

5 Reference Guide to Streams, Files, and I/O

symbolics



5 Reference Guide to Streams, Files, and I/O

symbolics

Reference Guide to Streams, Files, and I/O

996055

March 1985

This document corresponds to Release 6.0 and later releases.

The software, data, and information contained herein are proprietary to, and comprise valuable trade secrets of, Symbolics, Inc. They are given in confidence by Symbolics pursuant to a written license agreement, and may be used, copied, transmitted, and stored only in accordance with the terms of such license.

This document may not be reproduced in whole or in part without the prior written consent of Symbolics, Inc.

Copyright © 1985, 1984, 1983, 1982, 1981, 1980 Symbolics, Inc. All Rights Reserved.
Font Library Copyright © 1984 Bitstream Inc. All Rights Reserved.

Symbolics, Symbolics 3600, Symbolics 3670, Symbolics 3640, SYMBOLICS-LISP, ZETALISP, MACSYMA, S-GEOMETRY, S-PAINT, and S-RENDER are trademarks of Symbolics, Inc.

MS-DOS is a trademark of Microsoft, Inc. UNIX is a trademark of Bell Laboratories, Inc. VAX, VMS, and TOPS-20 are trademarks of Digital Equipment Corporation. TENEX is a registered trademark of Bolt, Beranek and Newman, Inc.

Restricted Rights Legend

Use, duplication, or disclosure by the government is subject to restrictions as set forth in subdivision (b)(3)(ii) of the Rights in Technical Data and Computer Software Clause at FAR 52.227-7013.

Text written and produced on Symbolics 3600-family computers by the Documentation Group of Symbolics, Inc.

Text typography: Century Schoolbook and Helvetica produced on Symbolics 3600-family computers from Bitstream, Inc., outlines; text masters printed on Symbolics LGP-1 Laser Graphics Printers.

Cover design: Schafer/LaCasse

Cover printer: W.E. Andrews Co., Inc.

Text printer: ZBR Publications, Inc.

Printed in the USA.

Printing year and number: 87 86 85 9 8 7 6 5 4 3 2 1

Table of Contents

	Page
I. Streams	1
1. Introduction to the I/O System	3
2. The Character Set	5
2.1 Effect of Character Set Translation on Direct Access File Streams	10
3. I/O Streams	11
3.1 Introduction to Streams	11
3.2 General-purpose Stream Operations	11
3.3 Special-purpose Stream Operations	19
3.4 Standard Streams	21
3.5 Making Your Own Stream	23
3.6 Coroutine Streams	25
4. Formatted Output	29
4.1 The Output Subsystem	44
5. Formatting Lisp Code	51
6. The Input Editor Program Interface	53
6.1 How the Input Editor Works	53
6.2 Invoking the Input Editor	54
6.3 Input Editor Options	59
6.4 Displaying Prompts in the Input Editor	65
6.5 Displaying Help Messages in the Input Editor	66
6.6 Examples of Use of the Input Editor	66
6.7 Input Editor Messages to Interactive Streams	70
7. The :read and :print Stream Operations	73
8. Input Functions	75
9. Output Functions	87
10. Accessing Files	91
10.1 Loading Files	108
10.2 File Attribute Lists	109

10.3	File Stream Operations	113
10.4	Direct Access File Streams	115
10.4.1	Stream Messages	115
11.	Accessing Directories	117
11.1	Functions for Accessing Directories	117
	II. Files	125
12.	Naming of Files	127
12.1	Pathnames	127
12.1.1	Simple Usage of the Pathname System	129
12.1.2	Host Determination in Pathnames	136
12.1.3	Interning of Pathnames	138
12.1.4	Printing Pathnames	139
12.1.5	Values of Pathname Components	140
12.1.6	Directory Pathnames and Directory Pathnames as Files	142
12.1.7	Case in Pathnames	145
12.2	Defaults and Merging	147
12.3	Generic Pathnames	149
12.4	Relative Pathnames	149
12.5	Canonical Types in Pathnames	151
12.5.1	Correspondence of Canonical Types and Editor Modes	153
12.6	Wildcard Pathname Mapping	153
12.6.1	Wildcard Directory Mapping	155
12.7	Pathname Functions	156
12.8	Pathname Messages	165
12.9	Pathnames on Supported Host File Systems	174
12.9.1	LMFS	174
12.9.2	FEP File System	176
12.9.3	UNIX	176
12.9.4	UNIX 4.2	178
12.9.5	VAX/VMS	180
12.9.6	TOPS-20 and TENEX	183
12.9.7	Multics	184
12.9.8	ITS	185
12.9.9	Ms-dos Pathnames	186
12.9.10	Logical Pathnames	186
12.10	Init File Naming Conventions	200
12.11	Maclisp Conversion	201
13.	Lisp Machine File System	203
13.1	Introduction to LMFS	203
13.2	Concepts	203

13.3	Properties	204
13.4	Deletion, Expunging, and Versions	209
13.5	LMFS Links	210
13.6	LMFS Backup	212
13.7	Multiple Partitions	213
13.7.1	Adding a Partition to LMFS	213
14.	FEP File System Overview	215
14.1	Microcode Loads	215
14.2	World Loads	215
14.3	Configuration Files	215
14.4	How LMFS Uses the FEP File System	216
14.5	Virtual Memory	216
14.6	FEP File Comment Properties	216
14.7	Installing Microcode	217
14.8	Renaming FEP Files	219
14.9	Using a Spare World Load for Paging	219
14.10	Adding a Spare World Load as LMFS File Space	219
15.	FEP File System	221
15.1	Naming of FEP Files	222
15.2	Accessing FEP Files	222
15.3	Operating on Disk Streams	224
15.4	Input and Output Disk Streams	226
15.5	Block Disk Streams	226
15.6	FEP File Properties	227
15.7	FEP File Locks	228
15.8	FEP File Types	228
16.	Fsedit	231
16.1	File System Editor	231
16.1.1	Entering the File System Editor	231
16.1.2	Using the File System Editor	232
16.1.3	Opening and Closing a Directory	232
16.1.4	Using Fsedit Commands	233
16.2	How to Interpret Directory Listings	236
17.	Creating More Room on the Local Disk	239
17.1	Allocating Extra Paging Space	240
18.	Putting Data in Compiled Code Files	241

III. The Serial I/O Facility	243
19. Introduction to Serial I/O	245
20. Hardware Description for Serial I/O	247
21. The Serial I/O Stream	251
22. Parameters for Serial I/O	253
23. Simple Example: Serial I/O	257
24. Troubleshooting: Serial I/O	259
25. Notes on Serial I/O	261
26. Hdlc Serial I/O	263
27. Using the Terminal Program with Hosts Connected to the Serial Line	265
IV. Writing Programs That Use Magnetic Tape	267
28. The tape:make-stream Function	269
29. Messages to Tape Streams	273
30. Tape Error Flavors	277
Index	279

List of Tables

Table 1.	Assignment of RS-232 Signals to Pins	248
Table 2.	Assignment of RS-232 Signals to Pins in Asynchronous Null Modems	248

PART I.

Streams

1. Introduction to the I/O System

Symbolics-Lisp provides a powerful and flexible system for performing input and output to peripheral devices. To allow device independent I/O (that is, to allow programs to be written in a general way so that the program's input and output may be connected with any device), the Symbolics-Lisp I/O system provides the concept of an "I/O stream". What streams are, the way they work, and the functions to create and manipulate streams, are described in this document. This document also describes the Lisp "I/O" operations **read** and **print**.

2. The Character Set

Zetalisp represents characters as fixnums. The Symbolics computer's mapping between these numbers and the characters is listed here. The mapping is similar to ASCII, but somewhat modified to allow the use of the so-called SAIL extended graphics, while avoiding certain ambiguities present in ITS. For a long time ITS treated the BACKSPACE, c-H, and λ keys on the keyboard identically as character code 10 octal; this problem is avoided from the start in the Symbolics computer's mapping.

It is worth pointing out that although the Zetalisp character set is different from the PDP-10 character set, when files are transferred between Symbolics computers and PDP-10s, the characters are automatically converted. Details of the mapping are explained below.

Fundamental characters are eight bits wide. Those less than 200 octal (with the 200 bit off), and only those, are "printing graphics"; when output to a device they are assumed to print a character and move the "cursor" one character position to the right. (All software provides for variable-width fonts, so the term "character position" should not be taken too literally.)

Characters in the range of 200 to 236 inclusive are used for special characters. Character 200 is a "null character", which does not correspond to any key on the keyboard. The null character is not used for anything much. Characters 201 through 236 correspond to the special function keys on the keyboard such as RETURN. Some characters are reserved for future expansion.

It should never be necessary for a user or a source program to know these numerical values. Indeed, they are likely to be changed in the future. There are symbolic names for all characters; see below.

Most of the special characters do not normally appear in files (although it is not forbidden for files to contain them). These characters exist mainly to be used as "commands" from the keyboard.

A few special characters, however, are "format effectors" which are just as legitimate as printing characters in text files. The following is a list of the names and meanings of these characters:

RETURN	The "carriage return" character which separates lines of text. Note that the PDP-10 convention that lines are ended by a pair of characters, "carriage return" and "line feed", is not used.
PAGE	The "page separator" character which separates pages of text.
TAB	The "tabulation" character which spaces to the right until the next "tab stop". Tab stops are normally every 8 character positions.

The Space character is considered to be a printing character whose printed image happens to be blank, rather than a format effector.

In some contexts, a fixnum can hold both a character code and a font number for that character. The following byte specifiers are defined:

%%ch-char *Variable*

The value of **%%ch-char** is a byte specifier for the field of a fixnum character which holds the character code.

Since **%%ch-char** will not exist after the switch to character objects has occurred, use **char-code** instead.

%%ch-font *Variable*

The value of **%%ch-font** is a byte specifier for the field of a fixnum character which holds the font number.

%%ch-font will not exist after the switch to character objects has occurred. There is no direct equivalent for this variable for you to use instead, since the concept of fonts will be reorganized.

Characters read in from the keyboard include a character code and modifier bits. A character cannot contain both a font number and modifier bits, since these data are both stored in the same bits. The following byte specifiers are provided:

%%kbd-char *Variable*

The value of **%%kbd-char** is a byte specifier for the field of a keyboard character which holds the normal eight-bit character code.

Since **%%kbd-char** will not exist after the switch to character objects has occurred, use **char-code** instead.

%%kbd-control *Variable*

The value of **%%kbd-control** is a byte specifier for the field of a keyboard character which is 1 if either CONTROL key was held down.

Since **%%kbd-control** will not exist after the switch to character objects has occurred, use **char-bit** instead.

%%kbd-meta *Variable*

The value of **%%kbd-meta** is a byte specifier for the field of a keyboard character which is 1 if either META key was held down.

Since **%%kbd-meta** will not exist after the switch to character objects has occurred, use **char-bit** instead.

%%kbd-super *Variable*

The value of **%%kbd-super** is a byte specifier for the field of a keyboard character which is 1 if either SUPER key was held down.

Since **%%kbd-super** will not exist after the switch to character objects has occurred, use **char-bit** instead.

%%kbd-hyper *Variable*

The value of **%%kbd-hyper** is a byte specifier for the field of a keyboard character which is 1 if either HYPER key was held down.

Since **%%kbd-hyper** will not exist after the switch to character objects has occurred, use **char-bit** instead.

%%kbd-control-meta *Variable*

The value of **%%kbd-control-meta** is a byte specifier for the four-bit field of a keyboard character which contains the control bits. The least significant bit is control. The most significant bit is hyper.

Since **%%kbd-control-meta** will not exist after the switch to character objects has occurred, use **char-bits** instead.

The following fields are used by some programs that encode signals from the mouse in a the format of a character. The generation of such characters is described elsewhere. See the document *Programming the User Interface*.

%%kbd-mouse *Variable*

The value of **%%kbd-mouse** is a byte specifier for the bit in a keyboard character which indicates that the character is not really a character, but a signal from the mouse.

Since **%%kbd-mouse** will not exist after the switch to character objects has occurred, use **mouse-char-p** instead.

%%kbd-mouse-button *Variable*

The value of **%%kbd-mouse-button** is a byte specifier for the field in a mouse signal which says which button was clicked. The value is **0**, **1**, or **2** for the left, middle, or right button, respectively.

Since **%%kbd-mouse-button** will not exist after the switch to character objects has occurred, use **char-mouse-button** instead.

%%kbd-mouse-n-clicks *Variable*

The value of **%%kbd-mouse-n-clicks** is a byte specifier for the field in a mouse signal which says how many times the button was clicked. The value is one less than the number of times the button was clicked.

Since **%%kbd-mouse-n-clicks** will not exist after the switch to character objects has occurred, use **char-mouse-n-clicks** instead.

When any of the modifier bits (control, meta, super, or hyper) is set in conjunction with a letter, the letter is always uppcased.

The Control-Shift- characters are encoded separately. `c-sh-A` is not a synonym for `h-c-A`; they are distinct compound keystrokes.

In addition to the four modifier keys HYPER, SUPER, CTRL, and META, the SHIFT key is a modifier key for letters when used in combination with one of the other modifiers. The CAPS LOCK key is not a modifier key and is always ignored in compound keystrokes. Thus typing CTRL and A at the same time gives `c-A`; typing CTRL and SHIFT and A at the same time gives `c-sh-A`. Typing CTRL and SHIFT and / at the same time gives `c-?` (not `c-sh-/?`).

The names for compound key strokes always show a letter as capitalized. This does not mean that you have to use the SHIFT key; use the SHIFT key as a modifier only when `sh-` appears in the key name.

In addition, printing names of characters have case in them.. Case is ignored on input. Some new synonyms for existing characters are accepted. In particular, names of the following form have these synonyms:

<i>Name</i>	<i>Equivalent to</i>
<code>#\c-sh-B</code>	<code>#\c-shift-B</code>
<code>#\mouse-L</code>	<code>#\mouse-L-1</code>

The names of the characters are in the table in `sys:io;rddefs.lisp`.

When characters are written to a file server computer that normally uses the ASCII character set to store text, Symbolics characters are mapped into an encoding that is reasonably close to an ASCII transliteration of the text. When a file is written, the characters are converted into this encoding, and the inverse transformation is done when a file is read back. No information is lost. Note that the length of a file, in characters, will not be the same measured in original Symbolics characters as it will measured in the encoded ASCII characters.

In TOPS-20, Tenex, and ITS, in the currently implemented ASCII file servers, the following encoding is used. All printing characters and any characters not mentioned explicitly here are represented as themselves. Codes 010 (lambda), 011 (gamma), 012 (delta), 014 (plus-minus), 015 (circle-plus), 177 (integral), 200 through 207 inclusive, 213 (delete/vt), and 216 and anything higher, are preceded by a 177; that is, 177 is used as a "quoting character" for these codes. Codes 210 (overstrike), 211 (tab), 212 (line), and 214 (page), are converted to their ASCII cognates, namely 010 (backspace), 011 (horizontal tab), 012 (line feed), and 014 (form feed) respectively. Code 215 (return) is converted into 015 (carriage return) followed by 012 (line feed). Code 377 is ignored completely, and so cannot be stored in files.

000 center-dot (·)	040 space	100 @	140 `
001 down arrow (↓)	041 !	101 A	141 a
002 alpha (α)	042 "	102 B	142 b
003 beta (β)	043 #	103 C	143 c
004 and-sign (^)	044 \$	104 D	144 d
005 not-sign (-)	045 %	105 E	145 e
006 epsilon (ε)	046 &	106 F	146 f
007 pi (π)	047 '	107 G	147 g
010 lambda (λ)	050 (110 H	150 h
011 gamma (γ)	051)	111 I	151 i
012 delta (δ)	052 *	112 J	152 j
013 up-arrow (↑)	053 +	113 K	153 k
014 plus-minus (±)	054 ,	114 L	154 l
015 circle-plus (⊕)	055 -	115 M	155 m
016 infinity (∞)	056 .	116 N	156 n
017 partial delta (∂)	057 /	117 O	157 o
020 left horseshoe (⋄)	060 0	120 P	160 p
021 right horseshoe (⋅)	061 1	121 Q	161 q
022 up horseshoe (⋆)	062 2	122 R	162 r
023 down horseshoe (⋇)	063 3	123 S	163 s
024 universal quantifier (∀)	064 4	124 T	164 t
025 existential quantifier (∃)	065 5	125 U	165 u
026 circle-X (⊗)	066 6	126 V	166 v
027 double-arrow (↔)	067 7	127 W	167 w
030 left arrow (←)	070 8	130 X	170 x
031 right arrow (→)	071 9	131 Y	171 y
032 not-equals (≠)	072 :	132 Z	172 z
033 diamond (altmode) (⋄)	073 ;	133 [173 {
034 less-or-equal (≤)	074 <	134 \	174
035 greater-or-equal (≥)	075 =	135]	175 }
036 equivalence (≡)	076 >	136 ^	176 ~
037 or (∨)	077 ?	137 _	177 f
200 null character	210 back space	220 stop-output	230 reserved
201 suspend	211 tab	221 abort	231 reserved
202 clear-input	212 line	222 resume	232 reserved
203 reserved	213 refresh	223 reserved	233 reserved
204 function	214 page	224 end	234 reserved
205 macro	215 return	225 square	235 select
206 help	216 quote	226 circle	236 network
207 rubout	217 hold-output	227 triangle	237 escape
240 reserved			
241 symbol-help			
242-377			

The Symbolics Character Set

2.1 Effect of Character Set Translation on Direct Access File Streams

The Symbolics generic file access protocol was designed to provide access to ASCII-based file systems for Symbolics computers. Symbolics machines support 8-bit characters and have 256 characters in their character set. This results in difficulties when communicating with ASCII machines, which have 7-bit characters.

The file server, on machines not using the Symbolics character set, is required to perform character translations for any character (not binary) opening. Some Symbolics characters expand to more than one ASCII character. Thus, for character files, when we speak of a given position in a file or the length of a file, we must specify whether we are speaking in *Symbolics units* or *server units*.

This causes major problems in file position reckoning. It is useless for the Symbolics machine (or other user side) to carefully monitor file position, counting characters, during output, when character translation is in effect. This is because the operating system interface for "position to point x in a file", which the server must use, operates in server units, but the Symbolics machine (or other user end) has counted in Symbolics units. The user end cannot try to second-guess the translation-counting process without losing host independence.

Since direct access file streams are designed for organized file position management, they are particularly susceptible to this problem. As with other file streams, it is only a problem when character files are used.

You can avoid this problem by always using binary files. If you must use character files, consider doing one of the following:

- Know the expansions of the Symbolics machine, that is, characters such as **Return** that do not expand into single host characters. Note that this sacrifices host independence.
- Do not use these characters. See the section "Qfile Character Set Translation" in *Networks*. This section explains which characters are expanded on the Symbolics computer.

3. I/O Streams

3.1 Introduction to Streams

Many programs accept input characters and produce output characters. Methods for performing input and output vary greatly from one device to another. Programs should be able to use any device available without each program having to know about each device.

The concept of *streams* solves this problem. A stream is a source and/or sink of characters. A set of *operations* is available with every stream; operations include such actions as "output a character" and "input a character". The way to perform an operation to a stream is the same for all streams, although what happens inside a stream depends on the kind of stream it is. Thus a program needs to know only how to deal with streams in general.

A stream is a message-receiving object, which means that you can apply it to arguments. The first argument is a keyword symbol that is the name of the operation you wish to perform. The remaining arguments depend on the operation. Message-passing is explained elsewhere: See the section "Flavors" in *Reference Guide to Symbolics-lisp*.

Some streams can do only input, some only output, and some can do both. Some streams support only some operations; however, unsupported operations might work, although slowly, because the **stream-default-handler** can handle them. An operation called **:which-operations** returns a list of the names of all operations that are supported "natively" by a stream. (All streams support **:which-operations**, so it might not be in the list itself.)

3.2 General-purpose Stream Operations

:tyo *char*

Message

The stream will output the character *char*. For example, if *s* is bound to a stream, then the following form will output a "B" to the stream:

```
(funcall s ' :tyo #\B)
```

For binary output streams, the argument is a nonnegative number rather than specifically a character.

:tyi &optional *eof*

Message

The stream will input one character and return it. For example, if the next character to be read in by the stream is a "C", then the following form returns the value of `#\C` (that is, 103 octal):

```
(funcall s ':tyi)
```

Note that the **:tyi** operation does not "echo" the character in any fashion; it only does the input. The **tyi** function echoes when reading from the terminal.

The optional *eof* argument to the **:tyi** message tells the stream what to do if it reaches the end of the file. If the argument is not provided or is **nil**, the stream returns **nil** at the end of file. Otherwise it signals an error and prints out the argument as the error message. Note that this is not the same as the *eof-option* argument to **read**, **tyi**, and related functions.

The **:tyi** operation on a binary input stream returns a nonnegative number, not necessarily to be interpreted as a character.

An EOF can be forced into the currently selected I/O buffer with the keystrokes **FUNCTION END**. The next **:tyi** message sent to a window taking input from that I/O buffer will return **nil**.

The EOF indicator is not "sticky," in that the next **:tyi** will take the next character from the I/O buffer. The reason for this is that some programs which read only from the terminal might not be prepared to encounter an EOF, and might loop trying to read input, since they wouldn't know to send the **:clear-eof** message.

This EOF feature makes it possible to fully test programs which use the **:line-in**, **:string-in**, and **:string-line-in** operations by taking input from a window instead of from a file. Typing **FUNCTION END** causes each of these operations to return. This is especially important when debugging programs which use the **:string-in** operation, since **:string-in** returns only when its buffer is full or an EOF is encountered.

FUNCTION END activates any input buffered in the input editor, since there is no representation for the EOF indicator within text strings.

:untyi *char*

Message

The stream will remember the character *char*, and the next time a character is input, it will return the saved character. In other words, **:untyi** means "put this character back into the input source". For example:

```
(funcall s ':untyi 120)
(funcall s ':tyi) ==> 120
```

This operation is used by **read**, and any stream that supports **:tyi** must support **:untyi** as well. Note that you are allowed to **:untyi** only one character before doing a **:tyi**, and you can **:untyi** only the last character you read from the stream. Some streams implement **:untyi** by saving the character, while others implement it by backing up the pointer to a buffer. You also cannot **:untyi** after you have peeked ahead with **:tyipeek**.

:which-operations*Message*

The object should return a list of the messages it can handle. The **:which-operations** method of **si:vanilla-flavor** generates the list once per flavor and remembers it, minimizing consing and compute time. If a new method is added, the list is regenerated the next time someone asks for it.

:operation-handled-p *operation**Message*

operation is a message name. The object should return **t** if it has a handler for the specified message, **nil** if it does not.

:send-if-handles *operation* &rest *arguments**Message*

operation is a message name and *arguments* is a list of arguments for that message. The object should send itself that message with those arguments if it handles the message. If it does not handle the message it should return **nil**.

:characters*Message*

Returns **t** if the stream is a character stream, **nil** if it is a binary stream.

:direction*Message*

Returns one of the keyword symbols **:input**, **:output**, or **:bidirectional**.

:interactive*Message*

The **:interactive** message to a stream returns **t** if the stream is interactive and **nil** if it is not. Interactive streams, built on **si:interactive-stream**, are streams designed for interaction with human users. They support input editing. Use the **:interactive** message to find out whether a stream supports the **:input-editor** message.

Any stream must either support **:tyo** or support both **:tyi** and **:untyi**. Several more advanced input and output operations work on any stream that can do input or output (respectively). Some streams support these operations themselves; you can tell by looking at the list returned by the **:which-operations** operation. Others are handled by the "stream default handler" even if the stream does not know about the operation itself. However, in order for the default handler to do one of the more advanced output operations, the stream must support **:tyo**, and for the input operations the stream must support **:tyi** (and **:untyi**).

Here is the list of such operations:

:input-wait &optional *whostate function* &rest *arguments**Message*

This message to an input stream causes the stream to **process-wait** with *whostate* until one of the following conditions is met:

- Applying *function* to *arguments* returns non-**nil**.
- The stream enters a state in which sending it a **:tyi** message would immediately return a value or signal an error.

When either of these conditions is met, **:input-wait** returns. If the stream enters a state in which sending it a **:tyi** message would signal an error, **:input-wait** returns instead of signalling the error. The returned value is not defined.

whostate is what to display in the status line while process-waiting. It can be a string or **nil**. A value of **nil** means to use the normal whostate for this stream, such as "Tyi", "Net In", or "Serial In". For interactive streams, the default whostate is "Tyi".

function can be a function or **nil**. A value of **nil** means that the stream just waits until sending it a **:tyi** message would immediately return a value or signal an error.

This message is intended for programs that need to wait until either input is available from some interactive stream or some other condition, such as the arrival of a notification, occurs. Any stream that can become the value of **terminal-io** must support **:input-wait**.

Following is a simple example of the use of **:input-wait** to wait for input or a notification to an interactive stream. The function just displays notifications and prints representations of characters or blips received as input.

```
(defun my-top-level (stream)
  (error-restart-loop ((error sys:abort) "My top level")
    (send stream :input-wait nil
      #'(lambda (note-cell)
          (not (null (location-contents note-cell))))
        (send stream :notification-cell)))
    (let ((note (send stream :receive-notification)))
      (if note
        (sys:display-notification stream note :stream)
        (let ((char (send stream :any-tyi-no-hang)))
          (cond ((null char)
                ((fixp char)
                 (format stream "~&Character: ~C" char))
                ((listp char)
                 (format stream "~&Blip: ~S" char))
                (t (format stream "~&Unknown object: ~S" char))))))))))
```

:listen

Message

On an interactive device, returns non-**nil** if any input characters are immediately available, or **nil** if no input is immediately available. On a noninteractive device, the operation always returns non-**nil** except at end-of-file, by virtue of the default handler. The main purpose of **:listen** is to test whether the user has pressed a key, perhaps trying to stop a program in progress.

- :tpeek** &optional *eof* *Message*
 On an input stream, returns the next character that is about to be read, or **nil** if the stream is at end-of-file. The *eof* argument has the same meaning as it does for **:tyi**. **:tpeek** is defined to have the same effect as a **:tyi** operation, followed by a **:untyi** operation if end-of-file is not reached. Note that this means that you cannot read some character, do a **:tpeek** to look at the next character, and then **:untyi** the original character.
- :fresh-line** *Message*
 Tells the stream to position itself at the beginning of a new line. If the stream is already at the beginning of a fresh line it does nothing; otherwise it outputs a carriage return. For streams that do not support this, the default handler always outputs a carriage return.
- :clear-rest-of-line** *Message*
 Erases from the current position to the end of the current line.
- :string-out** *string* &optional *start end* *Message*
 The characters of *string* are successively output to the stream. This operation is provided for two reasons: it saves the writing of a frequently used loop, and many streams can perform this operation much more efficiently than the equivalent sequence of **:tyo** operations. If the stream does not support **:string-out** itself, the default handler converts it to **:tyos**.
 If *start* and *end* are not supplied, the entire string is output. Otherwise a substring is output; *start* is the index of the first character to be output (defaulting to 0), and *end* is one greater than the index of the last character to be output (defaulting to the length of the string). Callers need not pass these arguments, but all streams that handle **:string-out** must check for them and interpret them appropriately.
- :line-out** *string* &optional *start end* *Message*
 The characters of *string*, followed by a carriage return character, are output to the stream. *start* and *end* optionally specify a substring, as with **:string-out**. If the stream does not support **:line-out** itself, the default handler converts it to **:tyos**.
- :string-in** *eof-option string* &optional (*start 0*) *end* *Message*
 Reads characters from an input stream into *string*, using the substring delimited with *start* and *end*.
 As is usual with strings, *start* defaults to 0 and *end* defaults to the length of the string. The difference between *end* and *start* constitutes a character count for this operation.
eof-option specifies stopping actions.

<i>Value</i>	<i>Meaning</i>
nil	Reading characters into the string stops either when it has transferred the specified character count or when it reaches end-of-file, whichever happens first. For strings with a fill pointer, it sets the fill pointer to point to the location following the last one filled by the read.
not nil	If the end-of-file is encountered while trying to transfer a specific number of characters, it signals sys:end-of-file , with the value of <i>eof</i> as the report string.

:string-in returns two values. The first value is one greater than the last location of *string* into which it stored a character. The second value is **t** if it reached end-of-file and **nil** if it did not. Using **:string-in** at the end of a file returns 0 and **t** and sets the fill pointer of *string* to *start* (if *string* has a fill pointer).

For example, suppose the file `my-host:>george>tiny.text` contains "Here is some tiny text."

```
(setq string (make-array 100 ':type 'art-string ':fill-pointer 0))
""

(with-open-file (stream "my-host:>george>tiny.text")
  (send stream ':string-in nil string))
23

string => "Here is some tiny text."
```

If *string* has an array-leader, the fill pointer is adjusted to *start* plus the number of characters stored into *string*.

string can be any kind of array, not necessarily a string; this is useful when reading from a binary input stream.

The **:string-in** message can be sent to windows. It interacts correctly with the input editor, including correct handling of activation characters.

The interface to this method for windows and the returned value is exactly the same as the equivalent methods for **si:input-stream** and **si:unbuffered-line-input-stream**.

:line-in &optional *leader*

Message

The stream should input one line from the input source and return it as a string with the carriage return character stripped off. Despite its name, this operation is not much like the **readline** function.

Many streams have a string that is used as a buffer for lines. If this string itself were returned, there would be problems if the caller of the stream attempted to save the string away somewhere, because the contents of the string would change when the next line was read in. To solve this problem,

the string must be copied. On the other hand, some streams do not reuse the string, and it would be wasteful to copy it on every **:line-in** operation. This problem is solved by using the *leader* argument to **:line-in**. If *leader* is **nil** (the default), the stream does not copy the string, and the caller should not rely on the contents of that string after the next operation on the stream. If *leader* is **t**, the stream makes a copy. If *leader* is an integer then the stream makes a copy with an array-leader *leader* elements long. (This is used by the editor, which represents lines of buffers as strings with additional information in their array-leaders, to eliminate an extra copy operation.)

If the stream reaches the end-of-file while reading in characters, it returns the characters it has read in as a string, and returns a second value of **t**. The caller of the stream should therefore arrange to receive the second value, and check it to see whether the string returned was an whole line or only the trailing characters after the last carriage return in the input source.

The **:line-in** message can be sent to windows. It interacts correctly with the input editor, including correct handling of activation characters.

:string-line-in *eof string &optional (start 0) end* *Message*

:string-line-in is a combination of **:string-in** and **:line-in**. It allows you to read many lines successively into the same buffer without creating strings.

:string-line-in reads a line from a file into a string (or other array) supplied by the user.

This message fills up a string as does **:string-in**, but reads only one line, as does **:line-in**. As with **:line-in**, the carriage return character at the end of the line is not stored into your buffer. **:line-in** reads a line from a stream and creates a string with that line in it. **:string-in** is given a string; it fills in the string (or other array) that you give it from the stream.

:string-line-in reads a line from a stream and fills the supplied array with that line. As with **:string-in**, if the string (or other array) has a fill pointer, it is set to the number of characters placed into the buffer.

:string-line-in returns three values:

- How many characters it read into your buffer. (This might be zero.)
- Whether the end of the input stream was encountered while trying to read in the string. *eof* is identical to the *eof* argument in **:string-in**.
- **nil** if the entire line fit in the buffer supplied, otherwise **t**. If **t** is returned for this value, as much of the line as could fit was stored in the buffer and more of the line is waiting to be read.

If the second and third values are both **nil**, a carriage return was read. If either is **t**, no carriage return was read from the stream.

:clear-input *Message*

The stream clears any buffered input. If the stream does not handle this, the default handler ignores it.

:clear-output *Message*

The stream clears any buffered output. If the stream does not handle this, the default handler ignores it.

:force-output *Message*

Causes any buffered output to be sent to a buffered asynchronous device, such as the Chaosnet. It does not wait for it to complete; use **:finish** for that. If a stream supports **:force-output**, then **:tyo**, **:string-out**, and **:line-out** might have no visible effect until a **:force-output** is done. If the stream does not handle this, the default handler ignores it.

:finish *Message*

Does a **:force-output** to a buffered asynchronous device, such as the Chaosnet, then waits until the currently pending I/O operation has been completed. If the stream does not handle this, the default handler ignores it.

For file output streams, **:finish** finalizes file content. It ensures that all data have actually been written to the file, and sets the byte count. It converts non-direct output openings into append openings. It allows other users to access the data that have been written before the **:finish** message was sent.

:close *&optional mode* *Message*

The stream is "closed", and no further operations should be performed on it; you can, however, **:close** a closed stream. If the stream does not handle **:close**, the default handler ignores it.

The *mode* argument is normally not supplied. If it is **:abort**, we are abnormally exiting from the use of this stream. If the stream is outputting to a file, and has not been closed already, the stream's newly created file is deleted, as if it were never opened in the first place. Any previously existing file with the same name remains, undisturbed.

:eof *Message*

Indicates the end of data on an output stream. This is different from **:close** because some devices allow multiple data files to be transmitted without closing. **:close** implies **:eof** when the stream is an output stream and the close mode is not **:abort**.

3.3 Special-purpose Stream Operations

See the section "General-purpose Stream Operations", page 11. There are several other defined operations that the default handler cannot deal with; if the stream does not support the operation itself, sending that message causes an error. This section describes the most commonly used, least device-dependent stream operations. Windows, files, and Chaosnet connections have their own special stream operations, which are documented separately.

:input-editor *function &rest arguments* *Message*

This is supported by interactive streams such as windows. It is described in its own section: See the section "The Input Editor Program Interface", page 53.

Most programs should not send this message directly. See the special form **with-input-editing**, page 57.

:beep *&optional type* *Message*

This is supported by interactive streams. It attracts the attention of the user by making an audible beep and/or flashing the screen. *type* is a keyword selecting among several different beeping noises. The allowed types have not yet been defined; *type* is currently ignored and should always be **nil**.

:tyi-no-hang *&optional eof* *Message*

Identical to **:tyi** except that if it would be necessary to wait in order to get the character, returns **nil** instead. This lets the caller efficiently check for input being available and get the input if there is any. **:tyi-no-hang** is different from **:listen** because it reads a character and because it is not simulated by the default handler for streams that do not support it.

:untyo-mark *Message*

This is used by the grinder if the output stream supports it. See the special form **grindef**, page 51. It takes no arguments. The stream should return some object that indicates where output has reached in the stream.

:untyo mark *Message*

This is used by the grinder in conjunction with **:untyo-mark**. See the special form **grindef**, page 51. It takes one argument, which is something returned by the **:untyo-mark** operation of the stream. The stream should back up output to the point at which the object was returned.

:read-cursorpos *&optional (units ':pixel)* *Message*

This operation is supported by windows. It returns two values, the current *x* and *y* coordinates of the cursor. It takes one optional argument, which is a symbol indicating in what units *x* and *y* should be; the symbols **:pixel** and **:character** are understood. **:pixel** means that the coordinates are measured

in display pixels (bits), while **:character** means that the coordinates are measured in characters horizontally and lines vertically.

This operation and **:set-cursorpos** are used by the format "**T**" request, which is why "**T**" does not work on all streams. Any stream that supports this operation must support **:set-cursorpos** as well.

:set-cursorpos *x y* &optional (*units* **:pixel**) *Message*

This operation is supported by the same streams that support **:read-cursorpos**. It sets the position of the cursor. *x* and *y* are similar to the values of **:read-cursorpos** and *units* is the same as the *units* argument to **:read-cursorpos**.

:clear-window *Message*

Erases the window on which this stream displays. Non-window streams do not support this operation.

There are many other special-purpose stream operations for graphics. See the section "Using the Window System" in *Programming the User Interface*.

The following operations are only implemented by streams to random-access devices, principally files.

:read-pointer *Message*

Returns the current position within the file, in characters (bytes in fixnum mode). For text files on PDP-10 file servers, this is the number of Symbolics characters, not PDP-10 characters. The numbers are different because of character-set translation.

:set-pointer *new-pointer* *Message*

Sets the reading position within the file to *new-pointer* (bytes in fixnum mode). For text files on PDP-10 file servers, this does not do anything reasonable unless *new-pointer* is 0, because of character-set translation. This operation is for input streams only.

The following operations are implemented by buffered input streams. They allow increased efficiency by making the stream's internal buffer available to the user.

:read-input-buffer &optional *eof* *Message*

Returns three values: a buffer array, the index in that array of the next input byte, and the index in that array just past the last available input byte. These values are similar to the *string*, *start*, *end* arguments taken by many functions and stream operations. If the end of the file has been reached and no input bytes are available, the stream returns **nil** or signals an error, based on the *eof* argument, just like the **:tyi** message. After reading as many bytes from the array as you care to, you must send the **:advance-input-buffer** message.

:advance-input-buffer &optional *new-pointer* *Message*
 If *new-pointer* is non-**nil**, it is the index in the buffer array of the next byte to be read. If *new-pointer* is **nil**, the entire buffer has been used up.

The following operations are provided for buffered output streams. They allow you to hand the stream's output buffer to a function that can fill it up.

:get-output-buffer *Message*
 Returns an array and starting and ending indices.

:advance-output-buffer &optional *index* *Message*
 Says that the array returned by the last **:get-output-buffer** operation was filled up through *index*. If *index* is omitted, the array was filled completely.

The following stream operations are obsolete and should no longer be used:

:rewind
:get-input-buffer

3.4 Standard Streams

Several variables whose values are streams are used by many functions in the Lisp system. By convention, variables that are expected to hold a stream capable of input have names ending with **-input**, and similarly for output. Those expected to hold a bidirectional stream have names ending with **-io**.

The variables **standard-input**, **standard-output**, **error-output**, **trace-output**, and **query-io** are initially bound to synonym streams that pass all operations on to the stream that is the value of **terminal-io**. Thus any operation performed on those streams goes to the terminal.

No user program should ever change the value of **terminal-io**. For example, a program to divert output to a file should do so by binding the value of **standard-output**; that way, error messages sent to **error-output** can still get to the user by going through **terminal-io**, which is usually what is desired.

standard-input *Variable*
 In the normal Lisp top-level loop, input is read from **standard-input** (that is, whatever stream is the value of **standard-input**). Many input functions, including **tyi** and **read**, take a stream argument that defaults to **standard-input**.

standard-output *Variable*
 In the normal Lisp top-level loop, output is sent to **standard-output** (that is, whatever stream is the value of **standard-output**). Many output functions, including **tyo** and **print**, take a stream argument that defaults to **standard-output**.

error-output*Variable*

The value of **error-output** is a stream to which error messages should be sent. Normally this is the same as **standard-output**, but **standard-output** might be bound to a file and **error-output** left going to the terminal.

query-io*Variable*

The value of **query-io** is a stream that should be used when asking questions of the user. The question should be output to this stream, and the answer read from it. The reason for this is that when the normal input to a program might be coming from a file, questions such as "Do you really want to delete all of the files in your directory??" should be sent directly to the user, and the answer should come from the user, not from the data file. **query-io** is used by **fquery** and related functions.

terminal-io*Variable*

The value of **terminal-io** is the stream that connects to the user's console. In an "interactive" program, it is the window from which the program is being run; I/O on this stream reads from the keyboard and displays on the terminal. However, in a "background" program that does not normally talk to the user, **terminal-io** defaults to a stream that does not ever expect to be used. If it is used, perhaps by an error notification, it turns into a "background" window and requests the user's attention.

trace-output*Variable*

The value of **trace-output** is the stream on which the **trace** function prints its output.

debug-io*Variable*

If not **nil**, this is the stream that the Debugger should use. The default value is a synonym stream that is synonymous with **terminal-io**. If the value of **dbg:*debug-io-override*** is not **nil**, the Debugger uses the value of that variable as the stream instead of the value of **debug-io**.

The value of **debug-io** can also be a string. This causes the debugger to use the cold-load stream; the string is the reason why the cold-load stream should be used.

No program other than the Debugger should do stream operations on the value of **debug-io**, since the value cannot be a stream. Other programs should use **query-io**, **error-output**, or **trace-output**.

dbg:*debug-io-override**Variable*

This is used during debugging to divert the Debugger to a stream that is known to work. If the value of this variable is **nil** (the default), the Debugger uses the stream that is the value of **debug-io**. But if the value of **dbg:*debug-io-override*** is not **nil**, the Debugger uses the stream that is the value of this variable instead. This variable should always be set (using **setq**), not bound, so all processes and stack groups can see it.

make-syn-stream *symbol* *Function*

make-syn-stream creates and returns a "synonym stream" (syn for short). *symbol* can be either a symbol or a locative.

If *symbol* is a symbol, the synonym stream is actually an uninterned symbol named `#:symbol-syn-stream`. This generated symbol has a property that declares it to be a legitimate stream. This symbol is the value of *symbol*'s **si:syn-stream** property, and its function definition is forwarded to the value cell of *symbol* using a **dtp-external-value-cell-pointer**. Any operations sent to this stream are redirected to the stream that is the value of *symbol*.

If *symbol* is a locative, the synonym stream is an uninterned symbol named `#:syn-stream`. This generated symbol has a property that declares it to be a legitimate stream. The function definition of this symbol is forwarded to the cell designated by *symbol*. Any operations sent to this stream are redirected to the stream that is the contents of the cell to which *symbol* points.

Synonym streams should not be passed between processes, since the streams to which they redirect operations are specific to a process.

make-broadcast-stream &rest *streams* *Function*

Returns a stream that works only in the output direction. Any output sent to this stream is sent to all of the streams given. The **:which-operations** is the intersection of the **:which-operations** of all of the streams. The value(s) returned by a stream operation are the values returned by the last stream in *streams*.

3.5 Making Your Own Stream

Here is a sample output stream that accepts characters and conses them onto a list.

```
(defvar the-list nil)
(defun list-output-stream (op &optional arg1 &rest rest)
  (selectq op
    (:tyo
     (setq the-list (cons arg1 the-list)))
    (:which-operations '(:tyo))
    (otherwise
     (stream-default-handler (function list-output-stream)
                             op arg1 rest))))
```

The lambda-list for a stream must always have one required parameter (**op**), one optional parameter (**arg1**), and a rest parameter (**rest**). This allows an arbitrary number of arguments to be passed to the default handler. This is an output stream, so it supports the **:tyo** operation. Note that all streams must support **:which-operations**. If the operation is not one that the stream understands (for

example, **:string-out**), it calls the **stream-default-handler**. The calling of the default handler is required, since the willingness to accept **:tyo** indicates to the caller that **:string-out** will work.

Here is a typical input stream that generates successive characters of a list.

```
(defvar the-list)      ;Put your input list here
(defvar untyied-char nil)
(defun list-input-stream (op &optional arg1 &rest rest)
  (selectq op
    (:tyi
     (cond ((not (null untyied-char))
            (prog1 untyied-char (setq untyied-char nil)))
          ((null the-list)
            (and arg1 (error arg1)))
          (t (prog1 (car the-list)
                   (setq the-list (cdr the-list)))))))
    (:untyi
     (setq untyied-char arg1))
    (:which-operations '(:tyi :untyi))
    (otherwise
     (stream-default-handler (function list-input-stream)
                             op arg1 rest))))
```

The important things to note are that **:untyi** must be supported, and that the stream must check for having reached the end of the information and do the right thing with the argument to the **:tyi** operation.

The above stream uses a free variable (**the-list**) to hold the list of characters, and another one (**untyied-char**) to hold the **:untyied** character (if any). You might want to have several instances of this type of stream, without their interfering with one another. This is a typical example of the usefulness of closures in defining streams. The following function will take a list and return a stream that generates successive characters of that list.

```
(defun make-a-list-input-stream (list)
  (let-closed ((list list) (untyied-char nil))
    (function list-input-stream)))
```

The above streams are very simple. When designing a more complex stream, it is useful to have some tools to aid in the task. The **defselect** function aids in defining message-receiving functions. The Flavor System provides powerful and elaborate facilities for programming message-receiving objects. See the section "Flavors" in *Reference Guide to Symbolics-lisp*.

stream-default-handler *stream op arg1 rest*

Function

Tries to handle the *op* operation on *stream*, given arguments of *arg1* and the elements of *rest*. The action taken for each of the defined operations is explained with the documentation on that operation. The handler sends the **:any-tyi** message for **:line-in** messages to streams that do not handle **:line-in** themselves.

si:null-stream *op &rest args*

Function

Can be used as a dummy stream object. As an input stream, it immediately reports end-of-file; as an output stream, it absorbs and discards arbitrary amounts of output. Note: **si:null-stream** is not a variable; it is defined as a function. Use its definition (or the symbol itself) as a stream, not its value.

Examples:

```
(stream-copy-until-eof a 'si:null-stream)
(stream-copy-until-eof a #'si:null-stream)
```

Either of the above two forms reads characters out of the stream that is the value of **a** and throws them away, until **a** reaches the end-of-file.

3.6 Coroutine Streams

Functions that produce data as output (output functions) are written in terms of **:tyo** and other output operations. Functions that receive data as input (input functions) are written in terms of **:tyi** and other input operations. Output functions operate on output streams, which handle the **:tyo** message. Input functions operate on input streams, which handle the **:tyi** message. Sometimes it is desirable to view an output function as an input stream, or an input function as an output stream. You can do this with coroutine streams.

Here is a simplified explanation of how coroutine streams work. A coroutine input stream can be built from an output function. Whenever that stream receives a **:tyi** message, it invokes the output function in a separate stack group so that the function can produce the data that the **:tyi** message returns. A coroutine output stream can be built out of an input function; it works in the opposite fashion. Whenever the output stream receives a **:tyo** message, it invokes the input function in a separate stack group so that the function can receive the data transmitted by the **:tyo** message. It is also possible to connect functions that do both input and output, by using bidirectional coroutine streams. Since you can use coroutine streams to connect two functions, they are the logical inverse of **stream-copy-until-eof**, a function used to connect two streams.

To create a coroutine stream, use one of three functions.

- If you want to make an input stream from an output function, use **si:make-coroutine-input-stream**.
- If you want to make an output stream to an input function, use **si:make-coroutine-output-stream**.
- If you want to make a bidirectional stream for a function that does both input and output, use **si:make-coroutine-bidirectional-stream**.

Following is an example using a coroutine input stream:


```
(setq input-stream
      (si:make-coroutine-input-stream
       #'(lambda (stream) (print-disk-label 0 stream))))

(send input-stream ':line-in) →
"1645 free, 260499//262144 used (99%)"
```

Following is an example using a coroutine output stream:

```
(setq output-stream
      (si:make-coroutine-output-stream
       #'(lambda (stream) (setf x (read stream)))))

(send output-stream ':string-out "(a b c)")

(send output-stream ':force-output)

x → (A B C)
```

Coroutine streams are implemented as buffered character streams. Each function that makes a coroutine stream actually creates two streams and one new stack group. One stream is associated with the new stack group and the other stream with the stack group that is current when the stream-making function is called. If you use **si:make-coroutine-input-stream** or **si:make-coroutine-output-stream**, one stream is an input stream and the other is an output stream; they share a common buffer. If you use **si:make-coroutine-bidirectional-stream**, both streams are bidirectional; the input buffer of each stream is the output buffer of the other.

With **si:make-coroutine-input-stream**, the output function runs in the new stack group. With **si:make-coroutine-output-stream**, the input function runs in the new stack group. With bidirectional streams, the function that does input or output runs in the new stack group.

In the case of **si:make-coroutine-input-stream**, for example, you typically send **:tyi** messages to the input stream that **si:make-coroutine-input-stream** returns. The input stream is associated with the new stack group. When the input stream receives a **:tyi** message, the new stack group is resumed, and the output function runs in that stack group. The output function typically sends **:tyo** messages to the output stream associated with the stack group from which **si:make-coroutine-input-stream** was called. When the output stream receives a **:tyo** message, the associated stack group is resumed. The data transmitted to the output stream become input to **:tyi** via the buffer that the two streams share. **si:make-coroutine-output-stream** and **si:make-coroutine-bidirectional-stream** work in analogous fashion.

In addition to **:tyi** and **:tyo**, coroutine streams support other standard input and output operations, such as **:line-in** and **:string-out**. Actually, the **:next-input-buffer** method of the input stream and the **:send-output-buffer** method of the output stream resume the new stack group, not the receipt of **:tyi**

and **:tyo** messages. Because the streams are buffered, you must send a **:force-output** message to an output stream to cause the new stack group to be resumed.

The instantiable flavors of coroutine streams are **si:coroutine-input-stream**, **si:coroutine-output-stream**, and **si:coroutine-bidirectional-stream**.

Do not confuse coroutine streams with pipes. Coroutine streams are used for intraprocess communication; pipes are used for interprocess communication. 3600-family machines do not currently support pipes.

si:make-coroutine-input-stream *function &rest arguments* *Function*

Creates two coroutine streams, an input stream and an output stream, with a shared buffer. **si:make-coroutine-input-stream** returns the input stream. The input stream is associated with a new stack group and the output stream with the stack group that is current when **si:make-coroutine-input-stream** is called. **:tyi** messages to the input stream cause the new stack group to be resumed and *function* to be called from that stack group. The first argument to *function* is the output stream; any additional arguments come from *arguments*. *function* should send **:tyo** messages to the output stream. These messages resume the stack group in which **si:make-coroutine-input-stream** was called. In this way, output from *function* becomes input to the caller of **si:make-coroutine-input-stream** through the shared buffer.

si:make-coroutine-output-stream *function &rest arguments* *Function*

Creates two coroutine streams, an output stream and an input stream, with a shared buffer. **si:make-coroutine-output-stream** returns the output stream. The output stream is associated with a new stack group and the input stream with the stack group that is current when **si:make-coroutine-output-stream** is called. **:tyo** messages to the output stream cause the new stack group to be resumed and *function* to be called from that stack group. The first argument to *function* is the input stream; any additional arguments come from *arguments*. *function* should send **:tyi** messages to the input stream. These messages resume the stack group in which **si:make-coroutine-output-stream** was called. In this way, output from the caller of **si:make-coroutine-output-stream** becomes input to *function* through the shared buffer.

si:make-coroutine-bidirectional-stream *function &rest arguments* *Function*

Creates two bidirectional coroutine streams. The input buffer of each stream is the output buffer of the other. One stream is associated with a new stack group and the other with the stack group that is current when **si:make-coroutine-bidirectional-stream** is called. **si:make-coroutine-bidirectional-stream** returns the stream associated with the new stack group.

:tyi and **:tyo** messages to the stream associated with the new stack group cause that stack group to be resumed and *function* to be called from that stack group. The first argument to *function* is the stream associated with the stack group from which **si:make-coroutine-bidirectional-stream** was called. Any additional arguments come from *arguments*. *function* should send **:tyi** or **:tyo** messages to the stream that is its first argument. These messages resume the stack group in which **si:make-coroutine-output-stream** was called. In this way *function* and the caller of **si:make-coroutine-bidirectional-stream** communicate through the shared buffers; output from one function becomes input to the other.

si:coroutine-input-stream *Flavor*

Coroutine input stream. Defines a **:next-input-buffer** method. Use this to construct an input stream from a function written in terms of output operations.

si:coroutine-output-stream *Flavor*

Coroutine output stream. Defines **:new-output-buffer** and **:send-output-buffer** methods. Use this to construct an output stream to a function written in terms of input operations.

si:coroutine-bidirectional-stream *Flavor*

Bidirectional coroutine stream. Defines **:next-input-buffer**, **:new-output-buffer**, and **:send-output-buffer** methods. Use this to construct a bidirectional stream to a function written in terms of input and output operations.

4. Formatted Output

There are two ways of doing general formatted output: the **format** function and the **output** subsystem. **format** uses a control string written in a special format specifier language to control the output format. **output** provides Lisp functions to do output in particular formats.

For simple tasks in which only the most basic format specifiers are needed, **format** is easy to use and has the advantage of brevity. For more complicated tasks, the format specifier language becomes obscure and hard to read. Then **output** becomes preferable because it works with ordinary Lisp control constructs.

Additional tools are available for formatting Lisp code (as opposed to text and tables). See the section "Formatting Lisp Code", page 51.

format *destination control-string &rest args* *Function*

Produces formatted output. **format** outputs the characters of *control-string*, except that a tilde (~) introduces a directive. The character after the tilde, possibly preceded by prefix parameters and modifiers, specifies the kind of formatting desired. Most directives use one or more elements of *args* to create their output; the typical directive puts the next element of *args* into the output, formatted in some special way.

The output is sent to *destination*. If *destination* is **nil**, a string is created that contains the output; this string is returned as the value of the call to **format**. In all other cases **format** returns no interesting value (generally **nil**). If *destination* is a stream, the output is sent to it. If *destination* is **t**, the output is sent to **standard-output**. If *destination* is a string with an array-leader, such as would be acceptable to **string-nconc**, the output is added to the end of that string.

A directive consists of a tilde, optional prefix parameters separated by commas, optional colon (:), and at-sign (@) modifiers, and a single character indicating the kind of directive. The alphabetic case of the character is ignored. The prefix parameters are generally decimal numbers. Examples of control strings:

```

"~S"      ; This is an S directive with no parameters.
"~3,4:es" ; This is an S directive with two parameters, 3 and 4,
          ; and both the colon and at-sign flags.
"~,4S"    ; The first prefix parameter is omitted and takes
          ; on its default value, while the second is 4.

```

format includes some extremely complicated and specialized features. It is not necessary to understand all or even most of its features to use **format** efficiently. The more sophisticated features are there for the convenience of programs with complicated formatting requirements.

Sometimes a prefix parameter is used to specify a character, such as the padding character in a right- or left-justifying operation. In this case a single quote (') followed by the desired character can be used as a prefix parameter, so that you do not have to know the decimal numeric values of characters in the character set. For example, you can use the following to print a decimal number in five columns with leading zeros.

```
"~5,'0d" instead of "~5,48d"
```

In place of a prefix parameter to a directive, you can put the letter **V**, which takes an argument from *args* as a parameter to the directive. Normally this should be a number but it does not have to be. This feature allows variable column-widths and the like. Also, you can use the character **#** in place of a parameter; it represents the number of arguments remaining to be processed.

Here are some relatively simple examples of how **format** is used.

```
(format nil "foo") => "foo"
(setq x 5)
(format nil "The answer is ~D." x) => "The answer is 5."
(format nil "The answer is ~3D." x) => "The answer is 5."

(setq y "elephant")
(format nil "Look at the ~A!" y) => "Look at the elephant!"
(format nil "The character ~:@C is strange." #o1003)
=> "The character Meta-Beta (Symbol-shift-B) is strange."

(setq n 3)
(format nil "~D item~:P found." n) => "3 items found."
(format nil "~R dog~:[s are~; is~] here." n (= n 1))
=> "three dogs are here."
(format nil "~R dog~:*~[~1; is~;;s are~] here." n)
=> "three dogs are here."
(format nil "Here ~[~1;is~;;are~] ~:*~R pupp~:@P." n)
=> "Here are three puppies."
```

~A The next element from the *args* of the **format** function, any Lisp object, is printed without slashification (as by **princ**). **~:A** prints () if the element is **nil**; this is useful when printing something that is always supposed to be a list. **~nA** inserts spaces on the right, if necessary, to make the column width at least *n*. The **@** modifier causes the spaces to be inserted on the left rather than the right. **~mincol,colinc,minpad,padcharA** is the full form of **~A**, which allows elaborate control of the padding. The string is padded on the right with at least *minpad* copies of *padchar*; padding characters are then inserted *colinc* characters at a time until the total width is at least *mincol*. The defaults are 0 for *mincol* and *minpad*, 1 for *colinc*, and space for *padchar*.

~S The next element from the *args* of the **format** function, any Lisp

object, is printed with slashification (as by **prin1**). **~:S** prints () if the element is **nil**; this is useful when printing something that is always supposed to be a list. **~nS** inserts spaces on the right, if necessary, to make the column width at least *n*. The **@** modifier causes the spaces to be inserted on the left rather than the right. **~mincol,colinc,minpad,padcharS** is the full form of **~S**, which allows elaborate control of the padding. The string is padded on the right with at least *minpad* copies of *padchar*; padding characters are then inserted *colinc* characters at a time until the total width is at least *mincol*. The defaults are 0 for *mincol* and *minpad*, 1 for *colinc*, and space for *padchar*.

- ~D** The next element from the *args* of the **format** function, normally a number, is printed as a decimal integer. Unlike **print**, **~D** never puts a decimal point after the number. **~nD** uses a column width of *n*; spaces are inserted on the left if the number requires fewer than *n* columns for its digits and sign. If the number does not fit in *n* columns, additional columns are used as needed. **~n,mD** uses *m* as the pad character instead of space. The **@** modifier causes the number's sign to be printed always; the default is to print it if only the number is negative. The **:** modifier causes commas to be printed between groups of three digits; the third prefix parameter can be used to change the character used as the comma. Thus the most general form of **~D** is **~mincol,padchar,commacharD**.

If the element is not an integer, it is printed in **~A** format and decimal base. Thus this directive can be used to print some list structure showing all fixnums in decimal.

- ~O** The next element from the *args* of the **format** function, normally a number, is printed as an octal integer. **~nO** uses a column width of *n*; spaces are inserted on the left if the number requires fewer than *n* columns for its digits and sign. If the number does not fit in *n* columns, additional columns are used as needed. **~n,mO** uses *m* as the pad character instead of space. The **@** modifier causes the number's sign to be printed always; the default is to print it only if the number is negative. The **:** modifier causes commas to be printed between groups of three digits; the third prefix parameter can be used to change the character used as the comma. Thus the most general form of **~O** is **~mincol,padchar,commacharO**.

If the element is not an integer, it is printed in **~A** format and octal base. Thus this directive can be used to print some list structure showing all fixnums in octal.

~B Formats a number in binary. For example:

```
(format t "~B" 10.)
1010
NIL
```

~X In Common Lisp, formats a number in hexadecimal; in Zetalisp, prints spaces. For example:

```
(cl:format t "~X" 50.)
32
NIL
```

```
(format t "~X" 50.)
```

```
NIL
```

~F The next element from the *args* of the **format** function is printed in floating-point format. **~nF** rounds the element to a precision of *n* digits. The minimum value of *n* is 2, since a decimal point is always printed. If the magnitude of the element is too large or too small, it is printed in exponential notation. If the element is not a number, it is printed in **~A** format. Note that the prefix parameter *n* is not *mincol*; it is the number of digits of precision desired.

The Common Lisp version of **~F** produces a different format. Examples:

```
(format nil "~2F" 5) => "5.0"
(format nil "~4F" 5) => "5.0"
(format nil "~4F" 1.5) => "1.5"
(format nil "~4F" 3.14159265) => "3.142"
(format nil "~3F" 1e10) => "1.0e10"

(cl:format nil "~2F" 5) ==> #"5."
(cl:format nil "~4F" 3.14159265) ==> #"3.14"
```

~E The next element from the *args* of the **format** function is printed in exponential format. **~nE** rounds the element to a precision of *n* digits. The minimum value of *n* is 2, since a decimal point is always printed. If the element is not a number, it is printed in **~A** format. Note that the prefix parameter *n* is not *mincol*; it is the number of digits of precision desired.

The Common Lisp version of **~E** is not supported.

~\$ The format for using it follows:

```
~rdig,ldig,field,padchar$
```

It expects a flonum argument. The modifiers for **~\$** are all optional.

rdig The number of digits after the decimal point. The default is 2.

<i>ldig</i>	The minimum number of digits before the decimal point. The default is 1. It pads on the left with leading zeros.
<i>field</i>	The full width of the field to print in. The default is the number of characters in the output. The field is padded to the left with <i>padchar</i> if necessary.
<i>padchar</i>	The character for padding the field if the field is wider than the number. The default is <code>#\space</code> . The sign character is to be at the beginning of the field, before the padding, rather than just to the left of the number.
@	The number must always appear signed.

Examples:

```
(format t "~&Pi is ~$" (atan 0 -1)) =>
Pi is 3.14
(format t "~&Pi is ~8$" (atan 0 -1)) =>
Pi is 3.14159265
(format t "~&Pi is ~8,2@:$" (atan 0 -1)) =>
Pi is +03.14159265
(format t "~&Pi is ~8,2,20$" (atan 0 -1)) =>
Pi is      03.14159265
(format t "~&Pi is ~8,,20,'x@$" (atan 0 -1)) =>
Pi is xxxxxxxx+3.14159265
```

It uses free format (`~@A`) for very large values of the argument.

`~C` (**character** *arg*) is put in the output, where *arg* is the next element from the *args* of the **format** function. *arg* is treated as a keyboard character and thus can contain extra modifier bits. See the variable `%%kbd-char`, page 6. The modifier bits are printed first, represented as appropriate prefixes: *c-* for Control, *m-* for Meta, *c-m-* for Control plus Meta, *h-* for Hyper, *s-* for Super.

With the colon flag (`~:C`), the names of the modifier bits are spelled out (for example, "Control-Meta-F"), and nonprinting characters are represented by their names (for example, "Return") rather than being output as themselves.

With both colon and at-sign (`~:@C`), the colon-only format is printed, and then if the character requires the SYMBOL or SHIFT shift key(s) to type it, this fact is mentioned (for example, "Symbol-1"). This is the format used for telling the user about a key he or she is expected to press, for instance, in prompt messages.

For all of these formats, if the character is not a keyboard character

but a mouse "character", it is printed as Mouse-, the name of the button, -, and the number of clicks.

With only an at-sign (`~@C`), the character is printed in such a way that the Lisp reader can understand it, using "#/" or "#\".

- `~◇` Takes a character as its argument and prints the name of the character inside a lozenge. The `~C` directive does this with some characters, but `~◇` does it with all of them.
- `~(` Format a string in lowercase. The `~(` directive must be matched by a corresponding `~)` directive. For example:


```
(format t "~(~S~)" 'fs:pathname)
fs:pathname
NIL

(format t "~S" 'fs:pathname)
FS:PATHNAME
NIL
```
- `~%` A carriage return is written to the output. `~n%` outputs *n* carriage returns. No argument is used. Simply putting a carriage return in the control string would work, but `~%` is usually used because it makes the control string look nicer in the Lisp source program.
- `~&` The **:fresh-line** operation is performed on the output stream. Unless the stream knows that it is already at the front of a line, this outputs a carriage return. `~n&` does a **:fresh-line** operation and then outputs *n-1* carriage returns.
- `~|` Outputs a page separator character (`#\page`). `~n|` does this *n* times. With a **:** modifier, if the output stream supports the **:clear-screen** operation this directive clears the screen; otherwise it outputs page separator character(s) as if no **:** modifier were present.
- `--` Outputs a tilde. `~n~` outputs *n* tildes.
- `~?` The next argument in *args* of the **format** function must be a string, and the argument after that must be a list. The string is processed as a **format** control string, with the elements of the list as the corresponding arguments. The processing of the format string containing `~?` resumes when the processing of `~?`'s string is finished. If the **@** modifier is supplied, the next argument in *args* must be a string; it is processed as part of the main format string, as if it were substituted for the `~@?` directive.

Examples:

```
(format nil "~? ~D" "<~A ~D)" ("Myname" 50.) 7) ==> "<Myname 50> 7"
```

```
(format nil "~@? ~D" "<~A ~D>" "Myname" 50. 7) ==> "<Myname 50> 7"
```

~<CR> Tilde immediately followed by a carriage return ignores the carriage return and any whitespace at the beginning of the next line. With a **:**, the whitespace is left in place. With an **@**, the carriage return is left in place. This directive is typically used when a format control string is too long to fit nicely into one line of the program.

~* The next element in the *args* of the **format** function is ignored. **~n*** ignores the next *n* arguments. **~:*** "ignores backwards"; that is, it backs up in the list of arguments so that the argument last processed will be processed again. **~n:*** backs up *n* arguments. When within a **{** construct, the ignoring (in either direction) is relative to the list of arguments being processed by the iteration.

~@* **~n@*** branches to the *n*th argument (0 is the first). **~@*** or **~0@*** goes back to the first argument in the *args* of the **format** function. Directives after a **~n@*** take sequential arguments after the one that is the target of the branch. When within a **{** construct, the branch is relative to the list of arguments being processed by the iteration. This is an "absolute branch". The directive for a relative branch is described elsewhere.

~nG In Zetalisp, "goes to" the *n*th argument. **~0G** goes back to the first argument in the *args* of the **format** function. Directives after this one correspond to the sequence of arguments following the argument that is the target of **~G**. Inside a **{** construct, the "goto" is relative to the list of arguments being processed by the iteration.

This is an "absolute" goto; for a relative goto.

The Common Lisp floating-point format specified by **~G** is not supported.

~P If the next element in the *args* of the **format** function is not 1, a lowercase **s** is output. ("P" is for "plural.") **~:P** does the same thing, after doing a **~:*** to back up one argument; that is, it prints a lowercase **s** if the last argument were not 1. **~@P** outputs "y" if the argument is 1 or "ies" if it is not. **~:@P** does the same thing but backs up first.

~T Spaces over to a given column. **~n,mT** outputs enough spaces to move the cursor to column *n*. If the cursor is already past column *n*, spaces are output to move it to column *n+mk*, for the smallest integer value *k* possible. *n* and *m* default to 1. Without the colon flag, *n* and *m* are in units of characters; with it, they are in units of pixels.

Note: This operation works properly only on streams that support

the `:read-cursorpos` and `:set-cursorpos` stream operations. On other streams, any `~T` operation simply outputs two spaces.

When `format` is creating a string, `~T` works, assuming that the first character in the string is at the left margin.

`~@T` `~@T` outputs a space. `~n@T` outputs n spaces.

`~R` `~R` prints `arg` as a cardinal English number, for example, four. `~:R` prints `arg` as an ordinal number, for example, fourth. `~@R` prints `arg` as a Roman numeral, for example, IV. `~:@R` prints `arg` as an old Roman numeral, for example, IIII.

`~nR` prints `arg` in radix n .

The full form is `~radix,mincol,padchar,commacharR`.

`~radix,nR` uses a column width of n ; spaces are inserted on the left if the number requires fewer than n columns for its digits and sign. If the number does not fit in n columns, additional columns are used as needed.

`~radix,n,mR` uses m instead of the space as the pad character.

The `@` modifier causes the number's sign to be printed always; the default is to print it only if the number is negative.

The `:` modifier causes commas to be printed between groups of three digits; the `commachar` parameter can be used to change the character used as the comma.

`~[str0~;str1~;...~;strn~]`

This is a set of alternative control strings. The alternatives (called *clauses*) are separated by `~;` and the construct is terminated by `~]`. `~]` is undefined elsewhere. `~;` is also used as a separator in justification (`~<`) constructions but is undefined elsewhere.

For example:

```
"~[Siamese ~;Manx ~;Persian ~;Tortoise-Shell ~
~;Tiger ~;Yu-Hsiang ~]kitty"
```

Where `arg` is the next element from the `args` of the `format` function, the `argth` alternative is selected; `0` selects the first. If a prefix parameter is given (that is, `~n[`), then the parameter is used instead of an argument (this is useful only if the parameter is `"#"`). If `arg` is out of range, no alternative is selected. After the selected alternative has been processed, the control string continues after the `~]`.

