



man pages section 9: DDI and DKI Kernel Functions

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-0702-10
December 2003

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 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, docs.sun.com, AnswerBook, AnswerBook2, 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.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. 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, docs.sun.com, AnswerBook, AnswerBook2, 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 REPOUDRE 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.



030827 @ 6671



Contents

Preface 13

Introduction 19

Intro(9F) 20

Kernel Functions for Drivers 39

adjmsg(9F) 40

alloca(9F) 41

anocancel(9F) 44

aphysio(9F) 45

ASSERT(9F) 47

backq(9F) 48

bcanput(9F) 49

bcmp(9F) 50

bcopy(9F) 51

bioclone(9F) 53

biodone(9F) 56

bioerror(9F) 58

biofini(9F) 59

bioinit(9F) 60

biomodified(9F) 61

bioreset(9F) 62

biosize(9F) 63

biowait(9F) 64

bp_mapin(9F) 65

bp_mapout(9F) 66
 btop(9F) 67
 btopr(9F) 68
 bufcall(9F) 69
 bzero(9F) 72
 canput(9F) 73
 canputnext(9F) 74
 clrbuf(9F) 75
 cmn_err(9F) 76
 condvar(9F) 81
 copyb(9F) 84
 copyin(9F) 86
 copymsg(9F) 88
 copyout(9F) 90
 csx_AccessConfigurationRegister(9F) 92
 csx_ConvertSize(9F) 94
 csx_ConvertSpeed(9F) 95
 csx_CS_DDI_Info(9F) 96
 csx_DeregisterClient(9F) 98
 csx_DupHandle(9F) 99
 csx_Error2Text(9F) 101
 csx_Event2Text(9F) 102
 csx_FreeHandle(9F) 103
 csx_Get8(9F) 104
 csx_GetFirstClient(9F) 105
 csx_GetFirstTuple(9F) 107
 csx_GetHandleOffset(9F) 109
 csx_GetMappedAddr(9F) 110
 csx_GetStatus(9F) 111
 csx_GetTupleData(9F) 115
 csx_MakeDeviceNode(9F) 117
 csx_MapLogSocket(9F) 119
 csx_MapMemPage(9F) 120
 csx_ModifyConfiguration(9F) 121
 csx_ModifyWindow(9F) 123
 csx_Parse_CISTPL_BATTERY(9F) 125
 csx_Parse_CISTPL_BYTEORDER(9F) 126
 csx_Parse_CISTPL_CFTABLE_ENTRY(9F) 128

csx_Parse_CISTPL_CONFIG(9F)	134
csx_Parse_CISTPL_DATE(9F)	137
csx_Parse_CISTPL_DEVICE(9F)	138
csx_Parse_CISTPL_DEVICEGEO(9F)	141
csx_Parse_CISTPL_DEVICEGEO_A(9F)	143
csx_Parse_CISTPL_FORMAT(9F)	145
csx_Parse_CISTPL_FUNCE(9F)	147
csx_Parse_CISTPL_FUNCID(9F)	155
csx_Parse_CISTPL_GEOMETRY(9F)	157
csx_Parse_CISTPL_JEDEC_C(9F)	158
csx_Parse_CISTPL_LINKTARGET(9F)	160
csx_Parse_CISTPL_LONGLINK_A(9F)	161
csx_Parse_CISTPL_LONGLINK_MFC(9F)	163
csx_Parse_CISTPL_MANFID(9F)	165
csx_Parse_CISTPL_ORG(9F)	166
csx_Parse_CISTPL_SPCL(9F)	167
csx_Parse_CISTPL_SWIL(9F)	169
csx_Parse_CISTPL_VERS_1(9F)	170
csx_Parse_CISTPL_VERS_2(9F)	171
csx_ParseTuple(9F)	172
csx_Put8(9F)	174
csx_RegisterClient(9F)	175
csx_ReleaseConfiguration(9F)	178
csx_RepGet8(9F)	180
csx_RepPut8(9F)	182
csx_RequestConfiguration(9F)	184
csx_RequestIO(9F)	188
csx_RequestIRQ(9F)	193
csx_RequestSocketMask(9F)	196
csx_RequestWindow(9F)	198
csx_ResetFunction(9F)	203
csx_SetEventMask(9F)	204
csx_SetHandleOffset(9F)	206
csx_ValidateCIS(9F)	207
datamsg(9F)	208
ddi_add_intr(9F)	209
ddi_add_softintr(9F)	212
ddi_binding_name(9F)	219

ddi_btop(9F) 220
 ddi_check_acc_handle(9F) 221
 ddi_copyin(9F) 223
 ddi_copyout(9F) 226
 ddi_create_minor_node(9F) 229
 ddi_device_copy(9F) 231
 ddi_device_zero(9F) 233
 ddi_devid_compare(9F) 234
 ddi_dev_is_needed(9F) 238
 ddi_dev_is_sid(9F) 240
 ddi_dev_nintrs(9F) 241
 ddi_dev_nregs(9F) 242
 ddi_dev_regsize(9F) 243
 ddi_dev_report_fault(9F) 244
 ddi_dma_addr_bind_handle(9F) 247
 ddi_dma_addr_setup(9F) 251
 ddi_dma_alloc_handle(9F) 253
 ddi_dma_buf_bind_handle(9F) 255
 ddi_dma_buf_setup(9F) 259
 ddi_dma_burstsizes(9F) 261
 ddi_dma_coff(9F) 262
 ddi_dma_curwin(9F) 263
 ddi_dma_devalign(9F) 264
 ddi_dmae(9F) 265
 ddi_dma_free(9F) 269
 ddi_dma_free_handle(9F) 270
 ddi_dma_get_attr(9F) 271
 ddi_dma_getwin(9F) 272
 ddi_dma_htoc(9F) 274
 ddi_dma_mem_alloc(9F) 275
 ddi_dma_mem_free(9F) 278
 ddi_dma_movwin(9F) 279
 ddi_dma_nextcookie(9F) 281
 ddi_dma_nextseg(9F) 283
 ddi_dma_nextwin(9F) 285
 ddi_dma_numwin(9F) 287
 ddi_dma_segtocookie(9F) 288
 ddi_dma_set_sbus64(9F) 290

ddi_dma_setup(9F)	291
ddi_dma_sync(9F)	293
ddi_dma_unbind_handle(9F)	295
ddi_driver_major(9F)	296
ddi_driver_name(9F)	297
ddi_enter_critical(9F)	298
ddi_ffs(9F)	299
ddi_get8(9F)	300
ddi_get_cred(9F)	302
ddi_get_devstate(9F)	303
ddi_get_driver_private(9F)	304
ddi_getiminor(9F)	305
ddi_get_instance(9F)	306
ddi_get_kt_did(9F)	307
ddi_get_lbolt(9F)	308
ddi_get_parent(9F)	309
ddi_get_pid(9F)	310
ddi_get_time(9F)	311
ddi_in_panic(9F)	312
ddi_intr_hilevel(9F)	313
ddi_io_get8(9F)	314
ddi_iomin(9F)	316
ddi_iopb_alloc(9F)	317
ddi_io_put8(9F)	319
ddi_io_rep_get8(9F)	321
ddi_io_rep_put8(9F)	323
ddi_log_sysevent(9F)	325
ddi_mapdev(9F)	328
ddi_mapdev_intercept(9F)	331
ddi_mapdev_set_device_acc_attr(9F)	333
ddi_map_regs(9F)	335
ddi_mem_alloc(9F)	337
ddi_mem_get8(9F)	339
ddi_mem_put8(9F)	340
ddi_mem_rep_get8(9F)	342
ddi_mem_rep_put8(9F)	344
ddi_mmap_get_model(9F)	346
ddi_model_convert_from(9F)	348

ddi_node_name(9F) 350
 ddi_no_info(9F) 351
 ddi_peek(9F) 352
 ddi_poke(9F) 354
 ddi_prop_create(9F) 356
 ddi_prop_exists(9F) 360
 ddi_prop_get_int(9F) 362
 ddi_prop_lookup(9F) 364
 ddi_prop_op(9F) 369
 ddi_prop_update(9F) 373
 ddi_put8(9F) 377
 ddi_regs_map_free(9F) 379
 ddi_regs_map_setup(9F) 380
 ddi_remove_minor_node(9F) 382
 ddi_removing_power(9F) 383
 ddi_rep_get8(9F) 385
 ddi_report_dev(9F) 387
 ddi_rep_put8(9F) 388
 ddi_root_node(9F) 390
 ddi_segmap(9F) 391
 ddi_slaveonly(9F) 394
 ddi_soft_state(9F) 395
 ddi_umem_alloc(9F) 400
 ddi_umem_iosetup(9F) 402
 ddi_umem_lock(9F) 404
 delay(9F) 406
 devmap_default_access(9F) 408
 devmap_devmem_setup(9F) 410
 devmap_do_ctxmgt(9F) 413
 devmap_set_ctx_timeout(9F) 416
 devmap_setup(9F) 417
 devmap_unload(9F) 419
 disksort(9F) 421
 drv_getparm(9F) 422
 drv_hztousec(9F) 424
 drv_priv(9F) 425
 drv_usectohz(9F) 426
 drv_usecwait(9F) 427

dupb(9F) 428
 dupmsg(9F) 431
 enableok(9F) 432
 esballoc(9F) 433
 esbbcall(9F) 435
 flushband(9F) 436
 flushq(9F) 437
 freeb(9F) 439
 freemsg(9F) 440
 freerbuf(9F) 441
 freezestr(9F) 442
 geterror(9F) 443
 gethrtime(9F) 444
 getmajor(9F) 445
 getminor(9F) 446
 get_pktiopb(9F) 447
 getq(9F) 449
 getrbuf(9F) 450
 gld(9F) 451
 hat_getkpfnum(9F) 454
 id32_alloc(9F) 455
 inb(9F) 456
 insq(9F) 458
 IOC_CONVERT_FROM(9F) 460
 kmem_alloc(9F) 461
 kmem_cache_create(9F) 463
 kstat_create(9F) 468
 kstat_delete(9F) 470
 kstat_install(9F) 471
 kstat_named_init(9F) 472
 kstat_queue(9F) 473
 linkb(9F) 475
 makecom(9F) 476
 makedevice(9F) 478
 max(9F) 479
 min(9F) 480
 mkiocb(9F) 481
 mod_install(9F) 484

msgdsz(9F) 485
 msgpullup(9F) 486
 mt-streams(9F) 487
 mutex(9F) 489
 nochpoll(9F) 492
 nodev(9F) 493
 noenable(9F) 494
 nulldev(9F) 495
 nvlist_add_boolean(9F) 496
 nvlist_alloc(9F) 498
 nvlist_lookup_boolean(9F) 500
 nvlist_next_nvpair(9F) 502
 nvlist_remove(9F) 504
 nvpair_value_byte(9F) 505
 OTHERQ(9F) 507
 outb(9F) 508
 pci_config_get8(9F) 510
 pci_config_setup(9F) 512
 pci_report_pmcaps(9F) 513
 pci_save_config_regs(9F) 515
 physio(9F) 517
 pm_busy_component(9F) 519
 pm_create_components(9F) 521
 pm_get_normal_power(9F) 522
 pm_power_has_changed(9F) 524
 pm_raise_power(9F) 526
 pm_trans_check(9F) 529
 pollwakeup(9F) 531
 proc_signal(9F) 532
 ptob(9F) 534
 pullupmsg(9F) 535
 put(9F) 537
 putbq(9F) 538
 putctl1(9F) 539
 putctl(9F) 540
 putnext(9F) 542
 putnextctl1(9F) 543
 putnextctl(9F) 544

putq(9F) 546
 qassociate(9F) 547
 qbufcall(9F) 549
 qenable(9F) 550
 qprocson(9F) 551
 qreply(9F) 552
 qsize(9F) 554
 qtimeout(9F) 555
 qunbufcall(9F) 556
 quntimeout(9F) 557
 qwait(9F) 558
 qwriter(9F) 560
 RD(9F) 561
 rmalloc(9F) 562
 rmallocmap(9F) 565
 rmalloc_wait(9F) 566
 rmfree(9F) 567
 rmvb(9F) 568
 rmvq(9F) 569
 rwlock(9F) 571
 SAMESTR(9F) 574
 scsi_abort(9F) 575
 scsi_alloc_consistent_buf(9F) 576
 scsi_cname(9F) 578
 scsi_destroy_pkt(9F) 580
 scsi_dmaget(9F) 581
 scsi_errmsg(9F) 583
 scsi_free_consistent_buf(9F) 586
 scsi_get_device_type_scsi_options(9F) 587
 scsi_hba_attach_setup(9F) 589
 scsi_hba_init(9F) 592
 scsi_hba_lookup_capstr(9F) 593
 scsi_hba_pkt_alloc(9F) 595
 scsi_hba_probe(9F) 597
 scsi_hba_tran_alloc(9F) 598
 scsi_ifgetcap(9F) 599
 scsi_init_pkt(9F) 603
 scsi_log(9F) 607

scsi_pktalloc(9F) 609
 scsi_poll(9F) 611
 scsi_probe(9F) 612
 scsi_reset(9F) 614
 scsi_reset_notify(9F) 615
 scsi_setup_cdb(9F) 616
 scsi_slave(9F) 617
 scsi_sync_pkt(9F) 619
 scsi_transport(9F) 620
 scsi_unprobe(9F) 621
 scsi_vu_errmsg(9F) 622
 semaphore(9F) 625
 sprintf(9F) 627
 stoi(9F) 629
 strchr(9F) 630
 strcmp(9F) 631
 strcpy(9F) 632
 strlen(9F) 633
 strlog(9F) 634
 strqget(9F) 636
 strqset(9F) 637
 STRUCT_DECL(9F) 638
 swab(9F) 643
 testb(9F) 644
 timeout(9F) 646
 uiomove(9F) 648
 unbufcall(9F) 649
 unlinkb(9F) 650
 untimeout(9F) 651
 ureadc(9F) 653
 uwritec(9F) 654
 va_arg(9F) 655
 vsprintf(9F) 657
 WR(9F) 660

Index 661

Preface

Both novice users and those familiar with the SunOS operating system can use online man pages to obtain information about the system and its features. 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.

Overview

The following contains a brief description of each man page section 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.
- 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 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 can 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 path name 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: <table> <tr> <td>[]</td><td>Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.</td></tr> <tr> <td>. . .</td><td>Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".</td></tr> <tr> <td> </td><td>Separator. Only one of the arguments separated by this character can be specified at a time.</td></tr> <tr> <td>{ }</td><td>Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.</td></tr> </table>	[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.	. . .	Ellipses. Several values can 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 a time.	{ }	Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.
[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.								
. . .	Ellipses. Several values can 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 a 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, and functions are described under USAGE.
IOCTL	This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the <code>ioctl(2)</code> system call is called <code>ioctl</code> and generates its own heading. <code>ioctl</code> calls for a specific device are listed alphabetically (on the man page for that specific device). <code>ioctl</code> calls are used for a particular class of devices all of which have an <code>io</code> ending, such as <code>mtio(7I)</code> .
OPTIONS	This section 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 <code>errno</code> 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 lists special rules, features, and commands that require in-depth explanations. The subsections listed here 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 <code>example%</code>, or if the user must be superuser, <code>example#</code>. 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 file names 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 <code>attributes(5)</code> 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.

Introduction

Intro(9F)

NAME	Intro – introduction to DDI/DKI functions
DESCRIPTION	Section 9F describes the kernel functions available for use by device drivers. See Intro(9E) for an overview of device driver interfaces. In this section, the information for each driver function is organized under the following headings: ■ NAME summarizes the function’s purpose. ■ SYNOPSIS shows the syntax of the function’s entry point in the source code. <code>#include</code> directives are shown for required headers. ■ INTERFACE LEVEL describes any architecture dependencies. ■ ARGUMENTS describes any arguments required to invoke the function. ■ DESCRIPTION describes general information about the function. ■ RETURN VALUES describes the return values and messages that can result from invoking the function. ■ CONTEXT indicates from which driver context (user, kernel, interrupt, or high-level interrupt) the function can be called. ■ 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. ■ A driver function has <i>kernel context</i> if it 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. ■ <i>Interrupt context</i> is kernel context, but also has an interrupt level associated with it. Driver interrupt routines have interrupt context. Note – A mutex acquired in user or kernel context that can also be acquired in interrupt context means that the user or kernel context thread holding that mutex is subject to all the restrictions imposed by interrupt context, for the duration of the ownership of that mutex. Please see the <code>mutex(9F)</code> man page for a more complete discussion of proper mutex handling for drivers. ■ <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>. ■ SEE ALSO indicates functions that are related by usage and sources, and which can be referred to for further information. ■ EXAMPLES shows how the function can be used in driver code.

**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
<code>adjmsg</code>	DDI/DKI
<code>allocb</code>	DDI/DKI
<code>backq</code>	DDI/DKI
<code>bcanput</code>	DDI/DKI
<code>bcanputnext</code>	DDI/DKI
<code>bufcall</code>	DDI/DKI
<code>canput</code>	DDI/DKI
<code>canputnext</code>	DDI/DKI
<code>clrbuf</code>	DDI/DKI
<code>copyb</code>	DDI/DKI
<code>copymsg</code>	DDI/DKI
<code>datamsg</code>	DDI/DKI
<code>dupb</code>	DDI/DKI
<code>dupmsg</code>	DDI/DKI
<code>enableok</code>	DDI/DKI
<code>esballoc</code>	DDI/DKI
<code>esbbcall</code>	DDI/DKI
<code>flushband</code>	DDI/DKI
<code>flushq</code>	DDI/DKI
<code>freeb</code>	DDI/DKI
<code>freemsg</code>	DDI/DKI
<code>freezestr</code>	DDI/DKI
<code>getq</code>	DDI/DKI
<code>insq</code>	DDI/DKI
<code>linkb</code>	DDI/DKI
<code>msgdsiz</code>	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

Routine	Type
strqget	DDI/DKI
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

Routine	Type
cmn_err	DDI/DKI
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

Routine	Type
ddi_dma_burstsizes	Solaris DDI
ddi_dma_coff	Solaris SPARC DDI
ddi_dma_curwin	Solaris SPARC DDI
ddi_dma_devalign	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

Routine	Type
<code>ddi_dmae_release</code>	Solaris x86 DDI
<code>ddi_dmae_stop</code>	Solaris x86 DDI
<code>ddi_enter_critical</code>	Solaris DDI
<code>ddi_exit_critical</code>	Solaris DDI
<code>ddi_ffs</code>	Solaris DDI
<code>ddi_fls</code>	Solaris DDI
<code>ddi_get16</code>	Solaris DDI
<code>ddi_get32</code>	Solaris DDI
<code>ddi_get64</code>	Solaris DDI
<code>ddi_get8</code>	Solaris DDI
<code>ddi_get_cred</code>	Solaris DDI
<code>ddi_get_driver_private</code>	Solaris DDI
<code>ddi_get_iblock_cookie</code>	Solaris DDI
<code>ddi_get_instance</code>	Solaris DDI
<code>ddi_get_name</code>	Solaris DDI
<code>ddi_get_parent</code>	Solaris DDI
<code>ddi_get_soft_iblock_cookie</code>	Solaris DDI
<code>ddi_get_soft_state</code>	Solaris DDI
<code>ddi_getb</code>	Solaris DDI
<code>ddi_getl</code>	Solaris DDI
<code>ddi_getll</code>	Solaris DDI
<code>ddi_getlongprop</code>	Solaris DDI
<code>ddi_getlongprop_buf</code>	Solaris DDI
<code>ddi_getprop</code>	Solaris DDI
<code>ddi_getproplen</code>	Solaris DDI
<code>ddi_getw</code>	Solaris DDI
<code>ddi_intr_hilevel</code>	Solaris DDI
<code>ddi_io_get16</code>	Solaris DDI
<code>ddi_io_get32</code>	Solaris DDI

Routine	Type
ddi_io_get8	Solaris DDI
ddi_io_getb	Solaris DDI
ddi_io_getl	Solaris DDI
ddi_io_getw	Solaris DDI
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

Routine	Type
<code>ddi_mapdev_set_device_acc_attr</code>	Solaris DDI
<code>ddi_mem_alloc</code>	Solaris DDI
<code>ddi_mem_free</code>	Solaris DDI
<code>ddi_mem_get16</code>	Solaris DDI
<code>ddi_mem_get32</code>	Solaris DDI
<code>ddi_mem_get64</code>	Solaris DDI
<code>ddi_mem_get8</code>	Solaris DDI
<code>ddi_mem_getb</code>	Solaris DDI
<code>ddi_mem_getl</code>	Solaris DDI
<code>ddi_mem_getll</code>	Solaris DDI
<code>ddi_mem_getw</code>	Solaris DDI
<code>ddi_mem_put16</code>	Solaris DDI
<code>ddi_mem_put32</code>	Solaris DDI
<code>ddi_mem_put64</code>	Solaris DDI
<code>ddi_mem_put8</code>	Solaris DDI
<code>ddi_mem_putb</code>	Solaris DDI
<code>ddi_mem_putl</code>	Solaris DDI
<code>ddi_mem_putll</code>	Solaris DDI
<code>ddi_mem_putw</code>	Solaris DDI
<code>ddi_mem_rep_get16</code>	Solaris DDI
<code>ddi_mem_rep_get32</code>	Solaris DDI
<code>ddi_mem_rep_get64</code>	Solaris DDI
<code>ddi_mem_rep_get8</code>	Solaris DDI
<code>ddi_mem_rep_getb</code>	Solaris DDI
<code>ddi_mem_rep_getl</code>	Solaris DDI
<code>ddi_mem_rep_getll</code>	Solaris DDI
<code>ddi_mem_rep_getw</code>	Solaris DDI
<code>ddi_mem_rep_put16</code>	Solaris DDI
<code>ddi_mem_rep_put32</code>	Solaris DDI

Routine	Type
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
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

Routine	Type
<code>ddi_prop_lookup</code>	Solaris DDI
<code>ddi_prop_lookup_byte_array</code>	Solaris DDI
<code>ddi_prop_lookup_int_array</code>	Solaris DDI
<code>ddi_prop_lookup_string</code>	Solaris DDI
<code>ddi_prop_lookup_string_array</code>	Solaris DDI
<code>ddi_prop_modify</code>	Solaris DDI
<code>ddi_prop_op</code>	Solaris DDI
<code>ddi_prop_remove</code>	Solaris DDI
<code>ddi_prop_remove_all</code>	Solaris DDI
<code>ddi_prop_undefine</code>	Solaris DDI
<code>ddi_prop_update</code>	Solaris DDI
<code>ddi_prop_update_byte_array</code>	Solaris DDI
<code>ddi_prop_update_int</code>	Solaris DDI
<code>ddi_prop_update_int_array</code>	Solaris DDI
<code>ddi_prop_update_string</code>	Solaris DDI
<code>ddi_prop_update_string_array</code>	Solaris DDI
<code>ddi_ptob</code>	Solaris DDI
<code>ddi_put16</code>	Solaris DDI
<code>ddi_put32</code>	Solaris DDI
<code>ddi_put64</code>	Solaris DDI
<code>ddi_put8</code>	Solaris DDI
<code>ddi_putb</code>	Solaris DDI
<code>ddi_putl</code>	Solaris DDI
<code>ddi_putll</code>	Solaris DDI
<code>ddi_putw</code>	Solaris DDI
<code>ddi_regs_map_free</code>	Solaris DDI
<code>ddi_regs_map_setup</code>	Solaris DDI
<code>ddi_remove_intr</code>	Solaris DDI
<code>ddi_remove_minor_node</code>	Solaris DDI

Routine	Type
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
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

Routine	Type
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
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

Routine	Type
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
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

Routine	Type
<code>mutex_owned</code>	Solaris DDI
<code>mutex_tryenter</code>	Solaris DDI
<code>nochpoll</code>	Solaris DDI
<code>nodev</code>	DDI/DKI
<code>nulldev</code>	DDI/DKI
<code>numtos</code>	Solaris DDI
<code>outb</code>	Solaris x86 DDI
<code>outl</code>	Solaris x86 DDI
<code>outw</code>	Solaris x86 DDI
<code>pci_config_get16</code>	Solaris DDI
<code>pci_config_get32</code>	Solaris DDI
<code>pci_config_get64</code>	Solaris DDI
<code>pci_config_get8</code>	Solaris DDI
<code>pci_config_getb</code>	Solaris DDI
<code>pci_config_getl</code>	Solaris DDI
<code>pci_config_getw</code>	Solaris DDI
<code>pci_config_put16</code>	Solaris DDI
<code>pci_config_put32</code>	Solaris DDI
<code>pci_config_put64</code>	Solaris DDI
<code>pci_config_put8</code>	Solaris DDI
<code>pci_config_putb</code>	Solaris DDI
<code>pci_config_putl</code>	Solaris DDI
<code>pci_config_putw</code>	Solaris DDI
<code>pci_config_setup</code>	Solaris DDI
<code>pci_config_teardown</code>	Solaris DDI
<code>physio</code>	Solaris DDI
<code>pollwakeup</code>	DDI/DKI
<code>proc_ref</code>	Solaris DDI
<code>proc_signal</code>	Solaris DDI

Routine	Type
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
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

Routine	Type
<code>scsi_errmsg</code>	Solaris DDI
<code>scsi_free_consistent_buf</code>	Solaris DDI
<code>scsi_hba_attach</code>	Solaris DDI
<code>scsi_hba_attach_setup</code>	Solaris DDI
<code>scsi_hba_detach</code>	Solaris DDI
<code>scsi_hba_fini</code>	Solaris DDI
<code>scsi_hba_init</code>	Solaris DDI
<code>scsi_hba_lookup_capstr</code>	Solaris DDI
<code>scsi_hba_pkt_alloc</code>	Solaris DDI
<code>scsi_hba_pkt_free</code>	Solaris DDI
<code>scsi_hba_probe</code>	Solaris DDI
<code>scsi_hba_tran_alloc</code>	Solaris DDI
<code>scsi_hba_tran_free</code>	Solaris DDI
<code>scsi_ifgetcap</code>	Solaris DDI
<code>scsi_ifsetcap</code>	Solaris DDI
<code>scsi_init_pkt</code>	Solaris DDI
<code>scsi_log</code>	Solaris DDI
<code>scsi_mname</code>	Solaris DDI
<code>scsi_pktalloc</code>	Solaris DDI
<code>scsi_pktfree</code>	Solaris DDI
<code>scsi_poll</code>	Solaris DDI
<code>scsi_probe</code>	Solaris DDI
<code>scsi_resalloc</code>	Solaris DDI
<code>scsi_reset</code>	Solaris DDI
<code>scsi_reset_notify</code>	Solaris DDI
<code>scsi_resfree</code>	Solaris DDI
<code>scsi_rname</code>	Solaris DDI
<code>scsi_slave</code>	Solaris DDI
<code>scsi_sname</code>	Solaris DDI

Routine	Type
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
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

Intro(9F)

SEE ALSO Intro(9E), mutex(9F)


Kernel Functions for Drivers

adjmsg(9F)

NAME	adjmsg – trim bytes from a message
SYNOPSIS	<pre>#include <sys/stream.h> int adjmsg(mblk_t *mp, ssize_t len);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Pointer to the message to be trimmed. <i>len</i> The number of bytes to be removed.
DESCRIPTION	The <code>adjmsg()</code> function removes bytes from a message. <code> len </code> (the absolute value of <code>len</code>) specifies the number of bytes to be removed. The <code>adjmsg()</code> function only trims bytes across message blocks of the same type. The <code>adjmsg()</code> 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, <code>adjmsg()</code> 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 <code>adjmsg()</code> function fails if <code> len </code> is greater than the number of bytes in the maximal leading sequence it finds. The <code>adjmsg()</code> 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 <code> len </code>.
RETURN VALUES	The <code>adjmsg()</code> function returns: 1 Successful completion. 0 An error occurred.
CONTEXT	The <code>adjmsg()</code> function can be called from user or interrupt context.
SEE ALSO	<i>STREAMS Programming Guide</i>

NAME	allocb – allocate a message block														
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *allocb(size_t size, uint_t pri);</pre>														
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).														
DESCRIPTION	allocb() tries to allocate a STREAMS message block. Buffer allocation fails only when the system is out of memory. If no buffer is available, the bufcall(9F) function can help a module recover from an allocation failure. A STREAMS message block is composed of three structures. The first structure is a message block (mblk_t). See msgb(9S). The mblk_t structure points to a data block structure (dblk_t). See datab(9S). 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. The fields in the mblk_t structure are initialized as follows: <table> <tr> <td>b_cont</td><td>set to NULL</td></tr> <tr> <td>b_rptr</td><td>points to the beginning of the data buffer</td></tr> <tr> <td>b_wptr</td><td>points to the beginning of the data buffer</td></tr> <tr> <td>b_datap</td><td>points to the dblk_t structure</td></tr> </table> The fields in the dblk_t structure are initialized as follows: <table> <tr> <td>db_base</td><td>points to the first byte of the data buffer</td></tr> <tr> <td>db_lim</td><td>points to the last byte + 1 of the buffer</td></tr> <tr> <td>db_type</td><td>set to M_DATA</td></tr> </table> The following figure identifies the data structure members that are affected when a message block is allocated.	b_cont	set to NULL	b_rptr	points to the beginning of the data buffer	b_wptr	points to the beginning of the data buffer	b_datap	points to the dblk_t structure	db_base	points to the first byte of the data buffer	db_lim	points to the last byte + 1 of the buffer	db_type	set to M_DATA
b_cont	set to NULL														
b_rptr	points to the beginning of the data buffer														
b_wptr	points to the beginning of the data buffer														
b_datap	points to the dblk_t structure														
db_base	points to the first byte of the data buffer														
db_lim	points to the last byte + 1 of the buffer														
db_type	set to M_DATA														

allocb(9F)



PARAMETERS *size* The number of bytes in the message block.

pri Priority of the request (no longer used).

RETURN VALUES Upon success, `allocb()` returns a pointer to the allocated message block of type `M_DATA`. On failure, `allocb()` returns a `NULL` pointer.

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

```

EXAMPLE 1 allocb() Code Sample (Continued)

```

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), esbbcall(9F), putnext(9F), testb(9F), datab(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.

anocancel(9F)

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

NAME	aphysio – perform asynchronous physical I/O												
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/buf.h> #include <sys/uio.h> #include <sys/aio_req.h> #include <sys/ddi.h> #include <sys/sunddi.h> int aphysio(int *stratstruct buf *, int *cancelstruct buf *, dev_t dev, int rw, void *mincntstruct buf *, struct aio_req *aio_reqp);</pre>												
PARAMETERS	<table> <tr> <td><i>strat</i></td><td>Pointer to device strategy routine.</td></tr> <tr> <td><i>cancel</i></td><td>Pointer to driver cancel routine. Used to cancel a submitted request. The driver must pass the address of the function <code>anocancel(9F)</code> because cancellation is not supported.</td></tr> <tr> <td><i>dev</i></td><td>The device number.</td></tr> <tr> <td><i>rw</i></td><td>Read/write flag. This is either <code>B_READ</code> when reading from the device or <code>B_WRITE</code> when writing to the device.</td></tr> <tr> <td><i>mincnt</i></td><td>Routine which bounds the maximum transfer unit size.</td></tr> <tr> <td><i>aio_reqp</i></td><td>Pointer to the <code>aio_req(9S)</code> structure which describes the user I/O request.</td></tr> </table>	<i>strat</i>	Pointer to device strategy routine.	<i>cancel</i>	Pointer to driver cancel routine. Used to cancel a submitted request. The driver must pass the address of the function <code>anocancel(9F)</code> because cancellation is not supported.	<i>dev</i>	The device number.	<i>rw</i>	Read/write flag. This is either <code>B_READ</code> when reading from the device or <code>B_WRITE</code> when writing to the device.	<i>mincnt</i>	Routine which bounds the maximum transfer unit size.	<i>aio_reqp</i>	Pointer to the <code>aio_req(9S)</code> structure which describes the user I/O request.
<i>strat</i>	Pointer to device strategy routine.												
<i>cancel</i>	Pointer to driver cancel routine. Used to cancel a submitted request. The driver must pass the address of the function <code>anocancel(9F)</code> because cancellation is not supported.												
<i>dev</i>	The device number.												
<i>rw</i>	Read/write flag. This is either <code>B_READ</code> when reading from the device or <code>B_WRITE</code> when writing to the device.												
<i>mincnt</i>	Routine which bounds the maximum transfer unit size.												
<i>aio_reqp</i>	Pointer to the <code>aio_req(9S)</code> structure which describes the user I/O request.												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).												
DESCRIPTION	<code>aphysio()</code> performs asynchronous I/O operations between the device and the address space described by <code>aio_reqp→aio_uio</code>. Prior to the start of the transfer, <code>aphysio()</code> 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. <code>aphysio()</code> does not wait for each transfer to complete, but returns as soon as the necessary requests have been made. <code>aphysio()</code> 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 <code>aphysio()</code> with <code>minphys(9F)</code>. <code>minphys(9F)</code> is the system <i>mincnt</i> routine. <code>minphys(9F)</code> ensures the transfer size does not exceed any system limits. If a driver supplies a local <i>mincnt</i> routine, this routine should perform the following actions: ■ If <code>bp→b_bcount</code> exceeds a device limit, set <code>bp→b_bcount</code> to a value supported by the device. ■ Call <code>minphys(9F)</code> to ensure that the driver does not circumvent additional system limits.												
RETURN VALUES	<code>aphysio()</code> returns:												

aphysio(9F)

0 Upon success.

non-zero 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.

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

backq(9F)

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

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

bcmp(9F)

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

NAME	bcopy – copy data between address locations in the kernel
SYNOPSIS	<pre>#include <sys/types.h> void bcopy(const void *from, void *to, size_t bcount);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>from</i> Source address from which the copy is made. <i>to</i> Destination address to which copy is made. <i>bcount</i> The number of bytes moved.
DESCRIPTION	bcopy() 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. Note that bcopy() 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 bcopy() always assumes that there will be no paging faults. If bcopy() attempts to access the user buffer when it is swapped out, the system will panic. It is safe to use bcopy() to move data within kernel space, since kernel space is never swapped out.
CONTEXT	bcopy() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Copying data between address locations in the kernel: 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). bcopy() is used since both the RAM disk and the buffer are part of the kernel address space. <pre> 1 #define RAMDNBLK 1000 /* blocks in the RAM disk */ 2 #define RAMDBSIZ 512 /* bytes per block */ 3 char ramdbls[RAMDNBLK][RAMDBSIZ]; /* blocks forming RAM /* disk ... 4 5 if (bp->b_flags & B_READ) /* if read request, copy data */ 6 /* from RAM disk data block */ 7 /* to system buffer */ 8 bcopy(&ramdbls[bp->b_blkno][0], bp->b_un.b_addr, 9 bp->b_bcount); 10 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, &ramdbls[bp->b_blkno][0], 16 bp->b_bcount);</pre>
SEE ALSO	copyin(9F), copyout(9F)

bcopy(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.

NAME	bioclone – clone another buffer																
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>																
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).																
PARAMETERS	<table> <tr> <td><i>bp</i></td><td>Pointer to the buf(9S) structure describing the original I/O request.</td></tr> <tr> <td><i>off</i></td><td>Offset within original I/O request where new I/O request should start.</td></tr> <tr> <td><i>len</i></td><td>Length of the I/O request.</td></tr> <tr> <td><i>dev</i></td><td>Device number.</td></tr> <tr> <td><i>blkno</i></td><td>Block number on device.</td></tr> <tr> <td><i>iodone</i></td><td>Specific biodone(9F) routine.</td></tr> <tr> <td><i>bp_mem</i></td><td>Pointer to a buffer structure to be filled in or NULL.</td></tr> <tr> <td><i>sleepflag</i></td><td>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.</td></tr> </table>	<i>bp</i>	Pointer to the buf(9S) structure describing the original I/O request.	<i>off</i>	Offset within original I/O request where new I/O request should start.	<i>len</i>	Length of the I/O request.	<i>dev</i>	Device number.	<i>blkno</i>	Block number on device.	<i>iodone</i>	Specific biodone(9F) routine.	<i>bp_mem</i>	Pointer to a buffer structure to be filled in or NULL.	<i>sleepflag</i>	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.
<i>bp</i>	Pointer to the buf(9S) structure describing the original I/O request.																
<i>off</i>	Offset within original I/O request where new I/O request should start.																
<i>len</i>	Length of the I/O request.																
<i>dev</i>	Device number.																
<i>blkno</i>	Block number on device.																
<i>iodone</i>	Specific biodone(9F) routine.																
<i>bp_mem</i>	Pointer to a buffer structure to be filled in or NULL.																
<i>sleepflag</i>	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.																
DESCRIPTION	bioclone() 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</i> + <i>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 biodone(9F) 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, bioclone() will allocate a new buffer using getrbuf(9F). If <i>sleepflag</i> is set to KM_SLEEP, the driver may sleep until space is freed up. If <i>sleepflag</i> is set to KM_NOSLEEP, 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 freerbuf(9F). If <i>bp_mem</i> is not NULL, it will be used as the space for the buffer structure. The driver has to ensure that <i>bp_mem</i> 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 bioclone(), 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.																

bioclone(9F)

	The driver has to ensure that the original buffer is not freed while any of the clone buffers is still performing I/O. The <code>biodone()</code> function has to be called on all clone buffers before it is called on the original buffer.
RETURN VALUES	The <code>bioclone()</code> function returns a pointer to the initialized buffer header, or <code>NULL</code> if no space is available.
CONTEXT	<code>bioclone()</code> can be called from user or interrupt context. Drivers must not allow <code>bioclone()</code> to sleep if called from an interrupt routine.
EXAMPLES	EXAMPLE 1 Using <code>bioclone()</code> for Disk Striping A device driver can use <code>bioclone()</code> 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. <pre>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); } }</pre>

EXAMPLE 1 Using `bioclone()` for Disk Striping *(Continued)*

```
        return (0);  
    }
```

SEE ALSO `biodone(9F)`, `bp_mapin(9F)`, `freerbuf(9F)`, `getrbuf(9F)`, `buf(9S)`

Writing Device Drivers

biodone(9F)

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

biodone(9F)

```

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.

bioerror(9F)

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

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

bioinit(9F)

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

NAME	biomodified – check if a buffer is modified						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int biomodified(struct buf *bp);</pre>						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).						
PARAMETERS	<i>bp</i> Pointer to the buffer header structure.						
DESCRIPTION	The <code>biomodified()</code> function returns status to indicate if the buffer is modified. The <code>biomodified()</code> 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 <code>buf(9S)</code> structure. The <code>biomodified()</code> 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 unmodified when it writes the pages to the backing store. The <code>biomodified()</code> function can be used to detect any modifications to the memory pages while I/O is in progress. A device driver can use <code>biomodified()</code> for disk mirroring. An application is allowed to <code>mmap</code> 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 <code>biomodified()</code> 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.						
RETURN VALUES	The <code>biomodified()</code> function returns the following values: <table> <tr> <td>1</td><td>Buffer is modified.</td></tr> <tr> <td>0</td><td>Buffer is not modified.</td></tr> <tr> <td>-1</td><td>Buffer is not used for paged I/O request.</td></tr> </table>	1	Buffer is modified.	0	Buffer is not modified.	-1	Buffer is not used for paged I/O request.
1	Buffer is modified.						
0	Buffer is not modified.						
-1	Buffer is not used for paged I/O request.						
CONTEXT	<code>biomodified()</code> can be called from any context.						
SEE ALSO	<code>bp_mapin(9F)</code>, <code>buf(9S)</code> <i>Writing Device Drivers</i>						

bioreset(9F)

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

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

biowait(9F)

NAME	biowait – suspend processes pending completion of block I/O				
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/buf.h> int biowait(struct buf *bp);</pre>				
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).				
PARAMETERS	<i>bp</i> Pointer to the buf structure describing the transfer.				
DESCRIPTION	Drivers allocating their own buf 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. 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.				
RETURN VALUES	<table><tr><td>0</td><td>Upon success</td></tr><tr><td>non-zero</td><td>Upon I/O failure. <code>biowait()</code> calls <code>geterror(9F)</code> to retrieve the error number which it returns.</td></tr></table>	0	Upon success	non-zero	Upon I/O failure. <code>biowait()</code> calls <code>geterror(9F)</code> to retrieve the error number which it returns.
0	Upon success				
non-zero	Upon I/O failure. <code>biowait()</code> calls <code>geterror(9F)</code> to retrieve the error number which it returns.				
CONTEXT	<code>biowait()</code> can be called from user context only.				
SEE ALSO	<code>biodone(9F)</code> , <code>geterror(9F)</code> , <code>getrbuf(9F)</code> , <code>buf(9S)</code> <i>Writing Device Drivers</i>				

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

bp_mapout(9F)

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

NAME	btop – convert size in bytes to size in pages (round down)
SYNOPSIS	<pre>#include <sys/ddi.h> unsigned long btop(unsigned long <i>numbytes</i>) ;</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>numbytes</i> Number of bytes.
DESCRIPTION	btop() 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 btop(4096) returns 2, and btop(4097) returns 2 as well. btop(0) returns 0.
RETURN VALUES	The return value is always the number of pages. There are no invalid input values, and therefore no error return values.
CONTEXT	btop() can be called from user or interrupt context.
SEE ALSO	btopr(9F), ddi_btop(9F), ptob(9F) <i>Writing Device Drivers</i>

btopr(9F)

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

NAME	bufcall – call a function when a buffer becomes available
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stream.h> bufcall_id_t bufcall(size_t <i>size</i>, uint_t <i>pri</i>, void *<i>func</i>void *<i>arg</i>, void *<i>arg</i>);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>size</i> Number of bytes required for the buffer. <i>pri</i> Priority of the allocb(9F) allocation request (not used). <i>func</i> Function or driver routine to be called when a buffer becomes available. <i>arg</i> Argument to the function to be called when a buffer becomes available.
DESCRIPTION	bufcall() serves as a timeout(9F) call of indeterminate length. When a buffer allocation request fails, bufcall() 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 allocb() or it may do something else.
RETURN VALUES	If successful, bufcall() returns a bufcall ID that can be used in a call to unbufcall() to cancel the request. If the bufcall() scheduling fails, <i>func</i> is never called and 0 is returned.
CONTEXT	bufcall() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Calling a function when a buffer becomes available: The purpose of this srv(9E) service routine is to add a header to all M_DATA messages. Service routines must process all messages on their queues before returning, or arrange to be rescheduled 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 canputnext(9F) succeeded, then send all but M_DATA messages to the next module with putnext(9F) (line 16). For M_DATA 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 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 <i>q</i> once a buffer of the specified size (sizeof (struct hdr)) becomes available. If it does, qenable(9F) will put <i>q</i> 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.

bufcall(9F)

EXAMPLE 1 Calling a function when a buffer becomes available: (Continued)

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     return (0);
41 }
```

SEE ALSO `srv(9E)`, `allocb(9F)`, `canputnext(9F)`, `esballoc(9F)`, `esbbcall(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)`.

bzero(9F)

NAME	bzero – clear memory for a given number of bytes
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> void bzero(void *addr, size_t bytes);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>addr</i> Starting virtual address of memory to be cleared. <i>bytes</i> The number of bytes to clear starting at <i>addr</i>.
DESCRIPTION	bzero() clears a contiguous portion of memory by filling it with zeros.
CONTEXT	bzero() can be called from user or interrupt context.
SEE ALSO	bcopy(9F), clrbuf(9F), kmem_zalloc(9F)
	<i>Writing Device Drivers</i>
WARNINGS	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.

NAME	canput – test for room in a message queue
SYNOPSIS	<pre>#include <sys/stream.h> int canput (queue_t *q) ;</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the message queue.
DESCRIPTION	canput () 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>→q_next) and bcanput (<i>q</i>→q_next , <i>pri</i>) respectively.
RETURN VALUES	1 If the message queue is not full. 0 If the queue is full.
CONTEXT	canput () can be called from user or interrupt context.
SEE ALSO	bcanput(9F), bcanputnext(9F), canputnext(9F), putbq(9F), putnext(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
WARNINGS	Drivers are responsible for both testing a queue with canput () and refraining from placing a message on the queue if canput () fails.

canputnext(9F)

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

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

cmn_err(9F)

NAME	cmn_err, vcmn_err – display an error message or panic the system						
SYNOPSIS	<pre>#include <sys/cmn_err.h> #include <sys/ddi.h> #include <sys/sunddi.h> void cmn_err(int <i>level</i>, char *<i>format</i>...); #include <sys/varargs.h> void vcmn_err(int <i>level</i>, char *<i>format</i>, va_list <i>ap</i>);</pre>						
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).						
PARAMETERS							
cmn_err()	<i>level</i> A constant indicating the severity of the error condition. <i>format</i> The message to be displayed.						
vcmn_err()	vcmn_err() takes <i>level</i> and <i>format</i> as described for cmn_err(), but its third argument is different: <i>ap</i> The variable argument list passed to the function.						
DESCRIPTION							
cmn_err()	cmn_err() displays a specified message on the console. cmn_err() 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 CE_PANIC <i>level</i> below. <i>level</i> is a constant indicating the severity of the error condition. The four severity levels are: <table> <tr> <td>CE_CONT</td><td>Used to continue another message or to display an informative message not associated with an error. Note that multiple CE_CONT messages without a newline may or may not appear on the system console or in the system log as a single line message. A single line message may be produced by constructing the message with sprintf(9F) or vsprintf(9F) before calling cmn_err().</td></tr> <tr> <td>CE_NOTE</td><td>Used to display a message preceded with NOTICE. This message is used to report system events that do not 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.</td></tr> <tr> <td>CE_WARN</td><td>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.</td></tr> </table>	CE_CONT	Used to continue another message or to display an informative message not associated with an error. Note that multiple CE_CONT messages without a newline may or may not appear on the system console or in the system log as a single line message. A single line message may be produced by constructing the message with sprintf(9F) or vsprintf(9F) before calling cmn_err().	CE_NOTE	Used to display a message preceded with NOTICE. This message is used to report system events that do not 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_CONT	Used to continue another message or to display an informative message not associated with an error. Note that multiple CE_CONT messages without a newline may or may not appear on the system console or in the system log as a single line message. A single line message may be produced by constructing the message with sprintf(9F) or vsprintf(9F) before calling cmn_err().						
CE_NOTE	Used to display a message preceded with NOTICE. This message is used to report system events that do not 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 log.

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,u	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.
p	The argument is taken to be a pointer; the value of the pointer is displayed in unsigned hexadecimal. The display format is equivalent to %lx. To avoid lint warnings, cast pointers to type void * when using the %p format specifier.

cmn_err(9F)

	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 <code>NULL</code>, the string <code><null string></code> is used in its place. % Copy a %; no argument is converted.
	The first character in <i>format</i> affects where the message will be written:
	! The message goes only to the system log.
	^ The message goes only to the console.
	? If <i>level</i> is also <code>CE_CONT</code>, the message is always sent to the system log, but is only written to the console when the system has been booted in verbose mode. See <code>kernel(1M)</code>. If neither condition is met, the '?' character has no effect and is simply ignored.
	Refer to <code>syslogd(1M)</code> to determine where the system log is written.
	<code>cmn_err()</code> appends a <code>\n</code> to each <i>format</i>, except when <i>level</i> is <code>CE_CONT</code>.
vcmn_err()	<code>vcmn_err()</code> is identical to <code>cmn_err()</code> except that its last argument, <i>ap</i>, is a pointer to a variable list of arguments. <i>ap</i> contains the list of arguments used by the conversion specifications in <i>format</i>. <i>ap</i> must be initialized by calling <code>va_start(9F)</code>. <code>va_end(9F)</code> is used to clean up and must be called after each traversal of the list. Multiple traversals of the argument list, each bracketed by <code>va_start(9F)</code> and <code>va_end(9F)</code>, are possible.
RETURN VALUES	None. However, if an unknown <i>level</i> is passed to <code>cmn_err()</code>, the following panic error message is displayed:
	<pre>panic: unknown level in cmn_err (level=level, msg=format)</pre>
CONTEXT	<code>cmn_err()</code> can be called from user, kernel, interrupt, or high-level interrupt context.
EXAMPLES	EXAMPLE 1 Using <code>cmn_err()</code> This first example shows how <code>cmn_err()</code> can record tracing and debugging information only in the system log (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 log (line 28). <pre> 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 }; </pre>

EXAMPLE 1 Using `cmn_err()` (Continued)

```

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\n", 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 log 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;

```

cmn_err(9F)

EXAMPLE 4 Error Routine *(Continued)*

```
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(3C), 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 Messages of arbitrary length can be generated using cmn_err(), but if the call to
cmn_err() is made from high-level interrupt context and insufficient memory is
available to create a buffer of the specified size, the message will be truncated to
LOG_MSGSIZE bytes (see <sys/log.h>). For this reason, callers of cmn_err() that
require complete and accurate message generation should post down from high-level
interrupt context before calling cmn_err().

NAME	condvar, cv_init, cv_destroy, cv_wait, cv_signal, cv_broadcast, cv_wait_sig, cv_timedwait, cv_timedwait_sig – condition variable routines												
SYNOPSIS	<pre>#include <sys/ksynch.h> 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>												
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>cvp</i></td><td>A pointer to an abstract data type kcondvar_t.</td></tr> <tr> <td><i>mp</i></td><td>A pointer to a mutual exclusion lock (kmutex_t), initialized by mutex_init(9F) and held by the caller.</td></tr> <tr> <td><i>name</i></td><td>Descriptive string. This is obsolete and should be NULL. (Non-NULL strings are legal, but they're a waste of kernel memory.)</td></tr> <tr> <td><i>type</i></td><td>The constant CV_DRIVER.</td></tr> <tr> <td><i>arg</i></td><td>A type-specific argument, drivers should pass arg as NULL.</td></tr> <tr> <td><i>timeout</i></td><td>A time, in absolute ticks since boot, when cv_timedwait() or cv_timedwait_sig() should return.</td></tr> </table>	<i>cvp</i>	A pointer to an abstract data type kcondvar_t.	<i>mp</i>	A pointer to a mutual exclusion lock (kmutex_t), initialized by mutex_init(9F) and held by the caller.	<i>name</i>	Descriptive string. This is obsolete and should be NULL. (Non-NULL strings are legal, but they're a waste of kernel memory.)	<i>type</i>	The constant CV_DRIVER.	<i>arg</i>	A type-specific argument, drivers should pass arg as NULL.	<i>timeout</i>	A time, in absolute ticks since boot, when cv_timedwait() or cv_timedwait_sig() should return.
<i>cvp</i>	A pointer to an abstract data type kcondvar_t.												
<i>mp</i>	A pointer to a mutual exclusion lock (kmutex_t), initialized by mutex_init(9F) and held by the caller.												
<i>name</i>	Descriptive string. This is obsolete and should be NULL. (Non-NULL strings are legal, but they're a waste of kernel memory.)												
<i>type</i>	The constant CV_DRIVER.												
<i>arg</i>	A type-specific argument, drivers should pass arg as NULL.												
<i>timeout</i>	A time, in absolute ticks since boot, when cv_timedwait() or cv_timedwait_sig() should return.												
DESCRIPTION	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 cv_init(), and must be deallocated by calling cv_destroy(). 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												

condvar(9F)

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_lbolt(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> , or due to a premature wakeup (see NOTES).

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.

```
mutex_enter(&un->un_lock);
while (un->un_flag & UNIT_BUSY)
    cv_wait(&un->un_cv, &un->un_lock);
un->un_flag |= UNIT_BUSY;
mutex_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.

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

NOTES

It is possible for `cv_wait()`, `cv_wait_sig()`, `cv_timedwait()`, and `cv_timedwait_sig()` to return prematurely, that is, not due to a call to `cv_signal()` or `cv_broadcast()`. This occurs most commonly in the case of `cv_wait_sig()` and `cv_timedwait_sig()` when the thread is stopped and restarted by job control signals or by a debugger, but can happen in other cases as well, even for `cv_wait()`. Code that calls these functions must always recheck the reason for blocking and call again if the reason for blocking is still true.

If your driver needs to wait on behalf of processes that have real-time constraints, use `cv_timedwait()` rather than `delay(9F)`. The `delay()` function calls `timeout(9F)`, which can be subject to priority inversions.

SEE ALSO

`kill(2)`, `ddi_get_lbolt(9F)`, `mutex(9F)`, `mutex_init(9F)`

Writing Device Drivers

copyb(9F)

NAME	copyb – copy a message block
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *copyb(mblk_t *bp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>bp</i> Pointer to the message block from which data is copied.
DESCRIPTION	copyb() 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. copyb() uses the <i>b_rptr</i> and <i>b_wptr</i> members of <i>bp</i> to determine how many bytes to copy.
RETURN VALUES	If successful, copyb() returns a pointer to the newly allocated message block containing the copied data. Otherwise, it returns a NULL pointer.
CONTEXT	copyb() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 : Using copyb For each message in the list, test to see if the downstream queue is full with the canputnext(9F) function (line 21). If it is not full, use copyb to copy a header message block, and dupmsg(9F) to duplicate the data to be retransmitted. If either operation fails, reschedule a timeout at the next valid interval. 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. <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 ... 10 int h_address; 11 ... 12 }; 13 14 void 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 */</pre>

EXAMPLE 1 : Using copyb *(Continued)*

```

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

copyin(9F)

NAME	copyin – copy data from a user program to a driver buffer
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> int copyin(const void *userbuf, void *driverbuf, size_t cn);</pre>
INTERFACE LEVEL	This interface is obsolete. ddi_copyin(9F) should be used instead.
PARAMETERS	<i>userbuf</i> User program source address from which data is transferred. <i>driverbuf</i> Driver destination address to which data is transferred. <i>cn</i> Number of bytes transferred.
DESCRIPTION	copyin() 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. 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.
RETURN VALUES	Under normal conditions, a 0 is returned indicating a successful copy. Otherwise, a –1 is returned if one of the following occurs: ■ Paging fault; the driver tried to access a page of memory for which it did not have read or write access. ■ Invalid user address, such as a user area or stack area. ■ Invalid address that would have resulted in data being copied into the user block. ■ Hardware fault; a hardware error prevented access to the specified user memory. For example, an uncorrectable parity or ECC error occurred. If a –1 is returned to the caller, driver entry point routines should return EFAULT.
CONTEXT	copyin() can be called from user context only.
EXAMPLES	EXAMPLE 1 An ioctl() Routine A driver ioctl(9E) routine (line 10) can be used to get or set device attributes or registers. In the XX_GETREGS 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. <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 rcv_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 */</pre>

EXAMPLE 1 An `ioctl()` Routine (Continued)

```

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     }

```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `attributes(5)`, `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.

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

copymsg(9F)

NAME	copymsg – copy a message
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *copymsg(mblk_t *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Pointer to the message to be copied.
DESCRIPTION	copymsg() forms a new message by allocating new message blocks, and copying the contents of the message referred to by <i>mp</i> (using the copyb(9F) function). It returns a pointer to the new message.
RETURN VALUES	If the copy is successful, copymsg() returns a pointer to the new message. Otherwise, it returns a NULL pointer.
CONTEXT	copymsg() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 : Using copymsg The routine lctouc() 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 copymsg() 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). <pre> 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 }</pre>

EXAMPLE 1 : Using copymsg *(Continued)*

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

Writing Device Drivers

STREAMS Programming Guide

copyout(9F)

NAME	copyout – copy data from a driver to a user program
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> int copyout(const void *driverbuf, void *userbuf, size_t cn);</pre>
INTERFACE LEVEL	This interface is obsolete. ddi_copyout(9F) should be used instead.
PARAMETERS	<i>driverbuf</i> Source address in the driver from which the data is transferred. <i>userbuf</i> Destination address in the user program to which the data is transferred. <i>cn</i> Number of bytes moved.
DESCRIPTION	copyout () copies data from driver buffers to user data space. 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.
RETURN VALUES	Under normal conditions, a 0 is returned to indicate a successful copy. Otherwise, a –1 is returned if one of the following occurs: ■ Paging fault; the driver tried to access a page of memory for which it did not have read or write access. ■ Invalid user address, such as a user area or stack area. ■ Invalid address that would have resulted in data being copied into the user block. ■ Hardware fault; a hardware error prevented access to the specified user memory. For example, an uncorrectable parity or ECC error occurred. If a –1 is returned to the caller, driver entry point routines should return EFAULT.
CONTEXT	copyout () can be called from user context only.
EXAMPLES	EXAMPLE 1 An ioctl () Routine A driver ioctl(9E) routine (line 10) can be used to get or set device attributes or registers. In the XX_GETREGS 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. <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 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,</pre>

EXAMPLE 1 An `ioctl()` Routine (Continued)

```
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 }
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `attributes(5)`, `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.

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

csx_AccessConfigurationRegister(9F)

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

`csx_AccessConfigurationRegister(9F)`

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

csx_ConvertSize(9F)

NAME	csx_ConvertSize – convert device sizes								
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ConvertSize(convert_size_t *cs);</pre>								
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)								
PARAMETERS	<i>cs</i> Pointer to a convert_size_t structure.								
DESCRIPTION	csx_ConvertSize() 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.								
STRUCTURE MEMBERS	The structure members of convert_size_t are: <pre>uint32_t Attributes; uint32_t bytes; uint32_t devsize;</pre> The fields are defined as follows: <table> <tr> <td>Attributes</td><td>This is a bit-mapped field that identifies the type of size conversion to be performed. The field is defined as follows:</td></tr> <tr> <td></td><td> <pre>CONVERT_BYTES_TO_DEVSIZE Converts <i>bytes</i> to <i>devsize</i> format. CONVERT_DEVSIZE_TO_BYTES Converts <i>devsize</i> format to <i>bytes</i>.</pre> </td></tr> <tr> <td>bytes</td><td>If CONVERT_BYTES_TO_DEVSIZE is set, the value in the bytes field is converted to a <i>devsize</i> format and returned in the devsize field.</td></tr> <tr> <td>devsize</td><td>If CONVERT_DEVSIZE_TO_BYTES is set, the value in the devsize field is converted to a <i>bytes</i> value and returned in the bytes field.</td></tr> </table>	Attributes	This is a bit-mapped field that identifies the type of size conversion to be performed. The field is defined as follows:		<pre>CONVERT_BYTES_TO_DEVSIZE Converts <i>bytes</i> to <i>devsize</i> format. CONVERT_DEVSIZE_TO_BYTES Converts <i>devsize</i> format to <i>bytes</i>.</pre>	bytes	If CONVERT_BYTES_TO_DEVSIZE is set, the value in the bytes field is converted to a <i>devsize</i> format and returned in the devsize field.	devsize	If CONVERT_DEVSIZE_TO_BYTES is set, the value in the devsize field is converted to a <i>bytes</i> value and returned in the bytes field.
Attributes	This is a bit-mapped field that identifies the type of size conversion to be performed. The field is defined as follows:								
	<pre>CONVERT_BYTES_TO_DEVSIZE Converts <i>bytes</i> to <i>devsize</i> format. CONVERT_DEVSIZE_TO_BYTES Converts <i>devsize</i> format to <i>bytes</i>.</pre>								
bytes	If CONVERT_BYTES_TO_DEVSIZE is set, the value in the bytes field is converted to a <i>devsize</i> format and returned in the devsize field.								
devsize	If CONVERT_DEVSIZE_TO_BYTES is set, the value in the devsize field is converted to a <i>bytes</i> value and returned in the bytes field.								
RETURN VALUES	<table> <tr> <td>CS_SUCCESS</td><td>Successful operation.</td></tr> <tr> <td>CS_BAD_SIZE</td><td>Invalid <i>bytes</i> or <i>devsize</i>.</td></tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td><td>No PCMCIA hardware installed.</td></tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_SIZE	Invalid <i>bytes</i> or <i>devsize</i> .	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.		
CS_SUCCESS	Successful operation.								
CS_BAD_SIZE	Invalid <i>bytes</i> or <i>devsize</i> .								
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								

NAME	csx_ConvertSpeed – convert device speeds								
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ConvertSpeed(convert_speed_t *cs);</pre>								
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)								
PARAMETERS	cs Pointer to a convert_speed_t structure.								
DESCRIPTION	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.								
STRUCTURE MEMBERS	The structure members of convert_speed_t are: <pre>uint32_t Attributes; uint32_t nS; uint32_t devspeed;</pre> The fields are defined as follows: <table> <tr> <td>Attributes</td><td>This is a bit-mapped field that identifies the type of speed conversion to be performed. The field is defined as follows:</td></tr> <tr> <td></td><td> <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> </td></tr> <tr> <td>nS</td><td>If CONVERT_NS_TO_DEVSPEED is set, the value in the nS field is converted to a <i>devspeed</i> format and returned in the devspeed field.</td></tr> <tr> <td>devspeed</td><td>If CONVERT_DEVSPEED_TO_NS is set, the value in the devspeed field is converted to an <i>nS</i> value and returned in the nS field.</td></tr> </table>	Attributes	This is a bit-mapped field that identifies the type of speed conversion to be performed. The field is defined as follows:		<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>	nS	If CONVERT_NS_TO_DEVSPEED is set, the value in the nS field is converted to a <i>devspeed</i> format and returned in the devspeed field.	devspeed	If CONVERT_DEVSPEED_TO_NS is set, the value in the devspeed field is converted to an <i>nS</i> value and returned in the nS field.
Attributes	This is a bit-mapped field that identifies the type of speed conversion to be performed. The field is defined as follows:								
	<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>								
nS	If CONVERT_NS_TO_DEVSPEED is set, the value in the nS field is converted to a <i>devspeed</i> format and returned in the devspeed field.								
devspeed	If CONVERT_DEVSPEED_TO_NS is set, the value in the devspeed field is converted to an <i>nS</i> value and returned in the nS field.								
RETURN VALUES	<table> <tr> <td>CS_SUCCESS</td><td>Successful operation.</td></tr> <tr> <td>CS_BAD_SPEED</td><td>Invalid <i>nS</i> or <i>devspeed</i>.</td></tr> <tr> <td>CS_BAD_ATTRIBUTE</td><td>Bad Attributes value.</td></tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td><td>No PCMCIA hardware installed.</td></tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_SPEED	Invalid <i>nS</i> or <i>devspeed</i> .	CS_BAD_ATTRIBUTE	Bad Attributes value.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.								
CS_BAD_SPEED	Invalid <i>nS</i> or <i>devspeed</i> .								
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) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>								

csx_CS_DDI_Info(9F)

NAME	csx_CS_DDI_Info – obtain DDI information						
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_CS_DDI_Info(cs_ddi_info_t *cdi);</pre>						
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)						
PARAMETERS	<i>cdi</i> Pointer to a <i>cs_ddi_info_t</i> structure.						
DESCRIPTION	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.						
STRUCTURE MEMBERS	The structure members of <i>cs_ddi_info_t</i> are: <pre>uint32_t Socket; /* socket number */ char* driver_name; /* unique driver name */ dev_info_t *dip; /* dip */ int32_t instance; /* instance */</pre> The fields are defined as follows: Socket This field must be set to the physical socket number that the client is interested in getting information about. driver_name This field must be set to a string containing the name of the client driver to get information about. If <i>csx_CS_DDI_Info()</i> is used in a client's <i>xx_getinfo</i> function, then the client will typically extract the <i>Socket</i> value from the <i>*arg</i> argument and it <i>must</i> set the <i>driver_name</i> field to the same string used with <i>csx_RegisterClient</i>(9F). If the <i>driver_name</i> is found on the <i>Socket</i>, the <i>csx_CS_DDI_Info()</i> function returns both the <i>dev_info</i> pointer and the <i>instance</i> fields for the requested driver instance.						
RETURN VALUES	<table> <tr> <td>CS_SUCCESS</td><td>Successful operation.</td></tr> <tr> <td>CS_BAD_SOCKET</td><td>Client not found on <i>Socket</i>.</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 <i>Socket</i> .	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.						
CS_BAD_SOCKET	Client not found on <i>Socket</i> .						
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.						
CONTEXT	This function may be called from user or kernel context.						
EXAMPLES	EXAMPLE 1 : Using <i>csx_CS_DDI_Info</i> The following example shows how a client might call the <i>csx_CS_DDI_Info()</i> in the client's <i>xx_getinfo</i> function to return the <i>dip</i> or the instance number: <pre>static int pcepp_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg, void **result) { int error = DDI_SUCCESS;</pre>						

EXAMPLE 1 : Using csx_CS_DDI_Info (Continued)

```

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

csx_DeregisterClient(9F)

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

NAME	csx_DupHandle – duplicate access handle																				
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_DupHandle(acc_handle_t <i>handle1</i>, acc_handle_t *<i>handle2</i>, uint32_t <i>flags</i>);</pre>																				
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)																				
PARAMETERS	<i>handle1</i> The access handle returned from <code>csx_RequestIO(9F)</code> or <code>csx_RequestWindow(9F)</code> that is to be duplicated. <i>handle2</i> A pointer to the newly-created duplicated data access handle. <i>flags</i> The access attributes that will be applied to the new handle.																				
DESCRIPTION	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. Both handles must be explicitly freed when they are no longer necessary. The <i>flags</i> argument is bit-mapped. The following bits are defined: <table> <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> 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 <i>flags</i> bits. Only one of the following bits may be specified: <table> <tr> <td>WIN_ACC_STRICT_ORDER</td><td>The data references must be issued by a CPU in program order. Strict ordering is the default behavior.</td></tr> <tr> <td>WIN_ACC_UNORDERED_OK</td><td>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).</td></tr> </table>	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_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_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_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).																				

csx_DupHandle(9F)

	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) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

NAME	csx_Error2Text – convert error return codes to text strings				
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Error2Text(error2text_t *er);</pre>				
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)				
PARAMETERS	<i>er</i> Pointer to an error2text_t structure.				
DESCRIPTION	This function is a Solaris-specific extension that provides a method for clients to convert Card Services error return codes to text strings.				
STRUCTURE MEMBERS	The structure members of error2text_t are: <pre>uint32_t item; /*the error code*/ char test[CS_ERROR_MAX_BUFSIZE]; /*the error code*/</pre> A pointer to the text for the Card Services error return code in the <i>item</i> field is returned in the <i>text</i> 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 <i>item</i> field is not found, the <i>text</i> field will be set to a string of the form: <i>"{unknown Card Services return code}"</i>				
RETURN VALUES	<table> <tr> <td>CS_SUCCESS</td><td>Successful operation.</td></tr> <tr> <td>CS_UNSUPPORTED_FUNCTION</td><td>No PCMCIA hardware installed.</td></tr> </table>	CS_SUCCESS	Successful operation.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.				
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.				
CONTEXT	This function may be called from user or kernel context.				
EXAMPLES	EXAMPLE 1 : Using the csxError2Text function <pre>if ((ret = csx_RegisterClient(&client_handle, & client_reg)) != CS_SUCCESS) { error2text_t error2text; error2text.item = ret; csx_Error2Text(&error2text); cmn_err(CE_CONT, "RegisterClient failed %s (0x%x)", error2text.text, ret); }</pre>				
SEE ALSO	csx_Event2Text(9F) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>				

csx_Event2Text(9F)

NAME	csx_Event2Text – convert events to text strings				
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Event2Text(event2text_t *ev);</pre>				
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)				
PARAMETERS	<i>ev</i> Pointer to an event2text_t structure.				
DESCRIPTION	This function is a Solaris-specific extension that provides a method for clients to convert Card Services events to text strings.				
STRUCTURE MEMBERS	The structure members of event2text_t are: <pre>event_t event; /*the event code*/ char text[CS_EVENT_MAX_BUFSIZE] /*the event code*/</pre> The fields are defined as follows: <table><tr><td>event</td><td>The text for the event code in the event field is returned in the text field.</td></tr><tr><td>text</td><td>The text string describing the name of the event.</td></tr></table>	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.
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.				
RETURN VALUES	<table><tr><td>CS_SUCCESS</td><td>Successful operation.</td></tr><tr><td>CS_UNSUPPORTED_FUNCTION</td><td>No PCMCIA hardware installed.</td></tr></table>	CS_SUCCESS	Successful operation.	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.
CS_SUCCESS	Successful operation.				
CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.				
CONTEXT	This function may be called from user or kernel context.				
EXAMPLES	EXAMPLE 1: Using csx_Event2Text() <pre>xx_event(event_t event, int priority, event_callback_args_t *eca) { event2text_t event2text; event2text.event = event; csx_Event2Text(&event2text); cmn_err(CE_CONT, "event %s (0x%x)", event2text.text, (int)event); }</pre>				
SEE ALSO	csx_event_handler(9E), csx_Error2Text(9F) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>				

NAME	csx_FreeHandle – free access handle
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_FreeHandle(acc_handle_t *handle);</pre>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>handle</i> The access handle returned from <code>csx_RequestIO(9F)</code>, <code>csx_RequestWindow(9F)</code>, or <code>csx_DupHandle(9F)</code>.
DESCRIPTION	This function frees the handle, <i>handle</i>. If the handle was created by the <code>csx_DupHandle(9F)</code> function, this function will free the storage associated with this handle, but will not modify any resources that the original handle refers to. If the handle was created by a common access setup function, this function will release the resources associated with this handle.
RETURN VALUES	<code>CS_SUCCESS</code> Successful operation. <code>CS_UNSUPPORTED_FUNCTION</code> No PCMCIA hardware installed.
CONTEXT	This function may be called from user or kernel context.
SEE ALSO	<code>csx_DupHandle(9F)</code>, <code>csx_RequestIO(9F)</code>, <code>csx_RequestWindow(9F)</code> <i>PC Card95 Standard</i>, <i>PCMCIA/JEIDA</i>

csx_Get8(9F)

NAME	csx_Get8, csx_Get16, csx_Get32, csx_Get64 – read data from device address
SYNOPSIS	<pre>#include <sys/pccard.h> 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>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>handle</i> The access handle returned from csx_RequestIO(9F), csx_RequestWindow(9F), or csx_DupHandle(9F). <i>offset</i> The offset in bytes from the base of the mapped resource.
DESCRIPTION	These functions generate a read of various sizes from the mapped memory or device register. The csx_Get8(), csx_Get16(), csx_Get32(), and csx_Get64() 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>. 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.
RETURN VALUES	These functions return the value read from the mapped address.
CONTEXT	These functions may be called from user, kernel, or interrupt context.
SEE ALSO	csx_DupHandle(9F), csx_GetMappedAddr(9F), csx_Put8(9F), csx_RepGet8(9F), csx_RepPut8(9F), csx_RequestIO(9F), csx_RequestWindow(9F) <i>PC Card 95 Standard</i>, <i>PCMCIA/JEIDA</i>

NAME	csx_GetFirstClient, csx_GetNextClient – return first or next client												
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_GetFirstClient(get_firstnext_client_t *fnc); int32_t csx_GetNextClient(get_firstnext_client_t *fnc);</pre>												
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)												
PARAMETERS	<i>fnc</i> Pointer to a get_firstnext_client_t structure.												
DESCRIPTION	The functions csx_GetFirstClient() and csx_GetNextClient() return information about the first or subsequent PC cards, respectively, that are installed in the system.												
STRUCTURE MEMBERS	The structure members of get_firstnext_client_t are: <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> The fields are defined as follows: <table> <tr> <td>Socket</td><td>If the CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY attribute is set, return information only on the PC card installed in this socket.</td></tr> <tr> <td>Attributes</td><td>This field indicates the type of client. The field is bit-mapped; the following bits are defined:</td></tr> <tr> <td></td><td>CS_GET_FIRSTNEXT_CLIENT_ALL_CLIENTS Return information on all clients.</td></tr> <tr> <td></td><td>CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY Return client information for the specified socket only.</td></tr> <tr> <td>client_handle</td><td>The client handle of the PC card driver is returned in this field.</td></tr> <tr> <td>num_clients</td><td>The number of clients is returned in this field.</td></tr> </table>	Socket	If the CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY attribute is set, return information only on the PC card installed in this socket.	Attributes	This field indicates the type of client. The field is bit-mapped; the following bits are defined:		CS_GET_FIRSTNEXT_CLIENT_ALL_CLIENTS Return information on all clients.		CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY Return client information for the specified socket only.	client_handle	The client handle of the PC card driver is returned in this field.	num_clients	The number of clients is returned in this field.
Socket	If the CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY attribute is set, return information only on the PC card installed in this socket.												
Attributes	This field indicates the type of client. The field is bit-mapped; the following bits are defined:												
	CS_GET_FIRSTNEXT_CLIENT_ALL_CLIENTS Return information on all clients.												
	CS_GET_FIRSTNEXT_CLIENT_SOCKET_ONLY Return client information for the specified socket only.												
client_handle	The client handle of the PC card driver is returned in this field.												
num_clients	The number of clients is returned in this field.												
RETURN VALUES	<table> <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_SOCKET</td><td>Socket number is invalid.</td></tr> <tr> <td>CS_NO_CARD</td><td>No PC Card in socket.</td></tr> <tr> <td>CS_NO_MORE_ITEMS</td><td>PC Card driver does not handle the CS_EVENT_CLIENT_INFO event.</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_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.
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)												

`csx_GetFirstClient(9F)`

PC Card 95 Standard, PCMCIA/JEIDA

NAME	csx_GetFirstTuple, csx_GetNextTuple – return Card Information Structure tuple
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_GetFirstTuple(client_handle_t ch, tuple_t *tu); int32_t csx_GetNextTuple(client_handle_t ch, tuple_t *tu);</pre>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure.
DESCRIPTION	The functions csx_GetFirstTuple() and csx_GetNextTuple() return the first and next tuple, respectively, of the specified type in the Card Information Structure (CIS) for the specified socket.
STRUCTURE MEMBERS	The structure members of tuple_t are: <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> 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 This field is bit-mapped. The following bits are defined: TUPLE_RETURN_LINK Return link tuples if set. The following are link tuples and will only be returned by this function if the TUPLE_RETURN_LINK bit in the Attributes field is set: <pre>CISTPL_NULL CISTPL_LOGLINK_MFC CISTPL_LOGLINK_A CISTPL_LINKTARGET CISTPL_LOGLINK_C CISTPL_NO_LINK CISTPL_LOGLINK_CB CISTPL_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,

csx_GetFirstTuple(9F)

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*

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

csx_GetMappedAddr(9F)

NAME	csx_GetMappedAddr – return mapped virtual address	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_GetMappedAddr(acc_handle_t <i>handle</i>, void **<i>addr</i>);</pre>	
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)	
PARAMETERS	<i>handle</i>	The access handle returned from csx_RequestIO(9F), csx_RequestWindow(9F), or csx_DupHandle(9F).
	<i>addr</i>	The virtual or I/O port number represented by the handle.
DESCRIPTION	This function returns the mapped virtual address or the mapped I/O port number represented by the handle, <i>handle</i> .	
RETURN VALUES	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.
CONTEXT	This function may be called from user, kernel, or interrupt context.	
SEE ALSO	csx_DupHandle(9F), csx_Get8(9F), csx_Put8(9F), csx_RepGet8(9F), csx_RepPut8(9F), csx_RequestIO(9F), csx_RequestWindow(9F) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

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

csx_GetStatus(9F)

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:

csx_GetStatus(9F)

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)

`csx_GetStatus(9F)`

PC Card 95 Standard, PCMCIA/JEIDA

NAME	csx_GetTupleData – return the data portion of a tuple													
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_GetTupleData(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>);</pre>													
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)													
PARAMETERS	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).												
	<i>tu</i>	Pointer to a tuple_t structure.												
DESCRIPTION	This function returns the data portion of a tuple, as returned by the csx_GetFirstTuple(9F) and csx_GetNextTuple(9F) functions.													
STRUCTURE MEMBERS	The structure members of tuple_t are: The fields are defined as follows: <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> <table><tr><td>Socket</td><td>Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</td></tr><tr><td>Attributes</td><td>Initialized by csx_GetFirstTuple(9F) or csx_GetNextTuple(9F); the client must not modify the value in this field.</td></tr><tr><td>DesiredTuple</td><td>Initialized by csx_GetFirstTuple(9F) or csx_GetNextTuple(9F); the client must not modify the value in this field.</td></tr><tr><td>TupleOffset</td><td>This field allows partial tuple information to be retrieved, starting anywhere within the tuple.</td></tr><tr><td>TupleDataMax</td><td>This field is the size of the tuple data buffer that Card Services uses to return raw tuple data from csx_GetTupleData(9F). It can be larger than the number of bytes in the tuple data body. Card Services ignores any value placed here by the client.</td></tr><tr><td>TupleDataLen</td><td>This field is the actual size of the tuple data body. It represents the number of tuple data body bytes returned.</td></tr></table>		Socket	Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.	Attributes	Initialized by csx_GetFirstTuple(9F) or csx_GetNextTuple(9F); the client must not modify the value in this field.	DesiredTuple	Initialized by csx_GetFirstTuple(9F) or csx_GetNextTuple(9F); the client must not modify the value in this field.	TupleOffset	This field allows partial tuple information to be retrieved, starting anywhere within the tuple.	TupleDataMax	This field is the size of the tuple data buffer that Card Services uses to return raw tuple data from csx_GetTupleData(9F). 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.
Socket	Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.													
Attributes	Initialized by csx_GetFirstTuple(9F) or csx_GetNextTuple(9F); the client must not modify the value in this field.													
DesiredTuple	Initialized by csx_GetFirstTuple(9F) or csx_GetNextTuple(9F); the client must not modify the value in this field.													
TupleOffset	This field allows partial tuple information to be retrieved, starting anywhere within the tuple.													
TupleDataMax	This field is the size of the tuple data buffer that Card Services uses to return raw tuple data from csx_GetTupleData(9F). 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.													

csx_GetTupleData(9F)

RETURN VALUES	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.
	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	<code>csx_GetFirstTuple(9F)</code> , <code>csx_ParseTuple(9F)</code> , <code>csx_RegisterClient(9F)</code> , <code>csx_ValidateCIS(9F)</code> , <code>tuple(9S)</code> <i>PC Card 95 Standard</i> , <i>PCMCIA/JEIDA</i>	

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

csx_MakeDeviceNode(9F)

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.
<code>CS_BAD_ATTRIBUTE</code>	The value of one or more arguments is invalid.
<code>CS_BAD_ARGS</code>	Action is invalid.
<code>CS_OUT_OF_RESOURCE</code>	Unable to create or remove device node.
<code>CS_UNSUPPORTED_FUNCTION</code>	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

NAME	csx_MapLogSocket – return the physical socket number associated with the client handle							
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_MapLogSocket(client_handle_t <i>ch</i>, map_log_socket_t *<i>ls</i>) ;</pre>							
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)							
PARAMETERS	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).						
	<i>ls</i>	Pointer to a map_log_socket_t structure.						
DESCRIPTION	This function returns the physical socket number associated with the client handle.							
STRUCTURE MEMBERS	The structure members of map_log_socket_t are: <pre>uint32_t LogSocket; /* logical socket number */ uint32_t PhyAdapter; /* physical adapter number */ uint32_t PhySocket; /* physical socket number */</pre> The fields are defined as follows: <table><tr><td>LogSocket</td><td>Not used by this implementation of Card Services and can be set to any arbitrary value.</td></tr><tr><td>PhyAdapter</td><td>Returns the physical adapter number, which is always 0 in the Solaris implementation of Card Services.</td></tr><tr><td>PhySocket</td><td>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.</td></tr></table>		LogSocket	Not used by this implementation of Card Services and can be set to any arbitrary value.	PhyAdapter	Returns the physical adapter number, which is always 0 in the Solaris implementation of Card Services.	PhySocket	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.
LogSocket	Not used by this implementation of Card Services and can be set to any arbitrary value.							
PhyAdapter	Returns the physical adapter number, which is always 0 in the Solaris implementation of Card Services.							
PhySocket	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.							
RETURN VALUES	CS_SUCCESS	Successful operation.						
	CS_BAD_HANDLE	Client handle is invalid.						
	CS_UNSUPPORTED_FUNCTION	No PCMCIA hardware installed.						
CONTEXT	This function may be called from user or kernel context.							
SEE ALSO	csx_RegisterClient(9F)							
	<i>PC Card 95 Standard, PCMCIA/JEIDA</i>							

csx_MapMemPage(9F)

NAME	csx_MapMemPage – map the memory area on a PC Card	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_MapMemPage(window_handle_t <i>wh</i>, map_mem_page_t *<i>mp</i>);</pre>	
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)	
PARAMETERS	<i>wh</i>	Window handle returned from csx_RequestWindow(9F).
	<i>mp</i>	Pointer to a map_mem_page_t structure.
DESCRIPTION	This function maps the memory area on a PC Card into a page of a window allocated with the csx_RequestWindow(9F) function.	
STRUCTURE MEMBERS	The structure members of map_mem_page_t are: <pre>uint32_t CardOffset; /* card offset */ uint32_t Page; /* page number */</pre> The fields are defined as follows: CardOffset The absolute offset in bytes from the beginning of the PC Card to map into system memory. Page Used internally by Card Services; clients must set this field to 0 before calling this function.	
RETURN VALUES	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.
CONTEXT	This function may be called from user or kernel context.	
SEE ALSO	csx_ModifyWindow(9F), csx_ReleaseWindow(9F), csx_RequestWindow(9F) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

NAME	csx_ModifyConfiguration – modify socket and PC Card Configuration Register
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ModifyConfiguration(client_handle_t <i>ch</i>, modify_config_t *<i>mc</i>);</pre>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>mc</i> Pointer to a modify_config_t structure.
DESCRIPTION	This function allows a socket and PC Card configuration to be modified. This function can only modify a configuration requested via csx_RequestConfiguration(9F).
STRUCTURE MEMBERS	The structure members of modify_config_t are: <pre>uint32_t Socket; /* socket number */ uint32_t Attributes; /* attributes to modify */ uint32_t Vpp1; /* Vpp1 value */ uint32_t Vpp2; /* Vpp2 value */</pre> 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 This field is bit-mapped. The following bits are defined: CONF_ENABLE_IRQ_STEERING Enable IRQ steering. Set to connect the PC Card IREQ line to a previously selected system interrupt. CONF_IRQ_CHANGE_VALID IRQ change valid. Set to request the IRQ steering enable to be changed. CONF_VPP1_CHANGE_VALID Vpp1 change valid. These bits are set to request a change to the corresponding voltage level for the PC Card. 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.)

csx_ModifyConfiguration(9F)

	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 <i>PC Card 95 Standard</i> , <i>PCMCIA/JEIDA</i> , systems must always support 5.0 volts for both Vcc and Vpp. (See WARNINGS.)
RETURN VALUES	CS_SUCCESS Successful operation. CS_BAD_HANDLE Client handle is invalid or csx_RequestConfiguration(9F) not done. CS_BAD_SOCKET Error getting/setting socket hardware parameters. CS_BAD_VPP Requested Vpp 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)	
	<i>PC Card 95 Standard, PCMCIA/JEIDA</i>	
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 Vpp to 0 volts may result in the loss of a PC Card's state. Any client setting Vpp 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). Vcc can not be changed using csx_ModifyConfiguration(). Vcc may be changed by first invoking csx_ReleaseConfiguration(9F), followed by csx_RequestConfiguration(9F) with a new Vcc value.	

NAME	csx_ModifyWindow – modify window attributes				
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ModifyWindow(window_handle_t wh, modify_win_t *mw);</pre>				
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)				
PARAMETERS	<i>wh</i> Window handle returned from csx_RequestWindow(9F). <i>mw</i> Pointer to a modify_win_t structure.				
DESCRIPTION	This function modifies the attributes of a window allocated by the csx_RequestWindow(9F) function. Only some of the window attributes or the access speed field may be modified by this request. The csx_MapMemPage(9F) function is also used to set the offset into PC Card memory to be mapped into system memory for paged windows. The csx_RequestWindow(9F) and csx_ReleaseWindow(9F) functions must be used to change the window base or size.				
STRUCTURE MEMBERS	The structure members of modify_win_t are: <pre>uint32_t Attributes; /* window flags */ uint32_t AccessSpeed; /* window access speed */</pre> The fields are defined as follows: Attributes This field is bit-mapped and defined as follows: WIN_MEMORY_TYPE_CM Window points to Common Memory area. Set this to map the window to Common Memory. WIN_MEMORY_TYPE_AM Window points to Attribute Memory area. Set this to map the window to Attribute Memory. WIN_ENABLE Enable Window. The client must set this to enable the window. 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 <i>PC Card 95 Standard</i>), the lower bits are a binary code representing a speed from the list below. Numbers in the first column are codes; items in the second column are speeds. <table> <tr> <td>0</td><td>Reserved: do not use</td></tr> <tr> <td>1</td><td>250 nsec</td></tr> </table>	0	Reserved: do not use	1	250 nsec
0	Reserved: do not use				
1	250 nsec				

csx_ModifyWindow(9F)

2	200 nsec
3	150 nsec
4	100 nsec
5 - 7	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

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

NAME	csx_Parse_CISTPL_BATTERY – parse the Battery Replacement Date tuple													
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_BATTERY(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_battery_t *<i>cb</i>);</pre>													
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)													
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).												
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).												
	<i>cb</i>	Pointer to a cistpl_battery_t structure which contains the parsed CISTPL_BATTERY tuple information upon return from this function.												
DESCRIPTION	This function parses the Battery Replacement Date tuple, CISTPL_BATTERY, into a form usable by PC Card drivers. 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.													
STRUCTURE MEMBERS	The structure members of cistpl_battery_t are: <pre>uint32_t rday; /* date battery last replaced */ uint32_t xday; /* date battery due for replacement */</pre> The fields are defined as follows: <table><tr><td>rday</td><td>This field indicates the date on which the battery was last replaced.</td></tr><tr><td>xday</td><td>This field indicates the date on which the battery should be replaced.</td></tr></table>		rday	This field indicates the date on which the battery was last replaced.	xday	This field indicates the date on which the battery should be replaced.								
rday	This field indicates the date on which the battery was last replaced.													
xday	This field indicates the date on which the battery should be replaced.													
RETURN VALUES	<table><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													

csx_Parse_CISTPL_BYTEORDER(9F)

NAME	csx_Parse_CISTPL_BYTEORDER – parse the Byte Order tuple														
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_BYTEORDER(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_byteorder_t *<i>cbo</i>);</pre>														
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>cbo</i> Pointer to a cistpl_byteorder_t structure which contains the parsed CISTPL_BYTEORDER tuple information upon return from this function.														
DESCRIPTION	This function parses the Byte Order tuple, CISTPL_BYTEORDER, into a form usable by PC Card drivers. 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.														
STRUCTURE MEMBERS	The structure members of cistpl_byteorder_t are: <pre>uint32_t order; /* byte order code */ uint32_t map; /* byte mapping code */</pre> The fields are defined as follows: <table> <tr> <td>order</td><td>This field specifies the byte order for multi-byte numeric data.</td></tr> <tr> <td>TPLBYTEORD_LOW</td><td>Little endian order</td></tr> <tr> <td>TPLBYTEORD_VS</td><td>Vendor specific</td></tr> <tr> <td>map</td><td>This field specifies the byte mapping for 16-bit or wider cards.</td></tr> <tr> <td>TPLBYTEMAP_LOW</td><td>Byte zero is least significant byte</td></tr> <tr> <td>TPLBYTEMAP_HIGH</td><td>Byte zero is most significant byte</td></tr> <tr> <td>TPLBYTEMAP_VS</td><td>Vendor specific mapping</td></tr> </table>	order	This field specifies the byte order for multi-byte numeric data.	TPLBYTEORD_LOW	Little endian order	TPLBYTEORD_VS	Vendor specific	map	This field specifies the byte mapping for 16-bit or wider cards.	TPLBYTEMAP_LOW	Byte zero is least significant byte	TPLBYTEMAP_HIGH	Byte zero is most significant byte	TPLBYTEMAP_VS	Vendor specific mapping
order	This field specifies the byte order for multi-byte numeric data.														
TPLBYTEORD_LOW	Little endian order														
TPLBYTEORD_VS	Vendor specific														
map	This field specifies the byte mapping for 16-bit or wider cards.														
TPLBYTEMAP_LOW	Byte zero is least significant byte														
TPLBYTEMAP_HIGH	Byte zero is most significant byte														
TPLBYTEMAP_VS	Vendor specific mapping														
RETURN VALUES	<table> <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.														

		csx_Parse_CISTPL_BYTEORDER(9F)
	CS_NO_CIS	No Card Information Structure (CIS) 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)	
	<i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

csx_Parse_CISTPL_CFTABLE_ENTRY(9F)

NAME	csx_Parse_CISTPL_CFTABLE_ENTRY – parse 16-bit Card Configuration Table Entry tuple
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_CFTABLE_ENTRY(client_handle_t ch, tuple_t *tu, cistpl_cftable_entry_t *cft);</pre>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>cft</i> Pointer to a cistpl_cftable_entry_t structure which contains the parsed CISTPL_CFTABLE_ENTRY tuple information upon return from this function.
DESCRIPTION	This function parses the 16 bit Card Configuration Table Entry tuple, CISTPL_CFTABLE_ENTRY, into a form usable by PC Card drivers. 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.
STRUCTURE MEMBERS	The structure members of cistpl_cftable_entry_t are: <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> The flags field is defined and bit-mapped as follows: CISTPL_CFTABLE_TPCE_DEFAULT This is a default configuration CISTPL_CFTABLE_TPCE_IF If configuration byte exists CISTPL_CFTABLE_TPCE_FS_PWR Power information exists CISTPL_CFTABLE_TPCE_FS_TD Timing information exists

csx_Parse_CISTPL_CFTABLE_ENTRY(9F)

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.

csx_Parse_CISTPL_CFTABLE_ENTRY(9F)

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_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_PEAKI`

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

csx_Parse_CISTPL_CFTABLE_ENTRY(9F)

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

csx_Parse_CISTPL_CFTABLE_ENTRY(9F)

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

csx_Parse_CISTPL_CONFIG(9F)

NAME	csx_Parse_CISTPL_CONFIG – parse Configuration tuple
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_CONFIG(client_handle_t ch, tuple_t *tu, cistpl_config_t *cc);</pre>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>cc</i> Pointer to a cistpl_config_t structure which contains the parsed CISTPL_CONFIG tuple information upon return from this function.
DESCRIPTION	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.
STRUCTURE MEMBERS	The structure members of cistpl_config_t are: <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> The fields are defined as follows: present This field indicates which configuration registers are present on the PC Card. CONFIG_OPTION_REG_PRESENT Configuration Option Register present CONFIG_STATUS_REG_PRESENT Configuration Status Register present 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

csx_Parse_CISTPL_CONFIG(9F)

	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 <code>regs</code> 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	<code>csx_GetFirstTuple(9F)</code>, <code>csx_GetTupleData(9F)</code>, <code>csx_Parse_CISTPL_CFTABLE_ENTRY(9F)</code>, <code>csx_RegisterClient(9F)</code>, <code>csx_ValidateCIS(9F)</code>, <code>tuple(9S)</code> <i>PC Card 95 Standard, PCMCIA/JEIDA</i>

csx_Parse_CISTPL_CONFIG(9F)

NOTES	PC Card drivers should not attempt to use configurations beyond the "last" member in the <code>cistpl_config_t</code> structure.

NAME	csx_Parse_CISTPL_DATE – parse the Card Initialization Date tuple
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_DATE(client_handle_t ch, tuple_t *tu, cistpl_date_t *cd);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>cd</i> Pointer to a cistpl_date_t structure which contains the parsed CISTPL_DATE tuple information upon return from this function.
DESCRIPTION	This function parses the Card Initialization Date tuple, CISTPL_DATE, into a form usable by PC Card drivers. 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.
STRUCTURE MEMBERS	The structure members of cistpl_date_t are: <pre>uint32_t time; uint32_t day;</pre> The fields are defined as follows: <i>time</i> This field indicates the time at which the PC Card was initialized. <i>day</i> This field indicates the date the PC Card was initialized.
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) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>

csx_Parse_CISTPL_DEVICE(9F)

NAME	csx_Parse_CISTPL_DEVICE, csx_Parse_CISTPL_DEVICE_A, csx_Parse_CISTPL_DEVICE_OC, csx_Parse_CISTPL_DEVICE_OA – parse Device Information tuples
SYNOPSIS	<pre>#include <sys/pccard.h> 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>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>cd</i> Pointer to a cistpl_device_t 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.
DESCRIPTION	csx_Parse_CISTPL_DEVICE() and csx_Parse_CISTPL_DEVICE_A() parse the 5 volt Device Information tuples, CISTPL_DEVICE and CISTPL_DEVICE_A, respectively, into a form usable by PC Card drivers. csx_Parse_CISTPL_DEVICE_OC() and csx_Parse_CISTPL_DEVICE_OA() parse the Other Condition Device Information tuples, CISTPL_DEVICE_OC and CISTPL_DEVICE_OA, respectively, into a form usable by PC Card drivers. 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. 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.
STRUCTURE MEMBERS	The structure members of cistpl_device_t are: <pre>uint32_t num_devices; /* number of devices found */ cistpl_device_node_t devnode[CISTPL_DEVICE_MAX_DEVICES];</pre> The structure members of cistpl_device_node_t are:

csx_Parse_CISTPL_DEVICE(9F)

