



man pages section 9: DDI and DKI Properties and Data Structures

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

Part No: 817-0703-10
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 REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



030827@6671



Contents

Preface 7

Introduction 13

Intro(9S) 14

Data Structures for Drivers 17

aio_req(9S) 18

buf(9S) 19

cb_ops(9S) 22

copyreq(9S) 24

copyresp(9S) 25

datab(9S) 26

ddi_device_acc_attr(9S) 27

ddi_dma_attr(9S) 31

ddi_dma_cookie(9S) 34

ddi_dmae_req(9S) 35

ddi_dma_lim_sparc(9S) 39

ddi_dma_lim_x86(9S) 41

ddi_dma_req(9S) 43

ddi-forceattach(9P) 46

ddi_idevice_cookie(9S) 47

ddi_mapdev_ctl(9S) 48

devmap_callback_ctl(9S) 49

dev_ops(9S) 51

fmodsw(9S) 52

free_rtn(9S) 53
 gld_mac_info(9S) 54
 gld_stats(9S) 57
 inquiry-device-type(9P) 59
 iocblk(9S) 60
 iovec(9S) 61
 kstat(9S) 62
 kstat_intr(9S) 64
 kstat_io(9S) 66
 kstat_named(9S) 67
 linkblk(9S) 68
 modldrv(9S) 69
 modlinkage(9S) 70
 modlstrmod(9S) 71
 module_info(9S) 72
 msgb(9S) 73
 no-involuntary-power-cycles(9P) 74
 pm(9P) 76
 pm-components(9P) 78
 qband(9S) 81
 qinit(9S) 82
 queclass(9S) 83
 queue(9S) 84
 removable-media(9P) 85
 scsi_address(9S) 86
 scsi_arq_status(9S) 87
 scsi_asc_key_strings(9S) 88
 scsi_device(9S) 89
 scsi_extended_sense(9S) 90
 scsi_hba_tran(9S) 93
 scsi_inquiry(9S) 95
 scsi_pkt(9S) 98
 scsi_status(9S) 102
 streamtab(9S) 104
 stroptions(9S) 105
 tuple(9S) 107
 uio(9S) 110

Index 113

Contents 5

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:
[]	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(9S)

NAME	Intro – introduction to kernel data structures and properties
DESCRIPTION	<p>Section 9P describes kernel properties used by device drivers. Section 9S describes the data structures used by drivers to share information between the driver and the kernel. See <code>Intro(9E)</code> for an overview of device driver interfaces.</p> <p>In Section 9S, reference pages contain the following headings:</p> <ul style="list-style-type: none"> ■ <code>NAME</code> summarizes the purpose of the structure or property. ■ <code>SYNOPSIS</code> lists the include file that defines the structure or property. ■ <code>INTERFACE LEVEL</code> describes any architecture dependencies. ■ <code>DESCRIPTION</code> provides general information about the structure or property. ■ <code>STRUCTURE MEMBERS</code> lists all accessible structure members (for Section 9S). ■ <code>SEE ALSO</code> gives sources for further information. <p>Of the preceding headings, Section 9P reference pages contain the <code>NAME</code>, <code>DESCRIPTION</code>, and <code>SEE ALSO</code> fields.</p> <p>Every driver MUST include <code><sys/ddi.h></code> and <code><sys/sunddi.h></code>, in that order, and as final entries.</p> <p>The following table summarizes the STREAMS structures described in Section 9S.</p>

Structure	Type
copyreq	DDI/DKI
copyresp	DDI/DKI
datab	DDI/DKI
fmodsw	Solaris DDI
free_rtn	DDI/DKI
iocblk	DDI/DKI
linkblk	DDI/DKI
module_info	DDI/DKI
msgb	DDI/DKI
qband	DDI/DKI
qinit	DDI/DKI
queclass	Solaris DDI
queue	DDI/DKI
streamtab	DDI/DKI
stroptions	DDI/DKI

The following table summarizes structures that are not specific to STREAMS I/O.

Structure	Type
aio_req	Solaris DDI
buf	DDI/DKI
cb_ops	Solaris DDI
ddi_device_acc_attr	Solaris DDI
ddi_dma_attr	Solaris DDI
ddi_dma_cookie	Solaris DDI
ddi_dma_lim_sparc	Solaris SPARC DDI
ddi_dma_lim_x86	Solaris x86 DDI
ddi_dma_req	Solaris DDI
ddi_dmae_req	Solaris x86 DDI
ddi_idevice_cookie	Solaris DDI
ddi_mapdev_ctl	Solaris DDI
devmap_callback_ctl	Solaris DDI
dev_ops	Solaris DDI
iovec	DDI/DKI
kstat	Solaris DDI
kstat_intr	Solaris DDI
kstat_io	Solaris DDI
kstat_named	Solaris DDI
map	DDI/DKI
modldrv	Solaris DDI
modlinkage	Solaris DDI
modlstrmod	Solaris DDI
scsi_address	Solaris DDI
scsi_arq_status	Solaris DDI
scsi_device	Solaris DDI
scsi_extended_sense	Solaris DDI

Intro(9S)

Structure	Type
<code>scsi_hba_tran</code>	Solaris DDI
<code>scsi_inquiry</code>	Solaris DDI
<code>scsi_pkt</code>	Solaris DDI
<code>scsi_status</code>	Solaris DDI
<code>uio</code>	DDI/DKI

SEE ALSO [Intro\(9E\)](#)

NOTES Do not declare arrays of structures as the size of the structures can change between releases. Rely only on the structure members listed in this chapter and not on unlisted members or the position of a member in a structure.

Data Structures for Drivers

aio_req(9S)

NAME	aio_req – asynchronous I/O request structure
SYNOPSIS	<pre>#include <sys/uio.h> #include <sys/aio_req.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	An aio_req structure describes an asynchronous I/O request.
STRUCTURE MEMBERS	<pre>struct uio*aio_uio; /* uio structure describing the I/O request */</pre> The aio_uio member is a pointer to a uio(9S) structure, describing the I/O transfer request.
SEE ALSO	aread(9E), awrite(9E), aphysio(9F), uio(9S)

NAME	buf – block I/O data transfer structure
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL DESCRIPTION	<p>Architecture independent level 1 (DDI/DKI)</p> <p>The buf structure is the basic data structure for block I/O transfers. Each block I/O transfer has an associated buffer header. The header contains all the buffer control and status information. For drivers, the buffer header pointer is the sole argument to a block driver strategy(9E) routine. Do not depend on the size of the buf structure when writing a driver.</p> <p>A buffer header can be linked in multiple lists simultaneously. Because of this, most of the members in the buffer header cannot be changed by the driver, even when the buffer header is in one of the driver's work lists.</p> <p>Buffer headers are also used by the system for unbuffered or physical I/O for block drivers. In this case, the buffer describes a portion of user data space that is locked into memory.</p> <p>Block drivers often chain block requests so that overall throughput for the device is maximized. The av_forw and the av_back members of the buf structure can serve as link pointers for chaining block requests.</p>
STRUCTURE MEMBERS	<pre>int b_flags; /* Buffer status */ struct buf *av_forw; /* Driver work list link */ struct buf *av_back; /* Driver work list link */ size_t b_bcount; /* # of bytes to transfer */ union { caddr_t b_addr; /* Buffer's virtual address */ } b_un; daddr_t b_blkno; /* Block number on device */ diskaddr_t b_lblkno; /* Expanded block number on device */ size_t b_resid; /* # of bytes not transferred */ size_t b_bufsize; /* size of allocated buffer */ int (*b_iodone)(struct buf *); /* function called */ /* by biodone */ int b_error; /* expanded error field */ void *b_private; /* "opaque" driver private area */ dev_t b_edev; /* expanded dev field */</pre> <p>The members of the buffer header available to test or set by a driver are as follows:</p> <p>b_flags stores the buffer status and indicates to the driver whether to read or write to the device. The driver must never clear the b_flags member. If this is done, unpredictable results can occur including loss of disk sanity and the possible failure of other kernel processes.</p> <p>All b_flags bit values not otherwise specified above are reserved by the kernel and may not be used.</p> <p>Valid flags are as follows:</p>

buf(9S)

B_BUSY	Indicates the buffer is in use. The driver must not change this flag unless it allocated the buffer with <code>getrbuf(9F)</code> and no I/O operation is in progress.
B_DONE	Indicates the data transfer has completed. This flag is read-only.
B_ERROR	Indicates an I/O transfer error. It is set in conjunction with the <code>b_error</code> field. <code>bioerror(9F)</code> should be used in preference to setting the B_ERROR bit.
B_PAGEIO	Indicates the buffer is being used in a paged I/O request. See the description of the <code>b_un.b_addr</code> field for more information. This flag is read-only.
B_PHYS	indicates the buffer header is being used for physical (direct) I/O to a user data area. See the description of the <code>b_un.b_addr</code> field for more information. This flag is read-only.
B_READ	Indicates that data is to be read from the peripheral device into main memory.
B_WRITE	Indicates that the data is to be transferred from main memory to the peripheral device. B_WRITE is a pseudo flag and cannot be directly tested; it is only detected as the NOT form of B_READ.

`av_forw` and `av_back` can be used by the driver to link the buffer into driver work lists.

`b_bcount` specifies the number of bytes to be transferred in both a paged and a non-paged I/O request.

`b_un.b_addr` is the virtual address of the I/O request, unless B_PAGEIO is set. The address is a kernel virtual address, unless B_PHYS is set, in which case it is a user virtual address. If B_PAGEIO is set, `b_un.b_addr` contains kernel private data. Note that either one of B_PHYS and B_PAGEIO, or neither, can be set, but not both.

`b_blkno` identifies which logical block on the device (the device is defined by the device number) is to be accessed. The driver might have to convert this logical block number to a physical location such as a cylinder, track, and sector of a disk. This is a 32-bit value. The driver should use `b_blkno` or `b_lblkno`, but not both.

`b_lblkno` identifies which logical block on the device (the device is defined by the device number) is to be accessed. The driver might have to convert this logical block number to a physical location such as a cylinder, track, and sector of a disk. This is a 64-bit value. The driver should use `b_lblkno` or `b_blkno`, but not both.

`b_resid` should be set to the number of bytes not transferred because of an error.

`b_bufsize` contains the size of the allocated buffer.

`b_iodone` identifies a specific `biodone` routine to be called by the driver when the I/O is complete.

`b_error` can hold an error code that should be passed as a return code from the driver. `b_error` is set in conjunction with the `B_ERROR` bit set in the `b_flags` member. `bioerror(9F)` should be used in preference to setting the `b_error` field.

`b_private` is for the private use of the device driver.

`b_edev` contains the major and minor device numbers of the device accessed.

SEE ALSO `strategy(9E)`, `aphysio(9F)`, `bioclone(9F)`, `biodone(9F)`, `bioerror(9F)`, `bioinit(9F)`, `clrbuf(9F)`, `getrbuf(9F)`, `physio(9F)`, `iovec(9S)`, `uio(9S)`

Writing Device Drivers

WARNINGS Buffers are a shared resource within the kernel. Drivers should read or write only the members listed in this section. Drivers that attempt to use undocumented members of the `buf` structure risk corrupting data in the kernel or on the device.

cb_ops(9S)

NAME	cb_ops – character/block entry points structure
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	<p>cb_ops contains all entry points for drivers that support both character and block entry points. All leaf device drivers supporting direct user process access to a device should declare a cb_ops structure.</p> <p>All drivers that safely allow multiple threads of execution in the driver at the same time must set the D_MP flag in the cb_flag field.</p> <p>If the driver properly handles 64-bit offsets, it should also set the D_64BIT flag in the cb_flag field. This specifies that the driver will use the uio_loffset field of the uio(9S) structure.</p> <p>mt-streams(9F) describes other flags that can be set in the cb_flag field.</p> <p>cb_rev is the cb_ops structure revision number. This field must be set to CB_REV.</p> <p>Non-STREAMS drivers should set cb_str to NULL.</p> <p>The following DDI/DKI or DKI-only or DDI-only functions are provided in the character/block driver operations structure.</p>

block/char	Function	Description
b/c	XXopen	DDI/DKI
b/c	XXclose	DDI/DKI
b	XXstrategy	DDI/DKI
b	XXprint	DDI/DKI
b	XXdump	DDI(Sun)
c	XXread	DDI/DKI
c	XXwrite	DDI/DKI
c	XXioctl	DDI/DKI
c	XXdevmap	DDI(Sun)
c	XXmmap	DKI
c	XXsegmap	DKI
c	XXchpoll	DDI/DKI

STRUCTURE MEMBERS

block/char	Function	Description
c	XXprop_op	DDI(Sun)
c	XXaread	DDI(Sun)
c	XXawrite	DDI(Sun)
int	(*cb_open)(dev_t *devp, int flag, int otyp, cred_t *credp);	
int	(*cb_close)(dev_t dev, int flag, int otyp, cred_t *credp);	
int	(*cb_strategy)(struct buf *bp);int(*cb_print)(dev_t dev, char *str);	
int	(*cb_print)(dev_t dev, char *str);	
int	(*cb_dump)(dev_t dev, caddr_t addr, daddr_t blkno, int nblk);	
int	(*cb_read)(dev_t dev, struct uio *uiop, cred_t *credp);	
int	(*cb_write)(dev_t dev, struct uio *uiop, cred_t *credp);	
int	(*cb_ioctl)(dev_t dev, int cmd, intptr_t arg, int mode, cred_t *credp, int *rvalp);	
int	(*cb_devmap)(dev_t dev, devmap_cookie_t dhp, offset_t off, size_t len, size_t *maplen, uint_t model);	
int	(*cb_mmap)(dev_t dev, off_t off, int prot);	
int	(*cb_segmap)(dev_t dev, off_t off, struct as *asp, caddr_t *addrp, off_t len, unsigned int prot, unsigned int maxprot, unsigned int flags, cred_t *credp);	
int	(*cb_chpoll)(dev_t dev, short events, int anyyet, short *reventsp, struct pollhead **phpp);	
int	(*cb_prop_op)(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op, int mod_flags, char *name, caddr_t valuep, int *length);	
	struct streamtab *cb_str; /* streams information */	
int	cb_flag;intcb_rev;	
int	(*cb_aread)(dev_t dev, struct aio_req *aio, cred_t *credp);	
int	(*cb_awrite)(dev_t dev, struct aio_req *aio, cred_t *credp);	

SEE ALSO

aread(9E), awrite(9E), chpoll(9E), close(9E), dump(9E), ioctl(9E), mmap(9E), open(9E), print(9E), prop_op(9E), read(9E), segmap(9E), strategy(9E), write(9E), nochpoll(9F), nodev(9F), nulldev(9F), dev_ops(9S), qinit(9S)

Writing Device Drivers

STREAMS Programming Guide

copyreq(9S)

NAME	copyreq – STREAMS data structure for the M_COPYIN and the M_COPYOUT message types
SYNOPSIS	#include <sys/stream.h>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	The data structure for the M_COPYIN and the M_COPYOUT message types.
STRUCTURE MEMBERS	<pre>int cq_cmd; /* ioctl command (from ioc_cmd) */ cred_t *cq_cr; /* full credentials */ uint_t cq_id; /* ioctl id (from ioc_id) */ uint_t cq_flag; /* must be zero */ mbk_t *cq_private; /* private state information */ caddr_t cq_addr; /* address to copy data to/from */ size_t cq_size; /* number of bytes to copy */</pre>
SEE ALSO	<i>STREAMS Programming Guide</i>

NAME	copyresp – STREAMS data structure for the M_IOCTLDATA message type
SYNOPSIS	#include <sys/stream.h>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	The data structure copyresp is used with the M_IOCTLDATA message type.
STRUCTURE MEMBERS	<pre> int cp_cmd; /* ioctl command (from ioc_cmd) */ cred_t *cp_cr; /* full credentials */ uint_t cp_id; /* ioctl id (from ioc_id) */ uint_t cp_flag; /* ioctl flags */ mbk_t *cp_private; /* private state information */ caddr_t cp_rval; /* status of request: 0 -> success; /* non-zero -> failure */ </pre>
SEE ALSO	<i>STREAMS Programming Guide</i>

datab(9S)

NAME	datab – STREAMS message data structure
SYNOPSIS	<pre>#include <sys/stream.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	<p>The <code>datab</code> structure describes the data of a STREAMS message. The actual data contained in a STREAMS message is stored in a data buffer pointed to by this structure. A <code>msgb</code> (message block) structure includes a field that points to a <code>datab</code> structure.</p> <p>Because a data block can have more than one message block pointing to it at one time, the <code>db_ref</code> member keeps track of a data block's references, preventing it from being deallocated until all message blocks are finished with it.</p>
STRUCTURE MEMBERS	<pre>unsigned char *db_base; /* first byte of buffer */ unsigned char *db_lim; /* last byte (+1) of buffer */ dbref_t db_ref; /* # of message pointers to this data */ unsigned char db_type; /* message type */</pre> <p>A <code>datab</code> structure is defined as type <code>dblk_t</code>.</p>
SEE ALSO	<p><code>free_rtn(9S)</code>, <code>msgb(9S)</code></p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

NAME	ddi_device_acc_attr – data access attributes structure												
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h></pre>												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).												
DESCRIPTION	The ddi_device_acc_attr structure describes the data access characteristics and requirements of the device.												
STRUCTURE MEMBERS	<pre>ushort_t devacc_attr_version; uchar_t devacc_attr_endian_flags; uchar_t devacc_attr_dataorder;</pre> <p>The devacc_attr_version member identifies the version number of this structure. The current version number is DDI_DEVICE_ATTR_V0.</p> <p>The devacc_attr_endian_flags member describes the endian characteristics of the device. Specify one of the following values:</p> <table border="0"> <tr> <td>DDI_NEVERSWAP_ACC</td> <td>Data access with no byte swapping</td> </tr> <tr> <td>DDI_STRUCTURE_BE_ACC</td> <td>Structural data access in big-endian format</td> </tr> <tr> <td>DDI_STRUCTURE_LE_ACC</td> <td>Structural data access in little endian format</td> </tr> </table> <p>DDI_STRUCTURE_BE_ACC and DDI_STRUCTURE_LE_ACC describes the endian characteristics of the device as big-endian or little-endian, respectively. Though most of the devices will have the same endian characteristics as their buses, examples of devices that have opposite endian characteristics of the buses do exist. When DDI_STRUCTURE_BE_ACC or DDI_STRUCTURE_LE_ACC is set, byte swapping is automatically performed by the system if the host machine and the device data formats have opposite endian characteristics. The implementation can take advantage of hardware platform byte swapping capabilities.</p> <p>When you specify DDI_NEVERSWAP_ACC, byte swapping is not invoked in the data access functions.</p> <p>The devacc_attr_dataorder member describes order in which the CPU will reference data. Specify one of the following values.</p> <table border="0"> <tr> <td>DDI_STRICTORDER_ACC</td> <td>The data references must be issued by a CPU in program order. Strict ordering is the default behavior.</td> </tr> <tr> <td>DDI_UNORDERED_OK_ACC</td> <td>The CPU can re-order the data references. This includes all kinds of re-ordering. For example, a load followed by a store may be replaced by a store followed by a load.</td> </tr> <tr> <td>DDI_MERGING_OK_ACC</td> <td>The CPU can merge individual stores to consecutive locations. For example, the CPU can turn two consecutive byte stores into</td> </tr> </table>	DDI_NEVERSWAP_ACC	Data access with no byte swapping	DDI_STRUCTURE_BE_ACC	Structural data access in big-endian format	DDI_STRUCTURE_LE_ACC	Structural data access in little endian format	DDI_STRICTORDER_ACC	The data references must be issued by a CPU in program order. Strict ordering is the default behavior.	DDI_UNORDERED_OK_ACC	The CPU can re-order the data references. This includes all kinds of re-ordering. For example, a load followed by a store may be replaced by a store followed by a load.	DDI_MERGING_OK_ACC	The CPU can merge individual stores to consecutive locations. For example, the CPU can turn two consecutive byte stores into
DDI_NEVERSWAP_ACC	Data access with no byte swapping												
DDI_STRUCTURE_BE_ACC	Structural data access in big-endian format												
DDI_STRUCTURE_LE_ACC	Structural data access in little endian format												
DDI_STRICTORDER_ACC	The data references must be issued by a CPU in program order. Strict ordering is the default behavior.												
DDI_UNORDERED_OK_ACC	The CPU can re-order the data references. This includes all kinds of re-ordering. For example, a load followed by a store may be replaced by a store followed by a load.												
DDI_MERGING_OK_ACC	The CPU can merge individual stores to consecutive locations. For example, the CPU can turn two consecutive byte stores into												

ddi_device_acc_attr(9S)

one halfword store. It can also batch individual loads. For example, the CPU might turn two consecutive byte loads into one halfword load. DDI_MERGING_OK_ACC also implies re-ordering.

DDI_LOADCACHING_OK_ACC

The CPU can cache the data it fetches and reuse it until another store occurs. The default behavior is to fetch new data on every load. DDI_LOADCACHING_OK_ACC also implies merging and re-ordering.

DDI_STORECACHING_OK_ACC

The CPU can 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. DDI_STORECACHING_OK_ACC 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.

EXAMPLES

The following examples illustrate the use of device register address mapping setup functions and different data access functions.

