

*Text Processing
with Regular Expressions*

2nd Edition



sed & awk

Pocket Reference

O'REILLY®

Arnold Robbins

sed & awk Pocket Reference



For people who create and modify text files, *sed* and *awk* are power tools for editing. *sed*, *awk*, and regular expressions allow programmers and system administrators to automate editing tasks that need to be performed on one or more files, to simplify the task of performing the same edits on multiple files, and to write conversion programs.

The *sed & awk Pocket Reference* is a concise summary of *sed*, *awk*, regular expressions, and pattern matching. This new edition has expanded coverage of *gawk* (GNU *awk*), and includes:

- An overview of *sed* and *awk*'s command-line syntax
- Alphabetical summaries of commands, including *nawk* and *gawk*
- Profiling with *pgawk*
- Coprocesses and sockets with *gawk*
- Internationalization with *gawk*
- A listing of resources for *sed* and *awk* users

Arnold Robbins is a professional programmer and technical author and coauthor of various O'Reilly Unix titles. He has been working with Unix systems since 1980, and currently maintains *gawk* and its documentation.

www.oreilly.com

US \$19.99

CAN \$26.99

ISBN: 978-0-596-00352-4



SECOND EDITION

sed & awk
Pocket Reference

Arnold Robbins

O'REILLY®

Beijing • Cambridge • Farnham • Köln • Paris • Sebastopol • Taipei • Tokyo

sed & awk Pocket Reference, Second Edition

by Arnold Robbins

Copyright © 2002, 2000 O'Reilly Media, Inc. All rights reserved.
Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North,
Sebastopol, CA 95472.

O'Reilly Media, Inc. books may be purchased for educational,
business, or sales promotional use. Online editions are also available
for most titles (*safari.oreilly.com*). For more information, contact our
corporate/institutional sales department: (800) 998-9938 or
corporate@oreilly.com.

Editor: Chuck Toporek
Production Editor: Jane Ellin
Cover Designer: Ellie Volckhausen
Interior Designer: David Futato

Printing History:

January 2000:	First Edition.
June 2002:	Second Edition.

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly Media, Inc. The *Pocket Reference* series designation, *sed & awk Pocket Reference, Second Edition*, the image of slender lorises, and related trade dress are trademarks of O'Reilly Media, Inc. Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc. was aware of a trademark claim, the designations have been printed in caps or initial caps. While every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

Contents

Introduction	1
Conventions Used in This Book	1
Matching Text	2
Filenames Versus Patterns	2
Metacharacters	3
Metacharacters, Listed by Unix Program	5
Examples of Searching	7
The sed Editor	10
Command-Line Syntax	11
Syntax of sed Commands	12
Group Summary of sed Commands	13
Alphabetical Summary of sed Commands	15
The awk Programming Language	18
Command-Line Syntax	20
Important gawk Options	21
Profiling with pgawk	23
Patterns and Procedures	24
Built-in Variables	26
Operators	28
Variable and Array Assignment	29
Octal and Hexadecimal Constants in gawk	30
User-Defined Functions	31

Group Listing of awk Functions and Commands	32
Coprocesses and Sockets with gawk	33
Implementation Limits	34
Alphabetical Summary of awk Functions and Commands	34
Internationalization with gawk	44
Additional Resources	45
Source Code	45
Books	46

sed & awk Pocket Reference

Introduction

This pocket reference is a companion volume to O'Reilly's *sed & awk*, Second Edition, by Dale Dougherty and Arnold Robbins, and to *Effective awk Programming*, Third Edition, by Arnold Robbins. It presents a concise summary of regular expressions and pattern matching, and summaries of *sed*, *awk*, and *gawk* (GNU *awk*).

Conventions Used in This Book

This pocket reference follows certain typographic conventions, outlined here:

Constant Width

Used for code examples, commands, directory names, and options.

Constant Width Italic

Used in syntax and command summaries to show replaceable text; this text should be replaced with user-supplied values.

Constant Width Bold

Used in code examples to show commands or other text that should be typed literally by the user.

Italic

Used to show generic arguments and options; these should be replaced with user-supplied values. Italic is also used to highlight comments in examples, to introduce new terms, and to indicate filenames.

\$

Used in some examples as the Bourne shell or Korn shell prompt.

[]

Surround optional elements in a description of syntax. (The brackets themselves should never be typed.)

Matching Text

A number of Unix text-processing utilities let you search for, and in some cases change, text patterns rather than fixed strings. These utilities include the editing programs *ed*, *ex*, *vi*, and *sed*, the *awk* programming language, and the commands *grep* and *egrep*. Text patterns (formally called *regular expressions*) contain normal characters mixed with special characters (called *metacharacters*).

Filenames Versus Patterns

Metacharacters used in pattern matching are different from metacharacters used for filename expansion. When you issue a command on the command line, special characters are seen first by the shell, then by the program; therefore, unquoted metacharacters are interpreted by the shell for filename expansion. For example, the command:

```
$ grep [A-Z]* chap[12]
```

could be transformed by the shell into:

```
$ grep Array.c Bug.c Comp.c chap1 chap2
```


and would then try to find the pattern *Array.c* in files *Bug.c*, *Comp.c*, *chap1*, and *chap2*. To bypass the shell and pass the special characters to *grep*, use quotes as follows:

```
$ grep "[A-Z]*" chap[12]
```

Double quotes suffice in most cases, but single quotes are the safest bet.

Note also that in pattern matching, `?` matches zero or one instance of a regular expression; in filename expansion, `?` matches a single character.

Metacharacters

Different metacharacters have different meanings, depending upon where they are used. In particular, regular expressions used for searching through text (matching) have one set of metacharacters, while the metacharacters used when processing replacement text have a different set. These sets also vary somewhat per program. This section covers the metacharacters used for searching and replacing, with descriptions of the variants in the different utilities.

Search patterns

The characters in the following table have special meaning only in search patterns:

Character	Pattern
.	Match any <i>single</i> character except newline. Can match newline in <i>awk</i> .
*	Match any number (or none) of the single character that immediately precedes it. The preceding character can also be a regular expression. For example, since <code>.</code> (dot) means any character, <code>.*</code> means "match any number of any character."
^	Match the following regular expression at the beginning of the line or string.
\$	Match the preceding regular expression at the end of the line or string.
\	Turn off the special meaning of the following character.

Character	Pattern
[]	Match any <i>one</i> of the enclosed characters. A hyphen (-) indicates a range of consecutive characters. A circumflex (^) as the first character in the brackets reverses the sense: it matches any one character <i>not</i> in the list. A hyphen or close bracket (]) as the first character is treated as a member of the list. All other metacharacters are treated as members of the list (i.e., literally).
{ <i>n,m</i> }	Match a range of occurrences of the single character that immediately precedes it. The preceding character can also be a metacharacter. { <i>n</i> } matches exactly <i>n</i> occurrences; { <i>n</i> ,} matches at least <i>n</i> occurrences; and { <i>n</i> , <i>m</i> } matches any number of occurrences between <i>n</i> and <i>m</i> . <i>n</i> and <i>m</i> must be between 0 and 255, inclusive.
\{ <i>n,m</i> \}	Just like { <i>n,m</i> }, but with backslashes in front of the braces.
\(\)	Save the pattern enclosed between \ (and \) into a special holding space. Up to nine patterns can be saved on a single line. The text matched by the subpatterns can be “replayed” in substitutions by the escape sequences \1 to \9.
\ <i>n</i>	Replay the <i>n</i> th sub-pattern enclosed in \ (and \) into the pattern at this point. <i>n</i> is a number from 1 to 9, with 1 starting on the left.
\< \>	Match characters at beginning (\<) or end (\>) of a word.
+	Match one or more instances of preceding regular expression.
?	Match zero or one instances of preceding regular expression.
	Match the regular expression specified before or after.
()	Apply a match to the enclosed group of regular expressions.

Many Unix systems allow the use of POSIX *character classes* within the square brackets that enclose a group of characters. These are typed enclosed in [: and :]. For example, [[:alnum:]] matches a single alphanumeric character.

