



man pages section 9: DDI and DKI Driver Entry Points

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

Part No: 817-0670-10
April 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.



030212@5115



Contents

Preface 7

Introduction 13

Intro(9E) 14

Driver Entry Points 21

aread(9E) 22

attach(9E) 24

awrite(9E) 27

chpoll(9E) 29

close(9E) 31

csx_event_handler(9E) 34

detach(9E) 40

devmap(9E) 43

devmap_access(9E) 47

devmap_contextmgt(9E) 50

devmap_dup(9E) 53

devmap_map(9E) 55

devmap_unmap(9E) 57

dump(9E) 60

_fini(9E) 61

getinfo(9E) 64

gld(9E) 66

gldm_get_stats(9E) 71

gldm_intr(9E) 76

gldm_ioctl(9E) 81
gldm_reset(9E) 86
gldm_send(9E) 91
gldm_set_mac_addr(9E) 96
gldm_set_multicast(9E) 101
gldm_set_promiscuous(9E) 106
gldm_start(9E) 111
gldm_stop(9E) 116
identify(9E) 121
_info(9E) 122
_init(9E) 125
ioctl(9E) 128
ks_snapshot(9E) 132
ks_update(9E) 134
mapdev_access(9E) 136
mapdev_dup(9E) 138
mapdev_free(9E) 140
mmap(9E) 141
open(9E) 145
power(9E) 148
print(9E) 150
probe(9E) 151
prop_op(9E) 152
put(9E) 154
read(9E) 156
segmap(9E) 158
srv(9E) 160
strategy(9E) 162
tran_abort(9E) 163
tran_bus_reset(9e) 164
tran_destroy_pkt(9E) 165
tran_dmafree(9E) 168
tran_getcap(9E) 169
tran_init_pkt(9E) 171
tran_quiesce(9e) 174
tran_reset(9E) 175
tran_reset_notify(9E) 177
tran_setcap(9E) 178

tran_start(9E)	180
tran_sync_pkt(9E)	183
tran_tgt_free(9E)	184
tran_tgt_init(9E)	185
tran_tgt_probe(9E)	186
tran_unquiesce(9e)	187
write(9E)	188

Index	189
--------------	------------

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	<p>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.</p> <p>The following special characters are used in this section:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 10px;">[]</td> <td>Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.</td> </tr> <tr> <td style="padding-right: 10px;">. . .</td> <td>Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".</td> </tr> <tr> <td style="padding-right: 10px;"> </td> <td>Separator. Only one of the arguments separated by this character can be specified at a time.</td> </tr> <tr> <td style="padding-right: 10px;">{ }</td> <td>Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.</td> </tr> </table>	[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.	. . .	Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".		Separator. Only one of the arguments separated by this character can be specified at a time.	{ }	Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.
[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.								
. . .	Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".								
	Separator. Only one of the arguments separated by this character can be specified at a time.								
{ }	Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.								

PROTOCOL	This section occurs only in subsection 3R to indicate the protocol description file.
DESCRIPTION	This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES. Interactive commands, subcommands, requests, macros, and functions are described under USAGE.
IOCTL	This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the <code>ioctl(2)</code> system call is called <code>ioctl</code> and generates its own heading. <code>ioctl</code> calls for a specific device are listed alphabetically (on the man page for that specific device). <code>ioctl</code> calls are used for a particular class of devices all of which have an <code>io</code> ending, such as <code>mtio(7I)</code> .
OPTIONS	This section lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.
OPERANDS	This section lists the command operands and describes how they affect the actions of the command.
OUTPUT	This section describes the output – standard output, standard error, or output files – generated by the command.
RETURN VALUES	If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.
ERRORS	On failure, most functions place an error code in the global variable <code>errno</code> indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than

	one condition can cause the same error, each condition is described in a separate paragraph under the error code.
USAGE	This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality: Commands Modifiers Variables Expressions Input Grammar
EXAMPLES	This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command-line entry and machine response is shown. Whenever an example is given, the prompt is shown as <code>example%</code> , or if the user must be superuser, <code>example#</code> . Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.
ENVIRONMENT VARIABLES	This section lists any environment variables that the command or function affects, followed by a brief description of the effect.
EXIT STATUS	This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion, and values other than zero for various error conditions.
FILES	This section lists all file names referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.
ATTRIBUTES	This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See <code>attributes(5)</code> for more information.
SEE ALSO	This section lists references to other man pages, in-house documentation, and outside publications.

DIAGNOSTICS	This section lists diagnostic messages with a brief explanation of the condition causing the error.
WARNINGS	This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.
NOTES	This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.
BUGS	This section describes known bugs and, wherever possible, suggests workarounds.

Introduction

Intro(9E)

NAME	Intro – overview of device driver interfaces and introduction to driver entry points
DESCRIPTION	This page provides an overview of device driver interfaces and all of the Section 9 man pages (9E, 9F, 9P, and 9S). This overview is followed by an introduction to Section 9E, the driver entry-point routines.
Overview of Device Driver Interfaces	<p>Section 9 provides reference information needed to write device drivers for the Solaris operating environment. It describes the interfaces provided by the Device Driver Interface and the Driver-Kernel Interface (DDI/DKI).</p> <p>Porting</p> <p>Software is usually considered portable if it can be adapted to run in a different environment more cheaply than it can be rewritten. The new environment may include a different processor, operating system, and even the language in which the program is written, if a language translator is available. Likewise the new environment might include multiple processors. More often, however, software is ported between environments that share an operating system, processor, and source language. The source code is modified to accommodate the differences in compilers or processors or releases of the operating system.</p> <p>In the past, device drivers did not port easily for one or more of the following reasons:</p> <ul style="list-style-type: none">■ To enhance functionality, members had been added to kernel data structures accessed by drivers, or the sizes of existing members had been redefined.■ The calling or return syntax of kernel functions had changed.■ Driver developers did not use existing kernel functions where available, or relied on undocumented side effects that were not maintained in the next release.■ Architecture-specific code had been scattered throughout the driver when it could have been isolated. <p>Operating systems are periodically reissued to customers as a way to improve performance, fix bugs, and add new features. This is probably the most common threat to compatibility encountered by developers responsible for maintaining software. Another common problem is upgrading hardware. As new hardware is developed, customers occasionally decide to upgrade to faster, more capable computers of the same family. Although they may run the same operating system as those being replaced, architecture-specific code may prevent the software from porting.</p> <p>Scope of Interfaces</p> <p>Although application programs have all of the porting problems mentioned, developers attempting to port device drivers have special challenges. Before describing the DDI/DKI, it is necessary to understand the position of device drivers in operating systems.</p> <p>Device drivers are kernel modules that control data transferred to and received from peripheral devices but are developed independently from the rest of the kernel. If the goal of achieving complete freedom in modifying the kernel is to be reconciled with</p>

the goal of binary compatibility with existing drivers, the interaction between drivers and the kernel must be rigorously regulated. This driver/kernel service interface is the most important of the three distinguishable interfaces for a driver, summarized as follows:

- Driver–Kernel. I/O System calls result in calls to driver entry point routines. These make up the kernel-to-driver part of the service interface, described in Section 9E. Drivers may call any of the functions described in Section 9F. These are the driver-to-kernel part of the interface.
- Driver–Hardware. All drivers (except software drivers) must include code for interrupt handling, and may also perform direct memory access (DMA). These and other hardware-specific interactions make up the driver/hardware interface.
- Driver–Boot/Configuration Software. The interaction between the driver and the boot and configuration software is the third interface affecting drivers.

Scope of the DDI/DKI

The primary goal of the DDI/DKI is to facilitate both source and binary portability across successive releases of the operating systems on a particular machine. In addition, it promotes source portability across implementations of UNIX on different machines, and applies only to implementations based on System V Release 4. The DDI/DKI consists of several sections:

- DDI/DKI Architecture Independent - These interfaces are supported on all implementations of System V Release 4.
- DKI-only - These interfaces are part of System V Release 4, and may not be supported in future releases of System V. There are only two interfaces in this class, `segmap(9E)` and `hat_getkpfnum(9F)`
- Solaris DDI - These interfaces specific to Solaris.
- Solaris SPARC specific DDI - These interfaces are specific to the SPARC processor, and may not be available on other processors supported by Solaris.
- Solaris IA specific DDI - These interfaces are specific to the IA processor, and may not be available on other processors supported by Solaris.

To achieve the goal of source and binary compatibility, the functions, routines, and structures specified in the DDI/DKI must be used according to these rules.

- Drivers cannot access system state structures (for example, `u` and `sysinfo`) directly.
- For structures external to the driver that may be accessed directly, only the utility functions provided in Section 9F should be used. More generally, these functions should be used wherever possible.
- The headers `<sys/ddi.h>` and `<sys/sunddi.h>` must be the last header files included by the driver.

Audience

Section 9 is for software engineers responsible for creating, modifying, or maintaining drivers that run on this operating system and beyond. It assumes that the reader is familiar with system internals and the C programming language.

PCMCIA Standard

The *PC Card 95 Standard* is listed under the *SEE ALSO* heading in some Section 9 reference pages. This refers to documentation published by the Personal Computer Memory Card International Association (PCMCIA) and the Japan Electronic Industry Development Association (JEIDA).

How to Use Section 9

Section 9 is divided into the following subsections:

- 9E Driver Entry Points – contains reference pages for all driver entry point routines.
- 9F Kernel Functions – contains reference pages for all driver support routines.
- 9P Driver Properties – contains reference pages for driver properties.
- 9S Data Structures – contains reference pages for driver-related structures.

Compatibility Note

Sun Microsystem's implementation of the DDI/DKI was designed to provide binary compatibility for third-party device drivers across currently supported hardware platforms across minor releases of the operating system. However, unforeseen technical issues may force changes to the binary interface of the DDI/DKI. We cannot therefore promise or in any way assure that DDI/DKI-compliant device drivers will continue to operate correctly on future releases.

Introduction to Section 9E

Section 9E describes the entry-point routines a developer can include in a device driver. These are called entry-point because they provide the calling and return syntax from the kernel into the driver. Entry-points are called, for instance, in response to system calls, when the driver is loaded, or in response to STREAMS events.

Kernel functions usable by the driver are described in section 9F.

In this section, reference pages contain the following headings:

- **NAME** describes the routine's purpose.
- **SYNOPSIS** summarizes the routine's calling and return syntax.
- **INTERFACE LEVEL** describes any architecture dependencies. It also indicates whether the use of the entry point is required, optional, or discouraged.
- **ARGUMENTS** describes each of the routine's arguments.
- **DESCRIPTION** provides general information about the routine.
- **RETURN VALUES** describes each of the routine's return values.
- **SEE ALSO** gives sources for further information.

Overview of Driver Entry-Point Routines and Naming Conventions

By convention, a prefix string is added to the driver routine names. For a driver with the prefix *prefix*, the driver code may contain routines named *prefixopen*, *prefixclose*, *prefixread*, *prefixwrite*, and so forth. All global variables associated with the driver should also use the same prefix.

All routines and data should be declared as `static`.

Every driver MUST include `<sys/ddi.h>` and `<sys/sunddi.h>`, in that order, and after all other include files.

The following table summarizes the STREAMS driver entry points described in this section.

Routine	Type
<code>put</code>	DDI/DKI
<code>srv</code>	DDI/DKI

The following table summarizes the driver entry points described in this section.

Routine	Type
<code>_fini</code>	Solaris DDI
<code>_info</code>	Solaris DDI
<code>_init</code>	Solaris DDI
<code>aread</code>	Solaris DDI
<code>attach</code>	Solaris DDI
<code>awrite</code>	Solaris DDI
<code>chpoll</code>	DDI/DKI
<code>close</code>	DDI/DKI
<code>detach</code>	Solaris DDI
<code>devmap</code>	Solaris DDI
<code>devmap_access</code>	Solaris DDI
<code>devmap_contextmgt</code>	Solaris DDI
<code>devmap_dup</code>	Solaris DDI
<code>devmap_map</code>	Solaris DDI

Intro(9E)

Routine	Type
devmap_unmap	Solaris DDI
dump	Solaris DDI
getinfo	Solaris DDI
identify	Solaris DDI
ioctl	DDI/DKI
ks_update	Solaris DDI
mapdev_access	Solaris DDI
mapdev_dup	Solaris DDI
mapdev_free	Solaris DDI
mmap	DKI only
open	DDI/DKI
power	Solaris DDI
print	DDI/DKI
probe	Solaris DDI
prop_op	Solaris DDI
read	DDI/DKI
segmap	DKI only
strategy	DDI/DKI
tran_abort	Solaris DDI
tran_destroy_pkt	Solaris DDI
tran_dmafree	Solaris DDI
tran_getcap	Solaris DDI
tran_init_pkt	Solaris DDI
tran_reset	Solaris DDI
tran_reset_notify	Solaris DDI
tran_setcap	Solaris DDI
tran_start	Solaris DDI
tran_sync_pkt	Solaris DDI
tran_tgt_free	Solaris DDI

Routine	Type
tran_tgt_init	Solaris DDI
tran_tgt_probe	Solaris DDI
write	DDI/DKI

The following table lists the error codes returned by a driver routine when it encounters an error. The error values are listed in alphabetic order and are defined in `sys/errno.h`. In the driver `open(9E)`, `close(9E)`, `ioctl(9E)`, `read(9E)`, and `write(9E)` routines, errors are passed back to the user by calling `bioerror(9F)` to set `b_flags` to the proper error code. In the driver `strategy(9E)` routine, errors are passed back to the user by setting the `b_error` member of the `buf(9S)` structure to the error code. For STREAMS `ioctl` routines, errors should be sent upstream in an `M_IOCNAK` message. For STREAMS `read()` and `write()` routines, errors should be sent upstream in an `M_ERROR` message. The driver `print` routine should not return an error code because the function that it calls, `cmn_err(9F)`, is declared as `void` (no error is returned).

Error Value	Error Description
EAGAIN	Kernel resources, such as the <code>buf</code> structure or cache memory, are not available at this time (device may be busy, or the system resource is not available). This is used in <code>open</code> , <code>ioctl</code> , <code>read</code> , <code>write</code> , and <code>strategy</code> .
EFAULT	An invalid address has been passed as an argument; memory addressing error. This is used in <code>open</code> , <code>close</code> , <code>ioctl</code> , <code>read</code> , <code>write</code> , and <code>strategy</code> .
EINTR	Sleep interrupted by signal. This is used in <code>open</code> , <code>close</code> , <code>ioctl</code> , <code>read</code> , <code>write</code> , and <code>strategy</code> .
EINVAL	An invalid argument was passed to the routine. This is used in <code>open</code> , <code>ioctl</code> , <code>read</code> , <code>write</code> , and <code>strategy</code> .
EIO	A device error occurred; an error condition was detected in a device status register (the I/O request was valid, but an error occurred on the device). This is used in <code>open</code> , <code>close</code> , <code>ioctl</code> , <code>read</code> , <code>write</code> , and <code>strategy</code> .
ENXIO	An attempt was made to access a device or subdevice that does not exist (one that is not configured); an attempt was made to perform an invalid I/O operation; an incorrect minor number was specified. This is used in <code>open</code> , <code>close</code> , <code>ioctl</code> , <code>read</code> , <code>write</code> , and <code>strategy</code> .
EPERM	A process attempting an operation did not have required permission. This is used in <code>open</code> , <code>ioctl</code> , <code>read</code> , <code>write</code> , and <code>strategy</code> .

Intro(9E)

Error Value	Error Description
EROFS	An attempt was made to open for writing a read-only device. This is used in open.

The table below cross references error values to the driver routines from which the error values can be returned.

open	close	ioctl	read, write and strategy
EAGAIN	EFAULT	EAGAIN	EAGAIN
EFAULT	EINTR	EFAULT	EFAULT
EINTR	EIO	EINTR	EINTR
EINVAL	ENXIO	EINVAL	EINVAL
EIO		EIO	EIO
ENXIO		ENXIO	ENXIO
EPERM		EPERM	
EROFS			

SEE ALSO [intro\(9F\)](#), [intro\(9S\)](#)

Driver Entry Points

aread(9E)

NAME	aread – asynchronous read from a device
SYNOPSIS	<pre>#include <sys/uio.h> #include <sys/aio_req.h> #include <sys/cred.h> #include <sys/ddi.h> #include <sys/sunddi.h> intprefix aread(dev_t dev, struct aio_req *aio_reqp, cred_t *cred_p);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). This entry point is <i>optional</i> . Drivers that do not support an aread() entry point should use nodev(9F)
PARAMETERS	<p><i>dev</i> Device number.</p> <p><i>aio_reqp</i> Pointer to the aio_req(9S) structure that describes where the data is to be stored.</p> <p><i>cred_p</i> Pointer to the credential structure.</p>
DESCRIPTION	<p>The driver's aread() routine is called to perform an asynchronous read. getminor(9F) can be used to access the minor number component of the <i>dev</i> argument. aread() may use the credential structure pointed to by <i>cred_p</i> to check for superuser access by calling drv_priv(9F). The aread() routine may also examine the uio(9S) structure through the aio_req structure pointer, <i>aio_reqp</i>. aread() must call aphysio(9F) with the aio_req pointer and a pointer to the driver's strategy(9E) routine.</p> <p>No fields of the uio(9S) structure pointed to by <i>aio_req</i>, other than <i>uio_offset</i> or <i>uio_loffset</i>, may be modified for non-seekable devices.</p>
RETURN VALUES	The aread() routine should return 0 for success, or the appropriate error number.
CONTEXT	This function is called from user context only.
EXAMPLES	<p>EXAMPLE 1 The following is an example of an aread() routine:</p> <pre>static int xxaread(dev_t dev, struct aio_req *aio, cred_t *cred_p) { int instance; struct xxstate *xsp; instance = getminor(dev); xsp = ddi_get_soft_state(statep, instance); /*Verify soft state structure has been allocated */ if (xsp == NULL) return (ENXIO); return (aphysio(xxstrategy, anocancel, dev, B_READ, xxminphys, aio)); }</pre>

aread(9E)

EXAMPLE 1 The following is an example of an aread () routine: *(Continued)*

SEE ALSO read(2), aioread(3AIO), awrite(9E), read(9E), strategy(9E), write(9E),
anocancel(9F), aphysio(9F), ddi_get_soft_state(9F), drv_priv(9F),
getminor(9F), minphys(9F), nodev(9F), aio_req(9S), cb_ops(9S), uio(9S)

Writing Device Drivers

BUGS There is no way other than calling aphysio(9F) to accomplish an asynchronous read.

attach(9E)

NAME	attach – Attach a device to the system, or resume it
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>attach(dev_info_t *dip, ddi_attach_cmd_t cmd);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>dip</i> A pointer to the device's dev_info structure. <i>cmd</i> Attach type. Possible values are DDI_ATTACH, DDI_PM_RESUME (obsolete), and DDI_RESUME. Other values are reserved. The driver must return DDI_FAILURE if reserved values are passed to it.
DESCRIPTION	The attach(9E) function is the device-specific initialization entry point. This entry point is <i>required</i> and must be written.
DDI_ATTACH	<p>The DDI_ATTACH command must be provided in the attach(9E) entry point. DDI_ATTACH is used to initialize a given device instance. When attach(9E) is called with <i>cmd</i> set to DDI_ATTACH, all normal kernel services (such as kmem_alloc(9F)) are available for use by the driver. Device interrupts are not blocked when attaching a device to the system.</p> <p>The attach(9E) function will be called once for each instance of the device on the system with <i>cmd</i> set to DDI_ATTACH. Until attach(9E) succeeds, the only driver entry points which may be called are open(9E) and getinfo(9E). See the <i>Writing Device Drivers</i> for more information. The instance number may be obtained using ddi_get_instance(9F).</p> <p>At attach time, all components of a power-manageable device are assumed to be at unknown levels. Before using the device, the driver needs to bring the required component(s) to a known power level. The pm_raise_power(9F) function can be used to set the power level of a component. This function must not be called before data structures referenced in power(9E) have been initialized.</p>
DDI_PM_RESUME	<p>The DDI_PM_RESUME command is required only if the device driver uses original Power Management interfaces (driver calls pm_create_components(9F)). This entry point is not needed if the device driver uses new automatic device Power Management interfaces (driver exports pm-components(9P) property instead of calling pm_create_components(9F)). The DDI_PM_RESUME command is obsolete and will be removed in a future release.</p> <p>The attach() function may be called with <i>cmd</i> set to DDI_PM_RESUME after detach(9E) has been successfully called with <i>cmd</i> set to DDI_PM_SUSPEND. When called with <i>cmd</i> set to DDI_PM_RESUME, attach() must restore the hardware state of a device (power may have been removed from the device), allow pending requests to continue, and service new requests.</p>

attach(9E)

The driver must not make any assumptions about the state of the hardware, but must restore it to the state it had when the detach(9E) entry point was called with DDI_PM_SUSPEND.

DDI_RESUME

The attach() function may be called with *cmd* set to DDI_RESUME after detach(9E) has been successfully called with *cmd* set to DDI_SUSPEND.

When called with *cmd* set to DDI_RESUME, attach() must restore the hardware state of a device (power may have been removed from the device), allow pending requests to continue, and service new requests. In this case, the driver must not make any assumptions about the state of the hardware, but must restore the state of the device except for the power level of components.

If the device driver uses original Power Management interfaces (driver calls pm_create_components(9F)) and device is still suspended by DDI_PM_SUSPEND, the only effect of DDI_RESUME is to allow the driver to call ddi_dev_is_needed(9F) for any new or pending requests, as a subsequent call to attach() will be made with *cmd* set to DDI_PM_RESUME to restore the hardware state.

If the device driver uses the new automatic device Power Management interfaces (driver exports pm-components(9P) property instead of calling pm_create_components(9F)), then while processing a DDI_RESUME command, the Power Management framework sets its notion of the power level of each component of a device to *unknown*.

The driver can deal with components during DDI_RESUME in one of the following ways:

1. If the driver can determine the power level of the component without having to power it up (for example, by calling ddi_peek(9F) or some other device-specific method) then it should notify the power level to the framework by calling pm_power_has_changed(9F).
2. The driver must also set its own notion of the power level of the component to *unknown*. The system will consider the component idle or busy based on the most recent call to pm_idle_component(9F) or pm_busy_component(9F) for that component. If the component is idle for sufficient time, the framework will call into the driver's power(9E) entry point to turn the component off. If the driver needs to access the device, then it must call pm_raise_power(9F) to bring the component up to the level needed for the device access to succeed. The driver must honor any request to set the power level of the component, since it cannot make any assumption about what power level the component has (or it should have called pm_power_has_changed(9F) as outlined above). As a special case of this, the driver may bring the component to a known state because it wants to perform an operation on the device as part of its DDI_RESUME processing (such as loading firmware so that it can detect hot-plug events).

RETURN VALUES

The attach() function returns:

DDI_SUCCESS Successful completion

attach(9E)

DDI_FAILURE Operation failed

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving (DDI_PM_RESUME is obsolete)

SEE ALSO `cpr(7)`, `pm(7D)`, `pm(9P)`, `pm-components(9P)`, `detach(9E)`, `getinfo(9E)`, `identify(9E)`, `open(9E)`, `power(9E)`, `probe(9E)`, `ddi_add_intr(9F)`, `ddi_create_minor_node(9F)`, `ddi_get_instance(9F)`, `ddi_map_regs(9F)`, `kmem_alloc(9F)`, `pm_create_components(9F)`, `pm_raise_power(9F)`

Writing Device Drivers

NAME	awrite – asynchronous write to a device
SYNOPSIS	<pre>#include <sys/uio.h> #include <sys/aio_req.h> #include <sys/cred.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>awrite(dev_t dev, struct aio_req *aio_req, cred_t *cred_p);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). This entry point is optional. Drivers that do not support an <code>awrite()</code> entry point should use <code>nodev(9F)</code>
PARAMETERS	<p><i>dev</i> Device number.</p> <p><i>aio_req</i> Pointer to the <code>aio_req(9S)</code> structure that describes where the data is stored.</p> <p><i>cred_p</i> Pointer to the credential structure.</p>
DESCRIPTION	<p>The driver's <code>awrite()</code> routine is called to perform an asynchronous write. <code>getminor(9F)</code> can be used to access the minor number component of the <i>dev</i> argument. <code>awrite()</code> may use the credential structure pointed to by <i>cred_p</i> to check for superuser access by calling <code>drv_priv(9F)</code>. The <code>awrite()</code> routine may also examine the <code>uio(9S)</code> structure through the <code>aio_req</code> structure pointer, <i>aio_req</i>. <code>awrite()</code> must call <code>aphysio(9F)</code> with the <code>aio_req</code> pointer and a pointer to the driver's <code>strategy(9E)</code> routine.</p> <p>No fields of the <code>uio(9S)</code> structure pointed to by <i>aio_req</i>, other than <code>uio_offset</code> or <code>uio_loffset</code>, may be modified for non-seekable devices.</p>
RETURN VALUES	The <code>awrite()</code> routine should return 0 for success, or the appropriate error number.
CONTEXT	This function is called from user context only.
EXAMPLES	<p>EXAMPLE 1 Using the <code>awrite</code> routine:</p> <p>The following is an example of an <code>awrite()</code> routine:</p> <pre>static int xxawrite(dev_t dev, struct aio_req *aio, cred_t *cred_p) { int instance; struct xxstate *xsp; instance = getminor(dev); xsp = ddi_get_soft_state(statep, instance); /*Verify soft state structure has been allocated */ if (xsp == NULL) return (ENXIO); return (aphysio(xxstrategy, anocancel, dev, B_WRITE, xxminphys, aio)); }</pre>

awrite(9E)

EXAMPLE 1 Using the `awrite` routine: *(Continued)*

SEE ALSO `write(2)`, `aiowrite(3AIO)`, `aread(9E)`, `read(9E)`, `strategy(9E)`, `write(9E)`, `anocancel(9F)`, `aphysio(9F)`, `ddi_get_soft_state(9F)`, `drv_priv(9F)`, `getminor(9F)`, `minphys(9F)`, `nodev(9F)`, `aio_req(9S)`, `cb_ops(9S)`, `uio(9S)`

Writing Device Drivers

BUGS There is no way other than calling `aphysio(9F)` to accomplish an asynchronous write.

NAME	chpoll – poll entry point for a non-STREAMS character driver
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/poll.h> #include <sys/ddi.h> #include <sys/sunddi.h> int prefixchpoll(dev_t dev, short events, int anyyet, short *reventsp, struct pollhead **phpp);</pre>
INTERFACE LEVEL PARAMETERS	<p>This entry point is optional. Architecture independent level 1 (DDI/DKI).</p> <p><i>dev</i> The device number for the device to be polled.</p> <p><i>events</i> The events that may occur. Valid events are:</p> <p style="margin-left: 40px;">POLLIN Data other than high priority data may be read without blocking.</p> <p style="margin-left: 40px;">POLLOUT Normal data may be written without blocking.</p> <p style="margin-left: 40px;">POLLPRI High priority data may be received without blocking.</p> <p style="margin-left: 40px;">POLLHUP A device hangup has occurred.</p> <p style="margin-left: 40px;">POLLERR An error has occurred on the device.</p> <p style="margin-left: 40px;">POLLRDNORM Normal data (priority band = 0) may be read without blocking.</p> <p style="margin-left: 40px;">POLLRDBAND Data from a non-zero priority band may be read without blocking</p> <p style="margin-left: 40px;">POLLWRNORM The same as POLLOUT.</p> <p style="margin-left: 40px;">POLLWRBAND Priority data (priority band > 0) may be written.</p> <p><i>anyyet</i> A flag that is non-zero if any other file descriptors in the <code>pollfd</code> array have events pending. The <code>poll(2)</code> system call takes a pointer to an array of <code>pollfd</code> structures as one of its arguments. See the <code>poll(2)</code> reference page for more details.</p> <p><i>reventsp</i> A pointer to a bitmask of the returned events satisfied.</p> <p><i>phpp</i> A pointer to a pointer to a <code>pollhead</code> structure.</p>

chpoll(9E)

DESCRIPTION

The `chpoll()` entry point routine is used by non-STREAMS character device drivers that wish to support polling. The driver must implement the polling discipline itself. The following rules must be followed when implementing the polling discipline:

1. Implement the following algorithm when the `chpoll()` entry point is called:

```
if (events_are_satisfied_now) {
    *reventsp = mask_of_satisfied_events;
} else {
    *reventsp = 0;
    if (!anyyet)
        *phpp = &my_local_pollhead_structure;
}
return (0);
```

2. Allocate an instance of the `pollhead` structure. This instance may be tied to the per-minor data structure defined by the driver. The `pollhead` structure should be treated as a "black box" by the driver. None of its fields should be referenced. However, the size of this structure is guaranteed to remain the same across releases.
3. Call the `pollwakeup()` function whenever an event of type `events` listed above occur. This function can be called with multiple events at one time. The driver must not hold any mutex across the call to `pollwakeup(9F)` that is acquired in its `chpoll()` entry point, or a deadlock may result.

RETURN VALUES

`chpoll()` should return 0 for success, or the appropriate error number.