EXAMPLE 1 Using ddi_device_acc_attr() in ddi_regs_map_setup(9F)

This example demonstrates the use of the ddi_device_acc_attr() structure in ddi_regs_map_setup(9F). It also shows the use of ddi_getw(9F) and ddi_putw(9F) functions in accessing the register contents.

```
dev_info_t *dip;
uint_t      rnumber;
ushort_t    *dev_addr;
offset_t    offset;
offset_t    len;
ushort_t    dev_command;
ddi_device_acc_attr_t dev_attr;
ddi_acc_handle_t handle;

. . .

/*
 * setup the device attribute structure for little endian,
 * strict ordering and 16-bit word access.
 */
dev_attr.devacc_attr_version = DDI_DEVICE_ATTR_V0;
dev_attr.devacc_attr_endian_flags = DDI_STRUCTURE_LE_ACC;
dev_attr.devacc_attr_dataorder = DDI_STRICTORDER_ACC;

/*
 * set up the device registers address mapping
```

EXAMPLE 1 Using `ddi_device_acc_attr()` in `ddi_regs_map_setup(9F)`
(Continued)

```

*/
ddi_regs_map_setup(dip, rnumber, (caddr_t *)&dev_addr, offset, len,
                  &dev_attr, &handle);

/* read a 16-bit word command register from the device      */
dev_command = ddi_getw(handle, dev_addr);

dev_command |= DEV_INTR_ENABLE;
/* store a new value back to the device command register    */
ddi_putw(handle, dev_addr, dev_command);

```

EXAMPLE 2 Accessing a Device with Different Apertures

The following example illustrates the steps used to access a device with different apertures. Several apertures are assumed to be grouped under one single "reg" entry. For example, the sample device has four different apertures, each 32 Kbyte in size. The apertures represent YUV little-endian, YUV big-endian, RGB little-endian, and RGB big-endian. This sample device uses entry 1 of the "reg" property list for this purpose. The size of the address space is 128 Kbyte with each 32 Kbyte range as a separate aperture. In the register mapping setup function, the sample driver uses the *offset* and *len* parameters to specify one of the apertures.

```

ulong_t      *dev_addr;
ddi_device_acc_attr_t dev_attr;
ddi_acc_handle_t handle;
uchar_t buf[256];

. . .

/*
 * setup the device attribute structure for never swap,
 * unordered and 32-bit word access.
 */
dev_attr.devacc_attr_version = DDI_DEVICE_ATTR_V0;
dev_attr.devacc_attr_endian_flags = DDI_NEVERSWAP_ACC;
dev_attr.devacc_attr_dataorder = DDI_UNORDERED_OK_ACC;

/*
 * map in the RGB big-endian aperture
 * while running in a big endian machine
 * - offset 96K and len 32K
 */
ddi_regs_map_setup(dip, 1, (caddr_t *)&dev_addr, 96*1024, 32*1024,
                  &dev_attr, &handle);

/*
 * Write to the screen buffer
 * first 1K bytes words, each size 4 bytes
 */
ddi_rep_putl(handle, buf, dev_addr, 256, DDI_DEV_AUTOINCR);

```

ddi_device_acc_attr(9S)

EXAMPLE 2 Accessing a Device with Different Apertures (Continued)

EXAMPLE 3 Functions That Call Out the Data Word Size

The following example illustrates the use of the functions that explicitly call out the data word size to override the data size in the device attribute structure.

```
struct device_blk {
    ushort_t    d_command;    /* command register */
    ushort_t    d_status;     /* status register */
    ulong       d_data;       /* data register */
} *dev_blkp;
dev_info_t *dip;
caddr_t dev_addr;
ddi_device_acc_attr_t dev_attr;
ddi_acc_handle_t handle;
uchar_t buf[256];

. . .

/*
 * setup the device attribute structure for never swap,
 * strict ordering and 32-bit word access.
 */
dev_attr.devacc_attr_version = DDI_DEVICE_ATTR_V0;
dev_attr.devacc_attr_endian_flags = DDI_NEVERSWAP_ACC;
dev_attr.devacc_attr_dataorder = DDI_STRICTORDER_ACC;

ddi_regs_map_setup(dip, 1, (caddr_t *)&dev_blkp, 0, 0,
    &dev_attr, &handle);

/* write command to the 16-bit command register */
ddi_putw(handle, &dev_blkp->d_command, START_XFER);

/* Read the 16-bit status register */
status = ddi_getw(handle, &dev_blkp->d_status);

if (status & DATA_READY)
    /* Read 1K bytes off the 32-bit data register */
    ddi_rep_getl(handle, buf, &dev_blkp->d_data,
        256, DDI_DEV_NO_AUTOINCR);
```

SEE ALSO ddi_getw(9F), ddi_putw(9F), ddi_regs_map_setup(9F)

Writing Device Drivers

NAME	ddi_dma_attr – DMA attributes structure
SYNOPSIS	<pre>#include <sys/ddidmareq.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	A <code>ddi_dma_attr_t</code> structure describes device- and DMA engine-specific attributes necessary to allocate DMA resources for a device. The driver might have to extend the attributes with bus-specific information, depending on the bus to which the device is connected.
STRUCTURE MEMBERS	<pre>uint_t dma_attr_version; /* version number */ uint64_t dma_attr_addr_lo; /* low DMA address range */ uint64_t dma_attr_addr_hi; /* high DMA address range */ uint64_t dma_attr_count_max; /* DMA counter register */ uint64_t dma_attr_align; /* DMA address alignment */ uint_t dma_attr_burstsizes; /* DMA burstsizes */ uint32_t dma_attr_minxfer; /* min effective DMA size */ uint64_t dma_attr_maxxfer; /* max DMA xfer size */ uint64_t dma_attr_seg; /* segment boundary */ int dma_attr_sgllen; /* s/g list length */ uint32_t dma_attr_granular; /* granularity of device */ uint_t dma_attr_flags; /* DMA transfer flags */</pre> <p><code>dma_attr_version</code> stores the version number of this DMA attribute structure. It should be set to <code>DMA_ATTR_V0</code>.</p> <p>The <code>dma_attr_addr_lo</code> and <code>dma_attr_addr_hi</code> fields specify the address range the device's DMA engine can access. The <code>dma_attr_addr_lo</code> field describes the inclusive lower 64-bit boundary. The <code>dma_attr_addr_hi</code> describes the inclusive upper 64-bit boundary. The system ensures that allocated DMA resources are within the range specified. See <code>ddi_dma_cookie(9S)</code>.</p> <p>The <code>dma_attr_count_max</code> describes an inclusive upper bound for the device's DMA counter register. For example, <code>0xFFFFFFFF</code> would describe a DMA engine with a 24-bit counter register. DMA resource allocation functions have to break up a DMA object into multiple DMA cookies if the size of the object exceeds the size of the DMA counter register.</p> <p>The <code>dma_attr_align</code> specifies alignment requirements for allocated DMA resources. This field can be used to force more restrictive alignment than imposed by <code>dma_attr_burstsizes</code> or <code>dma_attr_minxfer</code>, such as alignment at a page boundary. Most drivers set this field to 1, indicating byte alignment.</p> <p>Note that <code>dma_attr_align</code> only specifies alignment requirements for allocated DMA resources. The buffer passed to <code>ddi_dma_addr_bind_handle(9F)</code> or <code>ddi_dma_buf_bind_handle(9F)</code> must have an equally restrictive alignment (see <code>ddi_dma_mem_alloc(9F)</code>).</p>

ddi_dma_attr(9S)

The `dma_attr_burstsizes` field describes the possible burst sizes the device's DMA engine can accept. The format of the data sizes is binary encoded in terms of powers of two. When DMA resources are allocated, the system can modify the `burstsizes` value to reflect the system limits. The driver must use the allowable `burstsizes` to program the DMA engine. See `ddi_dma_burstsizes(9F)`.

The `dma_attr_minxfer` field describes the minimum effective DMA access size in units of bytes. DMA resources can be modified, depending on the presence and use of I/O caches and write buffers between the DMA engine and the memory object. This field is used to determine alignment and padding requirements for `ddi_dma_mem_alloc(9F)`.

The `dma_attr_maxxfer` field describes the maximum effective DMA access size in units of bytes.

The `dma_attr_seg` field specifies segment boundary restrictions for allocated DMA resources. The system allocates DMA resources for the device so that the object does not span the segment boundary specified by `dma_attr_seg`. For example, a value of `0xFFFF` means DMA resources must not cross a 64 Kbyte boundary. DMA resource allocation functions might have to break up a DMA object into multiple DMA cookies to enforce segment boundary restrictions. In this case, the transfer must be performed using scatter-gather I/O or multiple DMA windows.

The `dma_attr_sgllen` field describes the length of the device's DMA scatter/gather list. Possible values are as follows:

- < 0 Device DMA engine is not constrained by the size, for example, with DMA chaining.
- = 0 Reserved.
- = 1 Device DMA engine does not support scatter/gather such as third party DMA.
- > 1 Device DMA engine uses scatter/gather. `dma_attr_sgllen` is the maximum number of entries in the list.

The `dma_attr_granular` field describes the granularity of the device transfer size, in units of bytes. When the system allocates DMA resources, a single segment's size is a multiple of the device granularity. Or if `dma_attr_sgllen` is larger than 1 within a window, the sum of the sizes for a subgroup of segments is a multiple of the device granularity.

Note that all driver requests for DMA resources must be a multiple of the granularity of the device transfer size.

The `dma_attr_flags` field can be set to:

`DDI_DMA_FORCE_PHYSICAL`

Some platforms, such as SPARC systems, support what is called Direct Virtual Memory Access (DVMA). On these platforms, the device is provided with a virtual

address by the system in order to perform the transfer. In this case, the underlying platform provides an *IOMMU*, which translates accesses to these virtual addresses into the proper physical addresses. Some of these platforms also support DMA.

DDI_DMA_FORCE_PHYSICAL indicates that the system should return physical rather than virtual I/O addresses if the system supports both. If the system does not support physical DMA, the return value from `ddi_dma_alloc_handle(9F)` will be `DDI_DMA_BADATTR`. In this case, the driver has to clear `DDI_DMA_FORCE_PHYSICAL` and retry the operation.

EXAMPLES **EXAMPLE 1** Initializing the `ddi_dma_attr_t` Structure

Assume a device has the following DMA characteristics:

- Full 32-bit range addressable
- 24-bit DMA counter register
- Byte alignment
- 4- and 8-byte burst sizes support
- Minimum effective transfer size of 1 bytes
- 64 Mbyte maximum transfer size limit
- Maximum segment size of 32 Kbyte
- 17 scatter/gather list elements
- 512-byte device transfer size granularity

The corresponding `ddi_dma_attr_t` structure is initialized as follows:

```
static ddi_dma_attr_t dma_attrs = {
    DMA_ATTR_V0           /* version number */
    (uint64_t)0x0,        /* low address */
    (uint64_t)0xffffffff, /* high address */
    (uint64_t)0xffffffff, /* DMA counter max */
    (uint64_t)0x1        /* alignment */
    0x0c,                /* burst sizes */
    0x1,                 /* minimum transfer size */
    (uint64_t)0x3ffffff, /* maximum transfer size */
    (uint64_t)0x7fff,    /* maximum segment size */
    17,                  /* scatter/gather list lgth */
    512,                 /* granularity */
    0                     /* DMA flags */
};
```

SEE ALSO `ddi_dma_addr_bind_handle(9F)`, `ddi_dma_alloc_handle(9F)`,
`ddi_dma_buf_bind_handle(9F)`, `ddi_dma_burstsizes(9F)`,
`ddi_dma_mem_alloc(9F)`, `ddi_dma_nextcookie(9F)`, `ddi_dma_cookie(9S)`

Writing Device Drivers

ddi_dma_cookie(9S)

NAME	ddi_dma_cookie – DMA address cookie
SYNOPSIS	#include <sys/sunddi.h>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	The ddi_dma_cookie_t structure contains DMA address information required to program a DMA engine. The structure is filled in by a call to ddi_dma_getwin(9F), ddi_dma_addr_bind_handle(9F), or ddi_dma_buf_bind_handle(9F), to get device-specific DMA transfer information for a DMA request or a DMA window.
STRUCTURE MEMBERS	<pre>typedef struct { union { uint64_t _dmac_ll; /* 64 bit DMA address */ uint32_t _dmac_la[2]; /* 2 x 32 bit address */ } _dmu; size_t dmac_size; /* DMA cookie size */ uint_t dmac_type; /* bus specific type bits */ } ddi_dma_cookie_t;</pre> <p>You can access the DMA address through the #defines: dmac_address for 32-bit addresses and dmac_laddress for 64-bit addresses. These macros are defined as follows:</p> <pre>#define dmac_laddress _dmu._dmac_ll #ifdef _LONG_LONG_HTO #define dmac_notused _dmu._dmac_la[0] #define dmac_address _dmu._dmac_la[1] #else #define dmac_address _dmu._dmac_la[0] #define dmac_notused _dmu._dmac_la[1] #endif</pre> <p>dmac_laddress specifies a 64-bit I/O address appropriate for programming the device's DMA engine. If a device has a 64-bit DMA address register a driver should use this field to program the DMA engine. dmac_address specifies a 32-bit I/O address. It should be used for devices that have a 32-bit DMA address register. The I/O address range that the device can address and other DMA attributes have to be specified in a ddi_dma_attr(9S) structure.</p> <p>dmac_size describes the length of the transfer in bytes.</p> <p>dmac_type contains bus-specific type bits, if appropriate. For example, a device on a PCI bus has PCI address modifier bits placed here.</p>
SEE ALSO	pci(4), sbus(4), sysbus(4), ddi_dma_addr_bind_handle(9F), ddi_dma_buf_bind_handle(9F), ddi_dma_getwin(9F), ddi_dma_nextcookie(9F), ddi_dma_attr(9S)
	<i>Writing Device Drivers</i>

NAME	ddi_dmae_req – DMA engine request structure
SYNOPSIS	<pre>#include <sys/dma_engine.h></pre>
INTERFACE LEVEL	Solaris x86 DDI specific (Solaris x86 DDI).
DESCRIPTION	A device driver uses the <code>ddi_dmae_req</code> structure to describe the parameters for a DMA channel. This structure contains all the information necessary to set up the channel, except for the DMA memory address and transfer count. The defaults, as specified below, support most standard devices. Other modes might be desirable for some devices, or to increase performance. The DMA engine request structure is passed to <code>ddi_dmae_prog(9F)</code> .
STRUCTURE MEMBERS	<p>The <code>ddi_dmae_req</code> structure contains several members, each of which controls some aspect of DMA engine operation. The structure members associated with supported DMA engine options are described here.</p> <pre>uchar_t der_command; /* Read / Write * /uchar_t der_bufprocess; /* Standard / Chain */ uchar_t der_path; /* 8 / 16 / 32 */ uchar_t der_cycles; /* Compat / Type A / Type B / Burst */ uchar_t der_trans; /* Single / Demand / Block */ ddi_dma_cookie_t (*proc)(); /* address of nextcookie routine */ void* procparms; /* parameter for nextcookie call */</pre> <p>der_command Specifies what DMA operation is to be performed. The value <code>DMAE_CMD_WRITE</code> signifies that data is to be transferred from memory to the I/O device. The value <code>DMAE_CMD_READ</code> signifies that data is to be transferred from the I/O device to memory. This field must be set by the driver before calling <code>ddi_dmae_prog()</code>.</p> <p>der_bufprocess On some bus types, a driver can set <code>der_bufprocess</code> to the value <code>DMAE_BUF_CHAIN</code> to specify that multiple DMA cookies will be given to the DMA engine for a single I/O transfer. This action causes a scatter/gather operation. In this mode of operation, the driver calls <code>ddi_dmae_prog()</code> to give the DMA engine the DMA engine request structure and a pointer to the first cookie. The <code>proc</code> structure member must be set to the address of a driver <code>nextcookie</code> routine. This routine takes one argument, specified by the <code>procparms</code> structure member, and returns a pointer to a structure of type <code>ddi_dma_cookie_t</code> that specifies the next cookie for the I/O transfer. When the DMA engine is ready to receive an additional cookie, the bus nexus driver controlling that DMA engine calls the routine specified by the <code>proc</code> structure member to obtain the next cookie from the driver. The driver's <code>nextcookie</code> routine must then return the address of the next cookie (in static storage) to the bus nexus routine that called it. If there are no more segments in the current DMA window, then <code>(*proc)()</code> must return the <code>NULL</code> pointer.</p> <p>A driver can specify the <code>DMAE_BUF_CHAIN</code> flag only if the particular bus architecture supports the use of multiple DMA cookies in a single I/O transfer. A bus DMA engine can support this feature either with a fixed-length scatter/gather list, or by an interrupt chaining feature such as the one implemented in the EISA</p>

ddi_dmae_req(9S)

architecture. A driver must determine whether its parent bus nexus supports this feature by examining the scatter/gather list size returned in the `dlim_sgllen` member of the DMA limit structure returned by the driver's call to `ddi_dmae_getlim()`. (See `ddi_dma_lim_x86(9S)`.) If the size of the scatter/gather list is 1, then no chaining is available. The driver must not specify the `DMAE_BUF_CHAIN` flag in the `ddi_dmae_req` structure it passes to `ddi_dmae_prog()`, and the driver need not provide a `nextcookie` routine.

If the size of the scatter/gather list is greater than 1, then DMA chaining is available, and the driver has two options. Under the first option, the driver chooses not to use the chaining feature. In this case (a) the driver must set the size of the scatter/gather list to 1 before passing it to the DMA setup routine, and (b) the driver must not set the `DMAE_BUF_CHAIN` flag.

Under the second option, the driver chooses to use the chaining feature, in which case, (a) it should leave the size of the scatter/gather list alone, and (b) it must set the `DMAE_BUF_CHAIN` flag in the `ddi_dmae_req` structure. Before calling `ddi_dmae_prog()`, the driver must *prefetch* cookies by repeatedly calling `ddi_dma_nextseg(9F)` and `ddi_dma_segtocookie(9F)` until either (1) the end of the DMA window is reached (`ddi_dma_nextseg(9F)` returns `NULL`), or (2) the size of the scatter/gather list is reached, whichever occurs first. These cookies must be saved by the driver until they are requested by the nexus driver calling the driver's `nextcookie` routine. The driver's `nextcookie` routine must return the prefetched cookies in order, one cookie for each call to the `nextcookie` routine, until the list of prefetched cookies is exhausted. After the end of the list of cookies is reached, the `nextcookie` routine must return the `NULL` pointer.

The size of the scatter/gather list determines how many discontinuous segments of physical memory can participate in a single DMA transfer. ISA bus DMA engines have no scatter/gather capability, so their scatter/gather list sizes are 1. EISA bus DMA engines have a DMA chaining interrupt facility that allows very large scatter/gather operations. Other finite scatter/gather list sizes would also be possible. For performance reasons, drivers should use the chaining capability if it is available on their parent bus.

As described above, a driver making use of DMA chaining must prefetch DMA cookies before calling `ddi_dmae_prog()`. The reasons for this are:

- First, the driver must have some way to know the total I/O count with which to program the I/O device. This I/O count must match the total size of all the DMA segments that will be chained together into one DMA operation. Depending on the size of the scatter/gather list and the memory position and alignment of the DMA object, all or just part of the current DMA window might be able to participate in a single I/O operation. The driver must compute the I/O count by adding up the sizes of the prefetched DMA cookies. The number of cookies whose sizes are to be summed is the lesser of (a) the size of the scatter/gather list, or (b) the number of segments remaining in the window.

- Second, on some bus architectures, the driver's `nextcookie` routine can be called from a high-level interrupt routine. If the cookies were not prefetched, the `nextcookie` routine would have to call `ddi_dma_nextseg()` and `ddi_dma_segtocookie()` from a high-level interrupt routine, which is not recommended.

When breaking a DMA window into segments, the system arranges for the end of every segment whose number is an integral multiple of the scatter/gather list size to fall on a device-granularity boundary, as specified in the `dlim_granular` field in the `ddi_dma_lim_x86(9S)` structure.

If the scatter/gather list size is 1 (either because no chaining is available or because the driver does not want to use the chaining feature), then the total I/O count for a single DMA operation is the size of DMA segment denoted by the single DMA cookie that is passed in the call to `ddi_dmae_prog()`. In this case, the system arranges for each DMA segment to be a multiple of the device-granularity size.

`der_path`

Specifies the DMA transfer size. The default of zero (`DMAE_PATH_DEF`) specifies ISA compatibility mode. In that mode, channels 0, 1, 2, and 3 are programmed in 8-bit mode (`DMAE_PATH_8`), and channels 5, 6, and 7 are programmed in 16-bit, count-by-word mode (`DMAE_PATH_16`). On the EISA bus, other sizes can be specified: `DMAE_PATH_32` specifies 32-bit mode, and `DMAE_PATH_16B` specifies a 16-bit, count-by-byte mode.

`der_cycles`

Specifies the timing mode to be used during DMA data transfers. The default of zero (`DMAE_CYCLES_1`) specifies ISA compatible timing. Drivers using this mode must also specify `DMAE_TRANS_SINGL` in the `der_trans` structure member. On EISA buses, these other timing modes are available:

<code>DMAE_CYCLES_2</code>	Specifies type "A" timing;
<code>DMAE_CYCLES_3</code>	Specifies type "B" timing;
<code>DMAE_CYCLES_4</code>	Specifies "Burst" timing.