`~[str0~;str1~;...~;strn~;default~]` has a default case. If the *last* `~;` used to separate clauses is instead `~;`, then the last clause is an "else" clause, which is performed if no other clause is selected. For example:

```
"~[Siamese ~;Manx ~;Persian ~;Tiger ~
~;Yu-Hsiang ~;;Bad ~] kitty"
```

`~[tag00,tag01,...;str0~tag10,tag11,...;str1...~]` allows the clauses to have explicit tags. The parameters to each `~;` are numeric tags for the clause that follows it. That clause is processed that has a tag matching the argument. If `~a1,a2,b1,b2,...;` (note the colon) is used, the following clause is tagged not by single values but by ranges of values *a1* through *a2* (inclusive), *b1* through *b2*, and so on. `~;` with no parameters can be used at the end to denote a default clause. For example:

```
"~[~'+, '-, '*, '//;operator ~'A, 'Z, 'a, 'z;;letter ~
~'0, '9;;digit ~;;other ~]"
```

`~:[false~;true~]` selects the *false* control string if *arg* is **nil**, and selects the *true* control string otherwise.

`~@[true~]` tests the argument. If it is not **nil**, then the argument is not used up, but is the next one to be processed, and the one clause is processed. If it is **nil**, then the argument is used up, and the clause is not processed. For example:

```
(setq prinlevel nil prinlength 5)
(format nil "~@[ PRINLEVEL=~D~]~@[ PRINLENGTH=~D~]"
        prinlevel prinlength)
=> " PRINLENGTH=5"
```

The combination of `~[` and `#` is useful, for example, for dealing with English conventions for printing lists:

```
(setq foo "Items:~#[ none~; ~S~; ~S and ~
~S~;;~@{~#[~1; and~] ~S~^,~}~].")
(format nil foo)
=> "Items: none."
(format nil foo 'foo)
=> "Items: FOO."
(format nil foo 'foo 'bar)
=> "Items: FOO and BAR."
(format nil foo 'foo 'bar 'baz)
=> "Items: FOO, BAR, and BAZ."
(format nil foo 'foo 'bar 'baz 'quux)
=> "Items: FOO, BAR, BAZ, and QUUX."
```

`~{str~}` This is an iteration construct. The corresponding argument of the **format** function should be a list, which is used as a set of arguments as if for a recursive call to **format**. (The terminator `~}` is undefined elsewhere.)

The string *str* is used repeatedly as the control string. Each iteration can absorb as many elements of the list as it likes; if *str* uses up two arguments by itself, two elements of the list are used up each time around the loop.

If, before any iteration step, the list is empty, the iteration is terminated. Also, if a prefix parameter *n* is given, there will be at most *n* repetitions of processing of *str*. Here are some simple examples:

```
(format nil "Here it is:~{ ~S~}." '(a b c))
=> "Here it is: A B C."
(format nil "Pairs of things:~{ <~S,~S>~}." '(a 1 b 2 c 3))
=> "Pairs of things: <A,1> <B,2> <C,3>."
```

`~:{str~}` is similar, but the argument should be a list of sublists. At each repetition step, one sublist is used as the set of arguments for processing *str*; on the next repetition a new sublist is used, whether or not all of the last sublist had been processed. Example:

```
(format nil "Pairs of things:~:~{ <~S,~S>~}."
           '((a 1) (b 2) (c 3)))
=> "Pairs of things: <A,1> <B,2> <C,3>."
```

`~@{str~}` is similar to `~{str~}`, but instead of using one argument that is a list, all the remaining **format** arguments are used as the list of arguments for the iteration. Example:

```
(format nil "Pairs of things:~@~{ <~S,~S>~}."
           'a 1 'b 2 'c 3)
=> "Pairs of things: <A,1> <B,2> <C,3>."
```

`~:@{str~}` combines the features of `~:{str~}` and `~@{str~}`. All the remaining arguments are used, and each must be a list. On each iteration, the next argument is used as a list of arguments to *str*. Example:

```
(format nil "Pairs of things:~:@~{ <~S,~S>~}."
           '(a 1) '(b 2) '(c 3))
=> "Pairs of things: <A,1> <B,2> <C,3>."
```

Terminating the repetition construct with `~:~}` instead of `~}` forces *str* to be processed at least once even if the initial list of arguments is null (however, it will not override an explicit prefix parameter of zero).

If *str* is empty, an argument is used as *str*. It must be a string and precedes any arguments processed by the iteration. As an example, the following are equivalent:

```
(lexpr-funcall #'format stream string args)
(format stream "~1{~:~}" string args)
```

This will use **string** as a formatting string. The `~1{` says it will be processed at most once, and the `~:~}` says it will be processed at least once. Therefore it is processed exactly once, using **args** as the arguments.

As another example, the **format** function itself uses **format-error** (a routine internal to the **format** package) to signal error messages,

which in turn uses **error**, which uses **format** recursively. **format-error** takes a string and arguments, like **format**, but also prints some additional information: if the control string in **ctl-string** actually is a string (it might be a list), it prints the string and a small arrow showing where in the processing of the control string the error occurred. The variable **ctl-index** points one character after the place of the error.

```
(defun format-error (string &rest args)
  (if (stringp ctl-string)
      (ferror nil "~1{~:}~%~VT↓~%~3X/"~A/"~%"
              string args (+ ctl-index 3) ctl-string)
      (ferror nil "~1{~:}" string args)))
```

This first processes the given string and arguments using `~1{~:}`, then tabs a variable amount for printing the down-arrow, then prints the control string between double-quotes. The effect is something like this:

```
(format t "The item is a ~[Foo~;Bar~;Loser~]." 'quux)
>>ERROR: The argument to the FORMAT "~[" command
        must be a number
        ↓
        "The item is a ~[Foo~;Bar~;Loser~]."
...

```

`~<` *mincol, colinc, minpad, padchar*<text> justifies *text* within a field at least *mincol* wide. *text* can be divided into segments with `~;`; the spacing is evenly divided between the text segments. The terminator `~>` is undefined elsewhere.

With no modifiers, the leftmost text segment is left justified in the field, and the rightmost text segment right justified; if there is only one, as a special case, it is right justified.

The `:` modifier causes spacing to be introduced before the first text segment. The `@` modifier causes spacing to be added after the last. *Minpad*, default 0, is the minimum number of *padchar* padding characters (default is the space character) to be output between each segment. If the total width needed to satisfy these constraints is greater than *mincol*, then *mincol* is adjusted upwards in *colinc* increments. *colinc* defaults to 1. *mincol* defaults to 0. For example:

```
(format nil "~10<foo~;bar~>")      => "foo  bar"
(format nil "~10:<foo~;bar~>")     => "  foo bar"
(format nil "~10:@<foo~;bar~>")    => "  foo bar "
(format nil "~10<foobar~>")        => "   foobar"
(format nil "~10:<foobar~>")        => "   foobar"
(format nil "~10@<foobar~>")       => "foobar  "
(format nil "~10:@<foobar~>")      => "  foobar "
(format nil "$~10,,,'*~<3f~>" 2.59023) => "$*****2.59"
```

Note that *text* can include format directives. The last example illustrates how the `~<` directive can be combined with the `~f` directive to provide more advanced control over the formatting of numbers.

Here are some examples of the use of `^^` within a `~<` construct. `^^` eliminates the segment in which it appears and all following segments if there are no more arguments.

```
(format nil "~15<~S~;^^~S~;^^~S~>" 'foo)
=> "          FOO"
(format nil "~15<~S~;^^~S~;^^~S~>" 'foo 'bar)
=> "FOO          BAR"
(format nil "~15<~S~;^^~S~;^^~S~>" 'foo 'bar 'baz)
=> "FOO  BAR  BAZ"
```

If a segment contains a `^^`, and **format** runs out of arguments, it stops there instead of getting an error, and it as well as the rest of the segments are ignored.

If the first clause of a `~<` is terminated with `~;` instead of `~;`, it is used in a special way. All the clauses are processed (subject to `^^`, of course), but the first one is omitted in performing the spacing and padding. When the padded result has been determined, if it will fit on the current line of output, it is output, and the text for the first clause is discarded. If, however, the padded text will not fit on the current line, the text segment for the first clause is output before the padded text. The first clause should contain a carriage return (`~%`). The first clause is always processed, and so any arguments to which it refers are used; the decision is whether to use the resulting segment of text, not whether to process the first clause. If the `~;` has a prefix parameter *n*, the padded text must fit on the current line with *n* character positions to spare to avoid outputting the first clause's text. For example, the following control string can be used to print a list of items separated by commas, without breaking items over line boundaries, and beginning each line with `~;`:

```
"~%;; ~{~<~%;; ~1;; ~S~>~^,~}.~%"
```

The prefix parameter `1` in `~1;` accounts for the width of the comma that will follow the justified item if it is not the last element in the list, or the period if it is. If `~;` has a second prefix parameter, it is used as the width of the line, overriding the natural line width of the output stream. To make the preceding example use a line width of 50, you would write:

```
"~%;; ~{~<~%;; ~1,50;; ~S~>~^,~}.~%"
```

If the second argument is not specified, then **format** sees whether the stream handles the **:size-in-characters** message. If it does, then **format** sends that message and uses the first returned value as the line length in characters. If it does not, **format** uses **95** as the line length.

Rather than using this complicated syntax, you can often call the function **format:print-list**.

^^ This is an escape construct. If there are no more arguments remaining to be processed, then the immediately enclosing **~{** or **~<** construct is terminated. If there is no such enclosing construct, then the entire formatting operation is terminated. In the **~<** case, the formatting is performed, but no more segments are processed before doing the justification. The **^^** should appear only at the beginning of a **~<** clause, because it aborts the entire clause. **^^** can appear anywhere in a **~{** construct.

If a prefix parameter is given, then termination occurs if the parameter is zero. (Hence **^^** is the same as **~#^**.) If two parameters are given, termination occurs if they are equal. If three are given, termination occurs if the second is between the other two in ascending order. Of course, this is useless if all the prefix parameters are constants; at least one of them should be a **#** or a **V** parameter.

If **^^** is used within a **~{** construct, it merely terminates the current iteration step (because in the standard case it tests for remaining arguments of the current step only); the next iteration step commences immediately. To terminate the entire iteration process, use **~:^**.

~> **~>text~<** indents *text* at the cursor position that is current at the time of the **~>**. A **~>** must be terminated with a **~<**, which is undefined elsewhere. **~>** and **~<** can be nested like **~[]** and **~<>**; if they are nested, the indentation of an inner pair is relative to the margin set by the pair containing it. A numeric argument, if supplied, specifies how far to indent. This directive is especially useful in making error messages indent properly. For example:

```
(format t "~&Error: ~>~A~<" "File not found
for FOO.LISP.1")
```

prints

```
Error: File not found
for FOO.LISP.1
```

~Q An escape to arbitrary user-supplied code. *arg* is called as a function; its arguments are the prefix parameters to **~Q**, if any. *args* can be passed to the function by using the **V** prefix parameter. The function can output to **standard-output** and can look at the variables **format:colon-flag** and **format:atsign-flag**, which are **t** or **nil** to reflect the **:** and **@** modifiers on the **~Q**. For example,

```
(format t "~VQ" foo bar)
```

is a fancy way to say


```
(funcall bar foo)
```

and discard the value. Note the reversal of order; the **V** is processed before the **Q**.

~\date

Prints its argument as a date and time, assuming the argument is a universal time. It uses the function **time:print-universal-date**.

```
(format nil "Today is ~\date\" (time:get-universal-time))
=> "Today is Tuesday the fourteenth of May, 1985; 3:07:05 pm"
```

~\time

Prints its argument as a time, assuming the argument is a universal time. It uses the function **time:print-universal-time**.

```
(format nil "Today is ~\time\" (time:get-universal-time))
"Today is 5//14//85 15:08:41"
```

~\datetime

Prints the current time of day. It does not take an argument. It uses the function **time:print-current-time**.

```
(format nil "Today is ~\datetime\")
"Today is 5//14//85 15:19:06"
```

~\time-interval

Prints the length of a time interval. It uses the function **time:print-interval-or-never**.

```
(setq a (time:get-universal-time))
...
(format nil "It is ~\time-interval\ since I set this variable"
        (- (time:get-universal-time) a))
"It is 1 hour 5 minutes 9 seconds since I set this variable"
```

You can use the special form **format:defformat** to define your own directives.

format:defformat *directive* (*arg-type*) *arglist* *body* ... *Special Form*
 Defines a new **format** directive.

directive is a symbol that names the directive. If *directive* is longer than one character, it must be enclosed in backslashes in calls to **format**:

```
(format t "~\foo\" ...)
```

directive is usually in the **format** package; if it is in another package, the user must specify the package in calls to **format**:

```
(format t "~\foo:bar\" ...)
```

format:defformat defines a function to be called when **format** is

called using *directive*. *body* is the body of the function definition. *arg-type* is a keyword that determines the arguments to be passed to the function as *arglist*:

- :no-arg** The directive uses no arguments. The function is passed one argument, a list of parameters to the directive. The value returned by the function is ignored.
- :one-arg** The directive uses one argument. The function is passed two arguments: the argument associated with the directive and a list of parameters to the directive. The value returned by the function is ignored.
- :multi-arg** The directive uses a variable number of arguments. The function is passed two arguments. The first is a list of the first argument associated with the directive and all the remaining arguments to **format**. The second is a list of parameters to the directive. The function should **cdr** down the list of arguments, using as many as it wants, and return the tail of the list so that the remaining arguments can be given to other directives.

The function can examine the values of **format:colon-flag** and **format:atsign-flag**. If **format:colon-flag** is not **nil**, the directive was given a **:** modifier. If **format:atsign-flag** is not **nil**, the directive was given a **@** modifier.

The function should send its output to the stream that is the value of **format:*format-output***.

Here is an example of a **format** directive that takes one argument and prints a number in base 7:

```
(format:deformat format:base-7 (:one-arg) (argument parameters)
  parameters ; ignored
  (let ((base 7))
    (princ argument format:*format-output*)))
```

Now:

```
(format nil "> ~\base-7\ <" 8) => "> 11 <"
```

Note: **format** also allows *control-string* to be a list. If the list is a list of one element, which is a string, the string is simply printed. This is for the use of the **format:outfmt** function. In general, you should use **format:output** if you use lists as arguments.

format:print-list *destination element-format-string list* *Function*
 &optional (*separator-format-string* " , ")
 (*start-line-format-string* " ")
 (*tilde-brace-options* "")

Provides a simpler interface for the specific purpose of printing comma-separated lists with no list element split across two lines.

The *destination* argument tells where to send the output; it can be **t**, **nil**, a string suitable for **string-nconc**, or, as with **format**, a stream.

element-format-string is a **format** control string specifying how to print each element of *list*. It is used as the body of an iteration construction (as in `~{element-format-string}`).

separator-format-string, which defaults to " , " (comma, space), is a string that is placed after each element except the last. **format** control directives are allowed in this string but should not take arguments from the *list*.

start-line, which defaults to three spaces, is a **format** control string that is used as a prefix at the beginning of each line of output except the first.

tilde-brace-options is a string inserted before the opening brace (`{}`) of the iteration construct. It defaults to the null string but allows you to insert a colon or at-sign. The line width of the stream is computed in the same way as with the `~{str}` **format** directive. It is not possible to override the natural line width of the stream.

4.1 The Output Subsystem

The formatting functions associated with the **output** subsystem allow you to do formatted output using Lisp-style control structure. Instead of a directive in a **format** control string, there is one formatting function for each kind of formatted output.

The calling conventions of the formatting functions are all similar. The first argument is usually the datum to be output. The second argument is usually the minimum number of columns to use. The remaining arguments are options — alternating keywords and values.

Options that most functions accept include **:padchar**, followed by a character to use for padding; **:minpad**, followed by the minimum number of padding characters to output after the data; and **:tab-period**, followed by the distance between allowable places to stop padding. For example, if the value of **:tab-period** is 5, the minimum size of the field is 10, and the value of **:minpad** is 2, then a datum that takes 9 characters is padded to 15 characters. The requirement to use at least two

characters of padding means it cannot fit into 10 characters, and the **:tab-period** of 5 means the next allowable stopping place is at 10+5 characters. The default values for **:minpad** and **:tab-period**, if they are not specified, are 0 and 1. The default value for **:padchar** is Space.

The formatting functions always output to **standard-output** and do not require an argument to specify the stream. The macro **format:output** allows you to specify the stream or a string, just as **format** does, and also makes it convenient to concatenate constant and variable output.

format:output *stream string-or-form...* *Macro*

Makes it convenient to intersperse arbitrary output operations with printing of constant strings. **standard-output** is bound to *stream*, and each *string-or-form* is processed in succession from left to right. If it is a string it is printed; otherwise it is a form that is evaluated and the value is discarded. Presumably the forms send output to **standard-output**.

If *stream* is written as **nil**, the output is put into a string that is returned by **format:output**. If *stream* is written as **t**, then the output goes to the prevailing value of **standard-output**. Otherwise *stream* is a form that must evaluate to a stream.

Here is an example:

```
(format:output t "FOO is " (prin1 foo) " now." (terpri))
```

Because **format:output** is a macro, what matters about *stream* is not whether it *evaluates* to **t** or **nil**, but whether it is actually written as **t** or **nil**.

format:outfmt *string-or-form...* *Macro*

Some system functions ask for a **format** control string and arguments, to be printed later. If you wish to generate the output using the formatted output functions, you can use **format:outfmt**, which produces a control argument that eventually makes **format** print the desired output. (This is a list whose one element is a string containing the output.) A call to **format:outfmt** can be used as the first argument to **ferror**. For example:

```
(ferror (format:outfmt "Foo is " (format:onum foo)
" which is too large"))
```

format:onum *number &optional radix minwidth &rest options* *Function*

Outputs *number* in base *radix*, padding to at least *minwidth* columns and obeying the other padding options specified.

radix can be a number or **:roman**, **:english**, or **:ordinal**. The default *radix* is 10. (decimal).

Two special keywords are allowed as options: **:signed** and **:commas**. **:signed** with value **t** means print a sign even if the number is positive.

:commas with value **t** means print a comma every third digit in the customary way. These options are meaningful only with numeric radices.

format:ofloat *number* &optional *n-digits* *force-exponential-notation* *Function*
minwidth &rest *options*

Outputs *number* as a floating-point number using *n-digits* digits. If *force-exponential-notation* is non-**nil**, then an exponent is always used. *minwidth* and *options* are used to control padding as usual.

format:ostream *string* &optional *minwidth* &rest *options* *Function*

Outputs *string*, padding to at least *minwidth* columns if *minwidth* is not **nil**, and obeying the other padding options specified.

Normally the contents of the string are left justified; any padding follows the data. The special option **:right-justify** causes the padding to come before the data. The amount of padding is not affected.

The argument need not be a string. Any Lisp object is allowed, and it is output by **princ**.

format:oprint *object* &optional *minwidth* &rest *options* *Function*

Prints *object*, any Lisp object, padding to at least *minwidth* columns if *minwidth* is not **nil**, and obeying the padding options specified.

Normally the data are left justified; any padding follows. The special option **:right-justify** causes the padding to come before the data. The amount of padding is not affected.

The printing of the object is done with **prin1**.

format:ochar *character* &optional *style* *top-explain* *minwidth* *Function*
&rest *options*

Outputs *character* in one of three styles selected by the *style* argument. *minwidth* and *options* control padding as usual.

If *style* is **:read**, **nil**, or not specified, then the character is printed using **#/** or **#** so that it could be read back in.

If *style* is **:editor**, then the output is in the style of the string "Meta-Rubout".

If *style* is **:brief**, a somewhat more abbreviated style is used in which "c-", "m-", and the like are used to represent "Control" and "Meta", and shorter names for characters are also used when possible. See the section "The Character Set", page 5.

top-explain is useful with the **:editor** and **:brief** styles. It says that any character that has to be typed using the **SYMBOL** key should be followed by an explanation of how to type it. For example: "**α** (**Symbol-shift-A**)".

format:tab *mincol* &rest *options* *Function*

Outputs padding at least until column *mincol*. It is the only formatting function that bases its actions on the actual cursor position rather than the width of what is being output. The padding options **:padchar**, **:minpad**, and **:tab-period** are obeyed. Thus, at least the **:minpad** number of padding characters are output even if that goes past *mincol*, and once past *mincol*, padding can only stop at a multiple of **:tab-period** characters past *mincol*.

In addition, if the **:terpri** option is **t**, then if column *mincol* is passed, **format:tab** starts a new line and indents it to *mincol*.

The **:unit** option specifies the units of horizontal position. The default is to count in units of characters. If **:unit** is specified as **:pixel**, the computation (and the argument *mincol* and the **:minpad** and **:tab-period** options) are in units of pixels.

format:pad (*minwidth* *option...*) *body...* *Macro*

Prints several items in a fixed amount of horizontal space, padding between them to use up any excess space. Each of the *body* forms prints one item. The padding goes between items. The entire **format:pad** always uses at least *minwidth* columns; any unneeded columns are distributed as padding between the items. If that is not enough space, more space is allocated in units controlled by the **:tab-period** option until there is enough space. If it is more than enough, the excess is used as padding.

If the **:minpad** option is specified, then at least that many pad characters must go between each pair of items.

Padding goes only between items. If you want to treat several actual pieces of output as one item, put a **progn** around them. If you want padding before the first item or after the last, as well as between the items, include a dummy item **nil** at the beginning or the end.

If there is only one item, it is right justified. One item followed by **nil** is left justified. One item preceded and followed by **nil** is centered. Therefore, **format:pad** can be used to provide the usual padding options for a function that does not provide them itself.

format:plural *number* *singular* &optional *plural* *Function*

Outputs either the singular or the plural form of a word, depending on the value of *number*. The singular is used if and only if *number* is 1. *singular* specifies the singular form of the word. **string-pluralize** is used to compute the plural, unless *plural* is explicitly specified.

It is often useful for *number* to be a value returned by **format:onum**, which returns its argument. For example:

```
(format:plural (format:onum n-frobs) " frob")
```

prints 1 frob or 2 frobs.

format:breakline *line1 print-if-terpri print-always...* *Macro*

Goes to the next line if there is not enough room for something to be output on the current line. The *print-always* forms print the text that is supposed to fit on the line. *line1* is the column before which the text must end. If it does not end before that column, **format:breakline** moves to the next line and executes the *print-if-terpri* form before doing the *print-always* forms.

Constant strings are allowed as well as forms for *print-if-terpri* and *print-always*. A constant string is just printed.

To go to a new line unconditionally, simply call **terpri**.

Here is an example that prints the elements of a list, separated by commas, breaking lines between elements when necessary.

```
(defun pcl (list line1)
  (do ((l list (cdr l)) ((null l))
      (format:breakline line1 " "
        (princ (car l))
        (and (cdr l) (princ ", "))))))
```

sys:with-indentation (*stream-var relative-indentation*) &body *body* *Macro*

Within the body of **sys:with-indentation**, any output to *stream-var* is preceded by a number of spaces. At every recursion, the additional indentation is specified by *relative-indentation*. The macro does not work this way with the **:item** message used to display mouse-sensitive items; the items appear, but without indentation. (See the section "Interactive Streams and Mouse-sensitive Items" in *Programming the User Interface*.)

```
(defun traced-factorial (n)
  (format t "~%Argument: ~D" n)
  (sys:with-indentation (standard-output 2)
    (let ((value (if (<= n 1)
                    1
                    (* n (traced-factorial (1- n))))))
      (format t "~%Value: ~D" value)
      value)))
```

```
(traced-factorial 5)
```

```
Argument: 5
Argument: 4
Argument: 3
Argument: 2
Argument: 1
Value: 1
Value: 2
Value: 6
Value: 24
Value: 120
120
```


5. Formatting Lisp Code

grindef *function-spec...* *Special Form*

Prints the definitions of one or more functions, with indentation to make the code readable. Certain other "pretty-printing" transformations are performed:

- The **quote** special form is represented with the ' character.
- Displacing macros are printed as the original code rather than the result of macro expansion.
- The code resulting from the backquote (‘) reader macro is represented in terms of ‘.

The subforms to **grindef** are the function specs whose definitions are to be printed; ordinarily, **grindef** is used with a form such as (**grindef** **foo**) to print the definition of **foo**. When one of these subforms is a symbol, if the symbol has a value its value is prettily printed also. Definitions are printed as **defun** special forms, and values are printed as **setq** special forms.

If a function is compiled, **grindef** says so and tries to find its previous interpreted definition by looking on an associated property list. See the function **uncompile** in *Program Development Utilities*. This works only if the function's interpreted definition was once in force; if the definition of the function was simply loaded from a BIN file, **grindef** does not find the interpreted definition and cannot do anything useful.

With no subforms, **grindef** assumes the same arguments as when it was last called.

grind-top-level *obj* &optional *width* (*stream* **standard-output**) *Function*
 (*untyo-p* **nil**) (*displaced* 'si:displaced)
 (*terpri-p* **t**) *notify-fun* *loc*

Pretty-prints *obj* on *stream*, inserting up to *width* characters per line. This is the primitive interface to the pretty-printer. Note that it does not support variable-width fonts. If the *width* argument is supplied, it is how many characters wide the output is to be. If *width* is unsupplied or **nil**, **grind-top-level** tries to determine the "natural width" of the stream by sending a **:size-in-characters** message to the stream and using the first returned value. If the stream does not handle that message, a width of **95** characters is used instead.

The remaining optional arguments activate various features and usually should not be supplied. These options are for internal use by the system, and are documented here only for completeness. If *untyo-p* is **t**, the **:untyo** and **:untyo-mark** operations are used on *stream*, speeding up the algorithm somewhat. *displaced* controls the checking for displacing macros; it is the

symbol that flags a place that has been displaced, or **nil** to disable the feature. If *terpri-p* is **nil**, **grind-top-level** does not advance to a fresh line before printing.

If *notify-fun* is non-**nil**, it is a function of three arguments and is called for each "token" in the pretty-printed output. Tokens can be atoms, open and close parentheses, and reader macro characters such as '. The arguments to *notify-fun* are the token, its "location" (see next paragraph), and **t** if it is an atom or **nil** if it is a character.

loc is the "location" (typically a cons) whose **car** is *obj*. As the grinder recursively descends through the structure being printed, it keeps track of the location where each thing came from, for the benefit of the *notify-fun*, if any. This makes it possible for a program to correlate the printed output with the list structure. The "location" of a close parenthesis is **t**, because close parentheses have no associated location.

6. The Input Editor Program Interface

6.1 How the Input Editor Works

The input editor is a feature of all interactive streams, that is, streams that connect to terminals. Its purpose is to let you edit minor mistakes in typein. At the same time, it is not supposed to get in the way; Lisp is to see the input as soon as you have typed a syntactically complete form. The definition of "syntactically complete form" depends on the function that is reading from the stream; for **read**, it is a Lisp expression. This section describes the general protocol used for communication between the input editor and reading functions such as **read** and **readline**.

By *reading function* we mean a function that reads a number of characters from a stream and translates them into an object. For example, **read** reads a Lisp expression and returns an object. **readline** reads a line of characters and returns a string as its first value. Reading functions do not include the more primitive **:tyi** and **:any-tyi** stream operations, which take and return one character or blip from the stream.

The tricky thing about the input editor is the need for it to figure out when you are all done. The idea of an input editor is that as you type in characters, the input editor saves them up in an *input buffer* so that if you change your mind, you can edit them and replace them with different characters. However, at some point the input editor has to decide that the time has come to stop putting characters into the input buffer and let the reading function start processing the characters. This is called "activating".

The right time to activate depends on the function calling the input editor, and determining it may be very complicated. If the function is **read**, figuring out when one Lisp expression has been typed requires knowledge of all the various printed representations, what all currently defined reader macros do, and so on. The input editor should not have to know how to parse the characters in the input buffer to figure out what the caller is reading and when to activate; only the caller should have to know this. The input editor interface is organized so that the calling function can do all the parsing, while the input editor does all the handling of editing commands, and the two are kept completely separate.

Following is a summary of how the input editor works. The input editor used to be called the rubout handler, and some operations and variables still have "rubout-handler" in their names.

When a reading function is called to read from a stream that supports the **:input-editor** operation, that function "enters" the input editor. It then goes ahead **:tyi**'ing characters from the stream. Because control is inside the input editor, the stream echoes these characters so the user can see the input. (Normally echoing is

considered to be a higher-level function outside of the province of streams, but when the higher-level function tells the stream to enter the input editor it is also handing it the responsibility for echoing). The input editor is also saving all these characters in the input buffer, for reasons disclosed in the following paragraph. When the reading function decides it has enough input, it returns and control "leaves" the input editor. That was the easy case.

If you press RUBOUT or a keystroke that represents another editing command, the input editor processes the command and lets you insert characters before the last one in the line. The input editor modifies the input buffer and the screen accordingly. Then, when you type the next nonediting character at the end of the line, a **throw** is done, out of all recursive levels of **read**, reader macros, and so forth, back to the point where the input editor was entered. Now the **read** is tried over again, rereading all the characters you had typed and not rubbed out, but not echoing them this time. When the saved characters have been exhausted, additional input is read from you in the usual fashion.

The input editor has options that can cause the **throw** to occur at other times as well. With the **:activation** option, when you type an activation character a **throw** occurs, a rescan is done if necessary, and a final blip is returned to the reading function. With the **:preemptable** and **:command** options, a blip or special character in the input stream causes control to be returned from the input editor immediately, without a rescan. These options let you process mouse clicks or special keystroke commands as soon as they are read.

The effect of all this is a complete separation of the functions of input editing and parsing, while at the same time mingling the execution of these two functions in such a way that input is always "activated" at just the right time. It does mean that the parsing function (in the usual case, **read** and all macro-character definitions) must be prepared to be thrown through at any time and should not have nontrivial side-effects, since it may be called multiple times.

If an error occurs while inside the input editor, the error message is printed and then additional characters are read. When you press RUBOUT, it rubs out the error message as well as the last character. You can then proceed to type the corrected expression; the input is reparsed from the beginning in the usual fashion.

6.2 Invoking the Input Editor

The variable **rubout-handler** indicates the current state of input editing. This variable is not **nil** if the current process is already inside the input editor.

rubout-handler

Variable

Indicates the status of input editing within a process.

This variable is used internally by the **:input-editor** method and the input

editor. It should not be necessary for user programs to examine its value since the **with-input-editing** special form is provided for this purpose.

The possible values for this variable are:

<i>Value</i>	<i>Meaning</i>
nil	The process is outside the input editor.
:read	The process is inside the :input-editor method.
:tyi	The process is inside the editing portion of the :tyi method.

The input editor is invoked on a stream when the stream receives an **:input-editor** message. The **:input-editor** and **:tyi** methods of **si:interactive-stream** contain the code of the input editor. The **:input-editor** method initializes the input editor, establishes its **catch**, and then calls back to the reading function with **rubout-handler** bound to **:read**. When the reading function sends the **:tyi** or **:any-tyi** message, input is taken from the input buffer. If no input is available, the editing or **:tyi** portion of the input editor is invoked, and **rubout-handler** is bound to **:tyi**.

The first argument to the **:input-editor** message is the function that the input editor should call to do the reading, and the rest of the arguments are passed to that function. If the reading function returns normally, the values returned by the **:input-editor** message are just those returned by the reading function. If the input editor returns by throwing out of the reading function, the return values depend on which option caused the input editor to throw: See the option **:full-rubout**, page 59. See the option **:preemptable**, page 63. See the option **:command**, page 64.

The input editor can take a series of options. These are specified dynamically by the special forms **with-input-editing-options** and **with-input-editing-options-if**. For a description of the options: See the section "Input Editor Options", page 59.

with-input-editing-options *options &body body* *Special Form*
 Specifies input editing options and executes *body* with those options in effect. The scope of the option specifications is dynamic.

options is a list of input editor option specifications. Each element is a list whose car is an option-name specification and whose cdr is a list of forms to be evaluated to yield "arguments" for the option. The option-name specification is a keyword symbol or a list whose car is a keyword symbol. The symbol is the name of the option.

If the option-name specification is a list and if the symbol **:override** is an element of the cdr of the list, this option specification overrides any higher-level specifications for this option. Otherwise, the specification for each option that is dynamically outermost (that is, the specification from the highest-level caller) is in effect during the execution of *body*.

with-input-editing-options returns whatever values *body* returns.

In the following example, the user is prompted for a Lisp expression. Two input editor options are specified. The first says that the caller is also willing to receive mouse or menu blips. The second specifies a prompt.

```
(with-input-editing-options ((:preemptable :blip)
                             (:prompt "Form: "))
  (read))
```

In the following example, the user is prompted for a line of text. The text may be activated by any of the characters RETURN, END, or TRIANGLE. This might be useful if activating with TRIANGLE meant something different from activating with RETURN. This example also demonstrates the use of **:override** to make this **:activation** specification override any higher-level **:activation** specifications.

```
(with-input-editing-options
  (((:activation :override) 'memq '#\return #\end #\triangle)))
  (prompt-and-read :string "Name: "))
```

For a list of input editor options: See the section "Input Editor Options", page 59. See the special form **with-input-editing-options-if**, page 56.

with-input-editing-options-if *cond options &body body* *Special Form*

Executes *body*, possibly with specified input editing options in effect. The scope of the option specifications is dynamic.

cond is a form to be evaluated at run-time. If *cond* returns non-**nil**, the specified input editor options are in effect during the execution of *body*.

options is a list of input editor option specifications. Each element is a list whose car is an option-name specification and whose cdr is a list of forms to be evaluated to yield "arguments" for the option. The option-name specification is a keyword symbol or a list whose car is a keyword symbol. The symbol is the name of the option.

If the option-name specification is a list and if the symbol **:override** is an element of the cdr of the list, this option specification overrides any higher-level specifications for this option. Otherwise, the specification for each option that is dynamically outermost (that is, the specification from the highest-level caller) is in effect during the execution of *body*.

with-input-editing-options-if returns whatever values *body* returns.

For a list of input editor options: See the section "Input Editor Options", page 59. See the special form **with-input-editing-options**, page 55.

This example illustrates the use of the **:command**, **:preemptable**, and **:prompt** input editor options. It is a simple command loop that reads different kinds of commands -- typed Lisp expressions, single-keystroke commands, and mouse clicks.

The Lisp expressions are read using the **read-or-end** function. You can provide four kinds of input:

<i>Input</i>	<i>Action</i>
END	Exit the command loop
Lisp form	Print form on next line
Mouse click	Display type of click and mouse coordinates
Single-key command	Display keystroke

The predicate for detecting a single-keystroke command simply checks for the Super bit. In a more complex program, it might look up the character in a command table.

```
(defun command-char-p (c) (char-bit c :super))

(defun command-loop ()
  (loop
    do (multiple-value-bind (value flag)
        (with-input-editing-options
          ( (:command 'command-char-p)
            (:preemptable :blip)
            (:prompt "Command loop input: ")
            (read-or-end)))
      (selectq flag
        (:end
         (format t "Done")
         (return t))
        (:blip
         (selectq (car value)
           (:mouse-button
            (destructuring-bind (click nil x y) (cdr value)
              (format t "~C click at ~D, ~D" click x y)))
           (otherwise (format t "Random blip -- ~S" value))))
        (:command
         (format t "Execute ~:C command" (second value)))
        (otherwise
         (format t "~&Value is ~S" value))))))
```

To write a reading function that invokes the input editor, you should use the **with-input-editing** special form instead of sending the **:input-editor** message directly. Such functions as **read** and **readline** use this special form to provide input editing.

with-input-editing (&optional *stream keyword*) &body *body* *Special Form*

Provides a convenient way of invoking the input editor for use by a reading function. It establishes a context in which input editing should be provided. Use **with-input-editing** instead of sending an **:input-editor** message directly.

Both "arguments" are optional. *stream* is the stream from which characters are read; if *stream* is not provided or is **nil**, **standard-input** is used.

keyword determines the activation characters for the input editor:

<i>Value</i>	<i>Activation characters</i>
nil	None (unless specified at a higher level). This is the default.
:end-activation	#\end
:line-activation	#\end , #\return , and #\line
:line	#\end , #\return , and #\line . In addition, a Newline is echoed after the reading function returns.

To supply other input editor options: See the special form **with-input-editing-options**, page 55. See the special form **with-input-editing-options-if**, page 56.

with-input-editing defines an internal lexical closure with *body* as its body. When the **with-input-editing** form is evaluated from outside the input editor, the stream is sent an **:input-editor** message if it handles it. The argument to the **:input-editor** message is the lexical closure, except that if the **:line** keyword is supplied, **with-input-editing** also arranges to echo a Newline after the lexical closure returns. If the **with-input-editing** form is evaluated from inside the input editor or if the stream does not handle the **:input-editor** message, the lexical closure is called instead.

with-input-editing returns whatever values *body* returns.

The following example defines a simple sentence parser.

```

(defun read-sentence (&optional (stream standard-input))
  (with-input-editing-options ([:prompt "Type a sentence: "])
    (with-input-editing (stream)
      (loop named sentence
        with sentence = nil
        for word = (make-array 20. :type art-string :fill-pointer 0)
        do (loop for char = (send stream :tyi)
              do
                (cond ((memq char '(#\space #\return #\. #/? #/,))
                      (if (not (equal word ""))
                          (push word sentence))
                      (selectq char
                        ((#\space #\return #/,)
                         (return))
                        (#\.
                         (push :period sentence)
                         (return-from sentence (nreverse sentence)))
                        (#\?
                         (push :question-mark sentence)
                         (return-from sentence (nreverse sentence))))))
              (t (array-push-extend word char))))))

```

6.3 Input Editor Options

The input editor can take a series of options, specified by the special forms **with-input-editing-options** and **with-input-editing-options-if**. Following are descriptions of the options.

- :full-rubout** *token* *Option*
 If the user rubs out all the characters that were typed, control is returned from the input editor immediately. Two values are returned: **nil** and *token*. If the user does not rub out all the characters, the input editor propagates multiple values back from the function that it calls, as usual. In the absence of this option, the input editor simply waits for more characters to be typed and ignores any additional rubouts.
- :pass-through** *&rest characters* *Option*
 The characters in *characters* are not to be treated as special by the input editor. This option is used to pass format effectors (such as **HELP** or **CLEAR INPUT**) through to the reading function instead of interpreting them as input editor commands. **:pass-through** is allowed only for characters with no modifier bits set, that is, for character codes 0 through 377 (octal). For characters that have modifier bits set and must be visible to the reading function, use **:do-not-echo** or **:activation**.

:prompt &rest *prompt-option* *Option*

When it is time for the user to be prompted, the input editor displays *prompt-option*. *prompt-option* can have one element, which can be **nil**, a string, a function, or a symbol other than **nil**; or it can have more than one element: See the section "Displaying Prompts in the Input Editor", page 65.

The difference between **:prompt** and **:reprompt** is that the latter does not display the prompt when the input editor is first entered, but only when the input is redisplayed (for example, after a screen clear). If both options are specified, **:reprompt** overrides **:prompt** except when the input editor is first entered.

:reprompt &rest *prompt-option* *Option*

When it is time for the user to be reprompted, the input editor displays *prompt-option*. *prompt-option* can have one element, which can be **nil**, a string, a function, or a symbol other than **nil**; or it can have more than one element: See the section "Displaying Prompts in the Input Editor", page 65.

Unlike **:prompt**, **:reprompt** displays the prompt only when input is redisplayed (for example, after a screen clear), not when the input editor is first entered. If both **:prompt** and **:reprompt** are specified, **:reprompt** overrides **:prompt** except when the input editor is first entered.

:complete-help &rest *help-option* *Option*

When the user presses HELP, the input editor types out a message determined by *help-option*. None of the standard input editor help is displayed. If a **:brief-help** option has been specified, it overrides **:complete-help**. **:complete-help** overrides **:merged-help** and **:partial-help**.

help-option can have one element, which can be a string, a function, or a symbol; or it can have more than one element. For an explanation: See the section "Displaying Help Messages in the Input Editor", page 66.

This option is intended for programs that supply their own input editor help messages.

:partial-help &rest *help-option* *Option*

When the user presses HELP, the input editor first types out a message determined by *help-option*. It then types out a message describing how to invoke input editor commands and other information about the stream. If a **:brief-help**, **:complete-help**, or **:merged-help** option has been specified, it overrides **:partial-help**.

help-option can have one element, which can be a string, a function, or a symbol; or it can have more than one element. For an explanation: See the section "Displaying Help Messages in the Input Editor", page 66.

This option is intended for use when inexperienced users might be typing to

the input editor. Often *help-option* gives some information about the program to which the user is typing and what the user can do to exit from it.

:merged-help *function &rest arguments* *Option*

When the user presses HELP, the input editor types out a message determined by the arguments. *function* is a function that takes at least two arguments. The input editor calls the function to print the help message. The first argument is the stream. The second argument is a continuation (a list) to print a standard message describing how to invoke input editor commands and other information about the stream. When the function wants to print this message, it should apply the car of the continuation to the cdr. If any *arguments* are supplied, they are the remaining arguments to the function.

If a **:brief-help** or **:complete-help** option has been specified, it overrides **:merged-help**. **:merged-help** overrides **:partial-help**.

This option is intended for programs that want to decide when and where to display their own help messages and the standard help message.

:brief-help *&rest help-option* *Option*

When the user presses HELP, the input editor displays a message determined by *help-option* on the same line as the typein. The message is displayed in the default typeout font, and none of the usual conventions about input editor typeout apply. **:brief-help** overrides **:complete-help**, **:merged-help**, and **:partial-help**.

help-option can have one element, which can be a string, a function, or a symbol; or it can have more than one element. For an explanation: See the section "Displaying Help Messages in the Input Editor", page 66.

This option is intended for programs like **fquery** that need to supply only a brief help message, usually about expected typein.

:initial-input *string &optional begin end cursor-position* *Option*

When the input editor is entered, *string* is inserted into the input buffer as if the user had typed it. The user can edit the string before activating. *begin* and *end* are indices into *string* and mark the portion of the string to be copied into the input buffer. *begin* defaults to 0; *end* defaults to **(array-active-length string)**. *cursor-position* is an index into the string where the cursor should initially be placed. The default is to place the cursor at the end of the portion of the string copied into the input buffer. *string* can be **nil**, which is the same as not specifying the option.

In the following example, the user is prompted for a line of text. The input buffer initially contains the name of the user, and the cursor is placed at the beginning of the input buffer.

```
(with-input-editing-options
  ((:initial-input fs:user-personal-name nil nil 0))
  (prompt-and-read :string "Full name: "))
```

Placing a string in the input buffer is one style of input defaulting. Another style leaves the input buffer empty but allows a default to be yanked with `c-m-Y`. See the option `:input-history-default`, page 62.

:input-history-default *string* *Option*

Specifies *string* as the default to be yanked by `c-m-Y`. *string* is temporarily placed at the head of the input history. If the user types `c-m-Y m-Y`, the true first element of the input history is yanked. `c-m-0 c-m-Y` shows *string* at the head of the input history, and the entries in the input history are shifted down by one.

In the following example, the user is prompted for a line of text. The input buffer is initially empty, but the `c-m-Y` command yanks a default, which is the name of the user.

```
(with-input-editing-options
  ((:input-history-default fs:user-personal-name))
  (prompt-and-read :string "Full name: "))
```

This option is used by the `:pathname` option for `prompt-and-read`.

:blip-handler *function* *Option*

Specifies a function to handle blips received while inside the input editor. *function* must be a function of two arguments. The first argument is the blip; the second argument is the stream that received the blip. The handler is invoked when the input editor receives a blip. If the handler returns non-`nil`, no further action is taken. If it returns `nil` and a `:preemptable` option is in effect, the actions specified by that option are taken. Otherwise, the default blip handler is invoked.

In the following example, the user is prompted for a line of text. While entering this text, the user may also click the left or middle mouse buttons. If the left mouse button is clicked, the coordinates of the mouse with respect to the window are inserted into the input buffer. If the middle button is clicked, the name of the window is inserted.

```
(defun example-blip-handler (blip ignore)
  (destructuring-bind (type click window x y) blip
    (and (eq type :mouse-button)
      (selectq click
        (#\mouse-l-1
          (si:ie-insert-string (format nil " ~D ~D" x y))
          t)
        (#\mouse-m-1
          (si:ie-insert-string (format nil " ~A" window))
          t))))))
```

```
(with-input-editing-options ((:blip-handler 'example-blip-handler))
 (prompt-and-read :string "Blip handler test: "))
```

si:ie-insert-string is an internal function for inserting a string into the input buffer. Since the language for writing input editor commands has not been formalized, this example might not work in a later release.

:do-not-echo *&rest characters* *Option*

The characters in *characters* are interpreted as activation characters and are not echoed. The comparison is done with **char=**, not **char-equal**, so that the control and meta bits are not masked off. The characters are not inserted into the input buffer and are not interpreted as input editor commands. When one of these characters is typed, the final **:tyi** value returned is the character, not a blip.

This option exists only for compatibility with earlier releases. New programs should use the **:activation** option.

:activation *function &rest arguments* *Option*

For each character typed, the input editor invokes *function* with the character as the first argument and *arguments* as the remaining arguments. If the function returns **nil**, the input editor processes the character as it normally would. Otherwise, the cursor is moved to the end of the input buffer, a rescan of the input is forced (if one is pending), and the blip (**:activation character numeric-arg**) is returned by the final sending of the **:any-tyi** message to the stream. Activation characters are not inserted into the input buffer, nor are they echoed by the input editor. It is the responsibility of the reading function to do any echoing. For instance, **readline**, not the input editor, types a Newline at the end of the input buffer when RETURN, END, or LINE is pressed.

:preemptable *token* *Option*

A blip in the input stream causes control to be returned from the input editor immediately. Two values are returned: the blip and *token*, which is usually a keyword symbol. Any unscanned input typed before the blip remains in the input buffer, available to the next read operation from the stream.

:no-input-save *Option*

The input editor does not save the scanned contents of the input buffer on the input history when returning from the reading function. This is intended for use by functions such as **fquery** that use the input editor to ask simple questions whose responses are not worth saving. **yes-or-no-p** uses **:no-input-save** by default.

:command *function &rest arguments* *Option*

This option is used to implement nonediting single-keystroke commands. For each character typed, the input editor invokes *function* with the character as the first argument and *arguments* as the remaining arguments. If the function returns **nil**, the input editor processes the character as it normally would. Otherwise, control is returned from the input editor immediately. Two values are returned: a blip of the form (**:command** *character numeric-arg*) and the keyword **:command**. Any unscanned input typed before the command character remains in the input buffer, available to the next read operation from the stream.

:editor-command *&rest command-alist* *Option*

This option lets you specify your own input editor editing commands. Each element of *command-alist* is a cons whose car is a character and whose cdr is a symbol or a list. If the cdr is a symbol, it is a function to be called with no arguments when the user types the associated character. If the cdr is a list, the car of the list is a function to be applied to the cdr of the list when the user types the associated character. The function can examine the internal special variables that describe the state of the input editor.

If **:editor-command** specifies a command that is invoked by the same character as one of the standard input editor editing commands, the command specified by **:editor-command** overrides the standard command.

:input-wait *&optional whostate function &rest arguments* *Option*

When the input editor waits for input, it sends the stream an **:input-wait** message with the arguments to the **:input-wait** option as arguments. In addition, unless the **:suppress-notifications** option has been specified, **:input-wait** returns when a notification is received. See the message **:input-wait**, page 13.

:input-wait-handler *function &rest arguments* *Option*

When the input editor is waiting for input it sends the stream an **:input-wait** message. After **:input-wait** returns, the input editor applies *function* to *arguments*. The input editor does not process the input or display the notification until *function* returns.

:suppress-notifications *flag* *Option*

If *flag* is not **nil**, notifications received while in the input editor are ignored.

:notification-handler *function &rest arguments* *Option*

If a notification is received while in the input editor, *function* is called to handle it. *function* should take at least one argument, the notification (as returned by the **:receive-notification** message to the stream). *arguments* are the remaining arguments to *function*. *function* can do anything it wants with the notification. To display the notification, *function* would usually call **sys:display-notification**.

If this option is not specified, notifications appear one after the other using **:insert-style** typeout.

Following are two simple examples of notification handlers. The first handler assumes that you want each notification to overwrite the previous one. The second handler assumes that you want them to appear one after another. ***window*** should be bound to a window and ***stream*** to a stream where you want the notifications to appear.

```
(defun my-notification-handler-1 (notification)
  (send *window* :clear-window)
  (sys:display-notification *window* notification :window))

(defun my-notification-handler-2 (notification)
  (sys:display-notification *stream* notification :stream))
```

6.4 Displaying Prompts in the Input Editor

The input editor options **:prompt** and **:reprompt** and the functions **readline-no-echo** and **sys:read-character** take *prompt* arguments that let you specify an input editor prompt. *prompt* can be **nil**, a string, a function, a symbol other than **nil**, or a list (for the input editor options, the list is an **&rest** argument):

- nil** No prompt is displayed.
- string A **format** control string to be passed to **format** with one argument, the stream on which the prompt is displayed.
- function or symbol other than **nil** A function to display the prompt. The function should take two arguments: the first is the stream on which the prompt is displayed, and the second is a keyword that indicates the origin of the function call.
- list If the first element is **nil**, no prompt is displayed. If the first element is a string, it is a **format** control string to be passed to **format** with the remaining elements of the list as arguments. If the first element is a function or a symbol other than **nil**, it is a function to display the prompt. The first argument to the function is the stream on which the prompt is displayed. The second argument is a keyword that indicates the origin of the function call. The remaining arguments are the remaining elements of the list.

When a function is called to display the prompt, the second argument to the function is a keyword that indicates the origin of the function call:

<i>Keyword</i>	<i>Function called from</i>
:prompt	:input-editor method of si:interactive-stream , when the input editor is entered
:restore	:restore-input-buffer method of si:interactive-stream
:finish-typeout	:finish-typeout method of si:interactive-stream
:refresh	Body of the input editor, when the user presses REFRESH
:erase-typeout	Body of the input editor, when the user presses PAGE

6.5 Displaying Help Messages in the Input Editor

The input editor options **:brief-help**, **:partial-help**, and **:complete-help** and the functions **readline-no-echo** and **sys:read-character** take *help* arguments that let you specify input editor help messages. *help* can be a string, a function, a symbol, or a list (for the input editor options, the list is an **&rest** argument):

string	A format control string to be passed to format with one argument, the stream on which the help message is displayed.
function or symbol	A function to display the help message. The function should take one argument, the stream on which the help message is displayed.
list	If the first element is a string, it is a format control string to be passed to format with the remaining elements of the list as arguments. If the first element is a function or a symbol, it is a function to display the help message. The first argument to the function is the stream on which the help message is displayed, and the remaining arguments are the remaining elements of the list.

6.6 Examples of Use of the Input Editor

This series of examples shows several different ways of using the input editor, gradually increasing in complexity. The examples are also available in the file `sys: examples; interaction.lisp`.

We refer to functions whose names begin with "read-" as "reading functions" or "readers", since they read individual characters and construct a Lisp object as a returned value. Examples of readers the Lisp system provides are **read**, **readline**, and **read-delimited-string**. **read** returns Lisp objects of many types. **readline** and **read-delimited-string** return strings.

read-two-lines-1 reads two lines of input from the console. You type each line in its own editing context. After you enter the first line by pressing RETURN, LINE, or

END, you can no longer rub out or otherwise edit any of the characters in the first line. You can type and edit only the second line at that point.

```
(defun read-two-lines-1 () (list (readline) (readline)))
```

read-two-lines-2 lets you edit both lines in a single context by using the **with-input-editing** special form. Even after entering the first line you can edit it. For example, the `m-<` input editor command moves the cursor to the first character of the first line. **read-two-lines-2** also adds a stream parameter so that you can read from different streams without having to bind **standard-input**. You can also use this function for reading from noninteractive streams, such as file streams.

```
(defun read-two-lines-2 (&optional (stream standard-input))
  (with-input-editing (stream) (list (readline stream) (readline stream))))
```

read-two-lines-3 demonstrates the use of the **:prompt** input editor option and the **:end-activation** option for **with-input-editing**. When you invoke this function on an interactive stream you receive a prompt. This prompt is redisplayed if timeout to the stream occurs. This might happen if you press HELP or the window receives a notification.

The **:end-activation** option defines `#\end` as an activation character. This lets you activate previous input to **read-two-lines-3**, after yanking and editing it, by pressing END. The **:prompt** and **:end-activation** options have no effect on the behavior of the function for noninteractive streams.

```
(defun read-two-lines-3 (&optional (stream standard-input))
  (with-input-editing-options ([:prompt "Type two lines: "])
    (with-input-editing (stream :end-activation)
      (list (readline stream) (readline stream)))))
```

read-n-lines is like **read-two-lines** except that you specify the number of lines to be read using the **n-lines** argument. It also uses a prompt function instead of a string to generate the prompt.

```
(defun read-n-lines-prompt (stream ignore n-lines)
  (format stream "Type ~R line~:P::~%" n-lines))

(defun read-n-lines (n-lines &optional (stream standard-input))
  (with-input-editing-options ([:prompt 'read-n-lines-prompt n-lines])
    (with-input-editing (stream :end-activation)
      (loop repeat n-lines collect (readline stream)))))
```

Next is an example of a simple sentence parser. It builds a list of strings and symbols that represent the words and punctuation marks of the sentence. A sentence may be any number of lines long. It is delimited by a period or a question mark. Words are delimited by a space, newline, or punctuation mark. This is also an example of a reading function written entirely in terms of **:tyi** as the primitive input operation.

```

(defun read-sentence-1 (&optional (stream standard-input))
  (with-input-editing-options ([:prompt "Type a sentence: "])
    (with-input-editing (stream)
      (loop named sentence
        with sentence = nil
        for word = (make-array 20. :type art-string :fill-pointer 0)
        do (loop for char = (send stream :tyi)
          do
            (cond ((memq char '(#\space #\return #\. #/? #/,))
              (if (not (equal word ""))
                (push word sentence))
              (selectq char
                ((#\space #\return #/,)
                 (return))
                (#\.
                 (push :period sentence)
                 (return-from sentence (nreverse sentence)))
                (#/?
                 (push :question-mark sentence)
                 (return-from sentence (nreverse sentence))))))
            (t (array-push-extend word char))))))

```

Following is a different sentence parser that calls **read-delimited-string** to accumulate characters into a string. It uses the **:end-activation** option for **with-input-editing** so that previous input to **read-sentence-2** can be yanked, edited, and activated using the END key. When it detects incorrect uses of punctuation, it calls **sys:parse-error** to signal an error caught by the input editor.

```

(defun read-sentence-2 (&optional (stream standard-input))
  (with-input-editing-options ([:prompt "Type a sentence: "])
    (with-input-editing (stream :end-activation)
      (loop with sentence = nil
        do (multiple-value-bind (word nil delimiter)
            (read-delimited-string
              '(#\space #\return #/. #/? #/, #/: #/;) stream)
          (if (not (equal word ""))
              (push word sentence)
              (cond ((memq delimiter '(#\space #\return)))
                    ((null sentence)
                     (if (eq delimiter #\end)
                         (return nil)
                         (sys:parse-ferror
                          "The punctuation mark /"~C/" occurred at the ~
                          beginning of the sentence."
                          delimiter)))
                    ((symbolp (car sentence))
                     (sys:parse-ferror
                      "The punctuation mark /"~C/" was typed after a ~@^."
                      delimiter (car sentence)))
                    (t (selectq delimiter
                       (#/,
                        (push ':comma sentence))
                       (#/:
                        (push ':colon sentence))
                       (#/;
                        (push ':semicolon sentence))
                       (#/.,
                        (push ':period sentence)
                        (return (nreverse sentence)))
                       (#/?
                        (push ':question-mark sentence)
                        (return (nreverse sentence))))))))))

```

Sometimes an error in parsing is detected not by the function that invokes the input editor, but by some function that it calls. In the next example, **read-time** invokes **time:parse-universal-time** to do its parsing. If we did not use the **condition-case** form in **read-time**, we would enter the Debugger when **time:parse-universal-time** encountered incorrect input. The **condition-case** form encapsulates the original error in one of flavor **sys:parse-ferror** so that the input editor catches it. Alternately, we could define **time:parse-error** to be a subflavor of **sys:parse-error**.

```
(defun read-time (&optional (stream standard-input))
  (with-input-editing (stream :line)
    (let ((string (readline-or-nil stream)))
      (when string
        (condition-case (error)
          (time:parse-universal-time string)
          (time:parse-error
           (sys:parse-ferror "~A" error))))))))
```

6.7 Input Editor Messages to Interactive Streams

:input-editor *read-function* &rest *read-args* of **si:interactive-stream** *Method*

Apply *read-function* to *read-args* after invoking the input editor. For more information: See the section "The Input Editor Program Interface", page 53.

Normally a program does not send this message itself; it uses the special form **with-input-editing**. See the special form **with-input-editing**, page 57.

:start-typeout *type* &optional *spacing* of **si:interactive-stream** *Method*

Informs the input editor that typeout to the window will follow. The word "typeout" is used in the name of this message because this is very similar to typeout in the editor, even though typeout windows are not actually used. *type* can be one of the following keywords:

<i>Keyword</i>	<i>Action</i>
:insert	Typeout is inserted before the current input, as is done with notifications or input editor documentation.
:overwrite	Like :insert , but the next time :insert or :overwrite typeout is performed, this typeout is overwritten.
:append	Typeout appears after the current input, which remains visible before the typeout. This is the style used by break .
:temporary	Typeout appears after the current input and is erased after the user types a character.
:clear-window	The window is cleared, and typeout appears at the top.

spacing can be one of the following keywords:

<i>Keyword</i>	<i>Action</i>
:none	No spacing before typeout.
:fresh-line	Typeout begins at the beginning of a line.

:blank-line A blank line precedes typeout.

If *spacing* is not specified, a default that depends on *type* is computed.

si:*typeout-default* *Variable*

Controls the style of typeout performed by the input editor. Permissible values are the keywords acceptable as the *type* argument to the **:start-typeout** method of **si:interactive-stream**. These are **:insert**, **:overwrite**, **:append**, **:temporary**, and **:clear-window**. The default value is **:overwrite**.

:finish-typeout *&optional spacing erase?* of *Method*
si:interactive-stream

Completes typeout to the window and causes the input buffer to be refreshed. In the case of **:temporary** typeout, the *erase?* parameter is used to indicate whether or not the typeout overwrote part of the current input by wrapping around the screen. It is the responsibility of the program doing the typeout to keep track of how much is output.

spacing can be one of the following keywords:

<i>Keyword</i>	<i>Action</i>
:none	No spacing before typeout.
:fresh-line	Typeout begins at the beginning of a line.
:blank-line	A blank line precedes typeout.

If *spacing* is not specified, a default that depends on the *type* argument to the **:start-typeout** method is computed.

:rescanning-p of **si:interactive-stream** *Method*

This message can be sent by a read function that uses the input editor to determine whether the next character returned by **:tyi** will come from the input buffer or from the keyboard. If **t** is returned, the input is being rescanned and the next character will come from the input buffer. If **nil** is returned, the next character will come from the keyboard.

:force-rescan of **si:interactive-stream** *Method*

This message can be sent by a read function that uses the input editor to force a rescan of the current input. Before this message is sent, usually some global state has changed and the contents of the input buffer are interpreted differently.

:replace-input *n-chars string &optional (begin 0) end (rescan-mode* *Method*
:ignore) of **si:interactive-stream**

This message can be sent by a read function that uses the input editor to provide completion of the current input.

n-chars specifies the number of characters to be removed from the end of the input buffer and erased from the screen. It can be an integer, a string, or **nil**:

integer	Remove <i>n-chars</i> characters from immediately before the scan pointer
string	Remove as many characters as the string contains
nil	Remove characters from the beginning of the input buffer to the scan pointer

The substring of *string* determined by *begin* and *end* is then displayed on the screen. *end* defaults to (**string-length** *string*). The scan pointer is left after the string, and a rescan does not take place. If a rescan takes place at some later time, the characters in *string* are seen as input.

rescan-mode specifies what action to take if the **:replace-input** message is sent when the scan pointer is not at the end of the input buffer:

:ignore	Don't perform the :replace-input operation. This is the default.
:enable	Perform the operation.
:error	Signal an error.

:read-bp of **si:interactive-stream** *Method*
Returns the value of the scan pointer. This is for the benefit of read functions that might want to return a pointer into the input buffer when signalling an error of type **sys:parse-error**.

:noise-string-out *string* &optional (*rescan-mode* **:ignore**) of **si:interactive-stream** *Method*

This message can be sent by a read function to display a string that is not to be treated as input. For example, the string might prompt the user for a particular kind of input. *string* is displayed on the screen without changing the scan pointer, and a rescan does not take place. If a rescan takes place at some later time, the characters in *string* are ignored.

rescan-mode specifies what action to take if the **:noise-string-out** message is sent when the scan pointer is not at the end of the input buffer:

:ignore	Don't perform the :noise-string-out operation. This is the default.
:enable	Perform the operation.
:error	Signal an error.

7. The `:read` and `:print` Stream Operations

A stream can specially handle the reading and printing of objects by handling the `:read` and `:print` stream operations. Note that these operations are optional and that most streams do not support them.

If the `read` function is given a stream that has `:read` in its which-operations, then instead of reading in the normal way it sends the `:read` message to the stream with one argument, `read`'s *eof-option* if it had one or a magic internal marker if it did not. Whatever the stream returns is what `read` returns. If the stream wants to implement the `:read` operation by internally calling `read`, it must use a different stream that does not have `:read` in its which-operations.

If a stream has `:print` in its which-operations, it can intercept all object printing operations, including those due to the `print`, `prin1`, and `princ` functions, those due to `format`, and those used internally, for instance in printing the elements of a list. The stream receives the `:print` message with three arguments: the object being printed, the *prindepth* (for comparison against the `prinlevel` variable), and *slashify-p* (t for `prin1`, nil for `princ`). If the stream returns `nil`, then normal printing takes place as usual. If the stream returns non-`nil`, then `print` does nothing; the stream is assumed to have output an appropriate printed representation for the object. The two following functions are useful in this connection; however, they are in the `system-internals` package and might be changed without much notice.

si:print-object *object prindepth slashify-p stream &optional which-operations* *Function*

Outputs the printed representation of *object* to *stream*, as modified by *prindepth* and *slashify-p*. This is the guts of the Lisp printer. When a stream's `:print` handler calls this function, it should supply the list (**:string-out**) for *which-operations*, to prevent itself from being called recursively. It can supply `nil` if it does not want to receive **:string-out** messages.

Advising this function is the way to customize the behavior of all printing of Lisp objects. See the special form `advise` in *Program Development Utilities*.

si:print-list *list prindepth slashify-p stream which-operations* *Function*

This is the part of the Lisp printer that prints lists. A stream's `:print` handler can call this function, passing along its own arguments and its own which-operations, to arrange for a list to be printed the normal way and the stream's `:print` hook to get a chance at each of the list's elements.

8. Input Functions

Most of these functions take optional arguments called *stream* and *eof-option*. *stream* is the stream from which the input is to be read; if unsupplied it defaults to the value of **standard-input**. The special pseudostreams **nil** and **t** are also accepted, mainly for Maclisp compatibility. **nil** means the value of **standard-input** (that is, the default) and **t** means the value of **terminal-io** (that is, the interactive terminal). This is all more or less compatible with Maclisp, except that instead of the variable **standard-input** Maclisp has several variables and complicated rules. See the section "Introduction to Streams", page 11. Streams are documented in detail in that section.

eof-option controls what happens if input is from a file (or any other input source that has a definite end) and the end of the file is reached. If no *eof-option* argument is supplied, an error is signalled. If there is an *eof-option*, it is the value to be returned. Note that an *eof-option* of **nil** means to return **nil** if the end of the file is reached; it is *not* equivalent to supplying no *eof-option*.

Functions such as **read** that read an "object" rather than a single character always signal an error, regardless of *eof-option*, if the file ends in the middle of an object. For example, if a file does not contain enough right parentheses to balance the left parentheses in it, **read** complains. If a file ends in a symbol or a number immediately followed by end-of-file, **read** reads the symbol or number successfully and when called again, sees the end-of-file and obeys *eof-option*. If a file contains ignorable text at the end, such as blank lines and comments, **read** does not consider it to end in the middle of an object and obeys *eof-option*.

These end-of-file conventions are not completely compatible with Maclisp. Maclisp's deviations from this are generally considered to be bugs rather than features.

The functions below that take *stream* and *eof-option* arguments can also be called with the *stream* and *eof-option* in the other order. This functionality is only for compatibility with old Maclisp programs, and should never be used in new programs. The functions attempt to figure out which way they were called by seeing whether each argument is a plausible stream. Unfortunately, there is an ambiguity with symbols: a symbol might be a stream and it might be an *eof-option*. If there are two arguments, one a symbol and the other something that is a valid stream, or only one argument, which is a symbol, then these functions interpret the symbol as an *eof-option* instead of as a stream. To force them to interpret a symbol as a stream, give the symbol an **si:io-stream-p** property whose value is **t**.

Note that all of these functions except **readline-no-echo** echo their input if used on an interactive stream (one that supports the **:input-editor** operation). The functions that input more than one character at a time (**read**, **readline**) allow the input to be edited using rubout. **typeek** echoes all of the characters that were

skipped over if **tyi** would have echoed them; the character not removed from the stream is not echoed either.

read &optional (*stream standard-input*) *eof-option* *Function*

Reads in the printed representation of a Lisp object from *stream*, builds a corresponding Lisp object, and returns the object. For details: See the section "Input Functions", page 75.

(This function can take its arguments in the other order, for Maclisp compatibility only.)

read-preserve-delimiters *Variable*

Certain printed representations given to **read**, notably those of symbols and numbers, require a delimiting character after them. (Lists do not, because the matching close parenthesis serves to mark the end of the list.) Normally **read** throws away the delimiting character if it is "whitespace", but preserves it (with a **:untyi** stream operation) if the character is syntactically meaningful, since it might be the start of the next expression.

If **read-preserve-delimiters** is bound to **t** around a call to **read**, no delimiting characters are thrown away, even if they are whitespace. This might be useful for certain reader macros or special syntaxes.

sys:read-character &optional *stream* &key (*fresh-line t*) (*any-tyi nil*) *Function*
(eof nil) (*notification t*) (*prompt nil*) (*help nil*)
(refresh t) (*suspend t*) (*abort t*) (*status nil*)

Reads and returns a single character from *stream*. This function displays notifications and help messages and reprompts at appropriate times. It is used by **query** and the **:character** option for **prompt-and-read**.

stream must be interactive. It defaults to **query-io**.

Following are the permissible keywords:

- :fresh-line** If not **nil**, the function sends the stream a **:fresh-line** message before displaying the prompt. If **nil**, it does not send a **:fresh-line** message. The default is **t**.
- :any-tyi** If not **nil**, the function returns blips. If **nil**, blips are treated as the **:tyi** message to an interactive stream treats them. The default is **nil**.
- :eof** If not **nil** and the function encounters end-of-file, it returns **nil**. If **nil** and the function encounters end-of-file, it beeps and waits for more input. The default is **nil**.
- :notification** If not **nil** and a notification is received, the function displays the notification and reprompts. If **nil** and a notification is received, the notification is ignored. The default is **t**.

