXT</b>	<code>kstat_create()</code> can be called from user or kernel context.
<b>EXAMPLES</b>	<p><b>CODE EXAMPLE 1</b> Allocating and Initializing a <code>kstat</code> Structure</p> <pre> pkstat_t *ksp; ksp = kstat_create(module, instance, name, class, type, ndata, flags); if (ksp) {     /* ... provider initialization, if necessary */     kstat_install(ksp); } </pre>

**SEE ALSO**

kstat(3K), ddi\_get\_instance(9F), kstat\_delete(9F),  
kstat\_install(9F), kstat\_named\_init(9F), kstat(9S),  
kstat\_named(9S)

*Writing Device Drivers*

<b>NAME</b>	kstat_delete – remove a kstat from the system
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/kstat.h&gt;  void kstat_delete(kstat_t *ksp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI)
<b>PARAMETERS</b>	<b>ksp</b> Pointer to a currently installed <b>kstat(9S)</b> structure.
<b>DESCRIPTION</b>	<b>kstat_delete()</b> removes <i>ksp</i> from the <i>kstat</i> chain and frees all associated system resources.
<b>RETURN VALUES</b>	None.
<b>CONTEXT</b>	<b>kstat_delete()</b> can be called from any context.
<b>SEE ALSO</b>	<b>kstat_create(9F)</b> , <b>kstat_install(9F)</b> , <b>kstat_named_init(9F)</b> , <b>kstat(9S)</b>  <i>Writing Device Drivers</i>
<b>NOTES</b>	When calling <b>kstat_delete()</b> , the driver must not be holding that <i>kstat</i> 's <i>ks_lock</i> . Otherwise, it may deadlock with a <i>kstat</i> reader.

<b>NAME</b>	kstat_install – add a fully initialized kstat to the system
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/kstat.h&gt;  void kstat_install(kstat_t *ksp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI)
<b>PARAMETERS</b>	<b>ksp</b> Pointer to a fully initialized <b>kstat(9S)</b> structure.
<b>DESCRIPTION</b>	<p><b>kstat_install()</b> is used in conjunction with <b>kstat_create(9F)</b> to allocate and initialize a <b>kstat(9S)</b> structure.</p> <p>After a successful call to <b>kstat_create()</b> the driver must perform any necessary initialization of the data section (such as setting the name fields in a kstat of type <b>KSTAT_TYPE_NAMED</b>). Virtual kstats must have the <b>ks_data</b> field set at this time. The provider may also set the <b>ks_update</b>, <b>ks_private</b>, and <b>ks_lock</b> fields if necessary.</p> <p>Once the kstat is completely initialized, <b>kstat_install</b> is used to make the kstat accessible to the outside world.</p>
<b>RETURN VALUES</b>	None.
<b>CONTEXT</b>	<b>kstat_install()</b> can be called from user or kernel context.
<b>EXAMPLES</b>	<p><b>CODE EXAMPLE 1</b> Allocating and Initializing a kstat Structure</p> <p>The method for allocating and initializing a kstat structure is generally as follows:</p> <pre>kstat_t *ksp; ksp = kstat_create(module, instance, name, class, type, ndata, flags); if (ksp) {     /* ... provider initialization, if necessary */     kstat_install(ksp); }</pre>
<b>SEE ALSO</b>	<p><b>kstat_create(9F)</b>, <b>kstat_delete(9F)</b>, <b>kstat_named_init(9F)</b>, <b>kstat(9S)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	kstat_named_init – initialize a named kstat
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/kstat.h&gt;</pre> <pre>void kstat_named_init(kstat_named_t *knp, char *name, uchar_t data_type);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI)
<b>PARAMETERS</b>	<p><b>knp</b> Pointer to a <b>kstat_named(9S)</b> structure.</p> <p><b>name</b> The name of the statistic.</p> <p><b>data_type</b> The type of value. This indicates which field of the <b>kstat_named(9S)</b> structure should be used. Valid values are:</p> <p style="margin-left: 40px;">KSTAT_DATA_CHAR The "char" field.</p> <p style="margin-left: 40px;">KSTAT_DATA_LONG The "long" field.</p> <p style="margin-left: 40px;">KSTAT_DATA_ULONG The "unsigned long" field.</p> <p style="margin-left: 40px;">KSTAT_DATA_LONGLONG The "long long" field.</p> <p style="margin-left: 40px;">KSTAT_DATA_ULONGLONG The "unsigned long long" field.</p>
<b>DESCRIPTION</b>	<b>kstat_named_init()</b> associates a name and a type with a <b>kstat_named(9S)</b> structure.
<b>RETURN VALUES</b>	None.
<b>CONTEXT</b>	<b>kstat_named_init()</b> can be called from user or kernel context.
<b>SEE ALSO</b>	<b>kstat_create(9F)</b> , <b>kstat_install(9F)</b> , <b>kstat(9S)</b> , <b>kstat_named(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	kstat_queue, kstat_waitq_enter, kstat_waitq_exit, kstat_runq_enter, kstat_runq_exit, kstat_waitq_to_runq, kstat_runq_back_to_waitq - update I/O kstat statistics
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/kstat.h&gt;  void kstat_waitq_enter(kstat_io_t * kiop); void kstat_waitq_exit(kstat_io_t * kiop); void kstat_runq_enter(kstat_io_t * kiop); void kstat_runq_exit(kstat_io_t * kiop); void kstat_waitq_to_runq(kstat_io_t * kiop); void kstat_runq_back_to_waitq(kstat_io_t * kiop);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI)
<b>PARAMETERS</b>	<b>kiop</b> Pointer to a <b>kstat_io(9S)</b> structure.
<b>DESCRIPTION</b>	<p>A large number of I/O subsystems have at least two basic "lists" (or queues) of transactions they manage: one for transactions that have been accepted for processing but for which processing has yet to begin, and one for transactions which are actively being processed (but not done). For this reason, two cumulative time statistics are kept: wait (pre-service) time, and run (service) time.</p> <p>The <b>kstat_queue()</b> family of functions manage these times based on the transitions between the driver wait queue and run queue.</p> <p><b>kstat_waitq_enter()</b></p> <p><b>kstat_waitq_enter()</b> should be called when a request arrives and is placed into a pre-service state (such as just prior to calling <b>disksort(9F)</b> ).</p> <p><b>kstat_waitq_exit()</b></p> <p><b>kstat_waitq_exit()</b> should be used when a request is removed from its pre-service state. (such as just prior to calling the driver's <b>start</b> routine).</p> <p><b>kstat_runq_enter()</b></p>

**kstat\_runq\_enter()** is also called when a request is placed in its service state (just prior to calling the driver's start routine, but after **kstat\_waitq\_exit()** ).

#### **kstat\_runq\_exit()**

**kstat\_runq\_exit()** is used when a request is removed from its service state (just prior to calling **biodone(9F)** ).

#### **kstat\_waitq\_to\_runq()**

**kstat\_waitq\_to\_runq()** transitions a request from the wait queue to the run queue. This is useful wherever the driver would have normally done a **kstat\_waitq\_exit()** followed by a call to **kstat\_runq\_enter()** .

#### **kstat\_runq\_back\_to\_waitq()**

**kstat\_runq\_back\_to\_waitq()** transitions a request from the run queue back to the wait queue. This may be necessary in some cases (write throttling is an example).

#### **RETURN VALUES**

None.

#### **CONTEXT**

**kstat\_create()** can be called from user or kernel context.

#### **WARNINGS**

These transitions must be protected by holding the **kstat**'s **ks\_lock** , and must be completely accurate (all transitions are recorded). Forgetting a transition may, for example, make an idle disk appear 100% busy.

#### **SEE ALSO**

**biodone(9F)** , **disksort(9F)** , **kstat\_create(9F)** , **kstat\_delete(9F)** , **kstat\_named\_init(9F)** , **kstat(9S)** , **kstat\_io(9S)**

*Writing Device Drivers*

<b>NAME</b>	linkb – concatenate two message blocks
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  void linkb(mblk_t *mp1, mblk_t *mp2);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>mp1</b> The message to which <i>mp2</i> is to be added. <code>mblk_t</code> is an instance of the <code>msgb(9S)</code> structure.</p> <p><b>mp2</b> The message to be added.</p>
<b>DESCRIPTION</b>	<p><b>linkb()</b> creates a new message by adding <i>mp2</i> to the tail of <i>mp1</i>. The continuation pointer, <code>b_cont</code>, of <i>mp1</i> is set to point to <i>mp2</i>.</p> <pre> graph LR     mp1["mp1 b_datap b_cont"] --&gt; db_base1["db_base"]     db_base1 --&gt; data_buffer1["data buffer"]     mp1 --&gt; mp2["mp2 b_datap b_cont(0)"]     mp2 --&gt; db_base2["db_base"]     db_base2 --&gt; data_buffer2["data buffer"]   </pre> <p><code>linkb(mp1, mp2);</code></p>
<b>CONTEXT</b>	<b>linkb()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See <code>dupb(9F)</code> for an example of using <b>linkb()</b> .
<b>SEE ALSO</b>	<p><code>dupb(9F)</code>, <code>unlinkb(9F)</code>, <code>msgb(9S)</code></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	makecom, makecom_g0, makecom_g0_s, makecom_g1, makecom_g5 – make a packet for SCSI commands
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  void makecom_g0(struct scsi_pkt * pkt, struct scsi_device * devp, int flag, int cmd, int addr, int cnt);  void makecom_g0_s(struct scsi_pkt * pkt, struct scsi_device * devp, int flag, int cmd, int cnt, int fixbit);  void makecom_g1(struct scsi_pkt * pkt, struct scsi_device * devp, int flag, int cmd, int addr, int cnt);  void makecom_g5(struct scsi_pkt * pkt, struct scsi_device * devp, int flag, int cmd, int addr, int cnt);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>pkt</b> Pointer to an allocated <b>scsi_pkt(9S)</b> structure.</p> <p><b>devp</b> Pointer to the target's <b>scsi_device(9S)</b> structure.</p> <p><b>flag</b> Flags for the <b>pkt_flags</b> member.</p> <p><b>cmd</b> First byte of a group 0 or 1 or 5 SCSI CDB .</p> <p><b>addr</b> Pointer to the location of the data.</p> <p><b>cnt</b> Data transfer length in units defined by the SCSI device type. For sequential devices <b>cnt</b> is the number of bytes. For block devices, <b>cnt</b> is the number of blocks.</p> <p><b>fixbit</b> Fixed bit in sequential access device commands.</p>
<b>DESCRIPTION</b>	<p><b>makecom</b> functions initialize a packet with the specified command descriptor block, <b>devp</b> and transport flags. The <b>pkt_address</b> , <b>pkt_flags</b> , and the command descriptor block pointed to by <b>pkt_cdbp</b> are initialized using the remaining arguments. Target drivers may use <b>makecom_g0()</b> for Group 0 commands (except for sequential access devices), or <b>makecom_g0_s()</b> for Group 0 commands for sequential access devices, or <b>makecom_g1()</b> for Group 1 commands, or <b>makecom_g5()</b> for Group 5 commands. <b>fixbit</b> is used by sequential access devices for accessing fixed block sizes and sets the the tag portion of the SCSI CDB .</p>
<b>CONTEXT</b>	These functions can be called from user or interrupt context.

**EXAMPLES****CODE EXAMPLE 1** Using makecom Functions

```
if (blkno >= (1<<20)) {
    makecom_g1(pkt, SD SCSI_DEVP, pflag, SCMD_WRITE_G1,
              (int) blkno, nblk);
} else {
    makecom_g0(pkt, SD SCSI_DEVP, pflag, SCMD_WRITE,
              (int) blkno, nblk);
}
```

**SEE ALSO****scsi\_device(9S)**, **scsi\_pkt(9S)***ANSI Small Computer System Interface-2 (SCSI-2)**Writing Device Drivers*

<b>NAME</b>	makedevice – make device number from major and minor numbers
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/mkdev.h&gt; #include &lt;sys/ddi.h&gt;  dev_t makedevice(major_t majnum, minor_t minnum);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><i><b>majnum</b></i> Major device number.</p> <p><i><b>minnum</b></i> Minor device number.</p>
<b>DESCRIPTION</b>	<b>makedevice()</b> creates a device number from a major and minor device number. <b>makedevice()</b> should be used to create device numbers so the driver will port easily to releases that treat device numbers differently.
<b>RETURN VALUES</b>	The device number, containing both the major number and the minor number, is returned. No validation of the major or minor numbers is performed.
<b>CONTEXT</b>	<b>makedevice()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>getmajor(9F)</b> , <b>getminor(9F)</b>

<b>NAME</b>	max – return the larger of two integers
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  int max(int int1, int int2);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>int1</b>    The first integer.</p> <p><b>int2</b>    The second integer.</p>
<b>DESCRIPTION</b>	<b>max()</b> compares two signed integers and returns the larger of the two.
<b>RETURN VALUES</b>	The larger of the two numbers.
<b>CONTEXT</b>	<b>max()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>min(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	min – return the lesser of two integers
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  int min(int int1, int int2);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>int1</b>    The first integer.</p> <p><b>int2</b>    The second integer.</p>
<b>DESCRIPTION</b>	<b>min()</b> compares two signed integers and returns the lesser of the two.
<b>RETURN VALUES</b>	The lesser of the two integers.
<b>CONTEXT</b>	<b>min()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>max(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	mkiocb – allocates a STREAMS ioctl block for M_IOCTL messages in the kernel.
<b>SYNOPSIS</b>	#include <sys/stream.h>  mblk_t *mkiocb(uint_t <i>command</i> );
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b><i>command</i></b> The ioctl command for the <i>ioc_cmd</i> field.
<b>DESCRIPTION</b>	<p>STREAMS modules or drivers might need to issue an ioctl to a lower module or driver. The <b>mkiocb()</b> function tries to allocate (using <b>allocb(9F)</b>) a STREAMS M_IOCTL message block ( <b>iocblk(9S)</b>). Buffer allocation fails only when the system is out of memory. If no buffer is available, the <b>qbufcall(9F)</b> function can help a module recover from an allocation failure.</p> <p>The <b>mkiocb</b> function returns a <b>mblk_t</b> structure which is large enough to hold any of the ioctl messages ( <b>iocblk(9S)</b>, <b>copyreq(9S)</b> or <b>copyresp(9S)</b>), and has the following special properties:</p> <p><b>b_wptr</b>                Set to <b>b_rptr + sizeof(struct iocblk)</b> .</p> <p><b>b_cont</b>                Set to <b>NULL</b> .</p> <p><b>b_datap-&gt;db_type</b>    Set to <b>M_IOCTL</b>.</p> <p>The fields in the <b>iocblk</b> structure are initialized as follows:</p> <p><b>ioc_cmd</b>                Set to the command value passed in.</p> <p><b>ioc_id</b>                Set to a unique identifier.</p> <p><b>ioc_cr</b>                Set to point to a credential structure encoding the maximum system privilege and which does not need to be freed in any fashion.</p> <p><b>ioc_count</b>            Set to 0.</p> <p><b>ioc_rval</b>             Set to 0.</p> <p><b>ioc_error</b>            Set to 0.</p> <p><b>ioc_flags</b>            Set to <b>IOC_NATIVE</b> to reflect that this is native to the running kernel.</p>
<b>RETURN VALUES</b>	Upon success, the <b>mkiocb()</b> function returns a pointer to the allocated <b>mblk_t</b> of type <b>M_IOCTL</b> .

On failure, it returns a null pointer.

**CONTEXT**

The **mkiocb()** function can be called from user or interrupt context.

**EXAMPLES****EXAMPLE 1** M\_IOCTL Allocation

The first example shows an M\_IOCTL allocation with the ioctl command TEST\_CMD. If the **iocblk(9S)** cannot be allocated, NULL is returned, indicating an allocation failure (line 5). In line 11, the **putnext(9F)** function is used to send the message downstream.

```

1 test_function(queue_t *q, test_info_t *testinfo)
2 {
3     mblk_t *mp;
4
5     if ((mp = mkiocb(TEST_CMD)) == NULL)
6         return (0);
7
8     /* save off ioctl ID value */
9     testinfo->xx_iocid = ((struct iocblk *)mp->b_rptr)->ioc_id;
10
11     putnext(q, mp);      /* send message downstream */
12     return (1);
13 }
```

**EXAMPLE 2** The ioctl ID Value

During the read service routine, the ioctl ID value for M\_IOCACK or M\_IOCNAK should equal the ioctl that was previously sent by this module before processing.

```

1 test_lrsrv(queue_t *q)
2 {
3     ...
4
5     switch (DB_TYPE(mp)) {
6     case M_IOCACK:
7     case M_IOCNAK:
8         /* Does this match the ioctl that this module sent */
9         ioc = (struct iocblk*)mp->b_rptr;
10        if (ioc->ioc_id == testinfo->xx_iocid) {
11            /* matches, so process the message */
12            ...
13            freemsg(mp);
14        }
15        break;
16    }
17    ...
18 }
```

**EXAMPLE 3** An iocblk Allocation Which Fails

The next example shows an iocblk allocation which fails. Since the open routine is in user context, the caller may block using `qbufcall(9F)` until memory is available.

```

1 test_open(queue_t *q, dev_t devp, int oflag, int sflag, cred_t *credp)
2 {
3     while ((mp = mkiocb(TEST_IOCTL)) == NULL) {
4         int id;
5
6         id = qbufcall(q, sizeof (union ioctypes), BPRI_HI,
7             dummy_callback, 0);
8         /* Handle interrupts */
9         if (!qwait_sig(q)) {
10            qunbufcall(q, id);
11            return (EINTR);
12        }
13    }
14    putnext(q, mp);
15 }
```

**SEE ALSO**

`allocb(9F)`, `putnext(9F)`, `qbufcall(9F)`, `qwait_sig(9F)`, `copyreq(9S)`, `copyresp(9S)`, `iocblk(9S)`

*Writing Device Drivers*

*STREAMS Programming Guide*

**WARNINGS**

It is the module's responsibility to remember the ID value of the `M_IOCTL` that was allocated. This will ensure proper cleanup and ID matching when the `M_IOCACK` or `M_IOCNAK` is received.

<b>NAME</b>	mod_install, mod_remove, mod_info – add, remove or query a loadable module
<b>SYNOPSIS</b>	<pre>#include &lt;sys/modctl.h&gt;  int mod_install(struct modlinkage * modlinkage);  int mod_remove(struct modlinkage * modlinkage);  int mod_info(struct modlinkage * modlinkage, struct modinfo * modinfo);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>modlinkage</b>      Pointer to the loadable module's modlinkage structure which describes what type(s) of module elements are included in this loadable module.</p> <p><b>modinfo</b>          Pointer to the modinfo structure passed to _info(9E) .</p>
<b>DESCRIPTION</b>	<p><b>mod_install()</b> must be called from a module's _init(9E) routine.</p> <p><b>mod_remove()</b> must be called from a module's _fini(9E) routine.</p> <p><b>mod_info()</b> must be called from a module's _info(9E) routine.</p>
<b>RETURN VALUES</b>	<b>mod_install()</b> and <b>mod_remove()</b> return 0 upon success and non-zero on failure. <b>mod_info()</b> returns a non-zero value on success and 0 upon failure.
<b>EXAMPLES</b>	See _init(9E) for an example that uses these functions.
<b>SEE ALSO</b>	<p>_fini(9E) , _info(9E) , _init(9E) , modldrv(9S) , modlinkage(9S) , modlstrmod(9S)</p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	msgdsize – return the number of bytes in a message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  size_t msgdsize(mblk_t *mp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>mp</b> Message to be evaluated.
<b>DESCRIPTION</b>	<b>msgdsize()</b> counts the number of bytes in a data message. Only bytes included in the data blocks of type M_DATA are included in the count.
<b>RETURN VALUES</b>	The number of data bytes in a message, expressed as an integer.
<b>CONTEXT</b>	<b>msgdsize()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See <b>bufcall(9F)</b> for an example that uses <b>msgdsize()</b> .
<b>SEE ALSO</b>	<b>bufcall(9F)</b> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	msgpullup – concatenate bytes in a message
<b>SYNOPSIS</b>	#include <sys/stream.h>  mblk_t *msgpullup(mblk_t *mp, ssize_t len);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>mp</b> Pointer to the message whose blocks are to be concatenated.  <b>len</b> Number of bytes to concatenate.
<b>DESCRIPTION</b>	<b>msgpullup()</b> concatenates and aligns the first <i>len</i> data bytes of the message pointed to by <i>mp</i> , copying the data into a new message. Any remaining bytes in the remaining message blocks will be copied and linked onto the new message. The original message is unaltered. If <i>len</i> equals -1, all data are concatenated. If <i>len</i> bytes of the same message type cannot be found, <b>msgpullup()</b> fails and returns NULL.
<b>RETURN VALUES</b>	msgpullup returns the following values: Non-null       Successful completion. A pointer to the new message is returned.  NULL           An error occurred.
<b>CONTEXT</b>	<b>msgpullup()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	srv(9E), allocb(9F), pullupmsg(9F), msgb(9S)  <i>Writing Device Drivers</i>  <i>STREAMS Programming Guide</i>
<b>NOTES</b>	<b>msgpullup()</b> is a DKI-compliant replacement for the older pullupmsg(9F) routine. Users are strongly encouraged to use <b>msgpullup()</b> instead of pullupmsg(9F).

<b>NAME</b>	mt-streams – STREAMS multithreading								
<b>SYNOPSIS</b>	#include <sys/conf.h>								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).								
<b>DESCRIPTION</b>	<p>STREAMS drivers configures the degree of concurrency using the <code>cb_flag</code> field in the <code>cb_ops</code> structure (see <code>cb_ops(9S)</code>). The corresponding field for STREAMS modules is the <code>f_flag</code> in the <code>fmodsw</code> structure.</p> <p>For the purpose of restricting and controlling the concurrency in drivers/modules, we define the concepts of <i>inner</i> and <i>outer perimeters</i>. A driver/module can be configured either to have no perimeters, to have only an inner or an outer perimeter, or to have both an inner and an outer perimeter. Each perimeter acts as a readers-writers lock, that is, there can be multiple concurrent readers or a single writer. Thus, each perimeter can be entered in two modes: shared (reader) or exclusive (writer). The mode depends on the perimeter configuration and can be different for the different STREAMS entry points ( <code>open(9E)</code>, <code>close(9E)</code>, <code>put(9E)</code>, or <code>srv(9E)</code>).</p> <p>The concurrency for the different entry points is (unless specified otherwise) to enter with exclusive access at the inner perimeter (if present) and shared access at the outer perimeter (if present).</p> <p>The perimeter configuration consists of flags that define the presence and scope of the inner perimeter, the presence of the outer perimeter (which can only have one scope), and flags that modify the default concurrency for the different entry points.</p> <p>All MT safe modules/drivers specify the <code>D_MP</code> flag.</p>								
<b>Inner Perimeter Flags</b>	<p>The inner perimeter presence and scope are controlled by the mutually exclusive flags:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>D_MTPERQ</code></td> <td>The module/driver has an inner perimeter around each queue.</td> </tr> <tr> <td style="padding-right: 20px;"><code>D_MTQPAIR</code></td> <td>The module/driver has an inner perimeter around each read/write pair of queues.</td> </tr> <tr> <td style="padding-right: 20px;"><code>D_MTPERMOD</code></td> <td>The module/driver has an inner perimeter that encloses all the module's/driver's queues.</td> </tr> <tr> <td style="padding-right: 20px;"><b>None of the above</b></td> <td>The module/driver has no inner perimeter.</td> </tr> </table>	<code>D_MTPERQ</code>	The module/driver has an inner perimeter around each queue.	<code>D_MTQPAIR</code>	The module/driver has an inner perimeter around each read/write pair of queues.	<code>D_MTPERMOD</code>	The module/driver has an inner perimeter that encloses all the module's/driver's queues.	<b>None of the above</b>	The module/driver has no inner perimeter.
<code>D_MTPERQ</code>	The module/driver has an inner perimeter around each queue.								
<code>D_MTQPAIR</code>	The module/driver has an inner perimeter around each read/write pair of queues.								
<code>D_MTPERMOD</code>	The module/driver has an inner perimeter that encloses all the module's/driver's queues.								
<b>None of the above</b>	The module/driver has no inner perimeter.								
<b>Outer Perimeter Flags</b>	The outer perimeter presence is configured using:								

`D_MTOUTPERIM` In addition to any inner perimeter, the module/driver has an outer perimeter that encloses all the module's/driver's queues. This can be combined with all the inner perimeter options except `D_MTPERMOD`.

The default concurrency can be modified using:

`D_MTPUTSHARED` This flag modifies the default behavior when `put(9E)` procedure are invoked so that the inner perimeter is entered shared instead of exclusively.

`D_MTOCEXCL` This flag modifies the default behavior when `open(9E)` and `close(9E)` procedures are invoked so the the outer perimeter is entered exclusively instead of shared.

The module/driver can use `qwait(9F)` or `qwait_sig()` in the `open(9E)` and `close(9E)` procedures if it needs to wait "outside" the perimeters.

The module/driver can use `qwriter(9F)` to upgrade the access at the inner or outer perimeter from shared to exclusive.

The use and semantics of `qprocson()` and `qprocsoff(9F)` is independent of the inner and outer perimeters.

#### SEE ALSO

`close(9E)`, `open(9E)`, `put(9E)`, `srv(9E)`, `qprocsoff(9F)`, `qprocson(9F)`, `qwait(9F)`, `qwriter(9F)`, `cb_ops(9S)`

*STREAMS Programming Guide*

*Writing Device Drivers*

<b>NAME</b>	mutex, mutex_enter, mutex_exit, mutex_init, mutex_destroy, mutex_owned, mutex_tryenter – mutual exclusion lock routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ksynch.h&gt;  void mutex_init(kmutex_t * mp, char * name, kmutex_type_t type, void * arg); void mutex_destroy(kmutex_t * mp); void mutex_enter(kmutex_t * mp); void mutex_exit(kmutex_t * mp); int mutex_owned(kmutex_t * mp); int mutex_tryenter(kmutex_t * mp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>mp</b> Pointer to a kernel mutex lock ( kmutex_t ).</p> <p><b>name</b> Descriptive string. This is obsolete and should be NULL . (Non- NULL strings are legal, but they are a waste of kernel memory.)</p> <p><b>type</b> Type of mutex lock.</p> <p><b>arg</b> Type-specific argument for initialization routine.</p>
<b>DESCRIPTION</b>	<p>A mutex enforces a policy of mutual exclusion. Only one thread at a time may hold a particular mutex. Threads trying to lock a held mutex will block until the mutex is unlocked.</p> <p>Mutexes are strictly bracketing and may not be recursively locked. That is to say, mutexes should be exited in the opposite order they were entered, and cannot be reentered before exiting.</p> <p><b>mutex_init()</b> initializes a mutex. It is an error to initialize a mutex more than once. The <i>type</i> argument should be set to MUTEX_DRIVER .</p> <p><i>arg</i> provides type-specific information for a given variant type of mutex. When <b>mutex_init()</b> is called for driver mutexes, if the mutex is used by the interrupt handler, the <i>arg</i> should be the <code>ddi_iblock_cookie</code> returned from <b>ddi_get_iblock_cookie(9F)</b> or <b>ddi_get_soft_iblock_cookie(9F)</b> . If the mutex is never used inside an interrupt handler, the argument should be NULL .</p>

**mutex\_enter()** is used to acquire a mutex. If the mutex is already held, then the caller blocks. After returning, the calling thread is the owner of the mutex. If the mutex is already held by the calling thread, a panic will ensue.

**mutex\_owned()** should only be used in **ASSERT()** and may be enforced by not being defined unless the preprocessor symbol **DEBUG** is defined. Its return value is non-zero if the current thread (or, if that cannot be determined, at least some thread) holds the mutex pointed to by *mp*.

**mutex\_tryenter()** is very similar to **mutex\_enter()** except that it doesn't block when the mutex is already held. **mutex\_tryenter()** returns non-zero when it acquired the mutex and 0 when the mutex is already held.

**mutex\_exit()** releases a mutex and will unblock another thread if any are blocked on the mutex.

**mutex\_destroy()** releases any resources that might have been allocated by **mutex\_init()**. **mutex\_destroy()** must be called before freeing the memory containing the mutex, and should be called with the mutex unheld (not owned by any thread). The caller must somehow be sure that no other thread will attempt to use the mutex.

## RETURN VALUES

**mutex\_tryenter()** returns non-zero on success and zero of failure.

**mutex\_owned()** returns non-zero if the calling thread currently holds the mutex pointed to by *mp*, or when that cannot be determined, if any thread holds the mutex. **mutex\_owned()** returns zero otherwise.

## CONTEXT

These functions can be called from user, kernel, or high-level interrupt context, except for **mutex\_init()** and **mutex\_destroy()**, which can be called from user or kernel context only.

## EXAMPLES

### CODE EXAMPLE 1 Initializing a Mutex

A driver might do this to initialize a mutex that is part of its unit structure and used in its interrupt routine:

```
ddi_get_iblock_cookie(dip, 0, &iblock);
mutex_init(&un->un_lock, NULL, MUTEX_DRIVER,
\011      (void *)iblock);
ddi_add_intr(dip, 0, NULL, &dev_cookie, xxintr,
\011      (caddr_t)un);
```

### CODE EXAMPLE 2 Calling a Routine with a Lock

A routine that expects to be called with a certain lock held might have the following **ASSERT**:

```
xxstart(struct xxunit *un)
{
\011    ASSERT(mutex_owed(&un->un_lock));
...
}
```

**SEE ALSO**

**lockstat(1M)** , **condvar(9F)** , **ddi\_add\_intr(9F)** ,  
**ddi\_get\_iblock\_cookie(9F)** , **ddi\_get\_soft\_iblock\_cookie(9F)** ,  
**rwlock(9F)** , **semaphore(9F)**

*Writing Device Drivers*

**NOTES**

Compiling with `_LOCKTEST` or `_MPSTATS` defined no longer has any effect. To gather lock statistics, see **lockstat(1M)** .

<b>NAME</b>	nochpoll – error return function for non-pollable devices
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>int nochpoll(dev_t dev, short events, int anyyet, short *reventsp, struct pollhead **pollhdrp);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dev</b> Device number.</p> <p><b>events</b> Event flags.</p> <p><b>anyyet</b> Check current events only.</p> <p><b>reventsp</b> Event flag pointer.</p> <p><b>pollhdrp</b> Poll head pointer.</p>
<b>DESCRIPTION</b>	<b>nochpoll()</b> is a routine that simply returns the value <code>ENXIO</code> . It is intended to be used in the <code>cb_ops(9S)</code> structure of a device driver for devices that do not support the <code>poll(2)</code> system call.
<b>RETURN VALUES</b>	<b>nochpoll()</b> returns <code>ENXIO</code> .
<b>CONTEXT</b>	<b>nochpoll()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<code>poll(2)</code> , <code>chpoll(9E)</code> , <code>cb_ops(9S)</code> <i>Writing Device Drivers</i>

<b>NAME</b>	nodev – error return function
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt;  int nodev();</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>DESCRIPTION</b>	<b>nodev()</b> returns ENXIO. It is intended to be used in the <b>cb_ops(9S)</b> data structure of a device driver for device entry points which are not supported by the driver. That is, it is an error to attempt to call such an entry point.
<b>RETURN VALUES</b>	<b>nodev()</b> returns ENXIO.
<b>CONTEXT</b>	<b>nodev()</b> can be only called from user context.
<b>SEE ALSO</b>	<b>nulldev(9F)</b> , <b>cb_ops(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	noenable – prevent a queue from being scheduled
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void noenable(queue_t *q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue.
<b>DESCRIPTION</b>	<b>noenable()</b> prevents the queue <i>q</i> from being scheduled for service by <b>insq(9F)</b> , <b>putq(9F)</b> or <b>putbq(9F)</b> when enqueueing an ordinary priority message. The queue can be re-enabled with the <b>enableok(9F)</b> function.
<b>CONTEXT</b>	<b>noenable()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>enableok(9F)</b> , <b>insq(9F)</b> , <b>putbq(9F)</b> , <b>putq(9F)</b> , <b>qenable(9F)</b> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	nulldev – zero return function
<b>SYNOPSIS</b>	#include <sys/conf.h> #include <sys/ddi.h>  int nulldev();
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>DESCRIPTION</b>	<b>nulldev()</b> returns 0. It is intended to be used in the <b>cb_ops(9S)</b> data structure of a device driver for device entry points that do nothing.
<b>RETURN VALUES</b>	<b>nulldev()</b> returns a 0.
<b>CONTEXT</b>	<b>nulldev()</b> can be called from any context.
<b>SEE ALSO</b>	<b>nodev(9F)</b> , <b>cb_ops(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	OTHERQ, otherq – get pointer to queue’s partner queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  queue_t * OTHERQ(queue_t * q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue.
<b>DESCRIPTION</b>	The <b>OTHERQ()</b> function returns a pointer to the other of the two <code>queue</code> structures that make up a STREAMS module or driver. If <code>q</code> points to the read queue the write queue will be returned, and vice versa.
<b>RETURN VALUES</b>	<b>OTHERQ()</b> returns a pointer to a queue’s partner.
<b>CONTEXT</b>	<b>OTHERQ()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Setting Queues</p> <p>This routine sets the minimum packet size, the maximum packet size, the high water mark, and the low water mark for the read and write queues of a given module or driver. It is passed either one of the queues. This could be used if a module or driver wished to update its queue parameters dynamically.</p> <pre>1 void 2 set_q_params(q, min, max, hi, lo) 3     queue_t *q; 4     short min; 5     short max; 6     ushort_t hi; 7     ushort_t lo; 8 { 9\011     q-&gt;q_minpsz = min; 10\011    q-&gt;q_maxpsz = max; 11\011    q-&gt;q_hiwat = hi; 12 \011    q-&gt;q_lowat = lo; 13\011    OTHERQ(q)-&gt;q_minpsz = min; 14\011    OTHERQ(q)-&gt;q_maxpsz = max; 15\011    OTHERQ(q)-&gt;q_hiwat = hi; 16\011    OTHERQ(q)-&gt;q_lowat = lo; 17 }</pre>
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

*STREAMS Programming Guide*

<b>NAME</b>	outb, outw, outl, repoutsb, repoutsw, repoutsd – write to an I/O port				
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  void outb(int port, unsignedchar value); void outw(int port, unsignedshort value); void outl(int port, unsignedlong value); void repoutsb(int port, unsignedchar *addr, int count); void repoutsw(int port, unsignedshort *addr, int count); void repoutsd(int port, unsignedlong *addr, int count);</pre>				
<b>INTERFACE LEVEL</b>	Solaris x86 DDI specific (Solaris x86 DDI).				
<b>PARAMETERS</b>	<p><b>port</b>            A valid I/O port address.</p> <p><b>value</b>           The data to be written to the I/O port.</p> <p><b>addr</b>            The address of a buffer from which the values will be fetched.</p> <p><b>count</b>           The number of values to be written to the I/O port.</p>				
<b>DESCRIPTION</b>	<p>These routines write data of various sizes to the I/O port with the address specified by <i>port</i> .</p> <p>The <b>outb()</b> , <b>outw()</b> , and <b>outl()</b> functions write 8 bits, 16 bits, and 32 bits of data respectively, writing the data specified by <i>value</i> .</p> <p>The <b>repoutsb()</b> , <b>repoutsw()</b> , and <b>repoutsd()</b> functions write multiple 8-bit, 16-bit, and 32-bit values, respectively. <i>count</i> specifies the number of values to be written. <i>addr</i> is a pointer to a buffer from which the output values are fetched.</p>				
<b>CONTEXT</b>	These functions may be called from user or interrupt context.				
<b>ATTRIBUTES</b>	See <b>attributes(5)</b> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>x86</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	x86
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Architecture	x86				

**SEE ALSO**    `eisa(4)` , `isa(4)` , `attributes(5)` , `inb(9F)`

*Writing Device Drivers*

<b>NAME</b>	pci_config_get8, pci_config_get16, pci_config_get32, pci_config_get64, pci_config_put8, pci_config_put16, pci_config_put32, pci_config_put64, pci_config_getb, pci_config_getl, pci_config_getll, pci_config_getw, pci_config_putb, pci_config_putl, pci_config_putll, pci_config_putw – read or write single datum of various sizes to the PCI Local Bus Configuration space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  uint8_t pci_config_get8(ddi_acc_handle_t handle, off_t offset); uint16_t pci_config_get16(ddi_acc_handle_t handle, off_t offset); uint32_t pci_config_get32(ddi_acc_handle_t handle, off_t offset); uint64_t pci_config_get64(ddi_acc_handle_t handle, off_t offset); void pci_config_put8(ddi_acc_handle_t handle, off_t offset, uint8_t value); void pci_config_put16(ddi_acc_handle_t handle, off_t offset, uint16_t value); void pci_config_put32(ddi_acc_handle_t handle, off_t offset, uint32_t value); void pci_config_put64(ddi_acc_handle_t handle, off_t offset, uint64_t value);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>handle</b>            The data access handle returned from <code>pci_config_setup(9F)</code> .</p> <p><b>offset</b>            Byte offset from the beginning of the PCI Configuration space.</p> <p><b>value</b>             Output data.</p>
<b>DESCRIPTION</b>	<p>These routines read or write a single datum of various sizes from or to the PCI Local Bus Configuration space. The <code>pci_config_get8()</code> , <code>pci_config_get16()</code> , <code>pci_config_get32()</code> , and <code>pci_config_get64()</code> functions read 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively. The <code>pci_config_put8()</code> , <code>pci_config_put16()</code> , <code>pci_config_put32()</code> , and <code>pci_config_put64()</code> functions write 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively. The <i>offset</i> argument must be a multiple of the datum size.</p> <p>Since the PCI Local Bus Configuration space is represented in little endian data format, these functions translate the data from or to native host format to or from little endian format.</p>

`pci_config_setup(9F)` must be called before invoking these functions.

**RETURN VALUES**

`pci_config_get8()` , `pci_config_get16()` , `pci_config_get32()` , and `pci_config_get64()` return the value read from the PCI Local Bus Configuration space.