`der_trans`

Specifies the bus transfer mode that the DMA engine should expect from the device. The default value of zero (`DMAE_TRANS_SINGL`) specifies that the device performs one transfer for each bus arbitration cycle. Devices that use ISA compatible timing (specified by a value of zero, which is the default, in the `der_cycles` structure member) should use the `DMAE_TRANS_SINGL` mode. On EISA buses, a `der_trans` value of `DMAE_TRANS_BLK` specifies that the device perform a block of transfers for each arbitration cycle. A value of `DMAE_TRANS_DMND` specifies that the device perform the Demand Transfer Mode protocol.

ddi_dmae_req(9S)

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_segtocookie(9F)`, `ddi_dmae(9F)`, `ddi_dma_lim_x86(9S)`, `ddi_dma_req(9S)`

NAME	ddi_dma_lim_sparc, ddi_dma_lim – SPARC DMA limits structure
SYNOPSIS	<pre>#include <sys/ddidmareq.h></pre>
INTERFACE LEVEL	Solaris SPARC DDI specific (Solaris SPARC DDI).
DESCRIPTION	<p>This page describes the SPARC version of the <code>ddi_dma_lim</code> structure. See <code>ddi_dma_lim_x86(9S)</code> for a description of the x86 version of this structure.</p> <p>A <code>ddi_dma_lim</code> structure describes in a generic fashion the possible limitations of a device's DMA engine. This information is used by the system when it attempts to set up DMA resources for a device.</p>
STRUCTURE MEMBERS	<pre>uint_t dlim_addr_lo; /* low range of 32 bit addressing capability */ uint_t dlim_addr_hi; /* inclusive upper bound of addressing */ /* capability */ uint_t dlim_cntr_max; /* inclusive upper bound of dma engine's */ /* address limit */ uint_t dlim_burstsizes; /* binary encoded dma burst sizes */ uint_t dlim_minxfer; /* minimum effective dma transfer size */ uint_t dlim_dmaspeed; /* average dma data rate (kb/s) */</pre> <p>The <code>dlim_addr_lo</code> and <code>dlim_addr_hi</code> fields specify the address range the device's DMA engine can access. The <code>dlim_addr_lo</code> field describes the lower 32-bit boundary of the device's DMA engine, the <code>dlim_addr_hi</code> describes the inclusive upper 32-bit boundary. The system allocates DMA resources in a way that the address for programming the device's DMA engine (see <code>ddi_dma_cookie(9S)</code> or <code>ddi_dma_htoc(9F)</code>) is within this range. For example, if your device can access the whole 32-bit address range, you may use <code>[0,0xFFFFFFFF]</code>. If your device has just a 16-bit address register but will access the top of the 32-bit address range, then <code>[0xFFFF0000,0xFFFFFFFF]</code> is the right limit.</p> <p>The <code>dlim_cntr_max</code> field describes an inclusive upper bound for the device's DMA engine address register. This handles a fairly common case where a portion of the address register is only a latch rather than a full register. For example, the upper 8 bits of a 32-bit address register can be a latch. This splits the address register into a portion that acts as a true address register (24 bits) for a 16 Mbyte segment and a latch (8 bits) to hold a segment number. To describe these limits, specify <code>0xFFFFFFFF</code> in the <code>dlim_cntr_max</code> structure.</p> <p>The <code>dlim_burstsizes</code> field describes the possible burst sizes the device's DMA engine can accept. At the time of a DMA resource request, this element defines the possible DMA burst cycle sizes that the requester's DMA engine can handle. The format of the data is binary encoding of burst sizes assumed to be powers of two. That is, if a DMA engine is capable of doing 1-, 2-, 4-, and 16-byte transfers, the encoding is <code>0x17</code>. If the device is an SBus device and can take advantage of a 64-bit SBus, the lower 16 bits are used to specify the burst size for 32-bit transfers and the upper 16 bits are used to specify the burst size for 64-bit transfers. As the resource request is handled by the system, the <code>burstsizes</code> value can be modified. Prior to enabling DMA for the specific device, the driver that owns the DMA engine should check (using <code>ddi_dma_burstsizes(9F)</code>) what the allowed burst sizes have become and program the DMA engine appropriately.</p>

ddi_dma_lim_sparc(9S)

The `dlim_minxfer` field describes the minimum effective DMA transfer size (in units of bytes). It must be a power of two. This value specifies the minimum effective granularity of the DMA engine. It is distinct from `dlim_burstsizes` in that it describes the minimum amount of access a DMA transfer will effect.

`dlim_burstsizes` describes in what electrical fashion the DMA engine might perform its accesses, while `dlim_minxfer` describes the minimum amount of memory that can be touched by the DMA transfer. As a resource request is handled by the system, the `dlim_minxfer` value can be modified contingent upon the presence (and use) of I/O caches and DMA write buffers in between the DMA engine and the object that DMA is being performed on. After DMA resources have been allocated, the resultant minimum transfer value can be gotten using `ddi_dma_devalign(9F)`.

The field `dlim_dmaspeed` is the expected average data rate for the DMA engine (in units of kilobytes per second). Note that this should not be the maximum, or peak, burst data rate, but a reasonable guess as to the average throughput. This field is entirely optional and can be left as zero. Its intended use is to provide some hints about how much of the DMA resource this device might need.

SEE ALSO

`ddi_dma_addr_setup(9F)`, `ddi_dma_buf_setup(9F)`, `ddi_dma_burstsizes(9F)`, `ddi_dma_devalign(9F)`, `ddi_dma_htoc(9F)`, `ddi_dma_setup(9F)`, `ddi_dma_cookie(9S)`, `ddi_dma_lim_x86(9S)`, `ddi_dma_req(9S)`

NAME	ddi_dma_lim_x86 – x86 DMA limits structure
SYNOPSIS	<pre>#include <sys/ddidmareq.h></pre>
INTERFACE LEVEL	Solaris x86 DDI specific (Solaris x86 DDI)
DESCRIPTION	<p>A <code>ddi_dma_lim</code> structure describes in a generic fashion the possible limitations of a device or its DMA engine. This information is used by the system when it attempts to set up DMA resources for a device. When the system is requested to perform a DMA transfer to or from an object, the request is broken up, if necessary, into multiple sub-requests. Each sub-request conforms to the limitations expressed in the <code>ddi_dma_lim</code> structure.</p> <p>This structure should be filled in by calling the routine <code>ddi_dmae_getlim(9F)</code>. This routine sets the values of the structure members appropriately based on the characteristics of the DMA engine on the driver's parent bus. If the driver has additional limitations, it can <i>further restrict</i> some of the values in the structure members. A driver should <i>not relax</i> any restrictions imposed by <code>ddi_dmae_getlim()</code>.</p>
STRUCTURE MEMBERS	<pre>uint_t dlim_addr_lo; /* low range of 32 bit addressing capability */ uint_t dlim_addr_hi; /* inclusive upper bound of addressing capability */ uint_t dlim_minxfer; /* minimum effective dma transfer size */ uint_t dlim_version; /* version number of this structure */ uint_t dlim_adreg_max; /* inclusive upper bound of /* incrementing addr reg */ uint_t dlim_ctreg_max; /* maximum transfer count minus one */ uint_t dlim_granular; /* granularity (and min size) of transfer count */ short dlim_sgllen; /* length of DMA scatter/gather list */ uint_t dlim_reqsize; /* maximum transfer size in bytes of a single I/O */</pre> <p>The <code>dlim_addr_lo</code> and <code>dlim_addr_hi</code> fields specify the address range that the device's DMA engine can access. The <code>dlim_addr_lo</code> field describes the lower 32-bit boundary of the device's DMA engine. The <code>dlim_addr_hi</code> member describes the inclusive, upper 32-bit boundary. The system allocates DMA resources in a way that the address for programming the device's DMA engine will be within this range. For example, if your device can access the whole 32-bit address range, you can use <code>[0, 0xFFFFFFFF]</code>. See <code>ddi_dma_cookie(9S)</code> or <code>ddi_dma_segtocookie(9F)</code>.</p> <p>The <code>dlim_minxfer</code> field describes the minimum effective DMA transfer size (in units of bytes), which must be a power of two. This value specifies the minimum effective granularity of the DMA engine and describes the minimum amount of memory that can be touched by the DMA transfer. As a resource request is handled by the system, the <code>dlim_minxfer</code> value can be modified. This modification is contingent upon the presence (and use) of I/O caches and DMA write buffers between the DMA engine and the object that DMA is being performed on. After DMA resources have been allocated, you can retrieve the resultant minimum transfer value using <code>ddi_dma_devalign(9F)</code>.</p> <p>The <code>dlim_version</code> field specifies the version number of this structure. Set this field to <code>DMALIM_VER0</code>.</p>

ddi_dma_lim_x86(9S)

The `dlim_adreg_max` field describes an inclusive upper bound for the device's DMA engine address register. This bound handles a fairly common case where a portion of the address register is simply a latch rather than a full register. For example, the upper 16 bits of a 32-bit address register might be a latch. This splits the address register into a portion that acts as a true address register (lower 16 bits) for a 64-kilobyte segment and a latch (upper 16 bits) to hold a segment number. To describe these limits, you specify `0xFFFF` in the `dlim_adreg_max` structure member.

The `dlim_ctreg_max` field specifies the maximum transfer count that the DMA engine can handle in one segment or cookie. The limit is expressed as the maximum count minus one. This transfer count limitation is a per-segment limitation. Because the limitation is used as a bit mask, it must be one less than a power of two.

The `dlim_granular` field describes the granularity of the device's DMA transfer ability, in units of bytes. This value is used to specify, for example, the sector size of a mass storage device. DMA requests are broken into multiples of this value. If there is no scatter/gather capability, then the size of each DMA transfer will be a multiple of this value. If there is scatter/gather capability, then a single segment cannot be smaller than the minimum transfer value, but can be less than the granularity. However, the total transfer length of the scatter/gather list is a multiple of the granularity value.

The `dlim_sgllen` field specifies the maximum number of entries in the scatter/gather list. This value is the number of segments or cookies that the DMA engine can consume in one I/O request to the device. If the DMA engine has no scatter/gather list, set this field to one.

The `dlim_reqsize` field describes the maximum number of bytes that the DMA engine can transmit or receive in one I/O command. This limitation is only significant if it is less than $(dlim_ctreg_max + 1) * dlim_sgllen$. If the DMA engine has no particular limitation, set this field to `0xFFFFFFFF`.

SEE ALSO

`ddi_dmae(9F)`, `ddi_dma_addr_setup(9F)`, `ddi_dma_buf_setup(9F)`,
`ddi_dma_devalign(9F)`, `ddi_dma_segtocookie(9F)`, `ddi_dma_setup(9F)`,
`ddi_dma_cookie(9S)` `ddi_dma_lim_sparc(9S)`, `ddi_dma_req(9S)`

NAME ddi_dma_req – DMA Request structure

SYNOPSIS #include <sys/ddidmareq.h>

INTERFACE LEVEL Solaris DDI specific (Solaris DDI).

DESCRIPTION A `ddi_dma_req` structure describes a request for DMA resources. A driver can use it to describe forms of allocations and ways to allocate DMA resources for a DMA request.

STRUCTURE MEMBERS

```

ddi_dma_lim_t  *dmar_limits;          /* Caller's dma engine's */
                                           /* constraints */
uint_t         dmar_flags;           /* Contains information for */
                                           /* mapping routines */
int            (*dmar_fp)(caddr_t);  /* Callback function */
caddr_t        dmar_arg;             /* Callback function's argument */
ddi_dma_obj_t  dmar_object;         /* Description of the object */
                                           /* to be mapped */

```

For the definition of the DMA limits structure, which `dmar_limits` points to, see `ddi_dma_lim_sparc(9S)` or `ddi_dma_lim_x86(9S)`.

Valid values for `dmar_flags` are:

```

DDI_DMA_WRITE      /* Direction memory --> IO */
DDI_DMA_READ       /* Direction IO --> memory */
DDI_DMA_RDWR       /* Both read and write */
DDI_DMA_REDZONE    /* Establish an MMU redzone at end of mapping */
DDI_DMA_PARTIAL    /* Partial mapping is allowed */
DDI_DMA_CONSISTENT /* Byte consistent access wanted */
DDI_DMA_SBUS_64BIT /* Use 64 bit capability on SBus */

```

`DDI_DMA_WRITE`, `DDI_DMA_READ`, and `DDI_DMA_RDWR` describe the intended direction of the DMA transfer. Some implementations might explicitly disallow `DDI_DMA_RDWR`.

`DDI_DMA_REDZONE` asks the system 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 might not have the hardware ability to support it.

`DDI_DMA_PARTIAL` lets the system know that the caller can accept partial mapping. That is, if the size of the object exceeds the resources available, the system allocates only a portion of the object and returns status indicating this partial allocation. At a later point, the caller can use `ddi_dma_curwin(9F)` and `ddi_dma_movwin(9F)` to change the valid portion of the object that has resources allocated.

`DDI_DMA_CONSISTENT` gives a hint to the system that the object should be mapped for *byte consistent* access. Normal data transfers usually use a *streaming* mode of operation. They start at a specific point, transfer a fairly large amount of data sequentially, and then stop, usually on an aligned boundary. Control mode data transfers for memory-resident device control blocks (for example, Ethernet message descriptors) do not access memory in such a sequential fashion. Instead, they tend to modify a few words or bytes, move around and maybe modify a few more.

ddi_dma_req(9S)

Many machine implementations make this non-sequential memory access difficult to control in a generic and seamless fashion. Therefore, explicit synchronization steps using `ddi_dma_sync(9F)` or `ddi_dma_free(9F)` are required to make the view of a memory object shared between a CPU and a DMA device consistent. However, proper use of the `DDI_DMA_CONSISTENT` flag can create a condition in which a system will pick resources in a way that makes these synchronization steps as efficient as possible.

`DDI_DMA_SBUS_64BIT` tells the system that the device can perform 64-bit transfers on a 64-bit SBus. If the SBus does not support 64-bit data transfers, data will be transferred in 32-bit mode.

The callback function specified by the member `dmar_fp` indicates how a caller to one of the DMA resource allocation functions wants to deal with the possibility of resources not being available. (See `ddi_dma_setup(9F)`.) If `dmar_fp` is set to `DDI_DMA_DONTWAIT`, then the caller does not care if the allocation fails, and can deal with an allocation failure appropriately. Setting `dmar_fp` to `DDI_DMA_SLEEP` indicates the caller wants 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 later, when resources become available. When the specified function is called, it is passed the value set in the structure member `dmar_arg`. The specified callback function *must* return either:

- 0 Indicating that it attempted to allocate a DMA resource but failed to do so, again, in which case the callback function will be put back on a list to be called again later.
- 1 Indicating either success at allocating DMA resources or that it no longer wants to retry.

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

It is possible that a call to `ddi_dma_free(9F)`, which frees DMA resources, might cause a callback function to be called and, unless some care is taken, an undesired recursion can occur. This can cause an undesired recursive `mutex_enter(9F)`, which makes the system panic.

dmar_object Structure

The `dmar_object` member of the `ddi_dma_req` structure is itself a complex and extensible structure:

```
uint_t            dmao_size;        /* size, in bytes, of the object */
ddi_dma_atyp_t    dmao_type;       /* type of object */
ddi_dma_aobj_t    dmao_obj;       /* the object described */
```

The `dmao_size` element is the size, in bytes, of the object resources allocated for DMA.

The `dmao_type` element selects the kind of object described by `dmao_obj`. It can be set to `DMA_OTYP_VADDR`, indicating virtual addresses.

The last element, `dmao_obj`, consists of the virtual address type:

```
struct v_address virt_obj;
```

It is specified as:

```
struct v_address {
    caddr_t    v_addr;    /* base virtual address */
    struct as  *v_as;     /* pointer to address space */
    void      *v_priv;    /* priv data for shadow I/O */
};
```

SEE ALSO `ddi_dma_addr_setup(9F)`, `ddi_dma_buf_setup(9F)`, `ddi_dma_curwin(9F)`, `ddi_dma_free(9F)`, `ddi_dma_movwin(9F)`, `ddi_dma_setup(9F)`, `ddi_dma_sync(9F)`, `mutex(9F)`

Writing Device Drivers

ddi-forceattach(9P)

NAME	ddi-forceattach, ddi-no-autodetach – properties controlling driver attach/detach behavior
DESCRIPTION	<p>Solaris device drivers are attached by <code>devfsadm(1M)</code> and by the kernel in response to <code>open(2)</code> requests from applications. Drivers not currently in use can be detached when the system experiences memory pressure. The <code>ddi-forceattach</code> and <code>ddi-no-autodetach</code> properties can be used to customize driver attach/detach behavior.</p> <p>The <code>ddi-forceattach</code> is an integer property, to be set globally by means of the <code>driver.conf(4)</code> file. Drivers with this property set to 1 are loaded and attached to all possible instances during system startup. The driver will not be auto-detached due to system memory pressure.</p> <p>The <code>ddi-no-autodetach</code> is an integer property to be set globally by means of the <code>driver.conf(4)</code> file or created dynamically by the driver on a per-instance basis with <code>ddi_prop_update_int(9F)</code>. When this property is set to 1, the kernel will not auto-detach driver due to system memory pressure.</p> <p>Note that <code>ddi-forceattach</code> implies <code>ddi-no-autodetach</code>. Setting either property to a non-integer value or an integer value not equal to 1 produces undefined results. These properties do not prevent driver detaching in response to reconfiguration requests, such as executing commands <code>cfgadm(1M)</code>, <code>modunload(1M)</code>, <code>rem_drv(1M)</code>, and <code>update_drv(1M)</code>.</p>
SEE ALSO	<p><code>driver.conf(4)</code></p> <p><i>Writing Device Drivers</i></p>

NAME	ddi_idevice_cookie – device interrupt cookie
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	The <code>ddi_idevice_cookie_t</code> structure contains interrupt priority and interrupt vector information for a device. This structure is useful for devices having programmable bus-interrupt levels. <code>ddi_add_intr(9F)</code> assigns values to the <code>ddi_idevice_cookie_t</code> structure members.
STRUCTURE MEMBERS	<pre>u_short idev_vector; /* interrupt vector */ ushort_t idev_priority; /* interrupt priority */</pre> <p>The <code>idev_vector</code> field contains the interrupt vector number for vectored bus architectures such as VMEbus. The <code>idev_priority</code> field contains the bus interrupt priority level.</p>
SEE ALSO	<p><code>ddi_add_intr(9F)</code></p> <p><i>Writing Device Drivers</i></p>

ddi_mapdev_ctl(9S)

NAME	ddi_mapdev_ctl – device mapping-control structure
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/devops.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	<p>Future releases of Solaris will provide this structure for binary and source compatibility. However, for increased functionality, use <code>devmap_callback_ctl(9S)</code> instead. See <code>devmap_callback_ctl(9S)</code> for details.</p> <p>A <code>ddi_mapdev_ctl</code> structure describes a set of routines that allow a device driver to manage events on mappings of the device created by <code>ddi_mapdev(9F)</code>.</p> <p>See <code>mapdev_access(9E)</code>, <code>mapdev_dup(9E)</code> and <code>mapdev_free(9E)</code> for more details on these entry points.</p>
STRUCTURE MEMBERS	<pre>int mapdev_rev; int (*mapdev_access)(ddi_mapdev_handle_t handle, void *devprivate, off_t offset); void (*mapdev_free)(ddi_mapdev_handle_t handle, void *devprivate); int (*mapdev_dup)(ddi_mapdev_handle_t handle, void *devprivate, ddi_mapdev_handle_t new_handle, void **new_devprivate);</pre> <p>A device driver should allocate the device mapping control structure and initialize the following fields:</p> <p><code>mapdev_rev</code> Must be set to <code>MAPDEV_REV</code>.</p> <p><code>mapdev_access</code> Must be set to the address of the <code>mapdev_access(9E)</code> entry point.</p> <p><code>mapdev_free</code> Must be set to the address of the <code>mapdev_free(9E)</code> entry point.</p> <p><code>mapdev_dup</code> Must be set to the address of the <code>mapdev_dup(9E)</code> entry point.</p>
SEE ALSO	<p><code>exit(2)</code>, <code>fork(2)</code>, <code>mmap(2)</code>, <code>munmap(2)</code>, <code>mapdev_access(9E)</code>, <code>mapdev_dup(9E)</code>, <code>mapdev_free(9E)</code>, <code>segmap(9E)</code>, <code>ddi_mapdev(9F)</code>, <code>ddi_mapdev_intercept(9F)</code>, <code>ddi_mapdev_nointercept(9F)</code></p> <p><i>Writing Device Drivers</i></p>