SEE ALSO

`poll(2)`, `nochpoll(9F)`, `pollwakeup(9F)`

Writing Device Drivers

NAME	close – relinquish access to a device
Block and Character	<pre>#include <sys/types.h> #include <sys/file.h> #include <sys/errno.h> #include <sys/open.h> #include <sys/cred.h> #include <sys/ddi.h> #include <sys/sunddi.h> int prefixclose(dev_t dev, int flag, int otyp, cred_t *cred_p);</pre>
STREAMS	<pre>#include <sys/types.h> #include <sys/stream.h> #include <sys/file.h> #include <sys/errno.h> #include <sys/open.h> #include <sys/cred.h> #include <sys/ddi.h> #include <sys/sunddi.h> int prefixclose(queue_t *q, int flag, cred_t *cred_p);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI). This entry point is <i>required</i> for block devices.
Block and Character	<p><i>dev</i> Device number.</p> <p><i>flag</i> File status flag, as set by the <code>open(2)</code> or modified by the <code>fcntl(2)</code> system calls. The flag is for information only—the file should always be closed completely. Possible values are: <code>FEXCL</code>, <code>FNDELAY</code>, <code>FREAD</code>, <code>FKLYR</code>, and <code>FWRITE</code>. Refer to <code>open(9E)</code> for more information.</p> <p><i>otyp</i> Parameter supplied so that the driver can determine how many times a device was opened and for what reasons. The flags assume the <code>open()</code> routine may be called many times, but the <code>close()</code> routine should only be called on the last <code>close()</code> of a device.</p> <p><code>OTYP_BLK</code> Close was through block interface for the device.</p> <p><code>OTYP_CHR</code> Close was through the raw/character interface for the device.</p> <p><code>OTYP_LYR</code> Close a layered process (a higher-level driver called the <code>close()</code> routine of the device).</p> <p><i>*cred_p</i> Pointer to the user credential structure.</p>
STREAMS	<p><i>*q</i> Pointer to <code>queue(9S)</code> structure used to reference the read side of the driver. (A queue is the central node of a collection of structures and routines pointed to by a queue.)</p> <p><i>flag</i> File status flag.</p>

close(9E)

	<p><i>*cred_p</i> Pointer to the user credential structure.</p>
DESCRIPTION	<p>For STREAMS drivers, the <code>close()</code> routine is called by the kernel through the <code>cb_ops(9S)</code> table entry for the device. (Modules use the <code>fmodsw</code> table.) A non-null value in the <code>d_str</code> field of the <code>cb_ops</code> entry points to a <code>streamtab</code> structure, which points to a <code>qinit(9S)</code> containing a pointer to the <code>close()</code> routine. Non-STREAMS <code>close()</code> routines are called directly from the <code>cb_ops</code> table.</p> <p><code>close()</code> ends the connection between the user process and the device, and prepares the device (hardware and software) so that it is ready to be opened again.</p> <p>A device may be opened simultaneously by multiple processes and the <code>open()</code> driver routine is called for each open, but the kernel will only call the <code>close()</code> routine when the last process using the device issues a <code>close(2)</code> or <code>umount(2)</code> system call or exits. (An exception is a close occurring with the <code>otyp</code> argument set to <code>OTYP_LYR</code>, for which a close (also having <code>otyp = OTYP_LYR</code>) occurs for each open.)</p> <p>In general, a <code>close()</code> routine should always check the validity of the minor number component of the <code>dev</code> parameter. The routine should also check permissions as necessary, by using the user credential structure (if pertinent), and the appropriateness of the <code>flag</code> and <code>otyp</code> parameter values.</p> <p><code>close()</code> could perform any of the following general functions:</p> <ul style="list-style-type: none">■ disable interrupts■ hang up phone lines■ rewind a tape■ deallocate buffers from a private buffering scheme■ unlock an unsharable device (that was locked in the <code>open()</code> routine)■ flush buffers■ notify a device of the close■ deallocate any resources allocated on open <p>The <code>close()</code> routines of STREAMS drivers and modules are called when a stream is dismantled or a module popped. The steps for dismantling a stream are performed in the following order. First, any multiplexor links present are unlinked and the lower streams are closed. Next, the following steps are performed for each module or driver on the stream, starting at the head and working toward the tail:</p> <ol style="list-style-type: none">1. The write queue is given a chance to drain.2. The <code>close()</code> routine is called.3. The module or driver is removed from the stream.
RETURN VALUES	<p><code>close()</code> should return 0 for success, or the appropriate error number. Return errors rarely occur, but if a failure is detected, the driver should decide whether the severity of the problem warrants either displaying a message on the console or, in worst cases, triggering a system panic. Generally, a failure in a <code>close()</code> routine occurs because a problem occurred in the associated device.</p>

close(9E)

SEE ALSO | close(2), fcntl(2), open(2), umount(2), detach(9E), open(9E), cb_ops(9S),
qinit(9S), queue(9S)

Writing Device Drivers

STREAMS Programming Guide

csx_event_handler(9E)

NAME	csx_event_handler – PC Card driver event handler
SYNOPSIS	<pre>#include <sys/pccard.h> int32_t prefixevent_handler(event_t event, int32_t priority, event_callback_args_t *args);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI)</p> <p><i>event</i> The event.</p> <p><i>priority</i> The priority of the event.</p> <p><i>args</i> A pointer to the event_callback_t structure.</p>
DESCRIPTION	<p>Each instance of a PC Card driver must register an event handler to manage events associated with its PC Card. The driver event handler is registered using the event_handler field of the client_req_t structure passed to csx_RegisterClient(9F). The driver may also supply a parameter to be passed to its event handler function using the event_callback_args.client_data field. Typically, this argument is the driver instance's soft state pointer. The driver also registers which events it is interested in receiving through the EventMask field of the client_req_t structure.</p> <p>Each event is delivered to the driver with a priority, <i>priority</i>. High priority events with CS_EVENT_PRI_HIGH set in <i>priority</i> are delivered above lock level, and the driver must use its high-level event mutex initialized with the iblk_cookie returned by csx_RegisterClient(9F) to protect such events. Low priority events with CS_EVENT_PRI_LOW set in <i>priority</i> are delivered below lock level, and the driver must use its low-level event mutex initialized with a NULL interrupt cookie to protect these events.</p> <p>csx_RegisterClient(9F) registers the driver's event handler, but no events begin to be delivered to the driver until after a successful call to csx_RequestSocketMask(9F).</p> <p>In all cases, Card Services delivers an event to each driver instance associated with a function on a multiple function PC Card.</p>
Event Indications	<p>The events and their indications are listed below; they are always delivered as low priority unless otherwise noted:</p> <p>CS_EVENT_REGISTRATION_COMPLETE A registration request processed in the background has been completed.</p> <p>CS_EVENT_CARD_INSERTION A PC Card has been inserted in a socket.</p> <p>CS_EVENT_CARD_READY A PC Card's READY line has transitioned from the busy to ready state.</p>

CS_EVENT_CARD_REMOVAL

A PC Card has been removed from a socket. This event is delivered twice; first as a high priority event, followed by delivery as a low priority event. As a high priority event, the event handler should only note that the PC Card is no longer present to prevent accesses to the hardware from occurring. As a low priority event, the event handler should release the configuration and free all I/O, window and IRQ resources for use by other PC Cards.

CS_EVENT_BATTERY_LOW

The battery on a PC Card is weak and is in need of replacement.

CS_EVENT_BATTERY_DEAD

The battery on a PC Card is no longer providing operational voltage.

CS_EVENT_PM_RESUME

Card Services has received a resume notification from the system's Power Management software.

CS_EVENT_PM_SUSPEND

Card Services has received a suspend notification from the system's Power Management software.

CS_EVENT_CARD_LOCK

A mechanical latch has been manipulated preventing the removal of the PC Card from the socket.

CS_EVENT_CARD_UNLOCK

A mechanical latch has been manipulated allowing the removal of the PC Card from the socket.

CS_EVENT_EJECTION_REQUEST

A request that the PC Card be ejected from a socket using a motor-driven mechanism.

CS_EVENT_EJECTION_COMPLETE

A motor has completed ejecting a PC Card from a socket.

CS_EVENT_ERASE_COMPLETE

A queued erase request that is processed in the background has been completed.

CS_EVENT_INSERTION_REQUEST

A request that a PC Card be inserted into a socket using a motor-driven mechanism.

CS_EVENT_INSERTION_COMPLETE

A motor has completed inserting a PC Card in a socket.

CS_EVENT_CARD_RESET

A hardware reset has occurred.

CS_EVENT_RESET_REQUEST

A request for a physical reset by a client.

csx_event_handler(9E)

CS_EVENT_RESET_COMPLETE
A reset request that is processed in the background has been completed.

CS_EVENT_RESET_PHYSICAL
A reset is about to occur.

CS_EVENT_CLIENT_INFO
A request that the client return its client information data. If `GET_CLIENT_INFO_SUBSVC(args->client_info.Attributes)` is equal to `CS_CLIENT_INFO_SUBSVC_CS`, the driver should fill in the other fields in the `client_info` structure as described below, and return `CS_SUCCESS`. Otherwise, it should return `CS_UNSUPPORTED_EVENT`.

`args->client_data.Attributes`
Must be OR'ed with `CS_CLIENT_INFO_VALID`.

`args->client_data.Revision`
Must be set to a driver-private version number.

`args->client_data.CSLevel`
Must be set to `CS_VERSION`.

`args->client_data.RevDate`
Must be set to the revision date of the PC Card driver, using `CS_CLIENT_INFO_MAKE_DATE(day, month, year)`. *day* must be the day of the month, *month* must be the month of the year, and *year* must be the year, offset from a base of 1980. For example, this field could be set to a revision date of July 4 1997 with `CS_CLIENT_INFO_MAKE_DATE(4, 7, 17)`.

`args->client_data.ClientName`
A string describing the PC Card driver should be copied into this space.

`args->client_data.VendorName`
A string supplying the name of the PC Card driver vendor should be copied into this space.

`args->client_data.DriverName`
A string supplying the name of the PC Card driver will be copied into this space by Card Services after the PC Card driver has successfully processed this event; the driver does not need to initialize this field.

CS_EVENT_WRITE_PROTECT
The write protect status of the PC Card in the indicated socket has changed. The current write protect state of the PC Card is in the `args->info` field:

CS_EVENT_WRITE_PROTECT_WPOFF
Card is not write protected.

CS_EVENT_WRITE_PROTECT_WPON
Card is write protected.

STRUCTURE MEMBERS

The structure members of `event_callback_args_t` are:

```
void          *info;                /* event-specific information */
void          *client_data;         /* driver-private data */
```

```
client_info_t    client_info;    /* client information*/
```

The structure members of `client_info_t` are:

```
unit32_t        Attributes;      /* attributes */
unit32_t        Revisions;       /* version number */
uint32_t        CSLevel;         /* Card Services version */
uint32_t        RevDate;         /* revision date */
char            ClientName[CS_CLIENT_INFO_MAX_NAME_LEN];
                /*PC Card driver description */
char            VendorName[CS_CLIENT_INFO_MAX_NAME_LEN];
                /*PC Card driver vendor name */
char            DriverName[MODMAXNAMELEN];
                /* PC Card driver name */
```

RETURN VALUES	<code>CS_SUCCESS</code>	The event was handled successfully.
	<code>CS_UNSUPPORTED_EVENT</code>	Driver does not support this event.
	<code>CS_FAILURE</code>	Error occurred while handling this event.

CONTEXT This function is called from high-level interrupt context in the case of high priority events, and from kernel context in the case of low priority events.

EXAMPLES **EXAMPLE 1**

```
static int
xx_event(event_t event, int priority, event_callback_args_t *args)
{
    int    rval;
    struct xxx    *xxx = args->client_data;
    client_info_t    *info = &args->client_info;

    switch (event) {
    case CS_EVENT_REGISTRATION_COMPLETE:
        ASSERT(priority & CS_EVENT_PRI_LOW);
        mutex_enter(&xxx->event_mutex);
        xxx->card_state |= XX_REGISTRATION_COMPLETE;
        mutex_exit(&xxx->event_mutex);
        rval = CS_SUCCESS;
        break;

    case CS_EVENT_CARD_READY:
        ASSERT(priority & CS_EVENT_PRI_LOW);
        rval = xx_card_ready(xxx);
        mutex_exit(&xxx->event_mutex);
        break;

    case CS_EVENT_CARD_INSERTION:
        ASSERT(priority & CS_EVENT_PRI_LOW);
        mutex_enter(&xxx->event_mutex);
        rval = xx_card_insertion(xxx);
        mutex_exit(&xxx->event_mutex);
        break;
    }
```

csx_event_handler(9E)

EXAMPLE 1 (Continued)

```

case CS_EVENT_CARD_REMOVAL:
    if (priority & CS_EVENT_PRI_HIGH) {
        mutex_enter(&xxx->hi_event_mutex);
        xxx->card_state &= ~XX_CARD_PRESENT;
        mutex_exit(&xxx->hi_event_mutex);
    } else {
        mutex_enter(&xxx->event_mutex);
        rval = xx_card_removal(xxx);
        mutex_exit(&xxx->event_mutex);
    }
    break;

case CS_EVENT_CLIENT_INFO:
    ASSERT(priority & CS_EVENT_PRI_LOW);
    if (GET_CLIENT_INFO_SUBSVC_CS(info->Attributes) ==
        CS_CLIENT_INFO_SUBSVC_CS) {
        info->Attributes |= CS_CLIENT_INFO_VALID;
        info->Revision = 4;
        info->CSLevel = CS_VERSION;
        info->RevDate = CS_CLIENT_INFO_MAKE_DATE(4, 7, 17);
        (void)strncpy(info->ClientName,
            "WhizBang Ultra Zowie PC card driver",
            CS_CLIENT_INFO_MAX_NAME_LEN)

        "ACME PC card drivers, Inc.",
            CS_CLIENT_INFO_MAX_NAME_LEN);
        rval = CS_SUCCESS;
    } else {
        rval = CS_UNSUPPORTED_EVENT;
    }
    break;

case CS_EVENT_WRITE_PROTECT:
    ASSERT(priority & CS_EVENT_PRI_LOW);
    mutex_enter(&xxx->event_mutex);
    if (args->info == CS_EVENT_WRITE_PROTECT_WPOFF) {
        xxx->card_state &= ~XX_WRITE_PROTECTED;
    } else {
        xxx->card_state |= XX_WRITE_PROTECTED;
    }
    mutex_exit(&xxx->event_mutex);
    rval = CS_SUCCESS;
    break;

default:
    rval = CS_UNSUPPORTED_EVENT;
    break;
}

return (rval);
}

```

csx_event_handler(9E)

EXAMPLE 1 *(Continued)*

SEE ALSO csx_Event2Text(9F), csx_RegisterClient(9F), csx_RequestSocketMask(9F)
PC Card 95 Standard, PCMCIA/JEIDA

detach(9E)

NAME	detach – detach or suspend a device
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int prefix detach(dev_info_t *dip, ddi_detach_cmd_t cmd);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>dip</i> A pointer to the device's dev_info structure. <i>cmd</i> Type of detach; the driver should return DDI_FAILURE if any value other than DDI_DETACH, DDI_PM_SUSPEND (obsolete), or DDI_SUSPEND is passed to it.
DESCRIPTION	The detach() function complements the attach(9E) routine.
DDI_DETACH	<p>If <i>cmd</i> is set to DDI_DETACH, detach() is used to remove the state associated with a given instance of a device node prior to the removal of that instance from the system.</p> <p>The detach() function will be called once for each instance of the device for which there has been a successful attach(), once there are no longer any opens on the device. An attached instance of a driver can be successfully detached only once. The detach() function should clean up any per instance data initialized in attach(9E) and call kmem_free(9F) to free any heap allocations. For information on how to unregister interrupt handlers, see ddi_add_intr(9F). This should also include putting the underlying device into a quiescent state so that it will not generate interrupts.</p> <p>Drivers that set up timeout(9F) routines should ensure that they are cancelled before returning DDI_SUCCESS from detach().</p> <p>If detach() determines a particular instance of the device cannot be removed when requested because of some exceptional condition, detach() must return DDI_FAILURE, which prevents the particular device instance from being detached. This also prevents the driver from being unloaded. A driver instance failing the detach must ensure that no per instance data or state is modified or freed that would compromise the system or subsequent driver operation.</p> <p>The system guarantees that the function will only be called for a particular dev_info node after (and not concurrently with) a successful attach(9E) of that device. The system also guarantees that detach() will only be called when there are no outstanding open(9E) calls on the device.</p>
DDI_PM_SUSPEND	The DDI_PM_SUSPEND command is required only if the device driver uses original Power Management interfaces (driver calls pm_create_components(9F)). This entry point is not needed if the device driver uses new automatic device Power Management interfaces (driver exports pm-components(9P) property instead of calling pm_create_components(9F)). The DDI_PM_SUSPEND command is obsolete and will be removed in a future release.

detach(9E)

The `DDI_PM_SUSPEND cmd` is used to suspend all activity of a device before power is possibly removed from the device by setting component 0 to power level 0. In this case, `detach()` may be called with outstanding `open(9E)` requests. It must save the hardware state of the device to memory and block incoming or existing requests until `attach(9E)` is called with a command value of `DDI_PM_RESUME`. When the device is suspended using `DDI_PM_SUSPEND` and it receives a request which requires device to be powered on, it should call `ddi_dev_is_needed(9F)` to request the framework to resume the device.

A return of `DDI_FAILURE` will result in component 0 of the device not being set to power level 0.

DDI_SUSPEND

The `DDI_SUSPEND cmd` is issued when the entire system is being suspended and power removed from it or when the system must be made quiescent. It will be issued only to devices which have a `reg` property or which export a `pm-hardware-state` property with the value `needs-suspend-resume`.

If `cmd` is set to `DDI_SUSPEND`, `detach()` is used to suspend all activity of a device before power is (possibly) removed from the device. The steps associated with suspension must include putting the underlying device into a quiescent state so that it will not generate interrupts or modify or access memory. Once quiescence has been obtained, `detach()` can be called with outstanding `open(9E)` requests. It must save the hardware state of the device to memory and block incoming or existing requests until `attach()` is called with `DDI_RESUME`.

If the device is used to store file systems, then after `DDI_SUSPEND` is issued, the device should still honor `dump(9E)` requests as this entry point may be used by suspend-resume operation (see `cpr(7)`) to save state file. It must do this, however, without disturbing the saved hardware state of the device.

If the device driver uses original Power Management interfaces (driver calls `pm_create_components(9F)`) and it has also been suspended by `DDI_PM_SUSPEND`, it will need to call `ddi_dev_is_needed(9F)` to honor the `dump(9E)` request. If the device driver uses new automatic device Power Management interfaces (driver exports `pm-components(9P)` property instead of calling `pm_create_components(9F)`, it might need to call `pm_raise_power(9F)` if the current power level is lower than required to complete the `dump(9E)` request.

Before returning successfully from a call to `detach()` with a command of `DDI_SUSPEND`, the driver must cancel any outstanding timeouts and make any driver threads quiescent.

If `DDI_FAILURE` is returned for the `DDI_SUSPEND cmd`, either the operation to suspend the system or to make it quiescent will be aborted.

RETURN VALUES

`DDI_SUCCESS` For `DDI_DETACH`, the state associated with the given device was successfully removed. For `DDI_SUSPEND` and `DDI_PM_SUSPEND` (obsolete), the driver was successfully suspended.

detach(9E)

DDI_FAILURE The operation failed or the request was not understood. The associated state is unchanged.

CONTEXT This function is called from user context only.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving (DDI_PM_SUSPEND is obsolete.)

SEE ALSO `cpr(7)`, `pm(7D)`, `pm(9P)`, `pm-components(9P)`, `attach(9E)`, `dump(9E)`, `open(9E)`, `power(9E)`, `ddi_add_intr(9F)`, `ddi_dev_is_needed(9F)`, `ddi_map_regs(9F)`, `kmem_free(9F)`, `pm_create_components(9F)`, `pm_raise_power(9F)`, `timeout(9F)`

Writing Device Drivers

NAME	devmap – validate and translate virtual mapping for memory mapped device
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>devmap(dev_t dev, devmap_cookie_t dhp, offset_t off, size_t len, size_t *maplen, uint_t model);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris DDI specific (Solaris DDI).</p> <p><i>dev</i> Device whose memory is to be mapped.</p> <p><i>dhp</i> An opaque mapping handle that the system uses to describe the mapping.</p> <p><i>off</i> User offset within the logical device memory at which the mapping begins.</p> <p><i>len</i> Length (in bytes) of the mapping to be mapped.</p> <p><i>maplen</i> Pointer to length (in bytes) of mapping that has been validated. <i>maplen</i> is less than or equal to <i>len</i>.</p> <p><i>model</i> The data model type of the current thread.</p>
DESCRIPTION	<p>devmap () is a required entry point for character drivers supporting memory-mapped devices if the drivers use the devmap framework to set up the mapping. A memory mapped device has memory that can be mapped into a process's address space. The mmap(2) system call, when applied to a character special file, allows this device memory to be mapped into user space for direct access by the user applications.</p> <p>As a result of a mmap(2) system call, the system calls the devmap () entry point during the mapping setup when D_DEVMAP is set in the cb_flag field of the cb_ops(9S) structure, and any of the following conditions apply:</p> <ul style="list-style-type: none"> ■ ddi_devmap_segmap(9F) is used as the segmap(9E) entry point. ■ segmap(9E) entry point is set to NULL. ■ mmap(9E) entry point is set to NULL. <p>Otherwise EINVAL will be returned to mmap(2).</p> <p>Device drivers should use devmap () to validate the user mappings to the device, to translate the logical offset, <i>off</i>, to the corresponding physical offset within the device address space, and to pass the mapping information to the system for setting up the mapping.</p> <p><i>dhp</i> is a device mapping handle that the system uses to describe a mapping to a memory that is either contiguous in physical address space or in kernel virtual address space. The system may create multiple mapping handles in one mmap(2) system call (for example, if the mapping contains multiple physically discontinuous memory regions).</p> <p><i>model</i> returns the C Language Type Model which the current thread expects. It is set to DDI_MODEL_ILP32 if the current thread expects 32-bit (<i>ILP32</i>) semantics, or DDI_MODEL_LP64 if the current thread expects 64-bit (<i>LP64</i>) semantics. <i>model</i> is used</p>

devmap(9E)

in combination with `ddi_model_convert_from(9F)` to determine whether there is a data model mismatch between the current thread and the device driver. The device driver might have to adjust the shape of data structures before exporting them to a user thread which supports a different data model.

`devmap()` should return `EINVAL` if the logical offset, *off*, is out of the range of memory exported by the device to user space. If *off + len* exceeds the range of the contiguous memory, `devmap()` should return the length from *off* to the end of the contiguous memory region. The system will repeatedly call `devmap()` until the original mapping length is satisfied. The driver sets **maplen* to the validated length which must be either less than or equal to *len*.

The `devmap()` entry point must initialize the mapping parameters before passing them to the system through either `devmap_devmem_setup(9F)` (if the memory being mapped is device memory) or `devmap_umem_setup(9F)` (if the memory being mapped is kernel memory). The `devmap()` entry point initializes the mapping parameters by mapping the control callback structure (see `devmap_callback_ctl(9S)`), the device access attributes, mapping length, maximum protection possible for the mapping, and optional mapping flags. See `devmap_devmem_setup(9F)` and `devmap_umem_setup(9F)` for further information on initializing the mapping parameters.

The system will copy the driver's `devmap_callback_ctl(9S)` data into its private memory so the drivers do not need to keep the data structure after the return from either `devmap_devmem_setup(9F)` or `devmap_umem_setup(9F)`.

For device mappings, the system establishes the mapping to the physical address that corresponds to *off* by passing the register number and the offset within the register address space to `devmap_devmem_setup(9F)`.

For kernel memory mapping, the system selects a user virtual address that is aligned with the kernel address being mapped for cache coherence.

RETURN VALUES

0	Successful completion.
Non-zero	An error occurred.

EXAMPLES

EXAMPLE 1 Implementing the `devmap()` Entry Point

The following is an example of the implementation for the `devmap()` entry point. For mapping device memory, `devmap()` calls `devmap_devmem_setup(9F)` with the register number, *rnumber*, and the offset within the register, *roff*. For mapping kernel memory, the driver must first allocate the kernel memory using `ddi_umem_alloc(9F)`. For example, `ddi_umem_alloc(9F)` can be called in the `attach(9E)` routine. The resulting kernel memory cookie is stored in the driver soft state structure, which is accessible from the `devmap()` entry point. See `ddi_soft_state(9F)`. `devmap()` passes the cookie obtained from

EXAMPLE 1 Implementing the devmap() Entry Point (Continued)

ddi_umem_alloc(9F) and the offset within the allocated kernel memory to devmap_umem_setup(9F). The corresponding ddi_umem_free(9F) can be made in the detach(9E) routine to free up the kernel memory.

```

. . .
#define MAPPING_SIZE 0x2000      /* size of the mapping */
#define MAPPING_START 0x70000000 /* logical offset at beginning
                                   of the mapping */

static
struct devmap_callback_ctl xxmap_ops = {
    DEVMAP_OPS_REV,             /* devmap_ops version number */
    xxmap_map,                  /* devmap_ops map routine */
    xxmap_access,              /* devmap_ops access routine */
    xxmap_dup,                 /* devmap_ops dup routine */
    xxmap_unmap,               /* devmap_ops unmap routine */
};

static int
xxdevmap(dev_t dev, devmap_cookie_t dhp, offset_t off, size_t len,
         size_t *maplen, uint_t model)
{
    int instance;
    struct xxstate *xsp;
    struct ddi_device_acc_attr *endian_attr;
    struct devmap_callback_ctl *callbackops = NULL;
    ddi_umem_cookie_t cookie;
    dev_info_t *dip;
    offset_t roff;
    offset_t koff;
    uint_t rnumber;
    uint_t maxprot;
    uint_t flags = 0;
    size_t length;
    int err;

    /* get device soft state */
    instance = getminor(dev);
    xsp = ddi_get_soft_state(statep, instance);
    if (xsp == NULL)
        return (-1);

    dip = xsp->dip;
    /* check for a valid offset */
    if ( off is invalid )
        return (-1);
    /* check if len is within the range of contiguous memory */
    if ( (off + len) is contiguous.)
        length = len;
    else
        length = MAPPING_START + MAPPING_SIZE - off;

    /* device access attributes */
    endian_attr = xsp->endian_attr;

```

devmap(9E)

EXAMPLE 1 Implementing the devmap() Entry Point (Continued)

```
if ( off is referring to a device memory. ) {
    /* assign register related parameters */
    rnumber = XXX;          /* index to register set at off */
    roff = XXX;            /* offset of rnumber at local bus */
    callbackops = &xxmap_ops; /* do all callbacks for this mapping */
    maxprot = PROT_ALL;    /* allowing all access */
    if ((err = devmap_devmem_setup(dhp, dip, callbackops, rnumber, roff,
        length, maxprot, flags, endian_attr)) < 0)

        return (err);

} else if ( off is referring to a kernel memory.) {
    cookie = xsp->cookie;   /* cookie is obtained from
                            ddi_umem_alloc(9F) */
    koff = XXX;            /* offset within the kernel memory. */
    callbackops = NULL;    /* don't do callback for this mapping */
    maxprot = PROT_ALL;    /* allowing all access */
    if ((err = devmap_umem_setup(dhp, dip, callbackops, cookie, koff,
        length, maxprot, flags, endian_attr)) < 0)
        return (err);
}

*maplen = length;
return (0);
}
```

SEE ALSO mmap(2), attach(9E), detach(9E), mmap(9E), segmap(9E),
ddi_devmap_segmap(9F), ddi_model_convert_from(9F), ddi_soft_state(9F),
ddi_umem_alloc(9F), ddi_umem_free(9F), devmap_devmem_setup(9F),
devmap_setup(9F), devmap_umem_setup(9F), cb_ops(9S),
devmap_callback_ctl(9S)

Writing Device Drivers

NAME	devmap_access – device mapping access entry point
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>devmap_access(devmap_cookie_t <i>dhp</i>, void *<i>pvt</i>, offset_t <i>off</i>, size_t <i>len</i>, uint_t <i>type</i>, uint_t <i>rw</i>);</pre>
INTERFACE LEVEL ARGUMENTS	<p>Solaris DDI specific (Solaris DDI).</p> <p><i>dhp</i> An opaque mapping handle that the system uses to describe the mapping.</p> <p><i>pvt</i> Driver private mapping data.</p> <p><i>off</i> User offset within the logical device memory at which the access begins.</p> <p><i>len</i> Length (in bytes) of the memory being accessed.</p> <p><i>type</i> Type of access operation. Possible values are:</p> <p> DEVMAP_ACCESS Memory access.</p> <p> DEVMAP_LOCK Lock the memory being accessed.</p> <p> DEVMAP_UNLOCK Unlock the memory being accessed.</p> <p><i>rw</i> Direction of access. Possible values are:</p> <p> DEVMAP_READ Read access attempted.</p> <p> DEVMAP_WRITE Write access attempted.</p> <p> DEVMAP_EXEC Execution access attempted.</p>
DESCRIPTION	<p>The devmap_access() entry point is an optional routine. It notifies drivers whenever an access is made to a mapping described by <i>dhp</i> that has not been validated or does not have sufficient protection for the access. The system expects devmap_access() to call either devmap_do_ctxmgt(9F) or devmap_default_access(9F) to load the memory address translations before it returns. For mappings that support context switching, device drivers should call devmap_do_ctxmgt(9F). For mappings that do not support context switching, the drivers should call devmap_default_access(9F).</p> <p>In devmap_access(), drivers perform memory access related operations such as context switching, checking the availability of the memory object, and locking and unlocking the memory object being accessed. The devmap_access() entry point is set to NULL if no operations need to be performed.</p> <p><i>pvt</i> is a pointer to the driver's private mapping data that was allocated and initialized in the devmap_map(9E) entry point.</p>

devmap_access(9E)

off and *len* define the range to be affected by the operations in `devmap_access()`. *type* defines the type of operation that device drivers should perform on the memory object. If *type* is either `DEVMAP_LOCK` or `DEVMAP_UNLOCK`, the length passed to either `devmap_do_ctxmgt(9F)` or `devmap_default_access(9F)` must be same as *len*. *rw* specifies the direction of access on the memory object.