Class	Characters matched
alnum	Alphanumeric characters
alpha	Alphabetic characters
blank	Space or TAB
cntrl	Control characters
digit	Decimal digits
graph	Nonspace characters

Class	Characters matched
lower	Lowercase characters
print	Printable characters
space	Whitespace characters
upper	Uppercase characters
xdigit	Hexadecimal digits

Replacement patterns

The characters in the following table have special meaning only in replacement patterns:

Character	Pattern
\	Turn off the special meaning of the following character.
\n	Restore the text matched by the <i>n</i> th pattern previously saved by \ (and \). <i>n</i> is a number from 1 to 9, with 1 starting on the left.
&	Reuse the text matched by the search pattern as part of the replacement pattern.
~	Reuse the previous replacement pattern in the current replacement pattern. Must be the only character in the replacement pattern (<i>ex</i> and <i>vi</i>).
%	Reuse the previous replacement pattern in the current replacement pattern. Must be the only character in the replacement pattern (<i>ed</i>).
\u	Convert first character of replacement pattern to uppercase.
\U	Convert entire replacement pattern to uppercase.
\l	Convert first character of replacement pattern to lowercase.
\L	Convert entire replacement pattern to lowercase.
\E	Turn off previous \U or \L.
\e	Turn off previous \u or \l.

Metacharacters, Listed by Unix Program

Some metacharacters are valid for one program but not for another. Those that are available to a Unix program are marked by a bullet (•) in the following table. (This table is

correct for SVR4 and Solaris and most commercial Unix systems, but it's always a good idea to verify your system's behavior.) Items marked with a "P" are specified by POSIX; double check your system's version. Full descriptions were provided in the previous section.

Symbol	ed	ex	vi	sed	awk	grep	egrep	Action
.	•	•	•	•	•	•	•	Match any character.
*	•	•	•	•	•	•	•	Match zero or more preceding.
^	•	•	•	•	•	•	•	Match beginning of line/string.
\$	•	•	•	•	•	•	•	Match end of line/string.
\	•	•	•	•	•	•	•	Escape following character.
[]	•	•	•	•	•	•	•	Match one from a set.
\(\)	•	•	•	•		•		Store pattern for later replay. ^a
\n	•	•	•	•		•		Replay sub-pattern in match.
{ }					•P		•P	Match a range of instances.
\{ \}	•			•		•		Match a range of instances.
\< \>	•	•	•					Match word's beginning or end.
+					•		•	Match one or more preceding.
?					•		•	Match zero or one preceding.
					•		•	Separate choices to match.
()					•		•	Group expressions to match.

^a Stored sub-patterns can be "replayed" during matching. See the examples in the next table.

Note that in *ed*, *ex*, *vi*, and *sed*, you specify both a search pattern (on the left) and a replacement pattern (on the right). The metacharacters listed in this table are meaningful only in a search pattern.

In *ed*, *ex*, *vi*, and *sed*, the following metacharacters are valid only in a replacement pattern:

Symbol	ex	vi	sed	ed	Action
\	•	•	•	•	Escape following character.
\n	•	•	•	•	Text matching pattern stored in \ (\).
&	•	•	•	•	Text matching search pattern.
~	•	•			Reuse previous replacement pattern.
%				•	Reuse previous replacement pattern.
\u \U	•	•			Change character(s) to uppercase.
\l \L	•	•			Change character(s) to lowercase.
\E	•	•			Turn off previous \U or \L.
\e	•	•			Turn off previous \u or \l.

Examples of Searching

When used with *grep* or *egrep*, regular expressions should be surrounded by quotes. (If the pattern contains a \$, you must use single quotes; e.g., '*pattern*'.) When used with *ed*, *ex*, *sed*, and *awk*, regular expressions are usually surrounded by / although (except for *awk*), any delimiter works. Here are some example patterns:

Pattern	What does it match?
bag	The string <i>bag</i> .
^bag	<i>bag</i> at the beginning of the line.
bag\$	<i>bag</i> at the end of the line.

Pattern	What does it match?
<code>^bag\$</code>	<i>bag</i> as the only word on the line.
<code>[Bb]ag</code>	<i>Bag</i> or <i>bag</i> .
<code>b[aeiou]g</code>	Second letter is a vowel.
<code>b[^aeiou]g</code>	Second letter is a consonant (or uppercase or symbol).
<code>b.g</code>	Second letter is any character.
<code>^...\$</code>	Any line containing exactly three characters.
<code>^\.</code>	Any line that begins with a dot.
<code>^\.[a-z][a-z]</code>	Same as previous, followed by two lowercase letters (e.g., <i>troff</i> requests).
<code>^\.[a-z]\{2\}</code>	Same as previous; <i>ed</i> , <i>grep</i> and <i>sed</i> only.
<code>^[^.]</code>	Any line that doesn't begin with a dot.
<code>bugs*</code>	<i>bug</i> , <i>bugs</i> , <i>bugss</i> , etc.
<code>"word"</code>	A word in quotes.
<code>"*word"*</code>	A word, with or without quotes.
<code>[A-Z][A-Z]*</code>	One or more uppercase letters.
<code>[A-Z]+</code>	Same as previous; <i>egrep</i> or <i>awk</i> only.
<code>[[:upper:]]+</code>	Same as previous; POSIX <i>egrep</i> or <i>awk</i> .
<code>[A-Z].*</code>	An uppercase letter, followed by zero or more characters.
<code>[A-Z]*</code>	Zero or more uppercase letters.
<code>[a-zA-Z]</code>	Any letter, either lower- or uppercase.
<code>^[^0-9A-Za-z]</code>	Any symbol or space (not a letter or a number).
<code>^[[:alnum:]]</code>	Same, using POSIX character class.

egrep or awk pattern	What does it match?
<code>[567]</code>	One of the numbers <i>5</i> , <i>6</i> , or <i>7</i> .
<code>five six seven</code>	One of the words <i>five</i> , <i>six</i> , or <i>seven</i> .
<code>80[2-4]?86</code>	<i>8086</i> , <i>80286</i> , <i>80386</i> , or <i>80486</i> .
<code>80[2-4]?86 Pentium</code>	<i>8086</i> , <i>80286</i> , <i>80386</i> , <i>80486</i> , or <i>Pentium</i> .
<code>compan(y ies)</code>	<i>company</i> or <i>companies</i> .

ex or vi pattern	What does it match?
<code>\<the</code>	Words like <i>theater</i> , <i>there</i> , or <i>the</i> .
<code>the\></code>	Words like <i>breathe</i> , <i>seethe</i> , or <i>the</i> .
<code>\<the\></code>	The word <i>the</i> .

ed, sed, or grep pattern	What does it match?
<code>0\{5,\}</code>	Five or more zeros in a row.
<code>[0-9]\{3\}-[0-9]\{2\}-[0-9]\{4\}</code>	U.S. Social Security number (<i>nnn-nn-nnnn</i>).
<code>\(why\).*\1</code>	A line with two occurrences of <i>why</i> .
<code>\([[[:alpha:]]_][[:alnum:]]_.*\)</code> = <code>\1</code> ;	C/C++ simple assignment statements.

Examples of searching and replacing

The following examples show the metacharacters available to *sed* or *ex*. Note that *ex* commands begin with a colon. A space is marked by a ■; a TAB is marked by a →.

Command	Result
<code>s/.*(&)/</code>	Redo the entire line, but add parentheses.
<code>s/.*mv & &.old/</code>	Change a wordlist (one word per line) into <i>mv</i> commands.
<code>/^\$/d</code>	Delete blank lines.
<code>:g/^\$/d</code>	Same as previous, in <i>ex</i> editor.
<code>/^[■→]*\$/d</code>	Delete blank lines, plus lines containing only spaces or TABs.
<code>:g/^[■→]*\$/d</code>	Same as previous, in <i>ex</i> editor.
<code>s/■*/■/g</code>	Turn one or more spaces into one space.
<code>:%s/■*/■/g</code>	Same as previous, in <i>ex</i> editor.
<code>:s/[0-9]/Item &:/</code>	Turn a number into an item label (on the current line).
<code>:s</code>	Repeat the substitution on the first occurrence.
<code>:&</code>	Same as previous.
<code>:sg</code>	Same as previous, but for all occurrences on the line.

Command	Result
:&g	Same as previous.
:%&g	Repeat the substitution globally (i.e., on all lines).
., \$s/Fortran/\U&g	On current line to last line, change word to uppercase.
:%s/./\L&/	Lowercase entire file.
:s/\<.\>\u&/g	Uppercase first letter of each word on current line. (Useful for titles.)
:%s/yes/No/g	Globally change a word to <i>No</i> .
:%s/Yes/~>/g	Globally change a different word to <i>No</i> (previous replacement).

Finally, here are some *sed* examples for transposing words. A simple transposition of two words might look like this:

```
s/die or do/do or die/
```

The real trick is to use hold buffers to transpose variable patterns. For example, to transpose using hold buffers:

```
s/\([Dd]ie\)\ or \([Dd]o\)/\2 or \1/
```

The sed Editor

The stream editor, *sed*, is a noninteractive editor. It interprets a script and performs the actions in the script. *sed* is stream-oriented because, like many Unix programs, input flows through the program and is directed to standard output. For example, *sort* is stream-oriented; *vi* is not. *sed*'s input typically comes from a file or pipe, but it can also be directed from the keyboard. Output goes to the screen by default but can be captured in a file or sent through a pipe instead.

