

# SunOS Reference Manual

Sun Microsystems, Inc.  
2550 Garcia Avenue  
Mountain View, CA 94043  
U.S.A.



© 1994 Sun Microsystems, Inc. 2550 Garcia Avenue, Mountain View, California 94043-1100 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.

Portions of this product may be derived from the UNIX® system, licensed from UNIX System Laboratories, Inc., a wholly owned subsidiary of Novell, Inc., and from the Berkeley 4.3 BSD system, licensed from the University of California. Third-party software, including font technology in this product, is protected by copyright and licensed from Sun's suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

#### TRADEMARKS

Sun, Sun Microsystems, the Sun logo, SunSoft, the SunSoft logo, Solaris, SunOS, OpenWindows, DeskSet, ONC, ONC+, and NFS are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd. OPEN LOOK is a registered trademark of Novell, Inc. PostScript and Display PostScript are trademarks of Adobe Systems, Inc.

All SPARC trademarks are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. SPARCcenter, SPARCcluster, SPARCcompiler, SPARCdesign, SPARC811, SPARCengine, SPARCprinter, SPARCserver, SPARCstation, SPARCstorage, SPARCworks, microSPARC, microSPARC-II, and UltraSPARC are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK® and Sun™ Graphical User Interfaces were 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.

X Window System is a trademark of the X Consortium.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN. THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAMS(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

Portions © AT&T 1983-1990 and reproduced with permission from AT&T.

# *Preface*

---

## *OVERVIEW*

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

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

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2 of this volume.

- 
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
  - Section 5 contains miscellaneous documentation such as character set tables, etc.
  - Section 6 contains available games and demos.
  - Section 7 describes various special files that refer to specific hardware peripherals, and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.
  - Section 9 provides reference information needed to write device drivers in the kernel operating systems environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver–Kernel Interface (DKI).
  - Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer may include in a device driver.
  - Section 9F describes the kernel functions available for use by device drivers.
  - Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

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

## *NAME*

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

## *SYNOPSIS*

This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full pathname is shown. Literal characters (commands and options) are in **bold** font and variables (arguments, parameters and substitution characters) are in *italic* font. Options and

---

arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.

The following special characters are used in this section:

- [ ] The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument *must* be specified.
- ... Ellipses. Several values may be provided for the previous argument, or the previous argument can be specified multiple times, for example, '*filename ...*'.
- | Separator. Only one of the arguments separated by this character can be specified at time.
- { } Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.

## *PROTOCOL*

This section occurs only in subsection 3R to indicate the protocol description file. The protocol specification pathname is always listed in **bold** font.

## *AVAILABILITY*

This section briefly states any limitations on the availability of the command. These limitations could be hardware or software specific.

A specification of a class of hardware platform, such as **x86** or **SPARC**, denotes that the command or interface is applicable for the hardware platform specified.

In Section 1 and Section 1M, **AVAILABILITY** indicates which package contains the command being described on the manual page. In order to use the command, the specified package must have been installed with the operating system. If the package was not installed, see **pkgadd(1)** for information on how to upgrade.

## *MT-LEVEL*

This section lists the **MT-LEVEL** of the library functions described in the Section 3 manual pages. The **MT-LEVEL** defines the libraries' ability to support threads. See **Intro(3)** for more information.

---

## *DESCRIPTION*

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

## *IOCTL*

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

## *OPTIONS*

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

## *OPERANDS*

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

## *OUTPUT*

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

## *RETURN VALUES*

If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared as **void** do not return values, so they are not discussed in *RETURN VALUES*.

---

## *ERRORS*

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

## *USAGE*

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

- Commands**
- Modifiers**
- Variables**
- Expressions**
- Input Grammar**

## *EXAMPLES*

This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command line entry and machine response is shown. Whenever an example is given, the prompt is shown as

**example%**

or if the user must be super-user,

**example#**

Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS and USAGE sections.

## *ENVIRONMENT*

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 greater than zero for various error conditions.

## *FILES*

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

## *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. Messages appear in **bold** font with the exception of variables, which are in *italic* font.

## *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.



**NAME** Intro, intro – introduction to device driver entry points

**DESCRIPTION** Section 9E describes the entry-point routines a developer may 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. 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.

<i>Routine</i>	<i>Type</i>
<b>put</b>	DDI/DKI
<b>srv</b>	DDI/DKI

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

<i>Routine</i>	<i>Type</i>
<b>_fini</b>	Solaris DDI
<b>_info</b>	Solaris DDI
<b>_init</b>	Solaris DDI
<b>aread</b>	Solaris DDI
<b>attach</b>	Solaris DDI
<b>awrite</b>	Solaris DDI
<b>chpoll</b>	DDI/DKI
<b>close</b>	DDI/DKI
<b>detach</b>	Solaris DDI
<b>dump</b>	Solaris DDI

<b>getinfo</b>	Solaris DDI
<b>identify</b>	Solaris DDI
<b>ioctl</b>	DDI/DKI
<b>ks_update</b>	Solaris DDI
<b>mapdev_access</b>	Solaris DDI
<b>mapdev_dup</b>	Solaris DDI
<b>mapdev_free</b>	Solaris DDI
<b>mmap</b>	DKI only
<b>open</b>	DDI/DKI
<b>print</b>	DDI/DKI
<b>probe</b>	Solaris DDI
<b>prop_op</b>	Solaris DDI
<b>read</b>	DDI/DKI
<b>segmap</b>	DKI only
<b>strategy</b>	DDI/DKI
<b>tran_abort</b>	Solaris DDI
<b>tran_destroy_pkt</b>	Solaris DDI
<b>tran_dmafree</b>	Solaris DDI
<b>tran_getcap</b>	Solaris DDI
<b>tran_init_pkt</b>	Solaris DDI
<b>tran_reset</b>	Solaris DDI
<b>tran_reset_notify</b>	Solaris DDI
<b>tran_setcap</b>	Solaris DDI
<b>tran_start</b>	Solaris DDI
<b>tran_sync_pkt</b>	Solaris DDI
<b>tran_tgt_free</b>	Solaris DDI
<b>tran_tgt_init</b>	Solaris DDI
<b>tran_tgt_probe</b>	Solaris DDI
<b>write</b>	DDI/DKI

The table below lists the error codes that should be returned by a driver routine when an error is encountered. It lists the error values in alphabetic order. All the error values 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 returning the value. 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, as the function that it calls, **cmn\_err**(9F), is declared as **void** (no error is returned).

Error Value	Error Description	Use in these Driver Routines (9E)
<b>EAGAIN</b>	Kernel resources, such as the <b>buf</b> structure or cache memory, are not available at this time (device may be busy, or the system resource is not available).	<b>open, ioctl, read, write, strategy</b>
<b>EFAULT</b>	An invalid address has been passed as an argument; memory addressing error.	<b>open, close, ioctl, read, write, strategy</b>
<b>EINTR</b>	Sleep interrupted by signal.	<b>open, close, ioctl, read, write, strategy</b>
<b>EINVAL</b>	An invalid argument was passed to the routine.	<b>open, ioctl, read, write, strategy</b>
<b>EIO</b>	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).	<b>open, close, ioctl, read, write, strategy</b>
<b>ENXIO</b>	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.	<b>open, close, ioctl, read, write, strategy</b>
<b>EPERM</b>	A process attempting an operation did not have required permission.	<b>open, ioctl, read, write, close</b>
<b>EROFS</b>	An attempt was made to open for writing a read-only device.	<b>open</b>

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

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

<b>Name</b>	<b>Description</b>
<b>aread(9E)</b>	asynchronous read from a device
<b>attach(9E)</b>	attach a device to the system
<b>awrite(9E)</b>	asynchronous write to a device
<b>chpoll(9E)</b>	poll entry point for a non-STREAMS character driver
<b>close(9E)</b>	relinquish access to a device
<b>detach(9E)</b>	detach a device
<b>dump(9E)</b>	dump memory to device during system failure
<b>_fini(9E)</b>	loadable module configuration entry points
<b>getinfo(9E)</b>	get device driver information
<b>identify(9E)</b>	claim to drive a device
<b>_info(9E)</b>	See <b>_fini(9E)</b>
<b>_init(9E)</b>	See <b>_fini(9E)</b>
<b>ioctl(9E)</b>	control a character device
<b>ks_update(9E)</b>	dynamically update kstats
<b>mapdev_access(9E)</b>	device mapping access entry point
<b>mapdev_dup(9E)</b>	device mapping duplication entry point
<b>mapdev_free(9E)</b>	device mapping free entry point
<b>mmap(9E)</b>	check virtual mapping for memory mapped device
<b>open(9E)</b>	gain access to a device
<b>print(9E)</b>	display a driver message on system console
<b>probe(9E)</b>	determine if a non-self-identifying device is present
<b>prop_op(9E)</b>	report driver property information
<b>put(9E)</b>	receive messages from the preceding queue
<b>read(9E)</b>	read data from a device
<b>segmap(9E)</b>	map device memory into user space
<b>srv(9E)</b>	service queued messages
<b>strategy(9E)</b>	perform block I/O
<b>tran_abort(9E)</b>	abort a SCSI command
<b>tran_destroy_pkt(9E)</b>	See <b>tran_init_pkt(9E)</b>
<b>tran_dmafree(9E)</b>	SCSI HBA DMA deallocation entry point
<b>tran_getcap(9E)</b>	get/set SCSI transport capability
<b>tran_init_pkt(9E)</b>	SCSI HBA packet preparation and deallocation
<b>tran_reset(9E)</b>	reset a SCSI bus or target

<b>tran_reset_notify(9E)</b>	request to notify SCSI target of bus reset
<b>tran_setcap(9E)</b>	See <b>tran_getcap(9E)</b>
<b>tran_start(9E)</b>	request to transport a SCSI command
<b>tran_sync_pkt(9E)</b>	SCSI HBA memory synchronization entry point
<b>tran_tgt_free(9E)</b>	request to free HBA resources allocated on behalf of a target
<b>tran_tgt_init(9E)</b>	request to initialize HBA resources on behalf of a particular target
<b>tran_tgt_probe(9E)</b>	request to probe SCSI bus for a particular target
<b>write(9E)</b>	write data to a device