```
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 CISTPL_DEVICE_SPEED_SIZE_IGNORE, 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 CISTPL_DEVICE_SPEED_SIZE_IGNORE, 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:

csx_Parse_CISTPL_DEVICE(9F)

CISTPL_DEVICE_DTYPE_NULL
No device

CISTPL_DEVICE_DTYPE_ROM
Masked ROM

CISTPL_DEVICE_DTYPE_OTPROM
One Time Programmable ROM

CISTPL_DEVICE_DTYPE_EPROM
UV EPROM

CISTPL_DEVICE_DTYPE_EEPROM
EEPROM

CISTPL_DEVICE_DTYPE_FLASH
FLASH

CISTPL_DEVICE_DTYPE_SRAM
Static RAM

CISTPL_DEVICE_DTYPE_DRAM
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

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

csx_Parse_CISTPL_DEVICEGEO(9F)

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

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

csx_Parse_CISTPL_DEVICEGEO_A(9F)

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(9F), csx_RegisterClient(9F), tuple(9S)

PC Card 95 Standard, PCMCIA/JEIDA

NAME	csx_Parse_CISTPL_FORMAT – parse the Data Recording Format tuple														
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_FORMAT(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_format_t *<i>pt</i>);</pre>														
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)														
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>pt</i> Pointer to a cistpl_format_t structure which contains the parsed CISTPL_FORMAT tuple information upon return from this function.														
DESCRIPTION	This function parses the Data Recording Format tuple, CISTPL_FORMAT, into a form usable by PC Card drivers. The CISTPL_FORMAT tuple indicates the data recording format for a device partition.														
STRUCTURE MEMBERS	The structure members of cistpl_format_t are: <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> The fields are defined as follows: <table> <tr> <td>type</td><td>This field indicates the type of device:</td></tr> <tr> <td></td><td>TPLFMTTYPE_DISK disk-like device</td></tr> <tr> <td></td><td>TPLFMTTYPE_MEM memory-like device</td></tr> <tr> <td></td><td>TPLFMTTYPE_VS vendor-specific device</td></tr> <tr> <td>edc_length</td><td>This field indicates the error detection code length.</td></tr> <tr> <td>edc_type</td><td>This field indicates the error detection code type.</td></tr> <tr> <td>offset</td><td>This field indicates the offset of the first byte of data in this partition.</td></tr> </table>	type	This field indicates the type of device:		TPLFMTTYPE_DISK disk-like device		TPLFMTTYPE_MEM memory-like device		TPLFMTTYPE_VS 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.
type	This field indicates the type of device:														
	TPLFMTTYPE_DISK disk-like device														
	TPLFMTTYPE_MEM memory-like device														
	TPLFMTTYPE_VS 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.														

csx_Parse_CISTPL_FORMAT(9F)

	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: TPLFMTFLAGS_ADDR address is valid TPLFMTFLAGS_AUTO 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.
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) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

NAME	csx_Parse_CISTPL_FUNC – parse Function Extension tuple
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_FUNC(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_func_t *<i>cf</i>, uint32_t <i>fid</i>);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>cf</i> Pointer to a cistpl_func_t structure which contains the parsed CISTPL_FUNC tuple information upon return from this function. <i>fid</i> The function ID code to which this CISTPL_FUNC tuple refers. See csx_Parse_CISTPL_FUNCID(9F).
DESCRIPTION	This function parses the Function Extension tuple, CISTPL_FUNC, into a form usable by PC Card drivers. The CISTPL_FUNC tuple is used to describe information about a specific PCCard 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.
STRUCTURE MEMBERS	The structure members of cistpl_func_t are: <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 */ 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; }</pre>

csx_Parse_CISTPL_FUNC(9F)

```

    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;
        uint32_t  sr[16];      /* voice sampling rates (*100) */
        uint32_t  nss;
        uint32_t  ss[16];      /* voice sample sizes (*10) */
        uint32_t  nsc;
        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_FUNC 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

csx_Parse_CISTPL_FUNCE(9F)

fc	TPLFE_UC_STOP_15	1.5 stop bits supported
	TPLFE_UC_STOP_2	2 stop bits supported
	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	Transparent 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	2400bps (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) TPLFE_MS_V32BIS Up to 14400bps (V.32bis) TPLFE_MS_VFAST Up to 28800 V.FAST
em	This identifies modem error correction/detection protocols and is defined as follows: TPLFE_EM_MNP MNP levels 2-4 TPLFE_EM_V42 CCITT LAPM (V.42)
dc	This identifies modem data compression protocols and is defined as follows: TPLFE_DC_V42BI CCITT compression V.42 TPLFE_DC_MNP5 MNP compression (uses MNP 2, 3 or 4)
cm	This identifies modem command protocols and is defined as follows: TPLFE_CM_AT1 ANSI/EIA/TIA 602 "Action" commands TPLFE_CM_AT2 ANSI/EIA/TIA 602 "ACE/DCE IF Params" TPLFE_CM_AT3 ANSI/EIA/TIA 602 "Ace Parameters" TPLFE_CM_MNP_AT MNP specification AT commands TPLFE_CM_V25BIS V.25bis calling commands TPLFE_CM_V25A V.25bis test procedures TPLFE_CM_DMCL DMCL command mode
ex	This identifies the modem escape mechanism and is defined as follows:

csx_Parse_CISTPL_FUNCE(9F)

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

tech	TPLFE_FS_POLLING	Polling support
	TPLFE_FS_FTP	File transfer support
	TPLFE_FS_PASSWORD	Password support
	This identifies the LAN technology type and is defined as follows:	
	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

csx_Parse_CISTPL_FUNCE(9F)

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

NAME	csx_Parse_CISTPL_FUNCID – parse Function Identification tuple																
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_FUNCID(client_handle_t ch, tuple_t *tu, cistpl_funcid_t *cf);</pre>																
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)																
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>cf</i> Pointer to a cistpl_funcid_t structure which contains the parsed CISTPL_FUNCID tuple information upon return from this function.																
DESCRIPTION	This function parses the Function Identification tuple, CISTPL_FUNCID, into a form usable by PC Card drivers. 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_FUNCCE follow this tuple (see csx_Parse_CISTPL_FUNCCE(9F)).																
STRUCTURE MEMBERS	The structure members of cistpl_funcid_t are: <pre>uint32_t function; /* PC Card function code */ uint32_t sysinit; /* system initialization mask */</pre> The fields are defined as follows: <table> <tr> <td>function</td><td>This is the function type for CISTPL_FUNCID:</td></tr> <tr> <td>TPLFUNC_MULTI</td><td>Vendor-specific multifunction card</td></tr> <tr> <td>TPLFUNC_MEMORY</td><td>Memory card</td></tr> <tr> <td>TPLFUNC_SERIAL</td><td>Serial I/O port</td></tr> <tr> <td>TPLFUNC_PARALLEL</td><td>Parallel printer port</td></tr> <tr> <td>TPLFUNC_FIXED</td><td>Fixed disk, silicon or removable</td></tr> <tr> <td>TPLFUNC_VIDEO</td><td>Video interface</td></tr> <tr> <td>TPLFUNC_LAN</td><td>Local Area Network adapter</td></tr> </table>	function	This is the function type for CISTPL_FUNCID:	TPLFUNC_MULTI	Vendor-specific multifunction card	TPLFUNC_MEMORY	Memory card	TPLFUNC_SERIAL	Serial I/O port	TPLFUNC_PARALLEL	Parallel printer port	TPLFUNC_FIXED	Fixed disk, silicon or removable	TPLFUNC_VIDEO	Video interface	TPLFUNC_LAN	Local Area Network adapter
function	This is the function type for CISTPL_FUNCID:																
TPLFUNC_MULTI	Vendor-specific multifunction card																
TPLFUNC_MEMORY	Memory card																
TPLFUNC_SERIAL	Serial I/O port																
TPLFUNC_PARALLEL	Parallel printer port																
TPLFUNC_FIXED	Fixed disk, silicon or removable																
TPLFUNC_VIDEO	Video interface																
TPLFUNC_LAN	Local Area Network adapter																

csx_Parse_CISTPL_FUNCID(9F)

	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) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>

NAME	csx_Parse_CISTPL_GEOMETRY – parse the Geometry tuple	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_GEOMETRY(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_geometry_t *<i>pt</i>);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)	
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>pt</i>	Pointer to a cistpl_geometry_t structure which contains the parsed CISTPL_GEOMETRY tuple information upon return from this function.
DESCRIPTION	This function parses the Geometry tuple, CISTPL_GEOMETRY, into a form usable by PC Card drivers. The CISTPL_GEOMETRY tuple indicates the geometry of a disk-like device.	
STRUCTURE MEMBERS	The structure members of cistpl_geometry_t are: <pre>uint32_t spt; uint32_t tpc; uint32_t ncyl;</pre> The fields are defined as follows: spt This field indicates the number of sectors per track. tpc This field indicates the number of tracks per cylinder. ncyl This field indicates the number of cylinders.	
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)	
	<i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

csx_Parse_CISTPL_JEDEC_C(9F)

NAME	csx_Parse_CISTPL_JEDEC_C, csx_Parse_CISTPL_JEDEC_A – parse JEDEC Identifier tuples	
SYNOPSIS	<pre>#include <sys/pccard.h> 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>	
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)	
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>cj</i>	Pointer to a cistpl_jedec_t structure which contains the parsed CISTPL_JEDEC_C or CISTPL_JEDEC_A tuple information upon return from these functions, respectively.
DESCRIPTION	csx_Parse_CISTPL_JEDEC_C() and csx_Parse_CISTPL_JEDEC_A() parse the JEDEC Identifier tuples, CISTPL_JEDEC_C and CISTPL_JEDEC_A, respectively, into a form usable by PC Card drivers. 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.	
STRUCTURE MEMBERS	The structure members of cistpl_jedec_t are: <pre>uint32_t nid; /* # of JEDEC identifiers present */ jedec_ident_t jid[CISTPL_JEDEC_MAX_IDENTIFIERS];</pre> The structure members of jedec_ident_t are: <pre>uint32_t id; /* manufacturer id */ uint32_t info; /* manufacturer specific info */</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_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.	

`csx_Parse_CISTPL_JEDEC_C(9F)`

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

csx_Parse_CISTPL_LINKTARGET(9F)

NAME	csx_Parse_CISTPL_LINKTARGET – parse the Link Target tuple												
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_LINKTARGET(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_linktarget_t *<i>pt</i>);</pre>												
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>pt</i> Pointer to a cistpl_linktarget_t structure which contains the parsed CISTPL_LINKTARGET tuple information upon return from this function.												
DESCRIPTION	This function parses the Link Target tuple, CISTPL_LINKTARGET, into a form usable by PCCard drivers. 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.												
STRUCTURE MEMBERS	The structure members of cistpl_linktarget_t are: <pre>uint32_t length; char tpltg_tag[CIS_MAX_TUPLE_DATA_LEN];</pre> The fields are defined as follows: length This field indicates the number of bytes in tpltg_tag. tpltg_tag This field provides the Link Target tuple information.												
RETURN VALUES	<table> <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) <i>PC Card 95 Standard</i>, <i>PCMCIA/JEIDA</i>												

csx_Parse_CISTPL_LONGLINK_A(9F)

NAME	csx_Parse_CISTPL_LONGLINK_A, csx_Parse_CISTPL_LONGLINK_C – parse the Long Link A and C tuples									
SYNOPSIS	<pre>#include <sys/pccard.h> 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>									
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)									
PARAMETERS	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).								
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).								
	<i>pt</i>	Pointer to a cistpl_longlink_ac_t structure which contains the parsed CISTPL_LONGLINK_A or CISTPL_LONGLINK_C tuple information upon return from this function.								
DESCRIPTION	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. The CISTPL_LONGLINK_A and CISTPL_LONGLINK_C tuples provide links to Attribute and Common Memory.									
STRUCTURE MEMBERS	The structure members of cistpl_longlink_ac_t are: <pre>uint32_t flags; uint32_t tp11_addr;</pre> The fields are defined as follows: <table><tr><td>flags</td><td>This field indicates the type of memory:</td></tr><tr><td></td><td>CISTPL_LONGLINK_AC_AM long link to Attribute Memory</td></tr><tr><td></td><td>CISTPL_LONGLINK_AC_CM long link to Common Memory</td></tr><tr><td>tp11_addr</td><td>This field provides the offset from the beginning of the specified address space.</td></tr></table>		flags	This field indicates the type of memory:		CISTPL_LONGLINK_AC_AM long link to Attribute Memory		CISTPL_LONGLINK_AC_CM long link to Common Memory	tp11_addr	This field provides the offset from the beginning of the specified address space.
flags	This field indicates the type of memory:									
	CISTPL_LONGLINK_AC_AM long link to Attribute Memory									
	CISTPL_LONGLINK_AC_CM long link to Common Memory									
tp11_addr	This field provides the offset from the beginning of the specified address space.									
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.								

csx_Parse_CISTPL_LONGLINK_A(9F)

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

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

csx_Parse_CISTPL_LONGLINK_MFC(9F)

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*

NAME	csx_Parse_CISTPL_MANFID – parse Manufacturer Identification tuple	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_MANFID(client_handle_t ch, tuple_t *tu, cistpl_manfid_t *cm);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)	
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>cm</i>	Pointer to a cistpl_manfid_t structure which contains the parsed CISTPL_MANFID tuple information upon return from this function.
DESCRIPTION	This function parses the Manufacturer Identification tuple, CISTPL_MANFID, into a form usable by PC Card drivers. 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.	
STRUCTURE MEMBERS	The structure members of cistpl_manfid_t are: <pre>uint32_t manf; /* PCMCIA assigned manufacturer code */ uint32_t card; /* manufacturer information (part number and/or revision) */</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_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) <i>PC Card 95 Standard</i> , PCMCIA/JEIDA	

csx_Parse_CISTPL_ORG(9F)

NAME	csx_Parse_CISTPL_ORG – parse the Data Organization tuple												
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_ORG(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_org_t *<i>pt</i>);</pre>												
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>tu</i> Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F). <i>pt</i> Pointer to a cistpl_org_t structure which contains the parsed CISTPL_ORG tuple information upon return from this function.												
DESCRIPTION	This function parses the Data Organization tuple, CISTPL_ORG, into a form usable by PC Card drivers. The CISTPL_ORG tuple provides a text description of the organization.												
STRUCTURE MEMBERS	The structure members of cistpl_org_t are: <pre>uint32_t type; char desc[CIS_MAX_TUPLE_DATA_LEN];</pre> The fields are defined as follows: <i>type</i> This field indicates type of data organization. <i>desc</i> [CIS_MAX_TUPLE_DATA_LEN] This field provides the text description of this organization.												
RETURN VALUES	<table> <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) <i>PC Card 95 Standard</i>, <i>PCMCIA/JEIDA</i>												

NAME	csx_Parse_CISTPL_SPCL – parse the Special Purpose tuple	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_SPCL(client_handle_t ch, tuple_t *tu, cistpl_spcl_t *csp);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)	
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>csp</i>	Pointer to a cistpl_spcl_t structure which contains the parsed CISTPL_SPCL tuple information upon return from this function.
DESCRIPTION	This function parses the Special Purpose tuple, CISTPL_SPCL, into a form usable by PC Card drivers. 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.	
STRUCTURE MEMBERS	The structure members of cistpl_data_t are: <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> The fields are defined as follows: <i>id</i> 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. <i>seq</i> This field contains a data sequence number. CISTPL_SPCL_SEQ_END is the last tuple in sequence. <i>bytes</i> This field contains the number of data bytes in the data [CIS_MAX_TUPLE_DATA_LEN]. <i>data</i> The data component of this tuple.	
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.

csx_Parse_CISTPL_SPCL(9F)

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

NAME	csx_Parse_CISTPL_SWIL – parse the Software Interleaving tuple	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_SWIL(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cistpl_swil_t *<i>pt</i>);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)	
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>pt</i>	Pointer to a cistpl_swil_t structure which contains the parsed CISTPL_SWIL tuple information upon return from this function.
DESCRIPTION	This function parses the Software Interleaving tuple, CISTPL_SWIL, into a form usable by PC Card drivers. The CISTPL_SWIL tuple provides the software interleaving of data within a partition on the card.	
STRUCTURE MEMBERS	The structure members of cistpl_swil_t are: <pre>uint32_t intrlv;</pre> The fields are defined as follows: intrlv This field provides the software interleaving for a partition.	
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)	
	<i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

csx_Parse_CISTPL_VERS_1(9F)

NAME	csx_Parse_CISTPL_VERS_1 – parse Level-1 Version/Product Information tuple	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_VERS_1(client_handle_t ch, tuple_t *tu, cistpl_vers_1_t *cv1);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)	
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>cv1</i>	Pointer to a cistpl_vers_1_t structure which contains the parsed CISTPL_VERS_1 tuple information upon return from this function.
DESCRIPTION	This function parses the Level-1 Version/Product Information tuple, CISTPL_VERS_1, into a form usable by PC Card drivers. The CISTPL_VERS_1 tuple is used to describe the card Level-1 version compliance and card manufacturer information.	
STRUCTURE MEMBERS	The structure members of cistpl_vers_1_t are: <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>	
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) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

NAME	csx_Parse_CISTPL_VERS_2 – parse Level-2 Version and Information tuple	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_Parse_CISTPL_VERS_2(client_handle_t ch, tuple_t *tu, cistpl_vers_2_t *cv2);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI)	
	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>cv2</i>	Pointer to a cistpl_vers_2_t structure which contains the parsed CISTPL_VERS_2 tuple information upon return from this function.
DESCRIPTION	This function parses the Level-2 Version and Information tuple, CISTPL_VERS_2, into a form usable by PC Card drivers. The CISTPL_VERS_2 tuple is used to describe the card Level-2 information which has the logical organization of the card's data.	
STRUCTURE MEMBERS	The structure members of cistpl_vers_2_t are: <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>	
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) <i>PC Card 95 Standard, PCMCIA/JEIDA</i>	

csx_ParseTuple(9F)

NAME	csx_ParseTuple – generic tuple parser	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ParseTuple(client_handle_t <i>ch</i>, tuple_t *<i>tu</i>, cisparsedata_t *<i>cp</i>, cisdata_t <i>cd</i>);</pre>	
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)	
PARAMETERS	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>tu</i>	Pointer to a tuple_t structure (see tuple(9S)) returned by a call to csx_GetFirstTuple(9F) or csx_GetNextTuple(9F).
	<i>cp</i>	Pointer to a cisparsedata_t structure that unifies all tuple parsing structures.
	<i>cd</i>	Extended tuple data for some tuples.
DESCRIPTION	This function is the generic tuple parser entry point.	
STRUCTURE MEMBERS	The structure members of cisparsedata_t are: <pre>typedef union cisparsedata_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_funcname_t cistpl_funcname; 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; } cisparsedata_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.

csx_ParseTuple(9F)

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

csx_Put8(9F)

NAME	csx_Put8, csx_Put16, csx_Put32, csx_Put64 – write to device register
SYNOPSIS	<pre>#include <sys/pccard.h> void csx_Put8(acc_handle_t <i>handle</i>, uint32_t <i>offset</i>, uint8_t <i>value</i>); void csx_Put16(acc_handle_t <i>handle</i>, uint32_t <i>offset</i>, uint16_t <i>value</i>); void csx_Put32(acc_handle_t <i>handle</i>, uint32_t <i>offset</i>, uint32_t <i>value</i>); void csx_Put64(acc_handle_t <i>handle</i>, uint32_t <i>offset</i>, uint64_t <i>value</i>);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>handle</i> The access handle returned from csx_RequestIO(9F), csx_RequestWindow(9F), or csx_DupHandle(9F). <i>offset</i> The offset in bytes from the base of the mapped resource. <i>value</i> The data to be written to the device.
DESCRIPTION	These functions generate a write of various sizes to the mapped memory or device register. The <code>csx_Put8()</code>, <code>csx_Put16()</code>, <code>csx_Put32()</code>, and <code>csx_Put64()</code> 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 <i>offset</i>, <i>offset</i>. 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.
CONTEXT	These functions may be called from user, kernel, or interrupt context.
SEE ALSO	<code>csx_DupHandle(9F)</code>, <code>csx_Get8(9F)</code>, <code>csx_GetMappedAddr(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</i>, <i>PCMCIA/JEIDA</i>

NAME	csx_RegisterClient – register a client
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_RegisterClient(client_handle_t *ch, client_reg_t *cr);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Pointer to a client_handle_t structure. <i>mc</i> Pointer to a client_reg_t structure.
DESCRIPTION	This function registers a client with Card Services and returns a unique client handle for the client. The client handle must be passed to csx_DeregisterClient(9F) when the client terminates.
STRUCTURE MEMBERS	The structure members of client_reg_t are: <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> The fields are defined as follows: Attributes This field is bit-mapped and defined as follows: <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> <pre>INFO_CARD_EXCL Generate artificial CS_EVENT_CARD_INSERTION and CS_EVENT_REGISTRATION_COMPLETE events.</pre> <pre>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.</pre>

csx_RegisterClient(9F)

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.

Event Mask

This field is bit-mapped and specifies the client's global event mask. 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.

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.

csx_RegisterClient(9F)

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

csx_ReleaseConfiguration(9F)

NAME	csx_ReleaseConfiguration – release PC Card and socket configuration
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ReleaseConfiguration(client_handle_t <i>ch</i>, release_config_t *<i>rc</i>);</pre>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>rc</i> Pointer to a release_config_t structure.
DESCRIPTION	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.
STRUCTURE MEMBERS	The structure members of release_config_t are: <pre>uint32_t Socket; /* socket number */</pre> The <code>Socket</code> field is not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.
RETURN VALUES	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.

csx_ReleaseConfiguration(9F)

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

csx_RepGet8(9F)

NAME	csx_RepGet8, csx_RepGet16, csx_RepGet32, csx_RepGet64 – read repetitively from the device register
SYNOPSIS	<pre>#include <sys/pccard.h> 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>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>handle</i> The access handle returned from <code>csx_RequestIO(9F)</code>, <code>csx_RequestWindow(9F)</code>, or <code>csx_DupHandle(9F)</code>. <i>hostaddr</i> Source host address. <i>offset</i> The offset in bytes from the base of the mapped resource. <i>repcount</i> Number of data accesses to perform. <i>flags</i> Device address flags.
DESCRIPTION	These functions generate multiple reads of various sizes from the mapped memory or device register. 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>. 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. 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, when the <i>flags</i> argument is set to <code>CS_DEV_NO_AUTOINCR</code>, 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.

`csx_RepGet8(9F)`

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

csx_RepPut8(9F)

NAME	csx_RepPut8, csx_RepPut16, csx_RepPut32, csx_RepPut64 – write repetitively to the device register
SYNOPSIS	<pre>#include <sys/pccard.h> 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>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>handle</i> The access handle returned from <code>csx_RequestIO(9F)</code>, <code>csx_RequestWindow(9F)</code>, or <code>csx_DupHandle(9F)</code>. <i>hostaddr</i> Source host address. <i>offset</i> The offset in bytes from the base of the mapped resource. <i>repcount</i> Number of data accesses to perform. <i>flags</i> Device address flags.
DESCRIPTION	These functions generate multiple writes of various sizes to the mapped memory or device register. 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>. 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. 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, when the <i>flags</i> argument is set to <code>CS_DEV_NO_AUTOINCR</code>, 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.

`csx_RepPut8(9F)`

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

csx_RequestConfiguration(9F)

NAME	csx_RequestConfiguration – configure the PC Card and socket
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_RequestConfiguration(client_handle_t <i>ch</i>, config_req_t *<i>cr</i>) ;</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>cr</i> Pointer to a config_req_t structure.
DESCRIPTION	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. csx_RequestIO(9F) and csx_RequestIRQ(9F) must be used before calling this function to specify the I/O and IRQ requirements for the PC Card and socket if necessary. csx_RequestConfiguration() 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.
STRUCTURE MEMBERS	The structure members of config_req_t are: <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> 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 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 (CSR) 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

csx_RequestConfiguration(9F)

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.

csx_RequestConfiguration(9F)

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

csx_RequestIO(9F)

NAME	csx_RequestIO, csx_ReleaseIO – request or release I/O resources for the client
SYNOPSIS	<pre>#include <sys/pccard.h> 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>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>ir</i> Pointer to an io_req_t structure.
DESCRIPTION	The functions csx_RequestIO() and csx_ReleaseIO() request or release, respectively, I/O resources for the client. If a client requires I/O resources, csx_RequestIO() must be called to request I/O resources from Card Services; then csx_RequestConfiguration(9F) must be used to establish the configuration. csx_RequestIO() can be called multiple times until a successful set of I/O resources is found. csx_RequestConfiguration(9F) only uses the last configuration specified. csx_RequestIO() fails if it has already been called without a corresponding csx_ReleaseIO(). csx_ReleaseIO() 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.
STRUCTURE MEMBERS	The structure members of io_req_t are: <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> 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.

```
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.

csx_RequestIO(9F)

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*

`csx_RequestIO(9F)`

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 <code>csx_RequestIO()</code> first with any <code>IOAddrLines</code> 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.

NAME	csx_RequestIRQ, csx_ReleaseIRQ – request or release IRQ resource
SYNOPSIS	<pre>#include <sys/pccard.h> 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>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>ir</i> Pointer to an irq_req_t structure.
DESCRIPTION	The function csx_RequestIRQ() requests an IRQ resource and registers the client's IRQ handler with Card Services. If a client requires an IRQ, csx_RequestIRQ() 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 csx_RequestConfiguration(9F) or csx_ModifyConfiguration(9F) has successfully returned when either of these functions are called with the CONF_ENABLE_IRQ_STEERING bit set. The function csx_ReleaseIRQ() releases a previously requested IRQ resource. The Card Services IRQ resource list is adjusted by csx_ReleaseIRQ(). 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. csx_RequestIRQ() fails if it has already been called without a corresponding csx_ReleaseIRQ().
STRUCTURE MEMBERS	The structure members of irq_req_t are: <pre>uint32_t Socket; /* socket number */ uint32_t Attributes; /* IRQ attribute flags */ csfunction_t *irq_handler; /* IRQ handler */ void *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> 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 This field is bit-mapped. It specifies details about the type of IRQ desired by the client. The following bits are defined:

csx_RequestIRQ(9F)

IRQ_TYPE_EXCLUSIVE

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

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 not set, or an unsupported or reserved bit is 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

csx_RequestSocketMask(9F)

NAME	csx_RequestSocketMask, csx_ReleaseSocketMask – set or clear the client's client event mask						
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_RequestSocketMask(client_handle_t <i>ch</i>, request_socket_mask_t *<i>sm</i>); int32_t csx_ReleaseSocketMask(client_handle_t <i>ch</i>, release_socket_mask_t *<i>rm</i>);</pre>						
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>sm</i> Pointer to a request_socket_mask_t structure. <i>rm</i> Pointer to a release_socket_mask_t structure.						
DESCRIPTION	The function csx_RequestSocketMask() 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 csx_RegisterClient(9F) 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. csx_RequestSocketMask() must be used before calling csx_GetEventMask(9F) or csx_SetEventMask(9F) for the client event mask for this socket. The function csx_ReleaseSocketMask() clears the client's client event mask.						
STRUCTURE MEMBERS	The structure members of request_socket_mask_t are: <pre>uint32_t Socket; /* socket number */ uint32_t EventMask; /* event mask to set or return */</pre> The structure members of release_socket_mask_t are: <pre>uint32_t Socket; /* socket number */</pre> The fields are defined as follows: <table> <tr> <td>Socket</td><td>Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</td></tr> <tr> <td>EventMask</td><td>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.</td></tr> </table>	Socket	Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.	EventMask	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.						
EventMask	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.						
RETURN VALUES	<table> <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_IN_USE</td><td>csx_ReleaseSocketMask() has not been done.</td></tr> </table>	CS_SUCCESS	Successful operation.	CS_BAD_HANDLE	Client handle is invalid.	CS_IN_USE	csx_ReleaseSocketMask() has not been done.
CS_SUCCESS	Successful operation.						
CS_BAD_HANDLE	Client handle is invalid.						
CS_IN_USE	csx_ReleaseSocketMask() has not been done.						

		<code>csx_RequestSocketMask(9F)</code>
	<code>CS_BAD_SOCKET</code>	<code>csx_RequestSocketMask()</code> has not been done.
	<code>CS_UNSUPPORTED_FUNCTION</code>	No PCMCIA hardware installed.
CONTEXT	These functions may be called from user or kernel context.	
SEE ALSO	<code>csx_event_handler(9E)</code> , <code>csx_GetEventMask(9F)</code> , <code>csx_RegisterClient(9F)</code> , <code>csx_SetEventMask(9F)</code>	
	<i>PC Card 95 Standard</i> , PCMCIA/JEIDA	

csx_RequestWindow(9F)

NAME	csx_RequestWindow, csx_ReleaseWindow – request or release window resources
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_RequestWindow(client_handle_t <i>ch</i>, window_handle_t *<i>wh</i>, win_req_t *<i>wr</i>); int32_t csx_ReleaseWindow(window_handle_t <i>wh</i>);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>wh</i> Pointer to a window_handle_t structure. <i>wr</i> Pointer to a win_req_t structure.
DESCRIPTION	The function csx_RequestWindow() requests a block of system address space be assigned to a PC Card in a socket. The function csx_ReleaseWindow() releases window resources which were obtained by a call to csx_RequestWindow(). No adapter or socket hardware is modified by this function. The csx_MapMemPage(9F) and csx_ModifyWindow(9F) functions use the window handle returned by csx_RequestWindow(). This window handle must be freed by calling csx_ReleaseWindow() when the client is done using this window. The PC Card Attribute or Common Memory offset for this window is set by csx_MapMemPage(9F).
STRUCTURE MEMBERS	The structure members of win_req_t are: <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 */ /* or 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> 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 This field is bit-mapped. It is defined as follows:

csx_RequestWindow(9F)

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	Points to I/O space.
WIN_MEMORY_TYPE_CM	Points to common memory space.
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	Describes device as 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:
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).

csx_RequestWindow(9F)

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. 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. 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. 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

`CS_SUCCESS`

Successful operation.

`CS_BAD_ATTRIBUTE`

Attributes are invalid.

`CS_BAD_SPEED`

Speed is invalid.

`CS_BAD_HANDLE`

Client handle is invalid.

`CS_BAD_SIZE`

Window size is invalid.

csx_RequestWindow(9F)

CS_NO_CARD

No PC Card in socket.

CS_OUT_OF_RESOURCE

Unable to allocate window.

CS_UNSUPPORTED_FUNCTION

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*

NAME	csx_ResetFunction – reset a function on a PC card	
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ResetFunction(client_handle_t <i>ch</i>, reset_function_t *<i>rf</i>) ;</pre>	
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)	
PARAMETERS	<i>ch</i>	Client handle returned from csx_RegisterClient(9F).
	<i>rf</i>	Pointer to a reset_function_t structure.
DESCRIPTION	csx_ResetFunction() requests that the specified function on the PC card initiate a reset operation.	
STRUCTURE MEMBERS	The structure members of reset_function_t are: <pre>uint32_t Socket; /* socket number */ uint32_t Attributes; /* reset attributes */</pre> 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.	
RETURN VALUES	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.
CONTEXT	This function may be called from user or kernel context.	
SEE ALSO	csx_event_handler(9E), csx_RegisterClient(9F)	
	<i>PC Card 95 Standard</i> , PCMCIA/JEIDA	
NOTES	csx_ResetFunction() has not been implemented in this release and always returns CS_IN_USE.	

csx_SetEventMask(9F)

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

		csx_SetEventMask(9F)
	CS_BAD_SOCKET	csx_RequestSocketMask(9F) not called for CONF_EVENT_MASK_CLIENT.
	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_RegisterClient(9F), csx_ReleaseSocketMask(9F), csx_RequestSocketMask(9F)	
	<i>PC Card 95 Standard</i> , PCMCIA/JEIDA	

csx_SetHandleOffset(9F)

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

NAME	csx_ValidateCIS – validate the Card Information Structure (CIS)						
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t csx_ValidateCIS(client_handle_t <i>ch</i>, cisinfo_t *<i>ci</i>);</pre>						
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)						
PARAMETERS	<i>ch</i> Client handle returned from csx_RegisterClient(9F). <i>ci</i> Pointer to a cisinfo_t structure.						
DESCRIPTION	This function validates the Card Information Structure (CIS) on the PC Card in the specified socket.						
STRUCTURE MEMBERS	The structure members of cisinfo_t are: <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> The fields are defined as follows: <table> <tr> <td>Socket</td><td>Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.</td></tr> <tr> <td>Chains</td><td>This field returns the number of valid tuple chains located in the CIS. If 0 is returned, the CIS is not valid.</td></tr> <tr> <td>Tuples</td><td>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.</td></tr> </table>	Socket	Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.	Chains	This field returns the number of valid tuple chains located in the CIS. If 0 is returned, the CIS is not valid.	Tuples	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.
Socket	Not used in Solaris, but for portability with other Card Services implementations, it should be set to the logical socket number.						
Chains	This field returns the number of valid tuple chains located in the CIS. If 0 is returned, the CIS is not valid.						
Tuples	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.						
RETURN VALUES	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.						
CONTEXT	This function may be called from user or kernel context.						
SEE ALSO	csx_GetFirstTuple(9F), csx_GetTupleData(9F), csx_ParseTuple(9F), csx_RegisterClient(9F) <i>PC Card 95 Standard</i>, PCMCIA/JEIDA						

datamsg(9F)

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

NAME	ddi_add_intr, ddi_get_iblock_cookie, ddi_remove_intr – hardware interrupt handling routines
SYNOPSIS	<pre> #include <sys/types.h> #include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> 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_iddevice_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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	For ddi_get_iblock_cookie(): <i>dip</i> Pointer to dev_info structure. <i>inumber</i> Interrupt number. <i>iblock_cookiep</i> Pointer to an interrupt block cookie. For ddi_add_intr(): <i>dip</i> Pointer to dev_info structure. <i>inumber</i> Interrupt number. <i>iblock_cookiep</i> Optional pointer to an interrupt block cookie where a returned interrupt block cookie is stored. <i>idevice_cookiep</i> Optional pointer to an interrupt device cookie where a returned interrupt device cookie is stored. <i>int_handler</i> Pointer to interrupt handler. <i>int_handler_arg</i> Argument for interrupt handler. For ddi_remove_intr(): <i>dip</i> Pointer to dev_info structure. <i>inumber</i> Interrupt number. <i>iblock_cookie</i> Block cookie which identifies the interrupt handler to be removed.
DESCRIPTION	

ddi_add_intr(9F)

ddi_get_iblock_cookie() `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 for which interrupt specification to retrieve the cookie. *inumber* is associated with information provided either by the device (see `sbus(4)`) or the hardware configuration file (see `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()

`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, `ddi_add_intr()` will return `DDI_SUCCESS`. `DDI_INTR_NOTFOUND` is returned on i86pc and sun4m architectures if the interrupt information cannot be found. If the interrupt information cannot be found on the sun4u architecture, either `DDI_INTR_NOTFOUND` or `DDI_FAILURE` can be returned.

ddi_remove_intr()

`ddi_remove_intr()` removes an interrupt handler from the system. Unloadable drivers should call this routine during their `detach(9E)` routine to remove their interrupt handler from the system.

ddi_add_intr(9F)

	The device interrupt routine for this instance of the device will not execute after <code>ddi_remove_intr()</code> returns. <code>ddi_remove_intr()</code> 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 <code>ddi_remove_intr()</code> or deadlock may result.						
For All Three Functions:	For certain bus types, you can call these DDI functions from a high-interrupt context. These types include ISA, EISA, and SBus buses. See <code>sysbus(4)</code>, <code>isa(4)</code>, <code>eisa(4)</code>, and <code>sbus(4)</code> for details.						
RETURN VALUES	<code>ddi_add_intr()</code> and <code>ddi_get_iblock_cookie()</code> return: <table><tr><td><code>DDI_SUCCESS</code></td><td>On success.</td></tr><tr><td><code>DDI_INTR_NOTFOUND</code></td><td>On failure to find the interrupt.</td></tr><tr><td><code>DDI_FAILURE</code></td><td>On failure. <code>DDI_FAILURE</code> can also be returned on failure to find interrupt (<code>sun4u</code>).</td></tr></table>	<code>DDI_SUCCESS</code>	On success.	<code>DDI_INTR_NOTFOUND</code>	On failure to find the interrupt.	<code>DDI_FAILURE</code>	On failure. <code>DDI_FAILURE</code> can also be returned on failure to find interrupt (<code>sun4u</code>).
<code>DDI_SUCCESS</code>	On success.						
<code>DDI_INTR_NOTFOUND</code>	On failure to find the interrupt.						
<code>DDI_FAILURE</code>	On failure. <code>DDI_FAILURE</code> can also be returned on failure to find interrupt (<code>sun4u</code>).						
CONTEXT	<code>ddi_add_intr()</code>, <code>ddi_remove_intr()</code>, and <code>ddi_get_iblock_cookie()</code> can be called from user or kernel context.						
SEE ALSO	<code>driver.conf(4)</code>, <code>eisa(4)</code>, <code>isa(4)</code>, <code>sbus(4)</code>, <code>sysbus(4)</code>, <code>attach(9E)</code>, <code>detach(9E)</code>, <code>ddi_intr_hilevel(9F)</code>, <code>mutex(9F)</code>, <code>mutex_init(9F)</code>, <code>rw_init(9F)</code>, <code>rwlock(9F)</code>, <code>ddi_idevice_cookie(9S)</code> <i>Writing Device Drivers</i>						
NOTES	<code>ddi_get_iblock_cookie()</code> must not be called <i>after</i> the driver adds an interrupt handler for the interrupt specification corresponding to <i>inumber</i>. All consumers of these interfaces, checking return codes, should verify <code>return_code != DDI_SUCCESS</code>. Checking for specific failure codes can result in inconsistent behaviors among platforms.						
BUGS	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 the SBus and PCI buses are supported and a single data structure is used to describe both.						

ddi_add_softintr(9F)

NAME	ddi_add_softintr, ddi_get_soft_iblock_cookie, ddi_remove_softintr, ddi_trigger_softintr – software interrupt handling routines
SYNOPSIS	<pre> #include <sys/types.h> #include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	ddi_get_soft_iblock_cookie() <i>dip</i> Pointer to a dev_info structure. <i>preference</i> The type of soft interrupt to retrieve the cookie for. <i>iblock_cookiep</i> Pointer to a location to store the interrupt block cookie. ddi_add_softintr() <i>dip</i> Pointer to dev_info structure. <i>preference</i> A hint value describing the type of soft interrupt to generate. <i>idp</i> Pointer to a soft interrupt identifier where a returned soft interrupt identifier is stored. <i>iblock_cookiep</i> Optional pointer to an interrupt block cookie where a returned interrupt block cookie is stored. <i>idevice_cookiep</i> Optional pointer to an interrupt device cookie where a returned interrupt device cookie is stored (not used). <i>int_handler</i> Pointer to interrupt handler. <i>int_handler_arg</i> Argument for interrupt handler. ddi_remove_softintr() <i>id</i> The identifier specifying which soft interrupt handler to remove. ddi_trigger_softintr() <i>id</i> The identifier specifying which soft interrupt to trigger and which soft interrupt handler will be called.

DESCRIPTION	For <code>ddi_get_soft_iblock_cookie()</code>: <code>ddi_get_soft_iblock_cookie()</code> retrieves the interrupt block cookie associated with a particular soft interrupt preference level. This routine should be called before <code>ddi_add_softintr()</code> to retrieve the interrupt block cookie needed to initialize locks (<code>mutex(9F)</code>, <code>rwlock(9F)</code>) 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: <code>DDI_SOFTINT_LOW</code> Low priority soft interrupt. <code>DDI_SOFTINT_MED</code> Medium priority soft interrupt. <code>DDI_SOFTINT_HIGH</code> High priority soft interrupt. On a successful return, <i>iblock_cookiep</i> contains information needed for initializing locks associated with this soft interrupt (see <code>mutex_init(9F)</code> and <code>rw_init(9F)</code>). The driver can then initialize mutexes acquired by the interrupt routine before calling <code>ddi_add_softintr()</code> which prevents a possible race condition where the driver's soft interrupt handler is called immediately <i>after</i> the driver has called <code>ddi_add_softintr()</code> 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. For <code>ddi_add_softintr()</code>: <code>ddi_add_softintr()</code> 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 <code>ddi_get_soft_iblock_cookie()</code>. Refer to the description of <code>ddi_get_soft_iblock_cookie()</code> above. 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 <code>ddi_remove_softintr()</code> and <code>ddi_trigger_softintr()</code> to identify the soft interrupt and the soft interrupt handler. 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 <code>mutex_init(9F)</code> and <code>rw_init(9F)</code>). Note that the interrupt block cookie is normally obtained using <code>ddi_get_soft_iblock_cookie()</code> to avoid the race conditions described above (refer to the description of <code>ddi_get_soft_iblock_cookie()</code> above). For this reason, <i>iblock_cookiep</i> is no longer useful and should be set to <code>NULL</code>. <i>idevice_cookiep</i> is not used and should be set to <code>NULL</code>.

ddi_add_softintr(9F)

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`.

For `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.

For `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)`.

EXAMPLE 1 device using high-level interrupts (Continued)

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      id;
    ddi_iblock_cookie_t high_iblock_cookie;
    kmutex_t            high_mutex;
    ddi_iblock_cookie_t low_iblock_cookie;
    kmutex_t            low_mutex;
    int                 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)
{
    struct xxstate *xsp;
    . . .
    /* get high-level iblock cookie */
    if (ddi_get_iblock_cookie(dip, inumber,
        &xsp->high_iblock_cookie) != DDI_SUCCESS) {
        /* clean up */
        return (DDI_FAILURE); /* fail attach */
    }

    /* initialize high-level mutex */
    mutex_init(&xsp->high_mutex, "xx high mutex", MUTEX_DRIVER,
        (void *)xsp->high_iblock_cookie);

    /* add high-level routine - xxhighintr() */
    if (ddi_add_intr(dip, inumber, NULL, NULL,
        xxhighintr, (caddr_t) xsp) != DDI_SUCCESS) {
        /* cleanup */
        return (DDI_FAILURE); /* fail attach */
    }
}
```

ddi_add_softintr(9F)

EXAMPLE 2 sample attach() routine (Continued)

```

    }

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

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

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

    . . .
}

```

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;
    int need_softint;
    . . .
    mutex_enter(&xsp->high_mutex);
    /*
     * Verify this device generated the interrupt
     * and disable the device interrupt.
     * Enqueue data for xxsoftintr() processing.
     */

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

    /* read-only access to xsp->id, no mutex needed */
    if (need_softint)

```


EXAMPLE 3 High-level interrupt routine *(Continued)*

```

        ddi_trigger_softintr(xsp->id);
        . . .
        return (DDI_INTR_CLAIMED);
    }

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

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

    xsp->softint_running = 1;

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

        /* de-queue data */

        mutex_exit(&xsp->high_mutex);

        /* Process data on queue */

        mutex_enter(&xsp->high_mutex);

    }

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

    return (DDI_INTR_CLAIMED);
}

```

SEE ALSO ddi_add_intr(9F), ddi_in_panic(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

`ddi_add_softintr(9F)`

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()`.

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

ddi_btop(9F)

NAME	ddi_btop, ddi_btopr, ddi_ptob – page size conversions
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> unsigned long ddi_btop(dev_info_t *dip, unsigned long bytes); unsigned long ddi_btopr(dev_info_t *dip, unsigned long bytes); unsigned long ddi_ptob(dev_info_t *dip, unsigned long pages);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	This set of routines use the parent nexus driver to perform conversions in page size units. ddi_btop() 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. ddi_btopr() 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. ddi_ptob() converts the given number of pages to the number of bytes that it corresponds to. Because bus nexus may possess their own hardware address translation facilities, these routines should be used in preference to the corresponding DDI/DKI routines btop(9F), btopr(9F), and ptob(9F), which only deal in terms of the pagesize of the main system MMU.
RETURN VALUES	ddi_btop() and ddi_btopr() return the number of corresponding pages. ddi_ptob() returns the corresponding number of bytes. There are no error return values.
CONTEXT	This function can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Find the size (in bytes) of one page <pre>pagesize = ddi_ptob(dip, 1L);</pre>
SEE ALSO	btop(9F), btopr(9F), ptob(9F) <i>Writing Device Drivers</i>

ddi_check_acc_handle(9F)

NAME	ddi_check_acc_handle, ddi_check_dma_handle – Check data access and DMA handles				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_check_acc_handle(ddi_acc_handle_t acc_handle); int ddi_check_dma_handle(ddi_dma_handle_t dma_handle);</pre>				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)				
PARAMETERS	<table><tr><td><i>acc_handle</i></td><td>Data access handle obtained from a previous call to ddi_regs_map_setup(9F), ddi_dma_mem_alloc(9F), or similar function.</td></tr><tr><td><i>dma_handle</i></td><td>DMA handle obtained from a previous call to ddi_dma_setup(9F) or one of its derivatives.</td></tr></table>	<i>acc_handle</i>	Data access handle obtained from a previous call to ddi_regs_map_setup(9F), ddi_dma_mem_alloc(9F), or similar function.	<i>dma_handle</i>	DMA handle obtained from a previous call to ddi_dma_setup(9F) or one of its derivatives.
<i>acc_handle</i>	Data access handle obtained from a previous call to ddi_regs_map_setup(9F), ddi_dma_mem_alloc(9F), or similar function.				
<i>dma_handle</i>	DMA handle obtained from a previous call to ddi_dma_setup(9F) or one of its derivatives.				
DESCRIPTION	The ddi_check_acc_handle() and ddi_check_dma_handle() functions check for faults that can interfere with communication between a driver and the device it controls. Each function checks a single handle of a specific type and returns a status value indicating whether faults affecting the resource mapped by the supplied handle have been detected. If a fault is indicated when checking a data access handle, this implies that the driver is no longer able to access the mapped registers or memory using programmed I/O through that handle. Typically, this might occur after the device has failed to respond to an I/O access (for example, has incurred a bus error or timed out). The effect of programmed I/O accesses made after this happens is undefined; for example, read accesses (for example, ddi_get8(9F)) may return random values, and write accesses (for example, ddi_put8(9F)) may or may not have any effect. This type of fault is normally fatal to the operation of the device, and the driver should report it via ddi_dev_report_fault(9F) specifying DDI_SERVICE_LOST for the impact, and DDI_DATAPATH_FAULT for the location. If a fault is indicated when checking a DMA handle, it implies that a fault has been detected that has (or will) affect DMA transactions between the device and the memory currently bound to the handle (or most recently bound, if the handle is currently unbound). Possible causes include the failure of a component in the DMA data path, or an attempt by the device to make an invalid DMA access. The driver may be able to continue by falling back to a non-DMA mode of operation, but in general, DMA faults are non-recoverable. The contents of the memory currently (or previously) bound to the handle should be regarded as indeterminate. The fault indication associated with the current transaction is lost once the handle is (re-)bound, but because the fault may persist, future DMA operations may not succeed. Note – Some implementations cannot detect all types of failure. If a fault is not indicated, this does not constitute a guarantee that communication is possible. However, if a check fails, this is a positive indication that a problem <i>does</i> exist with respect to communication using that handle.				

ddi_check_acc_handle(9F)

RETURN VALUES The ddi_check_acc_handle() and ddi_check_dma_handle() functions return DDI_SUCCESS if no faults affecting the supplied handle are detected and DDI_FAILURE if any fault affecting the supplied handle is detected.

EXAMPLES

```
static int
xxattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    \\&...
    /* This driver uses only a single register-access handle */
    status = ddi_regs_map_setup(dip, REGSET_ZERO, &regaddr,
                                0, 0, , &acc_attrs, &acc_hdl);

    if (status != DDI_SUCCESS)
        return (DDI_FAILURE);
    \\&...
}

static int
xxread(dev_t dev, struct uio *uio_p, cred_t *cred_p)
{
    \\&...
    if (ddi_check_acc_handle(acc_hdl) != DDI_SUCCESS) {
        ddi_dev_report_fault(dip, DDI_SERVICE_LOST,
                            DDI_DATAPATH_FAULT, "register access fault during read");
        return (EIO);
    }
    \\&...
```

CONTEXT The ddi_check_acc_handle() and ddi_check_dma_handle() functions may be called from user, kernel, or interrupt context.

SEE ALSO ddi_regs_map_setup(9F), ddi_dma_setup(9F), ddi_dev_report_fault(9F), ddi_get8(9F), ddi_put8(9F)

NAME	ddi_copyin – copy data to a driver buffer
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_copyin(const void *buf, void *driverbuf, size_t cn, int flags);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>buf</i> Source address from which data is transferred. <i>driverbuf</i> Driver destination address to which data is transferred. <i>cn</i> Number of bytes transferred. <i>flags</i> Set of flag bits that provide address space information about <i>buf</i>.
DESCRIPTION	This routine is designed for use in driver <code>ioctl(9E)</code> routines for drivers that support layered <code>ioctls</code>. <code>ddi_copyin()</code> copies data from a source address to a driver buffer. The driver developer must ensure that adequate space is allocated for the destination address. The <i>flags</i> argument determines 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_copyin()</code> behaves like <code>bcopy(9F)</code>. Otherwise, <i>buf</i> is interpreted as a user buffer address, and <code>ddi_copyin()</code> behaves like <code>copyin(9F)</code>. 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.
RETURN VALUES	<code>ddi_copyin()</code> returns 0, indicating a successful copy. It returns -1 if one of the following occurs: ■ Paging fault; the driver tried to access a page of memory for which it did not have read or write access. ■ Invalid user address, such as a user area or stack area. ■ Invalid address that would have resulted in data being copied into the user block. ■ Hardware fault; a hardware error prevented access to the specified user memory. For example, an uncorrectable parity or ECC error occurred. If -1 is returned to the caller, driver entry point routines should return <code>EFAULT</code>.
CONTEXT	<code>ddi_copyin()</code> can be called from user or kernel context only.

ddi_copyin(9F)

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 }
```


EXAMPLE 1 ddi_copyin() example *(Continued)*

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.

ddi_copyin() should not be used from a streams driver. See M_COPYIN and M_COPYOUT in *STREAMS Programming Guide*.

ddi_copyout(9F)

NAME	ddi_copyout – copy data from a driver								
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_copyout(const void *driverbuf, void *buf, size_t cn, int flags) ;</pre>								
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).								
PARAMETERS	<table><tr><td><i>driverbuf</i></td><td>Source address in the driver from which the data is transferred.</td></tr><tr><td><i>buf</i></td><td>Destination address to which the data is transferred.</td></tr><tr><td><i>cn</i></td><td>Number of bytes to copy.</td></tr><tr><td><i>flags</i></td><td>Set of flag bits that provide address space information about <i>buf</i>.</td></tr></table>	<i>driverbuf</i>	Source address in the driver from which the data is transferred.	<i>buf</i>	Destination address to which the data is transferred.	<i>cn</i>	Number of bytes to copy.	<i>flags</i>	Set of flag bits that provide address space information about <i>buf</i> .
<i>driverbuf</i>	Source address in the driver from which the data is transferred.								
<i>buf</i>	Destination address to which the data is transferred.								
<i>cn</i>	Number of bytes to copy.								
<i>flags</i>	Set of flag bits that provide address space information about <i>buf</i> .								
DESCRIPTION	This routine is designed for use in driver <code>ioctl(9E)</code> routines for drivers that support layered ioctls. <code>ddi_copyout()</code> copies data from a driver buffer to a destination address, <i>buf</i>. The <i>flags</i> argument determines 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>. 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.								
RETURN VALUES	Under normal conditions, 0 is returned to indicate a successful copy. Otherwise, -1 is returned if one of the following occurs: ■ Paging fault; the driver tried to access a page of memory for which it did not have read or write access. ■ Invalid user address, such as a user area or stack area. ■ Invalid address that would have resulted in data being copied into the user block. ■ Hardware fault; a hardware error prevented access to the specified user memory. For example, an uncorrectable parity or ECC error occurred. If -1 is returned to the caller, driver entry point routines should return <code>EFAULT</code>.								
CONTEXT	<code>ddi_copyout()</code> 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 };

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

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

20     instance = getminor(dev);
21     sp = ddi_get_soft_state(statep, instance);
22     if (sp == NULL)
23         return (ENXIO);
24     rp = sp->regsp;
25     . . .
26     switch (cmd) {

27     case XX_GETREGS: /* copy registers to arg */
28         mutex_enter(&sp->reg_mutex);
29         /*
30          * Copy data from device registers to
31          * temporary device register buffer
32          * e.g. reg_buf.control = rp->control;
33          */
34         mutex_exit(&sp->reg_mutex);
35         if (ddi_copyout(&reg_buf, arg,
36             sizeof (struct device), mode) != 0) {
37             return (EFAULT);
38         }
39     }

40     break;
41 }

```

ddi_copyout(9F)

EXAMPLE 1 ddi_copyout () example *(Continued)*

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.

ddi_copyout () should not be used from a streams driver. See M_COPYIN and M_COPYOUT in *STREAMS Programming Guide*.

NAME	ddi_create_minor_node – create a minor node for this device																																						
SYNOPSIS	<pre>#include <sys/stat.h> #include <sys/sunddi.h> int ddi_create_minor_node(dev_info_t *dip, char *name, int spec_type, minor_t minor_num, char *node_type, int flag);</pre>																																						
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>dip</i></td><td>A pointer to the device's dev_info structure.</td></tr> <tr> <td><i>name</i></td><td>The name of this particular minor device.</td></tr> <tr> <td><i>spec_type</i></td><td>S_IFCHR or S_IFBLK for character or block minor devices respectively.</td></tr> <tr> <td><i>minor_num</i></td><td>The minor number for this particular minor device.</td></tr> <tr> <td><i>node_type</i></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></td><td>DDI_NT_SERIAL For serial ports</td></tr> <tr> <td></td><td>DDI_NT_SERIAL_MB For on board serial ports</td></tr> <tr> <td></td><td>DDI_NT_SERIAL_DO For dial out ports</td></tr> <tr> <td></td><td>DDI_NT_SERIAL_MB_DO For on board dial out ports</td></tr> <tr> <td></td><td>DDI_NT_BLOCK For hard disks</td></tr> <tr> <td></td><td>DDI_NT_BLOCK_CHAN For hard disks with channel or target numbers</td></tr> <tr> <td></td><td>DDI_NT_CD For CDROM drives</td></tr> <tr> <td></td><td>DDI_NT_CD_CHAN For CDROM drives with channel or target numbers</td></tr> <tr> <td></td><td>DDI_NT_FD For floppy disks</td></tr> <tr> <td></td><td>DDI_NT_TAPE For tape drives</td></tr> <tr> <td></td><td>DDI_NT_NET For DLPI style 1 or style 2 network devices</td></tr> <tr> <td></td><td>DDI_NT_DISPLAY For display devices</td></tr> <tr> <td></td><td>DDI_PSEUDO For pseudo devices</td></tr> <tr> <td><i>flag</i></td><td>If the device is a clone device then this flag is set to CLONE_DEV else it is set to 0.</td></tr> </table>	<i>dip</i>	A pointer to the device's dev_info structure.	<i>name</i>	The name of this particular minor device.	<i>spec_type</i>	S_IFCHR or S_IFBLK for character or block minor devices respectively.	<i>minor_num</i>	The minor number for this particular minor device.	<i>node_type</i>	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 DLPI style 1 or style 2 network devices		DDI_NT_DISPLAY For display devices		DDI_PSEUDO For pseudo devices	<i>flag</i>	If the device is a clone device then this flag is set to CLONE_DEV else it is set to 0.
<i>dip</i>	A pointer to the device's dev_info structure.																																						
<i>name</i>	The name of this particular minor device.																																						
<i>spec_type</i>	S_IFCHR or S_IFBLK for character or block minor devices respectively.																																						
<i>minor_num</i>	The minor number for this particular minor device.																																						
<i>node_type</i>	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 DLPI style 1 or style 2 network devices																																						
	DDI_NT_DISPLAY For display devices																																						
	DDI_PSEUDO For pseudo devices																																						
<i>flag</i>	If the device is a clone device then this flag is set to CLONE_DEV else it is set to 0.																																						
DESCRIPTION	ddi_create_minor_node() 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																																						

ddi_create_minor_node(9F)

	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 <code>/dev</code> hierarchy that refers to the names in the <code>/devices</code> hierarchy. See <code>disks(1M)</code>, <code>ports(1M)</code>, <code>tapes(1M)</code>, <code>devlinks(1M)</code>. Finally <i>flag</i> determines if this is a clone device or not, and what device class the node belongs to.				
RETURN VALUES	<code>ddi_create_minor_node()</code> returns: <table><tr><td>DDI_SUCCESS</td><td>Was able to allocate memory, create the minor data structure, and place it into the linked list of minor devices for this driver.</td></tr><tr><td>DDI_FAILURE</td><td>Minor node creation failed.</td></tr></table>	DDI_SUCCESS	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	Minor node creation failed.
DDI_SUCCESS	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	Minor node creation failed.				
EXAMPLES	EXAMPLE 1 Create Data Structure Describing 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 <code>DDI_NT_BLOCK</code> (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>				
SEE ALSO	<code>add_drv(1M)</code>, <code>devlinks(1M)</code>, <code>disks(1M)</code>, <code>drvconfig(1M)</code>, <code>ports(1M)</code>, <code>tapes(1M)</code>, <code>attach(9E)</code>, <code>ddi_remove_minor_node(9F)</code> <i>Writing Device Drivers</i>				
NOTES	If the driver is for a network device (<i>node_type</i> <code>DDI_NT_NET</code>), note that the driver name will undergo the driver name constraints identified in the NOTES section of <code>dlpi(7P)</code>. Additionally, the minor name must match the driver name for a DLPI style 2 provider. If the driver is a DLPI style 1 provider, the minor name must also match the driver name with the exception that the ppa is appended to the minor name.				
WARNING	Drivers must remove references to <code>GLOBAL_DEV</code>, <code>NODEBOUND_DEV</code>, <code>NODESPECIFIC_DEV</code>, and <code>ENUMERATED_DEV</code> to compile under Solaris 10 and later versions.				

NAME	ddi_device_copy – copy data from one device register to another device register																								
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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 bytcount, uint_t dev_datsz);</pre>																								
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>src_handle</i></td><td>The data access handle of the source device.</td></tr> <tr> <td><i>src_addr</i></td><td>Base data source address.</td></tr> <tr> <td><i>src_advcnt</i></td><td>Number of <i>dev_datsz</i> units to advance on every access.</td></tr> <tr> <td><i>dest_handle</i></td><td>The data access handle of the destination device.</td></tr> <tr> <td><i>dest_addr</i></td><td>Base data destination address.</td></tr> <tr> <td><i>dest_advcnt</i></td><td>Number of <i>dev_datsz</i> units to advance on every access.</td></tr> <tr> <td><i>bytcount</i></td><td>Number of bytes to transfer.</td></tr> <tr> <td><i>dev_datsz</i></td><td>The size of each data word. Possible values are defined as:</td></tr> <tr> <td></td><td>DDI_DATA_SZ01_ACC 1 byte data size</td></tr> <tr> <td></td><td>DDI_DATA_SZ02_ACC 2 bytes data size</td></tr> <tr> <td></td><td>DDI_DATA_SZ04_ACC 4 bytes data size</td></tr> <tr> <td></td><td>DDI_DATA_SZ08_ACC 8 bytes data size</td></tr> </table>	<i>src_handle</i>	The data access handle of the source device.	<i>src_addr</i>	Base data source address.	<i>src_advcnt</i>	Number of <i>dev_datsz</i> units to advance on every access.	<i>dest_handle</i>	The data access handle of the destination device.	<i>dest_addr</i>	Base data destination address.	<i>dest_advcnt</i>	Number of <i>dev_datsz</i> units to advance on every access.	<i>bytcount</i>	Number of bytes to transfer.	<i>dev_datsz</i>	The size of each data word. Possible values are defined as:		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
<i>src_handle</i>	The data access handle of the source device.																								
<i>src_addr</i>	Base data source address.																								
<i>src_advcnt</i>	Number of <i>dev_datsz</i> units to advance on every access.																								
<i>dest_handle</i>	The data access handle of the destination device.																								
<i>dest_addr</i>	Base data destination address.																								
<i>dest_advcnt</i>	Number of <i>dev_datsz</i> units to advance on every access.																								
<i>bytcount</i>	Number of bytes to transfer.																								
<i>dev_datsz</i>	The size of each data word. Possible values are defined as:																								
	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																								
DESCRIPTION	<code>ddi_device_copy()</code> copies <i>bytcount</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. 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. 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 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 <i>dev_datsz</i> 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	<code>ddi_device_copy()</code> returns: <table> <tr> <td>DDI_SUCCESS</td><td>Successfully transferred the data.</td></tr> </table>	DDI_SUCCESS	Successfully transferred the data.																						
DDI_SUCCESS	Successfully transferred the data.																								

ddi_device_copy(9F)

DDI_FAILURE The byte count is not a multiple *dev_datasz*.

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

NAME	ddi_device_zero – zero fill the device								
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>								
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>handle</i> The data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>dev_addr</i> Beginning of the device address. <i>bytecount</i> Number of bytes to zero. <i>dev_advcnt</i> Number of <i>dev_datsz</i> units to advance on every access. <i>dev_datsz</i> The size of each data word. Possible values are defined as: <table> <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								
DESCRIPTION	ddi_device_zero() function fills the given, <i>bytecount</i>, number of byte of zeroes to the device register or memory. The <i>dev_advcnt</i> argument determines the value of the device address, <i>dev_addr</i>, on each access. A value of 0 will use the same device address, <i>dev_addr</i>, 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 <i>dev_datsz</i> units. The <i>dev_datsz</i> argument determines the size of data word on each access.								
RETURN VALUES	ddi_device_zero() returns: DDI_SUCCESS Successfully zeroed the data. DDI_FAILURE The byte count is not a multiple of <i>dev_datsz</i>.								
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) <i>Writing Device Drivers</i>								

ddi_devid_compare(9F)

NAME	ddi_devid_compare, ddi_devid_free, ddi_devid_init, ddi_devid_register, ddi_devid_sizeof, ddi_devid_str_decode, ddi_devid_str_encode, ddi_devid_str_free, ddi_devid_unregister, ddi_devid_valid – kernel interfaces for device ids	
SYNOPSIS	<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); int ddi_devid_str_decode(char *devidstr, ddi_devid_t *retdevid, char **retminor_name); int ddi_devid_str_encode(ddi_devid_t devid, char *minor_name); int ddi_devid_str_free(char *devidstr); void ddi_devid_unregister(dev_info_t *dip); int ddi_devid_valid(ddi_devid_t devid);</pre>	
PARAMETERS	<i>devid</i>	The device id address.
	<i>devidstr</i>	The <i>devid</i> and <i>minor_name</i> represented as a string.
	<i>devid1</i>	The first of two device id addresses to be compared calling ddi_devid_compare().
	<i>devid2</i>	The second of two device id addresses to be compared calling ddi_devid_compare().
	<i>dip</i>	A dev_info pointer, which identifies the device.
	<i>devid_type</i>	The following device id types may be accepted by the ddi_devid_init() 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.
	<i>minor_name</i>	The minor name to be encoded.

INTERFACE LEVEL DESCRIPTION		ddi_devid_compare(9F)
	<i>nbytes</i>	The length in bytes of device ID.
	<i>retdevid</i>	The return address of the device ID.
	<i>retminor_name</i>	The return address of a minor name. Free string with <code>ddi_devid_str_free()</code> .
INTERFACE LEVEL DESCRIPTION	Solaris DDI specific (Solaris DDI).	
	The following routines are used to provide unique identifiers, device IDs, 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.	
	<code>ddi_devid_compare()</code> compares two device IDs byte-by-byte and determines both equality and sort order.	
	<code>ddi_devid_sizeof()</code> returns the number of bytes allocated for the passed in device ID (<i>devid</i>).	
	<code>ddi_devid_init()</code> allocates memory and initializes the opaque device ID structure. This function does not store the <i>devid</i> . If the device id is not derived from the device's firmware, it is the driver's responsibility to store the <i>devid</i> on some reliable store. When a <i>devid_type</i> of either <code>DEVID_SCSI3_WWN</code> , <code>DEVID_SCSI_SERIAL</code> , or <code>DEVID_ENCAP</code> is accepted, an array of bytes (<i>id</i>) must be passed in (<i>nbytes</i>).	
	When the <i>devid_type</i> <code>DEVID_FAB</code> is used, the array of bytes (<i>id</i>) must be NULL and the length (<i>nbytes</i>) must be zero. The fabricated device ids, <code>DEVID_FAB</code> will be initialized with the machine's host id and a timestamp.	
	Drivers must free the memory allocated by this function, using the <code>ddi_devid_free()</code> function.	
	<code>ddi_devid_free()</code> frees the memory allocated for the returned <i>devid</i> by the <code>ddi_devid_init()</code> and <code>devid_str_decode()</code> functions.	
	<code>ddi_devid_register()</code> registers the device ID address (<i>devid</i>) with the DDI framework, associating it with the <code>dev_info</code> passed in (<i>dip</i>). The drivers must register device IDs at attach time. See <code>attach(9E)</code> .	
	<code>ddi_devid_unregister()</code> removes the device ID address from the <code>dev_info</code> passed in (<i>dip</i>). 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 <code>ddi_devid_free()</code> function. See <code>detach(9E)</code> .	
	<code>ddi_devid_valid()</code> validates the device ID (<i>devid</i>) passed in. The driver must use this function to validate any fabricated device ID that has been stored on a device.	

ddi_devid_compare(9F)

The `ddi_devid_str_encode()` function encodes a *devid* and *minor_name* into a null-terminated ASCII string, returning a pointer to that string. If both a *devid* and a *minor_name* are non-null, then a slash (/) is used to separate the *devid* from the *minor_name* in the encoded string. If *minor_name* is null, then only the *devid* is encoded. If the *devid* is null, then the special string `id0` is returned. Note that you cannot compare the returned string against another string with `strcmp()` to determine *devid* equality. The returned string must be freed by calling `devid_str_free()`.

The `ddi_devid_str_decode()` function takes a string previously produced by the `devid_str_encode(3DEVID)` or `ddi_devid_str_encode()` function and decodes the contained device ID and *minor_name*, allocating and returning pointers to the extracted parts through the *retdevid* and *retminor_name* arguments. If the special *devidstr* `id0` was specified then the returned device ID and minor name will both be null. A non-null returned *devid* must be freed by the caller through the `ddi_devid_free()` function. A non-null returned minor name must be freed by calling `ddi_devid_str_free()`.

The `ddi_devid_str_free()` function is used to free all strings returned by the `ddi_devid` functions (the `ddi_devid_str_encode()` function return value and the returned *retminor_name* argument).

RETURN VALUES

`ddi_devid_init()` returns the following values:

DDI_SUCCESS	Success.
DDI_FAILURE	Out of memory. An invalid <i>devid_type</i> 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 first device ID is less than the second device ID.
0	The first device ID is equal to the second device ID.
1	The first device ID is greater than the second device ID.

ddi_devid_compare(9F)

ddi_devid_sizeof() returns the size of the *devid* in bytes. If called with a null, then the number of bytes that must be allocated and initialized to determine the size of a complete device ID is returned.

ddi_devid_str_encode() returns a value of null to indicate failure. Failure can be caused by attempting to encode an invalid *devid*. If the return value is non-null then the caller must free the returned string by using the *devid_str_free()* function.

ddi_devid_str_decode() returns the following values:

DDI_SUCCESS

Success.

DDI_FAILURE

Failure; the *devidstr* string was not valid.

CONTEXT These functions can be called from a user or kernel context.

