



# Device Driver Tutorial

---

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

Part No: 817-5789-10  
January 2005

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

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

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

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

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

---

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

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

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

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

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



041214@10536



# Contents

---

<b>Preface</b>	<b>7</b>
<b>1 Introduction to Device Drivers</b>	<b>13</b>
Solaris Operating System Definition	13
Kernel Overview	13
Differences Between Kernel Modules and User Programs	14
User and Kernel Address Spaces on x86 and SPARC Machines	16
Device Drivers	17
Driver Directory Organization	19
Devices as Files	20
Devices Directories	21
Device Tree	21
Character and Block Devices	22
Device Names	22
Device Numbers	23
<b>2 Template Driver Example</b>	<b>25</b>
Overview of the Template Driver Example	25
Writing the Template Driver	26
Writing the Loadable Module Configuration Entry Points	27
Writing the Autoconfiguration Entry Points	31
Writing the User Context Entry Points	38
Writing the Driver Data Structures	42
Writing the Device Configuration File	47
Building and Installing the Driver	48
Compiling and Linking the Driver	48

Installing the Driver in a Temporary Location	48
Testing the Template Driver	49
Loading the Driver	49
Reading and Writing the Device	51
Unloading the Driver	51
Complete Driver Source	52
<b>3 Reading and Writing Data in Kernel Memory</b>	<b>57</b>
Displaying Data Stored in Kernel Memory	57
Writing Quote Of The Day Version 1	58
Building, Installing, and Using Quote Of The Day Version 1	59
Displaying Data on Demand	60
Writing Quote Of The Day Version 2	60
Building, Installing, and Using Quote Of The Day Version 2	69
Modifying Data Stored in Kernel Memory	71
Writing Quote Of The Day Version 3	71
Building and Installing Quote Of The Day Version 3	91
Using Quote Of The Day Version 3	91
<b>4 Tips for Writing Device Drivers</b>	<b>97</b>
Steps in the Development Process	97
Device Driver Coding Tips	98
Device Driver Testing Tips	100
Device Driver Debugging and Tuning Tips	101
<b>Index</b>	<b>103</b>

# Examples

---

<b>EXAMPLE 3-1</b>	Quote Of The Day Version 1 Source File	58
<b>EXAMPLE 3-2</b>	Quote Of The Day Version 1 Configuration File	59
<b>EXAMPLE 3-3</b>	Quote Of The Day Version 2 Source File	64
<b>EXAMPLE 3-4</b>	Quote Of The Day Version 2 Configuration File	69
<b>EXAMPLE 3-5</b>	Quote Of The Day Version 3 Source File	80
<b>EXAMPLE 3-6</b>	Quote Of The Day Version 3 Header File	90
<b>EXAMPLE 3-7</b>	Quote Of The Day Version 3 Configuration File	90
<b>EXAMPLE 3-8</b>	Quote Of The Day I/O Control Command Source File	93



# Preface

---

This *Device Driver Tutorial* is a hands-on guide that shows you how to develop a simple device driver for the Solaris™ Operating System (Solaris OS). *Device Driver Tutorial* also explains how device drivers work in the Solaris OS. This book is a companion to *Writing Device Drivers*. *Writing Device Drivers* is a thorough reference document that discusses many types of devices and drivers. *Device Driver Tutorial* examines complete drivers but does not provide a comprehensive treatment of all driver types. *Device Driver Tutorial* often points to *Writing Device Drivers* and other books for further information.

---

**Note** – This Solaris release supports systems that use the SPARC® and x86 families of processor architectures: UltraSPARC®, SPARC64, AMD64, Pentium, and Xeon EM64T. The supported systems appear in the *Solaris 10 Hardware Compatibility List* at <http://www.sun.com/bigadmin/hcl>. This document cites any implementation differences between the platform types.

In this document, the term “x86” refers to 64-bit and 32-bit systems manufactured using processors compatible with the AMD64 or Intel Xeon/Pentium product families. For supported systems, see the *Solaris 10 Hardware Compatibility List*.

---

---

## Who Should Use This Book

You should read this tutorial if you need to develop, install, and configure device drivers for the Solaris OS. You also should read this book if you need to maintain existing drivers or add new functionality to existing Solaris OS drivers. Information about the kernel provided in this book also will help you troubleshoot any problems you might encounter installing or configuring Solaris systems.

---

## User Background

To write device drivers for the Solaris OS, you should have the following background:

- Be a confident C programmer
- Have experience with data structures, especially with linked lists
- Understand bit operations
- Understand indirect function calls
- Understand caching
- Understand multithreading (see the *Multithreaded Programming Guide*)
- Be familiar with a UNIX<sup>®</sup> shell
- Understand the basics of UNIX system and I/O architecture

The most important information you need to have to write a device driver are the characteristics of the device. Get a detailed specification for the device you want to drive.

Experience with Solaris OS compilers, debuggers, and other tools will be very helpful to you. You also need to understand where the file system fits with the kernel and the application layer. These topics are discussed in this tutorial.

---

## How This Book Is Organized

This book is organized into the following chapters:

[Chapter 1](#) provides an overview of the Solaris Operating System and kernel. This chapter also discusses the driver development environment and tools.

[Chapter 2](#) shows a simple template driver. This chapter shows in detail the steps to develop, build, install, load, and test this simple driver.

[Chapter 3](#) describes how to develop a driver that reads data from and writes data to kernel memory.

[Chapter 4](#) discusses some common errors in driver development and how to avoid them or handle them. This chapter also introduces driver analysis and debugging tools.



---

## Related Books

For detailed reference information about the device driver interfaces, see the section 9 man pages. Section 9E, `Intro(9E)`, describes DDI/DKI (Device Driver Interface, Driver-Kernel Interface) driver entry points. Section 9F, `Intro(9F)`, describes DDI/DKI kernel functions. Section 9S, `Intro(9S)`, describes DDI/DKI properties and data structures.

For information on other driver-related tools and issues, see these books from Sun Microsystems:

- *Writing Device Drivers*
- *Multithreaded Programming Guide*
- *STREAMS Programming Guide*
- *Solaris Modular Debugger Guide*
- *Solaris Dynamic Tracing Guide*
- *Application Packaging Developer's Guide*
- *Solaris 64-bit Developer's Guide*

---

## Accessing Sun Documentation Online

The docs.sun.com<sup>SM</sup> Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

---

## Ordering Sun Documentation

Sun Microsystems offers select product documentation in print. For a list of documents and how to order them, see "Buy printed documentation" at `http://docs.sun.com`.

---

# Typographic Conventions

The following table describes the typographic changes that are used in this book.

**TABLE P-1** Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
<b>AaBbCc123</b>	What you type, contrasted with onscreen computer output	<code>machine_name%</code> <b>su</b> Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . Perform a <i>patch analysis</i> . Do <i>not</i> save the file. [Note that some emphasized items appear bold online.]

---

# Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

**TABLE P-2** Shell Prompts

Shell	Prompt
C shell prompt	<code>machine_name%</code>
C shell superuser prompt	<code>machine_name#</code>
Bourne shell and Korn shell prompt	<code>\$</code>

**TABLE P-2** Shell Prompts (Continued)

Shell	Prompt
Bourne shell and Korn shell superuser prompt	#



## Introduction to Device Drivers

---

This chapter gives an overview of the Solaris Operating System and kernel. To design and code device drivers, you need a good understanding of the kernel and its I/O systems.

---

### Solaris Operating System Definition

The Solaris Operating System (Solaris OS) is implemented as an executable file that runs at boot time. The Solaris OS is referred to as the *kernel*. The kernel contains all of the routines that are necessary for the system to run. Because the kernel is essential for the running of the machine, the kernel runs in a special, protected mode that is called *kernel mode*. In contrast, user-level applications operate in a restricted mode called *user mode* that has no access to kernel instructions or to the kernel address space. Device drivers run in kernel mode and are prevented from directly accessing processes in user mode.

---

### Kernel Overview

The kernel manages the system resources, including file systems, processes, and physical devices. The kernel provides applications with system services such as I/O management, virtual memory, and scheduling. The kernel coordinates interactions of all user processes and system resources. The kernel assigns priorities, services resource requests, and services hardware interrupts and exceptions. The kernel schedules and switches threads, pages memory, and swaps processes.

## Differences Between Kernel Modules and User Programs

This section discusses several important differences between kernel modules and user programs.

### Execution Differences Between Kernel Modules and User Programs

The following characteristics of kernel modules highlight important differences between the execution of kernel modules and the execution of user programs:

- **Kernel modules have separate address space.** A module runs in *kernel space*. An application runs in *user space*. System software is protected from user programs. Kernel space and user space have their own memory address spaces. See [“User and Kernel Address Spaces on x86 and SPARC Machines”](#) on page 16 for important information about address spaces.
- **Kernel modules have higher execution privilege.** Code that runs in kernel space has greater privilege than code that runs in user space. Driver modules potentially have a much greater impact on the system than user programs. Test and debug your driver modules carefully and thoroughly to avoid adverse impact on the system. See [“Device Driver Testing Tips”](#) on page 100.
- **Kernel modules do not execute sequentially.** A user program typically executes sequentially and performs a single task from beginning to end. A kernel module does not execute sequentially. A kernel module registers itself in order to serve future requests.
- **Kernel modules can be interrupted.** More than one process can request your driver at the same time. An interrupt handler can request your driver at the same time that your driver is serving a system call. In a symmetric multiprocessor (SMP) system, your driver could be executing concurrently on more than one CPU.
- **Kernel modules must be preemptable.** You cannot assume that your driver code is safe just because your driver code does not block. Design your driver assuming your driver might be preempted.
- **Kernel modules can share data.** Different threads of an application program usually do not share data. By contrast, the data structures and routines that constitute a driver are shared by all threads that use the driver. Your driver must be able to handle contention issues that result from multiple requests. Design your driver data structures carefully to keep multiple threads of execution separate. Driver code must access shared data without corrupting the data. See Chapter 3, *“Multithreading,”* in *Writing Device Drivers and Multithreaded Programming Guide*.

## Structural Differences Between Kernel Modules and User Programs

The following characteristics of kernel modules highlight important differences between the structure of kernel modules and the structure of user programs:

- **Kernel modules do not define a main program.** Kernel modules, including device drivers, have no `main()` routine. Instead, a kernel module is a collection of subroutines and data. A device driver is a kernel module that forms a software interface to an input/output (I/O) device. The subroutines in a device driver provide *entry points* to the device. The kernel uses a device number attribute to locate the `open()` routine and other routines of the correct device driver. See “[Device Drivers](#)” on page 17 for more information on entry points. See “[Device Numbers](#)” on page 23 for a description of device numbers.
- **Kernel modules are linked only to the kernel.** Kernel modules do not link in the same libraries that user programs link in. The only functions a kernel module can call are functions that are exported by the kernel. If your driver references symbols that are not defined in the kernel, your driver will compile but will fail to load. Solaris OS driver modules should use prescribed DDI/DKI (Device Driver Interface, Driver-Kernel Interface) interfaces. When you use these standard interfaces you can upgrade to a new Solaris release or migrate to a new platform without recompiling your driver. For more information on the DDI, see “[DDI/DKI Interfaces](#)” in *Writing Device Drivers*. Kernel modules can depend on other kernel modules by using the `-N` option during link editing. See the `ld(1)` man page for more information.
- **Kernel modules use different header files.** Kernel modules require a different set of header files than user programs require. The required header files are listed in the man page for each function. See *man pages section 9: DDI and DKI Kernel Functions* for DDI/DKI functions, *man pages section 9: DDI and DKI Driver Entry Points* for entry points, and *man pages section 9: DDI and DKI Properties and Data Structures* for structures. Kernel modules can include header files that are shared by user programs if the user and kernel interfaces within such shared header files are defined conditionally using the `_KERNEL` macro.
- **Kernel modules should avoid global variables.** Avoiding global variables in kernel modules is even more important than avoiding global variables in user programs. As much as possible, declare symbols as `static`. When you must use global symbols, give them a prefix that is unique within the kernel. Using this prefix for private symbols within the module also is a good practice.
- **Kernel modules can be customized for hardware.** Kernel modules can dedicate process registers to specific roles. Kernel code can be optimized for a specific processor.
- **Kernel modules can be dynamically loaded.** The collection of subroutines and data that constitute a device driver can be compiled into a single loadable module of object code. This loadable module can then be statically or dynamically linked into the kernel and unlinked from the kernel. You can add functionality to the kernel while the system is up and running. You can test new versions of your

driver without rebooting your system.

## Data Transfer Differences Between Kernel Modules and User Programs

Data transfer between a device and the system typically is slower than data transfer within the CPU. Therefore, a driver typically suspends execution of the calling thread until the data transfer is complete. While the thread that called the driver is suspended, the CPU is free to execute other threads. When the data transfer is complete, the device sends an interrupt. The driver handles the interrupt that the driver receives from the device. The driver then tells the CPU to resume execution of the calling thread. See Chapter 8, “Interrupt Handlers,” in *Writing Device Drivers*.

Drivers must work with user process (virtual) addresses, system (kernel) addresses, and I/O bus addresses. Drivers sometimes copy data from one address space to another address space and sometimes just manipulate address-mapping tables. See “Bus Architectures” in *Writing Device Drivers*.

## User and Kernel Address Spaces on x86 and SPARC Machines

On Solaris systems on x86 machines, drivers can directly access user address space.

On SPARC machines, the system panics when a kernel module attempts to directly access user address space. You must make sure your driver does not attempt to directly access user address space on a SPARC machine.



---

**Caution** – A driver that works on an x86 machine might not work on a SPARC machine because the driver might access an invalid address.

---

Do not access user data directly. Use the `ddi_copyin(9F)` and `ddi_copyout(9F)` routines to transfer data to and from user address space. You must use these two routines to transfer user data in your drivers for SPARC machines. If you also use these two routines to transfer user data in your drivers for x86 machines, then you can more easily port the driver to SPARC machines. “[Modifying Data Stored in Kernel Memory](#)” on page 71 shows an example driver that uses `ddi_copyin(9F)` and `ddi_copyout(9F)`.


The `mmap(2)` system call maps pages of memory between a process’s address space and a file or shared memory object. In response to an `mmap(2)` system call, the system calls the `devmap(9E)` entry point to map device memory into user space. This information is then available for direct access by user applications.



## Device Drivers

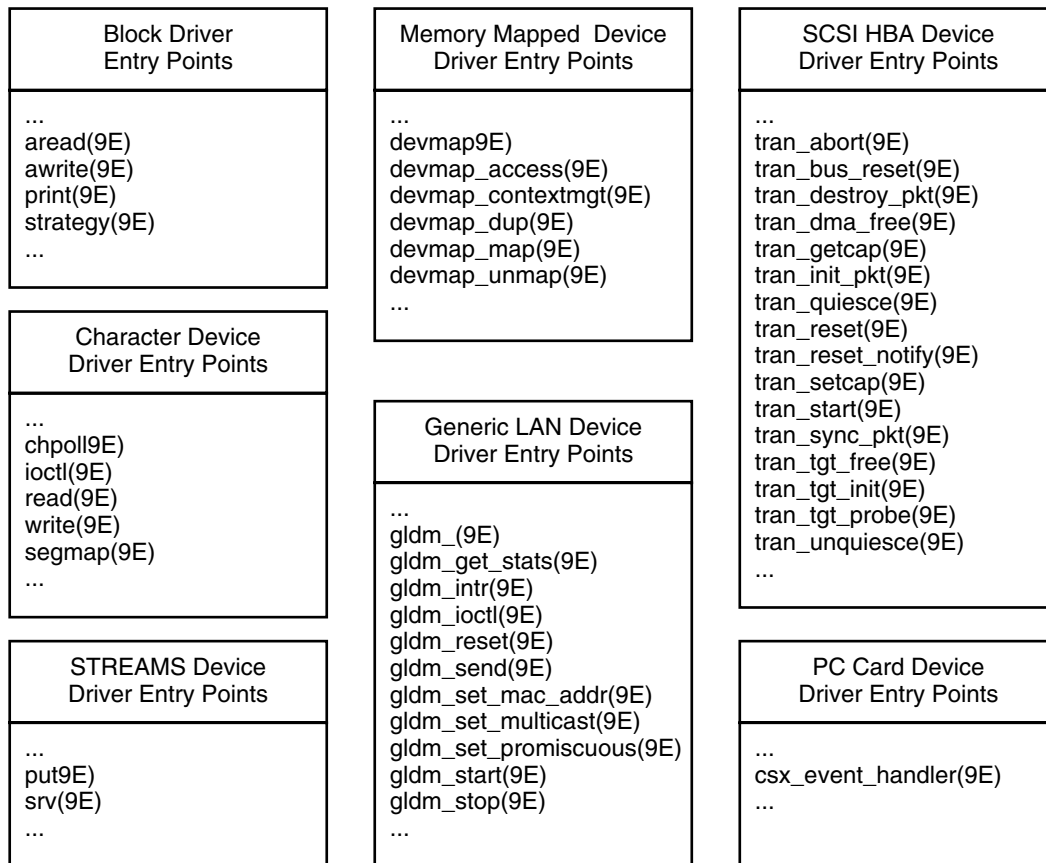
A device driver is a loadable kernel module that manages data transfers between a device and the OS. Loadable modules are loaded at boot time or by request and are unloaded by request. A device driver is a collection of C routines and data structures that can be accessed by other kernel modules. These routines must use standard interfaces called *entry points*. Through the use of entry points, the calling modules are shielded from the internal details of the driver. See “Device Driver Entry Points” in *Writing Device Drivers* for more information on entry points.

A device driver declares its general entry points in its `dev_ops(9S)` structure. A driver declares entry points for routines that are related to character or block data in its `cb_ops(9S)` structure. Some entry points and structures that are common to most drivers are shown in the following diagram.



**FIGURE 1-1** Typical Device Driver Entry Points

The Solaris OS provides many driver entry points. Different types of devices require different entry points in the driver. The following diagram shows some of the available entry points, grouped by driver type. No single device driver would use all the entry points shown in the diagram.



**FIGURE 1-2** Entry Points for Different Types of Drivers

In the Solaris OS, drivers can manage physical devices, such as disk drives, or software (pseudo) devices, such as bus nexus devices or ramdisk devices. In the case of hardware devices, the device driver communicates with the hardware controller that manages the device. The device driver shields the user application layer from the details of a specific device so that application level or system calls can be generic or device independent.

Drivers are accessed in the following situations:

- **System initialization.** The kernel calls device drivers during system initialization to determine which devices are available and to initialize those devices.
- **System calls from user processes.** The kernel calls a device driver to perform I/O operations on the device such as `open(2)`, `read(2)`, and `ioctl(2)`.
- **User-level requests.** The kernel calls device drivers to service requests from commands such as `prtconf(1M)`.

- **Device interrupts.** The kernel calls a device driver to handle interrupts generated by a device.
- **Bus reset.** The kernel calls a device driver to re-initialize the driver, the device, or both when the bus is reset. The bus is the path from the CPU to the device.

The following diagram illustrates how a device driver interacts with the rest of the system.

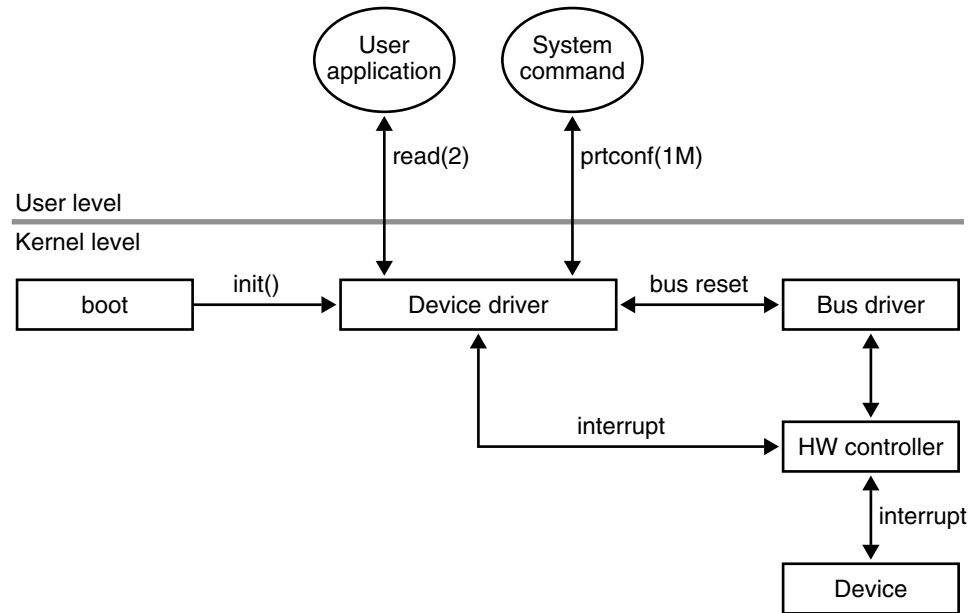


FIGURE 1-3 Typical Device Driver Interactions

## Driver Directory Organization

Device drivers and other kernel modules are organized into the following directories in the Solaris OS. See the `kernel(1M)` and `system(4)` man pages for more information about kernel organization and how to add directories to your kernel module search path.

`/kernel`

These modules are common across most platforms. Modules that are required for booting or for system initialization belong in this directory.

`/platform/`uname -i`/kernel`

These modules are specific to the platform identified by the command `uname -i`.

<code>/platform/`uname -m`/kernel</code>	These modules are specific to the platform identified by the command <code>uname -m</code> . These modules are specific to a hardware class but more generic than modules in the <code>uname -i kernel</code> directory.
<code>/usr/kernel</code>	These are user modules. Modules that are not essential to booting belong in this directory. This tutorial instructs you to put all your drivers in the <code>/usr/kernel</code> directory.

One benefit of organizing drivers into different directories is that you can selectively load different groups of drivers on startup when you boot interactively at the boot prompt as shown in the following example. See the `boot(1M)` man page for more information.

```
Type    b [file-name] [boot-flags] <ENTER>    to boot with options
or      i <ENTER>                             to enter boot interpreter
or      <ENTER>                               to boot with defaults
```

<<< timeout in 5 seconds >>>

```
Select (b)oot or (i)nterpreter: b -a
bootpath: /pci@0,0/pci8086,2545@3/pci8086,
Enter default directory for modules [/platform/i86pc/kernel /kernel
/usr/kernel]: /platform/i86pc/kernel /kernel
```

In this example, the `/usr/kernel` location is omitted from the list of directories to search for modules to load. You might want to do this if you have a driver in `/usr/kernel` that causes the kernel to panic during startup or on attach. Instead of omitting all `/usr/kernel` modules, a better method for testing drivers is to put them in their own directory. Use the `moddir` kernel variable to add this test directory to your kernel modules search path. The `moddir` kernel variable is described in `kernel(1M)` and `system(4)`. Another method for working with drivers that may have startup problems is described in [“Installing the Driver in a Temporary Location”](#) on page 48.

---

## Devices as Files

In UNIX, almost everything can be treated as a file. UNIX user applications access devices as if the devices were files. Files that represent devices are called *special files* or *device nodes*. Device special files are divided into two classes: *block* devices and *character* devices. See [“Character and Block Devices”](#) on page 22 for more information.

