

**Lucent Technologies**  
Bell Labs Innovations



**5ESS<sup>®</sup>-2000 Switch  
UNIX<sup>1</sup> RTR Operating System  
Reference Manual**

235-700-200  
Issue 7.00  
November 1998

---

1. Registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Limited.

**Copyright © 1998 Lucent Technologies. All Rights Reserved.**

This information product (IP) is protected by the copyright laws of the United States and other countries. It may not be reproduced, distributed, or altered in any fashion by any entity within or outside Lucent Technologies, except in accordance with applicable agreements, contracts, or licensing, without the express written consent of the Lucent Technologies Customer Training and Information Products organization and the business management owner of the IP.

This electronic document is protected by the copyright and trade secret laws of the United States and other countries. The complete document may not be reproduced, distributed, or altered in any fashion. Selected sections may be copied or printed with the utilities provided by the viewer software as set forth in the contract between the copyright owner and the licensee to facilitate use by the licensee, but further distribution of the data is prohibited.

For permission to reproduce or distribute, call:

1-888-LTINFO6 (1-888-584-6366) (Customers inside continental United States)

1-317-322-6848 (Customers outside continental United States).

**Notice**

Every effort was made to ensure that the information in this IP was complete and accurate at the time of printing. However, information is subject to change.

This IP describes certain hardware, software, features, and capabilities of Lucent products. The hardware, software, features, and capabilities described may not be the same as the equipment that you have. This IP is for information purposes, and you are cautioned that it may not describe your specific equipment.

**Mandatory Customer Information**

**Interference Information: Part 15 of FCC Rules**

Refer to the 5ESS<sup>®</sup>-2000 Switch Product Specification IP.

**Trademarks**

5ESS is a registered trademarks of Lucent Technologies.

ANSI is a registered trademark of American National Standards Institute.

AUTOPLEX is a registered trademarks of Lucent Technologies.

COMMON LANGUAGE is a registered trademark and CLEI, CLLI, CLCI, and CLFI are trademarks of Bell Communications Research, Inc.

ESS is a trademark of Lucent Technologies.

ETHERNET is a registered trademark of Xerox Corporation

KEYSTONE is a registered trademark of Control Data Corporation.

MC68030, MC68040, and MC68060 are trademarks of Motorola, Inc.

MOTOROLA is a registered trademark of Motorola, Inc.

SLC is a registered trademarks of Lucent Technologies.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Ltd.

**Warranty**

Warranty information applicable to the 5ESS<sup>®</sup>-2000 switch may be obtained from the Lucent Technologies Account Management organization. Customer-modified hardware and/or software is not covered by warranty.

**Ordering Information**

This IP is distributed by the Lucent Technologies Customer Information Center in Indianapolis, Indiana.

The order number for this IP is 235-700-200. To order, call:

1-888-LUCENT8 (1-888-582-3688) or fax to 1-800-582-9568 (Customers inside continental United States).

1-317-322-6848 (Customers outside continental United States).

**Support Telephone Numbers**

**Information Product Support Number:** To report errors or ask nontechnical questions about this or other IPs produced by Lucent Technologies, call 1-888-LTINFO6 (1-888-584-6366).

**Technical Support Numbers:** For initial technical assistance, call the Regional Technical Assistance Center (RTAC) at 1-800-225-RTAC (1-800-225-7822). For further assistance, call the Customer Technical Assistance Management Center (CTAM) at 1-800-225-4672 (Customers inside continental United States). For Customers outside continental United States, call 1-630-224-4672. The CTAM line is staffed 24 hours a day, 7 days a week.

**Acknowledgment**

Developed by Lucent Technologies Customer Training and Information Products.



# How Are We Doing?

Title: **5ESS<sup>®</sup>-2000 Switch UNIX RTR Operating System Reference Manual**

Identification No.: **235-700-200**

Issue No.: **7.00**

Date: **November 1998**

Lucent Technologies welcomes your feedback on this Customer Information Product (CIP). Your comments can be of great value in helping us improve our CIPs.

1. Please rate the effectiveness of this CIP in the following areas:

	Excellent	Good	Fair	Poor	Not Applicable
Ease of Use					////////////////////
Clarity					////////////////////
Completeness					////////////////////
Accuracy					////////////////////
Organization					////////////////////
Appearance					////////////////////
Examples					
Illustrations					
Overall Satisfaction					////////////////////

2. Please check the ways you feel we could improve this CIP.

- |  |   |
|--|---|
| <input type="checkbox"/> Improve the overview/introduction | <input type="checkbox"/> Make it more concise/brief                 |
| <input type="checkbox"/> Improve the table of contents     | <input type="checkbox"/> Add more step-by-step procedures/tutorials |
| <input type="checkbox"/> Improve the organization          | <input type="checkbox"/> Add more troubleshooting information       |
| <input type="checkbox"/> Include more figures              | <input type="checkbox"/> Make it less technical                     |
| <input type="checkbox"/> Add more examples                 | <input type="checkbox"/> Add more/better quick reference aids       |
| <input type="checkbox"/> Add more detail                   | <input type="checkbox"/> Improve the index                          |

Please provide details for the suggested improvement. \_\_\_\_\_  
 \_\_\_\_\_

3. What did you like most about this CIP?

\_\_\_\_\_  
 \_\_\_\_\_

4. Feel free to write any comments below or on an attached sheet.

\_\_\_\_\_  
 \_\_\_\_\_  
 \_\_\_\_\_

If we may contact you concerning your comments, please complete the following:

Name: \_\_\_\_\_ Telephone Number: \_\_\_\_\_

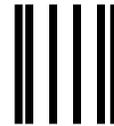
Company/Organization: \_\_\_\_\_ Date: \_\_\_\_\_

Address: \_\_\_\_\_

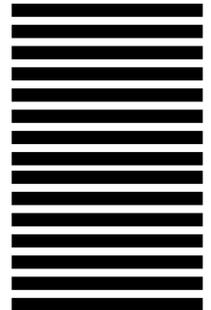
When you have completed this form, please fold, tape and return to address on back or Fax to: 336 727-3043.

Do Not Cut — Fold Here And Tape

**Lucent Technologies**  
Bell Labs Innovations



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES



**BUSINESS REPLY MAIL**

FIRST CLASS PERMIT NO. 1999 GREENSBORO, NC

POSTAGE WILL BE PAID BY ADDRESSEE

**DOCUMENTATION SERVICES**  
**2400 Reynolda Road**  
**Winston-Salem, NC 27199-2029**



# How Are We Doing?

Title: **5ESS<sup>®</sup>-2000 Switch UNIX RTR Operating System Reference Manual**

Identification No.: **235-700-200**

Issue No.: **7.00**

Date: **November 1998**

Lucent Technologies welcomes your feedback on this Customer Information Product (CIP). Your comments can be of great value in helping us improve our CIPs.

1. Please rate the effectiveness of this CIP in the following areas:

	Excellent	Good	Fair	Poor	Not Applicable
Ease of Use					////////////////////
Clarity					////////////////////
Completeness					////////////////////
Accuracy					////////////////////
Organization					////////////////////
Appearance					////////////////////
Examples					
Illustrations					
Overall Satisfaction					////////////////////

2. Please check the ways you feel we could improve this CIP.

- |  |   |
|--|---|
| <input type="checkbox"/> Improve the overview/introduction | <input type="checkbox"/> Make it more concise/brief                 |
| <input type="checkbox"/> Improve the table of contents     | <input type="checkbox"/> Add more step-by-step procedures/tutorials |
| <input type="checkbox"/> Improve the organization          | <input type="checkbox"/> Add more troubleshooting information       |
| <input type="checkbox"/> Include more figures              | <input type="checkbox"/> Make it less technical                     |
| <input type="checkbox"/> Add more examples                 | <input type="checkbox"/> Add more/better quick reference aids       |
| <input type="checkbox"/> Add more detail                   | <input type="checkbox"/> Improve the index                          |

Please provide details for the suggested improvement. \_\_\_\_\_  
 \_\_\_\_\_

3. What did you like most about this CIP?

\_\_\_\_\_  
 \_\_\_\_\_

4. Feel free to write any comments below or on an attached sheet.

\_\_\_\_\_  
 \_\_\_\_\_  
 \_\_\_\_\_

If we may contact you concerning your comments, please complete the following:

Name: \_\_\_\_\_ Telephone Number: \_\_\_\_\_

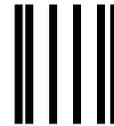
Company/Organization: \_\_\_\_\_ Date: \_\_\_\_\_

Address: \_\_\_\_\_

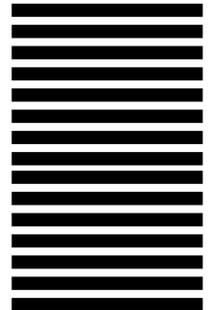
When you have completed this form, please fold, tape and return to address on back or Fax to: 336 727-3043.

Do Not Cut — Fold Here And Tape

**Lucent Technologies**  
Bell Labs Innovations



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES



**BUSINESS REPLY MAIL**

FIRST CLASS PERMIT NO. 1999 GREENSBORO, NC

POSTAGE WILL BE PAID BY ADDRESSEE

**DOCUMENTATION SERVICES**  
**2400 Reynolda Road**  
**Winston-Salem, NC 27199-2029**



**UNIX RTR Operating System Reference Manual**

	CONTENTS	PAGE
1.	INTRODUCTION . . . . .	1-1
2.	GETTING STARTED. . . . .	2-1
3.	ADMINISTRATION . . . . .	3-1
4.	EMACS DESCRIPTION. . . . .	4-1
5.	COMMANDS. . . . .	5-1
	GLOSSARY . . . . .	G-1
	INDEX . . . . .	I-1



## **UNIX RTR Operating System Reference Manual**

<b>CONTENTS</b>		<b>PAGE</b>
<b>1. INTRODUCTION.</b>		<b>1-1</b>
<b>1.1 PURPOSE</b>		<b>1-1</b>
<b>1.2 UPDATE INFORMATION</b>		<b>1-1</b>
<b>1.3 ORGANIZATION</b>		<b>1-1</b>
<b>1.4 USER FEEDBACK</b>		<b>1-3</b>
<b>1.5 DISTRIBUTION.</b>		<b>1-4</b>
<b>1.6 TECHNICAL ASSISTANCE</b>		<b>1-4</b>
<b>1.7 UNIX RTR OPERATING SYSTEM AVAILABILITY FEATURE.</b>		<b>1-4</b>
<b>1.7.1 DESCRIPTION</b>		<b>1-4</b>
<b>1.7.2 ORGANIZATION OF FILES.</b>		<b>1-5</b>



## 1. INTRODUCTION

### 1.1 PURPOSE

This manual is intended for use by operating telephone company personnel using the *UNIX* RTR operating system features of the Lucent 3B20D computer in the *5ESS*<sup>®</sup>-2000 switch. This manual provides the *UNIX* RTR operating system commands that are available, as well as instructions for setting up logins. This manual covers 5E11 and later software releases.

**IT IS ASSUMED THAT THE USER OF THIS MANUAL KNOWS HOW TO USE A *UNIX* RTR OPERATING SYSTEM ENVIRONMENT COMPUTER.** For those people who need to learn how to use the *UNIX* RTR operating system, formal training is available through one of Lucent's technical training locations. Contact your account representative for more information.

***Warning: The commands described in Chapter 5 of this manual are to be used as they are described. Any deviation from the procedures described, or misuse of these procedures, CAN result in a failure of the 5ESS-2000 switch.***

### 1.2 UPDATE INFORMATION

In accordance with the *5ESS-2000* Switch Software Support Plan, the 5E10 software release will be rated Discontinued Availability (DA) on December 4, 1998. The information supporting 5E10 and earlier has been removed from this information product.

This is an update to the **5ESS-2000 Switch *UNIX* RTR Operating System Reference Manual** 235-700-200, Issue 6.00 dated November 1997. This update covers software releases 5E11 through 5E13.

Since this is a complete reissue and every page has the same issue number (Issue 7.00), no page inventory is included with this update package.

It is recommended that the update instructions be retained for future reference. Retain your existing tabs and the binder with its cover and spine inserts, but replace all the existing text pages.

This manual has been updated to include new 5E13 software release information to add the following commands:

- vexpand (1)
- vcompress (1)

### 1.3 ORGANIZATION

This manual is divided into seven sections each separated by a tab. These sections are:

- Introduction
- Getting Started
- Administration
- EMACS Description
- Commands
- Glossary.

The *Introduction* gives a brief description of what this manual contains and a brief description of the *UNIX RTR* Operating System Availability feature.

The section on *Getting Started* contains a brief description of the following:

- Logging in
- Logging out
- How to communicate through your terminal
- How to run a program
- The current directory
- Pathnames
- Writing a program
- Text processing
- Surprises.

The section on *Administration* is intended for use by an administrator. This section describes procedures for securing the dial-up access to a supplementary trunk line work station or a text recent change terminal using new login software.

The *EMACS Description* section describes the EMACS screen editor and gives all the normally used commands and options that are available with the *UNIX RTR* operating system on the *5ESS-2000* switch.

The *Commands* section describes programs intended to be invoked directly by the user or by command language procedures, as opposed to subroutines, which are intended to be called by the user's programs. Commands generally reside in the directory */bin* (for **binary** programs). Some programs also reside in */usr/bin*, to save space in */bin*. The directories are also searched automatically by the command interpreter called the *shell*. Commands that reside in other directories are noted on the command page.

This section consists of many independent entries of a page or more each. The name of the entry appears in the upper corner of the page. Entries within the section are alphabetized. Some entries may describe several routines, commands, etc. In such cases, the entry appears only once, alphabetized under its "major" name.

All entries are based on a common format, not all of whose parts always appear.

- The **NAME** part gives the name(s) of the entry and briefly states its purpose.
- The **SYNOPSIS** part summarizes the use of the program being described.
- The **DESCRIPTION** part provides additional information about the program or facility outlined in the "Name" and "Synopsis" parts.
- The **EXAMPLES(S)** part gives example(s) of usage, where appropriate.
- The **FILES** part gives the filenames that are built into the program.
- The **SEE ALSO** part gives pointers to related information.

The **DIAGNOSTICS** part discusses the diagnostic indications that may be produced. Messages that are intended to be self-explanatory are not listed.

- The **WARNINGS** part points out potential pitfalls.

The **BUGS** part gives known bugs and sometimes deficiencies. Occasionally, the suggested fix is also described.

A few *conventions* are used in this section:

- **Boldface** strings are literals and are to be typed just as they appear.

*Italic* strings usually represent substitutable argument prototypes and program names found elsewhere in the manual. Note that this convention is not used in the "SYNOPSIS" or "SEE ALSO" part; regular print is used in place of italics.

Square brackets [ ] around an argument prototype show that the argument is optional. When an argument prototype is given as "name" or "file," it always refers to a *filename*.

Ellipses (...) are used to show that the previous argument prototype may be repeated.

A final convention is used by the commands themselves. An argument beginning with a ( -), plus (+), or equal sign (=) is often taken to be some sort of flag argument, even if it appears in a position where a filename could appear. Therefore, it is not wise to have files whose names begin with -, +, or =.

On most systems, all entries are available on-line via the *man*(1) command. As these on-line entries are updated from time to time, the on-line version may disagree with the entries in this manual. When they do disagree, the on-line version is the correct version to use.

#### 1.4 USER FEEDBACK

The producers of this IP are constantly striving to improve quality and usability. Please use the enclosed user feedback form for your comments and to advise us of any errors. If the form is missing or your comments will not fit, you can write to us at the following address:

Lucent Technologies Network Systems  
Documentation Services  
2400 Reynolda Road  
Winston-Salem, NC 27199-2029

Please include the issue number and/or date of this IP, your complete mailing address, and telephone number. We will attempt to answer all correspondence within 30 days, informing you of the disposition of your comments.

You may also call our documentation HOTLINE if you need an immediate answer to a documentation question. This HOTLINE is not intended to eliminate the use of the user feedback form, but rather to enhance the comment process. The HOTLINE number is **1-888-LTINFO6 (1-888-584-6366)** and it is available from 7:30 a.m. to 6:00 p.m., Eastern Standard time. Outside of those hours, the line is serviced by an answering machine. You can leave a message on the answering machine and someone will return your call the following business day.

Also, IP users who have access to *UNIX* system electronic mail facilities may send comments via electronic mail. The electronic address is **hotline5@wrddo.lucent.com**. Please make sure that the IP title, number, and issue number are included in the mail along with the sender's name, phone number, and address.

## 1.5 DISTRIBUTION

This document is distributed by the Lucent Technologies Customer Information Center (CIC) in Indianapolis, Indiana. Most operating telephone companies should place orders through their documentation coordinator. Some companies may allow customers to order directly from the CIC, however, the majority do not. Companies that use documentation coordinators to manage their orders receive a significant discount. If you do not know the name/number of the documentation coordinator for your company, you may call **1-888-LUCENT-8 (1-888-582-3688)** to obtain their name and telephone number.

Customers not represented by a documentation coordinator and Lucent employees can order the documentation for the 5ESS-2000 switch directly from the CIC. Proper billing information must be provided.

Mail these orders to the following address:

Lucent Technologies  
Customer Information Center  
Order Entry  
2855 N. Franklin Road  
Indianapolis, IN 46219

Orders may also be called in on **1-888-LUCENT-8 (1-888-582-3688)** or faxed in on **1-800-566-9568**.

## 1.6 TECHNICAL ASSISTANCE

Technical assistance for the 5ESS-2000 switch can be obtained by calling the Regional Technical Assistance Center (RTAC) at **1-800-225-RTAC**. This telephone number is monitored 24 hours a day, 7 days a week. During regular business hours, your call will be answered by your local RTAC. Outside of normal business hours, all calls will be answered at a centralized technical assistance center where service-affecting problems will be dispatched immediately to your local RTAC. All other problems will be referred to your local RTAC on the next regular business day.

## 1.7 UNIX RTR OPERATING SYSTEM AVAILABILITY FEATURE

### 1.7.1 DESCRIPTION

The UNIX RTR Operating System Availability feature allows the office craft to use some of the spare computing power of the 3B20D computer for small administrative tasks. For example, the craft could keep work schedules and spare circuit pack inventory lists on the 3B20D computer.

The following are features of the UNIX RTR Operating System Availability feature:

- Two new disk partitions; *unixa* and *unixabf*. The *unixa* partition is mounted as */unixa*. This partition contains the on-line manual pages and the control and spooling files for the *at*, *cron*, and *lpr* commands. The *unixabf* partition is mounted as */unixa/users*. This partition is provided as an area for the office craft to create files for their own use. The design of the system is such that excess file usage in these partitions will not overflow into the partitions needed for operation of the switch. In small disk configurations, the *unixabf* partition is not available. In these configurations, the *unixa* partition is also used for office craft files.

- A new *UNIX RTR* operating system command, *admin*, that allows the creation of nonroot logins. It is **strongly** advised that the root login only be used under direction of Lucent field support.
- A new *UNIX RTR* operating system command, *lpr*, to spool output to the ROP (receive-only printer).
- On-line manual pages and a *UNIX RTR* operating system command, *man*, to access them.
- New *UNIX RTR* operating system commands, *at* and *cron*, to allow scheduling of commands in the future.

### 1.7.2 ORGANIZATION OF FILES

The */unixa* or */unixa/users* directory system is the location that all the craft files will be created and stored. The */unixa* directory system is structured as follows:

```
/unixa/users/manager      # HOME directory for manager
|      /al                # HOME directory for al
|
|      /don               # HOME directory for don
|
|      /jim               # HOME directory for jim
|
|
|      /man/1.admin       # |
|      /...               # |On-line manual pages
|      /1.write           # |
|
|
|      /tmp                # Directory to hold temporary files
|
|
|      /spool/cron/crontabs/root # Root's crontab file
|                               |
|                               |      /don # Don's crontab file
|                               |
|                               |
|                               |      /atjobs/b.123456 # Batch job control file
|                               |      /b.6780 # Batch job control file
|                               |      /a.980 # At job control file
```

Unless directed by Lucent field support, craft *UNIX RTR* operating system users should create files only in (or below) their own HOME directory or the */unixa/tmp* directory. Also, no attempt should be made to delete or change any files outside these two directories.



**UNIX RTR Operating System Reference Manual**

	CONTENTS	PAGE
2.	GETTING STARTED . . . . .	2-1
2.1	BASIC INFORMATION . . . . .	2-1
2.2	LOGGING IN . . . . .	2-1
2.3	HOW TO COMMUNICATE THROUGH YOUR TERMINAL . . . . .	2-1
2.4	HOW TO RUN A PROGRAM. . . . .	2-2
2.5	THE CURRENT DIRECTORY . . . . .	2-3
2.6	PATHNAMES . . . . .	2-3
2.7	SURPRISES . . . . .	2-3



## 2. GETTING STARTED

### 2.1 BASIC INFORMATION

This section provides the basic information needed to get started, such as how to log in, how to communicate through your terminal, and how to run a program. The *UNIX* System Users' Handbook (320-042) and the *UNIX* System Quick Reference Guide (307-129) will provide you with general information on how the *UNIX* RTR operating system works. You should refer to the *UNIX* System Users' Handbook (320-042) if you need a more complete introduction to the system. For more information on manual pages referred to in this section but not found in the Commands Section, refer to the *UNIX* System Quick Reference Guide (307-129). Not all the commands and procedures found in the *UNIX* System Users' Handbook (320-042) and the *UNIX* System Quick Reference Guide (307-129) will work with the 5ESS<sup>®</sup>-2000 switch. Before using any commands and procedures on the 5ESS-2000 switch, be sure to read this manual thoroughly and use only the commands and procedures provided in this manual.

### 2.2 LOGGING IN

Most terminals have a speed switch that should be set to the appropriate speed and a half-/full-duplex switch that should be set to full-duplex. When a connection (at the speed of the terminal) has been established, the system types **login:**, and then you type your user name followed by the "return" key. If you have a password (and you should!), the system asks for it, but does not print ("echo") it on the terminal. After you have logged in, the "return," "new line," and "line feed" keys will give exactly the same result.

It is important that you type your login name in lowercase if possible; if you type uppercase letters, the system assumes that your terminal cannot generate lowercase letters and that you mean all further uppercase input to be treated as lowercase. When you have logged in successfully, the shell returns a \$. (The shell is described "How to Run a Program, Section 2.4").

For more information, consult *login*(1), which discusses the login sequence in more detail; and *stty*(1), which tells you how to describe the characteristics of your terminal to the system. *Profile*(4) explains how to accomplish this task automatically every time you log in. For information on the logout procedure, see Section 3.5.2.

### 2.3 HOW TO COMMUNICATE THROUGH YOUR TERMINAL

When you type input, the system gathers your characters and saves them. These characters are not given to a program until you type a "return" (or "new line"), as described in "Logging In, Section 2.2."

Your terminal is a full-duplex input/output terminal. It has full read-ahead capability, which means that you can type at any time, even while a program is prompting you. Of course, if you type during output, the output is interspersed with the input characters; however, your input is saved and interpreted in the correct sequence. There is a limit to the amount of read-ahead, but it is generous and not likely to be exceeded unless the system is in trouble. When the read-ahead limit is exceeded, the system throws away *all* the saved characters.

On an input line from a terminal, the character @ "kills" all the characters typed before it. The character # erases the last character typed. Successive uses of # erase characters back to, but not beyond, the beginning of the line; @ and # can be typed as

themselves by preceding them with \ (thus, to erase a \, you need two #s). These default erase and kill characters can be changed; see *stty(1)*.

The ASCII **DC3** (CONTROL-S) character can be used to temporarily stop output. It is useful with CRT terminals to prevent output from disappearing before it can be read. Output is resumed when a **DC1** (CONTROL-Q) or a second **DC3** (or any other character, for that matter) is typed. The **DC1** and **DC3** characters are not passed to any other program when used in this manner.

The ASCII **DEL** (a.k.a. "rubout") character is not passed to programs; but instead generates an *interrupt signal*, just like the "break," "interrupt," or "attention" signal. This signal, treated as lowercase, generally causes whatever program you are running to stop. It is typically used to stop a long printout that you do not want. However, programs can arrange either to ignore this signal altogether, or to be notified when it happens (instead of being stopped). The editor *ed(1)*, for example, catches interrupts and stops what it is doing, instead of terminating, so that an interrupt can be used to halt an editor printout without losing the file being edited.

The *quit* signal is generated by typing the ASCII **FS** character. It not only causes a running program to stop, but also generates a file with the "core image" of the stopped process. *Quit is useful for debugging.*

Besides adapting to the speed of the terminal, the *UNIX* RTR Operating System tries to determine whether you have a terminal with the "new-line" function, or whether it must be simulated with a "carriage-return" and "line-feed" pair. In the latter case, all *input* "carriage-return" characters are changed to "line-feed" characters (the standard line delimiter); and a "carriage-return" and "line-feed" pair is echoed to the terminal. If you get into the wrong mode, use the *stty(1)* command to correct it.

Tab characters are used freely in system source programs. If your terminal does not have the tab function, you can arrange to have tab characters changed into spaces during output and echoed as spaces during input. The *stty(1)* command can be used to set or reset this mode. The system assumes that tabs are set every eight-character positions.

## 2.4 HOW TO RUN A PROGRAM

When you have successfully logged into the system, a program called the shell is listening to your terminal. The shell reads the lines you type, splits them into a command name and its arguments, and executes the command. A command is simply an executable program. Normally, the shell looks first in your current directory (see "The Current Directory", Section 2.5 ) for a program with the given name. If none is there, it looks in system directories. There is nothing special about system-provided commands except that they are kept in directories where the shell can find them. You can also keep commands in your own directories and arrange for the shell to find them there.

The command name is the first word on an input line to the shell; the command and its arguments are separated from one another by space and/or tab characters.

When a program stops, the shell ordinarily regains control and returns a \$ to show that it is ready for another command. The shell has many other capabilities, which are described in detail in *insh(1)*.

## 2.5 THE CURRENT DIRECTORY

The *UNIX* RTR Operating System file system is arranged in a hierarchy of directories. When the system administrator gives you a user name, he or she also creates a directory for you (ordinarily with the same name as your user name). (This directory is known as your *login* or *home* directory.) When you log in, your home directory becomes your *current* or *working* directory; and any filename you type is by default assumed to be in that directory. Because you are the owner of this directory, you have full permissions to read, write, alter, or destroy its contents. Permissions for directories and files other than your own may be granted or denied to you by their owners, or by the system administrator. To change the current directory, use *cd* [explained under *sh*(1)].

## 2.6 PATHNAMES

To refer to files not in the current directory, you must use a pathname. Full pathnames begin with */*, which is the name of the *root* directory of the whole file system. Following the initial slash, each directory and subdirectory is named in succession, until the filename is reached. Each directory name is followed by */*. For example, */usr/ae/filex* refers to file *filex* in directory *ae*, while *ae* is itself a subdirectory of *usr*; *usr* springs directly from the root directory.

If your current directory contains subdirectories, the pathnames of files therein begin with the name of the corresponding subdirectory (*without* a prefixed */*). Without important exception, a pathname may be used anywhere a filename is required.

Important commands that change the contents of files are *cp*(1), *mv*(1), and *rm*(1), which copy, move (i.e., rename), and remove files. To find out the status of files or directories, use *ls*(1). Use *mkdir*(1) for making directories and *rmdir*(1) for destroying them.

For further discussion of the file system, see the *UNIX* System Users' Handbook (320-042).

## 2.7 SURPRISES

Certain commands provide inter-user communication. Even if you do not plan to use them, it would be well to learn something about them, because someone else may direct them to you. To communicate with another user currently logged in, *write*(1) is used; *mail*(1) leaves a message whose presence is announced to another user when he or she next logs in.

When you log in, a message-of-the-day may greet you before the first \$.



## UNIX RTR Operating System Reference Manual

CONTENTS		PAGE
3.	ADMINISTRATION . . . . .	3-1
3.1	INTRODUCTION . . . . .	3-1
3.1.1	DIAL-UP FACILITIES . . . . .	3-1
3.1.2	LOGIN CAPABILITY FEATURES . . . . .	3-1
3.2	HARDWARE. . . . .	3-1
3.3	SOFTWARE . . . . .	3-1
3.3.1	GENERAL . . . . .	3-1
3.3.2	ADDING NEW LOGINS . . . . .	3-1
3.3.3	ADDING A SECOND PASSWORD (OPTIONAL) . . . . .	3-2
3.3.4	DATA BASE . . . . .	3-3
3.4	INITIALIZE THE TELETYPEWRITER CONTROLLER . . . . .	3-4
3.5	USING THE LOGIN . . . . .	3-4
3.5.1	GENERAL . . . . .	3-4
3.5.2	LOGGING OFF . . . . .	3-4
3.6	PASSWORD PROTECTED COMMANDS . . . . .	3-4
3.7	MODIFYING THE <i>RM</i> COMMAND DEFAULT OPTIONS . . . . .	3-6

### LIST OF FIGURES

Figure 3-1	— Diagram of Cable From Modem to IOP . . . . .	3-7
Figure 3-2	— Description of an Entry in <i>/etc/passwd</i> . . . . .	3-8



### 3. ADMINISTRATION

#### 3.1 INTRODUCTION

##### 3.1.1 DIAL-UP FACILITIES

The use of remote dial-up facilities such as the dial-up supplementary trunk line work station (STLWS) and the dial-up text recent change terminal has greatly enhanced the effectiveness of the field support effort.

A shortcoming of these dial-up facilities is that, unless the data set has been turned off, there is nothing to prevent an unauthorized person from obtaining access to a *5ESS*<sup>®</sup>-2000 switch. For convenience, the data sets are usually left powered-on.

This section describes the procedures for securing the dial-up access to an STLWS or a text recent change terminal using new login software available with the *UNIX RTR* operating system in 5E11 and later software release offices.

##### 3.1.2 LOGIN CAPABILITY FEATURES

The login software provides for the conventional login and password security of a *UNIX RTR* operating system user environment. While this section details procedures for adding a login to the STLWS or the text recent change terminal, any terminal [except the Master control center (MCC) and remote maintenance facilities] can be modified to require a login by making the necessary equipment configuration (ECD) data changes.

The login code provides some additional flexibility in the configuration of the terminals. For users requiring only *UNIX RTR* operating system capability, the login software allows the STLWS and the recent change terminals to be logged in directly as *UNIX RTR* operating system terminals.

This section also describes an optional second password (similar to the External Security Code or Network Access Password as used on other systems) which may be required for all logins.

#### 3.2 HARDWARE

The only hardware concern for the login capability is the cable from the input/output processor (IOP) to the data set. This cable is slightly different from a standard IOP cable used with terminals in the office. The standard IOP cable, designed for use with terminals, will not work with a data set. A diagram of the correct cable is shown in Figure 3-1. The correct cable includes all necessary connections to allow the data set to signal the IOP when a user hangs up the telephone. If no login is required, this signaling is not necessary. If the port has been updated to require a login, and an incorrect cable is used, only the first user after each TTY restore, would be required to login.

#### 3.3 SOFTWARE

##### 3.3.1 GENERAL

Software changes for the login capability include updating the */etc/passwd* file to establish the login and recent change modifications to execute the login process.

##### 3.3.2 ADDING NEW LOGINS

As with any *UNIX RTR* operating system, new logins are added by updating the file */etc/passwd*. The user must be in *UNIX RTR* operating system to add new logins. Enter the *UNIX RTR* operating system with the command:

```
rcv:menu,sh;
```

Adding new logins is done by using the *admin(1)* tool available on *UNIX RTR* operating systems. [See the manual page for *admin(1)* in the Commands section of this manual.]

*Admin(1)* makes all required entries in the */etc/passwd* file while protecting other entries from possible corruption. *Admin(1)* does not provide a password. The user will be prompted for a password when the login is first used.

To secure the dial-up terminals, the following three logins may be established using *admin(1)*:

1. To set up a *UNIX RTR* operating system login, enter:

```
admin -a login name -u
```

2. To set up a Craft shell login, enter:

```
admin -a login name -p
```

In the examples above, the **-a** option is the login name; the **-u** option shows that this is a *UNIX RTR* operating system login; and the **-p** option shows that this login will use a *pds/mml* shell.

Note that *login name* is a variable that is replaced by the actual login.

Printing the */etc/passwd* should show the new logins. A description of an entry in */etc/passwd* is shown in Figure 3-2.

### 3.3.3 ADDING A SECOND PASSWORD (OPTIONAL)

More extensive security is desirable by requiring a second password. To add the second password, enter:

```
passwd daemon
```

The user will be prompted to enter the additional password.

When dial-up users log on, they will be prompted for the "Machine Password." The daemon password is entered in response.

Note there is only one "Machine Password" for all logins. This provides an effective method of securing all logins without changing individual passwords.

To clear the daemon password, enter:

```
admin -c daemon
```

This command will remove the daemon password, and future users will not be required to enter one.

When the new logins and the optional second password have been entered, exit the *UNIX RTR* operating system by entering a *CONTROL-D*.

**Note:** Procedures in sections 3.3.1 through 3.3.2, though still supported, are not recommended with the advent of the *5ESS-2000* Switch Command Restriction procedure. The preferred login and password administration scheme is called *authority management*. Authority management is documented in the Operations section of Lucent 235-105-210, *Routine Operations and Maintenance*.

### 3.3.4 DATA BASE

#### 3.3.4.1 General

Only a few data base changes are required to convert an STLWS or recent change terminal into a secure dial-up facility. It is assumed that the user is familiar with the use of recent change and with equipping these dial-up facilities. The following procedures assume that the teletypewriter controller and the teletypewriter are equipped and that all hardware has been connected to the proper location on the IOP backplane.

#### 3.3.4.2 Data Base Changes for a Dial-Up STLWS

Equipping a secure dial-up STLWS terminal requires four ECD changes. These changes lower the baud rate to 1200 baud and establish paths to execute the login process. The baud rate is lowered to 1200 baud because the dial-up modems normally operate at 1200 baud. The following ECD changes are required to set up a secure dial-up STLWS:

##### Lower the Baud Rate

Lowering the STLWS baud rate to 1200 baud requires one change on the **ciopt** form. If your terminal is on port /dev/ttyxx, then the correct **ciopt** form to select is ttyopxx. For example, if your terminal is on port /dev/tty9, select **ciopt** form ttyop9, /dev/tty10 select ttyop10, etc. On this form, change ttopt\_name (field 2) to **PDS12** (black & white terminal) or **PDS12C** (color terminal).

##### Large Scroll Region:

On Page 4 of the **cdopt** forms STLWSCG (color terminal) and VT100DAP (black & white terminal), change field 60, sub-fields 11 and 12, from '2', '3' to '0', '2'. Change field 94, first item from '23' to '2', also.

##### Set Up Paths for the Login Process

To execute the login process on an STLWS terminal, ECD changes are required on two **getty** forms. Two forms are involved since one form is required for the message section of the STLWS screen and a second form is used by the display region of the screen. The two "getty" form names differ only in that the TTY name in gettyres (field 1) is lowercase on one form and uppercase on the other form; for example: **gettyl** and **gettyL**. The same change is required on both **getty** forms. Change shlname (field 4) to **/etc/login**.

#### 3.3.4.3 Data Base Changes for a Dial-Up Recent Change Terminal

To equip a dial-up text recent change terminal requires only two ECD changes. These changes lower the baud rate to 1200 baud and establish a path to execute the login process.

##### Lower the Baud Rate

Lowering the recent change terminal baud rate to 1200 baud may require one change on the **ciopt** form. If your terminal is on port /dev/ttyxx, then the correct **ciopt** form to select is ttyopxx. For example, if your terminal is on port /dev/tty9, select **ciopt** form ttyop9, /dev/tty10 select ttyop10, etc. On this form, change ttopt\_name (field 2) to **PDS12**.

##### Set Up the Path for the Login Process

To execute the login process for a dial-up recent change terminal, an ECD change is required on one **getty** form. The TTY name in gettyrec (field 1) is lowercase (for example: **gettyw**). Only one form requires updating, since this is a text recent change terminal and there is no display region.

On the **getty** form, change shlname (field 4) to **/etc/login**.

#### 3.3.4.4 Computer Access Restriction

For increased security the Computer Access Restriction (CAR) feature may be used. This feature can be used to restrict access to your modem line(s) such that only callers from a listed set of telephone numbers can get through.

Instructions for setting up the CAR feature can be found in the *LASS User's Manual (235-190-300), Section 9, Issue 2.00*.

### 3.4 INITIALIZE THE TELETYPEWRITER CONTROLLER

After the ECD changes have been completed, the new options must be down loaded into the teletypewriter controller (TTYC) for the secured ports. This is done by restoring the TTYC. To restore the TTYC, enter:

```
RST:TTYC=aa: (mml) RST:TTYC aa! (pds) Where aa is the TTYC number.
```

### 3.5 USING THE LOGIN

#### 3.5.1 GENERAL

When the login process has been successfully completed, the user may continue with a normal session. The only exception may be a message, immediately after login, asking the user to hit the "break" key. This message will appear only when logging in as a *UNIX RTR* operating system shell system terminal.

#### 3.5.2 LOGGING OFF

When logging off a dial-up recent change terminal or a dial-up *UNIX RTR* operating system terminal, the user needs to only enter a CONTROL-D. Logging off the dial-up STLWS is more complex since the display and message regions are, in effect, two separate screens. The correct procedure for logging off an STLWS is:

1. Go to the message page (Page 120)
2. Go to "message mode"
3. Enter CONTROL-D.

These steps are necessary because the display region is frozen to prevent it from scrolling off the screen. Since the CONTROL-D will not be accepted in "command mode," the user must be in "message mode" to log off. If the next user logs on as a *UNIX RTR* operating system terminal, they will be limited to the bottom of the display. By going to the message page (Page 120) only, the top few lines are frozen giving a *UNIX RTR* operating system user more work space.

If a user logs on as a *UNIX RTR* operating system terminal and finds that he/she is stuck on the bottom of the display, the only easy solution is to reset the terminal. This can be done by pressing the "reset" key on the terminal, if it has one, or power cycling the terminal.

### 3.6 PASSWORD PROTECTED COMMANDS

This feature, if enabled by the site, restricts access to the *UNIX RTR* operating system shell, Office Data Base Editor (ODBE), and/or Access Editor (ACCED)

- RCV:MENU:SH
- RCV:MENU:ODBE
- RCV:MENU:ACCED

- RCV:MENU:SCREEN
- POKE 194.

The feature is enabled as follows:

From the *UNIX* RTR operating system, create new logins for each command you wish to protect using the `-R` option of the `admin` command. The login must be the name of the command followed by a dash (-). Enter

```
admin -a sh- -R    # To protect the UNIX RTR operating system
                  shell.
admin -a odbe- -R # To protect ODBE.
admin -a acced- -R # To protect ACCED.
```

Then add a password to each of these logins. Enter

```
passwd sh-      # To add a password to the UNIX RTR
                operating system shell.
passwd odbe-    # To add a password to ODBE.
passwd acced-   # To add a password to ACCED.
```

From this point, the selected command(s) will prompt for the corresponding login's password.

The feature may be disabled as follows:

From the *UNIX* RTR operating system, delete the login corresponding to the command. Enter

```
admin -d sh- -R    # To unprotect the UNIX RTR operating
                  system shell.
admin -d odbe- -R # To unprotect ODBE.
admin -d acced- -R # To unprotect ACCED.
```

An additional option is also provided that changes the access permissions of the *UNIX* RTR operating system shell such that the user has read-only permission except in the "tmp" directories. With this option selected, the shell user also will not be able to run some of the *UNIX* RTR operating system commands.

To select this option, create the "sh-" login with the `'-r'` option rather than the `'-R'` option and enter

```
admin -a sh- -r # Limited permission shell
```

If you select this option, enter the shell, and want full permissions, you must use the "su" command and know the password for the login 'root,' then enter

```
su root
<root's password>
```

Note that the "su" and "passwd" commands only operate when entered by RCV:MENU: SH. They will not operate from RCV:MENU:SCREEN nor POKE 194.

Password protected commands is **NOT** a security feature. It can be overridden by most experienced Lucent and telephone company field support personnel. It is provided, as requested, for the case where a site wants to block inexperienced personnel and/or clerks from accessing these commands.

### 3.7 MODIFYING THE *RM* COMMAND DEFAULT OPTIONS

The *UNIX*RTR operating system "rm" command supports an interactive option ("-i"), which prompts the user for confirmation before removing each file. This provides a second chance to avoid accidental removal of a file or files.

In some situations, it may be desirable to have the "rm" command invoked with the interactive option by default. This can be accomplished by setting up a shell alias for the "rm" command. The procedure for establishing the alias is as follows:

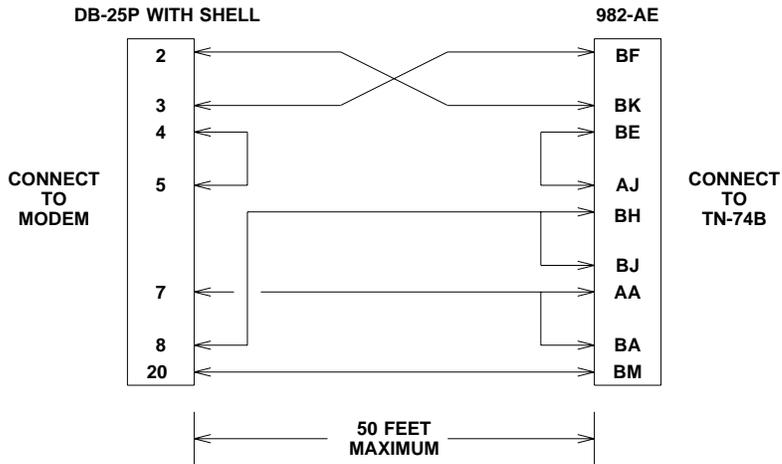
1. From the *UNIX*RTR operating system shell, copy the default shell alias file (provided) to the name expected by the *UNIX*RTR operating system shell. This file contains an alias for the "rm" command.  

```
# cp /etc/default.shrc /etc/shrc
```
2. Exit any active *UNIX*RTR operating system shells, including those invoked using RCV:MENU or SCREEN. (The alias file is only read when the shell is invoked.)

The shell "rm" alias is now active for all interactive shells, until the **etc/shrc** file created above is removed. When the "**rm**" command is used, the user will be prompted for each file before it is removed, just as if "rm -i" had been entered. This includes all shells established using RCV:MENU and the SCREEN program, but will not affect non-interactive shell scripts or applications.

The "rm" alias can be bypassed either by using the full path of the "rm" command (**/bin/rm**) or by executing the "**unalias**" command. Unalias removes the alias for the remainder of the shell session or until the alias is re-established by executing the "**alias**" command.

See the *UNIX*RTR operating system shell (sh) manual page for more information.



982-AE CONNECTOR  
LOOKING FROM THE FRONT

AN	BN
AM	BM
AL	BL
AK	BK
AJ	BJ
AH	BH
AF	BF
AE	BE
AD	BD
AC	BC
AB	BB
AA	BA

NOTES: →

1. Side A Has Solder Pads Around ALL Holes
2. Be Sure To Configure Your Modem Such That DCD (Pin 8) Drops When The Calling Party Disconnects.

↖ EQL NUMBER INDICATES THIS PIN

Figure 3-1 — Diagram of Cable From Modem to IOP

```
stlw:2xMm;YhM,MOXA:1329:1329:Dialup STLWS: /cft/sh1: /cft/bin/pdshel.app  
|_1_| |____2_____| |_3_| |_4_| |____5_____| |_6_| |____7_____|
```

- [1] Login id
- [2] Encrypted password
- [3] User id
- [4] Group id
- [5] Comment field (optional)
- [6] Home directory after successful login
- [7] Command to execute after successful login (optional)

**Figure 3-2 — Description of an Entry in /etc/passwd**

## UNIX RTR Operating System Reference Manual

	CONTENTS	PAGE
4.	EMACS DESCRIPTION . . . . .	4-1
4.1	INTRODUCTION . . . . .	4-1
4.2	BASIC CONCEPTS . . . . .	4-1
4.2.1	GENERAL . . . . .	4-1
4.2.2	THE CHARACTER SET . . . . .	4-1
4.2.3	THE DISPLAY. . . . .	4-2
4.2.4	THE TEXT IN THE BUFFER . . . . .	4-2
4.2.5	COMMAND STRUCTURE . . . . .	4-3
4.2.6	ARGUMENTS AND PARAMETERS. . . . .	4-3
4.3	BASIC EMACS COMMANDS . . . . .	4-3
4.3.1	GENERAL . . . . .	4-3
4.3.2	GETTING HELP OR GETTING OUT OF TROUBLE. . . . .	4-4
4.3.3	SIMPLE CURSOR MOVEMENT COMMANDS. . . . .	4-5
4.3.4	SIMPLE TEXT DELETING AND MOVING COMMANDS . . . . .	4-6
4.3.5	SIMPLE FILE AND BUFFER COMMANDS . . . . .	4-7
4.3.6	SIMPLE SEARCH AND REPLACE . . . . .	4-10
4.4	ADVANCED EDITING COMMANDS . . . . .	4-12
4.4.1	GENERAL . . . . .	4-12
4.4.2	INSERTING ODD CHARACTERS . . . . .	4-12
4.4.3	COMMANDS RELATED TO WINDOWS . . . . .	4-12
4.4.4	ADVANCED SEARCH AND REPLACE COMMANDS . . . . .	4-13
4.4.5	MACROS, KEYBOARD MACROS, AND INPUT FILES . . . . .	4-14
4.4.6	COMMANDS THAT ESCAPE TO THE <i>UNIX RTR</i> OPERATING SYSTEM . . . . .	4-15
4.4.7	MISCELLANEOUS COMMANDS. . . . .	4-16
4.5	MODES . . . . .	4-19
4.5.1	GENERAL . . . . .	4-19
4.5.2	DISPLAY MODES (PARAMETERS). . . . .	4-19
4.5.3	INTERFACE MODES . . . . .	4-20
4.5.4	COMMAND MODES . . . . .	4-22
4.5.5	TERMINAL MODES . . . . .	4-24
4.5.6	SPECIFYING DEFAULT MODES FOR A FILE. . . . .	4-25
4.6	GETTING STARTED. . . . .	4-25
4.6.1	GENERAL . . . . .	4-25
4.6.2	TERMINAL TYPE . . . . .	4-25
4.6.3	INITIALIZATION FILE . . . . .	4-25
4.6.4	COMMAND LINE ARGUMENTS . . . . .	4-26
4.6.5	HELPFUL HINTS . . . . .	4-26

4.6.6	LIMITATIONS OF THE EDITOR . . . . .	4-27
4.6.7	RECOVERING FROM VARIOUS PROBLEMS . . . . .	4-27
4.7	CONCLUSIONS . . . . .	4-27

LIST OF FIGURES

Figure 4-1	— EMACS Screen Display . . . . .	4-28
------------	----------------------------------	------

LIST OF TABLES

Table 4-1	— EMACS COMMAND SUMMARY . . . . .	4-28
-----------	-----------------------------------	------

## 4. EMACS DESCRIPTION

### 4.1 INTRODUCTION

EMACS is a screen editor that can be used to create or to edit files using a display terminal. The user interface to this editor is quite simple. The user is presented with a display of the contents of a portion of the buffer being edited. This display indicates *exactly* what is in the area being displayed, including any non-printing characters. The contents of the buffer being edited can be read from or written to a *UNIX* RTR operating system file. Characters typed by the user will be inserted into the buffer (and reflected in the display) at the point indicated by the terminal's cursor. This is the primary mechanism for entering and modifying text.

Control characters and escape sequences can be used to perform other editing functions, such as moving the cursor to a different position in the buffer, deleting text, replacing text, or searching. Thus there is only one mode of interpretation of characters typed to EMACS, in which either text to be entered or commands can be entered. This simple interface relieves users of the need to remember what mode they are in, and prevents the disastrous mistakes that can occur when text to be inserted is evaluated as an editor command. A simple mechanism is provided to allow a user to insert control and escape characters when needed.

Although there is a rich vocabulary of commands available, including commands that perform functions tailored to a particular application (such as indenting a C program), the most common way in which EMACS is used to edit is simply to position the cursor over the area to be changed, and enter the changes. The immediate feedback provided by the visual display appears to be very important to the user.

Table 4-1 provides a summary of EMACS commands.