**CONTEXT**

These routines can be called from user, kernel, or interrupt context.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	PCI Local Bus

**SEE ALSO**

`attributes(5)` , `pci_config_setup(9F)` , `pci_config_takedown(9F)`

**NOTES**

These functions are specific to PCI bus device drivers. For drivers using these functions, a single source to support devices with multiple bus versions may not be easy to maintain.

The functions described in this manual page previously used symbolic names which specified their data access size; the function names have been changed so they now specify a fixed-width data size. See the following table for the new name equivalents:

Previous Name	New Name
<code>pci_config_getb</code>	<code>pci_config_get8</code>
<code>pci_config_getw</code>	<code>pci_config_get16</code>
<code>pci_config_getl</code>	<code>pci_config_get32</code>
<code>pci_config_getll</code>	<code>pci_config_get64</code>
<code>pci_config_putb</code>	<code>pci_config_put8</code>
<code>pci_config_putw</code>	<code>pci_config_put16</code>
<code>pci_config_putl</code>	<code>pci_config_put32</code>
<code>pci_config_putll</code>	<code>pci_config_put64</code>

<b>NAME</b>	pci_config_setup, pci_config_tearardown – setup or tear down the resources for enabling accesses to the PCI Local Bus Configuration space				
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int pci_config_setup(dev_info_t *dip, ddi_acc_handle_t *handle); void pci_config_tearardown(ddi_acc_handle_t *handle);</pre>				
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).				
<b>PARAMETERS</b>	<p><b>dip</b>                    Pointer to the device's dev_info structure.</p> <p><b>handle</b>                Pointer to a data access handle.</p>				
<b>DESCRIPTION</b>	<p><b>pci_config_setup()</b> sets up the necessary resources for enabling subsequent data accesses to the PCI Local Bus Configuration space.</p> <p><b>pci_config_tearardown()</b> reclaims and removes those resources represented by the data access handle returned from <b>pci_config_setup()</b> .</p>				
<b>RETURN VALUES</b>	<p><b>pci_config_setup()</b> returns:</p> <p>DDI_SUCCESS    Successfully setup the resources.</p> <p>DDI_FAILURE    Unable to allocate resources for setup.</p>				
<b>CONTEXT</b>	<p><b>pci_config_setup()</b> must be called from user or kernel context.</p> <p><b>pci_config_tearardown()</b> can be called from any context.</p>				
<b>NOTES</b>	These functions are specific to PCI bus device drivers. For drivers using these functions, a single source to support devices with multiple bus versions may not be easy to maintain.				
<b>ATTRIBUTES</b>	See <b>attributes(5)</b> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>PCI Local Bus</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	PCI Local Bus
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Architecture	PCI Local Bus				
<b>SEE ALSO</b>	<p><b>attributes(5)</b></p> <p><i>IEEE 1275 PCI Bus Binding</i></p>				

<b>NAME</b>	physio, minphys – perform physical I/O
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/buf.h&gt; #include &lt;sys/uio.h&gt;  int <b>physio</b>(int (* <i>strat</i>)(struct buf *), struct buf *<i>bp</i>, dev_t <i>dev</i>, int <i>rw</i>, void (* <i>mincnt</i>)(struct buf *), struct uio *<i>uio</i>);  void <b>minphys</b>(struct buf *<i>bp</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
<b>physio()</b>	<p><b><i>strat</i></b>            Pointer to device strategy routine.</p> <p><b><i>bp</i></b>                Pointer to a <b>buf</b>(9S) structure describing the transfer. If <i>bp</i> is set to NULL then <b>physio()</b> allocates one which is automatically released upon completion.</p> <p><b><i>dev</i></b>                The device number.</p> <p><b><i>rw</i></b>                Read/write flag. This is either B_READ when reading from the device, or B_WRITE when writing to the device.</p> <p><b><i>mincnt</i></b>            Routine which bounds the maximum transfer unit size.</p> <p><b><i>uio</i></b>                Pointer to the <b>uio</b> structure which describes the user I/O request.</p>
<b>minphys()</b>	<p><b><i>bp</i></b>                Pointer to a <b>buf</b> structure.</p>
<b>DESCRIPTION</b>	<p><b>physio()</b> performs unbuffered I/O operations between the device <i>dev</i> and the address space described in the <b>uio</b> structure.</p> <p>Prior to the start of the transfer <b>physio()</b> verifies the requested operation is valid by checking the protection of the address space specified in the <b>uio</b> structure. It then locks the pages involved in the I/O transfer so they can not be paged out. The device strategy routine, <b>strat()</b>, is then called one or more times to perform the physical I/O operations. <b>physio()</b> uses <b>biowait</b>(9F) to block until <b>strat()</b> has completed each transfer. Upon completion, or detection of an error, <b>physio()</b> unlocks the pages and returns the error status.</p>

**physio()** uses **mincnt()** to bound the maximum transfer unit size to the system, or device, maximum length. **minphys()** is the system **mincnt()** routine for use with **physio()** operations. Drivers which do not provide their own local **mincnt()** routines should call **physio()** with **minphys()** .

**minphys()** limits the value of *bp* ->*b\_bcount* to a sensible default for the capabilities of the system. Drivers that provide their own **mincnt()** routine should also call **minphys()** to make sure they do not exceed the system limit.

**RETURN VALUES**

**physio()** returns:

0                    Upon success.

non-zero            Upon failure.

**CONTEXT**

**physio()** can be called from user context only.

**SEE ALSO**

**strategy(9E)** , **biodone(9F)** , **biowait(9F)** , **buf(9S)** , **uio(9S)**

*Writing Device Drivers*

**WARNINGS**

Since **physio()** calls **biowait()** to block until each buf transfer is complete, it is the drivers responsibility to call **biodone(9F)** when the transfer is complete, or **physio()** will block forever.

<b>NAME</b>	pm_busy_component, pm_idle_component – control device components' availability for power management
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int pm_busy_component(dev_info_t *dip, int component); int pm_idle_component(dev_info_t *dip, int component);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
pm_busy_component()	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>component</b> The number of the component to be power-managed.</p>
pm_idle_component()	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>component</b> The number of the component to be power-managed.</p>
<b>DESCRIPTION</b>	<p>The <b>pm_busy_component()</b> function sets <i>component</i> of <i>dip</i> to be busy. Calls to <b>pm_busy_component()</b> are stacked, requiring a corresponding number of calls to <b>pm_idle_component()</b> to make the component idle again. When a device is busy it will not be power-managed by the system.</p> <p>The <b>pm_idle_component()</b> function marks <i>component</i> idle, recording the time that <i>component</i> went idle. This function must be called once for each call to <b>pm_busy_component()</b>. A component which is idle is available to be power-managed by the system. The <b>pm_idle_component()</b> function has no effect if the component is already idle, except to update the system's notion of when the device went idle.</p>
<b>RETURN VALUES</b>	<p>The <b>pm_busy_component()</b> and <b>pm_idle_component()</b> functions return:</p> <p>DDI_SUCCESS Successfully set the indicated component busy or idle.</p> <p>DDI_FAILURE Invalid component number <i>component</i> or the device has no components.</p>
<b>CONTEXT</b>	These functions can be called from user or kernel context.

**SEE ALSO**

`power.conf(4)` , `pm(7D)` , `pm(9E)` , `pm_create_components(9F)` ,  
`pm_destroy_components(9F)`

*Writing Device Drivers*

<b>NAME</b>	pm_create_components, pm_destroy_components – create or destroy power-manageable components
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int pm_create_components(dev_info_t * dip, int components); void pm_destroy_components(dev_info_t * dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>                    Pointer to the device's dev_info structure.</p> <p><b>components</b>            The number of components to create.</p>
<b>DESCRIPTION</b>	<p>The <b>pm_create_components()</b> function creates power-manageable components for a device. It should be called from the driver's <b>attach(9E)</b> entry point if the device has power-manageable components.</p> <p>The correspondence of components to parts of the physical device controlled by the driver are the responsibility of the driver. Component 0 must represent the entire device. Components 1- <i>n</i> are driver-defined.</p> <p>The <b>pm_destroy_components()</b> function removes all components from the device. It should be called from the driver's <b>detach(9E)</b> entry point.</p>
<b>RETURN VALUES</b>	<p>The <b>pm_create_components()</b> function returns:</p> <p>DDI_SUCCESS    Components are successfully created.</p> <p>DDI_FAILURE    The device already has components.</p>
<b>CONTEXT</b>	These functions may be called from user or kernel context.
<b>SEE ALSO</b>	<p><b>power.conf(4)</b> , <b>pm(7D)</b> , <b>attach(9E)</b> , <b>detach(9E)</b> , <b>pm(9E)</b> , <b>pm_busy_component(9F)</b> , <b>pm_idle_component(9F)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	pm_get_normal_power, pm_set_normal_power – get or set a device component's normal power level
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int pm_get_normal_power(dev_info_t *dip, int component); void pm_set_normal_power(dev_info_t *dip, int component, int level);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
pm_get_normal_power()	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>component</b> Number of component to get normal power level of.</p>
pm_set_normal_power()	<p><b>dip</b> Pointer to the device's dev_info structure.</p> <p><b>component</b> Number of component to set normal power level for.</p> <p><b>level</b> Power level to become the component's new normal power level.</p>
<b>DESCRIPTION</b>	<p>The <b>pm_get_normal_power()</b> function returns the normal power level of <i>component</i> of the device <i>dip</i> .</p> <p>The <b>pm_set_normal_power()</b> function sets the normal power level of <i>component</i> of the device <i>dip</i> to <i>level</i> .</p> <p>When a device has been power-managed by <b>pm</b>(7D) and is being returned to a state to be used by the system, it will be brought to its normal power level. Except for a power level of 0, which is defined by the system to mean "powered off," or a power level in the range 1-15, which are reserved, the interpretation of the meaning of the power level is entirely up to the driver.</p>
<b>RETURN VALUES</b>	<p>The <b>pm_get_normal_power()</b> function returns:</p> <p><b>level</b> The normal power level of the specified component (a positive integer).</p>

DDI\_FAILURE Invalid component number *component* or the device has no components.

**CONTEXT**

These functions can be called from user or kernel context.

**SEE ALSO**

`power.conf(4)`, `pm(7D)`, `pm(9E)`, `power(9E)`, `pm_busy_component(9F)`, `pm_create_components(9F)`, `pm_destroy_components(9F)`, `pm_idle_component(9F)`

*Writing Device Drivers*

<b>NAME</b>	pollwakeupp – inform a process that an event has occurred
<b>SYNOPSIS</b>	<pre>#include &lt;sys/poll.h&gt;  void pollwakeupp(struct pollhead *php, short event);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>php</b> Pointer to a pollhead structure.</p> <p><b>event</b> Event to notify the process about.</p>
<b>DESCRIPTION</b>	<p><b>pollwakeupp()</b> wakes a process waiting on the occurrence of an event. It should be called from a driver for each occurrence of an event. The pollhead structure will usually be associated with the driver's private data structure associated with the particular minor device where the event has occurred. See <a href="#">chpoll(9E)</a> and <a href="#">poll(2)</a> for more detail.</p>
<b>CONTEXT</b>	<b>pollwakeupp()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<a href="#">poll(2)</a> , <a href="#">chpoll(9E)</a> <i>Writing Device Drivers</i>
<b>NOTES</b>	Driver defined locks should not be held across calls to this function.

<b>NAME</b>	proc_signal, proc_ref, proc_unref – send a signal to a process														
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt; #include &lt;sys/signal.h&gt;  void *proc_ref(void);  void proc_unref(void *pref);  int proc_signal(void *pref, int sig);</pre>														
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).														
<b>PARAMETERS</b>	<p><b>pref</b>            A handle for the process to be signalled.</p> <p><b>sig</b>             Signal number to be sent to the process.</p>														
<b>DESCRIPTION</b>	<p>This set of routines allows a driver to send a signal to a process. The routine <b>proc_ref()</b> is used to retrieve an unambiguous reference to the process for signalling purposes. The return value can be used as a unique handle on the process, even if the process dies. Because system resources are committed to a process reference, <b>proc_unref()</b> should be used to remove it as soon as it is no longer needed. <b>proc_signal()</b> is used to send signal <i>sig</i> to the referenced process. The following set of signals may be sent to a process from a driver:</p> <table border="0"> <tr> <td>SIGHUP</td> <td>The device has been disconnected.</td> </tr> <tr> <td>SIGINT</td> <td>The interrupt character has been received.</td> </tr> <tr> <td>SIGQUIT</td> <td>The quit character has been received.</td> </tr> <tr> <td>SIGPOLL</td> <td>A pollable event has occurred.</td> </tr> <tr> <td>SIGKILL</td> <td>Kill the process (cannot be caught or ignored).</td> </tr> <tr> <td>SIGWINCH</td> <td>Window size change.</td> </tr> <tr> <td>SIGURG</td> <td>Urgent data are available.</td> </tr> </table> <p>See <b>signal(5)</b> for more details on the meaning of these signals.</p> <p>If the process has exited at the time the signal was sent, <b>proc_signal()</b> returns an error code; the caller should remove the reference on the process by calling <b>proc_unref()</b> .</p>	SIGHUP	The device has been disconnected.	SIGINT	The interrupt character has been received.	SIGQUIT	The quit character has been received.	SIGPOLL	A pollable event has occurred.	SIGKILL	Kill the process (cannot be caught or ignored).	SIGWINCH	Window size change.	SIGURG	Urgent data are available.
SIGHUP	The device has been disconnected.														
SIGINT	The interrupt character has been received.														
SIGQUIT	The quit character has been received.														
SIGPOLL	A pollable event has occurred.														
SIGKILL	Kill the process (cannot be caught or ignored).														
SIGWINCH	Window size change.														
SIGURG	Urgent data are available.														

**RETURN VALUES**

The driver writer must ensure that for each call made to **proc\_ref()** , there is exactly one corresponding call to **proc\_unref()** .

**proc\_ref()** returns the following:

**pref** An opaque handle used to refer to the current process.

**proc\_signal()** returns the following:

0 The process existed before the signal was sent.

-1 The process no longer exists; no signal was sent.

**CONTEXT**

**proc\_unref()** and **proc\_signal()** can be called from user or interrupt context. **proc\_ref()** should only be called from user context.

**SEE ALSO**

**signal(5)** , **putnextctl1(9F)**

*Writing Device Drivers*

<b>NAME</b>	ptob – convert size in pages to size in bytes
<b>SYNOPSIS</b>	#include <sys/ddi.h>  unsigned long <b>ptob</b> (unsigned long <i>numpages</i> );
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b><i>numpages</i></b> Size in number of pages to convert to size in bytes.
<b>DESCRIPTION</b>	This function returns the number of bytes that are contained in the specified number of pages. For example, if the page size is 2048, then <b>ptob(2)</b> returns 4096. <b>ptob(0)</b> returns 0.
<b>RETURN VALUES</b>	The return value is always the number of bytes in the specified number of pages. There are no invalid input values, and no checking will be performed for overflow in the case of a page count whose corresponding byte count cannot be represented by an <code>unsigned long</code> . Rather, the higher order bits will be ignored.
<b>CONTEXT</b>	<b>ptob()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>btop(9F)</b> , <b>btopr(9F)</b> , <b>ddi_ptob(9F)</b>  <i>Writing Device Drivers</i>

<b>NAME</b>	pullupmsg – concatenate bytes in a message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  int pullupmsg(mblk_t *mp, ssize_t len);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>mp</b> Pointer to the message whose blocks are to be concatenated. <code>mblk_t</code> is an instance of the <code>msgb(9S)</code> structure.</p> <p><b>len</b> Number of bytes to concatenate.</p>
<b>DESCRIPTION</b>	<p><b>pullupmsg()</b> tries to combine multiple data blocks into a single block. <b>pullupmsg()</b> concatenates and aligns the first <i>len</i> data bytes of the message pointed to by <i>mp</i>. If <i>len</i> equals -1, all data are concatenated. If <i>len</i> bytes of the same message type cannot be found, <b>pullupmsg()</b> fails and returns 0.</p>
<b>RETURN VALUES</b>	On success, 1 is returned; on failure, 0 is returned.
<b>CONTEXT</b>	<b>pullupmsg()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1 Using pullupmsg()</b></p> <p>This is a driver write <code>srv(9E)</code> (service) routine for a device that does not support scatter/gather DMA. For all <code>M_DATA</code> messages, the data will be transferred to the device with DMA. First, try to pull up the message into one message block with the <b>pullupmsg()</b> function (line 12). If successful, the transfer can be accomplished in one DMA job. Otherwise, it must be done one message block at a time (lines 19–22). After the data has been transferred to the device, free the message and continue processing messages on the queue.</p> <pre> 1 xxxwsrv(q) 2   queue_t *q; 3   { 4     mblk_t *mp; 5     mblk_t *tmp; 6     caddr_t dma_addr; 7     ssize_t dma_len; 8 9     while ((mp = getq(q)) != NULL) { 10      switch (mp-&gt;b_datap-&gt;db_type) { 11        case M_DATA: 12          if (pullupmsg(mp, -1)) { 13            dma_addr = vtop(mp-&gt;b_rptr); 14            dma_len = mp-&gt;b_wptr - mp-&gt;b_rptr; 15            xxx_do_dma(dma_addr, dma_len);</pre>

```

16             freemsg(mp);
17             break;
18         }
19         for (tmp = mp; tmp; tmp = tmp->b_cont) {
20             dma_addr = vtop(tmp->b_rptr);
21             dma_len = tmp->b_wptr - tmp->b_rptr;
22             xxx_do_dma(dma_addr, dma_len);
23         }
24         freemsg(mp);
25         break;
26     . . .
27     }
28 }

```

**SEE ALSO** [srv\(9E\)](#), [alloca\(9F\)](#), [msgpullup\(9F\)](#), [msgb\(9S\)](#)

*Writing Device Drivers*

*STREAMS Programming Guide*

**NOTES** [pullupmsg\(\)](#) is not included in the DKI and will be removed from the system in a future release. Device driver writers are strongly encouraged to use [msgpullup\(9F\)](#) instead of [pullupmsg\(\)](#).

<b>NAME</b>	put – call a STREAMS put procedure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void put(queue_t *q, mblk_t *mp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>q</b>        Pointer to a STREAMS queue.</p> <p><b>mp</b>        Pointer to message block being passed into queue.</p>
<b>DESCRIPTION</b>	<b>put()</b> calls the put procedure ( <b>put</b> (9E) entry point) for the STREAMS queue specified by <i>q</i> , passing it the message block referred to by <i>mp</i> . It is typically used by a driver or module to call its own put procedure.
<b>CONTEXT</b>	<p><b>put()</b> can be called from a STREAMS module or driver put or service routine, or from an associated interrupt handler, timeout, bufcall, or esballoc call-back. In the latter cases the calling code must guarantee the validity of the <i>q</i> argument.</p> <p>Since <b>put()</b> may cause re-entry of the module (as it is intended to do), mutexes or other locks should not be held across calls to it, due to the risk of single-party deadlock ( <b>put</b>(9E), <b>putnext</b>(9F), <b>putctl</b>(9F), <b>qreply</b>(9F).) This function is provided as a DDI/DKI conforming replacement for a direct call to a put procedure.</p>
<b>SEE ALSO</b>	<p><b>put</b>(9E), <b>freezestr</b>(9F), <b>putctl</b>(9F), <b>putctl1</b>(9F), <b>putnext</b>(9F), <b>putnextctl</b>(9F), <b>putnextctl1</b>(9F), <b>qreply</b>(9F)</p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>
<b>NOTES</b>	<p>The caller cannot have the stream frozen when calling this function. See <b>freezestr</b>(9F).</p> <p>DDI/DKI conforming modules and drivers are no longer permitted to call put procedures directly, but must call through the appropriate STREAMS utility function, for example, <b>put</b>(9E), <b>putnext</b>(9F), <b>putctl</b>(9F), and <b>qreply</b>(9F). This function is provided as a DDI/DKI conforming replacement for a direct call to a put procedure.</p>

<b>NAME</b>	putbq – place a message at the head of a queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  int putbq(queue_t *q, mblk_t *bp);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue.  <b>bp</b> Pointer to the message block.
<b>DESCRIPTION</b>	<b>putbq()</b> places a message at the beginning of the appropriate section of the message queue. There are always sections for high priority and ordinary messages. If other priority bands are used, each will have its own section of the queue, in priority band order, after high priority messages and before ordinary messages. <b>putbq()</b> can be used for ordinary, priority band, and high priority messages. However, unless precautions are taken, using <b>putbq()</b> with a high priority message is likely to lead to an infinite loop of putting the message back on the queue, being rescheduled, pulling it off, and putting it back on.  This function is usually called when <b>bcanput(9F)</b> or <b>canput(9F)</b> determines that the message cannot be passed on to the next stream component. The flow control parameters are updated to reflect the change in the queue's status. If <b>QNOENB</b> is not set, the service routine is enabled.
<b>RETURN VALUES</b>	<b>putbq()</b> returns 1 upon success and 0 upon failure.
<b>CONTEXT</b>	<b>putbq()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See the <b>bufcall(9F)</b> function page for an example of <b>putbq()</b> .
<b>SEE ALSO</b>	<b>bcanput(9F)</b> , <b>bufcall(9F)</b> , <b>canput(9F)</b> , <b>getq(9F)</b> , <b>putq(9F)</b>  <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	putctl1 – send a control message with a one-byte parameter to a queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  int putctl1(queue_t *q, int type, int p);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Queue to which the message is to be sent. <b>type</b> Type of message. <b>p</b> One-byte parameter.
<b>DESCRIPTION</b>	<b>putctl1()</b> , like <b>putctl(9F)</b> , tests the <i>type</i> argument to make sure a data type has not been specified, and attempts to allocate a message block. The <i>p</i> parameter can be used, for example, to specify how long the delay will be when an M_DELAY message is being sent. <b>putctl1()</b> fails if <i>type</i> is M_DATA, M_PROTO, or M_PCPROTO, or if a message block cannot be allocated. If successful, <b>putctl1()</b> calls the <b>put(9E)</b> routine of the queue pointed to by <i>q</i> with the newly allocated and initialized message.
<b>RETURN VALUES</b>	On success, 1 is returned. 0 is returned if <i>type</i> is a data type, or if a message block cannot be allocated.
<b>CONTEXT</b>	<b>putctl1()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See the <b>putctl(9F)</b> function page for an example of <b>putctl1()</b> .
<b>SEE ALSO</b>	<b>put(9E)</b> , <b>allocb(9F)</b> , <b>datamsg(9F)</b> , <b>putctl(9F)</b> , <b>putnextctl1(9F)</b>  <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	putctl – send a control message to a queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  int putctl(queue_t *q, int type);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Queue to which the message is to be sent.  <b>type</b> Message type (must be control, not data type).
<b>DESCRIPTION</b>	<b>putctl()</b> tests the <i>type</i> argument to make sure a data type has not been specified, and then attempts to allocate a message block. <b>putctl()</b> fails if <i>type</i> is M_DATA, M_PROTO, or M_PCPROTO, or if a message block cannot be allocated. If successful, <b>putctl()</b> calls the <b>put(9E)</b> routine of the queue pointed to by <i>q</i> with the newly allocated and initialized messages.
<b>RETURN VALUES</b>	On success, 1 is returned. If <i>type</i> is a data type, or if a message block cannot be allocated, 0 is returned.
<b>CONTEXT</b>	<b>putctl()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<b>CODE EXAMPLE 1 Using putctl()</b>  The <b>send_ctl()</b> routine is used to pass control messages downstream. M_BREAK messages are handled with <b>putctl()</b> (line 11). <b>putctl1(9F)</b> (line 16) is used for M_DELAY messages, so that <i>parm</i> can be used to specify the length of the delay. In either case, if a message block cannot be allocated a variable recording the number of allocation failures is incremented (lines 12, 17). If an invalid message type is detected, <b>cmn_err(9F)</b> panics the system (line 21).  <pre> 1 void 2 send_ctl(wrq, type, parm) 3     queue_t *wrq; 4     uchar_t type; 5     uchar_t parm; 6 { 7     extern int num_alloc_fail; 8 9     switch (type) { 10        case M_BREAK: 11            if (!putctl(wrq-&gt;q_next, M_BREAK)) 12                num_alloc_fail++; 13            break; 14 15        case M_DELAY: </pre>

```
16         if (!putctl1(wrq->q_next, M_DELAY, parm))
17             num_alloc_fail++;
18         break;
19
20     default:
21         cmn_err(CE_PANIC, "send_ctl: bad message type passed");
22         break;
23     }
24 }
```

**SEE ALSO** [put\(9E\)](#), [cmn\\_err\(9F\)](#), [datamsg\(9F\)](#), [putctl1\(9F\)](#), [putnextctl1\(9F\)](#)

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	putnext – send a message to the next queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void putnext(queue_t *q, mblk_t *mp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>q</b>        Pointer to the queue from which the message <i>mp</i> will be sent.</p> <p><b>mp</b>        Message to be passed.</p>
<b>DESCRIPTION</b>	<b>putnext()</b> is used to pass a message to the <b>put(9E)</b> routine of the next queue in the stream.
<b>RETURN VALUES</b>	None.
<b>CONTEXT</b>	<b>putnext()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See <b>allocb(9F)</b> for an example of using <b>putnext()</b> .
<b>SEE ALSO</b>	<p><b>put(9E)</b>, <b>allocb(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	putnextctl1 – send a control message with a one-byte parameter to a queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  int putnextctl1(queue_t *q, int type, int p);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>q</b> Queue to which the message is to be sent.</p> <p><b>type</b> Type of message.</p> <p><b>p</b> One-byte parameter.</p>
<b>DESCRIPTION</b>	<p><b>putnextctl1()</b>, like <b>putctl1(9F)</b>, tests the <i>type</i> argument to make sure a data type has not been specified, and attempts to allocate a message block. The <i>p</i> parameter can be used, for example, to specify how long the delay will be when an M_DELAY message is being sent. <b>putnextctl1()</b> fails if <i>type</i> is M_DATA, M_PROTO, or M_PCPROTO, or if a message block cannot be allocated. If successful, <b>putnextctl1()</b> calls the <b>put(9E)</b> routine of the queue pointed to by <i>q</i> with the newly allocated and initialized message.</p> <p>A call to <b>putnextctl1(q, type, p)</b> is an atomic equivalent of <b>putctl1(q-&gt;q_next, type, p)</b>. The STREAMS framework provides whatever mutual exclusion is necessary to insure that dereferencing <i>q</i> through its <i>q_next</i> field and then invoking <b>putctl1(9F)</b> proceeds without interference from other threads.</p> <p><b>putnextctl1()</b> should always be used in preference to <b>putctl1(9F)</b>.</p>
<b>RETURN VALUES</b>	On success, 1 is returned. 0 is returned if <i>type</i> is a data type, or if a message block cannot be allocated.
<b>CONTEXT</b>	<b>putnextctl1()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See the <b>putnextctl1(9F)</b> function page for an example of <b>putnextctl1()</b> .
<b>SEE ALSO</b>	<p><b>put(9E)</b>, <b>allocb(9F)</b>, <b>datamsg(9F)</b>, <b>putctl1(9F)</b>, <b>putnextctl1(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	putnextctl – send a control message to a queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  int putnextctl(queue_t *q, int type);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>q</b> Queue to which the message is to be sent.</p> <p><b>type</b> Message type (must be control, not data type).</p>
<b>DESCRIPTION</b>	<p><b>putnextctl()</b> tests the <i>type</i> argument to make sure a data type has not been specified, and then attempts to allocate a message block. <b>putnextctl()</b> fails if <i>type</i> is M_DATA, M_PROTO, or M_PCPROTO, or if a message block cannot be allocated. If successful, <b>putnextctl()</b> calls the <b>put(9E)</b> routine of the queue pointed to by <i>q</i> with the newly allocated and initialized messages.</p> <p>A call to <b>putnextctl(q, type)</b> is an atomic equivalent of <b>putctl(q-&gt;q_next, type)</b>. The STREAMS framework provides whatever mutual exclusion is necessary to insure that dereferencing <i>q</i> through its <i>q_next</i> field and then invoking <b>putctl(9F)</b> proceeds without interference from other threads.</p> <p><b>putnextctl()</b> should always be used in preference to <b>putctl(9F)</b></p>
<b>RETURN VALUES</b>	On success, 1 is returned. If <i>type</i> is a data type, or if a message block cannot be allocated, 0 is returned.
<b>CONTEXT</b>	<b>putnextctl()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>CODE EXAMPLE 1</b></p> <p>The <b>send_ctl</b> routine is used to pass control messages downstream. M_BREAK messages are handled with <b>putnextctl()</b> (line 8). <b>putnextctl(9F)</b> (line 13) is used for M_DELAY messages, so that <i>parm</i> can be used to specify the length of the delay. In either case, if a message block cannot be allocated a variable recording the number of allocation failures is incremented (lines 9, 14). If an invalid message type is detected, <b>cmn_err(9F)</b> panics the system (line 18).</p> <pre>1 void 2 send_ctl(queue_t *wrq, uchar_t type, uchar_t parm) 3 { 4     extern int num_alloc_fail; 5 }</pre>

```
6         switch (type) {
7         case M_BREAK:
8             if (!putnextctl(wrq, M_BREAK))
9                 num_alloc_fail++;
10            break;
11
12         case M_DELAY:
13             if (!putnextctl1(wrq, M_DELAY, parm))
14                 num_alloc_fail++;
15            break;
16
17         default:
18             cmn_err(CE_PANIC, "send_ctl: bad message type passed");
19            break;
20        }
21    }
```

**SEE ALSO** [put\(9E\)](#), [cmn\\_err\(9F\)](#), [datamsg\(9F\)](#), [putctl\(9F\)](#), [putnextctl1\(9F\)](#)

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	putq – put a message on a queue
<b>SYNOPSIS</b>	#include <sys/stream.h>  int <b>putq</b> (queue_t *q, mblk_t *bp);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue to which the message is to be added. <b>bp</b> Message to be put on the queue.
<b>DESCRIPTION</b>	<b>putq()</b> is used to put messages on a driver's queue after the module's put routine has finished processing the message. The message is placed after any other messages of the same priority, and flow control parameters are updated. If QNOENB is not set, the service routine is enabled. If no other processing is done, <b>putq()</b> can be used as the module's put routine.
<b>RETURN VALUES</b>	<b>putq()</b> returns 1 on success and 0 on failure.
<b>CONTEXT</b>	<b>putq()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See the <b>datamsq(9F)</b> function page for an example of <b>putq()</b> .
<b>SEE ALSO</b>	<b>datamsq(9F)</b> , <b>putbq(9F)</b> , <b>qenable(9F)</b> , <b>rmvq(9F)</b>  <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	qbufcall – call a function when a buffer becomes available
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;</pre> <p>bufcall_id_t <b>qbufcall</b>(queue_t *q, size_t size, uint_t pri, void(*func)(void *arg), void *arg);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>q</b> Pointer to STREAMS queue structure.</p> <p><b>size</b> Number of bytes required for the buffer.</p> <p><b>pri</b> Priority of the <b>allocb(9F)</b> allocation request (not used).</p> <p><b>func</b> Function or driver routine to be called when a buffer becomes available.</p> <p><b>arg</b> Argument to the function to be called when a buffer becomes available.</p>
<b>DESCRIPTION</b>	<p><b>qbufcall()</b> serves as a <b>qtimeout(9F)</b> call of indeterminate length. When a buffer allocation request fails, <b>qbufcall()</b> can be used to schedule the routine <i>func</i> to be called with the argument <i>arg</i> when a buffer becomes available. <i>func</i> may call <b>allocb()</b> or it may do something else.</p> <p>The <b>qbufcall()</b> function is tailored to be used with the enhanced STREAMS framework interface, which is based on the concept of perimeters. (See <b>mt-streams(9F)</b>.) <b>qbufcall()</b> schedules the specified function to execute after entering the perimeters associated with the queue passed in as the first parameter to <b>qbufcall()</b>. All outstanding bufcalls should be cancelled before the close of a driver or module returns.</p> <p><b>qprocson(9F)</b> must be called before calling either <b>qbufcall()</b> or <b>qtimeout(9F)</b>.</p>
<b>RETURN VALUES</b>	If successful, <b>qbufcall()</b> returns a <code>qbufcall</code> ID that can be used in a call to <b>qunbufcall(9F)</b> to cancel the request. If the <b>qbufcall()</b> scheduling fails, <i>func</i> is never called and 0 is returned.
<b>CONTEXT</b>	<b>qbufcall()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>allocb(9F)</b> , <b>mt-streams(9F)</b> , <b>qprocson(9F)</b> , <b>qtimeout(9F)</b> , <b>qunbufcall(9F)</b> , <b>quntimeout(9F)</b>

*Writing Device Drivers*

*STREAMS Programming Guide*

**WARNINGS**

Even when *func* is called by **qbufcall()**, **allocb(9F)** can fail if another module or driver had allocated the memory before *func* was able to call **allocb(9F)**.

<b>NAME</b>	qenable – enable a queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void qenable(queue_t *q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue to be enabled.
<b>DESCRIPTION</b>	<b>qenable()</b> adds the queue pointed to by <i>q</i> to the list of queues whose service routines are ready to be called by the STREAMS scheduler.
<b>CONTEXT</b>	<b>qenable()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	See the <b>dupb(9F)</b> function page for an example of the <b>qenable()</b> .
<b>SEE ALSO</b>	<b>dupb(9F)</b> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	qprocson, qprocsoff – enable, disable put and service routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void qprocson(queue_t * q); void qprocsoff(queue_t * q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the RD side of a STREAMS queue pair.
<b>DESCRIPTION</b>	<p><b>qprocson()</b> enables the put and service routines of the driver or module whose read queue is pointed to by <i>q</i>. Threads cannot enter the module instance through the put and service routines while they are disabled.</p> <p><b>qprocson()</b> must be called by the open routine of a driver or module before returning, and after any initialization necessary for the proper functioning of the put and service routines.</p> <p><b>qprocson()</b> must be called before calling <b>qbufcall(9F)</b>, <b>qtimeout(9F)</b>, <b>qwait(9F)</b>, or <b>qwait_sig(9F)</b>,</p> <p><b>qprocsoff()</b> must be called by the close routine of a driver or module before returning, and before deallocating any resources necessary for the proper functioning of the put and service routines. It also removes the queue's service routines from the service queue, and blocks until any pending service processing completes.</p> <p>The module or driver instance is guaranteed to be single-threaded before <b>qprocson()</b> is called and after <b>qprocsoff()</b> is called, except for threads executing asynchronous events such as interrupt handlers and callbacks, which must be handled separately.</p>
<b>CONTEXT</b>	These routines can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><b>close(9E)</b>, <b>open(9E)</b>, <b>put(9E)</b>, <b>srv(9E)</b>, <b>qbufcall(9F)</b>, <b>qtimeout(9F)</b>, <b>qwait(9F)</b>, <b>qwait_sig(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>
<b>NOTES</b>	The caller may not have the STREAM frozen during either of these calls.