SEE ALSO *devid_get*(3DEVID), *libdevid*(3LIB), *attributes*(5), *attach*(9E), *detach*(9E), *kmem_free*(9F)

Writing Device Drivers

ddi_dev_is_needed(9F)

NAME	ddi_dev_is_needed – inform the system that a device's component is required
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dev_is_needed(dev_info_t *dip, int component, int level);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>dip</i> Pointer to the device's dev_info structure. <i>component</i> Component of the driver which is needed. <i>level</i> Power level at which the component is needed.
DESCRIPTION	The ddi_dev_is_needed() function is obsolete and will be removed in a future release. It is recommended that device drivers use pm_raise_power(9F) and pm_lower_power(9F). The ddi_dev_is_needed() function informs the system that a device component is needed at the specified power level. The <i>level</i> argument must be non-zero. This function sets a <i>component</i> to the required level and sets all devices which depend on this to their normal power levels. If <i>component</i> 0 of a device using original Power Management interfaces (calls pm_create_components(9F)) is at power level 0, the ddi_dev_is_needed() call will result in component 0 being returned to normal power and the device being resumed via attach(9E) before ddi_dev_is_needed() returns. The state of the device should be examined before each physical access. The ddi_dev_is_needed() function should be called to set a <i>component</i> 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. The ddi_dev_is_needed() function might cause re-entry of the driver. Deadlock may result if driver locks are held across the call to ddi_dev_is_needed().
RETURN VALUES	The ddi_dev_is_needed() function returns: DDI_SUCCESS Power successfully set to the requested level. DDI_FAILURE An error occurred.
EXAMPLES	EXAMPLE 1 disk driver code A hypothetical disk driver might include this code: <pre>static int xxdisk_spun_down(struct xxstate *xsp) { return (xsp->power_level[DISK_COMPONENT] < POWER_SPUN_UP); } static int xxdisk_strategy(struct buf *bp)</pre>

EXAMPLE 1 disk driver code (Continued)

```
{
. . .

    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.
- ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Obsolete

SEE ALSO pm(7D), pm-components(9P), attach(9E), detach(9E), power(9E), pm_busy_components(9F), pm_create_components(9F), pm_destroy_components(9F), pm_idle_component(9F)

Writing Device Drivers

ddi_dev_is_sid(9F)

NAME	ddi_dev_is_sid – tell whether a device is self-identifying
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dev_is_sid(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dip</i> A pointer to the device's dev_info structure.
DESCRIPTION	ddi_dev_is_sid() 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).
RETURN VALUES	DDI_SUCCESS Device is self-identifying. DDI_FAILURE Device is not self-identifying.
CONTEXT	ddi_dev_is_sid() can be called from user or interrupt context.
EXAMPLES	<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 && 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) <i>Writing Device Drivers</i>

NAME	ddi_dev_nintrs – return the number of interrupt specifications a device has
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dev_nintrs(dev_info_t *dip, int *resultp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_dev_nintrs() returns the number of interrupt specifications a device has in *resultp.
RETURN VALUES	ddi_dev_nintrs() 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.
CONTEXT	ddi_dev_nintrs() can be called from user or interrupt context.
SEE ALSO	isa(4), sbus(4), ddi_add_intr(9F), ddi_dev_nregs(9F), ddi_dev_regsize(9F) <i>Writing Device Drivers</i>

ddi_dev_nregs(9F)

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

NAME	ddi_dev_regsize – return the size of a device’s register						
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dev_regsize(dev_info_t *dip, uint_t rnumber, off_t *resultp);</pre>						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).						
PARAMETERS	<table> <tr> <td><i>dip</i></td><td>A pointer to the device’s dev_info structure.</td></tr> <tr> <td><i>rnumber</i></td><td>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).</td></tr> <tr> <td><i>resultp</i></td><td>Pointer to an integer that holds the size, in bytes, of the described register (if it exists).</td></tr> </table>	<i>dip</i>	A pointer to the device’s dev_info structure.	<i>rnumber</i>	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).	<i>resultp</i>	Pointer to an integer that holds the size, in bytes, of the described register (if it exists).
<i>dip</i>	A pointer to the device’s dev_info structure.						
<i>rnumber</i>	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).						
<i>resultp</i>	Pointer to an integer that holds the size, in bytes, of the described register (if it exists).						
DESCRIPTION	ddi_dev_regsize() 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.						
RETURN VALUES	ddi_dev_regsize() returns: <table> <tr> <td>DDI_SUCCESS</td><td>A successful return. The size, in bytes, of the specified register, is set in <i>resultp</i>.</td></tr> <tr> <td>DDI_FAILURE</td><td>An invalid (nonexistent) register number was specified.</td></tr> </table>	DDI_SUCCESS	A successful return. The size, in bytes, of the specified register, is set in <i>resultp</i> .	DDI_FAILURE	An invalid (nonexistent) register number was specified.		
DDI_SUCCESS	A successful return. The size, in bytes, of the specified register, is set in <i>resultp</i> .						
DDI_FAILURE	An invalid (nonexistent) register number was specified.						
CONTEXT	ddi_dev_regsize() can be called from user or interrupt context.						
SEE ALSO	ddi_dev_nintrs(9F), ddi_dev_nregs(9F) <i>Writing Device Drivers</i>						

ddi_dev_report_fault(9F)

NAME	ddi_dev_report_fault – Report a hardware failure
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_dev_report_fault (dev_info_t *dip, ddi_fault_impact_t impact, ddi_fault_location_t location, const char *message);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>dip</i> Pointer to the driver's dev_info structure to which the fault report relates. (Normally the caller's own dev_info pointer). <i>impact</i> One of a set of enumerated values indicating the impact of the fault on the device's ability to provide normal service. <i>location</i> One of a set of enumerated values indicating the location of the fault, relative to the hardware controlled by the driver specified by dip. <i>message</i> Text of the message describing the fault being reported.
DESCRIPTION	This function provides a standardized mechanism through which device drivers can report hardware faults. Use of this reporting mechanism enables systems equipped with a fault management system to respond to faults discovered by a driver. On a suitably equipped system, this might include automatic failover to an alternative device and/or scheduling replacement of the faulty hardware. The driver must indicate the impact of the fault being reported on its ability to provide service by passing one of the following values for the impact parameter: DDI_SERVICE_LOST Indicates a total loss of service. The driver is unable to implement the normal functions of its hardware. DDI_SERVICE_DEGRADED The driver is unable to provide normal service, but can provide a partial or degraded level of service. The driver may have to make repeated attempts to perform an operation before it succeeds, or it may be running at less than its configured speed. A driver may use this value to indicate that an alternative device should be used if available, but that it can continue operation if no alternative exists. DDI_SERVICE_UNAFFECTED The service provided by the device is currently unaffected by the reported fault. This value may be used to report recovered errors for predictive failure analysis. DDI_SERVICE_RESTORED The driver has resumed normal service, following a previous report that service was lost or degraded. This message implies that any previously reported fault condition no longer exists. The location parameter should be one of the following values:

ddi_dev_report_fault(9F)

DDI_DATAPATH_FAULT

The fault lies in the datapath between the driver and the device. The device may be unplugged, or a problem may exist in the bus on which the device resides. This value is appropriate if the device is not responding to accesses, (for example, the device may not be present) or if a call to `ddi_check_acc_handle(9F)` returns `DDI_FAILURE`.

DDI_DEVICE_FAULT

The fault lies in the device controlled by the driver. This value is appropriate if the device returns an error from a selftest function, or if the driver is able to determine that device is present and accessible, but is not functioning correctly.

DDI_EXTERNAL_FAULT

The fault is external to the device. For example, an Ethernet driver would use this value when reporting a cable fault.

If a device returns detectably bad data during normal operation (an "impossible" value in a register or DMA status area, for example), the driver should check the associated handle using `ddi_check_acc_handle(9F)` or `ddi_check_dma_handle(9F)` before reporting the fault. If the fault is associated with the handle, the driver should specify `DDI_DATAPATH_FAULT` rather than `DDI_DEVICE_FAULT`. As a consequence of this call, the device's state may be updated to reflect the level of service currently available. See `ddi_get_devstate(9F)`.

Note that if a driver calls `ddi_get_devstate(9F)` and discovers that its device is down, a fault should not be reported- the device is down as the result of a fault that has already been reported. Additionally, a driver should avoid incurring or reporting additional faults when the device is already known to be unusable. The `ddi_dev_report_fault()` call should only be used to report hardware (device) problems and should not be used to report purely software problems such as memory (or other resource) exhaustion.

EXAMPLES

An Ethernet driver receives an error interrupt from its device if various fault conditions occur. The driver must read an error status register to determine the nature of the fault, and report it appropriately:

```
static int
xx_error_intr(xx_soft_state *ssp)
{
    ...
    error_status = ddi_get32(ssp->handle, &ssp->regs->xx_err_status);
    if (ddi_check_acc_handle(ssp->handle) != DDI_SUCCESS) {
        ddi_dev_report_fault(ssp->dip, DDI_SERVICE_LOST,
            DDI_DATAPATH_FAULT, "register access fault");
        return DDI_INTR_UNCLAIMED;
    }
    if (ssp->error_status & XX_CABLE_FAULT) {
        ddi_dev_report_fault(ssp->dip, DDI_SERVICE_LOST,
            DDI_EXTERNAL_FAULT, "cable fault");
        return DDI_INTR_CLAIMED;
    }
    if (ssp->error_status & XX_JABBER) {
```

ddi_dev_report_fault(9F)

```
        ddi_dev_report_fault(ssp->dip, DDI_SERVICE_DEGRADED,  
                             DDI_EXTERNAL_FAULT, "jabbering detected")  
        return DDI_INTR_CLAIMED;  
    }  
    ...  
}
```

CONTEXT The `ddi_dev_report_fault()` function may be called from user, kernel, or interrupt context.

SEE ALSO `ddi_check_acc_handle(9F)`, `ddi_check_dma_handle(9F)`, `ddi_get_devstate(9F)`

NAME	ddi_dma_addr_bind_handle – binds an address to a DMA handle		
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>		
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).		
PARAMETERS	<i>handle</i>	The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).	
	<i>as</i>	A pointer to an address space structure. This parameter should be set to NULL, which implies kernel address space.	
	<i>addr</i>	Virtual address of the memory object.	
	<i>len</i>	Length of the memory object in bytes.	
	<i>flags</i>	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.
	<i>callback</i>	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.
	<i>arg</i>	Argument to be passed to the callback function, <i>callback</i> , if such a function is specified.	

ddi_dma_addr_bind_handle(9F)

	<i>cookiep</i> A pointer to the first <code>ddi_dma_cookie(9S)</code> structure.
	<i>ccountp</i> Upon a successful return, <i>ccountp</i> points to a value representing the number of cookies for this DMA object.
DESCRIPTION	<code>ddi_dma_addr_bind_handle()</code> 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 <code>ddi_dma_attr(9S)</code> (see <code>ddi_dma_alloc_handle(9F)</code>). <code>ddi_dma_addr_bind_handle()</code> 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 by calling <code>ddi_dma_nextcookie(9F)</code> the number of times specified by <i>*ccountp-1</i>. When a DMA transfer completes, the driver frees up system DMA resources by calling <code>ddi_dma_unbind_handle(9F)</code>. The <i>flags</i> argument contains information for mapping routines. <code>DDI_DMA_WRITE</code>, <code>DDI_DMA_READ</code>, <code>DDI_DMA_RDWR</code> These flags describe the intended direction of the DMA transfer. <code>DDI_DMA_STREAMING</code> 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 <i>minxfer</i> and <i>burstsizes</i> fields in the DMA attribute structure, <code>ddi_dma_attr(9S)</code> (see <code>ddi_dma_alloc_handle(9F)</code>) is used to allocate the most effective hardware support for large transfers. <code>DDI_DMA_CONSISTENT</code> This flag should be set if the device accesses memory randomly, or if synchronization steps using <code>ddi_dma_sync(9F)</code> need to be as efficient as possible. I/O parameter blocks used for communication between a device and a driver should be allocated using <code>DDI_DMA_CONSISTENT</code>. <code>DDI_DMA_REDZONE</code> 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. <code>DDI_DMA_PARTIAL</code> 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 <code>DDI_DMA_PARTIAL_MAP</code>. At a later point, the caller can use <code>ddi_dma_getwin(9F)</code> to change the valid portion of the object for which resources are allocated. If resources were allocated for only part of the object, <code>ddi_dma_addr_bind_handle()</code> returns resources for the first DMA window.

ddi_dma_addr_bind_handle(9F)

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:

<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.

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_umem_iosetup(9F)`, `ddi_dma_attr(9S)`, `ddi_dma_cookie(9S)`

Writing Device Drivers

`ddi_dma_addr_bind_handle(9F)`

NOTES	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> .

ddi_dma_addr_setup(9F)

NAME	ddi_dma_addr_setup – easier DMA setup for use with virtual addresses
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL PARAMETERS	This interface is obsolete. ddi_dma_addr_bind_handle(9F) should be used instead. <i>dip</i> A pointer to the device's dev_info structure. <i>as</i> A pointer to an address space structure. Should be set to NULL, which implies kernel address space. <i>addr</i> Virtual address of the memory object. <i>len</i> Length of the memory object in bytes. <i>flags</i> Flags that would go into the ddi_dma_req structure (see ddi_dma_req(9S)). <i>waitfp</i> 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. <i>arg</i> Argument to be passed to a callback function, if such a function is specified. <i>lim</i> 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. <i>handlep</i> Pointer to a DMA handle. See ddi_dma_setup(9F) for a discussion of handle.
DESCRIPTION	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.
RETURN VALUES	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.
ATTRIBUTES	See attributes(5) for a description of the following attributes:

ddi_dma_addr_setup(9F)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), 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

NAME	ddi_dma_alloc_handle – allocate DMA handle				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>				
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>dip</i> Pointer to the device's dev_info structure. <i>attr</i> Pointer to a DMA attribute structure for this device (see ddi_dma_attr(9S)). <i>callback</i> The address of a function to call back later if resources aren't available now. The following special function addresses may also be used. <table> <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> <i>arg</i> Argument to be passed to a callback function, if such a function is specified. <i>handlep</i> Pointer to the DMA handle to be initialized.	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.				
DESCRIPTION	ddi_dma_alloc_handle() allocates a new DMA handle. A DMA handle is an opaque object used as a reference to subsequently allocated DMA resources. ddi_dma_alloc_handle() accepts as parameters the device information referred to by <i>dip</i> and the device's DMA attributes described by a ddi_dma_attr(9S) structure. A successful call to ddi_dma_alloc_handle() 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. 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 DDI_DMA_DONTWAIT, then the caller does not care if the allocation fails, and can handle an allocation failure appropriately. If <i>callback</i> is set to DDI_DMA_SLEEP, then 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 a function to call at a later time when resources may become available. When the specified function is called, it is passed <i>arg</i> 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.				

ddi_dma_alloc_handle(9F)

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

NAME	ddi_dma_buf_bind_handle – binds a system buffer to a DMA handle		
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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);</pre>		
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI).		
	<i>handle</i>	The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).	
	<i>bp</i>	A pointer to a system buffer structure (see buf(9S)).	
	<i>flags</i>	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.
	<i>callback</i>	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.
	<i>arg</i>	Argument to be passed to the callback function, <i>callback</i> , if such a function is specified.	
	<i>cookiep</i>	A pointer to the first ddi_dma_cookie(9S) structure.	
	<i>ccountp</i>	Upon a successful return, <i>ccountp</i> points to a value representing the number of cookies for this DMA object.	

ddi_dma_buf_bind_handle(9F)

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 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 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

ddi_dma_buf_bind_handle(9F)

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:

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_buf_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 be called from user or kernel context only.

SEE ALSO

ddi_dma_addr_bind_handle(9F), ddi_dma_alloc_handle(9F), ddi_dma_free_handle(9F), ddi_dma_getwin(9F), ddi_dma_nextcookie(9F), ddi_dma_sync(9F), ddi_dma_unbind_handle(9F), buf(9S), ddi_dma_attr(9S), ddi_dma_cookie(9S)

Writing Device Drivers

ddi_dma_buf_bind_handle(9F)

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)`.

NAME	ddi_dma_buf_setup – easier DMA setup for use with buffer structures
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL PARAMETERS	This interface is obsolete. ddi_dma_buf_bind_handle(9F) should be used instead. <i>dip</i> A pointer to the device's dev_info structure. <i>bp</i> A pointer to a system buffer structure (see buf(9S)). <i>flags</i> Flags that go into a ddi_dma_req structure (see ddi_dma_req(9S)). <i>waitfp</i> 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. <i>arg</i> Argument to be passed to a callback function, if such a function is specified. <i>lim</i> 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. <i>handlep</i> Pointer to a DMA handle. See ddi_dma_setup(9F) for a discussion of handle.
DESCRIPTION	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.
RETURN VALUES	See ddi_dma_setup(9F) for the possible return values for this function.
CONTEXT	ddi_dma_buf_setup() can be called from user or interrupt context, except when <i>waitfp</i> is set to DDI_DMA_SLEEP, in which case it can be called from user context only.
ATTRIBUTES	See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

ddi_dma_buf_setup(9F)

SEE ALSO | attributes(5), 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

	ddi_dma_burstsizes(9F)
NAME	ddi_dma_burstsizes – find out the allowed burst sizes for a DMA mapping
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_burstsizes(ddi_dma_handle_t <i>handle</i>);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>handle</i> A DMA handle that was filled in by a successful call to ddi_dma_setup(9F).
DESCRIPTION	ddi_dma_burstsizes() returns the allowed burst sizes for a DMA mapping. This value is derived from the dlim_burstsizes member of the ddi_dma_lim_sparc(9S) 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.
RETURN VALUES	ddi_dma_burstsizes() returns a binary encoded value of the allowable DMA burst sizes. See ddi_dma_lim_sparc(9S) for a discussion of DMA burst sizes.
CONTEXT	This function can be called from user or interrupt context.
SEE ALSO	ddi_dma_dealign(9F), ddi_dma_setup(9F), ddi_dma_lim_sparc(9S), ddi_dma_req(9S) <i>Writing Device Drivers</i>

ddi_dma_coff(9F)

NAME	ddi_dma_coff – convert a DMA cookie to an offset within a DMA handle						
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_coff(ddi_dma_handle_t handle, ddi_dma_cookie_t *cookiep, off_t *offp);</pre>						
INTERFACE LEVEL	Solaris SPARC DDI (Solaris SPARC DDI).						
PARAMETERS	<table><tr><td><i>handle</i></td><td>The <i>handle</i> filled in by a call to ddi_dma_setup(9F).</td></tr><tr><td><i>cookiep</i></td><td>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.</td></tr><tr><td><i>offp</i></td><td>A pointer to an offset to be filled in.</td></tr></table>	<i>handle</i>	The <i>handle</i> filled in by a call to ddi_dma_setup(9F).	<i>cookiep</i>	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.	<i>offp</i>	A pointer to an offset to be filled in.
<i>handle</i>	The <i>handle</i> filled in by a call to ddi_dma_setup(9F).						
<i>cookiep</i>	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.						
<i>offp</i>	A pointer to an offset to be filled in.						
DESCRIPTION	ddi_dma_coff() 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. ddi_dma_coff() 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.						
RETURN VALUES	ddi_dma_coff() returns: <table><tr><td>DDI_SUCCESS</td><td>Successfully filled in <i>offp</i>.</td></tr><tr><td>DDI_FAILURE</td><td>Failed to successfully fill in <i>offp</i>.</td></tr></table>	DDI_SUCCESS	Successfully filled in <i>offp</i> .	DDI_FAILURE	Failed to successfully fill in <i>offp</i> .		
DDI_SUCCESS	Successfully filled in <i>offp</i> .						
DDI_FAILURE	Failed to successfully fill in <i>offp</i> .						
CONTEXT	ddi_dma_coff() can be called from user or interrupt context.						
SEE ALSO	ddi_dma_setup(9F), ddi_dma_sync(9F), ddi_dma_cookie(9S) <i>Writing Device Drivers</i>						

NAME	ddi_dma_curwin – report current DMA window offset and size					
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_curwin(ddi_dma_handle_t <i>handle</i>, off_t *<i>offp</i>, uint_t *<i>lenp</i>) ;</pre>					
INTERFACE LEVEL	This interface is obsolete. ddi_dma_getwin(9F) should be used instead.					
PARAMETERS	<i>handle</i>	The DMA handle filled in by a call to ddi_dma_setup(9F).				
	<i>offp</i>	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.				
	<i>lenp</i>	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.				
DESCRIPTION	ddi_dma_curwin() 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 ddi_dma_curwin().					
RETURN VALUES	ddi_dma_curwin() returns:					
	DDI_SUCCESS	The current length and offset can be established.				
	DDI_FAILURE	Otherwise.				
CONTEXT	ddi_dma_curwin() can be called from user or interrupt context.					
ATTRIBUTES	See attributes(5) for a description of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Stability Level</td><td>Obsolete</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Stability Level	Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Stability Level	Obsolete					
SEE ALSO	attributes(5), ddi_dma_getwin(9F), ddi_dma_movwin(9F), ddi_dma_setup(9F), ddi_dma_req(9S)					
	Writing Device Drivers					

ddi_dma_devalign(9F)

NAME	ddi_dma_devalign – find DMA mapping alignment and minimum transfer size						
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_devalign(ddi_dma_handle_t handle, uint_t *alignment, uint_t *minxfr);</pre>						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).						
PARAMETERS	<table><tr><td><i>handle</i></td><td>The DMA handle filled in by a successful call to ddi_dma_setup(9F).</td></tr><tr><td><i>alignment</i></td><td>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.</td></tr><tr><td><i>minxfr</i></td><td>A pointer to an unsigned integer to be filled in with the minimum effective transfer size (see ddi_iomin(9F), ddi_dma_lim_sparc(9S) and ddi_dma_lim_x86(9S)). This also is guaranteed to be a power of two.</td></tr></table>	<i>handle</i>	The DMA handle filled in by a successful call to ddi_dma_setup(9F).	<i>alignment</i>	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.	<i>minxfr</i>	A pointer to an unsigned integer to be filled in with the minimum effective transfer size (see ddi_iomin(9F), ddi_dma_lim_sparc(9S) and ddi_dma_lim_x86(9S)). This also is guaranteed to be a power of two.
<i>handle</i>	The DMA handle filled in by a successful call to ddi_dma_setup(9F).						
<i>alignment</i>	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.						
<i>minxfr</i>	A pointer to an unsigned integer to be filled in with the minimum effective transfer size (see ddi_iomin(9F), ddi_dma_lim_sparc(9S) and ddi_dma_lim_x86(9S)). This also is guaranteed to be a power of two.						
DESCRIPTION	ddi_dma_devalign() determines after a successful DMA mapping (see ddi_dma_setup(9F)) the minimum required data alignment and minimum DMA transfer size.						
RETURN VALUES	ddi_dma_devalign() returns: <table><tr><td>DDI_SUCCESS</td><td>The <i>alignment</i> and <i>minxfr</i> values have been filled.</td></tr><tr><td>DDI_FAILURE</td><td>The handle was illegal.</td></tr></table>	DDI_SUCCESS	The <i>alignment</i> and <i>minxfr</i> values have been filled.	DDI_FAILURE	The handle was illegal.		
DDI_SUCCESS	The <i>alignment</i> and <i>minxfr</i> values have been filled.						
DDI_FAILURE	The handle was illegal.						
CONTEXT	ddi_dma_devalign() can be called from user or interrupt context.						
SEE ALSO	ddi_dma_setup(9F), ddi_iomin(9F), ddi_dma_lim_sparc(9S), ddi_dma_lim_x86(9S), ddi_dma_req(9S) <i>Writing Device Drivers</i>						

NAME	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																				
SYNOPSIS	<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 *dmaareqp, 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>																				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). The ddi_dmae_getlim() interface, described below, is obsolete. Use ddi_dmae_getattr(), also described below, to replace it.																				
PARAMETERS	<table> <tr> <td><i>dip</i></td><td>A dev_info pointer that identifies the device.</td></tr> <tr> <td><i>chnl</i></td><td>A DMA channel number. On ISA or EISA buses this number must be 0, 1, 2, 3, 5, 6, or 7.</td></tr> <tr> <td><i>callback</i></td><td> The address of a function to call back later if resources are not currently available. The following special function addresses may also be used: <table> <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> </td></tr> <tr> <td><i>arg</i></td><td>Argument to be passed to the callback function, if specified.</td></tr> <tr> <td><i>dmaareqp</i></td><td>A pointer to a DMA engine request structure. See ddi_dmae_req(9S).</td></tr> <tr> <td><i>cookiep</i></td><td>A pointer to a ddi_dma_cookie(9S) object, obtained from ddi_dma_segtocookie(9F), which contains the address and count.</td></tr> <tr> <td><i>countp</i></td><td>A pointer to an integer that will receive the count of the number of bytes not yet transferred upon completion of a DMA operation.</td></tr> <tr> <td><i>limitsp</i></td><td>A pointer to a DMA limit structure. See ddi_dma_lim_x86(9S).</td></tr> </table>	<i>dip</i>	A dev_info pointer that identifies the device.	<i>chnl</i>	A DMA channel number. On ISA or EISA buses this number must be 0, 1, 2, 3, 5, 6, or 7.	<i>callback</i>	The address of a function to call back later if resources are not currently available. The following special function addresses may also be used: <table> <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>	DDI_DMA_SLEEP	Wait until resources are available.	DDI_DMA_DONTWAIT	Do not wait until resources are available and do not schedule a callback.	<i>arg</i>	Argument to be passed to the callback function, if specified.	<i>dmaareqp</i>	A pointer to a DMA engine request structure. See ddi_dmae_req(9S).	<i>cookiep</i>	A pointer to a ddi_dma_cookie(9S) object, obtained from ddi_dma_segtocookie(9F), which contains the address and count.	<i>countp</i>	A pointer to an integer that will receive the count of the number of bytes not yet transferred upon completion of a DMA operation.	<i>limitsp</i>	A pointer to a DMA limit structure. See ddi_dma_lim_x86(9S).
<i>dip</i>	A dev_info pointer that identifies the device.																				
<i>chnl</i>	A DMA channel number. On ISA or EISA buses this number must be 0, 1, 2, 3, 5, 6, or 7.																				
<i>callback</i>	The address of a function to call back later if resources are not currently available. The following special function addresses may also be used: <table> <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>	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.																				
<i>arg</i>	Argument to be passed to the callback function, if specified.																				
<i>dmaareqp</i>	A pointer to a DMA engine request structure. See ddi_dmae_req(9S).																				
<i>cookiep</i>	A pointer to a ddi_dma_cookie(9S) object, obtained from ddi_dma_segtocookie(9F), which contains the address and count.																				
<i>countp</i>	A pointer to an integer that will receive the count of the number of bytes not yet transferred upon completion of a DMA operation.																				
<i>limitsp</i>	A pointer to a DMA limit structure. See ddi_dma_lim_x86(9S).																				

ddi_dmae(9F)

	<i>attrp</i>	A pointer to a DMA attribute structure. See ddi_dma_attr(9S).
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 <i>callback</i> determines what action will be taken. If the value of <i>callback</i> 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 <i>arg</i>) 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 <code>ddi_dmae_prog()</code> function programs the DMA channel for a DMA transfer. The <code>ddi_dmae_req</code> 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 <code>ddi_dmae_prog()</code> may specify a value of <code>NULL</code> for <i>dmaereqp</i> 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 <code>ddi_dmae_prog()</code> a cookie obtained from <code>ddi_dma_segtocookie(9F)</code>. Other DMA engine parameters are specified by the DMA engine request structure passed in through <i>dmaereqp</i>. The fields of that structure are documented in <code>ddi_dmae_req(9S)</code>. Before using <code>ddi_dmae_prog()</code>, you must allocate system DMA resources using DMA setup functions such as <code>ddi_dma_buf_setup(9F)</code>. <code>ddi_dma_segtocookie(9F)</code> can then be used to retrieve a cookie which contains the address and count. Then this cookie is passed to <code>ddi_dmae_prog()</code>.
ddi_dmae_disable()	The <code>ddi_dmae_disable()</code> function disables the DMA channel so that it no longer responds to a device's DMA service requests.
ddi_dmae_enable()	The <code>ddi_dmae_enable()</code> function enables the DMA channel for operation. This may be used to re-enable the channel after a call to <code>ddi_dmae_disable()</code> . The channel is automatically enabled after successful programming by <code>ddi_dmae_prog()</code> .
ddi_dmae_stop()	The <code>ddi_dmae_stop()</code> function disables the channel and terminates any active operation.
ddi_dmae_getcnt()	The <code>ddi_dmae_getcnt()</code> 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.
ddi_dmae_1stparty()	In the case of ISA and EISA buses, <code>ddi_dmae_1stparty()</code> configures a channel in the system's DMA engine to operate in a "slave" ("cascade") mode. When operating in <code>ddi_dmae_1stparty()</code> mode, the DMA channel must first be allocated using <code>ddi_dmae_alloc()</code> and then configured using <code>ddi_dmae_1stparty()</code>. The driver then programs the device to perform the I/O, including the necessary DMA address and count values obtained from <code>ddi_dma_segtocookie(9F)</code>.
ddi_dmae_getlim()	Note that this function is obsolete. Use <code>ddi_dmae_getattr()</code>, described below, instead. The <code>ddi_dmae_getlim()</code> 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 <code>ddi_dmae_getlim()</code>. 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 <code>ddi_dma_nextseg(9F)</code> and <code>ddi_dma_nextwin(9F)</code>). If the device has any particular restrictions on transfer size or granularity (such as the size of disk sector), the driver

ddi_dmae(9F)

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 `ddi_dmae_getlim()`. After calling `ddi_dmae_getlim()`, 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 `ddi_dma_lim_x86(9S)` and `ddi_dmae_req(9S)` for additional information on scatter/gather DMA.

ddi_dmae_getattr The `ddi_dmae_getattr()` function fills in the DMA attribute structure, pointed to by *attrp*, 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 `ddi_dmae_getattr()`. 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 `ddi_dma_nextcookie(9F)` and `ddi_dma_getwin(9F)`.

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 <code>ddi_dmae_alloc()</code> if the requested resources are not available and the value of <i>dmae_waitfp</i> is not <code>DDI_DMA_SLEEP</code> .

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)`

NAME	ddi_dma_free – release system DMA resources				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_free(ddi_dma_handle_t <i>handle</i>);</pre>				
INTERFACE LEVEL	This interface is obsolete. ddi_dma_free_handle(9F) should be used instead.				
PARAMETERS	<i>handle</i> The handle filled in by a call to ddi_dma_setup(9F).				
DESCRIPTION	ddi_dma_free() releases system DMA resources set up by ddi_dma_setup(9F). When a DMA transfer completes, the driver should free up system DMA resources established by a call to ddi_dma_setup(9F). This is done by a call to ddi_dma_free(). ddi_dma_free() does an implicit ddi_dma_sync(9F) for you so any further synchronization steps are not necessary.				
RETURN VALUES	ddi_dma_free() returns: DDI_SUCCESS Successfully released resources DDI_FAILURE Failed to free resources				
CONTEXT	ddi_dma_free() can be called from user or interrupt context.				
ATTRIBUTES	See attributes(5) for a description of the following attributes: <table border="1" data-bbox="459 1047 1430 1138"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Stability Level</td><td>Obsolete</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Stability Level	Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Stability Level	Obsolete				
SEE ALSO	attributes(5), ddi_dma_addr_setup(9F), ddi_dma_buf_setup(9F), ddi_dma_free_handle(9F), ddi_dma_htoc(9F), ddi_dma_sync(9F), ddi_dma_req(9S) <i>Writing Device Drivers</i>				

ddi_dma_free_handle(9F)

NAME	ddi_dma_free_handle – free DMA handle		
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_dma_free_handle(ddi_dma_handle_t *handle) ;</pre>		
PARAMETERS	<table><tr><td><i>handle</i></td><td>A pointer to the DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).</td></tr></table>	<i>handle</i>	A pointer to the DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).
<i>handle</i>	A pointer to the DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).		
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).		
DESCRIPTION	ddi_dma_free_handle() destroys the DMA handle pointed to by <i>handle</i> . Any further references to the DMA handle will have undefined results. Note that ddi_dma_unbind_handle(9F) must be called prior to ddi_dma_free_handle() to free any resources the system may be caching on the handle.		
CONTEXT	ddi_dma_free_handle() can be called from user, kernel, or interrupt context.		
SEE ALSO	ddi_dma_alloc_handle(9F), ddi_dma_unbind_handle(9F) <i>Writing Device Drivers</i>		

NAME	ddi_dma_get_attr – get the device DMA attribute structure from a DMA handle
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_get_attr(ddi_dma_handle_t <i>handle</i>, ddi_dma_attr_t *attrp) ;</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>handle</i> The handle filled in by a call to ddi_dma_alloc_handle(9F). <i>attrp</i> Pointer to a buffer suitable for holding a DMA attribute structure. See ddi_dma_attr(9S).
DESCRIPTION	ddi_dma_get_attr() is used to get a ddi_dma_attr(9S) structure. This structure describes the attributes of the DMA data path to which any memory object bound to the given handle will be subject.
RETURN VALUES	DDI_SUCCESS Successfully passed back attribute structure in buffer pointed to by attrp. DDI_DMA_BADATTR A valid attribute structure could not be passed back.
CONTEXT	ddi_dma_get_attr() can be called from any context.
SEE ALSO	ddi_dma_alloc_handle(9F), ddi_dma_attr(9S)

ddi_dma_getwin(9F)

NAME	ddi_dma_getwin – activate a new DMA window												
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).												
PARAMETERS	<table><tr><td><i>handle</i></td><td>The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).</td></tr><tr><td><i>win</i></td><td>Number of the window to activate.</td></tr><tr><td><i>offp</i></td><td>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.</td></tr><tr><td><i>lenp</i></td><td>Upon a successful return, <i>lenp</i> will contain the size, in bytes, of the current window.</td></tr><tr><td><i>cookiep</i></td><td>A pointer to the first ddi_dma_cookie(9S) structure.</td></tr><tr><td><i>ccountp</i></td><td>Upon a successful return, <i>ccountp</i> will contain the number of cookies for this DMA window.</td></tr></table>	<i>handle</i>	The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).	<i>win</i>	Number of the window to activate.	<i>offp</i>	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.	<i>lenp</i>	Upon a successful return, <i>lenp</i> will contain the size, in bytes, of the current window.	<i>cookiep</i>	A pointer to the first ddi_dma_cookie(9S) structure.	<i>ccountp</i>	Upon a successful return, <i>ccountp</i> will contain the number of cookies for this DMA window.
<i>handle</i>	The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).												
<i>win</i>	Number of the window to activate.												
<i>offp</i>	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.												
<i>lenp</i>	Upon a successful return, <i>lenp</i> will contain the size, in bytes, of the current window.												
<i>cookiep</i>	A pointer to the first ddi_dma_cookie(9S) structure.												
<i>ccountp</i>	Upon a successful return, <i>ccountp</i> will contain the number of cookies for this DMA window.												
DESCRIPTION	ddi_dma_getwin() activates a new DMA window. If a DMA resource allocation request returns DDI_DMA_PARTIAL_MAP indicating that resources for less than the entire object were allocated, the current DMA window can be changed by a call to ddi_dma_getwin(). The caller must first determine the number of DMA windows, <i>N</i>, using ddi_dma_numwin(9F). ddi_dma_getwin() 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. ddi_dma_getwin() 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 ddi_dma_nextcookie(9F). ddi_dma_getwin() 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 ddi_dma_sync(9F). 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_dma_getwin(9F)

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

ddi_dma_htoc(9F)

NAME	ddi_dma_htoc – convert a DMA handle to a DMA address cookie						
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_htoc(ddi_dma_handle_t handle, off_t off, ddi_dma_cookie_t *cookiep);</pre>						
INTERFACE LEVEL	This interface is obsolete. ddi_dma_addr_bind_handle(9F) or ddi_dma_buf_bind_handle(9F) should be used instead.						
PARAMETERS	<table><tr><td><i>handle</i></td><td>The handle filled in by a call to ddi_dma_setup(9F).</td></tr><tr><td><i>off</i></td><td>An offset into the object that <i>handle</i> maps.</td></tr><tr><td><i>cookiep</i></td><td>A pointer to a ddi_dma_cookie(9S) structure.</td></tr></table>	<i>handle</i>	The handle filled in by a call to ddi_dma_setup(9F).	<i>off</i>	An offset into the object that <i>handle</i> maps.	<i>cookiep</i>	A pointer to a ddi_dma_cookie(9S) structure.
<i>handle</i>	The handle filled in by a call to ddi_dma_setup(9F).						
<i>off</i>	An offset into the object that <i>handle</i> maps.						
<i>cookiep</i>	A pointer to a ddi_dma_cookie(9S) structure.						
DESCRIPTION	ddi_dma_htoc() takes a DMA handle (established by ddi_dma_setup(9F)), 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.						
RETURN VALUES	ddi_dma_htoc() returns: <table><tr><td>DDI_SUCCESS</td><td>Successfully filled in the cookie pointed to by <i>cookiep</i>.</td></tr><tr><td>DDI_FAILURE</td><td>Failed to successfully fill in the cookie.</td></tr></table>	DDI_SUCCESS	Successfully filled in the cookie pointed to by <i>cookiep</i> .	DDI_FAILURE	Failed to successfully fill in the cookie.		
DDI_SUCCESS	Successfully filled in the cookie pointed to by <i>cookiep</i> .						
DDI_FAILURE	Failed to successfully fill in the cookie.						
CONTEXT	ddi_dma_htoc() can be called from user or interrupt context.						
ATTRIBUTES	See attributes(5) for a description of the following attributes: <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Stability Level</td><td>Obsolete</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Stability Level	Obsolete		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Stability Level	Obsolete						
SEE ALSO	attributes(5), ddi_dma_addr_bind_handle(9F), ddi_dma_addr_setup(9F), ddi_dma_buf_bind_handle(9F), ddi_dma_buf_setup(9F), ddi_dma_setup(9F), ddi_dma_sync(9F), ddi_dma_cookie(9S) <i>Writing Device Drivers</i>						

NAME	ddi_dma_mem_alloc – allocate memory for DMA transfer		
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>		
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI).		
	<i>handle</i>	The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).	
	<i>length</i>	The length in bytes of the desired allocation.	
	<i>accattrp</i>	Pointer to a device access attribute structure of this device (see ddi_device_acc_attr(9S)).	
	<i>flags</i>	Data transfer mode flags. Possible values are:	
		DDI_DMA_STREAMING	Sequential, unidirectional, block-sized, and block-aligned transfers.
		DDI_DMA_CONSISTENT	Nonsequential transfers of small objects.
	<i>waitfp</i>	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 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, <i>arg</i> 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.	

ddi_dma_mem_alloc(9F)

	The callback function must take whatever steps are necessary to protect its critical resources, data structures, queues, and so on.
<i>arg</i>	Argument to be passed to the callback function, if such a function is specified.
<i>kaddrp</i>	On successful return, <i>kaddrp</i> points to the allocated memory.
<i>real_length</i>	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> .
<i>handlep</i>	Pointer to a data access handle.
DESCRIPTION	<code>ddi_dma_mem_alloc()</code> 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 <code>ddi_dma_attr(9S)</code>) passed to <code>ddi_dma_alloc_handle(9F)</code> and the more restrictive attributes imposed by the system. <i>flags</i> should be set to <code>DDI_DMA_STREAMING</code> 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 <code>minxfer</code> and <code>burstsizes</code> fields in the DMA attribute structure, <code>ddi_dma_attr(9S)</code> (see <code>ddi_dma_alloc_handle(9F)</code>) 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, <code>ddi_dma_mem_alloc()</code> will align the memory at a cache line boundary and it will round up <i>real_length</i> to a multiple of the cache line size. <i>flags</i> should be set to <code>DDI_DMA_CONSISTENT</code> if the device accesses memory randomly, or if synchronization steps using <code>ddi_dma_sync(9F)</code> need to be as efficient as possible. I/O parameter blocks used for communication between a device and a driver should be allocated using <code>DDI_DMA_CONSISTENT</code>. The device access attributes are specified in the location pointed by the <i>accattrp</i> argument (see <code>ddi_device_acc_attr(9S)</code>). The data access handle is returned in <i>handlep</i>. <i>handlep</i> is opaque – drivers may not attempt to interpret its value. To access the data content, the driver must invoke <code>ddi_get8(9F)</code> or <code>ddi_put8(9F)</code> (depending on the data transfer direction) with the data access handle. DMA resources must be established before performing a DMA transfer by passing <i>kaddrp</i> and <i>real_length</i> as returned from <code>ddi_dma_mem_alloc()</code> and the flag <code>DDI_DMA_STREAMING</code> or <code>DDI_DMA_CONSISTENT</code> to <code>ddi_dma_addr_bind_handle(9F)</code>. 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 <code>ddi_dma_sync(9F)</code> or <code>ddi_dma_unbind_handle(9F)</code> are required.

ddi_dma_mem_alloc(9F)

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 <i>waitfp</i> 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) <i>Writing Device Drivers</i>
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.

ddi_dma_mem_free(9F)

NAME	ddi_dma_mem_free – free previously allocated memory
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_dma_mem_free(ddi_acc_handle_t *handlep);</pre>
PARAMETERS	<i>handlep</i> Pointer to the data access handle previously allocated by a call to ddi_dma_mem_alloc(9F).
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_dma_mem_free() deallocates the memory acquired by ddi_dma_mem_alloc(9F). In addition, it destroys the data access handle <i>handlep</i> associated with the memory.
CONTEXT	ddi_dma_mem_free() can be called from user, kernel, or interrupt context.
SEE ALSO	ddi_dma_mem_alloc(9F) <i>Writing Device Drivers</i>

NAME	ddi_dma_movwin – shift current DMA window
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_movwin(ddi_dma_handle_t handle, off_t *offp, uint_t *lenp, ddi_dma_cookie_t *cookiep);</pre>
INTERFACE LEVEL PARAMETERS	This interface is obsolete. ddi_dma_getwin(9F) should be used instead. <i>handle</i> The DMA handle filled in by a call to ddi_dma_setup(9F). <i>offp</i> 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. <i>lenp</i> A pointer to a value which must either be the current size of the DMA window (as known from a call to ddi_dma_curwin(9F) or from a previous call to ddi_dma_movwin()). Upon a successful return, it will be filled in with the size, in bytes, of the current window. <i>cookiep</i> A pointer to a DMA cookie (see ddi_dma_cookie(9S)). Upon a successful return, cookiep is filled in just as if an implicit ddi_dma_htoc(9F) had been made.
DESCRIPTION	ddi_dma_movwin() shifts the current DMA window. If a DMA request allows the system to allocate resources for less than the entire object by setting the DDI_DMA_PARTIAL flag in the ddi_dma_req(9S) structure, the current DMA window can be shifted by a call to ddi_dma_movwin(). The caller must first determine the current DMA window size by a call to ddi_dma_curwin(9F). Using the current offset and size of the window thus retrieved, the caller of ddi_dma_movwin() 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. ddi_dma_movwin() takes care of underlying resource synchronizations required to shift the window. However, if you want to <i>access</i> the data prior to or after moving the window, further synchronizations using ddi_dma_sync(9F) are required. 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 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.

ddi_dma_movwin(9F)

CONTEXT ddi_dma_movwin() can be called from user or interrupt context.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), ddi_dma_curwin(9F), ddi_dma_getwin(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.

NAME	ddi_dma_nextcookie – retrieve subsequent DMA cookie
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_dma_nextcookie(ddi_dma_handle_t <i>handle</i>, ddi_dma_cookie_t *<i>cookiep</i>) ;</pre>
PARAMETERS	<i>handle</i> The handle previously allocated by a call to ddi_dma_alloc_handle(9F). <i>cookiep</i> A pointer to a ddi_dma_cookie(9S) structure.
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_dma_nextcookie() retrieves subsequent DMA cookies for a DMA object. ddi_dma_nextcookie() fills in the ddi_dma_cookie(9S) structure pointed to by <i>cookiep</i>. The ddi_dma_cookie(9S) structure must be allocated prior to calling ddi_dma_nextcookie(). The DMA cookie count returned by ddi_dma_buf_bind_handle(9F), ddi_dma_addr_bind_handle(9F), or ddi_dma_getwin(9F) indicates the number of DMA cookies a DMA object consists of. If the resulting cookie count, <i>N</i>, is larger than 1, ddi_dma_nextcookie() must be called <i>N</i>-1 times to retrieve all DMA cookies.
CONTEXT	ddi_dma_nextcookie() can be called from user, kernel, or interrupt context.
EXAMPLES	EXAMPLE 1 process a scatter-gather list of I/O requests This example demonstrates the use of ddi_dma_nextcookie() to process a scatter-gather list of I/O requests. <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, &dmacookie, &ccount); if (status == DDI_DMA_MAPPED) { /* program DMA engine with first cookie */ while (--ccount > 0) { ddi_dma_nextcookie(handle, &dmacookie); /* program DMA engine with next cookie */ } } . . .</pre>

ddi_dma_nextcookie(9F)

EXAMPLE 1 process a scatter-gather list of I/O requests *(Continued)*

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

NAME	ddi_dma_nextseg – get next DMA segment
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_nextseg(ddi_dma_win_t win, ddi_dma_seg_t seg, ddi_dma_seg_t *nseg);</pre>
INTERFACE LEVEL PARAMETERS	This interface is obsolete. ddi_dma_nextcookie(9F) should be used instead. <i>win</i> A DMA window. <i>seg</i> The current DMA segment or NULL. <i>nseg</i> A pointer to the next DMA segment to be filled in. If <i>seg</i> is NULL, a pointer to the first segment within the specified window is returned.
DESCRIPTION	ddi_dma_nextseg() 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. A DMA segment is always required for a DMA window. A DMA segment is a contiguous portion of a DMA window (see ddi_dma_nextwin(9F)) which is entirely addressable by the device for a data transfer operation. 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 ddi_dma_lim_sparc(9S) or ddi_dma_lim_x86(9S)). In this example the object will be broken into smaller addressable DMA segments.
RETURN VALUES	ddi_dma_nextseg() returns: DDI_SUCCESS Successfully filled in the next segment pointer. DDI_DMA_DONE There is no next segment. The current segment is the final segment within the specified window. DDI_DMA_STALE <i>win</i> does not refer to the currently active window.
CONTEXT	ddi_dma_nextseg() can be called from user or interrupt context.
EXAMPLES	For an example, see ddi_dma_segtocookie(9F).
ATTRIBUTES	See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

ddi_dma_nextseg(9F)

SEE ALSO | attributes(5), ddi_dma_addr_setup(9F), ddi_dma_buf_setup(9F),
ddi_dma_nextcookie(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

NAME	ddi_dma_nextwin – get next DMA window
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_nextwin(ddi_dma_handle_t handle, ddi_dma_win_t win, ddi_dma_win_t *nwin);</pre>
INTERFACE LEVEL	This interface is obsolete. ddi_dma_getwin(9F) should be used instead.
PARAMETERS	<i>handle</i> A DMA handle. <i>win</i> The current DMA window or NULL. <i>nwin</i> 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.
DESCRIPTION	ddi_dma_nextwin() 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. 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 ddi_dma_buf_setup(9F) or ddi_dma_addr_setup(9F) or as part of a ddi_dma_req(9S) structure in a call to ddi_dma_setup(9F). 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 ddi_dma_nextwin(). Furthermore, because a call to ddi_dma_nextwin() will reallocate system resources to the new window, the previous window will become invalid. It is a <i>severe</i> error to call ddi_dma_nextwin() before any transfers into the current window are complete. ddi_dma_nextwin() takes care of underlying memory synchronizations required to shift the window. However, if you want to access the data before 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 <i>win</i> does not refer to the currently active window.
CONTEXT	ddi_dma_nextwin() can be called from user or interrupt context.

ddi_dma_nextwin(9F)

EXAMPLES For an example see ddi_dma_segtocookie(9F).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), ddi_dma_addr_setup(9F), ddi_dma_buf_setup(9F), ddi_dma_getwin(9F), ddi_dma_nextseg(9F), ddi_dma_segtocookie(9F), ddi_dma_sync(9F), ddi_dma_req(9S)

Writing Device Drivers

NAME	ddi_dma_numwin – retrieve number of DMA windows
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_numwin(ddi_dma_handle_t <i>handle</i>, uint_t *<i>nwinp</i>);</pre>
PARAMETERS	<i>handle</i> The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F). <i>nwinp</i> Upon a successful return, <i>nwinp</i> will contain the number of DMA windows for this object.
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_dma_numwin() returns the number of DMA windows for a DMA object if partial resource allocation was permitted.
RETURN VALUES	ddi_dma_numwin() returns: DDI_SUCCESS Successfully filled in the number of DMA windows. DDI_FAILURE DMA windows are not activated.
CONTEXT	ddi_dma_numwin() 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_unbind_handle(9F) <i>Writing Device Drivers</i>

ddi_dma_segtocookie(9F)

NAME	ddi_dma_segtocookie – convert a DMA segment to a DMA address cookie
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_segtocookie(ddi_dma_seg_t seg, off_t *offp, off_t *lenp, ddi_dma_cookie_t *cookiep);</pre>
INTERFACE LEVEL PARAMETERS	This interface is obsolete. ddi_dma_nextcookie(9F) should be used instead. <i>seg</i> A DMA segment. <i>offp</i> A pointer to an off_t. Upon a successful return, it is filled in with the offset. This segment is addressing within the object. <i>lenp</i> The byte length. This segment is addressing within the object. <i>cookiep</i> A pointer to a DMA cookie (see ddi_dma_cookie(9S)).
DESCRIPTION	ddi_dma_segtocookie() 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. ddi_dma_segtocookie() also fills in <i>*offp</i> and <i>*lenp</i> , which specify the range within the object.
RETURN VALUES	ddi_dma_segtocookie() returns: DDI_SUCCESS Successfully filled in all values. DDI_FAILURE Failed to successfully fill in all values.
CONTEXT	ddi_dma_segtocookie() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 ddi_dma_segtocookie() example <pre>for (win = NULL; (retw = ddi_dma_nextwin(handle, win, &nwin)) != DDI_DMA_DONE; win = nwin) { if (retw != DDI_SUCCESS) { /* do error handling */ } else { for (seg = NULL; (rets = ddi_dma_nextseg(nwin, seg, &nseg)) != DDI_DMA_DONE; seg = nseg) { if (rets != DDI_SUCCESS) { /* do error handling */ } else { ddi_dma_segtocookie(nseg, &off, &len, &cookie); /* program DMA engine */ } } } }</pre>

ddi_dma_segtocookie(9F)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), ddi_dma_nextcookie(9F), ddi_dma_nextseg(9F),
ddi_dma_nextwin(9F), ddi_dma_sync(9F), ddi_dma_cookie(9S)

Writing Device Drivers

ddi_dma_set_sbus64(9F)

NAME	ddi_dma_set_sbus64 – allow 64-bit transfers on SBus					
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_set_sbus64(ddi_dma_handle_t handle, uint_t burstsizes);</pre>					
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).					
PARAMETERS	<i>handle</i>	The handle filled in by a call to ddi_dma_alloc_handle(9F).				
	<i>burstsizes</i>	The possible burst sizes the device’s DMA engine can accept in 64-bit mode.				
DESCRIPTION	ddi_dma_set_sbus64 () 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 ddi_dma_alloc_handle(9F) with a ddi_dma_attr(9S) structure describing the DMA attributes for a 32-bit transfer mode. <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 ddi_dma_attr(9S) structure. The system will activate 64-bit SBus transfers if the SBus supports them. Otherwise, the SBus will operate in 32-bit mode. After DMA resources have been allocated (see ddi_dma_addr_bind_handle(9F) or ddi_dma_buf_bind_handle(9F)), the driver should retrieve the available burst sizes by calling ddi_dma_burstsizes(9F). 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.					
RETURN VALUES	ddi_dma_set_sbus64 () returns: DDI_SUCCESS Successfully set the SBus to 64-bit mode. DDI_FAILURE 64-bit mode could not be set.					
CONTEXT	ddi_dma_set_sbus64 () can be called from user, kernel, or interrupt context.					
ATTRIBUTES	See attributes(5) for descriptions of the following attributes: <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Architecture</td><td>SBus</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SBus
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.					

NAME	ddi_dma_setup – setup DMA resources
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_setup(dev_info_t *dip, ddi_dma_req_t *dmareqp, ddi_dma_handle_t *handlep);</pre>
INTERFACE LEVEL	This interface is obsolete. The functions ddi_dma_addr_bind_handle(9F), ddi_dma_alloc_handle(9F), ddi_dma_buf_bind_handle(9F), ddi_dma_free_handle(9F), and ddi_dma_unbind_handle(9F) should be used instead.
PARAMETERS	<i>dip</i> A pointer to the device's dev_info structure. <i>dmareqp</i> A pointer to a DMA request structure (see ddi_dma_req(9S)). <i>handlep</i> 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 ddi_dma_setup() 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.
DESCRIPTION	ddi_dma_setup() allocates resources for a memory object such that a device can perform DMA to or from that object. A call to ddi_dma_setup() 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 ddi_dma_req structure pointed to by <i>dmareqp</i>. A successful call to ddi_dma_setup() 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 ddi_dma_free(9F) is called. 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 ddi_dma_htoc(9F), which takes a pointer to a DMA handle and returns the appropriate DMA address. When DMA transfer completes, the driver should free up the the allocated DMA resources by calling ddi_dma_free().
RETURN VALUES	ddi_dma_setup() 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

ddi_dma_setup(9F)

using a flag setting in the `ddi_dma_req` structure (see `ddi_dma_req(9S)` and `ddi_dma_movwin(9F)`).

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 `dmar_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.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `attributes(5)`, `ddi_dma_addr_bind_handle(9F)`, `ddi_dma_alloc_handle(9F)`, `ddi_dma_buf_bind_handle(9F)`, `ddi_dma_free_handle(9F)`, `ddi_dma_unbind_handle(9F)`, `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.

NAME	ddi_dma_sync – synchronize CPU and I/O views of memory
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_sync(ddi_dma_handle_t handle, off_t offset, size_t length, uint_t type);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>handle</i> The <i>handle</i> filled in by a call to ddi_dma_alloc_handle(9F). <i>offset</i> The offset into the object described by the <i>handle</i>. <i>length</i> 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. <i>type</i> Indicates the caller's desire about what view of the memory object to synchronize. The possible values are DDI_DMA_SYNC_FORDEV, DDI_DMA_SYNC_FORCPU and DDI_DMA_SYNC_FORKERNEL.
DESCRIPTION	ddi_dma_sync() 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. This function need only be called under certain circumstances. When resources are allocated for DMA using ddi_dma_addr_bind_handle() or ddi_dma_buf_bind_handle(), an implicit ddi_dma_sync() is done. When DMA resources are deallocated using ddi_dma_unbind_handle(9F), an implicit ddi_dma_sync() 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>did not</i> do the modifying, a call to ddi_dma_sync() 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 ddi_dma_mem_alloc(9F)) or whether or not DMA resources have been allocated for consistent mode I/O (see ddi_dma_addr_bind_handle(9F) or ddi_dma_buf_bind_handle(9F)). 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 and the time they are deallocated, you <i>must</i> call ddi_dma_sync() to ensure that either a CPU or a DMA device has such a consistent view. What to set <i>type</i> 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 <i>written</i> to the memory object, and you are going to <i>read</i> (with a CPU) the object (using an extant virtual address mapping that you have to

ddi_dma_sync(9F)

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:

<code>DDI_SUCCESS</code>	Caches are successfully flushed.
<code>DDI_FAILURE</code>	The address range to be flushed is out of the address range established by <code>ddi_dma_addr_bind_handle(9F)</code> or <code>ddi_dma_buf_bind_handle(9F)</code> .

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

ddi_dma_unbind_handle(9F)

NAME	ddi_dma_unbind_handle – unbinds the address in a DMA handle
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_dma_unbind_handle(ddi_dma_handle_t handle);</pre>
PARAMETERS	<i>handle</i> The DMA handle previously allocated by a call to ddi_dma_alloc_handle(9F).
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_dma_unbind_handle() frees all DMA resources associated with an existing DMA handle. When a DMA transfer completes, the driver should call ddi_dma_unbind_handle() to free system DMA resources established by a call to ddi_dma_buf_bind_handle(9F) or ddi_dma_addr_bind_handle(9F). ddi_dma_unbind_handle() does an implicit ddi_dma_sync(9F) making further synchronization steps unnecessary.
RETURN VALUES	DDI_SUCCESS on success DDI_FAILURE on failure
CONTEXT	ddi_dma_unbind_handle() 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_free_handle(9F), ddi_dma_sync(9F) <i>Writing Device Drivers</i>

ddi_driver_major(9F)

NAME	ddi_driver_major – return driver's major device number
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> major_t ddi_driver_major(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	ddi_driver_major() returns the major device number for the driver associated with the supplied dev_info node. This value can then be used as an argument to makedevice(9F) to construct a complete dev_t.
PARAMETERS	<i>dip</i> A pointer to the device's dev_info structure.
RETURN VALUES	ddi_driver_major() returns the major number of the driver bound to a device, if any, or DDI_MAJOR_T_NONE otherwise.
CONTEXT	ddi_driver_major() can be called from kernel or interrupt context.
SEE ALSO	ddi_driver_name(9F) <i>Writing Device Drivers</i>

NAME	ddi_driver_name – return normalized driver name
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> const char *ddi_driver_name(dev_info_t *devi);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	dip A pointer to the device's dev_info structure.
DESCRIPTION	ddi_driver_name() returns the normalized driver name. This name is typically derived from the device name property or the device compatible property. If this name is a driver alias, the corresponding driver name is returned.
RETURN VALUES	ddi_driver_name() returns the actual name of the driver bound to a device.
CONTEXT	ddi_driver_name() can be called from kernel, or interrupt context.
SEE ALSO	ddi_get_name(9F) <i>Writing Device Drivers</i>
WARNINGS	The name returned by ddi_driver_name() is read-only.

ddi_enter_critical(9F)

NAME	ddi_enter_critical, ddi_exit_critical – enter and exit a critical region of control		
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> unsigned int ddi_enter_critical(void); void ddi_exit_critical(unsignedint ddic);</pre>		
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).		
PARAMETERS	<table><tr><td><i>ddic</i></td><td>The returned value from the call to ddi_enter_critical() must be passed to ddi_exit_critical().</td></tr></table>	<i>ddic</i>	The returned value from the call to ddi_enter_critical() must be passed to ddi_exit_critical().
<i>ddic</i>	The returned value from the call to ddi_enter_critical() must be passed to ddi_exit_critical().		
DESCRIPTION	Nearly all driver operations can be done without any special synchronization and protection mechanisms beyond those provided by, for example, mutexes (see mutex(9F)). 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 ddi_enter_critical() 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 ddi_exit_critical() is made (with an argument which was the returned value from ddi_enter_critical()). The driver may not call any functions external to itself in between the time it calls ddi_enter_critical() and the time it calls ddi_exit_critical().		
RETURN VALUES	ddi_enter_critical() returns an opaque unsigned integer which must be used in the subsequent call to ddi_exit_critical().		
CONTEXT	This function can be called from user or interrupt context.		
WARNINGS	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. Do not write code bracketed by ddi_enter_critical() and ddi_exit_critical() that can get caught in an infinite loop, as the machine may crash if you do.		
SEE ALSO	mutex(9F) <i>Writing Device Drivers</i>		

NAME	ddi_ffs, ddi_fls – find first (last) bit set in a long integer
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_ffs(long mask); int ddi_fls(long mask);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>mask</i> A 32-bit argument value to search through.
DESCRIPTION	The function ddi_ffs() takes its argument and returns the shift count that the first (least significant) bit set in the argument corresponds to. The function ddi_fls() does the same, only it returns the shift count for the last (most significant) bit set in the argument.
RETURN VALUES	0 No bits are set in mask. <i>N</i> Bit <i>N</i> is the least significant (ddi_ffs) or most significant (ddi_fls) 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.
CONTEXT	This function can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

ddi_get8(9F)

NAME	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						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).						
PARAMETERS	<table><tr><td><i>handle</i></td><td>The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).</td></tr><tr><td><i>dev_addr</i></td><td>Base device address.</td></tr></table>	<i>handle</i>	The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).	<i>dev_addr</i>	Base device address.		
<i>handle</i>	The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).						
<i>dev_addr</i>	Base device address.						
DESCRIPTION	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, <i>dev_addr</i>. 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. For certain bus types, you can call these DDI functions from a high-interrupt context. These types include ISA, EISA, and SBus buses. See sysbus(4), isa(4), eisa(4), and sbus(4) for details. For the PCI bus, you can, under certain conditions, call these DDI functions from a high-interrupt context. See pci(4).						
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_put8(9F), ddi_regs_map_free(9F), ddi_regs_map_setup(9F), ddi_rep_get8(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: <table><tr><th>Previous Name</th><th>New Name</th></tr><tr><td>ddi_getb</td><td>ddi_get8</td></tr><tr><td>ddi_getw</td><td>ddi_get16</td></tr></table>	Previous Name	New Name	ddi_getb	ddi_get8	ddi_getw	ddi_get16
Previous Name	New Name						
ddi_getb	ddi_get8						
ddi_getw	ddi_get16						

ddi_get8(9F)

Previous Name	New Name
ddi_get1	ddi_get32
ddi_get11	ddi_get64

ddi_get_cred(9F)

NAME	ddi_get_cred – returns a pointer to the credential structure of the caller
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> cred_t *ddi_get_cred(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_get_cred() returns a pointer to the user credential structure of the caller.
RETURN VALUES	ddi_get_cred() returns a pointer to the caller's credential structure.
CONTEXT	ddi_get_cred() can be called from user context only.
SEE ALSO	<i>Writing Device Drivers</i>

NAME	ddi_get_devstate – Check device state
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> ddi_devstate_t ddi_get_devstate(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
PARAMETERS	<i>dip</i> Pointer to the device's dev_info structure
DESCRIPTION	The ddi_get_devstate() function returns a value indicating the state of the device specified by dip, as derived from the configuration operations that have been performed on it (or on the bus on which it resides) and any fault reports relating to it.
RETURN VALUES	DDI_DEVSTATE_OFFLINE The device is offline. In this state, the device driver is not attached, nor will it be attached automatically. The device cannot be used until it is brought online. DDI_DEVSTATE_DOWN The device is online but unusable due to a fault. DDI_DEVSTATE_QUIESCED The bus on which the device resides has been quiesced. This is not a fault, but no operations on the device should be performed while the bus remains quiesced. DDI_DEVSTATE_DEGRADED The device is online but only able to provide a partial or degraded service, due to a fault. DDI_DEVSTATE_UP The device is online and fully operational.
CONTEXT	The ddi_get_devstate() function may be called from user, kernel, or interrupt context.
NOTES	A device driver should call this function to check its own state at each major entry point, and before committing resources to a requested operation. If a driver discovers that its device is already down, it should perform required cleanup actions and return as soon as possible. If appropriate, it should return an error to its caller, indicating that the device has failed (for example, a driver's read(9E) routine would return EIO). Depending on the driver, some non-I/O operations (for example, calls to the driver's ioctl(9E) routine) may still succeed; only functions which would require fully accessible and operational hardware will necessarily fail. If the bus on which the device resides is quiesced, the driver may return a value indicating the operation should be retried later (for example, EAGAIN). Alternatively, for some classes of device, it may be appropriate for the driver to enqueue the operation and service it once the bus has been unquiesced. Note that not all busses support the quiesce/unquiesce operations, so this value may never be seen by some drivers.
SEE ALSO	attach(9E), ioctl(9E), open(9E), read(9E), strategy(9E), write(9E), ddi_dev_report_fault(9F)

ddi_get_driver_private(9F)

NAME	ddi_get_driver_private, ddi_set_driver_private – get or set the address of the device's private data area
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> void ddi_set_driver_private(dev_info_t *dip, caddr_t data); caddr_t ddi_get_driver_private(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	ddi_get_driver_private() <i>dip</i> Pointer to device information structure to get from. ddi_set_driver_private() <i>dip</i> Pointer to device information structure to set. <i>data</i> Data area address to set.
DESCRIPTION	ddi_get_driver_private() returns the address of the device's private data area from the device information structure pointed to by <i>dip</i>. ddi_set_driver_private() 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>.
RETURN VALUES	ddi_get_driver_private() returns the contents of <code>devi_driver_data</code> . If ddi_set_driver_private() has not been previously called with <i>dip</i> , an unpredictable value is returned.
CONTEXT	These functions can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

NAME	ddi_getiminator – get kernel internal minor number from an external dev_t				
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mkdev.h> #include <sys/ddi.h> minor_t ddi_getiminator(dev_t dev);</pre>				
INTERFACE LEVEL	This interface is obsolete. getminor(9F) should be used instead.				
PARAMETERS	The following parameters are supported: <i>dev</i> Device number.				
DESCRIPTION	ddi_getiminator() extracts the minor number from a device number. This call should be used only for device numbers that have been passed to the kernel from the user space through opaque interfaces such as the contents of ioctl(9E) and putmsg(2). The device numbers passed in using standard device entry points must continue to be interpreted using the getminor(9F) interface. This new interface is used to translate between user visible device numbers and in kernel device numbers. The two numbers may differ in a clustered system. For certain bus types, you can call this DDI function from a high-interrupt context. These types include ISA, EISA, and SBus buses. See sysbus(4), isa(4), eisa(4), and sbus(4) for details.				
CONTEXT	ddi_getiminator() can be called from user context only.				
RETURN VALUES	The minor number or EMINOR_UNKNOWN if the minor number of the device is invalid.				
ATTRIBUTES	See attributes(5) for a description of the following attributes: <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Stability Level</td><td>Obsolete</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Stability Level	Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Stability Level	Obsolete				
SEE ALSO	attributes(5), getmajor(9F), getminor(9F), makedevice(9F) <i>Writing Device Drivers</i>				
WARNINGS	Validity checking is performed. If <i>dev</i> is invalid, EMINOR_UNKNOWN is returned. This behavior differs from getminor(9F).				

ddi_get_instance(9F)

NAME	ddi_get_instance – get device instance number
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_get_instance(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dip</i> Pointer to dev_info structure.
DESCRIPTION	ddi_get_instance() returns the instance number of the device corresponding to <i>dip</i>. 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. 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).
RETURN VALUES	ddi_get_instance() returns the instance number of the device corresponding to <i>dip</i> .
CONTEXT	ddi_get_instance() can be called from user or interrupt context.
SEE ALSO	path_to_inst(4) <i>Writing Device Drivers</i>

NAME	ddi_get_kt_did – get identifier of current thread
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> kt_did_t ddi_get_kt_did(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	The ddi_get_kt_did() function returns a unique 64-bit identifier for the currently running thread.
CONTEXT	This routine can be called from user, kernel, or interrupt context. This routine cannot be called from a high-level interrupt context.
RETURN VALUES	ddi_get_kt_did() always returns the identifier for the current thread. There are no error conditions.
SEE ALSO	<i>Writing Device Drivers</i>
NOTES	The value returned by this function can also be seen in adb or mdb as the did field displayed when using the thread macro. This interface is intended for tracing and debugging purposes.

ddi_get_lbolt(9F)

NAME	ddi_get_lbolt – returns the value of lbolt
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> clock_t ddi_get_lbolt(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_get_lbolt() returns the value of lbolt where lbolt 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 drv_usectohz(9F) which converts microseconds into clock ticks.
RETURN VALUES	ddi_get_lbolt() returns the value of lbolt.
CONTEXT	This routine can be called from any context.
SEE ALSO	ddi_get_time(9F), drv_getparm(9F), drv_usectohz(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	ddi_get_parent – find the parent of a device information structure
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> dev_info_t *ddi_get_parent(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dip</i> Pointer to a device information structure.
DESCRIPTION	ddi_get_parent() returns a pointer to the device information structure which is the parent of the one pointed to by <i>dip</i> .
RETURN VALUES	ddi_get_parent() returns a pointer to a device information structure.
CONTEXT	ddi_get_parent() can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

ddi_get_pid(9F)

NAME	ddi_get_pid – returns the process ID
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> pid_t ddi_get_pid(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_get_pid() obtains 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 whether a device context belongs to the current process.
RETURN VALUES	ddi_get_pid() returns the process ID.
CONTEXT	This routine can be called from user context only.
SEE ALSO	drv_getparm(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	ddi_get_time – returns the current time in seconds
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> time_t ddi_get_time(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_get_time() 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.
RETURN VALUES	ddi_get_time() returns the time in seconds.
CONTEXT	This routine can be called from any context.
SEE ALSO	ddi_get_lbolt(9F), drv_getparm(9F), drv_usectoh(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

ddi_in_panic(9F)

NAME	ddi_in_panic – determine if system is in panic state
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_in_panic(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	Drivers controlling devices on which the system may write a kernel crash dump in the event of a panic can call <code>ddi_in_panic()</code> to determine if the system is panicking. When the system is panicking, the calls of functions scheduled by <code>timeout(9F)</code> and <code>ddi_trigger_softintr(9F)</code> will never occur. Neither can <code>delay(9F)</code> be relied upon, since it is implemented via <code>timeout(9F)</code>. Drivers that need to enforce a time delay such as SCSI bus reset delay time must busy-wait when the system is panicking.
RETURN VALUES	<code>ddi_in_panic()</code> returns 1 if the system is in panic, or 0 otherwise.
CONTEXT	<code>ddi_in_panic()</code> may be called from any context.
SEE ALSO	<code>dump(9E)</code> , <code>delay(9F)</code> , <code>ddi_trigger_softintr(9F)</code> , <code>timeout(9F)</code> <i>Writing Device Drivers</i>

NAME	ddi_intr_hilevel – indicate interrupt handler type
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_intr_hilevel(dev_info_t *dip, uint_t inumber);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dip</i> Pointer to dev_info structure. <i>inumber</i> Interrupt number.
DESCRIPTION	ddi_intr_hilevel() returns non-zero if the specified interrupt is a "high level" interrupt. 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. A typical high level interrupt handler would put data into a circular buffer and schedule a soft interrupt by calling ddi_trigger_softintr(). The circular buffer could be protected by using a mutex that was properly initialized for the interrupt handler. ddi_intr_hilevel() can be used before calling ddi_add_intr() 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.
RETURN VALUES	non-zero indicates a high-level interrupt.
CONTEXT	These functions can be called from user or interrupt context.
SEE ALSO	ddi_add_intr(9F), mutex(9F) <i>Writing Device Drivers</i>

ddi_io_get8(9F)

NAME	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						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> uint8_t ddi_io_get8(ddi_acc_handle_t handle, uint8_t *dev_addr); uint16_t ddi_io_get16(ddi_acc_handle_t handle, uint16_t *dev_addr); uint32_t ddi_io_get32(ddi_acc_handle_t handle, uint32_t *dev_addr);</pre>						
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>handle</i> Data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>dev_addr</i> Device address.						
DESCRIPTION	These routines generate a read of various sizes from the device address, <i>dev_addr</i>, in I/O space. The ddi_io_get8(), ddi_io_get16(), and ddi_io_get32() functions read 8 bits, 16 bits, and 32 bits of data, respectively, from the device address, <i>dev_addr</i>. 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.						
CONTEXT	These functions can be called from user, kernel, or interrupt context.						
SEE ALSO	isa(4), ddi_io_put8(9F), ddi_io_rep_get8(9F), ddi_io_rep_put8(9F), ddi_regs_map_free(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: <table><thead><tr><th>Previous Name</th><th>New Name</th></tr></thead><tbody><tr><td>ddi_io_getb</td><td>ddi_io_get8</td></tr><tr><td>ddi_io_getw</td><td>ddi_io_get16</td></tr></tbody></table>	Previous Name	New Name	ddi_io_getb	ddi_io_get8	ddi_io_getw	ddi_io_get16
Previous Name	New Name						
ddi_io_getb	ddi_io_get8						
ddi_io_getw	ddi_io_get16						

ddi_io_get8(9F)

Previous Name	New Name
ddi_io_get1	ddi_io_get32

ddi_iomin(9F)

NAME	ddi_iomin – find minimum alignment and transfer size for DMA						
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_iomin(dev_info_t *dip, int initial, int streaming);</pre>						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).						
PARAMETERS	<table><tr><td><i>dip</i></td><td>A pointer to the device's dev_info structure.</td></tr><tr><td><i>initial</i></td><td>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.</td></tr><tr><td><i>streaming</i></td><td>This argument, if non-zero, indicates that the returned value should be modified to account for <i>streaming</i> mode accesses (see ddi_dma_req(9S) for a discussion of streaming versus non-streaming access mode).</td></tr></table>	<i>dip</i>	A pointer to the device's dev_info structure.	<i>initial</i>	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.	<i>streaming</i>	This argument, if non-zero, indicates that the returned value should be modified to account for <i>streaming</i> mode accesses (see ddi_dma_req(9S) for a discussion of streaming versus non-streaming access mode).
<i>dip</i>	A pointer to the device's dev_info structure.						
<i>initial</i>	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.						
<i>streaming</i>	This argument, if non-zero, indicates that the returned value should be modified to account for <i>streaming</i> mode accesses (see ddi_dma_req(9S) for a discussion of streaming versus non-streaming access mode).						
DESCRIPTION	ddi_iomin(), 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.						
RETURN VALUES	ddi_iomin() returns the minimum DMA transfer size for the calling device, or it returns zero, which means that you cannot get there from here.						
CONTEXT	This function can be called from user or interrupt context.						
SEE ALSO	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>						

NAME	ddi_iopb_alloc, ddi_iopb_free – allocate and free non-sequentially accessed memory	
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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 iopb);</pre>	
INTERFACE LEVEL	These interfaces are obsolete. Use ddi_dma_mem_alloc(9F) instead of ddi_iopb_alloc(). Use ddi_dma_mem_free(9F) instead of ddi_iopb_free().	
PARAMETERS		
ddi_iopb_alloc()	<i>dip</i>	A pointer to the device's dev_info structure.
	<i>limits</i>	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.
	<i>length</i>	The length in bytes of the desired allocation.
	<i>iopbp</i>	A pointer to a caddr_t. On a successful return, *iopbp points to the allocated storage.
ddi_iopb_free()	<i>iopb</i>	The iopb returned from a successful call to ddi_iopb_alloc().
DESCRIPTION	ddi_iopb_alloc() allocates memory for DMA transfers and should be used if the device accesses memory in a non-sequential fashion, or if synchronization steps using ddi_dma_sync(9F) 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. Note that you still must use DMA resource allocation functions (see ddi_dma_setup(9F)) to establish DMA resources for the memory allocated using ddi_iopb_alloc(). 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 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.	

ddi_iopb_alloc(9F)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), ddi_dma_free(9F), ddi_dma_mem_alloc(9F), ddi_dma_mem_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.

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_io_put8(ddi_acc_handle_t handle, uint8_t *dev_addr, uint8_t value); void ddi_io_put16(ddi_acc_handle_t handle, uint16_t *dev_addr, uint16_t value); void ddi_io_put32(ddi_acc_handle_t handle, uint32_t *dev_addr, uint32_t value);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>handle</i> Data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>dev_addr</i> Base device address. <i>value</i> Data to be written to the device.
DESCRIPTION	These routines generate a write of various sizes to the device address, <i>dev_addr</i>, in I/O space. The ddi_io_put8(), ddi_io_put16(), and ddi_io_put32() functions write 8 bits, 16 bits, and 32 bits of data, respectively, to the device address, <i>dev_addr</i>. 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.
CONTEXT	These functions can be called from user, kernel, or interrupt context.
SEE ALSO	isa(4), ddi_io_get8(9F), ddi_io_rep_get8(9F), ddi_io_rep_put8(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
ddi_io_putb	ddi_io_put8

ddi_io_put8(9F)

Previous Name	New Name
ddi_io_putw	ddi_io_put16
ddi_io_putl	ddi_io_put32

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_io_rep_get8(ddi_acc_handle_t handle, uint8_t *host_addr, uint8_t *dev_addr, , size_t repcount); void ddi_io_rep_get16(ddi_acc_handle_t handle, uint16_t *host_addr, uint16_t *dev_addr, , size_t repcount); void ddi_io_rep_get32(ddi_acc_handle_t handle, uint32_t *host_addr, uint32_t *dev_addr, , size_t repcount);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>handle</i> The data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>host_addr</i> Base host address. <i>dev_addr</i> Base device address. <i>repcount</i> Number of data accesses to perform.
DESCRIPTION	These routines generate multiple reads from the device address, <i>dev_addr</i>, in I/O space. <i>repcount</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 ddi_io_rep_get8(), ddi_io_rep_get16(), and ddi_io_rep_get32() functions read 8 bits, 16 bits, and 32 bits of data, respectively, from the device address. <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.
CONTEXT	These functions can be called from user, kernel, or interrupt context.
SEE ALSO	isa(4), ddi_io_get8(9F), ddi_io_put8(9F), ddi_io_rep_put8(9F), ddi_regs_map_free(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.

ddi_io_rep_get8(9F)

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_io_rep_getb	ddi_io_rep_get8
ddi_io_rep_getw	ddi_io_rep_get16
ddi_io_rep_getl	ddi_io_rep_get32

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_io_rep_put8(ddi_acc_handle_t handle, uint8_t *host_addr, uint8_t *dev_addr, size_t repcount); void ddi_io_rep_put16(ddi_acc_handle_t handle, uint16_t *host_addr, uint16_t *dev_addr, size_t repcount); void ddi_io_rep_put32(ddi_acc_handle_t handle, uint32_t *host_addr, uint32_t *dev_addr, size_t repcount);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>handle</i> Data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>host_addr</i> Base host address. <i>dev_addr</i> Base device address. <i>repcount</i> Number of data accesses to perform.
DESCRIPTION	These routines generate multiple writes to the device address, <i>dev_address</i>, in I/O space. <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 ddi_io_rep_put8(), ddi_io_rep_put16(), and ddi_io_rep_put32() functions write 8 bits, 16 bits, and 32 bits of data, respectively, to the device address. <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.
CONTEXT	These functions can be called from user, kernel, or interrupt context.
SEE ALSO	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:

ddi_io_rep_put8(9F)

Previous Name	New Name
ddi_io_rep_putb	ddi_io_rep_put8
ddi_io_rep_putw	ddi_io_rep_put16
ddi_io_rep_putl	ddi_io_rep_put32

NAME	ddi_log_sysevent – log system event for drivers														
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_log_sysevent(dev_info_t *dip, char *vendor, char *class, char *subclass, nvlist_t *attr_list, sysevent_id_t *eidp, int sleep_flag);</pre>														
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>dip</i></td><td>A pointer to the dev_info node for this driver.</td></tr> <tr> <td><i>vendor</i></td><td>A pointer to a string defining the vendor. Third-party drivers should use their company's stock symbol (or similarly enduring identifier). Sun-supplied drivers should use DDI_VENDOR_SUNW.</td></tr> <tr> <td><i>class</i></td><td>A pointer to a string defining the event class.</td></tr> <tr> <td><i>subclass</i></td><td>A pointer to a string defining the event subclass.</td></tr> <tr> <td><i>attr_list</i></td><td>A pointer to an nvlist_t, listing the name-value attributes associated with the event or NULL if there are no such attributes for this event.</td></tr> <tr> <td><i>eidp</i></td><td>The address of a sysevent_id_t structure in which the event's sequence number and timestamp are returned if the event is successfully queued. May be NULL if this information is not of interest. See below for the definition of sysevent_id_t.</td></tr> <tr> <td><i>sleep_flag</i></td><td>Indicates how a caller wants to handle the possibility of resources not being available. If <i>sleep_flag</i> is DDI_NOSLEEP, the caller does not care if the allocation fails or the queue is full and can handle a failure appropriately. If <i>sleep_flag</i> is DDI_SLEEP, the caller wishes to have the allocation and queuing routines wait for resources to become available.</td></tr> </table>	<i>dip</i>	A pointer to the dev_info node for this driver.	<i>vendor</i>	A pointer to a string defining the vendor. Third-party drivers should use their company's stock symbol (or similarly enduring identifier). Sun-supplied drivers should use DDI_VENDOR_SUNW.	<i>class</i>	A pointer to a string defining the event class.	<i>subclass</i>	A pointer to a string defining the event subclass.	<i>attr_list</i>	A pointer to an nvlist_t, listing the name-value attributes associated with the event or NULL if there are no such attributes for this event.	<i>eidp</i>	The address of a sysevent_id_t structure in which the event's sequence number and timestamp are returned if the event is successfully queued. May be NULL if this information is not of interest. See below for the definition of sysevent_id_t.	<i>sleep_flag</i>	Indicates how a caller wants to handle the possibility of resources not being available. If <i>sleep_flag</i> is DDI_NOSLEEP, the caller does not care if the allocation fails or the queue is full and can handle a failure appropriately. If <i>sleep_flag</i> is DDI_SLEEP, the caller wishes to have the allocation and queuing routines wait for resources to become available.
<i>dip</i>	A pointer to the dev_info node for this driver.														
<i>vendor</i>	A pointer to a string defining the vendor. Third-party drivers should use their company's stock symbol (or similarly enduring identifier). Sun-supplied drivers should use DDI_VENDOR_SUNW.														
<i>class</i>	A pointer to a string defining the event class.														
<i>subclass</i>	A pointer to a string defining the event subclass.														
<i>attr_list</i>	A pointer to an nvlist_t, listing the name-value attributes associated with the event or NULL if there are no such attributes for this event.														
<i>eidp</i>	The address of a sysevent_id_t structure in which the event's sequence number and timestamp are returned if the event is successfully queued. May be NULL if this information is not of interest. See below for the definition of sysevent_id_t.														
<i>sleep_flag</i>	Indicates how a caller wants to handle the possibility of resources not being available. If <i>sleep_flag</i> is DDI_NOSLEEP, the caller does not care if the allocation fails or the queue is full and can handle a failure appropriately. If <i>sleep_flag</i> is DDI_SLEEP, the caller wishes to have the allocation and queuing routines wait for resources to become available.														
DESCRIPTION	ddi_log_sysevent() causes a system event, of the specified class and subclass, to be generated on behalf of the driver and queued for delivery to syseventd, the user-land sysevent daemon. The publisher string for the event is constructed using the vendor name and driver name, with the format: <pre>"<vendor>:kern:<driver-name>"</pre> The two fields of eidp, eid_seq and eid_ts, are sufficient to uniquely identify an event.														
STRUCTURE MEMBERS	The structure members of sysevent_id_t are: <pre>uint64_t eid_seq; /* sysevent sequence number */ hrtime_t eid_ts; /* sysevent timestamp */</pre>														
RETURN VALUES	ddi_log_sysevent() returns:														

ddi_log_sysevent(9F)

DDI_SUCCESS

The event has been queued for delivery successfully.

DDI_ENOMEM

There is not enough memory to queue the system event at this time. DDI_ENOMEM cannot be returned when *sleep_flag* is DDI_SLEEP.

DDI_EBUSY

The system event queue is full at this time. DDI_EBUSY cannot be returned when *sleep_flag* is DDI_SLEEP.

DDI_ETRANSPORT

The syseventd daemon is not responding and events cannot be queued or delivered at this time. DDI_ETRANSPORT can be returned even when *sleep_flag* is DDI_SLEEP.

DDI_ECONTEXT

sleep_flag is DDI_SLEEP and the driver is running in interrupt context.

CONTEXT

ddi_log_sysevent() can be called from user or interrupt context, except when *sleep_flag* is DDI_SLEEP, in which case it can be called from user context only.

EXAMPLES

EXAMPLE 1 Logging System Event with No Attributes

```
if (ddi_log_sysevent(dip, DDI_VENDOR_SUNW, "class", "subclass",
    NULL, NULL, DDI_SLEEP) != DDI_SUCCESS) {
    cmn_err(CE_WARN, "error logging system event\n");
}
```

EXAMPLE 2 Logging System Event with Two Name/Value Attributes, an Integer and a String

```
nvlist_t    *attr_list;
sysevent_id_t eid;

if (nvlist_alloc(&attr_list, NV_UNIQUE_NAME_TYPE, NV_FLAG_KMSLEEP) == 0)
{
    err = nvlist_add_uint32(attr_list, int_name, int_value);
    if (err == 0)
        err = nvlist_add_string(attr_list, str_name, str_value);
    if (err == 0)
        err = ddi_log_sysevent(dip, DDI_VENDOR_SUNW,
            "class", "subclass", attr_list, &eid, DDI_SLEEP);
    if (err != DDI_SUCCESS)
        cmn_err(CE_WARN, "error logging system event\n");
    nvlist_free(attr_list);
}
```

EXAMPLE 3 Use Timeout to Handle nvlist and System Event Resource Allocation Failures

Since no blocking calls are made, this example would be useable from a driver needing to generate an event from interrupt context.