NAME	devmap_callback_ctl – device mapping-control structure
SYNOPSIS	<code>#include <sys/ddidevmap.h></code>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	<p>A <code>devmap_callback_ctl</code> structure describes a set of callback routines that are called by the system to notify a device driver to manage events on the device mappings created by <code>devmap_setup(9F)</code> or <code>ddi_devmap_segmap(9F)</code>.</p> <p>Device drivers pass the initialized <code>devmap_callback_ctl</code> structure to either <code>devmap_devmem_setup(9F)</code> or <code>devmap_umem_setup(9F)</code> in the <code>devmap(9E)</code> entry point during the mapping setup. The system makes a private copy of the structure for later use. Device drivers can specify different <code>devmap_callback_ctl</code> for different mappings.</p> <p>A device driver should allocate the device mapping control structure and initialize the following fields, if the driver wants the entry points to be called by the system:</p> <p><code>devmap_rev</code> Version number. Set this to <code>DEVMAP_OPS_REV</code>.</p> <p><code>devmap_map</code> Set to the address of the <code>devmap_map(9E)</code> entry point or to <code>NULL</code> if the driver does not support this callback. If set, the system calls the <code>devmap_map(9E)</code> entry point during the <code>mmap(2)</code> system call. The drivers typically allocate driver private data structure in this function and return the pointer to the private data structure to the system for later use.</p> <p><code>devmap_access</code> Set to the address of the <code>devmap_access(9E)</code> entry point or to <code>NULL</code> if the driver does not support this callback. If set, the system calls the driver's <code>devmap_access(9E)</code> entry point during memory access. The system expects <code>devmap_access(9E)</code> to call either <code>devmap_do_ctxmgt(9F)</code> or <code>devmap_default_access(9F)</code> to load the memory address translations before it returns to the system.</p> <p><code>devmap_dup</code> Set to the address of the <code>devmap_dup(9E)</code> entry point or to <code>NULL</code> if the driver does not support this call. If set, the system calls the <code>devmap_dup(9E)</code> entry point during the <code>fork(2)</code> system call.</p> <p><code>devmap_unmap</code> Set to the address of the <code>devmap_unmap(9E)</code> entry point or to <code>NULL</code> if the driver does not support this call. If set, the system will call the <code>devmap_unmap(9E)</code> entry point during the <code>munmap(2)</code> or <code>exit(2)</code> system calls.</p>
STRUCTURE MEMBERS	<pre>int devmap_rev; int (*devmap_map)(devmap_cookie_t dhp, dev_t dev, uint_t flags, offset_t off, size_t len, void **pvtp); int (*devmap_access)(devmap_cookie_t dhp, void *pvtp, offset_t off,</pre>

devmap_callback_ctl(9S)

```
        size_t len, uint_t type, uint_t rw);
int     (*devmap_dup)(devmap_cookie_t dhp, void *pvtp,
        devmap_cookie_t new_dhp, void **new_pvtp);
void    (*devmap_unmap)(devmap_cookie_t dhp, void *pvtp, offset_t off,
        size_t len, devmap_cookie_t new_dhp1, void **new_pvtp1,
        devmap_cookie_t new_dhp2, void **new_pvtp2);
```

SEE ALSO [exit\(2\)](#), [fork\(2\)](#), [mmap\(2\)](#), [munmap\(2\)](#), [devmap\(9E\)](#), [devmap_access\(9E\)](#), [devmap_dup\(9E\)](#), [devmap_map\(9E\)](#), [devmap_unmap\(9E\)](#), [ddi_devmap_segmap\(9F\)](#), [devmap_default_access\(9F\)](#), [devmap_devmem_setup\(9F\)](#), [devmap_do_ctxmgt\(9F\)](#), [devmap_setup\(9F\)](#), [devmap_umem_setup\(9F\)](#)

Writing Device Drivers

NAME	dev_ops – device operations structure																						
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/devops.h></pre>																						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).																						
DESCRIPTION	<p>dev_ops contains driver common fields and pointers to the bus_ops and cb_ops(9S).</p> <p>Following are the device functions provided in the device operations structure. All fields must be set at compile time.</p> <table border="0"> <tr> <td style="padding-right: 20px;">devo_rev</td> <td>Driver build version. Set this to DEVO_REV.</td> </tr> <tr> <td>devo_refcnt</td> <td>Driver reference count. Set this to 0.</td> </tr> <tr> <td>devo_getinfo</td> <td>Get device driver information (see getinfo(9E)).</td> </tr> <tr> <td>devo_identify</td> <td>Determine if a driver is associated with a device. See identify(9E).</td> </tr> <tr> <td>devo_probe</td> <td>Probe device. See probe(9E).</td> </tr> <tr> <td>devo_attach</td> <td>Attach driver to dev_info. See attach(9E).</td> </tr> <tr> <td>devo_detach</td> <td>Detach/prepare driver to unload. See detach(9E).</td> </tr> <tr> <td>devo_reset</td> <td>Reset device. (Not supported in this release.) Set this to nodev.</td> </tr> <tr> <td>devo_cb_ops</td> <td>Pointer to cb_ops(9S) structure for leaf drivers.</td> </tr> <tr> <td>devo_bus_ops</td> <td>Pointer to bus operations structure for nexus drivers. Set this to NULL if this is for a leaf driver.</td> </tr> <tr> <td>devo_power</td> <td>Power a device attached to system. See power(9E).</td> </tr> </table>	devo_rev	Driver build version. Set this to DEVO_REV.	devo_refcnt	Driver reference count. Set this to 0.	devo_getinfo	Get device driver information (see getinfo(9E)).	devo_identify	Determine if a driver is associated with a device. See identify(9E).	devo_probe	Probe device. See probe(9E).	devo_attach	Attach driver to dev_info. See attach(9E).	devo_detach	Detach/prepare driver to unload. See detach(9E).	devo_reset	Reset device. (Not supported in this release.) Set this to nodev.	devo_cb_ops	Pointer to cb_ops(9S) structure for leaf drivers.	devo_bus_ops	Pointer to bus operations structure for nexus drivers. Set this to NULL if this is for a leaf driver.	devo_power	Power a device attached to system. See power(9E).
devo_rev	Driver build version. Set this to DEVO_REV.																						
devo_refcnt	Driver reference count. Set this to 0.																						
devo_getinfo	Get device driver information (see getinfo(9E)).																						
devo_identify	Determine if a driver is associated with a device. See identify(9E).																						
devo_probe	Probe device. See probe(9E).																						
devo_attach	Attach driver to dev_info. See attach(9E).																						
devo_detach	Detach/prepare driver to unload. See detach(9E).																						
devo_reset	Reset device. (Not supported in this release.) Set this to nodev.																						
devo_cb_ops	Pointer to cb_ops(9S) structure for leaf drivers.																						
devo_bus_ops	Pointer to bus operations structure for nexus drivers. Set this to NULL if this is for a leaf driver.																						
devo_power	Power a device attached to system. See power(9E).																						
STRUCTURE MEMBERS	<pre>int devo_rev; int devo_refcnt; int (*devo_getinfo)(dev_info_t *dip, ddi_info_cmd_t infocmd, void *arg, void **result); int (*devo_identify)(dev_info_t *dip); int (*devo_probe)(dev_info_t *dip); int (*devo_attach)(dev_info_t *dip, ddi_attach_cmd_t cmd); int (*devo_detach)(dev_info_t *dip, ddi_detach_cmd_t cmd); int (*devo_reset)(dev_info_t *dip, ddi_reset_cmd_t cmd); struct cb_ops *devo_cb_ops; struct bus_ops *devo_bus_ops; int (*devo_power)(dev_info_t *dip, int component, int level);</pre>																						
SEE ALSO	attach(9E), detach(9E), getinfo(9E), identify(9E), probe(9E), power(9E), nodev(9F)																						
	<i>Writing Device Drivers</i>																						

fmodsw(9S)

NAME	fmodsw – STREAMS module declaration structure
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/conf.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	<p>The fmodsw structure contains information for STREAMS modules. All STREAMS modules must define a fmodsw structure.</p> <p>f_name must match mi_idname in the module_info structure. See module_info(9S). f_name should also match the module binary name. (See WARNINGS.)</p> <p>All modules must set the f_flag to D_MP to indicate that they safely allow multiple threads of execution. See mt-streams(9F) for additional flags.</p>
STRUCTURE MEMBERS	<pre>char f_name[FMNAMESZ + 1]; /* module name */ struct streamtab *f_str; /* streams information */ int f_flag; /* flags */</pre>
SEE ALSO	mt-streams(9F), modlstrmod(9S), module_info(9S) <i>STREAMS Programming Guide</i>
WARNINGS	If f_name does not match the module binary name, unexpected failures can occur.

NAME	free_rtn – structure that specifies a driver’s message-freeing routine
SYNOPSIS	<pre>#include <sys/stream.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	The free_rtn structure is referenced by the datab structure. When freeb(9F) is called to free the message, the driver’s message-freeing routine (referenced through the free_rtn structure) is called, with arguments, to free the data buffer.
STRUCTURE MEMBERS	<pre>void (*free_func)() /* user’s freeing routine */ char *free_arg /* arguments to free_func() */</pre> <p>The free_rtn structure is defined as type frtn_t.</p>
SEE ALSO	esballoc(9F), freeb(9F), datab(9S) <i>STREAMS Programming Guide</i>

gld_mac_info(9S)

NAME	gld_mac_info – Generic LAN Driver MAC info data structure
SYNOPSIS	#include <sys/gld.h>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
DESCRIPTION	<p>The Generic LAN Driver (GLD) Media Access Control (MAC) information (<code>gld_mac_info</code>) structure is the main data interface between the device-specific driver and GLD. It contains data required by GLD and a pointer to an optional additional driver-specific information structure.</p> <p>The <code>gld_mac_info</code> structure should be allocated using <code>gld_mac_alloc()</code> and deallocated using <code>gld_mac_free()</code>. Drivers can make no assumptions about the length of this structure, which might be different in different releases of Solaris and/or GLD. Structure members private to GLD, not documented here, should not be set or read by the device-specific driver.</p>
STRUCTURE MEMBERS	<pre> caddr_t gldm_private; /* Driver private data */ int (*gldm_reset)(); /* Reset device */ int (*gldm_start)(); /* Start device */ int (*gldm_stop)(); /* Stop device */ int (*gldm_set_mac_addr)(); /* Set device phys addr */ int (*gldm_set_multicast)(); /* Set/delete */ /* multicast address */ int (*gldm_set_promiscuous)(); /* Set/reset */ /* promiscuous mode */ int (*gldm_send)(); /* Transmit routine */ u_int (*gldm_intr)(); /* Interrupt handler */ int (*gldm_get_stats)(); /* Get device statistics */ int (*gldm_ioctl)(); /* Driver-specific ioctls */ char *gldm_ident; /* Driver identity string */ uint32_t gldm_type; /* Device type */ uint32_t gldm_minpkt; /* Minimum packet size */ /* accepted by driver */ uint32_t gldm_maxpkt; /* Maximum packet size */ /* accepted by driver */ uint32_t gldm_addrlen; /* Physical address */ /* length */ int32_t gldm_saplen; /* SAP length for */ /* DL_INFO_ACK */ unsigned char *gldm_broadcast_addr; /* Physical broadcast */ /* addr */ unsigned char *gldm_vendor_addr; /* Factory MAC address */ t_uscalar_t gldm_ppa; /* Physical Point of */ /* Attachment (PPA) number */ dev_info_t *gldm_devinfo; /* Pointer to device's */ /* dev_info node */ ddi_iblock_cookie_t gldm_cookie; /* Device's interrupt */ /* block cookie */ </pre> <p>Below is a description of the members of the <code>gld_mac_info</code> structure that are visible to the device driver.</p> <p><code>gldm_private</code> This structure member is private to the device-specific driver and is not used or modified by GLD.</p>

Conventionally, this is used as a pointer to private data, pointing to a driver-defined and driver-allocated per-instance data structure.

The following group of structure members must be set by the driver before calling `gld_register()`, and should not thereafter be modified by the driver; `gld_register()` can use or cache the values of some of these structure members, so changes made by the driver after calling `gld_register()` might cause unpredicted results.

<code>gldm_reset</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_start</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_stop</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_set_mac_addr</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_set_multicast</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_set_promiscuous</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_send</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_intr</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_get_stats</code>	Pointer to driver entry point; see <code>gld(9E)</code> .
<code>gldm_ioctl</code>	Pointer to driver entry point; can be NULL; see <code>gld(9E)</code> .
<code>gldm_ident</code>	Pointer to a string containing a short description of the device. It is used to identify the device in system messages.
<code>gldm_type</code>	The type of device the driver handles. The values currently supported by GLD are <code>DL_ETHER</code> (IEEE 802.3 and Ethernet Bus), <code>DL_TPR</code> (IEEE 802.5 Token Passing Ring), and <code>DL_FDDI</code> (ISO 9314-2 Fibre Distributed Data Interface). This structure member must be correctly set for GLD to function properly.
<code>gldm_minpkt</code>	Minimum <i>Service Data Unit</i> size — the minimum packet size, not including the MAC header, that the device will transmit. This can be zero if the device-specific driver can handle any required padding.
<code>gldm_maxpkt</code>	Maximum <i>Service Data Unit</i> size — the maximum size of packet, not including the MAC header, that can be transmitted by the device. For Ethernet, this number is 1500.

gld_mac_info(9S)

<code>gldm_addrlen</code>	The length in bytes of physical addresses handled by the device. For Ethernet, Token Ring, and FDDI, the value of this structure member should be 6.
<code>gldm_saplen</code>	The length in bytes of the Service Access Point (SAP) address used by the driver. For GLD-based drivers, this should always be set to -2, to indicate that two-byte SAP values are supported and that the SAP appears <i>after</i> the physical address in a DLSAP address. See the description under "Message DL_INFO_ACK" in the DLPI specification for more details.
<code>gldm_broadcast_addr</code>	Pointer to an array of bytes of length <code>gldm_addrlen</code> containing the broadcast address to be used for transmit. The driver must allocate space to hold the broadcast address, fill it in with the appropriate value, and set <code>gldm_broadcast_addr</code> to point at it. For Ethernet, Token Ring, and FDDI, the broadcast address is normally 0xFF-FF-FF-FF-FF-FF.
<code>gldm_vendor_addr</code>	Pointer to an array of bytes of length <code>gldm_addrlen</code> containing the vendor-provided network physical address of the device. The driver must allocate space to hold the address, fill it in with information read from the device, and set <code>gldm_vendor_addr</code> to point at it.
<code>gldm_ppa</code>	The Physical Point of Attachment (PPA) number for this instance of the device. Normally this should be set to the instance number, returned from <code>ddi_get_instance(9F)</code> .
<code>gldm_devinfo</code>	Pointer to the <code>dev_info</code> node for this device.
<code>gldm_cookie</code>	The interrupt block cookie returned by <code>ddi_get_iblock_cookie(9F)</code> , <code>ddi_add_intr(9F)</code> , <code>ddi_get_soft_iblock_cookie(9F)</code> , or <code>ddi_add_softintr(9F)</code> . This must correspond to the device's receive interrupt, from which <code>gld_recv()</code> is called.

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

Writing Device Drivers

NAME	gld_stats – Generic LAN Driver statistics data structure	
SYNOPSIS	#include <sys/gld.h>	
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).	
DESCRIPTION	<p>The Generic LAN Driver (GLD) statistics (<code>gld_stats</code>) structure is used to communicate statistics and state information from a GLD-based driver to GLD when returning from a driver's <code>gldm_get_stats()</code> routine as discussed in <code>gld(9E)</code> and <code>gld(7D)</code>. The members of this structure, filled in by the GLD-based driver, are used when GLD reports the statistics. In the tables below, the name of the statistics variable reported by GLD is noted in the comments. See <code>gld(7D)</code> for a more detailed description of the meaning of each statistic.</p> <p>Drivers can make no assumptions about the length of this structure, which might be different in different releases of Solaris and/or GLD. Structure members private to GLD, not documented here, should not be set or read by the device specific driver.</p>	
STRUCTURE MEMBERS	<p>The following structure members are defined for all media types:</p> <pre> uint64_t glds_speed; /* ifspeed */ uint32_t glds_media; /* media */ uint32_t glds_intr; /* intr */ uint32_t glds_norcvbuf; /* norcvbuf */ uint32_t glds_errrcv; /* ierrors */ uint32_t glds_errxmt; /* oerrors */ uint32_t glds_missed; /* missed */ uint32_t glds_underflow; /* uflo */ uint32_t glds_overflow; /* oflo */ </pre> <p>The following structure members are defined for media type <code>DL_ETHER</code>:</p> <pre> uint32_t glds_frame; /* align_errors */ uint32_t glds_crc; /* fcs_errors */ uint32_t glds_duplex; /* duplex */ uint32_t glds_nocarrier; /* carrier_errors */ uint32_t glds_collisions; /* collisions */ uint32_t glds_excoll; /* ex_collisions */ uint32_t glds_xmtlatecoll; /* tx_late_collisions */ uint32_t glds_defer; /* defer_xmts */ uint32_t glds_dot3_first_coll; /* first_collisions */ uint32_t glds_dot3_multi_coll; /* multi_collisions */ uint32_t glds_dot3_sqe_error; /* sqe_errors */ uint32_t glds_dot3_mac_xmt_error; /* macxmt_errors */ uint32_t glds_dot3_mac_rcv_error; /* macrcv_errors */ uint32_t glds_dot3_frame_too_long; /* toolong_errors */ uint32_t glds_short; /* runt_errors */ </pre> <p>The following structure members are defined for media type <code>DL_TPR</code>:</p> <pre> uint32_t glds_dot5_line_error /* line_errors */ uint32_t glds_dot5_burst_error /* burst_errors */ uint32_t glds_dot5_signal_loss /* signal_losses */ uint32_t glds_dot5_ace_error /* ace_errors */ uint32_t glds_dot5_internal_error /* internal_errors */ </pre>	

gld_stats(9S)

```
uint32_t    glds_dot5_lost_frame_error    /* lost_frame_errors */
uint32_t    glds_dot5_frame_copied_error  /* frame_copied_errors */
uint32_t    glds_dot5_token_error        /* token_errors */
uint32_t    glds_dot5_freq_error         /* freq_errors */
```

The following structure members are defined for media type DL_FDDI:

```
uint32_t    glds_fddi_mac_error;         /* mac_errors */
uint32_t    glds_fddi_mac_lost;         /* mac_lost_errors */
uint32_t    glds_fddi_mac_token;        /* mac_tokens */
uint32_t    glds_fddi_mac_tvx_expired;  /* mac_tvx_expired */
uint32_t    glds_fddi_mac_late;         /* mac_late */
uint32_t    glds_fddi_mac_ring_op;     /* mac_ring_ops */
```

Most of the above statistics variables are counters denoting the number of times the particular event was observed. Exceptions are:

glds_speed An estimate of the interface's current bandwidth in bits per second. For interfaces that do not vary in bandwidth or for those where no accurate estimation can be made, this object should contain the nominal bandwidth.

glds_media The type of media (wiring) or connector used by the hardware. Currently supported media names include GLDM_AUI, GLDM_BNC, GLDM_TP, GLDM_10BT, GLDM_100BT, GLDM_100BTX, GLDM_100BT4, GLDM_RING4, GLDM_RING16, GLDM_FIBER, and GLDM_PHYMII. GLDM_UNKNOWN can also be specified.

glds_duplex Current duplex state of the interface. Supported values are GLD_DUPLEX_HALF and GLD_DUPLEX_FULL. GLD_DUPLEX_UNKNOWN can also be specified.

SEE ALSO [gld\(7D\)](#), [gld\(9F\)](#), [gld\(9E\)](#), [gld_mac_info\(9S\)](#)

Writing Device Drivers

inquiry-device-type(9P)

NAME	inquiry-device-type, inquiry-vendor-id, inquiry-product-id, inquiry-revision-id – properties from SCSI inquiry data
DESCRIPTION	<p>These are optional properties created by the system for SCSI target devices.</p> <p><code>inquiry-device-type</code> is an integer property. When present, the least significant byte of the value indicates the device type as defined by the SCSI standard.</p> <p><code>inquiry-vendor-id</code> is a string property. When present, it contains the SCSI vendor identification inquiry data (from SCSI inquiry data bytes 8 - 15), formatted as a NULL-terminated string.</p> <p><code>inquiry-product-id</code> is a string property. When present, it contains the SCSI product identification inquiry data (from SCSI inquiry data bytes 16 - 31).</p> <p><code>inquiry-revision-id</code> is a string property. When present, it contains the SCSI product revision inquiry data (from SCSI inquiry data bytes 32 - 35).</p> <p>Consumers of these properties should compare the property values with <code>DTYPE_*</code> values defined in <code><sys/scsi/generic/inquiry.h></code>.</p>
SEE ALSO	<i>Writing Device Drivers</i>

iocblk(9S)

NAME	iocblk – STREAMS data structure for the M_IOCTL message type
SYNOPSIS	<pre>#include <sys/stream.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	The iocblk data structure is used for passing M_IOCTL messages.
STRUCTURE MEMBERS	<pre>int ioc_cmd; /* ioctl command type */ cred_t *ioc_cr; /* full credentials */ uint_t ioc_id; /* ioctl id */ uint_t ioc_flag; /* ioctl flags */ uint_t ioc_count; /* count of bytes in data field */ int ioc_rval; /* return value */ int ioc_error; /* error code */</pre>
SEE ALSO	<i>STREAMS Programming Guide</i>

NAME iovec – data storage structure for I/O using uio

SYNOPSIS #include <sys/uio.h>

INTERFACE LEVEL Architecture independent level 1 (DDI/DKI).

DESCRIPTION An iovec structure describes a data storage area for transfer in a uio(9S) structure. Conceptually, it can be thought of as a base address and length specification.

