

Formatting Disks and Memory Sticks

Types of memory device

I will refer to all means of data storage, external to your computer, as memory devices. These may be CDs, DVDs, memory “sticks”, external USB hard drives, and many other devices that can hold data. Even your digital camera probably can save files and be used as a memory store, because internally it has a memory device which can be accessed by a computer.

Bootable or not?

In order to understand why sometimes a device has to be “bootable”, it helps to understand what happens when you first switch on a PC. Briefly, the hardware must first find some instructions to execute. It is not given these in some magical manner, but in fact reads them off the first device it can find that has a suitable form of data at a specified location. Such a device will be known as bootable. If there is no data there, it is not.

So a bootable device has these data at the start of the device. If there is no data there, then the device is not bootable. Only bootable devices can be used to load a PC with new system software, be it Linux, Windows or Apple’s iOS.

Note: More accurately, it is the partition that has to be bootable, and in the right format that the BIOS of the PC can understand. Usually, data partitions are not bootable.

Partition table

Each device contains a description of how the data is laid out. But the software that reads and interprets this data has to be aware of the actual formats and where everything is. The first step is to identify where the various data areas are located. This is described in the partition table.

Whereas there are in theory several ways in which the table may be found, there is really only one you need to know about in practice, and that is called the MSDOS partition table.

The partition table is on addition to the bootable data. All devices that can be read by your PC will have a partition table.

Partitions

The partition table identifies the data areas that are treated as separate from one another. These areas are each known as a partition.

Each partition contains its data independently of all the other partitions, and can contain its data in different ways from them. This is the method to use when separating system programs and data from user data. Initially, it is sufficient to realise that one partition is needed and it will contain all the data on the device.

Some types of partition can be given labels which will then appear in a useful form when the device is plugged in. Labels are restricted in length, usually to 16 characters; some may also restrict the characters that can be used in the label.

File systems

To understand why we need file systems, consider how a system would find a file you have asked for, say, Story.odt. Somehow, the system (whether it be Linux, Apple iOS, or Windows, has to be able to find that file by looking through some sort of index information. Where that index is, and how it is laid out and to be interpreted is known as “meta-data”. A file system, in the current context, means the type of layout of the meta-data.

The different types of file system have different characteristics; they have been designed with different aims in mind. Some are sleek and fast for small numbers of smallish files. Such systems would be put onto the smaller devices. An example would be vfat.

Some systems are not expected to be re-written and modified after they have been constructed, and so need a simpler indexing system. For example, on a CD-R.

Some others have a “journalling” capability, which would enable them to recover in the event of a sudden power failure. These are used in secure databases and to ease the demand for frequent backups in other situations. A modern example is ext4. There is considerable work going on to improve these systems now, and more are being invented.

Others are designed for security of access, so that names of users who own and can access each file is essential information. All Unix file systems are traditionally capable of this as it is integral part of the Unix security model. A Windows file system with similar ability is ntfs.

For compatibility with all systems, only the vfat type is useful, but it has a limit to the amount of data it can store. Modern large devices normally come pre-formatted as ntfs.

As a word of caution, the driver that interprets the meta-data of your file system is the most important element in all of your software. If it has any bugs, the chance of it turning your hard-earned data into gobbledygook is increased considerably. My advice is generally to give a new system a couple of years in the hands of competent users to get the faults out before considering it for your own use. File system designers are among the most careful and particular engineers, and I have the greatest respect for them, but even so, there is no such thing as a bugless system of that complexity. But after a certain amount of time, the chance of hitting a bug under normal use decreases to an amount that is not worth bothering about.

A good rule would be to use a different file system for your backups to the one you use on your main data.

Formatting a device

Formatting a device is the name given to the process of preparing it to receive data.

The complete process will mean that the partition table must be defined and created; partitions must be defined and allocated space; each partition must then be given its file system; and it must be defined as bootable or not. Only when all these are in place can all the space on the device be used.

As you can probably guess, these things take up a little space of their own, which removes a small amount of space from all that could be available, but the proportion is minuscule.