A non-zero return value from `devmap_access()` may result in a `SIGSEGV` or `SIGBUS` signal being delivered to the process.

RETURN VALUES

`devmap_access()` returns the following values:

0	Successful completion.
Non-zero	An error occurred. The return value from <code>devmap_do_ctxmgt(9F)</code> or <code>devmap_default_access(9F)</code> should be returned.

EXAMPLES

EXAMPLE 1 `devmap_access()` entry point

The following is an example of the `devmap_access()` entry point. If the mapping supports context switching, `devmap_access()` calls `devmap_do_ctxmgt(9F)`. Otherwise, `devmap_access()` calls `devmap_default_access(9F)`.

```
. . .
#define OFF_DO_CTXMGT  0x40000000
#define OFF_NORMAL    0x40100000
#define CTXMGT_SIZE   0x100000
#define NORMAL_SIZE   0x100000

/*
 * Driver devmap_contextmgt(9E) callback function.
 */
static int
xx_context_mgt(devmap_cookie_t dhp, void *pvtp, offset_t offset,
               size_t length, uint_t type, uint_t rw)
{
    . . . . .
    /*
     * see devmap_contextmgt(9E) for an example
     */
}

/*
 * Driver devmap_access(9E) entry point
 */
static int
xxdevmap_access(devmap_cookie_t dhp, void *pvtp, offset_t off,
                size_t len, uint_t type, uint_t rw)
{
    offset_t diff;
    int err;

    /*
     * check if off is within the range that supports
     * context management.
     */
}
```


EXAMPLE 1 devmap_access() entry point *(Continued)*

```

*/
if ((diff = off - OFF_DOCTXMG) >= 0 && diff < CTXMGT_SIZE) {
/*
 * calculates the length for context switching
 */
if ((len + off) > (OFF_DOCTXMG + CTXMGT_SIZE))
return (-1);
/*
 * perform context switching
 */
err = devmap_do_ctxmgt(dhp, pvtp, off, len, type,
rw, xx_context_mgt);
/*
 * check if off is within the range that does normal
 * memory mapping.
 */
} else if ((diff = off - OFF_NORMAL) >= 0 && diff < NORMAL_SIZE) {
if ((len + off) > (OFF_NORMAL + NORMAL_SIZE))
return (-1);
err = devmap_default_access(dhp, pvtp, off, len, type, rw);
} else
return (-1);

return (err);
}

```

SEE ALSO devmap_map(9E), devmap_default_access(9F), devmap_do_ctxmgt(9F), devmap_callback_ctl(9S)

Writing Device Drivers

devmap_contextmgt(9E)

NAME	devmap_contextmgt – driver callback function for context management
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int devmap_contextmgt (devmap_cookie_t dhp, void *pvtp, offset_t off, size_t len, uint_t type, uint_t rw);</pre>
INTERFACE LEVEL ARGUMENTS	<p>Solaris DDI specific (Solaris DDI).</p> <p><i>dhp</i> An opaque mapping handle that the system uses to describe the mapping.</p> <p><i>pvtp</i> Driver private mapping data.</p> <p><i>off</i> User offset within the logical device memory at which the access begins.</p> <p><i>len</i> Length (in bytes) of the memory being accessed.</p> <p><i>type</i> Type of access operation. Possible values are:</p> <ul style="list-style-type: none">DEVMAP_ACCESS Memory access.DEVMAP_LOCK Lock the memory being accessed.DEVMAP_UNLOCK Unlock the memory being accessed. <p><i>rw</i> Direction of access. Possible values are:</p> <ul style="list-style-type: none">DEVMAP_READ Read access attempted.DEVMAP_WRITE Write access attempted.
DESCRIPTION	<p>devmap_contextmgt () is a driver-supplied function that performs device context switching on a mapping. Device drivers pass devmap_contextmgt () as an argument to devmap_do_ctxmgt(9F) in the devmap_access(9E) entry point. The system will call devmap_contextmgt () when memory is accessed. The system expects devmap_contextmgt () to load the memory address translations of the mapping by calling devmap_load(9F) before returning.</p> <p><i>dhp</i> uniquely identifies the mapping and is used as an argument to devmap_load(9F) to validate the mapping. <i>off</i> and <i>len</i> define the range to be affected by the operations in devmap_contextmgt ().</p> <p>The driver must check if there is already a mapping established at <i>off</i> that needs to be unloaded. If a mapping exists at <i>off</i>, devmap_contextmgt () must call devmap_unload(9F) on the current mapping. devmap_unload(9F) must be followed by devmap_load () on the mapping that generated this call to devmap_contextmgt (). devmap_unload(9F) unloads the current mapping so that a call to devmap_access(9E), which causes the system to call devmap_contextmgt (), will be generated the next time the mapping is accessed.</p>

pvtp is a pointer to the driver's private mapping data that was allocated and initialized in the `devmap_map(9E)` entry point. *type* defines the type of operation that device drivers should perform on the memory object. If *type* is either `DEVMAP_LOCK` or `DEVMAP_UNLOCK`, the length passed to either `devmap_unload(9F)` or `devmap_load(9F)` must be same as *len*. *rw* specifies the access direction on the memory object.

A non-zero return value from `devmap_contextmgt()` will be returned to `devmap_access(9E)` and will cause the corresponding operation to fail. The failure may result in a `SIGSEGV` or `SIGBUS` signal being delivered to the process.

RETURN VALUES

0	Successful completion.
Non-zero	An error occurred.

EXAMPLES **EXAMPLE 1** managing a device context

The following shows an example of managing a device context.

```

struct xxcontext cur_ctx;
static int
xxdevmap_contextmgt(devmap_cookie_t dhp, void *pvtp, offset_t off,
                    size_t len, uint_t type, uint_t rw)
{
    devmap_cookie_t cur_dhp;
    struct xxpvtdata *p;
    struct xxpvtdata *pvp = (struct xxpvtdata *)pvtp;
    struct xx_softc *softc = pvp->softc;
    int err;

    mutex_enter(&softc->mutex);

    /*
     * invalidate the translations of current context before
     * switching context.
     */
    if (cur_ctx != NULL && cur_ctx != pvp->ctx) {
        p = cur_ctx->pvt;
        cur_dhp = p->dhp;
        if ((err = devmap_unload(cur_dhp, off, len)) != 0)
            return (err);
    }
    /* Switch device context - device dependent */
    ...
    /* Make handle the new current mapping */
    cur_ctx = pvp->ctx;

    /*
     * Load the address translations of the calling context.
     */
    err = devmap_load(pvp->dhp, off, len, type, rw);

    mutex_exit(&softc->mutex);
}

```

devmap_contextmgt(9E)

EXAMPLE 1 managing a device context *(Continued)*

```
    return (err);  
}
```

SEE ALSO devmap_access(9E), devmap_do_ctxmgt(9F) devmap_load(9F),
devmap_unload(9F)

Writing Device Drivers

NAME	devmap_dup – device mapping duplication entry point
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>devmap_dup(devmap_cookie_t dhp, void *pvtp, devmap_cookie_t new_dhp, void **new_pvtp);</pre>
INTERFACE LEVEL ARGUMENTS	<p>Solaris DDI specific (Solaris DDI).</p> <p><i>dhp</i> An opaque mapping handle that the system uses to describe the mapping currently being duplicated.</p> <p><i>pvtp</i> Driver private mapping data for the mapping currently being duplicated.</p> <p><i>new_dhp</i> An opaque data structure that the system uses to describe the duplicated device mapping.</p> <p><i>new_pvtp</i> A pointer to be filled in by device drivers with the driver private mapping data for the duplicated device mapping.</p>
DESCRIPTION	<p>The system calls devmap_dup () when a device mapping is duplicated, such as during the execution of the fork(2) system call. The system expects devmap_dup () to generate new driver private data for the new mapping, and to set new_pvtp to point to it. new_dhp is the handle of the new mapped object.</p> <p>A non-zero return value from devmap_dup () will cause a corresponding operation such as fork () to fail.</p>
RETURN VALUES	<p>devmap_dup () returns the following values:</p> <p>0 Successful completion.</p> <p>Non-zero An error occurred.</p>
EXAMPLES	<p>EXAMPLE 1</p> <pre>static int xxdevmap_dup(devmap_cookie_t dhp, void *pvtp, devmap_cookie_t new_dhp, void **new_pvtp) { struct xpvtdata *prvtdata; struct xpvtdata *p = (struct xpvtdata *)pvtp; struct xx_softc *softc = p->softc; mutex_enter(&softc->mutex); /* Allocate a new private data structure */ prvtdata = kmem_alloc(sizeof (struct xpvtdata), KM_SLEEP); /* Return the new data */ prvtdata->off = p->off; prvtdata->len = p->len; prvtdata->ctx = p->ctx; prvtdata->dhp = new_dhp; prvtdata->softc = p->softc; *new_pvtp = prvtdata;</pre>

devmap_dup(9E)

EXAMPLE 1 *(Continued)*

```
    mutex_exit(&softc->mutex);  
    return (0);  
}
```

SEE ALSO fork(2), devmap_callback_ctl(9S)

Writing Device Drivers

NAME	devmap_map – device mapping create entry point
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>devmap_map(devmap_cookie_t <i>dhp</i>, dev_t <i>dev</i>, uint_t <i>flags</i>, offset_t <i>off</i>, size_t <i>len</i>, void **<i>pvtp</i>);</pre>
INTERFACE LEVEL ARGUMENTS	<p>Solaris DDI specific (Solaris DDI).</p> <p><i>dhp</i> An opaque mapping handle that the system uses to describe the mapping currently being created.</p> <p><i>dev</i> The device whose memory is to be mapped.</p> <p><i>flags</i> Flags indicating type of mapping. Possible values are:</p> <p> MAP_PRIVATE Changes are private.</p> <p> MAP_SHARED Changes should be shared.</p> <p><i>off</i> User offset within the logical device memory at which the mapping begins.</p> <p><i>len</i> Length (in bytes) of the memory to be mapped.</p> <p><i>pvtp</i> A pointer to be filled in by device drivers with the driver private mapping data.</p>
DESCRIPTION	<p>The devmap_map () entry point is an optional routine that allows drivers to perform additional processing or to allocate private resources during the mapping setup time. For example, in order for device drivers to support context switching, the drivers allocate private mapping data and associate the private data with the mapping parameters in the devmap_map () entry point.</p> <p>The system calls devmap_map () after the user mapping to device physical memory has been established. (For example, after the devmap(9E) entry point is called.)</p> <p>devmap_map () receives a pointer to the driver private data for this mapping in <i>pvtp</i>. The system expects the driver to allocate its private data and set *<i>pvtp</i> to the allocated data. The driver must store <i>off</i> and <i>len</i>, which define the range of the mapping, in its private data. Later, when the system calls devmap_unmap(9E), the driver will use the <i>off</i> and <i>len</i> stored in <i>pvtp</i> to check if the entire mapping, or just a part of it, is being unmapped. If only a part of the mapping is being unmapped, the driver must allocate a new private data for the remaining mapping before freeing the old private data. The driver will receive *<i>pvtp</i> in subsequent event notification callbacks.</p> <p>If the driver support context switching, it should store the mapping handle <i>dhp</i> in its private data *<i>pvtp</i> for later use in devmap_unload(9F).</p> <p>For a driver that supports context switching, <i>flags</i> indicates whether or not the driver should allocate a private context for the mapping. For example, a driver may allocate a memory region to store the device context if <i>flags</i> is set to MAP_PRIVATE.</p>

devmap_map(9E)

RETURN VALUES

devmap_map () returns the following values:

0 Successful completion.
Non-zero An error occurred.

EXAMPLES

EXAMPLE 1 devmap_map () implementation

The following shows an example implementation for devmap_map ().

```
static int
xxdevmap_map(devmap_cookie_t dhp, dev_t dev, uint_t flags, offset_t off,
             size_t len, void **pvtp)
{
    struct xx_resources *pvt;
    struct xx_context *this_context;
    struct xx_softc *softc;
    softc = ddi_get_softc(statep, getminor(dev));

    this_context = get_context(softc, off, len);

    /* allocate resources for the mapping - Device dependent */
    pvt = kmem_zalloc(sizeof (struct xx_resources), KM_SLEEP);

    pvt->off = off;
    pvt->len = len;
    pvt->dhp = dhp;
    pvt->ctx = this_context;
    *pvtp = pvt;
}
```

SEE ALSO

devmap_unmap(9E), devmap_unload(9F), devmap_callback_ctl(9S)

Writing Device Drivers

NAME	devmap_unmap – device mapping unmap entry point
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> void prefixdevmap_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);</pre>
INTERFACE LEVEL ARGUMENTS	<p>Solaris DDI specific (Solaris DDI).</p> <p><i>dhp</i> An opaque mapping handle that the system uses to describe the mapping.</p> <p><i>pvtp</i> Driver private mapping data.</p> <p><i>off</i> User offset within the logical device memory at which the unmapping begins.</p> <p><i>len</i> Length (in bytes) of the memory being unmapped.</p> <p><i>new_dhp1</i> The opaque mapping handle that the system uses to describe the new region that ends at (<i>off</i> - 1). <i>new_dhp1</i> may be NULL.</p> <p><i>new_pvtp1</i> A pointer to be filled in by the driver with the driver private mapping data for the new region that ends at (<i>off</i> - 1); ignored if <i>new_dhp1</i> is NULL.</p> <p><i>new_dhp2</i> The opaque mapping handle that the system uses to describe the new region that begins at (<i>off</i> + <i>len</i>); <i>new_dhp2</i> may be NULL.</p> <p><i>new_pvtp2</i> A pointer to be filled in by the driver with the driver private mapping data for the new region that begins at (<i>off</i> + <i>len</i>); ignored if <i>new_dhp2</i> is NULL.</p>
DESCRIPTION	<p>devmap_unmap () is called when the system removes the mapping in the range [<i>off</i>, <i>off</i> + <i>len</i>], such as in the munmap(2) or exit(2) system calls. Device drivers use devmap_unmap () to free up the resources allocated in devmap_map(9E).</p> <p><i>dhp</i> is the mapping handle that uniquely identifies the mapping. The driver stores the mapping attributes in the driver's private data, <i>pvtp</i>, when the mapping is created. See devmap_map(9E) for details.</p> <p><i>off</i> and <i>len</i> define the range to be affected by devmap_unmap (). This range is within the boundary of the mapping described by <i>dhp</i>.</p> <p>If the range [<i>off</i>, <i>off</i> + <i>len</i>] covers the entire mapping, the system passes NULL to <i>new_dhp1</i>, <i>new_pvtp1</i>, <i>new_dhp2</i>, and <i>new_pvtp2</i>. The system expects device drivers to free all resources allocated for this mapping.</p>

devmap_unmap(9E)

If *off* is at the beginning of the mapping and *len* does not cover the entire mapping, the system sets `NULL` to *new_dhp1* and to *new_pvtp1*. The system expects the drivers to allocate new driver private data for the region that starts at *off + len* and to set **new_pvtp2* to point to it. *new_dhp2* is the mapping handle of the newly mapped object.

If *off* is not at the beginning of the mapping, but *off + len* is at the end of the mapping the system passes `NULL` to *new_dhp2* and *new_pvtp2*. The system then expects the drivers to allocate new driver private data for the region that begins at the beginning of the mapping (for example, stored in *pvtp*) and to set **new_pvtp1* to point to it. *new_dhp1* is the mapping handle of the newly mapped object.

The drivers should free up the driver private data, *pvtp*, previously allocated in `devmap_map(9E)` before returning to the system.

EXAMPLES

EXAMPLE 1 devmap_unmap() implementation

```
static void
xxdevmap_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)
{
    struct xxpvtdata *ptmp;
    struct xxpvtdata *p = (struct xxpvtdata *)pvtp;
    struct xx_softc *softc = p->softc;
    mutex_enter(&softc->mutex);
    /*
     * If new_dhp1 is not NULL, create a new driver private data
     * for the region from the beginning of old mapping to off.
     */
    if (new_dhp1 != NULL) {
        ptmp = kmem_zalloc(sizeof (struct xxpvtdata), KM_SLEEP);
        ptmp->dhp = new_dhp1;
        ptmp->off = pvtp->off;
        ptmp->len = off - pvtp->off;
        *new_pvtp1 = ptmp;
    }

    /*
     * If new_dhp2 is not NULL, create a new driver private data
     * for the region from off+len to the end of the old mapping.
     */
    if (new_dhp2 != NULL) {
        ptmp = kmem_zalloc(sizeof (struct xxpvtdata), KM_SLEEP);
        ptmp->off = off + len;
        ptmp->len = pvtp->len - (off + len - pvtp->off);
        ptmp->dhp = new_dhp2;
        *new_pvtp2 = ptmp;
    }

    /* Destroy the driver private data - Device dependent */
    ...
    kmem_free(pvtp, sizeof (struct xxpvtdata));
    mutex_exit(&softc->mutex);
}
```

devmap_unmap(9E)

EXAMPLE 1 devmap_unmap() implementation (Continued)

SEE ALSO exit(2), munmap(2), devmap_map(9E), devmap_callback_ctl(9S)

Writing Device Drivers

dump(9E)

NAME	dump – dump memory to device during system failure
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> int prefixdump (dev_t dev, caddr_t addr, daddr_t blkno, int nblk);</pre>
INTERFACE LEVEL	Solaris specific (Solaris DDI). This entry point is required. For drivers that do not implement dump () routines, nodev(9F) should be used.
ARGUMENTS	<p><i>dev</i> Device number.</p> <p><i>addr</i> Address for the beginning of the area to be dumped.</p> <p><i>blkno</i> Block offset to dump memory.</p> <p><i>nblk</i> Number of blocks to dump.</p>
DESCRIPTION	<p>dump () is used to dump a portion of virtual address space directly to a device in the case of system failure. It can also be used for checking the state of the kernel during a checkpoint operation. The memory area to be dumped is specified by <i>addr</i> (base address) and <i>nblk</i> (length). It is dumped to the device specified by <i>dev</i> starting at offset <i>blkno</i>. Upon completion dump () returns the status of the transfer.</p> <p>When the system is panicking, the calls of functions scheduled by timeout(9F) and ddi_trigger_softintr(9F) will never occur. Neither can delay(9F) be relied upon, since it is implemented via timeout (). See ddi_in_panic(9F).</p> <p>dump () is called at interrupt priority.</p>
RETURN VALUES	dump () returns 0 on success, or the appropriate error number.
SEE ALSO	<p>cpr(7), nodev(9F)</p> <p><i>Writing Device Drivers</i></p>

NAME	<code>_fini</code> , <code>_info</code> , <code>_init</code> – loadable module configuration entry points
SYNOPSIS	<pre>#include <sys/modctl.h> int _fini(void); int _info(struct modinfo *modinfop); int _init(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). These entry points are required. You must write them.
_info()	<i>modinfop</i> A pointer to an opaque <code>modinfo</code> structure.
DESCRIPTION	<p><code>_init()</code> initializes a loadable module. It is called before any other routine in a loadable module. <code>_init()</code> returns the value returned by <code>mod_install(9F)</code>. The module may optionally perform some other work before the <code>mod_install(9F)</code> call is performed. If the module has done some setup before the <code>mod_install(9F)</code> function is called, then it should be prepared to undo that setup if <code>mod_install(9F)</code> returns an error.</p> <p><code>_info()</code> returns information about a loadable module. <code>_info()</code> returns the value returned by <code>mod_info(9F)</code>.</p> <p><code>_fini()</code> prepares a loadable module for unloading. It is called when the system wants to unload a module. If the module determines that it can be unloaded, then <code>_fini()</code> returns the value returned by <code>mod_remove(9F)</code>. Upon successful return from <code>_fini()</code> no other routine in the module will be called before <code>_init()</code> is called.</p>
RETURN VALUES	<p><code>_init()</code> should return the appropriate error number if there is an error, otherwise it should return the return value from <code>mod_install(9F)</code>.</p> <p><code>_info()</code> should return the return value from <code>mod_info(9F)</code></p> <p><code>_fini()</code> should return the return value from <code>mod_remove(9F)</code>. <code>_fini()</code> is permitted to return <code>EBUSY</code> prior to calling <code>mod_remove(9F)</code> if the driver should not be unloaded. Driver global resources, such as mutexes and calls to <code>ddi_soft_state_fini(9F)</code>, should only be destroyed in <code>_fini()</code> after <code>mod_remove()</code> returns successfully.</p>
EXAMPLES	<p>EXAMPLE 1 Initializing and Freeing a Mutex</p> <p>The following example demonstrates how to initialize and free a <code>mutex(9F)</code>.</p> <pre>#include <sys/modctl.h> #include <sys/ddi.h> #include <sys/sunddi.h> static struct dev_ops drv_ops; /* * Module linkage information for the kernel. */ static struct modldrv modldrv = {</pre>

`_fini(9E)`

EXAMPLE 1 Initializing and Freeing a Mutex (Continued)

```
    &mod_driverops,    /* Type of module. This one is a driver */
    "Sample Driver",
    &drv_ops           /* driver ops */
};

static struct modlinkage modlinkage = {
    MODREV_1,
    &modldrv,
    NULL
};

/*
 * Global driver mutex
 */
static kmutex_t    xx_global_mutex;

int
_init(void)
{
    int    i;

    /*
     * Initialize global mutex before mod_install'ing driver.
     * If mod_install() fails, must clean up mutex initialization
     */
    mutex_init(&xx_global_mutex, NULL,
               MUTEX_DRIVER, (void *)NULL);

    if ((i = mod_install(&modlinkage)) != 0) {
        mutex_destroy(&xx_global_mutex);
    }

    return (i);
}

int
_info(struct modinfo *modinfop)
{
    return (mod_info(&modlinkage, modinfop));
}

int
_fini(void)
{
    int    i;

    /*
     * If mod_remove() is successful, we destroy our global mutex
     */
    if ((i = mod_remove(&modlinkage)) == 0) {
        mutex_destroy(&xx_global_mutex);
    }
}
```

EXAMPLE 1 Initializing and Freeing a Mutex (Continued)

```
    }  
    return (i);  
}
```

SEE ALSO [add_drv\(1M\)](#), [mod_info\(9F\)](#), [mod_install\(9F\)](#), [mod_remove\(9F\)](#), [mutex\(9F\)](#), [modldrv\(9S\)](#), [modlinkage\(9S\)](#), [modlstrmod\(9S\)](#)

Writing Device Drivers

WARNINGS Do not change the structures referred to by the `modlinkage` structure after the call to `mod_install()`, as the system may copy or change them.

NOTES Even though the identifiers `_fini()`, `_info()`, and `_init()` appear to be declared as globals, their scope is restricted by the kernel to the module that they are defined in.

BUGS On some implementations `_info()` may be called before `_init()`.

getinfo(9E)

NAME	getinfo – get device driver information								
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg, void **resultp);</pre>								
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). This entry point is required for drivers which export <code>cb_ops(9S)</code> entry points.								
ARGUMENTS	<table><tr><td><i>dip</i></td><td>Do not use.</td></tr><tr><td><i>cmd</i></td><td>Command argument – valid command values are <code>DDI_INFO_DEVT2DEVINFO</code> and <code>DDI_INFO_DEVT2INSTANCE</code>.</td></tr><tr><td><i>arg</i></td><td>Command specific argument.</td></tr><tr><td><i>resultp</i></td><td>Pointer to where the requested information is stored.</td></tr></table>	<i>dip</i>	Do not use.	<i>cmd</i>	Command argument – valid command values are <code>DDI_INFO_DEVT2DEVINFO</code> and <code>DDI_INFO_DEVT2INSTANCE</code> .	<i>arg</i>	Command specific argument.	<i>resultp</i>	Pointer to where the requested information is stored.
<i>dip</i>	Do not use.								
<i>cmd</i>	Command argument – valid command values are <code>DDI_INFO_DEVT2DEVINFO</code> and <code>DDI_INFO_DEVT2INSTANCE</code> .								
<i>arg</i>	Command specific argument.								
<i>resultp</i>	Pointer to where the requested information is stored.								
DESCRIPTION	<p>When <i>cmd</i> is set to <code>DDI_INFO_DEVT2DEVINFO</code>, <code>getinfo()</code> should return the <code>dev_info_t</code> pointer associated with the <code>dev_t arg</code>. The <code>dev_info_t</code> pointer should be returned in the field pointed to by <i>resultp</i>.</p> <p>When <i>cmd</i> is set to <code>DDI_INFO_DEVT2INSTANCE</code>, <code>getinfo()</code> should return the instance number associated with the <code>dev_t arg</code>. The instance number should be returned in the field pointed to by <i>resultp</i>.</p> <p>Drivers which do not export <code>cb_ops(9S)</code> entry points are not required to provide a <code>getinfo()</code> entry point, and may use <code>nodev(9F)</code> in the <code>devo_getinfo</code> field of the <code>dev_ops(9S)</code> structure. A SCSI HBA driver is an example of a driver which is not required to provide <code>cb_ops(9S)</code> entry points.</p>								
RETURN VALUES	<p><code>getinfo()</code> should return:</p> <table><tr><td><code>DDI_SUCCESS</code></td><td>on success.</td></tr><tr><td><code>DDI_FAILURE</code></td><td>on failure.</td></tr></table>	<code>DDI_SUCCESS</code>	on success.	<code>DDI_FAILURE</code>	on failure.				
<code>DDI_SUCCESS</code>	on success.								
<code>DDI_FAILURE</code>	on failure.								
EXAMPLES	<p>EXAMPLE 1 <code>getinfo()</code> implementation</p> <pre>/*ARGSUSED*/ static int rd_getinfo(dev_info_t *dip, ddi_info_cmd_t infocmd, void *arg, void **resultp) { /* Note that in this simple example * the minor number is the instance * number. */ devstate_t *sp; int error = DDI_FAILURE; switch (infocmd) { case DDI_INFO_DEVT2DEVINFO: if ((sp = ddi_get_soft_state(statep, getminor((dev_t) arg))) != NULL) {</pre>								

EXAMPLE 1 getinfo() implementation (Continued)

```
        *resultp = sp->devi;
        error = DDI_SUCCESS;
    } else
        *result = NULL;
    break;

case DDI_INFO_DEVT2INSTANCE:
    *resultp = (void *) getminor((dev_t) arg);
    error = DDI_SUCCESS;
    break;
}

return (error);
}
```

SEE ALSO nodev(9F), cb_ops(9S), dev_ops(9S)

Writing Device Drivers

gld(9E)

NAME	gld, gldm_reset, gldm_start, gldm_stop, gldm_set_mac_addr, gldm_set_multicast, gldm_set_promiscuous, gldm_send, gldm_intr, gldm_get_stats, gldm_ioctl – Generic LAN Driver entry points														
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>														
INTERFACE LEVEL PARAMETERS	Solaris architecture specific (Solaris DDI). <table><tr><td><i>macinfo</i></td><td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>macaddr</i></td><td>Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multicastaddr</i></td><td>Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multiflag</i></td><td>A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</td></tr><tr><td><i>promiscflag</i></td><td>A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</td></tr><tr><td><i>mp</i></td><td>Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</td></tr><tr><td><i>stats</i></td><td>Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</td></tr></table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .	<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .	<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.	<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.														
<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .														
<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .														
<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.														
<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.														

DESCRIPTION	<p><i>q</i> Pointer to the <code>queue(9S)</code> structure to be used in the reply to the <code>ioctl</code>.</p> <p>These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).</p> <p>As described in <code>gld(7D)</code>, the main data structure for communication between the device-specific driver and the GLD module is the <code>gld_mac_info(9S)</code> structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its <code>attach(9E)</code> routine, initialize these function pointers before calling <code>gld_register()</code>.</p> <p><code>gldm_reset()</code> resets the hardware to its initial state.</p> <p><code>gldm_start()</code> enables the device to generate interrupts and prepares the driver to call <code>gld_recv()</code> for delivering received data packets to GLD.</p> <p><code>gldm_stop()</code> disables the device from generating any interrupts and stops the driver from calling <code>gld_recv()</code> for delivering data packets to GLD. GLD depends on the <code>gldm_stop()</code> routine to ensure that the device will no longer interrupt, and it must do so without fail.</p> <p><code>gldm_set_mac_addr()</code> sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address <i>macaddr</i>.</p> <p><code>gldm_set_multicast()</code> enables and disables device-level reception of specific multicast addresses. If the third argument <i>multiflag</i> is set to <code>GLD_MULTI_ENABLE</code>, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if <i>multiflag</i> is set to <code>GLD_MULTI_DISABLE</code>, the driver is allowed to disable reception of the specified multicast address.</p> <p>This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.</p> <p><code>gldm_set_multicast()</code> will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.</p>
--------------------	---

gld(9E)

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

<code>DDI_INTR_CLAIMED</code>	if and only if the device definitely interrupted.
<code>DDI_INTR_UNCLAIMED</code>	if the device did not interrupt.