<b>NAME</b>	qreply – send a message on a stream in the reverse direction
<b>SYNOPSIS</b>	#include <sys/stream.h>  void <b>qreply</b> (queue_t *q, mblk_t *mp);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue.  <b>mp</b> Pointer to the message to be sent in the opposite direction.
<b>DESCRIPTION</b>	<b>qreply()</b> sends messages in the reverse direction of normal flow. That is, <b>qreply(q, mp)</b> is equivalent to <b>putnext(OTHERQ(q), mp)</b> .
<b>CONTEXT</b>	<b>qreply()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>CODE EXAMPLE 1</b> Canonical Flushing Code for STREAMS Drivers.</p> <p>This example depicts the canonical flushing code for STREAMS drivers. Assume that the driver has service procedures so that there may be messages on its queues. See <b>srv(9E)</b>. Its write-side put procedure handles <b>M_FLUSH</b> messages by first checking the <b>FLUSHW</b> bit in the first byte of the message, then the write queue is flushed (line 8) and the <b>FLUSHW</b> bit is turned off (line 9). See <b>put(9E)</b>. If the <b>FLUSHR</b> bit is on, then the read queue is flushed (line 12) and the message is sent back up the read side of the stream with the <b>qreply(9F)</b> function (line 13). If the <b>FLUSHR</b> bit is off, then the message is freed (line 15). See the example for <b>flushq(9F)</b> for the canonical flushing code for modules.</p> <pre> 1 xxxwput(q, mp) 2   queue_t *q; 3   mblk_t *mp; 4   { 5       switch(mp-&gt;b_datap-&gt;db_type) { 6           case M_FLUSH: 7               if (*mp-&gt;b_rptr &amp; FLUSHW) { 8                   flushq(q, FLUSHALL); 9                   *mp-&gt;b_rptr &amp;= ~FLUSHW; 10              } 11              if (*mp-&gt;b_rptr &amp; FLUSHR) { 12                  flushq(RD(q), FLUSHALL); 13                  qreply(q, mp); 14              } else { 15                  freemsg(mp); 16              } 17              break; 18          . . . </pre>

```
18     }  
19 }
```

**SEE ALSO**

**put(9E), srv(9E), flushq(9F), OTHERQ(9F), putnext(9F)**

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	qsize – find the number of messages on a queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  int qsize(queue_t *q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Queue to be evaluated.
<b>DESCRIPTION</b>	<b>qsize()</b> evaluates the queue <i>q</i> and returns the number of messages it contains.
<b>RETURN VALUES</b>	If there are no message on the queue, <b>qsize()</b> returns 0. Otherwise, it returns the integer representing the number of messages on the queue.
<b>CONTEXT</b>	<b>qsize()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	qtimeout – execute a function after a specified length of time
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;</pre> <p>timeout_id_t <b>qtimeout</b>(queue_t *q, void (*func)(void *), void *arg, clock_t ticks);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>q</b> Pointer to STREAMS queue structure.</p> <p><b>func</b> Kernel function to invoke when the time increment expires.</p> <p><b>arg</b> Argument to the function.</p> <p><b>ticks</b> Number of clock ticks to wait before the function is called.</p>
<b>DESCRIPTION</b>	<p>The <b>qtimeout()</b> function schedules the specified function <i>func</i> to be called after a specified time interval. <i>func</i> is called with <i>arg</i> as a parameter. Control is immediately returned to the caller. This is useful when an event is known to occur within a specific time frame, or when you want to wait for I/O processes when an interrupt is not available or might cause problems. The exact time interval over which the timeout takes effect cannot be guaranteed, but the value given is a close approximation.</p> <p>The <b>qtimeout()</b> function is tailored to be used with the enhanced STREAMS framework interface which is based on the concept of perimeters. (See <b>mt-streams(9F)</b>.) <b>qtimeout()</b> schedules the specified function to execute after entering the perimeters associated with the queue passed in as the first parameter to <b>qtimeout()</b>. All outstanding timeouts should be cancelled before a driver closes or module returns.</p> <p><b>qprocson(9F)</b> must be called before calling <b>qtimeout()</b>.</p>
<b>RETURN VALUES</b>	<b>qtimeout()</b> returns an opaque non-zero timeout identifier that can be passed to <b>quntimeout(9F)</b> to cancel the request. Note: No value is returned from the called function.
<b>CONTEXT</b>	<b>qtimeout()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>mt-streams(9F)</b> , <b>qbufcall(9F)</b> , <b>qprocson(9F)</b> , <b>qunbufcall(9F)</b> , <b>quntimeout(9F)</b>
	<i>Writing Device Drivers</i>

*STREAMS Programming Guide*

<b>NAME</b>	qunbufcall – cancel a pending qbufcall request
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void qunbufcall(queue_t *q, bufcall_id_t id);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>q</b>        Pointer to STREAMS queue_t structure.</p> <p><b>id</b>       Identifier returned from qbufcall(9F)</p>
<b>DESCRIPTION</b>	<p><b>qunbufcall()</b> cancels a pending <b>qbufcall()</b> request. The argument <i>id</i> is a non-zero identifier of the request to be cancelled. <i>id</i> is returned from the <b>qbufcall()</b> function used to issue the cancel request.</p> <p>The <b>qunbufcall()</b> function is tailored to be used with the enhanced STREAMS framework interface which is based on the concept of perimeters. (See <b>mt-streams(9F)</b>.) <b>qunbufcall()</b> returns when the bufcall has been cancelled or finished executing. The bufcall will be cancelled even if it is blocked at the perimeters associated with the queue. All outstanding bufcalls should be cancelled before the driver closes or module returns.</p>
<b>CONTEXT</b>	<b>qunbufcall()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><b>mt-streams(9F)</b>, <b>qbufcall(9F)</b>, <b>qtimeout(9F)</b>, <b>quntimeout(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	quntimeout – cancel previous qtimeout function call
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  clock_t quntimeout(queue_t *q, timeout_id_t id);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>q</b>        Pointer to a STREAMS queue structure.</p> <p><b>id</b>        Opaque timeout ID a previous qtimeout(9F) call.</p>
<b>DESCRIPTION</b>	<p><b>quntimeout()</b> cancels a pending qtimeout(9F) request. The <b>quntimeout()</b> function is tailored to be used with the enhanced STREAMS framework interface, which is based on the concept of perimeters. (See <b>mt-streams(9F)</b>.) <b>quntimeout()</b> returns when the timeout has been cancelled or finished executing. The timeout will be cancelled even if it is blocked at the perimeters associated with the queue. <b>quntimeout()</b> should be executed for all outstanding timeouts before a driver or module close returns.</p>
<b>RETURN VALUES</b>	<b>quntimeout()</b> returns -1 if the <b>id</b> is not found. Otherwise, <b>quntimeout()</b> returns a 0 or positive value.
<b>CONTEXT</b>	<b>quntimeout()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<p><b>mt-streams(9F)</b>, <b>qbufcall(9F)</b>, <b>qtimeout(9F)</b>, <b>qunbufcall(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	qwait, qwait_sig – STREAMS wait routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  void qwait(queue_t * q); int qwait_sig(queue_t * q);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>qp</b>                    Pointer to the queue that is being opened or closed.</p>
<b>DESCRIPTION</b>	<p><b>qwait()</b> and <b>qwait_sig()</b> are used to wait for a message to arrive to the <b>put(9E)</b> or <b>srv(9E)</b> procedures. <b>qwait()</b> and <b>qwait_sig()</b> can also be used to wait for <b>qbufcall(9F)</b> or <b>qtimeout(9F)</b> callback procedures to execute. These routines can be used in the <b>open(9E)</b> and <b>close(9E)</b> procedures in a STREAMS driver or module. <b>qwait()</b> and <b>qwait_sig()</b> atomically exit the inner and outer perimeters associated with the queue, and wait for a thread to leave the module's <b>put(9E)</b>, <b>srv(9E)</b>, or <b>qbufcall(9F)</b> / <b>qtimeout(9F)</b> callback procedures. Upon return they re-enter the inner and outer perimeters.</p> <p>This can be viewed as there being an implicit wakeup when a thread leaves a <b>put(9E)</b> or <b>srv(9E)</b> procedure or after a <b>qtimeout(9F)</b> or <b>qbufcall(9F)</b> callback procedure has been run in the same perimeter.</p> <p><b>qprocson(9F)</b> must be called before calling <b>qwait()</b> or <b>qwait_sig()</b>.</p> <p><b>qwait()</b> is not interrupted by a signal, whereas <b>qwait_sig()</b> is interrupted by a signal. <b>qwait_sig()</b> normally returns non-zero, and returns zero when the waiting was interrupted by a signal.</p> <p><b>qwait()</b> and <b>qwait_sig()</b> are similar to <b>cv_wait()</b> and <b>cv_wait_sig()</b> except that the mutex is replaced by the inner and outer perimeters and the signalling is implicit when a thread leaves the inner perimeter. See <b>condvar(9F)</b>.</p>
<b>RETURN VALUES</b>	<p>0            For <b>qwait_sig()</b>, indicates that the condition was not necessarily signaled, and the function returned because a signal was pending.</p>
<b>CONTEXT</b>	These functions can only be called from an <b>open(9E)</b> or <b>close(9E)</b> routine.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b>    Using <b>qwait()</b></p> <p>The <b>open</b> routine sends down a <b>T_INFO_REQ</b> message and waits for the <b>T_INFO_ACK</b>. The arrival of the <b>T_INFO_ACK</b> is recorded by resetting a flag</p>

in the unit structure ( `WAIT_INFO_ACK` ). The example assumes that the module is `D_MTQPAIR` or `D_MTPERMOD` .

```

xxopen(qp, ...)
    queue_t *qp;
{
\011    struct xxdata *xx;
\011    /* Allocate xxdata structure */
    qprocson(qp);
    /* Format T_INFO_ACK in mp */
    putnext(qp, mp);
    xx->xx_flags |= WAIT_INFO_ACK;
    while (xx->xx_flags & WAIT_INFO_ACK)
\011\011        qwait(qp);
\011    return (0);
}

xxrput(qp, mp)
    queue_t *qp;
    mblk_t *mp;
{
    struct xxdata *xx = (struct xxdata *)q->q_ptr;

\011    ...

\011    case T_INFO_ACK:
\011\011        if (xx->xx_flags & WAIT_INFO_ACK) {
\011\011\011            /* Record information from info ack */
\011\011\011            xx->xx_flags &= ~WAIT_INFO_ACK;
\011\011\011            freemsg(mp);
\011\011\011            return;
\011\011        }
}

\011    ...
}

```

**SEE ALSO** `close(9E)` , `open(9E)` , `put(9E)` , `srv(9E)` `condvar(9F)` , `mt-streams(9F)` , `qbufcall(9F)` , `qprocson(9F)` , `qtimeout(9F)`

*STREAMS Programming Guide*

*Writing Device Drivers*

<b>NAME</b>	qwriter – asynchronous STREAMS perimeter upgrade
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;</pre> <pre>void <b>qwriter</b>(queue_t *qp, mblk_t *mp, void (*func, int perimeter);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>qp</b> Pointer to the queue.</p> <p><b>mp</b> Pointer to a message that will be passed in to the callback function.</p> <p><b>func</b> A function that will be called when exclusive (writer) access has been acquired at the specified perimeter.</p> <p><b>perimeter</b> Either PERIM_INNER or PERIM_OUTER.</p>
<b>DESCRIPTION</b>	<p><b>qwriter()</b> is used to upgrade the access at either the inner or the outer perimeter from shared to exclusive and call the specified callback function when the upgrade has succeeded. See <b>mt-streams(9F)</b>. The callback function is called as:</p> <pre>(*func)(queue_t *qp, mblk_t *mp);</pre> <p><b>qwriter()</b> will acquire exclusive access immediately if possible, in which case the specified callback function will be executed before <b>qwriter()</b> returns. If this is not possible, <b>qwriter()</b> will defer the upgrade until later and return before the callback function has been executed. Modules should not assume that the callback function has been executed when <b>qwriter()</b> returns. One way to avoid dependencies on the execution of the callback function is to immediately return after calling <b>qwriter()</b> and let the callback function finish the processing of the message.</p> <p>When <b>qwriter()</b> defers calling the callback function, the STREAMS framework will prevent other messages from entering the inner perimeter associated with the queue until the upgrade has completed and the callback function has finished executing.</p>
<b>CONTEXT</b>	<b>qwriter()</b> can only be called from an <b>put(9E)</b> or <b>srv(9E)</b> routine, or from a <b>qwriter()</b> , <b>qtimeout(9F)</b> , or <b>qbufcall(9F)</b> callback function.
<b>SEE ALSO</b>	<b>put(9E)</b> , <b>srv(9E)</b> , <b>mt-streams(9F)</b> , <b>qbufcall(9F)</b> , <b>qtimeout(9F)</b>

*STREAMS Programming Guide*

*Writing Device Drivers*

<b>NAME</b>	RD, rd – get pointer to the read queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  queue_t * RD(queue_t * q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 ( DDI/DKI) .
<b>PARAMETERS</b>	<b>q</b> Pointer to the <code>write</code> queue whose <code>read</code> queue is to be returned.
<b>DESCRIPTION</b>	<p>The <b>RD()</b> function accepts a <code>write</code> queue pointer as an argument and returns a pointer to the <code>read</code> queue of the same module.</p> <p><b>CAUTION:</b> Make sure the argument to this function is a pointer to a <code>write</code> queue. <b>RD()</b> will not check for queue type, and a system panic could result if it is not the right type.</p>
<b>RETURN VALUES</b>	The pointer to the <code>read</code> queue.
<b>CONTEXT</b>	<b>RD()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Function page reference</p> <p>See the <b>qreply(9F)</b> function page for an example of <b>RD()</b> .</p>
<b>SEE ALSO</b>	<p><b>qreply(9F)</b> , <b>wr(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	rmalloc – allocate space from a resource map
<b>SYNOPSIS</b>	<pre>#include &lt;sys/map.h&gt; #include &lt;sys/ddi.h&gt;</pre> <p>unsigned long <b>rmalloc</b>(struct map *mp, size_t size);</p>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>mp</b> Resource map from where the resource is drawn.</p> <p><b>size</b> Number of units of the resource.</p>
<b>DESCRIPTION</b>	<p><b>rmalloc()</b> is used by a driver to allocate space from a previously defined and initialized resource map. The map itself is allocated by calling the function <b>rmallocmap(9F)</b>. <b>rmalloc()</b> is one of five functions used for resource map management. The other functions include:</p> <p><b>rmalloc_wait(9F)</b> Allocate space from a resource map, wait if necessary.</p> <p><b>rmfree(9F)</b> Return previously allocated space to a map.</p> <p><b>rmallocmap(9F)</b> Allocate a resource map and initialize it.</p> <p><b>rmfreemap(9F)</b> Deallocate a resource map.</p> <p><b>rmalloc()</b> allocates space from a resource map in terms of arbitrary units. The system maintains the resource map by size and index, computed in units appropriate for the resource. For example, units may be byte addresses, pages of memory, or blocks. The normal return value is an <code>unsigned long</code> set to the value of the index where sufficient free space in the resource was found.</p>
<b>RETURN VALUES</b>	Under normal conditions, <b>rmalloc()</b> returns the base index of the allocated space. Otherwise, <b>rmalloc()</b> returns a 0 if all resource map entries are already allocated.
<b>CONTEXT</b>	<b>rmalloc()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Illustrating the principles of map management</p> <p>The following example is a simple memory map, but it illustrates the principles of map management. A driver allocates and initializes the map by calling both the <b>rmallocmap(9F)</b> and <b>rmfree(9F)</b> functions. <b>rmallocmap(9F)</b> is called to establish the number of slots or entries in the map, and <b>rmfree(9F)</b> to initialize the resource area the map is to manage.</p>

The following example is a fragment from a hypothetical `start` routine and illustrates the following procedures:

- Panics the system if the required amount of memory can not be allocated (lines 11–15).
- Uses `rmallocmap(9F)` to configure the total number of entries in the map, and `rmfree(9F)` to initialize the total resource area.

```

1  #define XX_MAPSIZE 12
2  #define XX_BUFSIZE 2560
3  static struct map *xx_mp;          /* Private buffer space map */
4  ...
5  xxstart()
6  /*
7   * Allocate private buffer.  If insufficient memory,
8   * display message and halt system.
9   */
10 {
11     register caddr_t bp;
12     ...
13     if ((bp = kmem_alloc(XX_BUFSIZE, KM_NOSLEEP) == 0) {
14         cmn_err(CE_PANIC, "xxstart: kmem_alloc failed before %d buffer"
15                "allocation", XX_BUFSIZE);
16     }
17     /*
18      * Initialize the resource map with number
19      * of slots in map.
20      */
21     xx_mp = rmallocmap(XX_MAPSIZE);
22     ...
23     /*
24      * Initialize space management map with total
25      * buffer area it is to manage.
26      */
27     rmfree(xx_mp, XX_BUFSIZE, bp);
28     ...

```

#### EXAMPLE 2 Allocating buffers

The `rmalloc()` function is then used by the driver's `read` or `write` routine to allocate buffers for specific data transfers. The `uiomove(9F)` function is used to move the data between user space and local driver memory. The device then moves data between itself and local driver memory through DMA.

The next example illustrates the following procedures:

- The size of the I/O request is calculated and stored in the `size` variable (line 10).

- Buffers are allocated through the **rmalloc()** function using the *size* value (line 15). If the allocation fails the system will panic.
- The **uiomove(9F)** function is used to move data to the allocated buffer (line 23).
- If the address passed to **uiomove(9F)** is invalid, **rmfree(9F)** is called to release the previously allocated buffer, and an EFAULT error is returned.

```

1  #define XX_BUFSIZE  2560
2  #define XX_MAXSIZE  (XX_BUFSIZE / 4)
3
4  static struct map *xx_mp;          /* Private buffer space map */
5  ...
6  xxread(dev_t dev, uio_t *uiop, cred_t *credp)
7  {
8      register caddr_t addr;
9      register int    size;
10     size = min(COUNT, XX_MAXSIZE); /* Break large I/O request */
11                                     /* into small ones */
12     /*
13      * Get buffer.
14      */
15     if ((addr = (caddr_t)rmalloc(xx_mp, size)) == 0)
16         cmn_err(CE_PANIC, "read: rmalloc failed allocation of size %d",
17                size);
18     /*
19      * Move data to buffer.  If invalid address is found,
20      * return buffer to map and return error code.
21      */
22     if (uiomove(addr, size, UIO_READ, uiop) == -1) {
23         rmfree(xx_mp, size, addr);
24         return(EFAULT);
25     }
26 }
27 }

```

**SEE ALSO**

**kmem\_alloc(9F)**, **rmalloc\_wait(9F)**, **rmap(9F)**, **rmallocmap(9F)**, **rmfree(9F)**, **rmfreemap(9F)**, **uiomove(9F)**

*Writing Device Drivers*

<b>NAME</b>	rmallocmap, rmallocmap_wait, rmfreemap – allocate and free resource maps
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  struct map * rmallocmap(size_t mapsize);  struct map * rmallocmap_wait(size_t mapsize);  void rmfreemap(struct map * mp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 ( DDI/DKI ).
<b>PARAMETERS</b>	<p><b>mapsize</b>            Number of entries for the map.</p> <p><b>mp</b>                    A pointer to the map structure to be deallocated.</p>
<b>DESCRIPTION</b>	<p><b>rmallocmap()</b> dynamically allocates a resource map structure. The argument <i>mapsize</i> defines the total number of entries in the map. In particular, it is the total number of allocations that can be outstanding at any one time.</p> <p><b>rmallocmap()</b> initializes the map but does not associate it with the actual resource. In order to associate the map with the actual resource, a call to <b>rmfree(9F)</b> is used to make the entirety of the actual resource available for allocation, starting from the first index into the resource. Typically, the call to <b>rmallocmap()</b> is followed by a call to <b>rmfree(9F)</b> , passing the address of the map returned from <b>rmallocmap()</b> , the total size of the resource, and the first index into the actual resource.</p> <p>The resource map allocated by <b>rmallocmap()</b> can be used to describe an arbitrary resource in whatever allocation units are appropriate, such as blocks, pages, or data structures. This resource can then be managed by the system by subsequent calls to <b>rmalloc(9F)</b> , <b>rmalloc_wait(9F)</b> , and <b>rmfree(9F)</b> .</p> <p><b>rmallocmap_wait()</b> is similar to <b>rmallocmap()</b> , with the exception that it will wait for space to become available if necessary.</p> <p><b>rmfreemap()</b> deallocates a resource map structure previously allocated by <b>rmallocmap()</b> or <b>rmallocmap_wait()</b> . The argument <i>mp</i> is a pointer to the map structure to be deallocated.</p>
<b>RETURN VALUES</b>	Upon successful completion, <b>rmallocmap()</b> and <b>rmallocmap_wait()</b> return a pointer to the newly allocated map structure. Upon failure, <b>rmallocmap()</b> returns a NULL pointer.

**CONTEXT** | **rmallocmap()** and **rmfreemap()** can be called from user, kernel, or interrupt context.

**rmallocmap\_wait()** can only be called from user or kernel context.

**SEE ALSO** | **rmalloc(9F)** , **rmalloc\_wait(9F)** , **rmfree(9F)**

*Writing Device Drivers*

<b>NAME</b>	rmalloc_wait – allocate space from a resource map, wait if necessary
<b>SYNOPSIS</b>	<pre>#include &lt;sys/map.h&gt; #include &lt;sys/ddi.h&gt;  unsigned long rmalloc_wait(struct map *mp, size_t size);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>mp</b> Pointer to the resource map from which space is to be allocated.</p> <p><b>size</b> Number of units of space to allocate.</p>
<b>DESCRIPTION</b>	<p><b>rmalloc_wait()</b> requests an allocation of space from a resource map. <b>rmalloc_wait()</b> is similar to the <b>rmalloc(9F)</b> function with the exception that it will wait for space to become available if necessary.</p>
<b>RETURN VALUES</b>	<b>rmalloc_wait()</b> returns the base of the allocated space.
<b>CONTEXT</b>	This function can be called from user or interrupt context. However, in most cases <b>rmalloc_wait()</b> should be called from user context only.
<b>SEE ALSO</b>	<b>rmalloc(9F)</b> , <b>rmallocmap(9F)</b> , <b>rmfree(9F)</b> , <b>rmfreemap(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	rmfree – free space back into a resource map
<b>SYNOPSIS</b>	<pre>#include &lt;sys/map.h&gt; #include &lt;sys/ddi.h&gt;  void <b>rmfree</b>(struct map *mp, size_t size, ulong_t index);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>mp</b> Pointer to the map structure.</p> <p><b>size</b> Number of units being freed.</p> <p><b>index</b> Index of the first unit of the allocated resource.</p>
<b>DESCRIPTION</b>	<p><b>rmfree()</b> releases space back into a resource map. It is the opposite of <b>rmalloc(9F)</b>, which allocates space that is controlled by a resource map structure.</p> <p>Drivers may define resource maps for resource allocation, in terms of arbitrary units, using the <b>rmallocmap(9F)</b> function. The system maintains the resource map structure by size and index, computed in units appropriate for the resource. For example, units may be byte addresses, pages of memory, or blocks. <b>rmfree()</b> frees up unallocated space for re-use.</p>
<b>CONTEXT</b>	<b>rmfree()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>rmalloc(9F)</b> , <b>rmalloc_wait(9F)</b> , <b>rmallocmap(9F)</b> , <b>rmfreemap(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	rmvb – remove a message block from a message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  mblk_t *rmvb(mblk_t *mp, mblk_t *bp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>mp</b> Message from which a block is to be removed. <code>mblk_t</code> is an instance of the <code>msgb(9S)</code> structure.</p> <p><b>bp</b> Message block to be removed.</p>
<b>DESCRIPTION</b>	<b>rmvb()</b> removes a message block ( <i>bp</i> ) from a message ( <i>mp</i> ), and returns a pointer to the altered message. The message block is not freed, merely removed from the message. It is the module or driver's responsibility to free the message block.
<b>RETURN VALUES</b>	If successful, a pointer to the message (minus the removed block) is returned. The pointer is <code>NULL</code> if <i>bp</i> was the only block of the message before <b>rmvb()</b> was called. If the designated message block ( <i>bp</i> ) does not exist, <code>-1</code> is returned.
<b>CONTEXT</b>	<b>rmvb()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p>This routine removes all zero-length <code>M_DATA</code> message blocks from the given message. For each message block in the message, save the next message block (line 10). If the current message block is of type <code>M_DATA</code> and has no data in its buffer (line 11), then remove it from the message (line 12) and free it (line 13). In either case, continue with the next message block in the message (line 16).</p> <pre> 1 void 2 xxclean(mp) 3     mblk_t *mp; 4 { 5     mblk_t *tmp; 6     mblk_t *nmp; 7 8     tmp = mp; 9     while (tmp) { 10        nmp = tmp-&gt;b_cont; 11        if ((tmp-&gt;b_datap-&gt;db_type == M_DATA) &amp;&amp; 12            (tmp-&gt;b_rptr == tmp-&gt;b_wptr)) { 13            (void) rmbv(mp, tmp); 14            freeb(tmp); 15        } 16        tmp = nmp; 17    }</pre>

**SEE ALSO** `freeb(9F)`, `msgb(9S)`  
*Writing Device Drivers*  
*STREAMS Programming Guide*

<b>NAME</b>	rmvq – remove a message from a queue
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  void rmvq(queue_t *q, mblk_t *mp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>q</b> Queue containing the message to be removed.</p> <p><b>mp</b> Message to remove.</p>
<b>DESCRIPTION</b>	<p><b>rmvq()</b> removes a message from a queue. A message can be removed from anywhere on a queue. To prevent modules and drivers from having to deal with the internals of message linkage on a queue, either <b>rmvq()</b> or <b>getq(9F)</b> should be used to remove a message from a queue.</p>
<b>CONTEXT</b>	<b>rmvq()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p>This code fragment illustrates how one may flush one type of message from a queue. In this case, only <code>M_PROTO T_DATA_IND</code> messages are flushed. For each message on the queue, if it is an <code>M_PROTO</code> message (line 8) of type <code>T_DATA_IND</code> (line 10), save a pointer to the next message (line 11), remove the <code>T_DATA_IND</code> message (line 12) and free it (line 13). Continue with the next message in the list (line 19).</p> <pre> 1  mblk_t *mp, *nmp; 2  queue_t *q; 3  union T_primitives *tp; 4 5  freezestr(q); 6  mp = q-&gt;q_first; 7  while (mp) { 8    if (mp-&gt;b_datap-&gt;db_type == M_PROTO) { 9      tp = (union T_primitives *)mp-&gt;b_rptr; 10     if (tp-&gt;type == T_DATA_IND) { 11       nmp = mp-&gt;b_next; 12       rmvq(q, mp); 13       freemsg(mp); 14       mp = nmp; 15     } else { 16       mp = mp-&gt;b_next; 17     } 18   } else { 19     mp = mp-&gt;b_next; 20   } 21 } 22 unfreezestr(q);</pre>

**SEE ALSO** `freemsg(9F)`, `freezestr(9F)`, `getq(9F)`, `insq(9F)`, `unfreezestr(9F)`

*Writing Device Drivers*

*STREAMS Programming Guide*

**WARNINGS** Make sure that the message *mp* is linked onto *q* to avoid a possible system panic.

**NOTES** The stream must be frozen using `freezestr(9F)` before calling `rmvq( )`.

<b>NAME</b>	rwlock, rw_init, rw_destroy, rw_enter, rw_exit, rw_tryenter, rw_downgrade, rw_tryupgrade, rw_read_locked – readers/writer lock functions
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ksynch.h&gt;  void rw_init(krwlock_t * rwp, char * name, krw_type_t type, void * arg);  void rw_destroy(krwlock_t * rwp);  void rw_enter(krwlock_t * rwp, krw_t enter_type);  void rw_exit(krwlock_t * rwp);  int rw_tryenter(krwlock_t * rwp, krw_t enter_type);  void rw_downgrade(krwlock_t * rwp);  int rw_tryupgrade(krwlock_t * rwp);  int rw_read_locked(krwlock_t * rwp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>rwp</b> Pointer to a <code>krwlock_t</code> readers/writer lock.</p> <p><b>name</b> Descriptive string. This is obsolete and should be <code>NULL</code>. (Non-null strings are legal, but they're a waste of kernel memory.)</p> <p><b>type</b> Type of readers/writer lock.</p> <p><b>arg</b> Type-specific argument for initialization function.</p> <p><b>enter_type</b> Indication of whether the lock is to be acquired non-exclusively or exclusively <code>RW_READER</code> or <code>RW_WRITER</code>.</p>
<b>DESCRIPTION</b>	<p>A multiple-readers, single-writer lock is represented by the <code>krwlock_t</code> data type. This type of lock will allow many threads to have simultaneous read-only access to an object. Only one thread may have write access at any one time. An object which is searched more frequently than it is changed is a good candidate for a readers/writer lock.</p> <p>Readers/writer locks are slightly more expensive than mutex locks, and the advantage of multiple read access may not occur if the lock will only be held for a short time.</p>

**rw\_init()** initializes a readers/writer lock. It is an error to initialize a lock more than once. The *type* argument should be set to `RW_DRIVER`. If the lock is used by the interrupt handler, the type-specific argument, *arg*, should be the `ddi_iblock_cookie` returned from `ddi_get_iblock_cookie(9F)` or `ddi_get_soft_iblock_cookie(9F)`. If the lock is not used by any interrupt handler, the argument should be `NULL`.

**rw\_destroy()** releases any resources that might have been allocated by **rw\_init()**. It should be called before freeing the memory containing the lock.

**rw\_enter()** acquires the lock, and blocks if necessary. If *enter\_type* is `RW_READER`, the caller blocks if there is a writer or a thread attempting to enter for writing. If *enter\_type* is `RW_WRITER`, the caller blocks if any thread holds the lock.

NOTE: It is a programming error for any thread to acquire an rwlock it already holds, even as a reader. Doing so can deadlock the system: if thread R acquires the lock as a reader, then thread W tries to acquire the lock as a writer, W will set write-wanted and block. When R tries to get its second read hold on the lock, it will honor the write-wanted bit and block waiting for W; but W cannot run until R drops the lock. Thus threads R and W deadlock.

**rw\_exit()** releases the lock and may wake up one or more threads waiting on the lock.

**rw\_tryenter()** attempts to enter the lock, like **rw\_enter()**, but never blocks. It returns a non-zero value if the lock was successfully entered, and zero otherwise.

A thread which holds the lock exclusively (entered with `RW_WRITER`), may call **rw\_downgrade()** to convert to holding the lock non-exclusively (as if entered with `RW_READER`). One or more waiting readers may be unblocked.

**rw\_tryupgrade()** can be called by a thread which holds the lock for reading to attempt to convert to holding it for writing. This upgrade can only succeed if no other thread is holding the lock and no other thread is blocked waiting to acquire the lock for writing.

**rw\_read\_locked()** returns non-zero if the calling thread holds the lock for read, and zero if the caller holds the lock for write. The caller must hold the lock. The system may panic if **rw\_read\_locked()** is called for a lock that isn't held by the caller.

## RETURN VALUES

0	<b>rw_tryenter()</b> could not obtain the lock without blocking.
0	<b>rw_tryupgrade()</b> was unable to perform the upgrade because of other threads holding or waiting to hold the lock.

0           **rw\_read\_locked()** returns 0 if the lock is held by the caller for write.

no n-zero    from **rw\_read\_locked()** if the lock is held by the caller for read.

non-zero     successful return from **rw\_tryenter()** or **rw\_tryupgrade()** .

**CONTEXT**    These functions can be called from user or interrupt context, except for **rw\_init()** and **rw\_destroy()** , which can be called from user context only.

**SEE ALSO**    **condvar(9F)** , **ddi\_add\_intr(9F)** , **ddi\_get\_iblock\_cookie(9F)** , **ddi\_get\_soft\_iblock\_cookie(9F)** , **mutex(9F)** , **semaphore(9F)**

*Writing Device Drivers*

**NOTES**      Compiling with **\_LOCKTEST** or **\_MPSTATS** defined no longer has any effect. To gather lock statistics, see **lockstat(1M)** .

<b>NAME</b>	SAMESTR, samestr – test if next queue is in the same stream
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  int SAMESTR(queue_t * q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the queue.
<b>DESCRIPTION</b>	The SAMESTR() function is used to see if the next queue in a stream (if it exists) is the same type as the current queue (that is, both are read queues or both are write queues). This function accounts for the twisted queue connections that occur in a STREAMS pipe and should be used in preference to direct examination of the <code>q_next</code> field of <code>queue(9S)</code> to see if the stream continues beyond <code>q</code> .
<b>RETURN VALUES</b>	SAMESTR() returns 1 if the next queue is the same type as the current queue. It returns 0 if the next queue does not exist or if it is not the same type.
<b>CONTEXT</b>	SAMESTR() can be called from user or interrupt context.
<b>SEE ALSO</b>	OTHERQ(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	scsi_abort – abort a SCSI command
<b>SYNOPSIS</b>	#include <sys/scsi/scsi.h>  int <b>scsi_abort</b> (struct scsi_address * <i>ap</i> , struct scsi_pkt * <i>pkt</i> );
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b><i>ap</i></b> Pointer to a <i>scsi_address</i> structure. <b><i>pkt</i></b> Pointer to a <b>scsi_pkt(9S)</b> structure.
<b>DESCRIPTION</b>	<b>scsi_abort()</b> terminates a command that has been transported to the host adapter driver. A NULL <i>pkt</i> causes all outstanding packets to be aborted. On a successful abort, the <i>pkt_reason</i> is set to CMD_ABORTED and <i>pkt_statistics</i> is OR'ed with STAT_ABORTED.
<b>RETURN VALUES</b>	<b>scsi_abort()</b> returns: 1        on success. 0        on failure.
<b>CONTEXT</b>	<b>scsi_abort()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<b>CODE EXAMPLE 1</b> Terminating a command.  <pre> if (scsi_abort(&amp;devp-&gt;sd_address, pkt) == 0) {     (void) scsi_reset(&amp;devp-&gt;sd_address, RESET_ALL); } </pre>
<b>SEE ALSO</b>	<b>tran_abort(9E)</b> , <b>scsi_reset(9F)</b> , <b>scsi_pkt(9S)</b>  <i>Writing Device Drivers</i>

<b>NAME</b>	scsi_alloc_consistent_buf – allocate an I/O buffer for SCSI DMA
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;</pre> <pre>struct buf *<b>scsi_alloc_consistent_buf</b>(structscsi_address*<i>ap</i>, struct buf *<i>bp</i>, size_t <i>datalen</i>, uint_t <i>bflags</i>, int (*<i>callback</i>, caddr_t),caddr_t <i>arg</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>ap</i></b> Pointer to the <b>scsi_address</b>(9S) structure.</p> <p><b><i>bp</i></b> Pointer to the <b>buf</b>(9S) structure.</p> <p><b><i>datalen</i></b> Number of bytes for the data buffer.</p> <p><b><i>bflags</i></b> Flags setting for the allocated buffer header.</p> <p><b><i>callback</i></b> A pointer to a callback function, <b>NULL_FUNC</b> or <b>SLEEP_FUNC</b>.</p> <p><b><i>arg</i></b> The callback function argument.</p>
<b>DESCRIPTION</b>	<p><b>scsi_alloc_consistent_buf()</b> allocates a buffer header and the associated data buffer for direct memory access (DMA) transfer. This buffer is allocated from the <b>iobp</b> space, which is considered consistent memory. For more details, see <b>ddi_dma_mem_alloc</b>(9F) and <b>ddi_dma_sync</b>(9F).</p> <p>For buffers allocated via <b>scsi_alloc_consistent_buf()</b>, and marked with the <b>PKT_CONSISTENT</b> flag via <b>scsi_init_pkt</b>(9F), the HBA driver must ensure that the data transfer for the command is correctly synchronized before the target driver's command completion callback is performed.</p> <p>If <b>bp</b> is <b>NULL</b>, a new buffer header will be allocated using <b>getrbuf</b>(9F). In addition, if <b>datalen</b> is non-zero, a new buffer will be allocated using <b>ddi_dma_mem_alloc</b>(9F).</p> <p><b>callback</b> indicates what the allocator routines should do when direct memory access (DMA) resources are not available; the valid values are:</p> <p><b>NULL_FUNC</b> Do not wait for resources. Return a <b>NULL</b> pointer.</p> <p><b>SLEEP_FUNC</b> Wait indefinitely for resources.</p> <p><b>Other Values</b> <b>callback</b> points to a function that is called when resources may become available. <b>callback</b> must return either 0 (indicating that it attempted to allocate resources but failed</p>

to do so), in which case it is put back on a list to be called again later, or 1 indicating either success in allocating resources or indicating that it no longer cares for a retry. The last argument *arg* is supplied to the *callback* function when it is invoked.

**RETURN VALUES**

**scsi\_alloc\_consistent\_buf()** returns a pointer to a **buf(9S)** structure on success. It returns **NULL** if resources are not available even if *waitfunc* was not **SLEEP\_FUNC**.

**CONTEXT**

If *callback* is **SLEEP\_FUNC**, then this routine may be called only from user-level code. Otherwise, it may be called from either user or interrupt level. The *callback* function may not block or call routines that block.

**EXAMPLES**

**EXAMPLE 1** Allocate a request sense packet with consistent DMA resources attached.

```
bp = scsi_alloc_consistent_buf(&devp->sd_address, NULL,
    SENSE_LENGTH, B_READ, SLEEP_FUNC, NULL);
rqpkt = scsi_init_pkt(&devp->sd_address,
    NULL, bp, CDB_GROUP0, 1, 0,
    PKT_CONSISTENT, SLEEP_FUNC, NULL);
```

**EXAMPLE 2** Allocate an inquiry packet with consistent DMA resources attached.

```
bp = scsi_alloc_consistent_buf(&devp->sd_address, NULL,
    SUN_INQSIZE, B_READ, canwait, NULL);
if (bp) {
    pkt = scsi_init_pkt(devp->sd_address, NULL, bp,
        CDB_GROUP0, 1, PP_LEN, PKT_CONSISTENT,
        canwait, NULL);
}
```

**SEE ALSO**

**ddi\_dma\_mem\_alloc(9F)**, **ddi\_dma\_sync(9F)**, **getrbuf(9F)**,  
**scsi\_destroy\_pkt(9F)**, **scsi\_init\_pkt(9F)**,  
**scsi\_free\_consistent\_buf(9F)**, **buf(9S)**, **scsi\_address(9S)**

*Writing Device Drivers*

<b>NAME</b>	scsi_cname, scsi_dname, scsi_mname, scsi_rname, scsi_sname – decode a SCSI name
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  char * <b>scsi_cname</b>(uchar_t <i>cmd</i>, char ** <i>cmdvec</i>); char * <b>scsi_dname</b>(int <i>dtype</i>); char * <b>scsi_mname</b>(uchar_t <i>msg</i>); char * <b>scsi_rname</b>(uchar_t <i>reason</i>); char * <b>scsi_sname</b>(uchar_t <i>sense_key</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>cmd</i></b>            A SCSI command value.</p> <p><b><i>cmdvec</i></b>         Pointer to an array of command strings.</p> <p><b><i>dtype</i></b>           Device type.</p> <p><b><i>msg</i></b>             A message value.</p> <p><b><i>reason</i></b>         A packet reason value.</p> <p><b><i>sense_key</i></b>      A SCSI sense key value.</p>
<b>DESCRIPTION</b>	<p><b>scsi_cname()</b> decodes SCSI commands. <i>cmdvec</i> is a pointer to an array of strings. The first byte of the string is the command value, and the remainder is the name of the command.</p> <p><b>scsi_dname()</b> decodes the peripheral device type (for example, direct access or sequential access) in the inquiry data.</p> <p><b>scsi_mname()</b> decodes SCSI messages.</p> <p><b>scsi_rname()</b> decodes packet completion reasons.</p> <p><b>scsi_sname()</b> decodes SCSI sense keys.</p>
<b>RETURN VALUES</b>	These functions return a pointer to a string. If an argument is invalid, they return a string to that effect.
<b>CONTEXT</b>	These functions can be called from user or interrupt context.

