Pnmtotiff User Manual

pnmtotiff

Updated: 27 March 2005
Table Of Contents

NAME

pnmtotiff - convert a PNM image to a TIFF file

SYNOPSIS

pnmtotiff [**-none** | **-packbits** | **-lzw** | **-g3** | **-g4** | **-flate** | **-adobeflate**] [**-2d**] [**-fill**] [**-predictor=***n*] [**-msb2lsb**|**-lsb2msb**] [**-rowsperstrip=***n*] [**-minisblack**|**-miniswhite**|**mb**|**mw**] [**-truecolor**] [**-color**] [**-indexbits=***bitwidthlist*] [**-xresolution=***xres*] [**-yresolution=***yres*] [**-resolutionunit=**{**inch** | **centimeter** | **none** | **in** | **cm** | **no**}] [**-indexbits=**[**1**[**2**[**4**[**8**]]]]] [**-append**] [*pnmfile*]

You can use the minimum unique abbreviation of the options. You can use two hyphens instead of one. You can separate an option name from its value with white space instead of an equals sign.

DESCRIPTION

This program is part of Netpbm.

pnmtotiff reads a PNM image as input and produces a TIFF file as output.

Actually, it handles multi-image PNM streams, producing multi-image TIFF streams (i.e. a TIFF stream with multiple "directories"). But before Netpbm 10.27 (March 2005), it ignored all but the first PNM image in the input stream.

The Output File

The output goes to Standard Output. **pnmtotiff** approaches this output file differently from Unix and Netpbm convention. This is entirely due to **pnmtotiff**'s use of the TIFF library to do all TIFF output.

> The output file must be seekable. **pnmtotiff** does not write it sequentially. Therefore, you can't use a pipe; you can't pipe the output of **pnmtotiff** to some other program. But any regular file should work.
>
> If the output file descriptor is readable, you must either specify **-append** so as to add to the existing file, or make sure the file is empty. Otherwise, **pnmtotiff** will fail with an unhelpful message telling you that a TIFF library function failed to open the TIFF output stream.
>
> If you are converting multiple images (your input stream contains multiple images), the output file must be both readable and writable.

If you're using a Unix command shell to run **pnmtotiff**, you use facilities of your shell to set up Standard Output. In Bash, for example, you would set up a write-only Standard Output to the file /tmp/myimage.tiff like this:

```
$ pnmtotiff myimage.pnm >/tmp/myimage.tiff
```

In Bash, you would set up a read/write Standard Output to the file /tmp/myimage.tiff like this:

```
$ pnmtotiff myimage.pnm 1<>/tmp/myimage.tiff
```

OPTIONS

Compression

By default, **pnmtotiff** creates a TIFF file with no compression. This is your best bet most of the time. If you want to try another compression scheme or tweak some of the other even more obscure output options, there are a number of options which which to play.

Before Netpbm 8.4 (April 2000), the default was to use LZW compression. But then new releases of the TIFF library started omitting the LZW compression capability due to concern about patents on LZW. So since then, the default has been no compression. The LZW patents have now expired and new TIFF libraries do LZW, but the **pnmtotiff** behavior remains the same for compatibility with older TIFF libraries and applications of **pnmtotiff**.

The **-none**, **-packbits**, **-lzw**, **-g3**, **-g4**, **-flate**, and **-adobeflate** options are used to override the default and set the compression scheme used in creating the output file. The CCITT Group 3 and Group 4 compression algorithms can be used only with bilevel data. The **-2d** and **-fill** options are meaningful only with Group 3 compression: **-2d** requests 2-dimensional encoding, while **-fill** requests that each encoded scanline be zero-filled to a byte boundary. The **-predictor** option is meaningful only with LZW compression: a predictor value of 2 causes each scanline of the output image to undergo horizontal differencing before it is encoded; a value of 1 forces each scanline to be encoded without differencing. By default, **pnmtotiff** creates a TIFF file with msb-to-lsb fill order. The **-msb2lsb** and **-lsb2msb** options are used to override the default and set the fill order used in creating the file.

