

# EXAM PREP

Your Complete Certification Solution

Exam **CX-310-200**

# Solaris 10

## System Administration

### Part I



CD Features ExamGear  
Practice Questions!

Bill Calkins

# Solaris 10 System Administration Exam Prep (Exam CX-310-200), Part I

Copyright © 2009 by Que Publishing

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-7897-3790-8

ISBN-10: 0-7897-3790-6

*Library of Congress Cataloging-in-Publication Data:*

Calkins, Bill.

Solaris 10 system administration exam prep (Exam CX-310-200) / Bill  
Calkins.

p. cm.

ISBN 978-0-7897-3790-8 (pbk. w/cd)

1. Electronic data processing personnel--Certification. 2. Operating  
systems (Computers)--Examinations--Study guides. 3. Solaris (Computer  
file) I. Title.

QA76.3.C34346 2008

005.4'32--dc22

2008031592

Printed in the United States on America

First Printing: September 2008

**ASSOCIATE PUBLISHER**  
David Dusthimer

**ACQUISITIONS EDITOR**  
Betsy Brown

**SENIOR DEVELOPMENT  
EDITOR**  
Christopher Cleveland

**TECHNICAL EDITOR**  
John Philcox

**MANAGING EDITOR**  
Patrick Kanouse

**PROJECT EDITOR**  
Mandie Frank

**COPY EDITOR**  
Gayle Johnson

**INDEXER**  
Ken Johnson

**PROOFREADER**  
Matt Purcell

**PUBLISHING COORDINATOR**  
Vanessa Evans

**MULTIMEDIA DEVELOPER**  
Dan Scherf

**DESIGNER**  
Gary Adair

**PAGE LAYOUT**  
Mark Shriar

## **Trademarks**

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Que Publishing cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

## **Warning and Disclaimer**

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the CD or programs accompanying it.

## **Bulk Sales**

Que Certification offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales. For more information, please contact

### **U.S. Corporate and Government Sales**

**1-800-382-3419**

**corpsales@pearsontechgroup.com**

For sales outside the United States, please contact

### **International Sales**

**international@pearsoned.com**

# Introduction

Bill Calkins has been training Solaris system administrators for more than 15 years. This book contains the training material he uses in his basic and advanced Solaris administration courses. Over the years, this material has helped thousands of Solaris administrators become certified. This is our first edition of the *Solaris 10 System Administration Exam Prep*. It began with the *Training Guide* for Solaris 2.6, 7, 8, and 9 and is now the *Exam Prep* for Solaris 10. Instructors from universities and training organizations around the world have used this book as courseware in their Solaris administration courses. In addition, administrators from around the world have used this book as a self-study guide when instruction from a Sun training center is either unavailable or not within their budget. Many people have written with their success stories, suggestions, and comments. Their suggestions are what keep making this guide more valuable.

The SCSA Solaris 10 OS CX-310-200 and CX-310-202 *Exam Prep* books provide training materials for anyone who is interested in becoming a Sun Certified System Administrator for Solaris 10. When used as a study guide, these two books will save you a great deal of time and effort searching for information you will need to know when taking the exam. Each book covers the exam objectives in enough detail for the inexperienced administrator to learn the objectives and apply the knowledge to real-life scenarios. Experienced readers will find the material in these books to be complete and concise, making it a valuable study guide for the Sun Certified System Administrator (SCSA) exams.

This book is not a cheat sheet or cram session for the exam; it is a training manual. In other words, it does not merely give answers to the questions you will be asked on the exam. We have made certain that this book addresses the exam objectives in detail, from start to finish. If you are unsure about the objectives on the exams, this book teaches you what you need to know. After reading each chapter, assess your knowledge of the material covered using the review questions at the end of the chapter. When you have completed reading a section, use the practice exam at the end of the book and the ExamGear test engine on the CD-ROM to assess your knowledge of the objectives covered on each exam. This CD-ROM contains sample questions similar to what you are likely to see on the real exam. More sample questions are available at <http://www.UnixEd.com>, so be sure to visit this site to find additional training and study materials.

# How This Book Helps You

This book teaches you how to administer the Solaris 10 operating system. It offers you a self-guided training course for all the areas covered on the CX-310-200 certification exam by showing you how to install, configure, and administer the Solaris 10 operating environment. You will learn the specific skills that are required to administer a system and to pass the first part of the Sun Certified System Administrator exam for Solaris 10 (CX-210-200). Experienced administrators who are upgrading an existing Solaris certification will find in-depth coverage of the new topics that they will need to learn for the CX-310-203 upgrade exam in both the SCSA Solaris 10 OS CX-310-200 and CX-310-202 *Exam Prep* books.

Throughout the book, we provide helpful tips and real-world examples that we have encountered as system administrators. In addition, we provide useful real-world exercises to help you practice the material you have learned. The following list describes this book's setup:

- ▶ **Organization:** This book is organized according to the individual exam objectives. Every objective you need to know for installing, configuring, and administering a Solaris 10 system is in this book. We have attempted to present the objectives in an order that is as close as possible to that listed by Sun. However, we have not hesitated to reorganize them as needed to make the material as easy as possible for you to learn. We have also attempted to make the information accessible in the following ways:
  - ▶ This Introduction includes the full list of exam topics and objectives.
  - ▶ Read the “Study and Exam Prep Tips” element early on to help you develop study strategies. This element provides valuable exam-day tips and information on exam/question formats.
  - ▶ Each chapter begins with a list of the objectives to be covered, exactly as they are defined by Sun. Throughout each section, material that is directly related to the exam objectives is identified.
  - ▶ Each chapter also begins with an outline that provides you with an overview of the material and the page numbers where particular topics can be found.
- ▶ **Instructional features:** This book is designed to provide you with multiple ways to learn and reinforce the exam material. The following are some of the helpful methods:
  - ▶ **Objective explanations:** As mentioned, each chapter begins with a list of the objectives covered in the chapter. In addition, immediately following each objective is an explanation in a context that defines it more meaningfully.
  - ▶ **Study Strategies:** The beginning of each chapter also includes strategies for studying and retaining the material in the chapter, particularly as it is addressed on the exam.

- ▶ **Exam Alerts:** Throughout each chapter you'll find exam tips that help you prepare for exam day. These tips were written by those who have already taken the Solaris 10 certification exams.
- ▶ **Notes:** Notes contain various types of useful information, such as tips on technology or administrative practices, historical background on terms and technologies, or side commentary on industry issues.
- ▶ **Cautions:** When you use sophisticated information technology, mistakes or even catastrophes are always possible because of improper application of the technology. Cautions alert you to such potential problems.
- ▶ **Step by Steps:** These are hands-on lab exercises that walk you through a particular task or function relevant to the exam objectives.
- ▶ **Key Terms:** A list of key terms appears near the end of each chapter.
- ▶ **Exercises:** Found at the end of the chapters in the "Summary" section, the exercises are performance-based opportunities for you to learn and assess your knowledge.
- ▶ **Suggested Readings and Resources:** At the end of each chapter is a list of additional resources that you can use if you are interested in going beyond the objectives and learning more about the topics presented in the chapter.
- ▶ **Extensive practice test options:** The book provides numerous opportunities for you to assess your knowledge and practice for the exam. The practice options include the following:
  - ▶ **Exam Questions:** These questions appear in the "Summary" section at the end of each chapter. They allow you to quickly assess your comprehension of what you just read. Answers to the questions are provided in the section "Answers to Exam Questions."
  - ▶ **Practice Exam:** A practice exam is included in Part II of this book, "Final Review," (as discussed in a moment).
  - ▶ **ExamGear:** The ExamGear software included on the CD-ROM provides further practice questions.

**NOTE**

**ExamGear software** For a complete description of the ExamGear test engine, see Appendix A, "What's on the CD-ROM."

- ▶ **Final Review:** This part of the book gives you two valuable tools for preparing for the exam:
  - ▶ **Fast Facts:** This condensed version of the information contained in the book will prove extremely useful for last-minute review.
  - ▶ **Practice Exam:** A full practice exam is included, with questions written in styles similar to those used on the actual exam. Use the practice exam to assess your readiness for the real exam.
- ▶ **Appendix and glossary:** This book also contains a glossary and a description of what is on the CD-ROM (Appendix A).

These and all the other book features will enable you to thoroughly prepare for the exam.

## Conventions Used in This Book

- ▶ **Commands:** In the text, steps, and examples, the commands you type are displayed in a special monospace font.
- ▶ **Arguments and options:** In command syntax, command options and arguments are enclosed in `< >`. (The italicized words within the `< >` symbols stand for what you actually type. You don't type the `< >`.)

```
lp -d<printer name> <filename> <return>
```
- ▶ **Using the mouse:** When using menus and windows, you select items using the mouse. Here is the default mapping for a three-button mouse:

Left button: Select

Middle button: Transfer/adjust

Right button: Menu

The Select button is used to select objects and activate controls. The middle mouse button is configured for either Transfer or Adjust. By default, it is set up for Transfer, which means that it is used to drag or drop list or text items. You use the left mouse button to highlight text, and then you use the middle button to move the text to another window or to reissue a command. The middle button can also be used to move around windows on the screen. The right mouse button, the Menu button, is used to display and choose options from pop-up menus.

- ▶ **Menu options:** The names of menus and the options that appear on them are separated by a comma. For example, “Select File, Open” means to pull down the File menu and choose the Open option.
- ▶ **Code-continuation character:** When a line of code is too long to fit on one line of a page, it is broken and continued to the next line. The continuation is preceded by a backslash.

## Audience

This book is designed for anyone who has a basic understanding of UNIX and wants to learn more about Solaris system administration. Whether or not you plan to become certified, this book is the starting point to becoming a Solaris system administrator. It’s the same training material that Bill uses in his Solaris 10 Intermediate System Administration course. This book covers the basics as well as intermediate system administration topics you need to know before you begin administering the Solaris operating system. Our goal was to present the material in an easy-to-follow format, with text that is easy to read and understand. The only prerequisite is that you have used UNIX, that you have attended a fundamentals of UNIX class for users, or that you have studied equivalent material so that you understand basic UNIX commands and syntax. Before you begin administering Solaris, it’s important that you have actually used UNIX.

This book is also intended for experienced system administrators who want to become certified, update their current Solaris certification, or simply learn about the features of the Solaris 10 operating environment. To pass the certification exams, you need a solid understanding of the fundamentals of administering Solaris. This book will help you review the fundamentals required to pass the certification exams.

## The Sun Certified System Administrator Exams

To become a Sun Certified System Administrator, you need to pass two exams: CX-310-200 (Part I) and CX-310-202 (Part II). This book covers the Part I exam, which is a prerequisite for Part II. You will not receive a certificate until you have passed both examinations. Also, if you are already certified in Solaris 2.6, 7, 8, or 9, you need to use materials found in volumes 1 and 2 of this series to take the upgrade exam, CX-310-203, to become certified on Solaris 10.

Beware of fakes. We have seen some websites that promote their own certification programs, so be sure to evaluate them carefully. Certification programs promoted by these sites are not

the same as the Sun certification program, and you will not receive a certificate from Sun until you pass Sun's exams from a certified Sun testing center. Go to my website (<http://www.UnixEd.com>) for links to the real exams and information on Sun's certification program if you are in doubt. In addition, feel free to visit our online Solaris certification discussion forum at <http://www.UnixEd.com>, where you can ask me questions directly.

## Exam CX-310-200

The following sections cover the Exam CX-310-200 objectives.

### Manage File Systems

List the different types of file systems and file types in the Solaris operating environment. Understand how to add disk devices to a system and the device files associated with each disk on each Solaris platform (SPARC and x86/x64). Understand how to use the `format` and `fdisk` utilities. Understand how to create, mount, and repair file systems. Understand all the configuration files associated with managing file systems.

### Install Software

Describe the methods used and the sequence of steps required to perform the Solaris 10 operating environment software installation on SPARC-, x64-, and x86-based systems. Identify the function of the package administration commands. Understand how to install, verify, and remove operating system patches.

### Perform System Boot Procedures

Understand the entire boot process, with knowledge of the various configuration files and startup scripts on SPARC-, x64-, and x86-based systems. Understand how to use and execute boot PROM commands. Understand the role of the Solaris Management Facility (SMF) in the boot process, and become familiar with SMF-related commands. Understand the function of the files or directories accessed during the boot process. Understand the commands used to change the run level of a system to a specified state.

## **Perform User and Security Administration**

Understand all aspects of administering users and groups. Understand how to set and verify file and directory permissions.

## **Manage Network Printers and System Processes**

Describe the purpose, features, and functionality of the print management tools available in the Solaris operating environment. Understand the LP print service directory structure and the Solaris operating environment printing process. Understand the commands that display information for all active processes on the system. Understand the effect of sending a specified signal to a process. Understand the various methods used to terminate an active process.

## **Perform System Backups and Restores**

Understand the functional capabilities of the various backup, archive, and restore utilities in Solaris 10. Identify the commands and steps required to back up and restore a file system. Given a specific scenario, be prepared to develop a strategy for scheduled backups and back up an unmounted file system using the appropriate commands.

## **Exam CX-310-203 (Solaris 10 Upgrade Exam)**

If you're already certified on Solaris 2.6, 7, 8, or 9, you only need to take the CX-310-203 upgrade exam to update your certification. Here are the current objectives for that exam:

- ▶ Install Software
- ▶ Manage File Systems
- ▶ Perform System Boot and Shutdown Procedures for SPARC-, x64-, and x86-Based Systems
- ▶ Perform User and Security Administration
- ▶ Perform Advanced Installation Procedures

## Summary

It's not uncommon for Sun to change the exam objectives or to shift them around after the exams have been published. I highly recommend that before you start this book, you visit my website at <http://www.UnixEd.com> to get the most up-to-date list of exam objectives, the errata for this book, up-to-date sample exam questions, and any other last-minute notes about these exams. We will provide all the information you need to pass the exam; all you need to do is devote the time. Learning the objectives is the first step; the next step is to practice. You need to have access to both SPARC- and x86/x64-based systems running Solaris 10 so that you can practice what you have learned. Unless you have a supernatural memory, it will be difficult to pass the exams without practice.

In the back of this book is the ExamGear software test CD that prepares you for the questions you might see on the exam. The CD-ROM-based test engine was designed by educational experts to help you learn as you test. It is a preview of the types of questions to expect on the exams, and it tests your knowledge on all the exam objectives. If you are weak in any area, the sample questions will help you identify that area so that you can go back to the appropriate chapter and study the topic. Each question on the CD-ROM has a flash card to help you in case you get stuck. This flash card contains brief, concise, textbook excerpts that explain why each answer is correct so that you can learn as you test.

Also, for an additional cost, you can purchase more questions for the ExamGear test engine from our website. You'll obtain hundreds of questions that will take you deep into each exam objective, providing a comprehensive skills assessment and helping you evaluate your readiness and retention of the materials.

## Advice on Taking the Exam

More extensive tips are found in the "Study and Exam Prep Tips" element and throughout the book, but keep in mind the following advice as you study for the exam:

- ▶ **Read all the material:** This book includes information not reflected in the exam objectives to better prepare you for the exam and for real-world experiences. Read all the material to benefit from this.
- ▶ **Do the step-by-step lab exercises, and complete the exercises in each chapter:** This will help you gain experience and prepare for the scenario-type questions that you will encounter.
- ▶ **Use the questions to assess your knowledge:** Each chapter contains exam questions. Use these to assess your knowledge and determine where you need to review material.

- ▶ **Review the exam objectives:** Develop your own questions and examples for each topic listed. If you can develop and answer several questions for each topic, you should not find it difficult to pass the exam.
- ▶ **Relax and sleep before taking the exam:** Your time for taking the exam is limited. However, if you have prepared and you know Solaris network administration, you will find plenty of time to answer all the questions. Be sure to rest well the night before so that you can handle the stress that time limitations put on you as you take the exam.
- ▶ **Review all the material in the “Fast Facts” element the night before or the morning you take the exam.**
- ▶ **If you don’t know the answer to a question, just skip it and don’t waste much time:** You need to complete the exam in the time allotted. Don’t be lazy during the exam; answer all the questions as quickly as possible. Any unfinished questions will be marked as incorrect.
- ▶ **Visit my website, <http://www.UnixEd.com>. It contains the following:**
  - ▶ Late-breaking changes that Sun might make to the exam or the objectives. You can expect Sun to change the exams frequently. Make sure you check my website before taking the exam.
  - ▶ A FAQs page with frequently asked questions and errata about this book or the exams.
  - ▶ Links to other informative websites.
  - ▶ Additional practice questions and sample exams for the ExamGear test engine. The ExamGear test engine has hundreds of questions that you can use to further assess your retention of the material presented in the book. The exams feature electronic flash cards that take the place of those sticky notes that you’ve used as bookmarks throughout the book. Don’t attempt the real exam until you can pass every section of the practice exams with a 95% or better score.
  - ▶ An online forum where you can discuss certification-related issues with me and other system administrators, including some who have already taken the exam.
  - ▶ Additional study materials, training programs, and online seminars related to Solaris certification.
  - ▶ You can also email me directly from this website with questions or comments about this book. I always try to answer each one.

When you feel confident, take the real exams and become certified. Don’t forget to drop me an email and let me know how you did ([guru@UnixEd.com](mailto:guru@UnixEd.com)).

# 5

CHAPTER FIVE

## Managing System Processes

---

### Objectives

The following test objectives for Exam CX-310-200 are covered in this chapter:

**Explain how to view system processes and clear frozen processes.**

- ▶ Managing system processes is a common task for any system administrator. You should know how to use the commands that display information for all active processes on the system, and how to terminate an active or deadlocked process.

**Explain how to schedule an automatic one-time execution of a command and the automatic recurring execution of a command.**

- ▶ Many processes compete for execution time so scheduling jobs to run at off-peak hours can dramatically improve system performance. The system administrator needs to understand how to use the Solaris batch processor to schedule execution of commands.



---

# Study Strategies

The following study strategies will help you prepare for the test:

- ▶ Understand each of the commands in this chapter enough so that you can match the command and option with a description. Practice them on a Solaris system so that you can become familiar with the output they produce.
- ▶ Know all the commands used to display information about a process. When viewing processes, understand each of the fields that are displayed in the output.
- ▶ Finally, understand how to schedule commands via the Solaris batch-processing facilities. Become familiar with all of the associated configuration files: what they do and how they are formatted.

## EXAM ALERT

**Managing System Processes** As of this writing, the topic of managing system processes is covered lightly on the CX-310-200 exam. This could change in the future as Sun keeps updating and changing its exams. The best approach is to be prepared and learn the material thoroughly. After all, it's a topic every system administrator needs to know to effectively perform the job.

# Introduction

This chapter covers Solaris processes—how to view processes, understand the effects signals have on processes, and how to manage processes.

