

# AWK (Aho, Kernighan, and Weinberger) Summary

## Predefined Variable Summary:

Variable	Description	Support:		
		AWK	NAWK	GAWK
<b>FS</b>	Input Field Separator, a space by default.	√	√	√
<b>OFS</b>	Output Field Separator, a space by default.	√	√	√
<b>NF</b>	The Number of Fields in the current input record.	√	√	√
<b>NR</b>	The total Number of input Records seen so far.	√	√	√
<b>RS</b>	Record Separator, a newline by default.	√	√	√
<b>ORS</b>	Output Record Separator, a newline by default.	√	√	√
<b>FILENAME</b>	The name of the current input file. If no files are specified on the command line, the value of <b>FILENAME</b> is "-". However, <b>FILENAME</b> is undefined inside the <b>BEGIN</b> block (unless set by <i>getline</i> ).	√	√	√
<b>ARGC</b>	The number of command line arguments (does not include options to gawk, or the program source). Dynamically changing the contents of <b>ARGV</b> control the files used for data.	–	√	√
<b>ARGV</b>	Array of command line arguments. The array is indexed from 0 to <b>ARGC</b> - 1.	–	√	√
<b>ARGIND</b>	The index in <b>ARGV</b> of the current file being processed.	–	–	√
<b>BINMODE</b>	On non-POSIX systems, specifies use of "binary" mode for all file I/O. Numeric values of 1, 2, or 3, specify that input files, output files, or all files, respectively, should use binary I/O. String values of "r", or "w" specify that input files, or output files, respectively, should use binary I/O. String values of "rw" or "wr" specify that all files should use binary I/O. Any other string value is treated as "rw", but generates a warning message.	–	–	√
<b>CONVFMT</b>	The <b>CONVFMT</b> variable is used to specify the format when converting a number to a string. Default: "%.6g"	–	–	√
<b>ENVIRON</b>	An array containing the values of the current environment.	–	–	√
<b>ERRNO</b>	If a system error occurs either doing a redirection for <i>getline</i> , during a read for <i>getline</i> , or during a <i>close()</i> , then <b>ERRNO</b> will contain a string describing the error. The value is subject to translation in non-English locales.	–	–	√
<b>FIELDWIDTHS</b>	A white-space separated list of fieldwidths. When set, <i>gawk</i> parses the input into fields of fixed width, instead of using the value of the <b>FS</b> variable as the field separator.	–	–	√
<b>FNR</b>	Contains number of lines read, but is reset for each file read.	–	√	√
<b>IGNORECASE</b>	Controls the case-sensitivity of all regular expression and string operations. If <b>IGNORECASE</b> has a non-zero value, then string comparisons and pattern matching in rules, field splitting with <b>FS</b> , record separating with <b>RS</b> , regular expression matching with <i>~</i> and <i>!~</i> , and the <i>gensub()</i> , <i>gsub()</i> , <i>index()</i> , <i>match()</i> , <i>split()</i> , and <i>sub()</i> built-in functions all ignore case when doing regular expression operations. NOTE: Array subscripting is not affected. However, the <i>asorti()</i> and <i>asort()</i> functions are affected.	–	–	√
<b>LINT</b>	Provides dynamic control of the <i>--lint</i> option from within an AWK program. When true, <i>gawk</i> prints lint warnings.	–	–	√
<b>OFMT</b>	The default output format for numbers. Default: "%.6g"	–	√	√
<b>PROCINFO</b>	The elements of this array provide access to information about the running AWK program. PROCINFO["egid"] the value of the getegid(2) system call. PROCINFO["euid"] the value of the geteuid(2) system call. PROCINFO["FS"] "FS" if field splitting with FS is in effect, or "FIELDWIDTHS" if field splitting with FIELDWIDTHS is in effect. PROCINFO["gid"] the value of the getgid(2) system call. PROCINFO["pgrp"] the process group ID of the current process. PROCINFO["pid"] the process ID of the current process. PROCINFO["ppid"] the parent process ID of the current process. PROCINFO["uid"] the value of the getuid(2) system call.	–	–	√
<b>RT</b>	The record terminator. Gawk sets <b>RT</b> to the input text that matched the character or regular expression specified by <b>RS</b> .	–	–	√
<b>RSTART</b>	The index of the first character matched by <i>match()</i> ; 0 if no match.	–	√	√
<b>RLENGTH</b>	The length of the string matched by <i>match()</i> ; -1 if no match.	–	√	√
<b>SUBSEP</b>	The character used to separate multiple subscripts in array elements. Default: "\034" (non-printable character, dec: 28, hex: 1C)	–	√	√
<b>TEXTDOMAIN</b>	The text domain of the AWK program; used to find the localized translations for the program's strings.	–	–	√

