

### Usage

```
awk [-v var=val] 'program' [file1 file2...]
awk [-v var=val] -f progfile [file1 file2...]
```

### Predefined Variable Summary

|                       |   |
|-----------------------|---|
| <b>FS</b>             | Input Field Separator, a space by default.  |
| <b>OFS</b>            | Output Field Separator, a space by default.   |
| <b>RS</b>             | Record Separator, a newline by default.   |
| <b>ORS</b>            | Output Record Separator, a newline by default.  |
| <b>NR</b>             | The total Number of input Records seen so far.  |
| <b>NF</b>             | The Number of Fields in the current input record.   |
| <b>FILENAME</b>       | The name of the current input file. If no files are specified on the command line, the value of FILENAME is "-". However, FILENAME is undefined inside the BEGIN block (unless set by getline). |
| <b>FNR</b>            | Contains number of lines read, but is reset for each file read.   |
| <b>\$0</b>            | The whole line.   |
| <b>\$1, \$2...\$n</b> | Fields from 1 to n.   |
| <b>ARGC</b>           | The number of command line arguments (does not include options to gawk, or the program source). Dynamically changing the contents of ARGV control the files used for data.                      |
| <b>ARGV</b>           | Array of command line arguments. The array is indexed from 0 to ARGC - 1.   |
| <b>ARGIND</b>         | The index in ARGV of the current file being processed.  |

### Numeric Functions

|                    |  |
|--------------------|--|
| <b>atan2(y, x)</b> | Returns the arctangent of y/x in radians.                              |
| <b>cos(expr)</b>   | Returns the cosine of expr, 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 N, between 0 and 1, such that $0 \leq N < 1$ . |

### Numeric Functions (cont)


|                      |  |
|----------------------|--|
| <b>sin(expr)</b>     | Returns the sine of expr, which is in radians.   |
| <b>sqrt(expr)</b>    | The square root function.  |
| <b>srand([expr])</b> | Uses expr as a new seed for the random number generator. If no expr is provided, the time of day is used. The return value is the previous seed for the random number generator. |

### Bit Manipulation Functions

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

### String Functions

|                       |   |
|-----------------------|---|
| <b>asort(s [, d])</b> | Returns the number of elements in the source array s. The contents of s are sorted using gawk's normal rules for comparing values, and the indexes of the sorted values of s are replaced with sequential integers starting with 1. If the optional destination array d is specified, then s is first duplicated into d, and then d is sorted, leaving the indexes of the source array s unchanged. |
|-----------------------|---|



By **TME520** (TME520)  
[cheatography.com/tme520/](https://cheatography.com/tme520/)  
[tme520.com](https://tme520.com)

Published 23rd April, 2015.  
 Last updated 12th May, 2016.  
 Page 1 of 5.

Sponsored by **Readable.com**  
 Measure your website readability!  
<https://readable.com>

### String Functions (cont)

**asorti(s [, d])** Returns the number of elements in the source array s. The behavior is the same as that of `asort()`, 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.

**gensub(r, s, h [, t])** Search the target string t for matches of the regular expression r. If h is a string beginning with g or G, then replace all matches of r with s. Otherwise, h is a number indicating which match of r to replace. If t is not supplied, \$0 is used instead. Within the replacement text s, the sequence `\n`, where n is a digit from 1 to 9, may be used to indicate just the text that matched the n'th parenthesized subexpression. The sequence `\0` represents the entire matched text, as does the character `&`. Unlike `sub()` and `gsub()`, the modified string is returned as the result of the function, and the original target string is not changed.

**gsub(r, s [, t])** For each substring matching the regular expression r in the string t, substitute the string s, and return the number of substitutions. If t is not supplied, use \$0. An `&` in the replacement text is replaced with the text that was actually matched. Use `\\&` to get a literal `&` (This must be typed as `\"\\&\"`).

### String Functions (cont)


**index(s, t)** Returns the index of the string t in the string s, or 0 if t is not present (this implies that character indices start at one).

**length([s])** Returns the length of the string s, or the length of \$0 if s is not supplied.

**match(s, r [, a])** Returns the position in s where the regular expression r occurs, or 0 if r is not present, and sets the values of RSTART and RLENGTH. Note that the argument order is the same as for the `~` operator: `str ~ re`. If array a is provided, a is cleared and then elements 1 through n are filled with the portions of s that match the corresponding parenthesized subexpression in r. The 0'th element of a contains the portion of s matched by the entire regular expression r. Subscripts `a[n, "start"]`, and `a[n, "length"]` provide the starting index in the string and length respectively, of each matching substring.

**split(s, a [, r])** Splits the string s into the array a on the regular expression r, and returns the number of fields. If r is omitted, FS is used instead. The array a is cleared first. Splitting behaves identically to field splitting.

**sprintf(fmt, expr-list)** Prints expr-list according to fmt, and returns the resulting string.



By TME520 (TME520)  
[cheatography.com/tme520/tme520.com](https://cheatography.com/tme520/tme520.com)

Published 23rd April, 2015.  
Last updated 12th May, 2016.  
Page 2 of 5.