With some older TIFF libraries, **-lzw** doesn't work because the TIFF library doesn't do LZW compression. This is because of concerns about Unisys's patent on LZW which was then in force. Actually, with very old TIFF libraries, **-lzw** works because no distributors of the TIFF library were sensitive yet to the patent issue.

-flate chooses "flate" compression, which is the patent-free compression common in the Unix world implemented by the "Z" library. It is what the PNG format uses.

Fill Order

The **-msb2lsb** and **lsb2msb** options controll the fill order.

The fill order is the order in which pixels are packed into a byte in the Tiff raster, in the case that there are multiple pixels per byte. msb-to-lsb means that the leftmost columns go into the most significant bits of the byte in the Tiff image. However, there is considerable confusion about the meaning of fill order. Some believe it means whether 16 bit sample values in the Tiff image are little-endian or big-endian. This is totally erroneous (The endianness of integers in a Tiff image is designated by the image's magic number). However, ImageMagick and older Netpbm both have been known to implement that interpretation. 2001.09.06.

If the image does not have sub-byte pixels, these options have no effect other than to set the value of the FILLORDER tag in the Tiff image (which may be useful for those programs that misinterpret the tag with reference to 16 bit samples).

Color Space

-color tells **pnmtotiff** to produce a color, as opposed to grayscale, TIFF image if the input is PPM, even if it contains only shades of gray. Without this option, **pnmtotiff** produces a grayscale TIFF image if the input is PPM and contains only shades of gray, and at most 256 shades. Otherwise, it produces a color TIFF output. For PBM and PGM input, **pnmtotiff** always produces grayscale TIFF output and this option has no effect.

The **-color** option can prevent **pnmtotiff** from making two passes through the input file, thus improving speed and memory usage. See Multiple Passes.

-truecolor tells **pnmtotiff** to produce the 24-bit RGB form of TIFF output if it is producing a color TIFF image. Without this option, **pnmtotiff** produces a colormapped (paletted) TIFF image unless there are more than 256 colors (and in the latter case, issues a warning).

The **-truecolor** option can prevent **pnmtotiff** from making two passes through the input file, thus improving speed and memory usage. See Multiple Passes.

The **-color** and **-truecolor** options did not exist before Netpbm 9.21 (December 2001).

If **pnmtotiff** produces a grayscale TIFF image, this option has no effect.

The **-minisblack** and **-miniswhite** options force the output image to have a "minimum is black" or "minimum is white" photometric, respectively. If you don't specify either, **pnmtotiff** uses minimum is black except when using Group 3 or Group 4 compression, in which case **pnmtotiff** follows CCITT fax standards and uses "minimum is white." This usually results in better compression and is generally preferred for bilevel coding.

Before February 2001, **pnmtotiff** always produced "minimum is black," due to a bug. In either case, **pnmtotiff** sets the photometric interpretation tag in the TIFF output according to which photometric is actually used.

The **-indexbits** option is meaningful only for a colormapped (paletted) image. In this kind of image, the raster contains values which are indexes into a table of colors, with the indexes normally taking less space that the color description in the table. **pnmtotiff** can generate indexes of 1, 2, 4, or 8 bits. By default, it will use 8, because many programs that interpret TIFF images can't handle any other width.

But if you have a small number of colors, you can make your image considerably smaller by allowing fewer than 8 bits per index, using the **-indexbits** option. The value is a comma-separated list of the bit widths you allow. **pnmtotiff** chooses the smallest width you allow that allows it to index the entire color table. If you don't allow any such width, **pnmtotiff** fails. Normally, the only useful value for this option is **1,2,4,8**, because a program either understands the 8 bit width (default) or understands them all.

In a Baseline TIFF image, according to the 1992 TIFF 6.0 specification, 4 and 8 are the only valid widths. There are no formal standards that allow any other values.

This option was added in June 2002. Before that, only 8 bit indices were possible.