**Command line argument summary:**

Argument	Description
<b>-F <i>fs</i></b> <b>--field-separator <i>fs</i></b>	Use <i>fs</i> for the input field separator (the value of the <b>FS</b> predefined variable).
<b>-v <i>var=val</i></b> <b>--assign <i>var=val</i></b>	Assign the value <i>val</i> to the variable <i>var</i> , before execution of the program begins. Such variable values are available to the <b>BEGIN</b> block of an <i>AWK</i> program.
<b>-f <i>program-file</i></b> <b>--file <i>program-file</i></b>	Read the <i>AWK</i> program source from the file <i>program-file</i> , instead of from the first command line argument. Multiple <b>-f</b> (or <b>--file</b> ) options may be used.
<b>-mf <i>NNN</i></b> <b>-mr <i>NNN</i></b>	Set various memory limits to the value <i>NNN</i> . The <b>f</b> flag sets the maximum number of fields, and the <b>r</b> flag sets the maximum record size. (Ignored by <i>gawk</i> , since <i>gawk</i> has no pre-defined limits)
<b>-W compat</b> <b>-W traditional</b> <b>--compat</b> <b>--traditional</b>	Run in compatibility mode. In compatibility mode, <i>gawk</i> behaves identically to UNIX <i>awk</i> ; none of the GNU-specific extensions are recognized.
<b>-W copyleft</b> <b>-W copyright</b> <b>--copyleft</b> <b>--copyright</b>	Print the short version of the GNU copyright information message on the standard output and exit successfully.
<b>-W dump-variables[=<i>file</i>]</b> <b>--dump-variables[=<i>file</i>]</b>	Print a sorted list of global variables, their types and final values to <i>file</i> . If no <i>file</i> is provided, <i>gawk</i> uses a file named <i>awkvars.out</i> in the current directory.
<b>-W help</b> <b>-W usage</b> <b>--help</b> <b>--usage</b>	Print a relatively short summary of the available options on the standard output.
<b>-W lint[=<i>value</i>]</b> <b>--lint[=<i>value</i>]</b>	Provide warnings about constructs that are dubious or non-portable to other <i>AWK</i> impl's. With argument <b>fatal</b> , lint warnings become fatal errors. With an optional argument of <b>invalid</b> , only warnings about things that are actually invalid are issued. (This is not fully implemented yet.)
<b>-W lint-old</b> <b>--lint-old</b>	Provide warnings about constructs that are not portable to the original version of Unix <i>awk</i> .
<b>-W gen-po</b> <b>--gen-po</b>	Scan and parse the <i>AWK</i> program, and generate a GNU <i>.po</i> format file on standard output with entries for all localizable strings in the program. The program itself is not executed.
<b>-W non-decimal-data</b> <b>--non-decimal-data</b>	Recognize octal and hexadecimal values in input data.
<b>-W posix</b> <b>--posix</b>	This turns on compatibility mode, with the following additional restrictions: <ul style="list-style-type: none"> <li>• \x escape sequences are not recognized.</li> <li>• Only space and tab act as field separators when <b>FS</b> is set to a single space, new-line does not.</li> <li>• You cannot continue lines after <b>?</b> and <b>:</b>.</li> <li>• The synonym <b>func</b> for the keyword function is not recognized.</li> <li>• The operators <b>**</b> and <b>**=</b> cannot be used in place of <b>^</b> and <b>^=</b>.</li> <li>• The <i>fflush()</i> function is not available.</li> </ul>
<b>-W profile[=<i>prof_file</i>]</b> <b>--profile[=<i>prof_file</i>]</b>	Send profiling data to <i>prof_file</i> . The default is <i>awkprof.out</i> . When run with <i>gawk</i> , the profile is just a "pretty printed" version of the program. When run with <i>pgawk</i> , the profile contains execution counts of each statement in the program in the left margin and function call counts for each user-defined function.
<b>-W re-interval</b> <b>--re-interval</b>	Enable the use of interval expressions in regular expression matching. Interval expressions were not traditionally available in the <i>AWK</i> language.
<b>-W source <i>program-text</i></b> <b>--source <i>program-text</i></b>	Use <i>program-text</i> as <i>AWK</i> program source code. This option allows the easy intermixing of library functions (used via the <b>-f</b> and <b>--file</b> options) with source code entered on the command line.
<b>-W version</b> <b>--version</b>	Print version information for this particular copy of <i>gawk</i> on the standard output.
<b>--</b>	Signal the end of options. This is useful to allow further arguments to the <i>AWK</i> program itself to start with a "-". This is mainly for consistency with the argument parsing convention used by most other POSIX programs.

