

FIRST STEPS LINUX BEGINNERS SERIES

Audacity Recording and editing

Follow the Pied Piper of Podcasting **Andy Channelle** and learn how to record, edit and mix audio.

LAST TIME

For calendars, email and opposable thumbs, we looked at *Evolution*, the first-class PIM from Novell. If you missed the issue, call 0870 8374773 or +44 1858 438975 for overseas orders.



Hats off to Ricky Gervais. The little man from Reading has created TV comedy gold with *The Office*, played a baddie in *Alias* and given us the most downloaded podcast in the (admittedly short) history of the medium (available at www.guardian.co.uk/rickygervais). If you want to emulate his success you'll need a few things: artistic genius, a good script... and the tools and technical know-how to prepare an audio stream for broadcast over the web.

Yes, this month's tutorial is all about recording and editing audio for use in podcasts (you could also use it for video presentations). We'll be working with *Audacity*, a free audio application that gives many commercial rivals a run for their money. *Audacity* was designed primarily as a musical tool, but increasingly it's being pressed into use as the best podcasting application available on all the main platforms.

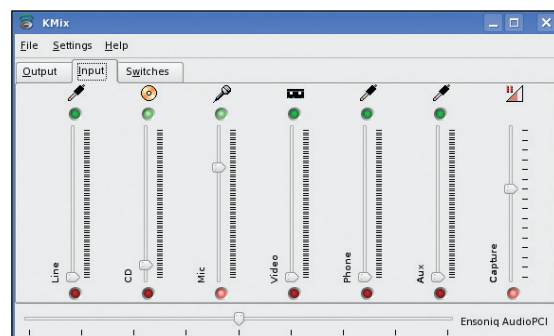
Your first job is to acquire and install the application itself. For most users this will involve nothing more strenuous than a trip to your package management application (*Yast*, *Yum*, *APT* etc) but if *Audacity* isn't available through this route the project's website (www.audacity.sourceforge.net) has binaries available for many distros – as well as the source code if you're feeling adventurous! Should all else fail, a search through *Rpmfind* (<http://rpmfind.net/linux/RPM>) will serve up binaries for every mainstream distro. You will also need the *Lame* MP3 encoder in order to export audio files that most listeners will be able to access. Some distros have specific *Lame* packages, but because of fear of litigation, most don't ship with the software,

so some users will have to seek out *Lame* from a separate source. A web search for 'Lame RPM' + distro name or 'Lame Deb' should find you the right file.

Any type of audio project can be broken down into four essential elements: capture, editing, assembly and output. In practice, the two middle parts will often bleed into each other, but when discussing any sort of broadcast or video accompaniment, they are distinct tasks. We'll begin at the beginning, which is recording a performance on to the PC.

Ready your sound system

If you want to follow my example, I'll be creating my inaugural podcast, which, inspired by a recent nostalgic TV programme about *Jackanory*, will feature a reading of Robert Browning's *The*

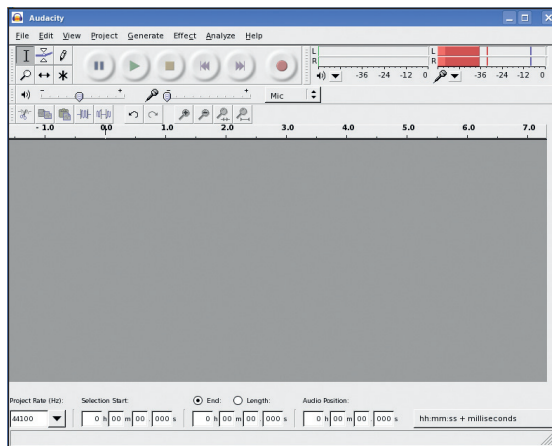


KMix is KDE's solution for managing soundcard connections.

Pied Piper of Hamelin (well out of copyright) with some suitable musical accompaniment and a few sound effects to add a bit of life. I'll assume you have a couple of essential items – a working soundcard with speakers and a microphone attached – and that they're in working order. You could record through another device such as a mixer or stereo system, in which case the signal will go into the PC through the 'line in' port. If this is the method you choose, substitute mentions of 'mic' with 'line in'.

The first step is to test the sound system, so open up either *KMix* or *Audio Mixer* from the panel. These are soundcard mixer programs, and the applet for them is usually hidden away under the Multimedia section, or else there will be a small icon in the system tray (the small section of the taskbar that contains running applets). Once the mixer window is open, choose the Input option and make sure that the mic (or line in) icon is selected – with both green and red lights lit – and that the slider is near the top. At this stage you should be able to tap the microphone and hear the taps through your speakers.

