



Finding obsolete files with Agedu

THE COMBER

Agedu helps you clean up your hard drive by identifying old files.

Free up wasted space instead of creating wasted time.

BY ERIK BÄRWALDT

view of directories containing files that have not been accessed for an extended period of time. Agedu additionally combines the timestamp function with *du* disk space statistics for the files in each directory.

Installation

Agedu is only available from the developer's website as a source code package [1]; as of this writing, the version number was 8442. After downloading the latest tarball, unpack it in a temporary directory by typing *tar -xzf agedu-rversion.tar.gz*, substituting the version number of your copy of agedu. Change to the directory and install the tool with the *./configure*, *make*, and *make install* commands; you need to be root for the last step. When you are done, delete the tarball along with the temporary *agedu-rversion* directory.

To receive meaningful output, you need to launch Agedu twice. In the first round, the software scans the disk to check directory usage. For this, type *agedu --scan /directory* to tell Agedu to scan a directory and its subdirectories and create an index file. Depending on the position of the directory in the directory tree and the number of files it contains, this can take a while. In the second round, Agedu evaluates the

index file before displaying the results in your web browser as a colorful bar chart. What's more, Agedu includes its own lean web server that even has its own authentication mechanisms. To evaluate the index file and then launch the web server, launch Agedu with the *agedu --web* command (Figure 1).

The software enables its internal web server and uses an IP address from the localhost address space (typically 127.164.152.163). The program uses a different port each time you launch it. This prevents unauthorized machines on the local network accessing the web server.

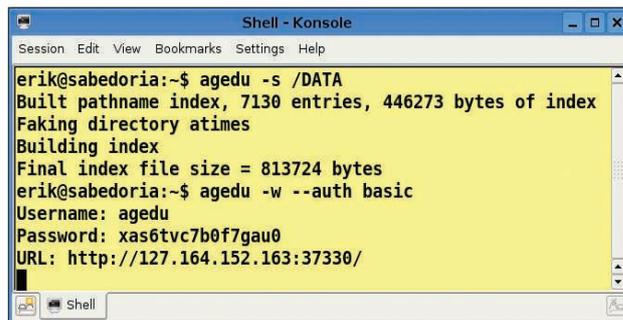
If the authentication feature is enabled, Agedu additionally generates a password whenever the web server is launched and displays the password in the terminal window.

The complete URL is displayed in the terminal and right-clicking with the mouse followed by a left-click on the *Open linkn* item launches the system's default web browser, which then displays the bar chart (Figure 2).

In this multimedia age, home computer hard disks always seem too small. Giant files quickly fill up drives that somehow seem to have a very short half-life. The consequence is too many time-consuming disk cleaning sessions.

Depending on your choice of window manager, Linux offers a variety of graphical tools that show your disk and directory space usage. Nearly all of these tools are based on the *du* command-line utility, which provides statistics for individual directories. The downside of *du* is that it is hard to tell which subdirectories have more than their share of obsolete files.

Simon Tatham, a developer from the UK, looked into this problem and came up with a tool named Agedu that reliably identifies unneeded file ballast. Agedu reads the Linux filesystem's internal last-access timestamp for each file and displays the results as an easy-to-read bar chart in a web browser. This display gives users a reliable at-a-glance over-



```

erik@sabedoria:~$ agedu -s /DATA
Built pathname index, 7130 entries, 446273 bytes of index
Faking directory atimes
Building index
Final index file size = 813724 bytes
erik@sabedoria:~$ agedu -w --auth basic
Username: agedu
Password: xas6tvc7b0f7gau0
URL: http://127.164.152.163:37330/

```

Figure 1: Two calls tell Agedu to generate an index and launch the web server.

Agedu supports a number of optional parameters that reflect the program's feature scope. Threshold values let you define how old files have to be for Agedu to consider them obsolete; instead of the last access date, you can choose the last change date as the criterion for the statistics. On request, the tool also performs a scan over multiple filesystems.

The IP address and port number for communicating with the web browser can be defined individually, and you can choose an authentication method for the internal web server. The user name and password can be defined by passing in parameters. If needed, Agedu will display a text version of the scan results in a terminal window, instead of sending

graphical output to the browser. Because the software lets you combine parameters, multiple steps can be combined in a single command; Agedu will process the steps sequentially.

Conclusions

Agedu gives you a quick overview of accumulations of obsolete files in the directories on your filesystem. The clear-cut graphical bar graphics save time searching.

Additionally, the software is customizable to reflect your requirements and preferences, but the huge range of options will see even experienced users scuttling to read the very detailed man page [2] before using the program. ■

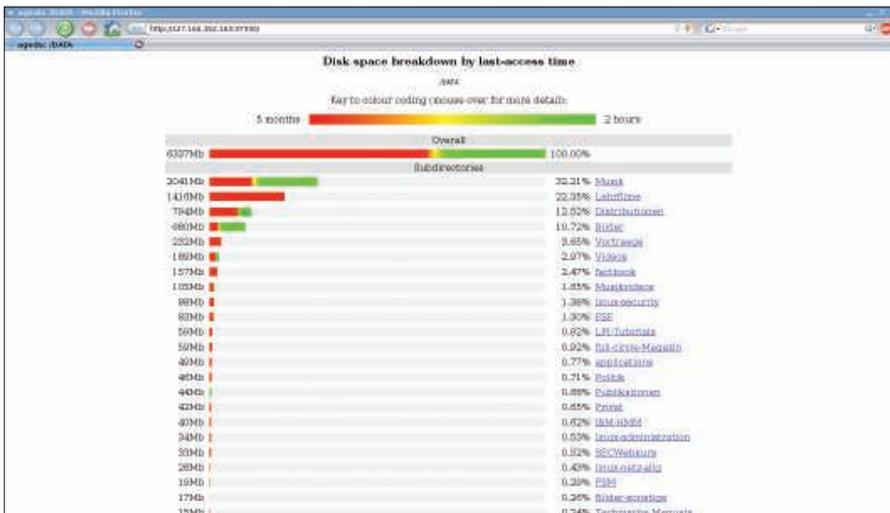


Figure 2: View file and directory statistics in Firefox.

INFO

- [1] Agedu homepage:
<http://www.chiark.greenend.org.uk/~sgtatham/agedu/>
- [2] Agedu man page:
<http://www.chiark.greenend.org.uk/~sgtatham/agedu/manpage.html>

WANT TO KNOW WHAT'S UP NEXT?

**SUBSCRIBE TO LINUX
MAGAZINE PREVIEW,
OUR FREE MONTHLY
EMAIL NEWSLETTER!**