```
static int
xx_se_timeout_handler(xx_state_t *xx)
{
```

EXAMPLE 3 Use Timeout to Handle nvlist and System Event Resource Allocation Failures (Continued)

```

        xx->xx_timeoutid = (xx_generate_event(xx) ?
            timeout(xx_se_timeout_handler, xx, 4) : 0);
    }

    static int
    xx_generate_event(xx_state_t *xx)
    {
        int err;

        err = nvlist_alloc(&xx->xx_ev_attrlist, NV_UNIQUE_NAME_TYPE, 0);
        if (err != 0)
            return (1);
        err = nvlist_add_uint32(&xx->xx_ev_attrlist,
            xx->xx_ev_name, xx->xx_ev_value);
        if (err != 0) {
            nvlist_free(xx->xx_ev_attrlist);
            return(1);
        }

        err = ddi_log_sysevent(xx->xx_dip, DDI_VENDOR_SUNW,
            xx->xx_ev_class, xx->xx_ev_sbclass,
            xx->xx_ev_attrlist, NULL, DDI_NOSLEEP);
        nvlist_free(xx->xx_ev_attrlist);
        if (err == DDI_SUCCESS || err == DDI_ETRANSPORT) {
            if (err == DDI_ETRANSPORT)
                cmn_err(CE_WARN, "cannot log system event\n");
            return (0);
        }
        return (1);
    }

```

SEE ALSO syseventd(1M), attributes(5), nvlist_add_boolean(9F), nvlist_alloc(9F)

Writing Device Drivers

ddi_mapdev(9F)

NAME	ddi_mapdev – create driver-controlled mapping of device																								
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_mapdev(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 *cred, struct ddi_mapdev_ctl *ctl, ddi_mapdev_handle_t *handlep, void *devprivate);</pre>																								
INTERFACE LEVEL PARAMETERS	This interface is obsolete. devmap_setup(9F) should be used instead. <table> <tr> <td><i>dev</i></td><td>The device whose memory is to be mapped.</td></tr> <tr> <td><i>offset</i></td><td>The offset within device memory at which the mapping begins.</td></tr> <tr> <td><i>as</i></td><td>An opaque pointer to the user address space into which the device memory should be mapped.</td></tr> <tr> <td><i>addrp</i></td><td>Pointer to the starting address within the user address space to which the device memory should be mapped.</td></tr> <tr> <td><i>len</i></td><td>Length (in bytes) of the memory to be mapped.</td></tr> <tr> <td><i>prot</i></td><td>A bit field that specifies the protections.</td></tr> <tr> <td><i>maxprot</i></td><td>Maximum protection flag possible for attempted mapping.</td></tr> <tr> <td><i>flags</i></td><td>Flags indicating type of mapping.</td></tr> <tr> <td><i>cred</i></td><td>Pointer to the user credentials structure.</td></tr> <tr> <td><i>ctl</i></td><td>A pointer to a ddi_mapdev_ctl(9S) structure. The structure contains pointers to device driver-supplied functions that manage events on the device mapping.</td></tr> <tr> <td><i>handlep</i></td><td>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.</td></tr> <tr> <td><i>devprivate</i></td><td>Driver private mapping data. This value is passed into each mapping call back routine.</td></tr> </table>	<i>dev</i>	The device whose memory is to be mapped.	<i>offset</i>	The offset within device memory at which the mapping begins.	<i>as</i>	An opaque pointer to the user address space into which the device memory should be mapped.	<i>addrp</i>	Pointer to the starting address within the user address space to which the device memory should be mapped.	<i>len</i>	Length (in bytes) of the memory to be mapped.	<i>prot</i>	A bit field that specifies the protections.	<i>maxprot</i>	Maximum protection flag possible for attempted mapping.	<i>flags</i>	Flags indicating type of mapping.	<i>cred</i>	Pointer to the user credentials structure.	<i>ctl</i>	A pointer to a ddi_mapdev_ctl(9S) structure. The structure contains pointers to device driver-supplied functions that manage events on the device mapping.	<i>handlep</i>	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.	<i>devprivate</i>	Driver private mapping data. This value is passed into each mapping call back routine.
<i>dev</i>	The device whose memory is to be mapped.																								
<i>offset</i>	The offset within device memory at which the mapping begins.																								
<i>as</i>	An opaque pointer to the user address space into which the device memory should be mapped.																								
<i>addrp</i>	Pointer to the starting address within the user address space to which the device memory should be mapped.																								
<i>len</i>	Length (in bytes) of the memory to be mapped.																								
<i>prot</i>	A bit field that specifies the protections.																								
<i>maxprot</i>	Maximum protection flag possible for attempted mapping.																								
<i>flags</i>	Flags indicating type of mapping.																								
<i>cred</i>	Pointer to the user credentials structure.																								
<i>ctl</i>	A pointer to a ddi_mapdev_ctl(9S) structure. The structure contains pointers to device driver-supplied functions that manage events on the device mapping.																								
<i>handlep</i>	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.																								
<i>devprivate</i>	Driver private mapping data. This value is passed into each mapping call back routine.																								
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 <i>ctl</i>. The user events that the driver is notified of are: <table> <tr> <td>access</td><td>User has accessed an address in the mapping that has no translations.</td></tr> </table>	access	User has accessed an address in the mapping that has no translations.																						
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 driver's `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 driver's `segmap()` entry point.

CONTEXT

This routine can be called from user or kernel context only.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO

`fork(2)`, `mmap(2)`, `munmap(2)`, `attributes(5)`, `mapdev_access(9E)`, `mapdev_dup(9E)`, `mapdev_free(9E)`, `mmap(9E)`, `segmap(9E)`, `devmap_setup(9F)`, `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

ddi_mapdev(9F)

NOTES	Only mappings of type <code>MAP_PRIVATE</code> should be used with <code>ddi_mapdev()</code> .

NAME	ddi_mapdev_intercept, ddi_mapdev_nointercept – control driver notification of user accesses
SYNOPSIS	<pre>#include <sys/sunddi.h> 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>
INTERFACE LEVEL	These interfaces are obsolete. Use devmap_load(9F) instead of ddi_mapdev_intercept(). Use devmap_unload(9F) instead of ddi_mapdev_nointercept().
PARAMETERS	<i>handle</i> An opaque pointer to a device mapping handle. <i>offset</i> An offset in bytes within device memory. <i>len</i> Length in bytes.
DESCRIPTION	Future releases of Solaris will provide these functions for binary and source compatibility. However, for increased functionality, use devmap_load(9F) or devmap_unload(9F) instead. See devmap_load(9F) and devmap_unload(9F) for details. The ddi_mapdev_intercept() and ddi_mapdev_nointercept() functions control whether or not user accesses to device mappings created by ddi_mapdev(9F) in the specified range will generate calls to the mapdev_access(9E) entry point. ddi_mapdev_intercept() tells the system to intercept the user access and notify the driver to invalidate the mapping translations. ddi_mapdev_nointercept() tells the system to not intercept the user access and allow it to proceed by validating the mapping translations. For both routines, the range to be affected is defined by the <i>offset</i> and <i>len</i> arguments. Requests affect the entire page containing the <i>offset</i> and all pages up to and including the page containing the last byte as indicated by <i>offset + len</i>. Supplying a value of 0 for the <i>len</i> argument affects all addresses from the <i>offset</i> to the end of the mapping. Supplying a value of 0 for the <i>offset</i> argument and a value of 0 for <i>len</i> argument affect all addresses in the mapping. To manage a device context, a device driver would call ddi_mapdev_intercept() on the context about to be switched out, switch contexts, and then call ddi_mapdev_nointercept() on the context switched in.
RETURN VALUES	ddi_mapdev_intercept() and ddi_mapdev_nointercept() return the following values: 0 Successful completion. Non-zero An error occurred.

ddi_mapdev_intercept(9F)

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)
{
    int err;
    /* enable access callbacks for the current mapping */
    if (cur_hdl != NULL) {
        if ((err = ddi_mapdev_intercept(cur_hdl, offset, 0)) != 0)
            return (err);
    }
    /* Switch device context - device dependent*/
    ...
    /* Make handle the new current mapping */
    cur_hdl = handle;
    /*
     * Disable callbacks and complete the access for the
     * mapping that generated this callback.
     */
    return (ddi_mapdev_nointercept(handle, offset, 0));
}
```

CONTEXT

These routines can be called from user or kernel context only.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO

attributes(5), mapdev_access(9E), devmap_load(9F), ddi_mapdev(9F)

Writing Device Drivers

NAME	ddi_mapdev_set_device_acc_attr – set the device attributes for the mapping										
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>										
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>mapping_handle</i></td><td>A pointer to a device mapping handle.</td></tr> <tr> <td><i>offset</i></td><td>The offset within device memory to which the device access attributes structure applies.</td></tr> <tr> <td><i>len</i></td><td>Length (in bytes) of the memory to which the device access attributes structure applies.</td></tr> <tr> <td><i>*accattrp</i></td><td>Pointer to a ddi_device_acc_attr(9S) structure. Contains the device access attributes to be applied to this range of memory.</td></tr> <tr> <td><i>rnumber</i></td><td>Index number to the register address space set.</td></tr> </table>	<i>mapping_handle</i>	A pointer to a device mapping handle.	<i>offset</i>	The offset within device memory to which the device access attributes structure applies.	<i>len</i>	Length (in bytes) of the memory to which the device access attributes structure applies.	<i>*accattrp</i>	Pointer to a ddi_device_acc_attr(9S) structure. Contains the device access attributes to be applied to this range of memory.	<i>rnumber</i>	Index number to the register address space set.
<i>mapping_handle</i>	A pointer to a device mapping handle.										
<i>offset</i>	The offset within device memory to which the device access attributes structure applies.										
<i>len</i>	Length (in bytes) of the memory to which the device access attributes structure applies.										
<i>*accattrp</i>	Pointer to a ddi_device_acc_attr(9S) structure. Contains the device access attributes to be applied to this range of memory.										
<i>rnumber</i>	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 devmap(9E) instead. See devmap(9E) for details. The ddi_mapdev_set_device_acc_attr() function assigns device access attributes to a range of device memory in the register set given by <i>rnumber</i>. <i>*accattrp</i> defines the device access attributes. See ddi_device_acc_attr(9S) for more details. <i>mapping_handle</i> is a mapping handle returned from a call to ddi_mapdev(9F). The range to be affected is defined by the <i>offset</i> and <i>len</i> arguments. Requests affect the entire page containing the <i>offset</i> and all pages up to and including the page containing the last byte as indicated by <i>offset+len</i>. Supplying a value of 0 for the <i>len</i> argument affects all addresses from the <i>offset</i> to the end of the mapping. Supplying a value of 0 for the <i>offset</i> argument and a value of 0 for the <i>len</i> argument affect all addresses in the mapping.										
RETURN VALUES	The ddi_mapdev_set_device_acc_attr() function returns the following values: <table> <tr> <td>DDI_SUCCESS</td><td>The attributes were successfully set.</td></tr> <tr> <td>DDI_FAILURE</td><td>It is not possible to set these attributes for this mapping handle.</td></tr> </table>	DDI_SUCCESS	The attributes were successfully set.	DDI_FAILURE	It is not possible to set these attributes for this mapping handle.						
DDI_SUCCESS	The attributes were successfully set.										
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.										

ddi_mapdev_set_device_acc_attr(9F)

SEE ALSO | segmap(9E), ddi_mapdev(9F), ddi_segmap_setup(9F),
ddi_device_acc_attr(9S)

Writing Device Drivers

NAME	ddi_map_regs, ddi_unmap_regs – map or unmap registers	
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> 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>	
INTERFACE LEVEL	These interfaces are obsolete. Use ddi_regs_map_setup(9F) instead of ddi_map_regs(). Use ddi_regs_map_free(9F) instead of ddi_unmap_regs().	
PARAMETERS		
ddi_map_regs()	<i>dip</i>	Pointer to the device's dev_info structure.
	<i>rnumber</i>	Register set number.
	<i>kaddrp</i>	Pointer to the base kernel address of the mapped region (set on return).
	<i>offset</i>	Offset into register space.
	<i>len</i>	Length to be mapped.
ddi_unmap_regs()	<i>dip</i>	Pointer to the device's dev_info structure.
	<i>rnumber</i>	Register set number.
	<i>kaddrp</i>	Pointer to the base kernel address of the region to be unmapped.
	<i>offset</i>	Offset into register space.
	<i>len</i>	Length to be unmapped.
DESCRIPTION	ddi_map_regs() 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 <code>sbus(4)</code> and for more information on register set descriptions. If <i>len</i> and <i>offset</i> 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. <i>rnumber</i>	

ddi_map_regs(9F)

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.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `attributes(5)`, `sbus(4)`, `ddi_regs_map_free(9F)`, `ddi_regs_map_setup(9F)`

Writing Device Drivers

NAME	ddi_mem_alloc, ddi_mem_free – allocate and free sequentially accessed memory	
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>	
INTERFACE LEVEL	These interfaces are obsolete. ddi_dma_mem_alloc(9F) and ddi_dma_mem_free(9F) should be used instead.	
PARAMETERS		
ddi_mem_alloc()	<i>dip</i>	A pointer to the device's dev_info structure.
	<i>limits</i>	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.
	<i>length</i>	The length in bytes of the desired allocation.
	<i>flags</i>	The possible flags 1 and 0 are taken to mean, respectively, wait until memory is available, or do not wait.
	<i>kaddrp</i>	On a successful return, *kaddrp points to the allocated memory.
	<i>real_length</i>	The length in bytes that was allocated. Alignment and padding requirements may cause ddi_mem_alloc() to allocate more memory than requested in <i>length</i> .
ddi_mem_free()	<i>kaddr</i>	The memory returned from a successful call to ddi_mem_alloc().
DESCRIPTION	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 <i>limits</i> argument and other limits imposed by the system. 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_mem_alloc(9F)

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.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `attributes(5)`, `ddi_dma_free(9F)`, `ddi_dma_mem_alloc(9F)`,
`ddi_dma_mem_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

ddi_mem_get8(9F)

NAME	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				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).				
PARAMETERS	<table><tr><td><i>handle</i></td><td>The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).</td></tr><tr><td><i>dev_addr</i></td><td>Base device address.</td></tr></table>	<i>handle</i>	The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).	<i>dev_addr</i>	Base device address.
<i>handle</i>	The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).				
<i>dev_addr</i>	Base device address.				
DESCRIPTION	These routines generate a read of various sizes from memory space or allocated DMA memory. The ddi_mem_get8(), ddi_mem_get16(), ddi_mem_get32(), and ddi_mem_get64() 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. 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.				
CONTEXT	These functions can be called from user, kernel, or interrupt context.				
SEE ALSO	ddi_mem_put8(9F), ddi_mem_rep_get8(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
ddi_mem_getb	ddi_mem_get8
ddi_mem_getw	ddi_mem_get16
ddi_mem_getl	ddi_mem_get32
ddi_mem_getll	ddi_mem_get64

ddi_mem_put8(9F)

NAME	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						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>						
PARAMETERS	<table><tr><td><i>handle</i></td><td>The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).</td></tr><tr><td><i>dev_addr</i></td><td>Base device address.</td></tr><tr><td><i>value</i></td><td>The data to be written to the device.</td></tr></table>	<i>handle</i>	The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).	<i>dev_addr</i>	Base device address.	<i>value</i>	The data to be written to the device.
<i>handle</i>	The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).						
<i>dev_addr</i>	Base device address.						
<i>value</i>	The data to be written to the device.						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).						
DESCRIPTION	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, <i>dev_addr</i>, in memory space. 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.						
CONTEXT	These functions can be called from user, kernel, or interrupt context.						
SEE ALSO	ddi_mem_get8(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: <table><tr><th>Previous Name</th><th>New Name</th></tr><tr><td>ddi_mem_putb</td><td>ddi_mem_put8</td></tr><tr><td>ddi_mem_putw</td><td>ddi_mem_put16</td></tr></table>	Previous Name	New Name	ddi_mem_putb	ddi_mem_put8	ddi_mem_putw	ddi_mem_put16
Previous Name	New Name						
ddi_mem_putb	ddi_mem_put8						
ddi_mem_putw	ddi_mem_put16						

ddi_mem_put8(9F)

Previous Name	New Name
ddi_mem_put1	ddi_mem_put32
ddi_mem_put11	ddi_mem_put64

ddi_mem_rep_get8(9F)

NAME	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		
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_mem_rep_get8(ddi_acc_handle_t handle, uint8_t *host_addr, uint8_t *dev_addr, size_t repcount, uint_t flags); void ddi_mem_rep_get16(ddi_acc_handle_t handle, uint16_t *host_addr, uint16_t *dev_addr, size_t repcount, uint_t flags); void ddi_mem_rep_get32(ddi_acc_handle_t handle, uint32_t *host_addr, uint32_t *dev_addr, size_t repcount, uint_t flags); void ddi_mem_rep_get64(ddi_acc_handle_t handle, uint64_t *host_addr, uint64_t *dev_addr, size_t repcount, uint_t flags);</pre>		
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI).		
	<i>handle</i>	The data access handle returned from setup calls, such as ddi_regs_map_setup(9F).	
	<i>host_addr</i>	Base host address.	
	<i>dev_addr</i>	Base device address.	
	<i>repcount</i>	Number of data accesses to perform.	
	<i>flags</i>	Device address flags:	
		DDI_DEV_AUTOINCR	Automatically increment the device address, <i>dev_addr</i> , during data accesses.
		DDI_DEV_NO_AUTOINCR	Do not advance the device address, <i>dev_addr</i> , during data accesses.
DESCRIPTION	These routines generate multiple reads from memory space or allocated DMA memory. <i>repcount</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 ddi_mem_rep_get8(), ddi_mem_rep_get16(), ddi_mem_rep_get32(), and ddi_mem_rep_get64() 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.		

ddi_mem_rep_get8(9F)

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
ddi_mem_rep_getb	ddi_mem_rep_get8
ddi_mem_rep_getw	ddi_mem_rep_get16
ddi_mem_rep_getl	ddi_mem_rep_get32
ddi_mem_rep_getll	ddi_mem_rep_get64

ddi_mem_rep_put8(9F)

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_mem_rep_put8(ddi_acc_handle_t handle, uint8_t *host_addr, uint8_t *dev_addr, size_t repcount, uint_t flags); void ddi_mem_rep_put16(ddi_acc_handle_t handle, uint16_t *host_addr, uint16_t *dev_addr, size_t repcount, uint_t flags); void ddi_mem_rep_put32(ddi_acc_handle_t handle, uint32_t *host_addr, uint32_t *dev_addr, size_t repcount, uint_t flags); void ddi_mem_rep_put64(ddi_acc_handle_t handle, uint64_t *host_addr, uint64_t *dev_addr, size_t repcount, uint_t flags);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>handle</i> The data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>host_addr</i> Base host address. <i>dev_addr</i> Base device address. <i>repcount</i> Number of data accesses to perform. <i>flags</i> Device address flags: DDI_DEV_AUTOINCR Automatically increment the device address, <i>dev_addr</i>, during data accesses. DDI_DEV_NO_AUTOINCR Do not advance the device address, <i>dev_addr</i>, during data accesses.
DESCRIPTION	These routines generate multiple writes to memory space or allocated DMA memory. <i>repcount</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 ddi_mem_rep_put8(), ddi_mem_rep_put16(), ddi_mem_rep_put32(), and ddi_mem_rep_put64() functions write 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, to the device address. <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.

ddi_mem_rep_put8(9F)

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
ddi_mem_rep_putb	ddi_mem_rep_put8
ddi_mem_rep_putw	ddi_mem_rep_put16
ddi_mem_rep_putl	ddi_mem_rep_put32
ddi_mem_rep_putll	ddi_mem_rep_put64

ddi_mmap_get_model(9F)

NAME	ddi_mmap_get_model – return data model type of current thread						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> uint_t ddi_mmap_get_model(void);</pre>						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).						
DESCRIPTION	ddi_mmap_get_model() returns the C Language Type Model which the current thread expects. ddi_mmap_get_model() is used in combination with ddi_model_convert_from(9F) in the mmap(9E) 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.						
RETURN VALUES	<table> <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 ddi_mmap_get_model() function was not called from the mmap(9E) 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 ddi_mmap_get_model() function was not called from the mmap(9E) 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 ddi_mmap_get_model() function was not called from the mmap(9E) entry point.						
CONTEXT	The ddi_mmap_get_model() function can only be called from the mmap(9E) driver entry point.						
EXAMPLES	EXAMPLE 1: Using ddi_mmap_get_model() The following is an example of the mmap(9E) entry point and how to support 32-bit and 64-bit applications with the same device driver. <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; switch (ddi_model_convert_from(ddi_mmap_get_model())) { case DDI_MODEL_ILP32: { struct data32 *da32p; da32p = (struct data32 *)shared_area; dp = &dtc; dp->len = da32p->len; dp->address = da32p->address; break; } }</pre>						

EXAMPLE 1: Using `ddi_mmap_get_model()` (Continued)

```

        case DDI_MODEL_NONE:
            break;
    }
    /* continues along using dp */
    ...
}

```

SEE ALSO `mmap(9E)`, `ddi_model_convert_from(9F)`

Writing Device Drivers

ddi_model_convert_from(9F)

NAME	ddi_model_convert_from – determine data model type mismatch	
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> uint_t ddi_model_convert_from(uint_t model);</pre>	
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).	
PARAMETERS	<i>model</i>	The data model type of the current thread.
DESCRIPTION	ddi_model_convert_from() 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 ioctl(9E) and mmap(9E) where it is necessary to identify the C Language Type Model of the user-mode originator of an 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. ddi_model_convert_from() is used to determine if data that is passed between the device driver and the application requires reformatting to any non-native data model.	
RETURN VALUES	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.
CONTEXT	ddi_model_convert_from() can be called from any context.	
EXAMPLES	EXAMPLE 1: Using ddi_model_convert_from() in the ioctl() entry point to support both 32-bit and 64-bit applications. The following is an example how to use ddi_model_convert_from() in the ioctl() entry point to support both 32-bit and 64-bit applications. <pre>struct passargs32 { int len; caddr32_t addr; }; 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; switch (ddi_model_convert_from(mode & FMODELS)) { case DDI_MODEL_ILP32: {</pre>	

EXAMPLE 1: Using `ddi_model_convert_from()` in the `ioctl()` entry point to support both 32-bit and 64-bit applications. (*Continued*)

```

        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;
    }

    do_ioctl(&pa);
    . . . .
}

```

SEE ALSO `ioctl(9E)`, `mmap(9E)`, `ddi_mmap_get_model(9F)`

Writing Device Drivers

ddi_node_name(9F)

NAME	ddi_node_name – return the devinfo node name
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> char *ddi_node_name(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dip</i> A pointer the device's dev_info structure.
DESCRIPTION	ddi_node_name() returns the device node name contained in the dev_info node pointed to by <i>dip</i> .
RETURN VALUES	ddi_node_name() returns the device node name contained in the dev_info structure.
CONTEXT	ddi_node_name() can be called from user or interrupt context.
SEE ALSO	ddi_binding_name(9F) <i>Writing Device Drivers</i>

NAME	ddi_no_info – stub for getinfo(9E)
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_no_info(dev_info_t *dip, ddi_info_cmd_t infocmd, void *arg, void **result);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) dev_info_t *dip Pointer to dev_info structure. ddi_info_cmd_t infocmd Command argument. Valid command values are: DDI_INFO_DEVT2DEVINFO and DDI_INFO_DEVT2INSTANCE. void *arg Command-specific argument. void **result Pointer to where the requested information is stored.
DESCRIPTION	The ddi_no_info() function always returns DDI_FAILURE. It is provided as a convenience routine for drivers not providing a cb_ops(9S) or for network drivers only providing DLPI-2 services. Such drivers can use ddi_no_info() in the devo_getinfo entry point (see getinfo(9E)) of the dev_ops(9S) structure.
RETURN VALUES	The ddi_no_info() function always returns DDI_FAILURE.
SEE ALSO	getinfo(9E), qassociate(9F), cb_ops(9S), dev_ops(9S)

ddi_peek(9F)

NAME	ddi_peek, ddi_peek8, ddi_peek16, ddi_peek32, ddi_peek64, ddi_peekc, ddi_peeks, ddi_peekl, ddi_peekd – read a value from a location
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). The ddi_peekc(), ddi_peeks(), ddi_peekl(), and ddi_peekd() functions are obsolete. Use, respectively, ddi_peek8(), ddi_peek16(), ddi_peek32(), and ddi_peek64(), instead.
PARAMETERS	<i>dip</i> A pointer to the device's dev_info structure. <i>addr</i> Virtual address of the location to be examined. <i>valuep</i> 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.
DESCRIPTION	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. If the address is not valid, or the value cannot be read without an error occurring, an error code is returned. The routines are most useful when first trying to establish the presence of a device on the system in a driver's probe(9E) or attach(9E) routines.
RETURN VALUES	DDI_SUCCESS The value at the given virtual address was successfully read, and if <i>valuep</i> is non-null, <i>*valuep</i> will have been updated. DDI_FAILURE An error occurred while trying to read the location. <i>*valuep</i> is unchanged.
CONTEXT	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: <pre>if (ddi_peek8(dip, csr, (int8_t *)0) != DDI_SUCCESS) { cmn_err(CE_WARN, "Status register not mapped"); return (DDI_FAILURE); }</pre>

EXAMPLE 1 Checking to see that the status register of a device is mapped into the kernel address space: *(Continued)*

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\n",
            ddi_get_name(dip), ddi_get_instance(dip), id_value);
    ...
    ...

    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

ddi_poke(9F)

NAME	ddi_poke, ddi_poke8, ddi_poke16, ddi_poke32, ddi_poke64, ddi_pokec, ddi_pokes, ddi_pokel, ddi_poked – write a value to a location						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). The ddi_pokec(), ddi_pokes(), ddi_pokel(), and ddi_poked() functions are obsolete. Use, respectively, ddi_poke8(), ddi_poke16(), ddi_poke32(), and ddi_poke64(), instead.						
PARAMETERS	<table><tr><td><i>dip</i></td><td>A pointer to the device's dev_info structure.</td></tr><tr><td><i>addr</i></td><td>Virtual address of the location to be written to.</td></tr><tr><td><i>value</i></td><td>Value to be written to the location.</td></tr></table>	<i>dip</i>	A pointer to the device's dev_info structure.	<i>addr</i>	Virtual address of the location to be written to.	<i>value</i>	Value to be written to the location.
<i>dip</i>	A pointer to the device's dev_info structure.						
<i>addr</i>	Virtual address of the location to be written to.						
<i>value</i>	Value to be written to the location.						
DESCRIPTION	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. If the address is not valid, or the value cannot be written without an error occurring, an error code is returned. These routines are most useful when first trying to establish the presence of a given device on the system in a driver's probe(9E) or attach(9E) routines. On multiprocessing machines these routines can be extremely heavy-weight, so use the ddi_peek(9F) routines instead if possible.						
RETURN VALUES	<table><tr><td>DDI_SUCCESS</td><td>The value was successfully written to the given virtual address.</td></tr><tr><td>DDI_FAILURE</td><td>An error occurred while trying to write to the location.</td></tr></table>	DDI_SUCCESS	The value was successfully written to the given virtual address.	DDI_FAILURE	An error occurred while trying to write to the location.		
DDI_SUCCESS	The value was successfully written to the given virtual address.						
DDI_FAILURE	An error occurred while trying to write to the location.						
CONTEXT	These functions can be called from user or interrupt context.						
SEE ALSO	attach(9E), probe(9E), ddi_peek(9F) <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: <table><tr><th>Previous Name</th><th>New Name</th></tr><tr><td>ddi_pokec</td><td>ddi_poke8</td></tr></table>	Previous Name	New Name	ddi_pokec	ddi_poke8		
Previous Name	New Name						
ddi_pokec	ddi_poke8						

ddi_poke(9F)

ddi_pokes	ddi_poke16
ddi_poke1	ddi_poke32
ddi_poked	ddi_poke64

ddi_prop_create(9F)

NAME	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
SYNOPSIS	<pre> #include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). The ddi_prop_create() and ddi_prop_modify() functions are obsolete. Use ddi_prop_update(9F) instead of these functions.
PARAMETERS	<pre> ddi_prop_create() dev dev_t of the device. dip dev_info_t pointer of the device. flags flag modifiers. The only possible flag value is DDI_PROP_CANSLEEP: Memory allocation may sleep. name name of property. valuep pointer to property value. length property length. ddi_prop_undefine() dev dev_t of the device. dip dev_info_t pointer of the device. flags flag modifiers. The only possible flag value is DDI_PROP_CANSLEEP: Memory allocation may sleep. name name of property. ddi_prop_modify() dev dev_t of the device. dip dev_info_t pointer of the device. </pre>

ddi_prop_create(9F)

<i>flags</i>	flag modifiers. The only possible flag value is DDI_PROP_CANSLEEP: Memory allocation may sleep.
<i>name</i>	name of property.
<i>valuep</i>	pointer to property value.
<i>length</i>	property length.
ddi_prop_remove()	
<i>dev</i>	dev_t of the device.
<i>dip</i>	dev_info_t pointer of the device.
<i>name</i>	name of property.
ddi_prop_remove_all()	
<i>dip</i>	dev_info_t pointer of the device.

DESCRIPTION

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 ddi_getprop(9F) to query whether or not a specific property exists.

Property creation is done by creating a new property definition in the driver's property list associated with *dip*.

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 returned to the caller.

The individual functions are described as follows:

ddi_prop_create()

ddi_prop_create() adds a property to the device's property list. If the property is not associated with any particular *dev* but is associated with the physical device itself, then the argument *dev* should be the special device DDI_DEV_T_NONE. If you do not have a *dev* for your device (for example during attach(9E) time), you can create one using makedevice(9F) with a major number of DDI_MAJOR_T_UNKNOWN. ddi_prop_create() will then make the correct *dev* for your device.

For boolean properties, you must set *length* to 0. For all other properties, the *length* argument must be set to the number of bytes used by the data structure representing the property being created.

Note that creating a property involves allocating memory for the property list, the property name and the property value. If *flags* does not contain DDI_PROP_CANSLEEP, ddi_prop_create() returns DDI_PROP_NO_MEMORY on memory allocation failure or DDI_PROP_SUCCESS if the allocation succeeded. If DDI_PROP_CANSLEEP was set, the caller may sleep until memory becomes available.

ddi_prop_create(9F)

ddi_prop_undefine()

ddi_prop_undefine() 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. However, **ddi_prop_undefine()** will not terminate a search when the **ddi_prop_get_int64(9F)** or **ddi_prop_lookup_int64_array(9F)** routines are used for lookup of 64-bit property value. See **ddi_prop_op(9F)**.

Note that undefining properties does involve memory allocation, and therefore, is subject to the same memory allocation constraints as **ddi_prop_create()**.

ddi_prop_modify()

ddi_prop_modify() modifies the length and the value of a property. If **ddi_prop_modify()** finds the property in the driver's property list, allocates memory for the property value and returns **DDI_PROP_SUCCESS**. If the property was not found, the function returns **DDI_PROP_NOT_FOUND**.

Note that modifying properties does involve memory allocation, and therefore, is subject to the same memory allocation constraints as **ddi_prop_create()**.

ddi_prop_remove()

ddi_prop_remove() unlinks a property from the device's property list. If **ddi_prop_remove()** finds the property (an exact match of both *name* and *dev*), it unlinks the property, frees its memory, and returns **DDI_PROP_SUCCESS**, otherwise, it returns **DDI_PROP_NOT_FOUND**.

ddi_prop_remove_all()

ddi_prop_remove_all() removes the properties of all the **dev_t**'s associated with the *dip*. It is called before unloading a driver.

RETURN VALUES

ddi_prop_create()	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_undefine()	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> DDI_DEV_T_ANY or if <i>name</i> is NULL or <i>name</i> is the NULL string.
ddi_prop_modify()	DDI_PROP_SUCCESS	On success.
	DDI_PROP_NO_MEMORY	On memory allocation failure.

ddi_prop_create(9F)

ddi_prop_remove()	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.
	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.

CONTEXT 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.

EXAMPLES **EXAMPLE 1** Creating a Property

The following example creates a property called *nblocks* for each partition on a disk.

```
int propval = 8192;

for (minor = 0; minor < 8; minor++) {
    (void) ddi_prop_create(makedevice(DDI_MAJOR_T_UNKNOWN, minor),
        dev, DDI_PROP_CANSLEEP, "nblocks", (caddr_t) &propval,
        sizeof (int));
    ...
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	ddi_prop_create() and ddi_prop_modify() are Obsolete

SEE ALSO driver.conf(4), attributes(5), attach(9E), ddi_getproplen(9F), ddi_prop_op(9F), ddi_prop_update(9F), makedevice(9F)

Writing Device Drivers

ddi_prop_exists(9F)

NAME	ddi_prop_exists – check for the existence of a property
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_prop_exists(dev_t match_dev, dev_info_t *dip, uint_t flags, char *name);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>match_dev</i> Device number associated with property or DDI_DEV_T_ANY. <i>dip</i> Pointer to the device info node of device whose property list should be searched. <i>flags</i> Possible flag values are some combination of: DDI_PROP_DONTPASS Do not pass request to parent device information node if the property is not found. DDI_PROP_NOTPROM Do not look at PROM properties (ignored on platforms that do not support PROM properties). <i>name</i> String containing the name of the property.
DESCRIPTION	ddi_prop_exists() checks for the existence of a property regardless of the property value data type. 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: 1. Search software properties created by the driver. 2. Search the software properties created by the system (or nexus nodes in the device info tree). 3. Search the driver global properties list. 4. If DDI_PROP_NOTPROM is not set, search the PROM properties (if they exist). 5. If DDI_PROP_DONTPASS is not set, pass this request to the parent device information node. 6. Return 0 if not found and 1 if found. 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 DDI_DEV_T_ANY, then ddi_prop_exists() will match the request regardless of the <i>match_dev</i> the property was created with. That is the first property whose name matches <i>name</i> will be returned. If a property was created with <i>match_dev</i> set to DDI_DEV_T_NONE then the only way to look up this property is with a <i>match_dev</i> set to DDI_DEV_T_ANY. PROM properties are always created with <i>match_dev</i> set to DDI_DEV_T_NONE. <i>name</i> must always be set to the name of the property being looked up.

ddi_prop_exists(9F)

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

ddi_prop_get_int(9F)

NAME	ddi_prop_get_int, ddi_prop_get_int64 – lookup integer property
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_prop_get_int(dev_t match_dev, dev_info_t *dip, uint_t flags, char *name, int defvalue); int64_t ddi_prop_get_int64(dev_t match_dev, dev_info_t *dip, uint_t flags, char *name, int64_t defvalue);</pre>
PARAMETERS	<i>match_dev</i> Device number associated with property or DDI_DEV_T_ANY. <i>dip</i> Pointer to the device info node of device whose property list should be searched. <i>flags</i> Possible flag values are some combination of: DDI_PROP_DONTPASS Do not pass request to parent device information node if property not found. DDI_PROP_NOTPROM Do not look at PROM properties (ignored on platforms that do not support PROM properties). <i>name</i> String containing the name of the property. <i>defvalue</i> An integer value that is returned if the property cannot be found.
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	The <code>ddi_prop_get_int()</code> and <code>ddi_prop_get_int64()</code> functions search for an integer property and, if found, returns the value of the 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). The property search order is as follows: 1. Search software properties created by the driver. 2. Search the software properties created by the system (or nexus nodes in the device info tree). 3. Search the driver global properties list. 4. If DDI_PROP_NOTPROM is not set, search the PROM properties (if they exist). 5. If DDI_PROP_DONTPASS is not set, pass this request to the parent device information node. 6. Return <i>defvalue</i>. 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 DDI_DEV_T_ANY, then <code>ddi_prop_get_int()</code> and <code>ddi_prop_get_int64()</code> will match the request

ddi_prop_get_int(9F)

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.

ddi_prop_get_int64() will not search the PROM for 64-bit property values.

RETURN VALUES

ddi_prop_get_int() and ddi_prop_get_int64() return the value of the property. If the property is not found, the argument *defvalue* is returned. If the property is found, but cannot be decoded into an int or an int64, then DDI_PROP_NOT_FOUND is returned.

CONTEXT

ddi_prop_get_int() and ddi_prop_get_int64() 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

ddi_prop_lookup(9F)

NAME	ddi_prop_lookup, ddi_prop_lookup_int_array, ddi_prop_lookup_int64_array, ddi_prop_lookup_string_array, ddi_prop_lookup_string, ddi_prop_lookup_byte_array, ddi_prop_free – look up property information
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_prop_lookup_int_array(dev_t match_dev, dev_info_t *dip, uint_t flags, char *name, int **datap, uint_t *nelementsp); int ddi_prop_lookup_int64_array(dev_t match_dev, dev_info_t *dip, uint_t flags, char *name, int64_t **datap, uint_t *nelementsp); int ddi_prop_lookup_string_array(dev_t match_dev, dev_info_t *dip, uint_t flags, char *name, char ***datap, uint_t *nelementsp); 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 *nelementsp); void ddi_prop_free(void *data);</pre>
PARAMETERS	<i>match_dev</i> Device number associated with property or DDI_DEV_T_ANY. <i>dip</i> Pointer to the device info node of device whose property list should be searched. <i>flags</i> Possible flag values are some combination of: - DDI_PROP_DONTPASS Do not pass request to parent device information node if the property is not found. - DDI_PROP_NOTPROM Do not look at PROM properties (ignored on platforms that do not support PROM properties). <i>name</i> String containing the name of the property. <i>nelementsp</i> 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. <i>datap</i> ddi_prop_lookup_int_array() The address of a pointer to an array of integers which, upon successful return, will point to memory containing the integer array property value.

ddi_prop_lookup(9F)

ddi_prop_lookup_int64_array()

The address of a pointer to an array of 64-bit integers which, upon successful return, will point to memory containing the integer array property value.

ddi_prop_lookup_string_array()

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 *argv* argument to *execve*(2).

ddi_prop_lookup_string()

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.

ddi_prop_lookup_byte_array()

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.

**INTERFACE
LEVEL
DESCRIPTION**

Solaris DDI specific (Solaris DDI).

The property look up routines search for and, if found, return the value of a given property. Properties are searched for based on the *dip*, *name*, *match_dev*, and the type of the data (integer, string, or byte). The property search order is as follows:

1. Search software properties created by the driver.
2. Search the software properties created by the system (or nexus nodes in the device info tree).
3. Search the driver global properties list.
4. If DDI_PROP_NOTPROM is not set, search the PROM properties (if they exist).
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, the property look up routines will match the request regardless of the actual *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.

ddi_prop_lookup(9F)

For the routines `ddi_prop_lookup_int_array()`, `ddi_prop_lookup_int64_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_int64_array()`

This routine searches for and returns an array of 64-bit integer property values. The array is defined to be **nelementsp* number of `int64_t` elements. *datap* should be set to the address of a pointer to an array of `int64_t`'s which, upon successful return, will point to memory containing the integer array value of the property. This routine will not search the PROM for 64-bit property values.

`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.

	<pre>ddi_prop_free()</pre> Frees the resources associated with a property previously allocated using <code>ddi_prop_lookup_int_array()</code>, <code>ddi_prop_lookup_int64_array()</code>, <code>ddi_prop_lookup_string_array()</code>, <code>ddi_prop_lookup_string()</code>, or <code>ddi_prop_lookup_byte_array()</code>.										
RETURN VALUES	The functions <code>ddi_prop_lookup_int_array()</code>, <code>ddi_prop_lookup_int64_array()</code>, <code>ddi_prop_lookup_string_array()</code>, <code>ddi_prop_lookup_string()</code>, and <code>ddi_prop_lookup_byte_array()</code> return the following values: <table> <tr> <td><code>DDI_PROP_SUCCESS</code></td><td>Upon success.</td></tr> <tr> <td><code>DDI_PROP_INVALID_ARG</code></td><td>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 <code>NULL</code> or <i>name</i> is the null string.</td></tr> <tr> <td><code>DDI_PROP_NOT_FOUND</code></td><td>Property not found.</td></tr> <tr> <td><code>DDI_PROP_UNDEFINED</code></td><td>Property explicitly not defined (see <code>ddi_prop_undefine(9F)</code>).</td></tr> <tr> <td><code>DDI_PROP_CANNOT_DECODE</code></td><td>The value of the property cannot be decoded.</td></tr> </table>	<code>DDI_PROP_SUCCESS</code>	Upon success.	<code>DDI_PROP_INVALID_ARG</code>	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 <code>NULL</code> or <i>name</i> is the null string.	<code>DDI_PROP_NOT_FOUND</code>	Property not found.	<code>DDI_PROP_UNDEFINED</code>	Property explicitly not defined (see <code>ddi_prop_undefine(9F)</code>).	<code>DDI_PROP_CANNOT_DECODE</code>	The value of the property cannot be decoded.
<code>DDI_PROP_SUCCESS</code>	Upon success.										
<code>DDI_PROP_INVALID_ARG</code>	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 <code>NULL</code> or <i>name</i> is the null string.										
<code>DDI_PROP_NOT_FOUND</code>	Property not found.										
<code>DDI_PROP_UNDEFINED</code>	Property explicitly not defined (see <code>ddi_prop_undefine(9F)</code>).										
<code>DDI_PROP_CANNOT_DECODE</code>	The value of the property cannot be decoded.										
CONTEXT	These functions can be called from user or kernel context.										
EXAMPLES	EXAMPLE 1 Using <code>ddi_prop_lookup_int_array()</code> The following example demonstrates the use of <code>ddi_prop_lookup_int_array()</code>. <pre>int *options; int noptions; /* * Get the data associated with the integer "options" property * array, along with the number of option integers */ if (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); }</pre>										
SEE ALSO	<code>execve(2)</code> , <code>ddi_prop_exists(9F)</code> , <code>ddi_prop_get_int(9F)</code> , <code>ddi_prop_remove(9F)</code> , <code>ddi_prop_undefine(9F)</code> , <code>ddi_prop_update(9F)</code>										

ddi_prop_lookup(9F)

Writing Device Drivers

NAME	ddi_prop_op, ddi_getprop, ddi_getlongprop, ddi_getlongprop_buf, ddi_getproplen – get property information for leaf device drivers
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). The ddi_getlongprop(), ddi_getlongprop_buf(), ddi_getprop(), and ddi_getproplen() functions are obsolete. Use ddi_prop_lookup(9F) instead of ddi_getlongprop(), ddi_getlongprop_buf(), and ddi_getproplen(). Use ddi_prop_get_int(9F) instead of ddi_getprop().
PARAMETERS	<i>dev</i> Device number associated with property or DDI_DEV_T_ANY as the wildcard device number. <i>dip</i> Pointer to a device info node. <i>prop_op</i> Property operator. <i>flags</i> Possible flag values are some combination of: - DDI_PROP_DONTPASS do not pass request to parent device information node if property not found - DDI_PROP_CANSLEEP the routine may sleep while allocating memory - DDI_PROP_NOTPROM do not look at PROM properties (ignored on architectures that do not support PROM properties) <i>name</i> String containing the name of the property. <i>valuep</i> If <i>prop_op</i> is PROP_LEN_AND_VAL_BUF, this should be a pointer to the users buffer. If <i>prop_op</i> is PROP_LEN_AND_VAL_ALLOC, this should be the address of a pointer.

ddi_prop_op(9F)

DESCRIPTION

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.

`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_op(9F)

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:

DDI_PROP_SUCCESS	Property found and returned.
DDI_PROP_NOT_FOUND	Property not found.
DDI_PROP_UNDEFINED	Property already explicitly undefined.
DDI_PROP_NO_MEMORY	Property found, but unable to allocate memory. <i>lengthp</i> points to the correct property length.
DDI_PROP_BUF_TOO_SMALL	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.

ATTRIBUTES

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

ddi_prop_op(9F)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	ddi_getlongprop(), ddi_getlongprop_buf(), ddi_getprop(), and ddi_getproplen() functions are Obsolete

SEE ALSO attributes(5), ddi_prop_create(9F), ddi_prop_get_int(9F), ddi_prop_lookup(9F), kmem_alloc(9F), kmem_free(9F)

Writing Device Drivers

NAME	ddi_prop_update, ddi_prop_update_int_array, ddi_prop_update_int, ddi_prop_update_string_array, ddi_prop_update_int64, ddi_prop_update_int64_array, ddi_prop_update_string, ddi_prop_update_byte_array – update properties
SYNOPSIS	<pre> #include <sys/ddi.h> #include <sys/sunddi.h> 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_int64_array(dev_t dev, dev_info_t *dip, char *name, int64_t *data, uint_t nelements); int ddi_prop_update_int64(dev_t dev, dev_info_t *dip, char *name, int64_t 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>
PARAMETERS	<i>dev</i> Device number associated with the device. <i>dip</i> Pointer to the device info node of device whose property list should be updated. <i>name</i> String containing the name of the property to be updated. <i>nelements</i> The number of elements contained in the memory pointed at by <i>data</i>. <i>ddi_prop_update_int_array()</i> <i>data</i> A pointer an integer array with which to update the property. <i>ddi_prop_update_int()</i> <i>data</i> An integer value with which to update the property. <i>ddi_prop_update_int64_array()</i> <i>data</i> An pointer to a 64-bit integer array with which to update the property. <i>ddi_prop_update_int64()</i> <i>data</i> A 64-bit integer value with which to update the property. <i>ddi_prop_update_string_array()</i>

ddi_prop_update(9F)

INTERFACE LEVEL DESCRIPTION

data 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()`

data A pointer to a string value with which to update the property.

`ddi_prop_update_byte_array()`

data A pointer to a byte array with which to update the property.

Solaris DDI specific (Solaris DDI).

The property update routines search for and, if found, modify the value of a given property. Properties are searched for based on the *dip*, *name*, *dev*, 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).

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 `ddi_prop_update_int_array()` or `ddi_prop_update_int()`. For a 64-bit integer, you must use `ddi_prop_update_int64_array()` or `ddi_prop_update_int64()`. Attempts to update a property with a function that does not correspond to the property data type that was used to create it results in an undefined state.

Usually, the *dev* argument should be set to the actual device number that this property is associated with. If the property is not associated with any particular *dev*, then the argument *dev* should be set to `DDI_DEV_T_NONE`. This property will then match a look up request (see `ddi_prop_lookup(9F)`) with the *match_dev* argument set to `DDI_DEV_T_ANY`. If no *dev* is available for the device (for example during `attach(9E)` time), one can be created using `makedevice(9F)` with a major number of `DDI_MAJOR_T_UNKNOWN`. The update routines will then generate the correct *dev* when creating or updating the property.

name must always be set to the name of the property being updated.

For the routines `ddi_prop_update_int_array()`, `ddi_prop_lookup_int64_array()`, `ddi_prop_update_string_array()`, `ddi_prop_update_string()`, and `ddi_prop_update_byte_array()`, *data* is a pointer which points to memory containing the value of the property. In each case

ddi_prop_update(9F)

**data* points to a different type of property value. See the individual descriptions of the routines below for details concerning the different values. *nelements* is an unsigned integer which contains the number of integer, string, or byte elements accounted for in the memory pointed at by **data*.

For the routines `ddi_prop_update_int()` and `ddi_prop_update_int64()`, *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 *nelements* 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_int64_array()`

Updates or creates an array of 64-bit integer property values. An array of integers is defined to be *nelements* of `int64_t` integer elements. *data* must be a pointer to a 64-bit integer array with which to update the property.

`ddi_prop_update_int64()`

Updates or creates a single 64-bit integer value of a property. *data* must be an `int64_t` 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 NULLterminated 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_update(9F)

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_array()`.

```
int    options[4];

/*
 * Create the "options" integer array with
 * our default values for these parameters
 */
options[0] = XX_OPTIONS0;
options[1] = XX_OPTIONS1;
options[2] = XX_OPTIONS2;
options[3] = XX_OPTIONS3;
i = ddi_prop_update_int_array(xx_dev, xx_dip, "options",
    &options, sizeof (options) / sizeof (int));
```

SEE ALSO `execve(2)`, `attach(9E)`, `ddi_prop_lookup(9F)`, `ddi_prop_remove(9F)`, `makedevice(9F)`

Writing Device Drivers

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>handle</i> The data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>value</i> The data to be written to the device. <i>dev_addr</i> Base device address.
DESCRIPTION	These routines generate a write of various sizes to the mapped memory or device register. The ddi_put8(), ddi_put16(), ddi_put32(), and ddi_put64() functions write 8 bits, 16 bits, 32 bits and 64 bits of data, respectively, to the device address, <i>dev_addr</i>. 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. For certain bus types, you can call these DDI functions from a high-interrupt context. These types include ISA, EISA, and SBus buses. See sysbus(4), isa(4), eisa(4), and sbus(4) for details. For the PCI bus, you can, under certain conditions, call these DDI functions from a high-interrupt context. See pci(4).
CONTEXT	These functions can be called from user, kernel, or interrupt context.
SEE ALSO	ddi_get8(9F), ddi_regs_map_free(9F), ddi_regs_map_setup(9F), ddi_rep_get8(9F), ddi_rep_put8(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:

ddi_put8(9F)

Previous Name	New Name
ddi_putb	ddi_put8
ddi_putw	ddi_put16
ddi_putl	ddi_put32
ddi_putll	ddi_put64

ddi_regs_map_free(9F)

NAME	ddi_regs_map_free – free a previously mapped register address space				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_regs_map_free(ddi_acc_handle_t *handle) ;</pre>				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).				
PARAMETERS	<table><tr><td><i>handle</i></td><td>Pointer to a data access handle previously allocated by a call to a setup routine such as ddi_regs_map_setup(9F).</td></tr></table>	<i>handle</i>	Pointer to a data access handle previously allocated by a call to a setup routine such as ddi_regs_map_setup(9F).		
<i>handle</i>	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 <i>handle</i> . 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: <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Architecture</td><td>PCI Local Bus, SBus, ISA, EISA</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	PCI Local Bus, SBus, ISA, EISA
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Architecture	PCI Local Bus, SBus, ISA, EISA				
SEE ALSO	attributes(5), ddi_regs_map_setup(9F) <i>Writing Device Drivers</i>				

ddi_regs_map_setup(9F)

NAME	ddi_regs_map_setup – set up a mapping for a register address space				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>				
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>dip</i> Pointer to the device's dev_info structure. <i>rnumber</i> Index number to the register address space set. <i>addrp</i> A platform-dependent value that, when added to an offset that is less than or equal to the <i>len</i> parameter (see below), is used for the dev_addr argument to the ddi_get, ddi_mem_get, and ddi_io_get/put routines. <i>offset</i> Offset into the register address space. <i>len</i> Length to be mapped. <i>accattrp</i> Pointer to a device access attribute structure of this mapping (see ddi_device_acc_attr(9S)). <i>handlep</i> Pointer to a data access handle.				
DESCRIPTION	ddi_regs_map_setup() maps in the register set given by <i>rnumber</i>. The register number determines which register set is mapped if more than one exists. <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>. The device access attributes are specified in the location pointed by the <i>accattrp</i> argument (see ddi_device_acc_attr(9S) for details). 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.				
RETURN VALUES	ddi_regs_map_setup() returns: <table> <tr> <td>DDI_SUCCESS</td><td>Successfully set up the mapping for data access.</td></tr> <tr> <td>DDI_FAILURE</td><td>Invalid register number <i>rnumber</i>, offset <i>offset</i>, or length <i>len</i>.</td></tr> </table>	DDI_SUCCESS	Successfully set up the mapping for data access.	DDI_FAILURE	Invalid register number <i>rnumber</i> , offset <i>offset</i> , or length <i>len</i> .
DDI_SUCCESS	Successfully set up the mapping for data access.				
DDI_FAILURE	Invalid register number <i>rnumber</i> , offset <i>offset</i> , or length <i>len</i> .				

ddi_regs_map_setup(9F)

DDI_ME_RNUMBER_RANGE

Invalid register number *rnumber* or unable to find *reg* property.

DDI_REGS_ACC_CONFLICT

Cannot enable the register mapping due to access conflicts with other enabled mappings.

Note that the return value DDI_ME_RNUMBER_RANGE is not supported on all platforms. Also, there is potential overlap between DDI_ME_RNUMBER_RANGE and DDI_FAILURE. Drivers should check for !=DDI_SUCCESS rather than checking for a specific failure value.

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

ddi_remove_minor_node(9F)

NAME	ddi_remove_minor_node – remove a minor node for this dev_info				
SYNOPSIS	<pre>void ddi_remove_minor_node(dev_info_t *dip, char *name);</pre>				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).				
PARAMETERS	<table><tr><td><i>dip</i></td><td>A pointer to the device's dev_info structure.</td></tr><tr><td><i>name</i></td><td>The name of this minor device. If <i>name</i> is NULL, then remove all minor data structures from this dev_info.</td></tr></table>	<i>dip</i>	A pointer to the device's dev_info structure.	<i>name</i>	The name of this minor device. If <i>name</i> is NULL, then remove all minor data structures from this dev_info.
<i>dip</i>	A pointer to the device's dev_info structure.				
<i>name</i>	The name of this minor device. If <i>name</i> is NULL, then remove all minor data structures from this dev_info.				
DESCRIPTION	ddi_remove_minor_node() removes a data structure from the linked list of minor data structures that is pointed to by the dev_info structure for this driver.				
EXAMPLES	EXAMPLE 1 Removing a minor node This will remove a data structure describing a minor device called dev1 which is linked into the dev_info structure pointed to by dip: <pre>ddi_remove_minor_node(dip, "dev1");</pre>				
SEE ALSO	attach(9E), detach(9E), ddi_create_minor_node(9F) <i>Writing Device Drivers</i>				

NAME	ddi_removing_power – check whether DDI_SUSPEND might result in power being removed from a device
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_removing_power(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	The ddi_removing_power() function indicates whether a currently pending call into a driver's detach(9E) entry point with a command of DDI_SUSPEND is likely to result in power being removed from the device. ddi_removing_power() can return true and power still not be removed from the device due to a failure to suspend and power off the system.
PARAMETERS	The ddi_removing_power() function supports the following parameter: <i>dip</i> pointer to the device's dev_info structure
RETURN VALUES	The ddi_removing_power() function returns: 1 Power might be removed by the framework as a result of the pending DDI_SUSPEND call. 0 Power will not be removed by the framework as a result of the pending DDI_SUSPEND call.
EXAMPLES	EXAMPLE 1 Protecting a Tape from Abrupt Power Removal A tape driver that has hardware that would damage the tape if power is removed might include this code in its detach(9E) code: <pre>int xxdetach(dev_info_t *dip, ddi_detach_cmd_t cmd) { ... case DDI_SUSPEND: /* * We do not allow DDI_SUSPEND if power will be removed and * we have a device that damages tape when power is removed * We do support DDI_SUSPEND for Device Reconfiguration, however. */ if (ddi_removing_power(dip) && xxdamages_tape(dip)) return (DDI_FAILURE); ... }</pre>

ddi_removing_power(9F)

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO `attributes(5)`, `cpr(7)`, `attach(9E)`, `detach(9E)`

Writing Device Drivers

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void ddi_rep_get8(ddi_acc_handle_t handle, uint8_t *host_addr, uint8_t *dev_addr, size_t repcount, uint_t flags); void ddi_rep_get16(ddi_acc_handle_t handle, uint16_t *host_addr, uint16_t *dev_addr, size_t repcount, uint_t flags); void ddi_rep_get32(ddi_acc_handle_t handle, uint32_t *host_addr, uint32_t *dev_addr, size_t repcount, uint_t flags); void ddi_rep_get64(ddi_acc_handle_t handle, uint64_t *host_addr, uint64_t *dev_addr, size_t repcount, uint_t flags);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>handle</i> The data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>host_addr</i> Base host address. <i>dev_addr</i> Base device address. <i>repcount</i> Number of data accesses to perform. <i>flags</i> Device address flags: DDI_DEV_AUTOINCR Automatically increment the device address, <i>dev_addr</i>, during data accesses. DDI_DEV_NO_AUTOINCR Do not advance the device address, <i>dev_addr</i>, during data accesses.
DESCRIPTION	These routines generate multiple reads from the mapped memory or device register. <i>repcount</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 ddi_rep_get8(), ddi_rep_get16(), ddi_rep_get32(), and ddi_rep_get64() 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.

ddi_rep_get8(9F)

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>

NAME	ddi_report_dev – announce a device
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> void ddi_report_dev(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dip</i> a pointer the device's dev_info structure.
DESCRIPTION	ddi_report_dev() 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 dmesg(1M)), but is only displayed on the console if the system was booted with the verbose (-v) argument.
CONTEXT	ddi_report_dev() can be called from user context.
SEE ALSO	dmesg(1M), kernel(1M) <i>Writing Device Drivers</i>

ddi_rep_put8(9F)

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>handle</i> The data access handle returned from setup calls, such as ddi_regs_map_setup(9F). <i>host_addr</i> Base host address. <i>dev_addr</i> Base device address. <i>repcount</i> Number of data accesses to perform. <i>flags</i> Device address flags: DDI_DEV_AUTOINCR Automatically increment the device address, <i>dev_addr</i>, during data accesses. DDI_DEV_NO_AUTOINCR Do not advance the device address, <i>dev_addr</i>, during data accesses.
DESCRIPTION	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 ddi_rep_put8(), ddi_rep_put16(), ddi_rep_put32(), and ddi_rep_put64() functions write 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively, to 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.

ddi_rep_put8(9F)

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
ddi_rep_putb	ddi_rep_put8
ddi_rep_putw	ddi_rep_put16
ddi_rep_putl	ddi_rep_put32
ddi_rep_putll	ddi_rep_put64

ddi_root_node(9F)

NAME	ddi_root_node – get the root of the dev_info tree
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> dev_info_t *ddi_root_node(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	ddi_root_node() returns a pointer to the root node of the device information tree.
RETURN VALUES	ddi_root_node() returns a pointer to a device information structure.
CONTEXT	ddi_root_node() can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

NAME	ddi_segmap, ddi_segmap_setup – set up a user mapping using seg_dev
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	These interfaces are obsolete. See devmap(9E) for an alternative to ddi_segmap(). Use devmap_setup(9F) instead of ddi_segmap_setup().
PARAMETERS	<i>dev</i> The device whose memory is to be mapped. <i>offset</i> The offset within device memory at which the mapping begins. <i>asp</i> An opaque pointer to the user address space into which the device memory should be mapped. <i>addrp</i> Pointer to the starting address within the user address space to which the device memory should be mapped. <i>len</i> Length (in bytes) of the memory to be mapped. <i>prot</i> A bit field that specifies the protections. Some combinations of possible settings are: 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. <i>maxprot</i> 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. <i>flags</i> Flags indicating type of mapping. Possible values are (other bits may be set):

ddi_segmap(9F)

	MAP_PRIVATE Changes are private. MAP_SHARED Changes should be shared. MAP_FIXED The user specified an address in <i>*addrp</i> rather than letting the system pick an address. <i>credp</i> Pointer to user credential structure.
ddi_segmap_setup()	<i>dev_acc_attr</i> Pointer to a <code>ddi_device_acc_attr(9S)</code> structure which contains the device access attributes to apply to this mapping. <i>rnumber</i> 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 <code>ddi_devmap_segmap(9F)</code> instead. See <code>ddi_devmap_segmap(9F)</code> for details. <code>ddi_segmap()</code> and <code>ddi_segmap_setup()</code> set up user mappings to device space. When setting up the mapping, the <code>ddi_segmap()</code> and <code>ddi_segmap_setup()</code> routines call the <code>mmap(9E)</code> entry point to validate the range to be mapped. When a user process accesses the mapping, the driver's <code>mmap(9E)</code> 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. <code>ddi_segmap()</code> is typically used as the <code>segmap(9E)</code> entry in the <code>cb_ops(9S)</code> structure for those devices that do not choose to provide their own <code>segmap(9E)</code> entry point. However, some drivers may have their own <code>segmap(9E)</code> entry point to do some initial processing on the parameters and then call <code>ddi_segmap()</code> to establish the default memory mapping. <code>ddi_segmap_setup()</code> is used in the driver's <code>segmap(9E)</code> entry point to set up the mapping and assign device access attributes to that mapping. <i>rnumber</i> specifies the register set representing the range of device memory being mapped. See <code>ddi_device_acc_attr(9S)</code> for details regarding what device access attributes are available. <code>ddi_segmap_setup()</code> cannot be used directly in the <code>cb_ops(9S)</code> structure and requires a driver to have a <code>segmap(9E)</code> entry point.
RETURN VALUES	<code>ddi_segmap()</code> and <code>ddi_segmap_setup()</code> return the following values: 0 Successful completion. Non-zero An error occurred. In particular, they return <code>ENXIO</code> if the range to be mapped is invalid.
CONTEXT	<code>ddi_segmap()</code> and <code>ddi_segmap_setup()</code> can be called from user or kernel context only.

ddi_segmap(9F)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO mmap(2), attributes(5), devmap(9E), mmap(9E), segmap(9E), devmap_setup(9F), 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.

ddi_slaveonly(9F)

NAME	ddi_slaveonly – tell if a device is installed in a slave access only location
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int ddi_slaveonly(dev_info_t *dip);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dip</i> A pointer to the device's dev_info structure.
DESCRIPTION	ddi_slaveonly() 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.
RETURN VALUES	DDI_SUCCESS The device has been installed in a slave access only location. DDI_FAILURE The device has not been installed in a slave access only location.
CONTEXT	ddi_slaveonly() can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

NAME	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										
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>										
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>state_p</i></td><td>Address of the opaque state pointer which will be initialized by ddi_soft_state_init() to point to implementation dependent data.</td></tr> <tr> <td><i>size</i></td><td>Size of the item which will be allocated by subsequent calls to ddi_soft_state_zalloc().</td></tr> <tr> <td><i>n_items</i></td><td>A hint of the number of items which will be preallocated; zero is allowed.</td></tr> <tr> <td><i>state</i></td><td>An opaque pointer to implementation-dependent data that describes the soft state.</td></tr> <tr> <td><i>item</i></td><td>The item number for the state structure; usually the instance number of the associated devinfo node.</td></tr> </table>	<i>state_p</i>	Address of the opaque state pointer which will be initialized by ddi_soft_state_init() to point to implementation dependent data.	<i>size</i>	Size of the item which will be allocated by subsequent calls to ddi_soft_state_zalloc().	<i>n_items</i>	A hint of the number of items which will be preallocated; zero is allowed.	<i>state</i>	An opaque pointer to implementation-dependent data that describes the soft state.	<i>item</i>	The item number for the state structure; usually the instance number of the associated devinfo node.
<i>state_p</i>	Address of the opaque state pointer which will be initialized by ddi_soft_state_init() to point to implementation dependent data.										
<i>size</i>	Size of the item which will be allocated by subsequent calls to ddi_soft_state_zalloc().										
<i>n_items</i>	A hint of the number of items which will be preallocated; zero is allowed.										
<i>state</i>	An opaque pointer to implementation-dependent data that describes the soft state.										
<i>item</i>	The item number for the state structure; usually the instance number of the associated devinfo node.										
DESCRIPTION	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. 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. 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. The routine ddi_soft_state_init() is usually called in the driver's _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 driver's 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										

ddi_soft_state(9F)

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. Returned memory is zeroed.

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 driver's `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 driver's `_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

`ddi_get_soft_state()`

`NULL` The requested state structure was not allocated at the time of the call.

pointer The pointer to the state structure.

`ddi_soft_state_init()`

`0` The allocation was successful.

`EINVAL` Either the size parameter was zero, or the *state_p* parameter was invalid.

`ddi_soft_state_zalloc()`

`DDI_SUCCESS` The allocation was successful.

`DDI_FAILURE` The routine failed to allocate the storage required; either the *state* 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 **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)
        return (error);
    if ((error = mod_install(&modlinkage)) != 0)
        ddi_soft_state_fini(&statep);
    return (error);
}

int
_fini(void)
{
    int error;

    if ((error = mod_remove(&modlinkage)) != 0)
        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:
        instance = ddi_get_instance(dip);
        if (ddi_soft_state_zalloc(statep, instance) != DDI_SUCCESS)
            return (DDI_FAILURE);
        softc = ddi_get_soft_state(statep, instance);
        softc->dip = dip;
        ...
        return (DDI_SUCCESS);
    default:
        return (DDI_FAILURE);
    }
}
```

ddi_soft_state(9F)

EXAMPLE 1 Creating and Removing Data Structures (Continued)

```
static int
xxdetach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    int instance;

    switch (cmd) {

    case DDI_DETACH:
        instance = ddi_get_instance(dip);
        ...
        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)
        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.

ddi_soft_state(9F)

```
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.

ddi_umem_alloc(9F)

NAME	ddi_umem_alloc, ddi_umem_free – allocate and free page-aligned kernel memory	
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/sunddi.h> void *ddi_umem_alloc(size_t <i>size</i>, int <i>flag</i>, ddi_umem_cookie_t *<i>cookiep</i>) ; void ddi_umem_free(ddi_umem_cookie_t <i>cookie</i>) ;</pre>	
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).	
PARAMETERS		
ddi_umem_alloc()	<i>size</i>	Number of bytes to allocate.
	<i>flag</i>	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.
	<i>cookiep</i>	Pointer to a kernel memory cookie.
ddi_umem_free()	<i>cookie</i>	A kernel memory cookie allocated in ddi_umem_alloc().
DESCRIPTION	ddi_umem_alloc() 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 devmap(9E) and devmap_umem_setup(9F) for further information. <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 ddi_umem_alloc() is zero-filled. *<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 devmap_umem_setup(9F) when the drivers want to export the kernel memory to a user application. To free the allocated memory, a driver calls ddi_umem_free() with the cookie obtained from ddi_umem_alloc(). ddi_umem_free() releases the entire buffer.	
RETURN VALUES	Non-null	Successful completion. ddi_umem_alloc() returns a pointer to the allocated memory.

ddi_umem_alloc(9F)

	NULL	Memory cannot be allocated by <code>ddi_umem_alloc()</code> because <code>DDI_UMEM_NOSLEEP</code> is set and the system is out of resources.
CONTEXT		<code>ddi_umem_alloc()</code> can be called from any context if <i>flag</i> is set to <code>DDI_UMEM_NOSLEEP</code> . If <code>DDI_UMEM_SLEEP</code> is set, <code>ddi_umem_alloc()</code> can be called from user and kernel context only. <code>ddi_umem_free()</code> can be called from any context.
SEE ALSO		<code>devmap(9E)</code> , <code>condvar(9F)</code> , <code>devmap_umem_setup(9F)</code> , <code>kmem_alloc(9F)</code> , <code>mutex(9F)</code> , <code>rwlock(9F)</code> , <code>semaphore(9F)</code> <i>Writing Device Drivers</i>
WARNINGS		Setting the <code>DDI_UMEM_PAGEABLE</code> flag in <code>ddi_umem_alloc()</code> 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 must not be used for the following purposes: ■ For synchronization objects such as locks and condition variables. See <code>mutex(9F)</code>, <code>semaphore(9F)</code>, <code>rwlock(9F)</code>, and <code>condvar(9F)</code>. ■ For driver interrupt routines. Memory allocated using <code>ddi_umem_alloc()</code> without setting <code>DDI_UMEM_PAGEABLE</code> 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. Excessive use of kernel memory is likely to effect overall system performance. Over-commitment of kernel memory may cause unpredictable consequences. 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. Do not call <code>ddi_umem_alloc()</code> within <code>DDI_SUSPEND</code> and <code>DDI_RESUME</code> operations. Memory acquired at these times is not reliable. In some cases, such a call can cause a system to hang.
NOTES		<code>ddi_umem_alloc(0, flag, cookiep)</code> always returns <code>NULL</code> . <code>ddi_umem_free(NULL)</code> has no effects on system.

ddi_umem_iosetup(9F)

NAME	ddi_umem_iosetup – Setup I/O requests to application memory																
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> struct buf *ddi_umem_iosetup(ddi_umem_cookie_t <i>cookie</i>, off_t <i>off</i>, size_t <i>len</i>, int <i>direction</i>, dev_t <i>dev</i>, daddr_t <i>blkno</i>, int (*iodone) (struct buf *), int <i>sleepflag</i>);</pre>																
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <table> <tr> <td><i>cookie</i></td><td>The kernel memory cookie allocated by ddi_umem_lock(9F).</td></tr> <tr> <td><i>off</i></td><td>Offset from the start of the cookie.</td></tr> <tr> <td><i>len</i></td><td>Length of the I/O request in bytes.</td></tr> <tr> <td><i>direction</i></td><td>Must be set to B_READ for reads from the device or B_WRITE for writes to the device.</td></tr> <tr> <td><i>dev</i></td><td>Device number</td></tr> <tr> <td><i>blkno</i></td><td>Block number on device.</td></tr> <tr> <td><i>iodone</i></td><td>Specific biodone(9F) routine.</td></tr> <tr> <td><i>sleepflag</i></td><td>Determines whether caller can sleep for memory. Possible flags are DDI_UMEM_SLEEP to allow sleeping until memory is available, or DDI_UMEM_NOSLEEP to return NULL immediately if memory is not available.</td></tr> </table>	<i>cookie</i>	The kernel memory cookie allocated by ddi_umem_lock(9F).	<i>off</i>	Offset from the start of the cookie.	<i>len</i>	Length of the I/O request in bytes.	<i>direction</i>	Must be set to B_READ for reads from the device or B_WRITE for writes to the device.	<i>dev</i>	Device number	<i>blkno</i>	Block number on device.	<i>iodone</i>	Specific biodone(9F) routine.	<i>sleepflag</i>	Determines whether caller can sleep for memory. Possible flags are DDI_UMEM_SLEEP to allow sleeping until memory is available, or DDI_UMEM_NOSLEEP to return NULL immediately if memory is not available.
<i>cookie</i>	The kernel memory cookie allocated by ddi_umem_lock(9F).																
<i>off</i>	Offset from the start of the cookie.																
<i>len</i>	Length of the I/O request in bytes.																
<i>direction</i>	Must be set to B_READ for reads from the device or B_WRITE for writes to the device.																
<i>dev</i>	Device number																
<i>blkno</i>	Block number on device.																
<i>iodone</i>	Specific biodone(9F) routine.																
<i>sleepflag</i>	Determines whether caller can sleep for memory. Possible flags are DDI_UMEM_SLEEP to allow sleeping until memory is available, or DDI_UMEM_NOSLEEP to return NULL immediately if memory is not available.																
DESCRIPTION	The ddi_umem_iosetup(9F) function is used by drivers to setup I/O requests to application memory which has been locked down using ddi_umem_lock(9F). The ddi_umem_iosetup(9F) function returns a pointer to a buf(9S) structure corresponding to the memory cookie <i>cookie</i>. Drivers can setup multiple buffer structures simultaneously active using the same memory cookie. The buf(9S) structures can span all or part of the region represented by the cookie and can overlap each other. The buf(9S) structure can be passed to ddi_dma_buf_bind_handle(9F) to initiate DMA transfers to or from the locked down memory. The <i>off</i> parameter specifies the offset from the start of the cookie. The <i>len</i> parameter represents the length of region to be mapped by the buffer. The <i>direction</i> parameter must be set to either B_READ or B_WRITE, to indicate the action that will be performed by the device. (Note that this direction is in the opposite sense of the VM system's direction of DDI_UMEMLOCK_READ and DDI_UMEMLOCK_WRITE.) The direction must be compatible with the flags used to create the memory cookie in ddi_umem_lock(9F). For example, if ddi_umem_lock() is called with the <i>flags</i> parameter set to DDI_UMEMLOCK_READ, the <i>direction</i> parameter in ddi_umem_iosetup() should be set to B_WRITE.																

ddi_umem_iosetup(9F)

The *dev* parameter specifies the device to which the buffer is to perform I/O. The *blkno* parameter represents the block number on the device. It will be assigned to the *b_blkno* field of the returned buffer structure. The *iodone* parameter enables the driver to identify a specific *biiodone(9F)* routine to be called by the driver when the I/O is complete. The *sleepflag* parameter determines if the caller can sleep for memory. *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 not available.

After the I/O has completed and the buffer structure is no longer needed, the driver calls *freerbuf(9F)* to free the buffer structure.

RETURN VALUES	The <i>ddi_umem_iosetup(9F)</i> function returns a pointer to the initialized buffer header, or <i>NULL</i> if no space is available.
CONTEXT	The <i>ddi_umem_iosetup(9F)</i> function can be called from any context only if flag is set to <i>DDI_UMEM_NOSLEEP</i> . If <i>DDI_UMEM_SLEEP</i> is set, <i>ddi_umem_iosetup(9F)</i> can be called from user and kernel context only.
SEE ALSO	<i>ddi_umem_lock(9F)</i> , <i>ddi_dma_buf_bind_handle(9F)</i> , <i>ddi_umem_unlock(9F)</i> , <i>freerbuf(9F)</i> , <i>physio(9F)</i> , <i>buf(9S)</i>

ddi_umem_lock(9F)

NAME	ddi_umem_lock, ddi_umem_unlock – lock and unlock memory pages
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int ddi_umem_lock(caddr_t <i>addr</i>, size_t <i>len</i>, int <i>flags</i>, ddi_umem_cookie_t *<i>cookiep</i>); void ddi_umem_unlock(ddi_umem_cookie_t <i>cookie</i>);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
PARAMETERS	
ddi_umem_lock	<i>addr</i> Virtual address of memory object <i>len</i> Length of memory object in bytes <i>flags</i> Valid flags include: DDI_UMEMLOCK_READ Memory pages are locked to be read from. (Disk write or a network send.) DDI_UMEMLOCK_WRITE Memory pages are locked to be written to. (Disk read or a network receive.) <i>cookiep</i> Pointer to a kernel memory cookie.
ddi_umem_unlock	<i>cookie</i> Kernel memory cookie allocated by ddi_umem_lock().
DESCRIPTION	The ddi_umem_lock(9F) function locks down the physical pages (including I/O pages) that correspond to the current process' virtual address range [<i>addr</i>, <i>addr</i> + <i>size</i>) and fills in a cookie representing the locked pages. This cookie can be used to create a buf(9S) structure that can be used to perform I/O (see ddi_umem_iosetup(9F) and ddi_dma_buf_bind_handle(9F)), or it can be used with devmap_umem_setup(9F) to export the memory to an application. The virtual address and length specified must be at a page boundary and the mapping performed in terms of the system page size. See pagesize(1). The flags argument indicates the intended use of the locked memory. Set flags to DDI_UMEMLOCK_READ if the memory pages will be read (for example, in a disk write or a network send.) Set flags to DDI_UMEMLOCK_WRITE if the memory pages will be written (for example, in a disk read or a network receive). You must choose one (and only one) of these values. To unlock the locked pages, the drivers call ddi_umem_unlock(9F) with the cookie obtained from ddi_umem_lock(9F).

ddi_umem_lock(9F)

The process is not allowed to `exec(2)` or `fork(2)` while its physical pages are locked down by the device driver.

The device driver must ensure that the physical pages have been unlocked after the application has called `close(2)`.

RETURN VALUES On success, a 0 is returned. Otherwise, one of the following `errno` values is returned.

`EFAULT`

User process has no mapping at that address range or does not support locking

`EACCES`

User process does not have the required permission.

`ENOMEM`

The system does not have sufficient resources to lock memory.

`EAGAIN`

Could not allocate system resources required to lock the pages. The `ddi_umem_lock()` could succeed at a later time.

`EINVAL`

Requested memory is not aligned on a system page boundary.

CONTEXT The `ddi_umem_lock()` function can only be called from user context; `ddi_umem_unlock()` from user, kernel, and interrupt contexts.

SEE ALSO `ddi_umem_iosetup(9F)`, `ddi_dma_buf_bind_handle(9F)`, `devmap_umem_setup(9F)`, `ddi_umem_alloc(9F)`

NOTES The `ddi_umem_lock(9F)` function consumes physical memory. The driver is responsible for a speedy unlock to free up the resources.