The other functions must return:

<code>GLD_SUCCESS</code>	on success. <code>gldm_stop()</code> and <code>gldm_ioctl()</code> should always return this value.
--------------------------	---

gld(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

NAME	gld, gldm_reset, gldm_start, gldm_stop, gldm_set_mac_addr, gldm_set_multicast, gldm_set_promiscuous, gldm_send, gldm_intr, gldm_get_stats, gldm_ioctl – Generic LAN Driver entry points
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>macinfo</i> Pointer to a gld_mac_info(9S) structure.</p> <p><i>macaddr</i> Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the gldm_addrlen element of the gld_mac_info(9S) structure.</p> <p><i>multicastaddr</i> Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the gldm_addrlen element of the gld_mac_info(9S) structure.</p> <p><i>multiflag</i> A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as GLD_MULTI_ENABLE or GLD_MULTI_DISABLE.</p> <p><i>promiscflag</i> A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as GLD_MAC_PROMISC_PHYS, GLD_MAC_PROMISC_MULTI, or GLD_MAC_PROMISC_NONE.</p> <p><i>mp</i> Pointer to a STREAMS message block containing the packet to be transmitted or the ioctl to be executed.</p> <p><i>stats</i> Pointer to a gld_stats(9S) structure to be filled in with the current values of statistics counters.</p>

gldm_get_stats(9E)

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION

These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

gldm_get_stats(9E)

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

`DDI_INTR_CLAIMED` if and only if the device definitely interrupted.

`DDI_INTR_UNCLAIMED` if the device did not interrupt.

The other functions must return:

`GLD_SUCCESS` on success. `gldm_stop()` and `gldm_ioctl()` should always return this value.

`gldm_get_stats(9E)`

<code>GLD_NORESOURCES</code>	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
<code>GLD_NOLINK</code>	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
<code>GLD_NOTSUPPORTED</code>	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
<code>GLD_BADARG</code>	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
<code>GLD_FAILURE</code>	on hardware failure.

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

Writing Device Drivers

gldm_intr(9E)

NAME	gld, gldm_reset, gldm_start, gldm_stop, gldm_set_mac_addr, gldm_set_multicast, gldm_set_promiscuous, gldm_send, gldm_intr, gldm_get_stats, gldm_ioctl – Generic LAN Driver entry points														
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>														
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).														
PARAMETERS	<table><tr><td><i>macinfo</i></td><td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>macaddr</i></td><td>Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrLen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multicastaddr</i></td><td>Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrLen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multiflag</i></td><td>A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</td></tr><tr><td><i>promiscflag</i></td><td>A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</td></tr><tr><td><i>mp</i></td><td>Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</td></tr><tr><td><i>stats</i></td><td>Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</td></tr></table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrLen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrLen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .	<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .	<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.	<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.														
<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrLen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrLen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .														
<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .														
<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.														
<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.														

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION

These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

gldm_intr(9E)

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

gldm_intr(9E)

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific `ioctl` commands. This element may be specified as `NULL` if the driver does not implement any `ioctl` functions. The driver is responsible for converting the message block into an `ioctl` reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

<code>DDI_INTR_CLAIMED</code>	if and only if the device definitely interrupted.
<code>DDI_INTR_UNCLAIMED</code>	if the device did not interrupt.

The other functions must return:

<code>GLD_SUCCESS</code>	on success. <code>gldm_stop()</code> and <code>gldm_ioctl()</code> should always return this value.
--------------------------	---

gldm_intr(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

NAME	<code>gld</code> , <code>gldm_reset</code> , <code>gldm_start</code> , <code>gldm_stop</code> , <code>gldm_set_mac_addr</code> , <code>gldm_set_multicast</code> , <code>gldm_set_promiscuous</code> , <code>gldm_send</code> , <code>gldm_intr</code> , <code>gldm_get_stats</code> , <code>gldm_ioctl</code> – Generic LAN Driver entry points
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>macinfo</i> Pointer to a <code>gld_mac_info(9S)</code> structure.</p> <p><i>macaddr</i> Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</p> <p><i>multicastaddr</i> Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</p> <p><i>multiflag</i> A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</p> <p><i>promiscflag</i> A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</p> <p><i>mp</i> Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</p> <p><i>stats</i> Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</p>

gldm_ioctl(9E)

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION

These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

gldm_ioctl(9E)

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

`DDI_INTR_CLAIMED` if and only if the device definitely interrupted.

`DDI_INTR_UNCLAIMED` if the device did not interrupt.

The other functions must return:

`GLD_SUCCESS` on success. `gldm_stop()` and `gldm_ioctl()` should always return this value.

gldm_ioctl(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

gldm_reset(9E)

NAME	<code>gld</code> , <code>gldm_reset</code> , <code>gldm_start</code> , <code>gldm_stop</code> , <code>gldm_set_mac_addr</code> , <code>gldm_set_multicast</code> , <code>gldm_set_promiscuous</code> , <code>gldm_send</code> , <code>gldm_intr</code> , <code>gldm_get_stats</code> , <code>gldm_ioctl</code> – Generic LAN Driver entry points														
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>														
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).														
PARAMETERS	<table><tr><td><i>macinfo</i></td><td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>macaddr</i></td><td>Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multicastaddr</i></td><td>Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multiflag</i></td><td>A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</td></tr><tr><td><i>promiscflag</i></td><td>A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</td></tr><tr><td><i>mp</i></td><td>Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</td></tr><tr><td><i>stats</i></td><td>Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</td></tr></table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .	<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .	<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.	<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.														
<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .														
<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .														
<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.														
<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.														

DESCRIPTION	<p><i>q</i> Pointer to the <code>queue(9S)</code> structure to be used in the reply to the <code>ioctl</code>.</p> <p>These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).</p> <p>As described in <code>gld(7D)</code>, the main data structure for communication between the device-specific driver and the GLD module is the <code>gld_mac_info(9S)</code> structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its <code>attach(9E)</code> routine, initialize these function pointers before calling <code>gld_register()</code>.</p> <p><code>gldm_reset()</code> resets the hardware to its initial state.</p> <p><code>gldm_start()</code> enables the device to generate interrupts and prepares the driver to call <code>gld_recv()</code> for delivering received data packets to GLD.</p> <p><code>gldm_stop()</code> disables the device from generating any interrupts and stops the driver from calling <code>gld_recv()</code> for delivering data packets to GLD. GLD depends on the <code>gldm_stop()</code> routine to ensure that the device will no longer interrupt, and it must do so without fail.</p> <p><code>gldm_set_mac_addr()</code> sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address <i>macaddr</i>.</p> <p><code>gldm_set_multicast()</code> enables and disables device-level reception of specific multicast addresses. If the third argument <i>multiflag</i> is set to <code>GLD_MULTI_ENABLE</code>, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if <i>multiflag</i> is set to <code>GLD_MULTI_DISABLE</code>, the driver is allowed to disable reception of the specified multicast address.</p> <p>This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.</p> <p><code>gldm_set_multicast()</code> will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.</p>
--------------------	---

gldm_reset(9E)

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

`gldm_reset(9E)`

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

<code>DDI_INTR_CLAIMED</code>	if and only if the device definitely interrupted.
<code>DDI_INTR_UNCLAIMED</code>	if the device did not interrupt.

The other functions must return:

<code>GLD_SUCCESS</code>	on success. <code>gldm_stop()</code> and <code>gldm_ioctl()</code> should always return this value.
--------------------------	---

gldm_reset(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

gldm_send(9E)

NAME	<code>gld</code> , <code>gldm_reset</code> , <code>gldm_start</code> , <code>gldm_stop</code> , <code>gldm_set_mac_addr</code> , <code>gldm_set_multicast</code> , <code>gldm_set_promiscuous</code> , <code>gldm_send</code> , <code>gldm_intr</code> , <code>gldm_get_stats</code> , <code>gldm_ioctl</code> – Generic LAN Driver entry points														
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>														
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).														
PARAMETERS	<table><tr><td><i>macinfo</i></td><td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>macaddr</i></td><td>Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multicastaddr</i></td><td>Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multiflag</i></td><td>A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</td></tr><tr><td><i>promiscflag</i></td><td>A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</td></tr><tr><td><i>mp</i></td><td>Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</td></tr><tr><td><i>stats</i></td><td>Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</td></tr></table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .	<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .	<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.	<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.														
<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .														
<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .														
<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.														
<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.														

gldm_send(9E)

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION

These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

gldm_send(9E)

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

`DDI_INTR_CLAIMED` if and only if the device definitely interrupted.

`DDI_INTR_UNCLAIMED` if the device did not interrupt.

The other functions must return:

`GLD_SUCCESS` on success. `gldm_stop()` and `gldm_ioctl()` should always return this value.

gldm_send(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

gldm_set_mac_addr(9E)

NAME	<code>gld</code> , <code>gldm_reset</code> , <code>gldm_start</code> , <code>gldm_stop</code> , <code>gldm_set_mac_addr</code> , <code>gldm_set_multicast</code> , <code>gldm_set_promiscuous</code> , <code>gldm_send</code> , <code>gldm_intr</code> , <code>gldm_get_stats</code> , <code>gldm_ioctl</code> – Generic LAN Driver entry points														
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>														
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).														
PARAMETERS	<table><tr><td><i>macinfo</i></td><td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>macaddr</i></td><td>Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multicastaddr</i></td><td>Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multiflag</i></td><td>A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</td></tr><tr><td><i>promiscflag</i></td><td>A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</td></tr><tr><td><i>mp</i></td><td>Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</td></tr><tr><td><i>stats</i></td><td>Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</td></tr></table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .	<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .	<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.	<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.														
<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .														
<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .														
<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.														
<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.														

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

gldm_set_mac_addr(9E)

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

`gldm_set_mac_addr(9E)`

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

<code>DDI_INTR_CLAIMED</code>	if and only if the device definitely interrupted.
<code>DDI_INTR_UNCLAIMED</code>	if the device did not interrupt.

The other functions must return:

<code>GLD_SUCCESS</code>	on success. <code>gldm_stop()</code> and <code>gldm_ioctl()</code> should always return this value.
--------------------------	---

gldm_set_mac_addr(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

NAME	gld, gldm_reset, gldm_start, gldm_stop, gldm_set_mac_addr, gldm_set_multicast, gldm_set_promiscuous, gldm_send, gldm_intr, gldm_get_stats, gldm_ioctl – Generic LAN Driver entry points
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>macinfo</i> Pointer to a gld_mac_info(9S) structure.</p> <p><i>macaddr</i> Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the gldm_addrlen element of the gld_mac_info(9S) structure.</p> <p><i>multicastaddr</i> Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the gldm_addrlen element of the gld_mac_info(9S) structure.</p> <p><i>multiflag</i> A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as GLD_MULTI_ENABLE or GLD_MULTI_DISABLE.</p> <p><i>promiscflag</i> A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as GLD_MAC_PROMISC_PHYS, GLD_MAC_PROMISC_MULTI, or GLD_MAC_PROMISC_NONE.</p> <p><i>mp</i> Pointer to a STREAMS message block containing the packet to be transmitted or the ioctl to be executed.</p> <p><i>stats</i> Pointer to a gld_stats(9S) structure to be filled in with the current values of statistics counters.</p>

gldm_set_multicast(9E)

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

gldm_set_multicast(9E)

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

`DDI_INTR_CLAIMED` if and only if the device definitely interrupted.

`DDI_INTR_UNCLAIMED` if the device did not interrupt.

The other functions must return:

`GLD_SUCCESS` on success. `gldm_stop()` and `gldm_ioctl()` should always return this value.

`gldm_set_multicast(9E)`

<code>GLD_NORESOURCES</code>	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
<code>GLD_NOLINK</code>	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
<code>GLD_NOTSUPPORTED</code>	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
<code>GLD_BADARG</code>	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
<code>GLD_FAILURE</code>	on hardware failure.

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

Writing Device Drivers

gldm_set_promiscuous(9E)

NAME	gld, gldm_reset, gldm_start, gldm_stop, gldm_set_mac_addr, gldm_set_multicast, gldm_set_promiscuous, gldm_send, gldm_intr, gldm_get_stats, gldm_ioctl – Generic LAN Driver entry points														
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>														
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>macinfo</i></td> <td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td> </tr> <tr> <td><i>macaddr</i></td> <td>Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td> </tr> <tr> <td><i>multicastaddr</i></td> <td>Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td> </tr> <tr> <td><i>multiflag</i></td> <td>A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</td> </tr> <tr> <td><i>promiscflag</i></td> <td>A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</td> </tr> <tr> <td><i>mp</i></td> <td>Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</td> </tr> <tr> <td><i>stats</i></td> <td>Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</td> </tr> </table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .	<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .	<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.	<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.														
<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .														
<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .														
<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.														
<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.														

gldm_set_promiscuous(9E)

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION

These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

`gldm_set_promiscuous(9E)`

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

`gldm_set_promiscuous(9E)`

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific `ioctl` commands. This element may be specified as `NULL` if the driver does not implement any `ioctl` functions. The driver is responsible for converting the message block into an `ioctl` reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

`DDI_INTR_CLAIMED` if and only if the device definitely interrupted.

`DDI_INTR_UNCLAIMED` if the device did not interrupt.

The other functions must return:

`GLD_SUCCESS` on success. `gldm_stop()` and `gldm_ioctl()` should always return this value.

gldm_set_promiscuous(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

NAME	gld, gldm_reset, gldm_start, gldm_stop, gldm_set_mac_addr, gldm_set_multicast, gldm_set_promiscuous, gldm_send, gldm_intr, gldm_get_stats, gldm_ioctl – Generic LAN Driver entry points
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>macinfo</i> Pointer to a gld_mac_info(9S) structure.</p> <p><i>macaddr</i> Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the gldm_addrlen element of the gld_mac_info(9S) structure.</p> <p><i>multicastaddr</i> Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the gldm_addrlen element of the gld_mac_info(9S) structure.</p> <p><i>multiflag</i> A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as GLD_MULTI_ENABLE or GLD_MULTI_DISABLE.</p> <p><i>promiscflag</i> A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as GLD_MAC_PROMISC_PHYS, GLD_MAC_PROMISC_MULTI, or GLD_MAC_PROMISC_NONE.</p> <p><i>mp</i> Pointer to a STREAMS message block containing the packet to be transmitted or the ioctl to be executed.</p> <p><i>stats</i> Pointer to a gld_stats(9S) structure to be filled in with the current values of statistics counters.</p>

gldm_start(9E)

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION

These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

gldm_start(9E)

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

`DDI_INTR_CLAIMED` if and only if the device definitely interrupted.

`DDI_INTR_UNCLAIMED` if the device did not interrupt.

The other functions must return:

`GLD_SUCCESS` on success. `gldm_stop()` and `gldm_ioctl()` should always return this value.

gldm_start(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

gldm_stop(9E)

NAME	<code>gld</code> , <code>gldm_reset</code> , <code>gldm_start</code> , <code>gldm_stop</code> , <code>gldm_set_mac_addr</code> , <code>gldm_set_multicast</code> , <code>gldm_set_promiscuous</code> , <code>gldm_send</code> , <code>gldm_intr</code> , <code>gldm_get_stats</code> , <code>gldm_ioctl</code> – Generic LAN Driver entry points														
SYNOPSIS	<pre>#include <sys/gld.h> int prefix_reset(gld_mac_info_t *macinfo); int prefix_start(gld_mac_info_t *macinfo); int prefix_stop(gld_mac_info_t *macinfo); int prefix_set_mac_addr(gld_mac_info_t *macinfo, unsigned char *macaddr); int prefix_set_multicast(gld_mac_info_t *macinfo, unsigned char *multicastaddr, int multiflag); int prefix_set_promiscuous(gld_mac_info_t *macinfo, int promiscflag); int prefix_send(gld_mac_info_t *macinfo, mblk_t *mp); uint_t prefix_intr(gld_mac_info_t *macinfo); int prefix_get_stats(gld_mac_info_t *macinfo, struct gld_stats *stats); int prefix_ioctl(gld_mac_info_t *macinfo, queue_t *q, mblk_t *mp);</pre>														
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).														
PARAMETERS	<table><tr><td><i>macinfo</i></td><td>Pointer to a <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>macaddr</i></td><td>Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multicastaddr</i></td><td>Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.</td></tr><tr><td><i>multiflag</i></td><td>A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code>.</td></tr><tr><td><i>promiscflag</i></td><td>A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code>, <code>GLD_MAC_PROMISC_MULTI</code>, or <code>GLD_MAC_PROMISC_NONE</code>.</td></tr><tr><td><i>mp</i></td><td>Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.</td></tr><tr><td><i>stats</i></td><td>Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.</td></tr></table>	<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.	<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.	<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .	<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .	<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.	<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.
<i>macinfo</i>	Pointer to a <code>gld_mac_info(9S)</code> structure.														
<i>macaddr</i>	Pointer to the beginning of a character array containing a valid MAC address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multicastaddr</i>	Pointer to the beginning of a character array containing a multicast, group, or functional address. The array will be of the length specified by the driver in the <code>gldm_addrlen</code> element of the <code>gld_mac_info(9S)</code> structure.														
<i>multiflag</i>	A flag indicating whether reception of the multicast address is to be enabled or disabled. This argument is specified as <code>GLD_MULTI_ENABLE</code> or <code>GLD_MULTI_DISABLE</code> .														
<i>promiscflag</i>	A flag indicating what type of promiscuous mode, if any, is to be enabled. This argument is specified as <code>GLD_MAC_PROMISC_PHYS</code> , <code>GLD_MAC_PROMISC_MULTI</code> , or <code>GLD_MAC_PROMISC_NONE</code> .														
<i>mp</i>	Pointer to a STREAMS message block containing the packet to be transmitted or the <code>ioctl</code> to be executed.														
<i>stats</i>	Pointer to a <code>gld_stats(9S)</code> structure to be filled in with the current values of statistics counters.														

q Pointer to the `queue(9S)` structure to be used in the reply to the `ioctl`.

DESCRIPTION These entry points must be implemented by a device-specific network driver designed to interface with the Generic LAN Driver (GLD).

As described in `gld(7D)`, the main data structure for communication between the device-specific driver and the GLD module is the `gld_mac_info(9S)` structure. Some of the elements in that structure are function pointers to the entry points described here. The device-specific driver must, in its `attach(9E)` routine, initialize these function pointers before calling `gld_register()`.

`gldm_reset()` resets the hardware to its initial state.

`gldm_start()` enables the device to generate interrupts and prepares the driver to call `gld_recv()` for delivering received data packets to GLD.

`gldm_stop()` disables the device from generating any interrupts and stops the driver from calling `gld_recv()` for delivering data packets to GLD. GLD depends on the `gldm_stop()` routine to ensure that the device will no longer interrupt, and it must do so without fail.

`gldm_set_mac_addr()` sets the physical address that the hardware is to use for receiving data. This function should program the device to the passed MAC address *macaddr*.

`gldm_set_multicast()` enables and disables device-level reception of specific multicast addresses. If the third argument *multiflag* is set to `GLD_MULTI_ENABLE`, then the function sets the interface to receive packets with the multicast address pointed to by the second argument; if *multiflag* is set to `GLD_MULTI_DISABLE`, the driver is allowed to disable reception of the specified multicast address.

This function is called whenever GLD wants to enable or disable reception of a multicast, group, or functional address. GLD makes no assumptions about how the device does multicast support and calls this function to enable or disable a specific multicast address. Some devices may use a hash algorithm and a bitmask to enable collections of multicast addresses; this is allowed, and GLD will filter out any superfluous packets that are not required. If disabling an address could result in disabling more than one address at the device level, it is the responsibility of the device driver to keep whatever information it needs to avoid disabling an address that GLD has enabled but not disabled.

`gldm_set_multicast()` will not be called to enable a particular multicast address that is already enabled, nor to disable an address that is not currently enabled. GLD keeps track of multiple requests for the same multicast address and only calls the driver's entry point when the first request to enable, or the last request to disable a particular multicast address is made.

gldm_stop(9E)

`gldm_set_promiscuous()` enables and disables promiscuous mode. This function is called whenever GLD wants to enable or disable the reception of all packets on the medium, or all multicast packets on the medium. If the second argument *promiscflag* is set to the value of `GLD_MAC_PROMISC_PHYS`, then the function enables physical-level promiscuous mode, resulting in the reception of all packets on the medium. If *promiscflag* is set to `GLD_MAC_PROMISC_MULTI`, then reception of all multicast packets will be enabled. If *promiscflag* is set to `GLD_MAC_PROMISC_NONE`, then promiscuous mode is disabled.

In the case of a request for promiscuous multicast mode, drivers for devices that have no multicast-only promiscuous mode must set the device to physical promiscuous mode to ensure that all multicast packets are received. In this case the routine should return `GLD_SUCCESS`. The GLD software will filter out any superfluous packets that are not required.

For forward compatibility, `gldm_set_promiscuous()` routines should treat any unrecognized values for *promiscflag* as though they were `GLD_MAC_PROMISC_PHYS`.

`gldm_send()` queues a packet to the device for transmission. This routine is passed a STREAMS message containing the packet to be sent. The message may comprise multiple message blocks, and the send routine must chain through all the message blocks in the message to access the entire packet to be sent. The driver should be prepared to handle and skip over any zero-length message continuation blocks in the chain. The driver should check to ensure that the packet does not exceed the maximum allowable packet size, and must pad the packet, if necessary, to the minimum allowable packet size. If the send routine successfully transmits or queues the packet, it should return `GLD_SUCCESS`.

The send routine should return `GLD_NORESOURCES` if it cannot immediately accept the packet for transmission; in this case GLD will retry it later. If `gldm_send()` ever returns `GLD_NORESOURCES`, the driver must, at a later time when resources have become available, call `gld_sched()` to inform GLD that it should retry packets that the driver previously failed to queue for transmission. (If the driver's `gldm_stop()` routine is called, the driver is absolved from this obligation until it later again returns `GLD_NORESOURCES` from its `gldm_send()` routine; however, extra calls to `gld_sched()` will not cause incorrect operation.)

If the driver's send routine returns `GLD_SUCCESS`, then the driver is responsible for freeing the message when the driver and the hardware no longer need it. If the send routine copied the message into the device, or into a private buffer, then the send routine may free the message after the copy is made. If the hardware uses DMA to read the data directly out of the message data blocks, then the driver must not free the message until the hardware has completed reading the data. In this case the driver will probably free the message in the interrupt routine, or in a buffer-reclaim operation at the beginning of a future send operation. If the send routine returns anything other than `GLD_SUCCESS`, then the driver must not free the message.

`gldm_stop(9E)`

`gldm_intr()` is called when the device might have interrupted. Since it is possible to share interrupts with other devices, the driver must check the device status to determine whether it actually caused an interrupt. If the device that the driver controls did not cause the interrupt, then this routine must return `DDI_INTR_UNCLAIMED`. Otherwise it must service the interrupt and should return `DDI_INTR_CLAIMED`. If the interrupt was caused by successful receipt of a packet, this routine should put the received packet into a STREAMS message of type `M_DATA` and pass that message to `gld_rcv()`.

`gld_rcv()` will pass the inbound packet upstream to the appropriate next layer of the network protocol stack. It is important to correctly set the `b_rptr` and `b_wptr` members of the STREAMS message before calling `gld_rcv()`.

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

The interrupt code should increment statistics counters for any errors. This includes failure to allocate a buffer needed for the received data and any hardware-specific errors such as CRC errors or framing errors.

`gldm_get_stats()` gathers statistics from the hardware and/or driver private counters, and updates the `gld_stats(9S)` structure pointed to by `stats`. This routine is called by GLD when it gets a request for statistics, and provides the mechanism by which GLD acquires device dependent statistics from the driver before composing its reply to the statistics request. See `gld_stats(9S)` and `gld(7D)` for a description of the defined statistics counters.

`gldm_ioctl()` implements any device-specific ioctl commands. This element may be specified as `NULL` if the driver does not implement any ioctl functions. The driver is responsible for converting the message block into an ioctl reply message and calling the `qreply(9F)` function before returning `GLD_SUCCESS`. This function should always return `GLD_SUCCESS`; any errors the driver may wish to report should be returned via the message passed to `qreply(9F)`. If the `gldm_ioctl` element is specified as `NULL`, GLD will return a message of type `M_IOCNAK` with an error of `EINVAL`.

RETURN VALUES

`gldm_intr()` must return:

<code>DDI_INTR_CLAIMED</code>	if and only if the device definitely interrupted.
<code>DDI_INTR_UNCLAIMED</code>	if the device did not interrupt.

The other functions must return:

<code>GLD_SUCCESS</code>	on success. <code>gldm_stop()</code> and <code>gldm_ioctl()</code> should always return this value.
--------------------------	---

gldm_stop(9E)

GLD_NORESOURCES	if there are insufficient resources to carry out the request at this time. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , <code>gldm_set_promiscuous()</code> , and <code>gldm_send()</code> may return this value.
GLD_NOLINK	if <code>gldm_send()</code> is called when there is no physical connection to a network or link partner.
GLD_NOTSUPPORTED	if the requested function is not supported. Only <code>gldm_set_mac_addr()</code> , <code>gldm_set_multicast()</code> , and <code>gldm_set_promiscuous()</code> may return this value.
GLD_BADARG	if the function detected an unsuitable argument, for example, a bad multicast address, a bad MAC address, or a bad packet or packet length.
GLD_FAILURE	on hardware failure.

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

Writing Device Drivers

NAME identify – determine if a driver is associated with a device

SYNOPSIS

```
#include <sys/conf.h>
#include <sys/ddi.h>
#include <sys/sunddi.h>

int prefixidentify(dev_info_t *dip);
```

INTERFACE LEVEL Solaris DDI specific (Solaris DDI). This entry point is obsolete and is no longer required. This entry point may not be supported in future releases. nulldev(9F) should be specified in the dev_ops(9S) structure.

ARGUMENTS *dip* A pointer to a dev_info structure.

DESCRIPTION identify() was used to determine whether a driver drives the device pointed to by *dip*. identify() is currently supported to provide backward compatibility with older drivers and should not be implemented. See the INTERFACE LEVEL section.

RETURN VALUES The return value from identify() is ignored.

SEE ALSO nulldev(9F), dev_ops(9S)

WARNINGS This routine may be called multiple times. It may also be called at any time. The driver should not infer anything from the the sequence or the number of times this entry point has been called.

This entry point may not be supported in future releases.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete*

* Schedule for removal in a minor release after 8/98.

`_info(9E)`

NAME	<code>_fini</code> , <code>_info</code> , <code>_init</code> – loadable module configuration entry points
SYNOPSIS	<pre>#include <sys/modctl.h> int _fini(void); int _info(struct modinfo *<i>modinfop</i>); int _init(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). These entry points are required. You must write them.
<code>_info()</code>	<i>modinfop</i> A pointer to an opaque <code>modinfo</code> structure.
DESCRIPTION	<p><code>_init()</code> initializes a loadable module. It is called before any other routine in a loadable module. <code>_init()</code> returns the value returned by <code>mod_install(9F)</code>. The module may optionally perform some other work before the <code>mod_install(9F)</code> call is performed. If the module has done some setup before the <code>mod_install(9F)</code> function is called, then it should be prepared to undo that setup if <code>mod_install(9F)</code> returns an error.</p> <p><code>_info()</code> returns information about a loadable module. <code>_info()</code> returns the value returned by <code>mod_info(9F)</code>.</p> <p><code>_fini()</code> prepares a loadable module for unloading. It is called when the system wants to unload a module. If the module determines that it can be unloaded, then <code>_fini()</code> returns the value returned by <code>mod_remove(9F)</code>. Upon successful return from <code>_fini()</code> no other routine in the module will be called before <code>_init()</code> is called.</p>
RETURN VALUES	<p><code>_init()</code> should return the appropriate error number if there is an error, otherwise it should return the return value from <code>mod_install(9F)</code>.</p> <p><code>_info()</code> should return the return value from <code>mod_info(9F)</code></p> <p><code>_fini()</code> should return the return value from <code>mod_remove(9F)</code>. <code>_fini()</code> is permitted to return <code>EBUSY</code> prior to calling <code>mod_remove(9F)</code> if the driver should not be unloaded. Driver global resources, such as mutexes and calls to <code>ddi_soft_state_fini(9F)</code>, should only be destroyed in <code>_fini()</code> after <code>mod_remove()</code> returns successfully.</p>
EXAMPLES	<p>EXAMPLE 1 Initializing and Freeing a Mutex</p> <p>The following example demonstrates how to initialize and free a <code>mutex(9F)</code>.</p> <pre>#include <sys/modctl.h> #include <sys/ddi.h> #include <sys/sunddi.h> static struct dev_ops drv_ops; /* * Module linkage information for the kernel. */ static struct modldrv modldrv = {</pre>

EXAMPLE 1 Initializing and Freeing a Mutex (Continued)

```

    &mod_driverops,      /* Type of module. This one is a driver */
    "Sample Driver",
    &drv_ops             /* driver ops */
};

static struct modlinkage modlinkage = {
    MODREV_1,
    &modldrv,
    NULL
};

/*
 * Global driver mutex
 */
static kmutex_t  xx_global_mutex;

int
_init(void)
{
    int    i;

    /*
     * Initialize global mutex before mod_install'ing driver.
     * If mod_install() fails, must clean up mutex initialization
     */
    mutex_init(&xx_global_mutex, NULL,
              MUTEX_DRIVER, (void *)NULL);

    if ((i = mod_install(&modlinkage)) != 0) {
        mutex_destroy(&xx_global_mutex);
    }

    return (i);
}

int
_info(struct modinfo *modinfop)
{
    return (mod_info(&modlinkage, modinfop));
}

int
_fini(void)
{
    int    i;

    /*
     * If mod_remove() is successful, we destroy our global mutex
     */
    if ((i = mod_remove(&modlinkage)) == 0) {
        mutex_destroy(&xx_global_mutex);
    }
}

```

`_info(9E)`

EXAMPLE 1 Initializing and Freeing a Mutex *(Continued)*

```
    }  
    return (i);  
}
```

SEE ALSO `add_drv(1M)`, `mod_info(9F)`, `mod_install(9F)`, `mod_remove(9F)`, `mutex(9F)`,
`modldrv(9S)`, `modlinkage(9S)`, `modlstrmod(9S)`

Writing Device Drivers

WARNINGS Do not change the structures referred to by the `modlinkage` structure after the call to `mod_install()`, as the system may copy or change them.