Typical uses of *sed* include:

- Editing one or more files automatically
- Simplifying repetitive edits to multiple files
- Writing conversion programs

sed operates as follows:

- Each line of input is copied into a *pattern space*, an internal buffer where editing operations are performed.
- All editing commands in a *sed* script are applied, in order, to each line of input.
- Editing commands are applied to all lines (globally) unless line addressing restricts the lines affected.
- If a command changes the input, subsequent commands and address tests will be applied to the current line in the pattern space, not the original input line.
- The original input file is unchanged because the editing commands modify a copy of each original input line. The copy is sent to standard output (but can be redirected to a file).
- *sed* also maintains the *hold space*, a separate buffer that can be used to save data for later retrieval.

Command-Line Syntax

The syntax for invoking *sed* has two forms:

```
sed [-n] [-e] 'command' file(s)
sed [-n] -f scriptfile file(s)
```

The first form allows you to specify an editing command on the command line, surrounded by single quotes. The second form allows you to specify a *scriptfile*, a file containing *sed* commands. Both forms may be used together, and they may be used multiple times. If no *file(s)* is specified, *sed* reads from standard input.

The following options are recognized:

-n

Suppress the default output; *sed* displays only those lines specified with the *p* command or with the *p* flag of the *s* command.

-e cmd

Next argument is an editing command. Useful if multiple scripts or commands are specified.

-f file

Next argument is a file containing editing commands.

If the first line of the script is *#n*, *sed* behaves as if *-n* had been specified.

Syntax of *sed* Commands

sed commands have the general form:

```
[address[,address]][!]command [arguments]
```

sed copies each line of input into the pattern space. *sed* instructions consist of *addresses* and editing *commands*. If the address of the command matches the line in the pattern space, then the command is applied to that line. If a command has no address, then it is applied to each input line. If a command changes the contents of the pattern space, subsequent commands and addresses will be applied to the current line in the pattern space, not the original input line.

addresses are described in the next section. *commands* consist of a single letter or symbol; they are described later, alphabetically and by group. *arguments* include the label supplied to *b* or *t*, the filename supplied to *r* or *w*, and the substitution flags for *s*.

Pattern addressing

A *sed* command can specify zero, one, or two addresses. An address can be a line number, the symbol *\$* (for last line), or a regular expression enclosed in slashes (*/pattern/*). Regular expressions are described in “Matching Text.” Additionally, *\n* can be used to match any newline in the pattern space (resulting from the *N* command), but not the newline at the end of the pattern space.

If the command specifies:	Then the command is applied to:
No address	Each input line.
One address	Any line matching the address. Some commands accept only one address: a, i, r, q, and =.
Two comma-separated addresses	First matching line and all succeeding lines up to and including a line matching the second address.
An address followed by !	All lines that do <i>not</i> match the address.

Examples

Command	Action performed
<code>s/xx/yy/g</code>	Substitute on all lines (all occurrences).
<code>/BSD/d</code>	Delete lines containing BSD.
<code>/^BEGIN/,/^END/p</code>	Print between BEGIN and END, inclusive.
<code>/SAVE/!d</code>	Delete any line that doesn't contain SAVE.
<code>/BEGIN/,/END/!s/xx/yy/g</code>	Substitute on all lines, except between BEGIN and END.

Braces (`{ }`) are used in *sed* to nest one address inside another or to apply multiple commands to the matched same address.

```
[/pattern/[,/pattern/]]{
  command1
  command2
}
```

The opening curly brace must end its line, and the closing curly brace must be on a line by itself. Be sure there are no spaces after the braces.

Group Summary of sed Commands

In the lists that follow, the *sed* commands are grouped by function and are described tersely. Full descriptions, including

syntax and examples, can be found afterward in the “Alphabetical Summary of sed Commands” section.

Basic editing

a\	Append text after a line.
c\	Replace text (usually a text block).
i\	Insert text before a line.
d	Delete lines.
s	Make substitutions.
y	Translate characters (like Unix <i>tr</i>).

Line information

=	Display line number of a line.
l	Display control characters in ASCII.
p	Display the line.

Input/output processing

n	Skip current line and go to the next line.
r	Read another file’s contents into the output stream.
w	Write input lines to another file.
q	Quit the <i>sed</i> script (no further output).

Yanking and putting

h	Copy into hold space; wipe out what’s there.
H	Copy into hold space; append to what’s there.
g	Get the hold space back; wipe out the destination line.
G	Get the hold space back; append to the pattern space.
x	Exchange contents of the hold and pattern spaces.

Branching commands

- b Branch to *label* or to end of script.
- t Same as b, but branch only after substitution.
- :*label* Label branched to by t or b.

Multiline input processing

- N Read another line of input (creates embedded newline).
- D Delete up to the embedded newline.
- P Print up to the embedded newline.

Alphabetical Summary of sed Commands

sed Command	Description
#	# Begin a comment in a <i>sed</i> script. Valid only as the first character of the first line. (Some versions allow comments anywhere, but it is better not to rely on this.) If the first line of the script is #n, <i>sed</i> behaves as if -n had been specified.
:	: <i>label</i> Label a line in the script for the transfer of control by b or t. <i>label</i> may contain up to seven characters.
=	[/ <i>pattern</i> /]= Write to standard output the line number of each line addressed by <i>pattern</i> .
a	[<i>address</i>]a\ <i>text</i> Append <i>text</i> following each line matched by <i>address</i> . If <i>text</i> goes over more than one line, newlines must be "hidden" by preceding them with a backslash. The <i>text</i> will be terminated by the first newline that is not hidden in this way. The <i>text</i> is not available in the pattern space, and subsequent commands cannot be applied to it. The results of this command are sent to standard output when the list of editing commands is finished, regardless of what happens to the current line in the pattern space.

sed Command	Description
b	<code>[address1[, address2]]b[label]</code> Unconditionally transfer control to <code>:label</code> elsewhere in script. That is, the command following the <code>label</code> is the next command applied to the current line. If no <code>label</code> is specified, control falls through to the end of the script, so no more commands are applied to the current line.
c	<code>[address1[, address2]]c\ text</code> Replace (change) the lines selected by the address(es) with <code>text</code> . (See a for details on <code>text</code> .) When a range of lines is specified, all lines are replaced as a group by a single copy of <code>text</code> . The contents of the pattern space are, in effect, deleted and no subsequent editing commands can be applied to the pattern space (or to <code>text</code>).
d	<code>[address1[, address2]]d</code> Delete the addressed line (or lines) from the pattern space. Thus, the line is not passed to standard output. A new line of input is read, and editing resumes with the first command in the script.
D	<code>[address1[, address2]]D</code> Delete the first part (up to embedded newline) of multi-line pattern space created by <code>N</code> command and resume editing with first command in script. If this command empties the pattern space, a new line of input is read, as if the <code>d</code> command had been executed.
g	<code>[address1[, address2]]g</code> Paste the contents of the hold space (see <code>h</code> and <code>H</code>) back into the pattern space, wiping out the previous contents of the pattern space.
G	<code>[address1[, address2]]G</code> Same as <code>g</code> , except that a newline and the hold space are pasted to the end of the pattern space instead of overwriting it.
h	<code>[address1[, address2]]h</code> Copy the pattern space into the hold space, a special temporary buffer. The previous contents of the hold space are obliterated. You can use <code>h</code> to save a line before editing it.
H	<code>[address1[, address2]]H</code> Append a newline and then the contents of the pattern space to the contents of the hold space. Even if the hold space is empty, <code>H</code> still appends a newline. <code>H</code> is like an incremental copy.
i	<code>[address]i\ text</code> Insert <code>text</code> before each line matched by <code>address</code> . (See <code>a</code> for details on <code>text</code> .)

