

NAME

etsh - enhanced Thompson shell (command interpreter)

SYNOPSIS

etsh [-V | -VV]

etsh [-nv] [- | -c [*string*] | -i | -l | -t |file [*arg1* ...]]

DESCRIPTION

Etsh is an enhanced, backward-compatible port of the standard command interpreter from Sixth Edition (V6) UNIX. It may be used either as an interactive shell or as a non-interactive shell. Throughout this manual, '(+)' indicates those cases where **etsh** is known to differ from the original *sh*(1), as it appeared in Sixth Edition (V6) UNIX.

The options are as follows:

- The shell reads and executes command lines from the standard input until end-of-file or **exit**.
- c [*string*]
If a *string* is specified, the shell executes it as a command line and exits. Otherwise, the shell treats it as the - option.
- i (+) The shell behaves as an interactive shell by reading and executing commands from the appropriate rc files if possible (see *Startup and shutdown* below) before prompting the user, reading, and executing command lines from the standard input. The shell prints a diagnostic and exits with a non-zero status if it is not connected to a terminal.
- l (+) The shell behaves as a login shell by reading and executing commands from the appropriate rc files if possible (see *Startup and shutdown* below) before prompting the user, reading, and executing command lines from the standard input. The shell prints a diagnostic and exits with a non-zero status if it is not connected to a terminal.
- n (+) *noexec*: The shell performs no command-line execution. It only checks the syntax of each command line, after performing parameter substitution and word splitting. The shell ignores this option for interactive and login shells.
- t The shell reads a single command line from the standard input, executes it, and exits.
- v (+) *verbose*: The shell prints the words of each command line to the standard error, after performing parameter substitution and word splitting, but before executing the resulting command line.

-V (+) The shell prints its version from $\$v$ to the standard output and exits.

-VV

(+) The shell prints its version from $\$v$, and build-time system information from ‘uname -srm’ (see *uname(1)*), to the standard output and exits.

The shell may also be invoked non-interactively to read, interpret, and execute a command file. The specified *file* and any arguments are treated as positional parameters (see *Parameter substitution* below) during execution of the command file.

Otherwise, if no arguments except for **-v** are specified and if both the standard input and standard error are connected to a terminal, the shell is interactive. By default, an interactive shell prompts a regular user with a ‘% ’ or with a ‘# ’ for the superuser before reading each command line from the terminal.

(+) Notice that the user can set a non-default shell prompt if desired (see *Prompt* below).

(+) When an interactive shell starts, it reads and executes commands from the appropriate rc files if possible (see *Startup and shutdown* below) before reading and executing command lines from the terminal.

Metacharacters

A *syntactic metacharacter* is any one of the following:

| ^ ; & () < > space tab

When such a character is unquoted, it has special meaning to the shell. The shell uses it to separate words (see *Commands* and *Command lines* below). A *quoting metacharacter* is any one of the following:

" ' \

See *Quoting* below. The *substitution metacharacter* is a:

\$

See *Parameter substitution* and *Variable substitution* below. Finally, a *pattern metacharacter* is any one of the following:

* ? [

See *File name generation* below.

Commands

Each command is a sequence of non-blank command arguments, or words, separated by one or more blanks (**spaces** or **tabs**). The first argument specifies the name of a command to be executed. Except for certain special arguments described below, the arguments other than the command name are passed without interpretation to the invoked command.

If the first argument names a special command, the shell executes it (see *Special commands* below).

(+) Otherwise, the shell treats it either as an alias (see *Aliases* below) or as an external command, which is located as follows.

(+) If the command name contains no `'/'` characters, the sequence of directories in the environment variable `PATH` is searched for the first occurrence of an executable file by that name, which the shell attempts to execute. However, if the command name contains one or more `'/'` characters, the shell attempts to execute it without performing any `PATH` search.

(+) If an executable file does not begin with the proper magic number or a `#!/shell` sequence, it is assumed to be a shell command file, and a new shell is automatically invoked to execute it. The environment variable `EXECShell` specifies the shell which is invoked to execute such a file.

If a command cannot be found or executed, a diagnostic is printed.

Command lines

Commands separated by `|` or `^` constitute a chain of *filters*, or a *pipeline*. The standard output of each command but the last is taken as the standard input of the next command. Each command is run as a separate process, connected by pipes (see *pipe(2)*) to its neighbors.

A *command line*, or *list*, consists of one or more pipelines separated, and perhaps terminated by `;` or `&`. The semicolon designates sequential execution. The ampersand designates asynchronous execution, which causes the preceding pipeline to be executed without waiting for it to finish. The process ID of each command in such a pipeline is reported, so that it may be used if necessary for a subsequent *kill(1)*.

