

Linux Graphical Disk Space Tools

By Whil Hentzen

While hard disk space is cheap, it isn't free, and it often comes up short at the most inconvenient times. Just what is taking up so much space on your hard disk and why the #&\$@! do you only have 26K left in ~? Here are some tools that provide a graphical view of the files and directories on your hard disk, and, in some cases, on all network drives as well.

1. Preface

1.1 Copyright

Copyright 2007 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	2007/11/10	Original	WH

1.2.2 New version

The newest version of this document will be found at 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 'booksales@hentzenwerke.com'. I also welcome suggestions for passages you find unclear.

1.3 Acknowledgments

Thanks to MLUG members Jonathan Detert, Chad Kittel, Tom Peters and Scott Zylstra for pointers to various tools they've used or run across, as well as the peanut gallery at the October MLUG meeting for additional info.

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 6.0 (and KDE), but the tools discussed here are available on most modern distros.

2. The need for a graphical disk space tool

There are thousands of files in hundreds of directories on a typical Linux distribution. When disk space runs short, it can be due to one or more reasons. One is a single file that's suddenly grown huge, perhaps a log file that's run amok, or a mail spool that filled up all of a sudden. Another is having a large, large, large number of small files in one or more folders, again

created by a process or program that has gone astray. Or perhaps your autodownloader found and grabbed every concert of your favorite band available on the Torrent, which would quickly fill up a few gig (well, if it's Pearl Jam, make that petas) without you knowing.

At any rate, it can be difficult to find the culprit.

You can get a listing of files and manually try to find the anomalies - the directory with 15,203 files, or the one file in a 'ls' listing that has a length that's four times as long as any other file in the list. This can get tedious, though, and is prone to error. You could try to automate the process with an intricate command or a short script, if you've got the knowledge and the patience to do so. But if your skillset is targeted elsewhere (that's French for "You don't know how to do it."), an alternative is desirable.

A graphical disk space tool can be just the ticket to help you find anomalies and reclaim your space.

3. KDiskFree

KDiskFree is a KDE applet that displays disk space on each partitions. Fire up KDiskFree via the

System | More Applications | KDiskFree

menu in Fedora Core 6 and you'll get a summary of each partitions, as shown in **Figure 1**.

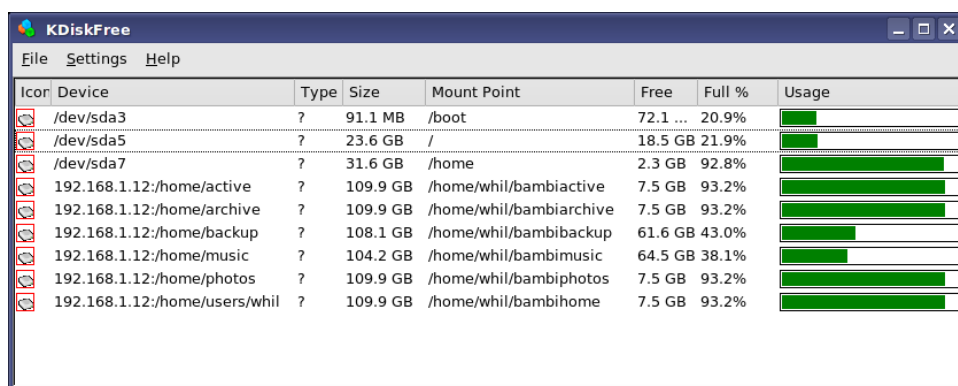


Figure 1. KDiskFree shows you space on each partition.

You can configure KDiskFree through the KwikDisk utility, also found under More Applications. It fires up the dialog shown in **Figure 2**.