`ddi_umem_unlock()` can defer unlocking of the pages to a later time depending on the implementation.

delay(9F)

NAME	delay – delay execution for a specified number of clock ticks
SYNOPSIS	<pre>#include <sys/ddi.h> void delay(clock_t ticks);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>ticks</i> The number of clock cycles to delay.
DESCRIPTION	<code>delay()</code> 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 <code>drv_usec2ohz(9F)</code> to convert microseconds into clock ticks. <code>delay()</code> uses <code>timeout(9F)</code> to schedule an internal function to be called after the specified amount of time has elapsed. <code>delay()</code> then waits until the function is called. Because <code>timeout()</code> is subject to priority inversion, drivers waiting on behalf of processes with real-time constraints should use <code>cv_timedwait(9F)</code> rather than <code>delay()</code>. <code>delay()</code> does not busy-wait. If busy-waiting is required, use <code>drv_usecwait(9F)</code>.
CONTEXT	<code>delay()</code> can be called from user and kernel contexts.
EXAMPLES	EXAMPLE 1 <code>delay()</code> Example 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. <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->status & 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) & 0xf)); 17 /* wait one minute and try again */ 18 delay(60 * drv_usec2ohz(1000000)); 19 }</pre>

delay(9F)

EXAMPLE 1 delay() Example *(Continued)*

SEE ALSO biodone(9F), biowait(9F), cv_timedwait(9F), ddi_in_panic(9F),
drv_hztousec(9F), drv_usectohz(9F), drv_usecwait(9F), timeout(9F),
untimeout(9F)

Writing Device Drivers

devmap_default_access(9F)

NAME	devmap_default_access – default driver memory access function
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int devmap_default_access(devmap_cookie_t dhp, void *pvtp, offset_t off, size_t len, uint_t type, uint_t rw);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dhp</i> An opaque mapping handle that the system uses to describe the mapping. <i>pvtp</i> Driver private mapping data. <i>off</i> User offset within the logical device memory at which the access begins. <i>len</i> Length (in bytes) of the memory being accessed. <i>type</i> Type of access operation. <i>rw</i> Type of access.
DESCRIPTION	devmap_default_access() is a function providing the semantics of devmap_access(9E). The drivers call devmap_default_access() to handle the mappings that do not support context switching. The drivers should call devmap_do_ctxmgt(9F) for the mappings that support context management. devmap_default_access() can either be called from devmap_access(9E) or be used as the devmap_access(9E) 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 devmap_access(9E) entry point and must not be modified.
RETURN VALUES	0 Successful completion. Non-zero An error occurred.
CONTEXT	devmap_default_access() must be called from the driver's devmap_access(9E) 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. <pre>. . . #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,</pre>

EXAMPLE 1 Using devmap_default_access in devmap_access. (Continued)

```

    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_DOCTXMG) >= 0 && diff < CTXMGT_SIZE) {
        /*
         * calculates the length for context switching
         */
        if ((len + off) > (OFF_DOCTXMG + 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

devmap_devmem_setup(9F)

NAME	devmap_devmem_setup, devmap_umem_setup – set driver memory mapping parameters																												
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>																												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).																												
PARAMETERS	devmap_devmem_setup() parameters: <table> <tr> <td><i>dhp</i></td><td>An opaque mapping handle that the system uses to describe the mapping.</td></tr> <tr> <td><i>dip</i></td><td>Pointer to the device's dev_info structure.</td></tr> <tr> <td><i>callbackops</i></td><td>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.</td></tr> <tr> <td><i>rnumber</i></td><td>Index number to the register address space set.</td></tr> <tr> <td><i>roff</i></td><td>Offset into the register address space.</td></tr> <tr> <td><i>len</i></td><td>Length (in bytes) of the mapping to be mapped.</td></tr> <tr> <td><i>maxprot</i></td><td>Maximum protection flag possible for attempted mapping. Some combinations of possible settings are: <table> <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> <tr> <td>PROT_ALL</td><td>All access is allowed.</td></tr> </table> </td></tr> <tr> <td><i>flags</i></td><td>Must be set to 0.</td></tr> <tr> <td><i>accattrp</i></td><td>Pointer to a ddi_device_acc_attr(9S) structure. The structure contains the device access attributes to be applied to this range of memory.</td></tr> </table>	<i>dhp</i>	An opaque mapping handle that the system uses to describe the mapping.	<i>dip</i>	Pointer to the device's dev_info structure.	<i>callbackops</i>	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.	<i>rnumber</i>	Index number to the register address space set.	<i>roff</i>	Offset into the register address space.	<i>len</i>	Length (in bytes) of the mapping to be mapped.	<i>maxprot</i>	Maximum protection flag possible for attempted mapping. Some combinations of possible settings are: <table> <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> <tr> <td>PROT_ALL</td><td>All access is allowed.</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_ALL	All access is allowed.	<i>flags</i>	Must be set to 0.	<i>accattrp</i>	Pointer to a ddi_device_acc_attr(9S) structure. The structure contains the device access attributes to be applied to this range of memory.
<i>dhp</i>	An opaque mapping handle that the system uses to describe the mapping.																												
<i>dip</i>	Pointer to the device's dev_info structure.																												
<i>callbackops</i>	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.																												
<i>rnumber</i>	Index number to the register address space set.																												
<i>roff</i>	Offset into the register address space.																												
<i>len</i>	Length (in bytes) of the mapping to be mapped.																												
<i>maxprot</i>	Maximum protection flag possible for attempted mapping. Some combinations of possible settings are: <table> <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> <tr> <td>PROT_ALL</td><td>All access is allowed.</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_ALL	All access is allowed.																		
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_ALL	All access is allowed.																												
<i>flags</i>	Must be set to 0.																												
<i>accattrp</i>	Pointer to a ddi_device_acc_attr(9S) structure. The structure contains the device access attributes to be applied to this range of memory.																												

devmap_umem_setup() parameters:

<i>dhp</i>	An opaque data structure that the system uses to describe the mapping.										
<i>dip</i>	Pointer to the device's dev_info structure.										
<i>callbackops</i>	Pointer to a devmap_callback_ctl(9S) structure. The structure contains pointers to device driver-supplied functions that manage events on the device mapping.										
<i>cookie</i>	A kernel memory <i>cookie</i> (see ddi_umem_alloc(9F)).										
<i>koff</i>	Offset into the kernel memory defined by <i>cookie</i> .										
<i>len</i>	Length (in bytes) of the mapping to be mapped.										
<i>maxprot</i>	Maximum protection flag possible for attempted mapping. Some combinations of possible settings are: <table> <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> <tr> <td>PROT_ALL</td><td>All access is allowed.</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_ALL	All access is allowed.
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_ALL	All access is allowed.										
<i>flags</i>	Must be set to 0.										
<i>accattrp</i>	Pointer to a ddi_device_acc_attr(9S) structure. Ignored in the current release. Reserved for future use.										

DESCRIPTION

devmap_devmem_setup() and devmap_umem_setup() are used in the devmap(9E) 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 *callbackops* 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).

devmap_devmem_setup(9F)

See `devmap_map(9E)`, `devmap_access(9E)`, `devmap_dup(9E)`, and `devmap_unmap(9E)` for details on these entry points.

By specifying a valid *callbackops* 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 *callbackops* 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	<code>devmap_devmem_setup()</code> and <code>devmap_umem_setup()</code> can be called from user, kernel, and interrupt context.

SEE ALSO	<code>exit(2)</code> , <code>fork(2)</code> , <code>mmap(2)</code> , <code>munmap(2)</code> , <code>devmap(9E)</code> , <code>ddi_umem_alloc(9F)</code> , <code>ddi_device_acc_attr(9S)</code> , <code>devmap_callback_ctl(9S)</code>

Writing Device Drivers

NAME	devmap_do_ctxmgt – perform device context switching on a mapping	
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>	
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI).	
	<i>dhp</i>	An opaque mapping handle that the system uses to describe the mapping.
	<i>pvtp</i>	Driver private mapping data.
	<i>off</i>	User offset within the logical device memory at which the access begins.
	<i>len</i>	Length (in bytes) of the memory being accessed.
	<i>devmap_contextmgt</i>	The address of driver function that the system will call to perform context switching on a mapping. See devmap_contextmgt(9E) for details.
	<i>type</i>	Type of access operation. Provided by devmap_access(9E). Should not be modified.
	<i>rw</i>	Direction of access. Provided by devmap_access(9E). Should not be modified.
DESCRIPTION	Device drivers call devmap_do_ctxmgt () in the devmap_access(9E) entry point to perform device context switching on a mapping. devmap_do_ctxmgt () passes a pointer to a driver supplied callback function, devmap_contextmgt(9E), to the system that will perform the actual device context switching. If devmap_contextmgt(9E) 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. devmap_do_ctxmgt () 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 devmap_access(9E) entry point and must not be modified. The range from <i>off</i> to <i>off+len</i> must support context switching. The system will pass through <i>dhp</i>, <i>pvtp</i>, <i>off</i>, <i>len</i>, <i>type</i>, and <i>rw</i> 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.

devmap_do_ctxmgt(9F)

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 *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.
     */
}
```

devmap_do_ctxmgt(9F)

EXAMPLE 1 Using devmap_do_ctxmgt in the devmap_access entry point. (Continued)

```
    } 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

devmap_set_ctx_timeout(9F)

NAME	devmap_set_ctx_timeout – set the timeout value for the context management callback
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void devmap_set_ctx_timeout(devmap_cookie_t <i>dhp</i>, clock_t <i>ticks</i>);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>dhp</i> An opaque mapping handle that the system uses to describe the mapping. <i>ticks</i> Number of clock ticks to wait between successive calls to the context management callback function.
DESCRIPTION	devmap_set_ctx_timeout() specifies the time interval for the system to wait between successive calls to the driver's context management callback function, devmap_contextmgt(9E). Device drivers typically call devmap_set_ctx_timeout() in the devmap_map(9E) routine. If the drivers do not call devmap_set_ctx_timeout() to set the timeout value, the default timeout value of 0 will result in no delay between successive calls to the driver's devmap_contextmgt(9E) callback function.
CONTEXT	devmap_set_ctx_timeout() can be called from user or interrupt context.
SEE ALSO	devmap_contextmgt(9E), devmap_map(9E), timeout(9F)

NAME	devmap_setup, ddi_devmap_segmap – set up a user mapping to device memory using the devmap framework																
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>																
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>dev</i> Device whose memory is to be mapped. <i>off</i> User offset within the logical device memory at which the mapping begins. <i>as</i> An opaque data structure that describes the address space into which the device memory should be mapped. <i>addrp</i> Pointer to the starting address in the address space into which the device memory should be mapped. <i>len</i> Length (in bytes) of the memory to be mapped. <i>prot</i> A bit field that specifies the protections. Some possible settings combinations are: <table> <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 mmap(2) system call).</td></tr> <tr> <td>PROT_ALL</td><td>All access is desired.</td></tr> </table> <i>maxprot</i> Maximum protection flag possible for attempted mapping; the PROT_WRITE bit may be masked out if the user opened the special file read-only. <i>flags</i> Flags indicating type of mapping. The following flags can be specified: <table> <tr> <td>MAP_PRIVATE</td><td>Changes are private.</td></tr> <tr> <td>MAP_SHARED</td><td>Changes should be shared.</td></tr> <tr> <td>MAP_FIXED</td><td>The user specified an address in *addrp rather than letting the system choose an address.</td></tr> </table> <i>cred</i> Pointer to the user credential structure.	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.	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.
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.																
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.																

devmap_setup(9F)

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 <i>off</i> and <i>len</i>. 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	<table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>Non-zero</td><td>An error occurred. The return value of devmap_setup() and ddi_devmap_segmap() should be used directly in the segmap(9E) entry point.</td></tr></table>	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.
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) <i>Writing Device Drivers</i>				

NAME	devmap_unload, devmap_load – control validation of memory address translations
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dhp</i> An opaque mapping handle that the system uses to describe the mapping. <i>off</i> User offset within the logical device memory at which the loading or unloading of the address translations begins. <i>len</i> Length (in bytes) of the range being affected. <i>type</i> Type of access operation. <i>rw</i> Direction of access.
devmap_load() only	
DESCRIPTION	devmap_unload() and devmap_load() are used to control the validation of the memory mapping described by <i>dhp</i> in the specified range. devmap_unload() invalidates the mapping translations and will generate calls to the devmap_access(9E) entry point next time the mapping is accessed. The drivers use devmap_load() to validate the mapping translations during memory access. A typical use of devmap_unload() and devmap_load() is in the driver's context management callback function, devmap_contextmgt(9E). To manage a device context, a device driver calls devmap_unload() on the context about to be switched out. It switches contexts, and then calls devmap_load() on the context switched in. devmap_unload() can be used to unload the mappings of other processes as well as the mappings of the calling process, but devmap_load() can only be used to load the mappings of the calling process. Attempting to load another process's mappings with devmap_load() will result in a system panic. 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, devmap_contextmgt(9E)) and should not be modified. Supplying a value of 0 for the <i>len</i> argument affects all addresses from the <i>off</i> to the end of the mapping. Supplying a value of 0 for the <i>off</i> argument and a value of 0 for <i>len</i> 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.

devmap_unload(9F)

	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. <pre>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)); }</pre>	
SEE ALSO	devmap_access(9E), devmap_contextmgt(9E) <i>Writing Device Drivers</i>	

NAME	disksort – single direction elevator seek sort for buffers
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> void disksort(struct diskhd *dp, struct buf *bp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>dp</i> A pointer to a diskhd structure. A diskhd structure is essentially identical to head of a buffer structure (see buf(9S)). The only defined items of interest for this structure are the av_forw and av_back structure elements which are used to maintain the front and tail pointers of the forward linked I/O request queue. <i>bp</i> A pointer to a buffer structure. Typically this is the I/O request that the driver receives in its strategy routine (see strategy(9E)). The driver is responsible for initializing the b_resid structure element to a meaningful sort key value prior to calling disksort().
DESCRIPTION	The function disksort() sorts a pointer to a buffer into a single forward linked list headed by the av_forw element of the argument *dp. 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 b_resid. 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. The head of the linked list is found by use of the av_forw structure element of the argument *dp. The tail of the linked list is found by use of the av_back structure element of the argument *dp. The av_forw element of the *bp argument is used by disksort() to maintain the forward linkage. The value at the head of the list presumably indicates the currently active disk area.
CONTEXT	This function can be called from user or interrupt context.
SEE ALSO	strategy(9E), buf(9S) <i>Writing Device Drivers</i>
WARNINGS	disksort() does no locking. Therefore, any locking is completely the responsibility of the caller.

drv_getparm(9F)

NAME	drv_getparm – retrieve kernel state information																		
SYNOPSIS	<pre>#include <sys/ddi.h> int drv_getparm(unsigned int <i>parm</i>, void *<i>value_p</i>);</pre>																		
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <table> <tr> <td><i>parm</i></td><td>The kernel parameter to be obtained. Possible values are:</td></tr> <tr> <td>LBOLT</td><td>Read the value of <code>lbolt</code>. <code>lbolt</code> is a <code>clock_t</code> that is unconditionally incremented by one at each clock tick. No special treatment is applied when this value overflows the maximum value of the signed integral type <code>clock_t</code>. 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.</td></tr> <tr> <td>PPGRP</td><td>Read the process group identification number. This number determines which processes should receive a HANGUP or BREAK signal when detected by a driver.</td></tr> <tr> <td>UPROCP</td><td>Read the process table token value.</td></tr> <tr> <td>PPID</td><td>Read process identification number.</td></tr> <tr> <td>PSID</td><td>Read process session identification number.</td></tr> <tr> <td>TIME</td><td>Read time in seconds.</td></tr> <tr> <td>UCRED</td><td>Return a pointer to the caller's credential structure.</td></tr> <tr> <td><i>value_p</i></td><td>A pointer to the data space in which the value of the parameter is to be copied.</td></tr> </table>	<i>parm</i>	The kernel parameter to be obtained. Possible values are:	LBOLT	Read the value of <code>lbolt</code> . <code>lbolt</code> is a <code>clock_t</code> that is unconditionally incremented by one at each clock tick. No special treatment is applied when this value overflows the maximum value of the signed integral type <code>clock_t</code> . 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.	PPGRP	Read the process group identification number. This number determines which processes should receive a HANGUP or BREAK signal when detected by a driver.	UPROCP	Read the process table token value.	PPID	Read process identification number.	PSID	Read process session identification number.	TIME	Read time in seconds.	UCRED	Return a pointer to the caller's credential structure.	<i>value_p</i>	A pointer to the data space in which the value of the parameter is to be copied.
<i>parm</i>	The kernel parameter to be obtained. Possible values are:																		
LBOLT	Read the value of <code>lbolt</code> . <code>lbolt</code> is a <code>clock_t</code> that is unconditionally incremented by one at each clock tick. No special treatment is applied when this value overflows the maximum value of the signed integral type <code>clock_t</code> . 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.																		
PPGRP	Read the process group identification number. This number determines which processes should receive a HANGUP or BREAK signal when detected by a driver.																		
UPROCP	Read the process table token value.																		
PPID	Read process identification number.																		
PSID	Read process session identification number.																		
TIME	Read time in seconds.																		
UCRED	Return a pointer to the caller's credential structure.																		
<i>value_p</i>	A pointer to the data space in which the value of the parameter is to be copied.																		
DESCRIPTION	Since the release of the Solaris 2.6 operating environment, the <code>drv_getparm()</code> function has been replaced by <code>ddi_get_lbolt(9F)</code>, <code>ddi_get_time(9F)</code>, and <code>ddi_get_pid(9F)</code>. <code>drv_getparm()</code> 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>. <code>drv_getparm()</code> 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 <i>value_p</i>. 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.																		

drv_getparm(9F)

RETURN VALUES	drv_getparm() returns 0 to indicate success, -1 to indicate failure. The value stored in the space pointed to by <i>value_p</i> 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	ddi_get_lbolt(9F), ddi_get_pid(9F), ddi_get_time(9F), buf(9S) <i>Writing Device Drivers</i>

drv_hztousec(9F)

NAME	drv_hztousec – convert clock ticks to microseconds
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> clock_t drv_hztousec(clock_t hertz);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>hertz</i> The number of clock ticks to convert.
DESCRIPTION	drv_hztousec() converts into microseconds the time expressed by <i>hertz</i>, which is in system clock ticks. The kernel variable <i>lbolt</i>, whose value should be retrieved by calling <i>ddi_get_lbolt</i>(9F), is the length of time the system has been up since boot and is expressed in clock ticks. Drivers often use the value of <i>lbolt</i> before and after an I/O request to measure the amount of time it took the device to process the request. <i>drv_hztousec</i>() can be used by the driver to convert the reading from clock ticks to a known unit of time.
RETURN VALUES	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 <i>clock_t</i> , then the maximum <i>clock_t</i> value will be returned.
CONTEXT	drv_hztousec() can be called from user or interrupt context.
SEE ALSO	<i>ddi_get_lbolt</i> (9F), <i>drv_usec2ohz</i> (9F), <i>drv_usecwait</i> (9F)
	<i>Writing Device Drivers</i>

NAME	drv_priv – determine driver privilege
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/cred.h> #include <sys/ddi.h> int drv_priv(cred_t *cr);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>cr</i> Pointer to the user credential structure.
DESCRIPTION	drv_priv() 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 suser() and any explicit checks for effective user ID = 0 in driver code.
RETURN VALUES	This routine returns 0 if it succeeds, EPERM if it fails.
CONTEXT	drv_priv() can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

drv_usectohz(9F)

NAME	drv_usectohz – convert microseconds to clock ticks
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> clock_t drv_usectohz(clock_t <i>microsecs</i>);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>microsecs</i> The number of microseconds to convert.
DESCRIPTION	drv_usectohz() converts a length of time expressed in microseconds to a number of system clock ticks. The time arguments to timeout(9F) and delay(9F) are expressed in clock ticks. drv_usectohz() is a portable interface for drivers to make calls to timeout(9F) and delay(9F) 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).
RETURN VALUES	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 clock_t, then the maximum clock_t value will be returned.
CONTEXT	drv_usectohz() can be called from user or interrupt context.
SEE ALSO	delay(9F), drv_hztousec(9F), timeout(9F) <i>Writing Device Drivers</i>

NAME	drv_usecwait – busy-wait for specified interval
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> void drv_usecwait(clock_t <i>microsecs</i>);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>microsecs</i> The number of microseconds to busy-wait.
DESCRIPTION	drv_usecwait() 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. delay(9F) 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, delay(9F) may only be invoked from user context and hence cannot be used at interrupt time or system initialization. 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, delay(9F) produces too long a wait period.
CONTEXT	drv_usecwait() can be called from user or interrupt context.
SEE ALSO	delay(9F), timeout(9F), untimeout(9F) <i>Writing Device Drivers</i>
NOTES	The driver wastes processor time by making this call since drv_usecwait() does not block but simply busy-waits. The driver should only make calls to drv_usecwait() as needed, and only for as much time as needed. drv_usecwait() does not mask out interrupts.

dupb(9F)

NAME	dupb – duplicate a message block descriptor
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *dupb(mblk_t *bp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	dupb() creates a new mblk_t structure (see msgb(9S)) to reference the message block pointed to by bp. Unlike copyb(9F), dupb() does not copy the information in the dblk_t structure (see datab(9S)), 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 bp. nbp=dupb(bp);
PARAMETERS	<i>bp</i> Pointer to the message block to be duplicated. mblk_t is an instance of the msgb(9S) structure.
RETURN VALUES	If successful, dupb() returns a pointer to the new message block. A NULL pointer is returned if dupb() cannot allocate a new message block descriptor or if the db_ref field of the data block structure (see datab(9S)) has reached a maximum value (255).
CONTEXT	dupb() 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) {
```

dupb(9F)

EXAMPLE 1 Using dupb() (Continued)

```
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
53 xxxclose(q, cflag, crp)
54     queue_t *q;
55     int cflag;
56     cred_t *crp;
57 {
58     struct xx *xx = (struct xx *)q->q_ptr;
59     ...
60     if (xx->xx_qbufcall_id)
61         qunbufcall(q, xx->xx_qbufcall_id);
62     ...
63 }
```

SEE ALSO [srv\(9E\)](#), [copyb\(9F\)](#), [qbufcall\(9F\)](#), [datab\(9S\)](#), [msgb\(9S\)](#)

Writing Device Drivers STREAMS Programming Guide

NAME	dupmsg – duplicate a message
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *dupmsg(mblk_t *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Pointer to the message.
DESCRIPTION	dupmsg() forms a new message by copying the message block descriptors pointed to by <i>mp</i> and linking them. dupb(9F) is called for each message block. The data blocks themselves are not duplicated.
RETURN VALUES	If successful, dupmsg() 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, db_ref (see datab(9S)), has reached a limit (255). See dupb(9F).
CONTEXT	dupmsg() can be called from user, kernel, or interrupt context.
EXAMPLES	EXAMPLE 1 Using dupmsg() See copyb(9F) for an example using dupmsg().
SEE ALSO	copyb(9F), copymsg(9F), dupb(9F), datab(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

enableok(9F)

NAME	enableok – reschedule a queue for service
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void enableok(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> A pointer to the queue to be rescheduled.
DESCRIPTION	enableok() enables queue <i>q</i> to be rescheduled for service. It reverses the effect of a previous call to noenable(9F) on <i>q</i> by turning off the QNOENB flag in the queue.
CONTEXT	enableok() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using enableok() The qrestart() routine uses two STREAMS functions to restart a queue that has been disabled. The enableok() function turns off the QNOENB flag, allowing the qenable(9F) to schedule the queue for immediate processing. <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>
SEE ALSO	noenable(9F), qenable(9F) <i>Writing Device Drivers STREAMS Programming Guide</i>

NAME	esballoc – allocate a message block using a caller-supplied buffer
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *esballoc (uchar *<i>base</i>, size_t <i>size</i>, uint_t <i>pri</i>, frtn_t *<i>fr_rtnp</i>) ;</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>base</i> Address of user supplied data buffer. <i>size</i> Number of bytes in data buffer. <i>pri</i> Priority of allocation request (to be used by allocb(9F) function, called by esballoc()). <i>fr_rtnp</i> Free routine data structure.
DESCRIPTION	esballoc() creates a STREAMS message and attaches a user-supplied data buffer in place of a STREAMS data buffer. It calls allocb(9F) to get a message and data block header only. The newly allocated message will have both the b_wptr and b_rptr set to the base of the buffer. As when using allocb(9F), the newly allocated message will have both b_wptr and b_rptr set to the base of the data buffer. The user-supplied data buffer, pointed to by <i>base</i>, is used as the data buffer for the message. When freeb(9F) is called to free the message, the driver's message freeing routine (referenced through the free_rtn structure) is called, with appropriate arguments, to free the data buffer. The free_rtn structure includes the following members: <pre>void (*free_func)(); /* user's freeing routine */ char *free_arg; /* arguments to free_func() */</pre> Instead of requiring a specific number of arguments, the free_arg field is defined of type char *. This way, the driver can pass a pointer to a structure if more than one argument is needed. The method by which free_func is called is implementation-specific. The module writer must not assume that free_func will or will not be called directly from STREAMS utility routines like freeb(9F) which free a message block. free_func 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. free_func must not access any dynamically allocated data structure that might no longer exist when it runs.
RETURN VALUES	On success, a pointer to the newly allocated message block is returned. On failure, NULL is returned.

esballoc(9F)

CONTEXT esballoc () can be called from user or interrupt context.

SEE ALSO allocb(9F), freeb(9F), datab(9S), free_rtn(9S)

Writing Device Drivers STREAMS Programming Guide

WARNINGS The free_func must be defined in kernel space, should be declared void and accept one argument. It has no user context and must not sleep.

NAME	esbbcall – call function when buffer is available
SYNOPSIS	<pre>#include <sys/stream.h> bufcall_id_t esbbcall(uint_t <i>pri</i>, void *<i>func</i>void *<i>arg</i>, void <i>arg</i>);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>pri</i> Priority of allocation request (to be used by allocb(9F) function, called by esbbcall()). <i>func</i> Function to be called when buffer becomes available. <i>arg</i> Argument to <i>func</i>.
DESCRIPTION	esbbcall(), like bufcall(9F), serves as a timeout(9F) call of indeterminate length. If esballoc(9F) is unable to allocate a message and data block header to go with its externally supplied data buffer, esbbcall() 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 esballoc(9F) or it may be another kernel function.
RETURN VALUES	On success, a bufcall ID is returned. On failure, 0 is returned. The value returned from a successful call should be saved for possible future use with unbufcall() should it become necessary to cancel the esbbcall() request (as at driver close time).
CONTEXT	esbbcall() can be called from user or interrupt context.
SEE ALSO	allocb(9F), bufcall(9F), esballoc(9F), timeout(9F), datab(9S), unbufcall(9F) <i>Writing Device Drivers STREAMS Programming Guide</i>

flushband(9F)

NAME	flushband – flush messages for a specified priority band
SYNOPSIS	<pre>#include <sys/stream.h> void flushband(queue_t *q, unsigned char pri, int flag);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue. <i>pri</i> Priority of messages to be flushed. <i>flag</i> 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.
DESCRIPTION	flushband() 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> .
CONTEXT	flushband() can be called from user or interrupt context.
SEE ALSO	flushq(9F)

Writing Device Drivers STREAMS Programming Guide

NAME	flushq – remove messages from a queue
SYNOPSIS	<pre>#include <sys/stream.h> void flushq(queue_t *q, int flag);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue to be flushed. <i>flag</i> 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.
DESCRIPTION	flushq() frees messages and their associated data structures by calling freemsg(9F). 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.
CONTEXT	flushq() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using flushq() 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 qreply(9F) 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->b_datap->db_type) { 9 case M_FLUSH: 10 if (*mp->b_rptr & FLUSHR) 11 flushq(RD(q), FLUSHALL); 12 if (*mp->b_rptr & FLUSHW) 13 flushq(q, FLUSHALL); 14 putnext(q, mp); 15 break; 16 . . . 16 } 17 }</pre>
SEE ALSO	flushband(9F), freemsg(9F), putq(9F), qreply(9F)

flushq(9F)

Writing Device Drivers STREAMS Programming Guide

NAME	freeb – free a message block
SYNOPSIS	<pre>#include <sys/stream.h> void freeb(mblk_t *bp);</pre>
PARAMETERS	<i>bp</i> Pointer to the message block to be deallocated. mblk_t is an instance of the msgb(9S) structure.
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	freeb() deallocates a message block. If the reference count of the db_ref member of the datab(9S) structure is greater than 1, freeb() decrements the count. If db_ref 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 esballoc(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, freeb() finds the free_rtn(9S) structure associated with the buffer. The free_rtn structure contains a pointer to the driver-dependent routine, which releases the buffer. Once this is accomplished, freeb() releases the STREAMS resources associated with the buffer.
CONTEXT	freeb() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using freeb() See copyb(9F) for an example of using freeb().
SEE ALSO	allocb(9F), copyb(9F), dupb(9F), esballoc(9F), free_rtn(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

freemsg(9F)

NAME	freemsg – free all message blocks in a message
SYNOPSIS	<pre>#include <sys/stream.h> void freemsg(mblk_t *mp);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>mp</i> Pointer to the message blocks to be deallocated. mblk_t is an instance of the msgb(9S) structure. If <i>mp</i> is NULL, freemsg() immediately returns.
DESCRIPTION	freemsg() calls freeb(9F) to free all message and data blocks associated with the message pointed to by <i>mp</i> .
CONTEXT	freemsg() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using freemsg() See copymsg(9F).
SEE ALSO	copymsg(9F), freeb(9F), msgb(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
NOTES	The behavior of freemsg() when passed a NULL pointer is Solaris-specific.

NAME	freerbuf – free a raw buffer header
SYNOPSIS	<pre>#include <sys/buf.h> #include <sys/ddi.h> void freerbuf(struct buf *bp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>bp</i> Pointer to a previously allocated buffer header structure.
DESCRIPTION	freerbuf () frees a raw buffer header previously allocated by getrbuf(9F). This function does not sleep and so may be called from an interrupt routine.
CONTEXT	freerbuf () can be called from user or interrupt context.
SEE ALSO	getrbuf(9F), kmem_alloc(9F), kmem_free(9F), kmem_zalloc(9F)

freezestr(9F)

NAME	freezestr, unfreezestr – freeze, thaw the state of a stream
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void freezestr(queue_t *q); void unfreezestr(queue_t *q);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>q</i> Pointer to the message queue to freeze/unfreeze.
DESCRIPTION	freezestr() 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 unfreezestr(). Drivers and modules must freeze the stream before manipulating the queues directly (as opposed to manipulating them through programmatic interfaces such as getq(9F), putq(9F), putbq(9F), etc.)
CONTEXT	These routines may be called from any stream open, close, put or service routine as well as interrupt handlers, callouts and call-backs.
SEE ALSO	<i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
NOTES	The freezestr() and unfreezestr() functions can have a serious impact on system performance. Their use should be very limited. In most cases, there is no need to use freezestr() and there are usually better ways to accomplish what you need to do than by freezing the stream. Calling freezestr() to freeze a stream that is already frozen by the caller will result in a single-party deadlock. The caller of unfreezestr() must be the thread who called freezestr(). STREAMS utility functions such as getq(9F), putq(9F), putbq(9F), and so forth, should not be called by the caller of freezestr() while the stream is still frozen, as they indirectly freeze the stream to ensure atomicity of queue manipulation.

NAME	geterror – return I/O error
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/buf.h> #include <sys/ddi.h> int geterror(struct buf *bp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>bp</i> Pointer to a buf(9S) structure.
DESCRIPTION	geterror() returns the error number from the error field of the buffer header structure.
RETURN VALUES	An error number indicating the error condition of the I/O request is returned. If the I/O request completes successfully, 0 is returned.
CONTEXT	geterror() can be called from user or interrupt context.
SEE ALSO	buf(9S) <i>Writing Device Drivers</i>

gethrtime(9F)

NAME	gethrtime – get high resolution time
SYNOPSIS	<pre>#include <sys/time.h> hrtime_t gethrtime(void);</pre>
DESCRIPTION	The <code>gethrtime()</code> function returns the current high-resolution real time. Time is expressed as nanoseconds since some arbitrary time in the past; it is not correlated in any way to the time of day, and thus is not subject to resetting or drifting by way of <code>adjtime(2)</code> or <code>settimeofday(3C)</code> . The hi-res timer is ideally suited to performance measurement tasks, where cheap, accurate interval timing is required.
RETURN VALUES	<code>gethrtime()</code> always returns the current high-resolution real time. There are no error conditions.
CONTEXT	There are no restrictions on the context from which <code>gethrtime()</code> can be called.
SEE ALSO	<code>proc(1)</code> , <code>gettimeofday(3C)</code> , <code>settimeofday(3C)</code> , <code>attributes(5)</code>
NOTES	Although the units of hi-res time are always the same (nanoseconds), the actual resolution is hardware dependent. Hi-res time is guaranteed to be monotonic (it does not go backward, it does not periodically wrap) and linear (it does not occasionally speed up or slow down for adjustment, as the time of day can), but not necessarily unique: two sufficiently proximate calls might return the same value. The time base used for this function is the same as that for <code>gethrtime(3C)</code>. Values returned by both of these functions can be interleaved for comparison purposes.

NAME	getmajor – get major device number
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mkdev.h> #include <sys/ddi.h> major_t getmajor(dev_t dev);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>dev</i> Device number.
DESCRIPTION	getmajor() extracts the major number from a device number.
RETURN VALUES	The major number.
CONTEXT	getmajor() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using getmajor() The following example shows both the getmajor() and getminor(9F) functions used in a debug cmn_err(9F) statement to return the major and minor numbers for the device supported by the driver. <pre>dev_t dev; #ifdef DEBUG cmn_err(CE_NOTE, "Driver Started. Major# = %d, Minor# = %d", getmajor(dev), getminor(dev)); #endif</pre>
SEE ALSO	cmn_err(9F), getminor(9F), makedevice(9F) <i>Writing Device Drivers</i>
WARNINGS	No validity checking is performed. If <i>dev</i> is invalid, an invalid number is returned.

getminor(9F)

NAME	getminor – get minor device number
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mkdev.h> #include <sys/ddi.h> minor_t getminor(dev_t dev);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>dev</i> Device number.
DESCRIPTION	getminor() extracts the minor number from a device number.
RETURN VALUES	The minor number.
CONTEXT	getminor() can be called from user or interrupt context.
EXAMPLES	See the getmajor(9F) manual page for an example of how to use getminor().
SEE ALSO	getmajor(9F), makedevice(9F) <i>Writing Device Drivers</i>
WARNINGS	No validity checking is performed. If <i>dev</i> is invalid, an invalid number is returned.

NAME	get_pktiopb, free_pktiopb – allocate/free a SCSI packet in the iopb map																
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> 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>																
INTERFACE LEVEL	These interfaces are obsolete. Use scsi_alloc_consistent_buf(9F) instead of get_pktiopb(). Use scsi_free_consistent_buf(9F) instead of free_pktiopb().																
PARAMETERS	<table> <tr> <td><i>ap</i></td><td>Pointer to the target's <code>scsi_address</code> structure.</td></tr> <tr> <td><i>datap</i></td><td>Pointer to the address of the packet, set by this function.</td></tr> <tr> <td><i>cdblen</i></td><td>Number of bytes required for the SCSI command descriptor block (CDB).</td></tr> <tr> <td><i>statuslen</i></td><td>Number of bytes required for the SCSI status area.</td></tr> <tr> <td><i>datalen</i></td><td>Number of bytes required for the data area of the SCSI command.</td></tr> <tr> <td><i>readflag</i></td><td>If non-zero, data will be transferred from the SCSI target.</td></tr> <tr> <td><i>callback</i></td><td>Pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code></td></tr> <tr> <td><i>pkt</i></td><td>Pointer to a <code>scsi_pkt</code>(9S) structure.</td></tr> </table>	<i>ap</i>	Pointer to the target's <code>scsi_address</code> structure.	<i>datap</i>	Pointer to the address of the packet, set by this function.	<i>cdblen</i>	Number of bytes required for the SCSI command descriptor block (CDB).	<i>statuslen</i>	Number of bytes required for the SCSI status area.	<i>datalen</i>	Number of bytes required for the data area of the SCSI command.	<i>readflag</i>	If non-zero, data will be transferred from the SCSI target.	<i>callback</i>	Pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code>	<i>pkt</i>	Pointer to a <code>scsi_pkt</code> (9S) structure.
<i>ap</i>	Pointer to the target's <code>scsi_address</code> structure.																
<i>datap</i>	Pointer to the address of the packet, set by this function.																
<i>cdblen</i>	Number of bytes required for the SCSI command descriptor block (CDB).																
<i>statuslen</i>	Number of bytes required for the SCSI status area.																
<i>datalen</i>	Number of bytes required for the data area of the SCSI command.																
<i>readflag</i>	If non-zero, data will be transferred from the SCSI target.																
<i>callback</i>	Pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code>																
<i>pkt</i>	Pointer to a <code>scsi_pkt</code> (9S) structure.																
DESCRIPTION	<code>get_pktiopb()</code> 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</code>(9F) for allocating the data area and <code>scsi_realloc</code>(9F) to allocate the packet and DMA resources. <i>callback</i> indicates what <code>get_pktiopb()</code> should do when resources are not available: <table> <tr> <td><code>NULL_FUNC</code></td><td>Do not wait for resources. Return a <code>NULL</code> pointer.</td></tr> <tr> <td><code>SLEEP_FUNC</code></td><td>Wait indefinitely for resources.</td></tr> <tr> <td>Other Values</td><td><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.</td></tr> </table> <code>free_pktiopb()</code> is used for freeing the packet and its associated resources.	<code>NULL_FUNC</code>	Do not wait for resources. Return a <code>NULL</code> pointer.	<code>SLEEP_FUNC</code>	Wait indefinitely for resources.	Other Values	<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.										
<code>NULL_FUNC</code>	Do not wait for resources. Return a <code>NULL</code> pointer.																
<code>SLEEP_FUNC</code>	Wait indefinitely for resources.																
Other Values	<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.																
RETURN VALUES	<code>get_pktiopb()</code> returns a pointer to the newly allocated <code>scsi_pkt</code> or a <code>NULL</code> pointer.																

get_pktiopb(9F)

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.

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

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `attributes(5)`, `ddi_iopb_alloc(9F)`, `scsi_alloc_consistent_buf(9F)`, `scsi_free_consistent_buf(9F)`, `scsi_pktalloc(9F)`, `scsi_realloc(9F)`, `scsi_pkt(9S)`

Writing Device Drivers

NOTES The `get_pktiopb()` and `free_pktiopb()` functions are obsolete and will be discontinued in a future release. These functions have been replaced by, respectively, `scsi_alloc_consistent_buf(9F)` and `scsi_free_consistent_buf(9F)`.

`get_pktiopb()` uses scarce resources. For this reason and its obsolescence (see above), its use is discouraged.

NAME	getq – get the next message from a queue
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *getq(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue from which the message is to be retrieved.
DESCRIPTION	getq() is used by a service (srv(9E)) routine to retrieve its enqueued messages. A module or driver may include a service routine to process enqueued messages. Once the STREAMS scheduler calls srv() it must process all enqueued messages, unless prevented by flow control. getq() obtains the next available message from the top of the queue pointed to by <i>q</i>. It should be called in a while 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, getq() back-enables (restarts) the service routine once it falls below the low water mark.
RETURN VALUES	If there is a message to retrieve, getq() returns a pointer to it. If no message is queued, getq() returns a NULL pointer.
CONTEXT	getq() can be called from user or interrupt context.
EXAMPLES	See dupb(9F).
SEE ALSO	srv(9E), bcanput(9F), canput(9F), dupb(9F), putbq(9F), putq(9F), qenable(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

getrbuf(9F)

NAME	getrbuf – get a raw buffer header
SYNOPSIS	<pre>#include <sys/buf.h> #include <sys/kmem.h> #include <sys/ddi.h> struct buf *getrbuf(int sleepflag);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>sleepflag</i> Indicates whether driver should sleep for free space.
DESCRIPTION	getrbuf() 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. getrbuf() calls kmem_alloc(9F) to perform the memory allocation. kmem_alloc() requires the information included in the <i>sleepflag</i> argument. If <i>sleepflag</i> is set to KM_SLEEP, the driver may sleep until the space is freed up. If <i>sleepflag</i> is set to KM_NOSLEEP, 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.
RETURN VALUES	getrbuf() returns a pointer to the allocated buffer header, or NULL if no space is available.
CONTEXT	getrbuf() can be called from user or interrupt context. (Drivers must not allow getrbuf() to sleep if called from an interrupt routine.)
SEE ALSO	bioinit(9F), freerbuf(9F), kmem_alloc(9F), kmem_free(9F) <i>Writing Device Drivers</i>

NAME	<code>gld</code> , <code>gld_mac_alloc</code> , <code>gld_mac_free</code> , <code>gld_register</code> , <code>gld_unregister</code> , <code>gld_recv</code> , <code>gld_sched</code> , <code>gld_intr</code> – Generic LAN Driver service routines								
SYNOPSIS	<pre>#include <sys/gld.h> gld_mac_info_t *gld_mac_alloc(dev_info_t *dip); void gld_mac_free(gld_mac_info_t *macinfo); int gld_register(dev_info_t *dip, char *name, gld_mac_info_t *macinfo); int gld_unregister(gld_mac_info_t *macinfo); void gld_recv(gld_mac_info_t *macinfo, mblk_t *mp); void gld_sched(gld_mac_info_t *macinfo); uint_t gld_intr(caddr_t);</pre>								
INTERFACE LEVEL PARAMETERS	Solaris architecture specific (Solaris DDI). <table> <tr> <td><i>macinfo</i></td><td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td></tr> <tr> <td><i>dip</i></td><td>Pointer to <code>dev_info</code> structure.</td></tr> <tr> <td><i>name</i></td><td>Device interface name.</td></tr> <tr> <td><i>mp</i></td><td>Pointer to a message block containing a received packet.</td></tr> </table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>dip</i>	Pointer to <code>dev_info</code> structure.	<i>name</i>	Device interface name.	<i>mp</i>	Pointer to a message block containing a received packet.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.								
<i>dip</i>	Pointer to <code>dev_info</code> structure.								
<i>name</i>	Device interface name.								
<i>mp</i>	Pointer to a message block containing a received packet.								
DESCRIPTION	<code>gld_mac_alloc()</code> allocates a new <code>gld_mac_info(9S)</code> structure and returns a pointer to it. Some of the GLD-private elements of the structure may be initialized before <code>gld_mac_alloc()</code> returns; all other elements are initialized to zero. The device driver must initialize some structure members, as described in <code>gld_mac_info(9S)</code>, before passing the <code>macinfo</code> pointer to <code>gld_register()</code>. <code>gld_mac_free()</code> frees a <code>gld_mac_info(9S)</code> structure previously allocated by <code>gld_mac_alloc()</code>. <code>gld_register()</code> is called from the device driver's <code>attach(9E)</code> routine, and is used to link the GLD-based device driver with the GLD framework. Before calling <code>gld_register()</code> the device driver's <code>attach(9E)</code> routine must first use <code>gld_mac_alloc()</code> to allocate a <code>gld_mac_info(9S)</code> structure, and initialize several of its structure elements. See <code>gld_mac_info(9S)</code> for more information. A successful call to <code>gld_register()</code> performs the following actions: ■ links the device-specific driver with the GLD system; ■ sets the device-specific driver's private data pointer (using <code>ddi_set_driver_private(9F)</code>) to point to the <code>macinfo</code> structure; ■ creates the minor device node. The device interface name passed to <code>gld_register()</code> must exactly match the name of the driver module as it exists in the filesystem.								

gld(9F)

The driver's `attach(9E)` routine should return `DDI_SUCCESS` if `gld_register()` succeeds. If `gld_register()` returns `DDI_FAILURE`, the `attach(9E)` routine should deallocate any resources it allocated before calling `gld_register()` and then also return `DDI_FAILURE`.

`gld_unregister()` is called by the device driver's `detach(9E)` function, and if successful, performs the following tasks:

- ensures the device's interrupts are stopped, calling the driver's `gldm_stop()` routine if necessary;
- removes the minor device node;
- unlinks the device-specific driver from the GLD system.

If `gld_unregister()` returns `DDI_SUCCESS`, the `detach(9E)` routine should deallocate any data structures allocated in the `attach(9E)` routine, using `gld_mac_free()` to deallocate the `macinfo` structure, and return `DDI_SUCCESS`. If `gld_unregister()` returns `DDI_FAILURE`, the driver's `detach(9E)` routine must leave the device operational and return `DDI_FAILURE`.

`gld_recv()` is called by the driver's interrupt handler to pass a received packet upstream. The driver must construct and pass a `STREAMS_M_DATA` message containing the raw packet. `gld_recv()` determines which `STREAMS` queues, if any, should receive a copy of the packet, duplicating it if necessary. It then formats a `DL_UNITDATA_IND` message, if required, and passes the data up all appropriate streams.

The driver should avoid holding mutex or other locks during the call to `gld_recv()`. In particular, locks that could be taken by a transmit thread may not be held during a call to `gld_recv()`: the interrupt thread that calls `gld_recv()` may in some cases carry out processing that includes sending an outgoing packet, resulting in a call to the driver's `gldm_send()` routine. If the `gldm_send()` routine were to try to acquire a mutex being held by the `gldm_intr()` routine at the time it calls `gld_recv()`, this could result in a panic due to recursive mutex entry.

`gld_sched()` is called by the device driver to reschedule stalled outbound packets. Whenever the driver's `gldm_send()` routine has returned `GLD_NORESOURCES`, the driver must later call `gld_sched()` to inform the GLD framework that it should retry the packets that previously could not be sent. `gld_sched()` should be called as soon as possible after resources are again available, to ensure that GLD resumes passing outbound packets to the driver's `gldm_send()` routine in a timely way. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

`gld_intr()` is GLD's main interrupt handler. Normally it is specified as the interrupt routine in the device driver's call to `ddi_add_intr(9F)`. The argument to the interrupt handler (specified as `int_handler_arg` in the call to `ddi_add_intr(9F)`) must be a pointer to the `gld_mac_info(9S)` structure. `gld_intr()` will, when appropriate,

call the device driver's `gldm_intr()` function, passing that pointer to the `gld_mac_info(9S)` structure. However, if the driver uses a high-level interrupt, it must provide its own high-level interrupt handler, and trigger a soft interrupt from within that. In this case, `gld_intr()` may be specified as the soft interrupt handler in the call to `ddi_add_softintr()`.

RETURN VALUES `gld_mac_alloc()` returns a pointer to a new `gld_mac_info(9S)` structure.

`gld_register()` and `gld_unregister()` return:

`DDI_SUCCESS` on success.

`DDI_FAILURE` on failure.

`gld_intr()` returns a value appropriate for an interrupt handler.

SEE ALSO `gld(7D)`, `gld(9E)`, `gld_mac_info(9S)`, `gld_stats(9S)`, `dlpi(7P)`, `attach(9E)`, `ddi_add_intr(9F)`.

Writing Device Drivers

hat_getkpfnum(9F)

NAME	hat_getkpfnum – get page frame number for kernel address		
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> pfnt_t hat_getkpfnum(caddr_t addr);</pre>		
INTERFACE LEVEL	Architecture independent level 2 (DKI only).		
PARAMETERS	<table><tr><td><i>addr</i></td><td>The kernel virtual address for which the page frame number is to be returned.</td></tr></table>	<i>addr</i>	The kernel virtual address for which the page frame number is to be returned.
<i>addr</i>	The kernel virtual address for which the page frame number is to be returned.		
DESCRIPTION	hat_getkpfnum() returns the page frame number corresponding to the kernel virtual address, <i>addr</i>. <i>addr</i> must be a kernel virtual address which maps to device memory. ddi_map_regs(9F) can be used to obtain this address. For example, ddi_map_regs(9F) can be called in the driver's attach(9E) routine. The resulting kernel virtual address can be saved by the driver (see ddi_soft_state(9F)) and used in mmap(9E). The corresponding ddi_unmap_regs(9F) call can be made in the driver's detach(9E) routine. Refer to mmap(9E) for more information.		
RETURN VALUES	The page frame number corresponding to the valid virtual address <i>addr</i> . Otherwise the return value is undefined.		
CONTEXT	hat_getkpfnum() can be called only from user or kernel context.		
SEE ALSO	attach(9E), detach(9E), mmap(9E), ddi_map_regs(9F), ddi_soft_state(9F), ddi_unmap_regs(9F) <i>Writing Device Drivers</i>		
NOTES	For some devices, mapping device memory in the driver's attach(9E) routine and unmapping device memory in the driver's detach(9E) routine is a sizeable drain on system resources. This is especially true for devices with a large amount of physical address space. Refer to mmap(9E) for alternative methods.		

NAME	id32_alloc, id32_free, id32_lookup – 32-bit driver ID management routines
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/id32.h> uint32_t id32_alloc(void *ptr, int flag); void id32_free(uint32_t token); void *id32_lookup(uint32_t token);</pre>
INTERFACE LEVEL PARAMETERS	Solaris architecture specific (Solaris DDI). <i>ptr</i> any valid 32- or 64-bit pointer <i>flag</i> determines whether caller can sleep for memory (see <code>kmem_alloc(9F)</code> for a description)
DESCRIPTION	These routines were originally developed so that device drivers could manage 64-bit pointers on devices that save space only for 32-bit pointers. Many device drivers need to pass a 32-bit value to the hardware when attempting I/O. Later, when that I/O completes, the only way the driver has to identify the request that generated that I/O is via a "token". When the I/O is initiated, the driver passes this token to the hardware. When the I/O completes the hardware passes back this 32-bit token. Before Solaris supported 64-bit pointers, device drivers just passed a raw 32-bit pointer to the hardware. When pointers grew to be 64 bits this was no longer possible. The <code>id32_*()</code> routines were created to help drivers translate between 64-bit pointers and a 32-bit token. Given a 32- or 64-bit pointer, the routine <code>id32_alloc()</code> allocates a 32-bit token, returning 0 if <code>KM_NOSLEEP</code> was specified and memory could not be allocated. The allocated token is passed back to <code>id32_lookup()</code> to obtain the original 32- or 64-bit pointer. The routine <code>id32_free()</code> is used to free an allocated token. Once <code>id32_free()</code> is called, the supplied token is no longer valid. Note that these routines have some degree of error checking. This is done so that an invalid token passed to <code>id32_lookup()</code> will not be accepted as valid. When <code>id32_lookup()</code> detects an invalid token it returns NULL. Calling routines should check for this return value so that they do not try to dereference a NULL pointer.
CONTEXT	These functions can be called from user or interrupt context. The routine <code>id32_alloc()</code> should not be called from interrupt context when the <code>KM_SLEEP</code> flag is passed in. All other routines can be called from interrupt or kernel context.
SEE ALSO	<code>kmem_alloc(9F)</code> <i>Writing Device Drivers</i>

inb(9F)

NAME	inb, inw, inl, repinsb, repinsw, repinsd – read from an I/O port						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> unsigned char inb(int port); unsigned short inw(int port); unsigned long inl(int port); void repinsb(int port, unsigned char *addr, int count); void repinsw(int port, unsigned short *addr, int count); void repinsd(int port, unsigned long *addr, int count);</pre>						
INTERFACE LEVEL	The functions described here are obsolete. For the <code>inb()</code> , <code>inw()</code> , and <code>inl()</code> functions, use, respectively, <code>ddi_get8(9F)</code> , <code>ddi_get16(9F)</code> , and <code>ddi_get32(9F)</code> instead. For <code>repinsb()</code> , <code>repinsw()</code> , and <code>repinsl()</code> , use, respectively, <code>ddi_rep_get8(9F)</code> , <code>ddi_rep_get16(9F)</code> , and <code>ddi_rep_get32(9F)</code> instead.						
PARAMETERS	<table><tr><td><i>port</i></td><td>A valid I/O port address.</td></tr><tr><td><i>addr</i></td><td>The address of a buffer where the values will be stored.</td></tr><tr><td><i>count</i></td><td>The number of values to be read from the I/O port.</td></tr></table>	<i>port</i>	A valid I/O port address.	<i>addr</i>	The address of a buffer where the values will be stored.	<i>count</i>	The number of values to be read from the I/O port.
<i>port</i>	A valid I/O port address.						
<i>addr</i>	The address of a buffer where the values will be stored.						
<i>count</i>	The number of values to be read from the I/O port.						
DESCRIPTION	These routines read data of various sizes from the I/O port with the address specified by <i>port</i>. The <code>inb()</code>, <code>inw()</code>, and <code>inl()</code> functions read 8 bits, 16 bits, and 32 bits of data respectively, returning the resulting values. The <code>repinsb()</code>, <code>repinsw()</code>, and <code>repinsd()</code> 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.						
RETURN VALUES	<code>inb()</code> , <code>inw()</code> , and <code>inl()</code> return the value that was read from the I/O port.						
CONTEXT	These functions may be called from user or interrupt context.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Architecture</td><td>x86</td></tr><tr><td>Stability Level</td><td>Obsolete</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	x86	Stability Level	Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Architecture	x86						
Stability Level	Obsolete						
SEE ALSO	<code>eisa(4)</code> , <code>isa(4)</code> , <code>attributes(5)</code> , <code>ddi_get8(9F)</code> , <code>ddi_get16(9F)</code> , <code>ddi_get32(9F)</code> , <code>ddi_rep_get8(9F)</code> , <code>ddi_rep_get16(9F)</code> , <code>ddi_rep_get32(9F)</code> , <code>outb(9F)</code>						

inb(9F)

Writing Device Drivers

insq(9F)

NAME	insq – insert a message into a queue
SYNOPSIS	<pre>#include <sys/stream.h> int insq(queue_t *q, mblk_t *emp, mblk_t *nmp);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>q</i> Pointer to the queue containing message <i>emp</i>. <i>emp</i> Enqueued message before which the new message is to be inserted. mbblk_t is an instance of the msgb(9S) structure. <i>nmp</i> Message to be inserted.
DESCRIPTION	insq() 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.
RETURN VALUES	insq() returns 1 on success, and 0 on failure.
CONTEXT	insq() can be called from user or interrupt context.
EXAMPLES	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, bp will be NULL and we fall out of the for loop (line 21). insq acts like putq(9F) in this case. <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->b_datap->db_type) { 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 /* Insert code here to protect queue and message block */ 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;</pre>

```

24             if (ntp->type != T_EXDATA_REQ)
25                 break;
26             }
27         }
28         (void)insq(q, bp, mp);
29         /* End of region that must be protected */
30         break;
31     . . .
32 }
33 }

```

When using `insq()`, you must ensure that the queue and the message block is not modified by another thread at the same time. You can achieve this either by using STREAMS functions or by implementing your own locking.

SEE ALSO `putq(9F)`, `rmvq(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.

IOC_CONVERT_FROM(9F)

NAME	IOC_CONVERT_FROM – determine if there is a need to translate M_IOCTL contents.
SYNOPSIS	<pre>#include <sys/stream.h> uint_t IOC_CONVERT_FROM(struct iocblk *iocp);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI Specific (Solaris DDI) <i>iocp</i> A pointer to the M_IOCTL control structure.
DESCRIPTION	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.
RETURN VALUES	IOC_CONVERT_FROM() returns the following values: IOC_ILP32 This is an LP64 kernel and the M_IOCTL originated in an ILP32 user process. IOC_NONE The M_IOCTL message uses the same C Language Type Model as this calling module or driver.
CONTEXT	IOC_CONVERT_FROM() can be called from user or interrupt context.
SEE ALSO	<i>ddi_model_convert_from(9F)</i> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	kmem_alloc, kmem_zalloc, kmem_free – allocate kernel memory
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kmem.h> 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>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>size</i> Number of bytes to allocate. <i>flag</i> 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. <i>buf</i> Pointer to allocated memory.
DESCRIPTION	kmem_alloc() 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 kmem_alloc() are random garbage. kmem_zalloc() is like kmem_alloc() but returns zero-filled memory. kmem_free() frees previously allocated kernel memory. The buffer address and size must exactly match the original allocation. Memory cannot be returned piecemeal.
RETURN VALUES	If successful, kmem_alloc() and kmem_zalloc() return a pointer to the allocated memory. If KM_NOSLEEP is set and memory cannot be allocated without sleeping, kmem_alloc() and kmem_zalloc() return NULL.
CONTEXT	kmem_alloc() and kmem_zalloc() 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>. kmem_free() can be called from user or interrupt context.
SEE ALSO	copyout(9F), freerbuf(9F), getrbuf(9F) <i>Writing Device Drivers</i>
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.

kmem_alloc(9F)

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.

NAME	kmem_cache_create, kmem_cache_alloc, kmem_cache_free, kmem_cache_destroy – kernel memory cache allocator operations
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kmem.h> kmem_cache_t *kmem_cache_create(char *name, size_t bufsize, size_t align, int (*constructor)(void *, void *, int), void (*destructor)(void *, void *), void (*reclaim)(void *), void *private, void *vmp, int cflags); void kmem_cache_destroy(kmem_cache_t *cp); void *kmem_cache_alloc(kmem_cache_t *cp, int kmflag); void kmem_cache_free(kmem_cache_t *cp, void *obj); [Synopsis for callback functions:] int (*constructor)(void *buf, void *un, int kmflags); void (*destructor)(void *buf, void *un);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) The parameters for the kmem_cache_* functions are as follows: <i>name</i> Descriptive name of a kstat(9S) structure of class kmem_cache. Only alphanumeric characters can be used in <i>name</i>. <i>bufsize</i> Size of the objects it manages. <i>align</i> Required object alignment. <i>constructor</i> Pointer to an object constructor function. Parameters are defined below. <i>destructor</i> Pointer to an object destructor function. Parameters are defined below. <i>reclaim</i> Drivers should pass NULL. <i>private</i> Pass-through argument for constructor/destructor. <i>vmp</i> Drivers should pass NULL. <i>cflags</i> Drivers must pass 0.

kmem_cache_create(9F)

kmflag

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.

**obj*

Pointer to the object allocated by `kmem_cache_alloc()`.

The parameters for the callback constructor function are as follows:

`void *buf`

Pointer to the object to be constructed.

`void *un`

The *private* parameter from the call to `kmem_cache_create()`; it is typically a pointer to the soft-state structure.

`int kmflags`

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.

The parameters for the callback destructor function are as follows:

`void *buf`

Pointer to the object to be deconstructed.

`void *un`

The *private* parameter from the call to `kmem_cache_create()`; it is typically a pointer to the soft-state structure.

DESCRIPTION

In many cases, the cost of initializing and destroying an object exceeds the cost of allocating and freeing memory for it. The functions described here address this condition.

Object caching is a technique for dealing with objects that are:

- frequently allocated and freed, and
- have setup and initialization costs.

The idea is to allow the allocator and its clients to cooperate to preserve the invariant portion of an object's initial state, or constructed state, between uses, so it does not have to be destroyed and re-created every time the object is used. For example, an object containing a mutex only needs to have `mutex_init()` applied once, the first time the object is allocated. The object can then be freed and reallocated many times without incurring the expense of `mutex_destroy()` and `mutex_init()` each time. An object's embedded locks, condition variables, reference counts, lists of other objects, and read-only data all generally qualify as constructed state. The essential requirement is that the client must free the object (using `kmem_cache_free()`) in its constructed state. The allocator cannot enforce this, so programming errors will lead to hard-to-find bugs.

kmem_cache_create(9F)

A driver should call `kmem_cache_create()` at the time of `_init(9E)` or `attach(9E)`, and call the corresponding `kmem_cache_destroy()` at the time of `_fini(9E)` or `detach(9E)`.

`kmem_cache_create()` creates a cache of objects, each of size *size* bytes, aligned on an *align* boundary. Drivers not requiring a specific alignment can pass 0. *name* identifies the cache for statistics and debugging. *constructor* and *destructor* convert plain memory into objects and back again; *constructor* can fail if it needs to allocate memory but cannot. *private* is a parameter passed to the constructor and destructor callbacks to support parameterized caches (for example, a pointer to an instance of the driver's soft-state structure). To facilitate debugging, `kmem_cache_create()` creates a `kstat(9S)` structure of class `kmem_cache` and name *name*. It returns an opaque pointer to the object cache.

`kmem_cache_alloc()` gets an object from the cache. The object will be in its constructed state. *kmflag* is either `KM_SLEEP` or `KM_NOSLEEP`, indicating whether it is acceptable to wait for memory if none is currently available.

`kmem_cache_free()` returns an object to the cache. The object must be in its constructed state.

`kmem_cache_destroy()` destroys the cache and releases all associated resources. All allocated objects must have been previously freed.

CONTEXT Constructors can be invoked during any call to `kmem_cache_alloc()`, and will run in that context. Similarly, destructors can be invoked during any call to `kmem_cache_free()`, and can also be invoked during `kmem_cache_destroy()`. Therefore, the functions that a constructor or destructor invokes must be appropriate in that context.

`kmem_cache_create()` and `kmem_cache_destroy()` must not be called from interrupt context.

`kmem_cache_alloc()` can be called from interrupt context only if the `KM_NOSLEEP` flag is set. It can be called from user or kernel context with any valid flag.

`kmem_cache_free()` can be called from user, kernel, or interrupt context.

EXAMPLES

EXAMPLE 1 Object Caching

Consider the following data structure:

```
struct foo {
    kmutex_t foo_lock;
    kcondvar_t foo_cv;
    struct bar *foo_barlist;
    int foo_refcnt;
};
```

EXAMPLE 1 Object Caching (Continued)

Assume that a `foo` structure cannot be freed until there are no outstanding references to it (`foo_refcnt == 0`) and all of its pending bar events (whatever they are) have completed (`foo_barlist == NULL`). The life cycle of a dynamically allocated `foo` would be something like this:

```
foo = kmem_alloc(sizeof (struct foo), KM_SLEEP);
mutex_init(&foo->foo_lock, ...);
cv_init(&foo->foo_cv, ...);
foo->foo_refcnt = 0;
foo->foo_barlist = NULL;
    use foo;
ASSERT(foo->foo_barlist == NULL);
ASSERT(foo->foo_refcnt == 0);
cv_destroy(&foo->foo_cv);
mutex_destroy(&foo->foo_lock);
kmem_free(foo);
```

Notice that between each use of a `foo` object we perform a sequence of operations that constitutes nothing more expensive overhead. All of this overhead (that is, everything other than use `foo` above) can be eliminated by object caching.

```
int
foo_constructor(void *buf, void *arg, int size)
{
    struct foo *foo = buf;
    mutex_init(&foo->foo_lock, ...);
    cv_init(&foo->foo_cv, ...);
    foo->foo_refcnt = 0;
    foo->foo_barlist = NULL;
}

void
foo_destructor(void *buf, void *arg, int size)
{
    struct foo *foo = buf;
    ASSERT(foo->foo_barlist == NULL);
    ASSERT(foo->foo_refcnt == 0);
    cv_destroy(&foo->foo_cv);
    mutex_destroy(&foo->foo_lock);
}

un = ddi_get_soft_state(foo_softc, instance);
(void) snprintf(buf, KSTAT_STRLEN, "foo%d_cache",
    ddi_get_instance(dip));
foo_cache = kmem_cache_create(buf,
    sizeof (struct foo), 0,
    foo_constructor, foo_destructor,
    NULL, un, 0);
```

To allocate, use, and free a `foo` object:

```
foo = kmem_cache_alloc(foo_cache, KM_SLEEP);
    use foo;
kmem_cache_free(foo_cache, foo);
```

EXAMPLE 1 Object Caching (Continued)

This makes `foo` allocation fast, because the allocator will usually do nothing more than fetch an already-constructed `foo` from the cache. `foo_constructor` and `foo_destructor` will be invoked only to populate and drain the cache, respectively.

RETURN VALUES

If successful, the constructor function must return a pointer to the allocated object. If `KM_NOSLEEP` is set and memory cannot be allocated without sleeping, the constructor must return `NULL`.

`kmem_cache_create()` returns a pointer to the allocated cache. If the `name` parameter contains non-alphanumeric characters, `kmem_cache_create()` returns `NULL`.

If successful, `kmem_cache_alloc()` returns a pointer to the allocated object. If `KM_NOSLEEP` is set and memory cannot be allocated without sleeping, `kmem_cache_alloc()` returns `NULL`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`condvar(9F)`, `kmem_alloc(9F)`, `mutex(9F)`, `kstat(9S)`

Writing Device Drivers

The Slab Allocator: An Object-Caching Kernel Memory Allocator, Bonwick, J.; USENIX Summer 1994 Technical Conference (1994).

kstat_create(9F)

NAME	kstat_create – create and initialize a new kstat														
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> kstat_t *kstat_create(char *module, int instance, char *name, char *class, uchar_t type, ulong_t ndata, uchar_t ks_flag);</pre>														
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <table> <tr> <td><i>module</i></td><td>The name of the provider's module (such as "sd", "esp", ...). The "core" kernel uses the name "unix".</td></tr> <tr> <td><i>instance</i></td><td>The provider's instance number, as from ddi_get_instance(9F). Modules which do not have a meaningful instance number should use 0.</td></tr> <tr> <td><i>name</i></td><td>A pointer to a string that uniquely identifies this structure. Only KSTAT_STRLEN – 1 characters are significant.</td></tr> <tr> <td><i>class</i></td><td>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.</td></tr> <tr> <td><i>type</i></td><td> The type of kstat to allocate. Valid types are: KSTAT_TYPE_NAMED Allows more than one data record per kstat. KSTAT_TYPE_INTR Interrupt; only one data record per kstat. KSTAT_TYPE_IO I/O; only one data record per kstat </td></tr> <tr> <td><i>ndata</i></td><td>The number of type-specific data records to allocate.</td></tr> <tr> <td><i>flag</i></td><td> A bit-field of various flags for this kstat. <i>flag</i> is some combination of: KSTAT_FLAG_VIRTUAL Tells kstat_create() not to allocate memory for the kstat data section; instead, the driver will set the ks_data field to point to the data it wishes to export. This provides a convenient way to export existing data structures. KSTAT_FLAG_WRITABLE Makes the kstat data section writable by root. KSTAT_FLAG_PERSISTENT Indicates that this kstat is to be persistent over time. For persistent kstats, kstat_delete(9F) simply marks the kstat as dormant; a subsequent kstat_create() reactivates the kstat. This feature is provided so that statistics are not lost across driver close/open (such as raw disk I/O on a disk with </td></tr> </table>	<i>module</i>	The name of the provider's module (such as "sd", "esp", ...). The "core" kernel uses the name "unix".	<i>instance</i>	The provider's instance number, as from ddi_get_instance(9F). Modules which do not have a meaningful instance number should use 0.	<i>name</i>	A pointer to a string that uniquely identifies this structure. Only KSTAT_STRLEN – 1 characters are significant.	<i>class</i>	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.	<i>type</i>	The type of kstat to allocate. Valid types are: KSTAT_TYPE_NAMED Allows more than one data record per kstat. KSTAT_TYPE_INTR Interrupt; only one data record per kstat. KSTAT_TYPE_IO I/O; only one data record per kstat	<i>ndata</i>	The number of type-specific data records to allocate.	<i>flag</i>	A bit-field of various flags for this kstat. <i>flag</i> is some combination of: KSTAT_FLAG_VIRTUAL Tells kstat_create() not to allocate memory for the kstat data section; instead, the driver will set the ks_data field to point to the data it wishes to export. This provides a convenient way to export existing data structures. KSTAT_FLAG_WRITABLE Makes the kstat data section writable by root. KSTAT_FLAG_PERSISTENT Indicates that this kstat is to be persistent over time. For persistent kstats, kstat_delete(9F) simply marks the kstat as dormant; a subsequent kstat_create() reactivates the kstat. This feature is provided so that statistics are not lost across driver close/open (such as raw disk I/O on a disk with
<i>module</i>	The name of the provider's module (such as "sd", "esp", ...). The "core" kernel uses the name "unix".														
<i>instance</i>	The provider's instance number, as from ddi_get_instance(9F). Modules which do not have a meaningful instance number should use 0.														
<i>name</i>	A pointer to a string that uniquely identifies this structure. Only KSTAT_STRLEN – 1 characters are significant.														
<i>class</i>	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.														
<i>type</i>	The type of kstat to allocate. Valid types are: KSTAT_TYPE_NAMED Allows more than one data record per kstat. KSTAT_TYPE_INTR Interrupt; only one data record per kstat. KSTAT_TYPE_IO I/O; only one data record per kstat														
<i>ndata</i>	The number of type-specific data records to allocate.														
<i>flag</i>	A bit-field of various flags for this kstat. <i>flag</i> is some combination of: KSTAT_FLAG_VIRTUAL Tells kstat_create() not to allocate memory for the kstat data section; instead, the driver will set the ks_data field to point to the data it wishes to export. This provides a convenient way to export existing data structures. KSTAT_FLAG_WRITABLE Makes the kstat data section writable by root. KSTAT_FLAG_PERSISTENT Indicates that this kstat is to be persistent over time. For persistent kstats, kstat_delete(9F) simply marks the kstat as dormant; a subsequent kstat_create() reactivates the kstat. 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 kstats cannot be virtual, since <code>ks_data</code> points to garbage as soon as the driver goes away.
DESCRIPTION	<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: <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. 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. 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.
RETURN VALUES	If successful, <code>kstat_create()</code> returns a pointer to the allocated <code>kstat</code> . <code>NULL</code> is returned upon failure.
CONTEXT	<code>kstat_create()</code> can be called from user or kernel context.
EXAMPLES	EXAMPLE 1 Allocating and Initializing a <code>kstat</code> Structure <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	<code>kstat(3KSTAT)</code>, <code>ddi_get_instance(9F)</code>, <code>kstat_delete(9F)</code>, <code>kstat_install(9F)</code>, <code>kstat_named_init(9F)</code>, <code>kstat(9S)</code>, <code>kstat_named(9S)</code> <i>Writing Device Drivers</i>

kstat_delete(9F)

NAME	kstat_delete – remove a kstat from the system
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> void kstat_delete(kstat_t *ksp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
PARAMETERS	<i>ksp</i> Pointer to a currently installed kstat(9S) structure.
DESCRIPTION	kstat_delete() removes <i>ksp</i> from the kstat chain and frees all associated system resources.
RETURN VALUES	None.
CONTEXT	kstat_delete() can be called from any context.
SEE ALSO	kstat_create(9F), kstat_install(9F), kstat_named_init(9F), kstat(9S) <i>Writing Device Drivers</i>
NOTES	When calling kstat_delete(), the driver must not be holding that kstat's <i>ks_lock</i> . Otherwise, it may deadlock with a kstat reader.

NAME	kstat_install – add a fully initialized kstat to the system
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> void kstat_install(kstat_t *ksp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
PARAMETERS	<i>ksp</i> Pointer to a fully initialized kstat(9S) structure.
DESCRIPTION	kstat_install() is used in conjunction with kstat_create(9F) to allocate and initialize a kstat(9S) structure. After a successful call to kstat_create() the driver must perform any necessary initialization of the data section (such as setting the name fields in a kstat of type KSTAT_TYPE_NAMED). Virtual kstats must have the ks_data field set at this time. The provider may also set the ks_update, ks_private, and ks_lock fields if necessary. Once the kstat is completely initialized, kstat_install is used to make the kstat accessible to the outside world.
RETURN VALUES	None.
CONTEXT	kstat_install() can be called from user or kernel context.
EXAMPLES	EXAMPLE 1 Allocating and Initializing a kstat Structure The method for allocating and initializing a kstat structure is generally as follows: <pre>kstat_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_create(9F), kstat_delete(9F), kstat_named_init(9F), kstat(9S) <i>Writing Device Drivers</i>

kstat_named_init(9F)

NAME	kstat_named_init, kstat_named_setstr – initialize a named kstat
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> void kstat_named_init(kstat_named_t *knp, char *name, uchar_t data_type); void kstat_named_setstr(kstat_named_t *knp, const char *str);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>knp</i> Pointer to a kstat_named(9S) structure. <i>name</i> The name of the statistic. <i>data_type</i> The type of value. This indicates which field of the kstat_named(9S) structure should be used. Valid values are: KSTAT_DATA_CHAR The "char" field. KSTAT_DATA_LONG The "long" field. KSTAT_DATA_ULONG The "unsigned long" field. KSTAT_DATA_LONGLONG The "long long" field. KSTAT_DATA_ULONGLONG The "unsigned long long" field. KSTAT_DATA_STRING Arbitrary length "long string" field. <i>str</i> Pointer to a NULL-terminated string.
DESCRIPTION	kstat_named_init() associates a name and a type with a kstat_named(9S) structure. kstat_named_setstr() associates <i>str</i> with the named kstat knp. It is an error for knp to be of type other than KSTAT_DATA_STRING. This is the only supported method of changing the value of long strings.
RETURN VALUES	None.
CONTEXT	kstat_named_init() and kstat_named_setstr() can be called from user or kernel context.
SEE ALSO	kstat_create(9F), kstat_install(9F), kstat(9S), kstat_named(9S) <i>Writing Device Drivers</i>

NAME	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
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> 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>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>kiop</i> Pointer to a kstat_io(9S) structure.
DESCRIPTION	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. The kstat_queue() family of functions manage these times based on the transitions between the driver wait queue and run queue. kstat_waitq_enter() kstat_waitq_enter() should be called when a request arrives and is placed into a pre-service state (such as just prior to calling disksort(9F)). kstat_waitq_exit() kstat_waitq_exit() should be used when a request is removed from its pre-service state. (such as just prior to calling the driver's start routine). kstat_runq_enter() 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

kstat_queue(9F)

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

NAME	linkb – concatenate two message blocks
SYNOPSIS	<pre>#include <sys/stream.h> void linkb(mblk_t *mp1, mblk_t *mp2);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	linkb() creates a new message by adding <i>mp2</i> to the tail of <i>mp1</i>. The continuation pointer, <i>b_cont</i>, of <i>mp1</i> is set to point to <i>mp2</i>. The following figure describes how the linkb(<i>m1</i>, <i>m2</i>) ; function concatenates two message blocks, <i>mp1</i> and <i>mp2</i>: <pre>graph LR subgraph mp1 [mp1] direction TB b_datap1[b_datap] b_cont1[b_cont] end subgraph mp2 [mp2] direction TB b_datap2[b_datap] b_cont2["b_cont (0)"] end db_base1[db_base] db_base2[db_base] data_buf1[data buffer] data_buf2[data buffer] b_datap1 --> db_base1 db_base1 --> data_buf1 b_cont1 --> b_datap2 b_datap2 --> db_base2 db_base2 --> data_buf2</pre> linkb(<i>mp1</i>, <i>mp2</i>);
PARAMETERS	<i>mp1</i> The message to which <i>mp2</i> is to be added. mblk_t is an instance of the msgb(9S) structure. <i>mp2</i> The message to be added.
CONTEXT	linkb() can be called from user or interrupt context.
EXAMPLES	See dupb(9F) for an example that uses linkb().
SEE ALSO	dupb(9F), unlinkb(9F), msgb(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

makecom(9F)

NAME	makecom, makecom_g0, makecom_g0_s, makecom_g1, makecom_g5 – make a packet for SCSI commands														
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> 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>														
INTERFACE LEVEL PARAMETERS	These interfaces are obsolete. <code>scsi_setup_cdb(9F)</code> should be used instead. <table> <tr> <td><i>pkt</i></td><td>Pointer to an allocated <code>scsi_pkt(9S)</code> structure.</td></tr> <tr> <td><i>devp</i></td><td>Pointer to the target's <code>scsi_device(9S)</code> structure.</td></tr> <tr> <td><i>flag</i></td><td>Flags for the <code>pkt_flags</code> member.</td></tr> <tr> <td><i>cmd</i></td><td>First byte of a group 0 or 1 or 5 SCSI CDB.</td></tr> <tr> <td><i>addr</i></td><td>Pointer to the location of the data.</td></tr> <tr> <td><i>cnt</i></td><td>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.</td></tr> <tr> <td><i>fixbit</i></td><td>Fixed bit in sequential access device commands.</td></tr> </table>	<i>pkt</i>	Pointer to an allocated <code>scsi_pkt(9S)</code> structure.	<i>devp</i>	Pointer to the target's <code>scsi_device(9S)</code> structure.	<i>flag</i>	Flags for the <code>pkt_flags</code> member.	<i>cmd</i>	First byte of a group 0 or 1 or 5 SCSI CDB.	<i>addr</i>	Pointer to the location of the data.	<i>cnt</i>	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.	<i>fixbit</i>	Fixed bit in sequential access device commands.
<i>pkt</i>	Pointer to an allocated <code>scsi_pkt(9S)</code> structure.														
<i>devp</i>	Pointer to the target's <code>scsi_device(9S)</code> structure.														
<i>flag</i>	Flags for the <code>pkt_flags</code> member.														
<i>cmd</i>	First byte of a group 0 or 1 or 5 SCSI CDB.														
<i>addr</i>	Pointer to the location of the data.														
<i>cnt</i>	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.														
<i>fixbit</i>	Fixed bit in sequential access device commands.														
DESCRIPTION	<code>makecom</code> functions initialize a packet with the specified command descriptor block, <i>devp</i> and transport flags. The <code>pkt_address</code>, <code>pkt_flags</code>, and the command descriptor block pointed to by <code>pkt_cdbp</code> are initialized using the remaining arguments. Target drivers may use <code>makecom_g0()</code> for Group 0 commands (except for sequential access devices), or <code>makecom_g0_s()</code> for Group 0 commands for sequential access devices, or <code>makecom_g1()</code> for Group 1 commands, or <code>makecom_g5()</code> for Group 5 commands. <i>fixbit</i> is used by sequential access devices for accessing fixed block sizes and sets the tag portion of the SCSI CDB.														
CONTEXT	These functions can be called from user or interrupt context.														
EXAMPLES	EXAMPLE 1 Using <code>makecom</code> Functions <pre>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); }</pre>														

EXAMPLE 1 Using makecom Functions *(Continued)*

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), scsi_setup_cdb(9F), scsi_device(9S), scsi_pkt(9S)

ANSI Small Computer System Interface-2 (SCSI-2)

Writing Device Drivers

NOTES The makecom_g0(), makecom_g0_s(), makecom_g1(), and makecom_g5() functions are obsolete and will be discontinued in a future release. These functions have been replaced by the scsi_setup_cdb() function. See scsi_setup_cdb(9F).

makedevice(9F)

NAME	makedevice – make device number from major and minor numbers				
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mkdev.h> #include <sys/ddi.h> dev_t makedevice(major_t <i>majnum</i>, minor_t <i>minnum</i>);</pre>				
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).				
PARAMETERS	<table><tr><td><i>majnum</i></td><td>Major device number.</td></tr><tr><td><i>minnum</i></td><td>Minor device number.</td></tr></table>	<i>majnum</i>	Major device number.	<i>minnum</i>	Minor device number.
<i>majnum</i>	Major device number.				
<i>minnum</i>	Minor device number.				
DESCRIPTION	makedevice() creates a device number from a major and minor device number. makedevice() should be used to create device numbers so the driver will port easily to releases that treat device numbers differently.				
RETURN VALUES	The device number, containing both the major number and the minor number, is returned. No validation of the major or minor numbers is performed.				
CONTEXT	makedevice() can be called from user or interrupt context.				
SEE ALSO	getmajor(9F), getminor(9F)				

NAME	max – return the larger of two integers
SYNOPSIS	<pre>#include <sys/ddi.h> int max(int <i>int1</i>, int <i>int2</i>);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>int1</i> The first integer. <i>int2</i> The second integer.
DESCRIPTION	max() compares two signed integers and returns the larger of the two.
RETURN VALUES	The larger of the two numbers.
CONTEXT	max() can be called from user or interrupt context.
SEE ALSO	min(9F) <i>Writing Device Drivers</i>

min(9F)

NAME	min – return the lesser of two integers
SYNOPSIS	<pre>#include <sys/ddi.h> int min(int <i>int1</i>, int <i>int2</i>);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>int1</i> The first integer. <i>int2</i> The second integer.
DESCRIPTION	min() compares two signed integers and returns the lesser of the two.
RETURN VALUES	The lesser of the two integers.
CONTEXT	min() can be called from user or interrupt context.
SEE ALSO	max(9F) <i>Writing Device Drivers</i>

NAME	mkioctl – allocates a STREAMS ioctl block for M_IOCTL messages in the kernel.
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *mkioctl (uint_t command) ;</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>command</i> The ioctl command for the ioc_cmd field.
DESCRIPTION	STREAMS modules or drivers might need to issue an ioctl to a lower module or driver. The mkioctl() function tries to allocate (using allocb(9F)) a STREAMS M_IOCTL message block (iocblk(9S)). Buffer allocation fails only when the system is out of memory. If no buffer is available, the qbufcall(9F) function can help a module recover from an allocation failure. The mkioctl function returns a mblk_t structure which is large enough to hold any of the ioctl messages (iocblk(9S), copyreq(9S) or copyresp(9S)), and has the following special properties: b_wptr Set to b_rptr + sizeof(struct iocblk) . b_cont Set to NULL . b_datap->db_type Set to M_IOCTL. The fields in the iocblk structure are initialized as follows: ioc_cmd Set to the command value passed in. ioc_id Set to a unique identifier. ioc_cr Set to point to a credential structure encoding the maximum system privilege and which does not need to be freed in any fashion. ioc_count Set to 0. ioc_rval Set to 0. ioc_error Set to 0. ioc_flags Set to IOC_NATIVE to reflect that this is native to the running kernel.
RETURN VALUES	Upon success, the mkioctl() function returns a pointer to the allocated mblk_t of type M_IOCTL . On failure, it returns a null pointer.
CONTEXT	The mkioctl() function can be called from user or interrupt context.

mkioctl(9F)

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 = mkioctl(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 = mkioctl(TEST_IOCTL)) == NULL) {
4         int id;
5     }
```

EXAMPLE 3 An iocblk Allocation Which Fails (Continued)

```

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.

mod_install(9F)

NAME	mod_install, mod_remove, mod_info – add, remove or query a loadable module				
SYNOPSIS	<pre>#include <sys/modctl.h> int mod_install(struct modlinkage *modlinkage); int mod_remove(struct modlinkage *modlinkage); int mod_info(struct modlinkage *modlinkage, struct modinfo *modinfo);</pre>				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).				
PARAMETERS	<table><tr><td><i>modlinkage</i></td><td>Pointer to the loadable module's modlinkage structure which describes what type(s) of module elements are included in this loadable module.</td></tr><tr><td><i>modinfo</i></td><td>Pointer to the modinfo structure passed to _info(9E).</td></tr></table>	<i>modlinkage</i>	Pointer to the loadable module's modlinkage structure which describes what type(s) of module elements are included in this loadable module.	<i>modinfo</i>	Pointer to the modinfo structure passed to _info(9E).
<i>modlinkage</i>	Pointer to the loadable module's modlinkage structure which describes what type(s) of module elements are included in this loadable module.				
<i>modinfo</i>	Pointer to the modinfo structure passed to _info(9E).				
DESCRIPTION	mod_install() must be called from a module's _init(9E) routine. mod_remove() must be called from a module's _fini(9E) routine. mod_info() must be called from a module's _info(9E) routine.				
RETURN VALUES	mod_install() and mod_remove() return 0 upon success and non-zero on failure. mod_info() returns a non-zero value on success and 0 upon failure.				
EXAMPLES	See _init(9E) for an example that uses these functions.				
SEE ALSO	_fini(9E), _info(9E), _init(9E), modldrv(9S), modlinkage(9S), modlstrmod(9S) <i>Writing Device Drivers</i>				

NAME	msgdsize – return the number of bytes in a message
SYNOPSIS	<pre>#include <sys/stream.h> size_t msgdsize(mblk_t *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Message to be evaluated.
DESCRIPTION	msgdsize() 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.
RETURN VALUES	The number of data bytes in a message, expressed as an integer.
CONTEXT	msgdsize() can be called from user or interrupt context.
EXAMPLES	See bufcall(9F) for an example that uses msgdsize().
SEE ALSO	bufcall(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

msgpullup(9F)

NAME	msgpullup – concatenate bytes in a message				
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *msgpullup(mblk_t *mp, ssize_t len);</pre>				
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).				
PARAMETERS	<i>mp</i> Pointer to the message whose blocks are to be concatenated. <i>len</i> Number of bytes to concatenate.				
DESCRIPTION	msgpullup() 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, msgpullup() fails and returns NULL.				
RETURN VALUES	msgpullup returns the following values: <table><tr><td>Non-null</td><td>Successful completion. A pointer to the new message is returned.</td></tr><tr><td>NULL</td><td>An error occurred.</td></tr></table>	Non-null	Successful completion. A pointer to the new message is returned.	NULL	An error occurred.
Non-null	Successful completion. A pointer to the new message is returned.				
NULL	An error occurred.				
CONTEXT	msgpullup() can be called from user or interrupt context.				
SEE ALSO	srv(9E), allocb(9F), pullupmsg(9F), msgb(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>				
NOTES	msgpullup() is a DKI-compliant replacement for the older pullupmsg(9F) routine. Users are strongly encouraged to use msgpullup() instead of pullupmsg(9F).				

NAME	mt-streams – STREAMS multithreading	
SYNOPSIS	#include <sys/conf.h>	
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).	
DESCRIPTION	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. 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>). 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). 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. All MT safe modules/drivers specify the <code>D_MP</code> flag.	
Inner Perimeter Flags	The inner perimeter presence and scope are controlled by the mutually exclusive flags:	
	D_MTPERQ	The module/driver has an inner perimeter around each queue.
	D_MTQPAIR	The module/driver has an inner perimeter around each read/write pair of queues.
	D_MTPERMOD	The module/driver has an inner perimeter that encloses all the module's/driver's queues.
	None of the above	The module/driver has no inner perimeter.
Outer Perimeter Flags	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 <code>D_MTPERMOD</code> .
	Note that acquiring exclusive access at the outer perimeter (that is, using <code>qwriter(9F)</code> with the <code>PERIM_OUTER</code> flag) can incur significant performance penalties, which grow linearly with the number of open instances of the module or driver in the system.	

mt-streams(9F)

The default concurrency can be modified using:

<code>D_MTPUTSHARED</code>	This flag modifies the default behavior when <code>put(9E)</code> procedure are invoked so that the inner perimeter is entered shared instead of exclusively.
<code>D_MTOCEXCL</code>	This flag modifies the default behavior when <code>open(9E)</code> and <code>close(9E)</code> procedures are invoked so the the outer perimeter is entered exclusively instead of shared. Note that drivers and modules using this flag can cause significant system performance degradation during stream open or close when many instances of the driver or module are in use simultaneously. For this reason, use of this flag is discouraged. Instead, since <code>open(9E)</code> and <code>close(9E)</code> both execute with user context, developers are encouraged to use traditional synchronization routines such as <code>cv_wait_sig(9F)</code> to coordinate with other open instances of the module or driver.

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

NAME	mutex, mutex_enter, mutex_exit, mutex_init, mutex_destroy, mutex_owned, mutex_tryenter – mutual exclusion lock routines								
SYNOPSIS	<pre>#include <sys/ksynch.h> 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>								
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).								
PARAMETERS	<table> <tr> <td><i>mp</i></td><td>Pointer to a kernel mutex lock (kmutex_t).</td></tr> <tr> <td><i>name</i></td><td>Descriptive string. This is obsolete and should be NULL. (Non-NULL strings are legal, but they are a waste of kernel memory.)</td></tr> <tr> <td><i>type</i></td><td>Type of mutex lock.</td></tr> <tr> <td><i>arg</i></td><td>Type-specific argument for initialization routine.</td></tr> </table>	<i>mp</i>	Pointer to a kernel mutex lock (kmutex_t).	<i>name</i>	Descriptive string. This is obsolete and should be NULL. (Non-NULL strings are legal, but they are a waste of kernel memory.)	<i>type</i>	Type of mutex lock.	<i>arg</i>	Type-specific argument for initialization routine.
<i>mp</i>	Pointer to a kernel mutex lock (kmutex_t).								
<i>name</i>	Descriptive string. This is obsolete and should be NULL. (Non-NULL strings are legal, but they are a waste of kernel memory.)								
<i>type</i>	Type of mutex lock.								
<i>arg</i>	Type-specific argument for initialization routine.								
DESCRIPTION	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. Mutexes are strictly bracketing and may not be recursively locked, meaning that mutexes should be exited in the opposite order they were entered, and cannot be reentered before exiting. <code>mutex_init()</code> initializes a mutex. It is an error to initialize a mutex more than once. The <i>type</i> argument should be set to <code>MUTEX_DRIVER</code>. <i>arg</i> provides type-specific information for a given variant type of mutex. When <code>mutex_init()</code> 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 <code>ddi_get_iblock_cookie(9F)</code> or <code>ddi_get_soft_iblock_cookie(9F)</code>. Note that <i>arg</i> should be the value of the <code>iblock</code> cookie casted to <code>(void *)</code>, not the address of the cookie. The arguments passed to <code>ddi_get_iblock_cookie(9F)</code> and <code>ddi_get_soft_iblock_cookie(9F)</code>, on the other hand, are the addresses of the cookie. If the mutex is never used inside an interrupt handler, the argument should be <code>NULL</code>.								

mutex(9F)

	<code>mutex_enter()</code> 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 ensues. <code>mutex_owned()</code> should only be used in <code>ASSERT()</code> and may be enforced by not being defined unless the preprocessor symbol <code>DEBUG</code> 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 <i>mp</i>. <code>mutex_tryenter()</code> is very similar to <code>mutex_enter()</code> except that it doesn't block when the mutex is already held. <code>mutex_tryenter()</code> returns non-zero when it acquired the mutex and 0 when the mutex is already held. <code>mutex_exit()</code> releases a mutex and will unblock another thread if any are blocked on the mutex. <code>mutex_destroy()</code> releases any resources that might have been allocated by <code>mutex_init()</code>. <code>mutex_destroy()</code> 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 be sure that no other thread attempts to use the mutex.
RETURN VALUES	<code>mutex_tryenter()</code> returns non-zero on success and zero of failure. <code>mutex_owned()</code> returns non-zero if the calling thread currently holds the mutex pointed to by <i>mp</i>, or when that cannot be determined, if any thread holds the mutex. <code>mutex_owned()</code> returns zero.
CONTEXT	These functions can be called from user, kernel, or high-level interrupt context, except for <code>mutex_init()</code> and <code>mutex_destroy()</code>, which can be called from user or kernel context only.
EXAMPLES	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: <pre>ddi_get_iblock_cookie(dip, 0, &iblock); mutex_init(&un->un_lock, NULL, MUTEX_DRIVER, (void *)iblock); ddi_add_intr(dip, 0, NULL, &dev_cookie, xxintr, (caddr_t)un);</pre> EXAMPLE 2 Calling a Routine with a Lock A routine that expects to be called with a certain lock held might have the following <code>ASSERT</code>: <pre>xxstart(struct xxunit *un) { ASSERT(mutex_owned(&un->un_lock)); ... }</pre>

EXAMPLE 2 Calling a Routine with a Lock *(Continued)*

SEE ALSO lockstat(1M), Intro(9F), 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 has no effect. To gather lock statistics, see lockstat(1M).

To write scalable, responsive drivers that do not hang, panic or deadlock the system, follow these guidelines:

Never return from a driver entry point with a mutex held.

Never hold a mutex when calling a service that may block, for example `kmem_alloc(9F)` with `KM_SLEEP` or `delay(9F)`.

Always acquire mutexes in a consistent order. If a critical section acquires mutex A followed by B, and elsewhere in the driver mutex B is acquired before A, the driver can deadlock with one thread holding A and waiting for B and another thread holding B while waiting for A.

Always use a mutex to enforce exclusive access to data, not instruction paths.

Acquiring a lock in user context that is also acquired in interrupt context means that, as long as that lock is held, the driver instance holding the lock is subject to all the rules and limitations of interrupt context.

In most cases, a mutex can and should be acquired and released within the same function.

Liberal use of debugging aids like `ASSERT(mutex_owned(&mutex))` can help find callers of a function which should be holding a mutex but are not. This means you need to test your driver compiled with `DEBUG`.

Do not use a mutex to set driver state. However, you should use a mutex to protect driver state data.

Use per-instance and automatic data where possible to reduce the amount of shared data. Per-instance data can be protected by a per-instance lock to improve scalability and reduce contention with multiple hardware instances.