**EXAMPLES****EXAMPLE 1** Decoding SCSI tape commands.**scsi\_cname()** decodes SCSI tape commands as follows:

```
static char *st_cmds[] = {
    "\\000test unit ready",
    "\\001rewind",\011
    "\\003request sense",
    "\\010read",
    "\\012write",
    "\\020write file mark",
    "\\021space",
    "\\022inquiry",
    "\\025mode select",
    "\\031erase tape",
    "\\032mode sense",
    "\\033load tape",
    NULL
};
..
cmn_err(CE_CONT, "st: cmd=%s", scsi_cname(cmd, st_cmds));\011
..
```

**SEE ALSO***Writing Device Drivers*

<b>NAME</b>	scsi_destroy_pkt – free an allocated SCSI packet and its DMA resource
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  void <b>scsi_destroy_pkt</b>(struct scsi_pkt *pkt);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>pkt</b> Pointer to a <b>scsi_pkt</b> (9S) structure.
<b>DESCRIPTION</b>	<b>scsi_destroy_pkt()</b> releases all necessary resources, typically at the end of an I/O transfer. The data is synchronized to memory, then the DMA resources are deallocated and <b>pkt</b> is freed.
<b>CONTEXT</b>	<b>scsi_destroy_pkt()</b> may be called from user or interrupt context.
<b>EXAMPLES</b>	<b>CODE EXAMPLE 1</b> Releasing resources. <pre>scsi_destroy_pkt(un-&gt;un_rqs);</pre>
<b>SEE ALSO</b>	<b>tran_destroy_pkt</b> (9E), <b>scsi_init_pkt</b> (9F), <b>scsi_pkt</b> (9S) <i>Writing Device Drivers</i>

<b>NAME</b>	scsi_dmaget, scsi_dmafree – SCSI dma utility routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  struct scsi_pkt * <b>scsi_dmaget</b>(struct scsi_pkt * <i>pkt</i>, opaque_t <i>dmatoken</i>, int (* <i>callback</i> )(void));  void <b>scsi_dmafree</b>(struct scsi_pkt * <i>pkt</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>pkt</i></b>                    A pointer to a <b>scsi_pkt</b>(9S) structure.</p> <p><b><i>dmatoken</i></b>                Pointer to an implementation dependent object</p> <p><b><i>callback</i></b>                Pointer to a callback function, or NULL_FUNC or SLEEP_FUNC.</p>
<b>DESCRIPTION</b>	<p><b>scsi_dmaget()</b> allocates DMA resources for an already allocated SCSI packet. <i>pkt</i> is a pointer to the previously allocated SCSI packet (see <b>scsi_pktalloc</b>(9F) ).</p> <p><i>dmatoken</i> is a pointer to an implementation dependent object which defines the length, direction, and address of the data transfer associated with this SCSI packet (command). The <i>dmatoken</i> must be a pointer to a <b>buf</b>(9S) structure. If <i>dmatoken</i> is NULL, no resources are allocated.</p> <p><i>callback</i> indicates what <b>scsi_dmaget()</b> should do when resources are not available:</p> <p>NULL_FUNC                Do not wait for resources. Return a NULL pointer.</p> <p>SLEEP_FUNC               Wait indefinitely for resources.</p> <p><b>Other Values</b>           <i>callback</i> points to a function which is called when resources may have become available. <i>callback</i> must return either 0 (indicating that it attempted to allocate resources but failed to do so again), in which case it is put back on a list to be called again later, or 1 indicating either success in allocating resources or indicating that it no longer cares for a retry.</p> <p><b>scsi_dmafree()</b> frees the DMA resources associated with the SCSI packet. The packet itself remains allocated.</p>
<b>RETURN VALUES</b>	<b>scsi_dmaget()</b> returns a pointer to a <b>scsi_pkt</b> on success. It returns NULL if resources are not available.

**CONTEXT** | If *callback* is `SLEEP_FUNC`, then this routine may only be called from user-level code. Otherwise, it may be called from either user or interrupt level. The *callback* function may not block or call routines that block.

**scsi\_dmafree()** can be called from user or interrupt context.

**SEE ALSO** | `scsi_pktalloc(9F)`, `scsi_pktfree(9F)`, `scsi_realloc(9F)`,  
`scsi_resfree(9F)`, `buf(9S)`, `scsi_pkt(9S)`

*Writing Device Drivers*

<b>NAME</b>	scsi_errmsg – display a SCSI request sense message																
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  void <b>scsi_errmsg</b>(struct scsi_device *<i>devp</i>, struct scsi_pkt *<i>pktp</i>, char *<i>drv_name</i>, int <i>severity</i>, daddr_t <i>blkno</i>, daddr_t <i>err_blkno</i>, struct scsi_key_strings *<i>cmdlist</i>, struct scsi_extended_sense *<i>sensep</i>);</pre>																
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).																
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><b><i>devp</i></b></td> <td>Pointer to the <b>scsi_device</b>(9S) structure.</td> </tr> <tr> <td><b><i>pktp</i></b></td> <td>Pointer to a <b>scsi_pkt</b>(9S) structure.</td> </tr> <tr> <td><b><i>drv_name</i></b></td> <td>String used by <b>scsi_log</b>(9F).</td> </tr> <tr> <td><b><i>severity</i></b></td> <td>Error severity level, maps to severity strings below.</td> </tr> <tr> <td><b><i>blkno</i></b></td> <td>Requested block number.</td> </tr> <tr> <td><b><i>err_blkno</i></b></td> <td>Error block number.</td> </tr> <tr> <td><b><i>cmdlist</i></b></td> <td>An array of SCSI command description strings.</td> </tr> <tr> <td><b><i>sensep</i></b></td> <td>A pointer to a <b>scsi_extended_sense</b>(9S) structure.</td> </tr> </table>	<b><i>devp</i></b>	Pointer to the <b>scsi_device</b> (9S) structure.	<b><i>pktp</i></b>	Pointer to a <b>scsi_pkt</b> (9S) structure.	<b><i>drv_name</i></b>	String used by <b>scsi_log</b> (9F).	<b><i>severity</i></b>	Error severity level, maps to severity strings below.	<b><i>blkno</i></b>	Requested block number.	<b><i>err_blkno</i></b>	Error block number.	<b><i>cmdlist</i></b>	An array of SCSI command description strings.	<b><i>sensep</i></b>	A pointer to a <b>scsi_extended_sense</b> (9S) structure.
<b><i>devp</i></b>	Pointer to the <b>scsi_device</b> (9S) structure.																
<b><i>pktp</i></b>	Pointer to a <b>scsi_pkt</b> (9S) structure.																
<b><i>drv_name</i></b>	String used by <b>scsi_log</b> (9F).																
<b><i>severity</i></b>	Error severity level, maps to severity strings below.																
<b><i>blkno</i></b>	Requested block number.																
<b><i>err_blkno</i></b>	Error block number.																
<b><i>cmdlist</i></b>	An array of SCSI command description strings.																
<b><i>sensep</i></b>	A pointer to a <b>scsi_extended_sense</b> (9S) structure.																
<b>DESCRIPTION</b>	<p><b>scsi_errmsg()</b> interprets the request sense information in the <i>sensep</i> pointer and generates a standard message that is displayed using <b>scsi_log</b>(9F). The first line of the message is always a CE_WARN, with the continuation lines being CE_CONT. <i>sensep</i> may be NULL, in which case no sense key or vendor information is displayed.</p> <p>The driver should make the determination as to when to call this function based on the severity of the failure and the severity level that the driver wants to report.</p> <p>The <b>scsi_device</b>(9S) structure denoted by <i>devp</i> supplies the identification of the device that requested the display. <i>severity</i> selects which string is used in the "Error Level:" reporting, according to the following table:</p> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 40px;">Severity Value:</td> <td>String:</td> </tr> <tr> <td>SCSI_ERR_ALL</td> <td>All</td> </tr> <tr> <td>SCSI_ERR_UNKNOWN</td> <td>Unknown</td> </tr> <tr> <td>SCSI_ERR_INFO</td> <td>Informational</td> </tr> </table>	Severity Value:	String:	SCSI_ERR_ALL	All	SCSI_ERR_UNKNOWN	Unknown	SCSI_ERR_INFO	Informational								
Severity Value:	String:																
SCSI_ERR_ALL	All																
SCSI_ERR_UNKNOWN	Unknown																
SCSI_ERR_INFO	Informational																

Severity Value:	String:
SCSI_ERR_RECOVERE	Recovered
SCSI_ERR_RETRYABL	Retryable
SCSI_ERR_FATAL	Fatal

*blkno* is the block number of the original request that generated the error. *err\_blkno* is the block number where the error occurred. *cmdlist* is a mapping table for translating the SCSI command code in *pkt* to the actual command string.

The *cmdlist* is described in the structure below:

```
struct scsi_key_strings {
    int key;
    char *message;
};
```

For a basic SCSI disk, the following list is appropriate:

```
static struct scsi_key_strings scsi_cmds[] = {
    0x00, "test unit ready",
    0x01, "rezero/rewind",
    0x03, "request sense",
    0x04, "format",
    0x07, "reassign",
    0x08, "read",
    0x0a, "write",
    0x0b, "seek",
    0x12, "inquiry",
    0x15, "mode select",
    0x16, "reserve",
    0x17, "release",
    0x18, "copy",
    0x1a, "mode sense",
    0x1b, "start/stop",
    0x1e, "door lock",
    0x28, "read(10)",
    0x2a, "write(10)",
    0x2f, "verify",
    0x37, "read defect data",
    0x3b, "write buffer",
    -1, NULL
};
```

**CONTEXT** `scsi_errmsg()` may be called from user or interrupt context.

**EXAMPLES** **EXAMPLE 1** Generating error information.

This entry:

```
scsi_errmsg(devp, pkt, "sd", SCSI_ERR_INFO, bp->b_blkno,  
            err_blkno, sd_cmds, rqsense);
```

**Generates:**

```
WARNING: /sbus@1,f8000000/esp@0,800000/sd@1,0 (sd1):  
Error for Command: read Error Level: Informational  
Requested Block: 23936 Error Block: 23936  
Vendor: QUANTUM Serial Number: 123456  
Sense Key: Unit Attention  
ASC: 0x29 (reset), ASCQ: 0x0, FRU: 0x0
```

**SEE ALSO**

**cmn\_err(9F)**, **scsi\_log(9F)**, **scsi\_device(9S)**,  
**scsi\_extended\_sense(9S)**, **scsi\_pkt(9S)**

*Writing Device Drivers*

<b>NAME</b>	scsi_free_consistent_buf – free a previously allocated SCSI DMA I/O buffer
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  void scsi_free_consistent_buf(struct buf *bp);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>bp</b> Pointer to the buf(9S) structure.
<b>DESCRIPTION</b>	scsi_free_consistent_buf() frees a buffer header and consistent data buffer that was previously allocated using scsi_alloc_consistent_buf(9F).
<b>CONTEXT</b>	scsi_free_consistent_buf() may be called from either the user or the interrupt levels.
<b>SEE ALSO</b>	freerbuf(9F), scsi_alloc_consistent_buf(9F), buf(9S) <i>Writing Device Drivers</i>
<b>WARNING</b>	scsi_free_consistent_buf() will call freerbuf(9F) to free the buf(9S) that was allocated before or during the call to scsi_alloc_consistent_buf(9F). If consistent memory is bound to a scsi_pkt(9S), the pkt should be destroyed before freeing the consistent memory.

<b>NAME</b>	scsi_hba_attach_setup, scsi_hba_attach, scsi_hba_detach – SCSI HBA attach and detach routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  int scsi_hba_attach_setup(dev_info_t * dip, ddi_dma_attr_t * hba_dma_attr, scsi_hba_tran_t * hba_tran, int hba_flags);  int scsi_hba_attach(dev_info_t * dip, ddi_dma_lim_t * hba_lim, scsi_hba_tran_t * hba_tran, int hba_flags, void * hba_options);  int scsi_hba_detach(dev_info_t * dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b> A pointer to the <code>dev_info_t</code> structure, referring to the instance of the HBA device.</p> <p><b>hba_lim</b> A pointer to a <code>ddi_dma_lim(9S)</code> structure.</p> <p><b>hba_tran</b> A pointer to a <code>scsi_hba_tran(9S)</code> structure.</p> <p><b>hba_flags</b> Flag modifiers. The only defined flag value is <code>SCSI_HBA_TRAN_CLONE</code>.</p> <p><b>hba_options</b> Optional features provided by the HBA driver for future extensions; must be <code>NULL</code>.</p> <p><b>hba_dma_attr</b> A pointer to a <code>ddi_dma_attr(9S)</code> structure.</p>
<b>DESCRIPTION</b>	<p><code>scsi_hba_attach_setup()</code> is the recommended interface over <code>scsi_hba_attach()</code>.</p> <p><code>scsi_hba_attach_setup()</code> registers the DMA limits <code>hba_lim</code> and the transport vectors <code>hba_tran</code> of each instance of the HBA device defined by <code>dip</code>.</p> <p><code>scsi_hba_attach_setup()</code> registers the DMA attributes <code>hba_dma_attr</code> and the transport vectors <code>hba_tran</code> of each instance of the HBA device defined by <code>dip</code>. The HBA driver can pass different DMA limits or DMA attributes, and transport vectors for each instance of the device, as necessary, to support any constraints imposed by the HBA itself.</p> <p><code>scsi_hba_attach()</code> and <code>scsi_hba_attach_setup()</code> use the <code>dev_bus_ops</code> field in the <code>dev_ops(9S)</code> structure. The HBA driver should initialize this field to <code>NULL</code> before calling <code>scsi_hba_attach()</code> or <code>scsi_hba_attach_setup()</code>.</p> <p>If <code>SCSI_HBA_TRAN_CLONE</code> is requested in <code>hba_flags</code>, the <code>hba_tran</code> structure will be cloned once for each target attached to the HBA. The cloning of the</p>
<code>scsi_hba_attach_setup()</code> <code>scsi_hba_attach()</code>	

structure will occur before the `tran_tgt_init(9E)` entry point is called to initialize a target. At all subsequent HBA entry points, including `tran_tgt_init(9E)`, the `scsi_hba_tran_t` structure passed as an argument or found in a `scsi_address` structure will be the 'cloned' `scsi_hba_tran_t` structure, thus allowing the HBA to use the `tran_tgt_private` field in the `scsi_hba_tran_t` structure to point to per-target data. The HBA must take care to free only the same `scsi_hba_tran_t` structure it allocated when detaching; all 'cloned' `scsi_hba_tran_t` structures allocated by the system will be freed by the system.

`scsi_hba_attach()` and `scsi_hba_attach_setup()` attach a number of integer-valued properties to `dip`, unless properties of the same name are already attached to the node. An HBA driver should retrieve these configuration parameters via `ddi_prop_get_int(9F)`, and respect any settings for features provided the HBA.

`scsi-options`                      Optional SCSI configuration bits

SCSI\_OPTIONS\_DR

If not set, the HBA should not grant Disconnect privileges to target devices.

SCSI\_OPTIONS\_LINK

If not set, the HBA should not enable Linked Commands.

SCSI\_OPTIONS\_TAG

If not set, the HBA should not operate in Command Tagged Queueing mode.

SCSI\_OPTIONS\_FAST

If not set, the HBA should not operate the bus in FAST SCSI mode.

SCSI\_OPTIONS\_FAST20

If not set, the HBA should not operate the bus in FAST20 SCSI mode.

SCSI\_OPTIONS\_WIDE

If not set, the HBA should not operate the bus in WIDE SCSI mode.

SCSI\_OPTIONS\_SYNC

If not set, the HBA should not operate the bus

	scsi-reset-delay	SCSI bus or device reset recovery time, in milliseconds.
<b>scsi_hba_detach()</b>	<b>scsi_hba_detach()</b> removes the reference to the DMA limits or attributes structure and the transport vector for the given instance of an HBA driver.	
<b>RETURN VALUES</b>	<b>scsi_hba_attach()</b> , <b>scsi_hba_attach_setup()</b> , and <b>scsi_hba_detach()</b> return DDI_SUCCESS if the function call succeeds, and return DDI_FAILURE on failure.	
<b>CONTEXT</b>	<b>scsi_hba_attach()</b> and <b>scsi_hba_attach_setup()</b> should be called from <b>attach(9E)</b> . <b>scsi_hba_detach()</b> should be called from <b>detach(9E)</b> .	
<b>SEE ALSO</b>	<b>attach(9E)</b> , <b>detach(9E)</b> , <b>tran_tgt_init(9E)</b> , <b>ddi_prop_get_int(9F)</b> , <b>ddi_dma_attr(9S)</b> , <b>ddi_dma_lim(9S)</b> , <b>dev_ops(9S)</b> , <b>scsi_address(9S)</b> , <b>scsi_hba_tran(9S)</b>	
	<i>Writing Device Drivers</i>	
<b>NOTES</b>	It is the HBA driver's responsibility to ensure that no more transport requests will be taken on behalf of any SCSI target device driver after <b>scsi_hba_detach()</b> is called.	

<b>NAME</b>	scsi_hba_init, scsi_hba_fini – SCSI Host Bus Adapter system initialization and completion routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  int scsi_hba_init(struct modlinkage * modlp);  void scsi_hba_fini(struct modlinkage * modlp);</pre>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>modlp</b>            Pointer to the Host Bus Adapters module linkage structure.</p>
<b>DESCRIPTION</b>	
<b>scsi_hba_init()</b>	<b>scsi_hba_init()</b> is the system-provided initialization routine for SCSI HBA drivers. The <b>scsi_hba_init()</b> function registers the HBA in the system and allows the driver to accept configuration requests on behalf of SCSI target drivers. The <b>scsi_hba_init()</b> routine must be called in the HBA's <b>_init(9E)</b> routine before <b>mod_install(9F)</b> is called. If <b>mod_install(9F)</b> fails, the HBA's <b>_init(9E)</b> should call <b>scsi_hba_fini()</b> before returning failure.
<b>scsi_hba_fini()</b>	<b>scsi_hba_fini()</b> is the system provided completion routine for SCSI HBA drivers. <b>scsi_hba_fini()</b> removes all of the system references for the HBA that were created in <b>scsi_hba_init()</b> . The <b>scsi_hba_fini()</b> routine should be called in the HBA's <b>_fini(9E)</b> routine if <b>mod_remove(9F)</b> is successful.
<b>RETURN VALUES</b>	<b>scsi_hba_init()</b> returns 0 if successful, and a non-zero value otherwise. If <b>scsi_hba_init()</b> fails, the HBA's <b>_init()</b> entry point should return the value returned by <b>scsi_hba_init()</b> .
<b>CONTEXT</b>	<b>scsi_hba_init()</b> and <b>scsi_hba_fini()</b> should be called from <b>_init(9E)</b> or <b>_fini(9E)</b> , respectively.
<b>SEE ALSO</b>	<b>_fini(9E)</b> , <b>_init(9E)</b> , <b>mod_install(9F)</b> , <b>mod_remove(9F)</b> , <b>scsi_pktalloc(9F)</b> , <b>scsi_pktfree(9F)</b> , <b>scsi_hba_tran(9S)</b>
<b>NOTES</b>	<i>Writing Device Drivers</i>
<b>NOTES</b>	The HBA is responsible for ensuring that no DDI request routines are called on behalf of its SCSI target drivers once <b>scsi_hba_fini()</b> is called.

<b>NAME</b>	scsi_hba_lookup_capstr – return index matching capability string																																
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  int scsi_hba_lookup_capstr(char *capstr);</pre>																																
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).																																
<b>PARAMETERS</b>	<p><b>capstr</b>            Pointer to a string.</p>																																
<b>DESCRIPTION</b>	<p><b>scsi_hba_lookup_capstr()</b> attempts to match <i>capstr</i> against a known set of capability strings, and returns the defined index for the matched capability, if found.</p> <p>The set of indices and capability strings is:</p> <table border="0"> <tr> <td>SCSI_CAP_DMA_MAX</td> <td>"dma-max" or "dma_max"</td> </tr> <tr> <td>SCSI_CAP_MSG_OUT</td> <td>"msg-out" or "msg_out"</td> </tr> <tr> <td>SCSI_CAP_DISCONNECT</td> <td>"disconnect"</td> </tr> <tr> <td>SCSI_CAP_SYNCHRONOUS</td> <td>"synchronous"</td> </tr> <tr> <td>SCSI_CAP_WIDE_XFER</td> <td>"wide-xfer" or "wide_xfer"</td> </tr> <tr> <td>SCSI_CAP_PARITY</td> <td>"parity"</td> </tr> <tr> <td>SCSI_CAP_INITIATOR_ID</td> <td>"initiator-id"</td> </tr> <tr> <td>SCSI_CAP_UNTAGGED_QING</td> <td>"untagged-qing"</td> </tr> <tr> <td>SCSI_CAP_TAGGED_QING</td> <td>"tagged-qing"</td> </tr> <tr> <td>SCSI_CAP_ARQ</td> <td>"auto-rqsense"</td> </tr> <tr> <td>SCSI_CAP_LINKED_CMDS</td> <td>"linked-cmds"</td> </tr> <tr> <td>SCSI_CAP_SECTOR_SIZE</td> <td>"sector-size"</td> </tr> <tr> <td>SCSI_CAP_TOTAL_SECTORS</td> <td>"total-sectors"</td> </tr> <tr> <td>SCSI_CAP_GEOMETRY</td> <td>"geometry"</td> </tr> <tr> <td>SCSI_CAP_RESET_NOTIFICATION</td> <td>"reset-notification"</td> </tr> <tr> <td>SCSI_CAP_QFULL_RETRIES</td> <td>"qfull-retries"</td> </tr> </table>	SCSI_CAP_DMA_MAX	"dma-max" or "dma_max"	SCSI_CAP_MSG_OUT	"msg-out" or "msg_out"	SCSI_CAP_DISCONNECT	"disconnect"	SCSI_CAP_SYNCHRONOUS	"synchronous"	SCSI_CAP_WIDE_XFER	"wide-xfer" or "wide_xfer"	SCSI_CAP_PARITY	"parity"	SCSI_CAP_INITIATOR_ID	"initiator-id"	SCSI_CAP_UNTAGGED_QING	"untagged-qing"	SCSI_CAP_TAGGED_QING	"tagged-qing"	SCSI_CAP_ARQ	"auto-rqsense"	SCSI_CAP_LINKED_CMDS	"linked-cmds"	SCSI_CAP_SECTOR_SIZE	"sector-size"	SCSI_CAP_TOTAL_SECTORS	"total-sectors"	SCSI_CAP_GEOMETRY	"geometry"	SCSI_CAP_RESET_NOTIFICATION	"reset-notification"	SCSI_CAP_QFULL_RETRIES	"qfull-retries"
SCSI_CAP_DMA_MAX	"dma-max" or "dma_max"																																
SCSI_CAP_MSG_OUT	"msg-out" or "msg_out"																																
SCSI_CAP_DISCONNECT	"disconnect"																																
SCSI_CAP_SYNCHRONOUS	"synchronous"																																
SCSI_CAP_WIDE_XFER	"wide-xfer" or "wide_xfer"																																
SCSI_CAP_PARITY	"parity"																																
SCSI_CAP_INITIATOR_ID	"initiator-id"																																
SCSI_CAP_UNTAGGED_QING	"untagged-qing"																																
SCSI_CAP_TAGGED_QING	"tagged-qing"																																
SCSI_CAP_ARQ	"auto-rqsense"																																
SCSI_CAP_LINKED_CMDS	"linked-cmds"																																
SCSI_CAP_SECTOR_SIZE	"sector-size"																																
SCSI_CAP_TOTAL_SECTORS	"total-sectors"																																
SCSI_CAP_GEOMETRY	"geometry"																																
SCSI_CAP_RESET_NOTIFICATION	"reset-notification"																																
SCSI_CAP_QFULL_RETRIES	"qfull-retries"																																

SCSI\_CAP\_QFULL\_RETRY\_INTERVAL "qfull-retry-interval"

**RETURN VALUES** **scsi\_hba\_lookup\_capstr()** returns a non-negative index value corresponding to the capability string, or `--1` if the string does not match any known capability.

**CONTEXT** **scsi\_hba\_lookup\_capstr()** can be called from user or interrupt context.

**SEE ALSO** `tran_getcap(9E)`, `tran_setcap(9E)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset_notify(9F)`

*Writing Device Drivers*

<b>NAME</b>	scsi_hba_pkt_alloc, scsi_hba_pkt_free – allocate and free a scsi_pkt structure																		
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  struct scsi_pkt * <b>scsi_hba_pkt_alloc</b>(dev_info_t * <i>dip</i>, struct scsi_address * <i>ap</i>, int <i>cmdlen</i>, int <i>statuslen</i>, int <i>tgtlen</i>, int <i>hbalen</i>, int (* <i>callback</i>, caddr_t <i>arg</i>, caddr_t <i>arg</i>);  void <b>scsi_hba_pkt_free</b>(struct scsi_address * <i>ap</i>, struct scsi_pkt * <i>pkt</i>);</pre>																		
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).																		
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="vertical-align: top;"><b><i>dip</i></b></td> <td>Pointer to a dev_info_t structure, defining the HBA driver instance.</td> </tr> <tr> <td style="vertical-align: top;"><b><i>ap</i></b></td> <td>Pointer to a <b>scsi_address(9S)</b> structure, defining the target instance.</td> </tr> <tr> <td style="vertical-align: top;"><b><i>cmdlen</i></b></td> <td>Length in bytes to be allocated for the SCSI command descriptor block ( CDB ).</td> </tr> <tr> <td style="vertical-align: top;"><b><i>statuslen</i></b></td> <td>Length in bytes to be allocated for the SCSI status completion block ( SCB ).</td> </tr> <tr> <td style="vertical-align: top;"><b><i>tgtlen</i></b></td> <td>Length in bytes to be allocated for a private data area for the target driver's exclusive use.</td> </tr> <tr> <td style="vertical-align: top;"><b><i>hbalen</i></b></td> <td>Length in bytes to be allocated for a private data area for the HBA driver's exclusive use.</td> </tr> <tr> <td style="vertical-align: top;"><b><i>callback</i></b></td> <td>Indicates what <b>scsi_hba_pkt_alloc()</b> should do when resources are not available: <p>NULL_FUNC Do not wait for resources. Return a NULL pointer.</p> <p>SLEEP_FUNC Wait indefinitely for resources.</p> </td> </tr> <tr> <td style="vertical-align: top;"><b><i>arg</i></b></td> <td>Must be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><b><i>pkt</i></b></td> <td>A pointer to a <b>scsi_pkt(9S)</b> structure.</td> </tr> </table>	<b><i>dip</i></b>	Pointer to a dev_info_t structure, defining the HBA driver instance.	<b><i>ap</i></b>	Pointer to a <b>scsi_address(9S)</b> structure, defining the target instance.	<b><i>cmdlen</i></b>	Length in bytes to be allocated for the SCSI command descriptor block ( CDB ).	<b><i>statuslen</i></b>	Length in bytes to be allocated for the SCSI status completion block ( SCB ).	<b><i>tgtlen</i></b>	Length in bytes to be allocated for a private data area for the target driver's exclusive use.	<b><i>hbalen</i></b>	Length in bytes to be allocated for a private data area for the HBA driver's exclusive use.	<b><i>callback</i></b>	Indicates what <b>scsi_hba_pkt_alloc()</b> should do when resources are not available: <p>NULL_FUNC Do not wait for resources. Return a NULL pointer.</p> <p>SLEEP_FUNC Wait indefinitely for resources.</p>	<b><i>arg</i></b>	Must be NULL.	<b><i>pkt</i></b>	A pointer to a <b>scsi_pkt(9S)</b> structure.
<b><i>dip</i></b>	Pointer to a dev_info_t structure, defining the HBA driver instance.																		
<b><i>ap</i></b>	Pointer to a <b>scsi_address(9S)</b> structure, defining the target instance.																		
<b><i>cmdlen</i></b>	Length in bytes to be allocated for the SCSI command descriptor block ( CDB ).																		
<b><i>statuslen</i></b>	Length in bytes to be allocated for the SCSI status completion block ( SCB ).																		
<b><i>tgtlen</i></b>	Length in bytes to be allocated for a private data area for the target driver's exclusive use.																		
<b><i>hbalen</i></b>	Length in bytes to be allocated for a private data area for the HBA driver's exclusive use.																		
<b><i>callback</i></b>	Indicates what <b>scsi_hba_pkt_alloc()</b> should do when resources are not available: <p>NULL_FUNC Do not wait for resources. Return a NULL pointer.</p> <p>SLEEP_FUNC Wait indefinitely for resources.</p>																		
<b><i>arg</i></b>	Must be NULL.																		
<b><i>pkt</i></b>	A pointer to a <b>scsi_pkt(9S)</b> structure.																		
<b>DESCRIPTION</b>																			

<b>scsi_hba_pkt_alloc()</b>	<p><b>scsi_hba_pkt_alloc()</b> allocates space for a <code>scsi_pkt</code> structure. HBA drivers should use this interface when allocating a <code>scsi_pkt</code> from their <code>tran_init_pkt(9E)</code> entry point.</p> <p>If <code>callback</code> is <code>NULL_FUNC</code>, <b>scsi_hba_pkt_alloc()</b> may not sleep when allocating resources, and callers should be prepared to deal with allocation failures.</p> <p><b>scsi_hba_pkt_alloc()</b> copies the <code>scsi_address(9S)</code> structure pointed to by <code>ap</code> to the <code>pkt_address</code> field in the <code>scsi_pkt(9S)</code>.</p> <p><b>scsi_hba_pkt_alloc()</b> also allocates memory for these <code>scsi_pkt(9S)</code> data areas, and sets these fields to point to the allocated memory:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>pkt_ha_private</code></td> <td>HBA private data area.</td> </tr> <tr> <td><code>pkt_private</code></td> <td>Target driver private data area.</td> </tr> <tr> <td><code>pkt_scbp</code></td> <td>SCSI status completion block.</td> </tr> <tr> <td><code>pkt_cdbp</code></td> <td>SCSI command descriptor block.</td> </tr> </table>	<code>pkt_ha_private</code>	HBA private data area.	<code>pkt_private</code>	Target driver private data area.	<code>pkt_scbp</code>	SCSI status completion block.	<code>pkt_cdbp</code>	SCSI command descriptor block.
<code>pkt_ha_private</code>	HBA private data area.								
<code>pkt_private</code>	Target driver private data area.								
<code>pkt_scbp</code>	SCSI status completion block.								
<code>pkt_cdbp</code>	SCSI command descriptor block.								
<b>scsi_hba_pkt_free()</b>	<b>scsi_hba_pkt_free()</b> frees the space allocated for the <code>scsi_pkt(9S)</code> structure.								
<b>RETURN VALUES</b>	<b>scsi_hba_pkt_alloc()</b> returns a pointer to the <code>scsi_pkt</code> structure, or <code>NULL</code> if no space is available.								
<b>CONTEXT</b>	<p><b>scsi_hba_pkt_alloc()</b> can be called from user or interrupt context. Drivers must not allow <b>scsi_hba_pkt_alloc()</b> to sleep if called from an interrupt routine.</p> <p><b>scsi_hba_pkt_free()</b> can be called from user or interrupt context.</p>								
<b>SEE ALSO</b>	<p><code>tran_init_pkt(9E)</code>, <code>scsi_address(9S)</code>, <code>scsi_pkt(9S)</code></p> <p><i>Writing Device Drivers</i></p>								

<b>NAME</b>	scsi_hba_probe – default SCSI HBA probe function
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  int scsi_hba_probe(struct scsi_device *sd, int (*waitfunc)(void));</pre>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>sd</b>                   Pointer to a <b>scsi_device</b>(9S) structure describing the target.</p> <p><b>waitfunc</b>            NULL_FUNC or SLEEP_FUNC.</p>
<b>DESCRIPTION</b>	<b>scsi_hba_probe()</b> is a function providing the semantics of <b>scsi_probe</b> (9F). An HBA driver may call <b>scsi_hba_probe()</b> from its <b>tran_tgt_probe</b> (9E) entry point, to probe for the existence of a target on the SCSI bus, or the HBA may set <b>tran_tgt_probe</b> (9E) to point to <b>scsi_hba_probe</b> directly.
<b>RETURN VALUES</b>	See <b>scsi_probe</b> (9F) for the return values from <b>scsi_hba_probe()</b> .
<b>CONTEXT</b>	<b>scsi_hba_probe()</b> should only be called from the HBA's <b>tran_tgt_probe</b> (9E) entry point.
<b>SEE ALSO</b>	<b>tran_tgt_probe</b> (9E), <b>scsi_probe</b> (9F), <b>scsi_device</b> (9S) <i>Writing Device Drivers</i>

<b>NAME</b>	scsi_hba_tran_alloc, scsi_hba_tran_free – allocate and free transport structures
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  scsi_hba_tran_t * <b>scsi_hba_tran_alloc</b>(dev_info_t * <i>dip</i>, int <i>flags</i>);  void <b>scsi_hba_tran_free</b>(scsi_hba_tran_t * <i>hba_tran</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>            Pointer to a dev_info structure, defining the HBA driver instance.</p> <p><b>flag</b>            Flag modifiers. The only possible flag value is SCSI_HBA_CANSLEEP (memory allocation may sleep).</p> <p><b>hba_tran</b>        Pointer to a <b>scsi_hba_tran</b>(9S) structure.</p>
<b>DESCRIPTION</b>	
<b>scsi_hba_tran_alloc()</b>	<p><b>scsi_hba_tran_alloc()</b> allocates a <b>scsi_hba_tran</b>(9S) structure for a HBA driver. The HBA must use this structure to register its transport vectors with the system by using <b>scsi_hba_attach_setup</b>(9F) .</p> <p>If the flag SCSI_HBA_CANSLEEP is set in <i>flags</i> , <b>scsi_hba_tran_alloc()</b> may sleep when allocating resources; otherwise it may not sleep, and callers should be prepared to deal with allocation failures.</p>
<b>scsi_hba_tran_free()</b>	<b>scsi_hba_tran_free()</b> is used to free the <b>scsi_hba_tran</b> (9S) structure allocated by <b>scsi_hba_tran_alloc()</b> .
<b>RETURN VALUES</b>	<b>scsi_hba_tran_alloc()</b> returns a pointer to the allocated transport structure, or NULL if no space is available.
<b>CONTEXT</b>	<p><b>scsi_hba_tran_alloc()</b> can be called from user or interrupt context. Drivers must not allow <b>scsi_hba_tran_alloc()</b> to sleep if called from an interrupt routine.</p> <p><b>scsi_hba_tran_free()</b> can be called from user or interrupt context.</p>
<b>SEE ALSO</b>	<p><b>scsi_hba_attach_setup</b>(9F) , <b>scsi_hba_tran</b>(9S)</p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	scsi_ifgetcap, scsi_ifsetcap – get/set SCSI transport capability										
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  int scsi_ifgetcap(struct scsi_address * ap, char * cap, int whom);  int scsi_ifsetcap(struct scsi_address * ap, char * cap, int value, int whom);</pre>										
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).										
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><b><i>ap</i></b></td> <td>Pointer to the <code>scsi_address</code> structure.</td> </tr> <tr> <td><b><i>cap</i></b></td> <td>Pointer to the string capability identifier.</td> </tr> <tr> <td><b><i>value</i></b></td> <td>Defines the new state of the capability.</td> </tr> <tr> <td><b><i>whom</i></b></td> <td>Determines if all targets or only the specified target is affected.</td> </tr> </table>	<b><i>ap</i></b>	Pointer to the <code>scsi_address</code> structure.	<b><i>cap</i></b>	Pointer to the string capability identifier.	<b><i>value</i></b>	Defines the new state of the capability.	<b><i>whom</i></b>	Determines if all targets or only the specified target is affected.		
<b><i>ap</i></b>	Pointer to the <code>scsi_address</code> structure.										
<b><i>cap</i></b>	Pointer to the string capability identifier.										
<b><i>value</i></b>	Defines the new state of the capability.										
<b><i>whom</i></b>	Determines if all targets or only the specified target is affected.										
<b>DESCRIPTION</b>	<p>The target drivers use <code>scsi_ifsetcap()</code> to set the capabilities of the host adapter driver. A <i>cap</i> is a name-value pair whose name is a null terminated character string and whose value is an integer. The current value of a capability can be retrieved using <code>scsi_ifgetcap()</code>. If <i>whom</i> is 0 all targets are affected, else the target specified by the <code>scsi_address</code> structure pointed to by <i>ap</i> is affected.</p> <p>A device may support only a subset of the capabilities listed below. It is the responsibility of the driver to make sure that these functions are called with a <i>cap</i> supported by the device.</p> <p>The following capabilities have been defined:</p> <table border="0"> <tr> <td style="padding-right: 20px;">dma-max</td> <td>Maximum dma transfer size supported by host adapter.</td> </tr> <tr> <td>msg-out</td> <td>Message out capability supported by host adapter: 0 disables, 1 enables.</td> </tr> <tr> <td>disconnect</td> <td>Disconnect capability supported by host adapter: 0 disables, 1 enables.</td> </tr> <tr> <td>synchronous</td> <td>Synchronous data transfer capability supported by host adapter: 0 disables, 1 enables.</td> </tr> <tr> <td>wide-xfer</td> <td>Wide transfer capability supported by host adapter: 0 disables, 1 enables.</td> </tr> </table>	dma-max	Maximum dma transfer size supported by host adapter.	msg-out	Message out capability supported by host adapter: 0 disables, 1 enables.	disconnect	Disconnect capability supported by host adapter: 0 disables, 1 enables.	synchronous	Synchronous data transfer capability supported by host adapter: 0 disables, 1 enables.	wide-xfer	Wide transfer capability supported by host adapter: 0 disables, 1 enables.
dma-max	Maximum dma transfer size supported by host adapter.										
msg-out	Message out capability supported by host adapter: 0 disables, 1 enables.										
disconnect	Disconnect capability supported by host adapter: 0 disables, 1 enables.										
synchronous	Synchronous data transfer capability supported by host adapter: 0 disables, 1 enables.										
wide-xfer	Wide transfer capability supported by host adapter: 0 disables, 1 enables.										

parity	Parity checking by host adapter: 0 disables, 1 enables.
initiator-id	The host's bus address is returned.
untagged-qing	The host adapter's capability to support internal queuing of commands without tagged queuing: 0 disables, 1 enables.
tagged-qing	The host adapter's capability to support tagged queuing: 0 disables, 1 enables.
auto-rqsense	The host adapter's capability to support auto request sense on check conditions: 0 disables, 1 enables.
sector-size	The target driver sets this capability to inform the HBA of the granularity, in bytes, of DMA breakup; the HBA's DMA limit structure will be set to reflect this limit (see <code>ddi_dma_lim_sparc(9S)</code> or <code>ddi_dma_lim_x86(9S)</code> ). It should be set to the physical disk sector size. This capability defaults to 512.
total-sectors	The target driver sets this capability to inform the HBA of the total number of sectors on the device, as returned from the SCSI <code>get capacity</code> command. This capability must be set before the target driver "gets" the geometry capability.
geometry	This capability returns the HBA geometry of a target disk. The target driver must set the <code>total-sectors</code> capability before "getting" the geometry capability. The geometry is returned as a 32-bit value: the upper 16 bits represent the number of heads per cylinder; the lower 16 bits represent the number of sectors per track. The geometry capability cannot be "set."
reset-notification	The host adapter's capability to support bus reset notification: 0 disables, 1 enables. Refer to <code>scsi_reset_notify(9F)</code> .

linked -cmds	The host adapter's capability to support linked commands: 0 disables, 1 enables.
qfull-retries	This capability enables/disables <code>QUEUE FULL</code> handling. If 0, the HBA will not retry a command when a <code>QUEUE FULL</code> status is returned. If greater than 0, then the HBA driver will retry the command at specified number of times at an interval determined by the "qfull-retry-interval". The range for qfull-retries is 0-255.
qfull-retry-interval	This capability sets the retry interval (in ms) for commands that were completed with a <code>QUEUE FULL</code> status. The range for qfull-retry-intervals is 0-1000 ms.