Formatting a partition may take a surprisingly long time due to a number of factors, one of which is that as technology has progressed, the way the data is laid out may no longer be suitable for efficient use of the particular device.

One of the clearest programs to manage partitions on a disk is Gparted under Linux. I'll use it

under Ubuntu 10.10 to re-partition a new external portable hard drive from Samsung.

When the drive was first plugged into the system, it showed some files on it. I do not want those, but I want to start from an empty disk.

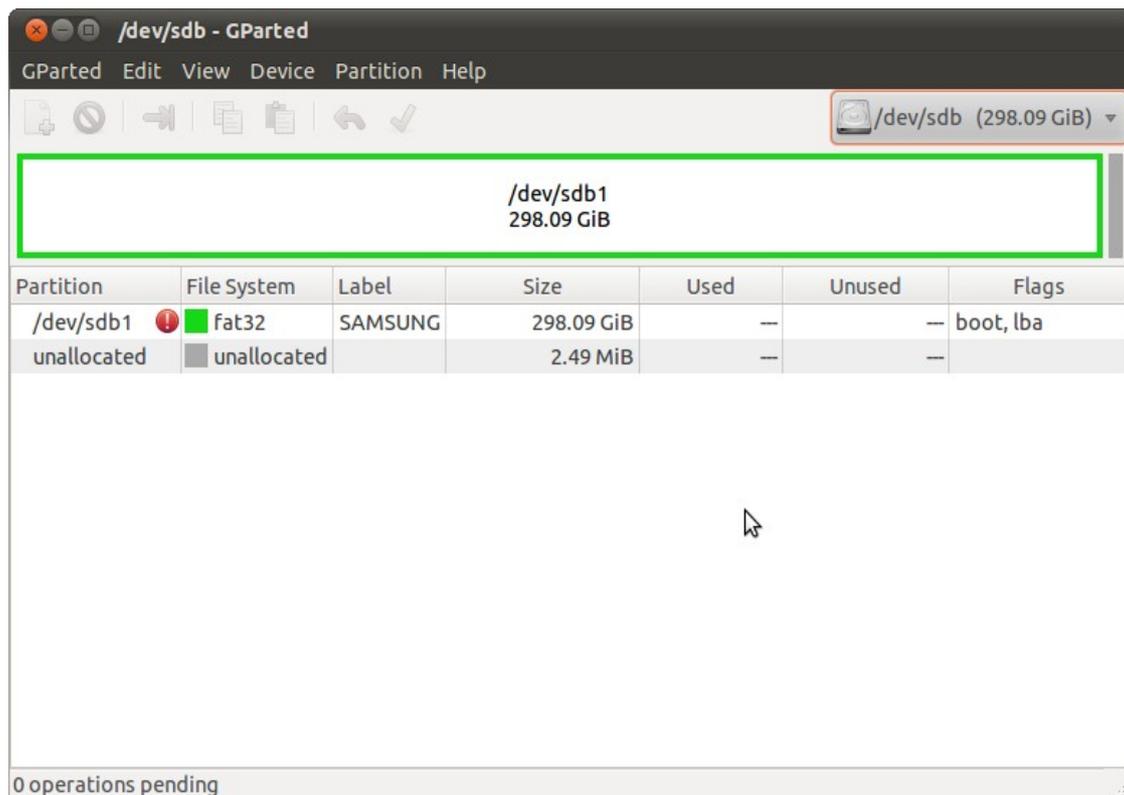
To format a disk and its partitions, it should not be “mounted” (see next section). So after checking the disk to see what it had, and that you do not want the contents of the partitions you intend to change, then close all windows that access it. You must not modify the partition containing an active file system or it will become unusable.

NOTE: DO NOT modify the file systems on the local disk in your machine from which it was currently booted. If you must do this, boot from a live CD or memory stick and run a system from there.