A list contained within parentheses such as `(list)` is executed in a subshell and may appear in place of a simple command as a filter.

If a command line is syntactically incorrect, a diagnostic is printed.

Termination reporting

All terminations other than `exit` and `interrupt` are considered to be abnormal. If a sequential process terminates abnormally, a message is printed. The termination report for an asynchronous process is

given upon execution of the first sequential command subsequent to its termination, or when the **wait** special command is executed. The following is a list of the possible termination messages:

- Hangup
- Quit
- Illegal instruction
- Trace/BPT trap
- IOT trap
- EMT trap
- Floating exception
- Killed
- Bus error
- Memory fault
- Bad system call
- Broken pipe

For an asynchronous process, its process ID is prepended to the appropriate message. If a core image is produced, ‘ -- Core dumped’ is appended to the appropriate message.

I/O redirection

Each of the following argument forms is interpreted as a *redirection* by the shell itself. Such a redirection may appear anywhere among the arguments of a simple command, or before or after a parenthesized command list, and is associated with that command or command list.

A redirection of the form `<arg` causes the file *arg* to be used as the standard input (file descriptor 0) for the associated command.

A redirection of the form `>arg` causes the file *arg* to be used as the standard output (file descriptor 1) for the associated command. If *arg* does not already exist, it is created; otherwise, it is truncated at the outset.

A redirection of the form `>>arg` is the same as `>arg`, except if *arg* already exists the command output is always appended to the end of the file.

For example, either of the following command lines:

```
% date >index.txt ; pwd >>index.txt ; ls -l >>index.txt
% ( date ; pwd ; ls -l ) >index.txt
```

creates on the file ‘index.txt’, the current date and time, followed by the name and a long listing of the

current working directory.

(+) A `<-` redirection causes input for the associated command to be redirected from the standard input which existed when the shell was invoked. This allows a command file to be used as a filter.

A `>arg` or `>>arg` redirection associated with any but the last command of a pipeline is ineffectual, as is a `<arg` redirection with any but the first.

The standard error (file descriptor 2) is never subject to redirection by the shell itself. Thus, commands may write diagnostics to a location where they have a chance to be seen. However, **fd2** provides a way to redirect the diagnostic output to another location.

If the file for a redirection cannot be opened or created, a diagnostic is printed.

Quoting

The shell treats all *single* (') and *backslash* (\) *quoted* characters literally, including characters which have special meaning to the shell (see *Metacharacters* above). If such characters are quoted, they represent themselves and may be passed as part of arguments.

(+) Like the quoting behavior described above, *double* (") quotes cause the shell to treat characters literally. However, double quotes also allow the shell to perform parameter and variable substitution via the dollar (\$) metacharacter, whereas *single* (') quotes and *backslash* (\) quotes do not.

Individual characters, and sequences of characters, are quoted when enclosed by a matched pair of *double* (") or *single* (') quotes. For example:

```
% awk '{ print NR "\t" $0 }' README ^ more
```

causes *awk*(1) to write each line from the 'README' file, preceded by its line number and a tab, to the standard output which is piped to *more*(1) for viewing. The outer single quotes prevent the shell from trying to interpret any part of the string, which is then passed as a single argument to *awk*.

An individual *backslash* (\) quotes, or *escapes*, the next individual character. A backslash followed by a newline is a special case which allows continuation of command-line input onto the next line. Each backslash-newline sequence in the input is translated into a blank.

If a double or single quote appears but is not part of a matched pair, a diagnostic is printed.

Parameter substitution

When the shell is invoked with arguments besides `-v`, it has additional string processing capabilities

which are not otherwise available. Such a shell may be invoked as follows:

```
etsh [-v] name [arg1 ...]
```

If the first character of *name* is not -, it is taken as the name of a *command file*, or *shell script*, which is opened as the standard input for a new shell instance. Thus, the new shell reads and interprets command lines from the named file.

Otherwise, *name* is taken as one of the shell options, and a new shell instance is invoked to read and interpret command lines from its standard input. However, notice that the **-c** option followed by a *string* is the one case where the shell does not read and interpret command lines from its standard input. Instead, the string itself is taken as a command line and executed.

In each command line, an unquoted character sequence of the form $\$N$, where N is a digit, is treated as a *positional parameter* by the shell. Each occurrence of a positional parameter in the command line is substituted with the value of the N th argument to the invocation of the shell (*argN*). $\$0$ is substituted with *name*.