sed Command	Description				
l	<code>[address1[, address2]]l</code> List the contents of the pattern space, showing nonprinting characters as ASCII codes. Long lines are wrapped.				
n	<code>[address1[, address2]]n</code> Read the next line of input into the pattern space. The current line is sent to standard output, and the next line becomes the current line. Control passes to the command following <code>n</code> instead of resuming at the top of the script.				
N	<code>[address1[, address2]]N</code> Append the next input line to contents of the pattern space; the new line is separated from the previous contents of the pattern space by a newline. (This command is designed to allow pattern matches across two lines.) By using <code>\n</code> to match the embedded newline, you can match patterns across multiple lines.				
p	<code>[address1[, address2]]p</code> Print the addressed line(s). Note that this can result in duplicate output unless default output is suppressed by using <code>#n</code> or the <code>-n</code> command-line option. Typically used before commands that change flow control (<code>d</code> , <code>n</code> , <code>b</code>), which might prevent the current line from being output.				
P	<code>[address1[, address2]]P</code> Print first part (up to embedded newline) of multiline pattern space created by <code>N</code> command. Same as <code>p</code> if <code>N</code> has not been applied to a line.				
q	<code>[address]q</code> Quit when <code>address</code> is encountered. The addressed line is first written to the output (if default output is not suppressed), along with any text appended to it by previous <code>a</code> or <code>r</code> commands.				
r	<code>[address]r file</code> Read contents of <code>file</code> and append after the contents of the pattern space. There must be exactly one space between the <code>r</code> and the filename.				
s	<code>[address1[, address2]]s/pat/repl/[flags]</code> Substitute <code>repl</code> for <code>pat</code> on each addressed line. If pattern addresses are used, the pattern <code>//</code> represents the last pattern address specified. Any delimiter may be used. Use <code>\</code> within <code>pat</code> or <code>repl</code> to escape the delimiter. The following flags can be specified: <table border="0"> <tr> <td>n</td> <td>Replace <i>n</i>th instance of <i>pat</i> on each addressed line. <i>n</i> is any number in the range 1 to 512; the default is 1.</td> </tr> <tr> <td>g</td> <td>Replace all instances of <i>pat</i> on each addressed line, not just the first instance.</td> </tr> </table>	n	Replace <i>n</i> th instance of <i>pat</i> on each addressed line. <i>n</i> is any number in the range 1 to 512; the default is 1.	g	Replace all instances of <i>pat</i> on each addressed line, not just the first instance.
n	Replace <i>n</i> th instance of <i>pat</i> on each addressed line. <i>n</i> is any number in the range 1 to 512; the default is 1.				
g	Replace all instances of <i>pat</i> on each addressed line, not just the first instance.				

sed	
Command	Description
s	<p>p Print the line if the substitution is successful. If several substitutions are successful, <i>sed</i> will print multiple copies of the line.</p> <p>w file Write the line to <i>file</i> if a replacement was done. A maximum of 10 different <i>files</i> can be opened.</p>
t	<p><i>[address1[,address2]]t [label]</i></p> <p>Test if successful substitutions have been made on addressed lines, and if so, branch to the line marked by <i>:label</i>. (See b and :.) If <i>label</i> is not specified, control branches to the bottom of the script. The t command is like a case statement in the C programming language or the various shell programming languages. You test each case; when it's true, you exit the construct.</p>
w	<p><i>[address1[,address2]]w file</i></p> <p>Append contents of pattern space to <i>file</i>. This action occurs when the command is encountered rather than when the pattern space is output. Exactly one space must separate the w and the filename. A maximum of 10 different files can be opened in a script. This command will create the file if it does not exist; if the file exists, its contents will be overwritten each time the script is executed. Multiple write commands that direct output to the same file append to the end of the file.</p>
x	<p><i>[address1[,address2]]x</i></p> <p>Exchange the contents of the pattern space with the contents of the hold space.</p>
y	<p><i>[address1[,address2]]y/abc/xyz/</i></p> <p>Translate characters. Change every instance of <i>a</i> to <i>x</i>, <i>b</i> to <i>y</i>, <i>c</i> to <i>z</i>, etc.</p>

The awk Programming Language

awk is a pattern-matching program for processing files, especially when they are databases. The new version of *awk*, called *nawk*, provides additional capabilities. (It really isn't so new. The additional features were added in 1984, and it was first shipped with System V Release 3.1 in 1987. Nevertheless, the name was never changed on many systems.) Every modern Unix system comes with a version of new *awk*, and its use is recommended over old *awk*. The GNU version of *awk*, called *gawk*, implements new *awk* and provides a number of additional features.

Different systems vary in what new and old *awk* are called. Some have *oawk* and *awk*, for the old and new versions, respectively. Others have *awk* and *nawk*. Still others only have *awk*, which is the new version. This example shows what happens if your *awk* is the old one:

```
$ awk 1 /dev/null
awk: syntax error near line 1
awk: bailing out near line 1
```

awk will exit silently if it is the new version.

Items described here as “common extensions” are often available in different versions of new *awk*, as well as in *gawk*, but should not be used if strict portability of your programs is important to you.

The freely available versions of *awk* described in “Additional Resources” all implement new *awk*. Thus, references in the following text such as “*nawk* only,” apply to all versions. *gawk* has additional features.

With original *awk*, you can:

- Think of a text file as made up of records and fields in a textual database
- Perform arithmetic and string operations
- Use programming constructs such as loops and conditionals
- Produce formatted reports

With *nawk*, you can also:

- Define your own functions
- Execute Unix commands from a script
- Process the results of Unix commands
- Process command-line arguments more gracefully
- Work more easily with multiple input streams
- Flush open output files and pipes (with the latest Bell Laboratories version of *awk*)

In addition, with GNU *awk* (*gawk*), you can:

- Use regular expressions to separate records, as well as fields
- Skip to the start of the next file, not just the next record
- Perform more powerful string substitutions
- Sort arrays
- Retrieve and format system time values
- Use octal and hexadecimal constants in your program
- Do bit manipulation
- Internationalize your *awk* programs, allowing strings to be translated into a local language at runtime
- Perform two-way I/O to a coprocess
- Open a two-way TCP/IP connection to a socket
- Dynamically add built-in functions
- Profile your *awk* programs

Command-Line Syntax

The syntax for invoking *awk* has two forms:

```
awk [options] 'script' var=value file(s)
awk [options] -f scriptfile var=value file(s)
```

You can specify a *script* directly on the command line, or you can store a script in a *scriptfile* and specify it with *-f*. *nawk* allows multiple *-f* scripts. Variables can be assigned a value on the command line. The value can be a string or numeric constant, a shell variable (*\$name*), or a command substitution (*`cmd`*), but the value is available only after the *BEGIN* statement is executed.

awk operates on one or more *files*. If none are specified (or if *-* is specified), *awk* reads from the standard input.

The recognized options are:

-F*fs*

Set the field separator to *fs*. This is the same as setting the built-in variable FS. Original *awk* only allows the field separator to be a single character. *nawk* allows *fs* to be a regular expression. Each input line, or record, is divided into fields by white space (spaces or TABs) or by some other user-definable field separator. Fields are referred to by the variables \$1, \$2,..., \$*n*. \$0 refers to the entire record.

-v *var=value*

Available in *nawk* only. Assign a *value* to variable *var*. This allows assignment before the script begins execution.

For example, to print the first three (colon-separated) fields of each record on separate lines:

```
awk -F: '{ print $1; print $2; print $3 }' /etc/passwd
```

Numerous examples are shown later in the “Simple pattern-procedure examples” section.

Important gawk Options

Besides the standard command-line options, *gawk* has a large number of additional options. This section lists those that are of most value in day-to-day use. Any unique abbreviation of these options is acceptable.

--dump-variables[=*file*]

When the program has finished running, print a sorted list of global variables, their types, and final values to *file*. The default is *awkvars.out*.

--gen-po

Read the *awk* program and print all strings marked as translatable to standard output in the form of a GNU

gettext Portable Object file. See “Internationalization with *gawk*” for more information.

- help
Print a usage message to standard error and exit.
- lint[=*fatal*]
Enable checking of nonportable or dubious constructs, both when the program is read, and as it runs. With an argument of *fatal*, lint warnings become fatal errors.
- non-decimal-data
Allow octal and hexadecimal data in the input to be recognized as such. This option is not recommended; use `strtonum()` in your program, instead.
- profile[=*file*]
With *gawk*, put a “prettyprinted” version of the program in *file*. Default is *awkprof.out*. With *pgawk* (see “Profiling with *pgawk*”), put the profiled listing of the program in *file*.
- posix
Turn on strict POSIX compatibility, in which all common and *gawk*-specific extensions are disabled.
- source=*'program text'*
Use *program text* as the *awk* source code. Use this option with `-f` to mix command-line programs with *awk* library files.
- traditional
Disable all *gawk*-specific extensions, but allow common extensions (e.g., the `**` operator for exponentiation).
- version
Print the version of *gawk* on standard error and exit.

Profiling with *pgawk*

When *gawk* is built and installed, a separate program named *pgawk* (*profiling gawk*) is built and installed with it. The two programs behave identically; however, *pgawk* runs more slowly since it keeps execution counts for each statement as it runs. When it is done, it automatically places an execution profile of your program in a file named *awkprof.out*. (You can change the filename with the `--profile` option.)

