



Full credit is given to the above companies including the OS that this PDF file was generated!

Red Hat Enterprise Linux Release 9.2 Manual Pages on 'perlrun.1' command

\$ man perlrun.1

PERLRUN(1) Perl Programmers Reference Guide PERLRUN(1)

NAME

perlrun - how to execute the Perl interpreter

SYNOPSIS

```
perl [ -sTtuUWX ] [ -hv ] [ -V[:configvar] ]
    [ -cw ] [ -d[t][:debugger] ] [ -D[number/list] ]
    [ -pna ] [ -Fpattern ] [ -l[octal] ] [ -O[octal/hexadecimal] ]
    [ -ldir ] [ -m[-]module ] [ -M[-]'module...' ] [ -f ]
    [ -C [number/list] ] [ -S ] [ -x[dir] ]
    [ -i[extension] ]
    [ [-e|-E] 'command' ] [ -- ] [ programfile ] [ argument ]...
```

DESCRIPTION

The normal way to run a Perl program is by making it directly executable, or else by passing the name of the source file as an argument on the command line. (An interactive Perl environment is also possible--see perldebug for details on how to do that.) Upon startup, Perl looks for your program in one of the following places:

1. Specified line by line via -e or -E switches on the command line.
2. Contained in the file specified by the first filename on the command line. (Note that systems supporting the "#!" notation invoke interpreters this way. See "Location of Perl".)
3. Passed in implicitly via standard input. This works only if there are no filename arguments--to pass arguments to a STDIN-read

program you must explicitly specify a "-" for the program name.

With methods 2 and 3, Perl starts parsing the input file from the beginning, unless you've specified a "-x" switch, in which case it scans for the first line starting with "#!" and containing the word "perl", and starts there instead. This is useful for running a program embedded in a larger message. (In this case you would indicate the end of the program using the "__END__" token.)

The "#!" line is always examined for switches as the line is being parsed. Thus, if you're on a machine that allows only one argument with the "#!" line, or worse, doesn't even recognize the "#!" line, you still can get consistent switch behaviour regardless of how Perl was invoked, even if "-x" was used to find the beginning of the program. Because historically some operating systems silently chopped off kernel interpretation of the "#!" line after 32 characters, some switches may be passed in on the command line, and some may not; you could even get a "-" without its letter, if you're not careful. You probably want to make sure that all your switches fall either before or after that 32-character boundary. Most switches don't actually care if they're processed redundantly, but getting a "-" instead of a complete switch could cause Perl to try to execute standard input instead of your program. And a partial -l switch could also cause odd results. Some switches do care if they are processed twice, for instance combinations of -l and -0. Either put all the switches after the 32-character boundary (if applicable), or replace the use of -0digits by "BEGIN{ \$/ = "\0digits"; }".

Parsing of the "#!" switches starts wherever "perl" is mentioned in the line. The sequences "-*" and "- " are specifically ignored so that you could, if you were so inclined, say

```
#!/bin/sh

#! -* perl -* -p

eval `exec perl -x -wS $0 ${1+"$@"}`

if 0;
```

to let Perl see the "-p" switch.

A similar trick involves the `env` program, if you have it.

```
#!/usr/bin/env perl
```

The examples above use a relative path to the perl interpreter, getting whatever version is first in the user's path. If you want a specific version of Perl, say, `perl5.14.1`, you should place that directly in the `"#!"` line's path.

If the `"#!"` line does not contain the word `"perl"` nor the word `"indir"`, the program named after the `"#!"` is executed instead of the Perl interpreter. This is slightly bizarre, but it helps people on machines that don't do `"#!"`, because they can tell a program that their SHELL is `/usr/bin/perl`, and Perl will then dispatch the program to the correct interpreter for them.

After locating your program, Perl compiles the entire program to an internal form. If there are any compilation errors, execution of the program is not attempted. (This is unlike the typical shell script, which might run part-way through before finding a syntax error.)

If the program is syntactically correct, it is executed. If the program runs off the end without hitting an `exit()` or `die()` operator, an implicit `exit(0)` is provided to indicate successful completion.

`#!` and quoting on non-Unix systems

Unix's `"#!"` technique can be simulated on other systems:

OS/2

Put

```
extproc perl -S -your_switches
```

as the first line in `"*.cmd"` file (`"-S"` due to a bug in `cmd.exe`'s ``extproc'` handling).

MS-DOS

Create a batch file to run your program, and codify it in `"ALTERNATE_SHEBANG"` (see the `dosish.h` file in the source distribution for more information).

Win95/NT

The Win95/NT installation, when using the ActiveState installer for Perl, will modify the Registry to associate the `.pl` extension with

the perl interpreter. If you install Perl by other means (including building from the sources), you may have to modify the Registry yourself. Note that this means you can no longer tell the difference between an executable Perl program and a Perl library file.

VMS Put

```
$ perl -mysw 'f$env("procedure")' 'p1' 'p2' 'p3' 'p4' 'p5' 'p6' 'p7' 'p8' !  
$ exit++ + ++$status != 0 and $exit = $status = undef;
```

at the top of your program, where -mysw are any command line switches you want to pass to Perl. You can now invoke the program directly, by saying "perl program", or as a DCL procedure, by saying @program (or implicitly via DCL\$PATH by just using the name of the program).

This incantation is a bit much to remember, but Perl will display it for you if you say "perl -V:startperl".

Command-interpreters on non-Unix systems have rather different ideas on quoting than Unix shells. You'll need to learn the special characters in your command-interpreter ("*", "\", and "" are common) and how to protect whitespace and these characters to run one-liners (see -e below).

On some systems, you may have to change single-quotes to double ones, which you must not do on Unix or Plan 9 systems. You might also have to change a single % to a %%.

For example:

```
# Unix  
perl -e 'print "Hello world\n"  
  
# MS-DOS, etc.  
perl -e "print \"Hello world\n\""  
  
# VMS  
perl -e "print ""Hello world\n"""
```

The problem is that none of this is reliable: it depends on the command and it is entirely possible neither works. If 4DOS were the command shell, this would probably work better:

```
perl -e "print <Ctrl-x>"Hello world\n<Ctrl-x>"
```

CMD.EXE in Windows NT slipped a lot of standard Unix functionality in when nobody was looking, but just try to find documentation for its quoting rules.