- :prompt** If **nil**, no prompt is displayed. Otherwise, the value should be a prompt option to be displayed at appropriate times. See the section "Displaying Prompts in the Input Editor", page 65. The default is **nil**.
- :help** If not **nil**, the value should be a help option. See the section "Displaying Help Messages in the Input Editor", page 66. Then, when the user presses HELP, the function displays the help option and reprompts. If **nil** and the user presses HELP, the function just returns **#\help**. The default is **nil**.
- :refresh** If not **nil** and the user presses REFRESH, the function sends the stream a **:clear-window** message and reprompts. If **nil** and the user presses REFRESH, the function just returns **#\refresh**. The default is **t**.
- :suspend** If not **nil** and the user types one of the **sys:kbd-standard-suspend-characters**, a **break** loop is entered. If **nil** and the user types a suspend character, the function just returns the character. The default is **t**.
- :abort** If not **nil** and the user types one of the **sys:kbd-standard-abort-characters**, **sys:abort** is signalled. If **nil** and the user types an abort character, the function just returns the character. The default is **t**.
- :status** This option takes effect only if the stream is a window. If the value is **:selected** and the window is no longer selected, the function returns **:status**. If the value is **:exposed** and the window is no longer exposed or selected, the function returns **:status**. If the value is **nil**, the function continues to wait for input when the window is deexposed or deselected. The default is **nil**.

tyi &optional *stream eof-option* *Function*

Inputs one character from *stream* and returns it. The character is echoed if *stream* is interactive, except that Rubout is not echoed. The Control, Meta, and so on shifts echo as prefix c-, m-, and so on.

The **:tyi** stream operation is preferred over the **tyi** function for some purposes. Note that it does not echo. See the message **:tyi**, page 11.

(This function can take its arguments in the other order, for Maclisp compatibility only)

read-for-top-level &optional (*stream standard-input*) *eof-option* *Function*

Differs from **read** only in that it ignores close parentheses seen at top level, and it returns the symbol **si:eof** if the stream reaches end-of-file if you have not supplied an *eof-option* (instead of signalling an error as **read** would).

This version of **read** is used in the system's "read-eval-print" loops.

read-expression &optional *stream* &key (*completion-alist* **nil**) *Function*
(*completion-delimiters* **nil**)

Like **read-for-top-level** except that if it encounters a top-level end-of-file, it just beeps and waits for more input. This function is used by the **:expression** option for **prompt-and-read**.

stream defaults to **standard-input**. This function is intended to read only from interactive streams.

If *completion-alist* is not **nil**, this function also sets up **COMPLETE** and **c-?** as input editor commands. When the user presses **COMPLETE**, the input editor tries to complete the current symbol over the set of possibilities defined by *completion-alist*. When the user presses **c-?**, the input editor displays the possible completions of the current symbol.

The style of completion is the same as that offered by **Zwei**. *completion-alist* can be **nil**, an alist, an **art-q-list** array, or a keyword:

nil	No completion is offered.
alist	The car of each alist element is a string representing one possible completion.
array	Each element is a list whose car is a string representing one possible completion. The array must be sorted alphabetically on the cars of the elements.
keyword	If the symbol is :zmacs , completion is offered over the definitions in Zmacs buffers. If the symbol is :flavors , completion is offered over all flavor names.

The default for *completion-alist* is **nil**.

completion-delimiters is **nil** or a list of characters that delimit "chunks" for completion. As in **Zwei**, completion works by matching initial substrings of "chunks" of text. If *completion-delimiters* is **nil**, the entire text of the current symbol is a single "chunk". The default is **nil**.

read-form &optional *stream* &key (*edit-trivial-errors-p* *Function*
read-form-edit-trivial-errors-p
(*completion-alist* ***read-form-completion-alist***)
(*completion-delimiters* ***read-form-completion-delimiters***)

Like **read-expression** except that it assumes that the returned value will be given immediately to **eval**. This function is used by the Lisp command loop and by the **:eval-form** and **:eval-form-or-end** options for **prompt-and-read**.

stream defaults to **standard-input**. This function is intended to read only from interactive streams.

If *edit-trivial-errors-p* is not **nil**, the function checks for two kinds of errors. If a symbol is read, it checks whether the symbol is bound. If a list whose first element is a symbol is read, it checks whether the symbol has a function definition. If it finds an unbound symbol or undefined function, it offers to use a lookalike symbol in another package or calls **parse-error** to let the user correct the input. *edit-trivial-errors-p* defaults to the value of ***read-form-edit-trivial-errors-p***. The default value is **t**.

If *completion-alist* is not **nil**, this function also sets up COMPLETE and c-? as input editor commands. When the user presses COMPLETE, the input editor tries to complete the current symbol over the set of possibilities defined by *completion-alist*. When the user presses c-?, the input editor displays the possible completions of the current symbol.

The style of completion is the same as that offered by *Zwei*. *completion-alist* can be **nil**, an alist, an **art-q-list** array, or a keyword:

nil	No completion is offered.
alist	The car of each alist element is a string representing one possible completion.
array	Each element is a list whose car is a string representing one possible completion. The array must be sorted alphabetically on the cars of the elements.
keyword	If the symbol is :zmacs , completion is offered over the definitions in Zmacs buffers. If the symbol is :flavors , completion is offered over all flavor names.

The default for *completion-alist* is the value of ***read-form-completion-alist***. The default value is **:zmacs**.

completion-delimiters is **nil** or a list of characters that delimit "chunks" for completion. As in *Zwei*, completion works by matching initial substrings of "chunks" of text. If *completion-delimiters* is **nil**, the entire text of the current symbol is a single "chunk". The default is the value of ***read-form-completion-delimiters***. The default value is (**#/- #/: #\space**).

read-form-edit-trivial-errors-p

Variable

If not **nil**, **read-form** checks for two kinds of errors. If a symbol is read, it checks whether the symbol is bound. If a list whose first element is a symbol is read, it checks whether the symbol has a function definition. If it finds an unbound symbol or undefined function, it offers to use a lookalike symbol in another package or calls **parse-error** to let the user correct the input. The default is **t**.

read-form-completion-alist*Variable*

If not **nil**, **read-form** sets up COMPLETE and c-? as input editor commands. When the user presses COMPLETE, the input editor tries to complete the current symbol over the set of possibilities defined by *completion-alist*. When the user presses c-?, the input editor displays the possible completions of the current symbol.

The style of completion is the same as that offered by Zwei.

read-form-completion-alist can be **nil**, an alist, an **art-q-list** array, or a keyword:

nil	No completion is offered.
alist	The car of each alist element is a string representing one possible completion.
array	Each element is a list whose car is a string representing one possible completion. The array must be sorted alphabetically on the cars of the elements.
keyword	If the symbol is :zmacs , completion is offered over the definitions in Zmacs buffers. If the symbol is :flavors , completion is offered over all flavor names.

The default value is **:zmacs**.

read-form-completion-delimiters*Variable*

The value is **nil** or a list of characters that delimit "chunks" for completion in **read-form**. As in Zwei, completion works by matching initial substrings of "chunks" of text. If ***read-form-completion-delimiters*** is **nil**, the entire text of the current symbol is a single "chunk". The default value is (#/- #/: #\space).

read-or-end &optional (*stream standard-input*) *reader**Function*

Like **read-expression** except that if it is reading from an interactive stream and the user presses END as the first character or the first character after only whitespace characters, it returns two values, **nil** and **:end**. If it encounters any nonwhitespace characters, it calls the *reader* function with an argument of *stream* to read the input. *reader* defaults to **read-expression**. *stream* defaults to **standard-input**.

The **:expression-or-end** and **:eval-form-or-end** options for **prompt-and-read** invoke **read-or-end**.

This function is intended to read only from interactive streams.

read-or-character &optional *delimiters stream reader**Function*

Like **read-expression**, except that if it is reading from an interactive stream and the user types one of the *delimiters* as the first character or the first

character after only whitespace characters, it returns four values: **nil**, **:character**, the character code of the delimiter, and any numeric argument to the delimiter. If it encounters any nonwhitespace characters, it calls the *reader* function with an argument of *stream* to read the input.

delimiters is a character, a list of characters, or **nil**. The default is **nil**. *reader* defaults to **read-expression**. *stream* defaults to **standard-input**. This function is intended to read only from interactive streams.

read-and-eval &optional *stream* (*catch-errors* *t*) *Function*

Calls **read-expression** to read a form, without completion. It then evaluates the form and returns the result. If *catch-errors* is not **nil**, it calls **parse-error** if an error occurs during the evaluation (but not the reading) so that the input editor catches the error.

stream defaults to **standard-input**. This function is intended to read only from interactive streams.

readline &optional (*stream* **standard-input**) *eof-option* *Function*

Reads in a line of text. If called from inside the input editor or if reading from a stream that does not support the input editor, the line is terminated by a Newline character. If the stream supports the input editor and **readline** is called from outside the input editor, the line is terminated by RETURN, LINE, or END.

This function is usually used to get a line of input from the user. If *stream* supports the input editor, **readline** calls **read-delimited-string**, and *input-editor-options* is passed as the list of options to the input editor.

readline returns two values:

- The line as a character string, without the Newline character, or if already at end-of-file, **nil**.
- An *eof* flag, if *eof-option* was **nil**. This is **t** if the line was terminated because end-of-file was encountered, or **nil** if it was terminated because of a RETURN, LINE, or END character.

See the function **read-delimited-string**, page 83.

readline-trim &optional (*stream* **standard-input**) *eof-option* *Function*

Trims leading and trailing whitespace from string input. "Whitespace" means spaces, tabs, or newlines. It takes the same arguments as the normal **readline** and returns the same values.

Examples:

```
(readline-trim) exciting option RETURN =>
"exciting option"
NIL
141
NIL
```

```
(readline-trim)RETURN =>
""
NIL
141
NIL
```

The **:string-trim** option for **prompt-and-read** and the **:string-trim tv:choose-variable-values** keyword use **readline-trim**.

readline-or-nil &optional (*stream standard-input*) *eof-option* *Function*

Like **readline-trim**, except that it returns a first value of **nil** instead of the empty string if the input string is empty.

The **:string-or-nil** option for **prompt-and-read** and the **:string-or-nil tv:choose-variable-values** keyword use **readline-or-nil**.

See the function **readline-trim**, page 81.

readline-no-echo &optional *stream* &key (*terminators* *Function*

'(#\return #\line #\end)) (*full-rubout nil*)
(*notification t*) (*prompt nil*) (*help nil*)

Reads a line of input from *stream* without echoing the input, and returns the input as a string, without the terminating character. This function is used to read passwords and encryption keys. It does not use the input editor but does allow input to be edited using RUBOUT.

stream must be interactive. It defaults to **query-io**.

Following are the permissible keywords:

:terminators A list of characters that terminate the input. If the user types **#\return**, **#\line**, or **#\end** as a terminator, the function echoes a Newline. If the user types any other character as a terminator, the function echoes that character. The default is (**#\return #\line #\end**).

:full-rubout If not **nil** and the user rubs out all characters on the line, the function returns **nil**. If **nil** and the user rubs out all characters on the line, the function waits for more input. The default is **nil**.

:notification If not **nil** and a notification is received, the function displays the notification and reprompts. If **nil** and a notification is received, the notification is ignored. The default is **t**.

- :prompt** If **nil**, no prompt is displayed. Otherwise, the value should be a prompt option to be displayed at appropriate times. See the section "Displaying Prompts in the Input Editor", page 65. The default is **nil**.
- :help** If not **nil**, the value should be a help option. See the section "Displaying Help Messages in the Input Editor", page 66. Then, when the user presses HELP, the function displays the help option and reprompts. If **nil** and the user presses HELP, the function just returns **#\help**. The default is **nil**.

read-delimited-string &optional (*delimiters* **#\end**) (*stream* **standard-input**) (*eof* **nil**) (*input-editor-options* **nil**) &rest (*make-array-args* '(100. :type **art-string**)) *Function*

delimiter is either a character or a list of characters. Characters are read from *stream* until one of the delimiter characters is encountered. The characters read up to the delimiter are returned as a string. This function can be invoked from inside or outside the input editor. If invoked from outside the input editor, the delimiter characters are set up as activation characters. The *eof* argument is treated the same way as the *eof* argument to the **:tyi** message to noninteractive streams. *input-editor-options* are passed on as the first argument to the **:rubout-handler** message, after having an **:activation** entry prepended. *make-array-args* are arguments to be passed to **make-array** when constructing the string to return.

read-delimited-string returns four values:

- The string
- An *eof* flag, if the *eof* parameter was **nil**
- The character that delimited the string
- Any numeric argument given the delimiter character

This function is used by **readline**, **qsend**, and the **:delimited-string** option for **prompt-and-read**.

Examples:

The following reads characters until **END** is typed and returns a string at least 200. characters long with a leader-length of 3:

```
(read-delimited-string #\end standard-input nil nil 200. :leader-length 3)
```

The following is the same as (**readline**), except that it does not echo a Newline after the string is activated:

```
(read-delimited-string '#\return #\line #\end)
```

A simple word parser:

```
(read-delimited-string '(#\space #/, #/. #/?))
```

For a more complex example of a sentence parser that uses **read-delimited-string**: See the section "Examples of Use of the Input Editor", page 66.

readch &optional *stream eof-option* *Function*

Provided only for Maclisp compatibility, since in Zetalisp characters are always represented as integers. **readch** is just like **tyi**, except that instead of returning an integer character, it returns a symbol whose print name is the character read in. The symbol is interned in the current package. This is just like a Maclisp "character object". (This function can take its arguments in the other order, for Maclisp compatibility only.)

tyipeek &optional *peek-type stream eof-option* *Function*

Provided mainly for Maclisp compatibility; the **:tyipeek** stream operation is usually clearer.

What **tyipeek** does depends on the *peek-type*, which defaults to **nil**. With a *peek-type* of **nil**, **tyipeek** returns the next character to be read from *stream*, without actually removing it from the input stream. The next time input is done from *stream* the character is still there; in general, (= (**tyipeek**) (**tyi**)) is **t**. See the message **:tyipeek**, page 15.

If *peek-type* is an integer less than 1000 octal, then **tyipeek** reads characters from *stream* until it gets one equal to *peek-type*. That character is not removed from the input stream.

If *peek-type* is **t**, then **tyipeek** skips over input characters until the start of the printed representation of a Lisp object is reached. As above, the last character (the one that starts an object) is not removed from the input stream.

The form of **tyipeek** supported by Maclisp, in which *peek-type* is an integer not less than 1000 octal, is not supported, since the readable formats of the Maclisp reader and the Symbolics-Lisp reader are quite different.

Characters passed over by **tyipeek** are echoed if *stream* is interactive.

The following functions are related functions that do not operate on streams. Most of the text at the beginning of this section does not apply to them.

read-from-string *string* &optional (*eof-option* **'si:no-eof-option**) *Function*
(*start* **0**) *end*

The characters of *string* are given successively to the reader, and the Lisp object built by the reader is returned. Macro characters and so on all take effect. If *string* has a fill-pointer it controls how much can be read.

eof-option is what to return if the end of the string is reached, as with other

reading functions. *start* is the index in the string of the first character to be read. *end*, if given, is used instead of (**array-active-length** *string*) as the integer that is one greater than the index of the last character to be read.

read-from-string returns two values: The first is the object read and the second is the index of the first character in the string not read. If the entire string was read, this is the length of the string.

Example:

```
(read-from-string "(a b c)") => (a b c) and 7
```

readlist *char-list*

Function

Provided mainly for Maclisp compatibility. *char-list* is a list of characters. The characters can be represented by anything that the function **character** accepts: integers, strings, or symbols. The characters are given successively to the reader, and the Lisp object built by the reader is returned. Macro characters and so on all take effect.

If there are more characters in *char-list* beyond those needed to define an object, the extra characters are ignored. If there are not enough characters, an "eof in middle of object" error is signalled.

See the special form **with-input-from-string** in *Reference Guide to Symbolics-lisp*.

9. Output Functions

These functions all take an optional argument called *stream*, which is where to send the output. If unsupplied *stream* defaults to the value of **standard-output**. If *stream* is **nil**, the value of **standard-output** (that is, the default) is used. If it is **t**, the value of **terminal-io** is used (that is, the interactive terminal). If *stream* is a list of streams, then the output is performed to all of the streams (this is not implemented yet, and an error is signalled in this case). This is all more or less compatible with Maclisp, except that instead of the variable **standard-output** Maclisp has several variables and complicated rules. See the section "Introduction to Streams", page 11. Streams are documented in detail in that section.

prin1 *x* &optional *stream* *Function*
 Outputs the printed representation of *x* to *stream*, with slashification. *x* is returned. See the section "What the Printer Produces" in *Reference Guide to Symbolics-lisp*.

prin1-then-space *x* &optional *stream* *Function*
 Like **prin1** except that output is followed by a space.

print *x* &optional *stream* *Function*
 Like **prin1** except that output is preceded by a carriage return and followed by a space. *x* is returned.

princ *x* &optional *stream* *Function*
 Like **prin1** except that the output is not slashified. *x* is returned.

tyo *char* &optional *stream* *Function*
 Outputs the character *char* to *stream*.

terpri &optional *stream* *Function*
 Outputs a carriage return character to **stream**.

The **format** function can do anything any of the above functions can do and is very useful for producing nicely formatted text. See the function **format**, page 29. **format** can generate a string or output to a stream.

The **grindef** function is useful for formatting Lisp programs. See the special form **grindef**, page 51.

See the special form **with-output-to-string** in *Reference Guide to Symbolics-lisp*.

stream-copy-until-eof *from-stream to-stream* &optional *leader-size* *Function*
 Inputs characters from *from-stream* and outputs them to *to-stream*, until it reaches the end-of-file on the *from-stream*. For example, if **x** is bound to a

stream for a file opened for input, then
(stream-copy-until-eof x terminal-io) prints the file on the console.

If *from-stream* supports the **:line-in** operation and *to-stream* supports the **:line-out** operation, then **stream-copy-until-eof** uses those operations instead of **:tyi** and **:tyo**, for greater efficiency. *leader-size* is passed as the argument to the **:line-in** operation.

beep &optional *beep-type* (*stream terminal-io*) *Function*

Tries to attract the user's attention by causing an audible beep, or flashing the screen, or something similar. If the stream supports the **:beep** operation, then this function sends it a **:beep** message, passing *type* along as an argument. Otherwise it just causes an audible beep on the terminal. *type* is a keyword selecting among several different beeping noises. The allowed types have not yet been defined; *type* is currently ignored and should always be **nil**. See the message **:beep**, page 19.

cursorpos &rest *args* *Function*

Exists primarily for Maclisp compatibility. Usually it is preferable to send the appropriate messages.

cursorpos normally operates on the **standard-output** stream; however, if the last argument is a stream or **t** (meaning **terminal-io**), **cursorpos** uses that stream and ignores it when doing the operations described below. Note that **cursorpos** only works on streams that are capable of these operations, such as windows. A stream is taken to be any argument that is not a number and not a symbol, or a symbol other than **nil** with a name more than one character long.

(cursorpos) => (*line . column*), the current cursor position.

(cursorpos line column) moves the cursor to that position. It returns **t** if it succeeds and **nil** if it does not.

(cursorpos op) performs a special operation coded by *op*, and returns **t** if it succeeds and **nil** if it does not. *op* is tested by string comparison, is not a keyword symbol, and can be in any package.

F Moves one space to the right.

B Moves one space to the left.

D Moves one line down.

U Moves one line up.

T Homes up (moves to the top left corner). Note that **t** as the last argument to **cursorpos** is interpreted as a stream, so a stream *must* be specified if the **T** operation is used.

Z Home down (moves to the bottom left corner).

- A** Advances to a fresh line. See the **:fresh-line** stream operation.
- C** Clears the window.
- E** Clear from the cursor to the end of the window.
- L** Clear from the cursor to the end of the line.
- K** Clear the character position at the cursor.
- X** B then K.

exploden *x* *Function*

Returns a list of characters (as integers) that are the characters that would be typed out by (**princ** *x*) (that is, the unslashified printed representation of *x*). Example:

```
(exploden '(+ /12 3)) => (50 53 40 61 62 40 63 51)
```

explodec *x* *Function*

Returns a list of characters represented by symbols that are the characters that would be typed out by (**princ** *x*) (that is, the unslashified printed representation of *x*). Example:

```
(explodec '(+ /12 3)) => ( / ( + / /1 /2 / /3 / ) )
```

(Note that there are slashified spaces in the above list.)

explode *x* *Function*

Returns a list of characters represented by symbols that are the characters that would be typed out by (**prin1** *x*) (that is, the slashified printed representation of *x*). Example:

```
(explode '(+ /12 3)) => ( / ( + / // /1 /2 / /3 / ) )
```

(Note that there are slashified spaces in the above list.)

flatsize *x* *Function*

Returns the number of characters in the slashified printed representation of *x*.

flatc *x* *Function*

Returns the number of characters in the unslashified printed representation of *x*.

10. Accessing Files

Symbolics-Lisp can access files on a variety of remote file servers, which are typically (but not necessarily) accessed through the Chaosnet, as well as accessing files on the Symbolics computer itself, if the machine has its own file system. This section tells you how to get a stream that reads or writes a given file, and what the device-dependent operations on that stream are. Files are named with *pathnames*. Since pathnames are quite complex they have their own chapter. See the section "Naming of Files", page 127.

with-open-file (*stream pathname options...*) *body...* *Special Form*
 Evaluates the *body* forms with the variable *stream* bound to a stream that reads or writes the file named by the value of *pathname*. The *options* forms evaluate to the file-opening options to be used.

When control leaves the body, either normally or abnormally (via **throw**), the file is closed. If a new output file is being written, and control leaves abnormally, the file is aborted and it is as if it were never written. Because it always closes the file, even when an error exit is taken, **with-open-file** is preferred over **open**. Opening a large number of files and forgetting to close them tends to break some remote file servers, ITS's for example.

pathname is the name of the file to be opened; it can be a pathname object, a string, a symbol, or a Maclisp-compatible "namelist". It can be anything acceptable to **fs:parse-pathname**. See the section "Naming of Files", page 127. The complete rules for parsing pathnames are explained there.

If an error occurs, such as file not found, the user is asked to supply an alternate pathname, unless this is overridden by *options*. At that point, the user can quit out or enter the Debugger, if the error was not due to a misspelled pathname.

If you are opening the file to read it with **read**, and you want to bind the package and so forth, see the special functions for handling file attributes. See the function **fs:read-attribute-list**, page 111. See the function **fs:file-attribute-bindings**, page 111.

The *options* used when opening a file are normally alternating keywords and values, like any other function that takes keyword arguments. In addition, for compatibility with the Maclisp **open** function, if only a single option is specified it is either a keyword or a list of keywords (not alternating with values).

The file-opening options control whether the stream is for input from an existing file or output to a new file, whether the file is text or binary, and so on.

The following option keywords are recognized. Unless otherwise noted, they are supported generically. Additional keywords can be implemented by particular file system hosts.

:direction	The :direction option allows the following values:
:input	The file is being opened for input. This is the default.
:output	The file is being opened for output.
:io	The file is being opened for intermixed input and output. Bidirectionality is supported only if the stream is to be a direct stream, that is, :direct t is given as well. See the section "Direct Access File Streams", page 115.
:probe	A "probe" opening; no data are to be transferred, and the file is being opened only to gain access to or change its properties. Returns the truenamename of the object at the end of a link or chain of links. (probef is usually preferable to an explicit probe opening.)
:probe-link	The same as :probe except that links are not chased. Returns the truenamename of the object named, even if it is a link.
:probe-directory	The pathname is being opened to find out about the existence of its <i>directory</i> component. Otherwise, the semantics are the same as :probe . If the directory is not found, a file lookup error is signalled.
:probe-link	The same as :probe except that links are not chased. Returns the truenamename of the object named, even if it is a link.
:direct	The default is nil . t specifies a direct access stream. See the section "Direct Access File Streams", page 115.
:submit	This is an option to open used to get batch jobs. Currently, this is implemented only for VAX/VMS. When the file you are writing is closed, the file is submitted as a batch job by using this option.
:characters	The possible values are t (the default), nil , which means that the file is a binary file, and :default , which means

that the file system should decide whether the file contains characters or binary data and open it in the appropriate mode.

:byte-size The possible values are **nil** (the default), a number in the range 1 to 16 inclusive, which is the number of bits per byte, and **:default**, which means that the file system should choose the byte size based on attributes of the file. If the file is being opened as characters, **nil** selects the appropriate system-dependent byte size for text files; it is usually not useful to use a different byte size. If the file is being opened as binary, **nil** selects the default byte size of 16 bits.

:error This option controls what happens when any **fs:file-operation-failure** condition is signalled. See the section "File-system Errors" in *Reference Guide to Symbolics-lisp*.

The option has three possible values:

<i>Value</i>	<i>Meaning</i>
t	Signals the error normally. t is both the default and the recommended value.
nil	Returns the condition object.
:reprompt	Reprompts the user for another file name and tries open again. When you use this option, remember that the :pathname message sent to the stream finds out what file name was really opened.

t is the recommended value for this option. The others have been provided for compatibility with previous systems to aid in converting programs.

The alternative to **:reprompt** is to use **:error t** and set up a condition handler for **fs:file-operation-failure** that explains the condition and prompts the user.

:deleted The default is **nil**. If **t** is specified, and the file system has the concept of deleted but not expunged files, it is possible to open a deleted file. Otherwise deleted files are invisible.

:temporary The default is **nil**. If **t** is specified, the file is marked as temporary, if the file system has that concept.

:preserve-dates The default is **nil**. If **t** is specified, the file's reference and modification dates are not updated.

:estimated-length

The value of the **:estimated-length** option can be **nil** (the default), which means there is no estimated length, or a number of bytes indicating the estimated length of a file to be written. Some file systems use this to optimize disk allocation.

:super-image

The value can be **nil** (the default), or **t** which disables the special treatment of Rubout in ASCII files. Normally Rubout is an escape that causes the following character to be interpreted specially, allowing all characters from 0 through 376 to be stored. This applies to PDP-10 file servers only.

:raw

The value can be **nil** (the default) or **t**, which disables all character set translation in ASCII files.

:if-exists

Specifies the action to be taken if the **:direction** is **:output** and a file of the specified name already exists. If the direction is **:input** or **:probe** (or any of the **:probe-like** directions), this argument is ignored.

The following values are allowed:

:error

Signals an error. This is the default when the version component of the filename is not **:newest**.

:new-version

Creates a new file with the same file name but a larger version number. This is the default when the version component of the filename is **:newest**. File systems without version numbers can choose to implement this by effectively treating it as **supersede**.

:rename

Renames the existing file to some other name, and then creates a new file with the specified name. On most file systems, this renaming happens at the time of a successful close.

:rename-and-delete

Renames the existing file to some other name and then deletes it (but does not expunge it, on those systems that distinguish deletion from expunging). Then creates a new file with the specified name. On most file systems, this renaming happens at the time of a successful close.

:overwrite	The existing file is used, and output operations on the stream destructively modify the file. The file pointer is initially positioned at the beginning of the file; however, the file is not truncated back to length zero when it is opened.
:truncate	The existing file is used, and output operations on the stream destructively modify the file. The file pointer is initially positioned at the beginning of the file; at that time, the file is truncated to length zero, and disk storage occupied by it is freed.
:append	The existing file is used, and output operations on the stream destructively modify the file. The file pointer is initially positioned at the current end of the file.
:supersede	Supersedes the existing file. If possible, the file system does not destroy the old file until the new stream is closed, against the possibility that the stream will be closed in "abort" mode. This differs from :new-version in that :supersede creates a new file with the same name as the old one, rather than a file name with a higher version number.
nil	Does not create a file or even a stream. Instead, simply returns nil to indicate failure.
:if-does-not-exist	Specifies the action to be taken if the file does not already exist. The following values are allowed:
:error	Signals an error. This is the default if the :direction is :input , :probe , or any of the :probe -like modes, or if the :if-exists argument is :overwrite , :truncate , or :append .
:create	Creates an empty file with the specified name, and then proceeds as if it had already existed. This is the default if

the **:direction** is **:output** and the **:if-exists** argument is anything but **:overwrite**, **:truncate**, or **:append**.

nil Does not create a file or even a stream. Instead, simply returns **nil** to indicate failure.

In the Maclisp compatibility mode, there is only one *option*, and it is either a symbol or a list of symbols. These symbols are recognized no matter what package they are in, since Maclisp does not have packages. The following symbols are recognized:

in, read Select opening for input (the default).

out, write, print Select opening for output; a new file is to be created.

binary, fixnum Select binary mode, otherwise character mode is used.

character, ascii The opposite of **fixnum**. This is the default.

single, block Ignored for compatibility with the Maclisp **open** function.

byte-size Must be followed by a number in the options list, and must be used in combination with **fixnum**. The number is the number of bits per byte, which can be from 1 to 16. On a PDP-10 file server these bytes are packed into words in the standard way defined by the **ILDB** instruction. The **:tyi** stream operation returns the bytes one at a time.

probe, error, noerror, raw, super-image, deleted, temporary These are not available in Maclisp. The corresponding keywords in the normal form of file-opening options are preferred over these.

with-open-file-case (*variable pathname . options*) &rest clauses *Macro*
 Opens a file, binding the input stream to *variable*, using the pathname and options given in the arguments. In the following example, it executes the first clause when the file is not found. When the file is found without error, it executes the second clause, which is the real reason for trying to open the file in the first place.

```
(with-open-file-case (x "f:>dla>foo.lisp" ':direction ':input)
  (fs:file-not-found (send x ':report error-output))
  (:no-error (stream-copy-until-eof x standard-output)))
```

Any errors other than **file-not-found** (for example, access violations or an unresponsive host) cause an error to be signalled normally.

with-open-stream (*stream-variable construction-form*) &body body *Special Form*
 Like **with-open-file** except that you specify a form whose value is the stream, rather than arguments to **open**. This is used with nonfile streams. See the special form **with-open-file**, page 91.

with-open-stream-case (*variable stream-creation-form*) &rest *clauses* *Macro*

Opens a stream and binds it to *variable*, using *stream-creation-form* to create it. It then executes whichever clause is appropriate, given the condition that resulted from the attempt to create the stream. Refer to the example shown for **with-open-file-case**.

sys:with-open-file-search (*stream-variable* **Macro**
(*operation defaults auto-retry*) (*type-list-function pathname . type-list-args*) .
open-options) *body...*

Performs a **with-open-file**, searching for a file with one of the types in a list of file types. **load** uses this macro when not given a specific file type to search first for a binary file and then for a source file.

The body is evaluated with *stream-variable* bound to a stream that reads or writes the file. *open-options* are alternating keywords and values to be passed to **open**.

type-list-function should be a function whose first argument is *pathname* and whose remaining arguments are *type-list-args*. The function should return two values: a list of file types to be searched, in order of preference, and a base pathname to be merged with the types and *defaults* in searching for the file. *defaults* can be a pathname or a defaults alist; if omitted, the defaults come from **fs:*default-pathname-defaults***. The macro uses **fs:merge-pathname-defaults** for merging.

If no file is found with any of the types in the list of types, **fs:multiple-file-not-found** is signalled. *operation* is the name of the operation that failed; usually this is the name of the function that contains the **sys:with-open-file-search** form. If *auto-retry* is not **nil** and the condition is not handled, the user is prompted for a new pathname.

open *pathname* &rest *options* *Function*

Returns a stream that is connected to the specified file. Unlike **Maclisp**, the **open** function only creates streams for *files*; streams for other devices are created by other functions. If an error occurs, such as file not found, the user is asked to supply an alternate pathname, unless this is overridden by *options*.

When the caller is finished with the stream, it should close the file by using the **:close** operation or the **close** function. The **with-open-file** special form does this automatically, and so is usually preferred. **open** should be used only when the control structure of the program necessitates opening and closing of a file in some way more complex than the simple way provided by **with-open-file**. Any program that uses **open** should set up **unwind-protect** handlers to close its files in the event of an abnormal exit. See the special form **unwind-protect** in *Reference Guide to Symbolics-lisp*.

For example:

```
(defun bliss-compile (file)
  (setq file (fs:parse-pathname file))
  (with-open-file (str "comet:usrd$:[mydir]tempfile.com"
                    ':direction ':output
                    ':characters t
                    ':submit t)
    (send str ':line-out (format nil "$ BLISS ~A" (send file ':string-for-host))))))
```

The *options* used when opening a file are normally alternating keywords and values, like any other function that takes keyword arguments. In addition, for compatibility with the Maclisp **open** function, if only a single option is specified it is either a keyword or a list of keywords (not alternating with values).

The file-opening options control whether the stream is for input from an existing file or output to a new file, whether the file is text or binary, and so on.

The following option keywords are recognized. Unless otherwise noted, they are supported generically. Additional keywords can be implemented by particular file system hosts.

:direction	The :direction option allows the following values:
:input	The file is being opened for input. This is the default.
:output	The file is being opened for output.
:io	The file is being opened for intermixed input and output. Bidirectionality is supported only if the stream is to be a direct stream, that is, :direct t is given as well. See the section "Direct Access File Streams", page 115.
:probe	A "probe" opening; no data are to be transferred, and the file is being opened only to gain access to or change its properties. Returns the truename of the object at the end of a link or chain of links. (probef is usually preferable to an explicit probe opening.)
:probe-link	The same as :probe except that links are not chased. Returns the truename of the object named, even if it is a link.
:probe-directory	The pathname is being opened to find

out about the existence of its *directory* component. Otherwise, the semantics are the same as **:probe**. If the directory is not found, a file lookup error is signalled.

:probe-link The same as **:probe** except that links are not chased. Returns the truename of the object named, even if it is a link.

:direct The default is **nil**. **t** specifies a direct access stream. See the section "Direct Access File Streams", page 115.

:submit This is an option to **open** used to get batch jobs. Currently, this is implemented only for VAX/VMS. When the file you are writing is closed, the file is submitted as a batch job by using this option.

:characters The possible values are **t** (the default), **nil**, which means that the file is a binary file, and **:default**, which means that the file system should decide whether the file contains characters or binary data and open it in the appropriate mode.

:byte-size The possible values are **nil** (the default), a number in the range 1 to 16 inclusive, which is the number of bits per byte, and **:default**, which means that the file system should choose the byte size based on attributes of the file. If the file is being opened as characters, **nil** selects the appropriate system-dependent byte size for text files; it is usually not useful to use a different byte size. If the file is being opened as binary, **nil** selects the default byte size of 16 bits.

:error This option controls what happens when any **fs:file-operation-failure** condition is signalled. See the section "File-system Errors" in *Reference Guide to Symbolics-lisp*.

The option has three possible values:

<i>Value</i>	<i>Meaning</i>
t	Signals the error normally. t is both the default and the recommended value.
nil	Returns the condition object.
:reprompt	Reprompts the user for another file name and tries open again. When you use this option, remember that the

:pathname message sent to the stream finds out what file name was really opened.

t is the recommended value for this option. The others have been provided for compatibility with previous systems to aid in converting programs.

The alternative to **:reprompt** is to use **:error t** and set up a condition handler for **fs:file-operation-failure** that explains the condition and prompts the user.

- :deleted** The default is **nil**. If **t** is specified, and the file system has the concept of deleted but not expunged files, it is possible to open a deleted file. Otherwise deleted files are invisible.
- :temporary** The default is **nil**. If **t** is specified, the file is marked as temporary, if the file system has that concept.
- :preserve-dates** The default is **nil**. If **t** is specified, the file's reference and modification dates are not updated.
- :estimated-length**
The value of the **:estimated-length** option can be **nil** (the default), which means there is no estimated length, or a number of bytes indicating the estimated length of a file to be written. Some file systems use this to optimize disk allocation.
- :super-image** The value can be **nil** (the default), or **t** which disables the special treatment of Rubout in ASCII files. Normally Rubout is an escape that causes the following character to be interpreted specially, allowing all characters from 0 through 376 to be stored. This applies to PDP-10 file servers only.
- :raw** The value can be **nil** (the default) or **t**, which disables all character set translation in ASCII files.
- :if-exists** Specifies the action to be taken if the **:direction** is **:output** and a file of the specified name already exists. If the direction is **:input** or **:probe** (or any of the **:probe-like** directions), this argument is ignored.

The following values are allowed:

- :error** Signals an error. This is the default when the version component of the filename is not **:newest**.
- :new-version** Creates a new file with the same file

name but a larger version number. This is the default when the version component of the filename is **:newest**. File systems without version numbers can choose to implement this by effectively treating it as **supersede**.

- :rename** Renames the existing file to some other name, and then creates a new file with the specified name. On most file systems, this renaming happens at the time of a successful close.
- :rename-and-delete** Renames the existing file to some other name and then deletes it (but does not expunge it, on those systems that distinguish deletion from expunging). Then creates a new file with the specified name. On most file systems, this renaming happens at the time of a successful close.
- :overwrite** The existing file is used, and output operations on the stream destructively modify the file. The file pointer is initially positioned at the beginning of the file; however, the file is not truncated back to length zero when it is opened.
- :truncate** The existing file is used, and output operations on the stream destructively modify the file. The file pointer is initially positioned at the beginning of the file; at that time, the file is truncated to length zero, and disk storage occupied by it is freed.
- :append** The existing file is used, and output operations on the stream destructively modify the file. The file pointer is initially positioned at the current end of the file.
- :supersede** Supersedes the existing file. If possible, the file system does not destroy the old file until the new stream is closed, against the possibility that the stream

will be closed in "abort" mode. This differs from **:new-version** in that **:supersede** creates a new file with the same name as the old one, rather than a file name with a higher version number.

nil Does not create a file or even a stream. Instead, simply returns **nil** to indicate failure.

:if-does-not-exist Specifies the action to be taken if the file does not already exist. The following values are allowed:

:error Signals an error. This is the default if the **:direction** is **:input**, **:probe**, or any of the **:probe**-like modes, or if the **:if-exists** argument is **:overwrite**, **:truncate**, or **:append**.

:create Creates an empty file with the specified name, and then proceeds as if it had already existed. This is the default if the **:direction** is **:output** and the **:if-exists** argument is anything but **:overwrite**, **:truncate**, or **:append**.

nil Does not create a file or even a stream. Instead, simply returns **nil** to indicate failure.

In the Maclisp compatibility mode, there is only one *option*, and it is either a symbol or a list of symbols. These symbols are recognized no matter what package they are in, since Maclisp does not have packages. The following symbols are recognized:

in, read Select opening for input (the default).

out, write, print Select opening for output; a new file is to be created.

binary, fixnum Select binary mode, otherwise character mode is used.

character, ascii The opposite of **fixnum**. This is the default.

single, block Ignored for compatibility with the Maclisp **open** function.

byte-size Must be followed by a number in the options list, and must be used in combination with **fixnum**. The number is the number of bits per byte, which can be from 1 to 16. On a PDP-10 file server these bytes are packed into words in the standard way defined by the **ILDB** instruction. The **:tyi** stream operation returns the bytes one at a time.

probe, error, noerror, raw, super-image, deleted, temporary

These are not available in Maclisp. The corresponding keywords in the normal form of file-opening options are preferred over these.

close *stream* &optional *abortp* *Function*
Sends the **:close** message to *stream*.

The *abortp* argument is normally not supplied. If it is **t**, we are abnormally exiting from the use of this stream. If the stream is outputting to a file, and has not been closed already, the stream's newly created file is deleted, as if it were never opened in the first place. Any previously existing file with the same name remains, undisturbed.

renamef *file new-name* &optional (*error-p t*) *Function*
Renames one file. The Rename File (r-X) command in the editor uses this function.

file can be a pathname, a string, or a stream that is open to a file. The specified file is renamed to *new-name* (a pathname or string). If *error-p* is **t**, when an error occurs it is signalled as a Lisp error. If *error-p* is **nil** and an error occurs, the error object is returned; otherwise the three values described below are returned.

file must refer to a unique file; it cannot contain any wild components. *new-name* can contain wild components, which are eliminated after merging the defaults by means of **:translate-wild-pathname**. **renamef** first attempts to open *file*. When that has happened successfully, it parses *new-name* and merges it (using **fs:merge-pathnames**) against the *link-opaque truename* of *file* and version of **:newest**. This has the following result for version numbers.

<i>Source</i>	<i>Target</i>	<i>Result</i>
>foo>a.b.newest	>bar>	Retains the version number
>foo>a.b.newest	>bar>x	Makes a new version of >bar>x.b

The defaults for *new-name* come from the *link-opaque truename* of *file*. For systems without links, this is indistinguishable from the truename. Otherwise, the *link-opaque truename* depends on whether *file* contains an **:oldest** or **:newest** version. If it does not and if it is fully defaulted, with no wild components, the pathname is its own *link-opaque truename*. If a pathname *x* contains an **:oldest** or **:newest** version, the *link-opaque truename* is the pathname of the file or link that corresponds to *x*, with the version number filled in. For example, renaming the LMFS file >a>p1.lisp to >b> results in >b>p1.lisp, with the version of >a>p1.lisp.newest inherited. This is so whether >a>p1.lisp.newest is a real file, a link, or a rename-through link.

renamef returns three values:

1. The pathname produced by merging and defaulting *new-name*. This is the attempted result of the renaming.
2. The pathname of the object that was actually renamed. This might not be the same as *file*. For example, *file* might have an **:oldest** or **:newest** version, or LMFS rename-through links might be involved. This pathname never has an **:oldest** or **:newest** version.
3. The actual pathname that resulted from the renaming. This might not be the same as *new-name*. For example, *new-name* might have an **:oldest** or **:newest** version, or LMFS create-through links might be involved.

The **:rename** message to streams and pathnames returns the second and third of these values.

Examples:

This example is as simple as possible. Using LMFS, on host johnny, with no links involved:

```
(renamef "johnny:>a>foo.lisp" "bar") =>
#<LMFS-PATHNAME "johnny:>a>bar.lisp">
#<LMFS-PATHNAME "johnny:>a>foo.lisp.17">
#<LMFS-PATHNAME "johnny:>a>bar.lisp.1">
```

This example is as complex as possible. Using LMFS, on host eddie, with links

```
>abel>moe.lisp.4 => >baker>larry.lisp (rename-through) (latest)
>baker>larry.lisp.4 =>
  >charlie>sam.lisp.19 (not rename- or create-through) (latest)
>david>jerry.lisp.5 => >earl>ted.lisp (create-through) (latest)

(renamef "eddie:>abel>moe.lisp.4" "eddie:>david>jerry") =>
#<LMFS-PATHNAME "eddie:>david>jerry.lisp">
#<LMFS-PATHNAME "eddie:>baker>larry.lisp.4">
#<LMFS-PATHNAME "eddie:>earl>ted.lisp.1">
```

deletef *file* &optional (*error-p* *t*) *Function*

Deletes the specified file. *file* can be a pathname or a stream that is open to a file. If *error-p* is *t*, then if an error occurs it is signalled as a Lisp error. If *error-p* is *nil* and an error occurs, the error object is returned; otherwise *t* is returned.

undeletef *file* &optional (*error-p* *t*) *Function*

Undeletes the specified file. *file* can be a pathname or a stream that is open to a file. If *error-p* is *t* and an error occurs, it is signalled as a Lisp error. If

error-p is **nil** and an error occurs, the error object is returned; otherwise **t** is returned. **undeletef** is like **deletef** except that it undeletes the file instead of deleting it. **undeletef** is meaningful only for files in file systems that support undeletion, such as TOPS-20 and the Lisp Machine File System.

fs:file-properties *pathname* &optional (*error-p* **t**) Function

Returns a disembodied property list for a single file (compare this to **fs:directory-list**). The car of the returned list is the truename of the file and the cdr is an alternating list of indicators and values. If *error-p* is **t** (the default) a Lisp error is signalled. If *error-p* is **nil** and an error occurs, the error object is returned.

fs:change-file-properties *pathname error-p* &rest *properties* Function

Some of the properties of a file can be changed, such as its creation date or its author. The properties that can be changed depend on the host file system; a list of the changeable property names is the **:settable-properties** property of the file system as a whole, returned by **fs:directory-list**. See the function **fs:directory-list**, page 117.

fs:change-file-properties changes one or more properties of a file. *pathname* names the file. The *properties* arguments are alternating keywords and values. If the *error-p* argument is **t**, a Lisp error is signalled. If *error-p* is **nil** and an error occurs, the error object is returned. If no error occurs, **fs:change-file-properties** returns **t**.

viewf *pathname* &optional (*stream standard-output*) *leader* Function

Prints the file named by *pathname* onto the *stream*. (The optional third argument is passed as the *leader* argument to **stream-copy-until-eof**.) The name **viewf** is analogous with **deletef**, **renamef**, and so on. Note: **viewf** should not be used for copying files; its output is not the same as the contents of the file (for example, it does a **:fresh-line** operation on the stream before printing the file).

copyf *from-path to-path* &key (*characters* **':default**) (*byte-size* **nil**) Function
 (*copy-creation-date* **t**) (*copy-author* **t**)
 (*report-stream* **nil**) (*create-directories* **':query**)

Copies one file to another. Copy File (M-X) in the editor uses this function.

from-path and *to-path* are the source and destination pathnames, which can be file specifications. *from-path* must refer to a unique file; it cannot contain any wild components. *to-path* can contain wild components, which are eliminated after merging the defaults by means of **:translate-wild-pathname**. **copyf** first attempts to open *from-path*. When that has happened successfully, it parses *to-path* and merges it (using **fs:merge-pathnames**) against the *link-opaque truename* of *from-path* and version of **:newest**. The output file specified by *to-path* is opened with **:if-exists :supersede**. The processing of *to-path* has the following result for version numbers.

<i>Source</i>	<i>Target</i>	<i>Result</i>
>foo>a.b.newest	>bar>	Retains the version number
>foo>a.b.newest	>bar>x	Makes a new version of >bar>x.b

The defaults for *to-path* come from the *link-opaque truename* of *from-path*. For systems without links, this is indistinguishable from the truename. Otherwise, the link-opaque truename depends on whether *from-path* contains an **:oldest** or **:newest** version. If it does not and if it is fully defaulted, with no wild components, the pathname is its own link-opaque truename. If a pathname *x* contains an **:oldest** or **:newest** version, the link-opaque truename is the pathname of the file or link that corresponds to *x*, with the version number filled in. For example, copying the LMFS file >a>p1.lisp to >b> results in >b>p1.lisp, with the version of >a>p1.lisp.newest inherited. This is so whether >a>p1.lisp.newest is a real file, a link, or a rename-through link.

By default, **copyf** copies the creation date and author of the file.

Following is a description of the other options:

:characters	Possible values: :default	copyf decides whether this is a binary or character transfer according to the canonical type of <i>from-path</i> . You do not need to supply this argument for standard file types. For types that are not known canonical types, it opens <i>from-path</i> in :default mode. In that case, the server for the file system containing <i>from-path</i> makes the character-or-binary decision.
	t	Specifies that the transfer must be in character mode.
	nil	Specifies that the transfer must be binary mode (in this case, you must supply <i>byte-size</i> if using a byte size other than 16).
:byte-size		Specifies the byte size with which both files are opened for binary transfers. You must supply :byte-size when :characters is nil and the byte size is other than 16. Otherwise, copyf determines the byte size from the file type for <i>from-path</i> . When <i>from-path</i> is a binary file with a known canonical type, it determines the byte size from the :binary-file-byte-size property of the type. When the file

does not have a known type, it requests the byte size for *from-path* from the file server. When the server for the file system containing *from-path* cannot supply the byte size, it assumes that the byte size is 16.

:report-stream When **:report-stream** is **nil** (the default), the copying takes place with no messages. Otherwise, the value must be a stream for reporting the start and successful completion of the copying. The completion message contains the truename of *to-path*.

:create-directories

Determines whether directories should be created, if needed, for the target of the copy. Permissible values are as follows:

- | | |
|---------------|---|
| t | Try to create the target directory of the copy and all superiors. Report directory creation to standard-output . |
| nil | Do not try to create directories. If the directory does not exist, handle this condition like any other error. |
| :query | If the directory does not exist, ask whether or not to create it. This is the default. |

probef *pathname*

Function

Returns **nil** if there is no file named *pathname*, or signals an error if anything else goes wrong (such as **sys:host-not-responding**). Otherwise, **probef** returns a pathname that is the truename of the file, which can be different from *pathname* because of file links, version numbers, and so on.

fs:close-all-files

Function

Closes all open files. This is useful when a program has run wild opening files and not closing them. It closes all the files in **:abort** mode, which means that files open for output will be deleted. Using this function is dangerous, because you might close files out from under various programs such as Zmacs and Zmail; only use it if you have to and if you feel that you know what you're doing.

fs:*remember-passwords*

Variable

If not **nil**, causes the first password for each file access path to be remembered. This suppresses prompting for passwords on subsequent attempts by the same user to use that access path. The default value is **nil**.

Note that if you set this variable in an init file, your first login password, typed before the init file is loaded, is not remembered.

Caution: Remembered passwords are accessible. Even after you log out the remembered password for each access path is accessible. If password security is important, you probably should not set this variable to a non-**nil** value.

10.1 Loading Files

To *load* a file is to read through the file, evaluating each form in it. Programs are typically stored in files; the expressions in the file are mostly special forms such as **defun** and **defvar** that define the functions and variables of the program.

Loading a compiled (or BIN) file is similar, except that the file does not contain text but rather predigested expressions created by the compiler that can be loaded more quickly.

These functions are for loading single files. There is a system for keeping track of programs that consist of more than one file: See the section "Maintaining Large Programs" in *Program Development Utilities*.

load *pathname* &optional *pkg nonexistent-ok-flag dont-set-default-p no-msg-p* *Function*

Loads the file named by *pathname* into the Lisp environment. The file can be either a Lisp source file or a binary file. If the *pathname* specifies the type, it is used; otherwise, **load** looks first for a binary file, then for a Lisp file. Normally, the file is read into its "home" package, but *pkg* can be supplied to specify the package. *pkg* can be either a package or the name of a package as a string or a symbol. If *pkg* is not specified, **load** prints a message saying what package the file is being loaded into.

nonexistent-ok controls the action of **load** if none of the files is found. If it is **nil** (the default), you are prompted for a new file unless the corresponding condition (**fs:multiple-file-not-found**) is handled. If it is not **nil**, it is the returned value if the file is not found. Other reasons for not finding the file, such as the host being down or the directory not existing, are signalled as different errors. For example, **load** fails when the host is down even when you specified the *nonexistent-ok* argument.

pathname can be anything acceptable to **fs:parse-pathname**. See the section "Naming of Files", page 127. *pathname* is defaulted from **fs:load-pathname-defaults**, which is the set of defaults used by **load** and similar functions. See the variable **fs:load-pathname-defaults**, page 149. Normally **load** updates the *pathname* defaults from *pathname*, but if *dont-set-default* is specified this is suppressed.

If an ITS *pathname* contains an FN1 but no FN2, **load** first looks for the file with an FN2 of BIN, then it looks for an FN2 of >. For non-ITS file systems, this generalizes to: if *pathname* specifies a type and/or a version,

load loads that file. Otherwise it first looks for a binary file, then a Lisp file, in both cases looking for the newest version.

If the value of **no-msg-p** is **t** (it defaults to **nil**), then **load** does not print out the message that it usually prints (that is, the message that tells you that a certain file is being loaded into a certain package).

readfile *pathname* &optional *pkg no-msg-p* *Function*
readfile is the version of **load** for text files. It reads and evaluates each expression in the file. As with **load**, *pkg* can specify what package to read the file into. Unless *no-msg-p* is **t**, a message is printed indicating what file is being read into what package. The defaulting of *pathname* is the same as in **load**.

10.2 File Attribute Lists

Any text file can contain an *attribute list* that specifies several attributes of the file. The functions that load files, the compiler, and the editor look at this attribute list. File attribute lists are especially useful in program source files, that is, a file that is intended to be loaded (or compiled and then loaded).

If the first nonblank line in the file contains the three characters "-*-", some text, and "-*-" again, the text is recognized as the file's attribute list. Each attribute consists of the attribute name, a colon, and the attribute value. If there is more than one attribute they are separated by semicolons. An example of such an attribute list is:

```
; -*- Mode:Lisp; Package:Cellophane; Base:10 -*-
```

The semicolon makes this line look like a comment rather than a Lisp expression. This defines three attributes: mode, package, and base.

The term *attribute list* applies not only to the *-*-* line in character files, but also to an analogous data structure in compiled files. For example, in both cases the attribute list tells **load** what package to load the file into.

An attribute name is made up of letters, numbers, and otherwise-undefined punctuation characters such as hyphens. An attribute value can be such a name, or a decimal number, or several such items separated by commas. Spaces can be used freely to separate tokens. Upper and lowercase letters are not distinguished. There is *no* quoting convention for special characters such as colons and semicolons. File attribute lists are different from Lisp property lists; attribute lists correspond to the text inside a file, while file properties are characteristics of the file itself, such as the creation date.

The file attribute list format actually has nothing to do with Lisp; it is just a convention for placing some information into a file that is easy for a program to interpret.

Symbolics-Lisp has a parser for file attribute lists that creates some Lisp data structure that corresponds to the file attribute list. When a file attribute list is read in and given to the parser (the **fs:read-attribute-list** function), it is converted into Lisp objects as follows: Attribute names are interpreted as Lisp symbols, and interned on the keyword package. Numbers are interpreted as Lisp fixnums, and are read in decimal. If an attribute value contains any commas, then the commas separate several expressions that are formed into a list.

When a file is edited, loaded, or compiled, its file attribute list is read in and the attributes are stored on the attribute list of the generic pathname for that file, where they can be retrieved with the **:get** and **:plist** messages. See the section "Generic Pathnames", page 149. So, to examine the attributes of a file, you usually use messages to a pathname object that represents the generic pathname of a file. Note that there other attributes there, too. The function **fs:read-attribute-list** reads the file attribute list of a file and sets up the attributes on the generic pathname; editing, loading, or compiling a file calls this function, but you can call it yourself if you want to examine the attributes of an arbitrary file.

If the attribute list text contains no colons, it is an old EMACS format, containing only the value of the **Mode** attribute.

The following are some of the attribute names allowed and what they mean.

Mode	The editor major mode to be used when editing this file. This is typically the name of the language in which the file is written. The most common values are Lisp and Text.
Package	The name of the package into which the file is to be loaded. See the section "The Need for Packages" in <i>Reference Guide to Symbolics-lisp</i> .
Base	The number base in which the file is written. This affects both ibase and base , since it is confusing to have different input and output bases. The most common values are 8 and 10 (the default).
Lowercase	If the attribute value is not nil , the file is written in lowercase letters and the editor does not translate to uppercase. (The editor does not translate to uppercase by default unless the user selects "Electric Shift Lock" mode.)
Fonts	The attribute value is a list of font names, separated by commas. The editor uses this for files that are written in more than one font.
Backspace	If the attribute value is not nil , the file can contain backspaces that cause characters to overprint on each other. The default is to disallow overprinting and display backspaces the way other special function keys are displayed. This default is to prevent the confusion that can be engendered by overstruck text.

Patch-File If the attribute value is not **nil**, the file is a "patch file". When it is loaded, the system does not complain about function redefinitions. Furthermore, the remembered source file names for functions defined in this file are changed to this file, but are left as whatever file the function came from originally. In a patch file, the **defvar** special-form turns into **defconst**; thus patch files always reinitialize variables.

You are free to define additional file attributes of your own. However, you should choose names that are different from all the names above, and from any names likely to be defined by anybody else's programs, to avoid accidental name conflicts.

The function **fs:pathname-attribute-list** is generally the most useful function for obtaining a file's attributes.

fs:pathname-attribute-list *pathname* *Function*
Returns the attribute list for a file designated by *pathname*.

fs:read-attribute-list *pathname stream* *Function*
Parses file attribute lists. *pathname* should be a pathname object (*not* a string or namelist, but an actual pathname); usually it is a generic pathname. See the section "Generic Pathnames", page 149.

stream should be a stream that has been opened and is pointing to the beginning of the file whose file attribute list is to be parsed. This function reads from the stream until it gets the file attribute list, parses it, puts corresponding attributes onto the attribute list of *pathname*, and finally sets the stream back to the beginning of the file by using the **:set-pointer** file stream operation. See the message **:set-pointer**, page 20.

The obsolete name of this function is **fs:file-read-property-list**.

Programs in Symbolics-Lisp generally react to the presence of attributes on a file's file attribute list by examining the attribute list in the generic pathname's property list. However, file attributes can also cause special variables to be bound whenever Lisp expressions are being read from the file—when the file is being loaded, when it is being compiled, when it is being read from by the editor, and when its QFASL file is being loaded. This is how the Package and Base attributes work. You can also deal with attributes this way, by using the following function.

fs:file-attribute-bindings *pathname* *Function*
Examines the property list of *pathname* and finds all those property names that have file-attribute bindings. Its obsolete name is **fs:file-property-bindings**.

Each such pathname-property name specifies a set of variables to bind and a set of values to which to bind them. This function returns two values: a list of all the variables, and a list of all the corresponding values. Usually you

call this function on a generic pathname whose attribute list has been parsed with **fs:read-attribute-list**. Then you use the two returned values as the first two subforms to a **progv** special form. Inside the body of the **progv** the specified bindings will be in effect.

Usually, *pathname* is a generic pathname. It can also be a locative, in which case it is interpreted to be the property list itself.

Of the standard names, the following ones have file-attribute bindings, with the following effects:

- **Package** binds the variable **package** to the package. See the variable **package** in *Reference Guide to Symbolics-lisp*.
- **Base** binds the variables **base** and **ibase** to the value. See the variable **base** in *Reference Guide to Symbolics-lisp*. See the variable **ibase** in *Reference Guide to Symbolics-lisp*.
- **Patch-file** binds **fs:this-is-a-patch-file** to the value.

Any properties whose names do not have file-attribute bindings are ignored completely.

You can also add your own pathname-property names that affect bindings. If an indicator symbol has a file-attribute binding, the value of that property is a function that is called when a file with a file attribute of that name is going to be read from. The function is given three arguments: the file pathname, the attribute name, and the attribute value. It must return two values: a list of variables to be bound and a list of values to bind them to. Both these lists must be freshly consed (using **list** or **ncons**). The function for the **Base** keyword could have been defined by:

```
(defun (:base file-attribute-bindings) (file ignore bse)
  (if (not (and (typep bse 'fixnum)
                (> bse 1)
                (< bse 37.)))
      (ferror nil "File ~A has an illegal -*- Base:~s -*-"
              file bse))
      (values (list 'base 'ibase) (list bse bse))))
```

Finally, the function **sys:dump-forms-to-file** offers, among other things, the option of manipulating the attribute list of a binary file. See the section "Putting Data in Compiled Code Files", page 241.

For example, the following form converts a Lisp file to a binary file, without compiling. The attribute list is obtained from the input stream and cached in the generic pathname. The function **fs:file-attribute-bindings** obtains the list of variables to bind from the generic pathname; these bindings are necessary to ensure that the file is read in the right base, syntax, and package. The **progv** actually accomplishes the binding of the variables.

```
(defun binify-file-internal (input-file output-file)
  (with-open-file (input input-file :direction :input :characters t)
    (let* ((generic-pathname (send input-file :generic-pathname))
           (attribute-list (fs:read-attribute-list generic-pathname input)))
      (multiple-value-bind (variables-list values-list)
        (fs:file-attribute-bindings generic-pathname)
        (progv variables-list values-list
              (loop with eof-val = (ncons 'eof)
                    for form = (read input eof-val)
                    while (neq form eof-val)
                    collect form into forms
                    finally
                      (sys:dump-forms-to-file output-file forms
                                             attribute-list)))))))
```

10.3 File Stream Operations

The following messages can be sent to file streams, in addition to the normal I/O messages that work on all streams. Note that several of these messages are useful to send to a file stream which has been closed. Some of these messages use pathnames. See the section "Naming of Files", page 127.

:pathname

Message

Returns the pathname that was opened to get this stream. This might not be identical to the argument to **open**, since missing components will have been filled in from defaults, and the pathname might have been replaced wholesale if an error occurred in the attempt to open the original pathname.

:truename

Message

Returns the pathname of the file actually open on this stream. This can be different from what **:pathname** returns because of file links, logical devices, mapping of "newest" version to a particular version number, and so on. For some systems (such as ITS) the truename of an output stream is not meaningful until after the stream has been closed, at least on an ITS file server.

:length

Message

Returns the length of the file, in bytes or characters. For text files on PDP-10 file servers, this is the number of PDP-10 characters, not Symbolics characters. The numbers are different because of character-set translation. (See the section "The Character Set", page 5.) For an output stream the length is not meaningful until after the stream has been closed, at least on an ITS file server.

:characters *Message*

Returns **t** if the stream is a character stream, **nil** if it is a binary stream.

:creation-date *Message*

Returns the creation date of the file, as a number which is a universal time. See the section "Dates and Times" in *Programming the User Interface*. See the function **fs:directory-list**, page 117.

:info *Message*

Returns a cons of the truename and creation date of the file. The creation date is a number that is a universal time. This can be used to tell if the file has been modified between two **opens**. For an output stream the info is not meaningful until after the stream has been closed, at least on an ITS file server.

:delete *Message*

Deletes the file open on this stream. The file does not really go away until the stream is closed. You should not use **:delete**. Instead, use **deletef**.

:rename *new-name* *Message*

Renames the file open on this stream. You should not use **:rename**. Instead, use **renamef**.

:properties *Message*

Returns two values:

- A list whose car is the pathname of the file and whose cdr is a list of the properties of the file; thus the element is a "disembodied" property list and **get** can be used to access the file's properties.
- A list of what properties of this file are "changeable".

:change-properties *Message*

Changes the file properties of the file open on this stream.

:change-properties signals an error rather than returning one.

:finish *Message*

Does a **:force-output** to a buffered asynchronous device, such as the Chaosnet, then waits until the currently pending I/O operation has been completed. If the stream does not handle this, the default handler ignores it.

For file output streams, **:finish** finalizes file content. It ensures that all data have actually been written to the file, and sets the byte count. It converts non-direct output openings into append openings. It allows other users to access the data that have been written before the **:finish** message was sent.

File output streams implement the **:finish** and **:force-output** messages.

10.4 Direct Access File Streams

Direct access file streams are supported by LMFS. They are designed to facilitate reading and writing data from many different points in a file. They are typically used to construct files organized into discrete extents or *records*, whose positions within a file are known by programs that access them in nonsequential order. Although this could be done with the **:set-pointer** message to input file streams, the direct access facility provides the following additional functions:

- Direct access to output files.
- Bidirectional file streams, which allow interspersed reading and writing of data to and from varied locations in a file.
- No use of network connections or file buffers during the time between data reading and the next call to position. In contrast, using the **:set-pointer** message with ordinary ("sequential") input file streams incurs a significant network and data transfer overhead if the program repeatedly positions, reads several bytes, and then computes for a time.

10.4.1 Stream Messages

The following messages are relevant to direct access file streams.

:read-bytes *n-bytes file-position* *Message*

Sent to a direct access input or bidirectional file stream, this requests the transfer of *n-bytes* bytes from position *file-position* of the file. The message itself does not return any data to the caller. It causes the stream to be positioned to that point in the file, and the transfer of *n-bytes* bytes to begin. An EOF is sent following the requested bytes. The bytes can then be read using **:tyi**, **:string-in**, or any of the standard input messages or functions.

The stream enforces the byte limit, and presents an EOF if you attempt to read bytes beyond that limit. You must actually read all the bytes and read past (that is, consume from the stream) the EOF.

It is also possible, before all the bytes have been read, to perform stream operations other than reading bytes. For example, an application might read several records at a time, to optimize transfer and buffering, and decide, after reading the first record, to position somewhere else. Direct access file streams handle this properly. Nevertheless, network and buffering resources allocated to the stream (both on the local machine and server machine) are not freed unless all the requested bytes (of the last **:read-bytes** request) and the EOF following them are read.

If you request more bytes than remain in the file, you receive the remaining bytes followed by EOF.

```

;;; Opens a file for direct I/O
(with-open-file (f "saratoga:>foo.lisp" :direct t :direction :io))

(send f :read-bytes 10. 25.) ; Requests 10 bytes, starting at byte 25.

(send f :tyi) ; Gets one byte from the stream.
.
.
(defvar record (make-array 10. :type art-string)) ; A record buffer.

(send f :string-in eof-option record) ; Reads 10-character record.
(send f :tyi) ; Consumes the following EOF.
...

```

10.4.1.1 Direct Access Output File Streams

You create direct access output to output and bidirectional direct access file streams by sending a **:set-pointer** message to the stream, and beginning to write bytes using standard messages, such as **:tyo**, **:string-out**, and so forth. The bytes are written to the file starting at the location requested, at successive file positions. Although you can extend the file in this manner, you cannot do a **:set-pointer** to *beyond* the current end of the file.

Direct access output, therefore, consists of sequences of **:set-pointer** messages and data output. Data are not guaranteed to actually appear in the file until either the stream is closed or a **:finish** message is sent to the stream. See the message **:finish**, page 18.

10.4.1.2 Direct Access Bidirectional File Streams

Bidirectional direct access file streams combine the features of direct access input and output file streams. Sequences of **:read-bytes** messages and reading data can be interspersed with sequences of **:set-pointer** messages and writing data. The stream is effectively switched between "input" and "output" states by the **:read-bytes** and **:set-pointer** messages. You cannot read data with **:tyi** or similar messages if a **:set-pointer** message has been sent to the stream since the last **:read-bytes** message. Similarly, you cannot write data with **:tyo** or similar messages unless a **:set-pointer** message has been sent to the stream since the last **:read-bytes** or **:tyi** messages, or similar operation.

When the EOF of a byte sequence requested with a **:read-bytes** message has been read for a bidirectional stream, the system frees network and buffering resources.

11. Accessing Directories

To understand the functions in this section, it is imperative to have read some other documentation. See the section "Naming of Files", page 127.

11.1 Functions for Accessing Directories

fs:directory-list *pathname* &rest *options* *Function*

Finds all the files that match *pathname* and returns a freshly consed list with one element for each file. Each element is a list whose car is the pathname of the file and whose cdr is a list of the properties of the file; thus the element is a "disembodied" property list and **get** can be used to access the file's properties. The car of one element is **nil**; the properties in this element are properties of the file system as a whole rather than of a specific file.

The matching is done using both host-independent and host-dependent conventions. Any component of *pathname* that is **:wild** matches anything; all files that match the remaining components of *pathname* are listed regardless of their values for the wild component. In addition, there is host-dependent matching. Typically, this uses the asterisk character (*) as a wild-card character. A pathname component that consists of just a * matches any value of that component (the same as **:wild**). *, appearing in a pathname component that contains other characters, matches any character (on ITS) or any string of characters (on TOPS-20, LMFS, UNIX, and Multics) in the starred positions and requires the specified characters otherwise. Other hosts follow similar but not necessarily identical conventions.

The *options* are keywords that modify the operation. The following options are currently defined:

- :noerror** If a file-system error (for example, no such directory) occurs during the operation, an error is normally signalled and the user is asked to supply a new pathname. However, if **:noerror** is specified and an error occurs, an error object describing the error is returned as the result of **fs:directory-list**. This is identical to the **:noerror** option to **open**.
- :deleted** This is for file servers with soft deletion, such as TOPS-20, LMFS, and FEP. It specifies that deleted (but not yet expunged) files are to be included in the directory listing. Normally, they are not included.

- :no-extra-info** This results in only enough information for listing the directory as in `Dired`.
- :sorted** This causes the directory to be sorted so that at least multiple versions of a file are consecutive in increasing version number.

The properties that might appear in the list of property lists returned by `fs:directory-list` are host-dependent to some extent. The following properties are defined for most file servers.