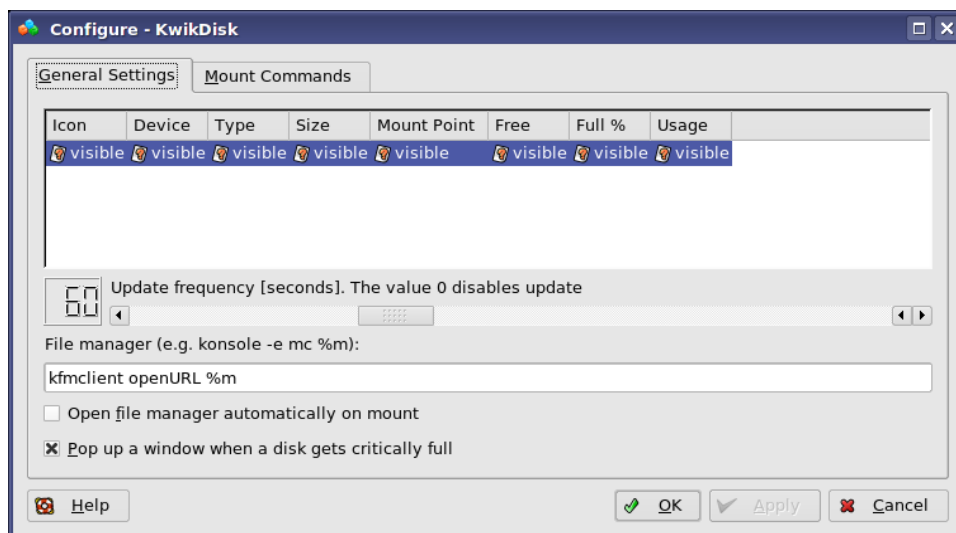


Figure 2. KwikDisk allows you to configure KDiskFree.

While KDiskFree is useful, it's still tough ("impossible") to see what is taking up the space on each of those partitions. So I kept looking. One day I ran across a nifty graphical disk space utility that showed all of the files in a folder - displayed graphically by relative size. I was looking for something similar to WinDirStat, a graphical disk utility for Windows:

<http://lifehacker.com/software/disk-space/geek-to-live--visualize-your-hard-drive-usage-219058.php>

The Linux version, as I recall, was uglier, but who cares. I mean, you've seen my wife....

Being a greybeard, I didn't write the name of the tool down, and quickly forgot what it was called, so one day when I wanted to use it, I posted a note on my friendly MLUG mailing list, and quickly received a half-dozen responses. Here is a synopsis of those responses.

4. Konqueror

Scott Zylstra says, "In Konqueror I can choose View Mode: 'File Size View'." **Figure 3** shows Konqueror with the View Mode set to Detailed List View, but being switched to File Size View.

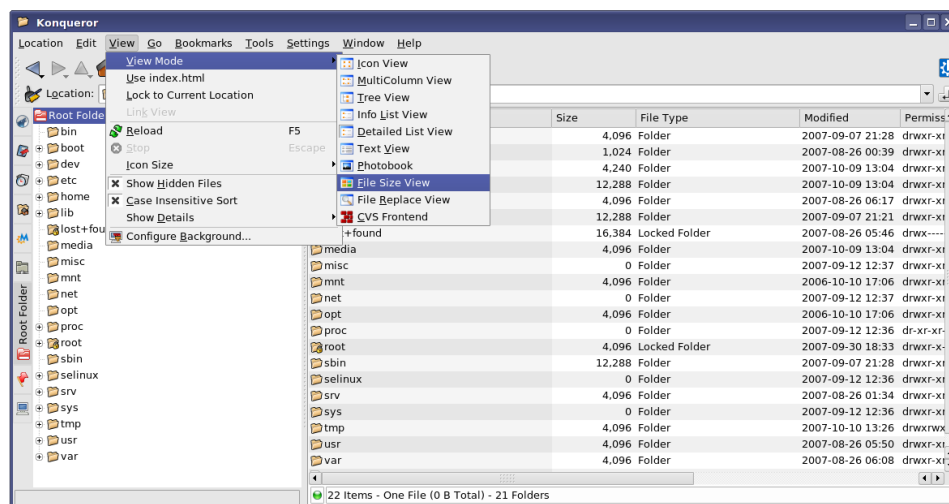


Figure 3. Switching Konqueror's View Mode to File Size View.

Figure 4 then shows the music folder displayed with File Size View.