**RETURN VALUES****scsi\_ifsetcap()** returns:

- 1 If the capability was successfully set to the new value.
- 0 If the capability is not variable.
- 1 If the capability was not defined, or setting the capability to a new value failed.

**scsi\_ifgetcap()** returns the current value of a capability, or:

- 1 If the capability was not defined.

**CONTEXT**

These functions can be called from user or interrupt context.

**EXAMPLES****EXAMPLE 1** Using **scsi\_ifgetcap()**

```

un->un_arq_enabled =
    ((scsi_ifsetcap(&devp->sd_address, "auto-rqsense", 1, 1) == 1)? 1: 0);
if (scsi_ifsetcap(&devp->sd_address, "tagged-qing", 1, 1) == 1) {
    un->un_dp->options |= SD_QUEUEING;
\011    un->un_throttle = MAX_THROTTLE;
} else if (scsi_ifgetcap(&devp->sd_address, "untagged-qing", 0) == 1) {
\011    un->un_dp->options |= SD_QUEUEING;
\011    un->un_throttle = 3;
} else {
\011un->un_dp->options &= ~SD_QUEUEING;
    un->un_throttle = 1;
}

```

**SEE ALSO**

`scsi_reset_notify(9F)`, `ddi_dma_lim_sparc(9S)`,  
`ddi_dma_lim_x86(9S)`, `scsi_address(9S)`, `scsi_arq_status(9S)`

*Writing Device Drivers*

<b>NAME</b>	scsi_init_pkt – prepare a complete SCSI packet
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  struct scsi_pkt *scsi_init_pkt(struct scsi_address *ap, struct scsi_pkt *pkt, struct buf *bp, int cmdlen, int statuslen, int privatelen, int flags, int (*callback)(caddr_t), caddr_t arg);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>ap</b> Pointer to a <b>scsi_address(9S)</b> structure.</p> <p><b>pkt</b> A pointer to a <b>scsi_pkt(9S)</b> structure.</p> <p><b>bp</b> Pointer to a <b>buf(9S)</b> structure.</p> <p><b>cmdlen</b> The required length for the SCSI command descriptor block (CDB) in bytes.</p> <p><b>statuslen</b> The required length for the SCSI status completion block (SCB) in bytes.</p> <p><b>privatelen</b> The required length for the <i>pkt_private</i> area.</p> <p><b>flags</b> Flags modifier.</p> <p><b>callback</b> A pointer to a callback function, <b>NULL_FUNC</b>, or <b>SLEEP_FUNC</b>.</p> <p><b>arg</b> The <i>callback</i> function argument.</p>
<b>DESCRIPTION</b>	<p>Target drivers use <b>scsi_init_pkt()</b> to request the transport layer to allocate and initialize a packet for a SCSI command which possibly includes a data transfer. If <i>pkt</i> is <b>NULL</b>, a new <b>scsi_pkt(9S)</b> is allocated using the HBA driver's packet allocator. The <i>bp</i> is a pointer to a <b>buf(9S)</b> structure. If <i>bp</i> is non-<b>NULL</b> and contains a valid byte count, the <b>buf(9S)</b> structure is also set up for DMA transfer using the HBA driver DMA resources allocator. When <i>bp</i> is allocated by <b>scsi_alloc_consistent_buf(9F)</b>, the <b>PKT_CONSISTENT</b> bit must be set in the <i>flags</i> argument to ensure proper operation. If <i>privatelen</i> is non-zero then additional space is allocated for the <i>pkt_private</i> area of the <b>scsi_pkt(9S)</b>. On return <i>pkt_private</i> points to this additional space. Otherwise <i>pkt_private</i> is a pointer that is typically used to store the <i>bp</i> during execution of the command. In this case <i>pkt_private</i> is <b>NULL</b> on return.</p> <p>The <i>flags</i> argument is a set of bit flags. Possible bits include:</p>

PKT_CONSISTENT	This must be set if the DMA buffer was allocated using <code>scsi_alloc_consistent_buf(9F)</code> . In this case, the HBA driver will guarantee that the data transfer is properly synchronized before performing the target driver's command completion callback.
PKT_DMA_PARTIAL	This may be set if the driver can accept a partial DMA mapping. If set, <code>scsi_init_pkt()</code> will allocate DMA resources with the <code>DDI_DMA_PARTIAL</code> bit set in the <code>dmar_flag</code> element of the <code>ddi_dma_req(9S)</code> structure. The <code>pkt_resid</code> field of the <code>scsi_pkt(9S)</code> structure may be returned with a non-zero value, which indicates the number of bytes for which <code>scsi_init_pkt()</code> was unable to allocate DMA resources. In this case, a subsequent call to <code>scsi_init_pkt()</code> may be made for the same <code>pktp</code> and <code>bp</code> to adjust the DMA resources to the next portion of the transfer. This sequence should be repeated until the <code>pkt_resid</code> field is returned with a zero value, which indicates that with transport of this final portion the entire original request will have been satisfied.

When calling `scsi_init_pkt()` to move already-allocated DMA resources, the `cmdlen`, `statuslen` and `privatelen` fields are ignored.

The last argument `arg` is supplied to the `callback` function when it is invoked.

`callback` indicates what the allocator routines should do when resources are not available:

`NULL_FUNC` Do not wait for resources. Return a `NULL` pointer.

`SLEEP_FUNC` Wait indefinitely for resources.

#### Other Values

`callback` points to a function which is called when resources may have become available. `callback` must return either 0 (indicating that it attempted to allocate resources but again failed to do so), in which case it is put back on a list to be called again later, or 1 indicating either success in allocating resources or indicating that it no longer cares for a retry.

When allocating DMA resources, **scsi\_init\_pkt()** returns the `scsi_pkt` field `pkt_resid` as the number of residual bytes for which the system was unable to allocate DMA resources. A `pkt_resid` of 0 means that all necessary DMA resources were allocated.

**RETURN VALUES**

**scsi\_init\_pkt()** returns `NULL` if the packet or DMA resources could not be allocated. Otherwise, it returns a pointer to an initialized **scsi\_pkt**(9S). If *pktp* was not `NULL` the return value will be *pktp* on successful initialization of the packet.

**CONTEXT**

If *callback* is `SLEEP_FUNC`, then this routine may only be called from user-level code. Otherwise, it may be called from either user or interrupt level. The *callback* function may not block or call routines that block.

**EXAMPLES****EXAMPLE 1** Allocating a Packet Without DMA Resources Attached

To allocate a packet without DMA resources attached, use:

```
pkt = scsi_init_pkt(&devp->sd_address, NULL, NULL, CDB_GROUP1,
                  STATUS_LEN, sizeof (struct my_pkt_private *), 0,
                  sd_runout, sd_unit);
```

**EXAMPLE 2** Allocating a Packet With DMA Resources Attached

To allocate a packet with DMA resources attached use:

```
pkt = scsi_init_pkt(&devp->sd_address, NULL, bp, CDB_GROUP1,
                  STATUS_LEN, 0, 0, NULL_FUNC, NULL);
```

**EXAMPLE 3** Attaching DMA Resources to a Preallocated Packet

To attach DMA resources to a preallocated packet, use:

```
pkt = scsi_init_pkt(&devp->sd_address, old_pkt, bp, 0,
                  0, 0, 0, sd_runout, (caddr_t) sd_unit);
```

**EXAMPLE 4** Allocating a Packet with Consistent DMA Resources Attached

Since the packet is already allocated the *cmdlen*, *statuslen* and *privatelen* are 0. To allocate a packet with consistent DMA resources attached, use:

```
bp = scsi_alloc_consistent_buf(&devp->sd_address, NULL,
                              SENSE_LENGTH, B_READ, SLEEP_FUNC, NULL);
pkt = scsi_init_pkt(&devp->sd_address, NULL, bp, CDB_GROUP0,
                  STATUS_LEN, sizeof (struct my_pkt_private *), PKT_CONSISTENT,
```

```
SLEEP_FUNC, NULL);
```

**EXAMPLE 5** Allocating a Packet with Partial DMA Resources Attached

To allocate a packet with partial DMA resources attached, use:

```
my_pkt = scsi_init_pkt(&devp->sd_address, NULL, bp, CDB_GROUP0,  
    STATUS_LEN, sizeof (struct buf *), PKT_DMA_PARTIAL,  
    SLEEP_FUNC, NULL);
```

**SEE ALSO**

[scsi\\_alloc\\_consistent\\_buf\(9F\)](#), [scsi\\_destroy\\_pkt\(9F\)](#),  
[scsi\\_dmaget\(9F\)](#), [scsi\\_pktalloc\(9F\)](#), [buf\(9S\)](#), [ddi\\_dma\\_req\(9S\)](#),  
[scsi\\_address\(9S\)](#), [scsi\\_pkt\(9S\)](#)

*Writing Device Drivers*

**NOTES**

If a DMA allocation request fails with `DDI_DMA_NOMAPPING`, the `B_ERROR` flag will be set in `bp`, and the `b_error` field will be set to `EFAULT`.

If a DMA allocation request fails with `DDI_DMA_TOOBIG`, the `B_ERROR` flag will be set in `bp`, and the `b_error` field will be set to `EINVAL`.

<b>NAME</b>	scsi_log – display a SCSI-device-related message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; #include &lt;sys/cmn_err.h&gt;  void scsi_log(dev_info_t *dip, char *drv_name, uint_t level, const char *fmt, . . .);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>dip</b>            Pointer to the dev_info structure.</p> <p><b>drv_name</b>       String naming the device.</p> <p><b>level</b>          Error level.</p> <p><b>fmt</b>            Display format.</p>
<b>DESCRIPTION</b>	<p><b>scsi_log()</b> is a utility function that displays a message via the <b>cmn_err(9F)</b> routine. The error levels that can be passed in to this function are <b>CE_PANIC</b>, <b>CE_WARN</b>, <b>CE_NOTE</b>, <b>CE_CONT</b>, and <b>SCSI_DEBUG</b>. The last level is used to assist in displaying debug messages to the console only. <i>drv_name</i> is the short name by which this device is known; example disk driver names are <b>sd</b> and <b>cmdk</b>. If the <i>dev_info_t</i> pointer is <b>NULL</b>, then the <i>drv_name</i> will be used with no unit or long name.</p> <p>If the first character in format is an '!' (exclamation point), the message goes only to the system buffer. If the first character in format is a '^CE_CONT', the message is always sent to the system buffer, but is only written to the console when the system has been booted in verbose mode. See <b>kernel(1M)</b>. If neither condition is met, the '?' character has no effect and is simply ignored.</p> <p>All formatting conversions in use by <b>cmn_err()</b> also work with <b>scsi_log()</b>.</p>
<b>CONTEXT</b>	<b>scsi_log()</b> may be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b></p> <pre>scsi_log(dev, "Disk Unit ", CE_PANIC, "Bad Value %d\n", foo);</pre> <p>generates:</p> <pre>PANIC: /eisa/aha@330,0/cmdk@0,0 (Disk Unit 0): Bad Value 5</pre> <p>This is followed by a PANIC.</p>

**EXAMPLE 2**

```
scsi_log(dev, "sd", CE_WARN, "Label Bad\n");
```

**generates:**

```
WARNING: /sbus@1,f8000000/esp@0,8000000/sd@1,0 (sd1): Label Bad
```

**EXAMPLE 3**

```
scsi_log((dev_info_t *) NULL, "Disk Unit ", CE_NOTE, "Disk Ejected\n");
```

**generates:**

```
Disk Unit: Disk Ejected
```

**EXAMPLE 4**

```
scsi_log(cmdk_unit, "Disk Unit ", CE_CONT, "Disk Inserted\n");
```

**generates:**

```
Disk Inserted
```

**EXAMPLE 5**

```
scsi_log(sd_unit, "sd", SCSI_DEBUG, "We really got here\n");
```

**generates (only to the console):**

```
DEBUG: sd1: We really got here
```

**SEE ALSO**

**kernel(1M)**, **sd(7D)**, **cmn\_err(9F)**, **scsi\_errmsg(9F)**

*Writing Device Drivers*

<b>NAME</b>	scsi_pktalloc, scsi_realloc, scsi_pktfree, scsi_resfree – SCSI packet utility routines
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  struct scsi_pkt * <b>scsi_pktalloc</b>(struct scsi_address* <i>ap</i>, int <i>cmdlen</i>, int <i>statuslen</i>, int (* <i>callback</i> )(void));  struct scsi_pkt * <b>scsi_realloc</b>(struct scsi_address* <i>ap</i>, int <i>cmdlen</i>, int <i>statuslen</i>, opaque_t <i>dmatoken</i>, int (* <i>callback</i> )(void));  void <b>scsi_pktfree</b>(struct scsi_pkt* <i>pkt</i>);  void <b>scsi_resfree</b>(struct scsi_pkt* <i>pkt</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>ap</i></b>                    Pointer to a <code>scsi_address</code> structure.</p> <p><b><i>cmdlen</i></b>                The required length for the SCSI command descriptor block (CDB) in bytes.</p> <p><b><i>statuslen</i></b>            The required length for the SCSI status completion block (SCB) in bytes.</p> <p><b><i>dmatoken</i></b>             Pointer to an implementation-dependent object.</p> <p><b><i>callback</i></b>             A pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code>.</p> <p><b><i>pkt</i></b>                    Pointer to a <code>scsi_pkt(9S)</code> structure.</p>
<b>DESCRIPTION</b>	<p><b>scsi_pktalloc()</b> requests the host adapter driver to allocate a command packet. For commands that have a data transfer associated with them, <b>scsi_realloc()</b> should be used.</p> <p><i>ap</i> is a pointer to a <code>scsi_address</code> structure. Allocator routines use it to determine the associated host adapter.</p> <p><i>cmdlen</i> is the required length for the SCSI command descriptor block. This block is allocated such that a kernel virtual address is established in the <code>pkt_cdbp</code> field of the allocated <code>scsi_pkt</code> structure.</p> <p><i>statuslen</i> is the required length for the SCSI status completion block. The address of the allocated block is placed into the <code>pkt_scbp</code> field of the <code>scsi_pkt</code> structure.</p>

*dmatoken* is a pointer to an implementation dependent object which defines the length, direction, and address of the data transfer associated with this SCSI packet (command). The *dmatoken* must be a pointer to a `buf(9S)` structure. If *dmatoken* is `NULL`, no DMA resources are required by this SCSI command, so none are allocated. Only one transfer direction is allowed per command. If there is an unexpected data transfer phase (either no data transfer phase expected, or the wrong direction encountered), the command is terminated with the `pkt_reason` set to `CMD_DMA_DERR`. *dmatoken* provides the information to determine if the transfer count is correct.

*callback* indicates what the allocator routines should do when resources are not available:

`NULL_FUNC` Do not wait for resources. Return a `NULL` pointer.

`SLEEP_FUNC` Wait indefinitely for resources.

**Other Values** *callback* points to a function which is called when resources may have become available. *callback* must return either 0 (indicating that it attempted to allocate resources but again failed to do so), in which case it is put back on a list to be called again later, or 1 indicating either success in allocating resources or indicating that it no longer cares for a retry.

`scsi_pktfree()` frees the packet.

`scsi_resfree()` free all resources held by the packet and the packet itself.

## RETURN VALUES

Both allocation routines return a pointer to a `scsi_pkt` structure on success, or `NULL` on failure.

## CONTEXT

If *callback* is `SLEEP_FUNC`, then this routine may only be called from user-level code. Otherwise, it may be called from either user or interrupt level. The *callback* function may not block or call routines that block. Both deallocation routines can be called from user or interrupt context.

## SEE ALSO

`scsi_dmafree(9F)`, `scsi_dmaget(9F)`, `buf(9S)`, `scsi_pkt(9S)`

*Writing Device Drivers*

<b>NAME</b>	scsi_poll – run a polled SCSI command on behalf of a target driver
<b>SYNOPSIS</b>	#include <sys/scsi/scsi.h>  int <b>scsi_poll</b> (struct scsi_pkt *pkt);
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>pkt</b> Pointer to the <b>scsi_pkt</b> (9S) structure.
<b>DESCRIPTION</b>	<b>scsi_poll()</b> requests the host adapter driver to run a polled command. Unlike <b>scsi_transport</b> (9F) which runs commands asynchronously, <b>scsi_poll()</b> runs commands to completion before returning. If the <b>pkt_time</b> member of <b>pkt</b> is 0, the value of <b>pkt_time</b> is defaulted to <b>SCSI_POLL_TIMEOUT</b> to prevent an indefinite hang of the system.
<b>RETURN VALUES</b>	<b>scsi_poll()</b> returns: 0        command completed successfully. --1     command failed.
<b>CONTEXT</b>	<b>scsi_poll()</b> can be called from user or interrupt level. This function should not be called when the caller is executing <b>timeout</b> (9F) in the context of a thread.
<b>SEE ALSO</b>	<b>makecom</b> (9F), <b>scsi_transport</b> (9F), <b>scsi_pkt</b> (9S)  <i>Writing Device Drivers</i>
<b>WARNINGS</b>	Since <b>scsi_poll()</b> runs commands to completion before returning, it may require more time than is desirable when called from interrupt context. Therefore, calling <b>scsi_poll</b> from interrupt context is not recommended.

<b>NAME</b>	scsi_probe – utility for probing a scsi device
<b>SYNOPSIS</b>	#include <sys/scsi/scsi.h>  int <b>scsi_probe</b> (struct scsi_device *devp, int (*waitfunc);
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>devp</b> Pointer to a <b>scsi_device(9S)</b> structure  <b>waitfunc</b> NULL_FUNC or SLEEP_FUNC
<b>DESCRIPTION</b>	<b>scsi_probe()</b> determines whether a target/lun is present and sets up the <b>scsi_device</b> structure with inquiry data.  <b>scsi_probe()</b> uses the SCSI Inquiry command to test if the device exists. It may retry the Inquiry command as appropriate. If <b>scsi_probe()</b> is successful, it will allocate space for the <b>scsi_inquiry</b> structure and assign the address to the <b>sd_inq</b> member of the <b>scsi_device(9S)</b> structure. <b>scsi_probe()</b> will then fill in this <b>scsi_inquiry(9S)</b> structure and return <b>SCSIPROBE_EXISTS</b> . If <b>scsi_probe()</b> is unsuccessful, it returns <b>SCSIPROBE_NOMEM</b> in spite of callback set to <b>SLEEP_FUNC</b> .  <b>scsi_unprobe(9F)</b> is used to undo the effect of <b>scsi_probe()</b> .  If the target is a non-CCS device, <b>SCSIPROBE_NONCCS</b> will be returned.  <b>waitfunc</b> indicates what the allocator routines should do when resources are not available; the valid values are: <b>NULL_FUNC</b> Do not wait for resources. Return <b>SCSIPROBE_NOMEM</b> or <b>SCSIPROBE_FAILURE</b>  <b>SLEEP_FUNC</b> Wait indefinitely for resources.
<b>RETURN VALUES</b>	<b>scsi_probe()</b> returns: <b>SCSIPROBE_BUSY</b> Device exists but is currently busy.  <b>SCSIPROBE_EXISTS</b> Device exists and inquiry data is valid.  <b>SCSIPROBE_FAILURE</b> Polled command failure.  <b>SCSIPROBE_NOMEM</b> No space available for structures.  <b>SCSIPROBE_NOMEM_CB</b> No space available for structures but callback request has been queued.

SCSIPROBE\_NONCCS      Device exists but inquiry data is not valid.  
 SCSIPROBE\_NORESP      Device does not respond to an INQUIRY.

**CONTEXT**

**scsi\_probe()** is normally called from the target driver's **probe(9E)** or **attach(9E)** routine. If *waitfunc* is **SLEEP\_FUNC**, then this routine may only be called from user-level code. Otherwise, it may be called from either user or interrupt level.

**EXAMPLES****CODE EXAMPLE 1 Using scsi\_probe()**

```
switch (scsi_probe(devp, NULL_FUNC)) {
default:
case SCSIPROBE_NORESP:
case SCSIPROBE_NONCCS:
case SCSIPROBE_NOMEM:
case SCSIPROBE_FAILURE:
case SCSIPROBE_BUSY:
    break;
case SCSIPROBE_EXISTS:
    switch (devp->sd_inq->inq_dtype) {
    case DTYPE_DIRECT:
        rval = DDI_PROBE_SUCCESS;
        break;
    case DTYPE_RODIRECT:
        rval = DDI_PROBE_SUCCESS;
        break;
    case DTYPE_NOTPRESENT:
    default:
        break;
    }
}
scsi_unprobe(devp);
```

**SEE ALSO**

**attach(9E)**, **probe(9E)**, **scsi\_slave(9F)**, **scsi\_unprobe(9F)**,  
**scsi\_unslave(9F)**, **scsi\_device(9S)**, **scsi\_inquiry(9S)**

*ANSI Small Computer System Interface-2 (SCSI-2)*

*Writing Device Drivers*

**NOTES**

A *waitfunc* function other than **NULL\_FUNC** or **SLEEP\_FUNC** is not supported and may have unexpected results.

<b>NAME</b>	scsi_reset – reset a SCSI bus or target
<b>SYNOPSIS</b>	#include <sys/scsi/scsi.h>  int <b>scsi_reset</b> (struct scsi_address *ap, int level);
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>ap</b> Pointer to the scsi_address structure. <b>level</b> The level of reset required.
<b>DESCRIPTION</b>	<b>scsi_reset()</b> asks the host adapter driver to reset the SCSI bus or a SCSI target as specified by <i>level</i> . If <i>level</i> equals RESET_ALL, the SCSI bus is reset. If it equals RESET_TARGET, <i>ap</i> is used to determine the target to be reset.  On a successful reset, the <code>pkt_reason</code> is set to <code>CMD_RESET</code> and <code>pkt_statistics</code> is OR'd with <code>STAT_BUS_RESET</code> or <code>STAT_DEV_RESET</code> .
<b>RETURN VALUES</b>	<b>scsi_reset()</b> returns: 1        Upon success. 0        Upon failure.
<b>CONTEXT</b>	<b>scsi_reset()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>tran_reset(9E)</b> , <b>tran_reset_notify(9E)</b> , <b>scsi_abort(9F)</b>  <i>Writing Device Drivers</i>

<b>NAME</b>	scsi_reset_notify – notify target driver of bus resets
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  void scsi_reset_notify(struct scsi_address *ap, int flag, void (*callback)(caddr_t), caddr_t arg);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>ap</b> Pointer to the <code>scsi_address</code> structure.</p> <p><b>flag</b> A flag indicating registration or cancellation of the notification request.</p> <p><b>callback</b> A pointer to the target driver's reset notification function.</p> <p><b>arg</b> The callback function argument.</p>
<b>DESCRIPTION</b>	<p><b>scsi_reset_notify()</b> is used by a target driver when it needs to be notified of a bus reset. The bus reset could be issued by the transport layer (e.g. the host bus adapter (HBA) driver or controller) or by another initiator.</p> <p>The argument <i>flag</i> is used to register or cancel the notification. The supported values for <i>flag</i> are as follows:</p> <p>SCSI_RESET_NOTIFY Register <i>callback</i> as the reset notification function for the target driver.</p> <p>SCSI_RESET_CANCEL Cancel the reset notification request.</p> <p>Target drivers can find out whether the HBA driver and controller support reset notification by checking the <code>reset-notification</code> capability using the <b>scsi_ifgetcap(9F)</b> function.</p>
<b>RETURN VALUES</b>	<p>If <i>flag</i> is SCSI_RESET_NOTIFY, <b>scsi_reset_notify()</b> returns:</p> <p>DDI_SUCCESS The notification request has been accepted.</p> <p>DDI_FAILURE The transport layer does not support reset notification or could not accept this request.</p> <p>If <i>flag</i> is SCSI_RESET_CANCEL, <b>scsi_reset_notify()</b> returns:</p> <p>DDI_SUCCESS The notification request has been canceled.</p> <p>DDI_FAILURE No notification request was registered.</p>

**CONTEXT** | `scsi_reset_notify()` can be called from user or interrupt context.

**SEE ALSO** | `scsi_address(9S)`, `scsi_ifgetcap(9F)`

*Writing Device Drivers*

<b>NAME</b>	scsi_setup_cdb – setup SCSI command descriptor block (CDB)
<b>SYNOPSIS</b>	int <b>scsi_setup_cdb</b> (union scsi_cdb *cdbp, uchar_t cmd, uint_t addr, uint_t cnt, uint_t othr_cdb_data);
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>cdbp</b>            Pointer to command descriptor block.</p> <p><b>cmd</b>             The first byte of the SCSI group 0, 1, 2, 4, or 5 CDB.</p> <p><b>addr</b>            Pointer to the location of the data.</p> <p><b>cnt</b>             Data transfer length in units defined by the SCSI device type. For sequential devices <i>cnt</i> is the number of bytes. For block devices, <i>cnt</i> is the number of blocks.</p> <p><b>othr_cdb_data</b>   Additional CDB data.</p>
<b>DESCRIPTION</b>	<p><b>scsi_setup_cdb()</b> function initializes a group 0, 1, 2, 4, or 5 type of command descriptor block pointed to by <i>cdbp</i> using <i>cmd</i>, <i>addr</i>, <i>cnt</i>, <i>othr_cdb_data</i>.</p> <p><i>addr</i> should be set to 0 for commands having no addressing information (for example, group 0 READ command for sequential access devices). <i>othr_cdb_data</i> should be additional CDB data for Group 4 commands; otherwise, it should be set to 0.</p> <p><b>scsi_setup_cdb()</b> function does not set the LUN bits in CDB[1] as the <b>makecom(9F)</b> functions do. Also, the fixed bit for sequential access device commands is not set.</p>
<b>RETURN VALUES</b>	<p><b>scsi_setup_cdb()</b> returns:</p> <p>1            Upon success.</p> <p>0            Upon failure.</p>
<b>CONTEXT</b>	These functions can be called from a user or interrupt context.
<b>SEE ALSO</b>	<p><b>makecom(9F)</b>, <b>scsi_pkt(9S)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>American National Standard Small Computer System Interface-2 (SCSI-2)</i></p> <p><i>American National Standard SCSI-3 Primary Commands (SPC)</i></p>

<b>NAME</b>	scsi_slave – utility for SCSI target drivers to establish the presence of a target
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  int scsi_slave(struct scsi_device *devp, int (*callback)(void));</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>devp</b>            Pointer to a <b>scsi_device(9S)</b> structure.</p> <p><b>callback</b>        Pointer to a callback function, <b>NULL_FUNC</b> or <b>SLEEP_FUNC</b>.</p>
<b>DESCRIPTION</b>	<p><b>scsi_slave()</b> checks for the presence of a SCSI device. Target drivers may use this function in their <b>probe(9E)</b> routines. <b>scsi_slave()</b> determines if the device is present by using a Test Unit Ready command followed by an Inquiry command. If <b>scsi_slave()</b> is successful, it will fill in the <b>scsi_inquiry</b> structure, which is the <b>sd_inq</b> member of the <b>scsi_device(9S)</b> structure, and return <b>SCSI_PROBE_EXISTS</b>. This information can be used to determine if the target driver has probed the correct SCSI device type. <b>callback</b> indicates what the allocator routines should do when DMA resources are not available:</p> <p><b>NULL_FUNC</b>        Do not wait for resources. Return a <b>NULL</b> pointer.</p> <p><b>SLEEP_FUNC</b>        Wait indefinitely for resources.</p> <p><b>Other Values</b>     <b>callback</b> points to a function which is called when resources may have become available. <b>callback</b> must return either 0 (indicating that it attempted to allocate resources but again failed to do so), in which case it is put back on a list to be called again later, or 1 indicating either success in allocating resources or indicating that it no longer cares for a retry.</p>
<b>RETURN VALUES</b>	<p><b>scsi_slave()</b> returns:</p> <p><b>SCSI_PROBE_NOMEM</b>                    No space available for structures.</p> <p><b>SCSI_PROBE_EXISTS</b>                    Device exists and inquiry data is valid.</p> <p><b>SCSI_PROBE_NONCCS</b>                    Device exists but inquiry data is not valid.</p> <p><b>SCSI_PROBE_FAILURE</b>                    Polled command failure.</p> <p><b>SCSI_PROBE_NORESP</b>                    No response to <b>TEST UNIT READY</b>.</p>

**CONTEXT** `scsi_slave()` is normally called from the target driver's `probe(9E)` or `attach(9E)` routine. If `callback` is `SLEEP_FUNC`, then this routine may only be called from user-level code. Otherwise, it may be called from either user or interrupt level. The `callback` function may not block or call routines that block.

**SEE ALSO** `attach(9E)`, `probe(9E)`, `ddi_iopb_alloc(9F)`, `makecom(9F)`, `scsi_dmaget(9F)`, `scsi_ifgetcap(9F)`, `scsi_pktalloc(9F)`, `scsi_poll(9F)`, `scsi_probe(9F)`, `scsi_device(9S)`

*ANSI Small Computer System Interface-2 (SCSI-2)*

*Writing Device Drivers*

<b>NAME</b>	scsi_sync_pkt – synchronize CPU and I/O views of memory
<b>SYNOPSIS</b>	#include <sys/scsi/scsi.h>  void <b>scsi_sync_pkt</b> (struct scsi_pkt *pktp);
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>pktp</b> Pointer to a <b>scsi_pkt</b> (9S) structure.
<b>DESCRIPTION</b>	<p><b>scsi_sync_pkt()</b> is used to selectively synchronize a CPU's or device's view of the data associated with the SCSI packet that has been mapped for I/O. This may involve operations such as flushes of CPU or I/O caches, as well as other more complex operations such as stalling until hardware write buffers have drained.</p> <p>This function need only be called under certain circumstances. When a SCSI packet is mapped for I/O using <b>scsi_init_pkt</b>(9F) and destroyed using <b>scsi_destroy_pkt</b>(9F), then an implicit <b>scsi_sync_pkt()</b> will be performed. However, if the memory object has been modified by either the device or a CPU after the mapping by <b>scsi_init_pkt</b>(9F), then a call to <b>scsi_sync_pkt()</b> is required.</p> <p>If the same <b>scsi_pkt</b> is reused for a data transfer from memory to a device, then <b>scsi_sync_pkt()</b> must be called before calling <b>scsi_transport</b>(9F). If the same packet is reused for a data transfer from a device to memory <b>scsi_sync_pkt()</b> must be called after the completion of the packet but before accessing the data in memory.</p>
<b>CONTEXT</b>	<b>scsi_sync_pkt()</b> may be called from user or interrupt context.
<b>SEE ALSO</b>	<b>tran_sync_pkt</b> (9E), <b>ddi_dma_sync</b> (9F), <b>scsi_destroy_pkt</b> (9F), <b>scsi_init_pkt</b> (9F), <b>scsi_transport</b> (9F), <b>scsi_pkt</b> (9S)
	<i>Writing Device Drivers</i>

<b>NAME</b>	scsi_transport – request by a SCSI target driver to start a command								
<b>SYNOPSIS</b>	#include <sys/scsi/scsi.h>  int <b>scsi_transport</b> (struct scsi_pkt *pkt);								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).								
<b>PARAMETERS</b>	<b>pkt</b> Pointer to a <b>scsi_pkt</b> (9S) structure.								
<b>DESCRIPTION</b>	Target drivers use <b>scsi_transport()</b> to request the host adapter driver to transport a command to the SCSI target device specified by <i>pkt</i> . The target driver must obtain resources for the packet using <b>scsi_init_pkt</b> (9F) prior to calling this function. The packet may be initialized using one of the <b>makecom</b> (9F) functions. <b>scsi_transport()</b> does not wait for the SCSI command to complete. See <b>scsi_poll</b> (9F) for a description of polled SCSI commands. Upon completion of the SCSI command the host adapter calls the completion routine provided by the target driver in the <code>pkt_comp</code> member of the <i>scsi_pkt</i> pointed to by <i>pkt</i> .								
<b>RETURN VALUES</b>	<b>scsi_transport()</b> returns: <table border="0" style="width: 100%;"> <tr> <td style="width: 30%;">TRAN_ACCEPT</td> <td>The packet was accepted by the transport layer.</td> </tr> <tr> <td>TRAN_BUSY</td> <td>The packet could not be accepted because there was already a packet in progress for this target/lun, the host adapter queue was full, or the target device queue was full.</td> </tr> <tr> <td>TRAN_BADPKT</td> <td>The DMA count in the packet exceeded the DMA engine's maximum DMA size.</td> </tr> <tr> <td>TRAN_FATAL_ERROR</td> <td>A fatal error has occurred in the transport layer.</td> </tr> </table>	TRAN_ACCEPT	The packet was accepted by the transport layer.	TRAN_BUSY	The packet could not be accepted because there was already a packet in progress for this target/lun, the host adapter queue was full, or the target device queue was full.	TRAN_BADPKT	The DMA count in the packet exceeded the DMA engine's maximum DMA size.	TRAN_FATAL_ERROR	A fatal error has occurred in the transport layer.
TRAN_ACCEPT	The packet was accepted by the transport layer.								
TRAN_BUSY	The packet could not be accepted because there was already a packet in progress for this target/lun, the host adapter queue was full, or the target device queue was full.								
TRAN_BADPKT	The DMA count in the packet exceeded the DMA engine's maximum DMA size.								
TRAN_FATAL_ERROR	A fatal error has occurred in the transport layer.								
<b>CONTEXT</b>	<b>scsi_transport()</b> can be called from user or interrupt context.								
<b>EXAMPLES</b>	<b>CODE EXAMPLE 1</b> Using <b>scsi_transport()</b>  <pre> if ((status = scsi_transport(rqpkt)) != TRAN_ACCEPT) {     scsi_log(devp, sd_label, CE_WARN,             "transport of request sense pkt fails (0x%x)\n", status); } </pre>								

**SEE ALSO** | `tran_start(9E)`, `makecom(9F)`, `scsi_init_pkt(9F)`, `scsi_pktalloc(9F)`,  
`scsi_poll(9F)`, `scsi_pkt(9S)`

*Writing Device Drivers*

<b>NAME</b>	scsi_unprobe, scsi_unslave – free resources allocated during initial probing
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  void <b>scsi_unslave</b>(struct scsi_device * <i>devp</i>); void <b>scsi_unprobe</b>(struct scsi_device * <i>devp</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b><i>devp</i></b> Pointer to a <b>scsi_device(9S)</b> structure.
<b>DESCRIPTION</b>	<b>scsi_unprobe()</b> and <b>scsi_unslave()</b> are used to free any resources that were allocated on the driver's behalf during <b>scsi_slave(9F)</b> and <b>scsi_probe(9F)</b> activity.
<b>CONTEXT</b>	<b>scsi_unprobe()</b> and <b>scsi_unslave()</b> may be called from either the user or the interrupt levels.
<b>SEE ALSO</b>	<b>scsi_probe(9F)</b> , <b>scsi_slave(9F)</b> , <b>scsi_device(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	scsi_vu_errmsg – display a SCSI request sense message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt;  void <b>scsi_vu_errmsg</b>(struct scsi_pkt *pkt, char *drv_name, int severity, int err_blkno, struct scsi_key_strings *cmdlist, struct scsi_extended_sense *sensep, struct scsi_asq_key_strings *asc_list, char *(*decode_fru)(struct scsi_device*, char *, int, char));</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p>The following parameters are supported:</p> <p><b>devp</b> Pointer to the <b>scsi_device</b>(9S) structure.</p> <p><b>pkt</b> Pointer to a <b>scsi_pkt</b>(9S) structure.</p> <p><b>drv_name</b> String used by <b>scsi_log</b>(9F).</p> <p><b>severity</b> Error severity level, maps to severity strings below.</p> <p><b>blkno</b> Requested block number.</p> <p><b>err_blkno</b> Error block number.</p> <p><b>cmdlist</b> An array of SCSI command description strings.</p> <p><b>sensep</b> A pointer to a <b>scsi_extended_sense</b>(9S) structure.</p> <p><b>asc_list</b> A pointer to a array of asc and ascq message list. The list must be terminated with -1 asc value.</p> <p><b>decode_fru</b> This is a function pointer that will be called after the entire sense information has been decoded. The parameters will be the scsi_device structure to identify the device. Second argument will be a pointer to a buffer of length specified by third argument. The fourth argument will be the FRU byte. decode_fru may be NULL if no special decoding is required. decode_fru is expected to return pointer to a char string if decoding possible and NULL if no decoding is possible.</p>
<b>DESCRIPTION</b>	This function is very similar to <b>scsi_errmsg</b> (9F) but allows decoding of vendor-unique ASC/ASCQ and FRU information.

**scsi\_vu\_errmsg()** interprets the request sense information in the *sensep* pointer and generates a standard message that is displayed using **scsi\_log(9F)**. It first searches the list array for a matching vendor unique code if supplied. If it does not find one in the list then the standard list is searched. The first line of the message is always a `CE_WARN`, with the continuation lines being `CE_CONT`. *sensep* may be `NULL`, in which case no sense key or vendor information is displayed.

The driver should make the determination as to when to call this function based on the severity of the failure and the severity level that the driver wants to report.

The **scsi\_device(9S)** structure denoted by *devp* supplies the identification of the device that requested the display. *severity* selects which string is used in the "Error Level:" reporting, according to the table below:

Severity Value:	String:
<code>SCSI_ERR_ALL</code>	All
<code>SCSI_ERR_UNKNOWN</code>	Unknown
<code>SCSI_ERR_INFO</code>	Information
<code>SCSI_ERR_RECOVERED</code>	Recovered
<code>SCSI_ERR_RETRYABLE</code>	Retryable
<code>SCSI_ERR_FATAL</code>	Fatal

*blkno* is the block number of the original request that generated the error. *err\_blkno* is the block number where the error occurred. *cmdlist* is a mapping table for translating the SCSI command code in *pktp* to the actual command string.