## Viewing a Process

---

Objective:

**Explain how to view system processes.**

Solaris is a multitasking environment in which a number of programs run at the same time. This means that many users can be active on the system at the same time, running many jobs (processes) simultaneously. Each Solaris program can start and stop multiple processes while it is running, but only one job is active per processor at any given time while the other jobs wait in a job queue. Because each process takes its turn running in very short time slices (much less than a second each), multitasking operating systems give the appearance that multiple processes are running at the same time. A parent process forks a child process, which, in turn, can fork other processes.

### NOTE

**Forks** The term *fork* is used to describe a process started from another process. As with a fork in the road, one process turns into two. You'll also see the term *spawn* used—the two words are interchangeable for the purposes of this subject.

A program can be made up of many processes. A *process* is part of a program running in its own address space. A process under Solaris consists of an *address space* and a set of data structures in the *kernel* to keep track of that process. The address space is divided into various sections that include the instructions that the process may execute, memory allocated during the execution of the process, the *stack*, and memory-mapped files. The kernel must keep track of the following data for each process on the system:

- ▶ Address space
- ▶ Current status of the process
- ▶ Execution priority of the process
- ▶ Resource usage of the process
- ▶ Current signal mask
- ▶ Ownership of the process

A process is distinct from a job, command, or program that can be composed of many processes working together to perform a specific task. For example, a computer-aided design application is a single program. When this program starts, it spawns other processes as it runs. When a user logs in to the program, it spawns yet other processes. Each process has a process ID associated with it and is referred to as a *PID*. You can monitor processes that are currently executing by using one of the commands listed in Table 5.1.

**TABLE 5.1 Commands to Display Processes**

| Command                       | Description   |
|-------------------------------|---|
| <code>mpstat</code>           | Executed from the command line, <code>mpstat</code> reports processor statistics in tabular form. Each row of the table represents the activity of one processor. |
| <code>ps</code>               | Executed from the command line to display information about active processes.   |
| <code>pgrep</code>            | Executed from the command line to find processes by a specific name or attribute.   |
| <code>prstat</code>           | Executed from the command line to display information about active processes on the system.   |
| <code>ptree</code>            | Prints the process trees containing the specified pids or users, with child processes indented from their respective parent processes.                            |
| <code>sdtprocess</code>       | A GUI used to display and control processes on a system. This utility requires the X Window System (also known as X Windows).                                     |
| <code>SMC process tool</code> | A GUI available in the Solaris Management Console used to monitor and manage processes on a system.   |
| <code>pargs</code>            | Executed from the command line to examine the arguments and environment variables of a process.   |
| <code>svcs</code>             | With the <code>-p</code> option, this Service Management Facility command will list processes associated with each service instance.                              |
| <code>time</code>             | Time a simple command.  |

Before getting into the commands used to monitor processes, you first need to become familiar with process attributes. A process has certain attributes that directly affect execution. These are listed in Table 5.2.

**TABLE 5.2 Process Attributes**

| Attribute             | Description   |
|-----------------------|---|
| <code>PID</code>      | The process identification (a unique number that defines the process within the kernel) |
| <code>PPID</code>     | The parent PID (the parent of the process)  |
| <code>UID</code>      | The user ID number of the user who owns the process                                     |
| <code>EUID</code>     | The effective user ID of the process  |
| <code>GID</code>      | The group ID of the user who owns the process   |
| <code>EGID</code>     | The effective group ID that owns the process  |
| <code>Priority</code> | The priority at which the process runs  |

Use the `ps` command to view processes currently running on the system. Use the `ps` command when you're on a character-based terminal and don't have access to a graphical display. Adding the `-l` option to the `ps` command displays a variety of other information about the processes currently running, including the state of each process (listed under `S`). The codes used to show the various process states are listed in Table 5.3.

**TABLE 5.3 Process States**

| Code | Process State | Description  |
|------|---------------|--|
| O    | Running       | The process is running on a processor.                                       |
| S    | Sleeping      | The process is waiting for an event to complete.                             |
| R    | Runnable      | The process is on the run queue.   |
| Z    | Zombie state  | The process was terminated and the parent is not waiting.                    |
| T    | Traced        | The process was stopped by a signal because the parent is tracing it.        |
| W    | Waiting       | The process is waiting for CPU usage to drop to the CPU-caps enforced limits |

To see all the processes that are running on a system, type the following:

```
ps -e1
```

The system responds with the following output:

```
# ps -e1
```

```
F S UID PID PPID C PRI NI ADDR SZ WCHAN TTY TIME CMD
19 T 0 0 0 0 0 SY ? 0 ? ? 0:18 sched
 8 S 0 1 0 0 40 20 ? 150 ? ? 0:00 init
19 S 0 2 0 0 0 SY ? 0 ? ? 0:00 pageout
19 S 0 3 0 0 0 SY ? 0 ? ? 0:01 fsflush
 8 S 0 309 1 0 40 20 ? 217 ? ? 0:00 sac
 8 S 0 315 1 0 40 20 ? 331 ? ? 0:00 sshd
 8 S 0 143 1 0 40 20 ? 273 ? ? 0:00 rpcbind
 8 S 0 51 1 0 40 20 ? 268 ? ? 0:00 sysevent
 8 S 0 61 1 0 40 20 ? 343 ? ? 0:01 picld
 8 S 0 453 403 0 50 20 ? 1106 ? ? 0:00 dtfile
 8 S 0 189 1 0 40 20 ? 509 ? ? 0:00 automoun
 8 S 0 165 1 0 40 20 ? 292 ? ? 0:00 inetd
 8 S 0 200 1 0 40 20 ? 415 ? ? 0:00 syslogd
 8 S 0 180 1 0 40 20 ? 266 ? ? 0:00 lockd
 8 S 0 219 1 0 40 20 ? 391 ? ? 0:00 lpsched
 8 S 1 184 1 0 40 20 ? 306 ? ? 0:00 statd
 8 S 0 214 1 0 40 20 ? 365 ? ? 0:00 nscd
 8 S 0 204 1 0 40 20 ? 254 ? ? 0:00 cron
 8 S 0 232 1 0 40 20 ? 173 ? ? 0:00 powerd
 8 S 0 255 254 0 40 20 ? 215 ? ? 0:00 smcboot
 8 S 0 258 1 0 40 20 ? 356 ? ? 0:02 vold
```

The manual page for the `ps` command describes all the fields displayed with the `ps` command, as well as all the command options. Table 5.4 lists some important fields.

**TABLE 5.4 Process Fields**

| Field | Description   |
|-------|---|
| F     | Flags associated with the process.  |
| S     | The state of the process.   |
| UID   | The user ID of the process owner. For many processes, this is 0 because they run <code>setuid</code> .  |
| PID   | The process ID of each process. This value should be unique. Generally, PIDs are allocated lowest to highest, but they wrap at some point. This value is necessary for you to send a signal, such as the <code>kill</code> signal, to a process.                            |
| PPID  | The parent process ID. This identifies the parent process that started the process. Using the PPID enables you to trace the sequence of process creation that took place.   |
| PRI   | The priority of the process. Without the <code>-c</code> option, higher numbers mean lower priority. With the <code>-c</code> option, higher numbers mean higher priority.  |
| NI    | The <code>nice</code> value, used in priority computation. This is not printed when the <code>-c</code> option is used. The process's <code>nice</code> number contributes to its scheduling priority. Making a process nicer means lowering its priority.                  |
| ADDR  | The memory address of the process.  |
| SZ    | The <code>SIZE</code> field. This is the total number of pages in the process. Page sizes are 8192 bytes on <code>sun4u</code> systems, but vary on different hardware platforms. Issue the <code>/usr/bin/pagesize</code> command to display the page size on your system. |
| WCHAN | The address of an event for which the process is sleeping (if it's -, the process is running).  |
| STIME | The starting time of the process (in hours, minutes, and seconds).  |
| TTY   | The terminal assigned to your process.  |
| TIME  | The cumulative CPU time used by the process in minutes and seconds.   |
| CMD   | The command that generated the process.   |

You often want to look at all processes. You can do this using the command `ps -e1`. A number of options available with the `ps` command control what information gets printed. A few of them are listed in Table 5.5.

**TABLE 5.5 ps Command Options**

| Option | Description   |
|--------|---|
| -A     | Lists information for all processes. Identical to the <code>-e</code> option.   |
| -a     | Lists information about all the most frequently requested processes. Processes not associated with a terminal will not be listed. |
| -e     | Lists information about every process on the system.  |

(continues)

**TABLE 5.5** *Continued*

| Option                     | Description  |
|----------------------------|--|
| -f                         | Generates a full listing.  |
| -l                         | Generates a long listing.  |
| -P                         | Prints the number of the processor to which the process is bound, if any, under an additional column header PSR. This is a useful option on systems that have multiple processors. |
| -u <i>&lt;username&gt;</i> | Lists only process data for a particular user. In the listing, the numerical user ID is printed unless you give the -f option, which prints the login name.                        |

For a complete list of options to the `ps` command, refer to the Solaris online manual pages.

**NOTE**

**sort Command** The `sort` command is useful when you're looking at system processes. Use the `sort` command as the pipe output to sort by size or PID. For example, to sort by the SZ field, use the command `ps -e1 | sort +9` (remember, `sort` starts numbering fields with 0).

## pgrep

The `pgrep` command replaces the combination of the `ps`, `grep`, `egrep`, and `awk` commands that were used to manage processes in earlier releases of Solaris. The `pgrep` command examines the active processes on the system and reports the process IDs of the processes whose attributes match the criteria you specify on the command line. The command syntax for the `pgrep` command is shown here:

```
pgrep <options> <pattern>
```

`pgrep` options are described in Table 5.6.

**TABLE 5.6** `pgrep` Options

| Option                     | Description  |
|----------------------------|--|
| -d <i>&lt;delim&gt;</i>    | Specifies the output delimiter string to be printed between each matching process ID. If no -d option is specified, the default is a newline character.  |
| -f                         | The regular expression pattern should be matched against the full process argument string. If no -f option is specified, the expression is matched only against the name of the executable file. |
| -g <i>&lt;pgrplist&gt;</i> | Matches only processes whose process group ID is in the given list.  |
| -G <i>&lt;gidlist&gt;</i>  | Matches only processes whose real group ID is in the given list. Each group ID may be specified as either a group name or a numerical group ID.  |

**TABLE 5.6** *Continued*

| Option        | Description  |
|---------------|--|
| -l            | Long output format. Prints the process name along with the process ID of each matching process.  |
| -n            | Matches only the newest (most recently created) process that meets all other specified matching criteria.  |
| -P <ppidlist> | Matches only processes whose parent process ID is in the given list.   |
| -s <sidlist>  | Matches only processes whose process session ID is in the given list.  |
| -t <termlist> | Matches only processes that are associated with a terminal in the given list. Each terminal is specified as the suffix following /dev/ of the terminal's device pathname in /dev (for example, term/a or pts/0). |
| -u <euidlist> | Matches only processes whose effective user ID is in the given list. Each user ID may be specified as either a login name or a numerical user ID.  |
| -U <uidlist>  | Matches only processes whose real user ID is in the given list. Each user ID may be specified as either a login name or a numerical user ID.   |
| -v            | Matches all processes except those that meet the specified matching criteria.  |
| -x            | Considers only processes whose argument string or executable filename exactly matches the specified pattern.   |
| <pattern>     | A pattern to match against either the executable filename or full process argument string.   |

For example, the following `pgrep` command finds all processes that have "dt" in the process argument string:

```
# pgrep -l -f "dt"
```

The system responds with this:

```
500 /usr/dt/bin/dtlogin -daemon
16224 ./dtterm
438 /usr/dt/bin/dtlogin -daemon
448 /usr/openwin/bin/Xsun :0 -defdepth 24 -nobanner -auth /var/dt/A:0-p_aW2a
520 dtgreet -display :0
```

To find the process ID for the `lpsched` process, issue this command:

```
# pgrep -l lpsched
```

The system responds with this:

```
6899 lpsched
```

## prstat

Use the `prstat` command from the command line to monitor system processes. Again, like the `ps` command, it provides information on active processes. The difference is that you can specify whether you want information on specific processes, UIDs, CPU IDs, or processor sets. By default, `prstat` displays information about all processes sorted by CPU usage. Another nice feature with `prstat` is that the information remains on the screen and is updated periodically. The information displayed by the `prstat` command is described in Table 5.7.

**TABLE 5.7 Column Headings for the `prstat` Command**

| Column Heading | Description  |
|----------------|--|
| PID            | The process identification (a unique number that defines the process within the kernel)  |
| USERNAME       | The login ID name of the owner of the process  |
| SIZE           | The total virtual memory size of the process in kilobytes (K), megabytes (M), or gigabytes (G)   |
| RSS            | The resident set size of the process in kilobytes, megabytes, or gigabytes   |
| STATE          | The state of the process:<br><code>cpu&lt;n&gt;</code> —Process is running on CPU.<br><code>sleep</code> —Process is waiting for an event to complete.<br><code>run</code> —Process is in the run queue.<br><code>zombie</code> —Process has terminated and parent is not waiting.<br><code>stop</code> —Process is stopped. |
| PRI            | The priority of the process  |
| NICE           | The value used in priority computation   |
| TIME           | The cumulative execution time for the process  |
| CPU            | The percentage of recent CPU time used by the process  |
| PROCESS        | The name of the process  |
| NLWP           | The number of lightweight processes (LWPs) in the process  |

This section will introduce some new terminology, so Table 5.8 defines a few terms related to processing in general.

**TABLE 5.8 Process Terminology**

| Term                      | Description   |
|---------------------------|---|
| Multitasking              | <p>A technique used in an operating system for sharing a single processor among several independent jobs.</p> <p>Multitasking introduces overhead because the processor spends some time choosing the next job to run and saving and restoring tasks' state. However, it reduces the worst-case time from job submission to completion compared with a simple batch system, in which each job must finish before the next one starts. Multitasking also means that while one task is waiting for some external event, the CPU is free to do useful work on other tasks.</p> <p>A multitasking operating system should provide some degree of protection of one task from another to prevent tasks from interacting in unexpected ways, such as accidentally modifying the contents of each other's memory areas.</p> <p>The jobs in a multitasking system may belong to one or many users. This is distinct from parallel processing, in which one user runs several tasks on several processors. Time sharing is almost synonymous with multitasking, but it implies that there is more than one user.</p> |
| Parallel processing       | <p>The simultaneous use of more than one CPU to solve a problem. The processors either may communicate to cooperate in solving a problem or may run completely independently, possibly under the control of another processor that distributes work to the others and collects results from them.</p>   |
| Multithreaded             | <p>Multithreaded is a process that has multiple flows (threads) of control. The traditional Unix process contained, and still contains, a single thread of control. Multithreading (MT) separates a process into many execution threads, each of which runs independently. For more information, see the Multithreaded Programming Guide at <a href="http://docs.sun.com/">http://docs.sun.com/</a> Part number 816-5137-10.</p>  |
| Lightweight process (LWP) | <p>A single-threaded subprocess. LWPs are scheduled by the kernel to use available CPU resources based on their scheduling class and priority. LWPs include a kernel thread, which contains information that must be in memory all the time, and a LWP, which contains information that is swappable. A process can consist of multiple LWPs and multiple application threads. A lightweight process is somewhere between a thread and a full process.</p>  |
| Application thread        | <p>A series of instructions with a separate stack that can execute independently in a user's address space. The threads can be multiplexed on top of LWPs.</p>  |
| Address space             | <p>The range of addresses that a processor or process can access, or at which a device can be accessed. The term may refer to either a physical address or a virtual address. The size of a processor's address space depends on the width of the processor's address bus and address registers. Processes running in 32-bit mode have a 4 gigabyte address space (<math>2^{32}</math> bytes) and processes running in 64-bit mode have a 16 terabyte (<math>2^{64}</math> bytes) address space.</p>  |
| Shared memory             | <p>Usually refers to RAM, which can be accessed by more than one process in a multitasking operating system with memory protection.</p>   |

The syntax for the `prstat` command is as follows:

```
prstat [options] <count> <interval>
```

Table 5.9 describes a few of the `prstat` command options and arguments.

**TABLE 5.9 prstat Options and Arguments**

| Option                | Description  |
|-----------------------|--|
| <b>prstat Options</b> |  |
| -a                    | Displays separate reports about processes and users at the same time.  |
| -c                    | Continuously prints new reports beneath previous reports instead of overwriting them.  |
| -j <projlist>         | Reports only processes or LWPs whose project ID is in the given list. Each project ID can be specified as either a project name or a numerical project ID.   |
| -J                    | Reports information about processes and projects.  |
| -k <tasklist>         | Reports only processes or LWPs whose task ID is in <i>tasklist</i> .   |
| -m                    | Reports microstate process accounting information. In addition to all fields listed in -v mode, this mode also includes the percentage of time the process has spent processing system traps, text page faults, and data page faults, and waiting for user locks and waiting for CPU (latency time). |
| -n <nproc>            | Restricts the number of output lines. The <nproc> argument specifies how many lines of process or LWP statistics are reported.   |
| -p <pidlist>          | Reports only processes that have a PID in the given list.  |
| -P <cpulist>          | Reports only processes or LWPs that have most recently executed on a CPU in the given list. The <cpulist> argument identifies each CPU by an integer as reported by <code>psrinfo</code> .   |
| -S <key>              | Sorts output lines by <key> in descending order. Values for <key> can be<br>cpu—Sorts by process CPU usage. This is the default.<br>time—Sorts by process execution time.<br>size—Sorts by size of process image.<br>rss—Sorts by resident set size.<br>pri—Sorts by process priority.               |
| -s <key>              | Sorts output lines by <key> in ascending order. See the -S option for a list of valid keys to use.   |
| -t                    | Reports total usage summary for each user.   |
| -u <uidlist>          | Reports only processes whose effective user ID is in the given list. The value for <uidlist> may be specified as either a login name or a numerical user ID.   |

TABLE 5.9 *Continued*

| Option                  | Description   |
|-------------------------|---|
| -U <uidlist>            | Reports only processes whose real user ID is in the given list. The value for <uidlist> may be specified as either a login name or a numerical user ID. |
| <b>prstat Arguments</b> |   |
| <count>                 | Specifies the number of times that the statistics are repeated. By default, prstat reports statistics until a termination signal is received.           |
| <interval>              | Specifies the sampling interval in seconds; the default interval is 5 seconds.  |

**NOTE**

**psrinfo Command** psrinfo displays one line for each configured processor, displaying whether it is online, non-interruptible, offline, or powered off, as well as when that status last changed.

The following example uses the prstat command to view the four most active root processes running. The -n option is used here to restrict the output to the top four processes. The next number, 5, specifies the sampling interval in seconds, and the last number, 3, runs the command three times:

```
# prstat -u root -n 4 5 3
```

The system displays the following output:

```
PID USERNAME  SIZE  RSS  STATE  PRI  NICE  TIME    CPU  PROCESS/NLWP
4375 root      4568K 4344K cpu0    59   0    0:00:00 0.4% prstat/1
4298 root      7088K 5144K sleep   59   0    0:00:02 0.2% dtterm/1
 304 root      2304K 1904K sleep   59   0    0:02:35 0.0% mibiisa/7
 427 root      1832K 1304K sleep   59   0    0:00:00 0.0% rpc.rstatd/1
Total: 53 processes, 111 lwps, load averages: 0.02, 0.01, 0.01
```

The output updates on your display five times every three seconds.