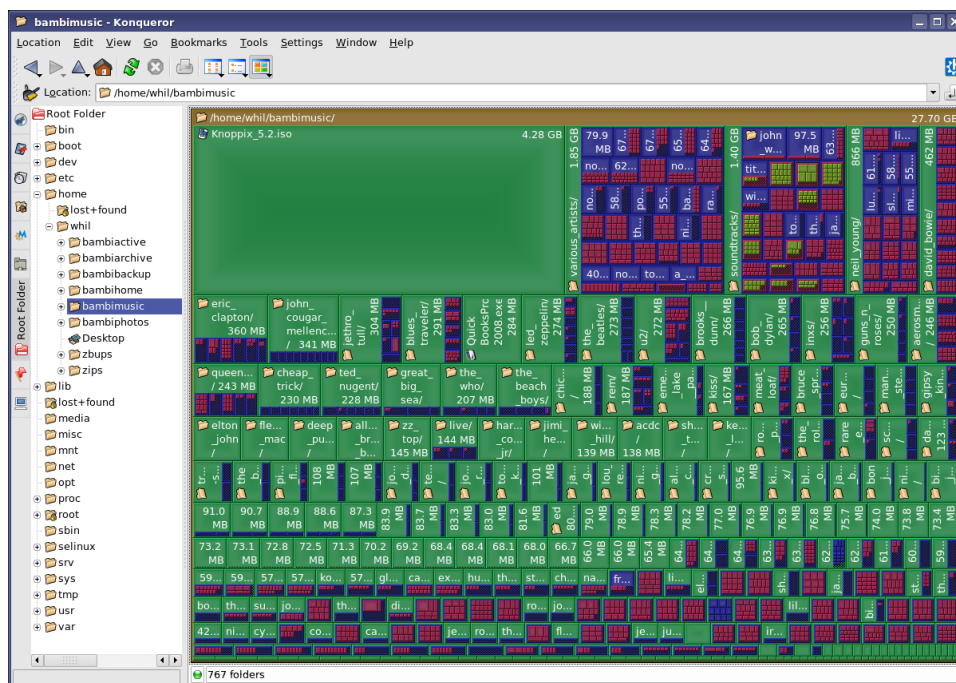


Figure 4. File Size View in Konqueror.

The colors show at what level down the directory tree a particular item resides. Figure 4 shows the contents of my music partition. You can see oodles and oodles of folders, each representing a single artist (although assigning that label to some of my kid's music creators is a questionable call.) I had temporarily parked the my Knoppix 5.2 DVD ISO file, all 4.3 GB of it, in the root of that partition, and you can see it taking up 15% of the partition. It's green alongside the folders of each artist because it's on the same level. You can see individual files under a green folder in some cases.

The folders are organized according to size, so you can see that "Various Artists", "Soundtracks", "Neil Young", "David Bowie", "Eric Clapton" and JCM have the largest shares in my collection.

You can drill down into a folder simply by double-clicking on the box of interest, as shown in **Figure 5**.

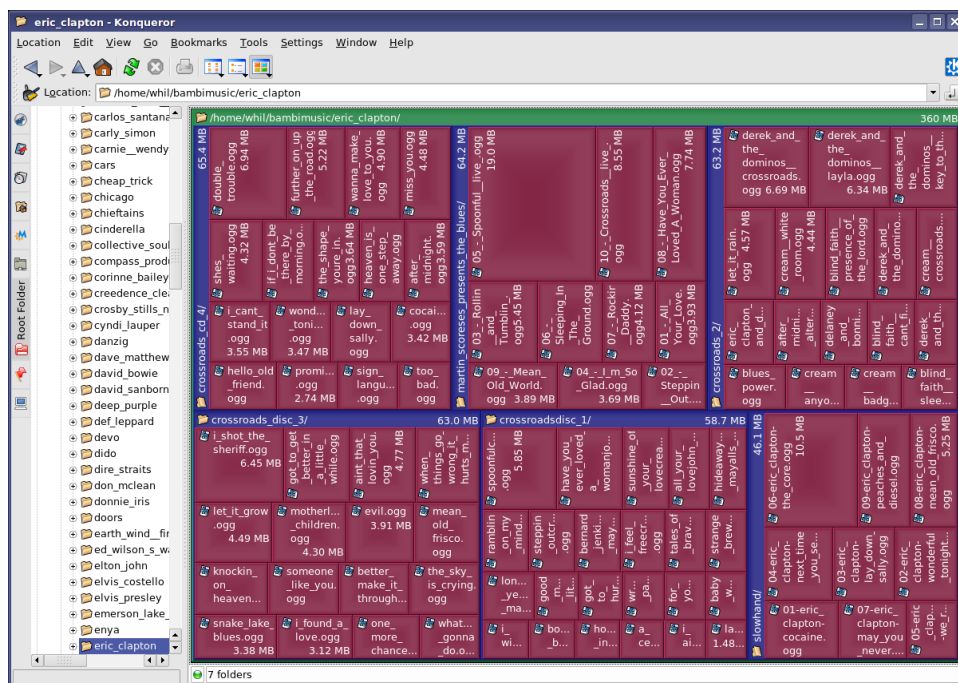


Figure 5. Drilling down into the Eric Clapton folder of the music partition.

As you can see, the Eric Clapton folder contains six folders; one each for the Crossroads four-CD set, and two more for "slowhand" and a blues compilation. In each folder, the red boxes represent files. You can see the live Spoonful cut in the blues compilation taking up a LOT of space. You can also see that files are organized according to size, just as folders were in Figure 4.

This tool works out well, because you can use the tree view in the left pane to navigate through the folder structure. It can be difficult sometimes to find the right place to drill down if you're just using the graphical view.

You know what's interesting? It turns out there's a standalone program, `fsview`, for which Konqueror simply acts as a wrapper. In other words, you can go to a terminal window, enter "`fsview`", and you'll get the same results - but in its own window, not inside Konqueror's right pane.

5. Kdirstat

Scott and Tom Peters both recommended `kdirstat`, which shows a lot of info! It can be found under the System menu. If it's not there, it can be added to your system via System | Add/Remove Software.