The *cmdlist* is described in the structure below:

```
struct scsi_key_strings {
    int key;
    char *message;
};
```

For a basic SCSI disk, the following list is appropriate:

```
static struct scsi_key_strings scsi_cmds[] = {
    0x00, "test unit ready",
    0x01, "rezero/rewind",
    0x03, "request sense",
    0x04, "format",
    0x07, "reassign",
    0x08, "read",
    0x0a, "write",
    0x0b, "seek",
    0x12, "inquiry",
    0x15, "mode select",
    0x16, "reserve",
    0x17, "release",
    0x18, "copy",
    0x1a, "mode sense",
    0x1b, "start/stop",
```

```

        0x1e, "door lock",
        0x28, "read(10)",
        0x2a, "write(10)",
        0x2f, "verify",
        0x37, "read defect data",
        0x3b, "write buffer",
        -1, NULL
};

```

**CONTEXT**

**scsi\_vu\_errmsg()** may be called from user or interrupt context.

**EXAMPLES****EXAMPLE 1** Using **scsi\_vu\_errmsg()**

```

struct scsi_asq_key_strings cd_slist[] = {
    0x81, 0, "Logical Unit is inaccessible",
    -1, 0, NULL,
};

scsi_vu_errmsg(devp, pkt, "sd",
               SCSI_ERR_INFO, bp->b_blkno, err_blkno,
               sd_cmds, rqsense, cd_slist,
               my_decode_fru);

```

This generates the following console warning:

```

WARNING: /sbus@1,f8000000/esp@0,800000/sd@1,0 (sd1):
Error for Command: read          Error Level: Informational
Requested Block: 23936          Error Block: 23936
Vendor: XYZ                      Serial Number: 123456
Sense Key: Unit Attention
ASC: 0x81 (Logical Unit is inaccessible), ASCQ: 0x0
FRU: 0x11 (replace LUN 1, located in slot 1)

```

**SEE ALSO**

**cmn\_err(9F)**, **scsi\_errmsg(9F)**, **scsi\_log(9F)**, **scsi\_errmsg(9F)**,  
**scsi\_asc\_key\_strings(9S)**, **scsi\_device(9S)**,  
**scsi\_extended\_sense(9S)**, **scsi\_pkt(9S)**

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	semaphore, sema_init, sema_destroy, sema_p, sema_p_sig, sema_v, sema_tryop – semaphore functions
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ksynch.h&gt;  void <b>sema_init</b>(ksema_t * <i>sp</i>, uint_t <i>val</i>, char * <i>name</i>, ksema_type_t <i>type</i>, void * <i>arg</i>); void <b>sema_destroy</b>(ksema_t * <i>sp</i>); void <b>sema_p</b>(ksema_t * <i>sp</i>); void <b>sema_v</b>(ksema_t * <i>sp</i>); int <b>sema_p_sig</b>(ksema_t * <i>sp</i>); int <b>sema_tryop</b>(ksema_t * <i>sp</i>);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI ).
<b>PARAMETERS</b>	<p><b><i>sp</i></b>            A pointer to a semaphore, type <code>ksema_t</code> .</p> <p><b><i>val</i></b>            Initial value for semaphore.</p> <p><b><i>name</i></b>           Descriptive string. This is obsolete and should be <code>NULL</code> . (Non- <code>NULL</code> strings are legal, but they are a waste of kernel memory.)</p> <p><b><i>type</i></b>           Variant type of the semaphore. Currently, only <code>SEMA_DRIVER</code> is supported.</p> <p><b><i>arg</i></b>            Type-specific argument; should be <code>NULL</code> .</p>
<b>DESCRIPTION</b>	<p>These functions implement counting semaphores as described by Dijkstra. A semaphore has a value which is atomically decremented by <b>sema_p()</b> and atomically incremented by <b>sema_v()</b> . The value must always be greater than or equal to zero. If <b>sema_p()</b> is called and the value is zero, the calling thread is blocked until another thread performs a <b>sema_v()</b> operation on the semaphore.</p> <p>Semaphores are initialized by calling <b>sema_init()</b> . The argument, <code>val</code> , gives the initial value for the semaphore. The semaphore storage is provided by the caller but more may be dynamically allocated, if necessary, by <b>sema_init()</b> . For this reason, <b>sema_destroy()</b> should be called before deallocating the storage containing the semaphore.</p>

**sema\_p\_sig()** decrements the semaphore, as does **sema\_p()** . However, if the semaphore value is zero, **sema\_p\_sig()** will return without decrementing the value if a signal (that is, from **kill(2)** ) is pending for the thread.

**sema\_tryop()** will decrement the semaphore value only if it is greater than zero, and will not block.

**RETURN VALUES**

- 0 **sema\_tryop()** could not decrement the semaphore value because it was zero.
- 1 **sema\_p\_sig()** was not able to decrement the semaphore value and detected a pending signal.

**CONTEXT**

These functions can be called from user or interrupt context, except for **sema\_init()** and **sema\_destroy()** , which can be called from user context only. None of these functions can be called from a high-level interrupt context. In most cases, **sema\_v()** and **sema\_p()** should not be called from any interrupt context.

If **sema\_p()** is used from interrupt context, lower-priority interrupts will not be serviced during the wait. This means that if the thread that will eventually perform the **sema\_v()** becomes blocked on anything that requires the lower-priority interrupt, the system will hang.

For example, the thread that will perform the **sema\_v()** may need to first allocate memory. This memory allocation may require waiting for paging I/O to complete, which may require a lower-priority disk or network interrupt to be serviced. In general, situations like this are hard to predict, so it is advisable to avoid waiting on semaphores or condition variables in an interrupt context.

**SEE ALSO**

**kill(2)** , **condvar(9F)** , **mutex(9F)**

*Writing Device Drivers*

<b>NAME</b>	printf – format characters in memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  char *printf(char *buf, const char *fmt, ...);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>buf</b>                    Pointer to a character string.</p> <p><b>fmt</b>                    Pointer to a character string.</p>
<b>DESCRIPTION</b>	<p><b>printf()</b> builds a string in <i>buf</i> under the control of the format <i>fmt</i>. The format is a character string with either plain characters, which are simply copied into <i>buf</i>, or conversion specifications, each of which converts zero or more arguments, again copied into <i>buf</i>. The results are unpredictable if there are insufficient arguments for the format; excess arguments are simply ignored. It is the user's responsibility to ensure that enough storage is available for <i>buf</i>.</p>
<b>Conversion Specifications</b>	<p>Each conversion specification is introduced by the % character, after which the following appear in sequence:</p> <p>An optional value specifying a minimum field width for numeric conversion. The converted value will be right-justified and, if it has fewer characters than the minimum, is padded with leading spaces unless the field width is an octal value, then it is padded with leading zeroes.</p> <p>An optional l (ll) specifying that a following d, D, o, O, x, X, or u conversion character applies to a long (long long) integer argument. An l (ll) before any other conversion character is ignored.</p> <p>A character indicating the type of conversion to be applied: d,D,o,O,x,X,u</p> <p>The integer argument is converted to signed decimal (d, D), unsigned octal (o, O), unsigned hexadecimal (x, X) or unsigned decimal (u), respectively, and copied. The letters abcdef are used for x and X conversion.</p> <p>c</p> <p>The character value of argument is copied.</p> <p>b</p>

This conversion uses two additional arguments. The first is an integer, and is converted according to the base specified in the second argument. The second argument is a character string in the form *<base> [ <arg> . . . ]*. The base supplies the conversion base for the first argument as a binary value; \10 gives octal, \20 gives hexadecimal. Each subsequent *<arg>* is a sequence of characters, the first of which is the bit number to be tested, and subsequent characters, up to the next bit number or terminating null, supply the name of the bit.

A bit number is a binary-valued character in the range 1-32. For each bit set in the first argument, and named in the second argument, the bit names are copied, separated by commas, and bracketed by *<* and *>*. Thus, the following function call would generate `reg=3<BitTwo,BitOne>\n` in *buf*.

```
sprintf(buf, "reg=%b\n", 3, "\10\2BitTwo\1BitOne")
```

s

The argument is taken to be a string (character pointer), and characters from the string are copied until a null character is encountered. If the character pointer is NULL, the string *<null string>* is used in its place.

%

Copy a %; no argument is converted.

#### RETURN VALUES

**sprintf()** returns its first argument, *buf*.

#### CONTEXT

**sprintf()** can be called from user or interrupt context.

#### SEE ALSO

*Writing Device Drivers*

<b>NAME</b>	stoi, numtos – convert between an integer and a decimal string
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  int stoi(char **str);  void numtos(unsigned long num, char * s);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>str</b> Pointer to a character string to be converted.</p> <p><b>num</b> Decimal number to be converted to a character string.</p> <p><b>s</b> Character buffer to hold converted decimal number.</p>
<b>DESCRIPTION</b>	
<b>stoi()</b>	<b>stoi()</b> returns the integer value of a string of decimal numeric characters beginning at <b>**str</b> . No overflow checking is done. <b>*str</b> is updated to point at the last character examined.
<b>numtos()</b>	<b>numtos()</b> converts a <code>long</code> into a null-terminated character string. No bounds checking is done. The caller must ensure there is enough space to hold the result.
<b>RETURN VALUES</b>	<b>stoi()</b> returns the integer value of the string <i>str</i> .
<b>CONTEXT</b>	<b>stoi()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>
<b>NOTES</b>	<b>stoi()</b> handles only positive integers; it does not handle leading minus signs.

<b>NAME</b>	strchr – find a character in a string
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  char *strchr(const char *str, int chr);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>str</b>     Pointer to a string to be searched.</p> <p><b>chr</b>     The character to search for.</p>
<b>DESCRIPTION</b>	<b>strchr()</b> returns a pointer to the first occurrence of <i>chr</i> in the string pointed to by <i>str</i> .
<b>RETURN VALUES</b>	<b>strchr()</b> returns a pointer to a character, or NULL, if the search fails.
<b>CONTEXT</b>	This function can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>strcmp(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	strcmp, strncmp – compare two null-terminated strings.
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt;  int strcmp(const char * s1, const char * s2); int strncmp(const char * s1, const char * s2, size_t n);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b>s1</b> , <b>s2</b>            Pointers to character strings.</p> <p><b>n</b>                    Count of characters to be compared.</p>
<b>DESCRIPTION</b>	
<b>strcmp()</b>	<b>strcmp()</b> returns 0 if the strings are the same, or the integer value of the expression <i>(*s1 - *s2)</i> for the last characters compared if they differ.
<b>strncmp()</b>	<b>strncmp()</b> returns 0 if the first <i>n</i> characters of <i>s1</i> and <i>s2</i> are the same, or <i>(*s1 - *s2)</i> for the last characters compared if they differ.
<b>RETURN VALUES</b>	<p><code>strcmp</code> () returns 0 if the strings are the same, or <i>(*s1 - *s2)</i> for the last characters compared if they differ.</p> <p><b>strncmp()</b> returns 0 if the first <i>n</i> characters of strings are the same, or <i>(*s1 - *s2)</i> for the last characters compared if they differ.</p>
<b>CONTEXT</b>	These functions can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	strcpy, strncpy – copy a string from one location to another.
<b>SYNOPSIS</b>	#include <sys/ddi.h>  char * <b>strcpy</b> (char * <i>dst</i> , char * <i>srs</i> ); char * <b>strncpy</b> (char * <i>dst</i> , char * <i>srs</i> , size_t <i>n</i> );
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b><i>dst</i></b> , <b><i>srs</i></b> Pointers to character strings.  <b><i>n</i></b> Count of characters to be copied.
<b>DESCRIPTION</b>	
<b>strcpy()</b>	<b>strcpy()</b> copies characters in the string <i>srs</i> to <i>dst</i> , terminating at the first null character in <i>srs</i> , and returns <i>dst</i> to the caller. No bounds checking is done.
<b>strncpy()</b>	<b>strncpy()</b> copies <i>srs</i> to <i>dst</i> , null-padding or truncating at <i>n</i> bytes, and returns <i>dst</i> . No bounds checking is done.
<b>RETURN VALUES</b>	<b>strcpy()</b> and <b>strncpy()</b> return <i>dst</i> .
<b>CONTEXT</b>	<b>strcpy()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	strlen – determine the number of non-null bytes in a string
<b>SYNOPSIS</b>	#include <sys/ddi.h>  size_t <b>strlen</b> (const char *s);
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<b>s</b> Pointer to a character string.
<b>DESCRIPTION</b>	<b>strlen()</b> returns the number of non-null bytes in the string argument <i>s</i> .
<b>RETURN VALUES</b>	<b>strlen()</b> returns the number of non-null bytes in <i>s</i> .
<b>CONTEXT</b>	<b>strlen()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

<b>NAME</b>	strlog – submit messages to the log driver														
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/strlog.h&gt; #include &lt;sys/log.h&gt;</pre> <p>int <b>strlog</b>(short <i>mid</i>, short <i>sid</i>, char <i>level</i>, unsigned short <i>flags</i>, char *<i>fmt</i>, ...);</p>														
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).														
<b>PARAMETERS</b>	<p><b><i>mid</i></b> Identification number of the module or driver submitting the message (in the case of a module, its <code>mi_idnum</code> value from <code>module_info(9S)</code>).</p> <p><b><i>sid</i></b> Identification number for a particular minor device.</p> <p><b><i>level</i></b> Tracing level for selective screening of low priority messages. Larger values imply less important information.</p> <p><b><i>flags</i></b> Valid flag values are:</p> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;">SL_ERROR</td> <td>Message is for error logger.</td> </tr> <tr> <td>SL_TRACE</td> <td>Message is for trace.</td> </tr> <tr> <td>SL_NOTIFY</td> <td>Mail copy of message to system administrator.</td> </tr> <tr> <td>SL_CONSOLE</td> <td>Log message to console.</td> </tr> <tr> <td>SL_FATAL</td> <td>Error is fatal.</td> </tr> <tr> <td>SL_WARN</td> <td>Error is a warning.</td> </tr> <tr> <td>SL_NOTE</td> <td>Error is a notice.</td> </tr> </table> <p><b><i>fmt</i></b> <code>printf(3S)</code> style format string. <code>%e</code>, <code>%g</code>, and <code>%G</code> formats are not allowed but <code>%s</code> is supported.</p>	SL_ERROR	Message is for error logger.	SL_TRACE	Message is for trace.	SL_NOTIFY	Mail copy of message to system administrator.	SL_CONSOLE	Log message to console.	SL_FATAL	Error is fatal.	SL_WARN	Error is a warning.	SL_NOTE	Error is a notice.
SL_ERROR	Message is for error logger.														
SL_TRACE	Message is for trace.														
SL_NOTIFY	Mail copy of message to system administrator.														
SL_CONSOLE	Log message to console.														
SL_FATAL	Error is fatal.														
SL_WARN	Error is a warning.														
SL_NOTE	Error is a notice.														
<b>DESCRIPTION</b>	<b>strlog()</b> expands the <code>printf(3S)</code> style format string passed to it, that is, the conversion specifiers are replaced by the actual argument values in the format string. The 32-bit representations of the arguments (up to <code>NLOGARGS</code> ) follow the string starting at the next 32-bit boundary following the string. Note that the 64-bit argument will be truncated to 32-bits here but will be fully represented in the string.														

The messages can be retrieved with the `getmsg(2)` system call. The *flags* argument specifies the type of the message and where it is to be sent. `strace(1M)` receives messages from the `log` driver and sends them to the standard output. `strerr(1M)` receives error messages from the `log` driver and appends them to a file called `/var/adm/streams/error.mm-dd`, where *mm-dd* identifies the date of the error message.

**RETURN VALUES**

`strlog()` returns 0 if it fails to submit the message to the `log(7D)` driver and 1 otherwise.

**CONTEXT**

`strlog()` can be called from user or interrupt context.

**FILES**

`/var/adm/streams/error.mm-dd`

Error messages dated *mm-dd* appended by `strerr(1M)` from the `log` driver

**SEE ALSO**

`strace(1M)`, `strerr(1M)`, `getmsg(2)`, `log(7D)`, `module_info(9S)`

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	strqget – get information about a queue or band of the queue																
<b>SYNOPSIS</b>	#include <sys/stream.h>  int <b>strqget</b> (queue_t *q, qfields_t what, unsigned char pri, void *valp);																
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).																
<b>PARAMETERS</b>	<p><b>q</b> Pointer to the queue.</p> <p><b>what</b> Field of the queue structure for (or the specified priority band) to return information about. Valid values are one of:</p> <table border="0" style="margin-left: 2em;"> <tr><td>QHIWAT</td><td>High water mark.</td></tr> <tr><td>QLOWAT</td><td>Low water mark.</td></tr> <tr><td>QMAXPSZ</td><td>Largest packet accepted.</td></tr> <tr><td>QMINPSZ</td><td>Smallest packet accepted.</td></tr> <tr><td>QCOUNT</td><td>Approximate size (in bytes) of data.</td></tr> <tr><td>QFIRST</td><td>First message.</td></tr> <tr><td>QLAST</td><td>Last message.</td></tr> <tr><td>QFLAG</td><td>Status.</td></tr> </table> <p><b>pri</b> Priority band of interest.</p> <p><b>valp</b> The address of where to store the value of the requested field.</p>	QHIWAT	High water mark.	QLOWAT	Low water mark.	QMAXPSZ	Largest packet accepted.	QMINPSZ	Smallest packet accepted.	QCOUNT	Approximate size (in bytes) of data.	QFIRST	First message.	QLAST	Last message.	QFLAG	Status.
QHIWAT	High water mark.																
QLOWAT	Low water mark.																
QMAXPSZ	Largest packet accepted.																
QMINPSZ	Smallest packet accepted.																
QCOUNT	Approximate size (in bytes) of data.																
QFIRST	First message.																
QLAST	Last message.																
QFLAG	Status.																
<b>DESCRIPTION</b>	<b>strqget()</b> gives drivers and modules a way to get information about a queue or a particular band of a queue without directly accessing STREAMS data structures, thus insulating them from changes in the implementation of these data structures from release to release.																
<b>RETURN VALUES</b>	On success, 0 is returned and the value of the requested field is stored in the location pointed to by <i>valp</i> . An error number is returned on failure.																
<b>CONTEXT</b>	<b>strqget()</b> can be called from user or interrupt context.																
<b>SEE ALSO</b>	<b>freezestr(9F)</b> , <b>strqset(9F)</b> , <b>unfreezestr(9F)</b> , <b>queue(9S)</b>  <i>Writing Device Drivers</i>																

*STREAMS Programming Guide*

NOTES

The stream must be frozen using `freezestr(9F)` before calling `strqget()`.

<b>NAME</b>	strqset – change information about a queue or band of the queue								
<b>SYNOPSIS</b>	#include <sys/stream.h>  int <b>strqset</b> (queue_t *q, qfields_t what, unsigned char pri, intptr_t val);								
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).								
<b>PARAMETERS</b>	<p><b>q</b> Pointer to the queue.</p> <p><b>what</b> Field of the queue structure (or the specified priority band) to return information about. Valid values are one of:</p> <table border="0"> <tr> <td>QHIWAT</td> <td>High water mark.</td> </tr> <tr> <td>QLOWAT</td> <td>Low water mark.</td> </tr> <tr> <td>QMAXPSZ</td> <td>Largest packet accepted.</td> </tr> <tr> <td>QMINPSZ</td> <td>Smallest packet accepted.</td> </tr> </table> <p><b>pri</b> Priority band of interest.</p> <p><b>val</b> The value for the field to be changed.</p>	QHIWAT	High water mark.	QLOWAT	Low water mark.	QMAXPSZ	Largest packet accepted.	QMINPSZ	Smallest packet accepted.
QHIWAT	High water mark.								
QLOWAT	Low water mark.								
QMAXPSZ	Largest packet accepted.								
QMINPSZ	Smallest packet accepted.								
<b>DESCRIPTION</b>	<b>strqset()</b> gives drivers and modules a way to change information about a queue or a particular band of a queue without directly accessing STREAMS data structures.								
<b>RETURN VALUES</b>	On success, 0 is returned. EINVAL is returned if an undefined attribute is specified.								
<b>CONTEXT</b>	<b>strqset()</b> can be called from user or interrupt context.								
<b>SEE ALSO</b>	<b>freezestr(9F)</b> , <b>strqget(9F)</b> , <b>unfreezestr(9F)</b> , <b>queue(9S)</b>  <i>Writing Device Drivers</i>  <i>STREAMS Programming Guide</i>								
<b>NOTES</b>	The stream must be frozen using <b>freezestr(9F)</b> before calling <b>strqset()</b>  To set the values of QMINPSZ and QMAXPSZ from within a single call to <b>freezestr(9F)</b> and <b>unfreezestr(9F)</b> : when lowering the existing values, set QMINPSZ before setting QMAXPSZ; when raising the existing values, set QMAXPSZ before setting QMINPSZ.								

<b>NAME</b>	STRUCT_DECL, SIZEOF_PTR, SIZEOF_SIZE, STRUCT_BUF, STRUCT_FADDR, STRUCT_FGET, STRUCT_FGETP, STRUCT_FSET, STRUCT_FSETP, STRUCT_HANDLE, STRUCT_INIT, STRUCT_SET_HANDLE – 32-bit application data access macros
<b>SYNOPSIS</b>	<pre> \011#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  STRUCT_DECL(struct_type handle, )  STRUCT_HANDLE(struct_type, handle);  void STRUCT_INIT(handle, model_t umodel, )  void STRUCT_SET_HANDLE(handle, model_t umodel, void *addr);  STRUCT_FGET(handle, field);  STRUCT_FGETP(handle, field);  STRUCT_FSET(handle, field, val);  STRUCT_FSETP(handle, field, val);  &lt;typeof field&gt; *STRUCT_FADDR(handle, field);  struct struct_type *STRUCT_BUF(handle);  size_t SIZEOF_STRUCT(struct_type, umodel);  size_t SIZEOF_PTR(umodel);  size_t SIZEOF_SIZE(handle); </pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p>The macros take the following parameters:</p> <p><b><i>struct_type</i></b>                   The structure name (as would appear <i>after</i> the C keyword “struct” )of the native form.</p> <p><b><i>umodel</i></b>                           A bit field containing either ILP32 model bit (DATAMODEL_ILP32), or the LP64 model get (DATAMODEL_ILP64). In an <code>ioctl(9E)</code> , these bits will be present in the flag parameter; in a <code>devmap(9E)</code> , they will be present in the model parameter <code>mmap(9E)</code> and can call <code>ddi_mmap_get_model(9F)</code> to get the data model of the current thread.</p>

<b>DESCRIPTION</b>	<p><b><i>handle</i></b>                    The variable name used to refer to a particular instance of a structure which is handled by these macros.</p> <p><b><i>field</i></b>                        The field name within the structure contain substructures. If the structures contain substructures, unions, or arrays, then <i>field</i> can be whether complex expression could occur after the first “.” or “-&gt;” .</p> <p>The above macros allow a device driver to access data consumed from a 32-bit application regardless whether the driver was compiled to the ILP32 or LP64 data model. These macros effectively hide the difference between the data model of the user application and the driver.</p> <p>The macros can be broken up into two main categories, the macros that declare and initialize structure handles and the macros that operate on these structures using the structure handles.</p>
<b>Declaration and Initialization Macros</b>	<p>The macros <b>STRUCT_DECL()</b> and <b>STRUCT_HANDLE()</b> declare structure handles on the stack, whereas the macros <b>STRUCT_INIT()</b> and <b>STRUCT_SET_HANDLE()</b> initialize the structure handles to point to an instance of the native form structure.</p> <p>The macros <b>STRUCT_HANDLE()</b> and <b>STRUCT_SET_HANDLE()</b> are used to declare and initialize a structure handle to an existing data structure, for example, ioctls within a STREAMS module.</p> <p>The macros <b>STRUCT_DECL()</b> and <b>STRUCT_INIT()</b> , on the other hand, are used in modules which declare and initialize a structure handle to a data structure allocated by <b>STRUCT_DECL()</b> , that is, any standard character or block device driver <b>ioctl(9E)</b> routine that needs to copy in data from a user-mode program.</p> <p><b>STRUCT_DECL(struct_type, handle)</b></p> <p style="padding-left: 40px;">Declares a “structure handle” for a “struct” and \011allocates an instance of its native form on the stack. It is assumed that the native form is larger than or equal to the ILP32 form. <i>handle</i> is a variable name and is declared as a variable by this macro.</p> <p><b>void STRUCT_INIT(handle, model_t umodel)</b></p> <p style="padding-left: 40px;">Initializes <i>handle</i> to point to the instance allocated by <b>STRUCT_DECL()</b> , it also sets data model for <i>handle</i> to <i>umodel</i> , and must be called before any access is made through the macros that operate on these structures. When</p>

used in an `ioctl(9E)` routine `umodel` is the flag parameter; in a `devmap(9E)` routine `umodel` is the model parameter and in a `mmap(9E)` routine, is the return value of `ddi_mmap_get_model(9F)`. This macro is intended for handles created with `STRUCT_DECL()` only.

#### **STRUCT\_HANDLE(struct\_type, handle)**

Declares a "structure handle" `handle` but unlike `STRUCT_DECL()` does not allocate an instance of "struct".

#### **void STRUCT\_SET\_HANDLE(handle, model\_t umodel, void \*addr)**

Initializes to point to the native form instance at `addr`, it also sets the data model for `handle` to `umodel`. This is intended for handles created with `STRUCT_HANDLE()`. Fields cannot be referenced via the `handle` until this macro has been invoked. Typically, `addr` is the address of the native form structure containing the user-mode programs data. When used in an `ioctl(9E)` routine `umodel` is the flag parameter, in a `devmap(9E)` routine is the model parameter and in a `mmap(9E)` routine, `umodel` is the return value of `ddi_mmap_get_model(9F)`.

### Operation Macros

#### **size\_t STRUCT\_SIZE(handle)**

Returns size of the structure referred to by `handle`. It will return the size depending upon the data model associated with `handle`. If the data model stored by `STRUCT_INIT()` or `STRUCT_SET_HANDLE()` was `DATAMODEL_ILP32`, it will return the size of the ILP32 form, else it will return the size of the native form.

#### **STRUCT\_FGET(handle, field)**

Returns the contents of `field` in the structure described by `handle` according to the data model associated with `handle`.

#### **STRUCT\_FGETP(handle, field)**

This is the same as `STRUCT_FGET()` except that the `field` in question is a pointer of some kind. This macro will cast `caddr32_t` to a `(void *)` when it is accessed. Failure to use this macro for a pointer will lead to compiler warnings or failures.

#### **STRUCT\_FSET(handle, field, val)**

\011Assigns *val* to the (non pointer) in the structure *handle* described by . It should not be used within any other expression, but rather only as a statement.

**STRUCT\_FSETP(handle, field, val)**

Returns a pointer to the in the structure described by *handle* .

**struct struct\_type \*STRUCT\_BUF(handle)**

Returns a pointer to the native mode instance of the structure described by *handle* .

**Macros Not Using  
Handles**

**size\_t SIZEOF\_STRUCT(struct\_type, umodel)**

Returns size of *struct\_type* based on *umodel* .

**size\_t SIZEOF\_PTR(umodel)**

Returns the size of a pointer based on *umodel* .

**EXAMPLES**

**EXAMPLE 1 Copying a Structure**

The following example uses an `ioctl(9E)` on a regular character device that copies a data structure that looks like this into the kernel:

```
struct opdata {
    size_t  size;
    uint_t  flag;
};
```

**EXAMPLE 2 Defining a Structure**

This data structure definition describes what the `ioctl(9E)` would look like in a 32-bit application using fixed width types.

```
#if defined(_MULTI_DATAMODEL)
struct opdata32 {
    size32_t  size;
    uint32_t  flag;
};
#endif
```

**EXAMPLE 3 Using STRUCT\_DECL() and STRUCT\_INIT()**

Note: This example uses the `STRUCT_DECL()` and `STRUCT_INIT()` macros to declare and initialize the structure handle.

```

int
xxioctl(dev_t dev, int cmd, intptr_t arg, int mode,
        cred_t *cr, int *rval_p);
{
    STRUCT_DECL(opdata, op);

    if (cmd != OPONE)
        return (ENOTTY);

    STRUCT_INIT(op, mode);

    if (copyin((void *)data,
              STRUCT_BUF(op), STRUCT_SIZE(op)))
        return (EFAULT);

    if (STRUCT_FGET(op, flag) != FACTIVE ||
        STRUCT_FGET(op, size) > sizeof (device_state))
        return (EINVAL);
    xxdowork(device_state, STRUCT_FGET(op, size));
    return (0);
}

```

This piece of code is an excerpt from a STREAMS module that handles `ioctl(9E)` data (`M_IOCTLDATA`) messages and uses the data structure defined above. This code has been written to run in the ILP32 environment only.

#### EXAMPLE 4 Using `STRUCT_HANDLE()` and `STRUCT_SET_HANDLE()`

The next example illustrates the use of the `STRUCT_HANDLE()` and `STRUCT_SET_HANDLE()` macros which declare and initialize the structure handle to point to an already existing instance of the structure.

The above code example can be converted to run in the LP64 environment using the `STRUCT_HANDLE()` and `STRUCT_SET_HANDLE()` as follows:

```

struct strbuf {
    int maxlen; /* no. of bytes in buffer */
    int len; /* no. of bytes returned */
    caddr_t buf; /* pointer to data */
};

static void
wput_iocdata(queue_t *q, mblk_t *msgp)
{
    mblk_t *data; /* message block descriptor */
    STRUCT_HANDLE(strbuf, sb);

    /* copyin the data */
    if (mi_copy_state(q, mp, &data) == -1) {
        return;
    }

    STRUCT_SET_HANDLE(sb, ((struct iocblk *)msgp->b_rptr)->ioc_flag,
                     (void *)data->b_rptr);
    if (STRUCT_FGET(sb, maxlen) < (int)sizeof (ipa_t)) {
        mi_copy_done(q, msgp, EINVAL);
    }
}

```

```
        return;  
    }  
}
```

**SEE ALSO** [devmap\(9E\)](#), [ioctl\(9E\)](#), [mmap\(9E\)](#), [ddi\\_mmap\\_get\\_model\(9F\)](#)

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	swab – swap bytes in 16-bit halfwords
<b>SYNOPSIS</b>	<pre>#include &lt;sys/sunddi.h&gt;  void swab(void *src, void *dst, size_t nbytes);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>src</b> A pointer to the buffer containing the bytes to be swapped.</p> <p><b>dst</b> A pointer to the destination buffer where the swapped bytes will be written. If <i>dst</i> is the same as <i>src</i> the buffer will be swapped in place.</p> <p><b>nbytes</b> Number of bytes to be swapped, rounded down to the nearest half-word.</p>
<b>DESCRIPTION</b>	<b>swab()</b> copies the bytes in the buffer pointed to by <i>src</i> to the buffer pointer to by <i>dst</i> , swapping the order of adjacent bytes in half-word pairs as the copy proceeds. A total of <i>nbytes</i> bytes are copied, rounded down to the nearest half-word.
<b>CONTEXT</b>	<b>swab()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>
<b>NOTES</b>	Since <b>swab()</b> operates byte-by-byte, it can be used on non-aligned buffers.

<b>NAME</b>	testb – check for an available buffer
<b>SYNOPSIS</b>	#include <sys/stream.h>  int testb(size_t size, uint_t pri);
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>size</b> Size of the requested buffer.  <b>pri</b> Priority of the allocb request.
<b>DESCRIPTION</b>	<b>testb()</b> checks to see if an <b>allocb(9F)</b> call is likely to succeed if a buffer of <b>size</b> bytes at priority <b>pri</b> is requested. Even if <b>testb()</b> returns successfully, the call to <b>allocb(9F)</b> can fail. The <b>pri</b> argument is no longer used, but is retained for compatibility.
<b>RETURN VALUES</b>	Returns 1 if a buffer of the requested size is available, and 0 if one is not.
<b>CONTEXT</b>	<b>testb()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<b>EXAMPLE 1 testb() example</b>  In a service routine, if <b>copymsg(9F)</b> fails (line 6), the message is put back on the queue (line 7) and a routine, <b>tryagain</b> , is scheduled to be run in one tenth of a second. Then the service routine returns.  When the <b>timeout(9F)</b> function runs, if there is no message on the front of the queue, it just returns. Otherwise, for each message block in the first message, check to see if an allocation would succeed. If the number of message blocks equals the number we can allocate, then enable the service procedure. Otherwise, reschedule <b>tryagain</b> to run again in another tenth of a second. Note that <b>tryagain</b> is merely an approximation. Its accounting may be faulty. Consider the case of a message comprised of two 1024-byte message blocks. If there is only one free 1024-byte message block and no free 2048-byte message blocks, then <b>testb()</b> will still succeed twice. If no message blocks are freed of these sizes before the service procedure runs again, then the <b>copymsg(9F)</b> will still fail. The reason <b>testb()</b> is used here is because it is significantly faster than calling <b>copymsg</b> . We must minimize the amount of time spent in a <b>timeout()</b> routine.  <pre>1 xxxsrv(q) 2     queue_t *q; 3     { 4 mblk_t *mp; 5 mblk_t *nmp;</pre>

```

6 if ((nmp = copymsg(mp)) == NULL) {
7   putbq(q, mp);
8   timeout(tryagain, (intptr_t)q, drv_usectohz(100000));
9   return;
10 }
11 }
12
13 tryagain(q)
14   queue_t *q;
15 {
16   register int can_alloc = 0;
17   register int num_blks = 0;
18   register mblk_t *mp;
19
20   if (!q->q_first)
21     return;
22   for (mp = q->q_first; mp; mp = mp->b_cont) {
23     num_blks++;
24     can_alloc += testb((mp->b_datap->db_lim -
25       mp->b_datap->db_base), BPRI_MED);
26   }
27   if (num_blks == can_alloc)
28     qenable(q);
29   else
30     timeout(tryagain, (intptr_t)q, drv_usectohz(100000));
31 }

```

**SEE ALSO**    `allocb(9F)`, `bufcall(9F)`, `copymsg(9F)`, `timeout(9F)`

*Writing Device Drivers*

*STREAMS Programming Guide*

**NOTES**    The *pri* argument is provided for compatibility only. Its value is ignored.

<b>NAME</b>	timeout – execute a function after a specified length of time
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/conf.h&gt;</pre> <p>timeout_id_t <b>timeout</b>(void (* <i>func</i>)(void *), void *<i>arg</i>, clock_t <i>ticks</i>);</p>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b><i>func</i></b>            Kernel function to invoke when the time increment expires.</p> <p><b><i>arg</i></b>              Argument to the function.</p> <p><b><i>ticks</i></b>             Number of clock ticks to wait before the function is called.</p>
<b>DESCRIPTION</b>	<p>The <b>timeout()</b> function schedules the specified function to be called after a specified time interval. The exact time interval over which the timeout takes effect cannot be guaranteed, but the value given is a close approximation.</p> <p>The function called by <b>timeout()</b> must adhere to the same restrictions as a driver soft interrupt handler.</p> <p>The function called by <b>timeout()</b> is run in interrupt context and must not sleep or call other functions which may sleep.</p>
<b>RETURN VALUES</b>	<b>timeout()</b> returns an opaque non-zero timeout identifier that can be passed to <b>untimeout(9F)</b> to cancel the request.
<b>CONTEXT</b>	<b>timeout()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using <b>timeout()</b></p> <p>In the following example, the device driver has issued an IO request and is waiting for the device to respond. If the device does not respond within 5 seconds, the device driver will print out an error message to the console.</p> <pre>static void xtimeout_handler(void *arg) {     struct xxstate *xsp = (struct xxstate *)arg;     mutex_enter(&gt;lock);     cv_signal(&gt;cv);     xsp-&gt;flags  = TIMED_OUT;     mutex_exit(&gt;lock);      xsp-&gt;timeout_id = 0;</pre>

```

}
static uint_t
xxintr(caddr_t arg)
{
    struct xxstate *xsp = (struct xxstate *)arg;
    .
    .
    mutex_enter(>lock);

    /* Service interrupt */

    cv_signal(>cv);
    mutex_exit(>lock);

    if (xsp->timeout_id != 0) {
        (void) untimeout(xsp->timeout_id);
        xsp->timeout_id = 0;
    }

    return(DDI_INTR_CLAIMED);
}

static void
xxcheckcond(struct xxstate *xsp)
{
    .
    .
    .
    xsp->timeout_id = timeout(xtimeout_handler,
        xsp, 5 * drv_usectohz(1000000));

    mutex_enter(>lock);
    while (/* Waiting for interrupt or timeout */)
        cv_wait(>cv, >lock);

    if (xsp->flags & TIMED_OUT)
        cmn_err(CE_WARN, "Device not responding");
    .
    .
    mutex_exit(>lock);
    .
    .
}

```

**SEE ALSO** [bufcall\(9F\)](#), [delay\(9F\)](#), [untimeout\(9F\)](#)

*Writing Device Drivers*

<b>NAME</b>	uiomove – copy kernel data using uio structure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/uio.h&gt;  int uiomove(caddr_t address, size_t nbytes, enum uio_rw rflag, uio_t *uio_p);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>address</b>            Source/destination kernel address of the copy.</p> <p><b>nbytes</b>             Number of bytes to copy.</p> <p><b>rflag</b>              Flag indicating read or write operation. Possible values are UIO_READ and UIO_WRITE.</p> <p><b>uio_p</b>              Pointer to the uio structure for the copy.</p>
<b>DESCRIPTION</b>	<p>The <b>uiomove()</b> function copies <i>nbytes</i> of data to or from the space defined by the <b>uio</b> structure (described in <b>uio(9S)</b>) and the driver.</p> <p>The <b>uio_segflg</b> member of the <b>uio(9S)</b> structure determines the type of space to or from which the transfer is being made. If it is set to <b>UIO_SYSSPACE</b>, the data transfer is between addresses in the kernel. If it is set to <b>UIO_USERSPACE</b>, the transfer is between a user program and kernel space.</p> <p><b>rflag</b> indicates the direction of the transfer. If <b>UIO_READ</b> is set, the data will be transferred from <i>address</i> to the buffer(s) described by <i>uio_p</i>. If <b>UIO_WRITE</b> is set, the data will be transferred from the buffer(s) described by <i>uio_p</i> to <i>address</i>.</p> <p>In addition to moving the data, <b>uiomove()</b> adds the number of bytes moved to the <b>iov_base</b> member of the <b>iovec(9S)</b> structure, decreases the <b>iov_len</b> member, increases the <b>uio_offset</b> member of the <b>uio(9S)</b> structure, and decreases the <b>uio_resid</b> member.</p> <p>This function automatically handles page faults. <i>nbytes</i> does not have to be word-aligned.</p>
<b>RETURN VALUES</b>	<b>uiomove()</b> returns 0 upon success or EFAULT on failure.
<b>CONTEXT</b>	User context only, if <b>uio_segflg</b> is set to <b>UIO_USERSPACE</b> . User or interrupt context, if <b>uio_segflg</b> is set to <b>UIO_SYSSPACE</b> .
<b>SEE ALSO</b>	<b>ureadc(9F)</b> , <b>uwritec(9F)</b> , <b>iovec(9S)</b> , <b>uio(9S)</b> <i>Writing Device Drivers</i>

**WARNINGS**

If `uio_segflg` is set to `UIO_SYSSPACE` and *address* is selected from user space, the system may panic.