Another way to check if the signal is getting through is to start *Audacity*, make sure that Mic is selected in the drop-down list on the Recording toolbar and use the drop-down list next to the microphone icon on the Monitoring toolbar to select Start Monitoring. Now when you tap or speak into the microphone you should see a visual response on the audio meter.



Once you've tested your sound system, it's time to start recording. Here's *Audacity*, ready to roll.

Press Record

Before we capture the audio in *Audacity*, it makes sense to go into its preferences (Edit > Preferences) and change the default audio format to 2 Channel so that it captures stereo sound, and then use the File Formats tab to tell the application where the previously downloaded *Lame* encoder is (my tip is, it will be wherever you saved it!). Now we can begin.

A couple of test recordings will let us sort out issues such as background noise, microphone positioning and input levels, so hit the Record button on the transport controls (see *Audacity At A Glance*, page 77) and begin speaking, singing or performing in the way you expect the final piece to be done. If all is going well, the position bar should begin moving across the screen (left to right), creating a waveform as it goes. If your recording is too quiet, the waveform will be small and very close to the centre line; if you're being too loud, it will be very fat and close to the edges of the window. Ideally, you need something in between, with the loudest part of the performance pushing the waveform close to the top and bottom edges.

You might have realised by now that when you hit Record, a new track is created that will run concurrently with previously recorded tracks. This is useful if you're recording two parts and need to interact with another take, and if this is the case, you'll

need to go into Preferences again and ensure that Play Other Tracks While Recording The New One is selected in the Audio I/O > Playthrough section. Incidentally, this is also how you would use *Audacity* as a multi-track music recorder. If multi-track recording is not part of your agenda, just use the track's Close button (the small X in each track's top-left) to delete it.

Now you're all ready to go. Don't worry if it takes a few goes to get things right – going over the same thing again and again can be rationalised as a rehearsal, an opportunity to refine the script or an exploration of different delivery styles. And by saving each take, even if at first you decide it's absolute rubbish, you may be able to salvage a line or two when it comes to assembling the project. If you are doing multiple takes, by the way, try to position yourself (in relation to the microphone) in the same place for each recording. Wild variations in volume can be corrected, but it's best to get decent material recorded than have to fix it later.

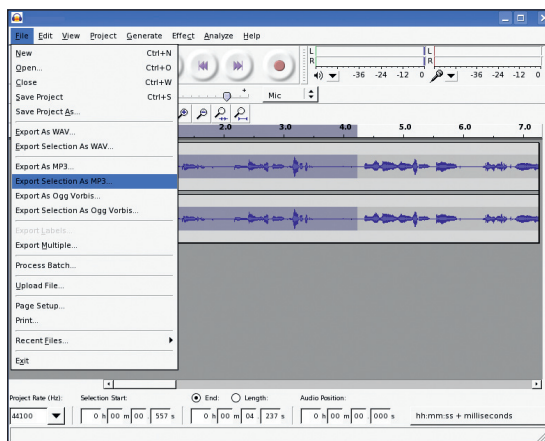
Audacity saves its raw audio files as uncompressed WAV files, which is not a very efficient storage method – each minute of uncompressed audio takes up about 10MB of hard disk space. However, unless you're constantly scrimping for disk space, I'd recommend keeping every take (or perhaps just the usable bits) until you've assembled the final piece.

If you're struggling for space, consider exporting each take (or partial take) as an MP3 or OGG file. Both formats compress the audio so that one minute will take up roughly 1MB of space, but OGG files give better sound quality than MP3s of the same size. To export the file (or a selection of a file) do File > Export As MP3.../Ogg Vorbis..., navigate to the correct location and give the file a name. Next time the file is needed, you can import it with the Project > Import Audio option. A word of warning though: exporting and importing the same audio file a number of times using one of these lossy compression formats (that is, some of the information is stripped out of the file as part of the compression process) will eventually lead to worse sound quality, so it's best to stick with uncompressed audio if possible.

Editing the recording

Editing is an amorphous term that I'm using to mean cleaning up files, adding any necessary special effects and generally getting the files ready for the assembly stage. This is the first time we'll get really acquainted with *Audacity*'s user interface.