STRUCTURE MEMBERS

```
caddr_t   iov_base; /* base address of the data storage area */
           /* represented by the iovec structure */
int        iov_len; /* size of the data storage area in bytes */
```

SEE ALSO uio(9S)

Writing Device Drivers

kstat(9S)

NAME	kstat – kernel statistics structure												
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)												
DESCRIPTION	<p>Each kernel statistic (<i>kstat</i>) exported by device drivers consists of a header section and a data section. The <i>kstat</i> structure is the header portion of the statistic.</p> <p>A driver receives a pointer to a <i>kstat</i> structure from a successful call to <i>kstat_create(9F)</i>. Drivers should never allocate a <i>kstat</i> structure in any other manner.</p> <p>After allocation, the driver should perform any further initialization needed before calling <i>kstat_install(9F)</i> to actually export the <i>kstat</i>.</p>												
STRUCTURE MEMBERS	<pre>void *ks_data; /* kstat type-specific data */ ulong_t ks_ndata; /* # of type-specific data records */ ulong_t ks_data_size; /* total size of kstat data section */ int (*ks_update)(struct kstat *, int); void *ks_private; /* arbitrary provider-private data */ void *ks_lock; /* protects this kstat's data */</pre> <p>The members of the <i>kstat</i> structure available to examine or set by a driver are as follows:</p> <table border="0"> <tr> <td style="vertical-align: top;"><i>ks_data</i></td> <td>Points to the data portion of the <i>kstat</i>. Either allocated by <i>kstat_create(9F)</i> for the drivers use, or by the driver if it is using virtual <i>kstats</i>.</td> </tr> <tr> <td style="vertical-align: top;"><i>ks_ndata</i></td> <td>The number of data records in this <i>kstat</i>. Set by the <i>ks_update(9E)</i> routine.</td> </tr> <tr> <td style="vertical-align: top;"><i>ks_data_size</i></td> <td>The amount of data pointed to by <i>ks_data</i>. Set by the <i>ks_update(9E)</i> routine.</td> </tr> <tr> <td style="vertical-align: top;"><i>ks_update</i></td> <td>Pointer to a routine that dynamically updates <i>kstat</i>. This is useful for drivers where the underlying device keeps cheap hardware statistics, but where extraction is expensive. Instead of constantly keeping the <i>kstat</i> data section up to date, the driver can supply a <i>ks_update(9E)</i> function that updates the <i>kstat</i> data section on demand. To take advantage of this feature, set the <i>ks_update</i> field before calling <i>kstat_install(9F)</i>.</td> </tr> <tr> <td style="vertical-align: top;"><i>ks_private</i></td> <td>Is a private field for the driver's use. Often used in <i>ks_update(9E)</i>.</td> </tr> <tr> <td style="vertical-align: top;"><i>ks_lock</i></td> <td>Is a pointer to a mutex that protects this <i>kstat</i>. <i>kstat</i> data sections are optionally protected by the per-<i>kstat</i> <i>ks_lock</i>. If <i>ks_lock</i> is non-NULL, <i>kstat</i> clients (such as <i>/dev/kstat</i>) will</td> </tr> </table>	<i>ks_data</i>	Points to the data portion of the <i>kstat</i> . Either allocated by <i>kstat_create(9F)</i> for the drivers use, or by the driver if it is using virtual <i>kstats</i> .	<i>ks_ndata</i>	The number of data records in this <i>kstat</i> . Set by the <i>ks_update(9E)</i> routine.	<i>ks_data_size</i>	The amount of data pointed to by <i>ks_data</i> . Set by the <i>ks_update(9E)</i> routine.	<i>ks_update</i>	Pointer to a routine that dynamically updates <i>kstat</i> . This is useful for drivers where the underlying device keeps cheap hardware statistics, but where extraction is expensive. Instead of constantly keeping the <i>kstat</i> data section up to date, the driver can supply a <i>ks_update(9E)</i> function that updates the <i>kstat</i> data section on demand. To take advantage of this feature, set the <i>ks_update</i> field before calling <i>kstat_install(9F)</i> .	<i>ks_private</i>	Is a private field for the driver's use. Often used in <i>ks_update(9E)</i> .	<i>ks_lock</i>	Is a pointer to a mutex that protects this <i>kstat</i> . <i>kstat</i> data sections are optionally protected by the per- <i>kstat</i> <i>ks_lock</i> . If <i>ks_lock</i> is non-NULL, <i>kstat</i> clients (such as <i>/dev/kstat</i>) will
<i>ks_data</i>	Points to the data portion of the <i>kstat</i> . Either allocated by <i>kstat_create(9F)</i> for the drivers use, or by the driver if it is using virtual <i>kstats</i> .												
<i>ks_ndata</i>	The number of data records in this <i>kstat</i> . Set by the <i>ks_update(9E)</i> routine.												
<i>ks_data_size</i>	The amount of data pointed to by <i>ks_data</i> . Set by the <i>ks_update(9E)</i> routine.												
<i>ks_update</i>	Pointer to a routine that dynamically updates <i>kstat</i> . This is useful for drivers where the underlying device keeps cheap hardware statistics, but where extraction is expensive. Instead of constantly keeping the <i>kstat</i> data section up to date, the driver can supply a <i>ks_update(9E)</i> function that updates the <i>kstat</i> data section on demand. To take advantage of this feature, set the <i>ks_update</i> field before calling <i>kstat_install(9F)</i> .												
<i>ks_private</i>	Is a private field for the driver's use. Often used in <i>ks_update(9E)</i> .												
<i>ks_lock</i>	Is a pointer to a mutex that protects this <i>kstat</i> . <i>kstat</i> data sections are optionally protected by the per- <i>kstat</i> <i>ks_lock</i> . If <i>ks_lock</i> is non-NULL, <i>kstat</i> clients (such as <i>/dev/kstat</i>) will												

acquire this lock for all of their operations on that `kstat`. It is up to the `kstat` provider to decide whether guaranteeing consistent data to `kstat` clients is sufficiently important to justify the locking cost. Note, however, that most statistic updates already occur under one of the provider's mutexes. If the provider sets `ks_lock` to point to that mutex, then `kstat` data locking is free. `ks_lock` is really of type `(kmutex_t*)` and is declared as `(void*)` in the `kstat` header. That way, users do not have to be exposed to all of the kernel's lock-related data structures.

SEE ALSO `kstat_create(9F)`

Writing Device Drivers

kstat_intr(9S)

NAME	kstat_intr – structure for interrupt kstats
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	<p>Interrupt statistics are kept in the <code>kstat_intr</code> structure. When <code>kstat_create(9F)</code> creates an interrupt <code>kstat</code>, the <code>ks_data</code> field is a pointer to one of these structures. The macro <code>KSTAT_INTR_PTR()</code> is provided to retrieve this field. It looks like this:</p> <pre>#define KSTAT_INTR_PTR(kptr) ((kstat_intr_t *) (kptr)->ks_data)</pre> <p>An interrupt is a hard interrupt (sourced from the hardware device itself), a soft interrupt (induced by the system through the use of some system interrupt source), a watchdog interrupt (induced by a periodic timer call), spurious (an interrupt entry point was entered but there was no interrupt to service), or multiple service (an interrupt was detected and serviced just prior to returning from any of the other types).</p> <p>Drivers generally report only claimed hard interrupts and soft interrupts from their handlers, but measurement of the spurious class of interrupts is useful for auto-vectored devices in order to pinpoint any interrupt latency problems in a particular system configuration.</p> <p>Devices that have more than one interrupt of the same type should use multiple structures.</p>
STRUCTURE MEMBERS	<pre>ulong_t intrs[KSTAT_NUM_INTRS]; /* interrupt counters */</pre> <p>The only member exposed to drivers is the <code>intrs</code> member. This field is an array of counters. The driver must use the appropriate counter in the array based on the type of interrupt condition.</p> <p>The following indexes are supported:</p> <pre>KSTAT_INTR_HARD Hard interrupt KSTAT_INTR_SOFT Soft interrupt KSTAT_INTR_WATCHDOG Watchdog interrupt KSTAT_INTR_SPURIOUS Spurious interrupt KSTAT_INTR_MULTSVC Multiple service interrupt</pre>
SEE ALSO	<code>kstat(9S)</code>

kstat_io(9S)

NAME	kstat_io – structure for I/O kstats
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	<p>I/O kstat statistics are kept in a <code>kstat_io</code> structure. When <code>kstat_create(9F)</code> creates an I/O kstat, the <code>ks_data</code> field is a pointer to one of these structures. The macro <code>KSTAT_IO_PTR()</code> is provided to retrieve this field. It looks like this:</p> <pre>#define KSTAT_IO_PTR(kptr) ((kstat_io_t *) (kptr)->ks_data)</pre>
STRUCTURE MEMBERS	<pre>u_longlong_t nread; /* number of bytes read */ u_longlong_t nwritten; /* number of bytes written */ ulong_t reads; /* number of read operations */ ulong_t writes; /* number of write operations */</pre> <p>The <code>nread</code> field should be updated by the driver with the number of bytes successfully read upon completion.</p> <p>The <code>nwritten</code> field should be updated by the driver with the number of bytes successfully written upon completion.</p> <p>The <code>reads</code> field should be updated by the driver after each successful read operation.</p> <p>The <code>writes</code> field should be updated by the driver after each successful write operation</p> <p>Other I/O statistics are updated through the use of the <code>kstat_queue(9F)</code> functions.</p>
SEE ALSO	<p><code>kstat_create(9F)</code>, <code>kstat_named_init(9F)</code>, <code>kstat_queue(9F)</code>, <code>kstat_runq_back_to_waitq(9F)</code>, <code>kstat_runq_enter(9F)</code>, <code>kstat_runq_exit(9F)</code>, <code>kstat_waitq_enter(9F)</code>, <code>kstat_waitq_exit(9F)</code>, <code>kstat_waitq_to_runq(9F)</code></p> <p><i>Writing Device Drivers</i></p>

NAME	kstat_named – structure for named kstats
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	Named kstats are an array of name-value pairs. These pairs are kept in the kstat_named structure. When a kstat is created by kstat_create(9F), the driver specifies how many of these structures will be allocated. The structures are returned as an array pointed to by the ks_data field.
STRUCTURE MEMBERS	<pre>union { char c[16]; long l; ulong_t ul; longlong_t ll; u_longlong_t ull; } value; /* value of counter */</pre> <p>The only member exposed to drivers is the value member. This field is a union of several data types. The driver must specify which type it will use in the call to kstat_named_init().</p>
SEE ALSO	kstat_create(9F), kstat_named_init(9F) <i>Writing Device Drivers</i>

linkblk(9S)

NAME	linkblk – STREAMS data structure sent to multiplexor drivers to indicate a link
SYNOPSIS	<pre>#include <sys/stream.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	The linkblk structure is used to connect a lower Stream to an upper STREAMS multiplexor driver. This structure is used in conjunction with the I_LINK, I_UNLINK, P_LINK, and P_UNLINK ioctl commands. See streamio(7I). The M_DATA portion of the M_IOCTL message contains the linkblk structure. Note that the linkblk structure is allocated and initialized by the Stream head as a result of one of the above ioctl commands.
STRUCTURE MEMBERS	<pre>queue_t *l_qtop; /* lowest level write queue of upper stream */ /* (set to NULL for persistent links) */ queue_t *l_qbot; /* highest level write queue of lower stream */ int l_index; /* index for lower stream. */</pre>
SEE ALSO	ioctl(2), streamio(7I) <i>STREAMS Programming Guide</i>

NAME	modldrv – linkage structure for loadable drivers
SYNOPSIS	<code>#include <sys/modctl.h></code>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	The modldrv structure is used by device drivers to export driver specific information to the kernel.
STRUCTURE MEMBERS	<pre> struct mod_ops *drv_modops; char *drv_link_info; struct dev_ops *drv_dev_ops; </pre> <p><code>drv_modops</code> Must always be initialized to the address of <code>mod_driverops</code>. This member identifies the module as a loadable driver.</p> <p><code>drv_linkinfo</code> Can be any string up to <code>MODMAXNAMELEN</code> characters (including the terminating NULL character), and is used to describe the module and its version number. This is usually the name of the driver and module version information, but can contain other information as well.</p> <p><code>drv_dev_ops</code> Pointer to the driver's <code>dev_ops(9S)</code> structure.</p>
SEE ALSO	<code>add_drv(1M)</code> , <code>dev_ops(9S)</code> , <code>modlinkage(9S)</code> <i>Writing Device Drivers</i>

modlinkage(9S)

NAME	modlinkage – module linkage structure
SYNOPSIS	#include <sys/modctl.h>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	The modlinkage structure is provided by the module writer to the routines that install, remove, and retrieve information from a module. See <code>_init(9E)</code> , <code>_fini(9E)</code> , and <code>_info(9E)</code> .
STRUCTURE MEMBERS	<pre>int ml_rev void *ml_linkage[4];</pre> <p><code>ml_rev</code> Is the revision of the loadable modules system. This must have the value <code>MODREV_1</code>.</p> <p><code>ml_linkage</code> Is a null-terminated array of pointers to linkage structures. Driver modules have only one linkage structure.</p>
SEE ALSO	<code>add_drv(1M)</code> , <code>_fini(9E)</code> , <code>_info(9E)</code> , <code>_init(9E)</code> , <code>modldrv(9S)</code> , <code>modlstrmod(9S)</code> <i>Writing Device Drivers</i>

NAME	modlstrmod – linkage structure for loadable STREAMS modules
SYNOPSIS	<code>#include <sys/modctl.h></code>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	The <code>modlstrmod</code> structure is used by STREAMS modules to export module specific information to the kernel.
STRUCTURE MEMBERS	<pre> struct mod_ops *strmod_modops; char *strmod_linkinfo; struct fmodsw *strmod_fmodsw; </pre> <p><code>strmod_modops</code> Must always be initialized to the address of <code>mod_strmodops</code>. This identifies the module as a loadable STREAMS module.</p> <p><code>strmod_linkinfo</code> Can be any string up to <code>MODMAXNAMELEN</code>, and is used to describe the module. This string is usually the name of the module, but can contain other information (such as a version number).</p> <p><code>strmod_fmodsw</code> Is a pointer to a template of a class entry within the module that is copied to the kernel's class table when the module is loaded.</p>
SEE ALSO	<p><code>modload(1M)</code></p> <p><i>Writing Device Drivers</i></p>

module_info(9S)

NAME	module_info – STREAMS driver identification and limit value structure
SYNOPSIS	<pre>#include <sys/stream.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	<p>When a module or driver is declared, several identification and limit values can be set. These values are stored in the <code>module_info</code> structure.</p> <p>The <code>module_info</code> structure is intended to be read-only. However, the flow control limits (<code>mi_hiwat</code> and <code>mi_lowat</code>) and the packet size limits (<code>mi_minpsz</code> and <code>mi_maxpsz</code>) are copied to the <code>QUEUE</code> structure, where they can be modified.</p> <p>For a driver, <code>mi_idname</code> must match the name of the driver binary file. For a module, <code>mi_idname</code> must match the <code>fname</code> field of the <code>fmodsw</code> structure. See <code>fmodsw(9S)</code> for details.</p>
STRUCTURE MEMBERS	<pre>ushort_t mi_idnum; /* module ID number */ char *mi_idname; /* module name */ ssize_t mi_minpsz; /* minimum packet size */ ssize_t mi_maxpsz; /* maximum packet size */ size_t mi_hiwat; /* high water mark */ size_t mi_lowat; /* low water mark */</pre> <p>The constant <code>FMNAMESZ</code>, limiting the length of a module's name, is set to eight in this release.</p>
SEE ALSO	<p><code>fmodsw(9S)</code>, <code>queue(9S)</code></p> <p><i>STREAMS Programming Guide</i></p>

NAME	msgb – STREAMS message block structure
SYNOPSIS	<code>#include <sys/stream.h></code>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	<p>A STREAMS message is made up of one or more message blocks, referenced by a pointer to a msgb structure. The <code>b_next</code> and <code>b_prev</code> pointers are used to link messages together on a QUEUE. The <code>b_cont</code> pointer links message blocks together when a message consists of more than one block.</p> <p>Each msgb structure also includes a pointer to a datab(9S) structure, the data block (which contains pointers to the actual data of the message), and the type of the message.</p>
STRUCTURE MEMBERS	<pre> struct msgb *b_next; /* next message on queue */ struct msgb *b_prev; /* previous message on queue */ struct msgb *b_cont; /* next message block */ unsigned char *b_rptr; /* 1st unread data byte of buffer */ unsigned char *b_wptr; /* 1st unwritten data byte of buffer */ struct datab *b_datap; /* pointer to data block */ unsigned char b_band; /* message priority */ unsigned short b_flag; /* used by stream head */ </pre> <p>Valid flags are as follows:</p> <p>MSGMARK Last byte of message is marked.</p> <p>MSGDELIM Message is delimited.</p> <p>The msgb structure is defined as type <code>mblk_t</code>.</p>
SEE ALSO	<p>datab(9S)</p> <p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>

no-involuntary-power-cycles(9P)

NAME	no-involuntary-power-cycles – device property to prevent involuntary power cycles
DESCRIPTION	<p>A device that might be damaged by power cycles should export the boolean (zero length) property <code>no-involuntary-power-cycles</code> to notify the system that all power cycles for the device must be under the control of the device driver.</p> <p>The presence of this property prevents power from being removed from a device or any ancestor of the device while the device driver is detached, unless the device was voluntarily powered off as a result of the device driver calling <code>pm_lower_power(9F)</code>.</p> <p>The presence of <code>no-involuntary-power-cycles</code> also forces attachment of the device driver during a CPR suspend operation and prevents the suspend from taking place, unless the device driver returns <code>DDI_SUCCESS</code> when its <code>detach(9E)</code> entry point is called with <code>DDI_SUSPEND</code>.</p> <p>The presence of <code>no-involuntary-power-cycles</code> does not prevent the system from being powered off due to a <code>halt(1M)</code> or <code>uadmin(1M)</code> invocation, except for CPR suspend.</p> <p>This property can be exported by a device that is not power manageable, in which case power is not removed from the device or from any of its ancestors, even when the driver for the device and the drivers for its ancestors are detached.</p>
EXAMPLES	<p>EXAMPLE 1 Use of Property in Driver's Configuration File</p> <p>The following is an example of a <code>no-involuntary-power-cycles</code> entry in a driver's <code>.conf</code> file:</p> <pre>no-involuntary-power-cycles=1; ...</pre> <p>EXAMPLE 2 Use of Property in <code>attach()</code> Function</p> <p>The following is an example of how the preceding <code>.conf</code> file entry would be implemented in the <code>attach(9E)</code> function of a driver:</p> <pre>xxattach(dev_info_t *dip, ddi_attach_cmd_t cmd) { ... if (ddi_prop_create(DDI_DEV_T_NONE, dip, DDI_PROP_CANSLEEP, "no-involuntary-power-cycles", NULL, 0) != DDI_PROP_SUCCESS) goto failed; ... }</pre>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

no-involuntary-power-cycles(9P)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Evolving

SEE ALSO [attributes\(5\)](#), [pm\(7D\)](#), [attach\(9E\)](#), [detach\(9E\)](#), [ddi_prop_create\(9F\)](#)

Writing Device Drivers

pm(9P)

NAME	pm – Power Management properties
DESCRIPTION	<p>The <code>pm-hardware-state</code> property can be used to influence the behavior of the Power Management framework. Its syntax and interpretation is described below.</p> <p>Note that this property is only interpreted by the system immediately after the device has successfully attached. Changes in the property made by the driver after the driver has attached will not be recognized.</p> <p><code>pm-hardware-state</code> is a string-valued property. The existence of the <code>pm-hardware-state</code> property indicates that a device needs special handling by the Power Management framework with regard to its hardware state.</p> <p>If the value of this property is <code>needs-suspend-resume</code>, the device has a hardware state that cannot be deduced by the framework. The framework definition of a device with hardware state is one with a <code>reg</code> property. Some drivers, such as SCSI disk and tape drivers, have no <code>reg</code> property but manage devices with "remote" hardware. Such a device must have a <code>pm-hardware-state</code> property with a value of <code>needs-suspend-resume</code> for the system to identify it as needing a call to its <code>detach(9E)</code> entry point with command <code>DDI_SUSPEND</code> when system is suspended, and a call to <code>attach(9E)</code> with command <code>DDI_RESUME</code> when system is resumed. For devices using original Power Management interfaces (which are now obsolete) <code>detach(9E)</code> is also called with <code>DDI_PM_SUSPEND</code> before power is removed from the device, and <code>attach(9E)</code> is called with <code>DDI_PM_RESUME</code> after power is restored.</p> <p>A value of <code>no-suspend-resume</code> indicates that, in spite of the existence of a <code>reg</code> property, a device has no hardware state that needs saving and restoring. A device exporting this property will not have its <code>detach()</code> entry point called with command <code>DDI_SUSPEND</code> when system is suspended, nor will its <code>attach()</code> entry point be called with command <code>DDI_RESUME</code> when system is resumed. For devices using the original (and now obsolete) Power Management interfaces, <code>detach(9E)</code> will not be called with <code>DDI_PM_SUSPEND</code> command before power is removed from the device, nor <code>attach(9E)</code> will be called with <code>DDI_PM_RESUME</code> command after power is restored to the device.</p> <p>A value of <code>parental-suspend-resume</code> indicates that the device does not implement the <code>detach(9E)</code> <code>DDI_SUSPEND</code> semantics, nor the <code>attach()</code> <code>DDI_RESUME</code> semantics, but that a call should be made up the device tree by the framework to effect the saving and/or restoring of hardware state for this device. For devices using original Power Management interfaces (which are now obsolete), it also indicates that the device does not implement the <code>detach(9E)</code> <code>DDI_PM_SUSPEND</code> semantics, nor the <code>attach(9E)</code> <code>DDI_PM_RESUME</code> semantics, but that a call should be made up the device tree by the framework to effect the saving and/or restoring the hardware state for this device.</p>

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Evolving

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

Writing Device Drivers

pm-components(9P)

NAME	pm-components – Power Management device property
DESCRIPTION	<p>A device is power manageable if the power consumption of the device can be reduced when it is idle. In general, a power manageable device consists of a number of power manageable hardware units called components. Each component is separately controllable and has its own set of power parameters.</p> <p>An example of a one-component power manageable device is a disk whose spindle motor can be stopped to save power when the disk is idle. An example of a two-component power manageable device is a frame buffer card with a connected monitor. The frame buffer electronics (with power that can be reduced when not in use) comprises the first component. The second component is the monitor, which can enter in a lower power mode when not in use. The combination of frame buffer electronics and monitor is considered as one device by the system.</p> <p>In the Power Management framework, all components are considered equal and completely independent of each other. If this is not true for a particular device, the device driver must ensure that undesirable state combinations do not occur.</p> <p>The <code>pm-components</code> property describes the Power Management model of a device driver to the Power Management framework. It lists each power manageable component by name and lists the power level supported by each component by numerical value and name. Its syntax and interpretation is described below.</p> <p>This property is only interpreted by the system immediately after the device has successfully attached, or upon the first call into Power Management framework, whichever comes first. Changes in the property made by the driver after the property has been interpreted will not be recognized.</p> <p><code>pm-components</code> is a string array property. The existence of the <code>pm-components</code> property indicates that a device implements power manageable components and describes the Power Management model implemented by the device driver. The existence of <code>pm-components</code> also indicates to the framework that device is ready for Power Management if automatic device Power Management is enabled. See <code>power.conf(4)</code>.</p> <p>The <code>pm-component</code> property syntax is:</p> <pre>pm-components="NAME=component name", "numeric power level=power level name", "numeric power level=power level name" [, "numeric power level=power level name" ...] [, "NAME=component name", "numeric power level=power level name", "numeric power level=power level name" [, "numeric power level=power level name"...]...];</pre> <p>The start of each new component is represented by a string consisting of <code>NAME=</code> followed by the name of the component. This should be a short name that a user would recognize, such as "Monitor" or "Spindle Motor." The succeeding elements in the string array must be strings consisting of the numeric value (can be decimal or 0x <hexadecimal number>) of a power level the component supports, followed by an equal sign followed by a short descriptive name for that power level. Again, the</p>

names should be descriptive, such as "On," "Off," "Suspend," "Standby," etc. The next component continues the array in the same manner, with a string that starts out `NAME=`, specifying the beginning of a new component (and its name), followed by specifications of the power levels the component supports.