There is no general solution to all of this. It's just a mess.

Location of Perl

It may seem obvious to say, but Perl is useful only when users can easily find it. When possible, it's good for both `/usr/bin/perl` and `/usr/local/bin/perl` to be symlinks to the actual binary. If that can't be done, system administrators are strongly encouraged to put (symlinks to) perl and its accompanying utilities into a directory typically found along a user's PATH, or in some other obvious and convenient place.

In this documentation, `#!/usr/bin/perl` on the first line of the program will stand in for whatever method works on your system. You are advised to use a specific path if you care about a specific version.

```
#!/usr/local/bin/perl5.14
```

or if you just want to be running at least version, place a statement like this at the top of your program:

```
use 5.014;
```

Command Switches

As with all standard commands, a single-character switch may be clustered with the following switch, if any.

```
#!/usr/bin/perl -spi.orig # same as -s -p -i.orig
```

A `--` signals the end of options and disables further option processing. Any arguments after the `--` are treated as filenames and arguments.

Switches include:

`-0[octal/hexadecimal]`

specifies the input record separator (`$/`) as an octal or hexadecimal number. If there are no digits, the null character is the separator. Other switches may precede or follow the digits.

For example, if you have a version of find which can print filenames terminated by the null character, you can say this:

```
find . -name '*.orig' -print0 | perl -n0e unlink
```

The special value 00 will cause Perl to slurp files in paragraph mode. Any value 0400 or above will cause Perl to slurp files whole, but by convention the value 0777 is the one normally used for this purpose.

You can also specify the separator character using hexadecimal notation: -0xHHH..., where the "H" are valid hexadecimal digits.

Unlike the octal form, this one may be used to specify any Unicode character, even those beyond 0xFF. So if you really want a record separator of 0777, specify it as -0x1FF. (This means that you cannot use the "-x" option with a directory name that consists of hexadecimal digits, or else Perl will think you have specified a hex number to -0.)

-a turns on autosplit mode when used with a "-n" or "-p". An implicit split command to the @F array is done as the first thing inside the implicit while loop produced by the "-n" or "-p".

```
perl -ane 'print pop(@F), "\n";'
```

is equivalent to

```
while (<>) {  
    @F = split(' ');  
    print pop(@F), "\n";  
}
```

An alternate delimiter may be specified using -F.

-a implicitly sets "-n".

-C [number/list]

The -C flag controls some of the Perl Unicode features.

As of 5.8.1, the -C can be followed either by a number or a list of option letters. The letters, their numeric values, and effects are as follows; listing the letters is equal to summing the numbers.

- O 2 STDOUT will be in UTF-8
- E 4 STDERR will be in UTF-8
- S 7 I + O + E
- i 8 UTF-8 is the default PerlIO layer for input streams
- o 16 UTF-8 is the default PerlIO layer for output streams
- D 24 i + o
- A 32 the @ARGV elements are expected to be strings encoded in UTF-8
- L 64 normally the "IOEioA" are unconditional, the L makes them conditional on the locale environment variables (the LC_ALL, LC_CTYPE, and LANG, in the order of decreasing precedence) -- if the variables indicate UTF-8, then the selected "IOEioA" are in effect
- a 256 Set \${^UTF8CACHE} to -1, to run the UTF-8 caching code in debugging mode.

For example, -COE and -C6 will both turn on UTF-8-ness on both STDOUT and STDERR. Repeating letters is just redundant, not cumulative nor toggling.

The "io" options mean that any subsequent open() (or similar I/O operations) in main program scope will have the ":utf8" PerlIO layer implicitly applied to them, in other words, UTF-8 is expected from any input stream, and UTF-8 is produced to any output stream. This is just the default set via "\${^OPEN}", with explicit layers in open() and with binmode() one can manipulate streams as usual. This has no effect on code run in modules.

-C on its own (not followed by any number or option list), or the empty string "" for the "PERL_UNICODE" environment variable, has the same effect as -CSDL. In other words, the standard I/O handles and the default "open()" layer are UTF-8-fied but only if the locale environment variables indicate a UTF-8 locale. This behaviour follows the implicit (and problematic) UTF-8 behaviour of Perl 5.8.0. (See "UTF-8 no longer default under UTF-8 locales" in perl581delta.)

You can use `-C0` (or `"0"` for `"PERL_UNICODE"`) to explicitly disable all the above Unicode features.

The read-only magic variable `"${^UNICODE}"` reflects the numeric value of this setting. This variable is set during Perl startup and is thereafter read-only. If you want runtime effects, use the three-arg `open()` (see `"open"` in `perlfunc`), the two-arg `binmode()` (see `"binmode"` in `perlfunc`), and the `"open"` pragma (see `open`). (In Perls earlier than 5.8.1 the `-C` switch was a Win32-only switch that enabled the use of Unicode-aware "wide system call" Win32 APIs. This feature was practically unused, however, and the command line switch was therefore "recycled".)

Note: Since perl 5.10.1, if the `-C` option is used on the `"#!"` line, it must be specified on the command line as well, since the standard streams are already set up at this point in the execution of the perl interpreter. You can also use `binmode()` to set the encoding of an I/O stream.

`-c` causes Perl to check the syntax of the program and then exit without executing it. Actually, it will execute any `"BEGIN"`, `"UNITCHECK"`, or `"CHECK"` blocks and any `"use"` statements: these are considered as occurring outside the execution of your program. `"INIT"` and `"END"` blocks, however, will be skipped.

`-d`

`-dt` runs the program under the Perl debugger. See `perldebug`. If `t` is specified, it indicates to the debugger that threads will be used in the code being debugged.