In all shell instances, $\$\$$ is substituted with the process ID of the current shell. The value is represented as a 5-digit ASCII string, padded on the left with zeros when the process ID is less than 10000.

(+) All shell instances attempt to set the special parameters in the following list. ‘(*)’ indicates those which are always set. Otherwise, the parameter is unset when the shell cannot determine its value.

$\$\#$ (*) The number of positional parameters currently available to the shell.

$\$*$ The values of the positional parameters currently available to the shell, from $\$1$ through the end of its argument list.

$\$\?$ (*) The exit status of the last sequential command from the *previous* command line.

$\$d$ The value of the environment variable ETSHDIR.

$\$e$ The value of the environment variable EXECHELL.

$\$g$ The effective group name of the current user, as determined by *getgrgid(3)*. The value (if any) is equivalent to that given by ‘id -gn’.

$\$h$ The value of the environment variable HOME.

- \$i** (*) The effective group ID of the current user, as determined by *getegid*(2). The value is equivalent to that given by 'id -g'.
- \$k** (*) The effective user ID of the current user, as determined by *geteuid*(2). The value is equivalent to that given by 'id -u'.
- \$m** The value of the environment variable MANPATH.
- \$p** The value of the environment variable PATH.
- \$t** The terminal name with which the standard input was associated when the shell was invoked, as determined by *ttyname*(3). The value (if any) is equivalent to that given by 'tty <-'.
- \$u** The effective user name of the current user, as determined by *getpwuid*(3). The value (if any) is equivalent to that given by 'id -un'.
- \$v** (*) The version of the current shell represented as a one-word, read-only string.
- \$w** The value of the environment variable CWD. It is the absolute physical path name of the current working directory, without symbolic links, as determined by *getcwd*(3). On error, the value of \$w is set to the empty string.

All substitution on a command line is performed *before* the line is interpreted. Thus, no action which alters the value of any parameter can have any effect on a reference to that parameter occurring on the *same* line.

A positional-parameter value may contain any number of metacharacters. Each one which is *unquoted*, or *unescaped*, within a positional-parameter value retains its special meaning when the value is substituted in a command line by the invoked shell.

Take the following two shell invocations for example:

```
% etsh -c '$1' 'echo Hello World! >/dev/null'
% etsh -c '$1' 'echo Hello World! \>/dev/null'
Hello World! >/dev/null
```

In the first invocation, the > in the value substituted by \$1 retains its special meaning. This causes all output from **echo** to be redirected to */dev/null*. However, in the second invocation, the meaning of > is *escaped* by \ in the value substituted by \$1. This causes the shell to pass '>/dev/null' as a single

argument to echo instead of interpreting it as a redirection.

Variable substitution (+)

The shell can substitute variables; a user may cause the shell to unset and set variables by using the **unset** and **set** special commands.

Variables may be used both in interactive shells and in non-interactive shells. However, notice that variables are not functional when a non-interactive shell is invoked either with the **-c** option followed by a *string* or with the **-t** option. Such a shell only executes one command line, but setting and using a variable requires executing two command lines in the same shell, one to set it and one to use it.

A variable can either be unset or set. When unset, a variable has no *value*. When set, a variable has both *name* and *value*. A valid variable *name* is a single ASCII character, which matches either the [A-Z] range or the [a-cfjlnq-sx-z] range, inclusive. A valid variable *value* can be anything from an empty string, denoted by "" or '', to whatever you can input into a syntactically correct command line with less than or equal to the maximum command line length of characters, 2048 for the sake of brevity.

For example, 'set C value' sets the variable C to value, as you can see in the following examples.

```
% : Example One
% unset C
% set C
% if $? -eq 1 -a "$C" = "C" echo 'C is unset.'
C is unset.

% : Example Two
% set C ''
% ( set C ) >/dev/null
% if $? -eq 0 -a -z "$C" echo 'C == ""$C""'
C == ""
% set W "Hello "
% set W "$WWorld!"
% if "$W" != "W" -a -n "$W" echo "$W"
Hello World!

% : Example Three
% alias now "date '+%A, %Y-%m-%d, %T %Z';;"
% alias loadavg "echo -n $H': '\
uptime|sed 's/^. *user[s,][,] */';;"
% set C '( now ; loadavg )' ; : 'C == Command Line (or List)'
```



```
% ( set C ) >/dev/null
% if $? -eq 0 -a -n "$C" echo "C == '$C'"
C == '( now ; loadavg )'
% $C | awk '{ print NR "\t" $0 }'
1      Saturday, 2018-02-24, 21:52:03 UTC
2      refugio: load averages: 0.54, 0.22, 0.08
```

As with parameters (see *Parameter substitution* above), all substitution on a command line is performed *before* the line is interpreted. Thus, no action which alters the value of any variable can have any effect on a reference to that variable occurring on the *same* line.