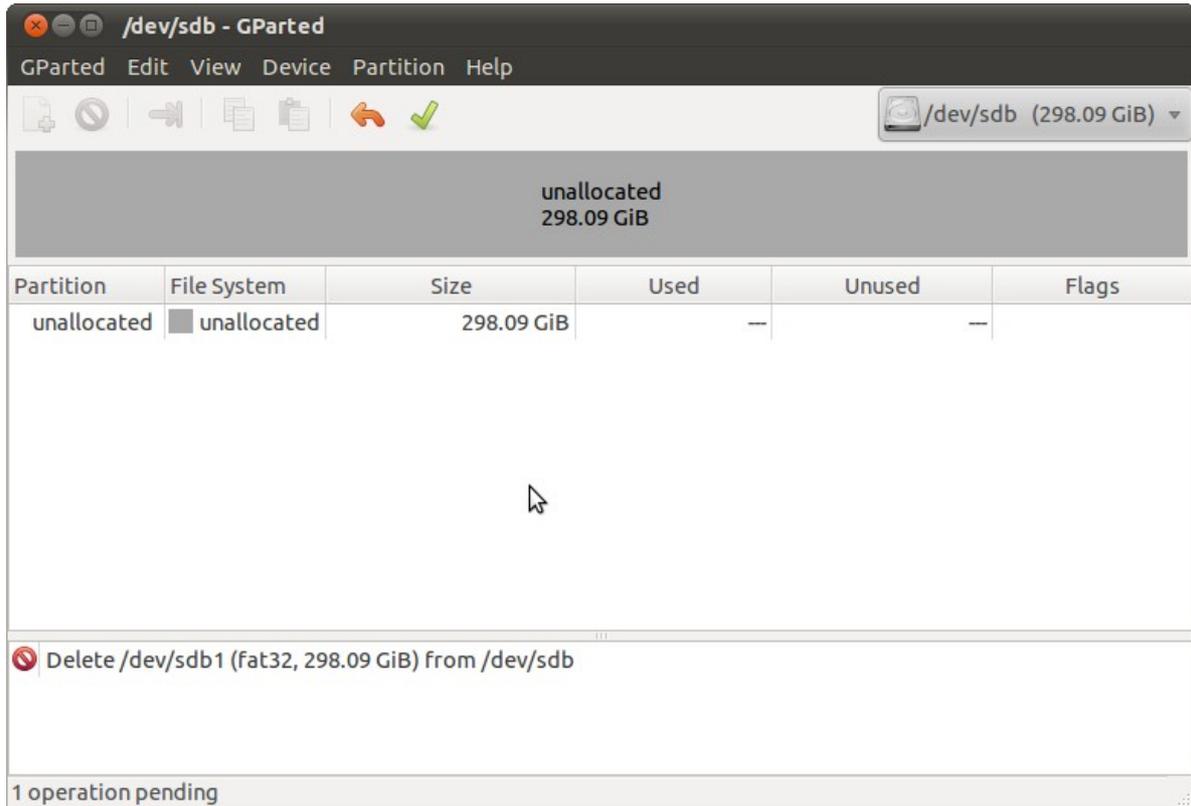
When all windows are closed, eject the device (e.g. right click on the desktop icon and Safely Remove), but do not detach from the USB port – leave it plugged in.

Then start GParted (under the System → Administration) and supply your administration password when asked.