Avoid global data and global mutexes whenever possible.

nochpoll(9F)

NAME	nochpoll – error return function for non-pollable devices										
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int nochpoll(dev_t <i>dev</i>, short <i>events</i>, int <i>anyyet</i>, short *<i>reventsp</i>, struct pollhead **<i>pollhdrp</i>);</pre>										
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table><tr><td><i>dev</i></td><td>Device number.</td></tr><tr><td><i>events</i></td><td>Event flags.</td></tr><tr><td><i>anyyet</i></td><td>Check current events only.</td></tr><tr><td><i>reventsp</i></td><td>Event flag pointer.</td></tr><tr><td><i>pollhdrp</i></td><td>Poll head pointer.</td></tr></table>	<i>dev</i>	Device number.	<i>events</i>	Event flags.	<i>anyyet</i>	Check current events only.	<i>reventsp</i>	Event flag pointer.	<i>pollhdrp</i>	Poll head pointer.
<i>dev</i>	Device number.										
<i>events</i>	Event flags.										
<i>anyyet</i>	Check current events only.										
<i>reventsp</i>	Event flag pointer.										
<i>pollhdrp</i>	Poll head pointer.										
DESCRIPTION	<code>nochpoll()</code> 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.										
RETURN VALUES	<code>nochpoll()</code> returns <code>ENXIO</code>.										
CONTEXT	<code>nochpoll()</code> can be called from user or interrupt context.										
SEE ALSO	<code>poll(2)</code>, <code>chpoll(9E)</code>, <code>cb_ops(9S)</code> <i>Writing Device Drivers</i>										

NAME	nodev – error return function
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> int nodev();</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	nodev() returns ENXIO. It is intended to be used in the cb_ops(9S) 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.
RETURN VALUES	nodev() returns ENXIO.
CONTEXT	nodev() can be only called from user context.
SEE ALSO	nulldev(9F), cb_ops(9S) <i>Writing Device Drivers</i>

noenable(9F)

NAME	noenable – prevent a queue from being scheduled
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void noenable(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue.
DESCRIPTION	<code>noenable()</code> prevents the queue <i>q</i> from being scheduled for service by <code>insq(9F)</code> , <code>putq(9F)</code> or <code>putbq(9F)</code> when enqueueing an ordinary priority message. The queue can be re-enabled with the <code>enableok(9F)</code> function.
CONTEXT	<code>noenable()</code> can be called from user or interrupt context.
SEE ALSO	<code>enableok(9F)</code> , <code>insq(9F)</code> , <code>putbq(9F)</code> , <code>putq(9F)</code> , <code>qenable(9F)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	nulldev – zero return function
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> int nulldev();</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	nulldev() returns 0. It is intended to be used in the cb_ops(9S) data structure of a device driver for device entry points that do nothing.
RETURN VALUES	nulldev() returns a 0.
CONTEXT	nulldev() can be called from any context.
SEE ALSO	nodev(9F), cb_ops(9S) <i>Writing Device Drivers</i>

nvlist_add_boolean(9F)

NAME	nvlist_add_boolean, nvlist_add_byte, nvlist_add_int16, nvlist_add_uint16, nvlist_add_int32, nvlist_add_uint32, nvlist_add_int64, nvlist_add_uint64, nvlist_add_string, nvlist_add_nvlist, nvlist_add_byte_array, nvlist_add_int16_array, nvlist_add_uint16_array, nvlist_add_int32_array, nvlist_add_uint32_array, nvlist_add_int64_array, nvlist_add_uint64_array, nvlist_add_string_array – value pair functions
SYNOPSIS	<pre>#include <sys/nvpair.h> int nvlist_add_boolean(nvlist_t *nvl, char *name); int nvlist_add_byte(nvlist_t *nvl, char *name, uchar_t val); int nvlist_add_int16(nvlist_t *nvl, char *name, int16_t val); int nvlist_add_uint16(nvlist_t *nvl, char *name, uint16_t val); int nvlist_add_int32(nvlist_t *nvl, char *name, int32_t val); int nvlist_add_uint32(nvlist_t *nvl, char *name, uint32_t val); int nvlist_add_int64(nvlist_t *nvl, char *name, int64_t val); int nvlist_add_uint64(nvlist_t *nvl, char *name, uint64_t val); int nvlist_add_string(nvlist_t *nvl, char *name, char *val); int nvlist_add_nvlist(nvlist_t *nvl, char *name, nvlist_t *val); int nvlist_add_byte_array(nvlist_t *nvl, char *name, uchar_t *val, uint_t nelem); int nvlist_add_int16_array(nvlist_t *nvl, char *name, int16_t *val, uint_t nelem); int nvlist_add_uint16_array(nvlist_t *nvl, char *name, uint16_t *val, uint_t nelem); int nvlist_add_int32_array(nvlist_t *nvl, char *name, int32_t *val, uint_t nelem); int nvlist_add_uint32_array(nvlist_t *nvl, char *name, uint32_t *val, uint_t nelem); int nvlist_add_int64_array(nvlist_t *nvl, char *name, int64_t *val, uint_t nelem); int nvlist_add_uint64_array(nvlist_t *nvl, char *name, uint64_t *val, uint_t nelem); int nvlist_add_string_array(nvlist_t *nvl, char *name, char **val, uint_t nelem); int nvlist_add_nvlist_array(nvlist_t *nvl, char *name, nvlist_t **val, uint_t nelem);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)

PARAMETERS	<i>nvlist</i>	The <code>nvlist_t</code> to be processed.
	<i>nvp</i>	The <code>nvpair_t</code> (name-value pair) to be processed.
	<i>name</i>	Name of the name-value pair (<code>nvpair</code>).
	<i>nelem</i>	Number of elements in value (that is, array size).
	<i>val</i>	Value or starting address of the array value.
DESCRIPTION	These functions add a new name-value pair to an <code>nvlist_t</code> . The memory allocation policy follows that specified in <code>nvlist_alloc()</code> , <code>nvlist_unpack()</code> , or <code>nvlist_dup()</code> . See <code>nvlist_alloc(9F)</code> . The uniqueness of <code>nvpair</code> name and data types follows the <i>nvflag</i> argument specified in <code>nvlist_alloc()</code> .	
	If <code>NV_UNIQUE_NAME</code> was specified for <i>nvflag</i> , existing <code>nvpairs</code> with matching names are removed before the new <code>nvpair</code> is added.	
	If <code>NV_UNIQUE_NAME_TYPE</code> was specified for <i>nvflag</i> , existing <code>nvpairs</code> with matching names and data types are removed before the new <code>nvpair</code> is added.	
	If neither was specified for <i>nvflag</i> , the new <code>nvpair</code> is unconditionally added at the end of the list. The library preserves the order of the name-value pairs across packing, unpacking, and duplication.	
RETURN VALUES	0	success
	<code>EINVAL</code>	invalid argument
	<code>ENOMEM</code>	insufficient memory
CONTEXT	These functions can be called from interrupt context only if the <code>nvlist_t</code> was allocated with the <code>KM_NOSLEEP</code> flag set. See <code>nvlist_alloc(9F)</code> for a description of <code>KM_NOSLEEP</code> . These functions can be called from user context in all cases.	

nvlist_alloc(9F)

NAME	nvlist_alloc, nvlist_free, nvlist_size, nvlist_pack, nvlist_unpack, nvlist_dup – manage a name-value pair list
SYNOPSIS	<pre>#include <sys/nvpair.h> int nvlist_alloc(nvlist_t **nvp, uint_t nvflag, int kmflag); void nvlist_free(nvlist_t *nvl); int nvlist_size(nvlist_t *nvl, size_t *size, int encoding); int nvlist_pack(nvlist_t *nvl, char **bufp, size_t *buflen, int encoding, int kmflag); int nvlist_unpack(char *buf, size_t buflen, nvlist_t **nvp, int kmflag); int nvlist_dup(nvlist_t *nvl, nvlist_t **nvp, int kmflag);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>nvp</i> Address of a pointer to list of name-value pairs (nvlist_t). <i>nvflag</i> Specify bit fields defining nvlist_t properties: NV_UNIQUE_NAME The nvpair names are unique. NV_UNIQUE_NAME_TYPE Name-data type combination is unique <i>kmflag</i> Kernel memory allocation policy, either KM_SLEEP or KM_NOSLEEP. <i>nvl</i> The nvlist_t to be processed. <i>size</i> Pointer to buffer to contain the encoded size. <i>bufp</i> Address of buffer to pack nvlist into. Must be 8-byte aligned. If NULL, library will allocate memory. <i>buf</i> Buffer containing packed nvlist_t. <i>buflen</i> Size of buffer <i>bufp</i> or <i>buf</i> points to. <i>encoding</i> Encoding method for packing.
DESCRIPTION	The nvlist_alloc() function allocates a new name-value pair list and updates <i>nvp</i> to point to the handle. The argument <i>nvflag</i> specifies nvlist_t properties to remain persistent across packing, unpacking, and duplication. The nvlist_free() function frees a name-value pair list. The nvlist_size() function returns the minimum size of a contiguous buffer large enough to pack <i>nvl</i>. The <i>encoding</i> parameter specifies the method of encoding when packing <i>nvl</i>. Supported encoding methods are:

		nvlist_alloc(9F)																				
	NV_ENCODE_NATIVE	Straight bcopy () as described in bcopy(9F).																				
	NV_ENCODE_XDR	Use XDR encoding, suitable for sending to another host.																				
	The nvlist_pack () function packs <i>nvlist</i> into contiguous memory starting at <i>*bufp</i>. The <i>encoding</i> parameter specifies the method of encoding (see above). ■ If <i>*bufp</i> is not NULL, <i>*bufp</i> is expected to be a caller-allocated buffer of size <i>*buflen</i>. The <i>kmflag</i> argument is ignored. ■ If <i>*bufp</i> is NULL, the library will allocate memory and update <i>*bufp</i> to point to the memory and update <i>*buflen</i> to contain the size of the allocated memory. The value of <i>kmflag</i> indicates the memory allocation policy The nvlist_unpack () function takes a buffer with a packed nvlist_t and unpacks it into a searchable nvlist_t. The library allocates memory for nvlist_t. The caller is responsible for freeing the memory by calling nvlist_free (). The nvlist_dup () function makes a copy of <i>nvlist</i> and updates <i>nvlistp</i> to point to the copy.																					
RETURN VALUES	For nvlist_alloc (), nvlist_dup (): <table><tr><td>0</td><td>success</td></tr><tr><td>EINVAL</td><td>invalid argument</td></tr><tr><td>ENOMEM</td><td>insufficient memory</td></tr></table> For nvlist_pack (), nvlist_unpack (): <table><tr><td>0</td><td>success</td></tr><tr><td>EINVAL</td><td>invalid argument</td></tr><tr><td>ENOMEM</td><td>insufficient memory</td></tr><tr><td>EFAULT</td><td>encode/decode error</td></tr><tr><td>ENOTSUP</td><td>encode/decode method not supported</td></tr></table> For nvlist_size (): <table><tr><td>0</td><td>success</td></tr><tr><td>EINVAL</td><td>invalid argument</td></tr></table>		0	success	EINVAL	invalid argument	ENOMEM	insufficient memory	0	success	EINVAL	invalid argument	ENOMEM	insufficient memory	EFAULT	encode/decode error	ENOTSUP	encode/decode method not supported	0	success	EINVAL	invalid argument
0	success																					
EINVAL	invalid argument																					
ENOMEM	insufficient memory																					
0	success																					
EINVAL	invalid argument																					
ENOMEM	insufficient memory																					
EFAULT	encode/decode error																					
ENOTSUP	encode/decode method not supported																					
0	success																					
EINVAL	invalid argument																					
CONTEXT	The nvlist_alloc (), nvlist_pack (), nvlist_unpack (), and nvlist_dup () functions can be called from interrupt context only if the KM_NOSLEEP flag is set. They can be called from user context with any valid flag.																					

nvlist_lookup_boolean(9F)

NAME	nvlist_lookup_boolean, nvlist_lookup_byte, nvlist_lookup_int16, nvlist_lookup_int32, nvlist_lookup_int64, nvlist_lookup_uint16, nvlist_lookup_uint32, nvlist_lookup_uint64, nvlist_lookup_byte_array, nvlist_lookup_uint16_array, nvlist_lookup_uint32_array, nvlist_lookup_uint64_array, nvlist_lookup_nvlist, nvlist_lookup_byte_array, nvlist_lookup_string_array, nvlist_lookup_string, nvlist_lookup_int16_array, nvlist_lookup_int32_array, nvlist_lookup_int64_array, nvlist_lookup_nvlist_array – match name and type indicated by the interface name and retrieve data value
SYNOPSIS	<pre>#include <sys/nvpair.h> int nvlist_lookup_boolean(nvlist_t *nvl, char *name); int nvlist_lookup_byte(nvlist_t *nvl, char *name, uchar_t *val); int nvlist_lookup_int16(nvlist_t *nvl, char *name, int16_t *val); int nvlist_lookup_uint16(nvlist_t *nvl, char *name, uint16_t *val); int nvlist_lookup_int32(nvlist_t *nvl, char *name, int32_t *val); int nvlist_lookup_uint32(nvlist_t *nvl, char *name, uint32_t *val); int nvlist_lookup_int64(nvlist_t *nvl, char *name, int64_t *val); int nvlist_lookup_uint64(nvlist_t *nvl, char *name, uint64_t *val); int nvlist_lookup_string(nvlist_t *nvl, char *name, char **val); int nvlist_lookup_nvlist(nvlist_t *nvl, char *name, nvlist_t **val); int nvlist_lookup_byte_array(nvlist_t *nvl, char *name, uchar_t **val, uint_t *nelem); int nvlist_lookup_int16_array(nvlist_t *nvl, char *name, int16_t **val, uint_t *nelem); int nvlist_lookup_uint16_array(nvlist_t *nvl, char *name, uint16_t **val, uint_t *nelem); int nvlist_lookup_int32_array(nvlist_t *nvl, char *name, int32_t **val, uint_t *nelem); int nvlist_lookup_uint32_array(nvlist_t *nvl, char *name, uint32_t **val, uint_t *nelem); int nvlist_lookup_int64_array(nvlist_t *nvl, char *name, int64_t **val, uint_t *nelem); int nvlist_lookup_uint64_array(nvlist_t *nvl, char *name, uint64_t **val, uint_t *nelem); int nvlist_lookup_string_array(nvlist_t *nvl, char *name, char ***val, uint_t *nelem);</pre>

		nvlist_lookup_boolean(9F)
		<pre>int nvlist_lookup_nvlist_array(nvlist_t *nvl, char *name, nvlist_t ***val, uint_t *nelem);</pre>
INTERFACE	Solaris DDI specific (Solaris DDI)	
LEVEL		
PARAMETERS	<i>nvl</i>	The list of name-value pairs (nvlist_t) to be processed.
	<i>name</i>	Name of the name-value pair (nvpair) to search.
	<i>nelem</i>	Address to store the number of elements in value.
	<i>val</i>	Address to store the value or starting address of the array value.
DESCRIPTION	These functions find the nvpair that matches the name and type as indicated by the interface name. If one is found, <i>nelem</i> and <i>val</i> are modified to contain the number of elements in value and the starting address of data, respectively. These interfaces work for nvlist_ts allocated with NV_UNIQUE_NAME or NV_UNIQUE_NAME_TYPE specified in nvlist_alloc(). (See nvlist_alloc(9F).) If this is not the case, the interface will return ENOTSUP because the list potentially contains multiple nvpairs with the same name and type. All memory required for storing the array elements, including string values, are managed by the library. References to such data remain valid until nvlist_free() is called on <i>nvl</i>.	
RETURN VALUES	0	success
	EINVAL	invalid argument
	ENOENT	no matching name-value pair found
	ENOTSUP	encode/decode method not supported
CONTEXT	These functions can be called from user or interrupt contexts.	

nvlist_next_nvpair(9F)

NAME	nvlist_next_nvpair, nvpair_name, nvpair_type – return data regarding name-value pairs
SYNOPSIS	<pre>#include <sys/nvpair.h> nvpair_t *nvlist_next_nvpair(nvlist_t *nvl, nvpair_t *nvpair); char *nvpair_name(nvpair_t *nvpair); data_type_t nvpair_type(nvpair_t *nvpair);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
PARAMETERS	<i>nvl</i> The list of name-value pairs (nvlist_t) to be processed. <i>nvpair</i> Handle to a name-value pair.
DESCRIPTION	The nvlist_next_nvpair() function returns a handle to the next name-value pair (nvpair) in the list following <i>nvpair</i>. If <i>nvpair</i> is NULL, the first pair is returned. If <i>nvpair</i> is the last pair in the nvlist_t, NULL is returned. The nvpair_name() function returns a string containing the name of <i>nvpair</i>. The nvpair_type() function retrieves the value of the <i>nvpair</i> in the form of enumerated type data_type_t. This is used to determine the appropriate nvpair_*() function to call for retrieving the value.
RETURN VALUES	For nvpair_name(), a string containing the name. For nvpair_type(), an enumerated data type data_type_t. Possible values for data_type_t are as follows: <pre>DATA_TYPE_BOOLEAN DATA_TYPE_BYTE DATA_TYPE_INT16 DATA_TYPE_UINT16 DATA_TYPE_INT32 DATA_TYPE_UINT32 DATA_TYPE_INT64 DATA_TYPE_UINT64 DATA_TYPE_STRING DATA_TYPE_BYTE_ARRAY DATA_TYPE_INT16_ARRAY DATA_TYPE_UINT16_ARRAY DATA_TYPE_INT32_ARRAY DATA_TYPE_UINT32_ARRAY DATA_TYPE_INT64_ARRAY DATA_TYPE_UINT64_ARRAY DATA_TYPE_STRING_ARRAY</pre> For nvlist_next_pair(): NULL Reached end of list. otherwise: Handle to next nvpair in the list.

`nvlist_next_nvpair(9F)`

CONTEXT	The functions described here can be called from user or interrupt context.

nvlist_remove(9F)

NAME	nvlist_remove, nvlist_remove_all – remove name-value pairs
SYNOPSIS	<pre>#include <sys/nvpair.h> void nvlist_remove(nvlist_t *nvl, char *name, data_type_t type); void nvlist_remove_all(nvlist_t *nvl, char *name);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>nvl</i> The list of name-value pairs (nvlist_t) to be processed. <i>name</i> Name of the name-value pair (nvpair) to be removed. <i>type</i> Data type of the nvpair to be removed.
DESCRIPTION	The nvlist_remove() function removes the first occurrence of <i>nvpair</i> that matches the name and the type. The nvlist_remove_all() function removes all occurrences of <i>nvpair</i> that match the name, regardless of type.
RETURN VALUES	None
CONTEXT	The nvlist_remove() and nvlist_remove_all() functions can be called from user or interrupt context.

nvpair_value_byte(9F)

NAME	nvpair_value_byte, nvpair_value_nvlist, nvpair_value_nvlist_array, nvpair_value_int16, nvpair_value_uint16, nvpair_value_int32, nvpair_value_uint32, nvpair_value_int64, nvpair_value_uint64, nvpair_value_string, nvpair_value_byte_array, nvpair_value_int16_array, nvpair_value_uint16_array, nvpair_value_int32_array, nvpair_value_uint32_array, nvpair_value_int64_array, nvpair_value_uint64_array, nvpair_value_string_array – retrieve value from a name-value pair
SYNOPSIS	<pre>#include <sys/nvpair.h> int nvpair_value_byte(nvpair_t *nvpair, uchar_t *val); int nvpair_value_int16(nvpair_t *nvpair, int16_t *val); int nvpair_value_uint16(nvpair_t *nvpair, uint16_t *val); int nvpair_value_int32(nvpair_t *nvpair, int32_t *val); int nvpair_value_uint32(nvpair_t *nvpair, uint32_t *val); int nvpair_value_int64(nvpair_t *nvpair, int64_t *val); int nvpair_value_uint64(nvpair_t *nvpair, uint64_t *val); int nvpair_value_string(nvpair_t *nvpair, char **val); int nvpair_value_nvlist(nvpair_t *nvpair, nvlist_t **val); int nvpair_value_byte_array(nvpair_t *nvpair, uchar_t **val, uint_t *nelem); int nvpair_value_int16_array(nvpair_t *nvpair, int16_t **val, uint_t *nelem); int nvpair_value_uint16_array(nvpair_t *nvpair, uint16_t **val, uint_t *nelem); int nvpair_value_int32_array(nvpair_t *nvpair, int32_t **val, uint_t *nelem); int nvpair_value_uint32_array(nvpair_t *nvpair, uint32_t **val, uint_t *nelem); int nvpair_value_int64_array(nvpair_t *nvpair, int64_t **val, uint_t *nelem); int nvpair_value_uint64_array(nvpair_t *nvpair, uint64_t **val, uint_t *nelem); int nvpair_value_string_array(nvpair_t *nvpair, char ***val, uint_t *nelem); int nvpair_value_nvlist_array(nvpair_t *nvpair, nvlist_t ***val, uint_t *nelem);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)

`nvpair_value_byte(9F)`

PARAMETERS	<i>nvpair</i>	Name-value pair (<i>nvpair</i>) to be processed.
	<i>nelem</i>	Address to store the number of elements in value.
	<i>val</i>	Address to store the value or starting address of array value.
DESCRIPTION	These functions retrieve the value of <i>nvpair</i> . The data type of <i>nvpair</i> must match the function name for the call to be successful.	
	There is no <code>nvpair_value_boolean()</code> ; the existence of the name implies the value is true.	
	For array data types, including string, the memory containing the data is managed by the library and references to the value remains valid until <code>nvlist_free()</code> is called on the <code>nvlist_t</code> from which <i>nvpair</i> is obtained. See <code>nvlist_free(9F)</code>	
RETURN VALUES	0	Success
	EINVAL	Either one of the arguments is NULL or type of <i>nvpair</i> does not match the interface name.
CONTEXT	These functions can be called from user or interrupt context.	

NAME	OTHERQ, otherq – get pointer to queue’s partner queue
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> queue_t *OTHERQ(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue.
DESCRIPTION	The OTHERQ() function returns a pointer to the other of the two queue structures that make up a STREAMS module or driver. If <i>q</i> points to the read queue the write queue will be returned, and vice versa.
RETURN VALUES	OTHERQ() returns a pointer to a queue’s partner.
CONTEXT	OTHERQ() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Setting Queues 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. <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 q->q_minpsz = min; 10 q->q_maxpsz = max; 11 q->q_hiwat = hi; 12 q->q_lowat = lo; 13 OTHERQ(q)->q_minpsz = min; 14 OTHERQ(q)->q_maxpsz = max; 15 OTHERQ(q)->q_hiwat = hi; 16 OTHERQ(q)->q_lowat = lo; 17 }</pre>
SEE ALSO	<i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

outb(9F)

NAME	outb, outw, outl, repoutsb, repoutsw, repoutsd – write to an I/O port						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void outb(int port, unsigned char value); void outw(int port, unsigned short value); void outl(int port, unsigned long value); void repoutsb(int port, unsigned char *addr, int count); void repoutsw(int port, unsigned short *addr, int count); void repoutsd(int port, unsigned long *addr, int count);</pre>						
INTERFACE LEVEL	The functions described here are obsolete. For the outb(), outw(), and outl() functions use, respectively, ddi_put8(9F), ddi_put16(9F), and ddi_put32(9F) instead. For repoutsb(), repoutsw(), and repoutsl(), use, respectively, ddi_rep_put8(9F), ddi_rep_put16(9F), and ddi_rep_put32(9F) instead.						
PARAMETERS	<i>port</i> A valid I/O port address. <i>value</i> The data to be written to the I/O port. <i>addr</i> The address of a buffer from which the values will be fetched. <i>count</i> The number of values to be written to the I/O port.						
DESCRIPTION	These routines write data of various sizes to the I/O port with the address specified by <i>port</i>. The outb(), outw(), and outl() functions write 8 bits, 16 bits, and 32 bits of data respectively, writing the data specified by <i>value</i>. The repoutsb(), repoutsw(), and repoutsd() 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.						
CONTEXT	These functions may be called from user or interrupt context.						
ATTRIBUTES	See attributes(5) 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> <tr> <td>Stability Level</td><td>Obsolete</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	x86	Stability Level	Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Architecture	x86						
Stability Level	Obsolete						
SEE ALSO	eisa(4), isa(4), attributes(5), ddi_put8(9F), ddi_put16(9F), ddi_put32(9F), ddi_rep_put8(9F), ddi_rep_put16(9F), ddi_rep_put32(9F), inb(9F)						

outb(9F)

Writing Device Drivers

pci_config_get8(9F)

NAME	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
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>handle</i> The data access handle returned from pci_config_setup(9F). <i>offset</i> Byte offset from the beginning of the PCI Configuration space. <i>value</i> Output data.
DESCRIPTION	These routines read or write a single datum of various sizes from or to the PCI Local Bus Configuration space. The pci_config_get8(), pci_config_get16(), pci_config_get32(), and pci_config_get64() functions read 8 bits, 16 bits, 32 bits, and 64 bits of data, respectively. The pci_config_put8(), pci_config_put16(), pci_config_put32(), and pci_config_put64() 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. 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. 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.

pci_config_get8(9F)

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_teardown(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
pci_config_getb	pci_config_get8
pci_config_getw	pci_config_get16
pci_config_getl	pci_config_get32
pci_config_getll	pci_config_get64
pci_config_putb	pci_config_put8
pci_config_putw	pci_config_put16
pci_config_putl	pci_config_put32
pci_config_putll	pci_config_put64

pci_config_setup(9F)

NAME	pci_config_setup, pci_config_teardown – setup or tear down the resources for enabling accesses to the PCI Local Bus Configuration space				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int pci_config_setup(dev_info_t *dip, ddi_acc_handle_t *handle); void pci_config_teardown(ddi_acc_handle_t *handle);</pre>				
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>dip</i> Pointer to the device's dev_info structure. <i>handle</i> Pointer to a data access handle.				
DESCRIPTION	pci_config_setup() sets up the necessary resources for enabling subsequent data accesses to the PCI Local Bus Configuration space. pci_config_teardown() reclaims and removes those resources represented by the data access handle returned from pci_config_setup().				
RETURN VALUES	pci_config_setup() returns: DDI_SUCCESS Successfully setup the resources. DDI_FAILURE Unable to allocate resources for setup.				
CONTEXT	pci_config_setup() must be called from user or kernel context. pci_config_teardown() can be called from any context.				
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.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes: <table><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				
SEE ALSO	attributes(5) <i>IEEE 1275 PCI Bus Binding</i>				

NAME	pci_report_pmcap – Report Power Management capability of a PCI device						
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int pci_report_pmcap(dev_info_t *dip, int cap, void *arg);</pre>						
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <table> <tr> <td><i>dip</i></td><td>Pointer to the device's dev_info structure</td></tr> <tr> <td><i>cap</i></td><td>Power management capability</td></tr> <tr> <td><i>arg</i></td><td>Argument for the capability</td></tr> </table>	<i>dip</i>	Pointer to the device's dev_info structure	<i>cap</i>	Power management capability	<i>arg</i>	Argument for the capability
<i>dip</i>	Pointer to the device's dev_info structure						
<i>cap</i>	Power management capability						
<i>arg</i>	Argument for the capability						
DESCRIPTION	Some PCI devices provide power management capabilities in addition to those provided by the PCI Power Management Specification. The pci_report_pmcap(9F) function reports those Power Management capabilities of the PCI device to the framework. Framework supports dynamic changing of the capability by allowing pci_report_pmcap(9F) to be called multiple times. Following are the supported capabilities as indicated by the cap: PCI_PM_IDLESPPEED — The PCI_PM_IDLESPPEED value indicates the lowest PCI clock speed that a device can tolerate when idle, and is applicable only to 33 MHz PCI bus. <i>arg</i> represents the lowest possible idle speed in KHz (1 KHz is 1000 Hz). The integer value representing the speed should be cast to (void *) before passing as <i>arg</i> to pci_report_pmcap(9F). The special values of <i>arg</i> are: <table> <tr> <td>PCI_PM_IDLESPPEED_ANY</td><td>The device can tolerate any idle clock speed.</td></tr> <tr> <td>PCI_PM_IDLESPPEED_NONE</td><td>The device cannot tolerate slowing down of PCI clock even when idle.</td></tr> </table> If the driver doesn't make this call, PCI_PM_IDLESPPEED_NONE is assumed. In this case, one offending device can keep the entire bus from being power managed.	PCI_PM_IDLESPPEED_ANY	The device can tolerate any idle clock speed.	PCI_PM_IDLESPPEED_NONE	The device cannot tolerate slowing down of PCI clock even when idle.		
PCI_PM_IDLESPPEED_ANY	The device can tolerate any idle clock speed.						
PCI_PM_IDLESPPEED_NONE	The device cannot tolerate slowing down of PCI clock even when idle.						
RETURN VALUES	The pci_report_pmcap(9F) function returns: <table> <tr> <td>DDI_SUCCESS</td><td>Successful reporting of the capability</td></tr> <tr> <td>DDI_FAILURE</td><td>Failure to report capability because of invalid argument(s)</td></tr> </table>	DDI_SUCCESS	Successful reporting of the capability	DDI_FAILURE	Failure to report capability because of invalid argument(s)		
DDI_SUCCESS	Successful reporting of the capability						
DDI_FAILURE	Failure to report capability because of invalid argument(s)						
CONTEXT	The pci_report_pmcap(9F) function can be called from user, kernel and interrupt context.						
EXAMPLES	1. A device driver knows that the device it controls works with any clock between DC and 33 MHz as specified in <i>Section 4.2.3.1: Clock Specification of the PCI Bus Specification Revision 2.1</i>. The device driver makes the following call from its attach(9E): <pre>if (pci_report_pmcap(dip, PCI_PM_IDLESPPEED, PCI_PM_IDLESPPEED_ANY) != DDI_SUCCESS) cmn_err(CE_WARN, "%s%d: pci_report_pmcap failed\n",</pre>						

pci_report_pmcap(9F)

```
ddi_driver_name(dip), ddi_get_instance(dip));
```

2. A device driver controls a 10/100 Mb Ethernet device which runs the device state machine on the chip from the PCI clock. For the device state machine to receive packets at 100 Mb, the PCI clock cannot drop below 4 MHz. The driver makes the following call whenever it negotiates a 100 Mb Ethernet connection:

```
if (pci_report_pmcap(dip, PCI_PM_IDLESPD, (void *)4000) !=
    DDI_SUCCESS)
    cmn_err(CE_WARN, "%s%d: pci_report_pmcap failed\n",
            ddi_driver_name(dip), ddi_get_instance(dip));
```

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

Writing Device Drivers

PCI Bus Power Management Interface Specification Version 1.1

PCI Bus Specification Revision 2.1

NAME	pci_save_config_regs, pci_restore_config_regs – save and restore the PCI configuration registers
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int pci_save_config_regs(dev_info_t *dip); int pci_restore_config_regs(dev_info_t *dip);</pre>
INTERFACE LEVEL ARGUMENTS	Solaris DDI-specific (Solaris DDI). <i>dip</i> Pointer to the device's dev_info structure.
DESCRIPTION	pci_save_config_regs() saves the current configuration registers on persistent system memory. pci_restore_config_regs() restores configuration registers previously saved by pci_save_config_regs(). pci_save_config_regs() should be called by the driver's power() entry point before powering a device off (to PCI state D3). Likewise, pci_restore_config_regs() should be called after powering a device on (from PCI state D3), but before accessing the device. See power(9E).
RETURN VALUES	pci_save_config_regs() and pci_restore_config_regs() return: DDI_SUCCESS Operation completed successfully. DDI_FAILURE Operation failed to complete successfully.
CONTEXT	Both these functions can be called from user or kernel context.
EXAMPLES	EXAMPLE 1 Invoking the save and restore functions <pre>static int xx_power(dev_info_t *dip, int component, int level) { struct xx *xx; int rval = DDI_SUCCESS; xx = ddi_get_soft_state(xx_softstate, ddi_get_instance(dip)); if (xx == NULL) { return (DDI_FAILURE); } mutex_enter(&xx->x_mutex); switch (level) { case PM_LEVEL_D0: XX_POWER_ON(xx); if (pci_restore_config_regs(dip) == DDI_FAILURE) { /* * appropriate error path handling here */ ... } } }</pre>

pci_save_config_regs(9F)

EXAMPLE 1 Invoking the save and restore functions (Continued)

```
        rval = DDI_FAILURE;
    }
    break;

case PM_LEVEL_D3:
    if (pci_save_config_regs(dip) == DDI_FAILURE) {
        /*
         * appropriate error path handling here
         */
        ...
        rval = DDI_FAILURE;
    }
    else {
        XX_POWER_OFF(xx);
    }
    break;

default:
    rval = DDI_FAILURE;
    break;
}

mutex_exit(&xx->x_mutex);
return (rval);
}
```

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO attributes(5), power(9E)

Writing Device Drivers

PCI Bus Power Management Interface Specification Version 1.1

PCI Bus Specification Revision 2.1

NAME	physio, minphys – perform physical I/O	
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/buf.h> #include <sys/uio.h> int physio(int (*strat) (struct buf *), struct buf *bp, dev_t dev, int rw, void (*mincnt) (struct buf *), struct uio *uio); void minphys(struct buf *bp);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI).	
physio()	<i>strat</i>	Pointer to device strategy routine.
	<i>bp</i>	Pointer to a buf(9S) structure describing the transfer. If <i>bp</i> is set to NULL then physio() allocates one which is automatically released upon completion.
	<i>dev</i>	The device number.
	<i>rw</i>	Read/write flag. This is either B_READ when reading from the device, or B_WRITE when writing to the device.
	<i>mincnt</i>	Routine which bounds the maximum transfer unit size.
	<i>uio</i>	Pointer to the uio structure which describes the user I/O request.
minphys()	<i>bp</i>	Pointer to a buf structure.
DESCRIPTION	physio() performs unbuffered I/O operations between the device <i>dev</i> and the address space described in the <i>uio</i> structure. Prior to the start of the transfer physio() verifies the requested operation is valid by checking the protection of the address space specified in the <i>uio</i> structure. 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. physio() uses biowait(9F) to block until <i>strat()</i> has completed each transfer. Upon completion, or detection of an error, physio() unlocks the pages and returns the error status. 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 <i>bp->b_bcount</i> 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:	

physio(9F)

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.

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Evolving

pm_busy_component(9F)

SEE ALSO | power.conf(4), pm(7D), pm_create_components(9F),
pm_destroy_components(9F), pm_raise_power(9F), pm(9P), pm-components(9P)

Writing Device Drivers

NAME	pm_create_components, pm_destroy_components – Create or destroy power-manageable components				
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int pm_create_components(dev_info_t *dip, int components); void pm_destroy_components(dev_info_t *dip);</pre>				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).				
PARAMETERS	<i>dip</i> Pointer to the device's dev_info structure <i>components</i> Number of components to create				
DESCRIPTION	The pm_create_components() and pm_destroy_components() functions are now obsolete and will be removed in a future release. It is recommended that the driver use pm-components(9) instead. The pm_create_components() function creates power-manageable components for a device. It should be called from the driver's attach(9E) entry point if the device has power-manageable components. The correspondence of components to parts of the physical device controlled by the driver are the responsibility of the driver. The pm_destroy_components() function removes all components from the device. It should be called from the driver's detach(9E) entry point.				
RETURN VALUES	The pm_create_components() function returns: DDI_SUCCESS Components are successfully created. DDI_FAILURE The device already has components.				
CONTEXT	These functions may be called from user or kernel context.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Interface stability</td><td>Obsolete</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface stability	Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface stability	Obsolete				
SEE ALSO	power.conf(4), pm(7D), attach(9E), detach(9E), pm_busy_component(9F), pm_idle_component(9F), pm(9P), pm-components(9P) <i>Writing Device Drivers</i>				

pm_get_normal_power(9F)

NAME	pm_get_normal_power, pm_set_normal_power – get or set a device component's normal power level
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> 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>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
PARAMETERS	pm_get_normal_power() parameters: <i>dip</i> Pointer to the device's <i>dev_info</i> structure <i>component</i> Number of component from which to get normal power level pm_set_normal_power() parameters: <i>dip</i> Pointer to the device's <i>dev_info</i> structure <i>component</i> Number of component for which to set normal power level <i>level</i> Component's new normal power level
DESCRIPTION	The pm_get_normal_power() and pm_set_normal_power() functions are now obsolete and will be removed in a future release. It is recommended that device drivers use new automatic device Power Management interfaces. The pm_get_normal_power() function returns the normal power level of <i>component</i> of the device <i>dip</i>. The pm_set_normal_power() function sets the normal power level of <i>component</i> of the device <i>dip</i> to <i>level</i>. When a device has been power managed 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," the interpretation of the meaning of the power level is entirely up to the driver.
RETURN VALUES	The pm_get_normal_power() function returns: <i>level</i> The normal power level of the specified component (a positive integer). DDI_FAILURE Invalid component number <i>component</i> or the device has no components.
CONTEXT	These functions can be called from user or kernel context.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

pm_get_normal_power(9F)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Obsolete

SEE ALSO `power.conf(4)`, `pm(7D)`, `pm(9P)`, `power(9E)`, `pm_busy_component(9F)`,
`pm_create_components(9F)`, `pm_destroy_components(9F)`,
`pm_idle_component(9F)`

Writing Device Drivers

pm_power_has_changed(9F)

NAME	pm_power_has_changed – Notify Power Management framework of autonomous power level change
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int pm_power_has_changed(dev_info_t *dip, int component, int level);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>dip</i> Pointer to the device dev_info structure <i>component</i> Number of the component that has changed power level <i>level</i> Power level to which the indicated component has changed
DESCRIPTION	The pm_power_has_changed(9) function notifies the Power Management framework that the power level of component of <i>dip</i> has changed to <i>level</i>. Normally power level changes are initiated by the Power Management framework due to device idleness, or through a request to the framework from the driver via pm_raise_power(9F) or pm_lower_power(9F), but some devices may change power levels on their own. For the framework to track the power level of the device under these circumstances, the framework must be notified of autonomous power level changes by a call to pm_power_has_changed(). Because of the asynchronous nature of these events, the Power Management framework might have called power(9E) between the device's autonomous power level change and the driver calling pm_power_has_changed(), or the framework may be in the process of changing the power level when pm_power_has_changed() is called. To handle these situations correctly, the driver should verify that the device is indeed at the level or set the device to the level if it doesn't support inquiring of power levels, before calling pm_power_has_changed(). In addition, the driver should prevent a power(9E) entry point from running in parallel with pm_power_has_changed().
RETURN VALUES	The pm_power_has_changed() function returns: DDI_SUCCESS The power level of component was successfully updated to <i>level</i>. DDI_FAILURE Invalid component <i>component</i> or power level <i>level</i>
CONTEXT	This function can be called from user or kernel context. This function can also be called from interrupt context, providing that it is not the first Power Management function called by the driver.
EXAMPLES	A hypothetical driver might include this code to handle pm_power_has_changed(9): <pre>static int xxusb_intr(struct buf *bp) { ... }</pre>

pm_power_has_changed(9F)

```
/*
 * At this point the device has informed us that it has
 * changed power level on its own. Inform this to framework.
 * We need to take care of the case when framework has
 * already called power() entry point and changed power level
 * before we were able to inform framework of this change.
 *   * Handle this by comparing the informed power level with
 * the actual power level and only doing the call if they
 * are same. In addition, make sure that power() doesn't get
 * run in parallel with this code by holding the mutex.
 */
ASSERT(mutex_owned(&xsp->lock));
if (level_informed == *(xsp->level_reg_addr)) {
    if (pm_power_has_changed(xsp->dip, XXUSB_COMPONENT,
        level_informed) != DDI_SUCCESS) {
        mutex_exit(&xsp->lock);
        return(DDI_INTR_UNCLAIMED);
    }
}

....
}

xxdisk_power(dev_info *dip, int comp, int level)
{
    mutex_enter(&xsp->lock);

    ...

    ...
}
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability level	Evolving

SEE ALSO power.conf(4), pm(7D), power(9E), pm_busy_components(9F), pm_idle_components(9F), pm_raise_power(9F), pm_lower_power(9F), pm(9P), pm-components(9P)

Writing Device Drivers

pm_raise_power(9F)

NAME	pm_raise_power, pm_lower_power – Raise or lower power of components	
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int pm_raise_power(dev_info_t *dip, int component, int level); int pm_lower_power(dev_info_t *dip, int component, int level);</pre>	
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI)	
pm_raise_power	<i>dip</i>	Pointer to the device's dev_info structure
	<i>component</i>	The number of the <i>component</i> for which a power level change is desired
	<i>level</i>	The power level to which the indicated <i>component</i> will be raised
pm_lower_power	<i>dip</i>	Pointer to the device's dev_info structure
	<i>component</i>	The number of the <i>component</i> for which a power level change is desired
	<i>level</i>	The power level to which the indicated <i>component</i> will be lowered
DESCRIPTION	The pm_raise_power(9F) function requests the Power Management framework to raise the power level of <i>component</i> of <i>dip</i> to at least <i>level</i>. The state of the device should be examined before each physical access. The pm_raise_power(9F) function should be called to set a <i>component</i> to the required power level if the operation to be performed requires the <i>component</i> to be at a power level higher than its current power level. When pm_raise_power(9F) returns with success, the <i>component</i> is guaranteed to be at least at the requested power level. All devices that depend on this will be at their full power level. Since the actual device power level may be higher than requested by the driver, the driver should not make any assumption about the absolute power level on successful return from pm_raise_power(9F). The pm_raise_power(9F) function may cause re-entry of the driver power(9E) to raise the power level. Deadlock may result if the driver locks are held across the call to pm_raise_power(9F). The pm_lower_power(9F) function requests the Power Management framework to lower the power level of <i>component</i> of <i>dip</i> to at most <i>level</i>. Normally, transitions to lower power levels are initiated by the Power Management framework based on <i>component</i> idleness. However, when detaching, the driver should also initiate reduced power levels by setting the power level of all device components to their lowest levels. The pm_lower_power(9F) function is intended for this use only, and will return DDI_FAILURE if the driver is not detaching at the time of the call.	

pm_raise_power(9F)

If automatic Power Management is disabled (see dtpower(1M) and power.conf(4)), pm_lower_power(9F) returns DDI_SUCCESS without changing the power level of the component. Otherwise, when pm_lower_power(9F) returns with success, the *component* is guaranteed to be at most at the requested power level. Since the actual device power level may be lower than requested by the driver, the driver should not make any assumption about the absolute power level on successful return from pm_lower_power(9F).

The pm_lower_power(9F) may cause re-entry of the driver power(9E) to lower the power level. Deadlock may result if the driver locks are held across the call to pm_raise_power(9F).

RETURN VALUES

The pm_raise_power(9F) function returns:

DDI_SUCCESS	<i>Component</i> is now at the requested power level or higher.
DDI_FAILURE	<i>Component</i> or <i>level</i> is out of range, or the framework was unable to raise the power level of the component to the requested level.

The pm_lower_power(9F) function returns:

DDI_SUCCESS	<i>Component</i> is now at the requested power level or lower, or automatic Power Management is disabled.
DDI_FAILURE	<i>Component</i> or <i>level</i> is out of range, or the framework was unable to lower the power level of the component to the requested level, or the device is not detaching.

EXAMPLES

A hypothetical disk driver might include this code to handle pm_raise_power(9F):

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

    /*
     * At this point we have determined that we need to raise the
     * power level of the device. Since we have to drop the
     * mutex, we need to take care of case where framework is
     * lowering power at the same time we are raising power.
     * We resolve this by marking the device busy and failing
     * lower power in power() entry point when device is busy.
     */
    ASSERT(mutex_owned(xsp->lock));
    if (xsp->pm_busycnt < 1) {
        /*
         * Component is not already marked busy
         */
        if (pm_busy_component(xsp->dip,
            XXDISK_COMPONENT) != DDI_SUCCESS) {
            bioerror(bp, EIO);
            biodone(bp);
            return (0);
        }
    }
}
```

pm_raise_power(9F)

```
        }
        xsp->pm_busycnt++;
    }
    mutex_exit(xsp->lock);
    if (pm_raise_power(xsp->dip,
        XXDISK_COMPONENT, XXPOWER_SPUN_UP) != DDI_SUCCESS) {
        bioerror(bp,EIO);
        biodone(bp);
        return (0);
    }
    mutex_enter(xsp->lock);

    ....
}

xxdisk_power(dev_info *dip, int comp, int level)
{
    ...

    /*
     * We fail the power() entry point if the device is busy and
     * request is to lower the power level.
     */
    ASSERT(mutex_owned( xsp->lock));
    if (xsp->pm_busycnt >= 1) {
        (level < xsp->cur_level) {
            mutex_exit( xsp->lock);
            return (DDI_FAILURE);
        }
    }

    ...
}
```

CONTEXT These functions can be called from user or kernel context.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Evolving

SEE ALSO power.conf(4), pm(7D), power(9E), pm_busy_component(9F),
pm_idle_component(9F), pm(9P), pm-components(9P)

Writing Device Drivers

NAME	pm_trans_check – Device power cycle advisory check
SYNOPSIS	<pre>#include <sys/sunddi.h> int pm_trans_check(struct pm_trans_data *datap, time_t *intervalp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
PARAMETERS	<i>datap</i> Pointer to a pm_trans_data structure <i>intervalp</i> Pointer to time difference when next power cycle will be advised
DESCRIPTION	The pm_trans_check(9F) function checks if a power-cycle is currently advised based on data in the pm_trans_data structure. This function is provided to prevent damage to devices from excess power cycles; drivers for devices that are sensitive to the number of power cycles should call pm_trans_check(9F) from their power(9E) function before powering-off a device. If pm_trans_check(9F) indicates that the device should not be power cycled, the driver should not attempt to power cycle the device and should fail the call to power(9E) entry point. If pm_trans_check(9F) returns that it is not advised to power cycle the device, it attempts to calculate when the next power cycle is advised, based on the supplied parameters. In such case, <i>intervalp</i> returns the time difference (in seconds) from the current time to when the next power cycle is advised. If the time for the next power cycle cannot be determined, <i>intervalp</i> indicates 0. To avoid excessive calls to the power(9E) entry point during a period when power cycling is not advised, the driver should mark the corresponding device component busy for the <i>intervalp</i> time period (if interval is not 0). Conveniently, the driver can utilize the fact that calls to pm_busy_component(9F) are stacked. If power cycling is not advised, the driver can call pm_busy_component(9F) and issue a timeout(9F) for the <i>intervalp</i> time. The timeout() handler can issue the corresponding pm_idle_component(9F) call. When the format field of pm_trans_data is set to DC SCSI FORMAT, the caller must provide valid data in svc_date[], lifemax, and ncycles. Currently, flag must be set to 0. <pre>struct pm_scsi_cycles { int lifemax; /* lifetime max power cycles */ int ncycles; /* number of cycles so far */ char svc_date[DC SCSI_MFR_LEN]; /* service date YYYYWW */ int flag; /* reserved for future */ }; struct pm_trans_data { int format; /* data format */ union { struct pm_scsi_cycles scsi_cycles; } un; };</pre>
RETURN VALUES	1 Power cycle is advised

pm_trans_check(9F)

0 Power cycle is not advised
-1 Error due to invalid argument.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO power.conf(4), attributes(5), power(9E)

Writing Device Drivers

Using Power Management

NAME	pollwakeupp – inform a process that an event has occurred
SYNOPSIS	<pre>#include <sys/poll.h> void pollwakeupp(struct pollhead *php, short event);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>php</i> Pointer to a pollhead structure. <i>event</i> Event to notify the process about.
DESCRIPTION	pollwakeupp() 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 chpoll(9E) and poll(2) for more detail.
CONTEXT	pollwakeupp() can be called from user or interrupt context.
SEE ALSO	poll(2), chpoll(9E) <i>Writing Device Drivers</i>
NOTES	Driver defined locks should not be held across calls to this function.

proc_signal(9F)

NAME	proc_signal, proc_ref, proc_unref – send a signal to a process														
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> #include <sys/signal.h> void *proc_ref(void); void proc_unref(void *pref); int proc_signal(void *pref, int sig);</pre>														
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).														
PARAMETERS	<i>pref</i> A handle for the process to be signalled. <i>sig</i> Signal number to be sent to the process.														
DESCRIPTION	This set of routines allows a driver to send a signal to a process. The routine <code>proc_ref()</code> 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, <code>proc_unref()</code> should be used to remove it as soon as it is no longer needed. <code>proc_signal()</code> 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: <table> <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> See <code>signal(3HEAD)</code> for more details on the meaning of these signals. If the process has exited at the time the signal was sent, <code>proc_signal()</code> returns an error code; the caller should remove the reference on the process by calling <code>proc_unref()</code>. The driver writer must ensure that for each call made to <code>proc_ref()</code>, there is exactly one corresponding call to <code>proc_unref()</code>.	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	<code>proc_ref()</code> returns the following: <table> <tr><td><i>pref</i></td><td>An opaque handle used to refer to the current process.</td></tr> </table> <code>proc_signal()</code> returns the following: <table> <tr><td>0</td><td>The process existed before the signal was sent.</td></tr> </table>	<i>pref</i>	An opaque handle used to refer to the current process.	0	The process existed before the signal was sent.										
<i>pref</i>	An opaque handle used to refer to the current process.														
0	The process existed before the signal was sent.														

proc_signal(9F)

-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(3HEAD), putnextctl1(9F)

Writing Device Drivers

ptob(9F)

NAME	ptob – convert size in pages to size in bytes
SYNOPSIS	<pre>#include <sys/ddi.h> unsigned long ptob(unsigned long <i>numpages</i>) ;</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>numpages</i> Size in number of pages to convert to size in bytes.
DESCRIPTION	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 <code>ptob(2)</code> returns 4096. <code>ptob(0)</code> returns 0.
RETURN VALUES	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 unsigned long. Rather, the higher order bits will be ignored.
CONTEXT	<code>ptob()</code> can be called from user or interrupt context.
SEE ALSO	<code>btop(9F)</code> , <code>btopr(9F)</code> , <code>ddi_ptob(9F)</code> <i>Writing Device Drivers</i>

NAME	pullupmsg – concatenate bytes in a message
SYNOPSIS	<pre>#include <sys/stream.h> int pullupmsg(mblk_t *mp, ssize_t len);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Pointer to the message whose blocks are to be concatenated. <i>mblk_t</i> is an instance of the <i>msgb(9S)</i> structure. <i>len</i> Number of bytes to concatenate.
DESCRIPTION	<i>pullupmsg()</i> tries to combine multiple data blocks into a single block. <i>pullupmsg()</i> 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, <i>pullupmsg()</i> fails and returns 0.
RETURN VALUES	On success, 1 is returned; on failure, 0 is returned.
CONTEXT	<i>pullupmsg()</i> can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using <i>pullupmsg()</i> This is a driver write <i>srv(9E)</i> (service) routine for a device that does not support scatter/gather DMA. For all <i>M_DATA</i> messages, the data will be transferred to the device with DMA. First, try to pull up the message into one message block with the <i>pullupmsg()</i> 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. <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->b_datap->db_type) { 11 case M_DATA: 12 if (pullupmsg(mp, -1)) { 13 dma_addr = vtop(mp->b_rptr); 14 dma_len = mp->b_wptr - mp->b_rptr; 15 xxx_do_dma(dma_addr, dma_len); 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); </pre>

pullupmsg(9F)

EXAMPLE 1 Using pullupmsg() (Continued)

```
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()`.

NAME	put – call a STREAMS put procedure
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void put (queue_t *q, mblk_t *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to a STREAMS queue. <i>mp</i> Pointer to message block being passed into queue.
DESCRIPTION	put () calls the put procedure (put(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.
CONTEXT	put () 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. Since put () 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 (put(9E), putnext(9F), putctl(9F), qreply(9F)). This function is provided as a DDI/DKI conforming replacement for a direct call to a put procedure.
SEE ALSO	put(9E), freezestr(9F), putctl(9F), putctl1(9F), putnext(9F), putnextctl(9F), putnextctl1(9F), qprocson(9F), qreply(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
NOTES	The caller cannot have the stream frozen when calling this function. See freezestr(9F). 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, put(9E), putnext(9F), putctl(9F), and qreply(9F). This function is provided as a DDI/DKI conforming replacement for a direct call to a put procedure. The put () and putnext () functions should be called only after qprocson () is finished.

putbq(9F)

NAME	putbq – place a message at the head of a queue
SYNOPSIS	<pre>#include <sys/stream.h> int putbq(queue_t *q, mblk_t *bp);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>q</i> Pointer to the queue. <i>bp</i> Pointer to the message block.
DESCRIPTION	<code>putbq()</code> 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. <code>putbq()</code> can be used for ordinary, priority band, and high priority messages. However, unless precautions are taken, using <code>putbq()</code> 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 <code>bcanput(9F)</code> or <code>canput(9F)</code> 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 <code>QNOENB</code> is not set, the service routine is enabled.
RETURN VALUES	<code>putbq()</code> returns 1 upon success and 0 upon failure. Note – Upon failure, the caller should call <code>freemsg(9F)</code> to free the pointer to the message block.
CONTEXT	<code>putbq()</code> can be called from user or interrupt context.
EXAMPLES	See the <code>bufcall(9F)</code> function page for an example of <code>putbq()</code> .
SEE ALSO	<code>bcanput(9F)</code> , <code>bufcall(9F)</code> , <code>canput(9F)</code> , <code>getq(9F)</code> , <code>putq(9F)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	putctl1 – send a control message with a one-byte parameter to a queue
SYNOPSIS	<pre>#include <sys/stream.h> int putctl1(queue_t *q, int type, int p);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Queue to which the message is to be sent. <i>type</i> Type of message. <i>p</i> One-byte parameter.
DESCRIPTION	<code>putctl1()</code>, like <code>putctl(9F)</code>, 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 <code>M_DELAY</code> message is being sent. <code>putctl1()</code> fails if <i>type</i> is <code>M_DATA</code>, <code>M_PROTO</code>, or <code>M_PCPROTO</code>, or if a message block cannot be allocated. If successful, <code>putctl1()</code> calls the <code>put(9E)</code> routine of the queue pointed to by <i>q</i> with the newly allocated and initialized message.
RETURN VALUES	On success, 1 is returned. 0 is returned if <i>type</i> is a data type, or if a message block cannot be allocated.
CONTEXT	<code>putctl1()</code> can be called from user or interrupt context.
EXAMPLES	See the <code>putctl(9F)</code> function page for an example of <code>putctl1()</code> .
SEE ALSO	<code>put(9E)</code>, <code>allocb(9F)</code>, <code>datamsg(9F)</code>, <code>putctl(9F)</code>, <code>putnextctl1(9F)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

putctl(9F)

NAME	putctl – send a control message to a queue
SYNOPSIS	<pre>#include <sys/stream.h> int putctl(queue_t *q, int type);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Queue to which the message is to be sent. <i>type</i> Message type (must be control, not data type).
DESCRIPTION	putctl() tests the <i>type</i> argument to make sure a data type has not been specified, and then attempts to allocate a message block. putctl() fails if <i>type</i> is M_DATA, M_PROTO, or M_PCPROTO, or if a message block cannot be allocated. If successful, putctl() calls the put(9E) routine of the queue pointed to by <i>q</i> with the newly allocated and initialized messages.
RETURN VALUES	On success, 1 is returned. If <i>type</i> is a data type, or if a message block cannot be allocated, 0 is returned.
CONTEXT	putctl() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using putctl() The send_ctl() routine is used to pass control messages downstream. M_BREAK messages are handled with putctl() (line 11). putctl1(9F) (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, cmn_err(9F) 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->q_next, M_BREAK)) 12 num_alloc_fail++; 13 break; 14 15 case M_DELAY: 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; </pre>

EXAMPLE 1 Using `putctl()` (Continued)

```
23         }  
24     }
```

SEE ALSO `put(9E)`, `cmn_err(9F)`, `datamsg(9F)`, `putctl1(9F)`, `putnextctl(9F)`

Writing Device Drivers

STREAMS Programming Guide

putnext(9F)

NAME	putnext – send a message to the next queue
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void putnext (queue_t *q, mblk_t *mp) ;</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue from which the message <i>mp</i> will be sent. <i>mp</i> Message to be passed.
DESCRIPTION	putnext () is used to pass a message to the put(9E) routine of the next queue in the stream.
RETURN VALUES	None.
CONTEXT	putnext () can be called from user or interrupt context.
EXAMPLES	See allocb(9F) for an example of using putnext () .
SEE ALSO	put(9E), allocb(9F), put(9F), qprocson(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
NOTES	The put () and putnext () functions should be called only after qprocson () is finished.

NAME	putnextctl1 – send a control message with a one-byte parameter to a queue
SYNOPSIS	<pre>#include <sys/stream.h> int putnextctl1(queue_t *q, int type, int p);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Queue to which the message is to be sent. <i>type</i> Type of message. <i>p</i> One-byte parameter.
DESCRIPTION	<code>putnextctl1()</code>, like <code>putctl1(9F)</code>, 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 <code>M_DELAY</code> message is being sent. <code>putnextctl1()</code> fails if <i>type</i> is <code>M_DATA</code>, <code>M_PROTO</code>, or <code>M_PCPROTO</code>, or if a message block cannot be allocated. If successful, <code>putnextctl1()</code> calls the <code>put(9E)</code> routine of the queue pointed to by <i>q</i> with the newly allocated and initialized message. A call to <code>putnextctl1(q, type, p)</code> is an atomic equivalent of <code>putctl1(q->q_next, type, p)</code>. The STREAMS framework provides whatever mutual exclusion is necessary to insure that dereferencing <i>q</i> through its <code>q_next</code> field and then invoking <code>putctl1(9F)</code> proceeds without interference from other threads. <code>putnextctl1()</code> should always be used in preference to <code>putctl1(9F)</code>
RETURN VALUES	On success, 1 is returned. 0 is returned if <i>type</i> is a data type, or if a message block cannot be allocated.
CONTEXT	<code>putnextctl1()</code> can be called from user or interrupt context.
EXAMPLES	See the <code>putnextctl1(9F)</code> function page for an example of <code>putnextctl1()</code> .
SEE ALSO	<code>put(9E)</code>, <code>allocb(9F)</code>, <code>datamsg(9F)</code>, <code>putctl1(9F)</code>, <code>putnextctl1(9F)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

putnextctl(9F)

NAME	putnextctl – send a control message to a queue
SYNOPSIS	<pre>#include <sys/stream.h> int putnextctl(queue_t *q, int type);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Queue to which the message is to be sent. <i>type</i> Message type (must be control, not data type).
DESCRIPTION	putnextctl() tests the <i>type</i> argument to make sure a data type has not been specified, and then attempts to allocate a message block. putnextctl() fails if <i>type</i> is M_DATA, M_PROTO, or M_PCPROTO, or if a message block cannot be allocated. If successful, putnextctl() calls the put(9E) routine of the queue pointed to by <i>q</i> with the newly allocated and initialized messages. A call to putnextctl(<i>q</i>, <i>type</i>) is an atomic equivalent of putctl(<i>q</i>->q_next, <i>type</i>). The STREAMS framework provides whatever mutual exclusion is necessary to insure that dereferencing <i>q</i> through its q_next field and then invoking putctl(9F) proceeds without interference from other threads. putnextctl() should always be used in preference to putctl(9F)
RETURN VALUES	On success, 1 is returned. If <i>type</i> is a data type, or if a message block cannot be allocated, 0 is returned.
CONTEXT	putnextctl() can be called from user or interrupt context.
EXAMPLES	The send_ctl routine is used to pass control messages downstream. M_BREAK messages are handled with putnextctl() (line 8). putnextctl1(9F) (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, cmn_err(9F) panics the system (line 18). <pre> 1 void 2 send_ctl(queue_t *wrq, uchar_t type, uchar_t parm) 3 { 4 extern int num_alloc_fail; 5 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:</pre>

putnextctl(9F)

```
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), putnextctl(9F)

Writing Device Drivers

STREAMS Programming Guide

putq(9F)

NAME	putq – put a message on a queue
SYNOPSIS	<pre>#include <sys/stream.h> int putq(queue_t *q, mblk_t *bp);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>q</i> Pointer to the queue to which the message is to be added. <i>bp</i> Message to be put on the queue.
DESCRIPTION	<code>putq()</code> 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 <code>QNOENB</code> is not set, the service routine is enabled. If no other processing is done, <code>putq()</code> can be used as the module's put routine.
RETURN VALUES	<code>putq()</code> returns 1 on success and 0 on failure. Note – Upon failure, the caller should call <code>freemsg(9F)</code> to free the pointer to the message block.
CONTEXT	<code>putq()</code> can be called from user or interrupt context.
EXAMPLES	See the <code>datamsg(9F)</code> function page for an example of <code>putq()</code> .
SEE ALSO	<code>datamsg(9F)</code> , <code>putbq(9F)</code> , <code>qenable(9F)</code> , <code>rmvq(9F)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	qassociate – associate STREAMS queue with driver instance
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stream.h> #include <sys/stropts.h> #include <sys/ddi.h> #include <sys/sunddi.h> int qassociate(queue_t *q, int instance);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). This entry point is required for drivers which export cb_ops(9S) entry points.
PARAMETERS	queue_t *q Pointer to a queue(9S) structure. Either the read or write queue can be used. int instance Driver instance number or -1.
DESCRIPTION	The qassociate() function associates the specified STREAMS queue with the specified instance of the bottom driver in the queue. Upon successful return, the stream is associated with the instance with any prior association dissolved. A DLPI style-2 driver calls qassociate() while processing the DL_ATTACH_REQ message. The driver is also expected to call this interface while performing stream associations through other means, such as ndd(1M) ioctl commands. If instance is -1, the stream is left unassociated with any hardware instance. If the interface returns failure, the stream is not associated with the specified instance. Any prior association is left untouched. The interface typically fails because of failure to locate and attach the device instance. The interface never fails if the specified instance is -1.
CONTEXT	qassociate() can be called from the stream's put(9E) entry point.
RETURN VALUES	0 Success. -1 Failure.
EXAMPLES	A Style-2 network driver's DL_ATTACH_REQ code would specify: <pre>if (qassociate(q, instance) != 0) goto fail;</pre> The association prevents Dynamic Reconfiguration (DR) from detaching the instance. A Style-2 network driver's DL_DETACH code would specify: <pre>(void) qassociate(q, -1);</pre>

qassociate(9F)

This dissolves the queue's association with any device instance.

A Style-2 network driver's open(9E) code must call:

```
qassociate(q, -1);
```

This informs the framework that this driver has been modified to be DDI-compliant.

SEE ALSO dlpi(7P), open(9E), put(9E), ddi_no_info(9F), queue(9S)

NAME	qbufcall – call a function when a buffer becomes available
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> bufcall_id_t qbufcall(queue_t *q, size_t size, uint_t pri, void*funcvoid *arg, void *arg);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>q</i> Pointer to STREAMS queue structure. <i>size</i> Number of bytes required for the buffer. <i>pri</i> Priority of the allocb(9F) allocation request (not used). <i>func</i> Function or driver routine to be called when a buffer becomes available. <i>arg</i> Argument to the function to be called when a buffer becomes available.
DESCRIPTION	qbufcall() serves as a qtimeout(9F) call of indeterminate length. When a buffer allocation request fails, qbufcall() 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 allocb() or it may do something else. The qbufcall() function is tailored to be used with the enhanced STREAMS framework interface, which is based on the concept of perimeters. (See mt-streams(9F).) qbufcall() schedules the specified function to execute after entering the perimeters associated with the queue passed in as the first parameter to qbufcall(). All outstanding timeouts and bufcalls must be cancelled (using, respectively, quntimeout(9F) and qunbufcall(9F)) before a driver close routine can block and before the close routine calls qprocsoff(9F). qprocson(9F) must be called before calling either qbufcall() or qtimeout(9F).
RETURN VALUES	If successful, qbufcall() returns a qbufcall ID that can be used in a call to qunbufcall(9F) to cancel the request. If the qbufcall() scheduling fails, <i>func</i> is never called and 0 is returned.
CONTEXT	qbufcall() can be called from user or interrupt context.
SEE ALSO	allocb(9F), mt-streams(9F), qprocson(9F), qtimeout(9F), qunbufcall(9F), quntimeout(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
WARNINGS	Even when <i>func</i> is called by qbufcall(), allocb(9F) can fail if another module or driver had allocated the memory before <i>func</i> was able to call allocb(9F).

qenable(9F)

NAME	qenable – enable a queue
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void qenable(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue to be enabled.
DESCRIPTION	<code>qenable()</code> 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.
CONTEXT	<code>qenable()</code> can be called from user or interrupt context.
EXAMPLES	See the <code>dupb(9F)</code> function page for an example of the <code>qenable()</code> .
SEE ALSO	<code>dupb(9F)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	qprocson, qprocsoff – enable, disable put and service routines
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void qprocson(queue_t *q) ; void qprocsoff(queue_t *q) ;</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the RD side of a STREAMS queue pair.
DESCRIPTION	qprocson() 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. qprocson() 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. qprocson() must be called before calling put(9F), putnext(9F), qbufcall(9F), qtimeout(9F), qwait(9F), or qwait_sig(9F), qprocsoff() 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. The module or driver instance is guaranteed to be single-threaded before qprocson() is called and after qprocsoff() is called, except for threads executing asynchronous events such as interrupt handlers and callbacks, which must be handled separately.
CONTEXT	These routines can be called from user or interrupt context.
SEE ALSO	close(9E), open(9E), put(9E), srv(9E), put(9F), putnext(9F), qbufcall(9F), qtimeout(9F), qwait(9F), qwait_sig(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
NOTES	The caller may not have the STREAM frozen during either of these calls.

qreply(9F)

NAME	qreply – send a message on a stream in the reverse direction
SYNOPSIS	<pre>#include <sys/stream.h> void qreply(queue_t *q, mblk_t *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the queue. <i>mp</i> Pointer to the message to be sent in the opposite direction.
DESCRIPTION	qreply() sends messages in the reverse direction of normal flow. That is, qreply(<i>q</i> , <i>mp</i>) is equivalent to putnext(OTHERQ(<i>q</i>), <i>mp</i>).
CONTEXT	qreply() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Canonical Flushing Code for STREAMS Drivers. 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 srv(9E). Its write-side put procedure handles M_FLUSH messages by first checking the FLUSHW bit in the first byte of the message, then the write queue is flushed (line 8) and the FLUSHW bit is turned off (line 9). See put(9E). If the FLUSHR 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 qreply(9F) function (line 13). If the FLUSHR bit is off, then the message is freed (line 15). See the example for flushq(9F) for the canonical flushing code for modules. <pre> 1 xxxwput(q, mp) 2 queue_t *q; 3 mblk_t *mp; 4 { 5 switch(mp->b_datap->db_type) { 6 case M_FLUSH: 7 if (*mp->b_rptr & FLUSHW) { 8 flushq(q, FLUSHALL); 9 *mp->b_rptr &= ~FLUSHW; 10 } 11 if (*mp->b_rptr & FLUSHR) { 12 flushq(RD(q), FLUSHALL); 13 qreply(q, mp); 14 } else { 15 freemsg(mp); 16 } 17 break; 18 . . . 19 }</pre>
SEE ALSO	put(9E), srv(9E), flushq(9F), OTHERQ(9F), putnext(9F) <i>Writing Device Drivers</i>

qsize(9F)

NAME	qsize – find the number of messages on a queue
SYNOPSIS	<pre>#include <sys/stream.h> int qsize(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Queue to be evaluated.
DESCRIPTION	<code>qsize()</code> evaluates the queue <i>q</i> and returns the number of messages it contains.
RETURN VALUES	If there are no message on the queue, <code>qsize()</code> returns 0. Otherwise, it returns the integer representing the number of messages on the queue.
CONTEXT	<code>qsize()</code> can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	qtimeout – execute a function after a specified length of time
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> timeout_id_t qtimeout(queue_t *q, void *funcvoid *, void *arg, clock_t ticks);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>q</i> Pointer to STREAMS queue structure. <i>func</i> Kernel function to invoke when the time increment expires. <i>arg</i> Argument to the function. <i>ticks</i> Number of clock ticks to wait before the function is called. Use <code>drv_usectohz(9F)</code> to convert microseconds to clock ticks.
DESCRIPTION	The <code>qtimeout()</code> 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. The <code>qtimeout()</code> function is tailored to be used with the enhanced STREAMS framework interface which is based on the concept of perimeters. (See <code>mt-streams(9F)</code>.) <code>qtimeout()</code> schedules the specified function to execute after entering the perimeters associated with the queue passed in as the first parameter to <code>qtimeout()</code>. All outstanding timeouts and bufcalls must be cancelled (using, respectively, <code>quntimeout(9F)</code> and <code>qunbufcall(9F)</code>) before a driver close routine can block and before the close routine calls <code>qprocsoff(9F)</code>. <code>qprocson(9F)</code> must be called before calling <code>qtimeout()</code>.
RETURN VALUES	<code>qtimeout()</code> returns an opaque non-zero timeout identifier that can be passed to <code>quntimeout(9F)</code> to cancel the request. Note: No value is returned from the called function.
CONTEXT	<code>qtimeout()</code> can be called from user or interrupt context.
SEE ALSO	<code>drv_usectohz(9F)</code>, <code>mt-streams(9F)</code>, <code>qbufcall(9F)</code>, <code>qprocson(9F)</code>, <code>qunbufcall(9F)</code>, <code>quntimeout(9F)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

qunbufcall(9F)

NAME	qunbufcall – cancel a pending qbufcall request
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void qunbufcall(queue_t *q, bufcall_id_t id);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>q</i> Pointer to STREAMS queue_t structure. <i>id</i> Identifier returned from qbufcall(9F)
DESCRIPTION	qunbufcall() cancels a pending qbufcall() request. The argument <i>id</i> is a non-zero identifier of the request to be cancelled. <i>id</i> is returned from the qbufcall() function used to issue the cancel request. The qunbufcall() function is tailored to be used with the enhanced STREAMS framework interface which is based on the concept of perimeters. (See mt-streams(9F).) qunbufcall() 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 timeouts and bufcalls must be cancelled before a driver close routine can block and before the close routine calls qprocsoff(9F).
CONTEXT	qunbufcall() can be called from user or interrupt context.
SEE ALSO	mt-streams(9F), qbufcall(9F), qtimeout(9F), quntimeout(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	quntimeout – cancel previous qtimeout function call
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> clock_t quntimeout(queue_t *q, timeout_id_t id);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>q</i> Pointer to a STREAMS queue structure. <i>id</i> Opaque timeout ID a previous qtimeout(9F) call.
DESCRIPTION	quntimeout() cancels a pending qtimeout(9F) request. The quntimeout() function is tailored to be used with the enhanced STREAMS framework interface, which is based on the concept of perimeters. (See mt-streams(9F).) quntimeout() 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. quntimeout() should be executed for all outstanding timeouts before a driver or module close returns. All outstanding timeouts and bufcalls must be cancelled before a driver close routine can block and before the close routine calls qprocsoff(9F).
RETURN VALUES	quntimeout() returns -1 if the id is not found. Otherwise, quntimeout() returns a 0 or positive value.
CONTEXT	quntimeout() can be called from user or interrupt context.
SEE ALSO	mt-streams(9F), qbufcall(9F), qtimeout(9F), qunbufcall(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

qwait(9F)

NAME	qwait, qwait_sig – STREAMS wait routines
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void qwait(queue_t *q); int qwait_sig(queue_t *q);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>qp</i> Pointer to the queue that is being opened or closed.
DESCRIPTION	qwait() and qwait_sig() are used to wait for a message to arrive to the put(9E) or srv(9E) procedures. qwait() and qwait_sig() can also be used to wait for qbufcall(9F) or qtimeout(9F) callback procedures to execute. These routines can be used in the open(9E) and close(9E) procedures in a STREAMS driver or module. qwait() and qwait_sig() atomically exit the inner and outer perimeters associated with the queue, and wait for a thread to leave the module's put(9E), srv(9E), or qbufcall(9F) / qtimeout(9F) callback procedures. Upon return they re-enter the inner and outer perimeters. This can be viewed as there being an implicit wakeup when a thread leaves a put(9E) or srv(9E) procedure or after a qtimeout(9F) or qbufcall(9F) callback procedure has been run in the same perimeter. qprocson(9F) must be called before calling qwait() or qwait_sig(). qwait() is not interrupted by a signal, whereas qwait_sig() is interrupted by a signal. qwait_sig() normally returns non-zero, and returns zero when the waiting was interrupted by a signal. qwait() and qwait_sig() are similar to cv_wait() and cv_wait_sig() 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 condvar(9F).
RETURN VALUES	0 For qwait_sig(), indicates that the condition was not necessarily signaled, and the function returned because a signal was pending.
CONTEXT	These functions can only be called from an open(9E) or close(9E) routine.
EXAMPLES	EXAMPLE 1 Using qwait() The open routine sends down a T_INFO_REQ message and waits for the T_INFO_ACK. The arrival of the T_INFO_ACK is recorded by resetting a flag in the unit structure (WAIT_INFO_ACK). The example assumes that the module is D_MTQPAIR or D_MTPERMOD. <pre>xxopen(qp, . . .) queue_t *qp; { struct xxdata *xx; /* Allocate xxdata structure */</pre>

EXAMPLE 1 Using qwait() (Continued)

```

        qprocson(qp);
        /* Format T_INFO_ACK in mp */
        putnext(qp, mp);
        xx->xx_flags |= WAIT_INFO_ACK;
        while (xx->xx_flags & WAIT_INFO_ACK)
            qwait(qp);
        return (0);
    }
xxrput(qp, mp)
    queue_t *qp;
    mblk_t *mp;
{
    struct xxdata *xx = (struct xxdata *)q->q_ptr;

    ...

    case T_INFO_ACK:
        if (xx->xx_flags & WAIT_INFO_ACK) {
            /* Record information from info ack */
            xx->xx_flags &= ~WAIT_INFO_ACK;
            freemsg(mp);
            return;
        }

    ...
}

```

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

qwriter(9F)

NAME	qwriter – asynchronous STREAMS perimeter upgrade								
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> void qwriter(queue_t *qp, mblk_t *mp, void (*func)(), int perimeter);</pre>								
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).								
PARAMETERS	<table><tr><td><i>qp</i></td><td>Pointer to the queue.</td></tr><tr><td><i>mp</i></td><td>Pointer to a message that will be passed in to the callback function.</td></tr><tr><td><i>func</i></td><td>A function that will be called when exclusive (writer) access has been acquired at the specified perimeter.</td></tr><tr><td><i>perimeter</i></td><td>Either PERIM_INNER or PERIM_OUTER.</td></tr></table>	<i>qp</i>	Pointer to the queue.	<i>mp</i>	Pointer to a message that will be passed in to the callback function.	<i>func</i>	A function that will be called when exclusive (writer) access has been acquired at the specified perimeter.	<i>perimeter</i>	Either PERIM_INNER or PERIM_OUTER.
<i>qp</i>	Pointer to the queue.								
<i>mp</i>	Pointer to a message that will be passed in to the callback function.								
<i>func</i>	A function that will be called when exclusive (writer) access has been acquired at the specified perimeter.								
<i>perimeter</i>	Either PERIM_INNER or PERIM_OUTER.								
DESCRIPTION	<code>qwriter()</code> 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 <code>mt-streams(9F)</code>. The callback function is called as: <pre>(*func)(queue_t *qp, mblk_t *mp);</pre> <code>qwriter()</code> will acquire exclusive access immediately if possible, in which case the specified callback function will be executed before <code>qwriter()</code> returns. If this is not possible, <code>qwriter()</code> 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 <code>qwriter()</code> returns. One way to avoid dependencies on the execution of the callback function is to immediately return after calling <code>qwriter()</code> and let the callback function finish the processing of the message. When <code>qwriter()</code> 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.								
CONTEXT	<code>qwriter()</code> can only be called from an <code>put(9E)</code> or <code>srv(9E)</code> routine, or from a <code>qwriter()</code> , <code>qtimeout(9F)</code> , or <code>qbufcall(9F)</code> callback function.								
SEE ALSO	<code>put(9E)</code>, <code>srv(9E)</code>, <code>mt-streams(9F)</code>, <code>qbufcall(9F)</code>, <code>qtimeout(9F)</code> <i>STREAMS Programming Guide</i> <i>Writing Device Drivers</i>								

NAME	RD, rd – get pointer to the read queue
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> queue_t *RD(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the write queue whose read queue is to be returned.
DESCRIPTION	The RD() function accepts a write queue pointer as an argument and returns a pointer to the read queue of the same module. CAUTION: Make sure the argument to this function is a pointer to a write queue. RD() will not check for queue type, and a system panic could result if it is not the right type.
RETURN VALUES	The pointer to the read queue.
CONTEXT	RD() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Function page reference See the qreply(9F) function page for an example of RD().
SEE ALSO	qreply(9F), WR(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

rmalloc(9F)

NAME	rmalloc – allocate space from a resource map
SYNOPSIS	<pre>#include <sys/map.h> #include <sys/ddi.h> unsigned long rmalloc(struct map *mp, size_t size);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Resource map from where the resource is drawn. <i>size</i> Number of units of the resource.
DESCRIPTION	rmalloc() 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 rmallocmap(9F). rmalloc() is one of five functions used for resource map management. The other functions include: rmalloc_wait(9F) Allocate space from a resource map, wait if necessary. rmfree(9F) Return previously allocated space to a map. rmallocmap(9F) Allocate a resource map and initialize it. rmfreemap(9F) Deallocate a resource map. rmalloc() 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.
RETURN VALUES	Under normal conditions, rmalloc() returns the base index of the allocated space. Otherwise, rmalloc() returns a 0 if all resource map entries are already allocated.
CONTEXT	rmalloc() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Illustrating the principles of map management 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 rmallocmap(9F) and rmfree(9F) functions. rmallocmap(9F) is called to establish the number of slots or entries in the map, and rmfree(9F) to initialize the resource area the map is to manage. 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. <pre>1 #define XX_MAPSIZE 12 2 #define XX_BUFSIZE 2560</pre>

EXAMPLE 1 Illustrating the principles of map management (Continued)

```

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)

```

rmalloc(9F)

EXAMPLE 2 Allocating buffers (Continued)

```
6  {
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     /*
20      * Move data to buffer.  If invalid address is found,
21      * return buffer to map and return error code.
22      */
23     if (uiomove(addr, size, UIO_READ, uiop) == -1) {
24         rmfree(xx_mp, size, addr);
25         return(EFAULT);
26     }
27 }
```

SEE ALSO kmem_alloc(9F), rmalloc_wait(9F), rmallocmap(9F), rmfree(9F),
rmfreemap(9F), uiomove(9F)

Writing Device Drivers

NAME	rmallocmap, rmallocmap_wait, rmfreemap – allocate and free resource maps
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> struct map *rmallocmap(size_t mapsize); struct map *rmallocmap_wait(size_t mapsize); void rmfreemap(struct map *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mapsize</i> Number of entries for the map. <i>mp</i> A pointer to the map structure to be deallocated.
DESCRIPTION	<code>rmallocmap()</code> 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. <code>rmallocmap()</code> 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 <code>rmfree(9F)</code> 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 <code>rmallocmap()</code> is followed by a call to <code>rmfree(9F)</code>, passing the address of the map returned from <code>rmallocmap()</code>, the total size of the resource, and the first index into the actual resource. The resource map allocated by <code>rmallocmap()</code> 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 <code>rmalloc(9F)</code>, <code>rmalloc_wait(9F)</code>, and <code>rmfree(9F)</code>. <code>rmallocmap_wait()</code> is similar to <code>rmallocmap()</code>, with the exception that it will wait for space to become available if necessary. <code>rmfreemap()</code> deallocates a resource map structure previously allocated by <code>rmallocmap()</code> or <code>rmallocmap_wait()</code>. The argument <i>mp</i> is a pointer to the map structure to be deallocated.
RETURN VALUES	Upon successful completion, <code>rmallocmap()</code> and <code>rmallocmap_wait()</code> return a pointer to the newly allocated map structure. Upon failure, <code>rmallocmap()</code> returns a NULL pointer.
CONTEXT	<code>rmallocmap()</code> and <code>rmfreemap()</code> can be called from user, kernel, or interrupt context. <code>rmallocmap_wait()</code> can only be called from user or kernel context.
SEE ALSO	<code>rmalloc(9F)</code>, <code>rmalloc_wait(9F)</code>, <code>rmfree(9F)</code> <i>Writing Device Drivers</i>

rmalloc_wait(9F)

NAME	rmalloc_wait – allocate space from a resource map, wait if necessary
SYNOPSIS	<pre>#include <sys/map.h> #include <sys/ddi.h> unsigned long rmalloc_wait(struct map *mp, size_t size);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>mp</i> Pointer to the resource map from which space is to be allocated. <i>size</i> Number of units of space to allocate.
DESCRIPTION	<i>rmalloc_wait</i> () requests an allocation of space from a resource map. <i>rmalloc_wait</i> () is similar to the <i>rmalloc</i> (9F) function with the exception that it will wait for space to become available if necessary.
RETURN VALUES	<i>rmalloc_wait</i> () returns the base of the allocated space.
CONTEXT	This function can be called from user or interrupt context. However, in most cases <i>rmalloc_wait</i> () should be called from user context only.
SEE ALSO	<i>rmalloc</i> (9F), <i>rmallocmap</i> (9F), <i>rmfree</i> (9F), <i>rmfreemap</i> (9F) <i>Writing Device Drivers</i>

NAME	rmfree – free space back into a resource map
SYNOPSIS	<pre>#include <sys/map.h> #include <sys/ddi.h> void rmfree(struct map *mp, size_t size, ulong_t index);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>mp</i> Pointer to the map structure. <i>size</i> Number of units being freed. <i>index</i> Index of the first unit of the allocated resource.
DESCRIPTION	rmfree() releases space back into a resource map. It is the opposite of rmalloc(9F), which allocates space that is controlled by a resource map structure. When releasing resources using rmfree() the size and index passed to rmfree() must exactly match the size and index values passed to and returned from a previous call to rmalloc(). Resources cannot be returned piecemeal. Drivers may define resource maps for resource allocation, in terms of arbitrary units, using the rmallocmap(9F) 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. rmfree() frees up unallocated space for re-use. rmfree() can also be used to initialize a resource map, in which case the size and index should cover the entire resource area.
CONTEXT	rmfree() can be called from user or interrupt context.
SEE ALSO	rmalloc(9F), rmalloc_wait(9F), rmallocmap(9F), rmfreemap(9F) <i>Writing Device Drivers</i>

rmvb(9F)

NAME	rmvb – remove a message block from a message
SYNOPSIS	<pre>#include <sys/stream.h> mblk_t *rmvb(mblk_t *mp, mblk_t *bp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Message from which a block is to be removed. mblk_t is an instance of the msgb(9S) structure. <i>bp</i> Message block to be removed.
DESCRIPTION	rmvb() 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.
RETURN VALUES	If successful, a pointer to the message (minus the removed block) is returned. The pointer is NULL if <i>bp</i> was the only block of the message before rmvb() was called. If the designated message block (<i>bp</i>) does not exist, -1 is returned.
CONTEXT	rmvb() can be called from user or interrupt context.
EXAMPLES	This routine removes all zero-length M_DATA 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 M_DATA 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). <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->b_cont; 11 if ((tmp->b_datap->db_type == M_DATA) && 12 (tmp->b_rptr == tmp->b_wptr)) { 13 (void) rmvb(mp, tmp); 14 freeb(tmp); 15 } 16 tmp = nmp; 17 }</pre>
SEE ALSO	freeb(9F), msgb(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	rmvq – remove a message from a queue
SYNOPSIS	<pre>#include <sys/stream.h> void rmvq(queue_t *q, mblk_t *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Queue containing the message to be removed. <i>mp</i> Message to remove.
DESCRIPTION	rmvq() 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 rmvq() or getq(9F) should be used to remove a message from a queue.
CONTEXT	rmvq() can be called from user or interrupt context.
EXAMPLES	This code fragment illustrates how one may flush one type of message from a queue. In this case, only M_PROTO T_DATA_IND messages are flushed. For each message on the queue, if it is an M_PROTO message (line 8) of type T_DATA_IND (line 10), save a pointer to the next message (line 11), remove the T_DATA_IND message (line 12) and free it (line 13). Continue with the next message in the list (line 19). <pre> 1 mblk_t *mp, *nmp; 2 queue_t *q; 3 union T_primitives *tp; 4 5 /* Insert code here to protect queue and message block */ 6 mp = q->q_first; 7 while (mp) { 8 if (mp->b_datap->db_type == M_PROTO) { 9 tp = (union T_primitives *)mp->b_rptr; 10 if (tp->type == T_DATA_IND) { 11 nmp = mp->b_next; 12 rmvq(q, mp); 13 freemsg(mp); 14 mp = nmp; 15 } else { 16 mp = mp->b_next; 17 } 18 } else { 19 mp = mp->b_next; 20 } 21 } 22 /* End of region that must be protected */</pre> When using rmvq(), you must ensure that the queue and the message block is not modified by another thread at the same time. You can achieve this either by using STREAMS functions or by implementing your own locking.
SEE ALSO	freemsg(9F), getq(9F), insq(9F)
	<i>Writing Device Drivers</i>

rmvq(9F)

STREAMS Programming Guide

WARNINGS Make sure that the message *mp* is linked onto *q* to avoid a possible system panic.