Also, a variable value may contain any number of metacharacters. Each one which is *unquoted*, or *unesaped*, within a variable value retains its special meaning when the value is substituted in a command line.

If a variable name passed as an argument to **set** or **unset** is invalid, a diagnostic is printed. Similarly, if a variable value causes an error, a diagnostic is printed.

File name generation

Prior to executing a command, the shell scans each argument for unquoted `*`, `?`, or `[` characters. If one or more of these characters appears, the argument is treated as a *pattern* and causes the shell to search for file names which *match* it. Otherwise, the argument is used as is.

The meaning of each pattern character is as follows:

- o The `*` character in a pattern matches any string of characters in a file name (including the null string).
- o The `?` character in a pattern matches any single character in a file name.
- o The `[...]` brackets in a pattern specifies a class of characters which matches any single file-name character in the class. Within the brackets, each character is taken to be a member of the class. A pair of characters separated by an unquoted `-` specifies the class as a range which matches each character lexically between the first and second member of the pair, inclusive. A `-` matches itself when quoted or when first or last in the class.

Any other character in a pattern matches itself in a file name.

Notice that the `'.'` character at the beginning of a file name, or immediately following a `'/'`, is always special in that it must be matched explicitly. The same is true of the `'/'` character itself.

If the pattern contains no `'/'` characters, the current directory is always used. Otherwise, the specified directory is the one obtained by taking the pattern up to the last `'/'` before the first unquoted `*`, `?`, or `[`. The matching process matches the remainder of the pattern after this `'/'` against the files in the specified directory.

In any event, a list of file names is obtained from the current (or specified) directory which match the given pattern. This list is sorted in ascending ASCII order, and the new sequence of arguments replaces the given pattern. The same process is carried out for each of the given pattern arguments; the resulting lists are *not* merged. Finally, the shell attempts to execute the command with the resulting argument list.

If a pattern argument refers to a directory which cannot be opened, a `'No directory'` diagnostic is printed.

If a command has only *one* pattern argument, a `'No match'` diagnostic is printed if it fails to match any files. However, if a command has more than one pattern argument, a diagnostic is printed only when they *all* fail to match any files. Otherwise, each pattern argument failing to match any files is removed from the argument list.

Startup and shutdown (+)

If the first character of the `argv[0]` used to invoke an interactive shell is `'-'` (e.g., `-etsh`), it is a login shell and tries to read and execute commands from the following four rc init files in sequence: `/usr/local/etc/etsh.login`, `/usr/local/etc/etsh.etchrc`, `$h/.etsh.login`, and `$h/.etchrc`. The same is true when the shell is invoked with the `-l` option, regardless of the value of `argv[0]`.

In the case where an interactive shell is not a login shell according to its `argv[0]`, it tries to read and execute commands from the following two rc init files in sequence: `/usr/local/etc/etsh.etchrc` and `$h/.etchrc`. The same is true when the shell is invoked with the `-i` option, regardless of the value of `argv[0]`.

In any case, after the shell finishes its startup actions, it then prompts the user, reads, and executes command lines from the standard input as usual.

If the shell is invoked as a login shell, it tries to read and execute commands from `/usr/local/etc/etsh.logout` and `$h/.etsh.logout` in sequence upon logout. These two rc logout files may be used if necessary for cleanup upon termination of a login session by an EOT (see *End of file* below) or a SIGHUP signal (see *Signals* below).

Notice that the shell only performs the startup and shutdown actions described above for readable, regular rc files. If any rc file is *not* readable, the shell ignores it and continues as normal. If any rc file

is *not* a regular file (or a link to a regular file), the shell ignores it, prints a diagnostic, and continues as normal.

In the normal case, a SIGINT or SIGQUIT signal received by the shell during execution of any rc file causes it to cease execution of that file without terminating. Thus, it may be desirable to use the **trap** special command to ignore these and other signals in some cases. For example, this is particularly true for */usr/local/etc/etsh.login*, */usr/local/etc/etsh.etshrc*, and */usr/local/etc/etsh.logout*.

The **exit** special command always causes the shell to terminate if it occurs in any rc file.