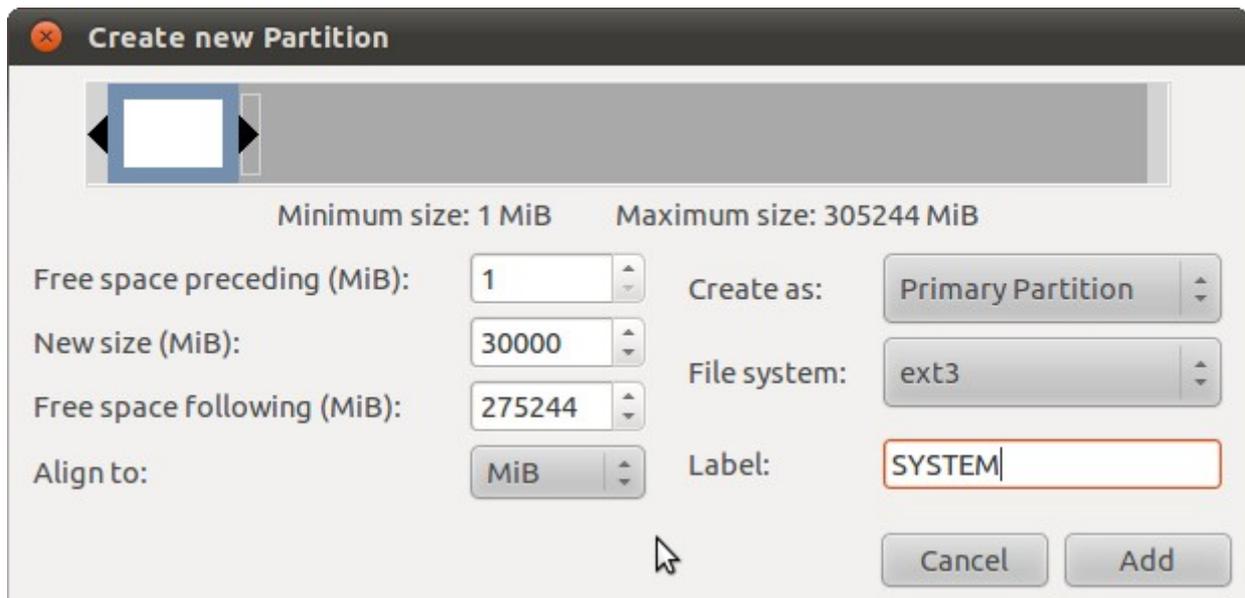
At the top right of the window is a drop-down list of devices. Select the one you want – probably something like /dev/sdb.



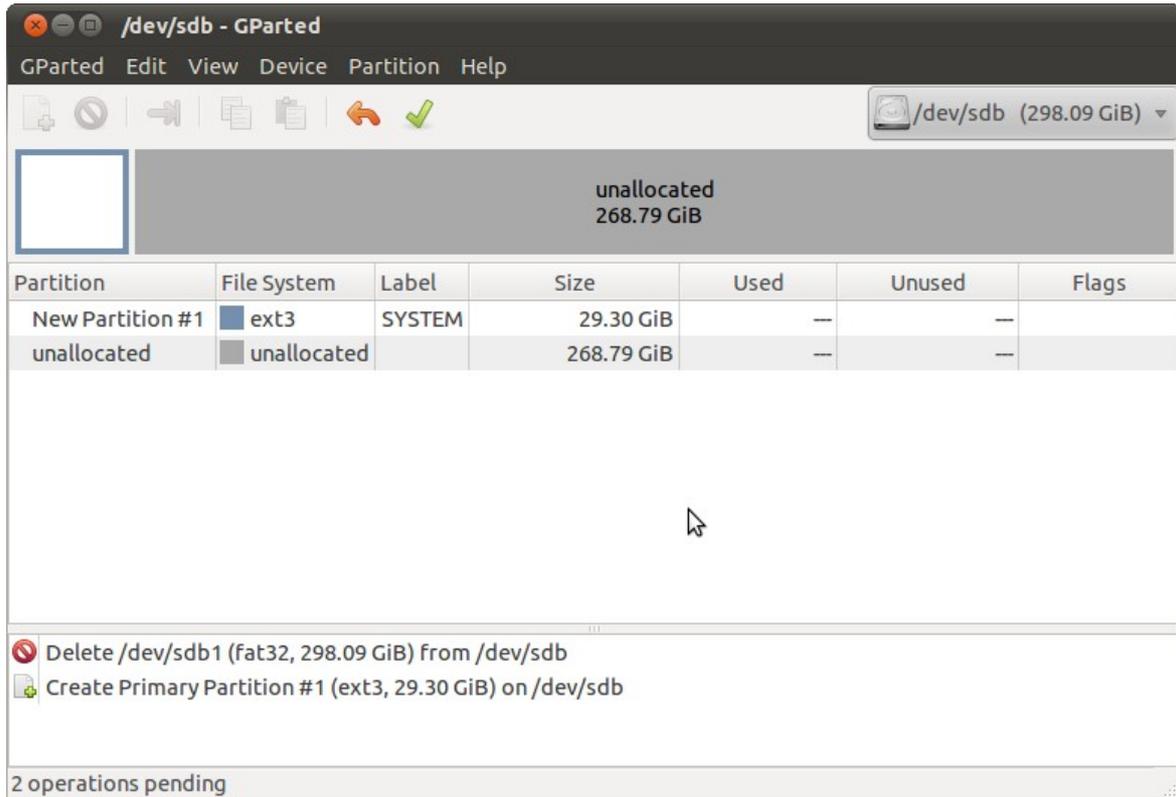
I don't want what is there, so I highlight the /dev/sdb1 partition, and go to Partition → Delete. A new pane appears at the bottom listing the tasks to be done. It is important to realise that no actions have yet been done; all that has happened is the Gparted has put this task in its to-do list, and will be done when I ask it to later. The data is still safe so far.



