

Hentzenwerke Whitepaper Series

# ***Capturing Screenshots on Linux***

**By Whil Hentzen**

**Making a copy of a screenshot on a Linux box and then placing it in a document or giving it to a printer for output is an operation that many people need to perform, but very few people know how to do properly. In this whitepaper, I'll discuss how to capture a screen, how to handle size and resolution issues, and how to prepare it for output. I'll be using The GIMP, an open source source tool that comes with most popular distributions of Linux.**

## 1. Preface

### 1.1 Copyright

Copyright 2004 Whil Hentzen. Some rights reserved. This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs License, which basically means that you can copy, distribute, and display only unaltered copies of this work, but in return, you must give the original author credit, you may not distribute the work for commercial gain, nor create derivative works based on it without first licensing those rights from the author. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/2.0/>.

### 1.2 Revisions

#### 1.2.1 History

Version	Date	Synopsis	Author
1.0.0	2004/1/31	Original	WH

#### 1.2.2 New version

The newest version of this document will be found at [www.hentzenwerke.com](http://www.hentzenwerke.com).

#### 1.2.3 Feedback and corrections

If you have questions, comments, or corrections about this document, please feel free to email me at 'books@hentzenwerke.com'. I also welcome suggestions for passages you find unclear.

### 1.3 References and acknowledgments

Thanks to MLUG members Phil Goembel, Rick Miller, George Moulton, and Doug Turner.

### 1.4 Disclaimer

No warranty! This material is provided as is, with no warranty of fitness for any particular purpose. Use the concepts, examples and other content at your own risk. There may be errors and inaccuracies that in some configurations may be damaging to your system. The author(s) disavows all liability for the contents of this document.

Before making any changes to your system, ensure that you have backups and other resources to restore the system to its state before making those changes.

All copyrights are held by their respective owners, unless specifically noted otherwise. Use of a term in this document should not be regarded as affecting the validity of any trademark or service mark. Naming of particular products or brands should not be seen as endorsements.

### 1.5 Prerequisites

This document was written using Fedora Core 1.0 and assumes a beginner's familiarity with use of Linux via the GUI and the Command Window.

## 2. Basic Concepts

Computer monitors display a grid of dots, or pixels. (I'll use the terms interchangeably, although some purists would probably take offense.) These dots are arranged in a rectangular pattern, in fixed proportions. Starting in the early 1980's, monitors used a grid of 640 dots across by 480 dots up and down. This proportion - 640 x 480 - is known as the resolution of the monitor. As technology advanced, the number of dots grew, resulting in higher resolutions - starting with 800 by 600, then 1024 x 768, 1152 x 864, 1280 x 1024, 1400 x 1050 and finally 1600 x 1200. Folks who work with monitor resolutions a lot often use a shorthand, referring to a resolution with just the first number of each dimension - so that "10 by 7" refers to 1024 x 768.

Monitor sizes also got larger. Monitors in early 80's were typically 14", which meant that the diagonal display was 14" - with a viewable area of somewhat less, perhaps 12 1/2 inches. As technology, and, more importantly, manufacturing prowess improved, monitor sizes grew, first to 15", then to 17", 19", 21", and even larger. Monitor size by itself didn't have anything to do with resolution, although a 1600 x 1200 display on a 14" monitor would be unreadable, because the pixels would be so small. (Some of us who are getting on in years find that 1600 x 1200 isn't all that readable even on a 21" monitor, but that's a different issue altogether.)

The factors of screen resolution and monitor size do combine for a third important item - image resolution, which is measured in dots per inch (or pixels per inch), known as "DPI". Most monitors display images at 72 DPI (although some vary a bit, perhaps up to 80 DPI or so.)

Images are created on a computer by turning some of the dots in the grid to a different color. Early monitors only had two colors - black and a contrasting color, usually green, white, or amber (yellow). In order to display an image, some of the dots were one color and the rest of the dots were the other color. Without contrasting colors, you'd only be able to see a limited number of images, like "black cat in a coal bin at midnight" and "polar bear in a snow drift drinking milk".

Eventually monitors developed the ability to display more than just two colors in a dot. First, the dot could be any one of eight colors, which provided color displays, although it seemed that you were living in kindergarten world, where primary colors ruled the roost. The number of colors ("color depth") increased over time until most monitors had choices of "256 colors", "thousands of colors", or "millions of colors".

Regardless of what software tool you use, capturing a screen (I'll get to the mechanics of doing so shortly) will result in an image that has 72 DPI. If your screen resolution is 640 x 480, you'll get an image that is about 9" across by 7" tall ( $640 \text{ pixels} / 72 = 8.89$  and  $480 / 72 = 6.67$ ). If your screen resolution is 1280 x 1024, you'll get an image that is 17.8 x 14.2.

