

NAME

`pcat` — expand compressed file to standard output

SYNOPSIS

`pcat name ...`

DESCRIPTION

Pcat does for packed files what *cat*(1) does for ordinary files. The specified files are unpacked and written to the standard output. Thus to view a packed file named *name.z* use:

`pcat name.z`

or just:

`pcat name`

To make an unpacked copy, say *nnn*, of a packed file named *name.z* (without destroying *name.z*) use the command:

`pcat name >nnn`

Pcat returns a value that is the number of files it was unable to unpack. Failure may occur if:

- the file name (exclusive of the *.z*) has more than 12 characters;
- the file cannot be opened;
- the file does not appear to be the output of *pack*(1).

SEE ALSO

`pack`(1), `unpack`(1).

NAME

pcstat - report statistics on output of *getpc* command

SYNOPSIS

pcstat [-nunix_n] [-cco%] [[-k] | [-uPID [usr_n]]] [-p] [-i] [-IPID ...]

DESCRIPTION

Pcstat reads its standard input, which is expected to be the output from the command *getpc*, and reports PC (Program Counter) hit statistics. Program counter distribution reports may be generated for either the operating system or a user process. Additionally, reports showing hits charged to each *PID* may also be given. The type of report asked for from *pcstat* must match the type of data collected by *getpc*. If there is a discrepancy between the two programs' options, an error message is printed and the program terminates. This means, for example, that trying to generate a report on pc hits in the kernel when only user pid data is collected is meaningless and is rejected. *Pcstat* may be used to statistically report on the distribution of cpu cycles on behalf of any one process. All reports contain information giving start time, stop time, % system hits, % user hits, % idle hits, % lost data and % of system time spent in interrupt. The options to *pcstat* have the following meanings:

- n specifies the kernel name list file. By convention, the kernel name list file is assumed to be */unix*, and this option is not needed unless */unix* is not the name list corresponding to the generated data from *getpc*. This option may be specified with any other option.
- c specifies the cut-off percentage at which printing data is terminated. For example, -c 3 would inhibit printing any data with a hit rate of less than 3%. All the data that is 'cut-off' is then lumped into a single category named **OTHER**. The default cut-off value is set to 1%. Fractional cut-offs, e.g., 1.25 may be specified. This option may be specified with any other option.

One and only one of the options -k, -u, or -p must be specified.

- k gives the pc hit distribution within the kernel only. A report is produced showing the distribution of hits in the system across all processes, as well as statistics about text switching if expanded text space is used. The text area for each name is identified by a number indicating its switchable text area or by a U indicating that this name is in a non-switchable text area. Statistics are also printed showing the percent of system hits at each priority level. *Getpc* must have been executed with the -k or -a option to permit this report.
- u *PID* [*usr_n*]
option reports pc distribution within the given *PID* number for a user process with namelist *usr_n*. The *PID* is the user process pid number, and if the name list for the pid is given, a user report is generated along with the system report on this pid. A system report is generated on behalf of the user process and gives the percent of hits for each system routine name charged to this *PID*. *Getpc* must have been executed with the -u, -g, or -a option to permit this report.
- p causes a report on pc hits within different pids to be generated. This option can only give a report when data is gathered using the -a option in *getpc*. Due to the number of *PIDs* that may be present, only the first 100 *PIDs* encountered are used to generate detailed output and all the remaining *PIDs* are lumped together into a category called **OTHER**.
- l gives the user the ability to specify up to 10 *PIDs* that are to be included in the detailed part of the -p report.
- i forces the -u and -p reports to include the system interrupt times.

DIAGNOSTICS

pcstat complains if its input was not generated by *getpc*, or if there is an option specification error.

SEE ALSO

getpc(1)

NAME

plot - graphics filters

SYNOPSIS

plot [-Tterminal [raster]]

DESCRIPTION

These commands read plotting instructions (see *plot(5)*) from the standard input and in general produce, on the standard output, plotting instructions suitable for a particular *terminal*. If no *terminal* is specified, the environment parameter **STERM** (see *environ(7)*) is used. Known *terminals* are:

300 DASI 300.

300S DASI 300s.

450 DASI 450.

4014 Tektronix 4014.

ver Versatec D1200A. This version of *plot* places a scan-converted image in */usr/tmp/raster* and sends the result directly to the plotter device, rather than to the standard output. The optional argument causes a previously scan-converted file *raster* to be sent to the plotter.

FILES

/usr/lib/t300

/usr/lib/t300s

/usr/lib/t450

/usr/lib/t4014

/usr/lib/vplot

/usr/tmp/raster

SEE ALSO

plot(5), *term(7)*.

BUGS

There is no lockout protection for */usr/tmp/raster*.