<b>NAME</b>	_fini, _info, _init – loadable module configuration entry points
<b>SYNOPSIS</b>	<pre>#include &lt;sys/modctl.h&gt; int _fini(void); int _info(struct modinfo *modinfo); int _init(void);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI). These entry points are required. You <b>must</b> write them.
<b>ARGUMENTS</b>	
_info()	<i>modinfo</i> A pointer to an opaque <b>modinfo</b> structure.
<b>DESCRIPTION</b>	<p><b>_init()</b> initializes a loadable module. It is called before any other routine in a loadable module. <b>_init()</b> returns the value returned by <b>mod_install(9F)</b>. The module may optionally perform some other work before the <b>mod_install(9F)</b> call is performed. If the module has done some setup before the <b>mod_install(9F)</b> function is called, then it should be prepared to undo that setup if <b>mod_install(9F)</b> returns an error.</p> <p><b>_info()</b> returns information about a loadable module. <b>_info()</b> returns the value returned by <b>mod_info(9F)</b>.</p> <p><b>_fini()</b> 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 <b>_fini()</b> returns the value returned by <b>mod_remove(9F)</b>. Upon successful return from <b>_fini()</b> no other routine in the module will be called before <b>_init()</b> is called.</p>
<b>RETURN VALUES</b>	<p><b>_init()</b> should return the appropriate error number if there is an error, else it should return the return value from <b>mod_install(9F)</b>.</p> <p><b>_info()</b> should return the return value from <b>mod_info(9F)</b></p> <p><b>_fini()</b> should return the return value from <b>mod_remove(9F)</b>.</p>
<b>EXAMPLES</b>	<p>The following example demonstrates how to initialize and free a <b>mutex(9F)</b>.</p> <pre>#include &lt;sys/modctl.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  static struct dev_ops drv_ops; /*  * Module linkage information for the kernel.  */ static struct modldrv modldrv = {     &amp;mod_driverops, /* Type of module. This one is a driver */     "Sample Driver",     &amp;drv_ops /* driver ops */</pre>

```
};

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, "XXX Global Mutex",
        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);  
}  
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()**.



<b>NAME</b>	aread – asynchronous read from a device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/uio.h&gt; #include &lt;sys/aio_req.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixaread(dev_t dev, struct aio_req *aio_reqp, cred_t *cred_p);</pre>
<b>ARGUMENTS</b>	<p><i>dev</i>            Device number.</p> <p><i>aio_reqp</i>       Pointer to the <b>aio_req</b>(9S) structure that describes where the data is to be stored.</p> <p><i>cred_p</i>          Pointer to the credential structure.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI). This entry point is <b>Optional</b> . Drivers that do not support an <b>aread</b> () entry point should use <b>nodev</b> (9F).
<b>DESCRIPTION</b>	The driver's <b>aread</b> () routine is called to perform an asynchronous read. <b>getminor</b> (9F) can be used to access the minor number component of the <i>dev</i> argument. <b>aread</b> () may use the credential structure pointed to by <i>cred_p</i> to check for superuser access by calling <b>drv_priv</b> (9F). The <b>aread</b> () routine may also examine the <b>uio</b> (9S) structure through the <b>aio_req</b> structure pointer, <i>aio_reqp</i> . However, the <b>aread</b> () routine must not modify <i>aio_reqp</i> or the contents of the <b>aio_req</b> (9S) structure. <b>aread</b> () must call <b>aphysio</b> (9F) with the <b>aio_req</b> pointer, and a pointer to the driver's <b>strategy</b> (9E) routine.
<b>RETURN VALUES</b>	The <b>aread</b> () routine should return <b>0</b> for success, or the appropriate error number.
<b>CONTEXT</b>	This function is called from user context only.
<b>EXAMPLES</b>	The following is an example of an <b>aread</b> () routine.

```

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));
}

```

**SEE ALSO** read(2), aioread(3), awrite(9E), read(9E), strategy(9E), write(9E), anocancel(9F), aphysio(9F), ddi\_get\_soft\_state(9F), drv\_priv(9F), getminor(9F), minphys(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.

<b>NAME</b>	attach – attach a device to the system
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixattach(dev_info_t *dip, ddi_attach_cmd_t cmd);</pre>
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Solaris DDI specific (Solaris DDI). This entry point is <i>required</i> and must be written.</p> <p><i>dip</i> A pointer to the device's <b>dev_info</b> structure.</p> <p><i>cmd</i> Attach type. Should be set to <b>DDI_ATTACH</b>. Other values are reserved. The driver should return <b>DDI_FAILURE</b> if reserved values are passed to it.</p>
<b>DESCRIPTION</b>	<p><b>attach()</b> is the device-specific initialization entry point. When <b>attach()</b> is called with <i>cmd</i> set to <b>DDI_ATTACH</b>, all normal kernel services (such as <b>kmem_alloc(9F)</b>) are available for use by the driver. Device interrupts are not blocked when attaching a device to the system.</p> <p><b>attach()</b> will be called once for each instance of the device on the system. Until <b>attach()</b> succeeds, the only driver entry points which may be called are <b>open(9E)</b> and <b>getinfo(9E)</b>. See the "Autoconfiguration" chapter in <i>Writing Device Drivers</i>. The instance number may be obtained using <b>ddi_get_instance(9F)</b>.</p> <p>Successful returns from <b>identify(9E)</b> and <b>probe(9E)</b> are required before a call to the driver's <b>attach()</b> entry point will be made.</p>
<b>RETURN VALUES</b>	<p><b>attach()</b> should return:</p> <p><b>DDI_SUCCESS</b> on success.</p> <p><b>DDI_FAILURE</b> on failure.</p>
<b>SEE ALSO</b>	<p><b>identify(9E)</b>, <b>probe(9E)</b>, <b>ddi_add_intr(9F)</b>, <b>ddi_create_minor_node(9F)</b>, <b>ddi_get_instance(9F)</b>, <b>ddi_map_regs(9F)</b>, <b>kmem_alloc(9F)</b>,</p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	awrite – asynchronous write to a device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/uio.h&gt; #include &lt;sys/aio_req.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixawrite(dev_t dev, struct aio_req *aio_reqp, cred_t *cred_p);</pre>
<b>ARGUMENTS</b>	<p><i>dev</i>            Device number.</p> <p><i>aio_reqp</i>       Pointer to the <b>aio_req</b>(9S) structure that describes where the data is stored.</p> <p><i>cred_p</i>          Pointer to the credential structure.</p>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI). This entry point is <b>Optional</b> . Drivers that do not support an <b>awrite</b> () entry point should use <b>nodev</b> (9F).
<b>DESCRIPTION</b>	The driver's <b>awrite</b> () routine is called to perform an asynchronous write. <b>getminor</b> (9F) can be used to access the minor number component of the <i>dev</i> argument. <b>awrite</b> () may use the credential structure pointed to by <i>cred_p</i> to check for superuser access by calling <b>drv_priv</b> (9F). The <b>awrite</b> () routine may also examine the <b>uio</b> (9S) structure through the <b>aio_req</b> structure pointer, <i>aio_reqp</i> . However, the <b>awrite</b> () routine must not modify <i>aio_reqp</i> or the contents of the <b>aio_req</b> (9S) structure. <b>awrite</b> () must call <b>aphysio</b> (9F) with the <b>aio_req</b> pointer, and a pointer to the driver's <b>strategy</b> (9E) routine.
<b>RETURN VALUES</b>	The <b>awrite</b> () routine should return <b>0</b> for success, or the appropriate error number.
<b>CONTEXT</b>	This function is called from user context only.
<b>EXAMPLES</b>	The following is an example of an <b>awrite</b> () routine.