Upon execution of `KDirStat`, you'll be prompted to select a folder, as shown in **Figure 6**.

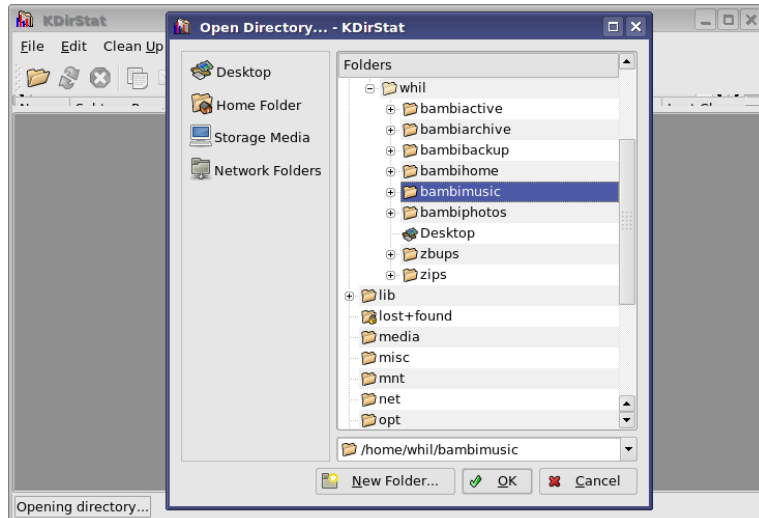


Figure 6. KdirStat asks for a starting directory, which can be bypassed if you like.

If you click Cancel, the window behind the Open Directory dialog will be displayed without data. However, selecting a directory will cause KDirStat to run and display the contents of the directory as shown in **Figure 7**.

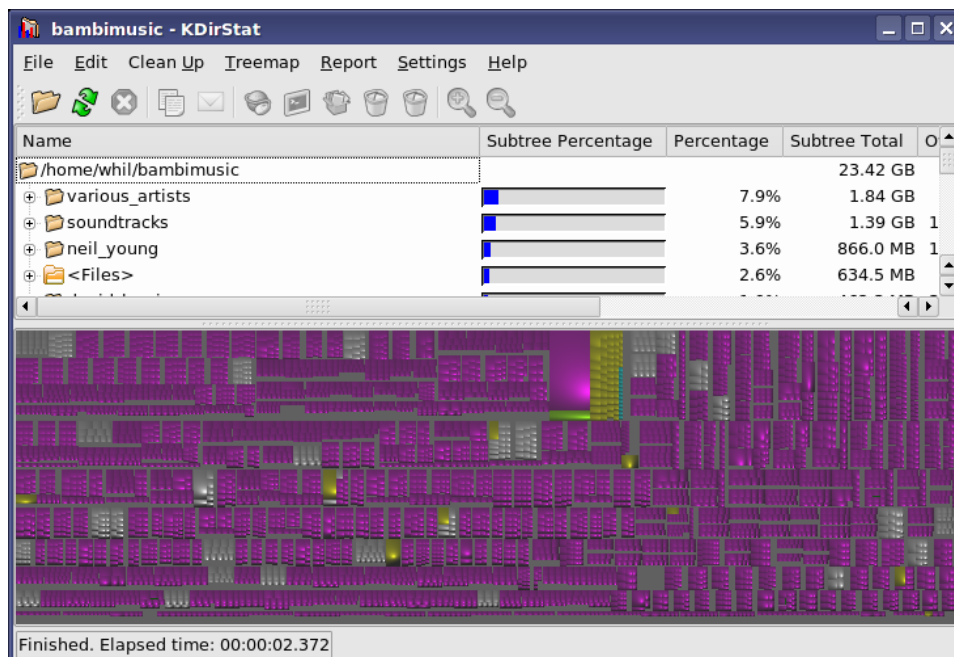


Figure 7. KdirStat's execution results in a combination tree view/bar-chart and graphical view of the selected folder.

The top half displays a list of folders and files while the bottom half shows a graphical representation of the same directory. Double-clicking on an item in the list in the top half drills down into the list, displaying the contents of that folder as shown in **Figure 8**.