In either case, your image still has the same resolution, but the number of pixels in the image varies. If you can resize the image that is on the screen, you can get more pixels to represent the same amount of image. For example, suppose you're capturing the display of a Web page. If your browser enables you to change the size of the Web page (in Mozilla, for example, you can use Ctrl++ and Ctrl-- to change the text zoom), you can

Suppose a Web page was designed for 640 x 480 displays. Using a 640x480 res on your monitor, you could do a screen capture that would be 9 wide by 7 tall at 72 DPI, and the page would take 0.3 million pixels. If you used 1600 x 1200, that Web page would take up less than a quarter of the display, and thus would still be 9 wide by 7 tall at 72 DPI. If you resized the Web page to fill the 1600 x 1200 display, it would be 22" wide by 17" tall. More importantly, there would be 1.9 million pixels to represent the page.


If you are capturing dialogs in software packages, you're probably stuck with objects that can't be resized. As a result, changing the resolution of the display won't do anything to the image you're capturing. A dialog that takes up most of the screen at 640 x 480 will only take up a small part of the screen at 1600 x 1200, and thus still end up with the same number of pixels.

### 3. Capturing a screen shot using The GIMP

The GIMP (GNU Image Manipulation Program) is found under the Graphics menu on Fedora Core. One of its millions of features is the ability to capture a screen shot and place it in a GIMP screen automatically. You can capture either a single window or the entire display.

In order to capture a single window, follow these steps:


Step 1. Select the File | Acquire | Screen Shot menu option to bring forth the Screen Shot dialog as shown in **Figure 1**.



**Figure 1.** The Acquire Screen Shot dialog in The GIMP.

Step 2. Bump up the value in the "after 0 Seconds Delay" spinner to a couple of seconds. You'll need the delay in order to move from the Screen Shot dialog to the window you want to capture and make it the active window. If you don't, when you

make the capture, the window will have an unselected title bar as shown in **Figure 2**. I don't know about you, but that's generally not what I want.



**Figure 2.** A window captured with an inactive title bar display.

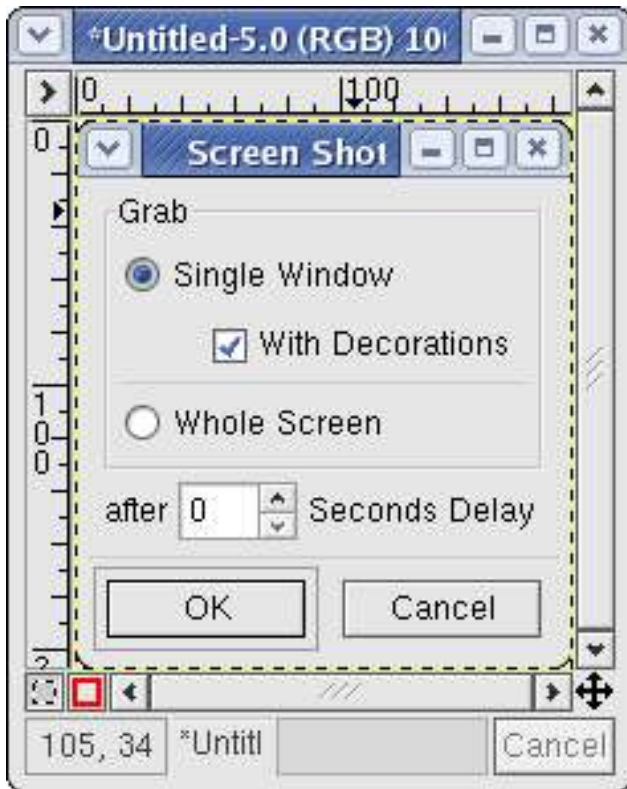
Step 3. Click the OK button in the Screen Shot dialog.

Step 4. Move your cursor to the title bar of the window you want to capture and click to make it the active window. You need to do this within the Delay interval you just set in Step 2. After a few captures, you'll either learn to set up the window you want to capture so that it's ready to be made active in a jiffy, or to give yourself more time with the Delay interval in order to set up the screen shot as you want it.

Step 5. Once the Delay interval is over, your mouse cursor will turn from its default shape into a cross hairs shape. Now click anywhere within the window to be captured.

Note that if you click in a different window, say the word processor you're using, or the file manager, or the terminal window that you've got open on your desktop, The GIMP will capture a screen shot of THAT window instead. Once you've got the cross hair cursor, the next object you click on will be captured.

Step 6. After you click in Step 5, a new GIMP-related window will appear, as shown in **Figure 3**.



**Figure 3.** A new GIMP window displays once a screen shot has been captured.

In Figure 3, you'll see a window within a window. The inside window, with the title bar that says "Screen Shot", is the dialog that I was capturing - as shown earlier. I am using dialogs from The GIMP as my examples, because I think it's appropriate to make this discussion self-referential. It's also fun in a geeky sort of way.