- :length-in-bytes** The length of the file expressed in terms of the basic units in which it is written (characters in the case of a text file and binary bytes for a binary file).
- :byte-size** The number of bits in a byte.
- :length-in-blocks** The length of the file in terms of the file system's unit of storage allocation.
- :block-size** The number of bits in a block.
- :creation-date** The date the file was created, as a universal time. This does not necessarily mean the time that the file itself was created, but rather, the time that the data in it were created. This property corresponds to the concept of "modification date" on many systems. See the section "Dates and Times" in *Programming the User Interface*.
- :directory** A boolean. If `t`, the object in question is a directory, as opposed to a file or a link. This property can only be returned as `t` in a hierarchical file system.
- :last-expunge-time** For directories, the date that the directory was last expunged. It is `nil` if the directory has never been expunged.
- :reference-date** The most recent date that the file was used, as a universal time.
- :author** The name of the person who created the data in the file, as a string.
- :account** A string. Highly system-dependent in format.
- :deleted** A boolean. `t` for a "deleted" file, in file systems supporting "soft deletion".
- :dont-delete** A boolean. If it is `t`, an error results if an attempt is made to delete the file.
- :dont-dump** A boolean. Suppresses backup dumping.

:dont-reap	A boolean. A flag used by directory maintenance tools.
:dumped	A boolean. t if and only if the file has been dumped to backup tape.
:generation-retention-count	A number that specifies how many versions of a file should be saved.
:link-to	A string. This is the target pathname of a link, as a string.
:offline	A boolean. The file has been moved to archival storage.
:physical-volume	A string.
:protection	A string.
:reader	A string. The last person to have read the file.
:temporary	A boolean.

fs:multiple-file-plists *filenames &rest options* *Function*

Returns a list of property lists, each property list corresponding to a file in **filenames**, which is a list of pathnames. For example:

```
(fs:multiple-file-plists
 (list "sys: doc; xm; xm1.sar" "sys: doc; xm; xm2.sar")) ==>

((#<LOGICAL-PATHNAME "SYS: DOC; XM; XM2.SAR.NEWEST"> :TRUENAME
 #<LMFS-PATHNAME "X:>current>sys>doc>xm>xm2.sar.19"> :LENGTH 46350
 :AUTHOR "BENSON" :BYTE-SIZE NIL :CREATION-DATE 2687555882)
 (#<LOGICAL-PATHNAME "SYS: DOC; XM; XM1.SAR.NEWEST"> :TRUENAME
 #<LMFS-PATHNAME "X:>current>sys>doc>xm>xm1.sar.17"> :LENGTH 33833
 :AUTHOR "BENSON" :BYTE-SIZE NIL :CREATION-DATE 2687544577))
```

fs:change-file-properties *pathname error-p &rest properties* *Function*

Some of the properties of a file can be changed, such as its creation date or its author. The properties that can be changed depend on the host file system; a list of the changeable property names is the **:settable-properties** property of the file system as a whole, returned by **fs:directory-list**. See the function **fs:directory-list**, page 117.

fs:change-file-properties changes one or more properties of a file. *pathname* names the file. The *properties* arguments are alternating keywords and values. If the *error-p* argument is **t**, a Lisp error is signalled. If *error-p* is **nil** and an error occurs, the error object is returned. If no error occurs, **fs:change-file-properties** returns **t**.

fs:file-properties *pathname &optional (error-p t)* *Function*

Returns a disembodied property list for a single file (compare this to **fs:directory-list**). The car of the returned list is the truename of the file

and the *cdr* is an alternating list of indicators and values. If *error-p* is **t** (the default) a Lisp error is signalled. If *error-p* is **nil** and an error occurs, the error object is returned.

fs:complete-pathname *defaults string type version &rest options* *Function*

string is a partially specified file name. (Presumably it was typed in by a user and terminated with the COMPLETE or END to request completion.)

fs:complete-pathname looks in the file system on the appropriate host and returns a new, possibly more specific string. Any unambiguous abbreviations are expanded in a host-dependent fashion.

string is completed relative to a default pathname constructed from *defaults*, the host (if any) specified by *string*, *type*, and *version*, using the function **fs:default-pathname**. See the function **fs:default-pathname**, page 162. If *string* does not contain a colon, the host comes from *defaults*; otherwise the host name precedes the first colon in *string*.

options are keywords (without following values) that control how the completion will be performed. The following option keywords are allowed. Their meanings are explained more fully below.

- :deleted** Look for files that have been deleted but not yet expunged. The default is to ignore such files.
- :read** or **:in** The file is going to be read. This is the default. The name **:in** is obsolete and should not be used in new programs.
- :write** or **:print** or **:out** The file is going to be written (that is, a new version is going to be created). The names **:print** and **:out** are obsolete and should not be used in new programs.
- :old** Look only for files that already exist. This is the default. **:old** is not meaningful when **:write** is specified.
- :new-ok** Allow either a file that already exists, or a file that does not yet exist. **:new-ok** is not meaningful when **:write** is specified. The **:new-ok** option is no longer used by any system software, because users found its effects (in the Zmacs command Find File (c-X c-F)) to be too confusing. It remains available, but programmers should consider this experience when deciding whether to use it.

The first value returned is always a string containing a file name; either the original string, or a new, more specific string. The second value returned indicates the status of the completion. It is non-**nil** if it was completely successful. The following values are possible:

- :old** The string completed to the name of a file that exists.
- :new** The string completed to the name of a file that could be created.
- nil** The operation failed for one of the following reasons:
- The file is on a file system that does not support completion. The original string is returned unchanged.
 - There is no possible completion. The original string is returned unchanged.
 - There is more than one possible completion. The string is completed up to the first point of ambiguity.
 - A directory name was completed. Completion was not successful because additional components to the right of this directory remain to be specified. The string is completed through the directory name and the delimiter that follows it.

Although completion is a host-dependent operation, the following guidelines are generally followed:

When a pathname component is left completely unspecified by *string*, it is generally taken from the default pathname. However, the name and type are defaulted in a special way described below and the version is not defaulted at all; it remains unspecified.

When a pathname component is specified by *string*, it can be recognized as an abbreviation and completed by replacing it with the expansion of the abbreviation. This usually occurs only in the rightmost specified component of *string*. All files that exist in a certain portion of the file system and match this component are considered. The portion of the file system is determined by the specified, defaulted, or completed components to the left of this component. A file's component *x* matches a specified component *y* if *x* consists of the characters in *y* followed by zero or more additional characters; in other words, *y* is a left substring of *x*. If no matching files are found, completion fails. If all matching files have the same component *x*, it is the completion. If there is more than one possible completion, that is, more than one distinct value of *x*, there is an ambiguity and completion fails unless one of the possible values of *x* is equal to *y*.

If completion of a component succeeds, the system attempts to complete any additional components to the right. If completion of a component fails, additional components to the right are not completed.

A blank component is generally treated the same as a missing component; for example, if the host is a LMFS, completion of the strings "foo" and "foo." deals with the type component in the same way. The strings are not completed identically; completion of "foo" attempts to complete the name component, but completion of "foo." leaves the name component alone since it is not the rightmost.

If *string* does not specify a name, then the name of the default pathname is *preferred* but is not necessarily used. The exact meaning of this depends on *options*:

- With the default options, if any files with the default name exist in the specified, defaulted, or completed directory, the default name is used. If no such files exist, but all files in the directory have the same name, that name is used instead. Otherwise, completion fails.
- With the **:write** option, the default name is always used when *string* does not specify a name, regardless of what files exist.
- With the **:new-ok** option, if any files with the default name exist in the specified, defaulted, or completed directory, the default name is used. If no such files exist, but all files in the directory have the same name, that name is used instead. Otherwise, the default name is used.

The special treatment of the case where all files in the directory have the same name is not very useful and is not implemented by all file systems.

If *string* does not specify a type, then the type of the default pathname is *preferred* but is not necessarily used. The exact meaning of this depends on *options*:

- With the default options, if a file with the specified, defaulted, or completed name and the default type exists, the default type is used. If no such file exists, but one or more files with that name and some other type do exist and all such files have the same type, that type is used instead. Otherwise, completion fails.
- With the **:write** option, the default type is always used when *string* does not specify a type, regardless of what files exist.
- With the **:new-ok** option, if a file with the specified, defaulted, or completed name and the default type exists, the default type is used. If no such file exists, but one or more files with that name and some other type do exist and all such files have the same type, that type is used instead. Otherwise, the default type is used.

In file systems such as LMFS and UNIX that require a trailing delimiter (>

or \wedge) to distinguish a directory component from a name component, the system heuristically decides whether the rightmost component was meant to be a directory or a name, and inserts the directory delimiter if necessary.

If *string* contains a relative directory specification for a host with a hierarchical file system, it is assumed to be relative to the directory in the default pathname and is expanded into an absolute directory specification.

The host and device components generally are not completed; they must be fully specified if they are specified at all. This might change in the future.

If *string* does not specify a version, the returned string does not specify a version either. This differs from file name completion in TOPS-20; TOPS-20 completes an implied version of "newest" to a specific number. This is possible in TOPS-20 because completing a file name also attaches a "handle" to a file. In Symbolics-Lisp, the version number of the newest file might change between the time the file name is completed and the time the actual file operation (open, rename, or delete) is performed.

A pathname component must satisfy the following rules in order to appear in a successful completion:

- The host, device, and directory must actually exist.
- The name must be the name of an existing file in the specified directory, unless **:write** or **:new-ok** is included in *options*.
- The type must be the type of an existing file with the specified name in the specified directory, unless **:write** or **:new-ok** is included in *options*.
- A pathname component always completes successfully if it is **:wild**.

When the rules are not satisfied by a component taken from the default pathname, completion fails and that component remains unspecified in the resulting string. When the rules are not satisfied by a component taken from *string*, completion fails and that part of *string* remains unchanged (other components of *string* can still be expanded).

listf *path* &optional (*output-stream standard-output*) *Function*

listf is a function for displaying an abbreviated directory listing. The default for name, type, and version of *path* is **:wild**.

```
(listf "f:>jwalker>mit-220")
```

The format of the listing varies with the operating system.

PART II.

Files

12. Naming of Files

A Symbolics computer generally has access to many file systems. While it can have its own file system on its own disks, a community of Lisp Machine users often has many shared file systems accessible by any of the Symbolics computers over a network. These shared file systems can be implemented by any computers that are capable of providing file system service. A *file server computer* might be a special-purpose computer that does nothing but service file system requests from computers on a network, or it might be an existing timesharing system.

Programs, at the behest of users, need to use names to designate files within these file systems. The main difficulty in dealing with names of files is that different file systems have different naming conventions and formats for files. For example, in the UNIX system, a typical name looks like:

```
/usr2/george/foo.bn
```

In this example, `/usr2/george` is the *directory name*, `foo` is the *file name* and `bn` is the *file type*. However, in TOPS-20, a similar file name is expressed as follows:

```
PS:<GEORGE>FOO.BIN
```

It would be unreasonable for each program that deals with file names to be expected to know about each different file name format that exists; in fact, new formats could be added in the future, and existing programs should retain their abilities to manipulate files in a system-independent fashion.

The functions, flavors, and messages described in this chapter exist to solve this problem. They provide an interface through which a program can deal with files and manipulate them without depending on their syntax. This lets a program deal with multiple remote file systems simultaneously, using a uniform set of conventions.

12.1 Pathnames

All file systems dealt with by the Symbolics computer are mapped into a common model, in which files are named by a conceptual object called a *pathname*. The Symbolics computer system, in fact, represents pathnames by objects of flavor **fs:pathname**, and the flavors built upon it. A pathname always has six conceptual components, described below. These components provide the common interface that allows programs to work the same way with different file systems; the mapping of the pathname components into the concepts peculiar to each file system is taken care of by the pathname software. This mapping is described elsewhere for each file system. See the section "The Character Set", page 5.

The following are the conceptual components of a pathname. They will be clarified by examples below.

Host	The computer system, the machine, on which the file resides.
Device	Corresponds to the "device" or "file structure" concept in many host file systems. Often, it designates a group of disks, or removable storage media, or one of several different media of differing storage densities or costs.
Directory	An organizational structure in which files are "contained" on almost all file systems. Files are "stored in", or "reside in" directories. The directories have names; the files' names are only valid within the context of a given directory. Some systems (<i>hierarchical</i> file systems) allow directories to be contained in other directories; others do not.
Name	The name of a group of files that can be thought of as conceptually the "same" file. In many systems, this is the "first name" of the file. For instance, source and object files for the same program generally have the same <i>name</i> , but differing <i>type</i> .
Type	Corresponds to the "filetype" or "extension" concept in many host file systems. This usually indicates the kind of data stored in the file, for example, binary object code, a Lisp source program, a FORTRAN source program, and so forth.
Version	<p>Corresponds to the "version number" concept in many host file systems. Some systems implement this concept, others do not. A version number is a number, part of the conceptual name of the file, that distinguishes succeeding versions of a file from each other. When a user of such a file system writes out a file he or she does not modify the file on the host computer but writes a <i>new version</i>, that is, one with a higher version number, automatically.</p> <p>The Symbolics computer system allows a version component of "newest" or "oldest", represented by the keyword symbols :newest and :oldest, respectively, to designate "the newest (oldest) version of the file, whichever that might be".</p>

As an example, consider a TOPS-20 user named "George", who writes a Lisp program that he thinks of as being named "conch". If George uses the TOPS-20 host named FISH, the source for his program might be in a file on the host FISH with the following name:

```
<GEORGE>CONCH.LISP.17
```

In this case, the host is FISH, the device would be some appropriate default, and the directory would be <GEORGE>. This directory would probably contain a number of files related to the "conch" program. The source code for this program would live in a file with name CONCH, type LISP, and versions 1, 2, 3, and so on. The compiled form of the program would live in a file named CONCH with type BIN.

Now suppose George is a UNIX user, using the UNIX host BIRD. The source for his program would probably be in a file on the host BIRD with the following name:

```
/usr2/george/conch.l
```

In this case, the host is BIRD, and the directory would be /usr2/george. This directory would probably contain a number of files related to the "conch" program. The source code for this program would live in a file with name conch, type l. The compiled form of the program would live in a file named conch, with type bn. There are no version numbers on UNIX.

Note that a pathname is not necessarily the name of a specific file. Rather, it is a way to get to a file; a pathname need not correspond to any file that actually exists, and more than one pathname can refer to the same file. For example, the pathname with a version of "newest" will refer to the same file as a pathname with the same components except a certain number as the version. In systems with links, multiple file names, logical devices, and so forth, two pathnames that look quite different can turn out to designate the same file. To get from a pathname to a file requires doing a file system operation such as **open**.

12.1.1 Simple Usage of the Pathname System

The pathname system can be very easy to use if you know a few simple techniques. It often seems that there are many different ways to do anything, and that only one of the is right for any circumstance, but most of these features only exist for special needs. This section shows you how to easily do some of the simple things.

12.1.1.1 Getting a Filename From the User

The simplest and most common application for using a pathname is simply to read or write a file. For example, a program to do some *very* simple processing of a database (it reads the file and ignores it):

```
(defun process-example-database (database-pathname)
  (with-open-file (database-stream database-pathname)
    (format t "~&Ignoring database ~A ..." (send database-stream :truename))
    (stream-copy-until-eof stream #'si:null-stream)
    (format t " ignored.~%")))
```

This simple example is adequate for a program interface, but for a user, it is rather awkward. The user must supply all components of the pathname, plus the quotation marks around the strings. Also, the user has no completion available. In this example, the user does not have to parse the pathname; **open** will do that for him. (Sometimes we won't be so lucky).

The user's job can be made easier by providing a function to read a pathname and pass it to **process-example-database**. To do this, **prompt-and-read** is used. See the function **prompt-and-read** in *Programming the User Interface*.

In our first version, we will just ask the user for the pathname.


```
(defun run-example ()
  (let ((pathname (prompt-and-read :pathname "Where is your database? ")))
    (process-example-database pathname)))
```

```
Where is your database? Y:>user>databases>dummy.database
Ignoring Y:>user>databases>dummy.database.7 ... ignored.
```

prompt-and-read does much of what we are looking for. It provides the following:

- Prompting, including reprompting when the user types `refresh`
- Parsing
- Completion
- Merging with defaults

In this case, we supplied no default, so the "default default", **fs:*default-pathname-defaults*** is used. But this default is not very helpful to the user, because it is not visible; it could even be confusing if the user expected one default and got another. Good practice dictates telling the user what the default is. **prompt-and-read** makes this easy with the **:visible-default** suboption to **:pathname**, **:pathname-or-nil**, and **:pathname-list**.

```
(defun run-example ()
  (let ((pathname (prompt-and-read
                   '(:pathname :visible-default ,fs:*default-pathname-defaults*)
                   "Where is your database? ")))
    (process-example-database pathname)))
Where is your database? (Default Y:>user>foo.lisp) databases>dummy.database
Ignoring Y:>user>databases>dummy.database.7 ... ignored.
```

Now that the user can see the defaults, he or she can make use of them. Note that in the above example, the user did not have to type the "Y:>user>", because the default was available.

Tailoring Pathname Defaults

fs:*default-pathname-defaults* is a global default, with nothing particularly appropriate to any specific application. Often, when an application is writing or reading a file, it knows more about the file than is implied by **fs:*default-pathname-defaults***. This information can be used to help prompt the user for a suitable filename and help reduce the amount of typing needed to specify a suitable filename.

For example, consider our example of reading a database. (See the section "Reading

a Filename".) In this example, we are just prompting for the filename and ignoring the actual database.

```
(defun run-example ()
  (let ((pathname (prompt-and-read
                   '(:pathname :visible-default ,fs:*default-pathname-defaults*)
                   "Where is your database? ")))
    (process-example-database pathname)))
Where is your database? (Default Y:>user>foo.lisp) databases>dummy.database
Ignoring Y:>user>databases>dummy.database.7 ... ignored.
```

First, if we are going to seriously use our own special file type, we need to define the type so that it can be used successfully on different systems. See the function **fs:define-canonical-type**, page 163.

```
(fs:define-canonical-type :database "DATABASE"
  ((:vms :vms4) "DBS")
  (:unix "DB"))
```

Now this type can be used as the default type for our example databases.

```
(defun run-example ()
  (let* ((default (fs:default-pathname fs:*default-pathname-defaults*
                                     nil ;Host
                                     :database)) ;Type
        (pathname (prompt-and-read '(:pathname :visible-default ,default)
                                   "Where is your database?~%")))
    (process-example-database pathname)))

Where is your database? (Default Y:>user>foo.database) databases>dummy
Ignoring Y:>user>databases>dummy.database.7 ... ignored.
```

12.1.1.2 More About Defaults

Most simple programs use **fs:*default-pathname-defaults*** as the source for their defaults. However, as a program makes more use of pathname reading and defaults, there are some things we can do to make things easier for the user.

- Provide a default based on other files in an operation, for example, defaulting an output file pathname from the input file.
- Provide "sticky" defaults, where the new default is based on the last file the user gave.
- Provide a default based on the current context, as in "pathname of the current buffer" in Zmacs.

Defaulting an Output File Pathname From an Input File

Perhaps the most common defaulting situation is that of defaulting an output file pathname from the input file. Usually, the output file differs from the input file only in file type and version, and we would like to have the user provide explicit information only when his or her desires differ from the usual case.

```
(defun my-compile-file (input-file output-file)
  (format t "~&Compiling ~A into ~A.~%"
          input-file output-file)
  (compiler:compile-file input-file output-file))

(defun comp-it ()
  (let* ((input-default (fs:default-pathname nil nil :lisp :newest))
         (input-file (prompt-and-read
                      '(:pathname :visible-default ,input-default)
                      "Input file: "))
         (output-default (fs:default-pathname input-file nil :bin :newest))
         (output-file (prompt-and-read
                      '(:pathname :visible-default ,output-default)
                      "Output file: ")))
    (my-compile-file input-file output-file)))
```

The above example works well for single files, but it does not handle wildcards. To handle wildcards, we need to introduce the use of **fs:translate-wild-pathname** and **fs:directory-link-opaque-dirlist**. **fs:translate-wild-pathname** does the work of interpreting how a given input file is to be mapped to its corresponding output file, and **fs:directory-link-opaque-dirlist** takes care of finding all the input files.

Note that we use **fs:directory-link-opaque-dirlist** rather than **fs:directory-list**. In general, this is necessary whenever the **fs:translate-wild-pathname** message is used. **fs:translate-wild-pathname** expects the input pathname to match the input pattern. **fs:directory-list**, in the presence of directory links or VAX/VMS logical devices, can have a different directory or a different device.

If the input pattern has wildcards in its directory component, **fs:directory-link-opaque-dirlist** currently does no better than **fs:directory-list**. This is a difficult problem still under investigation.

```
(defun comp-one-file (input-file-pattern output-file-pattern input-file)
  (let ((output-file (send input-file-pattern :translate-wild-pathname
                          output-file-pattern input-file)))
    (my-compile-file input-file output-file)))
```

```

(defun comp-files ()
  (let* ((input-default (fs:default-pathname nil nil :lisp :newest))
        (input-pattern (prompt-and-read
                        '(:pathname :visible-default ,input-default)
                        "Input file: "))
        (output-default (fs:default-pathname input-file nil :bin :newest))
        (output-pattern (prompt-and-read
                        '(:pathname :visible-default ,output-default)
                        "Output file: ")))
    (if (not (send input-file :wild-p))
        (comp-one-file input-pattern output-pattern input-pattern)
        (loop for (file) in (cdr (fs:directory-link-opaque-dirlist
                                input-pattern :fast))
              do (comp-one-file input-pattern output-pattern file))))))

```

Note that in the above example, we just call **comp-one-file** directly if the input pathname is not wild. While it is not strictly necessary to do this (**fs:directory-link-opaque-dirlist** works on non-wildcard pathnames), it does eliminate an unneeded operation.

Sticky Pathname Defaults

Often, when a single command or a related set of commands are to be repeated, the next command should operate on a file related to the one the current command is operating on. In this case, it would be most convenient for the default to be the previous pathname. This is called *sticky defaulting*.

For example, consider a simple user-written tool to either show or delete files.

```

(defun show-or-delete ()
  (loop with default = (fs:default-pathname)
        for ch = (prompt-and-read :character "Cmd>")
        do (multiple-value-bind (prompt function)
            (selector char-equal ch
                      (#\S (values "Show File" #'viewf))
                      (#\D (values "Delete File" #'deletef))
                      (#\Q (return nil))
                      (#\Help (format t "~&S = Show File~@
                                     D = Delete File~@
                                     Q = Quit~%"
                                     ch)))
            (values nil nil))
          (otherwise
            (tv:beep)
            (format t "~&~:C is an unknown command.~%" ch)))
        (when prompt
          (let ((file (prompt-and-read
                      '(:pathname :visible-default ,default)
                      prompt)))
            ; The following is done for us by prompt-and-read
            ;(setq default (fs:merge-pathnames file default))
            (funcall function file))))))

```

Each time around the loop, when the user specifies a file, it is remembered to serve as the default the next time around. Note the commented out **(setq default (fs:merge-pathnames file default))**. This isn't needed in this example, since **prompt-and-read** does this for us, but if we were reading pathnames via some other mechanism, it is important to keep the default as a fully specified pathname. Otherwise, the second time around the loop, we could end up with defaults like "Q:", which is not of much use if the user is then forced to type all the components of the pathname and may get an error if he or she does not.

If you wish to use a default such as this and not keep it in a local variable, you should use a defaults alist. This serves as a registered place to remember a pathname, so that if the world is moved to another site, it can be reset. Defaults alists can be passed to **fs:default-pathname** to extract a fully-merged default. See the function **fs:set-default-pathname**, page 163. See the function **fs:make-pathname-defaults**, page 162.

Pathname Defaulting From the Current Context

Often, an application program involves the user working on a single context for an extended time. For example, in the editor, the user is working on a single named buffer. In the font editor, the user is working on a single named font.

Often, the object being worked on was read in from a file. This file can serve as a default for further file operations, such as listing the directory, or resaving the object. Consider a picture editor, which lets the user edit multiple pictures, as the Zmacs editor lets the user edit multiple buffers. This picture editor stores its files in .BIN files.

```

(defflavor picture (name
                    (pathname sys:fdefine-file-pathname)
                    (array (make-array '(100. 100.) :type 'art-1b)))
  ()
  :gettable-instance-variables
  :settable-instance-variables
  :initable-instance-variables)

(defvar *pictures* nil
  "List of pictures being edited")

(defvar *current-picture* nil)

(defvar *picture-defaults* (fs:make-pathname-defaults))

(defun add-picture (picture)
  (setq *pictures* (del #'(lambda (p1 p2)
                           (string-equal (send p1 :name) (send p2 :name)))
                       picture
                       *pictures*))
  (push picture *pictures*)
  (setq *current-picture* picture))

(defmethod (picture :fasd-form) ()
  '(make-instance ',(typep self)
                  :name ',name
                  :array ',array))

(defun picture-default-pathname (&key type (version :newest))
  (let ((bare-default (fs:default-pathname *picture-defaults*
                                           nil type version))
        (path (when *current-picture*
                  (send *current-picture* :pathname))))
    (if (not *current-picture*)
        bare-default
        (if path
            (setq path (fs:merge-pathnames path bare-default version))
            ;; A new picture, so no pathname. Let's make a guess from the name.
            (let ((name (send *current-picture* :name)))
              (setq path
                    (condition-case ()
                      (fs:merge-pathnames name bare-default version)
                      ;; If name isn't parsable, just use the bare default.
                      (error bare-default))))))
        path)))

(defun com-create-picture ()
  (let ((name (prompt-and-read :string "Picture name: ")))
    (add-picture (make-instance 'picture :name name))))

```

```
(defun com-save-picture ()
  (let* ((default (picture-default-pathname :type :bin))
        (file (prompt-and-read
                '(:pathname :visible-default ,default)
                "Save to picture file: ")))
    ;; Remember the pathname given, so the next time we
    ;; get a new picture, we can have a better default.
    (fs:set-default-pathname file *picture-defaults*)
    (sys:dump-forms-to-file
     file
     '(add-picture ',*current-picture*))))
```

In this example, **picture-default-pathname** computes the default. If the current picture has a file associated with it, that serves as the default. If there is no pathname with the current picture, we attempt to make a pathname using the name. If that fails (or if there is no current picture), we just use the bare default.

Finally, the pathname we read is remembered, so the next time a default is needed for a new picture, we will have a more recent default.

Note that when the picture is loaded, **sys:fdefine-file-pathname** is used to get the file being loaded. This works well when the file being loaded is a .bin file, since **load** binds this variable. However, in other situations, you need to make other arrangements to set the pathname.

12.1.2 Host Determination in Pathnames

Two important operations of the pathname system are *parsing* and *merging*. Parsing is the conversion of a string, which might have been typed by the user when asked to supply the name of a file, into a pathname object. This involves finding out for which host the pathname is intended, using the file name syntax conventions of that host to parse the string into the standard pathname components, and constructing such a pathname. Merging is the operation that takes a pathname with missing components and supplies values for those components from a set of defaults.

Since each kind of file system has its own character string representation of names of its files, there has to be a different parser for each of these representations, capable of examining such a character string and determining the value of each component. The parsers, therefore, all work differently. How does the parsing operation know which parser to use? It determines for which host the pathname is intended, and uses the appropriate parser. A filename character string can specify a host explicitly, by having the name of the host, followed by a colon, at the beginning of the string, or it can assume a default, if there is no host name followed by a colon at the beginning of the string.

Here is how the pathname system determines for which host a pathname being

parsed is intended. The first colon in a pathname being parsed *always* delimits the host name. You can also enter pathname strings that are for a specific host and do not contain any host name. In that case, a *default host* is used. Normally, the identity of the default host is displayed to the user entering a pathname. See the section "Pathname Defaults and Merging", page 147.

However, pathnames can have colons in them that do not designate hosts, such as filenames constructed from clock times, and the like. Some systems use the colon character to delimit devices. This creates a problem in parsing such pathnames. See the function **fs:parse-pathname**, page 156. The standard Symbolics computer user interface does not use such pathnames, but they can be used by other programs, particularly those that deal with files whose format is defined by a foreign operating system.

The rule for parsing file names containing colons is, again, that any string used before a colon is *unconditionally* interpreted as a file computer. If the string cannot be interpreted as a host, an error is signalled.

If you must type a pathname that has an embedded colon *not* meaning a host, you omit the host and place a colon at the beginning of the string. This "null host" tells the parser that it should *not* look further for a colon, but instead assume the host from the defaults. Examples:

- SS:<FOO>BAR refers to a host named "SS". :SS:<FOO>BAR refers to no explicit host; if parsed relative to a TOPS-20 default, "SS" probably refers to a device.
- 09:25:14.data refers to a host named "09". :09:25:14.data refers to no explicit host.
- AI: COMMON; GEE WHIZ refers to a host named "AI".
- AI: ARC: USERS1; FOO BAR refers to a host named "AI". "ARC" is the name of a device in the ITS operating system.
- EE:PS:<COMMON>GEE.WHIZ.5 specifies host EE (TOPS-20).
- PS:<COMMON>GEE.WHIZ.5 specifies a host named PS, which is almost certainly not what is intended! The user probably intended the "PS" device on some TOPS-20 host.
- :PS:<COMMON>GEE.WHIZ.5, assuming that the default host is some TOPS-20, specifies a device named "PS" on that host.

There are a few "pseudohost" names, which are recognized as host names even though they are not actually the names of hosts:

"local" This pseudohost name always refers to the local file system

- (LMFS) of the machine that you are using. It does not matter whether or not a local file system actually exists on that machine; an attempt will be made to reference it. "Local" is always equivalent to the name of the local host.
- "FEP" This pseudohost name always refers to a FEP (front-end processor) file system on the machine you are using, specifically, the one on the disk unit from which the system was booted.
- "FEP n " This pseudo name always refers to a FEP file system on the machine you are using. The single digit n specifies the disk unit number; there is a separate FEP file system on each drive. This can access the boot unit, or any other disk unit, when multiple units are present.
- "*host*|FEP n " *host* must be a valid host name. This pseudohost name refers to a FEP file system on a remote 3600-family computer. The syntax "*host*|FEP" is not acceptable: you cannot access the "boot unit" of a remote machine in this fashion. You must know the disk unit number. The disk unit number of a host having only one disk unit is 0.

If the string to be parsed does not specify a host explicitly, the parser assumes that some particular host is the one in question, and it uses the parser for that host's file system. The optional arguments passed to the parsing function (`fs:parse-pathname`) tell it which host to assume.

12.1.3 Interning of Pathnames

Pathnames, like symbols, are *interned*. This means that there is only one pathname object with a given set of components. If a character string is parsed into components, and some pathname object with exactly those components already exists, then the parser returns the existing pathname object rather than creating a new one. The main reason for this is that a pathname has a property list. See the section "Property Lists" in *Reference Guide to Symbolics-lisp*. The system stores properties on pathnames to remember information about the file or family of files to which that pathname refers. (In fact, some of the *properties* stored on a generic pathname come from the file's *attribute* list when the file is edited or loaded, so they can be retrieved later without having to perform I/O on the file.) So you can parse a character string that represents a filename, and then look at its property list to get various information known about that pathname. The components of a pathname are never modified once the pathname has been created, just as the print name of a symbol is never modified. The only thing that can be modified is the property list.

When using property lists of pathnames, you have to be very careful which pathname you use to hold properties, in order to avoid a subtle problem: many different pathnames can refer to the same file, because of the `:newest` component,

file system links, multiple naming in the file system, and so on. If you put a property on one of these pathnames because you want to associate some information with the file itself, somebody else might look at another pathname that refers to the same file, and not find the information there. If you really want to associate information with the file itself rather than some particular pathname, you can get a canonical pathname for the file by using the **:truename** message to a stream opened to that file. See the message **:truename**, page 113. You might also want to store properties on "generic" pathnames. See the section "Generic Pathnames", page 149.

12.1.4 Printing Pathnames

A pathname can be converted back into a string, which is in the file name syntax of its host's file system. Although such a string (the *string for host*) can be produced from a pathname (by sending it the **:string-for-host** message), we discourage this practice. The Lisp Machine user interface prefers a string called the *string for printing*, which is the same as the string for host, except that it is preceded by the host name and a colon. This leaves no ambiguity about the host on which the file resides, when seen by a user. It is also capable of being reparsed, unambiguously, back into a pathname. **prin1** of a pathname (**~S in format**) prints it like a Lisp object (using the usual "#<" syntax), while **princ** of a pathname (**~A in format**) prints the string for printing. The **string** function, applied to a pathname, also returns the string for printing.

Not all the components of a pathname need to be specified. If a component pathname is missing, its value is **nil**. Before a file server can utilize a pathname to manipulate or otherwise access a file, all the pathname's missing components must be filled in from appropriate defaults. Pathnames with missing components are nevertheless often passed around by programs, since almost all pathnames typed by users do not specify all the components explicitly. The host is not allowed to be missing from any pathname; since the behavior of a pathname is host-dependent to some extent, it has to explicitly designate a host. Every pathname has a host attribute, even if the string that was parsed to create it did not specify one explicitly.

All pathname parsers support the cross-system convention that the double-shafted arrow character (↔) can be used to specify a null directory, name, type, or version component explicitly. Thus, for LMFS or TOPS-20, you can type the following:

```
↔.↔.5
```

This example specifies a version of 5, but no name or type. This is useful when typing against the default and attempting to change just the version of that default.

The keyword symbol **:unspecific** can also be a component of a pathname. This means that the component is not meaningful on the type of file system concerned. For example, UNIX pathnames do not have a concept of "version", so the version

component of every UNIX pathname is **:unspecific**. When a pathname is converted to a string, **nil** and **:unspecific** both cause the component not to appear in the string. The difference occurs in the merging operation, where **nil** is replaced with the default for that component, while **:unspecific** is left alone.

The special symbol **:wild** can also be a component of a pathname. This is only useful when the pathname is being used with a directory listing primitive such as **fs:directory-list** or **fs:all-directories**, where it means that this pathname component matches anything. See the function **fs:directory-list**, page 117. The printed representation of a pathname usually designates **:wild** with an asterisk; however, this is host-dependent.

:wild is one of several possible *wildcard* components, which are given to directory-listing primitives to filter file names. Many systems support other wildcard components, such as the string "foo*". This string, when supplied as a file name to a directory list operation on any of several system types, specifies all files whose name starts with "foo". In other contexts, it might not represent a wildcard at all. The component **:wild** matches all possible values for any component for which it appears. Other wildcard possibilities for directories exist, but they are more complicated, and are explained elsewhere. See the section "Values of Pathname Components", page 140. See the section "Directory Pathnames and Directory Pathnames as Files", page 142.

12.1.5 Values of Pathname Components

The set of permissible values for components of a pathname depends, in general, on the pathname's host. However, in order for pathnames to be usable in a system-independent way certain global conventions are adhered to. These conventions are stronger for the type and version than for the other components, since the type and version are actually understood by many programs, while the other components are usually treated as things chosen by the user that need to be preserved and passed around.

Most programs do not use or specify the components of a pathname explicitly, or only in a very limited way. In this way, they can remain operating-system-independent, while letting the pathname system take care of most issues of compatibility. In general, you should avoid where possible using specific values of pathname components in your programs. The descriptions here are illustrative but not complete, and programs should be written to expect component values other than those given here.

It is important to remember that not all pathname flavors accept all the values indicated here. For example, UNIX pathnames accept a type or version of **:unspecific**; few other pathnames do. Some systems do not allow certain characters or limit certain fields to a certain length.

It is generally *not* possible to simply copy components from one flavor of pathname to

another. It is often necessary to perform substitutions in order to produce a legal pathname. The **:new-default-pathname** message can be used instead of **:new-pathname** to get this substitution where necessary. The **:new-default-pathname** message attempts to substitute something as close as possible in meaning to the original component; however, the substitution can be arbitrary if necessary. For this reason, it is better to avoid copying components between pathnames of differing flavor, where possible.

The type is always a string (unless it is one of the special symbols **nil**, **:unspecific**, or **:wild**). Many programs that deal with files have an idea of what type they want to use. For example, Lisp source programs are "lisp", compiled Lisp programs (on, for example, a LMFS host) are "bin", text files are "text", and so on. The set of characters allowed in the type, and the number of characters, are system-dependent. In order to process file types in a system-independent way, the *canonical type* mechanism has been devised. A canonical type is a system-independent keyword symbol representing the conceptual type of a file. For instance, a Lisp source file on VMS has a file type of "LSP", and one on UNIX has a file type of "l". When we ask pathnames of either of these natures for their canonical type, we receive the keyword symbol **:lisp**. See the section "Canonical Types in Pathnames", page 151.

The version is either a number (specifically, a positive fixnum), or one of the symbols **nil**, **:unspecific**, **:wild**, **:newest**, or **:oldest**. **nil**, **:unspecific**, and **:wild** have been explained above. **:newest** refers to the largest version number that exists when reading a file, or that number plus one when writing a new file. **:oldest** refers to the smallest existing version number.

The host component of a pathname is always a host object. See the section "Namespace System Host Objects" in *Networks*.

The device component of a pathname can be one of the symbols **nil** or **:unspecific**, or a string designating some device, for those file systems that support such a notion.

The file name can be **nil** or a string, or **:wild**.

The directory component is highly system-specific. While it can be **nil** for any type of host, values designating actual directories, or partially wild specifications for directories, are more complicated. On nonhierarchical file systems, the directory component is usually a string such as "LMDOC", designating the name of the directory.

On hierarchical file systems, the directory component, when not **nil**, is a list of *directory level components*. For example:

<i>LMFS pathname</i>	<i>Directory component</i>
>sys>io>qfile.lisp.2357	("sys" "io")

"sys" and "io" are the directory level components. Since the "root directory" of

hierarchical file systems has no directory level components, it would be represented as **nil**, but this is impermissible, since **nil** already means that the directory component has not been specified. Thus, **:root** is used as the directory component in that case.

Relative pathnames on hierarchical file systems are represented by directory components having the level component **:relative**, followed by a number of occurrences of the symbol **:up** equal to the number of "upward relativization symbols", followed by the remaining directory level components. For example:

<i>LMFS relative pathname</i>	<i>Directory component</i>
<<x>y>z.lisp	(:relative :up :up "x" "y")

Directory components of pathnames for hierarchical file systems, on some systems, can also have the symbol **:wild** or a partially wild string (such as "foo*") as directory level components, to do level-by-level matching of level components. Also, on some systems, the level component **:wild-inferiors** (which is printed as "***" on LMFS and logical pathnames, and "..." on VMS, currently the only ones supporting it) to designate "any number, including zero of directory levels" to a directory list operation.

Note that some systems (currently VMS) do not allow using zero directory levels to denote their root directory. In this case, **:wild-inferiors** cannot stand alone, but must follow some other directory spec. For example: "[FOO...]" or "[*...]".

12.1.6 Directory Pathnames and Directory Pathnames as Files

In almost all systems having hierarchical directories, and certainly all the ones supported by the Symbolics computer as file server systems, the internal implementation of directories is as special files, known about by the operating system. The data in these files is not accessible to the user except through the defined operating system interfaces for dealing with directories.

Typically, listings of the contents of directories on hierarchical directory systems display names of both files and directories contained in the listed directory (as well as of links, on systems that support links).

Directories on hierarchical directory systems and files thus some things have in common. Appearing in directories is one. Another is that directories can usually be renamed, as can files, or, when the appropriate restrictions of the operating system are met (for instance, being empty), deleted. You can ask about the properties of a directory, or change some of them, with **fs:file-properties** and **fs:change-file-properties**, respectively, just as you do with a file.

Using LMFS as an example, consider the directory named "bar", which is contained in the directory named "foo", which itself is contained in the ROOT. A file in this directory named "tables.lisp.6" would have the following pathname:

```
>foo>bar>tables.lisp.6
```

The directory in which it is contained, bar, has the following pathname:

```
>foo>bar.directory.1
```

The file type of a directory, on LMFS, is "directory", and the version number of all directories is 1. The file types of directories, and their versions, if appropriate, vary among operating systems. If you wanted to rename, delete, or deal with the properties of the directory bar, you would have to present the above filename for this directory. A pathname of this type, which names a directory, as though it were a file, is called a *directory pathname as file*.

Directory pathnames as files are appropriate only to systems with hierarchical directories. On other systems, you cannot address directories directly.

The most common use of directories, however, is to reference files in them. The following pathname mentions the directory "bar" in this way:

```
>foo>bar>tables.lisp.6
```

This filename, when parsed into a pathname for the appropriate LMFS host, has a name component of "tables", a type component of "lisp", a version component of 6, and a directory component (in fact ("**foo**" "**bar**")) that designates the directory bar, inferior of foo, inferior of the ROOT. Such a pathname, which designates a given directory via its directory component, is called a *pathname as directory* for that directory. Of course, since the file name, type, and version are irrelevant to the specification of the directory, it is only one of many possible "pathnames as directory" for the directory bar.

The concept of pathname as directory is more general than the concept of directory pathname as file, since directories on nonhierarchical systems be described by their pathnames as directories as well. For instance, the following TENEX pathname, which describes a file in the "LMDOC" directory, is a pathname as directory for the LMDOC directory:

```
<LMDOC>CHFILE.TEXT;7
```

Note, also, that any pathname whose directory component is not **nil** is a pathname as directory for *some* directory.

Therefore, the Symbolics-Lisp primitives and operations that deal with directories explicitly (for example, **fs:expunge-directory** and **fs:all-directories**) expect pathnames of directories to be represented in the "pathname as directory" form. It is the canonical, system-independent way to represent pathnames of directories in the Symbolics system.

The following two messages convert between directory pathnames as files and pathnames as directories:

:directory-pathname-as-file of **fs:pathname** *Method*

Every pathname whose directory component is not **nil** is a pathname as directory for *some* directory. This method returns the directory pathname as file for that directory.

```
(setq p (fs:parse-pathname "Quabbin:>sys>lmfs>fsstr.lisp.243"))
#<LMFS-PATHNAME "Q:>sys>lmfs>fsstr.lisp.243">
(send p ':directory-pathname-as-file)
#<LMFS-PATHNAME "Q:>sys>lmfs.directory.1">
```

:pathname-as-directory of **fs:pathname** *Method*

This method is intended to be sent to a pathname that is the valid directory pathname as file for some directory. It produces one of many possible pathnames as directory for that directory, namely, the one whose name, type, and version are all **nil**.

```
(setq p1 (fs:parse-pathname "Quabbin:>sys>io.directory.1"))
#<LMFS-PATHNAME "Q:>sys>io.directory.1">
(setq p2 (send p1 ':pathname-as-directory))
#<LMFS-PATHNAME "Q:>sys>io">
(send p2 ':directory-pathname-as-file)
#<LMFS-PATHNAME "Q:>sys>io.directory.1">
```

If you are used to other systems' file-naming conventions, you may be confused by pathnames that have real directory components, but no name, type, or version. When typed in or printed, they look like the following:

```
>jones>book>examples>
```

Users who are familiar with Multics or UNIX immediately see such pathnames as invalid, even though they are often used on the Symbolics computer to access Multics and UNIX. When parsed for LMFS or Multics, the above filename string produces a pathname whose directory component designates the directory "examples", which is contained in "book", which itself is contained in "jones", an inferior of the ROOT. The name, type, and version components of this pathname are **nil**. This pathname is equivalent to the following:

```
>jones>book>examples>$.$.$.
```

Either of these is a canonical pathname as directory for the directory "examples". Typing such pathnames as input is exceedingly common, since the merging process, given such a pathname as its unmerged input, replaces the directory component of the default with a directory component specifying the directory named by the "pathname as directory". See the section "Pathname Defaults and Merging", page 147. For example:

```
Default:      Q:>abel>baker>cakes.list
User Typein:  >Romanoli>weddings>
Merged output: >Romanoli>weddings>cakes.list
```

Compare this with the following:

```

Default:      Q:>abel>baker>cakes.list
User Typein:  >Romanoli>weddings
Merged output: >Romanoli>weddings.list

Default:      Q:>abel>baker>cakes.list
User Typein:  >Romanoli>weddings>?.?.73
Merged output: >Romanoli>weddings>cakes.list.73

```

All the Symbolics hierarchical directory parsers recognize a trailing directory delimiter as an instruction to construct a pathname with **nil** name, type, and version, for the directory designated—a "pathname as directory". (The version component, however, remains **:unspecific** for systems not supporting file versions.) This is true even of the parsers for UNIX and Multics, on which systems such syntax is never seen.

This mode of directory naming is usually familiar to users of nonhierarchical systems. The following TENEX pathname results, when parsed, in a pathname as directory for the LMDOC directory (on the appropriate TENEX host), with name, type, and version of **nil**, that can be used in merging operations in a way similar to that shown in the above LMFS example.

```
<LMDOC>
```

As a side-effect of these conventions, the following kinds of pathnames occasionally occur on LMFS or Multics:

```
<lmdoc>
```

As explained above, this is a valid way of entering the following relative pathname:

```
<lmdoc>?.?.?
```

12.1.7 Case in Pathnames

The pathname system handles alphabetic case in pathnames and transferring of pathname components between hosts with different preferred alphabetic cases.

The components of a pathname (directory, name, type, and so on) have two possible representations for case, *raw* (also called *native*) and *interchange*. The raw case representation keeps the case in whatever form is normal for that system (for example, lowercase for UNIX, uppercase for TOPS-20). Interchange representation is a format for manipulating pathname components in a host-independent manner. All pathname defaulting and cross-host translation functions use the interchange form of pathname messages.

All the standard messages to pathnames (for example, **:directory**, **:name**) return pathname components in interchange case rather than raw case.

The components are stored internally in raw case, that is, the actual alphabetic case in which the names of the files are stored, or to be stored, in the host's file system. It is possible to access the raw case representation via the set of messages **:raw-directory**, **:raw-name**, and so forth. However, programs seeking to be system-independent should not use these messages, but the standard ones,

:directory, **:name**, and so forth. Doing so ensures that pathname components transferred between system types stay in the preferred case for each of the systems concerned.

The raw forms of the messages are provided for writing host-specific code or for manipulating several pathname objects known to be on the same host.

<i>Interchange case form</i>	<i>Raw case form</i>
:device	:raw-device
:directory	:raw-directory
:name	:raw-name
:type	:raw-type

The interchange form of the message specifies the following effect:

<i>Case of component</i>	<i>Translated case returned</i>
System default	Uppercase
Mixed case	Mixed case
Opposite to default	Lowercase

Uppercase was chosen as the interchange case because strings like "LISP", representing pathname components, appear in many programs. Either choice (upper or lower) would have been natural for some hosts and not for others.

This facility provides more features for dealing with pathname components independent of the case-sensitivity of file names of different hosts. The following table shows some examples for different host types.

<i>Host</i>	<i>Message</i>	<i>Applied to raw form</i>	<i>Returns interchange form</i>
UNIX	:name	"foo.bar"	"FOO"
	:name	"FOO.BAR"	"foo"
	:name	"Foo.Bar"	"Foo"
Lisp Machine	:name	"foo.bar"	"FOO"
File System	:name	"FOO.BAR"	"FOO"
	:name	"Foo.Bar"	"FOO"
TOPS-20	:name	"FOO.BAR"	"FOO"
	:name	"foo.bar"	"foo"
	:name	"Foo.Bar"	"Foo"

Note that the Lisp Machine File System (LMFS) appears not to follow the interchange case rules. This is because, for LMFS, case is usually maintained but is not significant ("foo", "Foo", and "FOO" are all the same). Thus any mixture of cases in a file name satisfies the "system default" condition and returns all uppercase for the interchange form.

Functions that manipulate pathnames, such as **fs:make-pathname**,

fs:merge-pathnames, and **fs:merge-pathname-and-set-defaults**, manipulate components in interchange case.

Pathname-constructing functions such as **fs:make-pathname** and `pathname` messages such as **:new-pathname** and **:new-default-pathname** accept both **:directory** and **:raw-directory**, to allow specification of components in either interchange case or raw case.

12.2 Defaults and Merging

It is unreasonable to require the user to type a complete pathname, containing all components. Instead the program is expected to supply a *default pathname*, from which values of components not specified by the user can be taken.

Every program that prompts the user for a pathname should maintain some default pathname, display it to the user when prompting for a pathname, and merge the parsed input from the user with that default. The function **prompt-and-read** provides easy ways to do all of these things. See the function **prompt-and-read** in *Programming the User Interface*. No program should use any pathname obtained from user input without merging it against *some* default. Since it is impossible for a user to type a pathname correctly without knowing against which default it will be merged, the default must be displayed to the user.

A *default default* is available for programs that have no better idea of a default pathname, and a function (**fs:default-pathname**) for customizing default pathnames.

Typically, a program might take the default default, customize it, perhaps by supplying a specific file type (usually via the canonical type mechanism), prompt the user for the name of a file, displaying that default, and merge the user's parsed input against that default.

A more complex program, one that requires an input file and an output file, might proceed as follows: It obtains the pathname of its input file as above, and prepares a default pathname for its output file by customizing the input file pathname, usually by supplying a new type, and presents and uses that as a default for the prompt for the output file pathname.

The *merging* operation is performed by the function **fs:merge-pathnames**. It takes as input an unmerged pathname and a default pathname and returns a *merged pathname*, which has no missing components. Basically, the missing components in the unmerged pathname are filled in from the default pathname. The merging operation also takes a *default version* argument, which specifies the version number of the output pathname, *if there is no version mentioned in the unmerged pathname*. That is, the version number is almost never defaulted from the default pathname. If the default version argument is not supplied, it is assumed to be **:newest**. The version number of the default is used as a default version in the following cases:

- Neither name, type, nor version is specified by the unmerged pathname.
- The unmerged pathname does not have a version, and the value of the default version argument is **:default**.

The full details of the merging rules are as follows.

1. If the unmerged pathname does not supply a device, the device is the default file device for that host.
2. If the unmerged pathname does not specify a host, device, directory, name, or type, that component comes from the defaults.
3. If the unmerged pathname supplies a version, it is used.
4. If it does not supply a version, the default version as explained above is used.

Thus, if the user supplies just a name, the host, device, directory and type will come from the default, with the default version argument (or **:newest** if there was none). If the user supplies nothing, or just a directory, the name, type, and version comes over from the default together. If the host's file name syntax provides a way to input a type or version without a name, the user can let the name default but supply a different type or version than the ones in the default.

The system also defines an object called a *defaults alist*. Functions are provided to create one, get the default pathname out of one, merge a pathname with one, and store a pathname back into one. A defaults alist is basically an object containing a replaceable pathname. **fs:merge-pathnames** accepts a defaults alist as its default pathname argument as well as a pathname.

fs:merge-pathnames-and-set-defaults is like **fs:merge-pathnames** but *requires* a defaults alist as its default pathname argument. When it has completed its merge, it stores the result back into the defaults alist before returning it. See the function **fs:merge-pathnames-and-set-defaults**, page 158. It is important that you do not attempt to construct a defaults alist, but instead use the primitives provided. See the function **fs:make-pathname-defaults**, page 162. See the function **fs:copy-pathname-defaults**, page 162. See the function **fs:set-default-pathname**, page 163.

The following special variables are parts of the pathname interface that are relevant to defaults.

fs:*default-pathname-defaults*

Variable

The default defaults alist; if the pathname primitives that need a set of defaults are not given one, they use this one. Most programs, however, should have their own defaults rather than using these.

fs:load-pathname-defaults*Variable*

The defaults alist for the **load** and **compiler:compile-file** functions. Other functions can share these defaults.

12.3 Generic Pathnames

A generic pathname stands for a whole family of files. The property list of a generic pathname is used to remember information about the family, some of which (such as the package) comes from the file attribute list line of a source file in the family. See the section "File Attribute Lists", page 109. All types of files with that name, in that directory, belong together. They are different members of the same family; for example, they might be source code, compiled code, and documentation for a program. All versions of files with that name, in that directory, belong together.

The generic pathname of pathname *p* has the same host, device, directory, and name as *p* does. However, it has a version of **nil**. Furthermore, if the canonical type of *p* is one of the elements of **fs:*known-types***, then it has a type of **nil**; otherwise it has the same type as *p*. The reason that the type of the generic pathname works this way is that in some file systems, such as that of ITS, the type component can actually be part of the file name; ITS files named "DIRECT IONS" and "DIRECT ORY" do not belong together.

The **:generic-pathname** message to a pathname returns its corresponding generic pathname. See the message **:generic-pathname**.

fs:*known-types**Variable*

This is a list of the canonical file types that are "not important"; constructing a generic pathname will strip off the file type if it is in this list. File types not in this list are really part of the name in some sense. The following is the initial list:

```
(:LISP :QBIN :BIN NIL :UNSPECIFIC)
```

Some users might need to add to this list. See the section "Canonical Types in Pathnames", page 151.

12.4 Relative Pathnames

Many operating systems support a notion called *relative pathnames* in order to simplify the typing of filenames by their users. Typically, a user on a system such as Multics or UNIX tells the system what directory on the system is his or her *working directory*. These operating systems assume the working directory as the default directory for filenames whose directory is not specified. For example, when the user types a filename, perhaps as an argument to a command (such as "print

foo") the system assumes that the name foo refers to a file named foo in the working directory, as long as the user did not specify another directory (for instance, by saying "print >sources>c>foo").

On hierarchical systems, such as UNIX and Multics, the working directory can often be several levels deep, and have a full name that is therefore cumbersome to type. The concept of working directory is all the more powerful in these cases. Since the hierarchical organization of directories exists to facilitate relating files by placing them in directories in common subtrees, it is common for users working on such systems to want to reference files in "siblings" of their working directory, or "uncles", or even "inferiors" or "inferiors of inferiors", that is, directories near in the directory hierarchy to their working directory.

In order to facilitate the referencing of files in directories "near" the working directory, without having to type full pathnames of directories, these systems support *relative pathnames*, which are interpreted relative to the working directory. Relative pathnames are always syntactically distinguishable from other pathnames. For instance, on Multics, if the working directory is >udd>Proj>Username, the pathname

```
<Othername>stuff>x.p11
```

refers to the file

```
>udd>Proj>Othername>stuff>x.p11
```

Although it supports relative pathnames, the Lisp Machine File System does not support a concept of working directory. One rationale for this is the fact that the user might be communicating with many systems at once, and might have several working directories to remember. The merging and defaulting system takes the place of the working directory concept. See the section "Pathname Defaults and Merging", page 147. The default pathname, which is displayed when a user is asked to enter a pathname, determines the default directory for a pathname having no directory explicitly specified. What is more, it specifies the default values of other components as well.

Systems supporting relative pathnames usually have some special syntax to indicate a pathname that is relative to a superior of the working directory, and another to indicate pathnames relative to superiors of the working directory. We call these "upward relativization" and "downward relativization". In this context, a pathname with an explicit directory specified is called an *absolute pathname*, and one without an explicit directory, a relative pathname. However, since specification of no directory at all is a very common case handled by systems that do not otherwise support relative directories, namely, by simply defaulting an entire directory component, this is not considered a relative pathname by the Symbolics system.

The Symbolics system supports relative directories for those hierarchical systems that support it themselves. As might be expected, the "resolution" of relative pathnames entered by the user is performed relative to the default pathname, as opposed to any working directory. Resolution of relative pathnames is performed by **fs:merge-pathnames** as part of its normal operation.

The following examples, using LMFS pathnames, show some examples of relative pathnames and their resolution via merging:

```

Default:    >sys>lmfs>new>xst.lisp
Unmerged:  test>xst.lisp
Merged:    >sys>lmfs>new>test>xst.lisp

Default:    >sys>lmfs>new>xst.lisp
Unmerged:  <test>thing.lisp           ;upward relativization
Merged:    >sys>lmfs>test>thing.lisp

Default:    >sys>lmfs>new>xst.lisp
Unmerged:  <<test>                   ;upward relativization
Merged:    >sys>test>xst.lisp

Default:    >sys>lmfs>new>xst.lisp
Unmerged:  test>best>                ;downward relativization
Merged:    >sys>new>test>best>xst.lisp

Default:    >sys>lmfs>new>xst.lisp
Unmerged:  <xst.lisp
Merged:    >sys>lmfs>xst.lisp

Default:    >sys>lmfs>new>xst.lisp
Unmerged:  <<abel>baker>foo.lisp
Merged:    >sys>abel>baker>foo.lisp

```

12.5 Canonical Types in Pathnames

A *canonical type* for a pathname is a symbol that indicates the nature of a file's contents. To compare the types of two files, particularly when they could be on different kinds of hosts, you compare their canonical types.

(**fs:*default-canonical-types*** and **fs:*canonical-types-alist*** show the canonical types and the default surface types for various hosts.)

Some terminology:

- canonical type* A host-independent name for a certain type of file, for example, Lisp compiled code files or LGP font files. A canonical type is a keyword symbol.
- file specification* What you type when you are prompted to supply a string for the system to build a pathname object.
- surface type* The appearance of the type component in a file specification. This is a string in native case.

default surface type

Each canonical type has as part of its definition a representation for the type when it has to be used in a string. Default surface type is the string (in interchange case) that would be used in a string being generated by the system and shown to the user. See the function **fs:define-canonical-type**, page 163.

preferred surface type

Some canonical types have several different possible surface representations. The definition for the type designates one of these as the preferred surface type. It is a string in interchange case. ("Default surface type" implies "preferred surface type" when one has been defined.)

Each canonical type has a default surface representation, which can be different from the surface file type actually appearing in a file specification. **:lisp** is a canonical type for files containing Lisp source code. For example, on UNIX, the default surface representation of the type for **:lisp** files is "L". (Remember, the default surface representation is kept in interchange case.) The surface type in a file specification containing lisp code is different on different systems, "LISP" for Lisp Machine file system, "l" for UNIX. You can find out from a pathname object both the canonical type for the pathname and the surface form of the type for the pathname by using the **:canonical-type** message. See the method **(:method fs:pathname :canonical-type)**, page 166.

The following tables illustrate the terminology.

		<i>UNIX</i>	
Surface type	"l"	"lisp"	"foo"
Raw type	"l"	"lisp"	"foo"
Type	"L"	"LISP"	"FOO"
Canonical type	:lisp	:lisp	"FOO"
Original type	nil	"LISP"	"FOO"

		<i>Lisp machine</i>	
Surface type	"l"	"lisp"	"foo"
Raw type	"l"	"lisp"	"foo"
Type	"L"	"LISP"	"FOO"
Canonical type	"L"	:lisp	"FOO"
Original type	"L"	nil	"FOO"

To translate the type field of a pathname from one host to another, determine the canonical type, using the surface type on the original host. Then find a surface type on the new host for that canonical type.

Copying operations can preserve the surface type of the file through translations and defaulting rather than by converting it to the surface form for the canonical type. For example:

```
(multiple-value-bind (ctype otype)
  (send p ':canonical-type)
  (send p ':new-pathname
    ':canonical-type ctype
    ':original-type otype
    ':name "temp-p"))
```

12.5.1 Correspondence of Canonical Types and Editor Modes

fs:*file-type-mode-alist* is an alist that associates canonical types (in the car) with editor major modes (in the cdr).

```
((:LISP . :LISP) (:SYSTEM . :LISP) (:TEXT . :TEXT) ...)
```

12.6 Wildcard Pathname Mapping

In the Symbolics system, as in some other systems, wildcard pathnames are used not only to specify groups of files, but to specify mappings between pairs of pathnames, for operations such as renaming and copying files.

For example, you might ask to copy **foo*.lisp* to **bar*.lisp*. All the files to be copied match the wildcard name **foo*.lisp*. **bar*.lisp* is a specification of how to construct the pathname of the new file. The two wildcard pathnames, as in the above example, are called the *source pattern* and *target pattern*. The original name of any file to be copied is called the *starting instance*. Here is an example:

```
Source pattern:      f:>fie>*old*.lisp
Target pattern:     vx:/usr2/fum/*older*.l
Starting instance:  f:>fie>--oldfoo.lisp
Target instance:    vx:/usr2/fum/--olderfoo.l
```

A more abstract description of this terminology:

Source pattern	A pathname containing wild components.
Target pattern	A pathname containing wild components.
Source instance	A pathname that matches the source pattern.
Target instance	A pathname specified by applying the common sequences between the source and target patterns to the source instance.

Two Zmacs commands accept pairs of wildcard file specifications:

```
Copy File (m-X)
Rename File (m-X)
```

The components of the target instance are determined component-by-component for all components except the host. (The host component is always determined literally

from the source and target patterns; it cannot be wild.) The mapping of pathnames is done in the native case of the target host. The source pattern and source instance are coerced to the target host via the **:new-default-pathname** message before the mapping takes place. See the method **(:method fs:pathname :new-default-pathname)**, page 170.

When the type of the target pattern is **:wild**, it uses the canonical type for the target, regardless of the surface form for the type in the source pattern and instance.

NOTE

In the Lisp Machine File System, * as the directory portion of a file specification specifies a relative pathname. You must use >** to indicate a wild directory component that matches any directory at all. See the section "LMFS Pathnames", page 174.

Here are the rules used in constructing a target instance, given the source and target patterns and a particular source instance. This set of rules is applied separately to each component in the pathname. In the mapping rules, a * character as the only contents of a component of a file spec is considered to be the same as the keyword symbol **:wild**. The rule uses the patterns from the example above.

1. If the target pattern does not contain *, copy the target pattern component literally to the target instance.
2. If the target pattern is **:wild**, copy the source component to the target literally with no further analysis. The type component is handled somewhat differently — when source and target hosts are of different system types, it uses the canonical-type mechanism to translate the type. This does not apply when the target pattern is **:wild-inferiors**, in directory specifications.
3. Find the positions of all * characters in the source and target patterns. Take the characters intervening between * characters as a literal value. Literal values for the name component:


```
Source: old
Target: older
```
4. Find each literal value from the source pattern in the source instance. Take the characters intervening between literal values as a matching value for the * from the source pattern. The matching value could be any number of characters, including zero. Matching values for the name component:


```
-- and foo
```
5. Create the component by assembling the literal and matching values in left to right order, substituting the matching values where * appears in the target pattern. For the name component:


```
--olderfoo
```

When not enough matching values are available (due to too few * in the source pattern) use the null string as the matching value. When the source pattern has too many *, ignore the first extra * and everything following it.

Some examples:

<i>Source pattern</i>	<i>Source instance</i>	<i>Target pattern</i>	<i>Target instance</i>
*report	6802-report	*summary	6802-summary
lmfs-*	lmfs-errors	*	lmfs-errors
l*	l	l*	l
l*	lisp	l*	lisp
OLD-DIR	OLD-DIR	NEW-PLACE	NEW-PLACE
*	doc	*-extract	doc-extract
doc	doc	doc-extract	doc-extract

12.6.1 Wildcard Directory Mapping

The rules for mapping directory components between two wildcard pathnames and a starting instance are parallel to the rules for single names. Directory-level components play roughly the roles of characters in the name-translating algorithm. See the section "Wildcard Pathname Mapping", page 153.

Consider a directory component as a sequence of directory level components. The levels are separated by level delimiters (> in LMFS). Example: In the pathname >foo>bar>*>mumble*>x>**>y>a.b.3, the directory-level components are foo, bar, *, mumble*, x, **, and y. The source and target patterns, as well as the starting instance, are considered as sequences of directory-level components, and are matched and translated level by level.

For this purpose, each directory-level component can be classified as one of three types:

<i>Type</i>	<i>Directory representation</i>
<i>constant</i>	String containing no *'s
<i>wild-inferiors</i>	** in LMFS, ... in VMS
<i>must-match</i>	* or string containing at least one * (but not the string representing wild-inferiors)

The matching and mapping of constant and wild-inferiors levels proceeds in a manner identical to the matching and mapping of constant substrings and *s for single names. See the section "Wildcard Pathname Mapping", page 153. Constant directory level components act as constant substrings in that algorithm, and wild-inferiors levels as *s. That is, wild-inferiors level components match and, on the target side, carry, zero to any number of constant directory-level components.

Examples:

```
Source pattern: >sys>**>*.*.newest
Target pattern: >old-systems>release-5>**>*.*.
Starting instance: >sys>lmfs>patch>lmfs-33.patch-dir.66
Target instance: >old-systems>release-5>lmfs>patch>lmfs-33.patch-dir.66
```

```
Source pattern: >a>b>c>**>d>e>**>x.y.*
Target pattern: >t>u>**>m>**>w>*.*.
Starting instance: >a>b>c>p>q>d>e>f>g>x.y.1
Target instance: >t>u>p>q>m>f>g>w>x.y.1
```

Must-match components are matched with exactly one directory-level component, which must be present. They are mapped according to the string-mapping rules in the name-translating algorithm. See the section "Wildcard Pathname Mapping", page 153.

Example:

```
Source pattern: >a>b>c>foo*>d>*>*.*.
Target pattern: >x>*bar>y>*man>*.*.
Starting instance: >a>b>c>foolish>d>yow>a.lisp.1
Target instance: >x>lishbar>y>yowman>a.lisp.1
```

You can intersperse constants, must-matches, and wild-inferiors directory-level components, as long as the sequence of wildcard types is the same in both patterns.

Example:

```
Source pattern: >a>*>c>**>*.lisp.*
Target pattern: >bsg>sub>new->q>**>*.*.
Starting instance: >a>bb>c>d>e>p1.lisp.6
Target instance: >bsg>sub>new-bb>q>d>e>p1.lisp.6
```

12.7 Pathname Functions

The following functions are what programs use to parse and default file names that have been typed in or otherwise supplied by the user.

fs:parse-pathname *thing* &optional *with-respect-to* (*defaults* *Function*
fs:*default-pathname-defaults*)

Turns *thing*, which can be a pathname, a string, or a Maclisp-style name list, into a pathname. Most functions that take a pathname argument call **fs:parse-pathname** on it so that they accept anything that can be turned into a pathname. Some, however, do it indirectly, by calling **fs:merge-pathnames**.

This function does *not* do defaulting, even though it has an argument named *defaults*; it only does parsing. The *with-respect-to* and *defaults* arguments are there because in order to parse a string into a pathname, it is necessary to know what host it is for so that it can be parsed with the file name syntax peculiar to that host.

If *with-respect-to* is supplied, it should be a host or a string to be parsed as the name of a host. If *thing* is a string, it is then parsed as a true string for that host; host names specified as part of *thing* are not removed. Thus, when *with-respect-to* is not *nil*, *thing* should not contain a host name.

If *with-respect-to* is not supplied or is *nil*, any host name inside *thing* is parsed and used as the host. If *with-respect-to* is *nil* and no host is specified as part of *thing*, the host is taken from *defaults*.

Examples, using a LMFS host named Q:

```
(fs:parse-pathname "a:>b.c" "q") => #<LMFS-PATHNAME "Q:a:>b.c"> ;(wrong)
(fs:parse-pathname "q:>b.c" "q") => #<LMFS-PATHNAME "Q:q:>b.c"> ;(wrong)
(fs:parse-pathname "q:>b.c")    => #<LMFS-PATHNAME "Q:>b.c">
(fs:parse-pathname ">b.c" "q")  => #<LMFS-PATHNAME "Q:>b.c">
```

Note that this causes correct parsing of a TOPS-20 pathname when *thing* contains a device but no host and when *with-respect-to* is not *nil*. (Warning: If *thing* contains a device but no host and if *with-respect-to* is *nil* or not supplied, the device is interpreted as a host.) In the following example, X is a TOPS-20 host and A is a device:

```
(fs:parse-pathname "a:<b>c.d" "x") => #<TOPS20-PATHNAME "X:A:<B>C.D">
(fs:parse-pathname "a:<b>c.d")    => Error: "a" is not a known file
                                server host.
```

In the same TOPS-20 example, if *with-respect-to* is *nil* and the host is taken from *defaults*, the pathname string must be preceded by a colon to be parsed correctly:

```
(fs:parse-pathname ":a:<b>c.d" nil "x:") => #<TOPS20-PATHNAME "X:A:<B>C.D">
(fs:parse-pathname "a:<b>c.d" nil "x:") => Error: "a" is not a known file
                                server host.
```

If *thing* is a list, *with-respect-to* is specified, and *thing* contains a host name, an error is signalled if the hosts from *with-respect-to* and *thing* are not the same.

fs:merge-pathnames *pathname* &optional (*defaults* *Function*
fs:*default-pathname-defaults*)
(*default-version* **'newest**)

Fills in unspecified components of *pathname* from the defaults, and returns a new pathname. This is the function that most programs should call to process a file name supplied by the user. *pathname* can be a pathname, a string, a symbol, or a Maclisp name list. The returned value is always a

pathname. The merging rules are documented elsewhere: See the section "Pathname Defaults and Merging", page 147.