### 4.2 BASIC CONCEPTS

#### 4.2.1 GENERAL

Before going into the editing commands of EMACS, some basic concepts should be learned. EMACS operates rather differently from line oriented editors, and even from other screen oriented editors in the way that it treats the screen and the keyboard. Some of its conventions for displaying and inputting characters are not like other *UNIX* RTR operating system tools, primarily because they were originally developed for another environment.

#### 4.2.2 THE CHARACTER SET

EMACS operates on characters from an alphabet of 256 different characters. These include the 128 ASCII characters that can be entered from a terminal, and 128 "Meta" characters. A Meta character is entered by preceding it with an escape (ESC key).

In this document and in the displays produced by EMACS, control characters are indicated by the character '^' followed by the equivalent printable character (usually capitalized). Thus '^X' represents a control-x, which is typed by hitting the control and 'x' keys simultaneously. For some unusual non-printing characters, the display is not obvious:

^?	Rubout or delete (ASCII 0177)
^@	Null (ASCII 0)
^[	Escape (ASCII 033)
^\	The "fs" character (ASCII 034)

- ^]           The "gs" character (ASCII 035)
- ^^           The "rs" character (ASCII 036)
- ^\_           The "us" character (ASCII 037)

Meta characters are typed to EMACS by hitting the escape character, and then any second character (including a control character). They are displayed by EMACS as "M-" followed by the equivalent ASCII character. Thus "M-a" (Meta - a) is the character obtained by typing escape followed by a, and "M-^B" (Meta - control-b) is the character obtained by typing escape followed by control-b.

#### 4.2.3 THE DISPLAY

The display screen contains a window showing a view of the buffer being edited, which contains about 20 lines on a typical display terminal. The terminal cursor is positioned at the editor cursor (the position where editing takes place). Each line of the buffer (delimited by a newline character) begins at the beginning of a display line. A line that exceeds the screen width is normally continued on the next screen line. Whenever a line must be continued on the next screen line an exclamation mark (!) is displayed in the last column of the first screen line. If the editor is in line number (*Inumb*) mode, then a line number is printed at the beginning of each line in the buffer.

Printable characters are displayed normally, while tabs are displayed as white space that fills up the space on the screen until the next position at a multiple of eight. Non printing control characters and meta characters are printed with the conventions outlined above. If you edit a file which contains characters that have the high order (parity) bit set, they are displayed as meta characters by EMACS. You will only run into this when trying to edit files containing binary information.

In addition to the display buffer, several lines of the screen are used for status information and for displaying parameters entered into EMACS, such as a file name. One of these lines known as the status line contains the editor name, editor version, buffer number and name, and file name. Some of the more recently introduced commands described in this document indicate the version in which they were introduced, so that you can determine whether or not a particular command is in the version that you are running. If the buffer has not been modified since the file was read or written, an '=' will be displayed between the buffer and file names. Otherwise, a '>' will appear.

The lines below the status line are used for the time of day display (if *time* mode is on), and for EMACS to prompt for parameters for commands. Some commands cause the buffer display to be erased in order to display other information in place of the buffer. The word "Continue?" will be displayed at the bottom of the screen when this happens. Typing 'y', ' ', or return will bring back the buffer display. Typing 'n' may allow you to re-execute the command producing the display.

Figure 4-1 shows a typical screen during a EMACS session. The buffer "Main", number 0, is being used to edit a program test.c. The buffer has been modified since the last write to the file test.c.

#### 4.2.4 THE TEXT IN THE BUFFER

Each buffer that you edit holds a sequence of characters. Any characters can be present in an EMACS buffer, including control and meta characters. The only limitation is on the number of characters that can be on one line in the buffer. Normally, EMACS treats the buffer just as a sequence of characters.

One difference between EMACS and many editors is that EMACS does not treat "newline" characters specially. Between each pair of adjacent lines of text in the buffer is an invisible "newline" character. If the cursor is at the end of one line, it is in front of the newline character, and deleting a single character will delete the newline, causing the text in the following line to be joined to the current line. Newlines can be inserted, deleted, and searched for like any other characters.

Some EMACS commands operate on units of text in the buffer, like words, lines, sentences, pages, etc. These work on top of the base level which still treats the buffer as a string of characters.

#### 4.2.5 COMMAND STRUCTURE

Unlike most other editors, EMACS does not have distinct "modes" for inserting text into the buffer and for entering commands. Thus there are no commands for inserting text, and no special convention to end a text insertion. Instead, ordinary characters can be inserted into the buffer at any time simply by typing them, while control and meta characters are used for editing commands.

Each character that is typed into EMACS is in fact interpreted as a command. All of the ordinary printing characters insert themselves into the buffer being edited at the point defined by the cursor. Thus the command invoked when you type the character 'x' inserts an x into the buffer at the point shown by the cursor. The control and meta characters are used for various editing functions.

#### 4.2.6 ARGUMENTS AND PARAMETERS

All commands, including the printing characters, take a numeric argument that has some effect on their interpretation. The default argument given to a command for which no argument is specified is 1. To specify some other argument to a command, you can enter escape, followed by a sequence of digits, and then the command. You can specify a negative value for an argument by entering escape followed by '-', followed by a sequence of digits. Numbers starting with a 0 are interpreted as octal, while numbers starting with any other digit are decimal. A second way of specifying the argument is to precede the command by one or more ^U (control-u) characters. Each ^U multiplies the value of the argument by 4.

For most commands, the effect of the argument is to multiply the number of times that the command is applied. Thus the sequence ^U^Ux inserts 16 x's into the buffer at the current location. The sequence ESC13^N moves forward 13 lines in the buffer.

In addition to the numeric argument given to all commands, some commands will prompt the user for additional character string parameters. The commands that take parameters and the method of entering parameters are described in the section on file and buffer commands, Section 4.3.5.

### 4.3 BASIC EMACS COMMANDS

#### 4.3.1 GENERAL

As noted above, every character you type to EMACS is interpreted as a command. This section describes a simple set of commands that will be sufficient for most editing that you do. Subsequent sections describe more advanced commands that are very useful in certain situations, and other aspects of EMACS.

EMACS has a large number of commands. Most of them have a mnemonic significance that should be obvious (like ^B for backwards or ^D for delete). Some, unfortunately, don't have any obvious meaning. As a general rule, control character commands

operate on characters and lines, while the corresponding meta character commands operate on words or sentences.

The user interface of EMACS was designed for touch typists. A deliberate choice was made to use control and meta characters for commands rather than special keys, such as the function or "arrow" keys on many terminals, since these keys are different on every terminal and generally cannot be reached without taking your fingers off of the home position of the keyboard. Macros can frequently be used to allow EMACS to respond to the arrow and function keys on a particular terminal, if desired. Typically the only difficulty faced by an EMACS user in adapting to a new terminal is locating the escape key, which is unfortunately located differently on every terminal. With a little practice, you will find that your fingers become adept at locating the keys for all of the basic commands with little thought and without having to look at the keyboard.

#### 4.3.2 GETTING HELP OR GETTING OUT OF TROUBLE

EMACS has many self-help features. The commands listed in this section are useful to know about because they can provide help or remedy mistakes.

- M-? Explain. This command prompts for a character and prints a brief explanation of what that character does.
- M-w Wall Chart. This command puts a listing of all commands (including user defined commands), and their help explanations into the current buffer. This command is a convenient way of producing a "wall chart" of the commands. The list is inserted into the buffer at the current position, so that normally one would want to execute it in an empty buffer. Table 4-1 contains a current copy of the wall chart.
- ^L Refresh. Refresh the display. Occasionally, some error may cause the display to become garbled. ^L re-creates it from scratch. If you give an argument to ^L, it is used to specify how many lines will appear on the screen before the current line. When invoked with an argument, this command does not re-create the display from scratch.
- ^G Abort. Typing ^G at any point that EMACS is asking for input will abort the current command. This applies at any step (specifying arguments, typing escape, entering parameters that EMACS asks for, etc.). (There are a couple of exceptions related to advanced editing commands, but even with these, typing ^G several times will always get you out with minimal damage). This is a convenient way of aborting anything that you are not sure that you want to complete, and may not know how you started.
- M-u Undo. This command undoes the effect of the last significant text modifying command that can be undone. Undo is its own inverse, so invoking undo twice in a row undoes the result of the undo command. Significant text modifying commands include all except insertion of individual characters.

If you give an argument to undo, it is taken as the number of changes to be undone. Each change undone appears as an independent change to the file. This allows you to "browse" through recent changes. If, for example, you had just deleted 10 lines from your file one at a time and typed M-5M-u, the last 5 lines deleted would reappear. If you decided you had instead wanted the last 7 lines to reappear, you could then type M-12M-u, undoing the 5 changes made by your first undo and then 7 more. There is

a limit on the total number of changes that can be undone that depends somewhat on the complexity of the change.

Almost all commands can be undone, but a few effects of commands cannot be undone. M-u will not undo the effects of reading or writing files, nor can it undo anything done by executing *UNIX* RTR operating system commands from EMACS. Undo cannot be used to bring back buffers that may have been inadvertently killed with the `^X^K` command (see Section 4.3.5). Undo can undo all other significant text modifications, with the exception that when the last significant text modification was a replace command, only a limited number of replacements can be undone. If more replacements were done with one command, undo prints a warning error message and if the user specifies then undoes what it can.

- `^X^C` Quit EMACS. If any buffers have been modified since the last write, EMACS will ask whether or not to write out each such buffer before exiting. EMACS will not ask whether or not to save an empty buffer.
- `BREAK` EMACS will not respond to your normal interrupt character (which is probably used for some EMACS command already), but it will respond if you hit the `BREAK` key on your terminal. Break causes EMACS to stop anything that it was doing at the next convenient point, and displays a message to you listing some options. When you hit break you can either continue what you were doing, abort what you were doing, exit from EMACS, or suspend what you were doing such that you can return to it later if you wish.

Many commands in EMACS can cause errors to be reported. When an error is reported, EMACS will display a message at the top of the screen followed by a line of underscores, ring the terminal bell, and wait for a response. In general, typing a space or 'y' will cause EMACS to continue trying to do whatever caused the error, typing 'n' will cause it to ignore the error, typing '^G' will cause it to abort the command causing the error, and '^Z' will cause EMACS to exit. The responses '^G' and '^Z' are always handled in the same way, while the interpretation of 'y' and 'n' depends on the error. Generally the message will describe what will be done with these responses for any error where it makes a difference.

### 4.3.3 SIMPLE CURSOR MOVEMENT COMMANDS

There are many ways to move the cursor around in the buffer without modifying the text in the buffer. Most of these use their argument to specify how many times the movement is to be repeated.

- `^F,^B` Move forward or backward one character. Recall that the end of each line counts as one character, so that `^F` at the end of one line will put you at the beginning of the next line.
- `^N,^P` Move to next or previous line. EMACS moves to the same character position in the line below (`^N`) or above (`^P`) the current line. Note that if the buffer contains tab or control characters, the same character position may display at different screen positions on different lines.
- `^A,^E` Move to the beginning (`^A`) or end (`^E`) of the current line of the buffer. Note that these work on one line of the buffer, not one line of the screen. If the current line is longer than will fit on one line of the screen display, these commands may move up or down on the screen to the real beginning or end of the line in the buffer.
- `M-<,M->` Move the cursor to the beginning or end of the buffer.

- M-f,M-b Move the cursor forward or backward by one word. In EMACS, words are delimited by non-alphabetic or non-numeric characters. Backspaces and underscores are considered legal characters in words.
- M-a,M-e Move the cursor to the beginning or end of the current sentence. The end of a sentence is defined as a punctuation mark followed by one or more whitespace characters (blanks or newlines). With an argument, these commands can be used to move forwards or backwards by a specified number of sentences.
- ^V,M-v Move to next or previous page. The cursor is moved forward or backward so that the display will show the text just before or just after the text now in the window on the screen.
- M-g Move the cursor to the line number specified by the argument given to the command.

#### 4.3.4 SIMPLE TEXT DELETING AND MOVING COMMANDS

Several commands are available to delete text from the buffer. All of these commands operate on text near the current cursor position. The deletion commands are:

- ^D,^H Delete forwards or backwards from the cursor. The ^H, or backspace deletes the character immediately before the cursor. The ^? or rubout, is a synonym for ^H. The ^D deletes the character on top of the cursor. If given arguments, these commands delete blocks of text forward or backward of the current cursor position.
- M-d,M-^H Delete words forwards or backwards from the cursor. These two commands delete words (as defined for M-f and M-b). If the current cursor position is in the middle of a word, M-d will delete from the cursor to the end of the word while M-^H will delete everything before the cursor. M-^? is a synonym for M-^H.
- ^K (Kill) Delete to the end of this line. If invoked without an argument, ^K deletes the remaining text on this line (if any). If no text follows the cursor on the current line, ^K deletes the newline. With an argument of 0, ^K deletes the text before the cursor on the current line. With an argument of n greater than zero, it deletes n lines forward from the cursor position. The text from the cursor up to and including the nth newline is deleted. With an argument less than zero, the deletion is backwards from the cursor position.
- M- (Meta space) The command Meta-space places an invisible mark on the current cursor position. This mark can be used in subsequent editing. Each mark is simply a position in the buffer (line number and character within the line). Thus if you add or delete text in front of a position where a mark was placed, the mark may not remain on the same character, but stays on the same position.
- EMACS actually maintains 16 different marks, normally allocated as one per buffer. (Thus if you set marks in different buffers, they are normally independent.) You can, however, alter this by specifying a mark number as an argument to Meta-space and other commands that work with marks. This allows you to mark up to 16 different positions in one buffer. The command ^@ (Control-@) is a synonym for Meta-space, but cannot be typed on all keyboards.
- ^W (Withdraw) The command '^W' deletes the text between the current cursor position and the mark. This is a convenient way to delete a well

defined block of text. If an argument is specified, it is used to select the mark number. The mark can be either before or after the cursor position and achieve the same effect.

- ^Y** (Yank back) Insert last killed text. All text that is deleted is saved in a "kill stack". The kill stack holds the last 16 deletions. There is also a limit on the total amount of text that can be held in the kill stack, but you are unlikely to encounter it. The ^Y retrieves the most recently deleted text. The most frequent use of this command is in moving text around. The procedure is: kill the text to be moved, move the cursor to where you want it, and enter ^Y. Another use of ^Y is to undo an unwanted deletion. The ^Y leaves the mark at the beginning of the inserted text, and puts the cursor at the end. The ^Y treats its argument (if any) as a count for the number of copies of the deleted text to bring back, and not a mark number. The ^Y operates only with the default mark.
- M-y** Replace last retrieved text. This command kills the text between the cursor and the mark and replaces it with the next to last item on the kill stack. This command can only be used immediately after ^Y, where it changes text that has just been retrieved. By entering ^Y followed by some number of M-y's, any text in the kill stack can be retrieved. If an argument is given, it is taken as the number of copies to retrieve.
- M-p** Pickup the region of text. This command picks up the text between the current position and the mark and puts it in the kill stack, without changing the buffer. This is useful for duplicating blocks of text in the buffer. An argument can be used to specify which mark to use.
- ^X^X** (Type control-X twice) Exchange the cursor position and the mark. An argument can be specified to indicate the mark to exchange with.

#### 4.3.5 SIMPLE FILE AND BUFFER COMMANDS

Commands that access files and buffers must ask for the name of the appropriate file or buffer. All of these commands (and some of the others) ask for the appropriate information at the bottom of the screen. You can use some of the simple editing commands described here, plus a couple of special commands, to edit a file or buffer name that you enter this way. The commands that you can use for editing are:

- ^F, ^B** Move forward or backward one character.
- ^A, ^E** Go to beginning or end of line.
- ^D, ^H** Delete forward or backward.
- ^U** Multiply the effect of the next command by 4.
- ^K** Kill (erase) the whole line.
- ^G** Abort the command asking for information.
- ^X** Enter the current line from the file at the cursor.
- ^Y** Enter the current file name at the cursor.
- ^L** Redisplay the prompt and the string being entered.
- ^T** Transpose the characters before and after the cursor.
- ^Q** Quote the following character (eliminates the special significance of the next character and just sticks it literally in the string being typed.)

**return** End the string being typed. The whole line is given to whatever command asked for it, even if the cursor is in the middle of the line when you hit return.

Most of the file accessing commands are invoked through the `^X` command. The `^X` is a prefix for several useful commands, most of which involve file or buffer access. These commands are invoked by a `^X` followed by a second character.

In the commands that ask for filenames, the normal shell conventions for partially specified names can be used. Any of the following sequences can be used in a filename and will be substituted properly.

**SVARIABLE**

Substitutes the value of the environment variable `$VARIABLE` or nothing if `VARIABLE` is not defined. Thus you can use path names like `$HOME/.profile`.

**\* and ?** These characters can be used to specify incomplete filenames and will be expanded. If more than one file matches the name given, then EMACS will pick only the first one.

**~USER** This translates into the home directory of the user `USER`. In addition, EMACS always translates the special name `~EMACS` into an EMACS library directory. This is a directory where special files needed by EMACS are stored and where standard macros are stored (in `~EMACS/macros`).

**'COMMAND'**

This causes `COMMAND` to be run and substitutes its standard output.

With these preliminaries out of the way, here are some of the commands that work with files and buffers.

**^X^R** Read file. EMACS will prompt for a file name, which you enter as described above. When the file name for `^X^R` has been entered, EMACS will read the specified file into the buffer. Hitting return in response to the prompt will cause EMACS to read from the current file (if any). EMACS will warn you if you are about to read over a buffer that has been modified since the last time you wrote it.

Many of the commands in this section, including `^X^R`, use their argument to specify minor variations on the basic action of the command, rather than specifying a count. In the normal case (with the default argument of one), `^X^R` clears the buffer before reading. If `^X^R` is invoked with an argument that is not 1 or -1 (i.e. `^U^X^R`) it does not clear out the buffer but instead, inserts the file into the buffer at the current cursor position. If `^X^R` is invoked with a negative argument, no error message is produced if the specified file cannot be read.

If the file contains a line that is too long for EMACS to handle, it will warn you with an error message. If you answer `^G` to the error, it will stop reading at that point. If you answer `n`, it will read the rest of the file, truncating lines that are too long, without further warnings. If you answer with `y` or a space, it will continue reading and warn you again if a subsequent line is too long.

**^X^W** Write file. Normally, if the specified file exists and has two or more links to it, EMACS will ask whether to overwrite the existing copy of the file or to unlink the specified file name and create a new file in its place, leaving the contents of the old file (which may be obtained through the other

names it was linked to) unchanged. If EMACS fails or the *UNIX* RTR operating system crashes during an attempted write (either `^X^W` or `^X^S`), the previous contents of the file are saved in a file `.EMACS` in your current working directory.

If you do not have write access to the file you are trying to write, EMACS will ask if it should try to write anyway. If you answer yes, EMACS will try to write the file by removing the old file and creating a new one. This is not a security hole, but few other *UNIX* RTR operating system programs will write files in this way.

Passing an argument to `^X^W` (i.e. `^U^X^W`) causes the contents of the current buffer to be appended to the specified file rather than replacing it.

`^X^S` Save buffer. This writes out the buffer to the last file read or written if the file has been modified. If the buffer was not read from a file, and has never been written to one, such that there is no file name associated with the buffer, EMACS will ask for a filename to save the buffer in.

`^X^B` Change buffer. EMACS allows up to 12 named buffers to be edited concurrently. Each buffer can hold a different file, and has its own current cursor position. In EMACS you work with one buffer at a time, although you can display two buffers on the screen at the same time. The `^X^B` command asks for the name of a buffer and makes that buffer the current buffer.

All of the commands that ask for buffer names accept either the text name of the buffer, or the buffer number (shown in parentheses after the editor name on the status line) for a buffer name. The number is convenient if you don't like typing long names. EMACS treats two buffer names specially. For any of the commands that ask for buffer name, if you enter an empty buffer name by just hitting return in response to the prompt, EMACS shows you a display of all of your currently defined buffers, indicating which one is current and which ones have been modified since they were last written. If you type space or 'y' in response to the "Continue" prompt that appears after the display, EMACS will abort whatever command asked for the buffer name and continue editing. If you type 'n' in response to the prompt, EMACS will ask again for the buffer name, and then complete whatever command asked for the buffer name. If you type a number in response to the prompt, EMACS will use the buffer with that number in whatever command asked for a buffer.

If you attempt to create more than 12 buffers, EMACS will indicate this and ask if you wish to kill one of your existing buffers that has not been modified. If you type 'y' or space, it will go ahead and re-use that buffer. If you type 'n' it will propose other buffers to be re-used until either you select one or there are no more, in which case you will get an error message and the command attempting to create the buffer will abort.

In all cases, if the buffer name `"..."` is entered, a new, empty buffer with a unique name is created.

`^X^F` Find file. This command prompts for a file name and switches to a buffer that holds the specified file. If the specified file has been read into a buffer, the effect of find file is to change to that buffer. If no buffer holds the specified file, the effect of find file is to create a new buffer and read

the specified file into it. Find file is a convenient way to switch between editing several files. If `^X^F` is invoked with a negative argument, no error message is produced if the specified file cannot be found. If it is given an argument `>1`, it will automatically read in a fresh copy of the file. If it is given a negative argument, it will not complain if the specified file doesn't exist, but will just create an empty buffer with that name.

`^X^K` Kill Buffer. This command prompts for a buffer name and destroys the specified buffer. You cannot kill the current buffer this way. Text in the buffer that is killed is lost and cannot be recovered.

#### 4.3.6 SIMPLE SEARCH AND REPLACE

EMACS provides several commands that search for text in the buffer, and also commands which allow you to specify global replacements, like change every instance of "football" to "baseball". The simplest forms of these commands are described here, along with a couple of miscellaneous commands that are useful. More complex versions of search and replace are described in Section 4.4.4.

`^S, ^R` Forward and Reverse Search. These commands allow you to look for text in your buffer. EMACS will prompt at the bottom of the screen with "Search" or "Reverse Search". In response to the prompt, you can type in characters, and EMACS will begin to look for the next match for what you type, going either forwards (`^S`) or backwards (`^R`) from the current position in the buffer. EMACS will show you the text that matches what you have typed as you type it, by moving the cursor to the text, possibly moving the display window in the buffer if the text you are looking for was not visible. As you type in the string to look for, you get immediate feedback about what EMACS has found. In addition to typing normal printing characters that become part of the string you are looking for, you can type some special characters to either edit the string you are looking for, or control the search in some other way.

`^H` This deletes the last character of the search string, and will cause the cursor to go back to whatever matches what is left.

escape Hitting escape stops the search, leaving the cursor on whatever you last found.

`^G` This quits from the search and goes back to the point in the buffer where you started the search from.

return or newline

These both cause a newline to become part of the search string. The newline is displayed as "`^J`" (which is the control character actually used by *UNIX* RTR operating system to indicate "newline") in the search string to allow you to see it. Thus if you type "`^Sthe~end<return>of`", EMACS will look for a spot where "the~end" appears at the end of one line and "of" appears at the start of the next line.

`^S` or `^R` These characters control the search. In general, if you are going forward, and type `^R`, or going backwards and type `^S`, the search changes direction, starting from the last thing you found. If you are going forward and type `^S` or backwards and type `^R`, the search proceeds to the next occurrence (in whatever direction you were going) of the search string. If you type `^S` or `^R` as the very first thing after starting a search, EMACS takes the *last* string that you successfully found with a search and makes it the

current search string. These characters provide a convenient way to navigate when looking for something that occurs many times in the buffer.

**^Q** Typing **^Q** "quotes" the next character, making it part of the search string, ignoring its usual significance. This is a way to look for strings that contain control characters.

other control characters

Typing any other control character causes the search to stop at whatever you found, and then executes the command corresponding to that control character.

This kind of search is called an incremental search in EMACS, because it shows you what you have matched incrementally as you type it. It is very easy to learn to use. For incremental search, the search string must *exactly* match whatever you are looking for. (There is a more complicated search available that allows some pattern matching, and is described in Section 4.4.4.) Incremental search stops, indicating that it fails if you reach the beginning or end of the buffer.

Normally, search considers uppercase and lowercase letters to be different, however you can override this with *caseless* mode.

**M-r** Query replace. You will be prompted for a *From* string and a *To* string. Each can be edited using the conventions described in the previous section for editing filenames. In general, Query replace will allow you to replace all of the strings in your buffer from the current cursor position to the end of the buffer that match the *From* string with the *To* string. In the *To* string, the **'&'** character can be used to designate replacement with the *From* string. To get a real **'&'**, prefix it with a **'**. To get a real **'**, prefix it with another **'**. If the single character **'%** is specified as the *From* string, the last string searched for (either from **^S** or **M-r**) is used. Note that you can specify a newline in the *To* string either by entering it explicitly as **"^Q^J"**, or by putting the string **"\n"** in the string where you want the newline to appear. EMACS then searches for the *From* string, positions the cursor in front of it, and prompts you. You can control the replacement of the item in question by what you type:

**<space>** or **y**

Replace this occurrence and move on to the next one.

**n** or **^?**

Skip this occurrence and move on to the next.

**.**

Replace this occurrence and exit query replace.

**<**

Go back to the last thing replaced and stop. (useful for correcting mistakes!)

**^G**

Quit. (Exit query replace without replacing the current match.)

**b**

Go back to the previous occurrence of the *To* string. It won't find one that you have already replaced!

**r**

Replace the rest without stopping to ask after each, and show the result of the replacement after each.

**R**

Replace the rest silently. (that is, don't show the result after each replacement.)

**<ESCAPE>** Causes EMACS to ask for a new string to replace the *To*

string with `with`. The current occurrence will be replaced with what you type, and EMACS will go on to the next occurrence.

Normally, query replace will show all occurrences of the search string. With an argument greater than 1 (that is, `^UM-r`), it behaves like the substitute command of the ed editor, looking at only the first match of the *From* string on each line. Query replace exits when the *From* string is no longer matched. `?` prints a summary of the options. As with incremental search, the *From* string must match exactly to something in the buffer. There is a more advanced form of query replace that allows pattern matching and is described in Section 4.4.4.

#### 4.4 ADVANCED EDITING COMMANDS

##### 4.4.1 GENERAL

The commands described in the previous section are sufficient to allow a user to perform most editing tasks efficiently. The commands in this section for the most part cover special situations, like inserting control characters into files, or provide more efficient ways to do things in certain situations.

##### 4.4.2 INSERTING ODD CHARACTERS

Because EMACS uses control and escape characters for commands, you cannot directly insert them into the buffer by typing them. The following three commands are useful for the occasional need to get such characters into a buffer.

- `^Q` Quote the next character(s). The `^Q` accepts one or more characters (the number of characters specified by its argument) from the terminal and inserts them "blindly" into the buffer without interpretation. Only the newline (line feed) character is interpreted. EMACS strips the parity bit from all characters read from the terminal, so all characters inserted this way have zero parity.
- `M-q` Quote characters and turn on parity bit. This acts just like `^Q`, however, it turns on the parity bit in the character before inserting. Characters inserted this way will be displayed as meta characters by EMACS.
- `M-\` Convert the argument to a character and insert into the buffer. This command takes its argument and converts it to a character and inserts it. This is occasionally useful for inserting odd characters for which the ASCII code is known.

##### 4.4.3 COMMANDS RELATED TO WINDOWS

EMACS provides a way to display two buffers on the screen at the same time. When this is done, the screen is split vertically, and one buffer is displayed in the top half and one in the bottom half. The status line will show status of the current buffer.

When EMACS displays two buffers like this, only the one that is the current buffer is actively updated. The display for the other just sits on the screen undisturbed until you return to that buffer. You can have the same buffer displayed in both windows, however, note that the current position is associated with a buffer, not a window, thus if you move the current position in the lower window, when you return to the upper window, the cursor will immediately move to wherever you were in the lower window. If you have different buffers in the two windows, the current positions in both buffers are independent, as they are with any two buffers.

- `^X2` Enter two window mode. EMACS will ask for a buffer to show in the

second window. The current buffer becomes the top window, while the buffer that you type in response to the prompt goes in the lower window and becomes the current buffer.

**^X^^** Grow window. Makes the current window grow by the number of lines specified by the argument to **^X^^**. You can use a negative argument to cause the current window to shrink.

You can also use this command to grow or shrink the buffer display with only one window on the screen. This can be useful in avoiding long delays when working from a low speed terminal port.

**^X1** Return to one window. The current window grows to fill the screen.

**^X^O** Switch windows. Make the dormant window current and the current window dormant.

#### 4.4.4 ADVANCED SEARCH AND REPLACE COMMANDS

The simple incremental versions of search and replace described above require that you match what you are looking for exactly. The commands described here allow pattern matching of regular expressions, like those used by the ed editor.

The description of regular expressions is too complex to reproduce here. Refer to the manual pages for ed(1) for a brief description. EMACS provides some additional special sequences for regular expressions. The character sequences "**\<**" and "**\>**" can be used to match the beginning and end of words. Thus the string **\<the\>** will match any occurrence of the word "the", but not any word containing the sequence of letters "the", such as "other". Note also that EMACS allows you to specify a newline as one of the characters to be matched anywhere in the expression. Specify a newline as "**\n**" either by itself or as an alternative in a set of characters (that is, **[a-z\n]**). Neither the special symbol "**.**" nor a negated set of character (that is, **[^a-z]**) will match a newline to avoid unexpected results of expressions such as **".\*"**.

In addition, a set of alternatives can be specified by expressions of the form: **\(exp1\|exp2\|exp3\)**. This expression matches the text matched by **exp1**, **exp2**, or **exp3**, and saves it for later reference as with other expressions enclosed in parentheses. Any number of alternatives may be specified. At present, **\***, **+**, and **\{n,m\}** can not be used after a set of alternatives to specify multiple matches. (Note that you can specify multiple copies of text matching one of the alternatives with an expression of the form: **\(exp1\|exp2\|exp3\)\1\***).

In constructing regular expressions, it is important to remember that in order to avoid the special significance of a character like **'.'** or **'\*'**, you must prefix it with a backslash **'\'**. If you must have control characters in regular expressions, you can quote them for EMACS by typing **^Q** before the control character.

Constructing regular expressions is tricky, and constructing ones that will match efficiently can be very difficult. Some things to keep in mind are that expressions beginning with a normal text character are generally matched much more efficiently than any other sort. An expression such as **[a-c]** will be matched much more efficiently than the equivalent form **\(a\|b\|c\)**. Expressions matching an indefinite number of alternatives by using **"\*"** or **"+"** are probably the slowest to execute, particularly when the expression is likely to match a large number of characters and is not the last expression being matched (that is, **.\*x**). You should be especially careful of expressions that match an indefinite number of characters including newlines (for example, **[a-z\n]\***), as these can trigger a very long search that extends over many lines.

M-^S Regular Expression Search. This will prompt for a regular expression to search for. You can edit the expression like editing filenames. Hitting return in response to the prompt will search for the last thing you searched for with M-^S, ^S, or ^R.

You can search forward or backward, either ending at the beginning or end of buffer, or wrapping around (like ed) depending on the argument given to the search command.

- 1 (default) Search forward, wrapping from the end of the buffer to the beginning, and failing only if the buffer contains no match for the given string.
- 1 Search backwards, wrapping around from the beginning to end of the buffer.
- > 1 Search forwards, stopping at the end of the buffer.
- < -1 Search backwards, stopping at the beginning of the buffer.
- 0 Just search the current line. The search will find the indicated expression anywhere on the current line and fail otherwise.

In all cases, you can have EMACS repeat the search, looking for the next (or previous) occurrence of the search string by typing ^S or ^R immediately after the regular expression search.

M-^R Regular expression query replace. This is just like query replace, except that a regular expression is allowed in the *From* string. You may also use the special character sequence \<digit> in the *To* string, to specify that the characters matched by the nth subexpression (delimited by nd \)) are to be used in the replacement string.

#### 4.4.5 MACROS, KEYBOARD MACROS, AND INPUT FILES

EMACS provides a number of ways for a user to construct editor programs from sequences of commands. The macro programming facility is the best way to construct substantial programs. The commands listed here deal with two other ways of saving a sequence of EMACS commands for later use, input files and keyboard macros, and with the commands that load and invoke "full" macros into EMACS for your use.

^X^I Redirect input. This command directs EMACS to take input from a file. The file is assumed to contain EMACS commands, and can be created by editing with EMACS, using ^Q to enter control and escape characters. You can also create an input file by using the commands to create keyboard macros described below, and then saving the resulting keyboard macro file.

This command can be used to perform a series of commands on the current buffer, or to set up a standard set of initializations. Thus the file should contain *exactly* what you would type from the keyboard to perform whatever task you wish to perform. Note that if the file contains only printable ASCII text, tabs, and newlines, ^X^I will effectively read the file into the buffer at the current location. Note, however, that this is very slow, and much better done with ^X^R.

A file suitable for executing with `^X^I` is known as a keyboard macro, because it is interpreted just as if it had been typed from the keyboard. The following commands provide a sensible way to create and execute keyboard macros.

- `^X(` Begin Keyboard Macro. This command starts remembering the keystrokes you enter so that they can later be executed as a keyboard macro.
- `^X)` End Keyboard Macro. This command stops remembering keystrokes for a keyboard macro.
- `^XE` Execute Keyboard Macro. This command retrieves and executes the keystrokes typed between `^X(` and `^X)`. EMACS executes them just like they came from your terminal. Keyboard macros are saved in the file `.emacs_kbd` in your home directory. These are saved between sessions, so that `^XE` is in fact the same as invoking `^X^I` (Input file) and giving `$HOME/.emacs_kbd` as the file name to execute. Note again that you can use `^X(` and `^X)` to create a keyboard macro, save it for later use by moving the file `$HOME/.emacs_kbd` to another file, and then invoke it by invoking `^X^I` with the name of the file you saved.
- `^Xd` Define macros. This command treats the current buffer as definitions of new macro commands. The commands are defined and become available for use. For a complete description, consult the macro programming manual. Note that you should not use `^Xd` with the file created from a keyboard macro.
- `^X^L` Load macros. This command allows you to load "full" macro definitions from a file. It is described in the macro programming manual, however even if you do not program your own macros, you may be interested in using those defined by others and will use `^X^L` to load the resulting files. On some systems on which EMACS is installed, there is a library directory of macros available for general use in `~EMACS/macros`, and the file `~EMACS/macros/CATALOG` gives a catalog to the available macros.
- `M-x` Execute Macro command. This command asks for a macro name and tries to execute it. If there is no macro currently loaded with the name you give and if *autoload* mode is on, EMACS will try to load the file with the name of the macro from the directory `$EMACS_LIB` (if this environment variable is defined) or in the directory `~EMACS/macros` if it hasn't been found yet. If the macro cannot be found, an error results.
- `^Z` Exit level. In an EMACS editing session, you may wind up in a nested level of EMACS. This can happen either by typing "break", and responding "y" to suspend whatever you were doing and invoke a new command interpreter, or by invoking a macro that uses the recursive edit command to allow you to edit something from inside of a macro. The `^Z` exits your current level of EMACS, returning to whatever called it.

The `^Z` command performs a function similar to `^X^C`. Ordinarily, both will exit from EMACS, however if you have hit break or type `^Z` while executing a macro that allows you to edit the buffer and return to the macro, `^Z` will return to the break or the macro instead of exiting.

#### 4.4.6 COMMANDS THAT ESCAPE TO THE *UNIX* RTR OPERATING SYSTEM

There are several commands that interact with *UNIX* RTR operating system, allowing you to run *UNIX* RTR operating system commands or send mail from inside of EMACS.

- M-!,M-\$ *UNIX* RTR operating system escape. These two commands implement five different ways to run *UNIX* RTR operating system commands from EMACS. In all cases, the commands prompt for the name of a *UNIX* RTR operating system command to be run and run it. The command is run through your normal shell as indicated by \$SHELL. If you enter the special command name "sh", it runs your normal shell instead of "sh". To facilitate writing programs that interact with EMACS, the environment variable "filename" is set to the name of the current file in EMACS when the command is run. The following summarizes the various flavors:
- M-!            Run the command, suspending EMACS while it runs.
- ^UM-!         Run the command and feed it the contents of your current buffer as standard input. When the buffer is exhausted, the command will see an end of file.
- M-\$            Run the command with standard input from your terminal but standard output and standard error are captured in the buffer ".exec", which is created if it doesn't already exist. (If it does exist, the command output replaces its current contents). Normally, output from the command is also displayed on the terminal as it is produced, but this can be overridden via the *usilent* or *noecho* modes described in Section 4.5.3. This is useful for saving a copy of the error messages produced by a C compilation of a file being edited, for example. The file name of the .exec buffer is set to the command line that produced it. This can be useful if you want to re-execute the same command, as you can make .exec your current buffer, enter M-\$, and enter ^Y followed by newline as the command line. ^Y gets the old command line back, and newline will execute it.
- ^UM-\$         Run the command and append the output to the buffer ".exec". This is just like the above except that the .exec buffer is not cleared and the output from the command is appended to it. It also inserts the command line into the .exec buffer when the command is run.
- ^X^D         Change Working Directory. This command directs EMACS to change the current working directory. Note that if you have buffers with filenames that are relative pathnames (not starting with '/'), change the working directory, and then save one of these buffers, the buffer will save into a different file because of the change in working directory.
- M-^M         Mail. This command takes the current buffer as *UNIX* RTR operating system mail, and sends it. The buffer must contain at least one line starting To: , which specifies the recipients of the mail. Each recipient is delimited by a space. Any number of recipients may be listed in a single line, however to improve readability, additional To: or Cc: lines may be used in specifying lists of recipients. Any errors encountered by mail are printed. If the environment variable \$MAILER is set, then it is taken as the name of the command to run to send the mail. Otherwise, EMACS runs "mail".

#### 4.4.7 MISCELLANEOUS COMMANDS

The remaining commands handle special situations that occur once in a while.

- ^O            Open up a line. This command creates one or more empty lines at the

- current cursor position. This is useful for inserting text in the middle of the buffer, while minimizing the amount of screen refresh needed.
- ^T** Transpose the next two characters. The cursor moves forward one character for each transposition, such that giving **^T** a count as an argument causes the character at the cursor to be dragged forward through the text.
- ^X^T** Transmit text to another buffer. This command sends the text between the mark and the current cursor position in the current buffer to another buffer. EMACS prompts for the name of the other buffer, and the text is inserted into that buffer at the current cursor position for that buffer. The current buffer remains unchanged. If an argument is given, it selects the mark to use. If the target buffer has a sub-process running under it, then the region is also sent to that process and is always put at the end of the buffer. (See the description of **M-\$** for more information.)
- M-s** This command displays some statistics about your editing session, such as how many characters EMACS has sent to you and how many characters you have typed. The information is normally not of much interest.
- ^X=** Status. Displays status information (current line, number of lines in buffer, current character position, number of characters in buffer, etc.).
- M-/** Begin comment. This command begins a C program comment by moving to the appropriate column (specified by *comcol* mode) and putting a */\** in the buffer. (If the cursor is in the first column, the comment is not indented). The next newline will close the comment and automatically append a *\*/*. If fill mode is on and the comment line is wrapped onto the next line as a result, the next line will begin with " \* ", and the comment will stay open until you close it with a newline.
- M-\_** Underline word. This command underlines the following word of text using backspaces and underscores. This can be used to generate underlined text that will be displayed properly by most printers and formatting software.
- ^C** Capitalize. This command capitalizes the letter under the cursor and moves the cursor forward one position. Lowercase alphabetic characters are converted to uppercase, while other letters are unchanged.
- M-c** Capitalize word. The letter under the cursor is capitalized, and the cursor is moved to the beginning of the next word.
- M-l** Lowercase letter. The letter under the cursor is converted to lowercase and the cursor is moved to the right.
- M-~** Unmodify Buffer. This command causes a buffer to be marked as being unmodified even if it has been modified since the last write. Doing this will avoid having EMACS ask whether or not to write the buffer when you exit if you know that you do not want to rewrite the buffer. With an argument greater than 1 (that is, **^UM-~**) this command marks the buffer as modified.
- M-^L** Redisplay top. Redisplay the window with the current line at the top. This is useful for viewing the lines that follow the current line. Note that this does not re-create the entire display, as does **^L**, so it will not necessarily clear up a garbled screen.
- M-"** Auto Fill Buffer. This command re-adjusts the lines in the buffer so that each line contains 72 or fewer characters. The adjustment is done by

moving words from one line to another. The nroff or mm command lines and blank lines are preserved as is. The adjustments made generally do not change how the file would look after being run through nroff/troff unless the text contains tables, displays, or other fixed format text. Lines are broken at whitespace (space or tab) characters.

If you give M-” an argument other than 1, it readjusts only that part of the buffer between the cursor and mark.

**M-:** Remap character. The command M-: allows you to remap character commands. It prompts for a character sequence (a single character or a meta or ^X sequence) and a command (also a specified by a character sequence) to put on that character. This allows you to reconfigure EMACS to your liking. You can remap any character you like, including characters like ^U, ^X, or escape, however, to re-map escape or ^X or to map other characters to these "commands", you must invoke M-: with an argument of -1 (M--M-:). With a negative argument, M-: will ask for only single characters for the sequences to map. The command sequence is always interpreted with the default bindings (as documented in this memo), and not with the bindings set up with earlier M-: commands. Thus M-:^A^B followed by M-:^B^A will swap the ^A and ^B commands, since the command ^A in the second M-: command refers to the default meaning of ^A, not that established by the first M-: command. M-: also changes the behavior of the control characters used to edit filenames and other string parameters, but does not change the behavior of some of the characters that have special meaning in response to prompts issued by various commands, such as the responses to query replace.

With an argument of 4 (^UM-:), this command asks for a character and a macro name to assign to that character. This allows you to bind macro commands to characters by their names instead of their current character binding.

With an argument of 0 (M-0M-:), this command resets all of the keyboard character bindings to their default values. This may be useful in recovering from trouble.

With an argument of -1 (M--M-:), the character sequence and command are single characters, allowing escape and ^X to be remapped to other commands, and allowing other single characters to assume their function. Note that ^X and escape "commands" can only be attached to single characters.

In all cases, the bindings established with this command and through defining macros apply to both character commands typed from the keyboard and to characters in keyboard macros and initialization files. They do not apply to the character commands executed in the body of "full" macros.

See Section 4.4.5 and the macro command manual for further information on bindings.

**^X^M** Set Mode. This command can be used to set parameters to customize the behavior of EMACS. For more information on modes, see Section 4.5.

## 4.5 MODES

### 4.5.1 GENERAL

EMACS has a variety of parameters that can be changed from commands entered in the terminal. These are referred to as "modes" in this document and in the messages printed by EMACS. There are two types of modes: on/off modes and integer modes. The `'^X^M'` command can be used to display or set these parameters. The `'^X^M'` will prompt for the name of the mode to set. If you enter a return in response to the prompt, the current mode settings are displayed. Normally, EMACS will display the value of each integer mode, and the name of each on/off mode that is currently on. If you enter *return* in response to `'^U^X^M'`, EMACS will also display the names of the on/off that are currently off, indicating for each mode whether it is on or off.

Modes are set by giving the name of the mode to set in response to `'^X^M'`. If an on/off mode is given, it is turned on if no argument is given to `'^X^M'`, and turns it off if an argument other than 1 is specified. (Thus `'^X^M'` turns on, `'^U^X^M'` turns off). For an integer mode, the mode is set to the value of the argument.

The modes and their types are listed in the following sections, along with their default values. For ON/OFF modes, the default is **highlighted**. The modes are grouped into 4 broad categories:

- Display modes change how the information in the buffer is displayed, but have no effect on its content.
- Interface modes change the command interface to EMACS in minor ways, but don't really change any of the behavior of the commands.
- Command modes change the way in which some of the commands work.
- Terminal modes change the way that EMACS uses the terminal, and exist mainly to get around special problems caused by certain kinds of terminals.

### 4.5.2 DISPLAY MODES (PARAMETERS)

<i>lnumb</i>	Line Number Mode ( <b>ON/off</b> ) This mode causes the current line number to display at the left of each line.
<i>lnowid</i>	line number width (INTEGER=4) This parameter specifies how many character positions are reserved for the line number when in line number mode.
<i>height</i>	Display height (INTEGER=<screen_size-4>) This parameter dictates how many lines from the buffer will be displayed on the screen. It is automatically set based on the terminal type whenever the terminal type is set, and is changed by the one window and two window commands. This mode and <i>width</i> mode can be set explicitly to restrict the display to a subset of the entire terminal screen, or can be used to allow you to use a terminal with a settable screen size for which EMACS does not have the right size built in.
<i>width</i>	Screen Width (INTEGER = <set based on terminal type>) This mode specifies the width of the display screen.
<i>tabstop</i>	Tabstop interval (INTEGER=8) This mode is the number of characters per tab that are displayed. A deeply indented C program may be more readable if <i>tabstop</i> is set to something smaller than the default value of 8.

- backspace* Display of backspaces (on/**OFF**).  
Turning on *backspace* mode causes backspace characters (^H) to display as moving back one column rather than as a ^H. This is very useful for viewing nroff output or manual pages, but editing the resulting text can be a bit tricky, because it is impossible to tell whether the character under the cursor is the one being displayed, one that has been overprinted, or a backspace. *Backspace* mode is set on automatically if your terminal can display underlined text.
- time* Display time (on/**OFF**)  
When *time* mode is on, EMACS will display the time of day below the mode line (the one that says EMACS and the buffer name). The time is updated every time a character is read. Using *time* mode when entering lots of text is expensive in processor cycles.
- display\_percent* Display cursor as percent of buffer (on/**OFF**)  
If set, EMACS will display the percentage of the current buffer beyond the current cursor position on the mode line.
- 7bit\_ascii* Display only 7-bit characters (on/**OFF**). This mode controls how characters with the high order bit set are displayed. With this mode off, they are displayed as "M-" followed by the character. With this mode ON, they are displayed as highlighted (underlined) characters. This mode is most useful for editing files used with personal computer word processing systems which use the high order bit for formatting control.
- leftmargin* Leftmost displayed column (INT=0). In picture mode (See Section 4.5.4), EMACS automatically scrolls the window left or right to keep the cursor on the screen. This mode allows you to alter this behavior. The value of *leftmargin* is the left most displayed column. Setting it will cause the screen to scroll left or right. In all cases, if the cursor wanders out of the window, EMACS will pick its own *leftmargin* to keep it on the screen.

### 4.5.3 INTERFACE MODES

- save* Automatic buffer saving (on/**OFF**)  
If *save* mode is on, EMACS will automatically write the current buffer after *savetype* characters have been entered since the last save. EMACS will also automatically save the current buffer before any command that changes buffers or runs a *UNIX* RTR operating system command, such as ^X^B, ^X^F, ^X^O, ^X^D, M-!, and M-\$. This mode reduces the chance of disaster in the event of a crash, but may delay editing by causing a lot of extra writing to the file.
- savetype* Save type ahead (INTEGER=256)  
If *save* mode is on, this is the number of keystrokes between saves.
- verbose* Verbose prompting (**ON**/off)  
When *verbose* mode is on, EMACS will prompt for more input when ^X, ^Q, escape, or ^U are entered. This makes it easier to keep track of where you are. *Verbose* mode also effects some error messages. In general turning *verbose* mode off causes some error conditions (such as using one of the commands that uses the mark when the mark has not been set) to be handled by supplying a default answer, rather than printing an error message.