The execution profile is a “prettyprinted” version of your program with execution counts listed in the left margin. For example, after running this program:

```
$ pgawk '/bash$/ { nusers++ }
> END { print nusers, "users use Bash." }' /etc/passwd
16 users use Bash.
```

the execution profile looks like this:

```
# gawk profile, created Wed Nov  1 14:34:38 2000

# Rule(s)

35 /bash$/ { # 16
16          nusers++
   }

# END block(s)

END {
1          print nusers, "users use Bash."
   }
```

If sent `SIGUSR1`, *pgawk* prints the profile and an *awk* function call stack trace, and then keeps going. Multiple `SIGUSR1` signals may be sent; the profile and trace will be printed each time. This facility is useful if your *awk* program appears to be looping, and you want to see if something unexpected is being executed.

If sent `SIGHUP`, *pgawk* prints the profile and stack trace, and then exits.

Patterns and Procedures

awk scripts consist of patterns and procedures:

```
pattern { procedure }
```

Both are optional. If *pattern* is missing, { *procedure* } is applied to all lines. If { *procedure* } is missing, the matched line is printed.

Patterns

A pattern can be any of the following:

```
/regular expression/  
relational expression  
pattern-matching expression  
BEGIN  
END
```

- Expressions can be composed of quoted strings, numbers, operators, function calls, user-defined variables, or any of the predefined variables described later in “Built-in Variables.”
- Regular expressions use the extended set of metacharacters and are described earlier in “Matching Text.”
- The ^ and \$ metacharacters refer to the beginning and end of a string (such as the fields), respectively, rather than the beginning and end of a line. In particular, these metacharacters will *not* match at a newline embedded in the middle of a string.
- Relational expressions use the relational operators listed in the section “Operators” later in this book. For example, \$2 > \$1 selects lines for which the second field is greater than the first. Comparisons can be either string or numeric. Thus, depending on the types of data in \$1 and \$2, *awk* will do either a numeric or a string comparison. This can change from one record to the next.
- Pattern-matching expressions use the operators ~ (match) and !~ (don’t match). See “Operators” later in this book.

- The BEGIN pattern lets you specify procedures that will take place *before* the first input line is processed. (Generally, you process the command line and set global variables here.)
- The END pattern lets you specify procedures that will take place *after* the last input record is read.
- In *nawk*, BEGIN and END patterns may appear multiple times. The procedures are merged as if there had been one large procedure.

Except for BEGIN and END, patterns can be combined with the Boolean operators `||` (or), `&&` (and), and `!` (not). A range of lines can also be specified using comma-separated patterns:

pattern,pattern

Procedures

Procedures consist of one or more commands, function calls, or variable assignments, separated by newlines or semicolons, and are contained within curly braces. Commands fall into five groups:

- Variable or array assignments
- Input/output commands
- Built-in functions
- Control-flow commands
- User-defined functions (*nawk* only)

Simple pattern-procedure examples

Print first field of each line:

```
{ print $1 }
```

Print all lines that contain *pattern*:

```
/pattern/
```

Print first field of lines that contain *pattern*:

```
/pattern/ { print $1 }
```

Select records containing more than two fields:

```
NF > 2
```

Interpret input records as a group of lines up to a blank line. Each line is a single field:

```
BEGIN { FS = "\n"; RS = "" }
```

Print fields 2 and 3 in switched order, but only on lines whose first field matches the string URGENT:

```
$1 ~ /URGENT/ { print $3, $2 }
```

Count and print the number of *pattern* found:

```
/pattern/ { ++x }  
END { print x }
```

Add numbers in second column and print total:

```
{ total += $2 }  
END { print "column total is", total }
```

Print lines that contain less than 20 characters:

```
length($0) < 20
```

Print each line that begins with Name: and that contains exactly seven fields:

```
NF == 7 && /^Name: /
```

Print the fields of each record in reverse order, one per line:

```
{  
    for (i = NF; i >= 1; i--)  
        print $i  
}
```

Built-in Variables

All *awk* variables are included in *nawk*. All *nawk* variables are included in *gawk*.

Version	Variable	Description
<i>awk</i>	FILENAME	Current filename.
	FS	Field separator (a space).

Version	Variable	Description
<i>awk</i>	NF	Number of fields in current record.
	NR	Number of the current record.
	OFMT	Output format for numbers ("% . 6g") and for conversion to string.
	OFS	Output field separator (a space).
	ORS	Output record separator (a newline).
	RS	Record separator (a newline).
	\$0	Entire input record.
	\$n	nth field in current record; fields are separated by FS.
<i>nawk</i>	ARGC	Number of arguments on the command line.
	ARGV	An array containing the command-line arguments, indexed from 0 to ARGC - 1.
	CONVFMT	String conversion format for numbers ("% . 6g"). (POSIX)
	ENVIRON	An associative array of environment variables.
	FNR	Like NR, but relative to the current file.
	RLENGTH	Length of the string matched by match() function.
	RSTART	First position in the string matched by match() function.
<i>gawk</i>	SUBSEP	Separator character for array subscripts ("\034").
	ARGIND	Index in ARGV of current input file.
	BINMODE	Controls binary I/O for input and output files. Use values of 1, 2, or 3 for input, output, or both kinds of files, respectively. Set on the command line to affect standard input, standard output, and standard error.
	ERRNO	A string indicating the error when a redirection fails for getline or if close() fails.
	FIELDWIDTHS	A space-separated list of field widths to use for splitting up the record, instead of FS.
	IGNORECASE	When true, all regular expression matches, string comparisons, and index() ignore case.
	LINT	Dynamically controls production of "lint" warnings. With a value of "fatal", lint warnings become fatal errors.

Version	Variable	Description
<i>gawk</i>	PROCINFO	An array containing information about the process, such as real and effective UID numbers, process ID number, and so on.
	RT	The text matched by RS, which can be a regular expression in <i>gawk</i> .
	TEXTDOMAIN	The text domain (application name) for internationalized messages ("messages").

Operators

The following table lists the operators, in order of increasing precedence, that are available in *awk*:

Symbol	Meaning
= += -= *= /= %= ^= **=	Assignment. ^a
?:	C conditional expression (<i>nawk</i> only).
	Logical OR (short-circuit).
&&	Logical AND (short-circuit).
in	Array membership (<i>nawk</i> only).
~ !~	Match regular expression and negation.
< <= > >= != ==	Relational operators.
(blank)	Concatenation.
+ -	Addition, subtraction.
* / %	Multiplication, division, and modulus (remainder).
+ - !	Unary plus and minus, and logical negation.
^ **	Exponentiation. ^a
++ --	Increment and decrement, either prefix or postfix.
\$	Field reference.

^a While ****** and ****=** are common extensions, they are not part of POSIX *awk*.

Variable and Array Assignment

Variables can be assigned a value with an = sign. For example:

```
FS = ","
```

Expressions using the operators listed in the previous table can be assigned to variables.

Arrays can be created with the `split()` function (described later), or they can simply be named in an assignment statement. Array elements can be subscripted with numbers (`array[1]`, ..., `array[n]`) or with strings. Arrays subscripted by strings are called *associative arrays*. (In fact, all arrays in *awk* are associative; numeric subscripts are converted to strings before using them as array subscripts. Associative arrays are one of *awk*'s most powerful features.)

For example, to count the number of widgets you have, you could use the following script:

```
/widget/ { count["widget"]++ }      Count widgets
END      { print count["widget"] }  Print the count
```

You can use the special for loop to read all the elements of an associative array:

```
for (item in array)
    process array[item]
```

The index of the array is available as `item`, while the value of an element of the array can be referenced as `array[item]`.

You can use the operator `in` to test that an element exists by testing to see if its index exists (*nawk* only). For example:

```
if (index in array)
    ...
```

tests that `array[index]` exists, but you cannot use it to test the value of the element referenced by `array[index]`.

You can also delete individual elements of the array using the `delete` statement (*nawk* only).

Escape sequences

Within string and regular expression constants, the following escape sequences may be used:

Sequence	Meaning
<code>\a</code>	Alert (bell)
<code>\b</code>	Backspace
<code>\f</code>	Form feed
<code>\n</code>	Newline
<code>\r</code>	Carriage return
<code>\t</code>	TAB
<code>\v</code>	Vertical tab
<code>\\</code>	Literal backslash
<code>\nnn</code>	Octal value <i>nnn</i>
<code>\xnn</code>	Hexadecimal value <i>nn</i>
<code>\"</code>	Literal double quote (in strings)
<code>\/</code>	Literal slash (in regular expressions)

NOTE

The `\x` escape sequence is a common extension, but it is not part of POSIX *awk*.