Every I/O service request initially refers to a named file. Most I/O operations that read or write data perform equally well on ordinary or special files. For example, the same `read(2)` system call reads bytes from a file created with a text editor and reads bytes from a terminal device.

Control signals also are handled as files. Use the `ioctl(9E)` function to manipulate control signals.

## Devices Directories

The Solaris OS includes both `/dev` and `/devices` directories for device drivers. Almost all the drivers in the `/dev` directory are links to the `/devices` directory. The `/dev` directory is UNIX standard. The `/devices` directory is specific to the Solaris OS.

By convention, file names in the `/dev` directory are more readable. For example, the `/dev` directory might contain files with names such as `kdb` and `mouse` that are links to files such as `/devices/pseudo/conskbd@0:kdb` and `/devices/pseudo/consms@0:mouse`. The `prtconf(1M)` command shows device names that are very similar to the file names in the `/devices` directory. In the following example, only selected output of the command is shown.

```
% prtconf -P
      conskbd, instance #0
      consms, instance #0
```

Entries in the `/dev` directory that are not links to the `/devices` directory are device nodes or special files created by `mknod(1M)` or `mknod(2)`. These are zero-length files that just have a major number and minor number attached to them. Linking to the physical name of the device in the `/devices` directory is preferred to using `mknod(1M)`.

Prior to the Solaris 10 OS, `/devices` was an on-disk filesystem composed of subdirectories and files. In the Solaris 10 OS, `/devices` is a virtual filesystem that creates these subdirectories and special files on demand.

For more information about the devices file system, see the `devfs(7FS)` man page.

## Device Tree

The device files in the `/devices` directory are also called the *device tree*.

The device tree shows relationships among devices. In the device tree, a directory represents a *nexus* device. A nexus is a device that can be a parent of other devices. In the following example, `pci@1f,0` is a nexus device. Only selected output from the command is shown.

```
# ls -l /devices
drwxr-xr-x  4 root    sys           512 date time pci@1f,0/
crw-----  1 root    sys       111,255 date time pci@1f,0:devctl
```

You can use `prtconf(1M)` or `prtpicl(1M)` to see a graphic representation of the device tree. See “Overview of the Device Tree” in *Writing Device Drivers* for more information about the device tree.

## Character and Block Devices

A file in the device tree that is not a directory represents either a *character* device or a *block* device.

A block device can contain addressable, reusable data. An example of a block device is a file system. Any device can be a character device. Most block devices also have character interfaces. Disks have both block and character interfaces. In your `/devices/pseudo` directory, you might find devices such as the following:

```
brw-r-----  1 root    sys           85,  0 Nov  3 09:43 md@0:0,0,blk
crw-r-----  1 root    sys           85,  0 Nov  3 09:43 md@0:0,0,raw
brw-r-----  1 root    sys           85,  1 Nov  3 09:43 md@0:0,1,blk
crw-r-----  1 root    sys           85,  1 Nov  3 09:43 md@0:0,1,raw
brw-r-----  1 root    sys           85,  2 Nov  3 09:43 md@0:0,2,blk
crw-r-----  1 root    sys           85,  2 Nov  3 09:43 md@0:0,2,raw
```

Block devices have a `b` as the first character of their file mode. Character devices have a `c` as the first character of their file mode. In this example, the block devices have `blk` in their names and the character devices have `raw` in their names.

The `md(7D)` device is a metadvice that provides disk services. The block devices access the disk using the system’s normal buffering mechanism. The character devices provide for direct transmission between the disk and the user’s read or write buffer.

## Device Names

This section shows a complex device name and explains the meaning of each part of the name in `/dev` and also in `/devices`. The following example is the name of a disk slice:

```
/dev/dsk/c0t0d0s7 -> ../../devices/pci@1c,600000/scsi@2/sd@0,0:h
```

First, examine the name of the file in the `/dev` directory. These names are managed by the `devfsadm(1M)` daemon.

- `c0` Controller 0
- `t0` Target 0. On SCSI controllers, this value is the disk number.
- `d0` SCSI LUN. This value indicates a virtual partitioning of a target or single physical device.
- `s7` Slice 7 on the target 0 disk.

For the same device, compare the name of the file in the `/devices` directory. These names show the physical structure and real device names. Note that some of the components of the device name in the `/devices` directory are subdirectories.

- `pci@1c,600000` PCI bus at address `1c,600000`. These addresses are meaningful only to the parent device.
- `scsi@2` SCSI controller at address 2 on the PCI bus at address `1c,600000`. This name corresponds to the `c0` in `/dev/dsk/c0t0d0s7`.
- `sd@0,0` SCSI disk at address `0,0` on the SCSI controller at address 2. This name corresponds to the `t0` in `/dev/dsk/c0t0d0s7`. The `sd` name and driver can also apply to IDE CD-ROM devices.
- `sd@0,0:h` Minor node `h` on the SCSI disk at address `0,0`. This name corresponds to the `s7` in `/dev/dsk/c0t0d0s7`.

## Device Numbers

A *device number* identifies a particular device and minor node in the device tree. The `dev_t` parameter that is required in many DDI/DKI routines is this device number.

Each device has a major number and a minor number. A device number is a *major,minor* pair. A long file listing shows the device number in the column where file sizes are usually listed. In the following example, the device number is 86,255. The device major number is 86, and the device minor number is 255.

```
% ls -l /devices/pci@0,0:devctl
crw----- 1 root    sys      86,255 date time /devices/pci@0,0:devctl
```

In the Solaris OS, the major number is chosen for you when you install the driver so that it will not conflict with any other major number. The kernel uses the major number to associate the I/O request with the correct driver code. The kernel uses this association to decide which driver to execute when the user reads or writes the device file. All devices and their major numbers are listed in the file `/etc/name_to_major`.

```
% grep 86 /etc/name_to_major
pci 86
```

The minor number is assigned in the driver. The minor number must map each driver to a specific device instance. Minor numbers usually refer to sub-devices. For example, a disk driver may communicate with a hardware controller device that has several disk drives attached. Minor nodes do not necessarily have a physical representation.

The following example shows instances 0, 1, and 2 of the md device. The numbers 0, 1, and 2 are the minor numbers.

```
brw-r----- 1 root    sys      85,  0 Nov  3 09:43 md@0:0,0,blk
crw-r----- 1 root    sys      85,  0 Nov  3 09:43 md@0:0,0,raw
brw-r----- 1 root    sys      85,  1 Nov  3 09:43 md@0:0,1,blk
crw-r----- 1 root    sys      85,  1 Nov  3 09:43 md@0:0,1,raw
brw-r----- 1 root    sys      85,  2 Nov  3 09:43 md@0:0,2,blk
crw-r----- 1 root    sys      85,  2 Nov  3 09:43 md@0:0,2,raw
```

In the name `sd@0,0:h`, `h` represents a minor node. When the driver receives a request for minor node `h`, the driver actually receives a corresponding minor number. The driver for the `sd` node interprets that minor number to be a particular section of disk, such as slice 7 mounted on `/export`.

[Chapter 2](#) shows how to use the `ddi_get_instance(9F)` routine in your driver to get an instance number for the device you are driving.



## Template Driver Example

---

This chapter shows you how to develop a very simple, working driver. This chapter explains how to write the driver and configuration file, compile the driver, load the driver, and test the driver.

The driver that is shown in this chapter is a pseudo device driver that merely writes a message to a system log every time an entry point is entered. This driver demonstrates the minimum functionality that any character driver must implement. You can use this driver as a template for building a complex driver.

This chapter discusses the following driver development steps:

- “Overview of the Template Driver Example” on page 25
- “Writing the Template Driver” on page 26
- “Writing the Device Configuration File” on page 47
- “Building and Installing the Driver” on page 48
- “Testing the Template Driver” on page 49
- “Complete Driver Source” on page 52

---

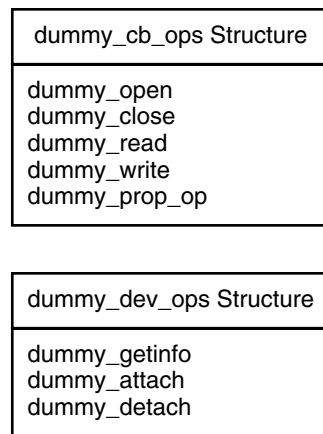
## Overview of the Template Driver Example

This example guides you through the following steps:

1. Create a directory where you can develop your driver and open a new text file named `dummy.c`.
2. Write the entry points for loadable module configuration: `_init(9E)`, `_info(9E)`, and `_fini(9E)`.
3. Write the entry points for autoconfiguration: `attach(9E)`, `detach(9E)`, `getinfo(9E)`, and `prop_op(9E)`.

4. Write the entry points for user context: `open(9E)`, `close(9E)`, `read(9E)`, and `write(9E)`.
5. Define the data structures: the character and block operations structure `cb_ops(9S)`, the device operations structure `dev_ops(9S)`, and the module linkage structures `modldrv(9S)` and `modlinkage(9S)`.
6. Create the driver configuration file `dummy.conf`.
7. Build and install the driver.
8. Test the driver by loading the driver, reading from and writing to the device node, and unloading the driver.

The entry points that are to be created in this example are shown in the following diagram.



**FIGURE 2-1** Entry Points for the dummy Example

---

## Writing the Template Driver

This section describes the entry points and data structures that are included in this driver and shows you how to define them. All of these data structures and almost all of these entry points are required for any character device driver.

This section describes the following entry points and data structures:

- Loadable module configuration entry points
- Autoconfiguration entry points
- User context entry points
- Character and block operations structure

- Device operations structure
- Module linkage structures

First, create a directory where you can develop your driver. This driver is named `dummy` because this driver does not do any real work. Next, open a new text file named `dummy.c`.

## Writing the Loadable Module Configuration Entry Points

Every kernel module of any type must define at least the following three loadable module configuration entry points:

- The `_init(9E)` routine initializes a loadable module. The `_init(9E)` routine must at least call the `mod_install(9F)` function and return the success or failure value that is returned by `mod_install(9F)`.
- The `_info(9E)` routine returns information about a loadable module. The `_info(9E)` routine must at least call the `mod_info(9F)` function and return the value that is returned by `mod_info(9F)`.
- The `_fini(9E)` routine prepares a loadable module for unloading. The `_fini(9E)` routine must at least call the `mod_remove(9F)` function and return the success or failure value that is returned by `mod_remove(9F)`. When `mod_remove(9F)` is successful, the `_fini(9E)` routine must undo everything that the `_init(9E)` routine did.

The `mod_install(9F)`, `mod_info(9F)`, and `mod_remove(9F)` functions are used in exactly the same way in every driver, regardless of the functionality of the driver. You do not need to investigate what the values of the arguments of these functions should be. You can copy these function calls from this example and paste them into every driver you write.

In this section, the following code is added to the `dummy.c` source file:

```
/* Loadable module configuration entry points */
int
_init(void)
{
    cmn_err(CE_NOTE, "Inside _init");
    return(mod_install(&ml));
}

int
_info(struct modinfo *modinfop)
{
    cmn_err(CE_NOTE, "Inside _info");
    return(mod_info(&ml, modinfop));
}
```

```

int
_fini(void)
{
    cmn_err(CE_NOTE, "Inside _fini");
    return(mod_remove(&ml));
}

```

## Declaring the Loadable Module Configuration Entry Points

The `_init(9E)`, `_info(9E)`, and `_fini(9E)` routine names are not unique to any particular kernel module. You customize the behavior of these routines when you define them in your module, but the names of these routines are not unique. These three routines are declared in the `modctl.h` header file. You need to include the `modctl.h` header file in your `dummy.c` file. Do not declare these three routines in `dummy.c`.

## Defining the Module Initialization Entry Point

The `_init(9E)` routine returns type `int` and takes no arguments. The `_init(9E)` routine must call the `mod_install(9F)` function and return the success or failure value that is returned by `mod_install(9F)`.

The `mod_install(9F)` function takes an argument that is a `modlinkage(9S)` structure. See [“Defining the Module Linkage Structures” on page 46](#) for information about the `modlinkage(9S)` structure.

This driver is supposed to write a message each time an entry point is entered. Use the `cmn_err(9F)` function to write a message to a system log. The `cmn_err(9F)` function usually is used to report an error condition. The `cmn_err(9F)` function also is useful for debugging in the same way that you might use print statements in a user program.

The `cmn_err(9F)` function requires you to include the `cmn_err.h` header file, the `ddi.h` header file, and the `sunddi.h` header file. The `cmn_err(9F)` function takes two arguments. The first argument is a constant that indicates the severity of the error message. The message written by this driver is not an error message but is simply a test message. Use `CE_NOTE` for the value of this severity constant. The second argument the `cmn_err(9F)` function takes is a string message.

The following code is the `_init(9E)` routine that you should enter into your `dummy.c` file. The `ml` structure is the `modlinkage(9S)` structure that is discussed in [“Defining the Module Linkage Structures” on page 46](#).

```

int
_init(void)
{
    cmn_err(CE_NOTE, "Inside _init");
}

```

```

        return(mod_install(&ml));
    }

```

## Defining the Module Information Entry Point

The `_info(9E)` routine returns type `int` and takes an argument that is a pointer to an opaque `modinfo` structure. The `_info(9E)` routine must return the value that is returned by the `mod_info(9F)` function.

The `mod_info(9F)` function takes two arguments. The first argument to `mod_info(9F)` is a `modlinkage(9S)` structure. See “[Defining the Module Linkage Structures](#)” on page 46 for information about the `modlinkage(9S)` structure. The second argument to `mod_info(9F)` is the same `modinfo` structure pointer that is the argument to the `_info(9E)` routine. The `mod_info(9F)` function returns the module information or returns zero if an error occurs.

Use the `cmn_err(9F)` function to write a message to the system log in the same way that you used the `cmn_err(9F)` function in your `_init(9E)` entry point.

The following code is the `_info(9E)` routine that you should enter into your `dummy.c` file. The `ml` structure is discussed in “[Defining the Module Linkage Structures](#)” on page 46. The `modinfo` argument is a pointer to an opaque structure that the system uses to pass module information.

```

int
_info(struct modinfo *modinfo)
{
    cmn_err(CE_NOTE, "Inside _info");
    return(mod_info(&ml, modinfo));
}

```

## Defining the Module Unload Entry Point

The `_fini(9E)` routine returns type `int` and takes no arguments. The `_fini(9E)` routine must call the `mod_remove(9F)` function and return the success or failure value that is returned by `mod_remove(9F)`.

When `mod_remove(9F)` is successful, the `_fini(9E)` routine must undo everything that the `_init(9E)` routine did. The `_fini(9E)` routine must call `mod_remove(9F)` because the `_init(9E)` routine called `mod_install(9F)`. The `_fini(9E)` routine must deallocate anything that was allocated, close anything that was opened, and destroy anything that was created in the `_init(9E)` routine.

The `_fini(9E)` routine can be called at any time when a module is loaded. In normal operation, the `_fini(9E)` routine often fails. This behavior is normal because the kernel allows the module to determine whether the module can be unloaded. If `mod_remove(9F)` is successful, the module determines that devices were detached, and the module can be unloaded. If `mod_remove(9F)` fails, the module determines that devices were not detached, and the module cannot be unloaded.

The following actions take place when `mod_remove(9F)` is called:

- The kernel checks whether this driver is busy. This driver is busy if one of the following conditions is true:
  - A device node that is managed by this driver is open.
  - Another module that depends on this driver is open. A module depends on this driver if the module was linked using the `-N` option with this driver named as the argument to that `-N` option. See the `ld(1)` man page for more information.
- If the driver is busy, then `mod_remove(9F)` fails and `_fini(9E)` fails.
- If the driver is not busy, then the kernel calls the `detach(9E)` entry point of the driver.
  - If `detach(9E)` fails, then `mod_remove(9F)` fails and `_fini(9E)` fails.
  - If `detach(9E)` succeeds, then `mod_remove(9F)` succeeds, and `_fini(9E)` continues its cleanup work.

The `mod_remove(9F)` function takes an argument that is a `modlinkage(9S)` structure. See “[Defining the Module Linkage Structures](#)” on page 46 for information about the `modlinkage(9S)` structure.

Use the `cmn_err(9F)` function to write a message to the system log in the same way that you used the `cmn_err(9F)` function in your `_init(9E)` entry point.

The following code is the `_fini(9E)` routine that you should enter into your `dummy.c` file. The `ml` structure is discussed in “[Defining the Module Linkage Structures](#)” on page 46.

```
int
_fini(void)
{
    cmn_err(CE_NOTE, "Inside _fini");
    return(mod_remove(&ml));
}
```

## Including Loadable Module Configuration Header Files

The `_init(9E)`, `_info(9E)`, `_fini(9E)`, and `mod_install(9F)` functions require you to include the `modctl.h` header file. The `cmn_err(9F)` function requires you to include the `cmn_err.h` header file, the `ddi.h` header file, and the `sunddi.h` header file.

The following header files are required by the three loadable module configuration routines that you have written in this section. Include this code near the top of your `dummy.c` file.

```
#include <sys/modctl.h> /* used by _init, _info, _fini */
#include <sys/cmn_err.h> /* used by all entry points for this driver */
#include <sys/ddi.h> /* used by all entry points for this driver */
#include <sys/sunddi.h> /* used by all entry points for this driver */
```

## Writing the Autoconfiguration Entry Points

Every character driver must define at least the following autoconfiguration entry points. The kernel calls these routines when the device driver is loaded.

- The `attach(9E)` routine must call `ddi_create_minor_node(9F)`. The `ddi_create_minor_node(9F)` function provides the information the system needs to create the device files.
- The `detach(9E)` routine must call `ddi_remove_minor_node(9F)` to deallocate everything that was allocated by `ddi_create_minor_node(9F)`. The `detach(9E)` routine must undo everything that the `attach(9E)` routine did.
- The `getinfo(9E)` routine returns requested device driver information through one of its arguments.
- The `prop_op(9E)` routine returns requested device driver property information through a pointer. You can call the `ddi_prop_op(9F)` function instead of writing your own `prop_op(9E)` entry point. Use the `prop_op(9E)` entry point to customize the behavior of the `ddi_prop_op(9F)` function.

In this section, the following code is added:

```
/* Device autoconfiguration entry points */
static int
dummy_attach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    cmn_err(CE_NOTE, "Inside dummy_attach");
    switch(cmd) {
    case DDI_ATTACH:
        dummy_dip = dip;
        if (ddi_create_minor_node(dip, "0", S_IFCHR,
            ddi_get_instance(dip), DDI_PSEUDO, 0)
            != DDI_SUCCESS) {
            cmn_err(CE_NOTE,
                "%s%d: attach: could not add character node.",
                "dummy", 0);
            return(DDI_FAILURE);
        } else
            return DDI_SUCCESS;
    default:
        return DDI_FAILURE;
    }
}

static int
dummy_detach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    cmn_err(CE_NOTE, "Inside dummy_detach");
    switch(cmd) {
    case DDI_DETACH:
        dummy_dip = 0;
        ddi_remove_minor_node(dip, NULL);
        return DDI_SUCCESS;
    default:

```

```

        return DDI_FAILURE;
    }
}

static int
dummy_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg,
              void **resultp)
{
    cmn_err(CE_NOTE, "Inside dummy_getinfo");
    switch(cmd) {
    case DDI_INFO_DEVT2DEVINFO:
        *resultp = dummy_dip;
        return DDI_SUCCESS;
    case DDI_INFO_DEVT2INSTANCE:
        *resultp = 0;
        return DDI_SUCCESS;
    default:
        return DDI_FAILURE;
    }
}

static int
dummy_prop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op,
              int flags, char *name, caddr_t valuep, int *lengthp)
{
    cmn_err(CE_NOTE, "Inside dummy_prop_op");
    return(ddi_prop_op(dev, dip, prop_op, flags, name, valuep, lengthp));
}

```

## Declaring the Autoconfiguration Entry Points

The `attach(9E)`, `detach(9E)`, `getinfo(9E)`, and `prop_op(9E)` entry point routines need to be uniquely named for this driver. Choose a prefix to use with each entry point routine.

---

**Note** – By convention, the prefix used for function and data names that are unique to this driver is either the name of this driver or an abbreviation of the name of this driver. Use the same prefix throughout the driver. This practice makes debugging much easier.

---

In the example shown in this chapter, `dummy_` is used for the prefix to each function and data name that is unique to this example.

The following declarations are the autoconfiguration entry point declarations you should have in your `dummy.c` file. Note that each of these functions is declared `static`.

```

static int dummy_attach(dev_info_t *dip, ddi_attach_cmd_t cmd);
static int dummy_detach(dev_info_t *dip, ddi_detach_cmd_t cmd);
static int dummy_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg,

```



```

void **resultp);
static int dummy_prop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op,
    int flags, char *name, caddr_t valuep, int *lengthp);

```

## Defining the Device Attach Entry Point

The `attach(9E)` routine returns type `int`. The `attach(9E)` routine must return either `DDI_SUCCESS` or `DDI_FAILURE`. These two constants are defined in `sunddi.h`. All of the autoconfiguration entry point routines except for `prop_op(9E)` return either `DDI_SUCCESS` or `DDI_FAILURE`.

The `attach(9E)` routine takes two arguments. The first argument is a pointer to the `dev_info` structure for this driver. All of the autoconfiguration entry point routines take a `dev_info` argument. The second argument is a constant that specifies the attach type. The value that is passed through this second argument is either `DDI_ATTACH` or `DDI_RESUME`. Every `attach(9E)` routine must define behavior for at least `DDI_ATTACH`.

The `DDI_ATTACH` code must initialize a device instance. In a realistic driver, you define and manage multiple instances of the driver by using a state structure and the `ddi_soft_state(9F)` functions. Each instance of the driver has its own copy of the state structure that holds data specific to that instance. One of the pieces of data that is specific to each instance is the device instance pointer. Each instance of the device driver is represented by a separate device file in `/devices`. Each device instance file is pointed to by a separate device instance pointer. See [“Managing Device State” on page 61](#) for information about state structures and `ddi_soft_state(9F)` functions. See [“Devices as Files” on page 20](#) for information about device files and instances.

This dummy driver allows only one instance. Because this driver allows only one instance, this driver does not use a state structure. This driver still must declare a device instance pointer and initialize the pointer value in the `attach(9E)` routine. Enter the following code near the beginning of `dummy.c` to declare a device instance pointer for this driver:

```

dev_info_t *dummy_dip; /* keep track of one instance */

```

The following code is the `dummy_attach()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in [“Declaring the Autoconfiguration Entry Points” on page 32](#).

```

static int
dummy_attach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    cmn_err(CE_NOTE, "Inside dummy_attach");
    switch(cmd) {
    case DDI_ATTACH:
        dummy_dip = dip;
        if (ddi_create_minor_node(dip, "0", S_IFCHR,

```

```

        ddi_get_instance(dip), DDI_PSEUDO, 0)
    != DDI_SUCCESS) {
    cmn_err(CE_NOTE,
        "%s%d: attach: could not add character node.",
        "dummy", 0);
    return(DDI_FAILURE);
} else
    return DDI_SUCCESS;
default:
    return DDI_FAILURE;
}
}
}

```

First, use `cmn_err(9F)` to write a message to the system log, as you did in your `_init(9E)` entry point. Then provide `DDI_ATTACH` behavior. Within the `DDI_ATTACH` code, first assign the device instance pointer from the `dummy_attach()` argument to the `dummy_dip` variable that you declared above. You need to save this pointer value in the global variable so that you can use this pointer to get information about this instance from `dummy_getinfo()` and detach this instance in `dummy_detach()`. In this `dummy_attach()` routine, the device instance pointer is used by the `ddi_get_instance(9F)` function to return the instance number. The device instance pointer and the instance number both are used by `ddi_create_minor_node(9F)` to create a new device node.