<b>NAME</b>	unbufcall – cancel a pending bufcall request
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  void unbufcall(bufcall_id_t id);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>id</b> Identifier returned from <b>bufcall(9F)</b> or <b>esbbscall(9F)</b> .
<b>DESCRIPTION</b>	<b>unbufcall</b> cancels a pending <b>bufcall()</b> or <b>esbbscall()</b> request. The argument <b>id</b> is a non-zero identifier for the request to be cancelled. <b>id</b> is returned from the <b>bufcall()</b> or <b>esbbscall()</b> function used to issue the request. <b>unbufcall()</b> will not return until the pending callback is cancelled or has run. Because of this, locks acquired by the callback routine should not be held across the call to <b>unbufcall()</b> or deadlock may result.
<b>RETURN VALUES</b>	None.
<b>CONTEXT</b>	<b>unbufcall()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>bufcall(9F)</b> , <b>esbbscall(9F)</b> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

<b>NAME</b>	unlinkb – remove a message block from the head of a message
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt;  mblk_t *unlinkb(mblk_t *mp);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>mp</b> Pointer to the message.
<b>DESCRIPTION</b>	<b>unlinkb()</b> removes the first message block from the message pointed to by <i>mp</i> . A new message, minus the removed message block, is returned.
<b>RETURN VALUES</b>	If successful, <b>unlinkb()</b> returns a pointer to the message with the first message block removed. If there is only one message block in the message, NULL is returned.
<b>CONTEXT</b>	<b>unlinkb()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1 unlinkb() example</b></p> <p>The routine expects to get passed an M_PROTO T_DATA_IND message. It will remove and free the M_PROTO header and return the remaining M_DATA portion of the message.</p> <pre>1 mblk_t * 2 makedata(mp) 3     mblk_t *mp; 4 { 5 mblk_t *nmp; 6 7 nmp = unlinkb(mp); 8 freeb(mp); 9 return(nmp); 10 }</pre>
<b>SEE ALSO</b>	<p><b>linkb(9F)</b></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

<b>NAME</b>	untimeout – cancel previous timeout function call
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/conf.h&gt;  clock_t untimeout(timeout_id_t id);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<i>id</i> Opaque timeout ID from a previous <code>timeout(9F)</code> call.
<b>DESCRIPTION</b>	<p><b>untimeout()</b> cancels a pending <code>timeout(9F)</code> request. <b>untimeout()</b> will not return until the pending callback is cancelled or has run. Because of this, locks acquired by the callback routine should not be held across the call to <b>untimeout()</b> or a deadlock may result.</p> <p>Since no mutex should be held across the call to <b>untimeout()</b>, there is a race condition between the occurrence of an expected event and the execution of the timeout handler. In particular, it should be noted that no problems will result from calling <b>untimeout()</b> for a timeout which is either running on another CPU, or has already completed. Drivers should be structured with the understanding that the arrival of both an interrupt and a timeout for that interrupt can occasionally occur, in either order.</p>
<b>RETURN VALUES</b>	<b>untimeout()</b> returns <code>-1</code> if the <i>id</i> is not found. Otherwise, it returns an integer value greater than or equal to <code>0</code> .
<b>CONTEXT</b>	<b>untimeout()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b></p> <p>In the following example, the device driver has issued an IO request and is waiting for the device to respond. If the device does not respond within 5 seconds, the device driver will print out an error message to the console.</p> <pre>static void xxtimeout_handler(void *arg) {     struct xxstate *xsp = (struct xxstate *)arg;     mutex_enter(&amp;xsp-&gt;lock);     cv_signal(&amp;xsp-&gt;cv);     xsp-&gt;flags  = TIMED_OUT;     mutex_exit(&amp;xsp-&gt;lock);     xsp-&gt;timeout_id = 0; } static uint_t xxintr(caddr_t arg)</pre>

```

{
    struct xxstate *xsp = (struct xxstate *)arg;
    .
    .
    .
    mutex_enter(&xsp->lock);
    /* Service interrupt */
    cv_signal(&xsp->cv);
    mutex_exit(&xsp->lock);
    if (xsp->timeout_id != 0) {
        (void) untimeout(xsp->timeout_id);
        xsp->timeout_id = 0;
    }
    return(DDI_INTR_CLAIMED);
}
static void
xxcheckcond(struct xxstate *xsp)
{
    .
    .
    .
    xsp->timeout_id = timeout(xtimeout_handler,
        xsp, (5 * drv_usectohz(1000000)));
    mutex_enter(&xsp->lock);
    while (/* Waiting for interrupt or timeout */)
        cv_wait(&xsp->cv, &xsp->lock);
    if (xsp->flags & TIMED_OUT)
        cmn_err(CE_WARN, "Device not responding");
    .
    .
    .
    mutex_exit(&xsp->lock);
    .
    .
    .
}

```

**SEE ALSO** [open\(9E\)](#), [cv\\_signal\(9F\)](#), [cv\\_wait\\_sig\(9F\)](#), [delay\(9F\)](#), [timeout\(9F\)](#)

*Writing Device Drivers*

<b>NAME</b>	ureadc – add character to a uio structure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/uio.h&gt; #include &lt;sys/types.h&gt;  int ureadc(int c, uio_t *uio_p);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<p><b>c</b>        The character added to the <b>uio(9S)</b> structure.</p> <p><b>uio_p</b>    Pointer to the <b>uio(9S)</b> structure.</p>
<b>DESCRIPTION</b>	<b>ureadc()</b> transfers the character <i>c</i> into the address space of the <b>uio(9S)</b> structure pointed to by <i>uio_p</i> , and updates the <b>uio</b> structure as for <b>uimove(9F)</b> .
<b>RETURN VALUES</b>	0 is returned on success and EFAULT on failure.
<b>CONTEXT</b>	<b>ureadc()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>uimove(9F)</b> , <b>uwritec(9F)</b> , <b>iovec(9S)</b> , <b>uio(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	uwritec – remove a character from a uio structure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/uio.h&gt;  int uwritec(uio_t *uio_p);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>uio_p</b> Pointer to the uio(9S) structure.
<b>DESCRIPTION</b>	<b>uwritec()</b> returns a character from the uio structure pointed to by <i>uio_p</i> and updates the uio structure as for <b>uiomove(9F)</b> .
<b>RETURN VALUES</b>	The next character for processing is returned on success, and -1 is returned if uio is empty or there is an error.
<b>CONTEXT</b>	<b>uwritec()</b> can be called from user or interrupt context.
<b>SEE ALSO</b>	<b>uiomove(9F)</b> , <b>ureadc(9F)</b> , <b>iovec(9S)</b> , <b>uio(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	va_arg, va_start, va_copy, va_end – handle variable argument list
<b>SYNOPSIS</b>	<pre>#include &lt;sys/varargs.h&gt;  void va_start(va_list pvar, void parmN);  (  type  *)  va_arg(va_list pvar, type);  void va_copy(va_list dest, va_list src);  void va_end(va_list pvar);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	
<b>va_start()</b>	<p><b>pvar</b>            Pointer to variable argument list.</p> <p><b>name</b>            Identifier of rightmost parameter in the function definition.</p>
<b>va_arg()</b>	<p><b>pvar</b>            Pointer to variable argument list.</p> <p><b>type</b>            Type name of the next argument to be returned.</p>
<b>va_copy()</b>	<p><b>dest</b>            Destination variable argument list.</p> <p><b>src</b>             Source variable argument list.</p>
<b>va_end()</b>	<p><b>pvar</b>            Pointer to variable argument list.</p>
<b>DESCRIPTION</b>	<p>This set of macros allows portable procedures that accept variable argument lists to be written. Routines that have variable argument lists but do not use the <b>varargs()</b> macros are inherently non-portable, as different machines use different argument-passing conventions. Routines that accept a variable argument list can use these macros to traverse the list.</p>

`va_list` is the type defined for the variable used to traverse the list of arguments.

`va_start()` is called to initialize `pvar` to the beginning of the variable argument list. `va_start()` must be invoked before any access to the unnamed arguments. The parameter `name` is the identifier of the rightmost parameter in the variable parameter list in the function definition (the one just before the “ , . . . ”). If this parameter is declared with the `register` storage class or with a function or array type, or with a type that is not compatible with the type that results after application of the default argument promotions, the behavior is undefined.

`va_arg()` expands to an expression that has the type and value of the next argument in the call. The parameter `pvar` must be initialized by `va_start()`. Each invocation of `va_arg()` modifies `pvar` so that the values of successive arguments are returned in turn. The parameter `type` is the type name of the next argument to be returned. The type name must be specified in such a way so that the type of a pointer to an object that has the specified type can be obtained simply by postfixing a `*` to `type`. If there is no actual next argument, or if `type` is not compatible with the type of the actual next argument (as promoted according to the default argument promotions), the behavior is undefined.

The `va_copy()` macro saves the state represented by the `va_list src` in the `va_list dest`. The `va_list` passed as `dest` should not be initialized by a previous call to `va_start()`, and must be passed to `va_end()` before being reused as a parameter to `va_start()` or as the `dest` parameter of a subsequent call to `va_copy()`. The behavior is undefined should any of these restrictions not be met.

The `va_end()` macro is used to clean up. It invalidates `pvar` for use (unless `va_start()` is invoked again).

Multiple traversals, each bracketed by a call to `va_start()` and `va_end()`, are possible.

## EXAMPLES

### EXAMPLE 1 Creating a Variable Length Command

The following example uses these routines to create a variable length command. This may be useful for a device which provides for a variable length command set. `ncmdbytes` is the number of bytes in the command. The new command is written to `cmdp`.

```
static void
xx_write_cmd(uchar_t *cmdp, int ncmdbytes, ...)
{
    va_list\01lap;
    int\01li;

\011    /*
```

```
\011      * Write variable-length command to destination
\011      */
\011      va_start(ap, ncmdbytes);
\011      for (i = 0; i < ncmdbytes; i++) {
\011          *cmdp++ = va_arg(ap, uchar_t);
\011      }
\011      va_end(ap);
\011  }
```

**SEE ALSO** [vcmn\\_err\(9F\)](#), [vsprintf\(9F\)](#)

**NOTES**

It is up to the calling routine to specify in some manner how many arguments there are, since it is not always possible to determine the number of arguments from the stack frame.

It is non-portable to specify a second argument of `char` or `short` to `va_arg`, because arguments seen by the called function are not `char` or `short`. C converts `char` and `short` arguments to `int` before passing them to a function.

<b>NAME</b>	vsprintf – format characters in memory
<b>SYNOPSIS</b>	<pre>#include &lt;sys/varargs.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;</pre> <p>char *<b>vsprintf</b>(char *<i>buf</i>, const char *<i>fmt</i>, va_list <i>ap</i>);</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>PARAMETERS</b>	<p><b><i>buf</i></b>            Pointer to a character string.</p> <p><b><i>fmt</i></b>            Pointer to a character string.</p> <p><b><i>ap</i></b>             Pointer to a variable argument list.</p>
<b>DESCRIPTION</b>	<p><b>vsprintf()</b> builds a string in <i>buf</i> under the control of the format <i>fmt</i>. The format is a character string with either plain characters, which are simply copied into <i>buf</i>, or conversion specifications, each of which converts zero or more arguments, again copied into <i>buf</i>. The results are unpredictable if there are insufficient arguments for the format; excess arguments are simply ignored. It is the user's responsibility to ensure that enough storage is available for <i>buf</i>.</p> <p><i>ap</i> contains the list of arguments used by the conversion specifications in <i>fmt</i>. <i>ap</i> is a variable argument list and must be initialized by calling <b>va_start(9F)</b>. <b>va_end(9F)</b> is used to clean up and must be called after each traversal of the list. Multiple traversals of the argument list, each bracketed by <b>va_start(9F)</b> and <b>va_end(9F)</b>, are possible.</p> <p>Each conversion specification is introduced by the % character, after which the following appear in sequence:</p> <p>An optional decimal digit specifying a minimum field width for numeric conversion. The converted value will be right-justified and padded with leading zeroes if it has fewer characters than the minimum.</p> <p>An optional l (ll) specifying that a following d, D, o, O, x, X, or u conversion character applies to a long (long long) integer argument. An l (ll) before any other conversion character is ignored.</p> <p>A character indicating the type of conversion to be applied: d,D,o,O,x,X,u</p>

The integer argument is converted to signed decimal (d, D), unsigned octal (o, O), unsigned hexadecimal (x, X) or unsigned decimal (u), respectively, and copied. The letters abcdef are used for x and X conversion.

c

The character value of the argument is copied.

b

This conversion uses two additional arguments. The first is an integer, and is converted according to the base specified in the second argument. The second argument is a character string in the form *<base>* [*<arg>* . . .]. The base supplies the conversion base for the first argument as a binary value; \10 gives octal, \20 gives hexadecimal. Each subsequent *<arg>* is a sequence of characters, the first of which is the bit number to be tested, and subsequent characters, up to the next bit number or terminating null, supply the name of the bit.

A bit number is a binary-valued character in the range 1-32. For each bit set in the first argument, and named in the second argument, the bit names are copied, separated by commas, and bracketed by < and >. Thus, the following function call would generate *reg=3<BitTwo,BitOne>\n* in *buf*.

```
vsprintf(buf, "reg=%b\n", 3, "\10\2BitTwo\1BitOne")
```

s

The argument is taken to be a string (character pointer), and characters from the string are copied until a null character is encountered. If the character pointer is NULL on SPARC, the string *<nullstring>* is used in its place; on x86, it is undefined.

%

Copy a %; no argument is converted.

**RETURN VALUES**

**vsprintf()** returns its first parameter, *buf*.

**CONTEXT**

**vsprintf()** can be called from user, kernel, or interrupt context.

**EXAMPLES****EXAMPLE 1** Using **vsprintf()**

In this example, **xxerror()** accepts a pointer to a `dev_info_t` structure *dip*, an error level *level*, a format *fmt*, and a variable number of arguments. The routine uses **vsprintf()** to format the error message in *buf*. Note that **va\_start(9F)** and **va\_end(9F)** bracket the call to **vsprintf()**. *instance*, *level*, *name*, and *buf* are then passed to **cmn\_err(9F)**.

```
#include <sys/varargs.h>
#include <sys/ddi.h>
#include <sys/sunddi.h>
#define MAX_MSG 256

void
xxerror(dev_info_t *dip, int level, const char *fmt, ...)
{
    va_list ap;
    int instance;
    char buf[MAX_MSG],
        *name;

    instance = ddi_get_instance(dip);
    name = ddi_binding_name(dip);

    /* format buf using fmt and arguments contained in ap */
    va_start(ap, fmt);
    vsprintf(buf, fmt, ap);
    va_end(ap);

    /* pass formatted string to cmn_err(9F) */
    cmn_err(level, "%s%d: %s", name, instance, buf);
}
```

**SEE ALSO**

**cmn\_err(9F)**, **ddi\_binding\_name(9F)**, **ddi\_get\_instance(9F)**, **va\_arg(9F)**

*Writing Device Drivers*

<b>NAME</b>	WR, wr – get pointer to the write queue for this module or driver
<b>SYNOPSIS</b>	<pre>#include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt;  queue_t * WR(queue_t * q);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI).
<b>PARAMETERS</b>	<b>q</b> Pointer to the <i>read</i> queue whose <i>write</i> queue is to be returned.
<b>DESCRIPTION</b>	<p>The <b>WR()</b> function accepts a <i>read</i> queue pointer as an argument and returns a pointer to the <i>write</i> queue of the same module.</p> <p>CAUTION: Make sure the argument to this function is a pointer to a <i>read</i> queue. <b>WR()</b> will not check for queue type, and a system panic could result if the pointer is not to a <i>read</i> queue.</p>
<b>RETURN VALUES</b>	The pointer to the <i>write</i> queue.
<b>CONTEXT</b>	<b>WR()</b> can be called from user or interrupt context.
<b>EXAMPLES</b>	<p><b>EXAMPLE 1 Using WR()</b></p> <p>In a STREAMS <code>close(9E)</code> routine, the driver or module is passed a pointer to the <i>read</i> queue. These usually are set to the address of the module-specific data structure for the minor device.</p> <pre>1 xxxclose(q, flag) 2   queue_t *q; 3   int flag; 4   { 5\011     q-&gt;q_ptr = NULL; 6         WR(q)-&gt;q_ptr = NULL; 7\011     . . . 7   }</pre>
<b>SEE ALSO</b>	<p><code>close(9E)</code> , <code>OTHERQ(9F)</code> , <code>RD(9F)</code></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>



# Index

---

## A

- activate a new DMA window —
  - ddi\_dma\_getwin, 315
- add a fully initialized kstat to the system —
  - kstat\_install, 499
- add a soft interrupt
  - ddi\_add\_softintr, 260
- add an interrupt handler
  - ddi\_add\_intr, 257
- address
  - return mapped virtual address —
    - csx\_GetMappedAddr, 128
- adjmsg — trim bytes from a message, 52
- allocate and free a scsi\_pkt structure —
  - scsi\_hba\_pkt\_alloc, 604, 606
  - scsi\_hba\_pkt\_free, 604
  - scsi\_hba\_tran\_free, 606
- allocate and free non-sequentially accessed memory
  - ddi\_iopb\_alloc, 356
  - ddi\_iopb\_free, 356
- allocate DMA handle —
  - ddi\_dma\_alloc\_handle, 296
- allocate kernel memory
  - ddi\_umem\_alloc, 436
  - ddi\_umem\_free, 436
  - ddi\_umem\_zalloc, 436
  - kmem\_alloc, 493
  - kmem\_free, 493

- kmem\_zalloc, 493
- allocate memory for DMA transfer —
  - ddi\_dma\_mem\_alloc, 319
- allocate space — rmalloc, 570
- allocate space from a resource map —
  - rmalloc\_wait, 575
- allow 64 bit transfers on SBus —
  - ddi\_dma\_set\_sbus64, 333
- anocancel — prevent cancellation of asynchronous I/O request, 56
- aphysio — perform asynchronous physical I/O, 57
- assert — expression verification, 59
- asynchronous physical I/O — aphysio, 57
- asynchronous STREAMS perimeter upgrade
  - qwriter, 567

## B

- bcopy — copy data between address locations in kernel, 63
- binds a system buffer to a DMA handle —
  - ddi\_dma\_buf\_bind\_handle, 299
- binds an address to a DMA handle —
  - ddi\_dma\_addr\_bind\_handle, 291
- bioclone — clone another buffer, 65
- bioerror — indicate error in buffer header, 70
- biofini — uninitialized a buffer structure, 71
- bioinit — initialize a buffer structure, 72

- biomodified — check if a buffer is modified, 73
- bioreset — reuse a private buffer header after I/O is complete, 74
- biosize — returns size of a buffer structure, 75
- bufcall — call a function when a buffer becomes available, 81, 662
  - call a function when a buffer becomes available, 81
- buffer header
  - indicate error — bioerror, 70
  - reuse a private buffer header after I/O is complete — bioreset, 74
- busy-wait for specified interval — drv\_usecwait, 464
- byte streams
  - compare two — bcmp, 62
- bytes, size
  - convert size in pages — ptob, 541
  - convert to size in memory pages (round down) — btop, 79
  - convert to size in memory pages (round up) — btopr, 80

## C

- call a function when a buffer becomes available — qbufcall, 554
  - bufcall, 81
- call a STREAMS put procedure — put, 544
- cancel a pending qbufcall request — qunbufcall, 563
- cancel previous timeout function call — quntimeout, 564
- cancellation of asynchronous I/O — anocancel, 56
- character strings
  - compare two null terminated strings — strcmp, strncmp, 641
  - convert between an integer and a decimal string — stoi, numtos, 639
  - copy a string from one location to another — strcpy, strncpy, 642
  - determine the number of non-null bytes in a string — strlen, 643
  - find a character in a string — strchr, 640
  - format in memory — sprintf, 637
- check for an available buffer — testb, 656
- check for the existence of a property — ddi\_prop\_exists, 398
- check if a buffer is modified — biomodified, 73
- CIS tuple
  - first tuple — csx\_GetFirstTuple, 125
  - next tuple — csx\_GetNextTuple, 125
- clear client event mask — csx\_ReleaseSocketMask, 240
- client
  - register client — csx\_RegisterClient, 216
- client event mask
  - return client event mask — csx\_GetEventMask, 249
  - set client event mask — csx\_SetEventMask, 249
- client return
  - csx\_GetFirstClient, 123
  - csx\_GetNextClient, 123
- clone another buffer — bioclone, 65
- condition variable routines, driver
  - condvar, 94
  - cv\_broadcast, 94
  - cv\_init, 94
  - cv\_signal, 94
  - cv\_timedwait, 94
  - cv\_timedwait\_sig, 94
  - cv\_wait, 94
  - cv\_wait\_sig, 94
- configure PC Card and socket — csx\_RequestConfiguration, 226
- control driver notification of user accesses — ddi\_mapdev\_intercept, 367
  - ddi\_mapdev\_nointercept, 367
- control device components' availability for power management
  - pm\_busy\_component, 533
  - pm\_idle\_component, 533
- control the validation of memory address translations
  - devmap\_load, 454
  - devmap\_unload, 454
- convert a DMA segment to a DMA address cookie — ddi\_dma\_segtocookie, 331
- convert clock ticks to microseconds — drv\_hztousec, 461
- convert device sizes — csx\_ConvertSize, 107

convert device speeds —  
     csx\_ConvertSpeed, 109  
 convert error return codes to text strings —  
     csx\_Error2Text, 117  
 convert events to text strings —  
     csx\_Event2Text, 119  
 convert microseconds to clock ticks —  
     drv\_usectohz, 463  
 copy data from one device register to another  
 device register —  
     ddi\_device\_copy, 276  
 create minor nodes for client —  
     csx\_MakeDeviceNode, 137  
 create a minor node for this device —  
     ddi\_create\_minor\_node, 275  
 create and initialize a new kstat —  
     kstat\_create, 496  
 create driver-controlled mapping of device —  
     ddi\_mapdev, 365  
 csx\_AccessConfigurationRegister — read or  
 write a PC Card  
     Configuration Register, 105  
 csx\_ConvertSize — convert device sizes, 107  
 csx\_ConvertSpeed — convert device  
 speeds, 109  
 csx\_CS\_DDI\_Info — obtain DDI  
 information, 111  
 csx\_DeregisterClient — remove client from  
 Card Services list, 113  
 csx\_DupHandle — duplicate access  
 handle, 114  
 csx\_Error2Text — convert error return codes to  
 text strings, 117  
 csx\_Event2Text — convert events to text  
 strings, 119  
 csx\_FreeHandle — free access handle, 121  
 csx\_Get16 — read from device register, 122  
 csx\_Get32 — read from device register, 122  
 csx\_Get64 — read from device register, 122  
 csx\_Get8 — read from device register, 122  
 csx\_GetEventMask — return client event  
 mask, 249  
 csx\_GetFirstClient — return first client, 123  
 csx\_GetFirstTuple — return first CIS tuple, 125  
 csx\_GetHandleOffset — return current access  
 handle offset, 127  
 csx\_GetMappedAddr — return mapped  
 virtual address, 128  
 csx\_GetNextClient — return next client, 123  
 csx\_GetNextTuple — return next CIS  
 tuple, 125  
 csx\_GetStatus — return status of PC Card and  
 socket, 129  
 csx\_GetTupleData — return data portion of  
 tuple, 134  
 csx\_MakeDeviceNode — create minor nodes  
 for client, 137  
 csx\_MapLogSocket — return physical socket  
 number, 140  
 csx\_MapMemPage — map memory area on  
 PC Card, 142  
 csx\_ModifyConfiguration — modify PC Card  
 configuration, 144  
 csx\_ModifyWindow — modify window  
 attributes, 147  
 csx\_ParseTuple — generic tuple parser, 213  
 csx\_Parse\_CISTPL\_BATTERY — parse Battery  
 Replacement Date tuple, 150  
 csx\_Parse\_CISTPL\_BYTEORDER — parse Byte  
 Order tuple, 152  
 csx\_Parse\_CISTPL\_CFTABLE\_ENTRY — parse  
 Card Configuration Table  
 tuple, 154  
 csx\_Parse\_CISTPL\_CONFIG — parse  
 Configuration tuple, 162  
 csx\_Parse\_CISTPL\_DATE — parse Card  
 Initialization Date tuple, 165  
 csx\_Parse\_CISTPL\_DEVICE — parse Device  
 Information tuple for  
 Common Memory, 167  
 csx\_Parse\_CISTPL\_DEVICEGEO — parse  
 Device Geo tuple, 171  
 csx\_Parse\_CISTPL\_DEVICEGEO\_A — parse  
 Device Geo A tuple, 173  
 csx\_Parse\_CISTPL\_DEVICE\_A — parse Device  
 Information tuple for  
 Attribute Memory, 167  
 csx\_Parse\_CISTPL\_DEVICE\_OA — parse  
 Other Condition Device  
 Information tuple for  
 Attribute Memory, 167  
 csx\_Parse\_CISTPL\_DEVICE\_OC — parse  
 Other Condition Device  
 Information tuple for  
 Common Memory, 167

csx\_Parse\_CISTPL\_FORMAT — parse Data Recording Format tuple, 175  
 csx\_Parse\_CISTPL\_FUNCE — parse Function Extension tuple, 178  
 csx\_Parse\_CISTPL\_FUNCID — parse Function Identification tuple, 188  
 csx\_Parse\_CISTPL\_GEOMETRY — parse Geometry tuple, 191  
 csx\_Parse\_CISTPL\_JEDEC\_A — parse JEDEC Identifier tuple for Attribute Memory, 193  
 csx\_Parse\_CISTPL\_JEDEC\_C — parse JEDEC Identifier tuple for Common Memory, 193  
 csx\_Parse\_CISTPL\_LINKTARGET — parse Link Target tuple, 195  
 csx\_Parse\_CISTPL\_LONGLINK\_A — parse Long Link A tuple, 197  
 csx\_Parse\_CISTPL\_LONGLINK\_C — parse Long Link C tuple, 197  
 csx\_Parse\_CISTPL\_LONGLINK\_MFC — parse Multi-Function tuple, 199  
 csx\_Parse\_CISTPL\_MANFID — parse Manufacturer Identification tuple, 201  
 csx\_Parse\_CISTPL\_ORG — parse Data Organization tuple, 203  
 csx\_Parse\_CISTPL\_SPCL — parse Special Purpose tuple, 205  
 csx\_Parse\_CISTPL\_SWIL — parse Software Interleaving tuple, 207  
 csx\_Parse\_CISTPL\_VERS\_1 — parse Level-1 Version/Product Information tuple, 209  
 csx\_Parse\_CISTPL\_VERS\_2 — parse Level-2 Version and Information tuple, 211  
 csx\_Put16 — write to device register, 215  
 csx\_Put32 — write to device register, 215  
 csx\_Put64 — write to device register, 215  
 csx\_Put8 — write to device register, 215  
 csx\_RegisterClient — register client, 216  
 csx\_ReleaseConfiguration — release configuration on PC Card, 220  
 csx\_ReleaseIO — release I/O resources, 231  
 csx\_ReleaseIRQ — release IRQ resource, 237  
 csx\_ReleaseSocketMask — clear client event mask, 240  
 csx\_ReleaseWindow — release window resources, 242  
 csx\_RepGet16 — read repetitively from device register, 222  
 csx\_RepGet32 — read repetitively from device register, 222  
 csx\_RepGet64 — read repetitively from device register, 222  
 csx\_RepGet8 — read repetitively from device register, 222  
 csx\_RepPut16 — write repetitively to device register, 224  
 csx\_RepPut32 — write repetitively to device register, 224  
 csx\_RepPut64 — write repetitively to device register, 224  
 csx\_RepPut8 — write repetitively to device register, 224  
 csx\_RequestConfiguration — configure PC Card and socket, 226  
 csx\_RequestIO — request I/O resources, 231  
 csx\_RequestIRQ — request IRQ resource, 237  
 csx\_RequestSocketMask — request client event mask, 240  
 csx\_RequestWindow — request window resources, 242  
 csx\_ResetFunction — reset a function on a PC card, 247  
 csx\_SetEventMask — set client event mask, 249  
 csx\_SetHandleOffset — set current access handle offset, 251  
 csx\_ValidateCIS — validate Card Information Structure (CIS), 252

## D

datamsg — test whether a message is a data message, 254  
 DDI access credential structure  
   — ddi\_get\_cred, 344  
 DDI announce a device  
   — ddi\_report\_dev, 423  
 DDI device access  
   slave access only — ddi\_slaveonly, 430

- DDI device critical region of control
  - enter — `ddi_enter_critical`, 340
  - exit — `ddi_exit_critical`, 340
- DDI device information structure
  - find parent — `ddi_get_parent`, 348
  - get the root of the `dev_info` tree — `ddi_root_node`, 426
  - remove a minor node for this `devinfo` — `ddi_remove_minor_node`, 420
- DDI device instance number
  - get — `ddi_get_instance`, 346
- DDI device mapping
  - `ddi_mapdev` — create driver-controlled mapping of device, 365
  - `ddi_mapdev_intercept` — control driver notification of user accesses, 367
  - `ddi_mapdev_nointercept` — control driver notification of user accesses, 367
  - `devmap_default_access` — device mapping access entry point, 440
- DDI device registers
  - map — `ddi_map_regs`, 371
  - return the number of register sets — `ddi_dev_nregs`, 288
  - return the size — `ddi_dev_regsize`, 289
  - unmap — `ddi_unmap_regs`, 371
- DDI device virtual address
  - read 16 bit — `ddi_peek16`, 388
  - write 16 bit — `ddi_poke16`, 391
- DDI device's private data area
  - get the address — `ddi_get_driver_private`, 345
  - set the address — `ddi_set_driver_private`, 345
- DDI `devinfo` node name
  - return — `ddi_binding_name`, 266, 387
- DDI direct memory access
  - convert DMA handle to DMA addressing cookie — `ddi_dma_htoc`, 317
- DDI direct memory access services
  - allocate consistent memory— `ddi_iopb_alloc`, 373
  - convert a DMA cookie — `ddi_dma_coff`, 305
  - easier DMA setup — `ddi_dma_addr_setup`, 294, 302
  - find minimum alignment and transfer size for device — `ddi_iomin`, 355
  - find post DMA mapping alignment and minimum effect properties — `ddi_dma_devalign`, 307
  - free consistent memory — `ddi_iopb_free`, 373
  - report current DMA window offset and size — `ddi_dma_curwin`, 306
  - setup DMA mapping — `ddi_dma_setup`, 326, 328, 331
  - setup DMA resources — `ddi_dma_setup`, 335
  - shift current DMA window — `ddi_dma_movwin`, 322
  - tear down DMA mapping — `ddi_dma_free`, 313
- DDI information — `csx_CS_DDI_Info`, 111
- DDI interrupt handling
  - add an interrupt — `ddi_add_intr`, 257
  - get interrupt block cookie — `ddi_get_iblock_cookie`, 257
  - indicate interrupt handler type — `ddi_intr_hilevel`, 352
  - remove an interrupt — `ddi_remove_intr`, 257
  - return the number of interrupt specifications — `ddi_dev_nintrs`, 287
- DDI memory mapping
  - map a segment — `ddi_segmap`, 428, 452
- DDI page size conversions
  - `ddi_btop`, 267
  - `ddi_btopr`, 267
  - `ddi_ptob`, 267
- DDI property management
  - create properties for leaf device drivers — `ddi_prop_create`, 394
  - `ddi_getlongprop`, 408
  - `ddi_getlongprop_buf`, 408
  - `ddi_getprop`, 408
  - `ddi_getpropflen`, 408
  - `ddi_prop_op`, 408

- modify properties for leaf device drivers
    - `ddi_prop_modify`, 394
  - remove all properties for leaf device drivers
    - `ddi_prop_remove_all`, 394
- DDI self identifying devices
  - tell whether a device is self-identifying
    - `ddi_dev_is_sid`, 285
- DDI soft interrupt handling
  - add a soft interrupt
    - `ddi_add_softintr`, 260
  - get soft interrupt block cookie
    - `ddi_get_soft_iblock_cookie`, 260
  - remove a soft interrupt
    - `ddi_remove_softintr`, 260
- DDI soft state utility routines
  - allocate state structure
    - `ddi_soft_state_zalloc`, 431
  - free soft state entry
    - `ddi_soft_state_free`, 431
  - get pointer to soft state
    - `ddi_get_soft_state`, 431
  - initialize state
    - `ddi_soft_state_init`, 431
  - remove all state info
    - `ddi_soft_state_fini`, 431
- `ddi_add_intr` — add an interrupt handler, 257
- `ddi_add_softintr` — add a soft interrupt, 260
- `ddi_binding_name` — return driver binding name, 266
- `ddi_create_minor_node` — create a minor node for this device, 275
- `ddi_device_copy` — copy data from one device register to another device register, 276
- `ddi_device_zero` — zero fill the device register, 278
- `ddi_devid_compare` — Kernel interfaces for device ids, 281
- `ddi_devid_free` — Kernel interfaces for device ids, 281
- `ddi_devid_init` — Kernel interfaces for device ids, 281
- `ddi_devid_register` — Kernel interfaces for device ids, 281
- `ddi_devid_sizeof` — Kernel interfaces for device ids, 281
- `ddi_devid_unregister` — Kernel interfaces for device ids, 281
- `ddi_devid_valid` — Kernel interfaces for device ids, 281
- `ddi_dev_is_needed` — inform the system that a device's component is required, 283
- `ddi_dmae` — system DMA engine functions, 309
- `ddi_dmae_1stparty` — system DMA engine functions, 309
- `ddi_dmae_alloc` — system DMA engine functions, 309
- `ddi_dmae_disable` — system DMA engine functions, 309
- `ddi_dmae_enable` — system DMA engine functions, 309
- `ddi_dmae_getattr` — system DMA engine functions, 309
- `ddi_dmae_getcnt` — system DMA engine functions, 309
- `ddi_dmae_getlim` — system DMA engine functions, 309
- `ddi_dmae_prog` — system DMA engine functions, 309
- `ddi_dmae_release` — system DMA engine functions, 309
- `ddi_dmae_stop` — system DMA engine functions, 309
- `ddi_dma_addr_bind_handle` — binds an address to a DMA handle, 291
- `ddi_dma_alloc_handle` — allocate DMA handle, 296
- `ddi_dma_buf_bind_handle` — binds a system buffer to a DMA handle, 299
- `ddi_dma_burstsizes` — find out the allowed burst sizes for a DMA mapping, 304
- `ddi_dma_free_handle` — free DMA handle, 314
- `ddi_dma_getwin` — activate a new DMA window, 315
- `ddi_dma_mem_alloc` — allocate memory for DMA transfer, 319
- `ddi_dma_mem_free` — free previously allocated memory, 321
- `ddi_dma_nextcookie` — retrieve subsequent DMA cookie, 324

**ddi\_dma\_nextseg** — get next DMA segment, 326  
**ddi\_dma\_nextwin** — get next DMA window, 328  
**ddi\_dma\_numwin** — retrieve number of DMA windows, 330  
**ddi\_dma\_segtocookie** — convert a DMA segment to a DMA address cookie, 331  
**ddi\_dma\_set\_sbus64** — allow 64 bit transfers on SBus, 333  
**ddi\_dma\_sync** — synchronize CPU and I/O views of memory, 337  
**ddi\_dma\_unbind\_handle** — unbinds the address in a DMA handle, 339  
**ddi\_ffs** — find first (last) bit set in a long integer, 341  
**ddi\_fls** — find first (last) bit set in a long integer, 341  
**ddi\_get16** — read data from the device, 342  
**ddi\_get32** — read data from the device, 342  
**ddi\_get64** — read data from the device, 342  
**ddi\_get8** — read data from the device, 342  
**ddi\_get\_iblock\_cookie** — get interrupt block cookie, 257  
**ddi\_get\_lbolt**  
     returns the value of lbolt, 347  
**ddi\_get\_name** — return driver binding name, 266  
**ddi\_get\_pid**  
     returns the process ID, 349  
**ddi\_get\_soft\_iblock\_cookie** — get soft interrupt block cookie, 260  
**ddi\_get\_time**  
     returns the current time in seconds, 350  
**ddi\_in\_panic** — determine if system is in panic state, 351  
**ddi\_iopb\_alloc** — allocate and free non-sequentially accessed memory, 356  
**ddi\_iopb\_free** — allocate and free non-sequentially accessed memory, 356  
**ddi\_io\_get16** — read data from the mapped device register in I/O space, 353  
**ddi\_io\_get32** — read data from the mapped device register in I/O space, 353  
**ddi\_io\_get8** — read data from the mapped device register in I/O space, 353  
**ddi\_io\_getb** — read data from the mapped device register in I/O space, 353  
**ddi\_io\_getl** — read data from the mapped device register in I/O space, 353  
**ddi\_io\_getw** — read data from the mapped device register in I/O space, 353  
**ddi\_io\_put16** — write data to the mapped device register in I/O space, 358  
**ddi\_io\_put32** — write data to the mapped device register in I/O space, 358  
**ddi\_io\_put8** — write data to the mapped device register in I/O space, 358  
**ddi\_io\_putb** — write data to the mapped device register in I/O space, 358  
**ddi\_io\_putl** — write data to the mapped device register in I/O space, 358  
**ddi\_io\_putw** — write data to the mapped device register in I/O space, 358  
**ddi\_io\_rep\_get16** — read multiple data from the mapped device register in I/O space, 360  
**ddi\_io\_rep\_get32** — read multiple data from the mapped device register in I/O space, 360  
**ddi\_io\_rep\_get8** — read multiple data from the mapped device register in I/O space, 360  
**ddi\_io\_rep\_getb** — read multiple data from the mapped device register in I/O space, 360