NAME

postnews - submit news articles

SYNOPSIS

postnews

DESCRIPTION

Postnews is a shell script that calls **inews(1)** to submit news articles to USENET. It will prompt the user for the title of the article (which should be a phrase suggesting the subject, so that persons reading the news can tell if they are interested in the article) and for the newsgroup. An omitted newsgroup (from hitting return) will default to **general**.

general is read by everyone on the local machine. Other possible newsgroups include, but are not limited to, **btl.general**, which is read by all users at all Bell Labs sites on USENET, **net.general**, which is read by all users at all sites on USENET, and **net.news**, which is read by users interested in the network news on all sites. There is often a local set of newsgroups, such as **ucb.all**, that circulate within a local set of machines. (In this case, ucb newsgroups circulate among machines at the University of California at Berkeley.)

After entering the title and newsgroup, the user should type the body of the article. To end the article, type control D at the beginning of a line.

For more sophisticated uses, such as posting from a file, see **inews(1)**.

FILES**SEE ALSO**

mail(1), **inews(1)**, **readnews(1)**, **uucp(1)**, **getdate(3)**, **msgs(1)**, **recnews(1)**, **sendnews(1)**, **uurec(1)**

NAME

SYNOPSIS

POSTNEWS

DESCRIPTION

... of a local network is usually to allow
 ... of USENET. It will search the net for the
 ... which should be a person suggesting the subject, so
 ... leading the news can be of the general
 ... and for the newspaper. An edited newspaper from
 ... will default to general.

general is read by everyone on the local network. Other
 ... not listed for `bfj.general`, `net.general`
 ... all users at the Bell Labs sites on USENET. `net.general`
 ... is read by all users at all sites on USENET, and `net.news`
 ... is read by users interested in the network news
 ... There is often a local set of newspapers
 ... that circulate within a local set of news
 ... newspapers circulate among machines at the University
 ... at Berkeley.

... the file and newspaper. The local site
 ... the article to end the article for
 ... of a line.

... which is being read & edited
 news

FILES
SEE ALSO

... `net.general`, `net.news`, `net.announce`, `net.announce`, `net.announce`
 ... `net.announce`, `net.announce`, `net.announce`

NAME

pr - print file

SYNOPSIS

```
pr [[-nDHtamLPnonlnbncnenwn] [-sc] [-h header]
[+nDHPatolLbceFwh] [ name ... ]]
```

DESCRIPTION

Pr produces a printed listing of one or more files. It's features provide the ability to:

1. insert *n* tabs at the beginning of each line, where *n* is a number between 1 and 9. The default is zero.
2. insert *n* blank lines after each printed line, where *n* is a number between 1 and 9. The default is zero.
3. change the length of the printed page to some length other than the default length of 66 lines.
4. print up to *n* copies of a file, where *n* is a number greater than 1. The default is one; i.e. only one listing is made for each file.
5. suppress the printing of the date, the header string, and/or the string "Page" and page number normally printed at the top of each page.
6. read header information from a file, rather than always being entered at the terminal.
7. change the length of the printed line to something other than the default length of 79 character positions.
8. start printing with page *x* and stop printing after page *y*. The defaults are to start printing with page one and to stop printing after the entire file has been printed.
9. suppress the five-line header and footer margins that are normally provided at the top and bottom of each page.
10. print diagnostic messages on the error output file, file descriptor 2, when files cannot be found or when values specified for certain options, such as *-o*, *-b* and *-c* are unrealistic.
11. permit lines that are longer than the specified line width to be continued to the next line. When determining the length of a line, tabs are counted as the number of 040's that are required to effect the tab and backspaces, 010's, and escape characters, 033's, are subtracted from the current length of the line. All other ASCII characters whose octal value is less than 040 (space) are not included in the character count for a line.
12. print the page number and the file name or optional header within the five line margin provided at the bottom of the page.
13. print the line number associated with each line of an input file at the beginning of the output line.

If *pr* is called without specifying any of the minus or plus options, one copy of each specified file is printed. Each printed page has:

1. a five-line header consisting of two blank lines; a printed line containing the date the file was last modified, the file name, and the page number; and two blank lines.
2. up to 56 printed lines of not more than 79 characters each. Each line will not be offset by any tabs and will be single-spaced unless the file contains blank lines.
3. five blank lines as a footer at the bottom of the page.

The minus and plus options permit the user to modify the standard, or default, output described in the preceding paragraph. When such options are specified with the *pr* command,

they apply to all following files unless reset between files. `-` options may be grouped, such as `-o1c2b2`. `+` options may also be grouped, such as `+HPDc`. A description of each option follows:

- `-n` produces n -column output. This option must appear at the beginning of a group of minus options or by itself.
- `-D` suppresses the printing of the date at the top of each page.
- `-H` suppresses the printing of the 'header string', or the file name if the `-h` option is not specified, at the top of each page.
- `-t` suppresses the five-line margins at the top and bottom of each page.
- `-a` The page number and the file name or optional header is printed within the five line margin provided at the bottom of the page. The `-t` option will cause this option to be ignored. The `-` option will cause only the file name or header to be printed within the margin at the bottom of the page.
- `-m` prints all files simultaneously, each in one column.
- `-L` prints the line number, associated with a line of an input file, at the beginning of each line. This feature is useful when debugging errors reported by the `cc` command.
- `-Pn` suppresses the printing of the string "Page" and page number at the top of each page if n is not specified or if n has the value 0. If, on the other hand, n has the value 1, the string "Page" is suppressed and the page number is printed in the center of the header line. If n has the value 1, then the options `-D` and `-H` must be specified; if they are not specified, the page number will not be centered.
- `-on` inserts n tabs before each printed line. (max. of 9)
- `-ln` uses n as the length of the page, rather than the default length of 66 lines. A length of 33, for example, would cause the physical page to be broken up into two actual pages.
- `-bn` inserts n blank lines after each printed line. (max. of 9)
- `-cn` prints n copies of each file.
- `-en` stops printing after page n .
- `-wn` uses n as the width of a line instead of the default width of 79 character positions. Lines which are longer than n will be continued on the next line. Such continued lines will be preceded by the character sequence "*" (this applies to single-column output only).
- `-sc` separates columns by the single character c instead of by the appropriate amount of white space. A missing c is taken to be a tab. This option must appear at the end of a group of minus options or by itself.
- `-h` treats the next argument as a header. If the 'header' argument begins with a '+', it is treated as the name of a file containing the header (50 characters or less) to be printed. This option must appear at the end of a group of minus options or by itself.
- `+n` begin printing with page n . This option must appear at the beginning of a group of plus options or by itself.
- `+D` restores the printing of the date at the top of each page.
- `+H` restores the printing of the 'header string' or the file name at the top of each page.
- `+P` restores the printing of the page and number at the top of each page.
- `+a` suppresses the printing of the page number and the file name or optional header within the five line margin provided at the bottom of each page.

- +t restores the top and bottom margins for each printed page to five lines each.
- +o suppresses the printing of tabs at the beginning of each printed line.
- +l restores page length to 66 lines.
- +L suppresses the printing of the line number at the beginning of each line.
- +b suppresses the printing of blank lines after each printed line.
- +c suppresses the printing of multiple copies of a file.
- +e makes the last page to be printed coincide with the end of the file being printed.
- +F makes the first page to be printed coincide with the beginning of the file being printed.
- +w restores the width of the line to 79 character positions.
- +h restores the printing of the file name as the header at the top of each page.

FILES

/dev/lm?? to suspend messages.

SEE ALSO

cat(1), cp(1), spr(1), mesg(1), stty(1), stty(2)

DIAGNOSTICS

The following output messages are routed to the user's error output device, file descriptor 2, for the following error conditions:

ERROR CONDITION*----- MESSAGE -----*

file not found "FILE <name> NOT FOUND."

tabs> 9 "TOO MANY TABS REQUESTED; TABS REDUCED TO 9."

blank lines> 9 "TOO MANY BLANK LINES REQUESTED; BLANK LINES REDUCED TO 9."

columns> lsize "TOO MANY COLUMNS REQUESTED; PROGRAM TERMINATED."

files> 9 "TOO MANY FILES FOR SIMULTANEOUS PRINTING."

temp. file "CANNOT CREATE TEMP FILE NAME; COPIES REDUCED TO 1."

"CANNOT CREATE TEMP FILE."

"CANNOT OPEN TEMP FILE."

"CANNOT WRITE TEMP FILE."

BUGS

Doesn't work correctly if the input file contains carriage control sequences like 'backspace' or 'back-one-line'.

NAME

prof - display profile data

SYNOPSIS

prof [-v] [-a] [-l] [name]

DESCRIPTION

Prof interprets the file *mon.out* produced by the *monitor* subroutine. Under default modes, the symbol table in the *name* object file (*a.out* default) is read and correlated with the *mon.out* profile file. For each external symbol, the percentage or time spent executing between that symbol and the next is printed (in decreasing order), together with the number of times that routine was called and the number of milliseconds per call.

If the *-a* option is used, all symbols are reported rather than just external symbols. If the *-l* option is used, the output is listed by symbol value rather than decreasing percentage. If the *-v* option is used, all printing is suppressed and a profile plot is produced on the 611 display.

In order for the number of calls to a routine to be tallied, the *-p* option of *cc* must have been given when the file containing the routine was compiled.

FILES

mon.out	for profile
a.out	for namelist
/dev/vt0	for plotting

SEE ALSO

monitor(3), profil(2), cc(1)

BUGS

The numeric results tend to vary considerably from run to run because of statistical variations in the actual clock interrupt times as correlated with the program run times. In addition, because a guess must be made as to which routine should be credited for a clock tick (16.67ms), invariably it happens that a routine next to a very active routine will get credited with more time than it deserves. It would be good to produce a histogram or a graph of time versus program address rather than simple numeric results.