NOTES Even though the identifiers `_fini()`, `_info()`, and `_init()` appear to be declared as globals, their scope is restricted by the kernel to the module that they are defined in.

BUGS On some implementations `_info()` may be called before `_init()`.

NAME	<code>_fini</code> , <code>_info</code> , <code>_init</code> – loadable module configuration entry points
SYNOPSIS	<pre>#include <sys/modctl.h> int _fini(void); int _info(struct modinfo *modinfop); int _init(void);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). These entry points are required. You must write them.
<code>_info()</code>	<i>modinfop</i> A pointer to an opaque <code>modinfo</code> structure.
DESCRIPTION	<p><code>_init()</code> initializes a loadable module. It is called before any other routine in a loadable module. <code>_init()</code> returns the value returned by <code>mod_install(9F)</code>. The module may optionally perform some other work before the <code>mod_install(9F)</code> call is performed. If the module has done some setup before the <code>mod_install(9F)</code> function is called, then it should be prepared to undo that setup if <code>mod_install(9F)</code> returns an error.</p> <p><code>_info()</code> returns information about a loadable module. <code>_info()</code> returns the value returned by <code>mod_info(9F)</code>.</p> <p><code>_fini()</code> prepares a loadable module for unloading. It is called when the system wants to unload a module. If the module determines that it can be unloaded, then <code>_fini()</code> returns the value returned by <code>mod_remove(9F)</code>. Upon successful return from <code>_fini()</code> no other routine in the module will be called before <code>_init()</code> is called.</p>
RETURN VALUES	<p><code>_init()</code> should return the appropriate error number if there is an error, otherwise it should return the return value from <code>mod_install(9F)</code>.</p> <p><code>_info()</code> should return the return value from <code>mod_info(9F)</code></p> <p><code>_fini()</code> should return the return value from <code>mod_remove(9F)</code>. <code>_fini()</code> is permitted to return <code>EBUSY</code> prior to calling <code>mod_remove(9F)</code> if the driver should not be unloaded. Driver global resources, such as mutexes and calls to <code>ddi_soft_state_fini(9F)</code>, should only be destroyed in <code>_fini()</code> after <code>mod_remove()</code> returns successfully.</p>
EXAMPLES	<p>EXAMPLE 1 Initializing and Freeing a Mutex</p> <p>The following example demonstrates how to initialize and free a <code>mutex(9F)</code>.</p> <pre>#include <sys/modctl.h> #include <sys/ddi.h> #include <sys/sunddi.h> static struct dev_ops drv_ops; /* * Module linkage information for the kernel. */ static struct modldrv modldrv = {</pre>

EXAMPLE 1 Initializing and Freeing a Mutex (Continued)

```
        &mod_driverops,    /* Type of module. This one is a driver */
        "Sample Driver",
        &drv_ops           /* driver ops */
    };

    static struct modlinkage modlinkage = {
        MODREV_1,
        &modldrv,
        NULL
    };

    /*
     * Global driver mutex
     */
    static kmutex_t    xx_global_mutex;

int
_init(void)
{
    int    i;

    /*
     * Initialize global mutex before mod_install'ing driver.
     * If mod_install() fails, must clean up mutex initialization
     */
    mutex_init(&xx_global_mutex, NULL,
               MUTEX_DRIVER, (void *)NULL);

    if ((i = mod_install(&modlinkage)) != 0) {
        mutex_destroy(&xx_global_mutex);
    }

    return (i);
}

int
_info(struct modinfo *modinfop)
{
    return (mod_info(&modlinkage, modinfop));
}

int
_fini(void)
{
    int    i;

    /*
     * If mod_remove() is successful, we destroy our global mutex
     */
    if ((i = mod_remove(&modlinkage)) == 0) {
        mutex_destroy(&xx_global_mutex);
    }
}
```

EXAMPLE 1 Initializing and Freeing a Mutex (Continued)

```
    }  
    return (i);  
}
```

SEE ALSO [add_drv\(1M\)](#), [mod_info\(9F\)](#), [mod_install\(9F\)](#), [mod_remove\(9F\)](#), [mutex\(9F\)](#), [modldrv\(9S\)](#), [modlinkage\(9S\)](#), [modlstrmod\(9S\)](#)

Writing Device Drivers

WARNINGS Do not change the structures referred to by the `modlinkage` structure after the call to `mod_install()`, as the system may copy or change them.

NOTES Even though the identifiers `_fini()`, `_info()`, and `_init()` appear to be declared as globals, their scope is restricted by the kernel to the module that they are defined in.

BUGS On some implementations `_info()` may be called before `_init()`.

ioctl(9E)

NAME	ioctl – control a character device
SYNOPSIS	<pre>#include <sys/cred.h> #include <sys/file.h> #include <sys/types.h> #include <sys/errno.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>ioctl(dev_t dev, int cmd, intptr_t arg, int mode, cred_t *cred_p, int *rval_p);</pre>
INTERFACE LEVEL ARGUMENTS	Architecture independent level 1 (DDI/DKI). This entry point is optional. <i>dev</i> Device number. <i>cmd</i> Command argument the driver <code>ioctl()</code> routine interprets as the operation to be performed. <i>arg</i> Passes parameters between a user program and the driver. When used with terminals, the argument is the address of a user program structure containing driver or hardware settings. Alternatively, the argument may be a value that has meaning only to the driver. The interpretation of the argument is driver dependent and usually depends on the command type; the kernel does not interpret the argument. <i>mode</i> A bit field that contains: <ul style="list-style-type: none">■ Information set when the device was opened. The driver may use it to determine if the device was opened for reading or writing. The driver can make this determination by checking the <code>FREAD</code> or <code>FWRITE</code> flags. See the <i>flag</i> argument description of the <code>open()</code> routine for further values.■ Information on whether the caller is a 32-bit or 64-bit thread.■ In some circumstances address space information about the <i>arg</i> argument. See below. <i>cred_p</i> Pointer to the user credential structure. <i>rval_p</i> Pointer to return value for calling process. The driver may elect to set the value which is valid only if the <code>ioctl()</code> succeeds.
DESCRIPTION	<p><code>ioctl()</code> provides character-access drivers with an alternate entry point that can be used for almost any operation other than a simple transfer of characters in and out of buffers. Most often, <code>ioctl()</code> is used to control device hardware parameters and establish the protocol used by the driver in processing data.</p> <p>The kernel determines that this is a character device, and looks up the entry point routines in <code>cb_ops(9S)</code>. The kernel then packages the user request and arguments as integers and passes them to the driver's <code>ioctl()</code> routine. The kernel itself does no processing of the passed command, so it is up to the user program and the driver to agree on what the arguments mean.</p>

I/O control commands are used to implement the terminal settings passed from `ttymon(1M)` and `stty(1)`, to format disk devices, to implement a trace driver for debugging, and to clean up character queues. Since the kernel does not interpret the command type that defines the operation, a driver is free to define its own commands.

Drivers that use an `ioctl()` routine typically have a command to “read” the current `ioctl()` settings, and at least one other that sets new settings. Drivers can use the *mode* argument to determine if the device unit was opened for reading or writing, if necessary, by checking the `FREAD` or `FWRITE` setting.

If the third argument, *arg*, is a pointer to a user buffer, the driver can call the `copyin(9F)` and `copyout(9F)` functions to transfer data between kernel and user space.

Other kernel subsystems may need to call into the drivers `ioctl()` routine. Drivers that intend to allow their `ioctl()` routine to be used in this way should publish the `ddi-kernel-ioctl` property on the associated devinfo node(s).

When the `ddi-kernel-ioctl` property is present, the *mode* argument is used to pass address space information about *arg* through to the driver. If the driver expects *arg* to contain a buffer address, and the `FKIOCTL` flag is set in *mode*, then the driver should assume that it is being handed a kernel buffer address. Otherwise, *arg* may be the address of a buffer from a user program. The driver can use `ddi_copyin(9F)` and `ddi_copyout(9F)` perform the correct type of copy operation for either kernel or user address spaces. See the example on `ddi_copyout(9F)`.

Drivers have to interact with 32-bit and 64-bit applications. If a device driver shares data structures with the application (for example, through exported kernel memory) and the driver gets recompiled for a 64-bit kernel but the application remains 32-bit, binary layout of any data structures will be incompatible if they contain longs or pointers. The driver needs to know whether there is a model mismatch between the current thread and the kernel and take necessary action. The *mode* argument has additional bits set to determine the C Language Type Model which the current thread expects. *mode* has `FILP32` set if the current thread expects 32-bit (*ILP32*) semantics, or `FLP64` if the current thread expects 64-bit (*LP64*) semantics. *mode* is used in combination with `ddi_model_convert_from(9F)` and the `FMODELS` mask to determine whether there is a data model mismatch between the current thread and the device driver (see the example below). The device driver might have to adjust the shape of data structures before exporting them to a user thread which supports a different data model.

To implement I/O control commands for a driver the following two steps are required:

1. Define the I/O control command names and the associated value in the driver's header and comment the commands.
2. Code the `ioctl()` routine in the driver that defines the functionality for each I/O control command name that is in the header.

ioctl(9E)

The `ioctl()` routine is coded with instructions on the proper action to take for each command. It is commonly a `switch` statement, with each case definition corresponding to an `ioctl()` name to identify the action that should be taken. However, the command passed to the driver by the user process is an integer value associated with the command name in the header.

RETURN VALUES `ioctl()` should return 0 on success, or the appropriate error number. The driver may also set the value returned to the calling process through `rval_p`.

EXAMPLES **EXAMPLE 1** `ioctl()` entry point

The following is an example of the `ioctl()` entry point and how to support 32-bit and 64-bit applications with the same device driver.

```
struct passargs32 {
    int len;
    caddr32_t addr;
};

struct passargs {
    int len;
    caddr_t addr;
};

xxioctl(dev_t dev, int cmd, intptr_t arg, int mode,
        cred_t *credp, int *rvalp) {
    struct passargs pa;

#ifdef _MULTI_DATAMODEL
    switch (ddi_model_convert_from(mode & FMODELS)) {
        case DDI_MODEL_ILP32:
            {
                struct passargs32 pa32;

                ddi_copyin(arg, &pa32, sizeof (struct passargs32), mode);
                pa.len = pa32.len;
                pa.address = pa32.address;
                break;
            }
        case DDI_MODEL_NONE:
            ddi_copyin(arg, &pa, sizeof (struct passargs), mode);
            break;
    }
#else /* _MULTI_DATAMODEL */
    ddi_copyin(arg, &pa, sizeof (struct passargs), mode);
#endif /* _MULTI_DATAMODEL */

    do_ioctl(&pa);
    . . . .
}
```

SEE ALSO `stty(1)`, `ttymon(1M)`, `dkio(7I)`, `fbio(7I)`, `termio(7I)`, `open(9E)`, `put(9E)`, `srv(9E)`, `copyin(9F)`, `copyout(9F)`, `ddi_copyin(9F)`, `ddi_copyout(9F)`, `ddi_model_convert_from(9F)`, `cb_ops(9S)`

WARNINGS	<p>Non-STREAMS driver <code>ioctl()</code> routines must make sure that user data is copied into or out of the kernel address space explicitly using <code>copyin(9F)</code>, <code>copyout(9F)</code>, <code>ddi_copyin(9F)</code>, or <code>ddi_copyout(9F)</code>, as appropriate.</p> <p>It is a severe error to simply dereference pointers to the user address space, even when in user context.</p> <p>Failure to use the appropriate copying routines can result in panics under load on some platforms, and reproducible panics on others.</p>
NOTES	<p>STREAMS drivers do not have <code>ioctl()</code> routines. The stream head converts I/O control commands to <code>M_IOCTL</code> messages, which are handled by the driver's <code>put(9E)</code> or <code>srv(9E)</code> routine.</p>

ks_snapshot(9E)

NAME	ks_snapshot – take a snapshot of kstat data
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> #include <sys/ddi.h> #include <sys/sunddi.h> int prefix_ks_snapshot(kstat_t *ksp, void *buf, int rw);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI). <i>ksp</i> Pointer to a kstat(9S) structure. <i>buf</i> Pointer to a buffer to copy the snapshot into. <i>rw</i> Read/Write flag. Possible values are: KSTAT_READ Copy driver statistics from the driver to the buffer. KSTAT_WRITE Copy statistics from the buffer to the driver.
DESCRIPTION	<p>The kstat mechanism allows for an optional ks_snapshot() function to copy kstat data. This is the routine that is called to marshall the kstat data to be copied to user-land. A driver can opt to use a custom snapshot routine rather than the default snapshot routine; to take advantage of this feature, set the ks_snapshot field before calling kstat_install(9F).</p> <p>The ks_snapshot() function must have the following structure:</p> <pre>static int xx_kstat_snapshot(kstat_t *ksp, void *buf, int rw) { if (rw == KSTAT_WRITE) { /* set the native stats to the values in buf */ /* return EACCES if you don't support this */ } else { /* copy the kstat-specific data into buf */ } return (0); }</pre> <p>In general, the ks_snapshot() routine might need to refer to provider-private data; for example, it might need a pointer to the provider's raw statistics. The ks_private field is available for this purpose. Its use is entirely at the provider's discretion.</p> <p>No kstat locking should be done inside the ks_update() routine. The caller will already be holding the kstat's ks_lock (to ensure consistent data) and will prevent the kstat from being removed.</p> <ol style="list-style-type: none">1. ks_snaptime must be set (via gethrtime(9F)) to timestamp the data.

2. Data gets copied from the `kstat` to the buffer on `KSTAT_READ`, and from the buffer to the `kstat` on `KSTAT_WRITE`.

RETURN VALUES

0
Success

`EACCESS`
If `KSTAT_WRITE` is not allowed

`EIO`
For any other error

CONTEXT This function is called from user context only.

EXAMPLES **EXAMPLE 1** Named `kstats` with Long Strings (`KSTAT_DATA_STRING`)

```
static int
xxx_kstat_snapshot(kstat_t *ksp, void *buf, int rw)
{
    if (rw == KSTAT_WRITE) {
        return (EACCESS);
    } else {
        kstat_named_t *knp = buf;
        char *end = knp + ksp->ks_ndata;
        uint_t i;

        bcopy(ksp->ks_data, buf,
              sizeof (kstat_named_t) * ksp->ks_ndata);
/*
 * Now copy the strings to the end of the buffer, and
 * update the pointers appropriately.
 */
        for (i = 0; i < ksp->ks_ndata; i++, knp++)
            if (knp->data_type == KSTAT_DATA_STRING &&
                KSTAT_NAMED_STR_PTR(knp) != NULL) {
                bcopy(KSTAT_NAMED_STR_PTR(knp), end,
                    KSTAT_NAMED_STR_BUFLLEN(knp));
                KSTAT_NAMED_STR_PTR(knp) = end;
                end += KSTAT_NAMED_STR_BUFLLEN(knp);
            }
    }
    return (0);
}
```

SEE ALSO `ks_update(9E)`, `kstat_create(9F)`, `kstat_install(9F)`, `kstat(9S)`

Writing Device Drivers

ks_update(9E)

NAME	ks_update – dynamically update kstats
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/kstat.h> #include <sys/ddi.h> #include <sys/sunddi.h> int prefix_ks_update(kstat_t *ksp, int rw);</pre>
INTERFACE LEVEL PARAMETERS	Solaris DDI specific (Solaris DDI) <i>ksp</i> Pointer to a kstat(9S) structure. <i>rw</i> Read/Write flag. Possible values are KSTAT_READ Update kstat structure statistics from the driver. KSTAT_WRITE Update driver statistics from the kstat structure.
DESCRIPTION	<p>The kstat mechanism allows for an optional ks_update () function to update kstat data. This is useful for drivers where the underlying device keeps cheap hardware statistics, but extraction is expensive. Instead of constantly keeping the kstat data section up to date, the driver can supply a ks_update () function which updates the kstat's data section on demand. To take advantage of this feature, set the ks_update field before calling kstat_install(9F).</p> <p>The ks_update () function must have the following structure:</p> <pre>static int xx_kstat_update(kstat_t *ksp, int rw) { if (rw == KSTAT_WRITE) { /* update the native stats from ksp->ks_data */ /* return EACCES if you don't support this */ } else { /* update ksp->ks_data from the native stats */ } return (0); }</pre> <p>In general, the ks_update () routine may need to refer to provider-private data; for example, it may need a pointer to the provider's raw statistics. The ks_private field is available for this purpose. Its use is entirely at the provider's discretion.</p> <p>No kstat locking should be done inside the ks_update () routine. The caller will already be holding the kstat's ks_lock (to ensure consistent data) and will prevent the kstat from being removed.</p>
RETURN VALUES	ks_update () should return 0 For success. EACCES If KSTAT_WRITE is not allowed. EIO For any other error.

ks_update(9E)

SEE ALSO kstat_create(9F), kstat_install(9F), kstat(9S)

Writing Device Drivers

mapdev_access(9E)

NAME	mapdev_access – device mapping access entry point
SYNOPSIS	<pre>#include <sys/sunddi.h> int <i>prefix</i>mapdev_access (ddi_mapdev_handle_t handle, void *devprivate, off_t offset);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<p><i>handle</i> An opaque pointer to a device mapping.</p> <p><i>devprivate</i> Driver private mapping data from ddi_mapdev(9F).</p> <p><i>offset</i> The offset within device memory at which the access occurred.</p>
DESCRIPTION	<p>Future releases of Solaris will provide this function for binary and source compatibility. However, for increased functionality, use devmap_access(9E) or devmap_contextmgt(9E) instead. See devmap_access(9E) or devmap_contextmgt(9E) for details.</p> <p>mapdev_access () is called when an access is made to a mapping that has either been newly created with ddi_mapdev(9F) or that has been enabled with a call to ddi_mapdev_intercept(9F).</p> <p>mapdev_access () is passed the <i>handle</i> of the mapped object on which an access has occurred. This handle uniquely identifies the mapping and is used as an argument to ddi_mapdev_intercept(9F) or ddi_mapdev_nointercept(9F) to control whether or not future accesses to the mapping will cause mapdev_access () to be called. In general, mapdev_access () should call ddi_mapdev_intercept () on the mapping that is currently in use and then call ddi_mapdev_nointercept () on the mapping that generated this call to mapdev_access (). This will ensure that a call to mapdev_access () will be generated for the current mapping next time it is accessed.</p> <p>mapdev_access () must at least call ddi_mapdev_nointercept () with <i>offset</i> passed in in order for the access to succeed. A request to allow accesses affects the entire page containing the <i>offset</i>.</p> <p>Accesses to portions of mappings that have been disabled by a call to ddi_mapdev_nointercept () will not generate a call to mapdev_access (). A subsequent call to ddi_mapdev_intercept () will enable mapdev_access () to be called again.</p> <p>A non-zero return value from mapdev_access () will cause the corresponding operation to fail. The failure may result in a SIGSEGV or SIGBUS signal being delivered to the process.</p>
RETURN VALUES	mapdev_access () should return 0 on success, -1 if there was a hardware error, or the return value from ddi_mapdev_intercept () or ddi_mapdev_nointercept ().

CONTEXT This function is called from user context only.

EXAMPLES **EXAMPLE 1** Managing a One Page Device Context

The following shows an example of managing a device context that is one page in length.

```

ddi_mapdev_handle_t cur_hdl;
static int
xxmapdev_access(ddi_mapdev_handle_t handle, void *devprivate,
                off_t offset)
{
    int    err;
    /* enable calls to mapdev_access for the current mapping */
    if (cur_hdl != NULL) {
        if ((err = ddi_mapdev_intercept(cur_hdl, off, 0)) != 0)
            return (err);
    }
    /* Switch device context - device dependent*/
    ...
    /* Make handle the new current mapping */
    cur_hdl = handle;

    /*
     * Disable callbacks and complete the access for the
     * mapping that generated this callback.
     */

    return (ddi_mapdev_nointercept(handle, off, 0));
}

```

SEE ALSO mmap(2), mapdev_dup(9E), mapdev_free(9E), segmap(9E), ddi_mapdev(9F), ddi_mapdev_intercept(9F), ddi_mapdev_nointercept(9F), ddi_mapdev_ctl(9S)

Writing Device Drivers

mapdev_dup(9E)

NAME	mapdev_dup – device mapping duplication entry point								
SYNOPSIS	<pre>#include <sys/sunddi.h> int prefix mapdev_dup(ddi_mapdev_handle_t handle, void *devprivate, ddi_mapdev_handle_t new_handle, void **new_devprivatep);</pre>								
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).								
PARAMETERS	<table><tr><td><i>handle</i></td><td>The handle of the mapping that is being duplicated.</td></tr><tr><td><i>devprivate</i></td><td>Driver private mapping data from the mapping that is being duplicated.</td></tr><tr><td><i>new_handle</i></td><td>An opaque pointer to the duplicated device mapping.</td></tr><tr><td><i>new_devprivatep</i></td><td>A pointer to be filled in by the driver with the driver private mapping data for the duplicated device mapping.</td></tr></table>	<i>handle</i>	The handle of the mapping that is being duplicated.	<i>devprivate</i>	Driver private mapping data from the mapping that is being duplicated.	<i>new_handle</i>	An opaque pointer to the duplicated device mapping.	<i>new_devprivatep</i>	A pointer to be filled in by the driver with the driver private mapping data for the duplicated device mapping.
<i>handle</i>	The handle of the mapping that is being duplicated.								
<i>devprivate</i>	Driver private mapping data from the mapping that is being duplicated.								
<i>new_handle</i>	An opaque pointer to the duplicated device mapping.								
<i>new_devprivatep</i>	A pointer to be filled in by the driver with the driver private mapping data for the duplicated device mapping.								
DESCRIPTION	<p>Future releases of Solaris will provide this function for binary and source compatibility. However, for increased functionality, use <code>devmap_dup(9E)</code> instead. See <code>devmap_dup(9E)</code> for details.</p> <p><code>mapdev_dup()</code> is called when a device mapping is duplicated such as through <code>fork(2)</code>. <code>mapdev_dup()</code> is expected to generate new driver private data for the new mapping, and set <i>new_devprivatep</i> to point to it. <i>new_handle</i> is the handle of the new mapped object.</p> <p>A non-zero return value from <code>mapdev_dup()</code> will cause the corresponding operation, such as <code>fork()</code> to fail.</p>								
RETURN VALUES	<code>mapdev_dup()</code> returns 0 for success or the appropriate error number on failure.								
CONTEXT	This function is called from user context only.								
EXAMPLES	<p>EXAMPLE 1</p> <pre>static int xxmapdev_dup(ddi_mapdev_handle_t handle, void *devprivate, ddi_mapdev_handle_t new_handle, void **new_devprivate) { struct xxpvtdata *pvtdata; /* Allocate a new private data structure */ pvtdata = kmem_alloc(sizeof (struct xxpvtdata), KM_SLEEP); /* Copy the old data to the new - device dependent*/ ... /* Return the new data */ *new_pvtdata = pvtdata; return (0); }</pre>								

mapdev_dup(9E)

EXAMPLE 1 *(Continued)*

SEE ALSO fork(2), mmap(2), mapdev_access(9E), mapdev_free(9E), segmap(9E),
ddi_mapdev(9F), ddi_mapdev_intercept(9F), ddi_mapdev_nointercept(9F),
ddi_mapdev_ctl(9S)

Writing Device Drivers

mapdev_free(9E)

NAME	mapdev_free – device mapping free entry point
SYNOPSIS	<pre>#include <sys/sunddi.h> void <i>prefix</i>mapdev_free(ddi_mapdev_handle_t handle, void *devprivate);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI).
PARAMETERS	<p><i>handle</i> An opaque pointer to a device mapping.</p> <p><i>devprivate</i> Driver private mapping data from ddi_mapdev(9F).</p>
DESCRIPTION	<p>Future releases of Solaris will provide this function for binary and source compatibility. However, for increased functionality, use devmap_unmap(9E) instead. See devmap_unmap(9E) for details.</p> <p>mapdev_free() is called when a mapping created by ddi_mapdev(9F) is being destroyed. mapdev_free() receives the <i>handle</i> of the mapping being destroyed and a pointer to the driver private data for this mapping in <i>devprivate</i>.</p> <p>The mapdev_free() routine is expected to free any resources that were allocated by the driver for this mapping.</p>
CONTEXT	This function is called from user context only.
EXAMPLES	<p>EXAMPLE 1 Using mapdev_free()</p> <pre>static void xxmapdev_free(ddi_mapdev_handle_t hdl, void *pvtdata) { /* Destroy the driver private data - Device dependent */ ... kmem_free(pvtdata, sizeof (struct xxpvtdata)); }</pre>
SEE ALSO	<p>exit(2), mmap(2), munmap(2), mapdev_access(9E), mapdev_dup(9E), segmap(9E), ddi_mapdev(9F), ddi_mapdev_intercept(9F), ddi_mapdev_nointercept(9F), ddi_mapdev_ctl(9S)</p> <p><i>Writing Device Drivers</i></p>

NAME	mmap – check virtual mapping for memory mapped device										
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/cred.h> #include <sys/mman.h> #include <sys/ddi.h> int <i>prefix</i>mmap(dev_t dev, off_t off, int prot);</pre>										
INTERFACE LEVEL PARAMETERS	<p>This interface is obsolete. devmap(9E) should be used instead.</p> <p><i>dev</i> Device whose memory is to be mapped.</p> <p><i>off</i> Offset within device memory at which mapping begins.</p> <p><i>prot</i> A bit field that specifies the protections this page of memory will receive. Possible settings are:</p> <table border="0"> <tr> <td>PROT_READ</td> <td>Read access will be granted.</td> </tr> <tr> <td>PROT_WRITE</td> <td>Write access will be granted.</td> </tr> <tr> <td>PROT_EXEC</td> <td>Execute access will be granted.</td> </tr> <tr> <td>PROT_USER</td> <td>User-level access will be granted.</td> </tr> <tr> <td>PROT_ALL</td> <td>All access will be granted.</td> </tr> </table>	PROT_READ	Read access will be granted.	PROT_WRITE	Write access will be granted.	PROT_EXEC	Execute access will be granted.	PROT_USER	User-level access will be granted.	PROT_ALL	All access will be granted.
PROT_READ	Read access will be granted.										
PROT_WRITE	Write access will be granted.										
PROT_EXEC	Execute access will be granted.										
PROT_USER	User-level access will be granted.										
PROT_ALL	All access will be granted.										
DESCRIPTION	<p>Future releases of Solaris will provide this function for binary and source compatibility. However, for increased functionality, use devmap(9E) instead. See devmap(9E) for details.</p> <p>The mmap() entry point is a required entry point for character drivers supporting memory-mapped devices. A memory mapped device has memory that can be mapped into a process's address space. The mmap(2) system call, when applied to a character special file, allows this device memory to be mapped into user space for direct access by the user application.</p> <p>The mmap() entry point is called as a result of an mmap(2) system call, and also as a result of a page fault. mmap() is called to translate the offset <i>off</i> in device memory to the corresponding physical page frame number.</p> <p>The mmap() entry point checks if the offset <i>off</i> is within the range of pages exported by the device. For example, a device that has 512 bytes of memory that can be mapped into user space should not support offsets greater than 512. If the offset does not exist, then -1 is returned. If the offset does exist, mmap() returns the value returned by hat_getkpfnum(9F) for the physical page in device memory containing the offset <i>off</i>.</p> <p>hat_getkpfnum(9F) accepts a kernel virtual address as an argument. A kernel virtual address can be obtained by calling ddi_regs_map_setup(9F) in the driver's attach(9E) routine. The corresponding ddi_regs_map_free(9F) call can be made in the driver's detach(9E) routine. Refer to Example 1 below for more information.</p>										

mmap(9E)

`mmap()` should only be supported for memory-mapped devices. See the `segmap(9E)` and `ddi_mapdev(9F)` reference pages for further information on memory-mapped device drivers.

If a device driver shares data structures with the application, for example through exported kernel memory, and the driver gets recompiled for a 64-bit kernel but the application remains 32-bit, the binary layout of any data structures will be incompatible if they contain longs or pointers. The driver needs to know whether there is a model mismatch between the current thread and the kernel and take necessary action. `ddi_mmap_get_model(9F)` can be used to get the C Language Type Model which the current thread expects. In combination with `ddi_model_convert_from(9F)` the driver can determine whether there is a data model mismatch between the current thread and the device driver. The device driver might have to adjust the shape of data structures before exporting them to a user thread which supports a different data model. See `ddi_mmap_get_model(9F)` for an example.

RETURN VALUES If the protection and offset are valid for the device, the driver should return the value returned by `hat_getkpfnum(9F)`, for the page at offset *off* in the device's memory. If not, -1 should be returned.