```

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));
}

```

**SEE ALSO** write(2), aiowrite(3), aread(9E), read(9E), strategy(9E), write(9E), anocancel(9F), aphysio(9F), ddi\_get\_soft\_state(9F), drv\_priv(9F), getminor(9F), minphys(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.

<b>NAME</b>	chpoll – poll entry point for a non-STREAMS character driver																		
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/poll.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixchpoll(dev_t dev, short events, int anyyet, short *reventsp,                  struct pollhead **phpp);</pre>																		
<b>INTERFACE LEVEL</b>	This entry point is <i>optional</i> . Architecture independent level 1 (DDI/DKI).																		
<b>ARGUMENTS</b>	<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> <table border="0"> <tr> <td style="padding-left: 2em;"><b>POLLIN</b></td> <td>Data other than high priority data may be read without blocking.</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLOUT</b></td> <td>Normal data may be written without blocking.</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLPRI</b></td> <td>High priority data may be received without blocking.</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLHUP</b></td> <td>A device hangup has occurred.</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLERR</b></td> <td>An error has occurred on the device.</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLRDNORM</b></td> <td>Normal data (priority band = 0) may be read without blocking.</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLRDBAND</b></td> <td>Data from a non-zero priority band may be read without blocking</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLWRNORM</b></td> <td>The same as <b>POLLOUT</b>.</td> </tr> <tr> <td style="padding-left: 2em;"><b>POLLWRBAND</b></td> <td>Priority data (priority band &gt; 0) may be written.</td> </tr> </table> <p><i>anyyet</i>     A flag that is non-zero if any other file descriptors in the <b>pollfd</b> array have events pending. The <b>poll(2)</b> system call takes a pointer to an array of <b>pollfd</b> structures as one of its arguments. See the <b>poll(2)</b> 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 <b>pollhead</b> structure.</p>	<b>POLLIN</b>	Data other than high priority data may be read without blocking.	<b>POLLOUT</b>	Normal data may be written without blocking.	<b>POLLPRI</b>	High priority data may be received without blocking.	<b>POLLHUP</b>	A device hangup has occurred.	<b>POLLERR</b>	An error has occurred on the device.	<b>POLLRDNORM</b>	Normal data (priority band = 0) may be read without blocking.	<b>POLLRDBAND</b>	Data from a non-zero priority band may be read without blocking	<b>POLLWRNORM</b>	The same as <b>POLLOUT</b> .	<b>POLLWRBAND</b>	Priority data (priority band > 0) may be written.
<b>POLLIN</b>	Data other than high priority data may be read without blocking.																		
<b>POLLOUT</b>	Normal data may be written without blocking.																		
<b>POLLPRI</b>	High priority data may be received without blocking.																		
<b>POLLHUP</b>	A device hangup has occurred.																		
<b>POLLERR</b>	An error has occurred on the device.																		
<b>POLLRDNORM</b>	Normal data (priority band = 0) may be read without blocking.																		
<b>POLLRDBAND</b>	Data from a non-zero priority band may be read without blocking																		
<b>POLLWRNORM</b>	The same as <b>POLLOUT</b> .																		
<b>POLLWRBAND</b>	Priority data (priority band > 0) may be written.																		
<b>DESCRIPTION</b>	The <b>chpoll()</b> 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 should only be called with one event at a time.

**RETURN VALUES**

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

**SEE ALSO**

**poll(2)**, **pollwakeup(9F)**

*Writing Device Drivers*

**NOTES**

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

<b>NAME</b>	close – relinquish access to a device
<b>SYNOPSIS</b> Block and Character	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/file.h&gt; #include &lt;sys/errno.h&gt; #include &lt;sys/open.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixclose(dev_t dev, int flag, int otyp, cred_t *cred_p);</pre>
<b>STREAMS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/stream.h&gt; #include &lt;sys/file.h&gt; #include &lt;sys/errno.h&gt; #include &lt;sys/open.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixclose(queue_t *q, int flag, cred_t *cred_p);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI). This entry point is <b>required</b> for block devices.
<b>ARGUMENTS</b> Block and Character	<p><i>dev</i> Device number.</p> <p><i>flag</i> File status flag, as set by the <b>open(2)</b> or modified by the <b>fcntl(2)</b> system calls. The flag is for information only—the file should always be closed completely. Possible values are: <b>FEXCL</b>, <b>FNDELAY</b>, <b>FREAD</b>, <b>FKLYR</b>, and <b>FWRITE</b>. Refer to <b>open(9E)</b> 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 <b>open()</b> routine may be called many times, but the <b>close()</b> routine should only be called on the last <b>close</b> of a device.</p> <p><b>OTYP_BLK</b> close was through block interface for the device</p> <p><b>OTYP_CHR</b> close was through the raw/character interface for the device</p> <p><b>OTYP_LYR</b> close a layered process (a higher-level driver called the <b>close()</b> routine of the device)</p> <p><i>*cred_p</i> Pointer to the user credential structure.</p>



**STREAMS**

*\*q* Pointer to **queue**(9S) 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.)

*flag* File status flag.

*\*cred\_p* Pointer to the user credential structure.

**DESCRIPTION**

For STREAMS drivers, the **close()** routine is called by the kernel through the **cb\_ops**(9S) table entry for the device. (Modules use the **fmodsw** table.) A non-null value in the **d\_str** field of the **cb\_ops** entry points to a **streamtab** structure, which points to a **qinit**(9S) containing a pointer to the **close()** routine. Non-STREAMS **close()** routines are called directly from the **cb\_ops** table.

**close()** 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.

A device may be opened simultaneously by multiple processes and the **open()** driver routine is called for each open, but the kernel will only call the **close()** routine when the last process using the device issues a **close(2)** or **umount(2)** system call or exits. (An exception is a close occurring with the *otyp* argument set to **OTYP\_LYR**, for which a close (also having *otyp* = **OTYP\_LYR**) occurs for each open.)

In general, a **close()** routine should always check the validity of the minor number component of the *dev* parameter. The routine should also check permissions as necessary, by using the user credential structure (if pertinent), and the appropriateness of the *flag* and *otyp* parameter values.

**close()** could perform any of the following general functions:

- 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 **open()** routine)
- flush buffers
- notify a device of the close
- deallocate any resources allocated on open

The **close()** 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:

1. The write queue is given a chance to drain.
2. The **close()** routine is called.
3. The module or driver is removed from the stream.

**RETURN VALUES**

**close()** 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 **close()** routine occurs because a problem occurred in the associated device.

**SEE ALSO**

**close(2)**, **detach(9E)**, **open(9E)**

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	detach – detach a device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt; int prefixdetach(dev_info_t *dip, ddi_detach_cmd_t cmd);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI). This entry point is <b>required</b> . It can be <b>nodev</b> .
<b>ARGUMENTS</b>	<p><i>dip</i>            A pointer to the device's <b>dev_info</b> structure.</p> <p><i>cmd</i>            Type of detach; the driver should return <b>DDI_FAILURE</b> if any value other than <b>DDI_DETACH</b> is passed to it.</p>
<b>DESCRIPTION</b>	<p><b>detach()</b> is the complement of the <b>attach(9E)</b> routine. It is used to remove all the states associated with a given instance of a device node prior to the removal of that instance from the system. The <b>dev_info</b> nodes that belong to a driver are removed as part of the process of unloading a device driver from the system.</p> <p>Depending on what was created in the drivers' <b>attach(9E)</b> routine, this might mean calling <b>ddi_unmap_regs()</b> (see <b>ddi_map_regs(9F)</b>) to remove mappings, calling <b>ddi_remove_intr()</b> (see <b>ddi_add_intr(9F)</b>) to unregister interrupt handlers, calling <b>kmem_free(9F)</b> to free any heap allocations, and so forth. This should also include putting the underlying device into a quiescent state so that it will not generate interrupts.</p> <p>If <b>detach()</b> determines that the particular instance of the device cannot be removed when requested, for example, because of some exceptional condition, <b>detach()</b> returns <b>DDI_FAILURE</b>, which prevents the particular device instance from being removed. This will also prevent the driver from being unloaded.</p> <p>Drivers that set up <b>timeout(9F)</b> routines should ensure that they are cancelled before returning <b>DDI_SUCCESS</b> from <b>detach()</b>.</p> <p>The system guarantees that the function will only be called for a particular <b>dev_info</b> node after (and not concurrently with) a successful <b>attach(9E)</b> of that device. The system also guarantees that <b>detach()</b> will only be called when there are no outstanding <b>open(9E)</b> calls on the device.</p>
<b>RETURN VALUES</b>	<p><b>DDI_SUCCESS</b>    The state associated with the given device was successfully removed.</p> <p><b>DDI_FAILURE</b>    The operation failed or the request was not understood. The associated state is unchanged.</p>
<b>CONTEXT</b>	This function is called from user context only.
<b>SEE ALSO</b>	<p><b>attach(9E)</b>, <b>open(9E)</b>, <b>ddi_add_intr(9F)</b>, <b>ddi_map_regs(9F)</b>, <b>kmem_free(9F)</b>, <b>timeout(9F)</b></p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	dump – dump memory to device during system failure
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixdump(dev_t dev, caddr_t addr, daddr_t blkno, int nblk);</pre>
<b>INTERFACE LEVEL</b>	Solaris specific (Solaris DDI). This entry point is <b>required</b> . For drivers that do not implement dump routines, <b>nodev</b> should be used.
<b>ARGUMENTS</b>	<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 to.</p> <p><i>nblk</i> Number of blocks to dump.</p>
<b>DESCRIPTION</b>	<p><b>dump()</b> is used to dump a portion of virtual address space directly to a device in the case of system failure. 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 <b>dump()</b> returns the status of the transfer.</p> <p><b>dump()</b> is called at interrupt priority.</p>
<b>RETURN VALUES</b>	<b>dump()</b> should return 0 on success, or the appropriate error number.
<b>SEE ALSO</b>	<i>Writing Device Drivers</i>

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

```
switch (infocmd) {
case DDI_INFO_DEVT2DEVINFO:
    if ((sp = ddi_get_soft_state(statep,
        getminor((dev_t) arg))) != NULL) {
        *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*

<b>NAME</b>	identify – claim to drive a device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt; int prefixidentify(dev_info_t *dip);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI). This entry point is <b>required</b> . You <b>must</b> write it.
<b>ARGUMENTS</b>	<i>dip</i> A pointer to a <b>dev_info</b> structure.
<b>DESCRIPTION</b>	<b>identify()</b> determines whether this driver drives the device pointed to by <i>dip</i> .
<b>RETURN VALUES</b>	<p><b>identify()</b> should return:</p> <p><b>DDI_IDENTIFIED</b> if it <b>claims</b> to drive this device.</p> <p><b>DDI_NOT_IDENTIFIED</b> if it does <b>not</b> claim to drive this device.</p>
<b>EXAMPLES</b>	<pre>#define XX_NAME "xx" static int xxidentify(dev_info_t *dip) {     if (strcmp(ddi_get_name(dip), XX_NAME) == 0) {         /*name matches device name*/         return(DDI_IDENTIFIED);     } else         return(DDI_NOT_IDENTIFIED); }</pre>
<b>SEE ALSO</b>	<p><b>attach(9E)</b>, <b>ddi_get_name(9F)</b>, <b>strcmp(9F)</b></p> <p><i>Writing Device Drivers</i></p>
<b>WARNINGS</b>	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.

<b>NAME</b>	ioctl – control a character device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/cred.h&gt; #include &lt;sys/file.h&gt; #include &lt;sys/types.h&gt; #include &lt;sys/errno.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixioctl(dev_t dev, int cmd, int arg, int mode, cred_t *cred_p, int *rval_p);</pre>
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Architecture independent level 1 (DDI/DKI). This entry point is <b>optional</b>.</p> <p><i>dev</i> Device number.</p> <p><i>cmd</i> Command argument the driver <b>ioctl</b> routine interprets as the operation to be performed.</p> <p><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 an integer 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.</p> <p><i>mode</i> Contains values set when the device was opened. Use of this mode is optional. However, 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 <b>FREAD</b> or <b>FWRITE</b> flags. See the <i>flag</i> argument description of the <b>open()</b> routine for further values for the <b>ioctl</b> routine's <i>mode</i> argument. In some circumstances, <i>mode</i> is used to provide address space information about the <i>arg</i> argument. See below.</p> <p><i>cred_p</i> Pointer to the user credential structure.</p> <p><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 <b>ioctl()</b> succeeds.</p>
<b>DESCRIPTION</b>	<p><b>ioctl()</b> 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, <b>ioctl()</b> 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 <b>cb_ops</b> (9S). The kernel then packages the user request and arguments as integers and passes them to the driver's <b>ioctl()</b> 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)**.

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.

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*.

#### 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)** **cb\_ops(9S)**

*Writing Device Drivers*

#### WARNINGS

Non-STREAMS driver **ioctl()** routines must make sure that user data is copied into or out of the kernel address space explicitly using **copyin(9F)**, **copyout(9F)**, **ddi\_copyin(9F)**, or **ddi\_copyout(9F)**, as appropriate.

It is a severe error to simply dereference pointers to the user address space, even when in user context.

Failure to use the appropriate copying routines can result in panics under load on some platforms, and reproducible panics on others.

**NOTES**

STREAMS drivers do not have **ioctl** routines. The stream head converts I/O control commands to **M\_IOCTL** messages, which are handled by the driver's **put(9E)** or **srv(9E)** routine.

<b>NAME</b>	ks_update – dynamically update kstats
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/kstat.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefix_ks_update(kstat_t *ksp, int rw);</pre>
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Solaris DDI specific (Solaris DDI)</p> <p><i>ksp</i>        Pointer to a <b>kstat</b>(9S) structure.</p> <p><i>rw</i>         Read/Write flag. Possible values are</p> <p>          <b>KSTAT_READ</b>    Update kstat structure statistics from the driver.</p> <p>          <b>KSTAT_WRITE</b>    Update driver statistics from the kstat structure.</p>
<b>DESCRIPTION</b>	<p>The kstat mechanism allows for an optional <b>ks_update()</b> 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 <b>ks_update()</b> function which updates the kstat's data section on demand. To take advantage of this feature, set the <b>ks_update</b> field before calling <b>kstat_install</b>(9F).</p> <p>The <b>ks_update()</b> 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-&gt;ks_data */         /* return EACCES if you don't support this */     } else {         /* update ksp-&gt;ks_data from the native stats */     }     return (0); }</pre> <p>In general, the <b>ks_update()</b> routine may need to refer to provider-private data; for example, it may need a pointer to the provider's raw statistics. The <b>ks_private</b> 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 <b>ks_update()</b> routine. The caller will already be holding the kstat's <b>ks_lock</b> (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.