`-d:MOD[=bar,baz]`

`-dt:MOD[=bar,baz]`

runs the program under the control of a debugging, profiling, or tracing module installed as `"Devel::MOD"`. E.g., `-d:DProf` executes the program using the `"Devel::DProf"` profiler. As with the `-M` flag, options may be passed to the `"Devel::MOD"` package where they will be received and interpreted by the `"Devel::MOD::import"` routine. Again, like `-M`, use `--d:-MOD` to call

"Devel::MOD::unimport" instead of import. The comma-separated list of options must follow a "=" character. If t is specified, it indicates to the debugger that threads will be used in the code being debugged. See perldebug.

-Dletters

-Dnumber

sets debugging flags. This switch is enabled only if your perl binary has been built with debugging enabled: normal production perls won't have been.

For example, to watch how perl executes your program, use -Dtls.

Another nice value is -Dx, which lists your compiled syntax tree, and -Dr displays compiled regular expressions; the format of the output is explained in perldebguts.

As an alternative, specify a number instead of list of letters (e.g., -D14 is equivalent to -Dtls):

- 1 p Tokenizing and parsing (with v, displays parse stack)
- 2 s Stack snapshots (with v, displays all stacks)
- 4 l Context (loop) stack processing
- 8 t Trace execution
- 16 o Method and overloading resolution
- 32 c String/numeric conversions
- 64 P Print profiling info, source file input state
- 128 m Memory and SV allocation
- 256 f Format processing
- 512 r Regular expression parsing and execution
- 1024 x Syntax tree dump
- 2048 u Tainting checks
- 4096 U Unofficial, User hacking (reserved for private, unreleased use)
- 16384 X Scratchpad allocation
- 32768 D Cleaning up
- 65536 S Op slab allocation

131072 T Tokenizing

262144 R Include reference counts of dumped variables

(eg when using -Ds)

524288 J show s,t,P-debug (don't Jump over) on opcodes within

package DB

1048576 v Verbose: use in conjunction with other flags to

increase the verbosity of the output. Is a no-op on

many of the other flags

2097152 C Copy On Write

4194304 A Consistency checks on internal structures

8388608 q quiet - currently only suppresses the "EXECUTING"

message

16777216 M trace smart match resolution

33554432 B dump suBroutine definitions, including special

Blocks like BEGIN

67108864 L trace Locale-related info; what gets output is very

subject to change

134217728 i trace PerlIO layer processing. Set PERLIO_DEBUG to

the filename to trace to.

268435456 y trace y///, tr/// compilation and execution

All these flags require -DDEBUGGING when you compile the Perl

executable (but see ":opd" in Devel::Peek or "'debug' mode" in re

which may change this). See the INSTALL file in the Perl source

distribution for how to do this.

If you're just trying to get a print out of each line of Perl code

as it executes, the way that "sh -x" provides for shell scripts,

you can't use Perl's -D switch. Instead do this

```
# If you have "env" utility
```

```
env PERLDB_OPTS="NonStop=1 AutoTrace=1 frame=2" perl -dS program
```

```
# Bourne shell syntax
```

```
$ PERLDB_OPTS="NonStop=1 AutoTrace=1 frame=2" perl -dS program
```

```
# csh syntax
```

```
% (setenv PERLDB_OPTS "NonStop=1 AutoTrace=1 frame=2"; perl -dS program)
```

See `perldebug` for details and variations.

`-e` commandline

may be used to enter one line of program. If `-e` is given, Perl will not look for a filename in the argument list. Multiple `-e` commands may be given to build up a multi-line script. Make sure to use semicolons where you would in a normal program.

`-E` commandline

behaves just like `-e`, except that it implicitly enables all optional features (in the main compilation unit). See `feature`.

`-f` Disable executing `$Config{sitelib}/sitecustomize.pl` at startup.

Perl can be built so that it by default will try to execute `$Config{sitelib}/sitecustomize.pl` at startup (in a `BEGIN` block). This is a hook that allows the sysadmin to customize how Perl behaves. It can for instance be used to add entries to the `@INC` array to make Perl find modules in non-standard locations.

Perl actually inserts the following code:

```
BEGIN {  
    do { local $!; -f "$Config{sitelib}/sitecustomize.pl"; }  
    && do "$Config{sitelib}/sitecustomize.pl";  
}
```

Since it is an actual "do" (not a "require"), `sitecustomize.pl` doesn't need to return a true value. The code is run in package "main", in its own lexical scope. However, if the script dies, `$@` will not be set.

The value of `$Config{sitelib}` is also determined in C code and not read from "Config.pm", which is not loaded.

The code is executed very early. For example, any changes made to `@INC` will show up in the output of ``perl -V``. Of course, "END" blocks will be likewise executed very late.

To determine at runtime if this capability has been compiled in your perl, you can check the value of `$Config{usesitecustomize}`.

`-Fpattern`

specifies the pattern to split on for `"-a"`. The pattern may be

surrounded by "/", "'", or ", otherwise it will be put in single quotes. You can't use literal whitespace or NUL characters in the pattern.

-F implicitly sets both "-a" and "-n".

-h prints a summary of the options.

-i[extension]

specifies that files processed by the "<>" construct are to be edited in-place. It does this by renaming the input file, opening the output file by the original name, and selecting that output file as the default for print() statements. The extension, if supplied, is used to modify the name of the old file to make a backup copy, following these rules:

If no extension is supplied, and your system supports it, the original file is kept open without a name while the output is redirected to a new file with the original filename. When perl exits, cleanly or not, the original file is unlinked.