I described projects in Chapter 4, “User and Security Administration,” where user accounts can be assigned to project groups. These projects can also be used to label workloads and separate projects and a project’s related processes from one another.

The project provides a networkwide administrative identifier for related work. A project consists of tasks, which collect a group of processes into a manageable entity that represents a workload component.

You can use the `prstat` command with the `-J` option to monitor the CPU usage of projects and the `-k` option to monitor tasks across your system. Therefore, you can have `prstat` report on the processes related to a project rather than just list all system processes. In addition, the system administrator can set processing limits on the project, such as setting a limit on the total amount of physical memory, in bytes, that is available to processes in the project. For more information on projects and resource capping, read the man pages on the following commands: `rcapd(1M)`, `project(4)`, `rcapstat(1)`, and `rcapadm(1M)`.

## mpstat

Use the `mpstat` command to report processor statistics on a multi-processor system. When executing the `mpstat` command, we'll usually want to see more than one result, so we specify the number of seconds between each `mpstat` as follows:

```
mpstat 30
```

The argument `30`, specifies that I want to get a report every 30 seconds. The system displays the following information every 30 seconds:

| CPU | minf | mjf | xcal | intr | ithr | csw | icsw | migr | smtx | srw | syscl | usr | sys | wt | idl |
|-----|------|-----|------|------|------|-----|------|------|------|-----|-------|-----|-----|----|-----|
| 0   | 6    | 0   | 0    | 114  | 14   | 25  | 0    | 6    | 3    | 0   | 48    | 1   | 2   | 25 | 72  |
| 1   | 6    | 0   | 0    | 86   | 85   | 50  | 0    | 6    | 3    | 0   | 66    | 1   | 4   | 24 | 71  |
| 2   | 7    | 0   | 0    | 42   | 42   | 31  | 0    | 6    | 3    | 0   | 54    | 1   | 3   | 24 | 72  |
| 3   | 8    | 0   | 0    | 0    | 0    | 33  | 0    | 6    | 4    | 0   | 54    | 1   | 3   | 24 | 72  |

The results are from a system with four processors. Typically, a system administrator will use the `mpstat` command to check CPU utilization. In this example, I look at the `idl` column (percent idle time) and see that the server's CPUs are approx 28% used. For more information on the other columns of information, refer to the `mpstat` man pages.

## ptree

The `ptree` command will display the process tree. The parent process is displayed with the respective child processes indented beneath it. Here is an example showing the processes that belong to the `inetd` process (PID 270):

```
# ptree 270<cr>
```

The system displays:

```
270  /usr/lib/inet/inetd start
    780  /usr/sbin/in.telnetd
      783  -sh
        1250  ptree 270
```

With no arguments, the `ptree` command will display every process along with the associated child processes.

## time

The `time` command is used to display the time that the system has spent executing a command. It's a useful command for benchmarking performance. Use this command to time a command on a particular system configuration and compare to another system. In the following example, I'll check the system processing time for a script I wrote named "longtime":

```
# time ./longtime<cr>
```

The system displays:

```
real    14.7
user    9.9
sys     2.3
```

The `real` time is the total time that has elapsed between invoking the script and its termination. The `user` time is the time the processor spends executing your user code. Finally, the `system` time is the time the processor spends executing Operating System code on behalf of your process.

## Process Manager

In the Desktop Environment (CDE & JAVA Desktop) you have access to the Process Manager GUI, `sdtprocess`, a graphical tool that provides a process manager window for monitoring and controlling system processes.

### EXAM ALERT

The exam will most likely ask you about the command-line tools used to manage system processes, such as `kill`, `killall`, `ps`, and `pgrep`. You only need to understand that GUI tools can be used to manage processes and you should be prepared to identify these GUI tools.

The advantage of using the Process Manager is that you can view and control processes without knowing all the complex options associated with the `ps` and `kill` commands. For example, you can display processes that contain specific character strings, and you can sort the process list alphabetically or numerically. You can initiate a search using the `find` command, or you can terminate a process simply by highlighting it and clicking `kill`.

To open the Process Manager, you need to log into the Desktop windowing environment. You can start the GUI by executing the command `sdtprocess`, as follows:

```
# sdtprocess &
```

Or, you can click Find Process on the Tools subpanel, as shown in Figure 5.1.

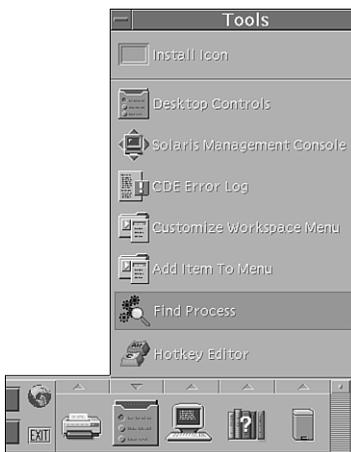


FIGURE 5.1 Front panel.

The Process Manager window opens, as shown in Figure 5.2.

| ID   | Name     | Owner | CPU% | RAM   | Size  | Started  | Parent | Command   |
|------|----------|-------|------|-------|-------|----------|--------|---|
| 4482 | sdtimage | root  | 2.4  | 15104 | 20096 | 12:57:57 | 4481   | /usr/dt/bin/sdtimage -snapshot                      |
| 314  | Xsun     | root  | 1.8  | 13192 | 20120 | Mar_18   | 293    | /usr/openwin/bin/Xsun :0 -nobanner -auth /var/dt/A: |
| 399  | dtwm     | root  | 0.6  | 7760  | 9440  | Mar_18   | 392    | dtwm  |
| 4493 | ps       | root  | 0.3  | 904   | 1136  | 12:59:39 | 4492   | /usr/bin/ps -A -o pid=ID -o fname=Name -o user=Ownr |
| 4492 | ksh      | root  | 0.1  | 1264  | 1888  | 12:59:39 | 4380   | ksh -c /usr/bin/ps -A -o pid=ID -o 'fname=Name' -o  |
| 4380 | sdtproce | root  | 0.1  | 6604  | 8952  | 12:51:08 | 418    | sdtprocess  |
| 304  | mb1isa   | root  | 0.1  | 1912  | 2312  | Mar_18   | 285    | mb1isa -r -p 32790                                  |
| 4481 | dtexec   | root  | 0.0  | 2360  | 3104  | 12:57:57 | 403    | /usr/dt/bin/dtexec -open 0 -ttprocid 3.yba-3 01 39: |
| 4300 | sh       | root  | 0.0  | 216   | 328   | 12:03:19 | 4298   | /sbin/sh  |
| 4298 | dtterm   | root  | 0.0  | 5144  | 7088  | 12:03:19 | 1      | /usr/dt/bin/dtterm -display 192.168.0.3:0           |
| 1659 | cat      | root  | 0.0  | 744   | 976   | Mar_19   | 406    | /bin/cat /tmp/.removable/notify406                  |

**Figure 5.2** Process Manager window.

Each process attribute in the header of the Process Manager window provides detailed information about the process and is described in Table 5.10.

**TABLE 5.10** Process Manager Window

| Column Heading | Description   |
|----------------|---|
| ID             | Process ID  |
| Name           | Process name  |
| Owner          | Login ID name of the owner of the process                                 |
| CPU%           | Ratio of CPU time available in the same period, expressed as a percentage |
| RAM            | Amount of RAM currently occupied by this process                          |
| Swap           | Total swap size in virtual memory   |
| Started        | Actual start time (or date, if other than current)                        |
| Parent         | Process ID of parent process, or PPID                                     |
| Command        | Actual Unix command (truncated) being executed                            |

Click any of the column headings to sort the processes by that attribute. For example, click the CPU heading to sort all processes by their CPU usage. The list updates every 30 seconds, but you can enter a value in the Sampling field to update the list as frequently as you like. Finally, you can enter a text string that is common to the process entries of all the processes you want to display in the Find drop-down menu. In Figure 5.3, I entered “root” in the Find field to display all processes owned by root. I also changed the sampling rate to every 5 seconds and clicked the CPU heading to sort processes by their CPU usage.

Another nice feature of the Process Manager is the capability to display the ancestry of a process. When a Unix process initiates one or more processes, these are *child processes*, or children. Child and parent processes have the same user ID. To view a parent process and all the child processes that belong to it, highlight the process in the Process Manager window. Click Process from the toolbar at the top of the window and select Show Ancestry, as shown in Figure 5.4.

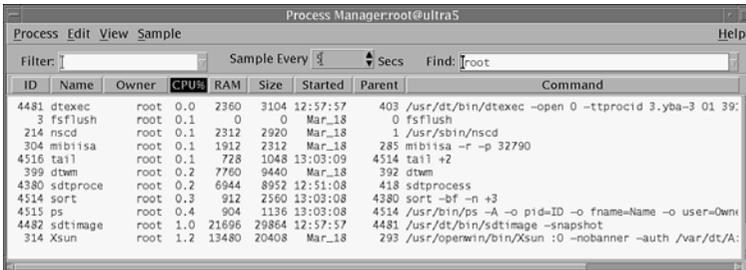


FIGURE 5.3 Sorted Process Manager window.

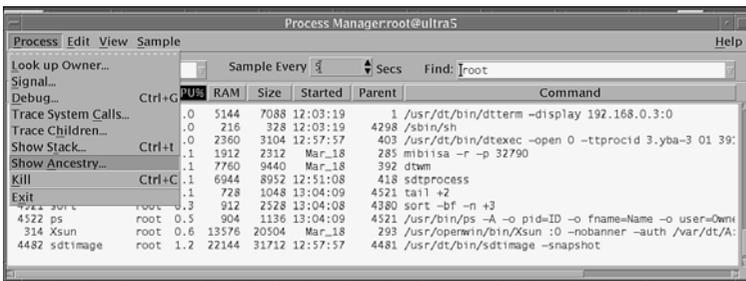


FIGURE 5.4 Selecting Show Ancestry.

The window shown in Figure 5.5 displays showing all the processes belonging to the parent. The command-line equivalent to the Ancestry selection in the Process Manager is the `ptree` command. Use this command when you don't have a graphical display terminal. The `ptree` command displays the process ancestry trees containing the specified PIDs or users. The child processes are displayed indented from their respective parent processes. For example, here is the process tree for the `-sh` process, which has a PID of 293:

```
# ptree 293
```

```

293 /usr/dt/bin/dtlogin -daemon
  316 /usr/dt/bin/dtlogin -daemon
    333 /bin/ksh /usr/dt/bin/Xsession
      376 /usr/dt/bin/sdt_shell -c unset DT; DISPLAY=:0; /u
        379 -sh -c unset DT; DISPLAY=:0; /usr/dt/bin/dtsses
          392 /usr/dt/bin/dtsession
            403 dtfile -session dtM/aGwj
              4481 /usr/dt/bin/dtexec -open 0 -ttprocid 3.yba-3 01 391 1289637
                4482 /usr/dt/bin/sdtimage -snapshot
  
```

**Figure 5.5** Show Ancestry window.

The system responds with this:

```

293 /usr/dt/bin/dtlogin -daemon
  316 /usr/dt/bin/dtlogin -daemon
    333 /bin/ksh /usr/dt/bin/Xsession
      376 /usr/dt/bin/sdt_shell -c unset DT;DISPLAY=:0;/usr/dt/bin/dt
        379 -sh -c unset DT; DISPLAY=:0; usr/dt/bin/dtsession_res - \
          merge
          392 /usr/dt/bin/dtsession
            402 /usr/dt/bin/dtterm -session dthIaGth -C -ls
              418 -sh
  
```

## SMC Process Tool

The Solaris Management Console (SMC) includes a GUI called the Process Tool, which is used for viewing and managing processes, similar to the Desktop Process Manager tool described in the previous section. You can use the job scheduler tool to

- ▶ Suspend a process
- ▶ Resume a suspended process
- ▶ Kill a process
- ▶ Display information about a process

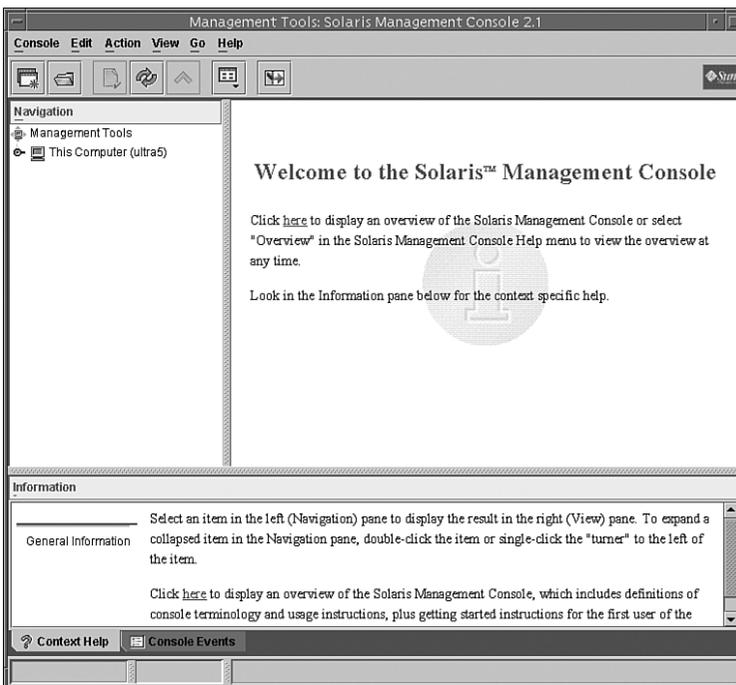
To open the Process Tool, follow Step by Step 5.1.

## STEP BY STEP

### 5.1 Opening the Process Tool

1. Start up the Solaris Management Console in the background by typing  

```
# smc &
```
2. The SMC Welcome window appears as shown in Figure 5.6.
3. In the SMC navigation pane, open the Process Tool by clicking on the This Computer icon, then click on the System Status icon, then click on the Processes icon as shown in Figure 5.7.
4. The Process Tool displays as shown in Figure 5.8



**FIGURE 5.6** SMC Welcome Window.

The Process Tool works much the same way as the Process Manager tool described earlier.

## pargs

The `pargs` command is used from the command line to examine the arguments and environment variables of a process (or number of processes). `pargs` can also be used to examine core files.

The syntax for the `pargs` command is as follows:

```
pargs [options] [pid | core]
```

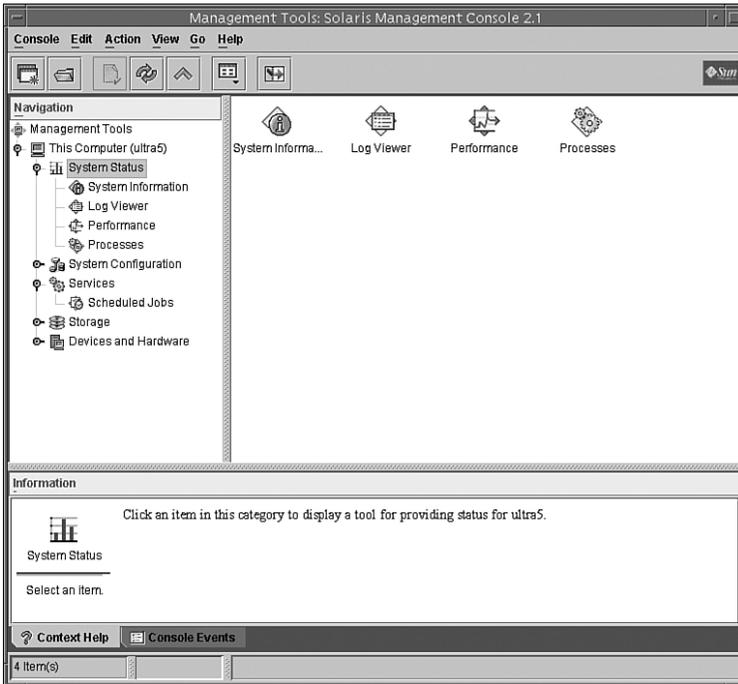


FIGURE 5.7 Opening the Job Scheduler.

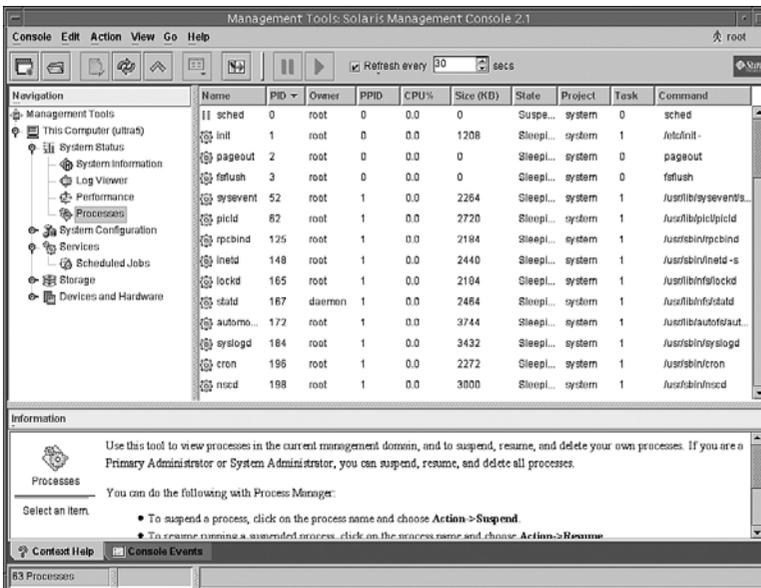


FIGURE 5.8 Process Tool.

Table 5.11 describes the `pargs` command options and arguments.

**TABLE 5.11** `pargs` Options and Arguments

| Option/Arguments | Description   |
|------------------|---|
| -a               | Prints the process arguments.   |
| -c               | Treats strings in the target process as though they were encoded in 7-bit ASCII.                |
| -e               | Prints process environment variables and values.  |
| -F               | Force. Grabs the target process even if another process has control.                            |
| -x               | Prints process auxiliary vector.  |
| <pid>            | Process ID list. The PID list can be a single process ID or multiple PIDs separated by a space. |
| core             | Processes a core file.  |

For example, to use the `pargs` command to view all of the environment variables associated with the `telnetd` process, I first need to find the PID of the `telnetd` process using `pgrep` as follows:

```
# pgrep telnetd
16173
```

Next, I issue the `pargs` command using the PID for the `telnetd` process as an argument:

```
# pargs -e 16173
```

The system responds with

```
16173: /usr/sbin/in.telnetd
envp[0]: SMF_RESTARTER=svc:/network/inetd:default
envp[1]: SMF_FMRI=svc:/network/telnet:default
envp[2]: SMF_METHOD=inetd_start
envp[3]: PATH=/usr/sbin:/usr/bin
envp[4]: TZ=US/Michigan
```

## SVCS

The Service Management Facility (SMF) was described in Chapter 3, “Perform System Boot and Shutdown Procedures for SPARC, x64, and x86-Based Systems,” so I won’t be redundant by describing it again here. However, this is just a reminder that you can use the `svcs` command with the `-p` option to list all processes associated with each service instance.