The first job in editing is cleaning up. This means selecting areas that are supposed to be silent, and inserting a period of silence. To do this, select the area to be altered by left-clicking at the start point and dragging the mouse to the end point. Now >>



You can export a selection of audio in a number of different formats. WAV files are best for professionals; MP3 is the most widespread music format. OGGs provide the best balance of compression and sound quality.

QUICK TIP

You'll find some useful information in the Track Settings part of the *Audacity* interface about the format of your recordings, such as the sample rate (44.1kHz by default) and the format, which is set up to be 32-bit floating point. It's nice to have this information available to you, but you don't really need to change things. If you're up to date and using version 1.3, these track settings will have a small Up arrow at the bottom. Hit this and the track will 'roll up' into a smaller version of itself, providing more space for subsequent tracks to be viewed and edited. For the super-organised, double-clicking on the track name will allow you to change it to something more descriptive than the default Audio Track 1, 2 and so on, which is vital when you're dealing with a large number of files.

do Generate > Silence... and hit Generate. For a really professional sound, make the selection slightly smaller than needed (leave a small gap at the beginning and the end) so you can fade out the pre-silence sounds (select a small section before the silence and do Effect > Fade Out...) and fade in the post-silence stuff (again, select a portion of the waveform after the silence and do Effect > Fade In...). Now the sound will smoothly blend into the silence without an abrupt transition.

It's also possible to change the timing of parts of your recording by removing or shortening silences (assuming we're

clipboard. Before deselecting the section, use the method that I introduced above to replace it with silence, then hit the left cursor key to move the playhead back to the beginning of the section. Now do Project > New Stereo Track to make a new track and then Edit > Paste or Ctrl+V to paste the previously copied section of dialog into the new track.

It's now possible to change the stereo position or volume of this section (or move it around with the Time Shift tool) without affecting the other parts. I have, for instance, put the mayor of Hamelin on the left side of the stereo field (at about 40% on the stereo slider) and the Pied Piper himself on the right so that when someone plays back the recording the two voices will seem more distinct from one another.

Though there's tons more we could do with Audacity in terms of editing and effects, this is really all we need for basic podcast construction. All of the effects options are housed under the Effect menu and are fun to play with. So play with them!

“THERE MAY BE PAUSES IN YOUR DELIVERY WHEN YOU WANT THE BACKGROUND MUSIC TO SWELL.”

not removing whole sections) or inserting a bit of extra space. To remove silences, simply select a part of the waveform and hit the Backspace or Delete key. To insert silence, move the playhead (Audacity's cursor) to the right place and, again, do Generate > Silence..., but this time add some time to the box in the middle of the dialog box. Don't overdo it though; three seconds of silence doesn't seem a lot, but it is.

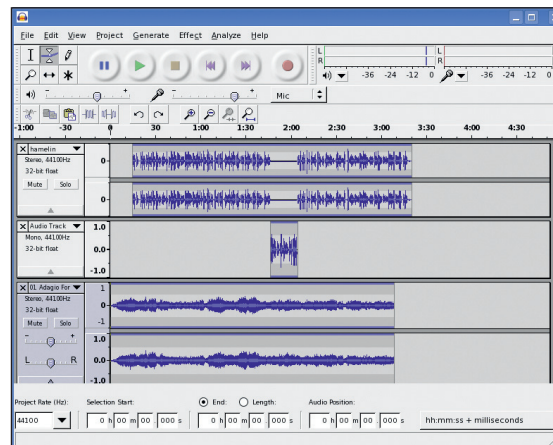
One other thing we'll do in here is shift a section of sound on to a different track, which makes it easy to move it around relative to other parts of the track. First select a section of track and do Edit > Copy (or hit Ctrl+C) to copy the selection to the

Assembling the piece

With the sounds cleaned up, edited and moved around (both in the track and the stereo field), it's time to do a bit of assembly. I am happy with the structure of the performance, but it needs some music. I have a (self-programmed) version of Barber's *Adagio for Strings* (still under copyright, sadly, so it must stay in the realm of LXF tutorials), which will make a suitably stirring accompaniment. Do Project > Import Audio File... and navigate to wherever your musical file is stored (see *Broadcasting And Copyright*, below, for information about open source music files) and select it. The file will be imported on to a new track, and it can be trimmed to match the length of the performance.