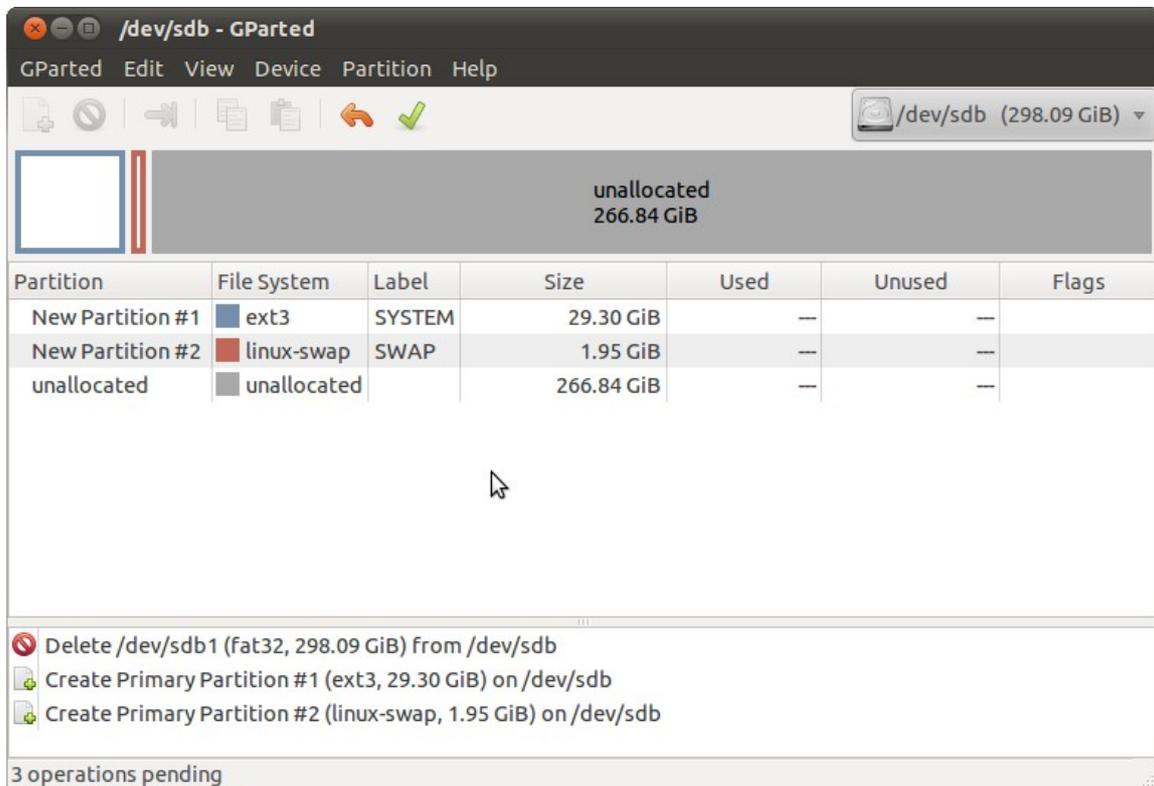
Next, I want to create a partition to be used to hold a system I can carry around and boot from when needed. This can be 30GB in size, be an ext3 file system, and I can label the partition SYSTEM. By selecting the unallocated area and then by going to Partition → New and filling in the details:



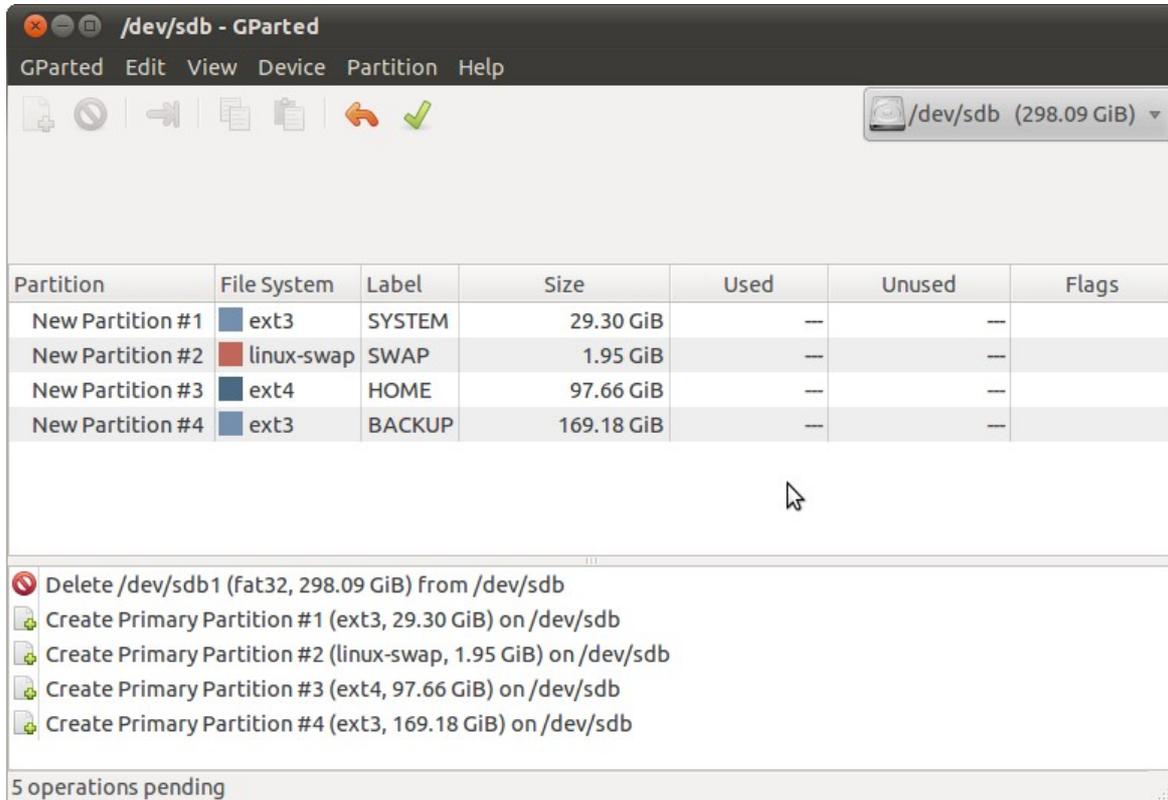
I see this:



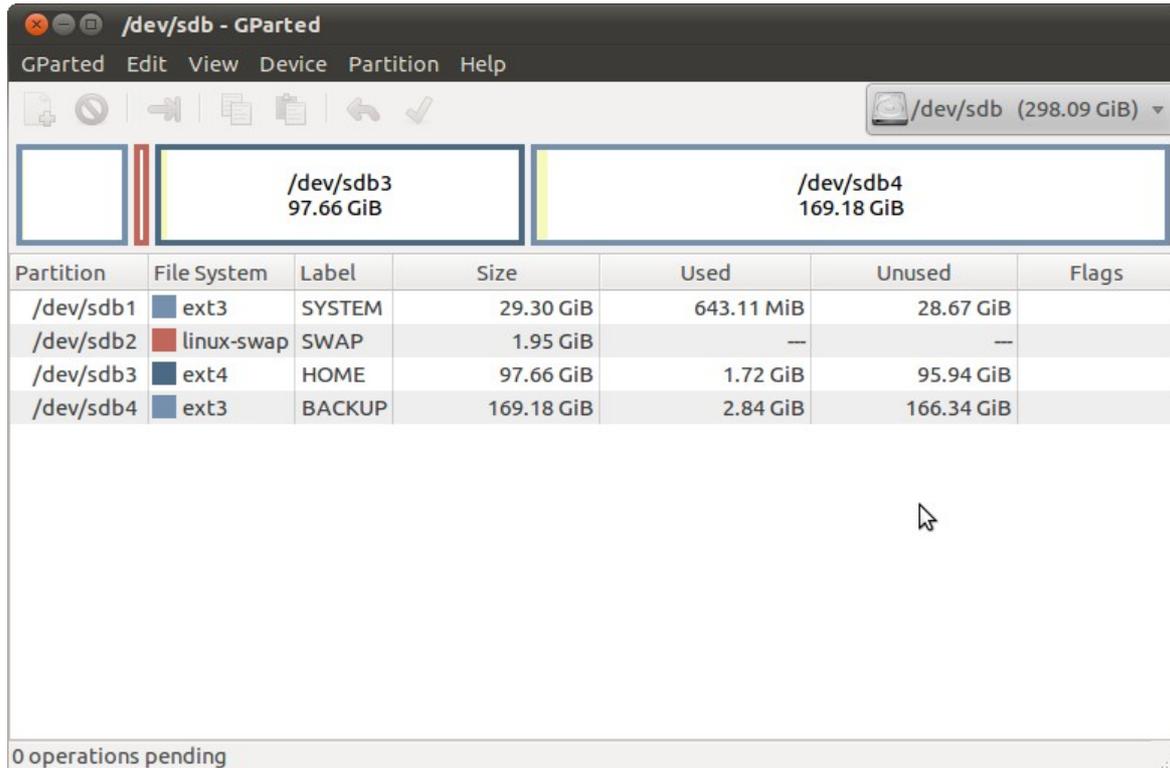
Any system needs a swap area, and that is what I do next, a new swap area partition, and 2GB is enough. Note that I've specified the type to be "linux-swap".



The home area can be another 100GB, the remainder I can set up to be a backup area to be used to back up data from my home machines to keep them separately.