- keepsroll* Lines kept when paging (INTEGER=0)  
This parameter specifies how many lines are to be preserved on the screen when forward page or backward page is invoked.
- smoothscroll* Smoothscrolling (ON/off)  
When set on, EMACS will try to scroll text onto the screen smoothly, one line at a time when moving forwards or backwards by less than one screen full or when inserting or deleting lines. With this mode off, any insertion, deletion, or screen movement is done all at once. The display will be somewhat slower with smoothscroll on, but may be easier to read when scrolling forward through text.
- caseless* Ignore case in searches (on/OFF)  
This mode causes either case characters in the search string to match either case in the buffer on all searches and query replace.
- mailtype* Check mail interval (INT=100)  
This parameter determines the number of input characters between checks of your mailbox (\$MAIL). When EMACS discovers something in your mailbox, a warning is displayed at the bottom of the screen.
- end\_newline* Behavior of ^N at end of buffer (on/OFF)  
This parameter determines whether executing the ^N command in the last line of the buffer adds a new line to the buffer (mode ON) or whether it signals an error (mode OFF).
- usilent* Silent UNIX RTR operating system commands (on/OFF)  
If set, this parameter causes EMACS not to display the command name or output for M-\$. This is useful for invoking a UNIX RTR operating system command in a macro and capturing the output without disturbing the display.
- noecho* Don't echo M-\$ output (on/OFF)  
If set, EMACS will not echo the output of commands run with M-\$ to the terminal. The difference between *noecho* and *usilent* is that *usilent* isolates commands run from EMACS from the terminal completely, so that EMACS does not clear the screen and does not have to redraw it when the command exits. *noecho* has no effect at all on commands run via M-!, but does eliminate the display of output from M-\$. It is most useful when you wish to run something that will produce lots of output, capturing the output in .exec.
- eofnl* Write newline at end of file. (ON/off)  
This mode when on causes EMACS to append a newline character to any file written by EMACS from a buffer that does not end in a newline. EMACS allows you to create files that do not end in newline. Unfortunately, many UNIX RTR operating system tools get confused when reading such a file. This mode is on by default, and prevents you from writing a file that will cause troubles. For editing files which you do not want to end in a newline, turn this mode off.
- savelink* Preserve links on write (on/OFF)  
Turning this mode on will cause EMACS to automatically write

into the existing file when writing to a file with multiple links, instead of asking the user what to do in this situation.

*search\_newline*

Newline ends search (on/OFF)

Turning this mode on causes incremental search to stop if a newline or carriage return is typed from the terminal.

*autoload*

Automatic macro loading (ON/off)

This mode controls the automatic loading of macro packages when an undefined macro is called. If *autoload* mode is on and an undefined macro by the name of <name> is called, EMACS will first attempt to load \$EMACS\_LIB/<name> (resolving the environment variable \$EMACS\_LIB), and then try to load ~EMACS/macros/<name>. If either of these files exists and defines a macro called <name>, that macro will be called and execution will continue. Otherwise, an error message is produced. This allows macros to be loaded incrementally, only when needed. With the mode off, any attempt to call an undefined macro results in an error.

#### 4.5.4 COMMAND MODES

*fill*

Automatic line filling (ON/off)

If this mode is on, EMACS will automatically move to the next line whenever the cursor moves to the end of the line, breaking the line at a word boundary (programmable by changing the character tables). This is very useful for entering text, as no newlines need be entered in the middle of the text.

*fillcol*

Right margin for *fill* mode (INTEGER=72)

This is the character position beyond which *fill* mode will cause the line to be broken.

*c*

C source indentation (INTEGER=0)

This mode controls automatic indentation. When set to 0 or 4, indentation is off and each new line will be started at the left-hand edge. When this mode is set to 1, each line will be indented with the same number of tabs as the one before it, adjusted for the number of opening and closing braces in the previous line. If, in addition, the file being edited has a name that ends in ".h", or ".c" (conventionally indicating a C program), the indentation of each line is adjusted for labels and case declarations (which are indented one level less) and preprocessor statements (which always start at the beginning of a line). The indentation of a line may be adjusted whenever any of the characters ':", '#', or '}' are typed to make it conform to the normal rules for C program indentation, allowing you to enter C source without adjusting indentation. The special behaviors for C source can be forced to occur in other files (not ending in .c or .h) by setting C mode to 2.

The C mode is useful both for text, where it can be used to maintain indentation in indented paragraphs, and source, where it allows you to enter text without indenting it and usually get a properly indented result.

*comcol*

Comment Column (INTEGER=40)

This is the column in which comments entered via M-/ begin.

*rigid\_newline*

Rigidly insert newlines (on/OFF)  
This mode causes any newline to insert a newline into the file. With this mode off (the default), a newline will not insert anything if the following line is empty, but will simply move to the next line.

*readonly*

Read only buffer (on/OFF)  
If this mode is turned on, saving the current buffer is disabled. The ^X^S will complain with an error message; autosaving will not take place, and EMACS will not complain if you try to exit without writing buffers.

*picture*

Tailor editing for two dimensional displays (on/OFF)  
If set, EMACS treats the buffer as an "electronic blackboard", rather than the exact contents of a *UNIX* RTR operating system file. The display shows a rectangular region of the blackboard through a window. Characters beyond the right margin are not displayed, but are indicated by a '!' in the right margin. The window is moved left or right to keep the cursor in the window. The *leftmargin* mode gives you some explicit control over the window, but in general it is self adjusting. If the window is not at the left edge of the blackboard, then the character offset of the left edge of the display is given on the mode line, in front of the editor name.

Picture mode changes the behavior of several basic commands to be more suitable for two dimensional editing.

- ^N/^P      These move to the same column in the target line, extending the target line with spaces if necessary.
- ^F/^B      These will not move off the current line. The ^B will stick at the left margin, while ^F will continue to extend the line to the right if necessary.
- ^W/^Y      These treat the region to be deleted as a rectangle, bounded by the mark and the cursor position at the corners. All text in the rectangle defined by these positions is killed by ^W. The ^Y brings text back in the same way. All text deleting commands behave like ^W, in that the start and end positions of the text region to be deleted are taken as corners of a rectangle. When the text region being deleted is all on one line, behavior is identical to normal EMACS, however, deletion across line boundaries via commands like M-d may not do anything sensible.

Picture mode is probably most useful with *nodelete* mode, *notabs* mode, and *overwrite* mode all set. You can use picture mode without the others set, but the presence of tabs or control characters in the buffer may give unexpected results when you navigate in the file. This mode is particularly useful for editing fixed format tables and picture images of "typewriter art".

*overwrite*

Overwrite instead of insert (on/OFF)  
In *overwrite* mode, text entered will overwrite text already there. Text entered when the cursor is at the end of a line will be inserted as before. *Overwrite* mode may be more natural for some people when making corrections.

- nodelete* Don't close deletions (on/OFF)  
With this mode set, text deleted in the file is not closed up but instead is overwritten with spaces. After any deletion, the cursor is moved to the first character of the region deleted. *nodelete* mode should probably be part of *overwrite* mode, however *overwrite* mode considerably pre-dated it and was left alone for upward compatibility.
- notabs* Eliminate tabs (on/OFF)  
With this mode set, EMACS does not display tab characters in the buffer as whitespace, but instead shows them as ^I (control-i), which is the ASCII code for a tab. When you enter a tab by typing ^I or tab, EMACS converts it to the proper number of spaces to come up to the next tabstop column. EMACS does not convert tabs already in the file, though there are macros that do this.
- The main use of this mode is in creating and editing fixed format information in *picture* mode, though it can also be useful in showing you where tabs are in your buffer.

#### 4.5.5 TERMINAL MODES

- nobell* No Bell (on/OFF)  
Ordinarily, unexpected conditions, such as errors or quitting out of commands, cause the terminal bell to ring. Turning on *nobell* mode prevents EMACS from ringing the terminal bell.
- tspeed* Terminal Speed (INT=<terminal dependent>)  
This parameter is the speed of your terminal in milliseconds per character. This parameter is set whenever you enter EMACS and is used in determining how to update the display most efficiently.
- controlify* Controlify mode (on/OFF)  
When on, this mode causes a particular character (normally ^^, but settable to another character with the *ctl\_char* described below), to act as a prefix specifying that the next character is to be interpreted as a control character. This mapping takes place at any time, not just when entering commands. Thus the sequence ^X^^s will be mapped into ^X^S, and cause the save command to be invoked. The sequence ^^q^^m typed in response to a request for a file name will be mapped into ^Q^M, which will be interpreted as asking for a file name of ^M. *Controlify* mode is primarily intended to allow you to enter characters which cannot be sent transparently from your terminal to EMACS. The most frequent examples are ^S and ^Q, which cannot be sent by some terminals and some local area networks. By setting *controlify* mode, you can send these characters with ^^s and ^^q.
- ctl\_char* Prefix character for *controlify* (INT=30)  
This parameter sets the value of the character used when in *controlify* mode to indicate that the next keystroke should be interpreted as a control character. It is the ASCII value of the character to be used.
- flow\_lim* Flow Control Limit (INT=0)  
This parameter enables the use of xon/xoff flow control during output to control the rate of output to the terminal. EMACS normally turns xon/xoff flow control off to allow ^S and ^Q to be passed to EMACS to be used in specifying commands. If *flow\_lim* is non-zero, then whenever more than *flow\_lim* characters are sent

to the terminal at once, xon/xoff flow control will be temporarily enabled to allow the terminal to send ^S to request that the *UNIX* RTR operating system stop output. The xon/xoff flow control is disabled when output is complete. The only effect you will notice is that you cannot type ahead while EMACS is updating the screen. Normally, EMACS supplies sufficient padding to allow it to run without xon/xoff flow control, however for some terminals at high speeds, xon/xoff flow control is required. In this case only, set *flow\_lim* to a number somewhat larger than the terminal's character buffer (typically 32, 64, or 128). In some cases it may be desirable to set up the terminal to always use ^S/^Q for flow control. This can be done by setting *flow\_lim* to -1 (escape '-' ^X^M *flow\_lim*). If you do this, you will not be able to type ^S or ^Q from the keyboard and must rebind them or use *controlify* mode to send them.

*no\_break* Suppress Breaks (on/OFF)  
Setting this mode disables the special handling of the break key to interrupt commands in progress. It may be useful for noisy lines that tend to generate breaks.

#### 4.5.6 SPECIFYING DEFAULT MODES FOR A FILE

You can specify the modes to be used while editing a particular file by putting the string "EMACS\_MODES:" somewhere in the first 10 lines of the file. The text on the same line following EMACS\_MODES: will be taken as names of modes to set on or off. A mode name preceded by '!' will be set off, while mode names just listed will be set on. If '=' immediately follows a mode name, then the characters immediately following the '=' will be taken as the value for the mode name. Any text on the line that does not correspond to a mode name will be ignored. Thus the line:

```
/* EMACS_MODES: !fill c, comcol=43 */
```

in a *c* source file will set *c* mode on, *fill* mode off, and set the column for starting comments to 43. These modes are set whenever the file is read, and whenever you switch buffers.

## 4.6 GETTING STARTED

### 4.6.1 GENERAL

When EMACS is invoked, it does a number of things to set up for your editing session. These include finding out the type of your terminal, processing an initialization file that allows you to customize EMACS for your needs, and processing command line arguments.

### 4.6.2 TERMINAL TYPE

The *5ESS*<sup>®</sup>-2000 switch EMACS only is supported for VT100\* compatible terminals. The VT100 compatible terminals include the Lucent 5425 and 4425 as well terminal. Baud rate should not exceed 4800~baud and slow-scroll mode must be turned off.

### 4.6.3 INITIALIZATION FILE

When you start EMACS, it consults an initialization file that allows you to initialize the various modes the way that you want them. This file is called ".emacs\_init" and should be put in your home directory. The file is read as a keyboard macro, and thus should contain *exactly* what you would type from the terminal in order to set it up. Thus the file:

```
^U^X^Mfill
^X^Mc
```

sets *c* mode and unsets *fill* mode. If you do not have a `.emacs_init` file, EMACS will run the standard initialization file for your system, which you can examine by looking at the file `~EMACS/.emacs_init`. It will not run both by default, but you can call for the default initialization file to be run from your private file by placing `^X^I~EMACS/.emacs_init` in your initialization file at the point at which you want the system initialization to be done. You can also use command line arguments to specify alternate initialization files as described in Section 4.6.4.

#### 4.6.4 COMMAND LINE ARGUMENTS

EMACS accepts a number of arguments on the command line that effect processing. These arguments are processed in order until an argument that is not recognized as a command line option is found. The first such argument is treated as the file to read, and any subsequent arguments can be accessed by the user using the `^X^A` command. This allows arguments following the file name to be treated in any way desired, specifying more options or filenames as interpreted by the user. The supported options are:

- i (init file) EMACS interprets the next argument as the name of an initialization file to run in addition to the standard initialization file (Yours or the system's). The specified file is run *after* `.emacs_init`.
- .i (init file) EMACS interprets the next argument as the name of an initialization file to run *instead* of the standard file. This option must be the first argument to have this effect.
- +n EMACS moves to line *n* of the specified file after it is read in. This argument must be the last of the options, immediately prior to the filename.
- <filename> The last argument on the command line should be the name of the file to be edited. This file (if present) is read into the buffer Main. If no filename is specified, EMACS puts you in the buffer Main with no associated file.

#### 4.6.5 HELPFUL HINTS

This editor is very easy to use, once you know a few of the basic commands. Here are some tips for making the best use of EMACS.

1. Learn a few of the basic commands at a time. You can accomplish a lot with just the basic commands.
2. EMACS tries to be reasonably efficient about the refreshing of the screen. Some sequences, however, will cause lots of text to be redisplayed. While you can insert anything into the middle of a buffer by typing it, if you intend to type a lot of text, it is frequently better to open up some blank space using `^O`, type your changes, and then kill any unneeded blank lines.
3. Use `M-?` when in doubt about a command. The explanations are brief, but should be sufficient to tell you what you want to know.
4. Like most editors, EMACS maintains a local buffer, so that changes made do not go into the file until the next write. Type `^X^S` reasonably frequently so as to avoid being wiped out by machine crashes, editor bugs, or other unpredictable events. You can set *save* mode to do this automatically.

5. Because EMACS tries to avoid unnecessary refreshing of the screen, it will get confused if characters are sent to your terminal from some other program while running EMACS. This can happen with output from a "background" program, or with text from a write command from another user. If you suspect that the display does not correspond to the buffer that you are editing, type ^L to refresh the screen. After typing ^L, the screen will match the buffer being edited.

#### 4.6.6 LIMITATIONS OF THE EDITOR

There are some limits that you may encounter:

- The maximum length of a line EMACS can handle is 511 characters. If you edit something that has longer lines, or create longer lines in the process of editing, EMACS will truncate at 511 characters.
- You can have at most 12 buffers.
- The kill stack contains the 16 most recent deletions, or a total of 256K characters.
- Filenames are limited to 128 characters.
- The undo command will undo only about 10 distinct replacements done in one replace command.

#### 4.6.7 RECOVERING FROM VARIOUS PROBLEMS

Because EMACS puts your terminal in "raw" mode, it does not respond to interrupt and quit characters the way that most *UNIX* RTR operating system programs do. If something goes wrong with EMACS, you can usually stop it with BREAK (typing ^Z to exit from EMACS in response to the message). You may have to kill it from another terminal or log out.

If EMACS runs into internal trouble, or if you kill EMACS or log out while running EMACS, it will try to save your buffers before terminating. The buffers are saved in the files emacs0-emacs11 in your home directory. You will get mail describing what files are there from EMACS after this happens, though if you do not clean up these files, they will continue to appear in the message each time you hang up on or kill an EMACS process.

#### 4.7 CONCLUSIONS

The EMACS editor provides an effective means of using a high-speed display terminal for text editing. EMACS provides a variety of unique features that are very useful in editing. User reaction indicates that EMACS has improved their productivity, however there are no quantitative measurements of this effect.

EMACS uses more computing resources than the standard *UNIX* RTR operating system editor, however the resources utilized do not appear to be a serious problem.

```

1  #include <stdio.h>
2  /* EMACS_MODES: c, !fill, comcol=43 */
3
4
5  /* This is a c program */
6
7  main()
8  {
9      int i;
10     char c;
11
12     for (i = 0; i > 0; i++) {
13         printf("i = %d\n", i); /* print i */
14     }
15 }
16
EMACS 4.8 (0) Main > test.c

```

Figure 4-1 — EMACS Screen Display

Table 4-1 — EMACS COMMAND SUMMARY

COMMAND	DEFINITION
^@	Sets the mark at the cursor position
^A	Moves to the beginning of the line
^B	Moves back one character
^C	Capitalizes the current character
^D	Deletes current character
^E	Moves to the end of the line
^F	Moves forward one character
^G	Quits from any command in progress
^H	Deletes backward one character
^I	Inserts a tab
^J	Opens a new line and moves to the beginning of it if the next line is non-empty, otherwise moves down one line
^K	Kills to end of line (with argument, kills multiple lines)
^L	Refreshes the screen
^M	Opens a new line and moves to the beginning of it if the next line is non-empty, otherwise moves down one line
^N	Moves down one line
^O	Opens up a new line
^P	Moves up one line
^Q	Quotes the next character
^R	Starts a reverse search
^S	Starts a search
^T	Transposes the next two characters
^U	Multiplies the argument by 4
^V	Moves to the next page
^W	Kills the current region (between cursor and mark)

Table 4-1 — EMACS COMMAND SUMMARY (Contd)

COMMAND	DEFINITION
^X	Is a prefix for more single character commands,
^Y	Restores last killed text (leaves cursor and mark around it)
^Z	Exits one level
^[	Makes the next character a meta character
^]	Makes a local variable of a macro invocation the argument to the next command
^^	Causes the last returned result to become the argument (space) is self-inserting and check for automatic word wrapping
#	Is self-inserting
-	Is self-inserting unless part of a numeric argument
.	Is self-inserting and check for automatic word wrapping
0	Is self-inserting unless part of a numeric argument
1	Is self-inserting unless part of a numeric argument
2	Is self-inserting unless part of a numeric argument
3	Is self-inserting unless part of a numeric argument
4	Is self-inserting unless part of a numeric argument
5	Is self-inserting unless part of a numeric argument
6	Is self-inserting unless part of a numeric argument
7	Is self-inserting unless part of a numeric argument
8	Is self-inserting unless part of a numeric argument
9	Is self-inserting unless part of a numeric argument
:	Is self-inserting, and readjusts indentation in C mode
}	Is self-inserting, and readjusts indentation in C mode
^?	Deletes backward one character
M-^H	Deletes the last word
M-^L	Redisplays with current line at top of page
M-^M	Mails the current buffer

Table 4-1 — EMACS COMMAND SUMMARY (Contd)

COMMAND	DEFINITION
M-^Q	Returns the next input character (in a macro)
M-^R	Regular expression query replace
M-^S	Regular expression search
M-^X	Executes argument 0 as a character command.
M-^]	Assigns the result of the next command to a macro local variable
M-	(Meta space) Sets the mark at the cursor position
M-!	Gets and executes a shell command
M-"	Auto fills the whole buffer
M-\$	Executes a command, saving the output in buffer .exec
M- -	Is self-inserting unless part of a numeric argument
M-/	Starts a comment
M-0	Is self-inserting unless part of a numeric argument
M-1	Is self-inserting unless part of a numeric argument
M-2	Is self-inserting unless part of a numeric argument
M-3	Is self-inserting unless part of a numeric argument
M-4	Is self-inserting unless part of a numeric argument
M-5	Is self-inserting unless part of a numeric argument
M-6	Is self-inserting unless part of a numeric argument
M-7	Is self-inserting unless part of a numeric argument
M-8	Is self-inserting unless part of a numeric argument
M-9	Is self-inserting unless part of a numeric argument
M-:	Maps a character to a command
M-<	Moves to top of file
M->	Moves to bottom of file
M-?	Explains the next character
M-E	Expands an environment variable and returns the result on the kill stack
M-X	Calls a macro by name

Table 4-1 — EMACS COMMAND SUMMARY (Contd)

COMMAND	DEFINITION
M-\	Converts its argument to a character and inserts it
M_	Underlines the next word
M-a	Moves to beginning of sentence
M-b	Moves back one word
M-c	Capitalizes the next word
M-d	Deletes the next word
M-e	Moves to end of sentence
M-f	Moves forward one word
M-g	Moves to a specific line (its argument)
M-l	Converts the next letter to lowercase
M-m	Displays active modes
M-p	Puts the current region in the kill buffer without killing it
M-q	Quotes the next character and adds the 0200 bit
M-r	Starts query replace
M-s	Gives EMACS statistics
M-u	Undoes the last significant text modification
M-v	Moves back one page
M-w	Puts a wall chart of explanations in the buffer
M-x	Calls a macro by name
M-y	Replaces the last restore(^Y) with the next text in the kill stack.
M-{	Enters a command sequence (in a macro)
M-}	Exits one level
M~~	Marks a buffer as being unmodified (up to date)
M-^?	Deletes the last word
<b>Control-X Commands</b>	
^X^A	Accesses the argument list to EMACS
^X^B	Changes buffers (Change to * lists active buffers)
^X^C	Exits gracefully (after asking whether or not to save the buffer)
^X^D	Changes the working directory
^X^E	Calls EMACS recursively taking input from the terminal
^X^F	Edits a file in its own buffer (if file has been read into a buffer, moves to it)
^X^I	Redirects input from a file
^X^K	Kills a buffer
^X^L	Loads a file full of macro definitions

Table 4-1 — EMACS COMMAND SUMMARY (Contd)

COMMAND	DEFINITION
^X^M	Sets mode from argument (prompts for mode name) and string if necessary
^X^N	Changes the buffer or file name
^X^O	Switches between windows
^X^Q	Returns the character under the cursor (in a macro)
^X^R	Reads a new file
^X^S	Saves the buffer in the current file (if modified)
^X^T	Prompts for a buffer name and inserts the text between the cursor and the mark into the named buffer.
^X^U	Updates the display and delays for a specified time
^X^V	Puts the current version on the kill stack.
^X^W	Writes a new or old file
^X^X	Exchanges the mark and the cursor
^X^^	Causes the current window to grow by one line
^X!	Begins a case statement (in a macro)
^X#	Reads or writes global variables
^X%	Exchanges the top of the kill stack with another item
^X&	Compares two strings
^X(	Starts a keyboard macro
^X)	Ends a keyboard macro
^X+	Causes the next entry to the kill stack to append to the previous entry
^X-	Pops the kill stack
^X1	Exits two-window mode
^X2	Enters two-window mode
^X<	Pushes a string from the TTY or macro text into the kill stack
^X=	Gives statistics about the buffer
^X>	Duplicates an item on the kill stack
^X@	Prompts the user with a string from the kill stack and returns the result
^XB	Puts the buffer name into the kill stack
^XE	Executes the keyboard macro
^XF	Puts the file name into the stack
^XL	Makes the next character a meta character
^XT	Traces the next command
^X^	Enters a "while" loop (in a macro)
^Xd	Defines macros from the current buffer
^Xg	Moves to a screen position (arg=128*y+x);

Table 4-1 — EMACS COMMAND SUMMARY (Contd)

COMMAND	DEFINITION
^Xm	Sets mode from argument (prompts for mode name) and string if necessary
^X	Begins a conditional execution sequence (in a macro)
^X~	Performs arithmetic or logical operations (in a macro)



**UNIX RTR Operating System Reference Manual**

	CONTENTS	PAGE
5. COMMANDS . . . . .		5-1



## 5. COMMANDS

The Commands describes programs intended to be invoked directly by the user or by command language procedures, as opposed to subroutines, which are intended to be called by the user's program.

This section consists of many independent entries of a page or more each. The name of the entry appears in the upper corner of the page and is in alphabetical order. Some entries may describe several routines, commands, etc. In such cases, the entry appears only once, alphabetized under its "major" name.



NAME

**admin** — administer logins

SYNOPSIS

```
admin -a nlogin [ -u | p ] [ -s shell ] [ -h home ]  
admin -a nlogin- [ -r | R ]  
admin -d ologin  
admin -c ologin  
admin -o ologin -n nlogin
```

DESCRIPTION

The *admin* command allows craft logins to be added, deleted, and renamed.

The *admin* command also allows passwords to be added or deleted for password-protected commands. In addition, it permits any login's password to be cleared. Only the super user (root) and the craft administrator (manager) may invoke this command.

Craft Login Creation

The format for craft login creation is:

```
admin -a nlogin -u | p [ -s shell ] [ -h home ]
```

where *nlogin* is the desired new login.

When a new login is created, a new entry is made in the password file with the desired name. The uid is unique and not within the range of system logins. The group id is the same for all new creations. This is controlled by the #define CFGID (10) in the source. The password is cleared. Password aging is set for 2 weeks minimum and 10 weeks maximum. A login directory is created in the directory specified by the #define of CFTDIR (/unixa/users) in the source. If the login is for the *UNIX* operating system use, "-u" must be specified. If the login is for PDS/MML use, "-p" must be specified. An option login shell and home directory may also be specified for special applications.

Command Password Creation

The format for command password creation is:

```
admin -a nlogin- [ -r | R ]
```

where *nlogin-* is the desired new login. This new login must be the name of a command to be password-protected with a trailing "-" dash. Thus to password protect the command "sh", the *nlogin-* should be "sh-".

When a new login is created, a new entry is made in the password file with the desired name. The uid is unique and within the range of system logins. The group id is the same for all creations. This is controlled by the #define CFGID (10) in the source. The password is cleared. Password aging is set for 2 weeks minimum and 10 weeks maximum. If "-r" is specified, the command may have limited access (only a few commands have a limited access mode). If "-R" is specified, the command will have full access.

**ADMIN(1)**

The *passwd* command can then be used with this *nlogin*- to set the password of the command.

## Login Deletion

The format for login deletion is:

**admin -d** *ologin*

where *ologin* is the login to delete.

System logins cannot be deleted. Only craft and command logins can be deleted. This procedure deletes all *at(1)* and *crontab(1)* jobs owned by the login. If the login home directory ends [right after a slash (/)] with the login name, then the login home directory and all files within it are deleted. Finally, the entry from the password file is deleted.

## Login Rename

The format for login rename is:

**admin -o** *ologin -n nlogin*

where *ologin* is the old login and *nlogin* is the new login name. System logins cannot be renamed. The procedure is to delete all *at(1)* and *crontab(1)* jobs owned by the old login.

If the old login home directory ends [right after a slash (/)] with the old login name, then the new home directory is the same as old except that the new login name replaces the old login name in the directory name.

In this case, the old home directory is renamed to be the new home directory. If the old home directory does not end with old login name, then the home directory is not changed. No files are deleted in either case.

## Clear Password

The format to clear a password is:

**admin -c** *ologin*

where *ologin* is the login to have its password cleared.

The clear procedure deletes the selected login password from the password file. The next time a login is attempted on this login name, a new password will be requested.

## SYSTEM LOGINS

A system login is a login with a uid of less than or equal to the #define LIMSID (10). System logins can only have their password cleared by this command.

## EXAMPLES

To add the pds/mml login *ralph* enter:

**admin -a ralph -p**

To add the *UNIX* system login *bill* enter:

**admin -a bill -u**

To clear the password of login *mary* enter:

**admin -c mary**

To rename the login *mary* as *ann* enter:

**admin -o mary -n ann**

To delete the login *richard* enter:

**admin -d richard**

To allow a password on the command *sh* enter:

**admin -a sh- -r**

#### FILE

/etc/passwd

#### SEE ALSO

at(1), crontab(1), passwd(1)

#### DIAGNOSTICS

Various diagnostics are issued if the login name is too short, too long, or wrong character set. Also, diagnostics are issued if *nlogin* already exists, or if *ologin* does not exist, or is a system login.



NAME

*at*,  
*batch* — execute commands at a later time

SYNOPSIS

**at** *time* [ *date* ] [ + *increment* ]  
**at** -*r* job ...  
**at** -*l* [ *job* ... ]  
**batch**

DESCRIPTION

*At* and *batch* read commands from standard input are to be executed at a later time. *At* allows you to specify when the commands should be executed, while jobs queued with *batch* will execute when system load level permits. *At -r* removes jobs previously scheduled with *at*. The *-l* option reports all jobs scheduled for the invoking user.

Standard output and standard error output are mailed to the user unless they are redirected elsewhere. The shell environment variables, current directory, and *umask* are retained when the commands are executed. Open file descriptors, traps, and priority are lost.

Users are permitted to use *at* if their name appears in the file */unixa/lib/cron/at.allow*. If that file does not exist, the file */unixa/lib/cron/at.deny* is checked to determine if the user should be denied access to *at*. If neither file exists, only root is allowed to submit a job. If either file is *at.deny*, global usage is permitted. The *allow/deny* files consist of one user name per line.

The *time* may be specified as 1, 2, or 4 digits. One- and two-digit numbers are taken to be hours; four digits to be hours and minutes. The time may alternately be specified as two numbers separated by a colon, meaning *hour : minute*. A suffix *am* or *pm* may be appended; otherwise, a 24-hour clock time is understood. The suffix *zulu* may be used to indicate GMT (Greenwich mean time). The special names *noon*, *midnight*, *now*, and *next* are also recognized.

An optional *date* may be specified as either a month name followed by a day number (and possibly year number preceded by an optional comma) or a day of the week (fully spelled or abbreviated to three characters). Two special "days," *today* and *tomorrow*, are recognized. If no *date* is given, *today* is assumed if the given hour is greater than the current hour and *tomorrow* is assumed if it is less. If the given month is less than the current month (and no year is given), next year is assumed.

The optional *increment* is simply a number suffixed by one of the following: *minutes*, *hours*, *days*, *weeks*, *months*, or *years*. (The singular form is also accepted.)

Thus, legitimate commands include:

**AT(1)**

```
at 0815am Jan 24
at 8:15am Jan 24
at now + 1 day
at 5 pm Friday
```

*At* and *batch* write the job number and schedule time to standard error.

*Batch* submits a batch job. It is almost equivalent to "at now." It is different in the following ways:

1. 1. It goes into a different queue.
2. 2. "at now" will respond with the error message too late.

*At -r* removes jobs previously scheduled by *at* or *batch*. The job number is the number given to you previously by the *at* or *batch* command. You can also get job numbers by typing *at -l*. Unless you are the super user, only you can remove your own jobs.

**EXAMPLES**

The *at* and *batch* commands read from standard input the commands to be executed at a later time. *Sh* (1) provides different ways of specifying standard input. Within your commands, it may be useful to redirect standard output.

This sequence can be used at a terminal:

```
batch
sort filename >outfile
<control-D> (hold down 'control' and depress 'D')
```

This sequence, which demonstrates redirecting standard error to a pipe, is useful in a shell procedure (the sequence of output redirection specifications is significant):

```
batch <<!
sort filename 2>&1 >outfile | mail loginid
!
```

To have a job reschedule itself, invoke *at* from within the shell procedure by including code similar to the following within the shell file:

```
echo "sh shellfile" | at 1900 thursday next week
```

**FILES**

```
/unixa/lib/cron -      main cron directory
/unixa/lib/cron/at.allow - list of allowed users
/unixa/lib/cron/at.deny - list of denied users
/unixa/lib/cron/queue - scheduling information
/unixa/spool/cron/atjobs - spool area
```

SEE ALSO

cron(1), kill(1), mail(1), nice(1), ps(1), sh(1)

DIAGNOSTICS

Complains about various syntax errors and times out of range.



NAME

atomsw — Atomic switch files

SYNOPSIS

**atomsw** file1 file2

DESCRIPTION

Atomic switch of two files. The contents, permissions, and owners of two files are switched in a single operation. In case of a system fault during the operation of this command, *file2* will either have its original contents, permissions and owner, or will have *file1's* contents, permissions and owner. Thus, *file2* is considered precious. *File1* may be truncated in case of a system fault.

RESTRICTIONS

Both files must exist. Both files must reside on the same file system. Neither file may be a "special device" (for example, a TTY port).

To enter this command from the craft shell, switching file "/tmp/abc" with file "/tmp/xyz", enter for MML:

```
EXC:ENVIR:UPROC,FN="/bin/atomsw",ARGS="/tmp/abc"/tmp/xyz";
```

For PDS enter:

```
EXC:ENVIR:UPROC,FN"/bin/atomsw",ARGS("/tmp/abc","/tmp/xyz")!
```

NOTE

File 1 may be lost during a system fault.

FILES

/bin/atomsw



NAME

awk — 1985 version of awk language

SYNOPSIS

```
awk [ -Fregex ] [ 'program' ] [ parameters ] [ files ]  
awk [ -Fregex ] [ -f 'programfile' ] [ parameters ] [ files ]
```

DESCRIPTION

The 1985 *awk* command is an expanded version of the previous standard *awk* pattern-scanning and processing language. The language has been augmented by the addition of dynamically-defined regular expressions, user-defined functions, improved input/output flexibility, and new built-in functions.

The *awk* command scans each input *file* for lines that match any of a set of patterns specified in *program*. With each pattern in *program* there can be an associated action that will be performed when a line of a file matches the pattern. The set of patterns may appear literally as *program*, or in a file specified as *-f programfile*. The *program* string should be enclosed in single quotes (') to protect it from the shell.

*Parameters*, in the form *x=... y=...* etc., may be passed to *awk*.

*Files* are read in order; if there are no *files*, the standard input is read. The file name *-* means the standard input. Each line is matched against the pattern portion of every pattern-action statement; the associated action is performed for each matched pattern.

An input line is made up of fields separated by white space (tabs and blanks). (This default can be changed by using *FS* or the *-F* parameter; see below). The fields are denoted *\$1*, *\$2* ...; *\$0* refers to the entire line.

A pattern-action statement has the form:

```
pattern { action }
```

A missing action means print the line; a missing pattern always matches.

A pattern can be one of the following:

**AWK(1)**

BEGIN	matches before any input is read
END	matches after all input is read
<i>relational expr</i>	matches if the relation holds (for example, NR == 3)
<i>/reg exp/</i> or <i>reg exp</i>	matches the regular expression
<i>pattern1 &amp;&amp; pattern2</i>	matches if both patterns match
<i>pattern1    pattern2</i>	matches if either pattern matches
<i>( pattern )</i>	grouping: matches <i>pattern</i>
<i>!pattern</i>	matches if <i>pattern</i> does not match (NOT operator)
<i>pattern1, pattern</i>	matches from the line where <i>pattern1</i> matched up to (andincluding) the line where <i>pattern2</i> matched
<i>func name (parameter list) { statement }</i>	defines a function called <i>name</i> .

An action is a sequence of statements. Statements are terminated by semicolons, new-lines, or right braces. Statements include the following that affect control-flow:

```

if ( conditional ) statement
if ( conditional ) statement else statement
if ( subscript in array ) statement else statement
while ( conditional ) statement
for ( expression ; expression ; expression ) statement
break
continue
{ [ statement ] ... }
next
exit
exit expression
function-name ( expr, expr, ... )
return
return expression

```

Input-output statements include:

<code>close (filename)</code>	close file
<code>getline</code>	set \$0 from next input record
<code>getline&lt;" file"</code>	set \$0 from next record of <i>file</i>
<code>getline var</code>	set <i>var</i> from next input record
<code>getline var &lt;"file"</code>	set <i>var</i> from next record of <i>file</i>
<code>print</code>	print current record
<code>print expression-list</code>	print expressions
<code>print expression-list &gt;"file"</code>	print expressions on <i>file</i>
<code>printf format , expression-list</code>	format and print
<code>printf format , expression-list&gt;"file"</code>	format and print on <i>file</i>
<code>system(cmd-line)</code>	execute command <i>cmd-line</i> , return its exit status.

Expressions take on string or numeric values as appropriate, and are built using the operators (in increasing precedence) *blank* (string concatenation), *+/-* (addition/subtraction), *\** or */* or *%* (multiplication/division/modulus), *++* or *--* (increment/decrement, prefix and postfix), and *\$* (field value). The C operators *+=*, *-=*, */=*, and *%=* are also available expressions, sharing the lowest precedence with the assignment operator *=*. Logical OR (*|*), and logical AND (*&&*), regular expression match (*()*) and non-match (*!()*), and relational operators (*<*, *<=*, *>*, *>=*, *==*, *!=*) are also available.

Variables may be scalars, array elements (denoted *x[i]*) or fields. Variables are initialized to the null string (treated as zero numerically). Array subscripts may be any string, not necessarily numeric; this allows for a form of associative memory. String constants are quoted (*"*).

printf format conversions that are supported are:

- %c* ASCII character
- %d* decimal number
- %o* unsigned octal number
- %s* string
- %x* unsigned hexadecimal number
- %%* print a *%*; no argument is converted .

Additional parameters may lie between the *%* and the control letter:

- left-justify expression in its field
- width* pad field to this width as needed; leading 0 pads with zeros .

Some of the built-in functions are:

`gsub(regex, string, target)`  
substitute *string* for each substring matching regular expression *regex* in string *target*, return number of substitutions. If *target* is omitted, use \$0. All occurrences of *&* in *string* are replaced by the substring matched by *regex*. The special meaning of *&* may be turned off by preceding it with a backslash, as in *\&* .

**AWK(1)**

*index(target , string)*  
return index of *string* in string *target*, or 0 if not present. (the first character is at index 1.)

*length(string)*  
return length of *string*

*match(string, regex)*  
return index of where *string* matches *regex* or 0 if there is no match; set RSTART and RLENGTH.

*split(string, array, regex)*  
split *string* into *array* on regular expression *regex*, return number of fields. If *regex* is omitted, FS is used in its place.

*sprintf(fmt, expr-list)*  
format *expr-list* according to *fmt*, and return the resulting string

*sub(regex, string, target)*  
like *gsub* except only the first matching substring is replaced

*substr(string, index, number)*  
return *number*-character substring of *string* starting at *index*. (The first character is at index 1.) If *number* is omitted, the substring goes to the end of *string*.

Patterns are regular expressions made of the following pieces (increasing precedence):

c	match non-metacharacter c
\c	match any literal character c
\.	matches any character but newline
^	matches beginning of line or string
\$	matched end of line of string
[abc...]	character class matches any one of abc...
[^abc...]	negated character class matches any but one of abc... and newline
r1 r2	alternation; matches either r1 or r2
r1r2	concatenation; matches r1, then r2
r+	matches one or more r's
r*	matches zero or more r's
r?	matches zero or one r
(r)	grouping; matches r

Regular expression constants must be surrounded by slashes; dynamic regular strings may be the values of variables or strings.

A pattern *regex* may be used to specify the field separator(s) by starting the program with:

```
BEGIN { FS = regex }
```

or by using the **-F** *regex* option.

Other variable names with special meanings include:

ARGC      number of command-line arguments  
ARGV      array of command-line arguments (0..ARGC-1)  
FILENAME the name of the current input file;  
FNR        input record number in current file  
NF        the number of fields in the current record;  
NR        the ordinal number of the current record;  
OFMT      the output format for numbers;  
OFS      the output field separator (default blank);  
ORS      the output record separator (default new-line);  
RLENGTH length of string matched by regular expression in match ()  
RS        the input record separator (default new-line);  
RSTART    beginning position of string matched in match ()  
SUBSEP    separator for array subscripts of form [i, j,...] (default "[0]34")

#### EXAMPLES

Print lines longer than 72 characters:

```
length > 72
```

Print first two fields in opposite order:

```
{ print $2, $1 }
```

Add up first column, print sum and average:

```
{ s += $1 }
```

```
END      { print "sum is", s, "average is", s/NR }
```

Print fields in reverse order:

```
{ for (i = NF; i > 0; --i) print $i }
```

Print all lines between start/stop pairs:

```
/start/, /stop/
```

Print all lines whose first field is different from previous one:

```
$1 != prev { print; prev = $1 }
```

Print file, filling in page numbers starting at 5:

```
/Page/ { $2 = n++; }
```

```
{ print }
```

```
command line: awk -f program n=5 input
```

*AWK(1)*

SEE ALSO

*The Awk Programming Language* by A. V. Aho, B. W. Kernighan, P. J. Weinberger. Published by Addison-Wesley, 1988, ISBN 0-201-07981-X.

LIMITATIONS

Input white space is not preserved on output if fields are involved.

There are no explicit conversions between numbers and strings. To force an expression to be treated as a number, add 0 to it; to force it to be treated as a string, concatenate the null string ("") to it.

The *5ESS*<sup>®</sup>-2000 switch version does not support floating point numbers nor any floating point functions.

235-700-200  
November 1998

COMMANDS

*BANNER(1)*

NAME

banner — make posters

SYNOPSIS

**banner** *strings*

DESCRIPTION

*Banner* prints its arguments (each up to 10 characters long) in large letters on the standard output.

SEE ALSO

echo(1)



## NAME

`basename`, `dirname` — deliver portions of path names

## SYNOPSIS

**basename** *string* [*suffix*] **dirname** *string*

## DESCRIPTION

*Basename* deletes any prefix ending in a slash (/) and the *suffix* (if present in **string**, and prints the results on the standard output. It is normally used inside substitution marks (") within shell procedures.

*Dirname* delivers all but the last level of the path name in *string*.

## EXAMPLES

The following example, invoked with the argument `/usr/src/cmd/cat.c`, compiles the named file and moves the output to a file named **cat** in the current directory:

```
cc $1 mv a.out basename $1 .c
```

The following example will set the shell variable `NAME` to `/usr/src/cmd:`

```
NAME=dirname /usr/src/cmd/cat.c
```

## SEE ALSO

`sh(1)`

## BUGS

The *basename* of `/` is null and considered an error.



235-700-200  
November 1998

**COMMANDS**

*Batch(1)*

NAME

batch

DESCRIPTION

*See at.*



NAME

`cat` — concatenate and print files

SYNOPSIS

`cat[ -b ] [ -s ] file ...`

DESCRIPTION

*Cat* reads each *file* in sequence and writes it on the standard output. Thus:

```
cat file
```

prints the file, and:

```
cat file 1 file 2 > file 3
```

concatenates the first two files and places the result on the third.

If no input file is given, or if the argument `—` is encountered, *cat* reads from the standard input file. Output is buffered unless the `-u` option is specified. The `-s` options makes *cat* silent about nonexistent files. No input file may be the same as the output file unless it is a special file.

WARNING

Command formats such as

```
cat file 1 file 2 > file 1
```

will cause the original data in *file 1* to be lost; therefore, take care when using shell special characters.

SEE ALSO

`cp(1)`, `pr(1)`



NAME

`cftshell` — execute craft shell

SYNOPSIS

**`cftshell`**command

DESCRIPTION

*cftshell* allows the *UNIX* operating system user to execute a single craft shell command. Control returns to the user after the command acknowledgement (e.g., OK, PF). Execution of the command may continue in the background if it is long running.

EXAMPLE

`cftshell 'OP:CLK'`



235-700-200  
November 1998

**COMMANDS**

***CHGRP(1)***

NAME

chgrp

DESCRIPTION

*See chown.*



NAME

chmod — change mode

SYNOPSIS

**chmod**mode files

DESCRIPTION

The permissions of the named *files* are changed according to *mode*, which may be absolute or symbolic. An absolute *mode* is an octal number constructed from the OR of the following modes:

4000 set user on execution  
2000 set group on execution  
1000 sticky bit, see *chmod(2)*  
0400 read by owner  
0200 write by owner  
0100 execute (search in directory) by owner  
0070 read, write, execute (search) by group  
0007 read, write, execute (search) by others

A symbolic *mode* has the form:

[ " who " ] " op permission " [ " op permission " ]

The *who* part is a combination of the letters **u** (for user's permissions), **g** (group) and **o** (other). The letter **a** stands for **ugo** ; the default if *who* is omitted.

*Op* can be **+** to add *permission* to the file mode, **-** to take away *permission*, or **=** to assign *permission* absolutely (all other bits will be reset).

*Permission* is any combination of the letters **r** (read), **w** (write), **x** (execute), **s** (set owner or group ID), and **t** (save text, or sticky); **u** , **g** , or **o** indicate that *permission* is to be taken from the current mode. Omitting *permission* is only useful with **=** to take away all permissions.

Multiple symbolic modes separated by commas may be given. Operations are performed in the order specified. The letter **s** is only useful with **u** or **g**, and **t** only works with **u**.

Only the owner of a file (or the super user) may change its mode.

EXAMPLES

The first example denies write permission to others; the second makes a file executable:

```
chmod o-w file
```

```
chmod +x file
```

SEE ALSO

ls(1)



NAME

`chown`, `chgrp` — change group or owner

SYNOPSIS

**chown**owner file ...  
**chgrp** group file ...

DESCRIPTION

*Chown* changes the owner of the *files* to *owner*. The owner may be either a decimal user ID or a login name found in the password file.

*Chgrp* changes the group of the *files* to *group*. The group may be either a decimal group ID or a group name found in the group file.

FILES

/etc/passwd  
/etc/group



NAME

`closewd` — close the window that is opened by *openwd*

SYNOPSIS

**`closewd`**

DESCRIPTION

*Closewd* closes the window opened by *openwd*. After *closewd*, processes may not open (with write) block devices for mounted file systems. This command should follow *openwd*. But its use is not mandatory since the window will be closed automatically by the file manager (20 minutes after open).

SEE ALSO

`openwd(1)`, `fsdb(1)`



NAME

`clrfs` — construct a file system

SYNOPSIS

`clrfs` special blocks [i-node blocks]

DESCRIPTION

*Clrfs* constructs a file system by writing on the special file according to the directions found in the remainder of the command line. If the second argument is given as a string of digits, *clrfs* builds a file system with a single empty directory on it.

The size of the file system is the value of *blocks* interpreted as a decimal number. The boot block is left uninitialized. If the optional number of *i-node blocks* is given, the i-list consists of eight times that value (that is, there are eight i-nodes per i-node block). If the optional number of *i-node blocks* is not given, the default is the number of *blocks* divided by four.



NAME

*clri* — clear i-node

SYNOPSIS

*/etc/clri* filesystem i-numbers

DESCRIPTION

*Clri* writes zeros on the 64 bytes occupied by the i-nodes numbered **i-number**. The **filesystem** argument must be a special filename referring to a device containing a file system. After *clri*, any blocks in the affected file will show up as "missing" in an *ichk* of the **filesystem**.

Read and write permission is required on the specified **filesystem** device. The i-node becomes allocatable.

The primary purpose of this routine is to remove a file which does not appear in any directory. If it is used to remove an i-node which appears in a directory, care should be taken to track down the entry and remove it. Otherwise, when the i-node is reallocated to some new file, the old entry will still point to that file. At that point, removing the old entry will destroy the new file. The new entry will again point to an unallocated i-node, so the whole cycle is likely to be repeated again and again.

SEE ALSO

*ichk(1)*

LIMITATIONS

If the file is open, *clri* is likely to be ineffective.



NAME

`cmp` — compare two files

SYNOPSIS

`cmp`[-l] [-s] file1 file2

DESCRIPTION

The two files are compared. (If *file1* is -, the standard input is used.) Under default options, *cmp* makes no comment if the files are the same; if they differ, it announces the byte and line number at which the difference occurred. If one file is an initial subsequence of the other, that fact is noted.

Options:

- l           Print the byte number (decimal) and the differing bytes (octal) for each difference.
- s           Print nothing for differing files; return codes only.

SEE ALSO

`diff(1)`

DIAGNOSTICS

Exit code 0 is returned for identical files, 1 for different files, and 2 for an inaccessible or missing argument.



NAME

`coflsb` — clear off-line superblock

SYNOPSIS

`coflsb`<partition name> <mount point>

DESCRIPTION

The *coflsb* command locates the given *partition name* in the off-line SG data base, and extracts the disk number, partition number, partition block count, and partition inode block count for that partition. This information is then used to update the superblock of the named off-line partition. The updated off-line partition is then mounted on `/tmp/ofl/ mount point`. When the partition is successfully mounted, the *coflsb* process goes to sleep for four (4) hours, allowing data transfer to take place.

DIAGNOSTICS

*Coflsb* will fail with an error code if the off-line SG data base cannot be attached or read, or if the requested off-line partition cannot be successfully mounted. Appropriate error messages are printed before *coflsb* exits.

FILES

`/tmp/dev/sgdbase` - temporary device file for attaching to SG data base

`/tmp/dev/attpar` - temporary device file for attaching requested off-line partition

`/tmp/ofl/attpar` - temporary mount point for off-line SG data base

`/tmp/ofl/<mount point>` - temporary mount point for requested off-line partition

EXAMPLES

The command:

```
coflsb no5sodd1 fubar
```

will mount the off-line `no5sodd1` partition on `/tmp/ofl/fubar`.

CAVEATS

Since the off-line partition is only accessible while *coflsb* is active, this command should be run asynchronously (in the background). Also, *coflsb* assumes that the odd numbered disks are the off-line disks.

SEE ALSO

`fsinit(1)`



## NAME

`compress`, `uncompress`, `zcat` — compress file, uncompress file, cat compressed file

## SYNOPSIS

```
compress[ -f ] [ -v ] [ -c ] [ -V ] [ -b bits ] [ name ... ]  
uncompress[ -f ] [ -v ] [ -c ] [ -V ] [ name ... ]  
zcat[ -V ] [ name ... ]
```

## DESCRIPTION

*Compress* reduces the size of the named files using adaptive Lempel-Ziv coding. Whenever possible, each file is replaced by one with the extension, `.Z`, while keeping the same ownership modes, access and modification times. If no files are specified, the standard input is compressed to standard output. Compressed files can be restored to their original form by using *uncompress* or *zcat*.