The outside window is The GIMP's window that contains the captured image. If I had captured something else, such as my email client, it would have looked something like **Figure 4**.

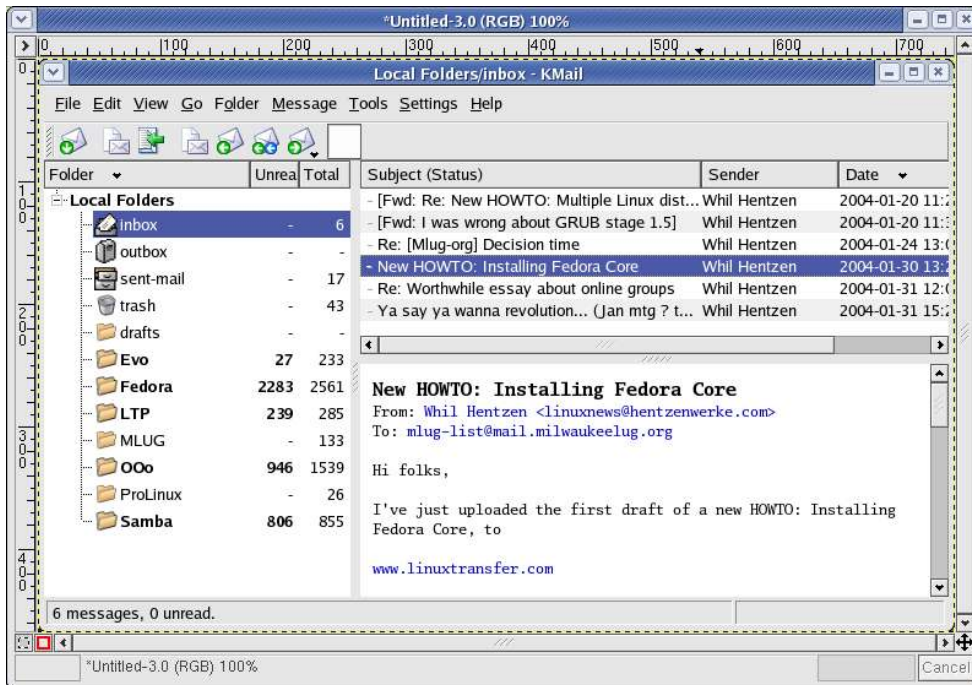


Figure 4. A different captured screen in The GIMP's image window.

This image is a 'life sized' image, captured at the resolution that your monitor is set at (such as 72 DPI). You can change the resolution inside The GIMP; I'll discuss how to do that later.

Step 7. Save the image to a file. Right-click in The GIMP's image window, and select the File | Save As menu option to bring forward the Save As dialog as shown in Figure 5.

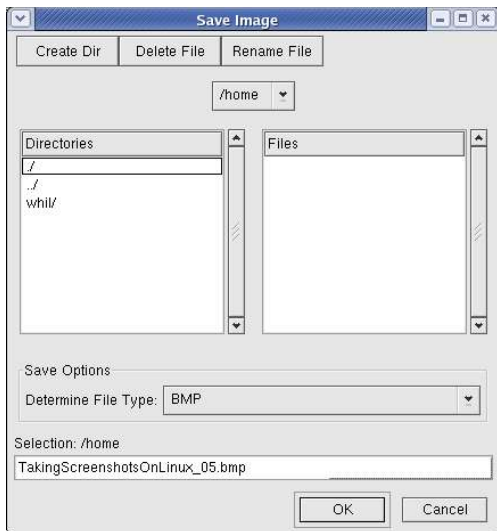


Figure 5. The GIMP's File | Save As dialog.

The Determine File Type drop down combo box will display "By extension" the first time you save a file. Once you select a File Type, that file type will be using for future file saves during the current GIMP session.

Step 8. Select the type of file you want to save the image as, using the Determine File Type drop down combo box.

Step 9. Type a file name in the Selection text box at the bottom of the screen, after navigating through your directory structure to find the appropriate directory, if desired.

Step 10. Click OK. A 'saving' thermometer bar will display and scroll across the screen as the file is saved.

So that's the mechanics of capturing a screen and saving it to a file. Now that you've got a capture, what do you do with it?

## 4. Working with a screen capture

Typically, you'll want to let other people see it (unless you're emulating Dominique from The Fountainhead). One common operation is to simply give it to another party to output. A second is to put it on a Web page. A third is to put it in a document yourself.