If *pathname* is a string or a symbol, it is parsed before merging. The default pathname is presented to **fs:parse-pathname** as a default pathname, from which the latter defaults the host if there is no explicit host named in the string.

defaults can be a pathname, a defaults alist, or a string. If it is a string, it is parsed against the default defaults. *defaults* defaults to the value of **fs:*default-pathname-defaults*** if unsupplied.

fs:merge-pathnames-and-set-defaults *pathname* &optional *Function*
 (*defaults* **fs:*default-pathname-defaults***)
 (*default-version* **'newest**)

The same as **fs:merge-pathnames** except that after it is done the result is stored back into *defaults*. This is handy for programs that have "sticky" defaults. (If *defaults* is a pathname rather than a defaults alist, then no storing back is done.) The optional arguments default the same way as in **fs:merge-pathnames**.

The following function is what programs use to complete a partially typed-in pathname.

fs:complete-pathname *defaults string type version* &rest *options* *Function*

string is a partially specified file name. (Presumably it was typed in by a user and terminated with the COMPLETE or END to request completion.)

fs:complete-pathname looks in the file system on the appropriate host and returns a new, possibly more specific string. Any unambiguous abbreviations are expanded in a host-dependent fashion.

string is completed relative to a default pathname constructed from *defaults*, the host (if any) specified by *string*, *type*, and *version*, using the function **fs:default-pathname**. See the function **fs:default-pathname**, page 162. If *string* does not contain a colon, the host comes from *defaults*; otherwise the host name precedes the first colon in *string*.

options are keywords (without following values) that control how the completion will be performed. The following option keywords are allowed. Their meanings are explained more fully below.

:deleted Look for files that have been deleted but not yet expunged. The default is to ignore such files.

:read or **:in** The file is going to be read. This is the default. The name **:in** is obsolete and should not be used in new programs.

:write or :print or :out

The file is going to be written (that is, a new version is going to be created). The names **:print** and **:out** are obsolete and should not be used in new programs.

:old

Look only for files that already exist. This is the default. **:old** is not meaningful when **:write** is specified.

:new-ok

Allow either a file that already exists, or a file that does not yet exist. **:new-ok** is not meaningful when **:write** is specified. The **:new-ok** option is no longer used by any system software, because users found its effects (in the Zmacs command Find File (c-X c-F)) to be too confusing. It remains available, but programmers should consider this experience when deciding whether to use it.

The first value returned is always a string containing a file name; either the original string, or a new, more specific string. The second value returned indicates the status of the completion. It is non-**nil** if it was completely successful. The following values are possible:

:old

The string completed to the name of a file that exists.

:new

The string completed to the name of a file that could be created.

nil

The operation failed for one of the following reasons:

- The file is on a file system that does not support completion. The original string is returned unchanged.
- There is no possible completion. The original string is returned unchanged.
- There is more than one possible completion. The string is completed up to the first point of ambiguity.
- A directory name was completed. Completion was not successful because additional components to the right of this directory remain to be specified. The string is completed through the directory name and the delimiter that follows it.

Although completion is a host-dependent operation, the following guidelines are generally followed:

When a pathname component is left completely unspecified by *string*, it is generally taken from the default pathname. However, the name and type are defaulted in a special way described below and the version is not defaulted at all; it remains unspecified.

When a pathname component is specified by *string*, it can be recognized as an abbreviation and completed by replacing it with the expansion of the abbreviation. This usually occurs only in the rightmost specified component of *string*. All files that exist in a certain portion of the file system and match this component are considered. The portion of the file system is determined by the specified, defaulted, or completed components to the left of this component. A file's component *x* matches a specified component *y* if *x* consists of the characters in *y* followed by zero or more additional characters; in other words, *y* is a left substring of *x*. If no matching files are found, completion fails. If all matching files have the same component *x*, it is the completion. If there is more than one possible completion, that is, more than one distinct value of *x*, there is an ambiguity and completion fails unless one of the possible values of *x* is equal to *y*.

If completion of a component succeeds, the system attempts to complete any additional components to the right. If completion of a component fails, additional components to the right are not completed.

A blank component is generally treated the same as a missing component; for example, if the host is a LMFS, completion of the strings "foo" and "foo." deals with the type component in the same way. The strings are not completed identically; completion of "foo" attempts to complete the name component, but completion of "foo." leaves the name component alone since it is not the rightmost.

If *string* does not specify a name, then the name of the default pathname is *preferred* but is not necessarily used. The exact meaning of this depends on *options*:

- With the default options, if any files with the default name exist in the specified, defaulted, or completed directory, the default name is used. If no such files exist, but all files in the directory have the same name, that name is used instead. Otherwise, completion fails.
- With the **:write** option, the default name is always used when *string* does not specify a name, regardless of what files exist.
- With the **:new-ok** option, if any files with the default name exist in the specified, defaulted, or completed directory, the default name is used. If no such files exist, but all files in the directory have the same name, that name is used instead. Otherwise, the default name is used.

The special treatment of the case where all files in the directory have the same name is not very useful and is not implemented by all file systems.

If *string* does not specify a type, then the type of the default pathname is *preferred* but is not necessarily used. The exact meaning of this depends on *options*:

- With the default options, if a file with the specified, defaulted, or completed name and the default type exists, the default type is used. If no such file exists, but one or more files with that name and some other type do exist and all such files have the same type, that type is used instead. Otherwise, completion fails.
- With the **:write** option, the default type is always used when *string* does not specify a type, regardless of what files exist.
- With the **:new-ok** option, if a file with the specified, defaulted, or completed name and the default type exists, the default type is used. If no such file exists, but one or more files with that name and some other type do exist and all such files have the same type, that type is used instead. Otherwise, the default type is used.

In file systems such as LMFS and UNIX that require a trailing delimiter (> or /) to distinguish a directory component from a name component, the system heuristically decides whether the rightmost component was meant to be a directory or a name, and inserts the directory delimiter if necessary.

If *string* contains a relative directory specification for a host with a hierarchical file system, it is assumed to be relative to the directory in the default pathname and is expanded into an absolute directory specification.

The host and device components generally are not completed; they must be fully specified if they are specified at all. This might change in the future.

If *string* does not specify a version, the returned string does not specify a version either. This differs from file name completion in TOPS-20; TOPS-20 completes an implied version of "newest" to a specific number. This is possible in TOPS-20 because completing a file name also attaches a "handle" to a file. In Symbolics-Lisp, the version number of the newest file might change between the time the file name is completed and the time the actual file operation (open, rename, or delete) is performed.

A pathname component must satisfy the following rules in order to appear in a successful completion:

- The host, device, and directory must actually exist.
- The name must be the name of an existing file in the specified directory, unless **:write** or **:new-ok** is included in *options*.
- The type must be the type of an existing file with the specified name in the specified directory, unless **:write** or **:new-ok** is included in *options*.
- A pathname component always completes successfully if it is **:wild**.

When the rules are not satisfied by a component taken from the default pathname, completion fails and that component remains unspecified in the resulting string. When the rules are not satisfied by a component taken from *string*, completion fails and that part of *string* remains unchanged (other components of *string* can still be expanded).

This function yields a pathname, given its components.

fs:make-pathname &rest *options* *Function*
options are alternating keywords and values that specify the components of the pathname. Missing components default to **nil**, except the host (all pathnames must have a host). The **:defaults** option specifies the defaults to get the host from if none are specified. The other options allowed are **:host**, **:device**, **:directory**, **:name**, **:type**, **:version**, **:raw-device**, **:raw-directory**, **:raw-name**, **:raw-type**, **:canonical-type**.

The following functions are used to manipulate defaults alists directly.

fs:make-pathname-defaults *Function*
 Creates a defaults alist initially containing no defaults. Asking this empty set of defaults for its default pathname before anything has been stored into it returns the file FOO on the user's home directory on the host to which the user logged in.

Defaults alists created with **fs:make-pathname-defaults** are remembered, and reset whenever the site is changed. This prevents remembered defaults from pointing to unknown hosts when world load files are moved between sites.

fs:copy-pathname-defaults *defaults* *Function*
 Creates a defaults alist, initially a copy of *defaults*.

fs:default-pathname &optional *defaults* *host* *default-type* *Function*
default-version *sample-p*

Obtains a pathname suitable for use as a default pathname and customizes it by modification of its *type* and *version*. It also extracts pathnames out of default alists.

The pathname returned by **fs:default-pathname** is always fully specified; that is, all components have non-**nil** values. This is needed when defaulting a pathname with **fs:merge-pathnames** to pass to **open** or other file-system operations, as these operations should always receive fully specified pathnames.

Specifying the optional arguments *host*, *default-type*, and *default-version* as not **nil** forces those fields of the returned pathname to contain those values. If *defaults*, which can be a pathname or a defaults alist, is not specified, the default defaults are used.

If *default-type* is a symbol representing a canonical type, that canonical type is used as the canonical type of the pathname returned. That is, the pathname has a type component that is the correct representation of that canonical type for the host.

Users should never supply the optional argument *sample-p*.

fs:set-default-pathname *pathname* &optional *defaults* *Function*
 Updates a defaults alist. It stores *pathname* into *defaults*. If *defaults* is not specified, the default defaults are used.

The following functions return useful information.

fs:user-homedir &optional (*host* **fs:user-login-machine**) *Function*
 Returns the pathname of the logged-in user's home directory on *host*, which defaults to the host the user logged in to. For a registered user (one who logged in without using the **:host** argument to **login**), the host is the user's **home-host** attribute. Home directory is a somewhat system-dependent concept, but from the point of view of the Symbolics computer it is usually the directory where the user keeps personal files such as init files and mail. This function returns a pathname without any name, type, or version component (those components are all **nil**).

fs:init-file-pathname *program-name* &optional (*canonical-type* **nil**) *Function*
 (*host* **fs:user-login-machine**)
 Returns the pathname of the logged-in user's init file for the program *program-name*, on the *host*, which defaults to the host the user logged in to. Programs that load init files containing user customizations call this function to find where to look for the file, so that they need not know the separate init file name conventions of each host operating system. The *program-name* "LISPM" is used by the **login** function. *canonical-type* is the canonical type of the init file. It should be **nil** when the returned pathname is being passed to **load** so that **load** can look for a file of the appropriate type.

The following function defines a canonical file type.

fs:define-canonical-type *canonical-type* *default* &body *specs* *Special Form*
 Defines a new canonical type. *canonical-type* is the symbol for the new type; *default* is a string containing the default surface type for any kind of host not mentioned explicitly. The body contains a list of specs that define the surface types that indicate the new canonical type for each host. The following example would define the canonical type **:lisp**.

```
(fs:define-canonical-type :lisp "LISP"
  ((:tops-20 :tenex) "LISP" "LSP")
  (:unix "L" "LISP")
  (:vms "LSP"))
```

For systems with more than one possible default surface form, the form that appears first becomes the preferred form for the type. Always use the interchange case.

Define new canonical types carefully so that they are valid for all host types. For example "com-map" would not be valid on VMS because it is both too long and contains an invalid character. You must define them so that the surface types are unique. That is, the same surface type cannot be defined to mean two different canonical types.

Canonical types that specify binary files must specify the byte size for files of the type. This helps **copyf** and other system tools determine the correct byte size and character mode for files. You specify the byte size by attaching a **:binary-file-byte-size** property to the canonical type symbol. For example, the system defines the byte size of press files as follows.

```
(defprop :press 8. :binary-file-byte-size)
```

The following function is useful when dealing with canonical types. Unlike other functions described here, this function actually accesses and searches a host file system. This description is provided here for completeness. For functions and messages that actually access host file systems: See the section "Streams", page 1.

fs:find-file-with-type *pathname canonical-type* *Function*

Searches the file system to determine the actual surface form for a *pathname* object. Like **probef**, it returns the truename for *pathname*. When no file can be found to correspond to a *pathname*, it returns **nil**.

If *pathname* is a string, it is parsed against the default defaults to obtain an actual *pathname* object before processing.

canonical-type applies only when *pathname* has **nil** as its type component. **fs:find-file-with-type** searches the file system for any matching file with *canonical-type*. For example, on a TOPS-20 host, this would look first for ps:<gcw>toolkit.lisp and then for ps:<gcw>toolkit.lsp:

```
(fs:find-file-with-type (fs:parse-pathname "sc:<gcw>toolkit") ':lisp)
```

If it finds more than one file, it returns the one with the preferred surface type for *canonical-type* (or chooses arbitrarily if none of the files has the preferred surface type).

If *pathname* already had a type supplied explicitly, that overrides *canonical-type*. You can ensure that *canonical-type* applies by first setting the type explicitly:

```
(fs:find-file-with-type (send p 'new-type nil) ':lisp)
```

System programs that supply a default type for input files (for example, **load**, **make-system**, and **qc-file**) could use this mechanism for finding their input files.

The following functions are useful for poking around.

fs:describe-pathname *pathname* *Function*

If *pathname* is a pathname object, this describes it, showing you its properties (if any) and information about files with that name that have been loaded into the machine. If *pathname* is a string, this describes all interned pathnames that match that string, ignoring components not specified in the string. This is useful for finding the directory of a file whose name you remember. Giving **describe** a pathname object does the same thing as this function.

fs:pathname-plist *pathname* *Function*

Parses and defaults *pathname* then returns the list of properties of that pathname.

12.8 Pathname Messages

This section documents some of the messages a user can send to a pathname object. These messages are known as the *passive messages* to pathnames. They deal with inspecting and extracting components, constructing new pathnames based on old pathnames and new components, matching pathnames, and so forth. None of these messages actually interact with any host file system; they deal only with pathname objects within the Symbolics computer.

The other common, useful class of messages to pathnames are those that open, delete, and rename files, list directories, find and change file properties, and so forth. These are the *active messages* to pathnames. You usually do not send these messages directly, but use interface functions, such as **open**, **probef**, **deletef**, **renamef**, **fs:directory-list**, **fs:file-properties**, and **fs:change-file-properties**. Neither these functions and messages, nor additional similar ones, are documented here. See the section "Streams", page 1.

Pathnames handle some additional messages that are intended to be sent only by the pathname system itself, and therefore are not documented here. Only someone who wanted to add a new type of file host to the system would need to understand those internal messages. This section also does not document messages that are peculiar to pathnames of a particular type of host.

:host of **fs:pathname** *Method*

Returns the host component of the pathname. The returned value is always a host object. If the pathname is a logical pathname, the logical host is returned. It is an error to send **:host** to a logical host.

- :device** of **fs:pathname** *Method*
Returns the device component of the pathname. The returned value can be **nil**, **:unspecific**, or a string. The string is in interchange case.
- :directory** of **fs:pathname** *Method*
Returns the directory component of the pathname. The returned value can be **nil**, **:wild**, or a list of strings and symbols, each representing a directory level. These symbols can be **:wild** or **:wild-inferiors**. Single names of directories in nonhierarchical file systems are returned as single element lists. The strings are in interchange case.
- :name** of **fs:pathname** *Method*
Returns the name component of the pathname. The returned value can be **nil**, **:wild**, or a string. The string is in interchange case.
- :type** of **fs:pathname** *Method*
Returns the type component of the pathname. The returned value is always be **nil**, **:unspecific**, **:wild**, or a string. The string is in interchange case.
- :version** of **fs:pathname** *Method*
Returns the version component of the pathname. The returned value is always be **nil**, **:wild**, **:unspecific**, **:oldest**, **:newest**, or a number.
- :raw-device** of **fs:pathname** *Method*
Returns the device component of the pathname. The returned value can be **nil**, **:unspecific**, or a string. The string is in its raw case.
- :raw-directory** of **fs:pathname** *Method*
Returns the directory component of the pathname. The returned value can be **nil**, **:wild**, or a list of strings and symbols, each representing a directory level. These symbols can be **:wild** or **:wild-inferiors**. Single names of directories in nonhierarchical file systems will be returned as single element lists. The strings are in their raw case.
- :raw-name** of **fs:pathname** *Method*
Returns the name component of the pathname. The returned value can be **nil**, **:wild**, or a string. The string is in its raw case.
- :raw-type** of **fs:pathname** *Method*
Returns the type component of the pathname. The returned value is always **nil**, **:unspecific**, **:wild**, or a string. The string is in its raw case.
- :canonical-type** of **fs:pathname** *Method*
Determines the canonical type of a pathname and a surface representation for the type. It returns two values:

<i>Value</i>	<i>Meaning</i>
canonical type	This is either a keyword symbol from the set of known canonical types or a string (when the type component of the pathname is not a known canonical type). The string contains the type component from the pathname, in interchange case.
original type	This is nil when the type of the pathname is the same as the preferred surface type for the canonical type. See the function fs:define-canonical-type , page 163. Otherwise, when the type differs from the preferred or default surface type, it is the original type in interchange case.

For example, for a UNIX pathname, sending the message **:canonical-type** to the following pathnames has these results:

<i>Pathname</i>	<i>Results from :canonical-type message</i>		
foo.l	:lisp	nil	Preferred surface type
foo.lisp	:lisp	"LISP"	Alternate surface type
foo.L	"l"	"l"	Not recognized
foo.LISP	"lisp"	"lisp"	Not recognized

Keep in mind that the **:canonical-type** message returns the type string in the interchange case rather than in the raw case.

:new-device *new-device* of **fs:pathname** *Method*

Returns a new pathname that is the same as the pathname it is sent to except that the value of the device component has been changed. The valid set of arguments to the **:new-device** message is the set of possible outputs of **:device**. See the method (**:method fs:pathname :device**), page 166. A string value is expected to be in interchange case.

:new-directory *new-directory* of **fs:pathname** *Method*

Returns a new pathname which is the same as the pathname it is sent to except that the value of the directory component has been changed. The valid set of arguments to the **:new-directory** message is the set of possible outputs of **:directory**. See the method (**:method fs:pathname :directory**), page 166. String values are expected to be in interchange case.

:new-name *new-name* of **fs:pathname** *Method*

Returns a new pathname which is the same as the pathname it is sent to except that the value of the name component has been changed. The valid set of arguments to the **:new-name** message is the set of possible outputs of **:name**. See the method (**:method fs:pathname :name**), page 166. String values are expected to be in interchange case.

- :new-type** *new-type* of **fs:pathname** *Method*
Returns a new pathname that is the same as the pathname it is sent to except that the value of the type component has been changed. The valid set of arguments to the **:new-type** message is the set of possible outputs of **:type**. See the method (**:method fs:pathname :type**), page 166. String values are expected to be in interchange case.
- :new-version** *new-version* of **fs:pathname** *Method*
Returns a new pathname that is the same as the pathname it is sent to except that the value of the version component has been changed. The valid set of arguments to the **:new-version** message is the set of possible outputs of **:version**. See the method (**:method fs:pathname :version**), page 166.
- :system-type** of **fs:pathname** *Method*
Returns the type of host that the pathname is intended for. This value is a keyword from the following set:
:its, :lispm, :multics, :tenex, :tops-20, :unix, :vms, :logical
This is the same set as returned by the **:system-type** message to a host object. It is not likely that you need to use this message directly.
- :new-raw-device** *dev* of **fs:pathname** *Method*
Returns a new pathname that is the same as the pathname it is sent to except that the value of the device component has been changed. The valid set of arguments to the **:new-raw-device** message is the set of possible outputs of **:raw-device**. See the method (**:method fs:pathname :raw-device**), page 166. A string value is expected to be in its raw case.
- :new-raw-directory** *new-directory* of **fs:pathname** *Method*
Returns a new pathname that is the same as the pathname it is sent to except that the value of the directory component has been changed. The valid set of arguments to the **:new-raw-directory** message is the set of possible outputs of **:raw-directory**. See the method (**:method fs:pathname :raw-directory**), page 166. String values are expected to be in their raw case.
- :new-raw-name** *new-name* of **fs:pathname** *Method*
Returns a new pathname which is the same as the pathname it is sent to except that the value of the name component has been changed. The valid set of arguments to the **:new-raw-name** message is the set of possible outputs of **:raw-name**. See the method (**:method fs:pathname :raw-name**), page 166. String values are expected to be in their raw case.

:new-raw-type *new-type* of **fs:pathname** *Method*

Returns a new pathname that is the same as the pathname it is sent to except that the value of the type component has been changed. The valid set of arguments to the **:new-raw-type** message is the set of possible outputs of **:raw-type**. See the method (**:method fs:pathname :raw-type**), page 166. String values are expected to be in their raw case.

:new-canonical-type *canonical-type* &optional *original-type* of **fs:pathname** *Method*

Returns a new pathname based on the old one but with a new canonical type. *canonical-type* specifies the canonical type for the new pathname. The surface type of the new pathname is based on the default surface type of the canonical type, unless the pathname already had the correct type.

When the pathname object receiving the message already has the correct canonical type, the surface type in the new pathname depends on the presence of *original-type*. When *original-type* is omitted, the new pathname type has the same surface type as the old pathname. When *original-type* is supplied, the surface type for the new pathname is *original-type*. This assumes that *original-type* is a valid representation for *canonical-type*; if that assumption is not met, the *canonical-type* prevails and its default surface type is used.

canonical-type is a symbol for a known type, **:unspecific**, **nil**, or a string. Use a string for *canonical-type* to make pathnames with types that are not known canonical types.

The following examples assume that a pathname object for the file specification "vixen:/usr2/jwalker/mild.new" is the value of **m**.

```
(send m 'new-canonical-type 'lisp) =>
#<UNIX-PATHNAME "VIXEN: //usr2//jwalker//mild.l">
(send m 'new-canonical-type 'lisp "LISP") =>
#<UNIX-PATHNAME "VIXEN: //usr2//jwalker//mild.lisp">
(send m 'new-canonical-type 'lisp "MSS") =>
#<UNIX-PATHNAME "VIXEN: //usr2//jwalker//mild.l">
(send m 'new-canonical-type "BAR" "BAR") =>
#<UNIX-PATHNAME "VIXEN: //usr2//jwalker//mild.bar">
(send m 'new-canonical-type 'lisp "lisp") =>
#<UNIX-PATHNAME "VIXEN: //usr2//jwalker//mild.l">
(send m 'new-canonical-type 'lisp nil) =>
#<UNIX-PATHNAME "VIXEN: //usr2//jwalker//mild.l">
```

:types-for-canonical-type *canonical-type* of **fs:pathname** *Method*

The internal primitive for finding which surface types correspond to *canonical-type*. Normally you would not use this directly. To determine what form of a pathname exists in a file system: See the function **fs:find-file-with-type**, page 164.

:new-pathname &rest *options* of **fs:pathname** *Method*

Returns a new pathname that is the same as the pathname it is sent to except that the values of some of the components have been changed. *options* is a list of alternating keywords and values. The keywords all specify values of pathname components; they are **:host**, **:device**, **:directory**, **:name**, **:type**, **:version**, **:raw-name**, **:raw-device**, **:raw-type**, **:raw-directory**, and **:canonical-type**. The **:type** argument also accepts a symbol as an argument, implying canonical type. See the section "Canonical Types in Pathnames", page 151.

:new-default-pathname &rest *options* of **fs:pathname** *Method*

Returns a new valid pathname based on the one receiving the message, using the pathname components supplied by *options*. The components do not need to be known to be valid on a particular host. The method uses the components "as suggestions" for building the new pathname; it is free to make substitutions as necessary to create a valid pathname. It is heuristic, not algorithmic, so it does not necessarily yield valid semantics. The heuristics used, however, seem to produce pathnames that match what many people expect from cross-host defaulting.

It always produces a pathname with valid syntax but not necessarily valid semantics. For example, when it tries to map between a hierarchical file system and a nonhierarchical file system, it uses the least significant of the hierarchical components as the directory component. Sometimes this is not correct, but in all cases it is syntactically valid. The main applications for **:new-default-pathname** are in producing defaults to offer to the user and in copying components from one kind of pathname to another.

Application notes: **:new-pathname** always does what its arguments specify; it never uses heuristics. Thus **:new-pathname** could signal an error in certain cross-host situations where **:new-default-pathname** would not have any problems. Usually, user programs should use **fs:default-pathname**, which sends **:new-default-pathname** as part of its operation. However, if you are copying a single component from one kind of pathname to another, **:new-default-pathname** is the right tool.

For example, the right way to copy the version from an input pathname to an output pathname is as follows:

```
(defun copy-version (input-pathname output-pathname)
  (send output-pathname :new-default-pathname
        :version (send input-pathname :version)))
```

If the above example used **:new-pathname** or **:new-version**, the input pathname were a UNIX pathname, and the output were a LMFS pathname, this example would signal an error, since **:unspecific** is not a valid version in a LMFS pathname. However, using **:new-default-pathname**, the closest equivalent is substituted, namely **:newest**.

:parse-truename *string* &optional (*from-filesystem t*) of *fs:pathname* *Method*

Returns the pathname corresponding to the *string* argument. The *string* is parsed, with the pathname supplying the defaults (notably, the host). The method is useful when, for example, a remote file system produces a string naming a file, and you want the corresponding pathname.

:generic-pathname of *fs:pathname* *Method*

Returns the generic pathname for the family of files of which this pathname is a member. See the section "Generic Pathnames", page 149.

The following messages get a pathname string out of a pathname object:

:string-for-printing of *fs:pathname* *Method*

Returns a string that is the printed representation of the pathname. This is the same as what you get if use **princ** or **string** on the pathname. It is the native host form of the pathname string, preceded by the name of the host and colon. This is the preferred user-visible printed representation of pathnames.

:string-for-wholine of *fs:pathname* *Method*

Returns a string that can be compressed in order to fit in the status line.

:string-for-editor of *fs:pathname* *Method*

Returns a string that is the pathname with its components rearranged so that the name is first. The editor uses this form to name its buffers.

:string-for-dired of *fs:pathname* *Method*

Returns a string to be used by the directory editor. The string contains only the name, type, and version.

:string-for-host of *fs:pathname* *Method*

Returns a string that is the pathname in the form preferred by the host file system.

:string-for-directory of *fs:pathname* *Method*

Returns a string suitable for describing the directory portion of the pathname, in the format that users of the host system are used to seeing it. The host name is not included.

The following messages manipulate the property list of a pathname:

:get indicator of *fs:pathname* *Method*

Manipulates the pathname's property list analogously to the function of the same name, which does not (currently) work on instances. See the section "Property Lists" in *Reference Guide to Symbolics-lisp*. Be careful using property lists of pathnames. See the section "Pathnames", page 127.

:getl *list-of-indicators* of **fs:pathname** *Method*

This manipulates the pathname's property list analogously to the function of the same name, which does not (currently) work on instances. See the section "Property Lists" in *Reference Guide to Symbolics-lisp*. Please take care in using property lists of pathnames. See the section "Pathnames", page 127.

:putprop *value indicator* of **fs:pathname** *Method*

This manipulates the pathname's property list analogously to the function of the same name, which does not (currently) work on instances. See the section "Property Lists" in *Reference Guide to Symbolics-lisp*. Please take care in using property lists of pathnames. See the section "Pathnames", page 127.

:remprop *indicator* of **fs:pathname** *Method*

This manipulates the pathname's property list analogously to the function of the same name, which does not (currently) work on instances. See the section "Property Lists" in *Reference Guide to Symbolics-lisp*. Please take care in using property lists of pathnames. See the section "Pathnames", page 127.

:plist of **fs:pathname** *Method*

This manipulates the pathname's property list analogously to the function of the same name, which does not (currently) work on instances. See the section "Property Lists" in *Reference Guide to Symbolics-lisp*.

The following messages can be sent to pathnames having wildcard components or suspected of having wildcard components:

:pathname-match *candidate-pathname* &optional (*match-host t*) *Method*
of **fs:pathname**

Determines whether *candidate-pathname* would satisfy the wildcard pattern of the pathname receiving the message. (The pathname receiving the message is assumed to be one that would satisfy **:wild-p**.) It compares corresponding components in the pattern pathname and *candidate-pathname*. It returns **nil** when *candidate-pathname* does not satisfy the pattern; otherwise it returns something other than **nil**.

match-host determines whether it requires the host component of the pattern to match as well. When *match-host* is **nil**, it ignores the host component. By default, it does require that the host component match.

A pattern pathname containing no wild components matches only itself.

If the *candidate-pathname* specifies a physical host, and the message is sent to a logical pathname, the physical host is "back-translated," if possible.

:wild-p of **fs:pathname** *Method*

A predicate that determines whether the pathname is syntactically a wildcard pathname. This means that any component is **:wild**, or, for most systems, contains the character *, or that the directory component has any of the

valid forms of directory wildcard in it. See the method (**:method fs:pathname :directory-wild-p**), page 173.

<i>Value</i>	<i>Meaning</i>
nil	No component of the name is syntactically a wildcard.
not nil	One or more components of the name are syntactically wild. The actual value in this case is the symbol for the most significant wild component: :device , :directory , and so on.

:device-wild-p of **fs:pathname** *Method*

If the device component of this pathname is a recognized wildcard for the system type concerned, or **:wild**, a non-**nil** is returned.

:directory-wild-p of **fs:pathname** *Method*

If the directory component of this pathname is a recognized wildcard for the system type concerned, or **:wild**, a non-**nil** is returned. All forms of wildcard at each directory level for hierarchical file systems, as well as **:wild-inferiors**, are recognized as constituting a wildcard directory component. Otherwise, **nil** is returned.

:name-wild-p of **fs:pathname** *Method*

If the name component of this pathname is a recognized wildcard for the system type concerned, or **:wild**, a non-**nil** is returned. Otherwise, **nil** is returned.

:type-wild-p of **fs:pathname** *Method*

If the type component of this pathname is a recognized wildcard for the system type concerned, or **:wild**, a non-**nil** is returned. Otherwise, **nil** is returned.

:version-wild-p of **fs:pathname** *Method*

If the version component of this pathname is a recognized wildcard for the system type concerned, or **:wild**, a non-**nil** is returned. Otherwise, **nil** is returned.

:translate-wild-pathname *target-pattern-pathname* *Method*
starting-pathname of **fs:pathname**

Produces a new pathname based on *starting-pathname* and the analogies between the pathname receiving the message and *target-pattern-pathname*.

:translate-wild-pathname examines the correspondences between *target-pattern-pathname* and the pathname receiving the message. It then does whatever is necessary to *starting-pathname* to transform it into the target pathname.

It checks to be sure *starting-pathname* matches the pathname receiving the

message and signals **error** if they do not match. A standard way for generating *starting-pathname* is to send **:directory-list** to the source pattern pathname to generate a set of starting pathnames.

12.9 Pathnames on Supported Host File Systems

This section lists the host file systems supported, gives an example of the pathname syntax for each system, and discusses any special idiosyncrasies.

12.9.1 LMFS

LMFS is an acronym for Lisp Machine File System, which is the native file system of the Symbolics computer. It is only one of many possible file systems accessible from the Symbolics computer.

LMFS is a hierarchical file system. It supports file versions. Every file has a name, type, and version. Names are virtually unlimited in length (hundreds of characters), but a performance penalty is imposed for names of over 30 characters. Types are limited to 14 characters. There is no limit to the depth of directories. There are no devices (**:device** to a LMFS pathname always returns **:unspecific**).

A LMFS pathname looks as follows:

```
>dir>ectory>name.type.version
```

The greater-than ("**>**") character separates directory levels. Absolute pathnames always start with greater-than's. Pathnames that specify no directory, relative or otherwise, contain no greater-than's, for example:

```
foo.bar.7
```

The topmost directory of the directory tree (the *ROOT* directory) is indicated by the absence of directory names but the continued presence of a greater-than. For example, the following is a file named `foo.bar`, version 7, in the *ROOT* directory:

```
>foo.bar.7
```

No file type abbreviations are needed for LMFS.

File and directory names in LMFS can be stored in upper, lower, or mixed case. Lowercase is the preferred case. Case is ignored on lookup.

Due to problems with interning of pathnames it is sometimes difficult to control the casing of a LMFS pathname, and it is almost always impossible to change it once established. See the section "Interning of Pathnames", page 138.

A version component of **:newest** is represented by the string "newest". A version component of **:oldest** is represented by the string "oldest".

Upward relativization in relative directory specifications is designated by a pathname

starting with the character less-than (" $<$ "). All and only all absolute pathnames start with the character greater-than (" $>$ "). Downward relativization is indicated by a pathname, which although it contains greater-than's, does not start with one. For example, the following specifies a directory named foo, inferior to the superior directory of the directory of the default pathname with which it is merged.

```
<foo>x.y
```

LMFS directories, when referenced as files, have a file type of "directory" and a version of 1. See the section "Directory Pathnames and Directory Pathnames as Files", page 142.

The following example specifies a directory named bar, inferior to the directory of the default pathname with which it is merged.

```
bar>x.y
```

LMFS supports recursive directory level matching (**:wild-inferiors**). The representation of **:wild-inferiors** in LMFS is ******. Any number of ****** components can appear in wildcard pathnames as directory levels, and need not be in trailing positions. (The further it gets from the trailing end of the directory name, however, the more expensive it gets to compute.) Here are some examples of the use of ******:

<i>Pathname</i>	<i>What it means</i>
-----------------	----------------------

>**>.lisp.newest	All the newest lisp files on the whole file system.
------------------	---

>**>>secret>*.*.newest	All files in subdirectories (but not top-level directories) named "secret".
------------------------	---

>lmach>**>*.*.newest	
----------------------	--

All the newest files in >lmach and all its subdirectories.

A component of **:wild**, in any component except the directory component, is represented by *****. *****, when accompanied by other characters, such as in `foo*bar*`, matches zero or more characters, as a wildcard. Although ***** or names containing ***** are valid as directory-level component names, a directory component of **:wild** cannot be specified through pathname syntax. This is because "any directory at all" is represented by (**:wild-inferiors**). A directory name given as ***** is a specification for a relative pathname, any subdirectory of the directory of the pathname which is merged. That is represented internally as (**:wild**), not **:wild**.

The name of the ROOT directory, as a file (its "directory pathname as file") is

```
>The Root Directory.directory.1
```

Names of files stored in the Lisp Machine File System can not contain *****. This restriction is necessary because ***** is used consistently to indicate wildcards in pathnames.

You can not access files whose names contain ***** as a character. A special function allows you to rename any file or directories whose names contain *****.

lmfs:rename-local-file-tool *from-path to-path* *Function*

Renames a file in which * appears in one of the pathname components. This function works locally only; you must run it on the machine in whose file system the file is stored. It does not rename a file across the network.

from-path and *to-path* must be pathnames or strings coercible to pathnames. *from-path* is parsed against a default on the local host. *to-path* is parsed against *from-path* as the default. The version number for *to-path* is inherited from the file being renamed. Any version number appearing in *to-path* is ignored.

```
(lmfs:rename-local-file-tool ">AUser>*secret-stuff*" "-secret-stuff-")
(lmfs:rename-local-file-tool ">*special*.directory.1" "-special-")
```

12.9.2 FEP File System

The syntax of FEP file system pathnames is identical to that of LMFS pathnames, and the semantics are the same as well. The following differences are to be noted.

- The maximum length of a file name is 32 characters.
- The maximum length of file types is 4 characters.
- The type of directories is "DIR".
- Recursive wildcards (:**wild-inferiors**) are not supported.

The name of the ROOT directory, as a file (its "directory pathname as file") is:

```
>ROOT-DIRECTORY.DIRECTORY.1
```

12.9.3 UNIX

Since UNIX file names can only be 14 characters long, the representations of most canonical types are stored in abbreviated form, according to the following table. Other values are represented as they are.

<i>Canonical type</i>	<i>UNIX abbreviation(s)</i>
:LISP	"l" "lisp"
:TEXT	"tx" "text" "txt"
:MIDAS	"md"
:QFASL	"qf" "qfasl"
:QBIN	"qb" "qbin"
:BIN	"bn" "bin"
:PRESS	"pr" "press"
:LGP	"lg" "lgp"
:PATCH-SYSTEM-DIRECTORY	"sd"

:PATCH-VERSION-DIRECTORY

	"pd"
:BABYL	"bb" "babyl"
:XMAIL	"xm" "xmail"
:MAIL	"ma" "mail"
:RMAIL	"rm"
:ZMAIL-TEMP	"_z" "_zmail"
:GMSGs-TEMP	"_g" "_gmsgs"
:UNFASL	"uf" "unfasl"
:OUTPUT	"ot" "output"
:ULOAD	"ul" "unload"
:MCR	"mc" "mcr"
:SYM	"sm" "sym"
:TBL	"tb" "tbl"
:MICROCODE	"mic"
:ERROR-TABLE	"err"
:FEP-LOAD	"flod"
:SYNC-PROGRAM	"sn" "sync"
:CWARNS	"cw" "cwarns"
:SYSTEM	"sy" "system"
:FONT-WIDTHS	"wd" "widths"
:BFD	"bfd"
:KST	"kt" "kst"
:AST	"at" "ast"
:PLT	"pl" "plt"
:DRW	"drw"
:WD	"wd"
:DIP	"dip"
:SAV	"sav"
:MAP	"map"
:CONSOLIDATED-MAP	"cm"
:TAGS	"tg" "tags"
:PALX-BIN	"pb" "pbin"
:XGP	"xg" "xgp"
:LIL	"ll" "lil"
:SAR	"sar"
:SAB	"sab"
:MSS	"mss" "ms"
:FORTRAN	"f"
:LOGICAL-PATHNAME-TRANSLATIONS	"lt" "logtran"
:LOGICAL-PATHNAME-DIRECTORY-TRANSLATIONS	"ld" "logdir"
:NULL-TYPE	:unspecific ""
:FILES	"fl"

:COLD-LOAD	"load"
:PXL	"px" "pxl"
:IMAGE	"im" "image"
:DUMP	"dm" "dump"

As is true with the canonical type mechanism in general, files having the canonical type spelled in full are also recognized as being of that canonical type.

Logical pathname translation must get around the restrictions in UNIX pathnames. When translating logical pathnames an extra translation step is invoked, in some cases, as for VAX/VMS pathnames.

The preferred case on UNIX is lowercase. Pathname components presented to **:new-directory**, **:new-name**, and so forth, are case-inverted in most instances. See the section "Case in Pathnames", page 145.

12.9.4 UNIX 4.2

UNIX 4.2 uses slightly different representations of some canonical types than do other versions of UNIX. In most cases, the representations are the same as for LMFS, but the UNIX versions are also allowed.

<i>Canonical type</i>	<i>UNIX 4.2 abbreviation(s)</i>
:LISP	"lisp" "l"
:TEXT	"text" "tx" "txt"
:MIDAS	"midas" "md"
:QFASL	"qfasl" "qf"
:QBIN	"qbin" "qb"
:BIN	"bin" "bn"
:PRESS	"pr" "press"
:LGP	"lgp" "lg"
:PATCH-SYSTEM-DIRECTORY	"system-dir" "sd"
:PATCH-VERSION-DIRECTORY	"patch-dir" "pd"
:BABYL	"babyl" "bb"
:XMAIL	"xmail" "xm"
:MAIL	"mail" "ma"
:RMAIL	"rmail" "rm"
:ZMAIL-TEMP	"_zmail" "_z"
:GMSGs-TEMP	"_gmsgs" "_g"
:UNFASL	"unfasl" "uf"
:OUTPUT	"output" "ot"
:ULOAD	"upload" "ul"
:MCR	"mcr" "mc"
:SYM	"sym" "sm"
:TBL	"tbl" "tb"

```

:MICROCODE      "mic"
:ERROR-TABLE    "err"
:FEP-LOAD       "flod"
:SYNC-PROGRAM  "sync" "sn"
:CWARNS        "cwarns" "cw"
:SYSTEM        "system" "sy"
:FONT-WIDTHS   "widths" "wd"
:BFD           "bfd"
:AC            "ac"
:AL            "al"
:KS            "ks"
:KST           "kst" "kt"
:AST           "ast" "at"
:PLT           "pl" "plt"
:DRW           "drw"
:WD            "wd"
:DIP           "dip"
:SAV           "sav"
:MAP           "map"
:CONSOLIDATED-MAP
                "con-map" "cm"
:TAGS          "tags" "tg"
:PALX-BIN      "palx_bin" "pbin" "pb"
:XGP           "xgp" "xg"
:LIL           "lil" "ll"
:FORTTRAN      "f"
:SAR           "sar"
:SAB           "sab"
:MSS           "mss" "ms"
:LOGICAL-PATHNAME-TRANSLATIONS
                "logtran" "lt"
:LOGICAL-PATHNAME-DIRECTORY-TRANSLATIONS
                "translations" "logdir" "ld"
:NULL-TYPE     :unspecific ""
:COLD-LOAD     "load"
:FILES         "files" "fl"
:PXL           "pxl" "px"
:IMAGE         "image" "im"
:DUMP          "dump" "dm"

```

As is true with the canonical type mechanism in general, files having the canonical type spelled in full are also recognized as being of that canonical type.

Logical pathname translation must get around the restrictions in UNIX pathnames. When translating logical pathnames, an extra translation step is invoked as for VAX/VMS pathnames.

The preferred case on UNIX is lowercase. Pathname components presented to **:new-directory**, **:new-name**, and so forth, are case-inverted in most instances. See the section "Case in Pathnames", page 145.

12.9.5 VAX/VMS

A VAX/VMS pathname looks as follows:

```
[DIR.ECTOR.OM.PONENTS]NAME.TYP;VERSION
```

The semicolon character is the standard delimiter for the version number. Because of it, a version can be specified even though the name and type are omitted. For compatibility with other Digital Equipment Corporation systems, however, a period is also accepted as a version delimiter when name and type are supplied.

Device is specified by a device name followed a colon preceding the pathname. You must take great caution with pathnames specifying devices so as not to confuse the pathname parser about host identity. See the section "Host Determination in Pathnames", page 136.

Uppercase is the only supported alphabetic case. Pathnames typed in lowercase are converted to uppercase on input.

File types are restricted to three characters, and names cannot contain hyphens. Here is a list of canonical types, their VMS representations, their default byte-size used for a binary transfer, and whether records are stored in fixed- or variable-length format:

<i>Canonical type</i>	<i>VMS representation</i>	<i>Byte-size</i>	<i>Format</i>
:LISP	"LSP"		
:TEXT	"TXT"		
:MIDAS	"MID"		
:QFASL	"QFS"	16	var
:QBIN	"QBN"	16	var
:BIN	"BIN"	16	var
:PRESS	"PRS"	8	fix
:PATCH-SYSTEM-DIRECTORY	"SPD"		
:PATCH-VERSION-DIRECTORY	"VPD"		

:BABYL	"BAB"		
:XMAIL	"XML"		
:MAIL	"MAI"		
:RMAIL	"RML"		
:ZMAIL-TEMP	"ZMT"		
:GMSGs-TEMP	"GMT"		
:UNFASL	"UNF"		
:OUTPUT	"OUT"		
:ULOAD	"ULD"		
:MCR	"MCR"		
:SYM	"SYM"		
:TBL	"TBL"		
:MICROCODE	"MIC"	8	var
:ERROR-TABLE	"ERR"		
:FEP-LOAD	"FLD"		
:SYNC-PROGRAM	"SYN"		
:CWARNS	"CWN"		
:SYSTEM	"SYD"		
:FONT-WIDTHS	"WID"	16	fix
:BFD	"BFD"	16	var
:KST	"KST"	9	
:AC	"AC"	16	
:AL	"AL"	16	

:KS	"KS"	16	
:AST	"AST"		
:PLT	"PLT"	9	
:DRW	"DRW"	12	
:WD	"WD"	12	
:DIP	"DIP"	12	
:SAV	"SAV"	12	
:MAP	"MAP"		
:CONSOLIDATED-MAP	"CON"		
:TAGS	"TAG"		
:PALX-BIN	"PXB"	8	var
:XGP	"XGP"		
:LIL	"LIL"		
:FOR	"FOR"		
:SAR	"SAR"		
:SAB	"SAB"	8	
:MSS	"MSS"		
:LOGICAL-PATHNAME-TRANSLATIONS	"LTR"		
:LOGICAL-PATHNAME-DIRECTORY-TRANSLATIONS	"LDT"		
:NULL-TYPE	""		
:COLD-LOAD	"LOD"	16	var
:FILES	"FLS"		
:PXL	"PXL"	8	

:IMAGE	"IMG"	
:DUMP	"IDM"	16

Logical pathname translation must get around the restrictions in VMS pathnames, including the prohibition against hyphens and restrictions on the lengths of components. When translating logical pathnames an extra translation step is performed before trying the usual translations. See the file sys:sys;sys.logtran.

The VMS pathname mechanism supports recursive directory matching (**:wild-inferiors**). The representation for a directory level component of **:wild-inferiors** is ".."; however it can appear only at the end of a directory name. Thus, the following matches any file in [A.B], or any subdirectory thereof:

```
[A.B...]*.*.*
```

Upward relativization in pathnames is specified by one or more minuses ("-") as the first directory name. Downward relativization is represented by a null (0-character) first directory name. For example, the following specifies a directory named FOO, inferior to the superior directory of the directory of the default pathname with which it is merged.

```
[-.FOO]X.Y
```

A pathname version component of **:newest** is specified by a version of 0 in the filename string. There is no VMS implementation of **:oldest**.

The percent sign (%) can be used in VMS wildcards to specify the matching of a single character.

The pathname system does not recognize logical device names. They are specified as device names, and are resolved by VMS, not the pathname system. There may be problems defaulting the directory specification of VAX/VMS pathnames when logical devices are used.

VMS directories, when referenced as files, have a type of "DIR", and a version of 1. See the section "Directory Pathnames and Directory Pathnames as Files", page 142.

12.9.6 TOPS-20 and TENEX

A TOPS-20 pathname has the form:

```
HOST:DEVICE:<DIRECTORY>NAME.TYPE.VERSION
```

The default device is PS:.

TOPS-20 pathnames are mapped to uppercase. Special characters (including lowercase letters) are quoted with the circle-x (⊗) character, which has the same character code in the Symbolics character set as control-V in the TOPS-20 character set.

TOPS-20 pathnames represent versions of **:oldest** and **:newest** by the strings `"..-2"` and `"..0"`, respectively.

The directory component of a TOPS-20 pathname is a list of directory level components. The directory `<FOO.BAR>` is represented as the list `("FOO" "BAR")`.

The TOPS-20 init file naming convention is `"<user>program.INIT"`.

When there not enough room in the status line to display an entire TOPS-20 file name, the name is truncated and followed by a center-dot character to indicate that there is more to the name than can be displayed.

TENEX pathnames are almost the same as TOPS-20 pathnames, except that the version is preceded by a semicolon instead of a period, the default device is DSK instead of PS, and the quoting requirements are slightly different.

12.9.7 Multics

Multics possesses a hierarchical file system. Every file has a name, and may or may not have a type. Multics does not support file versions. The sum of the lengths of name and type and the period required to separate them must not exceed 32 characters. A maximum of 16 directory levels is supported. There are no devices (**:device** to a Multics pathname always returns **:unspecific**). A Multics pathname looks as follows:

```
>dir>ectory>name.type
```

The greater-than ("`>`") character separates directory levels. Absolute pathnames always start with greater-than's. Pathnames that specify no directory, relative or otherwise, contain no greater-than's, for example:

```
foo.bar
```

The topmost directory of the directory tree (the *ROOT* directory) is indicated by the absence of directory names but the continued presence of a greater-than. For example, the following is a file named `foo.bar`, in the *ROOT* directory:

```
>foo.bar
```

There are no file type abbreviations needed for Multics.

File and directory names can be stored in upper, lower, or mixed case. Lower case is the preferred case. Case is significant. `foo`, `FOO`, and `foo` could well be the names of three different files in the same directory.

Upward relativization in relative directory specifications is designated by a pathname starting with the character less-than ("`<`"). All and only all absolute pathnames start with the character greater-than ("`>`"). Downward relativization is indicated by a pathname, which although it contains greater-than's, does not start with one. For example, the following specifies a directory named `foo`, inferior to the superior directory of the directory of the default pathname with which it is merged.

<foo>x.y

Multics directories, when referenced as files, have no specific type; they need not have any type at all. See the section "Directory Pathnames and Directory Pathnames as Files", page 142.

The following example specifies a directory named bar, inferior to the directory of the default pathname with which it is merged.

bar>x.y

Multics does not support **:wild-inferiors**, that is, recursive directory-level matching. For that matter, Multics does not support *any* form of wildcard in the directory component of a pathname. (Although **:pathname-match** matches such components, Multics does not support them in directory lists.) A component of **:wild**, in any component except the directory component, is represented by *****. *****, when accompanied by other characters, such as in **foo*bar***, matches zero or more characters, as a wildcard.

12.9.8 ITS

An ITS pathname looks like "HOST: DEVICE: DIR; FOO 69". The default device is DSK: but other devices such as ML:, ARC:, DVR:, or PTR: can be used.

ITS does not exactly fit the virtual file system model, in that a file name has two components (FN1 and FN2) rather than three (name, type, and version). Consequently to map any virtual pathname into an ITS filename, it is necessary to choose whether the FN2 will be the type or the version. The rule is that usually the type goes in the FN2 and the version is ignored; however, certain types (LISP and TEXT) are ignored and instead the version goes in the FN2. Also if the type is **:unspecific** the FN2 is the version.

An ITS filename is converted into a pathname by making the FN2 the version if it is "<", ">", or a number. Otherwise the FN2 becomes the type. ITS pathnames allow the special version symbols **:oldest** and **:newest**, which correspond to "<" and ">" respectively. If a version is specified, the type is always **:unspecific**. If a type is specified, the version is **:unspecific** so that it does not override the type.

Each component of an ITS pathname is mapped to uppercase and truncated to six characters.

Special characters (space, colon, and semicolon) in a component of an ITS pathname can be quoted by prefixing them with right horseshoe (>) or equivalence sign (≡). Right horseshoe is the same character code in the Symbolics character set as control-Q in the ITS character set.

The ITS init file naming convention is "homedir; user program".

fs:*its-uninteresting-types**Variable*

The ITS file system does not have separate file types and version numbers; both components are stored in the "FN2". This variable is a list of the file types that are "not important"; files with these types use the FN2 for a version number. Files with other types use the FN2 for the type and do not have a version number.

It is not possible to have two ITS pathnames with the same meaning that differ in an ignored component. **fs:*its-uninteresting-types*** controls which types are ignored in favor of retaining version numbers. The following table summarizes the interaction of type and version components for ITS pathnames.

<i>Type</i>	<i>Version</i>	<i>Result</i>
supplied	omitted	type is retained, version is :unspecific
omitted	supplied	type is :unspecific , version is retained
"interesting"	supplied	type is retained, version is :unspecific
"uninteresting"	supplied	type is :unspecific , version is retained

:fn1 of fs:its-pathname*Method*

This message returns a string that is the FN1 host-dependent component of the pathname.

:fn2 of fs:its-pathname*Method*

This message returns a string that is the FN2 host-dependent component of the pathname.

12.9.9 Ms-dos Pathnames

An MS-DOS pathname looks like this:

```
HOST:DEVICE:\DIRECTORY\NAME.TYPE
```

The default device is C:. Uppercase is the only supported case. Pathnames typed in lowercase are converted to uppercase on input.

File names and directory components are restricted to eight characters. File types are restricted to three characters. The canonical types for MS-DOS are the same as for VAX/VMS.

Relative pathnames are permitted. Upward-level changes are signalled with "..". For example:

```
PC:A:..\..\DIR\FILE.LSP
```

12.9.10 Logical Pathnames

A *logical pathname* does not correspond to a particular file server; its host is called a *logical host*. Every logical pathname can be translated into a corresponding "physical" pathname; a mapping from logical hosts into physical hosts is used to effect this translation.

Logical pathnames make it easy to move bodies of software to more than one file system. An important example is the body of software that constitutes the Symbolics-Lisp system. Any site may have a copy of all of the sources of the programs that are loaded into the initial Lisp environment. Some sites store the sources on a LMFS file system, while others store them on a VAX/VMS system. However, other software in the system must use pathnames for these files in such a way that the software will work correctly at all sites. This is accomplished with a logical host called SYS; all pathnames for system software files are actually logical pathnames with host SYS. At each site, SYS is defined as a logical host, but translation is different at each site. For example, at a site where the sources are stored on a certain VAX/VMS system, pathnames of the SYS host are translated into pathnames for that system.

You usually use logical pathnames when you are defining a system that you wish to be portable to other sites. All logical pathnames in your system should translate to a valid pathname on any kind of host to which the system might be distributed. (Currently, this includes LMFS (Symbolics), VAX/VMS, Unix, and TOPS-20). The converse is not true; logical pathnames make no attempt to provide a way to represent all pathnames on a particular host. For this reason, no way is provided to distinguish between "foo" and "foo.", or "foo" and "FOO" on UNIX. Your software will be much more portable if you choose good logical pathnames for your files rather than trying to make the logical pathnames conform to the limitations of whatever filesystem you happen to store your system on. For example, even though logical pathnames have a quoting character, it is good practice to avoid using it.

Here, roughly, is how translation is done: To translate a logical pathname, the system finds the mapping for that pathname's host and looks up that pathname's directory in the mapping. If the directory is found, a new pathname is created whose host is the physical host, and whose device and directory come from the mapping. The other components of the new pathname are left the same.

This means that when you invent a new logical host for a certain set of files, you also make up a set of logical directory names, one for each of the directories that the set of files is stored in. Now when you create the mappings at particular sites, you can choose any physical host for the files to reside on, and for each of your logical directory names, you can specify the actual directory name to use on the physical host. This gives you flexibility in setting up your directory names; if you used a logical directory name called fred and you want to move your set of files to a new file server that already has a directory called fred, being used by someone else, you can translate fred to some other name and so avoid getting in the way of the existing directory. Furthermore, you can set up your directories on each host to conform to the local naming conventions of that host.

However, a logical pathname host can have the same name as a physical host: See the section "Specifying a New Logical Host Name".

A logical pathname has the form HOST: DIRECTORY; NAME.TYPE.VERSION.

On input, spaces can separate the name, type, and version. There is no way to specify a device; parsing a logical pathname always returns a pathname whose device component is **:unspecific**. This is because devices have no meaning in logical pathnames. Logical pathnames can be hierarchical; directory levels are separated by semicolons.

Logical pathnames can be *relative*. That is, a pathname can have a directory component whose meaning is "when merging against a default, append these directories onto the end of any default directories." The syntax for this is **HOST: ; DIRECTORY; NAME.TYPE.VERSION**, that is, a leading bare **;** before the directory component. Thus, the above pathname, merged against a default of **HOST: USER; FOO.LISP.NEWEST** gives **HOST: USER; DIRECTORY; NAME.TYPE.VERSION**.

The equivalence-sign character (**=**) can be used for quoting special characters such as spaces and semicolons. (The use of this character is discouraged, however, as such files are unlikely to be transportable). The double-arrow character (**=**) can be used as a place-holder for unspecified components. Components are not mapped to uppercase. The **:newest**, **:oldest**, and **:wild** values for versions are specified with the strings **NEWEST**, **OLDEST**, and ***** respectively. On input, **:newest** can be represented by **>** and **:oldest** by **<**.

There is no init file naming convention for logical hosts; you cannot log into them. The **:string-for-host**, **:string-for-wholine**, **:string-for-dired**, and **:string-for-editor** messages are all passed on to the translated pathname, but the **:string-for-printing** is handled by the **fs:logical-pathname** flavor itself and shows the logical name.

12.9.10.1 Logical Pathname Wildcard Syntax

Logical pathnames support a wildcard syntax meaning "Match any directory, and any subdirectory, to any level." For example:

```
Show Directory SYS:**;*.BFD.*
```

Here, the Show Directory command lists all font files anywhere in the SYS hierarchy, to any level.

This corresponds to the **>**>** syntax for LMFS pathnames, and the **[name...]** syntax for VAX/VMS file specifications. See the section "LMFS Pathnames", page 174. See the section "VAX/VMS Pathnames", page 180.

This makes it easy to specify logical pathname translations on Lisp Machines and VAX/VMS. For example:

```
(fs:set-logical-pathname-host "SYS"
 :translations '(("SYS:**;*.*.*" "ACME-LISPM:>Rel-6>**>*.*.*)" ))
```

```
(fs:set-logical-pathname-host "SYS"
 :translations
 '(("SYS:**;*. *.*" "ACME-VMS:SYMBOLICS:[REL6...]*.*;*))
 :no-translate nil)
```

For more information about the argument **:no-translate**: See the section "Translation Rules" in *Release 6.0 Release Notes*.

It is important to note that wherever a "***;" appears in the logical-host pathname, there must be a corresponding "wild-inferiors" pathname on the physical-host pathname.

UNIX and TOPS-20 do not have a syntax with this meaning. For these hosts, it is necessary to list explicitly each level of directory to be translated. For example:

```
(fs:set-logical-pathname-host "SYS"
 :translations
 '(("SYS:*;*. *.*"
   "ACME-UNIX://usr//symbolics//rel-6//**/*.*")
 ("SYS:*;*;*. *.*"
   "ACME-UNIX://usr//symbolics//rel-6//**/**/*.*")
 ("SYS:*;*;*;*. *.*"
   "ACME-UNIX://usr//symbolics//rel-6//**/**/**/*.*")
 ("SYS:*;*;*;*;*. *.*"
   "ACME-UNIX://usr//symbolics//rel-6//**/**/**/**/*.*"))
 :no-translate nil)
```

12.9.10.2 Logical Pathname Translation

This section explains the format of the "translations" list of logical pathnames and the rules for translating a logical pathname to a physical pathname.

Each element of the list (one translation) specifies two wildcard pathnames, the first on the logical host and the second on the physical. In the Lisp form (in the file `sys:site;host.translations`) that specifies this form, they are given as strings to be parsed against these respective hosts. As they are parsed, they are merged with a pathname of wild name, wild type, and wild version.

Following is an example of a translations list. This is a sample LMFS translation table for the SYS host, slightly more complex than the default:

```
'(("SYS:DOC;**;*. *.*" "ACME-S:>Rel-6>doc>**>*. *.*")
 ("SYS:**;*. *.*" "ACME-Q:>Rel-6>**>*. *.*"))
```

There are two phases to the translation process. In the first phase, a logical/physical pathname pair is found in the translation table. This pair is called the *translation pair*.

In the second phase, this pair is presented to a *translation rule* to be processed. Normally, this rule uses **:translate-wild-pathname** to translate the pathname using the translation pair, but there is a wide variety of translation rules.

The first phase consists of matching the pathname to be translated against each first element of each translation, in succession. (The **:pathname-match** message is used.) The order in the list is thus very important. The first match is then taken to be the translation pair for the second phase.

When the physical host supports a syntax for **:wild-inferiors** (for example, **>*>** on LMFS), that syntax can be used to have a translation that matches "everything else", as in the example above. If no equivalent syntax is supported, a separate wildcard directory for each level of directory likely to be encountered serves the same purpose, as in the example below.

```
'(("SYS:*;*.*.*" "ACME-UNIX://usr//rel-6//**/*.*)"
  ("SYS:*;*;*.*.*" "ACME-UNIX://usr//rel-6//**/**.*)"
  ("SYS:*;*;*;*.*.*" "ACME-UNIX://usr//rel-6//**/**/**.*)" )
```

This example handles any SYS pathname with up to three directory levels. In the presence of such a translation, it is impossible to have an undefined translation.

The second phase is potentially more complex. In its simplest form, it reduces to producing the translated pathname by sending the **:translate-wild-pathname-reversible** message to the logical pathname, with the first element of the translation as the source pattern and the second element of the translation as the target pattern. See the section "Wildcard Pathname Mapping", page 153. See the section "Wildcard Directory Mapping", page 155. See the section "Reversible Wildcard Pathname Translation", page 191.

However, before deciding to using **:translate-wild-pathname-reversible**, a search is done to find a more suitable rule for performing the translation. With each logical host, there are three sets of translation rules. In addition, there is a global set of rules, and a default.

Here is the order in which these rules are searched:

Permanent	The permanent translation rules are special purpose rules that cannot be overridden. They provide for such things as the translation of patch file pathnames. This table is searched first.
Site	The site translation rules are provided to override the supplied translation rules at a specific site.
Supplied	The normal, supplied translation rules are normally supplied by the author of the software using the logical host.
Global	This is a set of rules independent of any particular host. This table is not currently used for anything, but it is provided for future extension.
Default	This is the rule used when no other rule is found. This is the :translate-wild rule, which uses L[:translate-wild-pathname-reversible] to translate according to the translation pair found in phase 1 of the translation process.

Back-translation is performed by searching the second elements of the translations list, and translating in the other direction. **:translate-wild-pathname-reversible** is always used for this, so it is not guaranteed to come up with the same logical pathname as might be expected.

Reversible Wildcard Pathname Translation

A special version of wild pathname translation, called "reversible wild pathname translation," is used. The difference between regular wild pathname translation and reversible translation is in the treatment of a target wildcard pattern consisting solely of *. In regular translation, a target pattern of **:wild** causes the source component to be copied verbatim. This is a useful user-interface feature, but it causes dropping of information and resultant noninvertibility of the transformation. In reversible mapping, this feature is not present. Logical pathname translation and back-translation is done in this mode.

Example:

<i>Type</i>	<i>Source pattern</i>	<i>Source instance</i>	<i>Target pattern</i>	<i>Result</i>
Regular	foo*	foolish	*	foolish
Reversible	foo*	foolish	*	lish
Either	*	bar	foo-*	foo-bar

Note that the inverse translation of foo-bar to bar cannot be accomplished under regular translation.

Defining a Translation Rule

Translation rules are defined using the **fs:set-logical-pathname-host** function, using the **:translation-rule** or **:site-translation-rule** argument. (The other rule tables are not normally set by the user). These arguments should be an alist of system type and translation rule specifications.

```
((:vms VMS rule specifications ...)
 (:vms4 VMS4 rule specifications ...)
 (:unix UNIX rule specifications ...))
```

Each rule specification consists of a *pattern*, a *rule type*, and optional arguments, as in the following example.

```
("PICTURE:EDITOR;LINE-DRAWING-COMMANDS.*.*" :vms-new-pathname :name "LINECMNDS")
```

In this example, **"PICTURE:EDITOR;LINE-DRAWING-COMMANDS.*.*"** is the pattern, **:vms-new-pathname** is the rule type, and **:name** and **"LINECMNDS"** form a keyword/value pair of arguments to the **:vms-new-pathname** rule type.

Normally, translation rules are defined in the system definition file before a **defsystem** form, so that the rules are loaded before they are needed. If you wish to override the translation rules provided either by Symbolics or another vendor, you

can use the **:site-rules** argument to the call to **fs:set-logical-pathname-host**, normally placed in the translation file.

The following sections describe the various rule types that exist and their arguments.

:translate-wild &rest *options*

Translation Rule

The default translation rule's type is **:translate-wild**. This simply sends the source pattern a **:translate-wild-pathname-reversible** with the target pattern as target and the pathname being translated as the source pathname. For example:

contents of sys.translations file:

```
(fs:set-logical-pathname-host "SYS"
 :translations '(("SYS:DOC;**/*.*)" "S:>Rel-6>doc>*.*)"
 ("SYS:**/*.*)" "Q:>Rel-6>**>*.*)"'))
```

pathname to translate:

SYS:IO;PATHNM.LISP.23

translation pair found in phase 1:

("SYS:**/*.*)" "ACME:>Rel-6>**>*.*)"

result of translation:

ACME:>Rel-6>io>pathnm.lisp.23

In other words, the default is for the translation to occur according to the wildcard mapping given in the translations.

:new-pathname &key *device directory name type version*

Translation Rule

The **:new-pathname** translation rule type is similar to **:translate-wild**, but replaces the *directory*, *name*, *type*, or *version*. Any components not specified in the argument list will not be replaced, and will be derived via **:translate-wild-pathname-reversible** as for the **:translate-wild** translation rule type.

:vms-heuristicate &optional *substitute*

Translation Rule

This translation rule tries hard to make understandable VMS pathnames out of longer, hyphenated filenames. It works for both :VMS and :VMS4 hosts. It produces usually understandable, hopefully unique, legal names and directories. In operation, it is similar to the **:translate-wild** type, but the components translated by wildcards are subjected to heuristics if needed to fit VMS's pathname syntax.

The *substitute* argument is used to perform character substitutions. For example, for VMS4, it can be used to substitute "_" for "-".

```
("SYS:**;*.*.*" :vms-heuristicate ((#\ - #\_)))
```

:vms-heuristicate-name &optional *substitute* *Translation Rule*

:vms-heuristicate-name is like **:vms-heuristicate**, but only heuristics the name.

:vms-heuristicate-directory &optional *substitute* *Translation Rule*

:vms-heuristicate-directory is like **:vms-heuristicate**, but only heuristics the directory name.

:vms-new-pathname &key *device directory name type version* *Translation Rule*

The **:vms-new-pathname** translation rule is a cross between **:new-pathname** and **:vms-heuristicate**. Components not explicitly specified in the argument list are supplied by wildcard mapping plus heuristics as for **:vms-heuristicate**.

:vms-font &optional *renamings* *Translation Rule*

The **:vms-font** translation rule *parses* the name component of the logical pathname as a font spec. For example, in **timesroman10b**, the *font name* is **timesroman**, the *font size* is **10**, and the *face code* is **b**. (The face code is optional). The font name is subjected to the VMS heuristics to fit in a smaller space (to allow room for the font size and face code). The result is concatenated with the font size and face code to construct a new name.

If the *renamings* argument is supplied, it is an alist of font names and replacement to be used instead of the one produced by the heuristics. This is useful in cases where the heuristic produces a confusing name, or where there would otherwise be name conflicts. For example, the following translation rule is used with the SYS: host for VMS hosts.


```

("SYS: FONTS; LGP-1; *.BFD.*" :vms-font
 ("DANG-MATH" "DANGM")
 ("GHELVETICA" "GHLVT")
 ("HELVETICA" "HELV")
 ("TIMESROMAN" "TIMSR")
 ("XGP-VGV" "XGPVV"))

```

This translation rule serves to encode the relevant information that makes each font distinct.

In addition, **:vms-font** performs full VMS heuristics on the directory.

:vms-microcode

Translation Rule

This translation rule encodes the microcode names in such a way as to be sure to retain the information that distinguishes different microcodes.

The name component of the logical pathname is parsed into words. Each word is looked up in the alist ***unix-microcode-translation-alist***. (The alist is shared with the equivalent translation for UNIX). If found, it is replaced with the replacement (a single character, except "MIC" maps to "") found in the second element of the alist bucket. This sequence of characters is then concatenated to produce the new filename.

The directory component is subject to full heuristics.

```

:tops20-heuristicate-directory &optional
  (levels fs:*default-tops20-directory-levels*)

```

Translation Rule

The **:tops20-heuristicate-directory** translation rule compensates for the fact that TOPS20 directories are limited to a size of 39, including the "." characters as directory-level separators. Each level of directory is allocated a share of the available space, and is compressed to fit in that space as needed. In determining how much space to allocate to each level, the rule assumes that no more than *levels* directory levels will be needed. The default is **fs:*default-tops20-directory-levels***, or 3 levels.