**SEE ALSO****kstat\_create(9F)**, **kstat\_install(9F)**, **kstat(9S)***Writing Device Drivers*

<b>NAME</b>	mapdev_access – device mapping access entry point						
<b>SYNOPSIS</b>	<pre>#include &lt;sys/sunddi.h&gt; int prefixmapdev_access(ddi_mapdev_handle_t handle, void *devprivate, off_t offset);</pre>						
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).						
<b>ARGUMENTS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><i>handle</i></td> <td>An opaque pointer to a device mapping.</td> </tr> <tr> <td><i>devprivate</i></td> <td>Driver private mapping data from <b>ddi_mapdev(9F)</b>.</td> </tr> <tr> <td><i>offset</i></td> <td>The offset within device memory at which the access occurred.</td> </tr> </table>	<i>handle</i>	An opaque pointer to a device mapping.	<i>devprivate</i>	Driver private mapping data from <b>ddi_mapdev(9F)</b> .	<i>offset</i>	The offset within device memory at which the access occurred.
<i>handle</i>	An opaque pointer to a device mapping.						
<i>devprivate</i>	Driver private mapping data from <b>ddi_mapdev(9F)</b> .						
<i>offset</i>	The offset within device memory at which the access occurred.						
<b>DESCRIPTION</b>	<p><b>mapdev_access()</b> is called when an access is made to a mapping that has either been newly created with <b>ddi_mapdev(9F)</b> or that has been enabled with a call to <b>ddi_mapdev_intercept(9F)</b>.</p> <p><b>mapdev_access()</b> 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 <b>ddi_mapdev_intercept(9F)</b> or <b>ddi_mapdev_nointercept(9F)</b> to control whether or not future accesses to the mapping will cause <b>mapdev_access()</b> to be called. In general, <b>mapdev_access()</b> should call <b>ddi_mapdev_intercept()</b> on the mapping that is currently in use and then call <b>ddi_mapdev_nointercept()</b> on the mapping that generated this call to <b>mapdev_access()</b>. This will ensure that a call to <b>mapdev_access()</b> will be generated for the current mapping next time it is accessed.</p> <p><b>mapdev_access()</b> must at least call <b>ddi_mapdev_nointercept()</b> 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 <b>ddi_mapdev_nointercept()</b> will not generate a call to <b>mapdev_access()</b>. A subsequent call to <b>ddi_mapdev_intercept()</b> will enable <b>mapdev_access()</b> to be called again.</p> <p>A non-zero return value from <b>mapdev_access()</b> will cause the corresponding operation to fail. The failure may result in a SIGSEGV or SIGBUS signal being delivered to the process.</p>						
<b>RETURN VALUES</b>	<b>mapdev_access()</b> should return <b>0</b> on success, <b>-1</b> if there was a hardware error, or the return value from <b>ddi_mapdev_intercept()</b> or <b>ddi_mapdev_nointercept()</b> .						
<b>CONTEXT</b>	This function is called from user context only.						

**EXAMPLES**

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*

<b>NAME</b>	mapdev_dup – device mapping duplication entry point
<b>SYNOPSIS</b>	<pre>#include &lt;sys/sunddi.h&gt; int prefixmapdev_dup(ddd_mapdev_handle_t handle, void *devprivate,     ddi_mapdev_handle_t new_handle, void **new_devprivatep);</pre>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>ARGUMENTS</b>	<p><i>handle</i>            The handle of the mapping that is being duplicated.</p> <p><i>devprivate</i>        Driver private mapping data from the mapping that is being duplicated.</p> <p><i>new_handle</i>        An opaque pointer to the duplicated device mapping.</p> <p><i>new_devprivatep</i>   A pointer to be filled in by the driver with the driver private mapping data for the duplicated device mapping.</p>
<b>DESCRIPTION</b>	<p><b>mapdev_dup()</b> is called when a device mapping is duplicated such as through <b>fork(2)</b>. <b>mapdev_dup()</b> 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 <b>mapdev_dup()</b> will cause the corresponding operation, such as <b>fork()</b> to fail.</p>
<b>RETURN VALUES</b>	<b>mapdev_dup()</b> returns <b>0</b> for success or the appropriate error number on failure.
<b>CONTEXT</b>	This function is called from user context only.
<b>EXAMPLES</b>	<pre>static int xxmapdev_dup(ddd_mapdev_handle_t handle, void *devprivate,     ddi_mapdev_handle_t new_handle, void **new_devprivate) {     struct xpvtdata        *pvtdata;     /* Allocate a new private data structure */     pvtdata = kmem_alloc(sizeof (struct xpvtdata), KM_SLEEP);     /* Copy the old data to the new - device dependent*/     ...     /* Return the new data */     *new_pvtdata = pvtdata;     return (0); }</pre>
<b>SEE ALSO</b>	<p><b>fork(2)</b>, <b>mmap(2)</b>, <b>mapdev_access(9E)</b>, <b>mapdev_free(9E)</b>, <b>segmap(9E)</b>, <b>ddd_mapdev(9F)</b>, <b>ddd_mapdev_intercept(9F)</b>, <b>ddd_mapdev_nointercept(9F)</b>, <b>ddd_mapdev_ctl(9S)</b>, <i>Writing Device Drivers</i></p>

<b>NAME</b>	mapdev_free – device mapping free entry point
<b>SYNOPSIS</b>	<b>#include</b> <sys/sunddi.h> <b>void prefixmapdev_free(ddd_mapdev_handle_t handle, void *devprivate);</b>
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI).
<b>ARGUMENTS</b>	<i>handle</i> An opaque pointer to a device mapping. <i>devprivate</i> Driver private mapping data from <b>ddd_mapdev(9F)</b> .
<b>DESCRIPTION</b>	<b>mapdev_free()</b> is called when a mapping created by <b>ddd_mapdev(9F)</b> is being destroyed. <b>mapdev_free()</b> 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> . The <b>mapdev_free()</b> routine is expected to free any resources that were allocated by the driver for this mapping.
<b>CONTEXT</b>	This function is called from user context only.
<b>EXAMPLES</b>	<pre>static void xxmapdev_free(ddd_mapdev_handle_t hdl, void *pvtdata) {     /* Destroy the driver private data - Device dependent */     ...     kmem_free(pvtdata, sizeof (struct xpvtdata)); }</pre>
<b>SEE ALSO</b>	<b>exit(2)</b> , <b>mmap(2)</b> , <b>munmap(2)</b> , <b>mapdev_access(9E)</b> , <b>mapdev_dup(9E)</b> , <b>segmap(9E)</b> , <b>ddd_mapdev(9F)</b> , <b>ddd_mapdev_intercept(9F)</b> , <b>ddd_mapdev_nointercept(9F)</b> , <b>ddd_mapdev_ctl(9S)</b> <i>Writing Device Drivers</i>



<b>NAME</b>	mmap – check virtual mapping for memory mapped device										
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/mman.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixmmap(dev_t dev, off_t off, int prot);</pre>										
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Architecture independent level 1 (DDI/DKI).</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" style="margin-left: 20px;"> <tr> <td><b>PROT_READ</b></td> <td>Read access will be granted.</td> </tr> <tr> <td><b>PROT_WRITE</b></td> <td>Write access will be granted.</td> </tr> <tr> <td><b>PROT_EXEC</b></td> <td>Execute access will be granted.</td> </tr> <tr> <td><b>PROT_USER</b></td> <td>User-level access will be granted.</td> </tr> <tr> <td><b>PROT_ALL</b></td> <td>All access will be granted.</td> </tr> </table>	<b>PROT_READ</b>	Read access will be granted.	<b>PROT_WRITE</b>	Write access will be granted.	<b>PROT_EXEC</b>	Execute access will be granted.	<b>PROT_USER</b>	User-level access will be granted.	<b>PROT_ALL</b>	All access will be granted.
<b>PROT_READ</b>	Read access will be granted.										
<b>PROT_WRITE</b>	Write access will be granted.										
<b>PROT_EXEC</b>	Execute access will be granted.										
<b>PROT_USER</b>	User-level access will be granted.										
<b>PROT_ALL</b>	All access will be granted.										
<b>DESCRIPTION</b>	<p>The <b>mmap(9E)</b> 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 <b>mmap(2)</b> 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 <b>mmap(9E)</b> entry point is called as a result of an <b>mmap(2)</b> system call, and also as a result of a page fault. <b>mmap(9E)</b> is called to translate the offset <i>off</i> in device memory to the corresponding physical page frame number.</p> <p>The <b>mmap(9E)</b> 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 <b>-1</b> is returned. If the offset does exist, <b>mmap(9E)</b> returns the value returned by <b>hat_getkpfnum(9F)</b> for the physical page in device memory containing the offset <i>off</i>.</p> <p><b>hat_getkpfnum(9F)</b> accepts a kernel virtual address as an argument. A kernel virtual address can be obtained by calling <b>ddi_map_regs(9F)</b> in the driver's <b>attach(9E)</b> routine. The corresponding <b>ddi_unmap_regs(9F)</b> call can be made in the driver's <b>detach(9E)</b> routine. Refer to the <b>EXAMPLES</b> section below for more information.</p> <p><b>mmap(9E)</b> should only be supported for memory-mapped devices. See the <b>segmap(9E)</b> and <b>ddi_mapdev(9F)</b> reference pages for further information on memory-mapped device drivers.</p>										

**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**

The following is an example of the **mmap**(9E) 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\_map\_regs**(9F) can be used to obtain this address. For example, **ddi\_map\_regs**(9F) can be called in the driver's **attach**(9E) routine. The resulting kernel virtual address is stored in the **xxstate** structure (see **ddi\_soft\_state**(9F)), which is accessible from the driver's **mmap**(9E) entry point. The corresponding **ddi\_unmap\_regs**(9F) call can be made in the driver's **detach**(9E) routine.

```

    struct reg {
        char    csr;
        char    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;

        /* 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));
    }

```

}

**SEE ALSO**

**mmap(2)**, **attach(9E)**, **ddi\_map\_regs(9F)**, **ddi\_unmap\_regs(9F)**, **detach(9E)**, **segmap(9E)**, **ddi\_btop(9F)**, **ddi\_get\_soft\_state(9F)**, **ddi\_mapdev(9F)**, **getminor(9F)**, **hat\_getkpfnum(9F)**

*Writing Device Drivers*

**NOTES**

For some devices, mapping device memory in the driver's **attach(9E)** routine and unmaping 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\_map\_regs(9F)** and **ddi\_unmap\_regs(9F)** in **mmap**. These function calls would bracket the call to **hat\_getkpfnum(9F)**.

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

<b>NAME</b>	open – gain access to a device
<b>SYNOPSIS</b> Block and Character	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/file.h&gt; #include &lt;sys/errno.h&gt; #include &lt;sys/open.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixopen(dev_t *devp, int flag, int otyp, cred_t *cred_p);</pre>
<b>STREAMS</b>	<pre>#include &lt;sys/file.h&gt; #include &lt;sys/stream.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixopen(queue_t *q, dev_t *devp, int oflag, int sflag, cred_t *cred_p);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI). This entry point is <b>required</b> , but it can be <b>nulldev(9F)</b> .
<b>ARGUMENTS</b> 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 <b>open(2)</b> system call that instructs the driver on how to open the file. Valid settings are:</p> <p><b>FEXCL</b> Open the device with exclusive access; fail all other attempts to open the device.</p> <p><b>FNDELAY</b> Open the device and return immediately (do not block the open even if something is wrong).</p> <p><b>FREAD</b> Open the device with read-only permission (if ORed with <b>FWRITE</b>, then allow both read and write access)</p> <p><b>FWRITE</b> Open a device with write-only permission (if ORed with <b>FREAD</b>, then allow both read and write access)</p> <p><i>otyp</i> Parameter supplied so that the driver can determine how many times a device was opened and for what reasons.</p> <p>For <b>OTYP_BLK</b> and <b>OTYP_CHR</b>, the <b>open()</b> routine may be called many times, but the <b>close(9E)</b> 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 <b>close(2)</b> or <b>exit(2)</b>. If the device is accessed through memory mapping, this is by a call to <b>munmap(2)</b> or <b>exit(2)</b>.</p> <p>For <b>OTYP_LYR</b>, there is exactly one <b>close(9E)</b> for each <b>open()</b> 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>

## STREAMS

<b>OTYP_BLK</b>	Open occurred through block interface for the device
<b>OTYP_CHR</b>	Open occurred through the raw/character interface for the device
<b>OTYP_LYR</b>	Open a layered process. This flag is used when one driver calls another driver's <b>open</b> or <b>close</b> (9E) 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.
<i>cred_p</i>	Pointer to the user credential structure.
<i>q</i>	A pointer to the read <b>queue</b> .
<i>devp</i>	Pointer to a device number. For STREAMS modules, <i>devp</i> always points to the device number associated with the driver at the end (tail) of the stream.
<i>oflag</i>	Valid <i>oflag</i> values are <b>FEXCL</b> , <b>FNDELAY</b> , <b>FREAD</b> , and <b>FWRITEL</b> , the same as those listed above for <i>flag</i> . For STREAMS modules, <i>oflag</i> is always set to <b>0</b> .
<i>sflag</i>	Valid values are as follows:
<b>CLONEOPEN</b>	Indicates that the <b>open</b> routine is called through the clone driver. The driver should return a unique device number.
<b>MODOPEN</b>	Modules should be called with <i>sflag</i> set to this value. Modules should return an error if they are called with <i>sflag</i> set to a different value. Drivers should return an error if they are called with <i>sflag</i> set to this value.
<b>0</b>	Indicates a driver is opened directly, without calling the clone driver.
<i>cred_p</i>	Pointer to the user credential structure.