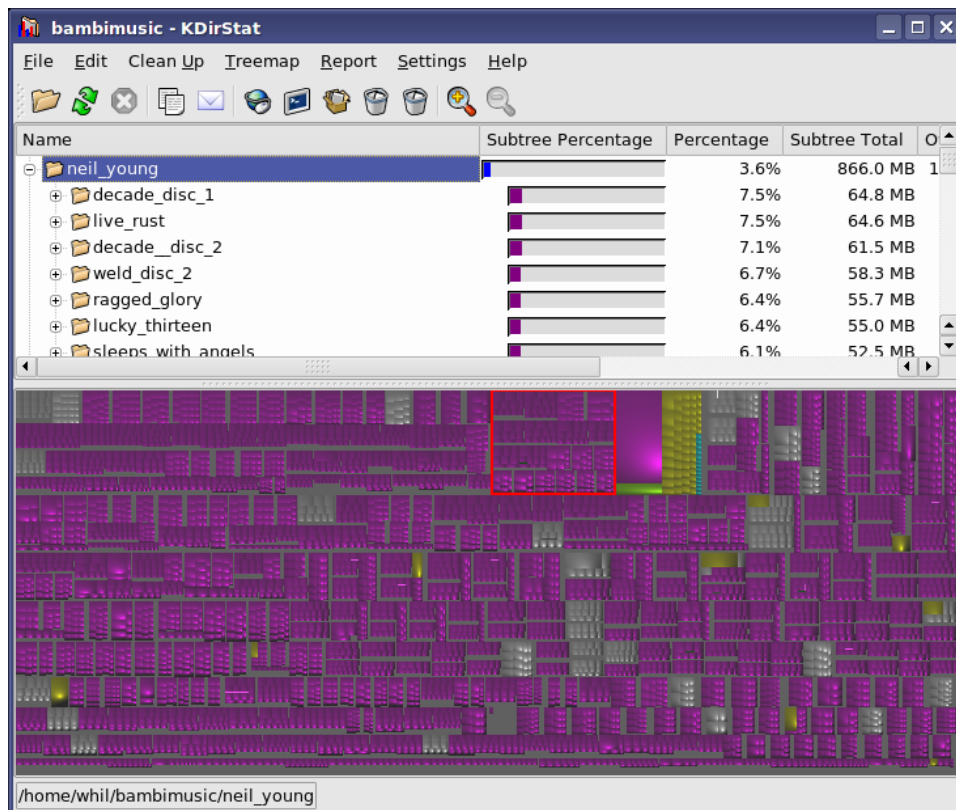


Figure 8. Selecting an item in the tree view in the top pane highlights that area in the bottom pane.

In addition, the area selected also is highlighted in the graphical view in the bottom half. As you can see, the Neil Young directory makes up a significant portion of the music collection, as it ought to for all true music lovers.

6. gdmmap

Jonathan Detert points to gdmmap: "I was intrigued by Whil's question, and did this:

```
apt-cache search disk | grep -iE 'space|map|visual|show|usage'
```

which told me about gdmmap, which I installed, and which also looks like it is what you want. The point is, you can use apt (or maybe yum) to help you find an app whose name you've forgotten."

gdmmap is also part of the standard FC distribution, also available through the Add/Remove Software menu item. As with kdirstat, installation is a matter of a mouse click and 30 seconds. (If you haven't used Fedora's Add/Remove Software to grab packages you're not familiar with, select System | Add/Remove Software, click on the Search tab, and enter the filename or a word that would be in the description - like 'visual' - as shown in **Figure 9**.)

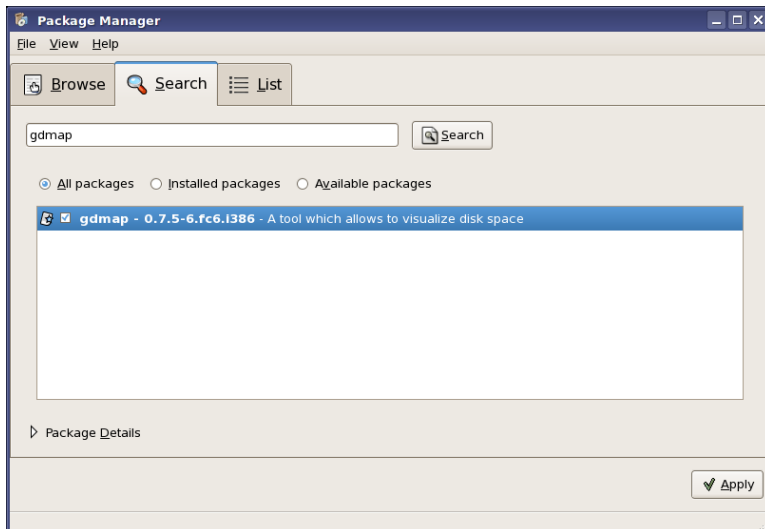


Figure 9. Searching for *gdmapp* in Fedora's Add/Remove Software dialog.

Doing so installs *gdmapp* and creates the System | Graphical Disk Map menu item. Run it, select File, Open, then choose a directory for which you want a map. It'll take a moment or two (or twelve). You'll see the 'Scanning' progress bar in the lower right, as shown in **Figure 10**.