# Process Types

When sitting at a terminal and typing in commands, the user is typically executing *foreground processes*. Commands such as `vi` are foreground processes—they read input from the keyboard and display output to the terminal. Foreground processes maintain control of the terminal, and the user cannot do anything else in that terminal window until the execution of that command is complete.

Some processes are not interactive and don't need to run in the foreground. These are referred to as *background processes* or *jobs*. A background process gets detached from the terminal, freeing up the terminal while it is running. When a user decides to run a process in the background, you must arrange for the process to get its input from another source. In addition, you need to arrange for the process to output to a device other than the terminal, such as a file.

To run a process in the background, enter an `&` (ampersand) after the command:

```
# find . -name core -print &
```

After typing in this command, you're returned to a command prompt. The `find` command executes in the background. One problem, however, is the standard output is still on your terminal. In other words, as the `find` command executes, the results still are displayed on your screen, which can become quite annoying. It's best to redirect the output to a file, as follows:

```
# find . -name core -print > /tmp/results &
```

After you put the `find` command in the background, the system displays two numbers associated with that process—the job number and the process ID number (PID) as follows:

```
[1] 14919
```

You use this job number to control background processes.

## NOTE

**No Job Control in the `sh` shell** The Bourne shell does not provide job control. Job control enables you to check and manage your background jobs. Thus, with the Bourne shell, you can submit jobs to the background, but you cannot manage them. Use `jsh` (job shell), which provides all the functionality of `sh` and enables job control. The Korn shell (`ksh`) and the C shell (`csh`) both allow for job control.

The shell maintains a table containing information about processes that are currently in the background. This is referred to as the *jobs table*. The jobs table is unique to the user, and each user has his own jobs table. Furthermore, the jobs table contains only entries for jobs that are running in your current shell. If you start a new shell, the jobs table for the new shell is empty. Each job in the table is assigned a number that is unique to that user only. In other words, two users can each have a job numbered 1. Don't confuse this job number with a process ID number; remember, process IDs are unique, and no two share the same number. Any jobs that the user has placed in the background are displayed here by typing in the `jobs` command, as follows:

```
# jobs
```

The system responds with this:

```
[3] + Running find / -name bill -print > /tmp/results3 &  
[2] - Running find / -name junk -print > /tmp/results2 &  
[1] Running find / -name core -print > /tmp/results1 &
```

The jobs table contains the following information:

- ▶ A numeric value for each job
- ▶ A + (plus) symbol to designate the current job that user commands will operate on
- ▶ A - (minus) symbol to designate the next job that the user commands will operate on
- ▶ The status of the job
- ▶ The name of the job

Each job in the job table has one of the following states:

- ▶ Running—An active job
- ▶ Stopped—A job that has been suspended
- ▶ Terminated—A job that has been killed
- ▶ Done—A completed job

When the job finishes, the following is displayed on your terminal:

```
[1] + Done find / -name core -print > /tmp/results &
```

Note the job number of 1 and the status of Done.

If you want to terminate a job, use the `kill` command followed by a % (percent sign) and then the job number, as follows:

```
# kill %1
```

### CAUTION

Pay special attention to the use of the % (percent) symbol—it's absolutely required. Without it, you could kill the wrong process and potentially crash the system. Get familiar with the `kill` command in the next section of this chapter before you use it.

If you do not enter a number following the % sign, the command acts upon the current job entry listed in the jobs table. For this example, you are going to kill job number 1, as follows:

```
# kill %1
```

The following message is displayed indicating successful termination:

```
[1] + Terminated find / -name core -print > /tmp/results &
```

You can also bring a job back into the foreground with the `fg` command. Typing `fg` with no arguments brings the current job (the job with the `+` sign next to it in the jobs table) into the foreground. You can also specify the job by typing `fg %<job number>`, as follows:

```
# fg %2
```

This brings job 2 back into the foreground on your terminal.

In a windowing environment such as Java Desktop System, placing jobs in the background is not an issue. Typically, you start a job in one window and open another window to continue working. Therefore, placing jobs into the background has all but disappeared unless you are working on a character-based terminal.

## Using Signals

Objective:

### Clearing frozen processes.

Solaris supports the concept of sending software *signals* to a process. These signals are ways for other processes to interact with a running process outside the context of the hardware. The `kill` command is used to send a signal to a process. System administrators most often use the signals `SIGHUP`, `SIGKILL`, `SIGSTOP`, and `SIGTERM`. The `SIGHUP` signal is used by some utilities as a way to notify the process to do something, such as re-read its configuration file. The `SIGHUP` signal is also sent to a process if the remote connection is lost or hangs up. The `SIGKILL` signal is used to abort a process, and the `SIGSTOP` signal is used to pause a process. The `SIGTERM` signal is the default signal sent to processes by commands such as `kill` and `pkill` when no signal is specified. Table 5.12 describes the most common signals an administrator is likely to use.

#### EXAM ALERT

Don't worry about remembering all of the signals listed; just be familiar with the more common signals, such as `SIGHUP`, `SIGKILL`, `SIGSTOP`, and `SIGTERM`.

**TABLE 5.12** Signals Available Under Solaris

| Signal              | Number | Description   |
|---------------------|--------|---|
| <code>SIGHUP</code> | 1      | Hangup. Usually means that the controlling terminal has been disconnected.                            |
| <code>SIGINT</code> | 2      | Interrupt. The user can generate this signal by pressing <code>Ctrl+C</code> or <code>Delete</code> . |

(continues)

**TABLE 5.12** *Continued*

| <b>Signal</b> | <b>Number</b> | <b>Description</b>   |
|---------------|---------------|--|
| SIGQUIT       | 3             | Quits the process and produces a core dump.                          |
| SIGILL        | 4             | Illegal instruction.   |
| SIGTRAP       | 5             | Trace or breakpoint trap.  |
| SIGABRT       | 6             | Abort.   |
| SIGEMT        | 7             | Emulation trap.  |
| SIGFPE        | 8             | Arithmetic exception. Informs a process of a floating-point error.   |
| SIGKILL       | 9             | Killed. Forces the process to terminate. This is a sure kill.        |
| SIGBUS        | 10            | Bus error.   |
| SIGSEGV       | 11            | Segmentation fault.  |
| SIGSYS        | 12            | Bad system call.   |
| SIGPIPE       | 13            | Broken pipe.   |
| SIGALRM       | 14            | Alarm clock.   |
| SIGTERM       | 15            | Terminated. A gentle kill that gives processes a chance to clean up. |
| SIGUSR1       | 16            | User signal 1.   |
| SIGUSR2       | 17            | User signal 2.   |
| SIGCHLD       | 18            | Child status changed.  |
| SIGPWR        | 19            | Power fail or restart.   |
| SIGWINCH      | 20            | Window size change.  |
| SIGURG        | 21            | Urgent socket condition.   |
| SIGPOLL       | 22            | Pollable event.  |
| SIGSTOP       | 23            | Stopped (signal). Pauses a process.                                  |
| SIGTSTP       | 24            | Stopped (user).  |
| SIGCONT       | 25            | Continued.   |
| SIGTTIN       | 26            | Stopped (tty input).   |
| SIGTTOU       | 27            | Stopped (tty output).  |
| SIGVTALRM     | 28            | Virtual timer expired.   |
| SIGPROF       | 29            | Profiling timer expired.   |
| SIGXCPU       | 30            | CPU time limit exceeded.   |
| SIGXFSZ       | 31            | File size limit exceeded.  |
| SIGWAITING    | 32            | Concurrency signal reserved by threads library.                      |
| SIGLWP        | 33            | Inter-LWP signal reserved by threads library.                        |
| SIGFREEZE     | 34            | Checkpoint freeze.   |
| SIGTHAW       | 35            | Checkpoint thaw.   |
| SIGCANCEL     | 36            | Cancellation signal reserved by the threads library.                 |

**NOTE**

Obtain a list of the signals by typing: `man signal.h`

In addition, you can write a *signal handler*, or *trap*, in a program to respond to a signal being sent. For example, many system programs, such as the name server daemon, respond to the SIGHUP signal by re-reading their configuration files. This signal can then be used to update the process while running, without having to terminate and restart the process. Signal handlers cannot be installed for SIGSTOP (23) or SIGKILL (9). Because the process cannot install a signal handler for signal 9, an otherwise well-behaved process may leave temporary files around or not be able to finish out critical operations that it is in the middle of. Thus, `kill -9` invites corruption of application data files and should only be used as a last resort.

Here's an example of how to trap a signal in a script:

```
# trap '/bin/rm tmp$$;exit 1' 1 2 3 15
```

As the name suggests, `trap` traps system interrupt until some command can be executed. The previous example traps the signals 1, 2, 3, and 15, and executes the `/bin/rm tmp$$` command before exiting the program. The example deletes all `tmp` files even if the program terminates abnormally.

The `kill` command sends a terminate signal (signal 15) to the process, and the process is terminated. Signal 15, which is the default when no options are used with the `kill` command, is a gentle kill that allows a process to perform cleanup work before terminating. Signal 9, on the other hand, is called a sure, unconditional kill because it cannot be caught or ignored by a process. If the process is still around after a `kill -9`, either it is hung up in the Unix kernel, waiting for an event such as disk I/O to complete, or you are not the owner of the process.

The `kill` command is routinely used to send signals to a process. You can kill any process you own, and the superuser can kill all processes in the system except those that have process IDs 0, 1, 2, 3, and 4. The `kill` command is poorly named because not every signal sent by it is used to kill a process. This command gets its name from its most common use—terminating a process with the `kill -15` signal.

**NOTE**

**Forking Problem** A common problem occurs when a process continually starts up new copies of itself—this is referred to as *forking* or *spawning*. Users have a limit on the number of new processes they can fork. This limit is set in the kernel with the MAXUP (maximum number of user processes) value. Sometimes, through user error, a process keeps forking new copies of itself until the user hits the MAXUP limit. As a user reaches this limit, the system appears to be waiting. If you kill some of the user's processes, the system resumes creating new processes on behalf of the user. It can be a no-win situation. The best way to handle these runaway processes is to send the STOP signal to all of the runaway processes to suspend the processes and then send a KILL signal to terminate the processes. Because the processes were first suspended, they can't create new ones as you kill them off.

You can send a signal to a process you own with the `kill` command. Many signals are available, as listed in Table 5.12. To send a signal to a process, first use the `ps` command to find the process ID (PID) number. For example, type `ps -ef` to list all processes and find the PID of the process you want to terminate:

```
# ps -ef

UID      PID  PPID  C   STIME  TTY  TIME  CMD
root      0      0  0   Nov 27  ?    0:01  sched
root      1      0  0   Nov 27  ?    0:01  /etc/init -
root      2      0  0   Nov 27  ?    0:00  pageout
root      3      0  0   Nov 27  ?    12:52 fsflush
root     101     1  0   Nov 27  ?    0:00  /usr/sbin/in.routed -q
root     298     1  0   Nov 27  ?    0:00  /usr/lib/saf/sac -t 300
root     111     1  0   Nov 27  ?    0:02  /usr/sbin/rpcbind
root     164     1  0   Nov 27  ?    0:01  /usr/sbin/syslogd -n -z 12
root     160     1  0   Nov 27  ?    0:01  /usr/lib/autofs/automountd
.
.
.
root    5497    433  1 09:58:02 pts/4  0:00  script psef
```

To kill the process with a PID number of 5497, type this:

```
# kill 5497
```

Another way to kill a process is to use the `pkill` command. `pkill` functions identically to `pgrep`, which was described earlier, except that instead of displaying information about each process, the process is terminated. A signal name or number may be specified as the first command-line option to `pkill`. The value for the signal can be any value described in Table 5.12. For example, to kill the process named `psef` with a `SIGKILL` signal, issue the following command:

```
# pkill -9 psef
```

## NOTE

**Killing a Process** If no signal is specified, `SIGTERM` (15) is sent by default. This is the preferred signal to send when trying to kill a process. Only when a `SIGTERM` fails should you send a `SIGKILL` signal to a process. As stated earlier in this section, a process cannot install a signal handler for signal 9 and an otherwise well-behaved process might not shut down properly.

In addition, the Desktop Process Manager, which was described earlier, can be used to kill processes. In the Process Manager window, highlight the process that you want to terminate, click Process from the toolbar at the top of the window, and then select Kill from the pull-down menu, as shown in Figure 5.9.

| Process               | PU%  | RAM   | Size  | Started  | Parent   | Command  |   |
|-----------------------|------|-------|-------|----------|----------|--|---|
| Trace System Calls... | .5   | 23784 | 31224 | 21:00:41 | 458      | /usr/bin/sdtimage -snapshot                          |   |
| Trace Children...     | .7   | 12344 | 19304 | 20:55:29 | 292      | /usr/operwin/bin/Xsun :D -nobanner -auth /var/dt/A   |   |
| Show Stack...         | .4   | 396   | 1136  | 21:01:52 | 468      | /usr/bin/ps -A -o pid-ID -o 'fname=Name' -o user=Own |   |
| Show Ancestry...      | .2   | 5696  | 8432  | 21:01:01 | 463      | /usr/bin/sdtimage -snapshot                          |   |
| Kill                  | .1   | 1256  | 1888  | 21:01:52 | 400      | ksh -c /usr/bin/ps -A -o pid-ID -o 'fname=Name' -o   |   |
| Exit                  | .1   | 2352  | 3104  | 21:01:01 | 398      | /usr/dt/bin/dtexec -open 0 -ttprocid 3.yd38x 01 38f  |   |
| 306 mib1isa           | root | 0.1   | 1896  | 2304     | 20:55:26 | 285  | mib1isa -r -p 32788                                 |
| 470 ksh               | root | 0.0   | 0     | 1888     | 21:01:52 | 468  | ksh -c /usr/bin/ps -A -o pid-ID -o 'fname=Name' -o  |
| 458 dtexec            | root | 0.0   | 2352  | 3104     | 21:00:41 | 398  | /usr/dt/bin/dtexec -open 0 -ttprocid 3.yd38x 01 38f |

**Figure 5.9** Killing processes.

The equivalent Unix command used by the Process Manager to terminate a process is shown here:

```
# kill -9 <PID>
```

<PID> is the process ID of the selected process.

The `preap` command forces the killing of a defunct process, known as a *zombie*. In previous Solaris releases, zombie processes that could not be killed off remained until the next system reboot. Defunct processes do not normally impact system operation; however, they do consume a small amount of system memory. See the `preap` manual page for further details of this command.

## Scheduling Processes

Processes compete for execution time. *Scheduling*, one of the key elements in a time-sharing system, determines which of the processes executes next. Although hundreds of processes might be present on the system, only one actually uses a given CPU at any given time. Time sharing on a CPU involves suspending a process and then restarting it later. Because the suspension and resumption of active processes occurs many times each second, it appears to the user that the system is performing many tasks simultaneously.

Unix attempts to manage the priorities of processes by giving a higher priority to those that have used the least amount of CPU time. In addition, processes that are waiting on an event, such as a keyboard press, get higher priority than processes that are purely CPU-driven.

On any large system with a number of competing user groups, the task of managing resources falls to the system administrator. This task is both technical and political. As a system administrator, you must understand your company goals to manage this task successfully. When you understand the political implications of who should get priority, you are ready to manage the technical details. As root, you can change the priority of any process on the system by using the `nice` or `priocntl` commands. Before you do this, you must understand how priorities work.

## Scheduling Priorities

All processes have assigned to them an execution priority—an integer value that is dynamically computed and updated on the basis of several different factors. Whenever the CPU is free, the scheduler selects the most favored process to resume executing. The process selected is the one with the lowest-priority number because lower numbers are defined as more favored than higher ones. Multiple processes at the same priority level are placed in the run queue for that priority level. Whenever the CPU is free, the scheduler starts the processes at the head of the lowest-numbered nonempty run queue. When the process at the top of a run queue stops executing, it goes to the end of the line and the next process moves up to the front. After a process begins to run, it continues to execute until it needs to wait for an I/O operation to complete, receives an interrupt signal, or exhausts the maximum execution time slice defined on that system. A typical time slice is 10 milliseconds.

A Unix process has two priority numbers associated with it. One of the priority numbers is its requested execution priority with respect to other processes. This value (its `nice` number) is set by the process's owner and by root; it appears in the `NI` column in a `ps -l` listing. The other priority assigned to a process is the execution priority. This priority is computed and updated dynamically by the operating system, taking into account such factors as the process's `nice` number, how much CPU time it has had recently, and other processes that are running and their priorities. The execution priority value appears in the `PRI` column on a `ps -l` listing.

Although the CPU is the most-watched resource on a system, it is not the only one. Memory use, disk use, I/O activity, and the number of processes all tie together in determining the computer's throughput. For example, suppose you have two groups, A and B. Both groups require large amounts of memory—more than is available when both are running simultaneously. Raising the priority of Group A over Group B might not help if Group B does not fully relinquish the memory it is using. Although the paging system does this over time, the process of swapping a process out to disk can be intensive and can greatly reduce performance. A better alternative might be to completely stop Group B with a signal and then continue it later, when Group A has finished.

## Changing the Priority of a Time-Sharing Process with `nice`

The `nice` command is supported only for backward compatibility with previous Solaris releases. The `prctl` command provides more flexibility in managing processes. The priority of a process is determined by the policies of its scheduling class and by its `nice` number. Each time-sharing process has a global priority that is calculated by adding the user-supplied priority, which can be influenced by the `nice` or `prctl` commands, and the system-calculated priority.

The execution priority number of a process is assigned by the operating system and is determined by several factors, including its schedule class, how much CPU time it has used, and its nice number. Each time-sharing process starts with a default nice number, which it inherits from its parent process. The nice number is shown in the NI column of the ps report.

A user can lower the priority of a process by increasing its user-supplied priority number. Only the superuser can increase the priority of a process by lowering its nice value. This prevents users from increasing the priorities of their own processes, thereby monopolizing a greater share of the CPU.

Two versions of the nice command are available: the standard version, /usr/bin/nice, and a version that is integrated into the C shell as a C shell built-in. /usr/bin/nice numbers range from 0 to +39 and the default value is 20, while the C-shell built-in version of nice has values that range from -20 to +20. The lower the number, the higher the priority and the faster the process runs.

Use the /usr/bin/nice command as described in Table 5.13 when submitting a program or command.

**TABLE 5.13** Setting Priorities with nice

| Command   | Description  |
|---|--|
| <b>Lowering the Priority of a Process Using /usr/bin/nice</b> |  |
| nice <process_name>   | Increases the nice number by 4 units (the default)             |
| nice -4 <process_name>  | Increases the nice number by 4 units                           |
| nice -10 <process_name>                                       | Increases the nice number by 10 units                          |
| <b>Increasing the Priority of a Process</b>                   |  |
| nice -n -10 <process_name>                                    | Raises the priority of the command by lowering the nice number |

#### NOTE

Root may run commands with a priority higher than normal by using a negative increment, such as -10. A negative increment assigned by an unprivileged user is ignored.