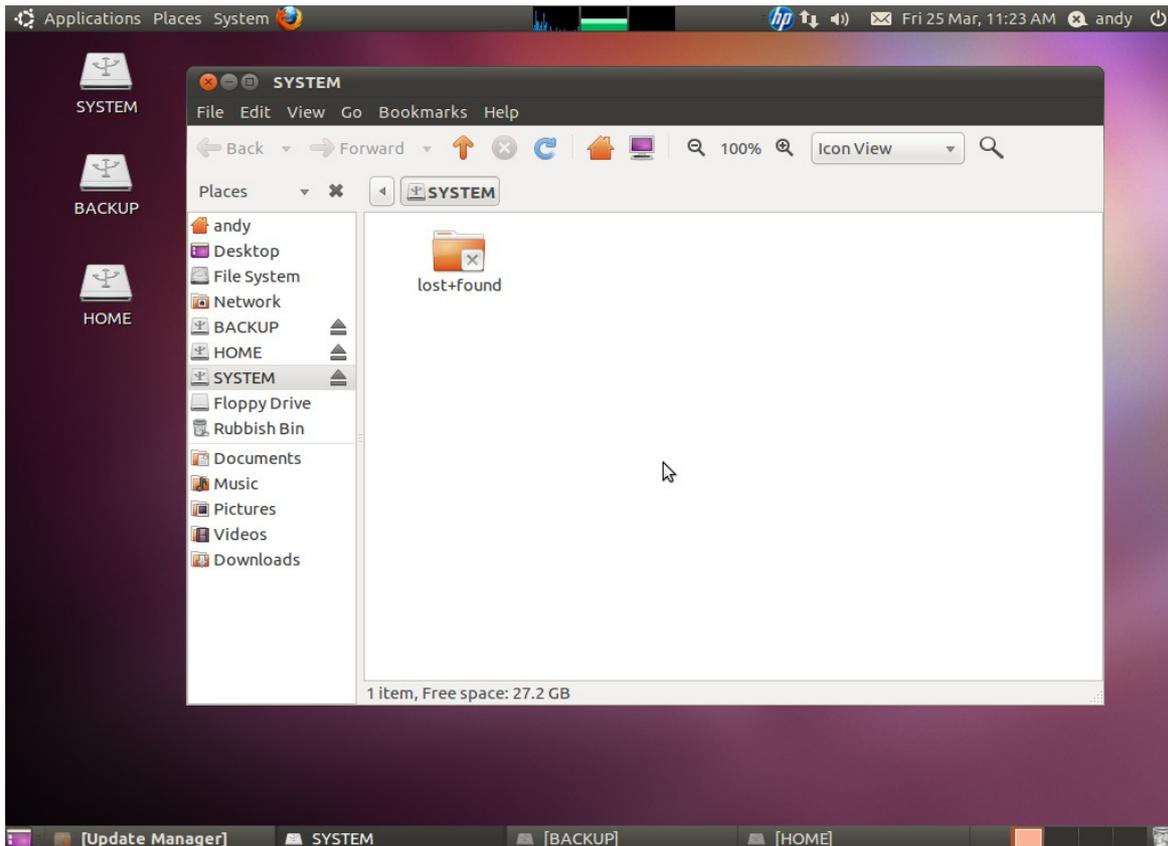
To complete the formatting, Edit → Apply All Operations. It asks you to make sure you really really really want to do this as it will destroy all data on the disk. Hitting Apply sets it going to do all the operations in the to-do list; but it could take a long time depending on the type of device, and the type of file system formatting that has been asked for.



Since I'll be trying to boot from this drive, I want the first partition to be bootable. By highlighting, the first partition and going to Partition → Manage flags, the “boot” flag can be set.

If I know that it will be used only for more modern systems, then setting the “lba” flag is also desirable as it enables larger areas to be managed. Older machines may not boot from a non-lba flagged device, however. Unfortunately, this decision should be taken when the partition table is created, and I did not do that in this case.

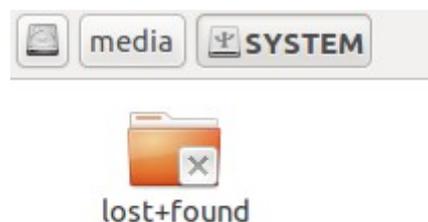
At the end of the partitioning and formatting, Gparted is closed, and the device is unplugged. When it is plugged back in, three of the partitions are found again (not the swap partition), and the result looks like this:



Mounting a device, or plug and play

When you plug a USB device into your PC, it will “mount” it. This means it will find what data it has on it, and each accessible partition will be made available. The details of the method depend on the operating system you are using. For Windows, a new device letter (e.g. H:) will be assigned for each partition. The way the OS will make it visible in a GUI will vary according to many things, like OS and Window managers, etc.

For Unix systems, the file system will appear in the directory structure somewhere; for Linux, it is typically under /media for almost all storage devices. In the previous picture, you can see that the label that was given to the partition is now the name of the directory that it finds. Clicking on the little left pointing triangle to the left of the the label SYSTEM, will show you where it is exactly, like this:



showing it is in /media/SYSTEM, and the others partitions likewise.

The alert will also notice that the swap partition has not appeared. That is because it has no file

system within it. It is formatted in a very special way to make the process of finding memory that has been moved out to disk very fast.

Further information

The Wikipedia page on tables is extensive with good external references for the fine detail for the serious engineers: http://en.wikipedia.org/wiki/Disk_partitioning