EXAMPLES **EXAMPLE 1** The `mmap()` Entry Point

The following is an example of the `mmap()` entry point. If offset *off* is valid, `hat_getkpfnum(9F)` is called to obtain the page frame number corresponding to this offset in the device's memory. In this example, `xsp→regp→csr` is a kernel virtual address which maps to device memory. `ddi_regs_map_setup(9F)` can be used to obtain this address. For example, `ddi_regs_map_setup(9F)` can be called in the driver's `attach(9E)` routine. The resulting kernel virtual address is stored in the `xxstate` structure, which is accessible from the driver's `mmap()` entry point. See `ddi_soft_state(9F)`. The corresponding `ddi_regs_map_free(9F)` call can be made in the driver's `detach(9E)` routine.

```
struct reg {
    uint8_t    csr;
    uint8_t    data;
};
struct xxstate {
    . . .
    struct reg    *regp
    . . .
};

struct xxstate *xsp;
. . .

static int
xxmmap(dev_t dev, off_t off, int prot)
{
    int        instance;
    struct xxstate *xsp;
```

EXAMPLE 1 The `mmap()` Entry Point (Continued)

```

        /* No write access */
        if (prot & PROT_WRITE)
            return (-1);

        instance = getminor(dev);
        xsp = ddi_get_soft_state(statep, instance);
        if (xsp == NULL)
            return (-1);

        /* check for a valid offset */
        if ( off is invalid )
            return (-1);
        return (hat_getkpfnum (xsp->regp->csr + off));
    }

```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Obsolete

SEE ALSO `mmap(2)`, `attributes(5)`, `attach(9E)`, `detach(9E)`, `devmap(9E)`, `segmap(9E)`, `ddi_btop(9F)`, `ddi_get_soft_state(9F)`, `ddi_mmap_get_model(9F)`, `ddi_model_convert_from(9F)`, `ddi_regs_map_free(9F)`, `ddi_regs_map_setup(9F)`, `ddi_soft_state(9F)`, `devmap_setup(9F)`, `getminor(9F)`, `hat_getkpfnum(9F)`

Writing Device Drivers

NOTES For some devices, mapping device memory in the driver's `attach(9E)` routine and unmapping device memory in the driver's `detach(9E)` routine is a sizeable drain on system resources. This is especially true for devices with a large amount of physical address space.

One alternative is to create a mapping for only the first page of device memory in `attach(9E)`. If the device memory is contiguous, a kernel page frame number may be obtained by calling `hat_getkpfnum(9F)` with the kernel virtual address of the first page of device memory and adding the desired page offset to the result. The page offset may be obtained by converting the byte offset *off* to pages. See `ddi_btop(9F)`.

Another alternative is to call `ddi_regs_map_setup(9F)` and `ddi_regs_map_free(9F)` in `mmap()`. These function calls would bracket the call to `hat_getkpfnum(9F)`.

mmap(9E)

However, note that the above alternatives may not work in all cases. The existence of intermediate nexus devices with memory management unit translation resources that are not locked down may cause unexpected and undefined behavior.

NAME	open – gain access to a device												
Block and Character	<pre>#include <sys/types.h> #include <sys/file.h> #include <sys/errno.h> #include <sys/open.h> #include <sys/cred.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre>												
STREAMS	<pre>int prefixopen(dev_t *devp, int flag, int otyp, cred_t *cred_p); #include <sys/file.h> #include <sys/stream.h> #include <sys/ddi.h> #include <sys/sunddi.h></pre> <pre>int prefixopen(queue_t *q, dev_t *devp, int oflag, int sflag, cred_t *cred_p);</pre>												
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI). This entry point is required, but it can be nulldev(9F)												
Block and Character	<p><i>devp</i> Pointer to a device number.</p> <p><i>flag</i> A bit field passed from the user program open(2) system call that instructs the driver on how to open the file. Valid settings are:</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="padding-right: 1em;">FEXCL</td> <td>Open the device with exclusive access; fail all other attempts to open the device.</td> </tr> <tr> <td style="padding-right: 1em;">FNDELAY</td> <td>Open the device and return immediately. Do not block the open even if something is wrong.</td> </tr> <tr> <td style="padding-right: 1em;">FREAD</td> <td>Open the device with read-only permission (if ORed with FWRITE, then allow both read and write access)</td> </tr> <tr> <td style="padding-right: 1em;">FWRITE</td> <td>Open a device with write-only permission (if ORed with FREAD, then allow both read and write access)</td> </tr> </table> <p><i>otyp</i> Parameter supplied so that the driver can determine how many times a device was opened and for what reasons. For OTYP_BLK and OTYP_CHR, the open() routine may be called many times, but the close(9E) routine is called only when the last reference to a device is removed. If the device is accessed through file descriptors, this is by a call to close(2) or exit(2) If the device is accessed through memory mapping, this is by a call to munmap(2) or exit(2) For OTYP_LYR, there is exactly one close(9E) for each open() called. This permits software drivers to exist above hardware drivers and removes any ambiguity from the hardware driver regarding how a device is used.</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="padding-right: 1em;">OTYP_BLK</td> <td>Open occurred through block interface for the device</td> </tr> <tr> <td style="padding-right: 1em;">OTYP_CHR</td> <td>Open occurred through the raw/character interface for the device</td> </tr> </table>	FEXCL	Open the device with exclusive access; fail all other attempts to open the device.	FNDELAY	Open the device and return immediately. Do not block the open even if something is wrong.	FREAD	Open the device with read-only permission (if ORed with FWRITE, then allow both read and write access)	FWRITE	Open a device with write-only permission (if ORed with FREAD, then allow both read and write access)	OTYP_BLK	Open occurred through block interface for the device	OTYP_CHR	Open occurred through the raw/character interface for the device
FEXCL	Open the device with exclusive access; fail all other attempts to open the device.												
FNDELAY	Open the device and return immediately. Do not block the open even if something is wrong.												
FREAD	Open the device with read-only permission (if ORed with FWRITE, then allow both read and write access)												
FWRITE	Open a device with write-only permission (if ORed with FREAD, then allow both read and write access)												
OTYP_BLK	Open occurred through block interface for the device												
OTYP_CHR	Open occurred through the raw/character interface for the device												

open(9E)

	<code>OTYP_LYR</code>	Open a layered process. This flag is used when one driver calls another driver's <code>open()</code> or <code>close(9E)</code> routine. The calling driver will make sure that there is one layered close for each layered open. This flag applies to both block and character devices.
STREAMS	<code>cred_p</code>	Pointer to the user credential structure.
	<code>q</code>	A pointer to the read queue.
	<code>devp</code>	Pointer to a device number. For STREAMS modules, <code>devp</code> always points to the device number associated with the driver at the end (tail) of the stream.
	<code>oflag</code>	Valid <code>oflag</code> values are <code>FEXCL</code> , <code>FNDELAY</code> , <code>FREAD</code> , and <code>FWRITE</code> , the same as those listed above for <code>flag</code> . For STREAMS modules, <code>oflag</code> is always set to 0.
	<code>sflag</code>	Valid values are as follows:
		<code>CLONEOPEN</code>
	<code>MODOPEN</code>	Modules should be called with <code>sflag</code> set to this value. Modules should return an error if they are called with <code>sflag</code> set to a different value. Drivers should return an error if they are called with <code>sflag</code> set to this value.
	0	Indicates a driver is opened directly, without calling the clone driver.
	<code>cred_p</code>	Pointer to the user credential structure.
DESCRIPTION	<p>The driver's <code>open()</code> routine is called by the kernel during an <code>open(2)</code> or a <code>mount(2)</code> on the special file for the device. The routine should verify that the minor number component of <code>*devp</code> is valid, that the type of access requested by <code>otyp</code> and <code>flag</code> is appropriate for the device, and, if required, check permissions using the user credentials pointed to by <code>cred_p</code>.</p> <p>The <code>open()</code> routine is passed a pointer to a device number so that the driver can change the minor number. This allows drivers to dynamically create minor instances of the device. An example of this might be a pseudo-terminal driver that creates a new pseudo-terminal whenever it is opened. A driver that chooses the minor number dynamically, normally creates only one minor device node in <code>attach(9E)</code> with <code>ddi_create_minor_node(9F)</code> then changes the minor number component of <code>*devp</code> using <code>makedevice(9F)</code> and <code>getmajor(9F)</code>. The driver needs to keep track of available minor numbers internally:</p> <pre>*devp = makedevice(getmajor(*devp), new_minor);</pre>	
RETURN VALUES	The <code>open()</code> routine should return 0 for success, or the appropriate error number.	

open(9E)

SEE ALSO close(2), exit(2), mmap(2), mount(2), munmap(2), open(2), intro(9E), attach(9E), close(9E), ddi_create_minor_node(9F), getmajor(9F), getminor(9F), makedevice(9F), nulldev(9F)

Writing Device Drivers

STREAMS Programming Guide

WARNINGS Do not attempt to change the major number.

power(9E)

NAME	power – power a device attached to the system
SYNOPSIS	<pre>#include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>power(dev_info_t *dip, int component, int level);</pre>
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). This entry point is required. If the driver writer does not supply this entry point, the value NULL must be used in the <code>cb_ops(9S)</code> structure instead.
PARAMETERS	<p><i>dip</i> Pointer to the device's <code>dev_info</code> structure.</p> <p><i>component</i> Component of the driver to be managed.</p> <p><i>level</i> Desired component power level.</p>
DESCRIPTION	<p>The <code>power(9E)</code> function is the device-specific Power Management entry point. This function is called when the system wants the driver to set the power level of <i>component</i> to <i>level</i>.</p> <p>The <code>power()</code> entry points behave as described under NEW BEHAVIOR if the device driver uses new automatic device Power Management interfaces (driver exports <code>pm-components(9P)</code> property). The <code>power()</code> entry point behaves as described under OBSOLETE BEHAVIOR if the device driver uses original Power Management interfaces (driver does not export the <code>pm-components</code> property but instead calls <code>pm_create_components(9F)</code>). The behavior described under OBSOLETE BEHAVIOR is obsolete and will be removed in a future release.</p>
NEW BEHAVIOR	<p>The <i>level</i> argument is the driver-defined power level to which the component needs to be set. Except for power level 0, which is interpreted by the framework to mean "powered off," the interpretation of <i>level</i> is entirely up to the driver.</p> <p>The <i>component</i> argument is the component of the device to be power-managed. The interpretation of <i>component</i> is entirely up to the driver.</p> <p>When a requested power transition would cause the device to lose state, the driver must save the state of the device in memory. When a requested power transition requires state to be restored, the driver must restore that state.</p> <p>If a requested power transition for one component requires another component to change power state before it can be completed, the driver must call <code>pm_raise_power(9F)</code> to get the other component changed, and the <code>power(9E)</code> entry point must support being re-entered.</p> <p>If the system requests an inappropriate power transition for the device (for example, a request to power down a device which has just become busy), then the power level should not be changed and power should return <code>DDI_FAILURE</code>.</p>

OBSOLETE BEHAVIOR

The *level* argument is the driver-defined power level to which *component* must be set. Except for power level 0 which is defined by the framework to mean “powered off,” the interpretation of *level* is entirely up to the driver.

The *component* argument is the component of the device to be power-managed. Except for component 0, which must represent the entire device, the interpretation of *component* is entirely up to the driver.

The `power()` function can assume that the driver will be suspended (using `detach(9E)` with command `DDI_PM_SUSPEND`), before a request is made to set component 0 to power level 0 and resumed (using `attach(9E)` with command `DDI_PM_RESUME`) after setting component 0 from power level 0 to a non-zero power level.

If the system requests an inappropriate power transition for the device (for example, a request to power down a device which has just become busy), then the power level should not be changed and the `power(9E)` function should return `DDI_FAILURE`.

RETURN VALUES

The `power()` function returns:

`DDI_SUCCESS` Successfully set the power to the requested *level*.

`DDI_FAILURE` Failed to set the power to the requested *level*.

CONTEXT

The `power()` function is called from user or kernel context only.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Evolving (Interfaces under OBSOLETE BEHAVIOR are obsolete.)

SEE ALSO

`attach(9E)`, `detach(9E)`, `pm_busy_component(9F)`, `pm_create_components(9F)`, `pm_destroy_components(9F)`, `pm_idle_component(9F)`, `pm_raise_power(9F)`, `pm-components(9P)`, `cb_ops(9S)`

Writing Device Drivers

Using Power Management

print(9E)

NAME	print – display a driver message on system console
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/errno.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>print(dev_t dev, char *str);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI). This entry point is required for block devices.
PARAMETERS	<i>dev</i> Device number. <i>str</i> Pointer to a character string describing the problem.
DESCRIPTION	The <code>print()</code> routine is called by the kernel when it has detected an exceptional condition (such as out of space) in the device. To display the message on the console, the driver should use the <code>cmn_err(9F)</code> kernel function. The driver should print the message along with any driver specific information.
RETURN VALUES	The <code>print()</code> routine should return 0 for success, or the appropriate error number. The <code>print</code> routine can fail if the driver implemented a non-standard <code>print()</code> routine that attempted to perform error logging, but was unable to complete the logging for whatever reason.
SEE ALSO	<code>cmn_err(9F)</code> <i>Writing Device Drivers</i>

NAME	probe – determine if a non-self-identifying device is present								
SYNOPSIS	<pre>#include <sys/conf.h> #include <sys/ddi.h> #include <sys/sunddi.h> static int <i>prefix</i>probe (dev_info_t *dip);</pre>								
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). This entry point is required for non-self-identifying devices. You must write it for such devices. For self-identifying devices, nulldev(9F) should be specified in the dev_ops(9S) structure if a probe routine is not necessary.								
ARGUMENTS	<i>dip</i> Pointer to the device's dev_info structure.								
DESCRIPTION	<p>probe() determines whether the device corresponding to <i>dip</i> actually exists and is a valid device for this driver. probe() is called after identify(9E) and before attach(9E) for a given <i>dip</i>. For example, the probe() routine can map the device registers using ddi_map_regs(9F) then attempt to access the hardware using ddi_peek(9F) or ddi_poke(9F) and determine if the device exists. Then the device registers should be unmapped using ddi_unmap_regs(9F).</p> <p>To probe a device that was left powered off after the last detach(), it might be necessary to power it up. If so, the driver must power up the device by accessing device registers directly. pm_raise_power(9F) will not be available until attach(9E). The framework ensures that the ancestors of the node being probed and all relevant platform-specific power management hardware is at full power at the time that probe() is called.</p> <p>probe() should only probe the device. It should not change any software state and should not create any software state. Device initialization should be done in attach(9E).</p> <p>For a self-identifying device, this entry point is not necessary. However, if a device exists in both self-identifying and non-self-identifying forms, a probe() routine can be provided to simplify the driver. ddi_dev_is_sid(9F) can then be used to determine whether probe() needs to do any work. See ddi_dev_is_sid(9F) for an example.</p>								
RETURN VALUES	<table border="0"> <tr> <td>DDI_PROBE_SUCCESS</td> <td>If the probe was successful.</td> </tr> <tr> <td>DDI_PROBE_FAILURE</td> <td>If the probe failed.</td> </tr> <tr> <td>DDI_PROBE_DONTCARE</td> <td>If the probe was unsuccessful, yet attach(9E) should still be called.</td> </tr> <tr> <td>DDI_PROBE_PARTIAL</td> <td>If the instance is not present now, but may be present in the future.</td> </tr> </table>	DDI_PROBE_SUCCESS	If the probe was successful.	DDI_PROBE_FAILURE	If the probe failed.	DDI_PROBE_DONTCARE	If the probe was unsuccessful, yet attach(9E) should still be called.	DDI_PROBE_PARTIAL	If the instance is not present now, but may be present in the future.
DDI_PROBE_SUCCESS	If the probe was successful.								
DDI_PROBE_FAILURE	If the probe failed.								
DDI_PROBE_DONTCARE	If the probe was unsuccessful, yet attach(9E) should still be called.								
DDI_PROBE_PARTIAL	If the instance is not present now, but may be present in the future.								
SEE ALSO	attach(9E), identify(9E), ddi_dev_is_sid(9F), ddi_map_regs(9F), ddi_peek(9F), ddi_poke(9F), nulldev(9F), dev_ops(9S)								
	<i>Writing Device Drivers</i>								

prop_op(9E)

NAME	prop_op – report driver property information														
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ddi.h> #include <sys/sunddi.h> int prefixprop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op, int flags, char *name, caddr_t valuelp, int *lengthp);</pre>														
INTERFACE LEVEL	Solaris DDI specific (Solaris DDI). This entry point is required, but it can be ddi_prop_op(9F).														
ARGUMENTS	<table><tr><td><i>dev</i></td><td>Device number associated with this device.</td></tr><tr><td><i>dip</i></td><td>A pointer to the device information structure for this device.</td></tr><tr><td><i>prop_op</i></td><td>Property operator. Valid operators are: PROP_LEN Get property length only. (<i>valuelp</i> unaffected). PROP_LEN_AND_VAL_BUF Get length and value into caller's buffer. (<i>valuelp</i> used as input). PROP_LEN_AND_VAL_ALLOC Get length and value into allocated buffer. (<i>valuelp</i> returned as pointer to pointer to allocated buffer).</td></tr><tr><td><i>flags</i></td><td>The only possible flag value is: DDI_PROP_DONTPASS Do not pass request to parent if property not found.</td></tr><tr><td><i>name</i></td><td>Pointer to name of property to be interrogated.</td></tr><tr><td><i>valuelp</i></td><td>If <i>prop_op</i> is PROP_LEN_AND_VAL_BUF, this should be a pointer to the user's buffer. If <i>prop_op</i> is PROP_LEN_AND_VAL_ALLOC, this should be the <i>address</i> of a pointer.</td></tr><tr><td><i>lengthp</i></td><td>On exit, <i>*lengthp</i> will contain the property length. If <i>prop_op</i> is PROP_LEN_AND_VAL_BUF then <i>lengthp</i> should point to an int that contains the length of caller's buffer, before calling prop_op().</td></tr></table>	<i>dev</i>	Device number associated with this device.	<i>dip</i>	A pointer to the device information structure for this device.	<i>prop_op</i>	Property operator. Valid operators are: PROP_LEN Get property length only. (<i>valuelp</i> unaffected). PROP_LEN_AND_VAL_BUF Get length and value into caller's buffer. (<i>valuelp</i> used as input). PROP_LEN_AND_VAL_ALLOC Get length and value into allocated buffer. (<i>valuelp</i> returned as pointer to pointer to allocated buffer).	<i>flags</i>	The only possible flag value is: DDI_PROP_DONTPASS Do not pass request to parent if property not found.	<i>name</i>	Pointer to name of property to be interrogated.	<i>valuelp</i>	If <i>prop_op</i> is PROP_LEN_AND_VAL_BUF, this should be a pointer to the user's buffer. If <i>prop_op</i> is PROP_LEN_AND_VAL_ALLOC, this should be the <i>address</i> of a pointer.	<i>lengthp</i>	On exit, <i>*lengthp</i> will contain the property length. If <i>prop_op</i> is PROP_LEN_AND_VAL_BUF then <i>lengthp</i> should point to an int that contains the length of caller's buffer, before calling prop_op().
<i>dev</i>	Device number associated with this device.														
<i>dip</i>	A pointer to the device information structure for this device.														
<i>prop_op</i>	Property operator. Valid operators are: PROP_LEN Get property length only. (<i>valuelp</i> unaffected). PROP_LEN_AND_VAL_BUF Get length and value into caller's buffer. (<i>valuelp</i> used as input). PROP_LEN_AND_VAL_ALLOC Get length and value into allocated buffer. (<i>valuelp</i> returned as pointer to pointer to allocated buffer).														
<i>flags</i>	The only possible flag value is: DDI_PROP_DONTPASS Do not pass request to parent if property not found.														
<i>name</i>	Pointer to name of property to be interrogated.														
<i>valuelp</i>	If <i>prop_op</i> is PROP_LEN_AND_VAL_BUF, this should be a pointer to the user's buffer. If <i>prop_op</i> is PROP_LEN_AND_VAL_ALLOC, this should be the <i>address</i> of a pointer.														
<i>lengthp</i>	On exit, <i>*lengthp</i> will contain the property length. If <i>prop_op</i> is PROP_LEN_AND_VAL_BUF then <i>lengthp</i> should point to an int that contains the length of caller's buffer, before calling prop_op().														
DESCRIPTION	prop_op() is an entry point which reports the values of certain properties of the driver or device to the system. Each driver must have a <i>prefix</i> prop_op entry point, but most drivers that do not need to create or manage their own properties can use ddi_prop_op() for this entry point. Then the driver can use ddi_prop_update(9F) to create properties for its device.														
RETURN VALUES	prop_op() should return: DDI_PROP_SUCCESS Property found and returned. DDI_PROP_NOT_FOUND Property not found.														

DDI_PROP_UNDEFINED	Prop explicitly undefined.
DDI_PROP_NO_MEMORY	Property found, but unable to allocate memory. <i>lengthp</i> has the correct property length.
DDI_PROP_BUF_TOO_SMALL	Property found, but the supplied buffer is too small. <i>lengthp</i> has the correct property length.

EXAMPLES **EXAMPLE 1** Using `prop_op()` to Report Property Information

In the following example, `prop_op()` intercepts requests for the *temperature* property. The driver tracks changes to *temperature* using a variable in the state structure in order to avoid frequent calls to `ddi_prop_update(9F)`. The *temperature* property is only updated when a request is made for this property. It then uses the system routine `ddi_prop_op(9E)` to process the property request. If the property request is not specific to a device, the driver does not intercept the request. This is indicated when the value of the *dev* parameter is equal to `DDI_DEV_T_ANY`.

```
int temperature;    /* current device temperature */
.
.
.
static int
xxprop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op,
          int flags, char *name, caddr_t valuep, int *lengthp)
{
    int instance;
    struct xxstate *xsp;
    if (dev == DDI_DEV_T_ANY)
        goto skip;
    instance = getminor(dev);
    xsp = ddi_get_soft_state(statep, instance);
    if (xsp == NULL)
        return (DDI_PROP_NOT_FOUND);
    if (strcmp(name, "temperature") == 0) {
        ddi_prop_update_int(dev, dip, "temperature", temperature);
    }
    /* other cases... */
skip:
    return (ddi_prop_op(dev, dip, prop_op, flags, name, valuep, lengthp));
}
```

SEE ALSO `Intro(9E)`, `ddi_prop_op(9E)`, `ddi_prop_update(9F)`

Writing Device Drivers

put(9E)

NAME	put – receive messages from the preceding queue
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stream.h> #include <sys/stropts.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefixrput</i>(queue_t *q, mblk_t *mp); /* read side */ int <i>prefixwput</i>(queue_t *q, mblk_t *mp); /* write side */</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI). This entry point is required for STREAMS.
ARGUMENTS	<p><i>q</i> Pointer to the queue(9S) structure.</p> <p><i>mp</i> Pointer to the message block.</p>
DESCRIPTION	<p>The primary task of the put () routine is to coordinate the passing of messages from one queue to the next in a stream. The put () routine is called by the preceding stream component (stream module, driver, or stream head). put () routines are designated “write” or “read” depending on the direction of message flow.</p> <p>With few exceptions, a streams module or driver must have a put () routine. One exception is the read side of a driver, which does not need a put () routine because there is no component downstream to call it. The put () routine is always called before the component’s corresponding srv(9E) (service) routine, and so put () should be used for the immediate processing of messages.</p> <p>A put () routine must do at least one of the following when it receives a message:</p> <ul style="list-style-type: none">■ pass the message to the next component on the stream by calling the putnext(9F) function;■ process the message, if immediate processing is required (for example, to handle high priority messages); or■ enqueue the message (with the putq(9F) function) for deferred processing by the service srv(9E) routine. <p>Typically, a put () routine will switch on message type, which is contained in the db_type member of the datab structure pointed to by mp. The action taken by the put () routine depends on the message type. For example, a put () routine might process high priority messages, enqueue normal messages, and handle an unrecognized M_IOCTL message by changing its type to M_IOCNAK (negative acknowledgement) and sending it back to the stream head using the qreply(9F) function.</p>

put(9E)

The putq(9F) function can be used as a module's put () routine when no special processing is required and all messages are to be enqueued for the srv(9E) routine.

RETURN VALUES Ignored.

CONTEXT put () routines do not have user context.

SEE ALSO srv(9E), putctl(9F), putctl1(9F), putnext(9F), putnextctl(9F), putnextctl1(9F), putq(9F), qreply(9F), queue(9S), streamtab(9S)

Writing Device Drivers

STREAMS Programming Guide

read(9E)

NAME	read – read data from a device
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/errno.h> #include <sys/open.h> #include <sys/uio.h> #include <sys/cred.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>read(dev_t dev, struct uio *uio_p, cred_t *cred_p);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). This entry point is <i>optional</i> . <i>dev</i> Device number. <i>uio_p</i> Pointer to the uio(9S) structure that describes where the data is to be stored in user space. <i>cred_p</i> Pointer to the user credential structure for the I/O transaction.
DESCRIPTION	The driver read() routine is called indirectly through cb_ops(9S) by the read(2) system call. The read() routine should check the validity of the minor number component of <i>dev</i> and the user credential structure pointed to by <i>cred_p</i> (if pertinent). The read() routine should supervise the data transfer into the user space described by the uio(9S) structure.
RETURN VALUES	The read() routine should return 0 for success, or the appropriate error number.
EXAMPLES	EXAMPLE 1 read() routine using physio() The following is an example of a read() routine using physio(9F) to perform reads from a non-seekable device: <pre>static int xxread(dev_t dev, struct uio *uiop, cred_t *credp) { int rval; offset_t off; int instance; xx_t xx; instance = getminor(dev); xx = ddi_get_soft_state(xxstate, instance); if (xx == NULL) return (ENXIO); off = uiop->uio_loffset; rval = physio(xxstrategy, NULL, dev, B_READ, xxmin, uiop); uiop->uio_loffset = off; return (rval); }</pre>
SEE ALSO	read(2), write(9E), physio(9F), cb_ops(9S), uio(9S)

read(9E)

Writing Device Drivers

segmap(9E)

NAME	segmap – map device memory into user space														
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mman.h> #include <sys/param.h> #include <sys/vm.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>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 *cred_p);</pre>														
INTERFACE LEVEL ARGUMENTS	<p>Architecture independent level 2 (DKI only).</p> <p><i>dev</i> Device whose memory is to be mapped.</p> <p><i>off</i> Offset within device memory at which mapping begins.</p> <p><i>asp</i> Pointer to the address space into which the device memory should be mapped.</p> <p><i>addrp</i> Pointer to the address in the address space to which the device memory should be mapped.</p> <p><i>len</i> Length (in bytes) of the memory to be mapped.</p> <p><i>prot</i> A bit field that specifies the protections. Possible settings are:</p> <table border="0" style="margin-left: 2em;"> <tr> <td>PROT_READ</td> <td>Read access is desired.</td> </tr> <tr> <td>PROT_WRITE</td> <td>Write access is desired.</td> </tr> <tr> <td>PROT_EXEC</td> <td>Execute access is desired.</td> </tr> <tr> <td>PROT_USER</td> <td>User-level access is desired (the mapping is being done as a result of a mmap(2) system call).</td> </tr> <tr> <td>PROT_ALL</td> <td>All access is desired.</td> </tr> </table> <p><i>maxprot</i> Maximum protection flag possible for attempted mapping; the PROT_WRITE bit may be masked out if the user opened the special file read-only.</p> <p><i>flags</i> Flags indicating type of mapping. Possible values are (other bits may be set):</p> <table border="0" style="margin-left: 2em;"> <tr> <td>MAP_SHARED</td> <td>Changes should be shared.</td> </tr> <tr> <td>MAP_PRIVATE</td> <td>Changes are private.</td> </tr> </table> <p><i>cred_p</i> Pointer to the user credentials structure.</p>	PROT_READ	Read access is desired.	PROT_WRITE	Write access is desired.	PROT_EXEC	Execute access is desired.	PROT_USER	User-level access is desired (the mapping is being done as a result of a mmap(2) system call).	PROT_ALL	All access is desired.	MAP_SHARED	Changes should be shared.	MAP_PRIVATE	Changes are private.
PROT_READ	Read access is desired.														
PROT_WRITE	Write access is desired.														
PROT_EXEC	Execute access is desired.														
PROT_USER	User-level access is desired (the mapping is being done as a result of a mmap(2) system call).														
PROT_ALL	All access is desired.														
MAP_SHARED	Changes should be shared.														
MAP_PRIVATE	Changes are private.														