The components must be listed in increasing order according to the component number as interpreted by the driver's `power(9E)` routine. (Components are numbered sequentially from 0). The power levels must be listed in increasing order of power consumption. Each component must support at least two power levels, or there is no possibility of power level transitions. If a power level value of 0 is used, it must be the first one listed for that component. A power level value of 0 has a special meaning (off) to the Power Management framework.

EXAMPLES

An example of a `pm-components` entry from the `.conf` file of a driver which implements a single power managed component consisting of a disk spindle motor is shown below. This is component 0 and it supports 2 power level, which represent spindle stopped or full speed.

```
pm-components="NAME=Spindle Motor", "0=Stopped", "1=Full Speed";
...
```

Below is an example of how the above entry would be implemented in the `attach(9E)` function of the driver.

```
static char *pmcomps[] = {
    "NAME=Spindle Motor",
    "0=Stopped",
    "1=Full Speed"
};
...

xxattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    ...
    if (ddi_prop_update_string_array(DDI_DEV_T_NONE, dip, "pm-components",
        &pmcomp[0], sizeof (pmcomps) / sizeof (char *)) != DDI_PROP_SUCCESS)
        goto failed;
}
}
```

Below is an example for a frame buffer which implements two components. Component 0 is the frame buffer electronics which supports four different power levels. Component 1 represents the state of Power Management of the attached monitor.

```
pm-components="NAME=Frame Buffer", "0=Off"
    "1=Suspend", "2=Standby", "3=On",
    "NAME=Monitor", "0=Off", "1=Suspend", "2=Standby,"
    "3=On;
```

pm-components(9P)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Evolving

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

Writing Device Drivers

NAME	qband – STREAMS queue flow control information structure
SYNOPSIS	<code>#include <sys/stream.h></code>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	<p>The qband structure contains flow control information for each priority band in a queue.</p> <p>The qband structure is defined as type qband_t.</p>
STRUCTURE MEMBERS	<pre> struct qband*qb_next; /* next band's info */ size_t qb_count /* number of bytes in band */ struct msgb *qb_first; /* start of band's data */ struct msgb *qb_last; /* end of band's data */ size_t qb_hiwat; /* band's high water mark */ size_t qb_lowat; /* band's low water mark */ uint_t qb_flag; /* see below */ </pre> <p>Valid flags are as follows:</p> <p>QB_FULL Band is considered full.</p> <p>QB_WANTW Someone wants to write to band.</p>
SEE ALSO	strqget(9F), strqset(9F), msgb(9S), queue(9S)
NOTES	<p><i>STREAMS Programming Guide</i></p> <p>All access to this structure should be through strqget(9F) and strqset(9F). It is logically part of the queue(9S) and its layout and partitioning with respect to that structure might change in future releases. If portability is a concern, do not declare or store instances of or references to this structure.</p>

qinit(9S)

NAME	qinit – STREAMS queue processing procedures structure
SYNOPSIS	<pre>#include <sys/stream.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	The <code>qinit</code> structure contains pointers to processing procedures for a <code>QUEUE</code> . The <code>streamtab</code> structure for the module or driver contains pointers to one <code>queue(9S)</code> structure for both upstream and downstream processing.
STRUCTURE MEMBERS	<pre>int (*qi_putp) (); /* put procedure */ int (*qi_srvp) (); /* service procedure */ int (*qi_qopen) (); /* open procedure */ int (*qi_qclose) (); /* close procedure */ int (*qi_qadmin) (); /* unused */ struct module_info *qi_minfo; /* module parameters */ struct module_stat *qi_mstat; /* module statistics */</pre>
SEE ALSO	<code>queue(9S)</code> , <code>streamtab(9S)</code> <i>Writing Device Drivers</i> <i>STREAMS Programming Guide</i>
NOTES	This release includes no support for module statistics.

NAME	queclass – a STREAMS macro that returns the queue message class definitions for a given message block
SYNOPSIS	<pre>#include <sys/stream.h> queclass(mblk_t *bp);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	<p>queclass returns the queue message class definition for a given data block pointed to by the message block <i>bp</i> passed in.</p> <p>The message can be either QNORM, a normal priority message, or QPCTL, a high priority message.</p>
SEE ALSO	<i>STREAMS Programming Guide</i>

queue(9S)

NAME	queue – STREAMS queue structure														
SYNOPSIS	#include <sys/stream.h>														
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)														
DESCRIPTION	<p>A STREAMS driver or module consists of two queue structures, one for upstream processing (read) and one for downstream processing (write). This structure is the major building block of a stream. It contains pointers to the processing procedures, pointers to the next and previous queues in the stream, flow control parameters, and a pointer defining the position of its messages on the STREAMS scheduler list.</p> <p>The queue structure is defined as type <code>queue_t</code>.</p>														
STRUCTURE MEMBERS	<pre>struct qinit*q_qinfo; /* module or driver entry points */ struct msgb*q_first; /* first message in queue */ struct msgb*q_last; /* last message in queue */ struct queue*q_next; /* next queue in stream */ struct queue*q_link; /* to next queue for scheduling*/ void *q_ptr; /* pointer to private data structure */ size_t q_count; /* approximate size of message queue */ uint_t q_flag; /* status of queue */ ssize_t q_minpsz; /* smallest packet accepted by QUEUE*/ ssize_t q_maxpsz; /*largest packet accepted by QUEUE */ size_t q_hiwat; /* high water mark */ size_t q_lowat; /* low water mark */</pre>														
	<p>Valid flags are as follows:</p> <table><tr><td>QENAB</td><td>Queue is already enabled to run.</td></tr><tr><td>QWANTR</td><td>Someone wants to read queue.</td></tr><tr><td>QWANTW</td><td>Someone wants to write to queue.</td></tr><tr><td>QFULL</td><td>Queue is considered full.</td></tr><tr><td>QREADR</td><td>This is the reader (first) queue.</td></tr><tr><td>QUSE</td><td>This queue is in use (allocation).</td></tr><tr><td>QNOENB</td><td>Do not enable queue by way of <code>putq()</code>.</td></tr></table>	QENAB	Queue is already enabled to run.	QWANTR	Someone wants to read queue.	QWANTW	Someone wants to write to queue.	QFULL	Queue is considered full.	QREADR	This is the reader (first) queue.	QUSE	This queue is in use (allocation).	QNOENB	Do not enable queue by way of <code>putq()</code> .
QENAB	Queue is already enabled to run.														
QWANTR	Someone wants to read queue.														
QWANTW	Someone wants to write to queue.														
QFULL	Queue is considered full.														
QREADR	This is the reader (first) queue.														
QUSE	This queue is in use (allocation).														
QNOENB	Do not enable queue by way of <code>putq()</code> .														
SEE ALSO	<code>strqget(9F)</code> , <code>strqset(9F)</code> , <code>module_info(9S)</code> , <code>msgb(9S)</code> , <code>qinit(9S)</code> , <code>streamtab(9S)</code>														
	<p><i>Writing Device Drivers</i></p> <p><i>STREAMS Programming Guide</i></p>														

NAME removable-media – removable media device property

DESCRIPTION A device that supports removable media—such as CDROM, JAZZ, and ZIP drives—and that supports power management and expects automatic mounting of the device via the volume manager should export the boolean (zero length) property `removable-media`. This property enables the system to make the power state of the device dependent on the power state of the frame buffer and monitor. See the `power.conf(4)` discussion of the `device-dependency-property` entry for more information.

Devices that behave like removable devices (such as PC ATA cards, where the controller and media both are removed at the same time) should also export this property.

EXAMPLES **EXAMPLE 1** `removable-media` Entry

An example of a `removable-media` entry from the `.conf` file of a driver is shown below.

```
# This entry keeps removable media from being powered down unless
# the console framebuffer and monitor are powered down
#
removable-media=1;
```