The `-f` option will force compression of *name*. This is useful for compressing an entire directory, even if some of the files do not shrink. If `-f` is not given and *compress* is run in the foreground, the user is prompted as to whether an existing file should be overwritten.

The `-c` option makes *compress/uncompress* write to the standard output; no files are changed. The nondestructive behavior of *zcat* is identical to that of *uncompress -c*.

*Compress* uses the modified Lempel-Ziv algorithm popularized in "A Technique for High Performance Data Compression," Terry A. Welch, *IEEE Computer*, Vol. 17, No. 6 (June 1984), Pages 8 through 19. Common substrings in the file are first replaced by 9-bit codes 257 and up. When code 512 is reached, the algorithm switches to 10-bit codes and continues to use more bits until the limit specified by the `-b` flag is reached (default 16). *Bits* must be between 9 and 16. The default can be changed in the source to allow *compress* to be run on a smaller machine.

After the *bits* limit is attained, *compress* periodically checks the compression ratio. If it is increasing, *compress* continues to use the existing code dictionary. However, if the compression ratio decreases, *compress* discards the table of substrings and rebuilds it from scratch. This allows the algorithm to adapt to the next "block" of the file.

Note that the `-b` flag is omitted for *uncompress*, since the *bits* parameter specified during compression is encoded within the output, along with a magic number to ensure that neither decompression of random data nor recompression of compressed data is attempted.

The amount of compression obtained depends on the size of the input, the number of *bits* per code, and the distribution of common substrings. Typically, text such as source code or English is reduced by 50 to 60 percent. Compression is generally much better than that achieved by Huffman coding (as used in *pack*), or adaptive Huffman coding (*compact*), and takes less time to compute.

**COMPRESS(1)**

Under the **-v** option, a message is printed yielding the percentage of reduction for each file compressed.

If the **-V** option is specified, the current version and compile options are printed on stderr.

Exit status is normally 0; if the last file is larger after (attempted) compression, the status is 2; if an error occurs, exit status is 1.

## DIAGNOSTICS

Usage: compress [-dfvcV] [-b maxbits] [file ...]  
Invalid options were specified on the command line.

Missing maxbits  
Maxbits must follow **-b**.

*file*: not in compressed format  
The file specified to *uncompress* has not been compressed.

*file*: compressed with *xx* bits, can only handle *yy* bits  
*File* was compressed by a program that could deal with more *bits* than the compress code on this machine. Recompress the file with smaller *bits*.

*File*: already has **.Z** suffix -- no change  
The file is assumed to be already compressed. Rename the file and try again.

*file*:filename too long to tack on **.Z**  
The file cannot be compressed because its name is longer than 12 characters. Rename and try again.

*file*:already exists; do you wish to overwrite (y or n)?  
Respond "y" if you want the output file to be replaced; "n" if not.

uncompress: corrupt input  
A SIGSEGV violation was detected which usually means that the input file has been corrupted.

Compression: *xx.xx*%  
Percentage of the input saved by compression. (Relevant only for **-v**.)

-- not a regular file: unchanged  
When the input file is not a regular file (e.g., a directory), it is left unaltered.

-- has *xx* other links: unchanged  
The input file has links; it is left unchanged. See *ln* (1) for more information.

-- file unchanged  
No savings is achieved by compression. The input remains unchanged.

To enter this command from the craft shell, compressing file `"/tmp/abc"` enter, for MML:

EXC:ENVIR:UPROC,FN="/bin/compress",ARGS="/tmp/abc";

For PDS enter:

EXC:ENVIR:UPROC,FN"/bin/compress",ARGS"/tmp/abc"!

#### BUGS

Although compressed files are compatible between machines with large memory, **-b 12** should be used for file transfer to architectures with a small process data space (64 kb or less, as exhibited by the Intel 80286, etc.).



NAME

`cp`, `ln`, `mv` — copy, link, or move files

SYNOPSIS

`cp` *file1* [*file2* ...] *target*  
`ln` *file1* [*file2* ...] *target*  
`mv` *file1* [*file2* ...] *target*

DESCRIPTION

*File1* is copied (linked, moved) to *target*. Under no circumstance can *file1* and *target* be the same [take care when using *sh* (1) metacharacters]. If *target* is a directory, then one or more files are copied (linked, moved) to that directory.

If *mv* determines that the mode of *target* forbids writing, it will print the mode and read the standard input for one line (if the standard input is a terminal). If the line begins with *y*, the move takes place; if not, *mv* exits.

Only *mv* will allow *file1* to be a directory, in which case the directory rename will occur only if the two directories have the same parent.

SEE ALSO

`cpio`(1), `rm`(1)

BUGS

If *file1* and *target* lie on different file systems, *mv* must copy the file and delete the original. In this case, the owner name becomes that of the copying process and any linking relationship with other files is lost.

*Ln* will not link across file systems.



NAME

`cpio` — copy file archives in and out

SYNOPSIS

```
cpio -o [ acBv ]
cpio -i [ BcdmrtuvfsSb6 ] [ patterns ]
cpio -p [ adlmruv ] directory
```

DESCRIPTION

**Cpio -o** (copy out) reads the standard input to obtain a list of path names and copies those files onto the standard output together with path name and status information.

**Cpio -i** (copy in) extracts files from the standard input which is assumed to be the product of a previous **cpio -o**. Only files with names that match *patterns* are selected. *Patterns* are given in the name-generating notation of *sh* (1). In *patterns*, meta-characters `?`, `*`, and `[ ... ]` match the slash `/` character. Multiple *patterns* may be specified and if no *patterns* are specified, the default for *patterns* is `*` (i.e., select all files). The extracted files are conditionally created and copied into the current directory tree based upon the options described below.

**Cpio -p** (pass) reads the standard input to obtain a list of path names of files that are conditionally created and copied into the destination *directory* tree based upon the options described below.

The meanings of the available options are:

- a** Reset access times of input files after they have been copied.
- B** Input/output is to be blocked 5,120 bytes to the record (does not apply to the *pass* option; meaningful only with data directed to or from `/dev/rmt?`).
- d** *Directories* are to be created as needed.
- c** Write *header* information in ASCII character form for portability.
- r** Interactively *rename* files. If the user types a null line, the file is skipped.
- t** Print a *table of contents* of the input. No files are created.
- u** Copy *unconditionally* (normally, an older file will not replace a newer file with the same name).
- v** *Verbose*: causes a list of file names to be printed. When used with the **t** option, the table of contents looks like the output of an `ls -l` command [see *ls* (1)].
- l** Whenever possible, link files rather than copying them. Usable only with the **-p** option.
- m** Retain previous file modification time. This option is ineffective on directories that are being copied.
- f** Copy in all files except those in *patterns*.
- s** Swap bytes. Use only with the **-i** option.

**CPIO(1)**

- S** Swap halfwords. Use only with the **-i** option.
- b** Swap both bytes and halfwords. Use only with the **-i** option.
- 6** Process an old (that is, *UNIX System Sixth Edition* format) file. Only useful with **-i** (copy in).

## EXAMPLES

The first example below copies the contents of a directory into an archive; the second duplicates a directory hierarchy:

```
ls | cpio -o >/dev/mt08
```

```
cd olddir
```

```
find . -depth -print | cpio -pdl newdir
```

The trivial case "find . -depth -print | cpio -oB >/dev/rmt08" can be handled more efficiently by:

```
find . -cpio /dev/mt08
```

## SEE ALSO

find(1)

## BUGS

Path names are restricted to 128 characters. If there are too many unique linked files, the program runs out of memory to keep track of them and, thereafter, linking information is lost. Only the super user can copy special files. The **-B** option does not work with certain magnetic tape drives.

NAME

`cron` — clock daemon

SYNOPSIS

`/etc/cron`

DESCRIPTION

*Cron* executes commands at specified dates and times. Regularly scheduled commands can be specified according to instructions found in crontab files; users can submit their own crontab file via the *crontab* command. Commands which are to be executed only once may be submitted via the *at* command. Since *cron* never exits, it should only be executed once.

*Cron* only examines crontab files and at command files during process initialization and when a file changes. This reduces the overhead of checking for new or changed files at regularly scheduled intervals.

Queue Definitions

Cron can limit dynamically the number of concurrently running jobs. It can also maintain up to 26 separate queues and control the number of jobs executed in each one. The file `queuedefs` is used to maintain definitions for all queues. Each line of the file must have the following format:

```
q.NNjNNnNNw
```

where

`q` is a letter a-z indicating the job queue.

`NNj`

is the limit on jobs that can be running at any one time for the job queue. (`NN` is an integer; its default is 100.)

`NNn` is the **nice (1)** value that is assigned each command executed for the job queue. (Default is 2.)

`NNw` is the time (in seconds) to wait before retrying to execute a command if all the criteria for running the command are not met. (Default is 60.)

Empty fields are initialized to the default values. The following is an example of a `queuedefs` file.

```
a.4j1n
b.2j2n90w
c.0n10j
n.120w4n1j
```

Changes to queue definitions take effect before the next job is executed by the `cron` daemon. If this file does not exist, the default values are used.

**CRON(1)**

## Prototype File

The prototype file provides a method to customize **at** command files by controlling what information is written into the **at** job file. If a file named **.proto.q** exists (where **q** indicates a queue name), this file is copied into the job file. Otherwise, the file **.proto** is used.

The following substitutions are made during creation of an **at** job file.

**\$m** user's current file creation mask.  
**\$l** user's current file size limit (not implemented)  
**\$d** name of the current working directory  
**\$t** time (in seconds) that the jobs is scheduled to execute  
**\$<** read standard input until EOF is reached

The following is an example of a prototype file:

```
cd $d
ulimit $l
umask $m
$<
```

The **at** command will exit with an error if no prototype file exists.

## Log Information

Cron logs all command invocations, terminations, and status information in the file **log**. Records that begin with the character **>** pertain to command invocations. Two invocation records are written for each command execution. The first displays the command that is being executed; the second contains the login name, process id, job queue, and a timestamp for when the command was invoked. Command termination records begin with the character **<** and are similar to the second invocation record, except that a nonzero termination status or exit status is also printed. Records that begin with an **!** indicate status information.

## Files to Limit Access to the Facility

There is potential for misuse of system resources through use of the **crontab** and **at** commands. Cron provides a method to restrict user access to it. The files **cron.allow** and **at.allow** contain login names of users (one per line) allowed access to the **crontab** and **at** commands, while the files **cron.deny** and **at.deny** contain login names of users denied access to the commands.

When a user submits a crontab file, **crontab** checks **cron.allow** for a list of users permitted to have a crontab. If that file does not exist, the file **cron.deny** is scanned for users who are denied crontabs. If neither file exists, only root is allowed to have a crontab file. The same scheme is used for determining access to the **at** command. The null file **cron.allow** would mean no user is allowed a crontab while a null file **cron.deny** would mean no user is denied a crontab.

### Defining Queues

The **at** command can queue jobs in one of 26 different queues, with the **cron** daemon controlling the number of executions for each queue. This mechanism can be used by system administrators to limit the number of simultaneous executions of high-resource commands. Running the **at** command with **-q c** as the first argument queues the command in queue *c*. The default queue is **a**. A special queue **b** is defined to be a batch queue; jobs in this queue run whenever the defined maximum level is not exceeded (specified in the **queuedefs** file). Cron executions are limited by the definition of the queue *c*. Jobs in all other queues run at the time specified on the command line.

The ability to define queues gives administrators a way to restrict executions for commands. An example of its use would be in providing a batch **sort (1)** facility. All that need be done is to define a queue in the **queuedefs** file to spool **sort** commands, create a prototype file (if necessary), and replace the **sort** command with a shell that creates an **at** job and invokes the real **sort**. The real **sort** command should be moved out of a standard bin location.

### FILES

**/unixa/lib/cron** main cron directory  
**/unixa/lib/cron/log** accounting information  
**/unixa/spool/cron** spool area

### SEE ALSO

**at(1)**, **crontab(1)**, **sh(1)**

### DIAGNOSTICS

A history of all actions taken by cron are recorded in **/unixa/lib/cron/log**.



NAME

`crontab` — user crontab file

SYNOPSIS

`crontab` [file]  
`crontab -r`  
`crontab -l`

DESCRIPTION

*Crontab* copies the specified file, or standard input if no file is specified, into a directory that holds all users' crontabs. The `-r` option removes a user's crontab from the crontab directory. *Crontab -l* will list the crontab file for the invoking user.

A user is permitted to use *crontab* if their name appears in the file `/unixa/lib/cron/cron.allow`. If that file does not exist, the file `/unixa/lib/cron/cron.deny` is checked to determine if the user should be denied access to *crontab*. If neither file exists, only root is allowed to submit a job. If either file is `at.deny`, global usage is permitted. The allow/deny files consist of one user name per line.

A crontab file consists of lines of six fields each. The fields are separated by spaces or tabs. The first five are integer patterns that specify the following:

minute (0-59),  
hour (0-23),  
day of the month (1-31),  
month of the year (1-12),  
day of the week (0-6 with 0=Sunday).

Each of these patterns may be either an asterisk (meaning all legal values), or a list of elements separated by commas. An element is either a number, or two numbers separated by a minus sign (meaning an inclusive range). Note that the specification of days may be made by two fields (day of the month and day of the week). If both are specified as a list of elements, both are adhered to. For example, `0 0 1,15 * 1` would run a command on the first and fifteenth of each month, as well as on every Monday. To specify days by only one field, the other field should be set to `*` (for example, `0 0 * * 1` would run a command only on Mondays).

The sixth field of a line in a crontab file is a string that is executed by the shell at the specified times. A percent character in this field (unless escaped by `\`) is translated to a new-line character. Only the first line (up to a `%` or end of line) of the command field is executed by the shell. The other lines are made available to the command as standard input.

The shell is invoked from your `$HOME` directory with an `arg0` of `sh`. Users who desire to have their `.profile` executed must explicitly do so in the crontab file.

*Cron* supplies a default environment for every shell, defining `HOME`, `LOGNAME`, `SHELL(=/bin/sh)`, and `PATH(=/bin:/usr/bin:)`.

**CRONTAB(1)**

**Note:** Users should remember to redirect the standard output and standard error of their commands! If this is not done, any generated output or errors will be mailed to the user.

## FILES

```
/unixa/lib/cron          # main cron directory
/unixa/spool/cron/crontabs # spool area
/unixa/lib/cron/log      # accounting information
/unixa/lib/cron/cron.allow # list of allowed users
/unixa/lib/cron/cron.deny # list of denied users
```

## SEE ALSO

cron(1), sh(1)

NAME

cut — cut out selected fields or each line of a file

SYNOPSIS

```
cut-clist [file1 file2 ...]
cut-flist [-d char] [-s] [file1 file2 ...]
```

DESCRIPTION

Use *cut* to cut out columns from a table or fields from each line of a file; in data base parlance, it implements the projection of a relation. The fields as specified by *list* can be fixed length, that is, character positions as on a punched card (**-c** option) or the length can vary from line to line and be marked with a field delimiter character like *tab* (**-f** option). *Cut* can be used as a filter; if no files are given, the standard input is used.

The meanings of the options are:

- list*           A comma-separated list of integer field numbers (in increasing order), with optional **-** to indicate ranges as in the **-o** option of *nroff* / *troff* for page ranges; for example, *1,4,7*; *1-3,8*; *-5,10* (short for **1-5,10**); or *3-* (short for third through last field).
- clist**       The *list* following **-c** (no space) specifies character positions (for example, *-c1-72* would pass the first 72 characters of each line).
- flist**       The *list* following **-f** is a list of fields assumed to be separated in the file by a delimiter character (see **-d**); for example, **-f1,7** copies the first and seventh field only. Lines with no field delimiters will be passed through intact (useful for table subheadings), unless **-s** is specified.
- dchar**       The character following **-d** is the field delimiter (**-f** option only). Default is *tab*. Space or other characters with special meaning to the shell must be quoted.
- s**           Suppresses lines with no delimiter characters in case of **-f** option. Unless specified, lines with no delimiters will be passed through untouched.

Either the **-c** or **-f** option must be specified.

HINTS

Use *grep* (1) to make horizontal "cuts" (by context) through a file, or *paste* (1) to put files together column-wise (that is, horizontally). To reorder columns in a table, use *cut* and *paste*.

EXAMPLES

```
cut -d: -f1,5 /etc/passwd
      mapping of user IDs to names
name=';who am i | cut -f1 -d" "'
      to set name to current login name.
```

## COMMANDS

235-700-200  
November 1998

### *CUT(1)*

#### DIAGNOSTICS

*line too long*

A line can have no more than 1023 characters or fields.

*bad list for c/f option*

Missing **-c** or **-f** option or incorrectly specified *list*. No error occurs if a line has fewer fields than the *list* calls for.

*no fields*

The *list* is empty.

#### SEE ALSO

grep(1), paste(1)

NAME

*cx* — 3B20D computer core file examiners

SYNOPSIS

**cx** [-!aphx] [-v *vaddr* [-n *nbytes*]] ... **file**

DESCRIPTION

*cx* prints the segment images of the specified file. *cx* resides on the 3B20D computer. When invoked without options, *cx* prints all data from all nonexecutable segments in the given file. Options to modify this behavior follow:

- !** prints descriptions of each flag's meaning. Illegal command lines generate a terse summary of usage without the descriptions.
- a** prints data for all nonexecutable segment images. This is the default action when no **-v**, **-h**, or **-x** flags are present.
- h** prints a header for every segment. This flag also suppresses the printing of segment data, but **-a**, **-v**, or **-x** cause the designated data to be printed.
- p** requests the *pfile* meanings, which are different in a small number of cases. Each segment's header contains a set of attribute flags. *cx* normally decodes the flags using their *execution* time meanings.
- x** prints data for all executable segments. Without **-x**, the data are suppressed.
- v *vaddr*** starts printing data from the given address in the segment(s) containing virtual address *vaddr*. Unless **-n** is also given, the rest of the segment will be printed. Using one or more **-v** specifications causes other segment images not to be printed, but **-a** and **-x** override that.
- n *nbytes*** prints *nbytes* bytes instead of printing all of a segment's data. A count of zero means print the remainder of the image. This option applies only to the preceding **-v**, which must be present.

Values for *addr* and *nbytes* are interpreted as hexadecimal numbers.

A core file's format must be the same as that written by the process manager, which is very similar to executable *pfiles*.

Usually the process manager places core files in */cdmp/name*, where *name* is the ASCII name in the PCB.

In actual core files, the segment images occur in the same order as they appeared in the process's segment list. The PCB is first, and the stack is second. Those segment images in *pfiles* are empty.

More than one segment can be mapped to the same virtual address, but only one of them can be active. **-v** directives apply to all mapped segments, not just the active one.



NAME

date — print and set the date

SYNOPSIS

date[ mmddhhmmyy ] [ +format ]

DESCRIPTION

If no argument is given, or if the argument begins with **+**, the current date and time are printed. Otherwise, the current date is set. The first *mm* is the month number; *dd* is the day number in the month; *hh* is the hour number (24-hour system); the second *mm* is the minute number; *yy* is the last 2 digits of the year number. For example:

```
date 1008004585
```

sets the date to Oct 8, 12:45 AM 1985. The system operates in GMT. *Date* takes care of the conversion to and from local standard and daylight time.

If the argument begins with **+**, the output of *date* is under the control of the user. All output fields are of fixed size (zero padded if necessary). Each field descriptor is preceded by **%** and will be replaced in the output by its corresponding value. A single **%** is encoded by **%%**. All other characters are copied to the output without change. The string is always terminated with a new-line character.

Field Descriptors:

<b>n</b>	insert a new-line character
<b>t</b>	insert a tab character
<b>m</b>	month of year - 01 to 12
<b>d</b>	day of month - 01 to 31
<b>y</b>	last 2 digits of year - 00 to 99
<b>D</b>	date as mm/dd/yy
<b>H</b>	hour - 00 to 23
<b>M</b>	minute - 00 to 59
<b>S</b>	second - 00 to 59
<b>T</b>	time as HH:MM:SS
<b>j</b>	day of year - 001 to 366
<b>w</b>	day of week - Sunday = 0
<b>a</b>	abbreviated weekday - Sun to Sat
<b>h</b>	abbreviated month - Jan to Dec
<b>r</b>	time in AM/PM notation

EXAMPLE

```
date '+DATE: %m/%d/%y%nTIME: %H:%M:%S'
```

## COMMANDS

235-700-200  
November 1998

### *DATE(1)*

would have generated as output:

DATE: 08/01/76

TIME: 14:45:05

## DIAGNOSTICS

*No permission*

if you are not the super user and you try to change the date;

*bad conversion*

if the date set is syntactically incorrect;

*bad format character*

if the field descriptor is not recognizable.

## FILES

/dev/kmem

## WARNING

If the system is call processing, use the PDS/MML command "SET:CLK" to change the date.

NAME

dc — desk calculator

SYNOPSIS

dc[ file ]

DESCRIPTION

*Dc* is an arbitrary precision arithmetic package. Ordinarily it operates on decimal integers, but one may specify an input base, output base, and a number of fractional digits to be maintained. The overall structure of *dc* is a stacking (reverse Polish) calculator. If an argument is given, input is taken from that file until its end, then from the standard input. The following constructions are recognized:

- |                                |   |
|--------------------------------|---|
| <i>number</i>                  | The value of the number is pushed on the stack. A number is an unbroken string of the digits 0-9. It may be preceded by an underscore ( <code>_</code> ) to input a negative number. Numbers may contain decimal points.  |
| <code>+ - / * % ^</code>       | The top two values on the stack are added ( <code>+</code> ), subtracted ( <code>-</code> ), multiplied ( <code>*</code> ), divided ( <code>/</code> ), remaindered ( <code>%</code> ), or exponentiated ( <code>^</code> ). The two entries are popped off the stack; the result is pushed on the stack in their place. Any fractional part of an exponent is ignored. |
| <code>sx</code>                | The top of the stack is popped and stored into a register named <i>x</i> , where <i>x</i> may be any character. If the <code>s</code> is capitalized, <i>x</i> is treated as a stack and the value is pushed on it.   |
| <code>lx</code>                | The value in register <i>x</i> is pushed on the stack. The register <i>x</i> is not altered. All registers start with zero value. If the <code>l</code> is capitalized, register <i>x</i> is treated as a stack and its top value is popped onto the main stack.  |
| <code>d</code>                 | The top value on the stack is duplicated.   |
| <code>p</code>                 | The top value on the stack is printed. The top value remains unchanged. <code>P</code> interprets the top of the stack as an ASCII string, removes it, and prints it.   |
| <code>f</code>                 | All values on the stack are printed.  |
| <code>q</code>                 | exits the program. If executing a string, the recursion level is popped by two. If <code>q</code> is capitalized, the top value on the stack is popped and the string execution level is popped by that value.  |
| <code>x</code>                 | treats the top element of the stack as a character string and executes it as a string of <i>dc</i> commands.  |
| <code>X</code>                 | replaces the number on the top of the stack with its scale factor.  |
| <code>[ ... ]</code>           | puts the bracketed ASCII string onto the top of the stack.  |
| <code>&lt; x &gt; x = x</code> | The top two elements of the stack are popped and compared. Register <i>x</i> is evaluated if they obey the stated relation.   |
| <code>v</code>                 | replaces the top element on the stack by its square root. Any existing fractional part of the argument is taken into account, but otherwise the scale factor is ignored.  |

*DC(1)*

- !** interprets the rest of the line as a UTS system command.
- c** All values on the stack are popped.
- i** The top value on the stack is popped and used as the number radix for further input.
- I** pushes the input base on the top of the stack.
- o** The top value on the stack is popped and used as the number radix for further output.
- O** pushes the output base on the top of the stack.
- k** The top of the stack is popped, and that value is used as a non-negative scale factor: the appropriate number of places are printed on output, and maintained during multiplication, division, and exponentiation. The interaction of scale factor, input base, and output base will be reasonable if all are changed together.
- z** The stack level is pushed onto the stack.
- Z** replaces the number on the top of the stack with its length.
- ?** A line of input is taken from the input source (usually the terminal) and executed.
- ;** **:** are used for array operations.

EXAMPLE

This example prints the first ten values of n!:

```
[!a1+dsa*pla10>y]sy
0sa1
lyx
```

DIAGNOSTICS

- x is unimplemented*  
where *x* is an octal number.
- stack empty* for not enough elements on the stack to do what was asked.
- Out of space*  
when the free list is exhausted (too many digits).
- Out of headers*  
for too many numbers being kept around.
- Out of pushdown*  
for too many items on the stack.
- Nesting Depth*  
for too many levels of nested execution.

NAME

**dd** — convert and copy a file

SYNOPSIS

**dd**[option=value] ...

DESCRIPTION

*Dd* copies the specified input file to the specified output with possible conversions. The standard input and output are used by default. The input and output block size may be specified to take advantage of raw physical I/O.

<i>option</i>	<i>values</i>
<b>if=</b> file	input file name; standard input is default
<b>of=</b> file	output file name; standard output is default
<b>ibs=</b> n	input block size <i>n</i> bytes (default 512)
<b>obs=</b> n	output block size (default 512)
<b>bs=</b> n	set both input and output block size, superseding <i>ibs</i> and <i>obs</i> ; also, if no conversion is specified, it is particularly efficient since no in-core copy need be done
<b>cbs=</b> n	conversion buffer size
<b>skip=</b> n	skip <i>n</i> "" input records before starting copy
<b>seek=</b> n	seek <i>n</i> records from beginning of output file before copying
<b>count=</b> n	copy only <i>n</i> "" input records
<b>conv=ascii</b>	convert EBCDIC to ASCII
<b>ebcdic</b>	convert ASCII to EBCDIC
<b>ibm</b>	slightly different map of ASCII to EBCDIC
<b>lcase</b>	map alphabetic to lowercase
<b>ucase</b>	map alphabetic to uppercase
<b>swab</b>	swap every pair of bytes
<b>noerror</b>	do not stop processing on an error
<b>sync</b>	pad every input record to <i>ibs</i>
... , ...	several comma-separated conversions

Where sizes are specified, a number of bytes is expected. A number may end with **k**, **b**, or **w** to specify multiplication by 1024, 512, or 2, respectively; a pair of numbers may be separated by **x** to indicate a product.

*Cbs* is used only if *ascii* or *ebcdic* conversion is specified. In the former case, *cbs* characters are placed into the conversion buffer, converted to ASCII, and trailing blanks trimmed and new-line added before sending the line to the output. In the latter case, ASCII characters are read into the conversion buffer, converted to EBCDIC, and blanks added to make up an output record of size *cbs*.

After completion, *dd* reports the number of whole and partial input and output blocks.

***DD(1)***

## EXAMPLE

This command will read an EBCDIC tape blocked ten 80-byte EBCDIC card images per record into the ASCII file **x**:

```
dd if=/dev/mt00 of=x ibs=800 cbs=80 conv=ascii,lcase
```

Note the use of raw magtape. *Dd* is especially suited to I/O on the raw physical devices because it allows reading and writing in arbitrary record sizes.

## SEE ALSO

*cp(1)*

## DIAGNOSTICS

*f+p records in(out)* numbers of full and partial records  
read(written)

## BUGS

The ASCII/EBCDIC conversion tables are taken from the 256-character standard in the CACM (Communications of the Association for Computing Machinery) Nov, 1968. The *ibm* conversion, while less blessed as a standard, corresponds better to certain IBM print train conventions. There is no universal solution.

New lines are inserted only on conversion to ASCII; padding is done only on conversion to EBCDIC. These should be separate options.

The magnetic tape (MT) unit driver program of the 3B20D has a bug in it that does not permit files of certain sizes to be written. Specifically, if the file size is  $(521*N)+1$  (that is, 1, 513, 1025 ..., the driver fails to write out the last byte of the file. This may or may not result in an error message.

This most commonly occurs when a file is being copied to tape by use of the "dd" command as part of a non-official Lucent procedure.

This problem can be avoided by using the

```
COPY:TAPE:OUT
```

command to copy files to tape.

NAME

**df** — report number of free disk blocks

SYNOPSIS

**df** [ file-systems ]

DESCRIPTION

*Df* prints out the number of free blocks available for on-line file systems by examining the counts kept in the in-core super-blocks; *file-systems* may be specified either by device name (for example, **/dev/bwm** ) or by mounted directory name (for example, **/etc/bwm** ). If the *file-systems* argument is unspecified, the free space on all of the mounted file systems is printed. Blocks are reported as 512 byte blocks.



NAME

**dgnnm** — assign special diagnostic filename

SYNOPSIS

**dgnnm** unit-name unit-number

DESCRIPTION

*Dgnnm* creates and returns a block special device file name that it has assigned to the unit specified in the command line.

The unit has been reserved for diagnostics when this program completes. After use, it must be released by *udgnnm*.

FILES

/dev/ecd

/dgn/xxx

/temp/ecdxxxxxxx

/tmp/xxx

SEE ALSO

**udgnnm(1)**

DIAGNOSTICS

Error numbers are returned. These numbers may be interpreted in *MOVEerrcod.h*.



NAME

`diff` — differential file comparator

SYNOPSIS

`diff [ -efbh ] file1 file2`

DESCRIPTION

*Diff* tells what lines must be changed in two files to bring them into agreement. If *file1* (*file2*) is -, the standard input is used. If *file1* (*file2*) is a directory, then a file in that directory with the name *file2* (*file1*) is used. The normal output contains lines of these forms:

`n1 a n3,n4`

`n1,n2 d n3`

`n1,n2 c n3,n4`

These lines resemble *ed* commands to convert *file1* into *file2*. The numbers after the letters pertain to *file2*. In fact, by exchanging **a** for **d** and reading backward one may ascertain equally how to convert *file2* into *file1*. As in *ed*, identical pairs where *n1* = *n2* or *n3* = *n4* are abbreviated as a single number.

Following each of these lines come all the lines that are affected in the first file flagged by **<**; then all the lines that are affected in the second file flagged by **>**.

The **-b** option causes trailing blanks (spaces and tabs) to be ignored and other strings of blanks to compare equal.

The **-e** option produces a script of *a*, *c*, and *d* commands for the editor *ed*, which will recreate *file2* from *file1*. The **-f** option produces a similar script, not useful with *ed*, in the opposite order. In connection with **-e**, the following shell program may help maintain multiple versions of a file. Only an ancestral file (*\$1*) and a chain of version-to-version *ed* scripts (*\$2*,*\$3*,...) made by *diff* need be on hand. A "latest version" appears on the standard output.

```
(shift; cat $*; echo `1,$p`) | ed - $1
```

Except in rare circumstances, *diff* finds a smallest sufficient set of file differences.

Option **-h** does a fast, half-hearted job. It works only when changed stretches are short and well separated, but does work on files of unlimited length. Options **-e** and **-f** are unavailable with **-h**.

FILES

`/tmp/d?????`  
`/usr/lib/diffh` for **-h**

SEE ALSO

`cmp(1)`, `ed(1)`

## COMMANDS

235-700-200  
November 1998

### *DIFF(1)*

## DIAGNOSTICS

Exit status is 0 for no differences, 1 for some differences, 2 for trouble.

## BUGS

Editing scripts produced under the **-e** or **-f** option are naive about creating lines consisting of a single period (.).

235-700-200  
November 1998

COMMANDS

*DIRNAME(1)*

NAME

dirname

DESCRIPTION

*See basename.*



NAME

`dls`um — “date-less” sum and block counts of a file

SYNOPSIS

`dls`um [-r] file

DESCRIPTION

*Dls*um calculates and prints a 16-bit checksum of the named file excluding the time-date stamp and header version information. A file block count is also printed. *Dls*um is typically used to check the "object" content of a file regardless of header and version time stamping.

SEE ALSO

`sum`(1)

DIAGNOSTICS

*Read error* is indistinguishable from end-of-file on most devices; check the block counts.



NAME

`du` — summarize disk usage

SYNOPSIS

`du [ -ars ] [ names ]`

DESCRIPTION

*Du* gives the number of blocks contained in all files and (recursively) directories within each directory and file specified by the *names* argument. The block count includes the indirect blocks of the file. If *names* is missing, `.` is used.

The optional argument `-s` causes only the grand total (for each of the specified *names*) to be given. The optional argument `-a` causes an entry to be generated for each file. Absence of either causes an entry to be generated for each directory only.

*Du* is normally silent about directories that cannot be read, files that cannot be opened, etc. The `-r` option will cause *du* to generate messages in such instances.

A file with two or more links is only counted once.

BUGS

If the `-a` option is not used, nondirectories given as arguments are not listed.

If there are too many distinct linked files, *du* will count the excess files more than once.

Files with holes in them will get an incorrect block count.



NAME

`echo` — echo arguments

SYNOPSIS

`echo` [ arg ] ...

DESCRIPTION

*Echo* writes its arguments separated by blanks and terminated by a new-line on the standard output. It also understands C-like escape conventions; beware of conflicts with the shell's use of `\`:

<code>\b</code>	backspace
<code>\c</code>	print line without new-line
<code>\f</code>	form-feed
<code>\n</code>	new-line
<code>\r</code>	carriage return
<code>\t</code>	tab
<code>\\</code>	backslash
<code>\n</code>	the 8-bit character whose ASCII code is the 1-, 2-, or 3-digit octal number <i>n</i> , which must start with a zero.

*Echo* is useful for producing diagnostics in command files and for sending known data into a pipe.

SEE

`sh(1)`



NAME

ed, red — text editor

SYNOPSIS

ed[ - ] [ file ]  
red[ - ] [ file ]

DESCRIPTION

*Ed* is the standard text editor. If the *file* argument is given, *ed* simulates an *e* command (see below) on the named file; that is to say, the file is read into *ed*'s buffer so that it can be edited. The optional *-* suppresses the printing of character counts by *e*, *r*, and *w* commands; of diagnostics from *e* and *q* commands; and of the *!* prompt after a *!* *shell command*. *Ed* operates on a copy of the file it is editing; changes made to the copy have no effect on the file until a *w* (write) command is given. The copy of the text being edited resides in a temporary file called the *buffer*. There is only one buffer.

*Red* is a restricted version of *ed*. It will only allow editing of files in the current directory. It prohibits executing shell commands via *!* *shell command*. Attempts to bypass these restrictions result in an error message ( *restricted shell* ).

Both *ed* and *red* support a formatting capability. After including a format specification as the first line of *file* and invoking *ed* with your terminal in "**stty -tabs**" or "**stty tab3**" mode [see *stty* (1)], the specified tab stops will automatically be used when scanning *file*. For example, if the first line of a file contained:

```
<:t5,10,15 s72:>
```

tab stops would be set at columns 5, 10, and 15, and a maximum line length of 72 would be imposed.

**Note:** While inputting text, tab characters when typed are expanded to every eighth column as is the default.

Commands to *ed* have a simple and regular structure: zero, one, or two *addresses* followed by a single-character *command*, possibly followed by parameters to that command. These addresses specify one or more lines in the buffer. Every command that requires addresses has default addresses, so that the addresses can very often be omitted.

In general, only one command may appear on a line. Certain commands allow the input of text. This text is placed in the appropriate place in the buffer.

While *ed* is accepting text, it is in *input mode*. In this mode, *no* commands are recognized; all input is merely collected. Input mode is left by typing a period (".") alone at the beginning of a line.

*Ed* supports a limited form of *regular expression* notation; regular expressions are used in addresses to specify lines and in some commands (for example, *s*) to specify portions of a line that are to be substituted. A regular expression (RE) specifies a set of character strings. A member of this set of strings is said to be *matched* by the RE.

The REs allowed by *ed* are constructed as follows:

## ED(1)

The following *1-character REs* match a *single* character:

- 1.1 An ordinary character (*not* one of those discussed in 1.2) is a 1-character RE that matches itself.
- 1.2 A backslash ( \ ) followed by any special character is a 1-character RE that matches the special character itself. The special characters are:
  - a. ".", \*, [, and \ (period, asterisk, left-square bracket, and backslash, respectively) which are always special, *except* when they appear within square brackets ( [] ); see 1.4).
  - b. ^ (caret or circumflex) which is special at the *beginning* of an *entire* RE (see 3.1 and 3.2), or when it immediately follows the left of a pair of square brackets ( [] ) (see 1.4).
  - c. \$ (currency symbol) which is special at the *end* of an entire RE (see 3.2).
  - d. The character used to bound (that is, delimit) an entire RE, which is special for that RE [for example, see how slash ( / ) is used in the *g* command, below].
- 1.3 A period ( "." ) is a 1-character RE that matches any character except new line.
- 1.4 A nonempty string of characters enclosed in square brackets ( [] ) is a 1-character RE that matches *any one* character in that string. If, however, the first character of the string is a circumflex ( ^ ), the 1-character RE matches any character *except* new line and the remaining characters in the string. The ^ has this special meaning *only* if it occurs first in the string. The minus ( - ) may be used to indicate a range of consecutive ASCII characters; for example, [0-9] is equivalent to [0123456789] . The - loses this special meaning if it occurs first (after an initial , if any) or last in the string. The right-square bracket ( ] ) does not terminate such a string when it is the first character within it (after an initial , if any); for example, [ ]a-f] matches either a right-square bracket ( ] ) or one of the letters **a** through **f** inclusive. The four characters listed in 1.2.a above stand for themselves within such a string of characters.

The following rules may be used to construct *REs* from 1-character REs:

- 2.1 A 1-character RE is an RE that matches whatever the 1-character RE matches.
- 2.2 A 1-character RE followed by an asterisk ( \* ) is an RE that matches *zero* or more occurrences of the 1-character RE. If there is any choice, the longest leftmost string that permits a match is chosen.
- 2.3 A 1-character RE followed by \ { *m* \ } , \ { *m*, \ } , or \ { *m*, *n* \ } is an RE that matches a *range* of occurrences of the 1-character RE. The values of *m* and *n* must be nonnegative integers less than 256; \ { *m* \ } matches *exactly m* occurrences; \ { *m*, \ } matches *at least m* occurrences; \ { *m*, *n* \ } matches *any number* of occurrences *between m* and *n*, inclusive.

Whenever a choice exists, the RE matches as many occurrences as possible.

- 2.4 The concatenation of REs is an RE that matches the concatenation of the strings matched by each component of the RE.
- 2.5 An RE enclosed between the character sequences `\ (` and `\ )` is an RE that matches whatever the unadorned RE matches.
- 2.6 The expression `n` matches the same string of characters as was matched by an expression enclosed between `\ (` and `\ )` earlier in the same RE. Here `n` is a digit; the subexpression specified is that beginning with the `n`th occurrence of `\ (` counting from the left. For example, the expression `^(. *)\1$` matches a line consisting of two repeated appearances of the same string.

Finally, an *entire RE* may be constrained to match only an initial segment or final segment of a line (or both):

- 3.1 A circumflex (^) at the beginning of an entire RE constrains that RE to match an *initial* segment of a line.
- 3.2 A currency symbol (\$) at the end of an entire RE constrains that RE to match a *final* segment of a line.

The construction `^entire\f RE$` constrains the entire RE to match the entire line.

The null RE (for example, `//`) is equivalent to the last RE encountered. Also see the last paragraph before *FILES*.

To understand addressing in *ed* it is necessary to know that at any time there is a *current line*. Generally speaking, the current line is the last line affected by a command; the exact effect on the current line is discussed under the description of each command. *Addresses* are constructed as follows:

1. The character `.` addresses the current line.
2. The character `$` addresses the last line of the buffer.
3. A decimal number `n` addresses the `n`th line of the buffer.
4. `ˆx` addresses the line marked with the mark name character `x`, which must be a lowercase letter. Lines are marked with the `k` command described below.
5. An RE enclosed by slashes (`/`) addresses the first line found by searching *forward* from the line *following* the current line toward the end of the buffer and stopping at the first line containing a string matching the RE. If necessary, the search wraps around to the beginning of the buffer and continues up to and including the current line, so that the entire buffer is searched. Also, read the last paragraph before *FILES*.
6. An RE enclosed in question marks (`?`) addresses the first line found by searching *backward* from the line *preceding* the current line toward the beginning of the buffer and stopping at the first line containing a string matching the RE. If necessary, the search wraps around to the end of the buffer and continues up to and including the current line. Also, read the last paragraph before *FILES*.

*ED(1)*

7. An address followed by a plus sign (+) or a minus sign (-) followed by a decimal number specifies that address plus (respectively minus) the indicated number of lines. The plus sign may be omitted.
8. If an address begins with + or -, the addition or subtraction is taken with respect to the current line, for example, -5 is understood to mean "-.5."
9. If an address ends with + or -, then 1 is added to or subtracted from the address, respectively. As a consequence of this rule and of rule 8, the address - refers to the line preceding the current line. (To maintain compatibility with earlier versions of the editor, the character in addresses is entirely equivalent to -.) Moreover, trailing + and - characters have a cumulative effect, so — refers to the current line less 2.
10. For convenience, a comma (,) stands for the address pair 1,\$, while a semicolon (;) stands for the pair ".,\$."

Commands may require zero, one, or two addresses. Commands that require no addresses regard the presence of an address as an error. Commands that accept one or two addresses assume default addresses when an insufficient number of addresses is given; if more addresses are given than such a command requires, the last one(s) are used.

Typically, addresses are separated from each other by a comma (,). Also, they may be separated by a semicolon (;). In the latter case, the current line (".") is set to the first address, and only then is the second address calculated. This feature can be used to determine the starting line for forward and backward searches (see rules 5. and 6.). The second address of any 2-address sequence must correspond to a line that follows, in the buffer, the line corresponding to the first address.

In the following list of *ed* commands, the default addresses are shown in parentheses. The parentheses are *not* part of the address; they show that the given addresses are the default.

It is generally illegal for more than one command to appear on a line. However, any command (except **e**, **f**, **r**, or **w**) may be suffixed by **l**, **n**, or **p**. In which case, the current line is either listed, numbered or printed, respectively, as discussed below under the **l**, **n**, and **p** commands.

- (.)**a** <text> . The **append** command reads the given text and appends it after the addressed line; "." is left at the last inserted line, or, if there were none, at the addressed line. Address 0 is legal for this command; it causes the "appended" text to be placed at the beginning of the buffer. The maximum number of characters that may be entered from a terminal is 256 per line (including the new-line character).
- (.)**c** <text> . The **change** command deletes the addressed lines and then accepts input text that replaces these lines. The "." is left at the last line input, or, if there were none, at the first line that was not deleted.
- (.,)**d** . The **delete** command deletes the addressed lines from the buffer. The line after the last line deleted becomes the current line; if the lines deleted were originally at the end of the buffer, the new last line becomes the current line.
- e** *file* . The **edit** command causes the entire contents of the buffer to be

deleted and then the named file to be read in. The "." is set to the last line of the buffer. If no file name is given, the currently remembered file name, if any, is used (see the **f** command). The number of characters read is typed; *file* is remembered for possible use as a default file name in subsequent **e**, **r**, and **w** commands. If *file* is replaced by **!**, the rest of the line is taken to be a shell [*sh*(1)] command whose output is to be read. Such a shell command is *not* remembered as the current file name. See *DIAGNOSTICS* .

**E** *file* The **E**dit command is like **e**, except that the editor does not check to see if any changes have been made to the buffer since the last **w** command.

**f** *file* If *file* is given, the **f**ile-name command changes the currently remembered file name to *file*; otherwise, it prints the currently remembered file name.

**(1,\$)g/RE** *command list*

In the **g**lobal command, the first step is to mark every line that matches the given RE. Then, for every such line, the given *command list* is executed with "." initially set to that line. A single command or the first of a list of commands appears on the same line as the global command. All lines of a multiline list except the last line must be ended with a **e**; **a**, **i**, and **c** commands and associated input are permitted; the "." terminating input mode may be omitted if it would be the last line of the *command list*. An empty *command list* is equivalent to the **p** command. The **g** , **G** , **v** , and **V** commands are *not* permitted in the *command list*. See also *BUGS* and the last paragraph before *FILES*.

**(1,\$)G/RE** In the interactive **G**lobal command, the first step is to mark every line that matches the given RE. Then, for every such line, that line is printed, the "." is changed to that line, and anyone command (other than one of the **a**, **c**, **i**, **g**, **G** , **v**, and **V** commands) may be input and is executed.

After the execution of that command, the next marked line is printed, and so on; a new line acts as a null command; an **&** causes the most recent command to be executed again within the current invocation of **G**. Note that the commands input as part of the execution of the **G** command may address and affect *any* lines in the buffer. The **G** command can be terminated by an interrupt signal (ASCII DEL or BREAK).

**h** The **h**elp command gives a short error message that explains the reason for the most recent ? diagnostic.

**H** The **H**elp command causes *ed* to enter a mode in which error messages are printed for all subsequent ? diagnostics. It will also explain the previous ? if there was one. The **H** command alternately turns this mode on and off; it is initially off.

**(.)i <text>** . The **i**nsert command inserts the given text before the addressed line; "." is left at the last inserted line, or, if there were none, at the addressed line. This command differs from the **a** command only in the placement of the input text. Address 0 is not legal for this command. The maximum number of characters that may be entered from a terminal is 256 per line (including the new-line character).

ED(1)

- (.,+1)j      The **join** command joins contiguous lines by removing the appropriate new-line characters. If exactly one address is given, this command does nothing.
- (.)kx        The **mark** command marks the addressed line with name *x*, which must be a lowercase letter. The address *ˆx* then addresses this line; "." is unchanged.
- (.,)l        The **list** command prints the addressed lines in an unambiguous way: a few nonprinting characters (for example, *tab*, *backspace*) are represented by (hopefully) mnemonic overstrikes, all other non-printing characters are printed in octal, and long lines are folded. An **l** command may be appended to any other command other than **e**, **f**, **r**, or **w**.
- (.,)ma      The **move** command repositions the addressed line(s) after the line addressed by **a**. Address 0 is legal for **a** and causes the addressed line(s) to be moved to the beginning of the file; it is an error if address **a** falls within the range of moved lines; "." is left at the last line moved.
- (.,)n        The **number** command prints the addressed lines, preceding each line by its line number and a tab character; "." is left at the last line printed. The **n** command may be appended to any other command other than **e**, **f**, **r**, or **w**.
- (.,)p        The **print** command prints the addressed lines; "." is left at the last line printed. The **p** command may be appended to any other command other than **e**, **f**, **r**, or **w**; for example, **dp** deletes the current line and prints the new current line.
- P**            The editor will prompt with **a\*** for all subsequent commands. The **P** command alternately turns this mode on and off; it is initially off.
- q**            The **quit** command causes *ed* to exit. No automatic write of a file is done (but see *DIAGNOSTICS*).
- Q**            The editor exits without checking if changes have been made in the buffer since the last **w** command.
- (\$)**r** *file*    The **read** command reads in the given file after the addressed line. If no file name is given, the currently remembered file name, if any, is used (see **e** and **f** commands). The currently remembered file name is *not* changed unless *file* is the very first file name mentioned since *ed* was invoked. Address 0 is legal for **r** and causes the file to be read at the beginning of the buffer. If the read is successful, the number of characters read is typed; the "." is set to the last line read in. If *file* is replaced by **!**, the rest of the line is taken to be a shell [*sh*(1)] command whose output is to be read. For example, "**Sr** **lls**" appends current directory to the end of the file being edited. Such a shell command is *not* remembered as the current file name.
- (.,)s/RE/replacement/ or  
(.,)s/RE/replacement/g  
The **substitute** command searches each addressed line for an occurrence of the specified RE. In each line in which a match is found, all (nonoverlapped) matched strings are replaced by the *replacement* if the global replacement indicator **g** appears after the command. If the global indicator does not appear, only the first occurrence of the matched string is replaced. It is an error for the substitution to fail on *all* addressed lines. Any character other than

space or new line may be used instead of */* to delimit the RE and the *replacement*; "." is left at the last line on which a substitution occurred. See also the last paragraph before *FILES*.