## Statements and Functions:

I/O Statements	Description
<b>close(file [, how])</b>	Close <i>file</i> , <i>pipe</i> or <i>co-process</i> . The optional <i>how</i> should only be used when closing one end of a two-way pipe to a co-process. It must be a string value, either "to" or "from".
<b>getline</b>	Set <b>\$0</b> from next input record; set <b>NF</b> , <b>NR</b> , <b>FNR</b> . Returns 0 on EOF and -1 on an error. Upon an error, <b>ERRNO</b> contains a string describing the problem.
<b>getline &lt;file</b>	Set <b>\$0</b> from next record of <i>file</i> ; set <b>NF</b> .
<b>getline var</b>	Set <i>var</i> from next input record; set <b>NR</b> , <b>FNR</b> .
<b>getline var &lt;file</b>	Set <i>var</i> from next record of <i>file</i> .
<b>command   getline [var]</b>	Run <b>command</b> piping the output either into <b>\$0</b> or <i>var</i> , as above. If using a pipe or co-process to <b>getline</b> , or from <b>print</b> or <b>printf</b> within a loop, you must use <b>close()</b> to create new instances
<b>command  &amp; getline [var]</b>	Run <b>command</b> as a co-process piping the output either into <b>\$0</b> or <i>var</i> , as above. Co-processes are a <i>gawk</i> extension.
<b>next</b>	Stop processing the current input record. The next input record is read and processing starts over with the first pattern in the AWK program. If the end of the input data is reached, the <b>END</b> block(s), if any, are executed.
<b>nextfile</b>	Stop processing the current input file. The next input record read comes from the next input file. <b>FILENAME</b> and <b>ARGIND</b> are updated, <b>FNR</b> is reset to 1, and processing starts over with the first pattern in the AWK program. If the end of the input data is reached, the <b>END</b> block(s), are executed.
<b>print</b>	Prints the current record. The output record is terminated with the value of the <b>ORS</b> variable.
<b>print expr-list</b>	Prints expressions. Each expression is separated by the value of the <b>OFS</b> variable. The output record is terminated with the value of the <b>ORS</b> variable.
<b>print expr-list &gt;file</b>	Prints expressions on file. Each expression is separated by the value of the <b>OFS</b> variable. The output record is terminated with the value of the <b>ORS</b> variable.
<b>printf fmt, expr-list</b>	Format and print.
<b>printf fmt, expr-list &gt;file</b>	Format and print on file.
<b>system(cmd-line)</b>	Execute the command <i>cmd-line</i> , and return the exit status.
<b>fflush([file])</b>	Flush any buffers associated with the open output file or pipe file. If file is missing, then stdout is flushed. If file is the null string, then all open output files and pipes have their buffers flushed.
<b>print ... &gt;&gt; file</b>	appends output to the file.
<b>print ...   command</b>	writes on a pipe.
<b>print ...  &amp; command</b>	sends data to a co-process.

Numeric Functions	Description
<b>atan2(y, x)</b>	Returns the arctangent of <i>y/x</i> in radians.
<b>cos(expr)</b>	Returns the cosine of <i>expr</i> , which is in radians.
<b>exp(expr)</b>	The exponential function.
<b>int(expr)</b>	Truncates to integer.
<b>log(expr)</b>	The natural logarithm function.
<b>rand()</b>	Returns a random number <i>N</i> , between 0 and 1, such that $0 \leq N < 1$ .
<b>sin(expr)</b>	Returns the sine of <i>expr</i> , which is in radians.
<b>sqrt(expr)</b>	The square root function.
<b>srand([expr])</b>	Uses <i>expr</i> as a new seed for the random number generator. If no <i>expr</i> is provided, the time of day is used. The return value is the previous seed for the random number generator.