## DESCRIPTION

The driver's **open()** routine is called by the kernel during an **open(2)** or a **mount(2)** on the special file for the device. The routine should verify that the minor number component of *\*devp* is valid, that the type of access requested by *otyp* and *flag* is appropriate for the device, and, if required, check permissions using the user credentials pointed to by *cred\_p*.

The **open()** 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 **attach(9E)** with **ddi\_create\_minor\_node(9F)**, then changes the minor number component of *\*devp* using **makedevice(9F)** and **getmajor(9F)**. The driver needs to keep track of available minor numbers internally.

**\*devp = makedevice(getmajor(\*devp), new\_minor);**

**RETURN VALUES**

The **open()** routine should return 0 for success, or the appropriate error number.

**SEE ALSO**

**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)**

*Writing Device Drivers*

*STREAMS Programming Guide*

**WARNINGS**

Do not attempt to change the major number.

<b>NAME</b>	print – display a driver message on system console
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/errno.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixprint(dev_t dev, char *str);</pre>
<b>INTERFACE LEVEL</b>	Architecture independent level 1 (DDI/DKI). This entry point is <b>required</b> for block devices.
<b>ARGUMENTS</b>	<p><i>dev</i> Device number.</p> <p><i>str</i> Pointer to a character string describing the problem.</p>
<b>DESCRIPTION</b>	The <b>print()</b> 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 <b>cmn_err(9F)</b> kernel function. The driver should print the message along with any driver specific information.
<b>RETURN VALUES</b>	The <b>print()</b> routine should return <b>0</b> for success, or the appropriate error number. The <b>print</b> routine can fail if the driver implemented a non-standard <b>print()</b> routine that attempted to perform error logging, but was unable to complete the logging for whatever reason.
<b>SEE ALSO</b>	<b>cmn_err(9F)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	probe – determine if a non-self-identifying device is present								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/conf.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  static int  prefixprobe(dev_info_t *dip);</pre>								
<b>INTERFACE LEVEL</b>	Solaris DDI specific (Solaris DDI). This entry point is <b>required</b> for non-self-identifying devices. You must write it for such devices. For self-identifying devices, <b>nulldev(9F)</b> should be specified in the <b>dev_ops(9S)</b> structure if a probe routine is not necessary.								
<b>ARGUMENTS</b>	<i>dip</i> Pointer to the device's <b>dev_info</b> structure.								
<b>DESCRIPTION</b>	<p><b>probe()</b> determines whether the device corresponding to <i>dip</i> actually exists and is a valid device for this driver. <b>probe()</b> is called after <b>identify(9E)</b> and before <b>attach(9E)</b> for a given <i>dip</i>. For example, the <b>probe()</b> routine can map the device registers using <b>ddi_map_regs(9F)</b> then attempt to access the hardware using <b>ddi_peek(9F)</b> and/or <b>ddi_poke(9F)</b> and determine if the device exists. Then the device registers should be unmapped using <b>ddi_unmap_regs(9F)</b>.</p> <p><b>probe()</b> should only probe the device – it should not create or change any software state. Device initialization should be done in <b>attach(9E)</b>.</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 <b>probe()</b> routine can be provided to simplify the driver. <b>ddi_dev_is_sid(9F)</b> can then be used to determine whether <b>probe()</b> needs to do any work. See <b>ddi_dev_is_sid(9F)</b> for an example.</p>								
<b>RETURN VALUES</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><b>DDI_PROBE_SUCCESS</b></td> <td>if the probe was successful.</td> </tr> <tr> <td><b>DDI_PROBE_FAILURE</b></td> <td>if the probe failed.</td> </tr> <tr> <td><b>DDI_PROBE_DONTCARE</b></td> <td>if the probe was unsuccessful, yet <b>attach(9E)</b> should still be called.</td> </tr> <tr> <td><b>DDI_PROBE_PARTIAL</b></td> <td>if the instance is not present now, but may be present in the future.</td> </tr> </table>	<b>DDI_PROBE_SUCCESS</b>	if the probe was successful.	<b>DDI_PROBE_FAILURE</b>	if the probe failed.	<b>DDI_PROBE_DONTCARE</b>	if the probe was unsuccessful, yet <b>attach(9E)</b> should still be called.	<b>DDI_PROBE_PARTIAL</b>	if the instance is not present now, but may be present in the future.
<b>DDI_PROBE_SUCCESS</b>	if the probe was successful.								
<b>DDI_PROBE_FAILURE</b>	if the probe failed.								
<b>DDI_PROBE_DONTCARE</b>	if the probe was unsuccessful, yet <b>attach(9E)</b> should still be called.								
<b>DDI_PROBE_PARTIAL</b>	if the instance is not present now, but may be present in the future.								
<b>SEE ALSO</b>	<p><b>attach(9E)</b>, <b>identify(9E)</b>, <b>ddi_dev_is_sid(9F)</b>, <b>ddi_map_regs(9F)</b>, <b>ddi_peek(9F)</b>, <b>ddi_poke(9F)</b>, <b>nulldev(9F)</b>, <b>dev_ops(9S)</b></p> <p><i>Writing Device Drivers</i></p>								



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

**RETURN VALUES****prop\_op()** should return:

<b>DDI_PROP_SUCCESS</b>	Property found and returned.
<b>DDI_PROP_NOT_FOUND</b>	Property not found.
<b>DDI_PROP_UNDEFINED</b>	Prop explicitly undefined.
<b>DDI_PROP_NO_MEMORY</b>	Property found, but unable to allocate memory. <i>lengthp</i> has the correct property length.
<b>DDI_PROP_BUF_TOO_SMALL</b>	Property found, but the supplied buffer is too small. <i>lengthp</i> has the correct property length.

**EXAMPLES**

```
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_NOTFOUND);

    if (!strcmp(name, "nblocks")) {
        ddi_prop_modify(dev, dip, "nblocks", flags,
            &xsp->nblocks, sizeof(int));
    }
    /*      other cases...      */

skip:
    return (ddi_prop_op(dev, dip, prop_op, flags, name,
        valuep, lengthp));
}
```

**SEE ALSO****ddi\_prop\_create(9F), ddi\_prop\_op(9F)***Writing Device Drivers*

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

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), **streamtab**(9S)

*Writing Device Drivers*

*STREAMS Programming Guide*

<b>NAME</b>	read – read data from a device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/errno.h&gt; #include &lt;sys/open.h&gt; #include &lt;sys/uio.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixread(dev_t dev, struct uio *uio_p, cred_t *cred_p);</pre>
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Architecture independent level 1 (DDI/DKI). This entry point is <b>optional</b>.</p> <p><i>dev</i>        Device number.</p> <p><i>uio_p</i>      Pointer to the <b>uio(9S)</b> structure that describes where the data is to be stored in user space.</p> <p><i>cred_p</i>     Pointer to the user credential structure for the I/O transaction.</p>
<b>DESCRIPTION</b>	The driver <b>read()</b> routine is called indirectly through <b>cb_ops(9S)</b> by the <b>read(2)</b> system call. The <b>read()</b> 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 <b>read()</b> routine should supervise the data transfer into the user space described by the <b>uio(9S)</b> structure.
<b>RETURN VALUES</b>	The <b>read()</b> routine should return <b>0</b> for success, or the appropriate error number.
<b>SEE ALSO</b>	<b>read(2), write(9E), cb_ops(9S), uio(9S)</b> <i>Writing Device Drivers</i>

<b>NAME</b>	segmap – map device memory into user space
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/mman.h&gt; #include &lt;sys/param.h&gt; #include &lt;sys/vm.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixsegmap(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>
<b>INTERFACE LEVEL ARGUMENTS</b>	<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> <ul style="list-style-type: none"> <li><b>PROT_READ</b> Read access is desired.</li> <li><b>PROT_WRITE</b> Write access is desired.</li> <li><b>PROT_EXEC</b> Execute access is desired.</li> <li><b>PROT_USER</b> User-level access is desired (the mapping is being done as a result of a <b>mmap(2)</b> system call).</li> <li><b>PROT_ALL</b> All access is desired.</li> </ul> <p><i>maxprot</i> Maximum protection flag possible for attempted mapping (the <b>PROT_WRITE</b> bit may be masked out if the user opened the special file read-only). If (<b>maxprot &amp; prot</b>) != <b>prot</b> then there is an access violation.</p> <p><i>flags</i> Flags indicating type of mapping. Possible values are (other bits may be set):</p> <ul style="list-style-type: none"> <li><b>MAP_SHARED</b> Changes should be shared.</li> <li><b>MAP_PRIVATE</b> Changes are private.</li> </ul> <p><i>cred_p</i> Pointer to the user credentials structure.</p>
<b>DESCRIPTION</b>	<p>The <b>segmap()</b> entry point is an optional routine for character drivers that support memory mapping. The <b>mmap(2)</b> 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>

Typically, a character driver that needs to support the **mmap(2)** system call supplies either an **mmap(9E)** entry point, or both an **mmap(9E)** and a **segmap()** entry point routine (see the **mmap(9E)** reference page). If no **segmap()** entry point is provided for the driver, **ddi\_segmap(9F)** is used as a default.

A driver for a memory-mapped device would provide a **segmap()** entry point if it:

- needs to maintain a separate context for each user mapping. See **ddi\_mapdev(9F)** for details.
- needs to assign device access attributes to the user mapping. See **ddi\_segmap\_setup(9F)** and **ddi\_mapdev\_set\_device\_acc\_attr(9F)** for details.