Figure 10. After startup, *gdmapp* will scan your drive to collect size information.

At first I was distressed at the amount of time it was taking to open up, and then I was presented with the results. It turns out it scanned my entire machine - including all network drives mapped under my /home directory - in about 3 minutes. Not bad for a first pass at nearly 200 GB of data.

The graphical view results are shown in **Figure 11**.

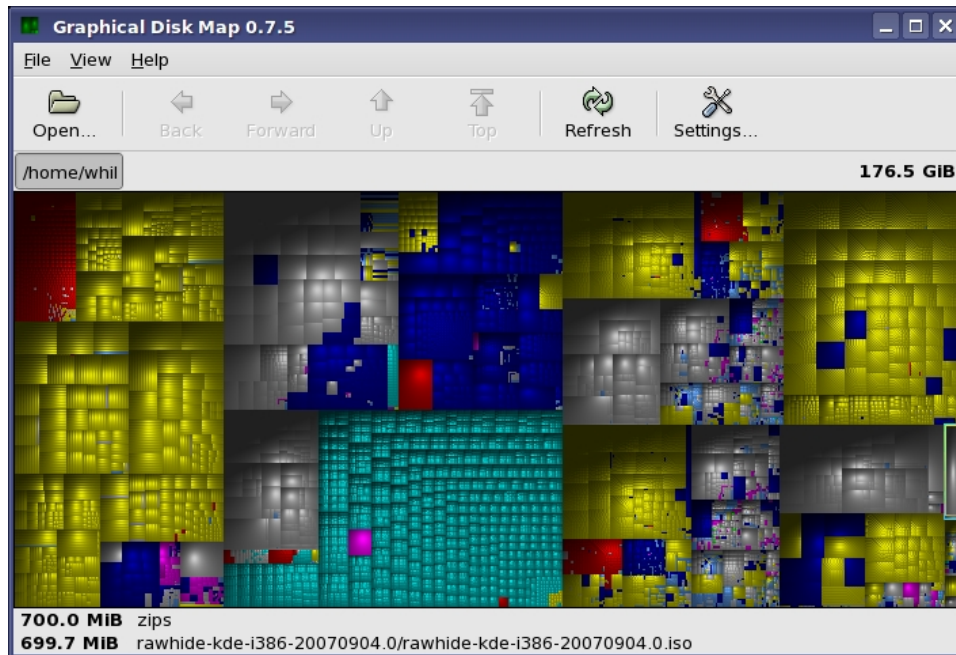


Figure 11. *gdmmap displays a visually busy, yet appealing, view of your disk.*

As with many Gnome applications, the interface is sparse and not always intuitive. So what to do? Just run your mouse around the screen and click randomly, that's what I always do. It turns out that you can select a box in the form to display its identity in the status bar underneath. In Figure 11, I've clicked on a box on the right edge somewhat below halfway. This file turns out to be a 700 MB .ISO, as shown in the grey status bar below.

This tool is probably the ticket if you're looking for a specific overly large file that's clobbering your usage percentages.

7. xdiskusage

Chad Kittel asked "Was it 'xdiskusage'?" and ran the following command:

```
chadk@roam ~ $ lvs what xdiskusage
```

According to the abstract on sourceforge, 'xdiskusage' is a user-friendly program to show you what is using up all your disk space. It is based on the design of xdu written by Phillip C. Dykstra. Changes have been made so it runs "du" for you, and can display the free space left on the disk, and produce a PostScript version of the display. It can be found at

<http://xdiskusage.sourceforge.net/>

I didn't venture into downloading and installing; the following screen shots are from the doc on the website. **Figure 12** shows that your first step will be to select a partition to examine.

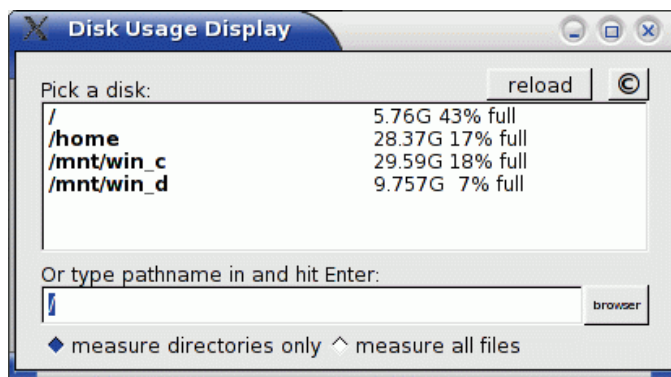


Figure 12. Select a partition to examine with *xdiskusage*.

After doing so, you can drill down into folders, as shown in **Figure 13**.