NAME	rwlock, rw_init, rw_destroy, rw_enter, rw_exit, rw_tryenter, rw_downgrade, rw_tryupgrade, rw_read_locked – readers/writer lock functions										
SYNOPSIS	<pre>#include <sys/ksynch.h> void rw_init(krwlock_t *rwlp, char *name, krw_type_t type, void *arg); void rw_destroy(krwlock_t *rwlp); void rw_enter(krwlock_t *rwlp, krw_t enter_type); void rw_exit(krwlock_t *rwlp); int rw_tryenter(krwlock_t *rwlp, krw_t enter_type); void rw_downgrade(krwlock_t *rwlp); int rw_tryupgrade(krwlock_t *rwlp); int rw_read_locked(krwlock_t *rwlp);</pre>										
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>rwlp</i></td><td>Pointer to a <code>krwlock_t</code> readers/writer lock.</td></tr> <tr> <td><i>name</i></td><td>Descriptive string. This is obsolete and should be NULL. (Non-null strings are legal, but they're a waste of kernel memory.)</td></tr> <tr> <td><i>type</i></td><td>Type of readers/writer lock.</td></tr> <tr> <td><i>arg</i></td><td>Type-specific argument for initialization function.</td></tr> <tr> <td><i>enter_type</i></td><td>One of the values <code>RW_READER</code> or <code>RW_WRITER</code>, indicating whether the lock is to be acquired non-exclusively (<code>RW_READER</code>) or exclusively (<code>RW_WRITER</code>).</td></tr> </table>	<i>rwlp</i>	Pointer to a <code>krwlock_t</code> readers/writer lock.	<i>name</i>	Descriptive string. This is obsolete and should be NULL. (Non-null strings are legal, but they're a waste of kernel memory.)	<i>type</i>	Type of readers/writer lock.	<i>arg</i>	Type-specific argument for initialization function.	<i>enter_type</i>	One of the values <code>RW_READER</code> or <code>RW_WRITER</code> , indicating whether the lock is to be acquired non-exclusively (<code>RW_READER</code>) or exclusively (<code>RW_WRITER</code>).
<i>rwlp</i>	Pointer to a <code>krwlock_t</code> readers/writer lock.										
<i>name</i>	Descriptive string. This is obsolete and should be NULL. (Non-null strings are legal, but they're a waste of kernel memory.)										
<i>type</i>	Type of readers/writer lock.										
<i>arg</i>	Type-specific argument for initialization function.										
<i>enter_type</i>	One of the values <code>RW_READER</code> or <code>RW_WRITER</code> , indicating whether the lock is to be acquired non-exclusively (<code>RW_READER</code>) or exclusively (<code>RW_WRITER</code>).										
DESCRIPTION	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. 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. <code>rw_init()</code> initializes a readers/writer lock. It is an error to initialize a lock more than once. The <i>type</i> argument should be set to <code>RW_DRIVER</code>. If the lock is used by the interrupt handler, the type-specific argument, <i>arg</i>, should be the <code>ddi_iblock_cookie</code> returned from <code>ddi_get_iblock_cookie(9F)</code> or <code>ddi_get_soft_iblock_cookie(9F)</code>. If the lock is not used by any interrupt handler, the argument should be NULL.										

rwlock(9F)

`rw_destroy()` releases any resources that might have been allocated by `rw_init()`. It should be called before freeing the memory containing the lock. The lock must not be held by any thread when it is destroyed.

`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	<code>rw_tryenter()</code> could not obtain the lock without blocking.
0	<code>rw_tryupgrade()</code> was unable to perform the upgrade because of other threads holding or waiting to hold the lock.
0	<code>rw_read_locked()</code> returns 0 if the lock is held by the caller for write.
non-zero	from <code>rw_read_locked()</code> if the lock is held by the caller for read.
non-zero	successful return from <code>rw_tryenter()</code> or <code>rw_tryupgrade()</code> .

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.

rwlock(9F)

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)`.

SAMESTR(9F)

NAME	SAMESTR, samestr – test if next queue is in the same stream
SYNOPSIS	<pre>#include <sys/stream.h> int SAMESTR(queue_t *q);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>q</i> Pointer to the queue.
DESCRIPTION	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 <i>q_next</i> field of queue(9S) to see if the stream continues beyond <i>q</i> .
RETURN VALUES	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.
CONTEXT	SAMESTR() can be called from user or interrupt context.
SEE ALSO	OTHERQ(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	scsi_abort – abort a SCSI command
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_abort(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>ap</i> Pointer to a scsi_address structure. <i>pkt</i> Pointer to a scsi_pkt(9S) structure.
DESCRIPTION	scsi_abort() 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.
RETURN VALUES	scsi_abort() returns: 1 on success. 0 on failure.
CONTEXT	scsi_abort() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Terminating a command. <pre>if (scsi_abort(&devp->sd_address, pkt) == 0) { (void) scsi_reset(&devp->sd_address, RESET_ALL); }</pre>
SEE ALSO	tran_abort(9E), scsi_reset(9F), scsi_pkt(9S) <i>Writing Device Drivers</i>

scsi_alloc_consistent_buf(9F)

NAME	scsi_alloc_consistent_buf – allocate an I/O buffer for SCSI DMA												
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> struct buf *scsi_alloc_consistent_buf(struct scsi_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>												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).												
PARAMETERS	<table> <tr> <td><i>ap</i></td><td>Pointer to the scsi_address(9S) structure.</td></tr> <tr> <td><i>bp</i></td><td>Pointer to the buf(9S) structure.</td></tr> <tr> <td><i>datalen</i></td><td>Number of bytes for the data buffer.</td></tr> <tr> <td><i>bflags</i></td><td>Flags setting for the allocated buffer header. This should either be B_READ or B_WRITE.</td></tr> <tr> <td><i>callback</i></td><td>A pointer to a callback function, NULL_FUNC or SLEEP_FUNC.</td></tr> <tr> <td><i>arg</i></td><td>The callback function argument.</td></tr> </table>	<i>ap</i>	Pointer to the scsi_address(9S) structure.	<i>bp</i>	Pointer to the buf(9S) structure.	<i>datalen</i>	Number of bytes for the data buffer.	<i>bflags</i>	Flags setting for the allocated buffer header. This should either be B_READ or B_WRITE.	<i>callback</i>	A pointer to a callback function, NULL_FUNC or SLEEP_FUNC.	<i>arg</i>	The callback function argument.
<i>ap</i>	Pointer to the scsi_address(9S) structure.												
<i>bp</i>	Pointer to the buf(9S) structure.												
<i>datalen</i>	Number of bytes for the data buffer.												
<i>bflags</i>	Flags setting for the allocated buffer header. This should either be B_READ or B_WRITE.												
<i>callback</i>	A pointer to a callback function, NULL_FUNC or SLEEP_FUNC.												
<i>arg</i>	The callback function argument.												
DESCRIPTION	scsi_alloc_consistent_buf() allocates a buffer header and the associated data buffer for direct memory access (DMA) transfer. This buffer is allocated from the iobp space, which is considered consistent memory. For more details, see ddi_dma_mem_alloc(9F) and ddi_dma_sync(9F). For buffers allocated via scsi_alloc_consistent_buf(), and marked with the PKT_CONSISTENT flag via scsi_init_pkt(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. If <i>bp</i> is NULL, a new buffer header will be allocated using getrbuf(9F). In addition, if <i>datalen</i> is non-zero, a new buffer will be allocated using ddi_dma_mem_alloc(9F). <i>callback</i> indicates what the allocator routines should do when direct memory access (DMA) resources are not available; the valid values are: <table> <tr> <td>NULL_FUNC</td><td>Do not wait for resources. Return a NULL pointer.</td></tr> <tr> <td>SLEEP_FUNC</td><td>Wait indefinitely for resources.</td></tr> <tr> <td>Other Values</td><td><i>callback</i> points to a function that is called when resources may become available. <i>callback</i> must return either 0 (indicating that it attempted to allocate resources but 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. The last argument <i>arg</i> is supplied to the <i>callback</i> function when it is invoked.</td></tr> </table>	NULL_FUNC	Do not wait for resources. Return a NULL pointer.	SLEEP_FUNC	Wait indefinitely for resources.	Other Values	<i>callback</i> points to a function that is called when resources may become available. <i>callback</i> must return either 0 (indicating that it attempted to allocate resources but 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. The last argument <i>arg</i> is supplied to the <i>callback</i> function when it is invoked.						
NULL_FUNC	Do not wait for resources. Return a NULL pointer.												
SLEEP_FUNC	Wait indefinitely for resources.												
Other Values	<i>callback</i> points to a function that is called when resources may become available. <i>callback</i> must return either 0 (indicating that it attempted to allocate resources but 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. The last argument <i>arg</i> is supplied to the <i>callback</i> function when it is invoked.												

scsi_alloc_consistent_buf(9F)

RETURN VALUES	<code>scsi_alloc_consistent_buf()</code> returns a pointer to a <code>buf(9S)</code> structure on success. It returns <code>NULL</code> if resources are not available even if <code>waitfunc</code> was not <code>SLEEP_FUNC</code> .
CONTEXT	If <i>callback</i> is <code>SLEEP_FUNC</code> , then this routine may be called only 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.
EXAMPLES	EXAMPLE 1 Allocate a request sense packet with consistent DMA resources attached. <pre>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);</pre> EXAMPLE 2 Allocate an inquiry packet with consistent DMA resources attached. <pre>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); }</pre>
SEE ALSO	<code>ddi_dma_mem_alloc(9F)</code> , <code>ddi_dma_sync(9F)</code> , <code>getrbuf(9F)</code> , <code>scsi_destroy_pkt(9F)</code> , <code>scsi_init_pkt(9F)</code> , <code>scsi_free_consistent_buf(9F)</code> , <code>buf(9S)</code> , <code>scsi_address(9S)</code> <i>Writing Device Drivers</i>

scsi_cname(9F)

NAME	scsi_cname, scsi_dname, scsi_mname, scsi_rname, scsi_sname – decode a SCSI name												
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> char *scsi_cname(uchar_t cmd, char **cmdvec); char *scsi_dname(int dtype); char *scsi_mname(uchar_t msg); char *scsi_rname(uchar_t reason); char *scsi_sname(uchar_t sense_key);</pre>												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).												
PARAMETERS	<table> <tr> <td><i>cmd</i></td><td>A SCSI command value.</td></tr> <tr> <td><i>cmdvec</i></td><td>Pointer to an array of command strings.</td></tr> <tr> <td><i>dtype</i></td><td>Device type.</td></tr> <tr> <td><i>msg</i></td><td>A message value.</td></tr> <tr> <td><i>reason</i></td><td>A packet reason value.</td></tr> <tr> <td><i>sense_key</i></td><td>A SCSI sense key value.</td></tr> </table>	<i>cmd</i>	A SCSI command value.	<i>cmdvec</i>	Pointer to an array of command strings.	<i>dtype</i>	Device type.	<i>msg</i>	A message value.	<i>reason</i>	A packet reason value.	<i>sense_key</i>	A SCSI sense key value.
<i>cmd</i>	A SCSI command value.												
<i>cmdvec</i>	Pointer to an array of command strings.												
<i>dtype</i>	Device type.												
<i>msg</i>	A message value.												
<i>reason</i>	A packet reason value.												
<i>sense_key</i>	A SCSI sense key value.												
DESCRIPTION	<code>scsi_cname()</code> 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. <code>scsi_dname()</code> decodes the peripheral device type (for example, direct access or sequential access) in the inquiry data. <code>scsi_mname()</code> decodes SCSI messages. <code>scsi_rname()</code> decodes packet completion reasons. <code>scsi_sname()</code> decodes SCSI sense keys.												
RETURN VALUES	These functions return a pointer to a string. If an argument is invalid, they return a string to that effect.												
CONTEXT	These functions can be called from user or interrupt context.												
EXAMPLES	EXAMPLE 1 Decoding SCSI tape commands. <code>scsi_cname()</code> decodes SCSI tape commands as follows: <pre>static char *st_cmds[] = { "\000test unit ready", "\001rewind", "\003request sense", "\010read", "\012write",</pre>												

EXAMPLE 1 Decoding SCSI tape commands. *(Continued)*

```
        "\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));
    ..
```

SEE ALSO *Writing Device Drivers*

scsi_destroy_pkt(9F)

NAME	scsi_destroy_pkt – free an allocated SCSI packet and its DMA resource
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void scsi_destroy_pkt(struct scsi_pkt *pktp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>pktp</i> Pointer to a <i>scsi_pkt</i> (9S) structure.
DESCRIPTION	<i>scsi_destroy_pkt</i> () 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 <i>pktp</i> is freed.
CONTEXT	<i>scsi_destroy_pkt</i> () may be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Releasing resources. <pre>scsi_destroy_pkt(un->un_rqs);</pre>
SEE ALSO	<i>tran_destroy_pkt</i> (9E), <i>scsi_init_pkt</i> (9F), <i>scsi_pkt</i> (9S) <i>Writing Device Drivers</i>

NAME	scsi_dmaget, scsi_dmafree – SCSI dma utility routines
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> struct scsi_pkt *scsi_dmaget(struct scsi_pkt *pkt, opaque_t dmatoken, int (*callback) (void)); void scsi_dmafree(struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL	These interfaces are obsolete. Use <code>scsi_init_pkt(9F)</code> instead of <code>scsi_dmaget()</code> . Use <code>scsi_destroy_pkt(9F)</code> instead of <code>scsi_dmafree()</code> .
PARAMETERS	<i>pkt</i> A pointer to a <code>scsi_pkt(9S)</code> structure. <i>dmatoken</i> Pointer to an implementation dependent object <i>callback</i> Pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code>.
DESCRIPTION	<code>scsi_dmaget()</code> allocates DMA resources for an already allocated SCSI packet. <i>pkt</i> is a pointer to the previously allocated SCSI packet (see <code>scsi_pktalloc(9F)</code>). <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 <code>buf(9S)</code> structure. If <i>dmatoken</i> is <code>NULL</code>, no resources are allocated. <i>callback</i> indicates what <code>scsi_dmaget()</code> should do when resources are not available: <code>NULL_FUNC</code> Do not wait for resources. Return a <code>NULL</code> pointer. <code>SLEEP_FUNC</code> Wait indefinitely for resources. Other Values <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. <code>scsi_dmafree()</code> frees the DMA resources associated with the SCSI packet. The packet itself remains allocated.
RETURN VALUES	<code>scsi_dmaget()</code> returns a pointer to a <code>scsi_pkt</code> on success. It returns <code>NULL</code> if resources are not available.
CONTEXT	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. <code>scsi_dmafree()</code> can be called from user or interrupt context.

scsi_dmaget(9F)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), scsi_pktalloc(9F), scsi_pktfree(9F), scsi_resalloc(9F), scsi_resfree(9F), buf(9S), scsi_pkt(9S)

Writing Device Drivers

NOTES The `scsi_dmaget()` and `scsi_dmafree()` functions are obsolete and will be discontinued in a future release. These functions have been replaced by, respectively, `scsi_init_pkt(9F)` and `scsi_destroy_pkt(9F)`.

NAME	scsi_errmsg – display a SCSI request sense message																
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void scsi_errmsg(struct scsi_device *devp, struct scsi_pkt *pktp, char *drv_name, int severity, daddr_t blkno, daddr_t err_blkno, struct scsi_key_strings *cmdlist, struct scsi_extended_sense *sensep) ;</pre>																
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>devp</i></td><td>Pointer to the <i>scsi_device</i>(9S) structure.</td></tr> <tr> <td><i>pktp</i></td><td>Pointer to a <i>scsi_pkt</i>(9S) structure.</td></tr> <tr> <td><i>drv_name</i></td><td>String used by <i>scsi_log</i>(9F).</td></tr> <tr> <td><i>severity</i></td><td>Error severity level, maps to severity strings below.</td></tr> <tr> <td><i>blkno</i></td><td>Requested block number.</td></tr> <tr> <td><i>err_blkno</i></td><td>Error block number.</td></tr> <tr> <td><i>cmdlist</i></td><td>An array of SCSI command description strings.</td></tr> <tr> <td><i>sensep</i></td><td>A pointer to a <i>scsi_extended_sense</i>(9S) structure.</td></tr> </table>	<i>devp</i>	Pointer to the <i>scsi_device</i> (9S) structure.	<i>pktp</i>	Pointer to a <i>scsi_pkt</i> (9S) structure.	<i>drv_name</i>	String used by <i>scsi_log</i> (9F).	<i>severity</i>	Error severity level, maps to severity strings below.	<i>blkno</i>	Requested block number.	<i>err_blkno</i>	Error block number.	<i>cmdlist</i>	An array of SCSI command description strings.	<i>sensep</i>	A pointer to a <i>scsi_extended_sense</i> (9S) structure.
<i>devp</i>	Pointer to the <i>scsi_device</i> (9S) structure.																
<i>pktp</i>	Pointer to a <i>scsi_pkt</i> (9S) structure.																
<i>drv_name</i>	String used by <i>scsi_log</i> (9F).																
<i>severity</i>	Error severity level, maps to severity strings below.																
<i>blkno</i>	Requested block number.																
<i>err_blkno</i>	Error block number.																
<i>cmdlist</i>	An array of SCSI command description strings.																
<i>sensep</i>	A pointer to a <i>scsi_extended_sense</i> (9S) structure.																
DESCRIPTION	<i>scsi_errmsg</i>() interprets the request sense information in the <i>sensep</i> pointer and generates a standard message that is displayed using <i>scsi_log</i>(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. 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 <i>scsi_device</i>(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: <table> <tr> <td>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> <tr> <td>SCSI_ERR_RECOVERE</td><td>Recovered</td></tr> <tr> <td>SCSI_ERR_RETRYABL</td><td>Retryable</td></tr> <tr> <td>SCSI_ERR_FATAL</td><td>Fatal</td></tr> </table>	Severity Value:	String:	SCSI_ERR_ALL	All	SCSI_ERR_UNKNOWN	Unknown	SCSI_ERR_INFO	Informational	SCSI_ERR_RECOVERE	Recovered	SCSI_ERR_RETRYABL	Retryable	SCSI_ERR_FATAL	Fatal		
Severity Value:	String:																
SCSI_ERR_ALL	All																
SCSI_ERR_UNKNOWN	Unknown																
SCSI_ERR_INFO	Informational																
SCSI_ERR_RECOVERE	Recovered																
SCSI_ERR_RETRYABL	Retryable																
SCSI_ERR_FATAL	Fatal																

scsi_errmsg(9F)

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_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
```


scsi_errmsg(9F)

SEE ALSO `cmn_err(9F)`, `scsi_log(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`,
`scsi_pkt(9S)`

Writing Device Drivers

scsi_free_consistent_buf(9F)

NAME	scsi_free_consistent_buf – free a previously allocated SCSI DMA I/O buffer
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void scsi_free_consistent_buf(struct buf *bp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>bp</i> Pointer to the buf(9S) structure.
DESCRIPTION	<code>scsi_free_consistent_buf()</code> frees a buffer header and consistent data buffer that was previously allocated using <code>scsi_alloc_consistent_buf(9F)</code> .
CONTEXT	<code>scsi_free_consistent_buf()</code> may be called from either the user or the interrupt levels.
SEE ALSO	<code>freerbuf(9F)</code> , <code>scsi_alloc_consistent_buf(9F)</code> , <code>buf(9S)</code> <i>Writing Device Drivers</i>
WARNING	<code>scsi_free_consistent_buf()</code> will call <code>freerbuf(9F)</code> to free the <code>buf(9S)</code> that was allocated before or during the call to <code>scsi_alloc_consistent_buf(9F)</code> . If consistent memory is bound to a <code>scsi_pkt(9S)</code> , the <code>pkt</code> should be destroyed before freeing the consistent memory.

	scsi_get_device_type_scsi_options(9F)
NAME	scsi_get_device_type_scsi_options – look up per-device-type scsi-options property
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_get_device_type_scsi_options(dev_info_t *dip, struct scsi_device *devp, int default_scsi_options);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>dip</i> Pointer to the device info node for this HBA driver. <i>devp</i> Pointer to a <code>scsi_device(9S)</code> structure of the target. <i>default_scsi_options</i> Value returned if no match is found.
DESCRIPTION	The <code>scsi_get_device_type_scsi_options()</code> function looks up the property <code>device-type-scsi-options-list</code>, which can be specified in the HBA's <code>driver.conf(4)</code> file. This property allows specification of <code>scsi-options</code> on a per-device-type basis. The formal syntax is: <pre>device-type-scsi-options-list = <duplet> [, <duplet> *];</pre> where: <pre><duplet> := <vid+pid>, <scsi-options-property-name></pre> and: <pre><scsi-options-property-name> = <value>;</pre> The string <code><vid+pid></code> is returned by the device on a SCSI inquiry command. This string can contain any character in the range 0x20-0x7e. Characters such as double quote (") or single quote ('), which are not permitted in property value strings, are represented by their octal equivalent (for example, \042 and \047). Trailing spaces can be truncated. For example: <pre>device-type-scsi-options-list= "SEAGATE ST32550W", "seagate-options", "EXABYTE EXB-2501". "exabyte-options", "IBM OEM DFHSS4S", "ibm-options"; seagate-options = 0x78; exabyte-options = 0x58; ibm-options = 0x378;</pre> The <code>scsi_get_device_type_scsi_options()</code> function searches the list of duplets for a matching INQUIRY string. If a match is found, <code>scsi_get_device_type_scsi_options()</code> returns the corresponding value.

scsi_get_device_type_scsi_options(9F)

RETURN VALUES	<code>scsi_get_device_type_scsi_options()</code> returns the <code>scsi-options</code> value found, or if no match is found the <code>default_scsi_options</code> value passed in.

CONTEXT	This function can be called from kernel or interrupt context.

SEE ALSO	<i>Writing Device Drivers</i>

NAME	scsi_hba_attach_setup, scsi_hba_attach, scsi_hba_detach – SCSI HBA attach and detach routines
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> 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>
INTERFACE LEVEL PARAMETERS	Solaris architecture specific (Solaris DDI). <i>dip</i> A pointer to the dev_info_t structure, referring to the instance of the HBA device. <i>hba_lim</i> A pointer to a ddi_dma_lim(9S) structure. <i>hba_tran</i> A pointer to a scsi_hba_tran(9S) structure. <i>hba_flags</i> Flag modifiers. The only defined flag value is SCSI_HBA_TRAN_CLONE. <i>hba_options</i> Optional features provided by the HBA driver for future extensions; must be NULL. <i>hba_dma_attr</i> A pointer to a ddi_dma_attr(9S) structure.
DESCRIPTION	scsi_hba_attach_setup() is the recommended interface over scsi_hba_attach(). For scsi_hba_attach_setup() and scsi_hba_attach(): scsi_hba_attach() registers the DMA limits <i>hba_lim</i> and the transport vectors <i>hba_tran</i> of each instance of the HBA device defined by <i>dip</i>. scsi_hba_attach_setup() registers the DMA attributes <i>hba_dma_attr</i> and the transport vectors <i>hba_tran</i> of each instance of the HBA device defined by <i>dip</i>. 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. scsi_hba_attach() and scsi_hba_attach_setup() use the dev_bus_ops field in the dev_ops(9S) structure. The HBA driver should initialize this field to NULL before calling scsi_hba_attach() or scsi_hba_attach_setup(). If SCSI_HBA_TRAN_CLONE is requested in <i>hba_flags</i>, the <i>hba_tran</i> structure will be cloned once for each target attached to the HBA. The cloning of the 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

scsi_hba_attach_setup(9F)

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_PARITY

If not set, the HBA should not operate in parity mode.

SCSI_OPTIONS_QAS

If not set, the HBA should not make use of the Quick Arbitration Select feature. Consult your Sun hardware documentation to determine whether your machine supports QAS.

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_FAST40

If not set, the HBA should not operate the bus in FAST40 SCSI mode.

SCSI_OPTIONS_FAST80

If not set, the HBA should not operate the bus in FAST80 SCSI mode.

SCSI_OPTIONS_FAST160

If not set, the HBA should not operate the bus in FAST160 SCSI mode.

SCSI_OPTIONS_FAST320

If not set, the HBA should not operate the bus in FAST320 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 in synchronous transfer mode.

scsi_hba_attach_setup(9F)

scsi-reset-delay

SCSI bus or device reset recovery time, in milliseconds.

scsi-selection-timeout

Default SCSI selection phase timeout value, in milliseconds. Please refer to individual HBA man pages for any HBA-specific information

For `scsi_hba_detach()`:

`scsi_hba_detach()` removes the reference to the DMA limits or attributes structure and the transport vector for the given instance of an HBA driver.

RETURN VALUES

`scsi_hba_attach()`, `scsi_hba_attach_setup()`, and `scsi_hba_detach()` return `DDI_SUCCESS` if the function call succeeds, and return `DDI_FAILURE` on failure.

CONTEXT

`scsi_hba_attach()` and `scsi_hba_attach_setup()` should be called from `attach(9E)`. `scsi_hba_detach()` should be called from `detach(9E)`.

SEE ALSO

`attach(9E)`, `detach(9E)`, `tran_tgt_init(9E)`, `ddi_prop_get_int(9F)`, `ddi_dma_attr(9S)`, `ddi_dma_lim(9S)`, `dev_ops(9S)`, `scsi_address(9S)`, `scsi_hba_tran(9S)`

Writing Device Drivers

NOTES

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 `scsi_hba_detach()` is called.

The `scsi_hba_attach()` function is obsolete and will be discontinued in a future release. This function is replaced by `scsi_hba_attach_setup()`.

scsi_hba_init(9F)

NAME	scsi_hba_init, scsi_hba_fini – SCSI Host Bus Adapter system initialization and completion routines
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_hba_init(struct modlinkage *modlp); void scsi_hba_fini(struct modlinkage *modlp);</pre>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
PARAMETERS	<i>modlp</i> Pointer to the Host Bus Adapters module linkage structure.
DESCRIPTION	
scsi_hba_init()	scsi_hba_init() is the system-provided initialization routine for SCSI HBA drivers. The scsi_hba_init() function registers the HBA in the system and allows the driver to accept configuration requests on behalf of SCSI target drivers. The scsi_hba_init() routine must be called in the HBA's _init(9E) routine before mod_install(9F) is called. If mod_install(9F) fails, the HBA's _init(9E) should call scsi_hba_fini() before returning failure.
scsi_hba_fini()	scsi_hba_fini() is the system provided completion routine for SCSI HBA drivers. scsi_hba_fini() removes all of the system references for the HBA that were created in scsi_hba_init() . The scsi_hba_fini() routine should be called in the HBA's _fini(9E) routine if mod_remove(9F) is successful.
RETURN VALUES	scsi_hba_init() returns 0 if successful, and a non-zero value otherwise. If scsi_hba_init() fails, the HBA's _init() entry point should return the value returned by scsi_hba_init() .
CONTEXT	scsi_hba_init() and scsi_hba_fini() should be called from _init(9E) or _fini(9E) , respectively.
SEE ALSO	_fini(9E) , _init(9E) , mod_install(9F) , mod_remove(9F) , scsi_pktalloc(9F) , scsi_pktfree(9F) , scsi_hba_tran(9S)
	<i>Writing Device Drivers</i>
NOTES	The HBA is responsible for ensuring that no DDI request routines are called on behalf of its SCSI target drivers once scsi_hba_fini() is called.

NAME	scsi_hba_lookup_capstr – return index matching capability string
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_hba_lookup_capstr(char *capstr);</pre>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
PARAMETERS	<i>capstr</i> Pointer to a string.
DESCRIPTION	<code>scsi_hba_lookup_capstr()</code> attempts to match <i>capstr</i> against a known set of capability strings, and returns the defined index for the matched capability, if found. The set of indices and capability strings is: <pre>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"</pre>

scsi_hba_lookup_capstr(9F)

SCSI_CAP_RESET_NOTIFICATION
"reset-notification"
SCSI_CAP_QFULL_RETRIES
"qfull-retries"
SCSI_CAP_QFULL_RETRY_INTERVAL
"qfull-retry-interval"
SCSI_CAP_LUN_RESET
"lun-reset"

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

NAME	scsi_hba_pkt_alloc, scsi_hba_pkt_free – allocate and free a scsi_pkt structure																		
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> struct scsi_pkt *scsi_hba_pkt_alloc(dev_info_t *dip, struct scsi_address *ap, int cmdlen, int statuslen, int tgtlen, int hbalen, int (*callback, caddr_t arg, caddr_t arg); void scsi_hba_pkt_free(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>																		
INTERFACE LEVEL PARAMETERS	Solaris architecture specific (Solaris DDI). <table> <tr> <td><i>dip</i></td><td>Pointer to a dev_info_t structure, defining the HBA driver instance.</td></tr> <tr> <td><i>ap</i></td><td>Pointer to a scsi_address(9S) structure, defining the target instance.</td></tr> <tr> <td><i>cmdlen</i></td><td>Length in bytes to be allocated for the SCSI command descriptor block (CDB).</td></tr> <tr> <td><i>statuslen</i></td><td>Length in bytes to be allocated for the SCSI status completion block (SCB).</td></tr> <tr> <td><i>tgtlen</i></td><td>Length in bytes to be allocated for a private data area for the target driver's exclusive use.</td></tr> <tr> <td><i>hbalen</i></td><td>Length in bytes to be allocated for a private data area for the HBA driver's exclusive use.</td></tr> <tr> <td><i>callback</i></td><td> Indicates what scsi_hba_pkt_alloc() should do when resources are not available: NULL_FUNC Do not wait for resources. Return a NULL pointer. SLEEP_FUNC Wait indefinitely for resources. </td></tr> <tr> <td><i>arg</i></td><td>Must be NULL.</td></tr> <tr> <td><i>pkt</i></td><td>A pointer to a scsi_pkt(9S) structure.</td></tr> </table>	<i>dip</i>	Pointer to a dev_info_t structure, defining the HBA driver instance.	<i>ap</i>	Pointer to a scsi_address(9S) structure, defining the target instance.	<i>cmdlen</i>	Length in bytes to be allocated for the SCSI command descriptor block (CDB).	<i>statuslen</i>	Length in bytes to be allocated for the SCSI status completion block (SCB).	<i>tgtlen</i>	Length in bytes to be allocated for a private data area for the target driver's exclusive use.	<i>hbalen</i>	Length in bytes to be allocated for a private data area for the HBA driver's exclusive use.	<i>callback</i>	Indicates what scsi_hba_pkt_alloc() should do when resources are not available: NULL_FUNC Do not wait for resources. Return a NULL pointer. SLEEP_FUNC Wait indefinitely for resources.	<i>arg</i>	Must be NULL.	<i>pkt</i>	A pointer to a scsi_pkt(9S) structure.
<i>dip</i>	Pointer to a dev_info_t structure, defining the HBA driver instance.																		
<i>ap</i>	Pointer to a scsi_address(9S) structure, defining the target instance.																		
<i>cmdlen</i>	Length in bytes to be allocated for the SCSI command descriptor block (CDB).																		
<i>statuslen</i>	Length in bytes to be allocated for the SCSI status completion block (SCB).																		
<i>tgtlen</i>	Length in bytes to be allocated for a private data area for the target driver's exclusive use.																		
<i>hbalen</i>	Length in bytes to be allocated for a private data area for the HBA driver's exclusive use.																		
<i>callback</i>	Indicates what scsi_hba_pkt_alloc() should do when resources are not available: NULL_FUNC Do not wait for resources. Return a NULL pointer. SLEEP_FUNC Wait indefinitely for resources.																		
<i>arg</i>	Must be NULL.																		
<i>pkt</i>	A pointer to a scsi_pkt(9S) structure.																		
DESCRIPTION	For scsi_hba_pkt_alloc(): scsi_hba_pkt_alloc() allocates space for a scsi_pkt structure. HBA drivers must use this interface when allocating a scsi_pkt from their tran_init_pkt(9E) entry point. If <i>callback</i> is NULL_FUNC, scsi_hba_pkt_alloc() may not sleep when allocating resources, and callers should be prepared to deal with allocation failures. scsi_hba_pkt_alloc() copies the scsi_address(9S) structure pointed to by <i>ap</i> to the pkt_address field in the scsi_pkt(9S).																		

scsi_hba_pkt_alloc(9F)

`scsi_hba_pkt_alloc()` also allocates memory for these `scsi_pkt(9S)` data areas, and sets these fields to point to the allocated memory:

<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.

For `scsi_hba_pkt_free()`:

`scsi_hba_pkt_free()` frees the space allocated for the `scsi_pkt(9S)` structure.

RETURN VALUES `scsi_hba_pkt_alloc()` returns a pointer to the `scsi_pkt` structure, or NULL if no space is available.

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

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

SEE ALSO `tran_init_pkt(9E)`, `scsi_address(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

NAME	scsi_hba_probe – default SCSI HBA probe function
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_hba_probe(struct scsi_device *sd, int (*waitfunc)(void));</pre>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
PARAMETERS	<i>sd</i> Pointer to a scsi_device(9S) structure describing the target. <i>waitfunc</i> NULL_FUNC or SLEEP_FUNC.
DESCRIPTION	scsi_hba_probe() is a function providing the semantics of scsi_probe(9F). An HBA driver may call scsi_hba_probe() from its tran_tgt_probe(9E) entry point, to probe for the existence of a target on the SCSI bus, or the HBA may set tran_tgt_probe(9E) to point to scsi_hba_probe directly.
RETURN VALUES	See scsi_probe(9F) for the return values from scsi_hba_probe().
CONTEXT	scsi_hba_probe() should only be called from the HBA's tran_tgt_probe(9E) entry point.
SEE ALSO	tran_tgt_probe(9E), scsi_probe(9F), scsi_device(9S) <i>Writing Device Drivers</i>

scsi_hba_tran_alloc(9F)

NAME	scsi_hba_tran_alloc, scsi_hba_tran_free – allocate and free transport structures
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> scsi_hba_tran_t *scsi_hba_tran_alloc(dev_info_t *dip, int flags); void scsi_hba_tran_free(scsi_hba_tran_t *hba_tran);</pre>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
PARAMETERS	<i>dip</i> Pointer to a dev_info structure, defining the HBA driver instance. <i>flag</i> Flag modifiers. The only possible flag value is SCSI_HBA_CANSLEEP (memory allocation may sleep). <i>hba_tran</i> Pointer to a scsi_hba_tran(9S) structure.
DESCRIPTION	For <code>scsi_hba_tran_alloc()</code>: <code>scsi_hba_tran_alloc()</code> allocates a <code>scsi_hba_tran(9S)</code> structure for a HBA driver. The HBA must use this structure to register its transport vectors with the system by using <code>scsi_hba_attach_setup(9F)</code>. If the flag <code>SCSI_HBA_CANSLEEP</code> is set in <i>flags</i>, <code>scsi_hba_tran_alloc()</code> may sleep when allocating resources; otherwise it may not sleep, and callers should be prepared to deal with allocation failures. For <code>scsi_hba_tran_free()</code>: <code>scsi_hba_tran_free()</code> is used to free the <code>scsi_hba_tran(9S)</code> structure allocated by <code>scsi_hba_tran_alloc()</code>.
RETURN VALUES	<code>scsi_hba_tran_alloc()</code> returns a pointer to the allocated transport structure, or NULL if no space is available.
CONTEXT	<code>scsi_hba_tran_alloc()</code> can be called from user or interrupt context. Drivers must not allow <code>scsi_hba_tran_alloc()</code> to sleep if called from an interrupt routine. <code>scsi_hba_tran_free()</code> can be called from user or interrupt context.
SEE ALSO	<code>scsi_hba_attach_setup(9F)</code>, <code>scsi_hba_tran(9S)</code> <i>Writing Device Drivers</i>

NAME	scsi_ifgetcap, scsi_ifsetcap – get/set SCSI transport capability																		
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> 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>																		
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>ap</i> Pointer to the <code>scsi_address</code> structure. <i>cap</i> Pointer to the string capability identifier. <i>value</i> Defines the new state of the capability. <i>whom</i> Determines if all targets or only the specified target is affected.																		
DESCRIPTION	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. 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. The following capabilities have been defined: <table> <tr> <td><code>dma-max</code></td><td>Maximum dma transfer size supported by host adapter.</td></tr> <tr> <td><code>msg-out</code></td><td>Message out capability supported by host adapter: 0 disables, 1 enables.</td></tr> <tr> <td><code>disconnect</code></td><td>Disconnect capability supported by host adapter: 0 disables, 1 enables.</td></tr> <tr> <td><code>synchronous</code></td><td>Synchronous data transfer capability supported by host adapter: 0 disables, 1 enables.</td></tr> <tr> <td><code>wide-xfer</code></td><td>Wide transfer capability supported by host adapter: 0 disables, 1 enables.</td></tr> <tr> <td><code>parity</code></td><td>Parity checking by host adapter: 0 disables, 1 enables.</td></tr> <tr> <td><code>initiator-id</code></td><td>The host's bus address is returned.</td></tr> <tr> <td><code>untagged-qing</code></td><td>The host adapter's capability to support internal queueing of commands without tagged queueing: 0 disables, 1 enables.</td></tr> <tr> <td><code>tagged-qing</code></td><td>The host adapter's capability to support tagged queueing: 0 disables, 1 enables.</td></tr> </table>	<code>dma-max</code>	Maximum dma transfer size supported by host adapter.	<code>msg-out</code>	Message out capability supported by host adapter: 0 disables, 1 enables.	<code>disconnect</code>	Disconnect capability supported by host adapter: 0 disables, 1 enables.	<code>synchronous</code>	Synchronous data transfer capability supported by host adapter: 0 disables, 1 enables.	<code>wide-xfer</code>	Wide transfer capability supported by host adapter: 0 disables, 1 enables.	<code>parity</code>	Parity checking by host adapter: 0 disables, 1 enables.	<code>initiator-id</code>	The host's bus address is returned.	<code>untagged-qing</code>	The host adapter's capability to support internal queueing of commands without tagged queueing: 0 disables, 1 enables.	<code>tagged-qing</code>	The host adapter's capability to support tagged queueing: 0 disables, 1 enables.
<code>dma-max</code>	Maximum dma transfer size supported by host adapter.																		
<code>msg-out</code>	Message out capability supported by host adapter: 0 disables, 1 enables.																		
<code>disconnect</code>	Disconnect capability supported by host adapter: 0 disables, 1 enables.																		
<code>synchronous</code>	Synchronous data transfer capability supported by host adapter: 0 disables, 1 enables.																		
<code>wide-xfer</code>	Wide transfer capability supported by host adapter: 0 disables, 1 enables.																		
<code>parity</code>	Parity checking by host adapter: 0 disables, 1 enables.																		
<code>initiator-id</code>	The host's bus address is returned.																		
<code>untagged-qing</code>	The host adapter's capability to support internal queueing of commands without tagged queueing: 0 disables, 1 enables.																		
<code>tagged-qing</code>	The host adapter's capability to support tagged queueing: 0 disables, 1 enables.																		

scsi_ifgetcap(9F)

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." If geometry is not relevant or appropriate for this target disk, because (for example) the HBA BIOS supports Logical Block Addressing for this drive, it is acceptable for <code>scsi_ifgetcap()</code> to return -1, indicating that the geometry is not defined. This will cause failure of attempts to retrieve the "virtual geometry" from the target driver (the <code>DKIOCG_VIRTGEOM</code> ioctl will fail). See <code>dkio(7I)</code> for more information about <code>DKIOCG_VIRTGEOM</code>.
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 QUEUE FULL handling. If 0, the HBA will not retry a command when a QUEUE FULL 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 <code>qfull-retries</code> is 0-255.

	<code>qfull-retry-interval</code> This capability sets the retry interval (in ms) for commands that were completed with a <code>QUEUE FULL</code> status. The range for <code>qfull-retry-intervals</code> is 0-1000 ms. <code>lun-reset</code> This capability is created with a value of zero by HBA drivers that support the <code>RESET_LUN</code> flag in their <code>tran_reset(9E)</code> routine. If it exists, its value can be set to 1 by target drivers that want to allow use of LOGICAL UNIT RESET on a specific target instance. If <code>lun-reset</code> does not exist or has a value of zero, <code>scsi_reset(9F)</code> is prevented from passing the <code>RESET_LUN</code> flag to the HBA driver's <code>tran_reset()</code> routine. If <code>lun-reset</code> exists and has a value of 1, then the HBA driver's <code>tran_reset()</code> routine can be called with the <code>RESET_LUN</code> flag.
RETURN VALUES	<code>scsi_ifsetcap()</code> 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. <code>scsi_ifgetcap()</code> 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 <code>scsi_ifgetcap()</code> <pre> if (scsi_ifgetcap(&sd->sd_address, "auto-rqsense", 1) == 1) { un->un_arq_enabled = 1; } else { un->un_arq_enabled = ((scsi_ifsetcap(&sd->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; un->un_throttle = MAX_THROTTLE; } else if (scsi_ifgetcap(&devp->sd_address, "untagged-qing", 0) == 1) { un->un_dp->options = SD_QUEUEING; un->un_throttle = 3; } else { un->un_dp->options &= ~SD_QUEUEING; un->un_throttle = 1; } </pre>

scsi_ifgetcap(9F)

EXAMPLE 1 Using `scsi_ifgetcap()` *(Continued)*

SEE ALSO `tran_reset(9E)`, `scsi_hba_lookup_capstr(9F)`, `scsi_reset(9F)`,
`scsi_reset_notify(9F)`, `ddi_dma_lim_sparc(9S)`, `ddi_dma_lim_x86(9S)`,
`scsi_address(9S)`, `scsi_arq_status(9S)`

Writing Device Drivers

NAME	scsi_init_pkt – prepare a complete SCSI packet
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> struct scsi_pkt *scsi_init_pkt(struct scsi_address *ap, struct scsi_pkt *pkt, struct buf *bp, int cmdlen, int statuslen, int privatenelen, int flags, int (*callback)(caddr_t), caddr_t arg);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>ap</i> Pointer to a scsi_address(9S) structure. <i>pkt</i> A pointer to a scsi_pkt(9S) structure. <i>bp</i> Pointer to a buf(9S) structure. <i>cmdlen</i> The required length for the SCSI command descriptor block (CDB) in bytes. <i>statuslen</i> The required length for the SCSI status completion block (SCB) in bytes. Valid values are: 0 No status back. 1 Return SCSI status byte. sizeof(scsi_arq_status) Return status information in a scsi_arq_status structure. This will include up to 20 bytes of sense data. Please refer to scsi_arq_status(9S) for more information. EXTCMDS_STATUS_SIZE Same as preceding. <i>privatenelen</i> The required length for the <i>pkt_private</i> area. <i>flags</i> Flags modifier. <i>callback</i> A pointer to a callback function, NULL_FUNC, or SLEEP_FUNC. <i>arg</i> The <i>callback</i> function argument.
DESCRIPTION	Target drivers use <code>scsi_init_pkt()</code> 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 NULL, a new <code>scsi_pkt(9S)</code> is allocated using the HBA driver's packet allocator. The <i>bp</i> is a pointer to a <code>buf(9S)</code> structure. If <i>bp</i> is non-NULL and contains a valid byte count, the <code>buf(9S)</code> structure is also set up for DMA transfer using the HBA driver DMA resources allocator. When <i>bp</i> is allocated by <code>scsi_alloc_consistent_buf(9F)</code>, the <code>PKT_CONSISTENT</code> bit must be set in the <i>flags</i> argument to ensure proper operation. If <i>privatenelen</i> is non-zero then additional

scsi_init_pkt(9F)

space is allocated for the *pkt_private* area of the *scsi_pkt*(9S). On return *pkt_private* points to this additional space. Otherwise *pkt_private* is a pointer that is typically used to store the *bp* during execution of the command. In this case *pkt_private* is NULL on return.

The *flags* argument is a set of bit flags. Possible bits include:

PKT_CONSISTENT

This must be set if the DMA buffer was allocated using *scsi_alloc_consistent_buf*(9F). 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, *scsi_init_pkt*() will allocate DMA resources with the *DDI_DMA_PARTIAL* bit set in the *dmr_flag* element of the *ddi_dma_req*(9S) structure. The *pkt_resid* field of the *scsi_pkt*(9S) structure may be returned with a non-zero value, which indicates the number of bytes for which *scsi_init_pkt*() was unable to allocate DMA resources. In this case, a subsequent call to *scsi_init_pkt*() may be made for the same *pktp* and *bp* to adjust the DMA resources to the next portion of the transfer. This sequence should be repeated until the *pkt_resid* 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	<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 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,
    1, 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,
    sizeof(struct scsi_arq_status), 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,
    sizeof(struct scsi_arq_status), 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,
    1, sizeof (struct buf *), PKT_DMA_PARTIAL,
    SLEEP_FUNC, NULL);
```

scsi_init_pkt(9F)

EXAMPLE 5 Allocating a Packet with Partial DMA Resources Attached (Continued)

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`.

NAME	scsi_log – display a SCSI-device-related message								
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> #include <sys/cmn_err.h> void scsi_log(dev_info_t *dip, char *drv_name, uint_t level, const char *fmt, ...);</pre>								
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <table> <tr> <td><i>dip</i></td><td>Pointer to the dev_info structure.</td></tr> <tr> <td><i>drv_name</i></td><td>String naming the device.</td></tr> <tr> <td><i>level</i></td><td>Error level.</td></tr> <tr> <td><i>fmt</i></td><td>Display format.</td></tr> </table>	<i>dip</i>	Pointer to the dev_info structure.	<i>drv_name</i>	String naming the device.	<i>level</i>	Error level.	<i>fmt</i>	Display format.
<i>dip</i>	Pointer to the dev_info structure.								
<i>drv_name</i>	String naming the device.								
<i>level</i>	Error level.								
<i>fmt</i>	Display format.								
DESCRIPTION	<code>scsi_log()</code> is a utility function that displays a message via the <code>cmn_err(9F)</code> routine. The error levels that can be passed in to this function are <code>CE_PANIC</code>, <code>CE_WARN</code>, <code>CE_NOTE</code>, <code>CE_CONT</code>, and <code>SCSI_DEBUG</code>. 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 <code>sd</code> and <code>cmdk</code>. If the <i>dev_info_t</i> pointer is <code>NULL</code>, then the <i>drv_name</i> will be used with no unit or long name. If the first character in <i>format</i> is: ■ An exclamation mark (!), the message goes only to the system buffer. ■ A caret (^), the message goes only to the console. ■ A question mark (?) and <i>level</i> is <code>CE_CONT</code>, the message is always sent to the system buffer, but is written to the console only when the system has been booted in verbose mode. See <code>kernel(1M)</code>. If neither condition is met, the ? character has no effect and is simply ignored. All formatting conversions in use by <code>cmn_err()</code> also work with <code>scsi_log()</code>.								
CONTEXT	<code>scsi_log()</code> may be called from user or interrupt context.								
EXAMPLES	<pre>scsi_log(dev, "Disk Unit ", CE_PANIC, "Bad Value %d\n", foo);</pre> generates: <pre>PANIC: /eisa/aha@330,0/cmdk@0,0 (Disk Unit 0): Bad Value 5</pre> This is followed by a PANIC. <pre>scsi_log(dev, "sd", CE_WARN, "Label Bad\n");</pre> generates: <pre>WARNING: /sbus@1,f8000000/esp@0,80000000/sd@1,0 (sd1): Label Bad</pre> <pre>scsi_log((dev_info_t *) NULL, "Disk Unit ", CE_NOTE, "Disk Ejected\n");</pre> generates: <pre>Disk Unit: Disk Ejected</pre>								

scsi_log(9F)

```
scsi_log(cmdk_unit, "Disk Unit ", CE_CONT, "Disk Inserted\n");  
generates:
```

```
Disk Inserted
```

```
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

NAME	scsi_pktalloc, scsi_resalloc, scsi_pktfree, scsi_resfree – SCSI packet utility routines
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> struct scsi_pkt *scsi_pktalloc(struct scsi_address*<i>ap</i>, int<i>cmdlen</i>, int<i>statuslen</i>, int (*<i>callback</i>) (void)); struct scsi_pkt *scsi_resalloc(struct scsi_address*<i>ap</i>, int<i>cmdlen</i>, int<i>statuslen</i>, opaque_t<i>dmationToken</i>, int (*<i>callback</i>) (void)); voidscsi_pktfree(struct scsi_pkt*<i>pkt</i>); voidscsi_resfree(struct scsi_pkt*<i>pkt</i>);</pre>
INTERFACE LEVEL	The <code>scsi_pktalloc()</code> , <code>scsi_pktfree()</code> , <code>scsi_resalloc()</code> , and <code>scsi_resfree()</code> functions are obsolete. The <code>scsi_pktalloc()</code> and <code>scsi_resalloc()</code> functions have been replaced by <code>scsi_init_pkt(9F)</code> . The <code>scsi_pktfree()</code> and <code>scsi_resfree()</code> functions have been replaced by <code>scsi_destroy_pkt(9F)</code> .
PARAMETERS	<i>ap</i> Pointer to a <code>scsi_address</code> structure. <i>cmdlen</i> The required length for the SCSI command descriptor block (CDB) in bytes. <i>statuslen</i> The required length for the SCSI status completion block (SCB) in bytes. <i>dmationToken</i> Pointer to an implementation-dependent object. <i>callback</i> A pointer to a callback function, or <code>NULL_FUNC</code> or <code>SLEEP_FUNC</code>. <i>pkt</i> Pointer to a <code>scsi_pkt(9S)</code> structure.
DESCRIPTION	<code>scsi_pktalloc()</code> requests the host adapter driver to allocate a command packet. For commands that have a data transfer associated with them, <code>scsi_resalloc()</code> should be used. <i>ap</i> is a pointer to a <code>scsi_address</code> structure. Allocator routines use it to determine the associated host adapter. <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. <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. <i>dmationToken</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>dmationToken</i> must be a pointer to a <code>buf(9S)</code> structure. If <i>dmationToken</i> is <code>NULL</code>, 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

scsi_pktalloc(9F)

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_DEERR`. *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.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `attributes(5)`, `scsi_dmafree(9F)`, `scsi_dmaget(9F)`, `buf(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

NOTES The `scsi_pktalloc()`, `scsi_pktfree()`, `scsi_resalloc()`, and `scsi_resfree()` functions are obsolete and will be discontinued in a future release. The `scsi_pktalloc()` and `scsi_resalloc()` functions have been replaced by `scsi_init_pkt(9F)`. The `scsi_pktfree()` and `scsi_resfree()` functions have been replaced by `scsi_destroy_pkt(9F)`.

NAME	scsi_poll – run a polled SCSI command on behalf of a target driver
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_poll(struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>pkt</i> Pointer to the scsi_pkt(9S) structure.
DESCRIPTION	scsi_poll() requests the host adapter driver to run a polled command. Unlike scsi_transport(9F) which runs commands asynchronously, scsi_poll() runs commands to completion before returning. If the pkt_time member of <i>pkt</i> is 0, the value of pkt_time is defaulted to SCSI_POLL_TIMEOUT to prevent an indefinite hang of the system.
RETURN VALUES	scsi_poll() returns: 0 command completed successfully. -1 command failed.
CONTEXT	scsi_poll() can be called from user or interrupt level. This function should not be called when the caller is executing timeout(9F) in the context of a thread.
SEE ALSO	makecom(9F), scsi_transport(9F), scsi_pkt(9S)
	<i>Writing Device Drivers</i>
WARNINGS	Since scsi_poll() runs commands to completion before returning, it may require more time than is desirable when called from interrupt context. Therefore, calling scsi_poll from interrupt context is not recommended.

scsi_probe(9F)

NAME	scsi_probe – utility for probing a scsi device														
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_probe(struct scsi_device *devp, int (*waitfunc);</pre>														
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).														
PARAMETERS	<i>devp</i> Pointer to a scsi_device(9S) structure <i>waitfunc</i> NULL_FUNC or SLEEP_FUNC														
DESCRIPTION	scsi_probe() determines whether a target/lun is present and sets up the scsi_device structure with inquiry data. scsi_probe() uses the SCSI Inquiry command to test if the device exists. It can retry the Inquiry command as appropriate. If scsi_probe() is successful, it will allocate space for the scsi_inquiry structure and assign the address to the sd_inq member of the scsi_device(9S) structure. scsi_probe() will then fill in this scsi_inquiry(9S) structure and return SCSI_PROBE_EXISTS. If scsi_probe() is unsuccessful, it returns SCSI_PROBE_NOMEM in spite of callback set to SLEEP_FUNC. scsi_unprobe(9F) is used to undo the effect of scsi_probe(). If the target is a non-CCS device, SCSI_PROBE_NONCCS will be returned. <i>waitfunc</i> indicates what the allocator routines should do when resources are not available; the valid values are: NULL_FUNC Do not wait for resources. Return SCSI_PROBE_NOMEM or SCSI_PROBE_FAILURE SLEEP_FUNC Wait indefinitely for resources.														
RETURN VALUES	scsi_probe() returns: <table> <tr> <td>SCSI_PROBE_BUSY</td><td>Device exists but is currently busy.</td></tr> <tr> <td>SCSI_PROBE_EXISTS</td><td>Device exists and inquiry data is valid.</td></tr> <tr> <td>SCSI_PROBE_FAILURE</td><td>Polled command failure.</td></tr> <tr> <td>SCSI_PROBE_NOMEM</td><td>No space available for structures.</td></tr> <tr> <td>SCSI_PROBE_NOMEM_CB</td><td>No space available for structures but callback request has been queued.</td></tr> <tr> <td>SCSI_PROBE_NONCCS</td><td>Device exists but inquiry data is not valid.</td></tr> <tr> <td>SCSI_PROBE_NORESP</td><td>Device does not respond to an INQUIRY.</td></tr> </table>	SCSI_PROBE_BUSY	Device exists but is currently busy.	SCSI_PROBE_EXISTS	Device exists and inquiry data is valid.	SCSI_PROBE_FAILURE	Polled command failure.	SCSI_PROBE_NOMEM	No space available for structures.	SCSI_PROBE_NOMEM_CB	No space available for structures but callback request has been queued.	SCSI_PROBE_NONCCS	Device exists but inquiry data is not valid.	SCSI_PROBE_NORESP	Device does not respond to an INQUIRY.
SCSI_PROBE_BUSY	Device exists but is currently busy.														
SCSI_PROBE_EXISTS	Device exists and inquiry data is valid.														
SCSI_PROBE_FAILURE	Polled command failure.														
SCSI_PROBE_NOMEM	No space available for structures.														
SCSI_PROBE_NOMEM_CB	No space available for structures but callback request has been queued.														
SCSI_PROBE_NONCCS	Device exists but inquiry data is not valid.														
SCSI_PROBE_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. In any case, this routine should not be called from interrupt context, because it can sleep waiting for memory to be allocated.														

EXAMPLES**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.

scsi_reset(9F)

NAME	scsi_reset – reset a SCSI bus or target
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_reset(struct scsi_address *ap, int level);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>ap</i> Pointer to the <code>scsi_address</code> structure. <i>level</i> The level of reset required.
DESCRIPTION	<code>scsi_reset()</code> 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 <code>RESET_ALL</code>, the SCSI bus is reset. If it equals <code>RESET_TARGET</code>, <i>ap</i> is used to determine the target to be reset. If it equals <code>RESET_LUN</code>, <i>ap</i> is used to determine the logical unit to be reset. When given the <code>RESET_LUN</code> level, <code>scsi_reset()</code> can return failure if the LOGICAL UNIT RESET message is not supported by the target device, or if the underlying HBA driver does not implement the ability to issue a LOGICAL UNIT RESET message. Note that, at the point when <code>scsi_reset()</code> resets the logical unit (case <code>RESET_LUN</code>), or the target (case <code>RESET_TARGET</code>), or the bus (case <code>RESET_ALL</code>), there might be one or more command packets outstanding. That is, packets have been passed to <code>scsi_transport()</code>, and queued or possibly transported, but the commands have not been completed and the target completion routine has not been called for those packets. The successful call to <code>scsi_reset()</code> has the side effect that any such commands currently outstanding are aborted, at which point the packets are marked with <code>pkt_reason</code> set to <code>CMD_RESET</code>, and the appropriate bit -- either <code>STAT_BUS_RESET</code> or <code>STAT_DEV_RESET</code> -- is set in <code>pkt_statistics</code>. Once thus appropriately marked, the aborted command packets are passed to the target driver command completion routine. Also note that, at the moment that a thread executing <code>scsi_reset()</code> actually resets the target or the bus, it is possible that a second thread may have already called <code>scsi_transport()</code>, but not yet queued or transported its command. In this case the HBA will not yet have received the second thread's packet and this packet will not be aborted.
RETURN VALUES	<code>scsi_reset()</code> returns: 1 Upon success. 0 Upon failure.
CONTEXT	<code>scsi_reset()</code> can be called from user or interrupt context.
SEE ALSO	<code>tran_reset(9E)</code>, <code>tran_reset_notify(9E)</code>, <code>scsi_abort(9F)</code> <i>Writing Device Drivers</i>

NAME	scsi_reset_notify – notify target driver of bus resets								
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void scsi_reset_notify(struct scsi_address *ap, int flag, void (*callback) (caddr_t), caddr_t arg);</pre>								
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>ap</i> Pointer to the <code>scsi_address</code> structure. <i>flag</i> A flag indicating registration or cancellation of the notification request. <i>callback</i> A pointer to the target driver's reset notification function. <i>arg</i> The callback function argument.								
DESCRIPTION	<code>scsi_reset_notify()</code> 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. The argument <i>flag</i> is used to register or cancel the notification. The supported values for <i>flag</i> are as follows: <table> <tr> <td>SCSI_RESET_NOTIFY</td><td>Register <i>callback</i> as the reset notification function for the target driver.</td></tr> <tr> <td>SCSI_RESET_CANCEL</td><td>Cancel the reset notification request.</td></tr> </table> Target drivers can find out whether the HBA driver and controller support reset notification by checking the reset-notification capability using the <code>scsi_ifgetcap(9F)</code> function.	SCSI_RESET_NOTIFY	Register <i>callback</i> as the reset notification function for the target driver.	SCSI_RESET_CANCEL	Cancel the reset notification request.				
SCSI_RESET_NOTIFY	Register <i>callback</i> as the reset notification function for the target driver.								
SCSI_RESET_CANCEL	Cancel the reset notification request.								
RETURN VALUES	If <i>flag</i> is <code>SCSI_RESET_NOTIFY</code>, <code>scsi_reset_notify()</code> returns: <table> <tr> <td>DDI_SUCCESS</td><td>The notification request has been accepted.</td></tr> <tr> <td>DDI_FAILURE</td><td>The transport layer does not support reset notification or could not accept this request.</td></tr> </table> If <i>flag</i> is <code>SCSI_RESET_CANCEL</code>, <code>scsi_reset_notify()</code> returns: <table> <tr> <td>DDI_SUCCESS</td><td>The notification request has been canceled.</td></tr> <tr> <td>DDI_FAILURE</td><td>No notification request was registered.</td></tr> </table>	DDI_SUCCESS	The notification request has been accepted.	DDI_FAILURE	The transport layer does not support reset notification or could not accept this request.	DDI_SUCCESS	The notification request has been canceled.	DDI_FAILURE	No notification request was registered.
DDI_SUCCESS	The notification request has been accepted.								
DDI_FAILURE	The transport layer does not support reset notification or could not accept this request.								
DDI_SUCCESS	The notification request has been canceled.								
DDI_FAILURE	No notification request was registered.								
CONTEXT	<code>scsi_reset_notify()</code> can be called from user or interrupt context.								
SEE ALSO	<code>scsi_address(9S)</code>, <code>scsi_ifgetcap(9F)</code> <i>Writing Device Drivers</i>								

scsi_setup_cdb(9F)

NAME	scsi_setup_cdb – setup SCSI command descriptor block (CDB)										
SYNOPSIS	<pre>int scsi_setup_cdb(union scsi_cdb *cdbp, uchar_t cmd, uint_t addr, uint_t cnt, uint_t othr_cdb_data);</pre>										
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).										
PARAMETERS	<table> <tr> <td><i>cdbp</i></td><td>Pointer to command descriptor block.</td></tr> <tr> <td><i>cmd</i></td><td>The first byte of the SCSI group 0, 1, 2, 4, or 5 CDB.</td></tr> <tr> <td><i>addr</i></td><td>Pointer to the location of the data.</td></tr> <tr> <td><i>cnt</i></td><td>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.</td></tr> <tr> <td><i>othr_cdb_data</i></td><td>Additional CDB data.</td></tr> </table>	<i>cdbp</i>	Pointer to command descriptor block.	<i>cmd</i>	The first byte of the SCSI group 0, 1, 2, 4, or 5 CDB.	<i>addr</i>	Pointer to the location of the data.	<i>cnt</i>	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.	<i>othr_cdb_data</i>	Additional CDB data.
<i>cdbp</i>	Pointer to command descriptor block.										
<i>cmd</i>	The first byte of the SCSI group 0, 1, 2, 4, or 5 CDB.										
<i>addr</i>	Pointer to the location of the data.										
<i>cnt</i>	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.										
<i>othr_cdb_data</i>	Additional CDB data.										
DESCRIPTION	<code>scsi_setup_cdb()</code> 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>. <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. <code>scsi_setup_cdb()</code> function does not set the LUN bits in CDB[1] as the <code>makecom(9F)</code> functions do. Also, the fixed bit for sequential access device commands is not set.										
RETURN VALUES	<code>scsi_setup_cdb()</code> returns: <table> <tr> <td>1</td><td>Upon success.</td></tr> <tr> <td>0</td><td>Upon failure.</td></tr> </table>	1	Upon success.	0	Upon failure.						
1	Upon success.										
0	Upon failure.										
CONTEXT	These functions can be called from a user or interrupt context.										
SEE ALSO	<code>makecom(9F)</code>, <code>scsi_pkt(9S)</code> <i>Writing Device Drivers</i> <i>American National Standard Small Computer System Interface-2 (SCSI-2)</i> <i>American National Standard SCSI-3 Primary Commands (SPC)</i>										

NAME	scsi_slave – utility for SCSI target drivers to establish the presence of a target											
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_slave(struct scsi_device *devp, int (*callback)(void));</pre>											
INTERFACE LEVEL	The scsi_slave() function is obsolete. This function has been replaced by scsi_probe(9F).											
PARAMETERS	devp	Pointer to a scsi_device(9S) structure.										
	callback	Pointer to a callback function, NULL_FUNC or SLEEP_FUNC.										
DESCRIPTION	scsi_slave() checks for the presence of a SCSI device. Target drivers may use this function in their probe(9E) routines. scsi_slave() determines if the device is present by using a Test Unit Ready command followed by an Inquiry command. If scsi_slave() is successful, it will fill in the scsi_inquiry structure, which is the sd_inq member of the scsi_device(9S) structure, and return SCSI_PROBE_EXISTS. This information can be used to determine if the target driver has probed the correct SCSI device type. callback indicates what the allocator routines should do when DMA resources are not available: <table><tr><td>NULL_FUNC</td><td>Do not wait for resources. Return a NULL pointer.</td></tr><tr><td>SLEEP_FUNC</td><td>Wait indefinitely for resources.</td></tr><tr><td>Other Values</td><td>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.</td></tr></table>		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.				
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.											
RETURN VALUES	scsi_slave() returns: <table><tr><td>SCSI_PROBE_NOMEM</td><td>No space available for structures.</td></tr><tr><td>SCSI_PROBE_EXISTS</td><td>Device exists and inquiry data is valid.</td></tr><tr><td>SCSI_PROBE_NONCCS</td><td>Device exists but inquiry data is not valid.</td></tr><tr><td>SCSI_PROBE_FAILURE</td><td>Polled command failure.</td></tr><tr><td>SCSI_PROBE_NORESP</td><td>No response to TEST UNIT READY.</td></tr></table>		SCSI_PROBE_NOMEM	No space available for structures.	SCSI_PROBE_EXISTS	Device exists and inquiry data is valid.	SCSI_PROBE_NONCCS	Device exists but inquiry data is not valid.	SCSI_PROBE_FAILURE	Polled command failure.	SCSI_PROBE_NORESP	No response to TEST UNIT READY.
SCSI_PROBE_NOMEM	No space available for structures.											
SCSI_PROBE_EXISTS	Device exists and inquiry data is valid.											
SCSI_PROBE_NONCCS	Device exists but inquiry data is not valid.											
SCSI_PROBE_FAILURE	Polled command failure.											
SCSI_PROBE_NORESP	No response to TEST UNIT READY.											
CONTEXT	scsi_slave() is normally called from the target driver's probe(9E) or attach(9E) routine. In any case, this routine should not be called from interrupt context, because it can sleep waiting for memory to be allocated.											
ATTRIBUTES	See attributes(5) for a description of the following attributes:											

scsi_slave(9F)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO attributes(5), 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

NOTES The `scsi_slave()` function is obsolete and will be discontinued in a future release. This function has been replaced by `scsi_probe(9F)`.

NAME	scsi_sync_pkt – synchronize CPU and I/O views of memory
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void scsi_sync_pkt(struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>pkt</i> Pointer to a scsi_pkt(9S) structure.
DESCRIPTION	<code>scsi_sync_pkt()</code> 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. This function need only be called under certain circumstances. When a SCSI packet is mapped for I/O using <code>scsi_init_pkt(9F)</code> and destroyed using <code>scsi_destroy_pkt(9F)</code>, then an implicit <code>scsi_sync_pkt()</code> will be performed. However, if the memory object has been modified by either the device or a CPU after the mapping by <code>scsi_init_pkt(9F)</code>, then a call to <code>scsi_sync_pkt()</code> is required. If the same <code>scsi_pkt</code> is reused for a data transfer from memory to a device, then <code>scsi_sync_pkt()</code> must be called before calling <code>scsi_transport(9F)</code>. If the same packet is reused for a data transfer from a device to memory <code>scsi_sync_pkt()</code> must be called after the completion of the packet but before accessing the data in memory.
CONTEXT	<code>scsi_sync_pkt()</code> may be called from user or interrupt context.
SEE ALSO	<code>tran_sync_pkt(9E)</code>, <code>ddi_dma_sync(9F)</code>, <code>scsi_destroy_pkt(9F)</code>, <code>scsi_init_pkt(9F)</code>, <code>scsi_transport(9F)</code>, <code>scsi_pkt(9S)</code> <i>Writing Device Drivers</i>

scsi_transport(9F)

NAME	scsi_transport – request by a SCSI target driver to start a command								
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int scsi_transport(struct scsi_pkt *pkt);</pre>								
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).								
PARAMETERS	<i>pkt</i> Pointer to a scsi_pkt(9S) structure.								
DESCRIPTION	Target drivers use <code>scsi_transport()</code> 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 <code>scsi_init_pkt(9F)</code> prior to calling this function. The packet may be initialized using one of the <code>makecom(9F)</code> functions. <code>scsi_transport()</code> does not wait for the SCSI command to complete. See <code>scsi_poll(9F)</code> 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 <code>scsi_pkt</code> pointed to by <i>pkt</i>.								
RETURN VALUES	<code>scsi_transport()</code> returns: <table> <tr> <td>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.								
CONTEXT	<code>scsi_transport()</code> can be called from user or interrupt context.								
EXAMPLES	EXAMPLE 1 Using <code>scsi_transport()</code> <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	<code>tran_start(9E)</code>, <code>makecom(9F)</code>, <code>scsi_init_pkt(9F)</code>, <code>scsi_pktalloc(9F)</code>, <code>scsi_poll(9F)</code>, <code>scsi_pkt(9S)</code> <i>Writing Device Drivers</i>								

NAME	scsi_unprobe, scsi_unslave – free resources allocated during initial probing
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void scsi_unslave(struct scsi_device *dev); void scsi_unprobe(struct scsi_device *dev);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). The <code>scsi_unslave()</code> interface is obsolete. Use <code>scsi_unprobe()</code> instead.
PARAMETERS	<i>dev</i> Pointer to a <code>scsi_device(9S)</code> structure.
DESCRIPTION	<code>scsi_unprobe()</code> and <code>scsi_unslave()</code> are used to free any resources that were allocated on the driver's behalf during <code>scsi_slave(9F)</code> and <code>scsi_probe(9F)</code> activity.
CONTEXT	<code>scsi_unprobe()</code> and <code>scsi_unslave()</code> must not be called from an interrupt context.
SEE ALSO	<code>scsi_probe(9F)</code> , <code>scsi_slave(9F)</code> , <code>scsi_device(9S)</code> <i>Writing Device Drivers</i>
NOTES	The <code>scsi_unslave()</code> function is obsolete and will be discontinued in a future release. This function has been replaced by <code>scsi_unprobe()</code> .

scsi_vu_errmsg(9F)

NAME	scsi_vu_errmsg – display a SCSI request sense message
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void scsi_vu_errmsg(struct scsi_pkt *pkt, char *drv_name, int severity, int err_blkno, struct scsi_key_strings *cmdlist, struct scsi_extended_sense *sensep, struct scsi_asc_key_strings *asc_list, char **decode_fru, struct scsi_device*, char *, int, char);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	The following parameters are supported: <i>devp</i> Pointer to the <code>scsi_device(9S)</code> structure. <i>pkt</i> Pointer to a <code>scsi_pkt(9S)</code> structure. <i>drv_name</i> String used by <code>scsi_log(9F)</code>. <i>severity</i> Error severity level, maps to severity strings below. <i>blkno</i> Requested block number. <i>err_blkno</i> Error block number. <i>cmdlist</i> An array of SCSI command description strings. <i>sensep</i> A pointer to a <code>scsi_extended_sense(9S)</code> structure. <i>asc_list</i> A pointer to a array of asc and ascq message list. The list must be terminated with -1 asc value. <i>decode_fru</i> This is a function pointer that will be called after the entire sense information has been decoded. The parameters will be the <code>scsi_device</code> 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. <code>decode_fru</code> might be NULL if no special decoding is required. <code>decode_fru</code> is expected to return pointer to a char string if decoding possible and NULL if no decoding is possible.
DESCRIPTION	This function is very similar to <code>scsi_errmsg(9F)</code> 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",
```

scsi_vu_errmsg(9F)

```
        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_asc_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_list,
               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

NAME	semaphore, sema_init, sema_destroy, sema_p, sema_p_sig, sema_v, sema_try_p – semaphore functions
SYNOPSIS	<pre>#include <sys/ksynch.h> void sema_init(ksema_t *sp, uint_t val, char *name, ksema_type_t type, void *arg); void sema_destroy(ksema_t *sp); void sema_p(ksema_t *sp); void sema_v(ksema_t *sp); int sema_p_sig(ksema_t *sp); int sema_try_p(ksema_t *sp);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>sp</i> A pointer to a semaphore, type ksema_t. <i>val</i> Initial value for semaphore. <i>name</i> Descriptive string. This is obsolete and should be NULL. (Non-NULL strings are legal, but they are a waste of kernel memory.) <i>type</i> Variant type of the semaphore. Currently, only SEMA_DRIVER is supported. <i>arg</i> Type-specific argument; should be NULL.
DESCRIPTION	These functions implement counting semaphores as described by Dijkstra. A semaphore has a value which is atomically decremented by <code>sema_p()</code> and atomically incremented by <code>sema_v()</code>. The value must always be greater than or equal to zero. If <code>sema_p()</code> is called and the value is zero, the calling thread is blocked until another thread performs a <code>sema_v()</code> operation on the semaphore. Semaphores are initialized by calling <code>sema_init()</code>. 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 <code>sema_init()</code>. For this reason, <code>sema_destroy()</code> should be called before deallocating the storage containing the semaphore. <code>sema_p_sig()</code> decrements the semaphore, as does <code>sema_p()</code>. However, if the semaphore value is zero, <code>sema_p_sig()</code> will return without decrementing the value if a signal (that is, from <code>kill(2)</code>) is pending for the thread. <code>sema_try_p()</code> will decrement the semaphore value only if it is greater than zero, and will not block.
RETURN VALUES	0 <code>sema_try_p()</code> could not decrement the semaphore value because it was zero.

semaphore(9F)

	1	<code>sema_p_sig()</code> 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 <code>sema_init()</code> and <code>sema_destroy()</code>, which can be called from user context only. None of these functions can be called from a high-level interrupt context. In most cases, <code>sema_v()</code> and <code>sema_p()</code> should not be called from any interrupt context. If <code>sema_p()</code> 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 <code>sema_v()</code> becomes blocked on anything that requires the lower-priority interrupt, the system will hang. For example, the thread that will perform the <code>sema_v()</code> 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		<code>kill(2)</code>, <code>condvar(9F)</code>, <code>mutex(9F)</code> <i>Writing Device Drivers</i>

NAME	sprintf, snprintf – format characters in memory
SYNOPSIS	<pre>#include <sys/ddi.h> char *sprintf(char *buf, const char *fmt, ...); size_t snprintf(char *buf, size_t n, const char *fmt, ...);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>buf</i> Pointer to a character string. <i>fmt</i> Pointer to a character string.
DESCRIPTION	sprintf() 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>. The snprintf() function is identical to sprintf() with the addition of the argument <i>n</i>, which specifies the size of the buffer referred to by <i>buf</i>. The buffer is always terminated with the null byte.
Conversion Specifications	Each conversion specification is introduced by the % character, after which the following appear in sequence: 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. 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,u 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 conversion. The letters ABCDEF are used for X conversion. c The character value of 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

sprintf(9F)

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")
```

p

The argument is taken to be a pointer; the value of the pointer is displayed in unsigned hexadecimal. The display format is equivalent to %lx. To avoid lint warnings, cast pointers to type `void *` when using the %p format specifier.

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*.

`snprintf()` returns the number of characters formatted, that is, the number of characters that would have been written to the buffer if it were large enough. If the value of *n* is less than or equal to 0 on a call to `snprintf()`, the function simply returns the number of characters formatted.

CONTEXT

`sprintf()` and `snprintf()` can be called from user or interrupt context.

SEE ALSO

Writing Device Drivers

NAME	stoi, numtos – convert between an integer and a decimal string
SYNOPSIS	<pre>#include <sys/ddi.h> int stoi(char **<i>str</i>); void numtos(unsigned long <i>num</i>, char *<i>s</i>);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>str</i> Pointer to a character string to be converted. <i>num</i> Decimal number to be converted to a character string. <i>s</i> Character buffer to hold converted decimal number.
DESCRIPTION	
stoi()	<i>stoi()</i> returns the integer value of a string of decimal numeric characters beginning at <i>**str</i> . No overflow checking is done. <i>*str</i> is updated to point at the last character examined.
numtos()	<i>numtos()</i> converts a long into a null-terminated character string. No bounds checking is done. The caller must ensure there is enough space to hold the result.
RETURN VALUES	<i>stoi()</i> returns the integer value of the string <i>str</i> .
CONTEXT	<i>stoi()</i> can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>
NOTES	<i>stoi()</i> handles only positive integers; it does not handle leading minus signs.

strchr(9F)

NAME	strchr, strrchr – find a character in a string
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> char *strchr(const char *str, int chr); char *strrchr(const char *str, int chr);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>str</i> Pointer to a string to be searched. <i>chr</i> The character to search for.
DESCRIPTION	strchr() The strchr() function returns a pointer to the first occurrence of <i>chr</i> in the string pointed to by <i>str</i>. strrchr() The strrchr() function returns a pointer to the last occurrence of <i>chr</i> in the string pointed to by <i>str</i>.
RETURN VALUES	strchr() and strrchr() return a pointer to a character, or NULL, if the search fails.
CONTEXT	These functions can be called from user or interrupt context.
SEE ALSO	strcmp(9F) <i>Writing Device Drivers</i>

NAME	strcmp, strcasecmp, strncasecmp, strncmp – compare two null-terminated strings.
SYNOPSIS	<pre>#include <sys/ddi.h> int strcmp(const char *s1, const char *s2); int strcasecmp(const char *s1, const char *s2); int strncasecmp(const char *s1, const char *s2, size_t n); int strncmp(const char *s1, const char *s2, size_t n);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>s1</i>, <i>s2</i> Pointers to character strings. <i>n</i> Count of characters to be compared.
DESCRIPTION	
strcmp()	strcmp() returns 0 if the strings are the same, or the integer value of the expression (*s1 - *s2) for the last characters compared if they differ.
strcasecmp(), strncasecmp()	The strcasecmp() and strncasecmp() functions are case-insensitive versions of strcmp() and strncmp(), respectively, described in this section. They assume the ASCII character set and ignore differences in case when comparing lowercase and uppercase characters.
strncmp()	strncmp() returns 0 if the first <i>n</i> characters of <i>s1</i> and <i>s2</i> are the same, or (*s1 - *s2) for the last characters compared if they differ.
RETURN VALUES	strcmp() returns 0 if the strings are the same, or (*s1 - *s2) for the last characters compared if they differ. strcasecmp() and strncasecmp() return values in the same fashion as strcmp() and strncmp(), respectively. strncmp() returns 0 if the first <i>n</i> characters of strings are the same, or (*s1 - *s2) for the last characters compared if they differ.
CONTEXT	These functions can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

strcpy(9F)

NAME	strcpy, strncpy – copy a string from one location to another.				
SYNOPSIS	<pre>#include <sys/ddi.h> char *strcpy(char *dst, const char *srs); char *strncpy(char *dst, const char *srs, size_t n);</pre>				
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).				
PARAMETERS	<table><tr><td><i>dst</i>, <i>srs</i></td><td>Pointers to character strings.</td></tr><tr><td><i>n</i></td><td>Count of characters to be copied.</td></tr></table>	<i>dst</i> , <i>srs</i>	Pointers to character strings.	<i>n</i>	Count of characters to be copied.
<i>dst</i> , <i>srs</i>	Pointers to character strings.				
<i>n</i>	Count of characters to be copied.				
DESCRIPTION					
strcpy()	strcpy() 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.				
strncpy()	strncpy() 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.				
RETURN VALUES	strcpy() and strncpy() return <i>dst</i> .				
CONTEXT	strcpy() can be called from user or interrupt context.				
SEE ALSO	<i>Writing Device Drivers</i>				

NAME	strlen – determine the number of non-null bytes in a string
SYNOPSIS	<pre>#include <sys/ddi.h> size_t strlen(const char *s);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>s</i> Pointer to a character string.
DESCRIPTION	strlen() returns the number of non-null bytes in the string argument <i>s</i> .
RETURN VALUES	strlen() returns the number of non-null bytes in <i>s</i> .
CONTEXT	strlen() can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>

strlog(9F)

NAME	strlog – submit messages to the log driver														
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/strlog.h> #include <sys/log.h> int strlog(short <i>mid</i>, short <i>sid</i>, char <i>level</i>, unsigned short <i>flags</i>, char *<i>fmt</i>, ...);</pre>														
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>mid</i> 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>). <i>sid</i> Identification number for a particular minor device. <i>level</i> Tracing level for selective screening of low priority messages. Larger values imply less important information. <i>flags</i> Valid flag values are: <table> <tr> <td><code>SL_ERROR</code></td><td>Message is for error logger.</td></tr> <tr> <td><code>SL_TRACE</code></td><td>Message is for trace.</td></tr> <tr> <td><code>SL_NOTIFY</code></td><td>Mail copy of message to system administrator.</td></tr> <tr> <td><code>SL_CONSOLE</code></td><td>Log message to console.</td></tr> <tr> <td><code>SL_FATAL</code></td><td>Error is fatal.</td></tr> <tr> <td><code>SL_WARN</code></td><td>Error is a warning.</td></tr> <tr> <td><code>SL_NOTE</code></td><td>Error is a notice.</td></tr> </table> <i>fmt</i> <code>printf(3C)</code> style format string. <code>%e</code>, <code>%g</code>, and <code>%G</code> formats are not allowed but <code>%s</code> is supported.	<code>SL_ERROR</code>	Message is for error logger.	<code>SL_TRACE</code>	Message is for trace.	<code>SL_NOTIFY</code>	Mail copy of message to system administrator.	<code>SL_CONSOLE</code>	Log message to console.	<code>SL_FATAL</code>	Error is fatal.	<code>SL_WARN</code>	Error is a warning.	<code>SL_NOTE</code>	Error is a notice.
<code>SL_ERROR</code>	Message is for error logger.														
<code>SL_TRACE</code>	Message is for trace.														
<code>SL_NOTIFY</code>	Mail copy of message to system administrator.														
<code>SL_CONSOLE</code>	Log message to console.														
<code>SL_FATAL</code>	Error is fatal.														
<code>SL_WARN</code>	Error is a warning.														
<code>SL_NOTE</code>	Error is a notice.														
DESCRIPTION	<code>strlog()</code> expands the <code>printf(3C)</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 <code>getmsg(2)</code> system call. The <i>flags</i> argument specifies the type of the message and where it is to be sent. <code>strace(1M)</code> receives messages from the log driver and sends them to the standard output. <code>strerr(1M)</code> receives error messages from the log driver and appends them to a file called <code>/var/adm/streams/error.mm-dd</code>, where <i>mm-dd</i> identifies the date of the error message.														
RETURN VALUES	<code>strlog()</code> returns 0 if it fails to submit the message to the <code>log(7D)</code> driver and 1 otherwise.														

strlog(9F)

CONTEXT	strlog() can be called from user or interrupt context.
FILES	<i>/var/adm/streams/error.mm-dd</i> Error messages dated <i>mm-dd</i> appended by strerr(1M) from the log driver
SEE ALSO	strace(1M), strerr(1M), getmsg(2), log(7D), module_info(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

strqget(9F)

NAME	strqget – get information about a queue or band of the queue																
SYNOPSIS	<pre>#include <sys/stream.h> int strqget(queue_t *q, qfields_t what, unsigned char pri, void *valp);</pre>																
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>q</i> Pointer to the queue. <i>what</i> Field of the queue structure for (or the specified priority band) to return information about. Valid values are one of: <table> <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> <i>pri</i> Priority band of interest. <i>valp</i> The address of where to store the value of the requested field.	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.																
DESCRIPTION	strqget () 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.																
RETURN VALUES	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.																
CONTEXT	strqget () can be called from user or interrupt context.																
SEE ALSO	strqset(9F), queue(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>																

NAME	strqset – change information about a queue or band of the queue								
SYNOPSIS	<pre>#include <sys/stream.h> int strqset(queue_t *q, qfields_t what, unsigned char pri, intptr_t val) ;</pre>								
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>q</i> Pointer to the queue. <i>what</i> Field of the queue structure (or the specified priority band) to return information about. Valid values are one of: <table> <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> <i>pri</i> Priority band of interest. <i>val</i> The value for the field to be changed.	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.								
DESCRIPTION	strqset () 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.								
RETURN VALUES	On success, 0 is returned. EINVAL is returned if an undefined attribute is specified.								
CONTEXT	strqset () can be called from user or interrupt context.								
SEE ALSO	strqget(9F), queue(9S) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>								
NOTES	When lowering existing values, set QMINPSZ before setting QMAXPSZ; when raising existing values, set QMAXPSZ before setting QMINPSZ.								