:unix-microcode

Translation Rule

This translation rule encodes the microcode names in such a way as to be sure to retain the information that distinguishes different microcodes.

The name component of the logical pathname is parsed into words. Each word is

looked up in the alist ***unix-microcode-translation-alist***. (The alist is shared with the equivalent translation for VMS). If found, it is replaced with the replacement (a single character, except "MIC" maps to "") found in the second element of the alist bucket. This sequence of characters is then concatenated to produce the new filename.

:unix-font &optional *renamings*

Translation Rule

The **:unix-font** translation rule *parses* the name component of the logical pathname as a font spec. For example, in **timesroman10b**, the *font name* is **timesroman**, the *font size* is **10**, and the *face code* is **b**. (The face code is optional). The font name is subjected to the UNIX heuristics to fit in a smaller space (to allow room for the font size and face code). The result is concatenated with the font size and face code to construct a new name.

If the *renamings* argument is supplied, it is an alist of font names and replacement to be used instead of the one produced by the heuristics. This is useful in cases where the heuristic produces a confusing name, or where there would otherwise be name conflicts. For example, the following translation rule is used with the SYS: host for UNIX hosts.

```
("SYS: FONTS; LGP-1; *.BFD.*" :unix-font
 ("DANG-MATH" "DANGMT")
 ("GHELVETICA" "GHELV")
 ("HELVETICA" "HELVET")
 ("TIMESROMAN" "TIMESR")
 ("XGP-VGV" "XGPPGV")))
```

This translation rule serves to encode the relevant information that makes each font distinct.

:unix-type-and-version &optional *renamings*

Translation Rule

The **:unix-type-and-version** translation rule is used for situations where you need to retain both the type and version. This is usually needed where differing versions of the file need to coexist.

The name component is matched against the *renamings* alist. If it is found, the second element of the alist bucket is used instead. Then, if the last character of the name (or the replacement) is a digit, a "+" is added to the end. Then, the version number (or "" if **nil** or "*" if **wild**) is added to the end. This is then used as the name component. The type is handled via the normal mechanisms.

The version is added to the name rather than the end of the type, so that the type field can be recognized by programs that look at the type (or canonical type).

:site-directory &key device directory name type version *Translation Rule*

The **:site-directory** translation rule substitutes the **:site-directory** attribute from the local site object for the host and directory. The arglist is like for **:new-pathname**. This is used to translate SYS:SITE;

As a special feature, this rule can be overridden by an explicit entry for SYS: SITE; in the translations. This can be useful when debugging, to get a different site directory without modifying your site namespace object.

fs:patch-file system-name &optional file-type *Translation Rule*

fs:patch-file rules, which will often be seen when doing a **fs:describe-logical-host**, are internal to the patch system. They provide for the translation of patch file logical names to physical files, in a system-dependent manner. These rules are added as a result of defining a system to be patchable.

fs:describe-logical-host host *Function*

The **fs:describe-logical-host** function takes a logical host (or the name of a logical host) and provides various information about the host, including:

- Default physical host.
- Translations.
- Translation rules sorted by search order.
- Translation rules sorted by group.

It is often useful for determining what went wrong with a translation file.

12.9.10.3 Splitting Logical Hosts Across Physical Hosts

It is possible to have a logical host translate to more than one physical host. All that is needed is an explicit specification of the hosts involved, in the translation list given to **fs:set-logical-pathname-host**. For example:

```
(fs:set-logical-pathname-host "SYS"
 :translations '(("SYS:DOC;**;*.*.*" "ACME-LISPM:>Rel-6>doc>**>*.*.*")
 ("SYS:**;*.*.*" "ACMEVAX:SYMBOLICS:[REL6...]*.*.*"))
 :no-translate nil)
```

Note that it is not necessary to specify the **:physical-host** argument to **fs:set-logical-pathname-host** as long as the host names are specified in the

translation list. If the argument is specified, it serves as a default when parsing those pathnames.

fs:make-logical-pathname-host *name* &key *no-search-for-shadowed-physical* *Function*

Defines *name*, which should be a string or symbol, to be the name of a logical pathname host. *name* should not conflict with the name of any existing host, logical or physical.

(fs:make-logical-pathname-host) also loads the file `sys:site;name.translations` and arranges for that file to be reloaded in the future. **load-patches** checks the translations file for each logical host that is defined in the current world; if any file has been changed it is reloaded. **load-patches** does this if and only if no specific systems are specified in its arguments.

fs:make-logical-pathname-host alters the **logical-pathnames-translation-files** system so that it contains the translations files for all logical hosts defined in the current world. **load-patches** loads updated translations files by calling **make-system** on this system.

An **fs:make-logical-pathname-host** form often appears in the file `sys:site;system-name.system`. **make-system** looks for this file when given the name of an unknown system. The **fs:make-logical-pathname-host** form must be the first form in the file, as the second form, a call to **si:set-system-source-file**, depends on the previous definition of the logical host.

Example: Following are the contents of the file `sys:site;cube.system`:

```
;;; -*- Mode: LISP; Package: USER -*-

(fs:make-logical-pathname-host "cube")
(si:set-system-source-file "cube" "cube: cube; cubpkg")
```

The argument **:no-search-for-shadowed-physical** (default **nil**) means to look only in the existing pathname hosts for a host with the same name as the logical host. This saves time by not asking the namespace server whether the name of the newly defined logical host conflicts with the names of any physical hosts, but it prevents you from seeing the following warnings:

```
Warning: the host ~A must now be referred to as ~A: in pathnames,
        since ~A is now a logical pathname host.
        This affects ~[no~:;~:*~D~] extant pathnames.
```

```
Warning: the nickname ~A: for the physical host ~A
        will now refer instead to the
        logical pathname host ~A.
        Use ~A: in pathnames.
```

fs:add-logical-pathname-host is an obsolete name for this function.

fs:set-logical-pathname-host *logical-host &key physical-host* *Function*
translations rulesite-rules (no-translate t)
no-search-for-shadowed-physical

Creates a logical host named *logical-host* if it does not already exist. It then establishes translations of logical directories on *logical-host* to physical directories on various hosts. (*physical-host* serves as a default.) *translations* is a "translations" list of two-element lists of strings representing associated logical directories (source patterns) and physical directories (target patterns). For the format of the lists and the translation rules: See the section "Logical Pathname Translation", page 189.

Source patterns are logical pathnames that are matched against the pathname being translated. The target patterns are physical pathnames and can be on any host.

If the physical pathname is on a TOPS-20 or VAX/VMS host, you should include the device name. In the case of VAX/VMS, it is important that this device name be either a physical device name or the name of a "concealed device." The simplest way to choose a device name is to connect to the VAX/VMS system in question. If you want to use FOO:[BAR...]*.*.* as the target, where FOO is a VAX/VMS system-wide logical name, connect to VAX/VMS and do the following:

```
$ DIRECTORY FOO:[BAR...]*.*.*

Directory USER$DISK:[BAR...]
...
```

In this example, you should use USER\$DISK:[BAR...] instead of FOO:[BAR...] in your translations.

If *no-translate* is **nil**, the translation of every interned logical pathname is checked. Properties are copied from the old physical pathname to the the new one, and logical pathnames that now have no corresponding physical pathnames are uninterned. If *no-translate* is not **nil** or not supplied, this mapping is suppressed, and some physical pathnames might not get the properties of the logical pathname. This is not normally of any consequence, so *no-translate* defaults to **t**.

A call to **fs:set-logical-pathname-host** is usually the only form in the file `sys:site;logical-host.translations`. This file is loaded by **fs:make-logical-pathname-host** (always in the **file-system** package), which also arranges for it to be reloaded in the future. **load-patches** checks this file for all logical hosts in the current world and reloads the file if it has changed. Similarly, changing the site object will cause each translation file to be reloaded from the new site directory.

The argument *no-search-for-shadowed-physical* (default **nil**) means to look only in the existing pathname hosts for a host with the same name as the logical host. This saves time by not asking the namespace server whether the name of the newly defined logical host conflicts with the names of any physical hosts, but it prevents you from seeing the following warnings:

Warning: the host ~A must now be referred to as ~A: in pathnames,
since ~A is now a logical pathname host.
This affects ~[no~:;~:*~D~] extant pathnames.

Warning: the nickname ~A: for the physical host ~A
will now refer instead to the
logical pathname host ~A.
Use ~A: in pathnames.

Example:

Following is a typical content of the file `sys:site;sys.translations`:

```
;; -*- Mode: LISP; Package: FILE-SYSTEM -*-

(set-logical-pathname-host "sys"
 :translations '(("**;*.*.*" ">Re1-6>**>")))
```

:translated-pathname of **fs:logical-pathname** *Method*

Converts a logical pathname to a physical pathname. It returns the translated pathname of this instance: a pathname whose **:host** component is the physical host that corresponds to this instance's logical host. See the section "Logical Pathnames", page 186.

If this message is sent to a physical pathname, it simply returns itself.

:back-translated-pathname *pathname* of **fs:logical-pathname** *Method*

Converts a physical pathname to a logical pathname. *pathname* should be a pathname whose host is the physical host corresponding to this instance's logical host. This returns a pathname whose host is the logical host and whose translation is *pathname*. See the section "Logical Pathnames", page 186.

This message might be used in connection with `truename`. Given a stream that was obtained by opening a logical pathname,

```
(send stream :pathname)
```

returns the logical pathname that was opened.

```
(send stream :truename)
```

returns the true name of the file that is open, which of course is a pathname on the physical host. To get this in the form of a logical pathname, one would do the following:

```
(send (send stream :pathname)
      :back-translated-pathname
      (send stream :truename))
```

If this message is sent to a physical pathname, it simply returns its argument. Thus the above example works no matter what kind of pathname was opened to create the stream. However, it is important to note two situations in which back translation can fail to do what you expect:

Links If opening the file involved following a link, the truename will no longer match, and back translation might not be able to convert it to a physical pathname at all.

File-system restrictions

If the translation involved compressing or modifying a name to adapt to a file-system's rules, the physical pathname may be translated to a logical pathname different from the one originally used.

Back translation is useful only in cases where the logical pathname is wanted for informational, not operational, purposes. For example, if you remember a back translation to reopen the file, you may end up with physical instead of logical pathnames in your program. Physical pathnames are not transportable between sites.

One way to avoid this problem is to avoid back translation. Often, all that is needed is the version number, in which case the following code will serve:

```
(send (send stream :pathname)
      :new-default-pathname
      :version (send (send stream :truename) :version))
```

Note that **:new-default-pathname** is used rather than **:new-pathname**. This is necessary because the logical host and the physical host are of different types. When copying components between host types, you need to allow for certain substitutions. In this case, if the physical host is a UNIX system, the version will be **:unspecific**, and **:new-default-pathname** will convert this to the nearest equivalent for logical pathnames: **:newest**.

12.10 Init File Naming Conventions

Init files are of canonical type **:lisp** for source files and **:bin** for compiled files. For hosts that support long file names, the init file name consists of *program-name* with "-INIT" appended. Thus, the standard file name for a 3600-family init file is LISPM-INIT; for a Zmail init file, it is ZMAIL-INIT. Hosts that do not support long file names have conventions peculiar to each system.

Following are the names of lisp init source files on some hosts:

<i>Host system</i>	<i>File name</i>
LMFS/TOPS-20	LISPM-INIT.LISP
UNIX	lispm-init.l
VMS	LISPMINI.LSP
Multics	lispm-init.lisp
ITS	If user has own directory: LISPM >. If user does not have own directory: <i>USER</i> LISPM.

12.11 Maclisp Conversion

This section briefly discusses how to convert from Maclisp I/O and filename functions to the corresponding but often more general Symbolics-Lisp versions.

The functions **load**, **open**, **probef**, **renamef**, and **deletef** are upward compatible. Most of them take optional additional arguments to do additional things, usually connected with error handling. Where Maclisp expects a file name in the form of a symbol or a list, the Symbolics system accepts those or a string or a pathname object. **probef** returns a pathname or **nil** rather than a namelist or **nil**.

load keeps defaults, which it updates from the file name it is given.

The old-I/O functions **uread**, **crunit**, and so on do not exist in the Symbolics system. **fasload** exists but is a function rather than a special form.

There is a special form, **with-open-file**, which should replace most calls to **open**. See the function **with-open-file**, page 91.

The functions for manipulating file names themselves are different. The system accepts a namelist as a pathname, but never creates a namelist. **mergef** is replaced by **fs:merge-pathname-defaults**. **defaultf** is replaced by **fs:default-pathname** or **fs:set-default-pathname**, depending on whether it is given an argument. **namestring** is replaced by the **:string-for-printing** message to a pathname, or the **string** function. **namelist** is approximately replaced by **fs:parse-pathname**. **(status udir)** and **(status homedir)** are approximately replaced by **fs:user-homedir**. The **truename** function is replaced by the **:truename** stream operation, which returns a pathname containing the actual name of the file open on that stream. The **directory** and **allfiles** functions are replaced by **fs:directory-list**.

13. Lisp Machine File System

13.1 Introduction to LMFS

The Lisp Machine File System (LMFS) provides a file system that runs on a Symbolics computer and stores information on the computer's disks. The information can be accessed locally (from that computer itself) or from other computers.

For information on performing I/O on files: See the section "Streams", page 1.

This discussion does not describe the internal program logic or organizational details of the file system. The methods for performing I/O on files are described elsewhere. See the section "Streams", page 1.

13.2 Concepts

Files are categorized as character files and binary files. Character files consist of a certain number (the *byte count*) of characters in the Symbolics character set. Binary files consist of a number (their byte count) of binary data bytes, which are unsigned binary numbers up to sixteen bits in length.

A file has a *name*, a *type*, and a *version*. The name is a character string of any length. The type is a character string of up to fourteen characters in length, and the version a positive integer up to 16777215. Names and types can consist of upper and lower case characters. However, searching for file names is not sensitive to case. This means that if you create a file whose name is "MyFile", the file has that name and appears that way in directory listings, but if you ask for "myfile" or "MYFILE" or "MYfile", the file is found. The characters ">" and "Return" cannot appear in names and types.

The name is an arbitrary user-chosen string describing the file. The type is supposed to indicate what type of data the file contains; a type of "lisp" is the system convention for files containing Lisp source programs, "bin" for compiled Lisp programs, and so forth. The version number distinguishes successive generations of a file; to change a file, you normally read the latest version of the file into the Symbolics computer, modify it, and write out a new version with the next highest version number. The general scheme for naming files is covered elsewhere. See the section "Naming of Files", page 127.

Files reside in *directories*. The combination of name, type, and version of any file is unique in the directory in which it is contained. With the exception of a single directory (the *ROOT*), directories also reside in other directories. The directory in

which a file or directory resides is called its *parent* directory, and these files and directories are said to be *inferior* to their parent directory. Directories and files thus form a strict tree (hierarchy); the ROOT directory is the root of this tree. Directories have a type of "directory" and a version of 1. Thus, the name of a directory alone identifies it in its containing (*superior*) directory. It is not possible to "fool" LMFS into thinking a file is a directory by giving it a type of "directory" and a version of 1, however.

Links are the third (and last) kind of object that can live in a directory. A link contains the character-string representation of the pathname of something else in the same file system, called the *target* of the link. This pathname specifies a directory, a name and a type, and it can specify a version. A link is conceptually an indirect pointer to something else; when certain operations are done on a link, the operation really gets done to the target instead of the link itself.

It is possible to have "directory links". See the section "LMFS Links", page 210.

The specific syntactic conventions, restrictions, and other information about LMFS pathnames are described elsewhere.. See the section "LMFS Pathnames", page 174.

LMFS also stores and maintains *properties* of objects. For example, for each file it stores the creation date, the author, whether the file has been backed up, and so on. Users can also create their own properties; each file has a property list that lets you store arbitrary associated information with the file. See the section "LMFS Properties", page 204.

The File System Editor is an interactive program that lets you manipulate the file system (the local LMFS system or the system on any host). You can invoke the program by typing SELECT F. See the section "File System Editor", page 231.

Before you use the file system on a machine, you must log into that machine. If you are using the file system locally, it is desirable to log out of the machine before you cold boot it or otherwise abandon it. This is especially desirable if you have created files on the file system or expunged directories (see below) while using it. If you do not follow this recommendation, you will run out free disk records at the rate of about 30 to 50 records per cold boot (in which files were created), and the free record salvager will have to be run.

You do not have to take any special action to access the local file system on your Symbolics computer. If you use the host name of the machine, or the special string "local" as a host name, it is accessed automatically, as with any host.

13.3 Properties

Files, directories, and links have various *properties*. There are *system* properties, which are defined and maintained by the file system itself, and *user* properties, which are defined and maintained by programs and people that use the file system.

Every property has a *name*, which is a keyword symbol, and a *value*, which is a Lisp object. The names of all of the system properties are listed below. (These properties should not be confused with the *file attribute list*, the *-** line in the beginning of the file.) See the section "File Attribute Lists", page 109.

You can examine the values of properties by using either the [View Properties] command in the File System Editor, or View File Properties (M-X) in Zmacs. Users alter the values by using either the [Edit Properties] command in the File System Editor, or Change File Properties (M-X) in Zmacs. See the section "File System Editor", page 231. Programs access the values of properties by using the **fs:directory-list** and **fs:file-properties** functions and alter the values by using the **fs:change-file-properties** function. See the section "Accessing Directories", page 117.

Some system properties apply to files, directories, and links alike; for example, all these objects have an *author* and a *creation time*. Other system properties are not defined for all kinds of object; for example, only files have a *length in bytes*, only directories have an *auto-expunge interval*, and only links have a *link-to*. The table of system properties tells you which kinds of objects each property applies to. User properties can always apply to any object.

The values of some system properties are determined by the file system and cannot be set by the user; for example, you cannot set the *length in bytes* nor the *byte size*. The values of other properties can be changed arbitrarily; for example, you can set the *generation retention count* or the *don't delete* property whenever you want to. The properties of the latter set are called *changeable* properties. The reason for the distinction is that the properties in the first group reflect facts about the file, whereas those in the second group represent the current state of user-settable options regarding the file.

When the **fs:change-file-properties** function is called for a changeable system property, the property is changed. When it is called for a non-changeable system property, an error is signalled. When it is called for any property name that is not the name of one of the system properties (listed below), it assumes that it is the name of a user property, and the property is established or changed.

When the **fs:file-properties** function is called for a LMFS file, it returns a second value: a list of the names of all the properties of the file that are changeable. This function lists all the system properties and all the user properties for the object it is given.

The names of user properties must be symbols on the keyword package, and must not be the same as any of the system property names. The value associated with a user property must be a string. The combined length of the name of the property and its value must not exceed 512 characters. To remove a user property from a file, you set the value of the property to **nil**. **fs:file-properties** returns all the user properties of a file, but **fs:directory-list** does not return any of them. You can create new user properties with the [New Property] command in the File System

Editor; after they are created, you can edit them with [Edit Properties]. Programs create and change user properties by using **fs:change-file-properties**.

User properties involve a subtle trap. If you misspell or otherwise misconstrue the name of a system property, LMFS assumes that you have given the name of a user property, and set it. Thus, LMFS can never admit of, nor diagnose, an unrecognized, or invalid, property name.

Here is a list of all of the standard properties that LMFS maintains. The standard, generic interpretation and representations of the system standard properties among them can be found elsewhere: See the section "Functions for Accessing Directories", page 117. Refer to the table below for the rest.

- :length-in-bytes
- :byte-size
- :author
- :creation-date
- :modification-date
- :reference-date
- :deleted
- :not-backed-up
- :dont-delete
- :dont-reap
- :open-for-writing (LMFS-specific)
- :length-in-blocks
- :generation-retention-count
- :directory
- :auto-expunge-interval
- :date-last-expunged
- :default-generation-retention-count
- :default-link-transparencies (LMFS-specific)
- :link-to
- :link-transparencies (LMFS-specific)

The following among them are changeable, that is, users can set their values by means of **fs:change-file-properties**:

```

:generation-retention-count
:modification-date
:reference-date
:creation-date
:author
:deleted
:dont-reap
:dont-delete
:auto-expunge-interval
:default-generation-retention-count
:default-link-transparencies
:link-transparencies

```

The following is a list of all the properties supported by LMFS that are either specific to LMFS or require other special comment.

:byte-size (*Files*)

For a character file, **8**. For a binary file, the byte size of the file (the number of bits in each byte), a fixnum between **1** and **16**, inclusive. LMFS maintains both the byte size and the binary/character quality of a file natively. It is not permitted to open a binary file with a byte size other than that with which it was written. This property is not currently a changeable one.

:length-in-blocks (*Files, directories, links*)

A *LMFS record* is 1152 32-bit words. This is the basic allocation unit of the file system. The name of the generic system property is confusing in the case of LMFS, for a LMFS record is composed of multiple disk blocks. This property cannot be meaningful for directories.

:creation-date (*Files, directories, links*) (*Changeable*)

LMFS allows setting of creation date by user programs. Creation date, when not set by a user program, is also updated when a file is appended to.

:modification-date (*Files*)

The most recent time at which this file was modified, expressed in Universal Time. This is the same as the creation date unless the file has been opened for appending. Operations such as renaming and property changing update this property, but do not update creation date. The dumper, for instance, is driven off this property.

:author (*Files, directories, links*) (*Changeable*)

This property is user-settable in LMFS.

:dont-delete (*Files, directories, links*) (*Changeable*)

If **t**, does not allow this object to be deleted. The purpose of this attribute is to prevent the accidental deletion of important files. An error results if an attempt is made to delete this file.

:dont-reap (*Files, directories, links*) (*Changeable*)

This attribute, although maintained internally by LMFS, is not interpreted by LMFS. Dired directory maintenance tools use this property.

:default-generation-retention-count (*Directories*) (*Changeable*)

The default value for the **:generation-retention-count** property of new objects created in this directory. See the section "LMFS Deletion, Expunging, and Versions", page 209. **nil** or a nonnegative fixnum.

:default-link-transparencies (*Directories*) (*Changeable*)

The initial value for the **:link-transparencies** attribute of links created in this directory. See the section "LMFS Links", page 210. To set this property, use the [Link Transparencies] command in the File System Editor rather than [Edit Properties].

:link-transparencies (*Links*) (*Changeable*)

The transparencies of this link. See the section "LMFS Links", page 210. To set this property, use either the [Edit Link Transparencies] or [Edit Properties] commands in the File System Editor, or Change File Properties (m-x) in Zmacs.

:complete-dump-date (*Files, directories, links*)

The most recent time at which this object was dumped on a complete dump tape, expressed in Universal Time. See the section "LMFS Backup", page 212. A positive bignum. If this object has never been dumped on a complete dump tape, this property is not present. This property does not appear in directory listings.

:complete-dump-tape (*Files, directories, links*)

The tape reel ID of the complete dump tape on which this object was most recently dumped. A string. If this object has never been dumped on a complete dump tape, this property is not present. This property does not appear in directory listings.

:incremental-dump-date (*Files, directories, links*)

The most recent time at which this object was dumped on an incremental or consolidated dump tape, expressed in Universal Time. A positive bignum. If this object has never been dumped on an incremental dump tape, this property is not present. This property does not appear in directory listings.

:incremental-dump-tape (*Files, directories, links*)

The tape reel ID of the incremental or consolidated dump tape on which this object was most recently dumped. A string. If this object has never been dumped on an incremental dump tape, this property is not present. This property does not appear in directory listings.

13.4 Deletion, Expunging, and Versions

When an object (file, directory, or link) in LMFS is deleted, it does not really cease to exist. Instead, it is marked as "deleted" and continues to reside in the directory. If you change your mind about whether the file should be deleted, you can *undelete* the file, which will bring it back. The deleted objects in a directory actually go away when the directory is *expunged*; this can happen by explicit user command or by means of the auto-expunge feature (see below). When a directory is expunged, the objects in it really disappear, and cannot be brought back (except from backup tapes.) See the section "LMFS Backup", page 212.

When a file is deleted, any attempts to open the file will fail as if the file did not exist. It is possible to open a deleted file by supplying the **:deleted** keyword to **open**, but this is rare.

Users normally delete and undelete objects with the Zmacs commands Delete File (**m-x**) and Undelete File (**m-x**), or [Delete] and [Undelete] commands in the File System editor, or **D** and **U** in Dired. Directories can be expunged with Dired or the File System Editor, also, and the Expunge Directory (**m-x**) command in Zmacs. See the section "File System Editor", page 231.

Programs normally delete files using the **deletef** function. See the function **deletef**, page 104. Whether a file is deleted or not also appears as the **:deleted** property of the file, and programs can delete or undelete files by using **fs:change-file-properties** to set this property to **t** or **nil**.

Directories can optionally be automatically expunged. Every directory has an **:auto-expunge-interval** property, whose value is a time interval. If any file system operation is done on a directory and the time since the last expunging of the directory is greater than this interval, the directory is immediately expunged. The default value for this property is **nil**, meaning that the directory should never be automatically expunged.

The normal way of writing files in the Symbolics-Lisp environment is to create a new version of the file each time a file is written. When you edit with Zmacs, for example, every so often the Save File command is issued, and a new version is written out. After a while, you end up with many versions of the same file, which clutters your directory and uses up disk space. Zmacs has some convenient commands that make it easy to identify and automatically delete the old versions.

LMFS also has a feature that deletes the old versions automatically. A file property called the *generation retention count* says how many generations (that is, new versions) of a file should be kept around. Suppose the generation retention count of a file is three, and versions 12, 13, and 14 exist. If you write out a new version of the file, then version 12 will be deleted, and now versions 13, 14, and 15 will exist. Actually, version 12 is only deleted and not expunged, so you can still get it back by undeleting it. If the generation retention count is zero, that means that no automatic deletion should take place.

The above explanation is simplified. You might wonder what would have happened if versions 2, 3, and 14 existed, and what might have happened if the different versions of the file had different generation retention counts. To be more exact: each file has its own generation retention count. When you create a new version of a file and some other version of the file already exists (that is, another file in the directory with the same name and type but some other version), then the new file's generation retention count is set to the generation retention count of the highest existing version of the file. If there is no other version of the file, it is set from the *default generation retention count* of the directory. (When a new directory is created, its default generation retention count is zero (no automatic deletion).) So if you want to change the generation retention count of a file, you should change the count of the highest-numbered version; new versions will inherit the new value. When the new file is closed, if the generation retention count is not zero, all versions of the file with a number less than or equal to the version number of the new file minus the generation retention count will be deleted.

When a file version is being created, it is marked with the property **:open-for-writing**. This property is removed when the file is successfully closed. While the file has this property, it is invisible to normal directory operations and to attempts to open or list it. Directory list operations that specify **:deleted** can see the file. Files in this state have the "open for writing" property when you use View Properties in the file system editor, or View File Properties (m-X) in Zmacs. Files left in this state by crashes have to be removed manually by deleting and expunging. For example, suppose versions 3, 4, and 5 exist, but 5 is open in this state. An attempt to read **:newest** would get version 4; an attempt to write **:newest** would create version 6.

13.5 LMFS Links

A link is a file system object that points to some other file system object. The idea is that if there is a file called >George>Sample.lisp and you want it to appear in the >Fred directory, with the name New.lisp, you can create a link by that name to the file. Then if you open >Fred>New.lisp, you really get >George>Sample.lisp. The object to which a link points is called the *target* of the link, and can be found from the **:link-to** property of the link.

The above explanation is simplified. You might wonder what happens if, for example, you try to rename >Fred>New.lisp: is the link renamed, or the target? Each link has a property called its **:link-transparencies**. The value of this property is a list of keyword symbols. Each symbol specifies an operation to which the link is transparent. If the link is transparent to an operation, then when the operation is performed, it really happens to the target. If the link is not transparent to the operation, the operation happens to the link itself. Here is a list of the keywords, and the operations to which they refer:

:read	Opening the file for :input .
:write	Opening the file for appending, via :if-exists :append .
:create	Opening the file for :output
:rename	Renaming the file.
:delete	Deleting the file.

You can create new links with the [Create Link] command in the File System Editor, or Create Link (m-X) in Zmacs. See the section "File System Editor", page 231. Programs can use the **:create-link** message to pathnames. See the section "Pathname Messages: Naming of Files", page 165. When a new link is created, its transparencies are set from the **:default-link-transparencies** property of its superior directory. When a new directory is created, its **:default-link-transparencies** property is set to (**:read :write**).

The value of the **:link-transparencies** property of a link is a list of keywords describing the transparency attributes of which this link is possessed. The value of the **:default-link-transparencies** attribute of a directory is, similarly, a list of all those transparencies to be possessed by newly created links in this directory. When changing the value of either of these properties with **fs:change-file-properties**, the new value of the property is such a list of transparency keywords, chosen from the table above. Transparencies not present in the new value are turned off, and they are not preserved. There is no way to change an individual transparency.

When you create a new link with the [Create Link] command, you have to specify both the name and the type component of the new link; the version defaults to being the newest version, as of the time when you create the link. When you specify the target, you have to give a complete pathname with the name and the type; the version can be left unspecified. Targets of links can have unspecified versions; whenever such a link is used, the version is treated as **:newest**.

There is a subtle point regarding "create-through" links (links transparent to **:create**): what happens when you try to create a new version of **foo.lisp** when the highest version of **foo.lisp** is a create-through link? Is a new version of **foo** created, or is a new version created in the directory of the target of the existing link? Here is the rule. If a pathname is opened for **:output**, which means that it is being created, and the pathname has version **:newest** or a version number that is, in fact, the newest version, and the newest version is actually a create-through link, then the link is transparent and the operation happens in the target's directory. If the target pathname has a version, it is as if that exact pathname were opened for **:output**; if the target has no version, it is as if the target pathname with a version of **:newest** were opened.

A directory link is a link whose type is "directory", whose version is 1, and whose target is a real directory or another directory link. The maximum permitted length of such chains of directory links is 10. The system respects a directory link when

looking for a directory. By means of directory links, "indirect pointers," or multiple names for directories, can be established. Simply naming a link in this fashion is sufficient; no special action need be taken to declare a link to be a directory link. Transparencies are not interpreted in directory links.

13.6 LMFS Backup

A file system can be damaged or destroyed in any number of ways. To guard against such a disaster, it is wise to periodically *dump* the file system, that is, write out the contents of the files, their properties, and the directory information onto magnetic tapes. If the file system is destroyed, it can then be *reloaded* from the tapes. Individual files can also be *retrieved* from tapes, in case a single file is destroyed, or just accidentally deleted (and expunged). Dump tapes can also be used to save a copy of all the files on a system for archival storage.

In a *complete dump*, all of the files, directories, and links in the file system are written out to tapes. This, obviously, saves all the information needed to reload the file system. However, a complete can take a long time and use a lot of tape, especially if the file system is large. In order to make it practical and convenient to dump the file system at short intervals, a second kind of dump can be done, called an *incremental dump*.

In an incremental dump, only those files and links that have been created or modified since the last dump (of either kind) are dumped; things that have stayed the same are not dumped. (All directories are always dumped in an incremental dump.) Now, if the file system is destroyed, you reload it by first reloading from the most recent complete dump and then reloading each of the incremental dump tapes made since that complete dump, in same the order they were created. Therefore, you do not need to retain incremental dump tapes made before the most recent complete dump was done; you can reuse those tapes for future dumps.

Since all incremental dumps done since a complete dump must be reloaded in order to restore the file system, doing a complete dump regularly makes recovery time faster. Doing complete dumps also lets you reuse incremental dump tapes, as described above. The more incremental dump tapes you must load at recovery time, the longer it takes to recover, and thus the more chance there is that something will go wrong. Thus, it is advantageous to take complete dumps regularly.

A *consolidated dump* is like an incremental dump, in that it only dumps files that have been created or changed recently. When a consolidated dump is requested, a *consolidation date* is specified. A consolidation date is always in the past. The consolidated dump dumps those and only those files that have been created or changed since the consolidation date. A consolidated dump is the appropriate kind to take if some event destroys recent incremental dump tapes, or they are found to be unreadable.

See the section "Dumping, Reloading, and Retrieving" in *Installation and Site Operations*.

13.7 Multiple Partitions

The Lisp Machine File System (LMFS) allows the use of multiple partitions residing on one or more disk drives. The selection of partitions to be used by LMFS is determined by a database called the *file system partition table* (FSPT).

The FSPT is a FEP file named >fspt.fspt on the boot drive. The FSPT is optional: If not present, LMFS behaves as previously and uses lmfs.file on the FEP boot drive. The FSPT is a simple character database containing the actual pathnames (in the FEP file system) of the partitions to be used for file system access.

Each partition in the file system knows how many partitions make up the file system. Only the FSPT, which is used only at LMFS startup time, indicates the locations of these partitions. That is, the file system databases in the actual partitions do not contain drive and partition numbers or FEP pathnames. Thus, when LMFS is down, partitions can be moved around using Copy File (m-x); as long as the FSPT is edited to indicate their new locations, LMFS comes up (when required) using the moved partitions.

The FSPT is edited only to move partitions around or to add a partition. When you add partitions to the file system, the file system automatically rewrites the FSPT database to include the locations of new partitions.

Do not delete file partitions from your LMFS. Each LMFS partition contains pointers to all other file partitions in the LMFS. Deleting a file partition leaves the other partitions with pointers to a nonexistent file.

If you want to reduce the size of your LMFS, you must completely backup your LMFS, delete the entire existing LMFS and create a new one. The user files can then be restored into this new LMFS from the backup tapes.

Never delete a FEP file that is entered in the FSPT. If you do so, even if you remove the corresponding FSPT entry, the file system will not be able to initialize itself properly and will not be able to be brought up. This means that your file system will be irrevocably inconsistent and inaccessible.

13.7.1 Adding a Partition to LMFS

Partitions can be added to LMFS with [Local FS Maint] on the File System Editor menu. Select this item to get a menu of file-system maintenance operations. The [Initialize (R)] command yields a menu of initialization options, which offers [New File System] and [Auxiliary Partition] as a choice. [New File System] is similar to [Initialize (L)]; it initializes a partition to be the basis of a file system.

When you add a new partition or a partition on another disk, the disk should be free of errors and properly initialized and formatted, and the partition should exist.

To add another partition, use [Auxiliary Partition]. Enter the pathname of the FEP file to be used as the new partition. The default presented, which is correct for [New File System], is never correct for adding a partition. Then use [Do It]. The system then performs much verification and error checking, roughly as much as when initializing a new partition. It must not be interrupted while performing these actions. When finished, it adds the partition and edits the FSPT automatically.

14. FEP File System Overview

The Symbolics computer disk has a file system called the *FEP file system*. The entire disk is divided up into *FEP files* (that is, files of the FEP file system). FEP files have names syntactically similar to those of files in the Symbolics computer's own local file system. However, the FEP file system and the Lisp Machine File System (LMFS) are completely distinct.

FEP files (for example, fep0:>Boot.boot) can be accessed from Lisp. The following files are part of the FEP file system and should never be disturbed.

```
>disk-label.fep
>root-directory.dir
>free-pages.fep
>bad-blocks.fep
>sequence-number.fep
```

14.1 Microcode Loads

By convention, files of type MIC are microcode loads. These files contain images of the microcode and the contents of other internal high-speed memories that are initialized when the computer is booted.

14.2 World Loads

By convention, FEP files of type LOAD contain world loads, or *bands* (images of entire Lisp worlds).

14.3 Configuration Files

Configuration files contain FEP commands tailored for a particular Lisp Machine configuration. The commands are executed if you specify the file as argument to a Boot command when cold booting the machine. See the section "FEP Commands" in *User's Guide to Symbolics Computers*.

The configuration file >Boot.boot usually contains FEP commands to:

- Clear the internal state of the machine
- Load the microcode
- Load a world

- Set the Chaosnet address
- Start the machine

To change the selection of microcode and world loads that are booted by default, simply use Zmacs to edit the file >Boot.boot. Be careful to avoid typographical errors; otherwise, you might have to type in the commands manually in order to boot the machine. Also, be sure that the last command in the file is followed by RETURN.

14.4 How LMFS Uses the FEP File System

The FEP file `feh0:>lmfs.file` is where LMFS normally keeps its files. It holds the machine's local file system. The entire Symbolics computer local file system normally resides inside one big file of the FEP file system.

The file `feh0:>fspt.fspt` tells LMFS which FEP files to use for file space, if not `feh0:>lmfs.file`.

14.5 Virtual Memory

The FEP file `feh0:>page.page` holds the virtual memory of the Lisp system while Lisp is running. To increase the effective size of virtual memory, you can add additional paging files. See the section "Allocating Extra Paging Space", page 240.

14.6 FEP File Comment Properties

Comment properties supply additional information about the contents of FEP files. In the Dired mode of Zmacs, they are listed inside square brackets, where the reference or expunge date appears for other file systems. You can list the contents of the FEP file system by using the function **print-disk-label**. The Zmacs command Dired (M-X) of `feh:>*`, or the form (dired "feh:>*") invokes the directory editor on the FEP file system. An example of the Zmacs Dired command output follows:

```

48150 free, 322330/370480 used (87%)
FEPO:>BAD-BLOCKS.FEP.1  776      0(8)  9/14/83 11:46:56 [List of bad blocks] rll
FEPO:>boot.boot.15     1      121(8) 1/15/84 12:19:15 [] DEG
FEPO:>boot.boot.16     1      121(8) 1/29/84 13:06:43 [] DEG
FEPO:>boot.boot.17     1      121(8) 2/21/84 13:35:28 [] whit
FEPO:>boot.boot.18     1      124(8) 2/21/84 13:39:20 [] whit
FEPO:>DISK-LABEL.FEP.1  24      0(8)  9/14/83 11:46:55 [The disk label] rll
FEPO:>FREE-PAGES.FEP.1  41      0(8)  9/14/83 11:46:56 [Free pages map] rll
FEPO:>fspt.fspt.1      1      0(8)  9/14/83 11:46:58 [A filesystem partition table] rll
FEPO:>LMFS.file.1 50000  0(8)  1/05/84 23:20:13 [] ptaylor
FEPO:>Microcode1.MIC.1 103 117020(8) 6/30/83 08:19:16 [TMC5-MIC 219] Feinberg
FEPO:>PAGE.PAGE.1 150000  0(8)  9/14/83 11:46:58 [Main paging area] rll
FEPO:>Release-5-0.load.1 19109 22013568(8) 11/02/83 17:02:31 [Release 5 Beta Test] joseph
FEPO:>ROOT-DIRECTORY.DIR.1 2      0(8)  9/14/83 11:46:55 [His highness] rll
FEPO:>sequence-number.fep.1 1      0(8)  9/14/83 11:49:39 [] rll
FEPO:>System-243-463.load.1 22348 25733376(8) 1/02/84 11:46:14 [Exp 243.463] Zippy
FEPO:>system-243-481.load.1 20544 23666688(8) 1/11/84 22:48:56 [Exp 243.481, Full-GC]
FEPO:>system-243-516.load.1 20754 23908608(8) 1/24/84 23:23:41 [Exp 243.516, Full-GC] Zippy
FEPO:>system-243-559.load.1 19157 22068864(8) 2/19/84 19:32:45 [Exp 243.559, Full-GC] Moon
FEPO:>TMC5-MIC.MIC.247  103 118018(8) 10/03/83 20:25:07 [TMC5-MIC 247, Beta Test] joseph
FEPO:>TMC5-MIC.MIC.262  101 115233(8) 12/27/83 21:15:16 [TMC5-MIC 262] whit
FEPO:>TMC5-MIC.MIC.273  101 115810(8) 2/19/84 15:13:56 [TMC5-MIC 273] whit
FEPO:>World1.load.1 19138 20318976(8) 10/07/83 12:09:08 [Re1 4.5] LISPMNIL

```

14.7 Installing Microcode

Use **si:install-microcode** to retrieve any new microcode from the file system of the sys host.

si:install-microcode *from-file-or-version* &optional *to-file-or-version* *Function*
boot-file-to-update

Installs microcode from a system file into a file in the FEP file system.

from-file-or-version is a microcode version number (in decimal). The file resides in the logical directory `sys:l-ucode;`.

to-file-or-version rarely needs to be supplied. It defaults to a file on `FEPO:>` (the root directory of the boot disk) whose name is based on the microcode name and version. If supplied, *to-file-or-version* is either a pathname (string) of a file on `FEPO:>`, or an integer *n*, which stands for the file `TMC5-MIC.MIC.n` on `FEPO:>`.

The logical directory `sys: l-ucode;` includes multiple types of microcode for each version number. The correct microcode to install depends upon the particular

hardware configuration of your machine. When your machine is shipped, the default microcode filename is correct, but if your machine is upgraded (for example, an FPA board is installed) you might need to override the default used by **si:install-microcode** to get the correct type for your configuration. Below is an example of how you would get the microcode for a 3600 running 6.0, with no console upgrade but an FPA board installed:

```
(si:install-microcode "tmc5-fpa-mic.mic.319")
```

The correct microcode types for each system and hardware configuration are shown below. The names in this table omit the suffix *mic.n* that you must include to indicate the version of the required microcode. The version number must be followed by a period. Microcode version 319. is required for Release 6.0.

3600

	tmc5	ifu
No FPA	tmc5-mic.	ifu-mic.
FPA	tmc5-fpa-mic.	ifu-fpa-mic.

3670/3600 with console upgrade

	tmc5	ifu
No FPA	tmc5-io4-mic.	ifu-io4-mic.
FPA	tmc5-io4-fpa-mic.	ifu-io4-fpa-mic.

3640

	tmc5	ifu
No FPA	tmc5-io4-st506-mic.	ifu-io4-st506-mic.
FPA	tmc5-io4-st506-fpa-mic.	ifu-io4-st506-fpa-mic.

If you use the wrong microcode for your configuration, your machine will not boot, except in the case where your system has an FPA and you use a non-FPA microcode. In this case, the machine functions normally, but does not make use of the FPA at all.

boot-file-to-update specifies whether to update the boot file with the new microcode version number. It accepts one of these values (the default value is **nil**):

<i>Value</i>	<i>Action</i>
nil	Prompts for a boot file to update.
pathname	Does not prompt but uses pathname as the boot file to update.
:no-boot-file-update	Does not prompt or update.

Initially, the Symbolics personnel who install your system establish these microcode files for you.

14.8 Renaming FEP Files

FEP files can be renamed. For example, if you save a world containing MACSYMA, you might want to rename the world file to >macsyml.load or >macsyma1.load. Be sure to update your boot file if you intend this to be the default world.

14.9 Using a Spare World Load for Paging

You can reuse FEP file space by specifying a FEP file, usually a spare world load file, to be used as an extension of the file >page.page. To do so, use the FEP Add Paging-file command. Note that this action overwrites the previous contents of the specified file.

You should rename this file to, say, >extra.page, so that other users do not attempt to use the file space (to hold customized world loads, for instance).

You can also create new FEP files and use them for extra paging space. See the section "Allocating Extra Paging Space", page 240.

14.10 Adding a Spare World Load as LMFS File Space

Partitions can be added to LMFS with [Local FS Maint] on the File System Editor menu. Select this item to get a menu of file-system maintenance operations. The [Initialize (R)] command yields a menu of initialization options, which offers [New File System] and [Auxiliary Partition] as a choice. [New File System] is similar to [Initialize (L)]; it initializes a partition to be the basis of a file system.

When you add a new partition or a partition on another disk, the disk should be free of errors and properly initialized and formatted, and the partition should exist.

To add another partition, use [Auxiliary Partition]. Enter the pathname of the FEP file to be used as the new partition. The default presented, which is correct for [New File System], is never correct for adding a partition. Then use [Do It]. The system then performs much verification and error checking, roughly as much as when initializing a new partition. It must not be interrupted while performing these actions. When finished, it adds the partition and edits the FSPT automatically.

15. FEP File System

The *FEP file system* manages the disk space available on a disk pack, grouping sets of data into named structures called *FEP files*. All the available space on a disk pack is described by the FEP file system. A single FEP file system cannot extend beyond a single disk pack; each disk pack has its own separate FEP file system.

The FEP file system supports all of the generic file system operations. It also supports multiple file versions, soft deletion and expunging, and hierarchical directories.

Although "FEP" is an acronym for *front-end processor*, the FEP file system is managed by the main Lisp processor. It is called the FEP file system because the FEP can read files stored in the FEP file system. For example, the FEP uses the FEP file system for booting the machine and running diagnostics.

Disk streams access FEP files. A disk stream is an I/O stream that performs input and output operations on the disk. (For information about streams: See the section "I/O Streams", page 11.) When disk streams are opened with a **:direction** keyword of **:input** or **:output**, the disk stream reads or writes bytes (respectively), buffering the data internally as required. When the **:direction** is **:block**, the disk stream can both read and write the specified disk blocks. Block mode disk streams address blocks with a block number relative to the beginning of the file, starting at file block number zero. This *file block number* is internally translated into the corresponding disk address.

The FEP file system is also used by the system for allocating system overhead files, such as the paging file. See the section "FEP File Types", page 228. This section lists some of these files and what they are used for.

The ability of the FEP to access FEP files and the use of FEP files by the system imposes some constraints on the design of the FEP file system. The internal data structures of the file system must be simple enough to permit the FEP to be able to read them, and a small amount of concurrent access by both the FEP and Lisp must be tolerated. A FEP file's data blocks should have a high degree of locality on the disk to minimize access times. And the FEP file system must be very reliable, since the FEP needs to use the file system for running diagnostics and for booting the machine.

Note: Because of these constraints, the FEP file system is not intended to be a replacement for LMFS. (See the section "Lisp Machine File System", page 203.) Allocating new blocks for FEP files is slow, so that creating many files, especially many small files, might impair the performance of the FEP file system, and ultimately the virtual memory system if paging files or world load files become highly fragmented.

15.1 Naming of FEP Files

See the section "Lisp Machine File System", page 203. The FEP filename format is similar to the LMFS filename format, with the following exceptions:

<i>host</i>	The name of the FEP file system host. The format for a FEP host is <i>host FEPdisk-unit</i> , where the <i>host</i> field specifies which machine's FEP file system is being referred to, and <i>disk-unit</i> specifies the disk unit number on the machine. The <i>host</i> field defaults to the local machine if it and the terminating vertical bar () are omitted. If both the <i>host</i> and <i>disk-unit</i> fields are omitted, the FEP host defaults to the disk unit the world was booted from on the local machine. For example:						
	<table> <tr> <td>Merrimack FEP0</td> <td>The FEP file system on Merrimack's unit 0.</td> </tr> <tr> <td>FEP2</td> <td>The FEP file system on the local machine's unit 2.</td> </tr> <tr> <td>FEP</td> <td>The FEP file system the booted world load file resides on.</td> </tr> </table>	Merrimack FEP0	The FEP file system on Merrimack's unit 0.	FEP2	The FEP file system on the local machine's unit 2.	FEP	The FEP file system the booted world load file resides on.
Merrimack FEP0	The FEP file system on Merrimack's unit 0.						
FEP2	The FEP file system on the local machine's unit 2.						
FEP	The FEP file system the booted world load file resides on.						
<i>directory</i>	The name of the directory. The FEP file system supports hierarchical directories in the same format as in LMFS. Each directory name is limited to a maximum of 32 characters; there is no limit on the total length of a hierarchical directory specification.						
<i>name</i>	The name of the FEP file, which cannot exceed 32 characters.						
<i>type</i>	The type of the FEP file, which cannot exceed 4 characters.						
<i>version</i>	The version number of the FEP file, which must be a positive integer or the characters "newest".						

15.2 Accessing FEP Files

FEP files are accessed by open disk streams. A disk stream is opened by the **open** function. (See the section "Accessing Files", page 91. That section contains more details on accessing files.) If a FEP file system residing on a remote host is referred to, a *remote stream* is returned with limited operations as specified by the remote file protocol.

In addition to the normal **open** options, the following keywords are recognized:

:direction	Specifies the type of disk stream to open.
:input	Open a buffered input disk stream. A buffered input disk stream can only read bytes of data;

write operations are not permitted. The number of disk blocks to buffer can be specified by the **:number-of-disk-blocks** keyword.

:output Open a buffered output disk stream. A buffered output disk stream can only write bytes of data; read operations are not permitted. The number of disk blocks to buffer can be specified by the **:number-of-disk-blocks** keyword.

:block Open a bidirectional block disk stream. Block disk streams are used to read and write random disk blocks of data as requested. Block disk streams do not internally buffer disk blocks.

Block disk streams are not supported by the remote file protocol.

:probe Open a probe disk stream. A probe stream can read and modify a FEP file's properties, but cannot read or modify the file's data.

:if-exists This keyword specifies the action to be taken if the specified file already exists and the **:direction** is **:output** or **:block**. This keyword is ignored when the **:direction** keyword is **:input** or **:probe**. Only the following values are supported:

:error Signal an error. This is the default when the version component of the file name is not **:newest**.

:new-version Create a new version of the file. This is the default when the version component of the file name is **:newest**.

:overwrite Use the existing file.

:supersede Supersede the existing file by deleting and expunging it.

nil Return **nil** if the file already exists without creating a file or a stream.

:if-does-not-exist This keyword specifies the action to be taken if the specified file does not exist.

:error Signal an error. This is the default if the **:direction** is **:input** or **:probe**, or if the **:if-exists** argument is **:overwrite**.

:create Create a new file with the specified file name.

- This is the default if the **:direction** is **:output** or **:block**, and the **:if-exists** argument is not **:overwrite**.
- nil** Return **nil** if the file does not already exist without creating a file or a stream.
- :if-locked** This keyword specifies the action to be taken if the specified file is locked. This keyword is not supported by the remote file protocol.
- :error** Signal an error. This is the default.
- :share** Open the specified file even if it is already locked, incrementing the file's lock count. This mode permits multiple processes to simultaneously write to the same file. (See the section "FEP File Locks", page 228. That section contains more information on file locks.)
- :estimated-length** The value of this keyword is the minimum number of bytes to preallocate for the file. If the file's block length is not large enough to accommodate **:estimated-length** bytes of data, disk blocks are allocated and appended to the file. If the file's block length is greater than is required to satisfy **:estimated-length**, its size is not adjusted. This keyword is ignored if the **:direction** keyword is **:input** or **:probe**.
- :number-of-disk-blocks** The value of this keyword is the number of disk blocks to buffer internally if the **:direction** keyword is **:input** or **:output**. This keyword is ignored for other values of **:direction** or for files on remote hosts. The default **:number-of-disk-blocks** is two.

15.3 Operating on Disk Streams

All disk streams to a local FEP file system handle the following messages:

- :grow** &optional *n-blocks* &key **:map-area** **:zero-p** *Message*
 This message allocates *n-blocks* of free disk blocks and appends them to the FEP file. The value of *n-blocks* defaults to one. If **:zero-p** is true the new blocks are filled with zeros; otherwise, they are not modified. The return value of **:grow** is the file's data map (the format of the data map is described in **:create-data-map**'s description below). The value of **:map-area** is the area to allocate the data map in, which defaults to **default-cons-area**.

:allocate *n-blocks* &key **:map-area** **:zero-p** *Message*

This message ensures that the FEP file is at least *n-blocks* long, allocating additional free blocks as required. Returns the file's data map (the format of the data map is described in **:create-data-map**'s description below).

:map-area specifies the area to create the data map in, and defaults to **default-cons-area**. The newly allocated blocks are filled with zeros if **:zero-p** is true. **:zero-p** defaults to **nil**.

:file-access-path *Message*

This message returns the disk stream's file access path.

For example, you can find out what unit number a FEP file resides on as follows:

```
(send (send stream :file-access-path) :unit)
```

:map-block-no *block-number* *grow-p* *Message*

This message translates the relative file *block-number* into a disk address, and returns two values: the first value is the disk address, and the second is the total number of disk blocks starting with *block-number* that are in consecutive disk addresses. *grow-p* specifies if the file should be extended if *block-number* addresses a block that does not exist. When *grow-p* is true, free disk blocks are allocated and appended to the FEP file to extend it to include *block-number*. Otherwise, if *grow-p* is false, **nil** is returned if *block-number* addresses a block that does not exist.

:create-data-map &optional *area* *Message*

This message returns a copy of the FEP file's data map allocated in area *area*, which defaults to **default-cons-area**. A FEP file data map is a one-dimensional **art-q** array. Each entry in the file data map describes a number of contiguous disk blocks, and requires two array elements: the first element is the number of disk blocks described by the entry, and the second element is the disk address for the first block described by the entry. The array's fill-pointer contains the number of active elements in the data map times two.

:write-data-map *new-data-map* *disk-event* *Message*

This message replaces the file's data map with *new-data-map*. *disk-event* is the disk event to associate with the disk writes when the disk copy of the file's data map is updated. This message overwrites the file's contents and should be used with caution.

15.4 Input and Output Disk Streams

Input and output disk streams are buffered streams. In addition to the standard buffered stream messages, local input and output disk streams also support the messages described elsewhere: See the section "Operating on Disk Streams", page 224.

Input disk streams read bytes of data starting at the current byte position in the FEP file, updating the byte position as the data is read. Output disk streams write bytes of data in the same way.

The bytes of data are stored in buffers internal to the stream. The **:number-of-disk-blocks open** keyword controls how many disk blocks the internal buffers can hold. When the current pointer moves beyond a disk block boundary, the buffered disk block is written to the file for an output stream, or the next unbuffered block is read in from the file for an input stream. Output streams also write out all the buffered disk blocks when the stream is sent a **:close** message without an **:abort** option.

15.5 Block Disk Streams

Block disk streams can both read and write disk blocks at specified file block numbers. A file block number is the relative block offset into the file. The first block in the file is at file block number zero, the second is at file block number one, and so on.

Block disk streams do not buffer any blocks internally. They are not supported by the remote file protocol.

See the section "Operating on Disk Streams", page 224. In addition to the messages described in that section, block disk streams support the following messages:

:block-length *Message*

The **:block-length** message returns the length of the FEP file in disk blocks.

:block-in *block-number n-blocks disk-arrays &key :hang-p* *Message*
:disk-event

The **:block-in** message causes the disk to start reading data from the disk into the disk arrays in *disk-arrays* starting with the file block number *block-number* for *n-blocks*. *disk-arrays* can be a disk array or a list of disk arrays. The value of *n-blocks* is the number of disk blocks to read. When *n-blocks* is greater than one, each disk array is completely filled before using the next disk array in *disk-arrays*. Unused disk arrays or portions of disk arrays remain unmodified.

When the value of **:hang-p** is true, which it is by default, the **:block-in**

message waits for all the reads to complete before returning. If the value of **:hang-p** is false, **:block-in** returns immediately upon enqueueing the disk reads without waiting for completion. In this case, all *disk-arrays* and the *disk-event* must be wired before sending the **:block-in** message, and must remain wired until the disk reads complete.

If the **:disk-event** keyword is supplied, its value is the disk event to associate with the disk reads. Otherwise the **:block-in** message allocates a disk event for its duration. A **:disk-event** must be supplied when **:hang-p** is false.

:block-out *block-number n-blocks disk-arrays &key :hang-p* Message
:disk-event

The **:block-out** message causes the disk to start writing the data in the disk arrays in *disk-arrays* onto the disk starting with the file block number *block-number* for *n-blocks*. The arguments to the **:block-out** message are identical to those of the **:block-in** message.

15.6 FEP File Properties

In addition to having a name and containing data, FEP files also have properties. These properties store information about the file itself, such as when it was last written and whether it can be deleted or not. File properties are read by the **fs:file-properties** function, and modified by the **fs:change-file-properties** function. The **fs:directory-list** function also returns the file properties of several files at once. (See the section "Accessing Directories", page 117.)

The following file properties can be both read and modified:

- :creation-date** The universal time the file was last written to. Universal times are integers. (See the section "Dates and Times" in *Programming the User Interface*.)
- :author** The user-id of the last writer. The user-id must be a string.
- :length-in-bytes** The length of the file expressed as an integer.
- :deleted** When **t** the file is marked as being deleted. A deleted file can then be marked as being undeleted by changing this property to be **nil**. The disk space used by a deleted file is not actually reclaimed for reuse until the file is expunged.
- :dont-delete** When **t**, attempting to delete or overwrite the file signals an error, otherwise **nil** indicating the file can be deleted or written to.
- :comment** A comment to be displayed in brackets in the directory listing. The comment must be a string.

The following file properties are returned by the **:properties** message, but cannot be modified by **:change-properties**:

- :byte-size** The number of bits in a byte. The value of this property is always 8.
- :length-in-blocks** The block length of the file expressed as an integer.
- :directory** If *t*, the file is a directory, otherwise **nil** if the file is not a directory.

15.7 FEP File Locks

A FEP file is *locked* for the interval from when it is opened for reading or writing until it is closed. If the **:direction** keyword is **:input**, the file is *read-locked*; if the **:direction** keyword is **:output** or **:block**, the file is *write-locked*.

When the **:if-locked** keyword is **:error**, which is its default, a file that is read-locked can still be opened for reading but signals an error if opened for writing; a file that is write-locked cannot be opened for reading or writing. This permits multiple readers to access a file concurrently, while prohibiting writing to the file being read.

When the **:if-locked** keyword is **:share** in an open call for write, it succeeds in opening the file even if it is already read- or write-locked.

An expunge operation on a file that is either read- or write-locked does not expunge the file. If expunging a directory fails to expunge a file, the file must be closed and the directory expunged again.

15.8 FEP File Types

By convention, the following file types are used by the FEP file system for files used by the system.

- BOOT** The file contains FEP commands that can be read by FEP's Boot command. BOOT files are text files, and can be manipulated by the editor.
- LOAD** The file contains a world load image that is used to boot the system. For example, `>Release-6.load.NEWEST` contains the release 6 world load image.
- MIC** The file contains a microcode image. For example, `>TMC5-MIC.MIC.234` contains version 234 of the microcode for version 5 of the TMC.
- FSPT** The file contains a LMFS partition table. For example, `>FSPT.FSPT.NEWEST` is the default partition table used by LMFS.

- FILE** The file contains a LMFS partition. For example, >LMFS.FILE.NEWEST is the default LMFS file partition.
- PAGE** The file contains disk space that can be used by the virtual memory system. For example, >PAGE.PAGE.NEWEST is the default file used by the virtual memory system as storage for swapping pages in and out of main memory.
- FLOD** The file contains a FEP Load file. FEP Load files contain binary code the FEP can load and execute.
- FEP** The file contains binary information used by the FEP file system. These files should not be written to by user programs. Some examples of these files are:
- >FREE-PAGES.FEP
Describes which blocks on the disk are allocated to existing files.
 - >BAD-BLOCKS.FEP
Owns all the blocks that contain a media defect and should not be used.
 - >SEQUENCE-NUMBER.FEP
Contains the highest sequence number in use. The FEP file system uses sequence numbers internally to uniquely identify files to assist in rebuilding the file system in case of a catastrophic disk failure.
 - >DISK-LABEL.FEP
Contains the disk pack's physical disk label. The label is used to identify the pack and describe its characteristics.
- DIR** The file contains a FEP directory. For example, FEP0:>ROOT-DIRECTORY.DIR.NEWEST contains the top-level root directory. The directory file for FEP0:>DanG>Examples> would reside in FEP0:>DanG>Examples.DIR.1.

16. Fsedit

16.1 File System Editor

The File System Editor (FSEdit) is an interactive program that lets you examine and modify the contents of a file system. You can create directories and links, view and edit the properties of file system objects, delete objects, and expunge directories. The File System Editor is part of the File System Maintenance program, and it is the only part that most users ever use.

16.1.1 Entering the File System Editor

To get the File System Maintenance program, press SELECT F. At the top of the frame is a menu of commands. Three commands in this menu invoke the File System Editor:

- [Tree edit root]
- [Tree edit any]
- [Tree edit Homedir]

When click on one of these three commands, the big window in the frame displays a particular tree of a particular file system; that is, it displays a certain directory (the *base* directory) and some of the objects under that directory. If you use:

[Tree edit root] The base directory is the root directory of the local file system; this lets you get at any file in the entire file system.

[Tree edit any] You can specify the base directory by typing in its wildcard pathname; after you click on this command you are prompted for a wildcard pathname.

[Tree edit Homedir] The base directory is your home directory. [Tree edit Homedir] prompts for a host instead of using only the "logged-in" host (the one designated during login). If you just want to try out the File System Editor, use [Tree edit Homedir].

These commands put you in the File System Editor. You never have to get "out" of it; if you want to do something unrelated to the file system, just select the window you want to use, and if you want to do something else with the File System Maintenance program, you just click on the appropriate command in the command menu.

16.1.2 Using the File System Editor

When you use [Tree edit root], at the top of the main window is a line reading `>*. *.*`. This line represents the root directory, which usually contains only directories. Below the root directory line is a set of indented lines, one representing each object in the root directory.

Move the mouse over any one of these directory lines and notice that the mouse documentation line reflects three actions that you can take:

- (L) Open object: See the section "Opening and Closing a Directory", page 232.
- (M) Close containing object: See the section "Opening and Closing a Directory", page 232.
- (R) Menu of operations: See the section "Using Fsededit Commands", page 233.

16.1.3 Opening and Closing a Directory

Now, suppose you move the mouse over the line that represents a directory, for instance `>sys`, and click left. That line changes to read `>sys>*. *.*`, and several lines are inserted just underneath it, one for each object in the `>sys` directory. You have just *opened* the `>sys` directory.

When you open a directory, a line is inserted in the display for each object in the directory. For every directory, there is a line with the pathname of the directory and nothing else; these directories are all closed. For every file, there is a line with the name, type, and version of the file, and other information about the file. For every link, there is a line with the name, type, and version of the link, followed by `=>` and the pathname of the target of the link, and other information. See the section "How to Interpret Directory Listings", page 236.

Whenever you click left on a closed directory, FSEdit opens it and displays its contents. By clicking on successive directories inside other directories, you can move around in the file system and see what is there. The base directory is automatically opened as soon as you start using the File System Editor.

When you are finished with a directory, you can *close* the directory by clicking middle on any of the objects inside that directory. So, if you click middle on a file, that file and everything at its level disappears from the display.

Using these commands, you can get at any part of the file system underneath the base directory, and see everything that is there.

It is easy for the display to become longer than the size of the window when you move around in large directories; you can use the usual mouse scrolling commands to move the display up and down in the window. See the section "Scrolling with the Mouse" in *User's Guide to Symbolics Computers*.

16.1.4 Using Fsededit Commands

To do something to an object, click right on the object. This pops up a menu of commands, each of which specifies an action to take regarding the object. Some commands make sense for all three kinds of objects (directories, files and links); others are specific to certain kinds of objects. The menu that appears when you click right on an object offers only the options that you can apply to that type of object on your host type. For example, the menu does not display [Expunge] as an option for files or links, only for directories, and it does not display [Expunge] as an option if the directory in question resides on a host that does not support soft deletion.

The following is a list of all these commands with the kind of object(s) to which each command applies:

[Delete] (*Files, directories, links*)

Marks this object for deletion. This command pertains to systems that support soft deletion, for example, Symbolics computers. This command is only displayed for objects that are not already deleted. You should not delete directories that have anything in them.

[Delete (immediate)] (*Files, directories, links*)

Deletes this object. This command pertains to systems that do not support soft deletion, for example, UNIX. This command asks for confirmation and then immediately removes the deleted object from the display. You should not delete directories that have anything in them.

[Wildcard Delete] (*Directories*)

Does wildcard deletion. This command prompts you with a default for deleting everything for the line to which the menu applies. It merges what you enter with * defaults. It lists the files it intends to delete, asks for confirmation, deletes them, reporting any errors, and updates the display.

[Undelete] (*Files, directories, links*)

Undeletes this object. This command pertains to systems that support soft deletion, for example, Symbolics computers, and is displayed only for objects that are deleted (are marked with a D).

[Rename] (*Files, directories, links*)

Renames this object; prompts for a new name. If the object is not a directory, you can optionally type in a whole pathname specifying a new directory, and the file or link will be moved to the new directory as well as being given the new name.

[View Properties] (*Files, directories, links*)

Types out one line for each property of the object, giving the name and the value of the property. Properties are the qualities of the file that are maintained by the file system on which it resides, such as creation date and time, author, time of last access, and length. For files on a Lisp Machine file system, this means user-defined properties as well. It prompts for the name of a file and pops up a choose-variable-

values window, allowing you to alter various properties of the file. The exact properties that can be altered depend on the file system, but they might include:

- Generation (version) retention count
- Author
- Creation, modification, and reference dates
- Protection flags
- Other file-associated information

This information types out on top of the display, and prompts you to type any character when you are ready to proceed. After you type this character, the properties vanish and the FSEdit window is redisplayed. You can also use [Flush Typeout] in the command menu to make the typeout vanish; this is convenient since you do not have to move from the mouse to the keyboard.

[Edit Properties] (*Files, directories, links*)

Pops up a Choose Variable Values window that lets you change the value of any changeable system property or user property of the object.

[New Property] (*Files, directories, links*)

Creates a new user property for the object. You are first prompted for the name of the property, and then the value. The name is uppercased. To remove a property, give an empty string as the value.

[View] (*Files, links*)

Displays the file. The file is typed out on top of the display, and you are prompted to type any character when you are ready to proceed. The **:reference-date** of the file of the file is not changed. See the section "LMFS Properties", page 204. If the object is a link, it must be transparent to **:read** and its target must be a file; the target is printed.

[Create Inferior Directory] (*Directories*)

Creates a new directory inside this directory. You are prompted for the name (just type in the name, not the whole pathname).

[Create Link] (*Directories*)

Creates a new link inside this directory. You are first prompted for the name of the link, and then for the full pathname of the target of the link. See the section "LMFS Links", page 210.

[Expunge] (*Directories*)

Expunges the directory. See the section "LMFS Deletion, Expunging, and Versions", page 209.

[Open] (*Directories*)

Opens the directory. This is the same as clicking left on the directory name. This command is only displayed for closed directories.

[Selective open] (*Directories*)

Prompts for a wildcard name, for example, a file name containing "*" characters to indicate a wild-card component. The directory is opened and displays only those objects in the directory that match this pattern. Unspecified components default to "*". The normal **[Open]** command is like a **[Selective open]** of *.*.*, displaying all files. For example, if you do a **[Selective open]** of *.lisp, only files whose type is "lisp" are displayed. (In this example, the version was unspecified and defaulted to "*".) The line in the display that corresponds to the directory shows this wildcard name.

[Close] (*Directories*)

Closes the directory. This is the same as clicking middle on one of the directory's inferiors. This command is only displayed for open directories.

[Link transparencies] (*Links, directories*)

Lets you change the **:link-transparencies** of a link, or the **:default-link-transparencies** of a directory.

Each link has a property called its **:link-transparencies**. The value of this property is a list of keyword symbols. Each symbol specifies an operation to which the link is transparent. If the link is transparent to an operation, that means that if the operation is performed, it will really happen to the target. If the link is not transparent to the operation, then the operation will happen to the link itself. See the section "LMFS Links", page 210.

This command displays a menu showing all of the operations to which a link can or can not be transparent. Each operation to which the link actually is transparent is highlighted with reverse video. By clicking on the name of any operation, you can turn the highlighting on or off. When you are done changing the transparencies, use **[Do It]**, and the transparencies (or default transparencies, if this is a directory) are set. You use **[Abort]** to abort the operation.