A realistic driver probably would use the `ddi_soft_state(9F)` functions to create and manage a device node. This dummy driver uses the `ddi_create_minor_node(9F)` function to create a device node. The `ddi_create_minor_node(9F)` function takes six arguments. The first argument to the `ddi_create_minor_node(9F)` function is the device instance pointer that points to the `dev_info` structure of this device. The second argument is the name of this minor node. The third argument is `S_IFCHR` if this device is a character minor device or is `S_IFBLK` if this device is a block minor device. This dummy driver is a character driver.

The fourth argument to the `ddi_create_minor_node(9F)` function is the minor number of this minor device. This number is also called the instance number. The `ddi_get_instance(9F)` function returns this instance number. The fifth argument to the `ddi_create_minor_node(9F)` function is the node type. The `ddi_create_minor_node(9F)` man page lists the possible node types. The `DDI_PSEUDO` node type is for pseudo devices. The sixth argument to the `ddi_create_minor_node(9F)` function specifies whether this is a clone device. This is not a clone device, so set this argument value to 0.

If the `ddi_create_minor_node(9F)` call is not successful, write a message to the system log and return `DDI_FAILURE`. If the `ddi_create_minor_node(9F)` call is successful, return `DDI_SUCCESS`. If this `dummy_attach()` routine receives any `cmd` other than `DDI_ATTACH`, return `DDI_FAILURE`.

## Defining the Device Detach Entry Point

The `detach(9E)` routine takes two arguments. The first argument is a pointer to the `dev_info` structure for this driver. The second argument is a constant that specifies the detach type. The value that is passed through this second argument is either `DDI_DETACH` or `DDI_SUSPEND`. Every `detach(9E)` routine must define behavior for at least `DDI_DETACH`.

The `DDI_DETACH` code must undo everything that the `DDI_ATTACH` code did. In the `DDI_ATTACH` code in your `attach(9E)` routine, you saved the address of a new `dev_info` structure and you called the `ddi_create_minor_node(9F)` function to create a new node. In the `DDI_DETACH` code in this `detach(9E)` routine, you need to reset the variable that pointed to the `dev_info` structure for this node. You also need to call the `ddi_remove_minor_node(9F)` function to remove this node. The `detach(9E)` routine must deallocate anything that was allocated, close anything that was opened, and destroy anything that was created in the `attach(9E)` routine.

The following code is the `dummy_detach()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in [“Declaring the Autoconfiguration Entry Points”](#) on page 32.

```
static int
dummy_detach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    cmn_err(CE_NOTE, "Inside dummy_detach");
    switch(cmd) {
        case DDI_DETACH:
            dummy_dip = 0;
            ddi_remove_minor_node(dip, NULL);
            return DDI_SUCCESS;
        default:
            return DDI_FAILURE;
    }
}
```

First, use `cmn_err(9F)` to write a message to the system log, as you did in your `_init(9E)` entry point. Then provide `DDI_DETACH` behavior. Within the `DDI_DETACH` code, first reset the `dummy_dip` variable that you set in `dummy_attach()` above. You cannot reset this device instance pointer unless you remove all instances of the device. This dummy driver supports only one instance.

Next, call the `ddi_remove_minor_node(9F)` function to remove this device node. The `ddi_remove_minor_node(9F)` function takes two arguments. The first argument is the device instance pointer that points to the `dev_info` structure of this device. The second argument is the name of the minor node you want to remove. If the value of the minor node argument is `NULL`, then `ddi_remove_minor_node(9F)` removes all instances of this device. Because the `DDI_DETACH` code of this driver always removes all instances, this dummy driver supports only one instance.

If the value of the `cmd` argument to this `dummy_detach()` routine is `DDI_DETACH`, remove all instances of this device and return `DDI_SUCCESS`. If this `dummy_detach()` routine receives any `cmd` other than `DDI_DETACH`, return `DDI_FAILURE`.

## Defining the Get Driver Information Entry Point

The `getinfo(9E)` routine takes a pointer to a device number and returns a pointer to a device information structure or returns a device instance number. The return value of the `getinfo(9E)` routine is `DDI_SUCCESS` or `DDI_FAILURE`. The pointer or instance number requested from the `getinfo(9E)` routine is returned through a pointer argument.

The `getinfo(9E)` routine takes four arguments. The first argument is a pointer to the `dev_info` structure for this driver. This `dev_info` structure argument is obsolete and is no longer used by the `getinfo(9E)` routine.

The second argument to the `getinfo(9E)` routine is a constant that specifies what information the `getinfo(9E)` routine must return. The value of this second argument is either `DDI_INFO_DEVT2DEVINFO` or `DDI_INFO_DEVT2INSTANCE`. The third argument to the `getinfo(9E)` routine is a pointer to a device number. The fourth argument is a pointer to the place where the `getinfo(9E)` routine must store the requested information. The information stored at this location depends on the value you passed in the second argument to the `getinfo(9E)` routine.

The following table describes the relationship between the second and fourth arguments to the `getinfo(9E)` routine.

**TABLE 2-1** Get Driver Information Entry Point Arguments

<i>cmd</i>	<i>arg</i>	<i>resultp</i>
<code>DDI_INFO_DEVT2DEVINFO</code>	Device number	Device information structure pointer
<code>DDI_INFO_DEVT2INSTANCE</code>	Device number	Device instance number

The following code is the `dummy_getinfo()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in [“Declaring the Autoconfiguration Entry Points”](#) on page 32.

```
static int
dummy_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg,
              void **resultp)
{
    cmn_err(CE_NOTE, "Inside dummy_getinfo");
    switch(cmd) {
        case DDI_INFO_DEVT2DEVINFO:
```

```

        *resultp = dummy_dip;
        return DDI_SUCCESS;
    case DDI_INFO_DEVT2INSTANCE:
        *resultp = 0;
        return DDI_SUCCESS;
    default:
        return DDI_FAILURE;
    }
}

```

First, use `cmn_err(9F)` to write a message to the system log, as you did in your `_init(9E)` entry point. Then provide `DDI_INFO_DEVT2DEVINFO` behavior. A realistic driver would use `arg` to get the instance number of this device node. A realistic driver would then call the `ddi_get_soft_state(9F)` function and return the device information structure pointer from that state structure. This dummy driver supports only one instance and does not use a state structure. In the `DDI_INFO_DEVT2DEVINFO` code of this `dummy_getinfo()` routine, simply return the one device information structure pointer that the `dummy_attach()` routine saved.

Next, provide `DDI_INFO_DEVT2INSTANCE` behavior. Within the `DDI_INFO_DEVT2INSTANCE` code, simply return 0. This dummy driver supports only one instance. The instance number of that one instance is 0.

## Defining the Report Driver Property Information Entry Point

The `prop_op(9E)` entry point is required for every driver. If your driver does not need to customize the behavior of the `prop_op(9E)` entry point, then your driver can use the `ddi_prop_op(9F)` function for the `prop_op(9E)` entry point. Drivers that create and manage their own properties need a custom `prop_op(9E)` routine. This dummy driver uses a `prop_op(9E)` routine to call `cmn_err(9F)` before calling the `ddi_prop_op(9F)` function.

The `prop_op(9E)` entry point and the `ddi_prop_op(9F)` function both require that you include the `types.h` header file. The `prop_op(9E)` entry point and the `ddi_prop_op(9F)` function both take the same seven arguments. These arguments are not discussed here because this dummy driver does not create and manage its own properties. See the `prop_op(9E)` man page to learn about the `prop_op(9E)` arguments.

The following code is the `dummy_prop_op()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in [“Declaring the Autoconfiguration Entry Points”](#) on page 32.

```

static int
dummy_prop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op,
              int flags, char *name, caddr_t valuep, int *lengthp)
{

```

```

        cmn_err(CE_NOTE, "Inside dummy_prop_op");
        return(ddi_prop_op(dev,dip,prop_op,flags,name,valuep,lengthp));
    }

```

First, use `cmn_err(9F)` to write a message to the system log, as you did in your `_init(9E)` entry point. Then call the `ddi_prop_op(9F)` function with exactly the same arguments as the `dummy_prop_op()` function.

## Including Autoconfiguration Header Files

All of the autoconfiguration entry point routines and all of the user context entry point routines require that you include the `ddi.h` and `sunddi.h` header files. You already included these two header files for the `cmn_err(9F)` function.

The `ddi_create_minor_node(9F)` function requires the `stat.h` header file. The `dummy_attach()` routine calls the `ddi_create_minor_node(9F)` function. The `prop_op(9E)` and the `ddi_prop_op(9F)` functions require the `types.h` header file.

The following code is the list of header files that you now should have included in your `dummy.c` file for the four autoconfiguration routines you have written in this section and the three loadable module configuration routines you wrote in the previous section.

```

#include <sys/modctl.h> /* used by _init, _info, _fini */
#include <sys/types.h> /* used by prop_op, ddi_prop_op */
#include <sys/stat.h> /* defines S_IFCHR used by ddi_create_minor_node */
#include <sys/cmn_err.h> /* used by all entry points for this driver */
#include <sys/ddi.h> /* used by all entry points for this driver */
/* also used by ddi_get_instance, ddi_prop_op */
#include <sys/sunddi.h> /* used by all entry points for this driver */
/* also used by ddi_create_minor_node, */
/* ddi_get_instance, and ddi_prop_op */

```

## Writing the User Context Entry Points

User context entry points correspond closely to system calls. When a system call opens a device file, then the `open(9E)` routine in the driver for that device is called.

All character and block drivers must define the `open(9E)` user context entry point. However, the `open(9E)` routine can be `nulldev(9F)`. The `close(9E)`, `read(9E)`, and `write(9E)` user context routines are optional.

- The `open(9E)` routine gains access to the device.
- The `close(9E)` routine relinquishes access to the device. The `close(9E)` routine must undo everything that the `open(9E)` routine did.
- The `read(9E)` routine reads data from the device node.
- The `write(9E)` routine writes data to the device node.

In this section, the following code is added:

```
/* Use context entry points */
static int
dummy_open(dev_t *devp, int flag, int otyp, cred_t *cred)
{
    cmn_err(CE_NOTE, "Inside dummy_open");
    return DDI_SUCCESS;
}

static int
dummy_close(dev_t dev, int flag, int otyp, cred_t *cred)
{
    cmn_err(CE_NOTE, "Inside dummy_close");
    return DDI_SUCCESS;
}

static int
dummy_read(dev_t dev, struct uio *uiop, cred_t *credp)
{
    cmn_err(CE_NOTE, "Inside dummy_read");
    return DDI_SUCCESS;
}

static int
dummy_write(dev_t dev, struct uio *uiop, cred_t *credp)
{
    cmn_err(CE_NOTE, "Inside dummy_write");
    return DDI_SUCCESS;
}
```

## Declaring the User Context Entry Points

The user context entry point routines need to be uniquely named for this driver. Use the same prefix for each of the user context entry points that you used for each of the autoconfiguration entry point routines. The following declarations are the entry point declarations you should have in your `dummy.c` file:

```
static int dummy_attach(dev_info_t *dip, ddi_attach_cmd_t cmd);
static int dummy_detach(dev_info_t *dip, ddi_detach_cmd_t cmd);
static int dummy_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg,
    void **resultp);
static int dummy_prop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op,
    int flags, char *name, caddr_t valuep, int *lengthp);
static int dummy_open(dev_t *devp, int flag, int otyp, cred_t *cred);
static int dummy_close(dev_t dev, int flag, int otyp, cred_t *cred);
static int dummy_read(dev_t dev, struct uio *uiop, cred_t *credp);
static int dummy_write(dev_t dev, struct uio *uiop, cred_t *credp);
```

## Defining the Open Device Entry Point

The `open(9E)` routine returns type `int`. The `open(9E)` routine should return either `DDI_SUCCESS` or the appropriate error number.

The `open(9E)` routine takes four arguments. This dummy driver is so simple that this `dummy_open()` routine does not use any of the `open(9E)` arguments. The examples in [Chapter 3](#) show the `open(9E)` routine in more detail.

The following code is the `dummy_open()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in “[Declaring the User Context Entry Points](#)” on page 39. Write a message to the system log and return success.

```
static int
dummy_open(dev_t *devp, int flag, int otyp, cred_t *cred)
{
    cmn_err(CE_NOTE, "Inside dummy_open");
    return DDI_SUCCESS;
}
```

## Defining the Close Device Entry Point

The `close(9E)` routine returns type `int`. The `close(9E)` routine should return either `DDI_SUCCESS` or the appropriate error number.

The `close(9E)` routine takes four arguments. This dummy driver is so simple that this `dummy_close()` routine does not use any of the `close(9E)` arguments. The examples in [Chapter 3](#) show the `close(9E)` routine in more detail.

The `close(9E)` routine must undo everything that the `open(9E)` routine did. The `close(9E)` routine must deallocate anything that was allocated, close anything that was opened, and destroy anything that was created in the `open(9E)` routine. In this dummy driver, the `open(9E)` routine is so simple that nothing needs to be reclaimed or undone in the `close(9E)` routine.

The following code is the `dummy_close()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in “[Declaring the User Context Entry Points](#)” on page 39. Write a message to the system log and return success.

```
static int
dummy_close(dev_t dev, int flag, int otyp, cred_t *cred)
{
    cmn_err(CE_NOTE, "Inside dummy_close");
    return DDI_SUCCESS;
}
```

## Defining the Read Device Entry Point

The `read(9E)` routine returns type `int`. The `read(9E)` routine should return either `DDI_SUCCESS` or the appropriate error number.

The `read(9E)` routine takes three arguments. This dummy driver is so simple that this `dummy_read()` routine does not use any of the `read(9E)` arguments. The examples in [Chapter 3](#) show the `read(9E)` routine in more detail.



The following code is the `dummy_read()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in “[Declaring the User Context Entry Points](#)” on page 39. Write a message to the system log and return success.

```
static int
dummy_read(dev_t dev, struct uio *uiop, cred_t *credp)
{
    cmn_err(CE_NOTE, "Inside dummy_read");
    return DDI_SUCCESS;
}
```

## Defining the Write Device Entry Point

The `write(9E)` routine returns type `int`. The `write(9E)` routine should return either `DDI_SUCCESS` or the appropriate error number.

The `write(9E)` routine takes three arguments. This dummy driver is so simple that this `dummy_write()` routine does not use any of the `write(9E)` arguments. The examples in [Chapter 3](#) show the `write(9E)` routine in more detail.

The following code is the `dummy_write()` routine that you should enter into your `dummy.c` file. You can copy the name portion of this function definition directly from the declaration you entered in “[Declaring the User Context Entry Points](#)” on page 39. Write a message to the system log and return success.

```
static int
dummy_write(dev_t dev, struct uio *uiop, cred_t *credp)
{
    cmn_err(CE_NOTE, "Inside dummy_write");
    return DDI_SUCCESS;
}
```

## Including User Context Header Files

The four user context entry point routines require your module to include several header files. You already have included the `types.h` header file, the `ddi.h` header file, and the `sunddi.h` header file. You need to include the `file.h`, `errno.h`, `open.h`, `cred.h`, and `uio.h` header files.

The following code is the list of header files that you now should have included in your `dummy.c` file for all the entry points you have written in this section and the previous two sections:

```
#include <sys/modctl.h> /* used by modlinkage, modldrv, _init, _info, */
/* and _fini */
#include <sys/types.h> /* used by open, close, read, write, prop_op, */
/* and ddi_prop_op */
#include <sys/file.h> /* used by open, close */
```

```

#include <sys/errno.h> /* used by open, close, read, write */
#include <sys/open.h> /* used by open, close, read, write */
#include <sys/cred.h> /* used by open, close, read */
#include <sys/uio.h> /* used by read */
#include <sys/stat.h> /* defines S_IFCHR used by ddi_create_minor_node */
#include <sys/cmn_err.h> /* used by all entry points for this driver */
#include <sys/ddi.h> /* used by all entry points for this driver */
/* also used by ddi_get_instance and */
/* ddi_prop_op */
#include <sys/sunddi.h> /* used by all entry points for this driver */
/* also used by ddi_create_minor_node, */
/* ddi_get_instance, and ddi_prop_op */

```

## Writing the Driver Data Structures

All of the data structures described in this section are required for every device driver. All drivers must define a `dev_ops(9S)` device operations structure. Because the `dev_ops(9S)` structure includes a pointer to the `cb_ops(9S)` character and block operations structure, you must define the `cb_ops(9S)` structure first. The `modldrv(9S)` linkage structure for loadable drivers includes a pointer to the `dev_ops(9S)` structure. The `modlinkage(9S)` module linkage structure includes a pointer to the `modldrv(9S)` structure.

Except for the loadable module configuration entry points, all of the required entry points for a driver are initialized in the character and block operations structure or in the device operations structure. Some optional entry points and other related data also are initialized in these data structures. Initializing the entry points in these data structures enables the driver to be dynamically loaded.

The loadable module configuration entry points are not initialized in driver data structures. The `_init(9E)`, `_info(9E)`, and `_fini(9E)` entry points are required for all kernel modules and are not specific to device driver modules.

In this section, the following code is added:

```

/* cb_ops structure */
static struct cb_ops dummy_cb_ops = {
    dummy_open,
    dummy_close,
    nodev, /* no strategy - nodev returns ENXIO */
    nodev, /* no print */
    nodev, /* no dump */
    dummy_read,
    dummy_write,
    nodev, /* no ioctl */
    nodev, /* no devmap */
    nodev, /* no mmap */
    nodev, /* no segmap */
    nochpoll, /* returns ENXIO for non-pollable devices */
    dummy_prop_op,
    NULL, /* streamtab struct; if not NULL, all above */
}

```

```

                                /* fields are ignored */
    D_NEW | D_MP,                /* compatibility flags: see conf.h */
    CB_REV,                      /* cb_ops revision number */
    nodev,                       /* no aread */
    nodev                        /* no awrite */
};

/* dev_ops structure */
static struct dev_ops dummy_dev_ops = {
    DEVO_REV,
    0,                          /* reference count */
    dummy_getinfo,
    nulldev,                    /* no identify - nulldev returns 0 */
    nulldev,                    /* no probe */
    dummy_attach,
    dummy_detach,
    nodev,                      /* no reset - nodev returns ENXIO */
    &dummy_cb_ops,
    (struct bus_ops *)NULL,
    nodev                        /* no power */
};

/* modldrv structure */
static struct modldrv md = {
    &mod_driverops,            /* Type of module. This is a driver. */
    "dummy driver",          /* Name of the module. */
    &dummy_dev_ops
};

/* modlinkage structure */
static struct modlinkage ml = {
    MODREV_1,
    &md,
    NULL
};

/* dev_info structure */
dev_info_t *dummy_dip; /* keep track of one instance */

```

## Defining the Character and Block Operations Structure

The `cb_ops(9S)` structure initializes standard character and block interfaces. See the `cb_ops(9S)` man page to learn what each element is and what the value of each element should be. This dummy driver does not use all of the elements in the `cb_ops(9S)` structure. See the description that follows the code sample.

When you name this structure, use the same `dummy_` prefix that you used for the names of the autoconfiguration routines and the names of the user context routines. Prepend the `static` type modifier to the declaration.

The following code is the `cb_ops(9S)` structure that you should enter into your `dummy.c` file:

```

static struct cb_ops dummy_cb_ops = {
    dummy_open,
    dummy_close,
    nodev,          /* no strategy - nodev returns ENXIO */
    nodev,          /* no print */
    nodev,          /* no dump */
    dummy_read,
    dummy_write,
    nodev,          /* no ioctl */
    nodev,          /* no devmap */
    nodev,          /* no mmap */
    nodev,          /* no segmap */
    nochpoll,      /* returns ENXIO for non-pollable devices */
    dummy_prop_op,
    NULL,          /* streamtab struct; if not NULL, all above */
                  /* fields are ignored */
    D_NEW | D_MP,  /* compatibility flags: see conf.h */
    CB_REV,        /* cb_ops revision number */
    nodev,          /* no aread */
    nodev          /* no awrite */
};

```

Enter the names of the `open(9E)` and `close(9E)` entry points for this driver as the values of the first two elements of this structure. Enter the names of the `read(9E)` and `write(9E)` entry points for this driver as the values of the sixth and seventh elements of this structure. Enter the name of the `prop_op(9E)` entry point for this driver as the value of the thirteenth element in this structure.

The `strategy(9E)`, `print(9E)`, and `dump(9E)` routines are for block drivers only. This dummy driver does not define these three routines because this driver is a character driver. This driver does not define an `ioctl(9E)` entry point because this driver does not use I/O control commands. This driver does not define `devmap(9E)`, `mmap(9E)`, or `segmap(9E)` entry points because this driver does not support memory mapping. This driver does not define `aread(9E)` or `awrite(9E)` entry points because this driver does not perform any asynchronous reads or writes. Initialize all of these unused function elements to `nodev(9F)`. The `nodev(9F)` function returns the `ENXIO` error code.

Specify the `nochpoll(9F)` function for the `chpoll(9E)` element of the `cb_ops(9S)` structure because this driver is not for a pollable device. Specify `NULL` for the `streamtab(9S)` STREAMS entity declaration structure because this driver is not a STREAMS driver.

The compatibility flags are defined in the `conf.h` header file. The `D_NEW` flag means this driver is a new-style driver. The `D_MP` flag means this driver safely allows multiple threads of execution. All drivers must be multithreaded-safe, and must specify this `D_MP` flag. The `D_64BIT` flag means this driver supports 64-bit offsets and block numbers. See the `conf.h` header file for more compatibility flags.

The `CB_REV` element of the `cb_ops(9S)` structure is the `cb_ops(9S)` revision number. `CB_REV` is defined in the `devops.h` header file.

## Defining the Device Operations Structure

The `dev_ops(9S)` structure initializes interfaces that are used for operations such as attaching and detaching the driver. See the `dev_ops(9S)` man page to learn what each element is and what the value of each element should be. This dummy driver does not use all of the elements in the `dev_ops(9S)` structure. See the description that follows the code sample.

When you name this structure, use the same `dummy_` prefix that you used for the names of the autoconfiguration routines and the names of the user context routines. Prepend the `static` type modifier to the declaration.

The following code is the `dev_ops(9S)` structure that you should enter into your `dummy.c` file:

```
static struct dev_ops dummy_dev_ops = {
    DEVO_REV,
    0, /* reference count */
    dummy_getinfo,
    nulldev, /* no identify - nulldev returns 0 */
    nulldev, /* no probe */
    dummy_attach,
    dummy_detach,
    nodev, /* no reset - nodev returns ENXIO */
    &dummy_cb_ops,
    (struct bus_ops *)NULL,
    nodev /* no power */
};
```

The `DEVO_REV` element of the `dev_ops(9S)` structure is the driver build version. `DEVO_REV` is defined in the `devops.h` header file. The second element in this structure is the driver reference count. Initialize this value to zero. The driver reference count is the number of instances of this driver that are currently open. The driver cannot be unloaded if any instances of the driver are still open.