As a system administrator, you can use the renice command to change the priority of a process after it has been submitted. The renice command has the following form:

```
renice priority -n <value> -p <pid>
```

Use the `ps -elf` command to find the PID of the process for which you want to change the priority. The process that you want to change in the following example is named `largejob`:

```
# ps -elf|grep largejob
9 S  0  8200  4100  0  84  20  f0274e38  193          Jun 04 ?    0:00 largejob
```

Issue the following command to increase the priority of PID `8200`:

```
renice -n -4 -p 8200
```

Issuing the `ps -elf` command again shows the process with a higher priority:

```
# ps -elf|grep largejob
9 S  0  8200  4100  0  60  16  f0274e38  193          Jun 04 ?    0:00 largejob
```

## Changing the Scheduling Priority of Processes with `prionctl`

The standard priority scheme has been improved since earlier versions of Solaris as part of its support for real-time processes. Real-time processes are designed to work in application areas in which a nearly immediate response to events is required. These processes are given nearly complete access to all system resources when they are running. Solaris uses time-sharing priority numbers ranging from `-20` to `20`. Solaris uses the `prionctl` command, intended as an improvement over the `nice` command, to modify process priorities. To use `prionctl` to change a priority on a process, type this:

```
prionctl -s -p <new-priority> -i pid <process-id>
```

*new-priority* is the new priority for the process, and *process-id* is the PID of the process you want to change.

The following example sets the priority level for process `8200` to `-5`:

```
prionctl -s -p -5 -i pid 8200
```

The following example is used to set the priority (nice value) for every process created by a given parent process:

```
prionctl -s -p -5 -I ppid 8200
```

As a result of this command, all processes forked from process `8200` have a priority of `-5`.

The priority value assigned to a process can be displayed using the `ps` command, which was described earlier in this chapter.

The functionality of the `prionctl` command goes much further than what is described in this section. Consult the online manual pages for more information about the `prionctl` command.

## Fair Share Scheduler (FSS) and the Fixed Scheduler (FX)

The *Fair Share Scheduler (FSS)* in Solaris 10 can be used to control the allocation of resources. The *Fixed Scheduler (FX)* is a fixed priority scheduler that provides an ensured priority for processes. Neither of these are objectives on the CX-310-200 exam and they are not covered in this chapter.

## Using the Solaris Batch-Processing Facility

A way to divide processes on a busy system is to schedule jobs so that they run at different times. A large job, for example, could be scheduled to run at 2:00 a.m., when the system would normally be idle. Solaris supports two methods of batch processing: the `crontab` and `at` commands. The `crontab` command schedules multiple system events at regular intervals, and the `at` command schedules a single system event.

## Configuring crontab

Objective:

**Explain how to schedule the automatic recurring execution of a command.**

`cron` is a Unix utility named after Chronos (“time”), the ancient Greek god of time. It enables you to execute commands automatically according to a schedule you define. The `cron` daemon schedules system events according to commands found in each `crontab` file. A `crontab` file consists of commands, one per line, that will be executed at regular intervals. The beginning of each line contains five date and time fields that tell the `cron` daemon when to execute the command. The sixth field is the full pathname of the program you want to run. These fields, described in Table 5.14, are separated by spaces.

**TABLE 5.14 The crontab File**

| Field | Description  | Values  |
|-------|--------------|---|
| 1     | Minute       | 0 to 59. A * in this field means every minute.                      |
| 2     | Hour         | 0 to 23. A * in this field means every hour.                        |
| 3     | Day of month | 1 to 31. A * in this field means every day of the month.            |
| 4     | Month        | 1 to 12. A * in this field means every month.                       |
| 5     | Day of week  | 0 to 6 (0 = Sunday). A * in this field means every day of the week. |
| 6     | Command      | Enter the command to be run.  |

Follow these guidelines when making entries in the crontab file:

- ▶ Use a space to separate fields.
- ▶ Use a comma to separate multiple values in any of the date or time fields.
- ▶ Use a hyphen to designate a range of values in any of the date or time fields.
- ▶ Use an asterisk as a wildcard to include all possible values in any of the date or time fields. For example, an asterisk (\*) can be used in the first five fields (time fields) to mean all legal values.
- ▶ Use a comment mark (#) at the beginning of a line to indicate a comment or a blank line.
- ▶ Each command within a crontab file must consist of one line, even if it is very long, because crontab does not recognize extra carriage returns.
- ▶ There can be no blank lines in the crontab file. Although this is not documented well, and some crontab files I've seen contain blank lines, the system will generate an email to root with a message that “there is an error in your crontab file.”

The following sample crontab command entry displays a reminder in the user's console window at 5:00 p.m. on the 1st and 15th of every month:

```
0 17 1,15 * * echo Hand in Timesheet > /dev/console
```

crontab files are found in the `/var/spool/cron/crontabs` directory. Several crontab files besides root are provided during the SunOS software installation process; they are also located in this directory. Other crontab files are named after the user accounts for which they are created, such as bill, glenda, miguel, or nicole. They also are located in the `/var/spool/cron/crontabs` directory. For example, a crontab file named root is supplied during software installation. Its contents include these command lines:

```
10 3 * * * /usr/sbin/logadm
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
#10 3 * * * /usr/lib/krb5/kprop_script __slave_kdcs__
```

The first command line instructs the system to run `/usr/sbin/logadmin` at 3:10 a.m. every day of the week. The second command line orders the system to execute `nfsfind` on Sunday at 3:15 a.m. The third command line runs each night at 3:30 a.m. and executes the `gsscred` command. The fourth command is commented out. The cron daemon never exits and is started via the `svc:/system/cron:default` service. The `/etc/cron.d/FIFO` file is used as a lock file to prevent running more than one instance of cron.

## Creating and Editing a crontab File

Creating an entry in the crontab file is as easy as editing a text file using your favorite editor. Use the steps described next to edit this file; otherwise, your changes are not recognized until the next time the cron daemon starts up. cron examines crontab configuration files only during

its own process-initialization phase or when the `crontab` command is run. This reduces the overhead of checking for new or changed files at regularly scheduled intervals.

Step by Step 5.2 tells you how to create or edit a crontab file.

---

## STEP BY STEP

### 5.2 Creating or Editing a crontab File

1. (Optional) To create or edit a crontab file belonging to root or another user, become superuser.
2. Create a new crontab file or edit an existing one by typing the following:

```
# crontab -e
```

#### NOTE

**crontab Default Editor** The `crontab` command chooses the system default editor, which is `ed`, unless you've set the `VISUAL` or `EDITOR` variable to `vi` (or another editor), as follows:

```
# EDITOR=vi;export EDITOR
```

3. Add command lines to the file, following the syntax described in Table 5.14. Because `cron` jobs do not inherit the users environment, such as `PATH`, you should specify the full pathname for commands.
4. Save the changes and exit the file. The crontab file is placed in `/var/spool/cron/crontabs`.
5. Verify the crontab file by typing the following:

```
# crontab -l
```

The system responds by listing the contents of the crontab file.

---

## Controlling Access to crontab

#### TIP

Many of the questions that you encounter regarding `cron` are related to controlling access to `cron` using the files described in this section. Make sure that you have a clear understanding of how you can control a user's access to `cron` by adding and removing entries in these files.

You can control access to crontab by modifying two files in the `/etc/cron.d` directory: `cron.deny` and `cron.allow`. These files permit only specified users to perform crontab tasks such as creating, editing, displaying, and removing their own crontab files. The `cron.deny` and `cron.allow` files consist of a list of usernames, one per line. These access control files work together in the following manner:

- ▶ If `cron.allow` exists, only the users listed in this file can create, edit, display, and remove crontab files.

- ▶ If `cron.allow` doesn't exist, all users may submit crontab files, except for users listed in `cron.deny`.
- ▶ If neither `cron.allow` nor `cron.deny` exists, superuser privileges are required to run crontab.

Superuser privileges are required to edit or create `cron.deny` and `cron.allow`.

During the Solaris software installation process, a default `/etc/cron.d/cron.deny` file is provided. It contains the following entries:

- ▶ `daemon`
- ▶ `bin`
- ▶ `nuucp`
- ▶ `listen`
- ▶ `nobody`
- ▶ `noaccess`

None of the users listed in the `cron.deny` file can access crontab commands. The system administrator can edit this file to add other users who are denied access to the crontab command. No default `cron.allow` file is supplied. This means that, after the Solaris software installation, all users (except the ones listed in the default `cron.deny` file) can access crontab. If you create a `cron.allow` file, only those users can access crontab commands.

## Scheduling a Single System Event (at)

---

Objective:

**Explain how to schedule an automatic one-time execution of a command using the command line.**

The `at` command is used to schedule jobs for execution at a later time. Unlike crontab, which schedules a job to happen at regular intervals, a job submitted with `at` executes once, at the designated time.

To submit an `at` job, type `at` followed by the time that you would like the program to execute. You'll see the `at>` prompt displayed and it's here that you enter the `at` commands. When you are finished entering the `at` command, press `control-d` to exit the `at` prompt and submit the job as shown in the following example:

```
# at 07:45am today
at> who > /tmp/log
at> <Press Control-d>
job 912687240.a at Thu Jun 6 07:14:00
```

When you submit an at job, it is assigned a job identification number, which becomes its filename along with the .a extension. The file is stored in the `/var/spool/cron/atjobs` directory. In much the same way as it schedules crontab jobs, the cron daemon controls the scheduling of at files.

The command syntax for at is shown here:

```
# at [-m -l -r] <time> <date>
```

The at command is described in Table 5.15.

**TABLE 5.15 at Command Syntax**

| Option | Description   |
|--------|---|
| -m     | Sends you mail after the job is completed.  |
| -l     | Reports all jobs for the user.  |
| -r     | Removes a specified job.  |
| <time> | The hour when you want to schedule the job. Add am or pm if you do not specify the hours according to a 24-hour clock. midnight, noon, and now are acceptable keywords. Minutes are optional. |
| <date> | The first three or more letters of a month, a day of the week, or the keywords today or tomorrow.   |

You can set up a file to control access to the at command, permitting only specified users to create, remove, or display queue information about their at jobs. The file that controls access to at is `/etc/cron.d/at.deny`. It consists of a list of usernames, one per line. The users listed in this file cannot access at commands. The default at.deny file, created during the SunOS software installation, contains the following usernames:

- ▶ daemon
- ▶ bin
- ▶ smtp
- ▶ nuucp
- ▶ listen
- ▶ nobody
- ▶ noaccess

With superuser privileges, you can edit this file to add other usernames whose at access you want to restrict.

## Checking Jobs in Queue (`atq` and `at -l`)

To check your jobs that are waiting in the `at` queue, use the `atq` command. This command displays status information about the `at` jobs you created. Use the `atq` command to verify that you have created an `at` job. The `atq` command confirms that `at` jobs have been submitted to the queue, as shown in the following example:

```
# atq
```

The system responds with this:

| Rank | Execution Date | Owner | Job         | Queue | Job Name |
|------|----------------|-------|-------------|-------|----------|
| 1st  | Jun 6, 08:00   | root  | 912690000.a | a     | stdin    |
| 2nd  | Jun 6, 08:05   | root  | 912690300.a | a     | stdin    |

Another way to check an `at` job is to issue the `at -l` command. This command shows the status information on all jobs submitted by a user, as shown in this example:

```
# at -l
```

The system responds with this:

```
user = root    912690000.a    Thu Jun 6 08:00:00
user = root    912690300.a    Thu Jun 6 08:05:00
```

## Removing and Verifying Removal of `at` Jobs

To remove the `at` job from the queue before it is executed, type this:

```
# at -r [job-id]
```

*job-id* is the identification number of the job you want to remove.

Verify that the `at` job has been removed by using the `at -l` (or `atq`) command to display the jobs remaining in the `at` queue. The job whose identification number you specified should not appear. In the following example, you'll remove an `at` job that was scheduled to execute at 8:00 a.m. on June 6. First, check the `at` queue to locate the job identification number:

```
# at -l
```

The system responds with this:

```
user = root    912690000.a    Thu Jun 6 08:00:00
user = root    912690300.a    Thu Jun 6 08:05:00
```

Next, remove the job from the `at` queue:

```
# at -r 912690000.a
```

Finally, verify that this job has been removed from the queue:

```
# at -l
```

The system responds with this:

```
# user = root    912690300.a    Thu Jun  6 08:05:00
```

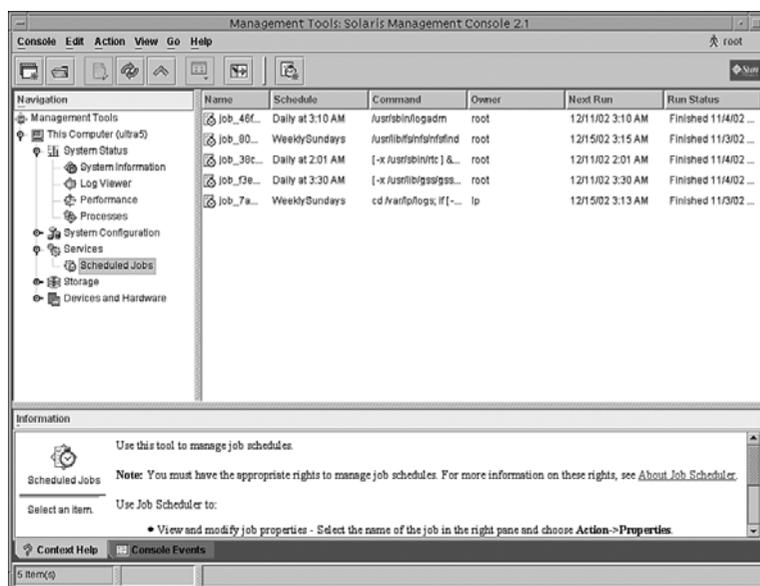
## Job Scheduler

The Solaris Management Console (SMC) includes a graphical tool to create and schedule jobs on your system. You can use the Job Scheduler Tool to

- ▶ View and modify job properties
- ▶ Delete a job
- ▶ Add a scheduled job
- ▶ Enable or disable job logging

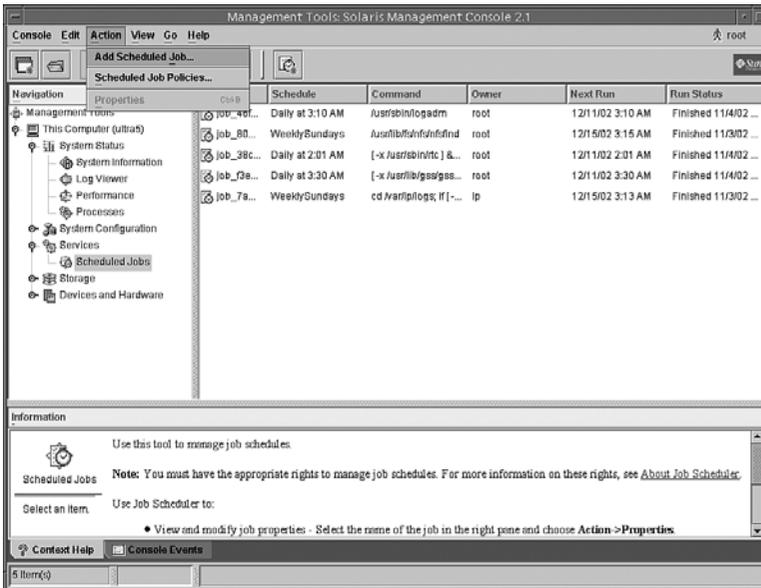
To open the Job Scheduler, follow the steps described in the “SMC Process Tool” section to start up the SMC using the `smc` command.

1. In the Navigation pane of the SMC Welcome window, open the Job Scheduler by clicking on the This Computer icon, then click on the Services icon, and then click on the Scheduled Jobs icon as shown in Figure 5.10.



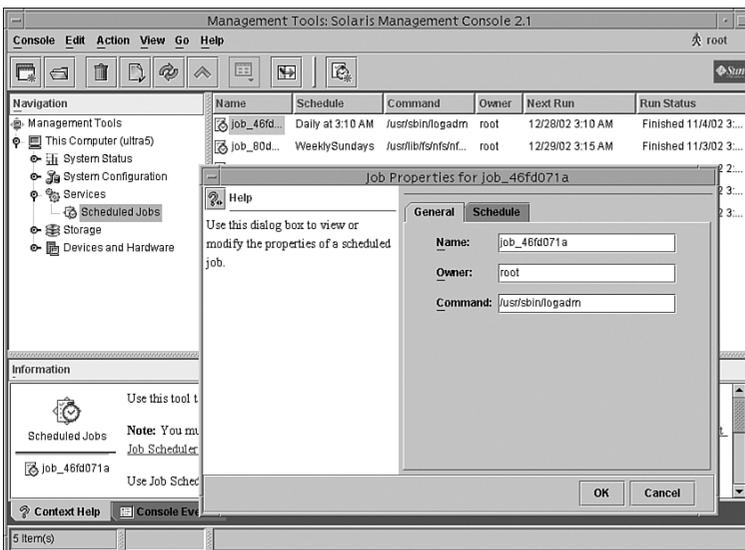
**FIGURE 5.10** Opening the Job Scheduler.

2. You can add jobs to the crontab by selecting Action from the top toolbar as shown in Figure 5.11.



**FIGURE 5.11** Adding a cron job.

3. Modify a cron job by double clicking on the job in the main window pane as shown in Figure 5.12.



**FIGURE 5.12** Modifying a cron job.

# Summary

This chapter described Solaris processes and the various Solaris utilities available to monitor them. Using commands such as `ps`, `prstat`, `pargs`, `sdtprocess`, and the SMC Process Tool, you can view all the attributes associated with a process. In addition, we described foreground and background jobs.

The concept of sending signals to a process was described. A signal is a message sent to a process to interrupt it and cause a response or action. You also learned how to send signals to processes to cause a response such as terminating a process.

Setting process priorities was described. We also described the concept of projects and tasks along with administrative commands used to administer them. The various commands, such as `nice` and `prionctl`, that are used to set and change process priorities were described. In addition, you learned how to use the `crontab` and `at` facilities. You can use these facilities to submit batch jobs and schedule processes to run when the system is less busy, to reduce the demand on resources such as the CPU and disks.

The system administrator needs to be aware of the processes that belong to each application. As users report problems, the system administrator can quickly locate the processes being used and look for irregularities. By keeping a close watch on system messages and processes, you'll become familiar with what is normal and what is abnormal. Don't wait for problems to happen—watch system messages and processes daily. Create shell scripts to watch processes for you and to look for irregularities in the system log files. By taking a proactive approach to system administration, you'll find problems before they affect the users.

In Chapter 6, “Managing the LP Print Service,” we'll explore another topic that you'll need to become acquainted with—the LP Print Service, the facility responsible for printing within the Solaris environment.