An ampersand (&) appearing in the *replacement* is replaced by the string matching the RE on the current line. The special meaning of & in this context may be suppressed by preceding it by \. As a more general feature, the characters \n, where *n* is a digit, are replaced by the text matched by then the regular subexpression of the specified RE enclosed between \ ( and \). When nested parenthesized subexpressions are present, *n* is determined by counting occurrences of \ ( starting from the left. When the character % is the only character in the *replacement*, the *replacement* used in the most recent substitute command is used as the *replacement* in the current substitute command. The % loses its special meaning when it is in a replacement string of more than one character or is preceded by a \.

A line may be split by substituting a new-line character into it. The new line in the *replacement* must be escaped by preceding it by [ ]. Such substitution cannot be done as part of **ag** or **v** command list.

(.,)ta This command acts just like the **m** command, except that a copy of the addressed lines is placed after address *a* (which may be 0); "." is left at the last line of the copy.

**u** The undo command nullifies the effect of the most recent command that modified anything in the buffer, namely the most recent **a**, **c**, **d**, **g**, **i**, **j**, **m**, **r**, **s**, **t**, **v**, **G**, or **V** command.

(1,\$)v/RE/command list This command is the same as the global command **g** except that the *command list* is executed with "." initially set to every line that does not match the RE.

(1,\$)V/RE/ This command is the same as the interactive global command **G** except that the lines that are marked during the first step are those that do not match the RE.

(1,\$)wfile The write command writes the addressed lines into the named file. If the file does not exist, it is created with mode 666 (readable and writable by everyone), unless your *umask* setting [see *sh* (1)] dictates otherwise. The currently remembered file name is not changed unless *file* is the very first file name mentioned since *ed* was invoked. If no file name is given, the currently remembered file name, if any, is used (see **e** and **f** commands); the "." is unchanged. If the command is successful, the number of characters written is typed. If *file* is replaced by !, the rest of the line is taken to be a shell [*sh*(1)] command whose standard input is the addressed lines. Such a shell command is not remembered as the current file name.

(\$)= The line number of the addressed line is typed; the "." is unchanged by this command.

!shell command

The remainder of the line after the ! is sent to the *UNIX* system shell [*sh*(1)] to be interpreted as a command. Within the text of that command, the unescaped character % is replaced with the

*ED(1)*

remembered file name; if *a!* appears as the first character of the shell command, it is replaced with the text of the previous shell command. Thus, *!!* will repeat the last shell command. If any expansion is performed, the expanded line is echoed; "." is unchanged.

**(.+1)<new line>**

An address alone on a line causes the addressed line to be printed. A new line alone is equivalent to ".+1p"; it is useful for stepping forward through the buffer.

If an interrupt signal (ASCII DEL or BREAK) is sent, *ed* prints a ? and returns to its command level.

Some size limitations: 512 characters per line, 256 characters per global command list, 64 characters per file name, and 128K characters in the buffer. The limit on the number of lines depends on the amount of user memory: each line takes one word.

When reading a file, *ed* discards ASCII NUL characters and all characters after the last new line. Files (for example, *a.out*) that contain characters not in the ASCII set (bit 8 on) cannot be edited by *ed*.

If the closing delimiter of an RE or of a replacement string (for example, */*) would be the last character before a new line, that delimiter may be omitted, in which case the addressed line is printed. The following pairs of commands are equivalent:

```
s/s1/s2 s/s1/s2/p
g/s1    g/s1/p
?s1     ?s1?
```

## FILES

*/tmp/e#* temporary; # is the process number.  
*ed.hup* work is saved here if the terminal is hung up.

## DIAGNOSTICS

? for command errors.  
? *file* for an inaccessible file.  
(use the **help** and **Help** commands for detailed explanations).

If changes have been made in the buffer since the last **w** command that wrote the entire buffer, *ed* warns the user if an attempt is made to destroy *ed*'s buffer via the **ee** or **oq** commands; it prints ? and allows one to continue editing. A second **oq** command at this point will take effect. The - command-line option inhibits this feature.

SEE ALSO

grep(1), sed(1), sh(1), stty(1)

CAVEATS AND BUGS

A! command cannot be subject to **ag** or **av** command.

The! command and the ! escape from the **e**, **r**, and **w** commands cannot be used if the editor is invoked from a restricted shell [see *sh(1)*].

The sequence **\n** in an RE does not match a new-line character.

The! command mishandles DEL.

Characters are masked to 7 bits on input.



NAME

`edobj` — object file editor

SYNOPSIS

`edobj` [`object[s]`]

DESCRIPTION

*Edobj* is a COFF (Common Object File Format) file dumper and editor. *Edobj* performs three separate functions. It is an interactive dumper used for dumping file header, section header, optional header, relocation, line number, and symbol table information. It is an overwriter used for overwriting the actual hexadecimal values in the file with new values. It is a file editor used to change and update the values in the file header, section header, optional header, relocation, line number, and symbol table information.

One or more COFF files may be specified on the command line. The interactive dumper is entered by executing the *edobj* command. It provides the ability to dump out any portion of the COFF file in a readable format. The overwriter is accessible through the interactive dumper by executing the *O* command. The file editor is entered from the interactive dumper by executing the *E* command. The interactive dumper options are:

**!** <command>

The *UNIX* system command is to be executed.

**?** The status of global variables is output.

**a** <symbol> All of the occurrences of *symbol* are searched for in the symbol table and output.

**b** [number] The number base is reset. If *number* is not specified, then the input number base is reset to the C standard. *Number* can be 8 for octal, 10 for decimal, and 16 for hexadecimal.

**c** The current archive member is closed. The next archive member is opened.

**d** <sclass> *Sclass* is a storage class. All of the symbols in the symbol table with the desired storage class are dumped out.

**e** <type> All of the symbol table entries of a given *type* are output.

**E** The object file editor is entered.

**F** <filename>

A search is performed to see if *filename* is a source file of the current object file.

**f** [ofilename]

The file header information is output. If an existing object file, *ofilename*, is specified, it becomes the currently opened file. Its file header information is then output.

**g** The current object file is closed, and the next file specified on the command line is opened. If the current file is the only or last file specified on the command line, then the current file remains open.

**h** [section] [secndx] [\*]

This command dumps out section header information. A "\*" dumps

## EDOBJ(1)

all of the section header information in the current file. A particular section header is output by specifying a section name, *section* or a section index number, *secndx*. If no specifications are made, the last section header specified will be output.

**l [function] [symndx] [\*]**

Line number information is output. A "\*" outputs all the line number entries. When a function is specified, the information for that function is output. If a symbol index is given, *symndx*, the line number information for the referenced symbol index is output. When the *l* command alone is given, the line number information for the last function specified will be output.

**m [mcommand]**

The command line menu is output. If a particular command is specified, only the menu information for that particular command is output.

**n [count]** The symbol table information for the next *count* symbols is produced. If *count* is not specified it is assumed to be one.

**o [a] [o] [p]** If no additional information is specified, then the entire optional header and patch list is dumped out. If *a* is specified, then the a.out header is output. *o* specifies the a.out with a library or pfile header. The patch list alone is specified by *p*.

**O [address] [symbol ['line]]**

The overwriter used to overwrite binary information in the object file is invoked. When specified, the overwriter will start at the given *address*, or at the address of the given *symbol*. If *symbol* is a function name, and *line* is an actual line number within that function, the overwriter can be specified to start at the address of a given line within a function.

**r [section] [secndx] [\*]**

With the "\*", all relocation information is dumped out. A specified *section*, dumps out the relocation information for a particular section. *Secndx*, a symbol index, outputs the relocation information for the referenced symbol index. If the *r* command alone is executed, the relocation information for the last specified section is output.

**s [section [start][:end] ]**

Raw section data is dumped out. A *section* can be specified to output the raw data of a particular section. A range of section data can be specified. If a section is not specified, then the raw section data for the last specified section will be dumped out.

**t [symbol [symbol]] [symndx [symndx]] [\*]**

The symbol table information is output. If no options are specified, then the information for the last specified symbol is output. A "\*" outputs all symbol table entries. *Symbol* is used to specify a particular symbol. A range of information, from one symbol to another may be output by specifying two *symbols* or symbol indexes, *symndx*.

**v [symbol [type]] [symndx [type]]**

The value of a particular symbol is output. If *symbol* is specified,

then the value of the named entry is output. A *type* may be specified. The *v* command alone will output the value of the last symbol specified.

**V** The version of **edobj** executing is printed along with the dmert generic object format it operates on.

**q** The **edobj** session is terminated.

The overwriter provides the ability to make changes to the currently opened object file by overwriting the binary information in the file with new values. The overwrites can only be performed on codes found within sections. The overwriter commands are:

**\n** The next 16 bytes after the current line are displayed.

**-** The previous 16 bytes before the current output line are displayed.

**! <command>**  
The *UNIX* system command is executed.

**?** The status and content of certain global variables are output.

**m** The program menu of the overwriter is listed.

**o <[address] [symbol ['line']]>**  
Change the current address to *address*, and then output 16 bytes starting at the new address position. If a *symbol* is specified, the current address is changed to the start of the symbol, and the 16 bytes are displayed. If *symbol* is a function name and *line* is an actual line number within that function, the current address is changed to the location of the line within the function and the 16 bytes are displayed.

**s <[hexnumber] ['char']>**  
A pattern is searched for starting at the current address. The pattern is made up of hexadecimal numbers or their character equivalent. If the pattern is found, the file pointer points to the start of the found pattern. The 16 bytes starting at this position are listed out. In each search, one byte of information must be specified. A hexadecimal value being searched for must start at the first half of a byte.

**u** This command will undo the last change made to the file.

**hexnumber** A *hexnumber* can be any hexadecimal number. These hexadecimal numbers will replace any of the displayed 16 bytes starting at the current address location. The desired changes are entered directly under the hexadecimal values to be changed. If a hexadecimal value is overwritten with a single space, ' ', the old value remains unchanged. At most, 16 bytes can be changed at any one time.

**\tcharacters** The tab specifies the use of characters. The input characters will replace the displayed characters and the hexadecimal representation of the characters within the object file. The replacement starts at the current address position. At most, 16 characters can be changed at any one time. If the user does not wish to overwrite a particular character, the character '.' is used to bypass the overwrite of that character.

**q** The *overwriter* is exited, and control is turned back to the main portion of *edobj*.

**EDOBJ(1)**

The file editor portion of *edobj* provides the capability of making quick and easy changes to the COFF headers, relocation information, line number information, and symbol table entries. The portion of the COFF file to change is specified. The changes are then prompted for. The time/date stamp in the file header is the only noneditable portion of the COFF file. The file editing commands are:

- ! <command>** The *UNIX* system *command* is to be executed.
- f** The current file header is edited.
- o [otype]** The optional file header is edited. The *otype* is *a* for an a.out header. It is *l* for a library header. It is *h* for a pfile header. The patch list is specified by *p*. A *u* specifies an update header. A *q* specifies a qfile header.
- h [[sectnum] [sectname]]**  
Editing is performed on a section header of the current file. The section number, *sectnum*, or the name of the section, *sectname*, must be specified to determine which particular section header to edit.
- l [[symndx] [funcname]]**  
Editing is performed on line number entries. The function the change is made in can be specified by its symbol table index, *symndx*, or the function name, *funcname*. The actual line entry to edit in the section is prompted for.
- m** A menu of the available commands is output.
- r [[sectnum] [sectname]]**  
A relocation entry in the current file is edited. The section which contains the relocation entry can be specified by *sectname*, the name of the current section, or *sectnum*, the number of the section. The individual relocation entry to be changed is prompted for.
- s [[sindex] [sname]]**  
An entry in the symbol table is edited. The symbol table index, *sindex*, or symbol name, *sname*, of the entry to change must be specified.
- V** The version of **edobj** executing is printed along with the dmert generic object format it operates on.
- q** The object file editing is terminated. Control is returned to the main portion of *edobj*.

## CAVEATS

The tool cannot distinguish between a qfile optional header and an update optional header. The user must know if he or she has a qfile or an update optional header.

When searching for a section header or symbol name, if more than one entry exists with that name, the first section header or symbol table entry with that name will always be retrieved.

A break executed in the file editing portion of *edobj* will terminate *edobj*.

Editing changes which expand or contract the COFF file will have side effects on other portions of the COFF file. These types of changes should be avoided.



NAME

EMACS — interactive screen editor

SYNOPSIS

**emacs** [-l *init\_file*] [+*line\_number*] [*file*]

DESCRIPTION

EMACS is an interactive screen editor which can be used to construct and edit files on the *UNIX* system. A window of text from the file being edited is displayed on the terminal screen, along with a status line describing the editor version and file being edited. Ordinary characters typed are inserted in the file, while escape sequences and control characters are used to invoke editing functions. Several files can be edited at the same time in different editing buffers, and two of the active buffers can be displayed on the same screen.

If given a *file* argument, EMACS reads the file into the buffer "Main". Otherwise, an empty buffer is created. If a *line\_number* is given, EMACS moves to that line number in the specified file. If an *init\_file* is specified, EMACS will treat the contents of that file as EMACS commands to be executed on startup. EMACS also looks in your home directory for a file named *.emacs\_init* and interprets commands from it before those in the specified *init\_file* are executed. This option can also be specified with *.i*, in which case it suppresses the processing of *.emacs\_init*. If an *.emacs\_init* does not exist, EMACS looks in the EMACS library directory (see below) for a standard *.emacs\_init*.

EMACS can be customized by the user through user-defined macro commands, which can redefine the effect of the basic editing commands. EMACS has a number of built in editing modes that customize some of the commands to support editing of particular types of files, such as C source programs or word processing source documents.

There are a number of self-help features in EMACS to aid in learning how to use the editor. Complete documentation appears in Tab 4 (EMACS Description).

MISCELLANEOUS CONVENTS

EMACS treats the following characters specially in filenames:

- \$NAME**     Substitute the environment variable NAME
- USER**     Substitute the home directory of USER.
- EMACS**    Specifies the EMACS library directory. (contains standard macros, etc.)
- \*,?**        Can be used for matching in partially specified filenames.
- 'COMMAND'**   Substitute the output of running COMMAND.
- ^X**        Control-X, where X is any character. These characters are input by holding down the control key and another key simultaneously on your terminal. EMACS also provides a special mode (controlify) to allow you to input control characters that your terminal cannot send to your system (see the discussion of modes).

*EMACS(1)*

{[ Used as they are with the shell. EMACS uses the following notation to display and input nonprinting characters:

Some of the control characters displayed are not very intuitive:

Some of the control characters displayed are not very intuitive:

^?	Rubout
^[	Excape
^]	Control-right-bracket
^\	Control-backslash
^_	Control underline
^@	Null
^H	Backspace (Displays as ^H when backspace mode is off, see below)
^I	Tab (Displays as a ^I when notabs mode is on)
^J	Newline (Displays as ^J in searches)
^M	Carriage return
M-x	Meta-x, where x is any character (including a control character). Meta characters are entered by typing escape followed by another character.

Many EMACS commands take an optional numeric argument. The argument for a command precedes the command itself, and can be specified a number of ways. Typing ^U specifies an argument of 4, or 4 times the current argument. Typing escape followed by a sequence of digits with or without a leading '.' specifies a decimal value. Some examples:

<b>^U^U^N</b>	Go forward 16 lines
<b>M-123^N</b>	Go forward 123 lines
<b>M—12^N</b>	Go forward -12 lines (goes back 12 lines).

Some EMACS commands prompt for a string parameter, such as a file name. Some of the normal EMACS commands can be used to edit the parameter while it is being entered. These include: ^F, ^B, ^D, ^H, ^A, ^E, ^K, ^U. In addition, typing your kill character (usually @) deletes the string, typing ^X substitutes the contents of the current line in the buffer, and ^Y substitutes the current file name. (The latter is a very convenient way of finding files with similar names.) Typing ^V while typing a string will cause EMACS to expand any shell meta characters (\$, \*, ?, etc.) in the string and show the first candidate in the result. If there is more than one expansion (e.g., foo.c and bar.c for \*.c), ^N and ^P will allow you to move backwards or forwards in them.

## REGULAR EXPRESSIONS

EMACS uses an extension of the regular expression syntax used by **ed(1)** and **grep(1)** for regular expression searches and query replace. The following character sequences have special meaning:

.	Matches any single character except newline.
[xyz]	Matches one character among the set enclosed in brackets. (If the

first character is ^, it matches all but the specified characters.) If a '-' appears in the brackets, it designates a range of character values (i.e. [a-ez] is equivalent to [abcdez]). The sequence \n can be used to specify a newline as one of the alternatives.

- \* Matches 0 or more of the preceding expression (a single character, specified as such or \* or [ ])
- + Matches 1 or more of the preceding expression
- \{x,y\} Matches x through y occurrences of the preceding expression. If y is omitted, it defaults to 255. If x is omitted it defaults to 0.
- \(expr\)
- \(ex1\|ex2\)
- \< Matches 0 characters at the beginning of a word.
- \> Matches 0 characters at the end of a word.
- ^ Matches 0 characters at the beginning of a line
- \$ Matches 0 characters at the end of a line.
- \n Matches on newline at end of a line.

The following special sequences apply in the strings to replace with in query replace.

- & Specifies the entire string matched by the from string.
- % Specifies the previous To string.
- \n Specifies the string matched by the *n*th occurrence of \ (expr\)  
(Regular expression replace only).
- ^J or \n Specifies a newline is to be inserted at this point.

## COMMAND SUMMARY

The following chart summarizes the available commands by category. Some commands appear in more than one category. Commands that are marked with '\*' take a numeric argument that indicates how many times to do the command. Commands that are marked with '+' take a numeric argument that changes the behavior of the command in some other way.

### General Commands

- ^G** Quit (Stops commands that prompt for things)
- ^Z** Exit one level (Usually exits EMACS)
- ^X^C** Exit EMACS
- M-u** Undo. Undoes the result of the last modification
- M-?** Help - Prompts for a command and displays its function.
- M-w** Put a wall chart of command explanations in the current buffer
- ^L** Refresh the screen. (Argument indicates where to put the current line)

*EMACS(1)*

## Character Oriented Commands

<b>^F</b>	Move forward one character
<b>^B</b>	Move backward one character
<b>^D</b>	Delete the character under the cursor
<b>^H,^?</b>	Delete the previous character
<b>^T</b>	Transpose the current and next character, move forward.
<b>^C</b>	Capitalize the current character

## Word Oriented Commands

<b>M-f</b>	Move forward one word
<b>M-b</b>	Move backward one word
<b>M-d</b>	Delete forward one word
<b>M-^?,M-^h</b>	Delete backwards one word
<b>M-c</b>	Capitalize the next word
<b>M-<u></u></b>	Underline the next word
<b>M-a</b>	Move to the beginning of the sentence
<b>M-e</b>	Move to the end of the sentence

## Line Oriented Commands

<b>^A</b>	Move to beginning of line
<b>^E</b>	Move to end of line
<b>^M-&lt;</b>	Move to beginning of file
<b>^M-&gt;</b>	Move to end of file
<b>^P</b>	Move back one line
<b>^N</b>	Move forward one line
<b>M-g</b>	Go to the line specified by the argument
<b>^O</b>	Create a blank line in front of the cursor
<b>^J,^M</b>	Make a new line (Just moves through empty lines).
<b>^K</b>	Kill (delete) to the end of line (with argument, kills whole lines)

## Delete Commands

<b>^D</b>	Delete the character under the cursor
<b>^H,^?</b>	Delete the previous character
<b>M-d</b>	Delete forward one word
<b>M-^?,M-^h</b>	Delete backwards one word
<b>^K</b>	Kill (delete) to the end of line (with argument, kills whole lines)
<b>^W</b>	Delete the marked region (argument specifies a mark number)
<b>^Y</b>	Restore the last deletion (sets mark in front of it).

**M-Y** Replace the marked region with the previous deletion (Use only immediately after ^Y)

#### Display Commands

**^L** Redraw the screen  
**^V** Display the next page  
**M-v** Display the previous page  
**M-<** Move to beginning of file  
**M->** Move to end of file  
**M-^L** Redraw with the current line at the top of the screen

#### Buffer Commands

(Most prompt for a buffer name, entering return gets a list of active buffers).

**^X^B** Change working buffer  
**^X^F** Find file (does change buffer if file is in one, creates a new buffer and reads the file if not).  
**^X^K** Kill buffer  
**^X^N** Change buffer name (with argument, changes file name)  
**^X^T** Send region to buffer  
**^X=** Display statistics on buffer  
**^X2** Enter two window mode (prompts for buffer name for second window)  
**^X1** Make current window the only window  
**^X^O** Switch windows.

#### File Commands

**^X^R** Read file into current buffer (with an argument, inserts the file at the current position)  
**^X^W** Write buffer to file (With an argument, appends to the file)  
**^X^S** Save current buffer into current file.  
**^X^F** Find file (does change buffer if file is in one, creates a new buffer and reads the file if not).  
**^X^N** Change buffer name (with argument, changes file name)  
**^X^L** Load macros from file. (With an argument, undefines all previously defined macros.)

#### Region Commands

**M-** (Meta space) Set mark at position (argument is the mark number)  
**^X^X** Exchange mark and cursor position (argument is the mark number)  
**^W** Delete the region and put it on the kill stack  
**M-p** Put the marked region in the kill stack without deleting it.

*EMACS(1)*

## Search and Replace Commands

(All prompt for search and replace strings.)

- ^S,^R** Forward and reverse incremental search. (Both display the string currently matched. ^S moves to next occurrence, ^R moves to previous occurrence. ^H deletes last character, ^G quits search, escape exits search at currently displayed position. See below on regular expression search.
- M-^S** Regular expression search. (waits for whole expression to be typed). ^S following M-^S goes to next occurrence.
- M-r,M-^R** Ordinary and regular expression query replace. (Prompts at each occurrence of from string for the following:
- y** Replace with "to" string and move on.
  - n** Do not replace this occurrence and move on.
  - r** Replace all of the rest, showing each replacement.
  - R** Replace the rest silently
  - p** Move to previous occurrence of from string.
  - .** Replace this one and stop
  - <** Go back to the last replacement and stop
  - ^G** Quit query replace
  - <escape>** Prompt for new to string, and replace this occurrence with it.

## Window Commands

- ^X2** Enter two window mode (prompts for buffer name for second window)
- ^X1** Make current window the only window
- ^X^O** Switch windows.
- ^X^^** Make current window grow by one line.

## Special Input Commands

- ^Q** Takes the next input character and inserts it, even if it is a control character
- M-Q** Takes the next input character, makes it a meta character, and inserts it.
- M-\** Converts it's argument to an ascii character and inserts it.

Interaction with *UNIX*

- M-!** Prompt for a *UNIX* command and execute it (with an argument, passes the buffer as standard input.
- M-\$** Execute *UNIX* command, put output in buffer .exec
- ^X^D** Change working directory
- M-^M** Send the current buffer as mail. (Lines starting To: or Cc: are taken as destinations.)

### Miscellaneous Commands

- ^X^M** Specifies modes (See below)
- M-s** Prints EMACS statistics
- M-** Re-adjusts line boundaries in the whole buffer to fill lines evenly. (With an argument, it works only on the current region.)
- M-/** Start a C comment.

### MODES

Mode parameters allow you to customize the behavior of certain commands. Some modes are switches, indicating only that something is either off or on, while others are numeric parameters. Modes can be set by the **^X^M** command. Typing **^X^M** followed by the name of a switch mode turns it on, typing **^U^X^M** followed by the name turns it off. To set numeric modes, give the value you want as an argument to **^X^M**. (i.e. **M-500^X^Msave**). Modes can be set automatically by putting **^X^M** commands in your `.emacs_init`. Modes can also be attached to a file by putting the string "EMACS\_MODES: " followed by a list of mode settings in the first 10 lines of the file. (The mode settings can be preceded or followed by anything, to allow you to make them look like a comment to other software processing the file.) The mode settings are separated by commas and can be of the following form:

- modename** Set this switch mode
- !modename** Turn this switch mode off
- modename=x**  
Set this numeric parameter to x.

The following indicates the modes and their default settings. Switches are listed as either ON or OFF, while numeric parameters have specified values. Note that the system default `.emacs_init` may alter these settings on your local machine.

- save** (OFF) Automatically saves each buffer after savetype characters of input or when you change buffers or run commands
- savetype** (256) Number of characters between automatic saves
- mailtype** (100) Number of characters between checks for new mail
- c** (OFF) Automatically indents during typing for C programs
- verbose** (ON) provides prompts for meta and control-x commands.
- fill** (ON) Automatically replaces a space with a newline when you type past the end of line or past fillcol characters.
- fillcol** (72) column beyond which lines are wrapped.
- eofnl** (ON) Causes a newline to be appended to any file that doesn't end in one.
- end\_newline**  
(OFF) Causes attempts to move beyond the end of the file to add newlines.
- keepscroll** (0) Number of lines kept from previous screen during **^V** and **M-v**

*EMACS(1)*

<b>smoothscroll</b>	(OFF) Updates display via scrolling instead of jumping
<b>readonly</b>	(OFF) Prevents saving the current buffer
<b>picture</b>	(OFF) Enables 2-dimensional editing (See the manual)
<b>tabstop</b>	(8) Width of a tab character
<b>overwrite</b>	(OFF) Causes input to replace characters already there, rather than insert.
<b>nodelete</b>	(OFF) Causes deletions to replace the characters with whitespace rather than deleting them.
<b>rigid_newline</b>	(OFF) Causes newline to always insert a new line, even if the next line is empty.
<b>notabs</b>	(OFF) Causes tabs to be expanded to spaces on input, and tabs in files to display as ^I
<b>comcol</b>	(40) Column where the M-/ commands starts a comment.
<b>backspace</b>	(?) Causes backspaces to appear as cursor motion, not ^H. This mode is set ON if your terminal handles underscored characters, OFF otherwise.
<b>nobell</b>	(OFF) Causes EMACS not to ring the terminal bell on an error
<b>caseless</b>	(OFF) Causes all searches to ignore upper/lower case distinctions
<b>usilent</b>	(OFF) Causes output from <i>UNIX</i> commands run from EMACS to be discarded.
<b>noecho</b>	(OFF) Causes output from M-\$ commands not to be echoed.
<b>controlify</b>	(OFF) Causes a sequence of <i>ctl_char</i> followed by another character to translate into the second character made a control character.
<b>ctl_char</b>	(30) Prefix for controlify (This is an integer specifying the ascii code of the character, the default is ^^).
<b>Inumb</b>	(ON) Displays line numbers
<b>Inowid</b>	(4) Width of line numbers.
<b>time</b>	(OFF) Displays a clock
<b>display_percent</b>	(OFF) Displays current position as a percentage of the whole file.
<b>flow_lim</b>	(0) If non-zero, flow control will be enabled whenever <i>flow_lim</i> characters are sent to the terminal in a burst.
<b>height</b>	(?) Height of screen area for buffer display (set automatically)
<b>width</b>	(?) Width of screen
<b>tspeed</b>	(?) Describes your terminal to host speed.
<b>autoload</b>	(ON) Automatically loads macros when first referenced.
<b>leftmargin</b>	(0) In picture mode, this is the number of characters to the left of the screen.
<b>display_euc</b>	(?) Determines whether characters with the high order bit are

considered extended ASCII or EMACS meta-characters. (Set on if your terminal will display them)

## ENVIRONMENT

The variable MAIL is the file name that EMACS looks at for newly arrived mail. If your mail is forwarded to some other system, MAIL should not be exported. The environment variable MAILER optionally specifies the name of a mail command to use for sending mail. The environment variable SHELL specifies the shell to use to execute shell commands. If the environment variable TEMP4EMACS is set it is taken as the pathname of a directory in which EMACS should place its temporary files.

## FILES

\$HOME/.emacs\_init

EMACS/.emacs\_init

\$HOME/emacs[0-11]

EMACS/macros

EMACS/terminals

EMACS/helpfile

EMACS/errfile

The .emacs\_init files, if present, contain a standard set of initializations to be made when you start EMACS. The characters in the file will be interpreted as if you had entered them as commands from your terminal. The most common application of this is to set modes different from the default modes.

When EMACS is killed by an internal error, the **kill(1)** command, or by hanging up during an editing session, it saves your buffers in the files emacs0-emacs11 in your home directory. You will receive **mail(1)** notifying you of what buffers were saved.

The directory EMACS is the EMACS library, the location of which depends on your local installation. Pathnames starting with EMACS will be translated to the local path of this directory by EMACS (but *not* by other tools). EMACS/macros may contain macro packages that can be loaded. EMACS/terminals contains terminal information for EMACS. The other two files contain internal data for EMACS.



NAME

ENV — set environment for command execution

SYNOPSIS

**env** [ - ] [ name=value ] ... [ command args ]

DESCRIPTION

*Env* obtains the current *environment*, modifies it according to its arguments, then executes the command with the modified environment. Arguments of the form *name = value* are merged into the inherited environment before the command is executed. The - flag causes the inherited environment to be ignored completely so that the command is executed with exactly the environment specified by the arguments.

If no command is specified, the resulting environment is printed, one name-value pair per line.

SEE ALSO

sh(1)



NAME

*errport* — interface to craft output spooler

SYNOPSIS

***/bin/errport***

DESCRIPTION

The error logging process, *errport*, is a 2-process system that receives error messages from both supervisor and kernel processes and the kernel itself. *Errport* adds a control string if necessary and forwards the error messages to the output spooler process for output on one or more of the craft terminals or log files.

Interfaces for the kernel, kernel processes, and supervisor processes to the *errport* system are provided in library *errpt* (see *Volume 4, System Libraries*, Chapter 7, "Operating System Libraries").

The kernel formats and sends a message to the system port PT\_ERRLOG. The FIFO driver (a kernel process */prc/fda*) attaches to this port and receives these messages. Their text is stored in an area of low core (starting at physical address 0x9000). A *UNIX* process (*/bin/errport*) reads this area by trapping to the FIFO driver. The text of each message it finds is placed in a message destined for the craft spooler.

Messages sent to PT\_ERRLOG and received prior to a bootstrap of any kind will be saved across the bootstrap.

Both processes are created at boot time. The FIFO driver is *pcreated* at boot time and ULARP brings up the *UNIX* process. This process then opens the low core FIFO (*/dev/errport*), which establishes the communication between the two processes.



NAME

*expr* — evaluate arguments as an expression

SYNOPSIS

**expr** arguments

DESCRIPTION

**DESCRIPTION**

The arguments are taken as an expression. After evaluation, the result is written on the standard output. Terms of the expression must be separated by blanks. Characters special to the shell must be escaped. Note that **0** is returned to indicate a zero value, rather than the null string. Strings containing blanks or other special characters should be quoted. Integer-valued arguments may be preceded by a unary minus sign. Internally, integers are treated as 32-bit, 2s complement numbers.

The operators and keywords are listed below. Characters that need to be escaped are preceded by **\**. The list is in order of increasing precedence; with equal precedence operators grouped within **{ }** symbols.

*expr \ | expr* returns the first *expr* if it is neither null nor **0**; otherwise returns the second *expr*.

*expr \& expr* returns the first *expr* if neither *expr* is null nor **0**; otherwise returns **0**.

*expr { =, \>, \>=, \<=, != } expr* returns the result of an integer comparison if both arguments are integers; otherwise returns the result of a lexical comparison.

*expr { +, - } expr* addition or subtraction of integer-valued arguments.

*expr { \ \*, /, % } expr* multiplication, division, or remainder of the integer-valued arguments.

*expr : expr* The matching operator **:** compares the first argument with the second argument which must be a regular expression; regular expression syntax is the same as that of *ed* (1), except that all patterns are “anchored” (that is, begin with **^**) and, therefore, **^** is not a special character in that context. Normally, the matching operator returns the number of characters matched (**0** on failure). Alternatively, the **\(...\)** pattern symbols can be used to return a portion of the first argument.

EXAMPLES

1. `a='expr $a + 1'`  
  
adds 1 to the shell variable **a** .
2. `# 'For $a equal to either "/usr/abc/file" or just "file",  
expr $a : '.*\/ \ ( .* \ )' \ | $a`

**EXPR(1)**

returns the last segment of a path name (i.e., **file**). Watch out for / alone as an argument; *expr* will take it as the division operator (see BUGS).

3. # A better representation of example 2.

```
expr // $a : '*.*/ \ (* \ )'
```

The addition of the // characters eliminates any ambiguity about the division operator and simplifies the whole expression.

4. `expr $VAR : '*'`

returns the number of characters in *\$VAR* .

## SEE ALSO

ed(1), sh(1)

## EXIT CODE

As a side effect of expression evaluation, *expr* returns the following exit values:

- |   |   |
|---|---|
| 0 | if the expression is neither null nor 0 |
| 1 | if the expression <i>is</i> null or 0   |
| 2 | for invalid expressions.                |

## DIAGNOSTICS

*syntax error* for operator/operand errors  
*non-numeric argument* if arithmetic is attempted on such a string

## BUGS

After argument processing by the shell, *expr* cannot tell the difference between an operator and an operand except by the value. If **\$a** is an = , the command:

```
expr $ a = '='
```

looks like:

```
expr = = =
```

as the arguments are passed to *expr* (and they will all be taken as the = operator). The following works:

```
expr X$a = X=
```

NAME

`falloc` — allocate a contiguous file

SYNOPSIS

**falloc** file name size

DESCRIPTION

A contiguous file of the specified file name is allocated to be of 'size' (512-byte) blocks.

DIAGNOSTICS

The command complains that a needed directory is not searchable, the final directory is not writable, the file already exists, or there is not enough space for the file.



235-700-200  
November 1998

COMMANDS

*FALSE(1)*

NAME

*false*

DESCRIPTION

*See true.*



NAME

*fgrep* — See *grep*.

DESCRIPTION

See *grep*.



NAME

**find** — find files

SYNOPSIS

**find** path-name-list expression

DESCRIPTION

*Find* recursively descends the directory hierarchy for each path name in the *path-name-list* (i.e., one or more path names) seeking files that match a boolean *expression* written in the primaries given below. In the descriptions, the argument *n* is used as a decimal integer where *+n* means more than *n*, *-n* means less than *n*, and *n* means exactly *n*.

**-name** *file* True if *file* matches the current file name. Normal shell argument syntax may be used if escaped (watch out for [, ? and \*).

**-perm** *onum*  
True if the file permission flags exactly match the octal number *onum* [see *chmod* (1)]. If *onum* is prefixed by a minus sign, the flags are compared: (flags&onum)==onum.

**-type** *c* True if the type of the file is *c*, where *c* is **b**, **c**, **d**, **p**, or **f** for block special file, character special file, directory, fifo (a.k.a named pipe), or plain file.

**-links** *n* True if the file has *n* links.

**-user** *uname*  
True if the file belongs to the user *uname*. If *uname* is numeric and does not appear as a login name in the */etc/passwd* file, it is taken as a user ID.

**-group** *gname*  
True if the file belongs to the group *gname*. If *gname* is numeric and does not appear in the */etc/group* file, it is taken as a group ID.

**-size** *n* True if the file is *n* blocks long (512 bytes per block).

**-atime** *n* True if the file has been accessed in *n* days.

**-mtime** *n* True if the file has been modified in *n* days.

**-ctime** *n* True if the file has been changed in *n* days.

**-exec** *cmd* True if the executed *cmd* returns a zero value as exit status. The end of *cmd* must be punctuated by an escaped semicolon. A command argument { } is replaced by the current path name.

**-ok** *cmd* Like **-exec** except that the generated command line is printed with a question mark first, and is executed only if the user responds by typing **y**.

**-print** Always true; causes the current path name to be printed.

**-cpio** *device*  
Write the current file on *device* in *cpio* (1) format (5120-byte records).

**-newer** *file* True if the current file has been modified more recently than the argument *file*.

***FIND(1)***

(*expression*) True if the parenthesized expression is true (parentheses are special to the shell and must be escaped).

The primaries may be combined using the following operators (in order of decreasing precedence):

- 1) The negation of a primary ( **!** is the unary *not* operator).
- 2) Concatenation of primaries (the *and* operation is implied by the juxtaposition of two primaries).
- 3) Alternation of primaries ( **-o** is the *or* operator).

**EXAMPLE**

To remove all files named **a.out** or **\*.o** that have not been accessed for a week:

```
find / \( -name a.out -o -name '*.o' \) -atime +7 -exec rm {} \;
```

**FILES**

/etc/passwd,/etc/group

**SEE ALSO**

chmod(1), cpio(1), sh(1)

NAME

**fmove** — make a file multiextent

SYNOPSIS

**fmove** filename

DESCRIPTION

The filename specified is made into a multiextent file.

DIAGNOSTICS

The command complains a needed directory is not searchable, the final directory is not writable, the file does not already exist, the file is open, or there is not enough contiguous space for the file.



NAME

`fsaudit` — file system mount auditor

SYNOPSIS

`fsaudit` <SG data base name><partition name>[mount point]

DESCRIPTION

**DESCRIPTION**

The *fsaudit* command locates the given *partition name* in the given *SG data base*, and extracts the disk number, partition number, partition block count, and partition inode block count for that partition. This information is then used to initialize the named partition, mounting it on the optional *mount point*. If no mount point name is given, then `/dev/<partition name>` is used. *Fsaudit* basically does an incremental mount of the named file system. If the file system is mounted, *fsaudit* will return immediately. Otherwise, the file system is mounted on the named device file. If the device file does not exist, then it is created. *Fsaudit* is similar to *fsinit*, but this command will not destroy the contents of the file system.

DIAGNOSTICS

*Fsaudit* will fail with an error code if the named SG data base cannot be attached or read, or if the requested partition cannot be successfully mounted. Appropriate error messages are printed before *fsaudit* exits.

FILES

`/etc/fsaudit`

EXAMPLES

The command:

```
fsaudit /data base/appdmert.sg no5text /tmp/no5mtpt
```

will read the application SG data base (`appdmert.sg`), extract the information needed to mount `no5text`, and mount the `no5text` partition on `/tmp/no5mtpt`.

The command:

```
fsaudit /data base/appdmert.sg no5text
```

will act the same as the above command, but the `no5text` partition will be mounted on `/dev/no5text` by default.

CAVEATS

Note that *fsaudit* can not always mount a corrupted file system, in which case *fsinit* may have to be used (with discretion).

SEE ALSO

`fsinit(1)`



NAME

*fsdb* — file system debugger

SYNOPSIS

*/etc/fsdb* special [ - ]  
*/etc/fsdb* [-w] special [ - ]

DESCRIPTION

*Fsdb* can be used to patch up a damaged file system. It contains conversions which translate block and i-numbers into their corresponding disk addresses. It also includes mnemonic offsets which can be used to access different parts of an i-node.

These features simplify correcting control block entries or descending the file system tree.

**Fsdb** contains several error checking routines to verify i-node and block addresses. These can be disabled if necessary by invoking *fsdb* with the optional - argument or by the use of the 0 symbol. (*fsdb* reads the i-size entries from the super-block of the file system in order to perform these checks.)

The **-w** option enables *fsdb* to write to mounted file systems. Without **-w**, *fsdb* defaults to the read only mode. To patch a mounted file system, the *openwd* command must have been issued prior to *fsdb* (see *openwd* in this section). *Openwd* allows *fsdb* to open special device files for mounted file systems with read/write.

Numbers are considered decimal by default. Octal numbers must be prefixed with a zero. Hexadecimal numbers must be prefixed by either x or 0x and must be terminated by a colon. During any assignment operation, numbers are checked for a possible truncation error due to a size mismatch between source and destination.

Since *fsdb* reads a block at a time, it is able to handle both raw and block I/O. A buffer management routine retains commonly used blocks of data and reduces the number of read system calls. Some assignment operations result in a delayed write-through of the corresponding block.

These symbols are recognized by *fsdb*:

<b>i</b>	convert from i-number to i-node address
<b>b</b>	convert to block address
<b>d</b>	directory slot offset
<b>+, -, */</b>	address arithmetic
<b>q</b>	quit
<b>&gt;, &lt;</b>	save, restore an address
<b>=</b>	numerical assignment
<b>+=</b>	incremental assignment
<b>-=</b>	decremental assignment

*FSDB(1)*

=	character string assignment
O	error checking flip-flop
p	general print facilities
f	file print facility
B	byte mode
W	word mode
S	half-word mode
!	escape to shell.

The print facilities generate a formatted output in various styles. The current address is normalized to an appropriate boundary before printing begins. It advances with the printing and is left at the address of the last item printed. The output can be terminated at any time by typing the delete character. If a number follows the *p* symbol, that many entries are printed. A check is made to detect block boundary overflows since logically sequential blocks are generally not physically sequential. If a count of zero is used, all entries to the end of the current block are printed. These print options are available:

i	print as i-nodes
d	print as directories
o	print as octal half words
e	print as decimal words
c	print as characters
b	print as hexadecimal bytes
h	print as hexadecimal words.

The *f* symbol is used to print data blocks associated with the current i-node. (Blocks are numbered starting with zero.) The desired print option letter follows the block number or the *f* symbol. It checks for special devices and for nonzero block pointers.

Dots, tabs, and spaces may be used as function delimiters but are not necessary. A line which contains only a newline character will increment the current address by the size of the data type last printed. That is, the address is set to the next byte, word, half-word, directory entry or i-node, allowing the user to step through a region of a file system. Information is printed in a format appropriate to the data type. Bytes, words, and double words are displayed with the hexadecimal address followed by the value in hexadecimal and decimal. A .B or .S is appended to the address for byte and half-word values, respectively. Directories are printed as a directory slot offset followed by the decimal i-number and the character representation of the entry name. I-nodes are printed with the labeled fields describing each element. The following mnemonics are used for i-node examination and refer to the current working i-node:

md	mode
ln	link count

<b>&amp;</b>	user id number
<b>gid</b>	group id number
<b>sz</b>	file size
<b>a#</b>	data block numbers (0-7)
<b>at</b>	access time
<b>mt</b>	modification time
<b>maj</b>	device class number
<b>min</b>	logical device identification number.

#### EXAMPLES

Examples of *fsdb* uses are:

- 386i** prints i-number 386 in an i-node format. This now becomes the current working i-node.
- a0b.p0x10:h** prints the first 16 words of the file for which a 0 is the starting block number.
- 2i.fd** prints the first 32 directory entries for the root i-node of this file system.
- d5i.fc** changes the current i-node to the i-node associated with the fifth directory. The first 512 bytes of the file are then printed in ASCII.
- lb.p0o** prints the superblock of this file system in octal.

#### LIMITATIONS

*fsdb* is a useful command which would be even more useful if it worked as planned (for example, the = operator when applied to an i-node value).



NAME

`fsinit` — file system initializer

SYNOPSIS

`fsinit` <SG data base name><partition name>[mount point]

DESCRIPTION

The *fsinit* command locates the given *partition name* in the given *SG data base*, and extracts the disk number, partition number, partition block count, and partition inode block count for that partition. This information is then used to initialize the named partition, mounting it on the optional *mount point*. If no mount point name is given, then `/dev/<partition name>` is used. *Fsinit* uses the *clrf*s system call to initialize the superblock of the named partition with the data extracted from the named SG data base. Basically, *fsinit* acts like a smart "*mount*" command - it tries to guarantee that the named partition is successfully mounted, and is correctly initialized as per the SG data base specifications.

DIAGNOSTICS

*Fsinit* will fail with an error code if the named SG data base cannot be attached or read, or if the requested partition cannot be successfully mounted. Appropriate error messages are printed before *fsinit* exits.

FILES

`/etc/fsinit`

EXAMPLES

The command:

```
fsinit /data base/appdmert.sg no5text /tmp/no5mtpt
```

will read the application SG data base (`appdmert.sg`), extract the information needed to initialize the superblock of `no5text`, clear and update the `no5text` superblock using the new information, and mount the `no5text` partition on `/tmp/no5mtpt`.

The command:

```
fsinit /data base/appdmert.sg no5text
```

will act the same as the above command, but the `no5text` partition will be mounted on `/dev/no5text` by default.

CAVEATS

Note that *fsinit* will wipe out the existing information in the named partition's superblock, thereby making that partition appear "empty" - as any link to the information stored therein is lost. For this reason, *fsinit* should be used with discretion and care.

SEE ALSO

`fsaudit(1)`



NAME

fsize — get size of contiguous file

SYNOPSIS

**fsize** filename

DESCRIPTION

The size of filename specified is returned.

LIMITATIONS

Fsize is meaningless on noncontiguous files.



## NAME

*grep*, *fgrep* — search a file for a pattern

## SYNOPSIS

**grep** [options] expression [files]  
**fgrep** [options] strings [files]

## DESCRIPTION

Commands of the *grep* family search the input *files* (standard input default) for lines matching a pattern. Normally, each line found is copied to the standard output. *Fgrep* patterns are fixed *strings*; it is fast and compact. *Grep* patterns are limited regular *expressions* in the style of *ed* (1); it uses a compact nondeterministic algorithm. The following *options* are recognized:

- v** All lines but those matching are printed.
- x** (Exact) only lines matched in their entirety are printed (*fgrep* only).
- c** Only a count of matching lines is printed.
- l** Only the names of files with matching lines are listed (once), separated by new lines.
- n** Each line is preceded by its relative line number in the file.
- b** Each line is preceded by the block number on which it was found. This is sometimes useful in locating disk block numbers by context.
- s** The error messages produced for nonexistent or unreadable files are suppressed (*grep* only).
- e expression** Same as a simple *expression* argument, but useful when the *expression* begins with a - (does not work with *grep*).
- f file** The *strings* list (*fgrep*) is taken from the *file*.

In all cases, the file name is output if there is more than one input file. Care should be taken when using the characters \$, \*, [, ^, |, (, ), and \ in *expression*; because they are also meaningful to the shell. It is safest to enclose the entire *expression* argument in single quotes '...'

*Fgrep* searches for lines that contain one of the *strings* separated by new lines.

## SEE ALSO

*ed*(1), *sed*(1), *sh*(1)

## DIAGNOSTICS

Exit status is 0 if any matches are found, 1 if none, and 2 for syntax errors or inaccessible files (even if matches were found).

## BUGS

Ideally there should be only one *grep*, but a single algorithm that spans a wide enough range of space-time tradeoffs does not exist.

Lines are limited to 256 characters; longer lines are truncated.



NAME

`ichk` — file system consistency check and interactive repair

SYNOPSIS

`ichk` [ `-bnsvy` ] [filesystem] ...

DESCRIPTION

*Ichk* audits DMERT file systems for consistency and allows the operator to correct any discrepancies. Note that a file system must be unmounted before corrections can be made (see *LIMITATIONS*). Since these corrections will, in general, result in a loss of data, the program will request operator concurrence for each such action. All questions should be answered by typing yes or no, followed by a newline character. Typing yes causes the correction to take place. However, if the program does not have write permission on the file system all questions will automatically be answered no.

The following options are recognized:

- b** runs only phases that check for block inconsistencies (phases 1, 2, and 7).
- n** causes all questions to be answered no automatically.
- s** salvages all blocks not found in some file and makes a new free list (if the bit map is bad, it too is recreated).
- v** turns the verbose mode off.
- y** causes all questions to be answered yes automatically.

The program consists of seven separate phases. Some phases are skipped if they are not needed. In phase one, *ichk* examines all block pointers in all files; checking for pointers which are outside of the file system (BAD), for blocks which appear in more than one file (DUP), and for blocks which are (incorrectly) included in the i-node's block pointer list but are actually beyond the size of the file. A table is made of all DUP blocks and all defective files are marked for clearing (BAD and DUP only). No correction takes place in this phase for BAD and DUP errors. The operator is, however, given an option of clearing the incorrect (BEYOND FILE SIZE) block pointers by replacing them with null (0) values.

The second phase is run only if DUP blocks were found in phase one. This phase finds the rest of the DUP blocks and marks them for clearing.

If the **-b** option has not been selected, then the third and fourth phases check the directory structure of the file system by descending the directory tree and examining each entry. A count of the number of references to each file is maintained. If any entry refers to an unallocated file, a file marked for clearing, or a file number outside the file system, the entry is printed, and, if the operator agrees, it is removed. Refusing to remove an entry to a marked file will clear the mark and preserve the file with its subsequent entries.

If the **-b** option has not been selected, then phase five lists all marked or unreferenced files. With concurrence from the operator, each of these files is

*ICLK(1)*

then cleared. In addition, any file with a link count that does not agree with the number of references is listed, and, if the operator agrees, the link count is adjusted.