Hitting the Play button should play back all of the tracks simultaneously, but I want to have a short section of music before the vocal performance comes in. If all of the latter were on one track, it would simply be a case of using the Time Shift tool to grab the file and shift it to the right. However, we have two vocal tracks, as you know. To move multiple tracks, we can use the normal selection method of clicking the left mouse button on each Track Settings box while holding down Shift. Each selected track will then go a darker shade of grey, and they are effectively linked until you click on another track. So now we can use the Time Shift tool to move every selected track 15 seconds forward. *Adagio* will begin when we hit the Play button, and the vocal will start soon afterwards.

You may notice that everything is playing at the same volume. This might sound OK for a few seconds, but it will soon begin to annoy. The solution to this problem is a process called ducking. In essence, ducking lowers the volume of the music



With audio shifted around and music added (on tracks four and five) we are almost ready for the airwaves.

BROADCASTING AND COPYRIGHT

Copyright law is a very fluid thing, and the restrictions on the use of a text (by which I mean any piece of work that can be authored – be it a novel, symphony, painting or whatever) are changing all the time.

The idea behind copyright is to give an author a certain period of exclusivity in which to make a living off their creation before it becomes 'public domain'. Public domain simply means that the text can be used without permission or payment to the original copyright holder. The length of this copyright period has been increasing since the idea was put into law, and currently stands (in most countries) at 70 years after the death of the author. Robert Browning, the author of the text I'm using here in my podcast, died in 1889, so his works are no longer subject to copyright restrictions.

Once a text has reached the public domain it cannot be returned to private ownership. There are some areas

where trip-ups can occur, however. When using a text such as a book or poem that was originally written in another language, you must make sure the translation that you're using is also out of copyright. With a piece of music, the music itself may be out of copyright, but the recording of the performance may be more recent and so would have the usual restrictions.

This is a complex area for amateur podcasters to immerse themselves in, but Larry Lessig's book *Free Culture* is a great place to start. A free (in both senses) edition is available at www.free-culture.cc.

Be sure that you have the right to use any music that you include in your podcast, or you might end up getting your collar felt for copyright violation. As an alternative you could create your own 'bed' music using something like *Rosegarden* or *LMMS*, but for the less musically minded there are thousands of sounds at

websites such as [Freesound](http://freesound.iua.upf.edu/searchText.php) (<http://freesound.iua.upf.edu/searchText.php>), the Creative Commons audio project (<http://creativecommons.org/audio>) or the Open Directory (http://dmoz.org/Arts/Music/Sound_Files/Samples_and_Loops), which are available either for a one-off payment (ie they are royalty free) or under various Creative Commons (CC) licences.

Regardless where you acquire an audio file, you must read, and abide by, the licence that the sound was published under. This may mean, for example, that under certain CC licences you will have to release your derivative work under a similar licence, or that you will be free to use sounds only on a non-commercial basis.

It's worth checking out the Creative Commons project anyway, and thinking about whether one of its liberal copyright licences will be appropriate for your own efforts.

AUDACITY AT A GLANCE

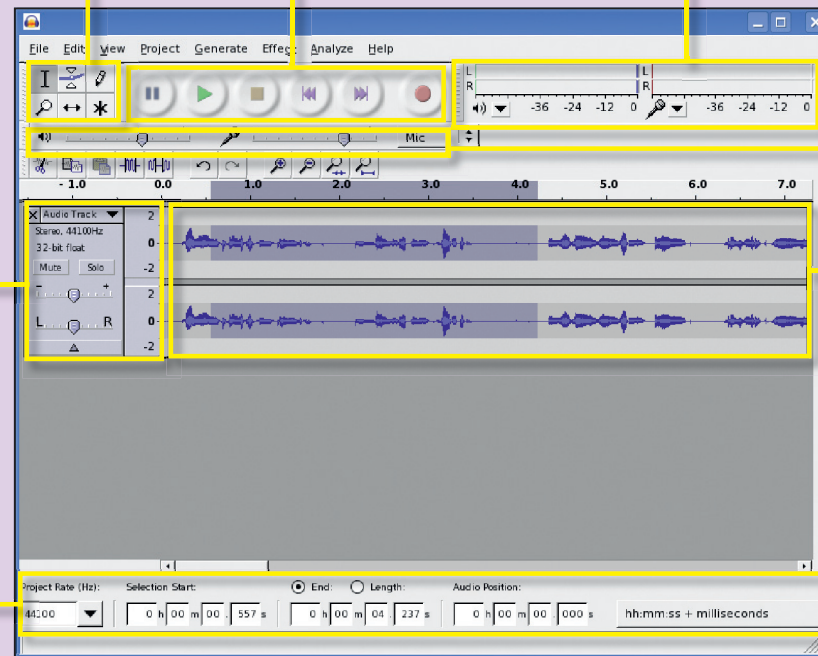
There are six tools in this little toolbox, but we're interested in just three: Selection, Envelope and Time Shift. For editing we need to stick with the Selection tool, which enables you to select different parts within the waveform window for copying, removing or isolating the selected range for the application of effects.