Octal and Hexadecimal Constants in *gawk*

gawk allows you to use octal and hexadecimal constants in your program source code. The form is as in C: octal constants start with a leading 0, and hexadecimal constants with a leading 0x or 0X. The hexadecimal digits a–f may be in either upper- or lowercase.

```
$ gawk 'BEGIN { print 042, 42, 0x42 }'  
34 42 66
```

Use the `strtonum()` function to convert octal or hexadecimal input data into numerical values.

User-Defined Functions

awk allows you to define your own functions. This makes it easy to encapsulate sequences of steps that need to be repeated into a single place, and re-use the code from anywhere in your program.

The following function capitalizes each word in a string. It has one parameter, named `input`, and five local variables, which are written as extra parameters:

```
# capitalize each word in a string
function capitalize(input, result, words, n, i, w)
{
    result = ""
    n = split(input, words, " ")
    for (i = 1; i <= n; i++) {
        w = words[i]
        w = toupper(substr(w, 1, 1)) substr(w, 2)
        if (i > 1)
            result = result " "
        result = result w
    }
    return result
}

# main program, for testing
{ print capitalize($0) }
```

With this input data:

```
A test line with words and numbers like 12 on it.
```

this program produces:

```
A Test Line With Words And Numbers Like 12 On It.
```

NOTE

For user-defined functions, no space is allowed between the function name and the left parenthesis when the function is called.

Group Listing of awk Functions and Commands

awk functions and commands may be classified as in the following table. For descriptions and examples of how to use these commands, see “Alphabetical Summary of awk Functions and Commands.”

Function type	All awk versions	nawk	gawk
Arithmetic	exp	atan2	
	int	cos	
	log	rand	
	sqrt	sin	
			srand
String	index	gsub	asort
	length	match	gensub
	split	sub	strtonum
	sprintf	tolower	
		toupper	
Control flow	break	do/while	
	continue	return	
	exit		
	for		
	if/else		
	while		
Input/output	next	close	fflush ^a
Processing	print	getline	nextfile ^a
	printf		
Programming		delete	extension
		function	
		system	

^a Also in Bell Labs *awk*.

The following functions are specific to *gawk*:

Function type	Functions		
Bit manipulation	and	lshift	rshift
	compl	or	xor
Time	mktime	strftime	systeme
Translation	bindtextdomain	dcgettext	dcngettext

Coprocesses and Sockets with *gawk*

gawk allows you to open a two-way pipe to another process, called a *coprocess*. This is done with the `|&` operator used with `getline` and `print` or `printf`.

```
print database command |& "db_server"  
"db_server" |& getline response
```

If the *command* used with `|&` is a filename beginning with `/inet/`, *gawk* opens a TCP/IP connection. The filename should be of the following form:

```
/inet/protocol/lport/hostname/rport
```

The parts of the filename are:

protocol

One of `tcp`, `udp`, or `raw`, for TCP, UDP, or raw IP sockets, respectively. Note: `raw` is currently reserved but unsupported.

lport

The local TCP or UPD port number to use. Use 0 to let the operating system pick a port.

hostname

The name or IP address of the remote host to connect to.

rport

The port (application) on the remote host to connect to. A service name (e.g., `tftp`) is looked up using the `C getservbyname()` function.

Implementation Limits

Many versions of *awk* have various implementation limits, on things such as:

- Number of fields per record
- Number of characters per input record
- Number of characters per output record
- Number of characters per field
- Number of characters per `printf` string
- Number of characters in literal string
- Number of characters in character class
- Number of files open
- Number of pipes open
- The ability to handle 8-bit characters and characters that are all zero (ASCII NUL)

gawk does not have limits on any of the above items, other than those imposed by the machine architecture and/or the operating system.

Alphabetical Summary of *awk* Functions and Commands

The following alphabetical list of keywords and functions includes all that are available in *awk* and *nawk*. *nawk* includes all old *awk* functions and keywords, plus some additional ones (marked as ①). Extensions that aren't part of POSIX *awk* but that are in both *gawk* and the Bell Laboratories *awk* are marked as ②. Cases where *gawk* has extensions are marked as ③. Items that aren't marked with a symbol are available in all versions.

Command	Description
<code>and</code>	<code>and(<i>expr1</i>, <i>expr2</i>)</code> ③ Return the bitwise AND of <i>expr1</i> and <i>expr2</i> , which should be values that fit in a C unsigned long.

Command	Description
asort	<code>asort(<i>src</i> [, <i>dest</i>])</code> ⑥ Sort the array <i>src</i> , destructively replacing the indexes with values from one to the number of elements in the array. If <i>dest</i> is supplied, copy <i>src</i> to <i>dest</i> and sort <i>dest</i> , leaving <i>src</i> unchanged. Returns the number of elements in <i>src</i> .
atan2	<code>atan2(<i>y</i>, <i>x</i>)</code> ⑥ Return the arctangent of <i>y/x</i> in radians.
bindtextdomain	<code>bindtextdomain(<i>dir</i> [, <i>domain</i>])</code> ⑥ Look in directory <i>dir</i> for message translation files for text domain <i>domain</i> (default: value of TEXTDOMAIN). Returns the directory where <i>domain</i> is bound.
break	<code>break</code> Exit from a <code>while</code> , <code>for</code> , or <code>do</code> loop.
close	<code>close(<i>expr</i>)</code> ⑥ <code>close(<i>expr</i>, <i>how</i>)</code> ⑥ In most implementations of <i>awk</i> , you can only have up to ten files open simultaneously and one pipe. Therefore, <i>nawk</i> provides a <code>close()</code> function that allows you to close a file or a pipe. It takes the same expression that opened the pipe or file as an argument. This expression must be identical, character by character, to the one that opened the file or pipe—even whitespace is significant. In the second form, close one end of either a TCP/IP socket or a two-way pipe to a coprocess. <i>how</i> is a string, either "from" or "to". Case does not matter.
compl	<code>compl(<i>expr</i>)</code> ⑥ Return the bitwise complement of <i>expr</i> , which should be a value that fits in a C unsigned long.
continue	<code>continue</code> Begin next iteration of <code>while</code> , <code>for</code> , or <code>do</code> loop.
cos	<code>cos(<i>x</i>)</code> ⑥ Return the cosine of <i>x</i> , an angle in radians.
dcgettext	<code>dcgettext(<i>str</i> [, <i>dom</i> [, <i>cat</i>]])</code> ⑥ Return the translation of <i>str</i> for the text domain <i>dom</i> in message category <i>cat</i> . Default text domain is value of TEXTDOMAIN. Default category is "LC_MESSAGES".
dcngettext	<code>dcngettext(<i>str1</i>, <i>str2</i>, <i>num</i> [, <i>dom</i> [, <i>cat</i>]])</code> ⑥ If <i>num</i> is one, return the translation of <i>str1</i> for the text domain <i>dom</i> in message category <i>cat</i> . Otherwise return the translation of <i>str2</i> . Default text domain is value of TEXTDOMAIN. Default category is "LC_MESSAGES". For <i>gawk</i> 3.1.1 and later.

Command	Description
delete	<p><code>delete array[element]</code> ⑧</p> <p><code>delete array</code> ⑥</p> <p>Delete <i>element</i> from <i>array</i>. The brackets are typed literally. The second form is a common extension, which deletes <i>all</i> elements of the array in one shot.</p>
do	<p><code>do</code></p> <p style="padding-left: 2em;"><i>statement</i></p> <p><code>while (expr)</code> ⑧</p> <p>Looping statement. Execute <i>statement</i>, then evaluate <i>expr</i> and if true, execute <i>statement</i> again. A series of statements must be put within braces.</p>
exit	<p><code>exit [expr]</code></p> <p>Exit from script, reading no new input. The END procedure, if it exists, will be executed. An optional <i>expr</i> becomes <i>awk</i>'s return value.</p>
exp	<p><code>exp(x)</code></p> <p>Return exponential of <i>x</i> (<i>e^x</i>).</p>
extension	<p><code>extension(lib, init)</code> ⑥</p> <p>Dynamically load the shared object file <i>lib</i>, calling the function <i>init</i> to initialize it. Return the value returned by the <i>init</i> function. This function allows you to add new built-in functions to <i>gawk</i>. See <i>Effective awk Programming</i>, Third Edition, for the details.</p>
fflush	<p><code>fflush([output-expr])</code> ⑥</p> <p>Flush any buffers associated with open output file or pipe <i>output-expr</i>.</p> <p><i>gawk</i> extends this function. If no <i>output-expr</i> is supplied, it flushes standard output. If <i>output-expr</i> is the null string (""), it flushes all open files and pipes.</p>
for	<p><code>for (init-expr; test-expr; incr-expr)</code></p> <p style="padding-left: 2em;"><i>statement</i></p> <p>C-style looping construct. <i>init-expr</i> assigns the initial value of a counter variable. <i>test-expr</i> is a relational expression that is evaluated each time before executing the <i>statement</i>. When <i>test-expr</i> is false, the loop is exited. <i>incr-expr</i> is used to increment the counter variable after each pass. All of the expressions are optional. A missing <i>test-expr</i> is considered to be true. A series of statements must be put within braces.</p>