If the salvage option, **-s**, and/or the block option, **-b**, is on, phase six is skipped. Otherwise, the bit map and free list are examined. If any blocks are found that are outside the file system, that have been previously encountered in a file or elsewhere in the free list, or that do not appear in the bit map, free list, or any file then the list or bit map is pronounced BAD and a salvage is called for. Operator agreement will set the salvage option and proceed to the next phase. If there are no defects in the free list and bit map and all blocks are accounted for, the check is finished. Otherwise, the number of missing blocks is printed and a salvage is requested.

The last phase is the salvage operation, where the free list and/or the bit map is recreated. It is run whenever the salvage option is on or a problem has been found in phase 6. Simply stated, a new free list is constructed containing all blocks not found in some file (if the bit map is bad, it too is recreated).

The system responses are, in general, self-explanatory and follow the sequence described above. In the description of the output which follows, this notation will be used:

**<b>**        block number  
**<i>**        i-node number  
**<fname>**    file pathname  
**<n>**        positive integer  
**<c>**        option character

*Ichk* begins with the following output.

```
<filesystem>{ (NO WRITE) }
```

## Phase 1 — Check Blocks

The (NO WRITE) message indicates that the program does not have write permission on the file system. Therefore, subsequent corrections will be suppressed, by automatically answering no to all questions. Phase 1 then proceeds to list any BAD, DUP, or BEYOND FILE SIZE blocks and their i-node number, as follows:

```
<b> BAD                    I=<i>
<b> DUP                    I=<i>
<b> EXCESSIVE DUPS        I=<i>
<b> BEYOND FILE SIZE      I=<i>
```

If no **-[yn]** option has been specified, the program will wait for a response of yes or no after each BEYOND FILE SIZE message. A yes answer will replace the block pointer with a null (0) value. A no answer will leave the block pointer unaffected. In either case, all other data associated with the i-node is not affected, nor is the i-node cleared.

*ICLK(1)*

If too many DUPs are encountered, the program will list all blocks, but it will not mark the excess DUPs for later processing. When phase 1 is finished, if any DUPs were encountered, then phase 2 is run. Otherwise, phase 2 is skipped.

Phase 2 — Rescan For More

<b> DUP I = <i>

Check descends the directory tree, asking whether to remove any defective entries.

Phase 3 — Check Pathnames

No text.

Phase 4 — Check Connectivity

I OUT OF RANGE I = <i> <fname> REMOVE?  
UNALLOCATED I = <i> <fname> REMOVE?  
BAD/DUP I = <i> <fname> REMOVE?

If no option has been specified, the program will wait for a response of yes or no after each question. A no answer to the BAD/DUP entry will unmark that i-node for clearing. This will suppress any subsequent correction to that file.

Phase 5 — Check Reference Counts

Now *ichk* will clear or adjust any defective files. Again, if no option has been specified, it will wait for a yes or no response to each question. The program will also indicate whether each entry is a file or a directory.

UNREFERENCED {FILE/DIRECTORY} I = <i> CLEAR?  
BAD/DUP {FILE/DIRECTORY} I = <i> CLEAR?  
LINK COUNT {FILE/DIRECTORY} I = <i> ADJUST?

Phase 6 — Check Free List

If the salvage option is not on, the program will now validate the free list and bit map. Otherwise, this phase is skipped. If there are any errors, it will specify them and request a salvage.

BAD FREE CHAIN SALVAGE?  
BAD BIT MAP SALVAGE?  
FREE LIST BOTCH SALVAGE?

Phase 7 — Salvage Free List

This operation is performed only upon request.

*ICLK(1)*

These totals are printed: the total number of allocated files (including directories and special files), the number of blocks in use, and the number of blocks in the free list.

<n> FILES <n> BLOCKS <n> FREE

If the file system has been modified, then the message, BOOT UNIX(NO SYNC!), is printed and the program goes into a loop. This is only a reminder to the operator since the program can be forced to terminate with a <DEL> character.

## DIAGNOSTICS

Exit code 0 is returned for no file system consistency errors; 1, if file system errors were found and all were corrected; 2, if file system errors were found but not all were corrected; and 3, for *ichk* process failure before completion (for example, because of inability to open device). A number of errors can terminate *ichk*. An illegal option is ignored, but the inability to open the file system is shown as:

CAN NOT OPEN <filesystem>

An I/O error on the file system will also return an error message. In this case, the operator is given the choice of exiting (yes) or continuing (no). This error is generally a hardware error, and continuing is rarely a good idea.

CAN NOT READ <filesystem> BLOCK <b> EXIT?

CAN NOT SEEK <filesystem> BLOCK <b> EXIT?

CAN NOT WRITE <filesystem> BLOCK <b> EXIT?

## LIMITATIONS

*ichk* has been known to produce core images on large file systems.

Since *ichk* diagnoses problems that exist on the disk version of the file system (such as superblock and bit map), two warnings are given.

*-File system errors found by*

*ichk* valid only if the file system was unmounted during the *ichk* process. If the file system cannot be unmounted and an *ichk* is mandatory then the operator should run several (three or more) *ichk* with *-n* option on a quiet (if possible) mounted file system. Any errors that appear in all the *ichks* probably exist. If the errors disappear in all the *ichks*, then they are probably temporary and of no concern.

*-Corrections that*

*ichk* makes to the mounted file system are overwritten (invalidated) with in-core file system information. Therefore, corrections should be made only when the file system is unmounted.

NAME

`id` — print user and group IDs and names

SYNOPSIS

`id`

DESCRIPTION

*Id* writes a message on the standard output giving the user and group IDs and the corresponding names of the invoking process. If the effective and real IDs do not match, both are printed.



NAME

kill — terminate a process

SYNOPSIS

kill [ - signo ] PID ...

DESCRIPTION

*Kill* sends signal 15 (terminate) to the specified processes. This will normally kill processes that do not catch or ignore the signal. The process number of each asynchronous process started with **&** is reported by the Shell (unless more than one process is started in a pipeline, in which case the number of the last process in the pipeline is reported). Process numbers can also be found by using *ps* (1).

If process number 0 is specified, all processes in the process group are signaled.

The killed process must belong to the current user unless the user is the super user.

If a signal number preceded by - is given as first argument, that signal is sent instead of terminate. In particular “kill -9 ...” is a sure kill.

SIGNALS

- 1 hang-up
- 2 interrupt
- 3 quit
- 4 illegal instruction
- 5 trace trap
- 6 IOT instruction
- 7 EMT instruction
- 8 floating point exception
- 9 kill (cannot be caught or ignored)
- 10 bus error
- 11 segmentation violation
- 12 bad argument to system call
- 13 write on a pipe with no one to read it
- 14 alarm clock
- 15 software termination signal

SEE ALSO

ps(1), sh(1)



NAME

*killp* — terminate user processes by invoking *pathname*

SYNOPSIS

**killp** [ **-signo** ] *full-pathname*

DESCRIPTION

*Killp* sends **signo** to all processes with an invocation path of **full-pathname**. **signo** is a decimal integer constant.

If **signo** is defaulted, the signal 9 (**SIGKILL**) is sent.

The path **full-pathname** must begin with a slash (/).

SEE ALSO

*kill(1)*, *pkill(1)*, *ps(1)*

DIAGNOSTICS

If an error occurs, a message is written to *stderr* and an exit status of 1 is returned.

LIMITATIONS

*Killp* will only kill processes whose *argv[0]* is identical to the full *pathname* with which the process was invoked. Therefore, processes invoked with relative *pathnames*, or processes that modify *argv[0]* will not be killable with this command. Also, processes invoked from linked files will only be killed if *argv[0]* is the full *pathname* of the appropriate linked filename (that is, the **full-pathname** argument to *killp*).

CAVEATS

The *killp* command will not terminate a suspended process because it results in a message being sent to the process, requesting it to kill itself. Since the process is suspended, it cannot handle these messages.



NAME

lf — list contents of directory

SYNOPSIS

lf [-1aAbcCdfFgilmnopqrRstuvVx] [files ...]

DESCRIPTION

For each directory argument, *lf* lists the contents of the directory; for each file argument, *lf* repeats its name and any other information requested. The output is sorted alphabetically by default. When no argument is given, the current directory is listed. When several arguments are given, the arguments are first sorted appropriately, but file arguments appear before directories and their contents.

There are several major listing formats. The format chosen depends on whether the output is going to a teletype, and may also be controlled by option flags. The default format for a teletype is to list the contents of directories in multi-column format, with the entries sorted down the columns. (Files which are not the contents of a directory being interpreted are always sorted across the page rather than down the page in columns. This is because the individual file names may be arbitrarily long. This format may be requested with *-x*.) If the standard output is not a teletype, the default format is to list one entry per line. (This is the *-1* option.) A long listing, detailing much information about the file can be requested with the *-l* option. Finally, there is a stream output format in which files are listed across the page, separated by *'* characters. The *-m* flag invocation enables this format.

There are an unbelievable number of options:

- 1* force one entry per line output format, e.g. to a teletype.
- a* List all entries; usually entries beginning with *'* are suppressed.
- A* List all entries except *'* and *'..*
- b* force printing of non-graphic characters to be in the *+dd* notation, in octal.
- c* Use time of file creation for sorting or printing.
- C* force multi-column output, e.g. to a file or a pipe.
- d* If argument is a directory, list only its name, not its contents (mostly used with *-l* to get status on directory.)
- f* Force each argument to be interpreted as a directory and list the name found in each slot. This option turns off *-l*, *-t*, *-s*, and *-r*, and turns on *-a*; the order is the order in which entries appear in the directory.
- F* cause directories to be marked with a leading and trailing square brackets *'[ dir ]'* and executable files to be marked with a leading *'\*'*; this is the default.
- g* Same as *-l* except the owner name is not printed.
- i* Print i-number in first column of the report for each file listed.

*LF(1)*

- l** List in long format, giving mode, number of links, owner, group, size in bytes, and time of last modification for each file. (See below.) If the file is a special file the size field will instead contain the major and minor device numbers.
- m** force stream output format.
- n** Same as **-l** except owner and group are printed as integers instead of searching the password and group files for translations.
- o** Same as **-l** except the group name is not printed.
- p** Same as **-F** except executable files are not marked.
- q** force printing of non-graphic characters in file names as the character '?'; this normally happens only if the output device is a teletype.
- r** Reverse the order of sort to get reverse alphabetic or oldest first as appropriate.
- R** recursively list subdirectories encountered.
- s** Give size in blocks, including indirect blocks, for each entry.
- t** Sort by time modified (latest first) instead of by name, as is normal.
- u** Use time of last access instead of last modification for sorting (**-t**) or printing (**-l**).
- v** on multi-column output, column width is fixed (non-variable). This is the default.
- V** Print version number and exit.
- x** force columnar printing to be sorted across rather than down the page; this is the default.

The mode printed under the **-l** option contains 11 characters which are interpreted as follows, the first character is:

- b** if the entry is a block-type special file;
- c** if the entry is a character-type special file;
- d** if the entry is a directory;
- p** if the entry is a FIFO (named pipe) special file.
- if the entry is a plain file.
- C** if the entry is a one contiguous extents file.
- x** if the entry is allocated by contiguous extents.

The next 9 characters are interpreted as three sets of three bits each. The first set refers to owner permissions; the next to permissions to others in the same user-group; and the last to all others. Within each set the three characters indicate permission respectively to read, to write, or to execute the file as a program. For a directory, 'execute' permission is interpreted to mean permission to search the directory for a specified file. The permissions are indicated as follows:

- r** if the file is readable;

- w** if the file is writable;
- x** if the file is executable;
- if the indicated permission is not granted.

The group-execute permission character is given as **s** if the file has set-group-ID mode; likewise the user-execute permission character is given as **s** if the file has set-user-ID mode. An **S** in either position indicates the set-ID bit is set but the corresponding execute bit is not.

The last character of the mode (normally 'x' or '-') is **t** if the 1000 (save text) bit of the mode is on. A **T** in this position indicates the save text bit is on but there is no execute permission for 'other.' See `chmod(1)` for the meaning of this mode.

When the sizes of the files in a directory are listed, a total count of blocks, including indirect blocks is printed.

## FILES

`/etc/passwd`  
`/etc/group`

## BUGS

Newline and tab are considered printing characters in file names.



NAME

`line` — read one line

SYNOPSIS

`line`

DESCRIPTION

*Line* copies one line (up to a new line) from the standard input and writes it on the standard output. It returns an exit code of 1 on **EOF** and always prints at least a new line. It is often used within shell files to read from the user's terminal.

SEE ALSO

`sh(1)`



235-700-200  
November 1998

COMMANDS

*LN(1)*

NAME

*ln*

DESCRIPTION

*See cp.*



NAME

`logdir` — get login directory

SYNOPSIS

`logdir` [`user ...`]

DESCRIPTION

*Logdir* returns the login directory of a user. If no arguments are specified, the contents of the environment variable **\$HOME** is listed. If the **\$HOME** variable is not set, ie:unset, or set to NULL, the login directory is taken from the password file. *User* argument(s), if specified, cause the login directory for that user to be taken from the password file.

FILES

`/etc/passwd`

SEE ALSO

`login(1)`, `env(1)`

DIAGNOSTICS

Returns the number of unknown users, which may be zero.



NAME

login — sign on

SYNOPSIS

**login** [ name ]

DESCRIPTION

The *login* command is used at the beginning of each terminal session and allows you to identify yourself to the system. It may be invoked as a command or by the system when a connection is first established. Also, it is invoked by the system when a previous user has terminated the initial shell by typing a *cntrl-d* to indicate an “end-of-file.”

*Login* may not be invoked as a command.

*Login* asks for your user name (if not supplied as an argument), and, if appropriate, your password. Echoing is turned off (where possible) during the typing of your password so it will not appear on the written record of the session.

At some installations, an option may be invoked that will require you to enter a second “machine” password. This will be prompted by the message “machine password:”. Both passwords are required for a successful login.

If you do not complete the login successfully within 5 tries, all later attempts will fail. The TTY will have to be removed and restored before it will again let any one login on it.

After a successful login, accounting files are updated; the procedure */etc/profile* is performed; the message-of-the-day, if any, is printed; the user-ID, the group-ID, the working directory, and the command interpreter [usually *sh* (1)] is initialized; and the file *.profile* in the working directory is executed, if it exists. These specifications are found in the */etc/passwd* file entry for the user. The name of the command interpreter is - followed by the last component of the interpreter’s pathname (i.e., *-sh* ). If this field in the password file is empty, then the default command interpreter, */bin/sh* is used.

The basic *environment* is initialized to:

HOME=*your-login-directory*

PATH=:/bin:/usr/bin

SHELL=*last-field-of-passwd-entry*

MAIL=/usr/mail/*your-login-name*

TZ=*timezone-specification*

**LOGIN(1)**

## FILES

/etc/utmp	accounting
/etc/wtmp	accounting
/usr/mail/ <i>your-name</i>	mailbox for user <i>your-name</i>
/etc/motd	message-of-the-day
/etc/passwd	password file
/etc/profile	system profile
.profile	user's login profile

When you want to end a terminal session, type a *cntrl-d* to indicate an "end-of-file." Note that if you are logged on to a DAP (display administration process) terminal, you must be in the message area at the bottom of the screen when you enter the *cntrl-d*. Always end the terminal session and wait for the subsequent prompt to appear before powering off a terminal.

## SEE ALSO

admin(1), mail(1), passwd(1), sh(1), su(1)

## DIAGNOSTICS

*Login incorrect* if the user name or the password cannot be matched.

*No shell, cannot open password file, or no directory*: consult a *UNIX* system expert.

## ROP OUTPUT

*REPT LOGIN TTYNAME x FAILED*. If *x* is the ttyname of the MCC (master control center) or the SCC (switch control center), change ECD so that login is not invoked from ttyname *x*.

*REPT LOGIN TTYNAME x UNABLE TO ACCESS ECD*. If the ECD is unavailable, try again later.

*REPT LOGIN TTYNAME x FAILED TO CREATE y*. If process *y* cannot be executed successfully, consult a *UNIX* system expert.

NAME

`lpr` — line printer spooler

SYNOPSIS

`lpr [ -s ] [ device ]`

DESCRIPTION

*Lpr* sends standard input to the local line printer [ROP (read-only printer)]. Output can be directed to other devices by either changing the device(s) associated with output class *UNIX* system or by specifying the desired **device** as the first parameter.

*Lpr* assigns a job number, divides the input into parts for the spooler, adds a header to each part, sends the parts to the spooler, and prints the job number and number of parts to standard output.

If the single part option **-s** is specified, the printout will not be divided into parts. This makes for a 'nicer' looking output but can cause alarms at the SCC.

NOTE

Many ROPs are configured to print lowercase letters as uppercase letters. If lowercase letters are desired, set switch 1 of switch block SWE8 of the ROP to the "off" state.

ERROR MESSAGES

**lpr: cancelled non-printable character @ X in input stream.**

Possible cause - attempt to print an object module. *X* is the location of the nonprintable.

**lpr: aborted - can't open temp file /unixa/tmp/UA....**

Possible causes - /unixa/tmp directory is missing, /unixa file system not mounted read/write, or file system full.

EXAMPLES

Print the output of the 'date' command on the recent change and verify terminal.

```
date | lpr ttyv
```

Print /etc/passwd on the ROP.

```
lpr </etc/passwd
```

FILES

/unixa/tmp/UA\* File(s) passed to spooler to print.

BUGS

Output to unknown devices is dropped by the spooler without notice.



NAME

ls — list contents of directories

SYNOPSIS

ls [ *-lgtasdrucif* ] names

DESCRIPTION

For each directory named, *ls* lists the contents of that directory; for each file named, *ls* repeats its name and any other information requested. By default, the output is sorted alphabetically. When no argument is given, the current directory is listed. When several arguments are given, the arguments are first sorted appropriately, but file arguments are processed before directories and their contents. There are several options:

- l** List in long format, giving mode, number of links, owner, size in bytes, and time of last modification for each file (see below). If the file is a special file, the size field will contain the major and minor device numbers; rather than a size.
- g** Print group rather than owner with **-l** option.
- t** Sort by time of last modification (latest first) instead of by name.
- a** List all entries; in the absence of this option, entries whose names begin with a period (.) are *not* listed.
- s** Give size in blocks (including indirect blocks) for each entry.
- d** If argument is a directory, list only its name; often used with **-l** to get the status of a directory.
- r** Reverse the order of sort to get reverse alphabetic or oldest first, as appropriate.
- u** Use time of last access instead of last modification for sorting (with the **-t** option) and/or printing (with the **-l** option).
- c** Use time of last modification of the inode (mode, etc.) instead of last modification of the file for sorting (**-t**) and/or printing (**-l**).
- i** For each file, print the i-number in the first column of the report.
- f** Force each argument to be interpreted as a directory and list the name found in each slot. This option turns off **-l**, **-t**, **-s**, and **-r**, and turns on **-a**; the order is the order in which entries appear in the directory.

The mode printed under the **-l** option consists of 11 characters that are interpreted as follows:

The first character is:

- d** if the entry is a directory;
- b** if the entry is a block special file;
- c** if the entry is a character special file;
- p** if the entry is a fifo (a.k.a. “named pipe”) special file;
- C** if the entry is a contiguous file;

*LS(1)*

- x** if the entry is a multiextent file;
- if the entry is an ordinary file.

The next nine characters are interpreted as three sets of three bits each. The first set refers to the owner's permissions; the next to permissions of others in the user-group of the file; and the last to all others. Within each set, the three characters indicate permission to read, to write, and to execute the file as a program, respectively. For a directory, "execute" permission is interpreted to mean permission to search the directory for a specified file.

The permissions are indicated as follows:

- r** if the file is readable;
- w** if the file is writable;
- x** if the file is executable;
- if the indicated permission is *not* granted.

The group-execute permission character is given as **s** if the file has set-group-ID mode; likewise, the user-execute permission character is given as **S** if the file has set-user-ID mode. The last character of the mode (normally **x** or **-**) is **t** if the 1000 (octal) bit of the mode is on; see *chmod(1)* for the meaning of this mode. The indications of set-ID and 1000 bit of the mode are capitalized (**S** and **T**, respectively) if the corresponding execute permission is *not* set.

When the sizes of the files in a directory are listed, a total count of blocks, including indirect blocks, is printed.

## FILES

`/etc/passwd` to get user IDs and  
`/etc/group` to get group IDs.

## SEE ALSO

`chmod(1)`, `find(1)`

NAME

mail, rmail — send mail to users or read mail

SYNOPSIS

**mail** [ **-pqr** ] [ **-f file** ]  
**mail** persons  
**rmail** persons

DESCRIPTION

*Mail* without arguments prints a user's mail, message-by-message, in last-in, first-out order. For each message, the user is prompted with a **?**, and a line is read from the standard input to determine the disposition of the message:

- <new line> Go on to next message.
- +** Same as <new line>.
- d** Delete message and go on to next message.
- p** Print message again.
- Go back to previous message.
- s** [ *files* ] Save message in the named *files* (**mbox** is default).
- w** [ *files* ] Save message, without its header, in the named *files* (**mbox** is default).
- m** [ *persons* ]  
Mail the message to the named *persons* (yourself is default).
- q** Put undeleted mail back in the *mailfile* and stop.
- EOT** (control-d)  
Same as **q**.
- x** Put all mail back in the *mailfile* unchanged and stop.
- !command** Escape to the shell to do *command*.
- \*** Print a command summary.

The optional arguments alter the printing of the mail:

- p** causes all mail to be printed without prompting for disposition.
- q** causes *mail* to terminate after interrupts. Normally an interrupt only causes the termination of the message being printed.
- r** causes messages to be printed in first-in, first-out order.
- f file** causes *mail* to use *file* (e.g., **mbox** ) instead of the default *mailfile*.

When *persons* are named, *mail* takes the standard input up to an end-of-file (or up to a line consisting of just a **.**) and adds it to each *person's mailfile*. The message is preceded by the sender's name and a postmark. Lines that look like postmarks in the message (that is, "From ...") are preceded with a **>**. A *person* is usually a user name recognized by *login* (1). If a *person* being sent mail is not recognized, or if *mail* is interrupted during input, the file **dead.letter** will be saved to allow editing and resending.

**MAIL(1)**

To denote a recipient on a remote system, prefix *person* by the system name and exclamation mark. Everything after the first exclamation mark in *persons* is interpreted by the remote system. In particular, if *persons* contains additional exclamation marks, it can denote a sequence of machines through which the message is to be sent on the way to its ultimate destination. For example, specifying **a!b!cde** as a recipient's name causes the message to be sent to user **b!cde** on system **a**. System **a** will interpret that destination as a request to send the message to user **cde** on system **b**. This might be useful, for instance, if the sending system can access system **a** but not system **b**, and system **a** has access to system **b**.

*Rmail* only permits the sending of mail.

When a user logs in, the presence of mail, if any, is indicated. Also, notification is made if new mail arrives while using *mail*.

## FILES

/etc/passwd	to identify sender and locate persons
/usr/mail/ <i>user</i>	incoming mail for <i>user</i> ; that is, the <i>manfile</i>
\$HOME/mbox	saved mail
\$MAIL	variable containing path name of <i>mailfile</i>
/tmp/ma*	temporary file
/usr/mail/*.lock	lock for mail directory
dead.letter	unmailable text

## SEE ALSO

login(1), write(1)

## BUGS

Race conditions sometimes result in a failure to remove a lock file.

After an interrupt, the next message may not be printed; printing may be forced by typing a **p**.

NAME

`man` — display manual pages

SYNOPSIS

`man` [`command ...`] **all.man.pages**

DESCRIPTION

*Man* attempts to locate the manual page(s) for each *command* specified and then displays them on the standard output; thus:

```
man at
```

displays the manual page for the 'at' command.

```
man cat man
```

displays the manual pages for the 'cat' and 'man' commands.

Some commands, such as 'true' and 'false,' share a single manual page. Thus, requesting either one results in displaying the same page.

If **all.man.pages** is specified, then the names of all manual pages are displayed.

FILES

`/unixa/man/*` Manual pages



NAME

`mesg` — permit or deny messages

SYNOPSIS

`msg [ n ] [ y ]`

DESCRIPTION

*Mesg* with argument **n** forbids messages via *write* (1) by revoking nonuser write permission on the user's terminal. *Mesg* with argument **y** reinstates permission. All by itself, *mesg* reports the current state without changing it.

FILES

`/dev/tty*`

SEE ALSO

`write(1)`

DIAGNOSTICS

Exit status is 0 if messages are receivable, 1 if not, or 2 on error.



NAME

`mkdir` — make a directory

SYNOPSIS

`mkdir` *dirname* ...

DESCRIPTION

*Mkdir* creates specified directories in mode *777*, possibly altered by *umask* [see *sh(1)*]. Standard entries, `.`, for the directory itself, and `..`, for its parent, are made automatically.

*Mkdir* requires write permission in the parent directory.

SEE ALSO

*sh(1)*, *rm(1)*, *umask* [see *sh(1)*]

DIAGNOSTICS

*Mkdir* returns exit code 0 if all directories were successfully made; otherwise, it prints a diagnostic and returns nonzero.



NAME

mkdisk — read in disk partition(s) from tape

SYNOPSIS

```
/etc/mkdisk -i <tape> -d <vtoc> -p <list>  
/etc/mkdisk -i <tape> -o <mhd>
```

DESCRIPTION

The mkdisk command allows one or more partitions from a Load Format (LDFT) tape to be read into disk. The first format allows selected partitions to be read into an on-line disk or disk pair. The second format reads in all partitions on the tape into a single Out Of Service (OOS) disk.

The parameters for the first format are:

<tape>      Tape device name, usually /dev/mt00 (high speed) or /dev/mt08 (low speed).  
<vtoc>      vtoc file name of destination disks. /dev/vtoc for MHDs 0&1, /dev/vtoc1 for MHDs 2&3, /dev/vtoc2 for MHDs 4&5 and so on.  
<list>      A file name containing the names of the partitions to be read in.

The parameters for the second format are:

<tape>      Tape device name, usually /dev/mt00 (high speed) or /dev/mt08 (low speed).  
<mhd>      Destination disk name. Examples: MHD=0, MHD=4, MHD=5.

Multi-Reel

If the tape is part of a multi-reel sequence, after each tape is read the system will rewind it and request that the tape next be mounted. After it is mounted, enter the command '/etc/mkstart'. This must be done within 5 minutes or the system will time out and abort the process.

Vtoc

Be aware the /etc/mkdisk will NOT read in if the vtoc on the tape does not agree with the vtoc on the disk. The craft command:

```
EXC:ENVIR:UPROC, FN="/etc/rcvtoc", ARG=x
```

Will place the correct vtoc on the MHD number specified by x.



NAME

`mknod` — build a special file

SYNOPSIS

```
/etc/mknod name [b|c|l|r] dcn part rid  
/etc/mknod name p
```

DESCRIPTION

**Mknod** makes a directory entry and corresponding i-node for a special file. The first argument is the **name** of the entry.

The second argument identifies the special file as one of the following types:

- b block-type (disk)
- c character-type
- i iop-type
- r record-type (magnetic tape)
- p FIFO-type (named pipe).

The last three arguments are numbers specifying the device class number (**dcn**), partition ( **part**), and record ID of the minor device chain table (**rid**) for the special file. The numbers may be expressed in either decimal octal notation with the latter being indicated by a leading zero. The default creation mode is 0666.

The assignment of the **dcn**, **part**, and **rid** is specific to each system and must be consistent with the equipment configuration data.

When creating a FIFO special file, the **dcn**, **part**, and **andrid** are unnecessary.



NAME

mkstart — continues a mkdsk run after a new tape has been mounted

SYNOPSIS

**mkstart**

DESCRIPTION

*Mkstart* sends a message to *mkdsk* in order to have *mkdsk* continue its run after a new tape has been mounted. *Mkstart* accepts no parameters.

HEADER FILES

umsg.h, splcl.h

SEE ALSO

mkdsk(1)

DIAGNOSTICS

Processing errors are echoed to the invoking terminal. The errors that *mkstart* reports are:

- A failure to get the process id of mkdsk
- Incorrect syntax of the mkstart command
- A failure to send the message to mkdsk.



NAME

mount, umount — mount and dismount file system

SYNOPSIS

*/etc/mount* [ special directory [ -r ] ]  
*/etc/umount* special

DESCRIPTION

*Mount* announces to the system that a file system is present on the device *special*. The *directory* must exist already; it becomes the name of the root of the newly mounted file system.

These commands maintain a table of mounted devices. If invoked with no arguments, *mount* prints the table.

The optional last argument indicates that the file is to be mounted read-only. Physically write-protected and magnetic tape file systems must be mounted in this way or errors occur when access times are updated, whether or not any explicit write is attempted.

*Umount* announces to the system that the removable file system previously mounted on device *special* is to be removed.

DIAGNOSTICS

*Mount* issues a warning if the file system to be mounted is currently mounted under another name.

*Umount* complains if the special file is not mounted or if it is busy. The file system is busy if it contains an open file or some user's working directory.

BUGS

Some degree of validation is done on the file system; however, it is generally unwise to mount garbage file systems.



NAME

msnap — memory usage snapshot program

SYNOPSIS

**msnap** -Ttype -P pid1 pid2 -U uid1 uid2 -Lfile -S -A -F

DESCRIPTION

*Msnap* is a user level program that collects run time information on the memory consumption of a process, shared library, the kernel, or an entire system.

Many options are provided by the tool to allow flexibility in collecting memory consumption data on a per process or system basis. Two major report types are currently supported for processes. The first is the short report (default). This report will print out: the process name, utility id, process id, and tty (if a supervisor). Sharable memory consumption is provided in pages for text, data, and stack segments. Private memory consumption is provided in pages for text, data, and stack segments. In the context of this document, sharable memory is defined by the presence of a flag in the process segment list, not by whether the segment is actually shared or not. All text segments are sharable by default. For a description of the long report, see the type arguments listed below.

Report Types

The -T argument controls the report type and format of data provided by the tool. The current report types are: sSkKIL. The meaning of each report type follows.

- s** This option will provide information on all active supervisor level processes.
- k** This option will provide information on all active kernel level processes (except kernel special processes).
- K** This option provides memory usage information on the kernel which includes the special processes.
- I** This option requests that shared library information be printed. The -L*libfile* argument must also be specified. The *libfile* keyword indicates the full or relative pathname to a file containing the shared library specifications. See the section on shared libraries for the format of this file.
- S** This option will provide summary information on the memory consumption of the processes snapshot. Total memory used, along with summary information on sharable, actually shared, swappable, and locked swappable memory, will be provided in segments, pages, and bytes.
- L** This option will provide a *long* listing of the memory usage for the specified processes. A line will be printed for each segment that the measured process owns. Segment information will be output in increasing order by segment index. The segment index (SNDX), segment list flags (SFLGS), segment id (SEGID), segment descriptor flags (SDEFLAGS), process lock count (PLC), I/O lock count (IOLC), nonswap count (NSWC), active count (ACTC),

**MSNAP(1)**

number of users of the segment (NUSRS), the system name for the segment (SEG\_NAME), the number of pages in the segment (PGS), the size of the last page in bytes (LSTPG), the pages in memory for supervisor processes (INM), and the swap block address for supervisor processes (SWAD) are provided. This information is collected from the processes segment list and the system segment descriptor.

**Shared Libraries**

Information on shared libraries can be obtained by this tool. For flexibility, a specfile (*-Lfile option*) can be specified for shared libraries. The file argument is the full or relative pathname of the specfile. The specfile contains information that will allow the tool to determine which processes include a particular shared library and which segments belong to the library.

Each shared library in the system has identifying attributes. The operating system allows two 32-bit words (library word 0 and 1) for library identification. Each bit in either word identifies a particular library. The msnap tool uses this means to determine which libraries a process is attached to.

The format of the specfile contains a star (\*) immediately followed by the library name (limit of 16 characters). Next, the library word and bit index are specified, with each value being separated by a space or tab. Additional lines will indicate the segment indexes associated with the library. The segment indexes are broken up into four groups. These are: the Text Segments (TS), Data Segments (DS), Patch Segments (PS), and Misc Segments (MS) for any additional segments. Each segment list entry must be separated by a space character. Each segment definition must be on its own line. At least one segment type definition with at least one segment index must be specified for each library. For text and data segments, the maximum number of segments that can be specified is 8. For patch and misc segments, the maximum is 4. There is a limit of 20 library definitions in the specfile.

An example of the library specifications file entry for libc follows:

```
*UNIX_LIBC 1 3
    TS 50
    DS 51
    PS 52
```

The library name follows the star (\*) character which is the library entry separator. The "1 3" indicates library word 1, bit position 3.

The "TS 50" line indicates one text segment at segment index 50. The "DS 51" line indicates one data segment at segment index 51. The "PS 52" line indicates one patch segment at segment index 52.

**Specific Processes**

The **-P** and **-U** arguments allow the selection of specific processes to be measured. The **-P** argument allows for a list of up to 20 process ids (in decimal) to be specified. The **-U** argument allows for a list of up to 20 utility ids to be

specified. Each utility id must be specified in hexadecimal in the following format: 0xffff, 0Xfff, or xfff.

### Memory Usage Summary

The **-S** argument will provide a summary of all memory used in the system. This option will provide a memory usage summary based on data derived from all the system sde's. Memory consumed, free, locked, and that on the swap device will be printed. A limitation with this option is that a process may be created or terminated during the collection of memory data. This may cause some inconsistencies in the data provided. Therefore, this option may need to be run a few times to verify the accuracy of the information provided.

### SDE Audit

The **-A** option will run a quick audit of all the system sde's. It will run through all process' pcbs and through the system sde segment and print a list of segments that are allocated but not owned by any process. Note that one segment will always show up. This segment is basically the low core segment and its segid is always fixed "0x1a0000". A limitation of this option is that, if a process is created or terminated during the audit, many segments may appear in the output. Under this case, the tool may need to be run several times and only the segments that repeat are the ones "really" not owned by any process.

### Process Capabilities & Performance

The **-F** option provides information about each running supervisor process. The information is data gathered from each process' PCB segment. The data output contains the number of capabilities owned by the process. Capabilities are used and returned by the FMGR in response to open file and change directory requests. Other information about the processes current state is also output. In earlier versions, the DCT flags yield was output. Starting in load 1.2.25, the flags will be presented as characters. Only four flags will be output after load 25. They are I - for in memory, 0 - for swapped out on the disk, R - for running, S - for roadblocked (sleeping), r - for the process has run since being swapped in, and N - to indicate that the process has the "NOFIT" dcte flag set.

### FILES

msnap - the program normally resides in the /usr/bin directory.

libfile - an optional file containing shared library information.

### SEE ALSO

Programmer's Manual Volume 1, System Architecture, Section 5, contains information that will aid the user in interpreting the data provided by this tool. The user can also reference the following system header files for definition of the flags and other data provided by the program.

These files reside in \$OFC/head:

pcb.h

kpcb.h

## COMMANDS

235-700-200  
November 1998

### *MSNAP(1)*

sde.h

pgt.h

pge.h

sge.h

const.h

dtype.h

va.h

## WARNINGS

Since this program is a low level supervisor program, it may be interrupted by higher level activity while collecting data from the kernel in heavily loaded systems. Therefore, some inconsistencies may be noted for some of the data provided by the tool. Under these conditions, more than one measurement may be required to get meaningful results. However, since the program does run at the user level, it is noninterfering to system operation and can be used in live field sites.

NAME

`mv`

DESCRIPTION

See *cp*.



NAME

`newgrp` — log in to a new group

SYNOPSIS

`newgrp` [ - ] [ group ]

DESCRIPTION

*Newgrp* changes the group identification of its caller, analogously to *login* (1). The same person remains logged in, and the current directory is unchanged, but calculations of access permissions to files are performed with respect to the new group ID.

*Newgrp* without an argument changes the group identification to the group in the password file; in effect, it changes the group identification back to the caller's original group.

An initial - flag causes the environment to be changed to the one that would be expected if the user actually logged in again.

A password is demanded if the group has a password and the user does not, or if the group has a password and the user is not listed in **/etc/group** as being a member of that group.

FILES

`/etc/group`  
`/etc/passwd`

SEE ALSO

`login(1)`

BUGS

There is no convenient way to enter a password into **/etc/group**. Use of group passwords is not encouraged, because, by their very nature, they encourage poor security practices. Group passwords may disappear in the future.



NAME

`news` — print news items

SYNOPSIS

`news [ -a ] [ -n ] [ -s ] [ items ]`

DESCRIPTION

*News* is used to keep the user informed of current events. By convention, these events are described by files in the directory `/usr/news` .

When invoked without arguments, *news* prints the contents of all current files in `/usr/news` , most recent first, with each preceded by an appropriate header. *News* stores the “currency” time as the modification date of a file named `.news_time` in the user’s home directory (the identity of this directory is determined by the environment variable `$HOME` ); only files more recent than this currency time are considered “current.”

The `-a` option causes *news* to print all items, regardless of currency. In this case, the stored time is not changed.

The `-n` option causes *news* to report the names of the current items without printing their contents, and without changing the stored time.

The `-s` option causes *news* to report how many current items exist, without printing their names or contents, and without changing the stored time. It is useful to include such an invocation of *news* in one’s `.profile` file, or in the system’s `/etc/profile` .

All other arguments are assumed to be specific news items that are to be printed.

If a *delete* is typed during the printing of a news item, printing stops and the next item is started. Another *delete* within 1 second of the first causes the program to terminate.

FILES

`/etc/profile`  
`/usr/news/*`  
`$HOME/.news_time`



NAME

`nice` — run a command at low priority

SYNOPSIS

`nice` [ - increment ] command [ arguments ]

DESCRIPTION

*Nice* executes *command* with a lower CPU scheduling priority. If the *increment* argument (in the range 1 through 19) is given, it is used; if not, an increment of 10 is assumed.

The super user may run commands with priority higher than normal by using a negative increment, e.g., `—10` .

SEE

`nohup(1)`

DIAGNOSTICS

*Nice* returns the exit status of the subject command.

BUGS

An *increment* larger than 19 is equivalent to 19.



NAME

`nohup` — run a command immune to hang-ups and quits

SYNOPSIS

**nohup** command [ arguments ]

DESCRIPTION

*Nohup* executes *command* with hang-ups and quits ignored. If output is not redirected by the user, it will be sent to **nohup.out**. If **nohup.out** is not writable in the current directory, output is redirected to **\$HOME/nohup.out** .

SEE ALSO

`nice(1)`



NAME

od — octal dump

SYNOPSIS

od [ -bcdosx ] [ file ] [ [ + ]offset[ . ][ b ]]

DESCRIPTION

*Od* dumps *file* in one or more formats as selected by the first argument. If the first argument is missing, **-o** is default. The meanings of the format options are:

- b** Interpret bytes in octal.
- c** Interpret bytes in ASCII. Certain nongraphic characters appear as C escapes: null=**\0**, backspace=**\b**, form-feed=**\f**, new-line=**\n**, return=**\r**, tab=**\t**; others appear as 3-digit octal numbers.
- d** Interpret words in unsigned decimal.
- o** Interpret words in octal.
- s** Interpret 16-bit words in signed decimal.
- x** Interpret words in hex.

The *file* argument specifies which file is to be dumped. If no file argument is specified, the standard input is used.

The offset argument specifies the offset in the file where dumping is to commence. This argument is normally interpreted as octal bytes. If **.** is appended, the offset is interpreted in decimal. If **b** is appended, the offset is interpreted in blocks of 512 bytes. If the file argument is omitted, the offset argument must be preceded by **+**.

Dumping continues until end-of-file.



NAME

`openwd` — allow to open (with write) block device files for mounted file systems

SYNOPSIS

**`openwd`**

DESCRIPTION

*Openwd* allows processes to open with write any block devices for mounted file systems. This command must precede the processes that modify file systems via the block device access method, i.e., *fsdb*. This is to protect the mounted file systems from trashing due to erroneous or malicious write attempts via the block device access method to file systems. The window can be closed by using *closewd* or will be closed automatically by the file manager in 20 minutes.

SEE ALSO

`closewd(1)`, `fsdb(1)`



NAME

`parchk` — partition name check

SYNOPSIS

`parchk` [filename]

DESCRIPTION

The *parchk* command reads the file *filename*, or stdin if no filename is given, for a list of partition names. Each partition name in the input stream is searched for in the on-line SG data base. For each partition name matching an SG data base entry, *parchk* extracts the pack number, partition number, and block size from the partition record. This information is printed on stdout in the following format:

*<partition name> rt <pack number> <partition number> <block count>.*

If any of the named partitions are not located, then *parchk* will list that partition name as "not found."

DIAGNOSTICS

*Parchk* will fail with an error code if the on-line SG data base cannot be attached or read. Appropriate error messages are printed to stderr before *parchk* exits.

EXAMPLES

The command:

```
echo "no5text" | parchk
```

or:

```
echo "no5text" > file; parchk file
```

would result in output similar to:

```
no5text rt 0 19 175000
```

Indicating that the no5text partition is on rt (pack #) 0, is partition number 19, and has 175,000 blocks allocated for it.

FILES

/etc/parchk

CAVEATS

*parchk* assumes that each input line contains only one partition name.



NAME

`passwd` — change login password

SYNOPSIS

`passwd` [ *name* ]

DESCRIPTION

This command changes (or installs) a password associated with the current login or a specified login *name*.

The program prompts for the old password (if any) and then for the new one (twice).

The following requirements will be enforced on nonsuper users to ensure nontriviality:

1. Each password must have at least six characters.
2. Each password must contain at least two alphabetic characters and one numeric or special character. In this case, "alphabetic" means both uppercase and lowercase letters.
3. The new password may not be similar to the old password nor the login name. Similar things have three or more consecutive characters in the same or reverse order. These comparisons do not distinguish between uppercase and lowercase.

Only the owner of *name* or the super user may change a password; the owner must prove ownership by entering the old password. Only the super user can create a null password or disobey the nontriviality tests.

The password file is not changed if the new password is the same as the old password, or if the password is less than 2 weeks old.

CONSIDERATIONS

The password 'ages' 1 week at 00:00 GMT Thursday. This occurs on Wednesday, either late afternoon or early evening within the 48 contiguous states.

Even though the super user may disobey the nontriviality tests, it is strongly recommended that they be obeyed for the security of the system.

Configurations with dial-up access are strongly urged to assign a password for the login name 'daemon'. This password will be the machine password.

FILE

`/etc/passwd`

SEE ALSO

`admin(1)`, `login(1)`



NAME

paste — merge same lines of several files or subsequent lines of one file

SYNOPSIS

```
paste file1 file2 ...
paste -dlist file1 file2 ...
paste -s [-dlist] file1 file2 ...
```

DESCRIPTION

In the first two forms, *paste* concatenates corresponding lines of the given input files *file1*, *file2*, etc. It treats each file as a column or columns of a table and pastes them together horizontally (parallel merging). If you will, it is the counterpart of *cat (1)* which concatenates vertically, that is, one file after the other. In the last form above, *paste* replaces the function of an older command with the same name by combining subsequent lines of the input file (serial merging). In all cases, lines are glued together with the *tab* character, or with characters from an optionally specified *list*. Output is to the standard output, so it can be used as the start of a pipe, or as a filter, if *-* is used in place of a file name.

The meanings of the options are:

- d** Without this option, the new-line characters of each but the last file (or last line in case of the **-s** option) are replaced by a *tab* character. This option allows replacing the *tab* character by one or more alternate characters (see below).
- list* One or more characters immediately following **-d** replace the default *tab* as the line concatenation character. The *list* is used circularly, that is, when exhausted, it is reused. In parallel merging (that is, no **-s** option), the lines from the last file are always terminated with a new-line character, not from the *list*. The *list* may contain the special escape sequences: `\n` (new-line), `\t` (tab), `\` (backslash), and `\0` (empty string, not a null character). Quoting may be necessary, if characters have special meaning to the shell (for example, to get one backslash, use *-d*).
- s** Merge subsequent lines rather than one from each input file. Use *tab* for concatenation, unless a *list* is specified with **-d** option. Regardless of the *list*, the very last character of the file is forced to be a new-line.
- May be used in place of any file name, to read a line from the standard input. (There is no prompting).

EXAMPLES

```
ls | paste -d" " -      list directory in one column
ls | paste - - - -     list directory in four columns
paste -s -d "\t\n" file combine pairs of lines into lines
```

SEE ALSO

cut(1), grep(1), pr(1)

## COMMANDS

235-700-200  
November 1998

### *PASTE(1)*

#### DIAGNOSTICS

*line too long*

Output lines are restricted to 511 characters.

*too many files*

Except for *-s* option, no more than 12 input files may be specified.

NAME

`pg` — file perusal filter for soft-copy terminals

SYNOPSIS

`pg [- number] [-p string] [-cefns] [+linenumber] [+pattern] [files ...]`

DESCRIPTION

The `pg` command is a filter which allows the examination of *files* one screen-full at a time on a soft-copy terminal.

The file name - and/or NULL arguments indicate that `pg` should read from the standard input. Each screenfull is followed by a prompt. If the user types a carriage return, another page is displayed; other possibilities are described below.

This command is different from previous paginators in that it allows you to back up and review something that has already passed. The method for doing this is explained below.

Only the *vt100* family of terminals is supported.

The options are as follows:

- `-number` Specifies the size (in lines) of the window that `pg` is to use instead of the default size of 24 lines. The default number of columns is 80, which cannot be changed. (On a terminal containing 24 lines, the default window size is 23.)
- `-p string` Causes `pg` to use *string* as the prompt. If the prompt string contains a `%d`, the first occurrence of `%d` in the prompt is replaced by the current page number when the prompt is issued. The default prompt string is `: .`
- `-c` Home the cursor and clear the screen before displaying each page.
- `-e` Causes `pg` *not* to pause at the end of each file.
- `-f` Normally, `pg` splits lines longer than the screen width, but some sequences of characters in the text being displayed (i.e., escape sequences for underlining) generate undesirable results. The `-f` option inhibits `pg` from splitting lines.
- `-n` Normally, commands must be terminated by a `< newline >` character. This option causes an automatic end of command as soon as a command letter is entered.
- `-s` Causes `pg` to print all messages and prompts in standout mode (usually inverse video).
- `+ linenumber`  
Start up at *linenumber* .
- `+/ pattern /` Start up at the first line containing the regular expression pattern.

The responses that may be typed when `pg` pauses can be divided into three categories: those causing further perusal, those that search, and those that modify the perusal environment.

*PG(1)*

Commands which cause further perusal normally take a preceding *address*, an optionally signed number indicating the point from which further text should be displayed. This *.I* address is interpreted in either pages or lines depending on the command. A signed *address* specifies a point relative to the current page or line, and an unsigned *address* specifies an address relative to the beginning of the file. Each command has a default address that is used if none is provided. The perusal commands and their defaults are as follows:

[+|-][num]<newline> or <blank>

Using absolute addressing, page *num* will be displayed. Using relative addressing, the page *num* forward or backward from the current page will be displayed.

[+|-][num] l Using absolute addressing displays a page beginning at the specified line. Using relative addressing simulates scrolling the screen, forward or backward, the specified number of lines.

[+|-] d or ^D

Simulates scrolling half a screen forward or backward.

The following perusal commands take no *address*.

. or ^L Typing a single period causes the current page of text to be redisplayed.

\$ Displays the last windowfull in the file. Use with caution when the input is a pipe.

The following commands are available for searching for text patterns in the text. The regular expressions described in *ed (1)* are available. They must always be terminated by a < *newline* >, even if the option is specified.

i / pattern / Searches forward for the *i* th (default *i* =1) occurrence of *pattern*. Searching begins immediately after the current page and continues to the end of the current file, without wraparound.

i ^ pattern ^  
and

i ? pattern ? Searches backwards for the *i* th (default *i* =1) occurrence of *pattern*. Searching begins immediately before the current page and continues to the beginning of the current file, without wraparound.

After searching, *pg* normally displays the line found at the top of the screen. This can be modified by appending *m* or *b* to the search command to leave the line found in the middle or at the bottom of the window from now on. The suffix *t* can be used to restore the original situation.

The perusal environment can be modified with the following commands:

i n Begin perusing the *i* th next file in the command line. The *i* is an unsigned number, default value is 1.

i p Begin perusing the *i* th previous file in the command line. The *i* is an unsigned number, default is 1.

i w Display another window of text. If *i* is present, set the window size to *i*.

s filename Save the input in the named file. Only the current file being perused is saved. The white space between the *s* and *filename* is

optional. This command must always be terminated by a *< newline >*, even if the option is specified.