History (+)

If the shell is invoked as an interactive shell, it tries to open the *\$h/.etsh.history* file to save the user's command-line history. Notice that the history file must already exist as a writable, regular file (or a link to a regular file) when the shell is invoked to save the user's command-line history. Otherwise, it will not do so.

An interactive shell reads each command line from its terminal and appends the words of each one to the history file as a line after performing parameter substitution and word splitting.

The shell does not read the history file or have any features that allow the user to make direct use of the saved history. Such features are available via standard external commands and also via the **history** command found in the */usr/local/libexec/etsh-5.3.1/etsh* directory. Execute 'history -h' to read its documentation.

Notice that the shell never creates or removes the *\$h/.etsh.history* file. It always leaves these actions to the user. For example:

```
% history -r ; history -c ; exec etsh -l
```

causes **history** to remove the existing history file (if any), to create a new (empty) one, and causes the current shell to replace itself with a new login shell, while opening the new history file. This, and future, interactive shells then save the user's command-line history as long as the history file exists.

If desired, the user can use the history file to repeat any command line as a command substitution with *sed(1)* and **etsh**. Taking the following command line and history entry for example:

```
% history -n 6171
Number      Command Line
-----
6171      uname -s | if { fd2 -ef/dev/null \
```

```
egrep '([ONF][a-z]{2,3}BSD|Darwin|Linux)' } \
echo '(Open|Net|Free)BSD || (Mac) OS X || OS == GNU/Linux'
```

and then doing a:

```
% sed -n 6171p <${h}/.etsh.history | etsh
OS == GNU/Linux || (Mac) OS X || (Free|Net|Open)BSD
(Open|Net|Free)BSD || (Mac) OS X || OS == GNU/Linux
```

causes sed to output the 6171st command line from the history file via pipe for repetition as a command substitution by etsh.

Aliases (+)

The shell can interpret command aliases set by the user. A user may cause the shell to set, print, and unset command aliases by using the **alias** and **unalias** special commands.

A command alias is a string that substitutes for a given command alias name set by the user. Command aliases provide a simple way to represent complex, long, or often-used commands as simple command names. Thus, if the first argument names an existing command alias, its alias string substitutes for the command alias name. Any remaining arguments are appended to the argument list.

Aliases may be used both in interactive shells and in non-interactive shells. However, notice that aliases are not functional when a non-interactive shell is invoked either with the **-c** option followed by a *string* or with the **-t** option. Such a shell only executes one command line, but setting and using an alias requires executing two command lines in the same shell, one to set it and one to execute it.

The shell parses each alias in a command line into a list of words from left to right, wraps it as a (*list*), re-parses it while parsing any nested aliases (up to three deep), and executes the resulting alias in a subshell on success. Three examples of alias usage follow:

```
% : Example One
% alias s 'echo $?;:' ; alias status 's'
% alias s ; alias status
(echo $?;:)
(s)
% false
% status
1

% : Example Two
```

```
% alias ll 'ls -AIF'
% alias ll
(ls -AIF)
% ll -d [A-Z]* | wc -l | tr -d ' \t'
10

% : Example Three
% alias loadavg "uname -n|sed 's/^([^.]*).*^1/'|tr -d \\n;\
echo -n ': ';uptime|sed 's/^.*user[s,][,] */';;"
% alias loadavg
(uname -n|sed 's/^([^.]*).*^1/'|tr -d \\n;\
echo -n ': ';uptime|sed 's/^.*user[s,][,] */';;)
% loadavg | awk '{ print NR "\t" $0 }'
1      serenity: load average: 0.49 0.39 0.29
```

If an alias, or its parsed result in a command line, is syntactically incorrect, a diagnostic is printed. If an alias loop error occurs, a ‘Too many nested aliases’ diagnostic is printed.

Prompt (+)

The state of the **P** *variable* determines how the shell prompt appears. The user can return to the default prompt by executing ‘unset P’ whenever desired. Since P is a variable, see *Variable substitution* above, as the same documentation applies here too.

Notice that ‘set P string’ sets the shell prompt to string. See **SetP** in the */usr/local/libexec/etsh-5.3.1/etsh* directory. You can **source** or **.** it if you wish to use it.