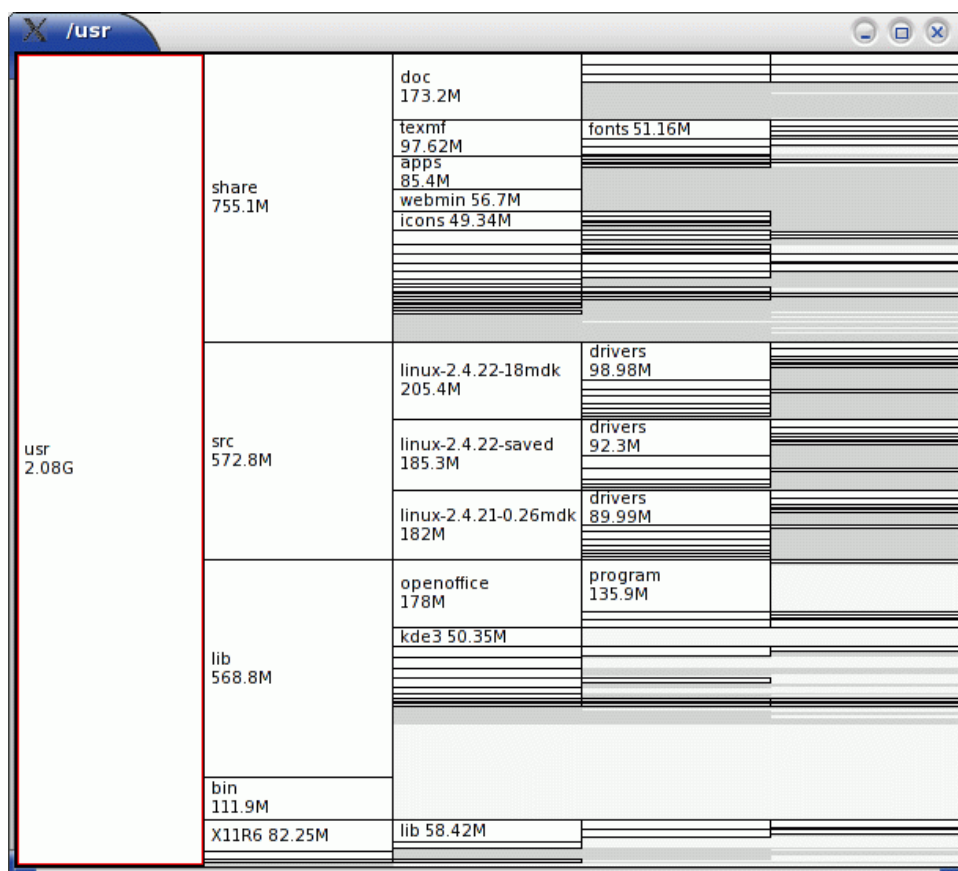


Figure 13. Drilling down into specific folders shows you a hierarchy of folders.

So for those of you who are color-averse, *xdiskusage* might be your ticket to disk space salvation.

8. fsv

And another entry was 'fsv', also found on sourceforge. This one, however, is pretty weird. Weird enough that it was used in the MOVIES. (I'll identify the movie at the end of this section.)

fsv ("file system visualizer"), pronounced 'eff-ess-vee', lays out files and directories in three dimensions, not just two. From the description on sourceforge: "It geometrically represents the file system hierarchy to allow visual overview and analysis. fsv can visualize a modest home directory, a workstation's hard drive, or any arbitrarily large collection of files, limited only by the host computer's memory and graphics hardware." Found at

<http://fsv.sourceforge.net/>

Two different 3D representation schemes are offered:

1. MapV view. This represents files and directories as rectangular blocks, all of equal height, and each with area proportional to the size of the corresponding file. Everything is laid out on top of its parent directory, somewhat like a Venn diagram in 3D. In the sample at left, the yellow blocks are regular files, and the grey blocks underlying them are directories. The root directory is the block at the bottom.

2. TreeV view. Built around a more conventional tree/leaf paradigm, this view represents directories as interconnected platforms, with leaves (files et al.) sitting on top. Everything is arranged concentrically, with the root directory closest to the center and subdirectories farther out. The leaf blocks all have the same footprint, and vary in height according to the corresponding file size.

fsv combines either of these 3D views with a standard 2D directory tree / file list interface, offering the best of both viewing paradigms. Directories can be expanded (shown) and collapsed (hidden) at will, allowing as much or as little of the file system to be visible at any given time.

The following screen shots are taken from the sourceforge documentation. **Figure 14** shows the contents of the linux folder on the box being analyzed.




Figure 14. A set of folders under a selected location.

Drilling down into the drivers area is shown in **Figure 15**.