If the extension doesn't contain a "*", then it is appended to the end of the current filename as a suffix. If the extension does contain one or more "*" characters, then each "*" is replaced with the current filename. In Perl terms, you could think of this as:

```
($backup = $extension) =~ s/^\*/$file_name/g;
```

This allows you to add a prefix to the backup file, instead of (or in addition to) a suffix:

```
$ perl -pi'orig_*' -e 's/bar/baz/' fileA # backup to  
# 'orig_fileA'
```

Or even to place backup copies of the original files into another directory (provided the directory already exists):

```
$ perl -pi'old/*.orig' -e 's/bar/baz/' fileA # backup to  
# 'old/fileA.orig'
```

These sets of one-liners are equivalent:

```
$ perl -pi -e 's/bar/baz/' fileA # overwrite current file  
$ perl -pi*' -e 's/bar/baz/' fileA # overwrite current file  
$ perl -pi'.orig' -e 's/bar/baz/' fileA # backup to 'fileA.orig'
```

```
$ perl -pi*.orig' -e 's/bar/baz/' fileA # backup to 'fileA.orig'
```

From the shell, saying

```
$ perl -p -i.orig -e "s/foo/bar; ... "
```

is the same as using the program:

```
#!/usr/bin/perl -pi.orig
```

```
s/foo/bar/;
```

which is equivalent to

```
#!/usr/bin/perl
```

```
$extension = '.orig';
```

```
LINE: while (<>) {
```

```
    if ($ARGV ne $oldargv) {
```

```
        if ($extension !~ /^/) {
```

```
            $backup = $ARGV . $extension;
```

```
        }
```

```
        else {
```

```
            ($backup = $extension) =~ s/^/$ARGV/g;
```

```
        }
```

```
        rename($ARGV, $backup);
```

```
        open(ARGVOUT, ">$ARGV");
```

```
        select(ARGVOUT);
```

```
        $oldargv = $ARGV;
```

```
    }
```

```
    s/foo/bar/;
```

```
}
```

```
continue {
```

```
    print; # this prints to original filename
```

```
}
```

```
select(STDOUT);
```

except that the `-i` form doesn't need to compare `$ARGV` to `$oldargv`

to know when the filename has changed. It does, however, use

`ARGVOUT` for the selected filehandle. Note that `STDOUT` is restored

as the default output filehandle after the loop.

As shown above, Perl creates the backup file whether or not any

output is actually changed. So this is just a fancy way to copy files:

```
$ perl -p -i'/some/file/path/*' -e 1 file1 file2 file3...
```

or

```
$ perl -p -i'.orig' -e 1 file1 file2 file3...
```

You can use "eof" without parentheses to locate the end of each input file, in case you want to append to each file, or reset line numbering (see example in "eof" in perlfunc).

If, for a given file, Perl is unable to create the backup file as specified in the extension then it will skip that file and continue on with the next one (if it exists).

For a discussion of issues surrounding file permissions and -i, see "Why does Perl let me delete read-only files? Why does -i clobber protected files? Isn't this a bug in Perl?" in perlfaq5.

You cannot use -i to create directories or to strip extensions from files.

Perl does not expand "~" in filenames, which is good, since some folks use it for their backup files:

```
$ perl -pi~ -e 's/foo/bar/' file1 file2 file3...
```

Note that because -i renames or deletes the original file before creating a new file of the same name, Unix-style soft and hard links will not be preserved.

Finally, the -i switch does not impede execution when no files are given on the command line. In this case, no backup is made (the original file cannot, of course, be determined) and processing proceeds from STDIN to STDOUT as might be expected.

-ldirectory

Directories specified by -I are prepended to the search path for modules (@INC).

-l[octnum]

enables automatic line-ending processing. It has two separate effects. First, it automatically chomps \$/ (the input record separator) when used with "-n" or "-p". Second, it assigns "\$\n

(the output record separator) to have the value of octnum so that any print statements will have that separator added back on. If octnum is omitted, sets "\$\" to the current value of \$/. For instance, to trim lines to 80 columns:

```
perl -lpe 'substr($_, 80) = ""'
```

Note that the assignment "\$\ = \$/" is done when the switch is processed, so the input record separator can be different than the output record separator if the -l switch is followed by a -O switch:

```
gnufind / -print0 | perl -ln0e 'print "found $_" if -p'
```

This sets "\$\" to newline and then sets \$/ to the null character.

-m[-]module

-M[-]module

-M[-]'module ...'

-[mM][=arg[,arg]...

-mmodule executes "use" module "();" before executing your program. This loads the module, but does not call its "import" method, so does not import subroutines and does not give effect to a pragma.

-Mmodule executes "use" module ";" before executing your program.

This loads the module and calls its "import" method, causing the module to have its default effect, typically importing subroutines or giving effect to a pragma. You can use quotes to add extra code after the module name, e.g., '-MMODULE qw(foo bar)'

If the first character after the -M or -m is a dash (-) then the 'use' is replaced with 'no'. This makes no difference for -m.

A little builtin syntactic sugar means you can also say

-mMODULE=foo,bar or -MMODULE=foo,bar as a shortcut for '-MMODULE qw(foo bar)'. This avoids the need to use quotes when importing symbols. The actual code generated by -MMODULE=foo,bar is "use module split(/,/q{foo,bar})". Note that the "=" form removes the distinction between -m and -M; that is, -mMODULE=foo,bar is the same as -MMODULE=foo,bar.

A consequence of the "split" formulation is that `-MMODULE=number` never does a version check, unless `"MODULE::import()"` itself is set up to do a version check, which could happen for example if `MODULE` inherits from `Exporter`.

`-n` causes Perl to assume the following loop around your program, which makes it iterate over filename arguments somewhat like `sed -n` or `awk`:

```
LINE:
    while (<>) {
        ...      # your program goes here
    }
```

Note that the lines are not printed by default. See `"-p"` to have lines printed. If a file named by an argument cannot be opened for some reason, Perl warns you about it and moves on to the next file.

Also note that `"<>"` passes command line arguments to `"open"` in `perlfunc`, which doesn't necessarily interpret them as file names.

See `perlop` for possible security implications.

Here is an efficient way to delete all files that haven't been modified for at least a week:

```
find . -mtime +7 -print | perl -nle unlink
```

This is faster than using the `-exec` switch of `find` because you don't have to start a process on every filename found (but it's not faster than using the `-delete` switch available in newer versions of `find`. It does suffer from the bug of mishandling newlines in pathnames, which you can fix if you follow the example under `-0`.

`"BEGIN"` and `"END"` blocks may be used to capture control before or after the implicit program loop, just as in `awk`.