STRUCT_DECL(9F)

NAME	STRUCT_DECL, SIZEOF_PTR, SIZEOF_STRUCT, STRUCT_BUF, STRUCT_FADDR, STRUCT_FGET, STRUCT_FGETP, STRUCT_FSET, STRUCT_FSETP, STRUCT_HANDLE, STRUCT_INIT, STRUCT_SIZE, STRUCT_SET_HANDLE – 32-bit application data access macros
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> STRUCT_DECL(<i>structname</i>, <i>handle</i>) ; STRUCT_HANDLE(<i>structname</i>, <i>handle</i>) ; void STRUCT_INIT(<i>handle</i>, model_t <i>umodel</i>) ; void STRUCT_SET_HANDLE(<i>handle</i>, model_t <i>umodel</i>, void *<i>addr</i>) ; STRUCT_FGET(<i>handle</i>, <i>field</i>) ; STRUCT_FGETP(<i>handle</i>, <i>field</i>) ; STRUCT_FSET(<i>handle</i>, <i>field</i>, <i>val</i>) ; STRUCT_FSETP(<i>handle</i>, <i>field</i>, <i>val</i>) ; <typeof <i>field</i>> *STRUCT_FADDR(<i>handle</i>, <i>field</i>) ; struct <i>structname</i> *STRUCT_BUF(<i>handle</i>) ; size_t SIZEOF_STRUCT(<i>structname</i>, <i>umodel</i>) ; size_t SIZEOF_PTR(<i>umodel</i>) ; size_t STRUCT_SIZE(<i>handle</i>) ;</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	The macros take the following parameters: <i>structname</i> The structure name (as would appear <i>after</i> the C keyword “struct”) of the native form. <i>umodel</i> A bit field containing either ILP32 model bit (DATAMODEL_ILP32), or the LP64 model get (DATAMODEL_LP64). 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. <i>handle</i> The variable name used to refer to a particular instance of a structure which is handled by these macros. <i>field</i> 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 "->".
DESCRIPTION	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. 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.
Declaration and Initialization Macros	The macros STRUCT_DECL() and STRUCT_HANDLE() declare structure handles on the stack, whereas the macros STRUCT_INIT() and STRUCT_SET_HANDLE() initialize the structure handles to point to an instance of the native form structure. The macros STRUCT_HANDLE() and STRUCT_SET_HANDLE() are used to declare and initialize a structure handle to an existing data structure, for example, ioctl within a STREAMS module. The macros STRUCT_DECL() and STRUCT_INIT(), on the other hand, are used in modules which declare and initialize a structure handle to a data structure allocated by STRUCT_DECL(), that is, any standard character or block device driver ioctl(9E) routine that needs to copy in data from a user-mode program. STRUCT_DECL(structname, handle) Declares a "structure handle" for a "struct" and allocates 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. void STRUCT_INIT(handle, model_t umodel) Initializes <i>handle</i> to point to the instance allocated by STRUCT_DECL(), 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 used in an ioctl(9E) routine <i>umodel</i> is the flag parameter; in adevmap(9E) routine <i>umodel</i> 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(structname, handle) Declares a "structure handle" <i>handle</i> 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 <i>addr</i>, it also sets the data model for <i>handle</i> to <i>umodel</i>. This is intended for handles created with STRUCT_HANDLE(). Fields cannot be referenced via the <i>handle</i> until this macro has been invoked. Typically, <i>addr</i> is the address of the native form structure containing the user-mode programs data. When used in an ioctl(9E) <i>umodel</i> is the flag parameter, in a devmap(9E) routine is the model parameter and in a mmap(9E) routine, <i>umodel</i> is the return value of ddi_mmap_get_model(9F).

STRUCT_DECL(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)`

Assigns *val* to the (non pointer) in the structure described by *handle*. 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 structname *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(structname, umodel)`

Returns size of *structname* 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_DECL(9F)

EXAMPLE 4 Using STRUCT_HANDLE() and STRUCT_SET_HANDLE() (Continued)

```
        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

NAME	swab – swap bytes in 16-bit halfwords
SYNOPSIS	<pre>#include <sys/sunddi.h> void swab(void *src, void *dst, size_t nbytes);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>src</i> A pointer to the buffer containing the bytes to be swapped. <i>dst</i> 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. <i>nbytes</i> Number of bytes to be swapped, rounded down to the nearest half-word.
DESCRIPTION	swab() 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.
CONTEXT	swab() can be called from user or interrupt context.
SEE ALSO	<i>Writing Device Drivers</i>
NOTES	Since swab() operates byte-by-byte, it can be used on non-aligned buffers.

testb(9F)

NAME	testb – check for an available buffer
SYNOPSIS	<pre>#include <sys/stream.h> int testb(size_t size, uint_t pri);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>size</i> Size of the requested buffer. <i>pri</i> Priority of the allocb request.
DESCRIPTION	testb() checks to see if an allocb(9F) call is likely to succeed if a buffer of <i>size</i> bytes at priority <i>pri</i> is requested. Even if testb() returns successfully, the call to allocb(9F) can fail. The <i>pri</i> argument is no longer used, but is retained for compatibility.
RETURN VALUES	Returns 1 if a buffer of the requested size is available, and 0 if one is not.
CONTEXT	testb() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 testb() example In a service routine, if copymsg(9F) fails (line 6), the message is put back on the queue (line 7) and a routine, tryagain, is scheduled to be run in one tenth of a second. Then the service routine returns. When the timeout(9F) 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 tryagain to run again in another tenth of a second. Note that tryagain 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 testb() will still succeed twice. If no message blocks are freed of these sizes before the service procedure runs again, then the copymsg(9F) will still fail. The reason testb() is used here is because it is significantly faster than calling copymsg. We must minimize the amount of time spent in a timeout() routine. <pre> 1 xxxsrv(q) 2 queue_t *q; 3 { 4 mblk_t *mp; 5 mblk_t *nmp; 6 . . . 7 if ((nmp = copymsg(mp)) == NULL) { 8 putbq(q, mp); 9 timeout(tryagain, (intptr_t)q, drv_usectohz(100000)); 10 return; 11 } 12 . . . 13 }</pre>

EXAMPLE 1 testb() example (Continued)

```

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.

timeout(9F)

NAME	timeout – execute a function after a specified length of time
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/conf.h> timeout_id_t timeout(void (* func)(void *), void *arg, clock_t ticks);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>func</i> Kernel function to invoke when the time increment expires. <i>arg</i> Argument to the function. <i>ticks</i> Number of clock ticks to wait before the function is called. Use <code>drv_usectohz(9F)</code> to convert microseconds to clock ticks.
DESCRIPTION	The <code>timeout()</code> 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. The function called by <code>timeout()</code> must adhere to the same restrictions as a driver soft interrupt handler. The function called by <code>timeout()</code> is run in interrupt context and must not sleep or call other functions that might sleep. The <code>delay(9F)</code> function calls <code>timeout()</code>. Because <code>timeout()</code> is subject to priority inversion, drivers waiting on behalf of processes with real-time constraints should use <code>cv_timedwait(9F)</code> rather than <code>delay()</code>.
RETURN VALUES	<code>timeout()</code> returns an opaque non-zero timeout identifier that can be passed to <code>untimeout(9F)</code> to cancel the request.
CONTEXT	<code>timeout()</code> can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using <code>timeout()</code> 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. <pre>static void xxtimeout_handler(void *arg) { struct xxstate *xsp = (struct xxstate *)arg; mutex_enter(&xsp->lock); cv_signal(&xsp->cv); xsp->flags = TIMED_OUT; mutex_exit(&xsp->lock); xsp->timeout_id = 0; } static uint_t xxintr(caddr_t arg) {</pre>

EXAMPLE 1 Using timeout() (Continued)

```

        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 bufcall(9F), cv_timedwait(9F), ddi_in_panic(9F), delay(9F),
drv_usectohz(9F), untimeout(9F)

Writing Device Drivers

uiomove(9F)

NAME	uiomove – copy kernel data using uio structure
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/uio.h> int uiomove(caddr_t address, size_t nbytes, enum uio_rw rwflag, uio_t *uio_p);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). <i>address</i> Source/destination kernel address of the copy. <i>nbytes</i> Number of bytes to copy. <i>rwflag</i> Flag indicating read or write operation. Possible values are UIO_READ and UIO_WRITE. <i>uio_p</i> Pointer to the uio structure for the copy.
DESCRIPTION	The uiomove() function copies <i>nbytes</i> of data to or from the space defined by the uio structure (described in uio(9S)) and the driver. The uio_segflg member of the uio(9S) structure determines the type of space to or from which the transfer is being made. If it is set to UIO_SYSSPACE, the data transfer is between addresses in the kernel. If it is set to UIO_USERSPACE, the transfer is between a user program and kernel space. <i>rwflag</i> indicates the direction of the transfer. If UIO_READ is set, the data will be transferred from <i>address</i> to the buffer(s) described by <i>uio_p</i>. If UIO_WRITE is set, the data will be transferred from the buffer(s) described by <i>uio_p</i> to <i>address</i>. In addition to moving the data, uiomove() adds the number of bytes moved to the iov_base member of the iovec(9S) structure, decreases the iov_len member, increases the uio_offset member of the uio(9S) structure, and decreases the uio_resid member. This function automatically handles page faults. <i>nbytes</i> does not have to be word-aligned.
RETURN VALUES	uiomove() returns 0 upon success or EFAULT on failure.
CONTEXT	User context only, if uio_segflg is set to UIO_USERSPACE. User or interrupt context, if uio_segflg is set to UIO_SYSSPACE.
SEE ALSO	ureadc(9F), uwritec(9F), iovec(9S), uio(9S) <i>Writing Device Drivers</i>
WARNINGS	If uio_segflg is set to UIO_SYSSPACE and <i>address</i> is selected from user space, the system may panic.

NAME	unbufcall – cancel a pending bufcall request
SYNOPSIS	<pre>#include <sys/stream.h> void unbufcall(bufcall_id_t id);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>id</i> Identifier returned from bufcall(9F) or esbbcall(9F).
DESCRIPTION	unbufcall cancels a pending bufcall() or esbbcall() request. The argument <i>id</i> is a non-zero identifier for the request to be cancelled. <i>id</i> is returned from the bufcall() or esbbcall() function used to issue the request. unbufcall() 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 unbufcall() or deadlock may result.
RETURN VALUES	None.
CONTEXT	unbufcall() can be called from user or interrupt context.
SEE ALSO	bufcall(9F), esbbcall(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

unlinkb(9F)

NAME	unlinkb – remove a message block from the head of a message
SYNOPSIS	<pre>#include <sys/stream.h> mbblk_t *unlinkb (mbblk_t *mp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>mp</i> Pointer to the message.
DESCRIPTION	unlinkb() removes the first message block from the message pointed to by <i>mp</i> . A new message, minus the removed message block, is returned.
RETURN VALUES	If successful, unlinkb() 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.
CONTEXT	unlinkb() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 unlinkb() example 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. <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>
SEE ALSO	linkb(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

NAME	untimeout – cancel previous timeout function call
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/conf.h> clock_t untimeout(timeout_id_t id);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>id</i> Opaque timeout ID from a previous timeout(9F) call.
DESCRIPTION	untimeout() cancels a pending timeout(9F) request. untimeout() 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 untimeout() or a deadlock may result. Since no mutex should be held across the call to untimeout(), 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 untimeout() 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.
RETURN VALUES	untimeout() returns -1 if the <i>id</i> is not found. Otherwise, it returns an integer value greater than or equal to 0.
CONTEXT	untimeout() can be called from user or interrupt context.
EXAMPLES	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. <pre>static void xxtimeout_handler(void *arg) { struct xxstate *xsp = (struct xxstate *)arg; mutex_enter(&xsp->lock); cv_signal(&xsp->cv); xsp->flags = TIMED_OUT; mutex_exit(&xsp->lock); xsp->timeout_id = 0; } static uint_t xxintr(caddr_t arg) { struct xxstate *xsp = (struct xxstate *)arg; . . . mutex_enter(&xsp->lock); /* Service interrupt */ cv_signal(&xsp->cv); mutex_exit(&xsp->lock);</pre>

untimeout(9F)

```
        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(xxtimeout_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

NAME	ureadc – add character to a uio structure
SYNOPSIS	<pre>#include <sys/uio.h> #include <sys/types.h> int ureadc(int <i>c</i>, uio_t *<i>uio_p</i>);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>c</i> The character added to the uio(9S) structure. <i>uio_p</i> Pointer to the uio(9S) structure.
DESCRIPTION	ureadc() transfers the character <i>c</i> into the address space of the uio(9S) structure pointed to by <i>uio_p</i> , and updates the uio structure as for uiomove(9F).
RETURN VALUES	0 is returned on success and EFAULT on failure.
CONTEXT	ureadc() can be called from user or interrupt context.
SEE ALSO	uiomove(9F), uwritec(9F), iovec(9S), uio(9S) <i>Writing Device Drivers</i>

uwritec(9F)

NAME	uwritec – remove a character from a uio structure
SYNOPSIS	<pre>#include <sys/uio.h> int uwritec(uio_t *uio_p);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
PARAMETERS	<i>uio_p</i> Pointer to the uio(9S) structure
DESCRIPTION	<code>uwritec()</code> returns a character from the uio structure pointed to by <i>uio_p</i> and updates the uio structure. See <code>uiomove(9F)</code> .
RETURN VALUES	The next character for processing is returned on success, and -1 is returned if uio is empty or if there is an error.
CONTEXT	<code>uwritec()</code> can be called from user or interrupt context.
SEE ALSO	<code>uiomove(9F)</code> , <code>ureadc(9F)</code> , <code>iovec(9S)</code> , <code>uio(9S)</code> <i>Writing Device Drivers</i>

NAME	va_arg, va_start, va_copy, va_end – handle variable argument list	
SYNOPSIS	<pre>#include <sys/varargs.h> void va_start(va_list <i>pvar</i>, void <i>parmN</i>) ; (<i>type</i> *) va_arg(va_list <i>pvar</i>, <i>type</i>) ; void va_copy(va_list <i>dest</i>, va_list <i>src</i>) ; void va_end(va_list <i>pvar</i>) ;</pre>	
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).	
PARAMETERS		
va_start()	<i>pvar</i>	Pointer to variable argument list.
	<i>name</i>	Identifier of rightmost parameter in the function definition.
va_arg()	<i>pvar</i>	Pointer to variable argument list.
	<i>type</i>	Type name of the next argument to be returned.
va_copy()	<i>dest</i>	Destination variable argument list.
	<i>src</i>	Source variable argument list.
va_end()	<i>pvar</i>	Pointer to variable argument list.
DESCRIPTION	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 <code>varargs()</code> 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. <code>va_list</code> is the type defined for the variable used to traverse the list of arguments. <code>va_start()</code> is called to initialize <i>pvar</i> to the beginning of the variable argument list. <code>va_start()</code> must be invoked before any access to the unnamed arguments. The parameter <i>name</i> is the identifier of the rightmost parameter in the variable parameter list in the function definition (the one just before the “<code>, . . .</code>”). If this parameter is declared with the <code>register</code> 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. <code>va_arg()</code> expands to an expression that has the type and value of the next argument in the call. The parameter <i>pvar</i> must be initialized by <code>va_start()</code>. Each invocation of <code>va_arg()</code> modifies <i>pvar</i> so that the values of successive arguments are returned in turn. The parameter <i>type</i> is the type name of the next argument to be returned. The type name must be specified in such a way that the type of pointer to an object that	

va_arg(9F)

has the specified type can be obtained 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()`. It then 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 if any of these restrictions are not 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 might be useful for a device that 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    ap;
    int        i;

    /*
     * Write variable-length command to destination
     */
    va_start(ap, ncmdbytes);
    for (i = 0; i < ncmdbytes; i++) {
        *cmdp++ = va_arg(ap, uchar_t);
    }
    va_end(ap);
}
```

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.

Specifying a second argument of `char` or `short` to `va_arg` makes your code non-portable, 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.

NAME	<code>vsnprintf</code> – format characters in memory
SYNOPSIS	<pre>#include <sys/varargs.h> #include <sys/ddi.h> #include <sys/sunddi.h> char *vsnprintf(char *buf, const char *fmt, va_list ap);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<i>buf</i> Pointer to a character string. <i>fmt</i> Pointer to a character string. <i>ap</i> Pointer to a variable argument list.
DESCRIPTION	<code>vsnprintf()</code> 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>. <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 <code>va_start(9F)</code>. <code>va_end(9F)</code> is used to clean up and must be called after each traversal of the list. Multiple traversals of the argument list, each bracketed by <code>va_start(9F)</code> and <code>va_end(9F)</code>, are possible. Each conversion specification is introduced by the <code>%</code> 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 <code>l</code> (<code>ll</code>) specifying that a following <code>d</code>, <code>D</code>, <code>o</code>, <code>O</code>, <code>x</code>, <code>X</code>, or <code>u</code> conversion character applies to a long (long long) integer argument. An <code>l</code> (<code>ll</code>) before any other conversion character is ignored. A character indicating the type of conversion to be applied: <code>d,D,o,O,x,X,u</code> The integer argument is converted to signed decimal (<code>d</code>, <code>D</code>), unsigned octal (<code>o</code>, <code>O</code>), unsigned hexadecimal (<code>x</code>, <code>X</code>) or unsigned decimal (<code>u</code>), respectively, and copied. The letters <code>abcdef</code> are used for <code>x</code> conversion. The letters <code>ABCDEF</code> are used for <code>X</code> conversion. <code>c</code> The character value of the argument is copied. <code>b</code> 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

vsprintf(9F)

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);
```

`vsprintf(9F)`

EXAMPLE 1 Using `vsprintf()` *(Continued)*

```
/* 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

WR(9F)

NAME	WR, wr – get pointer to the write queue for this module or driver
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/ddi.h> queue_t *WR(queue_t *q);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
PARAMETERS	<i>q</i> Pointer to the <i>read</i> queue whose <i>write</i> queue is to be returned.
DESCRIPTION	The WR() 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. CAUTION: Make sure the argument to this function is a pointer to a <i>read</i> queue. WR() will not check for queue type, and a system panic could result if the pointer is not to a <i>read</i> queue.
RETURN VALUES	The pointer to the <i>write</i> queue.
CONTEXT	WR() can be called from user or interrupt context.
EXAMPLES	EXAMPLE 1 Using WR() In a STREAMS close(9E) 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. <pre>1 xxxclose(q, flag) 2 queue_t *q; 3 int flag; 4 { 5 q->q_ptr = NULL; 6 WR(q)->q_ptr = NULL; 7 . . . 7 }</pre>
SEE ALSO	close(9E), OTHERQ(9F), RD(9F) <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>

Index

Numbers and Symbols

- 32-bit driver ID management routines —
 - id32_alloc, 455
- 32-bit driver ID management routines —
 - id32_free, 455
- 32-bit driver ID management routines —
 - id32_lookup, 455

A

- activate a new DMA window —
 - ddi_dma_getwin, 272
- add a fully initialized kstat to the system —
 - kstat_install, 471
- add a soft interrupt, — ddi_add_softintr, 213
- add an interrupt handler, — ddi_add_intr, 209
- address, return mapped virtual address —
 - csx_GetMappedAddr, 110
- adjmsg — trim bytes from a message, 40
- Device power cycle advisory check —
 - pm_trans_check, 529
- allocate and free a scsi_pkt structure —
 - scsi_hba_pkt_alloc, 595
- allocate and free transport structures —
 - scsi_hba_tran_alloc, 598
- allocate and free a scsi_pkt structure —
 - scsi_hba_pkt_alloc, scsi_hba_pkt_free, 595
- allocate and free transport structures —
 - scsi_hba_tran_alloc, scsi_hba_tran_free, 598
- allocate a message block — allocb, 41
- allocate and free non-sequentially accessed memory
 - allocate and free non-sequentially accessed memory (Continued)
 - ddi_iopb_alloc, 317
 - ddi_iopb_free, 317
 - allocate DMA handle —
 - ddi_dma_alloc_handle, 253
 - allocate kernel memory
 - ddi_umem_alloc, 400
 - ddi_umem_free, 400
 - ddi_umem_zalloc, 400
 - kmem_alloc, 461
 - kmem_free, 461
 - kmem_zalloc, 461
 - allocate memory for DMA transfer —
 - ddi_dma_mem_alloc, 276
 - allocate space — rmalloc, 562
 - allocate space from a resource map —
 - rmalloc_wait, 566
 - allocb — allocate a message block, 41
 - allow 64 bit transfers on SBus —
 - ddi_dma_set_sbus64, 290
 - anocancel — prevent cancellation of asynchronous I/O request, 44
 - aphysio — perform asynchronous physical I/O, 45
 - assert — expression verification, 47
 - associate STREAMS queue with bottom driver in queue — qassociate, 547
 - asynchronous physical I/O — aphysio, 45
 - asynchronous STREAMS perimeter upgrade —
 - qwriter, 560

B

bcopy — copy data between address locations in kernel, 51
binds a system buffer to a DMA handle — ddi_dma_buf_bind_handle, 256
binds an address to a DMA handle — ddi_dma_addr_bind_handle, 248
bioclose — clone another buffer, 53
bioerror — indicate error in buffer header, 58
biofini — uninitialize a buffer structure, 59
bioinit — initialize a buffer structure, 60
biomodified — check if a buffer is modified, 61
bioreset — reuse a private buffer header after I/O is complete, 62
biosize — returns size of a buffer structure, 63
bufcall — call a function when a buffer becomes available, 69, 649
bufcall, call a function when a buffer becomes available, 69
buffer header
 indicate error — bioerror, 58
 reuse a private buffer header after I/O is complete — bioreset, 62
busy-wait for specified interval — drv_usecwait, 427
byte streams, compare two — bcmp, 50
bytes, size
 convert size in pages — ptob, 534
 convert to size in memory pages (round down) — btob, 67
 convert to size in memory pages (round up) — btob, 68

C

call a function when a buffer becomes available — qbufcall, 549
call a function when a buffer becomes available, bufcall, 69
call a STREAMS put procedure — put, 537
cancel a pending qbufcall request — qunbufcall, 556
cancel previous timeout function call — quntimeout, 557
cancellation of asynchronous I/O — anocancel, 44

character strings

 compare two null terminated strings — strcmp, strncmp, 631
 convert between an integer and a decimal string — stoi, numtos, 629
 copy a string from one location to another — strcpy, strncpy, 632
 determine the number of non-null bytes in a string — strlen, 633
 find a character in a string — strchr, 630
 format in memory — sprintf, 627
check data access and DMA handles, 221
check device state, 303
check for an available buffer — testb, 644
check for the existence of a property — ddi_prop_exists, 360
check if a buffer is modified — biomodified, 61
CIS tuple
 first tuple — csx_GetFirstTuple, 107
 next tuple — csx_GetNextTuple, 107
clear client event mask — csx_ReleaseSocketMask, 196
client, register client — csx_RegisterClient, 175
client event mask
 return client event mask — csx_GetEventMask, 204
 set client event mask — csx_SetEventMask, 204
client return
 — csx_GetFirstClient, 105
 — csx_GetNextClient, 105
clone another buffer — bioclose, 53
concatenate two message blocks — linkb, 475
condition variable routines, driver
 — condvar, 81
 — cv_broadcast, 81
 — cv_init, 81
 — cv_signal, 81
 — cv_timedwait, 81
 — cv_timedwait_sig, 81
 — cv_wait, 81
 — cv_wait_sig, 81
configure PC Card and socket — csx_RequestConfiguration, 184
control driver notification of user accesses — ddi_mapdev_intercept, 331
 ddi_mapdev_nointercept, 331

control device components' availability for Power Management
 — pm_busy_component, 519
 — pm_idle_component, 519
 control the validation of memory address translations
 — devmap_load, 419
 — devmap_unload, 419
 convert a DMA segment to a DMA address
 cookie — ddi_dma_segtocookie, 288
 convert clock ticks to microseconds —
 drv_hztousec, 424
 convert device sizes — csx_ConvertSize, 94
 convert device speeds — csx_ConvertSpeed, 95
 convert error return codes to text strings —
 csx_Error2Text, 101
 convert events to text strings —
 csx_Event2Text, 102
 convert microseconds to clock ticks —
 drv_usecsthz, 426
 copy data from one device register to another
 device register — ddi_device_copy, 231
 create minor nodes for client —
 csx_MakeDeviceNode, 117
 create a minor node for this device —
 ddi_create_minor_node, 229
 create and initialize a new kstat —
 kstat_create, 469
 create driver-controlled mapping of device —
 ddi_mapdev, 328
 csx_AccessConfigurationRegister — read or
 write a PC Card Configuration Register, 92
 csx_ConvertSize — convert device sizes, 94
 csx_ConvertSpeed — convert device speeds, 95
 csx_CS_DDI_Info — obtain DDI
 information, 96
 csx_DeregisterClient — remove client from
 Card Services list, 98
 csx_DupHandle — duplicate access handle, 99
 csx_Error2Text — convert error return codes to
 text strings, 101
 csx_Event2Text — convert events to text
 strings, 102
 csx_FreeHandle — free access handle, 103
 csx_Get16 — read from device register, 104
 csx_Get32 — read from device register, 104
 csx_Get64 — read from device register, 104
 csx_Get8 — read from device register, 104
 csx_GetEventMask — return client event
 mask, 204
 csx_GetFirstClient — return first client, 105
 csx_GetFirstTuple — return first CIS tuple, 107
 csx_GetHandleOffset — return current access
 handle offset, 109
 csx_GetMappedAddr — return mapped virtual
 address, 110
 csx_GetNextClient — return next client, 105
 csx_GetNextTuple — return next CIS tuple, 107
 csx_GetStatus — return status of PC Card and
 socket, 111
 csx_GetTupleData — return data portion of
 tuple, 115
 csx_MakeDeviceNode — create minor nodes for
 client, 117
 csx_MapLogSocket — return physical socket
 number, 119
 csx_MapMemPage — map memory area on PC
 Card, 120
 csx_ModifyConfiguration — modify PC Card
 configuration, 121
 csx_ModifyWindow — modify window
 attributes, 123
 csx_Parse_CISTPL_BATTERY — parse Battery
 Replacement Date tuple, 125
 csx_Parse_CISTPL_BYTEORDER — parse Byte
 Order tuple, 126
 csx_Parse_CISTPL_CFTABLE_ENTRY — parse
 Card Configuration Table tuple, 128
 csx_Parse_CISTPL_CONFIG — parse
 Configuration tuple, 134
 csx_Parse_CISTPL_DATE — parse Card
 Initialization Date tuple, 137
 csx_Parse_CISTPL_DEVICE — parse Device
 Information tuple for Common Memory, 138
 csx_Parse_CISTPL_DEVICE_A — parse Device
 Information tuple for Attribute Memory, 138
 csx_Parse_CISTPL_DEVICE_OA — parse Other
 Condition Device Information tuple for
 Attribute Memory, 138
 csx_Parse_CISTPL_DEVICE_OC — parse Other
 Condition Device Information tuple for
 Common Memory, 138
 csx_Parse_CISTPL_DEVICEGEO — parse
 Device Geo tuple, 141
 csx_Parse_CISTPL_DEVICEGEO_A — parse
 Device Geo A tuple, 143

csx_Parse_CISTPL_FORMAT — parse Data Recording Format tuple, 145
 csx_Parse_CISTPL_FUNCE — parse Function Extension tuple, 147
 csx_Parse_CISTPL_FUNCID — parse Function Identification tuple, 155
 csx_Parse_CISTPL_GEOMETRY — parse Geometry tuple, 157
 csx_Parse_CISTPL_JEDEC_A — parse JEDEC Identifier tuple for Attribute Memory, 158
 csx_Parse_CISTPL_JEDEC_C — parse JEDEC Identifier tuple for Common Memory, 158
 csx_Parse_CISTPL_LINKTARGET — parse Link Target tuple, 160
 csx_Parse_CISTPL_LONGLINK_A — parse Long Link A tuple, 161
 csx_Parse_CISTPL_LONGLINK_C — parse Long Link C tuple, 161
 csx_Parse_CISTPL_LONGLINK_MFC — parse Multi-Function tuple, 163
 csx_Parse_CISTPL_MANFID — parse Manufacturer Identification tuple, 165
 csx_Parse_CISTPL_ORG — parse Data Organization tuple, 166
 csx_Parse_CISTPL_SPCL — parse Special Purpose tuple, 167
 csx_Parse_CISTPL_SWIL — parse Software Interleaving tuple, 169
 csx_Parse_CISTPL_VERS_1 — parse Level-1 Version/Product Information tuple, 170
 csx_Parse_CISTPL_VERS_2 — parse Level-2 Version and Information tuple, 171
 csx_ParseTuple — generic tuple parser, 172
 csx_Put16 — write to device register, 174
 csx_Put32 — write to device register, 174
 csx_Put64 — write to device register, 174
 csx_Put8 — write to device register, 174
 csx_RegisterClient — register client, 175
 csx_ReleaseConfiguration — release configuration on PC Card, 178
 csx_ReleaseIO — release I/O resources, 188
 csx_ReleaseIRQ — release IRQ resource, 193
 csx_ReleaseSocketMask — clear client event mask, 196
 csx_ReleaseWindow — release window resources, 198
 csx_RepGet16 — read repetitively from device register, 180

csx_RepGet32 — read repetitively from device register, 180
 csx_RepGet64 — read repetitively from device register, 180
 csx_RepGet8 — read repetitively from device register, 180
 csx_RepPut16 — write repetitively to device register, 182
 csx_RepPut32 — write repetitively to device register, 182
 csx_RepPut64 — write repetitively to device register, 182
 csx_RepPut8 — write repetitively to device register, 182
 csx_RequestConfiguration — configure PC Card and socket, 184
 csx_RequestIO — request I/O resources, 188
 csx_RequestIRQ — request IRQ resource, 193
 csx_RequestSocketMask — request client event mask, 196
 csx_RequestWindow — request window resources, 198
 csx_ResetFunction — reset a function on a PC card, 203
 csx_SetEventMask — set client event mask, 204
 csx_SetHandleOffset — set current access handle offset, 206
 csx_ValidateCIS — validate Card Information Structure (CIS), 207
 current thread, get id of, 307

D

datamsg — test whether a message is a data message, 208
 DDI access credential structure, —
 ddi_get_cred, 302
 ddi_add_intr — add an interrupt handler, 209
 ddi_add_softintr — add a soft interrupt, 213
 DDI announce a device, — ddi_report_dev, 387
 ddi_binding_name — return driver binding name, 219
 ddi_check_acc_handle, 221, 303
 ddi_check_dma_handle, 221
 ddi_create_minor_node — create a minor node for this device, 229

- ddi_dev_is_needed — inform the system that a device's component is required, 238
- ddi_dev_report_fault, 244
- DDI device access, slave access only —
 - ddi_slaveonly, 394
- ddi_device_copy — copy data from one device register to another device register, 231
- DDI device critical region of control
 - enter — ddi_enter_critical, 298
 - exit — ddi_exit_critical, 298
- DDI device information structure
 - find parent — ddi_get_parent, 309
 - get the root of the dev_info tree —
 - ddi_root_node, 390
 - remove a minor node for this devinfo —
 - ddi_remove_minor_node, 382
- DDI device instance number, get —
 - ddi_get_instance, 306
- DDI device mapping
 - ddi_mapdev — create driver-controlled mapping of device, 328
 - ddi_mapdev_intercept — control driver notification of user accesses, 331
 - ddi_mapdev_nointercept — control driver notification of user accesses, 331
 - devmap_default_access — device mapping access entry point, 408
- DDI device registers
 - map — ddi_map_regs, 335
 - return the number of register sets —
 - ddi_dev_nregs, 242
 - return the size — ddi_dev_regsize, 243
 - unmap — ddi_unmap_regs, 335
- DDI device's private data area
 - get the address —
 - ddi_get_driver_private, 304
 - set the address —
 - ddi_set_driver_private, 304
- DDI device virtual address
 - read 16 bit — ddi_peek16, 352
 - read 32 bit — ddi_peek32, 352
 - read 64 bit — ddi_peek64, 352
 - read 8 bit — ddi_peek8, 352
 - read a value — ddi_peek, 352
 - write 16 bit — ddi_poke16, 354
 - write 32 bit — ddi_poke32, 354
 - write 64 bit — ddi_poke64, 354
 - write 8 bit — ddi_poke8, 354
- DDI device virtual address (Continued)
 - write a value — ddi_poke, 354
- ddi_device_zero — zero fill the device register, 233
- ddi_devid_compare — Kernel interfaces for device ids, 235
- ddi_devid_free — Kernel interfaces for device ids, 235
- ddi_devid_init — Kernel interfaces for device ids, 235
- ddi_devid_register — Kernel interfaces for device ids, 235
- ddi_devid_sizeof — Kernel interfaces for device ids, 235
- ddi_devid_unregister — Kernel interfaces for device ids, 235
- ddi_devid_valid — Kernel interfaces for device ids, 235
- DDI devinfo node name
 - return — ddi_binding_name, 219
 - return — ddi_get_name, 219
 - return — ddi_node_name, 350
- DDI direct memory access, convert DMA handle to DMA addressing cookie —
 - ddi_dma_htoc, 274
- DDI direct memory access services
 - allocate consistent memory—
 - ddi_iopb_alloc, 337
 - convert a DMA cookie — ddi_dma_coff, 262
 - easier DMA setup —
 - ddi_dma_addr_setup, 251
 - easier DMA setup —
 - ddi_dma_buf_setup, 259
 - find minimum alignment and transfer size for device — ddi_iomin, 316
 - find post DMA mapping alignment and minimum effect properties —
 - ddi_dma_devalign, 264
 - free consistent memory —
 - ddi_iopb_free, 337
 - report current DMA window offset and size —
 - ddi_dma_curwin, 263
 - setup DMA mapping —
 - ddi_dma_setup, 283, 285, 288
 - setup DMA resources —
 - ddi_dma_setup, 291
 - shift current DMA window —
 - ddi_dma_movwin, 279

DDI direct memory access services (Continued)

- tear down DMA mapping —
 - ddi_dma_free, 269
- ddi_dma_addr_bind_handle — binds an address to a DMA handle, 248
- ddi_dma_alloc_handle — allocate DMA handle, 253
- ddi_dma_buf_bind_handle — binds a system buffer to a DMA handle, 256
- ddi_dma_burstsizes — find out the allowed burst sizes for a DMA mapping, 261
- ddi_dma_free_handle — free DMA handle, 270
- ddi_dma_get_attr, 271
- ddi_dma_getwin — activate a new DMA window, 272
- ddi_dma_mem_alloc — allocate memory for DMA transfer, 276
- ddi_dma_mem_free — free previously allocated memory, 278
- ddi_dma_nextcookie — retrieve subsequent DMA cookie, 281
- ddi_dma_nextseg — get next DMA segment, 283
- ddi_dma_nextwin — get next DMA window, 285
- ddi_dma_numwin — retrieve number of DMA windows, 287
- ddi_dma_segtocookie — convert a DMA segment to a DMA address cookie, 288
- ddi_dma_set_sbus64 — allow 64 bit transfers on SBUS, 290
- ddi_dma_sync — synchronize CPU and I/O views of memory, 293
- ddi_dma_unbind_handle — unbinds the address in a DMA handle, 295
- ddi_dmae — system DMA engine functions, 266
- ddi_dmae_1stparty — system DMA engine functions, 266
- ddi_dmae_alloc — system DMA engine functions, 266
- ddi_dmae_disable — system DMA engine functions, 266
- ddi_dmae_enable — system DMA engine functions, 266
- ddi_dmae_getattr — system DMA engine functions, 266

- ddi_dmae_getcnt — system DMA engine functions, 266
- ddi_dmae_getlim — system DMA engine functions, 266
- ddi_dmae_prog — system DMA engine functions, 266
- ddi_dmae_release — system DMA engine functions, 266
- ddi_dmae_stop — system DMA engine functions, 266
- ddi_driver_major, 296
- ddi_driver_name — return normalized driver name, 297
- ddi_ffs — find first (last) bit set in a long integer, 299
- ddi_fls — find first (last) bit set in a long integer, 299
- ddi_get_iblock_cookie — get interrupt block cookie, 209
- ddi_get_kt_did, 307
- ddi_get_lbolt, returns the value of lbolt, 308
- ddi_get_name — return driver binding name, 219
- ddi_get_pid, returns the process ID, 310
- ddi_get_soft_iblock_cookie — get soft interrupt block cookie, 213
- ddi_get_time, returns the current time in seconds, 311
- ddi_get16 — read data from the device, 300
- ddi_get32 — read data from the device, 300
- ddi_get64 — read data from the device, 300
- ddi_get8 — read data from the device, 300
- ddi_getiminoir, display a SCSI request sense message, 305
- ddi_in_panic — determine if system is in panic state, 312
- DDI information — csx_CS_DDI_Info, 96
- DDI interrupt handling
 - add an interrupt — ddi_add_intr, 209
 - get interrupt block cookie —
 - ddi_get_iblock_cookie, 209
 - indicate interrupt handler type —
 - ddi_intr_hilevel, 313
 - remove an interrupt — ddi_remove_intr, 209
 - return the number of interrupt specifications —
 - ddi_dev_nintrs, 241
- ddi_io_get16 — read data from the mapped device register in I/O space, 314

`ddi_io_get32` — read data from the mapped device register in I/O space, 314
`ddi_io_get8` — read data from the mapped device register in I/O space, 314
`ddi_io_getb` — read data from the mapped device register in I/O space, 314
`ddi_io_getl` — read data from the mapped device register in I/O space, 314
`ddi_io_getw` — read data from the mapped device register in I/O space, 314
`ddi_io_put16` — write data to the mapped device register in I/O space, 319
`ddi_io_put32` — write data to the mapped device register in I/O space, 319
`ddi_io_put8` — write data to the mapped device register in I/O space, 319
`ddi_io_putb` — write data to the mapped device register in I/O space, 319
`ddi_io_putl` — write data to the mapped device register in I/O space, 319
`ddi_io_putw` — write data to the mapped device register in I/O space, 319
`ddi_io_rep_get16` — read multiple data from the mapped device register in I/O space, 321
`ddi_io_rep_get32` — read multiple data from the mapped device register in I/O space, 321
`ddi_io_rep_get8` — read multiple data from the mapped device register in I/O space, 321
`ddi_io_rep_getb` — read multiple data from the mapped device register in I/O space, 321
`ddi_io_rep_getl` — read multiple data from the mapped device register in I/O space, 321
`ddi_io_rep_getw` — read multiple data from the mapped device register in I/O space, 321
`ddi_io_rep_put16` — write multiple data to the mapped device register in I/O space, 323
`ddi_io_rep_put32` — write multiple data to the mapped device register in I/O space, 323
`ddi_io_rep_put8` — write multiple data to the mapped device register in I/O space, 323
`ddi_io_rep_putb` — write multiple data to the mapped device register in I/O space, 323
`ddi_io_rep_putl` — write multiple data to the mapped device register in I/O space, 323
`ddi_io_rep_putw` — write multiple data to the mapped device register in I/O space, 323
`ddi_iopb_alloc` — allocate and free non-sequentially accessed memory, 317
`ddi_iopb_free` — allocate and free non-sequentially accessed memory, 317
`ddi_log_sysevent`, 325
`ddi_mapdev` — create driver-controlled mapping of device, 328
`ddi_mapdev_intercept` — control driver notification of user accesses, 331
`ddi_mapdev_set_device_acc_attr` — Set the device attributes for the mapping, 333
`ddi_mem_get16` — read data from mapped device in the memory space or allocated DMA memory, 339
`ddi_mem_get32` — read data from mapped device in the memory space or allocated DMA memory, 339
`ddi_mem_get64` — read data from mapped device in the memory space or allocated DMA memory, 339
`ddi_mem_put16` — write data to mapped device in the memory space or allocated DMA memory, 340
`ddi_mem_put32` — write data to mapped device in the memory space or allocated DMA memory, 340
`ddi_mem_put64` — write data to mapped device in the memory space or allocated DMA memory, 340
`ddi_mem_rep_get16` — read data from mapped device in the memory space or allocated DMA memory, 342
`ddi_mem_rep_get32` — read data from mapped device in the memory space or allocated DMA memory, 342
`ddi_mem_rep_get64` — read data from mapped device in the memory space or allocated DMA memory, 342
`ddi_mem_rep_get8` — read data from mapped device in the memory space or allocated DMA memory, 342
`ddi_mem_rep_put16` — write data to mapped device in the memory space or allocated DMA memory, 344
`ddi_mem_rep_put32` — write data to mapped device in the memory space or allocated DMA memory, 344
`ddi_mem_rep_put64` — write data to mapped device in the memory space or allocated DMA memory, 344

`ddi_mem_rep_put8` — write data to mapped device in the memory space or allocated DMA memory, 344
 DDI memory mapping
 map a segment — `ddi_segmap`, 392
 map a segment — `devmap_setup`, 418
`ddi_mmap_get_model` — return data model type of current thread, 346
`ddi_model_convert_from` — determine data model type mismatch, 348
`ddi_model_convert_from` — Determine if there is a need to translate shared data structure contents, 348
`ddi_no_info` — returns `DDI_FAILURE`, a convenience for drivers implementing DLPI style 2 services, 351
`ddi_node_name` — return the devinfo node name, 350
 DDI page size conversions
 — `ddi_btop`, 220
 — `ddi_btopr`, 220
 — `ddi_ptob`, 220
`ddi_prop_exists` — check for the existence of a property, 360
`ddi_prop_get_int` — look up integer property, 362
`ddi_prop_lookup` — lookup property information, 365
`ddi_prop_lookup_byte_array` — lookup property information, 365
`ddi_prop_lookup_int_array` — lookup property information, 365
`ddi_prop_lookup_string` — lookup property information, 365
`ddi_prop_lookup_string_array` — lookup property information, 365
`ddi_prop_update` — update property information., 374
`ddi_prop_update_byte_array` — update property information., 374
`ddi_prop_update_int` — update property information., 374
`ddi_prop_update_int_array` — update property information., 374
`ddi_prop_update_string` — update property information., 374
`ddi_prop_update_string_array` — update property information., 374
 DDI property management
 create properties for leaf device drivers —
 `ddi_prop_create`, 357
 — `ddi_getlongprop`, 370
 — `ddi_getlongprop_buf`, 370
 — `ddi_getprop`, 370
 — `ddi_getpropplen`, 370
 — `ddi_prop_op`, 370
 modify properties for leaf device drivers —
 `ddi_prop_modify`, 357
 remove all properties for leaf device drivers —
 `ddi_prop_remove_all`, 357
 remove properties for leaf device drivers —
 `ddi_prop_remove`, 357
 remove properties for leaf device drivers —
 `ddi_prop_undefine`, 357
`ddi_put16` — write data to the device, 377
`ddi_put32` — write data to the device, 377
`ddi_put64` — write data to the device, 377
`ddi_put8` — write data to the device, 377
`ddi_regs_map_free` — free mapped register address space, 379
`ddi_regs_map_setup` — set up a mapping for a register address space, 380
`ddi_remove_intr` — remove an interrupt handler, 209
`ddi_remove_softintr` — remove a soft interrupt, 213
`ddi_removing_power`, 383
`ddi_rep_get16` — read data from the mapped memory address, device register or allocated DMA memory address, 385
`ddi_rep_get32` — read data from the mapped memory address, device register or allocated DMA memory address, 385
`ddi_rep_get64` — read data from the mapped memory address, device register or allocated DMA memory address, 385
`ddi_rep_get8` — read data from the mapped memory address, device register or allocated DMA memory address, 385
`ddi_rep_getb` — read data from the mapped memory address, device register or allocated DMA memory address, 385
`ddi_rep_getl` — read data from the mapped memory address, device register or allocated DMA memory address, 385

`ddi_rep_getll` — read data from the mapped memory address, device register or allocated DMA memory address, 385
`ddi_rep_getw` — read data from the mapped memory address, device register or allocated DMA memory address, 385
`ddi_rep_put16` — write data to the mapped memory address, device register or allocated DMA memory address, 388
`ddi_rep_put32` — write data to the mapped memory address, device register or allocated DMA memory address, 388
`ddi_rep_put64` — write data to the mapped memory address, device register or allocated DMA memory address, 388
`ddi_rep_put8` — write data to the mapped memory address, device register or allocated DMA memory address, 388
`ddi_rep_putb` — write data to the mapped memory address, device register or allocated DMA memory address, 388
`ddi_rep_putl` — write data to the mapped memory address, device register or allocated DMA memory address, 388
`ddi_rep_putll` — write data to the mapped memory address, device register or allocated DMA memory address, 388
`ddi_rep_putw` — write data to the mapped memory address, device register or allocated DMA memory address, 388
DDI self identifying devices, tell whether a device is self-identifying — `ddi_dev_is_sid`, 240
DDI soft interrupt handling
 add a soft interrupt — `ddi_add_softintr`, 213
 get soft interrupt block cookie — `ddi_get_soft_iblock_cookie`, 213
 remove a soft interrupt — `ddi_remove_softintr`, 213
DDI soft state utility routines
 allocate state structure — `ddi_soft_state_zalloc`, 395
 free soft state entry — `ddi_soft_state_free`, 395
 get pointer to soft state — `ddi_get_soft_state`, 395
 initialize state — `ddi_soft_state_init`, 395
DDI soft state utility routines (Continued)
 remove all state info — `ddi_soft_state_fini`, 395
DDI_SUSPEND, 383
`ddi_trigger_softintr` — trigger a soft interrupt, 213
`ddi_umem_alloc` — allocate kernel memory, 400
`ddi_umem_free` — allocate kernel memory, 400
`ddi_umem_lock` — Locks and unlocks memory pages, 404
`ddi_umem_zalloc` — allocate kernel memory, 400
default SCSI HBA probe function — `scsi_hba_probe`, 597
delay — delay process execution for a specified number of clock ticks, 406
deregister client from Card Services list — `csx_DeregisterClient`, 98
determine data model type mismatch — `ddi_model_convert_from`, 348
Device Driver Interface, *See* DDI
device mapping access entry point — `devmap_default_access`, 408
device switch tables, return function for insignificant entries — `nulldev`, 495
devices
 get major device number — `getmajor`, 445
 get minor device number — `getminor`, 446
 make device number from major and minor numbers — `makedevice`, 478
devices, non-pollable, error return function — `nochpoll`, 492
`devmap_default_access` — device mapping access entry point, 408
`devmap_devmem_setup` — Set driver memory mapping parameters, 411
 `devmap_devmem_setup()`, 410
 `devmap_umem_setup()`, 411
`devmap_do_ctxmgt` — perform device context switching on a mapping, 413
`devmap_load` — control the validation of memory address translations, 419
`devmap_set_ctx_timeout` — set context management timeout value, 416
`devmap_umem_setup` — Set driver memory mapping parameters, 411

devmap_unload — control the validation of memory address translations, 419

disksort — single direction elevator seek sort for buffers, 421

display a SCSI request sense message, scsi_vu_errmsg, 622

DMA attribute structure, 271

DMA mapping, the allowed burst sizes for — ddi_dma_burstsizes, 261

driver buffers

- copy data— ddi_copyin, 223
- copy data from driver — ddi_copyout, 226
- copy data from driver to user program — copyout, 90
- copy data from user program — copyin, 86

driver error messages, display an error message or panic the system — cmn_err, 76

driver privilege — drv_priv, 425

drv_getparm — retrieve kernel state information, 422

drv_hztousec — convert clock ticks to microseconds, 424

drv_priv — determine driver privilege, 425

drv_usecsthz — convert microseconds to clock ticks, 426

drv_usecwait — busy-wait for specified interval, 427

dupb — duplicate a message block descriptor, 428

duplicate a message — dupmsg, 431

duplicate a message block descriptor — dupb, 428

duplicate access handle — csx_DupHandle, 99

dupmsg — duplicate a message, 431

E

enable/disable accesses to the PCI Local Bus Configuration space.

- pci_config_setup, 512
- pci_config_teardown, 512

error return codes converted to text strings — csx_Error2Text, 101

error return function for illegal entries — nodev, 493

event mask

- return client event mask — csx_GetEventMask, 204
- set client event mask — csx_SetEventMask, 204

events converted to text strings — csx_Event2Text, 102

expression verification, — assert, 47

F

find first (last) bit set in a long integer —

- ddi_ffs, 299
- ddi_fls, 299

first CIS tuple — csx_GetFirstTuple, 107

flushband — flush messages for specified priority band, 436

free access handle — csx_FreeHandle, 103

free DMA handle, —

- ddi_dma_free_handle, 270

free mapped register address space —

- ddi_regs_map_free, 379

free previously allocated memory —

- ddi_dma_mem_free, 278

free space — rmfree, 567

freerbuf — free a raw buffer header, 441

freeze, thaw the state of a stream —

- freezestr, 442
- unfreezestr, 442

freezestr — freeze, thaw the state of a stream, 442

G

generic tuple parser — csx_ParseTuple, 172

get interrupt block cookie, —

- ddi_get_iblock_cookie, 209

get kernel internal minor number from an external dev_t, scsi_vu_errmsg, 305

get next DMA segment —

- ddi_dma_nextseg, 283

get next DMA window —

- ddi_dma_nextwin, 285

get soft interrupt block cookie, —

- ddi_get_soft_iblock_cookie, 213

gethrtime, 444

getmajor — get major device number, 445
 getminor — get minor device number, 446
 getrbuf — get a raw buffer header, 450
 gld_mac_alloc — allocate a GLD mac_info structure, 451
 gld_mac_free — free a GLD mac_info structure, 451
 gld_rcv — pass the inbound packet upstream, 452
 gld_register — link the driver with the GLD framework, 451
 gld_sched — reschedule stalled outbound packets, 452
 gld_unregister — unlink the driver from the GLD framework, 452

H

handle variable argument list
 — va_arg, 655
 — va_copy, 655
 — va_end, 655
 — va_start, 655
 high resolution time, 444

I

I/O, block, suspend processes pending completion — biowait, 64
 I/O, buffer, release buffer and notify processes — biodone, 56
 I/O, paged request
 allocate virtual address space —
 bp_mapin, 65
 deallocate virtual address space —
 bp_mapout, 66
 I/O, physical
 — minphys, 517
 — physio, 517
 I/O error, return — geterror, 443
 I/O resources
 release I/O resources — csx_ReleaseIO, 188
 request I/O resources — csx_RequestIO, 188
 id32_alloc — 32-bit driver ID management routines, 455

id32_free — 32-bit driver ID management routines, 455
 id32_lookup — 32-bit driver ID management routines, 455
 inb — read from an I/O port, 456
 inform the system that a device's component is required. — ddi_dev_is_needed, 238
 initialize a named kstat —
 kstat_named_init, 472
 initialize a buffer structure — bioinit, 60
 inl — read from an I/O port, 456
 interrupt handling
 add an interrupt — ddi_add_intr, 209
 get interrupt block cookie —
 ddi_get_iblock_cookie, 209
 remove an interrupt — ddi_remove_intr, 209
 inw — read from an I/O port, 456
 IOC_CONVERT_FROM — Determine if there is a need to translate M_IOCTL contents, 460
 IRQ resource
 release IRQ resource — csx_ReleaseIRQ, 193
 request IRQ resource —
 csx_RequestIRQ, 193

K

kernel memory cache allocator operations —
 kmem_cache_alloc, 463
 kernel memory cache allocator operations —
 kmem_cache_create, 463
 kernel memory cache allocator operations —
 kmem_cache_destroy, 463
 kernel memory cache allocator operations —
 kmem_cache_free, 463
 kernel address locations, between locations —
 bcopy, 51
 kernel addresses, get page frame number —
 hat_getkpfnum, 454
 Kernel interfaces for device ids
 — ddi_devid_compare, 235
 — ddi_devid_free, 235
 — ddi_devid_init, 235
 — ddi_devid_register, 235
 — ddi_devid_sizeof, 235
 — ddi_devid_unregister, 235
 — ddi_devid_valid, 235

- kernel modules, dynamic loading
 - add loadable module — `mod_install`, 484
 - query loadable module — `mod_info`, 484
 - remove loadable module —
 - `mod_remove`, 484
- kernel state information — `drv_getparm`, 422
- `kmem_alloc` — allocate kernel memory, 461
- `kmem_cache_alloc` — kernel memory cache allocator operations, 463
- `kmem_cache_create` — kernel memory cache allocator operations, 463
- `kmem_cache_destroy` — kernel memory cache allocator operations, 463
- `kmem_cache_free` — kernel memory cache allocator operations, 463
- `kmem_free` — allocate kernel memory, 461
- `kmem_zalloc` — allocate kernel memory, 461
- `kstat_create` — create and initialize a new `kstat`, 469
- `kstat_delete` — remove a `kstat` from the system, 470
- `kstat_install` — add a fully initialized `kstat` to the system, 471
- `kstat_named_init` — initialize a named `kstat`, 472
- `kstat_queue` — update I/O `kstat` statistics, 473
- `kstat_runq_back_to_waitq` — update I/O `kstat` statistics, 473
- `kstat_runq_enter` — update I/O `kstat` statistics, 473
- `kstat_runq_exit` — update I/O `kstat` statistics, 473
- `kstat_waitq_enter` — update I/O `kstat` statistics, 473
- `kstat_waitq_exit` — update I/O `kstat` statistics, 473
- `kstat_waitq_to_runq` — update I/O `kstat` statistics, 473

L

- `linkb` — concatenate two message blocks, 475
- Locks and unlocks memory pages —
 - `ddi_umem_lock`, 404
- look up integer property —
 - `ddi_prop_get_int`, 362

- look up per-device-type `scsi-options` property
 - `scsi_get_device_type_scsi_options`, 587
- lookup property information
 - `ddi_prop_lookup`, 365
 - `ddi_prop_lookup_byte_array`, 365
 - `ddi_prop_lookup_int_array`, 365
 - `ddi_prop_lookup_string`, 365
 - `ddi_prop_lookup_string_array`, 365

M

- major device number, 296
- `makedevice` — make device number from major and minor numbers, 478
- map memory area on PC Card —
 - `csx_MapMemPage`, 120
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_boolean`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_byte`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_byte_array`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_int16`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_int16_array`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_int32`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_int32_array`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_int64`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_int64_array`, 500
- match name and type indicated by the interface name and retrieve data value —
 - `nvlist_lookup_string`, 500

match name and type indicated by the interface
 name and retrieve data value —
 nvlist_lookup_string_array, 500
 match name and type indicated by the interface
 name and retrieve data value —
 nvlist_lookup_uint16, 500
 match name and type indicated by the interface
 name and retrieve data value —
 nvlist_lookup_uint16_array, 500
 match name and type indicated by the interface
 name and retrieve data value —
 nvlist_lookup_uint32, 500
 match name and type indicated by the interface
 name and retrieve data value —
 nvlist_lookup_uint32_array, 500
 match name and type indicated by the interface
 name and retrieve data value —
 nvlist_lookup_uint64, 500
 match name and type indicated by the interface
 name and retrieve data value —
 nvlist_lookup_uint64_array, 500
 max — return the larger of two integers, 479
 memory, clear for a given number of bytes —
 bzero, 72
 min — return the lesser of two integers, 480
 minor node for device, create —
 ddi_create_minor_node, 229
 modify PC Card configuration —
 csx_ModifyConfiguration, 121
 modify window attributes —
 csx_ModifyWindow, 123
 mt-streams — STREAMS multithreading, 487
 mutex routines
 — mutex, 489
 — mutex_destroy, 489
 — mutex_enter, 489
 — mutex_exit, 489
 — mutex_init, 489
 — mutex_owned, 489
 — mutex_tryenter, 489
 mutual exclusion lock, *See* mutex

N

next CIS tuple — csx_GetNextTuple, 107
 nodes, create minor nodes for client —
 csx_MakeDeviceNode, 117

Notify pm framework of autonomous power
 level change — pm_power_has_changed, 524
 notify target driver of bus resets —
 scsi_reset_notify, 615
 nvlist_lookup_boolean — match name and type
 indicated by the interface name and retrieve
 data value, 500
 nvlist_lookup_byte — match name and type
 indicated by the interface name and retrieve
 data value, 500
 nvlist_lookup_byte_array — match name and
 type indicated by the interface name and
 retrieve data value, 500
 nvlist_lookup_int16 — match name and type
 indicated by the interface name and retrieve
 data value, 500
 nvlist_lookup_int16_array — match name and
 type indicated by the interface name and
 retrieve data value, 500
 nvlist_lookup_int32 — match name and type
 indicated by the interface name and retrieve
 data value, 500
 nvlist_lookup_int32_array — match name and
 type indicated by the interface name and
 retrieve data value, 500
 nvlist_lookup_int64 — match name and type
 indicated by the interface name and retrieve
 data value, 500
 nvlist_lookup_int64_array — match name and
 type indicated by the interface name and
 retrieve data value, 500
 nvlist_lookup_string — match name and type
 indicated by the interface name and retrieve
 data value, 500
 nvlist_lookup_string_array — match name and
 type indicated by the interface name and
 retrieve data value, 500
 nvlist_lookup_uint16 — match name and type
 indicated by the interface name and retrieve
 data value, 500
 nvlist_lookup_uint16_array — match name and
 type indicated by the interface name and
 retrieve data value, 500
 nvlist_lookup_uint32 — match name and type
 indicated by the interface name and retrieve
 data value, 500

nvlist_lookup_uint32_array — match name and type indicated by the interface name and retrieve data value, 500
 nvlist_lookup_uint64 — match name and type indicated by the interface name and retrieve data value, 500
 nvlist_lookup_uint64_array — match name and type indicated by the interface name and retrieve data value, 500
 nvlist_remove — remove name-value pairs, 504
 nvlist_remove_all — remove name-value pairs, 504

O

obtain DDI information —
 csx_CS_DDI_Info, 96
 OTHERQ — get pointer to queue's partner queue, 507
 outb — write to an I/O port, 508
 outl — write to an I/O port, 508
 outw — write to an I/O port, 508

P

panic state — ddi_in_panic, 312
 parse Battery Replacement Date tuple —
 csx_Parse_CISTPL_BATTERY, 125
 parse Byte Order tuple —
 csx_Parse_CISTPL_BYTEORDER, 126
 parse Card Configuration Table tuple —
 csx_Parse_CISTPL_CFTABLE_ENTRY, 128
 parse Card Initialization Date tuple —
 csx_Parse_CISTPL_DATE, 137
 parse Configuration tuple —
 csx_Parse_CISTPL_CONFIG, 134
 parse Data Organization tuple —
 csx_Parse_CISTPL_ORG, 166
 parse Data Recording Format tuple —
 csx_Parse_CISTPL_FORMAT, 145
 parse Device Geo A tuple —
 csx_Parse_CISTPL_DEVICEGEO_A, 143
 parse Device Geo tuple —
 csx_Parse_CISTPL_DEVICEGEO, 141

parse Device Information tuple
 for Attribute Memory —
 csx_Parse_CISTPL_DEVICE_A, 138
 for Common Memory —
 csx_Parse_CISTPL_DEVICE, 138
 parse Function Extension tuple —
 csx_Parse_CISTPL_FUNC, 147
 parse Function Identification tuple —
 csx_Parse_CISTPL_FUNCID, 155
 parse Geometry tuple —
 csx_Parse_CISTPL_GEOMETRY, 157
 parse JEDEC Identifier tuple
 for Attribute Memory —
 csx_Parse_CISTPL_JEDEC_A, 158
 for Common Memory —
 csx_Parse_CISTPL_JEDEC_C, 158
 parse Level-1 Version/Product Information tuple —
 csx_Parse_CISTPL_VERS_1, 170
 parse Level-2 Version and Information tuple —
 csx_Parse_CISTPL_VERS_2, 171
 parse Link Target tuple —
 csx_Parse_CISTPL_LINKTARGET, 160
 parse Long Link A tuple, —
 csx_Parse_CISTPL_LONGLINK_A, 161
 parse Long Link C tuple, —
 csx_Parse_CISTPL_LONGLINK_C, 161
 parse Manufacturer Identification tuple —
 csx_Parse_CISTPL_MANFID, 165
 parse Multi-Function tuple —
 csx_Parse_CISTPL_LONGLINK_MFC, 163
 parse Other Condition Device Information tuple
 for Attribute Memory —
 csx_Parse_CISTPL_DEVICE_OA, 138
 for Common Memory —
 csx_Parse_CISTPL_DEVICE_OC, 138
 parse Software Interleaving tuple —
 csx_Parse_CISTPL_SWIL, 169
 parse Special Purpose tuple —
 csx_Parse_CISTPL_SPCL, 167
 parser, for tuples (generic) —
 csx_ParseTuple, 172
 pci_config_get16 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510
 pci_config_get32 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_get64 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_get8 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_getb — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_getl — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_getll — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_getw — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_put16 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_put32 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_put64 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_put8 — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_putb — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_putl — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_putll — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_putw — read or write single datum of various sizes to the PCI Local Bus Configuration space, 510

pci_config_setup — enable/disable accesses to the PCI Local Bus Configuration space., 512

pci_config_teardown — enable/disable accesses to the PCI Local Bus Configuration space., 512

pci_report_pmcap—Report power management capability of a PCI device, 513

perform device context switching on a mapping — devmap_do_ctxmgt, 413

pm_busy_component — control device components' availability for Power Management, 519

pm_idle_component — control device components' availability for Power Management, 519

pm_power_has_changed – Notify pm framework of autonomous power level change, 524

pm_raise_power – Raise or lower power of components, 526

pm_trans_check – advisory check for device power cycles, 529

pollwake up — inform a process that an event has occurred, 531

proc_ref — send a signal to a process, 532

proc_signal — send a signal to a process, 532

proc_unref — send a signal to a process, 532

put — call a STREAMS put procedure, 537

Q

qassociate — associate STREAMS queue with bottom driver in queue, 547

qbufcall — call a function when a buffer becomes available, 549

qtimeout — execute a function after a specified length of time, 555

qunbufcall — cancel a pending qbufcall request, 556

quntimeout — cancel previous timeout function call, 557

qwait — STREAMS wait routines, 558

qwait_sig — STREAMS wait routines, 558

qwriter — asynchronous STREAMS perimeter upgrade, 560

R

Raise or lower power of components – pm_raise_power, 526

- raw buffer
 - free a raw buffer header — `freerbuf`, 441
 - get a raw buffer header — `getrbuf`, 450
- RD — get pointer to the read queue, 561
- read from an I/O port — `inb`, 456
- read from an I/O port — `inl`, `inl`, 456
- read from an I/O port — `inw`, `inw`, 456
- read from an I/O port — `repinsb`, `repinsb`, 456
- read from an I/O port — `repinsd`, `repinsd`, 456
- read from an I/O port — `repinsw`, `repinsw`, 456
- read data from mapped device in the memory space or allocated DMA memory
 - `ddi_mem_get16`, 339
 - `ddi_mem_get32`, 339
 - `ddi_mem_get64`, 339
 - `ddi_mem_get8`, 339
 - `ddi_mem_rep_get16`, 342
 - `ddi_mem_rep_get32`, 342
 - `ddi_mem_rep_get64`, 342
 - `ddi_mem_rep_get8`, 342
- read data from the device
 - `ddi_get16`, 300
 - `ddi_get32`, 300
 - `ddi_get64`, 300
 - `ddi_get8`, 300
- read data from the mapped device register in I/O space
 - `ddi_io_get16`, 314
 - `ddi_io_get32`, 314
 - `ddi_io_get8`, 314
 - `ddi_io_getb`, 314
 - `ddi_io_getl`, 314
 - `ddi_io_getw`, 314
- read data from the mapped memory address, device register or allocated DMA memory address
 - `ddi_rep_get16`, 385
 - `ddi_rep_get32`, 385
 - `ddi_rep_get64`, 385
 - `ddi_rep_get8`, 385
 - `ddi_rep_getb`, 385
 - `ddi_rep_getl`, 385
 - `ddi_rep_getll`, 385
 - `ddi_rep_getw`, 385
- read from device register
 - `csx_Get16`, 104
 - `csx_Get32`, 104
 - `csx_Get64`, 104
- read from device register (Continued)
 - `csx_Get8`, 104
- read multiple data from the mapped device register in I/O space
 - `ddi_io_rep_get16`, 321
 - `ddi_io_rep_get32`, 321
 - `ddi_io_rep_get8`, 321
 - `ddi_io_rep_getb`, 321
 - `ddi_io_rep_getl`, 321
 - `ddi_io_rep_getw`, 321
- read or write a PC Card Configuration Register
 - `csx_AccessConfigurationRegister`, 92
- read or write single datum of various sizes to the PCI Local Bus Configuration space
 - `pci_config_get16`, 510
 - `pci_config_get32`, 510
 - `pci_config_get64`, 510
 - `pci_config_get8`, 510
 - `pci_config_getb`, 510
 - `pci_config_getl`, 510
 - `pci_config_getll`, 510
 - `pci_config_getw`, 510
 - `pci_config_put16`, 510
 - `pci_config_put32`, 510
 - `pci_config_put64`, 510
 - `pci_config_put8`, 510
 - `pci_config_putb`, 510
 - `pci_config_putl`, 510
 - `pci_config_putll`, 510
 - `pci_config_putw`, 510
- read repetitively from device register
 - `csx_RepGet16`, 180
 - `csx_RepGet32`, 180
 - `csx_RepGet64`, 180
 - `csx_RepGet8`, 180
- readers/writer lock functions
 - `rw_destroy`, 571
 - `rw_downgrade`, 571
 - `rw_enter`, 571
 - `rw_exit`, 571
 - `rw_init`, 571
 - `rw_read_locked`, 571
 - `rw_tryenter`, 571
 - `rw_tryupgrade`, 571
 - `rwlock`, 571
- register client — `csx_RegisterClient`, 175
- release client event mask — `csx_ReleaseSocketMask`, 196

- release I/O resources — `csx_ReleaseIO`, 188
- release IRQ resource — `csx_ReleaseIRQ`, 193
- release window resources —
 - `csx_ReleaseWindow`, 198
- release configuration on PC Card —
 - `csx_ReleaseConfiguration`, 178
- remove name-value pairs —
 - `nvlist_remove`, 504
- remove name-value pairs —
 - `nvlist_remove_all`, 504
- remove a kstat from the system —
 - `kstat_delete`, 470
- remove a soft interrupt, —
 - `ddi_remove_softintr`, 213
- remove an interrupt handler, —
 - `ddi_remove_intr`, 209
- remove client from Card Services list —
 - `csx_DeregisterClient`, 98
- `repinsb` — read from an I/O port, 456
- `repinsd` — read from an I/O port, 456
- `repinsw` — read from an I/O port, 456
- Report a hardware failure, 244
- Report power management capability of a PCI
 - device-`pci_report_pmcap`, 513
- `repoutsb` — write to an I/O port, 508
- `repoutsd` — write to an I/O port, 508
- `repoutsw` — write to an I/O port, 508
- request client event mask —
 - `csx_RequestSocketMask`, 196
- request I/O resources — `csx_RequestIO`, 188
- request IRQ resource — `csx_RequestIRQ`, 193
- request window resources —
 - `csx_RequestWindow`, 198
- reset a function on a PC card —
 - `csx_ResetFunction`, 203
- resource map
 - allocate resource maps — `rmallocmap`, 565
 - free resource maps — `rmallocmap`, 565
- retrieve number of DMA windows —
 - `ddi_dma_numwin`, 287
- retrieve subsequent DMA cookie —
 - `ddi_dma_nextcookie`, 281
- return client event mask —
 - `csx_GetEventMask`, 204
- return client
 - `csx_GetFirstClient`, 105
 - `csx_GetNextClient`, 105
- return current access handle offset —
 - `csx_GetHandleOffset`, 109
- return data model type of current thread —
 - `ddi_mmap_get_model`, 346
- return data portion of tuple —
 - `csx_GetTupleData`, 115
- return driver binding name
 - `ddi_binding_name`, 219
 - `ddi_get_name`, 219
- return index matching capability string —
 - `scsi_hba_lookup_capstr`, 593
- return normalized driver name —
 - `ddi_driver_name`, 297
- return physical socket number —
 - `csx_MapLogSocket`, 119
- return status of PC Card and socket —
 - `csx_GetStatus`, 111
- return the devinfo node name —
 - `ddi_node_name`, 350
- return the larger of two integers — `max`, 479
- return the lesser of two integers — `min`, 480
- return tuple
 - first CIS tuple — `csx_GetFirstTuple`, 107
 - next CIS tuple — `csx_GetNextTuple`, 107
- returns `DDI_FAILURE`, a convenience for
 - drivers implementing DLPI style 2 services
 - `ddi_no_info`, 351
- returns size of a buffer structure — `biosize`, 63
- returns the current time in seconds,
 - `ddi_get_time`, 311
- returns the process ID, `ddi_get_pid`, 310
- returns the value of `lbolt`, returns the value of
 - `lbolt`, 308
- `rmalloc` — allocate space from a resource
 - map, 562
- `rmalloc_wait` — allocate space from a resource
 - map, 566
- `rmfree` — free space back into a resource
 - map, 567

S

- SAMESTR** — test if next queue is in the same
 - stream, 574
- SCSI Host Bus Adapter system initialization

 - and completion routines
 - `scsi_hba_init`, 592**

`scsi_abort` — abort a SCSI command, 575
`scsi_alloc_consistent_buf` — scsi dma utility for allocating an I/O buffer for SCSI DMA, 576
`scsi_cname` — decode SCSI commands, 578
 SCSI commands, make packet
 — `makecom`, 476
 — `makecom_g0`, 476
 — `makecom_g0_s`, 476
 — `makecom_g1`, 476
 — `makecom_g5`, 476
`scsi_destroy_pkt` — free an allocated SCSI packet and its DMA resource, 580
 SCSI dma utility routines
 — `scsi_dmafree`, 581
 — `scsi_dmaget`, 581
`scsi_dname` — decode SCSI peripheral device type, 578
`scsi_errmsg` — display a SCSI request sense message, 583
`scsi_free_consistent_buf` — free a previously allocated SCSI DMA I/O buffer, 586
`scsi_get_device_type_scsi_options` — look up per-device-type scsi-options property, 587
`scsi_hba_attach` — SCSI HBA attach and detach routines, 589
 SCSI HBA attach and detach routines
 — `scsi_hba_attach`, 589
 — `scsi_hba_attach_setup`, 589
 — `scsi_hba_detach`, 589
`scsi_hba_attach_setup` — SCSI HBA attach and detach routines, 589
`scsi_hba_detach` — SCSI HBA attach and detach routines, 589
`scsi_hba_fini` — SCSI Host Bus Adapter system completion routines, 592
`scsi_hba_init` — SCSI Host Bus Adapter system initialization routines, 592
`scsi_hba_lookup_capstr` — return index matching capability string, 593
`scsi_hba_pkt_alloc` — allocate and free a `scsi_pkt` structure, 595
`scsi_hba_pkt_free` — allocate and free a `scsi_pkt` structure, 595
`scsi_hba_probe` — default SCSI HBA probe function, 597
`scsi_hba_tran_alloc` — allocate and free transport structures, 598
`scsi_hba_tran_free` — allocate and free transport structures, 598
`scsi_ifgetcap` — get SCSI transport capability, 599
`scsi_ifsetcap` — set SCSI transport capability, 599
`scsi_init_pkt` — prepare a complete SCSI packet, 603
`scsi_log` — display a SCSI-device-related message, 607
`scsi_mname` — decode SCSI messages, 578
 SCSI packet
 allocate a SCSI packet in iopb map —
 `get_pktiopb`, 447
 free a packet in iopb map —
 `free_pktiopb`, 447
 free an allocated SCSI packet and its DMA resource — `scsi_destroy_pkt`, 580
 SCSI packet utility routines
 — `scsi_pktalloc`, 609
 — `scsi_pktfree`, 609
 — `scsi_resalloc`, 609
 — `scsi_resfree`, 609
`scsi_poll` — run a polled SCSI command on behalf of a target driver, 611
`scsi_probe` — utility for probing a scsi device, 612
`scsi_reset` — reset a SCSI bus or target, 614
`scsi_reset_notify` — notify target driver of bus resets, 615
`scsi_rname` — decode SCSI packet completion reasons, 578
`scsi_setup_cdb` — setup SCSI command descriptor block (CDB), 616
`scsi_slave` — utility for SCSI target drivers to establish the presence of a target, 617
`scsi_sname` — decode SCSI sense keys, 578
`scsi_sync_pkt` — synchronize CPU and I/O views of memory, 619
`scsi_transport` — request by a target driver to start a SCSI command, 620
`scsi_unprobe` — free resources allocated during initial probing, 621
`scsi_unslave` — free resources allocated during initial probing, 621
`scsi_vu_errmsg`, display a SCSI request sense message, 622

- semaphore functions
 - sema_destroy, 625
 - sema_init, 625
 - sema_p, 625
 - sema_p_sig, 625
 - sema_try, 625
 - sema_v, 625
 - semaphore, 625
- send a signal to a process
 - proc_ref, 532
 - proc_signal, 532
 - proc_unref, 532
- set client event mask —
 - csx_RequestSocketMask, 196
- set client event mask — csx_SetEventMask, 204
- set current access handle offset —
 - csx_SetHandleOffset, 206
- Set driver memory mapping parameters
 - devmap_devmem_setup, 411
 - devmap_umem_setup, 411
- Set the device attributes for the mapping —
 - ddi_mapdev_set_device_acc_attr, 333
- set up a mapping for a register address space —
 - ddi_regs_map_setup, 380
- setup SCSI command descriptor block (CDB) —
 - scsi_setup_cdb, 616
- single direction elevator seek sort for buffers —
 - disksort, 421
- size in bytes
 - convert size in pages — ptob, 534
 - convert to size in memory pages (round down) — btob, 67
 - convert to size in memory pages (round up) — btob, 68
- socket number, return physical socket number
 - csx_MapLogSocket, 119
- soft interrupt handling
 - add a soft interrupt — ddi_add_softintr, 213
 - get soft interrupt block cookie —
 - ddi_get_soft_iblock_cookie, 213
 - remove a soft interrupt —
 - ddi_remove_softintr, 213
 - trigger a soft interrupt —
 - ddi_trigger_softintr, 213
- sprintf — format characters in memory, 627
- status of PC Card and socket —
 - csx_GetStatus, 111
- STREAMS wait routines — qwait,
 - qwait_sig, 558
- STREAMS ioctl blocks, allocate — mkioch, 481
- STREAMS message blocks
 - attach a user-supplied data buffer in place —
 - esballoc, 433
 - call a function when a buffer becomes available — bufcall, 69, 649
 - call a function when a buffer becomes available — qbufcall, 549, 556
 - call function when buffer is available —
 - esbbcall, 435
 - concatenate bytes in a message —
 - msgpullup, 486
 - concatenate bytes in a message —
 - pullupmsg, 535
 - concatenate two — linkb, 475
 - copy — copyb, 84
 - erase the contents of a buffer — clrbuf, 75
 - free all message blocks in a message —
 - freemsg, 440
 - free one — freeb, 439
 - remove from head of message —
 - unlinkb, 650
 - remove one from a message — rmvb, 568
- STREAMS message queue, insert a message into a queue — insq, 458
- STREAMS message queues, 48
- STREAMS Message queues, get next message —
 - getq, 449
- STREAMS message queues
 - reschedule a queue for service —
 - enableok, 432
 - test for room — canputnext, 74
 - test for room — canput, 73
- STREAMS messages
 - copy a message — copymsg, 88
 - flush for specified priority band —
 - flushband, 436
 - remove from queue — flushq, 437
 - remove from queue — rmvq, 569
 - return the number of bytes in a message —
 - msgdsize, 485
 - submit messages to the log driver —
 - strlog, 634
 - test whether a message is a data message —
 - datamsg, 208
 - trim bytes — adjmsg, 40

STREAMS multithreading

- `mt-streams`, 487
- `qbufcall` — call a function when a buffer becomes available, 549
- `qtimeout` — execute a function after a specified length of time, 555
- `qunbufcall` — cancel a pending `qbufcall` request, 556
- `qtimeout` — cancel previous timeout function call, 557
- `qwait`, `qwait_sig` — STREAMS wait routines, 558
- `qwriter` — asynchronous STREAMS perimeter upgrade, 560

STREAMS put and service procedures

- disable — `qprocsoff`, 551
- enable — `qprocon`, 551

STREAMS queues

- change information about a queue or band of the queue — `strqset`, 637
- enable a queue — `qenable`, 550
- get pointer to queue's partner queue — `OTHERQ`, 507
- get pointer to the read queue — `RD`, 561
- get information about a queue or band of the queue — `strqget`, 636
- number of messages on a queue — `qsize`, 554
- place a message at the head of a queue — `putbq`, 538
- prevent a queue from being scheduled — `noenable`, 494
- put a message on a queue — `putq`, 546
- send a control message to a queue — `putctl`, 540
- send a control message to a queue — `putnextctl`, 544
- send a control message with a one-byte parameter to a queue — `putctl1`, 539
- send a control message with a one-byte parameter to a queue — `putnextctl1`, 543
- send a message on a stream in the reverse direction — `qreply`, 552
- send a message to the next queue — `putnext`, 542
- test if next queue is in the same stream — `SAMESTR`, 574

STREAMS queues (Continued)

- test for flow control in specified priority band — `bcanput`, 49
- STREAMS write queues, get pointer for this module or driver — `WR`, 660
- `STRUCT_DECL`, 32-bit application data access macros, 638
- `swab` — swap bytes in 16-bit halfwords, 643
- synchronize CPU and I/O views of memory — `ddi_dma_sync`, 293
- synchronize CPU and I/O views of memory — `scsi_sync_pkt`, 619
- system DMA engine functions
 - `ddi_dmae`, 266
 - `ddi_dmae_1stparty`, 266
 - `ddi_dmae_alloc`, 266
 - `ddi_dmae_disable`, 266
 - `ddi_dmae_enable`, 266
 - `ddi_dmae_getattr`, 266
 - `ddi_dmae_getcnt`, 266
 - `ddi_dmae_getlim`, 266
 - `ddi_dmae_prog`, 266
 - `ddi_dmae_release`, 266
 - `ddi_dmae_stop`, 266
- system event, logging of, 325

T

- `testb` — check for an available buffer, 644
- `timeout` — execute a function after a specified length of time, 646
- `timeout`, cancel previous timeout function call — `untimeout`, 651
- trigger a soft interrupt, — `ddi_trigger_softintr`, 213
- tuple
 - first CIS tuple — `csx_GetFirstTuple`, 107
 - next CIS tuple — `csx_GetNextTuple`, 107
 - return data portion of tuple — `csx_GetTupleData`, 115
- tuple entry
 - generic tuple parser — `csx_ParseTuple`, 172
 - parse Device Information tuple for Attribute Memory — `csx_Parse_CISTPL_DEVICE_A`, 138

tuple entry (Continued)

- parse Device Information tuple for Common Memory — `csx_Parse_CISTPL_DEVICE`, 138
- parse JEDEC Identifier tuple for Attribute Memory —
 - `csx_Parse_CISTPL_JEDEC_A`, 158
- parse JEDEC Identifier tuple for Common Memory —
 - `csx_Parse_CISTPL_JEDEC_C`, 158
- parse Long Link A tuple —
 - `csx_Parse_CISTPL_LONGLINK_A`, 161
- parse Long Link C tuple —
 - `csx_Parse_CISTPL_LONGLINK_C`, 161
- parse Other Condition Device Information tuple for Attribute Memory —
 - `csx_Parse_CISTPL_DEVICE_OA`, 138
- parse Other Condition Device Information tuple for Common Memory —
 - `csx_Parse_CISTPL_DEVICE_OC`, 138
- parse Battery Replacement Date tuple —
 - `csx_Parse_CISTPL_BATTERY`, 125
- parse Byte Order tuple —
 - `csx_Parse_CISTPL_BYTEORDER`, 126
- parse Card Configuration Table tuple —
 - `csx_Parse_CISTPL_CFTABLE_ENTRY`, 128
- parse Card Initialization Date tuple —
 - `csx_Parse_CISTPL_DATE`, 137
- parse Configuration tuple —
 - `csx_Parse_CISTPL_CONFIG`, 134
- parse Data Organization tuple —
 - `csx_Parse_CISTPL_ORG`, 166
- parse Data Recording Format tuple —
 - `csx_Parse_CISTPL_FORMAT`, 145
- parse Device Geo A tuple —
 - `csx_Parse_CISTPL_DEVICE_A`, 143
- parse Device Geo tuple —
 - `csx_Parse_CISTPL_DEVICEGEO`, 141
- parse Function Extension tuple —
 - `csx_Parse_CISTPL_FUNCCE`, 147
- parse Function Identification tuple —
 - `csx_Parse_CISTPL_FUNCID`, 155
- parse Geometry tuple —
 - `csx_Parse_CISTPL_GEOMETRY`, 157
- parse Level-1 Version/Product Information tuple — `csx_Parse_CISTPL_VERS_1`, 170
- parse Level-2 Version and Information tuple — `csx_Parse_CISTPL_VERS_2`, 171

tuple entry (Continued)

- parse Link Target tuple —
 - `csx_Parse_CISTPL_LINKTARGET`, 160
- parse Manufacturer Identification tuple —
 - `csx_Parse_CISTPL_MANFID`, 165
- parse Multi-Function tuple —
 - `csx_Parse_CISTPL_LONGLINK_MFC`, 163
- parse Software Interleaving tuple —
 - `csx_Parse_CISTPL_SWIL`, 169
- parse Special Purpose tuple —
 - `csx_Parse_CISTPL_SPCL`, 167

U

uio structure

- add character — `ureadc`, 653
- remove a character — `uwritec`, 654
- `uiomove` — copy kernel data using uio structure, 648
- unbinds the address in a DMA handle —
 - `ddi_dma_unbind_handle`, 295
- `unfreezestr` — freeze, thaw the state of a stream, 442
- uninitialize a buffer structure — `biofini`, 59
- update I/O kstat statistics
 - `kstat_queue`, 473
 - `kstat_runq_back_to_waitq`, 473
 - `kstat_runq_enter`, 473
 - `kstat_runq_exit`, 473
 - `kstat_waitq_enter`, 473
 - `kstat_waitq_exit`, 473
 - `kstat_waitq_to_runq`, 473
- update property information.
 - `ddi_prop_update`, 374
 - `ddi_prop_update_byte_array`, 374
 - `ddi_prop_update_int`, 374
 - `ddi_prop_update_int_array`, 374
 - `ddi_prop_update_string`, 374
 - `ddi_prop_update_string_array`, 374

V

- `va_arg` — handle variable argument list, 655
- `va_copy` — handle variable argument list, 655
- `va_end` — handle variable argument list, 655
- `va_start` — handle variable argument list, 655

validate Card Information Structure (CIS) —
 csx_ValidateCIS, 207
 virtual address, return mapped virtual address
 — csx_GetMappedAddr, 110
 vsprintf — format characters in memory, 657

W

window resources
 release window resources —
 csx_ReleaseWindow, 198
 request window resources —
 csx_RequestWindow, 198
 write data to mapped device in the memory
 space or allocated DMA memory
 — ddi_mem_put16, 340
 — ddi_mem_put32, 340
 — ddi_mem_put64, 340
 — ddi_mem_put8, 340
 — ddi_mem_rep_put16, 344
 — ddi_mem_rep_put32, 344
 — ddi_mem_rep_put64, 344
 — ddi_mem_rep_put8, 344
 write data to the device
 — ddi_put16, 377
 — ddi_put32, 377
 — ddi_put64, 377
 — ddi_put8, 377
 write data to the mapped device register in I/O
 space
 — ddi_io_put16, 319
 — ddi_io_put32, 319
 — ddi_io_put8, 319
 — ddi_io_putb, 319
 — ddi_io_putl, 319
 — ddi_io_putw, 319
 write data to the mapped memory address,
 device register or allocated DMA memory
 address
 — ddi_rep_put16, 388
 — ddi_rep_put32, 388
 — ddi_rep_put64, 388
 — ddi_rep_put8, 388
 — ddi_rep_putb, 388
 — ddi_rep_putl, 388
 — ddi_rep_putll, 388
 — ddi_rep_putw, 388

write multiple data to the mapped device
 register in I/O space
 — ddi_io_rep_put16, 323
 — ddi_io_rep_put32, 323
 — ddi_io_rep_put8, 323
 — ddi_io_rep_putb, 323
 — ddi_io_rep_putl, 323
 — ddi_io_rep_putw, 323
 write or read a PC Card Configuration Register
 — csx_AccessConfigurationRegister, 92
 write repetitively to device register
 — csx_RepPut16, 182
 — csx_RepPut32, 182
 — csx_RepPut64, 182
 — csx_RepPut8, 182
 write to an I/O port
 — outb, 508
 — outl, 508
 — outw, 508
 — repoutsb, 508
 — repoutsd, 508
 — repoutsw, 508
 write to device register
 — csx_Put16, 174
 — csx_Put32, 174
 — csx_Put64, 174
 — csx_Put8, 174

Z

zero fill the device register —
 ddi_device_zero, 233