### 4.1 Give it to a printer

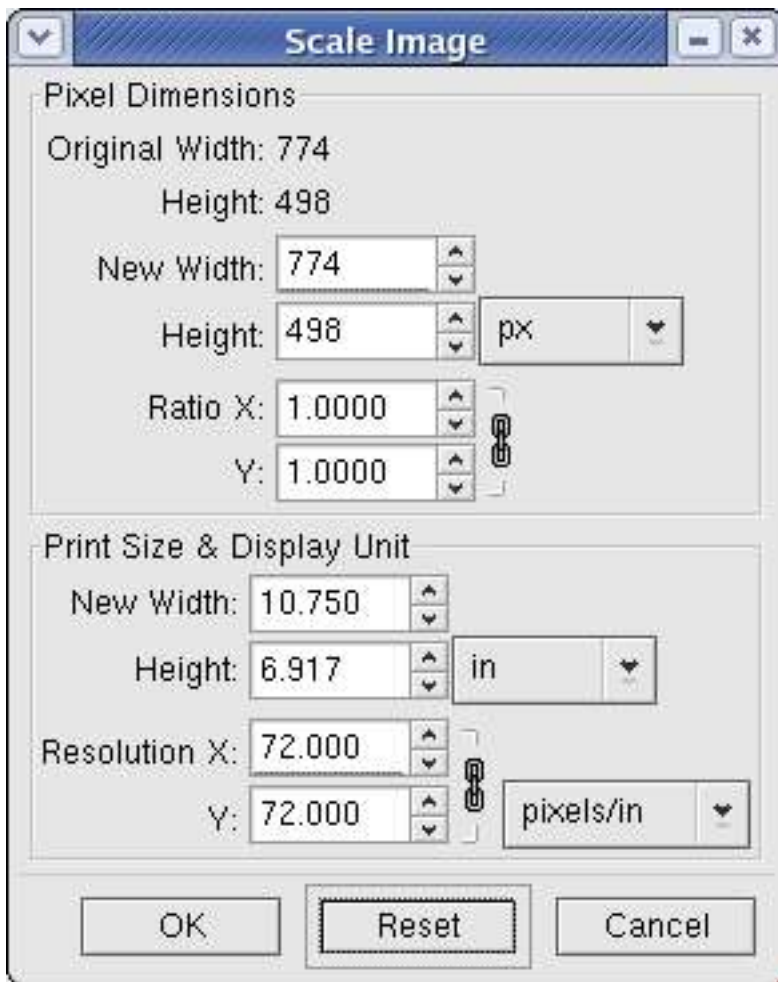
This is the easiest, actually. If they're going to deal with it, you might just give it to them 'as is' and have their experts do all the work. However, there are reasons you wouldn't do this. Their experts might not be as expert as they claim. Or they may not want to do the work - instead, they'll have specs that they require you to meet. A common requirement is to provide a higher resolution, such as 300 DPI, than the original capture is done in. Another common requirement is a specific file format. I'll cover file formats in another section.

To create a higher resolution, you'll end up making the image smaller. Remember, you've got X number of pixels in both directions. If they want you to change the DPI, then since the total "D" is going to be constant, the "I" will have to change. For example, suppose you've got an image that's 774 pixels across and 498 pixels tall. At 72 DPI, that's 10.75 inches across and 6.917 inches tall. If they want a 300 DPI image, then the result will still be 774 pixels across, but at 300 DPI, that's 2.58 inches (774/300) wide and 1.66 inches tall (498/300).

In order to create a higher resolution, follow these steps:

Step 1. Click on the Image window

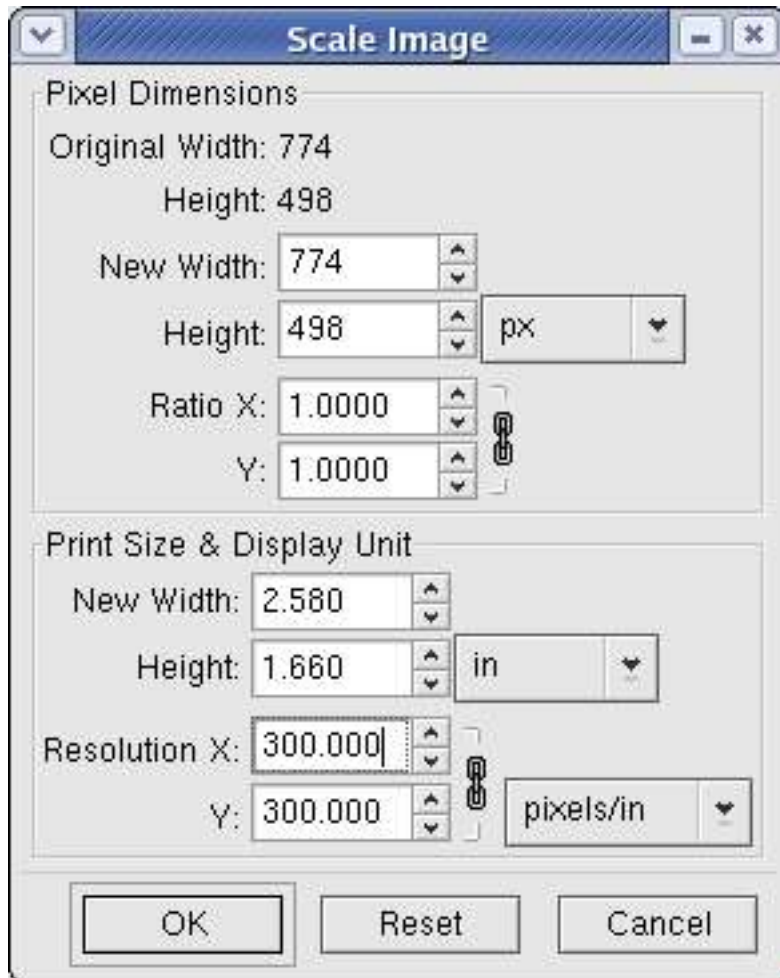
Step 2. Select the Image | Scale Image window to bring forward the Scale Image window as shown in **Figure 6**.