## Key Terms

- ▶ `at` command
- ▶ Process Manager GUI
- ▶ Child process
- ▶ `cron`
- ▶ `crontab`
- ▶ `crontab` file
- ▶ `mpstat`
- ▶ `ptree`
- ▶ `nice` command
- ▶ Parent process
- ▶ `pgrep` command
- ▶ `preap` command
- ▶ `prionctl` command
- ▶ Process
- ▶ Project (as it relates to process management)
- ▶ `prstat` command

- ▶ `ps` command
- ▶ Signals
- ▶ SMC Job Scheduler
- ▶ SMC Process Tool
- ▶ `time`
- ▶ Zombie process

## Exercises

---

### 5.1 Displaying Process Information

In this exercise, you'll use the various utilities described in this chapter to display information about active processes.

**Estimated time:** 10 minutes

1. Log in as root into the Java Desktop Environment or CDE.
2. Open a new window and display the active processes using the `ps` command:  

```
# ps -ef
```
3. Open another new window and display the active processes using the `prstat` command:  

```
# prstat
```

Notice how the `ps` command took a snapshot of the active processes, but the `prstat` command continues to update its display.

4. Type `q` to exit `prstat`.
5. Display the `dtlogin` process and all of its child processes. First obtain the PID of the `dtlogin` process with the `pgrep` command:  

```
# pgrep dtlogin
```

Now use the `ptree` command with the PID of the `dtlogin` process to display the ancestry tree:

```
# ptree <PID from dtlogin>
```

6. Now start the Process Manager.  

```
# sdtprocess &
```

Notice how the window updates periodically.

7. In the sample field at the top of the window, change the sample period from 30 to 5 seconds.
8. Sort the processes by ID by clicking on the ID button in the header.

## 5.2 Using the Batch Process

In this exercise, you'll use `crontab` to configure a process to execute everyday at a specified time.

**Estimated time:** 10 minutes

1. Log in as root into a Java Desktop or CDE session.
2. Make sure your default shell editor is set to `vi` (`EDITOR=vi;export EDITOR`) before beginning this exercise.
3. Open a new window and edit the `crontab` entry.

```
# crontab -e
```

4. Enter the following after the last line at the end of the file:

```
0 11 * * * echo Hand in Timesheet > /dev/console
```

5. Save and close the file.

Open a console window and at 11:00 a.m., you'll see the message `Hand in Timesheet` displayed.

## Exam Questions

1. Which of the following commands finds all processes that have `dt` in the process argument string? Choose all that apply.
  - A. `pgrep -l -f dt`
  - B. `ps -ef dt`
  - C. `ps -el dt`
  - D. `ps -ef|grep dt`
2. Which one of the following commands kills a process named `test`?
  - A. `pkill test`
  - B. `kill test`
  - C. `ps -ef||grep kill| kill -9`
  - D. `kill -test`

3. Which commands display active system processes and update at a specified interval? Choose all that apply.
- A. ps
  - B. prstat
  - C. sdtprocess
  - D. ptree
4. In output from the ps command, what does an R stand for in the S field?
- A. The process is on the run queue.
  - B. The process is receiving input.
  - C. It is a regular process.
  - D. The process is sleeping, so it must be restarted.
5. In output from the ps command, which of the following does the UID field display?
- A. The parent process
  - B. The process id
  - C. The process owner
  - D. The priority of the process
6. Which one of the following options to the ps command lists only processes for a particular user?
- A. -P
  - B. -f
  - C. -l
  - D. -u
7. Which one of the following commands lists all processes running on the local system?
- A. ps -e
  - B. ps -a
  - C. ps -f
  - D. ps -t

8. Which one of the following sends a terminate signal (signal 15) to a process with a PID of 2930?
- A. `kill 2930`
  - B. `stop 2930`
  - C. `Ctrl+C`
  - D. `cancel 2930`
9. Which one of the following signals kills a process unconditionally?
- A. 9
  - B. 0
  - C. 15
  - D. 1
10. Which of the following commands is used to change the priority on a process? Choose all that apply.
- A. `renice`
  - B. `priocntl`
  - C. `ps`
  - D. `hup`
11. Which one of the following commands is issued to increase the priority of PID 8200?
- A. `renice -n -4 -p 8200`
  - B. `nice -n -4 -p 8200`
  - C. `nice -i 8200`
  - D. `renice -I -p 8200`
12. Which utilities can be used to show the process ancestry tree? Choose all that apply.
- A. `ps`
  - B. `ptree`
  - C. `sdtprocess`
  - D. `prstat`

13. Which of the following commands schedules a command to run once at a given time?
- A. `crontab`
  - B. `priocntl`
  - C. `at`
  - D. `cron`
14. Which of the following commands show(s) the jobs queued up by the `at` command? Choose all that apply.
- A. `atq`
  - B. `at -l`
  - C. `ps`
  - D. `crontab`
15. Which one of the following `crontab` entries instructs the system to run `logchecker` at 3:10 on Sunday and Thursday nights?
- A. `0 4 * * 10,3 /etc/cron.d/logchecker`
  - B. `10 3 * * 0,4 /etc/cron.d/logchecker`
  - C. `* 10 3 0,4 /etc/cron.d/logchecker`
  - D. `10 3 * * 0-4 /etc/cron.d/logchecker`
16. Which one of the following logs keeps a record of all `cron` activity?
- A. `/var/cron/log`
  - B. `/var/spool/cron/log`
  - C. `/var/adm/cron`
  - D. `/var/adm/messages`
17. A user wants to execute a command later today, after leaving work. Which one of the following commands will allow him to do this?
- A. `runat`
  - B. `at`
  - C. `submit`
  - D. None of the above

18. You've added the user name `bcalkins` to the `/etc/cron.d/cron.allow` file. You've removed the name `bcalkins` from the `/etc/cron.d/cron.deny` file. Which statement is true regarding `crontab`?
- A. `bcalkins` cannot create `crontab` entries.
  - B. `bcalkins` can create `crontab` entries.
  - C. Only `root` can create `crontab` entries.
  - D. No one can create `crontab` entries.

## Answers to Exam Questions

1. **A, D.** Use the `pgrep` and `ps` commands to view processes running on your system. The commands `pgrep -l -f dt` and `ps -ef |grep dt` find all the processes that have `dt` in the process argument string and display them.
2. **A.** The command `kill test` kills a process named `test`.
3. **B, C.** The `prstat` and `sdtprocess` commands display active system processes and can be configured to update at a specified interval.
4. **A.** In output from the `ps` command, the `R` in the `S` field means that the process is on the run queue.
5. **C.** In output from the `ps` command, the `UID` field displays the process owner.
6. **D.** The `-u` option to the `ps` command lists only processes for a particular user.
7. **A.** The `-e` option to the `ps` command lists all processes currently running on the system. The other options only list processes for the local user.
8. **A.** The command `kill 2930` sends a terminate signal (signal 15) to a process with a PID of 2930.
9. **A.** Signal 9 stops a process unconditionally.
10. **A, B.** The commands `renice` and `priocntl` are used to change the priority on a process.
11. **A.** The `renice -n -4 -p 8200` command is issued to increase the priority of a process with a PID of 8200.
12. **B, C.** The utilities `ptree` and `sdtprocess` are used to show the process ancestry tree.
13. **C.** The `at` command schedules a command to run once at a given time.
14. **A, B.** The `atq` and `at -l` commands show the jobs queued up by the `at` command.
15. **B.** The `crontab` entry `10 3 * * 0,4 /etc/cron.d/logchecker` instructs the system to run `logchecker` at 3:10 on Sunday and Thursday nights.

16. **A.** The log file named `/var/cron/log` keeps a record of all cron activity.
17. **B.** Use the `at` command to execute a command or script at a later time.
18. **B.** Users can manage jobs if their name appears in the `/etc/cron.d/cron.allow` file and does not appear in the `/etc/cron.d/cron.deny` file.

## Suggested Reading and Resources

1. Calkins, Bill. *Inside Solaris 9*. New Riders Publishing. November 2002.

# Index

## A

---

**accept command, 569-570**

**access**

ACLs, 447

checking new file permissions, 450

copying to another file, 450-451

deleting entries, 451

modifying entries, 451

setting entries, 448-450

verifying entries, 450

file access

permissions, 443-445

setgid permission, 453-454

setting paths, 452

setuid permission, 452-453

sticky bits, 446

switching users, 445-446

user mask, 446

physical security and, 434

root access

monitoring, 467

restricting, 466-467

system access

controlling, 434

default files, 434-435

passwords, 435-437

restricted shells, 442-443

user account information storage, 438-442

**ACLs (Access Control Lists), 447**

copying to another file, 450-451

deleting entries, 451

file permissions, checking, 450

modifying entries, 451

- practice exams
  - answers, 489-492
  - questions, 480-488
- setting entries, 448-450
- verifying entries, 450

**Add User Wizard, 411****address spaces (processes), 496, 503****answers (practice exams), 685-690**

- ACL (Access Control Lists), 489-492
- ASET (Automated Security Enhancement Tool), 489-492
- backups/restorations, 640-641
- bootstrapping, 390-396
- file systems, 163-170, 673-641
- files, restoring, 637-641
- initialization files, 489-492
- printers, 586-587
- processes, 540
- shell, 489-492
- Solaris installations, 264-265
- system boot/shutdown procedures, 392-396

**application threads, definition of, 503****ASET (Automated Security Enhancement Tool), 452, 471-473, 480-492****at command, processes**

- queue, removing from, 530-531
- scheduling, 528-529

**atq command, 530****attributes (processes), 497****auditing users**

- logins, checking, 456-459
- system usage, monitoring, 454-456

**authentication (public-key), 469-470****autoconfiguration**

- device drivers, 33-35
- devices, 33-34, 142-143
- kernels, 332

**Automated Security Enhancement Tool. *See* ASET (Automated Security Enhancement Tool)****automatic system recovery, 274****B****background processes. *See* jobs****backing-store files, 624**

- disk space requirements, 625
- maxsize option, 626

**backups, 592**

- cpio command, 595
  - copy-in mode, 602
  - copy-out mode, 601
  - wildcards in, 602
- dd command, 595
- /export/home file systems, 606
- file compression, 595
- Flash archives, 595
- fssnap command, 624-628
- GRUB boot archives, 326-328
- incremental backups, UFS Snapshot backups, 628
- pax command, 595
  - copy mode, 603
  - file operand, 605
  - list mode, 603
  - options of, 603-604
  - read mode (-r), 603
  - usage examples, 605
  - write mode (-w), 603
- practice exam questions/answers, 637-641
- remote backups, 611
- root (/) file systems, 606
- SMF databases, 340
- tape drive, 623
- tar command, 595
- ufsdump command, 595, 606
  - dump schedules, 607-608
  - features of, 607
  - options of, 608-610
  - UFS Snapshot backups, 628
  - usage examples, 610-612

- ufsrestore command, 595, 612, 622
  - interactive mode commands, 614-615
  - options of, 613-614
  - restoresymtable files, 618
  - usage examples, 615-618
- /usr file systems, 606

**banner pages, 552****bash shell, 426****batch processing. *See* crontab or at command****BIOS (Basic Input/Output System), x64/x86-based system boots, 319-320****block devices, 46-47****blocks. *See also* cylinder groups**

- boot blocks, 88, 314, 322
- definition of, 49
- free blocks, 90
- logical block sizes, 93
- storage blocks, 90
- superblocks
  - description of, 88
  - locating, 103
- types of, 89

**boot archives, GRUB**

- failsafe archives, 326-328
- primary archives, 327

**boot blocks**

- description of, 88
- GRUB, 322
- installing, 314

**boot cdrom -s, 39****boot command, 309**

- interactive boot process, 311-312
  - bootblock installation, 314
  - kernel loads, 315
  - specifying device aliases, 313
  - standalone program specification, 314-315
- noninteractive boot process, 310

**bootadm command, 326-327****bootstrapping. *See also* systems, shutdown process**

- boot programs phase, 272-274
- boot PROM phase, 272-273
  - FPROM, 275
  - OpenBoot, 274-315
  - SEEPROM chips, 275
- init phase, 272, 335-370
- kernel initialization phase, 272, 331-335
- missing /etc/system files, 311
- OpenBoot startup process, 308, 314
- powering on systems, 273
- practice exam questions/answers, 384-396
- SPARC systems, 307
  - boot command (OpenBoot), 309-315
  - GRUB, 324-328
  - interactive boot process, 311-312
  - noninteractive boot process, 310
  - video display configuration, 329-331
- svc.startd phase, 272
- x64-based systems, 315-318
  - BIOS, 319-320
  - GRUB, 320-322
  - video display configuration, 329-331
- x86-based systems, 315-318
  - BIOS, 319-320
  - GRUB, 320-322
  - video display configuration, 329-331

**Bourne shells**

- initialization files, 426
- job control, 515

**Break key command, 277****BSD (Berkeley Software Distribution) versus SVR4, 552****buffered devices versus raw, 47**

## C

---

### C shell (csh)

- initialization files, 425-426
- job control, 515

### cable connections, 273

### CacheFS (Cache File System), 53

### callback systems, 462

### cancel command, 566, 570

### cancelling print requests, 566, 570

### CD-ROM drives

- identifying connections, 33
- managing. *See* volume management

### CD-ROMs

- booting from, 39
- UFS mounts, 122

### CDE (Common Desktop Environment), shell initialization files and, 427

### child processes, 510

### CLI (Command-Line Interpreter), Solaris installations, 186-187

### command line

- group accounts
  - adding, 423-424
  - deleting, 425
  - modifying, 424
- SMF commands, 343
- software management, 225
- user accounts
  - adding, 417-421
  - deleting, 423
  - managing, 417
  - modifying, 422-423
  - UID assignments, 420
- versus SMC, 417

### compression

- TAR files, 239
- zip/unzip commands, 595

### configuring

- crontab, 525-526
- device drivers, 33-35

- devices, 33-34, 142-143

- disk configuration information, displaying, 56-57, 144

- kernels, 332

- OpenBoot security, 299

- paths, 452

- printers

- network configuration, 551

- Print Manager, configuring via, 554-558

- software, configuring, 554-560

- reconfiguration startups, 34

- security mode (OpenBoot), 293

- video displays

- bootstrapping (SPARC system), 329-331

- x64-based systems, 329, 331

- x86-based systems, 329, 331

- Xorg servers, 329

### contract services (SMF), 363

### controllers (disks), 49

### converting RC scripts to SMF, 359, 361

### copy mode (pax command), 603

### copy-in mode (cpio command), file restoration via, 602

### copy-out mode (cpio command), backups via, 601

### copying file systems via ufsdump command, 606-612

### corrupted data, system shutdown process, 371

### cpio command, 595

- backups, 130

- wildcards in, 602

### crashed systems, interrupting

- SPARC systems, 375-376

- x64-based systems, 377

- x86-based systems, 377

### crontab

- access, controlling, 527-528

- configuring, 525-526

- files, creating/editing, 526-527

### .cshrc file

- C shell, 425

- default, 427

search path, 452  
user mask, 446

**CTFS (Contract File System), 53**  
**customizing SMF service scripts, 357-365**  
**cylinder groups, 89. *See also* inodes**  
**cylinders, 49**

## D

### databases

repository databases, 340  
SMF backups, 340  
terminfo, 557-558

**dd command, 595-599**

**default printers, setting, 572**

**defects lists (disks), 49**

**degraded service state (SMF), 344**

### deleting

ACL entries, 451  
print jobs, 566  
printers, 560-561  
user accounts with userdel command, 423

**Denied permission, 444**

**Desktop Process Manager, killing processes, 520**

**/dev directory, 46, 79**

**/dev/rdisk directory, 47**

**DEVFS (Device File System), 53**

**devfsadm command, 34, 39-40, 143-144**

**device drivers. *See also* devices, 26, 49**

autoconfiguration, 33-35  
physical device names, 27-30, 33  
unsupported, 35  
USB devices, 35-37  
-v option (physical device names), 30

### device trees

device-arguments parameter, 285  
driver-name parameter, 284  
OpenBoot, 280, 283-285  
browsing, 286  
device aliases, 287-288  
unit-address parameter, 284-285

**device-arguments parameter (device trees), 285**

**devices. *See also* device drivers**

adding, practice exercises, 143-144  
aliases  
creating in NVRAM, 295-298  
parameters versus, 289  
removing from NVRAM, 296  
specifying in interactive boot process, 313  
OpenBoot, 287-288  
autoconfiguration, 34  
block, 46-47  
block devices, 47  
configuration information, displaying, 30  
connections, identifying, 33  
logical device names, 42-47  
major/minor device numbers, 40-42  
pathnames. *See* device trees  
physical device names, 29, 32  
practice exercises, 143  
raw, 46-47