Command	Description
for	<pre>for (item in array) statement</pre> <p>Special loop designed for reading associative arrays. For each element of the array, the <i>statement</i> is executed; the element can be referenced by <i>array[item]</i>. A series of statements must be put within braces.</p>
function	<pre>function name(parameter-list) { statements }</pre> <p>Create <i>name</i> as a user-defined function consisting of <i>awk statements</i> that apply to the specified list of parameters. No space is allowed between <i>name</i> and the left parenthesis when the function is called.</p>
gensub	<pre>gensub(regex, str, how[, target])</pre> <p>General substitution function. Substitute <i>str</i> for matches of the regular expression <i>regex</i> in the string <i>target</i>. If <i>how</i> is a number, replace the <i>how</i>th match. If it is <i>g</i> or <i>G</i>, substitute globally. If <i>target</i> is not supplied, <i>\$0</i> is used. Return the new string value. The original <i>target</i> is <i>not</i> modified. (Compare with gsub and sub.)</p>
getline	<pre>getline getline [var] [< file] command getline [var] command & getline [var]</pre> <p>Read next line of input. Original <i>awk</i> does not support the syntax to open multiple input streams or use a variable.</p> <p>The second form reads input from <i>file</i> and the third form reads the output of <i>command</i>. All forms read one record at a time, and each time the statement is executed it gets the next record of input. The record is assigned to <i>\$0</i> and is parsed into fields, setting <i>NF</i>, <i>NR</i>, and <i>FNR</i>. If <i>var</i> is specified, the result is assigned to <i>var</i> and <i>\$0</i> and <i>NF</i> are not changed. Thus, if the result is assigned to a variable, the current record does not change. <i>getline</i> is actually a function and it returns 1 if it reads a record successfully, 0 if end-of-file is encountered, and -1 if for some reason it is otherwise unsuccessful.</p> <p>The fourth form reads the output from coprocess <i>command</i>. See “Coproceses and Sockets with <i>gawk</i>” for more information.</p>
gsub	<pre>gsub(regex, str [, target])</pre> <p>Globally substitute <i>str</i> for each match of the regular expression <i>regex</i> in the string <i>target</i>. If <i>target</i> is not supplied, defaults to <i>\$0</i>. Return the number of substitutions.</p>

Command	Description
if	<pre>if (condition) statement1 [else statement2]</pre> <p>If <i>condition</i> is true, do <i>statement1</i>, otherwise do <i>statement2</i> in optional <i>else</i> clause. The <i>condition</i> can be an expression using any of the relational operators <code><</code>, <code><=</code>, <code>=</code>, <code>!=</code>, <code>>=</code>, or <code>></code>, as well as the array membership operator <code>in</code>, and the pattern-matching operators <code>~</code> and <code>!~</code> (e.g., <code>if (\$1 ~ /[Aa].*/)</code>). A series of statements must be put within braces. Another <code>if</code> can directly follow an <code>else</code> in order to produce a chain of tests or decisions.</p>
index	<pre>index(str, substr)</pre> <p>Return the position (starting at 1) of <i>substr</i> in <i>str</i>, or zero if <i>substr</i> is not present in <i>str</i>.</p>
int	<pre>int(x)</pre> <p>Return integer value of <i>x</i> by truncating any fractional part.</p>
length	<pre>length([arg])</pre> <p>Return length of <i>arg</i>, or the length of <code>\$0</code> if no argument.</p>
log	<pre>log(x)</pre> <p>Return the natural logarithm (base <i>e</i>) of <i>x</i>.</p>
lshift	<pre>lshift(expr, count) ⑥</pre> <p>Return the result of shifting <i>expr</i> left by <i>count</i> bits. Both <i>expr</i> and <i>count</i> should be values that fit in a C unsigned long.</p>
match	<pre>match(str, regex) ⑧ match(str, regex[, array]) ⑥</pre> <p>Function that matches the pattern, specified by the regular expression <i>regex</i>, in the string <i>str</i> and returns either the position in <i>str</i> where the match begins, or 0 if no occurrences are found. Sets the values of <code>RSTART</code> and <code>RLENGTH</code> to the start and length of the match, respectively.</p> <p>If <i>array</i> is provided, <i>gawk</i> puts the text that matched the entire regular expression in <code>array[0]</code>, the text that matched the first parenthesized subexpression in <code>array[1]</code>, the second in <code>array[2]</code>, and so on.</p>
mktime	<pre>mktime(timespec) ⑥</pre> <p>Turns <i>timespec</i> (a string of the form "<code>YYYY MM DD HH MM SS [DST]</code>") representing a local time) into a time-of-day value in seconds since midnight, January 1, 1970, UTC.</p>

Command	Description
next	<code>next</code> Read next input line and start new cycle through pattern/procedures statements.
nextfile	<code>nextfile</code> ⑤ Stop processing the current input file and start new cycle through pattern/procedures statements, beginning with the first record of the next file.
or	<code>or(expr1, expr2)</code> ⑥ Return the bitwise OR of <i>expr1</i> and <i>expr2</i> , which should be values that fit in a C unsigned long.
print	<code>print [output-expr[, ...]] [dest-expr]</code> Evaluate the <i>output-expr</i> and direct it to standard output followed by the value of ORS. Each comma-separated <i>output-expr</i> is separated in the output by the value of OFS. With no <i>output-expr</i> , print \$0. The output may be redirected to a file or pipe via the <i>dest-expr</i> , which is described in the section “Output redirections” following this table.
printf	<code>printf(format [, expr-list]) [dest-expr]</code> An alternative output statement borrowed from the C language. It has the ability to produce formatted output. It can also be used to output data without automatically producing a newline. <i>format</i> is a string of format specifications and constants. <i>expr-list</i> is a list of arguments corresponding to format specifiers. As for <code>print</code> , output may be redirected to a file or pipe. See the section “printf formats” following this table for a description of allowed format specifiers.
rand	<code>rand()</code> ⑦ Generate a random number between 0 and 1. This function returns the same series of numbers each time the script is executed, unless the random number generator is seeded using <code>srand()</code> .
return	<code>return [expr]</code> ⑧ Used within a user-defined function to exit the function, returning value of expression. The return value of a function is undefined if <i>expr</i> is not provided.
rshift	<code>rshift(expr, count)</code> ⑥ Return the result of shifting <i>expr</i> right by <i>count</i> bits. Both <i>expr</i> and <i>count</i> should be values that fit in a C unsigned long.
sin	<code>sin(x)</code> ⑦ Return the sine of <i>x</i> , an angle in radians.

Command	Description
split	<code>split(string, array[, sep])</code> Split <i>string</i> into elements of array <i>array[1], ..., array[n]</i> . Return the number of array elements created. The string is split at each occurrence of separator <i>sep</i> . If <i>sep</i> is not specified, FS is used.
sprintf	<code>sprintf(format[, expressions])</code> Return the formatted value of one or more <i>expressions</i> , using the specified <i>format</i> . Data is formatted but not printed. See the section “printf formats” following this table for a description of allowed format specifiers.
sqrt	<code>sqrt(arg)</code> Return square root of <i>arg</i> .
srand	<code>srand([expr])</code> ⑨ Use optional <i>expr</i> to set a new seed for the random number generator. Default is the time of day. Return value is the old seed.
strftime	<code>strftime([format[, timestamp]])</code> ⑥ Format <i>timestamp</i> according to <i>format</i> . Return the formatted string. The <i>timestamp</i> is a time-of-day value in seconds since midnight, January 1, 1970, UTC. The <i>format</i> string is similar to that of <code>sprintf</code> . If <i>timestamp</i> is omitted, it defaults to the current time. If <i>format</i> is omitted, it defaults to a value that produces output similar to that of the Unix <i>date</i> command.
strtonum	<code>strtonum(expr)</code> ⑥ Return the numeric value of <i>expr</i> , which is a string representing an octal, decimal, or hexadecimal number in the usual C notations. Use this function for processing nondecimal input data.
sub	<code>sub(regex, str[, target])</code> ⑨ Substitute <i>str</i> for first match of the regular expression <i>regex</i> in the string <i>target</i> . If <i>target</i> is not supplied, defaults to \$0. Return 1 if successful; 0 otherwise.
substr	<code>substr(string, beg[, len])</code> Return substring of <i>string</i> at beginning position <i>beg</i> (counting from 1), and the characters that follow to maximum specified length <i>len</i> . If no length is given, use the rest of the string.