**Figure 6.** The Scale Image dialog allows you to change an image's size and/or resolution.

Step 3. Note the resolution of the output (Print Size) in the bottom half of the dialog, where it's set to 72 for both X and Y, and the two are linked (as denoted by the chain between the two spinners.) Change the X resolution from 72 to 300. When you tab out of the spinner, two things will happen.

First, the Y resolution will also change to 300. Second the New Width and Height will automatically be changed, as shown in **Figure 7**.



**Figure 7.** Changing the resolution of an image also changes the physical size.

You'll notice that the total number of pixels hasn't changed -  $10.75 * 72 = 774$  pixels =  $2.58 * 300$

Suppose your printer wants the image to be the original size, but at a higher resolution. In other words, they still want the image to be 10.75 inches wide by 6.9 inches tall, but at 300 DPI instead of 72 DPI?

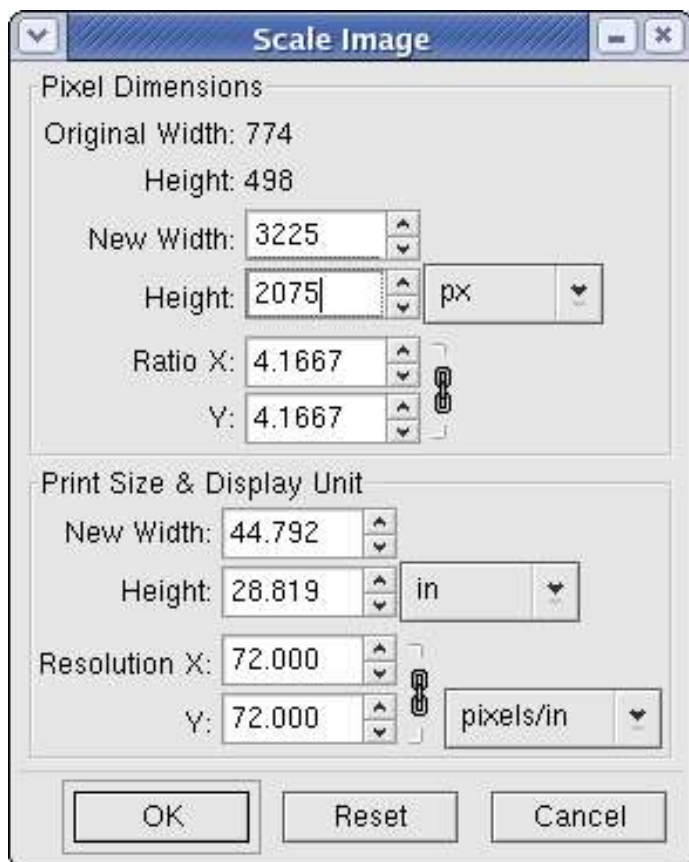
In essence, they're asking you to create something from nothing. You've got an image that's 774 pixels across. They want an image that will be 3225 pixels across (10.75 inches by 300 DPI). OK, you've got 774 pixels. Where do the other 2451 pixels come from? From thin air, of course! And The GIMP can do it, using a technique called "resampling".

In essence, what resampling does is create new pixels, based on the contents of nearby pixels. The mathematics and algorithms behind this operation are beyond the scope of this paper, but a simple example can illustrate the basic idea. Suppose you had an image where a completely white pixel was next to a completely black pixel. If you resampled the image such that a third pixel was created in between these two, the new pixel would be a medium grey. In reality, it's more complex than that, because more pixels than just the immediately adjacent pixels are included in the calculations, but you get the idea.