[Decache] (*Directories*)

When a directory is opened, the File System Editor examines the directory, sees what is there, and remembers it. If another user changes the contents of the directory while you are in the middle of editing that directory, the File System Editor does not know that anything has changed, and so what it shows you does not really correspond to the state of the file system. Using **[Decache]** tells the File System Editor to forget what it thinks it knows about the contents of the directory, and makes it go back to the file system to see what is really in the directory now.

[Hardcopy] (*Files*)

Hardcopies the file. Clicking on this command causes the system hardcopy menu to pop up.

[Edit] (*Files*)

Invokes the Zmacs editor on the file.

[Load] (*Files*)

Loads the file into the Lisp world.

16.2 How to Interpret Directory Listings

The system displays the contents of directories of file systems in three contexts:

- The File System Editor
- The View Directory (m-X) and Dired (m-X) Zmacs commands
- The Show Directory command

Contents of directories are displayed in a standard format, regardless of the context and regardless of what kind of file system (Lisp Machine, TOPS-20, UNIX) the directory came from. Since this format is designed to express a great deal of information in a single line, it is rather abbreviated. Some of the ways it expresses things might not be clear without an explanation.

The basic format usually looks something like the following:

```
pal.lisp.65 7 25548(8) 03/12/85 12:42:41 (05/13/85) dlw
```

The following is an explanation of the items in this listing:

<i>item</i>	<i>explanation</i>
pal	file name
lisp	file type
65	file version number
7	length of the file in blocks
25548	length of the file in bytes
8	byte-size of the file
03/12/85	date file created
12:42:41	time file created
5/13/85	date file last referred to
dlw	author

Many other things can appear in such a line; some of these things are seen only on certain types of file systems. If the first character in the line is a D, the file has been deleted (this makes sense only on file systems that support undeletion, such as the Lisp Machine and TOPS-20 file systems). After the D, if any, and before the name of the file, is the name of the physical volume that the file is stored on (on ITS, this is the disk-pack number).

On a line that describes a link rather than a file, the length numbers are replaced by an arrow (=>), followed by the name of the target of the link.

On a line that describes a subdirectory rather than a file, the length-in-blocks number is shown (if provided by the file system), but the length-in-bytes is replaced by the string DIRECTORY.

Next, before the dates, the line might contain any of several punctuation characters indicating things about the file. Only some of the file systems understand these flags. Following is a list of the various characters and the flags they indicate:

<i>character</i>	<i>flag</i>
!	not backed up
@	do not delete
\$	do not reap

For lines indicating subdirectories, the reference date can be replaced with a date preceded by x=, the date this directory was last expunged. The dates are followed by the file author's name, which is followed by the name of the last user to read the file.

Only certain file systems support certain features. Many file systems do not keep track of the last reader's name and do not have something comparable to a "do not delete" flag. Therefore, any of the above fields might be omitted on certain file systems. However, the same general format is followed for all file systems and so you can interpret the meaning of a line in a directory listing, even for a file system that you are not familiar with.

17. Creating More Room on the Local Disk

There are two file systems available on the Symbolics computer: the Lisp Machine File System (LMFS) and the FEP File System (FEP FS). LMFS is a general-purpose, highly flexible file system, suitable for everyday use. Currently, only the Symbolics processor understands how to operate on LMFS files. The FEP FS is a simple, basic file system that both the Symbolics computer and front-end processors understand how to access. The FEP FS is used mainly to store world loads, microcode loads, paging files, boot files, and file system partitions that LMFS uses to store its structure and data. The FEP FS is not a good place for users to store their files; that is what LMFS is for.

Sometimes the Save World command or the **disk-save** or **si:receive-band** function might inform you that you have run out of FEP file system space. For systems with 167-Mbyte or more of storage, you should delete and expunge old, unneeded world loads, and then resume from the Save World/**disk-save** "out of room" error or retry the **si:receive-band** operation. You should not delete any world loads from a 140-Mbyte system. See the section "Instructions for Managing Disk Space on the 3640" in *Installation and Site Operations*.

It is wise to keep a large (25K-30K), noncritical world load on the Symbolics computer's disk, where it is available for the FEP Disk Restore command to use in case all world loads become nonfunctional.

Sometimes, writing a file out to a Lisp Machine File System (LMFS) produces an "out of room" error. This means that the present allocation of that particular LMFS is not large enough to accommodate your request for space. It might help to expunge directories with deleted files in them to remove unneeded versions of files, using the Zmacs command **Dired** (*m-X*).

If you still do not have enough space after you have deleted and expunged the unnecessary files, you might consider creating an auxiliary file partition. You should only consider doing so on systems that have at least 280 Mbytes of storage. There is no room in the FEP file system to allocate an auxiliary file partition for 140-Mbyte systems, and allocating an auxiliary file partition on a 167-Mbyte system might inhibit the creation of large world loads.

Even for 280-Mbyte systems, you are trading off world load space for file space when you create auxiliary partitions. Be sure to reserve enough FEP file system space for two large world loads (about 60K blocks total): the world you are currently running from and a spare world load for the FEP Disk Restore command to use.

For details on how to create auxiliary file partitions: See the section "LMFS Multiple Partitions", page 213. Once you have created an auxiliary file partition, you should never delete it. Deleting it would lose all the data contained in that partition and make the entire Lisp Machine File System unusable.

If you run out of room while writing a LMFS file and then create a new partition to increase the LMFS space, you cannot resume the file operation that failed. Instead, you must ABORT that operation and then retry the operation.

17.1 Allocating Extra Paging Space

Programs that use large amounts of virtual memory might require you to allocate additional paging space, to perform better or to perform at all. Only systems with at least 280 Mbytes of disk storage have enough room to permit additional paging files to be allocated without adversely affecting the maintenance of worlds on the machine. In order to add a paging file to your virtual memory set, you must first create a FEP file. Below is a function that creates a FEP file of a given length.

```
(defun create-fep-file (name length)
  (with-open-file (stream name :direction :block :if-exists :error)
    (send stream :allocate length)))
```

The code below creates an additional 20K-block paging file on unit 0, using the above function:

```
(create-fep-file "fep0:>page1.page" 20000.)
```

After creating the extra paging file, any boot files should be modified to use this new paging file. A typical boot file might look something like this:

```
clear machine
load microcode >tmc5-mic.mic.319
load world >Dist-6-0.load
set chaos 401
start
```

After creating the new paging file, edit your boot file to look something like this:

```
clear machine
load microcode >tmc5-mic.mic.319
load world >Dist-6-0.load
clear paging
add paging >Page.page
add paging >Page1.page
set chaos 401
start
```

It is safe to delete extra paging files, but only if they are not in active use. You cannot change a paging file that is being used without booting. Anytime you change the paging area you have set up, first boot without adding the paging file to be deleted. Be sure to cold boot by hand, and *do not* type the Add Paging command for the extra paging partition you intend to delete.

18. Putting Data in Compiled Code Files

A compiled code file can contain data rather than a compiled program. This can be useful to speed up loading of a data structure into the machine, as compared with reading in printed representations. Also, certain data structures, such as arrays, do not have a convenient printed representation as text, but can be saved in compiled code files.

In compiled programs, the constants are saved in the compiled code file in this way. The compiler optimizes by making constants that are `equal` become `eq` when the file is loaded. This does not happen when you make a data file yourself; identity of objects is preserved. Note that when a compiled code file is loaded, objects that were `eq` when the file was written are still `eq`; this does not normally happen with text files.

The following types of objects can be represented in compiled code files:

- Symbols
- Numbers of all kinds
- Lists
- Strings
- Arrays of all kinds
- Instances (for example, hash tables)
- Compiled function objects

When an instance is put (dumped) into a compiled code file, it is sent a `:fasd-form` message, which must return a Lisp form that, when evaluated, will recreate the equivalent of that instance. This is because instances are often part of a large data structure, and simply dumping all of the instance variables and making a new instance with those same values is unlikely to work. Instances remain `eq`; the `:fasd-form` message is sent only the first time a particular instance is encountered during writing of a compiled code file. If the instance does not accept the `:fasd-form` message, it cannot be dumped.

sys:dump-forms-to-file *filename forms &optional file-attribute-list* *Function*
sys:dump-forms-to-file writes data to a file in binary form. *forms-list* is a list of Lisp forms, each of which is dumped in sequence. It dumps the forms, not their results. The forms are evaluated when you load the file.

For example, suppose `a` is a variable bound to any Lisp object, such as a list or array. The following example creates a compiled code file that recreates the variable `a` with the same value:

```
(sys:dump-forms-to-file "f:>foo>aval"
 (list '(setq a ',a)))
```


For the purposes of understanding what this function does, you can consider that it is the same as the following:

```
(defun sys:dump-forms-to-file (file forms)
  (with-open-file (s file ':direction ':output)
    (dolist (f forms)
      (print f s))))
```

The real definition writes a binary file so it will load faster. It can also dump arrays, which you cannot write to a Lisp source file.

attribute-list supplies an optional attribute list for the resulting compiled code file. It has basically the same result when loading the binary file as the file attribute list does for **compiler:compile-file**. Its most important application is for controlling the package that the file is loaded into.

```
(sys:dump-forms-to-file "foo" forms-list '(:package "user"))
```

sys:dump-forms-to-file always puts a package attribute into the binary file it writes. If you do not specify the *attribute-list* argument, or if *attribute-list* does not contain a **:package** attribute, the function uses the **user** package. This is to ensure that package prefixes on symbols are always interpreted when they are loaded as they were intended when the file was dumped.

The *file-attribute-list* argument can be used to store useful information (such as "headers" for special data structures) in the file's attribute list. The information can then be retrieved from the attribute list with **fs:pathname-attribute-list**, without reading the rest of the file.

PART III.

The Serial I/O Facility

19. Introduction to Serial I/O

3600-family computers have a serial input/output facility, which uses the EIA RS-232 protocol to receive and transmit serial data. Many computer peripherals can communicate using the RS-232 protocol, and so can be connected to the 3600-family computer through this facility. This chapter explains the capabilities of the facility, gives a brief description of the hardware performing the serial I/O and how to interface to it, and describes the 3600-family software driving that hardware.

Before reading this chapter, you should be familiar with the basic concepts of serial data communication, including the RS-232 standard. You should also be familiar with Symbolics-Lisp, which is the systems programming language for the 3600-family computer. In particular, you should understand what *streams* are. See the section "Streams", page 1.

20. Hardware Description for Serial I/O

This section gives a brief description of the hardware that performs serial I/O on 3600-family computers. You do not have to understand everything in this section to use the serial I/O facility.

3600-family computers support three serial I/O ports. The external data communication signals appear on three RS-232 25-pin D-type connectors on the rear bulkhead (in the back of the processor).

The gender and labeling of these connectors varies with the processor model:

- The 3600 I/O bulkhead presents 3 female connectors labelled "EIA 1", "EIA 2", and "EIA 3". (The male connector labelled "EIA 4" is not a serial port at all, but the connection to an inboard Vadic VA3450 modem, if present. See the section "Physical Connection to the Dial Network" in *Networks*.)
- The 3670 I/O bulkhead presents 3 male connectors labelled "EIA 1", "EIA 2", and "EIA 3".
- The 3640 I/O bulkhead presents 3 male connectors labelled "SERIAL 1", "SERIAL 2", and "SERIAL 3".

These connectors are internally connected to multiprotocol USART-like LSI chips controlled by the 3600-family computer's Front End Processor (FEP). Using the USARTs, the FEP does all of the direct I/O to the connectors, and, using direct memory access, communicates with the 3600-family computer to transmit data from and receive data into the Symbolics-Lisp system.

The RS-232 protocol provides for communication between Data Circuit Terminating Equipment (DCEs, also known as "data sets"; for example, modems), and Data Terminal Equipment (DTEs, also known as "data terminals"; for example, computer terminals, computers, or most devices that use serial lines). The 3600-family computer plays the part of a DTE. This means that if you want to connect the serial line to a DCE, a simple cable can be used, but if you want to connect the serial line to a DTE, you must supply a *null modem*.

The correspondence between connector pins on the rear bulkhead and RS-232 signals is given in Table 1.

<i>Rear bulkhead connector pin</i>	<i>RS-232 signal</i>
2	<i>Transmitted Data [Output]</i>
3	<i>Received Data [Input]</i>
4	<i>RTS (Request To Send) [Output]</i>
5	<i>CTS (Clear To Send) [Input]</i>
6	<i>DSR (Data Set Ready) [Input]</i>
8	<i>DCD (Data Carrier Detect) [Input]</i>
20	<i>DTR (Data Terminal Ready) [Output]</i>
1	<i>Chassis Ground</i>
7	<i>Signal Ground</i>

Table 1. Assignment of RS-232 Signals to Pins

To build a cable that includes a null modem for asynchronous communications, follow the wiring instructions in Table 2.

<i>One side</i>	<i>Other side</i>	<i>RS-232 signal</i>
3	2	<i>Data Out (from data set to terminal)</i>
2	3	<i>Data In (from terminal to data set)</i>
5	4	<i>RTS (Request To Send)</i>
4	5	<i>CTS (Clear To Send)</i>
20	6	<i>DSR (Data Set Ready)</i>
20	8	<i>DCD (Data Carrier Detect)</i>
6	20	<i>DTR (Data Terminal Ready)</i>
8	20	<i>DTR (Data Terminal Ready)</i>
1	1	<i>Chassis Ground</i>
7	7	<i>Signal Ground</i>

Table 2. Assignment of RS-232 Signals to Pins in Asynchronous Null Modems

Note that this null modem is suitable only for asynchronous communications; a synchronous null modem is considerably more complex.

When using the 3600-family computer with a device that does not supply RS-232 modem control signals, it is necessary to supply Clear To Send and Data Carrier Detect inputs to the 3600-family computer, for example by jumpering pin 4 to pin 5,

and pins 6, 8, and 20 together. This should be done in the cable or in the device connector, *not* in the 3600-family computer's connector or inside the 3600-family computer.

21. The Serial I/O Stream

The function of the serial I/O facility is to receive and transmit data over a serial communications channel. The unit of communication is the *character*; each character is represented as a binary number. The facility has two independent parts: a *receiver*, which receives a sequence of characters from the external device, and a *transmitter*, which transmits a sequence of characters to the external device.

A Symbolics-Lisp program uses the facility through an I/O stream. The output operations, such as **:tyo**, send characters to the transmitter and from there to the external device; the input operations, such as **:tyi**, read characters from the receiver, which gets them from the external device. In addition to regular I/O operations, the serial I/O stream also supports special operations that examine and alter parameters of the serial I/O facility. To perform serial I/O, a program should first get the serial I/O stream by calling the function **si:make-serial-stream**, setting up the parameters of the serial I/O facility as it needs them; then it can use normal stream operations to read and write characters. When the program is done with the serial I/O stream, it should close it; programs that use the serial I/O stream should include an **unwind-protect** form whose cleanup handler closes the stream. The **with-open-stream** special form is a good way to do this when the entire lifetime of the stream is to be enclosed in the body of one Symbolics-Lisp form. Closing the stream frees up a buffer in main memory and disables interrupts.

The serial I/O stream is different from most streams in that the characters you send to it and get from it are probably *not* interpreted as being in the Symbolics 8-bit character set. Of course, the interpretation of the characters depends completely on the external device, but most devices that are likely to use serial communications use the standard ASCII character set. You can tell the stream whether or not to convert between ASCII characters and Symbolics characters.

The serial I/O stream is also different from some streams in being buffered on the output side. If you send characters to the serial stream using, for example, **:tyo** or **:string-out**, the characters are placed into a buffer for eventual transmission over the serial line. They are not actually transmitted until the buffer fills up, the serial stream is closed, or a **:force-output** operation is done on the stream. The **:force-output** option to **si:make-serial-stream** causes characters to be transmitted immediately; this makes the serial stream easier to use but degrades its performance.

The serial I/O stream has several parameters. Each parameter is denoted by a keyword symbol. These keywords are passed to the **si:make-serial-stream** function and to the **:get** and **:put** operations to specify which parameter the caller is interested in. (Some parameters make sense only when creating a stream, or affect the flavor of the stream; these parameters are not valid for **:get** and **:put**.) For descriptions of the parameters: See the section "Parameters for Serial I/O", page 253.

si:make-serial-stream &rest *options* *Function*

Initializes the serial I/O facility and returns the serial I/O stream.

options are alternating keyword symbols, naming parameters, and initial values for those parameters. They let you initialize parameters when you start using the serial I/O stream. You can change most of them later with the **:put** operation.

si:make-serial-stream, which accesses a serial line, causes the accessing process to wait if all ports are in use. The command **c-m-SUSPEND** allows you to invoke a restart handler to close a line that you believe has been left open by mistake.

For documentation of parameters for serial I/O: See the section "Parameters for Serial I/O", page 253.

The serial I/O stream supports all standard stream operations. Of the optional input operations, it supports **:listen** and **:clear-input**; the latter is relevant because input from the serial port is buffered. There is also a **:reset** operation, which resets the state of the hardware and the FEP. The **:tyi-no-hang** special-purpose operation is supported as well. The **:force-output** and **:finish** optional output operations are supported, since output is buffered.

The serial I/O stream also supports two nonstandard operations: **:get** and **:put**. These two operations respectively allow you to examine and alter various properties of the serial I/O facility. The names of these operations are intended to suggest the **get** and **putprop** functions in Symbolics-Lisp.

:get *parameter* of **si:serial-stream** *Method*

parameter should be one of the symbols that name parameters of the serial I/O facility. This message returns the value of that parameter. See the section "Parameters for Serial I/O", page 253.

:put *parameter value* of **si:serial-stream** *Method*

parameter should be one of the symbols that name parameters of the serial I/O facility. The value of that parameter is set to *value*. See the section "Parameters for Serial I/O", page 253.

22. Parameters for Serial I/O

This section lists all parameters of the serial I/O facility. For each parameter, it lists the keyword symbol, the meaning of the parameter, and the default value. A few parameters can be examined but not altered; they are so marked in their descriptions. Parameters whose functions are similar are grouped together.

Parameters from the following group are used *only* when the stream is being created, as arguments to **si:make-serial-stream**. You cannot use the **:put** operation with them, and you can use the **:get** operation only with **:unit**.

- :unit** This parameter says which of the serial ports to create a stream to. Its value should be 1, 2, or 3. The default is 2.
- :ascii-characters** If the value of this parameter is **t**, characters are translated from ASCII to the Symbolics internal character set on input, and to ASCII on output. The default is **nil**.
- :flavor** The value of this parameter is the flavor of stream to create. Normally, the value is computed automatically, based on the values of the **ascii-characters** and **force-output** parameters; this parameter is needed only if you want to use some special flavor that includes the serial stream flavors and other mixins.
- :force-output** If the value of this is **t**, a **:force-output** stream operation is done after every **:tyo** and every **:string-out**. If it is **nil** (the default), output is not transmitted until the output buffer fills up, a **:force-output** is done explicitly, or the stream is closed (and the close mode is not **:abort**). The nonforcing mode is usually more efficient, although efficiency depends on the application.

The following group of parameters controls the format of the transmitted characters. It is important to set the parameters to be compatible with the external device, or else proper communication is impossible. These parameters apply to both the transmitter and the receiver.

- :mode** The kind of communications protocol used over the port. The two possible values are **:asynchronous**, for asynchronous operation, and **:hdlc**, for the HDLC-like bit-stuffing protocol. See the section "Hdlc Serial I/O", page 263. The default is **:asynchronous**.
- :baud** The data transmission rate, in bits per second. This should be one of the following integers (in decimal): **300, 600, 1200, 1800, 2000, 2400, 3600, 4800, 7200, 9600, 19200**. The default is **1200**.
- :number-of-data-bits** The number of bits in each character. This should be one of the following fixnums: **5, 6, 7, or 8**. The default is **7**.

:parity The kind of parity bit that should be sent. If the value of this parameter is **nil**, no parity bit is sent. If it is **:even**, even parity is transmitted. If it is **:odd**, odd parity is transmitted. The default is **:even**. This parameter also controls what kind of parity checking is done on received characters.

:number-of-stop-bits The number of "stop" bits transmitted after each character. It should be one of the following numbers: **1**, **1.5**, or **2**. The default is **1**.

The following parameters control error checking in the receiver. After a character is read by an input stream operation, the stream checks for error conditions detected by the receiver when the character arrived. If any of the enabled error conditions occurred, the stream signals an error.

:check-parity-errors If the value of this parameter is **nil**, parity errors are ignored; if it is **t**, a parity error causes an error to be signaled when the character is read. The default is **nil**. A parity error occurs when the parity of the data bits disagrees with the value of the received parity bit. This never happens if parity checking is not being used, that is, if the **:parity** option is **nil**.

:input-error-character The value is a character to be substituted for any input character in which a parity error is detected. This is independent of the **:check-parity-errors** flag. If the value is **nil** (the default), the character is left alone.

:check-over-run-errors If the value of this parameter is **nil**, over-run errors are ignored; if it is **t**, then an over-run error causes an error to be signaled when the character is read. The default is **nil**. An over-run error occurs if input arrives faster than it can be read.

:check-framing-errors If the value of this parameter is **nil**, framing errors are ignored; if it is **t**, then a framing error causes an error to be signaled when the character is read. The default is **nil**. A framing error occurs when the "stop" bit (the bit after all the data bits, and after the parity bit if parity is being checked) is not **1**. This indicates a line error, a baud rate mismatch between the external device and the receiver, or the sending of a "break".

The following parameters deal with the "modem control" signals (signals other than Data In and Data Out) defined by the RS-232 protocol.

:carrier-detect If the value of this parameter is **t**, the external device is asserting

the DCD ("data carrier detect") signal; otherwise it is not. This parameter can be examined but not altered.

:clear-to-send If the value of this parameter is **t**, the external device is asserting the CTS ("clear to send") signal; otherwise it is not. This parameter can be examined but not altered.

:request-to-send If the value of this parameter is **t**, assert the RTS ("request to send") signal; otherwise do not. The default is **nil**.

:data-terminal-ready

If the value of this parameter is **t**, assert the DTR ("data terminal ready") signal; otherwise do not. The default is **nil**.

The following parameters control the use of the XON/XOFF protocol.

:xon-xoff-protocol

If this is **t**, output to the serial stream is flow-controlled using the ASCII XON/XOFF (Control-S/Control-Q) protocol. While the stream is transmitting characters, it checks the receiver to see if any characters have arrived. If an ASCII XOFF or Control-S character (octal 23, decimal 19) has arrived, transmission is stopped. Then the stream reads characters from the receiver until an ASCII XON or Control-Q character (octal 21, decimal 17) arrives, and then proceeds with the transmission.

This feature allows the external device to limit the rate at which characters are transmitted to it by the serial I/O facility. The default is **nil** (XON/XOFF feature not enabled).

Interpretation of incoming XON/XOFF signals is done at interrupt level in the FEP, and is therefore quite fast. After an XOFF is received, the 3600-family computer ceases transmission after two or three characters (buffered in the multiprotocol chip).

:output-xoff-character

The value is a character that is used to control flow of data from the Symbolics Lisp Machine to the external device. It is used to suspend the flow of data when the **:xon-xoff-protocol** parameter is set. The default is **#o023**.

:output-xon-character

The value is a character that is used to control flow of data from the Symbolics Lisp Machine to the external device. It is used to resume the flow of data when the **:xon-xoff-protocol** parameter is set. The default is **#o021**.

:generate-xon-xoff

If the value of this parameter is **t**, then the serial port generates XON and XOFF controls itself. This can be used to accept input at high speed from devices that understand the XON/XOFF protocol. The default is **nil**.

The XON and XOFF characters are transmitted directly by the FEP, so the response time is excellent. After the FEP transmits an XOFF, the device is required to cease transmission after no more than about 100 characters, so the device is not required to act very quickly.

:input-xoff-character

The value is a character that is used to control flow of data from the external device to the Symbolics Lisp Machine. It is sent by the Symbolics Lisp Machine to suspend the flow of data when the **:generate-xon-xoff** flag is set. The default is **#o023**.

:input-xon-character

The value is a character that is used to control flow of data from the external device to the Symbolics Lisp Machine. It is sent by the Symbolics Lisp Machine to resume the flow of data when the **:generate-xon-xoff** flag is set. The default is **#o021**.

23. Simple Example: Serial I/O

The following function assumes that the serial I/O facility is hooked to a computer terminal operating on a normal RS-232 asynchronous connection at 300 baud, with one stop bit and odd parity. It types the characters "Hello there." on the terminal. A null modem is used between the serial port and the terminal, because both ends are acting as DTEs.

```
(defun type-greeting-message ()
  (with-open-stream (stream (si:make-serial-stream
                             :baud 300.
                             :number-of-stop-bits 1
                             :parity :odd))
    (send stream :string-out "Hello there.")))
```


24. Troubleshooting: Serial I/O

If you have trouble making your device communicate with the 3600-family computer through the serial port, there are several things to try.

- Make sure that the baud rate, the number of data bits, the parity checking, and the number of stop bits are set the same way on the device as they are in your serial stream parameters.
- Make sure that there is a null modem between your device and the serial connectors, if the device is a DTE. Since most devices are DTEs, the null modem is probably necessary.
- Make sure that the device is connected to the proper serial port. The serial ports are labelled "EIA1" (or on 3640s, "SERIAL 1"), "EIA2" ("SERIAL 2"), and "EIA3" ("SERIAL 3"). You must use the port corresponding to the value of the **:unit** keyword to **si:make-serial-stream**. The default value is **2**, so if you do not specify anything, the "EIA2" ("SERIAL 2") connector is the appropriate one.
- Try using a different port. Remember both to plug your device into a different connector, and to change the program to specify a different value for the **:unit** keyword.

25. Notes on Serial I/O

The receiver is implemented using the 3600-family computer's general front end processor (FEP) "channel" facility. When a character arrives at the serial port, the FEP buffers it and transfers it to the 3600-family computer over a "channel". Therefore, it is not necessary for the program doing input from the stream to read in characters as quickly as they arrive from the external device. The **:clear-input** operation to the serial stream resets this buffer (including the buffers in Symbolics-Lisp, and the buffers in the FEP). The buffering capacity is about 500 characters. If the buffer is full and another character arrives, an over-run error occurs; if the **:check-over-run-errors** parameter is used, this is reflected by the signalling of an error.

A useful debugging technique is to create a serial stream with the desired parameters and set a variable (say, `s`) to it, and do:

```
(stream-copy-until-eof s standard-output)
```

This prints received characters on the screen until you type `c-ABORT`. This technique works only with the **:number-of-data-bits** parameter set to **7**, so that the Symbolics computer does not see the ASCII parity bit. Unless character set translation is enabled (via the **:ascii-characters** parameter), ASCII control characters, including carriage return and line feed, are displayed as special symbols, such as circle-cross or delta, because of the differences between the Symbolics character set and ASCII. See the section "The Character Set", page 5.

26. HDLC Serial I/O

The 3600 family supports synchronous serial I/O using HDLC-like bit-stuffing protocols. The CCITT-16 CRC polynomial is used.

This facility requires that the computer be running with FEP version 14 or later. Also, some older 3600s might require that a special adapter cable be connected to serial port 1. Baud rates of 9600 or lower are recommended.

An HDLC stream is a stream of flavor **si:serial-hdlc-stream**. Use the function **si:make-serial-stream** to make one of these streams. HDLC can be used only on serial port 1, so you must supply a **:unit** argument to **si:make-serial-stream** with a value of 1 (it defaults to 2). HDLC streams accept **:read-frame** and **:write-frame** messages.

si:serial-hdlc-stream

Flavor

An HDLC serial I/O stream. This flavor is built on **si:serial-binary-stream** and **si:serial-hdlc-mixin**.

:read-frame *string* &optional (*start 0*) *end* of **si:serial-hdlc-mixin** *Method*
Reads an HDLC frame into *string*. Returns the length actually read.

:write-frame *string* &optional (*start 0*) *end* of **si:serial-hdlc-mixin** *Method*
Writes *string* as an HDLC frame. This method never calls **process-wait** and can be used in a simple process. If insufficient buffers are available, it returns a form that evaluates to **t** when buffers become available.

27. Using the Terminal Program with Hosts Connected to the Serial Line

You can connect a 3600-family machine to another host via the serial line. Specifically, you can use the terminal program to communicate with another host when the 3600-family computer's serial line is connected to a terminal port on the other host.

The network system treats the set of hosts connected to the serial lines of a 3600-family computer as a special network, a *pseudonet*. Before you can use the terminal program to talk to another host over the serial line, you must use the **tv:edit-namespace-object** or the Edit Namespace Object command to create this network and assign an address on that network to the 3600-family computer. You might want to create or modify the remote host as well.

1. Create the network. Give it a **name** attribute associated with the 3600-family computer and a **type** attribute of **serial-pseudonet**.

In the following example, Merrimack is the name of the 3600-family computer:

```
NETWORK MERRIMACK-SERIAL
TYPE SERIAL-PSEUDONET
```

2. Add an entry to the **address** attribute of the 3600-family computer to specify that the 3600-family computer is connected to the new network. Each **address** entry is usually a pair of the form (*network address*). By convention, the 3600-family computer is assigned address 0 on a serial pseudonet. Following is an example of a new **address** entry for the 3600-family computer Merrimack:

```
ADDRESS MERRIMACK-SERIAL 0
```

3. If the line rate of the serial line is other than 9600 baud, supply a **peripheral** entry for the 3600-family computer giving the correct baud rate. The peripheral type is **serial-pseudonet**, and the **unit** attribute is the unit number of the serial line. Following is an example of a **peripheral** entry for the 3600-family computer:

```
PERIPHERAL SERIAL-PSEUDONET UNIT 2 BAUD 4800
```

4. If you want the terminal program to start out simulating one of the supported terminal types, add a **terminal-type** attribute to the peripheral. Currently supported terminal types are the VT100 and Ann Arbor Ambassador. For example, to make the terminal program simulate an Ambassador, add to the 3600-family computer a **peripheral** entry of this form (note that the entry must actually be on one line):

PERIPHERAL SERIAL-PSEUDONET UNIT 2 BAUD 9600
TERMINAL-TYPE Ambassador

You can now use the terminal program to connect to the remote host. At the "Connect to host:" prompt, you must supply an address of the form MERRIMACK-SERIAL|2. If you want to type a name or nickname of the remote host instead, add **address** and **service** entries for the remote host's namespace object. If the remote host does not exist in the network database, use the Edit Namespace Object command or the function **tv:edit-namespace-object** to create it.

For the **address** entry, specify the serial pseudonet and an address that corresponds to the unit number of the serial line to which the host is connected. The **service** entry is a triple of the form (*service medium protocol*). For the regular host login server, *service* is **login**, *medium* is **serial-pseudonet**, and *protocol* is **tty-login**. Following is an example of **address** and **service** entries for the remote host Blue connected to the 3600-family computer Merrimack:

```
HOST BLUE
SYSTEM-TYPE TENEX
ADDRESS MERRIMACK-SERIAL 2
SERVICE LOGIN SERIAL-PSEUDONET TTY-LOGIN
```

You can also use the serial line to connect to servers other than normal login on a remote host. You must add a **service** entry for the remote host to specify the kind of service, the **serial-pseudonet** medium, and the protocol that the remote host uses. You must also add an **address** entry on the serial pseudonet for the remote host. In the **address** entry, specify the address in the form *protocol=unit* instead of just *unit*. Following are examples of **address** and **service** entries for a file server using protocol **myftp** on remote host Blue:

```
HOST BLUE
SYSTEM-TYPE TENEX
ADDRESS MERRIMACK-SERIAL MYFTP=2
SERVICE FILE SERIAL-PSEUDONET MYFTP
```

For information on the Terminal program: See the section "Using the Network" in *User's Guide to Symbolics Computers*.

For information on network and host attributes: See the section "Namespace System Object Definitions" in *Networks*.

For information on services, media, and protocols: See the section "The Lisp Machine Generic Network System" in *Networks*.

PART IV.

Writing Programs That Use Magnetic Tape

28. The `tape:make-stream` Function

`tape:make-stream`

Function

`tape:make-stream` is used to create streams that read or write magnetic tape. It handles both cartridge and industry-compatible tape. With `tape:make-stream`, you can access tape on the local machine, or on any machine with a tape server.

`tape:make-stream` creates a stream. `with-open-stream` and other standard tools for managing streams should be used to ensure proper closing of a stream made with `tape:make-stream`.

Tape streams accept (for output) and return (as input) 8-bit characters. Normal stream messages can be used to tape streams. See the section "Streams", page 1. There are a few other messages: See the section "Messages to Tape Streams", page 273.

`tape:make-stream` takes a large number of optional keyword arguments:

- `:host`** The host on which the tape drive to be used is located. This can be a string or a host object. The keyword `:local` is also accepted for the local host. If this argument is not provided, `tape:make-stream` prompts for the name of the host.

The host must already be registered in the network database for supporting TAPE service.
- `:unit`** The identifier of the tape drive on the selected host that is to be used. Hosts having only one tape drive generally do not require this information. The value of this argument is generally a character string. "" or `nil` specifies "don't care", which is the usual value.
- `:reel`** The name of the tape reel to be mounted. This information is needed by tape servers that have operators, who need to know the name of a tape in order to mount it. It is also needed by servers who have tape access control systems. Currently (Release 5.0) no such servers are supported. "" or `nil`, the usual default, means "don't care".
- `:direction`** Specifies whether reading, writing, or intermixed reading and writing are to be performed. The valid values of this argument are thus `:input`, `:output`, and `:bidirectional`, respectively.

:input-stream-mode

This argument, which is only valid if the **:direction** argument is **:input** or **:bidirectional**, controls whether record boundaries, on input, are reflected to you. The default is **t**, meaning that they are *not*. It is not meaningful for cartridge tapes: record boundaries are never visible to the user of cartridge tape.

In *input stream mode* (a value of **t**), input bytes are transferred from the tape records to you until a file mark (tape mark, EOF) is encountered, at which time you see an end-of-file in your stream.

In *input record mode* (a value of **nil**), input bytes are transferred from the tape records to you until a record boundary, at which time you see an end-of-file in your stream. To progress beyond the record boundary, the message **:discard-current-record** must be sent to the stream.

:record-length

Controls the maximum length, in bytes, of tape records. This is ignored for cartridge tape. For reading, it must provide for the largest record to be read. Not all input records need be this long, although in some cases the server decides whether to allow records of other than this size. See also the keywords **:minimum-record-length** and **:minimum-record-length-granularity**. The default is 4096.

:density

Density of the tape in bits per inch. This is ignored for cartridge tape. The default is 1600 for servers that have the capability of multiple densities.

:pad-char

A number that is the single character with which to pad records when short records are padded. (This is ignored for cartridge tape.) The default pad character is 0. For compatibility with previous releases, supplying this argument and *not* supplying a value for either **:minimum-record-length** or **:minimum-record-length-granularity** implies a value of **:full** for **:minimum-record-length**.

:minimum-record-length

A number that is the minimum record length, in bytes, to which all output records will be padded. (This is ignored for cartridge tape.) This ability is present because many tape controllers cannot read records shorter than some minimum. Arguments to this keyword can be:

not supplied If this argument is not supplied, a value of 64 is assumed.

<i>integer</i>	Some number smaller than the value of the :record-length argument. Short records are padded with 0, or the value of the :pad-char argument, if that is supplied.
:full	All records are padded to their maximum length, namely, the value of the :record-length argument. Short records are padded with 0, or the value of the :pad-char argument, if that is supplied.
nil	The Lisp Machine does not enforce any minimum record length. The tape server and/or the tape hardware on that server might enforce some minimum of its own.

:minimum-record-length-granularity

An integer, or **nil**, establishing a *granularity*, or enforced integral divisor, for the length of all tape records written. If non-**nil**, all records written are padded (with 0, or the value of the **:pad-char** argument, if that is supplied) to be multiples of this number in length. This value is ignored for cartridge tape. It is also ignored if short records are not to be written, that is, **:minimum-record-length** is given as **:full** or the same as **:record-length**.

All Lisp Machine tape applications (LMFS and distribution dumpers and carry tape) enforce a granularity of 4.

:prompt

This is an optional string that is formatted into **tape:make-stream**'s prompt for a host name, if one is issued. It should describe the tape to be mounted in terms of the application program running. For instance, if this string is supplied as "**billing master**", **tape:make-stream** might prompt

Type name of tape host for billing master:

:no-bot-prompt

Normally, **tape:make-stream** notices if the tape is offline, or not at BOT (beginning-of-tape) when it is called. If the tape is offline, **tape:make-stream** queries you to wait for it to become ready. If the tape is not at BOT, **tape:make-stream** queries you about rewinding it. Supplying a non-**nil** value for **:no-bot-prompt** suppresses these checks, allowing you to handle these exigencies in any way you choose. The message **:bot-p** can be sent to a tape stream to determine if it is at BOT, and **:check-ready** to wait for a tape to become ready.

- :norewind** Normally, **tape:make-stream** rewinds the tape at the time the stream is closed. Supplying a non-**nil** value for **:norewind** suppresses this behavior.
- :lock-reason** Another optional string describing the application. This string is used in error messages sent to other users who try to access the tape drive you are using. For instance, if it is supplied as "**daily billing run**", another user might see a message like:
- Cannot mount tape:
Drive 0 in use by daily billing run.

29. Messages to Tape Streams

The following messages to tape streams are important. Tape streams, of course, also support standard stream messages appropriate to input or output streams. See the section "Streams", page 1.

These are the messages relevant to any kind of tape stream:

:close (&optional (*abort-p* *nil*))

Closes the stream. Normally, causes a rewind, and all the operations associated with **:rewind** (see the description of **:rewind**) to take place. The **:norewind** argument suppresses this rewind, although, for an output stream, buffered output is written, along with two EOFs. The tape is left positioned between the two EOFs, for industry-compatible tape, or after them, for cartridge tape.

:rewind

Rewinds the tape. For input streams, buffered input is discarded before the rewind. For output streams, buffered output is written out, possibly padded, according to the current padding parameters, and then two EOFs written, before the rewind. No read-ahead is performed. This message does *not* wait for the rewind to complete.

:await-rewind

Waits for a previously started rewind to complete.

:set-offline

A **:rewind** is done, and the tape is set offline, or unloaded, as befits the controller and drive. The setting of the tape offline does not wait for the rewind to complete.

:clear-error

If a tape error occurs, and is handled by you, you must send this message before attempting to continue using the stream. Otherwise, it remains in the error state, where it can only be closed.

:skip-file (&optional (*n* *1*))

Skips to, and past, a file mark (EOF). *n* is how many to skip, and can be negative, indicating backward motion. For input streams, all buffered input is discarded before the motion. For output streams, this operation is not valid unless the last thing written was an EOF, not a data record. Cartridge tape cannot skip backward. Forward motion is not allowed immediately after output.

:host-name

The name of the host on which the tape is mounted.

:bot-p

Returns **t** if the tape is at BOT (beginning of tape), and **nil** if not.

:check-ready Checks to make sure the tape drive is ready, and informs you, waiting interactively, if not.

These are the messages specifically relevant to tape input streams. Most of them are relevant only to input record mode, which is the mode requested by a value of **nil** for **:input-stream-mode**. See the description of the **:input-stream-mode** argument to the function **tape:make-stream**.

:clear-eof This clears the EOF state that results from reading an EOF mark. When an EOF is encountered, all character-reading operations encounter an end-of-file indication until **:clear-eof** is sent. This is needed in input stream mode as well as input record mode.

:discard-current-record This discards the remainder of the current record, when in input record mode, and allows reading the next record. This message must be issued to progress past a record boundary in input record mode, even if all of the bytes in the record have been read. This is meaningless for cartridge tape.

:record-status (&optional (*error-p* *t*)) This is only valid in input record mode, and meaningless for cartridge tape. This call is only valid at the beginning of a record, that is, if no bytes have been read from the current record. It describes, via its return value, the record that is about to be read by the user. Here are the possible values:

<i>an error object</i>	The next record cannot be read, due to error. An error object is returned. If <i>error-p</i> is <i>t</i> , which is the default, an error is signalled in this case, instead of an error object being returned.
<i>integer</i>	The length of a good record, in bytes.
:eof	The next record is not a record at all, but an EOF (a file mark).

These are the messages relevant to tape output streams:

:write-eof Writes an EOF (a file mark). If a record is being built, it is written out. Whether or not it is padded depends upon the values of the arguments **:minimum-record-length** and **:minimum-record-length-granularity**.

:force-output Writes out any record being buffered. Whether or not it is padded depends upon the values of the arguments **:minimum-record-length** and **:minimum-record-length-granularity**. This is the normal way

to end a record when record boundaries are significant, or short records are written. Otherwise, records are written when they are full.

:write-error-status (&optional *error-p*)

Verifies that all records have been written correctly. Tape streams often buffer many records ahead. **:write-error-status** waits for all buffered I/O to complete. If there was no error, **nil** is returned. If there was an error, an error object is returned describing the error. If *error-p* is non-**nil**, an error is signalled instead. If the error is end of tape, however, and **error-p** is **nil**, **:end-of-tape** is returned.

30. Tape Error Flavors

tape:tape-error *Flavor*

This set includes all tape errors. This flavor is built on **error**.

tape:mount-error *Flavor*

A set of errors signalled because a tape could not be mounted. This includes problems such as no ring and drive not ready. Normally, **tape:make-stream** handles these errors and manages mount retry. This flavor is built on **tape:tape-error**.

tape:tape-device-error *Flavor*

A hardware data error, such as a parity error, controller error, or interface error, occurred. This flavor has **tape:tape-error** as a **:required-flavor**.

tape:end-of-tape *Flavor*

The end of the tape was encountered. When this happens on writing, the tape usually has a few more feet left, in which the program is expected to finish up and write two end-of-file marks. Normally, closing the stream does this automatically. Whether or not this error is ever seen on input depends on the tape controller. Most systems do not see the end of tape on reading, and rely on the software that wrote the tape to have cleanly terminated its data, with EOFs.

This flavor is built on **tape:tape-device-error** and **tape:tape-error**.

Index

A

UNIX pathname
 VAX/VMS pathname
 Direct
 Direct
 Direct
 Effect of Character Set Translation on Direct
 Functions for
 Direct
 fs:
 Translate relative file block number into disk
 Defaults
 MacIsp
 [Tree edit
 Disable all character set translation in
 Disable special treatment of Rubout in
 Output streams to buffered
 Assignment of RS-232 Signals to Pins in
 File
 Reading file
 File

A

abbreviations 180
 abbreviations 180
 Absolute branch format directive 35
 Access Bidirectional File Streams 116
 access bidirectional streams 116
 Access File Streams 115
 Access File Streams 10
 Accessing Directories 117
 Accessing Directories 117
 Accessing FEP Files 222
 Accessing Files 91
 Access Output File Streams 116
:activation option 63
 Adding a LMFS Partition 219
 Adding a Partition to LMFS 213
 Adding a Spare World Load as LMFS File
 Space 219
add-logical-pathname-host 197
 address 225
 Advanced General-purpose Stream Operations 13
:advance-input-buffer message 21
:advance-output-buffer message 21
 alist 133, 147, 148
allfiles function 201
:allocate message 225
 Allocating Extra Paging Space 240
 any] File System Maintenance menu item 231
:any-tyi message 24
 Arrays in compiled code files 241
 ASCII characters 5
 ASCII character set 251, 253
:ascii-characters serial I/O parameter 253
 ASCII files 91
 ASCII files 91
:ascii option for **with-open-file** 91
 asynchronous devices 18, 114
 Asynchronous Null Modems 247
 attribute list 109
 attributes 111
 attributes as pathname properties 138
 Audible beep 19
:author FEP file property 227
:author file property 117
:await-rewind message to tape streams 273

A

B

! character file not

LMFS

Audible

B

BABYL pathname type abbreviation 180
 backed up flag 236
Backspace file attribute 109
:back-translated-pathname method of
fs:logical-pathname 199
 Back-translation 189
 Backup 212
 Bands 215
Base file attribute 109
:baud serial I/O parameter 253
 beep 19
beep function 88
:beep message 19

B

Bidirectional disk streams 222, 226
 Bidirectional file streams 115
 Bidirectional File Streams 116
 Direct Access bidirectional streams 116
 Direct access
 :binary-file-byte-size 163
 Binary format directive 31
 Binary mode 91
 :bin canonical type 200
 Control
 Hyper bit 5
 Meta bit 5
 Super bit 5
 Modifier bits 5
 Number of bits per byte 91
 :blip-handler option 62
 :block disk stream 222
 Block disk stream messages 226
 Block Disk Streams 226
 :block-in message 226
 :block-length message 226
 Block mode disk streams 221
 File block number 221, 226
 Translate relative file block number into disk address 225
 :block option for **with-open-file** 91
 :block-out message 227
 :block-size file property 117
 :bot-p message to tape streams 273
 Absolute branch format directive 35
 Relative branch format directive 35
 format: **breakline** macro 48
 :brief-help option 61
 Output streams to buffered asynchronous devices 18, 114
 Clear buffered input 18
 Special-purpose Stream Operations for Buffered Input Streams 20
 Clear buffered output 18
 Buffered output stream 21
 Special-purpose Stream Operations for Buffered Output Streams 21
 Buffered serial I/O 251
 Buffering capacity 261
 Number of bits per byte 91
 :byte-size file-opening option 91
 :byte-size file property 117
 :byte-size option for **copyf** 105
 :byte-size option for **with-open-file** 91
 Character-related byte specifiers 5

C

C

C

:bin canonical type 200
:lisp canonical type 200
:qbin canonical type 200
 :canonical-type message 151
 :canonical-type method of **fs:pathname** 166
 Correspondence of Canonical Types and Editor Modes 153
 Canonical types example 130
 Canonical Types in Pathnames 151
 Canonical types in Unix 4.2 Pathnames 178
 Canonical types in Unix Pathnames 176
 Canonical types in VMS Pathnames 180
 Buffering capacity 261
 Cardinal format directive 36
 Carriage return character 5
 Carriage return format directive 34
 :carrier-detect serial I/O parameter 253
 Native pathname component case 145
 Raw pathname component case 145
 Lower case format directive 34

- Case in Pathnames 145
- Center-dot character in TOPS-20 pathname display 183
- fs:**
 - change-file-properties** function 105, 119
 - :change-properties** message 114
 - :change-property** message 227
- FEP
 - * channel facility 261
 - character 153
- Double-arrow (↔) character 185
- Double-arrow character (↔) in logical pathnames for **nil** 186
- Equivalence sign character (≡) in logical pathnames for quoting 186
- \$ character do not reap file flag 236
- ! character file not backed up flag 236
- :character** file-opening option 91
- Keyboard
 - character format directive 33
- ITS
 - character handling 5
- Equivalence sign (≡) character in ITS pathnames 185
- Right horseshoe (⤴) character in ITS pathnames 185
- Enclose character in lozenge 34
- Echo character input 75
- Center-dot character in TOPS-20 pathname display 183
- Circle-x (⊗) character in TOPS-20 pathnames 183
- Character mode 91
- Character object 84
- Character output 46
- Character-related byte specifiers 5
- character representation 5
- Special function keys
 - ASCII characters 5
 - Graphics characters 5
- Octal representation of
 - characters 5
 - character set 251, 253
- Lisp Machine
 - Character Set 5
 - Character Set 5
 - Character set translation 10
- Disable all
 - character set translation in ASCII files 91
- Effect of
 - Character Set Translation on Direct Access File Streams 10
- Reading
 - characters from an input stream 15
 - :characters** message 13, 114
 - :characters** option for **copyf** 105
 - %%ch-char** variable 6
 - :check-framing-errors** serial I/O parameter 253
 - :check-over-run-errors** serial I/O parameter 253
 - :check-parity-errors** serial I/O parameter 253
 - :check-ready** message to tape streams 273
 - %%ch-font** variable 6
- USART
 - chip 247
- Data
 - Circle-x (⊗) character in TOPS-20 pathnames 183
 - Circuit Terminating Equipment 247
 - Clear buffered input 18
 - Clear buffered output 18
 - :clear-eof** message to input tape streams 273
 - :clear-error** message to tape streams 273
 - :clear-input** message 18, 251, 261
 - :clear-output** message 18
 - :clear-rest-of-line** message 15
 - :clear-to-send** serial I/O parameter 253
 - :clear-window** message 20
- fs:**
 - close-all-files** function 107
- Fsedit
 - [Close] Command 235
 - [Close] File System Editor menu item 233
 - close** function 103
 - :close** message 18
 - :close** message to tape streams 273
 - :close** operation 97
- Opening and
 - Closing a Directory 232
 - :command** option 64
- Using Fsedit
 - Commands 233

- FEP File
- Disk label
- fs:**
- Details of
- Directory
- Filename
- Rules for Successful
- File server
- LMFS
- FEP
- Using the Terminal Program with Hosts
- Concatenating
- Pathname Defaulting From the Current
- Modem
- File name
- Init File Naming
- Maclisp
- :byte-size** option for
- :characters** option for
- :create-directories** option for
- :report-stream** option for
- fs:**
- si:**
- si:**
- si:**
- Functions for Creating
- Fsedit
- Fsedit
- Functions for
- Maclisp
- Pathname Defaulting From the
- :commas** option for **format:onum** 45
- Comment Properties 216
- comments 216
- Compiled code objects in compiled code files 241
- :complete-help** option 60
- complete-pathname** function 120, 158
- Completion 120, 158
- Completion 121, 159
- Completion 122, 161
- completion 117
- Completion 123, 161
- computer 127
- Concatenating constant and variable output 44
- Concepts 203
- Conditional format directive 36
- Configuration Files 215
- configuration files 215
- Connected to the Serial Line 265
- constant and variable output 44
- Context 134
- Continuation format directive 35
- Control bit 5
- control parameters 253
- conventions 127
- Conventions 200
- Conversion 201
- Converting logical pathname to physical
pathname 186, 199
- Converting physical pathname to logical
pathname 199
- copyf** 105
- copyf** 105
- copyf** 105
- copyf** 105
- copyf** function 105
- Copy File (m-X) Zmacs command 105, 153
- Copying files 105
- copy-pathname-defaults** function 162
- coroutine-bidirectional-stream** flavor 28
- coroutine-input-stream** flavor 28
- coroutine-output-stream** flavor 28
- Coroutine Streams 25
- Coroutine Streams 27
- Correspondence of Canonical Types and Editor
Modes 153
- :create-data-map** message 225
- :create-directories** option for **copyf** 105
- [Create Inferior Directory] Command 234
- [Create Inferior Directory] File System Editor menu
item 233
- [Create Link] Command 234
- [Create Link] File System Editor menu item 233
- Create new logical host 197
- :create** symbol in **:if-does-not-exist** option for
open 222
- Creating a serial I/O stream 251
- Creating Coroutine Streams 27
- Creating More Room on the Local Disk 239
- :creation-date** FEP file property 227
- :creation-date** file property 117
- :creation-date** message 114
- Cross-host translation functions 145
- crunit** function 201
- Current Context 134
- cursorpos** function 88
- CWARNS pathname type abbreviation 180

D

D

D

- Putting FEP file
- X-
- dbg:** ***debug-io-override*** variable 22
- Fsedit
 - [Decache] Command 235
 - [Decache] File System Editor menu item 233
 - Decimal format directive 31
 - default** function 201
 - default handler 11
 - Defaulting 130, 131
 - Defaulting an Output File Pathname From an Input File 132
 - Defaulting From the Current Context 134
 - Defaulting of the Name 122, 160
 - Defaulting of the Type 122, 160
 - fs:** ***default-pathname-defaults*** example 130
 - fs:** ***default-pathname-defaults*** variable 148
 - fs:** **default-pathname** example 130, 132, 134
 - fs:** **default-pathname** function 162
 - defaults 156
 - Defaults 131
 - defaults 147, 158
 - defaults 131, 133
 - Defaults 130
 - defaults 129
 - Defaults alist 133, 147, 148
 - Defaults and Merging 147
 - defaults example 133
 - Default surface type 151
 - Default translation rule 189
 - Default type 147
 - Default version 147
 - format:** **defformat** special form 42
 - fs:** **define-canonical-type** special form 163
 - Defining a Translation Rule 191
 - definitions 51
 - defsystem** 196
 - [Delete] Command 233
 - Delete] Command 233
 - :deleted** option for **fs:complete-pathname** 120, 158
 - :deleted** option for **fs:directory-list** 117
 - :deleted** option for **with-open-file** 91
 - :deleted** option to **open** 97
 - deletef** function 104
 - Delete file 104
 - [Delete] File System Editor menu item 233
 - Delete] File System Editor menu item 233
 - [Delete (Immediate)] Command 233
 - [Delete (immediate)] File System Editor menu item 233
 - :delete** message 114
 - Deletion, Expunging, and Versions 209
 - :density** option to **tape:make-stream** 269
 - fs:** **describe-logical-host** function 196
 - fs:** **describe-pathname** function 165
 - Hardware
 - Description for Serial I/O 247
 - Details of Completion 121, 159
- Maclisp Stream
- Run-example with Simple
- Pathname
 - Special
 - Special
 - fs:** ***default-pathname-defaults*** example 130
 - fs:** ***default-pathname-defaults*** variable 148
 - fs:** **default-pathname** example 130, 132, 134
 - fs:** **default-pathname** function 162
- Manipulating
- More About
- Pathname
- Sticky pathname
- Tailoring Pathname
- Using pathname
- Pathname
- Sticky pathname
- format:** **defformat** special form 42
- fs:** **define-canonical-type** special form 163
- Formatted print of function
 - Fsedit
 - Fsedit [Wildcard]
- [Wildcard] Fsedit
- LMFS
 - fs:** **describe-logical-host** function 196
 - fs:** **describe-pathname** function 165

- Host
 - Determination in Pathnames 136
 - Device generic pathname component 149
 - :device** message 145
 - :device** method of **fs:pathname** 166
 - devices 18, 114
 - :device-wild-p** method of **fs:pathname** 173
 - Direct Access Bidirectional File Streams 116
 - Direct access bidirectional streams 116
 - Direct Access File Streams 115
 - Direct Access File Streams 10
 - Direct Access Output File Streams 116
 - :direction** file-opening option 91
 - :direction** message 13
 - :direction** option for **open** 222, 228
 - :direction** option for **with-open-file** 91
 - :direction** option to **tape:make-stream** 269
- Output streams to buffered asynchronous
- Effect of Character Set Translation on
 - Absolute branch format directive 35
 - Binary format directive 31
 - Cardinal format directive 36
 - Carriage return format directive 34
 - Conditional format directive 36
 - Continuation format directive 35
 - Date and time format directive 42
 - Decimal format directive 31
 - Escape format directive 41
 - Exponential format directive 32
 - Fixed-point format directive 32
 - Floating-point format directive 32
 - Free format directive 30
 - Fresh line format directive 34
 - Hexadecimal format directive 32
 - Indentation format directive 41
 - Indirection format directive 34
 - Iteration format directive 37
 - Justification format directive 39
 - Keyboard character format directive 33
 - Lower case format directive 34
 - Lowercase format directive 34
 - Lozenge format directive 34
 - Octal format directive 31
 - Ordinal format directive 36
 - Page separator format directive 34
 - Plural format directive 35
 - Radix directive 36
 - Relative branch format directive 35
 - Roman numeral format directive 36
 - Slashification format directive 30
 - Space format directive 36
 - Tabulation format directive 35
 - Tilde format directive 34
 - Time format directive 42
 - Time interval format directive 42
 - Time-of-day format directive 42
 - User-function format directive 41
 - ~> format** Directive 41
 - ~♦ format** Directive 34
 - ~\$ format** Directive 32
 - ~<Cr> format** Directive 35
 - ~< format** Directive 39
 - ~? format** Directive 34
 - ~a format** Directive 30
 - ~B format** Directive 31
 - ~c format** Directive 33
 - ~d format** Directive 31
 - ~e format** Directive 32
 - ~f format** Directive 32
 - ~* format** Directive 35
 - ~% Format** Directive 34

- ~& format Directive 34
- ~@* format Directive 35
- ~G format Directive 35
- ~o format Directive 31
- ~P format Directive 35
- ~Q format Directive 41
- ~R format Directive 36
- ~s format Directive 30
- ~[Str0~;Str1~;...~;Strn~] format Directive 36
- ~T format Directive 35
- ~@T format Directive 36
- ~X format Directive 32
- ~\Date\ format Directive 42
- ~\Datetime\ format Directive 42
- ~\Time-interval\ format Directive 42
- ~\Time\ format Directive 42
- ~-- format Directive 41
- ~{Str~} format Directive 37
- ~} format Directive 34
- ~-- format Directive 34
- ~(and ~) format Directives 34
- Accessing Directories 117
- Functions for Accessing Directories 117
 - File directory 117
 - Home directory 163
 - Opening and Closing a Directory 232
 - Root directory 215
 - Sys: I-ucode; logical directory 217
 - Fsedit [Create Inferior Directory] Command 234
 - Directory Completion 122, 161
 - directory display 236
 - :directory FEP file property 227
 - [Create Inferior Directory] File System Editor menu item 233
 - Maclisp **directory** function 201
 - Directory generic pathname component 149
 - directory last expunged indicator 236
 - fs:** **directory-link-opaque-dirlist** instead of fs:directory-list example 132
 - fs:directory-link-opaque-dirlist** instead of **fs:** **directory-list** example 132
 - fs:** **directory-list** function 117
 - => link indicator in directory listings 236
 - How to Interpret Directory Listings 236
 - Wildcard Directory Mapping 155
 - :directory message 145
 - :directory method of **fs:pathname** 166
 - directory name 222
 - :directory-pathname-as-file method of **fs:pathname** 144
 - Directory Pathnames and Directory Pathnames as Files 142
 - Directory Pathnames as Files 142
 - Directory Pathnames and Directory Pathnames as Files 142
 - :relative relative directory specifier syntax in logical pathnames 186
 - Directory TOPS-20 pathname component 183
 - Site directory translation 195
 - :directory-wild-p method of **fs:pathname** 173
 - Disable all character set translation in ASCII files 91
 - Disable special treatment of Rubout in ASCII files 91
 - :discard-current-record message to input tape streams 273
 - Creating More Room on the Local Disk 239
 - Translate relative file block number into disk address 225
 - Block **:block** disk label comments 216
 - :input** disk stream 222
 - :output** disk stream 222
 - :probe** disk stream 222
 - Block Disk stream messages 224
 - Block disk stream messages 226

Bidirectional Disk streams 221
 Block disk streams 222, 226
 Block mode Disk Streams 226
 Input and Output disk streams 221
 Operating on Disk Streams 226
 FEP File Properties: 3600 Disk Streams 224
 Center-dot character in TOPS-20 pathname Disk System User Interface 227
 File system editor directory display 183
 display 236
 Displaying Help Messages in the Input Editor 66
 Displaying Prompts in the Input Editor 65
 :**do-not-echo** option 63
 do not reap file flag 236
 :**dont-delete** FEP file property 227
 Double-arrow (↔) character 185
 Double-arrow character (↔) in logical pathnames for nil 186
 DTE 247
 sys: **dump-forms-to-file** function 241

E

E

E

Echo character input 75
 [Tree edit any] File System Maintenance menu item 231
 Fsedit [Edit] Command 235
 [Edit] File System Editor menu item 233
 [Tree edit Homedir] File System Maintenance menu item 231
 Editing terminal input 53
 :**editor-command** option 64
 :**editor** output format style 46
 Fsedit [Edit Properties] Command 234
 [Edit Properties] File System Editor menu item 233
 [Tree edit root] File System Maintenance menu item 231
 Format effectors 5
 EIA RS-232 protocol 243
 Enclose character in lozenge 34
 Encoding keyboard input 5
 End-of-file marker 12
 End-of-file on input streams 75
 tape: **end-of-tape** flavor 277
 Entering the File System Editor 231
 :**eof** message 18
 Data Circuit Terminating Equipment 247
 Data Terminal Equipment 247
 Equivalence sign (≡) character in ITS pathnames 185
 Equivalence sign character (≡) in logical pathnames for quoting 186
 :**error** file-opening option 91
 Tape Error Flavors 277
 Error message stream 22
 :**error** Option to **open** 93, 99
 error-output variable 22
 :**error** symbol in **:if-does-not-exist** option for **open** 222
 :**error** symbol in **:if-exists** option for **open** 222
 :**error** symbol in **:if-locked** option for **open** 222
 Escape format directive 41
 :**estimated-length** option for **open** 222
 examples 129
 Example: Serial I/O 257
 Examples of Use of the Input Editor 66
 explodec function 89
 explode function 89
 exploden function 89
 Exponential format directive 32

Fsedit [Expunge] Command 234
 X= date directory last expunged indicator 236
 [Expunge] File System Editor menu item 233
 LMFS Deletion, Expunging, and Versions 209
 Allocating Extra paging space 219
 Extra Paging Space 240

F

F

F

FEP channel
 The Serial I/O
 SELECT
 LOAD file type
 MIC file type
 Increase size of
 Accessing
 Naming of
 Renaming
 How LMFS Uses the
 Translating pathname type
 %%kbd
 fs:
 Getting a
 :byte-size
 :character
 :direction
 :error
 :preserve-dates
 fs:
 Fs:
 Fs:
 fs:
 fs:
 fs:
 ! character file not backed up
 \$ character do not reap file
 fs:pathname
 facility 261
 Facility 243
 :fasd-form message 241
 Fasdump 241
 F command 213
 (FEP) 215
 (FEP) 215
 Fep0:)lms.file file 216
 FEP channel facility 261
 FEP configuration files 215
 FEP directory name 222
 FEP FEP file type 228
 FEP file 224
 FEP File Comment Properties 216
 FEP file data map 225
 FEP File Locks 228
 FEP filename format 222
 FEP file properties 216
 FEP File Properties: 3600 Disk System User
 Interface 227
 FEP Files 222
 FEP Files 222
 FEP Files 219
 FEP File System 221
 FEP File System 216
 FEP File System Overview 215
 FEP File System Pathnames 176
 FEP File Types 228
 FEP host 222
 field 151
 fields 5
 :file-access-path message 225
 file-attribute-bindings function 111
 Filename 127
 Filename completion 117
 filename format 222
 Filename From the User 129
 file-opening option 91
 file-opening option 91
 file-opening option 91
 file-opening option 91
 file-opening option 91
 file-properties function 105, 119
 file-property-bindings 111
 file-read-property-list 111
 file-type-mode-allst variable 153
 find-file-with-type function 164
 :finish 113
 :finish message 18, 114
 :finish-timeout method of si:interactive-stream 71
 Fixed-point format directive 32
 :fixnum option for with-open-file 91
 flag 236
 flag 236
 Flashing the screen 19
 flatc function 89
 flatsize function 89
 flavor 153

- List valid init-options for
 - si:coroutine-bidirectional-stream** flavor 104
 - si:coroutine-input-stream** flavor 28
 - si:coroutine-output-stream** flavor 28
 - si:serial-hdlc-stream** flavor 28
 - tape:end-of-tape** flavor 263
 - tape:mount-error** flavor 277
 - tape:tape-device-error** flavor 277
 - tape:tape-error** flavor 277
 - Tape Error Flavors 277
- Print number in
 - :flavor** serial I/O parameter 253
 - floating point 46
 - Floating-point format directive 32
 - FLOD FEP file type 228
 - :fn1** method of **fs:its-pathname** 186
 - :fn2** method of **fs:its-pathname** 186
 - Font number 5
 - Fonts** file attribute 109
 - :force-output** message 18
 - :force-output** messages 113
 - :force-output** message to output tape streams 273
 - :force-output** serial I/O parameter 251, 253
 - :force-rescan** method of **si:interactive-stream** 71
- format:deformat** special form 42
- fs:define-canonical-type** special form 163
- grindef** special form 51
- with-input-editing-options-if** special form 56
- with-input-editing-options** special form 55
- with-input-editing** special form 57
- with-open-file** special form 91
- with-open-stream** special form 96, 251
- FEP filename format 222
- format:breakline** macro 48
- format:deformat** special form 42
- format directive 35
- format directive 31
- format directive 36
- format directive 34
- format directive 36
- format directive 35
- format directive 42
- format directive 31
- format directive 41
- format directive 32
- format directive 32
- format directive 32
- format directive 30
- format directive 34
- format directive 32
- format directive 32
- format directive 41
- format directive 34
- format directive 37
- format directive 39
- format directive 33
- format directive 34
- format directive 34
- format directive 34
- format directive 31
- format directive 36
- format directive 34
- format directive 35
- format directive 35
- format directive 36
- format directive 30
- format directive 36
- format directive 35
- format directive 34
- format directive 34
- format directive 35
- format directive 36
- format directive 36
- format directive 36
- format directive 34
- format directive 34
- format directive 42

Time interval format directive 42
 Time-of-day format directive 42
 User-function format directive 41
 `@* format Directive 35
 `& format Directive 34
 `% Format Directive 34
 `* format Directive 35
 `→ format Directive 41
 `◇ format Directive 34
 `\$ format Directive 32
 `< format Directive 39
 `<Cr> format Directive 35
 `? format Directive 34
 `a format Directive 30
 `B format Directive 31
 `c format Directive 33
 `d format Directive 31
 `e format Directive 32
 `f format Directive 32
 `G format Directive 35
 `o format Directive 31
 `P format Directive 35
 `Q format Directive 41
 `R format Directive 36
 `s format Directive 30
 `[Str0";Str1";...";Strn`] format Directive 36
 `T format Directive 35
 `@T format Directive 36
 `X format Directive 32
 `\Date\ format Directive 42
 `\Datetime\ format Directive 42
 `\Time-interval\ format Directive 42
 `\Time\ format Directive 42
 `-- format Directive 41
 `{Str}` format Directive 37
 `-] format Directive 34
 `-- format Directive 34
 ` (and `) format Directives 34
 Format effectors 5
 format function 29, 87
 format:ochar function 46
 format:ofloat function 46
 format:onum 45
 format:onum 45
 format:onum function 45
 format:oprint function 46
 format:ostrng function 46
 format:outfmt macro 45
 format:output macro 45
 format:pad macro 47
 format:plural function 47
 format:print-list function 44
 :editor output format style 46
 :read output format style 46
 :sall output format style 46
 :terpri option for format:tab 47
 :unit option for format:tab 47
 format:tab function 47
 Formatted Output 29
 Formatted print of function definitions 51
 Formatting functions 44
 formatting functions 44
 formatting functions 44
 formatting functions 44
 Formatting Lisp Code 51
 Free format directive 30
 Fresh line format directive 34
 :fresh-line message 15