`ddi_io_rep_getl` — read multiple data from the mapped device register in I/O space, 360  
`ddi_io_rep_getw` — read multiple data from the mapped device register in I/O space, 360  
`ddi_io_rep_put16` — write multiple data to the mapped device register in I/O space, 362  
`ddi_io_rep_put32` — write multiple data to the mapped device register in I/O space, 362  
`ddi_io_rep_put8` — write multiple data to the mapped device register in I/O space, 362  
`ddi_io_rep_putb` — write multiple data to the mapped device register in I/O space, 362  
`ddi_io_rep_putl` — write multiple data to the mapped device register in I/O space, 362  
`ddi_io_rep_putw` — write multiple data to the mapped device register in I/O space, 362  
`ddi_mapdev` — create driver-controlled mapping of device, 365  
`ddi_mapdev_intercept` — control driver notification of user accesses, 367  
`ddi_mapdev_intercept` — control driver notification of user accesses, 367  
`ddi_mapdev_set_device_acc_attr` — Set the device attributes for the mapping, 369  
`ddi_mem_get16` — read data from mapped device in the memory space or allocated DMA memory, 375  
`ddi_mem_get32` — read data from mapped device in the memory space or allocated DMA memory, 375  
`ddi_mem_get64` — read data from mapped device in the memory space or allocated DMA memory, 375  
`ddi_mem_put16` — write data to mapped device in the memory space or allocated DMA memory, 377  
`ddi_mem_put32` — write data to mapped device in the memory space or allocated DMA memory, 377  
`ddi_mem_put64` — write data to mapped device in the memory space or allocated DMA memory, 377  
`ddi_mem_rep_get16` — read data from mapped device in the memory space or allocated DMA memory, 379  
`ddi_mem_rep_get32` — read data from mapped device in the memory space or allocated DMA memory, 379  
`ddi_mem_rep_get64` — read data from mapped device in the memory space or allocated DMA memory, 379  
`ddi_mem_rep_get8` — read data from mapped device in the memory space or allocated DMA memory, 379  
`ddi_mem_rep_put16` — write data to mapped device in the memory space or allocated DMA memory, 381  
`ddi_mem_rep_put32` — write data to mapped device in the memory space or allocated DMA memory, 381  
`ddi_mem_rep_put64` — write data to mapped device in the memory space or allocated DMA memory, 381  
`ddi_mem_rep_put8` — write data to mapped device in the memory space or allocated DMA memory, 381  
`ddi_mmap_get_model` — return data model type of current thread, 383

`ddi_model_convert_from` — Determine if there is a need to translate shared data structure contents, 385  
`ddi_node_name` — return the devinfo node name, 387  
`ddi_prop_exists` — check for the existence of a property, 398  
`ddi_prop_get_int` — look up integer property, 400  
`ddi_prop_lookup` — lookup property information, 403  
`ddi_prop_lookup_byte_array` — lookup property information, 403  
`ddi_prop_lookup_int_array` — lookup property information, 403  
`ddi_prop_lookup_string` — lookup property information, 403  
`ddi_prop_lookup_string_array` — lookup property information, 403  
`ddi_prop_update` — update property information., 412  
`ddi_prop_update_byte_array` — update property information., 412  
`ddi_prop_update_int` — update property information., 412  
`ddi_prop_update_int_array` — update property information., 412  
`ddi_prop_update_string` — update property information., 412  
`ddi_prop_update_string_array` — update property information., 412  
`ddi_put16` — write data to the device, 415  
`ddi_put32` — write data to the device, 415  
`ddi_put64` — write data to the device, 415  
`ddi_put8` — write data to the device, 415  
`ddi_regs_map_free` — free mapped register address space, 417  
`ddi_regs_map_setup` — set up a mapping for a register address space, 418  
`ddi_remove_intr` — remove an interrupt handler, 257  
`ddi_remove_softintr` — remove a soft interrupt, 260  
`ddi_rep_get16` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_get32` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_get64` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_get8` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_getb` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_getl` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_getll` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_getw` — read data from the mapped memory address, device register or allocated DMA memory address, 421  
`ddi_rep_put16` — write data to the mapped memory address, device register or allocated DMA memory address, 424  
`ddi_rep_put32` — write data to the mapped memory address, device register or allocated DMA memory address, 424  
`ddi_rep_put64` — write data to the mapped memory address, device register or allocated DMA memory address, 424  
`ddi_rep_put8` — write data to the mapped memory address, device register or allocated DMA memory address, 424  
`ddi_rep_putb` — write data to the mapped memory address, device

- register or allocated DMA memory address, 424
- ddi\_rep\_putl — write data to the mapped memory address, device register or allocated DMA memory address, 424
- ddi\_rep\_putll — write data to the mapped memory address, device register or allocated DMA memory address, 424
- ddi\_rep\_putw — write data to the mapped memory address, device register or allocated DMA memory address, 424
- ddi\_trigger\_softintr — trigger a soft interrupt, 260
- ddi\_umem\_alloc — allocate kernel memory, 436
- ddi\_umem\_free — allocate kernel memory, 436
- ddi\_umem\_zalloc — allocate kernel memory, 436
- default SCSI HBA probe function — scsi\_hba\_probe, 605
- delay — delay process execution for a specified number of clock ticks, 438
- deregister client from Card Services list — csx\_DeregisterClient, 113
- determine data model type mismatch — ddi\_model\_convert\_from, 385
- Device Driver Interface, *see* DDI,
- device mapping access entry point — devmap\_default\_access, 440
- device switch tables
  - return function for insignificant entries — nulldev, 523
- devices
  - get major device number — getmajor, 481
  - get minor device number — getminor, 482
  - make device number from major and minor numbers — makedevice, 506
- devices, non-pollable
  - error return function — nochpoll, 520
- devmap\_default\_access — device mapping access entry point, 440
- devmap\_devmem\_setup — Set driver memory mapping parameters, 444
  - devmap\_devmem\_setup(), 443
  - devmap\_umem\_setup(), 444
- devmap\_do\_ctxmgt — perform device context switching on a mapping, 447
- devmap\_load — control the validation of memory address translations, 454
- devmap\_set\_ctx\_timeout — set context management timeout value, 450
- devmap\_umem\_setup — Set driver memory mapping parameters, 444
- devmap\_unload — control the validation of memory address translations, 454
- disksort — single direction elevator seek sort for buffers, 457
- display a SCSI request sense message
  - scsi\_vu\_errmsg, 632
- DMA mapping, the allowed burst sizes for — ddi\_dma\_burstsizes, 304
- driver buffers
  - copy data— ddi\_copyin, 268
  - copy data from driver — ddi\_copyout, 271
  - copy data from driver to user program — copyout, 103
  - copy data from user program — copyin, 99
- driver error messages
  - display an error message or panic the system — cmn\_err, 88
- driver privilege — drv\_priv, 462
- drv\_getparm — retrieve kernel state information, 459
- drv\_hztousec — convert clock ticks to microseconds, 461
- drv\_priv — determine driver privilege, 462
- drv\_usecsthz — convert microseconds to clock ticks, 463
- drv\_usecwait — busy-wait for specified interval, 464
- dupb — duplicate a message block descriptor, 465
- duplicate a message — dupmsg, 468
- duplicate a message block descriptor — dupb, 465
- duplicate access handle — csx\_DupHandle, 114

dupmsg — duplicate a message, 468

## E

enable/disable accesses to the PCI Local Bus  
Configuration space.

— pci\_config\_setup, 530

— pci\_config\_takedown, 530

error return codes converted to text strings —  
csx\_Error2Text, 117

error return function for illegal entries —  
nodev, 521

event mask

return client event mask —  
csx\_GetEventMask, 249

set client event mask —  
csx\_SetEventMask, 249

events converted to text strings —  
csx\_Event2Text, 119

expression verification

— assert, 59

## F

find first (last) bit set in a long integer —  
ddi\_ffs, 341

ddi\_fls, 341

first CIS tuple — csx\_GetFirstTuple, 125

flushband — flush messages for specified  
priority band, 473

free access handle — csx\_FreeHandle, 121

free DMA handle

— ddi\_dma\_free\_handle, 314

free mapped register address space —

ddi\_regs\_map\_free, 417

free previously allocated memory —

ddi\_dma\_mem\_free, 321

free space — rmfree, 576

freerbuf — free a raw buffer header, 478

freeze, thaw the state of a stream —

freezestr, 479

unfreezestr, 479

freezestr — freeze, thaw the state of a  
stream, 479

## G

generic tuple parser — csx\_ParseTuple, 213

get interrupt block cookie

— ddi\_get\_iblock\_cookie, 257

get next DMA segment —

ddi\_dma\_nextseg, 326

get next DMA window —

ddi\_dma\_nextwin, 328

get soft interrupt block cookie

— ddi\_get\_soft\_iblock\_cookie, 260

getmajor — get major device number, 481

getminor — get minor device number, 482

getrbuf — get a raw buffer header, 486

## H

handle variable argument list

— va\_arg, 668

— va\_copy, 668

— va\_end, 668

— va\_start, 668

## I

I/O error

return — geterror, 480

I/O resources

release I/O resources —

csx\_ReleaseIO, 231

request I/O resources —

csx\_RequestIO, 231

I/O, block

suspend processes pending completion —  
biowait, 76

I/O, buffer

release buffer and notify processes —  
biodone, 68

I/O, paged request

allocate virtual address space —

bp\_mapin, 77

deallocate virtual address space —

bp\_mapout, 78

I/O, physical

— minphys, 531

— physio, 531

inb — read from an I/O port, 488

inform the system that a device's component  
is required. —

ddi\_dev\_is\_needed, 283

initialize a buffer structure — bioinit, 72

- initialize a named kstat —
  - kstat\_named\_init, 500
- inl — read from an I/O port, 488
- interrupt handling
  - add an interrupt — ddi\_add\_intr, 257
  - get interrupt block cookie —
    - ddi\_get\_iblock\_cookie, 257
  - remove an interrupt —
    - ddi\_remove\_intr, 257
- inw — read from an I/O port, 488
- IOC\_CONVERT\_FROM — Determine if there
  - is a need to translate
    - M\_IOCTL contents, 492
- IRQ resource
  - release IRQ resource —
    - csx\_ReleaseIRQ, 237
  - request IRQ resource —
    - csx\_RequestIRQ, 237

## K

- kernel address locations
  - between locations — bcopy, 63
- kernel addresses
  - get page frame number —
    - hat\_getkpfnum, 487
- Kernel interfaces for device ids
  - ddi\_devid\_compare, 281
  - ddi\_devid\_free, 281
  - ddi\_devid\_init, 281
  - ddi\_devid\_register, 281
  - ddi\_devid\_sizeof, 281
  - ddi\_devid\_unregister, 281
  - ddi\_devid\_valid, 281
- kernel modules, dynamic loading
  - add loadable module — mod\_install, 512
  - query loadable module — mod\_info, 512
  - remove loadable module —
    - mod\_remove, 512
- kernel state information — drv\_getparm, 459
- kmem\_alloc — allocate kernel memory, 493
- kmem\_free — allocate kernel memory, 493
- kmem\_zalloc — allocate kernel memory, 493
- kstat\_create — create and initialize a new
  - kstat, 496
- kstat\_delete — remove a kstat from the
  - system, 498

- kstat\_install — add a fully initialized kstat to
  - the system, 499
- kstat\_named\_init — initialize a named
  - kstat, 500
- kstat\_queue — update I/O kstat statistics, 501
- kstat\_runq\_back\_to\_waitq — update I/O kstat
  - statistics, 501
- kstat\_runq\_enter — update I/O kstat
  - statistics, 501
- kstat\_runq\_exit — update I/O kstat
  - statistics, 501
- kstat\_waitq\_enter — update I/O kstat
  - statistics, 501
- kstat\_waitq\_exit — update I/O kstat
  - statistics, 501
- kstat\_waitq\_to\_runq — update I/O kstat
  - statistics, 501

## L

- look up integer property —
  - ddi\_prop\_get\_int, 400
- lookup property information
  - ddi\_prop\_lookup, 403
  - ddi\_prop\_lookup\_byte\_array, 403
  - ddi\_prop\_lookup\_int\_array, 403
  - ddi\_prop\_lookup\_string, 403
  - ddi\_prop\_lookup\_string\_array, 403

## M

- makedevice — make device number from
  - major and minor
    - numbers, 506
- map memory area on PC Card —
  - csx\_MapMemPage, 142
- max — return the larger of two integers, 507
- memory
  - clear for a given number of bytes —
    - bzero, 84
- min — return the lesser of two integers, 508
- minor node for device
  - create — ddi\_create\_minor\_node, 275
- modify PC Card configuration —
  - csx\_ModifyConfiguration, 144
- modify window attributes —
  - csx\_ModifyWindow, 147

mt-streams — STREAMS multithreading, 515

mutex routines

- mutex, 517
- mutex\_destroy, 517
- mutex\_enter, 517
- mutex\_exit, 517
- mutex\_init, 517
- mutex\_owned, 517
- mutex\_tryenter, 517

mutual exclusion lock, *see* mutex,

## N

next CIS tuple — csx\_GetNextTuple, 125

nodes

- create minor nodes for client —  
csx\_MakeDeviceNode, 137

notify target driver of bus resets —

- scsi\_reset\_notify, 623

## O

obtain DDI information —

- csx\_CS\_DDI\_Info, 111

OTHERQ — get pointer to queue's partner

- queue, 524

outb — write to an I/O port, 526

outl — write to an I/O port, 526

outw — write to an I/O port, 526

## P

panic state — ddi\_in\_panic, 351

parse Battery Replacement Date tuple —

- csx\_Parse\_CISTPL\_BATTERY, 150

parse Byte Order tuple —

- csx\_Parse\_CISTPL\_BYTEORDER, 152

parse Card Configuration Table tuple —

- csx\_Parse\_CISTPL\_CFTABLE\_ENTRY, 151

parse Card Initialization Date tuple —

- csx\_Parse\_CISTPL\_DATE, 165

parse Configuration tuple —

- csx\_Parse\_CISTPL\_CONFIG, 162

parse Data Organization tuple —

- csx\_Parse\_CISTPL\_ORG, 203

parse Data Recording Format tuple —

- csx\_Parse\_CISTPL\_FORMAT, 175

parse Device Geo A tuple —

- csx\_Parse\_CISTPL\_DEVICEGEO\_A, 173

parse Device Geo tuple —

- csx\_Parse\_CISTPL\_DEVICEGEO, 171

parse Device Information tuple

for Attribute Memory —

- csx\_Parse\_CISTPL\_DEVICE\_A, 167

parse Function Extension tuple —

- csx\_Parse\_CISTPL\_FUNCE, 178

parse Function Identification tuple —

- csx\_Parse\_CISTPL\_FUNCID, 188

parse Geometry tuple —

- csx\_Parse\_CISTPL\_GEOMETRY, 191

parse JEDEC Identifier tuple

for Attribute Memory —

- csx\_Parse\_CISTPL\_JEDEC\_A, 193

parse Level-1 Version/Product Information

tuple —

- csx\_Parse\_CISTPL\_VERS\_1, 209

parse Level-2 Version and Information tuple —

- csx\_Parse\_CISTPL\_VERS\_2, 211

parse Link Target tuple —

- csx\_Parse\_CISTPL\_LINKTARGET, 195

parse Long Link A tuple

— csx\_Parse\_CISTPL\_LONGLINK\_A, 197

parse Long Link C tuple

— csx\_Parse\_CISTPL\_LONGLINK\_C, 197

parse Manufacturer Identification tuple —

- csx\_Parse\_CISTPL\_MANFID, 201

parse Multi-Function tuple —

- csx\_Parse\_CISTPL\_LONGLINK\_MFC, 199

parse Other Condition Device Information

tuple

for Attribute Memory —

- csx\_Parse\_CISTPL\_DEVICE\_OA, 167

parse Software Interleaving tuple —

- csx\_Parse\_CISTPL\_SWIL, 207

parse Special Purpose tuple —

- csx\_Parse\_CISTPL\_SPCL, 205

parser, for tuples (generic) —

- csx\_ParseTuple, 213

pci\_config\_get16 — read or write single datum

of various sizes to the PCI

Local Bus Configuration

space, 528

pci\_config\_get32 — read or write single datum

of various sizes to the PCI

Local Bus Configuration

space, 528

- pci\_config\_get64 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_get8 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_getb — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_getl — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_getll — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_getw — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_put16 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_put32 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_put64 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_put8 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_putb — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_putl — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_putll — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_putw — read or write single datum of various sizes to the PCI Local Bus Configuration space, 528
- pci\_config\_setup — enable/disable accesses to the PCI Local Bus Configuration space., 530
- pci\_config\_takedown — enable/disable accesses to the PCI Local Bus Configuration space., 530
- perform device context switching on a mapping —  
devmap\_do\_ctxmgt, 447
- pm\_busy\_component — control device components' availability for power management, 533
- pm\_idle\_component — control device components' availability for power management, 533
- pollwake up — inform a process that an event has occurred, 538
- proc\_ref — send a signal to a process, 539
- proc\_signal — send a signal to a process, 539
- proc\_unref — send a signal to a process, 539
- put — call a STREAMS put procedure, 544

## Q

- qbufcall — call a function when a buffer becomes available, 554
- qtimeout — execute a function after a specified length of time, 561
- qunbufcall — cancel a pending qbufcall request, 563
- quntimeout — cancel previous timeout function call, 564
- qwait — STREAMS wait routines, 565
- qwait\_sig — STREAMS wait routines, 565
- qwriter — asynchronous STREAMS perimeter upgrade, 567

## R

### raw buffer

- free a raw buffer header — `freerbuf`, 478
- get a raw buffer header — `getrbuf`, 486

### RD — get pointer to the read queue, 569

### read from an I/O port — `inb`, 488

`inl`, 488

`inw`, 488

`repinsb`, 488

`repinsd`, 488

`repinsw`, 488

### read data from mapped device in the memory space or allocated DMA memory

— `ddi_mem_get16`, 375

— `ddi_mem_get32`, 375

— `ddi_mem_get64`, 375

— `ddi_mem_get8`, 375

— `ddi_mem_rep_get16`, 379

— `ddi_mem_rep_get32`, 379

— `ddi_mem_rep_get64`, 379

— `ddi_mem_rep_get8`, 379

### read data from the device

— `ddi_get16`, 342

— `ddi_get32`, 342

— `ddi_get64`, 342

— `ddi_get8`, 342

### read data from the mapped device register in I/O space

— `ddi_io_get16`, 353

— `ddi_io_get32`, 353

— `ddi_io_get8`, 353

— `ddi_io_getb`, 353

— `ddi_io_getl`, 353

— `ddi_io_getw`, 353

### read data from the mapped memory address, device register or allocated DMA memory address

— `ddi_rep_get16`, 421

— `ddi_rep_get32`, 421

— `ddi_rep_get64`, 421

— `ddi_rep_get8`, 421

— `ddi_rep_getb`, 421

— `ddi_rep_getl`, 421

— `ddi_rep_getll`, 421

— `ddi_rep_getw`, 421

### read from device register

— `csx_Get16`, 122

— `csx_Get32`, 122

— `csx_Get64`, 122

— `csx_Get8`, 122

### read multiple data from the mapped device register in I/O space

— `ddi_io_rep_get16`, 360

— `ddi_io_rep_get32`, 360

— `ddi_io_rep_get8`, 360

— `ddi_io_rep_getb`, 360

— `ddi_io_rep_getl`, 360

— `ddi_io_rep_getw`, 360

### read or write a PC Card Configuration

Register —

`csx_AccessConfigurationRegister`, 105

### read or write single datum of various sizes to the PCI Local Bus

Configuration space

— `pci_config_get16`, 528

— `pci_config_get32`, 528

— `pci_config_get64`, 528

— `pci_config_get8`, 528

— `pci_config_getb`, 528

— `pci_config_getl`, 528

— `pci_config_getll`, 528

— `pci_config_getw`, 528

— `pci_config_put16`, 528

— `pci_config_put32`, 528

— `pci_config_put64`, 528

— `pci_config_put8`, 528

— `pci_config_putb`, 528

— `pci_config_putl`, 528

— `pci_config_putll`, 528

— `pci_config_putw`, 528

### read repetitively from device register

— `csx_RepGet16`, 222

— `csx_RepGet32`, 222

— `csx_RepGet64`, 222

— `csx_RepGet8`, 222

### readers/writer lock functions

— `rwlock`, 581

— `rw_destroy`, 581

— `rw_downgrade`, 581

— `rw_enter`, 581

— `rw_exit`, 581

— `rw_init`, 581

— `rw_read_locked`, 581

- rw\_tryenter, 581
- rw\_tryupgrade, 581
- register client — csx\_RegisterClient, 216
- release I/O resources — csx\_ReleaseIO, 231, 237, 240, 242
- release configuration on PC Card — csx\_ReleaseConfiguration, 220
- remove a kstat from the system — kstat\_delete, 498
- remove a soft interrupt
  - ddi\_remove\_softintr, 260
- remove an interrupt handler
  - ddi\_remove\_intr, 257
- remove client from Card Services list — csx\_DeregisterClient, 113
- repinsb — read from an I/O port, 488
- repinsd — read from an I/O port, 488
- repinsw — read from an I/O port, 488
- repoutsb — write to an I/O port, 526
- repoutsd — write to an I/O port, 526
- repoutsw — write to an I/O port, 526
- request I/O resources — csx\_RequestIO, 231, 237, 240, 242
- reset a function on a PC card — csx\_ResetFunction, 247
- resource map
  - allocate resource maps — rmallocmap, 573
  - free resource maps — rmallocmap, 573
- retrieve number of DMA windows — ddi\_dma\_numwin, 330
- retrieve subsequent DMA cookie — ddi\_dma\_nextcookie, 324
- return client event mask — csx\_GetEventMask, 249
- return client
  - csx\_GetFirstClient, 123
  - csx\_GetNextClient, 123
- return current access handle offset — csx\_GetHandleOffset, 127
- return data model type of current thread — ddi\_mmap\_get\_model, 383
- return data portion of tuple — csx\_GetTupleData, 134
- return driver binding name
  - ddi\_binding\_name, 266
  - ddi\_get\_name, 266
- return index matching capability string — scsi\_hba\_lookup\_capstr, 601

- return physical socket number — csx\_MapLogSocket, 140
- return status of PC Card and socket — csx\_GetStatus, 129
- return the devinfo node name — ddi\_node\_name, 387
- return the larger of two integers — max, 507
- return the lesser of two integers — min, 508
- return tuple
  - first CIS tuple — csx\_GetFirstTuple, 125
  - next CIS tuple — csx\_GetNextTuple, 125
- returns size of a buffer structure — biosize, 75
- returns the current time in seconds
  - ddi\_get\_time, 350
- returns the process ID
  - ddi\_get\_pid, 349
- returns the value of lbolt
  - returns the value of lbolt, 347
- rmalloc — allocate space from a resource map, 570
- rmalloc\_wait — allocate space from a resource map, 575
- rmfree — free space back into a resource map, 576

## S

- SAMESTR — test if next queue is in the same stream, 584
- SCSI Host Bus Adapter system initialization and completion routines
  - scsi\_hba\_init, 600
  - scsi\_hba\_init, 600
- SCSI commands, make packet
  - makecom, 504
  - makecom\_g0, 504
  - makecom\_g0\_s, 504
  - makecom\_g1, 504
  - makecom\_g5, 504
- SCSI dma utility routines
  - scsi\_dmafree, 591
  - scsi\_dmaget, 591
- SCSI HBA attach and detach routines
  - scsi\_hba\_attach, 597
  - scsi\_hba\_attach\_setup, 597
  - scsi\_hba\_detach, 597
- SCSI packet

- allocate a SCSI packet in iopb map —
  - get\_pktiopb, 483
- free a packet in iopb map —
  - free\_pktiopb, 483
- free an allocated SCSI packet and its DMA resource —
  - scsi\_destroy\_pkt, 590
- SCSI packet utility routines
  - scsi\_pktalloc, 617
  - scsi\_pktfree, 617
  - scsi\_resalloc, 617
  - scsi\_resfree, 617
- scsi\_abort — abort a SCSI command, 585
- scsi\_alloc\_consistent\_buf — scsi dma utility for allocating an I/O buffer for SCSI DMA, 586
- scsi\_cname — decode SCSI commands, 588
- scsi\_destroy\_pkt — free an allocated SCSI packet and its DMA resource, 590
- scsi\_dname — decode SCSI peripheral device type, 588
- scsi\_errmsg — display a SCSI request sense message, 593
- scsi\_free\_consistent\_buf — free a previously allocated SCSI DMA I/O buffer, 596
- scsi\_hba\_attach — SCSI HBA attach and detach routines, 597
- scsi\_hba\_attach\_setup — SCSI HBA attach and detach routines, 597
- scsi\_hba\_detach — SCSI HBA attach and detach routines, 597
- scsi\_hba\_fini — SCSI Host Bus Adapter system completion routines, 600
- scsi\_hba\_init — SCSI Host Bus Adapter system initialization routines, 600
- scsi\_hba\_lookup\_capstr — return index matching capability string, 601
- scsi\_hba\_pkt\_alloc — allocate and free a scsi\_pkt structure, 604
- scsi\_hba\_pkt\_free — allocate and free a scsi\_pkt structure, 604
- scsi\_hba\_probe — default SCSI HBA probe function, 605
- scsi\_hba\_tran\_alloc — allocate and free transport structures, 606
- scsi\_hba\_tran\_free — allocate and free transport structures, 606
- scsi\_ifgetcap — get SCSI transport capability, 607
- scsi\_ifsetcap — set SCSI transport capability, 607
- scsi\_init\_pkt — prepare a complete SCSI packet, 611
- scsi\_log — display a SCSI-device-related message, 615
- scsi\_mname — decode SCSI messages, 588
- scsi\_poll — run a polled SCSI command on behalf of a target driver, 619
- scsi\_probe — utility for probing a scsi device, 620
- scsi\_reset — reset a SCSI bus or target, 622
- scsi\_reset\_notify — notify target driver of bus resets, 623
- scsi\_rname — decode SCSI packet completion reasons, 588
- scsi\_setup\_cdb — setup SCSI command descriptor block (CDB), 625
- scsi\_slave — utility for SCSI target drivers to establish the presence of a target, 626
- scsi\_sname — decode SCSI sense keys, 588
- scsi\_sync\_pkt — synchronize CPU and I/O views of memory, 628
- scsi\_transport — request by a target driver to start a SCSI command, 629
- scsi\_unprobe — free resources allocated during initial probing, 631
- scsi\_unslave — free resources allocated during initial probing, 631
- scsi\_vu\_errmsg
  - display a SCSI request sense message, 632
- semaphore functions
  - semaphore, 635
  - sema\_destroy, 635
  - sema\_init, 635
  - sema\_p, 635
  - sema\_p\_sig, 635
  - sema\_try, 635
  - sema\_v, 635
- send a signal to a process
  - proc\_ref, 539
  - proc\_signal, 539

- proc\_unref, 539
- set client event mask —
  - csx\_RequestSocketMask, 240, 249
- set current access handle offset —
  - csx\_SetHandleOffset, 251
- Set driver memory mapping parameters
  - devmap\_devmem\_setup, 444
  - devmap\_umem\_setup, 444
- Set the device attributes for the mapping —
  - ddi\_mapdev\_set\_device\_acc\_attr, 369
- set up a mapping for a register address space
  - ddi\_regs\_map\_setup, 418
- setup SCSI command descriptor block (CDB)
  - scsi\_setup\_cdb, 625
- single direction elevator seek sort for buffers
  - disksort, 457
- size in bytes
  - convert size in pages — ptob, 541
  - convert to size in memory pages (round down) — btop, 79
  - convert to size in memory pages (round up) — btopr, 80
- socket number
  - return physical socket number —
    - csx\_MapLogSocket, 140
- soft interrupt handling
  - add a soft interrupt —
    - ddi\_add\_softintr, 260
  - get soft interrupt block cookie —
    - ddi\_get\_soft\_iblock\_cookie, 260
  - remove a soft interrupt —
    - ddi\_remove\_softintr, 260
  - trigger a soft interrupt —
    - ddi\_trigger\_softintr, 260
- sprintf — format characters in memory, 637
- status of PC Card and socket —
  - csx\_GetStatus, 129
- STREAMS wait routines — qwait,
  - qwait\_sig, 565
- STREAMS ioctl blocks
  - allocate — mkioch, 509
- STREAMS message blocks
  - allocate — allocb, 53
  - attach a user-supplied data buffer in place
    - esballoc, 470
  - call a function when a buffer becomes available — bufcall, 81, 554, 563, 662
  - call function when buffer is available —
    - esbcall, 472
  - concatenate bytes in a message —
    - msgpullup, 514, 542
  - concatenate two — linkb, 503
  - copy — copyb, 97
  - erase the contents of a buffer — clrbuf, 87
  - free all message blocks in a message —
    - freemsg, 477
  - free one — freeb, 476
  - remove from head of message —
    - unlinkb, 663
  - remove one from a message — rmvb, 577
- STREAMS message queue
  - insert a message into a queue — insq, 490
- STREAMS message queues, 60
- STREAMS Message queues
  - get next message — getq, 485
- STREAMS message queues
  - reschedule a queue for service —
    - enableok, 469
  - test for room — canputnext, 86
  - test for room — canput, 85
- STREAMS messages
  - copy a message — copymsg, 101
  - flush for specified priority band —
    - flushband, 473
  - remove from queue — flushq, 474, 579
  - return the number of bytes in a message
    - msgdsize, 513
  - submit messages to the log driver —
    - strlog, 644
  - test whether a message is a data message
    - datamsg, 254
  - trim bytes — adjmsg, 52
- STREAMS multithreading
  - mt-streams, 515
  - qbufcall — call a function when a buffer becomes available, 554
  - qtimeout — execute a function after a specified length of time, 561
  - qunbufcall — cancel a pending qbufcall request, 563

- quntimeout — cancel previous timeout function call, 564
- qwait, qwait\_sig — STREAMS wait routines, 565
- qwriter — asynchronous STREAMS perimeter upgrade, 567
- STREAMS put and service procedures
  - disable — qprocsoff, 557
  - enable — qprocson, 557
- STREAMS queues
  - change information about a queue or band of the queue — strqset, 648
  - enable a queue — qenable, 556
  - get pointer to queue's partner queue — OTHERQ, 524, 569
  - get information about a queue or band of the queue — strqget, 646
  - number of messages on a queue — qsize, 560
  - place a message at the head of a queue — putbq, 545
  - prevent a queue from being scheduled — noenable, 522
  - put a message on a queue — putq, 553
  - send a control message to a queue — putctl, 547, 551
  - send a control message with a one-byte parameter to a queue — putctl1, 546, 550
  - send a message on a stream in the reverse direction — greply, 558
  - send a message to the next queue — putnext, 549
  - test if next queue is in the same stream — SAMESTR, 584
  - test for flow control in specified priority band — bcanput, 61
- STREAMS write queues
  - get pointer for this module or driver — WR, 674
- STRUCT\_DECL
  - 32bit application data access macros, 649
- swab — swap bytes in 16-bit halfwords, 655
- synchronize CPU and I/O views of memory — ddi\_dma\_sync, 337, 628
- system DMA engine functions
  - ddi\_dmae, 309
  - ddi\_dmae\_1stparty, 309

- ddi\_dmae\_alloc, 309
- ddi\_dmae\_disable, 309
- ddi\_dmae\_enable, 309
- ddi\_dmae\_getattr, 309
- ddi\_dmae\_getcnt, 309
- ddi\_dmae\_getlim, 309
- ddi\_dmae\_prog, 309
- ddi\_dmae\_release, 309
- ddi\_dmae\_stop, 309

## T

- testb — check for an available buffer, 656
- timeout — execute a function after a specified length of time, 658
  - cancel previous timeout function call — untimeout, 664
- trigger a soft interrupt
  - ddi\_trigger\_softintr, 260
- tuple
  - first CIS tuple — csx\_GetFirstTuple, 125
  - next CIS tuple — csx\_GetNextTuple, 125
  - return data portion of tuple — csx\_GetTupleData, 134
- tuple entry
  - generic tuple parser — csx\_ParseTuple, 213
  - parse Device Information tuple for Attribute Memory — csx\_Parse\_CISTPL\_DEVICE\_A, 167, 193, 197
  - parse Battery Replacement Date tuple — csx\_Parse\_CISTPL\_BATTERY, 150
  - parse Byte Order tuple — csx\_Parse\_CISTPL\_BYTEORDER, 152
  - parse Card Configuration Table tuple — csx\_Parse\_CISTPL\_CFTABLE\_ENTRY, 154
  - parse Card Initialization Date tuple — csx\_Parse\_CISTPL\_DATE, 165
  - parse Configuration tuple — csx\_Parse\_CISTPL\_CONFIG, 162
  - parse Data Organization tuple — csx\_Parse\_CISTPL\_ORG, 203
  - parse Data Recording Format tuple — csx\_Parse\_CISTPL\_FORMAT, 175
  - parse Device Geo A tuple — csx\_Parse\_CISTPL\_DEVICE\_A, 173

- parse Device Geo tuple —
    - csx\_Parse\_CISTPL\_DEVICEGEO, 171
  - parse Function Extension tuple —
    - csx\_Parse\_CISTPL\_FUNC, 178
  - parse Function Identification tuple —
    - csx\_Parse\_CISTPL\_FUNCID, 188
  - parse Geometry tuple —
    - csx\_Parse\_CISTPL\_GEOMETRY, 191
  - parse Level-1 Version/Product Information tuple —
    - csx\_Parse\_CISTPL\_VERS\_1, 209
  - parse Level-2 Version and Information tuple —
    - csx\_Parse\_CISTPL\_VERS\_2, 211
  - parse Link Target tuple —
    - csx\_Parse\_CISTPL\_LINKTARGET, 195
  - parse Manufacturer Identification tuple —
    - csx\_Parse\_CISTPL\_MANFID, 201
  - parse Multi-Function tuple —
    - csx\_Parse\_CISTPL\_LONGLINK\_MFC, 199
  - parse Software Interleaving tuple —
    - csx\_Parse\_CISTPL\_SWIL, 207
  - parse Special Purpose tuple —
    - csx\_Parse\_CISTPL\_SPCL, 205
- U**
- uio structure
    - add character — ureadc, 666
    - remove a character — uwritec, 667
  - uiomove — copy kernel data using uio structure, 660
  - unbinds the address in a DMA handle —
    - ddi\_dma\_unbind\_handle, 339
  - unfreezestr — freeze, thaw the state of a stream, 479
  - uninitialize a buffer structure — biofini, 71
  - update I/O kstat statistics
    - kstat\_queue, 501
    - kstat\_runq\_back\_to\_waitq, 501
    - kstat\_runq\_enter, 501
    - kstat\_runq\_exit, 501
    - kstat\_waitq\_enter, 501
    - kstat\_waitq\_exit, 501
    - kstat\_waitq\_to\_runq, 501
  - update property information.
    - ddi\_prop\_update, 412
    - ddi\_prop\_update\_byte\_array, 412
    - ddi\_prop\_update\_int, 412
    - ddi\_prop\_update\_int\_array, 412
    - ddi\_prop\_update\_string, 412
    - ddi\_prop\_update\_string\_array, 412
- V**
- validate Card Information Structure (CIS) —
    - csx\_ValidateCIS, 252
  - va\_arg — handle variable argument list, 668
  - va\_copy — handle variable argument list, 668
  - va\_end — handle variable argument list, 668
  - va\_start — handle variable argument list, 668
  - virtual address
    - return mapped virtual address —
      - csx\_GetMappedAddr, 128
  - vsprintf — format characters in memory, 671
- W**
- write data to mapped device in the memory space or allocated DMA memory
    - ddi\_mem\_put16, 377
    - ddi\_mem\_put32, 377
    - ddi\_mem\_put64, 377
    - ddi\_mem\_put8, 377
    - ddi\_mem\_rep\_put16, 381
    - ddi\_mem\_rep\_put32, 381
    - ddi\_mem\_rep\_put64, 381
    - ddi\_mem\_rep\_put8, 381
  - write data to the device
    - ddi\_put16, 415
    - ddi\_put32, 415
    - ddi\_put64, 415
    - ddi\_put8, 415
  - write data to the mapped device register in I/O space
    - ddi\_io\_put16, 358
    - ddi\_io\_put32, 358
    - ddi\_io\_put8, 358
    - ddi\_io\_putb, 358

- ddi\_io\_putl, 358
- ddi\_io\_putw, 358

write data to the mapped memory address,  
device register or allocated  
DMA memory address

- ddi\_rep\_put16, 424
- ddi\_rep\_put32, 424
- ddi\_rep\_put64, 424
- ddi\_rep\_put8, 424
- ddi\_rep\_putb, 424
- ddi\_rep\_putl, 424
- ddi\_rep\_putll, 424
- ddi\_rep\_putw, 424

write multiple data to the mapped device  
register in I/O space

- ddi\_io\_rep\_put16, 362
- ddi\_io\_rep\_put32, 362
- ddi\_io\_rep\_put8, 362
- ddi\_io\_rep\_putb, 362
- ddi\_io\_rep\_putl, 362
- ddi\_io\_rep\_putw, 362

write or read a PC Card Configuration  
Register —  
csx\_AccessConfigurationRegister, 105

write repetitively to device register

- csx\_RepPut16, 224
- csx\_RepPut32, 224
- csx\_RepPut64, 224
- csx\_RepPut8, 224

write to an I/O port

- outb, 526
- outl, 526
- outw, 526
- repoutsb, 526
- repoutsd, 526
- repoutsw, 526

write to device register

- csx\_Put16, 215
- csx\_Put32, 215
- csx\_Put64, 215
- csx\_Put8, 215

## Z

zero fill the device register —  
ddi\_device\_zero, 278