You can use other controls in the Scale Image dialog to resample an image. In order to create that 10.75 inch wide image at 300 DPI, you'd first resample to create all 3225 pixels. The resulting image would still be 72 DPI, and thus would be about 44 inches wide. Then you'd change the resolution from 72 to 300, which would shrink the image back down to 10.75 inches wide, and you'd have your 10.75 inch image at 300 DPI.

In order to resample an image, open the Sale Image dialog, and change the 774 pixel width to 3225 as shown in **Figure 8**.





**Figure 8.** Changing the Pixel Dimensions New Width value resamples the image.

Once you tab out of the New Width spinner, the values in the Height, Ratio X, Ratio Y, Print Size & Display Unit New Width and Height spinners will change as well.

Ratio X and Y show the proportion of the old and new images ( $3225/774 = 4.1667$ ). Since the resolution in pixels per inch hasn't changed, the actual physical size of the image has, and that's what you see in the Print Size & Display Unit New Width and Height spinners.

You could have accomplished the same thing by entering 4.1667 in the Ratio X and Y spinners, by the way.

Since your printer (in this hypothetical example) wants the image to be 10.75 inches wide, your last step is to change the resolution from 72 to 300. The New Width and Height will change back to 10.75 and 6.9, and that's what you're looking for.

## 4.2 Web page placement

If you're going to put the image on a Web page, then you've got another set of challenges. These are resolution, image size, and file size.

### 4.2.1 Resolution

First of all, the resolution you've already got - 72 DPI - will be fine. That's about as good as the Web pages are going to render on computers screens anyway. Whew! Finally one issue we don't have to worry about.

### 4.2.2 Image size

The second issue, image size, can get tricky. If you're going to display an image that has a lot of detail, as shown in **Figure 9**, you may have to display it full size on the page for the detail to be seen.

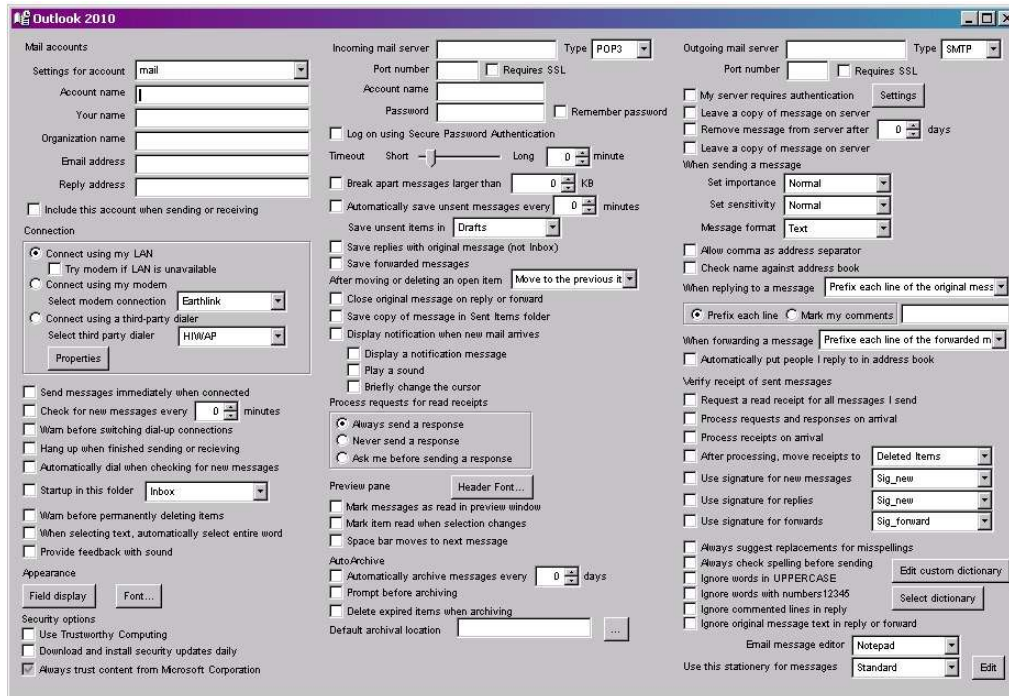


Figure 9. An image that's way too big to display in one piece.

However, the requirement of putting this entire image on one page begs the question as to whether or not that image is the proper one to display. A better method might be to create multiple images. One version of the image would display the image in toto, but where the detail isn't important. Subsequent versions would each show just a part of the image, so that details in that area could be seen clearly, sort of like Egyptian tablets that are discovered and translated in pieces. (The image shown in Figure 9, entitled "Microsoft's attempt at saving money by eliminating tabs in the Outlook preferences dialog" is a parody at jamming too much stuff into one screen.)