The responsibilities of a **segmap()** entry point are:

- Verify that the range, defined by *offset* and *len*, to be mapped is valid for the device. Typically, this task is performed by calling the **mmap(9E)** entry point. Note that if you are using **ddi\_segmap()** to set up the mapping, it will call your **mmap(9E)** entry point for you to validate the range to be mapped.
- Assign device access attributes to the mapping. See **ddi\_mapdev\_set\_device\_acc\_attr(9F)**, **ddi\_segmap\_setup(9F)**, and **ddi\_device\_acc\_attr(9S)** for details.
- Set up device contexts for the user mapping if your device requires context switching. See **ddi\_mapdev(9F)** for details.
- Perform the mapping with **ddi\_segmap()**, **ddi\_segmap\_setup()**, or **ddi\_mapdev()** and return the status if it fails.

## RETURN VALUES

The **segmap()** routine should return **0** if the driver is successful in performing the memory map of its device address space into the specified address space.

The **segmap()** must return an error number on failure. For example, valid error numbers would be **ENXIO** if the offset/length pair specified exceeds the limits of the device memory, or **EINVAL** if the driver detects an invalid type of mapping attempted.

If one of the mapping routines **ddi\_segmap()**, **ddi\_segmap\_setup()**, or **ddi\_mapdev()** fails, you must return the error number returned by the respective routine.

## SEE ALSO

**mmap(2)**, **mmap(9E)**, **ddi\_mapdev(9F)**, **ddi\_mapdev\_set\_device\_acc\_attr(9F)**, **ddi\_segmap(9F)**, **ddi\_segmap\_setup(9F)**, **ddi\_device\_acc\_attr(9S)**

*Writing Device Drivers*

<b>NAME</b>	srv – service queued messages
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/stream.h&gt; #include &lt;sys/stropts.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixrsrv(queue_t *q); /* read side */ int prefixwsrv(queue_t *q); /* write side */</pre>
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Architecture independent level 1 (DDI/DKI). This entry point is <b>required</b> for STREAMS.</p> <p><i>q</i>        Pointer to the <b>queue(9S)</b> structure.</p>
<b>DESCRIPTION</b>	<p>The optional service (<b>srv()</b>) routine may be included in a STREAMS module or driver for many possible reasons, including:</p> <ul style="list-style-type: none"> <li>• to provide greater control over the flow of messages in a stream</li> <li>• to make it possible to defer the processing of some messages to avoid depleting system resources</li> <li>• to combine small messages into larger ones, or break large messages into smaller ones</li> <li>• to recover from resource allocation failure. A module's or driver's <b>put(9E)</b> routine can test for the availability of a resource, and if it is not available, enqueue the message for later processing by the <b>srv</b> routine.</li> </ul> <p>A message is first passed to a module's or driver's <b>put(9E)</b> routine, which may or may not do some processing. It must then either:</p> <ul style="list-style-type: none"> <li>• Pass the message to the next stream component with <b>putnext(9F)</b>.</li> <li>• If a <b>srv</b> routine has been included, it may call <b>putq(9F)</b> to place the message on the queue</li> </ul> <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 <b>srv</b> 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 <b>srv</b> routines.</p> <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 <b>M_IOCACK</b>). High priority messages are always placed at the head of the <b>srv</b> routine's</p>



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)**, **queue(9S)**

*Writing Device Drivers*  
*STREAMS Programming Guide*

#### 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 of **nulldev(9F)** for a **srv()** routine may result in flow control errors.

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

<b>NAME</b>	tran_abort – abort a SCSI command
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; int prefixtran_abort(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Solaris architecture specific (Solaris DDI).</p> <p><i>ap</i>     Pointer to a <b>scsi_address</b>(9S) structure.</p> <p><i>pkt</i>     Pointer to a <b>scsi_pkt</b>(9S) structure.</p>
<b>DESCRIPTION</b>	<p>The <b>tran_abort()</b> vector in the <b>scsi_hba_tran</b>(9S) structure must be initialized during the HBA driver's <b>attach</b>(9E) to point to an HBA entry point to be called when a target driver calls <b>scsi_abort</b>(9F).</p> <p><b>tran_abort()</b> 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, <b>tran_abort()</b> must set the <b>pkt_reason</b> to <b>CMD_ABORTED</b>, and <b>pkt_statistics</b> must be OR'ed with <b>STAT_ABORTED</b>.</p>
<b>RETURN VALUES</b>	<p><b>tran_abort()</b> must return:</p> <p><b>1</b>     on success or partial success.</p> <p><b>0</b>     on failure.</p>
<b>SEE ALSO</b>	<p><b>attach</b>(9E), <b>scsi_abort</b>(9F), <b>scsi_hba_attach</b>(9F), <b>scsi_address</b>(9S), <b>scsi_hba_tran</b>(9S), <b>scsi_pkt</b>(9S)</p> <p><i>Writing Device Drivers</i></p>
<b>NOTES</b>	<p>If <b>pkt_reason</b> already indicates that an earlier error had occurred, <b>tran_abort()</b> should not overwrite <b>pkt_reason</b> with <b>CMD_ABORTED</b>.</p>

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

<b>NAME</b>	tran_getcap, tran_setcap – get/set SCSI transport capability
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; 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>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>ARGUMENTS</b>	<p><i>ap</i> Pointer to the <b>scsi_address</b>(9S) 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>
<b>DESCRIPTION</b>	<p>The <b>tran_getcap()</b> and <b>tran_setcap()</b> vectors in the <b>scsi_hba_tran</b>(9S) structure must be initialized during the HBA driver's <b>attach</b>(9E) to point to HBA entry points to be called when a target driver calls <b>scsi_ifgetcap</b>(9F) and <b>scsi_ifsetcap</b>(9F).</p> <p><b>tran_getcap()</b> 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 <b>scsi_address</b>(9S) structure pointed to by <i>ap</i>; if <i>whom</i> is <b>0</b> all targets are affected, else the target specified by the <b>scsi_address</b> structure pointed to by <i>ap</i> is affected.</p> <p><b>tran_setcap()</b> 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 <b>scsi_address</b>(9S) structure pointed to by <i>ap</i>; if <i>whom</i> is <b>0</b> all targets are affected, else the target specified by the <b>scsi_address</b> structure pointed to by <i>ap</i> is affected.</p> <p>A device may support only a subset of the defined capabilities.</p> <p>Refer to <b>scsi_ifgetcap</b>(9F) for the list of defined capabilities.</p> <p>HBA drivers should use <b>scsi_hba_lookup_capstr</b>(9F) to match <i>cap</i> against the canonical capability strings.</p>
<b>RETURN VALUES</b>	<p><b>tran_setcap()</b> must return <b>1</b> if the capability was successfully set to the new value, <b>0</b> if the HBA driver does not support changing the capability, and <b>-1</b> if the capability was not defined.</p> <p><b>tran_getcap()</b> must return the current value of a capability or <b>-1</b> if the capability was not defined.</p>
<b>SEE ALSO</b>	<p><b>attach</b>(9E), <b>scsi_hba_attach</b>(9F), <b>scsi_hba_lookup_capstr</b>(9F), <b>scsi_ifgetcap</b>(9F), <b>scsi_address</b>(9S), <b>scsi_hba_tran</b>(9S)</p> <p><i>Writing Device Drivers</i></p>

<b>NAME</b>	tran_init_pkt, tran_destroy_pkt – SCSI HBA packet preparation and deallocation
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; 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>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>ARGUMENTS</b>	<p><i>ap</i> Pointer to a <b>scsi_address</b>(9S) structure.</p> <p><i>pkt</i> Pointer to a <b>scsi_pkt</b>(9S) structure allocated in an earlier call, or NULL.</p> <p><i>bp</i> Pointer to a <b>buf</b>(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 <b>scsi_pkt</b> 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>
<b>DESCRIPTION</b>	<p>The <b>tran_init_pkt()</b> and <b>tran_destroy_pkt()</b> vectors in the <b>scsi_hba_tran</b> structure must be initialized during the HBA driver's <b>attach</b>(9E) to point to HBA entry points to be called when a target driver calls <b>scsi_init_pkt</b>(9F) and <b>scsi_destroy_pkt</b>(9F).</p> <p><b>tran_init_pkt()</b> is the entry point into the HBA which is used to allocate and initialize a <b>scsi_pkt</b> structure on behalf of a SCSI target driver. If <i>pkt</i> is NULL, the HBA driver must use <b>scsi_hba_pkt_alloc</b>(9F) to allocate a new <b>scsi_pkt</b> 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, via <b>ddi_dma_buf_setup</b>(9F).</p> <p>If the <b>PKT_CONSISTENT</b> bit is set in <i>flags</i>, the buffer was allocated by <b>scsi_alloc_consistent_buf</b>(9F). For packets marked with <b>PKT_CONSISTENT</b> the HBA driver must synchronize any cached data transfers before calling the target driver's command completion callback.</p> <p>If the <b>PKT_DMA_PARTIAL</b> bit is set in <i>flags</i>, the HBA driver should set up partial data transfers, such as setting the <b>DDI_DMA_PARTIAL</b> bit in the <i>flags</i> argument if interfaces such as <b>ddi_dma_buf_setup</b>(9F) are used.</p>

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 *pktp* and *bp* are non-NULL, the **PKT\_DMA\_PARTIAL** bit is set in *flags* and 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 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*.

*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.

If *pkt* is NULL on entry, and **tran\_init\_pkt()** allocated a *pkt* via **scsi\_hba\_pkt\_alloc(9F)** but was unable to allocate DMA resources, **tran\_init\_pkt()** must free the *pkt* via **scsi\_hba\_pkt\_free(9F)** before returning NULL.

#### SEE ALSO

**attach(9E)**, **tran\_sync\_pkt(9E)**, **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**, the **B\_ERROR** flag should be set in *bp*, and the **b\_error** field should be set to **EFAULT**.

If a DMA allocation request fails with **DDI\_DMA\_TOOBIG**, the **B\_ERROR** flag should be set in *bp*, and the **b\_error** field should be set to **EINVAL**.





<b>NAME</b>	tran_reset – reset a SCSI bus or target
<b>SYNOPSIS</b>	<b>#include</b> <sys/scsi/scsi.h> <b>int</b> <i>prefix</i> tran_reset(struct scsi_address *ap, int level);
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>ARGUMENTS</b>	<i>ap</i> Pointer to the scsi_address(9S) structure. <i>level</i> The level of reset required.
<b>DESCRIPTION</b>	The tran_reset() 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_reset(9F). <b>tran_reset()</b> must reset the SCSI bus or a SCSI target as specified by <i>level</i> . <i>level</i> must be one of the following: <b>RESET_ALL</b> reset the SCSI bus. <b>RESET_TARGET</b> reset the target specified by <i>ap</i> . <b>tran_reset</b> should set the <b>pkt_reason</b> field of all outstanding packets in the transport layer associated with each target that was successfully reset to <b>CMD_RESET</b> and the <b>pkt_statistics</b> field must be OR'ed with either <b>STAT_BUS_RESET</b> or <b>STAT_DEV_RESET</b> . The HBA driver should use a SCSI Bus Device Reset Message to reset a target device. Packets that are in the transport layer but not yet active on the bus should be returned with <b>pkt_reason</b> set to <b>CMD_RESET</b> , and <b>pkt_statistics</b> OR'ed with <b>STAT_ABORTED</b> .
<b>RETURN VALUES</b>	<b>tran_reset()</b> should return: <b>1</b> on success. <b>0</b> on failure.
<b>SEE ALSO</b>	<b>attach(9E)</b> , <b>ddi_dma_buf_setup(9F)</b> , <b>scsi_hba_attach(9F)</b> , <b>scsi_reset(9F)</b> , <b>scsi_address(9S)</b> , <b>scsi_hba_tran(9S)</b> <i>Writing Device Drivers</i>
<b>NOTES</b>	If <b>pkt_reason</b> already indicates that an earlier error had occurred for a particular <i>pkt</i> , <b>tran_reset()</b> should not overwrite <b>pkt_reason</b> with <b>CMD_RESET</b> .