Resolution

A Tiff image may contain information about the resolution of the image, which means how big in real dimensions (centimeters, etc.) each pixel in the raster is. You control that with the **-xresolution**, **-yresolution**, and **-resolutionunit** options.

These options do not control how many pixels **pnmtotiff** generates or how much information is in the pixels. They control only the value of tags that may or may not be used by whatever reads the image.

The value of the **-xresolution** option is a floating point decimal number that tells how many pixels there are per unit of distance in the horizontal direction. **-yresolution** is analogous for the vertical direction.

The unit of distance is given by the value of the **-resolutionunit** option. That value is either **inch**, **centimeter** or **none** (or abbreviations **in**, **cm**, or **no**). **none** means the unit is arbitrary or unspecified. This could mean that the creator and user of the image have a separate agreement as to what the unit is. But usually, it just means that the horizontal and vertical resolution values cannot be used for anything except to determine aspect ratio (because even though the unit is arbitrary or unspecified, it has to be the same for both resolution numbers).

If you don't specify **-xresolution**, the Tiff image does not contain horizontal resolution information. Likewise for **-yresolution**. If you don't specify **-resolutionunit**, the default is inches.

Before Netpbm 10.16 (June 2003), **-resolutionunit** did not exist and the resolution unit was always inches.

Other

You can use the **-rowsperstrip** option to set the number of rows (scanlines) in each strip of data in the output file. By default, the output file has the number of rows per strip set to a value that will ensure each strip is no more than 8 kilobytes long.

The **-append** option tells **pnmtotiff** to add images to the existing output file (a TIFF file may contain multiple images) instead of the default, which is to replace the output file.

-append was new in Netpbm 10.27 (March 2005).

NOTES

There are myriad variations of the TIFF format, and this program generates only a few of them. **pnmtotiff** creates a grayscale TIFF file if its input is a PBM (monochrome) or PGM (grayscale) file. **pnmtotiff** also creates a grayscale file if it input is PPM (color), but there is only one color in the image. If the input is a PPM (color) file and there are 256 colors or fewer, but more than 1, **pnmtotiff** generates a color palette TIFF file. If there are more colors than that, **pnmtotiff** generates an RGB (not RGBA) single plane TIFF file. Use **pnmtotiffcmyk** to generate the cyan-magenta-yellow-black ink color separation TIFF format.

The number of bits per sample in the TIFF output is determined by the maxval of the PNM input. If the maxval is less than 256, the bits per sample in the output is the smallest number that can encode the maxval. If the maxval is greater than or equal to 256, there are 16 bits per sample in the output.

Multiple Passes

pnmtotiff reads the input image once if it can, and otherwise twice. It needs that second pass (which happens before the main pass, of course) to analyze the colors in the image and generate a color map (palette) and determine if the image is grayscale. So the second pass happens only when the input is PPM. And you can avoid it then by specifying both the **-truecolor** and **-color** options.

If the input image is small enough to fit in your system's file cache, the second pass is very fast. If not, it requires reading from disk twice, which can be slow.

When the input is from a file that cannot be rewound and reread, **pnmtotiff** reads the entire input image into a temporary file which can, and works from that. Even if it needs only one pass.

Before Netpbm 9.21 (December 2001), **pnmtotiff** always read the entire image into virtual memory and then did one, two, or three passes through the memory copy. The **-truecolor** and **-color** options did not exist. The passes through memory would involve page faults if the entire image did not fit into real memory. The image in memory required considerably more memory (12 bytes per pixel) than the cached file version of the image would.

SEE ALSO

tifftopnm, **pnmtotiffcmyk**, **pnmdepth**, **pnm**

AUTHOR

Derived by Jef Poskanzer from ras2tiff.c, which is Copyright (c) 1990 by Sun Microsystems, Inc. Author: Patrick J. Naughton (naughton@wind.sun.com).

Table Of Contents

- NAME
- SYNOPSIS
- DESCRIPTION
- OPTIONS
- NOTES
- Multiple Passes
- SEE ALSO
- AUTHOR