Let's suppose you've got a more reasonably sized image, such as a 7" x 9" book cover, but you still don't want to display it on your Web page as a full sized image. Rather, you'd like a thumbnail-sized image.

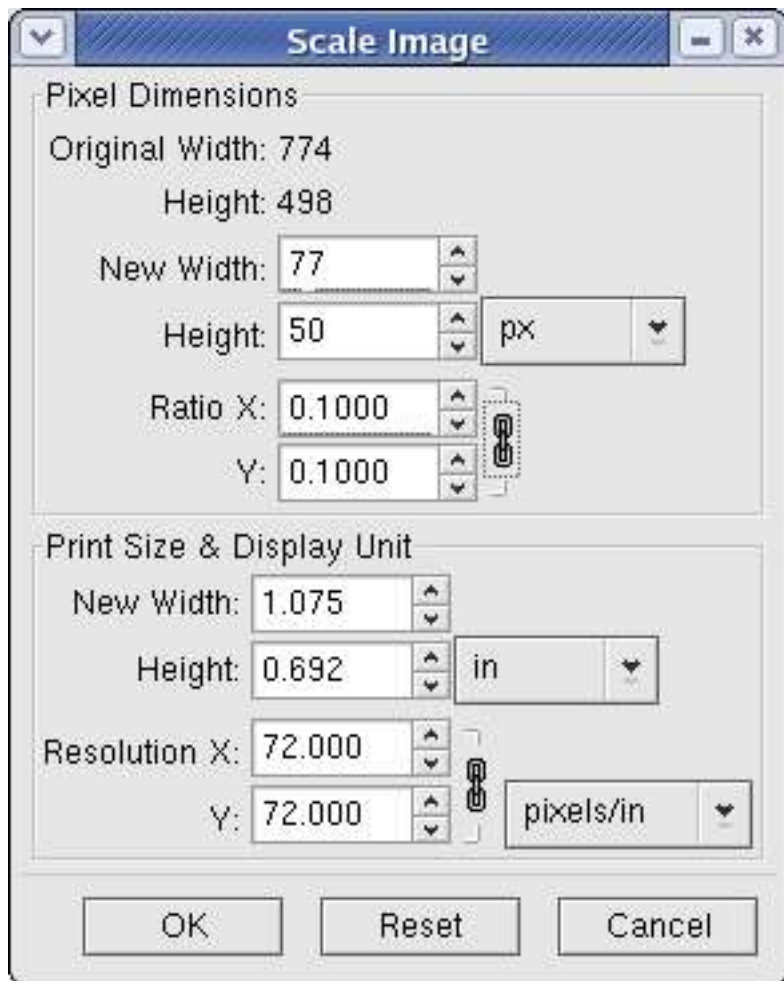
What you're trying to do is remove some of the pixels from the grid - instead of having, say, 774 pixels across at 72 DPI, you want 77 pixels across at 72 DPI. Naturally, you're going to lose a lot of detail (so if there's small text in the image, it'll get very blurry), but that's to be expected since you're making the image smaller.

Follow these steps:

Step 1. Click on the Image window

Step 2. Select the Image | Scale Image window to bring forward the Scale Image window as shown in Figure 9.

Step 3. Change the Ratio X value from 1.0 to 0.1, as shown in Figure 10.



**Figure 10.** Changing the ratio changes the number of pixels in the image.

Once you tab out of the Ratio X spinner, you'll see that the New Width and Height values change (from 774 to 77 and from 498 to 50). The Print Size New Width and Height also change, from 10.75 to 1.075 and from 6.92 to 0.692.

Step 4. Click OK and the image will be resized. You'll need to save the image, of course.

#### 4.2.3 File size

Until we've all got multi-gigabit Ethernet connections running helter skelter through our houses (as well as places of business), file size will be an issue. It takes longer to download a 400K image than a 20K image, and not all of us drink coffee (to give us something to do while waiting), so you'll want to keep the file size down to a minimum.

By resizing the image as shown in the previous section, you're removing pixels, and thus, data. The resulting image file will also be smaller, automatically.

### 4.3 Placing in a document

If you're going to put a screen shot into a document file (such as a OpenOffice.org Writer file) that will eventually be printed, you're probably going to want better than 72 DPI resolution. Typically, 300 DPI will suffice.

The novice will insert an image into their copy of OpenOffice.org using the Insert | Graphics | From File menu option, navigate to the file of choice, and click OK. Once the image is placed in the file, green sizing handles appear at all four corners and in the middle of each side., as shown in **Figure 11**.