`-p` causes Perl to assume the following loop around your program, which makes it iterate over filename arguments somewhat like `sed`:

```
LINE:
    while (<>) {
```



```

...      # your program goes here
} continue {
    print or die "-p destination: $!\n";
}

```

If a file named by an argument cannot be opened for some reason, Perl warns you about it, and moves on to the next file. Note that the lines are printed automatically. An error occurring during printing is treated as fatal. To suppress printing use the "-n" switch. A -p overrides a -n switch.

"BEGIN" and "END" blocks may be used to capture control before or after the implicit loop, just as in awk.

- s enables rudimentary switch parsing for switches on the command line after the program name but before any filename arguments (or before an argument of --). Any switch found there is removed from @ARGV and sets the corresponding variable in the Perl program. The following program prints "1" if the program is invoked with a -xyz switch, and "abc" if it is invoked with -xyz=abc.

```

#!/usr/bin/perl -s

if ($xyz) { print "$xyz\n" }

```

Do note that a switch like --help creates the variable "\${-help}", which is not compliant with "use strict 'refs'". Also, when using this option on a script with warnings enabled you may get a lot of spurious "used only once" warnings.

- S makes Perl use the "PATH" environment variable to search for the program unless the name of the program contains path separators. On some platforms, this also makes Perl append suffixes to the filename while searching for it. For example, on Win32 platforms, the ".bat" and ".cmd" suffixes are appended if a lookup for the original name fails, and if the name does not already end in one of those suffixes. If your Perl was compiled with "DEBUGGING" turned on, using the -Dp switch to Perl shows how the search progresses.

Typically this is used to emulate "#!" startup on platforms that

don't support "#!". It's also convenient when debugging a script that uses "#!", and is thus normally found by the shell's \$PATH search mechanism.

This example works on many platforms that have a shell compatible with Bourne shell:

```
#!/usr/bin/perl
eval 'exec /usr/bin/perl -wS $0 ${1+"$@"}'
    if $running_under_some_shell;
```

The system ignores the first line and feeds the program to /bin/sh, which proceeds to try to execute the Perl program as a shell script. The shell executes the second line as a normal shell command, and thus starts up the Perl interpreter. On some systems \$0 doesn't always contain the full pathname, so the "-S" tells Perl to search for the program if necessary. After Perl locates the program, it parses the lines and ignores them because the variable \$running_under_some_shell is never true. If the program will be interpreted by csh, you will need to replace "\${1+"\$@"}" with \$*, even though that doesn't understand embedded spaces (and such) in the argument list. To start up sh rather than csh, some systems may have to replace the "#!" line with a line containing just a colon, which will be politely ignored by Perl. Other systems can't control that, and need a totally devious construct that will work under any of csh, sh, or Perl, such as the following:

```
eval '(exit $?0)' && eval 'exec perl -wS $0 ${1+"$@"}'
    & eval 'exec /usr/bin/perl -wS $0 $argv:q'
    if $running_under_some_shell;
```

If the filename supplied contains directory separators (and so is an absolute or relative pathname), and if that file is not found, platforms that append file extensions will do so and try to look for the file with those extensions added, one by one.

On DOS-like platforms, if the program does not contain directory separators, it will first be searched for in the current directory

before being searched for on the PATH. On Unix platforms, the program will be searched for strictly on the PATH.

- t Like "-T", but taint checks will issue warnings rather than fatal errors. These warnings can now be controlled normally with "no warnings qw(taint)".

Note: This is not a substitute for "-T"! This is meant to be used only as a temporary development aid while securing legacy code: for real production code and for new secure code written from scratch, always use the real "-T".

- T turns on "taint" so you can test them. Ordinarily these checks are done only when running setuid or setgid. It's a good idea to turn them on explicitly for programs that run on behalf of someone else whom you might not necessarily trust, such as CGI programs or any internet servers you might write in Perl. See perlsec for details. For security reasons, this option must be seen by Perl quite early; usually this means it must appear early on the command line or in the "#!" line for systems which support that construct.

- u This switch causes Perl to dump core after compiling your program. You can then in theory take this core dump and turn it into an executable file by using the undump program (not supplied). This speeds startup at the expense of some disk space (which you can minimize by stripping the executable). (Still, a "hello world" executable comes out to about 200K on my machine.) If you want to execute a portion of your program before dumping, use the "CORE::dump()" function instead. Note: availability of undump is platform specific and may not be available for a specific port of Perl.

- U allows Perl to do unsafe operations. Currently the only "unsafe" operations are attempting to unlink directories while running as superuser and running setuid programs with fatal taint checks turned into warnings. Note that warnings must be enabled along with this option to actually generate the taint-check warnings.

- v prints the version and patchlevel of your perl executable.
- V prints summary of the major perl configuration values and the current values of @INC.
- V:configvar

Prints to STDOUT the value of the named configuration variable(s), with multiples when your "configvar" argument looks like a regex (has non-letters). For example:

```
$ perl -V:libc
    libc='/lib/libc-2.2.4.so';

$ perl -V:lib.
    libs='-lnsl -lgdbm -ldb -ldl -lm -lcrypt -lutil -lc';
    libc='/lib/libc-2.2.4.so';

$ perl -V:lib.*
    libpth='/usr/local/lib /lib /usr/lib';
    libs='-lnsl -lgdbm -ldb -ldl -lm -lcrypt -lutil -lc';
    lib_ext='.a';
    libc='/lib/libc-2.2.4.so';
    libperl='libperl.a';
    ....
```

Additionally, extra colons can be used to control formatting. A trailing colon suppresses the linefeed and terminator ";\n", allowing you to embed queries into shell commands. (mnemonic: PATH separator ":".)

```
$ echo "compression-vars: " `perl -V:z.*: ` " are here !"
compression-vars: zcat=" zip='zip' are here !
```

A leading colon removes the "name=" part of the response, this allows you to map to the name you need. (mnemonic: empty label)

```
$ echo "goodvfork=" `./perl -llib -V::usevfork`
goodvfork=false;
```

Leading and trailing colons can be used together if you need positional parameter values without the names. Note that in the case below, the "PERL_API" params are returned in alphabetical order.

```
$ echo building_on `perl -V::osname: -V::PERL_API_*:` now
building_on 'linux' '5' '1' '9' now
```

`-w` prints warnings about dubious constructs, such as variable names mentioned only once and scalar variables used before being set; redefined subroutines; references to undefined filehandles; filehandles opened read-only that you are attempting to write on; values used as a number that don't look like numbers; using an array as though it were a scalar; if your subroutines recurse more than 100 deep; and innumerable other things.