The next six elements of the `dev_ops(9S)` structure are the names of the `getinfo(9E)`, `identify(9E)`, `probe(9E)`, `attach(9E)`, `detach(9E)`, and `reset()` functions for this particular driver. The `identify(9E)` function is obsolete. Initialize this structure element to `nulldev(9F)`. The `probe(9E)` function determines whether the corresponding device exists and is valid. This dummy driver does not define a `probe(9E)` function. Initialize this structure element to `nulldev`. The `nulldev(9F)` function returns success. The `reset()` function is obsolete. Initialize the `reset()` function to `nodev(9F)`.

The next element of the `dev_ops(9S)` structure is a pointer to the `cb_ops(9S)` structure for this driver. You initialized the `cb_ops(9S)` structure for this driver in “[Defining the Character and Block Operations Structure](#)” on page 43. Enter `&dummy_cb_ops` for the value of the pointer to the `cb_ops(9S)` structure.

The next element of the `dev_ops(9S)` structure is a pointer to the bus operations structure. Only nexus drivers have bus operations structures. This dummy driver is not a nexus driver. Set this value to `NULL` because this driver is a leaf driver.

The last element of the `dev_ops(9S)` structure is the name of the `power(9E)` routine for this driver. The `power(9E)` routine operates on a hardware device. This driver does not drive a hardware device. Set the value of this structure element to `nodev`.

## Defining the Module Linkage Structures

Two other module loading structures are required for every driver. The `modlinkage(9S)` module linkage structure is used by the `_init(9E)`, `_info(9E)`, and `_fini(9E)` routines to install, remove, and retrieve information from a module. The `modldrv(9S)` linkage structure for loadable drivers exports driver-specific information to the kernel. See the man pages for each structure to learn what each element is and what the value of each element should be.

The following code defines the `modldrv(9S)` and `modlinkage(9S)` structures for the driver shown in this chapter:

```
static struct modldrv md = {
    &mod_driverops,      /* Type of module. This is a driver. */
    "dummy driver",     /* Name of the module. */
    &dummy_dev_ops
};

static struct modlinkage ml = {
    MODREV_1,
    &md,
    NULL
};
```

The first element in the `modldrv(9S)` structure is a pointer to a structure that tells the kernel what kind of module this is. Set this value to the address of the `mod_driverops` structure. The `mod_driverops` structure tells the kernel that the `dummy.c` module is a loadable driver module. The `mod_driverops` structure is declared in the `modctl.h` header file. You already included the `modctl.h` header file in your `dummy.c` file, so do not declare the `mod_driverops` structure in `dummy.c`. The `mod_driverops` structure is defined in the `modctl.c` source file.

The second element in the `modldrv(9S)` structure is a string that describes this module. Usually this string contains the name of this module and the version number of this module. The last element of the `modldrv(9S)` structure is a pointer to the `dev_ops(9S)` structure for this driver. You initialized the `dev_ops(9S)` structure for this driver in [“Defining the Device Operations Structure” on page 45](#).

The first element in the `modlinkage(9S)` structure is the revision number of the loadable modules system. Set this value to `MODREV_1`. The next element of the `modlinkage(9S)` structure is the address of a null-terminated array of pointers to linkage structures. Driver modules have only one linkage structure. Enter the address of the `md` structure for the value of this element of the `modlinkage(9S)` structure. Enter the value `NULL` to terminate this list of linkage structures.

## Including Data Structures Header Files

The `cb_ops(9S)` and `dev_ops(9S)` structures require you to include the `conf.h` and `devops.h` header files. The `modlinkage(9S)` and `modldrv(9S)` structures require you to include the `modctl.h` header file. You already included the `modctl.h` header file for the loadable module configuration entry points.

The following code is the complete list of header files that you now should have included in your `dummy.c` file:

```
#include <sys/devops.h> /* used by dev_ops */
#include <sys/conf.h> /* used by dev_ops and cb_ops */
#include <sys/modctl.h> /* used by modlinkage, modldrv, _init, _info, */
/* and _fini */
#include <sys/types.h> /* used by open, close, read, write, prop_op, */
/* and ddi_prop_op */
#include <sys/file.h> /* used by open, close */
#include <sys/errno.h> /* used by open, close, read, write */
#include <sys/open.h> /* used by open, close, read, write */
#include <sys/cred.h> /* used by open, close, read */
#include <sys/uio.h> /* used by read */
#include <sys/stat.h> /* defines S_IFCHR used by ddi_create_minor_node */
#include <sys/cmn_err.h> /* used by all entry points for this driver */
#include <sys/ddi.h> /* used by all entry points for this driver */
/* also used by cb_ops, ddi_get_instance, and */
/* ddi_prop_op */
#include <sys/sunddi.h> /* used by all entry points for this driver */
/* also used by cb_ops, ddi_create_minor_node, */
/* ddi_get_instance, and ddi_prop_op */
```

---

## Writing the Device Configuration File

This driver requires a configuration file. A driver that is not self-identifying must have a `node_name.conf` file. A self-identifying driver is a driver that can obtain all the property information it needs from the DDI property interfaces such as `ddi_prop_get_int(9F)` and `ddi_prop_lookup(9F)`. For information about device driver configuration files, see the `driver.conf(4)` man page.

The minimum information that a configuration file must contain is the name of the device node and the name or type of the device's parent. In this simple example, the node name of the device is the same as the file name of the driver. Create a file named `dummy.conf` in your working directory. Put the following single line of information into `dummy.conf`:

```
name="dummy" parent="pseudo";
```

---

## Building and Installing the Driver

This section tells you which compile options to use for different architectures and where to install the driver for different architectures.

### Compiling and Linking the Driver

Use the `-D_KERNEL` option to indicate that this code defines a kernel module.

- If you are compiling for a 64-bit SPARC architecture, use the `-xarch=v9` option:

```
% cc -D_KERNEL -xarch=v9 -c dummy.c
% ld -r -o dummy dummy.o
```

- If you are compiling for a 64-bit x86 architecture, use the `-xarch=amd64` option and the `-xmodel=kernel` option:

```
% cc -D_KERNEL -xarch=amd64 -xmodel=kernel -c dummy.c
% ld -r -o dummy dummy.o
```

- If you are compiling for a 32-bit architecture, use the following build commands:

```
% cc -D_KERNEL -c dummy.c
% ld -r -o dummy dummy.o
```

### Installing the Driver in a Temporary Location

Make sure you are user `root` when you install the driver.

Install drivers in the `/tmp` directory until you are finished modifying and testing the `_info()`, `_init()`, and `attach()` routines. Copy the driver binary to the `/tmp` directory. Link to the driver from the kernel driver directory.

If a driver has an error in its `_info()`, `_init()`, or `attach()` function, your machine could get into a state of infinite panic. The Solaris OS automatically reboots itself after a panic. The Solaris OS loads any drivers it can during boot. If you have an error in your `attach()` function that panics the system when you load the driver, then the system will panic again when it tries to reboot after the panic. The system will continue the cycle of panic, reboot, panic as it attempts to reload the faulty driver every time it reboots after panic.

To avoid an infinite panic, keep the driver in the `/tmp` area until it is well tested. Link to the driver in the `/tmp` area from the kernel driver area. The Solaris OS removes all files from the `/tmp` area every time the system reboots. If your driver causes a panic, the Solaris OS reboots successfully because the driver has been removed automatically from the `/tmp` area. The link in the kernel driver area points to nothing. The faulty



driver did not get loaded, so the system does not go back into a panic. You can modify the driver, copy it again to the /tmp area, and continue testing and developing. When the driver is well tested, copy it to the /usr/kernel/drv area so that it will remain available after a reboot.

```
# cp dummy /tmp
```

- On a 64-bit SPARC architecture, link to the sparcv9 directory:

```
# ln -s /tmp/dummy /usr/kernel/drv/sparcv9/dummy
```

- On a 64-bit x86 architecture, link to the amd64 directory :

```
# ln -s /tmp/dummy /usr/kernel/drv/amd64/dummy
```

- On a 32-bit architecture, create the link as follows:

```
# ln -s /tmp/dummy /usr/kernel/drv/dummy
```

Copy the configuration file to the kernel driver area of the system.

```
# cp dummy.conf /usr/kernel/drv
```

---

## Testing the Template Driver

This dummy driver merely writes a message to a system log each time an entry point routine is entered. To test this driver, watch for these messages to confirm that each entry point routine is successfully entered.

The `cmn_err(9F)` function writes low priority messages such as the messages defined in this dummy driver to /dev/log. The `syslogd(1M)` daemon reads messages from /dev/log and writes low priority messages to /var/adm/messages.

In a separate window, enter the following command and monitor the output as you perform the tests described in the remainder of this section:

```
% tail -f /var/adm/messages
```

## Loading the Driver

Make sure you are user `root` when you load the driver. Use the `add_drv(1M)` command to load the driver:

```
# add_drv dummy
```

You should see the following messages in the window where you are viewing /var/adm/messages:

```
date time machine dummy: [ID 513080 kern.notice] NOTICE: Inside _info
date time machine dummy: [ID 874762 kern.notice] NOTICE: Inside _init
```

```
date time machine dummy: [ID 678704 kern.notice] NOTICE: Inside dummy_attach
```

The `_info(9E)`, `_init(9E)`, and `attach(9E)` entry points are called in that order when you add a driver.

The dummy driver has been added to the `/devices` directory:

```
% ls -l /devices/pseudo | grep dummy
drwxr-xr-x  2 root    sys      512 date time dummy@0
crw-----  1 root    sys      92,  0 date time dummy@0:0
```

The dummy driver also is the most recent module listed by `modinfo(1M)`:

```
% modinfo
  Id Loadaddr  Size Info Rev Module Name
180 ed192b70   544 92  1  dummy (dummy driver)
```

The module name, `dummy driver`, is the value you entered for the second member of the `modldrv(9S)` structure. The value 92 is the major number of this module.

```
% grep dummy /etc/name_to_major
dummy 92
```

The `Loadaddr` address of `ed192b70` is the address of the first instruction in the dummy driver. This address might be useful, for example, in debugging.

```
% mdb -k
> dummy'_init $m
  BASE      LIMIT      SIZE NAME
ed192b70 ed192ff0    480 dummy
> $q
```

The dummy driver also is the most recent module listed by `prtconf(1M)` in the `pseudo` device section:

```
% prtconf -P
pseudo, instance #0
  dummy, instance #0 (driver not attached)
```

A driver is automatically unloaded when it is not in use. If your driver is in the `/devices` directory but `modinfo(1M)` does not list your driver, you can use either of the following methods to load your driver:

- Use the `modload(1M)` command.
- Access the device. The driver is loaded automatically when it is accessed. The following section describes how to access the dummy device.

## Reading and Writing the Device

Make sure you are user root when you perform the tests described in this section. If you are not user root, you will receive “Permission denied” error messages when you try to access the `/devices/pseudo/dummy@0:0` special file. Notice the permissions that are shown for `/devices/pseudo/dummy@0:0` in “Loading the Driver” on page 49.

Test reading from the device. Your dummy device probably is named `/devices/pseudo/dummy@0:0`. The following command reads from your dummy device even if it has a slightly different name:

```
# cat /devices/pseudo/dummy*
```

You should see the following messages in the window where you are viewing `/var/adm/messages`:

```
date time machine dummy: [ID 136952 kern.notice] NOTICE: Inside dummy_open
date time machine dummy: [ID 623947 kern.notice] NOTICE: Inside dummy_getinfo
date time machine dummy: [ID 891851 kern.notice] NOTICE: Inside dummy_prop_op
date time machine dummy: [ID 623947 kern.notice] NOTICE: Inside dummy_getinfo
date time machine dummy: [ID 891851 kern.notice] NOTICE: Inside dummy_prop_op
date time machine dummy: [ID 623947 kern.notice] NOTICE: Inside dummy_getinfo
date time machine dummy: [ID 709590 kern.notice] NOTICE: Inside dummy_read
date time machine dummy: [ID 550206 kern.notice] NOTICE: Inside dummy_close
```

Test writing to the device:

```
# echo hello > `ls /devices/pseudo/dummy*`
```

You should see the following messages in the window where you are viewing `/var/adm/messages`:

```
date time machine dummy: [ID 136952 kern.notice] NOTICE: Inside dummy_open
date time machine dummy: [ID 623947 kern.notice] NOTICE: Inside dummy_getinfo
date time machine dummy: [ID 891851 kern.notice] NOTICE: Inside dummy_prop_op
date time machine dummy: [ID 623947 kern.notice] NOTICE: Inside dummy_getinfo
date time machine dummy: [ID 891851 kern.notice] NOTICE: Inside dummy_prop_op
date time machine dummy: [ID 623947 kern.notice] NOTICE: Inside dummy_getinfo
date time machine dummy: [ID 672780 kern.notice] NOTICE: Inside dummy_write
date time machine dummy: [ID 550206 kern.notice] NOTICE: Inside dummy_close
```

As you can see, this output from the write test is almost identical to the output you saw from the read test. The only difference is in the seventh line of the output. Using the `cat(1)` command causes the kernel to access the `read(9E)` entry point of the driver. Using the `echo(1)` command causes the kernel to access the `write(9E)` entry point of the driver. The text argument that you give to `echo(1)` is ignored because this driver does not do anything with that data.

## Unloading the Driver

Make sure you are user root when you unload the driver. Use the `rem_drv(1M)` command to unload the driver and remove the device from the `/devices` directory:

```
# rem_drv dummy
```

You should see the following messages in the window where you are viewing `/var/adm/messages`:

```
date time machine dummy: [ID 513080 kern.notice] NOTICE: Inside _info
date time machine dummy: [ID 617648 kern.notice] NOTICE: Inside dummy_detach
date time machine dummy: [ID 812373 kern.notice] NOTICE: Inside _fini
```

The dummy device is no longer in the `/devices` directory:

```
# ls /devices/pseudo/dummy*
/devices/pseudo/dummy*: No such file or directory
```

The next time you want to read from or write to the dummy device, you must load the driver again using `add_drv(1M)`.

You can use the `modunload(1M)` command to unload the driver but not remove the device from `/devices`. Then the next time you read from or write to the dummy device, the driver is automatically loaded.

Press Control-C to stop tailing the `/var/adm/messages` messages.

---

## Complete Driver Source

The following code is the complete source for the dummy driver described in this chapter:

```
/*
 * Minimalist pseudo-device.
 * Writes a message whenever a routine is entered.
 *
 * Build the driver:
 *     cc -D_KERNEL -c dummy.c
 *     ld -r -o dummy dummy.o
 * Copy the driver and the configuration file to /usr/kernel/drv:
 *     cp dummy.conf /usr/kernel/drv
 *     cp dummy /tmp
 *     ln -s /tmp/dummy /usr/kernel/drv/dummy
 * Add the driver:
 *     add_drv dummy
 * Test (1) read from driver (2) write to driver:
 *     cat /devices/pseudo/dummy@*
 *     echo hello > `ls /devices/pseudo/dummy@*`
 * Verify the tests in another window:
 *     tail -f /var/adm/messages
 * Remove the driver:
 *     rem_drv dummy
 */
```

```

#include <sys/devops.h> /* used by dev_ops */
#include <sys/conf.h> /* used by dev_ops and cb_ops */
#include <sys/modctl.h> /* used by modlinkage, modldrv, _init, _info, */
/* and _fini */
#include <sys/types.h> /* used by open, close, read, write, prop_op, */
/* and ddi_prop_op */
#include <sys/file.h> /* used by open, close */
#include <sys/errno.h> /* used by open, close, read, write */
#include <sys/open.h> /* used by open, close, read, write */
#include <sys/cred.h> /* used by open, close, read */
#include <sys/uio.h> /* used by read */
#include <sys/stat.h> /* defines S_IFCHR used by ddi_create_minor_node */
#include <sys/cmn_err.h> /* used by all entry points for this driver */
#include <sys/ddi.h> /* used by all entry points for this driver */
/* also used by cb_ops, ddi_get_instance, and */
/* ddi_prop_op */
#include <sys/sunddi.h> /* used by all entry points for this driver */
/* also used by cb_ops, ddi_create_minor_node, */
/* ddi_get_instance, and ddi_prop_op */

static int dummy_attach(dev_info_t *dip, ddi_attach_cmd_t cmd);
static int dummy_detach(dev_info_t *dip, ddi_detach_cmd_t cmd);
static int dummy_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg,
void **resultp);
static int dummy_prop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op,
int flags, char *name, caddr_t valuep, int *lengthp);
static int dummy_open(dev_t *devp, int flag, int otyp, cred_t *cred);
static int dummy_close(dev_t dev, int flag, int otyp, cred_t *cred);
static int dummy_read(dev_t dev, struct uio *uiop, cred_t *credp);
static int dummy_write(dev_t dev, struct uio *uiop, cred_t *credp);

/* cb_ops structure */
static struct cb_ops dummy_cb_ops = {
    dummy_open,
    dummy_close,
    nodev, /* no strategy - nodev returns ENXIO */
    nodev, /* no print */
    nodev, /* no dump */
    dummy_read,
    dummy_write,
    nodev, /* no ioctl */
    nodev, /* no devmap */
    nodev, /* no mmap */
    nodev, /* no segmap */
    nochpoll, /* returns ENXIO for non-pollable devices */
    dummy_prop_op,
    NULL, /* streamtab struct; if not NULL, all above */
/* fields are ignored */
    D_NEW | D_MP, /* compatibility flags: see conf.h */
    CB_REV, /* cb_ops revision number */
    nodev, /* no aread */
    nodev /* no awrite */
};

```

```

/* dev_ops structure */
static struct dev_ops dummy_dev_ops = {
    DEVO_REV,
    0,                /* reference count */
    dummy_getinfo,
    nulldev,         /* no identify - nulldev returns 0 */
    nulldev,         /* no probe */
    dummy_attach,
    dummy_detach,
    nodev,           /* no reset - nodev returns ENXIO */
    &dummy_cb_ops,
    (struct bus_ops *)NULL,
    nodev            /* no power */
};

/* modldrv structure */
static struct modldrv md = {
    &mod_driverops, /* Type of module. This is a driver. */
    "dummy driver", /* Name of the module. */
    &dummy_dev_ops
};

/* modlinkage structure */
static struct modlinkage ml = {
    MODREV_1,
    &md,
    NULL
};

/* dev_info structure */
dev_info_t *dummy_dip; /* keep track of one instance */

/* Loadable module configuration entry points */
int
_init(void)
{
    cmn_err(CE_NOTE, "Inside _init");
    return(mod_install(&ml));
}

int
_info(struct modinfo *modinfop)
{
    cmn_err(CE_NOTE, "Inside _info");
    return(mod_info(&ml, modinfop));
}

int
_fini(void)
{
    cmn_err(CE_NOTE, "Inside _fini");
    return(mod_remove(&ml));
}

```

```

/* Device configuration entry points */
static int
dummy_attach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    cmn_err(CE_NOTE, "Inside dummy_attach");
    switch(cmd) {
    case DDI_ATTACH:
        dummy_dip = dip;
        if (ddi_create_minor_node(dip, "0", S_IFCHR,
            ddi_get_instance(dip), DDI_PSEUDO, 0)
            != DDI_SUCCESS) {
            cmn_err(CE_NOTE,
                "%s%d: attach: could not add character node.",
                "dummy", 0);
            return(DDI_FAILURE);
        } else
            return DDI_SUCCESS;
    default:
        return DDI_FAILURE;
    }
}

static int
dummy_detach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    cmn_err(CE_NOTE, "Inside dummy_detach");
    switch(cmd) {
    case DDI_DETACH:
        dummy_dip = 0;
        ddi_remove_minor_node(dip, NULL);
        return DDI_SUCCESS;
    default:
        return DDI_FAILURE;
    }
}

static int
dummy_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg,
    void **resultp)
{
    cmn_err(CE_NOTE, "Inside dummy_getinfo");
    switch(cmd) {
    case DDI_INFO_DEVT2DEVINFO:
        *resultp = dummy_dip;
        return DDI_SUCCESS;
    case DDI_INFO_DEVT2INSTANCE:
        *resultp = 0;
        return DDI_SUCCESS;
    default:
        return DDI_FAILURE;
    }
}

/* Main entry points */
static int

```

```

dummy_prop_op(dev_t dev, dev_info_t *dip, ddi_prop_op_t prop_op,
              int flags, char *name, caddr_t valuep, int *lengthp)
{
    cmn_err(CE_NOTE, "Inside dummy_prop_op");
    return(ddi_prop_op(dev,dip,prop_op,flags,name,valuep,lengthp));
}

static int
dummy_open(dev_t *devp, int flag, int otyp, cred_t *cred)
{
    cmn_err(CE_NOTE, "Inside dummy_open");
    return DDI_SUCCESS;
}

static int
dummy_close(dev_t dev, int flag, int otyp, cred_t *cred)
{
    cmn_err(CE_NOTE, "Inside dummy_close");
    return DDI_SUCCESS;
}

static int
dummy_read(dev_t dev, struct uio *uiop, cred_t *credp)
{
    cmn_err(CE_NOTE, "Inside dummy_read");
    return DDI_SUCCESS;
}

static int
dummy_write(dev_t dev, struct uio *uiop, cred_t *credp)
{
    cmn_err(CE_NOTE, "Inside dummy_write");
    return DDI_SUCCESS;
}

```



---

## Reading and Writing Data in Kernel Memory

---

In this chapter, you will extend the very simple prototype driver you developed in the previous chapter. The driver you will develop in this chapter displays data read from kernel memory. The first version of this driver writes data to a system log every time the driver is loaded. The second version of this driver displays data at user request. In the third version of this driver, the user can write new data to the device.

---

### Displaying Data Stored in Kernel Memory

The pseudo device driver presented in this section writes a constant string to a system log when the driver is loaded.

This first version of the Quote Of The Day driver (`qotd_1`) is even more simple than the dummy driver from the previous chapter. The dummy driver includes all functions that are required to drive hardware. This `qotd_1` driver includes only the bare minimum functions it needs to make a string available to a user command. For example, this `qotd_1` driver has no `cb_ops(9S)` structure. Therefore, this driver defines no `open(9E)`, `close(9E)`, `read(9E)`, or `write(9E)` function. If you examine the `dev_ops(9S)` structure for this `qotd_1` driver, you see that no `getinfo(9E)`, `attach(9E)`, or `detach(9E)` function is defined. This driver contains no function declarations because all the functions that are defined in this driver are declared in the `modctl.h` header file. You must include the `modctl.h` header file in your `qotd_1.c` file.

This `qotd_1` driver defines a global variable to hold its text data. The `_init(9E)` entry point for this driver uses the `cmn_err(9F)` function to write the string to a system log. The dummy driver also uses the `cmn_err(9F)` function to display messages. The `qotd_1` driver is different from the dummy driver because the `qotd_1` driver stores its string in kernel memory.