DESCRIPTION	<p>The <code>segmap()</code> entry point is an optional routine for character drivers that support memory mapping. The <code>mmap(2)</code> system call, when applied to a character special file, allows device memory to be mapped into user space for direct access by the user application.</p> <p>Typically, a character driver that needs to support the <code>mmap(2)</code> system call supplies either an <code>devmap(9E)</code> entry point, or both an <code>devmap(9E)</code> and a <code>segmap()</code> entry point routine (see the <code>devmap(9E)</code> reference page). If no <code>segmap()</code> entry point is provided for the driver, <code>devmap_setup(9F)</code> is used as a default.</p> <p>A driver for a memory-mapped device would provide a <code>segmap()</code> entry point if it:</p> <ul style="list-style-type: none"> ■ needs to maintain a separate context for each user mapping. See <code>devmap_setup(9F)</code> for details. ■ needs to assign device access attributes to the user mapping. <p>The responsibilities of a <code>segmap()</code> entry point are:</p> <ul style="list-style-type: none"> ■ Verify that the range, defined by <i>offset</i> and <i>len</i>, to be mapped is valid for the device. Typically, this task is performed by calling the <code>devmap(9E)</code> entry point. Note that if you are using <code>ddi_devmap_segmap(9F)</code> or <code>devmap_setup(9F)</code> to set up the mapping, it will call your <code>devmap(9E)</code> entry point for you to validate the range to be mapped. ■ Assign device access attributes to the mapping. See <code>ddi_devmap_segmap(9F)</code>, and <code>ddi_device_acc_attr(9S)</code> for details. ■ Set up device contexts for the user mapping if your device requires context switching. See <code>devmap_setup(9F)</code> for details. ■ Perform the mapping with <code>ddi_devmap_segmap(9F)</code>, or <code>devmap_setup(9F)</code> and return the status if it fails.
RETURN VALUES	<p>The <code>segmap()</code> routine should return 0 if the driver is successful in performing the memory map of its device address space into the specified address space.</p> <p>The <code>segmap()</code> must return an error number on failure. For example, valid error numbers would be <code>ENXIO</code> if the offset/length pair specified exceeds the limits of the device memory, or <code>EINVAL</code> if the driver detects an invalid type of mapping attempted.</p> <p>If one of the mapping routines <code>ddi_devmap_segmap()</code> or <code>devmap_setup()</code> fails, you must return the error number returned by the respective routine.</p>
SEE ALSO	<p><code>mmap(2)</code>, <code>devmap(9E)</code>, <code>devmap_setup(9F)</code>, <code>ddi_devmap_segmap(9F)</code>, <code>ddi_device_acc_attr(9S)</code></p> <p><i>Writing Device Drivers</i></p>

srv(9E)

NAME	srv – service queued messages
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stream.h> #include <sys/stropts.h> #include <sys/ddi.h> #include <sys/sunddi.h> intprefixrsrv(queue_t *q); /* read side */ intprefixwsrv(queue_t *q); /* write side */</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI). This entry point is required for STREAMS.
ARGUMENTS	<i>q</i> Pointer to the queue(9S) structure.
DESCRIPTION	<p>The optional service <i>srv</i> () routine may be included in a STREAMS module or driver for many possible reasons, including:</p> <ul style="list-style-type: none">■ to provide greater control over the flow of messages in a stream;■ to make it possible to defer the processing of some messages to avoid depleting system resources;■ to combine small messages into larger ones, or break large messages into smaller ones;■ to recover from resource allocation failure. A module's or driver's <i>put</i>(9E) routine can test for the availability of a resource, and if it is not available, enqueue the message for later processing by the <i>srv</i> () routine. <p>A message is first passed to a module's or driver's <i>put</i>(9E) routine, which may or may not do some processing. It must then either:</p> <ul style="list-style-type: none">■ Pass the message to the next stream component with <i>putnext</i>(9F).■ If a <i>srv</i> () routine has been included, it may call <i>putq</i>(9F) to place the message on the queue. <p>Once a message has been enqueued, the STREAMS scheduler controls the service routine's invocation. The scheduler calls the service routines in FIFO order. The scheduler cannot guarantee a maximum delay <i>srv</i> () routine to be called except that it will happen before any user level process are run.</p> <p>Every stream component (stream head, module or driver) has limit values it uses to implement flow control. Each component should check the tunable high and low water marks to stop and restart the flow of message processing. Flow control limits apply only between two adjacent components with <i>srv</i> () routines.</p>

STREAMS messages can be defined to have up to 256 different priorities to support requirements for multiple bands of data flow. At a minimum, a stream must distinguish between normal (priority zero) messages and high priority messages (such as `M_IOCACK`). High priority messages are always placed at the head of the `srv()` routine's queue, after any other enqueued high priority messages. Next are messages from all included priority bands, which are enqueued in decreasing order of priority. Each priority band has its own flow control limits. If a flow controlled band is stopped, all lower priority bands are also stopped.

Once the STREAMS scheduler calls a `srv()` routine, it must process all messages on its queue. The following steps are general guidelines for processing messages. Keep in mind that many of the details of how a `srv()` routine should be written depend of the implementation, the direction of flow (upstream or downstream), and whether it is for a module or a driver.

1. Use `getq(9F)` to get the next enqueued message.
2. If the message is high priority, process (if appropriate) and pass to the next stream component with `putnext(9F)`.
3. If it is not a high priority message (and therefore subject to flow control), attempt to send it to the next stream component with a `srv()` routine. Use `bcanputnext(9F)` to determine if this can be done.
4. If the message cannot be passed, put it back on the queue with `putbq(9F)`. If it can be passed, process (if appropriate) and pass with `putnext()`.

RETURN VALUES Ignored.

SEE ALSO `put(9E)`, `bcanput(9F)`, `bcanputnext(9F)`, `canput(9F)`, `canputnext(9F)`, `getq(9F)`, `nulldev(9F)`, `putbq(9F)`, `putnext(9F)`, `putq(9F)`, `qinit(9S)`, `queue(9S)`

WARNINGS Each stream module must specify a read and a write service `srv()` routine. If a service routine is not needed (because the `put()` routine processes all messages), a `NULL` pointer should be placed in module's `qinit(9S)` structure. Do not use `nulldev(9F)` instead of the `NULL` pointer. Use `ofnulldev(9F)` for a `srv()` routine can result in flow control errors.

strategy(9E)

NAME	strategy – perform block I/O
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/buf.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>strategy(struct buf *bp);</pre>
INTERFACE LEVEL	Architecture independent level 1 (DDI/DKI). This entry point is required for block devices.
PARAMETERS	<i>bp</i> Pointer to the buf(9S) structure.
DESCRIPTION	The strategy() routine is called indirectly (through cb_ops(9S)) by the kernel to read and write blocks of data on the block device. strategy() may also be called directly or indirectly to support the raw character interface of a block device (read(9E), write(9E) and ioctl(9E)). The strategy() routine's responsibility is to set up and initiate the transfer.
RETURN VALUES	The strategy() routine should always return 0. On an error condition, it should call bioerror(9f) to set b_flags to the proper error code, and call biodone(9f). Note that a partial transfer is not considered to be an error.
SEE ALSO	ioctl(9E), read(9E), write(9E), biodone(9F), bioerror(9F), buf(9S), cb_ops(9S) <i>Writing Device Drivers</i>

NAME	tran_abort – abort a SCSI command
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int <i>prefix</i>tran_abort(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL ARGUMENTS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i> Pointer to a scsi_address(9S) structure.</p> <p><i>pkt</i> Pointer to a scsi_pkt(9S) structure.</p>
DESCRIPTION	<p>The tran_abort () vector in the scsi_hba_tran(9S) structure must be initialized during the HBA driver's attach(9E) to point to an HBA entry point to be called when a target driver calls scsi_abort(9F).</p> <p>tran_abort () should attempt to abort the command <i>pkt</i> that has been transported to the HBA. If <i>pkt</i> is NULL, the HBA driver should attempt to abort all outstanding packets for the target/logical unit addressed by <i>ap</i>.</p> <p>Depending on the state of a particular command in the transport layer, the HBA driver may not be able to abort the command.</p> <p>While the abort is taking place, packets issued to the transported layer may or may not be aborted.</p> <p>For each packet successfully aborted, tran_abort () must set the pkt_reason to CMD_ABORTED, and pkt_statistics must be OR'ed with STAT_ABORTED .</p>
RETURN VALUES	<p>tran_abort () must return:</p> <p>1 upon success or partial success.</p> <p>0 upon failure.</p>
SEE ALSO	<p>attach(9E), scsi_abort(9F), scsi_hba_attach(9F), scsi_address(9S), scsi_hba_tran(9S), scsi_pkt(9S)</p> <p><i>Writing Device Drivers</i></p>
NOTES	<p>If pkt_reason already indicates that an earlier error had occurred, tran_abort () should not overwrite pkt_reason with CMD_ABORTED.</p>

tran_bus_reset(9e)

NAME tran_bus_reset – reset a SCSI bus

SYNOPSIS #include <sys/scsi/scsi.h> int *prefix*
tran_bus_reset(dev_info_t **hba_dip*, int *level*);

INTERFACE LEVEL Solaris DDI

PARAMETERS *hba_dip* The dev_info_t pointer associated with the SCSI HBA.
level The level of reset required.

DESCRIPTION The tran_bus_reset() vector in the scsi_hba_tran(9S) structure should be initialized during the HBA driver's attach(9E). It is an HBA entry point to be called when a user initiates a bus reset through device control interfaces.

tran_bus_reset() must reset the SCSI bus without resetting targets.

level will be one of the following:

RESET_BUS Reset the SCSI bus only, not the targets.

Implementation is hardware specific. If it is not possible to reset the SCSI bus without changing the state and operating mode of the targets, the HBA driver should not initialize this vector or return failure.

RETURN VALUES tran_bus_reset() should return:

1 on success.
0 on failure.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Evolving

SEE ALSO attributes(5), tran_quiesce(9E), scsi_hba_tran(9S)

NAME	tran_init_pkt, tran_destroy_pkt – SCSI HBA packet preparation and deallocation
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> struct scsi_pkt *prefixtran_init_pkt(struct scsi_address *ap, struct scsi_pkt *pkt, struct buf *bp, int cmdlen, int statuslen, int tgtlen, int flags, int (*callback, caddr_t), caddr_t arg); void prefixtran_destroy_pkt(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i> Pointer to a scsi_address(9S) structure.</p> <p><i>pkt</i> Pointer to a scsi_pkt(9S) structure allocated in an earlier call, or NULL.</p> <p><i>bp</i> Pointer to a buf(9S) structure if DMA resources are to be allocated for the <i>pkt</i>, or NULL.</p> <p><i>cmdlen</i> The required length for the SCSI command descriptor block (CDB) in bytes.</p> <p><i>statuslen</i> The required length for the SCSI status completion block (SCB) in bytes.</p> <p><i>tgtlen</i> The length of the packet private area within the <i>scsi_pkt</i> to be allocated on behalf of the SCSI target driver.</p> <p><i>flags</i> Flags for creating the packet.</p> <p><i>callback</i> Pointer to either NULL_FUNC or SLEEP_FUNC.</p> <p><i>arg</i> Always NULL.</p>
DESCRIPTION	<p>The tran_init_pkt() and tran_destroy_pkt() vectors in the scsi_hba_tran structure must be initialized during the HBA driver's attach(9E) to point to HBA entry points to be called when a target driver calls scsi_init_pkt(9F) and scsi_destroy_pkt(9F).</p> <p>tran_init_pkt() tran_init_pkt() is the entry point into the HBA which is used to allocate and initialize a scsi_pkt structure on behalf of a SCSI target driver. If <i>pkt</i> is NULL, the HBA driver must use scsi_hba_pkt_alloc(9F) to allocate a new scsi_pkt structure.</p> <p>If <i>bp</i> is non-NULL, the HBA driver must allocate appropriate DMA resources for the <i>pkt</i>, for example, through ddi_dma_buf_setup(9F) or ddi_dma_buf_bind_handle(9F).</p> <p>If the PKT_CONSISTENT bit is set in <i>flags</i>, the buffer was allocated by scsi_alloc_consistent_buf(9F). For packets marked with PKT_CONSISTENT, the HBA driver must synchronize any cached data transfers before calling the target driver's command completion callback.</p>

tran_destroy_pkt(9E)

If the `PKT_DMA_PARTIAL` bit is set in *flags*, the HBA driver should set up partial data transfers, such as setting the `DDI_DMA_PARTIAL` bit in the *flags* argument if interfaces such as `ddi_dma_buf_setup(9F)` or `ddi_dma_buf_bind_handle(9F)` are used.

If only partial DMA resources are available, `tran_init_pkt()` must return in the `pkt_resid` field of *pkt* the number of bytes of DMA resources not allocated.

If both *pkt* and *bp* are non-NULL, if the `PKT_DMA_PARTIAL` bit is set in *flags*, and if DMA resources have already been allocated for the *pkt* with a previous call to `tran_init_pkt()` that returned a non-zero `pkt_resid` field, this request is to move the DMA resources for the subsequent piece of the transfer.

The contents of `scsi_address(9S)` pointed to by *ap* are copied into the `pkt_address` field of the `scsi_pkt(9S)` by `scsi_hba_pkt_alloc(9F)`.

tgrlen is the length of the packet private area in the `scsi_pkt` structure to be allocated on behalf of the SCSI target driver.

statuslen is the required length for the SCSI status completion block. If the requested status length is greater than or equal to `sizeof(struct scsi_arq_status)` and the `auto_rqsense` capability has been set, automatic request sense (ARS) is enabled for this packet. If the status length is less than `sizeof(struct scsi_arq_status)`, automatic request sense must be disabled for this *pkt*.

If the HBA driver is not capable of disabling ARQ on a per-packet basis and `tran_init_pkt()` is called with a *statuslen* that is less than `sizeof(struct scsi_arq_status)`, the driver's `tran_init_pkt` routine should allocate at least `sizeof(struct scsi_arq_status)`. If an ARS is needed, upon successful ARS done by the HBA driver, the driver must copy the sense data over and set `STAT_ARQ_DONE` in `pkt_state`.

cmdlen is the required length for the SCSI command descriptor block.

Note: *tgrlen*, *statuslen*, and *cmdlen* are used only when the HBA driver allocates the `scsi_pkt(9S)`, in other words, when *pkt* is NULL.

callback indicates what the allocator routines should do when resources are not available:

`NULL_FUNC` Do not wait for resources. Return a NULL pointer.

`SLEEP_FUNC` Wait indefinitely for resources.

tran_destroy_pkt() `tran_destroy_pkt()` is the entry point into the HBA that must free all of the resources that were allocated to the `scsi_pkt(9S)` structure during `tran_init_pkt()`.

RETURN VALUES `tran_init_pkt()` must return a pointer to a `scsi_pkt(9S)` structure on success, or NULL on failure.

tran_destroy_pkt(9E)

If *pkt* is NULL on entry, and `tran_init_pkt()` allocated a packet through `scsi_hba_pkt_alloc(9F)` but was unable to allocate DMA resources, `tran_init_pkt()` must free the packet through `scsi_hba_pkt_free(9F)` before returning NULL.

SEE ALSO `attach(9E)`, `tran_sync_pkt(9E)`, `biodone(9F)`, `bioerror(9F)`, `ddi_dma_buf_bind_handle(9F)`, `ddi_dma_buf_setup(9F)`, `scsi_alloc_consistent_buf(9F)`, `scsi_destroy_pkt(9F)`, `scsi_hba_attach(9F)`, `scsi_hba_pkt_alloc(9F)`, `scsi_hba_pkt_free(9F)`, `scsi_init_pkt(9F)`, `buf(9S)`, `scsi_address(9S)`, `scsi_hba_tran(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

NOTES If a DMA allocation request fails with `DDI_DMA_NOMAPPING`, indicate the error by calling `bioerror(9F)` with *bp* and an error code of `EFAULT`.

If a DMA allocation request fails with `DDI_DMA_TOOBIG`, indicate the error by calling `bioerror(9F)` with *bp* and an error code of `EINVAL`.

tran_dmafree(9E)

NAME	tran_dmafree – SCSI HBA DMA deallocation entry point
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void <i>prefix</i>tran_dmafree(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL ARGUMENTS	Solaris architecture specific (Solaris DDI). <i>ap</i> A pointer to a <i>scsi_address</i> structure. See <i>scsi_address(9S)</i> . <i>pkt</i> A pointer to a <i>scsi_pkt</i> structure. See <i>scsi_pkt(9S)</i> .
DESCRIPTION	<p>The <i>tran_dmafree()</i> vector in the <i>scsi_hba_tran</i> structure must be initialized during the HBA driver's <i>attach()</i> to point to an HBA entry point to be called when a target driver calls <i>scsi_dmafree(9F)</i>. See <i>attach(9E)</i> and <i>scsi_hba_tran(9S)</i>.</p> <p><i>tran_dmafree()</i> must deallocate any DMA resources previously allocated to this <i>pkt</i> in a call to <i>tran_init_pkt(9E)</i>. <i>tran_dmafree()</i> should not free the structure pointed to by <i>pkt</i> itself. Since <i>tran_destroy_pkt(9E)</i> must also free DMA resources, it is important that the HBA driver keeps accurate note of whether <i>scsi_pkt(9S)</i> structures have DMA resources allocated.</p>
SEE ALSO	<i>attach(9E)</i> , <i>tran_destroy_pkt(9E)</i> , <i>tran_init_pkt(9E)</i> , <i>scsi_dmafree(9F)</i> , <i>scsi_dmaget(9F)</i> , <i>scsi_hba_attach(9F)</i> , <i>scsi_init_pkt(9F)</i> , <i>scsi_address(9S)</i> , <i>scsi_hba_tran(9S)</i> , <i>scsi_pkt(9S)</i>
	<i>Writing Device Drivers</i>
NOTES	A target driver may call <i>tran_dmafree()</i> on packets for which no DMA resources were allocated.

NAME	tran_getcap, tran_setcap – get/set SCSI transport capability
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int prefixtran_getcap(struct scsi_address *ap, char *cap, int whom); int prefixtran_setcap(struct scsi_address *ap, char *cap, int value, int whom);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i> Pointer to the <code>scsi_address(9S)</code> structure.</p> <p><i>cap</i> Pointer to the string capability identifier.</p> <p><i>value</i> Defines the new state of the capability.</p> <p><i>whom</i> Specifies whether all targets or only the specified target is affected.</p>
DESCRIPTION	<p>The <code>tran_getcap()</code> and <code>tran_setcap()</code> vectors in the <code>scsi_hba_tran(9S)</code> structure must be initialized during the HBA driver's <code>attach(9E)</code> to point to HBA entry points to be called when a target driver calls <code>scsi_ifgetcap(9F)</code> and <code>scsi_ifsetcap(9F)</code>.</p> <p><code>tran_getcap()</code> is called to get the current value of a capability specific to features provided by the HBA hardware or driver. The name of the capability <i>cap</i> is the NULL terminated capability string.</p> <p>If <i>whom</i> is non-zero, the request is for the current value of the capability defined for the target specified by the <code>scsi_address(9S)</code> structure pointed to by <i>ap</i>; if <i>whom</i> is 0, all targets are affected; else, the target specified by the <code>scsi_address</code> structure pointed to by <i>ap</i> is affected.</p> <p><code>tran_setcap()</code> is called to set the value of the capability <i>cap</i> to the value of <i>value</i>. If <i>whom</i> is non-zero, the capability should be set for the target specified by the <code>scsi_address(9S)</code> structure pointed to by <i>ap</i>; if <i>whom</i> is 0, all targets are affected; else, the target specified by the <code>scsi_address</code> structure pointed to by <i>ap</i> is affected. It is recommended that HBA drivers do not support setting capabilities for all targets, that is, <i>whom</i> is 0.</p> <p>A device may support only a subset of the defined capabilities.</p> <p>Refer to <code>scsi_ifgetcap(9F)</code> for the list of defined capabilities.</p> <p>HBA drivers should use <code>scsi_hba_lookup_capstr(9F)</code> to match <i>cap</i> against the canonical capability strings.</p>
RETURN VALUES	<code>tran_setcap()</code> must return 1 if the capability was successfully set to the new value, 0 if the HBA driver does not support changing the capability, and -1 if the capability was not defined.

tran_getcap(9E)

tran_getcap() must return the current value of a capability or -1 if the capability was not defined.

SEE ALSO attach(9E), scsi_hba_attach(9F), scsi_hba_lookup_capstr(9F), scsi_ifgetcap(9F), scsi_address(9S), scsi_hba_tran(9S)

Writing Device Drivers

NAME	tran_init_pkt, tran_destroy_pkt – SCSI HBA packet preparation and deallocation
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> struct scsi_pkt *prefixtran_init_pkt(struct scsi_address *ap, struct scsi_pkt *pkt, struct buf *bp, int cmdlen, int statuslen, int tgtlen, int flags, int (*callback, caddr_t), caddr_t arg); void prefixtran_destroy_pkt(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i> Pointer to a scsi_address(9S) structure.</p> <p><i>pkt</i> Pointer to a scsi_pkt(9S) structure allocated in an earlier call, or NULL.</p> <p><i>bp</i> Pointer to a buf(9S) structure if DMA resources are to be allocated for the <i>pkt</i>, or NULL.</p> <p><i>cmdlen</i> The required length for the SCSI command descriptor block (CDB) in bytes.</p> <p><i>statuslen</i> The required length for the SCSI status completion block (SCB) in bytes.</p> <p><i>tgtlen</i> The length of the packet private area within the scsi_pkt to be allocated on behalf of the SCSI target driver.</p> <p><i>flags</i> Flags for creating the packet.</p> <p><i>callback</i> Pointer to either NULL_FUNC or SLEEP_FUNC.</p> <p><i>arg</i> Always NULL.</p>
DESCRIPTION	<p>The tran_init_pkt() and tran_destroy_pkt() vectors in the scsi_hba_tran structure must be initialized during the HBA driver's attach(9E) to point to HBA entry points to be called when a target driver calls scsi_init_pkt(9F) and scsi_destroy_pkt(9F).</p> <p>tran_init_pkt() tran_init_pkt() is the entry point into the HBA which is used to allocate and initialize a scsi_pkt structure on behalf of a SCSI target driver. If <i>pkt</i> is NULL, the HBA driver must use scsi_hba_pkt_alloc(9F) to allocate a new scsi_pkt structure.</p> <p>If <i>bp</i> is non-NULL, the HBA driver must allocate appropriate DMA resources for the <i>pkt</i>, for example, through ddi_dma_buf_setup(9F) or ddi_dma_buf_bind_handle(9F).</p> <p>If the PKT_CONSISTENT bit is set in <i>flags</i>, the buffer was allocated by scsi_alloc_consistent_buf(9F). For packets marked with PKT_CONSISTENT, the HBA driver must synchronize any cached data transfers before calling the target driver's command completion callback.</p>

tran_init_pkt(9E)

If the `PKT_DMA_PARTIAL` bit is set in *flags*, the HBA driver should set up partial data transfers, such as setting the `DDI_DMA_PARTIAL` bit in the *flags* argument if interfaces such as `ddi_dma_buf_setup(9F)` or `ddi_dma_buf_bind_handle(9F)` are used.

If only partial DMA resources are available, `tran_init_pkt()` must return in the `pkt_resid` field of *pkt* the number of bytes of DMA resources not allocated.

If both *pkt* and *bp* are non-NULL, if the `PKT_DMA_PARTIAL` bit is set in *flags*, and if DMA resources have already been allocated for the *pkt* with a previous call to `tran_init_pkt()` that returned a non-zero `pkt_resid` field, this request is to move the DMA resources for the subsequent piece of the transfer.

The contents of `scsi_address(9S)` pointed to by *ap* are copied into the `pkt_address` field of the `scsi_pkt(9S)` by `scsi_hba_pkt_alloc(9F)`.

tgrlen is the length of the packet private area in the `scsi_pkt` structure to be allocated on behalf of the SCSI target driver.

statuslen is the required length for the SCSI status completion block. If the requested status length is greater than or equal to `sizeof(struct scsi_arq_status)` and the `auto_rqsense` capability has been set, automatic request sense (ARS) is enabled for this packet. If the status length is less than `sizeof(struct scsi_arq_status)`, automatic request sense must be disabled for this *pkt*.

If the HBA driver is not capable of disabling ARQ on a per-packet basis and `tran_init_pkt()` is called with a *statuslen* that is less than `sizeof(struct scsi_arq_status)`, the driver's `tran_init_pkt` routine should allocate at least `sizeof(struct scsi_arq_status)`. If an ARS is needed, upon successful ARS done by the HBA driver, the driver must copy the sense data over and set `STAT_ARQ_DONE` in `pkt_state`.

cmdlen is the required length for the SCSI command descriptor block.

Note: *tgrlen*, *statuslen*, and *cmdlen* are used only when the HBA driver allocates the `scsi_pkt(9S)`, in other words, when *pkt* is NULL.

callback indicates what the allocator routines should do when resources are not available:

`NULL_FUNC` Do not wait for resources. Return a NULL pointer.

`SLEEP_FUNC` Wait indefinitely for resources.

tran_destroy_pkt() `tran_destroy_pkt()` is the entry point into the HBA that must free all of the resources that were allocated to the `scsi_pkt(9S)` structure during `tran_init_pkt()`.

RETURN VALUES `tran_init_pkt()` must return a pointer to a `scsi_pkt(9S)` structure on success, or NULL on failure.

tran_init_pkt(9E)

If *pkt* is NULL on entry, and `tran_init_pkt()` allocated a packet through `scsi_hba_pkt_alloc(9F)` but was unable to allocate DMA resources, `tran_init_pkt()` must free the packet through `scsi_hba_pkt_free(9F)` before returning NULL.

SEE ALSO `attach(9E)`, `tran_sync_pkt(9E)`, `biodone(9F)`, `bioerror(9F)`, `ddi_dma_buf_bind_handle(9F)`, `ddi_dma_buf_setup(9F)`, `scsi_alloc_consistent_buf(9F)`, `scsi_destroy_pkt(9F)`, `scsi_hba_attach(9F)`, `scsi_hba_pkt_alloc(9F)`, `scsi_hba_pkt_free(9F)`, `scsi_init_pkt(9F)`, `buf(9S)`, `scsi_address(9S)`, `scsi_hba_tran(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

NOTES If a DMA allocation request fails with `DDI_DMA_NOMAPPING`, indicate the error by calling `bioerror(9F)` with *bp* and an error code of `EFAULT`.

If a DMA allocation request fails with `DDI_DMA_TOOBIG`, indicate the error by calling `bioerror(9F)` with *bp* and an error code of `EINVAL`.

tran_quiesce(9e)

NAME tran_quiesce, tran_unquiesce – quiesce and unquiesce a SCSI bus

SYNOPSIS #include <sys/scsi/scsi.h>

```
int prefixtran_quiesce(dev_info_t *hba_dip);
int prefixtran_unquiesce(dev_info_t *hba_dip);
```

INTERFACE LEVEL Solaris DDI

PARAMETERS *hba_dip* The dev_info_t pointer associated with the SCSI HBA.

DESCRIPTION The tran_quiesce() and tran_unquiesce() vectors in the scsi_hba_tran(9S) structure should be initialized during the HBA driver's attach(9E). They are HBA entry points to be called when a user initiates quiesce and unquiesce operations through device control interfaces.

tran_quiesce() should wait for all outstanding commands to complete and blocks (or queues) any I/O requests issued. tran_unquiesce() should allow I/O activities to resume on the SCSI bus.

Implementation is hardware specific.

RETURN VALUES tran_quiesce() and tran_unquiesce() should return:

- 1 on success.
- 0 on failure.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Evolving

SEE ALSO attributes(5), tran_bus_reset(9E), scsi_hba_tran(9S)

NAME	tran_reset – reset a SCSI bus or target						
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int prefixtran_reset(struct scsi_address *ap, int level);</pre>						
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i> Pointer to the <code>scsi_address(9S)</code> structure.</p> <p><i>level</i> The level of reset required.</p>						
DESCRIPTION	<p>The <code>tran_reset()</code> vector in the <code>scsi_hba_tran(9S)</code> structure must be initialized during the HBA driver's <code>attach(9E)</code> to point to an HBA entry point to be called when a target driver calls <code>scsi_reset(9F)</code>.</p> <p><code>tran_reset()</code> must reset either the SCSI bus, a SCSI target device, or a SCSI logical unit as specified by <i>level</i>.</p> <p><i>level</i> must be one of the following:</p> <table border="0"> <tr> <td>RESET_ALL</td> <td>Reset the SCSI bus.</td> </tr> <tr> <td>RESET_TARGET</td> <td>Reset the target specified by <i>ap</i>.</td> </tr> <tr> <td>RESET_LUN</td> <td>Reset the logical unit specified by <i>ap</i>.</td> </tr> </table> <p><code>tran_reset</code> should set the <code>pkt_reason</code> field of all outstanding packets in the transport layer associated with each target or logical unit that was successfully reset to <code>CMD_RESET</code> and the <code>pkt_statistics</code> field must be OR'ed with either <code>STAT_BUS_RESET</code> (if the SCSI bus was reset) or <code>STAT_DEV_RESET</code> (if the target or logical unit was reset).</p> <p>The HBA driver should use a SCSI Bus Device Reset Message to reset a target device. The HBA driver should use a SCSI Logical Unit Reset Message to reset a logical unit.</p> <p>Packets that are in the transport layer but not yet active on the bus should be returned with <code>pkt_reason</code> set to <code>CMD_RESET</code> and <code>pkt_statistics</code> OR'ed with <code>STAT_ABORTED</code>.</p> <p>Support for <code>RESET_LUN</code> is optional but strongly encouraged for new and updated HBA drivers. If an HBA driver provides <code>RESET_LUN</code> support, it must also create the <code>lun-reset</code> capability with a value of zero for each target device instance represented by a valid <i>ap</i>. The HBA is also required to provide the means to return the current value of the <code>lun-reset</code> capability in its <code>tran_getcap(9E)</code> routine, as well as the means to change the value of the <code>lun_reset</code> capability in its <code>tran_setcap(9E)</code> routine.</p>	RESET_ALL	Reset the SCSI bus.	RESET_TARGET	Reset the target specified by <i>ap</i> .	RESET_LUN	Reset the logical unit specified by <i>ap</i> .
RESET_ALL	Reset the SCSI bus.						
RESET_TARGET	Reset the target specified by <i>ap</i> .						
RESET_LUN	Reset the logical unit specified by <i>ap</i> .						
RETURN VALUES	<p><code>tran_reset()</code> should return:</p> <table border="0"> <tr> <td>1</td> <td>on success.</td> </tr> <tr> <td>0</td> <td>on failure.</td> </tr> </table>	1	on success.	0	on failure.		
1	on success.						
0	on failure.						

tran_reset(9E)

SEE ALSO | attach(9E), ddi_dma_buf_setup(9F), scsi_hba_attach(9F), scsi_reset(9F),
scsi_address(9S), scsi_hba_tran(9S)

Writing Device Drivers

NOTES | If `pkt_reason` already indicates that an earlier error had occurred for a particular *pkt*,
`tran_reset()` should not overwrite `pkt_reason` with `CMD_RESET`.