EXAMPLE 2 Implementation in `attach()`

Below is an example of how the entry above would be implemented in the `attach(9E)` function of the driver.

```
xxattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    ...
    if (ddi_prop_create(DDI_DEV_T_NONE, dip, DDI_PROP_CANSLEEP,
        "removable-media", NULL, 0) != DDI_PROP_SUCCESS)
        goto failed;
    ...
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Evolving

SEE ALSO `power.conf(4)`, `pm(7D)`, `attach(9E)`, `detach(9E)`, `ddi_prop_create(9F)`

Writing Device Drivers

scsi_address(9S)

NAME	scsi_address – SCSI address structure
SYNOPSIS	<pre>#include <sys/scsi/scsi.h></pre>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI)
DESCRIPTION	<p>A <code>scsi_address</code> structure defines the addressing components for a SCSI target device. The address of the target device is separated into two components: target number and logical unit number. The two addressing components are used to uniquely identify any type of SCSI device; however, most devices can be addressed with the target component of the address.</p> <p>In the case where only the target component is used to address the device, the logical unit should be set to 0. If the SCSI target device supports logical units, then the HBA must interpret the logical units field of the data structure.</p> <p>The <code>pkt_address</code> member of a <code>scsi_pkt(9S)</code> is initialized by <code>scsi_init_pkt(9F)</code>.</p>
STRUCTURE MEMBERS	<pre>scsi_hba_tran_t *a_hba_tran; /* Transport vectors for the SCSI bus */ ushort_t a_target; /* SCSI target id */ uchar_t a_lun; /* SCSI logical unit */</pre> <p><code>a_hba_tran</code> is a pointer to the controlling HBA's transport vector structure. The SCSSA interface uses this field to pass any transport requests from the SCSI target device drivers to the HBA driver.</p> <p><code>a_target</code> is the target component of the SCSI address.</p> <p><code>a_lun</code> is the logical unit component of the SCSI address. The logical unit is used to further distinguish a SCSI target device that supports multiple logical units from one that does not. The <code>makecom(9F)</code> family of functions use the <code>a_lun</code> field to set the logical unit field in the SCSI CDB, for compatibility with SCSI-1.</p>
SEE ALSO	<code>makecom(9F)</code> , <code>scsi_init_pkt(9F)</code> , <code>scsi_hba_tran(9S)</code> , <code>scsi_pkt(9S)</code> <i>Writing Device Drivers</i>

NAME	scsi_arq_status – SCSI auto request sense structure
SYNOPSIS	#include <sys/scsi/scsi.h>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)
DESCRIPTION	<p>When auto request sense has been enabled using <code>scsi_ifsetcap(9F)</code> and the "auto-rqsense" capability, the target driver must allocate a status area in the SCSI packet structure for the auto request sense structure (see <code>scsi_pkt(9S)</code>). In the event of a check <i>condition</i>, the transport layer automatically executes a request sense command. This check ensures that the request sense information does not get lost. The auto request sense structure supplies the SCSI status of the original command, the transport information pertaining to the request sense command, and the request sense data.</p>
STRUCTURE MEMBERS	<pre> struct scsi_status sts_status; /* SCSI status */ struct scsi_status sts_rqpkt_status; /* SCSI status of request sense cmd */ uchar_t sts_rqpkt_reason; /* reason completion */ uchar_t sts_rqpkt_resid; /* residue */ uint_t sts_rqpkt_state; /* state of command */ uint_t sts_rqpkt_statistics; /* statistics */ struct scsi_extended_sense sts_sensedata; /* actual sense data */ </pre> <p><code>sts_status</code> is the SCSI status of the original command. If the status indicates a check <i>condition</i>, the transport layer might have performed an auto request sense command.</p> <p><code>sts_rqpkt_status</code> is the SCSI status of the request sense command.</p> <p><code>sts_rqpkt_reason</code> is the completion reason of the request sense command. If the reason is not <code>CMD_CMPLT</code>, then the request sense command did not complete normally.</p> <p><code>sts_rqpkt_resid</code> is the residual count of the data transfer and indicates the number of data bytes that have not been transferred. The auto request sense command requests <code>SENSE_LENGTH</code> bytes.</p> <p><code>sts_rqpkt_state</code> has bit positions representing the five most important statuses that a SCSI command can go obtain.</p> <p><code>sts_rqpkt_statistics</code> maintains transport-related statistics of the request sense command.</p> <p><code>sts_sensedata</code> contains the actual sense data if the request sense command completed normally.</p>
SEE ALSO	<p><code>scsi_ifgetcap(9F)</code>, <code>scsi_init_pkt(9F)</code>, <code>scsi_extended_sense(9S)</code>, <code>scsi_pkt(9S)</code></p> <p><i>Writing Device Drivers</i></p>

scsi_asc_key_strings(9S)

NAME	scsi_asc_key_strings – SCSI ASC ASCQ to message structure
SYNOPSIS	#include <sys/scsi/scsi.h>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	The <code>scsi_asc_key_strings</code> structure stores the ASC and ASCQ codes and a pointer to the related ASCII string.
STRUCTURE MEMBERS	<pre>ushort_t asc; /* ASC code */ ushort_t ascq; /* ASCQ code */ char *message; /* ASCII message string */ asc Contains the ASC key code. ascq Contains the ASCQ code. message Points to the NULL terminated ASCII string describing the asc and ascq condition</pre>
SEE ALSO	<code>scsi_vu_errmsg(9F)</code> <i>ANSI Small Computer System Interface-2 (SCSI-2)</i> <i>Writing Device Drivers</i>

NAME	scsi_device – SCSI device structure
SYNOPSIS	<code>#include <sys/scsi/scsi.h></code>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	<p>The <code>scsi_device</code> structure stores common information about each SCSI logical unit, including pointers to areas that contain both generic and device specific information. There is one <code>scsi_device</code> structure for each logical unit attached to the system. The host adapter driver initializes part of this structure prior to <code>probe(9E)</code> and destroys this structure after a probe failure or successful <code>detach(9E)</code>.</p>
STRUCTURE MEMBERS	<pre> struct scsi_address sd_address; /* Routing information */ dev_info_t *sd_dev; /* Cross-reference */ /* to our dev_info_t */ kmutex_t sd_mutex; /* Mutex for this device */ struct scsi_inquiry *sd_inq; /* scsi_inquiry data structure */ struct scsi_extended_sense *sd_sense; /* Optional request */ /* sense buffer ptr */ caddr_t sd_private; /* Target drivers private data */ </pre> <p><code>sd_address</code> contains the routing information that the target driver normally copies into a <code>scsi_pkt(9S)</code> structure using the collection of <code>makecom(9F)</code> functions. The SCSI library routines use this information to determine which host adapter, SCSI bus, and target/logical unit number (lun) a command is intended for. This structure is initialized by the host adapter driver.</p> <p><code>sd_dev</code> is a pointer to the corresponding <code>dev_info</code> structure. This pointer is initialized by the host adapter driver.</p> <p><code>sd_mutex</code> is a mutual exclusion lock for this device. It is used to serialize access to a device. The host adapter driver initializes this mutex. See <code>mutex(9F)</code>.</p> <p><code>sd_inq</code> is initially <code>NULL</code> (zero). After executing <code>scsi_probe(9F)</code>, this field contains the inquiry data associated with the particular device.</p> <p><code>sd_sense</code> is initially <code>NULL</code> (zero). If the target driver wants to use this field for storing <code>REQUEST SENSE</code> data, it should allocate an <code>scsi_extended_sense(9S)</code> buffer and set this field to the address of this buffer.</p> <p><code>sd_private</code> is reserved for the use of target drivers and should generally be used to point to target specific data structures.</p>
SEE ALSO	<p><code>detach(9E)</code>, <code>probe(9E)</code>, <code>makecom(9F)</code>, <code>mutex(9F)</code>, <code>scsi_probe(9F)</code>, <code>scsi_extended_sense(9S)</code>, <code>scsi_pkt(9S)</code></p> <p><i>Writing Device Drivers</i></p>

scsi_extended_sense(9S)

NAME	scsi_extended_sense – SCSI extended sense structure
SYNOPSIS	#include <sys/scsi/scsi.h>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	The <code>scsi_extended_sense</code> structure for error codes 0x70 (current errors) and 0x71 (deferred errors) is returned on a successful REQUEST SENSE command. SCSI-2 compliant targets are required to return at least the first 18 bytes of this structure. This structure is part of <code>scsi_device(9S)</code> structure.
STRUCTURE MEMBERS	<pre> uchar_t es_valid :1; /* Sense data is valid */ uchar_t es_class :3; /* Error Class- fixed at 0x7 */ uchar_t es_code :4; /* Vendor Unique error code */ uchar_t es_segnum; /* Segment number: for COPY cmd only */ uchar_t es_filmk :1; /* File Mark Detected */ uchar_t es_eom :1; /* End of Media */ uchar_t es_ili :1; /* Incorrect Length Indicator */ uchar_t es_key :4; /* Sense key */ uchar_t es_info_1; /* Information byte 1 */ uchar_t es_info_2; /* Information byte 2 */ uchar_t es_info_3; /* Information byte 3 */ uchar_t es_info_4; /* Information byte 4 */ uchar_t es_add_len; /* Number of additional bytes */ uchar_t es_cmd_info[4]; /* Command specific information */ uchar_t es_add_code; /* Additional Sense Code */ uchar_t es_qual_code; /* Additional Sense Code Qualifier */ uchar_t es_fru_code; /* Field Replaceable Unit Code */ uchar_t es_skey_specific[3]; /* Sense Key Specific information */ </pre> <p><code>es_valid</code>, if set, indicates that the information field contains valid information.</p> <p><code>es_class</code> should be 0x7.</p> <p><code>es_code</code> is either 0x0 or 0x1.</p> <p><code>es_segnum</code> contains the number of the current segment descriptor if the REQUEST SENSE command is in response to a COPY, COMPARE, and COPY AND VERIFY command.</p> <p><code>es_filmk</code>, if set, indicates that the current command had read a file mark or set mark (sequential access devices only).</p> <p><code>es_eom</code>, if set, indicates that an end-of-medium condition exists (sequential access and printer devices only).</p> <p><code>es_ili</code>, if set, indicates that the requested logical block length did not match the logical block length of the data on the medium.</p> <p><code>es_key</code> indicates generic information describing an error or exception condition. The following sense keys are defined:</p> <p><code>KEY_NO_SENSE</code> Indicates that there is no specific sense key information to be reported.</p>

KEY_RECOVERABLE_ERROR

Indicates that the last command completed successfully with some recovery action performed by the target.

KEY_NOT_READY

Indicates that the logical unit addressed cannot be accessed.

KEY_MEDIUM_ERROR

Indicates that the command terminated with a non-recovered error condition that was probably caused by a flaw on the medium or an error in the recorded data.

KEY_HARDWARE_ERROR

Indicates that the target detected a non-recoverable hardware failure while performing the command or during a self test.

KEY_ILLEGAL_REQUEST

Indicates that there was an illegal parameter in the CDB or in the additional parameters supplied as data for some commands.

KEY_UNIT_ATTENTION

Indicates that the removable medium might have been changed or the target has been reset.

KEY_WRITE_PROTECT/KEY_DATA_PROTECT

Indicates that a command that reads or writes the medium was attempted on a block that is protected from this operation.

KEY_BLANK_CHECK

Indicates that a write-once device or a sequential access device encountered blank medium or format-defined end-of-data indication while reading or a write-once device encountered a non-blank medium while writing.

KEY_VENDOR_UNIQUE

This sense key is available for reporting vendor-specific conditions.

KEY_COPY_ABORTED

Indicates that a COPY, COMPARE, and COPY AND VERIFY command was aborted.

KEY_ABORTED_COMMAND

Indicates that the target aborted the command.

KEY_EQUAL

Indicates that a SEARCH DATA command has satisfied an equal comparison.

KEY_VOLUME_OVERFLOW

Indicates that a buffered peripheral device has reached the end-of-partition and data might remain in the buffer that has not been written to the medium.

KEY_MISCOMPARE

Indicates that the source data did not match the data read from the medium.

KEY_RESERVE

Indicates that the target is currently reserved by a different initiator. `es_info_{1, 2, 3, 4}` is device-type or command specific.

scsi_extended_sense(9S)

`es_add_len` indicates the number of additional sense bytes to follow.

`es_cmd_info` contains information that depends on the command that was executed.

`es_add_code` (ASC) indicates further information related to the error or exception condition reported in the sense key field.

`es_qual_code` (ASCQ) indicates detailed information related to the additional sense code.

`es_fru_code` (FRU) indicates a device-specific mechanism to unit that has failed.

`es_key_specific` is defined when the value of the sense-key specific valid bit (bit 7) is 1. This field is reserved for sense keys not defined above.

SEE ALSO `scsi_device(9S)`

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

Writing Device Drivers

NAME	scsi_hba_tran – SCSI Host Bus Adapter (HBA) driver transport vector structure
SYNOPSIS	#include <sys/scsi/scsi.h>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
DESCRIPTION	A <code>scsi_hba_tran_t</code> structure defines vectors that an HBA driver exports to SCSSA interfaces so that HBA specific functions can be executed.
STRUCTURE MEMBERS	<pre> dev_info_t *tran_hba_dip; /* HBAs dev_info pointer */ void *tran_hba_private; /* HBA softstate */ void *tran_tgt_private; /* HBA target private pointer */ struct scsi_device *tran_sd; /* scsi_device */ int (*tran_tgt_init)(); /* Transport target */ int (*tran_tgt_probe)(); /* Initialization */ int (*tran_tgt_free)(); /* Transport target probe */ void (*tran_start)(); /* Transport target free */ int (*tran_reset)(); /* Transport start */ int (*tran_abort)(); /* Transport reset */ int (*tran_getcap)(); /* Transport abort */ int (*tran_setcap)(); /* Capability retrieval */ struct scsi_pkt (*tran_init_pkt)(); /* Capability establishment */ void (*tran_destroy_pkt)(); /* Packet and DMA allocation */ void (*tran_dmafree)(); /* Packet and DMA */ void (*tran_sync_pkt)(); /* deallocation */ void (*tran_reset_notify)(); /* DMA deallocation */ int (*tran_bus_reset)(); /* Sync DMA */ int (*tran_quiesce)(); /* Bus reset notification */ int (*tran_unquiesce)(); /* Reset bus only */ </pre> <p><code>tran_hba_dip</code> <code>dev_info</code> pointer to the HBA supplying the <code>scsi_hba_tran</code> structure.</p> <p><code>tran_hba_private</code> Private pointer that the HBA driver can use to refer to the device's soft state structure.</p> <p><code>tran_tgt_private</code> Private pointer that the HBA can use to refer to per-target specific data. This field can only be used when the <code>SCSI_HBA_TRAN_CLONE</code> flag is specified in <code>scsi_hba_attach(9F)</code>. In this case, the HBA driver must initialize this field in its <code>tran_tgt_init(9E)</code> entry point.</p> <p><code>tran_sd</code> Pointer to <code>scsi_device(9S)</code> structure if cloning; otherwise NULL.</p> <p><code>tran_tgt_init</code> The function entry allowing per-target HBA initialization, if necessary.</p> <p><code>tran_tgt_probe</code> The function entry allowing per-target <code>scsi_probe(9F)</code> customization, if necessary.</p> <p><code>tran_tgt_free</code> The function entry allowing per-target HBA deallocation, if necessary.</p>

scsi_hba_tran(9S)

<code>tran_start</code>	The function entry that starts a SCSI command execution on the HBA hardware.
<code>tran_reset</code>	The function entry that resets a SCSI bus or target device.
<code>tran_abort</code>	The function entry that aborts one SCSI command, or all pending SCSI commands.
<code>tran_getcap</code>	The function entry that retrieves a SCSI capability.
<code>tran_setcap</code>	The function entry that sets a SCSI capability.
<code>tran_init_pkt</code>	The function entry that allocates a <code>scsi_pkt</code> structure.
<code>tran_destroy_pkt</code>	The function entry that frees a <code>scsi_pkt</code> structure allocated by <code>tran_init_pkt</code> .
<code>tran_dmafree</code>	The function entry that frees DMA resources that were previously allocated by <code>tran_init_pkt</code> .
<code>tran_sync_pkt</code>	Synchronize data in <code>pkt</code> after a data transfer has been completed.
<code>tran_reset_notify</code>	The function entry allowing a target to register a bus reset notification request with the HBA driver.
<code>tran_bus_reset</code>	The function entry that resets the SCSI bus without resetting targets.
<code>tran_quiesce</code>	The function entry that waits for all outstanding commands to complete and blocks (or queues) any I/O requests issued.
<code>tran_unquiesce</code>	The function entry that allows I/O activities to resume on the SCSI bus.

SEE ALSO

`tran_abort(9E)`, `tran_bus_reset(9E)`, `tran_destroy_pkt(9E)`, `tran_dmafree(9E)`, `tran_getcap(9E)`, `tran_init_pkt(9E)`, `tran_quiesce(9E)`, `tran_reset(9E)`, `tran_reset_notify(9E)`, `tran_setcap(9E)`, `tran_start(9E)`, `tran_sync_pkt(9E)`, `tran_tgt_free(9E)`, `tran_tgt_init(9E)`, `tran_tgt_probe(9E)`, `tran_unquiesce(9E)`, `ddi_dma_sync(9F)`, `scsi_hba_attach(9F)`, `scsi_hba_pkt_alloc(9F)`, `scsi_hba_pkt_free(9F)`, `scsi_probe(9F)`, `scsi_device(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

NAME	scsi_inquiry – SCSI inquiry structure																						
SYNOPSIS	#include <sys/scsi/scsi.h>																						
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).																						
DESCRIPTION	The <code>scsi_inquiry</code> structure contains 36 required bytes, followed by a variable number of vendor-specific parameters. Bytes 59 through 95, if returned, are reserved for future standardization. This structure is part of <code>scsi_device(9S)</code> structure and typically filled in by <code>scsi_probe(9F)</code> .																						
STRUCTURE MEMBERS	<pre> uchar_t inq_dtype; /* Peripheral qualifier, device type */ uchar_t inq_rmb :1; /* Removable media */ uchar_t inq_qual :7; /* Device type qualifier */ uchar_t inq_iso :2; /* ISO version */ uchar_t inq_ecma :3; /* ANSI version */ uchar_t inq_aenc :1; /* Async event notification cap. */ uchar_t inq_trmiop :1; /* Supports TERMINATE I/O PROC msg */ uchar_t inq_rdf :4; /* Response data format */ uchar_t inq_len; /* Additional length */ uchar_t inq_reladdr :1; /* Supports relative addressing */ uchar_t inq_wbus32 :1; /* Supports 32 bit wide data xfers */ uchar_t inq_wbus16 :1; /* Supports 16 bit wide data xfers */ uchar_t inq_sync :1; /* Supports synchronous data xfers */ uchar_t inq_linked :1; /* Supports linked commands */ uchar_t inq_cmd_que :1; /* Supports command queueing */ uchar_t inq_sftre :1; /* Supports Soft Reset option */ char inq_vid[8]; /* Vendor ID */ char inq_pid[16]; /* Product ID */ char inq_revision[4]; /* Revision level */ </pre> <p><code>inq_dtype</code> identifies the type of device. Bits 0 - 4 represent the Peripheral Device Type and bits 5 - 7 represent the Peripheral Qualifier. The following values are appropriate for Peripheral Device Type field:</p> <table border="0"> <tr> <td><code>DTYPE_ARRAY_CTRL</code></td> <td>Array controller device (for example, RAID).</td> </tr> <tr> <td><code>DTYPE_DIRECT</code></td> <td>Direct-access device (for example, magnetic disk).</td> </tr> <tr> <td><code>DTYPE_ESI</code></td> <td>Enclosure services device.</td> </tr> <tr> <td><code>DTYPE_SEQUENTIAL</code></td> <td>Sequential-access device (for example, magnetic tape).</td> </tr> <tr> <td><code>DTYPE_PRINTER</code></td> <td>Printer device.</td> </tr> <tr> <td><code>DTYPE_PROCESSOR</code></td> <td>Processor device.</td> </tr> <tr> <td><code>DTYPE_WORM</code></td> <td>Write-once device (for example, some optical disks).</td> </tr> <tr> <td><code>DTYPE_RODIRECT</code></td> <td>CD-ROM device.</td> </tr> <tr> <td><code>DTYPE_SCANNER</code></td> <td>Scanner device.</td> </tr> <tr> <td><code>DTYPE_OPTICAL</code></td> <td>Optical memory device (for example, some optical disks).</td> </tr> <tr> <td><code>DTYPE_CHANGER</code></td> <td>Medium Changer device (for example, jukeboxes).</td> </tr> </table>	<code>DTYPE_ARRAY_CTRL</code>	Array controller device (for example, RAID).	<code>DTYPE_DIRECT</code>	Direct-access device (for example, magnetic disk).	<code>DTYPE_ESI</code>	Enclosure services device.	<code>DTYPE_SEQUENTIAL</code>	Sequential-access device (for example, magnetic tape).	<code>DTYPE_PRINTER</code>	Printer device.	<code>DTYPE_PROCESSOR</code>	Processor device.	<code>DTYPE_WORM</code>	Write-once device (for example, some optical disks).	<code>DTYPE_RODIRECT</code>	CD-ROM device.	<code>DTYPE_SCANNER</code>	Scanner device.	<code>DTYPE_OPTICAL</code>	Optical memory device (for example, some optical disks).	<code>DTYPE_CHANGER</code>	Medium Changer device (for example, jukeboxes).
<code>DTYPE_ARRAY_CTRL</code>	Array controller device (for example, RAID).																						
<code>DTYPE_DIRECT</code>	Direct-access device (for example, magnetic disk).																						
<code>DTYPE_ESI</code>	Enclosure services device.																						
<code>DTYPE_SEQUENTIAL</code>	Sequential-access device (for example, magnetic tape).																						
<code>DTYPE_PRINTER</code>	Printer device.																						
<code>DTYPE_PROCESSOR</code>	Processor device.																						
<code>DTYPE_WORM</code>	Write-once device (for example, some optical disks).																						
<code>DTYPE_RODIRECT</code>	CD-ROM device.																						
<code>DTYPE_SCANNER</code>	Scanner device.																						
<code>DTYPE_OPTICAL</code>	Optical memory device (for example, some optical disks).																						
<code>DTYPE_CHANGER</code>	Medium Changer device (for example, jukeboxes).																						

scsi_inquiry(9S)

DTYPE_COMM	Communications device.
DTYPE_UNKNOWN	Unknown or no device type.
DTYPE_MASK	Mask to isolate Peripheral Device Type field.

The following values are appropriate for the Peripheral Qualifier field:

DPQ_POSSIBLE	The specified peripheral device type is currently connected to this logical unit. If the target cannot determine whether or not a physical device is currently connected, it uses this peripheral qualifier when returning the INQUIRY data. This peripheral qualifier does not imply that the device is ready for access by the initiator.
DPQ_SUPPORTED	The target is capable of supporting the specified peripheral device type on this logical unit. However, the physical device is not currently connected to this logical unit.
DPQ_NEVER	The target is not capable of supporting a physical device on this logical unit. For this peripheral qualifier, the peripheral device type shall be set to DTYPE_UNKNOWN to provide compatibility with previous versions of SCSI. For all other peripheral device type values, this peripheral qualifier is reserved.
DPQ_VUNIQ	This is a vendor-unique qualifier.

DTYPE_NOTPRESENT is the peripheral qualifier DPQ_NEVER and the peripheral device type DTYPE_UNKNOWN combined.

inq_rmb, if set, indicates that the medium is removable.

inq_qual is a device type qualifier.

inq_iso indicates ISO version.

inq_ecma indicates ECMA version.

inq_ansi indicates ANSI version.

inq_aenc, if set, indicates that the device supports asynchronous event notification capability as defined in SCSI-2 specification.

inq_trmiop, if set, indicates that the device supports the TERMINATE I/O PROCESS message.

inq_rdf, if reset, indicates the INQUIRY data format is as specified in SCSI-1.

`inq_inq_len` is the additional length field that specifies the length in bytes of the parameters.

`inq_reladdr`, if set, indicates that the device supports the relative addressing mode of this logical unit.

`inq_wbus32`, if set, indicates that the device supports 32-bit wide data transfers.

`inq_wbus16`, if set, indicates that the device supports 16-bit wide data transfers.

`inq_sync`, if set, indicates that the device supports synchronous data transfers.

`inq_linked`, if set, indicates that the device supports linked commands for this logical unit.

`inq_cmdque`, if set, indicates that the device supports tagged command queuing.

`inq_sftre`, if reset, indicates that the device responds to the RESET condition with the hard RESET alternative. If this bit is set, this indicates that the device responds with the soft RESET alternative.

`inq_vid` contains eight bytes of ASCII data identifying the vendor of the product.

`inq_pid` contains sixteen bytes of ASCII data as defined by the vendor.

`inq_revision` contains four bytes of ASCII data as defined by the vendor.

SEE ALSO `scsi_probe(9F)`, `scsi_device(9S)`

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

Writing Device Drivers

scsi_pkt(9S)

NAME	scsi_pkt – SCSI packet structure
SYNOPSIS	#include <sys/scsi/scsi.h>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
DESCRIPTION	A <code>scsi_pkt</code> structure defines the packet that is allocated by <code>scsi_init_pkt(9F)</code> . The target driver fills in some information, and passes it to <code>scsi_transport(9F)</code> for execution on the target. The host bus adapter (HBA) fills in some other information as the command is processed. When the command completes (or can be taken no further) the completion function specified in the packet is called, with a pointer to the packet as its argument. From fields within the packet, the target driver can determine the success or failure of the command.
STRUCTURE MEMBERS	<pre>opaque_t pkt_ha_private; /* private data for host adapter */ struct scsi_address pkt_address; /* destination packet */ opaque_t pkt_private; /* private data for target driver */ void (*pkt_comp)(struct scsi_pkt *); /* callback */ uint_t pkt_flags; /* flags */ int pkt_time; /* time allotted to complete command */ uchar_t *pkt_scbp; /* pointer to status block */ uchar_t *pkt_cdbp; /* pointer to command block */ ssize_t pkt_resid; /* number of bytes not transferred */ uint_t pkt_state; /* state of command */ uint_t pkt_statistics; /* statistics */ uchar_t pkt_reason; /* reason completion called */</pre> <p><code>pkt_ha_private</code> An opaque pointer that the Host Bus Adapter uses to reference a private data structure used to transfer <code>scsi_pkt</code> requests.</p> <p><code>pkt_address</code> Initialized by <code>scsi_init_pkt(9F)</code>; <code>pkt_address</code> records the intended route and recipient of a request.</p> <p><code>pkt_private</code> Reserved for the use of the target driver; <code>pkt_private</code> is not changed by the HBA driver.</p> <p><code>pkt_comp</code> Specifies the command completion callback routine. When the host adapter driver has gone as far as it can in transporting a command to a SCSI target, and the command has either run to completion or can go no further for some other reason, the host adapter driver</p>

scsi_pkt(9S)

	will call the function pointed to by this field and pass a pointer to the packet as argument. The callback routine itself is called from interrupt context and must not sleep or call any function that might sleep.
<code>pkt_flags</code>	Provides additional information about how the target driver expects the command to be executed. See <code>pkt_flag</code> Definitions.
<code>pkt_time</code>	Will be set by the target driver to represent the maximum time in seconds that this command is allowed to take to complete. Timeout starts when the command is transmitted on the SCSI bus. <code>pkt_time</code> may be 0 if no timeout is required.
<code>pkt_scbp</code>	Points to either a struct <code>scsi_status(9S)</code> or, if <code>auto-rqsense</code> is enabled, and <code>pkt_state</code> includes <code>STATE_ARQ_DONE</code> , a struct <code>scsi_arq_status</code> . If <code>scsi_status</code> is returned, the SCSI status byte resulting from the requested command is available; if <code>scsi_arq_status(9S)</code> is returned, the sense information is also available.
<code>pkt_cdbp</code>	Points to a kernel-addressable buffer whose length was specified by a call to the proper resource allocation routine, <code>scsi_init_pkt(9F)</code> .
<code>pkt_resid</code>	Contains a residual count, either the number of data bytes that have not been transferred (<code>scsi_transport(9F)</code>) or the number of data bytes for which DMA resources could not be allocated <code>scsi_init_pkt(9F)</code> . In the latter case, partial DMA resources may only be allocated if <code>scsi_init_pkt(9F)</code> is called with the <code>PKT_DMA_PARTIAL</code> flag.
<code>pkt_state</code>	Has bit positions that represent the six most important states that a SCSI command can go through (see <code>pkt_state</code> Definitions).
<code>pkt_statistics</code>	Maintains some transport-related statistics. (see <code>pkt_statistics</code> Definitions).
<code>pkt_reason</code>	Contains a completion code that indicates why the <code>pkt_comp</code> function was called. See <code>pkt_reason</code> Definitions, below.
	The host adapter driver will update the <code>pkt_resid</code> , <code>pkt_reason</code> , <code>pkt_state</code> , and <code>pkt_statistics</code> fields.

pkt_flags
Definitions:

The appropriate definitions for the structure member `pkt_flags` are:

scsi_pkt(9S)

FLAG_NOINTR	Run command with no command completion callback; command is complete upon return from <code>scsi_transport(9F)</code> .
FLAG_NODISCON	Run command without disconnects.
FLAG_NOPARITY	Run command without parity checking.
FLAG_HTAG	Run command as the head-of-queue-tagged command.
FLAG_OTAG	Run command as an ordered-queue-tagged command.
FLAG_STAG	Run command as a simple-queue —tagged command.
FLAG_SENSING	Indicates command is a request sense command.
FLAG_HEAD	Place command at the head of the queue.
FLAG_RENEGOTIATE_WIDE_SYNC	Before transporting this command, the host adapter should initiate the renegotiation of wide mode and synchronous transfer speed. Normally the HBA driver manages negotiations but under certain conditions forcing a renegotiation is appropriate. Renegotiation is recommended before Request Sense and Inquiry commands. (Refer to the SCSI 2 standard, sections 6.6.21 and 6.6.23.) This flag should not be set for every packet as this will severely impact performance.

pkt_reason Definitions:

The appropriate definitions for the structure member `pkt_reason` are:

CMD_CMPLT	No transport errors; normal completion.
CMD_INCOMPLETE	Transport stopped with abnormal state.
CMD_DMA_DERR	DMA direction error.
CMD_TRAN_ERR	Unspecified transport error.
CMD_RESET	SCSI bus reset destroyed command.
CMD_ABORTED	Command transport aborted on request.
CMD_TIMEOUT	Command timed out.
CMD_DATA_OVR	Data overrun.
CMD_CMD_OVR	Command overrun.
CMD_STS_OVR	Status overrun.
CMD_BADMSG	Message not command complete.

CMD_NOMSGOUT	Target refused to go to message out phase.
CMD_XID_FAIL	Extended identify message rejected.
CMD_IDE_FAIL	“Initiator Detected Error” message rejected.
CMD_ABORT_FAIL	Abort message rejected.
CMD_REJECT_FAIL	Reject message rejected.
CMD_NOP_FAIL	“No Operation” message rejected.
CMD_PER_FAIL	“Message Parity Error” message rejected.
CMD_BDR_FAIL	“Bus Device Reset” message rejected.
CMD_ID_FAIL	Identify message rejected.
CMD_UNX_BUS_FREE	Unexpected bus free phase.
CMD_TAG_REJECT	Target rejected the tag message.

**pkt_state
Definitions:**

The appropriate definitions for the structure member `pkt_state` are:

STATE_GOT_BUS	Bus arbitration succeeded.
STATE_GOT_TARGET	Target successfully selected.
STATE_SENT_CMD	Command successfully sent.
STATE_XFERRED_DATA	Data transfer took place.
STATE_GOT_STATUS	Status received.
STATE_ARQ_DONE	The command resulted in a check condition and the host adapter driver executed an automatic request sense command.

**pkt_statistics
Definitions:**

The definitions that are appropriate for the structure member `pkt_statistics` are:

STAT_DISCON	Device disconnect.
STAT_SYNC	Command did a synchronous data transfer.
STAT_PERR	SCSI parity error.
STAT_BUS_RESET	Bus reset.
STAT_DEV_RESET	Device reset.
STAT_ABORTED	Command was aborted.
STAT_TIMEOUT	Command timed out.

SEE ALSO

`tran_init_pkt(9E)`, `scsi_arq_status(9S)`, `scsi_init_pkt(9F)`,
`scsi_transport(9F)`, `scsi_status(9S)`

Writing Device Drivers

scsi_status(9S)

NAME	scsi_status – SCSI status structure												
SYNOPSIS	#include <sys/scsi/scsi.h>												
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI)												
DESCRIPTION	The SCSI-2 standard defines a status byte that is normally sent by the target to the initiator during the status phase at the completion of each command.												
STRUCTURE MEMBERS	<pre>uchar sts_scsi2 :1; /* SCSI-2 modifier bit */ uchar sts_is :1; /* intermediate status sent */ uchar sts_busy :1; /* device busy or reserved */ uchar sts_cm :1; /* condition met */ uchar sts_chk :1; /* check condition */</pre> <p>sts_chk indicates that a contingent allegiance condition has occurred.</p> <p>sts_cm is returned whenever the requested operation is satisfied</p> <p>sts_busy indicates that the target is busy. This status is returned whenever a target is unable to accept a command from an otherwise acceptable initiator (that is, no reservation conflicts). The recommended initiator recovery action is to issue the command again later.</p> <p>sts_is is returned for every successfully completed command in a series of linked commands (except the last command), unless the command is terminated with a check condition status, reservation conflict, or command terminated status. Note that host bus adapter drivers may not support linked commands (see <code>scsi_ifsetcap(9F)</code>). If <code>sts_is</code> and <code>sts_busy</code> are both set, then a reservation conflict has occurred.</p> <p>sts_scsi2 is the SCSI-2 modifier bit. If <code>sts_scsi2</code> and <code>sts_chk</code> are both set, this indicates a command terminated status. If <code>sts_scsi2</code> and <code>sts_busy</code> are both set, this indicates that the command queue in the target is full.</p> <p>For accessing the status as a byte, the following values are appropriate:</p> <table><tr><td>STATUS_GOOD</td><td>This status indicates that the target has successfully completed the command.</td></tr><tr><td>STATUS_CHECK</td><td>This status indicates that a contingent allegiance condition has occurred.</td></tr><tr><td>STATUS_MET</td><td>This status is returned when the requested operations are satisfied.</td></tr><tr><td>STATUS_BUSY</td><td>This status indicates that the target is busy.</td></tr><tr><td>STATUS_INTERMEDIATE</td><td>This status is returned for every successfully completed command in a series of linked commands.</td></tr><tr><td>STATUS_SCSI2</td><td>This is the SCSI-2 modifier bit.</td></tr></table>	STATUS_GOOD	This status indicates that the target has successfully completed the command.	STATUS_CHECK	This status indicates that a contingent allegiance condition has occurred.	STATUS_MET	This status is returned when the requested operations are satisfied.	STATUS_BUSY	This status indicates that the target is busy.	STATUS_INTERMEDIATE	This status is returned for every successfully completed command in a series of linked commands.	STATUS_SCSI2	This is the SCSI-2 modifier bit.
STATUS_GOOD	This status indicates that the target has successfully completed the command.												
STATUS_CHECK	This status indicates that a contingent allegiance condition has occurred.												
STATUS_MET	This status is returned when the requested operations are satisfied.												
STATUS_BUSY	This status indicates that the target is busy.												
STATUS_INTERMEDIATE	This status is returned for every successfully completed command in a series of linked commands.												
STATUS_SCSI2	This is the SCSI-2 modifier bit.												

scsi_status(9S)

STATUS_INTERMEDIATE_MET	This status is a combination of STATUS_MET and STATUS_INTERMEDIATE.
STATUS_RESERVATION_CONFLICT	This status is a combination of STATUS_INTERMEDIATE and STATUS_BUSY, and it is returned whenever an initiator attempts to access a logical unit or an extent within a logical unit is reserved.
STATUS_TERMINATED	This status is a combination of STATUS_SCSI2 and STATUS_CHECK, and it is returned whenever the target terminates the current I/O process after receiving a terminate I/O process message.
STATUS_QFULL	This status is a combination of STATUS_SCSI2 and STATUS_BUSY, and it is returned when the command queue in the target is full.

SEE ALSO scsi_ifgetcap(9F), scsi_init_pkt(9F), scsi_extended_sense(9S), scsi_pkt(9S)

Writing Device Drivers

streamtab(9S)

NAME	streamtab – STREAMS entity declaration structure
SYNOPSIS	<pre>#include <sys/stream.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI).
DESCRIPTION	<p>Each STREAMS driver or module must have a <code>streamtab</code> structure.</p> <p><code>streamtab</code> is made up of <code>qinit</code> structures for both the read and write queue portions of each module or driver. Multiplexing drivers require both upper and lower <code>qinit</code> structures. The <code>qinit</code> structure contains the entry points through which the module or driver routines are called.</p> <p>Normally, the read <code>QUEUE</code> contains the <code>open</code> and <code>close</code> routines. Both the read and write queue can contain <code>put</code> and service procedures.</p>
STRUCTURE MEMBERS	<pre>struct qinit *st_rdinit; /* read QUEUE */ struct qinit *st_wrinit; /* write QUEUE */ struct qinit *st_muxrinit; /* lower read QUEUE*/ struct qinit *st_muxwinit; /* lower write QUEUE*/</pre>
SEE ALSO	<p><code>qinit(9S)</code></p> <p><i>STREAMS Programming Guide</i></p>

NAME	stroptions – options structure for M_SETOPTS message
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/stropts.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	The M_SETOPTS message contains a stroptions structure and is used to control options in the stream head.
STRUCTURE MEMBERS	<pre>uint_t so_flags; /* options to set */ short so_readopt; /* read option */ ushort_t so_wroff; /* write offset */ ssize_t so_minpsz; /* minimum read packet size */ ssize_t so_maxpsz; /* maximum read packet size */ size_t so_hiwat; /* read queue high water mark */ size_t so_lowat; /* read queue low water mark */ unsigned char so_band; /* band for water marks */ ushort_t so_errropt; /* error option */</pre> <p>The following are the flags that can be set in the <code>so_flags</code> bit mask in the <code>stroptions</code> structure. Note that multiple flags can be set.</p> <pre>SO_READOPT Set read option. SO_WROFF Set write offset. SO_MINPSZ Set minimum packet size SO_MAXPSZ Set maximum packet size. SO_HIWAT Set high water mark. SO_LOWAT Set low water mark. SO_MREADON Set read notification ON. SO_MREADOFF Set read notification OFF. SO_NDELO Old TTY semantics for NDELAY reads and writes. SO_NDELOFFSTREAMS Semantics for NDELAY reads and writes. SO_ISTTY The stream is acting as a terminal. SO_ISNTTY The stream is not acting as a terminal. SO_TOSTOP Stop on background writes to this stream. SO_TONSTOP Do not stop on background writes to this stream. SO_BAND Water marks affect band. SO_ERROPT Set error option.</pre>

stroptions(9S)

When `SO_READOPT` is set, the `so_readopt` field of the `stroptions` structure can take one of the following values. See `read(2)`.

<code>RNORM</code>	Read message normal.
<code>RMSGD</code>	Read message discard.
<code>RMSGN</code>	Read message, no discard.

When `SO_BAND` is set, `so_band` determines to which band `so_hiwat` and `so_lowat` apply.

When `SO_ERROPT` is set, the `so_erropt` field of the `stroptions` structure can take a value that is either none or one of:

<code>RERRNORM</code>	Persistent read errors; default.
-----------------------	----------------------------------

<code>RERRNONPERSIST</code>	Non-persistent read errors.
-----------------------------	-----------------------------

OR'ed with either none or one of:

<code>WERRNORM</code>	Persistent write errors; default.
-----------------------	-----------------------------------

<code>WERRNONPERSIST</code>	Non-persistent write errors.
-----------------------------	------------------------------

SEE ALSO `read(2)`, `streamio(7I)`

STREAMS Programming Guide

NAME	tuple – card information structure (CIS) access structure
SYNOPSIS	<code>#include <sys/pccard.h></code>
INTERFACE LEVEL	Solaris DDI Specific (Solaris DDI)
DESCRIPTION	<p>The <code>tuple_t</code> structure is the basic data structure provided by card services to manage PC card information. A PC card provides identification and configuration information through its card information structure (CIS). A PC card driver accesses a PC card's CIS through various card services functions.</p> <p>The CIS information allows PC cards to be self-identifying: the CIS provides information to the system so that it can identify the proper PC card driver for the PC card, and provides configuration information so that the driver can allocate appropriate resources to configure the PC card for proper operation in the system.</p> <p>The CIS information is contained on the PC card in a linked list of tuple data structures called a CIS chain. Each tuple has a one-byte type and a one-byte link, an offset to the next tuple in the list. A PC card can have one or more CIS chains.</p> <p>A multi-function PC card that complies with the PC Card 95 MultiFunction Metaformat specification will have one or more global CIS chains that collectively are referred to as the global CIS. These PC Cards will also have one or more per-function CIS chains. Each per-function collection of CIS chains is referred to as a function-specific CIS.</p> <p>To examine a PC card's CIS, first a PC card driver must locate the desired tuple by calling <code>csx_GetFirstTuple(9F)</code>. Once the first tuple is located, subsequent tuples may be located by calling <code>csx_GetNextTuple(9F)</code>. See <code>csx_GetFirstTuple(9F)</code>. The linked list of tuples may be inspected one by one, or the driver may narrow the search by requesting only tuples of a particular type.</p> <p>Once a tuple has been located, the PC card driver may inspect the tuple data. The most convenient way to do this for standard tuples is by calling one of the number of tuple-parsing utility functions; for custom tuples, the driver may get access to the raw tuple data by calling <code>csx_GetTupleData(9F)</code>.</p> <p>Solaris PC card drivers do not need to be concerned with which CIS chain a tuple appears in. On a multi-function PC card, the client will get the tuples from the global CIS followed by the tuples in the function-specific CIS. The caller will not get any tuples from a function-specific CIS that does not belong to the caller's function.</p>
STRUCTURE MEMBERS	<p>The structure members of <code>tuple_t</code> are:</p> <pre> uint32_t Socket; /* socket number */ uint32_t Attributes; /* tuple attributes */ cisdata_t DesiredTuple; /* tuple to search for */ cisdata_t TupleOffset; /* tuple data offset */ cisdata_t TupleDataMax; /* max tuple data size */ cisdata_t TupleDataLen; /* actual tuple data length */ cisdata_t TupleData[CIS_MAX_TUPLE_DATA_LEN]; /* body tuple data */ </pre>

tuple(9S)

