

# Linux Fu: Scripting for Binary Files

[Originalartikel](#)

[Backup](#)

<html> <p>If you ever need to write a binary file from a traditional language like C, it isn't all that hard to do. About the worst thing you might have to deal with is attempts to fake line endings across Windows and Linux, but there's usually a way to turn that off if it is on by default. However, if you are using some type of scripting language, binary file support might be a bit more difficult. One answer is to use a tool like xxd or<a href="https://thosakwe.github.io/t2b/index.html" target=",\_blank">t2b</a> (text-to-binary) to handle the details. You can find the code for t2b on GitHub including prebuilt binaries for many platforms. You should be able to install xxd from your system repository.</p> <p>These tools take very different approaches. You might be familiar with tools like od or hexdump for producing readable representations of binary files. The xxd tool can actually do the same thing ; although it is not as flexible. What xxd can even reverse itself so that it can rebuild a binary file from a hex dump it creates (something other tools can't do). The t2b tool takes a much different approach. You issue commands to it that causes it to write an original hex file.</p> <p>Both of these approaches have some merit. If you are editing a binary file in a scripting language, xxd makes perfect sense. You can convert the file to text, process it, and then roll it back to binary using one program. On the other hand, if you are creating a binary file from scratch, the t2b program has some advantages, too.</p> <p>I decided to write a few test scripts using bash to show how it all works. These aren't production scripts so they won't be as hardened as they could be, but there is no reason they couldn't be made as robust as you were willing to make them.</p> <h2>Cheating a Little</h2> <p>I decided to write two shell scripts. One will generate an image file. I cheated in two ways there. First, I picked the PPM (Portable Pix Map) format which is very simple to create. And second I ignored the format that uses ASCII instead of binary. That's not strictly cheating because it does make a larger file, as you'd expect. So there is a benefit to using the binary format.</p> <p>The other script takes a file in the same format and cuts the color values within it by half. This shows off both tools since the first job is generating an image file from data and the second one is processing an image file and writing out a new one. I'll use t2b for the first job and xxd for the second.</p> <h2>PPM File Format</h2> <p><a href="https://hackaday.com.files.wordpress.com/2018/05/tux2.png" target=",\_blank">

<https://hackaday.com.files.wordpress.com/2018/05/tux2.png?w=308&h=400> 308w“ sizes=„(max-width: 192px) 100vw, 192px“/></a>The PPM format is part of a family of graphics formats from the 1980s. They are very simple to construct and deconstruct, although they aren't known for being small. However, if you needed to create a graphic from a Raspberry Pi program, it is sometimes handy to create them using this simple file format and then use ImageMagick or some other tool to convert to a nicer format like PNG.</p> <p>There are actually three variants of the format. One for black and white, one for grayscale, and another for color. In addition, each of them can contain ASCII data or binary data. There is a very simple header which is always in ASCII.</p> <p>We'll only worry about the color format. The header will start with the string &#8220;P6.&#8221; That usually ends with a newline, although defensively, you ought to allow for any whitespace character to end the header fields. Then the X and Y limits &#8212; in decimal and still in ASCII &#8212; appear separated by whitespace. This is usually really a space and a newline at the end. The next part of the header is another ASCII decimal value indicating the maximum value for the color components in the image. After that, the data is binary RGB (red/green/blue) triplets. By the way, if the P6 had been a P3, everything would remain the same, but the RGB triplets would be in ASCII, not binary. This could be handy in some cases but &#8212; as I mentioned &#8212; will result in a larger file.</p> <p>Here's a sample header with a little bit of binary data following it:</p> <p><a href="https://hackaday.com.files.wordpress.com/2018/06/header.png" target="blank"></a></p> <p>The green text represents hex numbers and the other boxes contain ASCII characters. You can see the first 15 bytes are header and after that, it is all image data.</p> <h2>T2B</h2> <p>The t2b program takes a variety of commands to generate output. You can write a string or various sizes of integers. You can also do things like repeat output a given number of times and even choose what to output based on conditions. There's a way to handle variables and even macros.</p> <p>As an example, my script will write out an image with three color bars in it. The background will be black with a white border. The color bars will automatically space to fit the box.&#160;I won't use too many of the t2b features, but I did like using the macros to make the resulting output easier to read.&#160;Here's the code for creating the header (with comments added):</p> <pre class="brush: bash; title: ; notranslate" title="">> strl P6 # Write P6 followed by a newline (no quotes needed because no whitespace in the string) str \$X # Write the X coordinate (no newline) u8 32 # a space strl \$Y # The Y coordinate (with newline) strl 255 #