Command	Description
system	<p><code>system(command)</code> ④</p> <p>Function that executes the specified <i>command</i> and returns its exit status. The status of the executed command typically indicates success or failure. A value of 0 means that the command executed successfully. A nonzero value indicates a failure of some sort. The documentation for the command you're running will give you the details.</p> <p>The output of the command is <i>not</i> available for processing within the <i>awk</i> script. Use <code>command getline</code> to read the output of a command into the script.</p>
systemtime	<p><code>systemtime()</code> ④</p> <p>Return a time-of-day value in seconds since midnight, January 1, 1970, UTC.</p>
tolower	<p><code>tolower(str)</code> ④</p> <p>Translate all uppercase characters in <i>str</i> to lowercase and return the new string.^a</p>
toupper	<p><code>toupper(str)</code> ④</p> <p>Translate all lowercase characters in <i>str</i> to uppercase and return the new string.^a</p>
while	<p><code>while (condition)</code> <i>statement</i></p> <p>Do <i>statement</i> while <i>condition</i> is true (see if for a description of allowable conditions). A series of statements must be put within braces.</p>
xor	<p><code>xor(expr1, expr2)</code> ④</p> <p>Return the bitwise XOR of <i>expr1</i> and <i>expr2</i>, which should be values that fit in a C unsigned long.</p>

^a Very early versions of *nawk* don't support `tolower()` and `toupper()`. However, they are now part of the POSIX specification for *awk*.

Output redirections

For `print` and `printf`, *dest-expr* is an optional expression that directs the output to a file or pipe.

> *file*

Directs the output to a file, overwriting its previous contents.

>> *file*

Appends the output to a file, preserving its previous contents. In both of these cases, the file will be created if it does not already exist.

| *command*

Directs the output as the input to a system command.

|\& *command*

Directs the output as the input to a coprocess. *gawk* only.

Be careful not to mix > and >> for the same file. Once a file has been opened with >, subsequent output statements continue to append to the file until it is closed.

Remember to call `close()` when you have finished with a file, pipe, or coprocess. If you don't, eventually you will hit the system limit on the number of simultaneously open files.

printf formats

Format specifiers for `printf` and `sprintf` have the following form:

`%[posn$][flag][width][.precision]letter`

The control *letter* is required. The format conversion control letters are given in the following table:

Character	Description
c	ASCII character.
d	Decimal integer.
i	Decimal integer. (Added in POSIX)
e	Floating-point format (<code>([-]d.precision[+-]dd)</code>).
E	Floating-point format (<code>([-]d.precisionE[+-]dd)</code>).
f	Floating-point format (<code>([-]ddd.precision)</code>).
g	e or f conversion, whichever is shortest, with trailing zeros removed.
G	E or f conversion, whichever is shortest, with trailing zeros removed.

Character	Description
o	Unsigned octal value.
s	String.
u	Unsigned decimal value.
x	Unsigned hexadecimal number. Uses a–f for 10 to 15.
X	Unsigned hexadecimal number. Uses A–F for 10 to 15.
%	Literal%.

gawk allows you to provide a *positional specifier* after the % (*posn\$*). A positional specifier is an integer count followed by a \$. The count indicates which argument to use at that point. Counts start at one, and don't include the format string. This feature is primarily for use in producing translations of format strings. For example:

```
$ gawk 'BEGIN { printf "%2$s, %1$s\n", "world", "hello" }'
hello, world
```

The optional *flag* is one of the following:

Character	Description
-	Left-justify the formatted value within the field.
<i>space</i>	Prefix positive values with a space and negative values with a minus.
+	Always prefix numeric values with a sign, even if the value is positive.
#	Use an alternate form: %o has a preceding 0 %x and %X are prefixed with 0x and 0X, respectively %e, %E, and %f always have a decimal point in the result %g and %G do not have trailing zeros removed
0	Pad output with zeros, not spaces. This only happens when the field width is wider than the converted result. This flag applies to all output formats, even non-numeric ones.

The optional *width* is the minimum number of characters to output. The result will be padded to this size if it is smaller. The 0 flag causes padding with zeros; otherwise, padding is with spaces.

The *precision* is optional. Its meaning varies by control letter, as shown in this table:

Conversion	Precision means
%d,%i,%o,%u,%x,%X	The minimum number of digits to print.
%e,%E,%f	The number of digits to the right of the decimal point.
%g,%G	The maximum number of significant digits.
%s	The maximum number of characters to print.

Internationalization with gawk

You can *internationalize* your programs if you use *gawk*. This consists of choosing a text domain for your program, marking strings that are to be translated and, if necessary, using the `bindtextdomain()`, `dcgettext()`, and `dcngettext()` functions.

Localizing your program consists of extracting the marked strings, creating translations, and compiling and installing the translations in the proper place. Full details are given in *Effective awk Programming*, Third Edition.

The internationalization features in *gawk* use GNU *gettext*. You may need to install these tools to create translations if your system doesn't already have them. Here is a very brief outline of the steps involved:

1. Set `TEXTDOMAIN` to your text domain in a `BEGIN` block:

```
BEGIN { TEXTDOMAIN = "whizprog" }
```

2. Mark all strings to be translated by prepending a leading underscore:

```
printf(_"whizprog: can't open /dev/telepath (%s)\n",  
      dcgettext(ERRNO)) > "/dev/stderr"
```

3. Extract the strings with the `--gen-po` option:

```
$ gawk --gen-po -f whizprog.awk > whizprog.pot
```

4. Copy the file for translating, and make the translations:

```
$ cp whizprog.pot esperanto.po  
$ ed esperanto.po
```

5. Use the *msgfmt* program from GNU *gettext* to compile the translations. The binary format allows fast lookup of the translations at runtime. The default output is a file named *messages*.

```
$ msgfmt esperanto.po
$ mv messages esperanto.mo
```

6. Install the file in the standard location. This is usually done at program installation. The location can vary from system to system.

That's it! *gawk* will automatically find and use the translated messages, if they exist.

Additional Resources

This section lists resources for further exploration.

Source Code

The following URLs indicate where to get source code for GNU *sed*, four freely available versions of *awk*, and GNU *gettext*.

<ftp://ftp.gnu.org/gnu/sed/sed-3.02.tar.gz>

The Free Software Foundation's version of *sed*. The somewhat older version, 2.05, is also available.

<http://cm.bell-labs.com/~bwk>

Brian Kernighan's home page, with links to the source code for the latest version of *awk* from Bell Laboratories.

<ftp://ftp.whidbey.net/pub/brennan/mawk1.3.3.tar.gz>

Michael Brennan's *mawk*. A very fast, very robust version of *awk*.

<ftp://ftp.gnu.org/gnu/gawk/gawk-3.1.1.tar.gz>

The Free Software Foundation's version of *awk*, called *gawk*.

<http://awka.sourceforge.net>

The home page for *awka*, a translator that turns *awk* programs into C, compiles the generated C, and then links the object code with a library that performs the core *awk* functions.

<ftp://ftp.gnu.org/gnu/gettext/gettext-0.11.2.tar.gz>

The source code for GNU *gettext*. Get this if you need to produce translations for your *awk* programs that use *gawk*.

Books

1. Dale Dougherty, Arnold Robbins, *sed & awk*, Second Edition (Sebastopol, Calif.: O'Reilly & Associates, 1997).
2. Alfred V. Aho, Brian W. Kernighan, Peter J. Weinberger, *The AWK Programming Language* (Reading, Mass.: Addison Wesley, 1988).
3. Arnold Robbins, *Effective awk Programming*, Third Edition (Sebastopol, Calif.: O'Reilly & Associates, 2001).
4. Brian W. Kernighan, Rob Pike, *The Unix Programming Environment* (Englewood Cliffs, N.J.: Prentice-Hall, 1984).
5. Arnold Robbins, *Unix In A Nutshell*, Third Edition (Sebastopol, Calif.: O'Reilly & Associates, 1999).
6. Jon Bentley, *Programming Pearls*, Second Edition (Reading, Mass.: Addison Wesley, 2000).
7. Jon Louis Bentley, *More Programming Pearls: Confessions of a Coder* (Reading, Mass.: Addison Wesley, 1988).