Sponsored by **Readable.com**  
Measure your website readability!  
<https://readable.com>

### String Functions (cont)

|                           |  |
|---------------------------|--|
| <b>strtonum(str)</b>      | Examines str, and returns its numeric value. If str begins with a leading 0, strtonum() assumes that str is an octal number. If str begins with a leading 0x or 0X, strtonum() assumes that str is a hexadecimal number. |
| <b>sub(r, s [, t])</b>    | Just like gsub(), but only the first matching substring is replaced.   |
| <b>substr(s, i [, n])</b> | Returns the at most n-character substring of s starting at i. If n is omitted, the rest of s is used.  |
| <b>tolower(str)</b>       | Returns a copy of the string str, with all the upper-case characters in str translated to their corresponding lower-case counterparts. Non-alphabetic characters are left unchanged.                                     |
| <b>toupper(str)</b>       | Returns a copy of the string str, with all the lower-case characters in str translated to their corresponding upper-case counterparts. Non-alphabetic characters are left unchanged.                                     |

### Operators


|   |                                  |
|---|----------------------------------|
| <b>&amp;&amp;    !</b>                  | Logical operators : AND, OR, NOT |
| <b>&lt; &lt;= == != &gt;= &gt; ~ !~</b> | Comparison operators.            |

### I/O Statements

|                            |  |
|----------------------------|--|
| <b>close(file [, how])</b> | Close file, pipe or co-process. The optional how 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 \$0 from next input record; set NF, NR, FNR. Returns 0 on EOF and 1 on an error. Upon an error, ERRNO contains a string describing the problem.                            |

### I/O Statements (cont)

|                                     |   |
|-------------------------------------|---|
| <b>getline &lt;file</b>             | Set \$0 from next record of file; set NF.   |
| <b>getline var</b>                  | Set var from next input record; set NR, FNR.  |
| <b>getline var &lt;file</b>         | Set var from next record of file.   |
| <b>command   getline [var]</b>      | Run command piping the output either into \$0 or var, as above. If using a pipe or co-process to getline, or from print or printf within a loop, you must use close() to create new instances.  |
| <b>command  &amp; getline [var]</b> | Run command as a co-process piping the output either into \$0 or var, as above. Co-processes are a gawk 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 END 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. FILENAME and ARGIND are updated, FNR 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 END block(s), are executed. |
| <b>print</b>                        | Prints the current record. The output record is terminated with the value of the ORS variable.  |
| <b>print expr-list</b>              | Prints expressions. Each expression is separated by the value of the OFS variable. The output record is terminated with the value of the ORS variable.  |



By **TME520** (TME520)  
[cheatography.com/tme520/](https://cheatography.com/tme520/)  
[tme520.com](https://tme520.com)

Published 23rd April, 2015.  
 Last updated 12th May, 2016.  
 Page 3 of 5.

Sponsored by **Readable.com**  
 Measure your website readability!  
<https://readable.com>

### I/O Statements (cont)

|                                       |   |
|---------------------------------------|---|
| <b>print expr-list &gt;file</b>       | Prints expressions on file. Each expression is separated by the value of the OFS variable. The output record is terminated with the value of the ORS 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 cmd-line, 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.   |

### Time Functions

|   |  |
|---|--|
| <b>systeme()</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).   |
| <b>mktime(datespec)</b>                 | Turns datespec into a time stamp of the same form as returned by systeme(). The datespec is a string of the form YYYY MM DD HH MM SS[ DST].  |
| <b>strftime([format [, timestamp]])</b> | Formats timestamp according to the specification in format. The timestamp should be of the same form as returned by systeme(). If timestamp is missing, the current time of day is used. If format is missing, a default format equivalent to the output of date(1) is used. |


### GNU AWK's Command Line Argument Summary

|   |  |
|---|--|
| <b>-F fs or --field-separator fs</b>          | Use fs for the input field separator (the value of the FS predefined variable).  |
| <b>-v var=val or --assign var=val</b>         | Assign the value val to the variable var, before execution of the program begins. Such variable values are available to the BEGIN block of an AWK program.                                     |
| <b>-f program-file or --file program-file</b> | Read the AWK program source from the file program-file, instead of from the first command line argument. Multiple -f (or -file) options may be used.   |
| <b>-mf NNN or -mr NNN</b>                     | Set various memory limits to the value NNN. The f flag sets the maximum number of fields, and the r flag sets the maximum record size (ignored by gawk, since gawk has no pre-defined limits). |

**-W compat** *or* **-W traditional** *or* **--compat--traditional** Run in compatibility mode. In compatibility mode, gawk behaves identically to UNIX awk; none of the GNU-specific extensions are recognized.

**-W dump-variables[=file]** *or* **--dump-variables[=file]** Print a sorted list of global variables, their types and final values to file. If no file is provided, gawk uses a file named awkvars.out in the current directory.

**-W help** *or* **-W usage** *or* **--help** *or* **--usage** Print a relatively short summary of the available options on the standard output.



By **TME520** (TME520)  
[cheatography.com/tme520/tme520.com](https://cheatography.com/tme520/tme520.com)

Published 23rd April, 2015.  
Last updated 12th May, 2016.  
Page 4 of 5.

Sponsored by **Readable.com**  
Measure your website readability!  
<https://readable.com>