This switch really just enables the global `^W` variable; normally, the lexically scoped "use warnings" pragma is preferred. You can disable or promote into fatal errors specific warnings using "`__WARN__`" hooks, as described in `perlvar` and "warn" in `perlfunc`. See also `perldiag` and `perltrap`. A fine-grained warning facility is also available if you want to manipulate entire classes of warnings; see `warnings`.

`-W` Enables all warnings regardless of "no warnings" or `^W`. See `warnings`.

`-X` Disables all warnings regardless of "use warnings" or `^W`. See `warnings`.

Forbidden in "`PERL5OPT`".

`-x`

`-xdirectory`

tells Perl that the program is embedded in a larger chunk of unrelated text, such as in a mail message. Leading garbage will be discarded until the first line that starts with "`#!`" and contains the string "perl". Any meaningful switches on that line will be applied.

All references to line numbers by the program (warnings, errors, ...) will treat the "`#!`" line as the first line. Thus a warning on the 2nd line of the program, which is on the 100th line in the file will be reported as line 2, not as line 100. This can be overridden by using the "`#line`" directive. (See "Plain Old

Comments (Not!) in perlsyn)

If a directory name is specified, Perl will switch to that directory before running the program. The `-x` switch controls only the disposal of leading garbage. The program must be terminated with `__END__` if there is trailing garbage to be ignored; the program can process any or all of the trailing garbage via the `"DATA"` filehandle if desired.

The directory, if specified, must appear immediately following the `-x` with no intervening whitespace.

ENVIRONMENT

HOME Used if `"chdir"` has no argument.

LOGDIR Used if `"chdir"` has no argument and `"HOME"` is not set.

PATH Used in executing subprocesses, and in finding the program if `"-S"` is used.

PERL5LIB A list of directories in which to look for Perl library files before looking in the standard library. Any architecture-specific and version-specific directories, such as `version/archname/`, `version/`, or `archname/` under the specified locations are automatically included if they exist, with this lookup done at interpreter startup time. In addition, any directories matching the entries in `$Config{inc_version_list}` are added. (These typically would be for older compatible perl versions installed in the same directory tree.)

If `PERL5LIB` is not defined, `"PERLLIB"` is used. Directories are separated (like in `PATH`) by a colon on Unixish platforms and by a semicolon on Windows (the proper path separator being given by the command `"perl -V:path_sep"`). When running taint checks, either because the program was running `setuid` or `setgid`, or the `"-T"` or `"-t"` switch was specified, neither `PERL5LIB` nor `"PERLLIB"` is consulted. The program should instead say:

```
use lib "/my/directory";
```

PERL5OPT Command-line options (switches). Switches in this variable are treated as if they were on every Perl command line. Only the `-[CDIMTUWdmtw]` switches are allowed. When running taint checks (either because the program was running `setuid` or `setgid`, or because the `-T` or `-t` switch was used), this variable is ignored. If `PERL5OPT` begins with `-T`, tainting will be enabled and subsequent options ignored. If `PERL5OPT` begins with `-t`, tainting will be enabled, a writable dot removed from `@INC`, and subsequent options honored.

PERLIO A space (or colon) separated list of PerlIO layers. If perl is built to use PerlIO system for IO (the default) these layers affect Perl's IO.

It is conventional to start layer names with a colon (for example, `":perlio"`) to emphasize their similarity to variable "attributes". But the code that parses layer specification strings, which is also used to decode the `PERLIO` environment variable, treats the colon as a separator.

An unset or empty `PERLIO` is equivalent to the default set of layers for your platform; for example, `":unix:perlio"` on Unix-like systems and `":unix:crlf"` on Windows and other DOS-like systems.

The list becomes the default for all Perl's IO.

Consequently only built-in layers can appear in this list, as external layers (such as `":encoding()"`) need IO in order to load them! See "open pragma" for how to add external encodings as defaults.

Layers it makes sense to include in the `PERLIO` environment variable are briefly summarized below. For more details see `PerlIO`.

`:crlf` A layer which does CRLF to `"\n"` translation distinguishing "text" and "binary" files in the

manner of MS-DOS and similar operating systems, and also provides buffering similar to ":perlio" on these architectures.

:perlio This is a re-implementation of stdio-like buffering written as a PerlIO layer. As such it will call whatever layer is below it for its operations, typically ":unix".

:stdio This layer provides a PerlIO interface by wrapping system's ANSI C "stdio" library calls. The layer provides both buffering and IO. Note that the ":stdio" layer does not do CRLF translation even if that is the platform's normal behaviour. You will need a ":crlf" layer above it to do that.

:unix Low-level layer that calls "read", "write", "lseek", etc.

:win32 On Win32 platforms this experimental layer uses native "handle" IO rather than a Unix-like numeric file descriptor layer. Known to be buggy in this release (5.30).

The default set of layers should give acceptable results on all platforms.

For Unix platforms that will be the equivalent of ":unix:perlio" or ":stdio". Configure is set up to prefer the ":stdio" implementation if the system's library provides for fast access to the buffer (not common on modern architectures); otherwise, it uses the ":unix:perlio" implementation.

On Win32 the default in this release (5.30) is ":unix:crlf". Win32's ":stdio" has a number of bugs/mis-features for Perl IO which are somewhat depending on the version and vendor of the C compiler. Using our own ":crlf" layer as the buffer avoids those issues and makes things more uniform.

This release (5.30) uses ":unix" as the bottom layer on Win32, and so still uses the C compiler's numeric file descriptor routines. There is an experimental native ":win32" layer, which is expected to be enhanced and may eventually become the default under Win32.

The PERLIO environment variable is completely ignored when Perl is run in taint mode.