# Writing Quote Of The Day Version 1

Enter the source code shown in the following example into a text file named `qotd_1.c`.

## EXAMPLE 3-1 Quote Of The Day Version 1 Source File

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

#define QOTD_MAXLEN    128

static const char qotd[QOTD_MAXLEN]
    = "Be careful about reading health books. \
You may die of a misprint. - Mark Twain\n";

static struct dev_ops qotd_dev_ops = {
    DEVO_REV,          /* devo_rev */
    0,                 /* devo_refcnt */
    ddi_no_info,       /* devo_getinfo */
    nulldev,           /* devo_identify */
    nulldev,           /* devo_probe */
    nulldev,           /* devo_attach */
    nulldev,           /* devo_detach */
    nodev,             /* devo_reset */
    (struct cb_ops *)NULL, /* devo_cb_ops */
    (struct bus_ops *)NULL, /* devo_bus_ops */
    nulldev            /* devo_power */
};

static struct modldrv modldrv = {
    &mod_driverops,
    "Quote of the Day 1.0",
    &qotd_dev_ops};

static struct modlinkage modlinkage = {
    MODREV_1,
    (void *)&modldrv,
    NULL
};

int
_init(void)
{
    cmn_err(CE_CONT, "QOTD: %s\n", qotd);
    return (mod_install(&modlinkage));
}

int
_info(struct modinfo *modinfop)
{
```

**EXAMPLE 3-1** Quote Of The Day Version 1 Source File (Continued)

```
        return (mod_info(&modlinkage, modinfo));
    }
    int
    _fini(void)
    {
        return (mod_remove(&modlinkage));
    }
}
```

Enter the configuration information shown in the following example into a text file named `qotd_1.conf`.

**EXAMPLE 3-2** Quote Of The Day Version 1 Configuration File

```
name="qotd_1" parent="pseudo" instance=0;
```

## Building, Installing, and Using Quote Of The Day Version 1

Compile and link the driver. Use the `-D_KERNEL` option to indicate that this code defines a kernel module. The following example shows compiling and linking for a 32-bit architecture:

```
% cc -D_KERNEL -c qotd_1.c
% ld -r -o qotd_1 qotd_1.o
```

Note that the name of the driver, `qotd_1`, must match the name property in the configuration file.

Make sure you are user `root` when you install the driver.

Copy the driver binary to the `/tmp` directory as discussed in [“Installing the Driver in a Temporary Location” on page 48](#).

```
# cp qotd_1 /tmp
# ln -s /tmp/qotd_1 /usr/kernel/drv/qotd_1
```

Copy the configuration file to the kernel driver area of the system.

```
# cp qotd_1.conf /usr/kernel/drv
```

This `qotd_1` driver writes a message to a system log each time the driver is loaded. The `cmn_err(9F)` function writes low priority messages such as the message defined in this `qotd_1` driver to `/dev/log`. The `syslogd(1M)` daemon reads messages from `/dev/log` and writes low priority messages to `/var/adm/messages`.

To test this driver, watch for the message in `/var/adm/messages`. In a separate window, enter the following command:

```
% tail -f /var/adm/messages
```

Make sure you are user `root` when you load the driver. Use the `add_drv(1M)` command to load the driver:

```
# add_drv qotd_1
```

You should see the following messages in the window where you are viewing `/var/adm/messages`:

```
date time machine pseudo: [ID 129642 kern.info] pseudo-device: devinfo0
date time machine genunix: [ID 936769 kern.info] devinfo0 is /pseudo/devinfo@0
date time machine qotd: [ID 197678 kern.notice] QOTD_1: Be careful about
reading health books. You may die of a misprint. - Mark Twain
```

This last line is the content of the variable output by the `cmn_err(9F)` function in the `_init(9E)` entry point. The `_init(9E)` entry point is called when the driver is loaded.

---

## Displaying Data on Demand

The sample code in this section creates a pseudo device that is controlled by the driver. The driver stores data in the device and makes the data available when the user accesses the device for reading.

This section first discusses the important code differences between these two versions of the Quote Of The Day driver. This section then shows you how you can access the device to cause the quotation to display.

## Writing Quote Of The Day Version 2

The driver that controls the pseudo device is more complex than the driver shown in the previous section. This section first explains some important features of this version of the driver. This section then shows all the source for this driver.

The following list summarizes the differences between the two versions of the Quote Of The Day driver:

- Version 2 of the driver defines a state structure that holds information about each instance of the device.
- Version 2 defines a `cb_ops(9S)` structure and a more complete `dev_ops(9S)` structure.
- Version 2 defines `open(9E)`, `close(9E)`, `read(9E)`, `getinfo(9E)`, `attach(9E)`, and `detach(9E)` entry points.

- Version 1 uses the `cmn_err(9F)` function to write a constant string to a system log in the `_init(9E)` entry point of the driver. The `_init(9E)` entry point is called when the driver is loaded. Version 2 uses the `uiomove(9F)` function to copy the quotation from kernel memory. The copied data is returned by the `read(9E)` entry point. The `read(9E)` entry point is called when the driver is accessed for reading.
- Version 2 of the driver uses `ASSERT(9F)` statements to check the validity of data.

The following sections provide more detail about the additions and changes in Version 2 of the Quote Of The Day driver.

## Managing Device State

The `_init(9E)` and `_fini(9E)` entry points and all six new entry points defined in this driver maintain a soft state for the device. Most device drivers maintain state information with each instance of the device they control. An instance usually is a sub-device. For example, a disk driver may communicate with a hardware controller device that has several disk drives attached. See “Retrieving Driver Soft State Information” in *Writing Device Drivers* for more information about soft states.

This sample driver allows only one instance. The instance number is assigned in the configuration file. See [Example 3–4](#). Most device drivers allow any number of instances of a device to be created. The system manages the device instance numbers, and the DDI soft state functions manage the instances.

The following flow gives an overview of how DDI soft state functions manage a state pointer and the state of a device instance:

1. The `ddi_soft_state_init(9F)` function initializes the state pointer. The state pointer is an opaque handle that enables allocation, deallocation, and tracking of a state structure for each instance of a device. The state structure is a user-defined type that maintains data specific to this instance of the device. In this example, the state pointer and state structure are declared after the entry point declarations. See `qotd_state_head` and `qotd_state` in [Example 3–3](#).
2. The `ddi_soft_state_zalloc(9F)` function uses the state pointer and the device instance to create the state structure for this instance.
3. The `ddi_get_soft_state(9F)` function uses the state pointer and the device instance to retrieve the state structure for this instance of the device.
4. The `ddi_soft_state_free(9F)` function uses the state pointer and the device instance to free the state structure for this instance.
5. The `ddi_soft_state_fini(9F)` function uses the state pointer to destroy the state pointer and the state structures for all instances of this device.

The `ddi_soft_state_zalloc(9F)`, `ddi_get_soft_state(9F)`, and `ddi_soft_state_free(9F)` functions coordinate access to the underlying data structures in a way that is safe for multithreading. No additional locks should be necessary.

## Initializing and Unloading

The `_init(9E)` entry point first calls the `ddi_soft_state_init(9F)` function to initialize the soft state. If the soft state initialization fails, that error code is returned. If the soft state initialization succeeds, the `_init(9E)` entry point calls the `mod_install(9F)` function to load a new module. If the module install fails, the `_init(9E)` entry point calls the `ddi_soft_state_fini(9F)` function and returns the error code from the failed module install.

Your code must undo everything that it does. You must call `ddi_soft_state_fini(9F)` if the module install fails because the `_init(9E)` call succeeded and created a state pointer.

The `_fini(9E)` entry point must undo everything the `_init(9E)` entry point did. The `_fini(9E)` entry point first calls the `mod_remove(9F)` function to remove the module that the `_init(9E)` entry point installed. If the module remove fails, that error code is returned. If the module remove succeeds, the `_fini(9E)` entry point calls the `ddi_soft_state_fini(9F)` function to destroy the state pointer and the state structures for all instances of this device.

## Attaching and Detaching

The `attach(9E)` entry point first calls the `ddi_get_instance(9F)` function to retrieve the instance number of the device information node. The `attach(9E)` entry point uses this instance number to call the `ddi_soft_state_zalloc(9F)`, `ddi_get_soft_state(9F)`, and `ddi_create_minor_node(9F)` functions.

The `attach(9E)` entry point calls the `ddi_soft_state_zalloc(9F)` function to create a state structure for this device instance. If creation of the soft state structure fails, `attach(9E)` writes an error message to a system log and returns failure. This device instance is not attached. If creation of the soft state structure succeeds, `attach(9E)` calls the `ddi_get_soft_state(9F)` function to retrieve the state structure for this device instance.

If retrieval of the state structure fails, `attach(9E)` writes an error message to a system log, calls the `ddi_soft_state_free(9F)` function to destroy the state structure that was created by `ddi_soft_state_zalloc(9F)`, and returns failure. This device instance is not attached. If retrieval of the state structure succeeds, `attach(9E)` calls the `ddi_create_minor_node(9F)` function to create the device node.

At the top of this driver source file, a constant named `QOTD_NAME` is defined that holds the string name of the device. This constant is one of the arguments that is passed to `ddi_create_minor_node(9F)`. If creation of the device node fails, `attach(9E)` writes an error message to a system log, calls the `ddi_soft_state_free(9F)` function to destroy the state structure that was created by `ddi_soft_state_zalloc(9F)`, calls the `ddi_remove_minor_node(9F)` function, and returns failure. This device instance is not attached.

If creation of the device node succeeds, this device instance is attached. The `attach(9E)` entry point assigns the instance number that was retrieved with `ddi_get_instance(9F)` to the instance member of the state structure for this instance. Then `attach(9E)` assigns the `dev_info` structure pointer that was passed in the `attach(9E)` call to the `dev_info` structure pointer member of the state structure for this instance. The `ddi_report_dev(9F)` function writes a message in the system log file when the device is added or when the system is booted. The message announces this device as shown in the following example:

```
% dmesg
date time machine pseudo: [ID 129642 kern.info] pseudo-device: qotd_20
date time machine genunix: [ID 936769 kern.info] qotd_20 is /pseudo/qotd_2@0
```

The `detach(9E)` entry point first calls the `ddi_get_instance(9F)` function to retrieve the instance number of the device information node. The `detach(9E)` entry point uses this instance number to call the `ddi_soft_state_free(9F)` function to destroy the state structure that was created by `ddi_soft_state_zalloc(9F)` in the `attach(9E)` entry point. The `detach(9E)` entry point then calls the `ddi_remove_minor_node(9F)` function to remove the device that was created by `ddi_create_minor_node(9F)` in the `attach(9E)` entry point.

## Opening the Device, Closing the Device, and Getting Module Information

The `open(9E)` and `close(9E)` entry points are identical in this sample driver. In each case, the entry point first calls the `getminor(9F)` function to retrieve the minor number of the device. Then each entry point uses this instance number to call the `ddi_get_soft_state(9F)` function to retrieve the state structure for this device instance. If no state structure is retrieved, an error code is returned. If a state structure is retrieved, the `open(9E)` and `close(9E)` entry points both verify the type of this device. If this device is not a character device, the `EINVAL` (invalid) error code is returned.

If the user wants device information for this device instance, the `getinfo(9E)` entry point returns the device information from the state structure. If the user wants the instance number of this device instance, the `getinfo(9E)` entry point uses the `getminor(9F)` function to return the minor number.

## Reading the Data

The `read(9E)` entry point first calls the `getminor(9F)` function to retrieve the minor number of the device. The `read(9E)` entry point uses this instance number to call the `ddi_get_soft_state(9F)` function to retrieve the state structure for this device instance. If no state structure is retrieved, `read(9E)` returns an error code. If a state structure is retrieved, `read(9E)` calls the `uiomove(9F)` function to copy the quotation from the driver to the `uio(9S)` I/O request structure.

## Checking Data Validity

Version 2 of the driver uses `ASSERT(9F)` statements to check the validity of data. If the asserted expression is true, the `ASSERT(9F)` statement does nothing. If the asserted expression is false, the `ASSERT(9F)` statement writes an error message to the console and causes the system to panic.

To use `ASSERT(9F)` statements, include the `sys/debug.h` header file in your source and define the `DEBUG` preprocessor symbol. If you do not define the `DEBUG` preprocessor symbol, then the `ASSERT(9F)` statements do nothing. Simply recompile to activate or inactivate `ASSERT(9F)` statements.

## Quote Of The Day Version 2 Source

Enter the source code shown in the following example into a text file named `qotd_2.c`.

### EXAMPLE 3-3 Quote Of The Day Version 2 Source File

```
#include <sys/types.h>
#include <sys/file.h>
#include <sys/errno.h>
#include <sys/open.h>
#include <sys/cred.h>
#include <sys/uio.h>
#include <sys/stat.h>
#include <sys/modctl.h>
#include <sys/conf.h>
#include <sys/devops.h>
#include <sys/debug.h>
#include <sys/cmn_err.h>
#include <sys/ddi.h>
#include <sys/sunddi.h>

#define QOTD_NAME      "qotd"
#define QOTD_MAXLEN   128

static const char qotd[QOTD_MAXLEN]
    = "You can't have everything. \
Where would you put it? - Steven Wright\n";

static void *qotd_state_head;

struct qotd_state {
    int          instance;
    dev_info_t   *devi;
};

static int qotd_getinfo(dev_info_t *, ddi_info_cmd_t, void *, void **);
static int qotd_attach(dev_info_t *, ddi_attach_cmd_t);
static int qotd_detach(dev_info_t *, ddi_detach_cmd_t);
static int qotd_open(dev_t *, int, int, cred_t *);
```



**EXAMPLE 3-3** Quote Of The Day Version 2 Source File (Continued)

```
static int qotd_close(dev_t, int, int, cred_t *);
static int qotd_read(dev_t, struct uio *, cred_t *);

static struct cb_ops qotd_cb_ops = {
    qotd_open,          /* cb_open */
    qotd_close,        /* cb_close */
    nodev,             /* cb_strategy */
    nodev,             /* cb_print */
    nodev,             /* cb_dump */
    qotd_read,        /* cb_read */
    nodev,             /* cb_write */
    nodev,             /* cb_ioctl */
    nodev,             /* cb_devmap */
    nodev,             /* cb_mmap */
    nodev,             /* cb_segmap */
    nochpoll,         /* cb_chpoll */
    ddi_prop_op,      /* cb_prop_op */
    (struct streamtab *)NULL, /* cb_str */
    D_MP | D_64BIT,   /* cb_flag */
    CB_REV,           /* cb_rev */
    nodev,            /* cb_aread */
    nodev,            /* cb_awrite */
};

static struct dev_ops qotd_dev_ops = {
    DEVO_REV,         /* devo_rev */
    0,                /* devo_refcnt */
    qotd_getinfo,    /* devo_getinfo */
    nulldev,         /* devo_identify */
    nulldev,         /* devo_probe */
    qotd_attach,     /* devo_attach */
    qotd_detach,     /* devo_detach */
    nodev,           /* devo_reset */
    &qotd_cb_ops,     /* devo_cb_ops */
    (struct bus_ops *)NULL, /* devo_bus_ops */
    nulldev,         /* devo_power */
};

static struct modldrv modldrv = {
    &mod_driverops,,
    "Quote of the Day 2.0",
    &qotd_dev_ops;};

static struct modlinkage modlinkage = {
    MODREV_1,
    (void *)&modldrv;,
    NULL
};

int
_init(void)
{
    int retval;
```

**EXAMPLE 3-3** Quote Of The Day Version 2 Source File (Continued)

```
        if ((retval = ddi_soft_state_init(&qotd_state_head,
            sizeof (struct qotd_state), 1)) != 0)
            return retval;
        if ((retval = mod_install(&modlinkage;)) != 0) {
            ddi_soft_state_fini(&qotd_state_head;);
            return (retval);
        }

        return (retval);
    }

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

int
_fini(void)
{
    int retval;

    if ((retval = mod_remove(&modlinkage;)) != 0)
        return (retval);
    ddi_soft_state_fini(&qotd_state_head;);

    return (retval);
}

/*ARGSUSED*/
static int
qotd_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg, void **resultp)
{
    struct qotd_state *qsp;
    int retval = DDI_FAILURE;

    ASSERT(resultp != NULL);

    switch (cmd) {
    case DDI_INFO_DEVT2DEVINFO:
        if ((qsp = ddi_get_soft_state(qotd_state_head,
            getminor((dev_t)arg))) != NULL) {
            *resultp = qsp->devi;
            retval = DDI_SUCCESS;
        } else
            *resultp = NULL;
        break;
    case DDI_INFO_DEVT2INSTANCE:
        *resultp = (void *)getminor((dev_t)arg);
        retval = DDI_SUCCESS;
        break;
    }
}
```

**EXAMPLE 3-3** Quote Of The Day Version 2 Source File (Continued)

```
        return (retval);
    }

static int
qotd_attach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    int instance = ddi_get_instance(dip);
    struct qotd_state *qsp;

    switch (cmd) {
    case DDI_ATTACH:
        if (ddi_soft_state_zalloc(qotd_state_head, instance)
            != DDI_SUCCESS) {
            cmn_err(CE_WARN, "Unable to allocate state for %d",
                instance);
            return (DDI_FAILURE);
        }
        if ((qsp = ddi_get_soft_state(qotd_state_head, instance))
            == NULL) {
            cmn_err(CE_WARN, "Unable to obtain state for %d",
                instance);
            ddi_soft_state_free(dip, instance);
            return (DDI_FAILURE);
        }
        if (ddi_create_minor_node(dip, QOTD_NAME, S_IFCHR, instance,
            DDI_PSEUDO, 0) != DDI_SUCCESS) {
            cmn_err(CE_WARN, "Cannot create minor node for %d",
                instance);
            ddi_soft_state_free(dip, instance);
            ddi_remove_minor_node(dip, NULL);
            return (DDI_FAILURE);
        }
        qsp->instance = instance;
        qsp->devi = dip;

        ddi_report_dev(dip);
        return (DDI_SUCCESS);
    case DDI_RESUME:
        return (DDI_SUCCESS);
    default:
        return (DDI_FAILURE);
    }
}

static int
qotd_detach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    int instance = ddi_get_instance(dip);

    switch (cmd) {
    case DDI_DETACH:
        ddi_soft_state_free(qotd_state_head, instance);
    }
}
```

**EXAMPLE 3-3** Quote Of The Day Version 2 Source File (Continued)

```
        ddi_remove_minor_node(dip, NULL);
        return (DDI_SUCCESS);
    case DDI_SUSPEND:
        return (DDI_SUCCESS);
    default:
        return (DDI_FAILURE);
    }
}

/*ARGSUSED*/
static int
gotd_open(dev_t *devp, int flag, int otyp, cred_t *credp)
{
    int instance = getminor(*devp);
    struct qotd_state *qsp;

    if ((qsp = ddi_get_soft_state(qotd_state_head, instance)) == NULL)
        return (ENXIO);

    ASSERT(qsp->instance == instance);

    if (otyp != OTYP_CHR)
        return (EINVAL);

    return (0);
}

/*ARGSUSED*/
static int
gotd_close(dev_t dev, int flag, int otyp, cred_t *credp)
{
    struct qotd_state *qsp;
    int instance = getminor(dev);

    if ((qsp = ddi_get_soft_state(qotd_state_head, instance)) == NULL)
        return (ENXIO);

    ASSERT(qsp->instance == instance);

    if (otyp != OTYP_CHR)
        return (EINVAL);

    return (0);
}

/*ARGSUSED*/
static int
gotd_read(dev_t dev, struct uio *uiop, cred_t *credp)
{
    struct qotd_state *qsp;
    int instance = getminor(dev);

    if ((qsp = ddi_get_soft_state(qotd_state_head, instance)) == NULL)
```

**EXAMPLE 3-3** Quote Of The Day Version 2 Source File (Continued)

```
        return (ENXIO);

        ASSERT(qsp->instance == instance);

        return (uiomove((void *)qotd, min(uiop->uio_resid, strlen(qotd)),
            UIO_READ, uiop));
    }
```

Enter the configuration information shown in the following example into a text file named `qotd_2.conf`.

**EXAMPLE 3-4** Quote Of The Day Version 2 Configuration File

```
name="qotd_2" parent="pseudo" instance=0;
```

## Building, Installing, and Using Quote Of The Day Version 2

Version 2 of the driver uses `ASSERT(9F)` statements to check the validity of data. To use `ASSERT(9F)` statements, include the `sys/debug.h` header file in your source and define the `DEBUG` preprocessor symbol.

Compile and link the driver. If you use `ASSERT(9F)` statements to check the validity of data, you must define the `DEBUG` preprocessor symbol:

```
% cc -D_KERNEL -DDEBUG -c qotd_2.c
% ld -r -o qotd_2 qotd_2.o
```

The following example shows compiling and linking for a 32-bit architecture if you are not using `ASSERT(9F)` statements:

```
% cc -D_KERNEL -c qotd_2.c
% ld -r -o qotd_2 qotd_2.o
```

Make sure you are user `root` when you install the driver.

Copy the driver binary to the `/tmp` directory as discussed in [“Building and Installing the Driver”](#) on page 48.

```
# cp qotd_2 /tmp
# ln -s /tmp/qotd_2 /usr/kernel/drv/qotd_2
```

Copy the configuration file to the kernel driver area of the system.

```
# cp qotd_2.conf /usr/kernel/drv
```

In a separate window, enter the following command:

```
% tail -f /var/adm/messages
```

Make sure you are user `root` when you load the driver. Use the `add_drv(1M)` command to load the driver:

```
# add_drv qotd_2
```

You should see the following messages in the window where you are viewing `/var/adm/messages`:

```
date time machine pseudo: [ID 129642 kern.info] pseudo-device: devinfo0
date time machine genunix: [ID 936769 kern.info] devinfo0 is /pseudo/devinfo@0
date time machine pseudo: [ID 129642 kern.info] pseudo-device: qotd_20
date time machine genunix: [ID 936769 kern.info] qotd_20 is /pseudo/qotd_2@0
```

When this version of the Quote Of The Day driver loads, it does not display its quotation. The `qotd_1` driver wrote a message to a system log through its `_init(9E)` entry point. This `qotd_2` driver stores its data and makes the data available through its `read(9E)` entry point.

You can use the `modinfo(1M)` command to display the module information for this version of the Quote Of The Day driver. The module name is the value you entered for the second member of the `modldr` structure. The value `96` is the major number of this module.