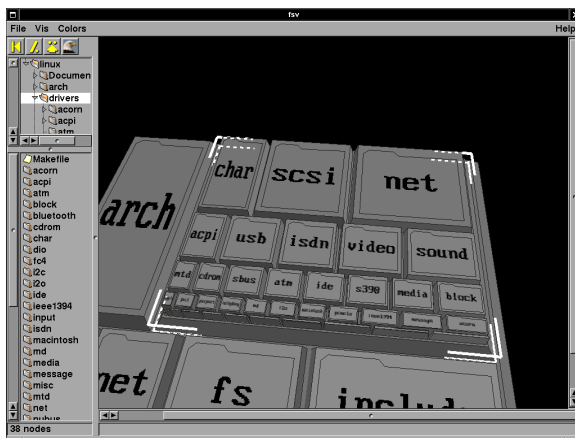


Figure 15. The drivers folder being opened.

Drilling down into the video folder is then shown in **Figure 16**.

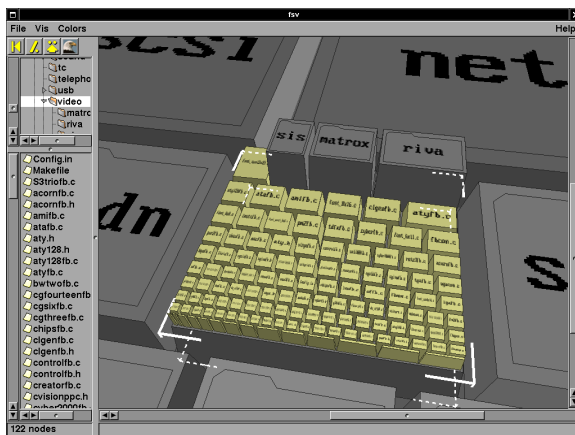


Figure 16. Drilling into the video folder under drivers.

Viewing the whole disk subsystem at once, shown in **Figure 17**.




Figure 17. A birds-eye view of the system.

But sometimes this type of view is boring. That's why fsv gives you another perspective on the whole thing. **Figure 18** shows you the disk space layout in a rounded tree view type of visual.

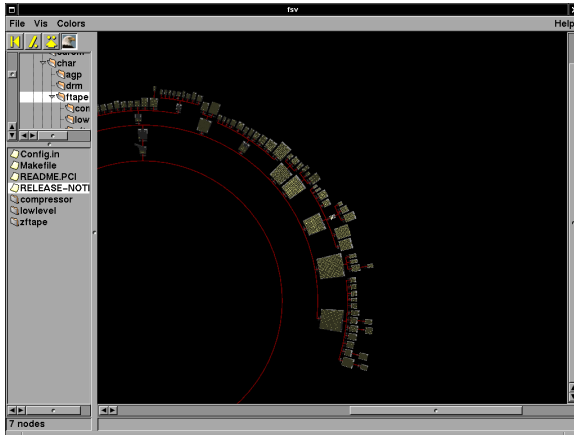


Figure 18. A different type of hierarchy.

If you've stayed with us this long, you may already recognize the parents of fsv. It was used on SGI boxen in the movie "Jurassic Park" - yes, that lame scene where the 11 year blonde girl exclaims, "I know how to use this! It's a UNIX system!" just like every 11 year old cheerleader-wannabe should. But watching the lawyer get eaten was recompense enough.

9. Where to go for more information

The first place to look is your very own hard disk - do a search on utilities for 'disk space' or other interesting terms in your Add/Remove Software mechanism. Source Forge is the number one destination for open source projects, and there are oodles of tools related to this feature; you may find one that fits your requirements better than these discussed. And go back through this article and try out some of the other search mechanisms that contributors offered.

10. About the author

Whil Hentzen started out life in the early '80's as a custom software developer using dBASE II (he still has the original 8 1/2 x 11 grey binder of documentation, much to the chagrin of his wife), and switched to FoxPro in 1990. Besides billing 15,000 hours in the 90's, he presented more than 70 papers at conferences throughout North America and Europe, edited FoxTalk, Pinnacle Publishing's high end technical journal for 7 years, hosted the Great Lakes Great Database Workshop since 1994. He's written 8 books and published 30 more on a variety of software development topics. He was a Microsoft Most Valuable Professional from 1995 through 2003 for his contributions to the FoxPro development community, and received the first Microsoft Lifetime Achievement Award for Visual FoxPro in 2001.

Whil began using Linux on the desktop when OpenOffice.org became a standard in the mainstream distributions, as it spelled potential for custom application development in the future, and has been a Linux user, developer, and evangelist ever since. His first book on Linux, Linux Transfer for Windows Power Users, was published in early 2004.

He is available for new and legacy Visual FoxPro application development as well as Web and desktop development on Linux.

11. A word from our sponsor

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 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