PERLIO_DEBUG

If set to the name of a file or device when Perl is run with the -Di command-line switch, the logging of certain operations of the PerLIO subsystem will be redirected to the specified file rather than going to stderr, which is the default. The file is opened in append mode. Typical uses are in Unix:

```
% env PERLIO_DEBUG=/tmp/perlio.log perl -Di script ...
```

and under Win32, the approximately equivalent:

```
> set PERLIO_DEBUG=CON  
perl -Di script ...
```

This functionality is disabled for setuid scripts, for scripts run with "-T", and for scripts run on a Perl built without "-DDEBUGGING" support.

PERLLIB A list of directories in which to look for Perl library files before looking in the standard library. If "PERL5LIB" is defined, PERLLIB is not used.

The PERLLIB environment variable is completely ignored when Perl is run in taint mode.

PERL5DB The command used to load the debugger code. The default is:

```
BEGIN { require "perl5db.pl" }
```

The PERL5DB environment variable is only used when Perl is started with a bare "-d" switch.

PERL5DB_THREADED

If set to a true value, indicates to the debugger that the

code being debugged uses threads.

PERL5SHELL (specific to the Win32 port)

On Win32 ports only, may be set to an alternative shell that Perl must use internally for executing "backtick" commands or `system()`. Default is `"cmd.exe /x/d/c"` on WindowsNT and `"command.com /c"` on Windows95. The value is considered space-separated. Precede any character that needs to be protected, like a space or backslash, with another backslash.

Note that Perl doesn't use COMSPEC for this purpose because COMSPEC has a high degree of variability among users, leading to portability concerns. Besides, Perl can use a shell that may not be fit for interactive use, and setting COMSPEC to such a shell may interfere with the proper functioning of other programs (which usually look in COMSPEC to find a shell fit for interactive use).

Before Perl 5.10.0 and 5.8.8, PERL5SHELL was not taint checked when running external commands. It is recommended that you explicitly set (or delete) `$ENV{PERL5SHELL}` when running in taint mode under Windows.

PERL_ALLOW_NON_IFS_LSP (specific to the Win32 port)

Set to 1 to allow the use of non-IFS compatible LSPs (Layered Service Providers). Perl normally searches for an IFS-compatible LSP because this is required for its emulation of Windows sockets as real filehandles. However, this may cause problems if you have a firewall such as McAfee Guardian, which requires that all applications use its LSP but which is not IFS-compatible, because clearly Perl will normally avoid using such an LSP.

Setting this environment variable to 1 means that Perl will simply use the first suitable LSP enumerated in the catalog, which keeps McAfee Guardian happy--and in that particular case Perl still works too because McAfee

Guardian's LSP actually plays other games which allow applications requiring IFS compatibility to work.

PERL_DEBUG_MSTATS

Relevant only if Perl is compiled with the "malloc" included with the Perl distribution; that is, if "perl -V:d_mymalloc" is "define".

If set, this dumps out memory statistics after execution.

If set to an integer greater than one, also dumps out memory statistics after compilation.

PERL_DESTRUCT_LEVEL

Controls the behaviour of global destruction of objects and other references. See "PERL_DESTRUCT_LEVEL" in perlhacktips for more information.

PERL_DL_NONLAZY

Set to "1" to have Perl resolve all undefined symbols when it loads a dynamic library. The default behaviour is to resolve symbols when they are used. Setting this variable is useful during testing of extensions, as it ensures that you get an error on misspelled function names even if the test suite doesn't call them.

PERL_ENCODING

If using the "use encoding" pragma without an explicit encoding name, the PERL_ENCODING environment variable is consulted for an encoding name.

PERL_HASH_SEED

(Since Perl 5.8.1, new semantics in Perl 5.18.0) Used to override the randomization of Perl's internal hash function. The value is expressed in hexadecimal, and may include a leading 0x. Truncated patterns are treated as though they are suffixed with sufficient 0's as required.

If the option is provided, and "PERL_PERTURB_KEYS" is NOT set, then a value of '0' implies "PERL_PERTURB_KEYS=0" and any other value implies "PERL_PERTURB_KEYS=2".

PLEASE NOTE: The hash seed is sensitive information. Hashes are randomized to protect against local and remote attacks against Perl code. By manually setting a seed, this protection may be partially or completely lost. See "Algorithmic Complexity Attacks" in perlsec, "PERL_PERTURB_KEYS", and "PERL_HASH_SEED_DEBUG" for more information.

PERL_PERTURB_KEYS

(Since Perl 5.18.0) Set to "0" or "NO" then traversing keys will be repeatable from run to run for the same "PERL_HASH_SEED". Insertion into a hash will not change the order, except to provide for more space in the hash. When combined with setting PERL_HASH_SEED this mode is as close to pre 5.18 behavior as you can get.

When set to "1" or "RANDOM" then traversing keys will be randomized. Every time a hash is inserted into the key order will change in a random fashion. The order may not be repeatable in a following program run even if the PERL_HASH_SEED has been specified. This is the default mode for perl.

When set to "2" or "DETERMINISTIC" then inserting keys into a hash will cause the key order to change, but in a way that is repeatable from program run to program run.

NOTE: Use of this option is considered insecure, and is intended only for debugging non-deterministic behavior in Perl's hash function. Do not use it in production.

See "Algorithmic Complexity Attacks" in perlsec and "PERL_HASH_SEED" and "PERL_HASH_SEED_DEBUG" for more information. You can get and set the key traversal mask for a specific hash by using the "hash_traversal_mask()" function from Hash::Util.

PERL_HASH_SEED_DEBUG

(Since Perl 5.8.1.) Set to "1" to display (to STDERR)

information about the hash function, seed, and what type of key traversal randomization is in effect at the beginning of execution. This, combined with "PERL_HASH_SEED" and "PERL_PERTURB_KEYS" is intended to aid in debugging nondeterministic behaviour caused by hash randomization.

Note that any information about the hash function, especially the hash seed is sensitive information: by knowing it, one can craft a denial-of-service attack against Perl code, even remotely; see "Algorithmic Complexity Attacks" in perlsec for more information. Do not disclose the hash seed to people who don't need to know it. See also "hash_seed()" and "hash_traversal_mask()".