NAME	tran_reset_notify – request to notify SCSI target of bus reset				
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int prefixtran_reset_notify(struct scsi_address *ap, int flag, void (*callback, caddr_t), caddr_t arg);</pre>				
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i> Pointer to the <code>scsi_address(9S)</code> structure.</p> <p><i>flag</i> A flag indicating registration or cancellation of a notification request.</p> <p><i>callback</i> A pointer to the target driver's reset notification function.</p> <p><i>arg</i> The callback function argument.</p>				
DESCRIPTION	<p>The <code>tran_reset_notify()</code> entry point is called when a target driver requests notification of a bus reset.</p> <p>The <code>tran_reset_notify()</code> vector in the <code>scsi_hba_tran(9S)</code> structure may be initialized in the HBA driver's <code>attach(9E)</code> routine to point to the HBA entry point to be called when a target driver calls <code>scsi_reset_notify(9F)</code>.</p> <p>The argument <i>flag</i> is used to register or cancel the notification. The supported values for <i>flag</i> are as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>SCSI_RESET_NOTIFY</code></td> <td>Register <i>callback</i> as the reset notification function for the target.</td> </tr> <tr> <td><code>SCSI_RESET_CANCEL</code></td> <td>Cancel the reset notification request for the target.</td> </tr> </table> <p>The HBA driver maintains a list of reset notification requests registered by the target drivers. When a bus reset occurs, the HBA driver notifies registered target drivers by calling the callback routine, <i>callback</i>, with the argument, <i>arg</i>, for each registered target.</p>	<code>SCSI_RESET_NOTIFY</code>	Register <i>callback</i> as the reset notification function for the target.	<code>SCSI_RESET_CANCEL</code>	Cancel the reset notification request for the target.
<code>SCSI_RESET_NOTIFY</code>	Register <i>callback</i> as the reset notification function for the target.				
<code>SCSI_RESET_CANCEL</code>	Cancel the reset notification request for the target.				
RETURN VALUES	<p>For <code>SCSI_RESET_NOTIFY</code> requests, <code>tran_reset_notify()</code> must return <code>DDI_SUCCESS</code> if the notification request has been accepted, and <code>DDI_FAILURE</code> otherwise.</p> <p>For <code>SCSI_RESET_CANCEL</code> requests, <code>tran_reset_notify()</code> must return <code>DDI_SUCCESS</code> if the notification request has been canceled, and <code>DDI_FAILURE</code> otherwise.</p>				
SEE ALSO	<p><code>attach(9E)</code>, <code>scsi_ifgetcap(9F)</code>, <code>scsi_reset_notify(9F)</code>, <code>scsi_address(9S)</code>, <code>scsi_hba_tran(9S)</code></p> <p><i>Writing Device Drivers</i></p>				

tran_setcap(9E)

NAME	tran_getcap, tran_setcap – get/set SCSI transport capability
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int prefixtran_getcap(struct scsi_address *ap, char *cap, int whom); int prefixtran_setcap(struct scsi_address *ap, char *cap, int value, int whom);</pre>
INTERFACE LEVEL PARAMETERS	Solaris architecture specific (Solaris DDI). <i>ap</i> Pointer to the <code>scsi_address(9S)</code> structure. <i>cap</i> Pointer to the string capability identifier. <i>value</i> Defines the new state of the capability. <i>whom</i> Specifies whether all targets or only the specified target is affected.
DESCRIPTION	<p>The <code>tran_getcap()</code> and <code>tran_setcap()</code> vectors in the <code>scsi_hba_tran(9S)</code> structure must be initialized during the HBA driver's <code>attach(9E)</code> to point to HBA entry points to be called when a target driver calls <code>scsi_ifgetcap(9F)</code> and <code>scsi_ifsetcap(9F)</code>.</p> <p><code>tran_getcap()</code> is called to get the current value of a capability specific to features provided by the HBA hardware or driver. The name of the capability <i>cap</i> is the NULL terminated capability string.</p> <p>If <i>whom</i> is non-zero, the request is for the current value of the capability defined for the target specified by the <code>scsi_address(9S)</code> structure pointed to by <i>ap</i>; if <i>whom</i> is 0, all targets are affected; else, the target specified by the <code>scsi_address</code> structure pointed to by <i>ap</i> is affected.</p> <p><code>tran_setcap()</code> is called to set the value of the capability <i>cap</i> to the value of <i>value</i>. If <i>whom</i> is non-zero, the capability should be set for the target specified by the <code>scsi_address(9S)</code> structure pointed to by <i>ap</i>; if <i>whom</i> is 0, all targets are affected; else, the target specified by the <code>scsi_address</code> structure pointed to by <i>ap</i> is affected. It is recommended that HBA drivers do not support setting capabilities for all targets, that is, <i>whom</i> is 0.</p> <p>A device may support only a subset of the defined capabilities.</p> <p>Refer to <code>scsi_ifgetcap(9F)</code> for the list of defined capabilities.</p> <p>HBA drivers should use <code>scsi_hba_lookup_capstr(9F)</code> to match <i>cap</i> against the canonical capability strings.</p>
RETURN VALUES	<code>tran_setcap()</code> must return 1 if the capability was successfully set to the new value, 0 if the HBA driver does not support changing the capability, and -1 if the capability was not defined.

tran_setcap(9E)

tran_getcap() must return the current value of a capability or -1 if the capability was not defined.

SEE ALSO attach(9E), scsi_hba_attach(9F), scsi_hba_lookup_capstr(9F), scsi_ifgetcap(9F), scsi_address(9S), scsi_hba_tran(9S)

Writing Device Drivers

tran_start(9E)

NAME	tran_start – request to transport a SCSI command
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int prefixtran_start(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL PARAMETERS	Solaris architecture specific (Solaris DDI). <i>pkt</i> Pointer to the scsi_pkt(9S) structure that is about to be transferred. <i>ap</i> Pointer to a scsi_address(9S) structure.
DESCRIPTION	<p>The tran_start() vector in the scsi_hba_tran(9S) structure must be initialized during the HBA driver's attach(9E) to point to an HBA entry point to be called when a target driver calls scsi_transport(9F).</p> <p>tran_start() must perform the necessary operations on the HBA hardware to transport the SCSI command in the <i>pkt</i> structure to the target/logical unit device specified in the <i>ap</i> structure.</p> <p>If the flag FLAG_NOINTR is set in <i>pkt_flags</i> in <i>pkt</i>, tran_start() should not return until the command has been completed. The command completion callback <i>pkt_comp</i> in <i>pkt</i> must not be called for commands with FLAG_NOINTR set, since the return is made directly to the function invoking scsi_transport(9F).</p> <p>When the flag FLAG_NOINTR is not set, tran_start() must queue the command for execution on the hardware and return immediately. The member <i>pkt_comp</i> in <i>pkt</i> indicates a callback routine to be called upon command completion.</p> <p>Refer to scsi_pkt(9S) for other bits in <i>pkt_flags</i> for which the HBA driver may need to adjust how the command is managed.</p> <p>If the auto_rqsense capability has been set, and the status length allocated in tran_init_pkt(9E) is greater than or equal to sizeof(struct scsi_arq_status), automatic request sense is enabled for this <i>pkt</i>. If the command terminates with a Check Condition, the HBA driver must arrange for a Request Sense command to be transported to that target/logical unit, and the members of the scsi_arq_status structure pointed to by <i>pkt_scbp</i> updated with the results of this Request Sense command before the HBA driver completes the command pointed by <i>pkt</i>.</p> <p>The member <i>pkt_time</i> in <i>pkt</i> is the maximum number of seconds in which the command should complete. Timeout starts when the command is transmitted on the SCSI bus. A <i>pkt_time</i> of 0 means no timeout should be performed.</p> <p>For a command which has timed out, the HBA driver must perform some recovery operation to clear the command in the target, typically an Abort message, or a Device or Bus Reset. The <i>pkt_reason</i> member of the timed out <i>pkt</i> should be set to CMD_TIMEOUT, and <i>pkt_statistics</i> OR'ed with STAT_TIMEOUT. If the HBA driver can successfully recover from the timeout, <i>pkt_statistics</i> must also be OR'ed with one of STAT_ABORTED, STAT_BUS_RESET, or STAT_DEV_RESET, as</p>

appropriate. This informs the target driver that timeout recovery has already been successfully accomplished for the timed out command. The `pkt_comp` completion callback, if not `NULL`, must also be called at the conclusion of the timeout recovery.

If the timeout recovery was accomplished with an Abort Tag message, only the timed out packet is affected, and the packet must be returned with `pkt_statistics` OR'ed with `STAT_ABORTED` and `STAT_TIMEOUT`.

If the timeout recovery was accomplished with an Abort message, all commands active in that target are affected. All corresponding packets must be returned with `pkt_reason`, `CMD_TIMEOUT`, and `pkt_statistics` OR'ed with `STAT_TIMEOUT` and `STAT_ABORTED`.

If the timeout recovery was accomplished with a Device Reset, all packets corresponding to commands active in the target must be returned in the transport layer for this target. Packets corresponding to commands active in the target must be returned with `pkt_reason` set to `CMD_TIMEOUT`, and `pkt_statistics` OR'ed with `STAT_DEV_RESET` and `STAT_TIMEOUT`. Currently inactive packets queued for the device should be returned with `pkt_reason` set to `CMD_RESET` and `pkt_statistics` OR'ed with `STAT_ABORTED`.

If the timeout recovery was accomplished with a Bus Reset, all packets corresponding to commands active in the target must be returned in the transport layer. Packets corresponding to commands active in the target must be returned with `pkt_reason` set to `CMD_TIMEOUT` and `pkt_statistics` OR'ed with `STAT_TIMEOUT` and `STAT_BUS_RESET`. All queued packets for other targets on this bus must be returned with `pkt_reason` set to `CMD_RESET` and `pkt_statistics` OR'ed with `STAT_ABORTED`.

Note that after either a Device Reset or a Bus Reset, the HBA driver must enforce a reset delay time of `'scsi-reset-delay'` milliseconds, during which time no commands should be sent to that device, or any device on the bus, respectively.

`tran_start()` should initialize the following members in `pkt` to 0. Upon command completion, the HBA driver should ensure that the values in these members are updated to accurately reflect the states through which the command transitioned while in the transport layer.

<code>pkt_resid</code>	For commands with data transfer, this member must be updated to indicate the residual of the data transferred.
<code>pkt_reason</code>	The reason for the command completion. This field should be set to <code>CMD_CMPLT</code> at the beginning of <code>tran_start()</code> , then updated if the command ever transitions to an abnormal termination state. To avoid losing information, do not set <code>pkt_reason</code> to any other error state unless it still has its original <code>CMD_CMPLT</code> value.
<code>pkt_statistics</code>	Bit field of transport-related statistics.

tran_start(9E)

pkt_state Bit field with the major states through which a SCSI command can transition. Note: The members listed above, and pkt_hba_private member, are the only fields in the scsi_pkt(9S) structure which may be modified by the transport layer.

RETURN VALUES tran_start() must return:

TRAN_ACCEPT The packet was accepted by the transport layer.

TRAN_BUSY The packet could not be accepted because there was already a packet in progress for this target/logical unit, the HBA queue was full, or the target device queue was full.

TRAN_BADPKT The DMA count in the packet exceeded the DMA engine's maximum DMA size, or the packet could not be accepted for other reasons.

TRAN_FATAL_ERROR A fatal error has occurred in the HBA.

SEE ALSO attach(9E), tran_init_pkt(9E), scsi_hba_attach(9F), scsi_transport(9F), scsi_address(9S), scsi_arq_status(9S), scsi_hba_tran(9S), scsi_pkt(9S)

Writing Device Drivers

NAME	tran_sync_pkt – SCSI HBA memory synchronization entry point
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void <i>prefix</i>tran_sync_pkt(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i> A pointer to a scsi_address(9S) structure.</p> <p><i>pkt</i> A pointer to a scsi_pkt(9S) structure.</p>
DESCRIPTION	<p>The tran_sync_pkt() vector in the scsi_hba_tran(9S) structure must be initialized during the HBA driver's attach(9E) to point to an HBA driver entry point to be called when a target driver calls scsi_sync_pkt(9F).</p> <p>tran_sync_pkt() must synchronize a HBA's or device's view of the data associated with the <i>pkt</i>, typically by calling ddi_dma_sync(9F). The operation may also involve HBA hardware-specific details, such as flushing I/O caches, or stalling until hardware buffers have been drained.</p>
SEE ALSO	<p>attach(9E), tran_init_pkt(9E), ddi_dma_sync(9F), scsi_hba_attach(9F), scsi_init_pkt(9F), scsi_sync_pkt(9F), scsi_address(9S), scsi_hba_tran(9S), scsi_pkt(9S)</p> <p><i>Writing Device Drivers</i></p>
NOTES	<p>A target driver may call tran_sync_pkt() on packets for which no DMA resources were allocated.</p>

tran_tgt_free(9E)

NAME	tran_tgt_free – request to free HBA resources allocated on behalf of a target								
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void prefixtran_tgt_free(dev_info_t *hba_dip, dev_info_t *tgt_dip, scsi_hba_tran_t *hba_tran, struct scsi_device *sd);</pre>								
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).								
PARAMETERS	<table><tr><td><i>hba_dip</i></td><td>Pointer to a dev_info_t structure, referring to the HBA device instance.</td></tr><tr><td><i>tgt_dip</i></td><td>Pointer to a dev_info_t structure, referring to the target device instance.</td></tr><tr><td><i>hba_tran</i></td><td>Pointer to a scsi_hba_tran(9S) structure, consisting of the HBA's transport vectors.</td></tr><tr><td><i>sd</i></td><td>Pointer to a scsi_device(9S) structure, describing the target.</td></tr></table>	<i>hba_dip</i>	Pointer to a dev_info_t structure, referring to the HBA device instance.	<i>tgt_dip</i>	Pointer to a dev_info_t structure, referring to the target device instance.	<i>hba_tran</i>	Pointer to a scsi_hba_tran(9S) structure, consisting of the HBA's transport vectors.	<i>sd</i>	Pointer to a scsi_device(9S) structure, describing the target.
<i>hba_dip</i>	Pointer to a dev_info_t structure, referring to the HBA device instance.								
<i>tgt_dip</i>	Pointer to a dev_info_t structure, referring to the target device instance.								
<i>hba_tran</i>	Pointer to a scsi_hba_tran(9S) structure, consisting of the HBA's transport vectors.								
<i>sd</i>	Pointer to a scsi_device(9S) structure, describing the target.								
DESCRIPTION	The tran_tgt_free() vector in the scsi_hba_tran(9S) structure may be initialized during the HBA driver's attach(9E) to point to an HBA driver function to be called by the system when an instance of a target device is being detached. The tran_tgt_free() vector, if not NULL, is called after the target device instance has returned successfully from its detach(9E) entry point, but before the dev_info node structure is removed from the system. The HBA driver should release any resources allocated during its tran_tgt_init() or tran_tgt_probe() initialization performed for this target device instance.								
SEE ALSO	attach(9E), detach(9E), tran_tgt_init(9E), tran_tgt_probe(9E), scsi_device(9S), scsi_hba_tran(9S) <i>Writing Device Drivers</i>								

NAME	tran_tgt_init – request to initialize HBA resources on behalf of a particular target
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> void prefixtran_tgt_init(dev_info_t *hba_dip, dev_info_t *tgt_dip, scsi_hba_tran_t *hba_tran, struct scsi_device *sd);</pre>
INTERFACE LEVEL PARAMETERS	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>hba_dip</i> Pointer to a dev_info_t structure, referring to the HBA device instance.</p> <p><i>tgt_dip</i> Pointer to a dev_info_t structure, referring to the target device instance.</p> <p><i>hba_tran</i> Pointer to a scsi_hba_tran(9S) structure, consisting of the HBA's transport vectors.</p> <p><i>sd</i> Pointer to a scsi_device(9S) structure, describing the target.</p>
DESCRIPTION	<p>The tran_tgt_init() vector in the scsi_hba_tran(9S) structure may be initialized during the HBA driver's attach(9E) to point to an HBA driver function to be called by the system when an instance of a target device is being created. The tran_tgt_init() vector, if not NULL, is called after the dev_info node structure is created for this target device instance, but before probe(9E) for this instance is called. Before receiving transport requests from the target driver instance, the HBA may perform any initialization required for this particular target during the call of the tran_tgt_init() vector.</p> <p>Note that <i>hba_tran</i> will point to a cloned copy of the scsi_hba_tran_t structure allocated by the HBA driver if the SCSI_HBA_TRAN_CLONE flag was specified in the call to scsi_hba_attach(9F). In this case, the HBA driver may choose to initialize the tran_tgt_private field in the structure pointed to by <i>hba_tran</i>, to point to the data specific to the particular target device instance.</p>
RETURN VALUES	<p>tran_tgt_init() must return:</p> <p>DDI_SUCCESS the HBA driver can support the addressed target, and was able to initialize per-target resources.</p> <p>DDI_FAILURE the HBA driver cannot support the addressed target, or was unable to initialize per-target resources. In this event, the initialization of this instance of the target device will not be continued, the target driver's probe(9E) will not be called, and the <i>tgt_dip</i> structure destroyed.</p>
SEE ALSO	<p>attach(9E), probe(9E), tran_tgt_free(9E), tran_tgt_probe(9E), scsi_hba_attach_setup(9F), scsi_device(9S), scsi_hba_tran(9S)</p> <p><i>Writing Device Drivers</i></p>

tran_tgt_probe(9E)

NAME	tran_tgt_probe – request to probe SCSI bus for a particular target
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int <i>prefix</i>tran_tgt_probe(struct scsi_device *sd, int (*waitfunc, void)); ;</pre>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
PARAMETERS	<p><i>sd</i> Pointer to a scsi_device(9S) structure.</p> <p><i>waitfunc</i> Pointer to either NULL_FUNC or SLEEP_FUNC.</p>
DESCRIPTION	<p>The tran_tgt_probe() vector in the scsi_hba_tran(9S) structure may be initialized during the HBA driver's attach(9E) to point to a function to be called by scsi_probe(9F) when called by a target driver during probe(9E) and attach(9E) to probe for a particular SCSI target on the bus. In the absence of an HBA-specific tran_tgt_probe() function, the default scsi_probe(9F) behavior is supplied by the function scsi_hba_probe(9F).</p> <p>The possible choices the HBA driver may make are:</p> <ul style="list-style-type: none">■ Initialize the tran_tgt_probe vector to point to scsi_hba_probe(9F), which results in the same behavior.■ Initialize the tran_tgt_probe vector to point to a private function in the HBA, which may call scsi_hba_probe(9F) before or after any necessary processing, as long as all the defined scsi_probe(9F) semantics are preserved. <p><i>waitfunc</i> indicates what tran_tgt_probe() should do when resources are not available:</p> <p>NULL_FUNC Do not wait for resources. See scsi_probe(9F) for defined return values if no resources are available.</p> <p>SLEEP_FUNC Wait indefinitely for resources.</p>
SEE ALSO	attach(9E), probe(9E), tran_tgt_free(9E), tran_tgt_init(9E), scsi_hba_probe(9F), scsi_probe(9F), scsi_device(9S), scsi_hba_tran(9S)

Writing Device Drivers

NAME	tran_quiesce, tran_unquiesce – quiesce and unquiesce a SCSI bus				
SYNOPSIS	<pre>#include <sys/scsi/scsi.h> int prefixtran_quiesce(dev_info_t *hba_dip); int prefixtran_unquiesce(dev_info_t *hba_dip);</pre>				
INTERFACE LEVEL	Solaris DDI				
PARAMETERS	<i>hba_dip</i> The dev_info_t pointer associated with the SCSI HBA.				
DESCRIPTION	<p>The tran_quiesce() and tran_unquiesce() vectors in the scsi_hba_tran(9S) structure should be initialized during the HBA driver's attach(9E). They are HBA entry points to be called when a user initiates quiesce and unquiesce operations through device control interfaces.</p> <p>tran_quiesce() should wait for all outstanding commands to complete and blocks (or queues) any I/O requests issued. tran_unquiesce() should allow I/O activities to resume on the SCSI bus.</p> <p>Implementation is hardware specific.</p>				
RETURN VALUES	<p>tran_quiesce() and tran_unquiesce() should return:</p> <p>1 on success.</p> <p>0 on failure.</p>				
ATTRIBUTES	See attributes(5) for a description of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Stability Level</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Stability Level	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Stability Level	Evolving				
SEE ALSO	attributes(5), tran_bus_reset(9E), scsi_hba_tran(9S)				

write(9E)

NAME	write – write data to a device
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/errno.h> #include <sys/open.h> #include <sys/cred.h> #include <sys/ddi.h> #include <sys/sunddi.h> int <i>prefix</i>write(dev_t dev, struct uio *uio_p, cred_t *cred_p);</pre>
INTERFACE LEVEL PARAMETERS	Architecture independent level 1 (DDI/DKI). This entry point is optional. <i>dev</i> Device number. <i>uio_p</i> Pointer to the uio(9S) structure that describes where the data is to be stored in user space. <i>cred_p</i> Pointer to the user credential structure for the I/O transaction.
DESCRIPTION	Used for character or raw data I/O, the driver write() routine is called indirectly through cb_ops(9S) by the write(2) system call. The write() routine supervises the data transfer from user space to a device described by the uio(9S) structure. The write() routine should check the validity of the minor number component of <i>dev</i> and the user credentials pointed to by <i>cred_p</i> , if pertinent.
RETURN VALUES	The write() routine should return 0 for success, or the appropriate error number.
EXAMPLES	EXAMPLE 1 The following is an example of a write() routine using physio(9F) to perform writes to a seekable device: <pre>static int xxwrite(dev_t dev, struct uio *uiop, cred_t *credp) { int instance; xx_t xx; instance = getminor(dev); xx = ddi_get_soft_state(xxstate, instance); if (xx == NULL) return (ENXIO); return (physio(xxstrategy, NULL, dev, B_WRITE, xxmin, uiop)); }</pre>
SEE ALSO	read(2), write(2), read(9E), physio(9F), cb_ops(9S), uio(9S) <i>Writing Device Drivers</i>

Index

A

aread — asynchronous read from a device, 22
asynchronous read — aread, 22
asynchronous write — awrite, 27
awrite — asynchronous write to a device, 27

C

character-oriented drivers, — ioctl, 128
csx_event_handler — PC Card driver event handler, 34

D

DDI device mapping
 devmap_access — device mapping access entry point, 47
 devmap_contextmgt — device mapping access entry point, 50
 devmap_dup — device mapping duplication entry point, 53
 devmap_map — device mapping access entry point, 55
 devmap_unmap — device mapping unmap entry point, 57
 mapdev_access — device mapping access entry point, 136
 mapdev_dup — device mapping duplication entry point, 138
 mapdev_free — device mapping free entry point, 140

dev_info structure, convert device number to —
 getinfo, 64
device access
 — close, 32
 — open, 146
device mapping access entry point —
 devmap_access, 47
device mapping access entry point —
 devmap_contextmgt, 50
device mapping access entry point —
 devmap_map, 55
device mapping access entry point —
 mapdev_access, 136
device mapping duplication entry point —
 devmap_dup, 53
device mapping duplication entry point —
 mapdev_dup, 138
device mapping free entry point —
 mapdev_free, 140
device mapping unmap entry point —
 devmap_unmap, 57
device number, convert to dev_info structure —
 getinfo, 64
devices
 attach to system — attach, 24
 claim to drive a device — identify, 121
 detach from system — detach, 40
 read data — read, 156
 write data to a device — write, 188
devices, memory mapped
 check virtual mapping — devmap, 43
 check virtual mapping — mmap, 141

devices, memory mapping, map device memory into user space — `segmap`, 159
 devices, non-self-identifying, determine if present — `probe`, 151
`devmap_access` — device mapping access entry point, 47
`devmap_contextmgt` — device mapping access entry point, 50
`devmap_dup` — device mapping duplication entry point, 53
`devmap_map` — device mapping access entry point, 55
`devmap_unmap` — device mapping unmap entry point, 57
 Driver entry point routines
 — `attach`, 24
 — `chpoll`, 30
 — `close`, 32
 — `detach`, 40
 — `devmap`, 43
 — `dump`, 60
 — `_fini`, 61
 — `getinfo`, 64
 — `identify`, 121
 — `_info`, 61
 — `_init`, 61
 — `ioctl`, 128
 — `mmap`, 141
 — `open`, 146
 — `print`, 150
 — `probe`, 151
 — `prop_op`, 152
 — `put`, 154
 — `read`, 156
 — `segmap`, 159
 — `srv`, 160
 — `strategy`, 162
 — `write`, 188
 driver messages, display on system console — `print`, 150
 driver property information, report
 — `prop_op`, 152
 drivers, character-oriented, — `ioctl`, 128
`dump` — dump memory to disk during system failure, 60
 dynamically update kstats — `ks_update`, 134

G

`get/set` SCSI transport capability —
 `tran_getcap`, 169
 `tran_setcap`, 169
`gldm_get_stats` — report hardware statistics, 69
`gldm_intr` — called when the device might have interrupted, 69
`gldm_ioctl` — implements device-specific ioctl commands, 69
`gldm_reset` — initialize the network device, 67
`gldm_send` — queue a packet to the device for transmission, 68
`gldm_set_mac_addr` — set the physical address of the network device, 67
`gldm_set_multicast` — enables/disables device-level multicast addresses, 67
`gldm_set_promiscuous` — enables/disables device-level promiscuous mode, 68
`gldm_start` — start the network device receiver, 67
`gldm_stop` — stop the network device receiver, 67

H

HBA resources
 request to free HBA resources allocated on behalf of a target — `tran_tgt_free`, 184
 request to initialize HBA resources on behalf of a particular target — `tran_tgt_init`, 185

I

`identify` — claim to drive a device, 121

K

kernel modules, dynamic loading
 initialize a loadable module — `_init`, 61
 prepare loadable module for unloading —
 `_fini`, 61
 return loadable module information —
 `_info`, 61
`ks_update` — dynamically update kstats, 134

M

mapdev_access — device mapping access entry point, 136
mapdev_dup — device mapping duplication entry point, 138
mapdev_free — device mapping free entry point, 140
memory mapping for devices
 check virtual mapping — devmap, 43
 check virtual mapping — mmap, 141
 map device memory into user space — segmap, 159

N

non-self-identifying devices, determine if present — probe, 151
non-STREAMS character device driver, poll entry point — chpoll, 30

P

PC Card driver event handler, — csx_event_handler, 34
power — power a device attached to the system, 148
power a device attached to the system — power, 148
put — receive messages from the preceding queue, 154

Q

quiesce and unquiesce a SCSI bus
 — tran_quiesce, 174
 — tran_unquiesce, 174

R

request to notify SCSI target of bus reset, — tran_reset_notify, 177
reset a SCSI bus — tran_bus_reset, 164
reset a SCSI bus or target — tran_reset, 175

S

SCSI HBA packet preparation and deallocation
 — tran_init_pkt, 171
 tran_destroy_pkt, 171
SCSI bus, request to probe SCSI bus for a particular target — tran_tgt_probe, 186
SCSI command
 abort — tran_abort, 163
 request to transport — tran_start, 180
SCSI HBA DMA deallocation entry point — tran_dmafree, 168
SCSI HBA memory synchronization entry point — tran_sync_pkt, 183
strategy — perform block I/O, 162
STREAMS message queues
 receive messages from the preceding queue — put, 154
 service queued messages — srv, 160

T

tran_abort — abort a SCSI command, 163
tran_bus_reset — reset a SCSI bus, 164
tran_destroy_pkt — SCSI HBA packet preparation and deallocation, 171
tran_dmafree — SCSI HBA DMA deallocation entry point, 168
tran_getcap — get/set SCSI transport capability, 169
tran_init_pkt — SCSI HBA packet preparation and deallocation, 171
tran_quiesce — quiesce and unquiesce a SCSI bus, 174
tran_reset — reset a SCSI bus or target, 175
tran_reset_notify — request to notify SCSI target of bus reset, 177
tran_setcap — get/set SCSI transport capability, 169
tran_start — request to transport a SCSI command, 180
tran_sync_pkt — SCSI HBA memory synchronization entry point, 183
tran_tgt_free — request to free HBA resources allocated on behalf of a target, 184
tran_tgt_init — request to initialize HBA resources on behalf of a particular target, 185

tran_tgt_probe — request to probe SCSI bus for a particular target, 186
tran_unquiesce — quiesce and unquiesce a SCSI bus, 174

V

virtual address space, dump portion of to disk in case of system failure — dump, 60

W

write — write data to a device, 188