NAME

pigz, unpigz – compress or expand files

SYNOPSIS

pigz [**–cdfhikKlLnNqrRtTz0..9**] [**-b** *blocksize*] [**-p** *threads*] [**-S** *suffix*] [*name ...*]
unpigz [**–cfhikKlLnNqrRtTz**] [**-b** *blocksize*] [**-p** *threads*] [**-S** *suffix*] [*name ...*]

DESCRIPTION

Pigz compresses using threads to make use of multiple processors and cores. The input is broken up into 128 KB chunks with each compressed in parallel. The individual check value for each chunk is also calculated in parallel. The compressed data is written in order to the output, and a combined check value is calculated from the individual check values.

The compressed data format generated is in the gzip, zlib, or single-entry zip format using the deflate compression method. The compression produces partial raw deflate streams which are concatenated by a single write thread and wrapped with the appropriate header and trailer, where the trailer contains the combined check value.

Each partial raw deflate stream is terminated by an empty stored block (using the Z_SYNC_FLUSH option of zlib), in order to end that partial bit stream at a byte boundary. That allows the partial streams to be concatenated simply as sequences of bytes. This adds a very small four to five byte overhead to the output for each input chunk.

The default input block size is 128K, but can be changed with the **-b** option. The number of compress threads is set by default to 4, which can be changed using the **-p** option. Specifying **-p 1** avoids the use of threads entirely.

The input blocks, while compressed independently, have the last 32K of the previous block loaded as a preset dictionary to preserve the compression effectiveness of deflating in a single thread. This can be turned off using the **-i** or **--independent** option, so that the blocks can be decompressed independently for partial error recovery or for random access.

Decompression can't be parallelized, at least not without specially prepared deflate streams for that purpose. As a result, *pigz* uses a single thread (the main thread) for decompression, but will create three other threads for reading, writing, and check calculation, which can speed up decompression under some circumstances. Parallel decompression can be turned off by specifying one process (**-dp 1** or **-tp 1**).

Compressed files can be restored to their original form using *pigz -d* or *unpigz.*

OPTIONS

-# --fast --best
Regulate the speed of compression using the specified digit *#*, where **–1** or **––fast** indicates the fastest compression method (less compression) and **–9** or **––best** indicates the slowest compression method (best compression). Level 0 is no compression.

-b --blocksize mmm
Set compression block size to mmmK (default 128KiB).

-c --stdout --to-stdout
Write all processed output to stdout (won't delete).

-d --decompress --uncompress
Decompress the compressed input.

-f --force
Force overwrite, compress .gz, links, and to terminal.

-h --help
Display a help screen and quit.

-i --independent
Compress blocks independently for damage recovery.

-k --keep
Do not delete original file after processing.

-K --zip
Compress to PKWare zip (.zip) single entry format.

-l --list List the contents of the compressed input.

-L --license
Display the *pigz* license and quit.

-n --no-name
Do not store or restore file name in/from header.

-N --name
Store/restore file name and mod time in/from header.

-p --processes n
Allow up to n processes (default is the number of online processors)

-q --quiet --silent
Print no messages, even on error.

-r --recursive
Process the contents of all subdirectories.

-S --suffix .sss
Use suffix .sss instead of .gz (for compression).

-t --test
Test the integrity of the compressed input.

-T --no-time
Do not store or restore mod time in/from header.

-v --verbose
Provide more verbose output.

-V --version
Show the version of pigz.

-z --zlib
Compress to zlib (.zz) instead of gzip format.

COPYRIGHT NOTICE