**h** Help by displaying an abbreviated summary of available commands.

*q* or *Q* Quit *pg*.

**!** *command* *Command* is passed to the shell, */bin/sh*. This command must always be terminated by a *< newline >*, even if the *-n* option is specified.

At any time when output is being sent to the terminal, the user can hit the quit key (normally control-*ü* or the interrupt (break) key. This causes *pg* to stop sending output, and display the prompt. Then the user may enter one of the above commands in the normal manner. Unfortunately, some output is lost when this is done, because any characters waiting in the terminal's output queue are flushed when the quit signal occurs.

If the standard output is not a terminal, then *pg* acts like the *cat* command, except that a header is printed before each file (if there is more than one).

#### EXAMPLE

An example of *pg* usage is:

```
lf -l | pg -p "(Page %d):"
```

#### NOTES

While waiting for terminal input, *pg* responds to **BREAK**, **DEL**, and **^** by terminating execution. Between prompts, however, these signals interrupt *pg's* current task and place the user in prompt mode. These should be used with caution when input is being read from a pipe, since an interrupt is likely to terminate the other commands in the pipeline.

#### SEE ALSO

*ed(1)*, *grep(1)*

#### CAVEATS

If terminal tabs are not set every eight positions, undesirable results may occur.

When using *pg* as a filter with another command that changes the terminal I/O options, terminal settings may not be restored correctly.



NAME

*pio* — physical I/O

SYNOPSIS

**pio** command

DESCRIPTION

*Pio* is executed using physical I/O to and from all files where possible. If physical I/O is not possible, the system does the appropriate side-buffering.

This command is useful for doing *check*, *dd*, and other *UNIX* system commands that make use of raw I/O.

EXAMPLE

*pio dd if = /dev/ofln of = /dev/mt00 obsize = 2048 count = 512*

SEE ALSO

*cp(1)*



NAME

**pkill** — terminate a process (superuser)

SYNOPSIS

**pkill** pid

DESCRIPTION

*Pkill* creates a new supervisor process from the file in */prc/pkill*. A message with the **pid** of the process to be killed is sent to the new supervisor process. The supervisor process sends a terminate message to the process manager to take down the process designated by pid.

SEE ALSO

*ps(1)*

LIMITATIONS

*Pkill* should only be used as a last resort to terminate a process. In cases where it is necessary to forcibly terminate a copy of a shared supervisor, some allocated resources may not be freed (buffers, semaphores, and so on). Therefore a degradation of service may occur for all users sharing that supervisor.

In short, *pkill* of a supervisor should only be done to force a core dump of a particular *UNIX* system process before you **reboot**. If you *pkill* a nonexistent or a nonkillable process ID, an error message is generated.



NAME

`pr` — print files

SYNOPSIS

`pr` [ options ] [ files ]

DESCRIPTION

*Pr* prints the named files on the standard output. If *file* is -, or if no files are specified, the standard input is assumed. By default, the listing is separated into pages, each headed by the page number, a date and time, and the name of the file.

By default, columns are of equal width, separated by at least one space; lines which do not fit are truncated. If the **-s** option is used, lines are not truncated and columns are separated by the separation character.

If the standard output is associated with a terminal, error messages are withheld until *pr* has completed printing.

The *options* below may appear singly or be combined in any order:

- +k**           Begin printing with page *k* (default is 1).
- k**           Produce *k* -column output (default is 1). The options **-e** and **-i** are assumed for multicolumn output.
- a**           Print multicolumn output across the page.
- m**           Merge and print all files simultaneously, one per column (overrides the **-k**, and **-a** options).
- d**           Double-space the output.
- eck**        Expand *input* tabs to character positions *k* "+1, 2\*" *k* "+1, 3\*" *k* +1, etc. If *k* is 0 or is omitted, default tab settings at every eighth position are assumed. Tab characters in the input are expanded into the appropriate number of spaces. If *c* (any nondigit character) is given, it is treated as the input tab character (default for *c* is the tab character).
- ick**        In *output*, replace white space wherever possible by inserting tabs to character positions *k* "+1, 2\*" *k* "+1, 3\*" *k* +1, etc. If *k* is 0 or is omitted, default tab settings at every eighth position are assumed. If *c* (any nondigit character) is given, it is treated as the output tab character (default for *c* is the tab character).
- nck**        Provide *k* -digit line numbering (default for *k* is 5). The number occupies the first *k* +1 character positions of each column of normal output or each line of **-m** output. If *c* (any nondigit character) is given, it is appended to the line number to separate it from whatever follows (default for *c* is a tab).
- wk**        Set the width of a line to *k* character positions (default is 72 for equal-width multi-column output, no limit otherwise).
- ok**        Offset each line by *k* character positions (default is 0). The number of character positions per line is the sum of the width and offset.
- lk**        Set the length of a page to *k* lines (default is 66).

*PR(1)*

- h** Use the next argument as the header to be printed instead of the file name.
- p** Pause before beginning each page if the output is directed to a terminal (*pr* will ring the bell at the terminal and wait for a carriage return).
- f** Use form-feed character for new pages (default is to use a sequence of line-feeds). Pause before beginning the first page if the standard output is associated with a terminal.
- r** Print no diagnostic reports on failure to open files.
- t** Print neither the 5-line identifying header nor the 5-line trailer normally supplied for each page. Quit printing after the last line of each file without spacing to the end of the page.
- sc** Separate columns by the single character *c* instead of by the appropriate number of spaces (default for *c* is a tab).

## EXAMPLES

Print **file1** and **file2** as a double-spaced, 3-column listing headed by "file list":

```
pr -3dh "file list" file1 file2
```

Write **file1** on **file2**, expanding tabs to columns 10, 19, 28, 37, ... :

```
pr -e9 -t <file1 >file2
```

## FILES

/dev/tty\* to suspend messages

## SEE ALSO

cat(1)

NAME

*ps* — report process status

SYNOPSIS

**ps** [ *aklxp* ] [ *s* *swapdev* ] *p* *tlist* ]

DESCRIPTION

*Ps* prints certain information about active processes. Without *options*, information is printed about processes associated with the current terminal. The output consists of a short listing containing only the terminal identifier, the process ID, and the command name. Otherwise, the information that is displayed is controlled by the selection of the following options:

- a** Print information about all supervisor processes associated with terminals.
- k** Print information about all kernel processes.
- l** Print a long listing
- x** Print information about all supervisor processes not associated with a terminal.
- p** Include parent process number (cannot be used with **k** or **l** flags).
- s***swapdev* Use the file *swapdev* in place of **/dev/swap**. This is useful when examining a *corefile*.
- tlist** Restrict listing to data about the processes associated with the terminals given in *tlist*. *Tlist* is a list of terminal identifiers separated from one another by a space.

The column headings and the meaning of the columns in a *ps* listing of supervisor processes are given below:

- PSTAT** status flags.
- PRI** current priority.
- KTIME** time spent in the kernel in milliseconds.
- KPTIM** time spent in kernel processes in milliseconds.
- STIME** time spent in supervisor processes in milliseconds.
- TTY** terminal associated with the process. Normally this is the same as the control channel for the process. If the control channel corresponds to a nonprinting character, a '?' is printed in this field.
- PID** process number of this process.
- PID** process number of the parent process.
- UTID** utility ID of the process.
- SIZE** size of process in bytes.
- SLEEP** bit pattern on which the process is sleeping.
- CMD** command line used to invoke the process.

The column headings and the meanings of the columns in a *ps* listing of kernel processes are given below:

## COMMANDS

235-700-200  
November 1998

### *PS(1)*

<b>PSTATE</b>	process status flag.
<b>PRI</b>	execution level.
<b>TOUT</b>	real-time clock value for next time-out.
<b>RTOUT</b>	interval between repetitive time-outs in milliseconds.
<b>EVENTS</b>	event word.
<b>CHAN</b>	control channel.
<b>PID</b>	process number.
<b>ADDR</b>	address of PCB segment descriptor entry.
<b>DCT</b>	dispatcher control table index for the process.
<b>UTID</b>	utility ID of the process.
<b>DEVICE</b>	name of controller, device or process.

### FILES

/dev/kmem  
/dev/pmem  
/dev/swap  
/dmrt/kprc

### SEE ALSO

kill(1)

### NOTE

Things can change while *ps* is running; the picture it gives is only a close approximation to reality.

NAME

*pst* — processor resource status and timing information (a performance measurement tool)

SYNOPSIS

**pst** {-Ddelay | -W} {-TkstafcL | -P pid pidn | -U uid1 uidn } [-Oks] [-I] [-pport]

DESCRIPTION

*Pst* is a tool that prints information on the status of system resources and/or the activity of running processes on an active 3B20D computer system. The process is a stand-alone users-level process that collects performance information from the *UNIX* RTR operating system kernel. It is noninterfering to application environments.

*Pst* can collect run time information on active processes. If requested, it will print the percentage of time spent in the process and in the kernel on behalf of the process (servicing ost requests). This percentage is based on the time consumed by the process over the snapshot interval. These reports will print the process pid and utility id. The process name and arguments must be obtained from a ps or other listing (see -L flag below). This is done to minimize overhead in the collection of data.

The tool works by collecting two samples of information from the kernel's address space. The two samples or snapshots are separated by a user definable time delay. The first set of data is collected when the tool is first started or after it receives the first E\_USR event (see -D and -W options). The tool then sleeps for the desired number of seconds or waits for a second E\_USR event. It then collects the second set of data. After the second set of data is collected, the process calculates the differences between the two sets of data and outputs the results to stdout.

The -T (type) arguments are:

- k** This option will provide timing information for all running kernel processes. The execution level, pid, uid, and timing information is provided.
- s** This option will provide timing information for all active supervisor level processes. The tty, initial and current priority, pid, uid, and timing information is provided.
- t** This option provides kernel timing and resource information. A global view of real time usage is provided in histogram format for each of the 16 execution levels supported by the operating system. The consumption of major system resources is provided over the requested snapshot interval.
- a** This option provides all of the t, s, and k data.
- c** This option provides a bar chart of real time consumption for each execution level.
- f** This option provides a bar chart of real time availability at each execution level.
- L** This option provides a *long format* report. For kernel processes, the

**PST(1)**

process name is output. For supervisor processes, it includes the process name, swap and dispatch counts, along with other information. This option must be used with caution since it causes *pst* to add all pcb segments to its address space. This can cause the *pst* program to grow to a large size. Under times of extreme memory overload, the *pst* program could hang if there is not sufficient swappable memory for it to fit in the processor memory.

The **-Ddelay** parameter is used to specify the measurement interval. *delay* is the time in seconds between the two snapshots.

The **-W** option is an alternative to the **-Ddelay** argument. This option allows the *pst* tool to be synchronized with other activity through E\_USR events. When this argument is used, the program will wait for a E\_USR event before taking each of the two snapshots.

The **-U** parameter allows the user to specify a list of up to 20 utility ids of processes to be snapshot. The data will be the same as the **-Ts** or **-Tk** arguments, but only the processes specified will be snapped. Therefore, this option cannot be used with the **-Ts** or **-Tk** parameters. All uid arguments must be in hex (prefixed with X, x, 0X, or 0x).

The **-P** parameter is the same as the **-U** parameter with the exception that process ids are specified instead of uids. The same limitation exists that the **-Ts** and **-Tk** options cannot be used with the **-P** option. These must be specified in decimal format.

The **-O** parameter specifies that OST usage counts are requested. The **k** and/or **s** flags specify that kernel and/or supervisor ost counts be reported. At least one of the **k** or **s** flags must follow the **-O** argument.

The **-I** parameter will provide interrupt usage information. Both software and hardware interrupt counts will be provided.

The **-pport** parameter specifies that *pst* should attach to the system port **port**. This parameter is useful when used with the **-W** argument or to guarantee that only one copy of the process is executing at any time.

## FILES

The *pst* command may be found in /usr/bin on the 3B20D computer.

NAME

`pwd` — working directory name

SYNOPSIS

`pwd`

DESCRIPTION

*Pwd* prints the path name of the working (current) directory.

DIAGNOSTICS

“Cannot open ..” and “Read error in ..” indicate possible file system trouble and should be referred to a *UNIX* system programming counselor.



235-700-200  
November 1998

COMMANDS

*RED(1)*

NAME

red

DESCRIPTION

See *ed*.



NAME

`rm, rmdir` — remove files or directories

SYNOPSIS

```
rm [-fri] file ...  
rmdir dir ...
```

DESCRIPTION

*Rm* removes the entries for one or more files from a directory. If an entry was the last link to the file, the file is destroyed. Removal of a file requires write permission in its directory, but neither read nor write permission on the file itself.

If a file has no write permission and the standard input is a terminal, its permissions are printed and a line is read from the standard input. If that line begins with **y**, the file is deleted; otherwise, the file remains. No questions are asked when the **-f** option is given or if the standard input is not a terminal.

If a designated file is a directory, an error comment is printed unless the optional argument **-r** has been used. In that case, *rm* recursively deletes the entire contents of the specified directory, and the directory itself.

If the **-i** (interactive) option is in effect, *rm* asks whether to delete each file; and, under **-r**, whether to examine each directory.

*Rmdir* removes entries for the named directories, which must be empty.

DIAGNOSTICS

Generally, it is self-explanatory. It is forbidden to remove the file `..` merely to avoid the antisocial consequences of inadvertently doing something like the following:

```
rm -r .*
```



235-700-200  
November 1998

COMMANDS

*RMAIL(1)*

NAME

rmail

DESCRIPTION

*See mail.*



235-700-200  
November 1998

COMMANDS

*RMDIR(1)*

NAME

`rmdir`

DESCRIPTION

See *rm*.



NAME

rsh

DESCRIPTION

See *sh*.



NAME

`run`, `urun` — run an environment (superuser), run a user level process

SYNOPSIS

`run [-b] [-f]file`  
`run [-b] file`

DESCRIPTION

*Run* starts up a new environment (task) either as a kernel process or a supervisor process. Although *run* can execute a user process, such use is not recommended, since the process will not be started with a user environment.

*Urun* forks and execs a *UNIX* system user-level process.

Two optional flags are as follows:

- b** indicates the process is to be spawned in the background and the death of the child process is not to be waited for. Normally, the death of a child process is waited for.
- f** indicates a nonzero value is to be passed in a message to the new process created from the file contents by the process manager.

LIMITATIONS

A process that is run using this will inherit the channel *id* of its parent (the shell). Thus, events generated for the shell's channel (such as the *E\_INT* caused by hitting the break key) will be sent to the process.



NAME

*sdiff* — side-by-side difference program

SYNOPSIS

**sdiff** [ options ... ] file1 file2

DESCRIPTION

*Sdiff* uses the output of *diff*(1) to produce a side-by-side listing of two files indicating those lines that are different. Each line of the two files is printed with a blank gutter between them if the lines are identical; a < in the gutter if the line only exists in *file1*; a > in the gutter if the line only exists in *file2*; and a | for lines that are different.

For example:

```
x | y
a a
b <
c <
d d
> c
```

The following options exist:

- wn** Use the next argument, *n*, as the width of the output line. The default line length is 130 characters.
- l** Only print the left side of any lines that are identical.
- s** Do not print identical lines.
- ooutput** Use the next argument, *output*, as the name of a third file that is created as a user-controlled merging of *file1* and *file2*. Identical lines of *file1* and *file2* are copied to *output*. Sets of differences, as produced by *diff*(1), are printed; where a set of differences share a common gutter character. After printing each set of differences, *sdiff* prompts the user with a % and waits for one of the following user-typed commands:
  - l** append the left column to the output file.
  - r** append the right column to the output file.
  - s** turn on silent mode; do not print identical lines.
  - v** turn off silent mode.
  - e l** call the editor with the left column.
  - e r** call the editor with the right column.
  - e b** call the editor with the concatenation of left and right.
  - e** call the editor with a zero length file.
  - q** exit from the program.

## COMMANDS

235-700-200  
November 1998

### *SDIFF(1)*

On exit from the editor, the resulting file is concatenated on the end of the *output* file.

### SEE ALSO

diff(1), ed(1)

NAME

`sed` — stream editor

SYNOPSIS

`sed [ -n ] [ -e script ] [ -f sfile ] [ files ]`

DESCRIPTION

*Sed* copies the named *files* (standard input default) to the standard output; they are edited according to a script of commands. The `-f` option causes the script to be taken from file *sfile*; these options accumulate. If there is just one `-e` option and no `-f` options, the flag `-e` may be omitted. The `-n` option suppresses the default output. A script consists of editing commands, one per line, of the following form:

```
[ address [ , address ] ] function [ arguments ]
```

In normal operation, *sed* cyclically copies a line of input into a *pattern space* (unless there is something left after a **D** command); applies in sequence all commands whose *addresses* select that pattern space; and, at the end of the script, copies the pattern space to the standard output (except under `-n`) and deletes the pattern space.

Some of the commands use a *hold space* to save all or part of the *pattern space* for subsequent retrieval.

An *address* is either a decimal number that counts input lines cumulatively across files; a **\$** that addresses the last line of input; or a context address, i.e., a */regular expression/* in the style of *ed*(1) modified thus:

- In a context address, the construction `\?regular expression?`, where *?* is any character, is identical to */regular expression/*. Note that in the context address `\xabc \xdefx`, the second **x** stands for itself so that the regular expression is **abcxdef**.
- The escape sequence `\n` matches a new line *embedded* in the pattern space.
- A period `.` matches any character except the *terminal* new line of the pattern space.
- A command line with no addresses selects every pattern space.
- A command line with one address selects each pattern space that matches the address.
- A command line with two addresses selects the inclusive range from the first pattern space that matches the first address through the next pattern space that matches the second. (If the second address is a number less than or equal to the line number first selected; only one line is selected.) Thereafter, the process is repeated, looking again for the first address.

Editing commands can be applied only to nonselected pattern spaces by use of the negation function `!`.

In the following list of functions, the maximum number of permissible addresses for each function is indicated in parentheses.

**SED(1)**

The *text* argument consists of one or more lines; all but the last one end with \ to hide the new-line. Backslashes in text are treated like backslashes in the replacement string of an **s** command, and may be used to protect initial blanks and tabs against the stripping that is done on every script line. The *rfile* or *wfile* argument must terminate the command line and must be preceded by exactly one blank. Each *wfile* is created before processing begins. At most, there can be ten distinct *wfile* arguments.

**(1)a\**

*text* Append. Place *text* on the output before reading the next input line.

**(2)b***label* Branch to the **:** command bearing the *label*. If *label* is empty, branch to the end of the script.

**(2)c\**

*text* Change. Delete the pattern space. With 0 or 1 address or at the end of a 2-address range, place *text* on the output. Start the next cycle.

**(2)d** Delete the pattern space. Start the next cycle.

**(2)D** Delete the initial segment of the pattern space through the first new line. Start the next cycle.

**(2)g** Replace the contents of the pattern space by the contents of the hold space.

**(2)G** Append the contents of the hold space to the pattern space.

**(2)h** Replace the contents of the hold space by the contents of the pattern space.

**(2)H** Append the contents of the pattern space to the hold space.

**(1) i\**

*text* Insert. Place *text* on the standard output.

**(2)l** List the pattern space on the standard output in an unambiguous form. Nonprinting characters are spelled in 2-digit ASCII and long lines are folded.

**(2)n** Copy the pattern space to the standard output. Replace the pattern space with the next line of input.

**(2)N** Append the next line of input to the pattern space with an embedded new line. (The current line number changes.)

**(2)p** Print. Copy the pattern space to the standard output.

**(2)P** Copy the initial segment of the pattern space through the first new line to the standard output.

**(1)q** Quit. Branch to the end of the script. Do not start a new cycle.

**(2)r** *rfile* Read the contents of *rfile*. Place them on the output before reading the next input line.

**(2)s** */regular expression/replacement/flags*

Substitute the *replacement* string for instances of the *regular expression* in the pattern space. Any character may be used instead of */*. For a fuller description, see *sed* (1). *Flags* is zero or more of:

- g** Global. Substitute for all nonoverlapping instances of the *regular expression* rather than just the first one.
- p** Print the pattern space if a replacement was made.
- wfile** Write. Append the pattern space to *wfile* if a replacement was made.
- (2)**t label** Test. Branch to the **:** command bearing the *label* if any substitutions have been made since the most recent reading of an input line or execution of **a**. If *label* is empty, branch to the end of the script.
- (2)**w wfile** Write. Append the pattern space to *wfile* .
- (2)**x** Exchange the contents of the pattern and hold spaces.
- (2)**y /string1/string2/**  
Transform. Replace all occurrences of characters in *string1* with the corresponding character in *string2* . The lengths of *string1* and *string2* must be equal.
- (2)**! function**  
Don't. Apply the *function* (or group, if *function* is **{}**) only to lines *not* selected by the address(es).
- (0): *label* This command does nothing; it bears a *label* for **b** and **t** commands to branch to.
- (1)= Place the current line number on the standard output as a line.
- (2){ Execute the following commands through a matching **}** only when the pattern space is selected.
- (0) An empty command is ignored.

SEE ALSO

ed(1), grep(1)



NAME

sh, rsh — shell, the standard/restricted command programming language

SYNOPSIS

sh [ -ceinrstuvx ] [ args ]  
rsh [-ceinrstuvx ] [ args ]

DESCRIPTION

*sh* it is used to set up login names and execution environments whose capabilities are more controlled than those of the standard shell. *Rsh* is a restricted version of the standard command interpreter. *Sh* is a command programming language that executes commands read from a terminal or a file. See *Invocation* for the meaning of arguments to the shell.

Commands

A *simple-command* is a sequence of nonblank *words* separated by *blanks* (a *blank* is a tab or a space). The first word specifies the name of the command to be executed. Except as specified below, the remaining words are passed as arguments to the invoked command. The command name is passed as argument 0. The *value* of a simple-command is its exit status if it terminates normally, or (octal) 200+*status* if it terminates abnormally [see *kill* (1) for a list of status values].

A *pipeline* is a sequence of one or more *commands* separated by ^ (or by a |). The standard output of each command but the last is connected by a pipe to the standard input of the next command. Each command is run as a separate process; the shell waits for the last command to terminate.

A *list* is a sequence of one or more pipelines separated by ;, &, &&, or ||, and optionally terminated by ; or &. Of these four symbols, ; and & have equal precedence, which is lower than that of && and ||. The symbols && and || also have equal precedence. A semicolon (;) causes sequential execution of the preceding pipeline; an ampersand (&) causes asynchronous execution of the preceding pipeline (that is, the shell does *not* wait for that pipeline to finish). The symbol && ( || ) causes the *list* following it to be executed only if the preceding pipeline returns a zero (nonzero) exit status. An arbitrary number of new lines may appear in a *list*, instead of semicolons, to delimit commands.

A *command* is either a simple-command or one of the following. Unless otherwise stated, the value returned by a command is that of the last simple-command executed in the command.

**for** *name in word ... do list done*

Each time a **for** command is executed, *name* is set to the next *word* taken from the **in word** list. If **in word ...** is omitted, then the **for** command executes the **do list** once for each positional parameter that is set (see *Parameter Substitution* below). Execution ends when there are no more words in the list.

**case** *word in pattern | pattern ...*) *list* ;; ... **esac**

A **case** command executes the *list* associated with the first *pattern* that matches *word*. The form of the patterns is the same as that used for file-name generation (see *File Name Generation*).

*SH(1)***if** *list* **then** *list* **elif** *list* **then** *list* ... **else** *list* **fi**

The *list* following **if** is executed; and, if it returns a zero exit status, the *list* following the first **then** is executed. Otherwise, the *list* following **elif** is executed; and, if its value is zero, the *list* following the next **then** is executed. Failing that, the **else** *list* is executed. If no **else** *list* or **then** *list* is executed, then the **if** command returns a zero exit status.

**while** *list* **do** *list* **done**

A **while** command repeatedly executes the **while** *list* ; and, if the exit status of the last command in the list is zero, executes the **do** *list*. Otherwise the loop terminates. If no commands in the **do** *list* are executed, then the **while** command returns a zero exit status; **until** may be used in place of **while** to negate the loop termination test.

(*list*)      Execute *list* in a subshell.

{*list*;}      *list* is simply executed.

The following reserved words are only recognized as the first word of the command and then not quoted:

```
if n if then else elif fi case esac for while until do done { }
```

## Comments

A word beginning with **#** causes that word and all the following characters up to a new line to be ignored.

## Aliasing

The first word of each command is replaced by the text of an alias if an alias for this word has been defined. The replacement string can contain any character excluding ";". Aliases can be used to redefine special built-in commands but cannot be used to redefine the reserved words. Aliases can be created, modified, and listed with the **alias** command and can be removed with the **unalias** command.

Aliasing is performed when scripts are read, not while they are executed. Therefore, for an alias to take effect the alias definition command has to be executed before the command which references the alias is read.

Aliases are frequently used as a short hand for full path names or to customize the behavior of other shell commands.

## Command Substitution

The standard output from a command enclosed in a pair of grave accents (`) may be used as part or all of a word; trailing new lines are removed.

## Parameter Substitution

The character **\$** is used to introduce substitutable *parameters*. Positional parameters may be assigned values by **set**. Variables may be set by writing:

```
name = value
```

```
name = value
```

...

Pattern-matching is not performed on *value* .

**`${parameter}`**

A *parameter* is a sequence of letters, digits, or underscores (a *name*); a digit, or any of the characters `*`, `@`, `#`, `?`, `-`, `$`, and `!`. The value, if any, of the parameter is substituted. The braces are required only when *parameter* is followed by a letter, digit, or underscore that is not to be interpreted as part of its name. A *name* must begin with a letter or underscore. If *parameter* is a digit, then it is a positional parameter. If *parameter* is `*` or `@` , then all the positional parameters, starting with `$1`, are substituted (separated by spaces). Parameter `$0` is set from argument zero when the shell is invoked.

**`${parameter :-word}`**

If *parameter* is set and is nonnull, then substitute its value; otherwise, substitute *word*.

**`${parameter :=word}`**

If *parameter* is not set or is null, then set it to *word*; the value of the parameter is then substituted. Positional parameters may not be assigned to in this way.

**`${parameter :?word}`**

If *parameter* is set and is nonnull, then substitute its value; otherwise, print *word* and exit from the shell. If *word* is omitted, then the message "parameter null or not set" is printed.

**`${parameter :+word}`**

If *parameter* is set and is nonnull, then substitute *word*; otherwise, substitute nothing.

In the above, *word* is not evaluated unless it is to be used as the substituted string; so that, in the following example, `pwd` is executed only if `d` is not set or is null:

```
echo ${d:-pwd}
```

If the colon (`:`) is omitted from the above expressions, then the shell only checks whether *parameter* is set or not.

The following parameters are automatically set by the shell:

- #** The number of positional parameters in decimal.
- Flags supplied to the shell on invocation or by the `set` command.
- ?** The decimal value returned by the last synchronously executed command.
- \$** The process number of this shell.
- !** The process number of the last background command invoked.

The following parameters are used by the shell:

- HOME** The default argument (home directory) for the `cd` command.

*SH(1)*

<b>PATH</b>	The search path for commands (see <i>Execution</i> below). The user may not change <b>PATH</b> if executing under <i>rsh</i> .
<b>MAIL</b>	If this variable is set to the name of a mail file, then the shell informs the user of the arrival of mail in the specified file.
<b>PS1</b>	Primary prompt string, by default “ \$ ”.
<b>PS2</b>	Secondary prompt string, by default “ > ”.
<b>IFS</b>	Internal field separators, normally <b>space</b> , <b>tab</b> , and <b>new-line</b> .
<b>ENV</b>	If this parameter is set, the value is used as the pathname to the script that will be executed when the shell is invoked. (See <i>invocation</i> .) This file is typically used for alias definitions.

The shell gives default values to **PATH**, **PS1**, **PS2**, and **IFS**, while **HOME** and **MAIL** are not set at all by the shell [although **HOME** is set by *login* (1)].

Blank Interpretation

After parameter and command substitution, the results of substitution are scanned for internal field separator characters (those found in **IFS**) and split into distinct arguments where such characters are found. Explicit null arguments (“” or ‘’) are retained. Implicit null arguments (those resulting from *parameters* that have no values) are removed.

File Name Generation

Following substitution, each command *word* is scanned for the characters \*, ?, and [. If one of these characters appears, then the word is regarded as a *pattern*. The word is replaced with alphabetically sorted file names that match the pattern. If no file name is found that matches the pattern, then the word is left unchanged. The character . at the start of a file name or immediately following a /, as well as the character / itself, must be matched explicitly.

*	Matches any string, including the null string.
?	Matches any single character.
[...]	Matches any one of the enclosed characters. A pair of characters separated by - matches any character lexically between the pair, inclusive.

Quoting

The following characters have a special meaning to the shell and cause termination of a word unless quoted:

; & ( ) | ^ < > new-line space tab

A character may be *quoted* (that is, made to stand for itself) by preceding it with a \. The pair **new-line** is ignored. All characters enclosed between a pair of single quote marks (’), except a single quote, are quoted.

Inside double quote marks (“”), parameter and command substitution occurs; \ quotes the characters \, , , and \$. "\$\*" is equivalent to "\$1 \$2 ..."; whereas "\$@" is equivalent to "\$1" "\$2" ....

## Prompting

When used interactively, the shell prompts with the value of **PS1** before reading a command. If, at any time a new line is typed and further input is needed to complete a command, then the secondary prompt (that is, the value of **PS2**) is issued.

## Input/Output

Before a command is executed, its input and output may be redirected using a special notation interpreted by the shell. The following may appear anywhere in a simple-command, or may precede or follow a *command*; and are *not* passed on to the invoked command. Substitution occurs before *word* or *digit* is used:

- <word**      Use file *word* as standard input (file descriptor 0).
- >word**      Use file *word* as standard output (file descriptor 1). If the file does not exist then it is created; otherwise, it is truncated to zero length.
- >>word**     Use file *word* as standard output. If the file exists, then output is appended to it (by first seeking to the end-of-file); otherwise, the file is created.
- <<-word**    The shell input is read up to a line that is the same as *word*, or to an end-of-file. The resulting document becomes the standard input. If any character of *word* is quoted, then no interpretation is placed upon the characters of the document; otherwise, parameter and command substitution occurs, (unescaped) **\new-line** is ignored, and **\** must be used to quote the characters **\**, **\$**, **,**, and the first character of *word*. If **-** is appended to **<<**, then all leading tabs are stripped from *word* and from the document.
- <&digit**     The standard input is duplicated from file descriptor *digit*. Similarly for the standard output using **>**.
- <&-**         The standard input is closed. Similarly for the standard output using **>**.

If one of the above is preceded by a digit, then the file descriptor created is that specified by the digit (instead of the default 0 or 1). For example:

```
... 2 >&1
```

creates file descriptor 2 that is a duplicate of file descriptor 1.

If a command is followed by **&**, then the default standard input for the command is the empty file **/dev/null**. Otherwise, the environment for the execution of a command contains the file descriptors of the invoking shell as modified by input/output specifications.

Redirection of output is not allowed in the restricted shell.

## Environment

The *environment* is a list of name-value pairs that is passed to an executed program in the same way as a normal argument list. The shell interacts with the environment in several ways. On invocation, the shell scans the environment and creates a parameter for each name found, giving it the

*SH(1)*

corresponding value. Executed commands inherit the same environment. If the user modifies the values of these parameters or creates new ones, none of these affects the environment unless the **export** command is used to bind the shell's parameter to the environment. The environment seen by any executed command is thus composed of any unmodified name-value pairs originally inherited by the shell, plus any modifications or additions, all of which must be noted in **export** commands.

The environment for any *simple-command* may be augmented by prefixing it with one or more assignments to parameters, thus:

```
TERM=450 cmd args          and
(export TERM; TERM=450; cmd args)
```

are equivalent (as far as the above execution of *cmd* is concerned).

If the **-k** flag is set, *all* keyword arguments are placed in the environment, even if they occur after the command name. The following first prints **a=b c** and then **c** :

```
echo a=b c
set -k
echo a=b c
```

## Signals

The INTERRUPT and QUIT signals for an invoked command are ignored if the command is followed by **&**; otherwise, signals have the values inherited by the shell from its parent, with the exception of signal 11 (but see also the **trap** command below).

## Execution

Each time a command is executed, the above substitutions are carried out. Except for the *Special Commands* listed below, a new process is created and an attempt is made to execute the command.

The shell parameter **PATH** defines the search path for the directory containing the command. Alternative directory names are separated by a colon (:). The default path is **:/bin:/usr/bin** (specifying the current directory, **/bin**, and **/usr/bin**, in that order). Note that the current directory is specified by a null path name, which can appear immediately after the equal sign or between the colon delimiters anywhere else in the path list. If the command name contains a **/**, then the search path is not used; such commands will not be executed by the restricted shell. Otherwise, each directory in the path is searched for an executable file. If the file has execute permission but is not an **a.out** file, it is assumed to be a file containing shell commands. A subshell (that is, a separate process) is spawned to read it. A parenthesized command is also executed in a subshell.

### Special Commands

The following commands are executed in the shell process and, except as specified, no input/output redirection is permitted for such commands:

- :** No effect; the command does nothing. A zero exit code is returned.
- . *file*** Read and execute commands from *file* and return. The search path specified by **PATH** is used to find the directory containing *file*.
- alias [*name*[=*value*] ... ]**  
Alias with no arguments prints the list of aliases in the form *name=value* on standard output. If a name is specified with no value, the name and value of the alias is printed. If both a name and value are specified, an alias is defined. There are some size limitations on the alias name and value. **name** can be no more than 40 characters. **value** can contain up to 10 space separated strings, each of which can be up to 40 characters. A value containing spaces or other shell special characters must be quoted.
- break *n*** Exit from the enclosing **for** or **while** loop, if any. If *n* is specified, then break *n* levels.
- continue *n*** Resume the next iteration of the enclosing **for** or **while** loop. If *n* is specified, then resume at the *n*th enclosing loop.
- cd *arg*** Change the current directory to *arg*. The shell parameter **HOME** is the default *arg*. The default path is **<null>** (specifying the current directory). Note that the current directory is specified by a null path name, which can appear immediately after the equal sign or between the colon delimiters anywhere else in the path list. If *arg* begins with a /, then the search path is not used. Otherwise, each directory in the path is searched for *arg*. The **cd** command may not be executed by *rsh*.
- eval *arg* ...** The arguments are read as input to the shell and the resulting command(s) executed.
- exec *arg* ...** The command specified by the arguments is executed in place of this shell without creating a new process. Input/output arguments may appear and, if no other arguments are given, cause the shell input/output to be modified.
- exit *n*** Causes a shell to exit with the exit status specified by *n*. If *n* is omitted, then the exit status is that of the last command executed (an end-of-file will also cause the shell to exit). The exit command is ignored in interactive (connected to a TTY) shells.
- export *name* ...**  
The given *names* are marked for automatic export to the *environment* of subsequently-executed commands. If no arguments are given, then a list of all names that are exported in this shell is printed.
- newgrp *arg* ...**  
Equivalent to **exec newgrp *arg* ...**
- read *name* ...**  
One line is read from the standard input and the first word is assigned to the first *name*, the second word to the second *name*; etc., with leftover words assigned to the last *name*. The return code is 0 unless an end-of-file is encountered.

## SH(1)

**readonly** *name* ...

The given *names* are marked *readonly* and the values of these *names* may not be changed by subsequent assignment. If no arguments are given, then a list of all *readonly* names is printed.

**set -ekntuvx** *arg* ...

- e** Exit immediately if a command exits with a nonzero exit status.
- k** All keyword arguments are placed in the environment for a command, not just those that precede the command name.
- n** Read commands but do not execute them.
- t** Exit after reading and executing one command.
- u** Treat unset variables as an error when substituting.
- v** Print shell input lines as they are read.
- x** Print commands and their arguments as they are executed.
- Do not change any of the flags; useful in setting **\$1** to **-**.

Using **+** rather than **-** causes these flags to be turned off. Also, these flags can be used upon invocation of the shell. The current set of flags may be found in **\$-**. The remaining arguments are positional parameters and are assigned, in order, to **\$1**, **\$2**, .... If no arguments are given, then the values of all names are printed.

**shift** *n* The positional parameters from **\$n+1** ... are renamed **\$1** .... If *n* is not given, it is assumed to be 1.

**test** **test** *expr* [ *expr* ]  
*Test* evaluates the expression *expr* and, if its value is true, returns a zero (true) exit status; otherwise, a nonzero (false) exit status is returned; *test* also returns a nonzero exit status if there are no arguments. The following primitives are used to construct *expr* :

- r** *file* true if *file* exists and is readable.
- w** *file* true if *file* exists and is writable.
- x** *file* true if *file* exists and is executable.
- f** *file* true if *file* exists and is a regular file.
- d** *file* true if *file* exists and is a directory.
- c** *file* true if *file* exists and is a character special file.
- b** *file* true if *file* exists and is a block special file.
- p** *file* true if *file* exists and is a named pipe (fifo).
- u** *file* true if *file* exists and its set-user-ID bit is set.
- g** *file* true if *file* exists and its set-group-ID bit is set.
- k** *file* true if *file* exists and its sticky bit is set.
- s** *file* true if *file* exists and has a size greater than zero.

- t** [ *fildev* ] true if the open file whose file descriptor number is *fildev* (1 by default) is associated with a terminal device.
- z** *s1* true if the length of string *s1* is zero.
- n** *s1* true if the length of the string *s1* is nonzero.
- s1* = *s2* true if strings *s1* and *s2* are identical.
- s1* != *s2* true if strings *s1* and *s2* are *not* identical.
- s1* true if *s1* is *not* the null string.
- n1* -*eq* *n2* true if the integers *n1* and *n2* are algebraically equal. Any of the comparisons **-ne**, **-gt**, **-ge**, **-lt**, and **-le** may be used in place of **-eq**.

These primaries may be combined with the following operators:

- !** unary negation operator.
- a** binary *and* operator.
- o** binary *or* operator (**-a** has higher precedence than **-o**).
- ( *expr* ) parentheses for grouping.

Notice that all the operators and flags are separate arguments to *test*. Notice also that parentheses are meaningful to the shell. Therefore, the parentheses must be escaped.

**times** Print the accumulated user and system times for processes run from the shell.

**trap** *arg n ...*

*arg* is a command to be read and executed when the shell receives signal(s) *n*. (Note that *arg* is scanned once when the trap is set and once when the trap is taken.) Trap commands are executed in order of signal number. Any attempt to set a trap on a signal that was ignored on entry to the current shell is ineffective. An attempt to trap on signal 11 (memory fault) produces an error. If *arg* is absent, then all trap(s) *n* are reset to their original values. If *arg* is the null string, then this signal is ignored by the shell and by the commands it invokes. If *n* is 0, then the command *arg* is executed on exit from the shell. The **trap** command with no arguments prints a list of commands associated with each signal number.

**unalias** [*name ...*]

The name given is removed from the alias list.

**umask** *nnn* The user file-creation mask is set to *nnn*. If *nnn* is omitted, the current value of the mask is printed.

**wait** *n* Wait for the specified process and report its termination status. If *n* is not given, then all currently active child processes are waited for and the return code is zero.

#### Invocation

If the shell is invoked through *login* and the first character of argument zero is -, commands are initially read from */etc/profile* and then from *\$HOME/.profile* (if such files exist). Next, commands are read from the file specified by the

*SH(1)*

environment parameter **ENV** if the file exists. Thereafter, commands are read as described below, which is also the case when the shell is invoked as `/bin/sh`. The flags below are interpreted by the shell on invocation only. Note that unless the **-c** or **-s** flag is specified, the first argument is assumed to be the name of a file containing commands, and the remaining arguments are passed as positional parameters to that command file. See below:

- c** *string*     If the **-c** flag is present, then commands are read from *string*.
- s**             If the **-s** flag is present or if no arguments remain, then commands are read from the standard input. Any remaining arguments specify the positional parameters. Shell output is written to file descriptor 2.
- i**             If the **-i** flag is present or if the shell input and output are attached to a terminal, then this shell is *interactive*. In this case, **TERMINATE** is ignored (so that **kill 0** does not kill an interactive shell) and **INTERRUPT** is caught and ignored (so that **wait** is interruptible). In all cases, **QUIT** is ignored by the shell.
- r**             If the **-r** flag is present, the shell is a restricted shell.

The remaining flags and arguments are described under the **set** command above.

## Rsh Only

*Rsh* is used to set up login names and execution environments whose capabilities are more controlled than those of the standard shell. The actions of *rsh* are identical to those of *sh*, except that the following are disallowed:

- changing directory (`cd`),
- setting the value of `$PATH`,
- specifying path or command names containing `.`
- redirecting output (`>` and `>>`)

The restrictions above are enforced after **.profile** is interpreted.

When a command to be executed is found to be a shell procedure, *rsh* invokes *sh* to execute it. Thus, it is possible to provide to the end-user with shell procedures that have access to the full power of the standard shell, while imposing a limited menu of commands; this scheme assumes that the end-user does not have write and execute permissions in the same directory.

The net effect of these rules is that the writer of the **.profile** has complete control over user actions by performing guaranteed setup actions and leaving the user in an appropriate directory (probably *not* the login directory).

The system administrator often sets up a directory of commands (that is, `/usr/rbin`) that can be safely invoked by *rsh*. Some systems also provide a restricted editor *red*.

## EXIT STATUS

Errors detected by the shell, such as syntax errors, cause the shell to return a nonzero exit status. If the shell is being used noninteractively, then execution of the shell file is abandoned. Otherwise, the shell returns the exit status of the last command executed (see also the **exit** command above).

## NOTE

See the **stty(1)** command about problems entering a | character.

## FILES

/etc/profile  
\$HOME/.profile  
/tmp/sh\*  
/dev/null

## SEE ALSO

**env(1)**, **login(1)**, **stty(1)**

## BUGS

The command **readonly** (without arguments) produces the same output as the command **export**. If **<<** is used to provide standard input to an asynchronous process invoked by **&**, the shell gets mixed up about naming the input document; a garbage file **/tmp/sh\*** is created and the shell complains about not being able to find that file by another name.



NAME

`sleep` — suspend execution for an interval

SYNOPSIS

**sleep***time*

DESCRIPTION

*Sleep* suspends execution for *time* seconds. It is used to execute a command after a certain amount of time as in:

```
(sleep 105; command)&
```

or to execute a command every so often, as in:

```
while true  
do  
  command  
  sleep 37  
done
```

BUGS

*Time* must be less than 65,536 seconds.



NAME

sort — sort and/or merge files

SYNOPSIS

**sort [-cmubdfinrtx] [+pos1 [-pos2]] ... [-o output] [names]**

DESCRIPTION

*Sort* sorts lines of all the named files together and writes the result on the standard output. The name *-* means the standard input. If no input files are named, the standard input is sorted.

The default sort key is an entire line. Default ordering is lexicographic by bytes in machine collating sequence. The ordering is affected globally by the following options, one or more of which may appear.

- b** Ignore leading blanks (spaces and tabs) in field comparisons.
- d** “Dictionary” order: only letters, digits, and blanks are significant in comparisons.
- f** Fold uppercase letters onto lowercase.
- i** Ignore characters outside the ASCII range 040-0176 in nonnumeric comparisons.
- n** An initial numeric string, consisting of optional blanks, optional minus sign, and zero or more digits with optional decimal point, is sorted by arithmetic value. Option **n** implies option **b**.
- r** Reverse the sense of comparisons.
- tx** “Tab character” separating fields is *x*.

The notation *+ "pos1" - pos2* restricts a sort key to a field beginning at *pos1* and ending just before *pos2*. *Pos1* and *pos2* each have the form *m.n*, optionally followed by one or more of the flags **bdfinr**, where *m* tells a number of fields to skip from the beginning of the line and *n* tells a number of characters to skip further. If any flags are present, they override all the global ordering options for this key. If the **b** option is in effect, *n* is counted from the first nonblank in the field; **b** is attached independently to *pos2*. A missing **.n** means **.0**; a missing **-pos2** means the end of the line. Under the **-tx** option, fields are strings separated by *x*; otherwise, fields are nonempty, nonblank strings separated by blanks.

When there are multiple sort keys, later keys are compared only after all earlier keys compare equal. Lines that otherwise compare equal are ordered with all bytes significant.

These option arguments are also understood:

- c** Check that the input file is sorted according to the ordering rules; give no output unless the file is out of sort.
- m** Merge only; the input files are already sorted.
- u** Suppress all but one in each set of equal lines. Ignored bytes and bytes outside keys do not participate in this comparison.

***SORT(1)***

- o The next argument is the name of an output file to use instead of the standard output. This file may be the same as one of the inputs.

## EXAMPLES

Print in alphabetical order all the unique spellings in a list of words (capitalized words differ from uncapitalized):

```
sort -u +0f +0 list
```

Print the password file (*/etc/passwd*) sorted by user ID (the third colon-separated field):

```
sort -t: +2n /etc/passwd
```

Print the first instance of each month in an already sorted file of (month-day) entries (the options **-um** with just one input file make the choice of a unique representative from a set of equal lines predictable):

```
sort -um +0 -1 dates
```

## FILES

*/usr/tmp/stm???*

## DIAGNOSTICS

Comments and exits with nonzero status for various trouble conditions and for disorder discovered under option **-c**.

## BUGS

Very long lines are silently truncated.

NAME

`split` — split a file into pieces

SYNOPSIS

`split [-n] [file [prefix] ]`

DESCRIPTION

*Split* reads *file* and writes it in *n* -line pieces (default 1000 lines) onto a set of output files. The name of the first output file is *prefix* with **aa** appended, and so on lexicographically, up to **zz** (a maximum of 676 files). *Name* cannot be longer than 12 characters. If no output prefix is given, **x** is default.

If no input file is given, or if **-** is given in its stead, then the standard input file is used.



NAME

*stty* — set the options for a terminal

SYNOPSIS

**stty** [-g] [options]

DESCRIPTION

*Stty* sets certain terminal I/O options for the device that is the current standard input; without arguments, it reports the settings of certain options; with the **-g** option, it reports current settings in a form that can be used as an argument to another *stty* command. Note that many combinations of options make no sense, but no sanity checking is performed. The options are selected from the following:

**even ( -even )**

allow (disallow) even parity.

**odd ( -odd )** allow (disallow) odd parity.

**0** hang up phone line immediately.

**50 75 110 134 150 200 300 600 1200 1800 2400 4800 9600 exta extb**  
set terminal baud rate to the number given, if possible. (All speeds are not supported by all hardware interfaces.)

**hup ( -hup)** hang up (do not hang up) *DATAPHONE*<sup>1</sup> service connection on last close.

**nl ( -nl )** allow only new-line (allow carriage return) to end input lines.

**lcase ( -lcase )**

map (do not map) uppercase alphabets to lowercase.

**LCASE ( -LCASE )**

map (do not map) uppercase alphabets to lowercase.

**cr0 cr1 cr2 cr3**

select delay for carriage returns (larger numbers are slower).

**nl0 nl1 nl2 nl3**

select delay for line-feeds (larger numbers are slower).

**tab0 tab1** select delay for horizontal tabs (larger numbers are slower).

**bs0 bs1** select delay for backspaces (larger numbers are slower).

**ff0 ff1** select delay for form-feeds (larger numbers are slower).

**cooked or -raw ( raw )**

enable (disable) canonical input (ERASE and KILL processing).

**echo ( -echo )**

echo back (do not echo back) every character typed.

**lfkc ( -lfkc )** echo (do not echo) LF after kill character.

**erase c** set erase character to c.

**kill c** set kill line character to c.

---

1. Registered trademark of Trademark of AT&T.

**STTY(1)****tabs (-tabs or tab3)**

preserve (expand to spaces) tabs when printing.

**ek** reset ERASE and KILL characters back to normal # and @.

**sane** resets all modes to some reasonable values.

**term** set all modes suitable for the terminal type *term*, where *term* is one of **tty33**, **tty37**, **vt05**, **tn300**, **ti700**, or **tek** .

## TN83 TN983

Terminals connected to TN83, TN983, and certain other specialized asynchronous interface controllers operate in an unusual manner. Certain ones have a line speed that is not selectable. Also, the START/STOP (^S/^Q) characters have only a temporary effect. You can use ^X/^Z on these terminals to START/STOP the output.