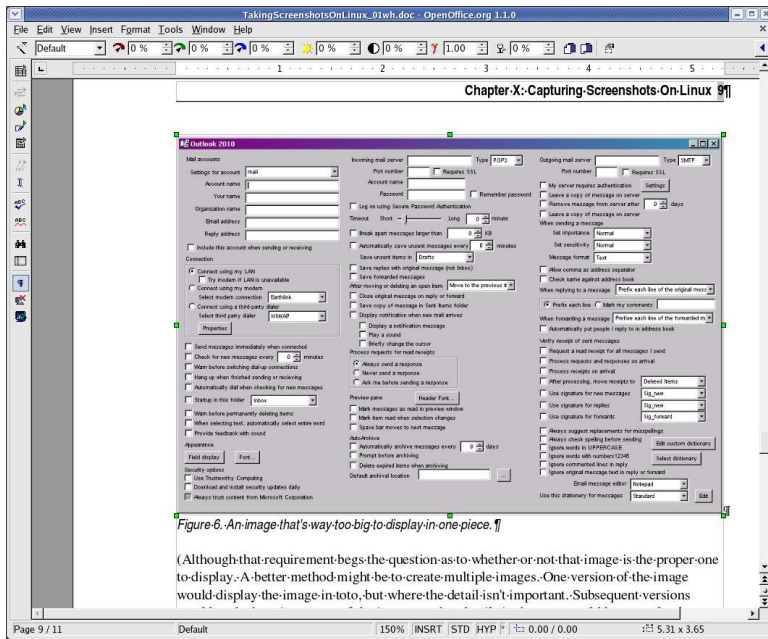


Figure 11. An image inserted into an OpenOffice.org document and then selected displays green 'sizing handles' at the corners and edges.

These can be used to change the size of the image right there in the document, as shown in Figure 12.

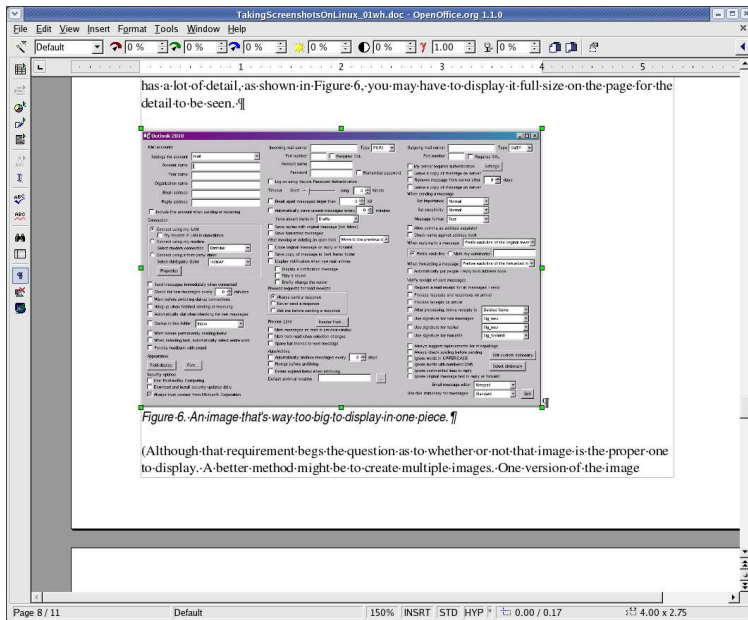


Figure 12. An image resized in an OpenOffice.org document via the green sizing handles.

However, this isn't going to accomplish the task. Instead, you'll need to use the Scale Image controls as described earlier to create a 300 DPI image that's the proper size for the document. For instance, the images in this whitepaper are usually 4 inches wide in order to make sure they fit in the margins but still are big enough to see the details.

Note that the Insert Graphics function in OpenOffice.org doesn't always place the image in the document at the original size. For instance, if you're inserting an image at the bottom of the page, occasionally OpenOffice.org will automatically resize the image for you in order to fit in the space remaining on the page, creating a very small image. In order to return the image to its original size, select the image so that the green sizing handles appear. Then right-click and select the "Graphics" menu, select the Type tab, and click the Original Size button on the left side of the dialog.

## 5. Where to go for more information

This free whitepaper is published and distributed by Hentzenwerke Publishing, Inc. We have the largest lists of “Moving to Linux”, OpenOffice.org, and Visual FoxPro books on the planet.

We also have oodles of free whitepapers on our website and more are being added regularly. Our Preferred Customer mailing list gets bi-monthly announcements of new whitepapers (and gets discounts on our books, first crack at special deals, and other stuff as we think of it.)

Click on “Your Account” at [www.hentzenwerke.com](http://www.hentzenwerke.com) to get on our Preferred Customer list.

**If you found this whitepaper helpful, check out these Hentzenwerke Publishing books as well:**

**Linux Transfer for Windows® Network Admins:  
A roadmap for building a Linux file and print server  
Michael Jang**

**Linux Transfer for Windows® Power Users:  
Getting started with Linux for the desktop  
Whil Hentzen**