Four examples of setting your prompt follow:

```
% : Example One
% : " This is the default prompt, but let's " ; unset P ;\
: " it to be sure. I am not root now, as you can see. "
%

% : Example Two
% : " OK, let's see others. "
% set P '=^\) hello \(\^=|>% '
=^) hello (^=>%

% : Example Three
=^) hello (^=>% : " Silly! Others? OK. "
```

```
=^) hello (^=>% . SetP "$u:p2\>"
jneitzel:p2>%
```

```
% : Example Four
jneitzel:p2>% : " Or with $w and 2 more lines? OK. "
jneitzel:p2>%      . $h/.etsh.prompt
```

```
~/src/git/v6shell
jneitzel@refugio
p2>%
```

```
~/src/git/v6shell
jneitzel@refugio
p2>% : " In effect, the '. $h/.etsh.prompt' above prepares \
      things before doing a special '. SetP string' . "
```

```
~/src/git/v6shell
jneitzel@refugio
p2>%
```

Notice that an *\$h/.etsh.prompt* file is not required in order to set and use a personal etsh prompt. That said, it can make things simpler or easier, in the same way that having an *\$h/.etshrc* file can do so. In any case, I trust you noticed that "Example Two" above did not mention or use *\$h/.etsh.prompt* at all; it was completely manual.

Now, while the included *.etsh.prompt* file is admittedly too complex with 74 lines, a simpler and shorter file can be quite effective if your prompt objective is more focused than mine is, or was, when I wrote it.

Documenting all of the appropriate bits here in the manual was quite a challenge. Hopefully, I have not added unnecessary complexity and confusion. If you find room for improvement, I invite you to send me your suggestions via email (see **AUTHORS** below for my email address).

End of file

An end-of-file in the shell's input causes it to exit. If the shell is interactive, this means it exits by default when the user types an EOT (^D) at the prompt. If desired, the user may change or disable interactive shell EOT exit behavior with *stty(1)*.

Special commands

The shell treats the following built-in commands specially.

: [*arg ...*]

Does nothing and sets the exit status to zero.

alias [*name [string]*] (+)

Sets the alias *name* to *string* in the current shell if both *name* and *string* are specified. Prints the current value of the alias *name*'s string if *name* is specified and set in the current shell. Prints the name and string of each alias set in the current shell if no arguments are specified.

cd [*dir ...*] (+)

Is a synonym for the **chdir** special command.

chdir [*dir ...*]

Changes the shell's current working directory to *dir*. (+) If *dir* is an unquoted -, the shell's previous working directory is used instead. Otherwise, if *dir* is not specified, the user's home directory is used by default.

echo [-**n**] [*string ...*] (+)

Writes its string arguments (if any) separated by blanks and terminated by a newline to the standard output. If **-n** is specified, the terminating newline is not written.

exec *command* [*arg ...*] (+)

Replaces the current shell with an instance of *command*, which must be external to the shell.

exit Causes the shell to cease execution of a file. This means exit has no effect at the prompt of an interactive shell.

fd2 [-**e**] [-**f** *file*] [--] *command* [*arg ...*] (+)

Redirects from/to file descriptor 2 for *command*. See the *fd2(1)* manual page for full details.

goto *label* [...] (+)

Transfers shell control to the '*:label*' line of the current command file. See the *goto(1)* manual page for full details.

history [-**c** | -**h** | -**l** | -**n** *number1* [, *number2*] | -**p** *pattern* | -**r**] (+)

Manages, prints, and searches the user's *\$h/.etsh.history* file. If no options are specified, the history command prints all history entries from the user's etsh history file to the standard output. Execute 'history -h' to read its documentation.

if [*expression [command [arg ...]]*] (+)

Evaluates *expression*, conditionally executes *command*, and sets the exit status to zero or non-zero

as appropriate. See the *if(1)* manual page for full details.

login [*arg ...*]

Replaces the current interactive shell with *login(1)*.

newgrp [*arg ...*]

Replaces the current interactive shell with *newgrp(1)*.

prompt [*on | debug | off*] (+)

Sets the shell's current prompt state to *on*, *debug*, or *off*, or prints its current value. The prompt command helps manage the shell's different prompt behaviors on the user's behalf. When *on*, it enables personal, non-default, shell prompts if the has... When set to *debug*, it enables some extra functionality that can help if you like to see what is going on to turn your \$P string into a personal prompt. When *off*, the shell uses its default prompts. An exit status of 2 indicates error.

pwd (+)

Prints the absolute physical pathname of the current working directory to the standard output. An exit status of 2 indicates that *getcwd(3)* detected an error, or that *pwd* itself detected a usage error.

set [*name [value]*] (+)

Sets the variable *name* to the string *value* in the current shell if both *name* and *value* are specified. Prints the current value of variable *name* if *name* is specified and set in the current shell. Prints the name and value of each variable set in the current shell if no arguments are specified.

setenv *name value* (+)

Sets the environment variable *name* to the string *value*.

shift

Shifts all positional-parameter values to the left by 1, so that the old value of \$2 becomes the new value of \$1 and so forth. The value of \$0 does not shift.