```
% modinfo | grep qotd
182 ed115948 754 96 1 qotd_2 (Quote of the Day 2.0)
% grep qotd /etc/name_to_major
qotd_1 94
qotd_2 96
```

This driver also is the most recent module listed by `prtconf(1M)` in the `pseudo` device section:

```
% prtconf -P | grep qotd
qotd_1, instance #0 (driver not attached)
qotd_2, instance #0
```

When you access this `qotd_2` device for reading, the command you use to access the device retrieves the data from the device node. The command then displays the data in the same way that the command displays any other input. To get the name of the device special file, look in the `/devices` directory:

```
% ls -l /devices/pseudo/qotd*
crw----- 1 root sys 96, 0 date time /devices/pseudo/qotd_2@0:qotd
```

This output shows that `qotd_2@0:qotd` is a character device. This listing also shows that only the `root` user has permission to read or write this device. Make sure you are user `root` when you test this driver. To test the `qotd_2` driver, you can use the `more(1)` command to access the device file for reading:

```
# more /devices/pseudo/qotd_2@0:qotd
You can't have everything. Where would you put it? - Steven Wright
You can't have everything. Where would you put it? - Steven Wright
```

---

# Modifying Data Stored in Kernel Memory

In this third version of the Quote Of The Day driver, the user can write to the data that is stored in kernel memory. The pseudo device that is created in this section is a pseudo-disk device or ramdisk device. A ramdisk device simulates a disk device by allocating kernel memory that is subsequently used as data storage. See `ramdisk(7D)` for more information about ramdisk devices.

As in Version 2 of the Quote Of The Day driver, this Version 3 driver stores its data and makes the data available through its `read(9E)` entry point. This Version 3 driver overwrites characters from the beginning of the data when the user writes to the device.

This section first discusses the important code differences between this version and the previous version of the Quote Of The Day driver. This section then shows you how you can modify and display the quotation.

In addition to changes in the driver, Quote Of The Day Version 3 introduces a header file and an auxiliary program. The header file is discussed in the following section. The utility program is discussed in [“Using Quote Of The Day Version 3” on page 91](#).

## Writing Quote Of The Day Version 3

This third version of the Quote Of The Day driver is more complex than the second version because this third version enables a user to change the text that is stored in the device.

This section first explains some important features of this version of the driver. This section then shows all the source for this driver, including the header file and the configuration file.

The following list summarizes the new features in Version 3 of the Quote Of The Day driver:

- Version 3 of the driver allocates and frees kernel memory.
- Version 3 uses condition variables and mutexes to manage thread synchronization.
- Version 3 copies data from user space to kernel space to enable the user to change the quotation.
- Version 3 adds two new entry points: `write(9E)` and `ioctl(9E)`.
- Version 3 adds a third new routine. The `gotd_rw()` routine is called by both the `read(9E)` entry point and the `write(9E)` entry point.

- As in Version 2, Version 3 of the driver uses the `uiomove(9F)` function to make the quotation available to the user. Version 3 uses the `ddi_copyin(9F)` function to copy the new quotation and the new device size from user space to kernel space. Version 3 uses the `ddi_copyout(9F)` function to report the current device size back to the user.
- Because the driver copies data between kernel space and user space, Version 3 of the driver uses the `ddi_model_convert_from(9F)` function to determine whether the data must be converted between 32-bit and 64-bit models. The 64-bit kernel supports both 64-bit and 32-bit user data.
- Version 3 defines one new constant to tell the driver whether the device is busy. Another new constant tells the driver whether the quotation has been modified. Version 3 defines four new constants to help the driver undo everything it has done.
- Version 3 includes a separate utility program to test the driver's I/O controls.

The following sections provide more detail about the additions and changes in Version 3 of the Quote Of The Day driver. The `dev_ops(9S)` structure and the `modlinkage(9S)` structure are the same as they were in Version 2 of the driver. The `modldrv(9S)` structure has not changed except for the version number of the driver. The `_init(9E)`, `_info(9E)`, `_fini(9E)`, `getinfo(9E)`, `open(9E)`, and `close(9E)` functions are the same as in Version 2 of the driver.

## Attaching, Allocating Memory, and Initializing a Mutex and a Condition Variable

The `gotd_attach()` entry point first allocates and gets the device soft state. The `gotd_attach()` routine then creates a minor node. All of this code is the same as in Version 2 of the Quote Of The Day driver. If the call to `ddi_create_minor_node(9F)` is successful, the `gotd_attach()` routine sets the `QOTD_DIDMINOR` flag in the new `flags` member of the `gotd_state` state structure.

Version 3 of the Quote Of The Day driver defines four new constants that keep track of four different events. A routine can test these flags to determine whether to deallocate, close, or remove resources. All four of these flags are set in the `gotd_attach()` entry point. All four of these conditions are checked in the `gotd_detach()` entry point, and the appropriate action is taken for each condition.

Note that operations are undone in the `gotd_detach()` entry point in the opposite order in which they were done in the `gotd_attach()` entry point. The `gotd_attach()` routine creates a minor node, allocates memory for the quotation, initializes a mutex, and initializes a condition variable. The `gotd_detach()` routine destroys the condition variable, destroys the mutex, frees the memory, and removes the minor node.



After the minor node is created, the `gotd_attach()` routine allocates memory for the quotation. For information on allocating and freeing memory in this driver, see [“Allocating and Freeing Kernel Memory” on page 74](#). If memory is allocated, the `gotd_attach()` routine sets the `QOTD_DIDALLOC` flag in the `flags` member of the state structure.

The `gotd_attach()` routine then calls the `mutex_init(9F)` function to initialize a mutex. If this operation is successful, the `gotd_attach()` routine sets the `QOTD_DIDMUTEX` flag. The `gotd_attach()` routine then calls the `cv_init(9F)` function to initialize a condition variable. If this operation is successful, the `gotd_attach()` routine sets the `QOTD_DIDCV` flag.

The `gotd_attach()` routine then calls the `strncpy(9F)` function to copy the initial quotation string to the new quotation member of the device state structure. Note that the `strncpy(9F)` function is used instead of the `strncpy(9F)` function. The `strncpy(9F)` function can be wasteful because it always copies *n* characters, even if the destination is smaller than *n* characters. Try using `strncpy(9F)` instead of `strlcpy(9F)` and note the difference in the behavior of the driver.

Finally, the initial quotation length is copied to the new quotation length member of the state structure. The remainder of the `gotd_attach()` routine is the same as in Version 2.

## Checking for Changes, Cleaning Up, and Detaching

The `gotd_detach()` routine is almost all new. The `gotd_detach()` routine must first get the soft state because the `gotd_detach()` routine needs to check the `flags` member of the state structure.

The first flag the `gotd_detach()` routine checks is the `QOTD_CHANGED` flag. The `QOTD_CHANGED` flag indicates whether the device is in the initial state. The `QOTD_CHANGED` flag is set in the `gotd_rw()` routine and in the `gotd_ioctl()` entry point. The `QOTD_CHANGED` flag is set any time the user does anything to the device other than simply inspect the device. If the `QOTD_CHANGED` flag is set, the size or content of the storage buffer has been modified. See [“Writing New Data” on page 78](#) for more information on the `QOTD_CHANGED` flag. When the `QOTD_CHANGED` flag is set, the detach operation fails because the device might contain data that is valuable to the user and the device should not be removed. If the `QOTD_CHANGED` flag is set, the `gotd_detach()` routine returns an error that the device is busy.

Once the quotation has been modified, the device cannot be detached until the user runs the `gotdctl` command with the `-r` option. The `-r` option reinitializes the quotation and indicates that the user is no longer interested in the contents of the device. See [“Exercising the Driver’s I/O Controls” on page 92](#) for more information about the `gotdctl` command.

The `gotd_detach()` routine then checks the four flags that were set in the `gotd_attach()` routine. If the `QOTD_DIDCV` flag is set, the `gotd_detach()` routine calls the `cv_destroy(9F)` function. If the `QOTD_DIDMUTEX` flag is set, the

`gotd_detach()` routine calls the `mutex_destroy(9F)` function. If the `QOTD_DIDALLOC` flag is set, the `gotd_detach()` routine calls the `ddi_umem_free(9F)` function. Finally, if the `QOTD_DIDMINOR` flag is set, the `gotd_detach()` routine calls the `ddi_remove_minor_node(9F)` function.

## Allocating and Freeing Kernel Memory

One of the new members of the device state structure supports memory allocation and deallocation. The `gotd_cookie` member receives a value from the `ddi_umem_alloc(9F)` function. The `gotd_cookie` value is then used by the `ddi_umem_free(9F)` function to free the memory.

Version 3 of the Quote Of The Day driver allocates kernel memory in three places:

- After the minor node is created
- In the `QOTDIOCSSZ` case of the `gotd_ioctl()` entry point
- In the `QOTDIOCDISCARD` case of the `gotd_ioctl()` entry point

The `gotd_attach()` routine allocates memory after the minor node is created. Memory must be allocated to enable the user to modify the quotation. The `gotd_attach()` routine calls the `ddi_umem_alloc(9F)` function with the `DDI_UMEM_NOSLEEP` flag so that the `ddi_umem_alloc(9F)` function will return immediately. If the requested amount of memory is not available, `ddi_umem_alloc(9F)` returns `NULL` immediately and does not wait for memory to become available. If no memory is allocated, `gotd_attach()` calls `gotd_detach()` and returns an error. If memory is allocated, `gotd_attach()` sets the `QOTD_DIDALLOC` flag so that this memory will be freed by `gotd_detach()` later.

The second place the driver allocates memory is in the `QOTDIOCSSZ` case of the `gotd_ioctl()` entry point. The `QOTDIOCSSZ` case sets a new size for the device. A new size is set when the user runs the `gotdctl` command with the `-s` option. See [“Exercising the Driver’s I/O Controls” on page 92](#) for more information about the `gotdctl` command. This time, the `ddi_umem_alloc(9F)` function is called with the `DDI_UMEM_SLEEP` flag so that `ddi_umem_alloc(9F)` will wait for the requested amount of memory to be available. When the `ddi_umem_alloc(9F)` function returns, the requested memory has been allocated.

Note that you cannot always use the `DDI_UMEM_SLEEP` flag. See the `CONTEXT` sections of the `ddi_umem_alloc(9F)`, `kmem_alloc(9F)`, and `kmem_zalloc(9F)` man pages. Also note the behavioral differences among these three functions. The `ddi_umem_alloc(9F)` function is used in this `gotd_3` driver because `ddi_umem_alloc(9F)` allocates whole pages of memory. The `kmem_zalloc(9F)` function might save memory because it might allocate smaller chunks of memory. This `gotd_3` driver demonstrates a ramdisk device. In a production ramdisk device, you would use `ddi_umem_alloc(9F)` to allocate page-aligned memory.

After the current quotation is copied to the new space, the `gotd_ioctl()` routine calls the `ddi_umem_free(9F)` function to free the memory that was previously allocated.

The third place the driver allocates memory is in the QOTDIOCDISCARD case of the `gotd_ioctl()` entry point. The QOTDIOCDISCARD case is called from the `gotdctl` command. The `gotdctl` command with the `-r` option sets the quotation back to its initial value. If the number of bytes allocated for the current quotation is different from the initial number of bytes, then new memory is allocated to reinitialize the quotation. Again, the `DDI_UMEM_SLEEP` flag is used so that when the `ddi_umem_alloc(9F)` function returns, the requested memory has been allocated. The `gotd_ioctl()` routine then calls the `ddi_umem_free(9F)` function to free the memory that was previously allocated.

## Managing Thread Synchronization

The Quote Of The Day Version 3 driver uses condition variables and mutual exclusion locks (mutexes) together to manage thread synchronization. See the *Multithreaded Programming Guide* for more information about mutexes, condition variables, and thread synchronization.

In this driver, the mutex and condition variable both are initialized in the `gotd_attach()` entry point and destroyed in the `gotd_detach()` entry point. The condition variable is tested in the `gotd_rw()` routine and in the `gotd_ioctl()` entry point.

The condition variable waits on the `QOTD_BUSY` condition. This condition is needed because the driver must do some operations that rely on exclusive access to internal structures without holding a lock. Accessing the storage buffer or its metadata requires mutual exclusion, but the driver cannot hold a lock if the operation might sleep. Instead of holding a lock in this case, the driver waits on the `QOTD_BUSY` condition.

The driver acquires a mutex when the driver tests the condition variable and when the driver accesses the storage buffer. The mutex protects the storage buffer. Failure to use a mutual exclusion when accessing the storage buffer could allow one user process to resize the buffer while another user process tries to read the buffer, for example. The result of unprotected buffer access could be data corruption or a panic.

The condition variable is used when functions are called that may need to sleep. The `ddi_copyin(9F)`, `ddi_copyout(9F)`, and `uiomove(9F)` functions can sleep. Memory allocation can sleep if you use the `SLEEP` flag. Functions must not hold a mutex while they are sleeping. Sleeping while holding a mutex can cause deadlock. When a function might sleep, set the `QOTD_BUSY` flag and take the condition variable, which drops the mutex. To avoid race conditions, the `QOTD_BUSY` flag can be set or cleared only when holding the mutex. For more information on deadlock, see “Using Mutual Exclusion Locks” in *Multithreaded Programming Guide* and “Avoiding Deadlock” in *Multithreaded Programming Guide*.

## Locking Rules for Quote Of The Day Version 3

The locking rules for this `qotd_3` driver are as follows:

1. You must have exclusive access to do any of the following operations. To have exclusive access, you must own the mutex or you must set `QOTD_BUSY`. Threads must wait on `QOTD_BUSY`.
  - Test the contents of the storage buffer.
  - Modify the contents of the storage buffer.
  - Modify the size of the storage buffer.
  - Modify variables that refer to the address of the storage buffer.
2. If your operation does not need to sleep, do the following actions:
  - a. Acquire the mutex.
  - b. Wait until `QOTD_BUSY` is cleared. When the thread that set `QOTD_BUSY` clears `QOTD_BUSY`, that thread also should signal threads waiting on the condition variable and then drop the mutex.
  - c. Perform your operation. You do not need to set `QOTD_BUSY` before you perform your operation.
  - d. Drop the mutex.

The following code sample illustrates this rule:

```
mutex_enter(&qsp->lock);
while (qsp->flags & QOTD_BUSY) {
    if (cv_wait_sig(&qsp->cv, &qsp->lock) == 0) {
        mutex_exit(&qsp->lock);
        ddi_umem_free(new_cookie);
        return (EINTR);
    }
}
memcpy(new_qotd, qsp->qotd, min(qsp->qotd_len, new_len));
ddi_umem_free(qsp->qotd_cookie);
qsp->qotd = new_qotd;
qsp->qotd_cookie = new_cookie;
qsp->qotd_len = new_len;
qsp->flags |= QOTD_CHANGED;
mutex_exit(&qsp->lock);
```

3. If your operation must sleep, do the following actions:
  - a. Acquire the mutex.
  - b. Set `QOTD_BUSY`.
  - c. Drop the mutex.
  - d. Perform your operation.
  - e. Reacquire the mutex.
  - f. Signal any threads waiting on the condition variable.
  - g. Drop the mutex.

These locking rules are very simple. These three rules ensure consistent access to the buffer and its metadata. Realistic drivers probably have more complex locking requirements. For example, drivers that use ring buffers or drivers that manage multiple register sets or multiple devices have more complex locking requirements.

### *Lock and Condition Variable Members of the State Structure*

The device state structure for Version 3 of the Quote Of The Day driver contains two new members to help manage thread synchronization:

- The `lock` member is used to acquire and exit mutexes for the current instance of the device. The `lock` member is an argument to each `mutex(9F)` function call. The `lock` member also is an argument to the `cv_wait_sig(9F)` function call. In the `cv_wait_sig(9F)` function call, the `lock` value ensures that the condition will not be changed before the `cv_wait_sig(9F)` function returns.
- The `cv` member is a condition variable. The `cv` member is an argument to each `condvar(9F)` (`cv_`) function call.

### *Creating and Destroying Locks and Condition Variables*

Version 3 of the Quote Of The Day driver defines two constants to make sure the mutex and condition variable are destroyed when the driver is finished with them. The driver uses these constants to set and reset the new `flags` member of the device state structure.

- The `QOTD_DIDMUTEX` flag is set in the `gotd_attach()` entry point immediately after a successful call to `mutex_init(9F)`. If the `QOTD_DIDMUTEX` flag is set when the `gotd_detach()` entry point is called, the `gotd_detach()` entry point calls the `mutex_destroy(9F)` function.
- The `QOTD_DIDCV` flag is set in the `gotd_attach()` entry point immediately after a successful call to `cv_init(9F)`. If the `QOTD_DIDCV` flag is set when the `gotd_detach()` entry point is called, the `gotd_detach()` entry point calls the `cv_destroy(9F)` function.

### *Waiting on Signals*

In the `gotd_rw()` and `gotd_ioctl()` routines, the `cv_wait_sig(9F)` calls `wait` until the condition variable is signaled to proceed or until a `signal(3C)` is received. Either the `cv_signal(9F)` function or the `cv_broadcast(9F)` function signals the `cv` condition variable to proceed.

A thread can wait on a condition variable until either the condition variable is signaled or a `signal(3C)` is received by the process. The `cv_wait(9F)` function waits until the condition variable is signaled but ignores `signal(3C)` signals. This driver uses the `cv_wait_sig(9F)` function instead of the `cv_wait(9F)` function because this driver responds if a signal is received by the process performing the operation. If a

`signal(3C)` is taken by the process, this driver returns an interrupt error and does not complete the operation. The `cv_wait_sig(9F)` function usually is preferred to the `cv_wait(9F)` function because this implementation offers the user program more precise response. The `signal(3C)` causes an effect closer to the point at which the process was executing when the `signal(3C)` was received.

## Writing New Data

The `cb_ops(9S)` structure for Version 3 of the Quote Of The Day driver declares two new entry points that support modifying the quotation. The two new entry points are `write(9E)` and `ioctl(9E)`. The `gotd_rw()` routine is a third new routine in Version 3 of the driver. The `gotd_rw()` routine is called by both the `read(9E)` entry point and the `write(9E)` entry point.

The device state structure for Version 3 of the Quote Of The Day driver contains two new members that are used to modify the quotation. The `gotd` string holds the quotation for the current instance of the device. The `gotd_len` member holds the length in bytes of the current quotation.

Version 3 of the driver also defines two new constants that support modifying the quotation. In place of `QOTD_MAXLEN`, Version 3 of the driver defines `QOTD_MAX_LEN`. `QOTD_MAX_LEN` is used in the `gotd_ioctl()` entry point to test whether the user has entered a string that is too long. Version 3 of the driver also defines `QOTD_CHANGED`. The `QOTD_CHANGED` flag is set in the `gotd_rw()` routine and in the `gotd_ioctl()` entry point when a new quotation is copied from the user.

When the `gotd_3` device is opened for writing, the kernel calls the `gotd_write()` entry point. The `gotd_write()` entry point then calls the `gotd_rw()` routine and passes a `UIO_WRITE` flag. The new `gotd_read()` entry point is exactly the same as the `gotd_write()` entry point, except that the `gotd_read()` entry point passes a `UIO_READ` flag. The `gotd_rw()` routine supports both reading and writing the device and thereby eliminates much duplicate code.

The `gotd_rw()` routine first gets the device soft state. Then the `gotd_rw()` routine checks the length of the I/O request in the `uio(9S)` I/O request structure. If this length is zero, the `gotd_rw()` routine returns zero. If this length is not zero, the `gotd_rw()` routine enters a mutex.

While the device is busy, the `gotd_rw()` routine checks whether the condition variable has been signaled or a `signal(3C)` is pending. If either of these conditions is true, the `gotd_rw()` routine exits the mutex and returns an error.

When the device is not busy, the `gotd_rw()` routine checks whether the data offset in the `uio(9S)` I/O request structure is valid. If the offset is not valid, the `gotd_rw()` routine exits the mutex and returns an error. If the offset is valid, the local length variable is set to the difference between the offset in the I/O request structure and the length in the device state structure. If this difference is zero, the `gotd_rw()` routine exits the mutex and returns. If the device was opened for writing, the `gotd_rw()` routine returns a space error. Otherwise, the `gotd_rw()` routine returns zero.

The `gotd_rw()` routine then sets the `QOTD_BUSY` flag in the `flags` member of the device state structure and exits the mutex. The `gotd_rw()` routine then calls the `uiomove(9F)` function to copy the quotation. If the `rw` argument is `UIO_READ`, then the quotation is transferred from the state structure to the I/O request structure. If the `rw` argument is `UIO_WRITE`, then the quotation is transferred from the I/O request structure to the state structure.

The `gotd_rw()` routine then enters a mutex again. If the device was opened for writing, the `gotd_rw()` routine sets the `QOTD_CHANGED` flag. The `gotd_rw()` routine then sets the device to not busy, calls `cv_broadcast(9F)` to unblock any threads that were blocked on this condition variable, and exits the mutex.

Finally, the `gotd_rw()` routine returns the quotation. The quotation is written to the device node.

## Reporting and Setting Device Size and Re-initializing the Device

The behavior of the `ioctl(9E)` entry point depends on the command value passed in to the entry point. These constants are defined in the new `gotd.h` header file. The `gotd_ioctl()` routine reports the size of the space allocated for the quotation, sets a new amount of space to allocate for the quotation, or resets the quotation back to its initial value.

If the request is to report the size of the space allocated for the quotation, then the `gotd_ioctl()` routine first sets a local size variable to the value of the quotation length in the state structure. If the device was not opened for reading, the `gotd_ioctl()` routine returns an error.

Because the `gotd_ioctl()` routine transfers data between kernel space and user space, the `gotd_ioctl()` routine must check whether both spaces are using the same data model. If the return value of the `ddi_model_convert_from(9F)` function is `DDI_MODEL_ILP32`, then the driver must convert to 32-bit data before calling `ddi_copyout(9F)` to transfer the current size of the quotation space. If the return value of the `ddi_model_convert_from(9F)` function is `DDI_MODEL_NONE`, then no data type conversion is necessary.

If the request is to set a new size for the space allocated for the quotation, then the `gotd_ioctl()` routine first sets local variables for the new size, the new quotation, and a new memory allocation cookie. If the device was not opened for writing, the `gotd_ioctl()` routine returns an error.

The `gotd_ioctl()` routine then checks again for data model mismatch. If the return value of the `ddi_model_convert_from(9F)` function is `DDI_MODEL_ILP32`, then the driver declares a 32-bit size variable to receive the new size from `ddi_copyin(9F)`. When the new size is received, the size is converted to the data type of the kernel space.

If the new size is zero or is greater than `QOTD_MAX_LEN`, the `gotd_ioctl()` routine returns an error. If the new size is valid, then the `gotd_ioctl()` routine allocates new memory for the quotation and enters a mutex.

While the device is busy, the `gotd_ioctl()` routine checks whether the condition variable has been signaled or a `signal(3C)` is pending. If either of these conditions is true, the `gotd_ioctl()` routine exits the mutex, frees the new memory it allocated, and returns an error.

When the device is not busy, the `gotd_ioctl()` routine uses `memcpy(9F)` to copy the quotation from the driver's state structure to the new space. The `gotd_ioctl()` routine then frees the memory currently pointed to by the state structure, and updates the state structure members to the new values. The `gotd_ioctl()` routine then sets the `QOTD_CHANGED` flag, exits the mutex, and returns.

If the request is to discard the current quotation and reset to the initial quotation, then the `gotd_ioctl()` routine first sets local variables for the new quotation and a new memory allocation cookie. If the device was not opened for writing, the `gotd_ioctl()` routine returns an error. If the space allocated for the current quotation is different from the space allocated for the initial quotation, then the `gotd_ioctl()` routine allocates new memory that is the size of the initial space and enters a mutex.

While the device is busy, the `gotd_ioctl()` routine checks whether the condition variable has been signaled or a `signal(3C)` is pending. If either of these conditions is true, the `gotd_ioctl()` routine exits the mutex, frees the new memory it allocated, and returns an error.

When the device is not busy, the `gotd_ioctl()` routine frees the memory currently pointed to by the state structure, updates the memory state structure members to the new values, and resets the length to its initial value. If the size of the current quotation space was the same as the initial size and no new memory was allocated, then `gotd_ioctl()` calls `bzero(9F)` to clear the current quotation. The `gotd_ioctl()` routine then calls the `strncpy(9F)` function to copy the initial quotation string to the quotation member of the state structure. The `gotd_ioctl()` routine then unsets the `QOTD_CHANGED` flag, exits the mutex, and returns.

Once the `QOTD_CHANGED` flag has been set, the only way to unset it is to run the `gotdctl` command with the `-r` option. See [“Exercising the Driver's I/O Controls” on page 92](#) for more information about the `gotdctl` command.

## Quote Of The Day Version 3 Source

Enter the source code shown in the following example into a text file named `gotd_3.c`.

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File