```
cisdata_t    TupleCode;    /* tuple type code */
cisdata_t    TupleLink;    /* tuple link */
```

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. TUPLE_RETURN_IGNORED_TUPLES Return ignored tuples if set. Ignored tuples are those tuples in a multi-function PC card's global CIS chain that are duplicates of the same tuples in a function-specific CIS chain. TUPLE_RETURN_NAME Return tuple name string using the <code>csx_ParseTuple(9F)</code> function if set.
DesiredTuple	This field is the requested tuple type code to be returned when calling <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code> . RETURN_FIRST_TUPLE is used to return the first tuple regardless of tuple type. RETURN_NEXT_TUPLE is used to return the next tuple regardless of tuple type.
TupleOffset	This field allows partial tuple information to be retrieved, starting at the specified offset within the tuple. This field must only be set before calling <code>csx_GetTupleData(9F)</code> .
TupleDataMax	This field is the size of the tuple data buffer that card services uses to return raw tuple data from <code>csx_GetTupleData(9F)</code> . It can be larger than the number of bytes in the tuple data body. Card services ignores any value placed here by the client.
TupleDataLen	This field is the actual size of the tuple data body. It represents the number of tuple data body bytes returned by <code>csx_GetTupleData(9F)</code> .
TupleData	This field is an array of bytes containing the raw tuple data body contents returned by <code>csx_GetTupleData(9F)</code> .
TupleCode	This field is the tuple type code and is returned by <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code> when a tuple matching the <code>DesiredTuple</code> field is returned.
TupleLink	This field is the tuple link, the offset to the next tuple, and is returned by <code>csx_GetFirstTuple(9F)</code> or <code>csx_GetNextTuple(9F)</code> when a tuple matching the <code>DesiredTuple</code> field is returned.

tuple(9S)

SEE ALSO | csx_GetFirstTuple(9F), csx_GetTupleData(9F), csx_ParseTuple(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)

PC Card 95 Standard, PCMCIA/JEIDA

uio(9S)

NAME	uio – scatter/gather I/O request structure
SYNOPSIS	#include <sys/uio.h>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI)
DESCRIPTION	<p>A uio structure describes an I/O request that can be broken up into different data storage areas (scatter/gather I/O). A request is a list of iovec structures (base-length pairs) indicating where in user space or kernel space the I/O data is to be read or written.</p> <p>The contents of uio structures passed to the driver through the entry points should not be written by the driver. The uiomove(9F) function takes care of all overhead related to maintaining the state of the uio structure.</p> <p>uio structures allocated by the driver should be initialized to zero before use, by bzero(9F), kmem_zalloc(9F), or an equivalent.</p>
STRUCTURE MEMBERS	<pre>iovec_t *uio_iov; /* pointer to the start of the iovec */ /* list for the uio structure */ int uio_iovcnt; /* the number of iovecs in the list */ off_t uio_offset; /* 32-bit offset into file where data is */ /* transferred from or to. See NOTES. */ offset_t uio_loffset; /* 64-bit offset into file where data is */ /* transferred from or to. See NOTES. */ uio_seg_t uio_segflg; /* identifies the type of I/O transfer: */ /* UIO_SYSSPACE: kernel <-> kernel */ /* UIO_USERSPACE: kernel <-> user */ short uio_fmode; /* file mode flags (not driver settable) */ daddr_t uio_limit; /* 32-bit ulimit for file (maximum block */ /* offset). not driver settable. See NOTES. */ diskaddr_t uio_llimit; /* 64-bit ulimit for file (maximum block */ /* offset). not driver settable. See NOTES. */ int uio_resid; /* residual count */</pre>
	<p>The uio_iov member is a pointer to the beginning of the iovec(9S) list for the uio. When the uio structure is passed to the driver through an entry point, the driver should not set uio_iov. When the uio structure is created by the driver, uio_iov should be initialized by the driver and not written to afterward.</p>
SEE ALSO	aread(9E), awrite(9E), read(9E), write(9E), bzero(9F), kmem_zalloc(9F), uiomove(9F), cb_ops(9S), iovec(9S)
	<i>Writing Device Drivers</i>
NOTES	<p>Only one structure, uio_offset or uio_loffset, should be interpreted by the driver. Which field the driver interprets is dependent upon the settings in the cb_ops(9S) structure.</p> <p>Only one structure, uio_limit or uio_llimit, should be interpreted by the driver. Which field the driver interprets is dependent upon the settings in the cb_ops(9S) structure.</p>

When performing I/O on a seekable device, the driver should not modify either the `uio_offset` or the `uio_loffset` field of the `uio` structure. I/O to such a device is constrained by the maximum offset value. When performing I/O on a device on which the concept of position has no relevance, the driver may preserve the `uio_offset` or `uio_loffset`, perform the I/O operation, then restore the `uio_offset` or `uio_loffset` to the field's initial value. I/O performed to a device in this manner is not constrained.

uio(9S)

Index

A

aio_req — asynchronous I/O request structure, 18
asynchronous I/O request structure — aio_req, 18

B

buf — block I/O data transfer structure, 19

C

Card Information Structure (CIS) access structure — tuple, 107
character/block entry points structure for drivers, — cb_ops, 22
copyreq — STREAMS data structure for the M_COPYIN and the M_COPYOUT message types, 24
copyresp — STREAMS data structure for the M_IOCTLDATA message type, 25

D

data access attributes structure — ddi_device_acc_attr, 27
ddi_device_acc_attr — data access attributes structure, 27

DDI device mapping

ddi_mapdev_ctl — device mapping-control structure, 48
devmap_callback_ctl — device mapping-control structure, 49
DDI direct memory access
DMA limits structure — ddi_dma_lim, 39, 41
DMA cookie structure — ddi_dma_cookie, 34
DMA Request structure — ddi_dma_req, 43
ddi_dma_attr — DMA attributes structure, 31
ddi_dmae_req — DMA engine request structure, 35
ddi_idevice_cookie — device interrupt cookie, 47
ddi_mapdev_ctl — device mapping-control structure, 48
device interrupt cookie — ddi_idevice_cookie, 47
device mapping-control structure — ddi_mapdev_ctl, 48
device mapping-control structure — devmap_callback_ctl, 49
device operations structure, — dev_ops, 51
devmap_callback_ctl — device mapping-control structure, 49
DMA attributes structure — ddi_dma_attr, 31
DMA cookie structure, — ddi_dma_cookie, 34
DMA engine request structure — ddi_dmae_req, 35
DMA limits structure — ddi_dma_lim, 39, 41

DMA Request structure, — `ddi_dma_req`, 43
driver's message-freeing routine, —
 `free_rtn`, 53
drivers, loadable, linkage structure, —
 `modldrv`, 69

F

`fmodsw` — STREAMS module declaration
structure, 52

G

`gld_mac_info` — GLD mac info data
structure, 54
`gld_stats` — GLD statistics data structure, 57

I

I/O, block, data transfer structure, — `buf`, 19
I/O data storage structure using `uio`, —
 `iovec`, 61
I/O request structure, scatter/gather, —
 `uio`, 110
`iocblk` — STREAMS data structure for the
 `M_IOCTL` message type, 60

K

kernel statistics structure — `kstat`, 62
`kstat` — kernel statistics structure, 62
`kstat_intr` — structure for interrupt kstats, 64
`kstat_io` — structure for I/O kstats, 66
`kstat_named` — structure for named kstats, 67

L

`linkblk` — STREAMS data structure sent to
multiplexor drivers to indicate a link, 68

M

`modlinkage` — module linkage structure, 70

O

options structure for `M_SETOPTS` message —
`stroptions`, 105

P

`pm-components`— Power Management device
property, 78
Power Management device property —
 `pm-component`, 78

Q

`queclass` — a STREAMS macro that returns the
queue message class definitions for a given
message block, 83

S

`scsi_address` — SCSI address structure, 86
SCSI address structure — `scsi_address`, 86
`scsi_arq_status` — SCSI auto request sense
structure, 87
SCSI ASC ASCQ to message structure,
 `scsi-vu-errmsg`, 88
`scsi_asc_key_strings`, SCSI ASC ASCQ to
message structure, 88
SCSI auto request sense structure —
 `scsi_arq_status`, 87
`scsi_device` — SCSI device structure, 89
SCSI device structure — `scsi_device`, 89
SCSI device structure — `scsi_inquiry`, 95
`scsi_extended_sense` — SCSI extended sense
structure, 90
SCSI extended sense structure —
 `scsi_extended_sense`, 90
`scsi_hba_tran` — SCSI Host Bus Adapter (HBA)
driver transport vector structure, 93
SCSI Host Bus Adapter (HBA) driver transport
vector structure — `scsi_hba_tran`, 93

scsi_inquiry — SCSI device structure, 95
 SCSI packet structure — scsi_pkt, 98
 scsi_pkt — SCSI packet structure, 98
 pkt_flags Definitions, 99
 pkt_reason Definitions, 100
 pkt_state Definitions, 101
 pkt_statistics Definitions, 101
 scsi_status — SCSI status structure, 102
 SCSI status structure — scsi_status, 102
 STREAMS data structure for the M_COPYIN
 and the M_COPYOUT message types —
 copyreq, 24
 STREAMS data structure for the M_IOCTLDATA
 message type — copyresp, 25
 STREAMS data structure for the M_IOCTL
 message type — iocblk, 60
 STREAMS data structure sent to multiplexor
 drivers to indicate a link — linkblk, 68
 STREAMS driver identification and limit value
 structure, — module_info, 72
 STREAMS entity declaration structure, —
 streamtab, 104
 STREAMS macro that returns the queue
 message class definitions for a given message
 block — queclass, 83
 STREAMS message block structure, —
 msgb, 73
 STREAMS message data structure, — datab, 26
 STREAMS module declaration structure —
 fmodsw, 52
 STREAMS modules, loadable, linkage structure,
 modlstrmod, 71
 STREAMS queue flow control information
 structure, — qband, 81
 STREAMS queue processing procedures
 structure, — qinit, 82
 STREAMS queue structure, — queue, 84
 stroptions — options structure for M_SETOPTS
 message, 105
 structure for I/O kstats — kstat_io, 66
 structure for interrupt kstats — kstat_intr, 64
 structure for named kstats — kstat_named, 67

T

tuple — Card Information Structure (CIS) access
 structure, 107

U

uio — scatter/gather I/O request structure, 110