**/devices directory, 28**

**/devices file system, device hierarchy creation, 34**

**devlinks command, 39**

**df command, 138**

**diagnostics key position (system control switches), 277**

**directories. *See also* root (/) file system**

access permissions, 444-446  
content listings, 81-83  
file types, identifying, 81  
home directories, 422, 431-432  
in-memory system directories, 81  
links  
definition of, 83  
hard links, 85-87  
removing, 87  
soft (symbolic) links, 83-85  
LP print service directories, 546-547  
size information, displaying, 135-136  
volume management list, 121

disable command

**disable command, 561**

**disabled service state (SMF), 344**

**disabling**

Break key command, 277

L1+A (Stop+A) command, 277

services

network services, SBD, 194-195

svc.startd daemon, 354

Stop+A (L1+A) command, 277

**disk controller, 49**

**disk-based file systems, 52**

**disks**

configuration information, displaying, 56-57, 144

controllers, 49

defects lists, 49

geometry of, 48

labels, 49-50

partition information, recovering, 74-75

physical disks, geometry of, 54

sapce usage, quot command, 139-140

slices, 296

space usage

backing-store files, 625

controlling, 138-139

displaying, 131-134

Solaris installation requirements, 194

spindles, 54

wipe disk standards, 91

**Disks Tool (SMC), 69-70, 73**

**dmesg command, 27-28, 38-39**

**driver not attached messages, 33**

**driver-name parameter (device trees), 284**

**drivers (device). *See also* devices**

autoconfiguration, 33-35

physical device names, 27-30, 33

unsupported, 35

USB devices, 35-37

-v option (physical device names), 30

**drvconfig command, 39**

**.dtprofile file, 427**

**dumping file systems, 607-612**

**DVD drives, UFS mounts, 122**

**dynamic kernels, 332**

## E

**eeeprom command**

boot behavior modification, 324

NVRAM, 293

OpenBoot security, configuring, 299

**EFI (Extensible Firmware Interface), 50-51**

**EFT (extended fundamental types), 76**

**enable command, 561**

**encryption (passwords), 438**

**environment variables, 428-429**

**.equiv file, 462-463**

**/etc directory, subdirectories of, 79-80**

**/etc files, 432**

**/etc/default directory, 434**

**/etc/default/login file, 462, 466**

**/etc/default/passwd file, 435**

**/etc/default/su file, 467**

**/etc/ftpd/ftpshosts file, 465**

**/etc/ftpd/ftpusers file, 464-465**

**/etc/group file, 403, 440-442**

**/etc/hosts.equiv file, 462-463**

**/etc/inetd.conf file, 460-461**

**/etc/inittab files**

fields in, 336

init bootstrapping phase, 335

**/etc/lp/interfaces directory, 558**

**/etc/mnttab file, 116-118**

**/etc/name\_to\_major file, 41-42**

**/etc/nologin files, 374**

**/etc/passwd file, 403, 438**

**/etc/path\_to\_inst file, 39**

**/etc/project file, 432**

**/etc/shadow file, 403, 438-440**

**/etc/shells file, 465**

**/etc/skel directory, 418**

**/etc/ssh/sshd\_config file, 468**

**/etc/telinit command, 374**

**/etc/vfstab entries, creating, 118-120**

**Ethernet, printer connections, 550**

**exams (practice)**

answers, 685-690

ACL (Access Control Lists), 489-492

ASET (Automated Security Enhancement Tool), 489-492

backups/restorations, 640-641

bootstrapping, 390-396

files/file system restoration, 637-641

initialization files, 489-492

printers, 586-587

processes, 540

shell, 489-492

Solaris installations, 264-265

system boot/shutdown procedures,  
392-396

user and security administration, 489-492

questions, 677-684

ACL (Access Control Lists), 480-488

ASET (Automated Security Enhancement Tool), 480-488

backups/restorations, 637-639

bootstrapping, 384-388

file system management, 147-163

initialization files, 480-488

printers, 582-585

processes, 535-539

shell, 480-488

Solaris installations, 258-263

system boot/shutdown procedures,  
384-391

user and security administration, 480-488

**exclude\_all service instance dependency (SMF),  
342**

**exercises (practice)**

devfsadm command, 143-144

devices

adding, 143-144

autoconfiguration, 142-143

displaying information, 143

disks, displaying configuration information,  
144

files/file systems

access permissions, 477-479

creating, 144-145

tuning, 145-146

initialization files, 476-477

printers, 576-581

processes, 534-535

user and security administration, 475-480

volume management, 146

**/export directory, 80**

**/export/home file systems, backups, 606**

## F

**failsafe boot archives, GRUB, 326-328**

**FDFS (File Descriptor File System), 54**

**fdisk partitions, creating, 65**

**file operand, pax command, 605**

**file systems**

administration tasks, 48

CacheFS (Cache File System), 53

CTFS (Contract File System), 53

creating, practice exercises, 144-145

current state, determining, 102

DEVFS (Device File System), 53

/devices file system, device hierarchy creation,  
34

directory size information, displaying,  
135-136

disk space usage, displaying, 131-134

dumping, 607-612

/etc/vfstab entries, creating, 118-120

/export/home file systems, backups, 606

FDFS (File Descriptor File System), 54

home directory, 431-432

/home file system, networked standalone  
systems, 181

## file systems

- HSFS (High Sierra File System), 52
- labeling, 104
- large files, 113
- laying out, 55-56
- LOFS (Loopback File System), 53
- Macintosh file systems, 130
- mounted file systems
  - creating, 108-112
  - /etc/mnttab file, 116-118
  - /etc/vfstab entries, creating, 118-120
  - fsck command, 100
  - large files, 113
  - SMC Mounts Tool, 113
  - UFS logging enabled systems, 114
  - unmounting, 114-115
- network file systems, 52
- newfs command, creating UFS, 91-92
- NFS (Network File System), 52
- OBJFS (Object File System), 54
- PCFS (PC File System), 52
- practice exams
  - answers, 163-170
  - questions, 147-163
- processes, listing, 115-116
- PROCFS (Process File System), 53
- pseudo file systems. *See* virtual file systems
- repairing, 100-102. *See also* fsck command
- restoring
  - practice exam questions/answers, 637-641
  - ufsrestore command, 612-618
- root (/) file system
  - backups, 606
  - directories, 79-81
  - location of, 76
  - networked standalone systems, 181
  - recovering, 618-622
- snapshots, 624
  - creating, 625-626
  - incremental backups, 628
  - mounting, 627
- SWAPFS (Swap File System), 53

- synchronizing, 99-100
- TMPFS (Temporary File System), 53
- tuning, 106-108, 145-146
- types, determining, 110
- unmounting, 114-115
- UFS (Unix File System), 52
  - block types, 87-90
  - CD/DVD drives, mounting, 122
  - creating with newfs command, 91-92
  - default file systems, 76-77
  - features, 76
  - fragment sizes, 93-94
  - fstyp command, 98
  - Inodes parameter, 95
  - locating blocks, 103
  - logical block sizes, 93
  - minimum free space parameter, 94-95
  - mkfs command, 96-99
  - mounting logging enabled systems, 114
  - optimization types, 95
  - parts of, 75-78
  - Snapshot command, 624-628
  - synchronizing, 99-100
  - volcopy command, 105-106
- ufsdump command, copying with, 606-612
- /usr file systems, 77
  - backups, 606
  - recovering, 618-622
- virtual file systems, 48, 53-54

**files**

- access
  - paths, setting, 452
  - permissions, 443-445, 477-479
  - setgid permission, 453-454
  - setuid permission, 452-453
  - sticky bits, 446
  - switching users, 445-446
  - user mask, 446
- backing-store files, 624
- disk space requirements, 625
- maxsize option, 626

restoring  
     practice exam questions/answers, 637-641  
     ufsrestore command, 612-618  
 types of, identifying, 81

### Flash archives, 595

floppy disk drive management. *See* volume management

FMRI (Fault Management Resource Identifiers), SMF service instances, 340

fmthard command, 74-75

forced off key position (system control switches), 277

foreground processes, 515

forks (processes), 496, 519

format command, 28

Format Main Menu, item descriptions, 61

format utility

    disks, labeling, 63  
     free hog slices, 66  
     slices  
         creating, 58-66  
         modifying, 66-68

fragment sizes, 93-94

free blocks, 90

free hog slices, 66

fsck command, 100

    logged file systems, 114  
     output of, 136-137  
     running manually, 102-103  
     -y option, 103

fsname command, 106

fssnap command, 624-628

fstyp command, 98

FTP (File Transfer Protocol), restricting, 464-465

ftp services

    disabling, 349  
     listing all associated properties of, 348

ftphosts file, 465

ftpusers file, 464-465

full device pathnames, 27-28. *See also* device trees; physical device names

fuser command, 115-116

## G

getfacl command, ACL entries

    copying to another file, 450-451  
     displaying, 449-450  
     verifying, 450

GIDs (group IDs)

    effective GIDs, 445-446  
     groupadd command options, 424  
     groupmod command options, 424  
     setgid permission, 453-454  
     values, 441

group accounts

    adding with  
         groupadd command, 423-424  
         SMC (Solaris Management Console), 414-417  
         smgroup command, 416-417  
     groupdel command, deleting with, 425  
     groupmod command, modifying with, 424  
     passwords, 446  
     projects, 432  
     switching, 446

groupadd command, 423-424

groupdel command, 425

groupmod command, 424

grpck command, 442

GRUB (Grand Unified Bootloader)

    boot archives  
         failsafe archives, 326-328  
         primary archives, 327  
     boot behavior modification, 322-324  
     boot blocks, 322  
     device names, 325  
     GRUB menu, 325-326  
     menu modification, 324  
     x64-based system boots, 320-321  
     x86-based system boots, 320-321

**GUI (Graphical User Interface)**

- software management, 225
- Solaris installations, 186

---

## H - I

**hard disks, 49****hard links, 85-87****home directories, 80, 432**

- modifying, 422
- mounting, 431
- naming, 431

**/home file system, networked standalone systems, 181****hosts (trusted), 462****hot-pluggable devices, 36, 122. *See also* USB devices****HSFS (High Sierra File System), 52****hung systems, interrupting**

- SPARC systems, 375-376
- x64-based systems, 377
- x86-based systems, 377

**-i option, 83****IDE disks, 45****in-memory system directories list, 81****in.ftpd command, 348****incremental backups, UFS Snapshot backups, 628****inetadm command, 343****inetd daemon, 460-461****inetd.conf file, 460-461****init command, changing run levels, 354****initialization files**

- bash shell, 426
- Bourne shell, 426
- C shell, 425-426
- CDE requirements, 427
- default files, 427
- environment variables, 428-429

## Korn shell, 426

## local initialization files, 428

## modifying, 428-431

## practice exams

- answers, 489-492

- questions, 480-488

## practice exercises, 476-477

## site initialization files, 428

## tsh shell, 426

**inodes. *See also* blocks**

## numbers

- allowed, 95

- displaying, 83

## overview, 89

**installboot command, 314, 619****installing**

## boot blocks, 314

## software groups, 212

## Solaris

- CLI, 186-187

- GUI, 186

- practice exam questions/answers, 258-265

- reinstallations, 196

**instance names (kernel), 37-40****Intel platforms, OpenBoot, 275****interactive boot process**

## bootblocks, installing, 314

## device aliases, specifying, 313

## kernel loads, 315

## messages, viewing, 312

## SPARC systems, 311-312

## standalone program specification, 314-315

---

## J - K

**jar command, 595, 632****Job Scheduler Tool, 531-532****jobs, 515-517. *See also* processes**

## at command

- removing from queue, 530-531

- scheduling with, 528-529

## atq command, 530

**jobs tables, 515-516**

**jsh (job shell), job control, 515**

**kdmconfig command, video display configuration in x64/x86-based systems, 329**

**Kerberos security, requirements for, 203**

## **kernel**

autoconfiguring, 33-35, 332

drivers. *See* device drivers

dynamic kernels, 332

instance names, 37-40

major/minor device numbers, 40-42

module subdirectories list, 35

x64-based system boots, 318

x86-based system boots, 318

**kernel command, boot behavior modification, 323**

**/kernel directory, 80**

**keyboard shortcuts, stopping systems, 278**

**kill command, 519-520**

## **killing processes**

Desktop Process Manager, 520

PID numbers, 520

ps command, 520

preap command, 521

## **Korn shell (ksh)**

initialization files, 426

job control, 515

**.kshrc file, 427**

# **L**

**L1+A (Stop+A) command, 277-278**

**labeling disks, 50**

**labeling file systems, 104**

**labelit command, 104**

**labels (disks), displaying disk configuration information, 56-57**

**large file-aware, 113**

**large file-safe, 113**

**large files, mounting file systems with, 113**

**last command, 459**

**LDAP (Lightweight Data Access Protocol), 432**

**legacy scripts, 368**

**legacy services, 344, 365-366**

**legacy\_run service state (SMF), 344**

## **links**

definition of, 83

hard links, 85-87

removing, 87

soft (symbolic) links, 83-85

volume management list, 121

**list command, viewing services in SMF service configuration repository, 348**

**list mode (pax command), 603**

**listprop command, listing all associated properties of ftp service, 348**

**local initialization files, 428**

**locked key position (system control switches), 277**

**LOFS (Loopback File System), 53**

**logging enabled systems (UFS), mounting, 114**

**logical block sizes, 93**

**logical device names, 42-47**

## **.login file**

C shell, 425

default, 427

search path, 452

**loginlog file, 455-456**

## **logins, initialization files**

bash shell, 426

Bourne shell, 426

C shell, 425-426

CDE requirements, 427

customizing, 428-431

default, 427

environment variables, 428-429

Korn shell, 426

local initialization files, 428

site initialization files, 428

tcsh shell, 426

logins command

**logins command, 454-455**

**.logout file, 426**

**logs (SMF message), 356**

**lp command**

print request priority levels, 567, 570-571

print styles, 553-554

**LP print daemons, 548-549**

**LP print service, 544-545. *See also* printers**

adding printers, 549

directories, 546-547

spooler, 547-548

**lpadmin command, printers**

classes, defining, 562-563

deleting, 560-561

limiting user access to, 567-568

setting up, 558-560

**lpd daemon, 549**

**LPDEST environmental variable, 554, 572**

**lpmove command, 562**

**lpsched daemon, 548-549**

print scheduler, restarting, 571

SVR4 print service, 552-553

**lpsched daemon, troubleshooting, 574**

**lpshut command, 571**

**lpstat command, 564-565, 568**

**ls -l command, 450**

**ls -la command, 443**

**LVMs (logical volumes), 75**

---

## M

**Macintosh file systems, 130**

**maintenance service state (SMF), 344**

**manifests (SMF), 339, 361-365**

**maxsize option, backing-store files, 626**

**MAXUP, 519**

**mdb command, 333**

**mdm command, 335**

**messages logs (SMF), 356**

**metadevice mirrors, 328**

**milestones (SMF), 339, 355**

booting systems via, 353-354

changing, 350-352

multi-user milestones, 345-346

system run states versus, 367

**minfree, setting, 107-108**

**minor device numbers, 41**

**mkfs command, 96-99**

**/mnt directory, 80**

**modems, security, 462**

**modify option (format utility), 66-68**

**mount command, 108-112**

**mount points, creating, 216**

**mounted file systems**

creating, 108-112

/etc/mnttab file, 116-118

/etc/vfstab entries, creating, 118-120

fsck command, 100

large files, 113

UFS logging enabled systems, 114

unmounting, 114-115

**mounting**

home directories, 431

partitions, 67

snapshots, 627

**Mounts Tool (SMC), 113**

**mt command, 623**

**multi-user milestones (SMF), 345-346**

**multitasking, 503**

**multithreading, 503**

---

## N

**name services, 432**

**net services command, 195**

**net services limited command, 356**

**net services open command, 356**

**networks**

file systems, 52

printers, 550-551

security, 459  
  /etc/default/login file, 462  
  /etc/hosts.equiv file, 462-463  
  FTP, restricting, 464-465  
  modems, 462  
  .rhosts file, 463-464  
  securing network services, 460-461  
  ssh (Secure Shell), 468-471  
  trusted hosts, 462

services  
  disabling, SBD, 194-195  
  /etc/default/login file, 462  
  securing, 460-461

**news command, creating UFS, 91-92**

**newgrp command, 446**

**newtask command, 432**

**NFS (Network File System), 52**

**NFS Server Service, sharing resources with SMF, 358-359**

**nice command, changing priorities (processes), 522-524**

**NIS (Network Information Service), 432**

**NIS+ (Network Information Service+), 432**

**nolargefiles mount option, 113**

**noninteractive boot process, SPARC systems, 310**

**normal key position (system control switches), 277**

**nvalias command, 289, 295**

**nvedit line editor, 296-298**

**NVRAM (Nonvolatile Random Access Memory)**  
  device aliases, 289  
  creating, 295-298  
  parameters versus, 289  
  removing, 296

**nvalias command, 289, 295**

**nvunalias command, 289, 295**

parameters versus device aliases, 289

**show-disks command, 295**

**nvunalias command, 289, 295**

## O - P

---

**OBJFS (Object File System), 54**

**offline service state (SMF), 344**

**online service state (SMF), 344**

**OpenBoot**  
  commands, 175  
  device aliases, 287-288  
  device trees, 280, 283-285  
  browsing, 286  
  device aliases, 287-288  
  monitor programs, 273  
  show-devs command, 285

**/opt directory, 77, 80**

**optional\_all service instance dependency (SMF), 342**

**parallel printer connections, 550**

**parallel processing, 503**

**pargs command, 512-514**

**partitions**  
  fdisk partitions, creating, 65  
  mounted partitions, modifying, 67  
  slice 2 as, 55  
  slices versus, 54  
  tables, 51

**pass mode (cpio command), 602**

**passwd command, 409, 421-422**

**password command, NVRAM, 291**

**passwords**  
  blank passwords, checking for, 455  
  changing, 409, 414  
  encryption, 438  
  /etc/default/passwd file, 435  
  group passwords, 446  
  guidelines, 436-437  
  options, setting, 414  
  password aging, 434  
  Solaris installations, Interactive Installation process (install-solaris), 207

**Patch Tool, 237****PATH variable, 430, 452****paths, configuring, 452****path\_to\_inst file, removing, 39****pax command, 595**

- copy mode, 603
- file operand, 605
- list mode, 603
- options of, 603-604
- read mode (-r), 603
- usage examples, 605
- write mode (-w), 603

**PCFS (PC File System), 52****percent sign (%), importance of, 516****permissions**

- ACLs, 447
  - checking new file permissions, 450
  - copying to another file, 450-451
  - deleting entries, 451
  - modifying entries, 451
  - setting entries, 448-450
  - verifying entries, 450
- displaying, 443
- exercises, 477-479
- modifying, 444-445
- setgid permission, 453-454
- setuid permission, 452-453
- sticky bits, 446
- user mask, 446

**pgrep command, 369, 500-501****physical device names, 27-30, 33****physical disks, geometry of, 54****PID (Process Identification) numbers, 520****pkgid command, 227****pkgtrans command, 225****pkill command, 520****/platform directory, 80****power failures, protecting against, 371****power-saving feature, Solaris installations, 218****practice exams**

- answers, 685-690
  - ACL (Access Control Lists), 489-492
  - ASET (Automated Security Enhancement Tool), 489-492
  - backups/restorations, 640-641
  - bootstrapping, 390-396
  - file systems, 163-170, 637-641
  - file restoration, 637-641
  - initialization files, 489-492
  - printers, 586-587
  - processes, 540
  - shell, 489-492
  - Solaris installations, 264-265
  - system boot/shutdown procedures, 392-396
  - user and security administration, 489-492
- questions, 677-684
  - ACL (Access Control Lists), 480-488
  - ASET (Automated Security Enhancement Tool), 480-488
  - backups/restorations, 637-639
  - bootstrapping, 384-388
  - file system management, 147-163
  - initialization files, 480-488
  - printers, 582-585
  - processes, 535-539
  - shell, 480-488
  - Solaris installations, 258-263
  - system boot/shutdown procedures, 384-391
  - user and security administration, 480-488

**practice exercises**

- devfsadm command, 143-144
- devices
  - adding, 143-144
  - autoconfiguration, 142-143
  - displaying information, 143
- disks, displaying configuration information, 144

- files/file systems
  - access permissions, 477-479
  - creating, 144-145
  - tuning, 145-146
- initialization files, 476-477
- printers, 576-581
- processes, 534-535
- user and security administration, 475-480
- volume management, 146

**preap, killing processes, 521**

**primary boot archives, GRUB, 327**

**print clients versus print servers, 553-554**

**Print Manager**

- configuring printers, 554-558
- deleting printers, 560
- lp commands, 552
- modifying printers, 572-574

**print scheduler. *See* lpsched daemon**

**print servers versus print clients, 553-554**

**printenv command, NVRAM, 291-292**

**PRINTER environmental variable, 554, 572**

**printers, 545. *See also* LP print service**

- adding, 549
- administering, 560
- BSD versus SVR4, 552
- classes, 562-563
- default, setting, 572
- deleting, 560-561
- Ethernet connections, 550
- exercises, 576-581
- modifying, 572-574
- networks, configuring on, 551
- parallel connections, 550
- practice exams
  - answers, 586-587
  - questions, 582-585
- print jobs
  - halting, 561-562
  - moving, 562