<b>NAME</b>	tran_reset_notify – request to notify SCSI target of bus reset								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; int prefixtran_reset_notify(struct scsi_address *ap, int flag, void (*callback)(caddr_t),     caddr_t arg);</pre>								
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).								
<b>ARGUMENTS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><i>ap</i></td> <td>Pointer to the <b>scsi_address</b>(9S) structure.</td> </tr> <tr> <td><i>flag</i></td> <td>A flag indicating registration or cancellation of a notification request.</td> </tr> <tr> <td><i>callback</i></td> <td>A pointer to the target driver's reset notification function.</td> </tr> <tr> <td><i>arg</i></td> <td>The callback function argument.</td> </tr> </table>	<i>ap</i>	Pointer to the <b>scsi_address</b> (9S) structure.	<i>flag</i>	A flag indicating registration or cancellation of a notification request.	<i>callback</i>	A pointer to the target driver's reset notification function.	<i>arg</i>	The callback function argument.
<i>ap</i>	Pointer to the <b>scsi_address</b> (9S) structure.								
<i>flag</i>	A flag indicating registration or cancellation of a notification request.								
<i>callback</i>	A pointer to the target driver's reset notification function.								
<i>arg</i>	The callback function argument.								
<b>DESCRIPTION</b>	<p>The <b>tran_reset_notify()</b> entry point is called when a target driver requests notification of a bus reset.</p> <p>The <b>tran_reset_notify()</b> vector in the <b>scsi_hba_tran</b>(9S) structure may be initialized in the HBA driver's <b>attach</b>(9E) routine to point to the HBA entry point to be called when a target driver calls <b>scsi_reset_notify</b>(9F).</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" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;"><b>SCSI_RESET_NOTIFY</b></td> <td>Register <i>callback</i> as the reset notification function for the target.</td> </tr> <tr> <td><b>SCSI_RESET_CANCEL</b></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>	<b>SCSI_RESET_NOTIFY</b>	Register <i>callback</i> as the reset notification function for the target.	<b>SCSI_RESET_CANCEL</b>	Cancel the reset notification request for the target.				
<b>SCSI_RESET_NOTIFY</b>	Register <i>callback</i> as the reset notification function for the target.								
<b>SCSI_RESET_CANCEL</b>	Cancel the reset notification request for the target.								
<b>RETURN VALUES</b>	<p>For <b>SCSI_RESET_NOTIFY</b> requests, <b>tran_reset_notify()</b> must return <b>DDI_SUCCESS</b> if the notification request has been accepted, and <b>DDI_FAILURE</b> otherwise.</p> <p>For <b>SCSI_RESET_CANCEL</b> requests, <b>tran_reset_notify()</b> must return <b>DDI_SUCCESS</b> if the notification request has been canceled, and <b>DDI_FAILURE</b> otherwise.</p>								
<b>SEE ALSO</b>	<b>attach</b> (9E), <b>scsi_ifgetcap</b> (9F), <b>scsi_reset_notify</b> (9F), <b>scsi_address</b> (9S), <b>scsi_hba_tran</b> (9S) <i>Writing Device Drivers</i>								

<b>NAME</b>	tran_start – request to transport a SCSI command
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; int prefixtran_start(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>ARGUMENTS</b>	<p><i>pkt</i>            Pointer to the <b>scsi_pkt</b>(9S) structure that is about to be transferred.</p> <p><i>ap</i>             Pointer to a <b>scsi_address</b>(9S) structure.</p>
<b>DESCRIPTION</b>	<p>The <b>tran_start()</b> vector in the <b>scsi_hba_tran</b>(9S) structure must be initialized during the HBA driver's <b>attach</b>(9E) to point to an HBA entry point to be called when a target driver calls <b>scsi_transport</b>(9F).</p> <p><b>tran_start()</b> 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>pkt</i> structure.</p> <p>If the flag <b>FLAG_NOINTR</b> is set in <b>pkt_flags</b> in <i>pkt</i>, <b>tran_start()</b> should not return until the command has been completed. The command completion callback <b>pkt_comp</b> in <i>pkt</i> must not be called for commands with <b>FLAG_NOINTR</b> set, since the return is made directly to the function invoking <b>scsi_transport</b>(9F).</p> <p>When the flag <b>FLAG_NOINTR</b> is not set, <b>tran_start()</b> must queue the command for execution on the hardware and return immediately. The member <b>pkt_comp</b> in <i>pkt</i> indicates a callback routine to be called upon command completion.</p> <p>Refer to <b>scsi_pkt</b>(9S) for other bits in <b>pkt_flags</b> for which the HBA driver may need to adjust how the command is managed.</p> <p>If the <b>auto_rqsense</b> capability has been set, and the status length allocated in <b>tran_init_pkt</b>(9E) is greater than or equal to <b>sizeof(struct scsi_arq_status)</b>, 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 <b>scsi_arq_status</b> structure pointed to by <b>pkt_scbp</b> updated with the results of this Request Sense command before the HBA driver completes the command pointed to by <i>pkt</i>.</p> <p>The member <b>pkt_time</b> in <i>pkt</i> is the maximum number of seconds in which the command should complete. A <b>pkt_time</b> of zero 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 <b>pkt_reason</b> member of the timed-out <i>pkt</i> should be set to <b>CMD_TIMEOUT</b>, and <b>pkt_statistics</b> OR'ed with <b>STAT_TIMEOUT</b>. If the HBA driver can successfully recover from the timeout, <b>pkt_statistics</b> must also be OR'ed with one of <b>STAT_ABORTED</b>, <b>STAT_BUS_RESET</b> or <b>STAT_DEV_RESET</b>, as appropriate. This informs the target driver that timeout recovery has already been successfully accomplished for the timed-out command. The <b>pkt_comp</b> completion callback, if not <b>NULL</b>, must also be</p>

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.

<b>pkt_resid</b>	For commands with data transfer, this member must be updated to indicate the residual of the data transferred.
<b>pkt_reason</b>	The reason for the command completion. This field should be set to <b>CMD_CMPLT</b> at the beginning of <b>tran_start()</b> , then updated if the command ever transitions to an abnormal termination state. To avoid losing information, do not set <b>pkt_reason</b> to any other error state unless it still has its original <b>CMD_CMPLT</b> value.
<b>pkt_statistics</b>	Bit field of transport-related statistics
<b>pkt_state</b>	Bit field with the major states through which a SCSI command can transition.
	<b>Note:</b> the members listed above, and <b>pkt_hba_private</b> member, are the only fields in the <b>scsi_pkt(9S)</b> structure which may be modified by the transport layer.

**RETURN VALUES****tran\_start()** must return:

<b>TRAN_ACCEPT</b>	The packet was accepted by the transport layer.
<b>TRAN_BUSY</b>	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.
<b>TRAN_BADPKT</b>	The DMA count in the packet exceeded the DMA engine's maximum DMA size, or the packet could not be accepted for other reasons.
<b>TRAN_FATAL_ERROR</b>	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\_hba\_tran(9S), scsi\_pkt(9S)***Writing Device Drivers*

<b>NAME</b>	tran_sync_pkt – SCSI HBA memory synchronization entry point
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; void prefixtran_sync_pkt(struct scsi_address *ap, struct scsi_pkt *pkt);</pre>
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).
<b>ARGUMENTS</b>	<p><i>ap</i>            A pointer to a <b>scsi_address</b>(9S) structure.</p> <p><i>pkt</i>            A pointer to a <b>scsi_pkt</b>(9S) structure.</p>
<b>DESCRIPTION</b>	<p>The <b>tran_sync_pkt()</b> vector in the <b>scsi_hba_tran</b>(9S) structure must be initialized during the HBA driver's <b>attach</b>(9E) to point to an HBA driver entry point to be called when a target driver calls <b>scsi_sync_pkt</b>(9F).</p> <p><b>tran_sync_pkt()</b> must synchronize a CPU's or device's view of the data associated with the <i>pkt</i>, typically by calling <b>ddi_dma_sync</b>(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>
<b>SEE ALSO</b>	<b>tran_init_pkt</b> (9E), <b>ddi_dma_sync</b> (9F), <b>scsi_hba_attach</b> (9F), <b>scsi_init_pkt</b> (9F), <b>scsi_sync_pkt</b> (9F), <b>scsi_hba_tran</b> (9S), <b>scsi_pkt</b> (9S) <i>Writing Device Drivers</i>
<b>NOTES</b>	A target driver may call <b>tran_sync_pkt()</b> on packets for which no DMA resources were allocated.

<b>NAME</b>	tran_tgt_free – request to free HBA resources allocated on behalf of a target								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; 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>								
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).								
<b>ARGUMENTS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><i>hba_dip</i></td> <td>Pointer to a <b>dev_info_t</b> structure, referring to the HBA device instance.</td> </tr> <tr> <td><i>tgt_dip</i></td> <td>Pointer to a <b>dev_info_t</b> structure, referring to the target device instance.</td> </tr> <tr> <td><i>hba_tran</i></td> <td>Pointer to a <b>scsi_hba_tran</b>(9S) structure, consisting of the HBA's transport vectors.</td> </tr> <tr> <td><i>sd</i></td> <td>Pointer to a <b>scsi_device</b>(9S) structure, describing the target.</td> </tr> </table>	<i>hba_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the HBA device instance.	<i>tgt_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the target device instance.	<i>hba_tran</i>	Pointer to a <b>scsi_hba_tran</b> (9S) structure, consisting of the HBA's transport vectors.	<i>sd</i>	Pointer to a <b>scsi_device</b> (9S) structure, describing the target.
<i>hba_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the HBA device instance.								
<i>tgt_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the target device instance.								
<i>hba_tran</i>	Pointer to a <b>scsi_hba_tran</b> (9S) structure, consisting of the HBA's transport vectors.								
<i>sd</i>	Pointer to a <b>scsi_device</b> (9S) structure, describing the target.								
<b>DESCRIPTION</b>	<p>The <b>tran_tgt_free()</b> vector in the <b>scsi_hba_tran</b>(9S) structure may be initialized during the HBA driver's <b>attach</b>(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 <b>tran_tgt_free()</b> vector, if not NULL, is called after the target device instance has returned successfully from its <b>detach</b>(9E) entry point, but before the <b>dev_info</b> node structure is removed from the system. The HBA driver should release any resources allocated during its <b>tran_tgt_init()</b> or <b>tran_tgt_probe()</b> initialization performed for this target device instance.</p>								
<b>SEE ALSO</b>	<p><b>attach</b>(9E), <b>detach</b>(9E), <b>tran_tgt_init</b>(9E), <b>tran_tgt_probe</b>(9E), <b>scsi_device</b>(9S), <b>scsi_hba_tran</b>(9S)</p> <p><i>Writing Device Drivers</i></p>								