.file [*arg1 ...*] (+)

Is a synonym for the **source** special command.

source file [*arg1 ...*] (+)

Causes the shell to read and execute commands from the specified *file* and return.

If the file name contains no '/' characters, the sequence of directories in the environment variable PATH is searched for the first occurrence of a file by that name. However, if the file name contains one or more '/' characters, the shell attempts to source it without performing any PATH

search. Notice that the file does not need to be executable.

The file and any arguments are treated as positional parameters (see *Parameter substitution* above) during execution of the file. The source command may be nested. As with command files, most shell-detected errors cause the shell to cease execution of the file. If the source command is nested and such an error occurs, all nested source commands terminate.

trap [*''* | *:* | *- signal_number ...*] (+)

'' or *:* causes the specified signals to be ignored if possible, and *-* causes the specified signals to be reset to the default action if possible. If a signal was already ignored when the shell was invoked, it cannot be reset with *-*. If no arguments are specified, a list is printed of those signals which are ignored by trap in the current shell context.

umask [*mask*] (+)

Sets the file creation mask (see *umask(2)*) to the octal value specified by *mask*. If the mask is not specified, its current value is printed.

unalias *name ...* (+)

Removes each of the specified alias *names* from the current shell instance.

unset *name ...* (+)

Removes each of the specified variable *names* from the current shell instance.

unsetenv *name ...* (+)

Removes each of the specified environment variable *names* from the environment of the current shell instance.

verbose [*true* | *false*] (+)

Sets the shell's current verbose state to *true* or *false*, or prints its current value. When true, the shell is verbose, as it is when invoked with the *-v* option. Otherwise, it is not verbose. Its exit status indicates the current value of 0 for true or 1 for false, with 2 indicating error. This may be tested with the **if** command.

wait

Waits for all asynchronous processes to terminate, reporting on abnormal terminations.

Signals (+)

An interactive or login shell always ignores the SIGINT, SIGQUIT, and SIGTERM signals (see *signal(3)*). A login shell also handles the SIGHUP signal, the receipt of which causes the shell to terminate the login session and to read and execute its rc logout files if possible.

If `SIGHUP`, `SIGINT`, `SIGQUIT`, or `SIGTERM` is already ignored when the shell starts, it is also ignored by the current shell and all of its child processes. Otherwise, `SIGINT` and `SIGQUIT` are reset to the default action for sequential child processes, whereas `SIGHUP` and `SIGTERM` are reset to the default action for all child processes.

When a non-interactive shell executes a command file, it does not handle or ignore any signal by default. Any other non-interactive shell ignores `SIGINT` and `SIGQUIT`.

For any signal not mentioned above, the shell inherits the signal action (default or ignore) from its parent process and passes it to its child processes. Remember that the `trap` special command may be used to ignore signals when the shell does not do so by default.

Asynchronous child processes always ignore both `SIGINT` and `SIGQUIT`. Also, if such a process has not redirected its input with a `<`, `|`, or `^`, the shell automatically redirects it to come from `/dev/null`.

EXIT STATUS (+)

The exit status of the shell is generally that of the last command executed prior to end-of-file or **exit**.

However, if the shell is interactive and detects an error, it exits with a non-zero status if the user types an EOT at the next prompt.

Otherwise, if the shell is non-interactive and is reading commands from a file, any shell-detected error causes the shell to cease execution of that file. This results in a non-zero exit status.

A non-zero exit status returned by the shell itself is always one of the values described in the following list, each of which may be accompanied by an appropriate diagnostic:

2 The shell detected a syntax, redirection, or other error not described in this list.

125 An external command was found but did not begin with the proper magic number or a `#!shell` sequence, and a valid shell was not specified by `EXECShell` with which to execute it.

126 An external command was found but could not be executed.

127 An external command was not found.

>128

An external command was terminated by a signal.

ENVIRONMENT (+)

Notice that the concept of ‘user environment’ was not defined in Sixth Edition (V6) UNIX. Thus, use of the following environment variables by this port of the shell is an enhancement:

CWD

The shell sets the value of this variable while setting the $\$w$ special parameter. Both represent the absolute physical path name of the current working directory, without symbolic links. On error, CWD is removed from the environment.

ETSHDIR

If set to a non-empty string, the value of this variable is taken as the path name of a directory which may be used for temporary files. Its value is available to the shell via the $\$d$ special parameter.