The transport controls replicate the buttons you would find on any recording device. The two back and forward buttons go to the end and beginning respectively of the longest section of audio in the project. The left and right keyboard cursor keys take you to the beginning or end of a selection. Shift+clicking Play will cause Audacity to loop around the current selection – that is, play it over and over again until you hit the Stop button.

This section shows you the levels of incoming and outgoing audio. By default the monitoring is switched off. To turn it on, click the small arrow by the mic icon and select Start Monitoring. You can save CPU power by slowing down the refresh rate of the metering in the Preferences dialog, but if your PC is really slow it may be better to leave monitoring turned off.

This is where you can set options for each individual track, such as the volume. There are two really useful extra buttons in here: the Mute button switches off the track while leaving everything else playing; and the Solo button switches everything off except the current track.

As well as selecting areas of the waveform graphically, you can enter measurement data in this selection bar for more accurate selections.

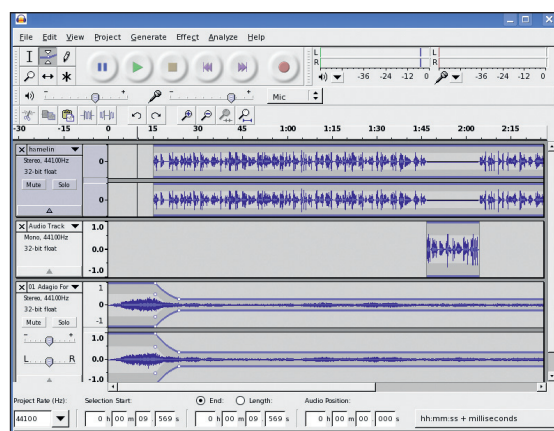


This Mixer section merely replicates the small section in the desktop's normal mixer and will influence the input and output volume.

The waveform is a visual representation of the sound recording and is where most of the work gets done. Focus in and out with the Zoom tool (left mouse button goes in, the right goes out) and go back and forth via the horizontal scroll bar at the base of the window.

when the vocal starts, and Audacity makes it really easy (in broadcasting terms this would be called a bed).

With the music track selected, choose the Envelope tool and add a control point on the track where the audio will begin to fade down by clicking with the left mouse button. This is the start of the fade. Now move a few seconds to the right and add another control point. This is where the fade will end. You should now be able to click and drag the second control point towards the centre of the waveform – making it narrower – and reduce the volume to a level where it doesn't detract from the vocal. There may be pauses in the delivery where you want the music to swell a little (this would be especially true of news



Here you can see three tracks, two dialogue (my *Pied Piper* reading) and one music (that's the *Adagio* bit), with the volume envelope changing as the vocal comes in. Podtastic!

podcasts where a quick stab of music can be used to break up the various stories), and so you could add more control points along the track to increase or decrease the volume as and when necessary.

Unleash your creation

We are at the final stage, which in recording parlance is known as the mixdown. This is where all of the tracks are blended into a stereo file ready for distribution. If the file were being burned to a CD, we'd just do File > Export As WAV, provide a name and location, and then hit OK. As I'm building this as a podcast for distribution over the internet, which puts a premium on small file sizes, we'll do the same but select Export As MP3... instead.

A warning will appear about mixing down to two channels – just hit OK – and you'll be asked to select a destination and filename. By default the exporter will output a file at 128kbps (kilobytes per second), which offers a good trade-off between quality and file size. However, if your performance is all vocal and you'd like to save some space, dropping the output down to 64kbps would not be too detrimental (and dial-up downloaders would probably thank you). To do this, go into Edit > Preferences and look under the File Formats tab. Within the MP3 Export Setup panel, you'll see a drop-down list labelled Bit Rate, which will enable you to set the appropriate value. Remember, smaller numbers make smaller sizes, but bigger numbers equals better quality.

Once the file is exported, you're ready to push it out on to the internet and into the ears of a potential audience of billions. [LXF](http://www.linuxformat.co.uk)

NEXT MONTH

Put some of your new Audacity skills to use in another area and add some soundtrack magic to your home videos.