- Defaulting an Output File Pathname
- Pathname Defaulting
- Getting a Filename
- From an Input File 132
- From the Current Context 134
- From the User 129
- Front-end Processor 221
- fs:add-logical-pathname-host** 197
- fs:change-file-properties** function 105, 119
- fs:close-all-files** function 107
- fs:complete-pathname** function 120, 158
- fs:copy-pathname-defaults** function 162
- fs:*default-pathname-defaults*** example 130
- fs:*default-pathname-defaults*** variable 148
- fs:default-pathname** example 130, 132, 134
- fs:default-pathname** function 162
- fs:define-canonical-type** special form 163
- fs:describe-logical-host** function 196
- fs:describe-pathname** function 165
- fs:directory-link-opaque-dirlist** instead of **fs:directory-list** example 132
- fs:directory-link-opaque-dirlist** instead of **fs:directory-list** example 132
- fs:directory-list** function 117
- Fsedit 231
- Fsedit [Close] Command 235
- Fsedit Commands 233
- Fsedit [Create Inferior Directory] Command 234
- Fsedit [Create Link] Command 234
- Fsedit [Decache] Command 235
- Fsedit [Delete] Command 233
- Fsedit [Delete (Immediate)] Command 233
- Fsedit [Edit] Command 235
- Fsedit [Edit Properties] Command 234
- Fsedit [Expunge] Command 234
- Fsedit [Hardcopy] Command 235
- Fsedit [Link Transparencies] Command 235
- Fsedit [Load] Command 236
- Fsedit [New Property] Command 234
- Fsedit [Open] Command 234
- Fsedit [Rename] Command 233
- Fsedit [Selective Open] Command 235
- Fsedit [Undelete] Command 233
- Fsedit [View] Command 234
- Fsedit [View Properties] Command 233
- Fsedit [Wildcard Delete] Command 233
- fs:file-attribute-bindings** function 111
- fs:file-properties** function 105, 119
- Fs:file-property-bindings 111
- Fs:file-read-property-list 111
- fs:*file-type-mode-alist*** variable 153
- fs:find-file-with-type** function 164
- fs:init-file-pathname** function 163
- fs:its-pathname** 186
- fs:its-pathname** 186
- fs:*its-uninteresting-types*** variable 186
- fs:*known-types*** variable 149
- fs:load-pathname-defaults** variable 149
- fs:logical-pathname** 199
- fs:logical-pathname** 199
- fs:make-logical-pathname-host** function 197
- fs:make-pathname-defaults** 131, 133
- fs:make-pathname-defaults** function 162
- fs:make-pathname** function 145, 162
- fs:merge-pathnames-and-set-defaults** function 158
- fs:merge-pathnames** example 133, 134
- fs:merge-pathnames** function 157
- fs:multiple-file-plists** function 119
- fs:parse-pathname** function 156
- fs:patch-file** Translation Rule 196
- fs:pathname** 166
- fs:pathname** 166
- :canonical-type** method of
- :device** method of
- :fn1** method of
- :fn2** method of
- :back-translated-pathname** method of
- :translated-pathname** method of

- :device-wild-p** method of **fs:pathname** 173
- :directory** method of **fs:pathname** 166
- :directory-pathname-as-file** method of **fs:pathname** 144
- :directory-wild-p** method of **fs:pathname** 173
- :generic-pathname** method of **fs:pathname** 171
- :getl** method of **fs:pathname** 172
- :get** method of **fs:pathname** 171
- :host** method of **fs:pathname** 165
- :name** method of **fs:pathname** 166
- :name-wild-p** method of **fs:pathname** 173
- :new-canonical-type** method of **fs:pathname** 169
- :new-default-pathname** method of **fs:pathname** 170
- :new-device** method of **fs:pathname** 167
- :new-directory** method of **fs:pathname** 167
- :new-name** method of **fs:pathname** 167
- :new-pathname** method of **fs:pathname** 170
- :new-raw-device** method of **fs:pathname** 168
- :new-raw-directory** method of **fs:pathname** 168
- :new-raw-name** method of **fs:pathname** 168
- :new-raw-type** method of **fs:pathname** 169
- :new-type** method of **fs:pathname** 168
- :new-version** method of **fs:pathname** 168
- :parse-truename** method of **fs:pathname** 171
- :pathname-as-directory** method of **fs:pathname** 144
- :pathname-match** method of **fs:pathname** 172
- :plist** method of **fs:pathname** 172
- :putprop** method of **fs:pathname** 172
- :raw-device** method of **fs:pathname** 166
- :raw-directory** method of **fs:pathname** 166
- :raw-name** method of **fs:pathname** 166
- :raw-type** method of **fs:pathname** 166
- :remprop** method of **fs:pathname** 172
- :string-for-directory** method of **fs:pathname** 171
- :string-for-dired** method of **fs:pathname** 171
- :string-for-editor** method of **fs:pathname** 171
- :string-for-host** method of **fs:pathname** 171
- :string-for-printing** method of **fs:pathname** 171
- :string-for-wholine** method of **fs:pathname** 171
- :system-type** method of **fs:pathname** 168
- :translate-wild-pathname** message to **fs:pathname** 192
- :translate-wild-pathname** method of **fs:pathname** 173
- :translate-wild-pathname-reversible** message to **fs:pathname** 192
- :type** method of **fs:pathname** 166
- :types-for-canonical-type** method of **fs:pathname** 169
- :type-wild-p** method of **fs:pathname** 173
- :version** method of **fs:pathname** 166
- :version-wild-p** method of **fs:pathname** 173
- :wild-p** method of **fs:pathname** 172
- fs:pathname-attribute-list** function 111
- fs:pathname** example 132
- fs:pathname** flavor 153
- fs:pathname-plist** function 165
- FSPT 213
- FSPT FEP file type 228
- fs:read-attribute-list** function 109, 111
- fs:*remember-passwords*** variable 107
- fs:set-default-pathname** 133
- fs:set-default-pathname** function 163
- fs:set-logical-pathname-host** 188, 196
- fs:set-logical-pathname-host** function 198
- fs:user-homedir** function 163
- :full-rubout** option 59
- Cross-host translation functions 145
- Formatting functions 44
- Input Functions 75
- :minpad** option for formatting functions 44
- Output Functions 87
- :padchar** option for formatting functions 44

Pathname Functions 156
 Slashification-related output functions 89
:tab-period option for formatting functions 44
 Functions for Accessing Directories 117
 Functions for Creating Coroutine Streams 27

G

G

G

Advanced
 Basic
 Device
 Directory
 Host
 Type
 Version
 General-purpose Stream Operations 11
 General-purpose Stream Operations 13
 General-purpose Stream Operations 11
:generate-xon-xoff serial I/O parameter 253
 generic pathname component 149
 generic pathname component 149
 generic pathname component 149
 generic pathname component 149
 generic pathname component 149
 generic pathname component 149
:generic-pathname message to pathnames 149
:generic-pathname method of **fs:pathname** 171
 Generic Pathnames 149
:getl method of **fs:pathname** 172
:get message 227
:get method of **fs:pathname** 171
:get method of **sl:serial-stream** 252
:get-output-buffer message 21
 Getting a Filename From the User 129
 Global translation rules 189
 Granularity 269
 Graphics characters 5
grindef special form 51
grind-top-level function 51
:grow message 224

H

H

H

Stream default handler 11
 ITS character handling 5
 Fsedit [Hardcopy] Command 235
 [Hardcopy] File System Editor menu item 233
 Hardware Description for Serial I/O 247
 Hdlc Serial I/O 263
 Displaying Help Messages in the Input Editor 66
 Hexadecimal format directive 32
 Home directory 163
 [Tree edit] Homedir File System Maintenance menu item 231
 Maclisp **status** **homedir** function 201
 Right horseshoe (⌞) character in ITS pathnames 185
 Create new logical host 197
 FEP host 222
 SYS logical host 186
 Pathnames on Supported Host Determination in Pathnames 136
 Host File Systems 174
 Host generic pathname component 149
 Host-independent file name 151
:host method of **fs:pathname** 165
:host-name message to tape streams 273
:host option to **tape:make-stream** 269
 Logical hosts 188
 Multiple physical hosts 196
 Splitting Logical Hosts Across Physical Hosts 196
 Translating logical hosts to physical hosts 186
 Splitting Logical Hosts Across Physical Hosts 196
 Using the Terminal Program with Hosts Connected to the Serial Line 265
 Translating logical hosts to physical hosts 186
 How LMFS Uses the FEP File System 216
 How the Input Editor Works 53
 How to Interpret Directory Listings 236

Hyper bit 5

Buffered serial	I/O	251
Hardware Description for Serial	I/O	247
Hdlc Serial	I/O	263
Introduction to Serial	I/O	245
Notes on Serial	I/O	261
Parameters for Serial	I/O	253
Simple Example: Serial	I/O	257
Troubleshooting: Serial	I/O	259
The Serial	I/O Facility	243
:ascii-characters serial	I/O parameter	253
:baud serial	I/O parameter	253
:carrier-detect serial	I/O parameter	253
:check-framing-errors serial	I/O parameter	253
:check-over-run-errors serial	I/O parameter	253
:check-parity-errors serial	I/O parameter	253
:clear-to-send serial	I/O parameter	253
:data-terminal-ready serial	I/O parameter	253
:flavor serial	I/O parameter	253
:force-output serial	I/O parameter	251, 253
:generate-xon-xoff serial	I/O parameter	253
:input-error-character serial	I/O parameter	253
:input-xoff-character serial	I/O parameter	253
:input-xon-character serial	I/O parameter	253
:mode serial	I/O parameter	253
:number-of-data-bits serial	I/O parameter	253
:number-of-stop-bits serial	I/O parameter	253
:output-xoff-character serial	I/O parameter	253
:output-xon-character serial	I/O parameter	253
:parity serial	I/O parameter	253
:request-to-send serial	I/O parameter	253
:unit serial	I/O parameter	253
:xon-xoff-protocol serial	I/O parameter	253
Reading serial	I/O parameters	251
Setting serial	I/O parameters	251
Creating a serial	I/O stream	251
The Serial	I/O Stream	251
Introduction to the	I/O Streams	11
Introduction to the	I/O System	3
:create symbol in	:if-does-not-exist option for open	222
:error symbol in	:if-does-not-exist option for open	222
nil symbol in	:if-does-not-exist option for open	222
:new-version symbol in	:if-does-not-exist option for open	222
nil symbol in	:if-exists option for open	222
:error symbol in	:if-exists option for open	222
:supersede symbol in	:if-exists option for open	222
:overwrite symbol in	:if-exists option for open	222
:share symbol in	:if-locked option for open	222, 228
:error symbol in	:if-locked option for open	222
Fsedit [Delete	(Immediate)] Command	233
[Delete	(immediate)] File System Editor menu item	233
	Increase size of FEP file	224
	Indentlon format directive	41
X- date directory last expunged	indicator	236
=> link	indicator in directory listings	236
	Indirection format directive	34
	Inferior Directory] Command	234
Fsedit [Create	Inferior Directory] File System Editor menu item	233
[Create	:info message	114
	Init file	163
ITS	init file	185

- TOPS-20
 - init file 183
 - Init File Naming Conventions 200
- fs:**
 - init-file-pathname** function 163
 - :initial-input** option 61
- List valid
 - init-options for flavor 104
 - INIT pathname type abbreviation 180
 - :in** option for **fs:complete-pathname** 120, 158
 - :in** option for **with-open-file** 91
- Clear buffered
 - input 18
- Echo character
 - input 75
- Editing terminal
 - input 53
- Encoding keyboard
 - input 5
- Stream
 - input 15
 - Input and Output Disk Streams 226
 - :input** disk stream 222
 - Input editor 53, 59, 60, 61, 62, 63, 64
 - Input Editor 66
 - Input Editor 65
 - Input Editor 66
 - Input Editor 54
 - input editor 57
 - Input Editor Example 56
 - :input-editor** message 19, 54
 - Input Editor Messages to Interactive Streams 70
 - :input-editor** method of **sl:interactive-stream** 70
 - Input Editor Options 59
 - Input Editor Program Interface 53
 - Input Editor Works 53
 - :input-error-character** serial I/O parameter 253
 - Input File 132
 - Input Functions 75
 - :input-history-default** option 62
 - Input operations on serial streams 251
 - input stream 15
 - :input-stream-mode** option to **tape:make-stream** 269
 - input streams 75
 - Input Streams 20
 - Input stream specification 75
 - input tape streams 273
 - input tape streams 273
 - input tape streams 273
 - :input-wait-handler** option 64
 - :input-wait** message 13
 - :input-wait** option 64
 - :input-xoff-character** serial I/O parameter 253
 - :input-xon-character** serial I/O parameter 253
 - Installing Microcode 217
 - install-microcode** function 217
 - Instances in compiled code files 241
 - instead of **fs:directory-list** example 132
 - :interactive** message 13
 - interactive-stream** 71
 - interactive-stream** 71
 - interactive-stream** 70
 - interactive-stream** 72
 - interactive-stream** 72
 - interactive-stream** 71
 - interactive-stream** 71
 - interactive-stream** 70
 - Interactive streams 53, 75
 - Interactive Streams 70
 - Interchange case in pathnames 145
 - Interchange case representation 145
 - Interface 227
 - Interface 53
 - Interning of Pathnames 138
 - Interpret Directory Listings 236
- Displaying Help Messages in the
 - Displaying Prompts in the
 - Examples of Use of the
 - Invoking the
 - Reading function to use
 - Command Loop
- The
 - How the
- Defaulting an Output File Pathname From an
 - Reading characters from an
- Reading characters from an
 - End-of-file on
- Special-purpose Stream Operations for Buffered
 - :clear-eof** message to
 - :discard-current-record** message to
 - :record-status** message to
- sl:**
 - fs:directory-link-opaque-dirlist**
 - :finish-typeout** method of **sl:**
 - :force-rescan** method of **sl:**
 - :input-editor** method of **sl:**
 - :noise-string-out** method of **sl:**
 - :read-bp** method of **sl:**
 - :replace-input** method of **sl:**
 - :rescanning-p** method of **sl:**
 - :start-typeout** method of **sl:**
- Input Editor Messages to
 - FEP File Properties: 3600 Disk System User
 - The Input Editor Program
- How to

Time Interval format directive 42
 Introduction to LMFS 203
 Introduction to Serial I/O 245
 Introduction to Streams 11
 Introduction to the I/O System 3
 Invoking the Input Editor 54
 Iteration format directive 37
 ITS 200
 ITS character handling 5
 ITS init file 185
:fn1 method of **fs:** **its-pathname** 186
:fn2 method of **fs:** **its-pathname** 186
 Type ITS pathname component 185
 Version ITS pathname component 185
 ITS Pathnames 185
 ITS pathnames 185
 Equivalence sign (\equiv) character in ITS pathnames 185
 Right horseshoe (\Rightarrow) character in ITS pathnames 185
fs: ***its-uninteresting-types*** variable 186

J

J

J

Justification format directive 39
 Left justify 46
 Right justify 46

K

K

K

%%kbd-char variable 6
%%kbd-control-meta variable 7
%%kbd-control variable 6
%%kbd fields 5
%%kbd-hyper variable 7
%%kbd-meta variable 6
%%kbd-mouse-button variable 7
%%kbd-mouse-n-clicks variable 7
%%kbd-mouse variable 7
%%kbd-super variable 6
 Keyboard character 5
 Keyboard character format directive 33
 keyboard input 5
 Encoding keys character representation 5
 Special function keyword 153
:wild
fs: ***known-types*** variable 149

L

L

L

Disk label comments 216
 X- date directory last expunged indicator 236
 Left justify 46
:length FEP file property 227
:length-in-blocks file property 117
:length-in-bytes file property 117
:length message 113
 Using the Terminal Program with Hosts Connected to the Serial Line 265
 Fresh line format directive 34
:line-in message 16, 24
 Testing **:line-in** operation 12
:line-out message 15
 Fsedit [Create Link] Command 234
 [Create Link] File System Editor menu item 233
 -> link indicator in directory listings 236
 LMFS Links 210
 Fsedit [Link Transparencies] Command 235
 [Link Transparencies] File System Editor menu item 233

- Formatting
 - :lisp** canonical type 200
 - Lisp Code 51
 - Lisp Machine Character Set 5
 - Lisp Machine File System 203
 - LISP pathname type abbreviation 180
- File attribute
 - list 109
- Parser for file attribute
 - list 109
 - :listen** message 14, 251
 - listf** function 123
- > link indicator in directory
 - listings 236
- How to Interpret Directory
 - Listings 236
- File Attribute
 - Lists 109
- Manipulating pathname property
 - lists 165
- Parsing file attribute
 - lists 109
- Printing
 - lists 44
 - Lists in compiled code files 241
 - List valid init-options for flavor 104
- Adding a Partition to
 - LMFS 200, 215
- Introduction to
 - LMFS 213
 - LMFS 203
- Local
 - LMFS 216
 - Lmfs.file file 213
 - LMFS Backup 212
 - LMFS Concepts 203
 - LMFS Deletion, Expunging, and Versions 209
- Renaming
 - LMFS files 176
- Adding a Spare World Load as
 - LMFS File Space 219
 - LMFS Links 210
 - LMFS Multiple Partitions 213
- Adding a
 - LMFS Partition 219
 - LMFS Pathnames 174
 - LMFS Properties 204
 - lmfs:rename-local-file-tool** function 176
- How
 - LMFS Uses the FEP File System 216
- Adding a Spare World
 - Load as LMFS File Space 219
 - Fsedit
 - [Load] Command 236
 - LOAD FEP file type 228
 - [Load] File System Editor menu item 233
 - LOAD file type (FEP) 215
- Using a Spare World
 - Load for Paging 219
 - load** function 108
 - Loading 108
 - Loading Files 108
 - Loading text files 109
- fs:**
 - load-pathname-defaults** variable 149
- Microcode
 - Loads 215
- World
 - Loads 215
- Creating More Room on the
 - Local Disk 239
 - Local LMFS 216
 - :lock-reason** option to **tape:make-stream** 269
- FEP File
 - Locks 228
- Sys: I-ucode;
 - logical directory 217
- Create new
 - logical host 197
 - logical host 186
 - Logical hosts 188
- SYS
 - Logical hosts 188
 - Logical Hosts Across Physical Hosts 196
- Splitting
 - logical hosts to physical hosts 186
- Translating
 - logical pathname 199
- Converting physical pathname to
 - logical-pathname** 199
- :back-translated-pathname** method of **fs:**
 - logical-pathname** 199
- :translated-pathname** method of **fs:**
 - Logical Pathnames 186
- :relative** relative directory specifier syntax in
 - logical pathnames 186
- :wild-inferiors** in
 - logical pathnames 188
- Double-arrow character (↔) in
 - logical pathnames for **nil** 186
- Equivalence sign character (≡) in
 - logical pathnames for quoting 186
- Converting
 - logical pathname to physical pathname 186, 199
 - Logical Pathname Translation 188, 189
- Patch file
 - logical pathname translation 196

VMS logical pathname translation rules 192
 Logical Pathname Wildcard Syntax 188
 Command Loop Input Editor Example 56
Lowercase file attribute 109
 Lower case format directive 34
 Lowercase format directive 34
 Lowercase in pathname components 145
 Enclose character in lozenge 34
 Lozenge format directive 34
 Sys: i-ucode; logical directory 217

M**M****M**

Lisp Machine Character Set 5
 Lisp Machine File System 203
 Lisp Maclisp **allfiles** function 201
 Maclisp Conversion 201
 Maclisp **crunit** function 201
 Maclisp **default** function 201
 Maclisp **directory** function 201
 Maclisp file manipulation 201
 Maclisp **mergef** function 201
 Maclisp **namelist** function 201
 Maclisp **namestring** function 201
 Maclisp **open** function 91
 Maclisp **status homedir** function 201
 Maclisp **status udir** function 201
 Maclisp **truenam** function 201
 Maclisp **uread** function 201
format:breakline macro 48
format:outfmt macro 45
format:output macro 45
format:pad macro 47
sys:with-indentation macro 49
sys:with-open-file-search macro 97
with-open-file-case macro 96
with-open-stream-case macro 97
 Streams that read or write magnetic tape 269
 Writing Programs That Use Magnetic Tape 267
 MAIL pathname type abbreviation 180
 [Tree edit any] File System Maintenance menu item 231
 [Tree edit Homedir] File System Maintenance menu item 231
 [Tree edit root] File System Maintenance menu item 231
 Editor major mode for file 109
 Editor Major modes 153
 Editor major modes 153
make-broadcast-stream function 23
si: make-coroutine-bidirectional-stream function 27
si: make-coroutine-input-stream function 27
si: make-coroutine-output-stream function 27
fs: make-logical-pathname-host function 197
fs: make-pathname-defaults 131, 133
fs: make-pathname-defaults function 162
fs: make-pathname function 145, 162
si: make-serial-stream function 252
:density option to **tape: make-stream** 269
:direction option to **tape: make-stream** 269
:host option to **tape: make-stream** 269
:input-stream-mode option to **tape: make-stream** 269
:lock-reason option to **tape: make-stream** 269
:minimum-record-length-granularity option to **tape: make-stream** 269
:minimum-record-length option to **tape: make-stream** 269
:no-bot-prompt option to **tape: make-stream** 269
:norewind option to **tape: make-stream** 269
:pad-char option to **tape: make-stream** 269
:prompt option to **tape: make-stream** 269

:record-length option to tape:	make-stream 269
:reel option to tape:	make-stream 269
:unit option to tape:	make-stream 269
tape:	make-stream function 269
The tape:	make-stream Function 269
	make-syn-stream function 23
	Making Your Own Stream 23
	Manipulating defaults 156
	Manipulating pathname property lists 165
	manipulation 201
Maclisp file	map 225
FEP file data	:map-block-no message 225
Wildcard Directory	Mapping 155
Wildcard Pathname	Mapping 153
End-of-file	marker 12
Virtual	Memory 216
	:merged-help option 61
Maclisp	mergef function 201
fs:	merge-pathnames-and-set-defaults function 158
fs:	merge-pathnames example 133, 134
fs:	merge-pathnames function 157
Pathname	merging 129
Pathname Defaults and	Merging 147
	Merging pathname components 147, 157
:advance-input-buffer	message 21
:advance-output-buffer	message 21
:allocate	message 225
:any-tyi	message 24
:beep	message 19
:block-in	message 226
:block-length	message 226
:block-out	message 227
:canonical-type	message 151
:change-properties	message 114
:change-property	message 227
:characters	message 13, 114
:clear-input	message 18, 251, 261
:clear-output	message 18
:clear-rest-of-line	message 15
:clear-window	message 20
:close	message 18
:create-data-map	message 225
:creation-date	message 114
:delete	message 114
:device	message 145
:direction	message 13
:directory	message 145
:eof	message 18
:fasd-form	message 241
:file-access-path	message 225
:finish	message 18, 114
:force-output	message 18
:fresh-line	message 15
:get	message 227
:get-output-buffer	message 21
:grow	message 224
:info	message 114
:input-editor	message 19, 54
:input-wait	message 13
:interactive	message 13
:length	message 113
:line-in	message 16, 24
:line-out	message 15
:listen	message 14, 251
:map-block-no	message 225
:name	message 145
:new-default-pathname	message 140, 145
:new-pathname	message 140, 145

:new-raw-device	message 145
:new-raw-directory	message 145
:new-raw-name	message 145
:new-raw-type	message 145
:operation-handled-p	message 13
:pathname	message 113
:print	message 73
:properties	message 114
:raw-device	message 145
:raw-directory	message 145
:raw-name	message 145
:raw-type	message 145
:read	message 73
:read-bytes	message 115
:read-cursorpos	message 19
:read-input-buffer	message 20
:read-pointer	message 20
:rename	message 114
:reset	message 251
:send-if-handles	message 13
:set-cursorpos	message 20
:set-pointer	message 20
:string-in	message 15
:string-line-in	message 17
:string-out	message 15
:truename	message 113
:tyi	message 11
:tyl-no-hang	message 19, 251
:typepeek	message 15
:tyo	message 11
:type	message 145
:untyi	message 12
:untyo	message 19
:untyo-mark	message 19
:which-operations	message 13
:write-data-map	message 225
Block disk stream	messages 226
Disk stream	messages 224
:force-output	messages 113
Stream	Messages 115
Displaying Help	Messages in the Input Editor 66
Pathname	Messages: Naming of Files 165
Input Editor	Messages to file streams 113
Error	Messages to Interactive Streams 70
:translate-wild-pathname	Messages to Tape Streams 273
:translate-wild-pathname-reversible	message stream 22
:translate-wild-pathname	message to fs:pathname 192
:clear-eof	message to fs:pathname 192
:discard-current-record	message to fs:pathname example 132
:record-status	message to input tape streams 273
:force-output	message to input tape streams 273
:write-eof	message to input tape streams 273
:write-error-status	message to output tape streams 273
:generic-pathname	message to output tape streams 273
:rename	message to output tape streams 273
:rename	message to pathnames 149
:rename	message to pathnames 103
:await-rewind	message to streams 103
:bot-p	message to tape streams 273
:check-ready	message to tape streams 273
:clear-error	message to tape streams 273
:close	message to tape streams 273
:host-name	message to tape streams 273
:rewind	message to tape streams 273
:set-offline	message to tape streams 273
:skip-file	message to tape streams 273
Meta bit	5

	:fn1	method of fs:its-pathname	186
	:fn2	method of fs:its-pathname	186
	:back-translated-pathname	method of fs:logical-pathname	199
	:translated-pathname	method of fs:logical-pathname	199
	:canonical-type	method of fs:pathname	166
	:device	method of fs:pathname	166
	:device-wild-p	method of fs:pathname	173
	:directory	method of fs:pathname	166
	:directory-pathname-as-file	method of fs:pathname	144
	:directory-wild-p	method of fs:pathname	173
	:generic-pathname	method of fs:pathname	171
	:get	method of fs:pathname	171
	:getl	method of fs:pathname	172
	:host	method of fs:pathname	165
	:name	method of fs:pathname	166
	:name-wild-p	method of fs:pathname	173
	:new-canonical-type	method of fs:pathname	169
	:new-default-pathname	method of fs:pathname	170
	:new-device	method of fs:pathname	167
	:new-directory	method of fs:pathname	167
	:new-name	method of fs:pathname	167
	:new-pathname	method of fs:pathname	170
	:new-raw-device	method of fs:pathname	168
	:new-raw-directory	method of fs:pathname	168
	:new-raw-name	method of fs:pathname	168
	:new-raw-type	method of fs:pathname	169
	:new-type	method of fs:pathname	168
	:new-version	method of fs:pathname	168
	:parse-truename	method of fs:pathname	171
	:pathname-as-directory	method of fs:pathname	144
	:pathname-match	method of fs:pathname	172
	:plist	method of fs:pathname	172
	:putprop	method of fs:pathname	172
	:raw-device	method of fs:pathname	166
	:raw-directory	method of fs:pathname	166
	:raw-name	method of fs:pathname	166
	:raw-type	method of fs:pathname	166
	:remprop	method of fs:pathname	172
	:string-for-directory	method of fs:pathname	171
	:string-for-dired	method of fs:pathname	171
	:string-for-editor	method of fs:pathname	171
	:string-for-host	method of fs:pathname	171
	:string-for-printing	method of fs:pathname	171
	:string-for-wholine	method of fs:pathname	171
	:system-type	method of fs:pathname	168
	:translate-wild-pathname	method of fs:pathname	173
	:type	method of fs:pathname	166
	:types-for-canonical-type	method of fs:pathname	169
	:type-wild-p	method of fs:pathname	173
	:version	method of fs:pathname	166
	:version-wild-p	method of fs:pathname	173
	:wild-p	method of fs:pathname	172
	:finish-typeout	method of si:interactive-stream	71
	:force-rescan	method of si:interactive-stream	71
	:input-editor	method of si:interactive-stream	70
	:noise-string-out	method of si:interactive-stream	72
	:read-bp	method of si:interactive-stream	72
	:replace-input	method of si:interactive-stream	71
	:rescanning-p	method of si:interactive-stream	71
	:start-typeout	method of si:interactive-stream	70
	:read-frame	method of si:serial-hdlc-mixin	263
	:write-frame	method of si:serial-hdlc-mixin	263
	:get	method of si:serial-stream	252
	:put	method of si:serial-stream	252
		MIC FEP file type	228
		MIC file type (FEP)	215
	Installing	Microcode	217
		Microcode Loads	215

MIDAS pathname type abbreviation 180
:minimum-record-length-granularity option to
tape:make-stream 269
:minimum-record-length option to **tape:make-**
stream 269
:minpad option for formatting functions 44
 Binary mode 91
 Character mode 91
 Block mode disk streams 221
Mode File Attribute 109
 Editor major mode for file 109
 Modem 247
 Null modem 247
 Modem control parameters 253
 Assignment of RS-232 Signals to Pins in Asynchronous Null
 Modems 247
 Modes 153
 Correspondence of Canonical Types and Editor modes 153
 Editor major modes 153
 Major modes 153
:mode serial I/O parameter 253
 Modifier bits 5
tape: mount-error flavor 277
 Mouse signals 5
 Ms-dos Pathnames 186
 Multics Pathnames 184
fs: multiple-file-plists function 119
 LMFS Multiple Partitions 213
 Multiple physical hosts 196
 Copy File (m-X) Zmacs command 105, 153
 Rename File (m-X) Zmacs command 153

N

N

N

FEP directory name 222
 Host-independent file name 151
 Special Defaulting of the Name 122, 160
 File name conventions 127
 Maclisp **namelist** function 201
:name message 145
:name method of **fs:pathname** 166
 name of file 107
 Maclisp **namestring** function 201
:name-wild-p method of **fs:pathname** 173
 Init File Naming Conventions 200
 Naming of FEP Files 222
 Naming of Files 127
 Pathname Messages: Naming of Files 165
 Native case representation 145
 Native pathname component case 145
:new-canonical-type method of **fs:pathname** 169
:new-default-pathname message 140, 145
:new-default-pathname method of
fs:pathname 170
:new-device method of **fs:pathname** 167
:new-directory method of **fs:pathname** 167
:newest pathname component 140, 170
:newest version specifier 186
 > new logical host 197
 Create **:new-name** method of **fs:pathname** 167
:new-ok option for **fs:complete-pathname** 120, 158
:new-pathname message 140, 145
:new-pathname method of **fs:pathname** 170
:new-pathname Translation Rule 192
 [New Properties] File System Editor menu item 233
 Fsedit [New Property] Command 234
:new-raw-device message 145
:new-raw-device method of **fs:pathname** 168

- Double-arrow character (↔) in logical pathnames for
 - :new-raw-directory** message 145
 - :new-raw-directory** method of **fs:pathname** 168
 - :new-raw-name** message 145
 - :new-raw-name** method of **fs:pathname** 168
 - :new-raw-type** message 145
 - :new-raw-type** method of **fs:pathname** 169
 - :new-type** method of **fs:pathname** 168
 - :new-version** method of **fs:pathname** 168
 - :new-version** symbol in **:if-exists** option for **open** 222
 - nil** 186
 - nil** pathname component 140
 - nil** symbol in **:if-does-not-exist** option for **open** 222
 - nil** symbol in **:if-exists** option for **open** 222
 - :no-bot-prompt** option to **tape:make-stream** 269
 - :noerror** option for **fs:directory-list** 117
 - :noerror** option for **with-open-file** 91
 - :no-extra-info** option for **fs:directory-list** 117
 - :no-input-save** option 63
 - :noise-string-out** method of **si:interactive-stream** 72
 - :norewind** option to **tape:make-stream** 269
 - :not-backed-up** file property 117
 - not backed up flag 236
 - Notes on Serial I/O 261
 - :notification-handler** option 64
 - not reap file flag 236
 - Null modem 247
 - Assignment of RS-232 Signals to Pins in Asynchronous Null Modems 247
 - si:** **null-stream** function 25
 - File block number 221, 226
 - Font number 5
 - Print number in floating point 46
 - Translate relative file block number into disk address 225
 - Number of bits per byte 91
 - :number-of-data-bits** serial I/O parameter 253
 - :number-of-disk-blocks** option for **open** 222, 226
 - :number-of-stop-bits** serial I/O parameter 253
 - Numbers in compiled code files 241
 - Roman numeral format directive 36
- O
- Character object 84
 - Compiled code objects in compiled code files 241
 - format:** **ochar** function 46
 - Octal format directive 31
 - Octal representation of characters 5
 - format:** **ofloat** function 46
 - :oldest** pathname component 140
 - :oldest** version specifier 186
 - :old** option for **fs:complete-pathname** 120, 158
 - onum** 45
 - onum** 45
 - onum** function 45
 - open** 222
 - open** 97
 - open** 222, 228
 - open** 93, 99
 - open** 222
 - open** 222
 - open** 222
 - open** 222
 - open** 222, 228
 - open** 222
- :commas** option for **format:**
 - :signed** option for **format:**
 - format:**
 - :create** symbol in **:if-does-not-exist** option for
 - :deleted** option to
 - :direction** option for
 - :error** Option to
 - :error** symbol in **:if-does-not-exist** option for
 - :estimated-length** option for
 - :if-does-not-exist** option for
 - :if-exists** option for
 - :if-locked** option for
 - :new-version** symbol in **:if-exists** option for

- nil** symbol in **:if-exists** option for
 - :number-of-disk-blocks** option for
 - :share** symbol in **:if-locked** option for
 - :submit** option to
 - :error** symbol in **:if-exists** option for
 - :error** symbol in **:if-locked** option for
 - nil** symbol in **:if-does-not-exist** option for
 - :overwrite** symbol in **:if-exists** option for
 - :supersede** symbol in **:if-exists** option for
- Fsedit
- Fsedit [Selective]
- [Selective]
- Maclisp
- :close**
- :read**
- Testing **:line-in**
- Testing **:string-in**
- Testing **:string-line-in**
- :tyi** stream
- Advanced General-purpose Stream
- Basic General-purpose Stream
- Basic Special-purpose Stream
- File Stream
- General-purpose Stream
- Special-purpose Stream
- The **:read** and **:print** Stream
- Special-purpose Stream
- Special-purpose Stream
- Input
- Output
- format:**
- format:**
- format:**
- Character
- Clear buffered
- Concatenating constant and variable
- Formatted
- Input and
- Defaulting an
- Direct Access
- :editor**
- :read**
- :sall**
- Slashification-related
- format:**
- Buffered
- Special-purpose Stream Operations for Buffered
- The
- :force-output** message to
- :write-eof** message to
- :write-error-status** message to
- open** 222
- open** 222, 226
- open** 222
- open** 97
- open** 222
- open** 222
- open** 222
- open** 222
- open** 222
- [Open] Command 234
- Open] Command 235
- [Open] File System Editor menu item 233
- Open] File System Editor menu item 233
- open** function 97, 222
- open** function 91
- Opening and Closing a Directory 232
- Operating on Disk Streams 224
- operation 97
- operation 73
- operation 12
- operation 12
- operation 12
- operation 12
- operation 91
- :operation-handled-p** message 13
- Operations 13
- Operations 11
- Operations 19
- Operations 113
- Operations 11
- Operations 19
- Operations 73
- Operations for Buffered Input Streams 20
- Operations for Buffered Output Streams 21
- Operations for Files 20
- operations on serial streams 251
- operations on serial streams 251
- oprint** function 46
- Ordinal format directive 36
- osting** function 46
- Other Components 123, 161
- outfmt** macro 45
- :out** option for **fs:complete-pathname** 120, 158
- :out** option for **with-open-file** 91
- output 46
- output 18
- output 44
- Output 29
- :output** disk stream 222
- Output Disk Streams 226
- Output File Pathname From an Input File 132
- Output File Streams 116
- output format style 46
- output format style 46
- output format style 46
- Output Functions 87
- output functions 89
- output** macro 45
- Output operations on serial streams 251
- Output padding 47
- OUTPUT pathname type abbreviation 180
- output stream 21
- Output Streams 21
- Output streams to buffered asynchronous devices 18, 114
- Output Subsystem 44
- output tape streams 273
- output tape streams 273
- output tape streams 273

- Converting logical pathname to physical
- Converting physical pathname to logical
 - :device** method of **fs:**
 - :device-wild-p** method of **fs:**
 - :directory** method of **fs:**
 - :directory-pathname-as-file** method of **fs:**
 - :directory-wild-p** method of **fs:**
 - :generic-pathname** method of **fs:**
 - :getl** method of **fs:**
 - :get** method of **fs:**
 - :host** method of **fs:**
 - :name** method of **fs:**
 - :name-wild-p** method of **fs:**
 - :new-canonical-type** method of **fs:**
 - :new-default-pathname** method of **fs:**
 - :new-device** method of **fs:**
 - :new-directory** method of **fs:**
 - :new-name** method of **fs:**
 - :new-pathname** method of **fs:**
 - :new-raw-device** method of **fs:**
 - :new-raw-directory** method of **fs:**
 - :new-raw-name** method of **fs:**
 - :new-raw-type** method of **fs:**
 - :new-type** method of **fs:**
 - :new-version** method of **fs:**
 - :parse-truename** method of **fs:**
 - :pathname-as-directory** method of **fs:**
 - :pathname-match** method of **fs:**
 - :plist** method of **fs:**
 - :putprop** method of **fs:**
 - :raw-device** method of **fs:**
 - :raw-directory** method of **fs:**
 - :raw-name** method of **fs:**
 - :raw-type** method of **fs:**
 - :remprop** method of **fs:**
 - :string-for-directory** method of **fs:**
 - :string-for-dired** method of **fs:**
 - :string-for-editor** method of **fs:**
 - :string-for-host** method of **fs:**
 - :string-for-printing** method of **fs:**
 - :string-for-wholine** method of **fs:**
 - :system-type** method of **fs:**
 - :translate-wild-pathname** message to **fs:**
 - :translate-wild-pathname** method of **fs:**
 - :translate-wild-pathname-reversible** message to **fs:**
 - :type** method of **fs:**
 - :types-for-canonical-type** method of **fs:**
 - :type-wild-p** method of **fs:**
 - :version** method of **fs:**
 - :version-wild-p** method of **fs:**
 - :wild-p** method of **fs:**
- UNIX
- VAX/VMS
- fs:**
 - Device generic
 - Directory generic
 - Directory TOPS-20
 - Host generic
 - :newest**
 - nil**
 - :oldest**
 - :relative**
 - Type generic
 - Type ITS
 - :unspecific**
 - :up**
- pathname 186, 199
- pathname 199
- pathname** 166
- pathname** 173
- pathname** 166
- pathname** 144
- pathname** 173
- pathname** 171
- pathname** 172
- pathname** 171
- pathname** 165
- pathname** 166
- pathname** 173
- pathname** 169
- pathname** 170
- pathname** 167
- pathname** 167
- pathname** 167
- pathname** 170
- pathname** 168
- pathname** 168
- pathname** 168
- pathname** 169
- pathname** 168
- pathname** 168
- pathname** 171
- pathname** 144
- pathname** 172
- pathname** 172
- pathname** 172
- pathname** 166
- pathname** 166
- pathname** 166
- pathname** 166
- pathname** 172
- pathname** 171
- pathname** 171
- pathname** 171
- pathname** 171
- pathname** 171
- pathname** 168
- pathname 192
- pathname** 173
- pathname 192
- pathname** 166
- pathname** 169
- pathname** 173
- pathname** 166
- pathname** 172
- pathname-abbreviations** 180
- pathname-abbreviations** 180
- :pathname-as-directory** method of **fs:pathname** 144
- pathname-attribute-list** function 111
- pathname component 149
- pathname component 149
- pathname component 183
- pathname component 149
- pathname component 140, 170
- pathname component 140
- pathname component 140
- pathname component 140
- pathname component 149
- pathname component 185
- pathname component 140, 170
- pathname component 140

- Version generic
- Version ITS
- Version TENEX
- Version TOPS-20
- :wild**
- :wild-inferiors**
- Native
- Raw
- Lowercase in
- Merging
- Uppercase in
- Values of
- * **:wild**
- Sticky
- Tailoring
- Using
- Sticky
- Center-dot character in TOPS-20
- :translate-wild-pathname** message to **fs:**
- fs:**
- Defaulting an Output File
- Wildcard
- fs:**
- File attributes as
- Manipulating
- Canonical Types in
- Canonical types in Unix
- Canonical types in Unix 4.2
- Canonical types in VMS
- Case in
- Circle-x (⊗) character in TOPS-20
- Equivalence sign (≡) character in ITS
- FEP File System
- Generic
- :generic-pathname** message to
- Host Determination in
- Interchange case in
- Interning of
- ITS
- LMFS
- Logical
- Ms-dos
- Multics
- Parsing
- Printing
- Raw case in
- Relative
- :relative** relative directory specifier syntax in logical
- :rename** message to
- Right horseshoe (⤵) character in ITS
- Surface type in
- TENEX
- TOPS-20
- TOPS-20 and TENEX
- UNIX
- UNIX 4.2
- Using
- Using **prompt-and-read** with
- pathname component 149
- pathname component 185
- pathname component 183
- pathname component 183
- pathname component 183
- pathname component 140
- pathname component 140
- pathname component case 145
- pathname component case 145
- pathname components 145
- pathname components 147, 157
- pathname components 145
- Pathname Components 140
- pathname component specifier 186
- Pathname Defaulting From the Current Context 134
- Pathname defaults 147, 158
- pathname defaults 131, 133
- Pathname Defaults 130
- pathname defaults 129
- Pathname Defaults and Merging 147
- pathname defaults example 133
- pathname display 183
- pathname** example 132
- Pathname examples 129
- pathname** flavor 153
- Pathname From an Input File 132
- Pathname Functions 156
- Pathname Mapping 153
- :pathname-match** method of **fs:pathname** 172
- Pathname merging 129
- :pathname** message 113
- Pathname Messages: Naming of Files 165
- pathname-plist** function 165
- pathname properties 138
- pathname property lists 165
- Pathnames 91, 127
- Pathnames 151
- Pathnames 176
- Pathnames 178
- Pathnames 180
- Pathnames 145
- pathnames 183
- pathnames 185
- Pathnames 176
- Pathnames 149
- pathnames 149
- Pathnames 136
- pathnames 145
- Pathnames 138
- Pathnames 185
- Pathnames 174
- Pathnames 186
- Pathnames 186
- Pathnames 184
- pathnames 156
- Pathnames 139
- pathnames 145
- pathnames 140
- pathnames 186
- pathnames 103
- pathnames 185
- pathnames 151
- pathnames 183
- pathnames 183
- Pathnames 183
- Pathnames 176
- Pathnames 178
- pathnames 129
- pathnames 129

print function 87
prin1-then-space function 87
princ function 87
print-disk-label function 216
Print file 105
print function 87
Printing lists 44
Printing Pathnames 139
format:
si: **print-list** function 44
print-list function 73
:print message 73
Print number in floating point 46
si: **print-object** function 73
Formatted **print** of function definitions 51
:print option for **fs:complete-pathname** 120, 158
:print option for **with-open-file** 91
:print Stream Operations 73
:probe disk stream 222
prober function 107
:probe option for **with-open-file** 91
Front-end Processor 221
The Input Editor Program Interface 53
Program source file 108
Writing Programs That Use Magnetic Tape 267
Using the Terminal Program with Hosts Connected to the Serial Line 265
prompt-and-read (:pathname :visible-default)
example 130, 132, 133, 134
prompt-and-read with pathnames 129
prompting for passwords 107
:prompt option 60
:prompt option to **tape:make-stream** 269
Displaying Prompts in the Input Editor 65
FEP file properties 216
FEP File Comment Properties 216
File properties 117, 204, 233
File attributes as pathname properties 138
LMFS Properties 204
FEP File Properties: 3600 Disk System User Interface 227
Fsedit [Edit Properties] Command 234
Fsedit [View Properties] Command 233
[Edit Properties] File System Editor menu item 233
[New Properties] File System Editor menu item 233
[View Properties] File System Editor menu item 233
:properties message 114
:author FEP file property 227
:author file property 117
:block-size file property 117
:byte-size file property 117
:creation-date FEP file property 227
:creation-date file property 117
:directory FEP file property 227
:dont-delete FEP file property 227
File property 204
:length FEP file property 227
:length-in-blocks file property 117
:length-in-bytes file property 117
:not-backed-up file property 117
:reference-date file property 117
:truename FEP file property 227
Fsedit [New Property] Command 234
Manipulating pathname property lists 165
EIA RS-232 protocol 243
XON/XOFF protocol parameters 253
:put method of **si:serial-stream** 252
:putprop method of **fs:pathname** 172
Putting Data in Compiled Code Files 241

Q

Q

Q

:qbin canonical type 200
QFASL pathname type abbreviation 180
query-io variable 22
 Equivalence sign character (\equiv) in logical pathnames for quoting 186
QWABL pathname type abbreviation 180

R

R

R

Specifying radix 45
 Radix directive 36
 Raw case in pathnames 145
 Raw case representation 145
:raw-device message 145
:raw-device method of **fs:pathname** 166
:raw-directory message 145
:raw-directory method of **fs:pathname** 166
:raw-name message 145
:raw-name method of **fs:pathname** 166
:raw option for **with-open-file** 91
 Raw pathname component case 145
:raw-type message 145
:raw-type method of **fs:pathname** 166
read-and-eval function 81
 The **:read** and **:print** Stream Operations 73
fs: **read-attribute-list** function 109, 111
:read-bp method of **si:interactive-stream** 72
:read-bytes message 115
sys: **read-character** function 76
readch function 84
:read-cursorpos message 19
read-delimited-string function 83
read-expression function 78
readfile function 109
read-form-completion-alist variable 80
read-form-completion-delimiters variable 80
read-form-edit-trivial-errors-p variable 79
read-form function 78
read-for-top-level function 77
:read-frame method of **si:serial-hdlc-mixin** 263
read-from-string function 84
read function 73, 76
 Reading characters from an Input stream 15
 Reading file attributes 111
 Reading function to use input editor 57
 Reading serial I/O parameters 251
:read-input-buffer message 20
readline function 81
readline-no-echo function 82
readline-or-nil function 82
readline-trim function 81
readlist function 85
 Read-locked 228
:read message 73
:read operation 73
:read option for **fs:complete-pathname** 120, 158
:read option for **with-open-file** 91
read-or-character function 80
read-or-end function 80
 Streams that **read** or write magnetic tape 269
:read output format style 46
:read-pointer message 20
read-preserve-delimiters variable 76
 * character do not reap file flag 236
:record-length option to **tape:make-stream** 269
:record-status message to input tape streams 273

:reel option to **tape:make-stream** 269
:reference-date file property 117
 Relative branch format directive 35
:relative relative directory specifier syntax in logical pathnames 186
 Translate relative file block number into disk address 225
:relative pathname component 140
 Relative pathnames 140
 Relative Pathnames: Files 149
fs: ***remember-passwords*** variable 107
 Remote file servers 91
:remprop method of **fs:pathname** 172
 Fsedit [Rename] Command 233
renamef function 103
 Rename file 103
 Rename File (m-X) Zmacs command 153
 [Rename] File System Editor menu item 233
lmfs: **rename-local-file-tool** function 176
:rename message 114
:rename message to pathnames 103
:rename message to streams 103
 Renaming FEP Files 219
 Renaming LMFS files 176
:replace-input method of **si:interactive-stream** 71
:report-stream option for **copyf** 105
 Interchange case representation 145
 Native case representation 145
 Raw case representation 145
 Special function keys character representation 5
 Octal representation of characters 5
:reprompt option 60
:request-to-send serial I/O parameter 253
:rescanning-p method of **si:interactive-stream** 71
:reset message 251
 Return character 5
 Carriage return character 5
 Values Returned 120, 159
 Carriage return format directive 34
 Reversible Wildcard Pathname Translation 191
 Reversible wild pathname translation 189
:rewind message to tape streams 273
 Right horseshoe (⤵) character in ITS pathnames 185
 Right justify 46
 Roman numeral format directive 36
 Creating More Room on the Local Disk 239
 Root directory 215
 [Tree edit root] File System Maintenance menu item 231
 EIA RS-232 protocol 243
 Assignment of RS-232 Signals to Pins 247
 Assignment of RS-232 Signals to Pins in Asynchronous Null Modems 247
rubout-handler variable 54
 Disable special treatment of Rubout in ASCII files 91
 Run-example with Simple Defaulting 130, 131

S

S

S

:sall output format style 46
 Flashing the screen 19
 SELECT F command 213
 Fsedit [Selective Open] Command 235
 [Selective Open] File System Editor menu item 233
:send-if-handles message 13
 Page separator character 5
 Page separator format directive 34
:read-frame method of **si:** **serial-hdlc-mixin** 263
:write-frame method of **si:** **serial-hdlc-mixin** 263

- si:** **serial-hdlc-stream** flavor 263
- Buffered serial I/O 251
- Hardware Description for Serial I/O 247
- Hdlc Serial I/O 263
- Introduction to Serial I/O 245
- Notes on Serial I/O 261
- Parameters for Serial I/O 253
- Simple Example: Serial I/O 257
- Troubleshooting: Serial I/O 259
- The Serial I/O Facility 243
- :ascii-characters** serial I/O parameter 253
- :baud** serial I/O parameter 253
- :carrier-detect** serial I/O parameter 253
- :check-framing-errors** serial I/O parameter 253
- :check-over-run-errors** serial I/O parameter 253
- :check-parity-errors** serial I/O parameter 253
- :clear-to-send** serial I/O parameter 253
- :data-terminal-ready** serial I/O parameter 253
- :flavor** serial I/O parameter 253
- :force-output** serial I/O parameter 251, 253
- :generate-xon-xoff** serial I/O parameter 253
- :input-error-character** serial I/O parameter 253
- :input-xoff-character** serial I/O parameter 253
- :input-xon-character** serial I/O parameter 253
- :mode** serial I/O parameter 253
- :number-of-data-bits** serial I/O parameter 253
- :number-of-stop-bits** serial I/O parameter 253
- :output-xoff-character** serial I/O parameter 253
- :output-xon-character** serial I/O parameter 253
- :parity** serial I/O parameter 253
- :request-to-send** serial I/O parameter 253
- :unit** serial I/O parameter 253
- :xon-xoff-protocol** serial I/O parameter 253
- Reading serial I/O parameters 251
- Setting serial I/O parameters 251
- Creating a serial I/O stream 251
- The Serial I/O Stream 251
- Using the Terminal Program with Hosts Connected to the Serial Line 265
- :get** method of **si:** **serial-stream** 252
- :put** method of **si:** **serial-stream** 252
- Input operations on serial streams 251
- Output operations on serial streams 251
- PDP-10 file server 91
- File server computer 127
- Remote file servers 91
- TOPS-20 file servers 91
- ASCII character set 251, 253
- Lisp Machine Character Set 5
- Symbolics Character Set 5
- The Character Set 5
- :set-cursorpos** message 20
- fs:** **set-default-pathname** 133
- fs:** **set-default-pathname** function 163
- fs:** **set-logical-pathname-host** 188, 196
- fs:** **set-logical-pathname-host** function 198
- :set-offline** message to tape streams 273
- :set-pointer** 115
- :set-pointer** message 20
- sets 247
- Setting serial I/O parameters 251
- Character set translation 10
- Disable all character set translation in ASCII files 91
- Effect of Character Set Translation on Direct Access File Streams 10
- Shared file system 127
- :share** symbol in **:if-locked** option for **open** 222
- si:****coroutine-bidirectional-stream** flavor 28
- si:****coroutine-input-stream** flavor 28

- Equivalence
- Mouse
- Assignment of RS-232
- Assignment of RS-232
- Equivalence
- :finish-typeout** method of
- :force-rescan** method of
- :input-editor** method of
- :noise-string-out** method of
- :read-bp** method of
- :replace-input** method of
- :rescanning-p** method of
- :start-typeout** method of
- Run-example with
- :read-frame** method of
- :write-frame** method of
- :get** method of
- :put** method of
- Increase
- Program
- Adding a Spare World Load as LMFS File
- Allocating Extra Paging
- Extra paging
- Swap
- Swap
- Adding a
- Using a
- format:deformat**
- fs:define-canonical-type**
- grindex**
- with-input-editing**
- with-input-editing-options**
- with-input-editing-options-if**
- with-open-file**
- with-open-stream**
- Basic
- si:coroutine-output-stream** flavor 28
- sign (\equiv) character in ITS pathnames 185
- signals 5
- Signals to Pins 247
- Signals to Pins in Asynchronous Null Modems 247
- sign character (\equiv) in logical pathnames for quoting 186
- :signed** option for **format:onum** 45
- si:install-microcode** function 217
- si:interactive-stream** 71
- si:interactive-stream** 71
- si:interactive-stream** 70
- si:interactive-stream** 72
- si:interactive-stream** 72
- si:interactive-stream** 71
- si:interactive-stream** 71
- si:interactive-stream** 70
- si:make-coroutine-bidirectional-stream** function 27
- si:make-coroutine-input-stream** function 27
- si:make-coroutine-output-stream** function 27
- si:make-serial-stream** function 252
- Simple Defaulting 130, 131
- Simple Example: Serial I/O 257
- Simple pathname usage 129
- Simple Usage of the Pathname System 129
- :single** option for **with-open-file** 91
- si:null-stream** function 25
- si:print-list** function 73
- si:print-object** function 73
- si:serial-hdlc-mixin** 263
- si:serial-hdlc-mixin** 263
- si:serial-hdlc-stream** flavor 263
- si:serial-stream** 252
- si:serial-stream** 252
- Site directory translation 195
- :site-directory** Translation Rule 195
- Site translation rules 189
- si:*typeout-default*** variable 71
- size of FEP file 224
- :skip-file** message to tape streams 273
- Slashification format directive 30
- Slashification-related output functions 89
- :sorted** option for **fs:directory-list** 117
- source file 108
- Space 219
- Space 240
- space 219
- space 216
- space 219
- Space format directive 36
- Spare World Load as LMFS File Space 219
- Spare World Load for Paging 219
- Special Defaulting of the Name 122, 160
- Special Defaulting of the Type 122, 160
- special form 42
- special form 163
- special form 51
- special form 57
- special form 55
- special form 56
- special form 91
- special form 96, 251
- Special function keys character representation 5
- Special-purpose Stream Operations 19
- Special-purpose Stream Operations 19
- Special-purpose Stream Operations for Buffered Input Streams 20
- Special-purpose Stream Operations for Buffered

- Disable
- Pathname
 - File
 - Input stream
 - Undelete
 - < :oldest version
 - < **:oldest** version
 - > **:newest** version
 - > :newest version
- * **:wild** pathname component
- Character-related byte
 - :relative relative directory
- Maclisp
- Maclisp
- :block** disk
 - Buffered output
 - Creating a serial I/O
 - Error message
 - :input** disk
 - Making Your Own
 - :output** disk
 - :probe** disk
- Reading characters from an input
 - Synonym
 - The Serial I/O
- Block disk
 - Disk
 - :tyi**
- Advanced General-purpose
- Basic General-purpose
- Basic Special-purpose
 - File
 - General-purpose
 - Special-purpose
- The **:read** and **:print**
 - Special-purpose
 - Special-purpose
 - Special-purpose
- Bidirectional disk
 - Bidirectional file
 - Block Disk
 - Block mode disk
 - Coroutine
 - Direct access bidirectional
 - Direct Access Bidirectional File
 - Direct Access File
 - Direct Access Output File
 - Disk
- Output Streams 21
- Special-purpose Stream Operations for Files 20
- special treatment of Rubout in ASCII files 91
- special variables 147
- specification 151
- specification 75
- specified file 104
- specifier 185
- specifier 186
- specifier 186
- specifier 185
- specifier 186
- specifier 186
- specifier 5
- specifier syntax in logical pathnames 186
- Specifying radix 45
- Splitting Logical Hosts Across Physical Hosts 196
- standard-input** variable 21
- standard-output** variable 21
- Standard pathname types 180
- Standard Streams 21
- :start-typeout** method of **sl:interactive-stream** 70
- status homedir** function 201
- status udir** function 201
- Sticky pathname defaults 131, 133
- Sticky pathname defaults example 133
- Stream 21, 22
- stream 222
- stream 21
- stream 251
- stream 22
- stream 222
- Stream 23
- stream 222
- stream 222
- stream 15
- stream 23
- Stream 251
- stream-copy-until-eof** function 87
- Stream default handler 11
- stream-default-handler** function 24
- Stream input 15
- Stream Messages 115
- stream messages 226
- stream messages 224
- stream operation 91
- Stream Operations 13
- Stream Operations 11
- Stream Operations 19
- Stream Operations 113
- Stream Operations 11
- Stream Operations 19
- Stream Operations 73
- Stream Operations for Buffered Input Streams 20
- Stream Operations for Buffered Output Streams 21
- Stream Operations for Files 20
- Streams 1, 243
- streams 222, 226
- streams 115
- Streams 226
- streams 221
- Streams 25
- streams 116
- Streams 116
- Streams 115
- Streams 116
- Streams 221
- Effect of Character Set Translation on Direct Access File Streams 10

- End-of-file on input streams 75
- Functions for Creating Coroutine Streams 27
- I/O Streams 11
- Input and Output Disk Streams 226
- Input Editor Messages to Interactive Streams 70
- Input operations on serial streams 251
 - Interactive streams 53, 75
- Introduction to Streams 11
- Messages to file streams 113
- Messages to Tape Streams 273
- Operating on Disk Streams 224
- Output operations on serial streams 251
- :rename** message to streams 103
- Special-purpose Stream Operations for Buffered Input Streams 20
 - Special-purpose Stream Operations for Buffered Output Streams 21
 - Standard Streams 21
 - Synonym streams 21
 - Tape streams 269
 - :await-rewind** message to tape streams 273
 - :bot-p** message to tape streams 273
 - :check-ready** message to tape streams 273
 - :clear-eof** message to input tape streams 273
 - :clear-error** message to tape streams 273
 - :close** message to tape streams 273
 - :discard-current-record** message to input tape streams 273
 - :force-output** message to output tape streams 273
 - :host-name** message to tape streams 273
 - :record-status** message to input tape streams 273
 - :rewind** message to tape streams 273
 - :set-offline** message to tape streams 273
 - :skip-file** message to tape streams 273
 - :write-eof** message to output tape streams 273
 - :write-error-status** message to output tape streams 273
 - Input stream specification 75
 - Streams that read or write magnetic tape 269
 - streams to buffered asynchronous devices 18, 114
 - :string-for-directory** method of **fs:pathname** 171
 - :string-for-dired** method of **fs:pathname** 171
 - :string-for-editor** method of **fs:pathname** 171
 - :string-for-host** method of **fs:pathname** 171
 - :string-for-printing** method of **fs:pathname** 171
 - :string-for-wholine** method of **fs:pathname** 171
 - :string-in** message 15
 - :string-in** operation 12
 - :string-line-in** message 17
 - :string-line-in** operation 12
 - :string-out** message 15
 - :editor** output format style 46
 - :read** output format style 46
 - :sall** output format style 46
 - :submit** option to **open** 97
 - Subsystem 44
 - Successful Completion 123, 161
 - Super bit 5
 - :super-image** option for **with-open-file** 91
 - :supersede** symbol in **:if-exists** option for **open** 222
 - Supplied translation rules 189
 - Supported Host File Systems 174
 - :suppress-notifications** option 64
 - Suppress prompting for passwords 107
 - Pathnames on
 - Default surface type 151
 - Preferred surface type 151
 - Surface type in pathnames 151
 - Swap space 216
 - Swap space 219
 - Symbolics Character Set 5

- :discard-current-record** message to input tape streams 273
- :force-output** message to output tape streams 273
- :host-name** message to tape streams 273
- :record-status** message to input tape streams 273
- :rewind** message to tape streams 273
- :set-offline** message to tape streams 273
- :skip-file** message to tape streams 273
- :write-eof** message to output tape streams 273
- :write-error-status** message to output tape streams 273
- tape:tape-device-error** flavor 277
- tape:tape-error** flavor 277
- :temporary** option for **with-open-file** 91
- Version TENEX pathname component 183
- TOPS-20 and Data TENEX pathnames 183
- Editing Terminal Equipment 247
- Using the terminal input 53
- Data **terminal-io** variable 22
- Data Circuit Terminal Program with Hosts Connected to the Serial Line 265
- terminals 247
- Data Terminating Equipment 247
- terpri function 87
- :terpri** option for **format:tab** 47
- Testing **:line-in** operation 12
- Testing **:string-in** operation 12
- Testing **:string-line-in** operation 12
- Loading text files 109
- Streams TEXT pathname type abbreviation 180
- Writing Programs that read or write magnetic tape 269
- Date and That Use Magnetic Tape 267
- Tilde format directive 34
- Time format directive 42
- time format directive 42
- Time interval format directive 42
- Time-of-day format directive 42
- TOPS-20 200
- TOPS-20 and TENEX Pathnames 183
- TOPS-20 file servers 91
- :tops20-heuristicate-directory** Translation Rule 194
- TOPS-20 init file 183
- Directory TOPS-20 pathname component 183
- Version TOPS-20 pathname component 183
- Center-dot character in TOPS-20 pathname display 183
- Circle-x (⊗) character in TOPS-20 pathnames 183
- TOPS-20 pathnames 183
- trace-output** variable 22
- :translated-pathname** method of **fs:logical-pathname** 199
- Translate relative file block number into disk address 225
- :translate-wild-pathname** message to **fs:pathname** 192
- :translate-wild-pathname** message to **fs:pathname** example 132
- :translate-wild-pathname** method of **fs:pathname** 173
- :translate-wild-pathname-reversible** message to **fs:pathname** 192
- :translate-wild** Translation Rule 192
- Translating logical hosts to physical hosts 186
- Translating pathname type field 151
- Character set translation 10
- Logical Pathname Translation 188, 189
- Patch file logical pathname translation 196
- Reversible Wildcard Pathname Translation 191
- Reversible wild pathname translation 189
- Site directory translation 195

OUTPUT pathname type abbreviation 180
 PATCH-DIRECTORY pathname type abbreviation 180
 (PDIR) pathname type abbreviation 180
 PRESS pathname type abbreviation 180
 QFASL pathname type abbreviation 180
 QWABL pathname type abbreviation 180
 TEXT pathname type abbreviation 180
 ULOAD pathname type abbreviation 180
 UNFASL pathname type abbreviation 180
 XMAIL pathname type abbreviation 180
 LOAD file type (FEP) 215
 MIC file type (FEP) 215
 Translating pathname type field 151
 Type generic pathname component 149
 Surface type in pathnames 151
 Type ITS pathname component 185
 :type message 145
 :type method of **fs:pathname** 166
 typeout-default variable 71
 si:
 FEP File Types 228
 Standard pathname types 180
 Correspondence of Canonical Types and Editor Modes 153
 Canonical types example 130
 :types-for-canonical-type method of
 fs:pathname 169
 Canonical Types in Pathnames 151
 Canonical types in Unix 4.2 Pathnames 178
 Canonical types in Unix Pathnames 176
 Canonical types in VMS Pathnames 180
 :type-wild-p method of **fs:pathname** 173

U

U

U

Macisp **status** **udlr** function 201
 ULOAD pathname type abbreviation 180
 Fsedit [Undelete] Command 233
 undeletef function 104
 [Undelete] File System Editor menu item 233
 Undelete specified file 104
 UNFASL pathname type abbreviation 180
 :unit option for **format:tab** 47
 :unit option to **tape:make-stream** 269
 :unit serial I/O parameter 253
 UNIX 200
 UNIX 4.2 Pathnames 178
 Canonical types in Unix 4.2 Pathnames 178
 :unix-font Translation Rule 195
 unix-microcode-translation-alist 194
 :unix-microcode Translation Rule 194
 UNIX pathname abbreviations 180
 UNIX Pathnames 176
 Canonical types in Unix Pathnames 176
 :unix-type-and-version Translation Rule 195
 :unspecific pathname component 140, 170
 :untyl message 12
 :untyo-mark message 19
 :untyo message 19
 ! character file not backed up flag 236
 :up pathname component 140
 Uppercase in pathname components 145
 Macisp **uread** function 201
 Simple pathname usage 129
 Simple Usage of the Pathname System 129
 USART chip 247
 Getting a Filename From the User 129
 User-defined file attributes 109
 User-function format directive 41

fs: **user-homedir** function 163
 FEP File Properties: 3600 Disk System User Interface 227
 How LMFS Uses the FEP File System 216

V

V

V

List
 valid init-options for flavor 104
 Values of Pathname Components 140
 Values Returned 120, 159
 variable 6
 variable 6
dbg:*debug-io-override* variable 22
debug-io variable 22
error-output variable 22
fs:*default-pathname-defaults* variable 148
fs:*file-type-mode-alist* variable 153
fs:*its-uninteresting-types* variable 186
fs:*known-types* variable 149
fs:load-pathname-defaults variable 149
fs:*remember-passwords* variable 107
 %%kbd-char variable 6
 %%kbd-control variable 6
 %%kbd-control-meta variable 7
 %%kbd-hyper variable 7
 %%kbd-meta variable 6
 %%kbd-mouse variable 7
 %%kbd-mouse-button variable 7
 %%kbd-mouse-n-clicks variable 7
 %%kbd-super variable 6
query-io variable 22
read-form-completion-alist variable 80
read-form-completion-delimiters variable 80
read-form-edit-trivial-errors-p variable 79
read-preserve-delimiters variable 76
rubout-handler variable 54
sl:*timeout-default* variable 71
standard-input variable 21
standard-output variable 21
terminal-io variable 22
trace-output variable 22
 Concatenating constant and variable output 44
 Pathname special variables 147
 VAX/VMS pathname abbreviations 180
 VAX/VMS Pathnames 180
 version 147
 Version generic pathname component 149
 Version ITS pathname component 185
:version method of **fs:pathname** 166
 Versions 209
 version specifier 185
 version specifier 186
 version specifier 186
 version specifier 185
 Version TENEX pathname component 183
 Version TOPS-20 pathname component 183
:version-wild-p method of **fs:pathname** 173
Fsedit [View] Command 234
viewf function 105
 [View] File System Editor menu item 233
Fsedit [View Properties] Command 233
 [View Properties] File System Editor menu item 233
 Virtual Memory 216
prompt-and-read (:pathname) **:visible-default** example 130, 132, 133, 134
 VMS 200
:vms-font Translation Rule 193
:vms-heuristicate-directory Translation Rule 193
:vms-heuristicate-name Translation Rule 193

Canonical types in

- :vms-heuristicate** Translation Rule 192
- VMS logical pathname translation rules 192
- :vms-microcode** Translation Rule 194
- :vms-new-pathname** Translation Rule 193
- VMS Pathnames 180

W**W****W**

Fsedit

Pathname

Reversible

Logical Pathname

Reversible

Reversible

sys:

sys:

How the Input Editor
Adding a Spare
Using a Spare

Streams that read or

- :which-operations** message 13
- [Wildcard Delete] Command 233
- [Wildcard Delete] File System Editor menu item 233
- Wildcard Directory Mapping 155
- wildcard example 132
- Wildcard Pathname Mapping 153
- Wildcard pathnames 153
- Wildcard Pathname Translation 191
- Wildcards 188
- Wildcard Syntax 188
- :wild-inferiors** in logical pathnames 188
- :wild-inferiors** pathname component 140
- :wild** keyword 153
- :wild** pathname component 140
- :wild** pathname component specifier 186
- Wild pathname translation 189
- wild pathname translation 189
- :wild-p** example 132
- :wild-p** method of **fs:pathname** 172
- with-indentation** macro 49
- with-input-editing-options-if** special form 56
- with-input-editing-options** special form 55
- with-input-editing** special form 57
- with-open-file-case** macro 96
- with-open-file-search** macro 97
- with-open-file** special form 91
- with-open-stream-case** macro 97
- with-open-stream** special form 96, 251
- Works 53
- World Load as LMFS File Space 219
- World Load for Paging 219
- World Loads 215
- :write-data-map** message 225
- :write-eof** message to output tape streams 273
- :write-error-status** message to output tape streams 273
- :write-frame** method of **si:serial-hdlc-mixin** 263
- Write-locked 228
- write magnetic tape 269
- :write** option for **fs:complete-pathname** 120, 158
- :write** option for **with-open-file** 91
- Writing Programs That Use Magnetic Tape 267

X**X****X**

- X= date directory last expunged indicator 236
- XMAIL pathname type abbreviation 180
- XON/XOFF protocol parameters 253
- :xon-xoff-protocol** serial I/O parameter 253

Y**Y****Y**

Making Your Own Stream 23

Z

Z
Copy File (m-X) Zmacs command 105, 153

Z