Bit Manipulations Functions	Description
<b>and(v1, v2)</b>	Return the bitwise AND of the values provided by <i>v1</i> and <i>v2</i> .
<b>compl(val)</b>	Return the bitwise complement of <i>val</i> .
<b>lshift(val, count)</b>	Return the value of <i>val</i> , shifted left by <i>count</i> bits.
<b>or(v1, v2)</b>	Return the bitwise OR of the values provided by <i>v1</i> and <i>v2</i> .
<b>rshift(val, count)</b>	Return the value of <i>val</i> , shifted right by <i>count</i> bits.
<b>xor(v1, v2)</b>	Return the bitwise XOR of the values provided by <i>v1</i> and <i>v2</i> .

I18N functions
<b>bindtextdomain(directory [, domain])</b>
Specifies the directory where <i>gawk</i> looks for the <i>.mo</i> files. It returns the directory where <i>domain</i> is ``bound.'' The default domain is the value of <b>TEXTDOMAIN</b> . If <i>directory</i> is the null string (""), then <b>bindtextdomain()</b> returns the current binding for the given <i>domain</i> .
<b>dcgettext(string [, domain [, category]])</b>
Returns the translation of <i>string</i> in text domain <i>domain</i> for locale category <i>category</i> . The default value for <i>domain</i> is the current value of <b>TEXTDOMAIN</b> . The default value for <i>category</i> is "LC_MESSAGES". If you supply a value for <i>category</i> , it must be a string equal to one of the known locale categories. You must also supply a text domain. Use <b>TEXTDOMAIN</b> if you want to use the current domain.
<b>dcngettext(string1 , string2 , number [, domain [, category]])</b>
Returns the plural form used for number of the translation of <i>string1</i> and <i>string2</i> in text domain <i>domain</i> for locale category <i>category</i> . The default value for <i>domain</i> is the current value of <b>TEXTDOMAIN</b> . The default value for <i>category</i> is "LC_MESSAGES". If you supply a value for <i>category</i> , it must be a string equal to one of the known locale categories. You must also supply a text domain. Use <b>TEXTDOMAIN</b> if you want to use the current domain.

String Functions	Description
<b>asort(<i>s</i> [, <i>d</i>])</b>	Returns the number of elements in the source array <i>s</i> . The contents of <i>s</i> are sorted using <i>gawk</i> 's normal rules for comparing values, and the indexes of the sorted values of <i>s</i> are replaced with sequential integers starting with 1. If the optional destination array <i>d</i> is specified, then <i>s</i> is first duplicated into <i>d</i> , and then <i>d</i> is sorted, leaving the indexes of the source array <i>s</i> unchanged.
<b>asorti(<i>s</i> [, <i>d</i>])</b>	Returns the number of elements in the source array <i>s</i> . The behavior is the same as that of <b>asort()</b> , except that the array indices are used for sorting, not the array values. When done, the array is indexed numerically, and the values are those of the original indices. The original values are lost; thus provide a second array if you wish to preserve the original.
<b>gensub(<i>r</i>, <i>s</i>, <i>h</i> [, <i>t</i>])</b>	Search the target string <i>t</i> for matches of the regular expression <i>r</i> . If <i>h</i> is a string beginning with <i>g</i> or <i>G</i> , then replace all matches of <i>r</i> with <i>s</i> . Otherwise, <i>h</i> is a number indicating which match of <i>r</i> to replace. If <i>t</i> is not supplied, <b>\$0</b> is used instead. Within the replacement text <i>s</i> , the sequence <code>\n</code> , where <i>n</i> is a digit from 1 to 9, may be used to indicate just the text that matched the <i>n</i> 'th parenthesized subexpression. The sequence <code>\0</code> represents the entire matched text, as does the character <code>&amp;</code> . Unlike <b>sub()</b> and <b>gsub()</b> , the modified string is returned as the result of the function, and the original target string is not changed.
<b>gsub(<i>r</i>, <i>s</i> [, <i>t</i>])</b>	For each substring matching the regular expression <i>r</i> in the string <i>t</i> , substitute the string <i>s</i> , and return the number of substitutions. If <i>t</i> is not supplied, use <b>\$0</b> . An <code>&amp;</code> in the replacement text is replaced with the text that was actually matched. Use <code>\&amp;</code> to get a literal <code>&amp;</code> . (This must be typed as <code>"\&amp;"</code> )
<b>index(<i>s</i>, <i>t</i>)</b>	Returns the index of the string <i>t</i> in the string <i>s</i> , or 0 if <i>t</i> is not present. (This implies that character indices start at one.)
<b>length([<i>s</i>])</b>	Returns the length of the string <i>s</i> , or the length of <b>\$0</b> if <i>s</i> is not supplied.
<b>match(<i>s</i>, <i>r</i> [, <i>a</i>])</b>	Returns the position in <i>s</i> where the regular expression <i>r</i> occurs, or 0 if <i>r</i> is not present, and sets the values of <b>RSTART</b> and <b>RLENGTH</b> . Note that the argument order is the same as for the <code>~</code> operator: <code>str ~ re</code> . If array <i>a</i> is provided, <i>a</i> is cleared and then elements 1 through <i>n</i> are filled with the portions of <i>s</i> that match the corresponding parenthesized subexpression in <i>r</i> . The 0'th element of <i>a</i> contains the portion of <i>s</i> matched by the entire regular expression <i>r</i> . Subscripts <code>a[n, "start"]</code> , and <code>a[n, "length"]</code> provide the starting index in the string and length respectively, of each matching substring.
<b>split(<i>s</i>, <i>a</i> [, <i>r</i>])</b>	Splits the string <i>s</i> into the array <i>a</i> on the regular expression <i>r</i> , and returns the number of fields. If <i>r</i> is omitted, <b>FS</b> is used instead. The array <i>a</i> is cleared first. Splitting behaves identically to field splitting.
<b>sprintf(<i>fmt</i>, <i>expr-list</i>)</b>	Prints <i>expr-list</i> according to <i>fmt</i> , and returns the resulting string.
<b>strtonum(<i>str</i>)</b>	Examines <i>str</i> , and returns its numeric value. If <i>str</i> begins with a leading 0, <b>strtonum()</b> assumes that <i>str</i> is an octal number. If <i>str</i> begins with a leading 0x or 0X, <b>strtonum()</b> assumes that <i>str</i> is a hexadecimal number.
<b>sub(<i>r</i>, <i>s</i> [, <i>t</i>])</b>	Just like <b>gsub()</b> , but only the first matching substring is replaced.
<b>substr(<i>s</i>, <i>i</i> [, <i>n</i>])</b>	Returns the at most <i>n</i> -character substring of <i>s</i> starting at <i>i</i> . If <i>n</i> is omitted, the rest of <i>s</i> is used.
<b>tolower(<i>str</i>)</b>	Returns a copy of the string <i>str</i> , with all the upper-case characters in <i>str</i> translated to their corresponding lower-case counterparts. Non-alphabetic characters are left unchanged.
<b>toupper(<i>str</i>)</b>	Returns a copy of the string <i>str</i> , with all the lower-case characters in <i>str</i> translated to their corresponding upper-case counterparts. Non-alphabetic characters are left unchanged.