EXECSHELL

If set to a non-empty string, the value of this variable is taken as the path name of the shell which is invoked to execute an external command when it does not begin with the proper magic number or a ‘#!shell’ sequence. Its value is available to the shell via the $\$e$ special parameter.

HOME

If set to a non-empty string, the value of this variable is taken as the user’s home directory. Its value is available to the shell via the $\$h$ special parameter and is the default directory for the **chdir** special command.

MANPATH

If set, the value of this variable is taken as the sequence of directories used by *man(1)* to search for manual page files. Its value is available to the shell via the $\$m$ special parameter.

PATH

If set to a non-empty string, the value of this variable is taken as the sequence of directories used by the shell to search both for external commands and for files to be executed by the **source** special command. Its value is available to the shell via the $\$p$ special parameter. Notice that the Sixth Edition (V6) UNIX shell always used the equivalent of ‘.:bin:/usr/bin’, not PATH.

FILES

/dev/null

default source of input for asynchronous processes

/usr/local/etc/etsh.login (+)

system-wide rc init file for login shells

/usr/local/etc/etsh.etchrc (+)

system-wide rc init file for *all* interactive shells

\$h/.etsh.history (+)

user history file for *all* interactive shells

\$h/.etsh.login (+)

user rc init file for login shells

\$h/.etchrc (+)

user rc init file for *all* interactive shells

/usr/local/etc/etsh.logout (+)

system-wide rc logout file for login shells

\$h/.etsh.logout (+)

user rc logout file for login shells

\$h/.etsh.prompt (+)

user etsh prompt preprocessor for interactive etsh instances. It supports the **prompt** special command. In effect, it serves as an interface between the shell and the SetP script.

SEE ALSO

awk(1), env(1), expr(1), fd2(1), goto(1), grep(1), if(1), kill(1), login(1), man(1), newgrp(1), sed(1), stty(1), tsh(1), uname(1)

Etsh home page: <https://etsh.io/>

‘The UNIX Time-Sharing System’ (CACM, July, 1974):

<https://etsh.io/history/unix/>

gives the theory of operation of both the system and the shell.

HISTORY

A **sh** command appeared as */bin/sh* in First Edition UNIX.

The Thompson shell was used as the standard command interpreter through Sixth Edition (V6) UNIX. Then, in the Seventh Edition (V7), it was replaced by the Bourne shell. However, the Thompson shell was still distributed with the system as **osh** because of known portability problems with the Bourne

shell's memory management in Seventh Edition (V7) UNIX.

AUTHORS

This port of the Thompson shell is derived from Sixth Edition (V6) UNIX `/usr/source/s2/sh.c`, which was principally written by Ken Thompson of Bell Labs. Jeffrey Allen Neitzel <jan@etsh.io> initially ported it in January 2004 and currently maintains it as `tsh(1)`. Additionally, I am the principal developer and maintainer of this enhanced port of the shell, which is hereby available as `etsh(1)`.

LICENSE

See either the LICENSE file which is distributed with **etsh** or <https://etsh.io/license/> for full details.

CAVEATS

Unlike the original, this port of the shell can handle 8-bit character sets, as well as the UTF-8 encoding. The original, on the other hand, can only handle 7-bit ASCII.

Notice that certain shell oddities were historically undocumented in this manual page. Particularly noteworthy is the fact that there is no such thing as a usage error. Thus, the following shell invocations are all perfectly valid:

```
etsh -cats_are_nice!!! ': "Good kitty =)"
etsh -tabbies_are_too!
etsh -s
```

The first two cases correspond to the `-c` and `-t` options respectively; the third case corresponds to the `-` option.

BUGS (+)

The shell differentiates between an end-of-file and an error returned by the `read(2)` system call. This allows the shell to print a diagnostic and exit with a non-zero status if appropriate. Previously, it did not recover from `read(2)` errors and exited when this system call failed, and this is still true today.

SECURITY CONSIDERATIONS

This port of the Thompson shell does not support set-ID execution. If the effective user (group) ID of the shell process is not equal to its real user (group) ID when an instance of it starts up, the shell prints the following diagnostic and exits with a non-zero status.

```
Set-ID execution denied
```

If the shell did support set-ID execution, it could possibly allow a user to violate the security policy on a host where the shell is used. For example, if the shell were running a `setuid-root` command file, a

regular user could possibly invoke an interactive root shell as a result.

This is *not* a bug. It is simply how the shell works. Thus, **etsh** does not support set-ID execution. This is a proactive measure to avoid problems, nothing more.