## ECD

Some terminal ports may be configured within the ECD such that backslash (\) does not work and the pipe character (|) hangs the terminal. This can be corrected by using RCV to change form 'ttop'. Change, on option\_name 'UNIX', item 47 (escape) to xff and item 48 (escape\_char) to x5c. For this to take effect, you must remove (RMV) and restore (RST) the terminal port (TTY).

NAME

`su` — become super user or another user

SYNOPSIS

`su [- ] [ name [ arg ... ] ]`

DESCRIPTION

*Su* allows one to become another user without logging off. The default user *name* is **root** (that is, super user).

To use *su*, the appropriate password must be supplied (unless one is already super user). If the password is correct, *su* will execute a new shell with the user ID set to that of the specified user. To restore normal user ID privileges, type an **EOF** to the new shell.

Any additional arguments are passed to the shell, permitting the super user to run shell procedures with restricted privileges (an *arg* of the form **-c string** executes *string* via the shell). When additional arguments are passed, */bin/sh* is always used. When no additional arguments are passed, *su* uses the shell specified in the password file.

An initial **-** flag causes the environment to be changed to the one that would be expected if the user actually logged in again. This is done by invoking the shell with an *arg0* of **-su** causing the **.profile** in the home directory of the new user ID to be executed. Otherwise, the environment is passed along with the possible exception of **\$PATH**, which is set to **/bin:/etc:/usr/bin** for root. Note that the **.profile** can check *arg0* for **-sh** or **-su** to determine how it was invoked.

FILES

`/etc/passwd` system password file  
`$HOME/. "profile` user profile"

SEE ALSO

`env(1)`, `login(1)`, `sh(1)`



NAME

`sum` — print checksum and block count of a file

SYNOPSIS

`sum [ -r ] file`

DESCRIPTION

*Sum* calculates and prints a 16-bit checksum for the named file, and also prints the number of blocks in the file. It is typically used to look for bad spots, or to validate a file communicated over some transmission line. The option `-r` causes an alternate algorithm to be used in computing the checksum.

SEE ALSO

`wc(1)`

DIAGNOSTICS

“Read error” is indistinguishable from end of file on most devices; check the block count.



NAME

*sync* — update the super block

SYNOPSIS

***sync***

DESCRIPTION

*Sync* executes the *sync* system primitive. If the system is to be stopped, *sync* must be called to insure file system integrity. It will flush all previously unwritten system buffers out to disk; thus assuring that all file modifications up to that point will be saved.



NAME

`tail` — deliver the last part of a file

SYNOPSIS

`tail [ +- [ number ][ ibc [ f ] ] ] [ file ]`

DESCRIPTION

*Tail* copies the named file to the standard output beginning at a designated place. If no file is named, the standard input is used.

Copying begins at distance `+ number` from the beginning, or `- number` from the end of the input (if *number* is null, the value 10 is assumed). *Number* is counted in units of lines, blocks, or characters, according to the appended option **l**, **b**, or **c** . When no units are specified, counting is by lines.

With the `-f` (“follow”) option, if the input file is not a pipe, the program will not terminate after the line of the input file has been copied, but will enter an endless loop; it sleeps for a second, and then attempts to read and copy further records from the input file. Thus, it may be used to monitor the growth of a file that is being written by some other process. For example, the command:

```
tail -f fred
```

will print the last ten lines of the file **fred**, followed by any lines that are appended to **fred** between the time *tail* is initiated and killed. As another example, the command:

```
tail -15cf fred
```

will print the last 15 characters of the file **fred**, followed by any lines that are appended to **fred** between the time *tail* is initiated and killed.

SEE ALSO

`dd(1)`

BUGS

Tails relative to the end of the file are treasured up in a buffer, and thus are limited in length. Various kinds of anomalous behavior may happen with character special files.



NAME

tee — pipe fitting

SYNOPSIS

**tee** [ -i ] [ -a ] [ file ] ...

DESCRIPTION

*Tee* transcribes the standard input to the standard output and makes copies in the *files* . The **-i** option ignores interrupts; the **-a** option causes the output to be appended to the *files* rather than overwriting them.



NAME

`time` — time a command

SYNOPSIS

**time** command

DESCRIPTION

The *command* is executed; after it is complete, *time* prints the elapsed time during the command, the time spent in the system, and the time spent in execution of the command. Times are reported in seconds.

The execution time can depend on what kind of memory the program happens to land in; the user time in MOS (Module Operating System) is often half what it is in core.

The times are printed on standard error.



NAME

timer — find processes with large CPU (central processing unit) usage

SYNOPSIS

**timer** seconds

DESCRIPTION

Sometimes, a process will, due to unexpected events, start using large amounts of CPU time. When this occurs, *UNIX* system level processes within the AM (3B20D), such as recent change, may slow down to an unacceptable speed. The **timer(1)** command aids in finding these large CPU using processes.

It operates by doing two **ps(1)** commands several minutes apart. **Timer(1)** then calculates the CPU time used by each process between the two **ps(1)** commands and prints a message showing the 12 highest CPU using processes.

The user specifies the time between the **ps(1)** commands, in seconds, as the argument to the command. Valid times are from 30 to 3600 seconds.

To enter this command from the craft shell, with a 120-second interval enter, for MML:

```
EXC:ENVIR:UPROC, FN="/bin/timer", ARG=120;
```

For PDS enter:

```
EXC:ENVIR:UPROC, FN"/bin/timer", ARG"120"!
```

FILES

/bin/timer  
/bin/ps



NAME

`touch` — update access and modification times of a file

SYNOPSIS

`touch [ -amc ] [ mmddhhmm[yy] ] files`

DESCRIPTION

*Touch* causes the access and modification times of each argument to be updated. If no time is specified [see *date* (1)] the current time is used. The **-a** and **-m** options cause *touch* to update only the access or modification times, respectively (default is **-am**). The **-c** option silently prevents *touch* from creating the file, if it did not previously exist.

The return code from *touch* is the number of files for which the times could not be successfully modified (including files that did not exist and were not created).

SEE ALSO

`date(1)`



NAME

tr — translate characters

SYNOPSIS

tr [ -c ds ] [ string1 [ string2 ] ]

DESCRIPTION

*Tr* copies the standard input to the standard output with substitution or deletion of selected characters. Input characters found in *string1* are mapped into the corresponding characters of *string2*. Any combination of the options -c ds may be used:

- c       Complements the set of characters in *string1* with respect to the universe of characters whose ASCII codes are 001 through 377 octal.
- d       Deletes all input characters in *string1*.
- s       Squeezes all strings of repeated output characters that are in *string2* to single characters.

The following abbreviation conventions may be used to introduce ranges of characters or repeated characters into the strings:

- [a-z]    Stands for the string of characters whose ASCII codes run from character **a** to character **z**, inclusive.
- [a \* n ]   Stands for *n* repetitions of **a**. If the first digit of *n* is **0**, *n* is considered octal; otherwise, *n* is taken to be decimal. A zero or missing *n* is taken to be huge; this facility is useful for padding *string2*.

The escape character [ ] may be used as in the shell to remove special meaning from any character in a string. In addition, [ ] followed by 1, 2, or 3 octal digits stands for the character whose ASCII code is given by those digits.

The following example creates a list of all the words in *file1* one per line in *file2*, where a word is taken to be a maximal string of alphabets. The strings are quoted to protect the special characters from interpretation by the shell; 012 is the ASCII code for newline.

```
tr -cs "[A-Z][a-z]" "[[0]12*]" <file1 >file2
```

SEE ALSO

ed(1), sh(1).  
ascii(5) in the *UTS Programmer Reference Manual*.

BUGS

Will not handle ASCII NUL in *string1* or *string2*; always deletes NUL from input.



NAME

*true*, *false* — provide truth values

SYNOPSIS

**true**  
**false**

DESCRIPTION

*True* does nothing successfully. They are typically used in input to *sh* (1) such as:

```
while true
do
    command
done
```

*False*

does nothing unsuccessfully.

SEE ALSO

*sh*(1)

DIAGNOSTICS

*True* has exit status zero; *false* nonzero.



NAME

`tty` — get the terminal name

SYNOPSIS

`tty [ -s ]`

DESCRIPTION

*tty* prints the path name of the user's terminal. The **-s** option inhibits printing of the terminal path name, allowing one to test only the exit code.

EXIT CODES

2 if invalid options were specified,  
0 if standard input is a terminal,  
1 otherwise.

DIAGNOSTICS

“not a tty” if the standard input is not a terminal and **-s** is not specified.



NAME

**udgnm** — generic special diagnostic filename

SYNOPSIS

**udgnm** unit-name unit-number

DESCRIPTION

*Udgnm* genericizes the special file set up by *dgnnm*. It also frees the unit.

FILES

/dev/mount  
/tmp/ecdxxxxxx  
/tmp/xxx

SEE ALSO

**dgnnm(1)**

DIAGNOSTICS

Error numbers are returned and may be found in *MOVEerrcod.h*.

LIMITATIONS

The interface buffer used by *dgnnm* is passed to this program in */tmp/ectbuf*. If this is lost, there is no way to free the unit.



235-700-200  
November 1998

**COMMANDS**

***UMOUNT(1)***

NAME

umount

DESCRIPTION

*See mount.*



NAME

uname — print name of current *UNIX* system

SYNOPSIS

uname [ -snrva ]

DESCRIPTION

*Uname* prints the current system name of the *UNIX* system on the standard output file. It is mainly useful to determine what system one is using. The options cause the selected information to be printed:

- s            print the system name (default).
- n            print the nodename (the nodename may be a name that the system is known by to a communications network).
- r            print the operating system release.
- v            print the operating system version.
- a            print all the above information.

Arguments, not recognized, default the command to the -s option.



235-700-200  
November 1998

COMMANDS  
*UNCOMPRESS(1)*

NAME

uncompress

DESCRIPTION

*See compress.*



235-700-200  
November 1998

## COMMANDS

*URUN(1)*

### NAME

`urun`

### DESCRIPTION

See *run*.



NAME

vcp — volume copy

SYNOPSIS

*/etc/vcp* from to [ startingblk ] nblocks -

DESCRIPTION

*Vcp* can be used to make a contiguous copy of a file. First use *falloc(1)* to allocate the file; then use *vcp* to copy a file into the contiguous space. The *from* parameter is an existing file to be the source of the data. The *to* parameter is the previously allocated contiguous file to receive the data. *Nblocks* is the number of blocks of data to copy. If a dash (-) is given, then a dot (.) will be printed on the terminal every time 256 blocks are copied.

The *startingblk* option is not normally useful.



NAME

Vexpand — Reproduce a target file from its' compressed representation.

SYNOPSIS

**vexpand** [-t] target [-d] delta [-s source][-l log]

DESCRIPTION

*vexpand* command is used in conjunction with the *vcompress* command to represent a file (the *target* in a compressed format that can be reproduced to its' original form. The *vcompress* command produces a compressed representation (the *deltafile*) of the *targetfile* in one or two ways. First, the *deltafile* may be created using a strict compression algorithm, similar to *compress (1)*. In this method, the target file is simply compressed and no source file is specified. The second method is *data differencing*, whereby the *targetfile* is compared to a specified *sourcefile*, nothing differences.

The *vexpand* command reproduces the *targetfile* from the *deltafile*. The optional *sourcefile* is used as a base for target creation and is specified only if the *vcompress* data differencing method was used when creating the *deltafile*.

The command line options are:

- t** *target* Specifies the name of the file to be created from the compressed format contained in the *deltafile*.
- d** *delta* Specifies the name of the file containing the compressed format of the target.
- s** *source* Optionally specifies the name of a source file to be used as a base for target creation. This option is required if *data differencing* was used when creating the *delta* file.
- l** *log* Optionally specifies the name of a file where output normally intended for *stdout stderr* is to be redirected.

**Exhibit 1 —**

File *new1* is a modified version of *previous1*. To compress *new1* into file *delta* by the data differencing method using *previous1* as a base for comparison execute the following:

- `vcompress -t new1 -s previous1 -d delta`

To reproduce *new1* from the *delta* file execute the following:

- `vexpand -t new1 -s previous1 -d delta`

**VEXPAND (1)****Exhibit 2 —**

Newly created file *newcmd.sh* is to be compressed into a file *delta* without data differencing (i.e. there is no existing file that can be used as a base for comparison). Execute the following:

- `vcompress -t newcmd.sh -d delta`

To reproduce *newcmd.sh* from the *delta* file execute the following:

- `vexpand -t newcmd.sh -d delta`

## HEADER FILES

*sfio.h*, *vdelta.h*.

## SEE ALSO

`vcompress(1)`, `diff(1)`, `compress(1)`.

## DIAGNOSTICS

*vexpand* returns 0 upon successful completion and 1 upon failure.

Upon successful completion, *vexpand* outputs the checksum for the newly created target file to *stdout* (or optionally the file specified by the `-1` option) in the following format:

- Target Sum= ttttt

Where ttttt is the checksum of the complete target file (equivalent to `sum -r`).

Upon failure, *stderr* (or optionally the file specified by the `-1` option). Special attention should be paid to the following messages:

- *vexpand*: Bad Delta

A read error occurred while attempting to read the *delta* file header data (e.g. magic number, target size, or window size). The *delta* file may be corrupted.

- *vexpand*: Bad magic number

The magic number in the *delta* file header conflicts with that in *vexpand*. The *delta* may be corrupted or was created by an incompatible version of *vcompress*.

- *vexpand*: No source file.

The *delta* file was created using a *source* file and data differencing, but no *source* file was specified on the *vexpand* command line.

- *vexpand*: Source file size mismatch.

The size of the *source* file used to create the *delta* disagrees with the size of the *source* file specified on the *vexpand* command line.

- *vexpand*: internal error <errnum>.

An internal error occurred within *vexpand*. The *errnum*, along with the *delta* file and any optional *source* file should be reported to your support organization.

LIMITATIONS

None

LIBRARIES

libvdelta.a, sfio.a.



**4PRESS (1)**

NAME

**vcompress** — produce a compressed representation of a target file.

SYNOPSIS

**vcompress -t target -d delta [-s source] [-l log] [-c]**

DESCRIPTION

The **vcompress** command is used in conjunction with the **vexpand (1)** command to represent a file (the *target*) in a compressed format that can be reproduced to its' original form. The **vcompress** command produces a compressed representation (the *delta* file) of the *target* file in one of two ways. First, the *delta* file may be created using a strict compression algorithm, similar to **compress (1)**. In this mode, *compression only*, the target file is simply compressed and no source file is specified. The second mode is *data differencing*, whereby the *target* file is compared to a specified *source* file, nothing differences.

The **vexpand (1)** command reproduces the *target* file from the *delta* file, and optionally the *source* file, if the method of compression by **vcompress** was data differencing.

The command line options are:

- t target** Specifies the name of the file to be represented in compressed format.
- d delta** Specifies the name of the file to be generated, which will contain the compressed format of the target.
- s source** Optionally specifies the name of a source file to be used for data differencing.
- l log** Optionally specifies the name of a file where output normally intended for *stdout* and *stderr* is to be redirected.
- c** Optionally specifies to perform special coff file header processing. This processing writes the entire header portion of a coff formatted file to the *delta* file. By doing this, **vexpand (1)** will reproduce the header portion of the file by copying it from the *delta* file, instead of using data differencing and the *source* file. This is necessary in the software update, SU, environment where the headers in the SU viewpath are not necessarily byte for byte comparable with the headers on a 5ESS<sup>®</sup>-2000 switch. Specifying this option for files that are not in coff format or for files being compressed using the *compression only* mode has no effect.

**4PRESS (1)****Exhibit 1 —**

File *new1* is a modified version of *previous1*. To compress *new1* into file *delta* by the data differencing method using *previous1* as a base for comparison execute the following:

- `vcompress -t new1 -s previous1 -d delta`

To reproduce *new1* from the *delta* file execute the following:

- `vexpand -t new1 -s previous1 -d delta`

**Exhibit 2 —**

Newly created file *newcmd.sh* is to be compressed into a file *delta* without data differencing (i.e. there is no existing file that can be used as a base for comparison). Execute the following:

- `vcompress -t newcmd.sh -d delta`

To reproduce *newcmd.sh* from the *delta* file execute the following:

- `vexpand -t newcmd.sh -d delta`

## HEADER FILES

`sfio.h`, `vdelta.h`.

## SEE ALSO

`vexpand(1)`, `diff(1)`, `compress(1)`.

## DIAGNOSTICS

**vcompress** returns 0 upon successful completion and 1 upon failure.

Upon successful completion, **vcompress** outputs the checksum for the *target* file and optional *source* file to *stdout* (or optionally the file specified by the `-l` option) in the following format:

- Target Sum=tttt Source Sum= sssss

where tttt and sssss are the complete checksums of the corresponding file (equivalent to `sum -r`).

Upon failure **vcompress** outputs error messages to *stderr* (or optionally the file specified by the `-l` option). Special attention should be paid to the following message:

- **vcompress:internal error <errnum>**

An internal error occurred within **vcompress**. The *errnum*, along with the *target* file and any optional *source* file should be reported to your support organization.

235-700-200  
November 1998

COMMANDS  
*4PRESS (1)*

LIBRARIES

libvdelta.a, sfio.a.



## NAME

*vi* — screen-oriented (visual) display editor based on *ex*

## SYNOPSIS

*vi* [-r file] p-L] [-wn] [-R] [+command] [-c command] file ...

## DESCRIPTION

*vi* (visual) is a display-oriented text editor based on an underlying line editor *ex* (1). It is possible to use the command mode of *ex* from within *vi*.

When using *vi*, changes you make to the file are reflected in what you see on your terminal screen. The position of the cursor on the screen indicates the position within the file.

## INVOCATION

The following invocation options are interpreted by *vi*:

- r file** Recovers *file* after an editor or system crash. If *file* is not specified, a list of all saved files is printed.
- L** Lists the name of all files saved as the result of an editor or system crash.
- w n** Sets the default window size to *n*. This is useful when using the editor over a slow speed line.
- R** Read only mode; the **readonly** flag is set, preventing accidental overwriting of the file.

**[-c|+] command**

The specified *ex* command is interpreted before editing begins.

The *file* argument indicates files to be edited.

If you invoke *vi* on a file and edit it, then discover the file is read-only when you try to save the changes, you can save the changes made to the file by doing the following:

Enter:

```
:!chmod u+rw filename<CR>
```

After the command executes, and you press the RETURN key to continue editing the file, you can save the changes made to the buffer by entering:

```
:w!<CR>
```

Now, you can add more changes to the file or quit the edit session. The changes will be saved.

## VI MODES

**Command** Normal and initial mode. Other modes return to command mode upon completion. ESC (escape) is used to cancel a partial command.

**Input** Entered by the following options **a i A I o O c C s S R**. Arbitrary

*VI(1)*

text may then be entered. Input mode is normally terminated with the ESC character, or abnormally with interrupt.

Last Line Reading input for : / ? or !; terminate with CR to execute, interrupt to cancel.

## COMMAND MODES

## Sample Commands

← ↓ ↑ → arrow keys move the cursor  
 h j k l same as arrow keys  
 i*text*ESC insert text *abc*  
 cw*new*ESC change word to *new*  
 easESC pluralize word  
 x delete a character  
 dw delete a word  
 dd delete a line  
 3dd ... 3 lines  
 u undo previous change  
 ZZ exit vi, saving changes  
 :q!CR quit, discarding changes  
 /*text*CR search for *text*  
 ^U ^D scroll up or down  
 :*ex cmd*CR any ex or ed command

## Counts Before vi Commands

Numbers may be typed as a prefix to some commands. They are interpreted in one of the following ways:

line/column number	z G
scroll amount	^D ^U
repeat effect	most of the rest

## Interrupting, Canceling

ESC end insert or incomplete cmd  
 DEL (delete or rubout) interrupts  
 ^L reprint screen if DEL scrambles it  
 ^R reprint screen if ^L is → key

## File Manipulation

**ZZ** if file modified, write and exit; otherwise, exit  
**:wCR** write back changes  
**:w!CR** forced write, if permission originally not valid  
**:qCR** quit  
**:q!CR** quit, discard changes  
**:e *name*CR** edit file *name*  
**:e!CR** reedit, discard changes  
**:e + *name*CR** edit, starting at end  
**:e +*n*CR** edit starting at line *n*  
**:e #CR** edit alternate file  
**:e! #CR** edit alternate file, discard changes synonym  
for **:e #**  
**:w *name*CR** write file *name*  
**:w! *name*CR** overwrite file *name*  
**:shCR** run shell, then return  
**:!*cmd*CR** run *cmd*, then return  
**:nCR** edit next file in arglist  
**:n *args*CR** specify new arglist  
**^G** show current file and line

#### Positioning Within File

**VI(1)**

**^F** forward screen  
**^B** backward screen  
**^D** scroll down half screen  
**^U** scroll up half screen  
**nG** go to the beginning of the specified line  
 (end default), where *n* is a line number  
**G** go to specified line (end default)  
*lpat* next line matching *pat*  
*?pat* previous line matching *pat*  
**n** repeat last / or ?  
**N** reverse last / or ?  
*lpat+ n* nth line after *pat*  
*?pat?- n* nth line before *pat*  
**]]** next section/function  
**[[** previous section/function  
**(** beginning of sentence  
**)** end of sentence  
**{** beginning of paragraph  
**}** end of paragraph  
**%** find matching ( ) { or }

Adjusting the Screen

**^L** clear and redraw  
**^R** retype, eliminate @ lines  
**zCR** redraw, current at window top  
**z-** ... at bottom  
**z.** ... at center  
*lpatz-* pat line at bottom  
*lpatz.* ... at center  
**zn.** use *n* line window  
**^E** scroll window down 1 line  
**^Y** scroll window up 1 line

#### Marking and Returning

**‘;’** move cursor to previous context  
**’’** ... at first non-white in line  
**m $x$**  mark current position with letter *x*  
**‘; $x$ ’** move cursor to mark *x*  
**‘ $x$ ’** ... at first non-white in line

#### Line Positioning

*VI(1)*

**H** top line on screen  
**L** last line on screen  
**M** middle line on screen  
**+** next line, at first non-white  
**-** previous line, at first non-white  
**CR** return, same as +  
**↓ or j** next line, same column  
**↑ or k** previous line, same column

Character Positioning

**^** first non white  
**0** beginning of line  
**\$** end-of-line  
**h or ←** backwards  
**l or →** forward  
**^H** same as ←  
*space* same as →  
**f*x*** find *x* forward  
**F*x*** f backward  
**t*x*** upto *x* forward  
**T*x*** back upto *x*  
**;** repeat last f F t or T  
**,** repeat inverse of last f F t or T  
***n*|** move to column *n*  
**%** find matching ( { ) or }

Words, Sentences, Paragraphs

w word forward  
b back word  
e end of word  
) to next sentence  
} to next paragraph  
( back sentence  
{ back paragraph  
W blank delimited word  
B back W  
E to end of W

#### Corrections During Insert

^H erase last character  
^W erase last word  
*erase* your erase, same as ^H  
*kill* your kill, erase input this line  
\ quotes ^H, your erase and kill  
ESC ends insertion, back to command  
DEL interrupt, terminates insert  
^D backtab over *autoindent*  
0^D backtab to beginning of line;  
reset left margin of *autoindent*  
^^D caret (^) followed by control-d (^D); backtab  
to beginning of line; do not reset left margin  
of *autoindent*  
^V quote non-printing character  
↑^D kill *autoindent*, save for next

#### Insert and Replace

a append after cursor  
i insert before cursor  
A append at end-of-line  
I insert before first non-blank  
o open line below  
O open above  
r*x* replace single char with *x*  
R*text* ESC replace characters

#### Operators

Operators are followed by a cursor motion, and affect all text that would have been moved over. For example, since **w** moves over a word, **dw** deletes the word that would be moved over. Double the operator, i.e., **dd** to affect whole lines.

*VI(1)*

**d** delete  
**c** change  
**y** yank lines to buffer  
**<** left shift  
**>** right shift  
**!** filter through command  
**=** indent for -2LISP

## Miscellaneous Operations

**C** change rest of line (c\$)  
**D** delete rest of line (d\$)  
**s** substitute characters (cl)  
**S** substitute lines (cc)  
**J** join lines  
**x** delete characters (dl)  
**X** ... before cursor (dh)  
**Y** yank lines (yy)

## Yank and Put

Put inserts the text most recently deleted or yanked. However, if a buffer is named, the text in that buffer is put instead.

**3yy** yank 3 lines  
**3yl** yank 3 characters  
**p** put back text after cursor  
**P** put before cursor  
**"xp** put from buffer *x*  
**"xy** yank to buffer *x*  
**"xd** delete into buffer *x*

## Undo, Redo, Retrieve

**u** undo last change  
**U** restore current line  
**.** repeat last change  
**"dp** retrieve *d*'th last delete

## AUTHOR

*vi* was developed by The University of California, Berkeley California, Computer Science Division, Department of Electrical Engineering and Computer Science.

## FILES

**/tmp/Ex?????**  
 temporary file used by *vi* during an edit session.

/tmp/Ex????

recoverable edited file preserved as a result of a killed edit session.

/usr/lib/exstrings

data file containing *vi* output messages.

#### WARNING

The **-x** and **-C** encryption options are not available under RTR, because the **crypt(1)** command is not supported.

Only **vt100** family of terminals is supported.

#### CAVEATS

Software tabs using **^T** works only immediately after the *autoindent*.

Left and right shifts on intelligent terminals do not make use of insert and delete character operations in the terminal.

If you edit a file and try to write the file when all it contains is a newline, the editor removes the newline and the file contains zero characters.

If the file you are editing has no characters, i.e. zero length, and you exit using **ZZ**, the file will not be written. Exiting using **:wq** does create a zero length file.



NAME

*wc* — word count

SYNOPSIS

**wc** [ **-lwc** ] [ names ]

DESCRIPTION

*Wc* counts lines, words, and characters in the named files or in the standard input if no *names* appear. It also keeps a total count for all named files. A word is a maximal string of characters delimited by spaces, tabs, or new lines.

The options **l**, **w**, and **c** may be used in any combination to specify that a subset of lines, words, and characters are to be reported. The default is **-lwc**.

When *names* are specified on the command line, they will be printed along with the counts.



NAME

*who* — who is on the system

SYNOPSIS

**who** [ *file* ]  
**who am i**

DESCRIPTION

*Who* can list the user's name, terminal line, login time, elapsed time since activity occurred on the line, and the process-ID of the command interpreter (shell) for each current *UNIX* system user. It examines the **/etc/utmp** file to obtain its information. If *file* is given, that file is examined. Usually, *file* will be **/etc/wtmp**, which contains a history of all the logins since the file was last created.

*Who* with the **am i** option identifies the invoking user.

FILES

**/etc/utmp**  
**/etc/wtmp**  
**/etc/inittab**

SEE ALSO

*date(1)*, *login(1)*, *mesg(1)*, *su(1)*



NAME

write — write to another user

SYNOPSIS

**write** user [ line ]

DESCRIPTION

*Write* copies lines from your terminal to that of another user. When first called, it sends the message:

**Message from** *your login (tty ??)* [ *date* ] ...

to the person you want to talk to. When it has successfully completed the connection, it also sends two bells to your own terminal to indicate that what you are typing is being sent.

The recipient of the message should write back at this point. Communication continues until an end of file is read from the terminal or an interrupt is sent. At that point, *write* writes **EOT** on the other terminal and exits.

If you want to write to a user who is logged in more than once, the *line* argument may be used to indicate which line or terminal to send to (for example, **tty00**); otherwise, the first instance of the user found in */etc/utmp* is assumed and the following message posted:

*user* is logged on more than one place.

You are connected to "*terminal*".

Other locations are:

*terminal*

Permission to write may be denied or granted by use of the *mesg(1)* command. Writing to others is normally allowed by default. Certain commands, in particular *pr(1)* disallows messages in order to prevent interference with their output. However, if the user has super user permissions, messages can be forced onto a write inhibited terminal.

If the character **!** is found at the beginning of a line, *write* calls the shell to execute the rest of the line as a command.

The following protocol is suggested for using *write*: when you first *write* to another user, wait for them to *write* back before starting to send. Each person should end a message with a distinctive signal [that is, **(o)** for "over"] so that the other person knows when to reply. The signal **(oo)** (for "over and out") is suggested when conversation is to be terminated.

FILES

*/etc/utmp* to find user  
*/bin/sh* to execute **!**

## COMMANDS

235-700-200  
November 1998

### *WRITE(1)*

#### SEE ALSO

mail(1), mesg(1), pr(1), sh(1), who(1)

#### DIAGNOSTICS

*“user not logged in”* if the person you are trying to *write* to is not logged in.

NAME

zcat

DESCRIPTION

*See compress.*



## GLOSSARY

This section provides acronyms, terms, and abbreviations used in this manual.

### GLOSSARY

**ACCED** — Access Editor

**Administration** — Administration is a number of related functions with the objective of ensuring the overall provision of service by the *5ESS*<sup>®</sup>-2000 switch. Administration includes the assignment of lines and trunks to the system, memory management, collection of traffic and plant data, provisions for additions and modifications to the switch, service evaluation, and capabilities to control and manage the *5ESS*-2000 switch. The primary objective of administration is to assure that the *5ESS*-2000 switch delivers a high level of quality service to the subscribing customers. This is accomplished by monitoring and evaluating system performance. Potential problems that could cause service deterioration are identified.

**AM** — Administrative Module

The AM is that part of the *5ESS*-2000 switch which performs the part of call processing, administration, and maintenance which cannot be economically distributed to the switching modules. The AM consists of the processor, disk storage, and tape backup units. The AM processor performs the centralized processing functions, high-speed tape, and controls the flow of data between the other dedicated processors distributed throughout the remaining units. The processor functions are fully duplicated (except for the port switch) in order to assure continued processing capability.

**ANSI** — American National Standards Institute

**AP** — Applications Processor

**ARS** — Automatic Route Selection

**ASCII** — American Standard Code for Information Interchange

**AT** — Access Tandem

**BOC** — Bell Operating Company

**BORSCHT** — Battery Feed, Overvoltage Protection, Ringing, Etc. Functions

**BPS** — Bits Per Second

A measure of the speed with which data communications can move over a line. Also abbreviated as bps, B/S, and b/s.

**BRI** — Basic Rate Interface

**BST** — Basic Services Terminal

**CC** — Control Console

**CE** — Control Equipment

**CI** — Critical Indicator

**CCITT** — International Telephone and Telegraph Consultative Committee

**CIU** — Craft Interface Unit

**CLEI** — Common Language Equipment Identification

**CM** — Communication Module

**CNI** — Common Network Interface

**CO** — Central Office

**COER** — Central Office Equipment Reports

**COT** — Central Office Terminal

**CPE** — Customer Premises Equipment

**CRT** — Cathode Ray Tube

**DA** — Discontinued Availability

**DAS** — Data Auxiliary Set

**DFC** — Disk File Controller

**DFI** — Digital Facility Interface

**DMA** — Direct Memory Access

**DN** — Directory Number

**DP** — Dial Pulse

**DSU** — Digital Service Unit

**DTMF** — Dual Tone Multifrequency

**EADAS** — Engineering and Administrative Data Acquisition System

A mechanized system that is part of Lucent supplied OSS and is a near real-time data collection and surveillance system. Data transmitted over the data link interface of the *5ESS-2000* switch is collected and summarized by a data processor system.

**EADAS/NM** — EADAS/Network Management

**EAI** — Emergency Action Interface

**ECD** — Equipment Configuration Data Base

**EIA** — Electronics Industry Association

**EMACS** — EMACS is a screen editor that can be used to create or to edit files using a display terminal.

**EOC** — Embedded Operations Channel

**ESF** — Extended Super Frame

**FIT** — Fully-Initializing Terminal

**Hz** — Hertz/Cycles per Second

**IC** — Interexchange Carrier

**IMR** — Initial Modification Request

**I/O** — Input/Output

**IOP** — Input Output Processor

**IPM** — Interruptions Per Minute

**ISDN** — Integrated Services Digital Network

**LED** — Light-Emitting Diode

**LTD** — Local Test Desk

**MCC** — Master Control Center

Provides craft interface of a *5ESS-2000* switch office. Consists of a video display terminal with keyboard, ROP, and a key telephone set. The TLWS is also part of the MCC.

**MFOS** — Multifunctional Operations System

**MFT** — Metallic Facility Terminal

**MIM** — Management Information Messages

**MML** — Man-Machine Language

**MMSU** — Modular Metallic Service Unit

**MTTY** — Maintenance Teletypewriter

**NCLK** — Network Clock

**NCT** — Network Control and Timing

**NDL** — New Data Link

**NIT** — Non-Initializing Terminal

**NM** — Network Management

Network Management is a set of real-time procedures aimed at optimizing network performance when the network is under stress due to adverse conditions. The NM provides and operates control and surveillance features that aid in maintaining the network integrity and security during overloads and failures.

**OA&M** — Operations, Administration, and Maintenance

**OAM&P** — Operations, Administration, Maintenance, and Provisioning

**ODA** — Office Data Administration

The mechanism by which initial translation information may be assembled for a *5ESS-2000* switch. Information from the *5ESS-2000* switch is entered via a video terminal and transferred into an ODA computer, assembled, then sent to the *5ESS-2000* switch.

**ODBE** — Office Data Base Editor

**ODD** — Office Dependent Data

**ORP** — Office Records Printer

**OS** — Operations Systems

**OSC** — Operator Service Center

**OSPS** — Operator Services Position System

**OSS** — Operations Support System

Service that provides the routine operations needed to maintain and engineer the telephone network.

**OTC** — Operating Telephone Company

A service organization using telephone equipment to provide communications for its customers.

**PBX** — Private Branch Exchange

Provides business customers service that allows a group of lines to make intragroup calls on an extension dialed basis plus allowing direct inward dialed calls from regular telephone network.

**PC** — Personal Computer**PC** — Peripheral Controller**PICB** — Peripheral Interface Control Bus**PIDB** — Peripheral Interface Data Bus**PIN** — Personal Identification Number

A number used to identify the customer using the ABC service.

**PODS** — Plain Old Digital Service**POTS** — Plain Old Telephone Service**PPS** — Pulses Per Second**PROM** — Programmable Read Only Memory**RB** — Ringback**RBOC** — Regional Bell Operating Companies**ROP** — Receive-Only Printer**RT** — Remote Terminal**RTR** — Real Time Reliable**SCANS** — Software Change Administration and Notification System**SCC** — Switching Control Center**SCCS** — Switching Control Center System

A centralized system that controls the switching operations of many switches.

**SCSI** — Small Computer System Interface**SD** — Schematic Drawing**SM** — Switching Module

Equipment consisting of the module controller unit, time slot interchange unit, local digital service unit, and analog and digital interface units. The SM performs 95 percent of all switching performed in the 5ESS-2000 switch. When an SM is remotely located, it is referred to as an RSM (Remote Switching Module). When an SM is used to support an RSM, it is referred to as an HSM (host SM).

**SPCS** — Stored Program Control System

**STLWS** — Supplementary Trunk and Line Work Station

**TELCO** — Telephone Company

**TLWS** — Trunk and Line Work Station

**TMT** — Transmission Maintenance Terminal

**TOD** — Time Of Day

**TOPAS** — Trunk, Operations, Provisioning, and Administrative System

**TSI** — Time Slot Interchange

**TSPS** — Traffic Service Position System

A type of Traffic Service System having stored program control that provides for the processing and recording of special calls requiring operator assistance.

**TT** — Touch-Tone

**TTF** — Transmission Test Facility

**TTY** — Teletypewriter

**TTYC** — Teletypewriter Controller

**TU** — Trunk Unit

**VDT** — Video Display Terminal

**VF** — Voice Frequency



## NUMERICAL

3B20D computer, 1-1, 1-4

### A

ACCED, 3-4

Access, 4-7, 4-9, 4-26

Access editor, 3-4

Access files, 4-7

Actual login, 3-2

Adding a second password, 3-2

Adding new logins, 3-1

Administration, 1-1, 1-2, 3-1

Advanced editing, 4-4, 4-10

Advanced editing commands, 4-4, 4-12

Advanced replace commands, 4-13

Advanced search commands, 4-13

Argument, 1-3, 4-4, 4-5, 4-6, 4-8, 4-12, 4-14, 4-17, 4-18, 4-19, 4-25, 4-26

Argument prototype, 1-3

Arguments, 2-2, 4-3

Arguments - command line, 4-26

ASCII, 4-1, 4-2, 4-12, 4-14, 4-24

ASCII 0, 4-1

ASCII 0177, 4-1

ASCII 033, 4-1

ASCII 034, 4-1

ASCII 035, 4-2

ASCII 036, 4-2

ASCII 037, 4-2

ASCII characters, 4-1

Assistance, technical, 1-4

### B

Basic action, 4-8

Basic commands, 4-23, 4-26

Basic concepts, 4-1

Basic EMACS commands, 4-3

Baud rate, 3-3

Break, 2-2

Buffer display, 4-2, 4-12, 4-13

Buffer name, 4-7, 4-9, 4-20

Buffer number, 4-2, 4-9

Buffer text, 4-2

Bugs, 1-3

### C

C program, 4-1, 4-17, 4-19, 4-22

Cable, 3-1

Carriage-return, 2-2

- Character sequence, 4-13, 4-14, 4-18
- Character set, 4-1
- Characters, 2-1, 2-2
- Ciopt form, 3-3
- Command language procedures, 1-2
- Command line, 4-16, 4-18, 4-25, 4-26
- Command line arguments, 4-26
- Command modes, 4-22
- Command name, 2-2
- Command structure, 4-3
- Commands, 1-1, 1-2, 1-3, 1-4, 1-5, 2-1, 3-2
- Commands related to windows, 4-12
- Commands that escape to the *UNIX* RTR operating system, 4-15
- Computer Access Restriction, 3-4
- Conclusions, 4-27
- Configuration, 3-1
- Constructing regular expressions, 4-13
- Control character, 4-2, 4-4, 4-10, 4-11, 4-12, 4-13, 4-18, 4-23, 4-24
- Control character commands, 4-4
- Control characters, 4-5
- Control-D, 3-2, 3-4
- Control-S, 2-2
- Conventional login, 3-1
- Conventions, 1-3
- Craft shell, 3-2
- Cron, 1-5
- CRT, 2-2
- Current directory, 2-3
- Cursor movement commands, 4-5

## D

- Data base, 3-3
- Data base changes for a dial-up recent change terminal, 3-3
- Data base changes for a dial-up STLWS, 3-3
- DC3, 2-2
- Daemon password, 3-2
- Default argument, 4-3, 4-8
- Default values, 4-18, 4-19
- Defined buffers, 4-9
- Defining macros, 4-18
- Deleting text, 4-1
- Deletion commands, 4-6
- Description, 1-2
- Description of *UNIX* RTR operating system availability feature, 1-4
- Diagnostics, 1-2
- Dial-up facilities, 3-1, 3-3
- Dial-up recent change terminal, 3-4
- Dial-up terminals, 3-2
- Dial-up text recent change terminal, 3-1, 3-3
- Dial-up users, 3-2
- Display, 3-3, 4-2

Display buffer, 4-2  
Display modes, 4-19  
Display region, 3-3, 3-4  
Display screen, 4-2, 4-19  
Display terminal, 4-1, 4-27  
Displaying parameters, 4-2  
Displays, 4-1, 4-5, 4-12, 4-17, 4-18, 4-23  
Duplex switch, 2-1

## E

ECD, 3-1, 3-3, 3-4  
Ed editor, 4-1, 4-12, 4-13  
Editing commands, 4-1, 4-3, 4-4  
Editor command, 4-1  
Editor cursor, 4-2  
Editor limitations, 4-27  
EMACS, 4-2, 4-5, 4-8, 4-9, 4-12, 4-13, 4-14, 4-15, 4-16, 4-17, 4-18, 4-19, 4-20, 4-22, 4-23, 4-24, 4-25, 4-26, 4-27  
EMACS commands, 4-1, 4-3, 4-14  
EMACS description, 1-1, 1-2, 4-1  
EMACS editor, 4-27  
EMACS/macros, 4-8, 4-15, 4-22  
EMACS/macros/CATALOG, 4-15  
Empty buffer name, 4-9  
Entering commands, 4-3, 4-24  
Entering parameters, 4-3, 4-4  
Entries, 1-2, 1-3  
Entry, 1-2  
Equipment configuration data, 3-1  
Equivalent printable character, 4-1  
Error message, 4-5, 4-8, 4-9, 4-16, 4-20, 4-22, 4-23  
Escape character, 4-1, 4-2, 4-12  
Escape commands to the *UNIX* RTR operating system, 4-15  
Escape key, 4-4  
Escape sequences, 4-1  
/etc/passwd, 3-1  
Examples, 1-2  
Exit, 3-2  
External security code, 3-1

## F

File accessing commands, 4-8  
File name, 4-2, 4-7, 4-8, 4-9, 4-15, 4-16, 4-24, 4-26  
Files, 1-2  
Files - organization, 1-5  
Formal training, 1-1  
Full-duplex, 2-1  
Full-duplex input/output terminal, 2-1  
Function keys, 4-4

**G**

Getting help, 4-4  
Getting out of trouble, 4-4  
Getting started, 1-1, 1-2, 2-1, 4-25  
Global replacements, 4-10  
Glossary, 1-1, 5.124-1

**H**

Hardware, 3-1, 3-3  
How to communicate through your terminal, 2-1  
How to run a program, 2-2

**I**

Incremental search, 4-11, 4-12, 4-22  
Indentation, 4-22  
Independent entries, 1-2  
Initialization file, 4-25, 4-26  
Initializations, 4-14  
*Initialize the teletypewriter controller*, 3-4  
Input files, 4-14  
Input line, 2-1, 2-2  
Input/output processor, 3-1  
Inserting control characters, 4-12  
Inserting odd characters, 4-12  
Inserting text, 4-3, 4-17  
Interface modes, 4-20  
Internal trouble, 4-27  
Interrupt, 2-2  
Introduction, 1-1, 1-2, 3-1, 4-1  
IOP, 3-1, 3-3  
IOP cable, 3-1

**K**

Keyboard character bindings, 4-18  
Keyboard macro, 4-14, 4-15, 4-18, 4-25  
Keyboard macros, 4-14  
Keyboards, 4-6  
Kill, 2-1

**L**

Limitations of the editor, 4-27  
Literals, 1-3  
Lnumb, 4-2, 4-19  
Logging in, 2-1  
Logging off, 3-4  
Login, 3-1, 3-4  
Login capability, 3-1

Login capability features, 3-1  
Login code, 3-1  
Login name, 2-1, 3-2  
Login process, 3-1, 3-3, 3-4  
Login sequence, 2-1  
Login software, 3-1  
Login name, 3-2  
Logins, 1-1, 1-5  
Logout procedure, 2-1  
Lucent Technologies Customer Information Center, 1-4

## M

Macro command, 4-18  
Macro name, 4-15, 4-18  
Macro programming facility, 4-14  
Macros, 4-14  
Man, 1-5  
Master control center, 3-1  
MCC, 3-1  
Meta character, 4-2, 4-3, 4-4  
Meta character commands, 4-4  
Miscellaneous commands, 4-10, 4-16  
Mode name, 4-25  
Modes, 4-19  
Modes - command, 4-22  
Modes - display', 4-19  
Modes - interface, 4-20  
Modes - terminal, 4-24  
More advanced commands, 4-3  
Moving commands, 4-6

## N

Name, 1-2  
Network access password, 3-1  
New logins, 3-1, 3-2  
Newline character, 4-2, 4-21  
Normal shell conventions, 4-8  
Numeric argument, 4-3

## O

ODBE, 3-4  
Office data base editor, 3-4  
ON/OFF modes, 4-19  
Optional second password, 3-1, 3-2  
Options, 4-5, 4-12, 4-26  
Ordinary printing characters, 4-3  
Organization, 1-1  
Organization of files, 1-5

**P**

Parameters, 4-2, 4-3, 4-4, 4-18, 4-19  
Passwd daemon, 3-2  
Password protected commands, 3-4, 3-6  
Password security, 3-1  
Pathnames, 2-3  
PDS12, 3-3  
PDS12C, 3-3  
Picture images, 4-23  
Picture mode, 4-20, 4-23  
Previous argument prototype, 1-3  
Printable ASCII text, 4-14  
Printing characters, 4-1, 4-3  
Profile, 2-1  
Program test, 4-2  
Prompt, 4-2, 4-3, 4-4, 4-7, 4-8, 4-9, 4-11, 4-13, 4-14, 4-16, 4-17, 4-18, 4-19, 4-20  
Purpose, 1-1

**Q**

Query, 4-11, 4-12, 4-14, 4-18, 4-21

**R**

Recent change terminal, 3-1, 3-3, 3-4  
Recovering from various problems, 4-27  
Regional Technical Assistance Center, 1-4  
Regular expression search, 4-14  
Regular expressions, 4-13  
Replace, 4-10  
Replace command, 4-5, 4-27  
Replacing text, 4-1  
ROP, 1-5  
Rubout, 2-2  
Running EMACS, 4-27

**S**

Screen editor, 4-1  
Screen width, 4-2  
Scroll region, 3-3  
Search, 4-10  
Search command, 4-14  
Search commands - advanced, 4-13  
Search string, 4-10, 4-11, 4-14, 4-21  
Second password, 3-1, 3-2  
Secure dial-up facility, 3-3  
Security, 3-2  
See also, 1-2, 1-3  
Shell, 1-2, 2-1, 2-2  
Simple cursor movement commands, 4-5

Simple file and buffer commands, 4-7  
Simple moving commands, 4-6  
Simple Search and Replace, 4-10  
Simple text deleting commands, 4-6  
Software, 3-1  
Special commands, 4-7  
Specifying default modes for a file, 4-25  
Status information, 4-2, 4-17  
Status line, 4-2, 4-9, 4-12  
STLWS, 3-1  
STLWSCG, 3-3  
stty command, 2-1, 2-2  
Subroutines, 1-2  
Substitute command, 4-12  
Supplementary trunk line work station, 3-1  
Surprises, 2-3  
Control-Q, 2-2  
DC1, 2-2  
Synopsis, 1-2, 1-3  
System source programs, 2-2

## T

Tab characters, 2-2  
Technical Assistance, 1-4  
Teletypewriter controller, 3-3, 3-4  
Terminal, 2-1  
Terminal bell, 4-5, 4-24  
Terminal cursor, 4-2  
Terminal modes, 4-24  
Terminal type, 4-19, 4-25  
Text deleting commands, 4-6  
Text editing, 4-27  
Text in the buffer, 4-2  
Text insertion, 4-3  
TTYC, 3-4  
Two dimensional editing, 4-23  
Typical screen, 4-2

## U

Undo, 4-4, 4-27  
/unixa, 1-4, 1-5  
/unixabf, 1-4  
/unixa/tmp, 1-5  
/unixa/users, 1-4, 1-5  
/unixa/users/manager, 1-5  
UNIX RTR, 3-1  
UNIX RTR operating system, 1-1, 1-2, 1-4, 1-5, 2-1, 2-2, 2-3, 3-1, 3-2, 3-4, 4-1, 4-5, 4-9,  
4-10, 4-15, 4-16, 4-20, 4-21, 4-23, 4-25, 4-27  
UNIX RTR operating system availability, 1-2, 1-4  
UNIX RTR operating system availability feature, 1-4

*UNIX* RTR operating system commands, 3-5  
*UNIX* RTR operating system escape commands, 4-15  
*UNIX* RTR operating system shell, 3-4  
*UNIX* RTR system, 3-1  
*UNIX* system, 3-1  
*UNIX* system quick reference guide - 307-129, 2-1  
*UNIX* system users' handbook - 320-042, 2-1  
*UNIX* system users' handbook - 320-042, 2-3  
Unusual non-printing characters, 4-1  
Update information, 1-1  
Updating, 3-1  
User Feedback, 1-3  
User handbook, 2-1, 2-3  
*Using the login*, 3-4

**V**

Variable, 3-2  
Various editing functions, 4-3  
Various modes, 4-25  
Visual display appears, 4-1  
VT100DAP, 3-3

**W**

Warning error message, 4-5  
Warnings, 1-2  
Window commands, 4-12