Maximum subpixel value (ASCII) </pre> <p>That's all there is to it. The RGB triples use the u8 command, although you could probably use a 24-bit command, too. I also set up some macros for the colors I used:</p> <pre class=„brush: bash; title: ; notranslate“ title=„“> macro RED

```
begin
  u8 255
  times 2 u8 0
  endtimes
```

endmacro </pre> <p>Once you have the t2b language down, the rest is just math. You can find the complete code on <a href=„<https://github.com/wd5gnr/binaryscript>“ target=„\_blank“>GitHub</a>, but you'll see it just computes 7 equal-sized regions and draws different colors as it runs through each pixel in a nested set of for loops. There's also a one-pixel white border around the edges for no good reason.</p> <p>When you want to run the code you can either specify the X and Y coordinates or take the 800&#215;600 default:</p> <pre> ./colorbar.sh 700 700 | t2b
&gt;outputfile.ppm </pre> <p>If you intercept the output before the t2b program, you'll see the commands rolling out of the script. Here's the default output to the ppm file:</p> <p><a href=„<https://hackaday.com.files.wordpress.com/2018/06/test1.png>“ target=„\_blank“><img data-attachment-id=„312253“
data-permalink=„<https://hackaday.com/2018/06/29/linux-fu-scripting-for-binary-files/test1/>“ data-original-file=„<https://hackaday.com.files.wordpress.com/2018/06/test1.png>“ data-original-size=„800,600“ data-comments-opened=„1“ data-image-
meta=„{“aperture“:„0“,”credit“:„“,”camera“:„
:“,”caption“:„“,”created\_timestamp“:„0“,”
:“,”copyright“:„“,”focal\_length“:„0“,”
:“,”iso“:„0“,”
:“,”shutter\_speed“:„0“,”
:“,”title“:„“,”
:“,”orientation“:„0“,”
:“}“ data-image-title=„test1“ data-image-description=„
data-medium-
file=„<https://hackaday.com.files.wordpress.com/2018/06/test1.png?w=400&h=300>“ data-large-
file=„<https://hackaday.com.files.wordpress.com/2018/06/test1.png?w=800>“ class=„aligncenter wp-
image-312253 size-medium“
src=„<https://hackaday.com.files.wordpress.com/2018/06/test1.png?w=400&h=300>“ alt=„
width=„400“ height=„300“
srcset=„<https://hackaday.com.files.wordpress.com/2018/06/test1.png?w=400&h=300> 400w,
<https://hackaday.com.files.wordpress.com/2018/06/test1.png?w=250&h=188> 250w,
<https://hackaday.com.files.wordpress.com/2018/06/test1.png?w=768&h=576> 768w,
<https://hackaday.com.files.wordpress.com/2018/06/test1.png> 800w“ sizes=„(max-width: 400px)
100vw, 400px“></a></p> <h2>Shades of Gray</h2> <p>The other script is a little different. The goal is to divide all the color values in a PPM file in half. If it were just binary data, that would be easy enough, but you need to skip the header so as not to corrupt it. That takes a little extra work. I used gawk (GNU awk) to make the work a little simpler.</p> <p>The code expects output from xxd, which looks like this:</p> <pre class=„brush: bash; title: ; notranslate“ title=„“> 00000000: 5036 0a38
3030 2036 3030 0a32 3535 0aff P6.800 600.255.. 00000010: ffff ffff ffff ffff ffff ffff ffff ffff ..... 00000020: ffff ffff ffff ffff ffff ffff ffff ffff ..... 00000030: ffff ffff ffff ffff ffff ffff ffff ffff ..... 00000040: ffff ffff ffff ffff ffff ffff ffff ffff ..... </pre> <p>The address isn't important to us. You can ask xxd to suppress it, but it is also easy to just skip it. The character representations to the right aren't important either. The xxd program will ignore that when it rebuilds the binary. Here's the code in awk (which is embedded in <a href=„<https://github.com/wd5gnr/binaryscript/blob/master/half.sh>“ target=„\_blank“>the shell
script</a>):</p> <pre class=„brush: bash; title: ; notranslate“ title=„“> # need to find 4 white
space fields BEGIN { noheader=4 }

```

{
lp=1
}
{
split($0, chars, "")

# skip initial address

while (chars[lp++]!=":");
n=0; # # of bytes read

# get two characters

while (n<16 && lp<length(chars)) { # heuristically two space
characters out of xxd ends the hex dump line (ascii follows) if (chars[lp] ~
/[ \t\n\r]/) { if (chars[++lp] ~ /[ \t\n\r]/) { break; # no need to look at
rest of line } } b=chars[lp++]; chars[lp++]; n++; # if header then skip white
space if (noheader>0) {
    if (b=="20" || b=="0a" || b=="0d" || b=="09") noheader--;
}
else {
# if not header then /2
    bn=strtonum("0x" b)/2;
    bs=sprintf("%02x",bn);
    chars[lp-2]=substr(bs,1,1);
    chars[lp-1]=substr(bs,2,1);
}
}

# recombine array and print

p=""
for (i=1;i<=length(chars);i++) p=p chars[i];
print p
}