<b>NAME</b>	tran_tgt_init – request to initialize HBA resources on behalf of a particular target								
<b>SYNOPSIS</b>	<pre>#include &lt;sys/scsi/scsi.h&gt; 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>								
<b>INTERFACE LEVEL</b>	Solaris architecture specific (Solaris DDI).								
<b>ARGUMENTS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><i>hba_dip</i></td> <td>Pointer to a <b>dev_info_t</b> structure, referring to the HBA device instance.</td> </tr> <tr> <td><i>tgt_dip</i></td> <td>Pointer to a <b>dev_info_t</b> structure, referring to the target device instance.</td> </tr> <tr> <td><i>hba_tran</i></td> <td>Pointer to a <b>scsi_hba_tran(9S)</b> structure, consisting of the HBA's transport vectors.</td> </tr> <tr> <td><i>sd</i></td> <td>Pointer to a <b>scsi_device(9S)</b> structure, describing the target.</td> </tr> </table>	<i>hba_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the HBA device instance.	<i>tgt_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the target device instance.	<i>hba_tran</i>	Pointer to a <b>scsi_hba_tran(9S)</b> structure, consisting of the HBA's transport vectors.	<i>sd</i>	Pointer to a <b>scsi_device(9S)</b> structure, describing the target.
<i>hba_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the HBA device instance.								
<i>tgt_dip</i>	Pointer to a <b>dev_info_t</b> structure, referring to the target device instance.								
<i>hba_tran</i>	Pointer to a <b>scsi_hba_tran(9S)</b> structure, consisting of the HBA's transport vectors.								
<i>sd</i>	Pointer to a <b>scsi_device(9S)</b> structure, describing the target.								
<b>DESCRIPTION</b>	<p>The <b>tran_tgt_init()</b> vector in the <b>scsi_hba_tran(9S)</b> structure may be initialized during the HBA driver's <b>attach(9E)</b> to point to an HBA driver function to be called by the system when an instance of a target device is being created. The <b>tran_tgt_init()</b> vector, if not <b>NULL</b>, is called after the <b>dev_info</b> node structure is created for this target device instance, but before <b>probe(9E)</b> 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 <b>tran_tgt_init()</b> vector.</p> <p>Note that <i>hba_tran</i> will point to a cloned copy of the <b>scsi_hba_tran_t</b> structure allocated by the HBA driver if the <b>SCSI_HBA_TRAN_CLONE</b> flag was specified in the call to <b>scsi_hba_attach(9F)</b>. In this case, the HBA driver may choose to initialize the <i>tran_tgt_private</i> field in the structure pointed to by <i>hba_tran</i>, to point to the data specific to the particular target device instance.</p>								
<b>RETURN VALUES</b>	<p><b>tran_tgt_init()</b> must return:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><b>DDI_SUCCESS</b></td> <td>the HBA driver can support the addressed target, and was able to initialize per-target resources.</td> </tr> <tr> <td><b>DDI_FAILURE</b></td> <td>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 <b>probe(9E)</b> will not be called, and the <i>tgt_dip</i> structure destroyed.</td> </tr> </table>	<b>DDI_SUCCESS</b>	the HBA driver can support the addressed target, and was able to initialize per-target resources.	<b>DDI_FAILURE</b>	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 <b>probe(9E)</b> will not be called, and the <i>tgt_dip</i> structure destroyed.				
<b>DDI_SUCCESS</b>	the HBA driver can support the addressed target, and was able to initialize per-target resources.								
<b>DDI_FAILURE</b>	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 <b>probe(9E)</b> will not be called, and the <i>tgt_dip</i> structure destroyed.								
<b>SEE ALSO</b>	<p><b>attach(9E)</b>, <b>probe(9E)</b>, <b>tran_tgt_free(9E)</b>, <b>tran_tgt_probe(9E)</b>, <b>scsi_device(9S)</b>, <b>scsi_hba_tran(9S)</b></p> <p><i>Writing Device Drivers</i></p>								



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

<b>NAME</b>	write – write data to a device
<b>SYNOPSIS</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;sys/errno.h&gt; #include &lt;sys/open.h&gt; #include &lt;sys/cred.h&gt; #include &lt;sys/ddi.h&gt; #include &lt;sys/sunddi.h&gt;  int prefixwrite(dev_t dev, struct uio *uio_p, cred_t *cred_p);</pre>
<b>INTERFACE LEVEL ARGUMENTS</b>	<p>Architecture independent level 1 (DDI/DKI). This entry point is <b>optional</b>.</p> <p><i>dev</i> Device number.</p> <p><i>uio_p</i> Pointer to the <b>uio</b>(9S) structure that describes where the data is to be stored in user space.</p> <p><i>cred_p</i> Pointer to the user credential structure for the I/O transaction.</p>
<b>DESCRIPTION</b>	<p>Used for character or raw data I/O, the driver <b>write()</b> routine is called indirectly through <b>cb_ops</b>(9S) by the <b>write</b>(2) system call. The <b>write()</b> routine supervises the data transfer from user space to a device described by the <b>uio</b>(9S) structure.</p> <p>The <b>write</b> 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).</p>
<b>RETURN VALUES</b>	The <b>write()</b> routine should return 0 for success, or the appropriate error number.
<b>SEE ALSO</b>	<p><b>read</b>(2), <b>read</b>(9E), <b>cb_ops</b>(9S), <b>uio</b>(9S)</p> <p><i>Writing Device Drivers</i></p>

# Index

---

## A

`aread` — asynchronous read from a device, 9E-13  
asynchronous read — `aread`, 9E-13  
asynchronous write — `awrite`, 9E-16  
`awrite` — asynchronous write to a device, 9E-16

## C

character-oriented drivers  
— `ioctl`, 9E-28

## D

DDI device mapping  
  `mapdev_access` — device mapping access  
  entry point, 9E-33  
  `mapdev_dup` — device mapping duplication  
  entry point, 9E-35  
  `mapdev_free` — device mapping free entry  
  point, 9E-36  
`dev_info` structure  
  convert device number to — `getinfo`, 9E-25  
device access  
  — `close`, 9E-21  
  — `open`, 9E-41  
device mapping access entry point —  
  `mapdev_access`, 9E-33  
device mapping duplication entry point —  
  `mapdev_dup`, 9E-35  
device mapping free entry point — `mapdev_free`,

9E-36

device number  
  convert to `dev_info` structure — `getinfo`,  
  9E-25

## devices

  attach to system — `attach`, 9E-15  
  claim to drive a device — `identify`, 9E-27  
  detach from system — `detach`, 9E-23  
  read data — `read`, 9E-49  
  write data to a device — `write`, 9E-70

## devices, memory mapped

  check virtual mapping — `mmap`, 9E-37

## devices, memory mapping

  map device memory into user space — `seg-`  
  `map`, 9E-50

## devices, non-self-identifying

  determine if present — `probe`, 9E-44

## Driver entry point routines

  — `_fini`, 9E-10  
  — `_info`, 9E-10  
  — `_init`, 9E-10  
  — `attach`, 9E-15  
  — `chpoll`, 9E-18  
  — `close`, 9E-21  
  — `detach`, 9E-23  
  — `dump`, 9E-24  
  — `getinfo`, 9E-25  
  — `identify`, 9E-27

---

Driver entry point routines, *continued*

- ioctl, 9E-28
- mmap, 9E-37
- open, 9E-41
- print, 9E-43
- probe, 9E-44
- prop\_op, 9E-45
- put, 9E-47
- read, 9E-49
- segmap, 9E-50
- srv, 9E-52
- strategy, 9E-54
- write, 9E-70

driver messages

- display on system console — print, 9E-43

driver property information

- report —prop\_op, 9E-45

drivers, character-oriented

- ioctl, 9E-28

dump — dump memory to disk during system

- failure, 9E-24

dynamically update kstats — ks\_update, 9E-31

## G

get/set SCSI transport capability — tran\_getcap,

9E-57

tran\_setcap, 9E-57

## H

HBA resources

- request to free HBA resources allocated on behalf of a target —  
tran\_tgt\_free, 9E-67

- request to initialize HBA resources on behalf of a particular target —

tran\_tgt\_init, 9E-68

## I

identify — claim to drive a device, 9E-27

## K

kernel modules, dynamic loading

- initialize a loadable module — \_init, 9E-10
- prepare loadable module for unloading —

\_fini,

kernel modules, dynamic loading, *continued*

9E-10

- return loadable module information — \_info, 9E-10

ks\_update — dynamically update kstats, 9E-31

## M

mapdev\_access — device mapping access entry point, 9E-33

mapdev\_dup — device mapping duplication entry point, 9E-35

mapdev\_free — device mapping free entry point, 9E-36

memory mapping for devices

- check virtual mapping — mmap, 9E-37

- map device memory into user space — segmap, 9E-50

## N

non-self-identifying devices

- determine if present — probe, 9E-44

non-STREAMS character device driver

- poll entry point — chpoll, 9E-18

## P

put — receive messages from the preceding queue, 9E-47

## R

request to notify SCSI target of bus reset

- tran\_reset\_notify, 9E-62

reset a SCSI bus or target — tran\_reset, 9E-61

## S

SCSI bus

- request to probe SCSI bus for a particular target — tran\_tgt\_probe, 9E-69

SCSI command

- abort — tran\_abort, 9E-55

- request to transport — tran\_start, 9E-63

SCSI HBA DMA deallocation entry point —

tran\_dmafree, 9E-56

---

SCSI HBA memory synchronization entry point —  
  tran\_sync\_pkt, 9E-66

SCSI HBA packet preparation and deallocation —  
  tran\_init\_pkt, 9E-58  
  tran\_destroy\_pkt, 9E-58

strategy — perform block I/O, 9E-54

STREAMS message queues  
  receive messages from the preceding queue —  
    put, 9E-47  
  service queued messages — *srv*, 9E-52

## T

tran\_abort — abort a SCSI command, 9E-55

tran\_destroy\_pkt — SCSI HBA packet preparation and deallocation, 9E-58

tran\_dmafree — SCSI HBA DMA deallocation entry point, 9E-56

tran\_getcap — get/set SCSI transport capability, 9E-57

tran\_init\_pkt — SCSI HBA packet preparation and deallocation, 9E-58

tran\_reset — reset a SCSI bus or target, 9E-61

tran\_reset\_notify — request to notify SCSI target of bus reset, 9E-62

tran\_setcap — get/set SCSI transport capability, 9E-57

tran\_start — request to transport a SCSI command, 9E-63

tran\_sync\_pkt — SCSI HBA memory synchronization entry point, 9E-66

tran\_tgt\_free — request to free HBA resources allocated on behalf of a target, 9E-67

tran\_tgt\_init — request to initialize HBA resources on behalf of a particular target, 9E-68

tran\_tgt\_probe — request to probe SCSI bus for a particular target, 9E-69

## V

virtual address space  
  dump portion of to disk in case of system failure — *dump*, 9E-24

## W

write — write data to a device, 9E-70