An example output might be:

```
HASH_FUNCTION = ONE_AT_A_TIME_HARD HASH_SEED = 0x652e9b9349a7a032 PERTURB_KEYS = 1
```

(RANDOM)

PERL_MEM_LOG

If your Perl was configured with `-Accflags=-DPERL_MEM_LOG`, setting the environment variable "PERL_MEM_LOG" enables logging debug messages. The value has the form "`<number>[m][s][t]`", where "number" is the file descriptor number you want to write to (2 is default), and the combination of letters specifies that you want information about (m)emory and/or (s)v, optionally with (t)imestamps. For example, "PERL_MEM_LOG=1mst" logs all information to stdout. You can write to other opened file descriptors in a variety of ways:

```
$ 3>foo3 PERL_MEM_LOG=3m perl ...
```

PERL_ROOT (specific to the VMS port)

A translation-concealed rooted logical name that contains Perl and the logical device for the @INC path on VMS only. Other logical names that affect Perl on VMS include PERLSHR, PERL_ENV_TABLES, and SYS\$TIMEZONE_DIFFERENTIAL, but are optional and discussed further in perlvms and in

README.vms in the Perl source distribution.

PERL_SIGNALS

Available in Perls 5.8.1 and later. If set to "unsafe", the pre-Perl-5.8.0 signal behaviour (which is immediate but unsafe) is restored. If set to "safe", then safe (but deferred) signals are used. See "Deferred Signals (Safe Signals)" in perlipc.

PERL_UNICODE

Equivalent to the -C command-line switch. Note that this is not a boolean variable. Setting this to "1" is not the right way to "enable Unicode" (whatever that would mean). You can use "0" to "disable Unicode", though (or alternatively unset PERL_UNICODE in your shell before starting Perl). See the description of the -C switch for more information.

PERL_USE_UNSAFE_INC

If perl has been configured to not have the current directory in @INC by default, this variable can be set to "1" to reinstate it. It's primarily intended for use while building and testing modules that have not been updated to deal with "." not being in @INC and should not be set in the environment for day-to-day use.

SYS\$LOGIN (specific to the VMS port)

Used if chdir has no argument and "HOME" and "LOGDIR" are not set.

PERL_INTERNAL_RAND_SEED

Set to a non-negative integer to seed the random number generator used internally by perl for a variety of purposes.

Ignored if perl is run setuid or setgid. Used only for some limited startup randomization (hash keys) if "-T" or "-t" perl is started with tainting enabled.

Perl may be built to ignore this variable.

Perl also has environment variables that control how Perl handles data specific to particular natural languages; see `perllocale`.

Perl and its various modules and components, including its test frameworks, may sometimes make use of certain other environment variables. Some of these are specific to a particular platform.

Please consult the appropriate module documentation and any documentation for your platform (like `perlsolaris`, `perllinux`, `perlmacosx`, `perlwin32`, etc) for variables peculiar to those specific situations.

Perl makes all environment variables available to the program being executed, and passes these along to any child processes it starts.

However, programs running `setuid` would do well to execute the following lines before doing anything else, just to keep people honest:

```
$ENV{PATH} = "/bin:/usr/bin"; # or whatever you need
$ENV{SHELL} = "/bin/sh" if exists $ENV{SHELL};
delete @ENV{qw(IFS CDPATH ENV BASH_ENV)};
```

ORDER OF APPLICATION

Some options, in particular `-I`, `-M`, `PERL5LIB` and `PERL5OPT` can interact, and the order in which they are applied is important.

Note that this section does not document what actually happens inside the perl interpreter, it documents what effectively happens.

`-I` The effect of multiple `-I` options is to "unshift" them onto `@INC` from right to left. So for example:

```
perl -I 1 -I 2 -I 3
```

will first prepend 3 onto the front of `@INC`, then prepend 2, and then prepend 1. The result is that `@INC` begins with:

```
qw(1 2 3)
```

`-M` Multiple `-M` options are processed from left to right. So this:

```
perl -Mlib=1 -Mlib=2 -Mlib=3
```

will first use the lib pragma to prepend 1 to `@INC`, then it will prepend 2, then it will prepend 3, resulting in an `@INC` that begins with:

```
qw(3 2 1)
```

the PERL5LIB environment variable

This contains a list of directories, separated by colons. The entire list is prepended to @INC in one go. This:

```
PERL5LIB=1:2:3 perl
```

will result in an @INC that begins with:

```
qw(1 2 3)
```

combinations of -I, -M and PERL5LIB

"PERL5LIB" is applied first, then all the "-I" arguments, then all the "-M" arguments. This:

```
PERL5LIB=e1:e2 perl -I i1 -Mlib=m1 -I i2 -Mlib=m2
```

will result in an @INC that begins with:

```
qw(m2 m1 i1 i2 e1 e2)
```

the PERL5OPT environment variable

This contains a space separated list of switches. We only consider the effects of "-M" and "-I" in this section.

After normal processing of "-I" switches from the command line, all the "-I" switches in "PERL5OPT" are extracted. They are processed from left to right instead of from right to left. Also note that while whitespace is allowed between a "-I" and its directory on the command line, it is not allowed in "PERL5OPT".

After normal processing of "-M" switches from the command line, all the "-M" switches in "PERL5OPT" are extracted. They are processed from left to right, i.e. the same as those on the command line.

An example may make this clearer:

```
export PERL5OPT="-Mlib=optm1 -lopti1 -Mlib=optm2 -lopti2"
```

```
export PERL5LIB=e1:e2
```

```
perl -I i1 -Mlib=m1 -I i2 -Mlib=m2
```

will result in an @INC that begins with:

```
qw(  
  optm2  
  optm1  
  m2  
  m1
```


opti2

opti1

i1

i2

e1

e2

)

Other complications

There are some complications that are ignored in the examples

above:

arch and version subdirs

All of "-I", "PERL5LIB" and "use lib" will also prepend arch

and version subdirs if they are present

sitecustomize.pl

perl v5.32.1

2023-01-18

PERLRUN(1)