```
#include <sys/types.h>
#include <sys/file.h>
```



**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
#include <sys/errno.h>
#include <sys/open.h>
#include <sys/cred.h>
#include <sys/uio.h>
#include <sys/stat.h>
#include <sys/ksynch.h>
#include <sys/modctl.h>
#include <sys/conf.h>
#include <sys/devops.h>
#include <sys/debug.h>
#include <sys/cmn_err.h>
#include <sys/ddi.h>
#include <sys/sunddi.h>

#include "qotd.h"

#define QOTD_NAME      "qotd_3"

static const char init_qotd[]
    = "On the whole, I'd rather be in Philadelphia. - W. C. Fields\n";
static const size_t init_qotd_len = 128;

#define QOTD_MAX_LEN    65536          /* Maximum quote in bytes */
#define QOTD_CHANGED    0x1           /* User has made modifications */
#define QOTD_DIDMINOR    0x2           /* Created minors */
#define QOTD_DIDALLOC    0x4           /* Allocated storage space */
#define QOTD_DIDMUTEX    0x8           /* Created mutex */
#define QOTD_DIDCV       0x10          /* Created cv */
#define QOTD_BUSY       0x20          /* Device is busy */

static void *qotd_state_head;

struct qotd_state {
    int          instance;
    dev_info_t   *devi;
    kmutex_t     lock;
    kcondvar_t   cv;
    char         *qotd;
    size_t       qotd_len;
    ddi_umem_cookie_t qotd_cookie;
    int          flags;
};

static int qotd_getinfo(dev_info_t *, ddi_info_cmd_t, void *, void **);
static int qotd_attach(dev_info_t *, ddi_attach_cmd_t);
static int qotd_detach(dev_info_t *, ddi_detach_cmd_t);
static int qotd_open(dev_t *, int, int, cred_t *);
static int qotd_close(dev_t, int, int, cred_t *);
static int qotd_read(dev_t, struct uio *, cred_t *);
static int qotd_write(dev_t, struct uio *, cred_t *);
static int qotd_rw(dev_t, struct uio *, enum uio_rw);
static int qotd_ioctl(dev_t, int, intptr_t, int, cred_t *, int *);
```

EXAMPLE 3-5 Quote Of The Day Version 3 Source File (Continued)

```
static struct cb_ops qotd_cb_ops = {
    qotd_open,          /* cb_open */
    qotd_close,        /* cb_close */
    nodev,             /* cb_strategy */
    nodev,             /* cb_print */
    nodev,             /* cb_dump */
    qotd_read,         /* cb_read */
    qotd_write,        /* cb_write */
    qotd_ioctl,        /* cb_ioctl */
    nodev,             /* cb_devmap */
    nodev,             /* cb_mmap */
    nodev,             /* cb_segmap */
    nochpoll,         /* cb_chpoll */
    ddi_prop_op,       /* cb_prop_op */
    (struct streamtab *)NULL, /* cb_str */
    D_MP | D_64BIT,    /* cb_flag */
    CB_REV,            /* cb_rev */
    nodev,             /* cb_aread */
    nodev,             /* cb_awrite */
};

static struct dev_ops qotd_dev_ops = {
    DEVO_REV,          /* devo_rev */
    0,                 /* devo_refcnt */
    qotd_getinfo,      /* devo_getinfo */
    nulldev,           /* devo_identify */
    nulldev,           /* devo_probe */
    qotd_attach,       /* devo_attach */
    qotd_detach,       /* devo_detach */
    nodev,             /* devo_reset */
    &qotd_cb_ops,       /* devo_cb_ops */
    (struct bus_ops *)NULL, /* devo_bus_ops */
    nulldev            /* devo_power */
};

static struct modldrv modldrv = {
    &mod_driverops;,
    "Quote of the day 3.0",
    &qotd_dev_ops;};

static struct modlinkage modlinkage = {
    MODREV_1,
    (void *)&modldrv;,
    NULL
};

int
_init(void)
{
    int retval;

    if ((retval = ddi_soft_state_init(&qotd_state_head;,
        sizeof (struct qotd_state), 1)) != 0)
```

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
        return retval;
    if ((retval = mod_install(&modlinkage;)) != 0) {
        ddi_soft_state_fini(&qotd_state_head;);
        return (retval);
    }

    return (retval);
}

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

int
_fini(void)
{
    int retval;

    if ((retval = mod_remove(&modlinkage;)) != 0)
        return (retval);
    ddi_soft_state_fini(&qotd_state_head;);

    return (retval);
}

/*ARGSUSED*/
static int
qotd_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void *arg, void **resultp)
{
    struct qotd_state *qsp;
    int retval = DDI_FAILURE;

    ASSERT(resultp != NULL);

    switch (cmd) {
    case DDI_INFO_DEVT2DEVINFO:
        if ((qsp = ddi_get_soft_state(qotd_state_head,
            getminor((dev_t)arg))) != NULL) {
            *resultp = qsp->devi;
            retval = DDI_SUCCESS;
        } else
            *resultp = NULL;
        break;
    case DDI_INFO_DEVT2INSTANCE:
        *resultp = (void *)getminor((dev_t)arg);
        retval = DDI_SUCCESS;
        break;
    }

    return (retval);
}
```

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
static int
gotd_attach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    int instance = ddi_get_instance(dip);
    struct gotd_state *qsp;

    switch (cmd) {
    case DDI_ATTACH:
        if (ddi_soft_state_zalloc(gotd_state_head, instance)
            != DDI_SUCCESS) {
            cmn_err(CE_WARN, "Unable to allocate state for %d",
                instance);
            return (DDI_FAILURE);
        }
        if ((qsp = ddi_get_soft_state(gotd_state_head, instance))
            == NULL) {
            cmn_err(CE_WARN, "Unable to obtain state for %d",
                instance);
            ddi_soft_state_free(dip, instance);
            return (DDI_FAILURE);
        }
        if (ddi_create_minor_node(dip, QOTD_NAME, S_IFCHR, instance,
            DDI_PSEUDO, 0) != DDI_SUCCESS) {
            cmn_err(CE_WARN, "Unable to create minor node for %d",
                instance);
            (void)gotd_detach(dip, DDI_DETACH);
            return (DDI_FAILURE);
        }
        qsp->flags |= QOTD_DIDMINOR;
        qsp->qotd = ddi_umem_alloc(init_qotd_len, DDI_UMEM_NOSLEEP,
            &qsp->qotd_cookie);
        if (qsp->qotd == NULL) {
            cmn_err(CE_WARN, "Unable to allocate storage for %d",
                instance);
            (void)gotd_detach(dip, DDI_DETACH);
            return (DDI_FAILURE);
        }
        qsp->flags |= QOTD_DIDALLOC;
        mutex_init(&qsp->lock, NULL, MUTEX_DRIVER, NULL);
        qsp->flags |= QOTD_DIDMUTEX;
        cv_init(&qsp->cv, NULL, CV_DRIVER, NULL);
        qsp->flags |= QOTD_DIDCV;

        (void)strncpy(qsp->qotd, init_qotd, init_qotd_len);
        qsp->qotd_len = init_qotd_len;
        qsp->instance = instance;
        qsp->devi = dip;

        ddi_report_dev(dip);
        return (DDI_SUCCESS);
    case DDI_RESUME:
        return (DDI_SUCCESS);
    }
```

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
        default:
            return (DDI_FAILURE);
    }
}

static int
qotd_detach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    int instance = ddi_get_instance(dip);
    struct qotd_state *qsp;

    switch (cmd) {
    case DDI_DETACH:
        qsp = ddi_get_soft_state(qotd_state_head, instance);
        if (qsp != NULL) {
            ASSERT(!(qsp->flags & QOTD_BUSY));
            if (qsp->flags & QOTD_CHANGED)
                return (EBUSY);
            if (qsp->flags & QOTD_DIDCV)
                cv_destroy(&qsp->cv);
            if (qsp->flags & QOTD_DIDMUTEX)
                mutex_destroy(&qsp->lock);
            if (qsp->flags & QOTD_DIDALLOC) {
                ASSERT(qsp->qotd != NULL);
                ddi_umem_free(qsp->qotd_cookie);
            }
            if (qsp->flags & QOTD_DIDMINOR)
                ddi_remove_minor_node(dip, NULL);
        }
        ddi_soft_state_free(qotd_state_head, instance);
        return (DDI_SUCCESS);
    case DDI_SUSPEND:
        return (DDI_SUCCESS);
    default:
        return (DDI_FAILURE);
    }
}

/*ARGSUSED*/
static int
qotd_open(dev_t *devp, int flag, int otyp, cred_t *credp)
{
    int instance = getminor(*devp);
    struct qotd_state *qsp;

    if ((qsp = ddi_get_soft_state(qotd_state_head, instance)) == NULL)
        return (ENXIO);

    ASSERT(qsp->instance == instance);

    if (otyp != OTYP_CHR)
        return (EINVAL);
}
```

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File    *(Continued)*

```
        return (0);
    }

/*ARGSUSED*/
static int
qotd_close(dev_t dev, int flag, int otyp, cred_t *credp)
{
    struct qotd_state *qsp;
    int instance = getminor(dev);

    if ((qsp = ddi_get_soft_state(qotd_state_head, instance)) == NULL)
        return (ENXIO);

    ASSERT(qsp->instance == instance);

    if (otyp != OTYP_CHR)
        return (EINVAL);

    return (0);
}

/*ARGSUSED*/
static int
qotd_read(dev_t dev, struct uio *uiop, cred_t *credp)
{
    return qotd_rw(dev, uiop, UIO_READ);
}

/*ARGSUSED*/
static int
qotd_write(dev_t dev, struct uio *uiop, cred_t *credp)
{
    return qotd_rw(dev, uiop, UIO_WRITE);
}

static int
qotd_rw(dev_t dev, struct uio *uiop, enum uio_rw rw)
{
    struct qotd_state *qsp;
    int instance = getminor(dev);
    size_t len = uiop->uio_resid;
    int retval;

    if ((qsp = ddi_get_soft_state(qotd_state_head, instance)) == NULL)
        return (ENXIO);

    ASSERT(qsp->instance == instance);

    if (len == 0)
        return (0);

    mutex_enter(&qsp->lock);
```

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
while (qsp->flags & QOTD_BUSY) {
    if (cv_wait_sig(&qsp->cv, &qsp->lock) == 0) {
        mutex_exit(&qsp->lock);
        return (EINTR);
    }
}

if (uiop->uio_offset < 0 || uiop->uio_offset > qsp->qotd_len) {
    mutex_exit(&qsp->lock);
    return (EINVAL);
}

if (len > qsp->qotd_len - uiop->uio_offset)
    len = qsp->qotd_len - uiop->uio_offset;

if (len == 0) {
    mutex_exit(&qsp->lock);
    return (rw == UIO_WRITE ? ENOSPC : 0);
}

qsp->flags |= QOTD_BUSY;
mutex_exit(&qsp->lock);

retval = uiomove((void *) (qsp->qotd + uiop->uio_offset), len, rw, uiop);

mutex_enter(&qsp->lock);
if (rw == UIO_WRITE)
    qsp->flags |= QOTD_CHANGED;
qsp->flags &= ~QOTD_BUSY;
cv_broadcast(&qsp->cv);
mutex_exit(&qsp->lock);

return (retval);
}

/*ARGSUSED*/
static int
qotd_ioctl(dev_t dev, int cmd, intptr_t arg, int mode, cred_t *credp,
           int *rvalp)
{
    struct qotd_state *qsp;
    int instance = getminor(dev);

    if ((qsp = ddi_get_soft_state(qotd_state_head, instance)) == NULL)
        return (ENXIO);

    ASSERT(qsp->instance == instance);

    switch (cmd) {
    case QOTDIOGSZ: {
        /* We are not guaranteed that ddi_copyout(9F) will read
         * atomically anything larger than a byte. Therefore we

```

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
    * must duplicate the size before copying it out to the user.
    */
    size_t sz = qsp->qotd_len;

    if (!(mode & FREAD))
        return (EACCES);

#ifdef _MULTI_DATAMODEL
    switch (ddi_model_convert_from(mode & FMODELS)) {
    case DDI_MODEL_ILP32: {
        size32_t sz32 = (size32_t)sz;
        if (ddi_copyout(&sz32;, (void *)arg, sizeof (size32_t),
            mode) != 0)
            return (EFAULT);
        return (0);
    }
    case DDI_MODEL_NONE:
        if (ddi_copyout(&sz;, (void *)arg, sizeof (size_t),
            mode) != 0)
            return (EFAULT);
        return (0);
    default:
        cmn_err(CE_WARN, "Invalid data model %d in ioctl\n",
            ddi_model_convert_from(mode & FMODELS));
        return (ENOTSUP);
    }
#else /* !_MULTI_DATAMODEL */
    if (ddi_copyout(&sz;, (void *)arg, sizeof (size_t), mode) != 0)
        return (EFAULT);
    return (0);
#endif /* !_MULTI_DATAMODEL */
}
case QOTDIOCSSZ: {
    size_t new_len;
    char *new_qotd;
    ddi_umem_cookie_t new_cookie;
    uint_t model;

    if (!(mode & FWRITE))
        return (EACCES);

#ifdef _MULTI_DATAMODEL
    model = ddi_model_convert_from(mode & FMODELS);

    switch (model) {
    case DDI_MODEL_ILP32: {
        size32_t sz32;
        if (ddi_copyin((void *)arg, &sz32;, sizeof (size32_t),
            mode) != 0)
            return (EFAULT);
        new_len = (size_t)sz32;
        break;
    }

```



**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
case DDI_MODEL_NONE:
    if (ddi_copyin((void *)arg, &new_len;, sizeof (size_t),
        mode) != 0)
        return (EFAULT);
    break;
default:
    cmn_err(CE_WARN, "Invalid data model %d in ioctl\n",
        mode);
    return (ENOTSUP);
}
#else /* ! _MULTI_DATAMODEL */
    if (ddi_copyin((void *)arg, &new_len;, sizeof (size_t),
        mode) != 0)
        return (EFAULT);
#endif /* _MULTI_DATAMODEL */

    if (new_len == 0 || new_len > QOTD_MAX_LEN)
        return (EINVAL);

    new_qotd = ddi_umem_alloc(new_len, DDI_UMEM_SLEEP, &new_cookie);

    mutex_enter(&qsp->lock);
    while (qsp->flags & QOTD_BUSY) {
        if (cv_wait_sig(&qsp->cv, &qsp->lock) == 0) {
            mutex_exit(&qsp->lock);
            ddi_umem_free(new_cookie);
            return (EINTR);
        }
    }
    memcpy(new_qotd, qsp->qotd, min(qsp->qotd_len, new_len));
    ddi_umem_free(qsp->qotd_cookie);
    qsp->qotd = new_qotd;
    qsp->qotd_cookie = new_cookie;
    qsp->qotd_len = new_len;
    qsp->flags |= QOTD_CHANGED;
    mutex_exit(&qsp->lock);

    return (0);
}
case QOTDIOCDISCARD: {
    char *new_qotd = NULL;
    ddi_umem_cookie_t new_cookie;

    if (!(mode & FWRITE))
        return (EACCES);

    if (qsp->qotd_len != init_qotd_len) {
        new_qotd = ddi_umem_alloc(init_qotd_len,
            DDI_UMEM_SLEEP, &new_cookie);
    }

    mutex_enter(&qsp->lock);
    while (qsp->flags & QOTD_BUSY) {
```

**EXAMPLE 3-5** Quote Of The Day Version 3 Source File (Continued)

```
        if (cv_wait_sig(&qsp->cv, &qsp->lock) == 0) {
            mutex_exit(&qsp->lock);
            if (new_qotd != NULL)
                ddi_umem_free(new_cookie);
            return (EINTR);
        }
    }
    if (new_qotd != NULL) {
        ddi_umem_free(qsp->qotd_cookie);
        qsp->qotd = new_qotd;
        qsp->qotd_cookie = new_cookie;
        qsp->qotd_len = init_qotd_len;
    } else {
        bzero(qsp->qotd, qsp->qotd_len);
    }
    (void)strncpy(qsp->qotd, init_qotd, init_qotd_len);
    qsp->flags &= ~QOTD_CHANGED;
    mutex_exit(&qsp->lock);

    return (0);
}
default:
    return (ENOTTY);
}
}
```

Enter the definitions shown in the following example into a text file named `qotd.h`.

**EXAMPLE 3-6** Quote Of The Day Version 3 Header File

```
#ifndef _SYS_QOTD_H
#define _SYS_QOTD_H

#define QOTDIOC          ('q' << 24 | 't' << 16 | 'd' << 8)

#define QOTDIOCGSZ      (QOTDIOC | 1)  /* Get quote buffer size */
#define QOTDIOCSSZ      (QOTDIOC | 2)  /* Set new quote buffer size */
#define QOTDIOCDISCARD  (QOTDIOC | 3)  /* Discard quotes and reset */

#endif /* _SYS_QOTD_H */
```

Enter the configuration information shown in the following example into a text file named `qotd_3.conf`.

**EXAMPLE 3-7** Quote Of The Day Version 3 Configuration File

```
name="qotd_3" parent="pseudo" instance=0;
```

## Building and Installing Quote Of The Day Version 3

Compile and link the driver. The following example shows compiling and linking for a 32-bit architecture:

```
% cc -D_KERNEL -c qotd_3.c
% ld -r -o qotd_3 qotd_3.o
```

Make sure you are user root when you install the driver.

Copy the driver binary to the /tmp directory as discussed in [“Building and Installing the Driver” on page 48](#).

```
# cp qotd_3 /tmp
# ln -s /tmp/qotd_3 /usr/kernel/drv/qotd_3
```

Copy the configuration file to the kernel driver area of the system.

```
# cp qotd_3.conf /usr/kernel/drv
```

In a separate window, enter the following command:

```
% tail -f /var/adm/messages
```

Make sure you are user root when you load the driver. Use the `add_drv(1M)` command to load the driver:

```
# add_drv qotd_3
```

You should see the following messages in the window where you are viewing /var/adm/messages:

```
date time machine pseudo: [ID 129642 kern.info] pseudo-device: qotd_30
date time machine genunix: [ID 936769 kern.info] qotd_30 is /pseudo/qotd_3@0
```

## Using Quote Of The Day Version 3

This section describes how to read and write the `qotd_3` device and how to test the driver’s I/O controls. The I/O controls include retrieving the size of the storage buffer, setting a new size for the storage buffer, and reinitializing the storage buffer size and contents.

### Reading the Device

When you access this `qotd_3` device for reading, the command you use to access the device retrieves the data from the device node. The command then displays the data in the same way that the `command` displays any other input. To get the name of the device special file, look in the /devices directory:

```
% ls -l /devices/pseudo/qotd*
crw----- 1 root sys 122, 0 date time /devices/pseudo/qotd_3@0:qotd_3
```

To read the qotd\_3 device, you can use the cat(1) command:

```
# cat /devices/pseudo/qotd_3@0:qotd_3
On the whole, I'd rather be in Philadelphia. - W. C. Fields
```

## Writing the Device

To write to the qotd\_3 device, you can redirect command-line input:

```
# echo "A life is not important except in the impact it has on others.
- Jackie Robinson" >> /devices/pseudo/qotd_3@0:qotd_3
# cat /devices/pseudo/qotd_3@0:qotd_3
A life is not important except in the impact it has on others. - Jackie
Robinson
```

## Exercising the Driver's I/O Controls

In addition to changes in the driver, Quote Of The Day Version 3 introduces a new utility program. The qotdctl command enables you to test the driver's I/O controls.

The source for this command is shown in [Example 3-8](#). Compile the qotdctl utility as follows:

```
% cc -o qotdctl qotdctl.c
```

The qotdctl command has the following options:

- g           Get the size that is currently allocated. Call the ioctl(9E) entry point of the driver with the QOTDIOCGSZ request. The QOTDIOCGSZ request reports the current size of the space allocated for the quotation.
- s *size*     Set the new size to be allocated. Call the ioctl(9E) entry point of the driver with the QOTDIOCSSZ request. The QOTDIOCSSZ request sets a new size for the quotation space.
- r           Discard the contents and reset the device. Call the ioctl(9E) entry point of the driver with the QOTDIOCDISCARD request.

Invoking qotdctl with the -r option is the only way to unset the QOTD\_CHANGED flag in the device. The device cannot be detached while the QOTD\_CHANGED flag is set. This protects the contents of the ramdisk device from being unintentionally or automatically removed. For example, a device might be automatically removed by the automatic device unconfiguration thread.

When you are no longer interested in the contents of the device, run the qotdctl command with the -r option. Then you can remove the device.

- h            Display help text.
- V            Display the version number of the `gotdctl` command.
- d *device*    Specify the device node to use. The default value is `/dev/gotd0`.

Use the `gotdctl` command to test the driver's I/O controls:

```
# ./gotdctl -v
gotdctl 1.0
# ./gotdctl -h
Usage: ./gotdctl [-d device] {-g | -h | -r | -s size | -v}
# ./gotdctl -g
open: No such file or directory
```

By default, the `gotdctl` command accesses the `/dev/gotd0` device. The `gotd_3` device in this example is `/devices/pseudo/gotd_3@0:gotd_3`. Either define a link from `/dev/gotd0` to `/devices/pseudo/gotd_3@0:gotd_3` or use the `-d` option to specify the correct device:

```
# ./gotdctl -d /devices/pseudo/gotd_3@0:gotd_3 -g
128
# ./gotdctl -d /devices/pseudo/gotd_3@0:gotd_3 -s 512
# ./gotdctl -d /devices/pseudo/gotd_3@0:gotd_3 -g
512
# ./gotdctl -d /devices/pseudo/gotd_3@0:gotd_3 -r
# cat /devices/pseudo/gotd_3@0:gotd_3
```

On the whole, I'd rather be in Philadelphia. - W. C. Fields

If you try to remove the device now, you will receive an error message:

```
# rem_drv gotd_3
Device busy
Cannot unload module: gotd_3
Will be unloaded upon reboot.
```

The device is still marked busy because you have not told the driver that you are no longer interested in this device. Run the `gotdctl` command with the `-r` option to unset the `QOTD_CHANGED` flag in the driver and mark the device not busy:

```
# ./gotdctl -r
```

Enter the source code shown in the following example into a text file named `gotdctl.c`.