- queues
  - accepting print requests, 569-570
  - cancelling print requests, 566, 570
  - deleting print jobs, 566
  - limiting user access, 567-569
  - overview, 565
  - rejecting print requests, 569
  - specifying priority levels, 567, 570-571
- requests, halting, 561-562
- serial connections, 550
- servers versus clients, 553-554
- software, configuring
  - lpadmin command, 558-560
  - overview, 554
  - Print Manager, 554-558
  - status, verifying, 564-565
  - SVR4 print service, 552-553
  - USB connections, 551

**prionctl command, 524**

**priorities (processes)**

- changing, 522-524
- scheduling, 522

**Process Manager, 508-511**

**Process Tool (SMC), 511-512**

**processes**

- address spaces, 496, 503
- application threads, 503
- at command
  - removing from queue, 530-531
  - scheduling with, 528-529
- atq command, 530
- attributes list, 497
- background processes. *See* jobs
- child processes, 510
- crontab
  - access, controlling, 527-528
  - configuring, 525-526
  - files, creating/editing, 526-527
- foreground processes, 515
- forking, troubleshooting, 519

Job Scheduler Tool, 531-532

killing with

- Desktop Process Manager, 520
- PID numbers, 520
- pkill command, 520
- preap, 521

listing, 115-116

monitoring, command for, 497

multitasking, 503

multithreading, 503

parallel processing, 503

pargs command, 512-514

percent sign (%), importance of, 516

pgrep command, 500-501

practice exams

- answers, 540
- questions, 535-539

practice exercises, 534-535

priorities

- changing, 522-524
- scheduling, 522

process states list, 498

prstat command

- column headings, 502
- options/arguments, 504-505
- syntax of, 504
- viewing with, 505-506

ps command, 498

psrinfo command, 505

scheduling, 521

shared memory, 503

signals

- kill command, 519-520
- overview, 517-518
- sending to processes, 520
- traps, 519

sort command, utilizing, 500

svcs command, listing processes with, 514

viewing, 498-500

### **PROCFS (Process File System), 53**

#### **.profile file**

- Bourne shell, 426
- CDE and, 427
- default, 427
- Korn shell, 426
- search path, 452

#### **project file, 432**

#### **projects command, 432**

#### **prompt variable, 429**

#### **prstat command**

- column headings, 502
- options/arguments, 504-505
- syntax of, 504
- viewing processes with, 505-506

#### **prtconf command, 28, 33, 294**

#### **prvtoc command, 54-57**

#### **ps command**

- manual pages, 498-499
- options, 499-500
- PID numbers, finding, 520

#### **PS1 variable, 429**

#### **pseudo file systems. *See* virtual file systems**

#### **psrinfo command, 505**

#### **ptree command, 510**

#### **public-key authentication, 469-470**

#### **public/private key pairs, 469-470**

#### **pwck command, 439**

#### **pwconv command, 440**

## **Q - R**

---

#### **questions (practice exams), 677-684**

- ACL (Access Control Lists), 480-488
- ASET (Automated Security Enhancement Tool), 480-488
- backups/restorations, 637-639
- bootstrapping, 384-388
- file system management, 147-163
- initialization files, 480-488
- printers, 582-585
- processes, 535-539

shell, 480-488  
 Solaris installations, 258-263  
 system boot/shutdown procedures, 384-391  
 user and security administration, 480-488

## queues

jobs queues  
   checking with atq command, 530  
   removing jobs from, 530-531  
 print queues, deleting print jobs from, 566

## quot command, 139-140

## quotas, utilizing, 138-139

## -r (read mode), pax command, 603

## raw devices, 46-47

## RC scripts, converting to SMF, 359-361

## re-preinstall command, 187

## read mode (-r), pax command, 603

## reconfiguration startups, 34

## /reconfigure file, 34

## recovery

root (/) file systems, 618-622  
 system recovery  
   automatic recovery, 274  
   interrupting SPARC systems, 375-376  
   interrupting x64-based systems, 377  
   interrupting x86-based systems, 377  
 /usr file systems, 618-622

## reinstalling Solaris, 196

## reject command, 561

## remote backups, /.rhosts files, 611

## removing device aliases, NVRAM aliases, 296

## repository databases, SMF, 340

## request command, 569

## require\_all service instance dependency (SMF), 342

## require\_any service instance dependency (SMF), 342

## restoresymtable files, file system recovery, 619-620

## restoring files/file systems

practice exam questions/answers, 637-641  
 ufsrestore command, 612-618

## rexid daemon, 461

## .rhosts file, 463-464, 611

## rksh shell, 442-443

## rmformat command, 122-123

## rmmount command, 127, 130

## root (/) file system. *See also* directories

backups, 606  
 directories, 79-81  
 location of, 76  
 networked standalone systems, 181  
 recovering, 618-622

## root access. *See also* superuser access

monitoring, 467  
 restricting, 466-467, 479-480

## RSC (Remote System Control) sessions, preventing restarts, 277

## rsh shell, 442-443

## run control scripts, 368

## run levels, changing, 354

## rusers command, 457

# S

## SBD (Secure by Default) feature, 194-195

## /sbin directory, 80

## /sbin/init command, 336

## /sbin/mountall, 109

## scheduling. *See also* Job Scheduler Tool

file systems dumps, 607-608  
 process priorities, 522-524

## scp command, 470-471

## scripts, 368

## SCSI disks, 45

## sdtprocess. *See* Process Manager

## secondary groups, 440

## sectors, 49

## Secure by Default Network Profile, 356

**Secure Shell (ssh), 468-471****security**

## ACLs, 447

- checking new file permissions, 450

- copying to another file, 450-451

- deleting entries, 451

- modifying entries, 451

- setting entries, 448-450

- verifying entries, 450

- ASET (Automated Security Enhancement Tool), 452, 471-473

## auditing users

- checking logins, 456-459

- monitoring system usage, 454-456

- common-sense techniques, 473

- exam question answers, 489-492

- exam questions, 480-488

## file access

- permissions, 443-445

- setgid permission, 453-454

- setting paths, 452

- setuid permission, 452-453

- sticky bits, 446

- switching users, 445-446

- user mask, 446

- Kerberos security, requirements for, 203

- network security, 459

- `/etc/default/login` file, 462

- `/etc/hosts.equiv` file, 462-463

- FTP restrictions, 464-465

- modems, 462

- `.rhosts` file, 463-464

- securing network services, 460-461

- ssh (Secure Shell), 468-471

- trusted hosts, 462

- physical security, 434

## root access

- monitoring, 467

- restricting, 466-467

## system access

- controlling, 434

- default files, 434-435

- passwords, 435-437

- restricted shells, 442-443

- user account information storage, 438-442

- system security, overview of, 433

**serial printer connections, 550****servers**

- print servers, 553-554

- Xorg servers, 329

**service configuration repository (SMF), 338**

- modifying, 347-349

- viewing all services in, 348

**service instances (SMF), 338-339**

- dependencies, 341-342

- FMRI, 340

**set-default command, NVRAM, 291****setenv command, NVRAM, 291****setfacl command, ACL entries**

- deleting, 451

- modifying, 451

- setting, 448-449

**setgid permission, 453-454****setuid permission, 452-453****shell**

- initialization files

- bash shell, 426

- Bourne shell, 426

- C shell, 425-426

- CDE requirements, 427

- default, 427

- environment variables, 428-429

- Korn shell, 426

- local initialization files, 428

- modifying, 428-431

- site initialization files, 428

- tcsh shell, 426

## practice exams

- answers, 489-492

- questions, 480-488

**shells file, 465**

**shortcuts (keyboard), stopping systems, 278**

**show-devs command, 285**

**show-disks command, 295**

**showrev -p command, 242**

**shutdown process (system).**

**See also bootstrapping**

commands list, 371

corrupted data, 371

crashed systems, interrupting

    SPARC systems, 375-376

    x64-based systems, 377

    x86-based systems, 377

power failures, 371

powering down hardware, 377

practice exam questions/answers, 384-396

/sbin/init command, 373, 375

/usr/sbin/halt command, 374

/usr/sbin/poweroff command, 374

/usr/sbin/reboot command, 374

/usr/sbin/shutdown command, 372-375

**SIGKILL, 520**

**signal handlers. See traps**

**signals, 517-518**

Desktop Process Manager, killing processes with, 520

kill command, 519-520

kill command, killing processes with, 520

preap command, killing processes with, 521

processes, sending to, 520

traps, 519

**signed patches, 239**

**SIGTERM, 520**

**site initialization files, 428**

**slice 2, as a partition, 55**

**slice 6. See free hog slices**

**slices**

boot blocks, 88

Disks Tool (SMC), 69-70, 73

format utility, creating with, 58-66

free blocks, 90

free hog slices, 66

modifying, 66-68

numbers, locations of, 77-78

overview, 54, 56

partition tables, 51

partitions versus, 54

recovering, 74-75

sizing, 217

storage blocks, 90

superblocks

    description of, 88

    locating, 103

x86 versus SPARC systems, 44-45

**SLP (Service Location Protocol), 545**

**SMC (Solaris Management Console), 402**

command line versus, 417

Disks Tool, 69-70, 73

exercises, 475-476

group accounts, adding, 414-417

Job Scheduler Tool, 531-532

Mounts Tool, 113

Process Tool, 511-512

user accounts

    adding, 403-404, 407-411

    deleting, 412

    modifying, 413-414

user fields, 408-409

**smgroup command, 416-417**

**smuser command, 420-425**

**snapshots (backups), 624**

creating, 625-626

incremental backups, creating via ufsdump command, 628

mounting, 627

**soft (symbolic) links, 83-85**

**software**

groups, installing, 212

managing via

    command-line, 225

    GUI, 225

printer software, configuring, 558

**Solaris**

## installing

CLI installations, 186-187

GUI installations, 186

practice exam questions/answers, 258-265

## reinstalling, 196

## upgrading, file space, 215

**Solaris LP print service. *See* LP print service****sort command, 136, 500****space optimization, 95****SPARC system**bootstrapping, 307. *See also* systems, shutdown process

boot command (OpenBoot), 309-315

boot programs phase, 272-274

boot PROM phase, 272-273

GRUB, boot archives, 326-328

GRUB, boot behavior modification, 324

GRUB, device names, 325

init phase, 272, 335-370

interactive boot process, 311-312

kernel initialization phase, 272, 331-335

missing /etc/system files, 311

noninteractive boot process, 310

OpenBoot startup process, 308, 314

powering on systems, 273

svc.startd phase, 272

video display configuration, 329-331

crashed systems, interrupting, 375-376

shutdown process. *See also* bootstrapping

commands list, 371

corrupted data, 371

interrupting crashed systems, 375-376

power failures, 371

powering down hardware, 377

/sbin/init command, 373-375

/usr/sbin/halt command, 374

/usr/sbin/poweroff command, 374

/usr/sbin/reboot command, 374

/usr/sbin/shutdown command, 372-375

slicing versus x86, 44-45, 54

spawn (processes), 496

spindles, 54

spoolers, LP print service, 547-548

spooling space, 547

ssh (Secure Shell), 468-471, 480-492

sshd daemon, 468

sshd\_config file, 468

start scripts, 368

startup scripts, environment variables, 428-429

state flags, 75

sticky bits, 446

stop scripts, 368

Stop+A (L1+A) command, 277-278

stopping systems, 278

storage blocks, 90

su command, 445

sulog file, 467

Sun Connection Services, 237

superblocks, 88, 103

superuser access. *See also* root access

monitoring, 467

restricting, 466-467

svc.config.d daemon, 347

svc.startd daemon, 337

disabling services, 354

milestones, 339

service configuration repository, 338

svcadm command, 343, 350-354

svccfg command, 342-343, 347

svccprop command, 343, 349

svcs command, 342-347, 351

displaying services, 353

listing processes with, 514

SVM (Solaris Volume Manager), 75

SVR4 print service, 552-553

swap space, 76

SWAPFS (Swap File System), 53

sync command, 89

synchronizing file systems, 99-100

**sysdef command, 28, 32, 333**

**sysidcfg files, 193**

## systems

### access

controlling, 434

default files, 434-435

passwords, 435-437

restricted shells, 442-443

user account information storage, 438-442

monitoring usage, 454-456

recovery, 274

security, 433

shutdown process. *See also* bootstrapping

commands list, 371

corrupted data, 371

interrupting crashed systems, 375-377

power failures, 371

powering down hardware, 377

practice exam questions/answers, 384-396

/sbin/init command, 373-375

/usr/sbin/halt command, 374

/usr/sbin/poweroff command, 374

/usr/sbin/reboot command, 374

/usr/sbin/shutdown command, 372-375

questions

file system management, 147-163

printers, 582-585

processes, 535-539

user and security administration, 480-488

Solaris installations, 258-265

system boot/shutdown procedures, 384-396

**third-party drivers, Solaris installations, 221**

**time optimization, 95**

**time-sharing processes, changing priorities, 522-524**

**TMPFS (Temporary File System), 53**

**tracks, 49**

**transient services (SMF), 363**

**traps, 519**

**Trojan horses, 452**

**troubleshooting**

power failures, 371

SMF boot problems, 355-356

volume management, 129-131

**trusted hosts, 462**

**tunefs command, 106-108, 145-146**

---

## T

**tape drives, mt command, 623**

**tar command, 595**

**TAR file format, 130, 239**

**tcsh shell, 426**

**terminfo database, 557-558**

## tests (practice)

### answers

file system management, 163-170

printers, 586-587

processes, 540

user and security administration, 489-492

backups/restorations, 637-641

---

## U

**UDF (Universal Disk Format), 52**

**UFS (Unix File System), 52. *See also* file systems**

### blocks

locating, 103

logical block sizes, 93

types of, 87-90

CD/DVD drives, mounting, 122

default file systems, 76-77

features, 76

fragment sizes, 93-94

fstyp command, 98

Inodes parameter, 95

logging enabled systems, mounting, 114

minimum free space parameter, 94-95

mkfs command, 96-99

newfs command, creating with, 91-92

- optimization types, 95
- parts of, 75-78
- Snapshot command, 624-628
- synchronizing, 99-100
- volcopy command, 105-106

**ufsdump command, 595, 606**

- dump schedules, 607-608
- features of, 607
- options of, 608-610
- UFS Snapshot backups, 628
- usage examples, 610-612

**ufsrestore command, 595, 612, 622**

- interactive mode commands, 614-615
- options of, 613-614
- restoresymtable files, 618
- usage examples, 615-618

**UIDs (user IDs), 420**

- effective UIDs, 445-446
- setuid permission, 452-453
- values, 441

**umask command, 429, 446****umount command, 114-115****uninitialized service state (SMF), 344****unit-address parameter (device trees), 284-285****Unix processes. *See* processes****upgrading Solaris file space, 215****USB devices, 35**

- disconnecting, 36-37
- hot-pluggable devices, 122
- managing drives. *See* volume management

**USB printer connections, 551****user accounts**

- adding
  - default settings, 411
  - from the command line, 417, 419
  - user fields for, 408-409
  - user information, 403-404
  - with SMC (Solaris Management Console), 403-404, 407-411
  - with smuser command, 420-425
  - with useradd command, 418-420

- administering, 402. *See also* SMC

- checking logins, 456-459

- deleting with

- SMC (Solaris Management Console), 412
  - userdel command, 423

- directory access permissions, 444

- exams

- answers, 489-492

- questions, 480-488

- expiration dates, setting, 423

- file access permissions, 443-445

- home directories, 431-432

- information storage, 438-442

- initialization files

- bash shell, 426

- Bourne shell, 426

- C shell, 425-426

- CDE requirements, 427

- default, 427

- environment variables, 428-429

- Korn shell, 426

- local initialization files, 428

- modifying, 428-431

- site initialization files, 428

- tcsh shell, 426

- managing from the command line, 417

- modifying

- with SMC (Solaris Management Console), 413-414

- with usermod command, 422-423

- monitoring

- exercises, 477

- system usage, 454-456

- name services, 432

- passwd command, 421-422

- projects, 432

- restricted shells, 442-443

- SMC versus the command line, 417

- switching users, 445

- UIDs, assigning, 420

**User Accounts tool, 405-407**

**user fields (SMC), 408-409**  
**user IDs (UIDs), 420**  
 effective UIDs, 445-446  
 setuid permission, 452-453  
 values, 441  
**user mask, 446**  
**useradd command, 417-420**  
**userdel command, 423**  
**usermod command, 422-423**  
**/usr directory**  
 assigning, 77  
 subdirectories of, 80-81  
**/usr file systems, 77**  
 backups, 606  
 recovering, 618-622  
**/usr/aset directory, 472-473**  
**/usr/lib/lp/model directory, 546**  
**/usr/share/lib/terminfo directory, 557**  
**utmpx file, 456**

## V

---

**-v option (physical device names), 30**  
**/var directory, 77, 81, 547-548**  
**/var/adm/loginlog file, 455-456**  
**/var/adm/messages file, 27**  
**/var/adm/sulog file, 467**  
**/var/adm/utmpx file, 456**  
**/var/adm/wtmpx file, 456**  
**/var/spool/lp directory, 547, 552**  
**Veritas Volume Manager. See VxVM (Veritas Volume Manager)**  
**vfstab (virtual file system table), 118-120**  
**video displays, configuring for x64/x86-based systems, 329-331**  
**virtual file systems, 48, 53-54**  
**volcancel command, 127**  
**volcheck command, 122, 127**  
**volcopy command, 105-106**  
**vold command, 127**

**vold daemon, 36, 120-128**  
**vold.conf file, 123-127**  
**volmissing command, 127**  
**volname, limitations of, 106**  
**volrmmount command, 36, 127**  
**volume management, 120-127**  
 directories/links, list of, 121  
 practice exercises, 146  
 troubleshooting, 129-131  
 utilizing effectively, 128  
**VTOC (volume table of contents), 50, 74-75**  
**VxVM (Veritas Volume Manager), 75**

## W - X

---

**-w (write mode), pax command, 603**  
**wait services (SMF), 363**  
**watchdog resets, 276**  
**Web Start Installer, 225**  
**who command, 456-457**  
**whoami command, 457-458**  
**whodo command, 458-459**  
**wildcards, cpio command, 602**  
**wipe disk standards, 91**  
**write mode (-w), pax command, 603**  
**wtmpx file, 456**

### x64-based systems

bootstrapping, 315-318  
 BIOS, 319-320  
 GRUB, 320-322  
 video display configuration, 329-331  
 crashed systems, interrupting, 377

### x86-based systems

bootstrapping, 315-318  
 BIOS, 319-320  
 GRUB, 320-322  
 crashed systems, interrupting, 377  
 slicing versus SPARC systems, 44-45, 54  
 video display configuration, 329-331

**738**

Xorg

**Xorg**

hardware management, 330-331

server configuration, 329

**Xsun, changing Xorg hardware management to,  
330-331**

---

**Y - Z**

**-y option (fsck), 103**

**zip command, 595**

**zombies, killing with preap command, 521**