```

</pre> <p>The awk code simply skips the address and then pulls up to 16 items from a line of data. The first task is to count whitespace characters to skip over the header. I made the assumption that there would not be runs of whitespace, although a more robust program would probably consume multiple spaces (easy to fix). After that, each byte gets divided and reassembled. This task is more character oriented and awk doesn't handle characters well without a trick.</p> <p><a href= „<https://hackaday.com/files.wordpress.com/2018/06/fifty.png>“ target= „\_blank“><img data-attachment-id= „312269“ data-permalink= „<https://hackaday.com/2018/06/29/linux-fu-scripting-for-binary-files/fifty-2/>“ data-orig-file= „<https://hackaday.com/files.wordpress.com/2018/06/fifty.png>“ data-orig-size= „800,600“ data-comments-opened= „1“ data-image- meta= „{“aperture“:„0“,”credit“:„“,”camera“:„“,”caption“:„“,”created\_timestamp“:„0“,”copyright“:„“,”focal\_length“:„0“,”iso“:„0“}“</p>

&quot;,&quot;shutter\_speed&quot;:&quot;0&quot;,&quot;title&quot;:&quot;,&quot;orientation&quot;:&quot;0&quot;};“ data-image-title=„fifty“ data-image-description=„ data-medium-file=„<https://hackadaycom.files.wordpress.com/2018/06/fifty.png?w=400>“ data-large-file=„<https://hackadaycom.files.wordpress.com/2018/06/fifty.png?w=800>“ class=„alignleft size-thumbnail wp-image-312269“ src=„<https://hackadaycom.files.wordpress.com/2018/06/fifty.png?w=250&h=188>“ alt=„ width=„250“ height=„188“ srcset=„<https://hackadaycom.files.wordpress.com/2018/06/fifty.png?w=250&h=188> 250w, <https://hackadaycom.files.wordpress.com/2018/06/fifty.png?w=500&h=376> 500w, <https://hackadaycom.files.wordpress.com/2018/06/fifty.png?w=400&h=300> 400w“ sizes=„,(max-width: 250px) 100vw, 250px“/></a>In particular, I used the split command to convert the current line into an array with each element containing a character. This includes any whitespace characters because I used an empty string as the split delimiter:</p> <pre> split(\$0, chars, „) </pre> <p>After processing the array &#8212; which isn&#8217;t hard to do &#8212; you can build a new string back like this:</p> <pre> p=„“ for (i=1;i<=length(chars);i++) p=p chars[i]; </pre> <p>The output file will feed back to xxd with the -r option and you are done:</p> <pre> xxd infile.ppm | ./half.sh | xxd -r &gt;outfile.ppm </pre> <h2>Two is the Loneliest</h2> <p>This is a great example of how the Unix philosophy makes it possible to build tools that are greater than the sum of their parts. A simple program changes a text-processing language like awk into a binary file manipulation language. Great. By the way, if your idea of manipulating binary is Intel hex or Motorola S records, be sure to check out the srec\_cat and related software which can manipulate those, too.</p> <p>Once you have a bunch of binary files, you might appreciate an <a href=„<https://hackaday.com/2017/07/27/edit-hex-in-the-browser/>“>online hex editor</a>. By the way, a couple of years ago, I mentioned using <a href=„<https://hackaday.com/2015/04/02/manual-data-recovery-with-a-hex-editor/>“>od to process binary files in awk</a>. That&#8217;s still legitimate, of course, but xxd allows you to go both ways, which is a lot more useful.</p> </html>

From:  
<https://schnipsl.qgelm.de/> - **Qgelm**



Permanent link:  
[https://schnipsl.qgelm.de/doku.php?id=wallabag:linux-fu\\_-scripting-for-binary-files](https://schnipsl.qgelm.de/doku.php?id=wallabag:linux-fu_-scripting-for-binary-files)

Last update: **2021/12/06 15:24**