**EXAMPLE 3-8** Quote Of The Day I/O Control Command Source File

```
#include <sys>
#include <fcntl.h>
#include <stdio.h>
#include <unistd.h>
#include <errno.h>
#include <stdlib.h>
```

**EXAMPLE 3-8** Quote Of The Day I/O Control Command Source File (Continued)

```
#include "qotd.h"

static const char *DEFAULT_DEVICE = "/dev/qotd0";
static const char *VERSION = "1.0";

static void show_usage(const char *);
static void get_size(const char *);
static void set_size(const char *, size_t);
static void reset_dev(const char *);

int
main(int argc, char *argv[])
{
    int op = -1;
    int opt;
    int invalid_usage = 0;
    size_t sz_arg;
    const char *device = DEFAULT_DEVICE;

    while ((opt = getopt(argc, argv,
        "d:(device)g(get-size)h(help)r(reset)s:(set-size)V(version)")
        != -1) {
        switch (opt) {
            case 'd':
                device = optarg;
                break;
            case 'g':
                if (op >= 0)
                    invalid_usage++;
                op = QOTDIOCGSZ;
                break;
            case 'h':
                show_usage(argv[0]);
                exit(0);
                /*NOTREACHED*/
            case 'r':
                if (op >= 0)
                    invalid_usage++;
                op = QOTDIOCDISCARD;
                break;
            case 's':
                if (op >= 0)
                    invalid_usage++;
                op = QOTDIOCSSZ;
                sz_arg = (size_t)atol(optarg);
                break;
            case 'V':
                (void) printf("gotdctl %s\n", VERSION);
                exit(0);
                /*NOTREACHED*/
            default:
                invalid_usage++;
        }
    }
}
```

**EXAMPLE 3-8** Quote Of The Day I/O Control Command Source File (Continued)

```
                break;
            }
        }

        if (invalid_usage > 0 || op < 0) {
            show_usage(argv[0]);
            exit(1);
        }

        switch (op) {
        case QOTDIOCGSZ:
            get_size(device);
            break;
        case QOTDIOCSSZ:
            set_size(device, sz_arg);
            break;
        case QOTDIOCDISCARD:
            reset_dev(device);
            break;
        default:
            (void) fprintf(stderr,
                "internal error - invalid operation %d\n", op);
            exit(2);
        }

        return (0);
    }

    static void
    show_usage(const char *execname)
    {
        (void) fprintf(stderr,
            "Usage: %s [-d device] {-g | -h | -r | -s size | -V}\n", execname);
    }

    static void
    get_size(const char *dev)
    {
        size_t sz;
        int fd;

        if ((fd = open(dev, O_RDONLY)) < 0) {
            perror("open");
            exit(3);
        }

        if (ioctl(fd, QOTDIOCGSZ, &sz) < 0) {
            perror("QOTDIOCGSZ");
            exit(4);
        }

        (void) close(fd);
    }
}
```

**EXAMPLE 3-8** Quote Of The Day I/O Control Command Source File *(Continued)*

```
        (void) printf("%zu\n", sz);
    }

static void
set_size(const char *dev, size_t sz)
{
    int fd;

    if ((fd = open(dev, O_RDWR)) < 0) {
        perror("open");
        exit(3);
    }

    if (ioctl(fd, QOTDIOCSSZ, &sz) < 0) {
        perror("QOTDIOCSSZ");
        exit(4);
    }

    (void) close(fd);
}

static void
reset_dev(const char *dev)
{
    int fd;

    if ((fd = open(dev, O_RDWR)) < 0) {
        perror("open");
        exit(3);
    }

    if (ioctl(fd, QOTDIOCDISCARD) < 0) {
        perror("QOTDIOCDISCARD");
        exit(4);
    }

    (void) close(fd);
}
```



---

## Tips for Writing Device Drivers

---

This chapter provides some general guidelines for writing device drivers. The guidelines are organized into the following categories:

- “Steps in the Development Process” on page 97
- “Device Driver Coding Tips” on page 98
- “Device Driver Testing Tips” on page 100
- “Device Driver Debugging and Tuning Tips” on page 101

---

### Steps in the Development Process

The general steps in writing a device driver are as follows:

1. Write, compile, and link the new code.
2. Create the necessary hardware configuration files.  
You need to create a hardware configuration file unique to the device called `xx.conf` where `xx` is the prefix for the device. This file is used to update the `driver.conf(4)` file. See “Writing the Device Configuration File” on page 47, which describes creation of a configuration file for the *dummy* example.
3. Copy the driver to the appropriate module directory.
4. Install the device driver using `add_drv(1M)`.
5. Load the driver.
6. Test the driver.

Drivers should be rigorously tested in the following areas:

- Configuration
- Functionality
- Error handling
- Loading and unloading

- Stress, performance, and interoperability
- DDI/DKI compliance
- Installation and packaging

Additional testing is specific to the type of driver.

7. Remove the driver using the `rem_drv(1M)`.

All drivers will need to be removed eventually so you need to make sure that your driver can be successfully removed.

For detailed descriptions of the development process, see “Driver Development Summary” in *Writing Device Drivers*.

---

## Device Driver Coding Tips

Use these guidelines when you are creating the software for your new driver:

- Use a prefix based on the name of your driver to give global variables and functions unique names.
- If you are basing your design on an existing driver, modify the configuration file before adding the driver.
 

The `-n` option in the `add_drv(1M)` command enables you to update the system configuration files for a driver without loading or attaching the driver.
- Use the `cmn_err()` function rather than `printf()` to log driver activity.
 

The `cmn_err(9F)` function is the recommended method for displaying information on the console. This function is more versatile than `printf(3C)` for displaying special formats used by drivers such as device register bits.
- Clean up allocations and other initialization activities when the driver exits.
 

When the driver exits, whether intentionally or prematurely, you need to perform such tasks as closing opened files, freeing allocated memory, releasing mutex locks, and destroying any mutexes that have been created. In addition, the system must be able to close all minor devices and detach driver instances even after the hardware fails. An orderly approach is to reverse `_init()` actions in the `_fini()` routine, reverse `open()` operations in the `close()` routine, and reverse `attach()` operations in the `detach()` routine.
- Use `assert(9F)` to catch unexpected error returns.
 

`ASSERT()` is a macro that halts the kernel execution if a condition that was expected to be true turns out to be false. To activate `ASSERT()`, you need to include the `sys/debug.h` header file and specify the `DEBUG` preprocessor symbol during compilation.
- Use `mutex_owned()` to validate and document locking requirements.

The `mutex_owned(9F)` function helps determine whether the current thread owns a specified mutex. To determine whether a mutex is held by a thread, use `mutex_owned()` within `ASSERT()`.

- Use conditional compilation to toggle “costly” debugging features.

The Solaris OS provides various debugging functions, such as `assert()` and `mutex-owned()`, that can be turned on by specifying the `DEBUG` preprocessor symbol when the driver is compiled. With conditional compilation, unnecessary code can be removed from the production driver. This approach can also be accomplished by using a global variable.

- Use a separate instance of the driver for each device to be controlled.
- Use DDI functions as much as possible in your device drivers.

These interfaces shield the driver from platform-specific dependencies such as mismatches between processor and device endianness and any other data order dependencies. With these interfaces, a single-source driver can run on the SPARC platform, x86 platform, and related processor architectures.

- Anticipate corrupted data.

Always check that the integrity of data before that data is used. The driver must avoid releasing bad data to the rest of the system.

- A device should only write to DMA buffers that are controlled solely by the driver. This technique prevents a DMA fault from corrupting an arbitrary part of the system’s main memory.

- Use the `ddi_umem_alloc(9F)` function when you need to make DMA transfers. This function guarantees that only whole, aligned pages are transferred.

- Set a fixed number of attempts before taking alternate action to deal with a stuck interrupt.

The device driver must not be an unlimited drain on system resources if the device locks up. The driver should time out if a device claims to be continuously busy. The driver should also detect a pathological (stuck) interrupt request and take appropriate action.

- Use care when setting the sequence for mutex acquisitions and releases so as to avoid unwanted thread interactions if a device fails.

See “Thread Interaction” in *Writing Device Drivers* for more information.

- Check for malformed `ioctl()` requests from user applications.

User requests can be potentially and intentionally destructive. The design of the driver should take into consideration the construction of each type of potential `ioctl()` request.

- Try to avoid situations where a driver continues to function without detecting a device failure.

A driver should switch to an alternative device rather than try to work around a device failure.

- All device drivers in the Solaris OS must support hotplugging.

All devices need to be able to be installed or removed without requiring a reboot of the system.

- All device drivers should support power management.  
Power management provides the ability to control and manage the electrical power usage of a computer system or device. Power management enables systems to conserve energy by using less power when idle and by shutting down completely when not in use.
- Apply the `volatile` keyword to any variable that references a device register.  
Without the `volatile` keyword, the compile-time optimizer can delete important accesses to a register.
- Perform periodic health checks to detect and report faulty devices.  
A periodic health check should include the following activities:
  - Check any register or memory location on the device whose value might have been altered since the last poll.
  - Timestamp outgoing requests such as transmit blocks or commands that are issued by the driver.
  - Initiate a test action on the device that should be completed before the next scheduled check.

---

## Device Driver Testing Tips

Testing a new device driver can cause irreparable harm to the kernel. The following tips can help avoid major problems:

- Use a serial connection to control your test machine from a separate host system.  
This technique is explained in “Testing With a Serial Connection” in *Writing Device Drivers*.
- Use an alternate kernel.  
Bootting from a copy of the kernel and the associated binaries rather than from the default kernel avoids inadvertently rendering the system inoperable.
- Use an additional kernel module to experiment with different kernel variable settings.  
This approach isolates experiments with the kernel variable settings. See “Setting Up Test Modules” in *Writing Device Drivers*.
- Make contingency plans for potential data loss on a test system.  
If your test system is set up as a client of a server, then you can boot from the network if problems occur. You could also create a special partition to hold a copy of a bootable root file system. See “Avoiding Data Loss on a Test System” in *Writing Device Drivers*.

- Capture system crash dumps if your test system panics.
- Use `fsck(1M)` to repair the damaged root file system temporarily if your system crashes during the `attach(9E)` process so that any crash dumps can be salvaged. See “Recovering the Device Directory” in *Writing Device Drivers*.
- Install drivers in the `/tmp` directory until you are finished modifying and testing the `_info()`, `_init()`, and `attach()` routines.

Keep a driver in the `/tmp` directory until the driver has been well tested. If a panic occurs, the driver will be removed from `/tmp` directory and the system will reboot successfully.

---

## Device Driver Debugging and Tuning Tips

The Solaris OS provides various tools for debugging and tuning your device driver:

- Use the `kmdb(1)` kernel debugger for runtime debugging.  
The `kmdb` debugger provides typical runtime debugger facilities, such as breakpoints, watch points, and single-stepping. For more information, see *Solaris Modular Debugger Guide*.
- Use the `mdb(1)` modular debugger for postmortem debugging.  
Postmortem debugging is performed on a system crash dump rather than on a live system. With postmortem debugging, the same crash dump can be analyzed by different people or processes simultaneously. In addition, `mdb` allows you to create special macros called *dmods* to perform rigorous analysis on the dump. For more information, see *Solaris Modular Debugger Guide*.
- Use the `kstat(3KSTAT)` facility to export module-specific kernel statistics for your device driver.
- Use the DTrace facility to add instrumentation to your driver dynamically so that you can perform tasks such as analyzing the system and measuring performance. For information on DTrace, see *Solaris Dynamic Tracing Guide*.



# Index

---

## A

add\_drv command, 49-50, 97  
    use in modifying existing drivers, 98  
alternate kernels, use in testing, 100  
ASSERT() kernel function, 61, 64, 69-70, 98  
attach() entry point, 31-38, 48-49, 62-63

## B

blk device, 22  
block device, 22  
boot command, 20  
bzero() kernel function, 80

## C

cat command, 51  
cb\_ops driver structure, 17-19, 42-47, 78-79  
cc command, 48  
character device, 22  
close() entry point, 38-42, 63  
cmn\_err() kernel function, 28-29, 49-52, 57, 98  
commands  
    add\_drv, 49-50, 97, 98  
    boot, 20  
    cat, 51  
    cc, 48  
    dmesg, 63  
    echo, 51  
    fsck, 101

## commands (Continued)

    kernel, 19-20  
    ld, 15, 30, 48  
    mknod, 21  
    modinfo, 50, 70  
    modload, 50  
    modunload, 52  
    more, 70  
    prtconf, 21, 22, 50, 70  
    prtpicl, 22  
    rem\_drv, 51-52, 98  
    syslogd, 49-52, 59  
compiling, 48  
condition variables, 75-78  
conditional compilation, 99  
condvar() kernel functions, 77  
configuration files, 47, 97  
crash dumps, use in testing, 101  
cv\_broadcast() kernel function, 77-78, 79  
cv\_destroy() kernel function, 73, 77  
cv\_init() kernel function, 73, 77  
cv\_signal() kernel function, 77-78  
cv\_wait() kernel function, 77-78  
cv\_wait\_sig() kernel function, 77

## D

data loss, avoiding while testing, 100  
data model, converting, 72  
ddi\_copyin() kernel function, 16, 72, 75, 79  
ddi\_copyout() kernel function, 16, 72, 79

- ddi\_create\_minor\_node() kernel function, 31-38, 62-63, 72-73
- ddi\_get\_instance() kernel function, 24, 34, 62-63
- ddi\_get\_soft\_state() kernel function, 61, 62-63, 63
- ddi\_model\_convert\_from() kernel function, 72, 79
- ddi\_prop\_get\_int() kernel function, 47
- ddi\_prop\_lookup() kernel function, 47
- ddi\_prop\_op() kernel function, 31-38
- ddi\_remove\_minor\_node() kernel function, 31-38, 62, 73
- ddi\_report\_dev() kernel function, 63
- ddi\_soft\_state\_fini() kernel function, 61, 62
- ddi\_soft\_state\_free() kernel function, 61, 62
- ddi\_soft\_state\_init() kernel function, 61, 62
- ddi\_soft\_state() kernel function, 33-34
- ddi\_soft\_state\_zalloc() kernel function, 61, 62-63
- ddi\_umem\_alloc() kernel function, 74-75
- ddi\_umem\_free() kernel function, 73, 74-75
- debugging device drivers, tips, 101
- detach() entry point, 30, 31-38, 63
- /dev directory, 21, 22-23
- dev\_info device structure, 34, 35-36, 36-37
- dev\_ops driver structure, 17-19, 42-47
- devfs devices file system, 21
- devfsadm devices file system administration command, 22-23
- device drivers, 17-19
  - coding tips, 98-100
  - compiling, 48
  - condition variables, 75-78
  - conditional compilation, 99
  - debugging tips, 101
  - development guidelines, 97
  - development steps, 97-98
  - directories, 19-20, 48
    - adding, 20
  - entry points, 17, 26
    - See also* entry points
  - how used, 18
  - I/O controls, 79-80, 92-96
  - installing, 48-49
- device drivers (Continued)
  - linking, 48
  - loading, 20, 49-50
  - mutexes, 75-78
  - naming conventions, 98
  - recommended housekeeping, 98
  - removing, 51-52
  - structures
    - See* driver structures
  - test areas, 97
  - testing tips, 100-101
  - thread synchronization, 75-78
  - tuning, 101
  - unloading, 52
- device instance pointer (dip), 34, 35-36, 36-37
- device number, 23-24
- device structures
  - dev\_info, 34, 35-36, 36-37
- device tree, 21-22
- devices
  - blk, 22
  - block, 20-24
  - character, 20-24, 26-47
  - configuration files, 47
  - device tree, 21-22
  - directories, 21, 22-23
  - exclusive access, 76
  - file system
    - devfs, 21
    - devfsadm, 22-23
  - instances, 24, 34, 35-36, 36-37
  - md metadvice, 22
  - names, 22-23
  - nexus, 18, 22
  - numbers, 21, 23-24, 34
  - properties, 37-38, 47
  - pseudo, 18, 25
  - ramdisk, 18, 71
  - raw, 22
  - reading, 51, 70, 91-92
  - special files, 20-24
  - state, 61
    - writing, 51, 71-96, 92
- /devices directory, 21, 23
- /devices/pseudo directory, 22, 50, 70
- devmap() entry point, 16
- dmesg command, 63
- driver.conf file, 47



driver structures  
 cb\_ops, 17-19, 42-47, 78-79  
 character and block operations  
   structure, 43-44  
 dev\_ops, 17-19, 42-47  
 device operations structure, 45-46  
 modinfo, 29  
 modldrv, 42-47  
 modlinkage, 28-29, 42-47  
 module linkage structures, 46  
 drivers, *See* device drivers  
 DTrace analyzer, 101

## E

echo command, 51  
 entry points  
 attach(), 31-38, 48-49, 62-63  
 autoconfiguration, 31-38  
 close(), 38-42, 63  
 detach(), 30, 31-38, 63  
 devmap(), 16  
 \_fini(), 27-30, 62  
 getinfo(), 31-38, 63  
 \_info(), 27-30, 48-49  
 \_init(), 27-30, 48-49, 62  
 ioctl(), 21, 78-79, 79-80, 92-96  
 loadable module configuration, 27-30  
 open(), 38-42, 63  
 prop\_op(), 31-38  
 read(), 38-42, 63  
 user context, 38-42  
 write(), 38-42, 78-79  
 /etc/name\_to\_major file, 50, 70

## F

files  
 driver.conf, 47  
 /etc/name\_to\_major, 50, 70  
 system, 19-20  
 /var/adm/messages, 49-52, 59  
 \_fini() entry point, 27-30, 62  
 fsck command, 101  
 functions  
 kstat(), 101

functions (Continued)  
 printf(), 98  
 signal(), 77-78, 78, 80

## G

getinfo() entry point, 31-38, 63  
 getminor() kernel function, 63

## H

hotplugging, 99

## I

I/O controls, 79-80, 92-96  
 \_info() entry point, 27-30, 48-49  
 \_init() entry point, 27-30, 48-49, 62  
 instance number, 24, 34, 35-36, 36-37  
 interrupts, avoiding problems, 99  
 ioctl() entry point, 21, 78-79, 79-80, 92-96  
 ioctl() requests, avoiding problems, 99

## K

kernel, 13  
   address space, 14, 16  
   privilege  
     *See also* kernel mode  
 kernel command, 19-20  
 kernel functions  
 ASSERT(), 61, 64, 69-70, 98  
 bzero(), 80  
 cmn\_err(), 28-29, 49-52, 57, 98  
 condvar(), 77  
 cv\_broadcast(), 77-78, 79  
 cv\_destroy(), 73, 77  
 cv\_init(), 73, 77  
 cv\_signal(), 77-78  
 cv\_wait(), 77-78  
 cv\_wait\_sig(), 77  
 ddi\_copyin(), 16, 72, 75, 79  
 ddi\_copyout(), 16, 72, 79  
 ddi\_create\_minor\_node(), 31-38, 62-63,  
   72-73

## kernel functions (Continued)

- ddi\_get\_instance(), 24, 34, 62-63
- ddi\_get\_soft\_state(), 61, 62-63, 63
- ddi\_model\_convert\_from(), 72, 79
- ddi\_prop\_get\_int(), 47
- ddi\_prop\_lookup(), 47
- ddi\_prop\_op(), 31-38
- ddi\_remove\_minor\_node(), 31-38, 62, 73
- ddi\_report\_dev(), 63
- ddi\_soft\_state(), 33-34
- ddi\_soft\_state\_fini(), 61, 62
- ddi\_soft\_state\_free(), 61, 62
- ddi\_soft\_state\_init(), 61, 62
- ddi\_soft\_state\_zalloc(), 61, 62-63
- ddi\_umem\_alloc(), 74-75
- ddi\_umem\_free(), 73, 74-75
- getminor(), 63
- kmem\_alloc(), 74
- kmem\_zalloc(), 74
- memcpy(), 80
- mod\_info(), 27-30
- mod\_install(), 27-30, 62
- mod\_remove(), 27-30, 62
- mutex(), 77
- mutex\_destroy(), 73, 77
- mutex\_init(), 73, 77
- mutex\_owned(), 98
- nochpoll(), 44
- nodev(), 44, 45
- nulldev(), 38-42, 45
- strncpy(), 73, 80
- strncpy(), 73
- uiomove(), 63, 72, 75, 79

kernel mode, 13

kernel modules, use in testing, 100

kernel statistics, 101

kernel structures

- uio, 63, 78

kmdb kernel debugger, 101

kmem\_alloc() kernel function, 74

kmem\_zalloc() kernel function, 74

kstat() function, 101

## L

- ld command, 15, 30, 48
- linking, 15, 30, 48

## M

- major number, 21, 23
- mdb modular debugger, 101
- memcpy() kernel function, 80
- metadvice, 22
- minor number, 21, 23, 34
- mknod command, 21
- mknod() system call, 21
- mmap() system call, 16
- mod\_info() kernel function, 27-30
- mod\_install() kernel function, 27-30, 62
- mod\_remove() kernel function, 27-30, 62
- moddir kernel variable, 20
- modinfo command, 50, 70
- modinfo driver structure, 29
- modldrv driver structure, 42-47
- modlinkage driver structure, 28-29, 42-47
- modload command, 50
- modunload command, 52
- more command, 70
- mutex\_destroy() kernel function, 73, 77
- mutex\_init() kernel function, 73, 77
- mutex() kernel function, 77
- mutex\_owned() kernel function, 98
- mutexes, 75-78
  - avoiding problems, 99

## N

- naming conventions, 98
- nexus device, 22
- nochpoll() kernel function, 44
- nodev() kernel function, 44, 45
- nulldev() kernel function, 38-42, 45

## O

- open() entry point, 38-42, 63

## P

- power management, 100
- printf() function, 98
- prop\_op() entry point, 31-38
- protected mode, 13

prtconf command, 21, 22, 50, 70  
prtpicl command, 22

## Q

QOTD\_BUSY condition, 75, 76

## R

raw device, 22  
read() entry point, 38-42, 63  
read() system call, 21  
rem\_drv command, 51-52, 98  
restricted mode, 13

## S

serial connections, use in testing, 100  
signal() function, 77-78, 78, 80  
soft state, 61  
SPARC  
    address space, 16  
    compiling, 48  
special files, 20-24  
state structures, 60, 61  
strncpy() kernel function, 73, 80  
strncpy() kernel function, 73  
syslogd command, 49-52, 59  
system calls  
    mknod(), 21  
    mmap(), 16  
    read(), 21  
system configuration information file, 19-20  
system crash dumps, use in testing, 101

## T

testing device drivers  
    test areas, 97  
    tips, 100-101  
thread synchronization, 75-78  
tuning device drivers, tips, 101

## U

uiio kernel structure, 63, 78  
uiomove() kernel function, 63, 72, 75, 79  
user mode, 13  
/usr/kernel directory, 20  
/usr/kernel/drv directory, 48

## V

/var/adm/messages file, 49-52, 59  
volatile keyword, 100

## W

write() entry point, 38-42, 78-79

## X

x86  
    address space, 16  
    compiling, 48