Time Functions	Description
<b>mktime(<i>datespec</i>)</b>	Turns <i>datespec</i> into a time stamp of the same form as returned by <b>systemtime()</b> . The <i>datespec</i> is a string of the form <code>YYYY MM DD HH MM SS[ DST]</code> . The contents of the string are six or seven numbers representing respectively the full year including century, the month from 1 to 12, the day of the month from 1 to 31, the hour of the day from 0 to 23, the minute from 0 to 59, and the second from 0 to 60, and an optional daylight saving flag. The values of these numbers need not be within the ranges specified; for example, an hour of -1 means 1 hour before midnight. The origin-zero Gregorian calendar is assumed, with year 0 preceding year 1 and year -1 preceding year 0. The time is assumed to be in the local timezone. If the daylight saving flag is positive, the time is assumed to be daylight saving time; if zero, the time is assumed to be standard time; and if negative (the default), <b>mktime()</b> attempts to determine whether daylight saving time is in effect for the specified time. If <i>datespec</i> does not contain enough elements or if the resulting time is out of range, <b>mktime()</b> returns -1.
<b>strftime([<i>format</i> [, <i>timestamp</i>]])</b>	Formats timestamp according to the specification in <i>format</i> . The timestamp should be of the same form as returned by <b>systemtime()</b> . If timestamp is missing, the current time of day is used. If <i>format</i> is missing, a default format equivalent to the output of <code>date(1)</code> is used. See the specification for the <code>strftime()</code> function in ANSI C for the format conversions that are guaranteed to be available. A public-domain version of <code>strftime(3)</code> and a man page for it come with <i>gawk</i> ; if that version was used to build <i>gawk</i> , then all of the conversions described in that man page are available to <i>gawk</i> .
<b>systemtime()</b>	Returns the current time of day as the number of seconds since the Epoch (1970-01-01 00:00:00 UTC on POSIX systems).